

Erik Cambria & Amir Hussain



Sentic Computing

Techniques, Tools, and Applications

2nd Edition

May 9, 2012

Springer

*To Laura,
For believing in me and fostering my
passions and dreams, always.*

Foreword

When the authors asked me to write the foreword to this book, I readily accepted since it is a subject that intrigues me and one for which I realized that I have much to learn. It was only as I got deeper into this work that I realized what a seminal and important contribution to information technology this book has the potential to be. This book will likely become a classic in the field of sentic computing and be adopted as a textbook for students of the subject. It provides not only a comprehensive overview of the foundational elements of the topic, but introduces new methods along with teaching and research challenges for the future.

There is no doubt of the impact that the World Wide Web has had on human communication and discourse and will continue to have for the unforeseeable future. As its communications capabilities grow, its skills must also grow. I do not mean the skills of Web users, but the skills of the Web itself. The future Web will not just tie together documents, data, and people, it will also become a vastly connected world of processes, objects, and things all of which will have unique identities and personalities. Like the humans who use the Web, these objects must have common sense, be affective, and perhaps even emotional. To be effective, they must be able to make decisions and judgements some of which are likely to be quite subtle. Such qualities of these objects will be embedded, hidden, and defined within the data used by these entities. This data will become the DNA that prescribes the behavior of these objects. Like its biological counterpart, this DNA will contain the programming instructions for data manipulation governing object actions. But what will be the source of this data?

Web technology was revolutionized when users discovered the ability to use it as a platform for expressing opinions and evaluations. The previously static Web became a dynamic corpus of user-generated information such as product reviews, statistical polling data, survey data, etc. Within this information, just waiting to be uncovered, is a wealth of conceptual, expressive, and emotional content reflecting user responses to natural language questions and options. Unfortunately, the majority of this content is unstructured text containing all of the ambiguities found in spoken communications. The challenge became how to convert this information to structured data formats required for efficient and accurate computer analysis.

It was soon realized that intelligent mining of this data was a difficult task requiring the definition of new paradigms and the development of innovative analysis tools. For example, is it possible to perform clause-level semantic analysis of essentially free-form textual data which accurately allows the inference of both conceptual and emotional information? Employing methods from artificial intelligence, data mining, and such diverse disciplines as behavioral and sentiment analysis, the authors of this book explain why the answer is yes. The result is Sentic Computing.

Potential applications of Sentic Computing are widespread. Analyzed data drawn from product ratings and reviews are a marketer's dream and provide the new currency for e-commerce websites and social media. The authors describe how Sentic Computing techniques can influence HCI (Human-Computer Interaction) and E-Health. Readers of the book will likely realize other possible applications of the methods and tools in a range of other disciplines.

I encourage the reader to carefully study this book because doing so will provide a stimulating introduction to an area of study with great future potential. It will change your mind about the value (and the potential usage) of the opinions and evaluations that you make every day. It did so for me.

Stanford, April 2012

Bebo White

Preface

The ways people express their opinions and sentiments have radically changed in the past few years thanks to the advent of social networks, web communities, blogs, wikis, and other online collaborative media. The distillation of knowledge from this huge amount of unstructured information can be a key factor for marketers who want to create an image or identity in the minds of their customers for their product, brand, or organisation. These online social data, however, remain hardly accessible to computers, as they are specifically meant for human consumption. The automatic analysis of online opinions, in fact, involves a deep understanding of natural language text by machines, from which we are still very far.

Hitherto, online information retrieval has been mainly based on algorithms relying on the textual representation of web pages. Such algorithms are very good at retrieving texts, splitting them into parts, checking the spelling, and counting their words. But when it comes to interpreting sentences and extracting meaningful information, their capabilities are known to be very limited. Existing approaches to opinion mining and sentiment analysis, in particular, can be grouped into three main categories: keyword spotting, in which text is classified into categories based on the presence of fairly unambiguous affect words; lexical affinity, which assigns arbitrary words a probabilistic affinity for a particular emotion; statistical methods, which calculate the valence of affective keywords and word co-occurrence frequencies on the base of a large training corpus.

Early works aimed to classify entire documents as containing overall positive or negative polarity, or rating scores of reviews. Such systems were mainly based on supervised approaches relying on manually labelled samples, such as movie or product reviews where the opinionist's overall positive or negative attitude was explicitly indicated. However, opinions and sentiments do not occur only at document level, nor they are limited to a single valence or target. Contrary or complementary attitudes toward the same topic or multiple topics can be present across the span of a document. In more recent works, text analysis granularity has been taken down to segment and sentence level, e.g., by using presence of opinion-bearing lexical items (single words or n-grams) to detect subjective sentences, or by exploiting association rule mining for a feature-based analysis of product reviews.

These approaches, however, are still far from being able to infer the cognitive and affective information associated with natural language as they mainly rely on knowledge bases that are still too limited to efficiently process text at sentence level. In this book, common-sense computing techniques are further developed and applied to bridge the semantic gap between word-level natural language data and the concept-level opinions conveyed by these. In particular, the ensemble application of graph mining and multi-dimensionality reduction techniques on two common-sense knowledge bases was exploited to develop a novel intelligent engine for open-domain opinion mining and sentiment analysis. The proposed approach, termed *semantic computing*, performs a clause-level semantic analysis of text, which allows the inference of both the conceptual and emotional information associated with natural language opinions and, hence, a more efficient passage from (unstructured) textual information to (structured) machine-processable data.

The engine was tested on three different resources, namely a Twitter hashtag repository, a LiveJournal database, and a PatientOpinion dataset, and its performance compared both with results obtained using standard sentiment analysis techniques and using different state-of-the-art knowledge bases such as Princeton's WordNet, MIT's ConceptNet, and Microsoft's Probase. Differently from most currently available opinion mining services, the developed engine does not base its analysis on a limited set of affect words and their co-occurrence frequencies, but rather on common-sense concepts and the cognitive and affective valence conveyed by these. This allows the engine to be domain-independent and, hence, to be embedded in any opinion mining system for the development of intelligent applications in multiple fields such as Social Web, HCI, and e-health. Looking ahead, the combined novel use of different knowledge bases and of common-sense reasoning techniques for opinion mining proposed in this work, will, eventually, pave the way for development of more bio-inspired approaches to the design of natural language processing systems capable of handling knowledge, retrieving it when necessary, making analogies, and learning from experience.

Acknowledgements

The completion of this book would not have been possible without the contribution of many individuals, to whom I express my appreciation and gratitude. Because the contents of this book are mainly drawn from my doctoral research work, I am firstly deeply indebted to my Ph.D. supervisors, namely: Amir Hussain, who guided me every step of the way and was an immense source of inspiration throughout; Catherine Havasi, whose guidance, encouragement, and support in the past three years has been simply invaluable; and Chris Eckl, who helped me to look at my research from many different points of view.

I am also grateful to all the other mentors and colleagues who supported me during my scientific missions and internships, in particular, Robert Speer, Kenneth Arnold, Dustin Smith, Jason Alonso, and Henry Lieberman, who assisted me in getting started, using, and further developing common-sense computing tools (both during my research visits at MIT Media Lab and remotely); Campbell Grant, for his critical opinions and sound advice about the commercial aspects of my research work; Anna Esposito, for her unwavering support within the COST 2102 programme; Thomas Mazzocco and Marco Grassi, for their invaluable research contributions; Evan Magill and Kang Li, for their insights and assistance in refining the contents of this book; Joseph Lyons and Carmen Tropeano, for their help in designing some of the graphics; and James Munro, for the support and the data provided within my work in the field of patient opinion mining.

Special thanks also go to Praphul Chandra and Sudhir Dixit, who helped me expand the horizons of my research during my internship at Hewlett-Packard Labs India; Tariq Durrani, Tieniu Tan, and Cheng-Lin Liu, for their guidance during my research visit at the National Laboratory of Pattern Recognition (NLPR) in the Institute of Automation of the Chinese Academy of Sciences; and Haixun Wang and Yangqiu Song, for helping me improve my skills and expertise in the field of knowledge-based systems during my internship at Microsoft Research Asia. A last, but not least, acknowledgement goes to my family and all the old and new friends who, in the past three years, have cheered me up in difficult moments, celebrated with me for my achievements, and never blamed me for being often far away from them.

My Ph.D. was one of the best experiences of my life as it gave me the possibility to: work shoulder to shoulder with scientists from top research institutes, get in touch with both far West and far East cultures, meet special people who will always be part of my life, and even exchange ideas with beautiful minds such as Tim Berners-Lee, the inventor of the Web, Bebo White, Web pioneer in the United States, and Marvin Minsky, one of the fathers of AI.

Singapore, March 2012

Erik Cambria

Contents

Introduction	1
1.1 Sentic Computing	2
1.1.1 Motivations	3
1.1.2 Aims	5
1.1.3 Methodology	7
Background	11
2.1 Opinion Mining and Sentiment Analysis	12
2.1.1 The Buzz Mechanism	12
2.1.2 Origins and Peculiarities	13
2.1.3 Sub-Tasks	14
2.2 Main Approaches to Opinion Mining	16
2.2.1 From Heuristics to Discourse Structure	16
2.2.2 From Coarse to Fine Grained	17
2.2.3 From Keywords to Concepts	18
2.3 Towards Machines with Common-Sense	19
2.3.1 The Importance of Common-Sense	20
2.3.2 Knowledge Representation	21
2.3.3 From Logical Inference to Digital Intuition	25
2.4 Conclusions	31
Techniques	33
3.1 Affective Blending: Enabling Emotion-Sensitive Inference	34
3.1.1 AffectNet	34
3.1.2 AffectiveSpace	39
3.2 Affective Categorisation: Modelling Human Emotions	45
3.2.1 Categorical Versus Dimensional Approaches	46
3.2.2 The Hourglass of Emotions	49
3.3 Sentic Medoids: Clustering Affective Common-Sense Concepts ...	54
3.3.1 Partitioning Around Medoids	54
3.3.2 Centroid Selection	55

3.4	Sentic Activation: A Two-Level Affective Reasoning Framework . .	57
3.4.1	Unconscious Reasoning	57
3.4.2	Conscious Reasoning	58
3.5	Sentic Panalogy: Switching Between Different Ways to Think	60
3.5.1	Changing Reasoning Strategies	60
3.5.2	Changing Reasoning Foci	61
3.6	Conclusions	62
Tools		63
4.1	SenticNet: A Semantic Resource for Opinion Mining	64
4.1.1	Building SenticNet	65
4.1.2	Working with SenticNet	67
4.2	Sentic Neural Networks: Brain-Inspired Affective Reasoning	67
4.2.1	Discrete Versus Continuous Approach	68
4.2.2	Affective Learning	69
4.3	Open Mind Common Sentic: An Emotion-Sensitive IUI	71
4.3.1	Games for Knowledge Acquisition	71
4.3.2	Collecting Affective Common-Sense Knowledge	75
4.4	IsaCore: A Common and Common-Sense Knowledge Base	80
4.4.1	Probase	80
4.4.2	Building the Instance-Concept Matrix	83
4.5	Opinion Mining Engine: Structuring the Unstructured	85
4.5.1	Constitutive Modules	86
4.5.2	Evaluation	88
4.6	Conclusions	93
Applications		95
5.1	Development of Social Web Systems	96
5.1.1	Troll Filtering	96
5.1.2	Social Media Marketing	99
5.1.3	Sentic Album	106
5.2	Development of HCI Systems	116
5.2.1	Sentic Avatar	116
5.2.2	Sentic Chat	122
5.2.3	Sentic Corner	124
5.3	Development of E-Health Systems	128
5.3.1	Crowd Validation	129
5.3.2	Sentic PROMs	131
5.4	Conclusions	134
Concluding Remarks		137
6.1	Summary of Contributions	137
6.1.1	Techniques	138
6.1.2	Tools	138
6.1.3	Applications	139

6.2	Limitations and Future Work	139
6.2.1	Limitations	140
6.2.2	Future Work	141
6.3	Conclusions	142
	References	144

List of Figures

2.1	Tree of Porphyry. Porphyry presented the basis of Aristotle's thought as a tree-like scheme of dichotomous divisions, in which the process continues until the lowest species is reached.	24
2.2	An Open Mind project on OCR: handwritten characters are presented to e-citizens whose judgements (here 4 versus 9) are returned to the Open Mind host and used to train the classifier.	29
2.3	The cumulative analogy process allows to perform comparisons between concepts in a knowledge base (represented as a matrix) and, hence, to infer new information about similar concepts.	30
3.1	ConceptNet represents the information in the Open Mind corpus as a directed graph where nodes are concepts and labelled edges are assertions of common-sense that interconnect them.	36
3.2	A sketch of AffectiveSpace. Affectively positive concepts (in the bottom-left corner) and affectively negative concepts (in the up-right corner) are floating in the multi-dimensional vector space. . .	41
3.3	Accuracy values achieved by testing AffectiveSpace on BACK, with dimensionality spanning from 1 to 250. The best trade-off between precision and efficiency is obtained around 100.	43
3.4	A two-dimensional projection (first and second eigenmoods) of AffectiveSpace. From this visualisation, it is evident that concept density is usually higher near the centre of the space.	44
3.5	Russell's circumplex model is one of the earliest examples of dimensional emotion representations. In the snippet, direct circular scaling coordinates are provided for 28 affect words.	47
3.6	Whissell's model is a bi-dimensional representation of emotions, in which affect words are displayed. The diagram shows the position of some of these words in the <activation, evaluation> space.	48
3.7	The Pleasantness emotional flow. The passage from a sentic level to another is regulated by a Gaussian function that models how stronger emotions induce higher emotional sensitivity.	50

3.8	The 3D model and the net of the Hourglass of Emotions. Since affective states go from strongly positive to null to strongly negative, the model assumes a hourglass shape.	52
3.9	Hourglass compound emotions of second level. By combining basic emotions pairwise, it is possible to obtain complex emotions resulting from the activation of two affective dimensions.	53
3.10	Sentic Medoids organise AffectiveSpace by iteratively selecting centroids, in order to minimise the sum of dissimilarities between these and the other concepts within the same cluster.	56
4.1	SenticNet is a Semantic Web aware AI resource in which Open Mind concepts are identified using the ConceptNet Web API, and statements are encoded in RDF/XML format on the base of HEO. . . .	66
4.2	The structure of the <i>SNN-C</i> model. In the ‘continuous’ approach, the SNN provides a single real-valued output corresponding to the best guess of the level of affective valence.	70
4.3	The structure of the <i>SNN-D</i> model. In the ‘discrete’ approach, the SNN returns seven different real-valued outputs, each showing the degree of belonging to a specified affective level.	70
4.4	A screenshot of PopVideo. Two or more players are shown the same video and have to timely describe the objects that appear in the video in order earn points.	72
4.5	A screenshot of Concept Game, a turn-based game taking advantage of the Facebook platform for adding competitive flavours and social aspects to the collection of common-sense knowledge. . . .	74
4.6	A typical output obtained by inserting some text into Open Mind Common Sentic interface. The system provides a list of extracted concepts, a list of sentics, and a polarity value.	77
4.7	A screenshot of the Hourglass Game. The GWAP aims to collect affective common-sense knowledge from the general public by engaging users in a game where precision and speed are awarded. . . .	78
4.8	A screenshot of Sentic Pet. Icons in the up-right corner specify the abilities of pets, while icons in the up-left corner allow players to train or test these according to different modalities.	79
4.9	Probase consists of concepts (e.g., <i>emerging markets</i>), instances (e.g., “China”), attributes and values (e.g., “China’s population is 1.3 billion”), and relationships between concepts.	81
4.10	Opinion mining engine block diagram. After performing a first skim of the input text, the engine extracts concepts from it and, hence, infers related semantics and sentics.	86
5.1	Troll filtering process. Once extracted, semantics and sentics are used to calculate blogposts’ level of trollness, which is then stored in the interaction database for the detection of malicious behaviours. .	97

- 5.2 Merging different ontologies. The combination of HEO, WNA, OMR and FOAF provides a comprehensive framework for the representation of social media affective information. 101
- 5.3 A screenshot of the social media marketing tool. The faceted classification interface allows the user to navigate through both the explicit and implicit features of the different products. 103
- 5.4 Sentic extraction evaluation. The process extracts sentics from posts in the LiveJournal database, and then compare inferred emotional labels with the relative mood tags in the database. 104
- 5.5 Sentic Album's annotation module. Online personal pictures are annotated at three different levels: content level (PIL), concept level (opinion mining engine) and context level (Context Deviser). . . 109
- 5.6 Sentic Album's storage module. Image statistics are saved into the Content DB, semantics and sentics are stored into the Concept DB, timestamp and geolocation are saved into the Context DB. 111
- 5.7 Sentic Album's search and retrieval module. The IUI allows to browse personal images both by performing keyword-based queries and by adding/removing constraints on the facet properties. 113
- 5.8 Sentic Avatar's architecture. The system mainly consists of two modules for managing the avatar's inputs and outputs, and two modules for performing affective common-sense reasoning. 118
- 5.9 Tracked facial feature points (on the left) and corresponding facial parameters (on the right). The variations of such distances and angles are exploited for affect recognition. 119
- 5.10 A screenshot of the Affective Analysis module output. The integration of the multi-modal emotional information takes place into the Whissell space, in terms of activation and evaluation. 121
- 5.11 A few screenshots of Sentic Chat IUI. Stage and actors gradually change, according to the semantics and sentics associated with the on-going conversation, to provide an immersive chat experience. . . . 123
- 5.12 Sentic Corner generation process. The semantics and sentics extracted from the user's micro-blogging activity are exploited to retrieve relevant audio, video, visual, and textual information. 126
- 5.13 Sentic Corner web interface. The multi-modal information obtained by means of Sentic Tuner, Sentic TV, Sentic Slideshow, and Sentic Library is encoded in RDF/XML for multi-faceted browsing. 127
- 5.14 The semantics and sentics stack. Semantics are built on the top of data and metadata. Sentics are built on the top of semantics, representing the affective information associated with these. 129
- 5.15 The crowd validation schema. PatientOpinion stories are encoded in a machine-accessible format, in a way that they can be compared with the ratings provided by NHS Choices and each NHS trust. 130
- 5.16 Sentic PROMs prototype on iPad. The new interface allows patients to assess their health status and health-care experience both in a structured (questionnaire) and unstructured (free text) way. 133

List of Tables

1.1	List of most popular companies that are leveraging sentiment analysis tools to track and dissect how consumers feel about products and services of their own and also of the competition.	6
3.1	Comparison between WordNet and ConceptNet. While WordNet synsets contain vocabulary knowledge, ConceptNet assertions convey knowledge about what concepts are used for.	35
3.2	Some examples of LiveJournal posts where affective information is not conveyed explicitly through affect words. Such implicit information can be inferred by analysing the semantics and sentics.	38
3.3	Distribution of concepts through the Pleasantness dimension. The affective information associated with most concepts concentrates around the centre of the Hourglass, rather than its extremes.	42
3.4	Some existing definition of basic emotions. The most widely adopted model for affect recognition is Ekman's, although is one of the poorest in terms of number of emotions.	46
3.5	The sentic levels of the Hourglass model. Labels are organised into four affective dimensions with six different levels each, whose combined activity constitutes the 'total state' of the mind.	51
3.6	The second-level emotions generated by pairwise combination of the sentic levels of the Hourglass model. The co-activation of different levels gives birth to different compound emotions.	53
4.1	SNNs performance comparison with a KNN model and a random classifier. SNN-C performs better in the 'relaxed accuracy' case, while SNN-D is more suitable for the 'strict accuracy' case.	71
4.2	Comparison between Probase and two popular common knowledge bases. Probase overcomes problems such as incompleteness of WordNet hypernyms and biased knowledge of Freebase.	82

4.3	An overview of recent model-based sentiment analysis systems. Studies are divided by techniques applied, number of categories of the model adopted, corpora and knowledge base used.	85
4.4	Structured output example of opinion mining engine. For each clause, the engine detects the opinion target, the category it belongs to, and the affective information associated with it.	89
4.5	Precision values relative to Twitter evaluation. Probase and IsaCore perform significantly better than WordNet and ConceptNet, as these usually lack factual knowledge concepts.	92
4.6	F-measures relative to LiveJournal evaluation. Probase does not perform as well as the other knowledge bases as it is based on semantic rather than affective relatedness of concepts.	92
4.7	F-measure values relative to PatientOpinion evaluation. The ensemble application of IsaCore and AffectNet represents the best trade-off between topic spotting and affect recognition.	92
5.1	Precision, recall, and F-measure values relative to the troll filter evaluation. The AffectiveSpace process performs consistently better than IsaCore and AnalogySpace in detecting troll posts.	99
5.2	Evaluation results of the sentics extraction process. Precision, recall, and F-measure rates are calculated for ten different moods by comparing the engine output with LiveJournal mood tags.	105
5.3	Assessment of Sentic Album's accuracy in inferring the cognitive (topic tags) and affective (mood tags) information associated with the conceptual metadata typical of personal photos.	114
5.4	Perceived utility of the different interface features by 18 Picasa regular users. Participants particularly appreciated the usefulness of concept facets and timeline, for search and retrieval tasks.	115
5.5	Confusion matrix obtained combining the five classifiers. Success rates for <i>neutral</i> , <i>joy</i> , and <i>surprise</i> are very high, but <i>disgust</i> with <i>fear</i> , <i>anger</i> with <i>disgust</i> , and <i>fear</i> with <i>surprise</i> tend to be confused.	120
5.6	Confusion matrix obtained after considering human assessment. Success ratios considerably increase, meaning that the adopted classification strategy is consistent with human classification.	120
5.7	Precision, recall, and F-measure rates associated with Ekman's basic emotions (plus neutral) obtained by concurrently employing the opinion mining engine and the Facial Expression Analyser.	122
5.8	Perceived consistency with chat text of stage change and actor alternation. The evaluation was performed on a 130-minute chat session operated by a pool of 6 regular chat users.	124
5.9	Relevance of audio, video, visual, and textual information assembled over 80 tweets. Because of their larger datasets, Sentic Tuner and Slideshow are the best-performing modules.	128

Acronyms

3NF	Third Normal Form
AI	Artificial Intelligence
AKAPU	Average Knowledge Acquired Per User
ANN	Artificial Neural Network
API	Application Programming Interface
BACK	Benchmark for Affective Common-sense Knowledge
BCNF	Boyce-Codd Normal Form
CF-IOF	Concept Frequency - Inverse Opinion Frequency
CMC	Computer Mediated Communication
CRF	Conditional Random Field
CRM	Customer Relationship Management
DAG	Directed Acyclic Graph
DAU	Daily Active User
DL	Description Logic
ECA	Embodied Conversational Agent
ELM	Extreme Learning Machine
EQ	Emotional Quotient
FMRI	Functional Magnetic Resonance Imaging
FOAF	Friend Of A Friend
FOL	First Order Logic
GUI	Graphic User Interface
GWAP	Game With A Purpose
HCI	Human Computer Interaction
HEO	Human Emotion Ontology
HRQoL	Health Related Quality of Life
HTML	Hyper Text Markup Language
HVS	Human Visual System
ICA	Independent Component Analysis
IM	Instant Messaging
IT	Information Technology
IUI	Intelligent User Interface

KNN	K-Nearest Neighbor
KR	Knowledge Representation
JSON	JavaScript Object Notation
LSA	Latent Semantic Analysis
LSR	Label Sequential Rules
MAU	Monthly Active User
MDS	Multi-Dimensional Scaling
MLP	Multi-Layer Perceptron
MMO	Massively Multiplayer Online
MVE	Minimum Volume Ellipsoid
NB	Naïve Bayes
NELL	Never-Ending Language Learning
NMF	Non-negative Matrix Factorisation
NLP	Natural Language Processing
NP	Nondeterministic Polynomial
OCR	Optical Character Recognition
OMCS	Open Mind Common Sense
OMR	Ontology for Media Resources
OWL	Ontology Web Language
PAM	Partitioning Around Medoids
PCA	Principal Component Analysis
POS	Part Of Speech
PROM	Patient Reported Outcome Measure
RDBMS	Relational Database Management Systems
RDF	Resource Description Framework
RDFS	Resource Description Framework Schema
SBoC	Small Bag of Concepts
SKOS	Simple Knowledge Organisation System
SQL	Structured Query Language
SNN	Sentic Neural Network
SVD	Singular Value Decomposition
SVM	Support Vector Machine
TF-IDF	Term Frequency - Inverse Document Frequency
TMS	Truth Maintenance System
TSVD	Truncated Singular Value Decomposition
UGC	User Generated Content
UML	Unified Modeling Language
VSM	Vector Space Model
XML	Extensible Markup Language
W3C	World Wide Web Consortium
WNA	WordNet-Affect

Chapter 1

Introduction

*We can understand almost anything,
but we can't understand how we understand.*

Albert Einstein

In a world in which millions of people express their opinions about commercial products in blogs, wikis, forums, chats, and social networks, the distillation of knowledge from this huge amount of unstructured information can be a key factor for marketers who want to create an image or identity in the minds of their customers for their product, brand, or organisation [283]. The automatic analysis of online opinions, however, involves a deep understanding of natural language text by machines, from which we are still very far [179]. Online information retrieval, in fact, is still mainly based on algorithms relying on the textual representation of web pages [189]. Such algorithms are very good at retrieving texts, splitting them into parts, checking the spelling, and counting their words. But when it comes to interpreting sentences and extracting useful information for users, their capabilities are still very limited.

In this book, common-sense computing techniques are further developed and applied to bridge the cognitive and affective gap between word-level natural language data and the concept-level opinions conveyed by these. In particular, two common-sense knowledge bases are designed, together with a novel emotion categorisation model, and graph mining and multi-dimensionality reduction techniques are applied on them in order to infer cognitive and affective information from natural language text and, hence, develop opinion-mining systems in fields such as Social Web, HCI, and e-health. The structure of the book, specifically, is as follows: this chapter presents motivations, aims, and methodology of the proposed approach; chapter 2 illustrates the state of the art of opinion mining and common-sense computing; chapters 3, 4, and 5 explain in detail the developed techniques, tools, and applications, respectively; chapter 6, finally, comprises concluding remarks and future work.

1.1 Sentic Computing

Sentic computing is a multi-disciplinary approach to sentiment analysis that exploits both computer and social sciences to better recognise, interpret, and process opinions and sentiments over the Web. The approach specifically brings together lessons from both affective computing and common-sense computing because, in the field of opinion mining, not only common-sense knowledge, but also emotional knowledge is important to grasp both the cognitive and affective information (termed semantics and sentics) associated with natural language opinions and sentiments. Although scientific research in the area of emotion stretches back to the 19th century when Charles Darwin and William James proposed theories of emotion that continue to influence thinking today [71, 143], the injection of affect into computer technologies is much more recent.

During most of the last century, research on emotions was conducted by philosophers and psychologists, whose work was based on a small set of emotion theories that continue to underpin research in this area. The first researchers to try linking text to emotions were actually social psychologists and anthropologists who tried to find similarities on how people from different cultures communicate [213]. This research was also triggered by a dissatisfaction with the dominant cognitive view centred around humans as ‘information processors’ [183].

Later on, in the 1980s, researchers such as Turkle [285] began to speculate about how computers might be used to study emotions. Systematic research programs along this front began to emerge in the early 1990s. For example, Scherer [249] implemented a computational model of emotion as an expert system. A few years later, Picard’s landmark book affective computing [228] prompted a wave of interest among computer scientists and engineers looking for ways to improve human-computer interfaces by coordinating emotions and cognition with task constraints and demands. Picard described three types of affective computing applications:

1. Systems that detect the emotions of the user;
2. Systems that express what a human would perceive as an emotion;
3. Systems that actually ‘feel’ an emotion.

Although touching upon HCI and affective modelling, sentic computing primarily focuses on affect detection from text. Affect detection is critical because an affect-sensitive interface can never respond to users’ affective states if it cannot sense their affective states. Affect detection need not be perfect, but must be approximately on target. Affect detection is, however, a very challenging problem because emotions are constructs (i.e., conceptual quantities that cannot be directly measured) with fuzzy boundaries and with substantial individual difference variations in expression and experience. To overcome such a hurdle, sentic computing builds upon a biologically-inspired and psychologically-motivated affective categorisation model [52] that can potentially describe the full range of emotional experiences in terms of four independent but concomitant dimensions, whose different levels of activation make up the total emotional state of the mind.

In sentic computing, whose term derives from the Latin *sentire* (root of words such as sentiment and sentience) and *sensus* (intended both as capability of feeling and as common-sense), the analysis of natural language is based on affective ontologies and common-sense reasoning tools, which enable the analysis of text not only at document, page, or paragraph level, but also at sentence and clause level. In particular, sentic computing involves the use of AI and Semantic Web techniques, for knowledge representation and inference; mathematics, for carrying out tasks such as graph mining and multi-dimensionality reduction; linguistics, for discourse analysis and pragmatics; psychology, for cognitive and affective modelling; sociology, for understanding social network dynamics and social influence; finally ethics, for understanding related issues about the nature of mind and the creation of emotional machines.

In this section, motivations for the development of sentic computing are illustrated (section 1.1.1), together with the main aims of the proposed approach (section 1.1.2), and the methodology adopted (section 1.1.3).

1.1.1 Motivations

Opinions play a primary role in decision-making processes. Whenever people need to make a choice, they are interested in hearing others' opinions. When this choice involves consuming valuable resources (e.g., spending time and money to buy products or services), in particular, people strongly rely on their peers' past experiences. Just a few years ago, the main sources for collecting such information were friends, acquaintances and, in some cases, specialised magazine or websites. The advent of Web 2.0 has provided people with new tools, e.g., forums, blogs, social networks, and content sharing services, that allow them to create and share, in a time and cost efficient way, their own contents, ideas, and opinions with virtually the millions of people connected to the World Wide Web. This has made available by click a new and oceanic source of information and opinions and has provided a powerful communication medium to share knowledge and to get advantage from others' experiences [153].

Currently, over 75,000 new blogs are created daily, along with 1.2 million new posts each day, and more and more people in the modern world rely on opinions, reviews, and recommendations collected from these and related websites. The Web has made available the opinions of a vast pool of people that are neither our personal acquaintances nor well-known professional critics. People, in fact, are not just naturally keen on listening to others' advice, but also naturally inclined to give others advice. Web users are often happy to share both their positive and negative real-world experiences for different reasons, e.g., because they benefited from others' reviews and want to give back to the community, because they seek for a sense of togetherness in adversity, for cathartic complaining, for supporting a product they really like, because it is a way to express themselves, because they think their opinions are important for others.

When people have a strong feeling about a specific product or service they tried, they feel like expressing it. If they loved it, they want others to enjoy it. If they hated it, they want to warn others away. This huge amount of useful information, however, is mainly unstructured, that is in natural language, as it is specifically produced for human consumption and, hence, it is not directly machine-processable. The opportunity to capture the opinions of the general public about social events, political movements, company strategies, marketing campaigns, and product preferences has raised more and more interest both in the scientific community, for the exciting open challenges, and in the business world, for the remarkable fallouts in marketing and financial market prediction.

This has led to the emerging fields of opinion mining and sentiment analysis, which deal with information retrieval and knowledge discovery from text using data mining and natural language processing (NLP) techniques to distil knowledge and opinions from the huge amount of information on the World Wide Web. Mining opinions and sentiments from natural language, however, is an extremely difficult task as it involves a deep understanding of most of the explicit and implicit, regular and irregular, syntactical and semantic rules proper of a language.

Opinion mining and sentiment analysis are branches of the broad field of text data mining [136] and refer generally to the process of extracting interesting and non-trivial patterns or knowledge from unstructured text documents. They can be viewed as an extension of data mining or knowledge discovery from (structured) databases [107, 256]. Although commonly used interchangeably to denote the same field of study, opinion mining and sentiment analysis actually focus on polarity detection and emotion recognition, respectively. Since the identification of sentiment is often exploited for detecting polarity, however, the two fields are usually combined under the same umbrella or even used as synonyms.

As the most natural form of storing information is text, opinion mining is believed to have a commercial potential higher than that of data mining. Opinion mining, however, is also a much more complex task as it involves dealing with text data that are inherently unstructured and fuzzy. It is a multi-disciplinary research area that involves the adoption of techniques in fields such as text analysis, information retrieval and extraction, auto-categorisation, machine learning, clustering, and visualisation. Most of the existing approaches to opinion mining rely on the extraction of a vector representing the most salient and important text features, which is later used for classification purposes. Some of the most commonly used features are term frequency [308] and presence [220]. The latter is a binary-valued feature vectors in which the entries merely indicate whether a term occurs (value 1) or not (value 0) formed a more effective basis for review polarity classification.

This is indicative of an interesting difference between typical topic-based text categorisation and polarity classification. While a topic is more likely to be emphasised by frequent occurrences of certain keywords, overall sentiment may not usually be highlighted through repeated use of the same terms. Other term-based features are often added to the features vector. Position is one of these, in consideration of how the position of a token in a text unit can affect the way in which the token affect the sentiment of the text. Also presence n-grams, typically bi-grams

and tri-grams, are often taken into account as useful features. Some methods also rely on the distance between terms. Part of speech (POS) information (nouns, adjectives, adverbs, verbs, etc.) is also commonly exploited in general textual analysis as a basic form of word sense disambiguation [303]. Certain adjectives, in particular, have been proved to be good indicators of sentiment and sometimes have been used to guide feature selection for sentiment classification. In other works, eventually, the detection of sentiments was performed through selected phrases, which were chosen via a number of pre-specified POS patterns, most including an adjective or an adverb [286]. All such approaches mainly rely on parts of text in which opinions and sentiments are explicitly expressed, e.g., polarity terms, affect words, and their co-occurrence frequencies. Opinions and sentiments, however, are often conveyed implicitly through context and domain dependent concepts, which make purely syntactical approaches ineffective.

To this end, novel approaches that go beyond mere word-level sentiment analysis are needed. Such approaches should employ new techniques capable to better grasp the conceptual rules that govern sentiment and the clues that can convey these concepts from realisation to verbalisation in the human mind. Next-generation opinion mining systems need broader and deeper common-sense knowledge bases and more cognitive and affective inspired reasoning methods, in order to better understand natural language opinions and sentiments and, hence, more efficiently bridge the gap between (unstructured) textual information and (structured) machine-processable data.

1.1.2 Aims

Today, opinion mining and sentiment analysis find applications in several different scenarios and there is a good number of companies, large and small, that include the analysis of opinions and sentiments as part of their mission. In current product review websites, such as Epinions¹, Yelp², and RateItAll³, feedback and reviews are explicitly solicited within the web interface. Opinion mining techniques can be exploited for the creation and automated upkeep of review and opinion aggregation websites, in which opinions are continuously gathered from the Web and not restricted to just product reviews, but also to wider topics such as political issues and brand perception. Opinion mining and sentiment analysis have also a great potential as sub-component technology for other system. They can enhance the capabilities of customer relationship management (CRM) and recommendation systems allowing, for example, to find out which features customers are particularly interested in or to exclude from the recommendations items that have received very negative feedbacks [277, 278].

¹ <http://epinions.com>

² <http://yelp.com>

³ <http://rateitall.com>

Similarly they can be used in email or other types of communication to detect and exclude ‘flames’, i.e., overly heated or antagonistic language, and to enhance anti-spam systems. Also, online systems that display advertisements as sidebars can use opinion mining techniques to detect web pages that contain sensitive content inappropriate for ads placement [144]. Business intelligence is also one of the main factors behind corporate interest in the field of sentiment analysis [219]. Nowadays, companies invest more and more money in marketing strategies and they are constantly interested in both collecting and predicting the opinions and the attitudes of the general public towards their products and brands. The design of automatic tools capable to crawl reviews and opinions over the Web in real-time and to create condensed versions of them represents one of the most active research and development area. Several companies, in fact, already provide tools to track public viewpoints on a large scale by offering graphical summarisations of trends and opinions in the blogosphere (Table 1.1).

The development of such systems, moreover, is not only important for commercial purposes, but also for government intelligence applications able to monitor increases in hostile or negative communications [1]. All of these tools, however, are still mainly keyword based and, hence, often fail to meet the gold standards of human annotators. The fundamental aim of this research work is to go beyond such approaches by developing common-sense knowledge bases to bridge the cognitive and affective gap between word-level natural language data and the concept-level opinions conveyed by these. Unlike keyword-based methods, sentic computing uses affective ontologies and common-sense reasoning tools for a concept-level analysis of natural language text.

Company	Founded	Headquarters	Web Link
Vocus	1992	USA	http://vocus.com
Kantar	1993	UK	http://www.kantar.com
Cymphony	1996	USA	http://www.cymfony.com
Alterian	1997	USA	http://alterian.com
Factiva	1999	USA	http://dowjones.com/factiva
Brandimensions	2001	Canada	http://brandprotect.com
Attensity	2000	USA	http://attensity.com
Converseon	2001	USA	http://converseon.com
Lithium	2001	USA	http://lithium.com
Lexalytics	2003	USA	http://lexalytics.com
MotiveQuest	2003	USA	http://www.motivequest.com
Visible Technologies	2003	USA	http://visibletechnologies.com
Evolve24	2004	USA	http://evolve24.com
Clarabridge	2005	USA	http://clarabridge.com
Collective Intellect	2005	USA	http://collectiveintellect.com
Radian6	2006	Canada	http://radian6.com
Rapid-I	2006	UK	http://rapid-i.com
Luminoso	2011	USA	http://lumino.so

Table 1.1 List of most popular companies that are leveraging sentiment analysis tools to track and dissect how consumers feel about products and services of their own and also of the competition.

Specifically, the ensemble application of graph mining and multi-dimensionality reduction techniques is employed, together with a novel emotion categorisation model, on two common-sense knowledge bases, in order to design an open-domain opinion mining engine capable to infer the cognitive and affective information associated with natural language text. Such engine has been exploited for the development of emotion-sensitive systems in fields such as social data mining, multimedia management, personalisation and persuasion, human-computer interaction, intelligent user interfaces, social media marketing, and patient-centred applications.

Evidence of the impact of the approach is found in the presence of sentic computing in high impact factor journals and top AI conferences, and in its adoption by several leading American, British, and Asian companies, including: Zoral Inc., Luminoso Inc., Abies Ltd., Patient Opinion Ltd., Sitekit Solutions Ltd., HP Labs India, and Microsoft Research Asia. For these reasons, sentic computing has also been recently put forward as impact case study to the UK Research Excellence Framework (REF) by the University of Stirling.

1.1.3 Methodology

Relying solely on traditional methods to develop computer systems with a new set of affect-sensitive functionalities is insufficient [26] because today user emotions are still far from being on the radar of computing methods. This is where insights gleaned from a century and a half of scientific study on human emotions can become useful for the development of affect-sensitive interfaces. Despite the extensive literature in emotion research, however, the affective computing literature has been primarily driven by computer scientists and AI researchers who have remained agnostic to the controversies inherent in the underlying psychological theory. Instead, they have focused their efforts on the technical challenges of developing emotion-sensitive computer interfaces. However, ignoring the important debates has significant limitations because a functional affective computing application can never be completely divorced from underlying emotion theory [36].

Blending scientific theories of emotion with the practical engineering goals of analysing sentiments in natural language text and developing affect-sensitive interfaces is one of the main contributions of this book. Recently, many research activities focusing on the extraction of cognitive and affective information from natural language text have gained ground under the umbrella of opinion mining and sentiment analysis. The reason of this trend lies on the ever-growing amount of valuable data available through the Web in the form of news, reviews, blogs, chats, tweets, etc. Sentiment analysis, however, is a multi-faceted and multi-disciplinary problem that requires a deep understanding of natural language. Existing reported solutions and currently available systems are still far from perfect or fail to meet the satisfaction level of the end users. The main issue may be that there are many conceptual rules that govern sentiment and the possibly unlimited clues that can convey these concepts from realisation to verbalisation in the brain.

Recent efforts in this context have been carried about through research works published in reputed conferences through special tracks and workshops, e.g., the *TREC-BLOG* tracks since 2006, the *Sentiment and Subjectivity in Text* workshop in COLING-ACL 2006, the *SemEval 2007 Task#14: Affective Text*, the *TAC 2008 Opinion Summarisation* task, *Emotion, Metaphor, Ontology and Terminology* in LREC 2008, the *Social Data on the Web (SDoW)* workshop from 2008, the *Workshop on Opinion Mining and Sentiment Analysis (WOMSA)* in 2009, *Topic Sentiment Analysis for Mass Opinion Measurement (TSA)* in CIKM 2009, *Computational Approaches to Analysis and Generation of Emotion in Text* in NAACL 2010, the *Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)* in ECAI 2010 and in ACL 2011, the ICDM workshop series on *Sentiment Elicitation from Natural Text for Information Retrieval and Extraction (SENTIRE)* since 2011, the *KDD Workshop on Issues of Sentiment Discovery and Opinion Mining (WISDOM)* since 2012, and the WWW workshop series on *Multidisciplinary Approaches to Big Social Data Analysis (MABSDA)* since 2013.

Sentic computing extended the findings expressed in such research works to develop new strategies for open-domain opinion mining, based on the implicit features associated with natural language concepts. The novelty of the approach, in particular, lies in:

1. its multi-disciplinarity (not only computational, but also biologically-inspired and psychologically-motivated);
2. its semantic-based analysis (not only based on word co-occurrence frequencies, but also on the cognitive and affective information associated with natural language);
3. its fine-grained classification (not only at document, page, or paragraph level, but also at sentence and clause level).

In order to evaluate the different facets of the engine from different perspectives, three different resources, namely a Twitter⁴ hashtag repository, a LiveJournal⁵ database and a PatientOpinion⁶ dataset, were used and results obtained using Princeton's WordNet⁷, MIT's ConceptNet⁸ and Microsoft's Probase⁹ were compared. The first resource is a collection of 3,000 tweets crawled from Bing¹⁰ web repository by exploiting Twitter hashtags as category labels, which is useful to test the engine's target spotting performance. In particular, hashtags about electronics (e.g., iPhone, XBox, Android and Wii), companies (e.g., Apple, Microsoft and Google), countries, cities, operative systems and cars were selected. In order to test the resource's consistency and reliability, a manual evaluation of 100 tweets was performed, which showed that hashtags are accurate to 89%.

⁴ <http://twitter.com>

⁵ <http://livejournal.com>

⁶ <http://patientopinion.org.uk>

⁷ <http://wordnet.princeton.edu>

⁸ <http://conceptnet5.media.mit.edu>

⁹ <http://research.microsoft.com/probase>

¹⁰ <http://bing.com>

The second resource is a 5,000 blogpost database extracted from LiveJournal, a virtual community of more than 23 millions users who keep a blog, journal or diary. An interesting feature of this website is that bloggers are allowed to label their posts with both a category and a mood tag, by choosing from predefined categories and mood themes. Since the indication of mood tags is optional, posts are likely to reflect the true mood of the authors, which is not always true for category tags. After a manual evaluation of 200 posts, in fact, the category tags turned out to be very noisy (53% accuracy). The mood tags, however, showed a good enough reliability (89% accuracy) so they were used to test the engine's affect recognition performance. The third resource, eventually, is a dataset obtained from PatientOpinion, a social enterprise pioneering an online feedback service for users of the UK national health service to enable people to share their recent experience of local health services online. It is a manually tagged dataset of 2,000 patient opinions that associates to each post a category (namely, clinical service, communication, food, parking, staff, and timeliness) and a positive or negative polarity. It was used to test the detection of opinion targets and the polarity associated with these.

There are no ethical issues involved in the data used in the experimentation as tweets, blogposts, and patient opinions were all anonymised. In order to guarantee full anonymity, moreover, the text associated with tweets, blogposts, and patient opinions has never been wholly reported in the proposed tables and examples.

Chapter 2

Background

*The good opinion of mankind,
like the lever of Archimedes,
with the given fulcrum, moves the world.*

Thomas Jefferson

The World Wide Web represents one of the most revolutionary applications in the history of computing and human communication, which is keeping on changing how information is disseminated and retrieved, how business is conducted and how people communicate with each other. As the dimension of the Web increases, the technologies used in its development and the services provided to its users are developing constantly. Even if just few years have passed, in fact, Web 1.0's static and read-only HTML pages seem now just an old memory. Today the Web has become a dynamic and interactive reality in which more and more people actively participate by creating, sharing, and consuming contents. In this way, the World Wide Web configures itself not only as a 'Web of data', but also as a 'Web of people' where data and users are interconnected in an unbreakable bond.

This chapter shows how and why online opinions are important in the Web 2.0 era (section 2.1) and illustrates existing approaches and depths of analysis in mining and characterising such opinions (section 2.2). Eventually, the chapter comprises a background section on common-sense knowledge representation and reasoning, which later will be further developed and applied to go beyond merely syntactical approaches to sentiment analysis (section 2.3), and some concluding remarks (section 2.4).

2.1 Opinion Mining and Sentiment Analysis

The passage from a read-only to a read-write Web made users more enthusiastic about interacting, sharing, and collaborating through social networks, online communities, blogs, wikis, and other online collaborative media. In the last years, this collective intelligence has spread to many different areas of the Web, with particular focus on fields related to our everyday life such as commerce, tourism, education, and health. The online review of commercial services and products, in particular, is an action that users usually perform with pleasure, to share their opinions about services they have received or products they have just bought, and it constitutes immeasurable value for other potential buyers.

This trend opened new doors to enterprises that want to reinforce their brand and product presence in the market by investing in online advertising and positioning, that is, in social media marketing. The reasons why opinion mining is attracting so much attention from both the academic and the business world, in particular, can be found in the dynamics behind the buzz mechanism (section 2.1.1), in the motivating factors that gave birth to the field (section 2.1.2), and in the sub-tasks that make it different from standard information retrieval (section 2.1.3).

2.1.1 *The Buzz Mechanism*

What mainly makes social media marketing work is the buzz mechanism [248]. A buzz replicates a message through user-to-user contact, rather than purchasing some advertising or promoting a press release. The message does not have to necessarily deal with the product. Many successful viral campaigns, in fact, have spread thanks to a compelling message, with the company logo included incidentally. At the heart of buzz is an understanding that the natural, spontaneous networks that comprise the social universe are the most effective means of reaching people in a meaningful way. The power of marketing lies, therefore, not in pushing information to the masses, but in effectively tapping those individuals who wield influence over others.

The marketers who are winning are the ones using consumers and culture to their advantage, crafting messages with consumers rather than throwing messages at them. In confirmation of the growing interest in this novel approach to marketing, several academic and commercial tools, e.g., OASYS¹ [66], ESSE [264], Lumino² [269], Factiva³, NM Incite⁴, Attensity⁵, and Converseon⁶, have been developed to provide companies (and users) with a way to analyse the blogosphere on a

¹ <http://oasys.umiacs.umd.edu/oasys>

² <http://lumino.so>

³ <http://dowjones.com/factiva>

⁴ <http://nmincite.com>

⁵ <http://attensity.com>

⁶ <http://converseon.com>

large scale, in order to extract information about the trend of the opinions relative to their products. Nevertheless most of the existing tools and the research efforts are limited to a polarity evaluation or a mood classification according to a very limited set of emotions. In addition, such methods mainly rely on parts of text in which emotional states are explicitly expressed and, hence, are unable to capture opinions and sentiments that are expressed implicitly.

2.1.2 Origins and Peculiarities

The term ‘opinion mining’ first appears in a paper by Dave et al. [90] dated 2003, which envisioned the ideal opinion mining tool as capable to “process a set of search results for a given item, generating a list of product attributes (quality, features, etc.) and aggregating opinion about each of them (poor, mixed, good)”. From this early definition, the term opinion mining has been later extended to refer more generally to the computational techniques for extracting, classifying, understanding, and assessing the opinions expressed in various online news sources, social media comments, and other user-generated contents (UGCs). The introduction of the term ‘sentiment’ to the automatic analysis of evaluative text and tracking of the predictive judgements was first introduced in 2001 by Das and Chen [87] and Tong [280], in the context of market sentiment analysis.

In the context of NLP, the term sentiment can be intended either as the emotions or the polarity conveyed by text. Strictly speaking, sentiment analysis consists in inferring affective information from text, while opinion mining mainly concerns polarity detection. However, since the identification of sentiment, affect, subjectivity, and other emotional states is often propaedeutic to polarity detection [140], opinion mining and sentiment analysis are strictly connected and, therefore, commonly used interchangeably to denote the same field of study.

The manifesto of opinion mining and sentiment analysis as a unified field can be seen in the extensive review paper published by Pang and Lee [219] in 2008. This survey covers techniques and approaches that promise to directly enable opinion-oriented information-seeking systems. The authors’ focus is on methods that seek to address the new challenges raised by sentiment-aware applications, as compared to those that are already present in more traditional fact-based analysis. They include material on summarisation of evaluative text and on broader issues regarding privacy, manipulation, and economic impact that the development of opinion-oriented information-access services gives rise to.

To the inexperienced eye, opinion mining and sentiment analysis might look like the same as fields such as traditional text mining or fact-based analysis. Moreover, since sentiment classification deals with a relatively small number of classes, it might look like an easy task compared to text auto-categorisation. Opinion mining, however, is a very complex task even at its more basic level of sentiment polarity classification, which is a case of binary classification. The extraction of opinion polarity from text can be performed by comparing words extracted from text with a set of keywords

with positive valence (e.g., love, wonderful, best, great, superb, still, beautiful) and negative valence (e.g., bad, worst, stupid, waste, boring), as in the case of topic-based binary classification. The identification of a right set of keywords for mining opinions, however, is not a trivial task. Even when machine learning techniques are employed to select keywords from training corpora, the level of accuracy is still very low if compared to the performance of typical topic-based binary classification [220]. The main reason is that, differently from topics, sentiments can often be expressed in a more subtle manner, making it difficult to be identified by any of a sentence or document's terms when considered in isolation.

In addition, sentiment and subjectivity are quite context and domain dependent. This is true not only for changes in vocabulary, but also because even the exact same expression can indicate different sentiment in different domains. The concept 'go read the book', for example, most likely indicates positive sentiment for book reviews, but negative sentiment for movie reviews; as well as the adjective unpredictable may have a negative orientation in a car review (e.g., 'unpredictable steering'), but it could have a positive orientation in a movie review (e.g., 'unpredictable plot').

2.1.3 Sub-Tasks

One of the most common sub-tasks of opinion mining is polarity classification and the assignment of degrees of positivity, that is, given an opinionated piece of text wherein it is assumed that the overall opinion is about one single issue or item, classify the opinion as falling under one of two opposing sentiment polarities, or locate its position on the continuum between these two polarities. Much work on sentiment polarity classification has been conducted in the context of reviews of evaluative opinions (e.g., 'thumbs up' versus 'thumbs down', or 'like' versus 'dislike').

In addition, polarity classification can be also applied to identifying 'pro and con' expressions that can be used in individual reviews to evaluate the pros and cons that have influenced the judgements of a product and that make such judgements more trustworthy. Another instance of binary sentiment classification is agreement detection, that is, given a pair of text documents, deciding whether they should receive the same or differing sentiment-related labels. The more general problem of rating inference, where one must determine the author's evaluation with respect to a multi-point scale (e.g., one to five stars for a review) can be viewed as a multi-class text categorisation problem.

Other common sub-tasks of opinion mining and sentiment analysis are subjectivity detection and opinion identification. The capability of distinguishing if a text, or parts of it, are subjective or objective can be particularly beneficial for a more effective sentiment classification. Mihalcea et al. showed evidence that the complexity of this task is superior than subsequent polarity classification [195]. Wilson et al. remarked how classifying a piece of text as expressing a neutral opinion for rating inference does not equal classifying that piece of text as objective [305].

A piece of text can also have a polarity without necessarily containing an opinion, for example a news article can be classified into good or bad news without being subjective. The classification of a piece of text as subjective or objective can be useful in several situations. For example, being able to distinguish in opinionated texts where the authors do explicitly express their sentiment through statements (e.g., “this laptop is great”) and where they provide objective information (e.g., “the laptop has long battery life”) is used to help determine the overall sentiment. Hatzivassiloglou and Wiebe examined the effects of adjective orientation and gradability on sentence subjectivity to detect if a sentence is subjective [130] while other projects address subjectivity detection at sub-sentence level. Wiebe et al. presented a comprehensive survey of subjectivity recognition using different clues and features [301].

Typically, sentiment analysis is performed over an on-topic document, e.g., on the result of a topic-based search engine. However, several studies suggested that managing these two task jointly can be beneficial for the overall performance. According to Riloff et al., topic-based text filtering and subjectivity filtering are complementary, in the context of experiments in information extraction [243]. For example, off-topic passages of a document could contain irrelevant affective information and result misleading for the global sentiment polarity about the main topic. Also, a document can contain material on multiple topics that may be of interest to the user. In this case, it is therefore necessary to identify the topics and separate the opinions associated with each of them. Several other researches in sentiment analysis focus on non-topic based categorisation, for example to classify documents according to their genre [109] and their style [24]. Also authorship and publisher identification are other relevant examples [202, 13]. Another problem that has been considered in intelligence and security settings is the detection of deceptive language. Affect detection, eventually, is also a task that is gaining a growing attention from different perspectives and for different applications.

Sentiment analysis has been traditionally more focused on the extraction of the valence of textual sample (i.e., positive/negative or bad/good) rather than assigning a particular emotion category to text. However, the classification of multimedia resources according to their mood and emotional content is also quite common. The advent of Web 2.0 has pushed the users at the centre of the Web universe, providing them revolutionary tools that have changed the way people communicate and express themselves, their ideas, and emotions. People spend more and more time using the Web not only for work, but also for expressing their opinions on blogs and forums, chatting, and organising events through social networks, and even for living a Second Life⁷. Therefore, the Web contains more and more affective content. The awareness that the capability to manage such affective content can be exploited for the development of next-generation web applications is dragging a growing attention also in sentiment analysis for affect extraction from textual Web content.

⁷ <http://secondlife.com>

2.2 Main Approaches to Opinion Mining

Several approaches have been developed for the general task of mapping a given piece of text to a label belonging to a predefined set of categories, or to a real number representative of a polarity degree. Such approaches and their performance, however, are strictly bound to the considered domain of application and to the related topics. Moreover, most of the literature on sentiment analysis has focused on text written in English and consequently most of the resources developed, such as lexicons with sentiment labels, are in English. Adapting such resources to other languages can be considered as a domain adaptation problem [310, 46]. This section discusses the evolution of different approaches and depths of analysis [58], i.e., from heuristics to discourse structure (section 2.2.1), from coarse to fine grained analysis (section 2.2.2), from keyword to concept level opinion mining (section 2.2.3).

2.2.1 From Heuristics to Discourse Structure

Several unsupervised learning approaches rely on the creation of a sentiment lexicon in an unsupervised manner that is later used to determine the degree of positivity (or subjectivity) of a text unit. The crucial component is, therefore, the creation of the lexicon via the unsupervised labelling of words or phrases with their sentiment polarity or subjectivity [219]. This lexicon can be used to identify the *prior polarity* or the *prior subjectivity* of terms or phrases, to use towards further identifying contextual polarity or subjectivity. Early works were mainly based on linguistic heuristics. For example, Hatzivassiloglou and McKeown's technique [129] was built on the fact that, in the case of polarity classification, the two classes of interest represent opposites, and 'opposition constraints' can be used to help labelling decisions.

Other works propagated the valence of seed words, for which the polarity is known, to terms that co-occur with them in general text or in dictionary glosses, or to synonyms and words that co-occur with them in other WordNet-defined relations. A collective labelling approach can also be applied to opinion about product features. Popescu and Etzioni [232] proposed an iterative algorithm that, starting from a global word label computed over a large collection of generic topic text, gradually tried to re-define such label first to one that is specific to a review corpus then to one that is specific to a given product feature, and finally to one that is specific to the particular context in which the word occurs.

Also Snyder and Barzilay [261] exploited the idea of utilising discourse information to aid the inference of relationships between product attributes. They designed a linear classifier for predicting whether all aspects of a product are given the same rating, and combined such prediction with that of individual-aspect classifiers, in order to minimise a certain loss function. Regression techniques are often employed for the prediction of the degree of positivity in opinionated documents such as product reviews. Regression, in fact, allows to implicitly model similarity relationships between classes that correspond to points on a scale, such as the number of 'stars'

given by a reviewer [219]. Modelling discourse structure, such as twists and turns in documents, contributes to a more effective overall sentiment labelling. Early works attempted to partially address this problem via incorporating location information in the feature set [217]. More recent studies have underlined that position is particularly relevant in the context of sentiment summarisation. In particular, in contrast to topic-based text summarisation, where the incipits of articles usually serve as a strong baseline, the last n sentences of a review have been shown to serve as a much better summary of the overall sentiment of the document, and to be almost as good as the n (automatically-computed) most subjective sentences [217]. Joshi and Rose [147], for example, explored how features based on syntactic dependency relations can be utilised to improve performance on opinion mining. Using a transformation of dependency relation triples, they convert them into ‘composite back-off features’ that generalise better than the regular lexicalised dependency relation features.

2.2.2 From Coarse to Fine Grained

The evolution of research works in the field of opinion mining and sentiment analysis can be seen not only in the use of more and more sophisticated techniques, but also in the different depths of analysis adopted. Early works, in fact, aimed to classify entire documents as containing overall positive or negative polarity [220] or rating scores (e.g., 1-5 stars) of reviews [218]. These were mainly supervised approaches relying on manually labelled samples, such as movie or product reviews where the opinionist’s overall positive or negative attitude was explicitly indicated. However, opinions and sentiments do not occur only at document level, nor are they limited to a single valence or target. Contrary or complementary attitudes toward the same topic or multiple topics can be present across the span of a document. Later works adopted a segment level opinion analysis aiming to distinguish sentimental from non-sentimental sections, e.g., by using graph-based techniques for segmenting sections of a document on the basis of their subjectivity [217], or by performing a classification based on some fixed syntactic phrases that are likely to be used to express opinions [286], or by bootstrapping using a small set of seed opinion words and a knowledge base such as WordNet [148].

In recent works, text analysis granularity has been taken down to sentence level, e.g., by using presence of opinion-bearing lexical items (single words or n-grams) to detect subjective sentences [154, 242], or by using semantic frames defined in FrameNet [16] for identifying the topics (or targets) of sentiment [155], or by exploiting association rule mining [4] for a feature-based analysis of product reviews [141]. Commonly, a certain degree of continuity exists in subjectivity labels of adjacent sentences, as an author usually does not switch too frequently between being subjective and being objective. Hence, some works also propose a collective classification of the document based on assigning preferences for pairs of nearby sentences [218, 314].

All such approaches, however, are still far from being able to infer the cognitive and affective information associated with natural language as they mainly rely on semantic knowledge bases which are still too limited to efficiently process text at sentence level. Moreover, such a text analysis granularity might still not be enough as a single sentence may express more than one opinion [305].

2.2.3 *From Keywords to Concepts*

Existing approaches can be grouped into three main categories, with few exceptions: keyword spotting, lexical affinity, and statistical methods. Keyword spotting is the most naïve approach and probably also the most popular because of its accessibility and economy. Text is classified into affect categories based on the presence of fairly unambiguous affect words like ‘happy’, ‘sad’, ‘afraid’, and ‘bored’. Elliott’s Affective Reasoner [100], for example, watches for 198 affect keywords, e.g., ‘distressed’ and ‘enraged’, plus affect intensity modifiers, e.g., ‘extremely’, ‘somewhat’, and ‘mildly’, plus a handful of cue phrases, e.g., ‘did that’ and ‘wanted to’. Other popular sources of affect words are Ortony’s Affective Lexicon [212], which groups terms into affective categories, and Wiebe’s linguistic annotation scheme [302]. The weaknesses of this approach lie in two areas: poor recognition of affect when negation is involved and reliance on surface features. About its first weakness, while the approach can correctly classify the sentence “today was a happy day” as being happy, it is likely to fail on a sentence like “today wasn’t a happy day at all”. About its second weakness, the approach relies on the presence of obvious affect words which are only surface features of the prose.

In practice, a lot of sentences convey affect through underlying meaning rather than affect adjectives. For example, the text “My husband just filed for divorce and he wants to take custody of my children away from me” certainly evokes strong emotions, but uses no affect keywords, and therefore, cannot be classified using a keyword spotting approach. Lexical affinity is slightly more sophisticated than keyword spotting as, rather than simply detecting obvious affect words; it assigns arbitrary words a probabilistic ‘affinity’ for a particular emotion. For example, ‘accident’ might be assigned a 75% probability of being indicating a negative affect, as in ‘car accident’ or ‘hurt by accident’. These probabilities are usually trained from linguistic corpora [304, 272, 262, 237].

Though often outperforming pure keyword spotting, there are two main problems with the approach. First, lexical affinity, operating solely on the word-level, can easily be tricked by sentences like “I avoided an accident” (negation) and “I met my girlfriend by accident” (other word senses). Second, lexical affinity probabilities are often biased toward text of a particular genre, dictated by the source of the linguistic corpora. This makes it difficult to develop a reusable, domain-independent model. Statistical methods, such as latent semantic analysis (LSA) and support vector machine (SVM), have been popular for affect classification of texts and have been used by researchers on projects such as Goertzel’s Webmind [117], Pang’s movie review

classifier [220], and many others [218, 141, 287, 2, 291, 209, 92]. By feeding a machine learning algorithm a large training corpus of affectively annotated texts, it is possible for the systems to not only learn the affective valence of affect keywords as in the keyword spotting approach, but such a system can also take into account the valence of other arbitrary keywords (like lexical affinity), punctuation, and word co-occurrence frequencies.

However, statistical methods are generally semantically weak, meaning that, with the exception of obvious affect keywords, other lexical or co-occurrence elements in a statistical model have little predictive value individually. As a result, statistical text classifiers only work with acceptable accuracy when given a sufficiently large text input. So, while these methods may be able to affectively classify user's text on the page or paragraph level, they do not work well on smaller text units such as sentences.

2.3 Towards Machines with Common-Sense

Communication is one of the most important aspects of human life. Communicating has always a cost in terms of energy and time, since information needs to be encoded, transmitted, and decoded, and sometimes such factors can even make the difference between life and death. This is why people, when communicating with each other, provide just the useful information and take the rest for granted. This 'taken for granted' information is what we call common-sense – obvious things people normally know and usually leave unstated. Common-sense is not the kind of knowledge that we can find in Wikipedia⁸, but it consists in all the basic relationships among words, concepts, phrases, and thoughts that allow people to communicate with each other and face everyday life problems. It is a kind of knowledge that sounds obvious and natural to us, but it is actually daedal and multi-faceted.

The illusion of simplicity comes from the fact that, as each new group of skills matures, we build more layers on top of them and tend to forget about the previous layers. Common-sense, in fact, is not a simple thing. Instead, it is an immense society of hard-earned practical ideas, of multitudes of life-learned rules and exceptions, dispositions and tendencies, balances and checks [197]. This section discusses the importance of common-sense for the development of intelligent systems (section 2.3.1) and illustrates different knowledge representation strategies (section 2.3.2). The section also refers to a recently proposed survey on common-sense computing [49] to present the evolution of such research field, from logic-based approaches to more recent methods based on natural language techniques (section 2.3.3).

⁸ <http://wikipedia.org>

2.3.1 *The Importance of Common-Sense*

Concepts are the glue that holds our mental world together [205]. Without concepts, there would be no mental world in the first place [25]. Doubtless to say, the ability to organise knowledge into concepts is one of the defining characteristics of human mind. Of the different sorts of semantic knowledge that are researched, arguably the most general and widely applicable kind is knowledge about the everyday world that is possessed by all people, i.e., common-sense knowledge. While to the average person the term common-sense is regarded as synonymous with good judgement, to the AI community it is used in a technical sense to refer to the millions of basic facts and understandings possessed by most people, e.g., “a lemon is sour”, “to open a door, you must usually first turn the doorknob”, “if you forget someone’s birthday, they may be unhappy with you”.

Common-sense knowledge, thus defined, spans a huge portion of human experience, encompassing knowledge about the spatial, physical, social, temporal, and psychological aspects of typical everyday life. Because it is assumed that every person possesses common-sense, such knowledge is typically omitted from social communications, such as text. A full understanding of any text then, requires a surprising amount of common-sense, which currently only people possess. Common-sense knowledge is what we learn and what we are taught about the world we live in during our formative years, in order to better understand and interact with the people and the things around us. Common-sense is not universal, but cultural and context dependent. The importance of common-sense can be particularly appreciated when travelling to far away places, where sometimes it is necessary to almost entirely reset oneself’s common-sense knowledge in order to get integrated.

Despite the language barrier, in fact, moving to a new place involves facing habits and situations that might go against what we consider basic rules of social interaction or things we were taught by our parents, such as eating with hands, eating from someone else’s plate, slurping on noodle-like food or while drinking tea, eating on the street, crossing the road despite the heavy traffic, squatting when tired, removing shoes at home, growing long nails on your last fingers, or bargaining on anything you need to buy. This can happen also the other way around, that is, when you do something perfectly in line with your common-sense that violates the local norms, e.g., cheek kissing as a form of greeting.

Common-sense is the knowledge (usually acquired in early stages of our lives) concerning all the social, political, economic, and environmental aspects of the society we live in. Machines, as they never got the chance to live a life, have no common-sense at all and, hence, they know nothing about us. To help us work, computers must get to know what our jobs are. To entertain us, they need to know what we like. To take care of us, they have to know how we feel. To understand us, they must think as we think. Today, in fact, computers do only what they are programmed to do. They only have one way to deal with a problem and, if something goes wrong, they get stuck. Nowadays we have programs that exceed the capabilities of world experts, but are not one able to do what a three years old child can do. It is because machines have no goals, no hopes, no fears; they do not know the meaning of things.

Computers can only do logical things, but meaning is an intuitive process – it cannot be reduced to zeros and ones. We need to transmit to computers our common-sense knowledge of the world because soon there will not be enough human workers to perform the necessary tasks for our rapidly ageing population. To face this AI emergency⁹, we will have to give them physical knowledge of how objects behave, social knowledge of how people interact, sensory knowledge of how things look and taste, psychological knowledge about the way people think, and so on. But having a database of millions of common-sense facts will not be enough: we will also have to teach computers how to handle this knowledge, retrieve it when necessary, learn from experience - in a word, we will have to give them the capacity for common-sense reasoning.

2.3.2 *Knowledge Representation*

From its very beginning, AI has rested on a foundation of formal representation of knowledge. Knowledge representation (KR) is a research area that directly addresses languages for representation and the inferences that go along with them. One of the central questions of KR research is in what form knowledge is to be expressed. One of the most popular representation strategies is first order logic (FOL), a deductive system that consists of axioms and rules of inferences and can be used to formalise relationally rich predicates and quantification [20].

FOL supports syntax, semantics and, to a certain degree, pragmatics expressions. Syntax specifies the way groups of symbols are to be arranged, so that the group of symbols is considered properly formed. Semantics specify what well-formed expressions are supposed to mean. Pragmatics specifies how contextual information can be leveraged to provide better correlation between different semantics, for tasks such as word sense disambiguation. Logic, however, is known to have the problem of monotonicity. The set of entailed sentences can only increase as information is added to the knowledge base. This violates a common property of human reasoning, i.e., changing one's mind. Solutions such as default and linear logic serve to address parts of these issues. Default logic is proposed by Raymond Reiter to formalise default assumptions, e.g., "all birds fly" [240]. However, issues arise when default logic formalise facts that are true in the majority of cases, but not always, e.g., "penguins do not fly".

Linear logic, or constructive logic, was developed by Arend Heyting [138]. It is a symbolic logical system that preserves justification, rather than truth, and supports rejecting the weakening and contraction rules. It excels in careful deductive reasoning and is suitable in situations that can be posed precisely. As long as a scenario is static and can be detailedly described, in fact, situation-specific rules can perfectly model it but, when it comes to capture a dynamic and uncertain real-world environment, logical representation usually fails for lack of generalisation capabil-

⁹ <http://mitworld.mit.edu/video/484>

ities. Accordingly, it is not natural for human to encode knowledge in logical formalisation. Another standard KR strategy, based on FOL, is the use of relational databases. The idea is to describe a database as a collection of predicates over a finite set of variables and describing constraints on the possible values. Structured query language (SQL) [88] is the database language designed for the retrieval and management of data in relational database management systems (RDBMS) [77]. Commercial (e.g., Oracle¹⁰, Sybase¹¹, Microsoft SQL Server¹²) and open-source (e.g., MySQL¹³) implementations of RDBMS are available and they are commonly used in the IT industry.

Relational database design requires a strict process called normalisation to ensure that the relational database is suitable for general purpose querying and the relational database is free of database operations anomalies. Third normal form (3NF) [78] is stricter than first and second normal forms and less strict as compared to Boyce-Codd normal form (BCNF) [79], fourth, and fifth normal forms. Stricter normal forms means that the database design is more structured and, hence, requires more database tables. The advantage is that the overall design looks more organised. The disadvantage is the performance trade-off when joint table SQL queries are invoked. Relational database design, moreover, does not directly address representation of parent-child relationship in the object-oriented paradigm, subjective degrees of confidence, and temporal dependent knowledge.

A popular KR strategy, especially among Semantic Web researchers, is production rule [74]. A production rule system keeps a working memory of on-going memory assertions. This working memory is volatile and keeps a set of production rules. A production rule comprises an antecedent set of conditions and a consequent set of actions (i.e., IF <conditions> THEN <actions>). The basic operation for a production rule system involves a cycle of three steps ('recognise', 'resolve conflict', and 'act') that repeats until no more rules are applicable to working memory. The step 'recognise' identifies the rules whose antecedent conditions are satisfied by the current working memory. The set of rules identified is also called the conflict set. The step 'resolve conflict' looks into the conflict set and selects a set of suitable rules to execute. The step 'act' simply executes the actions and updates the working memory. Production rules are modular. Each rule is independent from others, allowing rules to be added and deleted easily.

Production rule systems have simple control structure and the rules are easy for human to understand. This is because rules are usually derived from observation of expert behaviour or expert knowledge, thus the terminology used in encoding the rules tend to resonate with human understanding. However, there are issues with scalability when production rule systems get larger. Significant maintenance overhead is required to maintain systems with thousands of rules.

¹⁰ <http://oracle.com>

¹¹ <http://sybase.com>

¹² <http://microsoft.com/sqlserver>

¹³ <http://mysql.com>

Another prominent KR strategy among Semantic Web researchers is the ontology web language (OWL)¹⁴, an XML-based vocabulary that extends resource description framework (RDF)¹⁵ and resource description framework schema (RDFS)¹⁶ to provide a more comprehensive ontology representation, such as the definition of classes, relationships between classes, properties of classes, and constraints on relationships between classes and properties of classes. RDF supports subject-predicate-object model that makes assertion about a resource. Reasoning engines have been developed to check for semantic consistency and help to improve ontology classification. OWL is a W3C recommended specification and comprises three dialects: OWL-Lite, OWL-DL, and OWL-Full. Each dialect has a different level of expressiveness and reasoning capabilities. OWL-Lite is the least expressive compared to OWL-Full and OWL-DL. It is suitable for building ontologies that only require classification hierarchy and simple constraints and, for this reason, it provides the most computationally efficient reasoning. OWL-DL is more expressive than OWL-Full, but more expressive than OWL-Lite. It has restrictions on the use of some of the description tags, hence, computation formed by a reasoning engine on OWL-DL ontologies can be completed in a finite amount of time [160]. OWL-DL is so named due to its correspondence with description logic. It is also the most commonly used dialect for representing domain ontology for Semantic Web applications. OWL-Full is the complete language and is useful for modelling a full representation of a domain. However, the trade-off for OWL-Full is the high complexity of the model that can result in sophisticated computation that may not complete in finite time. In general, OWL requires strict definition of static structures, hence, it is not suitable for representing knowledge that requires subjective degrees of confidence, but rather for representing declarative knowledge. OWL, moreover, does not allow to easily represent temporal dependent knowledge.

Another well-known way to represent knowledge is to use networks. Bayesian networks [226], for example, provide a means of expressing joint probability distributions over many interrelated hypotheses. Bayesian network is also called belief network. All variables are represented using directed acyclic graph (DAG). The nodes of a DAG represent variables. Arcs are causal connections between two variables where the truth of the former directly affects the truth of the latter. A Bayesian network is able to represent subjective degrees of confidence. The representation explicitly explores the role of prior knowledge and combines evidence of the likelihood of events. In order to compute the joint distribution of the belief network, there is a need to know $\Pr(P|\text{parents}(P))$ for each variable P . It is difficult to determine the probability of each variable P in the belief network. Hence, it is also difficult to scale and maintain the statistical table for large scale information processing problem. Bayesian networks also have limited expressiveness, which is only equivalent to the expressiveness of proposition logic. For this reason, semantic networks are more often used for KR (Fig. 2.1).

¹⁴ <http://w3.org/TR/owl-overview>

¹⁵ <http://w3.org/TR/PR-rdf-syntax>

¹⁶ <http://w3.org/2001/sw/wiki/RDFS>

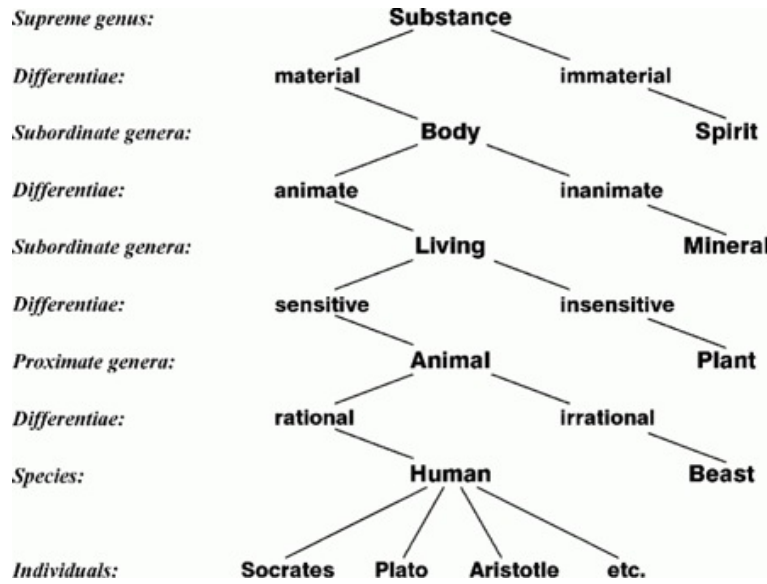


Fig. 2.1 Tree of Porphyry. Porphyry presented the basis of Aristotle's thought as a tree-like scheme of dichotomous divisions, in which the process continues until the lowest species is reached.

A semantic network [265] is a graphical notation for representing knowledge in patterns of interconnected nodes and arcs. There are six types of networks, namely definitional networks, assertional networks, implicational networks, executable networks, learning networks, and hybrid networks. A definitional network focuses on *IsA* relationships between a concept and a newly defined sub-type. The resulting network is called a generalisation, which supports the rule of inheritance for copying properties defined for a super-type to all of its sub-types. Definitions are true by definition and, hence, the information in definitional networks is often assumed to be true. Assertional networks are meant to assert propositions and the information is assumed to be contingently true. Contingent truth means that the proposition is true in some but not in all the worlds. The proposition also has sufficient reason in which the reason entails the proposition, e.g., "the stone is warm" with the sufficient reasons being "the sun is shining on the stone" and "whatever the sun shines on is warm". Contingent truth is not the same as the truth that is assumed in default logic. Contingent truth is closer to the truth assumed in model logic.

Implicational networks use implication as the primary relation for connecting nodes. They are used to represent patterns of beliefs, causality, or inferences. Methods for realising implicational networks include Bayesian networks and logic inferences used in a truth maintenance system (TMS). By combinations of forward and backward reasoning, a TMS propagates truth-values to nodes whose truth-value is unknown.

Executable networks contains mechanisms implemented in run-time environment such as message passing, attached procedure (e.g., data-flow graph), and graph transformation that can cause change to the network. Learning networks acquire knowledge from examples by adding and deleting nodes and links, or by modifying weights associated with the links. Learning networks can be modified in three ways: rote memory, changing weights, and restructuring. As for the rote memory, the idea is to add information without making changes to the current network. Exemplar methods can be found in relational database. For example, Patrick Winston used a version of relational graphs to describe structures, such as arches and towers [306]. When his program was given positive and negative examples of each type of structure, it would generalise the graphs to derive a definitional network for classifying all the types of structures that were considered. The idea of changing weights, in turn, is to modify the weights of links without changing the network structure for the nodes and links. Exemplar methods can be found in neural networks.

As for restructuring, finally, the idea is to create fundamental changes to the network structure for creative learning. Methods include case-based reasoning. The learning system uses rote memory to store various cases and associated action such as course of action. When a new case is encountered, the system finds those cases that are most similar to the new one and retrieves the outcome. To organise the search and evaluate similarity, the learning system must use restructuring to find common patterns in the individual cases and use those patterns as the keys for indexing the database. Hybrid networks combine two or more of the previous techniques. Hybrid networks can be a single network. They can also be separate but closely interacting networks.

Sowa used unified modelling language (UML) as an example to illustrate a hybrid semantic network. Semantic networks are very expressive. The representation is flexible and can be used to express different paradigm such as relational model and hierarchical relationship. The challenge is at implementation level. For example, it is difficult to implement hybrid semantic network, which requires an integration of different methods.

2.3.3 From Logical Inference to Digital Intuition

What magical trick makes us intelligent? - Marvin Minsky was wondering more than two decades ago - The trick is that there is no trick. The power of intelligence stems from our vast diversity, not from any single, perfect principle [197]. Human brain, in fact, is a very complex system, maybe the most complex in nature. The functions it performs are the product of thousands and thousands of different subsystems working together at the same time. Common-sense computing involves trying to emulate such mechanism and, in particular, exploiting common-sense knowledge to improve computers' understanding of the world. Before Minsky, many AI researchers started to think about the implementation of a common-sense reasoning machine.

The very first person who seriously started thinking about the creation of such a machine was perhaps Alan Turing when, in 1950, he first raised the question “can machines think?”. But he never managed to answer that question, he just provided a method to gauge artificial intelligence, the famous Turing test. The notion of common-sense in AI is actually dated 1958, when John McCarthy, in his seminal paper ‘Programs with Common Sense’ [192], proposed a program, termed the ‘advice taker’, for solving problems by manipulating sentences in formal language. The main aim of such a program was to try to automatically deduce for itself a sufficiently wide class of immediate consequences of anything it was told and what it already knew. In this paper, McCarthy stressed the importance of finding a proper method of representing expressions in the computer since, in order for a program to be capable of learning something, it must first be capable of being told. He also developed the idea of creating a property list for each object, in which the specific things people usually know about that object are listed. It was the first attempt to build a common-sense knowledge base but, more important, it was the epiphany of the need of common-sense to move forward in the technological evolution.

In 1959, McCarthy went to MIT and started, together with Minsky, the MIT Artificial Intelligence Project. They both were aware of the need for AI of a common-sense reasoning approach, but while McCarthy was more concerned with establishing logical and mathematical foundations for it, Minsky was more involved with theories of how we actually reason using pattern recognition and analogy. These theories were organised some years later with the publication of the Society of Mind [197], a masterpiece of AI literature, which consists in an illuminating vision of how the human brain might work. Minsky sees the mind made of many little parts, termed ‘agents’, each mindless by itself but able to lead to true intelligence when working together. These groups of agents, called ‘agencies’, are responsible to perform some type of function, such as remembering, comparing, generalising, exemplifying, analogising, simplifying, predicting, and so on.

The most common agents are the so called ‘K-lines’, whose task is simply to activate other agents: this is a very important issue since agents are all highly interconnected and activating a K-line can cause a significant cascade of effects. To Minsky, in fact, mental activity ultimately consists in turning individual agents on and off: at any time only some agents are active and their combined activity constitutes the ‘total state’ of the mind. K-lines are a very simple but powerful mechanism since they allow entering a particular configuration of agents that formed a useful society in a past situation. This is how we build and retrieve our problem solving strategies in our mind; this is how we should develop our problem solving strategies in our programs.

In 1990, McCarthy put together seventeen papers to try to define common-sense knowledge by using mathematical logic in such a way that common-sense problems could be solved by logical reasoning. Deductive reasoning in mathematical logic has the so-called monotonicity property: if we add new assumptions to the set of initial assumptions, there may be some new conclusions, but every sentence that was a deductive consequence of the original hypotheses is still a consequence of the enlarged set.

Much human reasoning is monotonic as well, but some important human common-sense reasoning is not. For example, if someone is asked to build a birdcage, this person concludes that it is appropriate to put a top on it, but when he/she learns the further fact that the bird is a penguin he/she no longer draws that conclusion. McCarthy formally described this assumption that things are as expected unless otherwise specified, with the ‘circumscription method’ of non-monotonic reasoning: a minimisation similar to the closed world assumption that what is not known to be true is false. In the same years, a similar attempt to give a shape to common-sense knowledge was done by Ernest Davis [103]. He tried to develop an ad hoc language for expressing common-sense knowledge and inference techniques for carrying out common-sense reasoning in specific domains such as space, time, quantities, qualities, flows, goals, plans, needs, beliefs, intentions, actions, and interpersonal relations. Thanks to his and McCarthy’s knowledge formalisations, the first steps were set towards the expression of common-sense facts in a way that would have been suitable for inclusion in a general purpose database and, hence, towards the development of programs with common-sense.

Minsky’s theory of human cognition, in particular, was welcomed with great enthusiasm by the AI community and gave birth to many attempts to build common-sense knowledge bases and develop systems capable of common-sense reasoning. The most representative projects are Cyc [170], Doug Lenat’s logic-based repository of common-sense knowledge, WordNet [108], Christiane Fellbaum’s universal database of word senses, and ThoughtTreasure [203], Erik Mueller’s story understanding system. Cyc is one of the first attempts to assemble a massive knowledge base spanning human common-sense knowledge.

Initially started by Doug Lenat in 1984, this project utilises knowledge engineers who hand-craft assertions and place them into a logical framework using CycL, Cyc’s proprietary language. Cyc’s knowledge is represented redundantly at two levels: a frame language distinction (epistemological level), adopted for its efficiency, and a predicate calculus representation (heuristic level), needed for its expressive power to represent constraints. While the first level keeps a copy of the facts in the uniform user language, the second level keeps its own copy in different languages and data structures suitable to be manipulated by specialised inference engines. Knowledge in Cyc is also organised into ‘microtheories’, resembling Minsky’s agencies, each one with its own knowledge representation scheme and sets of assumptions. These microtheories are linked via ‘lifting rules’ that allow translation and communication of expressions between them.

Begun in 1985 at Princeton University, WordNet is a database of words (primarily nouns, verbs, and adjectives). It has been one of the most widely used resources in computational linguistics and text analysis for the ease in interfacing it with any kind of application and system. The smallest unit in WordNet is the word/sense pair, identified by a ‘sense key’. Word/sense pairs are linked by a small set of semantic relations such as synonyms, antonyms, *IsA* superclasses, and words connected by other relations such as *PartOf*. Each synonym set, in particular, is called ‘synset’: it consists in the representation of a concept, often explained through a brief gloss, and represents the basic building block for hierarchies and other conceptual struc-

tures in WordNet. Erik Mueller's ThoughtTreasure is a story understanding system with a great variety of common-sense knowledge about how to read and understand children's stories. It was inspired by Cyc and is similar to Cyc in that it has both natural language and common-sense components. But whereas Cyc mostly uses logic, ThoughtTreasure uses multiple representations schemes: grids for stereotypical settings, finite automata for rules of device behaviour and mental processes, logical assertions for encyclopaedic facts and linguistic knowledge. ThoughtTreasure's lexicon is similar to WordNet but, while world knowledge is explicitly excluded from WordNet, ThoughtTreasure contains also concepts that are not lexicalised in English like 'going to the pub' or 'eating at the restaurant', which are very important for common-sense reasoning.

Using logic-based reasoning, in fact, can solve some problems in computer programming, but most real-world problems need methods better at matching patterns and constructing analogies, or making decisions based on previous experience with examples, or by generalising from types of explanations that have worked well on similar problems in the past [198]. In building intelligent systems we have to try to reproduce our way of thinking: we turn ideas around in our mind to examine them from different perspectives until we find one that works for us. From this arises the need of using several representations, each integrated with its set of related pieces of knowledge, to be able to switch from one to another when one of them fails. The key, in fact, is using different representations to describe the same situation.

Minsky blames our standard approach to writing a program for common-sense computing failures. Since computers appeared, our approach to solve a problem has always consisted in first looking for the best way to represent the problem, and then looking for the best way to represent the knowledge needed to solve it and finally looking for the best procedure for solving it. This problem-solving approach is good when we have to deal with a specific problem, but there is something basically wrong with it: it leads us to write only specialised programs that cope with solving only that kind of problem. This is why, today, we have millions of expert programs but not even one that can be actually defined intelligent.

From here comes the idea of finding a heterogeneous ways to represent common-sense knowledge and to link each unit of knowledge to the uses, goals, or functions that each knowledge-unit can serve. This non-monotonic approach reasserted by Minsky was adopted soon after by Push Singh within the Open Mind Common Sense (OMCS) project [257]. Initially born from an idea of David Stork [273], the project differs from previous attempts to build a common-sense database for the innovative way to collect knowledge and represent it (Fig. 2.2). OMCS is a second-generation common-sense database. Knowledge is represented in natural language, rather than using a formal logical structure, and information is not hand-crafted by expert engineers but spontaneously inserted by online volunteers. The reason why Lenat decided to develop an ad hoc language for Cyc is that vagueness and ambiguity pervade English and computer reasoning systems generally require knowledge to be expressed accurately and precisely. However, as expressed in the Society of Mind, ambiguity is unavoidable when trying to represent the common-sense world.




Fig. 2.2 An Open Mind project on OCR: handwritten characters are presented to e-citizens whose judgements (here 4 versus 9) are returned to the Open Mind host and used to train the classifier.

No single argument, in fact, is always completely reliable but, if we combine multiple types of arguments, we can improve the robustness of reasoning as well as we can improve a table stability by providing it with many small legs in place of just one very big leg. This way information is not only more reliable, but also stronger. If a piece of information goes lost, we can still access the whole meaning, exactly as the table keeps on standing up if we cut out one of the small legs. Diversity is, in fact, the key of OMCS' success: the problem is not choosing a representation in spite of another, but it is finding a way for them to work together in one system. The main difference between acquiring knowledge from the general public and acquiring it from expert engineers is that the general public is likely to leave as soon as they encounter something boring or difficult. The key is letting people do what they prefer to do. Different people, in fact, like to do different things: some like to enter new items, some like to evaluate items, others like to refine items. For this reason, OMCS is based on a distributed workflow model where the different stages of knowledge acquisition could be performed separately by different participants.

The system, in fact, was designed to allow users to insert new knowledge via both template-based input and free-form input, tag concepts, clarify properties, and validate assertions. But, since giving so much control to users can be dangerous, a fixed set of pre-validated sentences were meant to be presented to them from time to time, in order to assess their honesty, and the system was designed in a way that allowed users to reciprocally control each other by judging samples of each other's knowledge. OMCS exploits a method termed cumulative analogy [73], a class of analogy-based reasoning algorithms that leverage existing knowledge to pose knowledge acquisition questions to the volunteer contributors. When acquiring knowledge online, the stickiness of the website is of primary importance. The best way to involve users in this case is making them feel that they are contributing to the construction of a thinking machine and not just a static database. To do this, OMCS first determines what other topics are similar to the topic the user is currently inserting knowledge for, and then it uses cumulative analogy to generate and present

new specific questions about this topic. Because each statement consists of an object and a property, the entire knowledge repository can be visualised as a large matrix, with every known object of some statement being a row and every known property being a column. Cumulative analogy is performed by first selecting a set of nearest neighbours, in terms of similarity, of the treated concept and then by projecting known properties of this set onto not known properties of the concept and presenting them as questions (Fig. 2.3). The replies to the knowledge acquisition questions formulated by analogy are immediately added to the knowledge repository, affecting the similarity calculations. This way users can see the system’s behaviour improve or change as a result of the entered knowledge and be more tempted to participate.

A more generalised way to deal with the information contained in the Open Mind corpus is AnalogySpace [268], a process that applies singular value decomposition (SVD) on the matrix representation of the common-sense knowledge base, in order to reduce its dimensionality and capture the most important correlations. The entries in the resulting matrix are positive or negative numbers, depending on the reliability of the assertions, and their magnitude increases logarithmically with the confidence score. Applying SVD on this matrix causes it to describe other features that could apply to known concepts by analogy: if a concept in the matrix has no value specified for a feature owned by many similar concepts, then by analogy the concept is likely to have that feature as well.

A way to visualise and understand AnalogySpace is provided by Luminoso [269], a tool that allows to interactively explore the dimensionality-reduced semantic space of common-sense knowledge by ‘grabbing’ its data points and, hence, view their associated text and statistics. This operation also allows highlighting the point’s neighbourhood of semantically associated data points, providing clues for reasons as to why the points were classified along the dimensions they were.

Objects	Properties <i>(with simplified form)</i>			
	... contains knowledge <i>contain knowledge</i>	has pages <i>have page</i>	is cold <i>be cold</i>	is for reading <i>be read</i>
⋮	⋮	⋮	⋮	⋮
book	x	x		x
ice		-	x	
newspaper	x?	x		x
magazine	x	x		x
⋮	⋮	⋮	⋮	⋮

Fig. 2.3 The cumulative analogy process allows to perform comparisons between concepts in a knowledge base (represented as a matrix) and, hence, to infer new information about similar concepts.

The AnalogySpace process, eventually, is naturally extended by the ‘blending’ technique [134], a new method to perform inference over multiple sources of data simultaneously, taking advantage of the overlap between them. Blending consists in an alignment phase of two datasets and of a linear combination of them to be able to apply principal component analysis (PCA) on the resulting matrix. This enables common-sense to be used as a basis for inference in a wide variety of systems and applications so that they can achieve digital intuition about their own data, making assumptions and conclusions based on the connections between that specific data and the general common-sense that people have.

2.4 Conclusions

This chapter has shown how and why, today, online opinions are crucial both for companies to succeed in selling their products and services, and for individuals to properly choose among these. The dynamics behind the buzz mechanism were discussed, together with the motivating factors that gave birth to the field of opinion mining, and the tasks that make it different from standard information retrieval (section 2.1). The chapter also illustrated the approaches and depths of analysis in mining and characterising opinions, in order to map a given piece of text to a label belonging to a predefined set of categories, or to a real number representative of a polarity degree.

Specifically, the chapter discussed the evolution of different approaches from heuristics to discourse structure, from coarse to fine grained analysis, and from keyword to concept level opinion mining (section 2.2). Eventually, the chapter explained the importance of common-sense for the development of intelligent systems, illustrated different knowledge representation strategies, and presented the evolution of common-sense computing from logic-based methods to more recent approaches based on natural language techniques (section 2.3).

Chapter 3

Techniques

*A perfect intelligence would not confine itself to one order of thought,
but would simultaneously regard a group of objects
as classified in all the ways of which they are capable.*

Stanley Jevons

Providing a machine with physical knowledge of how objects behave, social knowledge of how people interact, sensory knowledge of how things look and taste and psychological knowledge about the way people think, is not enough to make it intelligent. Having a database of millions of concepts is not very useful for a computer, unless it is able to conveniently use such knowledge base. Our ability to use common-sense knowledge, in fact, highly depends on being able to do common-sense reasoning. Machines need to be taught not just common-sense knowledge itself but also strategies for handling it, retrieving it when necessary, and learning from experience.

To this end, adequately broad and deep common-sense knowledge bases are to be developed, as well as reasoning methods that exhibit the features of human thinking, including the ability to reason with knowledge that is true by default, reason rapidly across a broad range of domains, tolerate uncertainty in the available knowledge, take decisions under incomplete knowledge, and perhaps revise that belief or decision when complete knowledge becomes available. It is also important to develop new kinds of cognitive architectures able to support multiple reasoning methods and representations. If a machine is able to represent knowledge and perform reasoning in many different ways, in fact, it can switch among different points of view and find one that works, rather than getting stuck when something goes wrong.

This chapter shows how to build a semantic resource of affective common-sense knowledge and how to exploit it for enabling emotion-sensitive inference (section 3.1). In order to accordingly classify affect in such a knowledge base, moreover, a novel emotion categorisation model (section 3.2) and a new clustering method (section 3.3) are developed. In addition, a technique that combines dimensional-

ity reduction and graph mining methods on two different reasoning levels (section 3.4) is proposed, together with a technique that mimics the human capability to switch between different perspectives (section 3.5) and some concluding remarks (section 3.6).

3.1 Affective Blending: Enabling Emotion-Sensitive Inference

Opinions and sentiments are often conveyed implicitly through context and domain dependent concepts, rather than through specific affect words. Hence, in order to semantically and affectively analyse natural language text for opinion mining, we need to rely on a knowledge base containing both the cognitive and affective information (semantics and sentics) associated with concepts.

To this end, a semantic network of common-sense knowledge is merged with a linguistic resource for the lexical representation of affective knowledge, in order to obtain a new knowledge base in which concepts are interrelated by both common-sense and emotional features (section 3.1.1). Such a knowledge base is then exploited for the development of a novel affective common-sense knowledge visualisation and analysis system (section 3.1.2).

3.1.1 *AffectNet*

The affective common-sense knowledge base developed within this research work is built upon ConceptNet [181], the graph representation of the Open Mind corpus, which is structurally similar to WordNet, but whose scope of contents is general world knowledge, in the same vein as Cyc. Instead of insisting on formalising common-sense reasoning using mathematical logic [204], ConceptNet uses a new approach: it represents data in the form of a semantic network and makes it available to be used in natural language processing. The prerogative of ConceptNet, in fact, is contextual common-sense reasoning: while WordNet is optimised for lexical categorisation and word-similarity determination, and Cyc is optimised for formalised logical reasoning, ConceptNet is optimised for making practical context-based inferences over real-world texts.

In ConceptNet, WordNet's notion of node in the semantic network is extended from purely lexical items (words and simple phrases with atomic meaning) to include higher-order compound concepts, e.g., 'satisfy hunger' and 'follow recipe', to represent knowledge around a greater range of concepts found in everyday life (Table 3.1). Moreover WordNet's repertoire of semantic relations is extended from the triplet of synonym, *IsA* and *PartOf*, to a repertoire of twenty semantic relations including, for example, *EffectOf* (causality), *SubeventOf* (event hierarchy), *CapableOf* (agent's ability), *MotivationOf* (affect), *PropertyOf*, and *LocationOf*. ConceptNet's knowledge is also of a more informal, defeasible, and practically valued nature.

Term	WordNet Hypernyms	ConceptNet Assertions
Cat	Feline; Felid; Adult male; Man; Gossip; Gossiper; Gossipmonger; Rumormonger; Rumourmonger; Newsmonger; Woman; Adult female; Stimulant; Stimulant drug; Excitant; Tracked vehicle; ...	Cats can hunt mice; Cats have whiskers; Cats can eat mice; Cats have fur; cats have claws; Cats can eat meat; cats are cute; ...
Dog	Canine; Canid; Disagreeable woman; Chap; Fellow; Feller; Lad; Gent; Fella; Scoundrel; Sausage; Follow, ...	Dogs are mammals; A dog can be a pet; A dog can guard a house; You are likely to find a dog in kennel; An activity a dog can do is run; A dog is a loyal friend; A dog has fur; ...
Language	Communication; Auditory communication; Word; Higher cognitive process; Faculty; Mental faculty; Module; Text; Textual matter; ...	English is a language; French is a language; Language is used for communication; Music is a language; A word is part of language; ...
iPhone	N/A;	An iPhone is a kind of telephone; An iPhone is a kind of computer; An iPhone can display your position on a map; An iPhone can send and receive emails; An iPhone can display the time; ...
Birthday gift	Present;	Card is birthday gift; Present is birthday gift; Buying something for a loved one is for a birthday gift; ...

Table 3.1 Comparison between WordNet and ConceptNet. While WordNet synsets contain vocabulary knowledge, ConceptNet assertions convey knowledge about what concepts are used for.

For example, WordNet has formal taxonomic knowledge that ‘dog’ is a ‘canine’, which is a ‘carnivore’, which is a ‘placental mammal’; but it cannot make the practically oriented member-to-set association that ‘dog’ is a ‘pet’. ConceptNet also contains a lot of knowledge that is defeasible, i.e., it describes something that is often true but not always, e.g., EffectOf(‘fall off bicycle’, ‘get hurt’), which is something we cannot leave aside in common-sense reasoning. Most of the facts interrelating ConceptNet’s semantic network are dedicated to making rather generic connections between concepts.

This type of knowledge can be brought back to Minsky’s K-lines, as it increases the connectivity of the semantic network and makes it more likely that concepts parsed out of a text document can be mapped into ConceptNet. ConceptNet is produced by an automatic process, which first applies a set of extraction rules to the semi-structured English sentences of the OMCS corpus, and then applies an additional set of ‘relaxation’ procedures, i.e., filling in and smoothing over network gaps, to optimise the connectivity of the semantic network (Fig. 3.1).

In ConceptNet 2, a new system for weighting knowledge was implemented, which scores each binary assertion based on how many times it was uttered in the OMCS corpus, and on how well it can be inferred indirectly from other facts in ConceptNet. In ConceptNet 3 [131], users can also participate in the process of refining knowledge by evaluating existing statements on Open Mind Commons [266], the new interface for collecting common-sense knowledge from users over the Web.

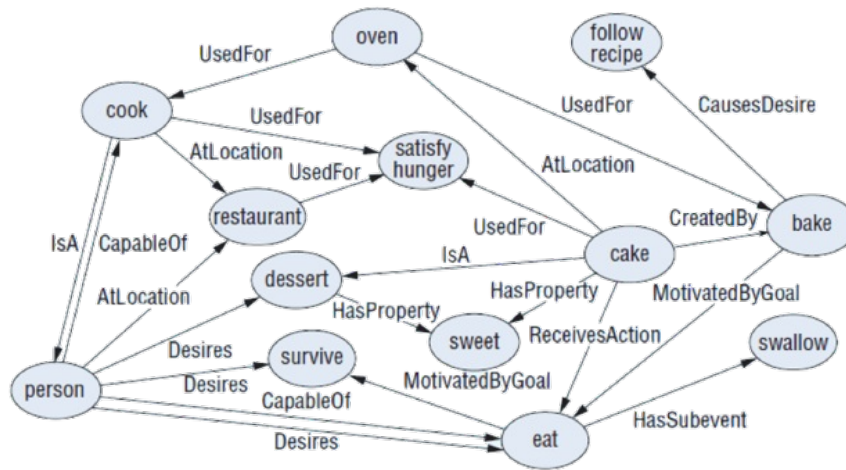


Fig. 3.1 ConceptNet represents the information in the Open Mind corpus as a directed graph where nodes are concepts and labelled edges are assertions of common-sense that interconnect them.

By giving the user many forms of feedback and using inferences by analogy to find appropriate questions to ask, Open Mind Commons can learn well-connected structures of common-sense knowledge, refine its existing knowledge, and build analogies that lead to even more powerful inferences. ConceptNet 4 includes data that was imported from the online game Verbosity. It also includes the initial import of the Chinese ConceptNet. ConceptNet 5 [267], eventually, contains knowledge from English Wikipedia, specifically from DBpedia, which extracts knowledge from the info-boxes that appear on articles, and ReVerb, a machine-reading project extracting relational knowledge from the actual text of each article. It also includes a large amount of content from the English Wiktionary, including synonyms, antonyms, translations of concepts into hundreds of languages, and multiple labelled word senses for many English words. ConceptNet 5 contains more dictionary-style knowledge coming from WordNet and some knowledge about people’s intuitive word associations coming from games with a purpose (GWAP). Previous versions of ConceptNet have been distributed as idiosyncratic database structures plus some software to interact with them. ConceptNet 5 is not a piece of software or a database, but a hypergraph, that is, a graph that has edges about edges. Each statement in ConceptNet, in fact, has justifications pointing to it, explaining where it comes from and how reliable the information seems to be.

In Chinese culture (and many others), the concepts of ‘heart’ and ‘mind’ used to be expressed by the same word (心) as it was believed that consciousness and thoughts came from the cardiac muscle. In human cognition, in fact, thinking and feeling are mutually present: emotions are often the product of our thoughts, as well

as our reflections are often the product of our affective states. Emotions are intrinsically part of our mental activity and play a key role in communication and decision-making processes. Emotion is a chain of events made up of feedback loops. Feelings and behaviour can affect cognition, just as cognition can influence feeling. Emotion, cognition, and action interact in feedback loops and emotion can be viewed in a structural model tied to adaptation [229]. There is actually no fundamental opposition between emotion and reason. In fact, it may be argued that reason consists of basing choices on the perspectives of emotions at some later time. Reason dictates not giving in to one's impulses because doing so may cause greater suffering later [114]. Reason does not necessarily imply exertion of the voluntary capacities to suppress emotion. It does not necessarily involve depriving certain aspects of reality of their emotive powers.

On the contrary, our voluntary capacities allow us to draw more of reality into the sphere of emotion. They allow one's emotions to be elicited not merely by the proximal, or the perceptual, or that which directly interferes with one's actions, but by that which, in fact, touches on one's concerns, whether proximal or distal, whether occurring now or in the future, whether interfering with one's own life or that of others. Cognitive functions serve emotions and biological needs. Information from the environment is evaluated in terms of its ability to satisfy or frustrate needs. What is particularly significant is that each new cognitive experience that is biologically important is connected with an emotional reaction such as fear, pleasure, pain, disgust, or depression [208].

Emotions, in fact, are special states shaped by natural selection to adjust various aspects of our organism in order to make it better face particular situations, e.g., anger evolved for reaction, fear evolved for protection, and affection evolved for reproduction. For these reasons, we cannot prescind from emotions in the development of intelligent systems: if we want computers to be really intelligent, not just have the veneer of intelligence, we need to give them the ability to recognise, understand, and express emotions. To this end, it is useful to build a knowledge base that contains not only common-sense concepts, but also the affective information associated with these.

ConceptNet is a good source of common-sense knowledge but alone is not enough for sentiment analysis tasks as it specifies how concepts are semantically related to each other but often lacks connections between concepts that convey the same kind of emotion or similar polarity. To overcome such a hurdle, WordNet-Affect (WNA) [275], a linguistic resource for the lexical representation of affective knowledge developed starting from WordNet, is used. WNA is built by assigning to a number of WordNet synsets one or more affective labels (a-labels). In particular, the affective concepts representing emotional states are identified by synsets marked with the a-label 'emotion', but there are also other a-labels for concepts representing moods, situations eliciting emotions, or emotional responses. WNA was developed in two stages. The first consisted of the identification of a first core of affective synsets. The second step consisted of the extension of the core with the relations defined in WordNet.

Mood	LiveJournal Posts	Concepts
Happy	Finally I got my student cap ! I am officially high school graduate now ! Our dog Tanja, me, Timo (our art teacher) and EmmaMe, Tanja, Emma and Tiia Only two weeks to Japan!!	student; school graduate; Japan
Happy	I got a kitten as an early birthday gift on Monday. Abby was smelly, dirty, and knawing on the metal bars of the kitten carrier though somewhat calm when I picked her up. We took her. She threw up on me on the ride home and repeatly keeps sneesing in my face.	kitten; birthday gift; metal bar; face
Sad	Hi. Can I ask a favor from you? This will only take a minute. Please pray for Marie, my friends' dog a labrador, for she has canine distemper. Her lower half is paralysed and she's having locked jaw. My friends' family is feeding her through syringe.	friends; dog; labrador; canine distemper; jaw; syringe
Sad	my uncle paul passed away on february 16, 2008. he lost his battle with cancer. i remember spending time with him and my aunt nina when they babysat me. we would go to taco bell and i would get nachos.	uncle; battle; cancer; aunt; taco bell; nachos

Table 3.2 Some examples of LiveJournal posts where affective information is not conveyed explicitly through affect words. Such implicit information can be inferred by analysing the semantics and sentics.

ConceptNet and WNA are blended together by combining the matrix representations of the two knowledge bases linearly into a single matrix, in which the information between the two initial sources is shared. The first step to create the affective blend is to transform the input data so that it can all be represented in the same matrix. To do this, the lemma forms of ConceptNet concepts are aligned with the lemma forms of the words in WNA and the most common relations in the affective knowledge base are mapped into ConceptNet's set of relations, e.g., *Hypernym* into *IsA* and *Holonym* into *PartOf*. In particular, ConceptNet is first converted into a matrix by dividing each assertion into two parts: a concept and a feature, where a feature is simply the assertion with the first or the second concept left unspecified such as 'a wheel is part of' or 'is a kind of liquid'.

The entries in the resulting matrix are positive or negative numbers, depending on the reliability of the assertions, and their magnitude increases logarithmically with the confidence score. WNA, similarly, is represented as a matrix where rows are affective concepts and columns are features related to these. The result of aligning the matrix representations of ConceptNet and WNA is a new affective semantic network, in which common-sense concepts are linked to a hierarchy of affective domain labels. In such a semantic network, termed *AffectNet*¹ [50], common-sense and affective knowledge are in fact combined, not just concomitant, i.e., everyday life concepts like 'have breakfast', 'meet people', or 'watch tv' are linked to affective domain labels like 'joy', 'anger', or 'surprise'. Such knowledge base results very useful when performing tasks such as emotion recognition or polarity detection from natural language text, as opinions and sentiments are often conveyed implicitly through context and domain dependent concepts, rather than through specific affect words (Table 3.2).

¹ <http://sentic.net/affectnet.zip>

3.1.2 *AffectiveSpace*

The best way to solve a problem is to already know a solution for it. But, if we have to face a problem we have never met before, we need to use our intuition. Intuition can be explained as the process of making analogies between the current problem and the ones solved in the past to find a suitable solution. Marvin Minsky attributes this property to the so called ‘difference-engines’ [197]. This particular kind of agents operates by recognising differences between the current state and the desired state, and acting to reduce each difference by invoking K-lines that turn on suitable solution methods. This kind of thinking is maybe the essence of our supreme intelligence since in everyday life no two situations are ever the same and have to perform this action continuously.

To emulate such a process, *AffectiveSpace*² [44], a novel affective common-sense knowledge visualisation and analysis system, is used. Human mind constructs intelligible meanings by continuously compressing over vital relations [106]. The compression principles aim to transform diffuse and distended conceptual structures to more focused versions so as to become more congenial for human understanding. To this end, principal component analysis (PCA) has been applied on the matrix representation of AffectNet. In particular, truncated singular value decomposition (TSVD) has been preferred to other dimensionality reduction techniques for its simplicity, relatively low computational cost, and compactness. TSVD, in fact, is particularly suitable for measuring the cross-correlations between affective common-sense concepts as it uses an orthogonal transformation to convert the set of possibly correlated common-sense features associated with each concept into a set of values of uncorrelated variables (the principal components of the SVD). By using Lanczos’ method [162], moreover, the generalisation process is relatively fast (a few seconds), despite the size and the sparseness of AffectNet. As the dimensions of such a matrix grow, however, PCA might cease to be a good solution in the future. To this end, different techniques, e.g., independent component analysis (ICA), random projections, and non-negative matrix factorisation (NMF) are being investigated.

At the present time, TSVD is applied over the concept-feature matrix in order to conveniently reduce its dimensionality and capture the most important correlations. The objective of such compression is to allow many details in the blend of ConceptNet and WNA to be removed such that the blend only consists of a few essential features that represent the global picture. Applying TSVD on AffectNet, in fact, causes it to describe other features that could apply to known affective concepts by analogy: if a concept in the matrix has no value specified for a feature owned by many similar concepts, then by analogy the concept is likely to have that feature as well. In other words, concepts and features that point in similar directions and, therefore, have high dot products, are good candidates for analogies. A pioneering work on understanding and visualising the affective information associated to natural language text was conducted by Osgood et al. [213].

² <http://sentic.net/affectivespace.zip>

Osgood used multi-dimensional scaling (MDS) to create visualisations of affective words based on similarity ratings of the words provided to subjects from different cultures. Words can be thought of as points in a multi-dimensional space and the similarity ratings represent the distances between these words. MDS projects these distances to points in a smaller dimensional space. Similarly, AffectiveSpace aims to grasp the semantic and affective similarity between different concepts by plotting them into a multi-dimensional vector space. Differently from Osgood’s space, however, the building blocks of AffectiveSpace are not simply a limited set of similarity ratings between affect words, but rather millions of confidence scores related to pieces of common-sense knowledge linked to a hierarchy of affective domain labels. Rather than merely determined by a few human annotators and represented as a word-word matrix, in fact, AffectiveSpace is built upon an affective common-sense knowledge base, namely AffectNet, represented as a concept-feature matrix. After performing TSVD on such matrix, hereby termed A for the sake of conciseness, a low-rank approximation of it is obtained, that is, a new matrix $\tilde{A} = U_k \Sigma_k V_k^T$. This approximation is based on minimising the Frobenius norm of the difference between A and \tilde{A} under the constraint $\text{rank}(\tilde{A}) = k$. For the Eckart–Young theorem [98], it represents the best approximation of A in the least-square sense, in fact:

$$\min_{\tilde{A} | \text{rank}(\tilde{A})=k} |A - \tilde{A}| = \min_{\tilde{A} | \text{rank}(\tilde{A})=k} |\Sigma - U^* \tilde{A} V| = \min_{\tilde{A} | \text{rank}(\tilde{A})=k} |\Sigma - S| \quad (3.1)$$

assuming that \tilde{A} has the form $\tilde{A} = USV^*$, where S is diagonal. From the rank constraint, i.e., S has k non-zero diagonal entries, the minimum of the above statement is obtained as follows:

$$\min_{\tilde{A} | \text{rank}(\tilde{A})=k} |\Sigma - S| = \min_{s_i} \sqrt{\sum_{i=1}^n (\sigma_i - s_i)^2} \quad (3.2)$$

$$\min_{s_i} \sqrt{\sum_{i=1}^n (\sigma_i - s_i)^2} = \min_{s_i} \sqrt{\sum_{i=1}^k (\sigma_i - s_i)^2 + \sum_{i=k+1}^n \sigma_i^2} = \sqrt{\sum_{i=k+1}^n \sigma_i^2} \quad (3.3)$$

Therefore, \tilde{A} of rank k is the best approximation of A in the Frobenius norm sense when $\sigma_i = s_i$ ($i = 1, \dots, k$) and the corresponding singular vectors are the same as those of A . If all but the first k principal components are discarded, common-sense concepts and emotions are represented by vectors of k coordinates. These coordinates can be seen as describing concepts in terms of ‘eigenmoods’ that form the axes of AffectiveSpace, i.e., the basis e_0, \dots, e_{k-1} of the vector space (Fig. 3.2). For example, the most significant eigenmood, e_0 , represents concepts with positive affective valence. That is, the larger a concept’s component in the e_0 direction is, the more affectively positive it is likely to be. Concepts with negative e_0 components, then, are likely to have negative affective valence. Thus, by exploiting the information sharing property of TSVD, concepts with the same affective valence are likely to have similar features – that is, concepts conveying the same emotion tend to fall near each other in AffectiveSpace.

The bigger is k , the more precisely AffectiveSpace represents AffectNet’s knowledge, but generating the vector space is slower, and so is computing dot products between concepts. The smaller is k , on the other hand, the more efficiently AffectiveSpace represents affective common-sense knowledge both in terms of vector space generation and of dot product computation. However, too few dimensions risk not to correctly represent AffectNet as concepts defined with too few features tend to be too close to each other in the vector space and, hence, not easily distinguishable and clusterable. In order to find a good k , AffectiveSpace was tested on a benchmark for affective common-sense knowledge (BACK) built by applying CF-IOF (concept frequency - inverse opinion frequency) [45] on the 5,000 posts of the LiveJournal corpus. CF-IOF is a technique that identifies common domain-dependent semantics in order to evaluate how important a concept is to a set of opinions concerning the same topic. Firstly, the frequency of a concept c for a given domain d is calculated by counting the occurrences of the concept c in the set of available d -tagged opinions and dividing the result by the sum of number of occurrences of all concepts in the set of opinions concerning d . This frequency is then multiplied by the logarithm of the inverse frequency of the concept in the whole collection of opinions, that is:

$$CF-IOF_{c,d} = \frac{n_{c,d}}{\sum_k n_{k,d}} \log \sum_k \frac{n_k}{n_c} \quad (3.4)$$

where $n_{c,d}$ is the number of occurrences of concept c in the set of opinions tagged as d , n_k is the total number of concept occurrences, and n_c is the number of occurrences of c in the whole set of opinions. A high weight in CF-IOF is reached by a high concept frequency in a given domain and a low frequency of the concept in the whole collection of opinions. Specifically, CF-IOF weighting was exploited to filter out common concepts in the LiveJournal corpus and to detect relevant mood-dependent semantics for the set of 24 emotions defined by Plutchik [229]. The result was a benchmark of 2000 affective concepts that were screened by 21 English-speaking students who were asked to map each concept to the 24 different emotional categories, which form the Hourglass of Emotions [52] (explained in next section). Results obtained were averaged (Table. 3.3).

Level	Label	Frequency
-G(1)	grief	14.3%
-G(2/3)	sadness	19.8%
-G(1/3)	pensiveness	11.4%
0	neutral	10.5%
+G(1/3)	serenity	20.6%
+G(2/3)	joy	18.3%
+G(1)	ecstasy	5.1%

Table 3.3 Distribution of concepts through the Pleasantness dimension. The affective information associated with most concepts concentrates around the centre of the Hourglass, rather than its extremes.




Fig. 3.3 Accuracy values achieved by testing AffectiveSpace on BACK, with dimensionality spanning from 1 to 250. The best trade-off between precision and efficiency is obtained around 100.

BACK's concepts were compared with the classification results obtained by applying the AffectiveSpace process using different values of k , from 1 to 250. As shown in Fig. 3.3, the best trade-off is achieved at 100, as selecting more than 100 singular values does not improve accuracy significantly. The distribution of the values of each AffectiveSpace dimension is bell-shaped, with different centres and different degree of dispersion around them. Affective common-sense concepts, in fact, tend to be close to the origin of the vector space (Fig. 3.4). In order to more uniformly distribute concept density in AffectiveSpace, an alternative strategy to represent the vector space was investigated. Such strategy consists in centring the values of the distribution of each dimension on the origin and in mapping dimensions according to a transformation $x \in \mathbb{R} \mapsto x^* \in [-1, 1]$. This transformation is often pivotal for better clustering AffectiveSpace as the vector space tends to have different grades of dispersion of data points across different dimensions, with some space regions more densely populated than others.

The switch to a different space configuration helps to distribute data more uniformly, possibly leading to an improved (or, at least, different) reasoning process. In particular, the transformation $x_{ij} \mapsto x_{ij} - \mu_i$ is first applied, being μ_i the average of all values of the i -th dimension. Then a normalisation is applied, combining the previous transformation with a new one $x_{ij} \mapsto \frac{x_{ij}}{a \cdot \sigma_i}$, where σ_i is the standard deviation calculated on the i -th dimension and a is a coefficient that can modify the same proportion of data that is represented within a specified interval. Finally, in order to ensure that all components of the vectors in the defined space are within $[-1, 1]$ (i.e., that the Chebyshev distance between the origin and each vector is smaller or equal to 1), a final transformation $x_{ij} \mapsto s(x_{ij})$ is needed, where $s(x)$ is a sigmoid function.




Fig. 3.4 A two-dimensional projection (first and second eigenmoods) of AffectiveSpace. From this visualisation, it is evident that concept density is usually higher near the centre of the space.

Different choices for the sigmoid function may be made, influencing how ‘fast’ the function approaches the unit value while the independent variable approaches infinity. Combining the proposed transformations, two possible mapping functions are expressed in the following formulae 3.5 and 3.6:

$$x_{ij}^* = \tanh\left(\frac{x_{ij} - \mu_i}{a \cdot \sigma_i}\right) \quad (3.5)$$

$$x_{ij}^* = \frac{x_{ij} - \mu_i}{a \cdot \sigma_i + |x_{ij} - \mu_i|} \quad (3.6)$$

This space transformation leads to two main advantages, which could be of notable importance depending on the problem being tackled. Firstly, this different space configuration ensures that each dimension is equally important by avoiding that the information provided by dimensions with higher (i.e., more distant from the origin) averages predominates. Secondly, normalising according to the standard deviations of each dimension allows a more uniform distribution of data around the origin, leading to a full use of information potential.

3.2 Affective Categorisation: Modelling Human Emotions

The study of emotions is one of the most confused (and still open) chapters in the history of psychology. This is mainly due to the ambiguity of natural language, which does not facilitate the description of mixed emotions in an unequivocal way. Love and other emotional words like anger and fear, in fact, are suitcase words (many different meanings packed in), not clearly defined and meaning different things to different people [199]. Hence, more than 90 definitions of emotions have been offered over the past century and there are almost as many theories of emotion, not to mention a complex array of overlapping words in our languages to describe them. Some categorisations include cognitive versus non-cognitive emotions, instinctual (from the amygdala) versus cognitive (from the prefrontal cortex) emotions, and also categorisations based on duration, as some emotions occur over a period of seconds (e.g., surprise), whereas others can last years (e.g., love).

The James-Lange theory posits that emotional experience is largely due to the experience of bodily changes [143]. Its main contribution is the emphasis it places on the embodiment of emotions, especially the argument that changes in the bodily concomitants of emotions can alter their experienced intensity. Most contemporary neuroscientists endorse a modified James-Lange view, in which bodily feedback modulates the experience of emotion [82]. In this view, emotions are related to certain activities in brain areas that direct our attention, motivate our behaviour, and determine the significance of what is going on around us. Pioneering works by Broca [30], Papez [223], and MacLean [187] suggested that emotion is related to a group of structures in the centre of the brain called limbic system (or paleomammalian brain), which includes the hypothalamus, cingulate cortex, hippocampi, and other structures. More recent research, however, has shown that some of these limbic structures are not as directly related to emotion as others are, while some non-limbic structures have been found to be of greater emotional relevance [167].

For tasks such as emotion recognition and polarity detection, it is key to have a model capable of finely describing the affective information associated with natural language concepts. To this end, a novel emotion categorisation model that goes beyond mere categorical and dimensional approaches (section 3.2.1) is proposed. Such a model represents affective states both through emotional labels and through four independent but concomitant dimensions that can potentially describe the full range of emotional experiences (section 3.2.2).

3.2.1 *Categorical Versus Dimensional Approaches*

Philosophical studies on emotions date back to ancient Greeks and Romans. Following the early Stoics, for example, Cicero enumerated and organised the emotions into four basic categories: *metus* (fear), *aegritudo* (pain), *libido* (lust), and *laetitia* (pleasure). Studies on evolutionary theory of emotions, in turn, were initiated in the late 19th century by Darwin [71]. His thesis was that emotions evolved via natural selection and, therefore, have cross-culturally universal counterparts. In the early 1970s, Ekman found evidence that humans share six basic emotions: happiness, sadness, fear, anger, disgust, and surprise [99]. Few tentative efforts to detect non-basic affective states, such as fatigue, anxiety, satisfaction, confusion, or frustration, have been also made [250, 225, 234, 96, 149, 64] (Table 3.4).

In 1980, Averill put forward the idea that emotions cannot be explained strictly on the basis of physiological or cognitive terms. Instead, he claimed that emotions are primarily social constructs; hence, a social level of analysis is necessary to truly understand the nature of emotion [14]. The relationship between emotion and language (and the fact that the language of emotion is considered a vital part of the experience of emotion) has been used by social constructivists and anthropologists to question the universality of Ekman's studies, arguably because the language labels he used to code emotions are somewhat US-centric. In addition, other cultures might have labels that cannot be literally translated to English (e.g., some languages do not have a word for fear [247]).

For their deep connection with language and for the limitedness of the emotional labels used, all such categorical approaches usually fail to describe the complex range of emotions that can occur in daily communication. The dimensional approach [214], in turn, represents emotions as coordinates in a multi-dimensional space. For both theoretical and practical reasons, more and more researchers like to define emotions according to two or more dimensions. An early example is Russell's circumplex model [246], which uses the dimensions of arousal and valence to plot 150 affective labels (Fig. 3.5).

Author	#Emotions	Basic Emotions
Ekman	6	anger, disgust, fear, joy, sadness, surprise
Parrot	6	anger, fear, joy, love, sadness, surprise
Frijda	6	desire, happiness, interest, surprise, wonder, sorrow
Plutchik	8	acceptance, anger, anticipation, disgust, joy, fear, sadness, surprise
Tomkins	9	desire, happiness, interest, surprise, wonder, sorrow
Matsumoto	22	joy, anticipation, anger, disgust, sadness, surprise, fear, acceptance, shy, pride, appreciate, calmness, admire, contempt, love, happiness, exciting, regret, ease, discomfort, respect, like

Table 3.4 Some existing definition of basic emotions. The most widely adopted model for affect recognition is Ekman's, although is one of the poorest in terms of number of emotions.




Fig. 3.5 Russell's circumplex model is one of the earliest examples of dimensional emotion representations. In the snippet, direct circular scaling coordinates are provided for 28 affect words.

Similarly, Whissell considers emotions as a continuous 2D space whose dimensions are evaluation and activation [300]. The evaluation dimension measures how a human feels, from positive to negative. The activation dimension measures whether humans are more or less likely to take some action under the emotional state, from active to passive (Fig. 3.6). In her study, Whissell assigns a pair of values $\langle \text{activation}, \text{evaluation} \rangle$ to each of the approximately 9,000 words with affective connotations that make up her Dictionary of Affect in Language.

Another bi-dimensional model is Plutchik's wheel of emotions, which offers an integrative theory based on evolutionary principles [229]. Following Darwin's thought, the functionalist approach to emotions holds that emotions have evolved for a particular function, such as to keep the subject safe [114, 113]. Emotions are adaptive as they have a complexity born of a long evolutionary history and, although we conceive emotions as feeling states, Plutchik says the feeling state is part of a process involving both cognition and behaviour and containing several feedback loops. In 1980, he created a wheel of emotions, which consisted of 8 basic emotions and 8 advanced emotions each composed of 2 basic ones. In such model, the vertical dimension represents intensity and the radial dimension represents degrees of similarity among the emotions.




Fig. 3.6 Whissell's model is a bi-dimensional representation of emotions, in which affect words are displayed. The diagram shows the position of some of these words in the <activation, evaluation> space.

Besides bi-dimensional approaches, a commonly used set for emotion dimension is the <arousal, valence, dominance> set, which is known in the literature also by different names, including <evaluation, activation, power> and <pleasure, arousal, dominance> [193]. Recent evidence suggests there should be a fourth dimension: Fontaine et al. reported consistent results from various cultures where a set of four dimensions is found in user studies, namely <valence, potency, arousal, unpredictability> [111]. Dimensional representations of affect are attractive mainly because they provide a way of describing emotional states that is more tractable than using words. This is of particular importance when dealing with naturalistic data, where a wide range of emotional states occurs. Similarly, they are much more able to deal with non-discrete emotions and variations in emotional states over time [76], since in such cases changing from one universal emotion label to another would not make much sense in real life scenarios.

Dimensional approaches, however, have a few limitations. Although the dimensional space allows to compare affect words according to their reciprocal distance, it usually does not allow making operations between these, e.g., for studying compound emotions. Most dimensional representations, moreover, do not model the fact that two or more emotions may be experienced at the same time. Eventually, all such approaches work at word level, which makes them unable to grasp the affective valence of multiple-word concepts.

3.2.2 *The Hourglass of Emotions*

The Hourglass of Emotions [52] is an affective categorisation model inspired by Plutchik's studies on human emotions [229]. It reinterprets Plutchik's model by organising primary emotions around four independent but concomitant dimensions, whose different levels of activation make up the total emotional state of the mind. Such a reinterpretation is inspired by Minsky's theory of the mind, according to which brain activity consists of different independent resources and that emotional states result from turning some set of these resources on and turning another set of them off [199]. This way, the model can potentially synthesise the full range of emotional experiences in terms of Pleasantness, Attention, Sensitivity, and Aptitude, as the different combined values of the four affective dimensions can also model affective states we do not have a specific name for, due to the ambiguity of natural language and the elusive nature of emotions.

The main motivation for the design of the model is the concept-level inference of the cognitive and affective information associated with text. Such faceted information is needed, within sentic computing, for a feature-based sentiment analysis, where the affective common-sense knowledge associated with natural language opinions has to be objectively assessed. Therefore, the Hourglass model systematically excludes what are variously known as self-conscious or moral emotions, e.g., pride, guilt, shame, embarrassment, moral outrage, or humiliation [166, 173, 251, 282]. Such emotions, in fact, present a blind spot for models rooted in basic emotions, because they are by definition contingent on subjective moral standards. The distinction between guilt and shame, for example, is based in the attribution of negativity to the self or to the act. So, guilt arises when believing to have done a bad thing, and shame arises when thinking to be a bad person. This matters because, in turn, these emotions have been shown to have different consequences in terms of action tendencies. Likewise, an emotion such as *schadenfreude* is essentially a form of pleasure, but it is crucially different from pride or happiness because of the object of the emotion (the misfortune of another that is not caused by the self), and the resulting action tendency (do not express).

However, since the Hourglass model currently focuses on the objective inference of affective information associated with natural language opinions, appraisal-based emotions are not taken into account within the present version of the model. The Hourglass model, in fact, is a biologically-inspired and psychologically-motivated

model based on the idea that emotional states result from the selective activation/disactivation of different resources in the brain. Each such selection changes how we think by changing our brain's activities: the state of anger, for example, appears to select a set of resources that help us react with more speed and strength while also suppressing some other resources that usually make us act prudently. Evidence of this theory is also given by several fMRI experiments showing that there is a distinct pattern of brain activity that occurs when people are experiencing different emotions. Zeki and Romaya, for example, investigated the neural correlates of hate with an fMRI procedure [312]. In their experiment, people had their brains scanned while viewing pictures of people they hated. The results showed increased activity in the medial frontal gyrus, right putamen, bilaterally in the premotor cortex, in the frontal pole, and bilaterally in the medial insula of the human brain. Also the activity of emotionally enhanced memory retention can be linked to human evolution [34]. During early development, in fact, responsive behaviour to environmental events is likely to have progressed as a process of trial-and-error.

Survival depended on behavioural patterns that were repeated or reinforced through life and death situations. Through evolution, this process of learning became genetically embedded in humans and all animal species in what is known as 'fight or flight' instinct [29]. The primary quantity we can measure about an emotion we feel is its strength. But, when we feel a strong emotion, it is because we feel a very specific emotion. And, conversely, we cannot feel a specific emotion like fear or amazement without that emotion being reasonably strong. For such reasons, the transition between different emotional states is modelled, within the same affective dimension, using the function $G(x) = -\frac{1}{\sigma\sqrt{2\pi}}e^{-x^2/2\sigma^2}$, for its symmetric inverted bell curve shape that quickly rises up towards the unit value (Fig. 3.7).




Fig. 3.7 The Pleasantness emotional flow. The passage from a sentic level to another is regulated by a Gaussian function that models how stronger emotions induce higher emotional sensitivity.

In particular, the function models how the level of activation of each affective dimension varies from the state of ‘emotional void’ (null value) to the state of ‘heightened emotionality’ (unit value). Justification for assuming that the Gaussian function (rather than a step or simple linear function) is appropriate for modelling the variation of emotion intensity is based on research into the neural and behavioural correlates of emotion, which are assumed to indicate emotional intensity in some sense. In fact, nobody genuinely knows what function subjective emotion intensity follows, because it has never been truly or directly measured [18]. For example, the so-called Duchenne smile (a genuine smile indicating pleasure) is characterised by smooth onset, increasing to an apex, and a smooth, relatively lengthy offset [157]. More generally, Klaus Scherer has argued that emotion is a process characterised by non-linear relations among its component elements - especially physiological measures, which typically look Gaussian [174]. Emotions, in fact, are not linear [229]: the stronger the emotion, the easier it is to be aware of it. Mapping this space of possible emotions leads to a hourglass shape (Fig. 3.8). It is worth to note that, in the model, the state of ‘emotional void’ is a-dimensional, which contributes to determine the hourglass shape. Total absence of emotion, in fact, can be associated with the total absence of reasoning (or, at least, consciousness) [80], which is not an envisaged mental state as, in human mind, there is never nothing going on.

The Hourglass of Emotions, in particular, can be exploited in the context of HCI to measure how much respectively: the user is amused by interaction modalities (Pleasantness), the user is interested in interaction contents (Attention), the user is comfortable with interaction dynamics (Sensitivity), the user is confident in interaction benefits (Aptitude). Each affective dimension, in particular, is characterised by six levels of activation (measuring the strength of an emotion), termed ‘sentic levels’, which represent the intensity thresholds of the expressed or perceived emotion. These levels are also labelled as a set of 24 basic emotions [229], six for each of the affective dimensions, in a way that allows the model to specify the affective information associated with text both in a dimensional and in a discrete form (Table 3.5). The dimensional form, in particular, is termed ‘sentic vector’ and it is a four-dimensional *float* vector that can potentially synthesise the full range of emotional experiences in terms of Pleasantness, Attention, Sensitivity, and Aptitude.

Interval	Pleasantness	Attention	Sensitivity	Aptitude
$[G(1), G(2/3))$	ecstasy	vigilance	rage	admiration
$[G(2/3), G(1/3))$	joy	anticipation	anger	trust
$[G(1/3), G(0))$	serenity	interest	annoyance	acceptance
$(G(0), -G(1/3)]$	pensiveness	distraction	apprehension	boredom
$(-G(1/3), -G(2/3)]$	sadness	surprise	fear	disgust
$(-G(2/3), -G(1)]$	grief	amazement	terror	loathing

Table 3.5 The sentic levels of the Hourglass model. Labels are organised into four affective dimensions with six different levels each, whose combined activity constitutes the ‘total state’ of the mind.




Fig. 3.8 The 3D model and the net of the Hourglass of Emotions. Since affective states go from strongly positive to null to strongly negative, the model assumes a hourglass shape.

In the model, the vertical dimension represents the intensity of the different affective dimensions, i.e., their level of activation, while the radial dimension represents K-lines [197] that can activate configurations of the mind, which can either last just a few seconds or years. The model follows the pattern used in colour theory and research in order to obtain judgements about combinations, i.e., the emotions that result when two or more fundamental emotions are combined, in the same way that red and blue make purple.

Hence, some particular sets of sentic vectors have special names, as they specify well-known compound emotions (Fig. 3.9). For example, the set of sentic vectors with a level of Pleasantness $\in [G(2/3), G(1/3))$, i.e., joy, a level of Aptitude $\in [G(2/3), G(1/3))$, i.e., trust, and a minor magnitude of Attention and Sensitivity, are termed 'love sentic vectors' since they specify the compound emotion of love (Table 3.6).

	Attention>0	Attention<0	Aptitude>0	Aptitude<0
Pleasantness>0	optimism	frivolity	love	gloat
Pleasantness<0	frustration	disapproval	envy	remorse
Sensitivity>0	aggressiveness	rejection	rivalry	contempt
Sensitivity<0	anxiety	awe	submission	coercion

Table 3.6 The second-level emotions generated by pairwise combination of the sentic levels of the Hourglass model. The co-activation of different levels gives birth to different compound emotions.

More complex emotions can be synthesised by using three, or even four, sentic levels, e.g., joy + trust + anger = jealousy. Therefore, analogous to the way primary colours combine to generate different colour gradations (and even colours we do not have a name for), the primary emotions of the Hourglass model can blend to form the full spectrum of human emotional experience. Beyond emotion detection, the Hourglass model is also used for polarity detection tasks. Since polarity is strongly connected to attitudes and feelings, in fact, it is defined in terms of the four affective dimensions, according to the formula:

$$p = \sum_{i=1}^N \frac{Pleasantness(c_i) + |Attention(c_i)| - |Sensitivity(c_i)| + Aptitude(c_i)}{3N} \quad (3.7)$$

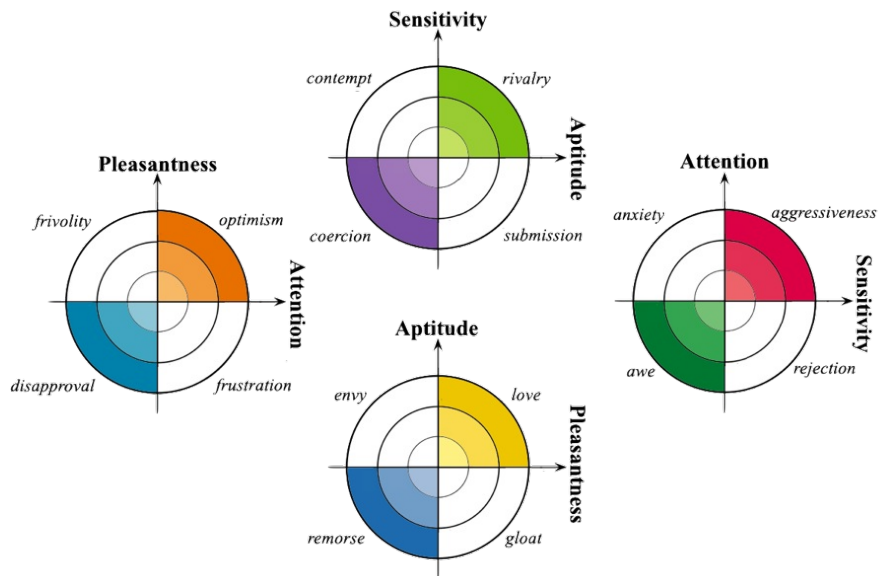


Fig. 3.9 Hourglass compound emotions of second level. By combining basic emotions pairwise, it is possible to obtain complex emotions resulting from the activation of two affective dimensions.

where c_i is an input concept, N the total number of concepts, and 3 the normalisation factor (as the Hourglass dimensions are defined as float $\in [-1,+1]$). In the formula, Attention is taken as absolute value since both its positive and negative intensity values correspond to positive polarity values (e.g., ‘surprise’ is negative in the sense of lack of Attention, but positive from a polarity point of view). Similarly, Sensitivity is taken as negative absolute value since both its positive and negative intensity values correspond to negative polarity values (e.g., ‘anger’ is positive in the sense of level of activation of Sensitivity, but negative in terms of polarity). The formula can be seen as one of the first attempts to show a clear connection between emotion recognition (sentiment analysis) and polarity detection (opinion mining).

3.3 Sentic Medoids: Clustering Affective Common-Sense Concepts

Sentic Medoids [54] is a technique that adopts a k -medoids approach [151] to partition the affective common-sense concepts of AffectiveSpace into k clusters around as many centroids, trying to minimise a given cost function. Differently from the k -means algorithm [128], which does not pose constraints on centroids, k -medoids do assume that centroids must coincide with k observed points. This section introduces the standard approach to k -medoids clustering (section 3.3.1) and describes the algorithm developed for clustering AffectiveSpace (section 3.3.2).

3.3.1 Partitioning Around Medoids

Clustering is the process of grouping a set of objects into classes or clusters so that objects within a cluster have similarity in comparison to one another, but are dissimilar to objects in other clusters [126]. Well known techniques for performing non-hierarchical clustering are k -means [188] and partitioning around medoids (PAM) [151]. The k -means approach finds the k centroids, where the coordinate of each centroid is the mean of the coordinates of the objects in the cluster and assigns every object to the nearest centroid. Unfortunately, k -means clustering is sensitive to the outliers and a set of objects closest to a centroid may be empty, in which case centroids cannot be updated. For this reason, k -medoids are often used when representative objects, rather than centroids, need to be considered.

In many clustering problems, in fact, one is interested in the characterisation of the clusters by means of typical objects, which represent the various structural features of objects under investigation. Because it uses the most centrally located object in a cluster, k -medoids clustering is less sensitive to outliers compared with k -means. Among many algorithms for k -medoids clustering, PAM is one of the most widely used. The algorithm, proposed by Kaufman and Rousseeuw [151], first computes k representative objects, termed medoids.

A medoid can be defined as that object of a cluster, whose average dissimilarity to all the objects in the cluster is minimal. PAM determines a medoid for each cluster selecting the most centrally located centroid within the cluster. After selection of medoids, clusters are rearranged so that each point is grouped with the closest medoid. Compared to k -means, PAM operates on the dissimilarity matrix of the given dataset. It is more robust because it minimises a sum of dissimilarities instead of a sum of squared Euclidean distances. A particularly nice property is that PAM allows clustering with respect to any specified distance metric. In addition, the medoids are robust representations of the cluster centres, which is particularly important in the common context that many elements do not belong well to any cluster.

3.3.2 Centroid Selection

Since k -medoids clustering is a NP-hard problem [115], different approaches based on alternative optimisation algorithms have been developed, all of which carry the risk of being trapped around local minima. Among many algorithms for k -medoids clustering, PAM is known to be most powerful. However, PAM also has a drawback that it works inefficiently for large data sets due to its complexity. To this end, a modified version of the algorithm recently proposed by Park and Jun [224] was used, which runs in a similar way to the k -means clustering algorithm. This has been shown to have similar performance when compared to the PAM algorithm while taking a significantly reduced computational time. In particular, AffectiveSpace contains N concepts ($N = 14,301$) encoded as points $x \in \mathbb{R}^p$ ($p = 100$). They need to be grouped into k clusters and, in this specific case, k can be fixed to 24, as one cluster for each sentic level s of the Hourglass model is being searched (Fig. 3.10). Generally, the initialisation of clusters for clustering algorithms is a problematic task as the process often risks getting stuck in local optimum points, depending on the initial choice of centroids [97]. In this work, we are able to conveniently use, as initial centroids, the concepts that are currently used as centroids for clusters - since they specify the emotional categories we want to organise AffectiveSpace into.

For this reason, what is usually seen as a limitation of the algorithm can be seen as an advantage for this particular approach, since what is being searched are not the 24 centroids leading to the best 24 clusters, but indeed for the 24 centroids identifying the required 24 sentic levels (that is, the centroids should not be ‘too far’ from the ones currently used). In particular, as the Hourglass affective dimensions are independent but concomitant, AffectiveSpace needs to be clustered four times, once for each dimension. According to the Hourglass categorisation model, in fact, each concept can convey, at the same time, more than one emotion (which is why we get compound emotions) and this information can be expressed via a sentic vector specifying the concept’s affective valence in terms of Pleasantness, Attention, Sensitivity, and Aptitude.

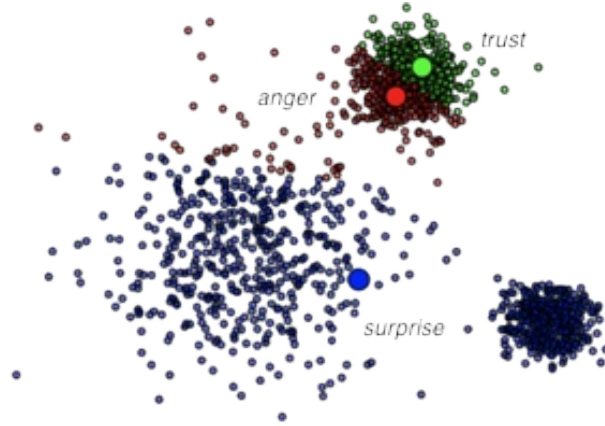


Fig. 3.10 Sentic Medoids organise AffectiveSpace by iteratively selecting centroids, in order to minimise the sum of dissimilarities between these and the other concepts within the same cluster.

Therefore, given that the distance between two points in AffectiveSpace is defined as $D(a, b) = \sqrt{\sum_{i=1}^p (a_i - b_i)^2}$ (note that the choice of Euclidean distance is arbitrary), the employed algorithm, applied for each of the four affective dimensions, can be summarised as follows:

1. Each centroid $C_n \in \mathbb{R}^{100}$ ($n = 1, 2, \dots, k$) is set as one of the six concepts corresponding to each s in the current affective dimension
2. Assign each record x to a cluster \mathcal{E} so that $x_i \in \mathcal{E}_n$ if $D(x_i, C_n) \leq D(x_i, C_m)$ $m = 1, 2, \dots, k$
3. Find a new centroid C for each cluster \mathcal{E} so that $C_j = x_i$ if $\sum_{x_m \in \mathcal{E}_j} D(x_i, x_m) \leq \sum_{x_m \in \mathcal{E}_j} D(x_h, x_m) \quad \forall x_h \in \mathcal{E}_j$
4. Repeat step 2 and 3 until no changes on centroids are observed

Note that condition posed on steps 2 and 3 may occasionally lead to more than one solution. Should this happen, the proposed model will randomly choose one of them. This clusterisation of AffectiveSpace allows to calculate, for each common-sense concept x , a four-dimensional sentic vector that defines its affective valence in terms of a degree of fitness $\mathbf{f}(x)$ where $f_a = D(x, C_j) \quad C_j | D(x, C_j) \leq D(x, C_k)$ $a = 1, 2, 3, 4 \quad k = 6a-5, 6a-4, \dots, 6a$.

3.4 Sentic Activation: A Two-Level Affective Reasoning Framework

Current thinking in cognitive psychology suggests that humans process information at a minimum of two distinct levels. There is extensive evidence for the existence of two (or more) processing systems within the human brain, one that involves fast, parallel, unconscious processing, and one that involves slow, serial, more conscious processing [156, 67, 259, 102]. Dual-process models of automatic and controlled social cognition have been proposed in nearly every domain of social psychology.

Evidence from neurosciences supports this separation, with identifiably different brain regions involved in each of the two systems [177]. Such systems, termed U-level (unconscious) and C-level (conscious), can operate simultaneously or sequentially, and are most effective in different contexts. The former, in particular, works intuitively, effortlessly, globally, and emotionally (section 3.4.1). The latter, in turn, works logically, systematically, effortfully, and rationally (section 3.4.2).

3.4.1 Unconscious Reasoning

In recent years, neuroscience has contributed a lot to the study of emotions through the development of novel methods for studying emotional processes and their neural correlates. In particular, new methods used in affective neuroscience, e.g., fMRI, lesion studies, genetics, electro-physiology, paved the way towards the understanding of the neural circuitry that underlies emotional experience and of the manner in which emotional states influence health and life outcomes. A key contribution in the last two decades has been to provide evidence against the notion that emotions are subcortical and limbic, whereas cognition is cortical.

This notion was reinforcing the flawed Cartesian dichotomy between thoughts and feelings [85]. There is now ample evidence that the neural substrates of cognition and emotion overlap substantially [83]. Cognitive processes, such as memory encoding and retrieval, causal reasoning, deliberation, goal appraisal, and planning, operate continually throughout the experience of emotion. This evidence points to the importance of considering the affective components of any human-computer interaction [36]. Affective neuroscience, in particular, has provided evidence that elements of emotional learning can occur without awareness [211] and elements of emotional behaviour do not require explicit processing [35]. Affective information processing, in fact, mainly takes place at unconscious level (U-level) [102].

Reasoning, at this level, relies on experience and intuition, which allow considering issues intuitively and effortlessly. Hence, rather than reflecting upon various considerations in sequence, the U-level forms a global impression of the different issues. In addition, rather than applying logical rules or symbolic codes (e.g., words or numbers), the U-level considers vivid representations of objects or events.

Such representations are laden with the emotions, details, features, and sensations that correspond to the objects or events. Such human capability of summarising the huge amount of inputs and outputs of previous situations, in order to find useful patterns that might work at the present time, is hereby implemented by means of AffectiveSpace. By reducing the dimensionality of the matrix representation of AffectNet, in fact, AffectiveSpace compresses the feature space of affective common-sense knowledge into one that allows to better gain global insight and human-scale understanding. In cognitive science, the term ‘compression’ refers to transforming diffuse and distended conceptual structures that are less congenial to human understanding so that they become better suited to our human-scale ways of thinking.

Compression is hereby achieved by balancing the number of singular values discarded when synthesising AffectiveSpace, in a way that the affective common-sense knowledge representation is neither too concrete nor too abstract with respect to the detail granularity needed for performing a particular task. The reasoning-by-analogy capabilities of AffectiveSpace, hence, are exploited at U-level to achieve digital intuition about the input data. In particular, the vector space representation of affective common-sense knowledge is clustered according the Hourglass model using the Sentic Medoids technique, in a way that concepts that are semantically and affectively related to the input data can be intuitively retrieved by analogy and unconsciously crop out to the C-level.

3.4.2 Conscious Reasoning

U-level and C-level are two conceptual systems that operate by different rules of inference. While the former operates emotionally and intuitively, the latter relies on logic and rationality. In particular, the C-level analyses issues with effort, logic, and deliberation rather than relying on intuition. Hence, while at U-level the vector space representation of AffecNet is exploited to intuitively guess semantic and affective relations between concepts, at C-level associations between concepts are made according to the actual connections between different nodes in the graph representation of affective common-sense knowledge. Memory is not a ‘thing’ that is stored somewhere in a mental warehouse and can be pulled out and brought to the fore. Rather, it is a potential for reactivation of a set of concepts that together constitute a particular meaning. Associative memory involves the unconscious activation of networks of association—thoughts, feelings, wishes, fears, and perceptions that are connected, so that activation of one node in the network leads to activation of the others [299].

Sentic Activation [55] aims to implement such a process through the ensemble application of multi-dimensionality reduction and graph mining techniques. Specifically, the semantically and affectively related concepts retrieved by means of AffectiveSpace at U-level are fed into AffectNet in order to crawl it according to how such seed concepts are interconnected to each other and to other concepts in the semantic network. To this end, spectral association [132] is employed.

Spectral association is a technique that assigns values, or activations, to seed concepts and spreads their values across the AffectNet graph. This operation, which is an approximation of many steps of spreading activation, transfers the most activation to concepts that are connected to the seed concepts by short paths or many different paths in affective common-sense knowledge. These related concepts are likely to have similar affective values. This can be seen as an alternate way of assigning affective values to all concepts, which simplifies the process by not relying on an outside resource such as WNA. In particular, a matrix A that relates concepts to other concepts, instead of their features, is built and the scores are added up over all relations that relate one concept to another, disregarding direction. Applying A to a vector containing a single concept spreads that concept's value to its connected concepts. Applying A^2 spreads that value to concepts connected by two links (including back to the concept itself). But the desired operation is to spread the activation through any number of links, with diminishing returns, so the operator wanted is:

$$1 + A + \frac{A^2}{2!} + \frac{A^3}{3!} + \dots = e^A \quad (3.8)$$

This odd operator, e^A , can be calculated because A can be factored. A is already symmetric, so instead of applying Lanczos' method [162] to AA^T and getting the SVD, it can be applied directly to A to obtain the spectral decomposition $A = V\Lambda V^T$. As before, this expression can be raised to any power and everything but the power of Λ cancelled. Therefore, $e^A = Ve^{\Lambda}V^T$. This simple twist on the SVD allows to calculate spreading activation over the whole matrix instantly. As with the SVD, these matrices can be truncated to k axes and, therefore, space can be saved while generalising from similar concepts. The matrix can also be rescaled so that activation values have a maximum of 1 and do not tend to collect in highly-connected concepts such as 'person', by normalising the truncated rows of $Ve^{\Lambda/2}$ to unit vectors, and multiplying that matrix by its transpose to get a rescaled version of $Ve^{\Lambda}V^T$.

Spectral association can spread not only positive, but also negative activation values. Hence, unconscious reasoning at U-level is exploited not only to retrieve concepts that are most semantically and affectively related, but also concepts that are most likely to be unrelated with the input data (lowest dot product). While the former are exploited to spread semantics and sentics across the AffectNet graph, the latter are used to contain such an activation in a way that potentially unrelated concepts (and their twins) do not get triggered. This brain-inspired ensemble application of dimensionality reduction and graph mining techniques (hereby referred as unconscious and conscious reasoning, respectively) allows Sentic Activation to more efficiently infer semantics and sentics from natural language text.

In fact, Sentic Activation was tested on the benchmark for affective common-sense knowledge (BACK) built by means of CF-IOF. In particular, BACK's concepts were compared with the classification results obtained by applying the AffectiveSpace process (U-level), spectral association (C-level), and Sentic Activation (ensemble of U-level and C-level). Results showed that Sentic Activation achieves +13.9% and +8.2% accuracy than the AffectiveSpace process and spectral association, respectively.

3.5 Sentic Panalogy: Switching Between Different Ways to Think

Emotions are different Ways to Think [199] that our mind triggers to deal with different situations we face in our lives. Our ability to reason and make decisions, in fact, is strictly dependent on both our common-sense knowledge about the world and our inner emotional states. This capability, termed affective common sense reasoning, is a fundamental component in human experience, cognition, perception, learning, and communication. For this reason, we cannot prescind from emotions in the development of intelligent user interfaces: if we want computers to be really intelligent, not just have the veneer of intelligence, we need to give them the ability to recognise, understand, and express emotions. There is extensive evidence for the existence of multiple processing systems within the human brain [156, 67, 259, 102, 177]. In order not to get stuck and to be able to tackle different problems from different perspectives, in fact, an intelligent machine should not have a unique way to deal with a task, but rather be endowed with different reasoning strategies and with the capability to switch among these.

This section discusses an approach to emulate the human capability to switch between different perspectives and find novel ways to look at things. Such approach is inspired by Minsky's notion of 'panalogy' (parallel analogy), which states that several analogous representations of the same problem should be maintained in parallel while trying to solve it so that, as soon as problem solving begins to fail when using one representation, the system can switch to one of the others [199]. In particular, multi-dimensionality reduction techniques are employed on AffectNet to dynamically configure it and, hence, model the switch between different reasoning strategies (section 3.5.1), while graph mining and clustering methods are applied to model the switch between the foci around which those strategies are developed [56], in order to accordingly exploit the different facets of the affective common-sense knowledge base (section 3.5.2).

3.5.1 Changing Reasoning Strategies

To some extent, our reasoning capability can be re-conducted to the identification of useful patterns in our acquired knowledge about the world. Our experience and common-sense knowledge is likely to be organised in our mind as interconnected concepts and events, and most of these links are weighted by affective information, as we tend to forget or hardly recall memories that are not associated with any kind of positive or negative emotion. Therefore, the human capacity to envision possible outcomes of a decision might lie both in the capability of crawling the semantic network of concepts we have acquired through experience (C-level), and in the capability of summarising the huge amount of inputs and outputs of previous situations to find useful patterns that might work at the present time (U-level). As shown in section 3.1.2, a way to perform reasoning on the affective common-sense knowledge base is to use multi-dimensionality reduction techniques.

In this context, the capability to look at things from a different perspective can be emulated by applying different space transformations to AffectiveSpace. As shown in section 3.1.2, the switch to a different space configuration helps to differently distribute data, possibly leading to an improved reasoning process. Switching between different space configurations, in fact, changes how each dimension is influential in the vector space representation of AffectNet and, hence, changes how we are looking at the affective common-sense knowledge because similarity in AffectiveSpace does not depend on concepts' absolute position, but rather on the angle they make with the origin of the vector space.

Moreover, the number k of singular values selected for building AffectiveSpace is a measure of the trade-off between precision and efficiency in the representation of the affective common-sense knowledge base. Switching between different vector space dimensionalities can be seen as looking at the data from many different points of view. Balancing the number of singular values discarded when synthesising AffectiveSpace, hence, corresponds to calibrate the affective common-sense knowledge representation in a way that it is neither too concrete nor too abstract with respect to the detail granularity needed for performing a particular task. Different k values, for example, work differently according to the affective dimension considered, e.g., for Pleasantness the best k appears to be closer to 100, while for Sensitivity a space of about 70 dimensions appears to be enough for precisely and efficiently represent affective common-sense knowledge.

3.5.2 *Changing Reasoning Foci*

The capability of switching among different Ways to Think can be thought as changing the foci around which we develop our different reasoning strategies. Such approach can be implemented in AffectiveSpace by changing the centroids around which the vector space is clustered. In section 3.3, it has been shown how this can be implemented through a k -medoids approach. According to the initial centroid chosen, the final clusterisation of AffectiveSpace can be very different. Hence, the way such initial medoids are selected can be re-conducted to the human capability to switch between different perspectives to grasp the different facets of a problem. As shown in section 3.4.2, another way to perform reasoning on AffectNet is to see it as a graph and exploit its connectivity to find semantically and affectively related concepts by means of spectral association [132].

As explained before, such a technique involves assigning activation to seed concepts and applying an operation that spreads their values across the graph. Seed concepts can also be associated with negative activation values in order to reduce the spreading operation in the parts of the graph we are specifically not interested in. If concepts semantically related to 'bank' as a financial institution have to be searched without getting concepts related to 'river bank', for example, concepts like 'money', 'savings', or 'investment' can be set as positive seeds and concepts like 'river', 'water', or 'shore' can be set as negative seeds.

The outcomes of spectral association can be very different according to which seeds are selected as starting points for the spreading activation steps. Since spectral association involves TSVD, results also depend on the number k of singular values selected. While choosing different k values can be seen as developing different reasoning strategies, choosing different seeds can be associated to changing the foci around which those strategies are developed. Through spectral association, positive and negative values of these concepts are spread across the graph representation of AffectNet, resulting in a set of contextually semantic related instances. Letting a machine switch between such seeds according to its own intuition (e.g., concepts obtained through AffectiveSpace) can be re-conducted to the human capability to change the foci around which different reasoning strategies are developed and, hence, to iterate on the ways to look at a problem until one that works is found.

3.6 Conclusions

This chapter has shown how ConceptNet was blended with a linguistic resource for the lexical representation of affective knowledge, in order to obtain a new knowledge base in which concepts are interrelated by both common-sense and affective features. The chapter has also explained how the application of dimensionality reduction techniques on such a knowledge base yields a vector space of affective common-sense knowledge, which can be accordingly configured depending on the desired trade-off between precision and efficiency (section 3.1).

So far, TSVD appeared to be a good method for generalising the information contained in AffectNet, but it is very expensive both in terms of computing time and storage, as it requires costly arithmetic operations such as division and square root in the computation of rotation parameters. This is a big issue because AffectNet is keeping on growing, in parallel with the continuously extended versions of ConceptNet. To this end, alternative multi-dimensionality reduction techniques, e.g., independent component analysis (ICA) and random projections, are currently being explored.

In order to accordingly categorise affect in the developed knowledge base, moreover, the chapter presented a novel emotion categorisation model that goes beyond mere categorical and dimensional approaches by representing affective states both through emotional labels and through four independent but concomitant dimensions that can potentially describe the full range of emotional experiences (section 3.2). The chapter, then, described a new PAM-based clustering method for organising and categorising AffectiveSpace according to such an emotion categorisation model (section 3.3). The rest of the chapter, finally, illustrated how the ensemble application of dimensionality reduction and graph mining techniques can be exploited to emulate conscious and unconscious reasoning processes (section 3.4) and the capability to switch between different perspectives in order to find novel ways to look at things (section 3.5).

Chapter 4

Tools

*You can know the name of a bird in all the languages of the world,
but when you're finished, you'll know absolutely nothing whatever about the bird.
So let's look at the bird and see what it's doing – that's what counts.
I learned very early the difference between
knowing the name of something and knowing something.*

Richard Feynman

In standard human-to-human communication, people usually refer to existing facts and circumstances and build new useful, funny, or interesting information on the top of those. This common knowledge comprehends information usually found in news, articles, debates, lectures, etc. (factual knowledge), but also principles and definitions that can be found in collective intelligence projects such as Wikipedia (vocabulary knowledge). Attempts to build a common knowledge base are countless and comprehend both resources crafted by human experts or community efforts, such as WordNet and Freebase [27], a large collaborative knowledge base consisting of metadata composed mainly by its community members, and automatically-built knowledge bases, such as WikiTaxonomy [230], a taxonomy extracted from Wikipedia's category links, YAGO [276], a semantic knowledge base derived from Wikipedia, WordNet, and GeoNames¹, and Never-Ending Language Learning (NELL), CMU's semantic machine learning system [63].

Common knowledge, however, is just a surface layer of human communication and alone is not enough for understanding natural language. In order to join in a conversation and actively participate in it, we need to have some information about the concepts and the context this conversation is based on but, before that, we need to be able to speak. Trying to understand natural language by simply relying on common knowledge would be pretty much like trying to join a conversation about food, as a one-year old child: we would almost perfectly know the debated topics,

¹ <http://geonames.org>

but we would not really know how to relate these concepts to each other in order to fully understand or form a sentence. These semantic relationships are what we call common-sense knowledge and consist in knowing that people are younger than their parents, that a butcher is unlikely to be a vegetarian, that people do not like being repeatedly interrupted, that if you hold a knife by its blade then it may cut you, that if you drop paper into a flame then it will burn, that people pay taxi drivers to drive them to places, that people generally sleep at night, and so forth.

Computers do not know such things, as they never had the opportunity to live a life and to experience or be taught the meaning of words, objects, and actions, and how these relate to each other. If we want machines to really understand natural language, hence, we need to provide them with such knowledge. Building a common-sense knowledge base, however, is not easy as common-sense is typically omitted from social communications. For Grice's theory of pragmatics [120], in fact, when people communicate with each other, they tend not to provide information that is obvious or extraneous. This is true both for face-to-face communication and for asynchronous communication (e.g., on the Web), which makes the collection of common-sense an extremely difficult task.

This chapter shows how sentic computing techniques can be employed for the development of a publicly available resource of affective common-sense knowledge (section 4.1), a brain-inspired model for affective reasoning (section 4.2), and an intelligent user interface (UI) for the collection of new affective common-sense knowledge (section 4.3). Additionally, the chapter describes the design of a knowledge base of common and common-sense knowledge (section 4.4), which is later embedded in an opinion mining engine (section 4.5) for improving its topic-spotting capabilities. The chapter, finally, provides a summary and some concluding remarks (section 4.6).

4.1 SenticNet: A Semantic Resource for Opinion Mining

SenticNet² [60] is a publicly available resource for opinion mining that exploits both AI and Semantic Web techniques to infer the polarity associated with common-sense concepts and represent it in a semantic-aware format. In particular, SenticNet uses dimensionality reduction to calculate the affective valence of a set of Open Mind concepts and represent it in a machine-accessible and machine-processable format. The result is a publicly available knowledge base (also accessible through an API³) for mining opinions from natural language text at a semantic, rather than just syntactic, level. This section shows motivations and techniques for building such a resource and encode it in a Semantic Web aware format (section 4.1.1), in order to be exploited for NLP tasks (section 4.1.2).

² <http://sentic.net/senticnet-1.0.zip>

³ <http://sentic.net/api>

4.1.1 Building SenticNet

The development of SenticNet was inspired by SentiWordNet [104], a lexical resource in which each WordNet synset is associated to three numerical scores describing how objective, positive and negative the terms contained in the synset are. Each of the three scores ranges from 0.0 to 1.0 and their sum is 1.0 for each synset. This means that a synset may have non-zero scores for all the three categories, which would indicate that the corresponding terms have, in the sense indicated by the synset, each of the three opinion-related properties only to a certain degree.

The method used to develop SentiWordNet is based on the quantitative analysis of the ‘glosses’ associated to synsets and on the use of the resulting vector representations for semi-supervised synset classification. The three scores are derived by combining the results produced by a committee of eight ternary classifiers, all characterised by similar accuracy levels but different classification behaviour. SentiWordNet currently represents a good resource for opinion mining, however, it contains a lot of noise and it mainly provides opinion polarity at syntactical level, leaving out polarity information for common-sense knowledge concepts such as ‘accomplish goal’, ‘bad feeling’, ‘celebrate special occasion’, ‘lose temper’, or ‘be on cloud nine’, which are usually found in natural language text to express positive and negative viewpoints. To this end, SenticNet was developed (Fig. 4.1).

SenticNet is a publicly available resource for opinion mining that aims to create a collection of commonly used ‘polarity concepts’, that is common-sense concepts with relatively strong positive or negative polarity. Differently from SentiWordNet (which also includes null polarity terms), in fact, SenticNet does not contain concepts with neutral or almost neutral polarity, i.e., concepts with polarity magnitude close to zero. Moreover, while SentiWordNet stores three values for each synset, in SenticNet each concept c is associated to just one value p_c , i.e., a float $\in [-1,1]$ representing its polarity, in order to avoid redundancy and more easily represent SenticNet as a semantic network. Therefore, in SenticNet, concepts like ‘make good impression’, ‘look attractive’, ‘show appreciation’, or ‘good deal’ are likely to have p_c very close to 1, while concepts such as ‘being fired’, ‘leave behind’, or ‘lose control’ are likely to have $p_c \approx -1$. Polarity values are assigned to Open Mind concepts by means of Sentic Activation. In particular, the sentic levels of the Hourglass model are used as input concepts for the U-level, which uses its digital intuition to infer semantically and affectively related concepts. Such seed concepts then crop out to the C-level that exploits spectral association to accordingly spread their activation through the graph structure of AffectNet.

After retrieving polarity concepts through Sentic Activation, they need to be reorganised in a way that they can be represented in a unique and consistent resource. Possible conflicts are handled by discarding duplicate concepts with smaller polarity magnitude since bigger concept polarity values usually correspond to more reliability (higher dot products in the vector space). Since concepts are usually strongly related to just one or two affective dimensions (most of compound emotions are, in fact, given by summing just two elementary emotions), the average magnitude is pretty low.

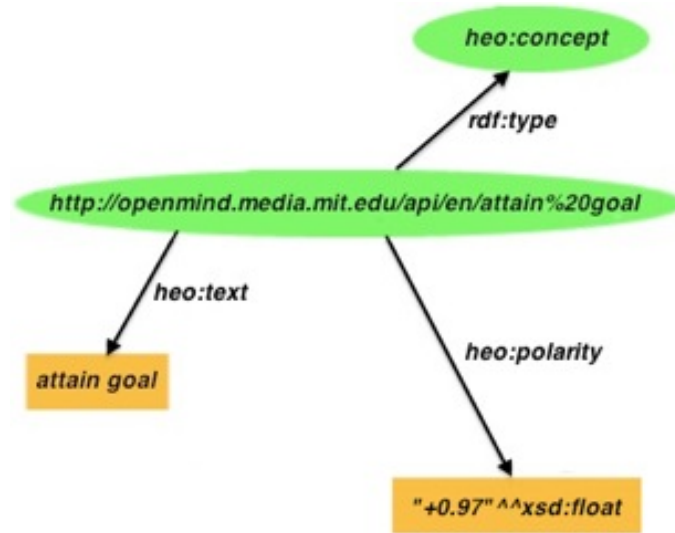


Fig. 4.1 SenticNet is a Semantic Web aware AI resource in which Open Mind concepts are identified using the ConceptNet Web API, and statements are encoded in RDF/XML format on the base of HEO.

Therefore, in order to obtain more homogeneous and intelligible polarity values, a normalisation process over SenticNet is conducted before storing its contents in a Semantic Web aware format. In order to represent SenticNet in a machine-accessible and machine-processable way, results are encoded in RDF triples using a XML syntax. In particular, concepts are identified using the ConceptNet Web API and statements, which have the form `concept-hasPolarity-polarityValue`, are encoded in RDF/XML format on the base of the human emotion ontology (HEO) [118], a high level ontology for human emotions that supplies the most significant concepts and properties which constitute the centrepiece for the description of every human emotion. The main purpose of HEO is to create a description framework that could grant at the same time enough flexibility (by allowing the use of a wide and extensible set of descriptors to represent all the main features of an emotion) and interoperability (by allowing to map concepts and properties belonging to different emotion representation models). HEO was developed in OWL description logic in order to allow a taxonomical organisation of emotion categories and properties restriction to link emotion description made both by category and dimension.

In HEO, for example, Ekman's archetypal emotion of 'joy' represents a superclass for Plutchik's emotions of 'ecstasy', 'joy', and 'serenity'. Using property restriction, Plutchik's emotion of 'joy' can also be defined as an emotion that 'has Pleasantness some float \in [$+G(1/3), +G(2/3)$]', 'interest' as an emotion that 'has Attention \in [$0, +G(1/3)$]', and 'love' as an emotion that 'has Pleasantness some float \in [$0, +G(1)$] and Aptitude some float \in [$0, +G(1)$]'.

In this way, querying a database that supports OWL description logic inference for basic emotions of type ‘joy’ will return not only the emotions expressly encoded as Ekman’s archetypal emotions of type ‘joy’, but also the emotions encoded as Plutchik’s basic emotion of type ‘joy’ and emotions that ‘have Pleasantness some float $\in [+G(1/3), +G(2/3)]$ ’.

4.1.2 Working with SenticNet

SenticNet is freely available online and currently contains more than 5,700 polarity concepts (nearly 40% of Open Mind corpus). It is very easy to interface SenticNet with any kind of opinion mining application and, especially if used within Open Mind software (for a full correspondence of concepts), it is a precise polarity detection tool. In particular, after deconstructing text into concepts through a semantic parser, SenticNet can be used to associate polarity values to these and, hence, infer the overall polarity of a clause, sentence, paragraph, or document by averaging such values.

SenticNet’s capacity of detecting opinion polarity was compared with SentiWordNet’s over a collection of 2,000 patient opinions, of which 57% are labelled as negative, 32% as positive, and the rest as neutral. After extracting concepts from each opinion, relative polarity values were searched both in SentiWordNet and SenticNet and compared with the dataset labels, in order to compute recall and precision rates as evaluation metrics. Results showed SenticNet to be much more accurate than SentiWordNet. The former, in particular, can identify positive opinions with much higher precision (79.1% against 53.8%) and significantly better recall rate (58.4% against 46.5%), for a total F-measure value of 67.1% versus 49.8%. In SenticNet 2⁴ [40], the whole Open Mind corpus is being labelled with polarity values and a list of mood and sentic values is being associated with each common-sense concept, in order to provide the public with a comprehensive semantic resource for easily extracting affective information from natural language text.

4.2 Sentic Neural Networks: Brain-Inspired Affective Reasoning

Sentic Neural Networks (SNNs) [53] represent a brain-inspired model based on the combined use of principal component analysis (PCA) and artificial neural networks (ANNs) on AffectNet. SNNs investigate if an emulation of the biological neural system, represented by two ANNs, could outperform the previously adopted k -nearest neighbour (KNN) clustering approach. PCA is used to organise the space where concepts lie and then an artificial neural network is trained to recognise emotions based on the presented Hourglass model.

⁴ <http://sentic.net/senticnet-2.0.zip>

Contrary to any clustering algorithm, the emotion recognition task is independent from both concepts' absolute and relative positions in the vector space. The eventual aim of the proposed ANN tool developed in this study is to predict which class each concept belongs to (i.e., its level of affective valence in a specific dimension of the Hourglass model). This section, in particular, discusses the model selection (section 4.2.1) and how SNNs can learn and recognise affective common-sense knowledge (section 4.2.2).

4.2.1 Discrete Versus Continuous Approach

Two different approaches may be adopted in order to set up an ANN: a 'discrete' approach (termed *SNN-D*) and a 'continuous' one (*SNN-C*). *SNN-D*, in particular, is expected to return seven different real-valued outputs $y_k \in [0, 1]$ for $k = 1, 2, \dots, 7$, each showing the degree of belonging to a specified affective level, while *SNN-C* provides a single real-valued output $y \in [-1, 1]$, corresponding to the best guess of the level of affective valence. In both cases, a further step is required in order to obtain a final classification output: for *SNN-D* the best selection strategy seems to be the choice of the class with the highest degree of belonging, while for *SNN-C* the easiest approach is to round off the output to get an integer corresponding to the class. Since the task of choosing from these two approaches is not easily solvable a priori, both approaches are adopted and compared in this study.

Therefore, two multi-layer perceptron (MLP) neural networks with three layers (one input layer, one hidden layer, and one output layer) were set up. The input vector $\mathbf{x}^{(k)}$ is built so that $x_{(k)0} = 1$ (the 'bias node') and $[x_{(k)1}, x_{(k)2}, \dots, x_{(k)100}] = \mathbf{a}^{(k)}$ for the k -th concept of the dataset. The target output, for *SNN-D*, is a vector $\mathbf{y}^{(k)}$ having $y_{(k)i} = 1$ if $b_{(k)} = i - 2$ and $y_{(k)i} = 0$ otherwise, for each $i = 1, 2, \dots, 7$. On the other hand, for *SNN-C* the target output is a single value $y_{(k)} = b_{(k)}$. Let us assume that the hidden layer has H neurons; the input and the hidden layers are then linked by equation (4.1) where $w_{mn} \in \mathbf{W}$, which is the matrix of weights (of dimensions $H \times 101$) defined during the network training:

$$h_j = \tanh \left(\sum_{i=0}^{100} w_{ji} x_i \right) \quad j = 1, 2, \dots, H \quad (4.1)$$

The relationships between the hidden and the output layers are expressed in equations (4.2) and (4.3) for *SNN-D* and *SNN-C* respectively, where $v_{mn} \in \mathbf{V}$, which is a matrix of weights (of dimensions $7 \times H$) and \mathbf{v} a vector of weights (of dimension H) built during the network training. The output is computed as follows:

$$y_j = \tanh \left(\sum_{i=1}^H v_{ji} h_i \right) \quad j = 1, 2, \dots, 7 \quad (4.2)$$

$$y = \tanh(\mathbf{v}^T \cdot \mathbf{h}) \quad (4.3)$$

It is worth noting that different choices of the activation functions, as well as other design choices, are possible; however, a definition of the best structure of the used neural networks besides being quite difficult to state (due to, for example, the dataset dependency) is beyond the scope of this study. Finally, further transformations of $\mathbf{y} \in \mathbb{R}^7 \mapsto b^* \in \Theta$ and $y \in \mathbb{R} \mapsto b^* \in \Theta$ are required. They are proposed in equations (4.4) and (4.5):

$$b^* = \mathbf{d}^T \cdot [-G(1), -G(2/3), -G(1/3), 0, +G(1/3), +G(2/3), +G(1)] \quad (4.4)$$

$$b^* = \text{round}(y) \quad (4.5)$$

where $d_i = \delta_{i,m}$ for $i = 1, 2, \dots, 7$, δ being the Kronecker's delta, $m = i \mid y_i = \max_j y_j$ and $\text{round}(x)$ a function $\mathbb{R} \mapsto \mathbb{Z}$ rounding x up to the closest integer. The final structure of the two networks is depicted in Figure 4.2 and Figure 4.3.

4.2.2 Affective Learning

Iterative gradient methods such as Levenberg-Marquardt (LM) [200] are in widespread use for MLP training. However, LM depends sensitively on the initial guess, necessitating repeated runs. This, combined with LM's high per-step cost, makes its computational burden quite high. Therefore, SNN training is conducted using the gradient descent optimisation algorithm with the inclusion of a momentum term, which has been proven to improve the algorithm speed of convergence [236].

The matrices (or vectors) of weights (\mathbf{U}) are then updated at each iteration τ according to equation (4.6), where \mathbf{E} is calculated with $e_j = b_j - b_j^*$ for the second layer of weights (where $j = 1, 2, \dots, 7$ for *SNN-D* and $j = 1$ for *SNN-C*), and $e_j = \sum_k \frac{b_j^*(\tau)}{b_j^*(\tau-1)} u_{jk}$ for the first layer of weights respectively.

$$\Delta \mathbf{U}_{(\tau)} = -\nu \nabla \mathbf{E}(\mathbf{U}_{(\tau)}) + \mu \Delta \mathbf{U}_{(\tau-1)} \quad (4.6)$$

In order to evaluate the designed system, SNNs were tested on the benchmark for affective common-sense knowledge (BACK). To avoid the risk of over-fitting, a cross-validation approach was adopted. The networks were trained 10 times (10-fold cross-validation), each of which excluded 10% of dataset entries that are used for evaluating the performance of the system; the excluded 10% is then cycled so that, at the end of all simulations, each dataset entry has been used exactly once to test the system. In order to evaluate the accuracy of the model, the percentage of entries where $b^* = b$ ('strict accuracy') was considered. However, since the used dataset can include noise and entries may incorporate a certain degree of subjectiveness, this criterion was relaxed by considering the accuracy of entries which have $|b^* - b| \leq 1$ ('relaxed accuracy').




Fig. 4.2 The structure of the *SNN-C* model. In the ‘continuous’ approach, the SNN provides a single real-valued output corresponding to the best guess of the level of affective valence.




Fig. 4.3 The structure of the *SNN-D* model. In the ‘discrete’ approach, the SNN returns seven different real-valued outputs, each showing the degree of belonging to a specified affective level.

The results of the proposed SNNs are tabulated in Table 4.1, where they are compared with the previously adopted KNN approach and a random classifier. A trial-and-error approach was adopted for the network parameters tuning: for *SNN-C* the best performance was obtained after 3 iterations (when the error stopped decreasing significantly) with $H = 10$, learning rate $\nu = 0.1$, momentum factor $\mu = 0.2$; for *SNN-D* the best set of parameters obtained were $H = 15$, $\nu = 0.05$, $\mu = 0.05$ with best performance reached after an average of 10 iterations.

Model	Strict Accuracy	Relaxed Accuracy
Random	14.3%	40.1%
KNN	41.9%	72.3%
<i>SNN-D</i>	43.8%	76.2%
<i>SNN-C</i>	38.3%	82.7%

Table 4.1 SNNs performance comparison with a KNN model and a random classifier. SNN-C performs better in the ‘relaxed accuracy’ case, while SNN-D is more suitable for the ‘strict accuracy’ case.

As shown in Table 4.1, the proposed SNN approaches outperform the previously adopted KNN model as well as the random classifier. Both proposed models improved the ‘relaxed accuracy’ (with the *SNN-C* producing a considerable 10% performance improvement) while the *SNN-D* was able to outperform the benchmark for the ‘strict accuracy’ case.

4.3 Open Mind Common Sentics: An Emotion-Sensitive IUI

In AffectNet, the general common-sense knowledge contained in ConceptNet is exploited to spread affective information from selected affect seeds to other concepts. Besides exploiting the emotional content of the Open Mind corpus, AffectNet is also enriched by collecting new affective common-sense knowledge through state-of-the-art crowd sourcing and games with a purpose (GWAP) techniques (section 4.3.1). In particular, Open Mind Common Sentics [61] is developed. Open Mind Common Sentics is an emotion-sensitive IUI that serves both as a platform for affective common-sense acquisition and as a publicly available NLP tool for extracting the cognitive and affective information associated with short texts (section 4.3.2).

4.3.1 Games for Knowledge Acquisition

The Casual Games Association⁵ reports more than 200 million casual gamers worldwide this year. People play games for different reasons, e.g., to relax, to be entertained, for the need of competition, and to be thrilled [231]. Additionally, they want to be challenged, both on a mental and on a skill-based level. Such army of gamers could be exploited for performing tasks that are relatively easy to complete by humans, but computationally rather infeasible to solve [6].

The idea is to integrate such tasks as goal of games [5] by producing a win-win situation where people have fun playing games while actually doing something useful. The nature of these games, in fact, focuses on exploiting player inputs to both create meaningful data and provide a funnier game experience [279].

⁵ <http://casualgamesassociation.org>




Fig. 4.4 A screenshot of PopVideo. Two or more players are shown the same video and have to timely describe the objects that appear in the video in order to earn points.

Such human-based computational power can be exploited for tasks such as video annotation, e.g., OntoTube⁶ [258], PopVideo⁷ (Fig. 4.4), Yahoo's Videotaggame [315], and Waisd⁸ [3], in which two players have to timely agree on a set of tags about the same streaming YouTube⁹ video. Similarly, in ESP game¹⁰ [7] and Google Image Labeler (before being discontinued last September) players have to consensually guess content objects or properties of random images by simultaneously typing what they see. Other games for image annotation include Matchin¹¹ [125], which focuses on image perceived quality by asking players to pairwise choose the picture they like better, Phetch [8], a game that collects explanatory descriptions of images in order to improve accessibility of the Web for the visually impaired by letting a player describe an image and others retrieve it using an image search en-

⁶ <http://ontogame.sti2.at/games>

⁷ <http://gwap.com/gwap/gamesPreview/popvideo>

⁸ <http://waisda.nl>

⁹ <http://youtube.com>

¹⁰ <http://gwap.com/gwap/gamesPreview/espgame>

¹¹ <http://gwap.com/gwap/gamesPreview/matchin>

gine, Peekaboom [10], which focuses on locating objects within images by letting a player reveal specific parts of an image in order for the other to guess the correct object name, Squigl¹², in which players have to spot objects in images previously annotated within ESP Game, and PictureThis, which asks players to choose, among a set of images, the one that best suits the given query. Among games for image annotation, there are also games for streamlining the robustness evaluation of CAPTCHAs, namely: Magic Bullet¹³ [311], a team game in which players need to agree on the meaning of CAPTCHAs, and TagCaptcha¹⁴ [201], where players are asked to quickly describe CAPTCHA images with one word each. GWAPs are also exploited to automatically tag music with deeper semantic labels. HerdIt¹⁵ [19], for example, asks players have to accomplish different tasks related to the song they are listening to, while in Tagatune¹⁶ [165] two players have to listen to an audio file and describe to the other what they are hearing, in order for him/her to decide whether the game has played the same soundtrack to both or not.

Several games have also been designed for text annotation. Verbosity¹⁷ [9], for example, is a real time quiz game for collecting common-sense facts. In the game, two players take different roles at different times: a narrator, who has to describe a word using templates, and a guesser, who has to guess such word in the shortest time possible. Sentiment Quiz, instead, gathers information about the polarity associated to words. It asks its players to evaluate random words on a five grade scale, from very negative over neutral to very positive. Phrase Detectives¹⁸ tries to identify relationships between phrases and other phrases or sentences.

Another approach to collecting common-sense knowledge is the FACTory Game¹⁹ [170], published by Cycorp. FACTory randomly chooses facts from Cyc and presents them to players, in order for them to guess whether a statement is true, false, or does not make sense. A variant of the FACTory game is the Concept Game on Facebook²⁰ [137], which collects common-sense knowledge by proposing random assertions to users in a slot machine fashion and asking them to decide whether this assertion is meaningful or not (Fig. 4.5). Page Hunt²¹ [185] is a GWAP for the annotation of websites. It allows to index websites and, hence, to improve the search index of a search engine (Microsoft Bing). The player gets assigned a random website and is asked to describe it with keywords. The game then shows players the top five page hits for the entered keywords and they are rewarded depending on how high ranked the previously assigned web page is in the result set.

¹² <http://gwap.com/gwap/gamesPreview/squigl>

¹³ <http://homepages.cs.ncl.ac.uk/jeff.yan/mb.htm>

¹⁴ <http://dolphin.unige.ch/tagcaptcha>

¹⁵ <http://herdit.org/music>

¹⁶ <http://gwap.com/gwap/gamesPreview/tagatune>

¹⁷ <http://gwap.com/gwap/gamesPreview/verbosity>

¹⁸ <http://anawiki.esex.ac.uk/phrasedetective>

¹⁹ <http://game.cyc.com>

²⁰ <http://apps.facebook.com/conceptgame>

²¹ <http://pagehunt.msrlivelab.com>



Fig. 4.5 A screenshot of Concept Game, a turn-based game taking advantage of the Facebook platform for adding competitive flavours and social aspects to the collection of common-sense knowledge.

Other GWAPs engage players in building ontologies. OntoPronto [258], for example, is a quiz game for vocabulary building that attempts to build a huge domain ontology from Wikipedia articles. This is achieved by mapping random articles to the most specific class of the Proton ontology²² using the *subClassOf* relationship. Virtual Pet Game²³ [159] aims to construct a semantic network that encodes common-sense knowledge. The game is built on top of PPT, a popular Chinese bulletin board system that is accessible through a terminal interface. Each player owns a pet, which they should take care of by asking and answering questions. The pet in this game is just a substitute for other players, who receive such questions and answers, and have to respond or validate them. Rapport Game²⁴ [159], similarly to Virtual Pet Game, exploits player labour for constructing a semantic network that encodes common-sense knowledge. Rapport Game, however, is built on top of Facebook and uses direct interaction between players. An interesting game for the creation of formal domain ontologies from Linked Open Data is Guess What?!²⁵ [190]. Given a seed concept, a player has to find a matching URI in DBpedia, Free-

²² <http://proton.semanticweb.org>

²³ http://agents.csie.ntu.edu.tw/commonsense/cate2_1_en.html

²⁴ <http://apps.facebook.com/conceptnet>

²⁵ <http://nitemaster.de/guesswhat/manual.html>

base, and OpenCyc. The resulting labels/URIs are analysed by simple NLP tools in order to identify expressions that can be translated into logical operators and break down complex descriptions into small fragments. The game starts with the most general fragment and, at each round, a more specific fragment is connected to it through a logical operator, with players having to guess the concept described by it.

There are GWAPs that try to align ontologies. Wordhunger²⁶, for example, is a web-based application that maps WordNet synsets to Freebase. Each game round consists of a WordNet term and up to three suggested possible Freebase articles, among which players have to select the most fitting (or pass, or select ‘no match’). SpotTheLink²⁷ is a two player game focusing on the alignment of random concepts from the DBpedia Ontology²⁸ to the Proton upper ontology. Each player has to select Proton concepts that are either the same as or more specific than a randomly selected DBpedia concept. The data generated by SpotTheLink is a SKOS mapping between the concepts of the two input ontologies.

Based on Wikipedia, there are three Wikiracing game, The Wiki Game²⁹, Wikispeedia³⁰, and WikipediaMaze³¹, where the objective is to find connections between two Wikipedia articles by clicking links within the text. WikipediaGame and Wikispeedia focus on completing the race faster and with fewer clicks than other players. In WikipediaMaze, instead, players are allowed to create races for each other and are incentivised to create and play races by earning badges.

4.3.2 Collecting Affective Common-Sense Knowledge

Human emotions and their modelling are increasingly understood to be a crucial aspect in the development of intelligent systems [228, 199]. Emotions are a basic part of human communication and have therefore to be taken into account for the development of more effective interfaces for human-machine communication such as chat systems, e-house, e-learning, e-health, or emphatic voice boxes. Besides general common-sense knowledge, the Open Mind corpus also contains affective information, e.g., “a gift is for celebrating a birthday” or “making a mistake causes embarrassment”, as common-sense encompasses, among many other aspects, also knowledge about the emotional facets of typical everyday life. However, the amount of affective information contained in the Open Mind corpus is still very limited, as relationships such as ArisesEmotion, MakesFeel, or AffectivelyRelated are missing from the set of properties that are currently used for collecting pieces of common-sense knowledge from the public.

²⁶ <http://wordhunger.freebaseapps.com>

²⁷ <http://semanticgames.org/2011/11/spotthelink>

²⁸ <http://dbpedia.org/Ontology>

²⁹ <http://thewikigame.com>

³⁰ <http://http://www.cs.mcgill.ca/~rwest/wikispeedia>

³¹ <http://www.wikimaze.me>

Since computers now have the ability to search vast amounts of data in little time, the use of a search engine to collect the affective information needed is pretty tempting. To this end, different lexical patterns, termed sentic patterns, are used for extracting affective information from the Web. Such patterns are built using label sequential rules (LSRs), which are generated from sequential patterns in data mining [180]. A rule is of the form $X \rightarrow Y$, where Y is a sequence and X is a sequence produced from Y by replacing some of its items with wildcards, denoted by a ‘*’, which can match any item. During the learning process, each segment is converted to a sequence. Each sequence element is a word, which is represented by both the word itself and its POS tag in a set. In the training data, all concepts are manually labelled and replaced by the label *\$concept*. A concept can be expressed with a noun (NN), adjective (JJ), verb (VB), or adverb (RB). The labels and their POS tags used in mining LSRs are $\{\$concept, NN\}$, $\{\$concept, JJ\}$, $\{\$concept, VB\}$, and $\{\$concept, RB\}$, where *\$concept* denotes a concept to be extracted. For example, the sentence segment “chocolate makes me feel happy” is turned into the sequence $\langle \{chocolate, NN\} \{make (me|you) feel, VB\} \{happy, JJ\} \rangle$. After labelling, it becomes $\langle \{\$concept, NN\} \{make (me|you) feel, VB\} \{happy, JJ\} \rangle$.

All the resulting sequences are then used to mine LSRs. A typical rule, for example, is $\langle \{*, NN\} \{put (me|you) on, VB\} \{cloud\ nine, NN\} \rangle \rightarrow \langle \{\$concept, NN\} \{put (me|you) on, VB\} \{cloud\ nine, NN\} \rangle$ confidence = 80%, where the confidence is the conditional probability, $\Pr(Y|X)$, which measures the accuracy of the rule. Concept extraction is performed by matching the patterns with each sentence segment in a new web page to extract affective information about concepts contained AffectNet. That is, the words in the sentence segment that both match *\$concept* in a pattern and an AffectNet concept are extracted. In the pattern match, only the right-hand side of each rule is used. In rule generation, both the right- and the left-hand sides are needed to compute the conditional probability or confidence.

Such patterns yield some useful data, however, they are not good enough for the purposes of the proposed study for three reasons. Firstly, most of the affective common-sense knowledge aimed to be collected is so obvious that no one has bothered to record it. Secondly, there exists incorrect knowledge on the Web (for example, the query “plastic can love” returns 33,200 results on Google, while “plastic cannot love” returns just a few links). Thirdly, the text on the Web is unstructured and turning it into a directly useful format is a non-trivial task. For these reasons, crowd sourcing techniques are mainly used. Distributed online knowledge acquisition projects have become quite popular in the past years. Examples include Freebase³², with its 1,450 concepts, WikiTaxonomy, counting 127,000 concepts, YAGO³³, with 149,162 instances, NELL³⁴, containing 959,654 beliefs, ProBase³⁵, Microsoft’s universal probabilistic ontology, and the different projects associated with the Open Mind Initiative, e.g., OMCS, Open Mind Word Expert [194], an ac-

³² <http://freebase.com>

³³ <http://mpi-inf.mpg.de/yago-naga/yago>

³⁴ <http://rtw.ml.cmu.edu/rtw>

³⁵ <http://research.microsoft.com/probase>

tive learning system that aims to create large annotated corpora, and Open Mind Indoor Common Sense [124], which aims to develop intelligent mobile robots for use in home and office environments. In a similar fashion to the Open Mind family of distributed knowledge capture projects, Open Mind Common Sentic, aims to collect affective common-sense knowledge for sentiment analysis (Fig. 4.6).

Whereas previous approaches have mainly relied on paid experts or unpaid volunteers, much stronger emphasis is hereby put on creating a system that is appealing to a large audience of people, regardless of whether or not they are interested in contributing to AI. The fundamental aim of Open Mind Common Sentic, in fact, is to transform as much as possible the activity of entering knowledge into an enjoyable interactive process. To this end, the system adopts a two-fold strategy: crowd sourcing, that is challenge volunteers over the Web through mood-spotting and fill-in-the-blank questions, in the same wake as Open Mind Commons [266], and GWAPs, in the same wake as Verbosity and ESP game. In particular, the mood-spotting questions consist in asking users to select an emoticon according the overall affect they can infer from a given sentence. The fill-in-the-blank questions, in turn, are sentences to be completed such as “opening a Christmas gift makes feel ___”. Sometimes, one or more taboo affect concepts are shown in the game window.

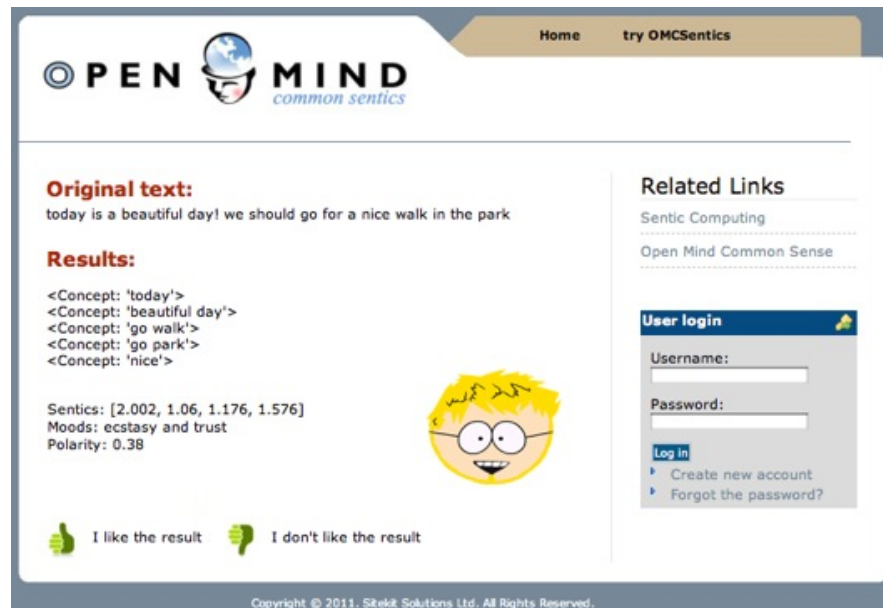


Fig. 4.6 A typical output obtained by inserting some text into Open Mind Common Sentic interface. The system provides a list of extracted concepts, a list of sentics, and a polarity value.

Such concepts are entries that have been validated a sufficient amount of times and, hence, they are not valid inputs any more (in order to collect synonyms or alternatives of a given affective common-sense concept). As for the GWAPs, the Hourglass Game was developed (Fig. 4.7). It is a speed game consisting in selecting, from the Hourglass model, the sentic level that is most likely associated with a given affective concept. Players earn points not only according to accuracy, but also quickness in clicking on the right area of the Hourglass. The game is quite engaging, although very simple, and players like to challenge each other to see who has higher emotional quotient (EQ), but users are not too keen on playing more than once. What is lacking from most of crowd-sourcing and GWAP techniques, in fact, is stickiness. GWAPs can be fun to play for a relatively short period of time, but then players are not too keen on returning.

In other words, GWAPs generally have a pretty low sticky factor. The sticky factor is defined as the amount of daily active users (DAUs) of an application divided by the number of monthly active users (MAUs). MAU is the most-quoted measure of a game's size, but it is effective only to discuss size or reach, not engagement. DAU, in turn, can be a very valuable number as it relates how much activity a game is seeing on a daily basis, but it falls into the same trap as MAU in that it does not discriminate between retention and acquisition.

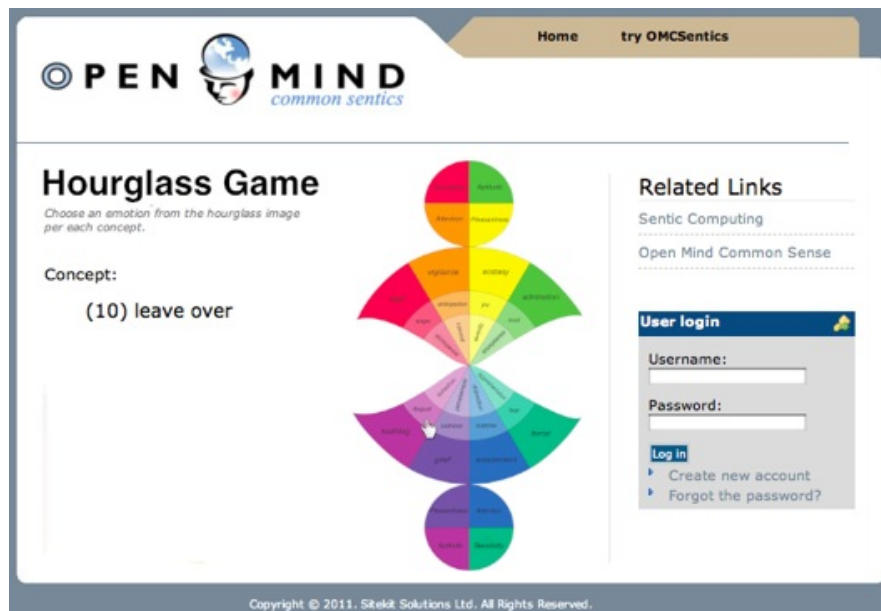


Fig. 4.7 A screenshot of the Hourglass Game. The GWAP aims to collect affective common-sense knowledge from the general public by engaging users in a game where precision and speed are awarded.

The single-most important metric for engagement is stickiness, i.e., DAU/MAU, which allows to more accurately calculate repeat visits and average knowledge acquired per user (AKAPU). A key for driving the sticky factor, besides great game play, is the ability of the application to prompt users to reach out to their friends, e.g., via stories and pictures about their gameplay. To this end, Sentic Pet is being developed (Fig. 4.8). Sentic Pet is a massively multiplayer online (MMO) game in which players have to raise and take care of their own pets. Unlike old-style tamagotchi games, in Sentic Pet, raising and caring pets is not about cleaning, feeding, and petting them, but rather training them, both at mental and skill level, by playing mini-GWAPs. Targeting players of a wide age range (10 to 50 year old), the game should appeal anyone who enjoyed and enjoys PetVille³⁶ or FarmVille³⁷.

Players start from level 1 with their pet being a baby born, having very little affective common-sense knowledge. The game involves balancing two main activities: training the pet and testing its skills by challenging other players. Training does not involve simply teaching the pet new knowledge, but also refining acquired knowledge. Challenges can be taken both at mental and skill level, which involve different kinds of activities.

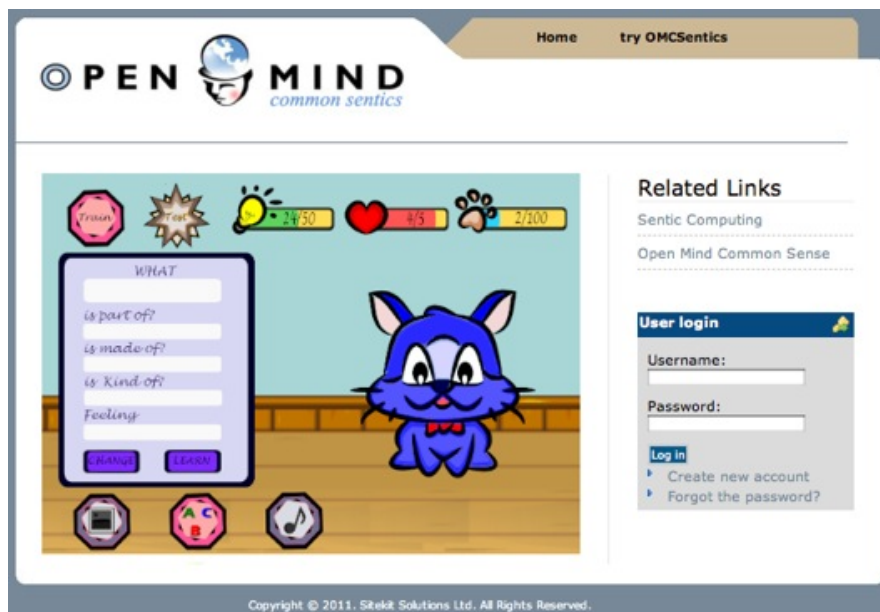


Fig. 4.8 A screenshot of Sentic Pet. Icons in the up-right corner specify the abilities of pets, while icons in the up-left corner allow players to train or test these according to different modalities.

³⁶ <http://petville.com>

³⁷ <http://farmville.com>

At mental level, for example, pets can be challenged according to different modalities, e.g., affective vocabulary learning (in the same wake of Verbosity) or affective meaning of images (in the same wake of ESP game). Pets can level-up according to the combination of IQ (light bulb icon) and EQ (heart icon) points earned playing the mini-GWAPs. Data are validated by majority and reputation, that is, the confidence score with which a piece of affective common-sense knowledge is saved into the knowledge base depends on how many players validated it and on the expertise level of these. Open Mind Common Sentic is an example of an emerging class of games that can be considered ‘human algorithms’, since humans act as processing nodes for problems that computers cannot yet solve. By providing an incentive for players, a large quantity of computing power can be gained and harnessed for multiple applications. Constructing a complete affective common-sense database would be extremely beneficial for many communities, e.g., sentiment analysis and HCI, and Open Mind Common Sentic can be highly effective in doing so.

4.4 IsaCore: A Common and Common-Sense Knowledge Base

Common-sense is the key to natural language understanding. However, it is not the only kind of knowledge involved in standard human-to-human communication. When communicating with each other, in fact, people usually refer to existing facts and circumstances and build new useful, funny, or interesting information on the top of those. This common knowledge comprehends information usually found in news, articles, debates, lectures, etc. (factual knowledge), but also principles and definitions that can be found in collective intelligence projects such as Wikipedia (vocabulary knowledge). Attempts to build a common knowledge base are countless and comprehend both resources crafted by human experts or community efforts, such as WordNet and Freebase, and automatically-built knowledge bases, such as Wiki-Taxonomy, YAGO, and ProBase, which is the most comprehensive resource with its 12 million concepts. In this work, in fact, the richness of Probase (section 4.4.1) is exploited to enhance the topic-spotting capabilities of the opinion mining engine. In particular, Probase is blended with ConceptNet in order to build possibly the most comprehensive resource of common and common-sense knowledge for sentiment analysis (section 4.4.2).

4.4.1 Probase

Probase is a research prototype that aims to build a unified taxonomy of worldly facts from web data and search log data [309]. The taxonomy consists of concepts (e.g., *painter*), instances (e.g., “Leonardo da Vinci”), attributes and values (e.g., “Leonardo’s birthday is April 15, 1452”), and relationships (e.g., “Mona Lisa is painted by Leonardo”). Such information is learned iteratively from 1.68 billion

web pages in Bing web repository. The taxonomy is probabilistic, which means that every claim in ProBase is associated with some probabilities that model the claim's correctness, ambiguity, and other characteristics. The probabilities are derived from evidences found in web data, search log data, and other available data. The core taxonomy consists of the *IsA* relationships extracted by using syntactic patterns such as the Hearst patterns [135]. For example, a part of text like "... artists such as Pablo Picasso ..." is considered as a piece of evidence for the claim that "Pablo Picasso" is an instance of the concept *artist*. Next, given a concept *C*, syntactic patterns such as "What is the A of B" are used to find its attributes (where B is an instance of C, and A is the attribute of interest). For example, sentences such as "What is the capital of China?" and "What is the GDP of Japan?" suggest that "capital" and "GDP" are candidate attributes of concept *country* (Fig. 4.9). Furthermore, every claim in Probase is associated with a few scores that model the consensus, typicality, ambiguity, and other characteristics of the claim.

In Table 4.2, Probase is compared with WordNet and Freebase. WordNet specialises in the linguistics of English words. For the word *cat*, WordNet has detailed descriptions of its various senses, although many of them are rarely used, or even unknown to many people (e.g., *gossip* and *woman* as concepts for *cat*). Also, it does not contain information for entities such as *IBM*, which is not considered as a word. Freebase, on the other hand, contains limited number of concepts for the word *cat*. In fact, the categories there are biased and sometimes inaccurate. For example, Freebase's concept space is biased toward entertainment, media related concepts.

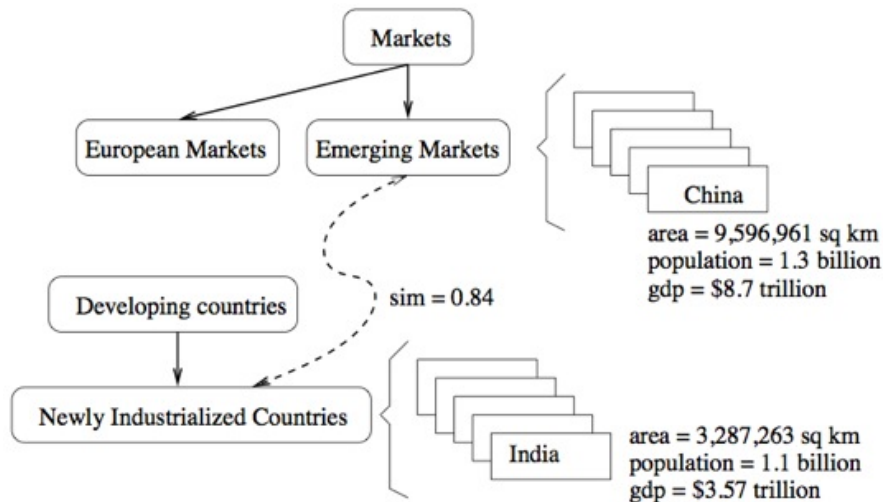


Fig. 4.9 Probase consists of concepts (e.g., *emerging markets*), instances (e.g., "China"), attributes and values (e.g., "China's population is 1.3 billion"), and relationships between concepts.

More importantly, the categories in WordNet and Freebase are not ranked or scored, and users cannot tell the difference in terms of their importance or typicality. In comparison, the concepts in Probase are more consistent with human's common knowledge. Concepts such as *gossip* and *woman* for *cat* are either not included or ranked very low because people rarely associate them with *cat*. In addition, for a word such as *language*, Probase indicates it can be both an instance on its own or an attribute of some concepts. Thus, Probase not only represents the largest taxonomy currently available, but it also provides information that is missing in commonly used resources such as WordNet and Freebase. For this reasons, Probase is selected as the common knowledge base to be exploited by the proposed opinion mining engine for the detection of topics in opinionated text.

Term	WordNet Hypernyms	Freebase Types	Probase Concepts
Cat	Feline; Felid; Gossip; Gossiper; Gossipmonger; Rumormonger; Rumour-monger; Newsmonger; Stimulant; Stimulant drug; Excitant; Tracked vehicle; ...	TV episode; Creative work; Musical recording; Organism classification; Musical release; Book; Musical album; Film character; Publication; Character species; Animal; Domesticated animal; ...	Animal; Pet; Species; Mammal; Small animal; Thing; Mammalian species; Small pet; Animal species; Carnivore; Domesticated animal; Companion animal; Exotic pet; Vertebrate; ...
IBM	N/A	Companies listed on the New York Stock Exchange; Cloud computing providers; Companies based in Westchester County, New York; Multinational companies; Software companies of the United States; Top 100 US Federal Contractors; ...	Company; Vendor; Client; Corporation; Organisation; Manufacturer; Industry leader; Firm; Brand; Large company; Fortune 500 company; Technology company; Supplier; Software vendor; Global company; Technology company; ...
Language	Communication; Auditory communication; Word; Higher cognitive process; Faculty; Mental faculty; Module; Text; Textual matter;	Written work; Musical recording; Musical artist; Musical album; Literature subject; Query; Periodical; Type profile; Journal; Quotation subject; Broadcast genre; Periodical subject; Video game content descriptor; ...	Instance of: Cognitive function; Knowledge; Cultural factor; Cultural barrier; Cognitive process; Cognitive ability; Cultural difference; Ability; Attribute of: Film; Area; Book; Publication; Magazine; Work; Program; Media; ...

Table 4.2 Comparison between Probase and two popular common knowledge bases. Probase overcomes problems such as incompleteness of WordNet hypernyms and biased knowledge of Freebase.

4.4.2 Building the Instance-Concept Matrix

Common knowledge is factual and vocabulary knowledge that is known by everyone, on which people base their arguments when communicating with each other. Common knowledge, however, is just a surface layer of human communication and alone is not enough for understanding natural language. To this end, common-sense knowledge is needed. Common-sense knowledge is very important as it spans a huge portion of human experience. It is, however, typically omitted from social communications and, hence, cannot be retrieved from the Web.

As a subsumption common knowledge base, in fact, Probase lacks information like “a dog is a best friend” (rather than simply an ‘animal’) or “a rose is a kind of meaningful gift” (rather than simply a kind of ‘flower’), i.e., common-sense that is not usually stated in web pages (or, at least, not that often to be extracted by Hearst patterns with a high enough confidence score). To overcome this problem, Probase is enriched with complementary hyponym-hypernym common-sense knowledge from ConceptNet. In particular, all the assertions involving *IsA* relationships with a non-null confidence score, such as “dog is man’s best friend” or “a birthday party is a special occasion”, are extracted from the Open Mind corpus. Such assertions are exploited to generate a directed graph of about 13,000 nodes (interconnected by *IsA* edges), representing subsumption common-sense knowledge.

One of the advantages of this new semantic network is that it allows to discriminate nodes between instance and concept, although in ConceptNet there is no such distinction. Instances, in fact, are represented by nodes with high out-degree and null in-degree, while concepts are nodes with high in-degree and null or low out-degree (as concepts can also be instances of other concepts). Another advantage of this representation is that, since the new network is an unlabelled directed graph, the semantic relatedness of concepts is retained when applying dimensionality reduction techniques. Building a vector space representation of standard ConceptNet, in fact, is good in a meta-heuristic sense, but we lose information about the relations governing how concepts self-organise in the multi-dimensional space, that is, we know that concepts close to each other are semantically related, but we cannot say which specific relations make these concepts similar. If we apply dimensionality reduction on the *IsA* version of ConceptNet, instead, we know that common-sense concepts self-organise according to the *IsA* relationship, e.g., if ‘dog’ and ‘cat’ are close to each other in the resulting vector space is because they are hypernyms of the same concepts.

The ensemble of Probase and ConceptNet, termed IsaCore [59] (beta version available for download³⁸), is obtained by first representing Microsoft’s knowledge base as a $2,715,218 \times 1,331,231$ matrix. Such hypernym-hyponym matrix is built out of 23,066,575 *IsA* triples with the form $\langle \text{instance, concept, confidence score} \rangle$, which are reorganised as rows (e.g., ‘pablo picasso’), columns (e.g., ‘artist’) and entries (e.g., ‘0.91’) of the matrix, respectively.

³⁸ <http://sentic.net/isacore.zip>

Performing reasoning on this matrix as it is would not be very convenient, because it is a very large and fat matrix that contains noise and multiple forms, since all of the evidences are automatically extracted from the Web. To this end, the matrix is firstly cleaned by applying different NLP techniques; secondly, its consistency is enhanced and its sparseness further reduced by adding complementary common-sense knowledge. In particular, three main issues need to be solved, namely: multiple word forms, multiple concept forms, and low connectivity. The first issue is addressed by processing both subjects and objects of triples with OMCS lemmatiser, which groups together the different inflected forms of words (different cases, plurals, verb tenses, etc.) so that they can be stored in IsaCore as a single item. In case of duplicates, the triple with higher confidence score is considered. To perform multiple concept form reconciliation, both word similarity and multi-dimensionality reduction techniques are exploited. The concept ‘barack obama’, for example, appears in the triples in many different forms such as ‘president obama’, ‘mr barack obama’, ‘president barack obama’, etc. Trying to disambiguate this kind of instances *a priori*, by simply using word similarity, could be dangerous as concepts like ‘buy christmas present’ and ‘present christmas event’ have very different meanings, although they have high word similarity. Hence, an *a posteriori* concept deduplication is performed by exploiting concept semantic relatedness, after IsaCore is built.

That is, concepts with high word similarity are merged together just if they are close enough to each other in the vector space generated from IsaCore. As for the multiple concept senses, dimensionality reduction techniques are employed for both *a priori* and *a posteriori* disambiguation. For the former, in particular, a vector space of concepts is built and dissimilarities between these are exploited to find instances likely to have multiple meanings. For the latter, the vectors corresponding to the instances found in an ambiguous sentence are averaged together to perform a context-dependent coarse sense disambiguation.

As for IsaCore’s connectivity, in order to apply dimensionality reduction techniques on it for finding similar patterns, the matrix is desired to be as less sparse as possible. To this end, hapax legomena, that is, instances/concepts with singular out-/in-degree, are to be discarded. These nodes can be useful for specific tasks such as finding the meaning of uncommon instances or give an example of a rare concept. For more general reasoning tasks, however, hapax legomena are very bad as they enlarge dimensionality without providing overlapping information that can be useful for finding similar patterns and perform analogies.

In this work, not only hapax legomena are discarded, but also the other nodes with low connectivity, in order to heavily reduce IsaCore’s sparseness. In particular, a trial-and-error approach is used and the best trade-off between size and sparseness is achieved by setting the minimum node connectivity equal to 10. This cut-off operation leaves out almost 40% of nodes and makes IsaCore a strongly connected core in which common and common-sense knowledge coexist, i.e., a matrix 340,000 × 200,000 whose rows are instances such as ‘birthday party’ and ‘china’, whose columns are concepts like ‘special occasion’ and ‘country’, and whose values indicate truth values of assertions.

4.5 Opinion Mining Engine: Structuring the Unstructured

In order to effectively mine and analyse opinions and sentiments, it is necessary to bridge the gap between unstructured natural language data and structured machine-processable data. To this end, an intelligent software engine has been proposed [39] that aims to extract the semantics and sentics, that is the cognitive and affective information, associated with natural language text, in a way that the opinions and sentiments in it contained can be more easily aggregated and interpreted. The engine exploits graph mining and multi-dimensionality reduction techniques on IsaCore and AffectNet respectively, and it is based on the Hourglass model (Fig. ??).

Several other affect recognition and sentiment analysis systems [184, 11, 178, 93, 86, 274, 94] are based on different emotion categorisation models, which generally comprise a relatively small set of categories (Table 4.3). The Hourglass of Emotions, in turn, allows the opinion mining engine to classify affective information both in a categorical way (according to a wider number of emotion categories) and in a dimensional format (which facilitates comparison and aggregation). Such engine, in particular, consists of four main components (section 4.5.1): a pre-processing module, which performs a first skim of the opinion; a sentic parser [57], whose aim is to extract concepts from the opinionated text; the IsaCore module, for inferring the semantics associated with the given concepts; the AffectiveSpace module, for the extraction of sentics. Eventually, this section illustrates an output example of the engine, given a short natural language sentence as input, and provides a thorough evaluation of the system (section 4.5.2).

Study	Techniques	Model	Corpora	Knowledge Base
[178]	NB, SVM	2 categories	Political articles	None
[93]	LSA, MLP, NB, KNN	3 categories	Dialogue turns	ITS interaction
[94]	Cohesion indices	4 categories	Dialogue logs	ITS interaction
[86]	VSM, NB, SVM	5 categories	ISEAR	ConceptNet
[274]	WN presence, LSA	6 categories	News stories	WNA
[184]	WN presence	6 categories	Chat logs	WNA
[11]	Winnow linear, C4.5	7 categories	Children stories	None
[48]	VSM, KNN	24 categories	LiveJournal	ConceptNet, WNA
[39]	VSM, <i>k</i> -means	24 categories	YouTube, LiveJournal	ConceptNet, WNA, HEO
[53]	VSM, ANN	24 categories	PatientOpinion	ConceptNet, WNA
[59]	VSM, <i>k</i> -medoids	24 categories	Twitter, LiveJournal, PatientOpinion	ConceptNet, Probase

Table 4.3 An overview of recent model-based sentiment analysis systems. Studies are divided by techniques applied, number of categories of the model adopted, corpora and knowledge base used.

4.5.1 Constitutive Modules

The pre-processing module firstly exploits linguistic dictionaries to interpret all the affective valence indicators usually contained in opinionated text, e.g., special punctuation, complete upper-case words, cross-linguistic onomatopoeias, exclamation words, degree adverbs, and emoticons. Secondly, the module detects negation and spreads it in a way that it can be accordingly associated to concepts during the parsing phase. Handling negation is an important concern in opinion- and sentiment-related analysis, as it can reverse the meaning of a statement. Such task, however, is not trivial as not all appearances of explicit negation terms reverse the polarity of the enclosing sentence and that negation can often be expressed in rather subtle ways, e.g., sarcasm and irony, which are quite difficult to detect. Lastly, the module converts text to lower-case and, after lemmatising it, splits the opinion into single clauses according to grammatical conjunctions and punctuation.

The sentic parser deconstructs text into concepts using a lexicon based on sequences of lexemes that represent multiple-word concepts extracted from AffectNet and IsaCore. These n-grams are not used blindly as fixed word patterns but exploited as reference for the module, in order to extract multiple-word concepts from information-rich sentences. So, differently from other shallow parsers, the module can recognise complex concepts also when irregular verbs are used or when these are interspersed with adjective and adverbs, e.g., the concept ‘buy christmas present’ in the sentence ‘I bought a lot of very nice Christmas presents’.

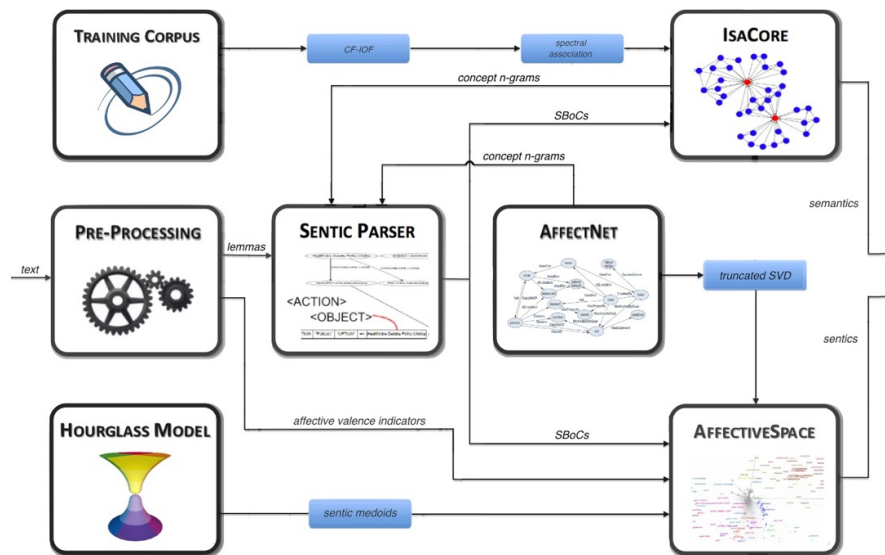


Fig. 4.10 Opinion mining engine block diagram. After performing a first skim of the input text, the engine extracts concepts from it and, hence, infers related semantics and sentic.

The sentic parser, additionally, provides, for each retrieved concept, its relative frequency, valence, and status, i.e., the concept's occurrence in the text, its positive or negative connotation, and the degree of intensity with which the concept is expressed, respectively. For each clause, the module outputs a small bag of concepts (SBoC), which is later on analysed separately by the IsaCore and AffectiveSpace modules to infer the cognitive and affective information associated with the input text, respectively. In case any of the detected concepts is found more than once in the vector space (that is, any of the concepts has multiple senses), all the SBoC concepts are exploited for a context-dependent coarse sense disambiguation. In particular, to represent the expected semantic value of the clause as a whole, the vectors corresponding to all concepts in the clause (in their ambiguous form) can be averaged together. The resulting vector does not represent a single meaning but the 'ad hoc category' of meanings that are similar to the various possible meanings of concepts in the clause [133]. Then, to assign the correct sense to the ambiguous concept, the concept sense with the highest dot product (and thus the strongest similarity) with the clause vector is searched.

Once natural language text is deconstructed into concepts, these are given as input to both the IsaCore and the AffectiveSpace modules. While the former exploits the graph representation of the common and common-sense knowledge base to detect semantics, the latter exploits the vector space representation of AffectNet to infer sentics. In particular, the IsaCore module applies spectral association for assigning activation to key nodes of the semantic network, which are used as seeds or centroids for classification. Such seeds can simply be the concepts corresponding to the class labels of interest plus their available synonyms and antonyms, if any. As shown in section 3.5.1, seeds can also be found by applying CF-IOF on a training corpus (when available), in order to perform a classification that is more relevant to the data under analysis. After seed concepts are identified, the module spreads their values across the IsaCore graph. This operation, an approximation of many steps of spreading activation, transfers the most activation to concepts that are connected to the seed concepts by short paths or many different paths in affective common-sense knowledge. Therefore, the concepts of each SBoC provided by the sentic parser are projected on the matrix resulting from spectral association in order to calculate their semantic relatedness to each seed concept and, hence, their degree of belonging to each different class. Such classification measure is directly proportional to the degree of connectivity between the nodes representing the retrieved concepts and the seed concepts in the IsaCore graph.

The concepts retrieved by the sentic parser are also given as input to the AffectiveSpace module, which, in turn, exploits dimensionality reduction techniques to infer the affective information associated with them. To this end, the concepts of each SBoC are projected into AffectiveSpace and, according to their position in the vector space representation of affective common-sense knowledge, they are assigned to an affective class defined through the Sentic Medoids technique. As well as in the IsaCore module, the categorisation does not consist in simply labelling each concept, but also in assigning a confidence score to each emotional label, which is directly proportional to the degree of belonging to a specific affective cluster (dot

product between the given concept and the relative Sentic Medoid). As shown in section 3.2.2, such affective information can also be exploited to calculate a polarity value associated with each SBoC provided by the sentic parser, as well as to detect the overall polarity associated with the opinionated text.

4.5.2 Evaluation

In order for the engine to perform fast real-time opinion mining, the sentic vector associated with each AffectNet concept can be calculated a priori and saved in an SQL database. At run-time, then, the sentic vectors relative to each of the concepts composing the SBoC can be retrieved from such a database and aggregated to compute the overall affective information associated with the specific SBoC. This allows the sentics extraction process to be faster than directly applying the AffectiveSpace process. Similarly, spectral association can be computed a priori on IsaCore and the semantic classification of each concept (that is, a set of different topic labels and the confidence associated with these) can be stored in an SQL database.

On average, in fact, while the processing of a 100-word opinion is on the order of tens of seconds when directly applying the AffectiveSpace and IsaCore processes, the extraction of semantics and sentics is on the order of seconds when using the corresponding SQL databases. Both resources are on the order of hundreds of megabytes and, hence, easily exportable and embeddable into bigger systems for the development of applications in fields such as Social Web, HCI, and e-health (as shown in the next chapter). As an example of how the software engine works, intermediate and final outputs obtained when a natural language opinion is given as input to the system are hereby examined. In particular, the following tweet was selected: "I think iPhone4 is the top of the heap! OK, the speaker is not the best i hv ever seen bt touchscreen really puts me on cloud 9... camera looks pretty good too!". After the pre-processing and semantic parsing operations, the following SBoCs are obtained:

SBoC#1:

<Concept: 'think'>
 <Concept: 'iphone4'>
 <Concept: 'top heap'>

SBoC#2:

<Concept: 'ok'>
 <Concept: 'speaker'>
 <Concept: '!good'++>
 <Concept: 'see'>

SBoC#3:

<Concept: 'touchscreen'>
 <Concept: 'put cloud nine'++>

SBoC#4:

<Concept: 'camera'>
 <Concept: 'look good'-->

Opinion Target	Category	Moods	Polarity
'iphone4'	'phones', 'electronics'	'ecstasy', 'interest'	+0.71
'speaker'	'electronics', 'music'	'annoyance'	-0.34
'touchscreen'	'electronics'	'ecstasy', 'anticipation'	+0.82
'camera'	'photography', 'electronics'	'acceptance'	+0.56

Table 4.4 Structured output example of opinion mining engine. For each clause, the engine detects the opinion target, the category it belongs to, and the affective information associated with it.

These are then concurrently processed by the IsaCore and the AffectiveSpace modules, which output the cognitive and affective information associated with each SBoC, both in a discrete way, with one or more labels, and in a dimensional way, with a polarity value $\in [-1,+1]$ (Table 4.4).

The developed engine overcomes the main problems faced by state-of-the-art keyword-based approaches. Thanks to the exploitation of common-sense knowledge bases that allow a concept-level analysis of text, in fact, the opinion mining engine is able to more effectively infer the cognitive and affective information associated with natural language. The validity of the proposed approach, hence, highly depends on two main factors: the richness of the knowledge bases and the accuracy of the sentic parser. If a concept for which there is no match in the knowledge bases is encountered in the opinionated text, the engine will not be able to infer the semantics and sentics associated with such concept and, hence, it might not be able to correctly perform the overall feature-based sentiment analysis of the input text.

For the same reason, the accuracy of the engine process also depends on the concept detection precision of the sentic parser. As shown in the previous section, the sentic parser can recognise complex concepts, e.g., the concept 'buy christmas present', in information-rich sentences, e.g., "I bought a lot of very nice Christmas presents". However, the parser would not be able to extract the same concept from a sentence like "I bought some nice Christmas bells for my mom". Since the current version of the parser is based on n-grams, in fact, it is not able to guess that a bell, in this case, is a Christmas present and, hence, does not detect the concept 'buy christmas present'. Since, moreover, the concept 'buy christmas bell' is not contained in AffectNet, the engine will not be able to infer the semantics and sentics associated with that concept, despite the affective common-sense knowledge base actually contains information about it. In this case, however, the sentic parser would be able to extract concepts such as 'buy', 'christmas', and 'bell' and, hence, make a good-enough guess about the semantics and sentics associated to the input text.

On average, in fact, SBoCs mostly consist of single-word or two-word concepts, hence, the validity of the proposed approach mainly depends on the richness of AffectNet and IsaCore. To this end, the evaluation process was performed at knowledge base level. An extensive corroborative evaluation of the opinion mining engine, in fact, was not possible because currently available test datasets in the field of sentiment analysis are still very few and there are no universally recognised benchmarks that can serve as a basis for comparing different techniques or different configurations of the same method.

A comparative evaluation of the engine with state-of-the-art techniques, moreover, was not feasible as most of current approaches are keyword-based, rather than concept-based, and aim for a document- or paragraph-level, rather than sentence- or clause-level, sentiment analysis. Hence, in order to evaluate the different facets of the opinion mining engine from various perspectives, three different resources were used, namely a Twitter hashtag repository, a LiveJournal database, and a PatientOpinion dataset, and results were compared with those obtained by using WordNet, ConceptNet, and Probase, in place of the two developed knowledge bases. Each of the evaluation resources has been selected to test a specific capability of the engine, according to the different properties of each test dataset. Specifically, the Twitter hashtag repository has been chosen for testing the precision of the engine in inferring semantics, the LiveJournal database has been selected to assess the engine's accuracy in extracting semantics, while the PatientOpinion dataset has been used for evaluating the inference of both semantics and semantics concurrently.

The Twitter hashtag repository, in particular, is a subset of the dataset used in Song et al. [263] for short text conceptualisation and consists of 3,000 tweets crawled from Bing web repository by exploiting Twitter hashtags as category labels. Such hashtags pertain to electronics (e.g., iPhone, Xbox, Android, and Wii), companies (e.g., Apple, Microsoft, and Google), countries, cities, operative systems, and cars. In order to test the resource's consistency and reliability, a manual evaluation of 100 tweets was performed, which showed that hashtags are accurate to 91%. The LiveJournal database, in turn, is a variation of the dataset used by Strapparava and Mihalcea [274] for emotion annotation and consists in a collection of 5,000 blogpost database extracted from LiveJournal, a virtual community of more than 23 millions users who keep a blog, journal, or diary. An interesting feature of this website is that bloggers are allowed to label their posts with both a category and a mood tag, by choosing from predefined categories and mood themes or by creating new ones. Since the indication of mood tags is optional, posts are likely to reflect the true mood of the authors, which is not always true for category tags.

After a manual evaluation of 200 posts, in fact, the category tags turned out to be very noisy (53% accuracy). The mood tags, however, showed a good enough reliability (89% accuracy) so they were used to test the engine's affect recognition performance. In order to have full correspondence between LiveJournal mood labels and the activation levels of the Hourglass model, moreover, a pool of 21 native English-speaking students was asked to map each of the 130 mood labels into the 24 emotional labels of the Hourglass model. In order to avoid any bias, students were randomly selected among different faculties in the university of Stirling and they were asked to perform the mapping by playing an ad hoc version of the Hourglass game in a secluded environment. Eventually, a reduced set of 80 moods (those with higher confidence level) was selected for inclusion in the blogpost database. The PatientOpinion dataset, finally, is a manually tagged evaluation resource kindly provided by James Munro, CEO of PatientOpinion, a social enterprise pioneering an online feedback service for users of the UK National Health Service to enable people to share their recent experience of local health services online. It is a collection of 2,000 patient opinions that associates to each post a category (namely,

clinical service, communication, food, parking, staff, and timeliness) and a positive or negative polarity. From a first manual evaluation of 50 opinions, the dataset appeared to be 100% accurate. It was used to test the detection of opinion targets and the polarity associated with these. Each of the used datasets has different strengths and weaknesses. The Twitter hashtag repository, for example, is written in very formal English and offers a great span of common knowledge and instances of concepts, but it is rather domain dependent and almost totally lacks the presence of common-sense concepts. The LiveJournal database, in turn, is more open-domain and affective common-sense oriented, but it often contains slang terms and grammatical mistakes, which lower the accuracy of clause chunking and concept parsing. The PatientOpinion database is very well-formed and the most accurate of the three datasets in terms of both affect and category labels, but it is very domain specific and rather limited in number of entries.

What led to the development of sentic computing is the need for better accuracy of sentiment analysis systems when switching between different domains. Currently available keyword-based approaches may perform nicely on a specific dataset, but they have very low accuracy if the domain changes. AffectNet and IsaCore, in turn, allow sentic computing to perform an open-domain opinion mining, in which text analysis is not simply based on word co-occurrence frequencies, but rather on the latent/implicit features associated with concepts. In order to assess the accuracy of such knowledge bases, hence, a comparison study was carried out by replacing AffectNet and IsaCore with state-of-the-art knowledge bases in the opinion mining engine. In particular, WordNet, ConceptNet, and Probase were firstly swapped with IsaCore to compare topic spotting performance of the engine on the Twitter hashtag repository. Secondly, the same knowledge bases were swapped with AffectNet to assess emotion recognition capabilities of the system on the LiveJournal database. Thirdly, WordNet, ConceptNet, and Probase were compared with the ensemble of IsaCore and AffectNet to concurrently evaluate the engine's topic spotting and polarity detection capabilities on the PatientOpinion dataset. As for the Twitter evaluation, results show that Probase and IsaCore perform significantly better than WordNet and ConceptNet, as these lack factual knowledge concepts such as Wii or Ford Focus (Table 4.5). Probase and IsaCore topic spotting precision, on the other hand, are comparable as Probase hyponym-hypernym common knowledge is enough for this kind of task. It actually even outperforms IsaCore sometimes as this contains just a subset of Probase instances (hub instances) and common-sense knowledge does not play a key role in this type of classification.

As for the LiveJournal evaluation, the capability of the software engine to properly categorise antithetical affective pairs from the Hourglass model (namely joy-sadness, anticipation-surprise, anger-fear, and trust-disgust) was tested. Results show that, in this case, Probase is consistently outperformed by WordNet, ConceptNet, and AffectNet, as it is based on semantic rather than affective relatedness of concepts (F-measure values are reported in Table 4.6). In Probase graph representation, in fact, instances like 'joy', 'surprise', and 'anger' are all close to each other, although they convey different affective valence, for being associated with the same hyponym-hypernym relationships.

Category	WordNet	ConceptNet	Probase	IsaCore
electronics	34.2%	45.4%	76.8%	79.3%
companies	26.5%	51.2%	82.3%	82.1%
countries	38.4%	65.1%	89.1%	86.4%
cities	25.9%	59.5%	81.5%	81.8%
operative systems	37.1%	51.6%	79.7%	78.3%
cars	13.3%	22.9%	74.8%	77.1%

Table 4.5 Precision values relative to Twitter evaluation. Probase and IsaCore perform significantly better than WordNet and ConceptNet, as these usually lack factual knowledge concepts.

Category	WordNet	ConceptNet	Probase	AffectNet
joy-sadness	47.1%	55.7%	33.9%	79.4%
anticipation-surprise	30.4%	41.1%	19.7%	68.2%
anger-fear	43.8%	49.4%	25.7%	74.6%
trust-disgust	27.2%	39.3%	12.8%	69.9%

Table 4.6 F-measures relative to LiveJournal evaluation. Probase does not perform as well as the other knowledge bases as it is based on semantic rather than affective relatedness of concepts.

Category	WordNet	ConceptNet	Probase	IsaCore+AffectNet
clinical service	35.7%	49.3%	56.6%	82.4%
communication	41.3%	50.6%	43.7%	79.1%
food	39.4%	45.7%	40.8%	81.7%
parking	47.8%	51.4%	49.1%	77.2%
staff	32.5%	37.5%	51.2%	73.8%
timeliness	44.1%	50.2%	41.9%	80.6%

Table 4.7 F-measure values relative to PatientOpinion evaluation. The ensemble application of IsaCore and AffectNet represents the best trade-off between topic spotting and affect recognition.

As for the PatientOpinion evaluation, eventually, the ensemble application of IsaCore and AffectNet turns out to be the best choice as it represents the best trade-off between common knowledge and affective common-sense knowledge, which is particularly needed when aiming to infer both the cognitive and affective information associated with text (F-measure values are reported in Table 4.7). As also shown by previous experiments, in fact, common knowledge is particularly functional for tasks such as open-domain text auto-categorisation while affective common-sense knowledge is notably useful for inferring the implicit emotional meaning underpinning words. More evidence of this is given in the next chapter, in which the opinion mining engine (or a sub-part of it) is employed for the realisation of different tasks and applications in fields such as Social Web, HCI, and e-health.

4.6 Conclusions

This chapter showed how sentic computing techniques are employed for the development of NLP tools such as a semantic resource for opinion mining (section 4.1) and a brain-inspired model for affective reasoning (section 4.2). Moreover, the chapter discussed how to enrich AffectNet through LSR, crowd sourcing, and GWAP techniques (section 4.3) and how to build possibly the most comprehensive resource of common and common-sense knowledge for sentiment analysis, from ConceptNet and Probase (section 4.4).

Finally, the chapter described how the ensemble application of the developed tools and techniques can be exploited for the design of an open-domain concept-level opinion mining engine (section 4.5), of which an output example and an evaluation are provided. In order to further assess how such an engine can be effectively exploited for tackling real-world problems, next chapter will explore multiple systems in which the engine can be embedded for enhancing their NLP capabilities.

Chapter 5

Applications

*Knowing is not enough; we must apply.
Willing is not enough; we must do.*

Johann Wolfgang von Goethe

The amount of data available on the Web is growing exponentially. These data, however, are mainly in an unstructured format and, hence, not machine-processable and machine-interpretable. What is called collective intelligence today is actually just collected intelligence as the value of user contributions is simply in their being collected together and aggregated into community or domain specific sites. True collective intelligence can emerge if the data collected from all those people is aggregated and recombined to create new knowledge and new ways of learning that individual humans cannot do by themselves [121].

So far, online information retrieval has mainly relied on keyword-based algorithms, which have proved to have important limitations, e.g., the inability to recognise topical authority that humans recognise effortlessly without the explicit words being in the content. In order to let machines better understand natural language and, hence, conveniently analyse and aggregate opinions and sentiments over the Web, we need to provide them with both adequately broad common-sense knowledge bases and reasoning methods to efficiently handle these. This chapter describes how sentic computing tools and techniques are employed for the development of applications in fields such as Social Web (section 5.1), HCI (section 5.2), and e-health (section 5.3). The chapter, eventually, ends with some concluding remarks (section 5.4).

5.1 Development of Social Web Systems

With the rise of the Social Web, there are now millions of humans offering their knowledge online, which means that the information is stored, searchable, and easily shared. This trend has created and maintained an ecosystem of participation, where value is created by the aggregation of many individual user contributions. Such contributions, however, are meant for human consumption and, hence, hardly accessible and processable by computers. Making sense of the huge amount of social data available on the Web requires the adoption of novel approaches to natural language understanding that can give a structure to such data, in a way that they can be more easily aggregated and analysed.

In this context, sentic computing can be exploited for NLP tasks requiring the inference of semantic and/or affective information associated with text, from big social data analysis [57] to management of online community data and metadata [119] to analysis of social network interaction dynamics [68]. This section, in particular, shows how the engine can be exploited for the development of a troll filtering system (section 5.1.1), a social media marketing tool (section 5.1.2), and an online personal photo management system (section 5.1.3).

5.1.1 Troll Filtering

The democracy of the Web is what made it so popular in the past decades, but such a high degree of freedom of expression also gave birth to negative side effects – the so called ‘dark side’ of the Web. Be it real or virtual world, in fact, existence of malicious faction among inhabitants and users is inevitable. An example of this, in the Social Web context, is the exploitation of anonymity to post inflammatory, extraneous, or off-topic messages in an online community, with the primary intent of provoking other users into a desired emotional response or of otherwise disrupting normal on-topic discussion.

Such a practice is usually referred as ‘trolling’ and the generator of such messages is called ‘a troll’. The term was first used in early 1990 and since then a lot of concern has been raised to contain or curb trolls. The trend of trolling appears to have spread a lot recently and it is alarming most of the biggest social networking sites since, in extreme cases such as abuse, has led some teenagers to commit suicide. These attacks usually address not only individuals, but also entire communities. For example, reports have claimed that a growing number of Facebook tribute pages had been targeted, including those in memory of the Cumbria shootings victims and soldiers who died in Afghanistan.

At present, users cannot do much rather than manually delete abusive messages. Current anti-trolling methods, in fact, mainly consist in identifying additional accounts that use the same IP address and blocking fake accounts based on name and anomalous site activity, e.g., users who send lots of messages to non-friends or whose friend requests are rejected at a high rate.

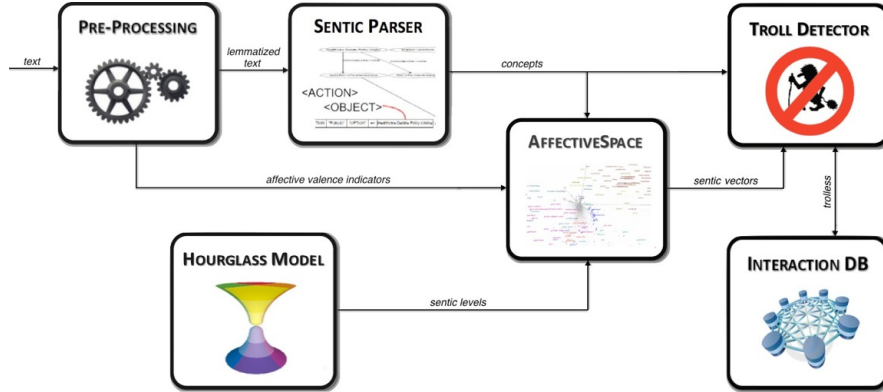


Fig. 5.1 Troll filtering process. Once extracted, semantics and sentics are used to calculate blog-posts' level of trollness, which is then stored in the interaction database for the detection of malicious behaviours.

In July 2010, Facebook launched an application that gives users a direct link to advice, help, and the ability to report cyber problems to the child exploitation and online protection centre (CEOP). Reporting trouble through a link or a button, however, is a too slow process since social networking websites usually cannot react instantly to these alarms. A button, moreover, does not stop users from being emotionally hurt by trolls and it is more likely to be pushed by people who actually do not need help rather than, for instance, children who are being sexually groomed and do not realise it.

A prior analysis of the trustworthiness of statements published on the Web has been presented by Rowe and Butters [245]. Their approach adopts a contextual trust value determined for the person who asserted a statement as the trustworthiness of the statement itself. Their study, however, does not focus on the problem of trolling, but rather on defining a contextual accountability for the detection of web, email, and opinion spam. The main aim of the troll filter [38] (Fig. 5.1) is to identify malicious contents in natural language text with a certain confidence level and, hence, automatically block trolls. To train the system, the concepts most commonly used by trolls are first identified by using the CF-IOF technique and, then, this set is expanded through spectral association. In particular, after analysing a set of 1000 offensive phrases extracted from Wordnik¹, it was found that, statistically, a post is likely to be edited by a troll when its average sentic vector has a high absolute value of Sensitivity and a very low polarity. Hence, the *trollness* t_i associated to a concept c_i is defined as a float $\in [0, 1]$ such that:

$$t_i(c_i) = \frac{s_i(c_i) + |Snsit(c_i)| - p_i(c_i)}{3} \quad (5.1)$$

¹ <http://wordnik.com>

where s_i (float $\in [0, 1]$) is the semantic similarity of c_i wrt any of the CF-IOF seed concepts, p_i (float $\in [-1, 1]$) is the polarity associated to the concept c_i , and 3 is the normalisation factor. Hence, the total *trollness* of a post containing N concepts is defined as:

$$t = \sum_{i=1}^N \frac{3 s_i(c_i) + 4 |Snsit(c_i)| - Plsnt(c_i) - |Attnt(c_i)| - Aptit(c_i)}{9N} \quad (5.2)$$

This information is stored, together with post type and content plus sender and receiver ID, in an interaction database that keeps trace of all the messages and comments interchanged between users within the same social network. Posts with a high level of *trollness* (current threshold has been set, using a trial-and-error approach, to 60%) are labelled as troll posts and, whenever a specific user addresses more than two troll posts to the same person or community, his/her sender ID is labelled as troll for that particular receiver ID. All the past troll posts sent to that particular receiver ID by that specific sender ID are then automatically deleted from the website (but kept in the database with the possibility for the receiver to either visualise them in an apposite *troll folder* and, in case, restore them). Moreover, any new post with a high level of *trollness* edited by a user labelled as troll for that specific receiver is automatically blocked, i.e., saved in the interaction database but never displayed in the social networking website.

This information, encoded as a sentic vector, is given as input to a troll detector which exploits it, together with the semantic information coming directly from the sentic parser, to calculate the post's *trollness* and, eventually, to detect and block the troll (according to the information stored in the interaction database). As an example of troll filtering process output, a troll post recently addressed to the Indian author, Chetan Bhagat, can be considered: "You can't write, you illiterate douchebag, so quit trying, I say!!!". In this case, there are a very high level of Sensitivity (corresponding sentic level 'rage') and a negative polarity, which give a high percentage of *trollness*, as shown below:

```
<Concept: !'write'>
<Concept: 'illiterate'>
<Concept: 'douchebag'>
<Concept: 'quit try'>
<Concept: 'say'>
Semantics: 0.69
Sentic: [0.0, 0.17, 0.85, -0.43]
Polarity: -0.38
Trollness: 0.75
```

Because the approach adopted by Rowe and Butters [245] is not directly comparable with the developed troll filtering system, a first evaluation was performed by considering a set of 500 tweets manually annotated as troll and non-troll posts, most of which were fetched from Wordnik. In particular, true positives were identified as posts with both a positive troll-flag and a *trollness* $\in [0.6, 1]$, or posts with both a negative troll-flag and a *trollness* $\in [0, 0.6]$.

Metric	IsaCore	AnalogySpace	AffectiveSpace
precision	57.1%	69.1%	82.5%
recall	40.0%	56.6%	75.1%
F-measure	47.0%	62.2%	78.6%

Table 5.1 Precision, recall, and F-measure values relative to the troll filter evaluation. The AffectiveSpace process performs consistently better than IsaCore and AnalogySpace in detecting troll posts.

The threshold has been set to 60% based on trial-and-error over a separate dataset of 50 tweets. Results show that, by using the troll filtering process, inflammatory and outrageous messages can be identified with good precision (82.5%) and decorous recall rate (75.1%). In particular, the F-measure value (78.9%) is significantly high compared to the corresponding F-measure rates obtained by using IsaCore and AnalogySpace in place of the AffectiveSpace process (Table 5.1).

However, much better results are expected for the process evaluation at interaction-level, rather than just at post-level. In the future, in fact, the troll filtering process will be evaluated by monitoring not just single posts, but also users' holistic behaviour, i.e., contents and recipients of their interaction, within the same social network.

5.1.2 Social Media Marketing

The advent of Web 2.0 made users more enthusiastic about interacting, sharing, and collaborating through social networks, online communities, blogs, wikis, and other online collaborative media. In the last years, this collective intelligence has spread to many different areas in the Web, with particular focus on fields related to our everyday life such as commerce, tourism, education, and health. The online review of commercial services and products, in particular, is an action that users usually perform with pleasure, to share their opinions about services they have received or products they have just bought, and it constitutes immeasurable value for other potential buyers.

This trend opened new doors to enterprises that want to reinforce their brand and product presence in the market by investing in online advertising and positioning. In confirmation of the growing interest in social media marketing, several commercial tools have been recently developed to provide companies with a way to analyse the blogosphere on a large scale in order to extract information about the trend of the opinions relative to their products. Nevertheless most of the existing tools and the research efforts are limited to a polarity evaluation or a mood classification according to a very limited set of emotions. In addition, such methods mainly rely on parts of text in which emotional states are explicitly expressed and, hence, they are unable to capture opinions and sentiments that are expressed implicitly. To this end, a novel social media marketing tool has been proposed [39] to provide marketers with an IUI for the management of social media information at semantic level, able

to capture both opinion polarity and affective information associated with UGCs. A polarity value associated to an opinion, in fact, sometimes can be restrictive. Enriching automatic analysis of social media with affective labels such as ‘joy’ or ‘disgust’ can help marketers to have a clearer idea of what their customers think about their products. In particular, YouTube was selected as a social media source since, with its over 2 billions views per day, 24 hours of video uploaded every minute, and 15 minutes a day spent by the average user, it represents more than 40% of online video market². Specifically, the focus was on video reviews of mobile phones because of the quantity and the quality of the comments usually associated with them.

The social media analysis is performed through three main steps: firstly, comments are analysed using the opinion mining engine; secondly, the extracted information is encoded on the base of different web ontologies; finally, the resulting knowledge base is made available for browsing through a multi-faceted classification website. Social Web resources represent a peculiar kind of data that is characterised for a deeply interconnected nature. The Web itself is, in fact, based on links that bound together different data and information, and community-contributed multimedia resources characterise themselves for the collaborative way in which they are created and maintained. An effective description of such resources needs therefore to capture and manage such interconnected nature, allowing to encode information not only about the resource itself, but also about the linked resources into an interconnected knowledge base. Encoding information relative to a market product to analyse its market trends represents a situation in which this approach is particularly suitable and useful. In this case, in fact, it is necessary not only to encode the information relative to product features, but also the information about the producer, the consumers, and their opinions.

The proposed framework for opinion description and management aims to be applicable to most of online resources (videos, images, text) coming from different sources, e.g., online video sharing services, blogs, and social networks. To such purpose, it is necessary to standardise as much as possible the descriptors used in encoding the information about multimedia resources and people to which the opinions refer (considering that every website uses its own vocabulary), in order to make it univocally interpretable and suitable to feed other applications. For this reason, the information relative to multimedia resources and people is encoded using the descriptors provided by OMR³ (Ontology for Media Resources) and FOAF⁴ (Friend of a Friend Ontology), respectively. OMR represents an important effort to help circumventing the current proliferation of audio/video metadata formats, currently carried on by the W3C Media Annotations Working Group. It offers a core vocabulary to describe media resources on the Web, introducing descriptors such as ‘title’, ‘creator’, ‘publisher’, ‘createDate’, and ‘rating’, and it defines semantic-preserving mappings between elements from existing formats in order to foster the interoperability among them.

² <http://viralblog.com/research/youtube-statistics>

³ <http://w3.org/TR/mediaont-10>

⁴ <http://www.foaf-project.org>

FOAF represents a recognised standard in describing people, providing information such as their names, birthdays, pictures, blogs, and especially other people they know, which makes it particularly suitable for representing data that appear in social networks and communities. OMR and FOAF together supply most of the vocabulary needed for describing media and people; other descriptors are added only when necessary. For example, OMR does not currently supply vocabulary for describing comments, which are hereby analysed to extract the affective information relative to media. Hence, the ontology is extended by introducing the ‘Comment’ class and by defining for it the ‘author’, ‘text’, and ‘publicationDate’ properties.

In HEO, properties to link emotions to multimedia resources and people were introduced. In particular, ‘hasManifestationInMedia’ and ‘isGeneratedByMedia’ were defined to describe emotions that occur and are generated in media, respectively, while the property ‘affectPerson’ was defined to connect emotions to people. Additionally, WNA was exploited as an ontology in order to improve the hierarchical organisation of emotions in HEO. Thus, the combination of HEO with WNA, OMR, and FOAF provides a complete framework to describe not only multimedia contents and the users that have created, uploaded, or interacted with them, but also the opinions and the affective content carried by the media and the way they are perceived by web users (Fig. 5.2).

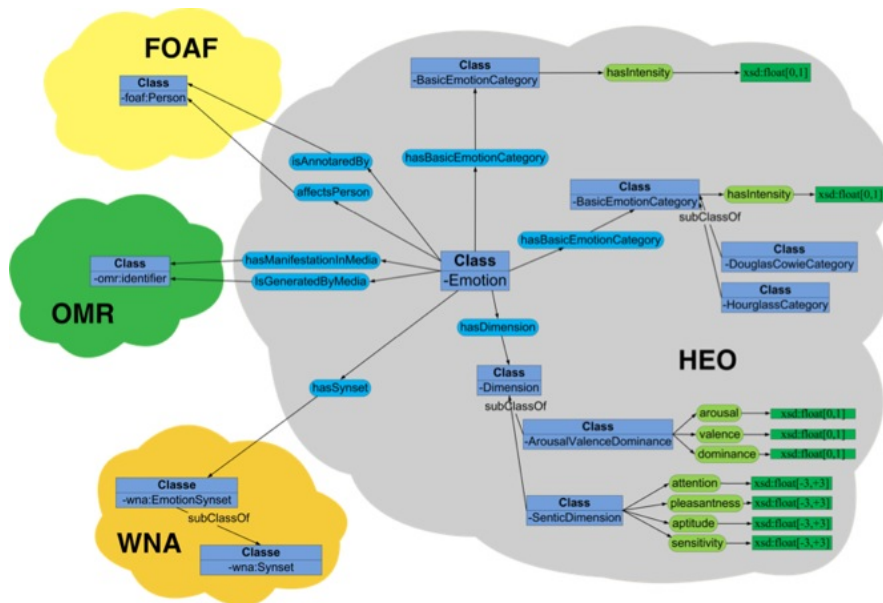


Fig. 5.2 Merging different ontologies. The combination of HEO, WNA, OMR and FOAF provides a comprehensive framework for the representation of social media affective information.

As remarked above, due to the way they are created and maintained, community-contributed multimedia resources are very different from standard web data. One fundamental aspect is the collaborative way in which such data is created, uploaded, and annotated. A deep interconnection emerges in the nature of these data and meta-data, allowing for example to associate videos of completely different genre, but uploaded by the same user, or different users, even living in opposite sides of the world, who have appreciated the same pictures. In the context of social media marketing, this interdependence can be exploited to find similar patterns in customer reviews of commercial products and, hence, to gather useful information for marketing, sales, public relations, and customer service. Online reviews of electronics products, in particular, usually offer substantial and reliable information about the perceived quality of the products because of the size of electronics online market and the type of customers related to it.

To visualise this information, the multi-faceted categorisation paradigm is exploited. Faceted classification allows the assignment of multiple categories to an object, enabling the classifications to be ordered in multiple ways, rather than in a single, pre-determined, and taxonomic order. This makes possible to perform searches combining the textual approach with the navigational one. Faceted search, in fact, enables users to navigate a multi-dimensional information space by concurrently writing queries in a text box and progressively narrowing choices in each dimension. For this application, specifically, the SIMILE Exhibit API⁵ is used. Exhibit consists of a set of Javascript files that allow to easily create rich interactive web pages including maps, timelines, and galleries, with very detailed client-side filtering. Exhibit pages use the multi-faceted classification paradigm to display semantically structured data stored in a Semantic Web aware format, e.g., RDF or JavaScript object notation (JSON). One of the most relevant aspects of Exhibit is that, once the page is loaded, the web browser also loads the entire data set in a lightweight database and performs all the computations (sorting, filtering, etc.) locally on the client-side, providing high performance.

Because they are one of the most prolific types of electronic products in terms of data reviews available on the Web, mobile phones were selected as a review target. In particular, a set of 220 models was considered. Such models were ranked as the most popular according to Kelkoo⁶, a shopping site featuring online shopping guides and user reviews, from which all the available information about each handset, such as model, brand, input type, screen resolution, camera type, standby time, and weight, was parsed. This information was encoded in RDF and stored in a Sesame⁷ triple-store, a purpose-built database for the storage and retrieval of RDF metadata. YouTube Data API was then exploited to retrieve from YouTube database the most relevant video reviews for each mobile phone and their relative metadata such as duration, rating, upload date and name, gender, and country of the uploaders.

⁵ <http://simile-widgets.org/exhibit>

⁶ <http://kelkoo.co.uk>

⁷ <http://openrdf.org>



Fig. 5.3 A screenshot of the social media marketing tool. The faceted classification interface allows the user to navigate through both the explicit and implicit features of the different products.

The comments associated with each video were also extracted and processed by means of sentic computing for emotion recognition and polarity detection. The extracted opinions in RDF/XML were then encoded using the descriptors defined by HEO, WNA, OMR, and FOAF, and inserted into the triple-store. Sesame can be embedded in applications and used to conduct a wide range of inferences on the information stored, based on RDFS and OWL type relations between data. In addition, it can also be used in a standalone server mode, much like a traditional database with multiple applications connecting to it. In this way, the knowledge stored inside Sesame can be easily queried; optionally, results can also be retrieved in a semantic aware format and used for other applications. For the developed demo, the information contained in the triple-store was exported into a JSON file, in order to make it available for being browsed as a unique knowledge base through Exhibit interface.

In the UI, mobile phones are displayed through a dynamic gallery that can be ordered according to different parameters, e.g., model, price, and rating, showing technical information jointly with their video reviews and the opinions extracted from the relative comments (Fig. 5.3). By using faceted menus, moreover, it is possible to explore such information both using the search box (to perform keyword-based queries) and filtering the results using the faceted menus (by adding or removing constraints on the facet properties). In this way, it becomes very easy and intuitive to search for mobile phones of interest: users can specify the technical features required using the faceted menus and compare different phones that match such requirements by consulting the video reviews and the opinions extracted from the relative comments.

In addition, it is possible to explore in detail the comments of each video review through a specific Exhibit page in which comments are organised in a timeline and highlighted in different colours, according to the value of their polarity. Moreover, faceted menus allow filtering the comments according to the reviewers' information, e.g., age, gender, and nationality. Using such a tool a marketer can easily get an insight about the trend of a product, e.g., at the end of an advertising campaign, by observing how the number of reviews and the relative satisfaction evolve in time and by monitoring this trend for different campaign targets.

In order to evaluate the proposed system both on the level of opinion mining and sentiment analysis, we separately tested its polarity detection accuracy with a set of like/dislike-rated video reviews from YouTube and evaluated its affect recognition capabilities with a corpus of mood-tagged blogs from LiveJournal. In order to evaluate the system in terms of polarity detection accuracy, we exploited YouTube Data API to retrieve from YouTube database the ratings relative to the 220 video reviews previously selected for displaying in the faceted classification interface. On YouTube, in fact, users can express their opinions about videos either by adding comments or by simply rating them using a like/dislike button. YouTube Data API makes this kind of information available by providing, for each video, number of raters and average rating, i.e., sum of likes and dislikes divided by number of raters. This information is expressed as a float $\in [1, 5]$ and indicates if a video is generally considered as bad (float $\in [1, 3)$) or good (float $\in [3, 5]$).

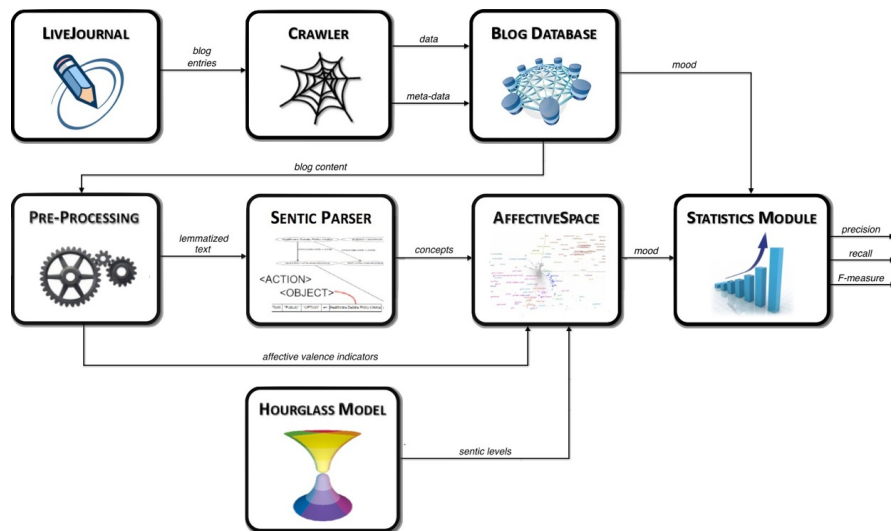


Fig. 5.4 Sentics extraction evaluation. The process extracts sentics from posts in the LiveJournal database, and then compare inferred emotional labels with the relative mood tags in the database.

This information was compared with the polarity values previously extracted by employing sentic computing on the comments relative to each of the 220 videos. True positives were identified as videos with both an average rating $\in [3, 5]$ and a polarity $\in [0, 1]$ (for positively rated videos), or videos with both an average rating $\in [1, 3]$ and a polarity $\in [-1, 0]$ (for negatively rated videos). The evaluation showed that, by using the system to perform polarity detection, negatively and positively rated videos (37.7% and 62.3% of the total respectively) can be identified with precision of 97.1% and recall of 86.3% (91.3% F-measure).

Since no mood-labelled dataset about commercial products is currently available, the LiveJournal database was used to test the system's affect recognition capabilities. For this test, a reduced set of 10 moods has been considered, specifically, 'ecstatic', 'happy', 'pensive', 'surprised', 'enraged', 'sad', 'angry', 'annoyed', 'scared', and 'bored'. All LiveJournal accounts have Atom, RSS, and other data feeds that show recent public entries, friend relationships, and interests. Unfortunately, it is not possible to get mood-tagged blogposts via data feeds, hence, an ad hoc crawler had to be designed. After retrieving and storing relevant data and meta-data for a total of 5,000 posts, the sentics extraction process was conducted on each of these and outputs were compared with the relative LiveJournal mood tags, in order to compute recall and precision rates as evaluation metrics (Fig. 5.4).

On average, each post contained around 140 words and, from it, about 4 affective valence indicators and 60 sentic vectors were extracted. According to this information, mood-labels were assigned to each post and compared with the corresponding LiveJournal mood tags, obtaining very good accuracy for each of the 10 selected moods (Table 5.2). Among these, 'happy' and 'sad' posts were identified with particularly high precision (89.2% and 81.8%, respectively) and decorous recall rates (76.5% and 68.4%). The F-measure values obtained, hence, were significantly good (82.3% and 74.5%, respectively), especially if compared to the corresponding F-measure rates of a standard keyword spotting system based on a set of 500 affect words (65.7% and 58.6%).

Mood	Precision	Recall	F-measure
ecstatic	73.1%	61.3%	66.6%
happy	89.2%	76.5%	82.3%
pensive	69.6%	52.9%	60.1%
surprised	81.2%	65.8%	72.6%
enraged	68.9%	51.6%	59.0%
sad	81.8%	68.4%	74.5%
angry	81.4%	53.3%	64.4%
annoyed	77.3%	58.7%	66.7%
scared	82.6%	63.5%	71.8%
bored	70.3%	55.1%	61.7%

Table 5.2 Evaluation results of the sentics extraction process. Precision, recall, and F-measure rates are calculated for ten different moods by comparing the engine output with LiveJournal mood tags.

5.1.3 *Sentic Album*

Efficient access to online personal pictures requires the ability to properly annotate, organise, and retrieve the information associated with them. While the technology to search personal documents has been available for some time, the technology to manage personal images is much more challenging. This is mainly due to the fact that, even if images can be roughly interpreted automatically, many salient features exist only in the user's mind. The only way for a system to accordingly index personal images, hence, is to try to capture and process such features.

Existing content based image retrieval (CBIR) systems such as QBIC [110], Virage [15], MARS [233], ImageGrouper [206], MediAssist [210], CIVR [253], EGO [288], ACQUINE [89], and K-DIME [23] have attempted to build IUIs capable of retrieving pictures according to their intrinsic content through statistics, pattern recognition, signal processing, computer vision, support vector machines, and neural networks, but all such techniques appeared too weak to bridge the gap between the data representation and the images' conceptual models in the user's mind. Image meta search engines such as Webseek [260], Webseer [112], PicASHOW [169], IGroup [145], or Google⁸, Yahoo⁹, and Bing¹⁰ Images, on the other hand, rely on tags associated with online pictures but, in the case of personal photo management, users are unlikely to expend substantial effort to manually classify and categorise images in the hopes of facilitating future retrieval. Moreover these techniques, as they depend on keyword-based algorithms, often miss potential connections between words expressed through different vocabularies or concepts that exhibit implicit semantic connectedness. In order to properly deal with photo metadata and, hence, effectively annotate images, in fact, it is necessary to work at a semantic, rather than syntactic, level.

A good effort in this sense has been made within the development of ARIA [175], a software agent which aims to facilitate the storytelling task by opportunistically suggesting photos that may be relevant to what the user is typing. ARIA goes beyond the naïve approach of suggesting photos by simply matching keywords in a photo annotation with keywords in the story, as it also takes into account semantically related concepts. A similar approach has been followed by Raconteur [72], a system for conversational storytelling that encourages people to make coherent points, by instantiating large-scale story patterns and suggesting illustrative media. It exploits a large common-sense knowledge base to perform NLP in real-time on a text chat between a storyteller and a viewer and recommends appropriate media items from a library. Both these approaches present a lot of advantages since concepts, unlike keywords, are not sensitive to morphological variation, abbreviations, or near synonyms. However, simply relying on a semantic knowledge base is not enough to infer the salient features that make different pictures more or less relevant in each user's mind.

⁸ <http://google.com/images>

⁹ <http://images.search.yahoo.com>

¹⁰ <http://bing.com/images>

To this end, Sentic Album [43] exploits AI and Semantic Web techniques to perform reasoning on different knowledge bases and, hence, infer both the cognitive and affective information associated with photo metadata. The system, moreover, supports this concept-level analysis with content and context based techniques, in order to capture all the different aspects of online pictures and, hence, provide users with an IUI that is navigable in real-time through a multi-faceted classification website. Much of what is called problem-solving intelligence, in fact, is really the ability to identify what is relevant and important in a context and to subsequently make that knowledge available just in time [176].

Cognitive and affective processes are tightly intertwined in everyday life [84]. The affective aspect of cognition and communication is recognised to be a crucial part of human intelligence and has been argued to be more fundamental in human behaviour for ensuring success in social life than intellect [292, 221]. Emotions, in fact, influence our ability to perform common cognitive tasks, such as forming memories and communicating with other people. A psychological study, for example, showed that people asked to conceal emotional facial expressions in response to unpleasant and pleasant slides remembered the slides less well than control participants [28]. Similarly, a study of conversations revealed that romantic partners who were instructed to conceal both facial and vocal cues of emotion while talking about important relationship conflicts with each other, remembered less of what was said than did partners who received no suppression instructions [241]. Many studies have indicated that emotions both seem to improve memory for the gist of an event and to undermine memory for more peripheral aspects of the event [32, 75, 298, 239].

The idea, broadly, is that arousal causes a decrease in the range of cues an organism can take in. This narrowing of attention leads directly to the exclusion of peripheral cues, and this is why emotionality undermines memory for information at the event's edge. At the same time, this narrowing allows a concentration of mental resources on more central materials, and this leads to the beneficial effects of emotion on memory for the event's centre [163]. Hence, rather than assigning particular cognitive and affective valence to a specific visual stimulus, we more often balance the importance of personal pictures is according to how much information contained in them is pertinent to our lives, goals, and values (or perhaps, the lives and values of people we care about). For this reason, a bad quality picture can be ranked high in the mind of a particular user, if it reminds him/her of a notably important moment or person of his/her life.

Events, in fact, are likely to be organised in the human mind as interconnected concepts and most of the links relating such concepts are probably weighted by affect, as we tend to better recall memories associated with either very positive or very negative emotions, as well as we usually tend to more easily forget about concepts associated with very little or null affective valence. The problem, when trying to emulate such cognitive and affective processes, is that, while cognitive information is usually objective and unbiased, affective information is rather subjective and argumentative. For example, while in the cognitive domain 'car' is always a car and there is usually not much discussion about the correctness of retrieving an image showing a tree in an African savannah under the label 'landscape', there might be

some discussion about whether the retrieved car is “cool” or just “nice”, or whether the found landscape is “peaceful” or “dull” [127]. To this end, Sentic Album applies sentic computing techniques on pictures data and metadata to infer what really matters to each user in different online photos. In particular, the Annotation Module mainly exploits metadata such as descriptions, tags, and comments, termed ‘conceptual metadata’, associated with each image to extract its relative semantics and sentics and, hence, enhance the picture specification with its intrinsic cognitive and affective information. This concept-level annotation procedure is performed through the opinion mining engine and it is supported with a parallel content and context level analysis.

The content based annotation, in particular, is performed through Python Imaging Library¹¹ (PIL), an external library for the Python¹² programming language that adds support for opening, manipulating, and saving many different image file formats. For every online personal picture, in particular, PIL is exploited to extract luminance and chrominance information and other image statistics, e.g., the total, mean, standard deviation, and variance of the pixel values.

The context-based annotation, in turn, exploits information such as timestamp, geolocation, and user interaction metadata. Such metadata, termed ‘contextual metadata’, are processed by the Context Deviser, a sub-module that extracts small bits of information suitable for storing in a relational database for re-use at a later time, i.e., time, date, city and country of caption, plus relevant user interaction metadata such as number and IDs of friends who viewed, commented, or liked the picture.

The concept-based annotation represents the core of the module as it allows Sentic Album to go beyond a mere syntactic analysis of the metadata associated with pictures. A big problem of manual image annotation, in fact, is the different vocabulary that different users (or even the same user) can use to describe the content of a picture. The different expertise and purposes of tagging users, in fact, may result in tags that use various levels of abstraction to describe a resource: a photo can be tagged at the ‘basic level’ of abstraction [161] as ‘cat’, or at a superordinate level as ‘animal’, or at various subordinate levels below the basic level as ‘Persian cat’ or ‘*Felis silvestris catus longhair Persian*’.

To overcome this problem, Sentic Album extends the set of available tags with related semantics and sentics and, to further expand the cognitive and affective metadata associated with each picture, it extracts additional common-sense and affective concepts from its description and comments. In particular, the conceptual metadata is processed by the opinion mining engine (Fig. 5.5). The IsaCore sub-module, specifically, finds matches between the retrieved concepts and those previously calculated using CF-IOF and spectral association. CF-IOF weighting is exploited to find seed concepts for a set of a-priori categories, extracted from Picasa¹³ popular tags, meant to cover common topics in personal pictures, e.g., art, nature, friends, travel, wedding, or holiday. Spectral association is then used to expand this set with

¹¹ <http://pythonware.com/products/pil>

¹² <http://python.org>

¹³ <http://picasa.google.com/>

semantically related common-sense concepts. The retrieved concepts are also processed by the AffectiveSpace sub-module, which projects them into the vector space representation of AffectNet, clustered by means of Sentic Medoids, in order to infer the affective valence and the polarity associated with them.

Providing a satisfactory visual experience is one of the main goals for present-day electronic multimedia devices. All the enabling technologies for storage, transmission, compression, and rendering should preserve, and possibly enhance, image quality; and to do so, quality control mechanisms are required. Systems to automatically assess visual quality are generally known as objective quality metrics. The design of objective quality metrics is a complex task because predictions must be consistent with human visual quality preferences.

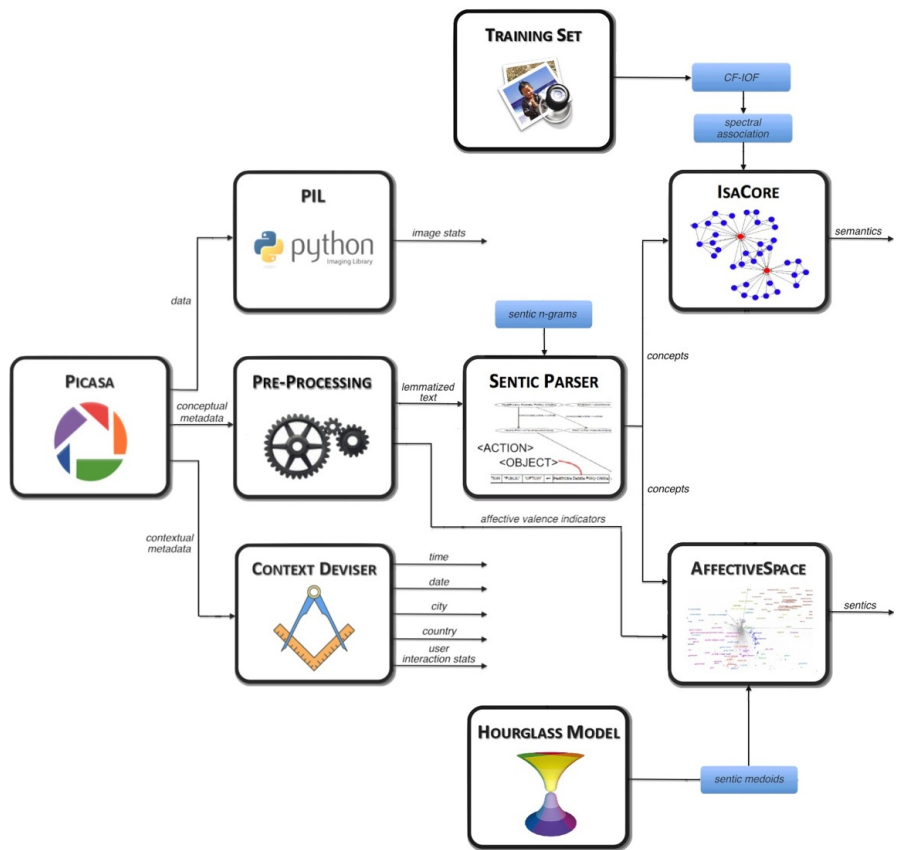


Fig. 5.5 Sentic Album’s annotation module. Online personal pictures are annotated at three different levels: content level (PIL), concept level (opinion mining engine) and context level (Context Deviser).

Human preferences are inherently quite variable and, by definition, subjective; moreover, in the field of visual quality, they stem from perceptual mechanisms that are not fully understood yet. A common choice is to design metrics that replicate the functioning of the human visual system (HVS) to a certain extent, or at least that take into account its perceptual response to visual distortions by means of numerical features [152]. Although successful, these approaches come with a considerable computational cost, which makes them impractical for most real-time applications. Computational intelligence paradigms allow to tackle the quality assessment task from a different perspective, since they aim at mimicking quality perception instead of designing an explicit model of the HVS [207, 182, 238].

In the special case of personal pictures, perceived quality metrics can be computed not only at content level, but also at concept and context level. One of the primary reasons why people take pictures is to remember the emotions they felt on special occasions of their lives. Extracting and storing such affective information can be a key factor in improving future searches, as users seldom want to find photos matching general requirements. Users' criteria in browsing personal pictures, in fact, are more often related to the presence of a particular person in the picture and/or its perceived quality (e.g., to find a good photo of your mother). Satisfying this type of requirement is a tedious task as chronological ordering or classification by event does not help much. The process usually involves repeatedly trying to think of a matching picture and then looking for it. An exhaustive search (looking through the whole collection for all of the photos matching a requirement) would normally only be carried out in exceptional circumstances, such as following a death in the family. In order to accordingly rank personal photos, Sentic Album exploits data and metadata associated with them to extract useful information at content, concept, and context level and, hence, calculate the perceived quality of online pictures (PQOP):

$$PQOP(p, u) = 3 \frac{Content(p) * Concept(p, u) * Context(p, u)}{Content(p) + Concept(p, u) + Context(p, u)} \quad (5.3)$$

where *Content*, *Concept*, and *Context* (3Cs) are float $\in [0,1]$ representing image quality assessment values associated with picture p and user u , in terms of visual, conceptual, and contextual information, respectively. In particular, *Content*(p) is computed from numerical features extracted through a reduced-reference framework for objective quality assessment based on extreme learning machine [91] and the colour correlogram [142] of p ; *Concept*(p, u) specifies how much the picture p is relevant to the user u in terms of cognitive and affective information; finally, *Context*(p, u) defines the degree of relevance of picture p for user u in terms of time, location, and user interaction. The 3Cs are all equally relevant for measuring how good a personal picture is to the eye of a user. According to the formula, in fact, if any of the 3Cs is null the PQOP is null as well, even though the remaining elements of the 3Cs have both maximum values, e.g., a perfect quality picture (*Content*(p) = 1) taken in the hometown of the user on the date of his birthday (*Context*(p, u) = 1), but depicting people he/she does not know and objects/places that are totally irrelevant for him/her (*Concept*(p, u) = 0).

The Storage Module is the middle-tier in which the outputs of the Annotation Module are stored, in a way that these can be easily accessible by the Search and Retrieval Module at a later time. The module stores information relative to photo data and metadata redundantly at three levels:

1. in a relational database fashion
2. in a Semantic Web format
3. in a matrix format

Sentic Album stores information in three main SQL databases (Fig. 5.6), that is a Content DB, for the information relative to data (image statistics), a Concept DB, for the information relative to conceptual metadata (semantics and sentsics), and a Context DB, for the information relative to contextual metadata (timestamp, geolocation, and user interaction metadata).

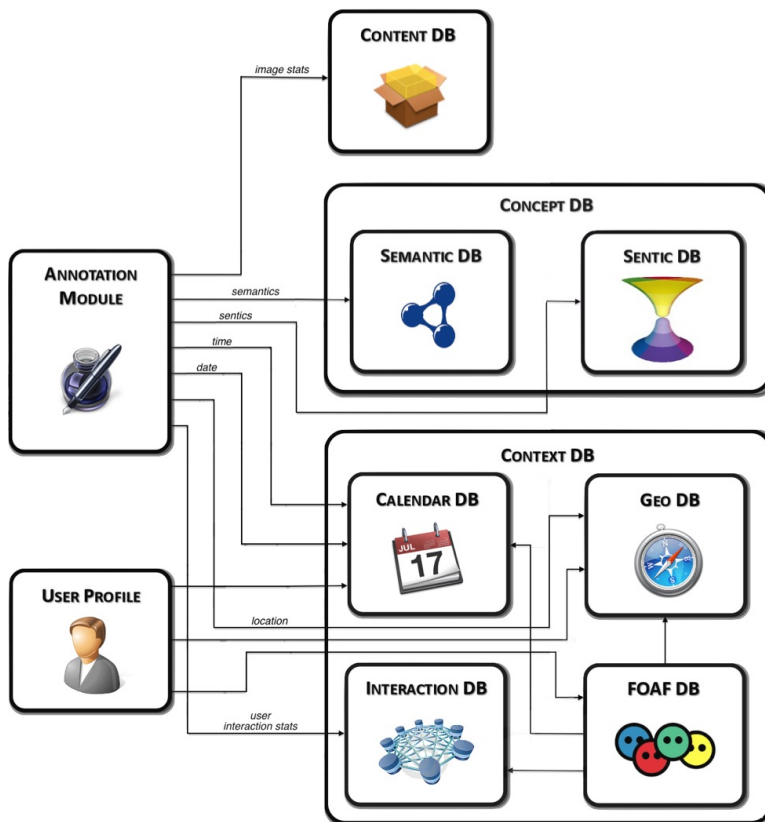


Fig. 5.6 Sentic Album's storage module. Image statistics are saved into the Content DB, semantics and sentsics are stored into the Concept DB, timestamp and geolocation are saved into the Context DB.

The Concept DB, in particular, consists of two databases, the Semantic DB and the Sentic DB, in which the cognitive and affective information associated with photo metadata are stored, respectively. The Context DB, in turn, is divided into four databases: the Calendar, Geo, FOAF (Friend Of A Friend), and Interaction DBs, which contain the information relative to timestamp, geolocation, social links, and social interaction, respectively. These databases are also integrated with information coming from the web profile of the user such as user's DOB (for the Calendar DB), user's current location (for the Geo DB), or user's list of friends (for the FOAF DB). The FOAF DB, in particular, plays an important role within the Context DB since it provides the other peer databases with information relative to user's social connections, e.g., relatives' birthdays or friends' location. Moreover, the Context DB receives extra contextual information from the inferred semantics. Personal names in the conceptual metadata are recognised by building a dictionary of first names from the Web and combining them with regular expressions to recognise full names. These are added to the database (in the FOAF DB) together with geographical places (in the Geo DB), which are also mined from databases on the Web and added to the parser's semantic lexicon.

As for the Semantic Web format [168], all the information related to pictures' metadata is stored in RDF/XML according to a set of predefined web ontologies. This operation aims to make the description of the semantics and sentics associated with pictures applicable to most online images coming from different sources, e.g., online photo sharing services, blogs, and social networks. To further this aim, it is necessary to standardise as much as possible the descriptors used in encoding the information about multimedia resources and people to which the images refer, in order to make it univocally interpretable and suitable to feed other applications. Hence, the ensemble of HEO, OMR, FOAF, and WNA is used again.

As for the storage of photo data and metadata in a matrix format, a dataset, termed '3CNet', is built, which integrates the information from the 3Cs in a unique knowledge base. The aim of this representation is to exploit principal component analysis (PCA) to later organise online personal images in a multi-dimensional vector space (as for AffectiveSpace) and, hence, reason on their similarity. 3CNet, in fact, is an $n \times m$ matrix whose rows are user's personal pictures IDs, whose columns are either content, concept, and context features (e.g., 'contains cold colours', 'conveys joy' or 'located in Italy'), and whose values indicate truth values of assertions. Therefore, in 3CNet, each image is represented by a vector in the space of possible features whose values are +1, for features that produce an assertion of positive valence, -1, for features that produce an assertion of negative valence, and 0 when nothing is known about the assertion.

The degree of similarity between two images, then, is the dot product between their rows in 3CNet. The value of such a dot product increases whenever two images are described with the same feature and decreases when they are described by features that are negations of each other. The main aim of the Search and Retrieval Module is to provide users with an IUI that allows them to easily manage, search, and retrieve their personal pictures online (Fig. 5.7).



Fig. 5.7 Sentic Album's search and retrieval module. The UI allows to browse personal images both by performing keyword-based queries and by adding/removing constraints on the facet properties.

Most of the existing photo management systems let users search for pictures through a keyword-based query, but results are hardly ever good enough since it is very difficult to come up with an ideal query from the user's initial request. The initial idea of an image the user has in mind before starting a search session, in fact, often deviates from the final results he/she will choose [289]. In order to let users start from a sketchy idea and then dynamically refine their search, the multi-faceted classification paradigm is adopted.

Personal images are displayed in a dynamic gallery that can be ordered according to different parameters, either textual or numeric, that is visual features (e.g., colour balance, hue, saturation, brightness, and contrast), semantics (i.e., common-sense concepts such as 'go jogging' and 'birthday party', but also people and objects contained in the picture), sentic (i.e., emotions conveyed by the picture and its polarity) and contextual information (e.g., time of capture, location, and social information such as users who viewed/commented on the picture).

In particular, NLP techniques similar to those used to process the image conceptual metadata are employed to analyse the text typed in the search box and, hence, perform queries on the SQL databases of the Storage Module. The order of visualisation of the retrieved images is given by the PQOP, so that images containing more relevant information at content, concept, and context level are first displayed. If, for example, the user is looking for pictures of his/her partner, Sentic Album initially proposes photos representing important events such as first date, first childbirth or honeymoon, that is, pictures with high PQOP.

Storage Module’s 3CNet is also exploited in the IUI, in order to find similar pictures. Towards the end of a search, the user sometimes may be interested in finding pictures similar to one of those so far obtained, even if this does not fulfill the constraints currently set via the facets. To serve this purpose, every picture is provided with a ‘like me’ button that opens a new Exhibit window displaying content, concept, and context related images, independently of any constraint.

Picture similarity is calculated by means of PCA and, in particular, through TSVD, as for AffectiveSpace. The number of singular values to be discarded (in order to reduce the dimensionality of 3CNet and, hence, reason on picture similarity) is chosen according to the total number of user’s online personal pictures and the amount of available metadata associated with them, i.e., according to size and density of 3CNet. Thus, by exploiting the information sharing property of TSVD, images specified by similar content, concept, and context are likely to have similar features and, hence, tend to fall near each other in the built-in vector space. Finally, the IUI also offers to display images according to date of caption on a timeline. Chronology, in fact, is a key categorisation concept for the management of personal pictures. Having the collection in chronological order is helpful for locating particular photos or events, since it is usually easier to remember when an event occurred relative to other events, as opposed to remembering its absolute date and time [164].

Many works dealing with object detection, scene categorisation, or content analysis on the cognitive level have been published, trying to bridge the semantic gap between represented objects and high-level concepts associated with them [172]. However, where affective retrieval and classification of digital media is concerned, publications, and especially benchmarks, are very few [186]. To overcome the lack of availability of relevant datasets, the performance and the user-friendliness of Sentic Album were tested on a topic and mood tagged evaluation dataset and through a usability test on a pool of 18 Picasa regular users, respectively.

For the system performance testing, in particular, 1,000 LiveJournal posts with labels matching Picasa tags such as ‘friends’, ‘travel’, and ‘holiday’, were selected in order to collect natural language text that is likely to have the same semantics as the conceptual metadata typical of personal photos. The classification test, hence, concurrently estimated the capacity of the system to infer both the cognitive and affective information (topic and mood tags, respectively) usually associated with online personal pictures (Table 5.3).

LiveJournal Tag	Precision	Recall	F-measure
art	62.9%	55.6%	59.0%
friends	77.2%	65.4%	70.8%
wedding	71.3%	60.4%	65.4%
holiday	68.9%	59.2%	63.7%
travel	81.6%	71.1%	75.9%
nature	67.5%	61.8%	64.5%

Table 5.3 Assessment of Sentic Album’s accuracy in inferring the cognitive (topic tags) and affective (mood tags) information associated with the conceptual metadata typical of personal photos.

For the usability test, users were asked to freely browse their online personal collections using Sentic Album IUI and to retrieve particular sets of pictures, in order to judge both usability and accuracy of the interface. Common queries included “find a funny picture of your best friend”, “search for the shots of your last summer holiday”, “retrieve pictures of you with animals”, “find an image taken on Christmas 2009”, “search for pictures of you laughing”, and “find a good picture of your mom”. From the test, it emerged that users really appreciate being able to dynamically and quickly set/remove constraints in order to display specific batches of pictures (which they cannot do in Picasa).

After the test session, participants were asked to fill-in an online questionnaire in which they were asked to rate, on a five-level scale, each single functionality of the interface according to their perceived utility. Concept facets and timeline, in particular, were found to be the most used by participants for search and retrieval tasks (Table 5.4). Users also really appreciated the ‘like me’ functionality, which was generally able to propose very relevant (semantically and affectively related) pictures (again not available in Picasa). When freely browsing their collections, users were particularly amused by the ability to navigate their personal pictures according to the emotion these conveyed, even though they did not always agree with the results.

Additionally, participants were not very happy with the accuracy of the search box, especially if they searched for one particular photo out of the entire collection. However, they always very much appreciated the order in which the pictures were proposed, which allowed them to quickly have all the most relevant pictures available as first results. 83.3% of test users declared that, despite not being as nifty as Picasa, Sentic Album is a very good photo management tool (especially for its novel semantic faceted search and PQOP functionalities) and they hope they could still be using it because, in the end, what really counts when browsing personal pictures is to find good matches in the shortest amount of time.

Feature	Not at all	Just a little	Somewhat	Quite a lot	Very much
Concept facets	0%	0%	5.6%	5.6%	88.8%
Content facets	77.8%	16.6%	5.6%	0%	0%
Context facets	16.6%	11.2%	5.6%	33.3%	33.3%
Search box	0%	11.2%	16.6%	33.3%	38.9%
Like me	0%	5.6%	5.6%	16.6%	72.2%
Timeline	0%	0%	0%	16.6%	83.4%
Sorting	11.2%	33.3%	33.3%	16.6%	5.6%

Table 5.4 Perceived utility of the different interface features by 18 Picasa regular users. Participants particularly appreciated the usefulness of concept facets and timeline, for search and retrieval tasks.

5.2 Development of HCI Systems

Human computer intelligent interaction is an emerging field aimed at providing natural ways for humans to use computers as aids. It is argued that, for a computer to be able to interact with humans, it needs to have the communication skills of humans. One of these skills is the affective aspect of communication, which is recognised to be a crucial part of human intelligence and has been argued to be more fundamental in human behaviour and success in social life than intellect [292, 221]. Emotions influence cognition and, therefore, intelligence, especially when it involves social decision-making and interaction.

The latest scientific findings indicate that emotions play an essential role in decision-making, perception, learning, and more. Most of the past research on affect sensing has considered each sense such as vision, hearing, and touch in isolation. However, natural human-human interaction is multi-modal: we communicate through speech and use body language (posture, facial expressions, gaze) to express emotion, mood, attitude, and attention. To this end, a novel fusion methodology is proposed, which is able to fuse any number of unimodal categorical modules, with very different time-scales, output labels, and recognition success rates, in a simple and scalable way. In particular, such a methodology is exploited to fuse the outputs of the opinion mining engine with the ones of a facial expression analyser for designing an embodied conversational agent with affective common-sense (section 5.2.1). This section, moreover, illustrates how the engine can be exploited for the development of HCI applications in fields such as instant messaging (IM) (section 5.2.2) and multimedia management (section 5.2.3).

5.2.1 *Sentic Avatar*

The capability of perceiving and expressing emotions through different modalities is a key issue for the enhancement of HCI. Natural human-human affective interaction is inherently multi-modal: people communicate emotions through multiple channels such as facial expressions, gestures, dialogues, etc. Although several studies prove that the multi-sensory fusion (e.g., audio, visual, and physiological responses) improves robustness and accuracy of human emotion analysis [149, 116, 313], most emotional recognition works still focus on increasing the success rates in sensing emotions from a single channel rather than merging complementary information across channels.

The multi-modal fusion of different affective channels is far from being solved and represents an active and open research issue [123]. Affect recognition from multiple modalities has a short historical background and is still in its first stage. It was not till 1998 that computer scientists attempted to use multiple modalities for recognition of emotions/affective states [244]. The results were promising: using multiple modalities improved the overall recognition accuracy helping the systems function in a more efficient and reliable way.

Following the findings in psychology, which suggested that the most significant channel for judging emotional cues of humans is the visual channel of face and body [12], a number of works combine facial expressions and body gestures for affect sensing [62, 122, 150]. Other approaches combine different biological information such as brain signals or skin conductivity for affect sensing [235, 33]. However, this research makes use of a single information channel, i.e., a single type of computer input device, and, therefore, must assume the reliability on this channel. For this reason, the trend in recent works is to consider and combine affective information coming from different channels. This way, eventual changes on the reliability of the different information channels are considered.

Recent literature on multi-modal affect sensing is focused on the fusion of data coming from the visual and audio channels. Most of those works make use of the visual channel for body gesture recognition [227] or facial expression classification [290] and the audio channel to analyse non-linguistic audio cues such as laughs [284], coughs [191], or cries [215]. However, very few works fuse information coming from the visual channel with linguistic-based (speech contents) audio affect sensing. With all these new areas of research in affect sensing, a number of challenges have arisen. In particular, the synchronisation and fusion of the information coming from different channels is a big problem to solve.

Previous studies fused emotional information either at a decision-level, in which the outputs of the unimodal systems are integrated by the use of suitable expert criteria [146] or at a feature-level, in which the data from both modalities are combined before classification [254]. In any case, the choice of fusion strategy depends on the targeted application. Accordingly, all available multi-modal recognisers have designed and/or used ad hoc solutions for fusing information coming from multiple modalities, but cannot accept new modalities without re-defining the whole system. In summary, there is not a general consensus when fusing multiple modalities and systems' scalability is not possible.

Sentic Avatar [42] is an embodied conversational agent (ECA) based on the multi-modal animation engine Maxine [17], which consists of four main modules: Perception, Affective Analysis, Deliberative/Generative, and Motor module (Fig. 5.8). The Perception module simply consists of the hardware necessary to gather the multi-modal information from the user, i.e., keyboard, microphone, and webcam. The Affective Analysis module aims to infer the user's affective state from the different inputs and integrate it. The Deliberative/Generative module is in charge of processing the extracted emotional information to manage the virtual agent's decisions and reactions, which are finally generated by the Motor module.

The Affective Analysis module is in charge of extracting emotions from the textual, vocal, and video inputs and integrating them. It consists of three main parts: the opinion mining engine, for inferring semantics and sentics associated with typed-in text and speech-to-text converted contents; the Facial Expression Analyser, for extracting affective information from video; and the Affective Integrator, for integrating the outputs coming from the two previous modules. The Facial Expression Analyser achieves an automatic classification of the shown facial expressions into discrete emotional categories. It is able to classify the user's emotion in terms of Ek-

man's six universal emotions (*fear, sadness, joy, disgust, surprise, and anger*) [99] plus *neutral*, giving a membership confidence value to each emotional category. The face modelling selected as input for the Facial Expression Analyser follows a feature-based approach: the inputs are a set of facial distances and angles calculated from feature points of the mouth, eyebrows, and eyes. In fact, the inputs are the variations of these angles and distances with respect to the neutral face. The points are obtained thanks to a real-time facial feature tracking program [65]. Fig. 5.9 shows, on the left side, the correspondence of these points with those defined by the MPEG4 standard. On the right side, in turn, the set of parameters obtained from these points is shown. In order to make the distance values consistent (independently of the scale of the image, the distance to the camera, etc.) and independent of the expression, all the distances are normalised with respect to the distance between the eyes, i.e., the MPEG4 Facial Animation Parameter Unit (FAPU). The choice of angles provides a size invariant classification and saves the effort of normalisation. As regards the classification process itself, the system intelligently combines the outputs of five different classifiers simultaneously.

In this way, the overall risk of making a poor selection with a given classifier for a given input is reduced. The classifier combination chosen follows a weighted majority voting strategy, where the voted weights are assigned depending on the performance of each classifier for each emotion. In order to select the best classifiers to combine, the Waikato environment for knowledge analysis (WEKA) tool was used [307]. This provides a collection of machine learning algorithms for data mining tasks. From this collection, five classifiers were selected after tuning: RIPPER, MLP, SVM, NB, and C4.5. The selection was based on their widespread use as well as on the individual performance of their WEKA implementation. To train the classifiers and evaluate the performance of the system, two different facial emotion databases were used: the FGNET database [293], which provides video sequences of 19 different Caucasian people, and the MMI Facial Expression Database [222], which holds 1280 videos of 43 different subjects from different races (Caucasian, Asian, and Arabic). Both databases are classified according to Ekman's six universal emotions plus *neutral*.

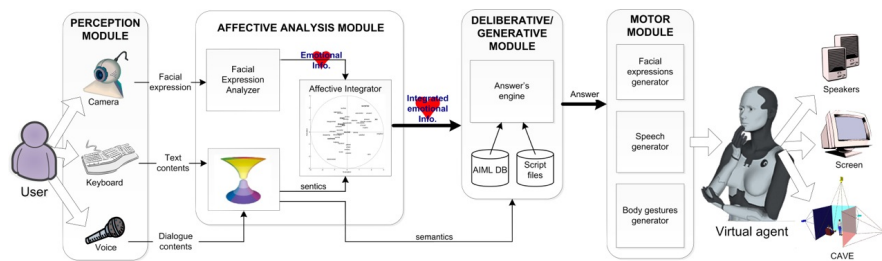


Fig. 5.8 Sentic Avatar's architecture. The system mainly consists of two modules for managing the avatar's inputs and outputs, and two modules for performing affective common-sense reasoning.




Fig. 5.9 Tracked facial feature points (on the left) and corresponding facial parameters (on the right). The variations of such distances and angles are exploited for affect recognition.

A new database has been built for testing this work with a total of 1500 static frames carefully selected from the apex of the video sequences from the FG-NET and MMI databases. The results obtained when applying the above-mentioned strategy for combining the scores of the five classifiers are shown in the form of confusion matrix in Table 5.5 (results have been obtained with a 10-fold cross-validation test over the 1500 database images). As it can be observed, the success rates for *neutral*, *joy*, and *surprise* are very high (84.44%–95.23%).

However, the system tends to confuse *disgust* with *fear*, *anger* with *disgust*, and *fear* with *surprise*; therefore, the performance for those emotions is slightly worse. The lowest result of the classification is for *sadness*: it is confused with *neutral* on 67.80% of occasions, due to the similarity of the facial expressions. Confusion between these pairs of emotions occurs frequently in the literature and for this reason many classification works do not consider some of them. Nevertheless, the results can be considered positive as two incompatible emotions (such as *sadness* and *joy* or *fear* and *anger*) are confused on less than 0.2% of occasions.

Another relevant aspect to be taken into account when evaluating the results is human opinion. The labels provided in the database for training classifiers correspond to the real emotions felt by users although they do not necessarily have to coincide with the perceptions other human beings may have about the facial expressions shown. Undertaking this kind of study is very important when dealing with human-computer interaction, since the system is proved to work in a similar way to the human brain. In order to take into account the human factor in the evaluation of the results, 60 persons were told to classify the 1500 images of the database in terms of emotions.

<i>classified as</i>	disgust	joy	anger	fear	sadness	neutral	surprise
disgust	79.41%	0%	2.39%	18.20%	0%	0%	0%
joy	4.77%	95.23%	0%	0%	0%	0%	0%
anger	19.20%	0%	74.07%	0%	3.75%	2.98%	0%
fear	9.05%	0%	0%	62.96%	8.53%	0%	19.46%
sadness	0.32%	0.20%	1.68%	0%	30.00%	67.80%	0%
neutral	0%	0%	1.00%	2.90%	4.10%	92.00%	0%
surprise	0%	0%	0%	11.23%	0%	4.33%	84.44%

Table 5.5 Confusion matrix obtained combining the five classifiers. Success rates for *neutral*, *joy*, and *surprise* are very high, but *disgust* with *fear*, *anger* with *disgust*, and *fear* with *surprise* tend to be confused.

<i>classified as</i>	disgust	joy	anger	fear	sadness	neutral	surprise
disgust	84.24%	0%	2.34%	13.42%	0%	0%	0%
joy	4.77%	95.23%	0%	0%	0%	0%	0%
anger	15.49%	0%	77.78%	0%	3.75%	2.98%	0%
fear	1.12%	0%	0%	92.59%	2.06%	0%	4.23%
sadness	0.32%	0.20%	1.68%	0%	66.67%	31.13%	0%
neutral	0%	0%	0%	0.88%	1.12%	98.00%	0%
surprise	0%	0%	0%	6.86%	0%	2.03%	91.11%


Table 5.6 Confusion matrix obtained after considering human assessment. Success ratios considerably increase, meaning that the adopted classification strategy is consistent with human classification.

As a result, each one of the frames was classified by 10 different people in 5 sessions of 50 images. The Kappa statistic obtained from raters' annotations is equal to 0.74 (calculated following the formula proposed in [255]), which indicates an adequate inter-rater agreement in the emotional images annotation. With this information, the evaluation of the results was repeated: the recognition was marked as good if the decision was consistent with that reached by the majority of the human assessors. The results (confusion matrix) of considering users' assessment are shown in Table 5.6. As it can be seen, the success ratios have considerably increased. Therefore, it can be concluded that the confusions of the algorithms go in the same direction as those of the users, which means that the adopted classification strategy is consistent with human classification.

The problem, now, is to fuse this affective classification with the emotion categorisation performed by the opinion mining engine, which outputs a list of sentic vectors that encompass the affective information associated with text and dialogue contents in terms of Pleasantness, Attention, Sensitivity, and Aptitude. To this end, a continuous 2D description of affect is considered [41]. Specifically, the emotion-related words corresponding to Ekman emotions and to the levels of the Hourglass of Emotions are mapped in the Whissell space [300], which allows to calculate a pair of values <activation, evaluation> from the obtained labels and, hence, to concurrently visualise and compare them (Fig. 5.10).

Spoken sentence: Wow! This is so great!

Video sequence:



Whissell output:




Fig. 5.10 A screenshot of the Affective Analysis module output. The integration of the multi-modal emotional information takes place into the Whissell space, in terms of activation and evaluation.

The evaluation dimension measures how humans feel, from positive to negative. The activation dimension measures whether humans are more or less likely to take some action under the emotional state, from active to passive. In particular, the process of affective integration is achieved through the following three steps. Firstly, each one of the emotional labels inferred by the opinion mining engine is mapped as a 2D point in the Whissell space. Secondly, the Facial Expression Analyser outputs the user's emotion in terms of Ekman's six universal emotions (plus *neutral*), giving a membership confidence value to each emotional category. The mapping of its output in the Whissell space is carried out considering each of Ekman's six basic emotions plus *neutral* as 2D weighted points in the \langle activation, evaluation \rangle space, where the weights are assigned depending on the confidence value obtained for each emotion in the classification process. The final detected emotion is calculated as the centre of mass of the seven weighted points in the Whissell space. This way, the Facial Expression Analyser outputs one emotional location in the Whissell space per frame of the studied video sequence. Finally, the whole set of 2D \langle activation, evaluation \rangle points obtained from both the opinion mining engine and the Facial Expression Analyser is fitted to the Minimum Volume Ellipsoid (MVE) that better covers the shape of the set of extracted points. The MVE is calculated following the algorithm described by Kumar and Yildirim [158].

metric	disgust	joy	anger	fear	sadness	neutral	surprise
precision	86.3%	95.4%	80.0%	94.0%	79.8%	95.6%	92.0%
recall	81.7%	99.1%	95.1%	85.4%	91.1%	72.8%	96.2%
F-measure	83.9%	97.2%	86.9%	89.5%	85.1%	82.7%	94.0%

Table 5.7 Precision, recall, and F-measure rates associated with Ekman’s basic emotions (plus neutral) obtained by concurrently employing the opinion mining engine and the Facial Expression Analyser.

The final emotional information outputted by the Affective Analysis module for the whole video sequence is given by the x-y coordinates of the centre of that MVE. In order to evaluate the performance of the Affective Analysis module, that is the F-measure rates of the opinion mining engine and the Facial Expression Analyser as an ensemble, an emotion recognition test on 50 selected videos from the HUMAINE database¹⁴ was performed. As shown in Table 5.7, the concurrent exploitation of different modalities generally leads to an improvement in affect detection, as the two subsystem components make uncorrelated errors. When multiple components make uncorrelated errors, in fact, the probability that they all make the same error is the product of their individual error probabilities, resulting in much lower error rates.

5.2.2 Sentic Chat

Online communication is an extremely popular form of social interaction. Unlike face-to-face communication, online IM tools are extremely limited in conveying emotions or the context associated with a communication. Users have adapted to this environment by inventing their own vocabulary, e.g., by putting actions within asterisks (“I just came from a shower *shivering*”), or by using emoticons, by addressing a particular user in a group communication (“@Ravi”). Such evolving workarounds clearly indicate a latent need for a richer, more immersive user experience in social communication. This problem is addressed by exploiting the semantics and sentics associated with the on-going communication to develop an adaptive user interface (UI) capable to change according to content and context of the online chat. Popular approaches to enhance and personalise computer-mediated communication (CMC) include emoticons, skins, avatars, and customisable status messages.

However, all these approaches require explicit user configuration or action: the user needs to select the emoticon, status-message, or avatar that best represents him/her. Furthermore, most of these enhancements are static - once selected by the user, they do not adapt themselves automatically. There is some related work on automatically updating the status of the user by analysing various sensor data available on mobile devices [196].

¹⁴ <http://humaine-db.sspnet.eu>



Fig. 5.11 A few screenshots of Sentic Chat UI. Stage and actors gradually change, according to the semantics and sentsics associated with the on-going conversation, to provide an immersive chat experience.

However, most of these personalisation approaches are static and do not automatically adapt. The approach of Sentic Chat [69] is unique in that it is: intelligent, as it analyses content and does not require explicit user configuration; adaptive, as the UI changes according to communication content and context; inclusive, as the emotions of one or more participants in the chat session are analysed to let the UI adapt dynamically. The module architecture can be deployed either on the cloud (if the client has low processing capabilities) or on the client (if privacy is a concern). Most IM clients offer a very basic UI for text communication. In Sentic Chat, the focus is on extracting the semantics and sentsics embedded in the text of the chat session to provide an UI that adapts itself to the mood of the communication.

For this prototype application, the weather metaphor was selected, as it is scalable and has previously been used effectively to reflect the subject's mood [70] or content's 'flavour' [216]. In the proposed UI, if the detected mood of the conversation is 'happy', the UI will reflect a clear sunny day. Similarly a gloomy weather reflects a melancholy tone in the conversation (Fig. 5.11). Of course, this is a subjective metaphor - one that supposedly scales well with conversation analysis. In the future, other relevant scalable metaphors could be explored, e.g., colours [132].

The adaptive UI primarily consists of three features: the stage, the actors, and the story. For any mapping, these elements pay a crucial role in conveying the feel and richness of the conversation mood, e.g., in the 'happy' conversation the weather 'clear sunny day' will be the stage, the actors will be lush green valley, the rainbow, and the cloud, which may appear or disappear as per the current conversation tone of the story. The idea is similar to a visual narrative of the mood the conversation is in; as the conversation goes on, the actors may come in or go off as per the tone of the thread. By analysing the semantics and sentsics associated with communication content (data) and context (metadata), the UI may adapt to include images of landmarks from remote-user's location (e.g., Times Square), images about concepts in the conversation (pets, education, etc.), or time of day of remote user (e.g., sunrise or dusk).

Feature	Not consistent	Consistent	Very consistent
stage change	0%	83.3%	16.7%
actor alternation	16.8%	66.6%	16.7%

Table 5.8 Perceived consistency with chat text of stage change and actor alternation. The evaluation was performed on a 130-minute chat session operated by a pool of 6 regular chat users.

The effectiveness of Sentic Chat was assessed through a usability test on a group of 6 regular chat users, who were asked to chat to each other pairwise for approximately 10 minutes (for a total of 130 minutes of chat data) and to rate the consistency with the story of both stage and actor alternation during the CMC (Table 5.8).

5.2.3 *Sentic Corner*

In a world in which web users are continuously blasted by ads and often compelled to deal with user-unfriendly interfaces, we sometimes feel like we want to evade from the sensory overload of standard web pages and take refuge in a safe web corner, in which contents and design are in harmony with our current frame of mind. Sentic Corner [47] is an IUI that dynamically collects audio, video, images, and text related to the user's current feelings and activities as an interconnected knowledge base, which is browsable through a multi-faceted classification website. In the new realm of Web 2.0 applications, the analysis of emotions has undergone a large number of interpretations and visualisations, e.g., We Feel Fine¹⁵, MoodView¹⁶, MoodStats¹⁷, and MoodStream¹⁸, which have often led to the development of emotion-sensitive systems and applications.

Nonetheless, today web users still have to almost continuously deal with sensory-overloaded web pages, pop-up windows, annoying ads, user-unfriendly interfaces, etc. Moreover, even for websites uncontaminated by web spam, the affective content of the page is often totally unsynchronised with the user's emotional state. Web pages containing multimedia information inevitably carry more than just informative content. Behind every multimedia content, in fact, there is always an emotion.

Sentic Corner exploits this concept to build a sort of parallel cognitive/affective digital world in which the most relevant multimedia contents associated to the users' current moods and activities are collected, in order to enable them, whenever they want to evade from sensory-rich, overwrought, and earnest web pages, to take refuge in their own safe web corner. To our knowledge, there is still no published study on the task of automatically retrieving and displaying multimedia contents according to user's moods and activities, although the affective and seman-

¹⁵ <http://wefeelfine.org>

¹⁶ <http://moodviews.com>

¹⁷ <http://moodstats.com>

¹⁸ <http://moodstream.gettyimages.com>

tic analysis of video, audio, and textual contents have been separately investigated extensively [127, 270, 252]. The most relevant commercial tool within this area is Moodstream, a mashup of several forms of media, designed to bring users music, images, and video according to the mood they manually select on the web interface. Moodstream aims to create a sort of audio-visual ambient mix that can be dynamically modified by users by selecting from the presets of ‘inspire’, ‘excite’, ‘refresh’, ‘intensify’, ‘stabilise’, and ‘simplify’, e.g., mixtures of mood spectra on the Moodstream mixer such as happy/sad, calm/lively, or warm/cool. Users can start with a preset and then mix things up including the type of image transition, whether they want more or less vocals in their music selection, and how long images and video will stay, among other settings. In Moodstream, however, songs are not played entirely but blended into one another every 30 seconds and, even if the user has control on the multimedia flow through the mood presets, he/she cannot actually set a specific mood and/or activity as a core theme for the audio-visual ambient mix.

Sentic Corner, on the contrary, uses sentic computing to automatically extract semantics and sentics associated with user’s status updates on micro-blogging websites and, hence, to retrieve relevant multimedia contents in harmony with his/her current emotions and motions. The module for the retrieval of semantically and affectively related music is termed Sentic Tuner. The relevant audio information is pulled from Stereomood¹⁹, an emotional online radio that provides music that best suits users’ mood and activities. In the web interface, music is played randomly through an online music player with the possibility for the user to play, stop, and skip tracks. In Stereomood, music tracks are classified according to some tags that users are supposed to manually choose in order to access a list of semantically or affectively related songs. These tags are either mood tags (e.g., ‘happy’, ‘calm’, ‘romantic’, ‘lonely’, and ‘reflective’) or activity tags (such as ‘reading’, ‘just woke up’, ‘dressing up’, ‘cleaning’, and ‘jogging’), the majority of which represent cognitive and affective knowledge contained in AffectiveSpace as common-sense concepts and emotional labels.

The Sentic Tuner uses the mood tags as centroids for affective blending and the activity tags as seeds for spectral association, in order to build a set of affectively and semantically related concepts respectively, which will be used at run-time to match the concepts extracted from user’s micro-blogging activity. The Sentic Tuner also contains a few hundreds *rāgas* (Sanskrit for moods), which are melodic modes used in Indian classical music meant to be played in particular situations (mood, time of the year, time of the day, weather conditions, etc.). It is considered inappropriate to play *rāgas* at the wrong time (it would be like playing Christmas music in July, lullabies at breakfast, or sad songs at a wedding) so these are played just when semantics and sentics exactly match time and mood specifications in the *rāgas* database. Hence, once semantics and sentics are extracted from natural language text through the opinion mining engine, Stereomood API and the *rāgas* database are exploited to select the most relevant tracks to user’s current feelings and activities.

¹⁹ <http://stereomood.com>

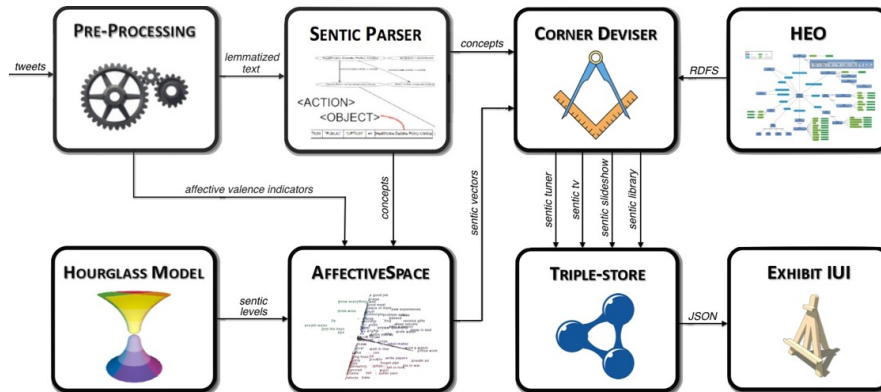


Fig. 5.12 Sentic Corner generation process. The semantics and sentics extracted from the user's micro-blogging activity are exploited to retrieve relevant audio, video, visual, and textual information.

Sentic TV is the module for the retrieval of semantically and affectively related videos. In particular, the module pulls information from Jinni²⁰, a new site that allows users to search for video entertainment in many specific ways. The idea behind Jinni is to reflect how people really think and talk about what they watch. It is based on an ontology developed by film professionals and new titles are indexed with an innovative NLP technology for analysing metadata and reviews. In Jinni, users can choose from movies, TV shows, short films, and online videos to find specific genres or what they are in the mood to watch. In particular, users can browse videos by topic, mood, plot, genre, time/period, place, audience, and praise. Similarly to the Sentic Tuner, Sentic TV uses Jinni's mood tags as centroids for affective blending and the topic tags as seeds for spectral association, in order to retrieve affectively and semantically related concepts respectively. Time tags and location tags are also exploited in case relevant time-stamp and/or geo-location information is available within user's micro-blogging activity.

Sentic Corner also offers semantically and affectively related images through the Sentic Slideshow module. Pictures related to the user's current mood and activity are pulled from Fotosearch²¹, a provider of royalty free and rights managed stock photography that claims to be the biggest repository of images on the Web. Since Fotosearch does not offer a priori mood tags and activity tags, the CF-IOF technique is used on a set of 1000 manually tagged (according to mood and topic) tweets, in order to find seeds for spectral association (topic-tagged tweets) and centroids for affective blending (mood-tagged tweets). Each of the resulting concepts is used to retrieve mood and activity related images through the Fotosearch search engine.

²⁰ <http://jinni.com>

²¹ <http://fotosearch.com>

The royalty free pictures, eventually, are saved in an internal database according to their mood and/or activity tag, in a way that they can be quickly retrieved at run-time, depending on user's current feelings and thoughts. The aim of Sentic Library is to provide book excerpts depending on user's current mood. The module proposes random book passages users should read according to the mood they should be in while reading it and/or what mood they will be in when they have finished. The excerpt database is built according to '1001 Books for Every Mood: A Bibliophile's Guide to Unwinding, Misbehaving, Forgiving, Celebrating, Commiserating' [101], a guide in which the novelist Hallie Ephron serves up a literary feast for every emotional appetite. In the guide, books are labelled with mood tags such as 'for a good laugh', 'for a good cry', and 'for romance', but also some activity tags such as 'for a walk on the wild side' or 'to run away from home'.

As for Sentic TV and Sentic Tuner, Sentic Library uses these mood tags as centroids for affective blending and the topic tags as seeds for spectral association. The Corner Deviser exploits the semantic and sentic knowledge bases previously built by means of blending, CF-IOF and spectral association to find matches for the concepts extracted by the sentic parser and their relative affective information inferred by AffectiveSpace. Such audio, video, visual, and textual information (namely Sentic Tuner, Sentic TV, Sentic Slideshow, and Sentic Library) is then encoded in RDF/XML according to HEO and stored in a triple-store (Fig. 5.12).

The screenshot displays the Sentic Corner web interface. At the top, a yellow header contains the title "Sentic Corner". Below the header, there is a search box and several filter panels on the left. The main content area shows 16 items, with the first two items displayed in detail. The first item is a picture titled "bubbles" showing a child running in a grassy field. The second item is a video titled "Le Fabuleux destin d'Amélie Poulain" showing a man and a woman on a bicycle. On the right side, there are several filter panels for "Sentic Slideshow", "Semantics", "Sentic", and "Location".

Sentic Corner

Search Box

Sentic TV

Semantics

- 2 Scrabble
- 2 auto-rickshaw
- 1 camel
- 1 children
- 1 cow
- 1 cycle-rickshaw

Sentic

- 1 admiration
- 1 anger
- 1 apprehension
- 2 distraction
- 1 interest
- 2 joy

Sentic Library

Semantics

- 1 cow
- 1 cycle-rickshaw
- 1 engine
- 1 food
- 1 mechanic
- 1 sushi

Sentic

- 1 distraction
- 1 interest

Sentic Corner retrieved 3 semantically related pictures
 Sentic Corner retrieved 7 affectively related pictures
 Sentic Corner retrieved 2 semantically related videos
 Sentic Corner retrieved 4 affectively related videos
 Sentic Corner is now playing: Yellow Submarine, The Beatles

16 Items sorted by: semantics and sentics grouped as sorted

1. **bubbles**

Type: Picture
 Source: Fotosearch
 Geotag: Avignon
 Description: Child running after bubbles
[Go to source](#)
[Like me](#)

2. **Le Fabuleux destin d'Amélie Poulain**

Type: Video
 Source: Itini
 Geotag: Paris
 Description: Amélie and Nino on a bicycle
[Go to source](#)
[Like me](#)

Sentic Slideshow

Semantics

- 2 engine
- 3 food
- 1 mechanic
- 1 sushi
- 2 temple
- 2 wedding

Sentic

- 3 distraction
- 1 interest
- 1 joy
- 2 pensiveness
- 2 sadness
- 1 serenity

Location

- 5 Paris
- 1 Avignon
- 1 Goa
- 3 Kochi
- 2 Mangalore
- 2 Mumbai

Fig. 5.13 Sentic Corner web interface. The multi-modal information obtained by means of Sentic Tuner, Sentic TV, Sentic Slideshow, and Sentic Library is encoded in RDF/XML for multi-faceted browsing.

Content	Not at all	Just a little	Somewhat	Quite a lot	Very much
audio	0%	11.1%	11.1%	44.5%	33.3%
video	11.1%	11.1%	44.5%	33.3%	0%
visual	0%	0%	22.2%	33.3%	44.5%
textual	22.2%	11.1%	55.6%	11.1%	0%

Table 5.9 Relevance of audio, video, visual, and textual information assembled over 80 tweets. Because of their larger datasets, Sentic Tuner and Slideshow are the best-performing modules.

In case the sentics detected belong to the lower part of the Hourglass, the multimedia contents searched will have an affective valence opposite to the emotional charge detected, as Sentic Corner aims to restore the positive emotional equilibrium of the user, e.g., if the user is angry he/she might want to calm down. The Exhibit IUI module, eventually, visualises the contents of the Sesame database exploiting the multi-faceted categorisation paradigm (Fig. 5.13).

In order to test the relevance of multimedia content retrieval, an evaluation based on the judgements of 8 regular Twitter users was performed. Specifically, users had to link Sentic Corner to their Twitter accounts and evaluate, over 10 different tweets, how the IUI would react to their status change in terms of relevance of audio, video, visual, and textual information assembled by Sentic Corner. The multimedia contents retrieved turned out to be pretty relevant in most cases, especially for tweets concerning concrete entities and actions (Table 5.9).

5.3 Development of E-Health Systems

In health-care, it has long been recognised that, although the health professional is the expert in diagnosing, offering help, and giving support in managing a clinical condition, the patient is the expert in living with that condition. Health-care providers need to be validated by somebody outside the medical departments but, at the same time, inside the health-care system. The best candidate for this is not the doctor, the nurse, or the therapist, but the real end-user of health-care – none other than the patient him/herself.

Patient 2.0 is central to understanding the effectiveness and efficiency of services and how they can be improved. The patient is not just a consumer of the health-care system but a quality control manager – his/her opinions are not just reviews of a product/service but more like small donations of experience, digital gifts which, once given, can be shared, copied, moved around the world, and directed to just the right people who can use them to improve health-care locally, regionally, or nationally. Web 2.0 dropped the cost of voice, of finding others ‘like me’, of forming groups, of obtaining and republishing information, to zero. As a result, it becomes easy and rewarding for patients and carers to share their personal experiences with the health-care system and to research conditions and treatments.

To bridge the gap between this social information and the structured information supplied by health-care providers, the opinion mining engine is exploited to extract the semantics and sentsics associated with patient opinions over the Web and, hence, provide the real end-users of the health system with a common framework to compare, validate, and select their health-care providers (section 5.3.1). This section, moreover, shows how the engine can be used as an embedded tool for improving patient reported outcome measures (PROMs) for health related quality of life (HRQoL), that is to record the level of each patient's physical and mental symptoms, limitations, and dependency (section 5.3.2).

5.3.1 Crowd Validation

As Web 2.0 dramatically reduced the cost of communication, today it is easy and rewarding for patients and carers to share their personal experiences with the health-care system. This social information, however, is often stored in natural language text and, hence, intrinsically unstructured, which makes comparison with the structured information supplied by health-care providers very difficult. To bridge the gap between these data, which though different at structure-level are similar at concept-level, a patient opinion mining tool has been proposed to provide the end-users of the health system with a common framework to compare, validate, and select their health-care providers.

In order to give structure to online patient opinions, both the semantics and sentsics associated with these are extracted in a way that it is possible to map them to the fixed structure of health-care data. This kind of data, in fact, usually consists of ratings that associate a polarity value to specific features of health-care providers such as communication, food, parking, service, staff, and timeliness. The polarity can be either a number in a fixed range or simply a flag (positive/negative). In the proposed approach, structure is added to unstructured data by building semantics and sentsics on top of it (Fig. 5.14).




Fig. 5.14 The semantics and sentsics stack. Semantics are built on the top of data and metadata. Sentsics are built on the top of semantics, representing the affective information associated with these.

In particular, given a textual resource containing a set of opinions O about a set of topics T with different polarity $p \in [-1, 1]$, the subset of opinions $o \subseteq O$ is extracted, for each $t \in T$, and p is determined for each o . In other words, since each opinion can regard more than one topic and the polarity values associated with each topic are often independent from each other, a set of topics is extracted from each opinion and then, for each topic detected, the polarity associated with it is inferred. Once natural language data are converted to a structured format, each topic expressed in each patient opinion and its related polarity can be aggregated and compared. These can then be easily assimilated with structured health-care information contained in a database or available through an API.

This process is termed crowd validation [51] (Fig. 5.15), because of the feedback coming from the masses, and it fosters next-generation health-care systems, in which patient opinions are crucial in understanding the effectiveness and efficiency of health services and how they can be improved. Within this work, in particular, the opinion mining engine is used to marshal PatientOpinion's social information in a machine-accessible and machine-processable format and, hence, compare it with the official hospital ratings provided by NHS Choices²² and each NHS trust. The inferred ratings are used to validate the information declared by the relevant health-care providers (crawled separately from each NHS trust website) and the official NHS ranks (extracted using NHS Choices API). At the present time, crowd validation cannot be directly tested because of the impossibility to objectively assess the truthfulness of both patient opinions and official NHS ratings.

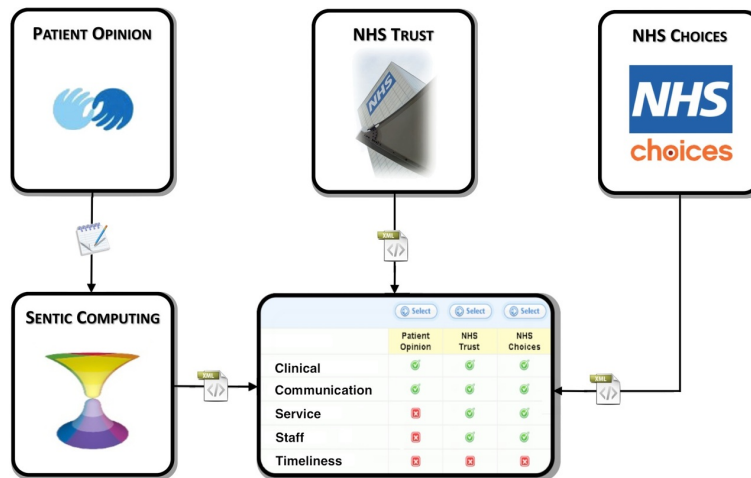


Fig. 5.15 The crowd validation schema. PatientOpinion stories are encoded in a machine-accessible format, in a way that they can be compared with the ratings provided by NHS Choices and each NHS trust.

²² <http://www.nhs.uk>

Given that the patient opinion mining performance of the system has already been tested (Table 4.7), however, an experimental investigation has been performed over a set of 200 patient opinions about three different NHS trusts, for which self-assessed ratings were crawled from each hospital website and official NHS ranks were obtained through NHS Choices API. Results showed an average discrepancy of 39% between official and unofficial ratings, which sounds plausible as, according to Panorama²³, 60% of hospitals inspected in 2010 gave inaccurate information to the government in assessing their own performance.

5.3.2 Sentic PROMs

Public health measures such as better nutrition, greater access to medical care, improved sanitation, and more widespread immunisation have produced a rapid decline in death rates in all age groups. Since there is no corresponding decline in birth rates, however, the average age of population is increasing exponentially. If we want health services to keep up with such monotonic growth, we need to automatise as much as possible the way patients access the health-care system, in order to improve both its service quality and timeliness. Everything we do that does not provide benefit to patients or their families, in fact, is waste. To this end, a new generation of short and easy-to-use tools to monitor patient outcomes and experience on a regular basis has been recently proposed by Benson et al. [21]. Such tools are quick, effective, and easy to understand, as they are very structured. However, they leave no space to those patients who would like to say something more.

Patients, in fact, are usually keen on expressing their opinions and feelings in free text, especially if driven by particularly positive or negative emotions. They are often happy to share their health-care experiences for different reasons, e.g., because they seek for a sense of togetherness in adversity, because they benefited from others' opinions and want to give back to the community, for cathartic complaining, for supporting a service they really like, because it is a way to express themselves, because they think their opinions are important for others. When people have a strong feeling about a specific service they tried, they feel like expressing it. If they loved it, they want others to enjoy it. If they hated it, they want to warn others away.

Standard PROMs allow patients to easily and efficiently measure their HRQoL but, at the same time, they limit patients' capability and will to express their opinions about particular aspects of the health-care service that could be improved or important facets of their current health status. Sentic PROMs [37], in turn, exploit the ensemble application of standard PROMs and sentic computing to allow patients to evaluate their health status and experience in a semi-structured way, i.e., both through a fixed questionnaire and through free text. PROMs provide a means of gaining an insight into the way patients perceive their health and the impact that treatments or adjustments to lifestyle have on their quality of life.

²³ <http://www.bbc.co.uk/programmes/b00rfqfm>

Pioneered by Donabedian [95], health status research began during the late 1960s with works focusing on health-care evaluation and resource allocation. In particular, early works mainly aimed to evaluate health states for policy and economic evaluation of health-care programmes, but devoted little attention to the practicalities of data collection [105, 281, 81]. Later works, in turn, aimed to develop lengthy health profiles to be completed by patients, leading to the term patient reported outcome [295, 22]. PROMs can provide a new category of real-time health information, which enables every level of the health service to focus on continuously improving those things that really matter to patients. The benefits of routine measurement of HRQoL include helping to screen for problems, promoting patient-centric care, aiding patients and doctors to take decisions, improving communication amongst multi-disciplinary teams, and monitoring progress of individual or groups of patients and the quality of care in a population. However, in spite of demonstrated benefits, routine HRQoL assessment in day-to-day practice remains rare as few patients are willing to spend the time needed to daily fill-in questionnaires, such as *SF-36* [297], *SF-12* [296], *Euroqol EQ-5D* [31], or the *Health Utilities Index* [139].

To overcome this problem, a new generic PROM, termed *howRu* [21], was recently proposed for recording the level of each patient's physical and mental symptoms, limitations, and dependency on four simple levels. The questionnaire was designed to take no more than a few seconds using electronic data collection and integration with electronic patient records as part of other routine tasks that patients have to do, such as booking appointments, checking in on arrival at clinic, and ordering or collecting repeat medication. The main aim of *howRu* is to use simple terms and descriptions to reduce the risk of ambiguity and to ensure that as many people as possible can use the measure reliably and consistently without training or support. The same approach has been employed to monitor also patient experience (*howRwe*) and staff satisfaction (*howRus*) on a regular basis. These questionnaires have been proved to be quick, effective, and easy to understand, as they are short, rigid, and structured. However, such structuredness can be very limiting, as it leaves no space to those patients who would like to say something more about their health or the service they are receiving. Patients, especially when driven by particularly positive or negative emotions, do want to express their opinions and feelings.

Sentic PROMs allow patients to assess their health status and health-care experience in a semi-structured way by enriching the functionalities of *howRu* with the possibility of adding free text (Fig. 5.16). This way, when patients are happy with simply filling-in the questionnaire, they can just leave the input text box blank but, when they feel like expressing their opinions and feelings, e.g., in the occasion of a particularly positive or negative situation or event, they can now do it in their own words. Hence, Sentic PROMs input data, although very similar at concept level, are on two completely different structure levels – structured (questionnaire selection) and unstructured (natural language). As we would like to extract meaningful information from such data, the final aim of Sentic PROMs is to format the unstructured input and accordingly aggregate it with the structured data, in order to perform statistical analysis and pattern recognition.

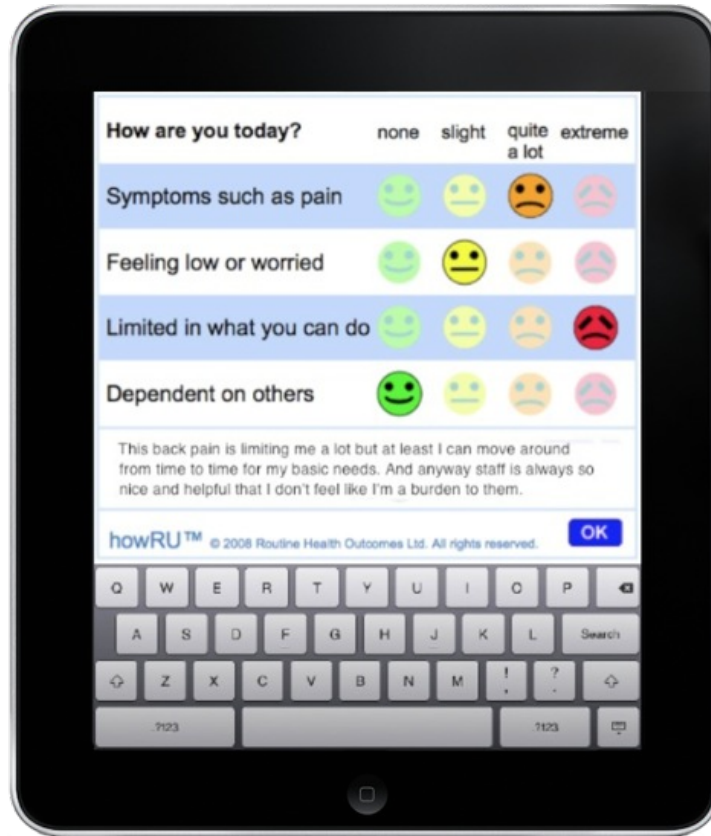


Fig. 5.16 Sentic PROMs prototype on iPad. The new interface allows patients to assess their health status and health-care experience both in a structured (questionnaire) and unstructured (free text) way.

In particular, the gap between unstructured and structured data is bridged by means of sentic computing. Among the benefits of questionnaires' structuredness, there are the quickness, effectiveness, and ease to use and understand. However, such structuredness involves some drawbacks. A questionnaire, in fact, can limit the possibility to discover new important patterns in the input data and can constrain users to omit important opinions that might be valuable for measuring service quality. In the medical sphere, in particular, patients driven by very positive or very negative emotions are usually willing to detailedly express their point of view, which can be particularly valuable for assessing uncovered points, raising latent problems, or redesigning the questionnaire. To this end, Sentic PROMs adopt a semi-structured approach that allows patients to assess their health status and health-care experience both by filling in a four-level questionnaire and by adding free text.

The analysis of free text, moreover, allows Sentic PROMs to measure patients' physio-emotional sensitivity. The importance of physio-emotional sensitivity in humans has been proven by recent health research, which has shown that individuals who feel loved and supported by friends and family, or even by a loving pet, tend to have higher survival rates following heart attacks than other cardiac patients who experience a sense of social isolation. Such concept is also reflected in natural language as we use terms such as 'heartsick', 'broken-hearted', and 'heartache' to describe extreme sadness and grief, idioms like 'full of gall' and 'venting your spleen' to describe anger, and expressions such as 'gutless', 'yellow belly', and 'feeling kicked in the gut' to describe shame. Human body contracts involuntarily when it feels emotional pain such as grief, fear, disapproval, shock, helplessness, shame, and terror, in the same reflex it does if physically injured. Such gripping reflex normally releases slowly, but if a painful experience is intense, or happens repeatedly, the physio-emotional grip does not release and constriction is retained in the body. Any repeated similar experience then layers on top of the original unreleased contraction, until we are living with layers of chronic tension, which constricts our bodies. The mind, in fact, may forget the origin of pain and tension, but the body does not.

Besides HRQoL measurement, Sentic PROMs aim to monitor also users' physio-emotional sensitivity on a regular basis, as a means of patient affective modelling. In particular, the dimensional affective information coming from both questionnaire data (*howRu* aggregated score) and natural language data (sentic vectors) is stored separately by the system every time patients conclude a Sentic PROM session and plotted on four different bi-dimensional diagrams. Such diagrams represent the pairwise fusion of the four dimensions of the Hourglass model and allow to detect more complex (compound) emotions that can be particularly relevant for monitoring patients' physio-emotional sensitivity, e.g., frustration, anxiety, optimism, disapproval, and rejection.

A preliminary validation study was undertaken to examine the psychometric properties and construct validity of Sentic PROMs and to compare these with *SF-12*. In particular, 2,751 subjects with long-term conditions (average age 62, female 62.8%), were classified by *howRu* score, primary condition, number of conditions suffered, age group, duration of illness, and area of residence. Across all six classifications, the correlation of the mean *howRu* scores with the mean values of the Physical Components Summary (*PCS-12*), the Mental Components Summary (*MCS-12*), and the sum of *PCS-12* + *MCS-12* were generally very high (0.91, 0.45, and 0.97, respectively).

5.4 Conclusions

In order to assess the capability of the opinion mining engine (or parts of it) to tackle real-world NLP tasks, the process for the extraction of semantics and sentics has been embedded in multiple systems for the development of intelligent applications in different domains. Specifically, sentic computing tools and techniques

have been exploited for the design of Social Web applications, i.e., a troll filtering system, a social media marketing tool, and an online personal photo management system (section 5.1); HCI applications, that is, an embodied conversational agent with affective common-sense, an adaptive IM tool, and an emotion-sensitive IUI (section 5.2); and e-health applications, i.e., a framework for assessing the quality of health-care providers and a tool for enhancing PROMs for HRQoL (section 5.3).

All these developed applications demonstrate how the opinion mining engine can be employed in nearly any domain, as it does not rely on domain-dependent keywords, but rather on common-sense knowledge bases that allow it to extrapolate the cognitive and affective information associated with natural language text. The tools and techniques employed for the development of such engine are summarised in the following chapter, where a discussion about limitations and future developments of these is also offered.

Chapter 6

Concluding Remarks

*It is difficult to be rigorous about whether a machine really knows and thinks,
because we are hard put to define these things.
We understand human mental processes only slightly better than
a fish understands swimming.*

John McCarthy

This book was the result of a doctoral work developed within an industrial research project born from the collaboration between the University of Stirling, the MIT Media Laboratory, and Sitekit Labs. The main aim of this book was to go beyond keyword-based approaches by further developing and applying common-sense computing techniques to bridge the cognitive and affective gap between word-level natural language data and the concept-level opinions conveyed by these. This has been pursued through a variety of novel tools and techniques that have been tied together to develop an opinion mining engine for the semantic analysis of natural language opinions and sentiments. Such engine has then been used for the development of intelligent web applications in fields such as Social Web, HCI, and e-health.

This chapter contains a summary of the contributions the book has introduced (section 6.1), a discussion about limitations and future developments of these (section 6.2), and conclusions (section 6.3).

6.1 Summary of Contributions

Despite significant progress, opinion mining and sentiment analysis are still finding their own voice as new inter-disciplinary fields. Engineers and computer scientists use machine learning techniques for automatic affect classification from video, voice, text, and physiology. Psychologists use their long tradition of emotion research with their own discourse, models, and methods.

This work has assumed that opinion mining and sentiment analysis are research fields inextricably bound to the affective sciences that attempt to understand human emotions. Simply put, the development of affect-sensitive systems cannot be divorced from the century-long psychological research on emotion. The emphasis on the multi-disciplinary landscape that is typical for emotion-sensitive applications and the need for common-sense sets this work apart from previous research on opinion mining and sentiment analysis.

In this book, a novel approach to opinion mining and sentiment analysis has been developed by exploiting both AI and Semantic Web techniques. In particular, two common-sense knowledge bases have been developed and an ensemble application of graph mining and multi-dimensionality reduction techniques has been employed to perform reasoning on them. Such tools and techniques have then been embedded in multiple systems in fields such as Social Web, HCI, and e-health. This section presents a summary of the sentic computing techniques (section 6.1.1), tools (section 6.1.2), and applications (section 6.1.3) developed within this work.

6.1.1 Techniques

In this book, a variety of techniques have been developed for the analysis of semantics and sentics in natural language text, namely:

1. **Affective Blending:** a process for reasoning by analogy on affective common-sense knowledge;
2. **Affective Categorisation:** a biologically-inspired and psychologically-motivated model for the representation and the analysis of human emotions;
3. **Sentic Medoids:** a novel clustering technique for organising affective common-sense concepts in AffectiveSpace;
4. **Sentic Activation:** a bio-inspired two-level framework that exploits an ensemble application of dimensionality reduction and graph mining techniques;
5. **Sentic Panalogy:** an approach that models the switch between different reasoning strategies and between the foci around which such strategies are developed.

6.1.2 Tools

The above-mentioned techniques have been employed within this research work for the design of a set of tools for the automatic analysis of opinions and sentiments:

1. **SenticNet:** a publicly available semantic resource for opinion mining built using both AI and Semantic Web techniques;
2. **Sentic Neural Networks:** an MLP-based architecture for a more brain-inspired affective common-sense recognition;

3. **Open Mind Common Sentic**: an emotion-sensitive IUI for collecting affective common-sense knowledge from general public;
4. **IsaCore**: a semantic network of common and common-sense knowledge for auto-categorisation built upon ConceptNet and Probase;
5. **Opinion Mining Engine**: an intelligent engine for concept-level open-domain opinion mining and sentiment analysis.

6.1.3 Applications

Finally, sentic computing tools and techniques have been exploited for the development of emotion-sensitive applications in fields such as Social Web, HCI, and e-health:

1. **Troll Filter**: a system for automatically filtering inflammatory and outrageous posts within online communities;
2. **Social Media Marketing Tool**: an intelligent web application for managing social media information about products and services through a faceted interface;
3. **Sentic Album**: a content, concept, and context based online personal photo management system;
4. **Sentic Avatar**: an emotion-sensitive avatar built through the integration of sentic computing with a facial emotional classifier;
5. **Sentic Chat**: an IM platform that enriches social communication through semantics and sentics;
6. **Sentic Corner**: an IUI that dynamically collects audio, video, images, and text related to the user's emotions and motions;
7. **Crowd Validation**: a process for mining patient opinions and bridging the gap between unstructured and structured health-care data;
8. **Sentic PROMs**: a new framework for measuring health care quality that exploits the ensemble application of standard PROMs and sentic computing.

6.2 Limitations and Future Work

The research work carried out in the past three years has put solid bases for the development of a variety of emotion-sensitive systems and applications in the fields of opinion mining and sentiment analysis. One of the main contributions of this book has also been the introduction of a novel approach to the analysis of opinions and sentiments, which goes beyond merely keyword-based methods by using common-sense reasoning tools and affective ontologies. The developed techniques, however, are still far from perfect as the common and common-sense knowledge base need to be further expanded and the reasoning tools built on the top of them adjusted accordingly. This last section discusses the limitations of such techniques (section 6.2.1) and their further development (section 6.2.2).

6.2.1 *Limitations*

As discussed in section 4.5.2, the validity of the proposed approach mainly depends on the richness of the developed knowledge bases. Without a comprehensive resource that encompasses human knowledge, in fact, it is not easy for an opinion mining system to get a hold of the ability to grasp the cognitive and affective information associated with natural language text and, hence, accordingly aggregate opinions in order to make statistics on them. Attempts to encode human common knowledge are countless and comprehend both resources generated by human experts (or community efforts) and automatically-built knowledge bases. The former kinds of resources are generally too limited, as they need to be hand-crafted, the latter too noisy, as they mainly rely on information available on the Web.

The span and the accuracy of knowledge available, however, is not the only limitation of opinion mining systems. Even though a machine “knows 50 million such things”¹, it needs to be able to accordingly exploit such knowledge through different types of associations, e.g., inferential, causal, analogical, deductive, or inductive. For the purposes of this work, singular value decomposition (SVD) appeared to be a good method for generalising the information contained in the common-sense knowledge bases, but it is very expensive in both computing time and storage, as it requires costly arithmetic operations such as division and square root in the computation of rotation parameters. This is a big issue because both AffectNet and IsaCore are keeping on growing, in parallel with the continuously extended versions of ConceptNet and Probase. Moreover, the eigenmoods of AffectiveSpace cannot be easily understood because they are linear combinations of all of the original concept features. Different strategies that clearly show various steps of reasoning might be preferable in the future.

Another limitation of the sentic computing approach is in its typicality. The clearly defined knowledge representation of AffectNet and IsaCore, in fact, does not allow to grasp different concept nuances as the inference of semantic and affective features associated with concepts is bounded. New features associated to a concept can be indeed inferred through the AffectiveSpace process, but the number of new features that can be discovered after reconstructing the concept-feature matrix is limited to the set of features associated with semantically related concepts (that is, concepts that share similar features). However, depending on the context, concepts might need to be associated with features that are not strictly pertinent to germane concepts.

The concept of book, for example, is typically associated to concepts such as newspaper or magazine, as it contains knowledge, has pages, etc. In a different context, however, a book could be used as paperweight, doorstop, or even as a weapon. Biased (context-dependent) association of concepts is possible through spectral association, in which spreading activation is concurrently determined by different nodes in the graph representation of IsaCore. Because concepts are hereby considered atomic and mono-faceted, however, it is not easy for the system to grasp

¹ <http://mitworld.mit.edu/video/484>

the many different ways a concept can be meaningful in a particular context, as the features associated with each concept identify just its typical qualities, traits, or characteristics. Finally, another limitation of the proposed approach is in the lack of time representation. Such an issue is not addressed by any of the currently available knowledge bases, including ConceptNet and Probase, upon which AffectNet and IsaCore are built. In the context of sentic computing, however, time representation is not specifically needed as the main aim of the opinion mining engine is the passage from unstructured natural language data to structured machine-processable information, rather than genuine natural language understanding. Every SBoC, in fact, is treated as independent from other SBoCs in the text data, as the goal is to simply infer a topic and a polarity associated with it, rather than understanding the whole meaning of the sentence in correlation with adjacent ones. In some cases, however, time representation might be needed for tasks such as comparative opinion analysis and co-reference resolution.

6.2.2 Future Work

In order to overcome some of the above-mentioned limitations, the current research work is focusing on expanding AffectNet and IsaCore with different kinds of knowledge (e.g., common-sense, affective knowledge, common knowledge) coming from external resources, e.g., Cyc, Freebase, and Yago. Such operation is not simply convenient for improving the accuracy of the opinion mining engine, but also for reducing the sparseness of the matrix representations of such knowledge bases and, hence, aid dimensionality reduction procedures.

In order to overcome the current weaknesses of the sentic parser, moreover, a new parsing tool is being developed. The parser is based on a construction grammar approach, that is, a constraint-based, generative, and mono-stratal grammatical model, committed to incorporating the cognitive and interactional foundations of language. Construction grammar is also inherently tied to a particular model of the ‘semantics of understanding’, known as frame semantics, which offers a way of structuring and representing meaning while taking into account the relationship between lexical meaning and grammatical patterning.

New graph-mining and multi-dimensionality reduction techniques are also being explored to perform reasoning on the common-sense knowledge bases. In particular, it is being investigated how AffectiveSpace can be built by means of independent component analysis (ICA) and random projections. Moreover, new classification techniques such as support and relevance vector machines are being experimented, together with the ensemble application of dimensionality reduction and extreme learning machine (ELM) techniques for emulating fast affective learning and reasoning.

6.3 Conclusions

The textual information available on the Web can be broadly grouped into two main categories: facts and opinions. Facts are objective expressions about entities or events. Opinions are usually subjective expressions that describe people's sentiments, appraisals, or feelings towards such entities and events. Much of the existing research on textual information processing has been focused on mining and retrieval of factual information, e.g., text classification [271], text recognition [294], text clustering [171], and many other text mining and NLP tasks. Little work had been done on the processing of opinions until only recently [179].

One of the main reasons for the lack of study on opinions is the fact that there was little opinionated text available before the recent passage from a read-only to a read-write Web. Before that, in fact, when people needed to make a decision, they typically asked for opinions from friends and family. Similarly, when organisations wanted to find the opinions or sentiments of the general public about their products and services, they had to specifically ask people by conducting opinion polls and surveys. However, with the advent of the Social Web, the way people express their views and opinions has dramatically changed. They can now post reviews of products at merchant sites and express their views on almost anything in Internet forums, discussion groups, and blogs. Such online word-of-mouth behaviour represents new and measurable sources of information with many practical applications. However, finding opinion sources and monitoring them can be a formidable task because there are a large number of diverse sources and each source may also have a huge volume of opinionated text. In many cases, in fact, opinions are hidden in long forum posts and blogs. It is extremely time consuming for a human reader to find relevant sources, extract related sentences with opinions, read them, summarise them, and organise them into usable forms. Thus, automated opinion discovery and summarisation systems are needed.

Sentiment analysis grows out of this need. It is a very challenging NLP or text mining problem. Due to its tremendous value for practical applications, there has been an explosive growth of both research in academia and applications in the industry. All the sentiment analysis tasks, however, are very challenging. Our understanding and knowledge of the problem and its solution are still limited. The main reason is that it is a NLP task, and NLP has no easy problems. Another reason may be due to our popular ways of doing research. So far, in fact, researchers have probably relied too much on machine learning algorithms. Some of the most effective machine learning algorithms, in fact, produce no human understandable results such that, although they may achieve improved accuracy, little about how and why is known, apart from some superficial knowledge gained in the manual feature engineering process. All such approaches, moreover, rely on syntactical structure of text, which is far from the way human mind processes natural language.

In this book, common-sense computing techniques were further developed and applied to bridge the semantic gap between word-level natural language data and the concept-level opinions conveyed by these. In particular, the ensemble application of graph mining and multi-dimensionality reduction techniques was exploited on

two common-sense knowledge bases to develop a novel intelligent engine for open-domain opinion mining and sentiment analysis. The proposed approach, termed *sentic computing*, performs a clause-level semantic analysis of text, which allows the inference of both the conceptual and emotional information associated with natural language opinions and, hence, a more efficient passage from (unstructured) textual information to (structured) machine-processable data.

Blending scientific theories of emotion with the practical engineering goals of analysing sentiments in natural language text and developing affect-sensitive interfaces is one of the main contributions of this book. Differently from most currently available opinion mining services, in fact, the developed engine does not base its analysis on a limited set of affect words and their co-occurrence frequencies, but rather on common-sense concepts and the cognitive and affective valence conveyed by these. This allows the engine to be domain-independent and, hence, to be embedded in any sentiment analysis system for the development of intelligent applications in multiple fields such as Social Web, HCI, and e-health.

Looking ahead, the combined novel use of different knowledge bases and of affective common-sense reasoning techniques for opinion mining proposed in this work will, eventually, pave the way for development of more bio-inspired approaches to the design of intelligent systems capable of handling knowledge, making analogies, learning from experience, perceiving, and expressing affect. The question, in fact, is not whether intelligent machines can have emotions, but whether machines can be intelligent without any emotions [199].

References

1. Abbasi, A.: Affect intensity analysis of dark web forums. In: IEEE ISI, pp. 282–288. New Brunswick (2007)
2. Abbasi, A., Chen, H., Salem, A.: Sentiment analysis in multiple languages: Feature selection for opinion classification in web forums. *ACM Transactions on Information Systems* **26**(3), 1–34 (2008)
3. Addis, M., Boch, L., Allasia, W., Gallo, F., Bailer, W., Wright, R.: 100 million hours of audiovisual content: Digital preservation and access in the prestoprime project. In: Digital Preservation Interoperability Framework Symposium. Dresden (2010)
4. Agrawal, R., Srikant, R.: Fast algorithm for mining association rules. In: VLDB. Santiago de Chile (1994)
5. von Ahn, L.: Games with a purpose. *IEEE Computer Magazine* **6**, 92–94 (2006)
6. von Ahn, L., Blum, M., Hopper, N., Langford, J.: Captcha: Using hard ai problems for security. In: E. Biham (ed.) EUROCRYPT, Lecture Notes in Computer Science, pp. 294–311. Springer, Berlin Heidelberg (2003)
7. von Ahn, L., Dabbish, L.: Labeling images with a computer game. In: CHI, pp. 319–326. Vienna (2004)
8. von Ahn, L., Ginosar, S., Kedia, M., Liu, R., Blum, M.: Improving accessibility of the web with a computer game. In: CHI, pp. 79–82. Quebec (2006)
9. von Ahn, L., Kedia, M., Blum, M.: Verbosity: A game for collecting common sense facts. In: CHI, pp. 75–78. Quebec (2006)
10. von Ahn, L., Liu, R., Blum, M.: Peekaboom: A game for locating objects in images. In: CHI, pp. 55–64 (2006)
11. Alm, C., Roth, D., Sproat, R.: Emotions from text: Machine learning for text-based emotion prediction. In: HLT/EMNLP, pp. 347–354 (2005)
12. Ambady, N., Rosenthal, R.: Thin slices of expressive behavior as predictors of interpersonal consequences: a meta-analysis. *Psychological Bulletin* **11**(2), 256–274 (1992)
13. Argamon, S., Koppel, M., Avneri, G.: Style-based text categorization: What newspaper am i reading? In: AAAI Workshop on Text Categorization. Madison (1998)
14. Averill, J.: A constructivist view of emotion. *Emotion: Theory, Research and Experience* pp. 305–339 (1980)
15. Bach, J., Fuller, C., Gupta, A., Hampapur, A., Horowitz, B., Humphrey, R., Jain, R., Shu, C.: Virage image search engine: An open framework for image management. In: I. Sethi, R. Jain (eds.) Storage and Retrieval for Still Image and Video Databases, vol. 2670, pp. 76–87 (1996)
16. Baker, C., Fillmore, C., Lowe, J.: The Berkeley FrameNet project. In: COLING/ACL, pp. 86–90. Montreal (1998)
17. Baldassarri, S., Cerezo, E., Seron, F.: Maxine: a platform for embodied animated agents. *Computers and Graphics* **32**(4), 430–437 (2008)
18. Barrett, L.: Solving the emotion paradox: Categorization and the experience of emotion. *Personality and Social Psychology Review* **10**(1), 20–46 (2006)
19. Barrington, L., O'Malley, D., Turnbull, D., Lanckriet, G.: User-centered design of a social game to tag music. In: ACM SIGKDD, pp. 7–10. Paris (2009)
20. Barwise, J.: An introduction to first-order logic. In: Handbook of Mathematical Logic. Studies in Logic and the Foundations of Mathematics. North-Holland (1977)
21. Benson, T., Sizmur, S., Whatling, J., Arikan, S., McDonald, D., Ingram, D.: Evaluation of a new short generic measure of health status. *Informatics in Primary Care* **18**(2), 89–101 (2010)
22. Bergner, M., Bobbitt, R., Kressel, S., Pollard, W., Gilson, B., Morris, J.: The sickness impact profile: Conceptual formulation and methodology for the development of a health status measure. *International Journal of Health Services* **6**, 393–415 (1976)
23. Bianchi-Berthouze, N.: K-DIME: An affective image filtering system. *IEEE Multimedia* **10**(3), 103–106 (2003)
24. Biber, D.: Variation across Speech and Writing. Cambridge University Press (1988)

25. Bloom, P.: Glue for the mental world. *Nature* **421**, 212–213 (2003)
26. Boehner, K., Depaula, R., Dourish, P., Sengers, P.: Affect: From information to interaction. In: *Critical Computing*, pp. 59–68. Aarhus (2005)
27. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: A collaboratively created graph database for structuring human knowledge. In: *SIGMOD*, pp. 1247–1250. Vancouver (2008)
28. Bonanno, G., Papa, A., O’Neill, K., Westphal, M., Coifman, K.: The importance of being flexible: The ability to enhance and suppress emotional expressions predicts long-term adjustment. *Psychological Science* **15**, 482–487 (2004)
29. Bradford Cannon, W.: *Bodily Changes in Pain, Hunger, Fear and Rage: An Account of Recent Researches into the Function of Emotional Excitement*. Appleton Century Crofts (1915)
30. Broca, P.: Anatomie comparée des circonvolutions cérébrales: Le grand lobe limbique. *Rev. Anthropol* **1**, 385–498 (1878)
31. Brooks, R.: EuroQoL - the current state of play. *Health Policy* **37**, 53–72 (1996)
32. Burke, A., Heuer, F., Reisberg, D.: Remembering emotional events. *Memory and Cognition* **20**, 277–290 (1992)
33. Burleson, W., Picard, R., Perlin, K., Lippincott, J.: A platform for affective agent research. In: *AAMAS*. New York (2004)
34. Cahill, L., McGaugh, J.: A novel demonstration of enhanced memory associated with emotional arousal. *Consciousness and Cognition* **4**(4), 410–421 (1995)
35. Calvo, M., Nummenmaa, L.: Processing of unattended emotional visual scenes. *Journal of Experimental Psychology: General* **136**, 347–369 (2007)
36. Calvo, R., D’Mello, S.: Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing* **1**(1), 18–37 (2010)
37. Cambria, E., Benson, T., Eckl, C., Hussain, A.: Sentic PROMs: Application of sentic computing to the development of a novel unified framework for measuring health-care quality. *Expert Systems with Applications* **39**(12), 10,533–10,543 (2012)
38. Cambria, E., Chandra, P., Sharma, A., Hussain, A.: Do not feel the trolls. In: *ISWC*. Shanghai (2010)
39. Cambria, E., Grassi, M., Hussain, A., Havasi, C.: Sentic computing for social media marketing. *Multimedia Tools and Applications* **59**(2), 557–577 (2012)
40. Cambria, E., Havasi, C., Hussain, A.: SenticNet 2: A semantic and affective resource for opinion mining and sentiment analysis. In: *FLAIRS*, pp. 202–207. Marco Island (2012)
41. Cambria, E., Howard, N., Hsu, J., Hussain, A.: Sentic blending: Scalable multimodal fusion for continuous interpretation of semantics and sentics. In: *IEEE SSCI*, pp. 108–117. Singapore (2013)
42. Cambria, E., Hupont, I., Hussain, A., Cerezo, E., Baldassarri, S.: Sentic avatar: Multimodal affective conversational agent with common sense. In: A. Esposito, A. Hussain, M. Faundez-Zanuy, R. Martone, N. Melone (eds.) *Toward Autonomous, Adaptive, and Context-Aware Multimodal Interfaces: Theoretical and Practical Issues*, *Lecture Notes in Computer Science*, vol. 6456, pp. 82–96. Springer-Verlag, Berlin Heidelberg (2011)
43. Cambria, E., Hussain, A.: Sentic album: Content-, concept-, and context-based online personal photo management system. *Cognitive Computation* **4**(4), 477–496 (2012)
44. Cambria, E., Hussain, A.: *Sentic Computing: Techniques, Tools, and Applications*. Springer, Dordrecht, Netherlands (2012)
45. Cambria, E., Hussain, A., Durrani, T., Havasi, C., Eckl, C., Munro, J.: Sentic computing for patient centered application. In: *IEEE ICSP*, pp. 1279–1282. Beijing (2010)
46. Cambria, E., Hussain, A., Durrani, T., Zhang, J.: Towards a chinese common and common sense knowledge base for sentiment analysis. In: H. Jiang, W. Ding, M. Ali, X. Wu (eds.) *Advanced Research in Applied Artificial Intelligence*, *Lecture Notes in Computer Science*, vol. 7345, pp. 437–446. Springer-Verlag, Berlin Heidelberg (2012)
47. Cambria, E., Hussain, A., Eckl, C.: Taking refuge in your personal sentic corner. In: *IJCNLP*, pp. 35–43. Chiang Mai (2011)
48. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: AffectiveSpace: Blending common sense and affective knowledge to perform emotive reasoning. In: *CAEPIA*, pp. 32–41. Seville (2009)

49. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: Common sense computing: From the society of mind to digital intuition and beyond. In: J. Fierrez, J. Ortega, A. Esposito, A. Drygajlo, M. Faundez-Zanuy (eds.) Biometric ID Management and Multimodal Communication, *Lecture Notes in Computer Science*, vol. 5707, pp. 252–259. Springer, Berlin Heidelberg (2009)
50. Cambria, E., Hussain, A., Havasi, C., Eckl, C.: SenticSpace: Visualizing opinions and sentiments in a multi-dimensional vector space. In: R. Setchi, I. Jordanov, R. Howlett, L. Jain (eds.) Knowledge-Based and Intelligent Information and Engineering Systems, *Lecture Notes in Artificial Intelligence*, vol. 6279, pp. 385–393. Springer, Berlin (2010)
51. Cambria, E., Hussain, A., Havasi, C., Eckl, C., Munro, J.: Towards crowd validation of the uk national health service. In: WebSci. Raleigh (2010)
52. Cambria, E., Livingstone, A., Hussain, A.: The hourglass of emotions. In: A. Esposito, A. Vinciarelli, R. Hoffmann, V. Muller (eds.) Cognitive Behavioral Systems, *Lecture Notes in Computer Science*, vol. 7403, pp. 144–157. Springer, Berlin Heidelberg (2012)
53. Cambria, E., Mazzocco, T., Hussain, A.: Application of multi-dimensional scaling and artificial neural networks for biologically inspired opinion mining. *Biologically Inspired Cognitive Architectures* 4, 41–53 (2013)
54. Cambria, E., Mazzocco, T., Hussain, A., Eckl, C.: Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space. In: D. Liu, H. Zhang, M. Polycarpou, C. Alippi, H. He (eds.) Advances in Neural Networks, *Lecture Notes in Computer Science*, vol. 6677, pp. 601–610. Springer-Verlag, Berlin (2011)
55. Cambria, E., Olsher, D., Kwok, K.: Sentic activation: A two-level affective common sense reasoning framework. In: AAAI, pp. 186–192. Toronto (2012)
56. Cambria, E., Olsher, D., Kwok, K.: Sentic panalogy: Swapping affective common sense reasoning strategies and foci. In: CogSci, pp. 174–179. Sapporo (2012)
57. Cambria, E., Rajagopal, D., Olsher, D., Das, D.: Big social data analysis. In: R. Akerkar (ed.) Big Data Computing, chap. 13. Chapman and Hall/CRC (2013)
58. Cambria, E., Schuller, B., Xia, Y., Havasi, C.: New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems* (2013). DOI 10.1109/MIS.2013.30
59. Cambria, E., Song, Y., Wang, H., Howard, N.: Semantic multi-dimensional scaling for open-domain sentiment analysis. *IEEE Intelligent Systems* (2013). DOI 10.1109/MIS.2012.118
60. Cambria, E., Speer, R., Havasi, C., Hussain, A.: SenticNet: A publicly available semantic resource for opinion mining. In: AAAI CSK, pp. 14–18. Arlington (2010)
61. Cambria, E., Xia, Y., Hussain, A.: Affective common sense knowledge acquisition for sentiment analysis. In: LREC, pp. 3580–3585. Istanbul (2012)
62. Camurri, A., Mazarino, B., Volpe, G.: Analysis of expressive gesture: The eyesweb expressive gesture processing library. In: Gesture Workshop. Genova (2003)
63. Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Hruschka, E., Mitchell, T.: Toward an architecture for never-ending language learning. In: AAAI, pp. 1306–1313. Atlanta (2010)
64. Castellano, G., Kessous, L., Caridakis, G.: Multimodal emotion recognition from expressive faces, body gestures and speech. In: Doctoral Consortium of ACII. Lisbon (2007)
65. Cerezo, E., Hupont, I., Manresa, C., Varona, J., Baldassarri, S., Perales, F., Seron, F.: Real-time facial expression recognition for natural interaction. In: Pattern Recognition and Image Analysis, *Lecture Notes in Computer Science*, vol. 4478, pp. 40–47. Springer-Verlag, Berlin Heidelberg (2007)
66. Cesarano, C., Dorr, B., Picariello, A., Reforgiato, D., Sagoff, A., Subrahmanian, V.: OASYS: An opinion analysis system. In: AAAI CAAW. Stanford (2006)
67. Chaiken, S., Trope, Y.: Dual-Process Theories in Social Psychology. New York: Guilford (1999)
68. Chandra, P., Cambria, E., Hussain, A.: Clustering social networks using interaction semantics and sentics. In: J. Wang, G. Yen, M. Polycarpou (eds.) Advances in Neural Networks, *Lecture Notes in Computer Science*, vol. 7367, pp. 379–385. Springer, Heidelberg (2012)
69. Chandra, P., Cambria, E., Pradeep, A.: Enriching social communication through semantics and sentics. In: IJCNLP, pp. 68–72. Chiang Mai (2011)
70. Chang, H.: Emotion barometer of reading: User interface design of a social cataloging website. In: International Conference on Human Factors in Computing Systems (2009)

71. Charles, D.: *The Expression of the Emotions in Man and Animals*. John Murray (1872)
72. Chi, P., Lieberman, H.: Intelligent assistance for conversational storytelling using story patterns. In: IUI. Palo Alto (2011)
73. Chklovski, T.: Learner: a system for acquiring commonsense knowledge by analogy. In: K-CAP (2003)
74. Chomsky, N.: Three models for the description of language. *IRE Transactions on Information Theory* **2**(3), 113–124 (1956)
75. Christianson, S., Loftus, E.: Remembering emotional events: The fate of detailed information. *Cognition and Emotion* **5**, 81–108 (1991)
76. Cochrane, T.: Eight dimensions for the emotions. *Social Science Information* **48**(3), 379–420 (2009)
77. Codd, E.: A relational model of data for large shared data banks. *Communications of the ACM* **13**(6), 377–387 (1970)
78. Codd, E.: Further normalization of the data base relational model. Tech. rep., IBM Research Report, New York (1971)
79. Codd, E.: Recent investigations into relational data base systems. Tech. Rep. RJ1385, IBM Research Report, New York (1974)
80. Csikszentmihalyi, M.: *Flow: The Psychology of Optimal Experience*. Harper Perennial (1991)
81. Culyer, A., Lavers, R., Williams, A.: Social indicators: Health. *Social Trends* **2**, 31–42 (1971)
82. Dalgleish, T.: The emotional brain. *Nature: Perspectives* **5**, 582–589 (2004)
83. Dalgleish, T., Dunn, B., Mobbs, D.: Affective neuroscience: Past, present, and future. *Emotion Review* pp. 355–368 (2009)
84. Damasio, A.: *Descartes' Error: Emotion, Reason, and the Human Brain*. Grossett/Putnam, New York (1994)
85. Damasio, A.: *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. Harcourt, Inc. (2003)
86. Danisman, T., Alpkocak, A.: Feeler: Emotion classification of text using vector space model. In: AISB. Aberdeen (2008)
87. Das, S., Chen, M.: Yahoo! for Amazon: Extracting market sentiment from stock message boards. In: APFA. Bangkok (2001)
88. Date, C., Darwen, H.: *A Guide to the SQL Standard*. Addison-Wesley, Reading, Mass (1993)
89. Datta, R., Wang, J.: ACQUINE: Aesthetic quality inference engine - real-time automatic rating of photo aesthetics. In: International Conference on Multimedia Information Retrieval. Philadelphia (2010)
90. Dave, K., Lawrence, S., Pennock, D.: Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In: WWW. Budapest (2003)
91. Decherchi, S., Gastaldo, P., Zunino, R., Cambria, E., Redi, J.: Circular-ELM for the reduced-reference assessment of perceived image quality. *Neurocomputing* **102**, 78–89 (2013)
92. Di Fabbriozio, G., Aker, A., Gaizauskas, R.: Starlet: Multi-document summarization of service and product reviews with balanced rating distributions. In: ICDM SENTIRE, pp. 67–74. Vancouver (2011)
93. D'Mello, S., Craig, S., Sullins, J., Graesser, A.: Predicting affective states expressed through an emote-aloud procedure from autotutor's mixed-initiative dialogue. *Int'l J. Artificial Intelligence in Education* **16**, 3–28 (2006)
94. D'Mello, S., Dowell, N., Graesser, A.: Cohesion relationships in tutorial dialogue as predictors of affective states. In: Conf. Artificial Intelligence in Education, pp. 9–16 (2009)
95. Donabedian, A.: Evaluating the quality of medical care. *The Millbank Memorial Fund Quarterly* **44**, 166–203 (1966)
96. Douglas-Cowie, E.: *Humaine deliverable d5g: Mid term report on database exemplar progress*. Tech. rep., Information Society Technologies (2006)
97. Duda, R., Hart, P.: *Pattern Classification and Scene Analysis*. John Wiley & Sons Inc, New York (1973)
98. Eckart, C., Young, G.: The approximation of one matrix by another of lower rank. *Psychometrika* **1**(3), 211–218 (1936)

99. Ekman, P., Dalgleish, T., Power, M.: *Handbook of Cognition and Emotion*. Wiley, Chichester (1999)
100. Elliott, C.D.: *The affective reasoner: A process model of emotions in a multi-agent system*. Ph.D. thesis, Northwestern University, Evanston (1992)
101. Ephron, H.: *1001 Books for Every Mood: A Bibliophile's Guide to Unwinding, Misbehaving, Forgiving, Celebrating, Commiserating*. Adams Media, Avon (2008)
102. Epstein, S.: *Cognitive-experiential self-theory of personality*. In: T. Millon, M. Lerner (eds.) *Comprehensive Handbook of Psychology*, vol. 5, pp. 159–184. Wiley & Sons, Hoboken, NJ (2003)
103. Ernest, D.: *Representations of Commonsense Knowledge*. Morgan Kaufmann (1990)
104. Esuli, A., Sebastiani, F.: *SentiWordNet: A publicly available lexical resource for opinion mining*. In: LREC (2006)
105. Fanshel, S., Bush, J.: *A health status index and its application to health-services outcomes*. *Operations Research* **18**, 1021–1066 (1970)
106. Fauconnier, G., Turner, M.: *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books (2003)
107. Fayyad, U., Piatetsky, G., Smyth, P.: *From data mining to knowledge discovery: An overview*. In: *Advances in Knowledge Discovery and Data Mining*, pp. 1–36. MIT Press, Cambridge (1996)
108. Fellbaum, C.: *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press (1998)
109. Finn, A., Kushmerick, N.: *Learning to classify documents according to genre*. *Journal of the American Society for Information Science and Technology* **7**(5) (2006)
110. Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D., Yanker, P.: *Query by image and video content: the QBIC system*. *Computer* **28**(9), 23–32 (1995)
111. Fontaine, J., Scherer, K., Roesch, E., Ellsworth, P.: *The world of emotions is not two-dimensional*. *Psychological Science* **18**(12), 1050–1057 (2007)
112. Frankel, C., Swain, M.J., Athitsos, V.: *WebSeer: An image search engine for the world wide web*. Tech. rep., University of Chicago (1996)
113. Freitas, A., Castro, E.: *Facial expression: The effect of the smile in the treatment of depression. empirical study with portuguese subjects*. In: *Emotional Expression: The Brain and The Face*, pp. 127–140. University Fernando Pessoa Press (2009)
114. Frijda, N.: *The laws of emotions*. *American Psychologist* **43**(5) (1988)
115. Garey, M., Johnson, D.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman (1979)
116. Gilroy, S., Cavazza, M., Niiranen, M., Andre, E., Vogt, T., Urbain, J., Benayoun, M., Seichter, H., Billinghurst, M.: *Pad-based multimodal affective fusion*. In: *ACII*, pp. 1–8. Amsterdam (2009)
117. Goertzel, B., Silverman, K., Hartley, C., Bugaj, S., Ross, M.: *The Baby Webmind project*. In: *AISB*. Birmingham (2000)
118. Grassi, M.: *Developing HEO human emotions ontology*. In: J. Fierrez, J. Ortega-Garcia, A. Esposito, A. Drygajlo, M. Faundez-Zanuy (eds.) *Biometric ID Management and Multimodal Communication, Lecture Notes in Computer Science*, vol. 5707, pp. 244–251. Springer, Berlin Heidelberg (2009)
119. Grassi, M., Cambria, E., Hussain, A., Piazza, F.: *Sentic web: A new paradigm for managing social media affective information*. *Cognitive Computation* **3**(3), 480–489 (2011)
120. Grice, P.: *Logic and Conversation*, vol. *Speech Acts*, pp. 41–58. Academic Press (1975)
121. Gruber, T.: *Collective knowledge systems: Where the social web meets the semantic web*. *Web Semantics: Science, Services and Agents on the World Wide Web* **6**(1), 4–13 (2007)
122. Gunes, H., Piccardi, M.: *Bi-modal emotion recognition from expressive face and body gestures*. *Network and Computer Applications* **30**(4), 1334–1345 (2007)
123. Gunes, H., Piccardi, M., Pantic, M.: *From the lab to the real world: Affect recognition using multiple cues and modalities*. *Affective Computing: Focus on Emotion Expression, Synthesis, and Recognition* pp. 185–218 (2008)

124. Gupta, R., Kochenderfer, M., McGuinness, D., Ferguson, G.: Common sense data acquisition for indoor mobile robots. In: AAAI, pp. 605–610. San Jose (2004)
125. Hacker, S., von Ahn, L.: Matchin: Eliciting user preferences with an online game. In: CHI. Boston (2009)
126. Han, J., Kamber, M., Tung, A.: Spatial clustering methods in data mining: A survey. In: T. Francis (ed.) *Geographic Data Mining and Knowledge Discovery* (2001)
127. Hanjalic, A.: Extracting moods from pictures and sounds: Towards truly personalized TV. *IEEE Signal Processing Magazine* **23**(2), 90–100 (2006)
128. Hartigan, J., Wong, M.: Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society* **28**(1), 100–108 (1979)
129. Hatzivassiloglou, V., McKeown, K.: Predicting the semantic orientation of adjectives. In: *ACL/EACL*. Madrid (1997)
130. Hatzivassiloglou, V., Wiebe, J.: Effects of adjective orientation and gradability on sentence subjectivity. In: *COLING*. Saarbrücken (2000)
131. Havasi, C., Speer, R., Alonso, J.: ConceptNet 3: A flexible, multilingual semantic network for common sense knowledge. In: *RANLP*. Borovets (2007)
132. Havasi, C., Speer, R., Holmgren, J.: Automated color selection using semantic knowledge. In: *AAAI CSK*. Arlington (2010)
133. Havasi, C., Speer, R., Pustejovsky, J.: Coarse word-sense disambiguation using common sense. In: *AAAI CSK*, pp. 46–51. Arlington (2010)
134. Havasi, C., Speer, R., Pustejovsky, J., Lieberman, H.: Digital intuition: Applying common sense using dimensionality reduction. *IEEE Intelligent Systems* **24**(4), 24–35 (2009)
135. Hearst, M.: Automatic acquisition of hyponyms from large text corpora. In: *COLING*, pp. 539–545. Nantes (1992)
136. Hearst, M.: Text data mining: Issues, techniques, and the relationship to information access. In: *UW/MS Workshop on Data Mining*. Berkeley (1997)
137. Herdagdelen, A., Baroni, M.: The concept game: Better commonsense knowledge extraction by combining text mining and game with a purpose. In: *AAAI CSK*. Arlington (2010)
138. Heyting, A.: Intuitionism. An introduction. North-Holland (1956)
139. Horsman, J., Furlong, W., Feeny, D., Torrance, G.: The health utility index (HUI): Concepts, measurement, properties and applications. *Health and Quality of Life Outcomes* **1**(54) (2003)
140. Hsinchun, C., Zimbra, D.: AI and opinion mining. *IEEE Intelligent Systems* **25**(3), 74–80 (2010)
141. Hu, M., Liu, B.: Mining and summarizing customer reviews. In: *KDD*. Seattle (2004)
142. Huang, J., Ravi, S., Mitra, M., Zhu, W., Zabih, R.: Image indexing using color correlograms. In: *IEEE CVPR*, pp. 762–768 (1997)
143. James, W.: What is an emotion? *Mind* **34**, 188–205 (1884)
144. Jin, X., Li, Y., Mah, T., Tong, J.: Sensitive webpage classification for content advertising. In: *ACM SIGKDD*, pp. 28–33. San Jose (2007)
145. Jing, F., Wang, C., Yao, Y., Deng, K., Zhang, L., Ma, W.Y.: IGroup: Web image search results clustering. In: *ACM Multimedia*. Santa Barbara (2006)
146. JongTae, J., SangWook, S., KwangEun, K., KweeBo, S.: Emotion recognition method based on multimodal sensor fusion algorithm. In: *ISIS*. Sokcho-City (2007)
147. Joshi, M., Rose, C.: Generalizing dependency features for opinion mining. In: *ACL/IJCNLP*. Singapore (2009)
148. Kamps, J., Marx, M., Mokken, R., de Rijke, M.: Using WordNet to measure semantic orientation of adjectives. In: *LREC*, pp. 1115–1118. Lisbon (2004)
149. Kapoor, A., Burleson, W., Picard, R.: Automatic prediction of frustration. *International Journal of Human-Computer Studies* **65**, 724–736 (2007)
150. Karpouzis, K., Caridakis, G., Kessous, L., Amir, N., Raouzaoui, A., Malatesta, L., Kollias, S.: Modeling naturalistic affective states via facial, vocal and bodily expressions recognition. In: *Lecture Notes in Artificial Intelligence*, vol. 4451, pp. 92–116. Springer (2007)
151. Kaufman, L., Rousseeuw, P.: *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley-Interscience (1990)

152. Keelan, B.: *Handbook of Image Quality*. Marcel Dekker, New York (2002)
153. Khan, K., Baharudin, B., Khan, A., Malik, F.: Mining opinion from text documents: A survey. In: *IEEE DEST*, pp. 217–222. Istanbul (2009)
154. Kim, S., Hovy, E.: Automatic detection of opinion bearing words and sentences. In: *IJCNLP*, pp. 61–66. Jeju Island (2005)
155. Kim, S., Hovy, E.: Extracting opinions, opinion holders, and topics expressed in online news media text. In: *Workshop on Sentiment and Subjectivity in Text*. Sydney (2006)
156. Kirkpatrick, L., Epstein, S.: Cognitive experiential self-theory and subjective probability: Further evidence for two conceptual systems. *Journal of Personality and Social Psychology* **63**, 534–544 (1992)
157. Krumhuber, E., Kappas, A.: Moving smiles: The role of dynamic components for the perception of the genuineness of smiles. *Journal of Nonverbal Behavior* **29**(1), 3–24 (2005)
158. Kumar, P., Yildirim, E.: Minimum-volume enclosing ellipsoids and core sets. *Journal of Optimization Theory and Applications* **126**, 1–21 (2005)
159. Kuo, Y., Lee, J., Chiang, K., Wang, R., Shen, E., Chan, C., Hu, J.Y.: Community-based game design: Experiments on social games for commonsense data collection. In: *ACM SIGKDD*, pp. 15–22. Paris (2009)
160. Lacy, L.: *OWL: Representing Information Using the Web Ontology Language*. Trafford Publishing (2005)
161. Lakoff, G.: *Women, Fire, and Dangerous Things*. University Of Chicago Press (1990)
162. Lanczos, C.: An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of The National Bureau of Standards* **45**(4), 255–282 (1950)
163. Laney, C., Campbell, H., Heuer, F., Reisberg, D.: Memory for thematically arousing events. *Memory and Cognition* **32**(7), 1149–1159 (2004)
164. Lansdale, M., Edmonds, E.: Using memory for events in the design of personal filing systems. *International Journal of Man-Machine Studies* **36**(1), 97–126 (1992)
165. Law, E., von Ahn, L., Dannenberg, R., Crawford, M.: Tagatune: A game for music and sound annotation. In: *International Conference on Music Information Retrieval*. Vienna (2007)
166. Lazarus, R.: *Emotion and Adaptation*. Oxford University Press, New York (1991)
167. Ledoux, J.: *Synaptic Self*. Penguin Books (2003)
168. Lee, B., Hendler, J., Lassila, O.: The semantic web. *Scientific American* (2001)
169. Lempel, R., Soffer, A.: PicASHOW: Pictorial authority search by hyperlinks on the web. In: *WWW*. Hong Kong (2001)
170. Lenat, D., Guha, R.: *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley, Boston (1989)
171. Leoncini, A., Sangiacomo, F., Peretti, C., Argentesi, S., Zunino, R., Cambria, E.: Semantic models for style-based text clustering. In: *IEEE ICSC*. Palo Alto (2011)
172. Lew, M., Sebe, N., Djeraba, C., Jain, R.: Content-based multimedia information retrieval: State of the art and challenges. *ACM Trans. Multimedia Comput. Commun. Appl.* **2**(1), 1–19 (2006)
173. Lewis, M.: Self-conscious emotions: Embarrassment, pride, shame, and guilt. In: *Handbook of Cognition and Emotion*, vol. 2, pp. 623–636. Guilford Press (2000)
174. Lewis, M., Granic, I.: *Emotion, Development, and Self-Organization: Dynamic Systems Approaches to Emotional Development*. Cambridge University Press (2002)
175. Lieberman, H., Rosenzweig, E., Singh, P.: ARIA: An agent for annotating and retrieving images. *IEEE Computer* **34**(7), 57–62 (2001)
176. Lieberman, H., Selker, T.: Out of context: Computer systems that adapt to, and learn from, context. *IBM Systems Journal* **39**(3), 617–632 (2000)
177. Lieberman, M.: Social cognitive neuroscience: A review of core processes. *Annual Review of Psychology* **58**, 259–89 (2007)
178. Lin, W., Wilson, T., Wiebe, J., Hauptmann, A.: Which side are you on? identifying perspectives at the document and sentence levels. In: *Conference on Natural Language Learning*, pp. 109–116 (2006)

179. Liu, B.: Sentiment analysis and subjectivity. In: Handbook of Natural Language Processing. Chapman and Hall/CRC (2010)
180. Liu, B., Hu, M., Cheng, J.: Opinion observer: Analyzing and comparing opinions on the web. In: WWW. Chiba (2005)
181. Liu, H., Singh, P.: ConceptNet: A practical commonsense reasoning toolkit. BT Technology Journal (2004)
182. Lu, W., Zeng, K., Tao, D., Yuan, Y., Gao, X.: No-reference image quality assessment in contourlet domain. Neurocomputing **73**(4-6), 784–794 (2012)
183. Lutz, C., White, G.: The anthropology of emotions. Ann. Rev. Anthropology **15**, 405–436 (1986)
184. Ma, C., Osherenko, A., Prendinger, H., Ishizuka, M.: A chat system based on emotion estimation from text and embodied conversational messengers. In: Int'l Conf. Active Media Technology, pp. 546–548 (2005)
185. Ma, H., Chandrasekar, R., Quirk, C., Gupta, A.: Page hunt: Improving search engines using human computation games. In: SIGIR, pp. 746–747. Boston (2009)
186. Machajdik, J., Hanbury, A.: Affective image classification using features inspired by psychology and art theory. In: International Conference on Multimedia. Florence (2010)
187. Maclean, P.: Psychiatric implications of physiological studies on frontotemporal portion of limbic system. Electroencephalogr Clin Neurophysiol Suppl **4**, 407–18 (1952)
188. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Fifth Berkeley Symposium on Mathematical Statistics and Probability, pp. 281–297 (1967)
189. Manning, C., Raghavan, P., Schütze, H.: An Introduction to Information Retrieval. Cambridge University Press, Cambridge (2009)
190. Markotschi, T., Volker, J.: GuessWhat?! - Human intelligence for mining linked data. In: EKAW. Lisbon (2010)
191. Matos, S., Biring, S., Pavord, I., Evans, D.: Detection of cough signals in continuous audio recordings using HMM. IEEE Trans. on Biomedical Engineering **53**(6), 1078–1083 (2006)
192. McCarthy, J.: Programs with common sense. In: Teddington Conference on the Mechanization of Thought Processes (1959)
193. Mehrabian, A.: Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. Current Psychology **14**(4), 261–292 (1996)
194. Mihalcea, R.: Open mind word expert: Creating large annotated data collections with web users' help. In: Workshop on Linguistically Interpreted Corpora. Budapest (2003)
195. Mihalcea, R., Banea, C., Wiebe, J.: Learning multilingual subjective language via cross-lingual projections. In: ACL. Prague (2007)
196. Milewski, A., Smith, T.: Providing presence cues to telephone users. In: ACM Conference on Computer Supported Cooperative Work (2000)
197. Minsky, M.: The Society of Mind. Simon and Schuster, New York (1986)
198. Minsky, M.: Commonsense-based interfaces. Communications of the ACM **43**(8), 67–73 (2000)
199. Minsky, M.: The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind. Simon & Schuster, New York (2006)
200. Moré, J., Watson, G.: The Levenberg-Marquardt Algorithm: Implementation and Theory, *Lecture Notes in Mathematics*, vol. 630, pp. 105–116. Springer, Berlin Heidelberg (1978)
201. Morrison, D., Maillet, S., Bruno, E.: Tagcaptcha: Annotating images with captchas. In: ACM SIGKDD, pp. 44–45. Paris (2009)
202. Mosteller, F., Wallace, D.: Applied Bayesian and Classical Inference: The Case of the Federalist Papers. Springer-Verlag, Berlin Heidelberg (1984)
203. Mueller, E.: Natural Language Processing with ThoughtTreasure. Signifonn (1998)
204. Mueller, E.: Commonsense Reasoning. Morgan Kaufmann (2006)
205. Murphy, G.: The Big Book of Concepts. The MIT Press, Cambridge (2004)
206. Nakazato, M., Manola, L., Huang, T.: ImageGrouper: Search, annotate and organize images by groups. In: S. Chang, Z. Chen, S. Lee (eds.) Recent Advances in Visual Information Systems, *Lecture Notes in Computer Science*, vol. 2314, pp. 93–105. Springer, Berlin (2002)

207. Narwaria, M., Lin, W.: Objective image quality assessment based on support vector regression. *IEEE Trans. on Neural Networks* **12**(3), 515–519 (2010)
208. Neisser, U.: *Cognitive Psychology*. Appleton Century Crofts (1967)
209. Nguyen, L., Wu, P., Chan, W., Peng, W., Zhang, Y.: Predicting collective sentiment dynamics from time-series social media. In: *KDD WISDOM*, 6. Beijing (2012)
210. O’Hare, N., Lee, H., Cooray, S., Gurrin, C., Jones, G., Malobabic, J., O’Connor, N., Smeaton, A., Uscilowski, B.: MediAssist: Using content-based analysis and context to manage personal photo collections. In: *CIVR*, pp. 529–532. Tempe (2006)
211. Ohman, A., Soares, J.: Emotional conditioning to masked stimuli: Expectancies for aversive outcomes following nonre- cognized fear-relevant stimuli. *Journal of Experimental Psychology: General* **127**(1), 69–82 (1998)
212. Ortony, A., Clore, G., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge (1988)
213. Osgood, C., May, W., Miron, M.: *Cross-Cultural Universals of Affective Meaning*. Univ. of Illinois Press (1975)
214. Osgood, C., Suci, G., Tannenbaum, P.: *The Measurement of Meaning*. University of Illinois Press (1957)
215. Pal, P., Iyer, A., Yantorno, R.: Emotion detection from infant facial expressions and cries. In: *International Conference on Acoustics, Speech and Signal Processing*. Dallas (2006)
216. Pampalk, E., Rauber, A., Merkl, D.: Content-based organization and visualization of music archives. In: *ACM International Conference on Multimedia* (2002)
217. Pang, B., Lee, L.: A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In: *ACL*, pp. 271–278. Barcelona (2004)
218. Pang, B., Lee, L.: Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In: *ACL*, pp. 115–124. Ann Arbor (2005)
219. Pang, B., Lee, L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* **2**, 1–135 (2008)
220. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? Sentiment classification using machine learning techniques. In: *EMNLP*, pp. 79–86. Philadelphia (2002)
221. Pantic, M.: Affective computing. In: *Encyclopedia of Multimedia Technology and Networking*, vol. 1, pp. 8–14. Idea Group Reference (2005)
222. Pantic, M., Valstar, M., Rademaker, R., Maat, L.: Web-based database for facial expression analysis. In: *ICME*. Singapore (2005)
223. Papez, J.: A proposed mechanism of emotion. *Neuropsychiatry Clin Neurosci* **7**, 103–112 (1937)
224. Park, H., Jun, C.: A simple and fast algorithm for k-medoids clustering. *Expert Systems with Applications* **36**(2), 3336–3341 (2009)
225. Parrott, W.: *Emotions in Social Psychology*. Psychology Press (2001)
226. Pearl, J.: Bayesian networks: A model of self-activated memory for evidential reasoning. Tech. Rep. CSD-850017, UCLA Technical Report, Irvine (1985)
227. Petridis, S., Pantic, M.: Audiovisual discrimination between laughter and speech. In: *ICASSP*. Las Vegas (2008)
228. Picard, R.: *Affective Computing*. The MIT Press, Boston (1997)
229. Plutchik, R.: The nature of emotions. *American Scientist* **89**(4), 344–350 (2001)
230. Ponzetto, S., Strube, M.: Deriving a large-scale taxonomy from Wikipedia. In: *AAAI*, pp. 1440–1445. Vancouver (2007)
231. PopCap: 2010 social gaming research. Tech. rep., Pop Cap Games (2010)
232. Popescu, A., Etzioni, O.: Extracting product features and opinions from reviews. In: *HLT/EMNLP*. Vancouver (2005)
233. Porkaew, K., Chakrabarti, K.: Query refinement for multimedia similarity retrieval in MARS. In: *ACM International Conference on Multimedia*, pp. 235–238. ACM, New York (1999)
234. Prinz, J.: *Gut Reactions: A Perceptual Theory of Emotion*. Oxford University Press (2004)
235. Pun, T., Alecu, T., Chanel, G., Kronegg, J., Voloshynovskiy, S.: Brain-computer interaction research at the computer vision and multimedia laboratory. *IEEE Trans. on Neural Systems and Rehabilitation Engineering* **14**(2), 210–213 (2006)

236. Qian, N.: On the momentum term in gradient descent learning algorithms. *Neural Networks* **12**, 145–151 (1999)
237. Rao, D., Ravichandran, D.: Semi-supervised polarity lexicon induction. In: *EACL*, pp. 675–682. Athens (2009)
238. Redi, J., Gastaldo, P., Heynderickx, I., Zunino, R.: Color distribution information for the reduced-reference assessment of perceived image quality. *IEEE Trans. on Circuits and Systems for Video Technology* **20**(12), 1757–1769 (2012)
239. Reisberg, D., Heuer, F.: Memory for emotional events. *Memory and Emotion* pp. 3–41 (2004)
240. Reiter, R.: A logic for default reasoning. *Artificial Intelligence* **13**, 81–132 (1980)
241. Richards, J., Butler, E., Gross, J.: Emotion regulation in romantic relationships: The cognitive consequences of concealing feelings. *Journal of Social and Personal Relationships* **20**, 599–620 (2003)
242. Riloff, E., Wiebe, J.: Learning extraction patterns for subjective expressions. In: *EMNLP*, pp. 105–112. Sapporo (2003)
243. Riloff, E., Wiebe, J., Phillips, W.: Exploiting subjectivity classification to improve information extraction. In: *AAAI*. Pittsburgh (2005)
244. Riseberg, J., Klein, J., Fernandez, R., Picard, R.: Frustrating the user on purpose: Using biosignals in a pilot study to detect the user’s emotional state. In: *CHI*. Los Angeles (1998)
245. Rowe, M., Butters, J.: Assessing trust: Contextual accountability. In: *ESWC*. Heraklion (2009)
246. Russell, J.: Affective space is bipolar. *Journal of Personality and Social Psychology* **37**, 345–356 (1979)
247. Russell, J.: Core affect and the psychological construction of emotion. *Psychological Rev.* **110**, 145–172 (2003)
248. Salzman, M., Matathia, I., O’Reilly, A.: *Buzz: Harness the Power of Influence and Create Demand*. John Wiley and Sons, New York (2003)
249. Scherer, K.: Studying the emotion-antecedent appraisal process: An expert system approach. *Cognition and Emotion* **7**, 325–355 (1993)
250. Scherer, K.: Psychological models of emotion. *The Neuropsychology of Emotion* Pages pp. 137–162 (2000)
251. Scherer, K., Shorr, A., Johnstone, T.: *Appraisal Processes in Emotion: Theory, Methods, Research*. Oxford University Press, Canary (2001)
252. Schleicher, R., Sundaram, S., Seebode, J.: Assessing audio clips on affective and semantic level to improve general applicability. In: *Fortschritte der Akustik - DAGA*. Berlin (2010)
253. Sebe, N., Tian, Q., Loupias, E., Lew, M.S., Huang, T.S.: Evaluation of salient point techniques. In: *International Conference on Image and Video Retrieval*, pp. 367–377. Springer-Verlag, London (2002)
254. Shan, C., Gong, S., McOwan, P.: Beyond facial expressions: Learning human emotion from body gestures. In: *BMVC*. Warwick (2007)
255. Siegel, S., Castellan, N.: *Nonparametric Statistics for the Social Sciences*. McGraw-Hill, New York (1988)
256. Simoudis, E.: Reality check for data mining. *IEEE Expert* **11**(5) (1996)
257. Singh, P.: The open mind common sense project. *KurzweilAI.net* (2002)
258. Siorpaes, K., Hepp, M.: Ontogame: Weaving the semantic web by online games. In: *ESWC*, pp. 751–766. Tenerife (2008)
259. Smith, E., DeCoster, J.: Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychological Review* **4**(2), 108–131 (2000)
260. Smith, J., Chang, S.: An image and video search engine for the world-wide web. In: *Symposium on Electronic Imaging: Science and Technology* (1997)
261. Snyder, B., Barzilay, R.: Multiple aspect ranking using the good grief algorithm. In: *HLT/NAACL*. Rochester (2007)
262. Somasundaran, S., Wiebe, J., Ruppenhofer, J.: Discourse level opinion interpretation. In: *COLING*, pp. 801–808. Manchester (2008)

263. Song, Y., Wang, H., Wang, Z., Li, H., Chen, W.: Short text conceptualization using a probabilistic knowledgebase. In: IJCAI. Barcelona (2011)
264. Sood, S., Vasserman, L.: ESSE: Exploring mood on the web. In: ICWSM. San Jose (2009)
265. Sowa, J.: Semantic networks. In: Encyclopedia of Artificial Intelligence. Stuart Shapiro (1987)
266. Speer, R.: Open Mind Commons: An inquisitive approach to learning common sense. In: Workshop on Common Sense and Interactive Applications. Honolulu (2007)
267. Speer, R., Havasi, C.: ConceptNet 5: A large semantic network for relational knowledge. In: E. Hovy, M. Johnson, G. Hirst (eds.) Theory and Applications of Natural Language Processing, chap. 6. Springer (2012)
268. Speer, R., Havasi, C., Lieberman, H.: Analogospace: Reducing the dimensionality of common sense knowledge. In: AAAI (2008)
269. Speer, R., Havasi, C., Treadway, N., Lieberman, H.: Finding your way in a multi-dimensional semantic space with Luminoso. In: IUI. Hong Kong (2010)
270. Srinivasan, U., Pfeiffer, S., Nepal, S., Lee, M., Gu, L., Barrass, S.: A survey of mpeg-1 audio, video and semantic analysis techniques. *Multimedia Tools and Applications* **27**(1), 105–141 (2005)
271. Sriram, B., Fuhry, D., Demir, E., Ferhatosmanoglu, H., Demirbas, M.: Short text classification in twitter to improve information filtering. In: ACM SIGIR. Geneve (2010)
272. Stevenson, R., Mikels, J., James, T.: Characterization of the affective norms for english words by discrete emotional categories. *Behavior Research Methods* **39**, 1020–1024 (2007)
273. Stork, D.: The open mind initiative. *IEEE Intelligent Systems* **14**(3), 16–20 (1999)
274. Strapparava, C., Mihalcea, R.: Learning to identify emotions in text. In: ACM Symp. Applied Computing, pp. 1556–1560 (2008)
275. Strapparava, C., Valitutti, A.: WordNet-Affect: An affective extension of WordNet. In: LREC, pp. 1083–1086. Lisbon (2004)
276. Suchanek, F., Kasneci, G., Weikum, G.: Yago: a core of semantic knowledge. In: WWW, pp. 697–706. Banff (2007)
277. Tatemura, J.: Virtual reviewers for collaborative exploration of movie reviews. In: IUI. New Orleans (2000)
278. Terveen, L., Hill, W., Amento, B., McDonald, D., Creter, J.: PHOAKS: A system for sharing recommendations. *Communications of the ACM* **40**(3) (1997)
279. Thaler, S., Siorpaes, K., Simperl, E., Hofer, C.: A survey on games for knowledge acquisition. Tech. rep., Semantic Technology Institute (2011)
280. Tong, R.: An operational system for detecting and tracking opinions in on-line discussion. In: SIGIR. New Orleans (2001)
281. Torrance, G., Thomas, W., Sackett, D.: A utility maximisation model for evaluation of health care programs. *Health Services Research* **7**, 118–133 (1972)
282. Tracy, J., Robins, R., Tangney, J.: *The Self-Conscious Emotions: Theory and Research*. The Guilford Press (2007)
283. Trout, J., Rivkin, S.: *The New Positioning: The Latest on the World's #1 Business Strategy*. McGraw Hill, New York (1996)
284. Truong, K., Van Leeuwen, D.: Automatic discrimination between laughter and speech. *Speech Communication* **49**, 144–158 (2007)
285. Turkle, S.: *The Second Self: Computers and the Human Spirit*. Simon & Schuster (1984)
286. Turney, P.: Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In: ACL, pp. 417–424. Philadelphia (2002)
287. Turney, P., Littman, M.: Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems* **21**(4), 315–346 (2003)
288. Urban, J., Jose, J.: EGO: A personalized multimedia management and retrieval tool. *International Journal of Intelligent Systems* **21**(7), 725–745 (2006)
289. Urban, J., Jose, J., Van Rijsbergen, C.: An adaptive approach towards content-based image retrieval. *Multimedia Tools and Applications* **31**, 1–28 (2006)
290. Valstar, M., Gunes, H., Pantic, M.: How to distinguish posed from spontaneous smiles using geometric features. In: ICMI. Nagoya (2007)

291. Velikovich, L., Goldensohn, S., Hannan, K., McDonald, R.: The viability of web-derived polarity lexicons. In: NAACL, pp. 777–785. Los Angeles (2010)
292. Vesterinen, E.: Affective computing. In: Digital Media Research Seminar. Helsinki (2001)
293. Wallhoff, F.: Facial expressions and emotion database. Tech. rep., Technische Universitat Munchen (2006)
294. Wang, Q., Cambria, E., Liu, C., Hussain, A.: Common sense knowledge for handwritten chinese recognition. *Cognitive Computation* (2012)
295. Ware, J.: Scales for measuring general health perceptions. *Health Services Research* **11**, 396–415 (1976)
296. Ware, J., Kosinski, M., Keller, S.: A 12-item short-form health survey: Construction of scales and preliminary tests of reliability and validity. *Medical Care* **34**(3), 220–233 (1996)
297. Ware, J., Sherbourne, C.: The MOS 36-item short-form health survey (SF-36). Conceptual framework and item selection. *Medical Care* **30**, 473–83 (1992)
298. Wessel, I., Merckelbach, H.: The impact of anxiety on memory for details in spider phobics. *Applied Cognitive Psychology* **11**, 223–231 (1997)
299. Westen, D.: Implications of developments in cognitive neuroscience for psychoanalytic psychotherapy. *Harv Rev Psychiatry* **10**(6), 369–73 (2002)
300. Whissell, C.: The dictionary of affect in language. *Emotion: Theory, Research, and Experience* **4**, 113–131 (1989)
301. Wiebe, J., Wilson, T., Bruce, R., Bell, M., Martin, M.: Learning subjective language. *Computational Linguistics* **30**(3) (2004)
302. Wiebe, J., Wilson, T., Cardie, C.: Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation* **39**(2), 165–210 (2005)
303. Wilks, Y., Stevenson, M.: The grammar of sense: Using part-of-speech tags as a first step in semantic disambiguation. *Journal of Natural Language Engineering* **4**(2) (1998)
304. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: HLT/EMNLP. Vancouver (2005)
305. Wilson, T., Wiebe, J., Hwa, R.: Just how mad are you? Finding strong and weak opinion clauses. In: AAAI, pp. 761–769. San Jose (2004)
306. Winston, P.: Learning structural descriptions from examples. *The Psychology of Computer Vision* pp. 157–209 (1975)
307. Witten, I., Frank, E.: *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, San Francisco (2005)
308. Wu, H., Luk, R., Wong, K., Kwok, K.: Interpreting TF-IDF term weights as making relevance decisions. *ACM Transactions on Information Systems* **26**(3) (2008)
309. Wu, W., Li, H., Wang, H., Zhu, K.: Probase: A probabilistic taxonomy for text understanding. In: SIGMOD, pp. 481–492. Scottsdale (2012)
310. Xia, R., Zong, C., Hu, X., Cambria, E.: Feature ensemble plus sample selection: A comprehensive approach to domain adaptation for sentiment classification. *IEEE Intelligent Systems* (2013). DOI 10.1109/MIS.2013.27
311. Yan, J., Yu, S.Y.: Magic bullet: a dual-purpose computer game. In: ACM SIGKDD, pp. 32–33. Paris (2009)
312. Zeki, S., Romaya, J.: Neural correlates of hate. *PLoS one* **3**(10), 35–56 (2008)
313. Zeng, Z., Pantic, M., Roisman, G., Huang, T.: A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **31**(1), 39–58 (2009)
314. Zirn, C., Niepert, M., Stuckenschmidt, H., Strube, M.: Fine-grained sentiment analysis with structural features. In: IJCNLP. Chiang Mai (2011)
315. van Zwol, R., Garcia, L., Ramirez, G., Sigurbjornsson, B., Labad, M.: Video tag game. In: WWW. Beijing (2008)