Deriving Explanations from Qualitative Models

Paulo Salles
University of Edinburgh
Inst. of Ecology & Res. Man.
Mayfield Road
Edinburgh EH9 3JU
United Kingdom
+44-131-6505408
psalles@holyrood.ed.ac.uk

Bert Bredeweg
University of Amsterdam
Dept of Social Science Inf.
Roetersstraat 15
1018 WB Amsterdam
The Netherlands
+31-20-525 6788
bert@swi.psy.uva.nl

Radboud Winkels
University of Amsterdam
Dept of Comp. Science & Law
PO BOX 1030
1000 BA Amsterdam
The Netherlands
+31-20-525 3485
winkels@lri.jur.uva.nl

Abstract

Qualitative computer simulations have great potential for teaching people to understand and interact with their physical environment. Prerequisite for using that potential, is that these simulations can be explained to humans in ways that they comprehend. Preferably, these explanations should be generated on the basis of the qualitative models that underly the simulations, to avoid having to handcraft the explanations for every new domain. The research that we describe in this paper deals with exactly that problem. It combines two lines of earlier research: representing qualitative models, GARP [2], and didactic discourse planning, DDP [13]. All qualitative models represented in GARP can be questioned by students, using an as yet limited set of questions, that will be answered by a generic didactic discourse planner. The overall interaction between students and systems is guided by a 'mental tour' through the successive states of the simulation (the 'envisionment graph'). At each state several questions can be asked. These questions are linked to 'information needs', the topics of discourse. On the basis of these topics, the discourse planner will plan sequences of utterances, taking into account such things as the students beliefs, and the state of the discourse process.

Publication details

P. Salles and B. Bredeweg and R. Winkels, Deriving Explanations from Qualitative Models. In: *Artificial Intelligence in Education: Knowledge and Media in Learning Systems*, B. du Bouley and R. Mizoguchi (eds.), pages 474-481, 1997 (proceedings of the 8th World Conference on Artificial Intelligence in Education, Japan, August 18-22, 1997).

1 Introduction

Early systems such as SOPHIE [4] and STEAMER [6] can be seen pioneering landmarks in trying to have computers communicate knowledge about the behaviour of (physical) systems with students. Efforts such as these gave raise to an area of AI research known as Qualitative Reasoning (QR). In this paper we investigate how qualitative models can be used for explanation purposes. Our work combines two lines of earlier research: representing qualitative models, GARP [2], and didactic discourse planning, DDP [13]. We have constructing a prototype of an Interactive Learning Environment (ILE), based on the qualitative models implemented in GARP, that can be consulted by students to investigate the behavioural characteristics of some (physical) system. The specific domain that we are dealing with in this paper is the ecology of the Brazilian cerrado vegetation and the effects of fire on this vegetation. It provides a diversified set of problems for exploring the possibilities of automatic generation of explanations.

2 Ecology in Brazilian Cerrado

The central region of Brazil is covered by a vegetation called cerrado. This huge area of almost 2 million square kilometres is characterised by a tropical climate, with two well marked seasons (wet and dry), and by soils that have low fertility. Within the cerrado vegetation it is possible to identify several types of cerrado physiognomies, spanning from open grasslands to more or less closed forests. These physiognomies have well defined floristic composition and are mainly determined by fire, soil fertility and the amount of water available in the soil during the dry season. Researchers have investigated the effects of fire on the cerrado (e.g. [5, 8]). It has been shown that fire can affect both physical and biological factors: fire causes changes on (1) the energy flux, (2) water and nutrient cycles, (3) the species composition of communities, (4) the biomass in many plants, and (5) stimulates flowering and germination of seeds in many species. Fire can therefore be used as a management tool, for example to stabilise the vegetation in certain areas, and to reduce the risk of big fire events [9]. Qualitative simulations are particularly important for ecological domains. Not only are quantitative data often almost non-existent, qualitative models also provide many additional features that are important for having learners interact with the simulation (see also [3]).

3 Modelling the Cerrado Vegetation

Following important research on the Brazilian cerrado vegetation [5, 9] our models should support explanations about the relation between fire frequency and the structure of the vegetation, expressed as follows: (1) If the fire frequency decreases, for example because of human actions, then succession will occur and as a consequence the vegetation will evolve and become denser, with more trees and shrubs and less grass. (2) If the fire frequency increases, then a degradation process is active and the vegetation tends to become more open, with less trees and shrubs and more grass. The models should provide access to the underlying causal models that represents how these changes follow from different responses of populations of trees, shrubs and grass to environmental influences such as light, humidity and temperature. Next to these more general modelling requirements we enforced a number of additional requirements in order to make the model useful for our teaching purposes (see also [10] for details):

Population as a key concept Reasoning about changes in communities requires knowledge about populations [11]. We have therefore developed a kernel of partial models that represents general knowledge about populations which can be used in different situations.

Model fragments as knowledge chunks Model fragments should represent stand-alone parts of the domain knowledge that students should master. The idea is that each relevant domain concept should be expressed in one model fragment. Model fragments will therefore be an important ingredient for deriving an explanation.

Assumptions as simplifying model fragments In order to reuse detailed model fragments in complex scenarios (problem situations) it is often necessary to reduce the range of the model. In other words, more complex scenarios usually require a more abstract view in order to keep the interpretation of the simulation tractable. In our model we realised more abstract views by defining model fragments that summarize certain variations at a lower level.

Aggregated processes and agent models Typically in ecology, different vocabulary exists for reasoning about changes in an ecological system at different levels of detail. In order to be able to represent this aspect we introduced the notion of an 'aggregated process'. It consists of the sum of a number processes at a lower level. In order to represent actions performed by humans within an ecological system we use the notion of an agent model [2]. Usually an agent model sets the value of some derivative, such as a decrease in fire frequency representing the notion of 'conservation'.

3.1 Models of Populations

Populations consist of groups of individuals of the same kind, living in a certain place, during a certain period of time. The size of the population is an important factor (Number-of), because it is an indication of the balance of the forces acting on the individuals. It can take on values from different quantity spaces, depending on the objectives of the model. For example, we used the quantity space {zero, low, medium, high, maximum} to describe the 'absolute' qualitative values, and make comparisons between populations.

Following our assumption that each relevant concept should be represented by a model fragment, we defined fragments for the concepts: small, medium, large, maximum sized, non-existent and extinct populations. Behaviour is expressed in model fragments representing the notions of increasing, decreasing and steady populations. Combining them we can say, for instance, "the population is small and increasing".

In order to predict changes in populations and to build explanations about the results of simulations, we need a vocabulary to express the basic processes affecting the individuals: natality, mortality, immigration, and emigration. These basic processes introduce the following quantities (rates): Born-flow, Dead-flow, Immigrated-flow and Emigrated-flow. These flows have the quantity space {zero, plus} (they cannot be negative). Note that these processes are independent of each other. We need the notion of population growth to express how the basic processes combine in a particular situation. The growth process is defined as an aggregation of the four basic processes and represents a unique flow (the addition of the individual flows). It introduces the quantity Growth-rate with the quantity space {minus, zero, plus}.

3.2 Models of Communities and Ecosystems

Communities are groups of populations. Communities in cerrado can be classified according to the Number-of trees, shrubs and grass. Typically, researchers classify the cerrado into: campo limpo, campo sujo, campo cerrado, cerrado sensu stricto and cerradao [5]. The cerrado communities are related to a general ecosystem, the cerrado sensu lato. The model fragment that defines the cerrado sensu lato specifies relevant properties of the micro-environment at the surface of the soil. It introduces the quantities Nutrient, Humidity, Light and Temperature. These quantities are related to the amount of Litter, the dead material that covers the ground (leaves, small pieces of wood and other parts of plants). We assume that these factors are always present in any scenario described by the models. Thus, their quantity space is: {plus}.

The canopy of the trees has an important influence on the factors mentioned above. In our models this is represented by the quantity Cover, with the same quantity space as used for the population of trees. It is assumed that there exists a direct correspondence between the Number-of trees and the amount of Cover: the value taken by the former is also assigned to the latter. For example, if the value of Number-of trees is low, then Cover is also low.

All the above mentioned factors are influenced by fire. Fire frequency is a component of the so called 'fire regime' [12]. It expresses how often a vegetation is burned. In the model this is represented by the quantity Fire-frequency, which can take on values from the quantity space {zero, plus}. Fire frequency changes as a consequence of human actions. This is modelled by agent models.

The influence from fire frequency on the community is indirect: it propagates through the described network of environmental factors that finally influences the basic processes underlying populations. Altogether, 13 direct influences (I) and 28 indirect influences (P) constitute the structure of the causal model, as shown in Figure 1.

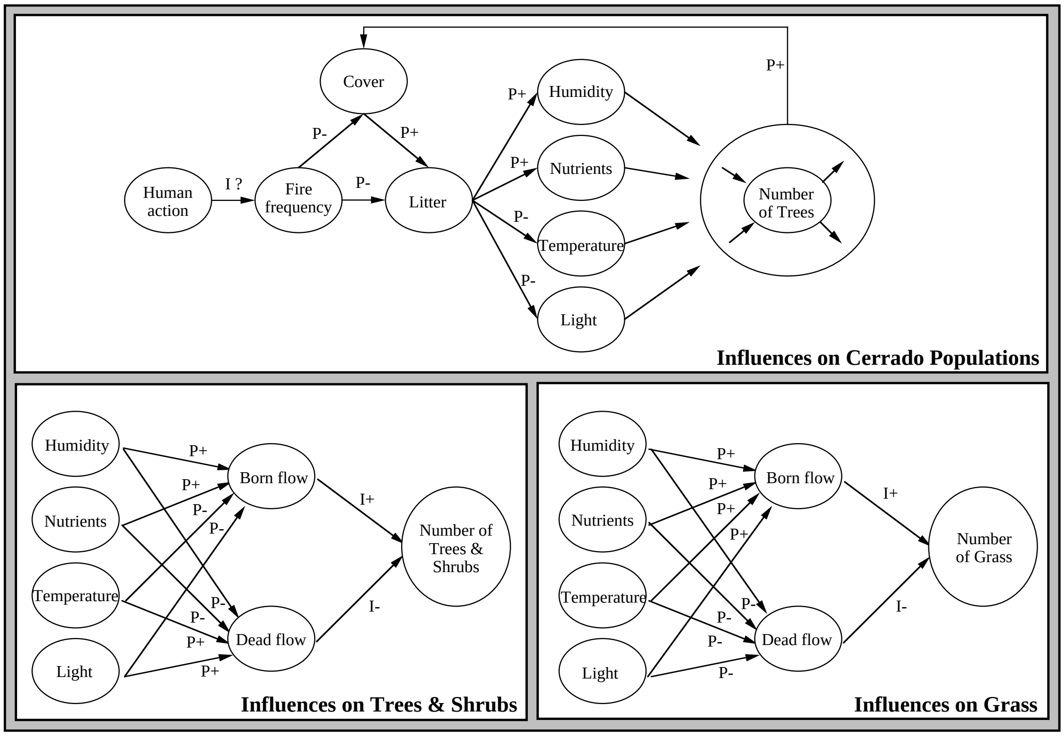

Figure 1: Causal model for the cerrado communities

3.3 Results from Simulations

A simulation with this model produced the envisonment graph depicted in Figure 2. It shows the successional changes in the cerrado predicted by the hypothesis presented in the beginning of this section. In order to reduce the number of ambiguities and possible states in the full simulation, we made some assumptions. The most important one was redefining the campo cerrado as a typical community with the values for trees, shrubs and grass equal to medium. We also removed the influences from Humidity and Nutrient on the population of grass in order to reduce ambiguity.

The resulting qualitative models offer several possibilities for tutoring. Using different initial scenarios, it is possible to explore selected parts of the causal path. For example, we can analyze the effect of each factor on the populations, or the effects of a group of factors on a specific population. Currently we are expanding the library, to be able to represent other aspects of the ecology of fire in the cerrado.

4 Generating explanations

In the literature there are at least two fundamentally different approaches to generating explanations. One is based on typical rhetorical structures of explanations, represented in so called *schemata*, and is exemplified by McKeown's TEXT system [7] for answering questions about a database. The second major approach to explanation is that of using planning formalisms to dynamically plan sequences of utterances to achieve certain communicative goals. It is exemplified by Appelt's work on KAMP [1].

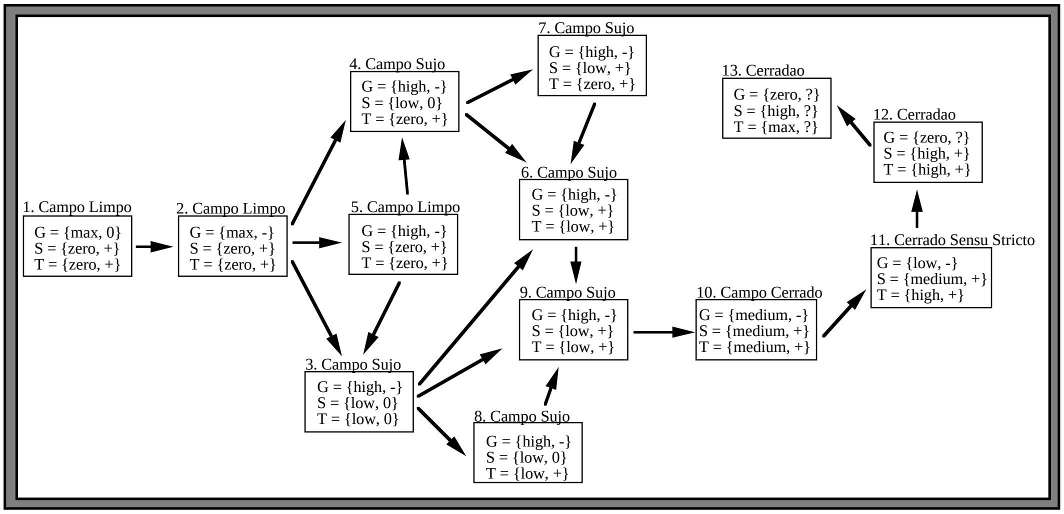

Figure 2: Succession in cerrado vegetation

In the EUROHELP project, Winkels used a combination of the schemata and planning techniques in his didactic discourse planner [13]. Discourse strategies are planned on the basis of communicative goals, but the starting point is a library of ‘skeletal plans’ that have worked in the past. Only when these strategies do not work, general fall-back strategies are refined to meet the current information need. The skeletal strategies can be viewed as hierarchical schemata, that still contain the communicative (sub)goals they are supposed to meet. It is this last approach we will follow in this research.

The overall interaction between students and the ILE is guided by a ‘mental tour’ trough the successive states of the simulation (the ‘envisionment graph’). At each state several questions can be asked. These questions are linked to ‘information needs’, the topics of discourse. On the basis of these topics, the discourse planner will plan sequences of utterances, taking into account such things as the students beliefs, and the state of the discourse process.

4.1 Asking questions

At present, the student can ask the following questions about a simulation in a specific state:

1. What are the system elements in the present state? This is asking about basic concepts, its instances, its attributes and values, and relations between concepts or instances.

2. What are the initial causes of change in the state? This is asking about processes and agents that can cause change. Causality is modelled as ‘influences’ between quantities. Agents will be given in the input system. Inequalities between quantities may trigger a process (e.g. quantity ‘Number-of’ greater than zero will trigger a mortality process).

3. How does change propagate to other quantities in the present state? This is asking about ‘proportionalities’ between quantities, and possibly new ‘influences’.

4. How does a particular quantity change over states? This is asking about values and derivatives of a quantity. For a specific quantity, it is easy to find its value and derivative in every state of the simulation.

Questions are formed by filling a text form or template with instances from the simulation. There is a logical order in these questions. One cannot ask (or explain) propagation of change, before one knows about the initial causes of change, etc. Therefore, questions have preconditions attached to it, that check whether the necessary prerequisite information is already available. Given a student’s question, we have to determine *what* to say, i.e. what his or her information need is, what the topic(s) of the interaction will be. This is what McKeown calls the 'relevant knowledge pool' [7], or Winkels calls the 'topicalization' process [13]. Since there is no task the student is supposed to do (except for understanding the 'cerrado' model), there is no generic diagnostic process that tries to infer the student's information need when he or she asks a question. A procedure to determine the initial topic is directly linked to the questions. These initial topics can be shortened or extended by the discourse planner when needed (see below). An example question is:

Type: Propagation of influence X within state
Conditions: Influence X has been introduced
Template: "How does X propagate in this state?"
Procedure: Find all proportionalities between the quantity that is being influenced X and other quantities. Look recursively for influences or proportionalities with these other quantities until no more can be found.

We will illustrate the workings of such a 'topicalization' procedure for the simulation shown in Figure 1. The scenario specifies a population of grass ('population3'), and a human agent called 'fire-decreaser1'. The initial change is caused by the human agent 'fire-decreaser1' which causes a decrease in 'fire-frequency1'. Now the question: "How does the negative influence of fire-decreaser1 on fire-frequency1 propagate in this state?", leads to a topic in the following way. First, find all proportionalities between 'fire-frequency1' and other quantities. There is only one in this state, a negative proportionality with 'litter1' (fire will burn litter on the ground). Now look for influences or proportionalities on this other quantity: positive proportionalities with 'moisture1' and 'nutrient1' (litter provides nutrients and will keep the soil moist), negative proportionalities with 'light1' and 'temperature1' (litter will block light and warmth). And again, for all these quantities, look for influences or proportionalities on them, etc.

```
[ inf-neg-by(fire-frequency1,fire-decreaser1),
  [ prop-neg(litter1,fire-frequency1),
    [ prop-pos(moisture1,litter1), [],
      prop-neg(light1,litter1),
        [ prop-pos(born-flow3,light1), [inf-pos-by(number-of-3,born-flow3), []],
          prop-neg(dead-flow3,light1), [inf-neg-by(number-of-3,dead-flow3), []]
        ],
      prop-pos(nutrient1,litter1), [],
      prop-neg(temperature1,litter1),
        [ prop-pos(born-flow3,temperature1),
            [inf-pos-by(number-of-3,born-flow3), []],
          prop-neg(dead-flow3,temperature1),
            [inf-neg-by(number-of-3,dead-flow3), []]
]]]]
```

4.2 Planning the explanation

Next, plan the interaction that is aimed at getting the needed information across to the user. This planning process is done by a generic *Didactic Discourse Planner* (DDP) [13]. Basically, DDP takes the information need ('local need') and first looks in a library of skeletal discourse strategies to see if one of those is applicable in the current situation. If it is, it is instantiated to the current situation, and the strategy will be transformed to natural language and presented to the student. If none can be found, general fall-back strategies will be refined to deal with the situation. The strategies take care of skipping or summarizing information, possibly extending parts, sequencing it, minimizing shifts of focus, etc.

DDP's strategies implement general principles for (didactic) discourse, of which the two most important ones are:

Given $\longrightarrow$ New. Always try to link new information to given, or "known" information. Practically, this means linking new information to: Something the student already knows (Student Model), something that has recently been taught (Coaching History), something that had just been mentioned (Discourse Model), or something that has just happened (Performance History).

Conciseness and relevance. Try to be to the point, do not explain things the student already knows, or is assumed to know. Do not introduce new topics, unless necessary for understanding the new information. Whenever possible, use references to existing ('given') information instead of describing a topic again (cf. 'given-new' principle above). An interesting example is explaining a topic when a *similar* topic is known, or has just been described. In that case, focus on the differences.

4.3 Explaining Qualitative Primitives

As mentioned before, the overall interaction between a student and the ILE is guided by a tour through the *envisionment graph* (see Figure 2). The student can ask questions that lead to topics (information needs) to be handled by the discourse planner. The discourse strategies take care of general rational, and didactic principles at the higher levels, but the utterances at the lowest, 'executable' level will have to map onto the knowledge representation of the qualitative simulator, i.c. GARP. In this section we will discuss the ways to explain the GARP primitives, and show how these primitives can be combined to assemble more complex explanations.

For the primitives we represent:

Type: Label to indicate its use;
KR: the GARP knowledge representation it maps onto;
Loc: Where the instantiations for the current simulation can be found in the state description;
Known: The information that should be known or given in order for the primitive to be used (a special type of condition). This means, the information should be either in the student model or the discourse model of the current session;
Cond: Other conditions that need to be met for this version of the primitive to be used. They may refer to the current state of the simulation (e.g. the value of a derivative), and the current state of the discourse (e.g. topic in focus);
NL: The natural language expression for the primitive when the conditions are met.

We distinguish the following primitives:

1. Explaining basic concepts: These building blocks are used to explain system elements of a simulation, their attributes, and their interrelations (the 'isa' hierarchy).
2. Explaining quantities of an instance, their values, value ranges, and derivatives.
3. Explaining a causal dependency between quantities: Used for explaining direct causality (modelled as influences) and indirect causality (modelled as proportionalities).
4. Explaining constraints between quantities and/or values.

An example primitive is:

```
Type:   explaining a quantity, its value, and derivative (no value scope)
KR:     Generic Name( Instance, Inst Qname, continuous, Qspace )
        value( Inst Qname, unk, Value, Derivative )
Known:  Instance
Where:  in list of parameters and par_values in a predicted state (SMD)
Cond:   Derivative is plus; Qspace is only Value; focus on Inst Qname
NL:     "Inst Qname is the quantity Generic Name of Instance.
         It has currently the value Value and is increasing.''
```

4.4 Combining primitives

In answering questions, the discourse planner will eventually combine the primitives. For the example question presented above, concerning the propagation of change of the decreasing fire frequency, this may result in an explanation as the following:

```
[context]
  [remind basic concept]
    You know: There is a cerrado referred to as cerrado1
  [remind quantity]
```

```
    You know: cerrado1 has a quantity fire frequency referred to as
      fire-frequency1. It has currently the value plus and is decreasing.
[new information]
  [explain causal dependency]
    The decrease in fire-frequency1 increases the litter1
    [explain quantity]
      litter1 is the quantity litter of cerrado1.
      It has currently the value plus and is increasing.
  [signalling]
    The increase in litter1 has four effects:
  [explain causal dependency]
    1. The increase in litter1 increases the moisture1
  [explain causal dependency]
    2. The increase in litter1 increases the nutrient1
  [explain causal dependency]
    3. The increase in litter1 decreases the light1
       [explain causal dependency]
         The decrease in light1 decreases the born-flow3
         [explain quantity]
           born-flow3 is the quantity born-flow of grass1.
           It has currently the value plus and is decreasing.
         [explain causal dependency]
           The decreasing amount of born-flow3 decreases the number-of3
           [explain quantity]
             number-of3 is the quantity number-of of grass1.
             It has currently the value max and is stable.
       [explain causal dependency]
         The decrease in light1 increases the dead-flow3
         [explain causal dependency]
           The increasing amount of dead-flow3 decreases the number-of3
  [explain causal dependency]
    [explain similar]
    4. The increase in litter1 decreases the temperature1
       [refer same]
         This propagates the same way as light1
       [explain difference]
         but now for temperature1
```

5 Discussion and Concluding Remarks

In this paper we have described the initial results of a research project designed to create an ILE combining previous work in representing qualitative models, GARP [2], and in didactic discourse planning, DDP [13]. Our objective was to explore how the primitives used for building the qualitative models can be used to support explanations.

We built a set of qualitative models that represent some widely accepted hypothesis about the effects of fire on the vegetation dynamics of the Brazilian cerrado. In modelling this domain requirements were formulated in order to make the resulting model more useful for teaching purposes. Particulary, the notion of having model fragments represent stand-alone parts of the domain knowledge that students should master was important in this respect.

The qualitative models and the results of the simulations were used by the DDP to create the explanatory discourse. Although the current state of the prototype does not allow for the full use of DDP's explanatory possibilities, the students can inspect the qualitative models and the reasoning process, and ask questions about them. Given a particular question, a set of procedures is triggered to decide what to say. The topics selected are worked out according to a set of general tactical and strategical principles by the discourse planner. Next, using GARP's primitives, the system collects the knowledge to compose the answer. The primitives are finally combined and an answer for the question is produced. Some examples of how the main elements in the discourse map into the knowledge represented in the qualitative models are presented in this paper.

We are currently extending the ILE prototype in order to cover a broader range of facilities already present both in GARP and in DDP. This includes a more detailed ontology to represent the ecology of the cerrado, the definition of tasks and problems to be solved by the students in order to provide more context for the generation of explanations, and a refinement in the process of mapping the primitives of qualitative models and the domain concepts.

Acknowledgements

This work has been carried out during a visit of Paulo Salles at the University of Amsterdam. We are grateful to the FUOS (96B20) for the financial support for this visit. Paulo Salles is also grateful to the Brazilian Conselho Nacional do Desenvolvimento Cientifico e Tecnologico (CNPq) for the scholarship to pursue a PhD program at the University of Edinburgh (process 201823/92-6). Paulo Salles is on leave from: Laboratory for Research on Biology and Science Teaching, Department of Genetics and Morphology (GEM), University of Brasilia, 70.910 - Brasilia - DF, Brazil. E-mail: psalles@guarany.unb.br. We would like to thank Joost Breuker, Kees de Koning and Bob Muetzelfeldt for their support and comments on the research presented in this paper.

References

[1] D. Appelt. Planning english referring expressions. *Artificial Intelligence*, 15(3):143–178, 1985.

[2] B. Bredeweg. *Expertise in qualitative prediction of behaviour*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands, March 1992.

[3] B. Bredeweg and R. Winkels. Qualitative models in interactive learning environments: An introduction. *Interactive Learning Environments*, 5:(in Press), 1996.

[4] J. S. Brown, R. R. Burton, and J. de Kleer. Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. In D. Sleeman and J. S. Brown, editors, *Intelligent Tutoring Systems*, pages 227–282, New York, 1982. Academic Press.

[5] L. M. Coutinho. Fire in the ecology of the brazilian cerrado. In J. G. Goldammer, editor, *Fire in the Tropical; Biota – Ecosystem Processes and Global Changes*, pages 82 – 105. Springer-Verlag, Berlin, Germany, 1990.

[6] J.D. Hollan, E.L. Hutchins, and L.M. Weitzman. Steamer: An interactive, inspectable, simulation-based training system. In G. Kearsley, editor, *Artificial intelligence and instruction: applications and methods*, pages 113–134. Addison-Wesley, Reading (Mass), 1987.

[7] D. M. McKeown, W. A. Harvey, and J. McDermott. Rule based interpretation of aerial imagery. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, PAMI–7(5):570–585, 1985.

[8] H. S. Miranda, C. H. Saito, and B. F. S. Dias. *Impactos de Queimadas em Areas de Cerrado e Caatinga*. Universidade de Brasilia, Brasilia, Brasil, 1996.

[9] V. R. Pivello. *An Expert System for the Use of Prescribed Fires in the Management of Brazilian Savannas*. PhD thesis, Imperial College, London, United Kingdom, 1992.

[10] P. Salles and B. Bredeweg. Building qualitative models in ecology. In L. Ironi, editor, *Proceedings of the International workshop on Qualitative Reasoning, QR'97*, pages 155–164, Italy, June 1997. Istituto di Analisi Numerica C.N.R. Pavia.

[11] P. S. B. A. Salles, R. I. Muetzelfeldt, and H. Pain. Qualitative models in ecology and their use in intelligent tutoring systems. In *Proceedings of the QR'96 Workshop*, pages 216 – 224, Menlo Park, California, May 1996. AAAI Press.

[12] R. J. Whelan. *The Ecology of Fire*. Cambridge University Press, Cambridge, 1995.

[13] R. G. F. Winkels. *Explorations in Intelligent Tutoring and Help*. IOS, Amsterdam, The Netherlands, 1992.