Learning MDL-guided Decision Trees for Constructor-Based Languages*

C. Ferri-Ramírez J. Hernández-Orallo M.J. Ramírez-Quintana

DSIC, UPV, Camino de Vera s/n, 46020 Valencia, Spain.
{cferri,jorallo,mramirez}@dsic.upv.es

Abstract. In this paper, we present a method for generating very expressive decision trees based on several partitions over a functional logic language. First, trees are generated top-down and partitions are selected according to the MDL principle. Secondly, not only a decision tree is obtained, but multiple decision trees can be produced guided as well by the MDL principle. This means that the overall search space is scoured in a short-to-long fashion, so allowing a better use of computational resources. The most important result of this paper is a framework for the efficient generation of constructor-based models, which are required for learning from semi-structured information, such as XML documents.

Keywords: Decision Trees, Inductive Functional Logic Programming (IFLP), Inductive Logic Programming (ILP), Minimum Description Length (MDL), eXtended Markup Language (XML).

1 Introduction

Decision trees are among the most popular tools for machine learning and data mining. There are several reason for this. First, they are an intuitive representation of the concept that is to be learned. Second, the trees can be constructed relatively fast compared with other methods. Nonetheless, their greedy behaviour jointly with the absence of backtracking make that sometimes the exploration of partitions may lead to bad solutions. However, to explore all the search space would be unfeasible. It is then necessary to sort out the search space in order to obtain the best subset of trees first.

The top-down induction of decision trees is performed in two phases: *building* phase and *pruning* phase. In the first phase a decision tree is grown by partitioning the input data into two or more subsets using a condition on an attribute (the *splitting attribute*) until a *stopping criterion* holds. Then, the pruning phase may remove some nodes to prevent *overfitting*. Some of these algorithms are CART [2], ID3 [13], C4.5 [15] and FOIL [14].

However, the existing decision trees algorithms are not capable to deal directly with attributes whose values belong to a constructor-based data type,

* This work has been partially supported by CICYT under grant TIC 98-0445-C03-C1 and Generalitat Valenciana under grant GV00-092-14.

such as the natural numbers defined by the 0 and $s(\_)$[1] constructor symbols, or lists defined by the constructor symbols *emptylist* and $\cdot(\_,\_)$[2]. This is due to the fact that the type definition of an attribute is not considered as a splitting criterion. This is usually overcome by the use of background predicates such as $suc(X,Y)$, $member(X,Y)$, etc., but this can only be solved when there are few types and they are known a priori. Nevertheless, handling constructor-based types is not only interesting in the definition of functions but are essential for some applications. For instance, the mining of XML documents, which are represented as terms of constructor-based types, requires to be able to work with their structure, usually without background knowledge to 'navigate' inside it.

In this paper we present a top-down method for inducing decision trees in a functional logic inductive framework (IFLP) [4]. We call this approach *Constructor-based Decision Tree Learning* (CDTL) since the idea is to use the type information of the functions to generate the trees. For the learning of a target function f, the training sample is composed by ground equations whose left hand sides are terms of the form $f(\ldots)$ and whose right hand sides are the function result (the *class* in the decision tree terminology). The root node of the tree is an equation of the form $f(X_1,\ldots,X_{n-1}) = X_n$. The attribute splits proposed here include tests on data types, on recursive definitions or on the definition of an attribute as the result of a call to a background knowledge function. Since we consider that the class is another attribute to test, the above split conditions allow the induction of nested and recursive functions as well as the use of background knowledge in the definition of the target function. In this way, this algorithm extends the IFLP framework making it conditional and capable of inducing more efficiently functions with a high arity.

Unlike other decision tree induction algorithms, for which general-to-specific or specific-to-general searches are used, in our approach we follow a short-to-long search. Hence, the splitting criterion is based on the Minimum Description Length (MDL) principle. The MDL principle has been previously used in the induction of decision trees, but only in the pruning phase [17, 10]. The use of the MDL principle at the generation phase is justified because other quality criteria based on discrimination such as the *information gain* [15], the *information gain ratio* [15] or the *Gini heuristic* [2] are not useful for problems such as $f(\ldots) = f(\ldots)$ or $f(\ldots) = g(\ldots)$ being g a function of the background knowledge. Another reason is that the guiding of the search by the MDL principle ensures a better use of computational resources [8, 19].

The paper is organised as follows. In Section 2 we introduce some basic notions about functions and the IFLP framework. Section 3 defines the set of partitions and compares them with those of some well-known algorithms. Sev-

[1] The s symbol stands for the successor function symbol.

[2] In functional programming and functional logic programming, it is usual to define a function as a program whose equations establish the function behavior w.r.t. the constructor terms of (some of) its arguments. For instance, the program $\{sum(0,Y) = Y, sum(s(X),Y) = s(sum(X,Y))\}$ define the addition of natural numbers.

eral criteria to guide the search are presented in Section 4. Some considerations about unary partitions are analised in Section 5. Finally, Section 6 presents the conclusions. [3]

2 Preliminaries and Notation

We briefly review some basic concepts about equations, $\mathcal{E}$-unification and the IFLP framework. Let S be a set (of *sorts*[4], also called *types*). An S-sorted signature Σ is an $S^* \times S$-sorted family $\langle \Sigma_{w,s} | w \in S^*, s \in S \rangle$. $f \in \Sigma_{w,s}$ is a function symbol of arity w and type s; the arity of a function symbol expresses what sorts of data it expects to see as inputs and in what order, and the sort of a function symbol expresses the type of data it returns. Given $f \in \Sigma_{(s_1,\ldots,s_n),s}$, we define the following auxiliary concepts:

- $Arity(f) = n$.
- $Arity^t(f)$ denotes the times that the type t is in the string $(s_1, \ldots, s_n)$.
- $Type(f) = s$.
- $Set\_Type(f) = \{s_i, 1 \leq i \leq n\}$

Also, we consider Σ as the disjoint union $\Sigma = \mathcal{C} \uplus \mathcal{F}$ of symbols $c \in \mathcal{C}$, called *constructors*, and symbols $f \in \mathcal{F}$, called *defined functions*. Let $\mathcal{X}$ be a countably infinite set of *variables*. Then $\mathcal{T}(\Sigma, \mathcal{X})$ denotes the set of *terms* built from Σ and $\mathcal{X}$, and $\mathcal{T}(\mathcal{C}, \mathcal{X})$ is the set of constructor terms. We extend the $Type$ function defined above to (sub)terms in a natural way: given a term $f(t_1, \ldots, t_n)$ where $f \in \Sigma_{(s_1,\ldots,s_n),s}$, then $Type(t_i) = s_i$. The set of variables occurring in a term t is denoted $Var(t)$. A term t is a *ground term* if $Var(t) = \emptyset$. Given a set of syntactic objects (like terms, variables or function symbols) $O = \{O_1, \ldots, O_n\}$, we define $O^s = \{O_i | Type(O_i) = s\}$. If S is a set of types, then $O^S = \bigcup_{s \in S} O^s$. A *substitution* is defined as a mapping from the set of variables $\mathcal{X}$ into the set of terms $\mathcal{T}(\Sigma, \mathcal{X})$. An equation is an expression of the form $l = r$ where l and r are terms. l is called the left hand side (lhs) of the equation and r is the right hand side (rhs). An equational theory $\mathcal{E}$ (which we call *program*) is a finite set of equational clauses of the form $l = r \Leftarrow e_1, \ldots, e_n$ with $n \geq 0$ where e_i is an equation, $1 \leq i \leq n$. The theory (and the clauses) are called *conditional* if $n > 0$ and *unconditional* if $n = 0$. In that follows, a program P defines a signature Σ_P which is composed by the function symbols that appear in P. An $\mathcal{E}$-unification algorithm defines a procedure for solving an equation $t = s$ within the theory $\mathcal{E}$. Narrowing is a sound and complete method for solving equations. For this reason, it is widely accepted that narrowing is the key to describing operational semantics of functional logic programming languages [5, 6].

IFLP [4] can be defined as the functional (or equational) extension of ILP. The goal is the inference of a theory (a functional logic program P) from evidence (a set of positive and optionally negative ground equations E) using a

[3] Appendices are also included to illustrate the application of our method. In the final version, this will be replaced by a reference to a web page, containing this and other examples, as well as some experimental results.

[4] A sort is a name for a set of objects.

background knowledge theory (a functional logic program B). The rhs of the equations in E are normalised w.r.t. the background theory B and the theory P which is meant to be discovered (target hypothesis). In [4] an approximation to the induction of unconditional functional logic programs has been presented based on two operators: *Consistent Restricted Generalisation* (CRG) and *Inverse Narrowing*. Informally, a CRG of an equation e is another equation which generalises e without introducing extra variables on the rhs and which is consistent with the evidence.

3 Constructor-Based Decision Tree Learning

The problem of learning decision tree classifiers lies in generating a decision-tree from a set of cases, each described by a vector of attribute values, and construct from them a mapping from attribute values to classes. The attributes can be continuous or discrete, whereas we consider that the class may have only discrete values.

In particular, the most popular systems for learning decision-trees, ID3/C4.5 [13, 15, 16] perform the following tests in order to split a tree node:

- If the attribute A is of a discrete type T, the node has as many outcomes (children) as possible values v_i of T, with condition $A = v_i$.
- If the attribute A is of a continuous type T, the node has two outcomes (children), $A < t$ and $A \geq t$ where t is a threshold that maximises the splitting criterion.

With these possible splits and a splitting criterion (which tells which tests are chosen first) a propositional decision tree learning can be constructed from here. In general, decision trees represent a disjunction of conditions over the attribute values. Hence, learned trees can also be represented as sets of conditional rules to improve their readability.

Some ILP systems, like FOIL [14] or TILDE [1], allow to introduce recursion in the definition tree by flattening the predicates, however they do not permit to deal with problems with (semi)-structured data based on constructors.

Given a node ν of a decision tree, we denote its set of conditions by C_ν. The Boolean function $leaf(\nu)$ returns true if ν is a leaf of the tree; $\Pi(\nu)$ denotes the partition in ν; $child_i(\nu)$ represents the *i-th* child of ν and $range(\nu)$ is the number of children of ν. Let us denote with E_ν the set of examples which are consistent with the conditions C_ν. The open variables OV_ν of ν are the variables that do not appear in the lhs of an equality of a condition of C_ν (note that disequalities). Let us denote with OVR_ν the set of variables of real type in OV_ν. $OVNR_\nu = OV_\nu - OVR_\nu$. With τ_i we denote each different type of the set $OVNR_\nu$.

First of all, let us define the possible types we are going to deal within our approach. Table 1 shows these types.

Next, let us define in Table 2 the possible partitions (splits) that can be done according to the types (consider an equation $f(X_1, .., X_{n-1}) = X_n$ and $(1 \leq i \leq n)$). Note that in IFLP the attribute tests are expressed as equations.

Kind	Description	Type Example	Attribute Ex.	Order
UDF	Unordered Discrete Finite	$\{red, green, blue\}$	$green$	-
UDI	Unordered Discrete Infinite	$Lists$	$cons(\lambda, a)$	-
ODF	Ordered Discrete Finite	$\{low, med, high\}$	low	$<$
ODI	Ordered Discrete Infinite	$Naturals$	4	$<$
C	(Ordered) Continuous (Infinite)	$[0.0 \ldots 360.0]$	47.34	$<$
U	Undefined Type	-	-	-
R	Restricted (Dummy) Type [5]	Elements of a List	a	-

Table 1. Kind of types allowed

#	Partition on Attribute X_i (Split)	Kinds of Types Applicable
1	$X_i = a_1 \mid X_i = a_2 \mid \ldots \mid X_i = a_k$	Finite (udf,odf)
2	$X_i = c_0 \mid \ldots \mid X_i = c_k(Y_1, \ldots, Y_{k_m})$	Constructor based (udi) and (u)
3	$X_i < t \mid X_i \geq t$[6]where t is a threshold	Ordered (odf,odi) and Continuous (c)
4	$X_i = Y$ where $Y \in \{X_1, \ldots, X_n\}$ and $Y \neq X_i$	Discrete (udf,udi,odf,odi) and (u,r)
5	$X_i = a \mid X_i \neq a$[7]	udf,odf,odi,c,u
6	$a_1 = f(Y_1, \ldots, Y_n) \mid \ldots \mid a_n = f(Y_1, \ldots, Y_n)$ where $\exists! Y_i \in \{X_1, \ldots, X_n\}$	all
7	$X_i = f(Y_1, \ldots, Y_n)$	all
8	$a_1 = g(Y_1, \ldots, Y_n) \mid \ldots \mid a_n = g(Y_1, \ldots, Y_n) \ldots$ where $\exists! Y_i \in \{X_1, \ldots, X_n\}$ and $g \in \Sigma_B$	all
9	$X_i = g(Y_1, \ldots, Y_n)$ where $g \in \Sigma_B$	all

Table 2. Splits allowed

All partitions, except 4, 7 and 9, split the example set in disjoint subsets. Note that partitions 4, 7 and 9 have only a child. Since the variables in conditions are existentially quantified, different children might give a non-disjoint partition.

These kinds of partitions entail an extension of ID3, FOIL and our old CRG method, as is shown in Table 3.

4 Guiding the Search

The previous partitions have to be well handled in order to convert such a variety of conditions and constructors into something useful. The adaptation of classical splitting criteria, such as those used by C4.5 / FOIL or CART would not be suitable, because these measurements are devised for rewarding partitions which discriminate well the class of the result, be it a predicate or a function. However, this may be misleading for recursive functions where this recursive call appears directly on the rhs of a rule; for instance, the lhs of the rule $sum(X, s(Y)) = s(sum(X, Y))$ does not make any bias on the distribution of the result of the function sum. Consequently, a criterion based on this dis-

[4] This type is useful for avoiding partitions on types which are irrelevant.

[5] $X_i < t$ represents the equation $<(X_i, t) = true$ being $<$ a built-in function.

[7] There are extensions of the operational mechanism of the functional logic programming able to handle disequalities under some constraints, as in [7, 18].

#	ID3	FOIL	CRG	CDTL
1	×	×	×	×
2	-	-	×	×
3	×	×	-	×
4	-	×	×	×
5	-	-	-	×
6	-	×	-	×
7	-	-	-	×
8	-	×	-	×
9	-	-	-	×

Table 3. Comparison between splits allowed by CDTL w.r.t. ID3, FOIL, and CRG methods

crimination or in a distribution change would never select the partitions which are needed to generate the previous rule.

As we have said in the introduction, one proper way to order the search space is by the description length of the hypothesis. In some way, a top-down construction of a decision tree is, by definition, short-to-long, in the way that it adds conditions, and after a partition is made, the tree (and hence the hypothesis) is longer to describe. However, this is not sufficient. The idea is to devise a splitting criterion such that partitions that presumably lead to shorter trees should be selected first.

There exists a suitable paradigm for performing this search: the MDL principle. If we assume $P(h) = 2^{-K(h)}$ where $K(\cdot)$ is the descriptional (Kolmogorov) complexity of h, and $P(E|h) = 2^{-K(E|h)}$, we can obtain the so-called maximum a posteriori (MAP) hypothesis as follows [9].

$$h_{MAP} = argmax_{h \in H} P(h|E) = argmin_{h \in H}(K(h) + K(E|h))$$

This last expression is the MDL principle, which means that the best hypothesis is the one which minimises the sum of description of the hypothesis and the description of the evidence using the hypothesis. The MDL principle has also been justified for first-order theories (see e.g. [3]).

The idea is that the two terms to be minimised in the expression of the MDL principle nicely represent the source and the sink of the information during the construction of a decision tree. Initially, when there is only the root of the tree, i.e. the function profile $f(X_1, X_2, ..., X_{n-1}) = X_n$, the length of the description of the hypothesis $K(h)$ is almost zero (just the profile) while the description of the data $K(E|h)$ is large, because the instance for each variable has to be provided for each example, since no pattern of the evidence has been captured yet. In the end of the construction of the decision tree, the description of the hypothesis $K(h)$ may be large, and each branch constitutes a rule of the program, while the description of the data by using the tree, i.e. $K(E|h)$, will have been reduced considerably. If the resulting tree is good, the term $K(h) + K(E|h)$ is smaller than initially.

The way to construct the tree is to select those partitions (whose description will swell the $K(h)$ part) such that the $K(E|h)$ is reduced considerably. In other

words, the best partition will be the one which minimises the term $K(h)+K(E|h)$ after the partition.

In what follows we will introduce an approximation for $K(E|h)$ and then an approximation for $K(h)$.

4.1 Estimate for $K(E|h)$

Given a node ν a table T_i is constructed for each different type τ_i of the set $OVNR_\nu$. The table contains an entry for each different term of type τ_i appeared in the evidence. With $|T_i|$ we denote the number of elements in table T_i. Each term is denoted by $term_{i,j}$, with j ranging from 1 to $|T_i|$. The information required for each entry in the table $Info(term_{i,j})$ is defined in the following way:

- for finite discrete types, $Info(term_{i,j}) = \log n$ with n being the number of possible values of the type[8].
- for constructor-based types, $Info(term_{i,j})$ is defined as the cost in bits of selecting the appropriate constructors (from the possible set of constructors applicable in each moment) to describe the term. For instance, given the list of naturals $[s(0), s(s(s(0))), 0]$, the information of this term is $\log 2$ (for choosing between the empty list or the constructor $\cdot(\_,\_)$) + $\log 2$ (for choosing between 0 and $s(\_)$) + $\log 2$ (for choosing between 0 and $s(\_)$ again) + $\log 2$ (for choosing between the empty list or the constructor $\cdot(\_,\_)$) + ... In the end, 4 (the list has three elements) + 2 for the first natural, 4 for the second one and 1 for the last natural, i.e. 11 bits.

From here, we can define the information which is required to describe the table:

$$InfoTable(T_i) = \log |T_i| + \sum_{j=1..|T_i|} Info(term_{i,j})$$

Note that a table is constructed for each different type, not for each different non-real open variable ($OVNR$).

Next we have to give a definition for the information required to give values to all the open variables in order to describe an example. Given an example e from E_ν, we have to code in a different way the substitution for non-real variables (just referring to the position in its corresponding table) and for real (continuous) variables. Let us denote the argument k of the lhs of example e of real type as $RealValue_k(e)$, and let us denote with $int(r)$ the integer part of r and $rat(r)$ the digits of the non-integer part of r in binary representation. We approximate the cost in bits for describing the real number as $InfoR(r) = 1 + \log int(r) + rat(r)$. The first bit is for the sign. For instance, the number 36.25, which can be represented as 10010.01 in binary notation has $InfoR = 1 + \log 36 + 2 = 8.17$ bits. Type ODI (e.g. integers) is made in a similar way (but only for the integer part). Another possibility would be the use of a constant value for real numbers described as floating-point format.

[8] All logarithms in this paper are binary logarithms.

#	Partition on Attribute X (Split)	Info
1	$X_i = a_1 \mid ... \mid X_i = a_k$	$\log OV_\nu$
2	$X_i = c_0 \mid ... \mid X_i = c_k(Y_l, ..., Y_{mi})$	$\log OV_\nu$
3	$X_i < t \mid X_i \geq t$	$\log OV_\nu + InfoR(t)$
4	$X_i = X_j i \neq j$	$1^* + \log OV_\nu + \log(OV_\nu - 1)$
5	$X_i = a \mid X_i \neq a$	$\log OV_\nu + \log TCard(Type(X_i))$
6	$a_1 = f(...X_i...) \mid ... \mid a_k = f(...X_i...)$	$1 + min[\log OV_\nu + \log Arity^{Type(X_i)}(f),$ $\log OV_\nu^{Type(f)} + \log Arity(f)]$
7	$X_i = f(Y_1, ..., Y_n)$	$1^* + \log OV_\nu^{SetType(f)}$
8	$a_1 = g(...X_i...) \mid ... \mid a_k = g(...X_i...)$	$1 + min[\log OV_\nu + \log Arity^{Type(X_i)}(g),$ $\log OV_\nu^{Type(g)} + \log Arity(g)] + \log \lvert\Sigma_B^\tau\rvert$
9	$X_i = g(Y_1, ..., Y_n)$	$1^* + 1 + min[\log OV_\nu + \log \lvert\Sigma_B^{Type(X_i)}\rvert ,$ $\log OV_\nu^{Type(g)} + \log \lvert\Sigma_B^\tau\rvert]$

Table 4. Information corresponding to splits allowed

From here, we can define the information which is required to code an example, given the node ν and using the tables, as:

$$Info(e|\nu) = \sum_{k \in OVNR_\nu} \log |T_i| + \sum_{k \in OVR_\nu} InfoR(RealValue_k(e))$$

Note that the first part is the same for all the examples.

And finally, we can define the cost of coding the whole set of examples that fall under ν, i.e. E_ν, as:

$$Info(E_\nu|\nu) = |E_\nu| + \sum_i InfoTable(T_i) + \sum_{e \in E_\nu} Info(e|\nu)$$

The first term codes the number of examples that will be described. An example is included in the appendices.
With this, we have an estimate for the second term of the definition of the MDL principle, $K(h) + K(E|h)$. In this way $K(E|h) \approx \sum_{\nu \in leaves(h)} Info(E_\nu|\nu)$.

4.2 Estimate for $K(h)$

Table 4 includes the cost in bits of each partition at a node ν, denoted by $InfoP(P, \nu)$. Note that unary partitions 4, 7 and 9 have an extra bit, denoted by 1^*. From this table, given a partition P, its information can be obtained as:

$$InfoPart(P, \nu) = \log 10 + InfoP(P, \nu)$$

the first term is used to select the partition from the 9 possible partitions (the tenth option corresponds to no split, i.e. a leaf). Note that leaves have also $InfoPart$, which is equal to $\log 10$, because it has to be said that the node is not further exploited.

And $K(h)$ can be estimated as follows:

$$K(h) \approx InfoHead(h) + \sum_{\nu \in nodes(h)} InfoPart(\pi(\nu), \nu)$$

where $InfoHead(h)$ captures the information which is required to code the profile of the function to be learned (which must include the arity and types of the function).

4.3 Information of the Tree

Finally, we have to introduce the information for the whole tree. Given a tree T with root node ν, the info of T is defined as:

$$InfoTree(E, T) = InfoHead(T) + InfoN(E, \nu)$$

where $InfoN(E, \nu)$ is obtained recursively in the following way:

$$InfoN(E, \nu) = \begin{cases} Info(E_\nu|\nu) & \text{if } leaf(\nu) \\ InfoPart(\pi(\nu)) + \sum_{i=1..range(\nu)} InfoN(E, child_i(\nu)) & \text{otherwise} \end{cases}$$

Initially, the tree with just one node has only information about the function profile. Since this node is a leaf, $InfoTree(E, T) = InfoHead(T) + Info(E_\nu|\nu)$. When the tree is being constructed, any exploited node ν (which is still a leaf) has an approximate value for the information, given by $Info(E_\nu|\nu)$, independently of the possible partitions that there could be underneath. But when this node is exploited, then the value is substituted by the sum of the information of the partition and the information required for the children subtrees.

Since the first approximations are useful when populating and pruning the tree, we will denote by $InfoN_d(E, \nu)$ the $InfoN$ of a node up to depth d. Obviously $InfoN_\infty(E, \nu) = InfoN(E, \nu)$ whereas $InfoN_0(E, \nu) = Info(E_\nu|\nu)$.

4.4 Constructing the Multi-Tree

Now we can establish the way in which one tree is constructed from the root $f(X_1, X_2, ..., X_{n-1}) = X_n$, which is an open node:

- NODE SELECTION CRITERION: From all the open nodes, we select the node with less $InfoN_0(E, \nu)$, i.e. with less $Info(E_\nu|\nu)$. This usually corresponds to nodes with less examples.
- SPLIT SELECTION CRITERION: The $InfoN_1(\nu)$ of the node ν is determined for every possible split. The split with less $InfoN_1$ is selected. Its children are new open nodes.
- STOPPING CRITERION: a node is closed when the class is consistent with all the examples that fall into that node (i.e. E_ν). The generation of the tree stops when all nodes are closed.

As can be seen in the examples in the appendices, these criteria usually lead to short trees, and in many propositional problems, the first tree which is generated matches the tree which is generated by other selection criteria based on gains.

However, in some cases, as any greedy method, an initial choice could make that the resulting tree is not optimal. Nonetheless, the contrary strategy, to explore all the possible splits, would be unfeasible. A trade-off can be found though. After a tree is finished, we will explore second-best splits up to dedicated space-time resources. In other words, *CDTL populates the tree up to a limit number of nodes or time*. For this we have to establish a new criterion:

- TREE SELECTION CRITERION: Consider the set of possible splits at depth 1 of a node ν, where σ_1 is the best split and σ_k is the best split that has not been exploited yet. Let us denote with ν_1 the node with split σ_1. Let us denote with ν_k the node with split σ_k. We define the 'rival ratio' $\rho(\nu) = InfoN_1(\nu_1)/InfoN_1(\nu_k)$. Once a tree has been constructed completely, a next tree can be explored, beginning with the split with more rival ratio. This next tree has to be fully completed before selecting another tree.

The result of this process is a **multi-tree** rather than a tree. Note that the use of information approximations at depth 1 avoids the overexploitment of splits at the bottom of the tree, where the information is usually reduced when good selections have been made. This generates a wider variety of trees than would be generated if the real value (depth ∞) were used instead.

The use of a multi-tree which is populated in a short-to-long way has some advantages:

- Each set of leaves of a solution is a program that covers the evidence.
- The memory resources and time invested can be better adjusted.
- The procedure generates several solutions.

Moreover, these hypotheses could be combined to give more precise predictions. According to the MAP hypothesis, we have by Bayes theorem:

$$P(h|E) = \frac{P(E|h)P(h)}{P(E)}$$

We can use the previous estimates for the three a priori probabilities:

$$P(E|h) = 2^{-K(E|h)} \approx 2^{-\sum_{\nu \in leaves(h)} Info(E_\nu|\nu)}$$

$$P(h) = 2^{-K(h)} \approx 2^{-[InfoHead(h)+\sum_{\nu \in nodes(h)} InfoPart(\pi(\nu),\nu)]} = 2^{-InfoTree(\emptyset|h)}$$

$$P(E) = P(E|h_0) = 2^{-K(E|h_0)} \approx 2^{-InfoTree(E|h_0)}$$

where h_0 represents the empty hypothesis (in our case the tree with one node, the function profile). This gives an approximation for the a posteriori probability. The use of $P(h|E)$ allows the combination of the different hypotheses provided by a multi-tree to give more accurate predictions (e.g. through weighted majority) and to give estimates for the reliability of a prediction.

5 Coding Exceptions of Unary Partitions

As said, some partitions have only one child, in particular the partitions 4, 7 and 9. The reason is that the complementary sibling node would not be evaluable under the usual semantics of functional logic programs. With this we avoid the generation of inappropriate programs.

However, the existence of just one child of a node ν which carries along with an extra condition can involve the fact that some examples are left uncovered (denoted by $E_{u,\nu}$), and the tree would not describe the entire evidence. These cases should not be avoided, because some unary partitions can be good for describing most of the evidence. Moreover, case 7, which is evaluated w.r.t. the evidence, could, in principle, leave uncovered examples, but when the base rules of the programs are discovered, these uncovered examples could fall into these nodes.

The idea is to code these exceptions. For this, partitions 4, 7 and 9 have an extra bit. If this bit is set to one, then the uncovered examples of the partition $E_{u,\nu}$ will be coded by using the parent node ν and its conditions by the previous equation, i.e.

$$Info(E_{u,\nu}|\nu) = |E_{u,\nu}| + \sum_i InfoTable(T_i) + \sum_{e \in E_{u,\nu}} Info(e|\nu)$$

Background knowledge is evaluated intensionally. Recursive calls to the function which is being learned are treated both extensionally and intensionally. Extensional evaluation uses the entire evidence (such as FOIL); the $Info$ is obtained by using the substitutions for the example which minimises the information. For the intensional evaluation the closed nodes can already be used (in a similar way as FOIDL [12]). The node selection criterion favours the appearance of base cases first. Note that this evaluation is limited and can leave some examples uncovered. Nonetheless, when the tree is completed, some uncovered examples could be recovered.

Appendix A contains examples of the learning process of simple problems by the CDTL mechanism.

6 Conclusions

In this paper we have presented a new splitting criterion, a node selection criterion, and a tree selection criterion, all of them based on the MDL principle. These criteria allow to deal with partitions with constructor-based types. This extends the applicability of decision tree learning to semi-structured data.

The MDL-guided search has three important advantages. First, MDL sorts out the search space providing an optimal use of computational resources. Secondly, it gives an estimate for $P(E|h)$, which allows a weighted combination of the different hypotheses obtained from the multi-tree. Last, the annotated *Infos* can be used directly to prune the trees for noisy evidence.

As future work, the expressiveness of the approach could be extended by allowing real numbers in the class, i.e. the result of the function. For this, the cost should be modified in order to code the description of the error in such a way that would be increasing for growing quadratic error.

References

1. H. Blockeel. Top-down Induction of First Order Logical Decision Trees. *AI Communications*, 12:119–20, 1999.
2. Leo Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth Publishing Company, 1984.
3. D. Conklin and I.H. Witten. Complexity-Based Induction. *Machine Learning*, 16:203–225, 1994.
4. J. Hernández and M.J. Ramírez. A Strong Complete Schema for Inductive Functional Logic Programming. In *Proc. of the Ninth International Workshop on Inductive Logic Programming, ILP'99*, volume 1634 of *LNAI*, pages 116–127, 1999.
5. S. Hölldobler. Equational Logic Programming. In *Proc. Second IEEE Symp. on Logic In Computer Science*, pages 335–346. IEEE Computer Society Press, 1987.
6. J. Jaffar, J.-L. Lassez, and M.J. Maher. A Logic Programming Language Scheme. In D. de Groot and G. Lindstrom, editors, *Logic Programming, Functions, Relations and Equations*, pages 441–468. Prentice Hall, Englewood Cliffs, NJ, 1986.
7. H. Kuchen, F. Lopez-Fraguas, J. J. Moreno-Navarro, and M. Rodriguez-Artalejo. Implementing a Lazy Functional Logic Language with Disequality Constraints. In *Joint Int. Conf. and Symp. on Logic Prog.*, pages 207–224. MIT Press, 1992.
8. L.A. Levin. Universal Search Problems. *Problems Inform. Transmission*, 9:265–266, 1973.
9. M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. 2nd Ed. Springer-Verlag, 1997.
10. M. Mehta, J. Rissanen, and R. Agrawal. MDL-Based Decision Tree Pruning. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining (KDD'95)*, pages 216–221, 1995.
11. T. M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
12. R. J. Mooney and M. E. Califf. Induction of First–Order Decision Lists: Results on Learning the Past Tense of English verbs. *J. of A.I. Research*, 3:1–24, 1995.
13. J. R. Quinlan. Induction of Decision Trees. In *Readings in Machine Learning*. Morgan Kaufmann, 1990.
14. J. R. Quinlan. Learning Logical Definitions from Relations. *Machine Learning*, 5(3):239–266, 1990.
15. J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.
16. J. R. Quinlan. Improved Use of Continuous Attributes in C4.5. *Journal of Artificial Intelligence*, 4:77–90, 1996.
17. J. R. Quinlan and R. L. Rivest. Inferring Decision Trees Using The Minimum Description Length Principle. *Information and Computation*, 80:227=–248, 1989.
18. M. J. Ramirez and M. Falaschi. Conditional Narrowing with Constructive Negation. *Lecture Notes in Computer Science*, 660, 1993.
19. Jürgen Schmidhuber, Jieyu Zhao, and Marco Wiering. Shifting Inductive Bias with Success-Story Algorithm, Adaptive Levin Search, and Incremental Self-Improvement. *Machine Learning*, 28:105–130, 1997.

A Examples

A.1 Example of $Info(e|\nu)$

Let us consider the node $\nu \equiv f(X, Y, Z, W, V) = R$ with associated conditions $Y = X$ and $V = d$, where X, Y and W are of type list of naturals, Z is a real and R can be one of the following three discrete values $\{a, b, c\}$. Given the initial entire evidence E:

$e_1 \equiv f([s(0)], [s(0)], 3.8, [], e) = a$
$e_2 \equiv f([0, s(0)], [0, s(0)], -2.75, [0], d) = b$
$e_3 \equiv f([], [], -3.25, [s(s(0)), 0, 0], d) = a$
$e_4 \equiv f([s(0), 0], [s(0), 0], -3.125, [], d) = a$
$e_5 \equiv f([], [s(0), 0], -85.23, [0], e) = c$
$e_6 \equiv f([0], [0], 1.5, [s(0), s(0)], d) = b$

we have that $E_\nu = \{e_2, e_3, e_4, e_6\}$ and $OV_\nu = \{X, Z, W, R\}$ with $OVR_\nu = \{Z\}$. Since there are two different non-real types (the type of X and W, and the type of R), we have to construct two tables. The table for the list of naturals must have the elements $\{[], [0], [s(0)], [0, s(0)], [s(0), 0], [s(s(0)), 0, 0], [s(0), s(0)]\}$.

Its $InfoTable = \log 7 + 1 + 3 + 4 + 6 + 6 + 9 + 7 = 38.81$ bits.

The table for the type of R must have the elements $\{a, b\}$ (note that c does not appear in E_ν). Consequently its $InfoTable = \log 2 + \log 3 + \log 3 = 4.17$ bits.

The Info of each example is:

$Info(e_2|\nu) = \log 7 + 5 + \log 7 + \log 2 = 11.61$ bits.
$Info(e_3|\nu) = \log 7 + 5.58 + \log 7 + \log 2 = 12.19$ bits.
$Info(e_4|\nu) = \log 7 + 6.58 + \log 7 + \log 2 = 13.19$ bits.
$Info(e_6|\nu) = \log 7 + 1 + \log 7 + \log 2 = 7.61$ bits.

Finally, we have:

$Info(E_\nu|\nu) = 4 + 38.81 + 4.17 + 11.61 + 12.19 + 13.19 + 7.61 = 92.11$ bits.

A.2 The playtennis problem

Let us begin the illustration of the previous procedure with the most classical example of decision trees [11]. The purpose is to classify the days in which the conditions for playing tennis are good or not. The evidence is:

$$E = \left\{ \begin{array}{l} e_1 : playtennis(overcast, hot, high, weak) = yes \\ e_2 : playtennis(rain, mild, high, weak) = yes \\ e_3 : playtennis(rain, cool, normal, weak) = yes \\ e_4 : playtennis(overcast, cool, normal, strong) = yes \\ e_5 : playtennis(sunny, cool, normal, weak) = yes \\ e_6 : playtennis(rain, mild, normal, weak) = yes \\ e_7 : playtennis(sunny, mild, normal, strong) = yes \\ e_8 : playtennis(overcast, mild, high, strong) = yes \\ e_9 : playtennis(overcast, hot, normal, weak) = yes \\ e_{10} : playtennis(sunny, hot, high, weak) = no \\ e_{11} : playtennis(sunny, hot, high, strong) = no \\ e_{12} : playtennis(rain, cool, normal, strong) = no \\ e_{13} : playtennis(sunny, mild, high, weak) = no \\ e_{14} : playtennis(rain, mild, high, strong) = no \end{array} \right\}$$

Initially, with this set of examples we need the following bits to describe it:

$$InfoN_0(E|\nu) = |E| + \sum_i InfoTable(T_i) + \sum_{e \in E} Info(e)$$

$$InfoN_0(E|\nu) = 14 + 21.68 + 86.38 = 122.06 \ bits$$

Consider we select the partition 5 with the split $\nu_1 : X = overcast \mid \nu_2 : X \neq overcast$. In this case ν_1 covers 4 examples and ν_2 covers 10, thus:

$$Info(E_{\nu_1}|\nu_1) = [4 + 12.34 + 14.34] = 30.68 \ bits$$

$$Info(E_{\nu_2}|\nu_2) = [10 + 19.51 + 55.85] = 85.36 \ bits$$

$$InfoPart(\nu) = 5.64 \ bits$$

$$InfoN_1(E_\nu|\nu) = 121.68 \ bits$$

Using the partition $\nu_1 : X = hot \mid \nu_2 : X \neq hot$ coverage=(4,10):

$$Info(E_{\nu_1}|\nu_1) = [4 + 13.17 + 16] = 33.17 \ bits$$

$$Info(E_{\nu_2}|\nu_2) = [10 + 19.51 + 58.5] = 87.65 \ bits$$

$$InfoPart(\nu) = 5.64 \ bits$$

$$InfoN_1(E_\nu|\nu) = 126.46 \ bits$$

And with the partition $\nu_1 : X = overcast \mid \nu_2 : X = sunny \mid \nu_3 : X = rain$, coverage=(4,5,5):

$$Info(E_{\nu_1}|\nu_1) = [12.34 + 14.34 + 4] = 30.68 \ bits$$

$$Info(E_{\nu_2}|\nu_2) = [5 + 13.17 + 20] = 38.17 \ bits$$

$$Info(E_{\nu_3}|\nu_3) = [5 + 13.17 + 20] = 38.17 \ bits$$

$$InfoPart(\nu) = 5.64 \ bits$$

$$InfoN_1(E_\nu|\nu) = 112.66 \ bits$$

The last partition needs less bits, so it is selected to open the tree. The partition ν_1 has only one value for the output class, thus it is a leaf. We continue by unfolding ν_2 with the partition $\nu_{21} : Y = normal \mid \nu_{22} : Y = high$, coverage=(2,3):

$$Info(E_{\nu_{21}}|\nu_{21}) = [2 + 7.17 + 4] = 13.17 \ bits$$

$$Info(E_{\nu_{22}}|\nu_{22}) = [3 + 7.17 + 6] = 16.17 \ bits$$

$$InfoPart(\nu_2) = 5.32 \ bits$$

$$InfoN_1(E_{\nu_2}|\nu_2) = 36.00 \ bits$$

As there is only one value for the output class in the two branches, these are leaves. Finally, we select the partition $\nu_{31} : Z = weak \mid \nu_{32} : Z = strong$ for the branch ν_3, coverage=(3,2):

$$Info(E_{\nu_{31}}|\nu_{31}) = [3 + 7.17 + 6] = 16.17$$

$$Info(E_{\nu_{32}}|\nu_{32}) = [2 + 7.17 + 4] = 13.17$$

$$InfoPart(\nu_3) = 5.32\ bits$$

$$InfoN_1(E_{\nu_3}|\nu_3) = 36.00\ bits$$

This pair of branches have also only one value for the output class, thus the tree is finished. Here you can see the final tree:

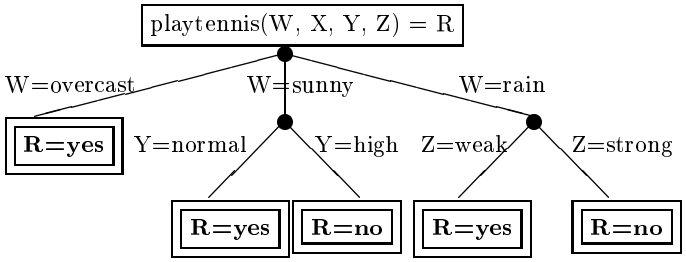

Fig. 0. Playtennis problem. Decision Tree Generated by the CDTL mechanism

The rules corresponding to this tree are:

$playtennis(overcast, X, Y, Z) = yes$
$playtennis(sunny, X, normal, Z) = yes$
$playtennis(rain, X, Y, weak) = yes$
$playtennis(sunny, X, high, Z) = no$
$playtennis(rain, X, Y, strong) = no$

A.3 The sum problem

The following example shows how the CDTL mechanism can be used to learn recursive problems and it can deal with complex data based on constructors. The goal is to learn a program for the addition of two natural numbers represented with the constructors 0 and $s(\_)$.

Consider the following evidence:

$$E = \left\{ \begin{array}{l} e_1 : sum(0,0) = 0 \\ e_2 : sum(s(0), s(0) = s(s(0)) \\ e_3 : sum(0, s(0) = s(0) \\ e_4 : sum(s(0), 0) = s(0) \\ e_5 : sum(s(s(0)), s(0)) = s(s(s(0))) \\ e_6 : sum(0, s(s(0))) = s(s(0)) \\ e_7 : sum(s(s(0)), 0) = s(s(0)) \\ e_8 : sum(s(s(0)), s(s(0))) = s(s(s(s(0)))) \\ e_9 : sum(s(s(s(0))), s(0) = s(s(s(s(0)))) \\ e_{10} : sum(0, s(s(s(s(0))))) = s(s(s(s(0)))) \end{array} \right\}$$

The root node is $sum(X,Y) = R$. Initially, with this set of examples we need the following bits to describe it:

$$InfoN_0(E|\nu) = 10 + 17.32 + 69.66 = 96.98\ bits$$

Consider that we select the partition 2 with the split $\nu_1 : X = 0 \mid \nu_2 : X \neq s(X')$. In this case ν_1 covers 4 examples and ν_2 covers 6:

$$Info(E_{\nu_1}|\nu_1) = [4 + 13 + 16] = 33\ bits$$

$$Info(E_{\nu_2}|\nu_2) = [6 + 17 + 41.8] = 64.8\ bits$$

$$InfoPart(\nu) = 4.9\ bits$$

$$InfoN_1(E_\nu|\nu) = 102.7\ bits$$

If we select the partition 4 with the split $\nu_1 : X = Y \mid \nu_2 : X \neq Y$. In this case ν_1 covers 3 examples and ν_2 covers 7:

$$Info(E_{\nu_1}|\nu_1) = [3 + 13 + 12] = 28\ bits$$

$$Info(E_{\nu_2}|\nu_2) = [7 + 17.32 + 48.72] = 73.04\ bits$$

$$InfoPart(\nu) = 6.9\ bits$$

$$InfoN_1(E_\nu|\nu) = 107.94\ bits$$

The best partition seems to be $\nu_1 : X = 0 \mid \nu_2 : X \neq s(X')$, so we continue developing ν_1 with $\nu_{11} : Y = R$ which covers the 4 examples, and represents a leaf:

$$Info(E_{\nu_{11}}|\nu_{11}) = [4 + 13 + 8] = 25\ bits$$

$$InfoPart(\nu_1) = 5.32\ bits$$

$$InfoN_1(E_{\nu_1}|\nu_1) = 30.32\ bits$$

And developing ν_2 with $\nu_{21} : R = 0 \mid \nu_{22} : R = s(R')$, we have that the branch ν_{22} covers the 6 examples:

$$Info(E_{\nu_{22}}|\nu_{22}) = [6 + 12 + 36] = 54\ bits$$

$$InfoPart(\nu_2) = 4.91\ bits$$

$$InfoN_1(E_{\nu_2}|\nu_2) = 58.91\ bits$$

Following with ν_{22}, we apply $\nu_{211} : R = sum(A,B)$. CDTL gets the solution with the partitions $\nu_{2111} : A = X'$ and $\nu_{21111} : A = X'$(For space reasons we do not develop these steps). The following figure shows the final tree:

The rules corresponding to this tree are:

$sum(0,Y) = Y$
$sum(s(X),Y) = s(sum(X,Y))$

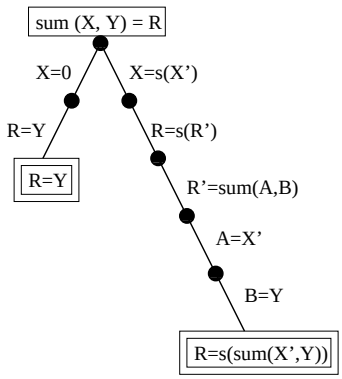

Fig. 1. Tree for the sum problem

A.4 The state problem

The following example shows how the CDTL mechanism can be used to learn from background knowledge and with ordered values. The goal is to learn a program to give the state of an object: liquid, solid or gaseous. The background knowledge is composed of one function, defined in the following way:

$$B = \left\{ \begin{array}{l} aquous(water) = true \\ aquous(lemonade) = true \\ aquous(juice) = true \\ aquous(iron) = false \\ aquous(alcohol) = false \end{array} \right\}$$

We have, as well, the following evidence:

$$E = \left\{ \begin{array}{l} e_1 : state(water, 50) = liquid \\ e_2 : state(lemonade, -30) = solid \\ e_3 : state(juice, 200) = gaseous \\ e_4 : state(water, 15) = liquid \\ e_5 : state(iron, 150) = solid \\ e_6 : state(alcohol, 80) = gaseous \\ e_7 : state(juice, 120) = gaseous \\ e_8 : state(lemonade, -10) = solid \\ e_9 : state(water, -90) = solid \\ e_{10} : state(juice, 60) = liquid \end{array} \right\}$$

Initially, given the root $state(X, Y) = R$ with this set of examples we need the following bits to describe it:

$$InfoN_0(E|\nu) = 10 + 20.27 + 107,34 = 137.61\ bits$$

Consider that we select the partition 8 with the split $\nu_1 : aquous(X) = true \mid \nu_2 : aquous(X) = false$. In this case ν_1 covers 8 examples and ν_2 covers 2:

$$Info(E_{\nu_1}|\nu_1) = [8 + 14.89 + 78.08] = 100.97\ bits$$

$$Info(E_{\nu_2}|\nu_2) = [2 + 8.81 + 19.55] = 30.36\ bits$$

$$InfoPart(\nu) = 3.32\ bits$$

$$InfoN_1(E_\nu|\nu) = 134.65\ bits$$

By partition 3, if we open ν_1 into $\nu_{11} : Y < t \mid \nu_{12} : Y \geq t$, where t is a threshold, we have that the $Info$ is minimised for threshold 2.5. We have coverage = (3,5).

$$Info(E_{\nu_{11}}|\nu_{11}) = [3 + 6.23 + 19.72] = 28.95\ bits$$

$$Info(E_{\nu_{12}}|\nu_{12}) = [5 + 9.81 + 45] = 59.81\ bits$$

$$InfoPart(\nu_1) = 5.9\ bits$$

$$InfoN_1(E_{\nu_1}|\nu_1) = 94.66\ bits$$

The node ν_{11} has only one class ($solid$). The node ν_{12} is further exploited into $\nu_{121} : Y < 90 \mid \nu_{122} : Y \geq 90$.

Finally, node ν_2 splits logically into $\nu_{21} : X = iron \mid \nu_{22} : X = alcohol$ because $(X = water|X = lemonade|X = juice)$ have no associated examples.

The resulting program is:

$$P = \left\{ \begin{array}{l} state(X,Y) = solid \Leftarrow aquous(X) = true, Y < 2.5 \\ state(X,Y) = liquid \Leftarrow aquous(X) = true, Y \geq 2.5, Y < 90 \\ state(X,Y) = solid \Leftarrow aquous(X) = true, Y \geq 90 \\ state(iron,Y) = solid \Leftarrow aquous(X) = false \\ state(alcohol,Y) = gaseous \Leftarrow aquous(X) = false \end{array} \right\}$$