Diterpene Structure Elucidation from ^{13}C NMR Spectra with Inductive Logic Programming

Sašo Džeroski

Department of Intelligent Systems, Jožef Stefan Institute
Jamova 39, 1000 Ljubljana, Slovenia
Email: `Saso.Dzeroski@ijs.si`

Steffen Schulze-Kremer and Karsten R. Heidtke
Max-Planck Institute for Molecular Genetics
Otto-Warburg-Laboratorium, Department Lehrach
Ihnestrasse 73, 14195 Berlin, Germany

Karsten Siems
AnalytiCon GmbH
Gustav-Meyer-Allee 25, 13335 Berlin-Wedding, Germany

Dietrich Wettschereck
GMD, FIT.KI, Schloss Birlinghoven, 53745 Sankt Augustin, Germany

Hendrik Blockeel
Department of Computer Science, Katholieke Universiteit Leuven
Celestijnenlaan 200A, 3001 Heverlee, Belgium

Abstract

We present a novel application of ILP to the problem of diterpene structure elucidation from ^{13}C NMR spectra. Diterpenes are organic compounds of low molecular weight with a skeleton of 20 carbon atoms. They are of significant chemical and commercial interest because of their use as lead compounds in the search for new pharmaceutical effectors. The interpretation of diterpene ^{13}C NMR-spectra normally requires specialists with detailed spectroscopic knowledge and substantial experience in natural products chemistry, more specifically knowledge on peak patterns and chemical structures. Given a database of peak patterns for diterpenes with known structure, we apply several ILP approaches to discover correlations between peak patterns and chemical structure. The approaches used include relational instance based learning, induction of logical decision trees and inductive constraint logic. Performance close to the one of domain experts is achieved, which suffices for practical use.

1 Introduction

1.1 NMR

Structure elucidation of compounds isolated from plants, fungi, bacteria or other organisms is a common problem in natural product chemistry. There are many useful spectroscopic methods of getting information about chemical structures, mainly nuclear magnetic resonance (NMR) and mass spectroscopy. The interpretation of these spectra normally requires specialists with detailed spectroscopic knowledge and experience in natural products chemistry. NMR-spectroscopy is the best method for complete structure elucidation (including stereochemistry) of non-crystalline samples. For structure elucidation of secondary natural products (not proteins) only ^1H -NMR- and ^{13}C -NMR-spectroscopy, including combined methods such as 2D-NMR-spectroscopy, are important because hydrogen and carbon are the most abundant atoms in natural products. In structure elucidation of peptides and proteins ^{15}N -NMR is sometimes helpful [2].

Figure 1: Multiplicity of Carbon Atoms.

^1H -NMR- and ^{13}C -NMR-spectroscopy are quite different: in a ^{13}C -NMR- spectrum every carbon atom occurs as a separate signal in most cases, while in ^1H -NMR-spectra many signals overlap and are therefore difficult to interpret [18]. ^1H -NMR- spectra are logically decomposable and the specialist can get direct information about the structure (including stereochemistry) from the resonance frequency and shape of the signals, provided a sufficiently high resolution of resonance signals can be experimentally achieved. In ordinary ^{13}C -NMR-spectra only the resonance frequency is observed and no signal splitting (decomposition) is possible. However, the absolute value of the resonance frequency provides more information about chemical structure and (important for prediction and database mainte-

nance purposes) is not as sensitive to varying experimental conditions (such as the magnetic strength of the NMR-spectrometer or the type of solvent) as ^1H -NMR-spectroscopy is.

Additional measurements can be used to determine the number of hydrogens directly connected to a particular carbon atom. This number is the so-called multiplicity of the signal: **s** stands for singlet, which means there is no proton (i.e., hydrogen) connected to the carbon; **d** stands for a doublet with one proton connected to the carbon; **t** stands for a triplet with two protons and **q** for a quartet with three protons bound to the carbon atom. Figure 1 shows each of the four situations in the order listed above. Because of the simpler nature of ^{13}C -NMR-data as compared to ^1H -NMR-data, the former are easier to handle and therefore remain the preferred basis for automatic structure elucidation [10].

1.2 Diterpenes

Diterpenes are one of a few fundamental classes of natural products with about 5000 members known [14]. The skeleton of every diterpene contains 20 carbon atoms. Sometimes there are additional groups linked to the diterpene skeleton by an oxygen atom with the possible effect of increasing the carbon atom count to more than 20 per diterpene. About 200 different diterpene skeletons are known so far, but some of them are only represented by one member compound. Most of the diterpenes belong to one of 20 common skeleton types.

The problem of structure elucidation of diterpenes requires knowledge about biosynthesis, the way in which biological organisms synthesize natural products. Not every combination of carbons, protons and oxygens that is feasible from a chemical point of view actually occurs in nature, as biosynthesis in biological organisms is limited to a characteristic subset of chemical reactions that constrain the structure space. Structure elucidation of diterpenes from ^{13}C -NMR-Spectra can be separated into three main stages: 1) identification of residues (ester and/or glycosides), 2) identification of the diterpene skeleton, and 3) arrangement of the residues on the skeleton.

This work deals with the second stage, the identification of the skeleton. A skeleton is a unique connection of carbon atoms each with a specific atom number and, normalized to a pure skeleton molecule without residues, a certain multiplicity (**s**, **d**, **t** or **q**). Figure 2 shows the ^1H -NMR- and ^{13}C -NMR-Spectra of a diterpene belonging to the skeleton type Labdan. The structure of the diterpene is also depicted, giving the arrangement and numbering of the 20 skeletal carbon atoms.

The task we address is to identify the skeleton (type) of diterpenoid compounds, given their ^{13}C -NMR-Spectra that include the multiplicities and the frequencies of the skeleton atoms. This task is usually done manually by human experts with specialized background knowledge on peak patterns and chemical structures. In the process, each of the 20 skeletal atoms is assigned an atom number that corresponds to its proper place in the skeleton and the diterpene is classified into one of the possible skeleton types. We address the problem of learning classification rules from a database of peak patterns for diterpenes with known structure. Recently, propositional methods were successfully applied to classifying spectra with assigned atom numbers [8]. As the assignment of atom numbers is a difficult process in itself (and possibly indistinguishable from the classification process), we apply ILP, i.e., relational learning, to the problem of classifying spectra without assigned atom numbers.

Figure 2: ^1H -NMR (top) and ^{13}C -NMR (bottom) Spectra of a diterpene belonging to the skeleton type Labdan, with acetyl as a residue. At positions 76.681, 76.999, and 77.316 of the latter is the signal of the solvent, the other signals correspond to carbon atoms. The additional measurements (above the ^{13}C -NMR Spectrum) show the multiplicity of carbon atoms.

Table 1: Facts describing the ^{13}C -NMR spectrum of a particular molecule that belongs to skeleton type Labdan (52).

```

nmr(v1,52,v1_1,1,t,39.00).
nmr(v1,52,v1_2,2,t,19.30).
nmr(v1,52,v1_3,3,t,42.20).
nmr(v1,52,v1_4,4,s,33.30).
nmr(v1,52,v1_5,5,d,55.60).
nmr(v1,52,v1_6,6,t,24.30).
nmr(v1,52,v1_7,7,t,38.30).
nmr(v1,52,v1_8,8,s,148.60).
nmr(v1,52,v1_9,9,d,57.20).
nmr(v1,52,v1_10,10,s,39.70).
nmr(v1,52,v1_11,11,t,17.60).
nmr(v1,52,v1_12,12,t,41.40).
nmr(v1,52,v1_13,13,s,73.50).
nmr(v1,52,v1_14,14,d,145.10).
nmr(v1,52,v1_15,15,t,111.40).
nmr(v1,52,v1_16,16,q,27.90).
nmr(v1,52,v1_17,17,t,106.30).
nmr(v1,52,v1_18,18,q,33.60).
nmr(v1,52,v1_19,19,q,21.60).
nmr(v1,52,v1_20,20,q,14.40).

```

The rest of the paper is organized as follows: Section 2 describes the database and the pre-processing of the data, then summarizes recent work on applying propositional learning in this domain that uses data with assigned atom numbers. When the atom numbers are not available, but only multiplicities and frequencies, we have an ILP problem. Section 3 describes several formulations of this problem and gives the results of applying the FOIL ILP system to these formulations. Section 4 gives a comparison to C4.5, nearest neighbor classification and backpropagation networks applied to different propositionalizations of the problem. Section 5 presents the results of applying RIBL, Section 6 the results of applying TILDE, and Section 7 the results of applying ICL to the ILP problem formulations. Finally, Section 8 concludes with a discussion and an outline of directions for further work.

2 The data and propositional learning results

2.1 The database

AnalytiCon GmbH maintains a database of diterpenoid compounds. This is done using the ISIS (Integrated Scientific Information System), a product of MDL Information Systems, Inc., on IBM PCs under MS Windows. The relational database contains information on 1503 diterpenes with known structure, stored in three relations - **atom**, **bond**, and **nmr**.

The first relation specifies to which element an atom in a given compound belongs. The second relation specifies which atoms are bound and in what way in a given compound.

The `nmr` relation stores the measured ^{13}C -NMR-Spectra. For each of the 20 carbon atoms in the diterpene skeleton, it contains the atom number, its multiplicity and frequency. For each compound, the skeleton type which it represents is also specified within the database. Table 1 gives an excerpt from the `nmr` relation describing molecule `v1`. The relation schema is `nmr(MoleculeID, SkeletonType, AtomID, AtomNumber, Multiplicity, Frequency)`.

2.2 Pre-processing

Every substituent or residue connected to a carbon atom exerts a characteristic shift on the resonance frequency signal of the carbon atom and sometimes also changes its multiplicity. Simple rules based on expert knowledge can be used to take this effect into account.

Table 2: Rules for generating reduced multiplicities from observed multiplicities.

IF ObservedM = s AND Frequency in [64.5, 95]	THEN ReducedM = d
IF ObservedM = s AND Frequency in [96, 114]	THEN ReducedM = t
IF ObservedM = s AND Frequency in [115, 165]	THEN ReducedM = d
IF ObservedM = s AND Frequency in [165, 188]	THEN ReducedM = q
IF ObservedM = s AND Frequency in [188, inf]	THEN ReducedM = t
IF ObservedM = d AND Frequency in [64.5, 95]	THEN ReducedM = t
IF ObservedM = d AND Frequency in [105, 180]	THEN ReducedM = t
IF ObservedM = d AND Frequency in [96, 104]	THEN ReducedM = q
IF ObservedM = d AND Frequency in [180, inf]	THEN ReducedM = q
IF ObservedM = t AND Frequency in [59, 90]	THEN ReducedM = q
IF ObservedM = t AND Frequency in [90, inf]	THEN ReducedM = q

They transform the raw, measured, NMR-multiplicities into the so-called reduced multiplicities, which carry more information about the skeleton types, as shown below. These rules are given in Table 2. They leave the measured frequencies unchanged. The reduced multiplicities that correspond to those in Table 1 are given in Table 3. Note that the multiplicities of atoms 8, 13, 14, 15, and 17 are changed.

2.3 Propositional learning: experiments and results

Given the above data, Dzeroski et al. [8] formulate several propositional learning problems. Twenty-three different skeleton types are represented in the whole set of 1503 compounds: there are thus 23 possible class values. The attributes are the multiplicities and frequencies of each of the 20 skeleton atoms. For instance, attribute `A1M` is the multiplicity of atom number 1, and `A7F` is the frequency of atom number 7. There are thus 40 attributes, twenty discrete (multiplicities) and 20 continuous (frequencies). At a suggestion from the domain experts, a version of the problem is also considered where only the multiplicities are used as attributes. Four series of experiments are performed, with the following sets of attributes: observed multiplicities and frequencies, reduced multiplicities and frequencies, observed multiplicities only and reduced multiplicities only.

Table 3: Pre-processed facts describing the ^{13}C -NMR spectrum of molecule v1.

```

red(v1,52,v1_1,1,t,39.00).
red(v1,52,v1_2,2,t,19.30).
red(v1,52,v1_3,3,t,42.20).
red(v1,52,v1_4,4,s,33.30).
red(v1,52,v1_5,5,d,55.60).
red(v1,52,v1_6,6,t,24.30).
red(v1,52,v1_7,7,t,38.30).
red(v1,52,v1_8,8,d,148.60).
red(v1,52,v1_9,9,d,57.20).
red(v1,52,v1_10,10,s,39.70).
red(v1,52,v1_11,11,t,17.60).
red(v1,52,v1_12,12,t,41.40).
red(v1,52,v1_13,13,d,73.50).
red(v1,52,v1_14,14,t,145.10).
red(v1,52,v1_15,15,q,111.40).
red(v1,52,v1_16,16,q,27.90).
red(v1,52,v1_17,17,q,106.30).
red(v1,52,v1_18,18,q,33.60).
red(v1,52,v1_19,19,q,21.60).
red(v1,52,v1_20,20,q,14.40).

```

Three different machine learning approaches were used: backpropagation networks [20, 19], nearest neighbor classification [5, 21], and decision tree induction [15, 17]. Table 4 gives a summary of the accuracies on unseen cases for the three approaches and the four different problems as estimated by ten-fold cross-validation on the same folds.

Table 4: Classification accuracy on unseen cases when learning from classified ^{13}C -NMR spectra with assigned atom numbers.

Problem/Approach	Backpropagation networks	Nearest neighbor	C4.5
Observed	89.6%	96.7%	96.4%
Reduced	95.5%	99.3%	97.6%
Observed - No Frequencies	83.9%	94.4%	93.1%
Reduced - No Frequencies	97.1%	98.8%	98.4%

While the accuracies achieved are very high and the problem may be considered solved at first sight, one should bear in mind that the propositional formulation specified above crucially depends on the atom number information. The assignment of atom numbers is a very difficult and important part of the classification process, and classification rules that do not rely on atom number assignments need to be derived for practical purposes. Without atom number information, there is no obvious propositional representation of the problem and more careful representation engineering (feature construction) or the use of more powerful techniques, such as ILP seems necessary. In the remainder of the paper, we describe experiments that apply each of these approaches separately, as well as their combination.

Table 5: The classes, their chemical names and number of instances in the database.

Code	Chemical name	Number of instances
c2	Trachyloban	9
c3	Kauran	353
c4	Beyeran	72
c5	Atisiran	33
c15	Ericacan	2
c18	Gibban	13
c22	Pimaran	155
c28	6,7-seco-Kauran	9
c33	Erythoxilan	9
c36	Spongian	10
c47	Cassan	12
c52	Labdan	448
c54	Clerodan	356
c71	Portulan	5
c79	5,10-seco-Clerodan	4
c80	8,9-seco-Labdan	6

3 Applying FOIL

As we are dealing with a learning problem where 23 different classes exist, there are 23 target relations `classC(MoleculeID)`, where `C` ranges over the possible classes. For example, the target relation `class52(MoleculeID)` corresponds to the skeleton type Labdan, most common among the 1503 diterpenes in the database (448 instances). The correspondence between the chemical names and the codes in the data is given in Table 5 together with the number of instances of each class. Only classes with more than one instance are listed.

3.1 Using multiplicities and frequencies

Given that much better results were obtained using the reduced multiplicities in earlier experiments, we decided to use only these (and not the observed multiplicities) for our experiments with ILP. After one takes the atom number information away, the background relation `red` can be simplified to a three argument relation `red(MoleculeID, Multiplicity, Frequency)`. A particular fact of this relation states that the ^{13}C -NMR spectrum of a given molecule contains a peak of a given multiplicity and frequency. For example, the fact `red(v1, t, 39.00)` states that the ^{13}C -NMR spectrum of molecule `v1` has a peak at frequency 39.00 ppm with multiplicity `t` (a triplet).

The background relation `red` is nondeterminate, as there are 20 tuples for each molecule. This prevents the use of ILP systems like GÖLEM [13] or DINUS [11]. While PROGOL [12] would be applicable, preliminary experiments showed it has prohibitive time complexity if longer rules/clauses are needed. Therefore, we used FOIL [16], in particular FOIL6.4.

Except for a variable depth limit of one, all other settings were left in their default state.

We first used FOIL to induce rules on the entire data set, and then performed ten-fold cross validations on the same partitions of training examples as used for the propositional case by Džeroski et al. [8]. For each experiment, FOIL was run 23 times, i.e., once for each target relation. The rules from the 23 runs were then taken together to produce a rule set. Before classification, the rules from the rule set were checked against the training set to record the number of examples of each class that they cover, as well as to note the majority class in the training set. This information is then used when classifying new examples, e.g., when testing the accuracy of the rule set.

For classification, we implemented a procedure analogous to that of CN2 [4]: when an example is classified, all clauses that cover it are taken into account. The distributions of examples of each class are summed for all rules that match and the majority class is assigned to the example. If no rule matches, the majority class (from the training set) is assigned.

Given the background relation `red(MoleculeID,Multiplicity,Frequency)`, FOIL induced 90 rules from the entire data set, comprising 607 literals in their bodies. At first sight, the rules are quite specific, covering relatively few examples. Also, many examples are left uncovered. Thus, even on the training set, these rules only achieve 51.6% accuracy. The ten-fold cross-validation yields 46.5% accuracy on unseen cases as the average over the ten runs. It seems that the background knowledge is sufficient to distinguish among the different classes (the rules typically cover examples of one class only), but FOIL induces too specific rules. The most general rule `class52(A) :- red(A,d,B), B<73.7, red(A,C,D), D>38.5, D<44.6, B>73.2` covers 43 examples of class 52.

3.2 Using multiplicities only

According to the domain expert (Karsten Siems), the multiplicities should suffice for correct classification, at least when atom numbers are available [8]. If we remove the `Frequency` argument from the relation `red`, all the information left about a particular molecule is captured by the number of atoms which have multiplicity `s`, `d`, `t`, and `q`, respectively. We store this information in the relation `prop(MoleculeID,SAtons,DAtoms,TAtoms,QAtoms)`. For our molecule `v1`, we have the fact `prop(v1,2,4,8,6)`.

Given the entire data set, the 23 target relations and the background relation `prop`, FOIL induced 17 rules with 52 literals in the bodies. Same settings were applied in FOIL as for the experiments with `red`. The rules induced in this case are much more general than the ones obtained with `red`. The most general rule `class52(A) :- prop(A,B,C,D,E), E>5, C>3, D>7, B>1` covers 429 examples (of which 355 of class 52, 72 of class 54). Many rules cover examples of several classes. It seems that the background knowledge is insufficient to completely distinguish among the different classes and FOIL is forced to induce overly general rules. This, however, has positive effect on accuracy as compared to the experiments with `red`. Using `prop`, FOIL achieves 69.0% accuracy on the entire data sets, and 70.1% accuracy on unseen cases (ten-fold cross-validation).

At this point, note that we have in fact applied FOIL to a propositional problem. Namely, from the nondeterminate representation with the `red` relation, we have constructed a four-feature representation of each molecule. We postpone the comparison to propositional learners to the next section.

3.3 Combining engineered features and a relational representation

Having introduced the four new features with the `prop` relation, which seem to capture some general properties of the molecules, we repeated the experiment with the nondeterminate relational representation. This time FOIL was given both the relation `red` and the relation `prop`. The same settings of FOIL were applied as in the previous two subsections.

Given the entire data set, the 23 target relations and the background relations `red` and `prop`, FOIL induced 68 rules with 401 literals in the bodies. The rules were more general than rules induced using `red` only, but were more specific than rules induced using `prop` only. In most cases, each rule covers examples of one class only. The most general rule `class52(A) :- prop(A,B,C,D,E), E>5, C>3, D>7, B>1, red(A,d,F), F>54.8, F<=72.3` covers 227 examples of the correct class. The induced rules achieve 83.2% accuracy on the entire data set and 78.3% accuracy on unseen cases (ten-fold cross-validation). Combining the engineered features with the relational representation thus has a positive effect.

4 Comparing FOIL to propositional approaches

4.1 C4.5 using multiplicities only

As mentioned above, the relation `prop` defines a propositional version of our classification problem. We therefore applied a propositional learner, C4.5 [17], to this problem. The same experimental setting (tree induced on whole data set first, then ten-fold cross-validation) and the default settings of C4.5 were used.

The induced tree achieves 80.4% accuracy on the entire data set. The leaves of the tree typically contain examples of several different classes, confirming the suspicion that the four features do not suffice for completely correct classification. The classification accuracy on unseen cases as measured by ten-fold cross-validation (same folds as above) is 78.5%, which is almost the same as the accuracy achieved by FOIL using both `red` and `prop`.

4.2 Nearest neighbor using multiplicities only

We also applied nearest neighbor classification [5, 21] to the propositional problem defined by the relation `prop`. The same experimental setting (cross-validation folds) was used. Training on the entire dataset gives 100% accuracy.

Cross-validation on the number of neighbors used in classification was tried out, as well as two different forms of feature weighting, but the basic nearest neighbor method performed best. The classification accuracy on unseen cases (average over the ten folds) is 79.0%, which is almost the same as the accuracy of C4.5 and the accuracy achieved by FOIL using both `red` and `prop`.

4.3 Backpropagation networks

Various network architectures were explored to see how in comparison backpropagation networks would perform at the classification of skeletons based on unassigned ^{13}C NMR data.

This was done independently of the above experimental setup, but using the same partitions for cross-validation.

A standard backpropagation network [20, 19] with 960 input neurons, no hidden units and 23 output units was trained with the same input data as for the ILP experiments. We divided the 960 input neurons into four equally large sets of 240 nodes, one each for singlets, doublets, triplets and quadruplets. The best representation was to have for each multiplicity 24 frequency intervals (0 - 10 ppm, 11 - 20 ppm, ..., 231 - 240 ppm) with 10 nodes each and to feed the value of 1 for each frequency in the spectrum into the next unoccupied input neuron of the appropriate interval. All remaining input nodes are filled with zeros. Thus, the artificial neural net sees a discretized distribution of multiplicity signals.

During cross-validation, accuracy on the training set (average for the 10 runs) reached 97.9%, while accuracy on unseen cases reached 79.9%. Other network variations, e.g., feeding the actual frequency value into the appropriate neuron instead of 1 or using a 4 x 15 architecture that receives the sorted frequencies for each multiplicity as input did not produce better results.

5 Applying RIBL

RIBL (relational instance-based learning) [9] generalizes the nearest neighbor method to a relational representation. RIBL first constructs cases by putting together all facts that relate to (in this case) a single molecule. Training cases are stored for further reference. When a new molecule has to be classified, RIBL calculates the similarities between it and each of the training cases, then assigns it the class of the nearest training case.

The similarity measure used in RIBL is a generalization of similarity measures used in propositional instance-based learners. In fact, given a propositional problem, RIBL becomes the classical nearest neighbor method and has the same performance as the latter. This is a very desirable property for a relational learner.

We used the same formulations and experimental methodology as for FOIL (Section 3). As it stores all training cases, RIBL achieves 100% accuracy given the entire data set for both training and testing. To estimate accuracy on unseen cases, ten-fold cross validations on the same folds as above were performed.

Given only the **red** relation, RIBL achieved 86.5% classification accuracy (average over the ten partitions) on unseen cases. This is an increase of 40% over the accuracy achieved by FOIL. Note that propositional approaches (C4.5) are not applicable to this formulation of the problem.

Given only the **prop** relation, RIBL behaves identically to the nearest neighbor method, thus yielding 79.0% accuracy on unseen cases, a performance equivalent to that of C4.5. When provided with both the **red** and the **prop** relations, RIBL achieves 91.2% accuracy on unseen cases. Using the engineered features improves RIBL’s performance by roughly 5%, pushing further the best result achieved at classifying diterpene NMR spectra without assigned atom numbers. Again, propositional approaches (C4.5) are not applicable to this formulation of the problem.

Table 6: Accuracies of different ILP approaches when classifying NMR spectra without assigned atom numbers.

Problem/System	FOIL	RIBL	TILDE	ICL
red	46.5%	86.5%	81.6%	65.3%
prop	70.1%	79.0%	78.5%	79.1%
red+prop	78.3%	91.2%	90.4%	86.0%

6 Applying TILDE

The TILDE system [3] is a first order logic upgrade of Quinlan’s C4.5 system [17] that induces logical decision trees. It learns from interpretations [6], which means that the `MoleculeID` argument does not occur in the predicates; except that, the same representation is used as for the FOIL experiment. We have also used the same experimental methodology: using multiplicities and frequencies, using multiplicities only by means of the engineered predicate `prop`, and combining both original and engineered predicates.

With purely propositional data, TILDE achieves exactly the same accuracies (80.4% for training on the whole set, 78.5% for tenfold crossvalidation) as C4.5, which is not unexpected as it is an upgrade of the latter. Table 6 lists the accuracies achieved by all ILP systems applied to all problem formulations (average accuracies of ten-fold crossvalidation).

For the non-propositional experiments, discretization was performed on the frequencies. Additionally, the multiplicity counts (arguments of the `prop` predicate) were discretized separately, but with the same number of discretization thresholds. The number of thresholds is a parameter of TILDE, and was varied. Best results were obtained with five discretization thresholds and this value was used for all subsequent experiments. Except for the number of discretization thresholds, default values were used for all parameters.

An accuracy of 81.6% is achieved on the `red` formulation and an accuracy of 90.4% on the `red+prop` formulation. The latter is essentially the same as the one of RIBL (91.2%), while the former is five percentage points lower than RIBL’s respective performance (86.5%).

Despite using a post-pruning mechanism similar to that of C4.5, TILDE with default parameters produces trees that are very large and consequently difficult to interpret. We thus varied the minimal number of examples that is to be covered by each leaf (by default, this is 2) in order to obtain smaller trees. Increasing this parameter lowers predictive accuracy, but also makes the trees simpler and more interpretable. Table 7 shows the average accuracy and tree complexity (number of internal nodes) on tenfold crossvalidation. Figure 3 shows the tree obtained with minimal coverage set to 20 when learning on the whole dataset.

While the trees become simpler, on the higher levels they are still very similar. We have taken a look at the most general leaf occurring in the tree and observed its evolution as the minimal coverage changes. The number of examples covered by this leaf does not change much, as can be seen in Table 7. The leaf has been indicated with `***` in Figure 3. The rule corresponding to this leaf is `c52 :- prop(A,B,C,D), A /= 4, A /= 3, red(q, I), 20.97 <= I <= 59.5, B <= 4, red(d, J), 49.12 <= J <= 93.8, D = 6`. It correctly predicts 347 of 357 examples to be of class 52.

Since the trees are quite similar, one could say that the simpler trees give an understandable approximate theory, while the complex trees refine this theory, making it much more

Table 7: Predictive accuracy of tree, average complexity of trees, and class+coverage of most general leaf, tabulated for different values of the minimal coverage parameter in TILDE.

minimal coverage	accuracy	complexity	most general leaf
2	90.4%	74.8	c52 (328/330)
5	88.9%	45.7	c52 (328/330)
10	86.4%	26.8	c52 (347/357)
20	83.0%	14.9	c52 (347/357)

detailed and less understandable, but more accurate.

7 Applying ICL

The ICL system [7] uses exactly the same representation as TILDE, but induces rule sets instead of trees. It is a first order logic upgrade of the CN2 algorithm [4].

For the ICL experiments, both frequencies and multiplicity counts (arguments of the **prop** predicate) were discretized. This was done mainly for efficiency reasons. Although these counts are discrete already, the discretization algorithm reduces the number of thresholds. Ten thresholds were used in our experiments.

On the propositional data, ICL achieves 79.1% accuracy with an average theory complexity of 22 rules. When learning from the whole data set, a theory of 22 rules is found, of which the most general rule covers practically the same examples as FOIL's: **c52** :- **prop(S,D,T,Q)**, **S=2**, **D=4**, **Q=6** (429 examples, mostly **c52** (355 examples) and **c54** (72 examples); in FOIL's rule, the tests are **S>1**, **D>3**, **T>7**, **Q>5**).

Using frequencies only, ICL achieves 65.3% accuracy on unseen cases (with an average complexity of 57 rules). Learning on the whole dataset yields a theory of 54 rules. Compared to FOIL, these rules are much more general (the most general rule covering 317 examples, of which 190 correct).

Using both frequencies and multiplicity counts, an average accuracy of 86.0% and average theory size of 64 rules was achieved. When learning a theory from all examples, 79 rules were found, the most general being **c3** :- **prop(S,4,9,Q)**, **Q<=4**, **red(s,B)**, **B>=42.5**, **red(s,B2)**, **38.9<=B2<=44.2**, which covers 262 examples (242 of class **c3**).

Table 8 gives an idea of the generalization power of the different systems by showing the coverage of the most general rule, as well as the total number of rules in the theory. Note that the complexity of each individual rule (e.g. the number of literals in it) is not relevant in this respect. If a theory has high predictive accuracy and few rules, the rules have to cover many examples and are in this respect general, whether or not they contain many literals.

It can be seen from the table that there are no large differences between the number of rules induced, but nevertheless FOIL has a tendency to induce more specific rules than the other systems. These two aspects, taken together, can explain the lower accuracy achieved by FOIL.

Figure 3: Tree with minimal coverage of 20 examples. The `red_X(Y)` literals stand for `red(X,Y)` (where $X=s, d, t$ or q). For each leaf, class prediction and coverage are given.

```

prop(A,B,C,D)
A = 4 ?
+--yes: c4 [72 / 95]
+--no:  A = 3 ?
    +--yes: D =< 4 ?
        |   +--yes: C =< 10 ?
            |   +--yes: red(t,E) , between(E,[99.7,148.15]) ?
                |   +--yes: red(s,F) , between(F,[ - inf,38.9]) ?
                    |   +--yes: c3 [24 / 35]
                    |   +--no:  c54 [17 / 26]
                    |   +--no:  B = 3 ?
                        |   +--yes: c3 [10 / 20]
                        |   +--no:  red(q,G) , between(G,[ - inf,15.2]) ?
                            |   +--yes: red(s,H) , between(H,[38.9,47.11]) ?
                                |   +--yes: c3 [50 / 62]
                                |   +--no:  c5 [18 / 20]
                                |   +--no:  c3 [246 / 248]
                            +--no:  c54 [30 / 30]
                        +--no:  B =< 3 ?
                            +--yes: c22 [139 / 161]
                            +--no:  c18 [11 / 34]
                    +--no:  red(q,I) , between(I,[20.97,59.5]) ?
                        +--yes: B =< 4 ?
                            |   +--yes: red(d,J) , between(J,[49.12,93.8]) ?
                                |   +--yes: D = 6 ?
                                    |   +--yes: c52 [347 / 357]
                                    |   +--no:  red(d,K) , between(K,[ - inf,49.12]) ?
                                        |   +--yes: red(L,M) , between(M,[167.4,176.6]) ?
                                            |   +--yes: c54 [28 / 32]
                                            |   +--no:  c52 [14 / 23]
                                            |   +--no:  c52 [40 / 43]
                                        +--no:  c54 [46 / 64]
                                    +--no:  c54 [20 / 45]
                                +--no:  c54 [188 / 208]
                            +--no:  c54 [188 / 208]
    +--no:  c54 [188 / 208]

```

Table 8: Comparison of generalization achieved with different ILP systems. Most general rule or leaf: predicted class (correctly covered/covered); number of rules/leaves.

		FOIL		TILDE		ICL
red	c52 (43/43)	90	c3 (233/238)	120	c3 (190/317)	57
prop	c52 (355/429)	17	c52 (365/444)	24	c52 (355/429)	22
red+prop	c52 (227/227)	68	c52 (328/330)	84	c3 (242/262)	64

8 Discussion

For practical purposes, ^{13}C -NMR spectra of diterpenes without assigned atom numbers have to be classified. This is a problem that is not directly transformable to propositional form and calls for either representation engineering or the use of inductive logic programming. We explored both approaches separately and in combination. Adding the engineered features to the natural relational representation improved the performance of both relational learning systems used.

Using the engineered features only, propositional approaches (in particular C4.5, nearest neighbor and neural networks) achieve around 79% accuracy on unseen cases. This is roughly 20% less than the accuracies achieved when classifying ^{13}C -NMR spectra of diterpenes with assigned atom numbers.

Using FOIL on the natural relational representation yields unsatisfactory results. Combining the relational representation with the engineered features greatly improves FOIL's performance. However, given the engineered features only FOIL performs much worse than C4.5, so that the best performance of FOIL (combining the relational representation and the engineered features) is comparable to that of C4.5. The reason for FOIL's poor performance is that the rules induced are overly specific as indicated by their coverage and confirmed by expert comments. However, the rules found by FOIL are quite short, indicating that the problem lies in the search heuristic: it directs the search to short rules with small coverage, despite the fact that longer rules with higher coverage exist.

From the above it is clear that a desirable property of relational learning systems is the following: given a propositional problem, a relational learning system should perform comparably to propositional systems. RIBL, TILDE, and ICL, which extend propositional learning paradigms to a relational framework, all have this property. Given the engineered features only, they achieve 79%, 78.5%, and 79.1%, accuracy on unseen cases, respectively. Given the relational representation only, RIBL performs best, followed by TILDE and ICL. All three perform much better than FOIL. The situation is similar for the problem formulation combining the relational representation and the engineered features: RIBL is followed closely by TILDE, which performs at the same level, ICL follows, and FOIL is far behind. The best (91%) accuracy on unseen cases is 11% better than the best propositional result of 80% (backpropagation networks with 960 features).

The accuracies achieved by RIBL and TILDE are in the range of the accuracy with which experts classify diterpenes into skeleton types given ^{13}C NMR spectra only. That number can actually only be estimated since it is expensive to have an expert carry out a statistically significant number of structure predictions without using other additional information that often becomes available from heterogeneous sources (such as literature, and ^1H NMR spectra). This basically means that ^{13}C NMR is not completely sufficient for classifying diterpenes and that great improvements of classification accuracy are not to be expected.

While RIBL does not yield understandable theories, FOIL, TILDE and ICL are able to do so, at least in principle. In practice, understandability of the theories is diminished when only complex theories are accurate. Results with TILDE show, however, that it may be advantageous to find a slightly less accurate theory that is more understandable. The latter theory can be considered a cruder version of the more accurate theory: indeed, in

the case of TILDE large similarities are found between the complex, accurate tree and the simpler, less accurate tree. This offers an interesting compromise between understandability and correctness.

It would be interesting to see if similar effects take place when using newer versions of RIBL (based on Aha's IB3 [1]), which store only a fraction of all training cases as prototypes. These prototypes may be considered as explanations/rules and a smaller number of prototypes would correspond to a shorter theory.

Another direction for further work is to carefully evaluate the rules produced by FOIL, TILDE and ICL from the understandability point of view (domain expert). In this way the trade-off between accuracy and understandability in this domain would be understood better. The explanatory power of these should be compared to that of prototypes offered as explanations by the newer versions of RIBL.

Finally, the use of additional information sources in the learning process, such as ^1H NMR spectra (where available), should be considered.

Acknowledgements

Part of this work was performed while Sašo Džeroski was an ERCIM (European Research Consortium for Informatics and Mathematics) fellow at GMD (German National Research Center for Information Technology), Sankt Augustin, Germany and ICS-FORTH (Institute of Computer Science, Foundation for Research and Technology - Hellas, Heraklion, Greece). This work was also supported in part by the ILP2 project (ESPRIT IV LTR Project 20237 Inductive Logic Programming 2).

References

- [1] Aha, D., Kibler, D., and Albert, M. Instance-based learning algorithms. *Machine Learning*, 6: 37–66, 1991.
- [2] Abraham, R.J., Loftus, P. *Proton and Carbon 13 NMR Spectroscopy, An Integrated Approach*. Heyden, London, 1978.
- [3] Blockeel, H., and De Raedt, L. Experiments with top-down induction of logical decision trees. Technical report CW 247, Dept. of Computer Science, Katholieke Universiteit Leuven, Heverlee, Belgium, 1997.
- [4] Clark, P. and Boswell, R. Rule induction with CN2: Some recent improvements. In *Proc. Fifth European Working Session on Learning*, pages 151–163. Springer, Berlin, 1991.
- [5] Cover, T.M., and Hart, P.E. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13: 21–27, 1968.
- [6] De Raedt, L., and Džeroski, S. First order jk -clausal theories are PAC-learnable. *Artificial Intelligence*, 70: 375–392, 1994.

- [7] De Raedt, L., and Van Laer, V. Inductive constraint logic. In *Proc. Sixth International Workshop on Algorithmic Learning Theory*, pages 80–94. Springer, Berlin, 1995.
- [8] Džeroski, S., Schulze-Kremer, S., Heidtke, K., Siems, K., Wettschereck, D. Diterpene structure elucidation from ^{13}C NMR spectra with machine learning. In *Proc. ECAI'96 Workshop on Intelligent Data Analysis in Medicine and Pharmacology*, 1996.
- [9] Emde, W., Wettschereck, D. Relational instance-based learning. In *Proc. Thirteenth International Conference on Machine Learning*, pages 122–130. Morgan Kaufmann, San Mateo, CA, 1996.
- [10] Gray, N. A. B. *Progress in NMR-spectroscopy, Vol. 15*, pp. 201–248, 1982.
- [11] Lavrač, N., Džeroski, S. *Inductive Logic Programming: Techniques and Applications*. Ellis Horwood, Chichester, 1994.
- [12] Muggleton, S. Inverse entailment and PROGOL. *New Generation Computing*, 13: 245–286, 1995.
- [13] Muggleton, S., and Feng, C. Efficient induction of logic programs. In *Proc. First Conference on Algorithmic Learning Theory*, pages 368–381. Ohmsha, Tokyo, 1990.
- [14] *Natural products on CD-ROM*. Chapman and Hall, London, 1995.
- [15] Quinlan, J.R. Induction of decision trees. *Machine Learning* 1(1): 81–106, 1986.
- [16] Quinlan, J.R. Learning logical definitions from relations. *Machine Learning*, 5(3): 239–266, 1990.
- [17] Quinlan, J.R. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.
- [18] Schulze-Kremer, S. *Molecular Bioinformatics - Algorithms and Applications*. de Gruyter, Berlin, 1995.
- [19] *Stuttgart Neural Network Simulator*. Computer code available from the University of Stuttgart, Germany, via anonymous ftp <ftp://ftp.informatik.uni-stuttgart.de/pub/SNNS>, 1995.
- [20] Tveter, D. R. *Fast-Backpropagation*. Computer code available from the author. Address: 5228 N Nashville Ave, Chicago, Illinois, 60656, drt@chinet.chi.il.us, 1995.
- [21] Wettschereck, D. A study of distance-based machine learning algorithms. PhD Thesis, Department of Computer Science, Oregon State University, Corvallis, OR, 1994.