1

A Non-Generative Constraint-Based Formalism

PHILIPPE BLACHE AND FRANK MORAWIETZ

LPL-CNRS, Université de Provence
29 Avenue Robert Schuman
13621 Aix-en-Provence, France

SFS, Universität Tübingen
Wilhelmstr. 113
72074 Tübingen, Germany

pb@lpl.univ-aix.fr
frank@sfs.nphil.uni-tuebingen.de

ABSTRACT. Most of modern linguistic theories rely on the notion of constraints. However, only few applications try to implement the parsing process directly with constraint programming: in most cases, constraints are interpreted in a passive sense, the parsing process itself not being different than the one of classical derivational approaches. We show in this paper that the problem comes from the representation of linguistic information and the generative interpretation of the relation between grammar and language. We propose a new approach, called *Property Grammar* in which all syntactic information is represented with constraints, the parsing process directly being constraint resolution. Our ideas are implemented using Constraint Handling Rules (Frühwirth; 1998).

1. INTRODUCTION

Modern linguistic theories often make use of the notion of *constraint*: most constraint-based theories (cf. Shieber (1992), Borsley (1996)) as well as other paradigms such as representational versions of the principle and parameters approach (including optimality theory, Archangeli and Langendoen (1997)) or dependency grammars (constraints dependency grammars, Maruyama (1990)) follow this tendency. A constraint approach offers several advantages, in particular a fine-grained representation of information, a clear distinction between linguistic objects and their properties, and a better declarativity.

In computational linguistics, the idea that constraints have a direct counterpart in the implementation is implicit. Unfortunately, this is only partly true and constraints only have local scope during the processing. In spite of several works trying to use constraint programming (see for example Günthner (1988), Maruyama (1990) or more recently Blache and Paquelin (1996), Götz and Meurers (1997), Duchier and Thater (1999)), the parsing process cannot be interpreted as actual constraint resolution. This doesn't follow the original idea (see in particular Sag and Wasow (1999)) which proposes the direct use of the linguistic information as a set of constraints over the categories. We show in this paper that this problem has two main origins: the representation of the linguistic information and the influence of the generative conception of the relation between grammars and their generated languages on the linguistic analysis.

We propose a new approach called *Property Grammar* representing all the linguistic information by means of simple constraints. Since these constraints constitute an actual constraint system (in the same sense as a set of equations forms a system), it is then possible to consider parsing as a pure constraint program.

2. Some Desiderata for Constraint-Based Approaches

Constraints in linguistic theories generally play only a local role[1]: they are used to control the instantiation of linguistic objects (the categories) and represent relations between these objects (or their components). Then, a clear distinction between the objects and their properties becomes possible, which favors declarativity. For example, in a theory like GPSG, one can find constraints at different levels of the description: linear precedence, feature cooccurrence restriction or universal instantiation principles. But not all the information is represented with constraints and during a parse, one has to select a local tree first and then verify that this tree satisfies the different constraints. This is typically a passive use of constraints.

The problem is different in HPSG in which most of the information is actually represented by means of constraints (cf. Sag and Wasow (1999) and especially for an implementation Blache and Paquelin (1996) or Meurers and Minnen (1998)). However, the hierarchical information remains predominant in the sense that verification of constraints can only be applied to an implicit local tree. The verification of the principles for example requires the construction of a hierarchical structure containing a mother and its daughters. This doesn't mean

[1]Constraints in this case are not active all along the parsing process, but only for specific structures built during the parse. Such constraints are local because their scope is reduced according to the knowledge of the constrained structure.

that such a structure has to be fully specified. Some delaying mechanisms can obviously be applied and active constraints can be used locally. But in this case too, we can globally characterize the use of constraints for the parsing process as passive and parsing doesn't amount to pure constraint resolution.

One of the reasons for these problems is the generative interpretation of the relation between grammars and languages. In this case, the notion of derivation is central and parsing an input consists in finding a derivation generating it. Such a conception is revealed in the approaches cited above with the role played by the local trees (which usually amounts to one derivation step). Parsing is then conceived as a derivation mechanism controlled by constraints (often reduced to unification). This certainly stems from the fact that licensing of constraints is implemented as improved versions of the generate-and-test paradigm to ensure some sort of efficiency (and sometimes even termination). So far, we do not know of an implementation which directly uses only constraints to drive the parsing (or the generation) process.

An optimal use of constraints, both from the linguistic and the computational point of view, should meet some requirements. We present in this section some basic desiderata for a genuine constraint-based approach.

The first important property concerns knowledge representation: all linguistic information has to be represented by means of constraints which form an actual system. This property has an important consequence on the declarativity of the approach. Indeed, it allows a clear distinction between the representation of information (in particular the form of the objects) and the mechanisms for calculating with them.

The second important point is the fact that, insofar as the set of constraints constitutes a system, all the constraints are at the same level. This means that the order of verification of the constraints is not relevant. Obviously, some heuristics can rely on a hierarchization of the constraints or more precisely on the control of their relaxation. But this doesn't alter the fact that the system of constraints as to be taken as a whole. Our approach is thus totally different from the Optimality Theory (see Archangeli and Langendoen (1997)) in which constraint ranking contains implicitly linguistic information.

Encapsulation is another important property for constraint-based theories: a constraint must represent homogeneous information. We can then distinguish different types of constraints, each one representing a specific part of linguistic information. Such a requirement avoids implicit information and each type of constraint becomes linguistically motivated.

The last requirement concerns the notion of grammaticality. One interesting property of constraint solving is the possibility to build an approximate solution (when no exact solution can be found) which is nothing but the state of the

constraint system after verifying its satisfiability for a given input. As for parsing, finding an exact solution consists in associating a syntactic structure with a given input, e.g. finding a derivation from a distinguished symbol to this input, answering the question of grammaticality: an input is grammatical if a solution can be found. However, the relevant question concerning parsing of real natural language should not be grammaticality, but rather the robustness of providing information about the input. Any kind of input has to be treated. Therefore, we propose to replace the notion of grammaticality with that of *characterization* which is much more general: a characterization is the state of the constraint system for a given input.

3. Property Grammar

Property Grammar (cf. Blache (1999), Bès and Blache (1999)) constitutes an approach implementing these requirements: declarativity, encapsulation, satisfiability. We present in this section their main characteristics. As specified above, all the linguistic information is represented here by means of *properties*.

3.1. The Basic Properties. We use for the representation of syntactic information 7 properties[2] defined as follows :

- **Constituency** (*const*): Set of categories constituting a phrase.
- **Obligation** (*oblig*): Set of compulsory, unique categories (heads).
- **Unicity** (*uniq*): Set of categories which cannot be repeated in a phrase.
- **Requirement** ($\Rightarrow$): Cooccurrence between sets of categories.
- **Exclusion** ($\not\Rightarrow$): Restriction of cooccurrence between sets of categories.
- **Linearity** ($\prec$): Linear precedence constraints.
- **Dependency** ($\rightsquigarrow$): Dependency relations between categories.

Let us take some examples (all our examples are from French):

- *Constituency: const(NP)* = $\{$*Det, N, AP, PP, Sup, Pro*$\}$
  All these constituents can appear within a sequence characterizing a *NP*.
- *Obligation*: *oblig(NP)* =$\{$*N, AP, Pro*$\}$
  In a NP, these categories can be the head.
- *Unicity*: *unic(NP)* = $\{$*Det, N, AP, PP, Sup, Pro*$\}$
  Set of unique categories in a *NP*.
- *Linearity*: *Det* $\prec$ *AP* *AP* $\prec$ *N*

[2]This list of properties can be extended in order to integrate other levels of linguistic analysis such as prosody or semantics.

- *Requirement*:
  - $\{le\ \hat{e}tre_{[N,P]}\} \overset{VP}{\Rightarrow} ClR_{[N,P]}$ (Je me le suis dit/*I told that to myself*)
    If the categories *le* (clitic) and a finite verb *être* (with agreement features *N* and *P*) cooccur (characterizing a *VP*), then a reflexive clitic (which agrees with the verb) also has to be present.
- *Exclusion*:
  - $ClR \overset{VP}{\not\Rightarrow} lui$ (*Je me lui dis / **I told him myself*)
    In a *VP*, a reflexive clitic cannot cooccur with the clitic *lui*.
- *Dependency*: $Det \rightsquigarrow N$ $\qquad\qquad$ $AP \rightsquigarrow N$

The first important remark is that characterizations are associated with sets of categories. This is interesting in the sense that a characterization can be given independently from the characterization of the non-lexical constituents belonging to this set. The second important point concerns the fact that properties can be expressed over sets of categories, allowing then to represent contextual information. It is the case in particular for the *requirement* and *exclusion* properties, see the indices.

4. Parsing as Constraint Resolution

Insofar as syntactic information forms a constraint system, parsing can then be implemented using constraint programming techniques such as constraint resolution. This approach has several advantages, in particular flexibility and reusability. For example, one can choose to verify only a subset of the constraint system. It is also possible to satisfy constraints with different granularity levels (e.g. a shallow description of the sentence level together with a fine-grained definition of the *NP*). The possibility of integrating several constraints coming from different levels of the linguistic analysis also constitutes an interesting property of using such an approach. But the most important advantage of constraint resolution lies in the fact that it allows a direct interpretation of the notion of characterization.

We present in this section the general architecture of the parsing process. We first describe the elementary trees constituting the basic information and then the constraint resolution process itself.

4.1. Elementary trees. Interpreting the parsing process as checking constraint satisfaction can offer a new point of view. Indeed, constraints are applied directly over the objects (the categories) rather than over structures (a local tree). It is thus possible to apply the constraints to the set of all the possible categories that can be associated with the input.

Insofar as constraints are applied over sets of categories (whatever their level, lexical or not), we can associate with each word its possible categories and their possible projections.

Such structures can be represented as an elementary tree as described in (1).

(1)

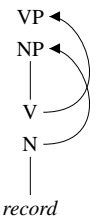

This example presents the elementary tree of the word *"record"* which can be a noun or a verb. These categories can respectively belong to the *constituent* property of the *NP* and *VP*. The elementary tree is then composed of three nodes: the lexical item, its categorizations and their projections. It is possible to associate such a tree with each lexical item, without any distinction between major and minor categories (i.e. our notion of projection is reduced to the membership of a category to a *const* property).

4.2. Satisfying Constraints to Build Characterizations. The parsing process takes as input the set of elementary trees that can be associated with the sentence and builds all the characterizations of this input. A characterization is the state of the constraint system for a particular set of categories, in other words the set of satisfied properties and that of unsatisfied ones.

A *characterization* is then defined over a set of categories by the set P^+ (satisfied properties) and P^- (unsatisfied properties).

Let us take an example. Fig. 1 presents a grammar of the noun phrase in French. It describes in particular the properties of the noun phrase, the adjective phrase and the superlative. In (2), the elementary trees associated with the words of the noun phrase *le livre le plus rare (the rarest book)* are given.

(2)

NP	.	NP	AP	AP
Sup	NP	Sup	Sup	Sup
|	▌	|	▌	|
Det	N	Det	Adv	Adj
|	▌	|	▌	|
le	*livre*	*le*	*plus*	*rare*
(1)	(2)	(3)	(4)	(5)

The first stage consists in determining the set of categories to be verified. The most general approach consists in characterizing all subsets of categories

Properties of the *NP*

(1) $Const = \{Det, N, AP, Sup\}$
(2) $Oblig = \{N, AP\}$
(3) $N[com] \Rightarrow Det$
(4) $Det \prec N$
(5) $Det \prec AP$
(6) $Det \prec Sup$
(8) $N \prec Sup$
(9) $AP \not\Leftrightarrow Sup$
(10) $Det \rightsquigarrow N$
(11) $AP \rightsquigarrow N$
(12) $Sup \rightsquigarrow N$

Properties of the *AP*

(1) $Const = \{Adj, Adv\}$
(2) $Oblig = \{Adj\}$
(3) $Adv \prec Adj$
(4) $Adv \rightsquigarrow Adj$

Properties of the *Sup*

(1) $Const = \{Det, Adv, Adj\}$
(2) $Oblig = \{Adj\}$
(3) $Adj \Rightarrow Det$
(4) $Adj \Rightarrow Adv$
(5) $Det \prec Adv$
(6) $Det \prec Adv$
(7) $Adv \prec Adj$
(8) $Det \rightsquigarrow Adj$
(9) $Adv \rightsquigarrow Adj$

FIGURE 1. A property grammar for the *NP*

belonging to the set of elementary trees. Considering all the possible subsets of categories has an exponential cost but some simple control mechanisms can be applied. In the following, we will only take into account the sets of juxtaposed categories (implementing then the projectivity constraint); in this case, sets of constraints are in fact lists.

Building the characterizations consists for each subset of categories of verifying the satisfiability of the properties. The properties that cannot be evaluated remain untouched. This means that a set of properties belonging to a characterization always specifies a phrase-level category. More precisely, it is possible to associate with a set of categories an upper-level category C (or several categories in case of ambiguity). We also say in this case that C is *characterized* by its P^+ and P^- sets.

The following example illustrates the construction of such characterizations for (2) according to the grammar presented in Fig. 1. Each category is associated with its position in the input. The characterization sets follow the sequence of categories with their properties being represented by their indices. For example, the characterization of the set $\{Det_1, N_2\}$ indicates that all the properties specified for these categories (i.e. constituency, obligation, requirement, precedence and dependency) are satisfied. The set $\{Adv_4, AP_5\}$ presents two characterizations: this sequence can specify an AP (all the properties of the grammar are verified) or a *Sup* in which a requirement property is not satisfied.

- Det_1/N_2 :	$P^+(NP) = \{1, 2, 3, 4, 10\}$
	$P^-(NP) = \emptyset$
- Sup_1/N_2 :	$P^+(NP) = \{1, 2, 12\}$
	$P^-(NP) = \{6,3\}$
- Det_3/Sup_4 :	$P^+(NP) = \{1, 6\}$
	$P^-(NP) = \{2\}$
- Det_3/AP_4 :	$P^+(NP) = \{1, 5\}$
	$P^-(NP) = \{2\}$
- Sup_3/AP_4 :	$P^+(NP) = \{1\}$
	$P^-(NP) = \{2, 9\}$
- Adv_4/AP_5 :	$P^+(AP) = \{1, 2, 3, 4\}$
	$P^-(AP) = \emptyset$
	$P^+(Sup) = \{1, 2, 4, 9\}$
	$P^-(Sup) = \{3\}$
- $Det_1/N_2/Sup_3$:	$P^+(NP) = \{1\text{-}4, 6, 8, 9, 10, 12\}$
	$P^-(NP) = \emptyset$
- $Sup_1/N_2/Det_3$:	$P^+(NP) = \{1, 2, 3, 9, 10, 12\}$
	$P^-(NP) = \{4, 6, 8\}$
- $Det_3/Adv_4/Adj_5$:	$P^+(Sup) = \{1, 2, 3, 4, 5, 6, 8, 9\}$
	$P^-(NP) = \emptyset$
- $Det_3/AP_4/Sup_5$:	$P^+(Sup) = \{1, 4, 5, 6\}$
	$P^-(NP) = \{2, 9\}$
- $Det_1/N_2/Sup_3/AP_4$:	$P^+(NP) = \{1\text{-}8, 10, 11, 12\}$
	$P^-(NP) = \{9\}$

FIGURE 2. Some characterizations of (2)

Note that some characterizations have an empty set of unsatisfied properties (let us call them *positive* characterizations). But no particular status is given to them. This allows the robust treatment of any kind of input, i.e., we do not restrict the analysis to the set of "grammatical" structures.

4.3. Determining Coverage. The core of the process is the construction of the characterization sets. A general mechanism makes use of these sets in order to find a characterization covering the entire input.

A characterization made up of lexical categories (cf. the superlative's characterization in the previous example), constitutes a covering of the segment of the input corresponding to the positions of the constituents. The question is different for embedded constituents. In this case, the set of categories covers the corresponding part of the input if there exists a characterization for all the embedded constituents describing non-lexical categories. In the previous example, the characterization $\{Det_1, N_2, Sup_3\}$ of the *NP* covers the input because there exists a characterization $\{Det_3, Adv_4, Adj_5\}$ for its embedded constituent, the *Sup*.

Let us be more precise on the evaluation of covering categorizations. Characterizations don't take into account the status of their embedded constituents. In other words, lexical and non-lexical categories play the same role. This means that each position corresponds to a category. In the case of embedded constituents, we then have several successive categories with the same type as described in (3).

(3)

NP	AP	AP	.
Sup	Sup	Sup	NP
|	|	|	▌
Det	Adv	Adj	N
|	|	|	▌
un	*très*	*vieux*	*livre*
(1)	(2)	(3)	(4)

In this example, one covering of the input will be the set of categories: $\{Det_1, AP_2, AP_3, N_4\}$ which characterizes a noun phrase. The condition for this set to be an actual covering is the presence of a characterization of the *AP* (which will be in this case the set of categories $\{Adv_2, Adj_3\}$).

4.4. Controls and Heuristics. A naive implementation of these mechanisms is obviously not efficient. Indeed, let m be the number of words of the input, c the maximum number of categorizations for each word, then the number n of categories to analyze is bounded by $n = 2mc$. Then, the number of sets of categories to analyze is bounded by 2^n.

Several mechanisms can be added in order to reduce the number of sets to analyze and control the resolution process. In the first case, we can choose to build only sets of juxtaposed categories. At a higher level, one can consider only sets corresponding to subsets of properties (i.e. filtering the sets using the *const* property for example). The satisfaction process itself can be controlled using several heuristics. In particular, it is possible to filter the satisfiability according to a threshold given by the cardinality of the set of unsatisfied properties P^-: it is possible to build only characterizations with less than a certain amount of unsatisfied constraints. At the extreme, we can reduce the satisfiability to positive characterizations.

Another kind of heuristic consists in defining a hierarchization within the set of properties. In this case, the technique of the threshold mentioned above can be modulated with such a hierarchization, some constraints playing a more important role than others.

5. IMPLEMENTING PROPERTY GRAMMARS

Constraint programming is an elegant, declarative way of solving problems with higher-level programming languages. It represents a generalization of logic programming in the sense that it allows the (logical) specification of a problem on an arbitrary domain and tries to solve it with a combination of inference and search. A complete introduction to constraint programming can be found for example in Marriott and Stuckey (1998).

We describe in this section an implementation of Property Grammars in a particular constraint programming paradigm. We present this paradigm, Constraint Handling Rules (CHR), in a first step and then explain how to take advantage of it for implementing the resolution process. We want to show with such an implementation that property grammar can be implemented directly using only constraint resolution, which is not actually possible with classical generative approaches.

5.1. Constraint Handling Rules. There are several constraint programming environments available. The most recent and maybe the most flexible is the Constraint Handling Rules package included in SICStus Prolog (Intelligent Systems Laboratory; 1995; Frühwirth; 1998). All of these systems provide mechanisms for solving constraints. They maintain a constraint base or store which is continually monitored for possible rule applications, i.e., whether there is enough information present to successfully use a rule to simplify constraints or to derive new constraints. Whereas usually one deals with a fixed constraint domain and a specialized solver, CHR is an extension of the Prolog language which allows for the specification of user-defined constraints and arbitrary solvers. The strength of the CHR approach lies in the fact that it allows for multiple (conjunctively interpreted) heads in rules, that it is flexible and that it is tightly and transparently integrated into the Prolog engine.

In CHR constraints are just distinguished sets of (atomic) formulas. CHR allow the definition of rule sets for constraint solving with three types of rules: Firstly simplification rules ($<=>$) which replace a number of constraints in the store with new constraints; secondly propagation rules ($==>$) which add new constraints to the store in case a number of constraints is already present; and thirdly "simpagation" rules ($<=>$ in combination with a $\backslash$ in the head of the rule) which replace only those constraints with new ones which are to the right of the backslash.

To improve efficiency, rules can have guards and pragma declarations. A guard (separated from the rest of the body by a $|$) is a condition which has to be met before the rule can be applied. A pragma declaration affects the way the rule is compiled. Our approach will only introduce the *passive* declaration,

so we ignore the other possibilities. Declaring a constraint in the head passive simply means that this constraint will not trigger the rule on its appearance in the constraint store.

We cannot go into the details of the formal semantics of CHR here. The interested reader is referred to Frühwirth (1998). Let us just note that logically, simplification rules are equivalences and propagation rules are implications if their guard is satisfied. Simpagation rules are special cases of simplification rules. Soundness and completeness results for CHR are available (Abdennadher et al.; 1996; Abdennadher; 1998).

5.2. The Implementation in CHR. For the implementation in CHR, we simply have to declare the basic properties as constraints, see Sec. 3. Additionally, we use constraints for the sets of satisfied/nonsatisfied constraints, i.e., we have constraints `pplus/4` and `pminus/4` which encode in their first argument the set of constraints we are considering and an identifier, in their second argument the type of constraint and in the last two the categories involved.

The program takes the input string, a Prolog list of words, and traverses it.[3] For each lexical entry, it posts the corresponding category constraint, i.e., `cat(X,Y)` where `X` is a category and `Y` its position in the numeration.

The CHR implementation immediately starts out by building the elementary trees of the initial categories from the posted `cat/2` constraints via constraint propagation. In general, the CHR engine, i.e., all constraint resolution steps, is triggered as soon as constraints matching the heads are posted.

Basically, the CHR program looks at two category constraints and – if they are in the set of constraints we are considering – uses them together with a constraint from the grammar, as the head of constraint propagation rules. The constraints propagated contain the information whether the pair of categories satisfy the constraint from the grammar or not. In this way, we build a complete table of results on all pairs of categories.[4]

So, we build a complete table of `pplus/4, pminus/4` constraints for all pairs of categories matching the heads of our propagation rules. These constraints store the information about the interaction of any two categories with the resulting rule and set of constraints which licensed them. Nothing more has to be implemented to achieve this behaviour since it is built into the CHR engine. All we have to do is specify the right propagation rules and to post the

[3] Currently from left-to-right although this is arbitrary and can be varied.

[4] Since in the end we are interested in bigger chunks of lexical items then just two, we have to do some post-processing to interpret the entries in the table in the right way. For example, if a precedence constraint succeeds on two pairs of categories, this does not imply that it also succeeds for the entire triple. So, one violation among all the pairs is sufficient to cause failure of that constraint.

constraints from the grammar and input which makes for a very simple and easy to understand program.

Generally, in all the rules dealing with the constraints from the grammar, we apply the rule if we can find two categories within the constraint store which are in the same constituent and a constraint which checks some configuration on them. Only this last constraint triggers the rule since we declared all other constraints to be passive. Note that this ensures that we did all possible projections in each constituent before we post other constraints.

Since, if two category constraints match the head of a rule, they do so in either order, we use the guard to force the linear order as it appears in the input string.

```
prec @ cat(X,N)#I1, cat(Z,N)#I2, cat(Y,M)#I3, cat(Z,M)#I4,
                                    prec(Z,Cat1,Cat2,Id) ==>
       N < M |
       ( ( X == Cat2, Y == Cat1 ) ->
             pminus(Z-Id,prec,X-N,Y-M)
       ;
             pplus(Z-Id,prec,X-N,Y-M) )
       pragma passive(I1), passive(I2),
              passive(I3), passive(I4).
```

FIGURE 3. An example CHR rule

An example of such a propagation rule is given in Fig. 3. We are considering a precedence constraint `Id` on the categories `X` and `Y` within the constituent `Z` with indices `N` and `M`. We check satisfiability of the constraint in the head of the implication and then add the corresponding `pplus` or `pminus` constraint to the store.[5] Note that all the constraints actually appearing in the constraint store are ground such that we can test for literal identity.

The post-processing then just collects all the sets, gets all pairs of appearing categories if they are within the same range and collects all positive and negative constraints attached to them. While doing so, we interpret its members in the following way:

- if we have a successful obligatory or implicational constraint on any two categories within a constituent, we can ignore all the failed ones;

[5]We need four `cat` constraints to ensure that both categories belong to the same constituent, i.e., that `Z` indeed covers `N` and `M`.

- if we have a failed precedence or co-occurrence constraint on any two categories within a constituent, we have to ignore all the successful ones.

Furthermore, the output of the program – consisting of the `pplus/pminus`-sets generated by the resolution process – can be varied according to the given threshold of failed constraints, in the extreme case reducing the process to finding only the positive characterizations.

The Appendix contains the full commented code of the implementation.

6. CONCLUSION

The representation of linguistic information by means of constraints is interesting both for knowledge representation (no implicit information, encapsulation) and from an implementational point of view as well (parsing is seen as constraint solving). Property Grammar offers a framework taking advantage of these characteristics and presenting several interesting insights for natural language processing. The first one concerns the generality and the robustness of the approach: constraints allow the introduction of the notion of characterization to replace that of grammaticality. The parsing process consists in verifying constraints for the possible set of categories associated with the input instead of trying to build a structure according to the grammar. One other interesting point concerns the integration of different information sources: the properties can concern all kind of information (prosodic, syntactic, pragmatic etc.).

A possible direct application of Property Grammars would be the development of tools to work with corpora. Since the modular and flexible nature of these simple constraints is close to the descriptions used by empirical linguists, it could be used to test or find examples of particular grammatical constructions (or even violations of principles) easily by focusing the analysis and using a shallow analysis for the rest. But obviously this would presuppose the development of a dedicated, more efficient constraint solver first.

BIBLIOGRAPHY

Abdennadher, S. (1998). *Analyse von regelbasierten Constraintlösern*, PhD thesis, Ludwig-Maximilians-Universität München.

Abdennadher, S., Fruehwirth, T. and Meuss, H. (1996). On confluence of constraint handling rules, *Lecture Notes in Computer Science* **1118**: 1–15.

Archangeli, D. and Langendoen, D. (eds) (1997). *Optimality Theory*, Blackwell.

Bès, G. and Blache, P. (1999). Propriétés et analyse d'un langage, *Actes de TALN'99*.

Blache, P. (1999). Filtering and fusion: A technique for parsing with properties, *Proceedings of NLPRS'99*.

Blache, P. and Paquelin, J.-L. (1996). Active constraints for a direct interpretation of hpsg, *Proceedings of HPSG'96*.

Borsley, R. (1996). *Modern Phrase Structure Grammars*, Blackwell.

Duchier, D. and Thater, S. (1999). Parsing with tree descriptions: a constraint based approach, *Proceedings of* NLULP'99.

Frühwirth, T. (1998). Theory and practice of constraint handling rules, *Journal of Logic Programming* **37**(1–3): 95–138. Special Issue on Constraint Logic Programming.

Götz, T. and Meurers, D. (1997). Interleaving universal principles and relational constraints over typed feature logic, *Proceedings of the ACL/EACL Conference '97*, Association for Computational Linguistics, Madrid, Spain, pp. 1–8.

Günthner, F. (1988). Features and München.

Intelligent Systems Laboratory (1995). SICStus Prolog user's manual.

Marriott, K. and Stuckey, P. J. (1998). *Programming with Constraints*, MIT Press.

Maruyama, H. (1990). Structural disambiguation with constraint propagation, *Proc. of the 28th Annual Meeting of the Association for Computational Linguistics. Pittsburgh, PA, 1990*, pp. 31–38.

Meurers, D. and Minnen, G. (1998). Off-line constraint propagation for efficient hpsg processing, *in* G. Webelhuth, J.-P. Koenig and A. Kathol (eds), *Lexical and Constructional Aspects of Linguistic Explanation*, CSLI.

Sag, I. and Wasow, T. (1999). *Syntactic Theory. A Formal Introduction*, CSLI.

Shieber, S. M. (1992). *Constraint–Based Grammar Formalisms*, MIT Press.

APPENDIX: THE IMPLEMENTATION OF PROPERTY GRAMMARS IN CHR

```
:- use_module(library(chr)).
:- use_module(library(lists),[sublist/2,member/2,append/3]).


%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Constraint "Solver"
%
% Input:  a list of lexical entries
%         predicate process/1
%         example: process([la,fille,la,plus,jolie]).
%
% Output: the sets P+ and P- for all subsets of categories
%         and all models and the final constraint store
%
% Currently the grammar is built into this file, but it is
% easy to change this to some general procedure or change
% the example. Look for the predicate post_grammar/0 and
% lex/2.
%
% Note that the dependency constraints from the grammar are
% not used for anything at the moment since there is no
% information which makes it worthwhile checking them. If we
% just test whether we could add a dependency link, we could
% simply use implicational constraints which are already
% present.
%
% We do not create P+ and P- sets for singleton sets of
% categories although this might be important in the future?
%
%============================================================
%
% Description of the program:
%
% First of all we build the projection lines of the
% categories gained by lexical lookup by propagation via the
% alphabet constraint.
%
% In general the CHR program looks at two catagory
```

```
% constraints and - if they are in the same model - uses
% them together with a constraint from the grammar, as the
% head of constraint propagation rules. The constraints
% propagated contain the information whether the pair of
% categories satisfy the constraint from the grammar or
% not. In this way, we build a complete table of results on
% all pairs of categories.
%
% Since in the end we are interested in bigger chunks of
% lexical items then just two, we have to do some
% post-processing to interpret the entries in the table in
% the right way. For example, if a precedence constraint
% succeeds on some pairs of categories, this does not imply
% that it also succeeds for the entire triple. So, one
% violation among all the pairs is sufficient to cause
% failure of that constraint.
%
% I implemented this in a pretty nasty Prolog routine which
% collects all failed and successful constraints on the
% elements of any subset of the numeration of the input and
% then massages them into shape. This predicate is highly
% wasteful and has to be reimplemented. Maybe the entire
% approach of looking at pairs of categories has to be
% abandoned?
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% name the constraint handler
handler propgrams.

% declare all constraints
constraints cat/2, pplus/4, pminus/4,
            obli/3, alph/3, impl/4, prec/4, cooc/4, depe/4.
% cat(X,Y)
%     X is a category and Y its position in the numeration
% pplus(A-Id,Name,X-N,Y-M)
% ppminus(A-Id,Name,X-N,Y-M)
%     pplus  are the successful constraints
%     pminus are the failed constraints
%     A is the module name, Id the identifier of the
```

```
%     constraint, Name is the type of constraint
%     X-N, Y-M are categories with their numeration
% The following are the "real" linguistic constraints from
% the grammar:
% obli(Mod,Cat,Id)
%     Mod is the module and Cat is the obligatory category,
%     Id the identifier
% alph(Mod,Cats)
%     Mod is the module and Cats is the list of possible
%     categories
% impl(Mod,Cat1,Cat2,Id)
%     Mod is the module and Cat1 is the antecedent, cat2 the
%     succedent of the implication, Id the identifier
% prec(Mod,Cat1,Cat2,Id)
% cooc(Mod,Cat1,Cat2,Id)
% depe(Mod,Cat1,Cat2,Id)
%     parallel to the above

% enable the code for CHR debugging
% option(debug_compile,on).

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% CONSTRAINT HANDLING RULES
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% this rules projects the input categories within their
% module if we have a corresponding alphabet constraint
project @ cat(X,N), alph(Cat,Alph,_) ==>
        member(X,Alph) |
        cat(Cat,N).

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% We build a complete table of pplus/4, pminus/4 constraints
% for all pairs of categories matching the heads of our
% propagation rules. These constraints store the information
% about the interaction of any two categories with the
% resulting rule and model which licensed them.
```

```
%
% Generally, in all the rules dealing with the constraints
% from the grammar, we apply the rule if we can find two
% categories within the constraint store which are in the
% same module and a constraint which checks some
% configuration on them. This last constraint triggers the
% rule since we declared all other constraints to be
% passive. Note that this assumes that we did all possible
% projections in a model before we post the other
% constraints.
%
% Since, if two category constraints match the head of a
% rule, they do so in either order, we use the guard to
% force the linear order as it appears in the input string.
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% Obligatory constraints:
% if we have a model which matches an obligatory constraint,
% then we also have to have a category constraint with the
% right Cat.
        obli @ cat(X,N)#I1, cat(Z,N)#I2, cat(Y,M)#I3,
        cat(Z,M)#I4, obli(Z,Cat,Id) ==>
        N < M |
        ( ( X == Cat ; Y == Cat ) ->
            pplus(Z-Id,obli,X-N,Y-M)
        ;
            pminus(Z-Id,obli,X-N,Y-M)
        )
        pragma passive(Id1),passive(Id2),
               passive(Id3),passive(Id4).
% (Classical) Implicational constraints:
% succeeds if both anteceedent and succedent are present
impl @ cat(X,N)#I1, cat(Z,N)#I2, cat(Y,M)#I3, cat(Z,M)#I4,
       impl(Z,Cat1,Cat2,Id) ==>
        N < M |
        ( ( (X == Cat1, Y \== Cat2);
            (X \== Cat2, Y == Cat1) )  ->
            pminus(Z-Id,impl,X-N,Y-M)
        ;
```

```
            pplus(Z-Id,impl,X-N,Y-M)
        )
        pragma passive(Id1),passive(Id2),
               passive(Id3),passive(Id4).
% Precedence constraints:
% fails only if both category constraints appear in the
% wrong order
prec @ cat(X,N)#I1, cat(Z,N)#I2, cat(Y,M)#I3, cat(Z,M)#I4,
       prec(Z,Cat1,Cat2,Id) ==>
        N < M |
        ( ( X == Cat2, Y == Cat1 )  ->
            pminus(Z-Id,prec,X-N,Y-M)
        ;
            pplus(Z-Id,prec,X-N,Y-M)
        )
        pragma passive(Id1),passive(Id2),
               passive(Id3),passive(Id4).
% Cooccurrence constraints:
% fail if both category constrainst appear at the same time
cooc @ cat(X,N)#I1, cat(Z,N)#I2, cat(Y,M)#I3, cat(Z,M)#I4,
       cooc(Z,Cat1,Cat2,Id) ==>
        N < M |
        ( ( ( X == Cat1, Y == Cat2 );
            ( X == Cat2, Y == Cat1 ) )   ->
            pminus(Z-Id,cooc,X-N,Y-M)
        ;
            pplus(Z-Id,cooc,X-N,Y-M)
        )
        pragma passive(Id1),passive(Id2),
               passive(Id3),passive(Id4).

% some rules for absorbing spurious double constraints
absorb1 @ cat(X,N) \ cat(X,N) <=> true.
absorb2 @ pplus(X,Y,A,B) \ pplus(X,Y,A,B) <=> true.
absorb3 @ pminus(X,Y,A,B) \ pminus(X,Y,A,B) <=> true.


%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% GRAMMAR
```

```
% A simple collection of basic constraints and a
% lexicon. The lexicon are standard prolog facts.
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% the lexicon
% lex(+Word,+Category)
lex(la,det).
lex(fille,n).
lex(plus,adv).
lex(jolie,adj).

% this predicate justs posts the grammar constraints. beware
% that the projections within a model have to be done before
% any other processing.  the constrainst can be identified
% uniquely by their first and last arguments: the model and
% an identifier, i.e., a number.
post_grammar :-
        alph(np,[det,n,adjp,sup],1),
        alph(adjp,[adj,adv],1),
        alph(sup,[det,adv,adj],1),
        obli(np,n,2),
        obli(adjp,adj,2),
        obli(sup,adj,2),
        impl(np,n,det,3),
        impl(sup,adj,det,3),
        impl(sup,adj,adv,4),
        prec(np,det,n,4),
        prec(np,det,adjp,5),
        prec(np,det,sup,6),
        prec(np,n,adjp,7),
        prec(np,n,sup,8),
        prec(adjp,adv,adj,3),
        prec(sup,det,adj,5),
        prec(sup,det,adv,6),
        prec(sup,adv,adj,7),
        cooc(np,adjp,sup,9),
        depe(np,det,n,10),
        depe(np,adjp,n,11),
        depe(np,sup,n,12),
```

```
        depe(adjp,adv,adj,4),
        depe(sup,det,adj,8),
        depe(sup,adv,adj,9).



%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% Processor (not CHR any more but standard Prolog):
% We enumerate and post the input category-constraints
% followed by searching the database of pplus/4 and pminus/4
% constrainst for the generation of answers
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% process(+ListOfLexicalItems)
        process(In) :-
        numeration(In,1),
        post_grammar,
        generate_output.

% numeration(+List,+Num)
% posts cat/2 constrainst for all lexical elements in List and
% associates their string position with them.
numeration([],_).
numeration([H|T],N):-
        lex(H,Cat),
        cat(Cat,N),
        N1 is N + 1,
        numeration(T,N1).

% finds all category constraints and transforms them into a
% list of Cat-Num pairs which is then used to generate the
% possible subsets which in turn triggers reporting the
% results.
%
% the alphabet constraints, i.e., identity 1 is never part
% of the solution set since it is always true.
generate_output :-
        findall_constraints(cat(_X,_Y),List),
```

```
        extract_categories(List,Cats),
        powerset(Cats,PowCats),
        nl,
        write('++++++++++++++++++++++++++++++++++++++++'),nl,
        report_results(PowCats).

% extract_categories(+ListOfCatConstraints,-ListOfCat-NumPairs)
% careful! we reverse the order of the Cat-Num pairs such
% that we can use keysort on the list.
extract_categories([],[]).
extract_categories([cat(X,Y)#_|T],[Y-X|R]):-
        extract_categories(T,R).

% powerset(+ListOfCat-NumPairs,-ListOfSubLists)

% creates a list of lists such that the memebers are the
% non-trivial subsets of the list such that no two items in
% any sublist have the same string position.
powerset(L,PowerSet) :-
        findall(Sub1,
                (   sublist(Sub,L),
%                   length(Sub,3),  % for tests only sets
                                    % with card N
                    length(Sub,K),
                    K > 1,
                    \+ (   member(N-X,Sub),
                           member(M-Y,Sub),
                           X \== Y,
                           N == M
                       ),
                    sort(Sub,Sub1)
                ),
                Pow),
        sort(Pow,PowerSet).

% report_results(+ListOfSubLists)

% we take each set, get all pairs of appearing categories if
% they are in the same model and collect all positive and
% negative constraints attached to them. All of these
```

```
% results are collected in a list, post processed, sorted,
% and passed to the predicate report_results/1 for output.
%
% nasty, nondeterministic => has to be changed!!
report_results([]).
report_results([Set|T]):-
        findall(Model,
          ( find_constraint(alph(Model,List,_),_),
            all_member(Set,List)
          ),
                ModList),
        ( ModList == [] ->          % No common model
            Solutions = []
        ;
            findall(ResList,
              ( member(Mod,ModList),
                findall(res(Mod,Set,LPlus-LMinus),
                  ( member(N-X,Set),
                    member(M-Y,Set),
                    \+ X-N == Y-M,
                    N < M,
                    findall_constraints(
                       pplus(Mod-_,_,X-N,Y-M),LPlus),
                    findall_constraints(
                       pminus(Mod-_,_,X-N,Y-M),LMinus),
                    \+ (LPlus == [], LMinus == [])
                   ),
                        ResList)
                ),
                    Sols),
            post_process(Sols,Solutions)
        ),
        ( Solutions == [] ->
            true                    % switch reporting of
                                    % empty solutions
%           report_empty(Set)       % on or off
        ;
            report(Solutions)
        ),
        report_results(T).
```

```
% all_member(+ListOfCats,+AlphabetList)
% checks wether all elements of the ListOfCats are members
% of one model
all_member([],_).
all_member([_-H|T],L):-
        member(H,L),
        all_member(T,L).


% post_process(+ListOfResults,-ListOfResults)
% this is the awful predicate ;-)
% we flatten the list of results to one result which has
% been massaged into shape....
% we recurse on the list of results from report_results/1.
post_process([],[]).
post_process([H|T],[res(Mod,Set,LPlus-LMinus)|Out]) :-
        post_proc(H,Mod,Set,[],LPlus,[],LMinus),
        post_process(T,Out).

% post_proc(+ListOfResults,-Model,-Set,
%           -ListOfPosCons,-ListOfNegCons)
% the recursive case just flattens the results into two
% lists: the successfull constraints and the failed ones.
% the base case takes each lists and interprets its members
% in the following way:
% - if we have a successful obligatory or implicational
%   constraint on any two categories from the set, we can
%   ignore all the failed ones
% - if we have a failed precedence or co-occurrence
%   constraint on any two categories from the set, we have
%   to ignore all the successful ones
% finally, we strip all the ids generated by the CHR part
% from the constraints and sort the resulting lists.
post_proc([],_,_,LPlus,NewLPlus,LMinus,NewLMinus) :-
        findall(N1,
                member(pplus(_-N1,obli,_,_)#_,LPlus),
               LN1),
        remove(obli-LN1,LMinus,LMinus1),
        findall(N2,
```

```
                member(pplus(_-N2,impl,_,_)#_,LPlus),
                LN2),
        remove(impl-LN2,LMinus1,LMinus2),
        findall(N3,
                member(pminus(_-N3,prec,_,_)#_,LMinus2),
                LN3),
        remove(prec-LN3,LPlus,LPlus1),
        findall(N4,
                member(pminus(_-N4,cooc,_,_)#_,LMinus2),
                LN4),
        remove(cooc-LN4,LPlus1,LPlus2),
        strip_ids(LPlus2,[],LPlus3),
        strip_ids(LMinus2,[],LMinus3),
        sort(LPlus3,NewLPlus),
        sort(LMinus3,NewLMinus).
post_proc([res(Mod,Set,P1-M1)|T],Mod,Set,Pin,Pout,
                                              Min,Mout) :-
        append(Pin,P1,P2),
        append(Min,M1,M2),
        post_proc(T,Mod,Set,P2,Pout,M2,Mout).

% remove(+Keyword-RuleId,+ListOfCons,-ListOfCons)
% takes as input a list of constraints, a keyword and the
% list of identities of constraints from the grammar and
% removes them from the list.
remove(_-[],In,In) :- !.
remove(obli-N,In,Out) :-
        removem(In,obli,N,Out).
remove(impl-N,In,Out) :-
        removem(In,impl,N,Out).
remove(prec-N,In,Out) :-
        removep(In,prec,N,Out).
remove(cooc-N,In,Out) :-
        removep(In,cooc,N,Out).

% removem(+ListOfConstraints,+Keyword,
%         +ListOfIdentities-LotsOfCons)
% removes negative constraints matching the keyword and an
% identity
removem([],_,_,[]).
```

```
removem([pminus(_-N,X,_,_)#_|T],X,LN,Out) :-
        member(N,LN),
        !,
        removem(T,X,LN,Out).
removem([pminus(A-M,Y,C1,C2)#I|T],X,LN,
        [pminus(A-M,Y,C1,C2)#I|Out]) :-
        ( X \== Y ; \+ member(M,LN) ),
        !,
        removem(T,X,LN,Out).
% removep(+ListOfConstraints,+Keyword,
%         +ListOfIdentities-LotsOfCons)
% removes positive constraints matching the keyword and an
% identity
removep([],_,_,[]).
removep([pplus(_-N,X,_,_)#_|T],X,LN,Out) :-
        member(N,LN),
        !,
        removep(T,X,LN,Out).
removep([pplus(A-M,Y,C1,C2)#I|T],X,LN,
        [pplus(A-M,Y,C1,C2)#I|Out]) :-
        ( X \== Y ; \+ member(M,LN) ),
        !,
        removep(T,X,LN,Out).

% strip_ids(+ListOfCons,+Acc,-Acc)
% removes the trailing Identities form the CHR constraints
strip_ids([],_,[]).
strip_ids([H#_|R],In,Out):-
        H =.. [_,_-N|_],
        ( member(N,In) ->
            NewIn = In,
            Out = Out1
        ;
            NewIn = [N|In],
            Out = [H|Out1]
        ),
        strip_ids(R,NewIn,Out1).


% report(+ListOfResults)
```

```
% outputs the Model, the categories and lists of the
% positive and negative constraints.
report([]) :-
        write('++++++++++++++++++++++++++++++++++++++++'),
        nl,nl.
report([res(Mod,Set,H1-H2)|T]) :-
        keysort(Set,Set1),
        write('Categories: '),write_set(Set1),nl,
        write('P+( '),write(Mod), write(' ) = [ '),
        write_listp(H1),
        write(' ]'),nl,
        write('P-( '),write(Mod), write(' ) = [ '),
        write_listm(H2),
        write(' ]'),nl,nl,
        report(T).

% report_empty(+L)
% this usually means that we had a subset of categories
% which did not have a common model
report_empty(L) :-
        write('The following categories did not'),
        write(' produce any results: '),nl,
        write('    '),write(L),nl,nl,
        write('++++++++++++++++++++++++++++++++++++++++'),
        nl,nl.

% outputs the list of categories we are looking at.
% reverses the elements back.
write_set([]) :- write('ERROR....'),!,fail.
write_set([H-N]):-
        write(N-H),!.
write_set([H-N|T]):-
        write(N-H), write(' and '),
        write_set(T).

% outputs the identities of positive constraints
write_listp([]).
write_listp([pplus(_-K,_,_,_)]) :-
        write(K), !.
write_listp([pplus(_-K,_,_,_)|T]):-
```

```
        write(K), write(', '),
        write_listp(T).
% outputs the identities of negative constraints
write_listm([]).
write_listm([pminus(_-K,_,_,_)]) :-
        write(K), !.
write_listm([pminus(_-K,_,_,_)|T]):-
        write(K), write(', '),
        write_listm(T).
%
% end of file
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```