Extended Dependency Structures and their Formal Interpretation

Marc Dymetman Max Copperman

Rank Xerox Research Centre

6, chemin de Maupertuis, Meylan 38240, France

{dymetman, copperman}@xerox.fr

Abstract

We describe two “semantically-oriented” dependency-structure formalisms, U-forms and S-forms. U-forms have been previously used in machine translation as interlingual representations, but without being provided with a formal interpretation. S-forms, which we introduce in this paper, are a scoped version of U-forms, and we define a compositional semantics mechanism for them. Two types of semantic composition are basic: complement incorporation and modifier incorporation. Binding of variables is done at the time of incorporation, permitting much flexibility in composition order and a simple account of the semantic effects of permuting several incorporations.

1 INTRODUCTION

U-forms (Unscoped dependency form) are a representation formalism which has been used (under a different name) as the basis for the intermediary language in the machine translation system CRITTER (Isabelle et al., 1988; Dymetman, 1992; Isabelle, 1995). U-forms account for two central aspects of linguistic structure: predicate-argument relations and headedness (complements vs. modifiers), and so form a middle ground between a “semantic” and a “syntactic” representation. This, combined with their formal simplicity, accounts for much of the popularity of U-forms or related formalisms — such as the semantic and deep syntactic representations used in Mel’čuk’s Meaning-Text Theory (Mel’čuk, 1981) — in applications such as machine translation and text generation.

Although U-forms are strongly “meaning-oriented”, their interpretation is never made explicit but is left to the computational linguist’s intuition. This has two consequences:

- Operations performed on U-forms and related formalisms cannot be controlled for semantic va-

lidity. So, for instance, it is common practice to define graph rewriting rules on these representations which are believed to produce semantically equivalent expressions. Without the check of formal interpretation, these rules may work in some cases, but produce wrong results in other cases. So for instance, a rule rewriting (the representation of) “John’s salary is \$25000 higher this year than last year” into “John’s salary was \$25000 lower last year than this year” would seem intuitively valid until one considered the case of “John’s salary is 50% higher this year than last year”, where it does not work any more.

- U-forms are not directly adapted to applications putting emphasis on denotational semantics and formal reasoning, like for instance some natural language generation systems in well-formalized domains (Huang and Fiedler, 1995; Ranta, 1995; Levine, 1990), see also (Alshawi, 1992).

A basic obstacle to providing a formal interpretation for U-forms is the fact that these representations leave the relative scopes of dependents implicit. The S-form representation (Scoped dependency form), which we introduce here, is an extension of U-form notation which makes scope explicit, by allowing dependents to be ordered relative to one another. Dependents (complements or modifiers) can move freely relative to one another in the S-form structure, under certain binding-site constraints.

We then go on to provide a compositional interpretation mechanism for S-forms. Free variables (generalizations of the arg_1 , arg_2 , arg_3 annotations of standard dependency formalisms) are used to connect an argument to its binding-site inside a predicate. Binding of variables is done at the time of incorporation, permitting much flexibility in composition order and a simple account of the semantic effects of permuting several incorporations. This liberal use of free variables is contrasted to the approach of Montague grammar, where the requirement that semantic expressions entering into a composition are closed (do not contain

free variables) leads to a certain rigidity in the order of composition.

Two kinds of semantic composition are basic: complement incorporation, where the complement fills a semantic role inside the head, and modifier incorporation, where the head fills a semantic role inside the modifier. The mechanism of actually deriving the semantic translation of the composition from the semantic translations of its two components is handled through a list of type-sensitive composition rules, which determine the action to be taken on the basis of the component types. The flexibility of the approach is illustrated on an example involving proper names, quantified noun phrases, adverbials and relative clauses.

2 U-FORMS

Formally, U-forms are unordered labelled n-ary trees such as the one shown in Fig. 1, corresponding to the sentence: (S1) "John does not like every woman hated by Peter".

Figure 1: A U-form.

The edge labels are members of the set $\{\text{det}, 1, 2, 3, \dots, -1, -2, -3, \dots\}$, and correspond either to determiners (label 'det') or to argument positions relative to a predicate node (other labels).

The U-form of Fig. 1 expresses three predicate-argument relations among the nodes:

Figure 2: Predicate-argument relations in a U-form.

In order to extract the predicate-argument relations encoded into the U-form, one needs to apply the following "rule". Let's notate (A,L,B) an edge of the

tree, where A is the upper vertex, B the lower vertex, and L the edge label. With each node A in the tree, one associates its set of *predication edges*, that is the set PA_A of edges of the form (A,i,X) or $(X,-i,A)$. One then considers the *predication tree* T_A made by forming the collection of edges (A,L,X) where L is positive and either (A,L,X) or $(X,\text{inverse}(L),A)$ is a predication edge of A. Each predication tree denotes a predicate-argument relation among U-form nodes. So for instance, the tree T_{hate} is formed by forming the edges $(\text{hate},1,\text{peter})$ and $(\text{hate},2,\text{woman})$, and this corresponds to the predicate-argument relation $\text{hate}(\text{peter},\text{woman})$.

WELL-FORMEDNESS CONDITIONS ON U-FORMS In order to be well-formed, a U-form UF has to respect the following condition. For any node A of UF, the predication tree T_A must be such that:

1. [No holes condition] If (A,i,B) is an edge of T_A , then for any number j between 1 and i, T_A must contain a node of form (A,j,C) .
2. [No repetition condition] No two edges of T_A can have the same label i.

MORE ON U-FORMS Negative labels are a device which permits to reconcile the notation of predicate-argument structure with the notation of syntactic dependency. So, in the U-form considered above, while "semantically" the 'woman' node is an argument of the 'hate' node, "syntactically" the 'hate' node is a dependent of the 'woman' node. Cases such as this one, where there is a conflict between predicate-argument directionality and dependency directionality are notated in the U-form through negative labels, and correspond to *modifiers*. Cases where the directionality is parallel correspond to *complements*.

When used as interlingual representations in machine translation systems, U-forms have several advantages. The first is that they neutralize certain details of syntactic structure that do not carry easily between languages. For instance, French and English express negation in syntactically different ways: "Rachel does not like Claude" vs. "Rachel n'aime pas Claude"; this difference is neutralized in the U-form representation, for both negations are expressed through a single negation predicate in the U-form.

A second advantage is that they represent a good compromise between paraphrasing potential and semantic precision. So, for instance, in the CRITTER system, the three sentences:

John does not like every woman that Peter hates
John does not like every woman hated by Peter
Every woman whom Peter hates is not liked by John

would be assigned the U-form of Fig. 1. On the other hand, the sentence:

Peter hates every woman that John does not like
 would be assigned the U-form of Fig. 3, which is different from the previous U-form, although the predicate-argument relations are exactly the same in both cases.

Figure 3: A different U-form.

One can take advantage of such paraphrasing potential in certain cases of syntactic divergence between languages. For instance, French does not have a syntactic equivalent to the dative-movement + passive configuration of:

Rachel was given a book by Claude

so that a direct syntactic translation is not possible. However, at the level of U-form, this sentence is equivalent to the French sentence:

Claude a donné un livre à Rachel

and this equivalence can be exploited to provide a translation of the first sentence.

One serious problem with U-forms, however, is that they do not have unambiguous readings in cases where the relative scopes of constituents can result in different semantic interpretations. So, in the case of sentence (S1), the two readings: "it is not the case that John likes every woman hated by Peter", and "John dislikes every woman that Peter hates" are not distinguished by the U-form of Fig. 1.

3 S-FORMS

INTRODUCING SCOPE Let's consider the tree represented in Fig. 4.

The only difference between this tree and the U-form of Fig. 1 is that the nodes of our new tree are considered ordered whereas they were considered unordered in the U-form. The convention is now that dependent sister nodes are interpreted as having different scopes, with narrower scope corresponding to a position more to the right.

The tree of Fig. 4 can be glossed in the following way:

Figure 4: Introducing scope by ordering the nodes.

John, it is not the case that he likes every woman that Peter hates

If we consider the six permutations of the nodes under like, we can produce six different scopings. Because John refers to an individual, not a quantified NP, these six permutations really correspond to only the two interpretations given above. The tree of Fig. 4 corresponds to the first of these interpretations, which is the preferred interpretation for sentence (S1).

Our discussion of scope being represented by node order has been informal so far. In order to make it formal, we need to encode our representation into a binary-tree format on which a compositional semantics can be defined. To do that, in a first step we replace the argument numbers of Fig. 4 by explicit argument names; in a second step we encode the resulting ordered n-ary tree into a binary format which makes explicit the order in which dependents are incorporated into their head.

S-FORMS Consider the n-ary tree of Fig. 4. For any node A in this tree, take the set of predication edges associated with A, that is the set of edges $(A, +i, B_i)$ and $(B_i, -i, A)$. By renaming each such node A into $A(X_1, \dots, X_n)$, where $X_1, \dots, X_n$ are fresh identifiers, and by renaming each such label $+i$ (resp. $-i$) into $+X_i$ (resp. $-X_i$), one obtains a new tree where argument numbers have been replaced by argument names. For instance the previous representation now becomes the tree of Fig. 5.

This representation is called a *scoped dependency form*, or *S-form*.

BINARY TREE ENCODING OF S-FORMS: B-FORMS In order to encode the ordered n-ary tree into a binary tree, we need to apply recursively the transformation illustrated in Fig. 6, which consists in forming a "head-line", projecting in a north-west direction from the head H, and in "attaching" to this line "dependent-lines" $D_1, D_2, \dots, D_n$, with D_1 the right-

Figure 5: An S-form.

most dependent (narrowest scope) and D_n the leftmost dependent (widest scope) in the original tree.

Figure 6: The transformation between S-forms and B-forms.

Applying this encoding to our example, we obtain the binary tree of Fig. 7, which is called a *B-form*.

The B-form makes explicit the order of incorporation of dependents into the head-line. By permuting several dependent-lines along their head-line, this incorporation order is changed and gives rise to different scopings.

S-forms and B-forms are completely equivalent representations. Clearly, the encoding, called the *S-form/B-form encoding*, which has just been defined is reversible. The S-form is more compact and makes the dependency relations more conspicuous, whereas the B-form makes the compositionality more explicit.

WELL-FORMEDNESS CONDITIONS ON B-FORMS AND S-FORMS Starting from the U-form and enriching it, we have informally introduced the notions of S-form and B-form. We now define them formally.

We start by giving a recursive definition of IBFs (incomplete B-forms), that is, B-forms which may contain unresolved free variables. We use the notation $((D, \text{Label}), H)$ the labelled binary tree obtained by taking H as the right subtree, D as the left subtree, and by labelling the left edge with Label . We also use the notation $\text{fv}(\text{IBF})$ for the set of the free variables in IBF.

DEFINITION OF INCOMPLETE B-FORMS

Figure 7: A B-form.

1. A node N of the form $\text{Pred}(x_1, \dots, x_n)$ is an IBF with the set of free variables $\text{fv}(N) = \{x_1, \dots, x_n\}$;
2. If D and H are IBFs, $\text{fv}(D)$ and $\text{fv}(H)$ are disjoint, and $x \in \text{fv}(H)$ then $H' = ((D, +x), H)$ is an IBF with $\text{fv}(H') = \text{fv}(D) \cup \text{fv}(H) \setminus \{x\}$;
3. If D and H are IBFs, $\text{fv}(D)$ and $\text{fv}(H)$ are disjoint, and $x \in \text{fv}(D)$ then $H' = ((D, -x), H)$ is an IBF with $\text{fv}(H') = \text{fv}(D) \cup \text{fv}(H) \setminus \{x\}$;
4. If D and H are IBFs, and $\text{fv}(D)$ and $\text{fv}(H)$ are disjoint, then $H' = ((D, \text{det}), H)$ is an IBF with $\text{fv}(H') = \text{fv}(D) \cup \text{fv}(H)$.

DEFINITION OF B-FORMS A B-form is an IBF with an empty set of free variables.

The notion of S-form can now be defined through the use of the S-form/B-form encoding.

DEFINITION OF S-FORMS A S-form is an ordered labelled n -ary tree which can be obtained from a B-form through the inverse application of the S-form/B-form encoding.

It can be easily verified that the representation of Fig. 7 is indeed a B-form, and, consequently, the representation of Fig. 5 is a valid S-form. More generally, it can be easily verified that enriching a U-form by ordering its nodes, and then replacing argument variables by argument names always results in a valid S-form.¹

¹The converse is not true: not all S-forms can be obtained in this way from a U-form. For instance, there exists a S-form corresponding to the preferred reading for "Fido visited most trashcans on every street", which has "every street"

4 THE INTERPRETATION PROCESS

We now describe the interpretation process on B-forms. Interpretation proceeds by propagating semantic translations and their types bottom-up.

The first step consists in typing the leaves of the tree, while keeping track of the types of free variables, as in Fig. 8.

Figure 8: Typing the leaves. The free variables and their types are indicated in brackets.

The types given to the leaves of the tree are the usual functional types formed starting with e (entities) and t (truth values). In the case where the leaf entity contains free variable arguments, the types of these free variables are indicated, and the type of the leaf takes into account the fact that these free variables have already been included in the functional form of the leaf. Thus $\text{hate}(h1, h2)$, which can be glossed as: “ $h1$ hates $h2$ ”, is given type t , while $h1$ and $h2$ are constrained to be free variables of type e .

VARIABLE-BINDING RULES According to the well-formedness conditions for B-forms, a complement incorporation $((D, +x), H)$ is only possible when H contains x among its free variables; the “syntactic dependent” D is seen as semantically “filling” the place that x occupies in the “syntactic head” H . In the same way, a modifier incorporation $((D, -x), H)$ is only possible when D contains x among its free variables;

outscoping “most trasheans”, and which is not obtained from a U-form in this simple way. However, there exists a mapping from S-forms to U-forms, the *scope-forgetting mapping*, which permits to define equivalence classes among S-forms “sharing” the same U-form. This relation between S-forms and U-forms can be used to give a (non-deterministic) formal interpretation to U-forms, by considering the interpretations of the various S-forms associated with it (see the technical report companion to this paper.)

in this case the “syntactic” head H is seen as semantically “filling” the place that x occupies in the “syntactic dependent” D . (This difference corresponds to the opposition which is sometimes made between syntactic and semantic heads and dependents: complements are dependents both syntactically and semantically, while modifiers are syntactically dependents but semantically heads.)

In order to make formal sense of the informal notion “filling the place of x in A_x ” (where the notation A_x means that A contains the free variable x), we introduce the variable-binding rules of Fig. 9.

Figure 9: Variable-binding rules. D' and H' correspond to the semantic translation of the subtrees rooted in D and H respectively.

These rules tell us how to “get rid” of the free variable being bound during complement or modifier incorporation, namely by forming the abstraction $\lambda x. A_x$ before actually performing the semantic composition between the dependent and the head. For completeness, determiner incorporation, which does not involve variable binding, is given along with complement and modifier incorporation.

Two things should be noted about this way of “delaying” variable-binding until the relevant dependent is incorporated:

- Suppose that we had bound the variables appearing in the head predicate locally, that is to say, that, in the style of Montague grammar (Gamut, 1991), we had written $\lambda I2 I1. \text{like}(I1, I2)$ instead of $\text{like}(I1, I2)$, and so forth, in Fig. 7. Then each incorporation of a dependent into the “head-line” would have changed the type of the head; thus ‘not’ would have had to combine either with a head of type $e \rightarrow e \rightarrow t$, or $e \rightarrow t$, or t , depending on its scope relative to the other dependents; with the scheme adopted here, the type of the head remains invariant along the head-line;
- Under the same hypothesis, the incorporation of the second argument first and of the first argument second would have been much simpler than

the reverse incorporation order, and some mechanism would have had to be found to distinguish the two orders. Then permuting the relative order of two dependents along the head-line — corresponding to different scope possibilities — would have had complex computational consequences. In the scheme adopted here, these cases are handled in a uniform way.

The way free variables are used in our scheme is somewhat reminiscent of the use of *syntactic variables* he_n in Montague grammar. Montague grammar has the general requirement that only closed lambda-terms (lambda terms containing only bound variables) are composed together. This requirement, however, is difficult to reconcile with the flexibility needed for handling quantifier scope ambiguities. Syntactic variables are a device which permit to “quantify into” clauses at an arbitrary time, bypassing the normal functional composition of lambda-terms, which requires a strict management of incorporation order. In our scheme, by contrast, this secondary mechanism of Montague grammar is graduated to a central position. Composition is always done between two lambda-terms one of which at least contains a free variable which gets bound at the time of incorporation.

TYPE SENSITIVE COMPOSITION RULES If we apply the variable-binding rules to the subtree $PH = ((peter.h1).hate(h1.h2))$ of Fig. 8, we find that we must compose the semantic translations $peter$ and $\lambda h1.hate(h1.h2)$ in “complement” (+) mode. The first function is of type e , while the second function is of type $e \rightarrow t$ (for $hate(h1.h2)$ is of type t , and $h1$ of type e).

How do we compose two such functions? A first solution, in the spirit of Lambek calculus (Morrill, 1994) or of linear logic (Dalrymple et al., 1995), would be to define a general computational mechanism which would be able, through a systematic discipline of type-changing operations, to “adapt” automatically to the types of the functions undergoing composition.

Such mechanisms are powerful, but they tend to be algorithmically complex, to be non-local, and also to give rise to spurious ambiguities (superficial variations in the proof process which do not correspond to different semantic readings).

Here, we will prefer to use a less general mechanism, but one which has two advantages. First, it is local, simple, and efficient. Second, it is flexible and can be extended to handle the semantics of sentences extracted from a real corpus of texts, which it might be perilous to constrain too strongly from the start.

The mechanism is the following. We establish a list of acceptable “type-sensitive composition rules”, which tell us how to compose two functions according

to their types. Such a (provisory) list is given below:²

- (C1) composition(+, L:T $\rightarrow$ S, R:T, L(R):S)
- (C2) composition(+, L:e, R:e $\rightarrow$ t, R(L):t)
- (C3) composition(det, L:T $\rightarrow$ S, R:T, L(R):S)
- (C4) composition(-, L:T $\rightarrow$ S, R:T, L(R):S)
- (C5) composition(-, L:e $\rightarrow$ t, R:e $\rightarrow$ t,
 $\lambda x.R(x) \wedge L(x):e \rightarrow t$)

The entries in this list have the following format. The first argument indicates the type of composition (+ for complement incorporation, - for modifier incorporation, det for determiner incorporation); the second argument is of the form Left:LeftType, where Left is the left translation entering the composition, and LeftType is its type; similarly, the second argument Right:RightType corresponds to the right subtree entering the composition; finally the third argument gives the result Result:ResultType of the composition, where the notation $A(B)$ has been used to indicate standard functional application of function A on argument B . Uppercase letters indicate unifiable variables.

It may be remarked that if, in these rules, we neglect the functions themselves (Left, Right, Result) and concentrate on their types (LeftType, RightType, ResultType), then the rules can be seen as imposing constraints on what can count as validly typed trees: these constraints can flow from mother to daughters as well as in the opposite direction. Thus, through these rules, knowing that the head-line functions projecting from a verbal head must be of type t imposes some constraints on what are the possible types for the dependents: this can be useful in particular for constraining the types of semantically ambiguous lexical elements.

If we now go back to our example, we have to compose in complement mode (+) the function $peter$, of type e , with the function $\lambda h1.hate(h1.h2)$, of type $e \rightarrow t$. Consulting the list of composition rules, we see that the only applicable rule is (C2), and that the result is $\lambda h1.hate(h1.h2)(peter) = hate(peter.h2)$, of type t .

Now that we have the semantic translation $hate(peter.h2)$ for the subtree PH , we can compute the translation for the subtree $((PH.h2).woman)$. By the variable-binding rule for modifiers, we need first to form the abstraction $\lambda h2.hate(peter.h2)$, of type $e \rightarrow t$, and compose it in ‘-’ mode with $woman$, of type $e \rightarrow t$. Consulting the list of composition rules, we find that the only applicable rule is (C5), and that the result of this application is $\lambda h2.woman(h2) \wedge hate(peter.h2)$.³

²It is a matter for further research to propose principles for producing such rules. Some of them can be seen as special cases of general type-raising principles, others (such as C5) are necessary if one accepts that the type of intersective adjectives and restrictive relative clauses has to be $e \rightarrow t$.

³The rule (C5) differs from the previous rules in the list in that it introduces the logical connective $\wedge$ which does not originate in functional material already present in either of the arguments. A possible justification for the rule, however,

Figure 10: B-form interpretation. For ‘every’, we make use of the generalized quantifier notation *quant(restriction,scope)*.

The process of semantic translation proceeds in this way bottom-up on the B-form. The end result is shown in Fig. 10.

Acknowledgments

Thanks to Alain Lecomte and Frédérique Segond for comments and discussions.

References

- Hiyan Alshawi, editor. 1992. *The Core Language Engine*. The MIT Press.
- J. Barwise and R. Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4.
- Mary Dalrymple, John Lamping, Fernando C. N. Pereira, and Vijay Saraswat. 1995. A deductive account of quantification in LFG. In Makoto Kanazawa, Christopher J. Piñón, and Henriette de Swart, editors, *Quantifiers, Deduction, and Context*. Center for the Study of Language and Information, Stanford, California.

Marc Dymetman. 1992. Transformations de grammaires logiques et réversibilité en Traduction Automatique. Thèse

is that it allows conferring the “natural” type $e \rightarrow t$ to an (intersective) adjective such as “black”, or for a relative modifier such as “hated by peter”, and also that there does not seem to exist any good reason why type composition should be restricted to “functionally matching” types only. Semantic type coercions abound in natural language, as in the case of “glass elephant”, “short win”, etc., and these require complex composition operations on the elements combined.

d’Etat. Université Joseph Fourier (Grenoble 1), Grenoble, France.

L. T. F. Gamut. 1991. *Logic, Language, and Meaning*, volume 2: Intensional Logic and Logical Grammar. The University of Chicago Press, Chicago.

Xiaorong Huang and Armin Fiedler. 1995. Generating multilingual proofs. In Richard Kittredge, editor, *IJCAI Workshop on Multilingual Text Generation*, pages 54–63, August.

Pierre Isabelle, Marc Dymetman, and Elliott Macklovitch. 1988. CRITTER: a translation system for agricultural market reports. In *Proceedings of the 12th International Conference on Computational Linguistics*, pages 261–266, Budapest, August.

Pierre Isabelle. 1995. *Un modèle linguistique calculable pour la traduction automatique bidirectionnelle*. Ph.D. thesis, Université de Montréal.

John Levine. 1990. Pragma – a flexible bidirectional dialogue system. In *Proceedings, Eight National Conference on Artificial Intelligence*, pages 964–969.

Igor A. Mel’čuk. 1987. *Dependency Syntax: Theory and Practice*. State University of New York Press.

Igor Mel’čuk. 1981. Meaning-text models. *Journal of Anthropology*, 10:27–62.

Glyn V. Morrill. 1994. *Type Logical Grammar: Categorical Logic of Signs*. Kluwer Academic Publishers, Dordrecht, Holland.

Aarne Ranta. 1995. *Type Theoretical Grammar*. Oxford University Press.