

Translating near-synonyms: Possibilities and preferences in the interlingua

Philip Edmonds

Department of Computer Science
University of Toronto
Toronto, Canada M5S 3G4
pedmonds@cs.toronto.edu

1 Introduction

This paper argues that an interlingual representation must explicitly represent some parts of the meaning of a situation as *possibilities* (or preferences), not as necessary or definite components of meaning (or constraints). Possibilities enable the analysis and generation of nuance, something required for faithful translation. Furthermore, the representation of the meaning of words is crucial, because it specifies which nuances words can convey in which contexts.

In translation it is rare to find the exact word that faithfully and directly translates a word of another language. Often, the target language will provide many near-synonyms for a source language word that differ (from the target word and among themselves) in nuances of meaning. For example, the French *fournir* could be translated as *provide*, *supply*, *furnish*, *offer*, *volunteer*, *afford*, *bring*, and so on, which differ in fine-grained aspects of denotation, emphasis, and style. (Figures 1 and 2 show some of the distinctions.) But none of these options may carry the right nuances to match those conveyed by *fournir* in the source text; unwanted extra nuances may be conveyed, or a desired nuance may be left out. Since an exact match is probably impossible in many situations, faithful translation will require uncovering the nuances conveyed by a source word and then determining how the nuances can be conveyed in the target language by appropriate word choices in any particular context. The inevitable mismatches that occur are one type of *translation mismatch*—differences of meaning, but not of form, in the source and target language (Kameyama et al., 1991).¹

¹A separate class of difference, *translation divergence*, involves differences in the form of the source and target texts and results from lexical gaps in the target language (in which no single word lexicalizes the meaning of a source word), and from syntactic and collocational constraints imposed by the source language. ‘Paraphrasing’ the source text in the target language is required in order to preserve the meaning as much as possible (Dorr, 1994; Stede, 1996; Elhadad et al., 1997). But even when paraphrasing, choices between near-synonyms will have to be made, so, clearly, translation mismatches and translation divergences are not independent phenomena. Just as standard semantic content can be incorporated or spread around in different ways, so can nuances of meaning.

Provide may suggest foresight and stress the idea of making adequate preparation for something by stocking or shipping ...

Supply may stress the idea of replacing, of making up what is needed, or of satisfying a deficiency.

Furnish may emphasize the idea of fitting something or someone with whatever is necessary, or sometimes, normal or desirable.

Figure 1: An abridged entry from *Webster’s New Dictionary of Synonyms* (Gove, 1973).

Offer and **volunteer** may both refer to a generous extending of aid, services, or a desired item. Those who *volunteer* agree by free choice rather than by submission to selection or command.

Figure 2: An abridged entry from *Choose the Right Word* (Hayakawa, 1994).

2 Near-synonyms across languages

This section examines how near-synonyms can differ within and across languages. I will discuss some of the specific problems of lexical representation in an interlingual MT system using examples drawn from the French and English versions of the multi-lingual text provided for this workshop.

To be as objective as possible, I’ll rely on several dictionaries of synonym discrimination including, for English, Gove (1973) and Hayakawa (1994), and for French, Bailly (1970), Bénac (1956), and Batchelor and Offord (1993). Unless otherwise stated, the information on differences below comes from one of these reference books.

Notation: Below, ‘*english::french*’ indicates that the pair of words or expressions *english* and *french* correspond to one another in the multi-lingual text (i.e., they are apparent translations of each other).

Fine-grained denotational mismatches

If a word has near-synonyms, then they most likely differ in fine-grained aspects of denotation. Consider the following pairs:

- 1a. *provides* :: *fournit*
- b. *provided* :: *apportaient*
- c. *provide* :: *offrir*
- d. *brought* :: *fournissait*
- e. *brought* :: *se chargeait*

These all share the basic meaning of giving or making available what is needed by another, but each adds its own nuances. And these are not the only words that the translator could have used: in English, *furnish*, *supply*, *offer*, and *volunteer* would have been possibilities; in French, *approvisionner*, *muni*, *pourvoir*, *nantir*, *présenter*, among others, could have been chosen. The differences are complex and often language-specific. Figures 1 and 2 discuss some of the differences between the English words, and figures 3 and 4 those between the French words. And this is the problem for translation: none of the words match up exactly, and the nuances they carry when they are actually used are context-dependent. (Also notice that the usage notes are vague in many cases, using words like ‘may’ and ‘idée’.)

Consider this second example:

- 2a. *began* :: *amorcé*
- b. *began* :: *commença*
- c. *started* :: *au début*

Amorcer implies a beginning that prepares for something else; there is no English word that carries the same nuance, but *begin* appears to be the closest match. *Commencer* also translates as *begin*, although *commencer* is a general word in French, implying only that the thing begun has a duration. In English, *begin* differs from *start* in that the latter can imply a setting out from a certain point after inaction (in opposition to *stop*).

More pairings that exhibit similar fine-grained denotational differences include these:

- 3a. *broaden* :: *élargir*
- b. *expand* :: *étendre*
- c. *increase* :: *accroître*
- 4a. *transformation* :: *passer*
- b. *transition* :: *transition*
5. *enable* :: *permettre*
6. *opportunities* :: *perspectives*
7. *assistance* :: *assistance*

There are two main problems in representing the meanings of these words. First, although some of the nuances could be represented by simple features, such as ‘foresight’ or ‘generous’, most of them cannot because they are complex and have an ‘internal’ structure. They are concepts that relate aspects of the situation. For example, for *furnish*, ‘fitting someone with what is necessary’ is not a simple feature; it involves a concept of ‘fitting’,

Fourni a rapport à la quantité et ce dit de ce qui à suffisamment ou en abondance le nécessaire.

Muni et **armé** sont relatifs à l’état d’une chose rendue forte ou capable, *muni*, plus générale, annonçant un secours pour faire quoi que ce soit.

Pourvu comporte un idée de précaution et ce dit bien en parlant des avantages naturels donnés par une sorte de finalité ...

Nanti, muni d’un gage donné par un débiteur à son créancier, par ext. muni par précaution et, absolument, assez enrichi pour ne pas craindre l’avenir.

Figure 3: An abridged entry from Bénac (1956).

Offrir, c’est faire hommage d’une chose à quelqu’un, en manifestant le désir qu’il l’accepte, afin que l’offre devienne un don.

Présenter, c’est offrir une chose que l’on tient à la main ou qui est là sous les yeux et dont la personne peut à l’instant prendre possession.

Figure 4: An abridged entry from Bailly (1970).

a patient (the same patient that the overall situation has), a thing that is provided, and the idea of the necessity of that thing to someone. Thus, many nuances must be represented as fully-fledged concepts (or instances thereof) in an interlingua.

Second, many of the nuances are merely suggested or implied, if they are conveyed at all. That is, they are conveyed indirectly—the reader has the license to decide that such a nuance was unintended—and as such are not necessary conditions for the definition of the words. This has ramifications for both the analysis of the source text and the generation of the target text because one has to determine how strongly a certain nuance is intended, if at all (in the source), and then how it should be conveyed, if it can be, in the target language. One should seek to translate indirect expressions as such, and avoid making them direct. One must also avoid choosing a target word that might convey an unwanted implication. In any case, aspects of word meaning that are indirect must be represented as such in the lexicon.

Coarse-grained denotational mismatches

Sometimes the translator chooses a target word that is semantically quite different from the source word, yet still conveys the same basic idea. Considering pair 1e, above: *bring* seems to mean to carry as a contribution, and *se charger* to take responsibility for. Perhaps there are no good equivalents in the opposite languages for these terms, or alternatively, the words might have been chosen because of syntactic or collocational preferences—they co-occur with *leadership* :: *l’administration*, which are not close translations either.

In fact, the desire to use natural-sounding syntactic

and collocational structures is probably responsible for many of these divergences. In another case, the pair *factors :: raisons* occurs perhaps because the translator did not want to literally translate the expressions *Many factors contributed to :: Parmi les raisons de*. Such mismatches are outside the scope of this paper, because they fall more into the area of translation divergences. (See Smadja et al. (1996) for research on translating collocations.)

Stylistic mismatches

Words can also differ on many stylistic dimensions, but formality is the most recognized dimension.² Consider the following pairs:

- 8a. *plans :: entend bien*
- b. *plan :: envisagent de*

While the French words differ in formality (*entend bien* is formal, and *envisagent de* is neutral), the same word was chosen in English. Note that the other French words that could have been chosen also differ in formality: *se proposent de* has intermediate formality, and *comptent, avont l'intention, and projettent de* are all neutral.

Similarly, in 2, above, *amorcer* is more formal than *commencer*. Considering the other near-synonyms: the English *commence* and *initiate* are quite formal, as is the French *initier*. *Débuter* and *démarrer* are informal, yet both are usually translated by *begin*, a neutral word in English. (Notice also that the French cognate of the formal English *commence, commencer*, is neutral.)

Style, which can be conveyed by both the words and the structure of a text, is best represented as a global property in an interlingual representation. That way, it can influence all decisions that are made. (It is probably not always necessary to preserve the style of particular words across languages.)

A separate issue of style in this text is its use of technical or domain-specific vocabulary. Consider the following terms used to refer to the subject of the text:

- 9a. *institution :: institution*
- b. *institution :: établissement*
- c. *institution :: association*
- d. *joint venture :: association*
- e. *programme :: association*
- f. *bank :: établissement*
- g. *bank :: banque*

In French, it appears that *association* must be used to refer to non-profit companies and *établissement* or *banque* for their regulated (for-profit) counterparts. In English *institution*, among other terms, is used for both. Consider also the following pairs:

- 10a. *seed capital :: capital initial*
- b. *working capital :: fonds de roulement*
- c. *equity capital :: capital social*

Attitudinal mismatches

Words also differ in the attitude that they express. For example, of *poor :: démunis*, *poor* can express a derogatory attitude, but *démunis* (which can be translated as *impoverished*) probably expresses a neutral attitude. Consider also *people of indigenous background :: Indiens*. Attitudes must be included in the interlingual representation of an expression, and they must refer to the specific participant(s) about whom the speaker is expressing an attitude.

3 Representing near-synonyms

Before I discuss the requirements of the interlingual representation, I must first discuss how the knowledge of near-synonyms ought to be modelled if we are to account for the complexities of word meaning in an interlingua. In the view taken here, the lexicon is given the central role as bridge between natural language and interlingua.

The conventional model of lexical knowledge, used in many computational systems, is not suitable for representing the fine-grained distinctions between near-synonyms (Hirst, 1995). In the conventional model, knowledge of the world is represented by ostensibly language-neutral concepts that are often organized as an ontology. The denotation of a lexical item is represented as a concept, or a configuration of concepts, and amounts to a direct word-to-concept link. So except for polysemy and (absolute) synonymy, there is no logical difference between a lexical item and a concept. Therefore, words that are nearly synonymous have to be linked each to their own slightly different concepts. The problem comes in trying to represent these slightly different concepts and the relationships between them. Hirst (1995) shows that one ends up with an awkward proliferation of language-dependent concepts, contrary to the interlingual function of the ontology. And this assumes we can even build a representative taxonomy from a set of near-synonyms to begin with.

Moreover, the denotation of a word is taken to embody the necessary and sufficient conditions for defining the word. While this has been convenient for text analysis and lexical choice, since a denotation can be used as an applicability condition of the word, the model is inadequate for representing the nuances of meaning that are conveyed indirectly, which, clearly, are not necessary conditions.

An alternative representation is suggested by the principle behind Gove's (1973) synonym usage notes. Words are grouped into an entry if they have the same essential meaning, i.e., that they "can be defined in the same terms up to a certain point" (p. 25a) and differ only in terms of minor ideas involved in their meanings. We combine this principle with Saussure's paradigmatic view that "each of

²Hovy (1988) suggests others including force and floridity, and Di-Marco et al. (1993) suggest concreteness or vividness. Actually, it seems that the French text is more vivid—if a text on banking can be considered vivid at all—than the English, using words such as *baptisée, éclatant, contagieux, and démunis*.

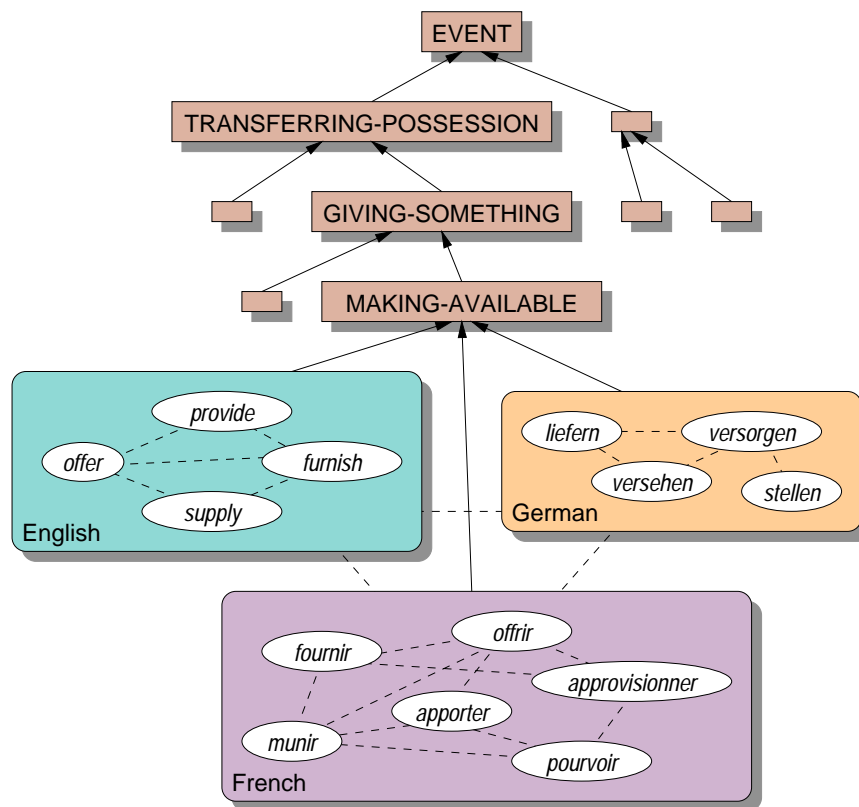


Figure 5: The clustered model of lexical knowledge.

a set of synonyms ... has its particular value only because they stand in contrast with one another” (Saussure, 1983, p. 114) and envision a representation in which the meaning of a word arises out of a combination of its essential denotation (shared with other words) and a set of explicit differences to its near-synonyms.

Thus, I propose a *clustered model of lexical knowledge*, depicted in figure 5. A cluster has two levels of representation: a core concept and peripheral concepts. The *core concept* is a denotation as in the conventional model—a configuration of concepts (that are defined in the ontology) that functions as a necessary applicability condition (for choice)—but it is shared by the near-synonyms in the cluster. In the figure, the ontological concepts are shown as rectangles; in this case all three clusters denote the concept of MAKING-AVAILABLE. All of the *peripheral concepts* that the words may differ in denoting, suggesting, or emphasizing are also represented as configurations of concepts, but they are explicitly distinguished from the core concept as indirect meanings that can be conveyed or not depending on the context. In the figure, the differences between words (in a single language) are shown as dashed lines; not all words need be differentiated. Stylistic, attitudinal, and collocational factors are also encoded in the cluster.

Each language has its own set of clusters. Corresponding clusters (across languages) need not have the same peripheral concepts since languages may differentiate their synonyms in entirely different terms. Differences across languages are represented, for convenience, by dashed lines between clusters, though these would not be used in pure interlingual MT. Essentially, a cluster is a language-specific *formal usage note*, an idea originated by DiMarco et al. (1993) that Edmonds (forthcoming) is formalizing.

4 Interlingual representation

Crucially, an interlingual representation should not be tied to any particular linguistic structure, whether lexical or syntactic.

Assuming that one has constructed an ontology or domain model (of language-neutral concepts), an interlingual representation of a situation is, for us, an instantiation of part of the domain knowledge. Both Stede (1996) and Elhadad et al. (1997) have developed such formalisms for representing the input to natural language generation applications (the former to multilingual generation), but they are applicable to interlingual MT as well. The formalisms allow their applications to paraphrase the same input in many ways including realizing information

at different syntactic ranks and covering/incorporating the input in different ways. For them, generation is a matter of satisfying two types of constraints: (1) covering the whole input structure with a set of word denotations (thereby choosing the words), and (2) building a well-formed syntactic structure out of the words. But while their systems can provide many options to choose from, they lack the complementary ability to actually choose which is the most appropriate.

Now, finding the most appropriate translation of a word involves a tradeoff between many possibly conflicting desires to express certain nuances in certain ways, to establish the right style, to observe collocational preferences, and to satisfy syntactic constraints. This suggests that lexical choice is not a matter of satisfying constraints (i.e., of using the necessary applicability conditions of a word), but rather of attempting to meet a large set of *preferences*. Thus, a distinction must be made between knowledge that should be treated as preferences as opposed to constraints in the interlingual representation. In the generation stage of MT, one attempts to choose the near-synonym from a cluster (activated because of the constraints) whose peripheral concepts best meet the most preferences.

Turning to the analysis stage of MT, since many nuances are expressed indirectly and are influenced by the context, one cannot know for sure whether they have been expressed unless one performs a very thorough analysis. Indeed, it might not be possible for even a thorough analysis to decide whether a nuance was expressed, or how indirectly it was expressed, given the context-dependent nature of word meaning. Thus, on the basis of the knowledge of what words can express, stored in the clusters, the analysis stage would output an interlingual representation that includes *possibilities* of what was expressed. The possibilities then become preferences during generation.

5 Examples

Figures 6–9 give examples of interlingual representations for four segments of the text that involve some of the words discussed in section 2. Since my focus is on word meanings, I will not give complete representations of the expressions. Also note that while I use specific ontological concepts in these descriptions, this in no way implies that I claim these are the right concepts to represent—in fact, some are quite crude. A good ontology is crucial to MT, and I assume that such an ontology will in due course be constructed.

I have used attribute-value structures, but any equivalent formalism would do. Square brackets enclose recursive structures of instantiations of ontological concepts. Names of instances are in lowercase; concepts are capitalized; relations between instances are in uppercase; and cross-reference is indicated by a digit in a square. A whole interlingual representation is surrounded by brace brackets and consists of exactly one

specification of the situation and any number of possibilities, attitudes, and stylistic preferences. The ‘situation’ encodes the information one might find in a traditional interlingual representation—the definite portion of meaning to be expressed. A ‘possibility’ takes as a value a four-part structure of (1) frequency (never, sometimes, or always), which represents the degree of possibility; (2) strength (weak, medium, or strong), which represents how strongly the nuance is conveyed; (3) type (emphasis, suggestion, implication, or denotation), which represents how the nuance is conveyed; and (4) an instance of a concept. The ‘style’ and ‘attitude’ attributes should be self-explanatory. As for content, some of the meanings were discussed in section 2, and the rest are derived from the aforementioned dictionaries. Comments (labelled with ‘%’) are included to indicate which words gave rise to which possibilities.

6 Conclusion

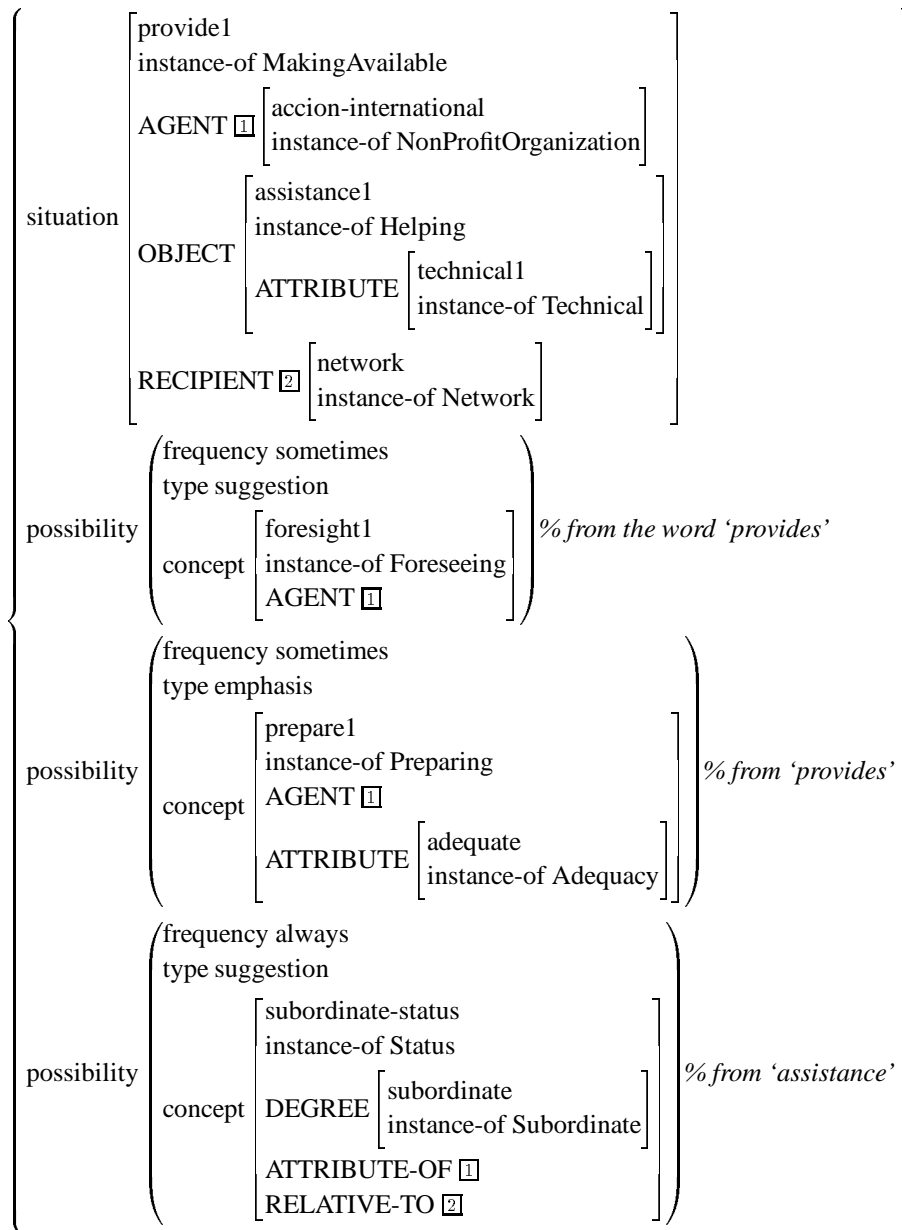
This paper has motivated the need to represent possibilities (or preferences) in addition to necessary components (or constraints) in the interlingual representation of a situation. Possibilities are required because words can convey a myriad of sometimes indirect nuances of meaning depending on the context. Some examples of how one could represent possibilities were given.

Acknowledgements

For comments and advice, I thank Graeme Hirst. This work is financially supported in part by the Natural Sciences and Engineering Research Council of Canada.

References

- René Bailly. 1970. *Dictionnaire des synonymes de la langue française*. Paris Larousse.
- Ronald E. Batchelor and Malcolm H. Offord. 1993. *Using French Synonyms*. Cambridge University Press.
- Henri Bénac. 1956. *Dictionnaire des synonymes*. Paris Hachette.
- Chrysanne DiMarco, Graeme Hirst, and Manfred Stede. 1993. The semantic and stylistic differentiation of synonyms and near-synonyms. In *AAAI Spring Symposium on Building Lexicons for Machine Translation*, pages 114–121, Stanford, CA, March.
- Bonnie J. Dorr. 1994. Machine translation divergences: A formal description and proposed solution. *Computational Linguistics*, 20(4):597–634.
- Philip Edmonds. forthcoming. *Semantic Representations of Near-Synonyms for Automatic Lexical Choice*. Ph.D. thesis, Department of Computer Science, University of Toronto.
- Michael Elhadad, Kathleen McKeown, and Jacques Robin. 1997. Floating constraints in lexical choice. *Computational Linguistics*, 2(23):195–240.
- Philip B. Gove, editor. 1973. *Webster's New Dictionary of Synonyms*. G. & C. Merriam Co.



“ACCION International ... provides technical assistance to a network ...”
 “ACCION International ... fournit une assistance technique à un réseau ...”

Figure 6: Interlingual representation of the ‘equivalent’ sentences shown above. Includes four possibilities of what is expressed.

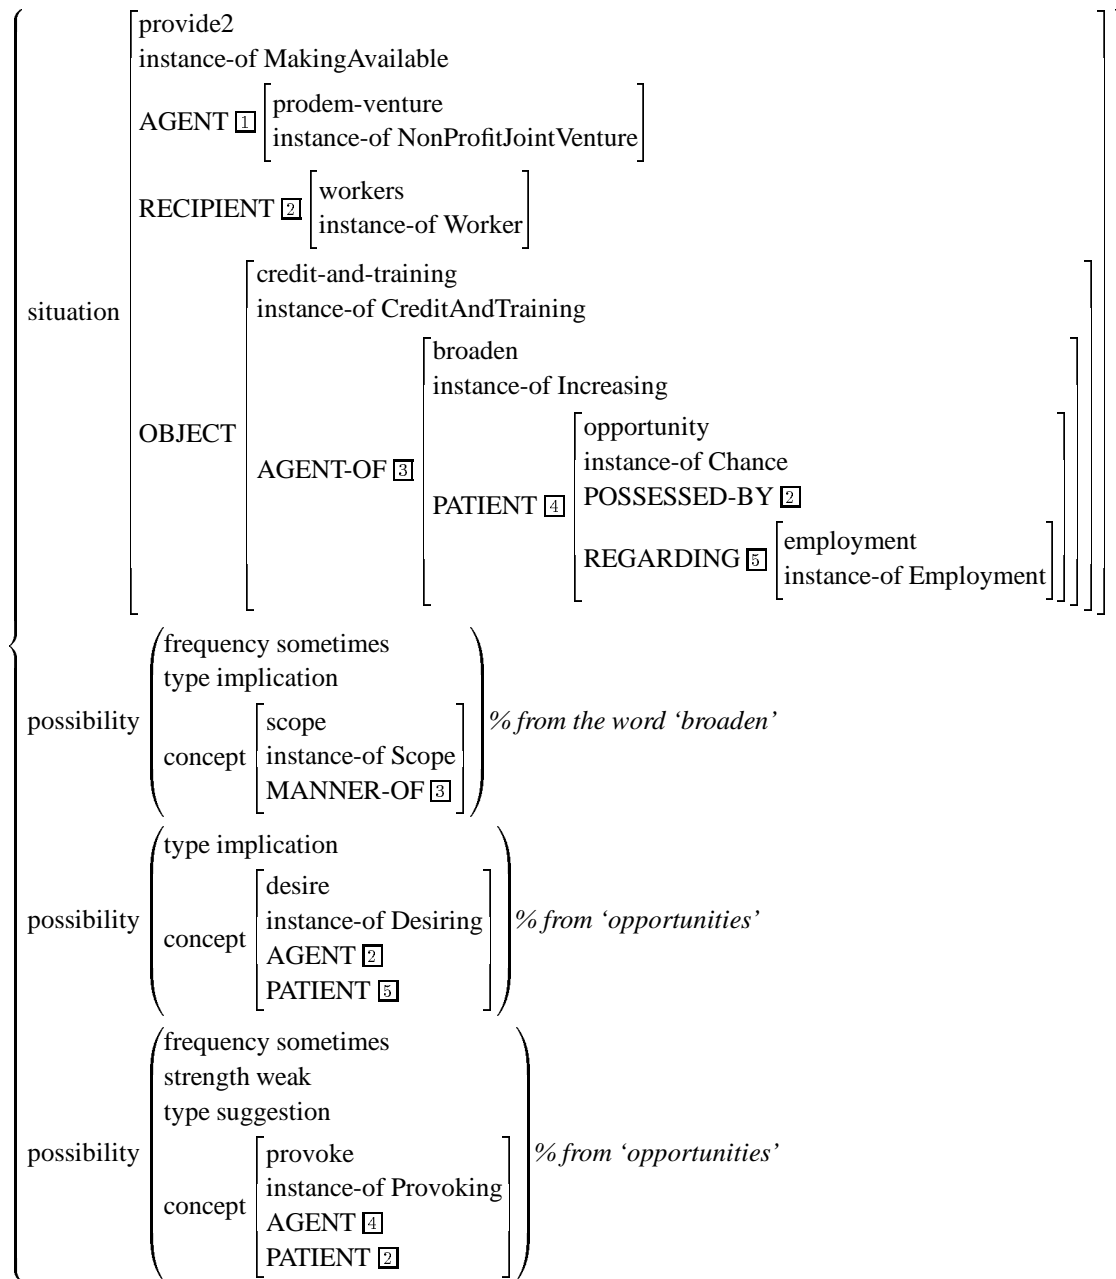
S. I. Hayakawa, editor. 1994. *Choose the Right Word: A Contemporary Guide to Selecting the Precise Word for Every Situation*. HarperCollins Publishers, New York.

Graeme Hirst. 1995. Near-synonymy and the structure of lexical knowledge. In *AAAI Symposium on Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity*, pages 51–56, Stan-

ford, CA, March.

Eduard Hovy. 1988. *Generating Natural Language Under Pragmatic Constraints*. Lawrence Erlbaum Associates.

Megumi Kameyama, Ryo Ochitani, Stanley Peters, and Hidetoshi Sirai. 1991. Resolving translation mismatches with information flow. In *Proceedings of the*



“PRODEM ... provided credit and training to broaden employment opportunities ...”
 “PRODEM ... d’offrir ... des possibilités de crédit et de formation pour élargir leurs perspectives d’emploi”

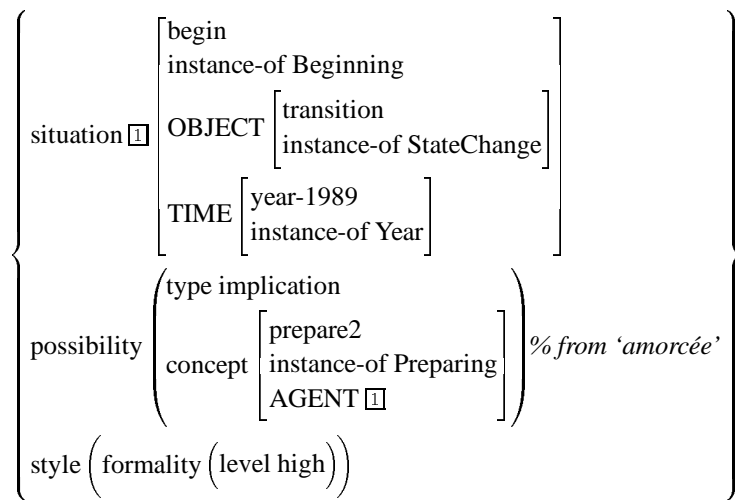
Figure 7: Another interlingual representation with possibilities of what is expressed.

29th Annual Meeting of the Association for Computational Linguistics, pages 193–200.

Ferdinand de Saussure. 1983. *Course in General Linguistics*. G. Duckworth, London. Translation by Roy Harris of *Cours de linguistique générale*, 1916.


Frank Smadja, Kathleen McKeown, and Vasileios Hatzivassiloglou. 1996. Translating collocations for bilingual lexicons: A statistical approach. *Computational Linguistics*, 22(1):1–38.

Manfred Stede. 1996. Lexical paraphrases in multilingual sentence generation. *Machine Translation*, 11:75–107.



“The transition ... began in 1989.”
 “La transition, amorcée en 1989 ...”

Figure 8: Interlingual representation with a stylistic preference (for high formality).



“the very poor self-employed”
 “travailleurs indépendents les plus d’émunis”

Figure 9: Interlingual representation with an expressed attitude.
