TRANSLATIONS AS A SEMANTIC KNOWLEDGE SOURCE

Helge Dyvik
University of Bergen (Norway)

Abstract

Theories of meaning are sometimes used to throw light on the phenomenon of translation. We argue that light can fruitfully be thrown in the opposite direction: we can use translations to get a handle on meaning. More specifically, we will motivate and present a method for the automatic extraction of wordnet-type information from translational data, and review some results. The basic insight behind the method is that much information about the semantic relations among the words in a language resides in the way in which the sets of their possible translations into some other language overlap. Therefore, if we take the translational relation between two languages as a theoretical primitive, languages can serve as each other's "semantic mirrors".

Keywords: translation, parallel corpora, wordnets, lexical semantics, thesaurus derivation.

1. Introduction

Students of translation sometimes ask what a study of meanings may teach us about translation. In this paper I want to reverse the question and ask what a study of translation may teach us about meanings.

From the perspective of the descriptive linguist, or the developer of language resources, this question seems at least as reasonable as the first one. After all, meanings appear to be far more elusive phenomena than translations: we generally feel that we know more or less what translations are, while answers tend to get much vaguer when we are asked what meanings are, or how we should distinguish them. The latter questions require theory-bound reflection, while translation is a practical task. Translations come about when translators, usually with no theoretical concern in mind, evaluate the interpretational possibilities of linguistic expressions in specific contexts, within texts with specific purposes, and then try to recreate the same interpretational possibilities in a target text serving a comparable purpose in another language. This is a normal and common kind of linguistic activity in multilingual societies – an activity which provides an empirical basis for talking about a *translational relation* between languages. Given its basis in the ubiquitous activity of practical translation, the translational relation emerges as epistemologically prior to more abstract and theory-bound notions such as 'meaning', 'synonymy' and 'inference'. What this suggests is taking the translational relation between languages as a theoretical primitive – a concept not to be defined in terms of other concepts, but assumed to be extractable from translational data by interpretive methods – and then investigate to what extent other concepts can be defined in its terms. By this move, we may hope to give semantic

description more of an intersubjective basis. Besides, semantics becomes an essentially multilingual concern.

Two questions may spring to mind at this point. The first one is: Is translation really possible? and the second one: Even if it is, how can it tell us more about the semantics of each language involved than the monolingual approach? The answers, in my opinion, are: No, in a certain sense translation is impossible, and yes, precisely because perfect translation is impossible, actual translations can tell us a lot about semantics. Translation is impossible because meanings and interpretations are not like soft and pliant substances extractable from one expression in one language and mouldable without loss or modification into another expression in another language. Languages, on the contrary, are discrete structures, and meanings are entwined in the structures themselves. Therefore, during translation, things crack and snap, things disappear, and things are added, and there is hardly ever a unique correct solution to a translational task. Instead, actual translations provide a host of alternative approximations to the unattainable ideal, and this is a potential source of information: semantic insights may emerge from the way the sets of alternatives are structured. Semantic studies always depend on paraphrases, or alternative ways of saying the same thing; translations provide such alternatives from a theoretically untainted source.

There is an increasing interest in exploring the potential of translations to provide semantic insights, see e.g. Resnik and Yarowsky 1997, Ide 1999a, Ide 1999b, Diab and Resnik 2002, Ide et al. 2002, Tufis and Ion 2003, Tufis et al. 2003, Tufis et al. 2004, Priss and Old 2005. In Tufis et al. 2004 the authors combine the use of parallel corpora and aligned wordnets from different languages in order to achieve improved word sense disambiguation. In the following we will also look at the relationship between parallel corpora and wordnets, but in a different way: Rather than presupposing them as independent resources – still a rather rare luxury – we will consider to what extent one could reasonably expect to derive wordnets and similar semantic resources from parallel corpora consisting of originals aligned with their translations, based on the general ideas just sketched.

2. Semantic fields and translation

The traditional notion of a ‘semantic field’ stands for a conceptual continuum which is carved up in a certain way by a subset of the vocabulary in a given language, but which may be carved up in different ways in other languages or in different historical stages of the same language. The concept goes back to structuralist studies of lexical semantics by Jost Trier and others. The meanings of words belonging to the same semantic field are supposed to be to some extent interdependent, so that, for instance, the meaning change of a word over time has to be seen in connection with the meaning development of the words around it in the semantic field. A classical structuralist approach to the description of word meanings within a field is the use of *componential analysis*, expressed by assigning *semantic features* to the words, capturing their interrelations. This is closely related to the modern work on ontologies, in which concepts may be structured in lattices defined by feature inheritance, as in the simple example in Figure 1.

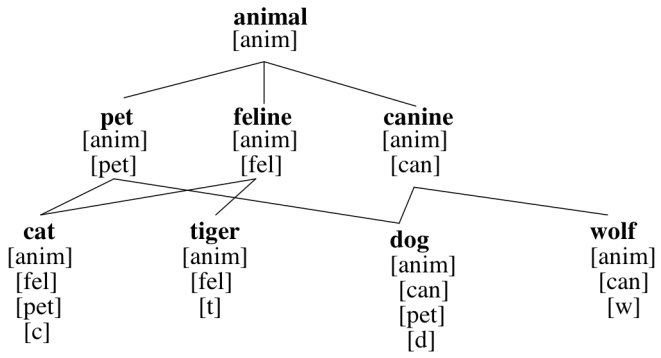

Figure 1. A simple semantic field.

In Figure 1 the most general concept is *animal*, whose single intrinsic feature [anim] is inherited by all the other concepts. One level down *pet* is a common hyperonym of *cat* and *dog*, which inherit *pet*'s intrinsic feature [pet] in addition to its inherited feature [anim], while features inherited from *feline* and *canine* similarly distinguish cats and tigers from dogs and wolves.

The point to note here is that the lattice structure can be read off from the inclusion and overlap relations among the resulting feature sets: the mother/daughter relation in the lattice is a subset/superset relation, while all nodes with intersecting feature sets are dominated by a node carrying the intersection. We should also note in passing that the intuitive semantic content of each feature is unimportant for the characterization of the lattice structure; only the distinctness and distribution of the features matter.

A difference between ontologies and semantic fields is that work on ontologies typically intends to capture constant, language-independent conceptual structures, while work on semantic fields typically intends to bring out the variability and language-specificity of the sets of terms and their interrelations: different languages may carve up the same field in different ways. Without going into the philosophical question of what the 'sameness' of semantic fields across different languages consists in, we may at least observe that the corresponding sets of terms in two languages are connected by a relation of translation. The differences between the ways in which different languages carve up the 'same' field is then reflected in the fact that this translational relation is not one-to-one; consider the classical example in Figure 2.

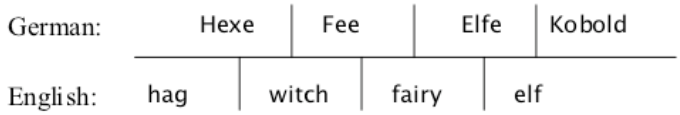

Figure 2. Different partitionings of the 'same' semantic field

In Figure 2, German *Hexe* corresponds translationally both to *hag* (an old repulsive woman, with no presupposition of magical powers) and *witch* (a woman of any age endowed with magical powers), etc. This does not imply that *Hexe* is ambiguous, only that its denotation spans the denotations of two words in the other language. Ambiguity might be involved, but that would have to be independently

established; the existence of more than one translation is not enough. We may also observe how the non-transitive translational connections may tie together semantically distant words in the same semantic field: *hag* and *elf* have little to do with each other semantically, but there is a way, documented by the translational mirror image in German, of getting from the one to the other by small steps from one word to a semantically close word.

Analyses of semantic fields by means of features have also been used in a translational context; one example can be found in an article by one of the pioneers of translation theory, Eugene A. Nida (1958). Here the perspective is that of the translator faced with heavily culture-specific semantic fields of which he has scant knowledge. Hence the question is the traditional one about what a study of meanings may teach us about translation, rather than the reverse. The task is to find translational correspondences between a variety of terms for 'shaman' in two Mayan languages. The method, called 'Componential Plotting', was to make a table with the terms along one axis, and all the different functions of a shaman – healing sick, casting spells, etc. – along the other (Nida 1958:15). Then informants were asked what they would call a person performing each function, and the correspondences between terms and functions were plotted in the table. This is a nice example of an empirical semantic investigation, applied mono-lingually, leading to the assignment of semantic features (denoting shaman functions) to a set of words across two languages. A network of translational correspondences between terms in each language could then be established on the basis of shared features.

In our context this example illustrates the connection between feature sharing and translational correspondence, but we want to use that connection in order to go in the other direction – from translational correspondences to semantic features – since we are taking the translational relation as a primitive. After all, the normal case is that translation is performed without any previous, theoretically sophisticated analysis like Componential Plotting, but rather based on the existing cross-cultural competence of translators. Treating the output of translators as data is therefore not much different from treating any kind of output from language users as data for linguistic studies.

3. Translationally based representations

In reversing the direction of inference from the 'shaman' case to the case of deriving semantic features from translational data, the basic question becomes: What minimal set of semantic features, and which distribution of them, would motivate this given network of translational relations? In order to answer this question we need not consider the possible semantic interpretation of the features themselves; they are simply translationally derived formal devices whose distribution among a set of words is the only thing that matters. We may consider a simple example.

As we saw earlier, the German noun *Hexe* can be found translated into English as *hag* and *witch*; cf. Figure 3.

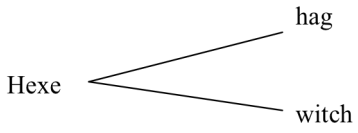

Figure 3. A simple translational correspondence

These alternative translations are obviously related to different 'aspects', or related subsenses, of the meaning of *Hexe*. The two English words indicate one way, undoubtedly among many, of dividing up the semantic potentiality of *Hexe*. In fact, we could conceive of lexical subsenses as corresponding to ordered pairs like <Hexe, hag> and <Hexe, witch> – or to sets rather than pairs, if we take several languages into account simultaneously. A translational approach to semantics sees such sets of translationally corresponding items across languages as *the primitives of semantic descriptions*. (This idea is related to the idea behind Martin Kay's 'triangulation' approach to translation.) Pairs like <Hexe, hag> can then be treated as a kind of semantic features, written [Hexe|hag] and assignable to lexical items, both to the items they were derived from (as in Figure 4), and to others, which may inherit them – a point to which we will return.

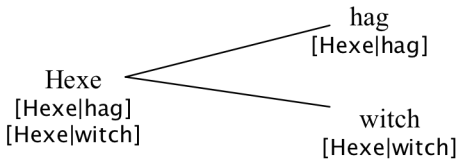

Figure 4. Assignment of translationally derived features

Intuitively, the features encode subsenses that the lexical items share with each other. In this way the features become classificatory devices, grouping lexical items together according to shared semantic properties.

4. The Semantic Mirrors method

4.1. Assumptions

Given a word-aligned parallel corpus, we may extract the set of alternative translations for each lemma in the corpus. The result is an intricate network of translational correspondences uniting the vocabularies of the two languages. This network allows us to treat each language as the 'semantic mirror' of the other, based on the ideas sketched above, in conjunction with the following assumptions:

(1) Semantically closely related words tend to have strongly overlapping sets of translations.
(2) Words with wide meanings tend to have a higher number of translations than words with narrow meanings.
(3) If a word *a* is a hyponym of a word *b* (such as *tasty* of *good*, for example), then the possible translations of *a* will probably be a subset of the possible translations of *b*.
(4) Contrastive ambiguity, i.e., ambiguity between two unrelated senses of a word, such as the two senses of the English noun *band* ('orchestra' and 'piece of tape'), tends to be a historically accidental and idiosyncratic property of individual words. Hence we don't expect to find instances of the same contrastive ambiguity replicated by other words in the language or by words in other languages. (More precisely, we should talk about ambiguous *phonological/graphic* words here, since such ambiguity is normally analysed as homonymy and hence as involving two lemmas.)

(5) Words with unrelated meanings will not share translations into another language, except in cases where the shared (graphic/phonological) word is contrastively ambiguous between the two unrelated meanings. By assumption (4) there should then be at most one such shared word.

4.2. Isolating word senses

The first step in applying the method is to use assumptions (4) and (5) to identify the set of alternative, mutually unrelated senses of each word.

We will refer to the set of translations in L2 of a word *w* in L1 as 'the first *t*-image' of *w*. Taking the first *t*-images back in L1 of all the members of *w*'s first *t*-image gives us a set of intersecting sets of words in L1; this will be referred to as *w*'s 'inverse *t*-image'. We may then make a third translational move, finding the first *t*-images in L2 of all the members of the union of *w*'s inverse *t*-image; this gives us a set of intersecting sets of words in L2, which we will call *w*'s 'second *t*-image'.

We may exemplify sense individuation by means of *t*-images with a corpus example taken from The English-Norwegian Parallel Corpus (ENPC), a corpus which comprises approximately 2.6 million words, originals and translations included. The corpus contains fiction as well as non-fiction and English originals translated into Norwegian as well as the other way around (Johansson et al. 1996). The example[1] is based on manual word alignment. Figure 5 shows the contrastively ambiguous Norwegian noun *rett* (which can mean, i.a., 'dish' and 'court of law') with its first and inverse *t*-images. Obviously, *rett* vill be a member of all the sets in its inverse *t*-image, but this is not shown in the figure, to keep it reasonably simple. However, it should be kept in mind, since it means that all the sets that are shown as intersecting in the inverse *t*-image actually contain *rett* as well in their intersections. This is crucial, given assumption (5) above, because it means that all the intersections contain at least two members, which in turn means that the sets are assumed to contain semantically related words. With only one word in the intersection, the chances are that this word may be contrastively ambiguous between the senses represented by each set. For example, this is the case with the first *t*-images of *law* and *food*, which only contain *rett* itself in the intersection (not shown in the figure). Given that the two *t*-images are not indirectly connected by means of intersections with other sets, this leads to the conclusion that *rett* is contrastively ambiguous between a *law* sense and a *food* sense.

To put it more carefully: The sets in the inverse *t*-image are divided in groups based on intersections containing words in addition to *rett* itself, and each such group is assumed to correspond to a distinct sense of *rett*. Mapping these groups back on the first *t*-image gives a partitioning of it into *sense partitions*, indicated by horizontal lines in Figure 5. Thus we individuate four senses *rett1, rett2, rett3* and *rett4*, each associated with its own first *t*-image.

While the result looks plausible as far as the separation of the *food* and *law* senses is concerned, it also illustrates the inevitable limitations of using a finite corpus: *course* really belongs in the *food* partition, but constitutes its own spurious sense here because the corpus happens not to contain any translations of *course*, apart from *rett*, shared with any of the other food-related English words.

---

[1] The example is taken from Lyse (2003).

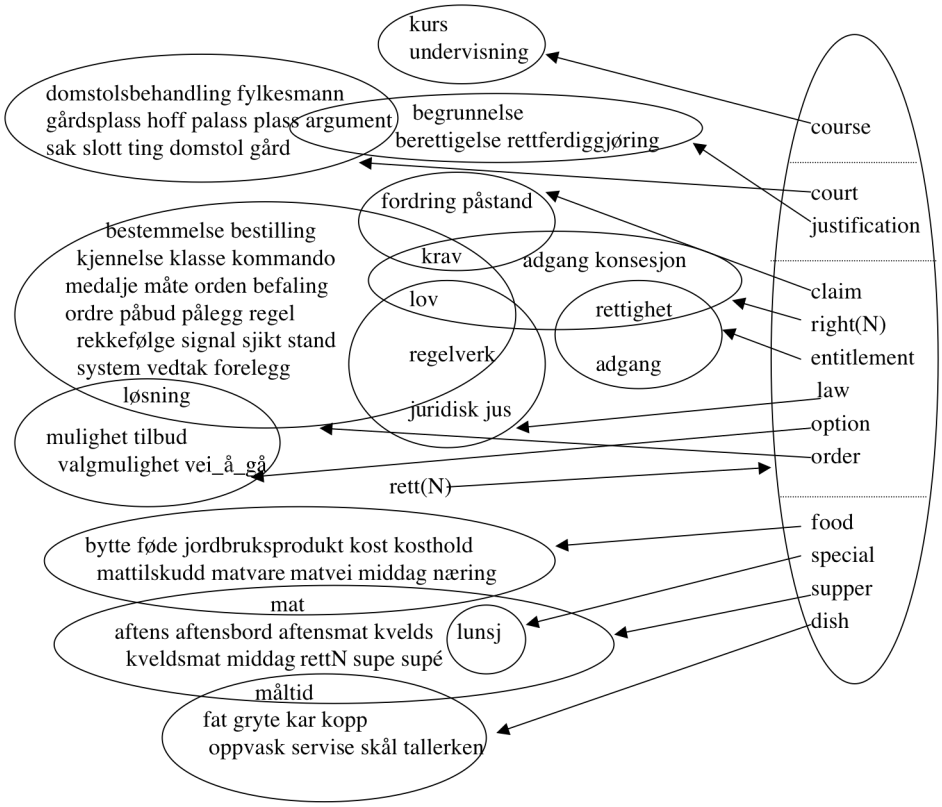

Figure 5. The first (on the right) and inverse (on the left) *t*-images of the noun *rett*

4.3. Semantic fields and feature assignment

Once senses are individuated in the manner described in both languages, they can be grouped into *semantic fields*. In our translational approach, the semantic fields are isolated on the basis of overlapping *t*-images: two senses belong to the same semantic field if they have intersecting first *t*-images (after sense individuation one member in the intersection is sufficient), or if there is a sequence of such intersecting *t*-images joining them.

We treat translational correspondence as a symmetric relation (disregarding the direction of translation), and as a consequence we get paired semantic fields in the two languages involved. Each field f1 and f2 in such a pair imposes a subset structure on the other, since all the *t*-images of the members of f1 will be subsets of f2, and *vice versa*. By assumptions (1-3) above, rich information about the semantic relations among the senses can be derived from this subset structure.

Taking the food-related sense of *rett* (*rett4*) as a starting point, we can collect its semantic field by finding all other senses with directly or indirectly intersecting *t*-images. This is shown on the left in Figure 6. The corresponding field in English is shown on the right. Furthermore, the subset structures imposed by the *t*-images are also indicated.

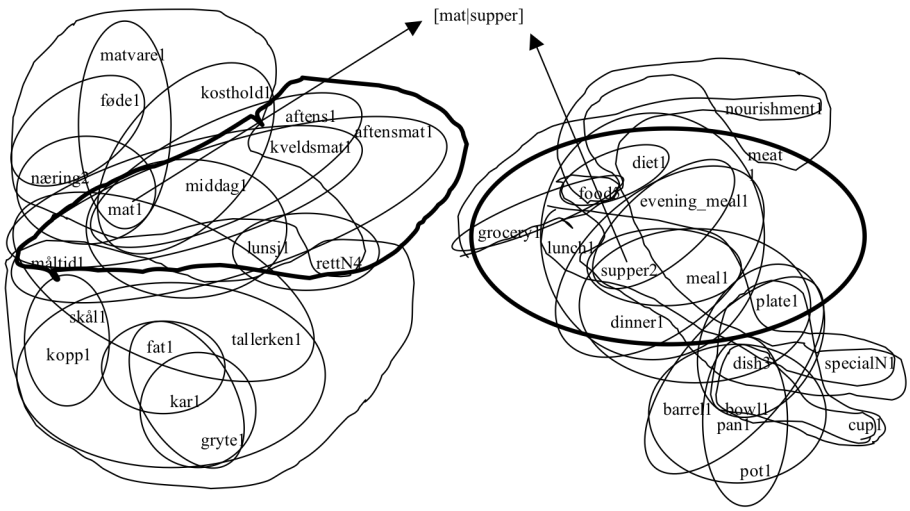

Figure 6. Paired semantic fields from Norwegian and English

The fact that a sense is a member of many subsets, i.e., of many *t*-images, indicates that it has many translational partners in the other field. By assumption (2) such senses are expected to have wide meanings as compared to other senses in the field. As expected, senses such as *food5* and *mat1* ('food') in Figure 6 constitute such peaks in the subset structures (although *supper2* happens to outrank *food5* in the English field, being a member of an even higher number of subsets). Furthermore, the fact that two senses are co-members of many subsets means that they share many translations and hence ought to be closely related semantically.

In this way the subset structures contain rich information about the semantic relations among the senses, and the next step is to encode this information in feature sets associated with the senses. The procedure[2] is to start from the 'peaks', i.e., from the pair of senses that are both translationally related and members of the highest number of subsets – *mat1* and *supper2* in the example. A feature is constructed from these two senses, as also illustrated in Figure 6. The feature is assigned to the two senses *mat1* and *supper2*, and is then inherited by 'lower' senses, i.e., by all senses ranked lower than *mat1* within the first *t*-image of *supper2*, and by all senses ranked lower than *supper2* within the first *t*-image of *mat1*. The *t*-images in question are marked by bold lines in Figure 6. Then the procedure moves on iteratively to the next highest peaks – *middag1* ('dinner') and *food5* in the example – constructing the feature [middag1|food5] and assigning it according to the same principles. The final result is feature sets assigned to all the senses in the two fields. By hypothesis, feature set inclusion now expresses a hyperonymy/hyponymy relation, e.g. as in the two senses *food5* : *lunch1*:

food5	**lunch1**
[mat1|supper2]	[mat1|supper2]
[middag1|food5]	[middag1|food5]
	[lunsj1|meal1]
	[lunch1]

[2] The procedure is described in more detail in Dyvik 1998:80ff.

The full set of senses in a field is thus partially ordered by set inclusion. We can construct an upper semilattice from this set, allowing us to compare the distances between all the senses in the field. An upper semilattice is a partially ordered set in which each pair of elements has a least upper bound. Applied to our case this means that for each pair of feature sets, either one set includes the other or there is a third feature set consisting of the intersection of the two sets. By adding elements with such intersections whenever they don't exist already, we construct an upper semilattice from a semantic field. Intuitively, the added elements are 'virtual hyperonyms' of the intersecting elements – potential senses that happen not to be lexicalized in the language (or at least not to occur in the corpus). We label the added elements as indexed X'es. Thus, given two intersecting sets such as the sets for *busy2* and *alive2* in Figure 7, we construct the node *X1* carrying the intersection of the feature sets:

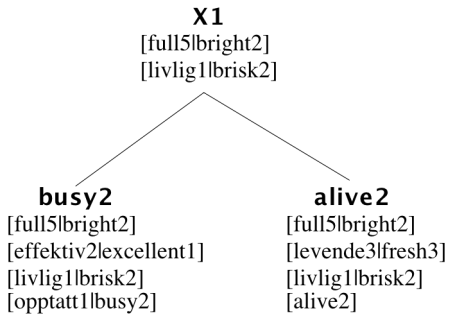

Figure 7. Adding X-nodes to construct a semilattice.

Figure 8 shows a small part of a semilattice for adjectives, based on manual word alignment of the ENPC, in which the adjective *brilliant* unites senses related to cleverness and radiance.

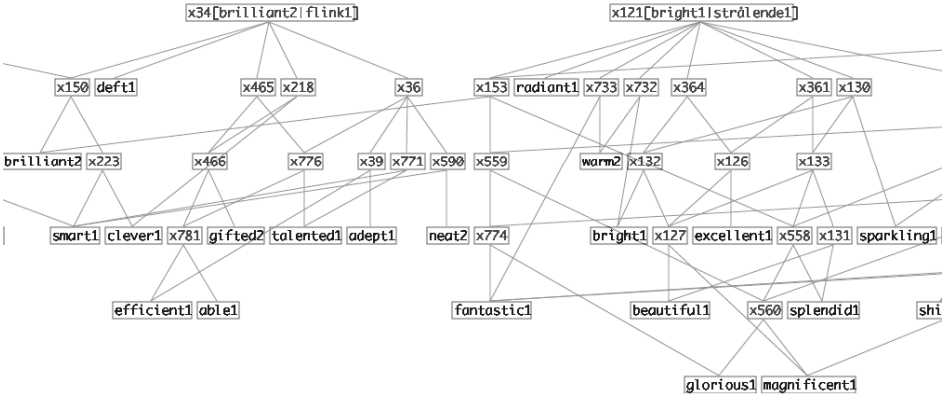

Figure 8. Part of a semilattice for adjectives

4.4. Deriving thesaurus entries

The feature lattices contain some of the information represented in thesaurus entries, and we may derive rudimentary thesaurus-like entries from them. Derivation of a thesaurus entry for a sense *s* involves collecting senses that are sufficiently related to *s* from the semilattice, and sort them into hyperonyms, hyponyms and synonyms of *s*. Basically, a related sense of *s* is a sense sharing features (and hence translations) with *s*. A hyperonym of *s* is then a sense *h* from which *s* has inherited a feature, provided that the number of senses having inherited this feature exceeds a certain threshold (called *SynsetLimit*); the latter provision ensures that hyperonyms have sufficiently wide meanings. Hyponyms of *s* are, conversely, senses which have inherited an inherent feature of *s*, with the same provision about the number of heirs. Synonyms and 'related words' are also identified on the basis of certain kinds of feature sharing.

Furthermore the sense *s* can be divided into mutually related subsenses. Each feature assigned to *s* potentially represents a distinct subsense; whether two features *f1* and *f2* should be considered as belonging to the same subsense or not, can be determined on the basis of the sets of senses to which *f1* and *f2* are assigned. If the intersection of these sets of senses exceeds a certain theshold (called *OverlapThreshold*), the features are not considered as representing distinct subsenses.

For example, with a certain setting of the thresholds the following entry is derived for one sense of the adjective *brilliant*:

brilliant
Hyperonyms: bright‹1›.
Subsense (i)
(**Translation:** skarp, flink.)
Synonyms: able, adept, clever, deft‹1›, efficient‹1›, fierce, gifted‹2›, neat‹2›, smart‹1›, talented‹1›.
Related words: burning‹1›, harsh‹1›, hot‹1›, keen‹1›, piercing‹2›, sharp‹1›, shrill‹1›, spiny‹1›, stark‹1›, steep‹1›, stinging‹1›.
Subsense (ii)
(**Translation:** fantastisk, strålende.)
Synonyms: amazing‹1›, enormous‹1›, exceptional, extraordinary‹1›, fantastic, glorious‹1›, magnificent‹1›, marvellous, remarkable‹1›, spectacular‹1›, splendid‹1›, startling‹1›, surprising‹1›, unusual‹1›.

Increasing the *OverlapThreshold* leads to a splitting up of Subsense (i) in two subsenses, separating the 'hot' and 'sharp' aspect of the sense from the 'clever' and 'efficient' aspect:

brilliant
Hyperonyms: bright‹1›.
Subsense (i)
(**Translation:** flink.)
Synonyms: able, adept, clever, deft‹1›, efficient‹1›, gifted‹2›, neat‹2›, smart‹1›, talented‹1›.
Subsense (ii)
(**Translation:** fantastisk.)
Synonyms: amazing‹1›, enormous‹1›, exceptional, extraordinary‹1›, fantastic, glorious‹1›, magnificent‹1›, marvellous, remarkable‹1›, spectacular‹1›, splendid‹1›, startling‹1›, surprising‹1›, unusual‹1›.
Subsense (iii)
(**Translation:** skarp.)
Synonyms: fierce.
Related words: burning‹1›, clever, harsh‹1›, hot‹1›, keen‹1›, piercing‹2›, sharp‹1›, shrill‹1›, smart‹1›, spiny‹1›, stark‹1›, steep‹1›, stinging‹1›.

5. Empirical findings

The Semantic Mirrors method has been explored in a project involving automatic word alignment[3] of the ENPC and comparison of results from manually aligned and automatically aligned data, comparison of the output of our method with existing resources such as the Princeton Wordnet and Merriam-Webster's Thesaurus[4], and testing of the method as a basis for word sense disambiguation[5] (presently only with preliminary results). We may briefly summarize some of our findings so far as follows:

- The method is vulnerable to the increased noise introduced by automatic word alignment: precision and recall in the thesaurus output from automatically aligned data as compared with the output from manually aligned data seems to be lower than the precision and recall of the automatic word alignment itself as compared with manual word alignment.
- It is hard to find a suitable gold standard for the evaluation of the thesaurus output. When using Merriam-Webster's Thesaurus or Princeton Wordnet as gold standards for the sets of semantically related words associated with the thesaurus entries, precision and recall is low, but not very much worse than the results obtained when we compare the established resources Merriam-Webster and Princeton Wordnet with each other.
- There is a distinct difference between different parts of speech: the method gives better results for adjectives than for nouns and verbs, and abstract nouns give better results than concrete nouns. With concrete nouns very few hyperonym/hyponym-relations are discovered, probably because translational relations between hyperonyms and hyponyms are more rare with concrete nouns than with, e.g. adjectives: translating *dog* with a word meaning 'animal' doesn't happen as often as translating *tasty* with a word meaning 'good', for example. Besides, adjectives, typically denoting single properties, tend to form tighter groups of closely related members than nouns, which typically denote clusters of properties; this may explain why adjectives tend to have more alternative translations than nouns.

In sum, if high-quality translational data can be provided, the method clearly seems to provide some useful results.

6. References

Diab, Mona & Philip Resnik 2002. An Unsupervised Method for Word Sense Tagging using Parallel Corpora. *40th Anniversary Meeting of the Association for Computational Linguistics (ACL-02)*, Philadelphia, July, 2002.

Dyvik, Helge 1998a. A translational basis for semantics. In Stig Johansson and Signe Oksefjell (eds.): *Corpora and Crosslinguistic Research: Theory, Method, and Case Studies*. Amsterdam: Rodopi 1998, 51-86.

Dyvik, Helge 1998b. Translations as semantic mirrors. In *Proceedings of Workshop W13: Multilinguality in the lexicon II*, pp. 24.44, Brighton, UK. The 13th biennial European Conference on Artificial Intelligence ECAI 98.

[3] The algorithm for automatic word alignment was developed by Sindre Sørensen.

[4] This evaluation is carried out by Martha Thunes.

[5] Word sense disambiguation is explored by Gunn Inger Lyse.

Dyvik, Helge 2002. Translations as semantic mirrors: from parallel corpus to wordnet. *Language and Computers*, 1 April 2004, vol. 49, iss. 1, 311-326(16) Rodopi.

Ide, Nancy 1999. Word sense disambiguation using cross-lingual information. In *Proceedings of ACH-ALLC '99 International Humanities Computing Conference*, Charlottesville, Virginia. http://jefferson.village.virginia.edu/ach-allc.99/proceedings

Ide, Nancy 1999. Parallel translations as sense discriminators. In *SIGLEX99: Standardizing Lexical Resources, ACL99 Workshop*, College Park, Maryland, 52-61.

Ide, Nancy, Tomas Erjavec & Dan Tufis 2002. Sense Discrimination with Parallel Corpora. In *Proceedings of ACL'02 Workshop on Word Sense Disambiguation: Recent Successes and Future Directions*, Philadelphia, 54-60.

Johansson, Stig, Jarle Ebeling, and Knut Hofland 1996. Coding and aligning the English-Norwegian Parallel Corpus. In K. Aijmer, B. Altenberg, and M. Johansson (eds.) 1996. *Languages in contrast. Papers from a symposium on text-based cross-linguistic studies in Lund, 4-5 March 1994*, 73-85. Lund: Lund University Press, 87-112.

Lyse, Gunn Inger 2003. *Fra speilmetoden til automatisk ekstrahering av et betydningstagget korpus for WSD-formål.* Masters thesis, University of Bergen.

Nida, Eugene A. 1958. Analysis of Meaning and Dictionary Making. In *Language Structure and Translation. Essays by Eugene A. Nida*. Selected and Introduced by Anwar S. Dil, Stanford University Press 1975, 1-23.

Priss, Uta and John Old 2005. Conceptual Exploration of Semantic Mirrors. In Ganter, Godin (eds.) *Formal Concept Analysis: Third International Conference, ICFCA 2005*, Springer Verlag.

Tufis, Dan and Radu Ion 2003. Word sense clustering based on translation equivalence in parallel texts; a case study in Romanian. In *Proceedings of the International Conference on Speech and Dialog – SPED*, 13-26, Bucharest.

Tufis, Dan, Radu Ion, Nancy Ide 2004. Fine-Grained Word Sense Disambiguation Based on Parallel Corpora, Word Alignment, Word Clustering and Aligned Wordnets. In *Proceedings of the 20th International Conference on Computational Linguistics, COLING2004*, 1312-1318, Geneva.

Resnik, Philip Stuart & David Yarowsky 1997. A perspective on word sense disambiguation methods and their evaluation, position paper presented at the ACL SIGLEX Workshop on *Tagging Text with Lexical Semantics: Why, What, and How?*, held April 4-5, 1997 in Washington, D.C., USA in conjunction with ANLP-97.

HELGE DYVIK is professor of general linguistics at the University of Bergen, Norway. He received his dr.philos. degree in 1981 based on a study of syntactic theory and linguistic methodology. His present research interests comprise semantics and translation, machine translation, and the development of computational grammars. He has previously published studies on Old Norse, Old English and Vietnamese, and on the interpretation of runic inscriptions. He is presently head of the programme committee for the research programme on language technology under The Nordic Council of Ministers.