Translations as a Semantic Knowledge Source

Helge Dyvik

Abstract

The first part of this paper discusses the relationship between meaning theory, meaning descriptions and translation, arguing that translations are a plausible source of information about semantics. The second part describes a method, based on these assumptions, for deriving semantic descriptions in the form of wordnet- and thesaurus-like entries from translational text corpora.

1. Meaning and Translation

1.1 Meaning Theory and Meaning Descriptions

There are many alternative answers to the question what meanings are. They are variously viewed as simply a useful theoretical construct in the account of certain kinds of behaviour or knowledge, as concepts or mental representations in the brains of human beings, as residing in the norm-governed use of language, as existing in the external world as relations between types of situations, or as functions from informational states to new informational states in a discourse. There are also other positions.

An interesting observation, though, is that it is possible for linguists to disagree completely about what and where meanings are – or even whether they are at all – and still to a large extent agree on the way they should be described: Meaning description seems to be possible across conflicting philosophical persuasions. This might suggest that we should try to derive some hints about the nature of meanings from considering our consensus about meaning descriptions.

Admittedly, 'consensus' may be putting it too strongly, since meaning descriptions after all come in many shapes. One recently popular shape within the field of lexical semantics is the *wordnet*, with the Princeton WordNet for English as a prototype (Fellbaum 1998). The wordnet is a semantically classified lexical database, in which different senses of words are separated and then grouped into sets of similar senses – socalled *synsets*, or sets of approximate synonyms. Furthermore, synsets are interrelated by various types of semantic relations: hypero-/hyponymy (i.e., super- and subconcepts), opposites, part-whole relations, etc. As an example we may consider the online Princeton WordNet entry for the hyperonyms of the noun 'society':

4 senses of *society*

Sense 1
society -- (an extended social group having a distinctive cultural and economic organization)
    => social group -- (people sharing some social relation)
        => group, grouping -- (any number of entities (members) considered as a unit)

Sense 2
club, society, guild, gild, lodge, order -- (a formal association of people with similar interests; "he joined a golf club"; "they formed a small lunch society"; "men from the fraternal order will staff the soup kitchen today")

=> association -- (a formal organization of people or groups of people; "he joined the Modern Language Association")
=> organization, organisation -- (a group of people who work together)
=> social group -- (people sharing some social relation)
=> group, grouping -- (any number of entities (members) considered as a unit)

Sense 3
company, companionship, fellowship, society -- (the state of being with someone; "he missed their company"; "he enjoyed the society of his friends")
=> friendship, friendly relationship -- (the state of being friends)
=> relationship -- (a state involving mutual dealings between people or parties or countries)
=> state -- (the way something is with respect to its main attributes; "the current state of knowledge"; "his state of health"; "in a weak financial state")

Sense 4
society, high society, beau monde, smart set, bon ton -- (the fashionable elite)
=> elite -- (a group or class of persons enjoying superior intellectual or social or economic status)
=> upper class, upper crust -- (the class occupying the highest position in the social hierarchy)
=> class, social class, socio-economic class -- (people having the same social or economic status; "the working class"; "an emerging professional class")
=> people -- ((plural) any group of human beings (men or women or children) collectively; "old people"; "there were at least 200 people in the audience")
=> group, grouping -- (any number of entities (members) considered as a unit)

In this entry meanings are described in different ways. We have a paraphrase (e.g., «a formal association of people with similar interests»); we have examples of use (e.g., «they formed a small lunch society»); and first and foremost in a wordnet: we have other words in the same language, with specified semantic relations to the word in question – approximate synonyms (e.g., «club, society, guild, gild, lodge, order») and a hierarchy of hyperonyms (e.g., «association => organisation => social group => group»).

A related kind of meaning description is the thesaurus, which also characterises senses by means of related words from the same language. As an example we may consider the entry for the adjective *conspicuous* in the Merriam-Webster Collegiate Thesaurus (http://www.m-w.com/home.htm):

Entry Word: **conspicuous**
Function: *adjective*
Text: **1**
Synonyms CLEAR 5, apparent, distinct, evident, manifest, obvious, open-and-shut, openhanded, patent, plain
2
Synonyms NOTICEABLE, arresting, arrestive, marked, outstanding, pointed, prominent, remarkable, salient, striking
Related Word celebrated, eminent, illustrious; showy
Contrasted Words common, everyday, ordinary; covert, secret; concealed, hidden
Antonyms inconspicuous

Again we have a separation into senses, and a characterisation of each of them by means of synonyms, related words, antonyms and contrasted words.

In anticipation of the second part of this paper we may compare the Merriam-Webster entry with the entry below, which has been derived automatically from translational data extracted from a Norwegian-English parallel corpus according to the method to be described:

conspicuous
Sense 1
Hyperonyms: large, hard, great.
Subsense (i) (Translation: synlig, tydelig.)
Synonyms: apparent, evident, pervasive, substantial, visible.
Related words: clear, definite, distinct, distinctive, obvious, plain, unmistakable, vivid.
Subsense (ii) (Translation: fremtredende, kraftig, sterk, stor.)
Synonyms: primary.
Subsense (iii) (Translation: oppsiktsvekkende.)
Synonyms: amazing.
Related words: spectacular, startling, surprising, unusual.
Sense 2
(Translation: avstikkende.)

1.2 Semantic Representations

Wordnets, thesauri and dictionaries in general are meaning descriptions designed for practical use, and mostly confined to the description of individual words. The meaning descriptions in theoretical studies on semantics, on the other hand, usually have a different character. In that context meaning descriptions have traditionally been referred to as «semantic representations». (The term may be somewhat out of vogue today, but we will stick to it in the present context.) A semantic representation of a sentence is a structure taken to display the meaning of that sentence. In some approaches (e.g., cognitive semantics) semantic representations have pictorial qualities, whereas in other approaches they are expressions in formal languages. (1)-(6) below is a selection of conceivable types of semantic representations, constructed for the sake of illustration, of the sentence «The boy is repairing the bicycle in the garage»:

(1)

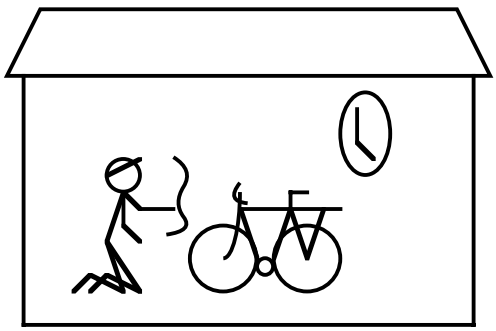

(2)

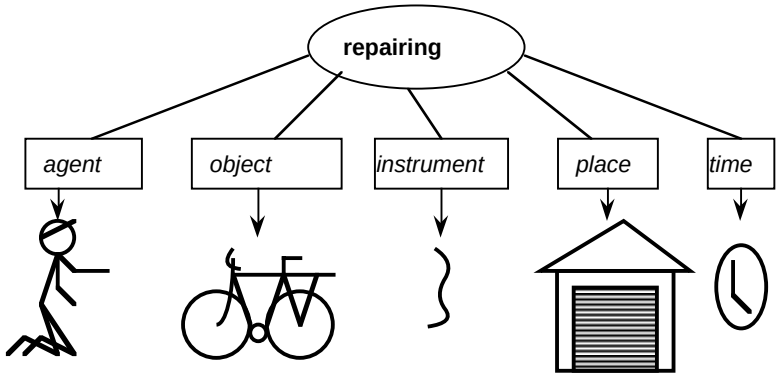

(3)

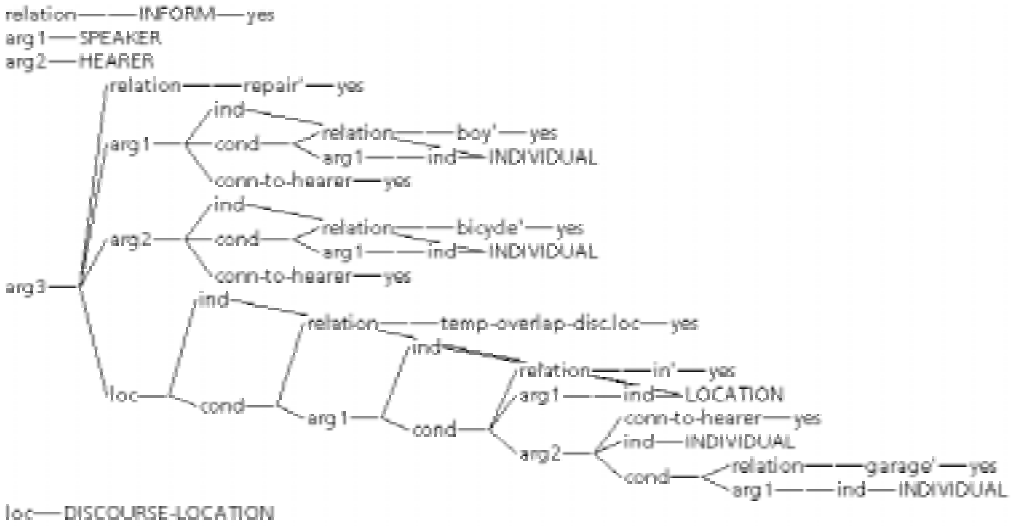

(4) ∃!x ∃!y ∃!z ∃e [event(e) & boy'(e,x) & bicycle'(e,y) & garage'(e,z) & repair' (e,x,y) & in' (e, z)]

(5) The boy is repairing the bicycle in the garage

(1) is intended as a simplistic picture of the described situation. Thus it seems to be an iconic representation, and we might feel that it describes the meaning of the sentence by relating it to its denotation in the world in virtue of an unproblematic relation of similarity. But this feeling is deceptive. Depending on how the picture came about, it is either totally uninformative about the meaning of the sentence, or it is not really iconic, but rather like a piece of language itself. Under the first alternative the artist has drawn the picture unconstrained by anything but his own artistic talents. In that case the picture says far too much, and far too little. It says far too much because there are many arbitrary details in the picture that are not described in the sentence. Therefore there is no way of extracting the exact message in the sentence from the picture alone. Furthermore, the picture says far too little – for one thing, it is

hard to make out that the vague activity depicted should be seen as an instance of *repairing*. Creative drawings may be pedagogically valuable, but they are theoretically useless.

On the other hand, the picture might have come about in a different way. Its iconic appearance might be incidental and misleading; it might actually have been put together by strict rules from a well-defined repertoire of primitives, conventionally drawn as kneeling boy-shapes, bicycle-shapes, and a string-shape stipulated to represent the act of repairing. To someone knowing the formal system defining the set of possible drawings, and its interpretation, the drawing would then be eminently informative about the meaning of the sentence – but perhaps (2) would then be a more practical layout. In (2) the unifying activity *repairing* has been singled out and arbitrarily symbolised by the English verb, while the roles of the other participants are explicitly indicated. A figure with a cap is stipulated to represent the class of boys, etc. Now the representation starts to look like something that could be used in the study of semantic relations between expressions, for instance inference. But then, of course, the shapes of all the elements in the drawing are entirely arbitrary, since the meaning of the elements resides in their interpretation and not in their shapes. Hence the elements could just as well have been replaced with linguistic symbols – as in (3). (3) is a feature structure (in the form of a directed graph) with hardly any iconic properties. Relations are represented by relational primitives like *repair'*, and classes of objects are represented by indeterminates constrained to be arguments of relations like *boy'*, etc. There would be no point in replacing these symbols with little pictures. The same is true of the more common kind of semantic representation: The formula in first-order logic (4). In other words, a semantic representation is theoretically interesting only to the extent that it is like a piece of language itself, rather than like a picture.

This is usually taken to suggest that assigning a semantic representation to a sentence is like an act of translation: The sentence has just been translated into a form which we find more useful for some purpose. We have now transferred the question of the meaning of words and sentences to that of the meaning of the semantic representations – but the question is still there. Under this perspective the meaning of a semantic representation is of the same kind as the meaning of the sentence which it represents. If that is the case, the sentence itself could count as its own semantic representation (5). The only reason for choosing one of the others would then be that they serve our purpose better (for instance the characterisation of inference).

1.3 Models and Relations between Expressions

The account above is more or less the picture we find in model theoretic semantics after Richard Montague. Within model theoretic semantics the semantic representations get their meaning through being interpreted in a semantic model, usually a set-theoretic structure with individuals and possibly other elements:

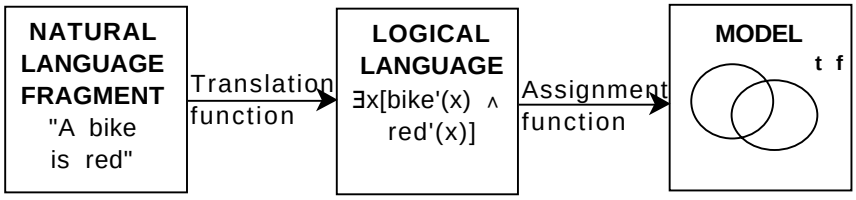

Figure 1 The components of model-theoretic semantics

With the semantic models, something potentially picture-like has turned up again. The models have set-theoretic and sometimes geometrical structures which intuitively seem to correspond to certain properties of the actual world talked about. However, such a possible relation of similarity is usually not a primary concern among formal semanticists. The models are constructed, not discovered through inspection of the world, and they are constructed not so much in order to look like the world as in order to *capture relations of inference among expressions in the language*. The immediate concern of the models is the linguistic expressions and their inference relations: that is the domain against which their adequacy is tested.

Accordingly, interpreting a semantic representation in a model is an indirect way of interpreting it in a domain of natural language expressions entering into relations of inference and synonymy with each other. Of the two – the model and the natural language – it is the natural language which is empirically given, while the model is constructed in order to capture some of its properties. Hence, it must basically be the properties of the natural language itself which give meaning to the semantic representations. Through denoting entities in a model the representations classify linguistic expressions, which can be seen as their ultimate denotations, rather than entities in the world talked about. In this way semantic representations stand for equivalence classes of linguistic signs. The element *repair'* in the feature structure or logical formula does not ultimately stand for a relation or a set in the described world, but for a set of linguistic signs which are equivalent with respect to inferential, denotational or other semantic properties.

Under this assumption it is no longer motivated to consider the semantic representation as a *translation* of the sentence, since the sentence and the representation denote different things. The English verb *repair* denotes a kind of activity in the actual world, while the expression *repair'* in the representations denotes the English verb *repair* and its synonyms. Then we can no longer claim that a sentence is its own semantic representation, since a sentence denotes something in the world, and not itself, while a semantic representation denotes a class of semantically equivalent expressions, of which the sentence is a member.

Saying that semantic representations are assigned to linguistic expressions, which they in turn denote, may sound circular. However, a semantic representation does not in general denote only the single linguistic expression to which it is assigned, but a class of such expressions, held together by some kind of equivalence relation. It is this relation which gives the representation empirical content. In monolingual semantic studies the relation is usually one of *paraphrase*, while relations of *inference* may be characterised in terms of formal properties of the representations.

Relations of paraphrase and inference can be elicited from informants. Still, one might feel that these rather sophisticated and laboratory-like relations are not entirely satisfactory as the only raw and basic empirical data for linguistic semantics,

removed as they are from situated language use. They are more sufficient for logicians than for linguists, one might suggest.

1.4 Translation

Where else, then, do we find information about the semantic relations between expressions, as untainted as possible by philosophical or theoretical considerations about meaning? One answer is of course: *In translations* from one natural language into another. 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', 'paraphrase' 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 terms of it. By this move, semantics also becomes essentially multilingual. In a multilingual setting – for instance, in interlingua-based machine translation – the element *repair'* in the semantic representations then does not only denote a set of synonyms in one particular language, but a set of signs across languages, held together by a relation of intertranslatability.

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 we will suggest 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. Perfect 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 inextricably entwined in the structures themselves; the message is enmeshed in the medium. Therefore, during translation, things crack and snap, things disappear, and things are added. The target language is like a Procrustean bed for the source language. The source language is never quite comfortable – it tosses and turns and finds that it can rest its shoulder in this position and its leg in that position, but no position will rest it all at once. It is in this set of alternative positions that we see what it is made of: The anatomy of meaning emerges in the translational tension between languages. Actual translations provide a host of alternative approximations to the unattainable ideal, and semantic insights emerge from the structure of the sets of alternatives.

1.5 Translationally based semantic representations

In a translational approach, the semantic representations must be designed so as to capture the intricate network of translational approximations. Simple elements like *repair'* will seldom do; since the translational relations across languages typically are not one-to-one, the elements must be decomposed into more primitive elements capable of capturing the translational alternatives. We may consider a simple example.

The Norwegian adjective *lekker* can be found translated into English as, e.g., *delicious* and *pretty*:

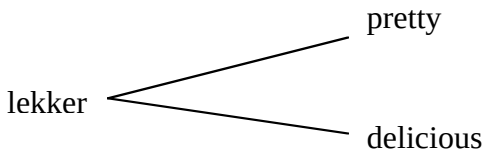

These alternative translational partners obviously are related to different aspects, or related subsenses, of *lekker*. The two English words indicate one way, undoubtedly among many, of dividing up the semantic potentiality of *lekker*. In fact, we could conceive of lexical subsenses as corresponding to pairs like [lekker|delicious] and [lekker|pretty] (or to sets rather than pairs, if we take several languages into account simultaneously). A translational approach to semantics would see such sets of translationally corresponding items across languages as *the primitives of semantic descriptions*. (This idea is similar to the idea behind Martin Kay's triangulation approach to translation.) Pairs like [lekker|delicious] can then be treated as a kind of semantic feature assignable to lexical items – both to the items it was derived from, and to others, which may inherit it – encoding subsenses that they may share with each other:

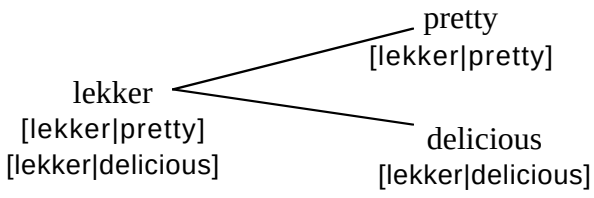

In this way the features become classificatory devices, grouping lexical items together based on shared semantic properties: they are semantic representations, according to the discussion in in 1.3 above.

By such a method (to be further developed below), the granularity of the semantic representations will be dictated by the choice of translational partners. If only very closely related languages are taken into account – such as Norwegian and Danish, for instance – there will probably be a relatively high incidence of one-to-one correspondences between lexical items and grammatical categories across the two languages, and a corresponding relatively low degree of granularity in the semantic representations. The translational relation between very closely related languages is therefore probably not very informative about semantics. At the other end of the scale, the incongruent lexical and grammatical distinctions in unrelated, typologically

distant languages will necessitate a granularity going far beyond the lexicalisations and grammaticalisations in each language. In the theoretical limiting case – taking the full set of possible natural languages into account as translational partners – the degree of granularity in the required semantic representations will be maximal, limited only by the granularity of the accumulatively superimposed lexical and grammatical distinctions in the full set of languages.

Hence we expect there to be more semantic information to be gleaned from the translational relations between distantly related languages than from the relations between closely related languages. The worse the Procrustean bed, the more we should learn.

The idea that semantic insights can be derived from the translational relation has also been explored by others; see for instance Diab (2002), Ide (1999a, 1999b), Ide & al. (2002), Resnik & Yarowsky (1997).

1.6 Literalness and Word-Alignment

However, there is a difficult step to be taken in order to get from translational data to insights about meanings. Meanings are properties of linguistic signs as types; they belong to *la langue*. The immediately accessible translational relation, on the other hand, is a relation between situated texts: it concerns tokens or occurrences and hence belongs to *la parole*. In a translational corpus many perfectly good translations reflect specific properties of the particular texts and their circumstances more than they reflect linguistic properties of the expressions involved. For instance, some aspects of the interpretation of a text may be expressed linguistically in the original while left for the reader to infer from the context or from general world knowledge in the translation, or *vice versa*. In order to conclude anything about translational relations between elements of the languages on a de-contextualised type level, based on the translational relations between texts in a parallel corpus, we have to make at least the assumption that relations on the type level can be teased out by applying interpretive methods to translational corpus data. In other words, we have to assume that it is meaningful to talk about a relation of 'literal translation', and that a relation of literal translation can be extracted through an interpretive sifting of data from actual translations, excluding situation-specific, non-literal correspondences from consideration. The Procrustean considerations above make it clear that 'literal translation' does not mean 'perfect translation', but rather something like 'translation predicted as possible by the properties of the expression translated seen in isolation'.

We will not go more deeply into the involved question of literalness in the present context, except for noting that it is hardly possible to draw a borderline once and for all between instances of literal and instances of non-literal translation. Rather, the distinction must be drawn relative to the delimitation of the languages (general languages, sublanguages etc.) in which we assume that the texts are composed. However, the distinction is inevitable – it has to be drawn somewhere – if translation and meaning are to be at all accessible to systematic study.

A concept of literalness is a hidden assumption behind word-alignment. Word-alignment of a parallel corpus consists in connecting each word occurrence which can be said to have an identifiable translational partner in the other text, with its partner. Attempts to align parallel corpora on the level of words can be seen as attempts to extricate a word-level relation of literal translation from the corpus data.

Word-alignment, whether done manually or automatically, is a precondition for the use of parallel corpora in the service of lexical semantics. It is therefore a starting point for our project *From Parallel Corpus to Wordnet* at the University of Bergen, which tests the approach to lexical semantics sketched above, and which will be described in the remainder of this article.

2. The Semantic Mirrors Method

2.1 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:

i. Semantically closely related words tend to have strongly overlapping sets of translations.
ii. Words with wide meanings tend to have a higher number of translations than words with narrow meanings.
iii. 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*.
iv. Contrastive ambiguity, i.e., ambiguity between two unrelated senses of a word, such as the two senses of the English noun *bank* ('money institution' and 'riverside'), 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.)
v. Words with unrelated meanings will not share translations into another language, except in cases where the shared word is contrastively ambiguous between the two unrelated meanings. By assumption *iv* there should then be at most one such shared word.

2.2 Isolating word senses

The first step in applying the method is to use assumptions *iv* and *v* to identify the set of alternative senses for each word. We may exemplify this with a simple corpus example. The corpus is The English-Norwegian Parallel Corpus (ENPC), 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. ENPC is aligned at sentence level (Johansson & al. 1996), while it is a part of our project to align the ENPC at word level.

We refer to the set of translations of a word *a* as 'the first *t*-image' of *a*. The following example is based on manually extracted translational correspondences from the corpus. The Norwegian noun *tak* is contrastively ambiguous between the meanings 'roof' and 'grip'. Figure 2 shows the first *t*-image of *tak* in the right-hand box, and the first *t*-images of each of those English words again in the left-hand box. We refer to the last-mentioned set of sets as the 'inverse *t*-image' of *tak*.

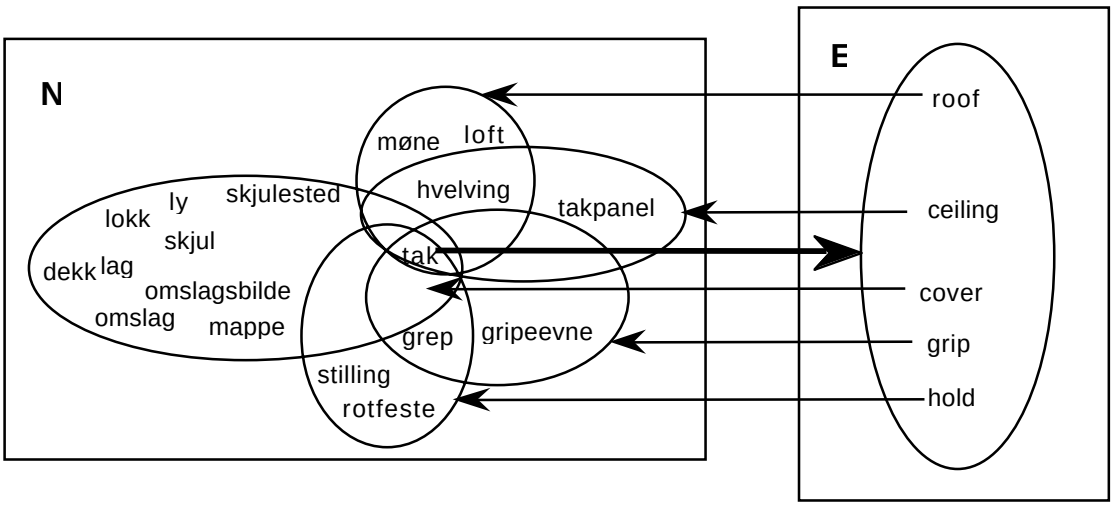

Figure 2 The first and inverse *t*-images of *tak*.

The point worth noticing is that the images of *roof* and *ceiling* overlap in *hvelving* in addition to *tak*, while the images of *grip* and *hold* overlap in *grep* in addition to *tak*. By assumption *v* above, this indicates that *roof* and *ceiling* are semantically related (they share more than one translation), and similarly *grip* and *hold*, while no overlap (apart from *tak*) unites *grip/hold* and *roof/ceiling*. *Grip/hold* and *roof/ceiling* hence seem to represent unrelated meanings, and the conclusion is that *tak* is contrastively ambiguous between those meanings. Perhaps less convincingly the method also separates out *cover* as representing a third meaning.

On this basis the first *t*-image of *tak* can be split into three 'sense partitions':

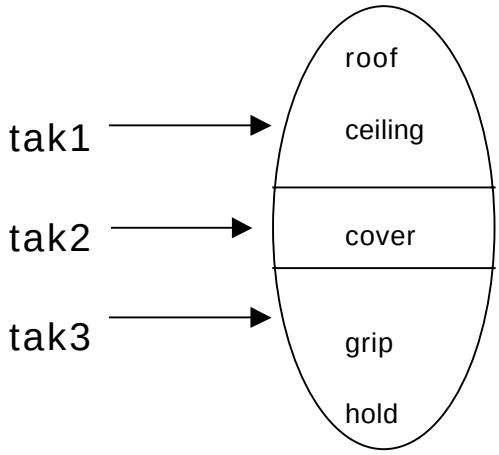

Figure 3 The sense partitions of *tak*'s first *t*-image

Thus the main senses of lemmas are individuated, and each sense is associated with its own *t*-image in the other language.

2.3 Semantic Fields

Once senses are individuated in the manner described, they can be grouped into *semantic fields*. Traditionally, a semantic field is a set of senses that are directly or indirectly related to each other by a relation of semantic closeness. 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 at least one sense in the other language corresponds translationally with both of 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 *i-iii* above, rich information about the semantic relations among the senses can be derived from this subset structure. Figure 4 shows an example of paired fields from the corpus, with the subset structure indicated.[1]

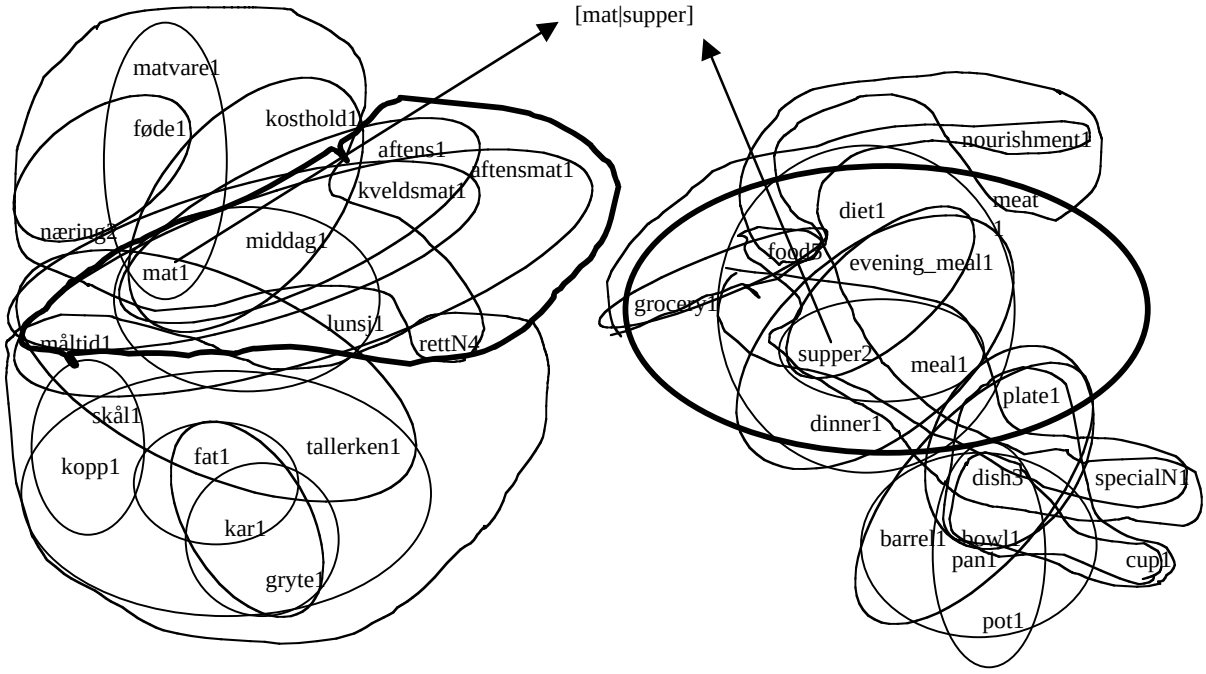

Figure 4 Paired, structured semantic fields from Norwegian and English

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 *ii* 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 4 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.

[1] The translational data from which these fields were derived were extracted and analysed by Gunn Inger Lyse in connection with her Masters thesis (Lyse 2003).

In this way the subset structures contain rich information about the semantic relations among the senses (given that our assumptions are correct), and the next step is to encode this information in semantic representations in the form of feature sets associated with the senses. The procedure 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 4. 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 4. Then the procedure moves on 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. The assignments in the Norwegian field is shown in table 1:

aftens1 [mat1|supper2] [aftens1|evening_meal1] **aftensmat1** [mat1|supper2] [lunsj1|meal1] [kveldsmat1|meal1] [aftensmat1] **fat1** [fat1|dish3] **føde1** [middag1|food5] [føde1|grocery1] **gryte1** [måltid1|dish3] [fat1|dish3] [kar1|dish3] [gryte1|bowl1] **kar1** [kar1|dish3]	**kopp1** [måltid1|dish3] [fat1|dish3] [kar1|dish3] [skål1|bowl1] [gryte1|bowl1] [kopp1|cup1] **kosthold1** [middag1|food5] [kosthold1|diet1] **kveldsmat1** [mat1|supper2] [kveldsmat1|meal1] **lunsj1** [mat1|supper2] [lunsj1|meal1] **mat1** [mat1|supper2] **matvare1** [middag1|food5] [føde1|grocery1] [matvare1] **middag1** [mat1|supper2] [middag1|food5]	**måltid1** [mat1|supper2] [måltid1|dish3] **næring2** [middag1|food5] [næring2|nourishment1] [næring2|meat1] **rettN4** [mat1|supper2] [middag1|food5] [måltid1|dish3] [fat1|dish3] [kar1|dish3] [rettN4|specialN1] **skål1** [måltid1|dish3] [fat1|dish3] [kar1|dish3] [skål1|bowl1] **tallerken1** [måltid1|dish3] [fat1|dish3] [kar1|dish3] [skål1|bowl1] [gryte1|bowl1] [tallerken1|plate1]

Table 1 Feature assignments to the Norwegian senses

The hierarchical relations in the field are now expressed through inclusion and overlap relations among the feature sets. Thus, for instance, wide senses like *mat1* have small feature sets, whereas their assumed hyponyms have supersets of those sets.

In this way the feature sets form a semilattice which can be represented by the graph in figure 5:

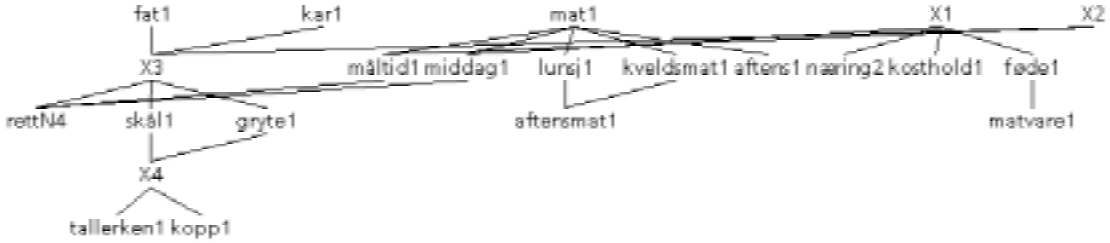

Figure 5 A semilattice formed by the feature sets.

According to the hypothesis, then, senses on dominating nodes are hyperonyms of senses on dominated nodes. Thus, a sense of *mat* 'food' dominates senses of *rett* 'course', 'dish', *middag* 'dinner', *måltid* 'meal', *lunsj* 'lunch', *kveldsmat* 'supper', *aftensmat* 'supper', and *aftens* 'supper', all of which are plausible hyponyms of *mat*. Furthermore, we find *næring* 'nutrition', *kosthold* 'fare', 'food regimen', *føde* 'nourishment' and *matvare* 'food article' clustered together, with *føde* dominating *matvare.* Similarly, *kar* 'vessel' dominates words for more specific vessels like *skål* 'bowl', *gryte* 'pan', *kopp* 'cup', etc. Less convincingly, *lunsj* also dominates *aftensmat.*

2.4 Deriving Thesaurus Entries

The feature lattices contain some of the information represented in thesaurus entries like the Merriam-Webster entry for *conspicuous* cited initially, and we may derive rudimentary thesaurus-like entries from them. Derivation of thesaurus entries involves determining subsenses, hyperonyms, synonyms and hyponyms of each sense on the basis of the information in the semilattices. The semilattices are in some cases extremely complex, showing intricate networks of connections between the word senses. In the transition to a wordnet database or a thesaurus we want to abstract away from much detail in the lattices, and this can obviously be done in more than one way.

We may illustrate the procedure by means of an example: the adjective *sweet*. Figure 6 shows a small sublattice of the large lattice including the sense *sweet1*:

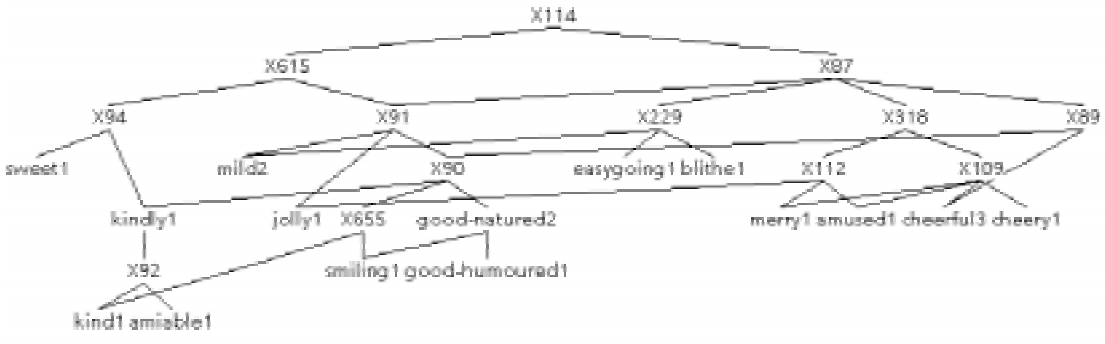

Fig 6 A sublattice containing *sweet1*

Sweet1 is also dominated by several nodes outside this sublattice; size limitations prevent displaying a more complete graph. The node *sweet1* is associated with the following feature set:

[god2|good1]
[fin1|nice2]
[pen1|gentle2]
[vakker1|soft1]
[snill1|pleasant1]
[deilig1|splendid2]
[frisk4|sweet1]
[blid2|sweet1]

The features in a feature set assigned to a sense *s* are in general of two types: inherited features (i-features) and own features (o-features). I-features are inerited from other senses according to the procedure sketched in 2.3 above. An o-feature has been constructed from the sense *s* itself (and one of its translational correspondents), and may be inherited by other senses. In the example, [frisk4|sweet1] and [blid2|sweet1] are the o-features of the sense *sweet1*. Figure 6 above shows just the sublattice dominated by the feature [blid2|sweet1].

Finding hyperonyms, near-synonyms and hyponyms of *sweet1* now first involves considering what other senses in the complete lattice share features with *sweet1*. The features in question have been assigned to the following senses in the complete lattice (we refer to the sets of senses as the *denotations* of the features):

[god2|good1]:
(able1 accurate1 adept1 adequate1 affectionate1 all_right1 amiable1 appropriate3 attractive3 beautiful1 beneficial1 benign2 bright1 burning2 charming1 clean1 clear1 close3 comfortable1 comforting2 competent2 confident2 correct1 cozy2 cute1 decent1 delicious1 delightful1 detailed2 dishy1 easy1 efficient1 elegant2 excellent2 fair1 fancy1 favourable1 fine1 firmA1 first-class2 first-rate1 fit2 fortunate1 fresh2 friendly1 full1 genuine1 good1 handsome1 happy2 healthy1 high2 hot1 joyful1 kind1 kindly1 long2 lovely1 lucky1 magnificent2 marvellous1 neat2 nice2 okay1 peaceful1 perfect2 picturesque1 placid1 pleasant1 pleased1 pleasing1 pleasurable1 plentiful1 plenty1 polite2 positive1 pretty2 proficient1 quite_certain1 real1 reassuring1 respectable2 right2 ripe1 safe1 satisfactory1 satisfying1 secure1 sizeable1 skilled1 smart1 smooth3 soft1 solid1 sound2 spectacular1 steady1 strong1 successful1 suited1 superb1 superior4 sure1 sweet1 talented1 thorough1 tidy1 well2 whole1 wholesome1 wonderful2 worthy2)

[fin1|nice2]:
(attractive3 beautiful1 breathtaking2 charming1 comfortable1 cute1 delicate3 dishy1 easy1 elegant2 enchanting1 excellent2 fancy1 fine1 first-class2 gentle2 glorious3 graceful1 handsome1 impressive1 lovely1 magnificent2 marvellous1 neat2 nice2 okay1 perfect2 picturesque1 pleasurable1 polite2 pretty2 pure1 slight3 smart1 soft1 splendid2 sweet1 thin1 wonderful2)

[pen1|gentle2]:
(attractive3 beautiful1 charming1 clean1 cute1 dishy1 elegant2 enchanting1 fancy1 fine1 first-class2 formal2 gentle2 graceful1 handsome1 lovely1 mild2 neat2 picturesque1 pleasant1 polite2 pretty2 soft1 sweet1 tidy1)

[vakker1|soft1]:

(attractive3 beautiful1 charming1 cute1 delightful1 dishy1 enchanting1 fair1 fancy1 graceful1 handsome1 lovely1 magnificent2 mild2 ornate2 picturesque1 pleasant1 pleasurable1 pretty2 soft1 sweet1)

[snill1|pleasant1]:
(all_right1 amiable1 benign2 friendly1 good-humoured1 good-natured2 jolly1 kind1 kindly1 mild2 pleasant1 pleasing1 polite2 smiling1 sweet1)

[deilig1|splendid2]:
(beautiful1 charming1 cute1 delicious1 delightful1 enchanting1 pleasurable1 splendid2 sweet1)

[frisk4|sweet1]:
(all_right1 brisk4 eager1 fit2 fresh2 healthy1 new1 pert1 sweet1 well2)

[blid2|sweet1]:
(amiable1 amused1 blithe1 cheerful3 cheery1 easygoing1 good-humoured1 good-natured2 jolly1 kind1 kindly1 merry1 mild2 smiling1 sweet1)

We use two parameters to regulate the generation of thesaurus entries from this information: *SynsetLimit* and *OverlapThreshold.*

SynsetLimit:
The most general features, [god2|good1], [fin1|nice2] and [pen1|gentle2], denote a high number of senses each – especially [god2|good1]. This reflects the fact that they are constructed from wide senses such as *god2* and *good1*. As a result, many of the senses carrying those features are not sufficiently close to *sweet1* to be called «related words». Therefore we do not want to consider all the senses sharing such general features as related words or synonyms of each other. The value of the parameter *SynsetLimit* defines the maximal size which the set denoted by a feature can have in order to be included among the synonyms or related words. With *SynsetLimit* = 20, the sets of senses denoted by [god2|good1], [fin1|nice2] and [pen1|gentle2] are not included among the related words of *sweet1* (unless they are denoted by other features as well). On the other hand, *good1*, *nice2* and *gentle2* – the English senses from which the wide features were constructed – are recorded as hyperonyms of *sweet1*.

More precisely, the following principles are the ones presently being explored in the project:

(1) A *hyperonym* of *s* is a sense which has an o-feature with a denotation greater than SynsetLimit which is inherited by *s* (i.e., is an i-feature of *s*).

(2) A *synonym* of *s* is a sense which is not a hyperonym or a hyponym of *s* and furthermore either (a) has an o-feature with a denotation less than or equal to SynsetLimit which is inherited by *s*, or (b) has inherited one of *s*'s o-features *f* iff the denotation of *f* is less than or equal to SynsetLimit, or (c) shares two or more features with *s*.

(3) A *related word* of *s* is a sense which is not a hyperonym, a synonym or a hyponym of *s*, and which is denoted by one of *s*'s i-features whose denotation is less than or equal to SynsetLimit.

(4) A *hyponym* of *s* is a sense which has inherited one of *s*'s o-features *f* iff the denotation of *f* is greater than SynsetLimit.

The intention is that 'synonyms' should be more closely related semantically than 'related words'. The principles ensure that synonymy and 'related word' each are symmetric relations, and that hyperonymy and hyponymy are inverses. The higher the value of SynsetLimit, the 'flatter' the semantic structure: as we increase the SynsetLimit, hyperonymy/hyponymy relations are replaced by synonymy relations, and the sets of synonyms and related words increase in size. It seems recommendable to let the value of SynsetLimit vary with the size of the semantic fields. Thus, in order to have *mat1* 'food' emerge as a hyperonym in the thesaurus entry derived from the small semantic field illustrated in figure 5, SynsetLimit has to be set to 5.

OverlapThreshold:
The value of the parameter *OverlapThreshold* decides the granularity of the division into *subsenses* in the thesaurus entry. This does not concern the division into main *senses* described in 2.2 above (*tak1, tak2, tak3* etc.) – those senses usually end up in different semantic fields and hence in different lattices. Division into *subsenses* is a further subdivision of each sense into related shades of meaning. We assume that there is no final and universal answer to the question of how many related subsenses a word sense has (cf. Kilgarriff 1997). By means of the parameter *OverlapThreshold* we may attune that kind of semantic granularity to our purposes.

Intuitively, the features assigned to *sweet1* represent different 'aspects' of the sense, and the question now is whether those 'aspects' are sufficiently different from each other to be considered different subsenses. Their distinctness can be measured in terms of the degree of overlap among the sets of senses they denote. If the set of features denote strongly overlapping sets of senses, the favoured conclusion is that there is no division into subsenses. On the other hand, the less the denotations of the features overlap, the more a division into subsenses is motivated. The degree of overlap in a set of sets can be measured as a value between 0 and 1, with 0 indicating no overlap and 1 full overlap (full overlap meaning that for each set *s*, every set either includes *s* or is included in *s*). In calculating the overlap degree among feature denotations we disregard the sense *sweet1* itself, since it is necessarily a member of all the feature denotations.

Accordingly, the value of the parameter *OverlapThreshold* is a number between 0 and 1. The parameter is used to group the feature set of a sense into subsenses. A feature belongs to subsense *n* if the degree of overlap between its denotation and the denotation of at least one other feature in subsense *n* is equal to or greater than *OverlapThreshold*. Hence, the lower the OverlapThreshold, the more features will be gouped together in subsenses, while the higher the OverlapThreshold, the more subsenses tend to be distinguished, with one sense per feature as the limiting case. We may consider the consequences for the example *sweet1*.

With SynsetLimit = 20 the three first feature denotations ([god2|good1], [fin1|nice2] and [pen1|gentle2]) are excluded from consideration as sources of synonyms and related words. As for the remaining features, their denotations intersect as shown below (removing the sense *sweet1* itself):

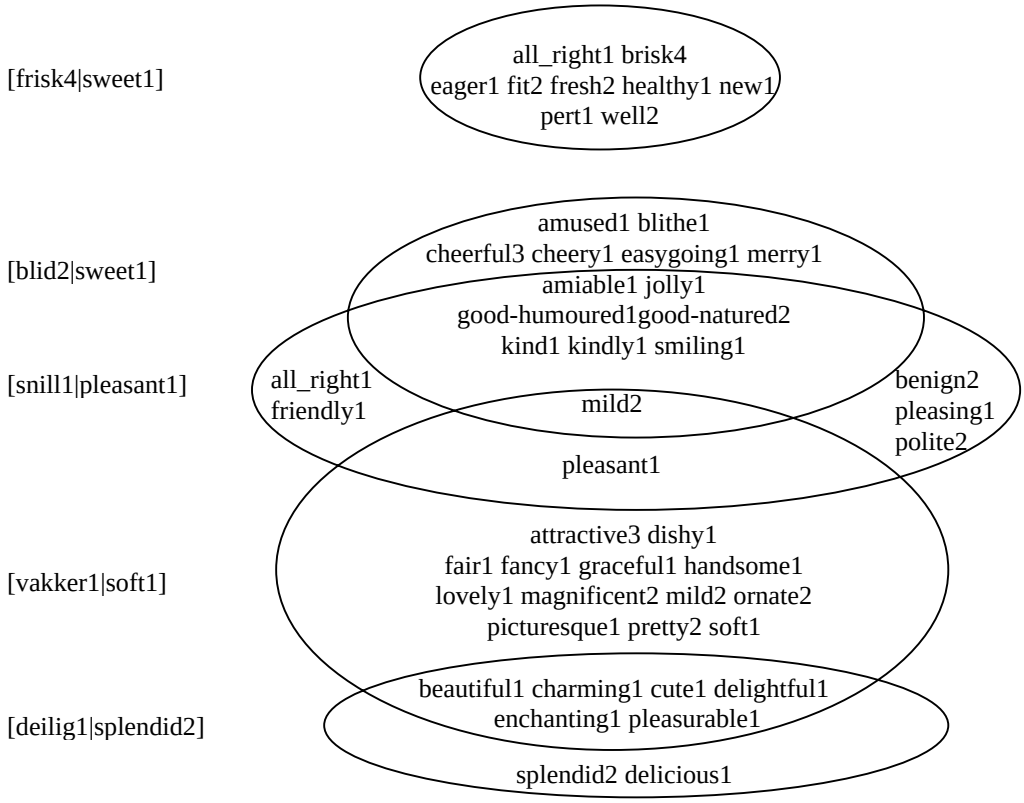

Figure 7 Intersections among the denotations of the features of *sweet1*.

The denotation of the feature [frisk4|sweet1] does not intersect with any other denotation. As a consequence this feature will necessarily be isolated as one subsense in the thesaurus entry. The other features intersect to varying degrees. The smallest overlap – the set {*mild2, pleasant1*} – unites the denotations of [snill1|pleasant1] and [vakker1|soft1].

If we set the value of *SynsetLimit* to 20 and the value of *OverlapThreshold* to 0.05, the Mirrors program consequently generates the following entry for *sweet*:

Version I:

OverlapThreshold: 0.05
SynsetLimit: 20

sweet
Hyperonyms: gentle, nice, good.
Subsense (i) (Translation: blid, deilig, fin, god, pen, snill, vakker.)
Synonyms:
amiable, amused, beautiful, blithe, charming, cheerful, cheery, cute, delightful, easygoing, enchanting, good-humoured, good-natured, jolly, kind, kindly, merry, mild, pleasant, pleasurable, smiling, soft, splendid.
Related words:
all_right, attractive, benign, delicious, dishy, fair, fancy, friendly, graceful, handsome, lovely, magnificent, ornate, picturesque, pleasing, polite, pretty.
Subsense (ii) (Translation: frisk.)
Synonyms: all_right, brisk, eager, fit, fresh, healthy, new, pert, well.

Given the low OverlapThreshold, the intersection {*mild2, pleasant1*} is sufficient to unite all the overlapping feature denotations in one subsense (i). This subsense includes synonyms and related words referring to personal character (e.g., *amiable*) as well as words referring to appearance (e.g., *beautiful*). Raising the OverlapThreshold to 0.1 leads to the separation of those two kinds of synonyms in two subsenses (i) and (iii):

Version II:

OverlapThreshold: 0.1
SynsetLimit: 20

sweet
Hyperonyms: gentle, nice, good.
Subsense (i) (Translation: blid, snill.)
Synonyms:
amiable, amused, blithe, cheerful, cheery, easygoing, good-humoured, good-natured, jolly, kind, kindly, merry, mild, pleasant, smiling.
Related words: all_right, benign, friendly, pleasing, polite.
Subsense (ii) (Translation: frisk.)
Synonyms: all_right, brisk, eager, fit, fresh, healthy, new, pert, well.
Subsense (iii) (Translation: deilig, fin, god, pen, vakker.)
Synonyms: beautiful, charming, cute, delightful, enchanting, pleasurable, soft, splendid.
Related words:
attractive, delicious, dishy, fair, fancy, graceful, handsome, lovely, magnificent, mild, ornate, picturesque, pleasant, pretty.

Furthermore, lowering the SynsetLimit to 13 will alter the status of the features [vakker1|soft1], [snill1|pleasant1] and [blid2|sweet1], whose denotations will now exceed SynsetLimit. The consequences are that *soft* and *pleasant* are added to the hyperonyms (cf. Principle (1) above), while the senses denoted by the o-feature [blid2|sweet1] change status from synonyms to hyponyms (cf. Principle (4) above). This is shown in Version III:

Version III:

OverlapThreshold: 0.1
SynsetLimit: 13

sweet
Hyperonyms: pleasant, soft, gentle, nice, good.
Subsense (i) (Translation: blid, snill.)
Hyponyms:
amiable, amused, blithe, cheerful, cheery, easygoing, good-humoured, good-natured, jolly, kind, kindly, merry, mild, smiling.
Subsense (ii) (Translation: frisk.)
Synonyms: all_right, brisk, eager, fit, fresh, healthy, new, pert, well.
Subsense (iii) (Translation: deilig, fin, god, pen, vakker.)
Synonyms: splendid.
Related words: beautiful, charming, cute, delicious, delightful, enchanting, pleasurable.

On the other hand, raising the SynsetLimit to 25 makes the feature [pen1|gentle2] fall below the limit, changing the status of *gentle* from a hyperonym to a synonym (cf.

Principle (2a) above), and adding the senses in its denotation to the sets of synonyms and related words (cf. Principles (2c) and (3) above). This is shown in Version IV:

Version IV:

OverlapThreshold: 0.1
SynsetLimit: 25

sweet
Hyperonyms: nice, good.
Subsense (i) (Translation: blid, snill.)
Synonyms:
amiable, amused, blithe, cheerful, cheery, easygoing, good-humoured, good-natured, jolly, kind, kindly, merry, mild, pleasant, smiling.
Related words: all_right, benign, friendly, pleasing, polite.
Subsense (ii) (Translation: frisk.)
Synonyms: all_right, brisk, eager, fit, fresh, healthy, new, pert, well.
Subsense (iii) (Translation: deilig, fin, god, pen, vakker.)
Synonyms:
attractive, beautiful, charming, cute, delightful, dishy, enchanting, fancy, gentle, graceful, handsome, lovely, mild, picturesque, pleasant, pleasurable, pretty, soft, splendid.
Related words:
clean, delicious, elegant, fair, fine, first-class, formal, magnificent, neat, ornate, polite, tidy.

2.5 Interpreting the Results

The corpus data underlying the examples in this article are extracted on the basis of manual word alignment, while the project also studies results derived from data based on automatic word alignment. A central problem is finding ways to characterise the degree of success of the various methods. It is probably futile to search for a universal gold standard for the 'correct' semantic relations among the words in a language. The appropriate relations are subject to variation according to context, text type, speaker intuitions, and first and foremost: the purpose of the semantic description. We only need to compare published resources like the Princeton WordNet and Merriam-Webster's Thesaurus to find ample evidence for variation in the division into subsenses and the choice of related words for the same entries.

To some this might raise the question of the utility of such resources. Still, if we consider individual domains of discourse rather that the global language, the possibilities of finding intersubjectively valid semantic relations among the words seem much more promising. But the task of manually building large, domain-specific thesauri for many different discourse domains has forbidding proportions. That is why automatic methods like the Mirrors method are of some interest. To the extent that the Mirrors method is viable, the task before us is then reduced to providing tagged and word-aligned parallel corpora for the domains in question – in itself also a highly demanding task, but one which would serve other purposes as well, and which would be less dependent on subjective semantic evaluations than manual thesaurus construction.

On the other hand, the absence of independent gold standards poses a problem for the evaluation of the results of the automatic methods. We may conceive of at least three possible approaches to evaluation: (I) measuring precision and recall with

respect to one or more published resources like Princeton Wordnet and Merriam-Webster's Thesaurus, and interpreting the results in ways that take into account the doubtful status of these resources as gold standards; (II) consulting informants in controlled circumstances; (III) using the generated thesaurus resource in the performance of some other task, e.g., word sense disambiguation (WSD) or meaning-based information retrieval, and then measure the difference between the results achieved with and without the resource.

The last-mentioned approach (III) has been attempted for limited material in a masters thesis by Lyse (2003), who studied the use of the Mirrors method for WSD, and this approach will be pursued in further work. So will the first-mentioned approach (I), which so far has been applied to a selection of entries derived by manual word alignment (Thunes, to appear). At this point, still awaiting more comprehensive evaluation results, we will restrict ourselves to stating that the method appears clearly promising.

References

Aijmer, Karin, Bengt Altenberg, and Mats 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.

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.) 1998, pp. 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 Artyificial Intelligence ECAI 98.

Dyvik, Helge (2002): Translations as semantic mirrors: from parallel corpus to wordnet. To appear in a volume published by Rodopi. URL: http://www.hf.uib.no/i/LiLi/SLF/ans/Dyvik/ICAMEpaper.pdf

Fellbaum, Christiane (ed.) (1998): *WordNet. An Electronic Lexical Database.*Cambridge: The MIT Press.

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

Ide, Nancy, Tomas Erjavec & Dan Tufis (2002): Sense Discrimination with Parallel Corpora. *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 (1996), 87-112.

Johansson, Stig and Signe Oksefjell (eds.) (1998): *Corpora and Crosslinguistic Research: Theory, Method, and Case Studies*. Amsterdam: Rodopi.

Kilgarriff, Adam (1997): I don't believe in word senses. In: *Computers and the Humanities* 31 (2), pp. 91-113.

Lyse, Gunn Inger (2003): *Fra speilmetoden til automatisk ekstrahering av et betydningstagget korpus for WSD-formål*. Masters thesis, University of Bergen. URL: http://www.ub.uib.no/elpub/2003/h/516001/

Merriam-Webster online. URL: http://www.m-w.com/home.htm

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.

Thunes, Martha (Forthcoming): Evaluating thesaurus entries derived from translational features. To appear in: *Proceedings of the 14th Nordic Conference on Computational Linguistics. May 30-31, 2003. Reykjavík.*

Wordnet: a lexical database for the English language. URL: http://www.cogsci.princeton.edu/~wn/