How is Grammatical Gender Processed?

Christer Johansson

Section for Computational Linguistics, Bergen University

Abstract

Paradigmatic gaps are a problem for computational models of language acquisition, as most models that generalize online (eager learners, such as rule based learning and neural networks) will not notice systematically missing input. This is mainly a problem for the plausibility of the model, since the missing forms and structures will not deteriorate performance on recognition (because they will not be found often enough to matter). We are looking not only for a descriptive model of paradigmatic gaps, but also an explanatory model of why they emerge. The use for computational linguistics is that we can show how a linguistically motivated feature makes it possible to notice a negative regularity (i.e. that forms are missing), and this suggests that a hypothesis driven approach may be combined with statistical techniques (e.g. a memory-based learner) in interesting ways.

1 Introduction

This paper will introduce a problem for computational models that process language by learning and using generalizations. The essence of the problem is the existence of paradigmatic gaps. One case is the failure of an extremely frequent and simple regularity in the Swedish adjective paradigm—simply add t to make a neuter from a common gender form. Many neuter forms of adjectives are missing, regardless of the frequency of use for the common gender form. Two factors collaborate in causing this absence of neuters. First, all the problematic forms are monosyllabic, or have stress on the last syllable, and the gender marking is assimilated to that syllable making a good candidate for a free morpheme. The second necessary property is that the semantics of the adjective imply a property that cannot be verified by sensory experience only (such as internal states: fear, laziness, true happiness).

Swedish has two gender classes: common and neuter gender with agreement within the noun phrase. Semantically, the neuter gender in Swedish suggests a 'non-individualized', 'non-animate' noun. More semantic cues to gender assignment have been discussed in the literature (Trosterud 2001, Steinmetz 1986). Trosterud argues for a rule based account of gender assignment, in which gender classes are assigned on the basis of semantic cues. Trosterud's large number of rules, many with a small number of exceptions, shapes a good argument for a memory-based account of gender assignment, as the rules can be implemented in a decision tree structure.

A simple and very frequent regularity in the Swedish adjective paradigm is to add a /t/ to the common gender form to produce the neuter gender form. (This is only slightly complicated by predictable assimilation, such as devoicing.) The phenomenon also exists in Norwegian. Students learning Norwegian as a second language are sometimes given a list of adjectives with no proper neuter form to

memorize, but the phenomena is, as we will show, systematic in an interesting way (Pettersson 1990, Johansson 1999).

This article presents a corpus study, and a reaction time experiment. The corpus study estimated how exclusive the problematic context is. The reaction time experiment shows that the problematic adjectives have significantly longer decision times than congruent or non-congruent, existing or non-existing adjectives in both genders. Processing must encounter conflicting cues for gender agreement and/or the existence of the problematic forms. Similar studies are in progress for Norwegian.

Let me first begin by giving you a list of Swedish adjectives without a proper neuter gender form, mostly from Pettersson (1990): *kry, glad, god, ked, kåt, led, blid, pryd, snöd, strid, vred, flat, lat, rät, såt, fadd, rädd, disträ, kavat, brydd, gladd, gravid* and their rough translations: well/healthy, happy, morally good, bored, horny, bored, blithe, prudish, sordid, swift, angry, stale / not standing out, lazy, right-angled, great (pals), stale, forced, afraid, absent-minded, plucky/cocky, puzzled, made happy, pregnant.

If these words followed the grammatical rule, the neuter forms would have been: *krytt, glatt, gott, kett, kått, lett, blitt, prytt, snött, stritt, vrett, flatt, latt, rätt, sått, fatt, rätt, disträtt, kavatt, brytt, glatt, gravitt.* Most of these have phonetically, and sometimes also semantically, close neighbors that are accepted: *lytt, satt, —, fett, rått, hett, skitt, —, insnöat, stridigt, vredgat, flackt, lagt, skrämt, grått, matt, nött, skrämt, —, kavaljerskt, brutet, glatt, rigitt.* (In rough translation: handicapped, plump, —, greasy, raw, hot, dirty, —, snowed-in, fierce, angry, flat (landscape), put, scared, grey, faded, worn out, scared, —, chivalrous, broken, slippery or smooth, rigid).

Examples of unproblematic adjectives with either the same syllable structure or the same semantic field as the problematic adjectives are common, if only one of the conditions is present. The next section will give more detail on the anatomy of the missing neuter forms.

1.1 The Problem

It is not only that creating a neuter fails, but also that the failure occurs in a context that is unlikely to occur. Therefore it is very unlikely that we acquire knowledge about the oddness of these adjectives by memorizing a number of 'bad' exemplars. It has been argued (Johansson 1999) that the problem has its root in the fact that neuter gender nouns are very seldom words for individual beings. The problem is not directly observable in examples, and it arise for structural reasons. The problem is neither exclusively *phonological* nor *semantic*.

The problematic adjectives have a semantic reference to hidden non-verifiable properties, such as *lazy*, and *morally good*. A homonym of 'morally good', meaning 'tasty' is generally unproblematic. *'Ett gott barn'* (a good child) would therefore be heard as a tasty child, despite its semantic absurdity. Native speakers avoid the problematic constructions; e.g. *en god unge* (a warm-hearted brat).

An example of a successful neuter gender is *rö:d+t = rött* / red. Examples

of problematic adjectives are *la:t+t = ?latt*/lazy, and *rädd+t = ?rätt*/frightened. Cases where the denoted property can be immediately verified from sensory experience are all unproblematic. There are some accepted cases with reference to hidden (non-verifiable) properties with either a) one form in common to both genders (e.g. *trött*/tired, or *lätt*/easy (light)) or b) a t-ending that is not assimilated by the adjective stem (e.g. *skrämd+t = skrämt*/scared (by someone)).

The problematic adjectives range from very low to *high* frequency of the common gender form. The problem is systematic, and productive, and is not restricted to a handful of bad examples that could be memorized.

The problem exists in Swedish, Norwegian, and to some extent also in Danish. Distinctive forms for verifiable or non-verifiable qualities exist in, at least, one Slavic language: Russian *krasiva* is intrinsically beautiful, whereas *krasivaja* is observably beautiful (Pettersson 1990). The verifiability dimension can alternatively be thought of as a public/private distinction: Is the quality open for observation or not?

Additionally, there is neurolinguistic evidence. A specific kind of aphasia/agnosia exists, where the affected patient has lost the ability to name either living or non-living entities (Warrington and Shallice 1984). It has been argued that these covert noun classes (in English without grammatical gender) may be formed by how words are described, which links back to adjectives. Noun classes and adjectives seem to help define each other.

As always there is some complicating factor. For example, one word *glad/?glatt* is not only polysemous (happy or slippery), but it is sometimes possible to immediately verify either property. A happy face, or a happy message, can easily be verified, and are used idiomatically. On the other hand, we sometimes show a happy face when we are not really happy, and a happy message is not really happy in itself. A message can also be verified for its happy implications to us. To me, the first interpretation of *ett glatt barn* is a slippery child, not a happy child. The congruency is also grammatically correct for the sense of a slippery child. The consequence of this is that the adjective *glad* was omitted from the following reaction time experiment, as it is not clear which category the word would belong to. It could be that verifiability is really a matter of degree, but for practical purposes the experiment need to concentrate on the clearest cases.

2 Background

Paradigmatic gaps is a much cleaner version of the poverty-of-stimulus argument (Chomsky 1975) than the discussion on learning past-tense forms, which has gone back-and-forth within connectionism for more than a decade (Rumelhart et al. 1986, Pinker and Prince 1988, Daugherty and Seidenberg 1992, Pinker 1989, Johansson 1997, inter al.). Some of the problems for connectionist models were thought to be solved by adding more structure (Daugherty and Seidenberg 1992) and limitations on early perception (Elman 1993). Others suggested that more-or-less free variation between two or more past tense forms was a real challenge to neural network learning (Johansson 1997). General neural network mod-

els also had inherent problems, such as catastrophic forgetting (Carpenter and Grossberg 1992, French 1999), and difficulties separating supplied facts from knowledge invented by generalization.

The task of learning past tense forms was not a clear case of poverty-of-stimulus. One common assumption was that there is one (and only one) correct past tense form (one exception is (Johansson 1997), where alternating strong and weak forms, as well as possibly missing verb forms, are discussed). Instead of poverty of stimulus, we find a sometimes too rich stimulus, as a correct past tense form can be used as negative evidence to all other forms.

A paradigmatic gap is another matter altogether: The correct forms are entirely missing. Ways to solve this problem involve either using a regular form, despite knowing that it is perceived to be incorrect, or to rephrase and avoid the missing construction. When informally trying to elicit missing neuter forms from young children, I found the child would prefer to use a safer construction rather than produce the neuter. "And then the hunter scared the *lion*$_{neu}$ with a loud noise. *It* quickly ran away, how do you think *it* felt?" "The lion *he* was afraid$_{com}$."

One characterization of a paradigmatic gap is that a normally productive generalization is blocked, in other words generalization is not general, and may depend on, possibly innate, factors that are deep in our perception of the world around us.

It is quite well-known that the back-propagation is able to learn any one-to-one function. Johansson (1997) pointed out that the Swedish past tense is more like a one-to-many mapping, which is not possible to learn with back-propagation. Paradigmatic gaps are problematic because they are a many-to-nothing mapping (i.e. there are no examples of the failing forms).

Some mechanisms are cautious about generalization. A lazy, memory-based, learner (Daelemans et al. 1994) would ideally not generalize to cases that are not supported in a larger database. The model still depends on the strength of the features that are used to characterize the indata. Instance-based models of learning and generalization have been investigated as a plausible alternative to approaches to learning based on function approximation and algorithmic processing (Logan 1988, Wattenmaker 1993, Lamberts 1994). It is unclear what the limits of instance based learning are, but such models are, to my knowledge, those least vulnerable to over-generalization, which is one of the recurring, general problems in computational linguistics.

3 Methods

3.1 Corpus Study

The distribution of syllabic structure and semantic reference can be found from very large corpora. In 1.2 million words from the Stockholm Umeå Corpus (Stockholm-Umeå Corpus 1.0. 1997, Ejerhed et al. 1992), we found 357 *different* monosyllabic adjectives, out of which 152 were neuter forms of which only 27 had a non-detachable t. The 27 forms could be divided into two groups: 1) those that had the same form in both genders (e.g. *trött*/tired), and 2) those that

had reference to observable or verifiable properties (e.g. *sött*/sweet).

This corpus study has many flaws. It is based on quite a small corpus (although it is one of the largest Swedish part-of-speech tagged corpora available), and it is based on material that has been edited and corrected (such as newspaper texts), and professional writers originally wrote the material.

The most interesting fact that can be derived from the study is that there is support for the paradigmatic gap in the neuter adjective forms. Let me just try to characterize the data for you in distributions. If we start with all adjectives, and divide them up in common and neuter gender forms we see them distributed as we would expect from the distribution of modified nouns. Let us divide the neuter forms into two categories: 1) those with a discernible morphology ("detachable t"), and 2) those without such internal morphology, i.e. good candidates for free morphemes. These two groups also seem to behave more or less as expected. Let us divide these two groups into subgroups based on how we can verify that the property is present. We will clearly see that there are missing forms, namely the subgroup of neuter adjectives lacking internal morphology and denoting properties that cannot be verified by sensory experience. We might see occasional examples that would fit this subgroup, but at a frequency comparable to misspelled or wrongly marked adjectives.

A model of the data, based on Royal Skousen's AML model (Skousen 1989), showed that there is a tendency for the problematic adjectives to be more supported as common gender forms. It is still an open question if we can find a measurable response similar to a gender mismatch for these failing forms.

The corpus study cannot objectively prove that there is a paradigmatic gap, but we can test if the failing adjective forms are reacted to differently than other existing or non-existing adjectives. This could add more objective evidence to the reality of the missing adjective forms. One criticism is otherwise that the discussed paradigmatic gaps do not really exist, they are just the invention of introspection. I intend to show that the phenomenon is real. If it is real, computational models of language learning with pretensions to psychological relevance should deal with it.

3.2 Reaction Time Experiment

We have performed a reaction time experiment in Swedish to further investigate the problem (Johansson and Zlatev 2002). A similar study is in progress for Norwegian. The reaction time for a lexical decision (good–no good) was measured for 25 subjects from a population of linguistics students and faculty at Lund University (6 subjects with too slow responses were excluded). The subjects, according to themselves, had acquired Swedish as a first language. Some of the excluded subjects had acquired several languages in childhood. The subjects were further instructed that it was possible to answer 'no good' to all non-words, and that the test was not intended to measure their knowledge of the norms of grammar.

The test material consisted of 240 Swedish noun phrases of the type: *article adjective noun* (all items must agree in gender). The examples were presented in 4 blocks with pause in between. The test material was classified into 10 test

Table 1: Pairwise comparisons. C=common N=neuter, congruency=YES/NO, PROB=problematic congruency ($** = p < 0.01$ $* = p < 0.05$ $-$ =not sign).

Test conditions - existing adjectives

	CYES	CNO	NYES	NNO	NPROB	CPROB
cases	760	380	380	760	304	304
mean	0.96	1.13	1.02	1.22	1.33	1.23
CYES		**	**	**	**	**
CNO			**	**	**	**
NYES				**	**	**
NNO					**	-
NPROB						**

conditions: NPROB (problematic forms in neuter context), CPROB (likewise in common gender context), CYES and NYES (normal congruency in common and neuter context), CNO and NNO (incongruent adjectives in respective gender context), and non-existing adjectives in both genders and types of congruency.

The conditions [+/-congruent, common/neuter gender, +/-problematic, +/-real adjective] were balanced for phonological complexity and semantic field (if possible) of both nouns and adjectives. For example the neuter *barn*/child was matched by common *unge*/kid and *rädd*/frightened was matched by *skrämd*/scared. All scores were normalized by the overall *median* score. An analysis of variance (ANOVA) found a very significant overall difference between conditions, and specific effects were found using pairwise t-tests (table 1 shows the results for existing adjectives and table 2 shows the results for made-up adjectives). (I recently checked the experiment and found a few items in the CYES condition for nonsense words that had accidentally been given the wrong article. This might explain why there was only one star for NYES (1.02) compared to CYES (1.08) in table 2).

The results show that the reactions to the problematic forms give 30% slower response times than correct agreement, and 10% slower responses than the incongruent case (table 1). In fact, the problematic forms are significantly slower than any of the other conditions. In general, neuter contexts are slower to decide, also for nonsense adjectives. CPROB should be comparable to CNO, but it is significantly slower. This shows that that the problem is objectively observable in reaction time experiments. For existing adjectives, the first factor is that congruency is faster to decide than non-congruency, and secondly that common gender is faster than neuter gender. The order $CYES < NYES < CNO < NNO <= CPROB < NPROB$ was extremely noticeable ($p < 0.0001$).

For nonsense adjectives the results are somewhat less clear, but the cases with congruent gender are *slower* to decide than their corresponding CYES or NYES, and similarly cases with *incongruent* gender for nonsense adjectives are *faster* to decide than their corresponding CNO or NNO. The likely explanation for this is that a decision is faster if there are consistent cues pointing to the same answer,

Table 2: Pairwise comparisons for invented words. C=common N=neuter, congruency=YES/NO, PROB=problematic congruency. The top 4 rows are existing vs non-existing words. The last 3 rows are non-existing vs non-existing words ($** = p < 0.01$ $* = p < 0.05$ $- =$not sign).

Test conditions - nonsense adjectives

	CYES	CNO	NYES	NNO
cases	304	304	304	304
mean	1.08	1.05	1.12	1.11
CYES 0.96	**	**	**	**
CNO 1.13	-	**	-	-
NYES 1.02	*	-	**	**
NNO 1.22	**	**	**	**
NPROB 1.33	**	**	**	**
CPROB 1.23	**	**	**	**
CYES 1.08		-	-	-
CNO 1.05			*	-
NYES 1.12				-

and slower if the cues are leading in different directions. The effect is rather large, 10 percent faster or slower. The observation that reactions to common gender generally are faster is likely a frequency effect, as common gender is the most frequently occurring gender in Swedish.

Problematic adjectives (NPROB) may have a similar status to nonsense adjectives, or they may cause gender disagreement (Johansson 1999). If they were non-words there is no reason why they were significantly slower than nonsense adjectives (table 2). Similarly, if it were only gender disagreement we would expect similar reaction times for NPROB as for NNO, but in fact NPROB was significantly slower.

Do we have conflicting evidence for gender? Is it that the added *t*-morpheme points to neuter gender, and the support from examples point to common gender? We can conclude that there need to be more conflict involved for the problematic adjectives, since they are clearly slower to decide than both incongruency and non-words in general.

4 Discussion

One reasonable question was put to us at an early presentation of the experimental results: *Why doesn't German have similarly missing adjectives?* German is morphologically similar to Scandinavian, but has three gender classes. Neuter gender in German is similarly correlated with non-sentient or non-living.

One good reason is that in German the neuter gender is a default gender; for example many exotic animals are neuter gender: *das Dromedar, Kamel, Krokodil,*

Oppossum, Zebra, as are many domestic animals: *das Pferd, Schwein*, and diminutives: *das Mädchen, Fräulein.*

An informal inquiry found that German speakers more easily than Scandinavian speakers find sentient neuter nouns. Sentient beings are good examples of entities with internal states that are normally hidden from our direct perception. Scandinavian speakers showed severe difficulties finding sentient neuter nouns, typically some 5 minutes to list 10 such nouns, compared to about half that time for the Germans. A baseline for normal English speakers is about 20 seconds. Scandinavians were also slow to find sentient nouns for common gender (typically about 2 minutes). We can conclude that, in these cases, grammatical gender does not facilitate lexical lookup.

4.1 Situational Priming and Finite State Morphology: Mind the Gap

Another question has to do with why there is a gap. Could it be that the gap is, in some sense, functional? Morphology is thought to be finite state (Jurafsky and Martin 2000), which implicates reversible morphology. We have discussed a process that may have lost reversibility if it was not for the semantic cue. If the property is immediately verifiable (i.e. primed) then neutralizing the /t/-ending retrieves the correct common form. If the property is not verifiable, then nothing speaks against treating a problematic adjective as a free morpheme that could either be in the lexicon, or be independently added to it. This supports an idea of a tight link between syllable and morpheme, and acquisition guided by situational priming.

There is a rather strong prediction coming out of this idea that acquisition of the paradigmatic gap relies on situational priming. Certain handicaps could prevent us from direct experience of some dimensions of our world. One example: people who are born blind would have no experience of colors. A similar problem to the discussed problem would then arise if a color word has a neuter gender form that is a good candidate for a free morpheme.

One such word is the Swedish word for blue (*blå – blått*). The prediction is that a blind person who was unable to verify the blueness of a sentient being (i.e. an entity capable of internal, hidden states) would react slower to judging the neuter gender form of blue (blått) than to the neuter gender form of yellow (*gul+t*). We will also have to establish that the blind subjects have acquired the usual paradigmatic gaps, as discussed in this article. We can thus expect an additional significant difference between the neuter forms of 1) colors like: blått/blue, grått/gray, and rött/red, and 2) colors like: gult/yellow, grönt/green and brunt/brown for blind people (but not for seeing people). If we get that result it would give extremely strong support for the explanation of the paradigmatic gap outlined throughout this article. The only caveats are that the surrounding language is shaped by people with the ability to see, and that the neuter gender forms of colors are commonly used to modify both neuter and common gender nouns; thus a negative result is not so much a contradiction of the explanation. I hope to have shown that there are testable hypotheses coming out of the presented theoretical insight. Is computa-

Table 3: Distribution of examples. 1=non-assimilated t ending; 2=assimilated t ending. Adjectives with the same form in both genders are excepted.

	All data	
	1	**2**
Neuter	many (100 000)	many (50 000)
Common	many (200 000)	many (100 000)

tional linguistics to be an engineering project or a scientific field?

4.2 The Challenge

The challenge to the increasingly popular, but linguistically irrelevant, general data mining approaches to automated learning of language is to show how a paradigmatic gap can be acquired without using the linguistically relevant feature(s). Without the verifiability feature, the data would look like in table 3. Note that this feature cannot be changed without changing the meaning of the adjective. A second challenge is to provide interesting and testable hypotheses. Data mining has so far provided very little of that sort.

You might argue that the feature could be detected by co-occurrence with nouns denoting living beings. There are, of course, living beings in both gender classes. In addition, the non-problematic adjectives may be used with all kinds of nouns without anybody protesting that the adjective form is ungrammatical.

I have actually simplified the above characterization to make the task easier for the data miners. There is also phonological complexity to consider. Each adjective would take several dozen phonological features to describe accurately.

If we use the verifiability feature (which we have discovered by fundamental insights into the problem, not by automated data mining) the data easily sorts into the data in table 4. The numbers have been provided solely for pedagogical purposes. In the corpus study, problematic adjectives were found at a frequency lower than ungrammatical congruency. This is a good example of how a linguistic insight can reveal a pattern in the data that would be hidden otherwise.

What kind of features would we need in our data collections to facilitate machine learning? This is a good question. We have shown that quite unexpected features may turn out to be useful. Another question could be: What kind of sensory perception would facilitate learning?

5 Conclusion and Implications for Computational Models

I have shown that the problematic forms are missing in a (small) sized corpus, and that they are reacted to differently than congruent or incongruent adjectives; for both existing adjectives and made up adjectives with correct phonotactics. Additionally, Skousens general AML-model of analogical support has difficulties in supporting the problematic forms as the intended neuter gender. A reasonable account for the underlying factors of the phenomena has been given, which fits pos-

Table 4: Distribution of examples. 1 = non-assimilated t ending; 2 = assimilated t ending. Adjectives with the same form in both genders are excepted.

	Non-verifiable	
	1	**2**
Neuter	many (20 000)	*nothing* (10)
Common	many (40 000)	many (20 000)
	Verifiable	
	1	**2**
Neuter	many (80 000)	many (49 990)
Common	many (160 000)	many (80 000)

itive data on the phenomenon in Swedish and Norwegian, and explains the lack of the phenomenon in German.

I have experimentally observed that the problematic adjectives are different from 'nonsense' words, as well as different from ordinary discongruency. Missing gender forms might be central to investigations of the underlying semantics of grammatical gender, as we are looking not only for a descriptive model of what is, but also an explanatory model of why some forms are missing.

The Swedish missing forms are a problem for some rule-based models, since they imply a cross-modular dependency between syllable structure and semantics. Most neural networks, eager learners, and prototype models without stored examples, will not notice that a form is systematically missing. The irregular form is a necessary negative feedback in the classic case of learning irregular past tense forms with neural networks. What is really special about paradigmatic gaps is that there are no irregular forms either.

A lazy learner is a more plausible model, as it first stores positive exemplars, and later it may find out that there are no examples of a specific combination of factors, some of which factors may have 'emerged' after exemplars are collected. Most models of lazy learners have mechanisms that may over-generalize, but it can be argued that they are less prone to over-generalization, since they are always looking for support from a database of instances, and their generalizations are (only) made at the time when generalization is needed. Models that do not store exemplars will have problems detecting that some forms and expressions are systematically missing.

The discussed observations support a cautious approach to generalization. Models that readily generalize and store their generalizations (e.g. most neural networks) would find it difficult to restrict generalization so that the model reflects the problems human native speakers experience.

The discussed formation rule (add t) is very simple, and very commonly applied, and, to make matters worse, the problematic adjectives occur in a very restricted context that makes negative examples very rare. It takes a special kind of learner to find the gaps in that paradigm.

Acknowledgements

I thank three anonymous reviewers, and Tanja Gaustad for valuable comments in preparing this document. I thank Jordan Zlatev for help with designing and conducting the presented reaction time experiments, and Janne Cecilie von Koss Torkildsen for her enthusiasm and her efforts with the yet to appear Norwegian reaction time experiment. Finally, I thank Jill Walker for her kind suggestions that made the article clearer and hopefully more enjoyable for the reader.

References

Carpenter, G. and Grossberg, S. (1992), A self-organizing neural network for supervised learning, recognition, and prediction: Can neural networks learn to recognize new objects without forgetting familiar ones?, *IEEE Communications Magazine* **30**, 38–49.

Chomsky, N. (1975), *Reflections on Language*, Pantheon Books, New York.

Daelemans, W., Gillis, S. and Durieux, G. (1994), Skousen's analogical modeling algorithm: A comparison with lazy learning, *Proceedings of NeMLap*, Manchester, pp. 1–7.

Daugherty, K. and Seidenberg, M. (1992), Rules or connections? the past tense revisited, *Proceedings of the 14th Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum Associates, Hillsdale, pp. 259–264.

Ejerhed, E., Källgren, G., Wennstedt, O. and Åström, M. (1992), The linguistic annotation system of the Stockholm-Umeå corpus project, *Technical Report 4.31*, Department of Linguistics, Umeå University, Umeå.

Elman, J. (1993), Learning and development in neural networks: The importance of starting small, *Cognition* **48**, 71–99.

French, R. (1999), Catastrophic forgetting in connectionist networks, *Trends in Cognitive Sciences* **3** (4), 128–135.

Johansson, C. (1997), Connecting Swedish verb forms, *Nordic Journal of Linguistics* **20** (1), 3–30.

Johansson, C. (1999), Learning what cannot be by failing expectations, *Nordic Journal of Linguistics* **22**, 61–76.

Johansson, C. and Zlatev, J. (2002), Rule or analogy in gender (dis-)agreement?, Presentation at Brain and Cognition in Århus, Denmark. Abstract available at `http://www.brain-cog.au.dk/`.

Jurafsky, D. and Martin, J. (2000), *Speech and Language Processing*, Prentice Hall, Englewood Cliffs.

Lamberts, K. (1994), Flexible tuning in exemplar-based categorization, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **20** (5), 1003–1021.

Logan, G. (1988), Towards an instance theory of automatization, *Psychological Review* **95** (4), 492–527.

Pettersson, T. (1990), Varför barnet inte kan vara latt [why the child cannot be lazy], *in* E. Andersson and M. Sundman (eds), *Svenskans Beskrivning*, Åbo

Academic Press, Åbo, pp. 293–302.

Pinker, S. (1989), *Learnability and Cognition: The acquisition of argument structure*, MIT Press, Cambridge.

Pinker, S. and Prince, A. (1988), On language and connectionism: Analysis of a parallel distributed processing model of language acquisition, *Cognition* **28**, 73–193.

Rumelhart, D., Hinton, G. and Williams, R. (1986), Learning internal representations by error propagation, *in* D. Rumelhart, J. McClelland and the PDP research group (eds), *Parallel Distributed Processing: Explorations in the microstructure of cognition*, Vol. 1: Foundations, MIT Press, Cambridge, pp. 318–362.

Skousen, R. (1989), *Analogical Modeling of Language*, Kluwer Academic Publishers, Dordrecht.

Steinmetz, D. (1986), Two principles and some rules for gender in German: Inanimate nouns, *Nordic Journal of Linguistics* **37**, 189–217.

Stockholm-Umeå Corpus 1.0. (1997), CD-ROM. Department of Linguistics, Umeå University. `http://www.ling.su.se/DaLi/`.

Trosterud, T. (2001), Genustilldelning i norsk er regelstyrt [gender assignment in Norwegian is regular], *Norsk Lingvistisk Tidskrift* **19**, 29–58.

Warrington, E. and Shallice, T. (1984), Category specific semantic impairments, *Behavioral and Brain Sciences* **107**, 829–853.

Wattenmaker, W. (1993), Incidental concept learning, feature frequency, and correlated properties, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **19** (1), 203–222.