Corrections and Higher–Order Unification

Claire Gardent, Michael Kohlhase, Noor van Leusen

Abstract

We propose an analysis of corrections which models some of the requirements corrections place on context. We then show that this analysis naturally extends to the interaction of corrections with pronominal anaphora on the one hand, and (in)definiteness on the other. The analysis builds on previous unification–based approaches to NL semantics and relies on Higher–Order Unification with Equivalences, a form of unification which takes into account not only syntactic $\beta\eta$-identity but also denotational equivalence.

Wir schlagen eine Analyse vor, die einige der Anforderungen von Korrekturen an den Kontext modelliert und sich natürlich auf die Interaktion von Korrekturen mit Pronominalanaphern und Undefiniertheit erweitern läßt. Die Analyse basiert auf bekannten unifikationsbasierten Ansätzen für die Semantik natürlicher Sprache und benutzt eine Erweiterung der Unifikation höherer Stufe. Diese berücksichtigt nicht nur strukturelle $\beta\eta$-Gleichheit, sondern auch logische Äquivalenz.

1 Introduction

Corrections are utterances such as (1b) where a discourse participant corrects the utterance of some other discourse participant[1].

(1) a. *A: Jon likes Mary.*
  b. *B: No, PETER likes Mary.*

Although there is much literature on corrections (e.g. [SJS77, Nor91, RB82]), a thorough investigation of their linguistics is still outstanding. In this paper, we build up on [vL94] and examine some of the requirements corrections place on context or in other words, the relationship between correction (the correcting utterance) and correctee (the utterance being corrected). For instance, it is clear that the pair of utterances in (2) does not form a well–formed dialog.

(2) a. *A: Jon likes Mary.*

[1] Here and in what follows, we use capital letters to indicate prosodic prominence.

b. $\star$ *B: No, PETER likes Sarah.*

On the other hand, it is also clear that a simple equality requirement between the semantic representation of the deaccented part of the correction and that of its parallel counterpart in the source is not appropriate either:

(3) a. *A: Jon likes [the woman with the red hat]*$_1$

b. *B: No, PETER likes Sarah*$_1$

Here the correction contains an NP *Sarah* whose semantic representation is not identical with that of its source parallel element *the woman with the red hat.* In other words, a requirement such as [Sag76]'s alphabetical variant constraint would fail[2]. At this stage one could be tempted to conclude that the equality requirement is a semantic one: the deaccented part of the correction must be semantically equivalent with its parallel correlate in the source utterance. However, this is also incorrect. Thus in (4), the property denoted by the VP in (4b) need not be the same as the property denoted by its parallel counterpart in (4a): whereas the VP in (4a) denotes the property of loving Jon's wife, the VP in (4b) may denote the property of loving Peter's wife[3].

(4) a. *A: Jon*$_1$ *loves his*$_1$ *wife.*

b. *B: No, PETER loves his wife.*

In short, it is clear that some identity requirement is needed to appropriately characterise the relation between correctum and correction (cf. example 2). On the other hand, it is less clear what this identity requirement should be (cf. examples 3,4). In this paper, we contend that the correct notion of identity is given by Higher–Order Unification with equivalences, a form of Unification which takes into account not only syntactic identity, but also denotational equivalence. We show that the HOUE–based analysis of corrections we propose, not only captures some of the contextual requirements of corrections, but also makes appropriate predictions about the interaction of corrections with both pronominal anaphora and (in)definiteness.

2 HOU with Equivalences

Now we will briefly review higher-order unification and its properties, for details we refer the reader to [Sny91]. Higher-order unification solves the problem

[2] Sag proposes an analysis of VP ellipsis which requires that the semantic representation of a VP ellipsis be an alphabetical variant of the semantic representation of its antecedent. The basic assumption is that semantic representations are λ–terms. Two terms are alphabetical variants of each other iff they are identical up to renaming of bound variables.

[3] This is of course similar to the sloppy/strict ambiguity characteristic of VP ellipsis. Indeed, as we shall later see, our treatment is very similar to [DSP91]'s treatment of VPE.

of finding substitutions σ that for a given equation $A = B$ make both sides equal in the theory of $\beta\eta$-equality ($\sigma(A) =_{\beta\eta} \sigma(B)$). Huet's well-known algorithm [Hue75] solves the problem by recursively decomposing formulae and binding Function variables to most general formulae of a given type and given head.

However, even though HOU considers $\beta\eta$-equality of formulae, it does not take into account the semantics of the logical connectives and quantifiers contained in the logical representation of natural language utterances. For this we need a unification algorithm for $\beta\eta$-equality augmented by logical equivalence. Obviously, such an algorithm has to generalize theorem proving methods for higher-order logic, since the task of unifying an equation $(A \vee \neg A) = T$, where T is a sentence, is equivalent to proving the validity of the theorem T[4]. An algorithm that solves this problem is described in [Koh95]. It is a generalization of the first-order Tableau method [Fit90] for automated theorem proving, which refutes a negated theorem by analyzing the connectives in an and/or tree and finding instantiations that close each branch of the tree by finding elementary contradictions on it.

Instead of a formal recapitulation of the tableau method, we discuss the example of the logical theorem $(p(a) \vee p(b) \Rightarrow \exists x.p(x))$. The negation of this is equivalent[5] to the formula at the root of the following tableau.

$$\begin{array}{c}
p(a) \vee p(b) \wedge \forall x.\neg p(x) \\
p(a) \vee p(b) \\
\forall x.\neg p(x) \\
\begin{array}{c|c}
p(a) & p(b) \\
\neg p(y) & \neg p(z) \\
* \, [y = a] & * \, [z = b]
\end{array}
\end{array}$$

Here we see that conjuncts are simply added to the branch, whereas disjuncts are analyzed in separate branches of the tree. The scopes of universal quantifications (with new variables) can be inserted at the end of branches, the same is possible with the scopes of existential quantifications (with the bound variables replaced by Skolem[6] terms). Finally, both branches of the tableau are closed, i.e. the last formula can be instantiated (by the substitution in brackets) so that it contradicts a formula in the branch above.

These instantiations are computed by unification, and in the case of higher-order logic by HOU. The distinguishing feature of the HOUE algorithm [Koh95] is that intermediate equations $(A = B)$ of type t (generated either by unifying two formulae on the branch to make them contradictory or by processing other unification problems) can be transformed into negated equivalences (which can

[4] The formula $(A \vee \neg A)$ must be true in all models, so T can only be equivalent to it, if it is a theorem.

[5] In addition to the de Morgan laws we use the identity $\exists x.A = \neg\forall x.\neg A$.

[6] Skolem terms serve as witnesses for the objects whose existence is claimed by the existential formula A. Since this object may depend on the values of free variables $x_1, \ldots, x_n$ occurring in A, they have the form $f(x_1, \ldots, x_n)$ where f is a new function.

then be treated by the theorem proving component). Actually, tableau development for the negated equivalence $\neg(A \Leftrightarrow B)$ contains trivial branches, so we use the following (optimized) rule, which splits an equation of type t into two tableau branches

$$\begin{array}{c} A = B \\ \begin{array}{c|c} A & B \\ \neg B & \neg A \end{array} \end{array}$$

This way, HOU and tableau theorem proving recursively call each other in HOUE, until a refutation is found (all branches of the tableau are closed).

3 The basic analysis

Typically, a correction partially or completely repeats a previous utterance and one of its characteristic properties is that the repeated material is deaccented, that is, it is characterised by an important reduction in pitch, amplitude and duration (cf. [Bar95]). Our proposal is to analyse corrections as involving a deaccented anaphor which consists of the repeated material. Furthermore, we require that the semantic representation of a deaccented anaphor unify with the semantic representation of its antecedent.

More precisely, let $SSem$ and $TSem$ be the semantic representations of the source (i.e. antecedent) and target (i.e. anaphoric) clause respectively, and $TP^1 \ldots TP^n, SP^1 \ldots SP^n$ be the target and source parallel elements[7], then the interpretation of an SOE must respect the following equations:

$$\begin{array}{l} An(SP^1, \ldots, SP^n) = SSem \\ An(TP^1, \ldots, TP^n) = TSem \end{array}$$

Intuitively, these two equations require that target and source clause share a common semantics: An, the semantics of the deaccented anaphor. We illustrate the workings of the analysis by a simple example. Given the dialog in (1), the equations to be solved are:

$$\begin{array}{l} An(j) = like(j, m) \\ An(p) = like(p, m) \end{array}$$

Given these equations, HOU yields a unique solution $An = \lambda x.like(x, m)$. In contrast, the equations required for the analysis of example (2) are:

$$\begin{array}{l} An(j) = like(j, m) \\ An(p) = like(p, s) \end{array}$$

[7] As in [DSP91], we take the identification of parallel elements as given.

Since there is no substitution of values for free variables which simultaneously makes $An(j)$ $\alpha\beta\eta$–identical with $like(j, m)$ and $An(p)$ $\alpha\beta\eta$–identical with $like(p, s)$, unification fails thereby indicating the ill–formedness of (2).

4 Corrections and pronominal anaphora

The resolution of pronouns occurring in the destressed part of a correction appears to be subject to very strong parallelism constraints. For instance in (5b), the pronoun *her* can only be understood as referring to its source parallel element *Sarah* – else it must be stressed.

(5) a. *Jon loves* $Sarah_1$.
b. *No, PETER loves her.*

Intuitively, there is a simple explanation for this: if the destressed part of a correction is a repeat of its parallel element in the source utterance, then pronouns occurring in it must necessarily resolve to their parallel counterpart in the source expression. As we shall see, the picture is somewhat more complex however. In some cases, a destressed pronoun in the correction may be ambiguous. In other cases, it functions as a paycheck pronoun. Finally, extraneous factors such as scope constraints and world knowledge interact with the semantics of corrections in determining the resolution of destressed pronouns. In what follows, we show how HOUE allows us to correctly predict this array of empirical facts.

4.1 Pronouns

Let us start with example (5) above. Given the analysis of corrections described in section 3, the equations to be resolved are[8]:

$$\begin{array}{lcl} An(j) & = & love(j, s) \\ An(p) & = & love(p, x) \end{array}$$

By unification, the only possible values for An and x are $\lambda y\ love(y, s)$ and s respectively. That is, the destressed pronoun is resolved by unification to its parallel element in the source utterance, *Sarah.* As required.

In some cases however, a destressed pronoun in the correction is ambiguous. For instance in (6b), the pronoun *his* may resolve either to *Jon* or to *Peter.*

(6) a. Jon_1 *loves* his_1 *wife.*

[8]Unresolved pronouns are represented by free variables i.e. variables whose value is determined by unification. Alternatively, pronouns could be resolved first and unification would then function as a filter on admissible resolutions.

b. *No, PETER $_2$ loves his$_{1,2}$ wife.*

Interestingly, such cases are similar to the sloppy/strict ambiguity[9] characteristic of VP ellipsis and as [DSP91] have shown, HOU straightforwardly captures such cases because of its ability to yield several solutions. In the case of (6), the analysis proceeds as follows. First, the following equations must be resolved:

$$\begin{aligned} An(j) &= love(j,\ wof(j)) \\ An(p) &= love(p,\ wof(x)) \end{aligned}$$

Resolution of the first equation yields two values for An[10], $\lambda y\ love(y,\ wof(j))$ and $\lambda y\ love(y,\ wof(y))$. By applying An to p, we then get two possible values for $An(p)$: $love(p,\ wof(j))$ and $love(p,\ wof(p))$. As a side effect, the pronoun *his* represented by x is resolved either to *Jon* or to *Peter*. In short, for such cases, the multiple solutions delivered by HOU match the ambiguity of natural language.

4.2 Paycheck pronouns

Destressed pronouns whose source parallel element is a pronominal possessive NP are particularly interesting. At first sight, they seem to behave just like any other destressed pronouns occurring in a correction, that is, they seem to resolve unambiguously to their parallel source element. For instance, in (7b), the most likely resolution of *her* is *Jon's wife.*

(7) a. *Jon$_1$ likes his$_1$ wife.*

b. *No, PETER likes her (= his$_1$ wife)*

However, a closer investigation of the data suggests that this reading is a kind of default reading which is preferred out of a pair of two grammatically possible interpretations. To see this, consider examples (8) and (9).

(8) a. *Jon$_1$ broke his$_1$ arm yesterday.*

b. *No, PETER$_2$ broke it (= his$_{1,2}$ arm) yesterday.*

[9] The terminology *sloppy/strict* originated with [Ros67]. Intuitively, a pronoun has a strict interpretation if it denotes as its antecedent. By contrast, a pronoun which denotes differently from its antecedent is said to have a sloppy interpretation.

[10] Unification yields a third value for An, namely $\lambda y\ love(j,\ wof(y))$. This solution however is ruled out by the second equation. More generally, we assume a restriction similar to [DSP91]'s **Primary Occurrence Restriction** (POR): the occurrences directly associated with the contrastive elements are primary occurrences and any solution containing a primary occurrence is discarded as linguistically invalid. For instance, in $An(j) = love(j,\ wof(j))$, the first occurrence of j is a primary occurrence so that the solution $An = \lambda y\ love(j,\ wof(y))$ is ruled out. For a proposal of how the POR can be formally modelled, see [GK96].

(9) a. *Jon$_1$ had his$_1$ nose remodelled in Paris.*

b. *No, PETER$_2$ had it (= his$_2$ nose) remodelled in Paris.*

Although these examples are structurally identical with (7), they differ in the interpretation of the destressed pronoun occurring in the correction. Whereas (7) only allows for a strict interpretation of this pronoun, (8) permits both a strict and a sloppy interpretation whilst (9) only admits of a sloppy reading.

Our contention is that a destressed pronoun in the correction whose source parallel element is a possessive definite, is systematically ambiguous between a strict and a sloppy interpretation. However extraneous factors may have the effect that only one reading is available. For instance, in (9) the strict reading is ruled out by our world knowledge that one can only have one's own nose remodelled. As for (7), the absence of sloppy reading can be explained if we assume that the interpretation of a destressed anaphor follows a default strategy geared toward maximal semantic identity between the destressed anaphor and its antecedent. Under this assumption, the strict reading is the most natural since it establishes a strict denotational identity between the antecedent VP *likes Jon's wife* and the destressed anaphor *likes her*.

The behaviour of these pronouns is simply explained once they are viewed as *paycheck pronouns* as illustrated by Karttunen's famous example (cf. [Kar69]):

(10) *The man who gave his paycheck to his wife was wiser than the man who gave it to his mistress*

Paycheck pronouns differ from other pronouns in that they can neither be seen as coreferential constants nor as bound variables – instead they pick up the definite description introduced by their antecedent and reanchor its possessive pronoun in its immediate context. For instance in (10) above, the paycheck pronoun *it* picks up the description *his paycheck* and reanchors its possessive pronoun *his* to the second occurrence of *the man*.

There are various ways in which paycheck pronouns can be accounted for but essentially, the idea is that their denotation is fixed by a definite description containing either an unresolved pronoun or an unresolved property. As [Coo79] convincingly argues, the second solution is methodologically more satisfactory. We will therefore assume that paycheck pronouns are definite NPs whose representation includes a free variable of type $(e \rightarrow t)$ i.e. a property. More specifically, we assume that a paycheck pronoun is assigned the following representation:

$$\lambda Q.\exists x[P(x) \wedge \forall y[P(y) \leftrightarrow y = x] \wedge Q(x)]$$

where $P \in wff_{(e \rightarrow t)}$. Given this, the analysis of (7) runs as follows. The equations to be resolved to check the well-formedness of the destressed anaphor *likes her* are:[11].

$$\begin{array}{ll} An(j) & = \exists x[\ wof(x,j) \wedge unique(x) \wedge love(j,x)] \\ An(p) & = \exists x[P(x) \wedge unique(x) \wedge love(p,x)] \end{array}$$

Resolution of the first equation yields the two values $\lambda y.\exists x[\ wof(x,j) \wedge unique(x) \wedge\ love(y,x)]$ and $\lambda y.\exists x[\ wof(x,y) \wedge unique(x) \wedge\ love(y,x)]$ for An and thus, the values $\exists x[\ wof(x,j) \wedge unique(x) \wedge\ love(p,x)]$, and $\exists x[\ wof(x,p) \wedge unique(x) \wedge\ love(p,x)]$ for $An(p)$. The first result yields the strict reading (Peter loves Jon's wife) whereas the second yields the sloppy reading (Peter loves Peter's wife).

5 Corrections and definiteness

So far, we have only considered cases where the semantic representation of the destressed anaphor could syntactically unify with that of its antecedent. That is, in each case it was possible to find a substitution of values for free variables which made the two semantic representations $\alpha\beta\eta$–identical. In this section, we turn to more semantic cases, cases in which the relation between destressed anaphor and source parallel element is one of denotational – rather than syntactical – identity. Definites are a primary example of such a phenomenon: since one and the same individual can be referred to by several, distinct definite descriptions, it often happens that the definite description used in the destressed part of a correction is not structurally identical with the description used in its source parallel element. This is illustrated in example (11) where the source utterance contains the definite *the woman with the red hat.* As illustrated by (11a–d), the parallel element in the correction can be *his wife, her, the neighbour's daughter* or *Sarah.* In each case, the description does not syntactically unify with the source description *the woman with the red hat.* Note however that the correction is only well–formed when the parallel descriptions are interpreted as referring to one and the same individual (cf. the ill–formedness of (11e–g)). That is, when they are semantically equivalent.

(11) *Jon*$_2$ *likes [the woman with the red hat]*$_1$

a. *No, PETER*$_3$ *likes his wife (=* NP_1*)*

b. *No, PETER likes her*$_1$*.*

c. *No, PETER likes [the neighbour's daughter]*$_1$*.*

d. *No, PETER likes Sarah*$_1$*.*

e. $\star$ *No, PETER likes her*$_4$*.*

f. $\star$ *No, PETER likes Mary*$_4$*.*

g. $\star$ *No, PETER likes him.*

[11] In what follows, we abbreviate $\lambda Q.\exists x[P(x) \wedge \forall y[P(y) \leftrightarrow y = x] \wedge Q(x)]$ to $\lambda Q.\exists x[P(x) \wedge unique(x) \wedge Q(x)]$.

How does HOUE account for such examples? To show this, we now sketch the main steps of the unification process for example (11d) with equations:

$$\begin{array}{ll} An(p) & = like(p,s) \\ An(j) & = \exists x(w(x) \wedge wrh(x) \wedge unique(x) \wedge like(j,x)) \end{array}$$

These are solved in a context, where Sarah is the only woman with a red hat. The HOUE method is given access to the hypotheses $unique(s)$, $w(s)$ and $wrh(s)$ by adding them to the initial tableau. In a first step, we solve the first equation to $An = \lambda z.like(z,s)$ and obtain the following tableau:

$$\begin{array}{c} unique(s) \\ w(s) \\ wrh(s) \\ An(p) = like(p,s) \\ \vdots \\ like(j,s) = \exists x(w(x) \wedge wrh(x) \wedge unique(x) \wedge like(j,x)) \end{array}$$

The HOUE rule discussed in section 2 now splits the initial equation into two branches. The first one has the form

$$\begin{array}{c} like(j,s) \\ \neg\exists x(w(x) \wedge \ldots \wedge like(j,x)) \\ \begin{array}{c|c|c|c} \neg w(z) & \neg wrh(z) & \neg unique(z) & \neg like(j,z) \\ *\,[z=s] & *[z=s] & *[z=s] & *[z=s] \end{array} \end{array}$$

and contains the formulae $like(j,s)$ and $(\neg\exists x(w(x) \wedge wrh(x) \wedge unique(x) \wedge like(j,x)))$. The latter is universally quantified[12] and can therefore be developed into four branches $\neg w(z)$, $\neg wrh(z)$, $\neg unique(z)$, and $\neg like(j,z)$. The first three branches can be closed using the hypotheses on Sarah and the last one with the first formula, all by binding the new variable z to s. The second branch has the form

$$\begin{array}{c} \neg like(j,s) \\ \exists x(\ldots \wedge unique(x) \wedge like(j,x)) \\ unique(c) \\ like(j,c)) \\ \vdots \\ c = s \\ like(j,s) \\ *\,[] \end{array}$$

and consists of the formulae $\neg like(j,s)$ and $\exists x(w(x) \wedge wrh(x) \wedge unique(x) \wedge like(j,x))$, which is developed into the single branch containing the conjuncts $w(c)$, $wrh(c)$, $unique(c)$, and $like(j,c))$, where c is a Skolem constant for x. Here an expansion of the definition of uniqueness

$$unique(x) \Leftrightarrow \forall z(w(z) \wedge wrh(z) \leftrightarrow x = z)$$

[12] We use that $\neg\exists x.A$ is equivalent to $\forall x.\neg A$ here.

closes the branch (if Sarah and c are unique, then $s = c$).

By now, it should be clear that our treatment will also encounter no particular problem in dealing with examples such as (12) and (13) below. The first example relies on the world-knowledge that marrying is a symmetric relation (both partners have to say "yes I do"), whereas the second relies on the fact that getting wounded is synonymous to being hurt by someone/thing. Once these equivalences are taken into account, the HOUE analysis of corrections will correctly predict that these examples are well–formed.

(12) a. *A: Jon married Sarah*

  b. *B: No, Sarah married PETER*

(13) a. *A: Sarah hurt Paul.*

  b. *B: No, PETER was wounded.*

We have seen that a deaccented anaphor must either have a semantic representation which syntactically unifies with that of its antecedent, or be semantically equivalent to this antecedent. To show that this is a necessary condition, we need to provide some ill–formed examples in which neither condition holds. Such examples are given when the correction contains a destressed pronoun whose source parallel element is either an indefinite (14) or a quantifier (15).

(14) a. *Jon eats an_1 apple.*

  b. $*$ *No, PETER eats it_1.*

(15) a. *Jon kissed $most_1$ women at the party yesterday.*

  b. $*$ *No, PETER kissed $them_1$.*

In both cases, the semantic representation of the pronoun in the correction fails to syntactically unify with the semantic representation of its antecedent. Neither can it be proved that *it* and *them* are semantically equivalent to *an apple* and *most women at the party* respectively. Therefore, unification fails correctly ruling out (14) and (15). The logical reason for this e.g. in (14), is that while the second equation $An(p) = eat(p, y)$ can be solved to $An = \lambda x.eat(x, y)$ yielding the negated $\neg(eat(j, y) \Leftrightarrow \exists x(ap(x) \wedge eat(j, x))$, this cannot be refuted [13].

[13] Example (14) is in fact ambiguous between a specific reading of the indefinite *an apple* and a non-specific one. In the first case, the indefinite denote uniquely so that *it* in (14b) refers to this unique apple. Since it is denotationally equivalent with its antecedent, HOUE will succeed. In the second case, there is no unique apple salient in the context, hence *it* and *an apple* cannot be denotationally equivalent. Therefore HOUE fails. The above discussion focuses on this second possibility.

6 Conclusion

In a sense, it would be much more natural to express the proposed analysis in a dynamic setting (cf. [Kam81]). The data discussed in section 5 clearly shows that definite, indefinites and quantifiers behave differently wrt. corrections. The intuition is that whereas, a definite can bind a pronoun in the correction (cf. example 11), indefinites and quantifiers cannot (cf. examples 14,15). These are of course precisely the sort of facts dynamic semantics was designed to deal with: if we assume that the correctee–correction pair is semantically represented by a disjunction ($\Phi \vee \Psi$), then a definite in the correctee will be able to bind an anaphor in the correction (because definites have global scope) whereas indefinites and quantifiers won't (because traditionally disjunction is static and the discourse referents introduced by one disjunct are not accessible to the other disjunct). In this paper, we've shown that such facts could be modelled by means of HOUE on static semantic representations; it would be interesting to see how the analysis would transpose to a more dynamic setting. This however must await the development of Higher–Order Unification for a dynamic lambda–calculus.

Another question worth investigating is whether the interleaving of anaphora resolution and quantification proposed in [DSP91] could account for the data considered here. The approach has the advantage that it does not resort to equivalences, thus permitting better computational properties. However, unless definites are treated in a special way, it is unlikely that the approach will be able to capture examples such as (11) where denotational equivalence, rather than strict unification, is required.

Finally, an interesting issue concerns the relationship between HOUE and accommodation. A simple way to model accomodation would be to posit that, as theorem proving hits a dead-end, accomodation can be used to close off a branch: the accomodated fact is the fact needed to derive a contradiction and close off this tableau branch. Naturally, this idea is too simplistic in that some model must be defined which constrains accomodation. This we leave as an open research issue.

References

[Bar95] Christine Bartels. Second occurrence test. Ms., 1995.

[Coo79] Robin Cooper. The interpretation of pronouns. In F. Heny and H.S. Schnelle, editors, *Syntax and Semantics*, number 10, pages 61–93. 1979.

[DSP91] Mary Dalrymple, Stuart Shieber, and Fernando Pereira. Ellipsis and higher-order-unification. *Linguistics and Philosophy*, 14:399–452, 1991.

[Fit90] Melvin Fitting. *First-Order Logic and Automated Threorem Proving*. Springer Verlag, 1990.

[GK96] Claire Gardent and Michael Kohlhase. Higher–order coloured unification and natural language semantics. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*. Santa Cruz, USA, 1996.

[Hue75] Gérard P. Huet. A unification algorithm for typed λ-calculus. *Theoretical Computer Science*, 1:27–57, 1975.

[Kam81] Hans Kamp. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Formal Methods in the study of language*. Amsterdam, 1981.

[Kar69] Lauri Karttunen. Pronouns and variables. In R. Binnick et al., editor, *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society*, University of Chicago, USA, 1969.

[Koh95] Michael Kohlhase. Higher-Order Tableaux. In R. Hähnle P. Baumgartner and J. Posegga, editors, *Theorem Proving with Analytic Tableaux and Related Methods*, volume 918 of *Lecture Notes in Artificial Intelligence*, pages 294–309, 1995.

[Nor91] Neal R. Norrick. On the organization of corrective exchanges in conversation. *Journal of Pragmatics*, (16):59–83, 1991.

[RB82] Martin H. Ringle and Bertram C. Bruce. Conversation failure. In W.G. Lehnert and M.H. Ringle, editors, *Strategies for Natural Language Processing*, chapter 7, pages 203–221. Lawrence Erlbaum Associates, Inc., Hillsdale, New Yersey, 1982.

[Ros67] J. Ross. *Constraints on Variables in Syntax*. PhD thesis, MIT, Cambridge, Massachusetts, 1967.

[Sag76] Ivan Sag. *Deletion and Logical Form*. PhD thesis, MIT, 1976.

[SJS77] Emanuel Schegloff, Gail Jefferson, and Harvey Sacks. The preference for self correction in the organization of repair in conversation. *Language*, (53):361–382, 1977.

[Sny91] Wayne Snyder. *A Proof Theory for General Unification*. Progress in Computer Science and Applied Logic. Birkhäuser, 1991.

[vL94] Noor van Leusen. The interpretation of corrections. In Peter Bosch and Rob van der Sandt, editors, *Focus and Natural Language Processing*, number 7 in IBM Working Papers of the Institute for Logic and Linguistics. IBM Deutschland GmbH, December 1994. IBM TR-80.94-007.