Learning to reason about speakers' alternatives in sentence comprehension: A computational account

Petra Hendriks [a,1], Hedderik van Rijn [b], Bea Valkenier [b]

[a] Center for Language and Cognition Groningen, University of Groningen, PO Box 716, 9700 AS Groningen, The Netherlands

[b] Department of Artificial Intelligence, University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands

[1] Corresponding author. Tel.: +31 50 363 5863; fax: +31 50 363 6855.
E-mail addresses: P.Hendriks@rug.nl (P.Hendriks), D.H.van.Rijn@rug.nl (H. van Rijn), bea@ai.rug.nl (B. Valkenier).

Abstract

We present a computational simulation study of the acquisition of pronouns and reflexives. The computational simulation is based on an Optimality Theory analysis, and is shown to account for the well-known observation that in English and many other languages the correct comprehension of pronouns lags behind that of reflexives (the so-called Delay of Principle B Effect). Comprehension is modelled as a two-step process involving optimization from a given form to its corresponding meaning followed by optimization from this meaning to its corresponding form. This model is implemented using plausible assumptions with respect to the cognitive architecture. The computational simulation shows that lack of processing speed causes the model to produce an output before both steps of the comprehension process have been completed. Because, according to the Optimality Theory analysis, the adult interpretation of pronouns is dependent on reasoning about alternative forms and hence on completion of both steps of the comprehension process, whereas the interpretation of reflexives is not, this results in comprehension errors with pronouns but not with reflexives. We conclude that speed of processing may be an essential factor in explaining the Delay of Principle B Effect and other comprehension delays in language acquisition.

Keywords: Delay of Principle B Effect, language acquisition, Optimality Theory, bi-directional optimization, cognitive modelling, computational simulation

1. Introduction

The acquisition of discourse-semantic aspects of language is characterized by several well- and less well-documented delays in comprehension. Recently, a number of studies have argued that these delays can be accounted for by assuming that children are only able to consider their own perspective, whereas adult hearers are able to simultaneously take into account the perspective of the speaker (Hendriks and Spenader, 2004, in press; de Hoop and Krämer, 2006). Thus, if the child's grammar is modelled as a process of unidirectional optimization, the adult grammar should be modelled as a process of bi-directional optimization. In this paper, we develop a computational cognitive model of the transition from unidirectional to bi-directional optimization, applied to the well-known Delay of Principle B Effect. Using this cognitive model, we can shed more light on the interaction between linguistic knowledge, rule application and memory processes. As we show, it follows that children will start to optimize bi-directionally once the unidirectional process can be performed fast enough. Thus, no qualitative difference needs to be assumed between children and adults with respect to their knowledge of the grammar. Moreover, the cognitive model yields several testable predictions with respect to the course of language acquisition in individual children as well as with respect to the speed of acquisition for different linguistic forms.

2. Delays in comprehension

One of the most well-documented delays in comprehension is the delay occurring with respect to children's acquisition of the binding principles A and B. Children interpret reflexives (governed by Principle A) in an adult-like manner from the age of 3, but they continue to perform poorly on the interpretation of pronouns (governed by Principle B) even up to the age of 6;6 (Chien and Wexler, 1990; see also Grimshaw and Rosen, 1990 for a review). For example, presented in a context with two male referents, say Bert and Ernie, sentences like (1) are correctly understood from a young age (95% of the time according to some studies). However, children misinterpret *him* in (2) as co-referring with the subject about half the time, which seems to be the result of chance performance.

(1) Bert washed himself.

(2) Bert washed him.

This pattern is generally referred to as the Delay of Principle B Effect, and has been the topic of investigation in many acquisition studies over the past two decades. The pattern is highly unexpected, especially under a nativist view on language according to which the binding principles A and B are innate and therefore should emerge simultaneously. Why then would children master one of these principles at least three years later than the other one?

Children's production data complicate the picture even more. Acquisition studies of children's language production suggest that children do not have problems in producing reflexives or pronouns correctly. De Villiers, Cahillane and Altreuter (in press) studied the production as well as the comprehension of third person reflexives and pronouns in 68 English speaking children between the ages of 4;6 and 7;2. In their study, production was significantly better than comprehension for all sentence types studied. Children showed superior performance in producing pronouns correctly and almost never produced a reflexive when a pronoun was the target. These results are supported by Bloom, Barss, Nicol and Conway's (1994) study of naturalistic data, which looked at the spontaneous production of the English first person forms *me* and *myself* in the CHILDES database.[2]

To summarize, 5- and 6-year-old children produce pronouns in an adult-like manner, but these same children seem to resort to guessing when interpreting the same forms. Similar delays in comprehension have been observed in other areas of language, such as with respect to the comprehension of sentence stress (Cutler and Swinney, 1987) the interpretation of indefinite subjects and objects (de Hoop and Krämer, 2006) and the use of scalar implicatures (Noveck, 2001; Papafragou and Musolino, 2003).

[2] As one of the reviewers remarked, *me* and *myself* are special because they do not give rise to ambiguity. In fact, this was why Bloom et al. (1994) were able to determine from a corpus of spoken forms whether children produced pronouns and reflexives correctly. A similar study could not have been carried out with *him* and *himself*. The form "Bert washed him" doesn't provide any information about whether *him* is used correctly here because *him* can in principle refer to any male referent present in the discourse, including the subject *Bert*. Only when the interpretation is taken into account is it possible to tell whether the speaker indeed obeyed Principle B by not using the pronoun to refer to the subject.

How can it be explained that children are able to correctly produce forms which they are not yet able to correctly understand? In the next section, we will discuss an explanation for this phenomenon based on the assumption that children are only able to consider their own perspective, whereas adult hearers simultaneously take into account the perspective of the speaker. Only when children have learned to take into account their conversational partner's perspective will they interpret pronouns in an adult-like manner. But how do child hearers learn to consider the speaker's perspective in comprehension? In section 4, we develop a computational cognitive model which simulates this process. The results of the computational simulation study are presented in section 5. On the basis of the correspondence between the simulation data and the actual acquisition data we can make several predictions, which are discussed in section 6. Section 7, finally, summarizes the conclusions of this simulation study.

3. Unidirectional vs. bi-directional optimization

The Delay of Principle B Effect has been attributed to extra-grammatical factors, such as problems with real-world knowledge (the 'pragmatic account'; e.g. Thornton and Wexler, 1999) or lack of processing resources (the 'processing account'; e.g. Reinhart, 2004, in press). In contrast, Hendriks and Spenader (2004, in press) propose that the Delay of Principle B Effect should be accounted for within the grammar (in other words, they propose a 'grammatical account'). Their grammatical account not only correctly predicts the observed pattern in the production and comprehension of pronouns in young children, but at the same time yields a number of other predictions for which there is some suggestive evidence (see Hendriks and Spenader (in press) for a discussion of the differences between their account and the pragmatic and processing accounts). However, a grammatical account of the Delay of Principle B Effect and other acquisition delays is feasible only if the grammar itself is asymmetrical. In section 3.1, we discuss such a grammar. In section 3.2, it is argued that bi-directional reasoning should be an essential part of such an asymmetrical grammar.

3.1 An asymmetrical grammar

A linguistic framework which is inherently asymmetrical is Optimality Theory (Prince and Smolensky, 2004). In Optimality Theory (henceforth OT), production can be modelled as optimization from meaning to form on the basis of a set of ranked constraints. Comprehension can be modelled as optimization in the other direction, that is, from form to meaning (Hendriks and de Hoop, 2001). OT distinguishes between constraints punishing ‘marked’ output forms (so-called markedness constraints), and constraints punishing dissimilarity between input and output (so-called faithfulness constraints). Because the output becomes the input when the direction of optimization changes, markedness constraints (which are only concerned with distinguishing between potential outputs) only have an effect in one direction of optimization, and not in the opposite direction. As a result, production may yield different optimal form-meaning combinations than comprehension. This property of OT, discussed in Smolensky (1996), is illustrated by means of the following two abstract OT tableaux, which are based on an example presented by Smolensky:

Input: m	*f	*<f,m’> *<f’,m>
<f,m>	*!	
☞ <f’,m>		*

Tableau 1: Production under non-adult constraint ranking

Input: f	*f	*<f,m’> *<f’,m>
☞ <f,m>	*	
<f,m’>	*	*!

Tableau 2: Comprehension under non-adult constraint ranking

In optimization from meaning to form and from form to meaning, the input is fixed. In production tableau 1, the input is a given meaning m. In comprehension tableau 2, the input is a given form f. What competes are the different output candidates for these inputs, which are listed in the first column of each tableau. That is, in tableau 1 the forms f and f’ compete for expressing meaning m. In tableau 2, the meanings m and m’ compete for being expressed by input form f. We assume that the same grammar (i.e., the same set of constraints under the same ranking) is used for production and comprehension. The constraints are listed in order of descending strength from left to right across the first row of each tableau. So *f is stronger than *<f,m’> and *<f’,m>, which means that it is more important to satisfy *f than it is to satisfy the other two constraints. An asterisk indicates a constraint violation. The pair <f,m> violates the constraint *f (read: 'avoid f'), which punishes all pairs containing form f. This violation is fatal in the first tableau: because of this violation another candidate is optimal, namely the pair <f',m>. This is indicated by the exclamation mark. The optimal candidate (the output for the given input) is indicated by the pointing hand.

As a comparison of production tableau 1 and comprehension tableau 2 shows, changing the direction of optimization can have dramatic effects on what wins. The optimal pair in production tableau 1 is <f',m>, whereas in comprehension tableau 2 it is <f,m>. Crucial is the behaviour of the markedness constraint *f, which causes the pair <f,m> to be suboptimal in tableau 1, but does not distinguish between candidates in tableau 2. In the latter case, the form f is already given in the input, so all candidates contain this form f and hence violate *f to the same degree.

If grammars are rankings of universal constraints, as is commonly assumed in OT, then acquiring a grammar must involve learning the adult constraint ranking (Boersma and Hayes, 2001; Tesar and Smolensky, 1998). In children who still entertain a non-adult constraint ranking, with one or more markedness constraints being ranked too high, the inherent asymmetry of the grammar may give rise to errors in production but at the same time result in adult-like performance in comprehension (Smolensky, 1996). For example, assume *f to represent markedness constraints such as NoCoda ('syllables do not have codas') and *Dors ('segments do not have the feature [dorsal]', which punishes the pronunciation of speech segments such as [k]), and *<f,m'> and *<f',m> to represent the faithfulness constraints Parse and Fill. Potential surface forms are f: [kæt] and f': [ta], and potential underlying forms are m: /kæt/ and m': /skæti/. As Smolensky (1996) shows, in such an OT model optimization in production yields <[ta], /kæt/> as the optimal pair, resulting in errors in production. That is, the child will say *ta* when referring to a cat. Optimization in comprehension yields <[kæt], /kæt/> as the optimal pair, combining the adult surface form with the adult underlying form. So the child will interpret the word *cat* correctly. Thus, the inherent asymmetry of OT may explain why the comprehension of early word forms precedes their production.

If the child re-ranks the constraints in such a way that the markedness constraints are dominated by the faithfulness constraints, the adult pattern of production and comprehension emerges. Production will now yield the pair <[kæt],

/kæt/> too, because under the adult ranking it is more important to satisfy the faithfulness constraints *<f,m’> and *<f’,m> than it is to satisfy the markedness constraints represented by *f. Several learning algorithms have been proposed within OT to drive such constraint re-ranking on the basis of positive evidence only (Boersma and Hayes, 2001; Tesar and Smolensky, 1998). These algorithms proceed from the view that, by hearing the surface form [kæt] and comparing this form to the output of the child’s current grammar, [ta], the child will be able to conclude that the constraint ranking is incorrect and re-rank the constraints accordingly. This process continues until no mismatches are detected anymore between the heard forms and the forms produced by the grammar. Thus the adult constraint ranking is obtained.

However, an adult constraint ranking does not always result in a symmetrical pattern. Again, this is illustrated by means of two abstract OT tableaux to emphasize the similarity between this example and the previous one:

	Input: m	*<f’,m>	*f
☞	<f,m>		*
	<f’,m>	*!	

Tableau 3: Production under adult constraint ranking

	Input: f	*<f’,m>	*f
☞	<f,m>		*
☞	<f,m’>		*

Tableau 4: Comprehension under adult constraint ranking

Because neither constraint *<f’,m> nor constraint *f distinguishes between the candidates in tableau 2, both pairs are optimal in comprehension. In other words, form f is ambiguous between meaning m and meaning m’.

Now assume *<f’,m> to represent the violable constraint Principle A (‘reflexives do not have a disjoint interpretation’), and *f to represent the markedness constraint Referential Economy (‘avoid pronouns’). Note that in this example, m and m’ refer to actual interpretations (a disjoint and a co-referential interpretation, respectively) rather than underlying forms, as in the previous example. The forms f and f’ represent the overt form of a pronoun and a reflexive, respectively. The markedness constraint adopted here is comparable to the markedness constraints used in the previous example. However, the constraint Principle A does not relate an underlying form to an overt form, but rather a form (a reflexive) to a meaning (a co-referential interpretation). For this reason, the term ‘faithfulness constraint’ may not be entirely suitable for such a constraint.

As tableau 3 shows, a pronoun f is the optimal form for a disjoint input meaning m. So if we wish to describe a situation in which Bert washed Ernie, with the two arguments of the verb *washed* being disjoint, we should use a pronoun: “Bert washed him”. In comprehension, however, a pronoun may receive both a disjoint and a co-referential interpretation, as is shown in tableau 4. In other words, according to tableau 4 a pronoun is ambiguous. Tableaux 3 and 4 thus illustrate the pattern discussed in section 2 that children show adult-like performance in producing pronouns, but seem to resort to guessing when comprehending pronouns, selecting the adult meaning half of the time and the non-adult meaning the other half of the time.

In contrast to Smolensky’s example illustrated by tableaux 1 and 2, however, under the assumption of a total ranking of the constraints the standard OT mechanism of constraint re-ranking is not able to explain the acquisition of the adult pattern of pronoun interpretation. This is true even if we consider other constraints than the ones used here, which are taken from Hendriks and Spenader (2004, in press). This can easily be seen by comparing the constraint systems required to model children’s and adults’ patterns of interpretation. If pronouns are ambiguous for children, then the disjoint interpretation and the co-referential interpretation must have the same constraint profile, i.e., they must satisfy and violate the same constraints. Obviously, re-ranking these constraints will never result in one of the interpretations being preferred to the other, which is necessary to model the adult pattern. Crucially, no combination of constraints will allow for an OT explanation of the acquisition pattern in terms of constraint re-ranking.

Weakening the OT model to allow for tied constraints may result in a constraint system that accounts for children’s as well as adults’ pattern, but such a model still does not provide any explanation for why it takes children so much time (until at least the age of 7) to learn the adult constraint ranking. Most other grammatical knowledge is acquired well before the age of 4 or 5. This suggests that learning to interpret pronouns in an adult-like way involves something more than constraint re-ranking.

For this reason we hypothesize that the constraints in tableaux 3 and 4 accurately model children’s linguistic competence. Re-ranking these constraints will make things worse. Placing the markedness constraint above the faithfulness constraint will still result in incorrect comprehension. But, in addition, it will lead to incorrect production, with reflexives being used for disjoint interpretations. So if the constraints in tableaux 3 and 4 accurately model children's linguistic competence, the ranking in these tableaux is the best ranking possible. Nevertheless, this ranking does not give rise to the symmetrical adult pattern. According to this ranking pronouns are predicted to be ambiguous, but for adults pronouns are not.

3.2 Reasoning about the speaker's alternatives

Assuming that the OT analyses in tableaux 3 and 4 accurately model children's production and comprehension of pronouns and reflexives, why are pronouns not ambiguous for adults? Hendriks and Spenader (2004, in press; see also de Hoop and Krämer, 2006) argue that adults not only optimize in the intended direction (i.e., from form to meaning, or from meaning to form), but also consider the results of optimization in the opposite direction. In other words, they not only take into account their own alternatives but also the alternatives available to their conversational partner. This type of bi-directional optimization is formalized in the definition in (3).

(3) Bi-directional optimization (adapted from Blutner, 2000):
A form-meaning pair <f,m> is bi-directionally optimal iff:
a. there is no optimal pair <f',m> such that <f',m> is more harmonic than <f,m>.
b. there is no optimal pair <f,m'> such that <f,m'> is more harmonic than <f,m>.

The term "harmonic" refers to the notion of harmony, which is taken from neural network theory and is a numerical measure of how well the pairs conform to the constraints of the grammar (Smolensky, 1986). Bi-directional optimization has the effect that if a form or a meaning already is part of an optimal form-meaning pair, its use is blocked for another form-meaning pair. For example, if pronouns are ambiguous between a co-referential and a disjoint meaning, but the co-referential meaning already is part of the optimal form-meaning pair <reflexive, co-referential meaning>, the co-referential meaning is blocked for the pronoun. This explains why pronouns are not ambiguous for adults. Thus, a mature language user, when hearing a pronoun, implicitly reasons that if the speaker would have wanted to express a co-referential meaning, the speaker would have used a reflexive. Since the speaker did not use a reflexive, a mature hearer will conclude that the speaker did not mean to express a co-referential meaning. Consequently, the pronoun is interpreted as expressing a disjoint meaning.[3]

Although the grammar is asymmetrical, the resulting adult system of form-meaning pairs is symmetrical because of the additional step of bi-directional optimization. Similar explanations in terms of bi-directional optimization have been proposed for the adult comprehension of indefinite subjects and objects (de Hoop and Krämer, 2006), sentence stress with free focus (Aloni et al., in press) and sentence stress with bound focus (Hendriks et al., 2005). The symmetrical adult pattern in these studies reflects the observation that whatever language users can produce they are generally able to understand, and vice versa.

If adult's optimization strategy differs from children's, an important question is how children arrive at the adult strategy of bi-directional optimization. In the next section, we develop a computational model based on plausible assumptions with respect to human cognitive processing which provides a possible explanation for the transition from unidirectional to bi-directional optimization. Note that the computational model we develop is not intended as an alternative to the OT approach to language acquisition, but must rather be seen as a necessary embedding of the linguistic system in a cognitively motivated architecture. The first step in language acquisition, learning the adult constraint ranking, can be modelled using linguistic learning algorithms such as proposed by Boersma and Hayes (2001) and Tesar and Smolensky (1998). The second step in language acquisition, learning to take into account the opposite perspective, however, goes beyond the power of these linguistic algorithms and seems to require a more general cognitive developmental approach.

[3] The version of bidirectional optimization defined in (3) is non-recursive and is known as 'strong bidirectional optimization'. In addition, a recursive version has been proposed (see section 6 for discussion), which is termed 'weak bidirectional optimization'. This recursive version seems extremely suited to account for conversational implicatures (Blutner, 2006). However, as the analysis of pronouns presented here shows, the non-recursive version in (3) is already sufficient for explaining phenomena such as pronoun interpretation that are not traditionally seen as involving implicatures and are not defeasible.

4. An ACT-R model of optimization in language

ACT-R (Anderson and Lebiere, 1998; Anderson et al., 2005) is a cognitive architecture designed for computationally simulating and understanding human cognition. The computational simulation models are constrained by the architecture of ACT-R in the way they retrieve, store and process information. These architectural constraints of ACT-R are derived from numerous experiments on human cognition. ACT-R is a hybrid architecture, which means that it operates at a symbolic as well as a sub-symbolic level. At the symbolic level, two kinds of memory can be distinguished. Declarative memory contains chunks of information representing facts. Procedural memory contains production rules (IF-THEN rules) representing actions. Production rules compete with each other at the sub-symbolic level. When more than one production rule can be applied, there will be competition among these rules. The production rule with the highest expected utility, which can be seen as a measure weighing costs and benefits, will be executed. Retrieval of chunks from declarative memory also is dependent on properties at the sub-symbolic level, such as recency and frequency of usage.

4.1 Bi-directional optimization as a serial process

By running computational simulations which adhere to the architectural constraints of ACT-R, more insights can be obtained with respect to how humans perform certain cognitive tasks. To this end, however, several decisions have to be made as to how to implement the cognitive task in ACT-R. Because bi-directional optimization is dependent on unidirectional optimization, it may be useful to describe bi-directional optimization in terms of two unidirectional processes. In particular, we will assume bi-directional optimization to involve the serial application of two unidirectional processes of optimization. In the case of comprehension, this amounts to a first step of optimization from form to meaning, followed by a second optimization step from meaning to form. In production, which we do not discuss here, the directions of optimization in the two steps are exactly the other way around. So comprehension and production are each other's mirror images.

In the second optimization step in comprehension, the meaning which was found to be optimal in the first step serves as the input. If the resulting output form of the second optimization step equals the original input form of the first optimization step, the unidirectionally optimal meaning is bi-directionally optimal as well. If not, and if there was another optimal meaning in the first step which could also have been chosen, the model will start over with this other optimal meaning. Thus, in the first step a possible output is determined. In the second step, it is checked whether this output is part of a coherent form-meaning pair (i.e., whether the input-output mapping is symmetrical).

A crucial property of ACT-R is the assumption that actions take time to perform, and that for most of central cognition, performance is limited by the serial processing bottleneck. This means that if unidirectional optimization takes a certain amount of time, bi-directional optimization will need about twice this amount of time. Given that speakers need to adhere to conversational rules that state that speech needs to be reasonably fluent, the amount of time available for constructing a form from a meaning is limited. Because ACT-R limits the totally available processing time, it is expected that bi-directional optimization needs higher processing efficiency to be performed than unidirectional optimization does. Higher processing efficiency can be gained by means of learning through production compilation (Taatgen and Anderson, 2002). In production compilation, two existing production rules are integrated into one new production rule. Because fewer production rules are needed, the result is more automatic processing that requires less processing time. Production compilation occurs when two existing production rules are repetitively executed in sequence. This process also removes any dependencies on retrieval from memory, making it possible that multiple constraints are evaluated in parallel (cf. Misker and Anderson, 2003).

Under these assumptions, the ACT-R model will perform unidirectional optimization as long as processing efficiency is not high enough to perform bi-directional optimization within the given amount of time. Thus, we assume that, potentially, children have the ability to optimize bi-directionally and strive to optimize bi-directionally already from the start. However, they do not manage to do so because they lack processing efficiency. Bi-directional optimization is performed as soon as unidirectional optimization from form to meaning and unidirectional optimization from meaning to form are performed fast enough, so that in the total time available to the speaker, both processes can be completed successfully. Failure to optimize bi-directionally results in optimizing unidirectionally.

4.2 ACT-R and the comprehension of pronouns and reflexives

In this subsection we will show how the comprehension of pronouns and reflexives is modelled within ACT-R. Figure 1 shows the high level design of the ACT-R model of optimization from form to meaning (see Valkenier, 2006, for details).

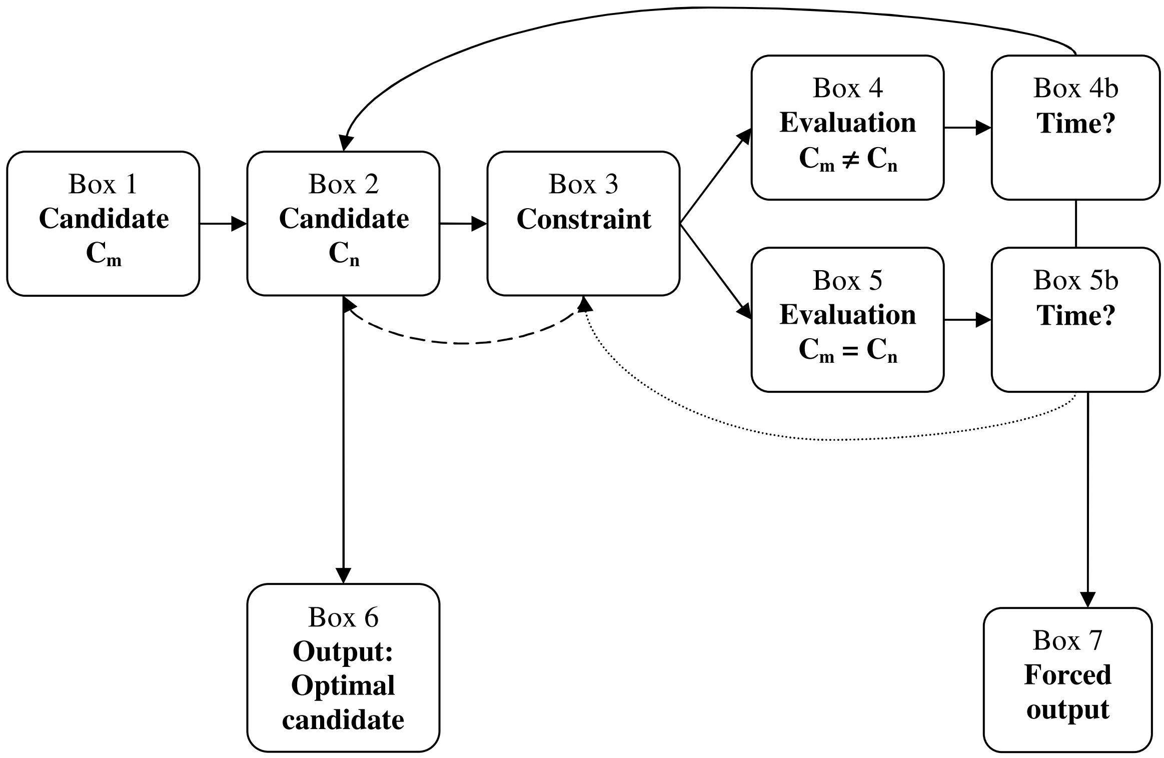

Figure 1: Structure of the ACT-R model.

The input of the model is either a pronoun or a reflexive. The first step in the model is the retrieval of one candidate meaning (box 1), another candidate meaning (box 2), and a constraint (box 3). In the current model, candidates are retrieved randomly. Both retrieved candidates are evaluated on the basis of the first retrieved constraint. If only one of the candidates violates this first constraint (box 4), and if there still is processing time left (box 4b), a new candidate will be retrieved which replaces the candidate violating the first constraint (solid arrow from box 4b to box 2). The candidate which did not violate the first constraint will then be compared to the newly retrieved candidate (cf. Misker and Anderson, 2003, for an alternative implementation of OT in ACT-R). If either both or none of the candidates violate the first constraint (box 5), and there still is processing time left (box 5b), the next constraint will be retrieved and the same two candidates will be evaluated against this new constraint (dotted arrow). The dashed arrow reflects the case in which the candidates have been evaluated with respect to all constraints but nevertheless no single optimal candidate emerges. In this case, one of the two candidates is chosen at random and a next candidate will be retrieved. These two processes of retrieving constraints and retrieving candidates will be repeated until there are no candidates left that have not been evaluated (box 6). The optimal candidate at this point is returned as the output. This process of optimization will be terminated earlier if there is no processing time left (box 7).

Each box in the figure represents one or more ACT-R production rules. The OT process of candidate generation is reflected by the boxes 1 and 2. The OT process of candidate evaluation is reflected by the boxes 3, 4 and 5. Candidates and constraints are implemented as chunks in declarative memory, in contrast to Misker and Anderson (2003), who implemented the processes of candidate generation and evaluation outside the scope of ACT-R. In our model the number of candidates that can be evaluated is limited by the amount of time available. Candidates in memory are activated one by one. If there would not be any limitations on processing time, this process could go on until all candidates have been retrieved from memory. In practice, however, there are severe limitations on processing time. Although this approach seems incompatible with the model presented by Misker and Anderson (2003), in which all constraints were implemented as production rules, the declarative constraints are compiled into production rules by means of production compilation. After compilation, different production rules exist for different constraints, making the resulting model similar to Misker and Anderson's model.

Production compilation may occur at several stages in the model. For example, in box 1 a production rule retrieves a candidate from declarative memory, and in box 2 another production rule retrieves another candidate from declarative memory. These two production rules can be integrated into a single production rule retrieving two different candidates from memory in one step. This new production rule will be faster than the two old production rules together. This process can in principle be repeated until most of the optimization process is integrated in one production rule. In the next section we discuss the results of our simulations on the basis of this model.

5. Learning to optimize bi-directionally

The model described in section 4 is implemented in the ACT-R computational architecture. Computer simulations based on the model are run many times. Each model run simulates the comprehension of one pronoun or reflexive. The model archives whether the interpretation is correct or not (figure 2)[4], and whether this interpretation was achieved through unidirectional or bi-directional optimization (figure 3). The data simulate the learning process of one child. The proportions presented in figures 2 and 3 are computed by taking a running average over 30 presentations of a pronoun or reflexive. A presentation is one instance of a pronoun or reflexive as the input to optimization.

In figure 2, the proportions of correct interpretation of a pronoun or reflexive is plotted as a function of time.

[4] In our ACT-R model, we simply stipulate that the correct interpretation of a pronoun is a disjoint interpretation and the correct interpretation of a reflexive is a co-referential interpretation. Of course, as one of the reviewers noted, this is an oversimplification. However, no explicit negative evidence seems to be required for successful acquisition. To drive constraint re-ranking, robust interpretive parsing offers the implicit feedback necessary for production (Tesar and Smolensky, 1998). In the example we focus on in this study, constraint re-ranking already led to the adult constraint ranking. Although the adult constraint ranking results in a reduction of the error rate, still no symmetrical adult pattern of forms and meanings has emerged (see section 3.1). At this stage, the only way to minimize the error rate is to optimize bi-directionally and block suboptimal form-meaning pairs. A mismatch between the input form in comprehension and the output form in production will then indicate an incorrect interpretation, and a correct match a correct interpretation. So although Tesar and Smolensky's (1998) combination of production optimization and robust interpretive parsing and Blutner's (2000) bidirectionality are different types of bidirectional optimization having different effects on language acquisition (constraint re-ranking and blocking, respectively), they both allow for acquisition on the basis of positive evidence only.

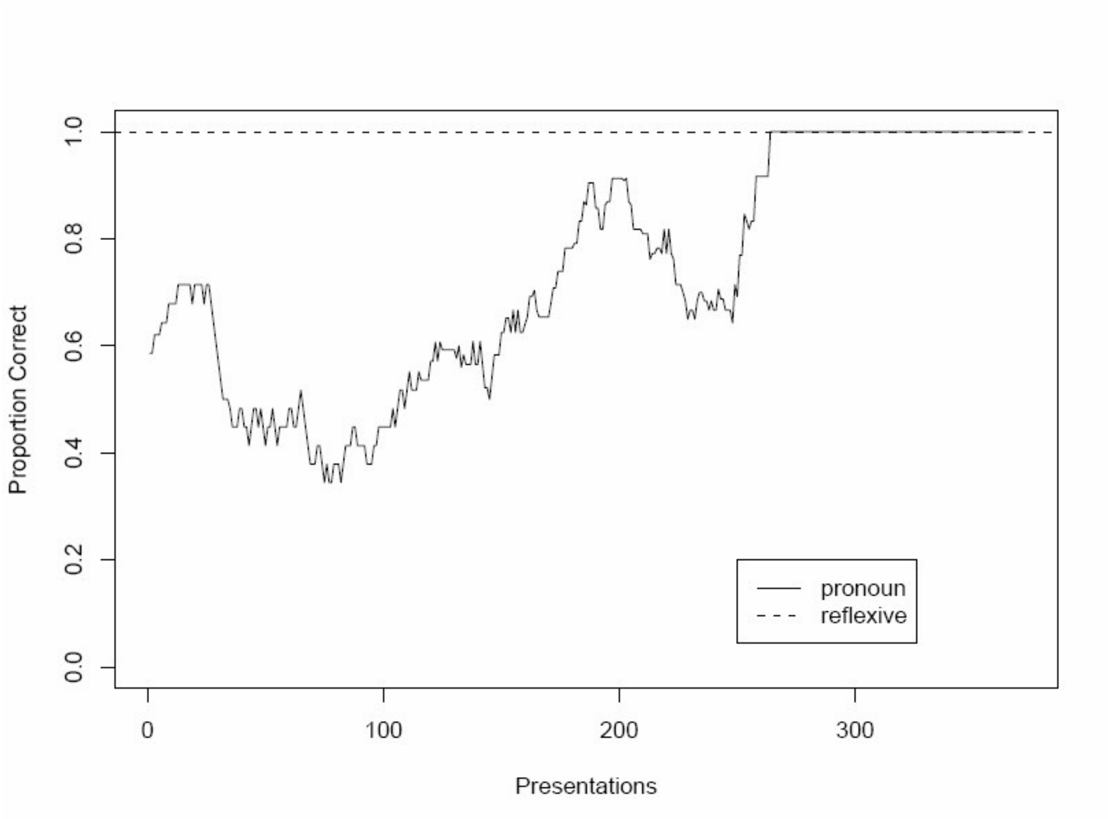

Figure 2: Performance of the ACT-R model on the comprehension of pronouns and reflexives.

The interpretation of reflexives is 100% correct already from the beginning (dashed line). This is not surprising: both unidirectional and bi-directional optimization lead to a correct interpretation of reflexives. For pronouns, the proportion of correct interpretations hovers around 50% during the first half of the learning period (solid line). In the second half of the learning period, the proportion of correct interpretations increases to about 80%. As it is in general impossible to link development in a computation model on a one-to-one timescale to development in children, the results should be seen as indicative of the relative speed of acquisition.

Figure 3 shows the proportion of bi-directional optimization compared to unidirectional optimization as a function of time.

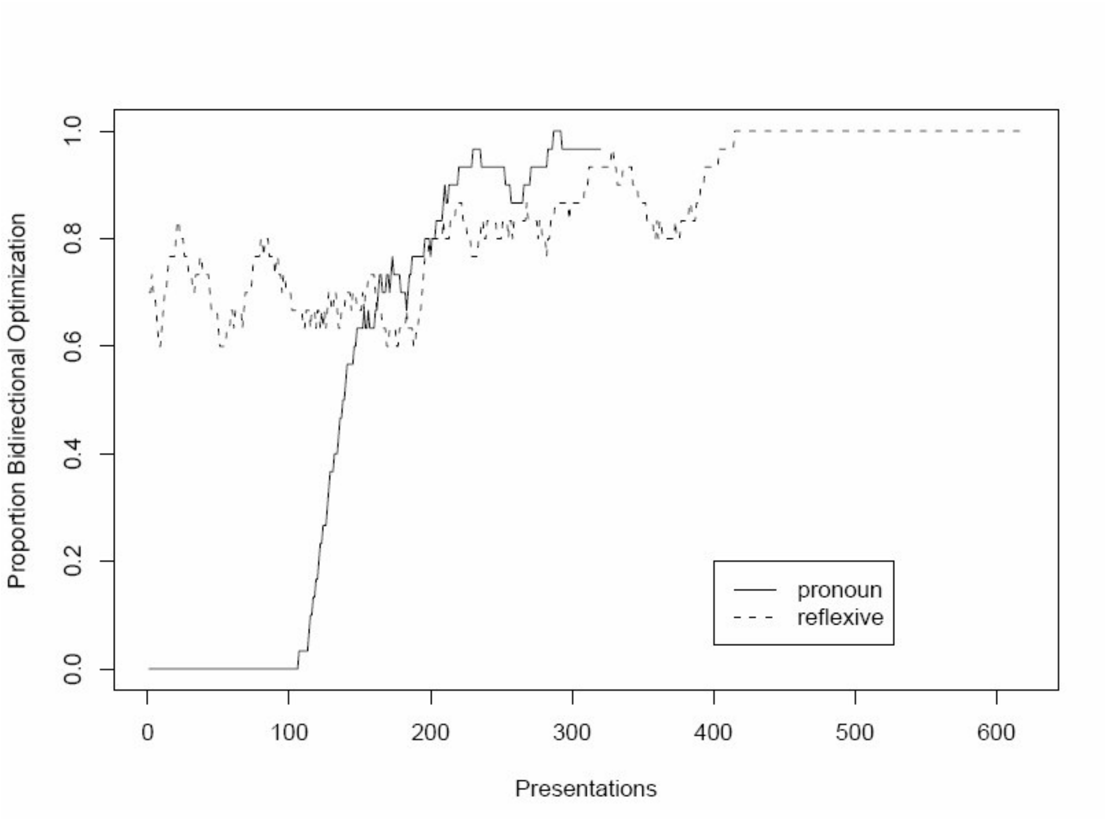

Figure 3: Percentage of bi-directional rather than unidirectional optimization by the ACT-R model.

This proportion is different for pronouns and reflexives. As a general tendency, the use of bi-directional optimization increases with time and reaches a 100% result in the second half of the learning period. This means that the process of bi-directional optimization is mastered before the end of the learning period.

With respect to the development of bi-directional optimization for pronoun comprehension (solid line), three periods can be distinguished. During the first period, only unidirectional optimization is used. The proportion of bi-directional optimization is 0. The intermediate period shows a steady upwards-sloping line, which suggests that bi-directional optimization is mastered gradually. During the last period, only bi-directional optimization is used, resulting in a 100% score. The proportion of bi-directional optimization for reflexives (dashed line) develops differently. For reflexive comprehension, bi-directional optimization is used in a significant proportion of cases already from the start (see section 6, prediction 2, for an explanation of this difference between pronouns and reflexives). The intermediate period is less steady than with pronouns, and it takes more presentations to reach the 100% level.

Note that the data in figures 2 and 3 represent the learning data of one simulated child. Level of performance and learning slope are particular to this simulated child. These properties of the data are the result of particular parameters in the model such as the amount of received data and the adaptability of the system; different parameter settings would yield different "simulated children". Other properties of the data are general results of the model. In the next section, we discuss these general results and their implications for language acquisition.

6. Discussion

Our cognitive model interprets reflexives correctly already from the start. This corresponds to the data from language acquisition experiments discussed in section 2. With respect to pronouns, our cognitive model shows around 50% performance in the first half of the learning process, with correct interpretations increasing to about 80% in the second half of the learning process. Again, this corresponds to the language acquisition data discussed earlier. Similarly to young children, in the beginning our model selects the adult meaning half of the time and the non-adult meaning the other half of the time. From the general correspondence between the language acquisition data described in the literature and the simulation data of our model, we may conclude that the cognitive model is a plausible model of the acquisition of bi-directional optimization in language.

Our cognitive model is based on the assumption that bi-directional optimization involves the serial application of two unidirectional processes of optimization. Once the unidirectional processes can be performed fast enough, bi-directional optimization is possible. Because of the similarities between the learning curves of children and our cognitive model, children's limited speed of linguistic processing may yield an explanation for the comprehension problems they experience with pronouns. To be able to interpret pronouns in an adult-like manner, children must optimize bi-directionally and consider the other forms the speaker could have used. In particular, the child hearer must conclude that a co-referential interpretation is blocked for the pronoun, because if the speaker would have wanted to express this meaning, she would have used a reflexive instead. However, bi-directional optimization requires sufficient processing time. If children's processing is still too slow, they are predicted not yet to be able to optimize bi-directionally. As a consequence, pronouns are still ambiguous for children. With experience, their unidirectional processing speeds up until, during the course of language acquisition, it is fast enough to allow for bi-directional optimization. From this point on, children will be able to consider the alternative form of a reflexive, which the speaker could also have used, block the co-referential meaning for the pronoun and consequently assign an adult interpretation to the pronoun.

This answers our central question of how child hearers learn to consider the speaker's perspective in comprehension. In addition, our cognitive model yields several predictions that we will discuss one by one:

1. Gradual increase of correct responses for pronoun comprehension.

As can be seen in figure 2, the proportion of correct responses for pronouns increases gradually. This prediction derives from the property of the model that more production rules are used in the beginning of the learning process than at the end. The time needed for each application of a production rule is slightly variable due to stochastic aspects at the sub-symbolic level of the model. Because of the larger number of production rules in the beginning in combination with the variance in time associated with each application of a production rule, there is substantial variation in the total amount of time needed for comprehension in the beginning of the learning process. As a result, in some cases (namely in those cases in which the application of the production rules happens to take relatively little time) bi-directional optimization will already be possible. In most cases, however, it will not. At this early stage in the learning process, therefore, there will be a relatively large period of time in which children alternate between unidirectional and bi-directional optimization.

During the course of learning, fewer production rules are used because of the mechanism of production compilation (see section 4.2). This has two consequences. First of all, the speed of processing increases because fewer production rules have to apply. Second, because fewer production rules apply, the variance in processing time decreases. As a result of this, there is a smaller period of time in which children alternate between unidirectional and bi-directional optimization. If the processing speed is large enough, bi-directional optimization will be the only strategy that is used because the available time is sufficient for optimizing in both directions in all cases. So the model predicts a gradual increase of bi-directional (= correct) responses for pronoun comprehension for each individual child. Whether this prediction is correct can only be determined through a longitudinal study of children's comprehension of pronouns. As far as we know, no such study has yet been done.

2. Early use of bi-directional optimization in reflexive comprehension.

Figure 3 shows that the comprehension of reflexives already involves bi-directional optimization from an early point on. This contrasts with the use of bi-directional optimization for the comprehension of pronouns. This prediction can also be explained from the properties of the model. For the interpretation of reflexives, fewer production rules are needed than for the interpretation of pronouns. Because fewer production rules are needed, less time is needed for the comprehension of reflexives and chances are higher that bi-directional optimization is possible.[5]

The use of fewer production rules for reflexives than for pronouns is a side effect of the way we implemented the OT constraint hierarchy in ACT-R. In OT, it is assumed that all constraints apply simultaneously. This, however, is not how we implemented the process of constraint evaluation in this model. Here, we assume that constraints are applied serially, and that the time it takes to apply the constraints plays a critical role in performance. Given a sufficient amount of time, the results of parallel constraint application and serial constraint application are equivalent. For this reason, we chose to apply the constraints one by one, in order of descending strength. The first constraint to apply is the strongest constraint. Only when this constraint does not distinguish between the candidates is the next constraint in the hierarchy retrieved from memory. Because Principle A ('reflexives do not have a disjoint interpretation') is stronger than Referential Economy ('avoid pronouns') (if not, pronouns would incorrectly be predicted to be totally impossible in English, see Hendriks and Spenader, 2004, in press), this constraint is retrieved first. If the input is a reflexive,

[5] A reviewer raised the question whether bidirectional optimization is necessary for reflexives if unidirectional and bidirectional optimization produce the same results. Although bidirectional optimization will not change the results, it may be that hearers, as a kind of permanent feedback mechanism that is also necessary for language acquisition (see footnote 4), automatically check whether the selected meaning will give rise to the heard form in production. Indeed, hearers seem to be constantly aware of the alternative forms the speaker could have used, as is suggested by the observation that people are able to create new scales for interpreting conversational implicatures and reason about alternative forms on those scales (Hirschberg, 1985). On the other hand, it may be that the type of linguistic reasoning involved in interpreting conversational implicatures is much more under cognitive control than the type of reasoning involved in pronoun interpretation. But since production compilation will result in higher processing efficiency, interpreting reflexives bi-directionally will eventually be about as fast as interpreting reflexives unidirectionally, thus reducing the advantage of unidirectional over bidirectional optimization.

Principle A is already able to distinguish between the two relevant candidate meanings. This is illustrated in tableau 5. Here, f’ stands for a reflexive, f for a pronoun, m for a disjoint meaning and m’ for a co-referential meaning (see section 3).

Input: f’	*<f’,m>	*f
<f’,m>	*!	
☞ <f’,m’>		

Tableau 5: Comprehension of reflexives

Input: f	*<f’,m>	*f
☞ <f,m>		*
☞ <f,m’>		*

Tableau 6: Comprehension of pronouns

Because the form-meaning pair <f’,m> violates the constraint Principle A (*<f’,m>), the form-meaning pair <f’,m’> (the pair consisting of a reflexive and a co-referential meaning) is the optimal pair. So the application of Principle A suffices to determine the optimal meaning for a reflexive. At this point, optimization can proceed in the opposite direction with the aim of providing the bi-directionally optimal pair for reflexive comprehension.

With respect to the comprehension of pronouns, in contrast, not only Principle A but also Referential Economy (*f) is insufficient for distinguishing between the two candidate meanings. This is illustrated by tableau 6, which is identical to tableau 4 from section 3.1. Because the comprehension of pronouns requires both Principle A and Referential Economy to apply, whereas the comprehension of reflexives only requires Principle A to apply, unidirectional optimization takes more time for pronouns than for reflexives. Hence, bi-directional optimization will emerge later for pronouns. This prediction might be tested by studying which candidate meanings are activated in children during the online comprehension of pronouns and reflexives.

3. Bi-directional meanings are not necessarily acquired at the same time for different linguistic forms.

The cognitive model predicts that high frequent forms will be learned faster than low frequent forms. Each time a form or meaning is met, production compilation can take place, resulting in a higher speed of processing. In this paper we focused on pronouns and reflexives. However, in section 3.2 we mentioned several other phenomena that have been argued to require bi-directional optimization. Our model predicts that, because of the dependency of speed of learning on frequency of use, these different forms are not necessarily learned at the same time during language development. Indeed, children seem to differ in the ages at which they provide adult-like responses for particular linguistic forms. Whereas from the age of 6 or 7 on children start to interpret pronouns correctly, children until roughly 11 years old select a non-adult meaning for indefinite objects (Unsworth, 2005), and many 10- and 11-year-olds do not draw a scalar implicature where most adults would (Noveck, 2001). This suggests that bi-directional optimization is not a general strategy that has to be learned by children in one step, but rather that the possibility of bi-directional optimization is dependent on the frequency of use of the relevant production rules.

In this paper we implemented Hendriks and Spenader's grammatical account of pronoun comprehension. An alternative account is Reinhart's (2004, in press) processing account, which is based on the notion of working memory capacity. Reinhart argues that children lack sufficient working memory capacity to perform reference-set computation. Reinhart takes reference-set computation to be a process performed by the parser rather than the grammar. It involves constructing for a given form a reference set consisting of form-meaning pairs, and determining whether the given form is appropriate or whether the form-meaning pair could be obtained more economically. Without a clear theory of what is considered working memory and which load different linguistic forms impose on working memory, however, it is difficult to assess Reinhart's claim. If Reinhart assumes a static view on working memory which holds that working memory has a fixed capacity over different tasks, her account will incorrectly predict that, once working memory capacity is high enough, reference-set computation should be possible for all forms. On the other hand, if working memory is supposed to vary according to the task it performs, this account introduces a large number of new, free parameters, making it difficult to generate precise predictions.

4. Bi-directional optimization can but need not be automatized.

According to our cognitive model, the speed of application of production rules can increase due to the mechanism of production compilation. Production compilation integrates two existing production rules into a new production rule. Through this mechanism of production compilation, most and perhaps even all of the process of bi-directional optimization can be automatized. Since production compilation integrates two existing production rules only if these existing production rules have been used, it is to be expected that linguistic forms differ in the degree of automization they exhibit in comprehension. For the very frequently occurring pronouns, the entire process of bi-directional optimization may be fully automatized. On the other hand, relatively infrequent forms such as newly construed scalar implicatures may have to be calculated on the spot by applying each production rule separately.

Our cognitive model implemented strong bi-directional optimization. Strong bi-directional optimization has the effect of blocking one of the meanings of an ambiguous form in comprehension, and of freezing one of the word orders in the case of variable word order in production. Strong bi-directional optimization is sufficient for explaining adult pronoun comprehension. The recursive variant of strong bi-directional optimization is called weak bi-directional optimization (Blutner, 2000). Weak bi-directional optimization has the effect that not only pairs for which there does not exist a pair with either a better form or a better meaning are optimal, but also pairs for which there does not exist a *bi-directionally optimal* pair with either a better form or a better meaning. Thus weak bi-directional optimization allows for additional optimal pairs. Blutner and Zeevat (2004) have argued that weak bi-directional optimization cannot be seen as an online mechanism of language comprehension because it makes several incorrect predictions. However, for indefinite subjects and objects it has been claimed that strong bi-directional optimization is not sufficient and weak bi-directional optimization is required (de Hoop and Krämer, 2006). According to ACT-R, all steps taken to reach a certain goal remain weakly available in declarative memory. If one assumes that these traces of previous steps might be retrieved later on, this could explain weak bi-directional optimization findings. Once it is discovered that a pair that was previously considered optimal turns out to be suboptimal for this particular case, the system might try to retrieve previously rejected suboptimal pairs to construct an additional optimal pair. Although this is not implemented in the current model, there is no theoretical reason why this would be impossible.

7. Conclusions

The acquisition of discourse-semantic aspects of language is characterized by several delays in comprehension. Recently, it has been argued that these delays in language acquisition can be attributed to children's inability to optimize bi-directionally. For this explanation to hold, it must be shown that children's comprehension abilities are able to expand from performing unidirectional optimization to performing bi-directional optimization. In this paper, we developed a cognitive model of optimization in language with the aim of studying the transition from unidirectional to bi-directional optimization. We showed that, by implementing plausible assumptions with respect to the cognitive architecture, our model started to optimize bi-directionally as soon as the unidirectional process could be performed fast enough. Given the similarities between the language acquisition data described in the literature and the simulation data of our model, a cognitively plausible explanation for children's inability to optimize bi-directionally is their relatively low speed of processing. With experience, children's processes of unidirectional optimization from form to meaning and from meaning to form gain in efficiency, until children's processing is fast enough to allow for bi-directional optimization. Thus, no qualitative differences need to be assumed between children's and adults' knowledge of the grammar or the nature of their parser.

Acknowledgements

This investigation was supported in part by grants from the Netherlands Organisation for Scientific Research, NWO (grants no. 051-02-070 and 015-001-103 for Petra Hendriks). The authors thank the participants of the workshop "Language Acquisition between Sentence and Discourse", Irene Krämer, Niels Taatgen, and two anonymous Lingua reviewers for valuable comments and discussion.

References

Aloni, M., Butler, A., Hindsill, D., in press. Nuclear Accent in Bi-directional Optimality Theory. In: Aloni, M., Butler, A., Dekker, P. (Eds.), *Questions in Dynamic Semantics*, Elsevier (CRiSPI Series), Amsterdam.

Anderson, J. R., Lebiere, C., 1998. *The atomic components of thought*. Erlbaum, Mahwah, NJ.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y., 2004. An integrated theory of the mind. *Psychological Review* 111, 1036-1060.

Bloom, P., Barss, A., Nicol, J., Conway, L., 1994. Children's Knowledge of Binding and Coreference: Evidence from Spontaneous Speech, *Language* 70, 53-71.

Blutner, R., 2000. Some aspects of optimality in natural language interpretation, *Journal of Semantics* 17, 189-216.

Blutner, R., 2006. Embedded Implicatures and Optimality Theoretic Pragmatics. In: Solstad, T., Grønn, A., Haug, D. (Eds.), *A Festschrift for Kjell Johan Sæbø*. (http://folk.uio.no/torgriso/fs/)

Blutner, R., Zeevat, H., 2004. Editors' Introduction: Pragmatics in Optimality Theory. In: Blutner, R., Zeevat, H. (Eds.), *Optimality Theory and Pragmatics*. Palgrave Macmillan, Houndmills, Basingstoke, Hampshire, and New York, pp. 1-24.

Boersma, P., Hayes, B., 2001. Empirical tests of the gradual learning algorithm, *Linguistic Inquiry* 32: 45-86.

Chien, Y.-C., Wexler, K., 1990. Children's Knowledge of Locality Conditions on Binding as Evidence for the Modularity of Syntax and Pragmatics, *Language Acquisition* 13, 225-295.

Cutler, A., Swinney, D.A., 1987. Prosody and the development of comprehension, *Journal of Child Language* 14, 145-167.

De Hoop, H., Krämer, I., 2006. Children's optimal interpretations of indefinite subjects and objects. *Language Acquisition* 13, 103-123.

De Villiers, J.G., Cahillane, J., Altreuter, E., in press. What can production reveal about Principle B? In: Deen, K.U., Nomura, J., Schulz, B., Schwartz, B.D. (Eds.), *The Proceedings of the Inaugural Conference on Generative Approaches to Language Acquisition–North America* (GALANA), Honolulu, HI. Occasional Papers in Linguistics 4, University of Connecticut, pp. 89-100.

Grimshaw, J., Rosen, S.T., 1990. Knowledge and Obedience: The Developmental Status of the Binding Theory, *Linguistic Inquiry* 21, 187-222.

Hendriks, P., Hendrickx, S., Looije, R., Pals, C., 2005. Hoe perfect is ons taalsysteem? Een bidirectionele OT-analyse van de verwerving van klemtoonverschuiving, *Tabu* 34, 71-97.

Hendriks, P., de Hoop, H., 2001. Optimality Theoretic Semantics, *Linguistics and Philosophy* 24, 1-32.

Hendriks, P., Spenader, J., 2004. A bi-directional explanation of the pronoun interpretation problem. In: Schlenker, P., Keenan, E. (Eds.), *Proceedings of the ESSLLI '04 Workshop on Semantic Approaches to Binding Theory*, Nancy, France.

Hendriks, P., Spenader, J., in press. When production precedes comprehension: An optimization approach to the acquisition of pronouns. *Language Acquisition*.

Hirschberg, J., 1985. *A Theory of Scalar Implicature*. PhD Dissertation, University of Pennsylvania.

Misker, J.M.V., Anderson, J.R., 2003. Combining Optimality Theory and a Cognitive Architecture. In: Detje, F., Doerner, D., Schaub, H. (Eds.), *Proceedings of the Fifth International Conference on Cognitive Modeling*. Universitaetsverlag Bamberg, Bamberg, pp. 165-170.

Noveck, I.A., 2001. When children are more logical than adults: experimental investigations of scalar implicature, *Cognition* 78, 165-188.

Papafragou, A., Musolino, J., 2003. Scalar implicatures: experiments at the semantics-pragmatics interface. *Cognition* 86, 253-282.

Prince, A., Smolensky, P., 2004. *Optimality Theory: Constraint Interaction in Generative Grammar*, Blackwell, Malden, MA. Appeared as Technical Report in 1993.

Reinhart, T., 2004. The Processing Cost of Reference-Set Computation: Acquisition of Stress Shift and Focus, *Language Acquisition* 12, 109-155.

Reinhart, T., in press. Processing or Pragmatics? Explaining the Coreference Delay. In: Gibson, T., Pearlmutter, N. (Eds.), *The Processing and Acquisition of Reference*, MIT Press, Cambridge, MA.

Smolensky, P., 1986. Information processing in dynamical systems: Foundations of Harmony Theory. In: Rumelhart, D.E., McClelland, J.L., & the PDP Research Group, *Parallel Distributed Processing. Explorations in the Microstructure of Cognition*, vol. 1: Foundations. MIT Press, Cambridge, MA, pp. 194-281.

Smolensky, P., 1996. On the comprehension/production dilemma in child language, *Linguistic Inquiry* 27, 720-731.

Taatgen, N., Anderson, J. 2002. Why do children learn to say “broke”? A model of learning the past tense without feedback. *Cognition* 86, 123-155.

Tesar, B., Smolensky, P., 1998. Learnability in Optimality Theory, *Linguistic Inquiry* 29, 229-268.

Thornton, R., Wexler, K., 1999. *Principle B, VP Ellipsis and Interpretation in Child Grammar*, MIT Press, Cambridge, MA.

Unsworth, S., 2005. *Child L2, Adult L2, Child L1: Differences and Similarities*. PhD Dissertation, Utrecht University.

Valkenier, B., 2006. *Bi-directionality in language acquisition*. MSc Thesis, University of Groningen.