Proceedings of the Twentieth Annual Conference of the Cognitive Science Society.
Hillsdale, NJ: Lawrence Erlbaum Associates.

Dimensions of Grammatical Coreference

Peter C. Gordon (pcg@email.unc.edu)
Department of Psychology; University of North Carolina
Chapel Hill, NC 27599-3270 USA

Randall Hendrick (hendrick@email.unc.edu)
Department of Linguistics, University of North Carolina
Chapel Hill, NC 27599-3155 USA

Abstract

The correlational structure of judgments of grammatical coreference is examined using factor analysis and the results are used to identify the dimensions of grammatical variation in competent speakers of English. The dimensions that are discovered do not correspond to those typically discussed in generative linguistics but they can be explained very naturally by a model in which coreference is achieved through a process in which linguistic expressions are mapped onto a model of discourse.

Intuitions of grammaticality constitute the most basic data to be explained by theory in generative linguistics. This has been so since Chomsky (1965) argued that a competent speaker-hearer's implicit knowledge of a language provides the best path to characterizing the essential, generative capacity of a grammar. While generative linguists have accepted intuitions of grammaticality as their basic data, they have by-and-large eschewed the development of formal methods for assessing those intuitions. There have been occasional, interesting attempts to apply more formal methodology to the study of grammatically significant intuitions, but these have not had much impact on linguistic theory (Shutze, 1996). Psychologists studying intuitions of grammatical (and other types of) well-formedness have come to characterize such studies as "offline" and to regard them primarily as ways of validating materials being used in online studies designed to reveal moment-to-moment processing of language. We believe that more systematic use of formal methods for studying intuitions of grammaticality can be of real value both to formal theories of grammar and to models of language processing. The present paper provides support for this belief by showing how applying scaling techniques to judgments of grammaticality can reveal how different types of linguistic forms give rise to dimensions of grammaticality in the domain of coreference.

In previous work (Gordon & Hendrick, 1997; Gordon & Hendrick, in press*b*), we have applied elementary techniques of experimental psychology to the study of judgments of when coreference between two noun phrases (NPs) is grammatically acceptable. The results of these studies were analyzed by calculating the mean acceptability of coreference for different types of NPs in different syntactic relations. The pattern of acceptable coreference in some cases provided support for basic claims presented in the Binding Theory (Chomsky, 1981; 1986) but in other cases did not. In the present work we use factor analysis as a scaling tool for revealing the dimensions underlying grammatical coreference in a community of competent speakers of English. The dimensions that emerge provide information about which forms of referring expressions vary together in their ease of coreferential interpretation. The resulting classification of types of referring expressions is not consistent with central theoretical principles in generative linguistics (Chomsky, 1981; 1986; Evans, 1980; May, 1985), but is consistent with a model that treats the acceptability of coreferential interpretation as emerging from the ease with which a discourse model can be dynamically constructed from linguistic input containing different types of referring expressions (Gordon & Hendrick, in press*a*).

Coreference With Names and Pronouns

In Gordon and Hendrick (1997) we report a series of surveys of intuitive judgments of coreference designed to test the adequacy of Principle C of the Binding Theory. Those surveys systematically investigated coreference possibilities in sentences with coreference between name-pronoun, name-name and pronoun-name sequences; the structures were systematically varied as to whether the coreferential elements were in a c-command relation or not. (A constituent α is said to c-command another constituent β if the first branching node that dominates α dominates β as well. Principle C states that a name or definite description cannot have a c-commanding antecedent.) Table 1 shows the types of coreferential configurations examined and sample stimuli from the fourth experiment in that work. Subjects were asked to use a six-point scale to rate the grammatical acceptability of coreferential interpretation of the two boldfaced words. The experiment manipulated the linear order of names and pronouns, and whether a c-command relation existed between them.

The implications of the pattern of means are discussed in Gordon and Hendrick (1997). Here we examine how individual variation in grammaticality judgments can reveal underlying dimensions of grammaticality. Factor analysis is a statistical tool for capturing the correlational structure in a set of data by determining how linear combinations of observed variables can account for the pattern of observed

Table 1. Sample stimuli and summary results (grammatical acceptability of a coreferential interpretation of the boldfaced words on a 1 to 6 scale) from Experiment 4 of Gordon and Hendrick (1997).

Type of Sequence	Example Stimuli	C-Command	Average
Name-Pronoun	**Lisa's** brother visited **her** at college.	No	4.47
Name-Pronoun	**Lisa** visited **her** brother at college.	Yes	5.32
Name-Name	**Lisa's** brother visited **Lisa** at college.	No	4.12
Name-Name	**Lisa** visited **Lisa's** brother at college.	Yes	3.50
Pronoun-Name	**Her** brother visited **Lisa** at college.	No	2.70
Pronoun-Name	**She** visited **Lisa's** brother at college.	Yes	2.44

correlations between these variables. It can be used either in an exploratory or confirmatory manner. For current purposes, we have performed an exploratory factor analysis of the data from this experiment. Table 2 shows the correlations in subjects' ratings between sentences in the six basic coreference conditions. A very substantial positive correlation of .76 was observed between ratings for the Name-Name sentences in the c-command and no c-command conditions. Another very substantial positive correlation of .68 was observed between ratings for the Pronoun-Name sentences in the c-command and no c-command conditions. Two other correlations were smaller but still significant: Name-Pronoun sentences in the two c-command conditions (.32) and c-commanded Name-Pronoun sentences with c-commanded Pronoun-Name sentences (-.30). No other correlations approached significance. The strong correlations suggest that our subjects showed reliable individual differences along clear syntactic dimensions.

A factor analysis simplifies the pattern in this correlation matrix by determining how linear combinations of the observed variables can account for the pattern of observed correlations between those variables. Factor analysis of these data revealed three factors with eigenvalues greater than one, which together accounted for 82.3 percent of the predictable variance in the matrix. Table 3 shows the resulting factor matrix. The absolute value of the numerical entries for the factors indicates on a scale of zero to one the extent to which individual coreference conditions contribute to the factors. Interpretation of the pattern of these weights provides the basis for labeling the factors. Accordingly, Factor 1 can be called the "Name-Pronoun" factor because it depends on Name-Pronoun sequences of NPs. Factor 2 can be called the "Pronoun-Name" factor because it very clearly depends on the "Pronoun-Name" sequences of NPs, with a smaller contribution of the c-commanded Name-Pronoun sequences. Factor 3 can be called the "Name-Name" factor, because it depends on the Name-Name sequences of NPs. The very clear cut pattern of these factors suggests that different principles govern the grammaticality of these different sequences of NPs.

The correlational results (Table 2) and the subsequent factor analysis (Table 3) provide a direct window on the systematic variation of acceptable coreference in a community of competent users of English. They show that there are three independent factors along which individuals reliably vary in their willingness to accept coreference between two NPs. These factors are related to the three sequences of types of NPs that we explored: Name-Pronoun, Pronoun-Name and Name-Name. On each factor a subject's criterion for accepting coreference with or without c-

Table 2. Correlations in subjects' mean grammaticality ratings for the different configurations of referring expressions for the data in Table 1. *N = Name, P = Pronoun, Yes = C-Command, No = no C-Command.*

	N-P: Yes	N-N: No	N-N: Yes	P-N: No	P-N: Yes
N-P: No	.32**	-.07	-.13	.15	.15
N-P: Yes		.15	.06	-.19	-.30**
N-N: No			.76***	-.09	-.17
N-N: Yes				.01	.01
P-N: No					.68***

n = 94, * p < .05; ** p < .005; *** p < .001

Table 3. Factor matrix for subjects' grammaticality ratings in Table 1. *N = Name, P = Pronoun, Yes = C-Command, No = no C-Command.* This factor matrix is shown after varimax rotation, which is a standard factor analytic technique for assisting in interpreting the factors (Gorsuch, 1983). Varimax rotation will tend to promote a representation of the factor-analytic solution that emphasizes separate factors and minimizes an overall group factor.

	Factor 1	Factor 2	Factor 3
N-P: No	.84	.26	-.13
N-N: No	.05	-.11	.93
P-N: No	.04	.89	.01
N-P: Yes	.78	-.34	.14
N-N: Yes	-.06	.05	.94
P-N: Yes	-.06	.91	-.05

command is quite similar. Accordingly, a subject who was conservative in accepting coreference in the Name-Name condition without c-command would also tend to be conservative in accepting it with c-command. Likewise, a subject who was liberal about coreference in the Name-Name condition would be that way regardless of whether there was a c-command relation. Such strong relations do not exist across factors. Accordingly, each factor can be seen as a *principle* of linguistic knowledge while a subject's criterion for acceptance can be seen as a *parameter* setting for that principle.

Of course, the analysis of language structure in terms of principles and parameters has been an important part of syntactic theory for the last decade (Rizzi, 1982; Huang, 1982). This approach has been applied across languages in an effort to adduce universal principles of human language, and to explain language variation in terms of the setting of parameters. Linguists often try to explain gross differences between languages and major dialects this way. In principle, it should extend to idiolectical variation as well. Applying this approach to experimental results obtained from a collection of speakers of a single language is novel as far as we know. We do so because the correlational structure of the subjects' judgments is strikingly clear in its delineation of independent linguistic dimensions, which constrain the kinds of variation that exist in individual subject's judgments. That is what the principles and parameters approach is meant to do. The principles that our correlational analysis reveals are not consistent with those of binding theory. Classical binding theory (Chomsky, 1981) offers only a single Principle C that should apply to both Name-Name and Pronoun-Name sequences. That is not consistent with our finding that these dimensions are independent in our subjects' acceptability judgments. Further, our findings lead to a conception of parameter setting in terms of specifying the value of a criterion on a continuous dimension. That conception also differs from the standard approach to principles and parameters.

From this point of view correlational studies of linguistic variability within a speech community bear directly on claims about hypothesized principles and parameters. Our correlational analyses show that there is systematic variation across individuals in their acceptance of different kinds of linguistic coreference that can be modeled in terms of principles and parameters of grammar. We should note, however, that these analyses cannot by themselves provide evidence about why that variation exists. Gerken and Bever (1986) have examined individual variation in the acceptance of coreference in pronoun-name sequences, and have shown that it is related to differences among individuals in how their comprehension strategies handle syntactic closure. We are not in a position to offer this kind of basis for interpreting the differences among our individual subjects. However, our examination of a large variety of linguistic constructions does allow us to provide compelling evidence about the dimensions of linguistic variation that are psychologically salient.

Coreference With Non Definite NPs

In Gordon and Hendrick (in press*b*) we examine patterns of acceptable coreference involving non definite NPs (such as *no one*, *everyone*, and *who*) and compare them to patterns of acceptable coreference involving definite NPs (such as names and definite descriptions). It is commonly held that there are special structural restrictions on the relation between an non definite NP and a pronoun that it binds (cf. (Hornstein, 1995)). These differences have been used to support claims that non definite NPs semantically bind pronouns in a different fashion than definite NPs (Evans, 1980) and have been crucial to the development of claims about the existence of Logical Form as an abstract level of syntactic structure (May, 1985). These theoretical claims have also depended on Principle C of the Binding Theory with its reliance on c-command. Because our results on coreference involving names and pronouns (Gordon & Hendrick, 1997) challenge Principle C we undertook a formal study of judgments of coreference involving non definite expressions (Gordon & Hendrick, in press*b*).

The first experiment in Gordon and Hendrick (in press*b*) examined the assumption, common since (Evans, 1980), that

Table 4. A summary of conditions, sample stimuli and results used in Gordon and Hendrick (in press*b*). Subjects were asked to rate the grammatical acceptability of the sentences given a coreferential interpretation of the expressions in boldface.

C-Command	Type of Sequence	Sample Stimulus	Mean Rating & (95% Confidence Interval)
Yes	Quantifier-Pronoun	**Each girl** decided what **she** could do.	4.83 ± .18
No	Quantifier-Pronoun	**Each girl's** parents decided what **she** could do.	4.05 ± .24
Yes	Pronoun-Quantifier	**She** decided what **each girl** could do.	2.43 ± .30
No	Pronoun-Quantifier	**Her** parents decided what **each girl** could do.	2.42 ± .24
Yes	Definite Description-Pronoun	**The girl** decided what **she** could do.	5.43 ± .14
No	Definite Description-Pronoun	**The girl's** parents decided what **she** could do.	4.98 ± .20
Yes	Pronoun-Definite Description	**She** decided what **the girl** could do.	2.27 ± .26
No	Pronoun-Definite Description	**Her** parents decided what **the girl** could do.	3.20 ± .22
Yes	Name-Pronoun	**Jill** decided what **she** could do.	5.45 ± .16
No	Name-Pronoun	**Jill's** parents decided what **she** could do.	5.23 ± .21
Yes	Pronoun-Name	**She** decided what **Jill** could do.	2.17 ± .27
No	Pronoun-Name	**Her** parents decided what **Jill** could do.	3.19 ± .24

Table 5. Results of the factor analysis of the judgments obtained in the first experiment of Gordon and Hendrick (in press*b*).

C-Command	Type of Sequence	Factor 1	Factor 2
Yes	Quantified NP-Pronoun	.72	-.14
No	Quantified NP -Pronoun	.58	.09
Yes	Pronoun-Quantified NP	-.11	.84
No	Pronoun-Quantified NP	.13	.77
Yes	Definite NP -Pronoun	.79	-.12
No	Definite NP -Pronoun	.83	.13
Yes	Pronoun-Definite NP	-.32	.83
No	Pronoun-Definite NP	.32	.61
Yes	Name-Pronoun	.83	.00
No	Name-Pronoun	.89	.08
Yes	Pronoun-Name	-.26	.85
No	Pronoun-Name	.25	.58

quantified expressions must c-command pronouns that they bind. We also evaluated the long-standing claim that quantified expressions exhibit weak and strong crossover. (Crossover refers to configurations in which a question word and its trace are separated by a coreferential pronoun.) The second experiment in Gordon and Hendrick (in press*b*) explored the nature of crossover phenomena more deeply, examining how and why structures involving syntactically dislocated arguments exhibit the crossover restriction. The method was similar to that in Gordon and Hendrick (1997) and involved ratings of whether coreference between two boldfaced expressions was acceptable. The experimental conditions as well as sample stimuli and average acceptability ratings are shown in Table 4 for the first experiment and Table 6 for the second experiment. The implications of these patterns of average acceptability are discussed Gordon and Hendrick (in press*b*). Here we use factor analysis to examine the correlational structure of these ratings in order to provide evidence about whether non definite and definite NPs contribute to the same or different dimensions of grammatical variation.

Our exploratory factor analysis (presented above) of the results of the results of Experiment 4 from Gordon and Hendrick (1997) yielded two factors, "pronoun-name order" and "name-pronoun order", that could be expected to be found in the results of our current experimental data set if quantified expressions are psychologically grouped with definite expressions and names. Accordingly, a factor analysis restricted to two factors was performed on the results of the two experiments from Gordon and Hendrick (in press*b*) on coreference involving non definite expressions. The results for the first experiment are shown in Table 5.

Examination of the absolute values of the entries for Factor 1 shows that it clearly depends on quantified NP-pronoun sequences, definite NP - pronoun sequences and name-pronoun sequences regardless of the presence or absence of a c-command relation. In contrast, Factor 2 depends on the remaining types of NP sequences, pronoun-quantified NP sequences, pronoun- definite NP sequences and pronoun - name sequences, again showing no special sensitivity to c-command. This pattern strongly suggests that linear order rather than hierarchical relations lies behind the intuitive judgments of our subjects.

A factor analysis of the data in Table 6 is given in Table 7 and provides further support for the importance of linear order in these intuitive judgments by yielding two factors whose loadings are consistent with the Name-Pronoun and Pronoun-Name factors that were observed in the preceding analyses. The general pattern in Table 7 accords well with our earlier claim that quantified expression group with definite expressions and names in how they bind pronouns.

The general tendency in Table 7 for the linear order NP-

Table 6. Experimental conditions, sample stimuli and results from Experiment 2 of Gordon & Hendrick (in press*b*). Subjects were asked to rate the grammatical acceptability of the sentences given a coreferential interpretation of the expressions in boldface and italics. Mean grammaticality ratings (on a scale of 1 to 6), along with the lower and upper bounds of the 95% confidence intervals, are shown for the different kinds of coreferential expressions.

Type of Sequence	Crossover	C-Command	Sample Sentence	Mean Rating (95% C.I.)
WH trace -pronoun	No	No	Brian knows **which girl** asked the teacher to help *her* classmates.	4.41 ± .21
WH trace - pronoun	No	No	Brian knows **which girl** asked the teacher to help *her*.	4.45 ± .24
pronoun -WH trace	Weak	No	Brian knows **which girl** *her* classmates asked the teacher to help.	2.97 ± .23
pronoun - WH trace	Strong	Yes	Brian knows **which girl** *she* asked the teacher to help.	2.62 ± .27
OP trace - pronoun	No	No	Brian knows **the girl** that asked the teacher to help *her* classmates.	4.69 ± .19
OP trace- pronoun	No	No	Brian knows **the girl** that asked the teacher to help *her*.	4.91 ± .20
pronoun -OP trace	Weak	No	Brian knows **the girl** that *her* classmates asked the teacher to help.	2.95 ± .22
pronoun -OP trace	Strong	Yes	Brian knows **the girl** that *she* asked the teacher to help.	2.48 ± .20
name - pronoun		No	Brian knows that **Karen** asked the teacher to help *her* classmates.	5.27 ± .14
name - pronoun		No	Brian knows that **Karen** asked the teacher to help *her*.	5.40 ± .14
pronoun - name		No	Brian knows that *her* classmates asked the teacher to help **Karen**.	3.55 ± .17
pronoun - name		Yes	Brian knows that *she* asked the teacher to help **Karen**.	2.66 ± .18

Table 7. Results of the factor analysis of the judgments obtained in the second experiment of Gordon and Hendrick (in press*b*).

Type of structure	c-command	Factor 1	Factor 2
wh-trace pronoun	yes	.73	.19
wh-trace pronoun	yes	.76	-.26
pronoun wh-trace	no	.02	.77
pronoun wh-trace	yes	-.16	.84
OP-trace pronoun	yes	.75	.16
OP-trace pronoun	yes	.78	-.26
pronoun OP-trace	no	.40	.35
pronoun OP-trace	yes	-.10	.76
name pronoun	yes	.69	.00
name pronoun	yes	.67	-.33
pronoun name	no	.25	.70
pronoun name	yes	-.17	.70

pronoun and pronoun-NP to be inversely related accords well with our earlier claim that quantified expression group with definite expressions and names in how they bind pronouns.

In this analysis, relative clauses that have the pattern pronoun OP-trace show the weakest loadings on the relevant factor when there is no c-command relation present. From our perspective on the correlational structure of idiolectical variability, weak loading of pronoun OP-trace sequences without c-command should reflect a parameter of grammatical knowledge. A possible candidate is the often proposed parameter of Universal Grammar that gives grammars of particular languages an option between traces and resumptive pronouns. Many languages (e.g., Irish, McCloskey, 1990) allow resumptive pronouns where standard English requires a trace and such languages do not exhibit weak crossover effects. Accordingly, the weak loading of sentences with relative clauses and pronoun - OP trace sequences without c-command could emerge from the use, by some subjects, of the resumptive pronoun parameter. The resumptive pronoun parameter is unavailable for WH questions because they involve movement (eschewing resumptive pronouns in all languages) and it is inapplicable to pronoun-name sequences. This provides an explanation of the judgments of pronoun OP-trace sequences in our data, but it does not provide a good account of the dependence of these judgments on c-command. Other possible explanations of this finding could draw on the fact that in all of our stimuli where a pronoun c-commands an OP-trace, the pronoun is in subject position adjacent to the head of a relative clause. Assessing the validity of these possibilities will require further investigations into the grammaticality of coreference in these types of NP sequences.

General Discussion

While a complete explanation of the details of the grammaticality judgments requires further work, the larger pattern is very clear in revealing dimensions of grammaticality in coreference. The first factor analysis revealed three clear factors which we dubbed (Name-Pronoun, Pronoun-Name, and Name-Name). The second and third factor analyses showed that different types of full NPs (definite descriptions, quantified expressions, and WH phrases) grouped together with names in terms of the variation they elicited in subjects' grammaticality judgments; these analyses provided further support for the Name-Pronoun and Pronoun-Name factors (with the label "name" being broadened to include these other types of NPs). These findings provide a basis for treating these different types of NPs as equivalent with respect to coreference and for focusing on how coreference is achieved between these different types of NPs and pronouns that either follow them or precede them.

In Gordon and Hendrick (in press*a*) we have developed a model of the interpretation of referring expressions as deriving from the way in which syntactic representations are mapped onto discourse representations. This explanation is expressed using a formalism created by Kamp and Reyle (1993) which seeks to apply model theoretic semantics to phenomena in natural language, particularly the semantics of discourse. Our model aims to provide an explanation for the distribution and processing of coreferential expressions within and between sentences. Here we offer a summary of the model and how it accounts for the ease of coreferential interpretation of different types of referring expressions. Then we show how the model can account for the dimensions of grammatical coreference revealed by the results presented in this paper.

Models based on the Kamp and Reyle formalism use *Construction Rules* to map syntactic representations onto *Discourse Representation Structures*. Each construction rule is composed of a *triggering condition* that delineates the linguistic input that causes the rule to apply, and specific instructions for replacing part of the linguistic representation with a information in the Discourse Representation Structure. The material added to a Discourse Representation Structure by a Construction Rule is termed a *condition set*. What we take to be important in this formalism is that syntactic structures trigger the building of a part of a discourse representation structure in a dynamic fashion. Entities are introduced into the universe of the discourse by the dynamic rules that construct these representations. These entities are termed *discourse referents*.

Our model includes three construction rules to handle reference and coreference. (1) The construction rule for names is triggered by the occurrence of a name in the syntactic representation of the input sentence; it then posits a new discourse entity in the discourse model predicated on the name. The same basic mechanism of interpretation applies to the definite and non definite expressions. (2) The construction rule for pronouns is triggered by a pronoun; it then attempts to find a *suitable antecedent* in the discourse model. Failing that, it posits a new discourse entity. (3) The final construction rule (for equivalence) is triggered by the presence of two discourse entities predicated on the same name. It adds a condition set to the discourse representation equating the two entities.

These three rules provide a straightforward account of our results concerning the acceptability of coreference in the

three types of NP sequences shown in Table 1. Name-pronoun sequences trigger successive application of the construction rule for names and then for pronouns. The rule for names posits a discourse entity predicated on the name, this subsequently provides a suitable antecedent when the construction rule for pronouns is triggered, yielding coreferential interpretation of the name and pronoun. Name-name sequences trigger two applications of the construction rule for names, resulting in two discourse entities predicated on the same name. This situation triggers the construction rule for equivalence, which equates the two discourse entities. Compared to name-pronoun sequences, establishing coreference in name-name sequences requires an additional construction rule (equivalence) and results in a more complex discourse representation. This is in accord with our subjects' judgments of coreference. Pronoun-name sequences first trigger application of the construction rule for pronouns. As it can find no suitable antecedent, it posits a new entity in the discourse model. Subsequently, the construction rule for names is triggered, resulting in the positing of a name. However, the equivalence rule is not triggered at this point because predication of the entity created by the pronoun rule has no identifying information. This accounts for the low acceptability of coreference in pronoun-name sequences.

Because the construction rules for interpreting definite and non definite expressions share the essential features of the construction rule for names, the principles outlined above also explain the relative difficulty of establishing coreference in sequences containing those types of NPs as shown in Tables 4 and 6. Thus, the three basic construction rules can be seen as providing an account of the relative acceptability of coreference in sequences of different types of NPs.

The ways in which coreference is established for the three different kinds of NP sequences leads to a straightforward account of the finding that these kinds of coreference correspond to separate dimensions of grammaticality in a community of competent language users. Coreference in each case is determined by how different linguistic configurations trigger different construction operations. If individual subjects have thresholds for triggering that are independently set for the different construction rules, then we will find that subjects' judgments of different reference conditions are highly correlated when coreference requires the same construction rule, and uncorrelated when they do not. This is the pattern shown by the correlation matrix in Table 2 and by the factor analyses in Tables 3, 5 and 7. The Name-Pronoun factor (Factor 1 in Tables 3, 5 & 7) is created by individual differences in sensitivity to triggering the construction rule for pronouns, which is the basis for coreferential interpretation in these cases. The Name-Name factor (Factor 3 in Table 3) is similarly created by differing thresholds for triggering the construction rule for equivalence. The Pronoun-Name factor (Factor 2 in Tables 3, 5 and 7) emerges from the fact that establishing coreference in such sequences requires that the NPs in a sentence not be interpreted in a strict left-to-right fashion, something that clearly happens in cases where an NP is present in a fronted adjunct expression. The Pronoun-Name factor can be explained by the supposition that some subjects are more inclined than others to interpret NPs in the verb phrase before NPs in the subject.

In this way, our model provides a straightforward account for the dimensions of grammatical variation revealed by our analyses of the correlational structure of judgments of coreference.

Acknowledgments

This research was supported by NSF Grant IRI-94-04756.

References

Gorsuch, R. L., (1983). *Factor analysis*. Hillsdale, NJ: Erlbaum Associates.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge: MIT Press.

Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht, The Netherlands: Foris.

Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin and Use*. New York. Praeger.

Evans, G. 1980. Pronouns. *Linguistic Inquiry, 11*, 337-62.

Gordon P.C. & R. Hendrick. (1997). Intuitive Knowledge of Linguistic Coreference. *Cognition, 62*, 325-370.

Gordon P.C., & Hendrick, R.. (in press*a*). The Representation and Processing of Coreference in Discourse. *Cognitive Science*.

Gordon P.C., & Hendrick, R. (in press*b*). Intuitive knowledge of non definite NP anaphora. *Proceedings of the Chicago Linguistics Society*. Chicago, IL: University of Chicago Press.

Hornstein, N. (1995). *Logical Form*. *From GB to Minimalism*. Oxford, UK: Blackwell.

Huang, C.T. (1982). *Logical relations in Chinese and the theory of grammar*. Doctoral dissertation. Department of Linguistics and Philosophy, MIT, Cambridge, MA.

Kamp, H., & Reyle, U. (1993). *From Discourse to Logic. Introduction to Model Theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic Publishers.

May, R. (1985). *Logical form*. Cambridge, MA: MIT Press.

McCloskey, J. (1990). Irish resumptive pronouns. In R. Hendrick (Ed.), *Syntax and Semantics*, *23*, 199-248. San Diego, CA: Academic Press.

Rizzi, L. 1982. *Issues in Italian syntax*. Dordrecht, The Netherlands: Foris.

Schutze, C. 1996. *The empirical base of linguistics : grammaticality judgments and linguistic methodology*. Chicago, IL: University of Chicago Press.