

A Critical Look at the Foundations of Autonomous Syntactic Analysis

Lawrence Birnbaum

Yale University
Department of Computer Science
New Haven, Connecticut

I argue that the claim of autonomous syntactic processing is irrefutable if non-determinism is permitted, e.g., by the use of arbitrary choice and backtracking. The vast majority of existing models of syntactic analysis, therefore, cannot support such a claim if it is to be considered an *empirical* claim. More recent deterministic theories of syntactic analysis—in particular, Marcus (1980) and Marcus *et al.* (1983)—seem at first glance more promising. However, by their repeated failure to address many of the problems which make language analysis so difficult in the first place—such as lexical ambiguity and genuine structural ambiguity—these theories too fail to assert or support an empirically significant claim of autonomous syntactic processing. Moreover, the very problems in language analysis that these theories ignore provide strong evidence that such claims are in fact false.

AUTONOMY AND DETERMINISM

The autonomy of syntax—or the lack of it—has long been a contentious issue within the cognitive sciences. Part of the contentiousness has been due to a certain degree of vagueness, probably unavoidable, as to what exactly would constitute “autonomy.” A rather modest claim might be that semantics does not dictate the content of syntactic facts, for example the fact that adjectives precede nouns in English noun phrases. While such a claim might be interesting from a linguistic perspective, however, it is the question of whether syntactic *processing* is autonomous that is of greatest interest to artificial intelligence. Within the realm of processing, the claim that is generally identified with syntactic autonomy is the following: That the syntactic structure of linguistic input can and should be determined by a process that makes little or no reference to semantic and contextual information.

Now, whether syntactic analysis *should* be autonomous in this way depends on whether that would, on balance, be the most useful approach. Such utility might be defined with respect to the ability to explain characteristics of human language processing, or more functionally, with respect to facilitating the construction of language processing systems that are more efficient, more easily extended, more robust, etc. The more basic question, however, and the one that is addressed most directly by computational modelling, is this: Is it even *possible* to construct models of syntactic analysis that are autonomous in the sense described above?

On the face of it, obviously, this question seems rather silly. After all, quite a large number of autonomous syntactic analyzers have been written and are in use at the present time. However, the aim of this paper is to show that current models of syntactic analysis in fact fail to support an

empirically significant claim of autonomous syntactic processing. My argument will involve making explicit what is in fact required in order for such a claim to be “empirically significant.”

The first part of the argument concerns *non-deterministic* models of syntactic analysis, which constitute by far the overwhelming majority of such models. When confronted with a possible or genuine ambiguity in the syntactic analysis of an input, non-deterministic parsers either pursue all possible analyses, or decide which analysis to pursue *arbitrarily*, while making some provision for backtracking should that decision prove erroneous. Leaving aside the obvious issue of computational efficiency, or even tractability, there is a deeper problem with such an approach. For if non-determinism is permitted, then *if it is possible to construct a process model of syntactic analysis at all, it will always be possible to construct one that is autonomous*: Whenever a syntactic decision seems to require contextual information of the sort that is forbidden under the assumption of autonomy, a non-deterministic model can always be constructed that simply pursues all possible choices, or else chooses arbitrarily, and then backtracks if that choice proves mistaken. Thus, the claim that the syntactic analysis of linguistic inputs can be accomplished without appeal to semantic or contextual information is *irrefutable* if non-determinism is permitted. As a result, non-deterministic models of syntactic analysis cannot support the claim that syntactic analysis is possible without the use of semantic and contextual information—at least, not if that claim is intended as an empirical one. This includes, to repeat, the vast majority of syntactic analyzers, including obviously all ATN-like models (e.g., Thorne, Bratley, and Dewar, 1968; Bobrow and Fraser, 1969; Woods, 1970; and numerous descendant models), as well as Prolog-based parsers (e.g., Colmerauer, 1978; Pereira and Warren, 1980).

In order to support an empirically meaningful claim of syntactic autonomy, then, a model of syntactic analysis must be *deterministic*. This makes Marcus’s (1980) attempt to construct such a model, and subsequent efforts, particularly significant and noteworthy. However, I will argue below—using Marcus’s theory as a prototype case—that these models, too, systematically fail to address issues that must be addressed in order to support an empirically significant claim of autonomous syntactic processing.¹ In particular, I will articulate the functional requirements that must be met by the output of a syntactic analyzer in order to sustain such a claim, show how current theories fail to address those requirements, and show further that there is good evidence that in fact those requirements cannot be met.

One final comment: It is important to be clear that the main point of this paper is *not* to refute the possibility of any empirically significant claim of autonomous syntactic processing being true, although I believe that the preponderance of evidence argues for such a conclusion. Rather, my point is that current models of syntactic analysis simply fail to support such an empirically significant claim.

DETERMINISM AND GENUINE STRUCTURAL AMBIGUITY

Genuine structural ambiguity poses a severe problem to a theory of deterministic, autonomous syntactic analysis. The dedication to autonomy largely forecloses the use of semantic and

¹My use of Marcus’s theory as the prototypical case should not be taken to imply that it suffers from any special problem. Quite the contrary, it reflects my view that his work constitutes the most comprehensive and explicit defense of syntactic autonomy from a computational perspective.

contextual information to resolve such ambiguity within the language analyzer itself, while the dedication to determinism militates against generating all possible analyses automatically. What Marcus (1980) decided to do about this quandary, therefore, was to preserve autonomy, but shift the burden of non-determinism to the language understanding process as a whole. In particular, he proposed that the syntactic analyzer would produce only one analysis of an input sentence at a time—even if others were possible—and that if this analysis proved incorrect, the analyzer would be called on the same input again, with some provision made to block the original, erroneous analysis.

Such an approach does not immediately render the theory empirically insignificant: After all, it might prove impossible to construct a deterministic autonomous analyzer capable of dealing even with those ambiguities that remain. Nevertheless, it is highly questionable for several reasons. First, I think it is fair to say that such an approach greatly reduces the scope of Marcus's claims: His parser no longer needs to be concerned about resolving those potential structural ambiguities that turn out to reflect actual structural ambiguities of the input. Even more, it no longer needs to deal with any subsequent potential ambiguity that might be caused by prior, genuine ambiguity—that is, potential ambiguity that might have arisen if some prior genuine ambiguity had not been resolved, expeditiously, by fiat.²

Second, on this account, the rest of the understanding system must in some way be able to prevent the syntactic analyzer from producing the original, and erroneous, analysis, so that some other analysis will be produced instead. But in order to do that, the language understanding process as a whole needs some knowledge of, and access to, the internals of the syntactic processor. It is, of course, exactly to avoid this kind of interaction that autonomous theories are proposed in the first place.

Now, what Marcus (1980) seems to imply, presumably in an attempt to forestall such criticism, is that the rest of the understanding system need not know very *much* in order to control the syntactic analyzer's behavior in this way. All that seems necessary is a communication protocol by which the rest of the understanding system could send a message indicating that the syntactic analyzer should "reparse the input, taking a different analysis path, if the other consistent analyses are desired," (Marcus, 1980, note 10, p. 13). However, the apparent simplicity of such an interaction belies the underlying additional complexity of the syntactic analyzer which it presupposes. In particular, in order to make such a scheme work, the syntactic processor would need to keep a record of the decisions which, if made differently, would result in an alternate analysis, including its state at each of those decision points, the order in which those decisions arose, and some way of keeping track of which alternatives it had already chosen in producing previous analyses. In a word, that is, it would need all of the machinery that is needed to implement non-determinism.

The main defect of this approach, however, does not lie in the additional costs, in terms of increased complexity, that are imposed on the syntactic processor itself. The real problem here is that the theory is robbing Peter, repeatedly, to pay Paul once: Determinism and autonomy are preserved for the syntactic analyzer at the expense of non-determinism in the language understanding process as a whole. On this account, determining the correct interpretation of

²These same objections can be raised about Marcus's failure to address the issue of lexical ambiguity; see Birnbaum (1985).

genuinely ambiguous sentences requires the use of arbitrary choice and backtracking not by the syntactic processor alone, but by the syntactic and semantic processors in conjunction. That is, if the knowledge of and access to the internals of the syntactic analyzer on the part of the understanding system as a whole are to be kept to a minimum—as they must be, in order to uphold the claim of autonomy—then not only can semantic and contextual processes play no role in determining the syntactic analysis of an input utterance, but the only information that they can transmit to the syntactic module about an inappropriate analysis is that it is inappropriate. As a result, no information about the particular *way* in which the analysis happens to be inappropriate can be used to help produce the correct analysis. In sum, the cost of minimizing the bandwidth of communication between syntactic processing and other processing in this case is that the language understander as a whole is reduced to the crudest and most expensive possible method for producing an appropriate interpretation of the input, namely, arbitrary choice and backtracking.³

Perhaps most troubling, however, is not the fact that Marcus's theory avoids the problems posed by genuine structural ambiguity, but the manner in which it does so. Simply putting such problems aside—shifting the burden of non-determinism to the understanding process as a whole—places this entire approach on a slippery slope. One can, obviously, continue to put aside each problem that seems to require sacrificing either determinism or autonomy, and narrow the scope of the theory further and further—what has been dubbed by one of my colleagues “the incredible shrinking module.” But, just as with the use of non-determinism, this leads ultimately to theories which are irrefutable. It is tautological that the subset of syntactic decisions that can be resolved deterministically, using only syntactic information, can in fact be resolved within a deterministic, autonomous model of syntactic analysis. If the claim of autonomy is to have any empirical force, therefore, it is necessary to show that the set of such decisions, and the rules needed to resolve them, can be characterized in advance. And it seems to me that the fact that exactly the same potential ambiguity can, in some cases, be resolved within a sentence on syntactic grounds, and in others leads to genuine structural ambiguity, immediately refutes this possibility in the case of most, if not all, syntactic decisions.

This point also bears on one of the most common arguments made in favor of autonomous syntactic analysis. There is a common-sense grain of truth at the root of the idea, and it is often put as follows: “What’s the problem here? You let syntax do what it can deterministically, and then semantics takes care of the rest.” Although even this claim is arguable, let’s suppose that it were true. It simply does not follow that there must exist an independent syntactic processor which makes those particular decisions that can in fact be made deterministically using only syntactic information. It is perfectly compatible with almost any view of language processing that decisions that don’t happen, in a particular instance, to require the use of semantics can be made, in that instance, without such information. Once again, therefore, the fact that particular *instances* of syntactic decisions can be made deterministically using only syntactic information does not support an empirically meaningful claim of autonomous syntactic analysis. What is

³It is also worth pointing out that this approach to genuinely ambiguous sentences seems inconsistent with Marcus's account of garden path sentences. That account rests on the claim that people do not backtrack and produce an alternate analysis for an input without becoming consciously aware of that fact. However, this raises the question of why people are not consciously aware of backtracking when confronted with genuinely ambiguous sentences. In any case, Crain and Steedman (1985) have demonstrated that the phenomenon of garden path sentences is crucially dependent upon semantic and pragmatic factors, and that no purely syntactic account, such as Marcus's, can be correct.

necessary is that classes of such decisions, and the rules which are necessary to make them, be characterized and specified in advance as the proper domain of an autonomous syntactic analyzer.

REDEFINING THE OUTPUT

Consider now the problems raised in determining, and representing, the correct role of a prepositional phrase within a sentence—one of the most common instances of genuine structural ambiguity in English, and moreover a problem that is widely acknowledged, even by Marcus, to require heavy use of semantic and contextual information. Because a syntactic analyzer is by itself incapable of correctly determining where to attach a prepositional phrase, Marcus's original parser simply attached all such phrases to the closest available constituent that was syntactically acceptable, regardless of whether or not that is correct (personal communication). For example, it would analyze the sentence "I kissed the girl in the park on the lips," as if the prepositional phrase "on the lips" modified the noun "park."

More recently, Marcus, Hindle, and Fleck (1983) have attempted to face up to this sort of problem by proposing that the output of a syntactic analyzer should not, in fact, be a phrase marker of the sort employed in linguistics. Instead, they propose that it should be something vaguer and less informative, a description corresponding to a set of such phrase markers.⁴ In particular, the relation "immediately dominates" in the structural description of a sentence is replaced by the less informative relation "dominates," and constituents are referred to by non-unique names, so that two symbols may—if the facts that are known about the constituents to which they refer are compatible—turn out to refer to the same constituent.

Unfortunately, however, once again no functional justification is offered for this sort of output, other than the fact that it may prove possible to produce it without the use of semantic and contextual information, and once again this failure places the entire approach on a slippery slope. As long as the output of the syntactic analyzer can be redefined to be less informative—without being subject to any constraints of functional utility—whenever the attempt to produce a more informative output appears to threaten either autonomy or non-determinism, then of course it will always be possible to maintain them both. But if the output of a syntactic analyzer is defined as that structural information about a sentence which can be deterministically derived without appeal to semantic or contextual information, then it is once again tautological that deterministic autonomous syntactic analysis is possible. Unless it can be demonstrated that such an output will be functionally useful, the claim of autonomy is without empirical content. Just as much as permitting arbitrary choice and back-up, defining the output in this way results in autonomous theories which are irrefutable. To take this to its logical conclusion, one can quite simply write a deterministic parser that does not require semantics: It need only take in strings and output them without alteration. Without applying the constraints of functional utility, we have no guarantee that the output of a process is anything other than a trivial transformation of the input.

Moreover, although a functionally justified output is a *necessary* condition for an empirically significant claim of autonomy, it is by no means *sufficient*. It must also be shown that the information contained in such an output can actually be exploited without violating the original

⁴A similar proposal was subsequently made by Barton and Berwick (1985). The comments that follow apply equally to this work.

claim of syntactic autonomy, and this is not a foregone conclusion: The utilization of such a representation may itself turn out to require the highly integrated application of syntactic, semantic, and contextual information. That is, even if it were possible to produce useful syntactic representations in a deterministic autonomous fashion, such representations might prove too impoverished to support the application of useful—or even necessary—syntactic rules. Indeed, it is quite possible that a given syntactic rule could sometimes be applied by the syntactic processor—given a simple and unambiguous enough input—while at other times the information necessary to apply the same rule would not be available. If such rules are genuinely useful, then it will be necessary to apply them after sufficient information about the structure of the input has been recovered—which is to say, after semantic and contextual information have been applied. There seems little point to the claim of autonomy under such circumstances.

It is important to understand that the above argument is not merely theoretical. The output produced according to Marcus *et al.* (1983) is in fact insufficiently informative to support the use of the putative syntactic rules constraining pronoun reference cited in Marcus (1984), in response to Schank and Birnbaum (1984), as arguing for the need to compute explicit, independent syntactic representations. Such rules depend on knowing, rather precisely, how high in the structural description of a sentence a noun phrase is with respect to a potentially co-referent noun phrase, and this is exactly the sort of information that has been discarded by Marcus *et al.*'s later theory.

In particular, the rule cited by Marcus (1984) is as follows:⁵

Given two noun phrases NP1 and NP2 in a sentence, if (1) NP1 precedes NP2 in the sentence, (2) NP1 commands NP2—i.e., the first noun phrase or sentence node above NP1 is also above NP2—and (3) NP2 is not a pronoun, then NP1 and NP2 cannot be coreferential.

The problem here, of course, is with the notion of “command” used in the second condition of this rule. Consider, for example, the following sentence: “I recognized the spirit in him by the boy’s behavior.” Determining that “him” and “the boy” can be coreferential according to this rule depends on knowing that the prepositional phrase “in him” is attached to the noun phrase “the spirit,” while the prepositional phrase “by the boy’s behavior” is attached to the verb phrase “recognized.” In that case, “him” does not command “the boy,” and so the rule permits them to be coreferential—as indeed they seem to be. If instead, for example, they were both attached to the verb phrase, then “him” *would* command “the boy,” and coreference would be blocked.⁶ However, since prepositional phrase attachment depends on semantic and contextual information, on Marcus *et al.*'s account this determination would not be made by the syntactic analyzer. Thus, if such syntactic rules for pronoun reference were in fact to be applied in understanding, that could only occur *after* semantic and contextual information had been employed to recover sufficiently explicit information about the structure of the input utterance. Indeed, on this account, such syntactic rules would not even seem to be within the province of the syntactic

⁵Marcus references Lasnik's (1976) formulation of this rule. However, the original insight that “precede and command” relations might play a role in explaining pronominal reference is due to Langacker (1966).

⁶This mistaken reading corresponds, as a matter of fact, to what Marcus *et al.* call the “default” interpretation of their representation.

BIRNBAUM

processor itself. Alternatively, one might consider using the above pronominal reference rule to help determine the appropriate syntactic analysis of an input. That, however, would depend on a prior decision as to whether the noun phrases in question should be viewed as coreferential or not—and *that* decision could only be made using semantic and contextual information.

Now the fact is that I don't believe that these rules are completely correct, or that a purely syntactic account of the phenomena in question is actually possible. (A convincing critique of such claims can, in any event, be found in Bolinger, 1979.) Still, I am not sure that I would go so far as to say that syntax plays no role in the matter, and unless Marcus is prepared to say that that is the case, he must concede either (1) that sophisticated syntactic knowledge must exist and be applied outside of the syntactic processor, (2) that semantic and contextual preferences about coreference actively constrain syntactic representations, rather than the other way around as Marcus (1984) contends, or else (3) he must abandon determinism. It is not clear, under these circumstances, what the claim of syntactic autonomy amounts to.

THE NEXT MOVE

It seems clear that job of developing models of language analysis capable of supporting an empirically significant claim of autonomous syntactic processing is far more difficult than has been recognized by proponents of such claims. In particular, the use of vague syntactic representations is not the panacea that it might at first appear to be. One possible alternative is Tomita's (1985) proposal to use a "shared-packed forest" representation as the output of a syntactic analyzer. This proposal shares with Marcus *et al.*'s the idea of an output that represents a set of possible syntactic descriptions rather than just one, but differs radically in the way that it does so. Rather than using a vague description corresponding to a set of possible syntactic analyses, a shared-packed forest representation encodes all of the possible interpretations of an ambiguous sentence in an extremely compact and efficient way. Although Tomita's overall approach has severe limitations as currently formulated—in particular, from both a computational and a linguistic perspective, the restriction to context-free grammars, and from a psychological perspective, the difficulty of accounting for garden path sentences in a model that carries forward all possible analyses⁷—his representation seems superior to Marcus *et al.*'s as far as the application of the coreference rule described above is concerned. That is, it seems quite feasible to apply such rules to a shared-packed forest representation.⁸

The larger point remains, however, that rules of this sort cannot unambiguously be applied as filters to narrow the choices available to semantic and contextual analysis, as Marcus (1984) contends. What they might decide depends on what the correct syntactic analysis is determined to be, and this determination will involve arbitrarily complex inferential processing. Such a decision may, in a given case, make use of potential syntactic constraints on coreference; but then again, it may not. The picture that emerges is one in which syntactic, semantic, and contextual constraints are exploited in an interleaved manner.

⁷Although, in fairness to Tomita, it is not at all clear how Marcus *et al.*'s theory could account for them either.

⁸There is however, some question as to the functional suitability of such representations in a larger language processing system: At least one project attempting to use Tomita's approach was led to discard local ambiguity packing, the most novel aspect of his representation, as being more trouble than it was worth (A. Fano, personal communication).

BIRNBAUM

Acknowledgments: This paper is based on portions of my Ph.D. thesis (Birnbaum, 1986). I owe a special debt to my advisor, Roger Schank, as well as the other members of my committee, Robert Abelson, Drew McDermott, and Elliot Soloway, for their help and support. Chris Riesbeck first taught me about language analysis. I also thank Gregg Collins, Ed Hovy, Steve Lytinen, Mitch Marcus, Charlie Martin, Rod McGuire, and Mallory Selfridge for many spirited discussions on these topics. Special thanks to Eric Jones for helping me to understand the implications of Tomita's "shared-packed forest" representation. This work was supported in part by the Defense Advanced Research Projects Agency, monitored by the Office of Naval Research under contract N0014-85-K-0108.

REFERENCES

- Barton, E., and Berwick, R. 1985. Parsing with assertion sets and information monotonicity. *Proceedings of the Ninth IJCAI*, Los Angeles, CA, pp. 769-771.
- Birnbaum, L. 1985. Lexical ambiguity as a touchstone for theories of language analysis. *Proceedings of the Ninth IJCAI*, Los Angeles, CA, pp. 815-820.
- Birnbaum, L. 1986. Integrated processing in planning and understanding. Research report no. 489, Yale University, Dept. of Computer Science, New Haven, CT.
- Bobrow, D., and Fraser, B. 1969. An augmented state transition network analysis procedure. *Proceedings of the First IJCAI*, Washington, DC, pp. 557-567.
- Bolinger, D. 1979. Pronouns in discourse. In T. Givon, ed., *Syntax and Semantics, Vol. 12: Discourse and Syntax*, Academic Press, New York, pp. 289-309.
- Colmerauer, A. 1978. Metamorphosis grammars. In L. Bolc, ed., *Natural Language Communication with Computers*, Springer, Berlin, pp. 133-189.
- Crain, S., and Steedman, M. 1985. On not being led down the garden path: The use of context by the psychological parser. In D. Dowty, L. Karttunen, and A. Zwicky, eds., *Natural Language Parsing: Psychological, Computational, and Theoretical Perspectives*, Cambridge University Press, Cambridge, England.
- Langacker, R. 1966. On pronominalization and the chain of command. In W. Reibel and S. Schane, eds., *Modern Studies in English*, Prentice Hall, Englewood Cliffs, NJ.
- Lasnik, H. 1976. Remarks on coreference. *Linguistic Analysis*, vol. 2, no. 1.
- Marcus, M. 1980. *A Theory of Syntactic Recognition for Natural Language*. MIT Press, Cambridge, MA.
- Marcus, M. 1984. Some inadequate theories of human language processing. In T. Bever, J. Carroll, and L. Miller, eds., *Talking Minds: The Study of Language in Cognitive Science*, MIT Press, Cambridge, MA, pp. 253-278.
- Marcus, M., Hindle, D., and Fleck, M. 1983. D-theory: Talking about talking about trees. *Proceedings of the 21st ACL Conference*, Cambridge, MA, pp. 129-136.
- Pereira, F., and Warren, D. 1980. Definite clause grammars for language analysis—A survey of the formalism and a comparison with augmented transition networks. *Artificial Intelligence*, vol. 13, pp. 231-278.
- Schank, R., and Birnbaum, L. 1984. Memory, meaning, and syntax. In T. Bever, J. Carroll, and L. Miller, eds., *Talking Minds: The Study of Language in Cognitive Science*, MIT Press, Cambridge, MA, pp. 209-251.
- Thorne, J., Bratley, P., and Dewar, H. 1968. The syntactic analysis of English by machine. In D. Michie, ed., *Machine Intelligence, Vol. 3*, American Elsevier, New York, pp. 281-309.
- Tomita, M. 1985. An efficient context-free parsing algorithm for natural languages. *Proceedings of the Ninth IJCAI*, Los Angeles, CA, pp. 756-764.
- Woods, W. 1970. Transition network grammars for natural language analysis. *Communications of the ACM*, vol. 13, pp. 591-606.