

Failure-Driven Acquisition of Figurative Phrases by Second Language Speakers*

Uri Zernik
Michael G. Dyer

Artificial Intelligence Laboratory
Computer Science Department
University of California
Los Angeles, CA 90024

Abstract

The problem of manually modifying the lexicon appears with any natural language processing computer program. Ideally, a program should be able to acquire new lexical entries from context, the way people learn. We address the problem of acquiring entire phrases, specifically figurative phrases, through augmenting a phrasal lexicon. We show that idiosyncratic behavior of certain phrases can be encoded in the lexicon only by modeling the learning process. We also show how metaphor mappings are acquired through parsing. The acquisition of novel figurative phrases, encountered in context, involves three problems: First, since all the constituents of the phrase are known, the existence of a novel phrase needs to be detected. Second, the scope and the generality of the linguistic pattern of the new phrase needs to be determined. Third, the meaning of the phrase needs to be extracted from the context. Our model is based on second language speakers' behavior and observation of their errors. We have designed and implemented a program called RINA which receives new figurative phrases in context and through the application of a sequence of failure-driven rules, creates and refines both the patterns and the concepts which hold syntactic and semantic information about phrases.

1. Introduction

The lexical approach to language processing [Becker75, Searle79, Bresnan82, Pawley83, Fillmore84] emphasizes the role of the lexicon as a knowledge source. Rather than maintaining a single "generic" lexical entry for each word e.g., *take*, the lexicon contains many phrases, e.g., *take on*, *take to the streets*, *take to swimming*, *take over*, etc. This approach proves effective in parsing and in generation [Wilensky84]. However a huge lexicon must be acquired, which cannot be done manually, especially when considering subtle phrase meanings and idiosyncratic behavior of phrases. Moreover, we show in this paper that there are phrases whose behavior can be captured and encoded in the lexicon only by modeling the acquisition process. Therefore, for phrase-based programs to communicate effectively in natural language they must be able to augment their own lexicon, by simulating human learning. The program RINA [Zernik85] which is used as an experimental model, simulates a second language speaker learning new phrases.

Learning Figurative Phrases

Within language learning, our task domain is verb-phrase acquisition. The modeled phenomenon is described in the dialogues below. In the first dialogue, RINA is introduced to an unknown phrase: *take on*. The words *take* and *on* are familiar to RINA, who also remembers the biblical story of David and Goliath. RINA, modeling a language learner, interacts with a native speaker as follows:

David vs. Goliath

Native:	Remember the story of David and Goliath? David took on Goliath.
Learner:	David took Goliath somewhere?
Native:	No. David took on Goliath.
Learner:	He took on him. He won the fight?

* This work was made possible in part by a grant from the Keck Foundation.

Native: No. He took him on. David attacked him.
Learner: He took him on. He accepted the challenge?
Native: Right.

Native: Here is another story. John took on the third exam question.
Learner: He took on a hard problem.

Another dialogue involves *put one's foot down*. Again, the phrase is unknown while its constituents are known:

Going Punk

Native: Jenny wanted to go punk. but her father put his foot down.
Learner: He moved his foot down? It does not make sense.
Native: No. He put his foot down.
Learner: He put his foot down. He refused to let her go punk.

A figurative phrase such as: *put one's foot down* is a linguistic pattern whose associated meaning cannot be produced from the composition of its constituents. Indeed, an interpretation of the phrase based on the meanings of its constituents often exists, but it carries a different meaning. The fact that this literal interpretation of the figurative phrase exists is a misleading clue in learning. Furthermore, the learner may not even notice that a novel phrase has been introduced since she is familiar with *down* as well as with *foot*. Following Becker [Becker75], we describe a space of phrases ranging in their generality from fixed proverbs: *charity begins at home* through idioms: *lay down the law* and phrasal verbs: *put up with one's spouse*, *look up the name*, to literal verb phrases such as: *sit on the chair*. He suggested employing a *phrasal lexicon* to capture this entire range of linguistic structures.

Issues in Phrase Acquisition

Four issues must be addressed when learning phrases in context.

- (1) **Detecting fallures:** What are the indications that the initial interpretation of the phrase: *take him on* as "to take a person to a location" is incorrect? Since all the words in the sentence are known, the problem is detected both as a *conceptual discrepancy* (why would he take his enemy anywhere?) and as a *syntactic failure* (the expected location of the assumed physical transfer is missing).
- (2) **Determining scope and generality of patterns:** The linguistic pattern of a phrase may be perceived by the learner at various levels of *generality*. For example, in the second dialogue, incorrect generalizations could yield patterns accepting sentences such as:
Her boss put his left foot down.
He moved his foot down.
They put down their feet.
He put it.
A decision is also required about the *scope* of the pattern. For instance, the scope of the pattern in *John put up with Mary* could be (1) *?x:person put:verb up* where *with* is associated with *Mary* or (2) *?x:person put:verb up with ?y:person*, where *with* is associated with *put up*. This issue is described in greater detail in [Zernik85].
- (3) **Finding appropriate meanings:** The conceptual meaning of the phrase must be extracted from the context which contains many concepts, both appropriate and inappropriate for hypothesis formation. Thus there must be strategies for focusing on appropriate elements in the context.
- (4) **Example generation:** The learner must generate examples to convey her hypothesis about the phrase. Memory organization is required to allow (1) easy access to episodic examples and (2) limit the number of episodes accessible for each phrase.

2 Our Approach and Its Background

Past work in language learning emphasized either learning of linguistic patterns or learning of conceptual representations. There are three models for learning linguistic patterns:

- PST [Reeker76] operated by GPS principles [Newell57] and similarly used a table of *difference-action* pairs. PST learned grammar by acting upon differences between the input sentence and an internally generated sentence. Six types of differences were classified and the detection of a difference which belonged to a class caused the associated alteration of the grammar.
- LAS [Anderson77] learned ATNs (Augmented Transition Networks) from sample sentence/meaning pairs. LAS presented one element in a larger cognitive model which accounted for general human inference and memory access. This work intended to demonstrate that language learning could be modeled using general learning principles.
- AMBER [Langley82] modeled learning of basic sentence structure and function words. The learning process was directed by mismatches between input sentences and sentences generated by the program. Learning involved recovery from both *errors of omission* (omitting a function word such as *is* and *the* in *daddy bouncing ball*) and *errors of commission* (producing *daddy is lik¹ing dinner*). Like LAS, AMBER's main thrust was to apply general learning principles in language learning.

On the other hand, two models emphasized learning conceptualizations:

- FOUL-UP [Granger77] learned meanings of single unknown words from context. The meaning was extracted from the script [Schank77] which provided the context. A typical learning situation was *The car was driving on Hwy 66, when it swiveled off the road. The unknown verb was guessed from the \$accident script.* FOUL-UP introduced three important elements: (1) Learning was invoked by parsing failures. However, there was only one possible failure—the absence of a word in the lexicon—thus failure analysis was not required. (2) Word meanings were figured out from currently active scripts. (3) Linguistic clues, such as preposition senses, took part in forming meanings.
- CHILD [Selfridge80] modeled a one-year old child learning native language. At that age, concepts rather than language word-order conventions account for comprehension. A sentence such as *Joshua, put the ball in the box* is understood from the conceptual relationships and conceptual clues. Thus CHILD was able to learn basic word meanings starting only with a minimal linguistic knowledge. CHILD introduced heuristics to identify the unknown word in a sentence and to identify the intended concept in a context. Learning was accomplished by associating the new word with that concept.

The integration of these two aspects is at the focus of our approach. For each new phrase, RINA acquires both the linguistic pattern and the conceptual meaning. It turns out that these two aspects are not independent: For example, learning the conceptual meaning of *take on* depends on knowledge of the words and their combination, while learning the linguistic pattern of that phrasal verb depends on the concepts involved.

Literal or Figurative, Dead or Alive

Gibbs [Gibbs84] has addressed the issue of figurative-phrase comprehension, arguing against Searle's notion of *literal interpretation*. Searle [Searle79] had proposed the existence of literal interpretation as a default comprehension mechanism. According to Searle, in comprehension of figurative phrases (speech acts such as that of *Can you please pass the salt?*), people try literal interpretation ("are you able to do that?") first and only when this attempt fails do they resort to the intended figurative interpretation ("pass it, please!"). Gibbs proved using experimental data that people spend no more time processing figurative phrases than they do processing analogous literal phrases. Thus, he claims, there is no extra effort in processing these phrases as implied by Searle. However, Gibbs does not propose any constructive model for the interpretation of figurative phrases. Moreover, he obviously concerns his discussion merely with the interpretation of *dead phrases*, namely *familiar* figurative phrases, while he does not explain how people interpret novel phrases—phrases they have not encountered before. Our contention is twofold. First, a uniform mechanism, the phrasal lexicon, accounts for comprehension of figura-

tive as well as literal phrases, provided that the phrase is already known. Second, an unknown phrase requires a special treatment which actually constitutes the learning of that new phrase.

Regularity and Idiomaticity in Phrases

Idiosyncratic behavior of phrases is difficult to capture in the lexicon. For example, read the next two sentences:

- Peace was struck between Israel and Egypt. The hatchet was buried.
- Finally, death prevailed. The bucket was kicked.

As opposed to the second phrase *kick the bucket*, the first phrase *bury the hatchet* can take the passive voice and still maintain its figurative meaning. What is the reason for this difference, and how can it be predicted for idiomatic phrases in general? Fillmore, Kay and O'Connor [Fillmore84] address the problematic behavior of idiomatic phrases, classifying them into categories according to the knowledge required for the understanding of each idiom. Our contention is that the only way to predict phrase behavior is by modeling the learning process.

Modeling Second Language Acquisition

Learning in general, and specifically learning linguistic knowledge, is an ongoing process. Two groups in particular experience extensive language learning: children learning native language and adults learning a second language [Hatch83, Gasser85, Ulm75]. Adults, as opposed to children, may augment their linguistic knowledge while, to a large extent, maintaining otherwise unchanging world knowledge. Three aspects of second language acquisition are investigated in our research:

- (1) The various types of errors committed: acquiring incorrect concepts, acquiring incorrect linguistic patterns and performing incorrectly while having the correct linguistic knowledge.
- (2) The processes underlying these errors (observing errors is the only way to expose these processes).
- (3) Strategies for error-recovery based on failure-analysis.

The implications of this study, therefore, are not confined only to second language speakers. Rather, observing second language speakers may reveal general learning processes which are used more frequently by second language speakers.

3. The Program

RINA is a computer program designed to learn English phrases. It takes as input English sentences which may include unknown phrases and conveys as output its hypotheses about novel phrases. The program consists of four components:

- (1) **Phrasal lexicon:** This is a list of phrases where each phrase is a declarative *pattern-concept* pair [Wilensky84].
- (2) **Case-frame parser:** In the parsing process, case-frame expectations are handled by spawning *demons* [Dyer83]. The parser detects comprehension failures which are used in learning.
- (3) **Pattern Constructor:** Learning of phrase patterns is accomplished by analyzing parsing failures. Each failure situation is associated with a pattern-modification action.
- (4) **Concept Constructor:** Learning of phrase concepts is accomplished by a set of strategies which are selected according to the context.

Schematically, the program receives a sequence of *sentence/context* pairs from which it refines its current *pattern/concept* pair. The pattern is derived from the sentence and the concept is derived from the context. However, the two processes are not independent since the context influences construction of patterns while linguistic clues in the sentence influence formation of concepts.

Phrasal Representation of the Lexicon

RINA uses a declarative phrasal lexicon suggested by Wilensky, where a lexical phrase is a *pattern-concept* pair. RINA's patterns are similar to those in [Arens82]. The notation is explained by three example patterns:

P1:	?x:animate nibble:verb <on ?y:food>
P2:	?x:person take:verb on ?y:person
P3:	?x:person <put:verb foot:body-part down>

Figure 1: The Pattern Notation

- (1) A *token* is a literal unless otherwise specified. For example, *on* is a literal in the patterns above.
- (2) *?x:sort* denotes a variable called *?x* of a *semantic class sort*. *?y:food* above is a variable which stands for references to objects of the *semantic class food*.
- (3) *Act:verb* denotes any form of the *verb syntactic class* with the root *act*. *nibbles:verb* above stands for expressions such as: *nibbled*, *did not nibble*, *will never be nibbled*.
- (4) By default, a pattern sequence does not specify the order of its tokens*. However, based on general English knowledge also provided as patterns, it is the *actor* which is expected to precede the verb in the active form of a sentence.
- (5) Tokens delimited by *<* and *>* are restricted to their specified order. In P1 above, *on* must directly precede *?y:food*.

Each pattern has an associated meaning. Meaning representations are not discussed here; they are specified using Dyer's [Dyer83] *i-link* notation which defines a set of intentional links connecting primitive actions, plans, and goals [Schank77].

Case-Frame Parser

Three tasks in phrasal parsing are identified, ordered by degree of difficulty:

- (1) **Phrase disambiguation:** When more than one lexical phrase matches the input sentence, the phrase intended by the speaker must be selected by the parser. For example, *John took to the streets* could mean: "he led a criminal life", "he demonstrated" or "he was fond of the streets".
- (2) **Ill-formed input comprehension:** Even when an input sentence is not well phrased according to textbook grammar, it may be comprehensible by people and so must be comprehensible to the parser. For example, *John took Mary school* is somehow telegraphic, but comprehensible, while *John took Mary to* conveys only a partial concept.
- (3) **Error Detection:** when the hypothesized phrase does not match the input sentence/context pair, the parser is required to detect the failure and return with an indication of its nature which is geared to the construction of a more accurate hypothesis.

The key element in accomplishing these tasks is the use of case frames for pattern representation, as elaborated in [Zernik85].

* In order to derive phrases with a definite word order, *lexical patterns* must interact with *ordering patterns* [Arens82] which hold general English word-order conventions.

Failure-Driven Pattern Constructor

Learning of phrases is an iterative process. The input is a sequence of sentence-context pairs, through which the program refines its current hypothesis about the new phrase. The hypothesis pertains to both the pattern and the concept of the phrase. The basic cycle in the process is:

- (a) A sentence is parsed on the background of a conceptual context.
- (b) Using the current hypothesis, either the sentence is comprehended smoothly, or a failure is detected.
- (c) The analysis of a failure directs the update of the current hypothesis.

The crucial point in this scheme is to obtain from the parser an intelligible analysis of the failure. As an example, consider this part of the first dialog:

```
1   Program:   He took on him. He won the fight?
2   User:      No. He took him on. David attacked him.
3   Program:   He took him on. He accepted the challenge?
```

The first hypothesis is shown in Figure 4.

pattern:	?x:person take:verb <on ?y:person >
concept:	?x win the conflict with ?y

Figure 4: First Hypothesis

Notice that the preposition *on* is attached to the object *?y*, thus assuming that the phrase is similar to *He looked at Mary* which cannot produce the following sentence: *He looked her at*. This hypothesis underlies Sentence 1 which is erroneous in both its form and its meaning. Two observations should be made by comparing this pattern to Sentence 2:

- The object is not preceded by the preposition *on*.
- The preposition *on* does not precede any object.

These comments direct the construction of the new hypothesis:

pattern:	?x:person take:verb on ?y:person
concept:	?x win the conflict with ?y

Figure 5: Second Hypothesis

where the preposition *on* is taken as a modifier of the verb itself, thus correctly generating Sentence 3. In Figure 5 the conceptual hypothesis is still incorrect and must itself be modified.

4. Strategies for Learning Concepts

In the first dialog (Section 1.1), the context has been presented by the biblical story of *David and Goliath*. RINA comes up with the wrong hypothesis, assuming that *take on* means "to win a fight"*. What is the process underlying this error? We explain such errors by the theory of *story points* [Wilensky82] and salient expectations. Story points encapsulate our impression of the story beyond the level of mundane details (e.g., David's hair color, the weapons involved, the various moves in the fight) which get forgotten with time. When the learner is required to select a conceptualization from the context, she attempts to use the dominant story point. Since "David won the fight in spite of his physical inferiority" is the salient point of the story, this point serves as the learner's first guess.

* This is a real error recorded by the authors from second language speakers.

When the context is given as a story, RINA uses story points to construct the meanings of unknown phrases. For other contexts, where a variety of knowledge structures are dominant we have developed a set of strategies for extraction of meanings from context. Interestingly, all these strategies demonstrate the integration of learning and comprehension. Each learning strategy presented here is derived from an existing comprehension strategy.

Story Points

In the context of *David and Goliath*, RINA selects the concepts from a set of given story points. Initially, the first story point ("David won the fight in spite of his physical inferiority") is selected. The factors in the selection of this point are:

- Prefer a story point which determines the outcome of the situation.
- Use linguistic clues. For instance, the preposition *on* carries the sense of winning (in turn on, put on, hang on, etc., *on* depicts a "positive" state). It matches the outcome of the fight.

Later in the dialogue, when the native speaker negates this interpretation by stating he attacked him, the learner reverts to the second story point (he accepted the challenge) which depicts the correct meaning.

Script-Based Expectations

In a way similar to FOUL-UP [Granger77], RINA is able to associate the new phrase with script-based expectation. For example consider the following text: *After a long illness, his beloved wife finally kicked the bucket.* The first sentence invokes *\$disease* (the disease script) which incorporates a chain of events such as feeling sick, staying in bed, seeing a doctor and eventually recovering or dying. From linguistic clues (e.g., *finally*) and from the causal clue (*a long disease*) RINA selects the last event in the script as the meaning to refer to. It is interesting to notice what triggered the learning process: The failure in the literal interpretation of *the bucket* suggests the existence of an unknown idiomatic phrase. There is a similarity between this method and general expectation-based parsing [Riesbeck74]. In parsing, outstanding expectations are used in disambiguation of meanings. In learning, outstanding expectations are taken as conceptual meanings for novel phrases.

Goal/Plan-Based Expectations

An expectation at the plan/goal level dominates the following text (in Section 1.1, *Going Punk*): *Jenny wanted to go punk, but her father put his foot down.* The context describes a goal conflict between Jenny and her father. Jenny is expected to implement a certain goal (satisfy-vanity) when the word *but* is encountered, indicating that the implementation of the goal is blocked. The act denoted by the phrase is expected to cause this goal to become blocked. However, the metaphor has not yet been resolved.

Metaphor Mapping

A *metaphor* [Carbonell83] is defined as a *mapping* between *patterns* such as *put your foot down*, *shoot one's foot*, *put one's foot in one's mouth*, *climb the walls*, etc., and specific *episodes* in memory. In parsing an already familiar metaphor, the lexicon provides the mapping from the elements of the pattern to the elements of the episode. However, while the mapping is still unknown, the episode cannot be accessed from the sentence itself since "putting one's foot on the floor" does not necessarily convey resistance. The episode may be accessed only by indirect memory search through links from affects and goal/plan situations. For example, the process in *Going Punk* proceeds as follows: Since the act (moving one's foot down) does not conform with the goal/plan situation, and the reference (*foot*) cannot be resolved in the context, a search for a metaphor is initiated. A link is found from the goal-conflict situation and the resistance stance to the episode of stamping one's foot. The phrase is interpreted successfully in that episode, and as a result, the mapping between the pattern and the episode is established in the lexicon.

Generalizing Word Meanings

In idioms such as: *They buried the hatchet*, *He threw the book at him*, and *He laid down the law*, the meaning is constructed by generalizing single words' meanings. After a long dispute, the couple buried the hatchet. The literal interpretation of the action does not make sense (burying some tool in the ground), and no episode is found to constitute a metaphor. Therefore, a sequence of generalizations is initiated:

hatchet -> weapon -> fight-plan
bury -> disenable use -> cease-a-plan

Through these generalizations, the act is realized as the cease-a-plan of the fight-plan which finally relates to the context. This process is triggered by sensing that the word *hatchet* is *too specific* (according to Rosch's [Rosch78] *basic-level* principle of categorization). Such a word is expected to appear in a specific script, otherwise it suggests a figurative use.

5. Connotations

RINA's lexicon indexes episodic as well as generic elements. Indexed episodes which facilitate example-generation are built up in the learning process. RINA generates sentences which exemplify the correct use (or the current hypothesis) of newly acquired phrases. For instance, RINA could generate: *The Lakers took on The Celtics*, using a familiar episode as an example of the phrase *take on*. Example generation is the preferred way for people to discuss phrases (rather than speak in terms of syntax and semantics). Upon request, a person may easily generate a few examples of correct phrase use:

David took on Goliath.
My brother took on a new job.
The Lakers took on The Celtics.

relating to specific episodes. However, beyond a limited set of examples, generation of additional examples becomes more difficult. Computationally, example generation seems to be a complex task. In a database system, example generation for a phrase would require scanning the entire database for an appropriate situation which satisfies the constraints imposed by the meaning of the phrase.

Connotations Indexed by Phrases

RINA simulates this phenomenon in terms of memory organization. A phrase in the lexicon is a *pattern-concept* pair. However, in addition, a phrase accumulates links to episodes in which RINA has encountered the phrase. When required, RINA generates an example by selecting an episode which is linked to that phrase. How are links to episodes created? If on each encounter with the phrase *take on*, an episode were linked to the phrase, then the linkage system would grow out of bounds. Therefore, a new link is created only in situations where a learning effort is required. Thus only one episode of a kind is linked, avoiding redundant links. An illustrative learning scenario for *take on* is:

Native: *My brother took on a new job.*
Learner: <link phrase to episode *career-challenge1*>

(The learner is not familiar with the phrase. She learns the phrase, assuming that it means to start a fight and she links it to the episode.)

Native: *The Lakers took on The Celtics.*
Learner: <link phrase to episode *sporting-event1*>

(Again a new sense of the phrase is learned while the episode is linked.)

Native: *The Redskins took on The Whiteskins.*
Learner: <no linking>

(This time there is no need to learn a new sense. The sentence was comprehended using existing phrases.) Following this scenario, the learner can generate only two examples for *take on*, which are *career-challenge1* and *sporting-event1*. The third episode is not indexed to the phrase.

6. Conclusions

The issue addressed in this paper is the construction of the phrasal lexicon. As opposed to other systems (PHRAN [Arens82] and PHRED [Jacobs85], for example) where the lexicon is constructed manually, in RINA new phrases are acquired through learning. This is significant not only for robustness and flexibility, but also for the correct encoding of lexical phrases. For example, observe phrases such as *bury the hatchet* and *kick the bucket*. On the outset, it is difficult to determine that the first phrase takes the passive voice while the second one does not. (How can the systems programmer, who sets up the lexicon, predict the behavior of other such idiomatic phrases?) Moreover, how can such information be encoded in the lexicon? One way is to mark the phrase *kick the bucket* by the *syntactic* feature: **does not take the passive voice**. Our contention is that such a feature is *conceptual* rather than syntactic and that word-order restrictions actually reflect phrase concepts. The methodology for acquiring this feature is through learning: The first phrase is acquired through word-sense generalization (see Section 4), where the entire phrase has an associated concept, but each individual word also stands for a certain concept (specifically, *hatchet* stands for "war"). Thus, the passive voice serves a communicative discourse function. However, in the second phrase which is acquired as a whole unit, the single words do not stand for concepts (on the contrary, the phrase was actually acquired by noticing that neither *bucket* nor *kick* could be interpreted in the context). Not only does the passive voice serve no function, but it is also misleading, causing the listener to interpret the words literally, and search for their individual meanings.

We have shown in this paper a computational method for metaphor resolution. A metaphor is defined as a mapping from a pattern to an episode. Familiar metaphors are comprehended through the mappings themselves which are provided by the phrasal lexicon. A new mapping is constructed when the phrase is heard in context. The successful construction depends on: (1) the accessibility of a metaphor episode through salient elements in the context, and (2) the interpretation of the phrase in the episode which establishes the mapping of pattern elements to episode concepts.

7. Future Work

We have shown the steps in learning new phrases: a novel phrase is detected, and its pattern is shaped by the given sentences: detected comprehension failures cause pattern modifications. The concept of the phrase is formed by (a) the given context, (b) linguistic clues, (c) generalized word meanings, and (d) by metaphor mappings. In the future we intend to investigate more concept-forming strategies and focus on generalization of concepts. For example, two different episodes are associated with *take on*:

David took on Goliath.

He took on a new job.

The associated concepts extracted from the corresponding contexts are:

"deciding to fight"

"undertaking a responsibility"

The dilemma a learner is facing is whether there is a general common concept which encompasses these two concepts: do they share one lexical entry or are they two separate entries? In this task, similar to learning linguistic patterns, people make errors by overgeneralizing and by undergeneralizing. Therefore, analysis of erroneous hypotheses both for linguistic patterns and for conceptual meanings, accounts for learning.

References

- [Anderson77] Anderson, J. R., "Induction of Augmented Transition Networks," *Cognitive Science* 1, pp.125-157 (1977).
- [Arens82] Arens, Y., "The Context Model: Language Understanding in a Context," in *Proceedings 4th Annual Conference of the Cognitive Science Society*, Ann-Arbor MI (1982).
- [Becker75] Becker, J., "The Phrasal Lexicon," in *Proceedings Interdisciplinary Workshop on Theoretical Issues in Natural Language Processing*, Cambridge MA (1975).
- [Bresnan82] Bresnan, J. and R. Kaplan, "Lexical-Functional Grammar," in *The Mental Representation of Grammatical Relations*, ed. J. Bresnan, MIT Press, MA (1982).
- [Carbonell83] Carbonell, J. and S. Minton, "Metaphor and Common-Sense Reasoning," cmu-cs-83-110, Carnegie Mellon University, Pittsburgh PA (1983).

- [Dyer83] Dyer, Michael G., *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*, MIT Press, Cambridge, MA (1983).
- [Fillmore84] Fillmore, C., P. Kay, and M. O'Connor, "Regularity and Idiomaticity in Grammar," Cognitive Science Working Paper, University of California (1984).
- [Gasser85] Gasser, M., "Second Language Production: Coping with Gaps in Linguistics Knowledge," UCLA-AI-15, LA CA (July 1985).
- [Gibbs84] Gibbs, R., "Literal meaning and psychological theory," *Cognitive Science* 8(3) (1984).
- [Granger77] Granger, R., "FOUL-UP: A Program That Figures Out Meanings of Words from Context," in *Proceedings Fifth IJCAI*, Cambridge, Massachusetts (1977).
- [Hatch83] Hatch, E. M., *Psycholinguistics: A Second Language Perspective*, Newbury House, Rowley MA (1983).
- [Jacobs85] Jacobs, Paul S., "PHRED: A Generator for Natural Language Interfaces," UCB/CSD 85/198, University of California Berkeley, Berkeley CA (Jan 1985).
- [Langley82] Langley, P., "Language Acquisition Through Error Recovery," *Cognition and Brain Theory* 5(3) (1982).
- [Newell57] Newell, A., C. Shaw, and A. Simon, "Preliminary Description of GPS-I," CIP Working Paper 7, Carnegie Institute of Technology, Pittsburgh, PA (1957).
- [Pawley83] Pawley, A. and H. Syder, "Two Puzzles for Linguistic Theory: Nativelike Selection and Nativelike Fluency," in *Language and Communication*, ed. J. C. Richards R. W. Schmidt, Longman, London (1983).
- [Reeker76] Reeker, L. H., "The Computational Study of Language Learning," in *Advances in Computers*, ed. M. Yovits M. Rubinoff, Academic Press, New York (1976).
- [Riesbeck74] Riesbeck, C., "Analysis of Sentences and Context," AI-238, Yale (1974).
- [Rosch78] Rosch, E., "Principles of Categorization," in *Cognition and Categorization*, ed. B. Lloyd, LEA (1978).
- [Schank77] Schank, R. and R. Abelson, *Scripts, Plans, Goals, and Understanding*, LEA, Hillsdale, NJ (1977).
- [Searle79] Searle, J., "Speech Acts and Recent Linguistics," in *Expression and Meaning*, ed. J. Searle, Cambridge University Press, Cambridge (1979).
- [Selfridge80] Selfridge, M., "A Process Model of Language Acquisition," 172, Yale, New Haven CT (1980). Ph.D. Dissertation.
- [Ulm75] Ulm, S., *The Separation Phenomenon in English Phrasal Verbs, Double trouble*, UCLA (1975). M.A. Thesis.
- [Wilensky82] Wilensky, R., "Points: A Theory of Structure of Stories in Memory," pp. 345-375 in *Strategies for Natural Language Processing*, ed. W. G. Lehnert M. H. Ringle, Laurence Erlbaum Associates, New Jersey (1982).
- [Wilensky84] Wilensky, R., Y. Arens, and D. Chin, "Talking to UNIX in English," *CACM* 27(6) (1984).
- [Zernik85] Zernik, U. and M. G. Dyer, "Towards a Self-Extending Phrasal Lexicon," in *Proceedings 29rd Annual Meeting of the Association for Computational Linguistics*, Chicago, Ill. (July 1985).