

Encoding and Acquiring Meanings for Figurative Phrases *

Michael G. Dyer
Uri Zernik

Artificial Intelligence Laboratory
Computer Science Department
3531 Boelter Hall
University of California
Los Angeles, California 90024

Abstract

Here we address the problem of mapping phrase meanings into their conceptual representations. Figurative phrases are pervasive in human communication, yet they are difficult to explain theoretically. In fact, the ability to handle idiosyncratic behavior of phrases should be a criterion for any theory of lexical representation. Due to the huge number of such phrases in the English language, phrase representation must be amenable to parsing, generation, and also to learning. In this paper we demonstrate a semantic representation which facilitates, for a wide variety of phrases, both learning and parsing.

1. Introduction

The phrasal approach to language processing [Becker75, Pawley83, Fillmore86] 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 over*, *take it or leave it*, *take it up with*, *take it for granted*, etc. Although this approach proves effective in parsing and in generation [Wilensky84], there are three problems which require further investigation. First, phrase *interaction*: the lexicon provides representation for single phrases, such as *take to task* and *make up one's mind*. Yet it is required to analyze complex clauses such as *he made up his mind to take her to task*. The problem lies with the way the *meanings* of the two phrases interact to form the compound meaning. Second, phrase *ambiguity*: [Zernik86] phrasal parsing shifts the task from single-word selection to the selection of entire lexical phrases. When a set of lexical phrases appear syntactically equivalent, i.e.: *he ran into a friend*, *he ran into an 1986 Mercedes*, *he ran into the store*, and *he ran into trouble again*, disambiguation must be performed by semantic means. The conditions which facilitate phrase discrimination reside within each lexical entry itself. Third, phrase *idiosyncrasy*: the meaning representation of phrases such as: *lay down the law* vs. *put one's foot down*, must distinguish the special use of each phrase. This paper is concerned in the representation of phrase meanings and the process of acquiring these meanings from examples in context.

* This research was supported in part by a grant from the ITA Foundation.

1.1 The Task Domain

Consider the figurative phrases in the sentences below, as they are parsed by the program RINA [Zernik85a].

- S1: The Democrats in the house **carried the water** for Reagan's tax-reform bill.**
S2: The famous mobster evaded prosecution for years. Finally, they **threw the book** at him for tax evasion.

Depending on the contents of the given lexicon, the program may interpret these sentences in one of two ways. On the one hand, assuming that the meaning of a phrase exists in the lexicon, the program applies that meaning in the comprehension of the sentence. In S1, the program understands that the Democratic representatives did the "dirty" work in passing the bill for Reagan. On the other hand, if the figurative phrase does not exist in the lexicon, an additional task is performed: the program must figure out the meaning of the new phrase, using existing knowledge: First, the meanings given for the single words *carry* and *water* are processed literally. Second, the context which exists prior to the application of the phrase, provides a hypothesis for the formation of the phrase meaning. A dialog with RINA proceeds as follows:

- RINA: They moved water?
User: No. The Democrats carried the water for Reagan.
RINA: They helped him pass the bill?

Thus, RINA detects the metaphor underlying the phrase, and using the context, it learns that *carry the water* means helping another person do a hard job. Consider encounters with three other phrases:

- Jenny wanted to go punk but her father
S3: **laid down the law**.
S4: **put his foot down**.
S5: **read her the riot act**.

In all these cases, it is understood from the context that Jenny's father objected to her plan of going punk (aided by the word *but* which suggests that something went wrong with Jenny's goals). However, what is the meaning of each one of the phrases, and in particular do all these phrases convey identical concepts?

** This sentence was recorded off the ABC television program *Nightline*, December 12, 1985.

1.2 The Issues

In encoding meanings of figurative phrases, we must address the following issues.

Underlying Knowledge

What is the knowledge required in order to encode the phrase *throw the book*? Clearly, this knowledge includes the situation and the events that take place in court, namely the judge punishing the defendant.

The phrase *carry the water*, for example, requires two kinds of knowledge:

- (a) Knowledge about the act of carrying water which can support the analysis of the phrase metaphor.
- (b) Knowledge about general plans and goals, and the way one person agrees to serve as an agent in the execution of the plans of another person. This knowledge supports the analysis of the context.

While the phrases above could be denoted in terms of plans and goals, other phrases, i.e.: *rub one's nose in it*, *climb the walls*, and *have a chip on one's shoulder* require knowledge about emotions, such as *embarrassment* and *frustration*. Unless the program maintains knowledge about *resentment*, the phrase *have a chip on the shoulder*, for example, cannot be represented. Thus, a variety of knowledge structures take place in encoding figurative phrases.

Representing Phrase Meanings and Connotations

The appearance of each phrase carries certain implications. For example, *John put his foot down* implies that John refused a request, and on the other hand, *John read the riot act* implies that he reacted angrily about a certain event in the past. *John gave Mary a hard time* implies that he refused to cooperate, and argued with Mary since he was annoyed, while *John laid down the law* implies that John imposed his authority in a discussion. The representation of each phrase must account for such implications.

Three different phrases in sentences S3-S5 are applied in the same context. However, not any phrase may be applied in every context. For example, consider the context established by this paragraph:

S6: Usually, Mary put up with her husband's cooking, but when he served her cold potatoes for breakfast, she **put her foot down**.

Could the phrase in this sentence be replaced by the other two phrases: (a) *lay down the law*, or (b) *read the riot act*? While understandable, these two phrases are not appropriate in that context. The sentence *she read him the riot act* does not make sense in the context of debating food taste. The sentence *she laid down the law* does not make as much sense since there is no argument between individuals with non-equal authority. Thus, there are conditions for the applicability of each lexical phrase in various contexts. These conditions support phrase disambiguation, and must be included as part of a phrase meaning.

Phrase Acquisition

Phrase meanings are learned from examples given in context. Suppose the structure and meaning of *put one's foot down* is acquired through the analysis of the following sentences:

- S6: Usually, Mary put up with her husband's cooking, but when he served her cold potatoes for breakfast, she **put her foot down**.
- S7: Jenny was dating a new boyfriend and started to show up after midnight. When she came at 2am on a weekday, her father **put his foot down**: no more late dates.
- S8: From time to time I took money from John, and I did not always remember to give it back to him. He **put his foot down** yesterday when I asked him for a quarter.

Since each example contains many concepts, both appropriate and inappropriate, the appropriate concepts must be identified and selected. Furthermore, although each example provides only a specific episode, the ultimate meaning must be generalized to encompass further episodes.

Literal Interpretation

Single-word senses (e.g.: the sense of the particle *into* in *run into another car*), as well as entire metaphoric actions (e.g.: *carry the water* in the Democratic representatives carried the water for Reagan's tax-reform bill) take part in forming the meaning of unknown figurative phrases. Can the meaning of a phrase be acquired in spite of the fact that its original metaphor is unknown, as is the case with *read the riot act* (what act exactly?) or *carry the water* (carry what water)?

2. The Program

The program RINA [Zernik85b] is designed to parse sentences which include figurative phrases. When the meaning of a phrase is given, that meaning is used in forming the concept of the sentence. However, when the phrase is unknown, the figurative phrase should be acquired from the context. The program consists of three components: phrasal parser, phrasal lexicon, and phrasal acquisition module.

2.1 Phrasal Parser

A lexical entry, a *phrase*, is a triple associating a linguistic *pattern* with its *concept* and a *situation*. A clause in the input text is parsed in three steps:

- (1) Matching the phrase *pattern* against the clause in the text.
- (2) Validating in the context the relations specified by the phrase *situation*.
- (3) If both (1) and (2) are successful then instantiating the phrase *concept* using variable bindings computed in (1) and (2).

For example, consider the sentence:

S9: Fred wanted to marry Sheila, but she ducked the issue for years. Finally he put her on the spot.

The figurative phrase is parsed relative to the context established by the first sentence. Assume that the lexicon contains a single phrase, described informally as:

phrase
pattern: Person1 put Person2 on the spot
situation: Person2 avoids making a certain tough decision
concept: Person1 prompts Person2 to make that decision

The steps in parsing the clause using this phrase are:

- (1) The pattern is matched successfully against the text. Consequently, Person1 and Person2 are bound to Fred and Sheila respectively.
- (2) The situation associated with the pattern is validated in the context. After reading the first phrase the context contains two concepts: (a) Fred wants to marry Sheila, and (b) she avoids a decision. The situation matches the input.
- (3) Since both (1) and (2) are successful, then the pattern itself is instantiated, adding to the context:

Fred prompted Sheila to make up her mind.

Phrase *situation*, distinguished from phrase *concept*, is introduced in our representation, since it help solve three problems: (a) in *disambiguation* it provides a discrimination condition for phrase selection, (b) in *generation* it determines if the phrase is applicable, and (c) in *acquisition* it allows the incorporation of the input context as part of the phrase.

2.2 Phrasal Lexicon

RINA uses a declarative phrasal lexicon which is implemented through GATE [Mueller84] using *unification* [Kay79] as the grammatic mechanism. Below are some sample phrasal patterns.

P1: ?x <lay down> <the law>
P2: ?x throw <the book> <at ?y>

These patterns actually stand for the slot fillers given below:

P1:	(subject ?x (class person))
	(verb (root lay) (modifier down))
	(object (determiner the) (noun law))
P2:	(subject ?x (class person))
	(verb (root throw))
	(object ?z (marker at) (class person))
	(object (determiner the) (noun book))

This notation is described in greater detail in [Zernik85b].

2.3 Phrase Acquisition through Generalization and Refinement

Phrases are acquired in a process of hypothesis formation and error correction. The program generates and refines hypotheses about both the linguistic pattern, and the conceptual meaning of phrases. For example, in acquiring the phrase

carry the water, RINA first uses the phrase already existing in the lexicon, but it is too general a pattern and does not make sense in the context.

?x carry:verb ?z:phys-obj <for ?y>

Clearly, such a syntactic error stems from a conceptual error. Once corrected, the hypothesis is:

?x carry:verb <the water> <for ?y>

The meaning of a phrase is constructed by identifying salient features in the context. Such features are given in terms of scripts, relationships, plan/goal situations and emotions. For example, carry the water is given in terms of *agency* goal situation (?x executes a plan for ?x) on the background of *rivalry* relationship (?x and ?y are opponents). Only by detecting these elements in the context can the program learn the meaning of the phrase.

3. Conceptual Representation

The key for phrase acquisition is appropriate conceptual representation, which accounts for various aspects of phrase meanings.

Consider the phrase to throw the book in the following paragraph:

S2: The famous mobster avoided prosecution for years. Finally they **threw the book** at him for tax evasion.

We analyze here the components in the representation of this phrase.

3.1 Scripts

Basically, the figurative phrase depicts the *trial script* which is given below:

(a)	The prosecutor says his arguments to the judge
(b)	The defendant says his arguments to the judge
(c)	The judge determines the outcome, either:
	(1) to punish the defendant
	(2) not to punish the defendant

This script involves a Judge, a Defendant, and a Prosecutor, and it describes a sequence of events. Within the script, the phrase points to a single event, the decision to punish the defendant. However, this event presents only a rough approximation of the real meaning which requires further refinement.

- (a) The phrase may be applied in situations that are more general than the trial script itself. For example:

S10: When they caught him cheating in an exam for the third time, the dean of the school decided to **throw the book** at him.

Although the context does not contain the specific *trial script*, the *social authority* which relates the judge and the defendant exists also between the dean and John.

- (b) The phrase in S2 asserts not only that the mobster was punished by the judge, but also that a certain prosecution strategy was applied against him.

3.2 Specific Plans and Goals

In order to accommodate such knowledge, scripts incorporate specific planning situations. For example, in prosecuting a person, there are three options, a basic rule and two deviations:

- (a) Basically, for each law violation, assign a penalty as prescribed in the book.
- (b) However, in order to loosen a prescribed penalty, mitigating circumstances may be taken into account.
- (c) And on the other hand, in order to toughen a prescribed penalty, additional violations may be thrown in.

In S2 the phrase conveys the concept that the mobster is punished for tax evasion since they cannot prosecute him for his more serious crimes. It is the selection of this particular prosecution plan which is depicted by the phrase. The phrase representation is given below,

```
phrase
pattern    ?x:person throw:verb
           <the book> <at ?y:person>
situation ($trial (prosecution ?x)
           (defendant ?y))
concept    (act (select-plan
               (actor prosecution)
               (plan (ulterior-crime
                     (crime ?c)
                     (crime-of ?y))))))
           (result (thwart-goal
                   (goal ?g)
                   (goal-of ?y)))
```

where *ulterior-crime* is the third prosecution plan above.

3.3 Relationships

The *authority* relationship [Schank78, Carbonell79] is pervasive in phrase meanings, appearing in many domains: judge-defendant, teacher-student, employer-employee, parent-child, etc. The existence of *authority* creates certain expectations; if X presents an *authority* for Y, then:

- (a) X issues rules which Y has to follow.
- (b) Y is expected to follow these rules.
- (c) Y is expected to support goals of X.
- (d) X may punish Y if Y violates the rules in (a).
- (e) X cannot dictate actions of Y; X can only appeal to Y to act in a certain way.
- (f) X can delegate his authority to Z which becomes an authority for Y.

In S10, the dean of the school presents an authority for John. John violated the rules of the school and is punished by the dean. More phrases involving authority are given by the following examples.

- S11: I thought that parking ticket was unfair so I **took it up with** the judge.
- S12: My boss wanted us to stay in the office until 9pm every evening to finish the project on time. Everybody was upset, but nobody **stood up to** the boss.
- S13: Jenny's father **laid down the law**: no more late dates.

The representation of the phrase *take it up with*, for example, is given below:

```
phrase
pattern    ?x:person <take:verb up>
           ?z:problem <with ?y:person>
situation (authority (high ?y) (low ?x))
concept    (act (auth-appeal(actor ?x)
               (to ?y) (object ?z))
           (purpose (act (auth-decree
                         (actor ?y)
                         (to ?x)
                         (object ?z))))
           (result (support-plan
                   (plan-of ?x))))
```

The underlying situation is an *authority* relationship between X and Y. The phrase implies that X appeals to Y so that Y will act in favor of X.

3.4 Abstract Planning Situations

General planning situations, such as *agency*, *agreement*, *goal-conflict* and *goal-coincidence* [Wilensky83] are addressed in the examples below.

S1: The Democrats in the house **carried the water** for Reagan in his tax-reform bill.

The phrase in S1 is described using both *rivalry* and *agency*. In contrast to expectations stemming from *rivalry*, the actor serves as an agent in executing his opponent's plans. The representation of the phrase is given below:

```
phrase
pattern    ?x:person carry:verb
           <the water ?z:plan> <for ?y:person>
situation (rivalry (actor1 ?x) (actor2 ?y))
concept    (agency (agent ?x)
           (plan ?z)
           (plan-of ?y))
```

Many other phrases describe situations at the abstract goal/plan level. Consider S14:

S14: I planned to do my CS20 project with Fred. I **backed out of** it when I heard that he had flunked CS20 twice in the past.

Back out of depicts an agreed plan which is cancelled by one party in contradiction to expectations stemming from the agreement.

S15: John's strongest feature in arguing is his ability to **fall back on** his quick wit.

Fall back on introduces a *recovery* of a goal through an alternative plan, in spite of a *failure* of the originally selected plan.

S16: My standing in the tennis club deteriorated since I was **bogged down with** CS20 assignments the whole summer.

In *bog down*, a goal competition over the actor's time exists between a major goal (tennis) and a minor goal (CS20). The major goal fails due to the efforts invested in the minor goal.

3.5 Emotions and Attitudes

In text comprehension, emotions [Dyer83, Mueller85] and attitudes are accounted for in two ways: (a) they are generated by goal/planning situations, such as goal failure and goal achievement, and (b) they generate goals, and influence plan selection. Some examples of phrases involving emotions are given below. *Humiliation* is experienced by a person when other people achieve a goal which he fails to achieve. The phrase in S17 depicts humiliation which is caused when John reminds the speaker of his goal situation:

S17: I failed my CS20 class. My friend John rubbed my nose in it by telling me that he got an A+.

Resentment is experienced by a person when a certain goal of his is not being satisfied. This goal situation causes the execution of plans by that person to deteriorate. The phrase in S18 depicts such an attitude:

S18: Since clients started to complain about John, his boss asked him if he had a chip on his shoulder.

Embarrassment is experienced by a person when his plan failure is revealed to other people. The phrase in S19, depicts embarrassment which is caused when a person is prompted to make up his mind between several bad options.

S19: Ted Koppel put his guest on the spot when he asked him if he was ready to denounce apartheid in South Africa.

In all the examples above, it is not the emotion itself which is conveyed by the phrase. Rather, the concept conveys a certain goal situation which causes that emotion. For example, in S20 (rub one's nose) a person does something which causes the speaker to experience *humiliation*.

4. Learning Phrase Meanings

Consider the situation when a new phrase is first encountered by the program:

User: The Democrats in the house carried the water for Reagan's tax-reform bill.

RINA: They moved water?

User: No. They carried the water for him.

RINA: They helped him pass the bill.

Three sources take part in forming the new concept, (a) the linguistic clues, (b) the context, and (c) the metaphor.

4.1 The Context

The context prior to reading the phrase includes two concepts:

- (a) Reagan has a goal of passing a law.
- (b) The Democrats are Reagan's rivals—they are expected to thwart his goals, his legislation in particular.

These concepts provide the phrase *situation* which specifies the context required for the application of the phrase.

4.2 The Literal Interpretation

The literal interpretation of carried the water as "moved water" does not make sense given the goal/plan situation in the context. As a result, RINA generates the literal interpretation and awaits confirmation from the user. If the user repeats the utterance or generates a negation, then RINA generates a number of utterances, based on the current context, in hypothesizing a novel phrase interpretation.

4.3 The Metaphor

Since the action of moving water does not make sense literally, it is examined at the level of plans and goals: Moving water from location A to B is a low-level plan which supports other high-level plans (i.e., using the water in location B). Thus, at the goal/plan level, the phrase is perceived as: "they executed a low-level plan as his agents" (the *agency* is suggested by the prepositional phrase: for his tax-reform bill; i.e., they did an act for his goal). This is taken as the phrase *concept*.

4.4 The Constructed Meaning

The new phrase contains three parts:

- (a) The phrase *pattern* is extracted from the example sentence:

?x carry:verb <the water> <for ?y>

- (b) The phrase *situation* is extracted from the underlying context:

(rivalry (actor1 ?x) (actor2 ?y))

- (c) The phrase *concept* is taken from the metaphor:

(plan-agency (actor ?x) (plan ?z) (plan-of ?y))

Thus, the phrase means that in a rivalry situation, an opponent served as an agent in carrying out a plan.

5. Future Work and Conclusions

The *phrasal* approach elevates language processing from interaction among single words to interaction among entire phrases. Although it increases substantially the size of the lexicon, this *chunking* simplifies the complexity of parsing since clauses in the text include fewer modules which interact in fewer ways. The phrasal approach does reduce the power of the program in handling non-standard uses of phrases. For example, consider the situation where a mobster kidnaps a judge, points the gun at him, and says: No funny book you could throw at me now would do you any good!*. Our current parser would certainly fail in matching the syntactic pattern and inferring the ironic meaning. The analysis of such a sentence would require that the program *associate* the two existing phrases, the general throw something and the figurative throw the book, and make inferences about the pun meant by the mobster. Such examples show that it is difficult to capture human behavior through a single parsing paradigm.

* This example is attributed to an anonymous referee.

Parsing text is a futile task unless it addresses the ultimate objective of language processing, namely mapping text into conceptual representation. To this end, we have shown the structure of a lexicon which provides the association between syntactic patterns with their semantic concepts. However, due to the huge size of the English language, not all phrases can be given at the outset. A parsing program is required to handle unknown phrases as they are encountered in the text. In RINA we have shown how new phrases can be acquired from examples in context.

Phrase acquisition from context raises questions regarding the volume of knowledge required for language processing. A phrase such as *throw the book* requires highly specialized knowledge involving sentencng strategies in court. Now, this is only one figurative phrase out of many. Thus, in order to handle figurative phrases in general, a program must ultimately have access to all the knowledge of a socially mature person. Fortunately, learning makes this problem more tractible. In the process of phrase acquisition, phrase meaning is elevated from the specific domain in which the phrase has originated to a level of abstract goal situations. For example, once *throw the book* is understood as the act of authority-decree, then knowledge of the trial situation no longer needs to be accessed. The phrase is well comprehended in other domains: *my boss threw the book at me, his parents threw the book at him, her teacher threw the book at her, etc.* At that level, a finite number of goal situations can support the application of figurative phrases across a very large number of domains.

References

- [Becker75] Becker, Joseph D., "The Phrasal Lexicon," pp. 70-73 in *Proceedings Interdisciplinary Workshop on Theoretical Issues in Natural Language Processing*, Cambridge, Massachusetts (June 1975).
- [Carbonell79] Carbonell, J. G., "Subjective Understanding: Computer Models of Belief Systems," TR-150, Yale, New Haven CT (1979). Ph.D. Dissertation.
- [Dyer83] Dyer, Michael G., *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*, MIT Press, Cambridge, MA (1983).
- [Fillmore86] Fillmore, C., P. Kay, and M. O'Connor, *Regularity and Idiomaticity in Grammatical Constructions: The Case of Let alone*, UC Berkeley, Department of Linguistics (1986). Unpublished Manuscript.
- [Kay79] Kay, Martin, "Functional Grammar," pp. 142-158 in *Proceedings 5th Annual Meeting of the Berkeley Linguistic Society*, Berkeley, California (1979).
- [Mueller84] Mueller, E. and U. Zernik, "GATE Reference Manual," UCLA-AI-84-5, Computer Science, AI Lab (1984).
- [Mueller85] Mueller, E. and M. Dyer, "Daydreaming in Humans and Computers," in *Proceedings 9th International Joint Conference on Artificial Intelligence*, Los Angeles CA (1985).
- [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).
- [Schank78] Schank, R. and J. Carbonell, "The Gettysburg Address: Representing Social and Political Acts," TR-127, Yale University, Department of Computer Science, New Haven CT (1978).
- [Wilensky83] Wilensky, Robert, *Planning and Understanding*, Addison-Wesley, Massachusetts (1983).
- [Wilensky84] Wilensky, R., Y. Arens, and D. Chin, "Talking to UNIX in English: an Overview of UC," *Communications of the ACM* 27(6), pp.574-593 (June 1984).
- [Zernik85a] Zernik, Uri and Michael G. Dyer, "Learning Phrases in Context," in *Proceedings The 3rd Machine Learning Workshop*, New-Brunswick NJ (June 1985).
- [Zernik85b] Zernik, Uri and Michael G. Dyer, "Towards a Self-Extending Phrasal Lexicon," in *Proceedings 23rd Annual Meeting of the Association for Computational Linguistics*, Chicago IL (July 1985).
- [Zernik86] Zernik, U. and M. G. Dyer, "Disambiguation and Acquisition using the Phrasal Lexicon," in *Proceedings 11th International Conference on Computational Linguistics*, Bonn Germany (1986).