Computer Science Department Technical Report
University of California
Los Angeles, CA 90024-1596

SYMBOL PROCESSING,CONNECTIONISM AND ARTIFICIAL
NEURAL NETWORKS FOR HIGH-LEVEL COGNITIVE TASKS:
NATURAL LANGUAGE RESEARCH AT THE UCLA
ARTIFICIAL INTELLIGENCE LABORATORY

Michael G. Dyer, Director   November 1990
                                             CSD-900037

Symbol Processing, Connectionism and Artificial Neural Networks for High-Level Cognitive Tasks: Natural Language Research at the UCLA Artificial Intelligence Laboratory

Michael G. Dyer, Director

January 1990

Technical Report UCLA-AI-90-01

Symbol Processing, Connectionism and Artificial Neural Networks for High-Level Cognitive Tasks:
Natural Language Research at the
UCLA Artificial Intelligence Laboratory

Michael G. Dyer[*]
(UCLA AI Lab Director)

3532 Boelter Hall
Computer Science Department
UCLA, Los Angeles CA 90024

1. Introduction

The field of *natural language processing* (NLP) deals with the construction of computer programs that take natural language text[1] as input and produce natural language text as output. Unlike linguistics, which often is restricted to the nature of grammars, NLP research is mainly concerned with mapping natural language text into conceptual representations (NL analysis or conceptual "parsing"), manipulating those representations (as the result of reasoning, planning, invention, persuasion, learning) and then mapping from conceptual representations back into language (NL generation). By its very nature, NLP research cuts across a great many disciplines, since NL text of just 50 words in length can include:

- argument fragments and editorials,
- literary short stories and newspaper articles,
- headlines and advertisements,
- NL descriptions of mechanical devices,
- zen koans, parables, and religious texts,
- advice columns (e.g. Dear Abby),
- jokes and sarcastic or ironic utterances,
- everyday conversations and other forms of dialog,
- scientific and philosophical articles,
- game descriptions by sports announcers,
- legal case descriptions, and so on.

An attempt to understand mechanical device descriptions, for example, leads the NLP researcher into a the subfield of AI that is called "naive physics", while research in editorial text comprehension leads to theories of belief maintenance, attack/support relations among beliefs, and processes of persuasion.

---

[1] As opposed to acoustic input. Researchers who deal with acoustic analysis are usually referred to as being in the area of "speech processing". NLP research, in constrast, focusses on the conceptual representations built as the result of understanding the semantic content of textual input.

1.1. NLP Tasks

In addition to varying the domains of discourse, NLP research can be differentiated by the nature of the task being addressed. For example, NLP tasks include:

1. *Representation of words and phrases..* The English language contains thousands of words and phrases, each loaded with semantic content. What are the representations of the words "irresponsible", "hit", "milk", "unfortunately", "threaten", "gratitude", or the phrases "speak of the devil", "put one over on", and so on? To understand NL text, a computer program must have representations for the semantic and syntactic content of each word/phrase, along with processes for combining them into larger conceptualizations.

2. *Disambiguation and reference.* Words are highly ambiguous and their meanings can change as context changes. "John smoked the ham" versus "John is a ham." "John smoked the pot" versus "John put the pot on the stove." Interpretations can change: "John put the pot in the dishwasher. It was the nearest place to hide it before the police raided his apartment." Here, pot is initially interpreted as cooking-pot, and then reinterpreted to be marijuana. "John hit Bill and *his* knuckles hurt for a week." John hit Bill and *his* jaw hurt for a week." Reference requires access to world knowledge.

3. *Dynamic inference and application of world knowledge.* Most NL utterances omit a great deal of information and this missing information must be inferred. For instance, in "John picked up the hammer and hit Bill", the sentence does not actually state that John hit Bill *with* the hammer. The statement "John was listening to the guitar" is not actually correct, since one cannot listen to objects; one can only listen to sounds. Unless we realize that John is listening to the *sounds* being made by the guitar, we won't be able to make sense of the sentence that follows, namely, "He cursed when the string broke", and this sentence does not explicitly state that it is a *guitar* string that broke (e.g., why isn't it a kite string?).

4. *Relation of syntax to semantics.* Consider "John hit the hippie with a brick." and "John hit the hippie with long hair." In these examples, semantics and world influence the interpretation of the syntac structure of the sentences. Since long hair is normally attached to hippies' heads, it is most likely not being used as the instrument of the hitting. In contrast, formal (programming) languages are specifically designed to have their semantic content directed by the syntactic structure, not vice versa.

5. *Automatic database construction and learning from NL input.* Comprehension is closely related to learning as the result of being told. When we comprehend, we built conceptual databases "on the fly." Such NL databases are very different from standard relational databases. NL databases consist of conceptual relationships among planners and their beliefs, hopes, emotional states, and the results of their plans and goals. In addition, any NLP system that can comprehend NL text is, by its very nature, learning new information. Thus, NLP cannot be separated from theories of learning.

6. *Memory organization and generalization from interaction with multiple texts.* As text is read, what is maintained in long-term memory and how is new knowledge integrated with old? E.g., what general insights can be distilled from reading many economic sanctions descriptions from around the world?

7. *Question comprehension, question answering and advice giving.* How do we understand a question and search episodic memory to find an answer? How do we decide what is the most appropriate answer? In general, selection of an appropriate answer depends on maintaining a model of the current knowledge state and interests of the individual asking the question.

8. *Planning.* Most short stories, plays and novels are about planners and their goals. Consider the Aesop Fable of "the Fox and the Crow". The planning here is complex, involving the use of flatter and deception to achieve a goal. In stories planners form coalitions and cooperate. They counter-plan and block the actions of other planners. They retaliate and make planning errors. They maintain and manage multiple and often conflicting goals. For a computer program to understand such stories, it must track complex goal/plan structures.

9. *Argumentation, reasoning, and persuasion.* Consider the following fragment:

Korean: The USSR should be condemned for shooting down our KL-10.
Soviet: But it was a spy plane.
Korean: With over 300 passengers aboard?

Would this arguent still be coherent if the Soviet had issued the (technically correct) statement "But it was a *white* plane"? Why is "300 passengers aboard" a better response than "two bathrooms aboard"? How do we recognize that a given statement violates one of our beliefs and how do we generate a rebuttal? Why do certain arguments persuade us to change our beliefs? Many arguments are supported by analogical reasoning, which leads an NL researcher to examine such forms of reasoning.

10. *Generation and invention.* Consider the utterance: "When you've had it you've had it". What is expressed semantically by the repetition of a given phrase and in what circumstances do such repetitions occur? How are thoughts expressed in language? Unlike many linguists (who want to understand syntax first, and then semantics), we take the position at the UCLA AI Lab that a representational system for a vaste number of concepts must be in place *before* the processes of syntactically correct sentence generation can be addressed, mainly because generation is the communication of concepts. Given a large number of concepts, how do we decide how much to say and how to say it? Consider story telling. How do story tellers create the stories they tell? How do they select the point of the story, weave the plot, develop the characters, and maintain story coherence?

1.2. Levels of NLP Research

Finally, NLP research can be differentiated by levels of model building. Some NL researchers build models at the symbolic level and restrict themselves to capturing the functional aspects of NLP. Others attempt to construct models that are as close to neural models as possible, and there are many levels of abstraction in between. At this point, we distinguish 5 levels at which NLP models are being constructed in the UCLA AI lab:

1. *Symbolic and logical/functional models.* At this level, knowledge is represented in terms of symbol structures. Processes of comprehension, planning, argumentation, etc. are modelled in terms of the creation and modification of symbolic structures.

2. *Marker-passing networks.* These networks implement inferences via propagation in parallel of pointer-like structures. Marker-passing networks allow one to explore parallel algorithms for NL comprehension and generation while retaining the binding and structure-building capabilities of symbolic approaches.

3. *Localist connectionist networks.* These networks spread numeric values ("activation") over weighted links (modifiable "synapses"). The processing units are very simple and perform nonlinear (e.g.thresholding) operations over the summation of their inputs. Each node in the network represents a given semantic or syntactic "primitive".

4. *Parallel distributed processing (PDP) models.* Here, long-term knowledge is encoded, not in the nodes, but in the connection weights and transient information is encoded as a pattern of activation over an ensemble of processing units.

5. *Artifical neural networks.* Here, an attempt is made to model actual neurons as closely as possible (above the neurochemical level), taking into account the dynamics of neural firing (e.g. decay, refractory periods after firing, phase locking of firing rates).

Hybrid models may also be built, consisting of modules with differing levels of behavior.

Why is it important to perform NL research at many levels? Currently, AI and neural paradigms supply different theoretical and technological capabilities. A comparison of symbolic and connectionist paradigms reveals that each has strengths where the other has weaknesses and vice versa (Dyer 1989). For example, symbolic systems supply variables, bindings, virtual pointers, logical rules, constituent structure, tokens (instances) versus types (templates), hierarchies, and inheritance. In contrast, distributed connectionist and artificial neural network systems supply statistically based associative retrieval, reconstructive memories from partial inputs, graceful error degradation in the face of noise or damage, automatic category/prototype formation through adaptive learning, and generalization to novel instances. High-level cognitive tasks, such as natural language processing (NLP), appear to require symbolic processing capabilities; however, the brain seems to gain its robustness from its distributed connectionist nature. Clearly, a synthesis is desirable and an important goal of NLP research is to increase the amount of parallelism, learning capabilities and robustness in NLP systems by implementing symbol processing capabilities within artificial neural networks and other connectionist models.

2. UCLA AI Lab Dissertations

An central way of characterizing a research lab is in terms of its publications, especially its doctoral disserations. Since the AI lab was founded in 1984, seven Ph.Ds. (chair: Dyer) have been produced:

1. Uri Zernik. *Strategies in Language Acquisitions: Learning Phrases from Examples in Context.* March 1987.

Zernik's dissertation concerns the acquisition of the syntax and semantics of figurative phrases from context. For example, given knowledge of the separate words "take" and "on", Zernik's system, RINA, recognizes that there is a novel phrase, "take on" whose structure

4

and meaning must be acquired. RINA incrementally acquires the correct syntax and semantics for "take on" through interaction with the user:

> Native: David took on Goliath.
> Learner: David took him somewhere?
> Native: No. David took on Goliath.
> Learner: David won the fight. He took on him?
> Native: No David took him on. He decided to fight.
> Learner: David accepted the challenge. He took him on?
> Native: Ok.   (Zernik 1987, p. 3)

Zernik is currently a research scientist at General Electric Research and Development Center, Schenectady NY. His dissertation is to be published as book by LEA Press.

   2. Erik T. Mueller. *Daydreaming and Computation: A Computer Model of Everday Creativity, Learning, and Emotions in the Human Stream of Thought.* March 1987.

Mueller's model, DAYDREAMER, produces a continual stream of conceptual representations, emulating the structure of daydreams. His research is an exploration into (a) the nature of invention (since the task is related to that of story invention), (b) the role of emotions in controlling thought patterns in a system with multiple conflicting goals, and (c) heuristics for generating hypothetical scenarios and learning from them. An example of a portion of output from DAYDREAMER (produced by an English generator from conceptual representations, and in female mode) is:

> I feel really interested in going out with someone... I have to go see a movie... I go to the Nuart [a theatre]... (User input: Harrison Ford is at the Nuart.)... What do you know! ... I have to have a conversation with him... I tell him I like his movies... Maybe he wants to be going out with me... I tell him I would like to have dinner with him at a restaurant... (User input: He declines.)... I feel really angry at him.... [daydreaming:] I study to be an actor... I am a movie star even more famous than he is... I feel pleased... He is interested in me... He breaks up with his girlfriend... He wants to be going out with me... He calls me up... I turn him down... I get even with him... I am pleased... What if I were going out with him? ... I remember the time he had a job with Paramount Pictures in Cairo... He would have to be in Cairo... He would go to Cairo... Our relationship would be in trouble... I would go to Cairo. I would lose my job at May Company. I feel relieved. ... ...   (Mueller 1987, pp. 235-284).

Most computers sit idle, awaiting input from the keyboard. In constrast, humans are constantly generating a continual stream of thoughts, in which they examine past events; altering those events to see how they might have come out differently. They generate hypothetical future scenarios and modify their current planning based on the outcomes produced by these imaginings. Even if only a small percentage of one's daydreams turn out to be useful in future planning situations, this will supply and important advantage over any system that remains idle and only computes when an actual situation is encountered.

Mueller is Member of Technical Staff, Morgan Stanley, NYC. Dissertation to be published as book by Ablex.

3. Michael J. Pazzani. *Learning Causal Relationships: An Integration of Empirical and Analytic Learning Methods.* May 1988.

Pazzani's system, OCCAM, deals with generalizing from event descriptions concerning threats and extortions, in the domains of children playing, kidnappings, and international economic sanctions. For example, OCCAM takes as input a sequence of economic sanctions events from around the world, during the period of 1921 to 1983. OCCAM organizes these economic sanction events into an episodic memory such that novel generalizations are formed. As a result of these generalizations, OCCAM can answer questions concerning hypothetical future events.

Question: What would happen if the US refused to sell computers to South Korea unless South Korea stopped exporting automobiles to Canada?

OCCAM: The goal of the United States that South Korea not sell automobiles to Canada will fail and South Korea will agree to purchase computers from a country which exports computers.   (Pazzani 1988, p. 14)

OCCAM starts out learning very slowly, by gathering statistical information from the input events. Based on this information, OCCAM forms causal relations that it then uses to produce causally based explanations of subsequent data. These explanations are then generalized in memory through processes of abstraction and coallescing of related experiences. OCCAM integrates two important forms of learning in the field of machine learning: SBL (similarity based learning) and EBL (explanation-based learning).

Pazzani is an Assist. Professor in the Computer Science Dept., Univ. of Calif. Irvine CA. His dissertation was nominated by the UCLA CS Dept. for the 1988 ACM Best Dissertation Award. His dissertation is being published as book by LEA Press.

4. Michael E. Gasser. *A Connectionist Model of Sentence Generation in a First and Second Language.* July 1988. (Applied Linguistics Ph.D., co-chaired with Prof. Evelyn Hatch).

Gasser's dissertation concerns the representation of concepts in memory and their relationship to speech-act structures for generation. Gasser was particularly concerned with how knowledge of one language interferes with the generation of another language. Gasser developed a localist connectionist model, called CHIE, that captures aspects of language generation errors produced by Japanese speakers attempting to learn English. Interesting features of his model are: (a) Although the final stream of language is sequential in nature, the process by which CHIE produces that stream is completely parallel. (b) Although CHIE doe not use markers (i.e. symbolic pointers) it can still generate pronouns; this requires an ability to termporarily bind syntactic structures. CHIE accomplishes this by creating highly primed paths of nodes in its network. (c) To control spread of activation, CHIE incorporates a refractory period in the firing of its nodes.

Gasser is an Assist. Professor in the Computer Science Dept. at Indiana University, Bloomington IN.

5. Richard G. Feifer. *An Intelligent Tutoring System for Graphic Mapping Strategies.* Feb. 1989. (Education Ph.D., co-chaired with Prof. Eva L. Baker).

The task of graphic mapping involves building a network of icons and labelled links to represent a segment of text. Feifer constructed a system, SHERLOCK, to help tutor students learning to construct graphic maps of a fragment of legal text concerning contract law. Using SHERLOCK's graphic interface, the student constructs a network that links icons (legal terms in this case) with part, is-a, leads, equiv, prop and not links. Unlike programmed learning systems, where the user is essentially passing through a series of branches within the program, SHERLOCK makes use of both a general rule-based inference engine and a spreading activation mechanism. As a result, SHERLOCK can dynamically examine student responses against its own knowledge in order to provide advice.

Feifer is currently a Research Associate at the Institute for Learning Science at Northwestern University, Evanston, IL.

6. Charles P. Dolan. *Tensor Manipulation Networks: Connectionist and Symbolic Approaches to Comprehension, Learning, and Planning.* May 1989.

Dolan's dissertation examines the underlying computational properties of two competing cognitive modeling paradigms: (a) Newell's Physical Symbol System Hypothesis (PSSH) and (b) Rumelhart and McClelland's Parallel Distributed Processing (PDP) approach. Dolan constructs a hybrid model, CRAM, designed to comprehend and acquire novel planning knowledge through reading Aesop's Fables and other Aesop-like stories. For example, CRAM takes as input a story in which a boss promises several secretaries each individually that they will be hired for a single, special position. The secretaries compare notes and then no longer trust the boss. CRAM reads the story and recognizes that a planning error has occurred. CRAM analyzes the planning error and applies heuristics to generate plan fixes. (For example, one fix is for the boss in the future to tell each secretary that the offer is a secret.) After reading this story, CRAM is able to more rapidly recognize the moral/point of a new story in which a professor promises each of several graduate students that each graduate student will be the one to receive the (single) post-doc position available.

CRAM is a hybrid model, in which the parser is symbolic, but the memory and binding modules consist of PDP networks. In order to construct and represent stories involving deception and flatter, CRAM must implement symbol structures and bindings in its PDP networks. To accomplish this, Dolan has created a new computational device: tensor manipulation networks (TMNs). TMNs are a class of PDP networks that allow the design of computers which obey the constraints (and exhibit many of the nice features) of PDP computation and still obey knowledge-level structures as in PSSH computation.

Dolan is a Research Scientist and Project Leader at the Hughes Research Labs. Malibu, CA. His dissertation has been nominated by the UCLA CS Dept. for the 1989 ACM Best Dissertation Award.

7. Sergio J. Alvarado. *Understanding Editorial Text: A Computer Model of Argument Comprehension.* June 1989.

Alvarado's dissertation is a purely symbolic model of argument comprehension in the domain of editorials concerning economic protectionism. For example, the program, OpEd, takes the following text, ED-JOBS (a fragment from an article by Milton Friedman) as input:

> Recent protectionist measures by the Reagan administration have disappointed us. Voluntary limits on Japanese automobiles and coluntary limits on steel by the Common Market are bad for the nation. They do not promote the long-run health of the industries affected. The problem with automobile and stell industries is: in both industries, average wage rates are twice as high as the average. Far from saving jobs, the limitations on imports will cost jobs. If we import less, foreign countries will earn fewer dollars. They will have less to spend on American exports. The result will be fewer jobs in export industries.

As OpEd reads the editorial, each word (or phrase) in its lexicon causes the creation of frames in working memory. Attached to each frame are one of more demons, each with parameter values (received from the lexical entries). As demons "fire", they perform various tasks, including: (a) disambiguation of word senses, (b) binding of one frame to the slot within another frame, and (c) recognition of argument structures. For example, the pattern <"far from" x, y> where x and y are opposing goals, causes OpEd to infer in ED-JOBS that it is Friedman (versus Reagan) who believes that imports will cost jobs (versus save jobs). As beliefs are encountered, OpEd constructs an argument graph of attack/support relations among beliefs. OpEd then uses this graph to answer questions:

> Q5: Why does Milton Friedman believe that the limitations on imports will cost jobs?

> A5: MILTON FRIEDMAN BELIEVES THAT PROTECTIONIST POLICIES BY THE REAGAN ADMINISTRATION WILL THWART THE PRESERVATION OF JOBS FOR U.S. BECAUSE ... IF U.S. IMPORTS FEWER PRODUCTS, THEN THERE IS A DECREASE IN PROFITS OF FOREIGN COUNTRIES; ... THEN FOREIGN COUNTRIES BUY FEWER AMERICAN EXPORTS; ... THEN THERE IS A DECREASE IN PROFITS FOR EXPORT INDUSTRIES; ... THEN THERE IS A DECREASE IN EXPORT JOBS; AND A DECREASED IN EXPORT JOBS ... THWARTS THE PRESERVATION OF JOBS FOR U.S.
> (Alvarado 1989, p.3)

Alvarado is currently an Assist. Professor in the Computer Science Department at the Unversity of California at Davis, and Director of the UC Davis CS Dept. AI Lab.

3. Ongoing Research Projects

Currently, there are over a dozen advanced graduate students pursuing NLP-related Ph.D. research in the UCLA AI Lab. These students are organized into related projects. A representative sample of these projects is included below.

3.1. LACQ Project

This project explores issues in language acquisition. Both symbolic and distributed connectionist approaches are taken in representing and acquiring the structure and meaning of NL words and phrases. At the symbolic level, special emphasis is placed on heuristic

reasoning methods for acquiring novel symbol representations for phrases in context. Consider interpreting the sentence "After years of fighting, Israel and Egypt *buried the hatchet*". First, an NLP system must know that Israel and Egypt are nations and that these nations have been at war. The system must also know that a hatchet can be used as a weapon in a fight and that burying places an object under the ground. Finally, the system must have the ability to make inferences, i.e. that objects underground are inaccessible and thus their normal uses are disabled. Since "the hatchet" refers to a specific hatchet, and since this unbound reference constitutes a discrepancy, the system must search for a semantic hypothesis which matches events from the known event sequence. Since weapons enable fighting, the system must generalize "hatchet" to weapons of war, and "burying" to disablement of warfare. Thus, learning this phrase involves discrepancy analysis, access of world knowledge, generalization and analogical reasoning. Zernik's dissertation is one result of this project; see also Dyer and Zernik (1986) and Zernik and Dyer (1985, 1986, 1987).

At the distributed connectionist level, research focusses on how word/phrase content can be represented as a pattern of activation over a set of processing units. How should symbols be represented and formed in distributed connectionist networks? In von Neumann machines, symbols are implemented as bit patterns residing in separate memory registers. The bit patterns are specified by a predetermined coding scheme, such as ASCII. The encoding scheme is both arbitrary and static; e.g., ASCII code was invented by engineers and the ASCII code for, say, "CAT", remains the same throughout all system executions. In purely symbolic systems, the arbitrary and static nature of symbol representations are not viewed as any problem, since it is assumed that the semantics of a given symbol develops only in terms of the structured relationships it enters into with other symbols. While it is the case that symbols enter into structured relationships with other symbols, the arbitrary and static nature of von Neumann symbol representations results in the inability of standard symbolic models to perform associative inference, handle noise and damage, complete partial patterns, or generalize to novel cases. In contrast, distributed connectionist systems can represent symbols as patterns of activation over a set of processing units in the input/output layers of a given network. These patterns of activation are intimately involved in the success (or failure) of the associative operations that are demanded of them. The generalization and noise handling capabilities of distributed connectionist networks depend on similar patterns in the input layer reconstructing related patterns in the output layer.

What we want is a method by which symbols can enter into structured, recursive relationships with one another; while at the same time forming distributed patterns of activation. The general technique for accomplishing this goal we refer to as *symbol recirculation*. Symbols are maintained in a separate connectionist network that acts as a global symbol memory, where each symbol is composed of a pattern of activation. Symbol representations start out as random patterns of activation. Over time they are "recirculated" through multiple tasks being demanded of them, and as a result, gradually form distributed representations that aid in the performance of these multiple tasks in distinct networks.

As a result of symbol recirculation methods, the symbols formed have their own "microsemantics"; i.e., *words with similar semantics* (as defined by word usage) *end up forming similar distributed representations in the lexicon*. The resulting theory of semantics in distributed connectionist models is very different from that of traditional NLP, in which word meanings are represented in terms of symbolic structures and their expectations in terms of explicit inference rules, e.g. (Dyer 1983). In symbol recirculation, *the representation*

9

of each word carries a memory trace of all the contexts of use that serve to define it. For more discussion of these methods, see (Dyer 1989) (Dyer, Flowers, Wang in press), (Miikulainen and Dyer 1988, 1989a,b), and (Lee, Flowers, Dyer 1989b).

3.2. PDS Project

This project explores issues in integrating the structure of semantic networks in AI with the processing capabilities of localist and distributed connectionist networks. The resulting *Parallel Distributed Semantic (PDS)* networks have two levels of behavior and analysis. At the macroscopic level, each node has semantic content and is linked to other nodes to form structured representations. At the microscopic level, however, each semantic-level node actually consists of an ensemble of PDP units, and links between semantic nodes are realized by full connectivity between PDP units across distinct ensembles. In standard semantic networks, instances are represented by uniquely "gensym-ed" symbols (e.g. INGEST14, JOHN3, PIZZA7) and types (e.g. HUMAN, FOOD). Such instances and types do not exist explicitly in PDS networks, but rather are reconstructed as patterns of activation over PDP ensembles. For example, instead of INGEST14 to represent "John ate a pizza" there is a pattern of activation that, when propagated to the ACTOR and OBJECT ensembles in the PDS network, will cause the reconstruction of patterns of activation that represent JOHN and PIZZA, respectively (Sumida and Dyer 1989). Another method of representing and propagating bindings is to associate a unique pattern of activation (called a "signature") with each instance in a network. A frame-slot (e.g. the ACTOR of the INGEST) is then bound to that instance (e.g. JOHN) when the frame-slot receives an activation pattern that corresponds to the instance's signature (Lange and Dyer 1989a,b). For more discussion of PDS networks and their features, see (Dyer 1989d) .

3.3. Grounding Project

In most AI systems, primitive symbols and relations (of enablement, causality, motivation, etc.) are defined and then used to construct knowledge. For example, Schank (1973) defined a set of basic predicates, including PTRANS, which represents motion through space. Once PTRANS is given, one can define concepts in terms of it. For instance, "walk", "fly", "throw" all involve PTRANS (in configurations with other basic predicates). But where did this predicate/symbol come from? In humans, knowledge of motion-through-space is acquired incrementally during infancy, as the infant is carried about, atempts to move, and observes other objects and agents move about. The "grounding problem" is that of mapping basic concepts to sensory experiences in the perceptual world.

We have been designing a system, called DETE (Nenov and Dyer 1988), that receives, on an artificial retinal of artificial neurons, a sequence of images of a simple moving object. Projections from this retina extract the direction of motion, mass of the object, change in direction, speed of motion, etc. As the object moves, a sequence of words, such as "ball... moves... up... bounces... off... wall... moves... down", is placed in a word buffer, and the two sequences are fed to a polymodal cortex, where visual and verbal information are then associated. In DETE, the time dynamics of the neurons are important, since the words and images are arriving with varying frequencies. DETE has a "mind's eye" retina and after training, DETE can be given a sequence of words describing an object's motion and a fuzzy image of a moving object will be generated as output on the mind's eye. Likewise, given a visual sequence, DETE can generate a simple verbal description of that sequence. The result

of this research should be a richer representation of basica conceptual predicates, such as PTRANS, which can then be used during symbolic/perceptual reasoning.

3.4. F.L.E. Project

This project explores the foundations of legal expertise (Dyer and Flowers 1985). A program, STARE (Goldman, Dyer, Flowers 1985, 1987, 1988) has been designed that takes simplified legal case descriptions as input, in the domain of contract law, and recalls cases from legal episodic memory that are conceptually relevant. The program also generates some legal advice, based on the cases recalled. Major research issues being explored here include: (a) cased-based, precedent-based, and analogical reasoning in law, (b) legal language analysis and comprehension, (c) representation of legal concepts, and (d) organization of legal cases in a conceptual memory. STARE makes use of conceptual predicates concerning, e.g., rights, duties, and liabilities in order to represent contractual situations.

3.5. Mentor Project

This project explores the tasks of intelligent advice giving and user modeling. Feifer's Ph.D. (i.e. the SHERLOCK graphic mapping tutor) is one result of this project. Another result is AQUA (Quilici, Dyer, Flowers 1988), a Prolog program designed to represent both tutor and learner goal/plan/belief knowledge. AQUA models advice giving in the domain of Unix file system use. AQUA has knowledge of patterns of belief concerning goals and plans. AQUA reasons about this knowledge in order to generate appropriate advice to Unix users whose problems stem, not only from lack of Unix knowledge, but mistaken beliefs concerning the behavior and/or purpose of various Unix file commands.

3.6. EDISON Project

EDISON stands for "an engineering design invention system operating on naive mechanical principles". EDISON (Dyer and Flowers 1984), (Dyer, Flowers, Hodges 1987, 1986) (Hodges 1989) takes natural language descriptions of simple mechanical devices (e.g. nail clipper, crank can opener, toy dart gun) and builds symbolic representations in memory. Given these representations, EDISON can simulate device function. EDISON also applies mutation and analogical operations to generate novel devices. A major goal of this project is to test out primitive symbolic predicates and relations for representing (a) the interaction between mechanical components and (b) naive reasoning concerning primitive machines and their principles (e.g. mechanical leverage, springs, troughs, hinges).

3.7. OpEd Project

The OpEd (Opinions to/from the Editor) project involves modeling argumentation, such as that which arises in the "oped" page in newspapers and magazines. A result of this research is Alvarado's Ph.D. The OpEd program makes use of knowledge structures called *argument units (AUs)*. An example of an AU is:

> x believes B1 (that plan P1 will achieve goal G1)
> therefore, x believes B2 (that P1 should be executed)
> y believes that G1 is desirable
> But y does not believe B2

because y believes B3 (that executing P1 will cause
          a side-effect (failure for goal G2)
and y believes that G2 is as important (or more important) than G1

AUs allow OpEd to represent arguments across many domains of discourse. *Meta-AUs*
support the representation of arguments about the validity of argument methods (Alvarado,
Dyer, Flowers 1985, 1986, in press).

3.8. MORRIS Project

This project concerns moral reasoning and the use of themes in story comprehension and
story invention. One result of this project is Dolan's Ph.D., which (in addition to an
exploration of tensor manipulation networks) examines heuristics for learning novel planning
knowledge through reading theme-based stories. In general, the thematic level includes
abstract constructs, such as irony (Dyer, Flowers, Reeves in press), ethics in planning
(Reeves 1988), and the recognition of planning errors (Dyer 1983d). Based on a
representation of themes, a story invention system, MINSTREL (Turner and Dyer 1985)
produces stories that have a moral or point to them. An example of a planning error is that of
inadvertently deceiving a co-planner when using deception against a common foe. In Romeo
and Juliet, for instance, Juliet inadvertently deceives Romeo into thinking that she is dead.
Issues of irony and ethical comprehension arise in stories where evil characteres receive
"poetic justice". Computationally, morality greatly complicates planning systems, since the
planner must take into account the effect of his plans on the goals and plans of other planners
in the environment. The MORRIS project explores these complications.

4. UCLA AI Lab Facilities

The AI lab contains 6 Hewlitt-Packard model 9000 series color workstations, each with 8-16
Mbytes memory and 165 to 260 Mbyte disk storage. These HPs are networked (with 2 571
MB HP disk drives) to other CS department HP workstations. This equipment was obtained
as part of a $2.5 million equipment grant (PI: Dyer) by Hewlitt Packard to the Computer
Science Department. In addition, there are 4 DN4000/3000 series color Apollo workstations.
These machines were obtained from a JPL contract and from a $1.2 million equipment grant by
Apollo Computer Inc. to CSRP (the Cognitive Science Research Program, an interdisciplinary,
organized research program in the cognitive sciences at UCLA).

Student projects are written in C, T, Prolog, and Common Lisp, which are supplied on these
workstations. In addition, there are in-house simulation tools, for modeling connectionist and
symbolic processes. The UCLA AI lab also houses the monitor to the CM2 Connection
Machine. The CM2 was obtained by the CSRP via an NSF grant. The CM2 is a massively
parallel SIMD machine with 16K processors and resides on COGNET, the CSRP network.
The CM2 can be accessed via any workstation in the AI lab, CS department, or on COGNET.

5. Conclusions

Research in natural language processing at the UCLA AI lab offers AI students an
entry into the major issues facing AI researchers and other cognitive scientists: (a)
the nature of knowledge: its representation, organization in memory, acquisition,
indexing and access, acquisition, and application, (b) the structure and function of

12

high-level cognitive tasks: comprehension and advice giving, production and invention, generalization and learning, planning and reasoning, (c) the role of parallelism in high-level cognitive tasks, through the use of spreading activation and distributed connectionist models, (d) the role of language with perception, through grounding symbols with visual input, and (e) the relationship of mind to brain, through implementing symbol processing capabilities in artificial neural networks.

6. References

Alvarado, S. J., Dyer, M. G. and M. Flowers. Natural Language Processing: Computer Comprehension of Editorial Text. To appear in: H. Adeli (Ed.). *Knowledge Engineering.* McGraw-Hill, NY. (in press).

Alvarado, S., Dyer, M.G. and M. Flowers. Editorial Comprehension in OpEd through Argument Units. *Proceedings of American Association of Artificial Intelligence (AAAI-86),* Philadelphia, PA 1986.

Dolan, C. P. and M. G. Dyer. Parallel Retrieval and Application of Conceptual Knowledge. *Proceedings of the 1988 Connectionist Models Summer School.* Morgan Kaufmann, 1989.

Dolan, C. and M. G. Dyer. Symbolic Schemata, Role Binding and the Evolution of Structure in Connectionist Memories. *Proceedings of the IEEE First Annual International Conference on Neural Networks.* San Diego, CA, June 1987.

Dolan, C. and M.G. Dyer Encoding Planning Knowledge for Recognition, Construction, and Learning. *Proceedings of Annual Conference of the Cognitive Science Society (CogSci-86),* Amherst, MA 1986.

Dyer, M.G. Knowledge Interactions and Integrated Parsing for Narrative Comprehension. In D. Waltz (ed). *Semantic Structures: Advances in Natural Language Processing,* Lawrence Erlbaum Assoc. (LEA Press), Hillsdale, NJ. 1989a.

Dyer, M. G. Symbolic NeuroEngineering for Natural Language Processing: A Multilevel Research Approach. In J. Barnden and J. Pollack (Eds.) *Advances in Connectionist and Neural Computation Theory.* Vol. 1, 1989b. (in press).

Dyer, M. G. Comprehension and Acquisition of Figurative Expressions with Phrasal/Lexical Memory. *Metaphor and Symbolic Activity.* Vol. 4, No. 3, 1989c. (in press).

Dyer, M. G. Symbol Processing Techniques in Connectionist Networks and their Application to High-Level Cognitive Tasks. In W. Brauer and C. Freksa (eds.). *Proceedings of 3rd international GI Congress on Knowledge-Based Systems.* Springer-Verlag, Informatik-Fachberichte series. 1989d (in press).

Dyer, M. G. The Promise and Problems of Connectionism. (Open Peer Commentary on Smolensky's "On the Proper Treatment of Connectionism"), *The Behavioral and Brain Sciences,* Vol. 11, No. 1, 1988a.

13

Dyer, M. G. Is Connectionism Cognitive Science? *AI Magazine*, Vol. 9, no. 4, 1988b. pp. 75-76. (Panel Report on Fifth Annual Theoretical Issues in Conceptual Information Processing Workshop, Wash. D.C. June 1987)

Dyer, M.G. Emotions and Their Computations: Three Computer Models. *Cognition and Emotion*, Vol. 1, No. 3, 323-347, 1987.

Dyer, M.G. (ed.) Artificial Intelligence Research at the University of California, Los Angeles. *AI Magazine*, Volume 6, No. 3, 1985.

Dyer, M.G. The Point of Thematic Abstraction Units. *The Behavioral and Brain Sciences*. Vol. 6, No 4, 1983b. (Open Peer Commentary on R. Wilensky's "Story Grammars versus Story Points").

Dyer, M.G. The Role of Affect in Narratives. *Cognitive Science*, Vol. 7, No. 3, 1983c.

Dyer, M.G. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press. Cambridge, MA 1983d.

Dyer, M.G.. and M. Flowers. Toward Automating Legal Expertise. In C. Walter (ed.) *Computer Power and Legal Reasoning*. West Publishing Co. 1985.

Dyer, M. and Lehnert, W. Question Answering for Narrative Memory. In *Language and Comprehension*, J.F. Le Ny and W. Kintsch (eds), North-Holland Publishing Co., NY., 1982.

Dyer, M.G. and U. Zernik. Encoding and Acquiring Meanings for Figurative Phrases. *Proceedings of the 24th Annual Meeting of Association for Computational Linguistics (ACL-86)*, New York, NY. 1986.

Dyer, M.G., Cullingford, R. & Alvarado, S. SCRIPTS. In Shapiro (Ed.) *Encyclopedia of Artificial Intelligence*. John Wiley & Sons, 1987.

Dyer, M.G., Flowers, M. and J.B. Hodges. EDISON: An Engineering Design System Operating Naively, *International Journal of Artificial Intelligence in Engineering*, Vol. 1, No. 1, July 1986.

Dyer, M.G., Flowers, M. and J. Hodges. Naive Mechanics Comprehension and Invention in EDISON. *Proceedings of the 10th Intl. Joint Conference on Artificial Intelligence (IJCAI-87)*. Milan, Italy, August 1987.

Dyer, M. G., Flowers, M. and Y. A. Wang. Distributed Symbol Discovery through Symbol Recirculation: Toward Natural Language processing in Distributed Connectioinist Networks. In R. Reilly and N. Sharkey (Eds.). *Connectionist Approaches to Natural Language Understanding*. North Holland Publ. (in press).

Dyer, M., Hodges, J. and Flowers, M. Computer Comprehension of Mechanical Device Descriptions. In J. Gero (ed.) *Knowledge-Based Systems in Engineering and Architecture*. Addison-Wesley (in press).

Dyer, M.G. Reeves, J. and M. Flowers. A Computer Model of Irony Recognition in Narrative Understanding. *Advances in Computing and the Humanities*, Vol. 1, No. 1 (in press).

Feifer, R. G.; Dyer, M. G. and E. Baker. Learning Procedural and Declarative Knowledge. *Proceedings of International Conference on Intelligent Tutoring Systems (ITS-88)*, sponsored by ACM, AFCET and CSCSI. Montreal, Canada, June, 1988.

Flowers, M. and Dyer, M.G. Really Arguing With Your Computer. *Proceedings of the National Computer Conference (NCC-84)*, Las Vegas, Nevada. July, 1984.

Gasser, M. and M. G. Dyer. Sequencing in a Connectionist Model of Language Processing. *Proceedings of 12th International Conference on Computational Linguistics (COLING-88)*. Budapest, Hungary, August 1988.

Gasser, M. and M.G. Dyer. Speak of the Devil: Representing Deictic and Speech Act Knowledge in an Integrated Lexical Memory. *Proceedings of Annual Conference of the Cognitive Science Society (CogSci-86)*, Amherst, MA. August 1986.

Goldman, S., Dyer, M.G. and M. Flowers. Representing Contractual Situations. In C. Walter (Ed.) *Computer Power and Legal Language: The Use Of Computational Linguistics, Artificial Intelligence, and Expert systems in Law*. Quorum Books, NY. 1988.

Goldman, S.R., Dyer, M.G. and M. Flowers. Precedent-based Legal Reasoning and Knowledge Acquisition in Contract Law. *Proceedings of the First Intl. Conference on Artificial Intelligence and Law (ICAIL-87)*, Boston, MA, May 1987. (sponsored by ACM SIGart and Northeastern Univ. Center for Law and Computer Science).

Hodges, J., Flowers, M. and M. G. Dyer. Knowledge Representation for Design Creativity. In C. R. Liu, A. Requicha and S. Chandrasekar (Eds.), *Intelligent and Integrated Manufacturing Analysis and Synthesis*. The American Society of Mechanical Engineers, NY, NY. 1987.

Hodges, J. Device Representation for Modeling Improvisation in Mechanical Use Situations. *Proceedings of the 11th Annual Conference of the Cognitive Science Society (Cog-Sci-89)*. Ann Arbor MI 1989.

Lange, T. E. and M. G. Dyer. Frame Selection in a Connectionist Model of High-Level Inferencing. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*. Ann Arbor, Michigan, 1989.

Lange, T. and M. G. Dyer. Dynamic, Non-Local Role Bindings and Inferencing in a Localist Network for Natural Language Understanding. In D. S. Touretzky (Ed.). *Advances in Neural Information Processing Systems 1*. San Mateo, CA: Morgan Kauffman, 1989.

Lee, G., Flowers, M. and M. G. Dyer. A Symbolic/Connectionist Script Applier Mechanism. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*. Ann Arbor, Michigan, 1989.

Lehnert, W., Dyer, M.G., Johnson, P., Yang, C.J. and S. Harley. BORIS -- An Experiment in In-Depth Understanding of Narratives. *Artificial Intelligence,* Vol. 20, No. 1, 1983.

Miikkulainen, R. and M. G. Dyer. Encoding Input/Output Representations in Connectionist Cognitive Systems. *Proceedings of the 1988 Connectionist Models Summer School.* Morgan Kaufmann, 1989.

Miikkulainen, R. and M. G. Dyer. A Modular Neural Network Architecture for Sequential Paraphrasing of Script-Based Stories. *Proceedings of the International Joint Conference on Neural Networks (IJCNN-89).* Wash. D.C. June, 1989.

Miikkulainen, R. and M. G. Dyer. Forming Global Representations with Extended Backpropagation. *Proceedings of the IEEE Second Annual International Conference on Neural Networks (ICNN-88).* San Diego, CA, July 1988.

Mueller, E. and M.G. Dyer. Toward a Computational Theory of Human Daydreaming. *Proceedings of CogSci-85,* UC Irvine, CA, August 1985.

Mueller, E. and M. G. Dyer. Daydreaming in Humans and Computers. *Proceedings of IJCAI-85,* UCLA, CA, August 1985.

Nenov, V. I. and M. G. Dyer. DETE: Connectionist/Symbolic Model of Visual and Verbal Association. *Proceedings of the IEEE Second Annual International Conference on Neural Networks (ICNN-88).* San Diego, CA, July 1988.

Pazzani, M. J. and M. G. Dyer. Memory Organization and Explanation-Based Learning. *Expert Systems: Research & Applications.* Vol. 2, No. 3. 1989. (in press).

Pazzani, M., Dyer, M.G. and M. Flowers. Using Prior Learning to Facilitate the Learning of New Causal Theories. *Proceedings of 10th Intl. Joint Conference on Artificial Intelligence (IJCAI-87).* Milan, Italy, Aug. 1987.

Pazzani, M., Dyer, M.G. and M. Flowers. The Role of Prior Causal Theories in Generalization. *Proceedings of American Association of Artificial Intelligence (AAAI-86),* Philadelphia, PA. August 1986.

Quilici, A., Dyer, M. G., and M. Flowers. Recognizing and Responding to Plan-Oriented Misconceptions. *Computational Linguistics,* Vol. 14, No. 3, 1988.

Reeves, John F. Ethical Understanding: Recognizing and Using Belief Conflict in Narrative Processing. *Proceedings of the National Association for Artificial Intelligence (AAAI-88).* ST. Paul, MN, 1988.

Seifert, C.M., Dyer, M.G. and J.B. Black. Thematic Knowledge in Story Understanding. *Text,* Vol. 6, No. 4, 1986, pp. 393-426.

Sumida, R. A. and M. G. Dyer. Storing and Generalizing Multiple Instances while Maintaining Knowledge-Level Parallelism. *Proceedings of Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89).* Detroit, Michigan. August 1989.

Sumida, R. A., Dyer, M. G. and M. Flowers. Integrating Marker Passing and Connectionism for Handling Conceptual and Structural Ambiguities. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society (CogSci-88)*. Montreal, Quebec, August 17-19, 1988.

Turner, S. and M. G. Dyer. Thematic Knowledge, Episodic Memory and Analogy in MINSTREL, a Story Invention System. *Proceedings of CogSci-85*, UC Irvine, CA, August 1985.

Zernik, U. and M.G. Dyer. The Self-Extending Phrasal Lexicon. *Computational Linguistics*, Vol. 13, Nos. 3-4, 308-327, 1987.

Zernik, U. and M.G. Dyer. Disambiguation and Language Acquisition through the Phrasal Lexicon. *Proceedings of Computational Linguistics (COLING-86)*, Bonn, W. Germany, August 1986.