

Of Frames, Scripts, and Stories

by

Choong Huei Seow

Submitted to the Department of Electrical Engineering and Computer Science
in Partial Fulfillment of
the Requirements for the Degree of
Bachelor of Science in Computer Science and Engineering

at the

Massachusetts Institute of Technology

May 1990

© Choong Huei Seow 1990

The author hereby grants to MIT permission to reproduce and to
distribute copies of this thesis document in whole or in part.

Signature redacted

Signature of Author _____

Department of Electrical Engineering and Computer Science

May 21st, 1990

Signature redacted

Certified by _____

Marvin L. Minsky
Donner Professor of Science
Thesis Supervisor

Signature redacted

Accepted by _____

Leonard A. Gould

Chairman, Department Committee on Undergraduate Theses
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

SEP 18 1990

LIBRARIES
ARCHIVES

Of Frames, Scripts, and Stories

by

Choong Huei Seow

Submitted to the Department of Electrical Engineering and Computer Science
on May 21st, 1990
in partial fulfillment of the requirements for the degree of
Bachelor of Science in Computer Science and Engineering

Abstract

An essential part of research in Artificial Intelligence (AI) involves the development of systems which can exhibit common sense and understanding. Such systems are proposed to be built upon platforms containing rich and different domains of everyday knowledge. This paper discusses the issues which dominate the framework of such systems. Such issues include the frame problem, the recognition of patterns and analogies, organization and usage of knowledge domains, etc.

Thesis Supervisor: Marvin L. Minsky

Title: Donner Professor of Science

Acknowledgements

I would like to express my sincere thanks to my thesis supervisor, Marvin Minsky. The opportunity to work with Marvin has always been very exciting, fun, and stimulating. I wish to also thank all my past UROP project supervisors, Dr. Andy Boughton, Professor Arvind, Professor Tommy Poggio, and Professor Steven Pinker, for the breadth and variety of perspectives and good discussions. My thanks to the members of the Learning and Epistemology Group of the MIT Media Laboratory, for initiating my curiosity in how we learn and think. To the MIT Artificial Intelligence Laboratory and all its members, thank you for the unusual, informal, and great research environment.

To my family, my heartfelt thanks for their continuous support, love, and encouragement. To Anne and Bob Berg, thank you for your love and support, and for your kindness of inviting me to experience life as part of an American family. My friends, thank you for being there when I needed it most.

Chapter 1

Introduction

1.1 Intelligence

What is intelligence ? The dictionary definition of intelligence is the capacity to apprehend facts and propositions and their relations, and to reason about them. The human ability to exhibit intelligence has been an area of scientific interest, in particular, scientists have been studying the implementation of intelligent “behavior” in artificial systems, namely computers.

1.2 Motivation

A major part of a human’s daily life involves the role of understanding. Our ability to grasp new information, learn new concepts and ideas, are all essential parts of understanding. It seems that intelligence can only exist if we have the ability to understand. We learn new knowledge when we understand the information which is presented to us. At this point, we still do not know how the process of understanding really takes place. However, research related to the process of understanding has been carried out in the field of Artificial Intelligence in the past such as the work which have been carried out on story-understanding models [Charniak, 1972][Wilensky, 1978], concept-learning programs [Winston, 1970], etc.

1.3 AI According to Marr

According to Marr, the implementation of a method to solve a particular information processing problem should be left to be the final stage of the project [Marr, 76]. One should begin with first identifying the information processing problem, and formulate an abstract method for solving the problem. This method can be decomposed and described by the following

- Algorithms which will be needed by the method.
- Functional descriptions of the algorithms.
- Information which is captured by this method.

Marr describes the above method as a *Type I* theory. According to him however, one of the principal difficulties of the study of A.I. is that one can never be sure if the problem which is to be solved has a theory of this type. He describes most A.I. programs which have been carried out are of the *Type II* theory. While a program of this theory type solves the proposed problem, it often performs the solution in a very clumsy and fragile manner. Thus, any candidate with a *Type II* theory will attach greater importance to the performance of the program. As Marr explains further, even though a *Type II* theory can be successfully implemented, one cannot still assume that a simpler underlying method does not exist. Marr states that a large AI system which has been built without a *Type I* theory would not be considered substantial, thus this system will not have any importance significance to AI.

This thesis discusses the underlying processes of understanding, in particular how it works, and how it is used in intelligent systems. This thesis will focus on the issues which need to be addressed in the design of systems which organize large and different domains of knowledge, systems which exhibit “common-sense knowledge”. Most of the thesis will focus on the heuristics of understanding, and the development of a framework for building such systems of knowledge.

1.4 The Role of Knowledge in Intelligence

Intelligence is being able to learn and understand.

1.4.1 Problems of Knowledge Representation

The issue of modeling knowledge is tricky. Workers in the field of A.I. have been trying to build models to represent knowledge but have had limited success. A fairly common problem is the amount and variety of knowledge we are most likely to use in the system, how much is needed and what levels of detail are needed ? One might argue that we need to represent explicitly, a large number of assertions to correctly represent an event or a concept. But, just how much is needed ? Do we want to be bogged down with details of knowledge which is not relevant to us at the time where we reference this knowledge structure? Also, if you make a comparison between AI models and the human memory model, you will notice that the human memory model poses a constraint on the amount of processing capability. We cannot recall everything at one time, but yet, we are able to recollect relevant information.

1.4.2 Useful Knowledge is Important

A good model of knowledge representation should represent knowledge which is useful to the problem-solver in an explicit manner. However, most models of knowledge representational which have been developed have a fairly limited domain of problem-solving. Once you move out slightly from this domain, the whole system becomes very fragile. How do they compare this to the human “model” of representing knowledge? We learn new things every moment, concepts, experiences, skills, etc. The things which we learn are very different, do we really have a single model of knowledge representation ? Does this indicate that we do not possess a single model of knowledge representation? Or is one model of knowledge representation sufficient for performing all cognitive skills which we have ?

In comparison to human cognition, we seem to use the many different domains of knowledge that we have in a wide variety of applications. At times, humans seem to

be able to make “*connections* or “*links*” to events which by themselves seem to be very different in content. The main question arises, and remains an important one, “*How do we organize and integrate these domains of knowledge?*” The following section will describe an area of cognition which is in need of a theory of knowledge integration.

1.5 Connectionist Models

1.5.1 Perceptrons and PDP

The field of artificial intelligence has experienced many different forms of growth and change. In the 1950’s the theory of *perceptrons* and *connectionist models* gained popularity among the workers in this field, and “suffered” a great setback upon a detailed study carried out by Minsky and Papert [Minsky, 68] which described the computational limitations of *perceptrons*. The revival of the popularity of *connectionist models* was initiated by the work of Rumerlhart and McClelland [McClelland, 86] who successfully described and implemented a connectionist model consisting of multiple layers of *neural units* which was “capable” of generating past tense forms of English language verbs. While this connectionist model did show forms of performance which were comparable to human ability, one should treat the success of the implementation of this connectionist model with caution. Although such a network model may actually work for a very particular cognitive skill, it does not necessarily imply that these forms of cognitive architecture will be successful for other domains of cognition. I believe that scientists should carefully evaluate the potential and functional extents of neural architectures before setting out to build neural networks to perform various cognitive skills. We have yet to learn in detail the true dynamics of neural networks (their detailed behaviour is still unknown), yet a a considerable number of researchers in the field are setting out to build more complex and larger neural network architectures to perform more complicated cognitive skills. In short, while the work carried out by McClelland and Rumelhart on multi-layered neural networks has brought new insight and scope to revitalize the field of connectionism once again, researchers have failed to recognize that a careful evaluation

of the prospects of neural networks should be carried out before building more complex forms of the networks.

Some of the limitations of connectionist architectures currently being developed is that they have cannot effectively store past states of the network, such that they could be automatically retrieved later for cognitive usage. Can we then safely conclude that if connectionist models of cognitive processes were to be successful, they would either be implemented in a system where other models of cognitive architecture are present too ? I posed this question since that if connectionist models are poor at maintaining knowledge of past states, it thus cannot be a good model for the human memory “system” as a standalone.

1.6 What’s Holding Back Computational Vision

The ideas of Marr described in the previous section were greatly influenced by his pioneering work in formulating a framework for solving the vision problem. Marr was convinced that a mathematical model of vision was correct since there had been successful implementations of *early vision* processes, such as Horn’s *shape-from-shading* theory for object reconstruction, etc.

At this point, we do not still understand how human object recognition works. Computational vision workers argue that it is composed of various “early vision” processes such as edge-detection, stereopsis, color vision and segmentation, etc. While most of the “early vision” modules have been successfully implemented, there still remains the difficult and complicated problem of *integrating* these processes into a model of high-level visual recognition. Have we progressed in the domain of object recognition from the implementations of “early vision” modules ? I must admit that I have the tendency to solve problems using a “bottom-up” design methodology. This design methodology is sound if the overall design structure of the entire high-level system has been carefully analyzed and the low level modules implemented according to the “global” high-level design. Instead, vision scientists have implemented “early low-level” vision modules and are currently working their way upwards to solve the problem of intermediate and high-

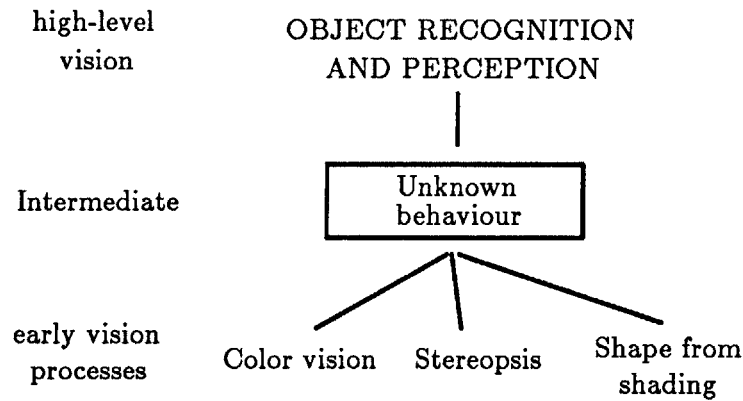


Figure 1.1: Segmentation of the Vision Tree

level vision processes. I would not be entirely surprised if the future system of a high level visual process such as visual recognition is not entirely dependent upon these “early vision” modules. While the study and implementation of such processes are important, it should be done in the manner and perspective of how such processes combine, interact and contribute to more abstract and complex processes. Consider Figure 1.1 as a rough representation of the current situation that current workers in vision are facing. While there has been successful implementations of such “early vision” modules, there still remains the open question of integrating these vision modules to construct higher-level vision theories such as object recognition and perception. The main problem which I believe is plaguing further development is due to the fact that the workers are using a mathematically based, computational approach to develop a theory of object recognition. While mathematical theories has been successfully used to solve early vision problems, I do not believe that they can hold as the very basis for the development of high-level vision processes.

1.6.1 Object Recognition Needs Knowledge

There have been numerous cases where AI scientists who began working on developing a model or a system which “successfully” imitated a particular cognitive function. Some of the projects were mainly motivated by an interest to solve particular problem, for example “*How do humans recognize objects ?*” So we set out on a project, trying to solve the problem from an engineer’s perspective (as a large number of workers in AI




Figure 1.2: Initial Learning Environment

have done). After a period of substantial effort, such a “system” is completed. The system performs “quite well” implementing a simple human cognitive mechanism, for example a system for navigating robots in an environment. The design ideas behind the development of the system was carefully planned and implemented. However, the problem of such a system is that its domain of application is very limited. How is the implementation of such a system relevant to Artificial Intelligence ? Minsky argues that the key essence of developing intelligent systems lie in the domain of understanding how we organize domains of knowledge which we refer to as being “common-sense” knowledge.

1.7 Looking Further Ahead

Previously, it was discussed that a successful system would have to have a good model of representing knowledge. However, such a model is not the only essential criterion for the successful implementation of a intelligent system. It is also essential for such a system to possess powerful learning mechanisms. Such learning mechanisms may be built out of more primitive mechanisms of reasoning, differentiation, etc. Do infants initially rely upon supervised learning to obtain new concepts and knowledge. The reliance on a supervisor for the concept learning is then slowly decreased until the period where the learner has understood the concept which was presented. Once this understanding has taken place, this individual can apply the knowledge of this new concept whenever the appropriate situation arises.

Referring to Figure 1.2, this diagram gives a general description of the interaction between the teacher and the student. The student is initially assumed to have very little or no background of the concept to be presented. The teacher is also assumed to have the goal of wanting the student to understand and learn the new concept. This

process proceeds further with the teacher presenting the student with information which is relevant to the concept. Such a presentation of information is consistent with the viewpoint of a good knowledge representational model, which was discussed earlier as being an essential part of AI systems. The teacher should explicitly show examples where the concept is correct, and slight variations of it where the concept is almost right. This propagation of information is referred to as *Process A* in Figure 1.2.

Such a methodology was followed by Winston's model for learning structural descriptions, essentially by presenting models which were correct or which were near misses of the correct examples [Winston, 1970]. While Winston's model was one of the successful AI systems which were built, I still see a strong potential of such a model failing. For instance, the order of which the examples are presented to the student observer is very important, since different conclusions about the structure to be learned can be obtained via different combination sequences of the same examples. Here, the sequence of information input has caused the failure, not the representation of the structural examples. It is very unlikely that the student learner can learn the correct structural description. From the viewpoint of the teacher, the teacher must take the following heuristics for teaching the student learner a new concept/idea

- What forms of knowledge does this student learner have ?
- What forms of reasoning mechanisms does it have ?
- How can I present examples of correct and near-misses of a concept in a fashion such that the student learner can infer the correct representation of it?

Following Winston's model of learning structural descriptions from examples, the teacher must present the student learner with the correct sequence/combination of examples in such a manner that the learner system could correctly derive the intended model's structural description. However, this heuristic discussed is not commonly experienced with everyday world situations. The sequence of which we receive knowledge about particular areas of the environment can be quite disorganized. However, we still have the ability to reconstruct all the information needed to learn a new concept. What allows

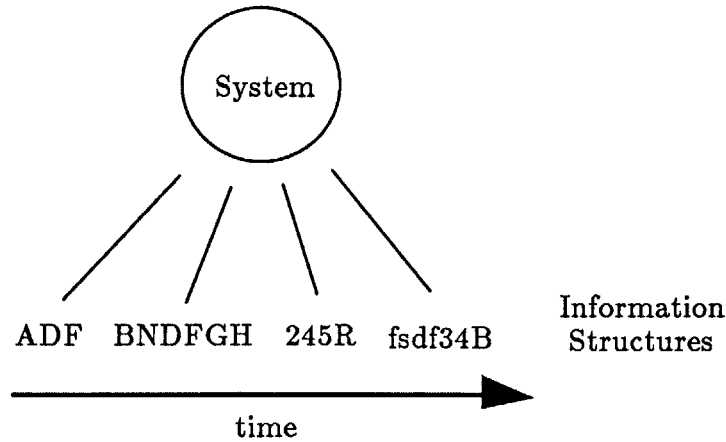


Figure 1.3: Streams of Information

us to do have this ability ? Figure 1.3 represents a more accurate situation of everyday information processing.

You don't know what you're looking for most of the time. If you were, it would be much easier to find a solution.

Obviously, this is quite a serious deficiency. How do we cope with such a situation when information is flowing into the system in an incoherent manner. We shall now turn to focus on several other AI models which were developed in the area of understanding.

1.7.1 Representations and Cognitive Skills

Humans use many different forms of information representations for facilitating our cognitive skills. When we try to find a solution for a particular situational problem, such as trying to find a path to go from *Location A in State M* to *Location B in State N*, we often use a form of a roadmap (a spatial map) to facilitate us in finding a solution. Work carried out by Kuipers [Kuipers, 1976] on spatial cognition describes different models of representations for representing spatial knowledge. Compare the map's model of spatial representation with an "identical" representation such as a complete written description of the same entire geographical area covered in the situation. While both forms of representations can be semantically "identical", the latter form of knowledge representation does not serve its purpose in efficiently aiding human spatial cognitive skills. This

leads me to the following general direction. I believe that the constraints imposed by human limitations in information processing capabilities caused us to develop models of knowledge representations which provide “excellent performance” for performance compensation. The forms of the knowledge representation models that we develop are most likely to be influenced by the type of knowledge being represented. If you had followed the example just given, the written form describing the spatial environment was not suitable for usage in trying to arrange a journey between the two geographical locations.

It is quite possible that the nature of human knowledge representation models have been influenced by the the functions and limitations of human cognitive skills.

1.8 Summary

I would like to end this introductory chapter with the intention of leaving the reader with a sense of what the remaining sections of this thesis will be focusing on. First, I have introduced certain aspects of AI which are considered to be key issues, such as the problem of organizing knowledge, and the notion of understanding information. The following sections of the thesis will remain focusing on the problem of organizing large domains of everyday common-sense knowledge.

Chapter 2

Organization of Knowledge

2.1 Introduction

As described previously, there remains a problem of integrating the different domains of knowledge that a system (and even humans) may possess into a body which is both useful and efficient. These domains of knowledge are by comparison, extremely large, since they contain all the facts and knowledge that the person has to date. While there have been systems which have been developed in the past to contain domains of knowledge, such systems are considered to be too specific in their applications, and their domains of knowledge considered to be too small in comparison to a general description of everyday knowledge which is contained in individuals. A good solution to the process of integrating knowledge of different domains will be beneficial to both builders of systems which exhibit common-sense knowledge.

2.2 Recognizing Knowledge Structures

2.2.1 Abstraction of Events

Can we build a representation tree which describes the event *Walking* ? If we take a not-too-short temporal subabstraction of this event, we will still get an event classified as *Walking*. What will happen if I took a much smaller temporal slice of this event. Would this small slice of this temporal event still be considered the same event, e.g. if I

were to show you a freeze frame of an event which displays a set of states, would it still represent the same event. What cues and features do we use to describe an event? If the freeze frame we had taken previously included only states which display an arm and a leg moving, do we actually classify this set of states as being the event *Walking*?

Furthermore, there are certain classes of events which are very difficult to represent and describe. For instance, how would one describe the event *JohnIsDrivingACar* in a set of assertions ? If we observe *John*, he could be experiencing a large number of smaller events which are happening throughout the period of this event. I strongly believe that such an event can only be described over a course of a certain period of time, i.e. there is a minimal interval during which this event takes place over, and to consider that the event actually happened, this interval should be completed. What follows is that when we take a temporal slice during this interval, we do not consider that temporal slice as an event of the same type also, but rather, we should assert that that event of *JohnIsDrivingACar* is in the midst of the process.

2.2.2 Important Properties of Events

There exist a large set of events which contain properties which will be called *static*. These so-called *static* properties are the ones which are constant and are of little significance or importance to the entire structure of the event. For instance, is the property *John's car is blue in color* in the event *JohnisDrivingACar* of any important significance to the event ? For most properties like these which are part of the world which the event takes place in, I would expect these properties to be static, assuming that event is the only "event" dynamically taking place at that time. However, some of the properties which are of significant importance to the event can themselves be static. Thus I came up with the following questions which would probably be the major design issues to be considered while building this "mechanical" system.

- How does a system determine if a property of an event is of importance to the event itself ?
 - What would the representation look like then ?

- Can a particular property of a subevent which is considered to be of little importance to the subevent be of much greater importance to a “parent” event ? If so, how would this be detected, and how would the status be changed ?

A system which exhibits intelligence will have to use various different domains of knowledge that it has. However, a significant problem which is facing the symbolic AI researchers is that the symbolic AI systems exhibit very poor performance in the area of *pattern recognition* and the *retrival of knowledge*. Such mechanisms are considered very crucial, at first glance humans seem to use such mechanisms always for problem solving. The following sections introduce a wide variety of abilities which are dependent upon pattern-recognition.

2.2.3 Building Analogies

One of the fundamental aspects of intelligence that humans display is the ability to draw analogies between various forms of skills, events, and knowledge. A similar mechanism, if implemented in a AI system, should be able to draw similar analogies between events which happen. In order to draw such analogies, the mechanism should be presented with knowledge which is represented in a form which possess certain similar structures. However, in order to perform the task of drawing analogies, the mechanism will have to examine the knowledge structures at varying levels of “detail”. In the course of searching for analogies, a more detailed analysis of subabstractions of the event should be done. In the case of temporal events, we would have to examine and identify the temporal relations which occur between subabstractions in an event. To an experienced individual, the partial set of states presented in the freeze frame is sufficient for the conclusion that the event *Walking* was actually being depicted in the freeze frame. Consider this case, where the individual with no knowledge of what *Walking* is, can this individual learn the concept of *Walking* which has just been presented to him ? Consider that the human being has quite a wide domain of varying forms of knowledge which is obtained, is it quite possible to state that we are not aware of the how we represent new information in memory. Thus, this question arises, should an AI system be even aware on how its

knowledge-based library of information be stored, and if not, how do we represent such information in a manner which is explicit and beneficial for the “intelligent engines” in this system to use. A genuine reason for an interest in this question, is that a large knowledge-based system will eventually contain a massive library of assertions. The mechanisms of the system which use this library of assertions must be able to use the information efficiently. Furthermore, there will probably be a wide and varying variety of different levels of information. It is important that the various mechanisms of the system which manipulate this library of assertions are able to “figure out” what is going on, what was stored in the database, and how to use these knowledge structures accordingly.

In the domain of building systems for understanding stories, the role of analogy is considered to be very important in the process of describing the thematic content of the story. For instance, the Broadway musical *“The West Side Story”* is very much based upon the story of *“Romeo and Juliet”*. An observer who has read both stories will have no trouble finding an “analogous flavor” relating them, although if the reader analyzes both stories are a lower level of detail, he will find a fairly large number of differences between the two stories.

2.2.4 The importance of inference

The ability to make appropriate inferences is important in any intelligent system, be it a natural or an artificial one. As mentioned in the previous section, the reader of a story must be able to deduce and make certain inferences based upon the story being read. By using inferences, a system can learn “much more”, in the sense that it can make implicit inferences which lead to new ideas and concepts, by also depending on prior default knowledge.

2.2.5 Consistency of Knowledge Structures

Another question which stimulates curiosity with is the notion of maintaining information consistency. Assuming we have a very large database of knowledge which we have accumulated across a long period of time, e.g. our own experiences that we have had

stored in memory, how do we maintain a consistent view of the experiences we had, in temporal terms ? In my opinion, I don't believe that we actually maintain total temporal consistency of events which had occurred in our memory. Let me give an example which was given by Minsky [Minsky, 90], that shows inconsistency among answers which we may give to questions about famous history events.

Person A: Which event do you think happened first, the American revolution or the French revolution ?

Person B: The French revolution.

Someone who has studied history will know that the American Revolution had taken place earlier (in 1776) than the French Revolution (in 1789). Obviously, Person B was wrong in choosing the latter. Why did this person choose the wrong answer? I carried out this experiment with a number of people, and among those who stated the wrong answer mentioned that France seemed much 'older' than the United States. Let me now proceed with this argument. I boldly state that in this case, the number of persons who initially given the correct answer will decrease. The reason, I believe, is that many of these individuals had plainly used the dates of the occurrences of both events to help them in making a guess. Now, if we hypothetically erase both dates from their memories, they will have to resort to using other heuristics. If we look back at the group who seemed to think that the revolution in France had taken place earlier since it was an 'older' country. Here, they have made the error in judgement, by using a simple heuristic of temporal reasoning. There is no doubt that an error was made, but the decision was made within a short period of time. Is the knowledge that we represented in our memory incorrect, or did we use an incorrect heuristic to come to the conclusion ?

2.3 A Flavor of Knowledge Structures

This part of the thesis will cover the motivations which led to the development of the *frame-system theory* by Minsky [Minsky, 75] and another theory of knowledge structures quite similar to frame-systems, the notion of *scripts and plans* developed by Schank

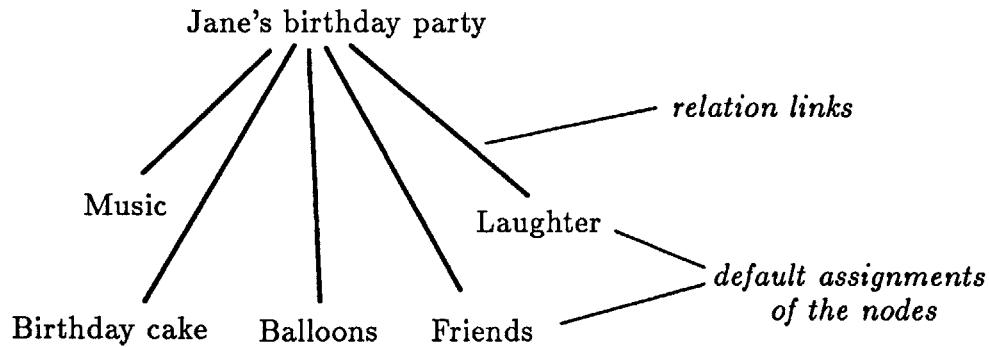


Figure 2.1: Frame representation for *Jane's birthday party*

[Schank, 1977]. Both these theories of knowledge structures are considered to be significant and essential in building AI systems which exhibit common-sense reasoning.

2.4 Frames

The idea of *frames* was first proposed by Minsky [Minsky, 75]. The frame theory is considered to be an key essence in the construction very large and rich databases which contain encyclopedic domains of knowledge needed in a system used for common-sense reasoning. In particular, the main idea of the theory was to construct such a database which encoded real-world knowledge in a structured and flexible manner [Shapiro, 1987]. A frame is considered to be a structure for representing knowledge of stereotyped situations such as a children's birthday party, being in the living room of your Aunt Mary's home in New Jersey, etc. Such a frame structure allows a computer system which is processing input from the external environment to exhibit coherence and understanding in similar situations, and would allow the system to access information which is needed for novel situations which are not anticipated prior to processing. Attached to such frames are default assignments of information which are present in those situations, and also information on usage and functions of that particular frame system. The theory of frames was constructed with the essence that knowledge had to consist of "larger and structured chunks", and that their procedural and factual contents were linked in intimate manners.

2.5 Charniak's Model of Story Understanding

2.5.1 Charniak's Demons

The following is a definition of “demons” given by Charniak [Charniak, 1972].

demons Facts which are introduced by “concepts” occurring in the story are called “demons” since in many cases they must wait for further information. In such cases we can think of them “looking” for the appropriate fact. So “not being willing to trade” might put in a demon looking for another offer.

Mary was invited to Jack's party. She wondered if he would like a kite.

Charniak [Charniak, 1972] was interested in learning how we interpret text sentences, in particular how the processing of a prior sentence effected the way we “processed” the sentences that followed it. Charniak suggested that the processing of a sentence or an event was immediately followed by the activation of recognition-agents known as “*demons*”. The activated demons then will proceed with analyzing further data input (the remaining sentences of the story), and will dynamically alter the further processing of text sentences. Charniak's system introduces “concepts” while processing sentences in a story. These concepts influences the the processing of future sentences of the story. In this system, “*demons*” are introduced into the story whenever the proper concept (relative to a particular demon) has been mentioned in the story.

1. *He plunked down \$5 at the window.*
2. *She tried to give him \$2.50, but he wouldn't take it.*
3. *So when they got inside, she bought him a large bag of popcorn.*

Consider the example above which is given as an example of how we “activate” concepts while processing text sentences. If we first read the first sentence, we will probably invoke a concept of horse race frame [Shapiro, 1987]. However, if we carry on reading the second sentence, we may conclude that the event was an attempt to return change. However, with the processing of the third sentence, that the whole sequence of event was probably describing a date.

2.5.2 Filling in Details

Consider the following short story (given by Miikkulainen and Dyer [Miikkulainen, 89]).

John went to Leone's. John asked the waiter for lobster. John gave a large tip.

Based upon our own personal experience, we can probably come up with a few other details and events to the story.

John went to Leone's. The waiter seated John. John asked the waiter for lobster. The waiter brought John the lobster. John ate the lobster. The lobster tasted good. John paid the waiter. John gave a large tip. John left Leone's.

How easy should it be to activate demons ? How long should they remain active ? How many of them should be activated ? Charniak studied how questions which were related to a story could be answered by a story understanding system. It seemed that in order to do this, there was a need for a substantial amount of knowledge which is not directly contained in the story. Charniak's model of children's story understanding was one of the first AI systems developed which followed a trend in AI representation theory to use "large chunks" of knowledge [Charniak, 1972]. Charniak's model is considered to be an early idea of what Minsky defines as being a *frame system* [Minsky, 75]. In comparison to Winston's model for learning structural descriptions, Charniak's model was quite different. Winston's system lacked a rich containment of default knowledge. Instead, it relied on examples of correct structural descriptions, and examples which were considered to be near misses of the structural descriptions by the teacher supervisor. In Winston's case, the student learner had a goal, that is to find the correct structural description for a object block. However with the proper sequence presentation of carefully chosen structural examples, Winston's system could deduce the correct default properties of the block structure. On the other hand, a system which is build to understand a story does not necessarily have definite goals to work on. Thus when such a system is processing a story, it is difficult for the system to deduce indirect consequences of such a story

without having default assignments of values. However, there is a strong resemblance of Charniak's system and Winston's system. Each system, if given the proper sequence of events and information can build and construct a general idea of what was being presented to them. This is quite true in individuals who write stories. Such individuals write stories in a fashion such that readers are able to grasp the ideas which the authors are to present. Most importantly, the writers do take into account the reader audience that the stories are intended for. In this sense, the writers themselves conform to a type of *default structure* for the type of audience the stories are intended for.

2.5.3 Good ideas from Winston's system

Winston's system for recognizing structural descriptions from examples exhibits some heuristics which are considered to be excellent, and should be used in conjunction with the frame theory. Winston's system was able to learn the correct description of an object by systematically adding and deleting features/relations that were present in the objects or the *near-misses*. Thus after going through a series of near-misses and correct descriptions, the system was able to create a correct structural description of the object. The system exhibited a behavior which is considered to be a *refinement process*. At each stage while the system was processing the structural descriptions given to it, it selectively constructed a database containing the object's structural description by using *pattern matching* or a *graph-matching* process.

2.5.4 Groupings

Winston's system exhibited the ability to form generalizations about the underlying important features of a certain structure. For instance, the system could create a generalization for a *Tower* structure, which it converged with the generalization of a three block object in the BLOCKS world with the following characteristic relations

$$Clear(A) \tag{2.1}$$

$$Between(x, (A, C)) \text{ for any number of } x \tag{2.2}$$

$$On(C, Table) \tag{2.3}$$

The method of learning the general description of a *Tower* in Winston's system can be viewed as being very similar to the creation of a frame description. In this case, the three most important structural features which categorized the object *Tower* was captured by the system. These structural descriptions can be viewed as being the relational links between slots in a frame description.

However, Winston's system was fortunate and successful in the sense that it had an excellent teacher who could provide it with good examples of what were considered to be correct samples of objects, and near-misses of object samples. In addition to this, the teacher has also provided only key descriptions (data) which it considered to be essential in learning the structural description of an object. Such a problem is evident while constructing a system which processes information from the outside world, a process considered to be similar to human perception. For example, consider a system which was built for object recognition. If we were just to consider the amount of raw data which is to be processed from the visual image, the amount of computation needed will be inherently be large, expensive, and slow. In such a case, I would render that the frame theory will provide an excellent solution since it eliminates the need to process all the raw data which is being fed into the system. A system which uses the frame theory as a platform will only need to selectively choose the input which the particular frame structure considers "relevant" for common-sense reasoning.

Although Winston's system displayed some considerable success in learning the structural properties, it relied on the teacher to provide it with features which were considered important to the learning of structural descriptions. In other words, the teacher had chosen a representation which had exploited the important aspects in the problem of learning structural descriptions. In the case of Winston's system, the teacher explicitly emphasized the spatial and geometric relational links between the objects and their features. Winston himself acknowledged the possibility of complications if the number of object features and their relations grew larger in size, than the ones he had used for the system. In the case of more complex and larger sets of features and links, the system would have a much more difficult time in matching the features between the different examples given to it. I would expect that the number of examples of correct and near-misses which have

to be presented to the system in order for it to learn could easily increase by an *exponential* factor. I have reason to believe that such an expectation is a lower-bound estimate on the scalability of the problem since we assume that all features are “equal” in nature. There are strong reasons to believe that more complex features are in existence within the framework of a certain problem. The notion of “intelligence” is one good example of a very complex feature which humans possess. How does one categorize “intelligence”? Can we teach a system to recognize the existence of “intelligence”? The Turing test is one of the classic tests which were proposed by AI researchers that could be used to test if a machine was intelligent. A machine which passed the Turing test is thus said to be indistinguishable from a human being, relative to another human being’s observation.

2.5.5 Folklore Beliefs on Relational Links

Human languages seem magical and mysterious. How can the manipulation of a formal set of symbols and notations convey an seemingly infinite group of information ? The ways of which we express these notations (in forms such as voice tones, emphasis, etc.) seem to also convey various different meanings and interpretations. At this point, let us take a look again at Charniak’s model of understanding children’s stories. Why wasn’t this system as successful as Winston’s model for learning structural descriptions from examples ? Besides the point of not having a supervisor, Charniak’s system attempted to implement natural language processing in the context of understanding children’s stories. In such domains of natural language, there are strong reasons which indicate the task of analyzing the semantic content of sentences is very difficult and ambiguous indeed. Such stories are very much different from the problem of learning structural descriptions. The latter is very well-defined, with laws of geometry and spatial relations governing the problem. On the other hand, Charniak’s system tries to understand children’s stories. The laws of processing, creating, and understanding the semantic contents of such stories is very subjective and ill-defined in comparison to the processing of the traditional informative text found in science articles, financial reports, etc. Can we actually build a correct semantic representation of the former ?

2.5.6 Inherent Difficulty of Relational Links

One of the principal difficulties which can be observed in data interpretation lies in the domain of the relational links. It is quite evident that the context in which we interpret data from the environment lies in the way which we introduce the relational links which connect the various domains of information available from the input. Once again, we take a look at Winston's system and see why his system displayed success in learning structural descriptions. In the case of the examples used, it only contained *prepositions* which "captured" the important aspects of learning structural descriptions. However, this is not quite true if a teacher which can give the "correct hints" is not available. In such cases, how do we generate and assign relational links to the data which we receive then? This problem is very evident in problem domains which are very complicated (many control variables and parameters), a good and relatively humorous example are the dynamics of the stock market. No one seems to be able to correctly predict the trends of the market with great accuracy. A same analogy could be applied to other domains as well, if we cannot learn and process the relational links which exist between different domains of the data, we may never learn the concept being presented. Could this problem be very similar to ones Piaget observed in young children then? Children of the ages 3-5 years old cannot grasp the *principle of conservation*. In this case, the same amount of liquid is poured into two different containers. The first container is a thin and tall container, while the second container is shaped very low and wide instead. Almost all of the children in this age group will testify that the first container has a larger amount of liquid than the second one.

It seems quite clear the children do not have the ability to recognize and observe the relational function which connected the initial frame to the final frame. What was very important in this relational function of *Pouring Water into the Container* was that the amount of water from the initial frame was the same as the final frame. Instead of observing the transformational function, the children had instead just focused on the comparison of the two final frames (referring to the water in the first container, and the water in the second container after the pouring of water). During this comparison, it is quite possible that children were relying on the relational links and the observational

data solely present in the final frames. In such a case, the children would have used the relation of *Higher Level* meaning “more” water to make their judgement. Thus they had failed to observe the importance of the conservation of liquid in the transformational function/relation. One more example is another observation made by Piaget on young children. The following is a reconstruction of the experiment conducted by him.

The little girl was playing with a toy cat. There was a small blanket which was near them. I came closer, and took the cat away from the little girl. I placed the cat under the blanket while the girl was looking. The girl started crying when she couldn't find the toy..

In this case, we could apply a similar explanation for the observation. The child was not able to grasp the fact that the toy had not disappeared, but was merely under the blanket. Further, this could be due to the fact that the child had only processed the information of the final frame, which did not explicitly indicate the presence of the toy under the blanket. In this case, the child had ignored (or failed to recognize) the “important” characteristic of the transformational function *HideTheToyUnderTheBlanket*, the characteristic that the object would still be in the environment of reference in the final frame, only that it was under the blanket this time around.

This phenomena also exists as a problem in a number of AI systems which have been developed. This problem of not being able to use the transformational relations between frames is evident in the STRIPS problem solver. In the case of STRIPS, its mechanisms run problem-solving methods on only the current state of events which appear in the particular frame which it is processing, and it does not have any database which contains the history of past events. Thus, it is quite possible that the system may repeat procedures which have (in the past) been carried out for problem-solving, and end up have the same results, on the negative side this results as failures to solve the problem. Furthermore, if such a system does not reach its success goal, it cannot resort to using heuristics which in the past history have been known to be successful. STRIPS seem to act like the young child given in the example above. It probably cannot realize that the transformational event of putting the toy under the blanket maintains that the toy is under the blanket. Instead, it processes the final frame, where the notion of the toy being under the blanket

is not present explicitly (if you assume that the vision mechanism is the one providing the input from the frame).

2.6 The Problem of Frame Recognition

“Things that which are so important to us are easiest to retrieve.”

Frame recognition is one important problem in the whole essence of the frame theory which has yet to draw any solid answers. As examples which were given in the previous sections, we did not have any trouble activating the frames for the *date* event. However, the method on how we perform this skill is still not known clearly. I have reason to believe that the following “things” are important towards building frame recognizers. These “things” are just as essential towards the whole notion of the frame theory.

- refinement process of learning
- organization of knowledge and relations
- intentions/goals of parties

2.6.1 The Problem of Too Many Frames

Frames are a collection of knowledge structures which were constructed as a result of events, concepts, and experiences that we have accumulated throughout time. While the concept of frames seem natural to us, a side effect of having such forms of knowledge structures is their size and numbers. It is quite possible that we may have a typical long-term memory storage which contains a million different types of frames. The large number of frames is accounted for by the large number of stereotypical events, concepts, ideas, etc. which an adult accumulates during an average lifetime. This relatively large number of such structures poses issues which question the theory of frames. How do we recognize frames which need to be activated from this large number of frames? In the previous sections, a stereotypical situation of *Aunt Mary’s home in New Jersey* (frame) was given as an example. How is this frame chosen, recognized, and activated from a set

of one million other different frames ? There could be several ways which we recognize objects, and they can be listed as follows

- data slots of frames
- relational links of frames
- the combination of slots and relational links

A recognition mechanism based on only the first heuristic can be very simple. In this case, we recognize and familiarize a situation with a certain stereotypical frame based upon the base components which are taken from the situation/event currently being processed. However, there are a number of disadvantages of this heuristic. This heuristic allows the possible activation of a large number of frames residing in memory. For instance, if we were to activate all possible frames which have an object like *Car*, the number of such activated frames can be large. Another example in which this recognition heuristic fails takes place in visual perception. Such an example is once again the case of features of a face (eyes, ears, nose, etc.) which are presented in a non-facial form (all these features are disjointed and separated). We obviously do not activate a *face* frame using this heuristic, rather this observation points towards the validity of the second heuristic used, the relational links of the frames.

The usage of relational links in frames for recognition purposes are more suited towards recognizing possible analogies between the observed situation and a particular stored frame. As the earlier example of *West Side Story* and *Romeo & Juliet*, the “similarities” which exist between them exist at a higher level than the level of the base components. In this case, the “similarities” in the relational links of these two frame events led to the recognition that the *West Side Story* musical is merely a modern version of *Romeo & Juliet*. Scientist classify this level of recognition as being more “abstract”. At this point, there seem to exist a difficult question which needs to be answered. Does the ability to draw such abstract analogies a good indication that we understand the functions and the structures of the relational links in a frame? Or can we actually draw correct analogies without understanding or being aware of all the relational links in the frame?

Looking back again at Winston's system, the problem of pattern recognition was somewhat more easier. This was due to the paradigm model which was used for the system. It assumed that the presenter (teacher) was correct with presenting the system with correct and near-misses (incorrect) structural descriptions of objects.

At this point, it would be interesting to reintroduce an important feature of frames. Knowledge which are contained in frames are most likely linked in a graph-like representation, rather than a feature vector form of representation. This conclusion is mainly due to the observation that such facts and knowledge which are present in a frame structure must be "connected" in certain manners, for instance, if A is an example of B , then such a concept may be represented and connected in the frame with a IS-A link. The importance of these relational links in frames can be justified as the this following example shows. Suppose I present you with two figures. The first figure contains an image of your face. The second image will contain all the components of your face (e.g. nose, eyes, ears, mouth), however they are all "jumbled" up. When presented with the two figures, we can immediately recognize the former, but not the latter. This show that the relational links between the components play a key role in the description and thus recognition of a frame.

2.6.2 Prototypical Frame Features

This section will introduce the issue of prototypical features which may appear as default slots in frames. This issue is very much related to the notion of scripts which will be discussed later. Relating back to the frame which covers the event of Jane's birthday party. In this event frame, we have specific default slots, such as party balloons, birthday cake, presents, friends and relatives, etc. However, we have also left out some finer details. For instance, when we mentioned balloons as being present in the event, we do also have a general notion of what type of balloons we have present, what color they have, etc. Curiously however, these "micro-features" are not considered to be important to the frame, however yet we still have a certain "sense" of their prescene.

2.7 Scripts

The relational links in frames are very much related to the notion of script which were proposed by Schank and Abelson [Schank, 1977]. They proposed that knowledge that we have of everyday common events are organized into knowledge structures which are labelled as *scripts*. Events which appear in scripts are considered as the primary entries of this knowledge structure. These events are connected by *causal links* which relate these events, in a serial-type manner. For example, the story of John dining at the restaurant [Miikkulainen, 89] mentioned previously is considered to be a script. Scripts are constructed with emphasis given to the ordering of events in accordance to stereotypical events which are commonly encountered.

2.7.1 The Relevance of Scripts in Frame Theory

Scripts seem to be a naturally occurring event of the frame theory. In the theory of frames, we have default slots and relations which describe a particular event or concept. If the representation of the frame is constructed properly, the roles and functional descriptions of the agents participating in the event will be captured. Thus, by using the functional descriptions and roles of these “agents”, we can successfully predict their sequence of “role-playing” in the event. With the goals and intentions of the agents relative to a particular frame made available for processing, one can construct a general-type script to describe the event.

Like frames, some events which are considered to be unimportant to the particular context which it is being processed, are not “activated”.

2.7.2 What Demons Should and Shouldn't Do

The problem of having a very large number of possible frames for all events of possible stories is related closely to the question of object-recognition systems which rely on a knowledge base. Such an object recognition scheme will basically use its prior knowledge (residing in the knowledge base) to perform object matching and recognition. Obviously, if we had all possible variations of even a commonly known object, for instance a *Auto-*

mobile, the number of *Car* templates which are stored in the knowledge base is extremely large. As a result, the process of matching the currently observed *automobile* with an entry in the knowledge base will be computationally slow, if carried out without any search heuristics.

Charniak's model of story understanding would be ideal if all stories were written in the same format as Charniak's model would expect [Wilensky, 1978], that is a knowledge structure which was a *frame* type. However, such a system cannot adapt to novel situations. Such situations involves children's stories like the one give below which Wilensky [Wilensky, 1978] and Charniak himself pointed out as being problematic.

Jack was going to paint. He washed the brush.

In this story, Jack is presumed to have washed the brush so that it was clean. Charniak's system could not find this interpretation since its *painting frame* did not possess the particular event describing the brush cleaning at the beginning of its activated frame. One could argue then that this particular event could be added to the *painting frame*. However, if we were to account for all the possible stories and events which could be fed to the system, the number of such stories is enormously large. This would be of no advantage since the system would have to search through a large database of frames to find a correct match. Such a search is computationally slow and inefficient, and it is most likely that we do not have such a methodology. However, I believe that the frame theory is an strong example of an excellent Knowledge Representation model of AI, thus the problem of having large sets of frames and irrelevant inferences should be addressed.

2.7.3 A Counter Argument

From the previous section, we obtained Wilensky's argument that Charniak's system would fail upon the example given by him. There seems to be a solution to this problem. Suppose we assumed that the event of washing the paint brush was linked to the painting frame via a relational link such as the following. In this case it is assumed that the event of cleaning the paint brush (we call this event A) is "small" compared to the event of

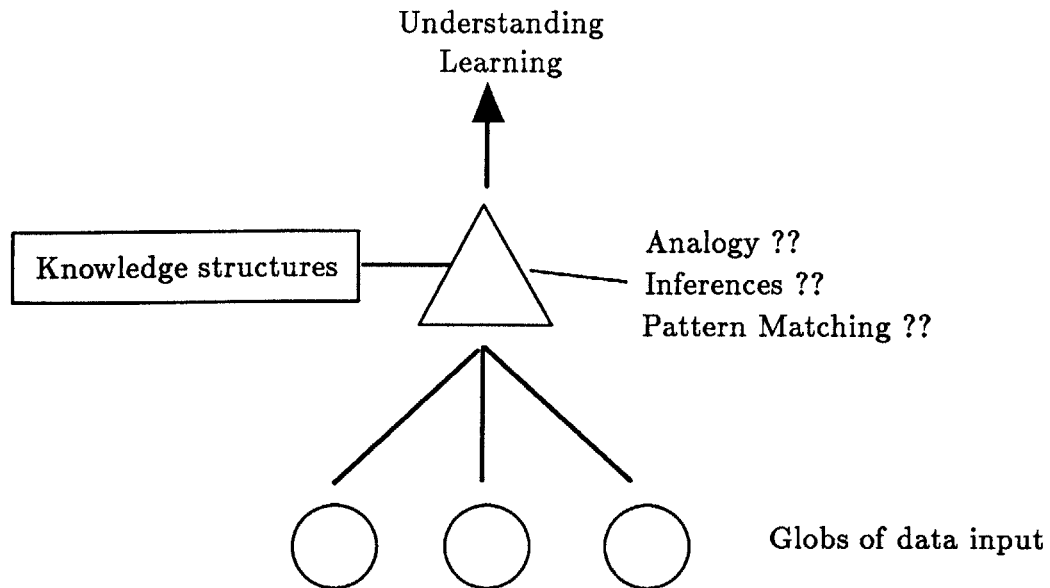


Figure 2.2: A Framework of an AI system ?

painting (we call this event B).

$$A \longrightarrow B \quad (2.4)$$

however, the relation is not symmetric, as shown below.

$$B \not\rightarrow A \quad (2.5)$$

Thus, can one say that the reason the *Painting* frame did not “contain” the event of *washing the brush*, is because our system could not trace it to the latter ? In the example given by Wilensky to criticize Charniak’s system, we can use the following explanation. After processing the first sentence “*Jack was going to paint.*”, we did not arrive with the activation of the *brushing cleaning* expectation, since we have no way of reaching out and activating the latter, since there was no causal link present in that direction. However, the latter event was still contained in the general frame of the *painting* event. If so, why do we do this ? This problem is evident in the design of relational databases in computer systems. Designers of such databases have pointers relating one entry to another via the method of using *pointers*. However, since pointers only travel in one direction, there cannot be any pathway which the pointed can go back to see who the pointer is.

2.7.4 How Large are Frames?

At this point, we don't really have a definite bound on the size of frames which can be generated and used. However, it is possible that we use the following heuristics for determining the size of frames.

- The type of event which is being represented
- The context of which the event appears in
- The information processing capabilities of a system

The first two heuristics are very closely related in determining the size of a frame. It is not clear yet if there exists subframes which are used in building a larger frame. Frames of different types can share a common set of slots/features, therefore, it is very difficult and complicated to draw boundaries between frames of this type. Such an observational difficulty poses the following question about the strength of the frame theory itself. Do we actually have these frames in our knowledge bases, or do we use a certain set of operations/transformations on a common database of raw knowledge such that the result of these transformations ? If we were to assume that we do not use the latter type of representation, we must then address the important issue which arises from the common sharing of data. This issue once again deals with the importance of the relational links connecting the data. How do we distinguish the different frames with ease, with the difficult situation which arises with the sharing of common knowledge assertions. It seems that the ability to do this is either caused by the following

- The frames which are constructed in memory do not share any particular common knowledge, i.e. there are no intersections between the "frames".
- Alternatively, the mechanism which can recognize and activate the correct frame must depend heavily upon the relational/structural links which connect the common sets of knowledge

The third heuristic mentioned is based upon the fact psychological observation that human beings have a limited capacity of information processing. This is based upon a

cognitive argument that we have a limited amount of short-term memory. With this limitation in existence, we resorted to (by choice or by evolution) the construction of frames which facilitated in the processing of information. The assistance which was provided by the formation of frames was quite similar to the concept of using heuristics for problem solving. We used these heuristics to serve as “guidelines” to the direction which we should head to. Similarly, we can activate a certain frame to facilitate the processing of a certain situation.

2.7.5 Information Processing Methods

Winston’s model of learning structural examples is a system which uses *data-driven*) processing. In data-driven processing, the inferences that the system make are just strictly dependent upon the input it receives, in the case of Winston’s system, the input are the sequence structural examples of an object. On the other hand, Charniak’s model is an example of *top-down* processing. In this type of system, the program will have a knowledge structure activated. The system will thus have to determine how it will match the data input it receives (in Charniak’s system, this consists of sentences in the story) with the “data slots” that are present in the knowledge structure. While matching takes place between the data and the knowledge structure Thus top-down processing can be described as a mechanism of prediction [Wilensky, 1978].

Chapter 3

Conclusion

In summary, there still remain open questions in the frame theory which need attention. These questions cover specific areas, such as the problem of organizing large domains of knowledge. The problem of organizing large domains of knowledge can substantially be reduced if a good set of heuristics can be developed to explore and “understand” the relational links which exist between different knowledge structures. AI systems which have displayed success, such as Winston’s model for learning structural descriptions from examples, were able to develop general sets of concepts through a process of refinement. However, in the case of Winston’s system, it relied upon a teacher to present it with examples which emphasized the important relational links between different parts of the information.

While it is crucial to recognize the importance of such relational links, it is still a very difficult task to understand the structure of such relational links, ones which represent a “concept”. The idea of a concept remains mysterious, we seem to be able to grasp such high-level knowledge structures, but we cannot seem to be able to understand fully the underlying composition and operations of the relational links. This problem is more apparent in knowledge which are categorized to be of “common-sense” nature, such as the knowledge required to understand children’s stories. Such situations are much more difficult than building a system to understand scientific text. In the case of scientific text, most of the text processing involves building a knowledge structure which contains objective data. This differs greatly from the processing of understanding children’s stories. The text contained in children’s stories have “thematic” contents

which are very subject indeed, dependent upon the goals and themes which the write and the reader have. In short, there does not seem to be an objective method for learning “concepts” which are presented in a child’s story. These concepts have very subjective interpretations, thus their network of relational links between the different knowledge structures in the “concept” is very context-dependent. This is due to the observation that concepts which are generated in context of common-sense knowledge seldom contain objective forms of information. In particular, the relational links which are present in the structural representation of a “concept” can be non-objective themselves (subject to context). To build a successful system which can organize “common-sense” knowledge from different domains, the questions which are expected to be of importance are the following

- The objectivity of relational links
- The order of importance of such links

An important aspect of relational links which should be given attention are the temporal relational links which exist between knowledge segments in a frame. It is still unclear on how these relations are built, and how they are used in relation to common-sense knowledge. Another issue which needs further study is the nature of accessing and categorizing relational links. This question is in part raised by the organizational hierarchy of frames. How are the slots in the frames organized ? There does not seem to be a definite manner for organizing the different “levels” of a frame. Rather, the manner of which the structure of the frame which we use seem to be very context-dependent. Thus it is quite possible that frames themselves are a form of graph representation which can be described as “non-hierarchical”. In such a case, it is possible that all the slots which exist in the frame can be accessed since there is no usage of “directional” pointers.

It is our hope that future work in AI will focus on studying and building systems which exhibit common-sense knowledge. Current research in AI places emphasis on problems which are very specialized and domain limited. These trends in AI research are considered not to be of fundamental importance to the field of AI. Instead, we should focus our efforts on solving problems which are considered by humans to be straight

forward and easy. Good solutions on how we solve these “easy” problems will provide fruitful insights to the advancement of AI.

Bibliography

- [Bobrow, 1975] Bobrow, Daniel G. and Donald A. Norman. Some Principles of Memory Schemata. In *Language, Thought, and Culture: Advances in the Study of Cognition*. Bobrow, Daniel G. and Collins, Allan (editors) Academic Press 1975.
- [Charniak, 1972] Charniak, Eugene. *Towards a Model of Children's Story Comprehension*. MIT Artificial Intelligence Laboratory AI-TR-266.
- [Dyer, 1983] Dyer, Michael G. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press 1983.
- [Guha, 1989a] Guha, R.V. and Douglas B. Lenat. *The World According to Cyc, Part 2: Agenthood, Institutions, and Agreements*. MCC Technical Report No. ACT-AI-453-89
- [Guha, 1989b] Guha, R.V. and Douglas B. Lenat. *The World According to Cyc, Part 3*. MCC Technical Report No. ACT-AI-455-89
- [Kahn, 1975] Kahn, Kenneth M. *Mechanization of Temporal Knowledge*. MIT Project Mac MAC-TR-155
- [Kuipers, 1976] Kuipers, Benjamin. *Spatial Knowledge*. MIT Artificial Intelligence Laboratory AI-Memo No. 359
- [Lenat, 90] Lenat, Douglas B. and R.V. Guha. *Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project*. Addison-Wesley, 1990
- [Marr, 76] Marr, David. *Artificial Intelligence - A personal view*. MIT AI Laboratory Memo No. 355.
- [McClelland, 86] McClelland, John and David Rumelhart (editors). *PDP: Explorations in Parallel Distributed Computing Vol. I, II and III* The MIT Press, 1986.
- [Miikkulainen, 89] Miikkulainan, Risto and Michael G. Dyer. "A Modular Neural Network Architecture for Sequential Paraphrasing of Script-Based Stories." *Proceedings of the 1989 IEEE Conference on Neural Networks* pp. (II)49-56.

- [Miller, 56] Miller, G.A. "The magical number seven, plus or minus two: Some limits on our capacity for processing information." *Psychology Review* 63:pp. 81-97.
- [Minsky, 61] Minsky, Marvin L. *Steps Towards Artificial Intelligence*. Proceedings of the IRE Vol. 49, No. 1, 1961.
- [Minsky, 68] Minsky, Marvin L. and Seymour Papert. *Perceptrons*. MIT Press 1968.
- [Minsky, 75] Minsky, Marvin L. "A Framework for Representing Knowledge." In *Readings in Knowledge Representation* R.J. Brachman and H.J. Levesque (editors). Morgan Kaufmann Publishers 1985.
- [Minsky, 85] Minsky, Marvin L. *The Society of Mind*. Simon and Schuster, New York 1985.
- [Minsky, 90] Minsky, Marvin L. *Personal communications*. 1990
- [Schank, 1977] Schank, Roger C. and Robert Abelson. *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum Press, Hillsdale, New Jersey.
- [Shapiro, 1987] *The Encyclopedia of Artificial Intelligence* Edward Shapiro (editor). John Wiley 1987.
- [Simon, 1979] Simon, Herbert A. *Models of Thought*. Yale University Press 1979.
- [Van Baalen, 88] Van Baleen, Jeffrey and Randall Davis. "Overview of an Approach to Representation Design." In *Proceedings of the National Conference on Artificial Intelligence*, pp. 392-397, Minneapolis, MN, 1988.
- [Wilensky, 1978] Wilensky, Robert. *Understanding Goal-based Stories*. Yale University Research Report No. 140.
- [Winston, 1970] Winston, Patrick H. *Learning Structural Descriptions from Examples*. MIT Artificial Intelligence Laboratory AI TR-231.