

Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2

Second Semester 2020/2021

Lesson 17

7 April 2021

Knowledge Structures: Frames, Scripts and Situation Models

Structured Knowledge Representations	2
Frames	2
Scripts.....	4
Situation Models	6
Situation Models as State Spaces	8
Narrative Reasoning	10
Narrative Reasoning with Situation Models	10
Events.....	11
Observability	11
Narratives	12
Entailment	12
Causal reasoning.	13
Reasoning about Causality with Narratives.	14
Problems with Structured Knowledge Representations...	15

Bibliography:

- 1) M. Minsky, "A Framework for Representing Knowledge", MIT AI Lab no 306, 1974.
- 2) R. C. Schank and R. P. Abelson, (1977). Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.
- 3) Tomkins, S. S. "Script theory: Differential magnification of affects." In Nebraska symposium on motivation. University of Nebraska Press, 1978.
- 4) P. N Johnson-Laird, "Mental models", MIT Press Cambridge, MA, USA, 1989.

Structured Knowledge Representations

Structured Knowledge Representations are declarative structures for reasoning and communication. A key property of Structured Representations is the association of symbolic information (including numerical values) with procedures for perception and reasoning. Structured knowledge representations are typically represented as a form of schema, which associate (or relate) information with slot and filler structures, and serve as a template to associate reified entities with properties, procedures and other information. Structured knowledge is generally organized in a class hierarchy with both slots and procedures inherited within the hierarchy.

Today we will review 3 forms of schema: Frames, Scripts, and Situation Models. One can say that these three techniques enable automated answers to the famous 5W+H questions: Who, What, When, Where, Why and How posed by Aristotle. (Also known as the Quintilien Questions).

Frames: What, Where, How?

Scripts: Who, Where, When, Why?

Situations Models: Who, What, When, Where, Why and How?

Frames

Frames were proposed by Marvin Minsky in 1974, as data structures to guide perception and reasoning. Minsky's insight was that it is much easier to **see** if you know **what** to look for and **how** to find it. Minsky proposed Frames as a structure to guide visual interpretation of scenes in a top down manner, telling a vision system what to look for and where to find it. Frames organize knowledge as hierarchical descriptions of entities, relations and events, and associate procedures and operations with data and goals.

A frame describes an entity as a set of properties and relations, represented by slots. This description includes a collection of procedures for perceiving, reasoning and acting with the entity. Key innovations included the association of procedures for detection and interpretation of data, and provision for default values for data when perception is not possible or fails.

A frame contains information on how to use the frame, what to expect next, and what to do when these expectations are not met. Top-level frames provide factual information that is always true in a particular context. This can be associated with other information that must be observed or inferred.


Structured Knowledge: Frames, Scripts and Situations Models

Information is stored in slots. A slot may contain:


- Data (Facts)
- Procedures:
 - IF-NEEDED procedures for deferred evaluation
 - IF-ADDED procedures for updates linked information
- Default values for Data and Procedures
- Pointers to other instances of other Frames or Sub-frames

Frames are composed using **relations**, represented by slots that contain pointers to other frames. Relations represent information about the object, such as part relations (composed of, part-of), position relations (above, below, beside, inside, contains), time relations (before, after, during), as well as specific properties of the entity (size, position, color, orientation).

Frames represent situations hierarchically, using PART-OF and IS-A relations to describe complex entities. Compositional (Part-whole) hierarchies provide visual context to guide scene interpretation by telling a program what to look for and where to look for it.




For example, the concept Triangle has the part relation "composed of" with three segments. The triangle can also have an is-a relation (category membership) with different triangle types such as equilateral, isosceles, right angle, etc.



Ultimately, some slots point to raw perceptions (phenomena).

Structured Knowledge: Frames, Scripts and Situations Models

Frames are instantiated by entities. When a slot points to entity, the entity is said to play a "role" in the current frame. Frames typically come with methods (procedures) for searching for the entities that can play roles in the frame. Typically a slot-filling procedure will apply a set of acceptance tests to an entity to see if it satisfies the requirements for the role.



A Frame for a Cube

(from E. Rich "Artificial Intelligence", Fig 7-13, p231)

Frames generally include typical examples (prototypes) that can serve as examples in reasoning, and default values that are used if no entity has been found to fill a slot. In this way, frames can be used for abstract reasoning or for default reasoning when perception is not possible.

Discovering the appropriate frame for reasoning about a situation is called "The Frame Problem". The Frame Problem was often identified as one of the hard, unsolved problems in symbolic AI.

Scripts

Script theory, originally proposed by Tomkins, posits that human behavior largely falls into pre-learned patterns that can provide a program for action and guide interaction with others. Tomkins created script theory as a further development of his affect theory, which describes human emotional responses to stimuli as falling into categories called "affects" that drive awareness and action. Tomkins proposed that scripts inform behavior so as to maximize positive affect and to minimize negative affect.

Schank and Abelson extended Tomkins' scripts as a method of representing procedural knowledge. In their work, scripts are a form of frame representation, where the values that fill the slots can be ordered. Thus a script can be used as a structured representation describing a stereotyped sequence of events in a particular context.

Structured Knowledge: Frames, Scripts and Situations Models

Scripts have been used in natural language understanding systems to organize the knowledge in terms of a sequence of situations. For example, scripts have been used to interpret and extract information from Newspaper Stories. Scripts are also used to observe subjects and to describe (or recognize) what they are doing, not simply listing actions but explaining goals and intentions. This includes plan-recognition as well as activity description. Scripts have also been used to represent procedural knowledge for planning as we will see in a future lesson on planning.

Scripts are schema that include slots that represent a sequence of situations. In script theory, the basic unit of analysis are scenes: specific regions of space and time, where situations and events take place.

A script is composed of

- 1) Scenes: Regions of time and space where the situation that compose the script take place.
- 2) Props: Entities (objects) involved in the script.
- 3) Actors: Agents that can provoke changes in the scenes.
Actors are typically people, but may be artificial.
- 4) Events (action): Changes in situation.

In each scene, one or more actors perform actions, typically using props. The actions of actors cause events that provoke changes in situation and eventually can advances the script to a new scene. The script can be represented as a tree or network of scenes, driven by events. As with Frames, scripts drive interpretation by telling the system what to look for and where to look next. The script can predict events and provide information about why actors perform actions.

Example of a script: Restaurant Script.

The classic example of a script involves the typical sequence of events that occur when a person drinks in a restaurant: finding a seat, reading the menu, ordering food and drinks from the restaurant staff.

The classic restaurant script includes

Props: A restaurant with an entrance, tables, chairs, plates, eating utensils, glasses, menu, food, drinks, the check.

Actors: The host (Maitre d'Hotel), clients, servers, chef, bus-boy, etc.

Scenes: Entering, seating, reading the menu, ordering, serving, requesting the check, paying, leaving, etc.

Structured Knowledge: Frames, Scripts and Situations Models

Scripts provide context for default reasoning.


As with Frames, scripts drive interpretation by providing procedures that tell a system what to look for and where to look for it. Scripts also provide default information for reasoning about stories or actions.

For example, for story understanding, a story will typically only provide sparse detail of what happened. The reader is expected to fill in the missing information using context. Scripts provide a means to represent contextual information.

Situation Models

Situations models were originally proposed by P. Johnson-Laird to capture the activity of working memory. Situation models are used in cognitive science to describe the mental models that people use to understand and reason. Situations models have also been adapted as means of structuring frames for observing and reasoning about human actions.

Working memory elements are called "entities". Entities are instances of concepts. The structure (the set of associations) for an entity are defined by its concept.




A situation is a set of relations between entities. In informatics terms, a situation is a form of "State" and a situation model is a "state space model". A situation model is composed of a set of entities (defined as instance of concepts with properties), a set of relations between entities, a set of behavior (event-condition-action rules) that can be associated to situations, and set of state transitions that organize the situations as a graph.

Structured Knowledge: Frames, Scripts and Situations Models

- Entities:** Anything that can be named or designated; People, things, etc.
(Entities can be defined using concept schema or frames)
- Properties:** Descriptions of entities such as position, size, color, etc
- Relations:** N-ary predicates (N=1,2,3 ...) that relate entities.
(Relations are defined as tests on the properties of entities).
- Situation:** A set of relations between entities

Situations can be organized into a state space referred to as a situation network. Each situation (or state) corresponds to a specific configuration of relations between entities. A change in relation results in a change in situation (or state).



The situation graph, along with the set of entities and relations define a **Context**. Thus situation models provide contextual knowledge for understanding.

Each situation can prescribe and proscribe behaviors, in the form of a list of actions and reactions that are required, allowed or forbidden for a situation. Behaviors are commonly encoded as Condition-Action rules, and can be encoded as tree structure.

Behaviors include


- 1) methods for sensing and perception, and
- 2) appropriateness of actions
- 3) changes in situation that result from perception or events

Situations can direct attention by listing entities and relations that should be observed along with methods for observation. As with frames, situation models can include default values for properties of entities and for relations between entities.

Situation models include a graph of possible situations with probabilities and conditions for state transitions. This graph can be seen as a form of state space, with transitions representing events.

Situation Models as State Spaces

A Situation model includes a directed graph of situations. Nodes are situations. Arcs are events.



The set of Concepts, Relations, and Events compose the Ontology for the Situation.

A situation model can be defined as a set of 4 sets: $\{X\}\{R\}\{E\}\{S\}$


- 1) A set of Entities $\{X\}$: Concrete instances of concepts.
- 3) A set of Relations $\{R\}$: Predicates over subsets of the Entities in $\{X\}$
- 4) A set of Events $\{E\}$: Operators that can change Situations by changing Relations.
- 5) A set of Situations $\{S\}$: A set of situations that share entities and relations

The situations model captures what partial descriptions used to reason. Situation models are typically limited to a finite set of entities in order to avoid an exponential explosion in the set of possible relations. Fortunately, for most problem domains, we can generally limit the number of entities, the number of concepts and the number of relations to small sets this makes the situation model tractable.

The sets of entities, relations, behaviors, and situations define a Context model and can be used to construct context aware systems. For example the RocamaRoll System installed in Amphi D of ENSIMAG is based on a context aware meeting recording system system developed at INRIA for meeting recording, by P. Reignier, D. Vaufreydaz and colleagues as part of a European FAME project in 2004. The systems was extended and improved in a second European project, CHIL - Computers in the Human Interaction Loop from 2005 to 2008.

Structured Knowledge: Frames, Scripts and Situations Models

The original FAME meeting room recording system was composed hierarchy of components shown below.



The system behavior was driven by a situation model.



For example, the above scene includes:

Entities: Patrick, Jerome, Sonia and Stan, table, chairs, agenda

Roles: Moderator, Speaker, Participant, current-agenda-item, etc

Relations:

Moderator(Patrick) speaks-to participants(...)

Participant(Jerome) talks-to Participant(Stan)

Participant(Sonia) looks-at Participant(Patrick)

Situation models can be used for narrative reasoning.

Narrative Reasoning



Calvin and Hobbes: Dad explains science

Why is it raining? "A cold front is approaching from the north-west drawing warm moist air from the Mediterranean and causing it to rise into cold air from the north."

Why are you late for class? "I missed the bus and had to walk."

These are examples of Narratives.

A Narrative is a temporal sequence of Situations.

Humans use Narratives to understand phenomena.

Narratives enable humans to provide rich descriptions for situations that are not directly observable, including situations that occurred in the past, and hypothetical or abstract situations. Narratives enable humans to make predictions for situation that may occur in the future.

Frames, scripts and Situation Models provide a formal foundation for narrative reasoning. Temporal sequences of situations can be used to provide rich descriptions of unobserved phenomena, and predict the likelihood of future phenomena. Situation models indicate how to collect evidence for and against the likelihood of competing narratives. Narratives can also be used to reason about causality.

Narrative Reasoning with Situation Models

A situation can be defined as a: $\{X\}, \{R\}$

- 1) A set of Entities $\{X\}$: Instances of concepts that represent phenomena.
- 2) A set of Relations $\{R\}$: Predicates over subsets of Entities.

A Situation is a set of relations over entities (a state).

Structured Knowledge: Frames, Scripts and Situations Models

Situation models can be used to

- (1) Providing meaning and explanations for phenomena
- (2) Describe phenomena that cannot be observed
- (3) Predict phenomena that have not yet occurred.

A situation model is a directed graph where the nodes represent situations and the arcs represent events. A situation model is a form of state space.

In most real-world situations the number of observable entities exceeds Working Memory. A human is required to create possible models with subsets of relevant entities and relations.

Generally humans develop reasoning skills in the form of processes to "cycle" through partial situation models looking for configurations of entities and relations that trigger associations with additional information.

Certain partial descriptions in WM "match" with narrative stories that allow the reasoner to predict phenomena or propose an explanation. Such narratives are a form of "story" recalled from episodic memory. Narratives provide a form of meaning.

Events

Events are external phenomena that change a Situation (S) by changing the set of entities or the relations between entities. Events can be the results of actions by the reasoner, by other agents, or they can be natural phenomena.

$$S_1 - E_{12} \rightarrow S_2 - E_{23} \rightarrow S_3 - E_{34} \rightarrow S_4$$

Observability

In the real world many situations are not directly observable! For example, we cannot observe situations that occurred in the past, hypothetical situations, or situations that will occur in the future.

In this case, a situation model is said to be "Partially Observable". In addition, situation models can be defined for concepts that are not directly observable!

Examples include radio waves, phenomena that are too slow or too fast for human perception (global warming, ultrasound, infrared light), or too large or too small for human perception (atoms and molecules, the Universe).

Humans use Narratives to reason and communicate about unobservable phenomena and partially observable phenomena.


Structured Knowledge: Frames, Scripts and Situations Models

Narratives

A Narrative is a temporally ordered sequence of Situations.

Situations in a narrative have a temporal duration and are ordered by time relations.

A Narrative represents a path through a situation graph.



Narratives are an essential part of human reasoning.

For example, a personal narrative is an essential part of our self-identity.

Narratives are commonly used in public discourse to explain and predict social phenomena. Scientific theories are a form of narrative for explanation and prediction of phenomena.

A Narrative provides an explanation for a situation by describing previous (non-observable) situations. A narrative can provide predictions of situations that will occur in the sequence after the current situation. Narratives provide explanations by describing how events can change the relations between entities.

A Situation model can be used to "infer" the occurrence of non-observable entities, relations and events.

Entailment

An entailment is a predicate (relation) that is implied by a statement.

A situation "entails" the predicates that define it. If the situation occurs, then the relations of that situation are assumed to be true even if not observed.

$$S \vdash r_1(X_1) \wedge r_2(X_2) \wedge \dots \wedge r_n(X_n)$$

Where S is a situation.

$r_1(X_1), r_2(X_2), r_n(X_n)$ are relations in the situation,

X_1, X_2, X_n : are sets of entities in the situation.


Each entity in a situation is represented by a predicate $Exists(X)$. $S \vdash \exists(X)$

Structured Knowledge: Frames, Scripts and Situations Models

Exists(X) is included among the relations that can be entailed by a partially observable situation.

Event if a relation is not directly observable, it can be entailed by a situation.

We can use entailment to specify evidence to collect for alternative narratives.



For example, when considering two competing narratives for a situation, the relations that define the situations can be used as evidence for or against each narrative.

Unfortunately, some relations will be in both narratives (ambiguous).

Other relations may not be observable. We need a mechanism to reason with unobservable and ambiguous evidence.

Causal reasoning.

Causality is an abstraction that indicates how the world progresses.

Defining Causality is difficult to define as a technical term because the word has divergent definitions in different scientific and philosophical domains, in particular because of different notions of agency or action.

Causality assumes notions of time, state and process. A process is a temporally ordered sequence of states (situations).

Causality can be defined as actions or events (the cause) that change a process (the effect). The Cause is said to be partly responsible for the effect, and the effect is partly dependent on the cause.

Processes can have many causes, which are said to be causal factors. All causal factors are prior (before) the process. An effect can in turn be a cause of, or causal factor for, many other effects, which all lie in its future.

Narratives allow us to reason about alternative causes for a situation, where the causes may be random events or deliberate actions.

Reasoning about Causality with Narratives.

In terms of narratives, a Cause is an event or action that changes the evolution of the narrative.

Consider the following example.

Situation 1: John walks to the tramway-stop from 7:20 to 7:30

EventA: John boards the tram at 7:30

EventB: John fails to board the tram at 7:30.

Situation 2A: John is at campus from 7:50 to 9:00

Situation 2B: John is at campus 8:00 to 9:00

We can say that not boarding the 7:30 tram **caused** john to be arrive at 8:00.

Events, in this formalism, can be replaced with actions.

For example, maybe John decided not to board the tramway at 7:30 because it was too full.

Events can also be random and non-deliberate.

For example, maybe John tried and failed to board the tramway at 7:30 because it was too full.

Problems with Structured Knowledge Representations.

Structured knowledge representations were invented as a programming tool for intelligent systems. This approach suffers from a number of open problems:

- 1) Top down reasoning: Frames (and most schema systems) are designed for top-down reasoning. Most human reasoning is both top-down and bottom-up (active), with associations flowing both ways.
- 2) Knowledge Acquisition: Building a Frame system by hand is long, tedious, and ad hoc process. There is a temptation to overload the system with useless information, "just in case". Automatic acquisitions (learning) of frame systems for recognition and reasoning is a notoriously hard research problem for classical AI.
- 3) Context Recognition (The Frame problem): Many problems are easily solved once the appropriate frame is known. Recognizing the correct context can be very difficult.
- 4) Semantic Alignment: Two Frame systems describing the same concepts, may not have the same relations. Meanings of similar concepts might be slightly different. However, communication and integration of conceptual knowledge from different sources requires a shared ontology.