

Intelligent Systems: Reasoning and Recognition

James L. Crowley

ENSIMAG 2

Second Semester 2019/2020

Lesson 19

15 April 2020

Narrative Reasoning

Narratives	2
Situation Models (revisited)	3
Reasoning	3
Events	3
Situation Models	4
Observability	4
Narratives	5
Entailment	6
Bayesian Inference from Partially Observable Narratives.....	7
Evidence for a Situation	9
Affective Bias	10
Causal reasoning with narratives.	11
Reasoning about Causality with Narratives.	11
Causal Reasoning with Graphical Models	12

Bibliography:

- 1) P. N Johnson-Laird, Mental models, MIT Press Cambridge, MA, USA, 1989.
- 2) P. Abell, Narratives, Bayesian narratives and narrative actions. *Sociologica*, 1(3), 2007

Narratives



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.

In this lecture we will

- 1) Revisit Situation Modeling to provide a formal foundation for narrative reasoning.
- 2) Show how temporal sequences of situations can be used to provide rich descriptions of unobserved phenomena, and predict the likelihood of future phenomena.
- 3) Show how to collect evidence for and against the likelihood of competing narratives using Bayesian reasoning
- 4) Show how narratives can be used to reason about causality.

Situation Models (revisited)

Situations models are used in cognitive science to describe the mental models that people use to understand and reason. 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.

Situations

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

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

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

Reasoning

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.


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

Narrative Reasoning

Situation Models

A Situation Model is a directed graph of situations.

Nodes are situations. Arcs are events.



The set of Concepts, Relations, and Events are said to compose the Ontology for the Situation. This is sometimes called the "context".

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.

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.

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 provides 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.

Narrative Reasoning

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 $Exists(X)$. $S \vdash \exists(X)$

$Exists(X)$ 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.

Narrative Reasoning

Bayesian Inference from Partially Observable Narratives



Let N be a partially observable narrative composed of a sequence of situations $\{S_n\}$ defined with an Ontology composed of a set of relations $\{R\}$ defined over entities $\{X\}$ represented with Concepts $\{C\}$. We are interested in calculating the likelihood for the narrative given observations of Relations from partially observable Situations within the narratives.

Let $\{X\}$ be the set the entities (instances of concepts) that define the situations. The set of relations $r(X) \in \{R\}$ are defined over subsets of one or more entities $X \subset \{X\}$ from the situation.

Observations of relations from $\{R\}$ provide evidence for or against the narrative. We can include in $\{R\}$ a relation Exists(X) for each entity.

In order to apply Bayesian reasoning, we will replace the Boolean predicates that define relations with probabilistic predicates, $P(r)$. This allows us to define:

$$P(S|r) = P(S) \frac{P(r|S)}{P(r)}$$

The probability for a narrative is the product of the probability for its component relations.

$$P(S|\{R\}) = \prod_{r_n \in \{R\}} P(S|r_n)$$

Our task is made difficult because some of the entities may be unobservable, rendering the relations unobservable. In some cases entire situations are unobservable.

We need to account for both the unobservability of a relation and the observation that the relation is not valid.

Absence of evidence is not evidence of absence!

Narrative Reasoning

Products are not so convenient for accumulating sparse evidence.
To handle this, we will work with Logarithms.

Let $r \in R_s$ be the relations that define the situation.

We will drop the argument (X), for notational simplicity. You can read r as $r(X)$

Let $P(r)$ be the probability that we have observed relation r .

Let $P(S)$ be the probability of situation S .

$$P(r) \cdot P(S|r) = P(S) \cdot P(r|S)$$

and

$$P(r) \cdot P(\neg S|r) = P(\neg S) \cdot P(r|\neg S)$$

thus
$$\frac{P(S|r)}{P(\neg S|r)} = \frac{P(S) \cdot P(r|S)}{P(\neg S) \cdot P(r|\neg S)}$$

The a-priori odds of a situation S are defined as $Odds(S:\neg S) = \frac{P(S)}{P(\neg S)}$

The conditional odds for the situation are $Odds((S:\neg S)|r) = \frac{P(S|r)}{P(\neg S|r)}$

Thus: $Odds((S:\neg S)|r) = Odds((S:\neg S)) \cdot \frac{P(r|S)}{P(r|\neg S)}$

The ratio $\frac{P(r|S)}{P(r|\neg S)}$ is called the conditional likelihood of r .

$$L_r = \frac{P(r|S)}{P(r|\neg S)}$$

Thus the conditional odds for a situation are

$$Odds((S:\neg S)|r) = Odds((S:\neg S)) \cdot L_r$$

The conditional odds for a situation are the a-priori odds times the conditional likelihood for observing the relation.

Narrative Reasoning

Evidence for a Situation

Let us define the “evidence” E_r for a situation as the log of the conditional likelihood.

$$E_r = \text{Log}(L_r) = \text{Log}\left(\frac{P(r|S)}{P(r|\neg S)}\right) = \text{Log}(P(r|S)) - \text{Log}(P(r|\neg S))$$

Observation that an entity or relation is present is positive evidence FOR the situation.

Positive Evidence: if $\text{Log}(P(r|S)) > \text{Log}(P(r|\neg S))$ then $E_r > 0$

Observation that a relation is NOT present is negative evidence against the situation.

Negative Evidence: if $\text{Log}(P(r|S)) < \text{Log}(P(r|\neg S))$ then $E_r < 0$

Inability to observe evidence has no impact.

Absence of Evidence: if $\text{Log}(P(r|S)) = \text{Log}(P(r|\neg S))$ then $E_r = 0$

Notice the difference between “absence of evidence”, and “negative evidence”.

“Negative Evidence” means that an observation indicates that the relation is NOT present in the situations of the situation.

“Absence of evidence” means that it was not possible to observe whether the relation was present or not. This is the case for non-observable situations.

Accumulating evidence for all of the observable relations provides evidence for the situation.

Let us define the a-priori Belief for a situation as the log of the odds for the situation.

$$B_S = \text{Log}(\text{Odds}(S : \neg S))$$

The a-priori Belief is a sort of “bias” for or against believing that the situation has occurred and thus for believing the narrative.

This belief can be influenced by accumulation of evidence: $B(S|\{E_r\}) = B_S + \sum_{r \in \{R\}} E_r$

Narrative Reasoning

Finally, the belief for a Narrative is the sum of the beliefs (log-odds) for the situations that make up the narrative

$$B(N|\{E_r\}) = B_N + \sum_{S \in N} B(S|\{E_r\})$$

Where B_n is an a-priori belief for the narrative.

The "weight" of the evidence depends on the uniqueness of the relation or entity. If a relation is unique to a single situation, then the probability $P(r|\neg S) \rightarrow 0$

And $L_r = \frac{P(r|S)}{P(r|\neg S)} \rightarrow \infty$ and thus $E_r \rightarrow \infty$ (the situation is certainly true)

If the relation is in many situations, $L_r = \frac{P(r|S)}{P(r|\neg S)} \rightarrow 1$. $E_r \rightarrow 0$

If the relation is never in the situation, $L_r = \frac{P(r|S)}{P(r|\neg S)} \rightarrow 0$. $E_r \rightarrow -\infty$

(the situation is certainly false)

However, certainty is dangerous. There may be unknown situations in which the relation exists. To protect from unknown situations that also include the evidence, it is recommended to bound E_r within finite limits. This is a form of skepticism.

Affective Bias

If humans were rational, than the a-priori belief for a narrative would be based only on probability.

$$B_n = \text{Log}(\text{Odds}(N : \neg N))$$

Unfortunately, humans tend to add an additional bias based on whether the narrative causes pleasure or displeasure. This is sometimes referred to as an affective bias: B_A

$$B(N|\{E_r\}) = B_A + B_N + \sum_{S \in N} B(S|\{E_r\})$$

The affective bias can lead to choosing false narratives because they please, and rejecting likely narratives because they displease.

Narrative Reasoning

Causal reasoning with narratives.

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.

Causal Reasoning with Graphical Models

A causal diagram is a graphical tool that enables the visualisation of causal relationships between variables in a causal model. A typical causal diagram will comprise a set of variables represented by nodes. Any variable in the diagram should be connected by an arrow to another variable with which it has a causal influence - the arrowhead delineates the direction of this causal relationship, e.g., an arrow connecting variables A and B with the arrowhead at B indicates a relationship whereby (all other factors being equal) a qualitative or quantitative change in A may cause change in B.