Embodiment in Language understanding: Modeling the semantics of causal narratives

Srini Narayanan
SNARAYAN@CS.BERKELEY.EDU
International Computer Science Institute and University of California, Berkeley
1947 Center Street Suite 560, Berkeley CA 94704 USA

1 Introduction

This abstract describes a computational model that cashes out the belief that metaphor interpretation is grounded in embodied primitives. The specific task addressed is the interpretation of simple causal narratives in the domains of Politics and Economics. The stories are taken from newspaper articles in these domains. When presented with a pre-parsed version of these narratives as input, the system described is able to generate commonsense inferences consistent with the input.

Recent work in Cognitive Semantics (LJ 1980; Lakoff 1994; Talmy 1985; Talmy 1987) suggests that the structure of abstract actions (such as states, causes, purposes, means, etc.) are characterized cognitively in terms of embodied concepts from the domains of force, motion, and space. Such observations have resulted in proposals that abstract reason is grounded in recurring patterns of experience called *Image Schemas* (Lakoff 1987). However, the work in Cognitive Semantics lacks any computational model for such theories, and consequently such ideas cannot currently be used in natural language understanding or problem solving systems.

This work describes an attempt to develop computational models of causal semantics based on a synthesis of results in sensory-motor control (Kandel *et al.* 1991; Latash 1993; Sternberg 1978), insights from structured connectionist systems (Feldman 1989) and linguistic research in Cognitive Semantics. Furthermore, we exploit the fact that the deep semantics of the causal narratives are **dynamic** and arise from a *continuous interaction* between input and memory. This enables the high degree of context sensitivity required since changing input context can dramatically affect the correlation between input and memory and thereby the set of possible

[1] Slightly modified version of the paper in AAAI 1996 Fall Symposium on Embodied Cognition and Action, TR FS-96-02, AAAI Press.

expectations, goals, and inferences.

This abstract gives an overview of the central ideas involved in building an embodied computational model of causal narratives. Details of the various aspects of the design and implementation and connections to the larger **NTL** project under way at UC Berkeley and ICSI can be obtained from the web site *http:://www.icsi.berkeley.edu/NTL* or from (Narayanan 1997).

2 Motivation

Consider the following paragraph that appeared in the NYT in August, 1995.

> In 1991, in response to World Bank pressure, India boldly set out on a path of liberalization. The government loosened its strangle-hold on business, and removed obstacles to international trade. While great strides were made in the first few years, the Government is currently stumbling in its efforts to implement the liberalization plan.

In any interpretation of the narrative above, there are several points to be noted.

1. Institutions are conceptualized as causal agents, causes are mapped as forces, actions are mapped as motions, and goals are mapped as states in a spatial terrain. These mappings are part of a larger composite metaphorical mapping, referred to by Lakoff (Lakoff 1994) as the Event Structure Metaphor.
2. From the fact that the Governments/Institutions are conceptualized as agents, we note the specific causal events that are used in the article, such as apply pressure, respond to pressure, loosen strangle-hold, remove obstacles, stride, and stumble.
3. More interestingly, note that some commonsense inferences that are required for interpreting the article *rely* on the embodied domain of force dynamics and agent motions in space. For instance, the experiential inference that stumbling *leads to* falling can felicitously be transferred to the abstract domain of economic policy through a conventionalized metaphor that *falling maps to failure.* This enables the interpreter to conclude that the Government is likely to fail in its liberalization plan. Many other inferences that are embodied (consider the implications of strangle-hold).
4. While embodied inferences contribute significantly to interpretation, they are asymmetric, context-sensitive and may be overridden by target domain knowledge. For instance, stumble $\Rightarrow$ fall (and the corresponding metaphoric inference of plan failure) is only a default causal inference that is made in the absence of information to the contrary. Such an inference may be non-monotonically disabled in the face of evidence that the liberalization plan is succeeding. As far as we know, no previous effort at modeling metaphoric interpretation has taken this fact seriously and consequently such systems are unable to satisfactorily account for many of the known cognitive phenomena.

3 Computational model

We have built a computational model that can interpret simple causal narratives including the one in Section 2. I will only attempt a crude description of the interpretation process in this abstract. Further information can be obtained from (Narayanan 1997)

The central hypothesis pursued here is that the meaning of embodied causal terms is grounded in patterns generated by our sensory and motor systems as we interact in the world. In this theory, linguistic devices may refer to *schematized processes* that recur in sensory-motor control (such as goal, periodicity, iteration, final state, duration, and parameters such as force and effort). An *x-schema* is a computational model that encodes such patterns.

Results from research on high-level sensory-motor control (Sternberg 1978; Latash 1993; Kandel *et al.* 1991; Bernstein 1967; Arbib 1992) suggest that x-schemas should be active structures capable of modeling coordinated control programs with sequential, conditional, hierarchical and concurrent actions. They should also be capable of modeling asynchronous event-based control with interrupts.

We have built a computational model of x-schemas that satisfies these requirements. The model is a variant of high-level Petri Nets (Reisig 1985). Extensions include the use of parameters and individuated tokens. Technical details of our representation can be found in (Narayanan 1997).

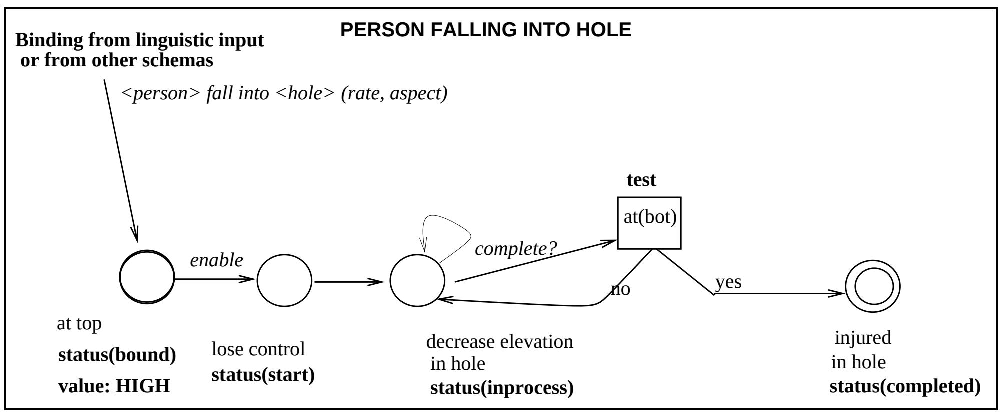

Figure 1: Falling into a hole x-schema

Figure 1 shows a simplified version of the x-schema that encodes the experience of "falling into holes". When active, the "fall into hole" x-schema executes and returns bindings that specify the state of the agent falling into the hole. For instance, at the end of the fall, the agent is in the hole, is injured, and has a goal to get out.

X-schemas have already been shown to model a fundamental problem in interpreting expressions of *abstract actions*, namely the problem of *aspect*.[1] The x-schema based model of aspect, grounded in sensory-motor primitives is able to model cross-linguistic variation in aspectual expressions while resolving paradoxes and problems

[1]Aspectual expressions in a language allow a speaker to direct attention to the internal structure of an event. For instance the use of the *progressive* (as in the *X-ing* construction in English) focuses on the ongoing nature of the underlying process while allowing for inferences that the process has started and that it has not yet completed. Our use of x-schemas helped us solve an outstanding problem in Aspect which relates to how the inherent aspect of a verb (such as the inherent iterativity of the verb *tap*) interacts with grammaticalized aspect such as the *X-ing* construction to produce the proper inferences (*tapping* involves repeated taps, in contrast to climbing or walking).

in model-theoretic and other traditional accounts (Narayanan 1997).

To extend the model to causal narrative interpretation, we assume that the reader's epistemic state consists of a set of features each of which can take on one of a discrete set of values with varying degrees of belief. Our implementation uses a temporally extended probabilistic network (Pearl 1988). Linguistic input and world knowledge jointly condition the reader's epistemic state by clamping features to specific values. In future discussion, we will refer to these feature value pairs as f–structs.

Furthermore, we assume the presence of a systematic correspondence across conceptual domains, that is characterized by a fixed system of conceptual metaphors (LJ 1980; Lakoff 1994) which govern how abstract concepts and domains are understood in terms of the more concrete and "experiential" domains. In our model, such metaphors are first class objects that are belief net nodes capturing inter-domain constraints and correlations.

Table 1 and Table 2 illustrate the I/0 behavior of the implemented system interpreting the newspaper headline *Liberalization plan stumbling*. The input to the system is a set of feature-value pairs (called "f–structs") resulting from a partial parse.

Table 1: Input is a set of F-structs

Feature	**Value**
Event	stumble
Domain	Ec. Policy
Ec. Policy	Liberalization
Aspect	Present-Prog

Comprehending a story corresponds to finding the set of trajectories that satisfy the constraints of the story and are consistent with the domain knowledge. This may involve *filling in* missing values for the abstract domain features, as well as inferring values for unmentioned target features *implied* by the story. The most probable trajectory can then be retrieved as the most likely explanation of the story. Features with highly selective posterior distributions are likely to be present in the recall of the story.

Table 2: Output is a new set of F-structs

Feature	**Value**
Event	stumble
Domain	Ec. Policy
Ec. Policy	Liberalization
Aspect	Present-Prog
Context	**ongoing-plan ∧ difficulty**
Status	**suspended(.8)**
Outcome	**fail (.7)**
Goal	free-trade ∧ deregulation

The result of processing the input in Table 1 is a set of new bindings asserted in the target domain resulting in an updated posterior for other variables. This is the situation shown in Table 2. **Bold** entries correspond to cases where the change from the prior is a result of metaphoric inference. Of particular interest is the *context setting inference* which projects the embodied knowledge that stumble occurs as a result of an obstacle while executing a step (causing an interruption to forward motion) to the target as plan difficulty (causing a temporary suspension). Another interesting binding occurs as a result of the embodied domain knowledge that stumble **may** lead to a *fall* which is mapped onto the target as an enhanced likelihood of *plan failure*. Thus we note that while *stumble* is **not directly** mapped in our system as a meaningful concept in the domain of Economics, through **inferential projection** from maps such as *Falling MAPS TO Plan Failure* and *Obstacle MAPS TO Plan Difficulty* the system is able to assert a target context where an ongoing plan is experiencing difficulty increasing the chance of failure as the outcome. This kind of real-time defeasable inferential projection is a novel feature of our model and we believe a crucial requirement for any model of metaphor interpretation.

Of course, many possible x-schema bindings, especially those that don't activate any conventional metaphor are invalid are thus have no impact on the agent's epistemic state (for example the source inference stumble $\Rightarrow$ losing balance). Thus the inferences that are actually made are context-sensitive and depend on the target domain and the associated set of metaphoric maps.

The resultant target network state shown in Table 2 is now a prior for processing the next input at stage $t = 2$. Background knowledge is encoded as the network state at $t = 0$. Potentially target inferences can go forward and backward in time in the estimation of the most probable explanation of the input story.

4 Results

The program was developed using a database of concepts and stories collected from various on-line sources including the Economist, the Wall Street Journal and the New York Times. The Metaphor Reasoning System is designed to work as a component of an incremental interpreter, accepting partial interpretations (as input f–structs) as input and returning more informative ones (as output f–structs) as output. In its current form, the Metaphor Reasoning System should be viewed as a stage in the interpretation process, after certain parts of the input (namely syntactic, surface semantic, and morphological structures) have been identified.

Figure 2 shows the response of the system to the input *Brazil fell into recession.* (ref. # Fin. Times, 1992). The input f–struct for this story is

```
(Input ((Context EconomicState)(Type recession)(Agent Brazil)
(Event fall-into) (Aspect Perfective)(Ut-type description)))
```

For the output (thresholded at $P \geq 0.6$) refer to Figure 2. Note especially the projection from the **uncontrollable** and **unintentional** nature of *falling* onto the domain of Economic states as the recession being cased by external factors beyond direct control or influence of the specific policy being pursued ($\neg control(loc) \Rightarrow \neg control(Ec.State)$). Of course, this kind of information is often useful to assign responsibility or blame to various agents in the described scenario. For instance, with *fall into recession*, the speaker is inherently indicating that he does not hold any specific administration or policy to blame (probably a normal business cycle recession or some external causative factor). Contrast this to the choice of *walk into recession*, where the speaker is quite likely assigning some responsibility to a failed policy. In fact, quite often intentional aspects of embodied terms get transferred onto more abstract domains as causative factors. While the explicit casual connection is not modeled, even our prototype is able to model some issues of controllability as shown in Figure 2. [2]

[2] In cases where the results are displayed as graphs, the x-axis represents the time step

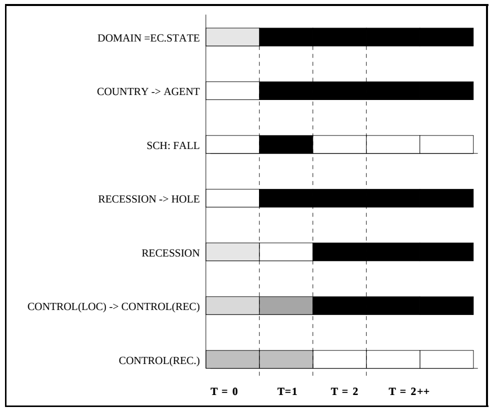

Figure 2: Simple inferences are transfered from embodied to abstract domains for the input "Brazil fell into recession". Of specific interest is the spatial inference that falling into a hole results in being there; this results in setting evidence for the target domain node that Brazil is in recession at the end of the period specified by the input. Also interesting is the inference that the recession was probably unanticipated and uncontrollable, an inference from falling. No such inference is intended or available from processing *Germany has walked into recession (another sentence in the database).*

Our results show that the implemented system is sensitive to both target and source domain context. Using a temporally extended Bayesian network allows us to store the state of the results of previous processing and propagate influences backward and forward in time (shown as the degree of activation in Figure 2 (black is highly active, white highly inactive)). The x-schema representation of the concrete domain is inherently **stateful**, so the current execution state of the x-schema influences future evolutions making the x-schema based inference and the overall system completely context sensitive (both to the previous parse and world states).

Currently our embodied domain theory has about 50 linked x-schemas, while the abstract domain theory is relatively sparse with a belief net of 12 multi-valued variables with at most 5 temporal stages. We have also encoded about 50 metaphor maps from the domains of *health* and *spatial motion*. These were developed using a database of around 30 $2-3$ line newspaper stories all of which have been successfully interpreted by the program. The database of stories used during program development and the full I/O behavior of the implemented system can be found in (Narayanan 1997).

In summary, our results suggest that a large proportion of commonplace descriptions of abstract events project familiar motion and manipulation concepts onto more abstract domains such as economics and politics. This allows non-experts to comprehend and reason about such abstract policies and actions in terms of more universal and commonplace concepts. Our results also provide evidence for our hypothesis that familiar and essential domain of spatial motion is encoded as highly accessible **compiled** knowledge required both for action monitoring and failure re-

---

of interest (0 is the prior, 1 - 3 are the the result of processing the input at that time step, and 4 is the predicted future state of variable in question). In the case of a graph the shading of the relevant cell indicates the degree to which the variable is believed to be *true*, the darker the shading the higher the $degree\_of\_belief$.

covery but also used for fast, parallel, real-time reflex inference in interpretation. Furthermore, our implemented model is able make these inferences in real-time. While the results pertain to how information from embodied terms can be useful in interpreting descriptions of events in the domain of international economics, we believe that the model should generalize to abstract plans, goals, and intentions in any domain. However, this remains future work.

5 Conclusion

This abstract describes a new model for modeling the understanding of narratives such as newspaper stories involving political or economic causation. The central novel ideas investigated are

1. A computational model of narrative understanding by metaphoric mapping from abstract domains, such as politics and economics to the concrete and embodied spatial domain.
2. Representation of the epistemic state as a temporally extended belief net that allows for a single system to be query answering, generative and capable of counterfactual reasoning.
3. The grounding of the deep semantics of the abstract causal terms in body-based active models.

5.1 Acknowledgments

Jerry Feldman and George Lakoff provided significant advice and encouragement and made substantial contributions to the work presented here. I have also benefited greatly from discussions with David Bailey, Dan Jurafsky, Lokendra Shastri, and other members of the Lzero group at ICSI and UC Berkeley. I would like to thank the members of the graduate seminar in Cognitive Linguistics and the RUGS group at UC Berkeley.

References

Arbib, M.A. (1992). *Schema Theory*, In the Encyclopedia Of Artificial Intelligence, 2nd. Edition, edited by Stuart Shapiro, 2:1427-1443, Wiley, 1992.

Bernstein, N. A. (1967). *The Co-ordination and Regulation of Movement*. New York: Pergamon Press.

Feldman, J.A. (1989) *Neural Representation Of Conceptual Knowledge*, Nadel, L. etal eds., Neural Connections, Mental Computation, 68-103, MIT Press, 1989.

Kandel, Eric R., James H. Schwartz, and Thomas M. Jessell(eds.) (1991) *Principles of Neural Science*. New York: Elsevier, third edition.

Lakoff, George. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press.

George Lakoff (1994). *Metaphor*, Internal Document, UC Berkeley, 1994.

Mark L. Latash (1993). *Control of Human Movement*, Human Kinetics Publishers, 1993.

G. Lakoff and M. Johnson (1980). *Metaphors we live by.* Chicago, University Of Chicago Press.

Narayanan, S (1997). *KARMA: Knowledge-based Action Representations for Metaphor and Aspect* Ph.D. Dissertation, Dept. Of Computer Science, University of California, Berkeley, 1997.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, San Mateo, Ca.

Reisig, Wolfgang (1985). *Petri Nets: An Introduction*. Berlin: Springer–Verlag.

S. Sternberg and S. Monsell and R. L. Knoll and C. E. Wright (1978), *The latency and duration of rapid movement sequences: comparisons of speech and typewriting*, Information Processing in Motor Control and Learning, G. E. Stelmach, Academic , New York, "117-152", 1978.

Talmy, Len (1985). *Lexicalization Patterns, Semantic Structures in Lexical Form*, in Language Typology And Symbolic Description, v3: Cambridge University Press, 1985.

Talmy, Len (1987) *Force Dynamics in Language*, Tech Report. Institute For Cognitive Science, UC Berkeley, 1987.