Moving Right Along: A Computational Model of Metaphoric Reasoning about Events

Srini Narayanan
snarayan@icsi.berkeley.edu
International Computer Science Institute and University of California, Berkeley
1947 Center Street Suite 600, Berkeley CA 94704 USA

Abstract

This paper describes the results of an implemented 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.

Introduction

This paper describes the results of an implemented computational model that cashes out the belief that metaphor interpretation is grounded in the embodied primitives of motion and manipulation. 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.

Work in Cognitive Semantics (Lakoff & Johnson 1980; Talmy 1987; Sweetser 1990; Johnson 1987; Langacker 1987; Lakoff 1994) suggests that the structure of abstract actions (such as states, causes, purposes, means) are characterized cognitively in terms of *image schemas* which are *schematized* recurring patterns from the embodied domains of force, motion, and space. However, the work in Cognitive Semantics lacks any computational model for such theories, and consequently these ideas cannot currently be used in natural language understanding or problem solving systems.

This paper describes a project that provides evidence through computer simulation that a key reason for using motion words and phrases is that it allows for the deep semantics of causal narratives to be *dynamic* and arise from a *continuous interaction* between input and memory. Since knowledge of moving around or manipulating objects is essential for survival, it has to be highly compiled and readily accessible knowledge. Representations meeting these criteria must be context sensitive and allow changing input context to dramatically affect the correlation between input and memory and thereby the set of possible expectations, goals, and inferences. Speakers are able to felicitously exploit this context-sensitivity in specifying important information about abstract actions and plans that take place in complex, uncertain and dynamically changing environments.

Motivation

Consider the following narrative about India's march toward liberalized economics.[1]

Example 1 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 Example 1, note that institutions are conceptualized as causal agents, causes as forces, actions as motions, and goals as states in a spatial terrain. These mappings are part of a crosslinguistic metaphor system called the Event Structure Metaphor (Lakoff 1994) which is the general name for projections from the concrete experiential domain of forces and spatial motion (source domain) to the abstract domain of causes, actions, and events (target domain). Following from the fact that institutions are conceptualized as agents, specific causal events are attributed as effected by or affecting the institution; such as apply pressure, respond to pressure, loosen strangle-hold, remove obstacles, stride, and stumble. Commonsense inferences that are required for interpreting the article often *rely* on our experience of force dynamics and motion in space. For instance, the inference that stumbling *leads to* falling can felicitously be transferred to the abstract domain of economic policy through a conventionalized metaphor that *falling* $\mapsto$ *failure*. This enables the interpreter to conclude that the government is likely to fail in its liberalization plan. Many other inferences rely on the source domain (consider the implications of strangle-hold).

[1] While this story appeared in the New York Times in 1995, the reader is invited to convince herself of the ubiquity of the mappings discussed (albeit at the risk of severely impaired newspaper reading pleasure).

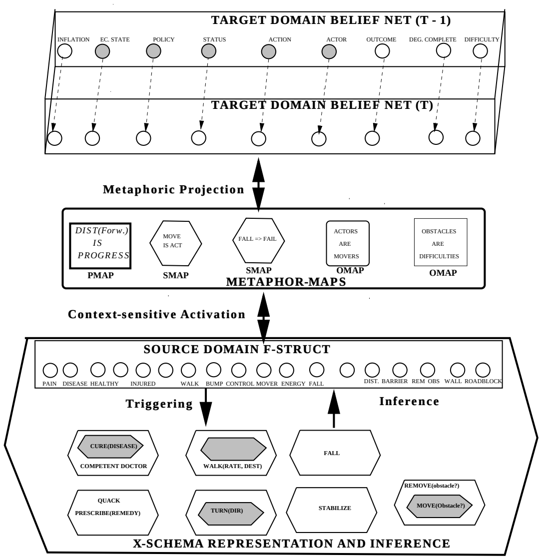

Figure 1: Metaphors capture systematic correlations between features of different domains.

While source domain 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 target domain evidence that the liberalization plan is succeeding.

In summary, we note that a large proportion of commonplace descriptions of abstract events seem to project embodied, familiar, 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 familiar and universal embodied concepts. The fact that the metaphoric inferences are context-sensitive, immediate, and defeasible set up fairly strong representational requirements for a metaphor interpretation system.

Model

The specific hypothesis pursued here is that the meaning of motion and manipulation terms is grounded in patterns generated by our sensory and motor systems as we interact in the world. Systematic metaphors project these features onto abstract domains such as Economics enabling linguistic devices to use motion terms to describe abstract actions and processes. Figure 1 shows the basic computational architecture of the implemented system. As shown in the figure the system has three main components, namely the **source domain**, the **target domain** and the **metaphor maps**. These components are discussed below.

The source domain

We hypothesize that the causal theory of the familiar and essential domain of embodied motion is encoded as highly accessible *compiled* knowledge used both for action monitoring and failure recovery and for fast, parallel, real-time reflexive inference in interpretation. In previous work (Bailey *et al.*, 1997; Narayanan, 1997), we have referred to our basic model of events as **x-schemas**. Our model is based on results in sensory-motor control (Pearson 1993; Bernstein, 1967) and linguistic research in Cognitive Semantics. Formally, the computational model is an extension to Stochastic Petri Nets (Murata 1989). A Petri net is a bipartite graph containing *places* (drawn as circles) and *transitions* (rectangles). Places hold *tokens* and represent predicates about the world state or internal state. Transitions are the active component. When all of the places pointing into a transition contain an adequate number of tokens (usually 1) the transition is *enabled* and may *fire*, removing its input tokens and depositing a new set of tokens in its output places. The most relevant features of Petri nets for our purposes are their ability to model events and states in a distributed system and cleanly capture sequentiality, concurrency and event-based asynchronous control. Our extensions to the basic Petri net formalism include typed arcs, hierarchical control, durative transitions, parameterization, typed (individual) tokens and stochasticity. For this paper, the crucial fact about our representation is that it is *active* with a well specified real-time execution semantics that can be used for acting and reacting in dynamic environments or for context sensitive simulative inference in language understanding.

The central idea behind our model is that the reader interpreting a phrase that corresponds to a motion term is in fact performing a mental simulation of the entailed event in the current context. The basic idea is simple. We assume that people can execute x-schemas with respect to structures that are not linked to the body, the here and the now. In this case, x-schema actions are not carried out directly, but instead trigger simulations of what they would do in the imagined situation. We model the physical world as other x-schemas that have i/o links to the x-schema representing the planned action.

In our implementation, source domain structure is encoded as connected x-schemas. Our model of the source domain is a dynamic system based on inter-x-schema *activation*, *inhibition* and *interruption*. In the simulation framework, whenever an executing x-schema makes a control transition, it potentially modifies state, leading to asynchronous and parallel triggering or inhibition of other x-schemas. The notion of state as a graph marking is inherently distributed over the network, so the working memory of an x-schema-based inference system is distributed over the entire set of x-schemas and source

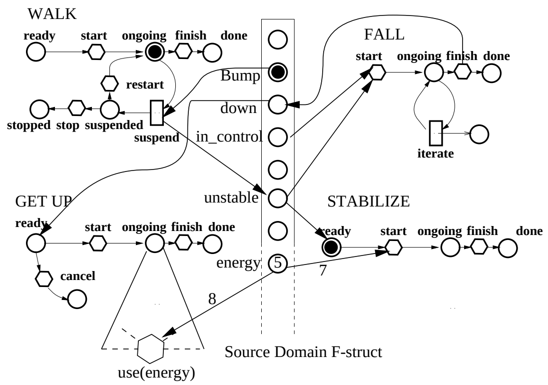

Figure 2: Source Domain is a x-schema simulation environment used for inference.

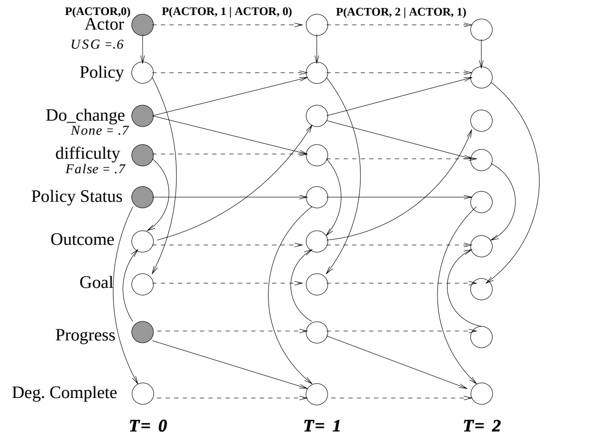

Figure 3: Target Domain is a temporally extended Belief net.

domain f-structs (see Figure 1). Of course, this is intended to model the massively parallel computation of the brain.

Figure 2 depicts a simplified x-schema model of walking and reacting to obstacles. For instance, during a walk (specified by a token in the *ongoing* phase of the WALK x-schema) encountering an unanticipated **bump**, you become *unstable*. [2] This *may lead* to a FALL unless you are able to simultaneously *expend energy* and STABILIZE, in which case you may *resume* the *interrupted* walk. If you are unable to STABILIZE, and thus FALL, you will be **down** and **hurt**. In order to *start walking* again you will have to GET UP and be standing and in control again.

An important and novel aspect of our source domain representation is that the same system is able to respond to either direct sensory-motor input **or** other ways of setting the agent state (such as linguistic devices). This allows for the same mechanism to perform simulative reasoning and generate inferences from linguistic input as well as be used for high-level control and reactive planning. There is some biological evidence to support this view (Rizzolatti et al 1996; Jeannerod 1997; Tanji & Shima 1994) that planning, recognition and imagination share a common representational substrate.

Target domain representation

The structure of the abstract domain (the domain of international economic policies) encodes knowledge about Economic Policies. We require that our representation be capable of a) representing background knowledge (such as US is a market economy), b) modeling inherent target domain structure and constraints (high-growth may result in higher inflation), and c) be capable of computing the impact of new observations which may from direct input ("US economy is experiencing high-growth"), or from metaphoric (or other) inferences ("Economy stumbling"). Furthermore, these different sources of evidence have different degrees of believability, and the representation must provide a framework for their combination. For all these reasons, we chose to represent the target domain as a Belief network (Pearl 1988; Jensen 1996). Belief networks are the dominant methodology for reasoning with uncertain knowledge sources. They provide a principled and coherent semantics based on probability theory, which allows us to study the joint impact of metaphoric inference, background knowledge and inherent target structure using well understood, off-the-shelf algorithms.

Our model of the target domain consists of multiple copies (up to 4) of a temporally extended Belief net (Dean & Wellman 1991), representing different time slices. The structure of the target domain for three temporal slices of the Belief network is shown in Figure 3. Within a single temporal slice, the nodes of the network correspond to economic variables which can take on different values. For instance, in Figure 1, we have a node corresponding the the economic actor which can be instantiated to be the US government, IMF, Indian Government, etc. Links within a single time slice model the probabilistic dependence between variables. For instance, there is a link between the actor variable and the policy variable, which models the fact that if we knew the actor in question (US Government) we would have a good idea of the policy (free-market economy). The strength of this belief is quantified as the conditional probability table $P(Policy|Actor)$. Links between nodes at different time slices encode the conditional probability of a variable's value at time t , given its value at $t-1$. For instance, the link $P((Actor,1)|(Actor,0))$ (ref. to the top of Figure 3) results in the conditional probability table (CPT) that corresponds to the probability of a specific actor being instantiated at time $t=1$, given the value of the

[2] In fact, the simulation is of finer granularity in that it is during an ongoing STEP (subschema of WALK), that the interruption occurs. This is not shown to simplify exposition.

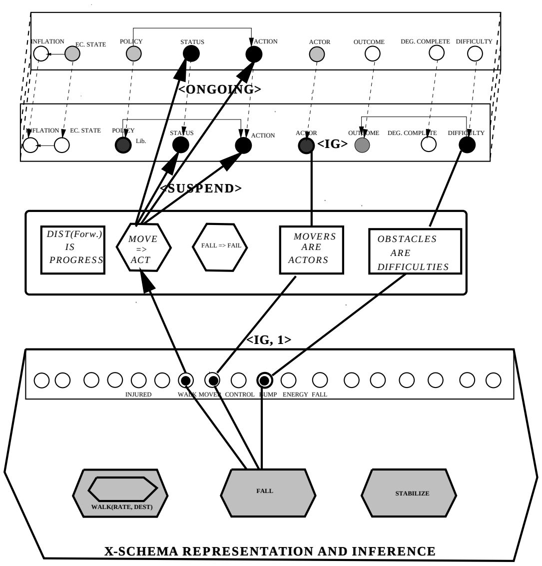

Figure 4: Metaphor maps project source f-struct values as target domain evidence on one or more slice of the target domain Belief net.

actor at time $t = 0$. These values are default values and are often overridden by specific assertions as we will soon see in detail in the next section. From such local conditional probability tables, BELIEF UPDATE algorithms (Jensen 1996) compute the global posterior probabilities for the entire network propagating influences backward and forward in time.

Metaphor maps

In our model, metaphor maps connect the x-schema based representations to the belief network representing knowledge about international economics. Such maps project specific results of x-schema executions by projecting specific source domain f–struct values to the target domain by asserting new *evidence* at one or more time slices of the temporally extended Belief net. Figure 4 shows the projection of "stumbling" onto the target domain. We will return to this example in the next section.

Our model currently includes three different types of embodied maps. One type of map corresponds to **ontological** maps (Lakoff & Johnson 1980) which map entities and objects between embodied and abstract domains. Such maps are called OMAPS. In general, one central function of OMAPS is to map the fillers of various case-roles of an *event phrase* across domains. A second type of map projects *events*, *actions*, and *processes* from embodied to abstract domains. In keeping with our representation, we will call such maps Schema maps or SMAPS. An important function of SMAP projection is to **invariantly** map the *aspect* of the embodied domain event onto the target domain. A third type of map projects x-schema parameters from source to target domains. Such maps are called x-schema parameter maps (PMAPS). Examples include maps that project velocities onto the abstract domain as the rate of progress made; or distance traveled onto the abstract domain as degree of completion of a plan.

I/O behavior

In this model, a story represents the specification of a *partial* trajectory over epistemic states. This is done by clamping some of the Belief network nodes to specific values. The remaining features are estimated using known target domain inter-feature correlations as well as from metaphoric projections from the highly compiled embodied domain knowledge (x-schemas). Metaphoric projections of x-schema executions may clamp target features to specific values (by creating new *evidence* on the target domain belief net shown in Figure 3).

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 or creating new evidence on the Belief network. 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 $\wedge$ difficulty**
Status	**suspended (.8)**
Outcome	**fail (.7)**
Goal	free-trade $\wedge$ 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 stumbling 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 stumbling leads 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.

Of course, many possible x-schema bindings, especially those that don't activate any conventional metaphor are invalid and 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.

Results

Currently our embodied domain theory has about 100 linked x-schemas, while the abstract domain theory is relatively sparse with a belief net of about 20 multi-valued variables with at most 4 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 30 $2-3$ phrase fragments from newspaper stories all of which have been successfully interpreted by the program. All the examples in this section have been taken from our database. Both the database and details of the programs behavior can be found in (Narayanan, 1997).

X-schema parameters

Distances, *speeds*, *force-values*, *sizes* and *energy-levels* are obviously important perceptual and motor control parameters, but with PMAP projections, they become important descriptive features of events in abstract domains including impacting early parsing decisions of inferring semantic role assignments as shown in Figure 5.

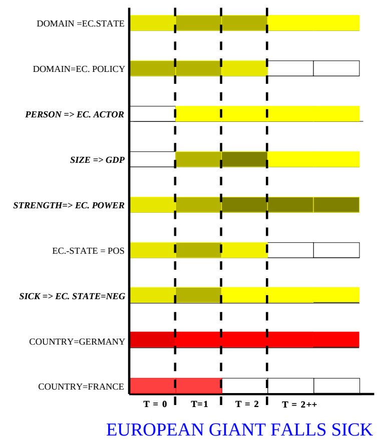

Figure 5: Processing the input "European Giant falls sick". Note that the OMAP *Person* $\Rightarrow$ *Economic Actor* coupled with the PMAP that projects size onto *GDP* in the Economic context, results in asserting evidence in the target that identifies Germany as the country with the largest *GDP* and thus as the referent of the *subject* of the input sentence. Note also that France has some posterior probability of being the referent as well (darker implies higher belief values, lighter implies higher belief in the falsity of the feature-value).

In our examples, we were able to use PMAPs to map size parameters like *giant steps*, *large step*, *small steps*, *great leap forward* (including the Chinese Economic Reform); speed parameters in expressions like *slow progress*, *slowed down*, *sprint*, *jog*, and *long, painful slide into recession*; rate and manner parameters in *crawl*, *leap*, *trod*, *plod*, *slog*, *lurch* and *slither*; distance related parameters in expressions like *almost there*, *long way to go*, *halfway there*, and *a little further*. Force magnitudes and durations were also routinely projected as in *grip*,*tear down hold back*.

Aspectual inferences

(Narayanan 1997) outlined an x-schema based model of of *aspect* (the internal temporal structure of events) which is able to detect and model subtle interactions between grammatical devices (such as morphological modifiers like *be + V-ing* (progressive aspect) versus *has V-ed* (perfect aspect)) and the inherent aspect of events (such as the inherent iterativity of *tap* or *rub*, or the punctuality of *cough* or *hit*). In examining our metaphor database, we found aspectual distinctions to be *invariantly projected* across domains. Furthermore, the high-frequency of aspectual references in describing events makes it important to model the relevant semantic distinctions.

The following examples from our database were successfully processed by the metaphor reasoning system. We already saw in detail how the concept *is stumbling*

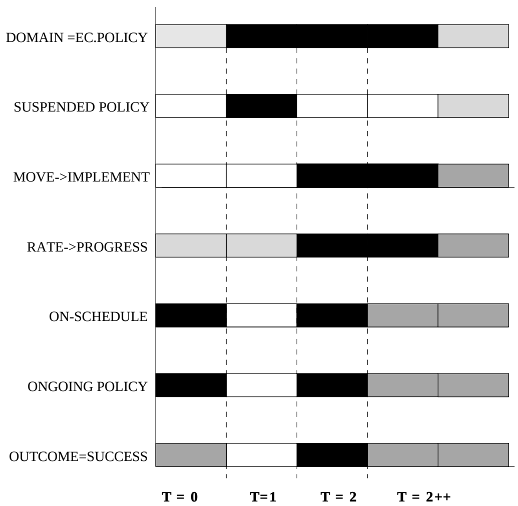

Figure 6: Processing back-on-track. Note the inference that the policy was ongoing, on-schedule, and successful before going off-schedule and being suspended, and then resuming on-track.

is projected onto the domain of international economic policies as an *interruption* to an *ongoing* policy. Other easily modeled cases include the focus on the consequent state that is signaled by the use of the *perfect* aspect such as *have robbed*, *has been lurching forward*, *has sidestepped*. We could also nicely model several other high frequency aspectual expressions such as *start to pullout*, *on the verge of*, *still trying to climb out of recession* and metaphoric expressions of aspect such as *set out*, *remain stuck in recession*, *on-track*, and the interesting phrase *back-on-track* (shown in Figure 6). In summary, almost every event description had an aspectual component, and so we believe attention to the details of the semantics of verbal aspect is essential even to interpret the simplest of event phrases and distinctions. We believe our model is unique in integrating the semantics of aspect with metaphoric interpretation.

Goals, resources

It is well known that narratives are generally about *goals* (their accomplishment, abandonment, etc.) and *resources* (their presence, absence, levels, etc.) (Wilensky 1983; Schank & Abelson 1977; Carbonell 1982). However, in our experiments, we found that embodied motion and manipulation terms may in fact be *compactly coding* for these features as well. Note that this assertion is quite different from (Carbonell 1982) who hypothesized that entire strategies and plans (proof trees) ware invariantly transferred. In our theory, it is the *key-event* such as the thwarting of a goal, or the absence or production of a needed resource, etc. that is asserted using embodied terms. Narratives are able to exploit this feature of x-schema representations to assert

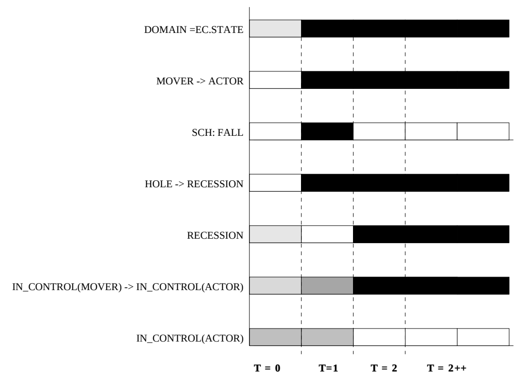

Figure 7: Simple inferences are transfered from embodied to abstract domains for the input "Brazil fell into recession".

changing goals and resources. Amount of *energy* usually maps to *resource* levels as in *slog*, *anemic*, *sluggish* or *bruised and bloodied*, or *stagger to their feet*. Similarly *tearing barriers* or *lightening burdens* are able to assert conditions where an impediment to goal achievement has now been removed. Compare this to the expression *go around* or *sidestep* where the strategy is one of avoidance rather than direct confrontation. Similarly *slippery slopes*, *slipperiest stones*, *slide into recessions*, get projected through SMAPS as the possible thwarting of goals due to unanticipated circumstances. *Falling* is interesting in this regard (ref. Figure 7). In all the cases where a country was described as *falling* into recession, we never saw a case in which the country's administration was directly blamed as being able to control the downturn, a fact directly projectable from the fact that falling is not controllable (an obvious and easy inference about fall). This is shown in Figure 7. No such inference is intended or available from processing *Germany has walked into recession*.

In general, we found stories in the abstract domain to often be about the complex notion of controllability, monitoring problems, and policy adjustments. Again, monitoring, changing directions, rates, etc. are obviously common in sensory-motor activity, and so again using these features and appropriate projections allows the speaker to communicate monitoring and control problems in abstract plans. Other successfully interpreted examples include *taking a cautious step in the right direction* or the beautiful example *Economic reform is like crossing a river by feeling for the stones*, including the concept of *inflation is the slipperiest stone*.

Novel expressions

As (Lakoff 1994; Gibbs 1994) and other researchers point out, a variety of novel expressions in ordinary discourse as well as in poetry make use of highly conventionalized mappings such as the ones described here. In fact, the implemented system is able to interpret novel expres-

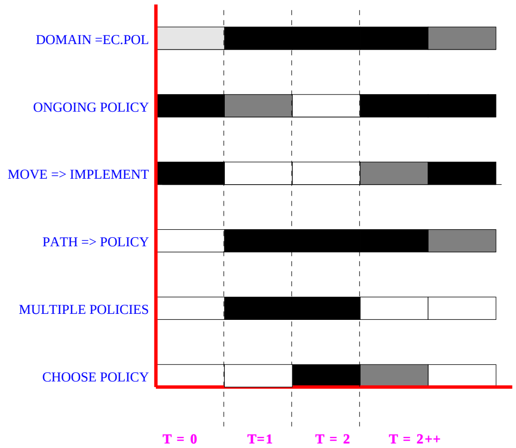

Figure 8: Novel expressions may use conventional conceptual mappings.

sions which it has never seen in the context of abstract actions and plans. Figure 8 shows how knowing the concrete domain meaning of crossroads (multiple possible paths) and the event structure metaphor mapping conceptual features from the domain of motions to abstract actions allows for a reasonable interpretation that there are multiple plans and the agent is at a choice point.

Other examples of novel expressions (in our database) correctly interpreted by our program include *roadblocks*, *anemic recovery*, *lurching forward*, *long, painful slide*, *treading on toes*, and the beautiful *stumble over rocky relationship*.

Multiple source domains

Multiple source domains pose no problem for the system, as long as they are interpretable and coherent in the *target*. Figure 9 shows the system's response to the input *Stocks down. Healthy again.*

Agent attitudes and affects

We found agent attitudes to be essential ways of encoding anticipatory conditions, motivation and determination of agents involved. We have implemented some of this in the prototype system. For instance *bold* (Example 1) encodes determination in the face of anticipated obstacles/counterforces ahead. In the current model this is directly encoded as the semantics of *bold* in the context of the embodied domain (anticipating some counterforces at future time steps). As in the case of stumbling, obstacle at the next time step gets translated to anticipated difficulty at the $t + 1$ temporal slice. Determination to keep on the path gets translated as a reduced *prior* chance of policy change.

The point to note here is that the embodied term *bold* codes for possible future obstacles, and the readiness to deal with them. This projects onto the target as the possibility of future difficulty and asserts the status of removing difficulty to be READY, indicating the readiness of the policy maker to deal with such a future difficulty.[3]

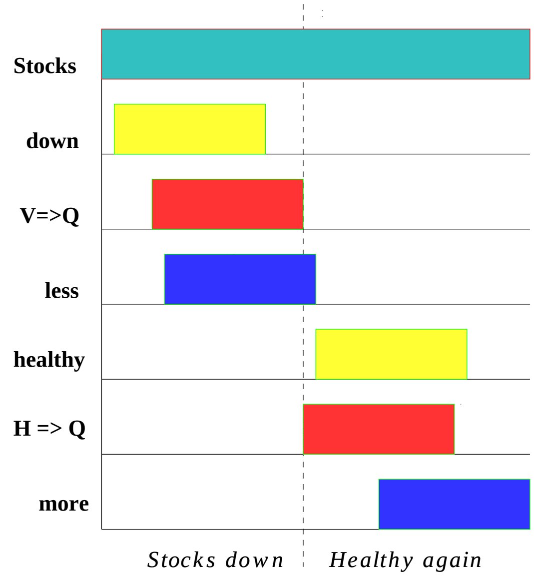

Figure 9: Multiple source domains can be serially used as long as they are coherent in the target domain. Here the first input *Stocks down* activates the *Less IS Down* metaphor, while the second input *Healthy again* activates the *More IS Healthy* metaphor.

Communicative intent and metaphor

One of the important aspects of communication involves specifying evaluative judgments of situations to communicate speaker intentions and attitudes. We hypothesize that the cross-linguistic prevalence of the use of embodied notions of force and motion to communicate aspects of situations and events is linked to the ease with which evaluative aspects can be communicated in experiential terms. To study this phenomenon, we enhanced the target domain Belief network (see Figure 3) to include information about the interpreter's bias toward specific actors and policies. We can now set the interpreter to be biased favorably toward a specific actor (like World Bank) or a specific policy (liberalization). This directly influences both conditional belief of some outcome variables (so a free-market biased interpreter would consider tariff reduction as a successful policy) or could result in different source domain inferences as in the example below.

With these additions, our implemented system was able to distinguish between the following sentences (second is from Example 1).

Example 2

1. Government loosened strangle-hold on business.

[3] Another example where linguistic devices are able to exploit the CONTROLLER distinctions between READY and START, a fine-grained control distinction that is useful for motor control but proving quite indispensable for language.

2. Government deregulated business.

■

Both sentences communicate the same fact in the domain of economics, namely the the situation corresponding to business deregulation. But the source domain inference of "stranglehold" is able to assert the detrimental nature of Government control leading to the possible eventual "demise" of business.

In another example, we tested the program with the example "World Bank prescribed Structural Adjustment Program (SAP) bleeding Indian Economy" under different prior speaker attitudes toward World Bank. In the three cases, we set the prior belief of the speaker to be positive, neutral or negative with respect to the World Bank. In the positive case, the prior belief of the interpreter activates the CURE x-schema. Here, the target domain inferences is one of *ongoing therapy*. In the negative case, the prior belief of the interpreter activates the HARMER schema. where the source domain inferences is one of systemic harm and eventual death. In the neutral case, the prior of the interpreter activates the TREAT x-schema. Here, there is a conflict between CURE and MISTAKEN THERAPY, where the cure is not working.

One crucial difference in the three cases is in the positive case, the the outcome of a *cure* is asserted as succeeding for India, in the negative case the outcome of the policy is asserted as *unsuccessful* for India, while in the neutral case it is *ambiguous*. Thus in the three cases, we are able to model how changes in prior evaluation of a situation can be used to compute what the *meaning* of an utterance is. Crucially, the difference seems to be in which **source domain** schema gets invoked, and the resulting inferences. We know of no other implemented model of metaphor understanding that can reason about these phenomena.

Discussion

It is now generally accepted that metaphor interpretation requires the ability to explicitly represent the source and target domains as well as the metaphor maps themselves. Metaphoric reasoning with knowledge-rich sources and targets and explicit maps have been the primary method of choice for several implemented metaphor interpretation systems (Martin 1990; Barnden *et al.* 1994; Carbonell 1982; Indurkhya 1992). These approaches share many goals and bear some similarities with the work described here. However, there are some crucial differences as well.

First, our representation of actions and events with durations is more fine-grained than other systems we are aware of. Specifically, we believe our system to be novel in being able to model rich temporal and aspectual inferences across domains. Such fine-grained semantic distinctions are routinely exploited by metaphors found in ordinary discourse. Second, our use of a temporally extended Belief network to represent target domain knowledge allows us to uniformly combine direct linguistic input and background knowledge with results of metaphoric projections in a single normative framework. It allows us to study the evidential interaction of these different sources in interpretation, while previous efforts have focussed on isolating one or more of these components. Third, while most approaches require extra resources to process novel expressions, our approach explains why *some* novel expressions can be processed with no additional resources (consistent with psychological observations (Gibbs 1994)). Fourth, our approach is quite unique in being able to exploit **implicit** *evaluative* information and speaker *intent* which we believe is often the reason to choose embodied expressions in the first place. Finally, evidence from a recent study by Joe Grady (Grady 1997), suggests that complex metaphoric maps are composed from simple experiential correlations, consistent with the work reported here.

Conclusion

This paper outlined an implemented computational model for interpreting simple narratives such as newspaper story fragments and headlines involving political or economic causation. The central novel ideas investigated are a) a model of narrative understanding by metaphoric mapping from abstract domains to concrete and embodied domains and b) the grounding of the deep semantics of the abstract causal terms in body-based *active* models.

It is somewhat interesting that even our simplistic model is able to detect rather subtle differences in speaker intent and communicative goals. We believe the choice of the motion term is often a compact and efficient way to encode such information. Conversely, the unconscious choice by a speaker of an embodied term can give the hearer significant clues as to the prior belief and intent of the speaker, obviously something that needs far more exploration.

Acknowledgements

Thanks to Jerry Feldman, George Lakoff, Stuart Russell, Robert Wilensky, and the NTL and RUGS groups at UC Berkeley and ICSI.

References

Bailey, D., Feldman, J.A., Narayanan, S., & Lakoff, G. (1997). Modelling Embodied Lexical Development. *Proc. 19th Conference of the Cognitive Science Society*, 1997.

Barnden, J. Helmreich, L. Iverson, E. Stein, G.C. (1994). An integrated implementation of simulative, uncertain, and metaphorical reasoning about mental states. *Priniciples of KR and Reasoning: Proceedings of the Fourth International Conference* (Bonn, Germany, 24-27 May 1994), San Mateo, CA: Morgan Kaufmann.

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

Carbonell, J. (1982) Metaphor Comprehension. *Strategies for Natural Language Processing*, 413–433, Lawrence Earlbaum, 1982.

Dean, Tom and Wellman, Michael, (1991). *Planning and Control.* Morgan Kaufman Series in Representation and Reasoning, 1991.

Gibbs, R. Jr. (1994). *The Poetics Of Mind.* Cambridge University Press, 1994.

Grady J. (1997). *Foundations of meaning: Primary metaphors and primary scenes.* UC Berkeley Dissertation, Dept. of Linguistics, Fall 1997.

Indurkhya, B. (1992). *Metaphor and Cognition.* Kluwer Academic Publishers.

Jeannerod, M.(1997). *The Cognitive Neuroscience of Action*, Blackwell, 1997.

Jensen, F. (1996). *An Introduction to Bayesian Networks.* Springer-Verlag ISBN 0-387-91502-8.

Johnson, M. (1987). *The Body In The Mind: The Bodily Basis of Meaning, Imagination, and Reason.* University Of Chicago Press, ISBN 0-226-40318-1.

Langacker, R. (1987). *Foundations of Cognitive Grammar I: Theoretical Prerequisites.* Stanford University Press, Stanford.

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

—— (1994). What is Metaphor?. *Advances in Connectionist Theory. V3 : Analogical Connections*, V3,1994.

Martin, J. (1990). *A Computational Model of Metaphor Interpretation.* Academic Press, NY, 1990.

Murata, T. (1989). Petri Nets: Properties, Analysis, and Applications. In *Proc. IEEE–89, V77, Number 4, April 1989, pp. 541-576*.

Narayanan, S. (1997). *Knowledge-based Action Representations for Metaphor and Aspect (KARMA).* PhD thesis, Computer Science Division, EECS Department, University of California at Berkeley.

—— (1997). Talking The Talk is Like Walking the Walk: A Computational Model of Verbal Aspect. *Proceedings of the Nineteenth Annual Conference on Cognitive Science Society* 1997.

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

Pearson K.G. (1993). Common Principles of Motor Control in Vertebrates and Invertebrates. *Ann. Review Of Neuroscience*, 1993, 16:265-97.

Rizzolatti, et al (1996). Premotor Cortex and the recognition of motor actions. *Cognitive Brain Research*, 3 (1996) 131-141.

Schank, R.C. & and Abelson, R.P. 1977. *Scripts, Plans, Goals, and Understanding: An inquiry into human knowledge structures.* Hillsdale, NJ:Erlbaum 1977.

Sweetser, E. (1990). *From Etymology to Pragmatics: The mind-as-body metaphor in semantic structure and semantic change.* Cambridge, UK: Cambridge University Press, 1990.

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

Tanji J. and Shima S. (1994). The supplementary motor area in the cerebral cortex. *Nature*, vol. 371, issue 6496, (SEP 29, 1994) : pp. 413-416.

Wilensky, R. (1983). *Planning and Natural Language Understanding.* Addison Wesley, 1983.