LifeNet: A Propositional Model of Ordinary Human Activity

Push Singh
Media Lab, MIT
Cambridge, MA, 02139
push@media.mit.edu

William Williams
Media Lab, MIT
Cambridge, MA, 02139
whw@media.mit.edu

ABSTRACT

We describe *LifeNet*, a new common sense knowledge base that captures a first-person model of human experience in terms of a propositional representation. LifeNet represents knowledge as an undirected graphical model relating 80,000 egocentric propositions with 415,000 temporal and atemporal links between these propositions. We explain how we built LifeNet by extracting its propositions and links from the Open Mind Common Sense corpus of common sense assertions, present a method for reasoning with the resulting knowledge base, evaluate the knowledge in LifeNet and the quality of inference, and describe a knowledge acquisition system that lets people interact with LifeNet to extend it further.

INTRODUCTION

We are interested in building 'common sense' models of the structure and flow of human life. Today's computer systems lack such models—they know almost nothing about the kinds of activities people engage in, the actions we are capable of and their likely effects, the kinds of places we spend our time and the things that are found there, the types of events we enjoy and types we loathe, and so forth. By finding ways to give computers the ability to represent and reason about ordinary life, we believe they can be made more helpful participants in the human world.

An adequate common sense model should include knowledge about a wide range of objects, states, events, and situations. For example, a common sense model of human life should enable the following kinds of predictions:

- When someone is thirsty, it is likely that they will soon be drinking a liquid beverage.
- When someone is at an airport, it is likely they possess a plane ticket.
- When someone is typing at a computer, it is possible that they are composing an e-mail.
- When someone is crying, it is likely that they feel sad or are in pain.
- After someone wakes up, they are likely to get out of bed.

K-CAP'03, October 23–25, 2003, Sanibel Island, Florida, USA.

Most previous efforts to encode common sense knowledge have made use of relational representations such as frames or predicate logics. However, while such representations have proven expressive enough to describe a wide range of common sense knowledge (see Davis [1] for many examples of how types of common sense knowledge can be formulated in first-order logic, or the Cyc upper level ontology [2]), it has been challenging finding methods of default reasoning that can both make use of such powerful representations and also scale to the number of assertions that are needed to encompass a reasonably broad range of common sense knowledge. In addition, as a knowledge base grows, it is increasingly likely that individual pieces of knowledge will suffer from bugs of various kinds; it seems necessary that we find methods of common sense reasoning that are tolerant to some errors and uncertainties in the knowledge base.

However, in recent years there has been much progress in finding ways to reason in uncertain domains using less expressive propositional representations, for example, with Bayesian networks and other types of graphical models. Could such methods be applied to the common sense reasoning problem? Is it possible to take an approach to common sense reasoning that begins not with an ontology of predicates and individuals, but rather with a large set of propositions linked by their conditional or joint probabilities? Propositional representations are less expressive than relational ones, and so it may take a great many propositional rules to express the same constraint as a single relational rule, but such costs in expressivity often come with potential gains in tractability, and in the case of common sense domains, this trade-off seems to be rather poorly understood.

The potential benefits of a proposition representation go beyond just matters of efficiency. From the perspective of knowledge acquisition, interfaces for browsing and entering propositional knowledge are potentially much easier to use because they do not require that the user learn to read and write some complex syntax. From the perspective of applying common sense reasoning within applications, propositional representations have such a simple semantics that they are likely quite easy to interface to. Thus, while propositional representations may be less expressive and require a larger ontology of propositions than relational representations for the same domain, they are in many ways easier to build, understand and use.

In this paper we explore such questions by describing *LifeNet*, a new common sense knowledge base that captures a first-person model of human experience in terms of a propositional representation. LifeNet represents knowledge as a graphical model relating 80,000 egocentric propositions with 415,000 temporal and atemporal links between these propositions, e.g.

- I-put-my-foot-on-the-brake-pedal → I-stop-a-car
- I-pour-detergent-into-wash → I-clean-clothes
- I-put-quarter-in-washing-machine → I-clean-clothes
- I-am-at-a-zoo → I-see-a-monkey
- I-put-on-a-seat-belt → I-drive-a-car
- I-put-a-key-in-the-ignition → I-drive-a-car

We explain how we built LifeNet by extracting its propositions and links from the Open Mind Common Sense corpus of common sense assertions supplied by thousands of members of the general public, present a method for reasoning with the resulting knowledge base, evaluate the knowledge in LifeNet and the quality of inference, and describe a knowledge acquisition system that lets people interact with LifeNet to extend it further.

LIFENET

LifeNet is a large-scale temporal graphical model expressed in terms of 'egocentric' propositions, e.g. propositions of the form:

- I-am-at-a-restaurant
- I-am-eating-a-sandwich
- It-is-3-pm
- It-is-raining-outside
- I-feel-frightened
- I-am-drinking-coffee

Each of these propositions is a statement that a person could say was true or not true of their situation, perhaps with some probability. In LifeNet these propositions are arranged into two columns representing the state at two consecutive moments in time, and these propositions are linked by joint probability tables representing both the probability that one proposition follows another, and also the probability of two propositions being true at the same time. A small sample of LifeNet is shown in Figure 1 below:

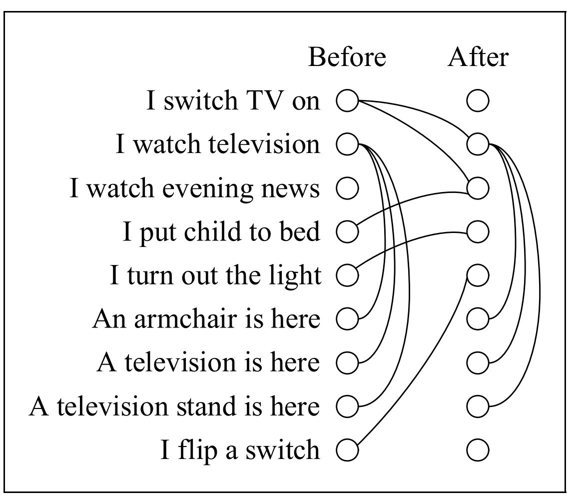

Figure 1. A sample of LifeNet

Figure 1 barely hints at the extent of LifeNet. Presently LifeNet consist of a total of 80,000 propositional nodes linked by 415,000 joint probability tables between pairs of nodes. More details about how the LifeNet model is generated are given later in this paper.

What kinds of questions might we ask of LifeNet? Given such a model, we can engage in several types of temporal reasoning:

- **Prediction**: $P(Y_{t+L} \mid x_t = a)$

  Guess what might be true in a moment, e.g.

  t1: I-start-my-car
  ⇒ t2: I-drive-down-the-highway

- **Elaboration**: $P(Y_t \mid x_t = a)$

  Guess what else might be true now, e.g.

  t1: I-am-at-the-zoo
  ⇒ t1: A-monkey-is-here

- **Explanation**: $P(Y_{t-L} \mid x_t = a)$

  Guess at what happened prior to the current event, e.g.

  t2: I-am-at-a-movie-theater
  t2: I-am-watching-a-movie
  ⇒ t1: I-buy-a-movie-ticket

- **Projection**: $P(Y_{t:t+L} \mid x_t = a)$

  Guess what series of events might follow, e.g.

  t1: I-wake-up
  ⇒ t2: I-brush-my-teeth
  t3: I-take-shower
  t4: I-get-dressed
  t5: I-eat-breakfast

- **Filtering**: $P(Y_t \mid x_0 = a_0, ..., x_t = a_t)$

  Filter unlikely current states or events, e.g.

  t1: I-am-skateboarding
  t1: I-am-at-school
  t1: I-am-100-years-old
  ⇒ t1: ¬I-am-100-years-old

- **Fixed-lag smoothing**: $P(Y_{t-L} \mid x_0 = a_0, ..., x_t = a_t)$

  Filter unlikely past states or events, e.g.

  t1: I-am-outside
  t1: It-is-raining
  t1: ¬I-have-umbrella
  t2: I-am-dry
  ⇒ t1: ¬It-is-raining

This list was partly drawn from Murphy's review [3] of the kinds of reasoning that could be enabled with Dynamic Bayes Nets, a type of Bayes net with a temporal structure similar to that of LifeNet.

REASONING WITH LIFENET

We had to confront two major issues in finding a way to reason with LifeNet:

Noise in the data. LifeNet was derived from the largely unedited contributions of 12,000 people via several stages of processing, and so the knowledge contains problems of various types: links may go in the wrong direction, links may be inconsistent, the same proposition may be expressed in multiple ways, and so forth. Is there an inference procedure that is tolerant to these kinds of errors?

Scale of the data. The scale of LifeNet is a difficult challenge for even the most modern inference methods. Currently, the LifeNet graph contains about a ½ million links. The problem can only get worse, as this network is meant only a starting point. An adequate common sense model in terms of a propositional representation will likely contain at least tens of millions and possibly billions of nodes and links. Is there an inference procedure that can run quickly over such a network?

Treating LifeNet as a graphical model

To deal with these two issues, we built LifeNet not as a set of propositional rules, but rather as a probabilistic graphical model. Statistical methods for reasoning are somewhat more tolerant than ordinary logical reasoning methods to the uncertainty in our knowledge of the situation, as well as to the uncertainty in the reliability of the rules themselves. Also there are simple and well-known inference procedures that can run fairly fast if you are willing to accept approximate solutions.

Our early experiments reasoning with LifeNet treat it as *Markov network*, an undirected graphical model where the nodes represent random variables and the edges joint probability constraints relating those variables. We convert the rules of LifeNet into joint probabilities (the details of this process are described later this paper), and we reason with the resulting network using local belief updating techniques. We engage in 'loopy' belief propagation as described by Pearl [4]. Belief propagation in a Markov network is straightforward. We use the following belief updating rules, as described in [5]:

$$m_{ij}(x_i) \leftarrow \alpha \sum_{x_i} \psi_{ij}(x_i, x_j) \psi_i(x_i) \prod_{k \in N(i) \setminus j} m_{ki}(x_i) \quad (1)$$

$$b_i(x_i) \leftarrow \alpha \psi_i(x_i) \prod_{k \in N(i)} m_{ki}(x_i) \quad (2)$$

In these rules x_i represents the random variable at node i. The current belief in node i is denoted by b_i, the local evidence for node i by ψ_i and the joint probability of a pair of linked nodes i and j by ψ_{ij}. The message sent from node i to j is denoted by m_{ij}. $N(i)$ is the set of all neighbors of node i, and $N(i) \setminus j$ represents the set of all neighbors of node i except for node j. α is a normalization constant.

These simple updating rules run fairly quickly even on a knowledge base as large as LifeNet. In our optimized Common Lisp implementation, on a 1.8 GHz Pentium 4 with 512 MB ram, a single iteration of belief updating runs in 15 seconds. Inference is further sped up by restricting the belief updating to run only on nodes within a fixed distance from the evidence nodes. Given a single evidence node and using only those nodes within a depth of three edges away, a single iteration of belief updating runs in as little as 0.5 seconds for some nodes; on average it takes about 3 seconds. This is of course an approximate inference, but later we present an evaluation that suggests it may be adequate.

Problems with our model

Why do we not use a directed model? The main problem is that we are uncertain about the direction of the links—often, the causality implied by a LifeNet link really flows in the opposite direction. This is a problem both with the original assertions from which LifeNet was derived (raising the interesting question of whether the general public can be trusted to supply good 'causal theories'), and also bugs in the way we generated LifeNet from those assertions. There is also the problem that many of the nodes of LifeNet have too many parents to efficiently encode the conditional probability tables that relate child nodes to their parents. In the present network some nodes have dozens or even hundreds of parents.

A disadvantage of using an undirected over a directed model is that we lose the phenomenon of 'explaining away', where if there is evidence for a single cause, the probability of other causes is reduced, a heuristic that is surely used in human commonsense reasoning. One possibility is to 'marry' the parents of nodes so that they reduce each other's probabilities, but when there are many parents this requires adding a great many new edges to the graph. Another possible solution is to use a compact formulation for conditional probability tables such as the noisy-or function; Pearl [4] describes an analytic formulation of his belief updating rules for Bayes nets with noisy-or CPTs that avoids summing over the exhaustive combination of parent states.

GENERATING LIFENET

LifeNet is generated from the Open Mind Common Sense (OMCS) corpus, a collection of about 600,000 'common sense' assertions supplied by the general public [6,7] via the Open Mind Common Sense web site[1]. The contributed knowledge is expressed in natural language, and consists largely of the kinds of simple assertions shown in Table 1.

Table 1. Sample of OMCS corpus

A person wants to eat when hungry.
Things often found together: light bulb, contact, glass.
Coffee helps wake you up.
The effect of going for a swim is getting wet.
The first thing you do when you wake up is open your eyes.

[1] http://openmind.media.mit.edu

Starting with this knowledge, we generate LifeNet in several stages of processing.

Generating the propositions

The first thing we do is generate a list of egocentric propositions that will serve as the nodes of LifeNet. We do this in two ways: (a) by using propositions that people entered directly in the OMCS web site, and (b) by generating propositional sentence forms by filling in sentence patterns with phrases from the OMCS corpus.

The First-Person Situation activity

We make use of the propositions that were entered directly by contributors to the original OMCS web site. The site's *First-Person Situation* activity explicitly collected first-person propositions of the form:

- I am having trouble keeping awake
- The sky is getting dark
- It is the early morning

We collected about 14,000 such propositions from OMCS contributors.

Generating propositions from OMCSNet

In addition to these explicitly collected propositions, we derive further first-person propositions from the OMCSNet knowledge base, a collection of about 280,000 binary relations of the form:

- UsedFor("shampoo", "washing hair")
- LocatedAt("bottle of shampoo", "in bathroom")
- EffectOf("drinking coffee", "feeling awake")
- SubEventOf("playing baseball", "swinging bat")

The OMCSNet database was built by spotting patterns in the raw assertions collected by the OMCS web site. Many of the assertions in the original OMCS corpus were supplied by users filling in the blanks of templates of the kind *One of the things you do when you [play baseball] is [swing a bat]*. It is a simple matter to then extract binary relations from the assertions gathered via such activities. More details about how the OMCSNet knowledge base was built are provided in [8].

To generate LifeNet propositions, we group the arguments of the OMCSNet relations into *actions, objects, locations,* and *states*. For example, for the relation *UsedFor(X,Y)*, the argument *X* is added to the collection of *objects*, and the argument Y is added to the collection of *actions*. We then generate propositions by filling in the following sentence templates:

1. I am [PLACE]
2. I am [STATE]
3. I [ACTION]
4. A/An [OBJECT] is here

We run the resulting sentences through a set of filters to repairs the tense of verbs and fix determiners. We also filter out those sentences that use words that are not in a custom list of 5,000 common English words. The final results are sentences of the form:

1. I am in a bathroom
2. I am thirsty
3. I drink coffee
4. A bottle of shampoo is here

These propositions still have many types of problems, not least of which is that many of the sentences are not very likely or impossible, e.g. "I am inside an oven". Thus we have done some scanning of the resulting list of propositions by eye and we deleted many of the most egregious such errors, but given the size of the list we can only repair some such problems. Ultimately, repairing and extending this list of propositions is a task we wish to give to the participants in the web-based distributed project to grow LifeNet, discussed later in this paper.

Generating the propositional rules

The next step is to generate a set of propositional rules. Given our initial list of 40,000 propositions, we can go back to OMCSNet to generate a large collection of propositional rules. We do this in a straightforward way, by iterating over a subset of the OMCSNet knowledge base, and from each useful assertion applying one of the rule generation rules shown in Table 2. We use ⇒ to refer to a temporal rule, and → to refer to an atemporal rule, e.g. *I am yelling → I feel angry,* vs. *I drink coffee ⇒ I feel alert.*

Table 2. Rules for generating LifeNet from OMCSNet

EffectOf(e1, e2):
I [*BestMatch(e1, actions)*] ⇒ I [*BestMatch(e2, actions)*]

SubeventOf(e1, e2):
I [*BestMatch(e1)*] ⇒ I [*BestMatch(e2)*]

LocatedAt(object, place):
I am [*BestMatch(place, places)*] → A/an [*object*] is here

OftenNear(object1, object2):
A/An [*object1*] is here → A/An [*object2*] is here

FirstSubeventOf(e1, e2):
I [*BestMatch(e1, actions)*] ⇒ I [*BestMatch(e2, actions)*]

LastSubeventOf(e1, e2):
I [*BestMatch(e1, actions)*] ⇒ I [*BestMatch(e2,actions)*]

UsedFor(object, event):
A/An [*object*] is here → I [*BestMatch(event, actions)*]

Requires(event, object):
A/An [*object*] is here → I [*BestMatch(event, actions)*]

BestMatch is a simple matching procedure that finds the proposition that best matches the corresponding argument of the OMCSNet relation by correlating the non-stopwords between the arguments. We presently don't use synonyms or taxonomic information to compute more accurate semantic distances between propositions and the arguments of the OMCSNet relations.

Generating propositional rule plausibilities

The third step is to assign to each resulting rule a probability, or rather a 'plausibility', a number that represents a belief in the accuracy of the statement. Because of errors in the steps involved in producing the propositional rules, some of the final rules will not make sense. To these we wish to assign a low plausibility. We presently compute such plausibilities using a simple method. We first compute a correlation between pairs of words in the Open Mind Common Sense corpus, by counting for every word pair how many assertions in which they co-occur. We then take the log of this number and scale it so that the resulting correlation lies between zero and one. Finally, for each propositional rule, we assign it a plausibility by computing the average of the correlation between every non-stopword in the antecedent proposition and the consequent proposition.

Generating a temporal graphical model

Finally, we convert these propositional rules with their plausibilities into a large-scale temporal graphical model as shown in Figure 1. The previous three steps produced 130,000 simple propositional rules, each with a single antecedent and consequent clause. From this set we generate a total of 415,000 links because (a) the atemporal rules are repeated in each column of LifeNet, (b) for each link we add 'reversed' links for temporal rules with a reduced joint probability (not shown in Figure 1) because of the uncertainty in the direction of the links, and (c) we add propositional 'persistence' links (not shown in Figure 1) that relate each 'before' node to its corresponding 'after' node with a joint probability of 0.5. Obviously, not all states persist or persist with a probability 0.5, but this seemed like a reasonable initial estimate.

THE STRUCTURE OF LIFENET

The result of all of this processing is a very large graph. Some of its statistics are given in Table 3.

Table 3. LifeNet graph statistics

Total nodes	78332
Total edges	415248
Average edges per node	10.6
Average fan-in per node	5.3
Average fan-out per node	5.3
Nodes with > 50 parents	1109
Nodes with > 10 parents	4730
Nodes with > 5 parents	12454
Nodes with > 50 children	1108
Nodes with > 10 children	5294
Nodes with > 5 children	12592
Maximum fan-in per node	632
Maximum fan-out per node	589
Maximum edges per node	1219

In other words, most nodes are connected to only a small number of other nodes, but about 10% of the nodes are 'hubs' with a substantial number of connections. This is because the original OMCS web site was 'seeded' with about a thousand initial concepts around which grew a great many relations.

LIFENET QUALITY

How can one evaluate a commonsense knowledge base and inference system? Presently, there is no 'gold standard' against which a common sense inference system can be compared, other than the judgement of people themselves. Thus we asked a human judge to assess both samples of knowledge from the system, and samples of the inferences that it made. The judge evaluated the knowledge base according to several criteria:

- Do individual propositions make sense?
- Do individual rules make sense? In other words, for atemporal rules, when X holds, does Y often hold? For temporal rules, when X holds, does Y often follow?
- Do individual inferences work? In other words, does Y often follow from X?

The results of this evaluation are shown in Table 4.

Table 4. Quality of LifeNet knowledge and inferences

(1) Propositions that make sense	89%
(2) Reasonable temporal links	76%
(3) Reasonable atemporal links	61%
(4) Reasonable inferences	78%

In (1) a judge examined 500 random propositions from LifeNet. In (2) a judge examined 500 temporal links. In (3) the judge examined 500 atemporal links. In (4) the judge queried the inference engine to make 50 separate predictions. To do this, 50 'good' propositions (ones that made sense) were first chosen at random. Then, for each of these propositions, the network was initialized to set that proposition to be true at time t=1, and the inference engine run for 5 iterations of belief updating, and the new beliefs at time t=2 were measured. The judge examined the 5 most probable resulting beliefs at time t=2 and measured how many of these were plausible. We then averaged these to estimate the precision of inferencing.

We believe these results are not unreasonable for a knowledge base that was built by a large community of contributors. However, the current LifeNet knowledge base is only a initial starting point. In the next section we describe a new web site that builds on the LifeNet knowledge base and reasoning methods to further debug and grow the knowledge base.

OPEN MIND LIFENET

We see the LifeNet we have built as only the starting point for a substantially better common sense knowledge base. Propositional rules with single antecedents and consequents are a very coarse model, and much of the knowledge in LifeNet is buggy. Thus, as we did with the original Open Mind Common Sense web site, we are turning to the general public to build a better model, by allowing them to add new propositions and links as well as repair the existing knowledge. The new web site, to be called Open Mind LifeNet, will let people enter new propositions, modify existing rules, and add new rules that include multiple antecedents and consequents, e.g.

t1: I-am-thirsty
t1: I-have-a-full-glass-of-water
t1: I-drink-a-glass-of-water
→ t2: ¬I-am-thirsty
t2: ¬I-have-a-full-glass-of-water

Open Mind LifeNet Activities

The Open Mind LifeNet web site is presently still under development, but we have built a prototype non-web-based interface to the knowledge and inference engine to explore ideas about the kinds of knowledge elicitation activities we might need. At present there are just three types of activities: (a) acquiring knowledge about new propositions, (b) acquiring knowledge about atemporal links, and (c) acquiring knowledge about temporal links.

Acquiring new propositions

While tens of thousands of propositions may seem like a great many, this represents only the tiniest a fraction of what one might say about a given situation. For example, just listing the number of possible places one is at, e.g., "I am in St. Hubert, Quebec", "I am in San Francisco, California", etc. would require thousands of propositions in itself. Thus we allow people to add new propositions if they wish to extend the range of the system. A sketch of this activity is shown in Figure 2.

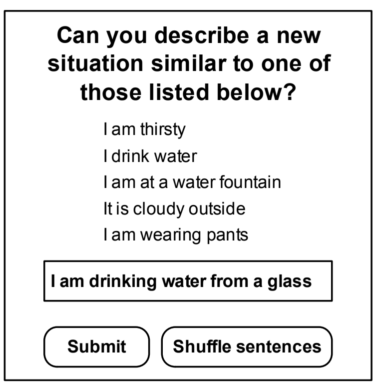

Figure 2. Interface for acquiring new propositions

Presently we simply as the user to enter a new statement, although in the future we plan to generate the propositions by using a simple grammar for egocentric propositions in combination with an ontology of objects, events, states, and other entities.

Acquire new atemporal links

This activity, shown in Figure 3, allows the user to add atemporal links by selecting groups of propositions that are likely to be true or not true at the same time. The major challenge in building such an interface is in making it easy for people to retrieve relevant propositions to select. We take the approach of using the LifeNet inference engine itself to generate plausible propositions, by asking it to elaborate the situation by guessing what else might be true at the same time as the selected propositions. Because the existing inferences are only partly reliable, many of the generated propositions are reasonable new candidates.

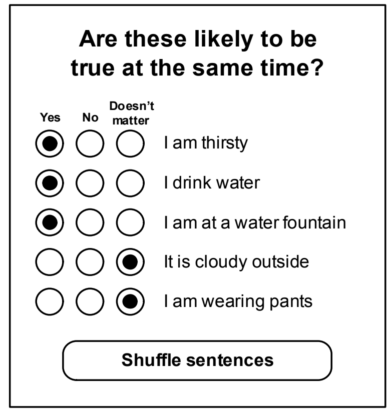

Figure 3. Interface for acquiring atemporal links

Acquire new temporal links

This activity, shown in Figure 4, allows the user to add temporal links. As in the atemporal activity, the candidate propositions are generated by using the inference engine.

FUTURE WORK

There are many possible directions to go in the future.

Refine the probabilistic model. The current model is a rather simple undirected network, which does not efficiently support explaining away or more complex joint or conditional probability tables. In the next version of LifeNet we are considering using dynamic Bayes nets with a noisy-or combination rule for parents. However, to make use of the more complex conditional probability tables that Open Mind LifeNet will let us build from rules with multiple antecedents, it is likely we will need to go beyond such simple tricks and develop new ways make inferences in large-scale graphical models.

Group the LifeNet propositions. One way to improve the quality of LifeNet inference is to organize the network structure so that propositions that are highly interdependent are grouped into subnetworks in which exact inference can be performed. For example, we are considering grouping the LifeNet nodes along 'mental realms' like 'location' (I-am-the-zoo, I-am-at-home), 'weather' (It-is-raining, The-sky-is-clear), 'emotion' (I-feel-scared, I-feel-angry), and so forth. This is the beginning of a more 'agent-based' view of LifeNet where the subnetworks are specialists in subdomains of common sense.

Grow LifeNet by learning from people's experiences. We have developed an interface that lets a wearable computer listen in on human conversations as well as observe other types of context like location as people are engaging in their day-to-day activities [9]. We hope this will let us collect a large amount of data from which we can learn a common sense model of human activity. We have also built the Open Mind Experiences web site, which lets peo-

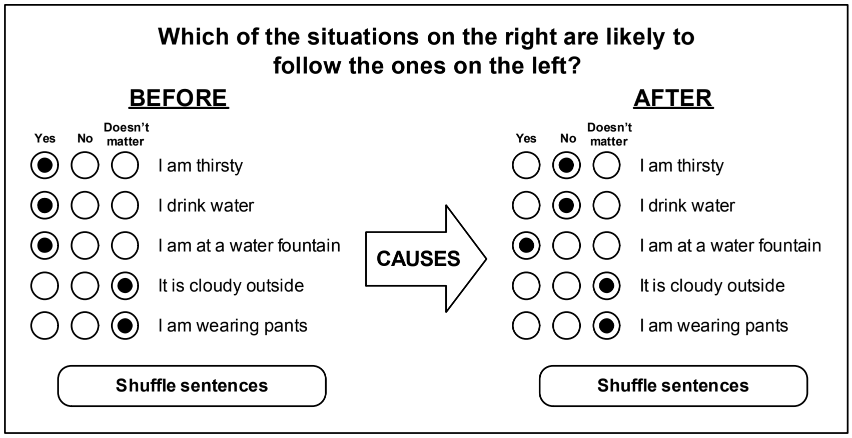

Figure 4. Interface for acquiring temporal rules

ple more directly supply descriptions of common sense experiences [10].

Build a social model by coupling several LifeNets. The present 'egocentric' LifeNet captures common sense from the point of view of an individual. We would like to extend LifeNet so that it could capture common sense about multiple interacting people. We are considering 'coupling' several LifeNets to model the interaction of several people, e.g. if one person is a driver and in a car, then if the other person is in a car they are likely to be a passenger.

Turn to more relational representations. While we have argued in this paper for the value of propositional representations, there has been much progress recently in combining the advantages of statistical and relational representations, for example, Probabilistic Relational Models [11] or their temporal extension Dynamic Probabilistic Relational Models [12]. We may be able to return to a relational representation while retaining the ability to deal with uncertain domains.

CONCLUSIONS

LifeNet is the first large-scale common sense reasoning system that represents knowledge as a graphical model and that draws inferences by probabilistic belief updating. While the current database is fairly noisy, we plan to further debug, develop, and grow LifeNet over time. We hope that this work will suggest to those developing statistical inference techniques that their methods may be applicable to the challenging domain of common sense reasoning.

ACKNOWLEDGMENTS

We would like to thank Barbara Barry, Nathan Eagle and Sandy Pentland for their input on this project.

REFERENCES

[1] E. Davis (1990). *Representations of commonsense knowledge*. San Mateo, Calif.: Morgan Kaufmann.

[2] The Upper Cyc Ontology. Available at http://www.cyc.com

[3] K. Murphy (2002). *Dynamic Bayesian Networks: Representation, Inference and Learning* (PhD Thesis). Department of Computer Science, University of California, Berkeley.

[4] J. Pearl (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Mateo, CA: Morgan Kaufman.

[5] J. S. Yedidia, W. T. Freeman, and Y. Weiss (2000). Generalized belief propagation. *Advances in Neural Information Processing Systems (NIPS)*, 13, 689-695.

[6] P. Singh (2002). The public acquisition of commonsense knowledge. In *Proceedings of AAAI Spring Symposium: Acquiring (and Using) Linguistic (and World) Knowledge for Information Access*. Palo Alto, CA: AAAI.

[7] P. Singh, T. Lin, E. T. Mueller, G. Lim, T. Perkins, and W. L. Zhu (2002). Open Mind Common Sense: Knowledge acquisition from the general public. *Proceedings of the First International Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems*. *Lecture Notes in Computer Science* (Volume 2519). Heidelberg: Springer-Verlag.

[8] H. Liu, and P. Singh (2003). OMCSNet: A commonsense inference toolkit. Retrieved from http://web.media.mit.edu/~push/OMCSNet.pdf

[9] N. Eagle, P. Singh, and A. Pentland (2003) Common sense conversations: understanding casual conversation using a common sense database. In *Proceedings of the Artificial Intelligence, Information Access, and Mobile Computing Workshop (IJCAI 2003)*. Acapulco, Mexico.

[10] P. Singh and B. Barry (2003). Collecting commonsense experiences. In *Proceedings of the Second International Conference on Knowledge Capture*. Florida, USA.

[11] N. Friedman, L. Getoor, D. Koller, and A. Pfeffer (1999). Learning probabilistic relational models. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*. Stockholm, Sweden.

[12] S. Sanghai, P. Domingos, and D. Weld (2003). Dynamic probabilistic relational models. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*. Acapulco, Mexico.