

A Framework for Commonsense Knowledge Retrieval

Preliminary Report

Phil Oertel

Computer Science Department
University of Illinois at Urbana-Champaign
Urbana, IL 61801
oertel@uiuc.edu

Eyal Amir

Computer Science Department
University of Illinois at Urbana-Champaign
Urbana, IL 61801
eyal@cs.uiuc.edu

Abstract

An autonomous agent that explores and acts in a rich world needs knowledge to act effectively. This agent can use knowledge that is available in *Commonsense Knowledge Bases* (CSKBs), when the agent designer cannot encode all the information the agent might need. CSKBs include general-purpose information about the everyday world in a formal language, but this information is not always correct, relevant, or useful for the agent's purpose.

In this paper we present an approach to retrieving commonsense knowledge for autonomous decision making. We consider agents whose formal language is different from that of the CSKB, and can use multiple CSKBs of various expressivity and coverage. We present a retrieval framework with algorithms for mapping languages and selection of knowledge. We report on preliminary experimental results of these algorithms for the ConceptNet CSKB.

1 Introduction

There is growing interest in using knowledge about the world to aid autonomous decision making, e.g., [Bacchus and Kabanza, 2000; Doherty and Kvarnström, 2001]. This knowledge is crafted carefully by a knowledge engineer to fit the domain and goal. However, autonomous agents that explore and acts in rich worlds cannot receive all their information from the agent designer because he does not know the target environment in design time or the environment is too complex. In these cases this agent can use knowledge that is available in *Commonsense Knowledge Bases* (CSKBs).

Typically, CSKBs comprise of ontologies and hierarchical concept information together with many thousands of axioms that bring these concepts together. Examples include CYC [Lenat, 1995], SUMO [Niles and Pease, 2001], the HPKB project at SRI [Cohen *et al.*, 1998], the HPKB project at Stanford's KSL [Fikes and Farquhar, 1999], ConceptNet [Singh *et al.*, 2002], and WordNet [Miller, 1995].

The challenges posed by decision making are not addressed by current CSKB research. Research on CSKBs has focused so far on applications for natural language processing

and for question answering. The first typically uses the concept hierarchy information that is embedded in the CSKB, while the second uses axioms of the CSKB that are tuned carefully for the specific topic and queries together with a theorem prover. In contrast, autonomous exploration and decision making require knowledge that is versatile and can aid in answering questions like “*what will action a do in my situation?*” or “*how do I get closer to X?*”. The two challenges posed by such applications are selecting knowledge that is correct, relevant, and useful for the current task, and then using this knowledge. We cannot expect fine tuning of axioms or manual selection for an agent that explores an unknown territory. Furthermore, our agent has to be able to select and use knowledge that is not complete or accurate.

In this paper, we present an approach to retrieving commonsense knowledge for autonomous decision making, and a retrieval framework with several algorithms for mapping languages and selection of knowledge. We consider agents whose formal language is different from that of the CSKB, and can use multiple CSKBs of various expressivity and coverage. Our algorithms translate an agent's knowledge base (AKB) to the language of a CSKB, even when the CSKB is sparse. We discuss briefly how an agent may use the AKB and knowledge from a CSKB to choose actions.

The usage of knowledge by our agent gives rise to two types of queries that we can ask from a CSKB. These are *region queries*, which find relevant concepts and axioms given a set of concepts that the agent considers *active*, and *path queries*, which find relevant concepts and axioms given the current situation and a goal description. The corresponding queries are processed using a graph that we associate with the CSKB, and using methods from information retrieval [Salton and Buckley, 1988] and automated reasoning. We report on preliminary experimental results of these algorithms for the ConceptNet CSKB.

This is work in progress, and some details are omitted with only little experimental evaluation.

2 Knowledge and Decision Making

Our running example is that of an agent playing an adventure game [Amir and Doyle, 2002]. This problem isolates the commonsense reasoning problem from external challenges like vision and motor control. Far from the arrows and breezy

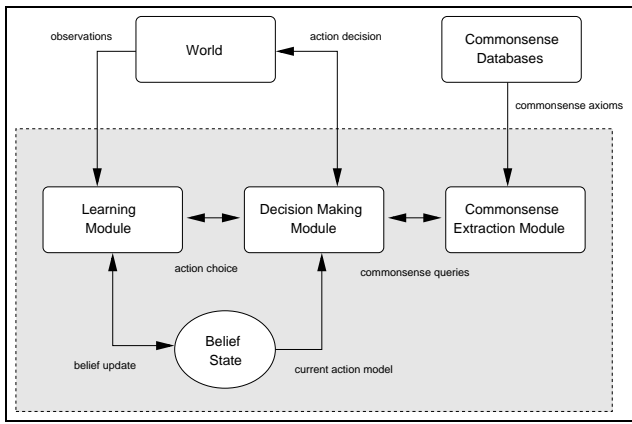


Figure 1: This is the adventure game architecture in which our framework operates. The current paper describes the commonsense extraction module. Thanks to Hlubocky and Amir for use of the diagram.

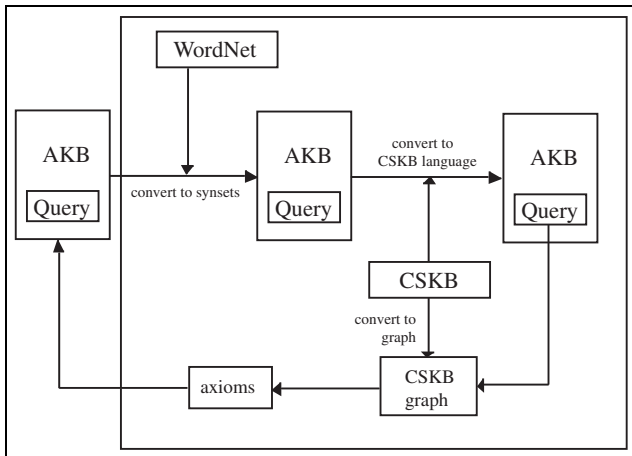


Figure 2: A high-level perspective of the retrieval framework

pits of the Wumpus world, these adventure games are intended to be abstractions of real-world scenarios. The agent is designed to track its world in a compact belief state, use commonsense to reason about the semantics of its knowledge, learn the effects of actions through observation, and use its knowledge to progress toward some goal. Importantly, the agent is required to learn the action model, which is unknown initially¹.

2.1 Problem Setting

We assume that there is a single agent in the partially observable world, with knowledge about the current state of the world (and its effects of actions) that is represented by a set of logical sentences. We also assume that the agent can observe its environment using its sensors and act upon it. The

¹This paper does not concern learning action models or maintaining a compact belief state; the focus is on knowledge retrieval only.

world conforms to the assumptions of the basic Situation Calculus [McCarthy and Hayes, 1969; Reiter, 2001]: the Markov property holds, actions have explicit preconditions and effects and are deterministic, and there is a set of unique names axioms. Simple adventure games already have these properties, so they are not too restrictive, and they allow us to represent actions using a well-understood language.

The agent is given a goal that is achievable from any state via a finite sequence of executable actions. (In many adventure games, a state satisfying the goal is not specified initially but must be *discovered* through exploration (e.g., by reading a note), but we ignore this case here.) Thus, in such environments the agent can perform some default exploratory behavior until a goal is reached. However, this exploration is infeasible because it may take time that is linear (or worse) in the number of states in our system, which is in $\Omega(2^n)$ for n propositional features in the domain.

Thus, our decision maker is faced with the choice of a action policy A such that executing A in the current state s produces the goal state g . Methods that address this problem to some degree include planning under partial observability [Bertoli and Pistore, 2004] and reinforcement learning in partially observable domains [Even-Dar *et al.*, 2004]. They maintain a belief state and compute a plan or a policy for a given problem domain and goal.

2.2 Using Knowledge

There are four ways in which we wish to use knowledge: (1) control knowledge for planning and search, (2) information about the way actions change the world (eliminate some of the uncertainty about the transition model), (3) better guidance to exploration on the way to the goal (e.g., using information gain), and (4) better guidance to acting on the way to the goal (using projection paths).

Control knowledge for planning (e.g., [Bacchus and Kambanza, 2000; Doherty and Kvarnström, 2001]) can be represented as a restriction strategy on the choices that the agent can make. For example, one can state that if our agent is faced by a locked door, then it should not try to unlock the door with a banana (this may jam the door).

Information about actions' effects and preconditions can help in making our planning problem more feasible, and should allow us to reach islands in our search space. For example, if we know that unlocking the door with the correct key will have the effect that the door is unlocked, then we will try this action before others, to find if the key that we have is the right one, and possibly get into the room on the other side of the door as a result.

Knowledge about the way our actions change (or do not change) the world, and knowledge about the way they are *likely* or *unlikely* to change the world is important for exploration. Information gain (see [Mitchell, 1997]) is a measure that is typically used to perform this exploration. It is particularly attractive in adventure and exploration games because there one can use the information gain measure with hill climbing to reach the goal (the goal is achieved exactly when we discover an action that achieves it).

Finally, a chain of plausible, high-level concept landmarks or actions (e.g., a chain of concepts connecting a house with

a grocery store) can guide both planning and reinforcement learning. Planning techniques can use such chains to guide a rough forward plan search with a subsequent process that corrects the plan where needed (e.g., [Hoffmann and Nebel, 2001]). Reinforcement learning can use such information to re-shape the reward structure by transferring some of the final reward (of reaching the goal) to various positions in the state space [Ng *et al.*, 1999].

2.3 Representing the Agent’s Knowledge and Interaction Model

In this paper we assume a logical representation for actions (e.g., in the situation calculus), having met the preconditions for doing so with the previous section’s restrictions. As mentioned above, the action model is unknown initially. However, we assume that the game environment provides the agent with a list of relevant actions.

We use the intuition that action $a(\vec{x})$ is relevant if it is executable (with a non-nil effect) in some state s' that is reachable from the initial state. Such a list of relevant actions is usually available to a human agent in an adventure game, either through a help system or through feedback that notifies the player of an unrecognized action. The rest of the action model is initially hidden from the agent.

Ideally, an agent should be able to receive all observations in natural language, as humans do, but converting natural language sentences to logical ones is tantamount to a solution to the general knowledge extraction problem. Thus, we assume that observations contain only a limited natural language component, and we associate this natural language component with a set of logical sentences that represent the observations of the agent. We use the text that is associated with nonlogical symbols in our First-Order Logical language L as semantic information that allows us to connect L with the language of any CSKB.

3 Framework Overview

Figure 2 shows how a commonsense retrieval system augments an agent’s knowledge base with relevant commonsense axioms. The upper and lower halves of the diagram represent the two broad tasks being performed. The first task is to create a mapping from symbols in the AKB to those in the CSKB. The second task is to use a subset of the AKB (now mapped to the language of the CSKB) as a query with which to find useful axioms in the CSKB. To determine relevance as generally as possible given the unstructured nature of some CSKBs, relevance is only calculated based on distance in a graph representation of the CSKB. Every CSKB has ground terms and axioms over those terms, so a simple graph representation has a node for each term and an edge between two nodes if and only if there is an axiom that mentions both of those terms. The astute reader will note that the algorithms as presented retrieve concepts rather than axioms. These concepts are intended to be used along with the supplied query concepts as endpoints of a path along which all axioms are retrieved.

3.1 Mapping Agent Language to CSKB Language

To reason with commonsense knowledge, a mapping must be made from the symbols in the agent’s language to those in the

CSKB. The most important thing to realize when constructing the mapping is that it won’t be exact, because no formal ontology exists that can represent all the information in English words (even when these words occur in isolation). Worse, it is not clear that such a formalization is even possible. With that limitation in mind, we propose the baseline system seen in Figure 3.

PROCEDURE MapSymbols($AKB, WordNetKB, CSKB$)
 AKB Agent’s KB (in FOL), $WordNetKB$ WordNet KB (Ontology and English words), $CSKB$ Commonsense KB (in FOL)

1. Set $Mapping \leftarrow \emptyset$
2. Loop for every constant symbol s_a in AKB :
 - (a) Search for s_a in WordNet, putting the resulting set of synsets^a in $senses$
 - (b) If $senses = \emptyset$, then set $Mapping \leftarrow Mapping \cup \{(s_a, nil)\}$. Else:
 - i. Set $syn \leftarrow argmax_{s \in senses} P(s|AKB)$
 - ii. Set s_c to be a CSKB concept (predicate, function, or constant) symbol corresponding to syn ’s WordNet ID.
 - iii. If s_c is non-nil, then set $Mapping \leftarrow Mapping \cup \{(s_a, s_c)\}$. Otherwise, loop until $syn = TOP$:
 - A. Set $syn \leftarrow hypernym(syn)$, syn ’s nearest hypernym^b in WordNet
 - B. If $CSKB$ has a concept s_c corresponding to syn ’s WordNet ID, then set $Mapping \leftarrow Mapping \cup \{(s_a, s_c)\}$, and break
3. Return $Mapping$

^aA *synset* is a basic unit in WordNet - the set of all synonyms of a meaning.
^bA *hypernym* is a parent concept in the WordNet ontology.

Figure 3: Algorithm for mapping agent’s symbols to CSKB symbols

MapSymbols maps every constant symbol in the AKB to a matching entry in WordNet, based on a textual match only. It performs word sense disambiguation based only on bag-of-words co-occurrence data, which could be obtained from a WordNet sense-tagged corpus. Then it attempts to find a matching concept in the CSKB, requiring an existing mapping from CSKB symbols to WordNet synsets. If a match cannot be found in the CSKB, the next nearest hypernym in WordNet is checked for a match, and so on until the most specific match available is found.

Consider the following example, a portion of a belief state appearing in [Hlubocky and Amir, 2004]:

(AND (AT ME TREE-BRANCHES) (AT TREE_TRUNK TREE-BRANCHES))

“ME” refers to the agent itself, and this statement represents the belief that the agent is at location “TREE-BRANCHES”. There is an object called “TREE_TRUNK” in the same location. A human player would understand that the agent is

in a tree among the branches, and that there is a tree trunk there. Assuming dashes and underscores can be parsed easily and translated to spaces, the agent has to try to translate “TREE TRUNK” to a symbol in a CSKB. To understand why our WordNet solution is needed and how it helps, we’ll try to translate that symbol into Cyc symbols manually.

We look for “tree trunk” but find no matching symbol in any publicly available version of Cyc. The next term up the WordNet hierarchy is “stem”. This term matches a symbol in ResearchCyc. If we were restricted to OpenCyc, we could go up one level higher to “plant organ”, which matches a term there. This method works particularly well for restricted versions of Cyc, in which there is a sparse knowledge base with a full upper ontology. But it’s not clear that the technique will still work well for CSKBs without such an ontology, such as ConceptNet.

For symbols other than constants, there may be more dimensions to match than simple textual similarity. In Cyc, for instance, functions and predicates take typed arguments, and functions themselves have types. Types and arity represent syntactic restrictions that must be checked before an agent symbol can be mapped to a CSKB symbol. To map functions and predicates using MapSymbols, we can perform unification, recursively matching terms textually or - if the terms are functions - through unification.

3.2 Retrieval Tasks

We want our knowledge retrieval system to work well on each of the CSKBs, but their differences make this goal difficult to achieve. To minimize the effects of these differences, we create a simplified but uniform representation for CSKBs.

The retrieval system is not an algorithm for deciding which axioms are relevant but a system for facilitating the application of such algorithms on different types of CSKBs. Our aim is to be able to retrieve useful information from any of them, but their differences make that a difficult task. To minimize the effects of these differences, we simplify the CSKBs, converting them to a weighted graph. The simple procedure is given in Figure 4. The following sections describe *region queries* and *path queries*, the two retrieval options allowed in the framework.

Region Query

This type of search is intended to help the retriever find out more about its current knowledge. The intent is for an agent to select a small number of concepts from its KB, and the result should be the set of axioms most relevant to them. These axioms are found using spreading activation, a technique developed in the cognitive science community. Spreading activation was originally developed to retrieve nodes from a semantic graph in a way believed similar to the human memory’s own retrieval mechanisms. As such, it is not a well-founded technique but rather a well-studied standard for semantic information retrieval [Pirulli *et al.*, 1996].

Conceptually, spreading activation starts with a set of nodes with some *activation weight* and proceeds to activate neighboring nodes recursively over a series of time steps. The activation weight $a_i^{(t)}$ of a node i at time t is $f(\sum_j w_{ij} a_j^{(t-1)})$, where j varies over the neighbors of i ,

PROCEDURE RetrieveFromGraph($CSKB, S, T, q$)
 $CSKB$, Commonsense KB (in FOL); $S \subseteq CSKB$, a relevant subset of the agent’s KB, mapped to $CSKB$ concepts; $T \subseteq CSKB$, a set of concepts appearing in the goal state, mapped to $CSKB$ concepts; $q \in \{region, path\}$, the type of query

1. Set $A \leftarrow \emptyset$, the set of retrieved axioms.
2. Construct a weighted graph, G , from $CSKB$, with a node for every concept and an edge between two nodes iff the corresponding concepts are mentioned in any axioms together. The weight on each edge corresponds to the number of axioms it represents.
3. Set $S' \leftarrow$ the nodes in G corresponding to S
4. Set $T' \leftarrow$ the nodes in G corresponding to T
5. if $query_type = region$ then set $A \leftarrow RegionQuery(G, S')$
6. else $A \leftarrow PathQuery(G, S', T')$
7. Return A

Figure 4: Algorithm for retrieving relevant commonsense axioms

and w_{ij} is the edge weight from j to i . $f()$ is usually a decay function that has the effect of decreasing activation as distance increases from the activation sources. Activation weights are ranked after a halting condition is reached, based either on time or nearness to some asymptotic distribution. Spreading activation is only a general procedure, but the “leaky capacitor” model used in [Pirulli *et al.*, 1996] has been analyzed parametrically in [Huberman and Hogg, 1987] and can be used on our network with little modification. $RegionQuery(G, S)$, then, is a straightforward application of the leaky capacitor model, parameterized for a rich, connected network.

Path Query

This type of search is intended to find paths between two concepts (or sets of concepts). They might represent two regions of the belief state that the agent wants to connect (e.g., a key and a locked door), or one of them might represent the goal or part of the goal. A goal-directed search returns axioms lying on the shortest paths between the supplied concept sets because it is assumed that the retrieving agent is trying to find the shortest path to its goal. The algorithm is given in Figure 5.

After the shortest paths are found from the source to the destination, each node is given a score based on the number of those paths on which it lies. Allowing for the possibility that there is significant variance in the length of the shortest paths, preference is given to nodes along shorter paths. The source and destination can be given as sets of nodes, and the sets will be collapsed into single source and destination nodes.

Figure 6 gives the top ten results returned by the path heuristic on ResearchCyc between the terms “MovieTheaterSpace” and “ActorActress”. These results are clearly not what we want. Restricting the result set to actions doesn’t

help, because no actions even receive any score. Clearly by ignoring the semantics of the KB we lose too much information and get useless results. It seems likely that this heuristic will work better in ConceptNet since it lacks an upper ontology, but we have not yet applied the heuristic to that KB.

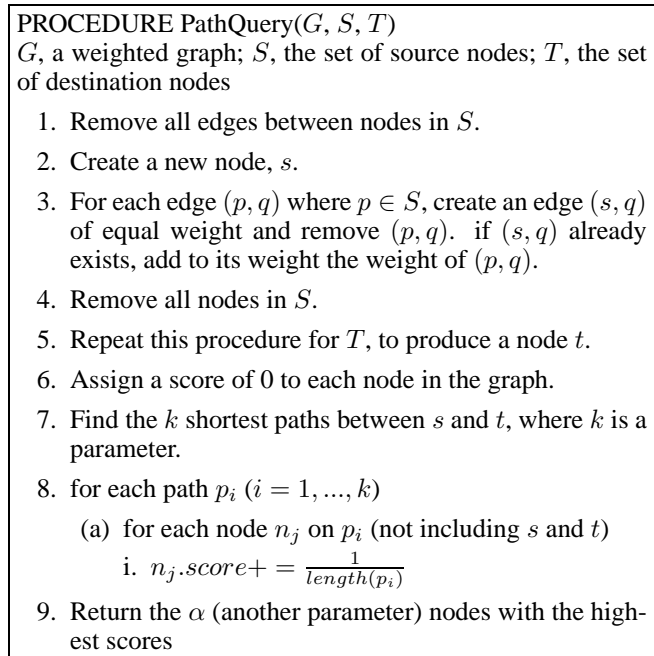


Figure 5: Algorithm for returning nodes on paths between two regions

Rank	Concept Returned	Score (sum of $\frac{path_i}{length_i}$)
1	mtVisible	263.5
2	Gens	125.3
3	Individual	94.64
4	ResearchCycConstant- NotFullyReviewed	57.4
5	termOfUnit	57.4
6	CycSecureFORT	52
7	definingMt	37.6
8	posForms	32.2
9	LexicalWord	4.54
10	arg1Isa	4.53

Figure 6: Top ten results for path query between terms “MovieTheaterSpace” and “ActorActress” in Cyc

4 Region Query with ConceptNet

We have implemented a limited, simplified version of the retrieval system for ConceptNet. This version approximates spreading activation in *RegionQuery* with ConceptNet’s own *GetContext* function. This function measures two nodes’ relatedness as the number of paths between them, weighted to favor shorter paths. As mentioned previously,

all axioms on the paths between query concepts and retrieved concepts are returned.

We did not use the retrieved axioms for reasoning yet, but we can already make a few observations about the results of these first retrieval experiments. First, only the top few (e.g., 5) retrieved concepts appear related enough to help an agent select a correct next action. For example, querying the terms “door” and “key” (manually selected and translated from an agent KB) returns the top ten concepts found in Figure 7. Note that the top result, “metal”, is very unlikely to be important in unlocking a door, but the terms “unlock door”, “open lock”, and “open door” may actually suggest the correct next action for the agent.

We do not yet have a good method for selecting which axioms to return based on the retrieved concepts. Figure 8 shows how many axioms are returned for the query in figure 7 according to one simple procedure. This procedure identifies all axioms lying on paths among query concepts, among retrieved concepts, and between query and retrieved concepts. As the figure shows, even if only the top result is used to select axioms by this method, over two hundred axioms are returned. What’s worse, our top concept, “metal”, is irrelevant. We would have found much more relevance in the top five concepts, but this would have retrieved around seven hundred axioms. Furthermore, the procedure defines no ranking over axioms. The number of axioms can be reduced by an automatic pruning of axioms that cannot be used in inference (e.g., identifying and removing actions that are not part of our action model). We plan to perform such refinements when we have a set of test cases over which we can try to reason with the retrieved axioms. This is our current work.

<ol style="list-style-type: none"> 1. metal 2. handle 3. unlock door 4. open lock 5. open door 6. keyboard 7. lock 8. doorway 9. room access 10. knob

Figure 7: Top ten results for region query on terms “door” and “key” using *GetContext* ranking function

5 Related Work

Background knowledge is increasingly important for decision making. Work on reinforcement learning [Kaelbling *et al.*, 1996; Andre and Russell, 2000; Ng and Russell, 2000; Ng *et al.*, 1999] uses background knowledge to structure the state space, update the reward function, and approximate the value function. Also, work on planning uses background knowledge to guide the search for a plan [Levesque *et*

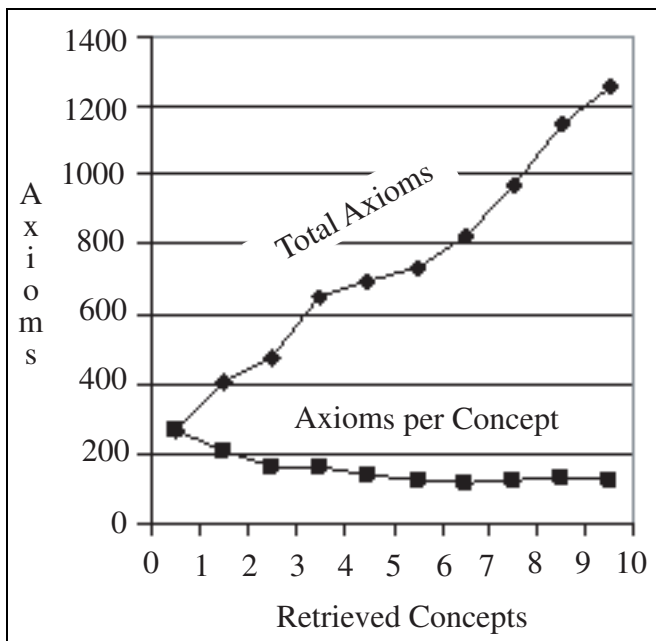


Figure 8: Number of axioms retrieved versus number of concepts kept from the top ten

al., 1997; Bacchus and Kabanza, 2000; Doherty and Kvarnström, 2001]. Finally, there has been some work on discovering knowledge to aid planning, and also about using non-monotonic reasoning to speed up reasoning [Ginsberg, 1991]. However, no work known to us has approached the problem of retrieving the right knowledge for a specific task.

The topic of matching symbols between KBs has attracted much attention in recent years. There have been some successes in the case of matching database schemas that have common underlying instances and language [Doan, 2002; Doan *et al.*, 2003], and some investigation was made in the case of more general AI knowledge bases. However, work on merging KBs and matching ontologies between KBs remains manually driven [Noy *et al.*, 2001].

6 Discussion

Ideally, we would like a CSKB to contain every action precondition and effect relevant to our agent’s world, along with any axioms necessary for inference. Of course, if such a complete and correct KB existed in a format we knew how to reason with, we would simply merge the agent’s KB with the CSKB.

The spirit of our current approach to commonsense knowledge retrieval, then, is to find ways to use what knowledge is available to us now. Retrieving relevant and consistently correct axioms from ConceptNet might not be possible now, but it is possible to get recommendations of which actions will be the most interesting.

So, we take advantage of such a recommendation by pairing it with an agent that can learn from failed actions and placing it in a world in which the goal state is always reachable. ConceptNet appears to be the KB best suited to provid-

ing these recommendations because it contains many simple assertions stating direct relationships between truly commonsense concepts.

Trying to extract similar recommendations from Cyc has been difficult because of its formal upper ontology, and the overly simple heuristic that we took. We hope to overcome those problems soon by using more semantics of the Cyc KB. But, by characterizing these KBs as ways of finding some kind of mutual information between actions, we open up the possibility of finding that knowledge from less structured sources, like Google. If evaluations show that our recommendation heuristics help, then broadening our information sources could be a fruitful avenue to pursue.

7 Acknowledgements

We wish to acknowledge help from Pace Reagan, Keith Goolsbey, and others at CYCorp for help with furnishing a graph representation of the ResearchCyc knowledge base. We also wish to thank Push Singh and Adam Pease for help with answers about ConceptNet and SUMO, respectively. We also thank Deepak Ramachandran for help with working with and installing ResearchCyc.

References

- [Amir and Doyle, 2002] Eyal Amir and Patrick Doyle. Adventure games: A challenge for cognitive robotics (full version). AAAI’02 workshop on Cognitive Robotics. Also, available at the author’s website (<http://www.cs.uiuc.edu/~eyal/papers/>), 2002.
- [Andre and Russell, 2000] David Andre and Stuart J. Russell. Programmable reinforcement learning agents. In *Proceedings of the 13th Conference on Neural Information Processing Systems (NIPS’00)*, pages 1019–1025. MIT Press, 2000.
- [Bacchus and Kabanza, 2000] Fahiem Bacchus and Froduald Kabanza. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1-2):123–191, 2000.
- [Bertoli and Pistore, 2004] Piergiorgio Bertoli and Marco Pistore. Planning with extended goals and partial observability. In *Proceedings of the 14th Int’l Conf. on Automated Planning and Scheduling (ICAPS’04)*, 2004.
- [Cohen *et al.*, 1998] Paul Cohen, Robert Schrag, Eric Jones, Adam Pease, Albert Lin, Barbara Starr, David Gunning, and Murray Burke. The darpa high-performance knowledge bases project. *AI Magazine*, 19(4):25–49, 1998.
- [Doan *et al.*, 2003] AnHai Doan, Pedro Domingo, and Alon Halevi. Learning to match the schemas of databases: A multistrategy approach. *Machine Learning*, 50:279–301, 2003.
- [Doan, 2002] AnHai Doan. *Learning to map between structured representations of data*. PhD thesis, University of Washington, Seattle, 2002.
- [Doherty and Kvarnström, 2001] Patrick Doherty and Jonas Kvarnström. Talplanner: A temporal logic based planner. *AI Magazine*, 2001. Accepted for publication.

- [Even-Dar *et al.*, 2004] Eyal Even-Dar, Sham M. Kakade, and Yishay Mansour. Reinforcement learning in POMDPs. In *Proceedings of the 17th Conference on Neural Information Processing Systems (NIPS'04)*, 2004.
- [Fikes and Farquhar, 1999] Richard Fikes and Adam Farquhar. Large-scale repositories of highly expressive reusable knowledge. *IEEE Intelligent Systems*, 14(2), 1999.
- [Ginsberg, 1991] Matthew L. Ginsberg. The computational value of nonmonotonic reasoning. In James Allen, Richard Fikes, and Erik Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 262–268, San Mateo, CA, USA, April 1991. Morgan Kaufmann Publishers.
- [Hlubocky and Amir, 2004] Brian Hlubocky and Eyal Amir. Knowledge-gathering agents in adventure games. In *AAAI-04 Workshop on Challenges in Game AI*. AAAI Press, 2004.
- [Hoffmann and Nebel, 2001] Jörg Hoffmann and Bernhard Nebel. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.
- [Huberman and Hogg, 1987] B.A. Huberman and T. Hogg. Phase transitions in artificial intelligence systems. *Artificial Intelligences*, 1987.
- [Kaelbling *et al.*, 1996] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: a survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [Lenat, 1995] Douglas B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995.
- [Levesque *et al.*, 1997] H.J. Levesque, R. Reiter, Y. Lesprance, F. Lin, and R. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31:59–84, 1997.
- [McCarthy and Hayes, 1969] John McCarthy and Patrick J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.
- [Miller, 1995] George A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.
- [Mitchell, 1997] Tom Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [Ng and Russell, 2000] Andrew Y. Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *Proc. 17th International Conf. on Machine Learning*, pages 663–670. Morgan Kaufmann, 2000.
- [Ng *et al.*, 1999] Andrew Y. Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: theory and application to reward shaping. In *Proc. 16th International Conf. on Machine Learning*, pages 278–287. Morgan Kaufmann, 1999.
- [Niles and Pease, 2001] Ian Niles and Adam Pease. Towards a standard upper ontology. In *International conference on Formal Ontology in Information Systems*, pages 2–9. ACM Press, 2001.
- [Noy *et al.*, 2001] Natalya F. Noy, Michael Sintek, Stefan Decker, Monica Crubzy, Ray W. Ferguson, and Mark A. Musen. Creating semantic web contents with protege-2000. *IEEE Intelligent Systems special issue on Semantic Web Teconology*, 16(2):60–71, 2001.
- [Pirolli *et al.*, 1996] Peter Pirolli, James Pitkow, and Ramana Rao. Silk from a sow's ear: Extracting usable structures from the web. In *Proc. ACM Conf. Human Factors in Computing Systems, CHI*. ACM Press, 1996.
- [Reiter, 2001] Raymod Reiter. *Knowledge In Action: Logical Foundations for Describing and Implementing Dynamical Systems*. MIT Press, 2001.
- [Salton and Buckley, 1988] Gerard Salton and Chris Buckley. On the use of spreading activation methods in automatic information retrieval. In *SIGIR '88: Proceedings of the 11th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 147–160. ACM Press, 1988.
- [Singh *et al.*, 2002] Push Singh, Thomas Lin, Erik T. Mueller, Grace Lim, Travell Perkins, and Wan Li Zhu. Open mind common sense: Knowledge acquisition from the general public. In *Proceedings of the First International Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems*, LNCS. Springer-Verlag, 2002.