

Situated Structure Learning of a Bayesian Logic Network for Commonsense Reasoning

Haley Garrison and Sonia Chernova

Georgia Institute of Technology
Atlanta, GA

hgarrison3@gatech.edu, chernova@cc.gatech.edu

Abstract

This paper details the implementation of an algorithm for automatically generating a high-level knowledge network to perform commonsense reasoning, specifically with the application of robotic task repair. The network is represented using a Bayesian Logic Network (BLN) (Jain, Waldherr, and Beetz 2009), which combines a set of directed relations between abstract concepts, including *IsA*, *AtLocation*, *HasProperty*, and *UsedFor*, with a corresponding probability distribution that models the uncertainty inherent in these relations. Inference over this network enables reasoning over the abstract concepts in order to perform appropriate object substitution or to locate missing objects in the robot's environment. The structure of the network is generated by combining information from two existing knowledge sources: ConceptNet (Speer and Havasi 2012), and WordNet (Miller 1995). This is done in a "situated" manner by only including information relevant to a given context. Results show that the generated network is able to accurately predict object categories, locations, properties, and affordances in three different household scenarios.

Introduction

Imagine a world in which autonomous robots are available for everyday people: you could go to the store, pick up a robot, and place it in your home. You could ask the robot to make dinner, do your laundry, or clean the house. However, in order for a robot to execute such high-level tasks in new or uncertain environments, it must be able to adapt the learned tasks to its local environment and repair any missing information from the tasks. For example, say a robot is cooking a known recipe in a new kitchen. The cookware and other objects it originally used no longer exist. It must instead reason about high level concepts (ex. pots and pans), determine which ones are suitable for the task at hand (ex. an object that can be used as a container), and find those objects in the new kitchen based on knowledge of their likely locations (ex. pans can be found in cabinets).

For this type of abstract reasoning to be possible, the robot must be able to consult a commonsense knowledge network

and make inferences over the concepts in this network. To achieve good performance on such inference tasks, the network must have the following properties:

1. The network is situated - it contains only information relevant to the current context
2. The size of the network is small enough for fast, online inference

In order to prevent excessive noise in the network and reduce its size, contextually irrelevant concepts and relations must be excluded. However, even if irrelevant information is excluded, the size of the network may still be too large to facilitate fast inference, so it will also be necessary to exclude concepts and relations that are redundant or carry little valuable information. Since no existing knowledge network holds both of these properties, an automated structure learning algorithm was implemented to combine information from existing sources in a situated manner based on the current context of the robot.

Related Work

Two existing commonsense knowledge networks that are widely used for a variety of applications are WordNet (Miller 1995) and ConceptNet (Speer and Havasi 2012). WordNet consists of a collection of synsets, which connect concepts hierarchically through the *IsA* relation. WordNet also distinguishes between different senses of the same word and provides glosses, or definitions, for each sense. While WordNet is clean and hand-coded, it also lacks diversity in the types of relations it contains. ConceptNet, on the other hand, contains a wide variety of different relations, but it does not distinguish between word senses and it is not hand coded, leading to a large amount of noise.

The closest known work to that proposed in this paper is the KnowRob project (Tenorth and Beetz 2013). In this work, the authors created a knowledge network from a variety of encyclopedic sources and represented the network using Prolog rules and the Web Ontology Language (OWL). This network was used to perform plan repair by filling in missing low-level details from high-level task descriptions. However, this representation resulted in a large network without contextual refinement. It also consisted of many separate components and lacked a unified model. Furthermore, the concepts used in the network were manually

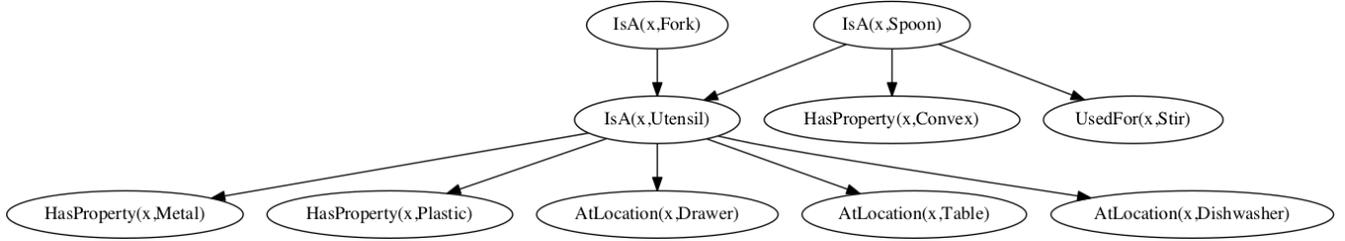


Figure 1: An example of high-level knowledge representation using a BLN.

selected according to the perceived relevance to robotic applications rather than automatically generating the network.

Other related works have had similar shortcomings. Zhu, et al. (Zhu, Fathi, and Fei-Fei 2014) performed affordance prediction on a set of images by using a Markov Logic Network (MLN) (Richardson and Domingos 2006) to represent affordance knowledge. Like KnowRob, this work did not deal with context and used hand-selected objects and affordances in the network. In (Chen and Liu 2011), contextual noise was addressed by disambiguating the concepts in ConceptNet to enrich the WordNet senses with more diverse knowledge for improved performance on word sense disambiguation tasks. While disambiguating ConceptNet helped provide context for each of its concepts, the resulting knowledge base was not further limited in size based on the context of any particular domain. In contrast to this approach, (Stoica and Hearst 2004) did construct a situated knowledge hierarchy in a (nearly) automated way. However, it only included the *IsA* relation and did not enrich this information with other relations from sources like ConceptNet.

Bayesian Logic Networks

The knowledge network generated by this work is represented using a Bayesian Logic Network (BLN) (Jain, Waldherr, and Beetz 2009). BLN's are a type of directed statistical relational model that serves as a template for a Bayesian Network by representing each node as a function/predicate with arguments rather than a single random variable. Additionally, BLN's allow logical constraints, represented as first-order logic rules, to be imposed on the network. A BLN is formally defined as a tuple, $\mathcal{B} = (\mathcal{D}, \mathcal{F}, \mathcal{L})$, such that:

- $\mathcal{D} = (\mathcal{T}, \mathcal{S}, \mathcal{E}, t)$ is the declaration, where \mathcal{T} consists of the declared types, \mathcal{S} is a set of function signatures, and \mathcal{E} is a set of abstract entities where $t : \mathcal{E} \rightarrow 2^{\mathcal{T}} \setminus \{\emptyset\}$ is a mapping from each entity its respective types.
- \mathcal{F} defines a set of "fragments" of a conditional probability distribution. Each fragment represents a directed conditional dependence between two abstract random variables (a parent and a child). These random variables consist of a function $f(p_1, \dots, p_n)$, $f \in \mathcal{S}$, where each of the parameters of f can either be a "meta-variable" or an entity $e \in \mathcal{E}$. The fragments are represented by a conditional probability function (CPF) that specifies a distribution over the child variable for each configuration of the parent variables.

- \mathcal{L} is a set of deterministic constraints described as first-order logic formulas over the abstract random variables.

Before inference can be performed on the network, a mixed ground instantiation $M = ((X, D, G, P), (X, D, C))$ of the BLN, \mathcal{B} , is generated. In this case, X is the set of grounded random variables, $f(e_1, \dots, e_n), e_i \in \mathcal{E}$, D is the domain of the random variables produced by each grounding, G specifies the connectivity of the graph given by the fragments in \mathcal{F} , and P is the conditional probability function for each random variable, as determined by \mathcal{F} . The first-order logic formulas in \mathcal{L} are grounded by substituting the abstract random variables with their groundings and applying constraints C that specify the configurations of the random variables required to satisfy the logic formula.

The inference process proceeds on the grounded network by conditioning the query variables, Q , on the set of evidence variables, V , and marginalizing over the non-query variables, N :

$$P(Q | V = v) = \int_N P(Q, N | V = v) dN$$

Since this marginalization grows exponentially with the number of non-query variables, approximate inference algorithms that have been applied to traditional Bayesian Networks, such as Likelihood Weighting (Fung and Chang 2013) or Gibbs Sampling (Geman and Geman 1984), can be used as an alternative.

To handle the logical constraints, boolean auxiliary variables are added to X for each constraint in C with parent nodes in G for each of the random variables involved in the constraint. This allows the inference process described above to remain unchanged with the addition of logical constraints.

Although some similar works such as (Zhu, Fathi, and Fei-Fei 2014) use undirected Markov Logic Networks (MLN) (Richardson and Domingos 2006), a directed network was chosen for this work because it more explicitly models the directed nature of the relations between concepts. In preliminary tests, the BLN representation was able to perform complex reasoning in both the causal and diagnostic directions, while the MLN suffered poor performance when trying to reason in both directions. Furthermore, MLNs require a gradient descent on the pseudo-log-likelihood to learn the network weights. As a result, the learning process for MLNs is much slower than BLNs which use a simple maximum likelihood frequency count.

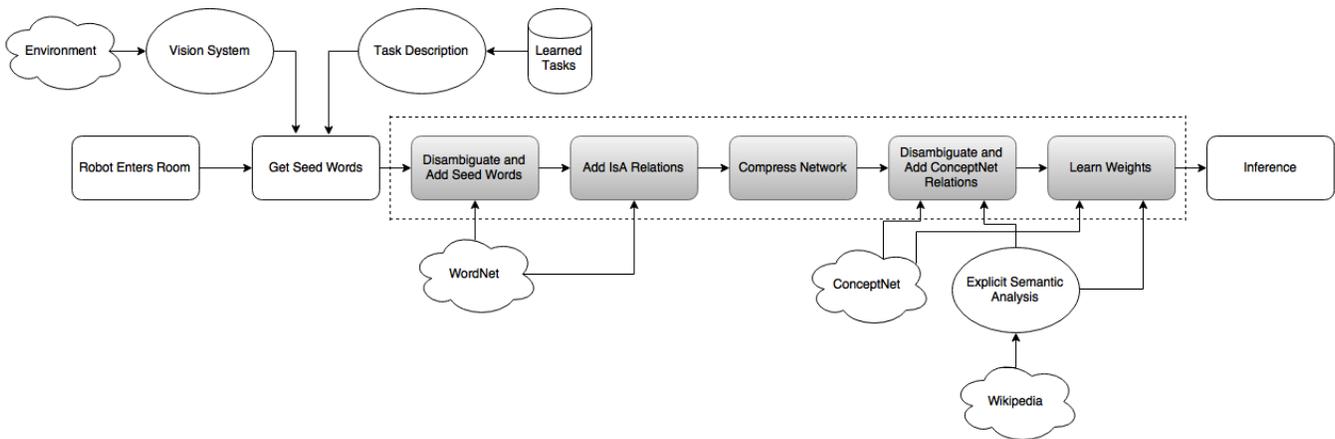


Figure 2: A flowchart showing the overall approach taken to generate the network. The dashed box shows the components of the algorithm, while white boxes are from external sources.

Network Representation

For the proposed knowledge network, the predicates, f , were chosen to be boolean with the following signatures and associated parameter types in \mathcal{T} :

- $\text{IsA}(\text{object}, \text{concept})$
- $\text{HasProperty}(\text{object}, \text{property})$
- $\text{AtLocation}(\text{object}, \text{location})$
- $\text{UsedFor}(\text{object}, \text{affordance})$

The *IsA*, *HasProperty*, and *UsedFor* relations were chosen because they can be used to perform object substitution for plan repair by finding objects that are similar to the original object, or objects that can perform the same function as the missing object. The *AtLocation* relation will allow the robot to reason about possible locations of objects to enable it to find missing objects. For each predicate, the "object" parameter will be a meta-variable that will represent a grounded instance of an object over which the robot can reason, and the "concept," "property," "location," and "affordance" parameters will represent abstract entities.

The general structure of the network fragments in \mathcal{F} will consist of connections such as those shown in Figure 1. Each of these fragments will be associated with a discrete CPF that represents the likelihood of their occurrence. For example, utensils may have some likelihood that they are metal versus plastic, and they have some likelihood that they will be found in a drawer compared to a table or a dishwasher. For the purposes of this paper, no logical constraints were imposed on the network, as experiments so far have shown that high probability relations can effectively be modeled by assigning a probability of one to the corresponding fragment.

Network Generation

An overview of the approach taken for the network generation can be found in Figure 2. The dashed line indicates the components that were implemented as part of the network generation algorithm.

Getting Seed Words

Before the network can be generated, a set of seed words must first be obtained. These seed words should be related to the domain in which the robot is operating and could come from the robot's vision system (objects it sees in its environment), or from the task description. For testing purposes, a set of objects from three different household tasks was extracted and used as input to the network generation algorithm.

Seed Word Disambiguation

After the seed words have been collected, they must be disambiguated to determine the contextually correct senses of the words. For example, the word "pan" has the following four senses in WordNet:

1. pan, cooking pan – cooking utensil consisting of wide metal vessel
2. Pan, goat god – (Greek mythology) god of fields and woods and shepherds and flocks
3. pan – shallow container made of metal
4. Pan, genus Pan – chimpanzees; more closely related to Australopithecus than to other pongids

Given a particular environment, not all of the above senses will be contextually relevant. To keep the size of the network small and contextually accurate, the seed words can be disambiguated and the irrelevant senses can be excluded from the generated network.

Since WordNet provides information on the different word senses, it can be used to perform this disambiguation. The approach used for this paper is similar to that in (Tsatsaronis, Varlamis, and Vazirgiannis 2008). Given that the seed words originate from the same context, they are likely to be semantically similar. Therefore, disambiguation can be performed by finding the sense of each word that maximizes the overall similarity between the seed words. To

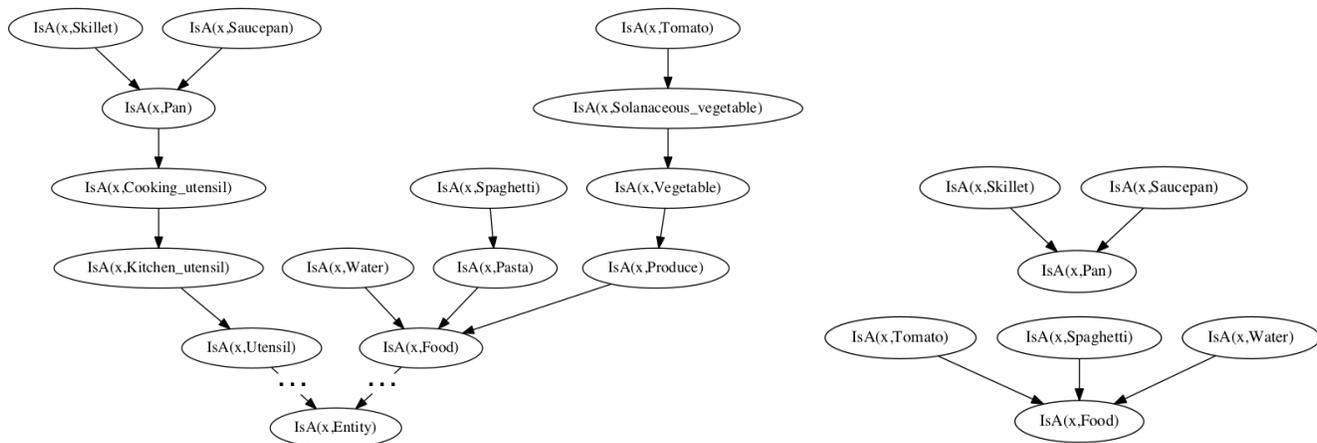


Figure 3: An example of the *IsA* relation before (left) and after (right) compression. Compression reduces the amount of redundant and high-level information in the network.

do so, the disambiguation algorithm finds a Minimum Spanning Tree (MST), where each node represents the most relevant sense of one of the seed words. Given a set of n seed words, $W = \{w_1, w_2, \dots, w_n\}$, and a set of possible senses (synset), $S_i = \{s_i^1, s_i^2, \dots\}$, $i \in \{1, 2, \dots, n\}$, for each seed word, the MST, $T = \{s_1^*, s_2^*, \dots, s_n^*\}$, is computed, where s_i^* indicates the chosen sense of the word, w_i , that minimizes the cost of the tree. The cost metric used for this algorithm is $C_{ij} = \min_{k,l} 1 - wup(s_i^k, s_j^l)$, where $wup()$ is the Wu & Palmer similarity measure (Wu and Palmer 1994) for the senses, k and l , of the i^{th} and j^{th} words, respectively. This measure is based on the length of the path between the two senses of the two words in the WordNet hierarchy, where a longer path generally indicates less semantic similarity.

Since it is possible for a given set of seed words to have multiple minimum spanning trees, the starting node for the MST is chosen as the word, w_i , for which the number of senses, $|S_i|$, is a minimum. Then for each of the senses of this word, the MST is computed, and the MST with the lowest overall cost, determined by the sum of the costs of all edges in the MST, is chosen. While this does not guarantee that the best possible MST is found, it avoids having to compute all possible MSTs, thereby reducing computation time. Using the MST approach also assumes that all of the seed words are connected. While this might not always be true, it is likely that the seed words are related by context, so the MST approach should yield good results in most cases.

IsA Relation and Compression

After the seed words have been disambiguated, the *IsA* relation is added to the network by traversing the WordNet hypernym hierarchy from each of the disambiguated seed words to the root node and adding each node along this path to the network. Although WordNet is hand-coded, it does contain a large amount of redundant and high-level concepts that convey little information, as shown in Figure 3. If these nodes are not removed from the network, this can lead to rapid expansion in the size of the network when other re-

lations are added. To reduce the size of the network to a manageable level and remove the high-level and redundant nodes, the compression strategy implemented in (Stoica and Hearst 2004) is employed. The compression uses the following three rules:

1. Eliminate selected top-level (very general) categories, like abstraction, entity.
2. Starting from the leaves, eliminate a parent that has fewer than n children, unless the parent is the root.
3. Eliminate a child whose name appears within the parent's.

For the first rule, "top-level" categories are defined as words with an Information Content (IC) of less than 5.0 when evaluated against the Brown corpus.

Adding ConceptNet Relations

In addition to the *IsA* relation, the *UsedFor*, *HasProperty*, and *AtLocation* relations are added from ConceptNet. To do so, the relations in ConceptNet are first disambiguated to remove contextually irrelevant relations. An approach similar to that in (Chen and Liu 2011) is used for this purpose. For each ConceptNet relation, $\langle c, relation, d \rangle$, where d is an ambiguous word and c is disambiguated, the Word Sense Profile, $WSP(d_i) = w_1, w_2, \dots$ is generated for each sense, d_i , of the word, d . Each w_j in the WSP is a word from one of the following sources in WordNet:

1. All synonyms of d_i
2. All words (excluding stop words) in the gloss/definition for d_i
3. All direct hypernyms (parent nodes) and hyponyms (child nodes) of d_i in WordNet
4. All meronyms/holonyms (has part or part of) relations in WordNet
5. All words (excluding stop words) in the glosses of the direct hyponyms of d_i

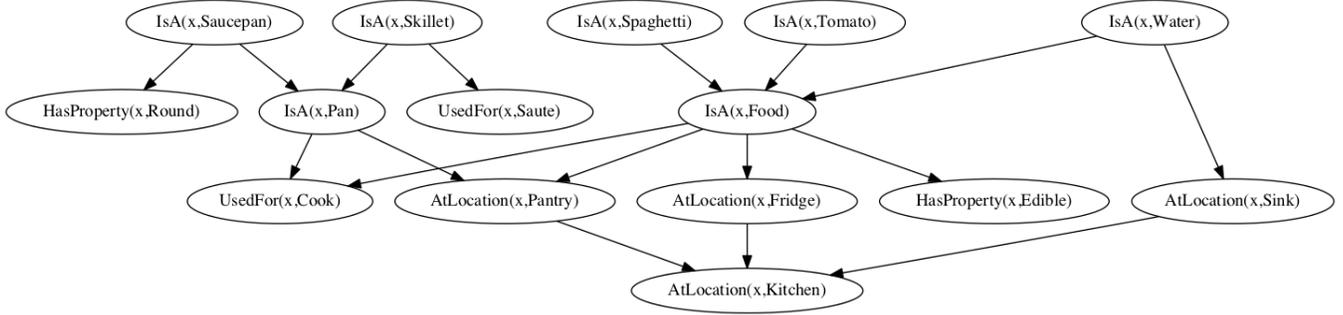


Figure 4: An example of a network after adding all four types of relations from WordNet and ConceptNet.

After the WSP has been generated for each sense, a score is computed for each of the WSPs. This score is equal to the sum of the semantic relatedness between the non-ambiguous word, c , and each word in $WSP(d_i)$. The i^{th} sense is chosen if the score is maximal for that value of i . In (Chen and Liu 2011), the relatedness is measured using the Normalized Google Distance (NGD), which is based on the number of Google hits returned for the two words together. However, since the current version of the Google Search API limits the number of queries per day, a different semantic relatedness measure called Explicit Semantic Analysis (ESA) (Gabrilovich and Markovitch 2007) was used instead. ESA uses a pre-processed dump of Wikipedia to generate a large table, where the columns, $C = \{c_i\}, i \in [1, N]$, represent each of the concepts (pages) in Wikipedia, and the rows, $W = \{w_j\}, j \in [1, M]$, represent the words on those pages. The entries, $E = \{e_{ij}\}$, in the table represent the frequency count of the words, w_j , in each Wikipedia page, c_i . The semantic relatedness between two words is computed by taking each corresponding row, E_j , as a weighted vector of concepts and computing the cosine distance between the two vectors.

After each of the ConceptNet relations has been disambiguated, only the relations corresponding to the correct senses of each word are added to the BLN. This helps prevent contextually irrelevant information from being added to the network. Additionally, relations, $\langle c, relation, d \rangle$, where d consists of more than one word were excluded from the network. Since ConceptNet is not hand-coded, it contains a significant amount of noisy or erroneous relations. Excluding such relations was found to significantly reduce the size of the network without removing a large number of correct relations.

Due to the hierarchical nature of locations, two hops in ConceptNet were allowed when adding the *AtLocation* relation to the network. However, to prevent the size of the network from expanding rapidly with the increased number of hops, and to prevent contextually irrelevant locations from being added to the network, any locations that were added to the network were constrained to locations within the robot’s current environment. For example, $IsA(x,Food) \rightarrow AtLocation(x,Store)$, would be excluded if the robot’s current environment is in the kitchen, since stores are not located within kitchens. A sample of the output of the al-

gorithm after adding the relations from ConceptNet can be seen in Figure 4.

Weight Learning

The final component of the network generation is to learn the CPF for each fragment in the network. To do so, a set of training evidence is generated with a likelihood equal to a linear combination of the weights assigned to each relation in ConceptNet and the ESA relatedness measure between the two concepts in the relation. In future work, this evidence will be augmented with evidence collected by the robot, but generating simulated evidence will provide an initial estimate of the real-world probabilities and enable inference results to be ranked according to their relative likelihoods. Once the evidence has been collected, the CPFs can be learned via maximum likelihood by counting the frequency of each child node being true for each configuration of the parent nodes.

Evaluation

To evaluate the network generation algorithm, three sets of seed words were collected, from three different task scenarios related to typical household chores. The three scenarios included cooking a recipe, doing laundry, and cleaning the house, with 19, 15, and 11 seed words, respectively. An example of some of the relations the network generated in each case can be seen in Table 1. For each seed word, s_i , inference was run over the network where the evidence variable was $IsA(Object_i, s_i)$ with a value set to true. Queries were then made for the variables $IsA(Object_i, x)$, $AtLocation(Object_i, x)$, $HasProperty(Object_i, x)$, $AtLocation(Object_i, x)$, and $UsedFor(Object_i, x)$ for each seed word. The output of the inference process was then compared to a gold standard. This gold standard was created by hand labeling each of the possible inference outputs as either true or false. Each query result was assumed to be true if the associated probability was greater than 0.5, and false otherwise. Table 2 shows the results from each of the three scenarios on each of the four query types.

In all three cases, the highest accuracy occurred for the *IsA* relation. Since WordNet is hand-coded, the majority of the *IsA* relations were correct when compared to the gold standard. Most of the errors that occurred with the *IsA* rela-

Source	Relations
Recipe	$\text{IsA}(x, \text{Garlic}) \rightarrow \text{IsA}(x, \text{Flavorer})$ $\text{IsA}(x, \text{Food}) \rightarrow \text{AtLocation}(x, \text{Container})$ $\text{IsA}(x, \text{Container}) \rightarrow \text{HasProperty}(x, \text{Plastic})$ $\text{IsA}(x, \text{Stove}) \rightarrow \text{UsedFor}(x, \text{Heat})$
Laundry	$\text{IsA}(x, \text{Washer}) \rightarrow \text{IsA}(x, \text{Appliance})$ $\text{IsA}(x, \text{Sock}) \rightarrow \text{AtLocation}(x, \text{Dresser})$ $\text{IsA}(x, \text{Towel}) \rightarrow \text{HasProperty}(x, \text{Cotton})$ $\text{IsA}(x, \text{Shirt}) \rightarrow \text{UsedFor}(x, \text{Dress})$
Cleaning	$\text{IsA}(x, \text{Rag}) \rightarrow \text{IsA}(x, \text{Piece_of_cloth})$ $\text{IsA}(x, \text{Soap}) \rightarrow \text{AtLocation}(x, \text{Sink})$ $\text{IsA}(x, \text{Paper_towel}) \rightarrow \text{HasProperty}(x, \text{Paper})$ $\text{IsA}(x, \text{Broom}) \rightarrow \text{UsedFor}(x, \text{Sweep})$

Table 1: An example of relations added to each of the three test networks.

tion corresponded to the seed words that had not been disambiguated correctly. As shown in Table 3, the lowest disambiguation accuracy occurs for the recipe scenario, and the highest accuracy is achieved for the cleaning scenario. This corresponds to the inference performance for each of the three tasks, with recipe achieving the lowest accuracy and cleaning achieving the highest.

Overall, accuracies for the *AtLocation*, *HasProperty*, and *UsedFor* relations ranged from the low to mid seventies to upper eighties. Several sources of error limited the inference accuracy of the network. Errors made during earlier stages in the network tended to propagate through the rest of the network. For example, if a seed word was incorrectly disambiguated, the relations added from ConceptNet would often be related to the incorrect sense of the seed word. Other sources of error came from the noise present in ConceptNet. In some cases, the relations themselves are inaccurate – the recipe network, for example, included the relation $\text{IsA}(x, \text{Container}) \rightarrow \text{UsedFor}(x, \text{Wash})$. In other cases, this noise appeared in the form of missing connections within the network. Although the $\text{IsA}(x, \text{Saucepan}) \rightarrow \text{UsedFor}(x, \text{Saute})$ connection existed in the recipe network, the $\text{IsA}(x, \text{Frying_pan}) \rightarrow \text{UsedFor}(x, \text{Saute})$ relation was missing, though both objects are arguably equally suited to the task of sauteing. This error occurred because the $\text{IsA}(x, \text{Frying_pan}) \rightarrow \text{UsedFor}(x, \text{Saute})$ does not exist at all in ConceptNet.

Future Work

One of the main goals of future work is to perform more extensive evaluation on the network generation algorithm. This could include use of crowdsourcing to develop a gold standard that more accurately reflects the uncertainty of the relations. Additionally, the algorithm could be tested across several different domains to determine whether it generalizes beyond the household scenarios presented in this paper.

Another future goal is to perform grounding to allow a robot to associate the real objects it encounters in its environment with the abstract concepts over which it can perform inference. Grounding the network will enable the robot to

Source	IsA	AtLocation	HasProperty	UsedFor
Recipe	97.6	86.8	82.0	88.1
Laundry	98.3	77.3	88.9	89.5
Cleaning	98.6	72.7	94.7	79.2

Table 2: Total inference accuracy (as a percentage) over the set of seed words from each source when compared to the gold standard.

Source	Accuracy
Recipe	73.7
Laundry	80.0
Cleaning	81.8
Mean	78.5

Table 3: Seed word disambiguation accuracy for each of the three testing scenarios.

use this high-level knowledge to perform plan repair by locating missing objects or finding suitable substitutes. At the point where task execution fails, the appropriate query can be formulated, and the ranked inference results can then be grounded to allow the robot to attempt to continue execution.

The last goal of this work is to enable both the network structure and associated probability distribution to be updated online without the need to regenerate the network. Updates to the conditional probabilities could be performed as the robot collects new evidence on its own or through interaction with humans such as question asking. The structure of the network could also be modified by adding or removing nodes as the robot encounters new objects or determines that portions of the network have been unused and are unnecessary.

References

- [Chen and Liu 2011] Chen, J., and Liu, J. 2011. Combining ConceptNet and WordNet for Word Sense Disambiguation. In *International Joint Conference on Natural Language Processing*, 686–694.
- [Fung and Chang 2013] Fung, R., and Chang, K.-C. 2013. Weighing and Integrating Evidence for Stochastic Simulation in Bayesian Networks.
- [Gabrilovich and Markovitch 2007] Gabrilovich, E., and Markovitch, S. 2007. Computing Semantic Relatedness using Wikipedia-based Explicit Semantic Analysis. In Veloso, M. M., ed., *International Joint Conference on Artificial Intelligence*, 1606–1611.
- [Geman and Geman 1984] Geman, S., and Geman, D. 1984. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6(6):721–741.
- [Jain, Waldherr, and Beetz 2009] Jain, D.; Waldherr, S.; and Beetz, M. 2009. Bayesian Logic Networks. Technical report, Technische Universität München, München.

- [Miller 1995] Miller, G. A. 1995. WordNet: A Lexical Database for English. *Communications of the ACM* 38(11):39–41.
- [Richardson and Domingos 2006] Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine Learning* 62(1-2):107–136.
- [Speer and Havasi 2012] Speer, R., and Havasi, C. 2012. Representing General Relational Knowledge in ConceptNet 5. In *Proceedings of the Eight International Conference on Language Resources and Evaluation*.
- [Stoica and Hearst 2004] Stoica, E., and Hearst, M. A. 2004. Nearly-Automated Metadata Hierarchy Creation. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 117–120.
- [Tenorth and Beetz 2013] Tenorth, M., and Beetz, M. 2013. KnowRob – Knowledge Processing for Autonomous Personal Robots. *International Journal of Robotics Research* 32.
- [Tsatsaronis, Varlamis, and Vazirgiannis 2008] Tsatsaronis, G.; Varlamis, I.; and Vazirgiannis, M. 2008. Word Sense Disambiguation with Semantic Networks. In Sojka, P.; Horák, A.; Kopeček, I.; and Pala, K., eds., *Text, Speech, and Dialogue*, 219–226. Brno: Springer.
- [Wu and Palmer 1994] Wu, Z., and Palmer, M. 1994. Verbs semantics and lexical selection. In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics -*, 133–138. Morristown, NJ, USA: Association for Computational Linguistics.
- [Zhu, Fathi, and Fei-Fei 2014] Zhu, Y.; Fathi, A.; and Fei-Fei, L. 2014. Reasoning About Object Affordances in a Knowledge Base Representation. In *European Conference on Computer Vision*.