

SQuARE: Semantics-based Question Answering and Reasoning Engine*

Kinjal Basu, Sarat Chandra Varanasi, Farhad Shakerin and Gopal Gupta
The University of Texas at Dallas, Richardson, USA

{Kinjal.Basu, Sarat-Chandra.Varanasi, Farhad.Shakerin, Gopal.Gupta}@utdallas.edu

Understanding the meaning of a text is a fundamental challenge of natural language understanding (NLU) and from its early days, it has received significant attention through question answering (QA) tasks. We introduce a general semantics-based framework for natural language QA and also describe the SQuARE system, an application of this framework. The framework is based on the denotational semantics approach widely used in programming language research. In our framework, valuation function maps syntax tree of the text to its commonsense meaning represented using basic knowledge primitives (the semantic algebra) coded using answer set programming (ASP). We illustrate an application of this framework by using VerbNet primitives as our semantic algebra and a novel algorithm based on partial tree matching that generates an answer set program that represents the knowledge in the text. A question posed against that text is converted into an ASP query using the same framework and executed using the s(CASP) goal-directed ASP system. Our approach is based purely on (commonsense) reasoning. SQuARE achieves 100% accuracy on all the five datasets of bAbI QA tasks that we have tested. The significance of our work is that, unlike other machine learning based approaches, ours is based on “understanding” the text and does not require any training. SQuARE can also generate an explanation for an answer while maintaining high accuracy.

KEY WORDS: Natural Language Understanding, Answer Set Programming, Question Answering, Commonsense Reasoning, VerbNet

1 Introduction

Semantic representation of natural language is a long term goal of natural language understanding (NLU) research. To understand the meaning of a natural language sentence, humans first process the syntactic structure of the sentence and then infer the meaning next. Also, humans use commonsense knowledge to understand the often complex and ambiguous meaning of natural language sentences. Humans interpret a passage as a sequence of sentences and will normally process the events in the story in the same order as the sentences. Once humans understand the meaning of a passage, they can answer questions posed, along with an explanation for the answer. We believe that an automated question answering (QA) system should work in a similar way. Thus, an ideal automated QA framework should map each sentence to the knowledge it represents, then augment it with commonsense knowledge related to the concepts involved (e.g., if the passage involves a lion, the system should know that a lion is a living being, an animal, carnivore, etc.), then use the combined knowledge to automatically answer questions. Questions should be first automatically converted into a query that is executed against this knowledge. We develop such a framework in this paper that relies on denotational semantics to specify the mapping from syntax to knowledge (meaning). This denotational, automated QA framework has been employed to realize the SQuARE engine described in this paper.

Researchers have created several QA datasets to test a QA system, such as SQuAD-1.0 [19], bAbI [23], MCTest [20], SQuAD-2.0 [18], WikiQA [24], etc. However, most of the state-of-the-art QA systems have been

*Work partially supported by NSF awards IIS 1718945, IIS 1910131, IIP 1916206, and a DARPA award.

developed using machine learning (ML) techniques. The common drawback among all of these ML-based systems is that they answer questions by exploiting the correlation among words and without properly understanding the content. These systems are also not explainable, i.e., cannot generate a justification of their conclusions. They may perform well for a particular dataset or a domain, but will fail on other tasks. Humans can easily generalize from one type of question to another, even if they have not encountered it before. This leads to another problem for all these machine learning based systems: they fail to answer questions for which they have not seen training data. Most of such ML-based QA systems work as information retrieval agents that can answer questions against a textual passage with either shallow (or, no) reasoning capabilities. In this paper, we present a general semantics-based framework for automated QA that is a step towards mimicking the human question answering mechanism.

We also introduce our novel algorithm for automatically generating the knowledge (semantics) corresponding to each English sentence. The algorithm leverages the comprehensive verb-lexicon for English verbs—VerbNet [6]. For each English verb, VerbNet gives the syntactic and semantic patterns. The algorithm employs partial syntactic matching between parse-tree of a sentence and a verb’s *frame syntax* from VerbNet to obtain the meaning of the sentence in terms of VerbNet’s primitive predicates. This matching is motivated by denotational semantics and can be thought of as mapping parse-trees of sentences to knowledge that is constructed out of semantics provided by VerbNet. The VerbNet semantics is expressed using a set of primitive predicates that can be thought of as the *semantic algebra* of the denotational semantics [21].

SQuARE has been tested on five datasets (*single supporting fact, two supporting facts, three supporting facts, two argument relations* and *positional reasoning*) of the bAbI QA tasks, that is created by Facebook-Research for testing text understanding and reasoning of a QA system. The bAbI QA datasets consist of stories. Each story is a collection of simple sentences and the questions are asked in between a set of sentences interactively or at the end of the story. Answer for each question can be computed by performing commonsense reasoning over the information captured from the earlier sentences. We choose bAbI QA dataset to test SQuARE, as the sentences involved are simpler and require less amount of pre-processing work so we can concentrate on the semantic representation and reasoning tasks. The SQuARE system achieves 100% accuracy on all five tasks tried, outperforming state-of-the-art systems. It also surpasses other systems wrt explainability as it can provide proper justification for each answer.

This paper makes the following novel contributions: (i) furnishes a fully explainable semantic-based general framework for textual question answering, (ii) demonstrates that a QA application—SQuARE—developed using the framework can reach 100% accuracy without training; it also outperforms the machine learning based models in terms of explainability, (iii) showcases the novel partial tree matching algorithm using VerbNet primitives as semantic algebra for the generation of a sentence semantics, (iv) shows how commonsense reasoning can be included for QA in dynamic domains, and (v) provides a method that relies on goal-directed ASP that guarantees a correct answer as long as the semantic representation and the query generation are correct. Our work is purely based on reasoning and does not require any manual intervention other than providing reusable commonsense knowledge coded in ASP. It paves the way for developing advanced NLU systems based on “truly understanding” the text.

2 Background

We next describe some of the key technologies we employ. Our effort is based on answer set programming (ASP) technology [4, 3], specifically, its goal-directed implementation in the s(ASP) and s(CASP) systems [13, 1]. The semantics of a sentence and *frame semantics* of VerbNet that we use are represented as answer set programs. The use of a goal-directed implementation of ASP is critical to our project. We assume that the reader is familiar with ASP. Expositions of ASP can be found elsewhere [3].

The s(ASP) System: s(ASP) [13] is a query-driven, goal-directed implementation of ASP. The s(CASP) system [1] is an extension of s(ASP) that adds constraint solving over reals. Goal-directed execution of s(ASP)/s(CASP)

is indispensable for automating commonsense reasoning, as traditional grounding and SAT-solver based implementations of ASP may not be scalable. There are three major advantages of using s(ASP) (and s(CASP)) system: (i) s(ASP) does not ground the program, which makes our framework scalable, (ii) it only explores the parts of the knowledge base that are needed to answer a query, and (iii) it provides justification (proof tree) for an answer.

Stanford CoreNLP Tool Set: Stanford CoreNLP [11] is a set of tools for natural language processing. SQuARE only uses *parts-of-speech (POS) tagger*, *named-entity recognizer (NER) tagger*, *constituency parser* and *dependency parser* from this set. Based on the context, POS tagger generates the necessary parts of speech such as noun, verb, adjective, etc., for each statement and also identifies the question type (e.g., what, where, who). The NER tagger identifies named-entities, such as a person’s name. Constituency parser produces the syntactic parse tree (shown in figure 5) of a sentence where the node’s value follows Penn-Treebank’s tags [12]. The dependency graph provides the grammatical relations between words in a sentence. Dependency relations follow enhanced universal dependency annotation [17]. Figure 1 shows an example of the POS tagging and the dependency graph of a statement from the *two argument relations* task from the bAbI dataset [23].




Figure 1: Example of POS tagging and dependency graph

VerbNet: VerbNet [6] is the largest online network of English verbs. It is inspired by Beth Levins classification of verb classes and their syntactic alternations [9]. A verb class in VerbNet is mainly expressed by *syntactic frames*, *thematic roles*, and *semantic representation*. The VerbNet lexicon identifies thematic roles and syntactic patterns of each verb class and infers the common syntactic structure and semantic relations for all the member verbs. Figure 2 shows an example of a VerbNet frame of the verb class *discard*. Note that there can be multiple frames for a given verb class, one for each type of usage of the verb. The first line in the figure (*NP V NP PP.Destinations*) shows the syntax tree pattern for this particular usage. The second line shows an example of English sentence that matches the usage. The third line (we call it *frame syntax*) shows the constraints on the syntactic categories NP, PP, etc. (e.g., the first noun phrase should be of type AGENT). The final line (we call it *frame semantics*) shows the meaning expressed using basic primitive predicates of VerbNet (*exert_force*, *contact*, *motion*, etc.). A verb in an English sentence generally corresponds to an event. The meaning expressed using primitive predicates represents the semantic relations between event participants across the course of the event.

3 General Architecture for Semantics-based NLP QA

Denotational semantics has been extensively used in the study of programming languages. Broadly, denotational semantics is concerned with defining the meaning of a program in terms of mathematical objects called domains, such as integers, truth values, tuples of values, and mathematical functions [21]. Typically, there are three components of denotational semantics of a programming language:

- *Syntax*: specified as abstract syntax trees.

	NP V NP PP.Destinations
Example	“ Steve tossed the ball to the garden.”
Syntax	Agent V Theme { {+Dest— +Loc} } Destination
Semantics	<i>Exert_Force(During(E₀),Agent,Theme), Contact(End(E₀),Agent,Theme), Motion(During(E₁),Theme), - Contact(During(E₁), Agent, Theme), Path_Rel(Start(E₀), Theme, Initial_Location, Ch.of_Loc, Prep), Path_Rel(End(E₁), Theme, Destination, Ch.of_Loc, Prep), Cause(Agent, E₁), Meets(E₀, E₁), Equals(Agent, Initial_Location)</i>

Figure 2: VerbNet frame instance for the verb class *discard*

- *Semantic Algebra*: these are the basic domains along with the associated operations; meaning of a program is expressed in terms of these basic domains.
- *Valuation Function*: these are mappings from abstract syntax trees and semantic algebra to values in the domains in the semantic algebra.

To “truly understand” a text, an ideal NLU system should transform the natural language text into knowledge expressed using the semantic algebra of well-understood concepts. Then, this semantics can be used for different NLU tasks like question answering, summarization, information extraction, etc. With the formalized textual information, the semantics, and commonsense knowledge, an NLU system can easily perform basic reasoning such as counting, comparing, inference, spatial reasoning. Figure 3 illustrates a general framework for textual question answering and is explained next.




Figure 3: General Framework for Semantic-based QA System

Natural Language Processor: performs the grammatical (syntax) analysis (e.g., lemmatization, parts-of-speech tagging, parsing, etc.) of a natural language sentence and generates the syntactic parse tree via constituency parsing. Many PCFG (probabilistic context free grammar) or deep learning based state-of-the-art parsers are available that find the most probable parse tree from all the trees generated due to the ambiguous nature of the English language.

Valuation Function: takes the parse tree as an input and maps it to the knowledge entailed in the sentence. This knowledge is represented in terms of knowledge primitives (semantic algebra) whose meaning is well-understood.

Semantics Generator: The semantics generator applies the valuation function to parse trees using the partial tree matching algorithm to generate the knowledge encapsulated in a sentence. This knowledge is represented as an answer set program. *Partial* tree matching is needed due to the complexity and ambiguity of natural languages.

Commonsense Knowledge: This module represents the background commonsense knowledge needed for reasoning. When some information is not explicitly mentioned in a passage and the system needs it to make inferences, this commonsense knowledge becomes crucial. The commonsense knowledge is also coded in ASP for consistency of notation. Also, our inference algorithms require that all knowledge be coded in ASP. This handcrafted commonsense knowledge has been represented in a generic manner and can be reused in future. This method is very similar to a human knowledge gathering process. We acquire/refine our commonsense knowledge incrementally from information we encounter in our everyday life and we use it for reasoning later when needed. Commonsense knowledge can be thought of as augmenting the semantic algebra of knowledge primitives. It is also employed to augment knowledge output by valuation functions. Note that all commonsense knowledge is generally represented as rules that define defaults, exceptions and preferences coded in ASP. Our goal is to keep the inference algorithm—default reasoning with exceptions and preferences—as simple as possible. This simplicity is important as otherwise it is possible that our encoding contains the knowledge needed




Figure 4: Sentence to semantic generation process in SQuARE

to answer a question but we are not able to answer it because of the complexity of our decision procedures and not knowing which decision procedure to apply and when.

s(CASP) ASP Engine: The s(CASP) goal-directed ASP engine executes the query generated from the question against the knowledge generated from the text. The s(CASP) system performs a goal-directed search of the knowledge base while applying defaults, resolving exceptions, and retracting defeasible worlds. The (partial) stable models produced in response to the query represents the answer to the question. Moreover, s(CASP) performs all these tasks without grounding the program, avoiding exponential blowup in the grounded program size common in SAT-based systems.

4 Semantics-driven QA: the SQuARE System

Unlike programming languages, the denotation of a natural language can be quite ambiguous. English is no exception and the meaning of a word or sentence may depend on the context. The generation of correct knowledge from a sentence, hence, is quite hard. We next present a novel approach to automatically mapping parse trees of simple English sentences to their denotations, i.e., knowledge they represent.

4.1 Semantics-driven translation of English sentences

The verb plays a significant role in an English sentence as it represents the event encapsulated in the sentence. The verb also constrains the relation among event participants. As described earlier (Section 2), the VerbNet system captures all of this information through verb classes that represent a group of verbs that share common

Algorithm 1 Semantic Knowledge Generation

Input: p_t : constituency parse tree of a sentence
Output: *semantics*: sentence semantics

- 1: **procedure** GETSENTENCESEMANTICS(p_t)
- 2: $verbs \leftarrow getVerbs(p_t)$ ▷ returns list of verbs present in the sentence
- 3: $semantics \leftarrow \{\}$ ▷ initialization
- 4: **for each** $v \in verbs$ **do**
- 5: $classes \leftarrow getVNClasses(v)$ ▷ get the VerbNet classes of the verb
- 6: **for each** $c \in classes$ **do**
- 7: $frames \leftarrow getVNFrames(c)$ ▷ get the VerbNet frames of the class
- 8: **for each** $f \in frames$ **do**
- 9: $thematicRoles \leftarrow getThematicRoles(p_t, f.syntax, v)$ ▷ see Algorithm 2
- 10: $semantics \leftarrow semantics \cup getSemantics(thematicRoles, f.semantics)$
- 11: ▷ map the thematic roles into the frame semantics
- 12: **end for**
- 13: **end for**
- 14: **end for**
- 15: **return** *semantics*
- 16: **end procedure**

meaning (e.g., toss and discard). As discussed earlier, VerbNet also provides a skeletal parse tree (frame syntax) for different usages of a verb and the corresponding semantics (frame semantics) that represent the knowledge. Thus, VerbNet can be thought of as a very large valuation function that maps syntax tree patterns to their respective meanings: essentially, a giant case statement represented through the large collection of VerbNet frames, where each frame represents the mapping of one specific type of usage of a verb class to its meaning. Thus, in Fig. 2, knowledge contained in the sentence “Steve tossed the ball to the garden.” is represented by the conjunction of predicates in the frame semantics (line 4), where predicate argument AGENT is instantiated to ‘Steve’, THEME to ‘the ball’, DESTINATION to ‘the garden’, etc. Fig. 4 shows another example.

Thus, given the parse tree of a sentence, p_i , the SQuARE system matches VerbNet’s syntax tree skeletons of the verbs occurring in the sentence with sub-trees of p_i . The predicates corresponding to the matching skeletons are assembled together to produce the knowledge represented in the sentence, as discussed earlier and illustrated in figure 4. The knowledge is represented in terms of VerbNet primitive predicates (semantic algebra). For example, one relation in the semantic algebra of the verb *discard* is *CAUSE*. The relation *CAUSE* is a binary relation qualifying the actor of *discard* with the verb *discard* itself. That is, *CAUSE* is represented as *CAUSE(Agent, event(discard))* where *Agent* is a thematic role from VerbNet.

In the SQuARE system, first, the English sentence is parsed using Stanford’s CoreNLP parser to obtain the syntactic parse tree of the sentence. In the parse tree (illustrated in Figure 5), there will always be a root node (blue colored), named *Root*, from where the tree expands. All the internal nodes (green colored) possess the *word*, *phrase* or *clause* tags from Penn-Treebank [12]. The word-level tags are the parts-of-speech tag of each word, such as NN (nominal noun), VBD (verb: past tense), etc. The example of phrase-level tags are NP (*noun phrase*), VP (*verb phrase*), PP (*preposition phrase*), etc., and clause-level tags are S (*simple declarative clause*), SBARQ (*Direct question introduced by a wh-word or a wh-phrase*), etc. The leaf nodes (orange colored) of the parse tree are the words from the sentence with the parts-of-speech node as their parent. Figure 5 shows the syntactic parse tree of the sentence - “Mary discarded the milk there.”

Next, the generated syntactic parse tree of a sentence is passed as an input to our *Semantic Knowledge Generation* algorithm 1 that maps the sentence to its meaning. To accomplish this, the list of verbs mentioned in the sentence are collected and their syntactic and semantic information is determined. As discussed earlier, each verb is a part of one or more VerbNet classes and each class has multiple frames. A frame provides information about the syntactic structure of the verb along with thematic roles. The semantic definition of each frame uses pre-defined predicates of VerbNet that have thematic-roles (AGENT, THEME, etc.) as arguments. These predicates are represented in ASP. *Our goal is to find the partial matching between the parse tree and the frame’s syntactic structure and ground the thematic-role variables so that we can get the semantics of the sentence from the frame semantics and represent it in ASP.*

The syntax of a frame consists of word- and phrase-level tags from the Penn Treebank, described earlier. For each frame, VerbNet also provides the mapping between the thematic roles and the tags (Figure 2 illustrates an example). The algorithm starts from the verb node, performs a bottom-up search and finds the matching frame at each level through depth-first traversal. The matching mechanism can also skip words at a level if it does not fit the VerbNet frame. For instance, Algorithm 2 ignores prepositions if it is not used in the frame syntax. If an exact or partial match is found, the algorithm returns the thematic roles to the parent Algorithm

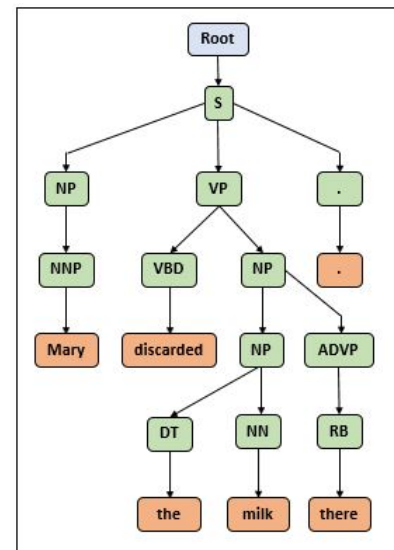


Figure 5: Parse-Tree

1. If a match is not found then the algorithm looks for a match at the parent level. The algorithm terminates when we reach the root without a match. In such a case, we return an empty set. This partial tree matching is summarized in Algorithm 2.

4.2 Passage translation

A passage or a story is a collection of sentences. Humans formulate the meaning of a passage by combining the meaning of each individual sentence. In our mind, we can quickly order the events that occur in the passage and with our commonsense knowledge, we can relate the participants involved (e.g., actor, recipient, or object) in the state before, after, and during each event. When reading a passage, humans construct the knowledge (semantics) of each sentence. The knowledge constructed is non-monotonic, for example, containing defaults and exceptions. The knowledge accumulated as we read the sentences in the passage captures the information of the events in the passage until the last sentence read. For example, when processing the sentence - “*John moved to the bathroom*”, a human would assume that *John* would be at *the bathroom* at the next instant, unless we encounter an exception in the next sentence, such as “*But the door was locked.*”. In other words, we arrange events in a passage in the temporal order corresponding to the order of the sentences containing these events. Predicates associated with events thus will have to have a timestamp argument to capture this ordering. We assume the events in the story have occurred sequentially unless we find exceptions. The meaning of a passage is composed of the meaning of each of its individual sentences. Each sentence is assigned a unique identifier. The meaning of each sentence is represented as an answer set program. The meaning of a passage is the union of the answer set programs of its sentences.

4.3 Commonsense Concepts

Similar to humans, SQuARE needs to use commonsense knowledge to understand the environments and the relationships between objects and persons in the story. Each verb will have associated commonsense knowledge. As an example, consider the verb *possess* arising in the bAbI dataset. SQuARE will have to compute, for example, whether an object is possessed by someone at some point in time and also compute the location of the objects from the knowledge of events in the passage. The property *possession* captures the ownership of an object by a person of the story. Similarly, the property *location* captures the position or place of a person in the story. The commonsense meaning of *possess* is described below as a default rule in ASP.

```
property(possession, T, Per, Obj) :- path_rel(T, start(X), source(_), theme(Obj)),
    cause(T, agent(Per), event(X)), not ab(possess_verb(X)), not unknown_value(Obj).
```

The *path_rel* or *path-relation* predicate is part of the VerbNet semantic algebra that depicts the relationship between the start of an event/action, the *source*, and the *theme*. From the definition of the *path_rel* predicate, the relationship tells that the starting of event *X* moves the *theme* from the *source* (which we do not care about now

Algorithm 2 Partial Tree Matching

Input: p_t : constituency parse tree of a sentence; s : frame syntax; v : verb
Output: tr : thematic role set or *empty-set*: $\{\}$

- 1: **procedure** GETTHEMATICROLES(p_t, s, v)
- 2: $root \leftarrow getSubTree(node(v), p_t)$ ▷ returns the sub-tree from the parent of the verb node
- 3: **while** $root$ **do**
- 4: $tr \leftarrow getMatching(root, s)$ ▷ if s matches the tree return thematic-roles, else $\{\}$
- 5: **if** $tr \neq \{\}$ **then return** tr
- 6: **end if**
- 7: $root \leftarrow getSubTree(root, p_t)$ ▷ returns *false* if $root$ equals p_t
- 8: **end while**
- 9: **return** $\{\}$
- 10: **end procedure**

as it is unknown). *Cause* predicate illustrates that in the same timestamp of the *path_rel*, the *agent* causes the event. The rule has also a weak exception [3] that falsifies *possession* if there is any abnormality present about the possessive verb *X*. Also, the rule confirms the variable *object* is not *unknown* through the negated call `not unknown_value(Obj)`.

Commonsense law of inertia is also needed for the verb *possess*. Thus, the system assumes that the object/person stays in the same state where it was in the previous time-stamp if there is no evidence about the state change of that object/person in the current time-stamp. The reverse is also true, that is, if an object is not possessed in an earlier timestamp, it is not possessed by any person in the current time-stamp, unless there is evidence to the contrary. These default rules are given below (their meaning is fairly obvious and is omitted due to lack of space; note that `after(T2,T1)` stands for time $T2 > \text{time } T1$):

```
(1) property(possession, T2, Per, Obj) :- after(T2, T1),
    property(possession, T1, Per, Obj), not neg_property(possession, T2, Per, Obj).
(2) neg_property(possession, T2, Per, Obj) :- after(T2, T1),
    neg_property(possession, T1, Per, Obj), not property(possession, T2, Per, Obj).
```

4.4 Question translation

Based on the natural language question, SQuARE generates the ASP query, which runs on the goal directed s(CASP) ASP engine to compute the answer. It mainly follows the ASP query generation mechanism described in [2]. For a task, a general query rule is defined that is a collection of sub-queries needed to answer the question. This query rule can be thought of as part of commonsense knowledge. Unlike machine learning based systems, SQuARE is capable of handling all the datasets simultaneously, with no loss in accuracy. Incremental and scalable design of SQuARE helps to use one reasoning component inside another, when needed. For example, the reasoning component to solve the *single supporting fact* task becomes a sub-goal in the *two supporting facts* task.

To answer questions in the bAbI data sets, the systems also need to have commonsense understanding of various concepts. In the case of bAbI, commonsense knowledge of finding a location is needed, which is shown below:

```
(1) get_location(Per, Loc) :- get_all_times(Ts), get_location(Per, Loc, Ts).
(2) get_location(Per, Loc, Ts) :- filter_times(Per, Ts, Fil_Ts),
    get_max_time(Fil_Ts, MaxT), property(location, MaxT, Per, Loc).
(3) get_object_location(Obj, Loc) :- get_all_times(Ts),
    filtered_possession_times(Obj, Ts, Fil_Ts), get_max_time(Fil_Ts, MaxT),
    get_sublist_times(MaxT, Ts), property(possession, MaxT, Per, Obj),
    get_location(Per, Loc, Ts).
```

The rules are self-explanatory. However, from a commonsense reasoning perspective, such rules can be thought of as a process template that a human may follow to achieve a specific task. Humans either learn these patterns on their own or someone imparts the knowledge to them. Supplying these rules assumes the latter.

Initially, the natural language questions are parsed using the Stanford's CoreNLP parser to get the parts-of-speech, the dependency graph, and the named-entities (if any). The named-entities help the system to identify that the query is about a person or a physical object. For example, if there are questions like - "*Where is Sandra?*" and "*Where is the football?*", then named-entity recognizer identifies 'Sandra' as a person. In the process of parsing, the system also detects the question type and the answer type, similar to the process described in [16] and [2]. For instance, '*where*' wh-word normally asks for a location. Then the ASP query is formulated based on the information gathered by parsing the question. The queries for the questions - '*where is Sandra?*' and '*where is the football?*' are as follows respectively:

```
?- get_location(sandra, Location).
?- get_object_location(the_football, Location).
```


5 Example

To illustrate the capability of SQuARE further, we discuss a full-fledged example in this section showing the data-flow and the intermediate outputs from each step. Customized section of a story is taken from Task-2.

1 Mary moved to the bedroom.	4 Mary discarded the milk there.
2 Mary got the milk there.	5 Mary journeyed to the bathroom.
3 Mary travelled to the hallway.	

The story tells that: *Mary moved to the bedroom, got the milk and moved to the hallway, where she left the milk and moved to the bathroom.*

Parsed Output: The Stanford’s CoreNLP parser generates the syntactic parse tree of each statement, that is next input to the semantic generation algorithm (Algorithm 1). Due to lack of space, only one parse tree is shown (statement 4 for story in Figure 5).

Matching VerbNet Frame: Algorithm 1 finds partial matching in the parse tree with the frame syntax from VerbNet and grounds the frame semantics. Figure 2 illustrates an example of a partial matching frame for the verb *discard* that is used in statement (4) of the story. It depicts that if the parse tree partially matches the syntactic structure - *NP V NP PP* then the frame syntax variables — *Agent, Theme, Destination* will be grounded respectively with concrete objects from the sentence. These frame syntax variables are nothing but the thematic roles pre-defined in the VerbNet repository [6]. As we can see, in this scenario, there is a partial matching between the parse tree and the frame structure. The frame syntax variables are the predicate terms in the frame semantics.

Semantic Generation: After processing the whole story and due to the partial matching, the program generates multiple semantics for different verbs. Because of space constraints, we show one of the many predicates of the instantiated semantics of the verb *discard* below. The process of semantic generation is discussed in Section 4.1.

```
exert_force(t4, agent(mary), theme(the_milk)).
```

The variables *Agent* and *Theme* of the frame semantics (shown above) are grounded with the terms *mary* and *the milk*, respectively. Recall that frame semantics is coded in ASP.

Question: Now, the following natural language question is asked against the story. The actual answer is also given in the same line after the question.

```
Where is the milk? hallway
```

ASP Query: Following is the ASP query along with the query rule for this type of question generated by our SQuARE system.

```
get_object_location(Obj, Loc) :- get_all_times(Ts),
    filtered_possession_times(Obj, Ts, Fil_Ts), get_max_time(Fil_Ts, MaxT),
    get_sublist_times(MaxT, Ts), property(possession, MaxT, Person, Obj),
    get_location(Person, Loc, Ts).
?- get_object_location(the_milk, Loc).
```

Answer: The answer found by the SQuARE system matches with the actual answer (*hallway*).

Justification: One major advantage of our framework is that it can provide a justification. This is possible because of the use of query-driven execution under the s(CASP) system. The formatted justification tree for the query corresponding to “*Where is the milk?*” is given below. The answer to the question is *the hallway*, the second argument in the query *get_object_location*. Providing a justification is as simple as printing the s(CASP) proof tree for the query (indentation represents call depth). In the future, we plan to work on converting this justification into English, similar to an explanation that a human would give.

```

BEGIN JUSTIFICATION
  get_object_location(the_milk, the_hallway)
  get_all_times([t1, t2, t3, t4, t5])
  filtered_possession_times(the_milk, [t1, t2, t3, t4, t5], [t2, t3])
  get_max_time([t2, t3], t3)
  get_sublist_times(t3, [t1, t2, t3])
  property(possession, t3, mary, the_milk)
  get_location(mary, the_hallway, [t1, t2, t3])
END JUSTIFICATION

```

6 Experiments and Results

The bAbI Dataset: We have tested our SQuARE system on the bAbI dataset. The bAbI (pronounced as “baby”) question answering tasks [23] were created by the Facebook AI research team to help NLU research. As the name indicates, the tasks are also very easy for humans to reason about and answer with proper justification (though it should be noted that some stories have as many as 71 sentences, so the human effort can be non-trivial). The bAbI QA tasks comprise 20 different datasets for 20 independent reasoning tasks such as basic induction, deduction, temporal reasoning, chaining facts, and many more.

This dataset was mainly created for train-and-test based deep learning models for natural language understanding. For each question asked after a set of declarative statements, the datasets provide the supporting facts for the actual answer as well. Not only it expects a learner to memorize the pattern in between a question and an answer, but also to grasp the supporting sentence for that answer. We mainly choose this dataset to show that a model can be created without training that can “*truly understand*” the story and produce the correct answer to a question with proper justification. Moreover, another purpose of choosing the dataset is that it is actually composed of simple sentences and mainly focuses on reasoning. This helps us to concentrate more on knowledge representation and modeling than pre-processing/parsing of the English sentences.

Tasks Description: SQuARE was tested on five bAbI tasks:

- **Single Supporting Fact:** consists of simple questions that can be answered using single supporting fact given earlier in the story for that particular question. The other information given in the story is irrelevant.
- **Two and Three Supporting Facts:** contains questions that require multi-hop reasoning to find the answer from two or three supporting facts that are chained together with information possibly not relevant to that question.
- **Two Argument Relations:** checks the ability of a system to recognize and differentiate the subject and object of a sentence and use commonsense knowledge to answer a question (e.g. *if room A is south of room B, then B is north of A*).
- **Positional Reasoning:** tests the spatial reasoning of a system. A story describes a few object’s positions and features (e.g., *color of a box*), and asks questions about positional relationship between objects. ML-based system have the most difficulty in this task.

<p>1 John travelled to the hallway. 2 Mary journeyed to the bathroom. 3 Where is John? hallway</p> <p>Example: Single Supporting Facts</p>	<p>1 Mary took the apple there. 2 Mary moved to the bedroom. 3 Mary moved to the garden. 4 Where was the apple before the garden? bedroom</p> <p>Example: Three Supporting Facts</p>
<p>1 The bedroom is west of the kitchen. 2 The hallway is west of the bedroom. 3 What is west of the kitchen? Bedroom</p> <p>Example: Two Argument Relations</p>	<p>1 The red square is below the blue square. 2 The red square is to the left of the pink rectangle. 3 Is the blue square below the pink rectangle? no 4 Is the blue square to the left of the pink rectangle? yes</p> <p>Example: Positional Reasoning</p>

Figure 6: Examples of stories in various tasks

Attributes Tasks	No. of Stories	No. of Questions per Story	Avg. Story Size (No. of Sentences)	Accuracy (%)	Avg. Time per Question (Seconds)
Single Supporting Facts	2200	5	10	100	15.4
Two Supporting Facts	2200	5	22	100	42.8
Three Supporting Facts	1800	5	31	100	147.1
Two Argument Relations	11000	1	2	100	0.6
Positional Reasoning	1375	8	2	100	0.3

Table 1: Performance Results

Example stories for four of the tasks are shown in Fig. 6. Example story for the remaining task (*two supporting facts*) was shown in Sec. 5. Note that the verb ‘*is*’ is not supported in VerbNet. Thus, commonsense knowledge describing the verb ‘*is*’, coded in ASP, had to be added.

Experiment: As discussed earlier, the stories are converted to an answer set program. The s(CASP) engine executes the translated ASP query to get the answer. In our experiments, this process is repeated for each story and for each question posed to that story. As SQuARE does not require any sort of training, we merged the training and testing datasets together to obtain one single repository for each task to test the SQuARE system. Table 1 summarizes our results. By achieving 100% accuracy on all five datasets and providing the proof tree for each answer, SQuARE outperforms all the other machine learning and neural-net based systems in terms of accuracy and explainability.

Average time to answer a question can range from 0.3 seconds to around 2.5 minutes. However, we are confident that these times can be greatly improved by careful analysis of the code based on execution profiling as well as improvements to s(CASP) engine that are continually being made. Note that, all the benchmarks are tested on a *intel i9-9900 CPU* with *16G RAM*.

Comparison of Results & Related Work: Table 2 compares our result in terms of accuracy with other machine learning based models for the five bAbI datasets. N-gram classifier is the baseline model [23] for this dataset whereas the LSTM (long-short-term-memory recurrent neural network) [5] is the most basic deep-learning model tested. The accuracies for both the N-gram classifier and the LSTM are the result of weak supervision used [23], that means the models are trained only on the training dataset given. However, the Structured SVM [23], dynamic memory network [7], and memory-neural-network (MemNN) with adaptive memory (AM), N-grams (NG), and non-linearity (NL) [23] are provided with external supporting facts for training. Multitasking with MemNN shows the accuracy if all the data from all the 20 tasks are used for training. For these tasks, in terms of accuracy, MemNN (AM, NG, NL) and SQuARE achieved 100%. The SQuARE system can provide a justification for the answer, which other ML-based systems cannot.

ASP has also been applied to the bAbI dataset by Mitra and Baral [14, 15] and has been an inspiration for our work, however, theirs is still a machine learning-based system as it employs inductive LP and statistical methods along with ASP. The system achieves great accuracy, however, it requires annotated examples to be fed into the ILP system to derive the hypothesis. This process is iterative and requires “manually finding the group size” [15]. Our SQuARE system, in contrast, relies completely on automated reasoning. No manual intervention

Models Tasks	Single Supporting Facts	Two Supporting Facts	Three Supporting Facts	Two Arguments Relations	Positional Reasoning
N-gram Classifier	36	2	7	50	46
LSTM	50	20	20	61	51
Structured SVM	99	74	17	98	61
DMN	100	98.2	95.2	100	59.6
Multi-Tasking (MemNN)	100	100	98	80	72
MemNN (AM+NG+NL)	100	100	100	100	65
SQuARE	100	100	100	100	100

Table 2: Accuracy (%) comparison of SQuARE with other models for bAbI tasks

is required other than the necessary task of providing commonsense knowledge. Also, in the ILP based approach, the ‘mode’ declaration needs to be manually encoded for each verb used in each task, whereas SQuARE can generate ASP code automatically using our novel semantic knowledge generation algorithm (discussed in section 4.1). Mitra and Baral’s system learns the commonsense knowledge using ILP which requires manual annotation of the examples. In contrast, (reusable) commonsense knowledge has to be hand-coded in the SQuARE system. In that sense, one could argue that Mitra and Baral’s approach is better as commonsense knowledge is produced automatically. In their case, however, generation of knowledge from the scenario expressed in natural language *also* requires manual annotation. That part is completely automated in our approach, thanks to our partial matching algorithm and VerbNet. Thus, the SQuARE system further advances state of the art in automated reasoning based approaches for QA.

With respect to commonsense knowledge, we take a different approach: we want to mimic a human in that all commonsense knowledge that an average human possesses should be coded in ASP using predefined simple generic predicates (such as `property`, `cause`, `after`, etc.) and made available to the question answering system. Use of generic predicates simplifies the development of commonsense knowledge, making it reusable and resulting in a scalable QA system. Of course, the amount of knowledge to be coded may be daunting, but it can be constructed incrementally on an as-needed basis. The amount of work would not be any less than annotating examples and generating commonsense knowledge automatically using a machine learning algorithm such as ILP. There may be additional work needed as the commonsense knowledge learned using ILP may have to be manually reviewed for errors and omissions and then manually fixed. Also, the generated commonsense knowledge should follow a standard naming convention for predicates, as we discussed earlier, something that is harder to realize in the ILP approach. If standardized conventions are not followed, the generated commonsense knowledge cannot be of any use to other applications without re-annotation and re-training.

An action language based QA methodology using VerbNet has been developed by Lierler et al [10]. The project aims to extend frame semantics with ALM, *an action language* [10], to provide interpretable semantic annotations. Unlike SQuARE, it is not an end-to-end automated QA system.

Cyc [8] is one of the earliest AI projects that aims to capture commonsense knowledge and reason over it. In Cyc, the acquired commonsense knowledge about the world is represented in the form of a vast collection of ontologies that consist of basic concepts and implicit rules about how the world works. To solve an inference problem, Cyc uses multiple reasoning agents in collaboration and these agents rely on more than 1000 pre-modeled heuristic modules. This is one of the challenges of Cyc as it is very hard to decide which model to apply and when, and which heuristic to use. A commonsense reasoning system should be designed in a simpler way, otherwise we may have acquired all the individual pieces of knowledge to answer and query, but may not be able to compose the answer from the pieces. For this reason, SQuARE uses minimal number of generic predicates and simple ASP-based reasoning rules (default reasoning is used as much as possible). Cyc did not support natural language question answering, however, recently some efforts have been initiated in that direction.

7 Discussion

Our goal is to create a semantic-based textual QA framework that mimics the way humans answer questions. Humans understand a passage semantics and use commonsense knowledge to logically reason to find answers. We believe that this process is the most effective way to develop an automated QA system. Intelligent behavior in humans is due to both learning and reasoning. At present, machine learning dominates AI. Machine learning is useful for many tasks (e.g., in parsing English sentences, detecting sentiment, etc.), however, it is not sufficient by itself for completely automating tasks that require intelligence. For automating intelligent behavior we need automated commonsense reasoning as well. The SQuARE system is a step in that direction. Note that it does not require any manual coding other than providing commonsense knowledge.

The SQuARE system has many advantages over a machine learning based QA systems. It can answer

questions without training. For example, none of the ML-based systems discussed in Table 2 will be able to answer the question - “*Is milk a liquid?*” without appropriate training data. SQuARE can answer such questions based on its commonsense knowledge, just like a human. SQuARE is also capable of providing explanation for each computed answer. Providing explanation shows a system’s true “understanding” of the matter, which is a necessary feature of a truly intelligent system. Interpretability makes SQuARE transparent, that means the system can be well understood, debugged, and improved. Also, our approach is incremental in nature. The system can be easily expanded to cover more NLU tasks without hampering the earlier accuracy. On the contrary, expanding the capabilities of a machine learning system often requires hyper-parameter tuning, which often results in reduced accuracy.

8 Future Work and Conclusions

This paper proposes a semantic based general question answering framework that uses commonsense knowledge to generate an answer with proper justification. This robust framework for automated QA preserves scalability, interpretability and explainability. We also demonstrated an end-to-end QA system called SQuARE that uses this framework to convert textual knowledge to semantic representation and reason over it with goal-directed ASP to answer questions posed in English. SQuARE relies on the s(CASP) goal-directed system and produced 100% accuracy on five bAbI tasks on which it was tested. The SQuARE system is a step towards general QA systems that can answer natural language questions against any text, automated conversational AI systems, etc.

As part of our future work, we plan to complete the remaining 15 bAbI tasks. The bAbI suite has many other more advanced tasks as well (*understanding children’s books, understanding movie dialogs*, etc.). We plan to extend SQuARE to implement those tasks too. Eventually, our goal is to develop a conversational AI system based on automated commonsense reasoning (using the goal-directed s(CASP) system) that can “converse” with a person based on “truly understanding” that person’s dialogs. We also plan to automate the commonsense knowledge generation process using state-of-the-art ILP systems [22].

References

- [1] Joaquin Arias, Manuel Carro, Elmer Salazar, Kyle Marple & Gopal Gupta (2018): *Constraint answer set programming without grounding*. *TPLP* 18(3-4), pp. 337–354, doi:10.1017/S1471068418000285.
- [2] Kinjal Basu, Farhad Shakerin & Gopal Gupta (2020): *AQuA: ASP-Based Visual Question Answering*. In: *Practical Aspects of Declarative Languages*, Springer International Publishing, Cham, pp. 57–72, doi:10.1007/978-3-030-39197-3_4.
- [3] Michael Gelfond & Yulia Kahl (2014): *Knowledge representation, reasoning, and the design of intelligent agents: The answer-set programming approach*. Cambridge University Press, doi:10.1017/CBO9781139342124.
- [4] Michael Gelfond & Vladimir Lifschitz (1988): *The stable model semantics for logic programming*. In: *ICLP/SLP*, 88, pp. 1070–1080.
- [5] Sepp Hochreiter & Jürgen Schmidhuber (1997): *Long short-term memory*. *Neural computation* 9(8), pp. 1735–1780, doi:10.1162/neco.1997.9.8.1735.
- [6] Karin Kipper, Anna Korhonen, Neville Ryant & Martha Palmer (2008): *A large-scale classification of English verbs*. *Language Resources and Evaluation* 42(1), pp. 21–40, doi:10.1007/s10579-007-9048-2.
- [7] Ankit Kumar et al. (2016): *Ask me anything: Dynamic memory networks for natural language processing*. In: *ICML*, pp. 1378–1387, doi:10.1109/ASYU48272.2019.8946411.
- [8] Douglas B Lenat (1995): *CYC: A large-scale investment in knowledge infrastructure*. *Communications of the ACM* 38(11), pp. 33–38, doi:10.1145/219717.219745.
- [9] Beth Levin (1993): *English verb classes and alternations: A preliminary investigation*. U. Chicago Press, doi:10.1075/fol.2.1.16noe.

- [10] Yuliya Lierler, Daniela Incezan & Michael Gelfond (2017): *Action languages and question answering*. In: *IWCS 2017 - 12th International Conference on Computational Semantics - Short papers*.
- [11] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard & David McClosky (2014): *The Stanford CoreNLP NLP Toolkit*. In: *ACL System Demonstrations*, pp. 55–60, doi:10.3115/v1/P14-5010.
- [12] Mitch Marcus, Beatrice Santorini & Mary Ann Marcinkiewicz (1993): *Building a Large Annotated Corpus of English: The Penn Treebank*. *Computational Linguistics* 19(2), pp. 313–330, doi:10.21236/ada273556.
- [13] Kyle Marple, Elmer Salazar & Gopal Gupta (2017): *Computing stable models of normal logic programs without grounding*. *arXiv preprint arXiv:1709.00501*.
- [14] Arindam Mitra & Chitta Baral (2016): *Addressing a question answering challenge by combining statistical methods with inductive rule learning and reasoning*. In: *Proc. AAAI*.
- [15] Arindam Mitra & Chitta Baral (2018): *Incremental and iterative learning of answer set programs from mutually distinct examples*. *Theory and Practice of Logic Programming* 18(3-4), pp. 623–637, doi:10.1017/S1471068418000248.
- [16] D. Pendharkar & G. Gupta (2019): *An ASP Based Approach to Answering Questions for Natural Language Text*. In: *International Symposium on PADL*, Springer, pp. 46–63, doi:10.1007/978-3-030-05998-9_4.
- [17] Peng Qi, Timothy Dozat, Yuhao Zhang & Christopher D Manning (2019): *Universal dependency parsing from scratch*. *arXiv preprint arXiv:1901.10457*, doi:10.18653/v1/K18-2016.
- [18] Pranav Rajpurkar, Robin Jia & Percy Liang (2018): *Know what you don't know: Unanswerable questions for SQuAD*. *arXiv preprint arXiv:1806.03822*, doi:10.18653/v1/P18-2124.
- [19] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev & Percy Liang (2016): *SQuAD: 100,000+ Questions for Machine Comprehension of Text*. *arXiv preprint arXiv:1606.05250*, doi:10.18653/v1/D16-1264.
- [20] Matthew Richardson, Christopher JC Burges & Erin Renshaw (2013): *MCTest: A Challenge Dataset for the Open-Domain Machine Comprehension of Text*. In: *Proceedings of the 2013 EMNLP*, pp. 193–203.
- [21] David A Schmidt (1986): *Denotational semantics: a methodology for language development*, William C. Brown Publishers, Dubuque, IA, USA.
- [22] Farhad Shakerin, Elmer Salazar & Gopal Gupta (2017): *A new algorithm to automate inductive learning of default theories*. *Theory and Practice of Logic Programming* 17(5-6), pp. 1010–1026, doi:10.1017/s1471068417000333.
- [23] Jason Weston et al. (2015): *Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks*. *arXiv preprint arXiv:1502.05698*.
- [24] Yi Yang, Wen-tau Yih & Christopher Meek (2015): *WikiQA: A Challenge Dataset for Open-Domain Question Answering*. In: *Proceedings of the 2015 EMNLP, ACL, Lisbon, Portugal*, pp. 2013–2018, doi:10.18653/v1/D15-1237. Available at <https://www.aclweb.org/anthology/D15-1237>.