

Conversational Neuro-Symbolic Commonsense Reasoning

Forough Arabshahi^{1*}, Jennifer Lee¹, Mikayla Gawarecki²
Kathryn Mazaitis², Amos Azaria³, Tom Mitchell²

¹Facebook,

²Carnegie Mellon University,

³Ariel University

{forough, jenniferlee98}@fb.com, {mgawarec, krivard}@cs.cmu.edu,
amos.azaria@ariel.ac.il, tom.mitchell@cs.cmu.edu

Abstract

In order for conversational AI systems to hold more natural and broad-ranging conversations, they will require much more commonsense, including the ability to identify unstated *presumptions* of their conversational partners. For example, in the command “If it snows at night then wake me up early because I don’t want to be late for work” the speaker relies on commonsense reasoning of the listener to infer the implicit presumption that they wish to be woken only if it snows enough to cause traffic slowdowns. We consider here the problem of understanding such imprecisely stated natural language commands given in the form of **if-(state), then-(action), because-(goal)** statements. More precisely, we consider the problem of identifying the unstated *presumptions* of the speaker that allow the requested action to achieve the desired goal from the given state (perhaps elaborated by making the implicit presumptions explicit). We release a benchmark data set for this task, collected from humans and annotated with commonsense presumptions. We present a neuro-symbolic theorem prover that extracts multi-hop reasoning chains, and apply it to this problem. Furthermore, to accommodate the reality that current AI commonsense systems lack full coverage, we also present an interactive conversational framework built on our neuro-symbolic system, that conversationally evokes commonsense knowledge from humans to complete its reasoning chains.

Introduction

Despite the remarkable success of artificial intelligence (AI) and machine learning in the last few decades, commonsense reasoning remains an unsolved problem at the heart of AI (Levesque, Davis, and Morgenstern 2012; Davis and Marcus 2015; Sakaguchi et al. 2020). Common sense allows us humans to engage in conversations with one another and to convey our thoughts efficiently, without the need to specify much detail (Grice 1975). For example, if Alice asks Bob to “wake her up early whenever it snows at night” so that she can get to work on time, Alice assumes that Bob will wake her up only if it snows enough to cause traffic slowdowns, and only if it is a working day. Alice does not explicitly state these conditions since Bob makes such *presumptions* without much effort thanks to his common sense.

*work done when FA and JL were at Carnegie Mellon University. Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

A study, in which we collected such **if(state)-then(action)** commands from human subjects, revealed that humans often under-specify conditions in their statements; perhaps because they are used to speaking with other humans who possess the common sense needed to infer their more specific intent by making presumptions about their statement.

The inability to make these presumptions makes it challenging for computers to engage in natural sounding conversations with humans. While conversational AI systems such as Siri, Alexa, and others are entering our daily lives, their conversations with us humans remains limited to a set of pre-programmed tasks. We propose that handling unseen tasks requires conversational agents to develop common sense.

Therefore, we propose a new commonsense reasoning benchmark for conversational agents where the task is to infer *commonsense presumptions* in commands of the form “If state holds Then perform action Because I want to achieve goal.” The **if-(state), then-(action)** clause arises when humans instruct new conditional tasks to conversational agents (Azaria, Krishnamurthy, and Mitchell 2016; Labutov, Srivastava, and Mitchell 2018). The reason for including the **because-(goal)** clause in the commands is that some presumptions are ambiguous without knowing the user’s purpose, or goal. For instance, if Alice’s goal in the previous example was to see snow for the first time, Bob would have presumed that even a snow flurry would be excuse enough to wake her up. Since humans frequently omit details when stating such commands, a computer possessing common sense should be able to infer the hidden *presumptions*; that is, the additional unstated conditions on the If and/or Then portion of the command. Please refer to Tab. 1 for some examples. In this paper, in addition to the proposal of this novel task and the release of a new dataset to study it, we propose a novel initial approach that infers such missing presumptions, by extracting a chain of reasoning that shows how the commanded action will achieve the desired goal when the state holds. Whenever any additional reasoning steps appear in this reasoning chain, they are output by our system as assumed implicit presumptions associated with the command. For our reasoning method we propose a neuro-symbolic interactive, conversational approach, in which the computer combines its own commonsense knowledge with conversationally evoked knowledge provided by a human user. The reasoning chain

Domain	if clause	then clause	because clause	Example	Annotation: Commonsense Presumptions	Count
Restricted domain	state	action	goal	If it's going to rain in the afternoon (↑) then remind me to bring an umbrella (↑) because I want to remain dry	(8, and I am outside) (15, before I leave the house)	76
Restricted domain	state	action	anti-goal	If I have an upcoming bill payment (↑) then remind me to pay it (↑) because I don't want to pay a late fee	(7, in the next few days) (13, before the bill payment deadline)	3
Restricted domain	state	action	modifier	If my flight (↑) is from 2am to 4am then book me a supershuttle (↑) because it will be difficult to find ubers.	(3, take off time) (13, for 2 hours before my flight take off time)	2
Restricted domain	state	action	conjunction	If I receive emails about sales on basketball shoes (↑) then let me know (↑) because I need them and I want to save money.	(9, my size) (13, there is a sale)	2
SUM						83
Everyday domain	state	action	goal	If there is an upcoming election (↑) then remind me to register (↑) and vote (↑) because I want my voice to be heard.	(6, in the next few months) (6, and I am eligible to vote) (11, to vote) (13, in the election)	55
Everyday domain	state	action	anti-goal	If it's been two weeks since my last call with my mentee and I don't have an upcoming appointment with her (↑) then remind me to send her an email (↑) because we forgot to schedule our next chat	(21, in the next few days) (29, to schedule our next appointment)	4
Everyday domain	state	action	modifier	If I have difficulty sleeping (↑) then play a lullaby because it soothes me.	(5, at night)	12
Everyday domain	state	action	conjunction	If the power goes out (↑) then when it comes back on remind me to restart the house furnace because it doesn't come back on by itself and I want to stay warm	(5, in the Winter)	6
SUM						77

Table 1: Statistics of if-(state), then-(action), because-(goal) commands collected from a pool of human subjects. The table shows four distinct types of *because*-clauses we found, the count of commands of each type, examples of each and their corresponding commonsense presumption annotations. *Restricted domain* includes commands whose *state* is limited to checking email, calendar, maps, alarms, and weather. *Everyday domain* includes commands concerning more general day-to-day activities. Annotations are tuples of (index, presumption) where index shows the starting word index of where the missing presumption should be in the command, highlighted with a red arrow. Index starts at 0 and is calculated for the original command.

is extracted using our neuro-symbolic theorem prover that learns sub-symbolic representations (embeddings) for logical statements, making it robust to variations of natural language encountered in a conversational interaction setting.

Contributions This paper presents three main contributions. 1) We propose a benchmark task for commonsense reasoning in conversational agents and release a data set containing **if-(state), then-(action), because-(goal)** commands, annotated with commonsense presumptions. 2) We present CORGI (COMmonsense ReasoninG by Instruction), a system that performs soft logical inference. CORGI uses our proposed neuro-symbolic theorem prover and applies it to extract a multi-hop reasoning chain that reveals commonsense presumptions. 3) We equip CORGI with a conversational interaction mechanism that enables it to collect just-in-time commonsense knowledge from humans. Our user-study shows (a) the plausibility of relying on humans to evoke commonsense knowledge and (b) the effectiveness of our theorem prover, enabling us to extract reasoning chains for up to 45% of the

studied tasks¹.

Related Work

The literature on commonsense reasoning dates back to the very beginning of the field of AI (Winograd 1972; Mueller 2014; Davis and Marcus 2015) and is studied in several contexts. One aspect focuses on building a large knowledge base (KB) of commonsense facts. Projects like CYC (Lenat et al. 1990), ConceptNet (Liu and Singh 2004; Havasi, Speer, and Alonso 2007; Speer, Chin, and Havasi 2017) and ATOMIC (Sap et al. 2019; Rashkin et al. 2018) are examples of such KBs (see (Davis and Marcus 2015) for a comprehensive list). Recently, Bosselut et al. (2019) proposed COMET, a neural knowledge graph that generates knowledge tuples by learning on examples of structured knowledge. These KBs provide background knowledge for tasks that require common sense. However, it is known that knowledge bases are incomplete,

¹The code and data are available here: <https://github.com/ForoughA/CORGI>

and most have ambiguities and inconsistencies (Davis and Marcus 2015) that must be clarified for particular reasoning tasks. Therefore, we argue that reasoning engines can benefit greatly from a *conversational interaction strategy* to ask humans about their missing or inconsistent knowledge. Closest in nature to this proposal is the work by Hixon, Clark, and Hajishirzi (2015) on relation extraction through conversation for question answering and Wu et al. (2018)’s system that learns to form simple concepts through interactive dialogue with a user. The advent of intelligent agents and advancements in natural language processing have given learning from conversational interactions a good momentum in the last few years (Azaria, Krishnamurthy, and Mitchell 2016; Labutov, Srivastava, and Mitchell 2018; Srivastava 2018; Goldwasser and Roth 2014; Christmann et al. 2019; Guo et al. 2018; Li et al. 2018, 2017; Li, Azaria, and Myers 2017).

A current challenge in commonsense reasoning is lack of benchmarks (Davis and Marcus 2015). Benchmark tasks in commonsense reasoning include the Winograd Schema Challenge (WSC) (Levesque, Davis, and Morgenstern 2012), its variations (Kocijan et al. 2020), and its recently scaled up counterpart, Winogrande (Sakaguchi et al. 2020); ROCStories (Mostafazadeh et al. 2017), COPA (Roemmele, Bejan, and Gordon 2011), Triangle COPA (Maslan, Roemmele, and Gordon 2015), and *ART* (Bhagavatula et al. 2020), where the task is to choose a plausible outcome, cause or explanation for an input scenario; and the TimeTravel benchmark (Qin et al. 2019) where the task to revise a story to make it compatible with a given counterfactual event. Other than TimeTravel, most of these benchmarks have a multiple choice design format. However, in the real world the computer is usually not given multiple choice questions. None of these benchmarks targets the extraction of unspoken details in a natural language statement, which is a challenging task for computers known since the 1970’s (Grice 1975). Note that inferring commonsense presumptions is different from intent understanding (Janíček 2010; Tur and De Mori 2011) where the goal is to understand the intent of a speaker when they say, e.g., “pick up the mug”. It is also different from implicature and presupposition (Sbisà 1999; Simons 2013; Sakama and Inoue 2016) which are concerned with what can be presupposed or implicated by a text.

CORGI has a neuro-symbolic logic theorem prover. Neuro-symbolic systems are hybrid models that leverage the robustness of connectionist methods and the soundness of symbolic reasoning to effectively integrate learning and reasoning (Garcez et al. 2015; Besold et al. 2017). They have shown promise in different areas of logical reasoning ranging from classical logic to propositional logic, probabilistic logic, abductive logic, and inductive logic (Mao et al. 2019; Manhaeve et al. 2018; Dong et al. 2019; Marra et al. 2019; Zhou 2019; Evans and Grefenstette 2018). To the best of our knowledge, neuro-symbolic solutions for commonsense reasoning have not been proposed before. Examples of commonsense reasoning engines are: AnalogySpace (Speer, Havasi, and Lieberman 2008; Havasi et al. 2009) that uses dimensionality reduction and Mueller (2014) that uses the event calculus formal language. TensorLog (Cohen 2016) converts a first-order logical database into a factor graph and proposes a

differentiable strategy for belief propagation over the graph. DeepProbLog (Manhaeve et al. 2018) developed a probabilistic logic programming language that is suitable for applications containing categorical variables. Contrary to our approach, both these methods do not learn embeddings for logical rules that are needed to make CORGI robust to natural language variations. Therefore, we propose an end-to-end differentiable solution that uses a Prolog (Colmerauer 1990) proof trace to learn rule embeddings from data. Our proposal is closest to the neural programmer interpreter (Reed and De Freitas 2015) that uses the trace of algorithms such as addition and sort to learn their execution. The use of Prolog for performing multi-hop logical reasoning has been studied in Rocktäschel and Riedel (2017) and Weber et al. (2019). These methods perform Inductive Logic Programming to learn rules from data, and are not applicable to our problem. DeepLogic (Cingillioglu and Russo 2018), Rocktäschel et al. (2014), and Wang and Cohen (2016) also learn representations for logical rules using neural networks. Very recently, transformers were used for temporal logic (Finkbeiner et al. 2020) and to do multi-hop reasoning (Clark, Tafjord, and Richardson 2020) using logical facts and rules stated in natural language. A purely connectionist approach to reasoning suffers from some limitations. For example, the input token size limit of transformers restricts Clark, Tafjord, and Richardson (2020) to small knowledge bases. Moreover, generalizing to arbitrary number of variables or an arbitrary inference depth is not trivial for them. Since symbolic reasoning can inherently handle all these challenges, a hybrid approach to reasoning takes the burden of handling them off of the neural component.

Proposed Commonsense Reasoning Benchmark

The benchmark task that we propose in this work is that of uncovering hidden *commonsense presumptions* given commands that follow the general format “if ⟨state holds⟩ then ⟨perform action⟩ because ⟨I want to achieve goal⟩”. We refer to these as if-then-because commands. We refer to the *if*-clause as the *state*, the *then*-clause as the *action* and the *because*-clause as the *goal*. These natural language commands were collected from a pool of human subjects (more details in the Appendix). The data is annotated with unspoken commonsense presumptions by a team of annotators. Tab. 1 shows the statistics of the data and annotated examples from the data. We collected two sets of if-then-because commands. The first set contains 83 commands targeted at a *state* that can be observed by a computer/mobile phone (e.g. checking emails, calendar, maps, alarms, and weather). The second set contains 77 commands whose *state* is about day-to-day events and activities. 81% of the commands over both sets qualify as “if ⟨state⟩ then ⟨action⟩ because ⟨goal⟩”. The remaining 19% differ in the categorization of the *because*-clause (see Tab. 1); common alternate clause types included anti-goals (“...because I don’t want to be late”), modifications of the state or action (“... because it will be difficult to find an Uber”), or conjunctions including at least one non-goal type. Note that we did not instruct the subjects to give us data from these categories, rather we uncovered them

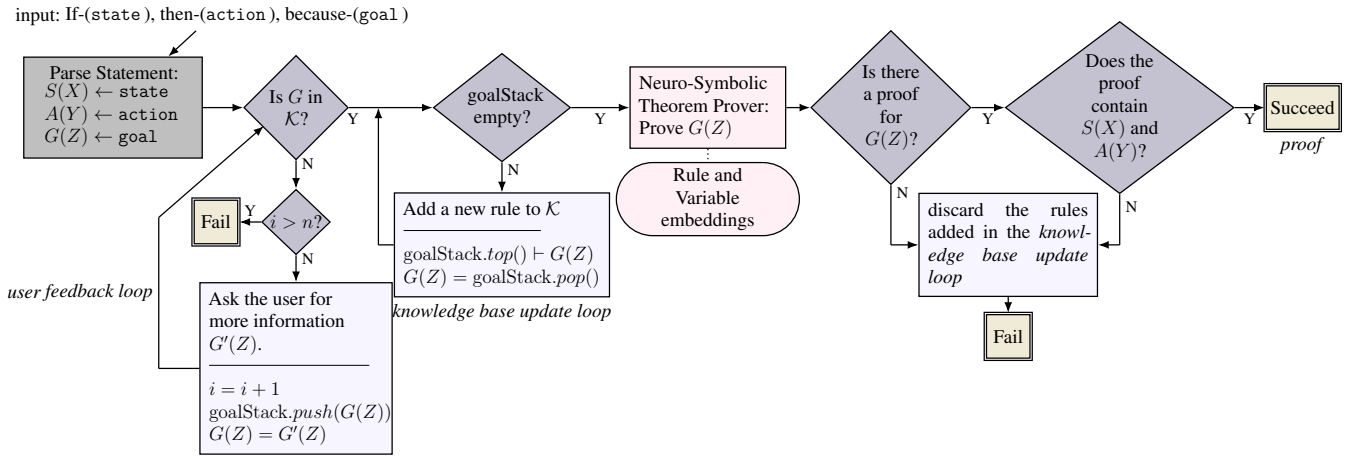


Figure 1: CORGI’s flowchart. The input is an if-then-because command e.g., “if it snows tonight then wake me up early because I want to get to work on time”. The input is parsed into its logical form representation (for this example, $S(X)$ = weather(snow, Precipitation)). If CORGI succeeds, it outputs a proof tree for the because-clause or goal (parsed into $G(Z)=get(i,work,on_time)$). The output proof tree contains commonsense presumptions for the input statement (Fig 2 shows an example). If the predicate G does not exist in the knowledge base, \mathcal{K} , (Is G in \mathcal{K} ?), we have missing knowledge and cannot find a proof. Therefore, we extract it from a human in the *user feedback loop*. At the heart of CORGI is a neuro-symbolic theorem prover that learns rule and variable embeddings to perform a proof (Appendix). *goalStack* and the loop variable i are initialized to empty and 0 respectively, and $n = 3$. *italic text* in the figure represents descriptions that are referred to in the main text.

after data collection. Also, commonsense benchmarks such as the Winograd Schema Challenge (Levesque, Davis, and Morgenstern 2012) included a similar number of examples (100) when first introduced (Kocijan et al. 2020).

Lastly, the if-then-because commands given by humans can be categorized into several different logic templates. The discovered logic templates are given in the Appendix ². Our neuro-symbolic theorem prover uses a general reasoning strategy that can address all reasoning templates. However, in an extended discussion in the Appendix, we explain how a reasoning system, including ours, could potentially benefit from these logic templates.

Method

Background and notation The system’s commonsense knowledge is a KB, denoted \mathcal{K} , programmed in a Prolog-like syntax. We have developed a modified version of Prolog, which has been augmented to support several special features (types, soft-matched predicates and atoms, etc). Prolog (Colmerauer 1990) is a declarative logic programming language that consists of a set of predicates whose arguments are atoms, variables or predicates. A predicate is defined by a set of rules (Head $:-$ Body.) and facts (Head.), where Head is a predicate, Body is a conjunction of predicates, and $:-$ is logical implication. We use the notation $S(X)$, $A(Y)$ and $G(Z)$ to represent the logical form of the state, action and goal, respectively where S , A and G are predicate names and X , Y and Z indicate the list of arguments of each predicate. For example, for goal = “I want to get to work on time”, we have $G(Z) = get(i, work, on_time)$. Prolog can be used to logically “prove” a query (e.g., to prove $G(Z)$ from $S(X)$, $G(Z)$

and appropriate commonsense knowledge (see the Appendix - Prolog Background)).

CORGI: Commonsense Reasoning by Instruction

CORGI takes as input a natural language command of the form “if (state) then (action) because (goal)” and infers commonsense presumptions by extracting a chain of commonsense knowledge that explains how the commanded action achieves the goal when the state holds. For example from a high level, for the command in Fig. 2 CORGI outputs (1) if it snows more than two inches, then there will be traffic, (2) if there is traffic, then my commute time to work increases, (3) if my commute time to work increases then I need to leave the house earlier to ensure I get to work on time (4) if I wake up earlier then I will leave the house earlier. Formally, this reasoning chain is a proof tree (proof trace) shown in Fig.2. As shown, the proof tree includes the commonsense presumptions.

CORGI’s architecture is depicted in Figure 1. In the first step, the if-then-because command goes through a parser that extracts the state, action and goal from it and converts them to their logical form representations $S(X)$, $A(Y)$ and $G(Z)$, respectively. For example, the action “wake me up early” is converted to wake(me, early). The parser is presented in the Appendix (Sec. Parsing).

The proof trace is obtained by finding a proof for $G(Z)$, using \mathcal{K} and the context of the input if-then-because command. In other words, $S(X) \cap A(Y) \cap \mathcal{K} \rightarrow G(Z)$. One challenge is that even the largest knowledge bases gathered to date are incomplete, making it virtually infeasible to prove an arbitrary input $G(Z)$. Therefore, CORGI is equipped with a conversational interaction strategy, which enables it to prove

²The appendix is available at <https://arxiv.org/abs/2006.10022>

a query by combining its own commonsense knowledge with conversationally evoked knowledge provided by a human user in response to a question from CORGI (*user feedback loop* in Fig.1). There are 4 possible scenarios that can occur when CORGI asks such questions:

- Ⓐ The user understands the question, but does not know the answer.
- Ⓑ The user misunderstands the question and responds with an undesired answer.
- Ⓒ The user understands the question and provides a correct answer, but the system fails to understand the user due to:
 - Ⓒ.1 limitations of natural language understanding.
 - Ⓒ.2 variations in natural language, which result in misalignment of the data schema in the knowledge base and the data schema in the user’s mind.
- Ⓓ The user understands the question and provides the correct answer and the system successfully parses and understands it.

CORGI’s different components are designed such that they address the above challenges, as explained below. Since our benchmark data set deals with day-to-day activities, it is unlikely for scenario Ⓐ to occur. If the task required more specific domain knowledge, Ⓐ could have been addressed by choosing a pool of domain experts. Scenario Ⓑ is addressed by asking informative questions from users. Scenario Ⓒ.1 is addressed by trying to extract small chunks of knowledge from the users piece-by-piece. Specifically, the choice of what to ask the user in the *user feedback loop* is deterministically computed from the user’s `goal`. The first step is to ask how to achieve the user’s stated `goal`, and CORGI expects an answer that gives a `sub-goal`. In the next step, CORGI asks how to achieve the `sub-goal` the user just mentioned. The reason for this piece-by-piece knowledge extraction is to ensure that the language understanding component can correctly parse the user’s response. CORGI then adds the extracted knowledge from the user to \mathcal{K} in the *knowledge update loop* shown in Fig.1. Missing knowledge outside this `goal/sub-goal` path is not handled, although it is an interesting future direction. Moreover, the model is user specific and the knowledge extracted from different users are not shared among them. Sharing knowledge raises interesting privacy issues and requires handling personalized conflicts and falls out of the scope of our current study.

Scenario Ⓒ.2, caused by the variations of natural language, results in semantically similar statements to get mapped into different logical forms, which is unwanted. For example, “make sure I am awake early morning” vs. “wake me up early morning” will be parsed into different logical forms `awake(i,early_morning)` and `wake(me, early_morning)`, respectively although they are semantically similar. This mismatch prevents a logical proof from succeeding since the proof strategy relies on exact match in the unification operation (see Appendix). This is addressed by our neuro-symbolic theorem prover (Fig.1) that learns vector representations (embeddings) for logical rules and variables and uses them to perform a logical proof through soft unification. If the theorem prover can prove the user’s `goal`, $G(Z)$, CORGI outputs

the proof trace (Fig.2) returned by its theorem prover and succeeds. In the next section, we explain our theorem prover in detail. We revisit scenarios Ⓐ – Ⓓ in detail in the discussion section and show real examples from our user study.

Neuro-Symbolic Theorem Proving

Our Neuro-Symbolic theorem prover is a neural modification of *backward chaining* and uses the vector similarity between rule and variable embeddings for unification. In order to learn these embeddings, our theorem prover learns a general proving strategy by training on proof traces of successful proofs. From a high level, for a given query our model maximizes the probability of choosing the correct rule to pick in each step of the backward chaining algorithm. This proposal is an adaptation of Reed et al.’s Neural Programmer-Interpreter (Reed and De Freitas 2015) that learns to execute algorithms such as addition and sort, by training on their execution trace.

In what follows, we represent scalars with lowercase letters, vectors with bold lowercase letters and matrices with bold uppercase letters. $\mathbf{M}^{rule} \in \mathbb{R}^{n_1 \times m_1}$ denotes the embedding matrix for the rules and facts, where n_1 is the number of rules and facts and m_1 is the embedding dimension. $\mathbf{M}^{var} \in \mathbb{R}^{n_2 \times m_2}$ denotes the variable embedding matrix, where n_2 is the number of all the atoms and variables in the knowledge base and m_2 is the variable embedding dimension. Our knowledge base is type coerced, therefore the variable names are associated with their types (e.g., `alarm(Person,Time)`)

Learning The model’s core consists of an LSTM network whose hidden state indicates the next rule in the proof trace and a proof termination probability, given a query as input. The model has a feed forward network that makes variable binding decisions. The model’s training is fully supervised by the proof trace of a query given in a depth-first-traversal order from left to right (Fig. 2). The trace is sequentially input to the model in the traversal order as explained in what follows. In step $t \in [0, T]$ of the proof, the model’s input is $\epsilon_t^{inp} = (\mathbf{q}_t, \mathbf{r}_t, (\mathbf{v}_t^1, \dots, \mathbf{v}_t^\ell))$ and T is the total number of proof steps. \mathbf{q}_t is the query’s embedding and is computed by feeding the predicate name of the query into a character RNN. \mathbf{r}_t is the concatenated embeddings of the rules in the parent and the left sister nodes in the proof trace, looked up from \mathbf{M}^{rule} . For example in Fig.2, \mathbf{q}_3 represents the node at proof step $t = 3$, \mathbf{r}_3 represents the rule highlighted in green (parent rule), and \mathbf{r}_4 represents the fact `alarm(i, 8)`. The reason for including the left sister node in \mathbf{r}_t is that the proof is conducted in a left-to-right depth first order. Therefore, the decision of what next rule to choose in each node is dependent on both the left sisters and the parent (e.g. the parent and the left sisters of the node at step $t = 8$ in Fig. 2 are the rules at nodes $t = 1$, $t = 2$, and $t = 6$, respectively). The arguments of the query are presented in $(\mathbf{v}_t^1, \dots, \mathbf{v}_t^\ell)$ where ℓ is the arity of the query predicate. For example, \mathbf{v}_3^1 in Fig 2 is the embedding of the variable `Person`. Each \mathbf{v}_t^i for $i \in [0, \ell]$, is looked up from the embedding matrix \mathbf{M}^{var} . The output of the model in step t is $\epsilon_t^{out} = (c_t, \mathbf{r}_{t+1}, (\mathbf{v}_{t+1}^1, \dots, \mathbf{v}_{t+1}^\ell))$

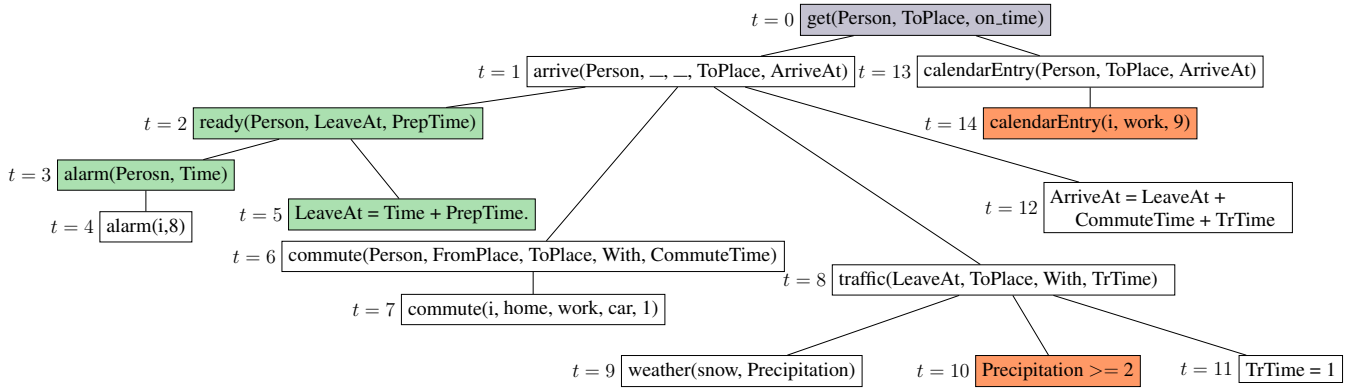


Figure 2: Sample proof tree for the *because*-clause of the statement: “If it snows tonight then wake me up early because I want to get to work on time”. Proof traversal is depth-first from left to right (t gives the order). Each node in the tree indicates a rule’s head, and its children indicate the rule’s body. For example, the nodes highlighted in green indicate the rule $\text{ready}(\text{Person}, \text{LeaveAt}, \text{PrepTime}) \text{ :- } \text{alarm}(\text{Person}, \text{Time}) \wedge \text{LeaveAt} = \text{Time} + \text{PrepTime}$. The goal we want to prove, $G(Z) = \text{get}(\text{Person}, \text{ToPlace}, \text{on_time})$, is in the tree’s root. If a proof is successful, the variables in $G(Z)$ get grounded (here Person and ToPlace are grounded to i and work, respectively). The highlighted orange nodes are the uncovered commonsense presumptions.

and is computed through the following equations

$$\mathbf{s}_t = f_{enc}(\mathbf{q}_t), \quad \mathbf{h}_t = f_{lstm}(\mathbf{s}_t, \mathbf{r}_t, \mathbf{h}_{t-1}), \quad (1)$$

$$c_t = f_{end}(\mathbf{h}_t), \quad \mathbf{r}_{t+1} = f_{rule}(\mathbf{h}_t), \quad \mathbf{v}_{t+1}^i = f_{var}(\mathbf{v}_t^i), \quad (2)$$

where \mathbf{v}_{t+1}^i is a probability vector over all the variables and atoms for the i^{th} argument, \mathbf{r}_{t+1} is a probability vector over all the rules and facts and c_t is a scalar probability of terminating the proof at step t . f_{enc} , f_{end} , f_{rule} and f_{var} are feed forward networks with two fully connected layers, and f_{lstm} is an LSTM network. The trainable parameters of the model are the parameters of the feed forward neural networks, the LSTM network, the character RNN that embeds \mathbf{q}_t and the rule and variable embedding matrices \mathbf{M}^{rule} and \mathbf{M}^{var} .

Our model is trained end-to-end. In order to train the model parameters and the embeddings, we maximize the log likelihood probability given below

$$\theta^* = \underset{\theta}{\text{argmax}} \sum_{\epsilon^{out}, \epsilon^{in}} \log(P(\epsilon^{out} | \epsilon^{in}; \theta)), \quad (3)$$

where the summation is over all the proof traces in the training set and θ is the trainable parameters of the model. We have

$$\log(P(\epsilon^{out} | \epsilon^{in}; \theta)) = \sum_{t=1}^T \log P(\epsilon_t^{out} | \epsilon_1^{in} \dots \epsilon_{t-1}^{in}; \theta), \quad (4)$$

$$\begin{aligned} \log P(\epsilon_t^{out} | \epsilon_1^{in} \dots \epsilon_{t-1}^{in}; \theta) &= \log P(\epsilon_t^{out} | \epsilon_{t-1}^{in}; \theta) \\ &= \log P(c_t | \mathbf{h}_t) + \\ &\quad \log P(\mathbf{r}_{t+1} | \mathbf{h}_t) + \\ &\quad \log \sum_i P(\mathbf{v}_{t+1}^i | \mathbf{v}_t^i). \end{aligned} \quad (5)$$

Where the probabilities in Equation (5) are given in Equations (2). The inference algorithm for proving is given in the Appendix, section Inference.

Experiment Design

The knowledge base, \mathcal{K} , used for all experiments is a small handcrafted set of commonsense knowledge that reflects the incompleteness of SOTA KBs. Some examples of our KB entries are available in the Appendix. \mathcal{K} includes general information about time, restricted-domains such as setting alarms and notifications, emails, and so on, as well as commonsense knowledge about day-to-day activities. \mathcal{K} contains a total of 228 facts and rules. Among these, there are 189 everyday-domain and 39 restricted domain facts and rules. We observed that most of the if-then-because commands require everyday-domain knowledge for reasoning, even if they are restricted-domain commands (see Table 3 for example).

Our Neuro-Symbolic theorem prover is trained on proof traces (proof trees similar to Fig. 2) collected by proving automatically generated queries to \mathcal{K} using sPyrolog³. \mathbf{M}^{rule} and \mathbf{M}^{var} are initialized randomly and with GloVe embeddings (Pennington, Socher, and Manning 2014), respectively, where $m_1 = 256$ and $m_2 = 300$. Since \mathcal{K} is type-coerced (e.g. Time, Location, . . .), initializing the variables with pre-trained word embeddings helps capture their semantics and improves the performance. The neural components of the theorem prover are implemented in PyTorch (Paszke et al. 2017) and the prover is built on top of sPyrolog.

User Study

In order to assess CORGI’s performance, we ran a user study. We selected 10 goal-type if-then-because commands from the dataset in Table 1 and used each as the prompt for a reasoning task. We had 28 participants in the study, 4 of which were experts closely familiar with CORGI and its capabilities. The rest were undergraduate and graduate students with the majority being in engineering or computer science fields and some that majored in business administration or psychology.

³<https://github.com/leonweber/spyrolog>

CORGI variations	Novice User	Expert User
No-feedback	0%	0%
Soft unification	15.61%	35.00%
Oracle unification	21.62%	45.71%

Table 2: percentage of successful reasoning tasks for different user types. In no-feedback, user responses are not considered in the proof attempt. In soft unification CORGI uses our proposed neuro-symbolic theorem prover. In the Oracle scenario, the theorem prover has access to oracle embeddings and soft unification is 100% accurate.

These users had never interacted with CORGI prior to the study (novice users). Each person was issued the 10 reasoning tasks, taking on average 20 minutes to complete all 10. Solving a reasoning task consists of participating in a dialog with CORGI as the system attempts to complete a proof for the goal of the current task; see sample dialogs in Tab. 3. The task succeeds if CORGI is able to use the answers provided by the participant to construct a reasoning chain (proof) leading from the goal to the state and action. We collected 469 dialogues in our study.

The user study was run with the architecture shown in Fig. 1. We used the participant responses from the study to run a few more experiments. We (1) Replace our theorem prover with an *oracle* prover that selects the optimal rule at each proof step in our inference Algorithm (Appendix) and (2) attempt to prove the goal without using any participant responses (*no-feedback*). Tab. 2 shows the success rate in each setting.

Discussion

In this section, we analyze the results from the study and provide examples of the 4 scenarios in Section that we encountered. As hypothesized, scenario \mathfrak{A} hardly occurred as the commands are about day-to-day activities that all users are familiar with. We did encounter scenario \mathfrak{B} , however. The study’s dialogs show that some users provided means of *sensing* the goal rather than the *cause* of the goal. For example, for the reasoning task “*If there are thunderstorms in the forecast within a few hours then remind me to close the windows because I want to keep my home dry*”, in response to the system’s prompt “*How do I know if ‘I keep my home dry’?*” a user responded “*if the floor is not wet*” as opposed to an answer such as “*if the windows are closed*”. Moreover, some users did not pay attention to the context of the reasoning task. For example, another user responded to the above prompt (same reasoning task) with “*if the temperature is above 80*”! Overall, we noticed that CORGI’s ability to successfully reason about an if-then-because statement was heavily dependent on whether the user knew how to give the system what it needed, and not necessarily what it asked for; see Table 3 for an example. As it can be seen in Table 2, expert users are able to more effectively provide answers that complete CORGI’s reasoning chain, likely because they know that regardless of what CORGI asks, the object of the dialog is to connect the because goal back to the knowledge base in some series of if-then rules (goal/sub-goal path in

Successful task
If it’s going to rain in the afternoon then remind me to bring an umbrella because I want to remain dry. <i>How do I know if “I remain dry”?</i>
If I have my umbrella. <i>How do I know if “I have my umbrella”?</i>
If you remind me to bring an umbrella. <i>Okay, I will perform “remind me to bring an umbrella” in order to achieve “I remain dry”.</i>
Failed task
If it’s going to rain in the afternoon then remind me to bring an umbrella because I want to remain dry. <i>How do I know if “I remain dry”?</i>
If I have my umbrella. <i>How do I know if “I have my umbrella”?</i>
If it’s in my office. <i>How do I know if “it’s in my office”?</i>
...

Table 3: Sample dialogs of 2 novice users in our study. CORGI’s responses are noted in italics.

Sec.). Therefore, one interesting future direction is to develop a dynamic context-dependent Natural Language Generation method for asking more effective questions.

We would like to emphasize that although it seems to us, humans, that the previous example requires very simple background knowledge that likely exists in SOTA large commonsense knowledge graphs such as ConceptNet⁴, ATOMIC⁵ or COMET (Bosselut et al. 2019), this is not the case (verifiable by querying them online). For example, for queries such as “*the windows are closed*”, COMET-ConceptNet generative model⁶ returns knowledge about blocking the sun, and COMET-ATOMIC generative model⁷ returns knowledge about keeping the house warm or avoiding to get hot; which while being correct, is not applicable in this context. For “*my home is dry*”, both COMET-ConceptNet and COMET-ATOMIC generative models return knowledge about house cleaning or house comfort. On the other hand, the fact that 40% of the novice users in our study were able to help CORGI reason about this example with responses such as “*If I close the windows*” to CORGI’s prompt, is an interesting result. This tells us that conversational interactions with humans could pave the way for commonsense reasoning and enable computers to extract just-in-time commonsense knowledge, which would likely either not exist in large knowledge bases or be irrelevant in the context of the particular reasoning task. Lastly, we re-iterate that as conversational agents (such as Siri and Alexa) enter people’s lives, leveraging conversational interactions for learning has become a more realistic opportunity than ever before.

In order to address scenario $\mathfrak{C}.1$, the conversational

⁴<http://conceptnet.io/>

⁵https://mosaickg.apps.allenai.org/kg_atomic

⁶https://mosaickg.apps.allenai.org/comet_conceptnet

⁷https://mosaickg.apps.allenai.org/comet_atomic

prompts of CORGI ask for specific small pieces of knowledge that can be easily parsed into a predicate and a set of arguments. However, some users in our study tried to provide additional details, which challenged CORGI’s natural language understanding. For example, for the reasoning task “If I receive an email about water shut off then remind me about it a day before because I want to make sure I have access to water when I need it.”, in response to the system’s prompt “How do I know if ‘I have access to water when I need it.’?” one user responded “If I am reminded about a water shut off I can fill bottles”. This is a successful knowledge transfer. However, the parser expected this to be broken down into two steps. If this user responded to the prompt with “If I fill bottles” first, CORGI would have asked “How do I know if ‘I fill bottles’?” and if the user then responded “if I am reminded about a water shut off” CORGI would have succeeded. The success from such conversational interactions are not reflected in the overall performance mainly due to the limitations of natural language understanding.

Table 2 evaluates the effectiveness of conversational interactions for proving compared to the no-feedback model. The 0% success rate there reflects the incompleteness of \mathcal{K} . The improvement in task success rate between the no-feedback case and the other rows indicates that when it is possible for users to contribute useful common-sense knowledge to the system, performance improves. The users contributed a total number of 96 rules to our knowledge base, 31 of which were unique rules. Scenario C.2 occurs when there is variation in the user’s natural language statement and is addressed with our neuro-symbolic theorem prover. Rows 2-3 in Table 2 evaluate our theorem prover (*soft unification*). Having access to the optimal rule for unification does still better, but the task success rate is not 100%, mainly due to the limitations of natural language understanding explained earlier.

Conclusions

In this paper, we introduced a benchmark task for common-sense reasoning that aims at uncovering unspoken intents that humans can easily uncover in a given statement by making presumptions supported by their common sense. In order to solve this task, we propose CORGI (COmmon-sense Reasoning by Instruction), a neuro-symbolic theorem prover that performs commonsense reasoning by initiating a conversation with a user. CORGI has access to a small knowledge base of commonsense facts and completes it as she interacts with the user. We further conduct a user study that indicates the possibility of using conversational interactions with humans for evoking commonsense knowledge and verifies the effectiveness of our proposed theorem prover.

Acknowledgements

This work was supported in part by AFOSR under research contract FA9550201.

References

Azaria, A.; Krishnamurthy, J.; and Mitchell, T. M. 2016. Instructable intelligent personal agent. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Besold, T. R.; Garcez, A. d.; Bader, S.; Bowman, H.; Domingos, P.; Hitzler, P.; Kühnberger, K.-U.; Lamb, L. C.; Lowd, D.; Lima, P. M. V.; et al. 2017. Neural-symbolic learning and reasoning: A survey and interpretation. *arXiv preprint arXiv:1711.03902*.

Bhagavatula, C.; Bras, R. L.; Malaviya, C.; Sakaguchi, K.; Holtzman, A.; Rashkin, H.; Downey, D.; Yih, S. W.-t.; and Choi, Y. 2020. Abductive commonsense reasoning. In *International Conference on Learning Representations (ICLR)*.

Bosselut, A.; Rashkin, H.; Sap, M.; Malaviya, C.; Celikyilmaz, A.; and Choi, Y. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. *arXiv preprint arXiv:1906.05317*.

Christmann, P.; Saha Roy, R.; Abujabal, A.; Singh, J.; and Weikum, G. 2019. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 729–738.

Cingillioglu, N.; and Russo, A. 2018. DeepLogic: Towards End-to-End Differentiable Logical Reasoning. *arXiv preprint arXiv:1805.07433*.

Clark, P.; Tafjord, O.; and Richardson, K. 2020. Transformers as soft reasoners over language. *arXiv preprint arXiv:2002.05867*.

Cohen, W. W. 2016. Tensorlog: A differentiable deductive database. *arXiv preprint arXiv:1605.06523*.

Colmerauer, A. 1990. An introduction to Prolog III. In *Computational Logic*, 37–79. Springer.

Davis, E.; and Marcus, G. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM* 58(9): 92–103.

Dong, H.; Mao, J.; Lin, T.; Wang, C.; Li, L.; and Zhou, D. 2019. Neural logic machines. In *International Conference on Learning Representations (ICLR)*.

Evans, R.; and Grefenstette, E. 2018. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research* 61: 1–64.

Finkbeiner, B.; Hahn, C.; Rabe, M. N.; and Schmitt, F. 2020. Teaching Temporal Logics to Neural Networks. *arXiv preprint arXiv:2003.04218*.

Garcez, A. d.; Besold, T. R.; De Raedt, L.; Földiák, P.; Hitzler, P.; Icard, T.; Kühnberger, K.-U.; Lamb, L. C.; Miikkulainen, R.; and Silver, D. L. 2015. Neural-symbolic learning and reasoning: contributions and challenges. In *2015 AAAI Spring Symposium Series*.

Goldwasser, D.; and Roth, D. 2014. Learning from natural instructions. *Machine learning* 94(2): 205–232.

Grice, H. P. 1975. Logic and conversation. In *Speech acts*, 41–58. Brill.

Guo, D.; Tang, D.; Duan, N.; Zhou, M.; and Yin, J. 2018. Dialog-to-action: Conversational question answering over a large-scale knowledge base. In *Advances in Neural Information Processing Systems*, 2942–2951.

- Havasi, C.; Speer, R.; and Alonso, J. 2007. ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In *Recent advances in natural language processing*, 27–29. Citeseer.
- Havasi, C.; Speer, R.; Pustejovsky, J.; and Lieberman, H. 2009. Digital intuition: Applying common sense using dimensionality reduction. *IEEE Intelligent systems* 24(4): 24–35.
- Hixon, B.; Clark, P.; and Hajishirzi, H. 2015. Learning knowledge graphs for question answering through conversational dialog. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 851–861.
- Janiček, M. 2010. Abductive reasoning for continual dialogue understanding. In *New Directions in Logic, Language and Computation*, 16–31. Springer.
- Kocijan, V.; Lukaszewicz, T.; Davis, E.; Marcus, G.; and Morgenstern, L. 2020. A Review of Winograd Schema Challenge Datasets and Approaches. *arXiv preprint arXiv:2004.13831*.
- Labutov, I.; Srivastava, S.; and Mitchell, T. 2018. LIA: A natural language programmable personal assistant. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 145–150.
- Lenat, D. B.; Guha, R. V.; Pittman, K.; Pratt, D.; and Shepherd, M. 1990. Cyc: toward programs with common sense. *Communications of the ACM* 33(8): 30–49.
- Levesque, H.; Davis, E.; and Morgenstern, L. 2012. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- Li, T. J.-J.; Azaria, A.; and Myers, B. A. 2017. SUGILITE: creating multimodal smartphone automation by demonstration. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 6038–6049. ACM.
- Li, T. J.-J.; Labutov, I.; Li, X. N.; Zhang, X.; Shi, W.; Ding, W.; Mitchell, T. M.; and Myers, B. A. 2018. APPINITE: A Multi-Modal Interface for Specifying Data Descriptions in Programming by Demonstration Using Natural Language Instructions. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 105–114. IEEE.
- Li, T. J.-J.; Li, Y.; Chen, F.; and Myers, B. A. 2017. Programming IoT devices by demonstration using mobile apps. In *International Symposium on End User Development*, 3–17. Springer.
- Liu, H.; and Singh, P. 2004. ConceptNet—a practical commonsense reasoning tool-kit. *BT technology journal* 22(4): 211–226.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, 3749–3759.
- Mao, J.; Gan, C.; Kohli, P.; Tenenbaum, J. B.; and Wu, J. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *International Conference on Learning Representations (ICLR)*.
- Marra, G.; Giannini, F.; Diligenti, M.; and Gori, M. 2019. Integrating Learning and Reasoning with Deep Logic Models. *arXiv preprint arXiv:1901.04195*.
- Maslan, N.; Roemmele, M.; and Gordon, A. S. 2015. One hundred challenge problems for logical formalizations of commonsense psychology. In *AAAI Spring Symposium Series*.
- Mostafazadeh, N.; Roth, M.; Louis, A.; Chambers, N.; and Allen, J. 2017. Lsdsem 2017 shared task: The story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, 46–51.
- Mueller, E. T. 2014. *Commonsense reasoning: an event calculus based approach*. Morgan Kaufmann.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *NIPS 2017 Workshop on Autodiff*.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Qin, L.; Bosselut, A.; Holtzman, A.; Bhagavatula, C.; Clark, E.; and Choi, Y. 2019. Counterfactual Story Reasoning and Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5046–5056.
- Rashkin, H.; Sap, M.; Allaway, E.; Smith, N. A.; and Choi, Y. 2018. Event2mind: Commonsense inference on events, intents, and reactions. *arXiv preprint arXiv:1805.06939*.
- Reed, S.; and De Freitas, N. 2015. Neural programmer-interpreters. *arXiv preprint arXiv:1511.06279*.
- Rocktäschel, T.; Bošnjak, M.; Singh, S.; and Riedel, S. 2014. Low-dimensional embeddings of logic. In *Proceedings of the ACL 2014 Workshop on Semantic Parsing*, 45–49.
- Rocktäschel, T.; and Riedel, S. 2017. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, 3788–3800.
- Roemmele, M.; Bejan, C. A.; and Gordon, A. S. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *2011 AAAI Spring Symposium Series*.
- Sakaguchi, K.; Le Bras, R.; Bhagavatula, C.; and Choi, Y. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 8732–8740.
- Sakama, C.; and Inoue, K. 2016. Abduction, conversational implicature and misleading in human dialogues. *Logic Journal of the IGPL* 24(4): 526–541.
- Sap, M.; Le Bras, R.; Allaway, E.; Bhagavatula, C.; Lourie, N.; Rashkin, H.; Roof, B.; Smith, N. A.; and Choi, Y. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3027–3035.

- Sbisà, M. 1999. Presupposition, implicature and context in text understanding. In *International and Interdisciplinary Conference on Modeling and Using Context*, 324–338. Springer.
- Simons, M. 2013. On the conversational basis of some presuppositions. In *Perspectives on linguistic pragmatics*, 329–348. Springer.
- Speer, R.; Chin, J.; and Havasi, C. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Speer, R.; Havasi, C.; and Lieberman, H. 2008. AnalogySpace: Reducing the Dimensionality of Common Sense Knowledge. In *AAAI*, volume 8, 548–553.
- Srivastava, S. 2018. *Teaching Machines to Classify from Natural Language Interactions*. Ph.D. thesis, Samsung Electronics.
- Tur, G.; and De Mori, R. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Wang, W. Y.; and Cohen, W. W. 2016. Learning First-Order Logic Embeddings via Matrix Factorization. In *IJCAI*, 2132–2138.
- Weber, L.; Minervini, P.; Münchmeyer, J.; Leser, U.; and Rocktäschel, T. 2019. NLprolog: Reasoning with Weak Unification for Question Answering in Natural Language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, ACL 2019, Florence, Italy, Volume 1: Long Papers*, volume 57. ACL (Association for Computational Linguistics).
- Winograd, T. 1972. Understanding natural language. *Cognitive psychology* 3(1): 1–191.
- Wu, B.; Russo, A.; Law, M.; and Inoue, K. 2018. Learning Commonsense Knowledge Through Interactive Dialogue. In *Technical Communications of the 34th International Conference on Logic Programming (ICLP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Zhou, Z.-H. 2019. Abductive learning: towards bridging machine learning and logical reasoning. *Science China Information Sciences* 62(7): 76101.