ASTRID: Bootstrapping Commonsense Knowledge

Hans Peter Willems
MIND|CONSTRUCT
March 2021

Abstract

The need for Commonsense Knowledge in the machine becomes more and more apparent, as we try to move forward in the development of Artificial (General) Intelligence. It is becoming evident that this kind of knowledge is paramount in the human cognitive capacity and therefore also crucial for machine intelligence to ever reach any level of performance nearing that of a human brain. However, attaining a sufficient amount, and qualitative level, of human Commonsense Knowledge in the machine, appears to be an 'AI-hard' or 'AI-complete' problem.

How do humans do this? There is a lot to be learned from child development, and although there are AI-projects that try to use a developmental model to 'grow' intelligence, there have not been any (relevant) Commonsense Knowledge projects that leveraged the child development paradigm. That is, until now.

I present ASTRID (Analysis of Systemic Tagging Results in Intelligent Dynamics), a real-world implementation of a Cognitive Architecture based on human developmental models. The current state of this project is the result of a full decade of research and development. This paper describes the project background, underlying philosophies, objectives and current results and insights.

Keywords: Commonsense Knowledge, Unsupervised Transfer Learning, Machine Intelligence, Semantics, Natural Language Processing, Deep Inference.

The case for Commonsense Knowledge

As early as 1959, John McCarthy argued for the need of Commonsense Knowledge to attain human level Artificial Intelligence (McCarthy, 1959), currently referred to as Artificial General Intelligence (AGI). In subsequent decades of research into Symbolic solutions to Artificial Intelligence, Commonsense Knowledge has played an important role. With the rising popularity of Artificial Neural Networks (ANNs), and more specifically Deep Learning, the role of Commonsense Knowledge as an intrinsic part of (human) intelligence became understated.

Over the last few years, Commonsense Knowledge has regained interest in the AI-research field, predominantly in discussions about the current shortcomings of Deep Learning solutions (Marcus, 2017; Willems, 2021). While Deep Learning has been pitched as the solution for reaching Artificial General Intelligence (Hassabis, 2018), real progress in this direction has failed to materialize. Recently, it has been argued that Deep Learning might need some sort of hybrid coupling with Symbolic approaches to be able to attain Commonsense Knowledge (Marcus, 2002), and subsequently the capability of higher forms of reasoning.

The need for Commonsense Knowledge can easily be illustrated by the concept of 'context' (McCarthy, 2007). The lack of Commonsense Knowledge in current AI-approaches results in very shallow capabilities for determination, because both the problem-space and the solution-space lack context. Without context, the determinant information is too sparse to arrive at any robust results outside a narrow problem space. To be able to find the best solution for any real-world problem, a system needs access to a Commonsense Knowledge base that is actively maintained and current, and that is free of limiting biases. It also needs the capability to do several levels of inference across this knowledge base, to be able to find new information and handle novel situations.

Another argument for the need of Commonsense Knowledge comes from Case-based reasoning. This form of reasoning depends on prior experiences to handle novel situations in a structural way. Among other capabilities, this kind of reasoning relies on the storage of experiences and the capability to handle analogies (Kolodner, 1992). Experiences lean heavily on Commonsense Knowledge and the handling of analogies which, as stated before, needs contexts.

Commonsense-related problems in the AI-research field

There are several problems that have been defined as either scientific or technological obstacles in the quest for Artificial Intelligence. I will argue that many of these problems can be solved, at least in part, by implementing Commonsense Knowledge into the AI-system. I'll discuss the most obvious problems with their perceived solution based on Commonsense Knowledge.

Word Sense Disambiguation: Ambiguity of language is an ongoing field of research, related to Natural Language Processing and specific applications like Machine Translation, Sentiment Mining, Document Search and Text Classification (Wang et al., 2020). There have been many proposed solutions, mainly grammar-based, but thus far this has not been solved. A limited context is (mainly) available within the single sentence that contains the word to be determined. But to understand this limited context within one single sentence, a wider context about the use of this specific word and the other words in the current sentence is needed. Humans use Commonsense Knowledge, combined with several forms of higher reasoning, to determine such wider contexts. It stands to reason that having (enough) Commonsense Knowledge in the system, together with sufficient levels of reasoning capability, will solve the Word Sense Disambiguation problem. However, this hints at the notion that Word Sense Disambiguation might be part of the AI-complete problem (see below).

The Frame problem: Creating a usable description of the current (limited) reality wherein a certain problem needs to be solved, with enough detail to support a solution, is called the Frame problem (McCarthy and Hayes, 1969; Hayes, 1981). As a solution might involve several steps, and every step might depend on a (yet) unknown number of prerequisites, it is impossible to predict a solution. The Frame problem has (initially) been defined as being mainly related to the use of first-order logic and is also closely related to the subject of planning, but looking at it from a broader view it is evident that this has strong links to the Commonsense Knowledge problem. As the perceived state of the problem space relies extensively on ones beliefs and insights about that state, it becomes clear that having (up to date) Commonsense Knowledge in the system facilitates the re-evaluation of those beliefs and insights at any moment. This won't cancel the Frame problem entirely, but it can reduce the search space for possible solutions dramatically, to the point of feasibility.

Reasoning with uncertainty: The human brain relies largely on context when dealing with uncertainty. Contextual knowledge supports the discovery of useful analogies, that facilitate closing of the inference gap when directly related knowledge is absent. It is commonly accepted that the human brain uses analogies when confronted with novel problems, so it stands to reason that we should have a similar implementation in Artificial Intelligence. Instead of using a statistics-based algorithm to calculate the best option in cases where the correct answer is uncertain, we might expect Deep Inference across a rich context to be able to serve up the best answer. This is not to say that statistic models could not be used for the evaluation of the results given by Deep Inference.

The Symbol Grounding problem: The Symbol Grounding problem has been argued rather successfully, stating that adding related symbols to another symbol doesn't create understanding (Harnad, 1990). This statement has shown to be hard to refute. However, relating to this statement we should consider the actual need for Symbol Grounding. Humans are very capable of handling ungrounded symbols: Outside the scientific field there are many people that 'know' the fact that '$E=MC^2$' has something to do with 'relativity' and that this was discovered by Albert Einstein, without having any notion what $E=MC^2$ entails to, what relativity actually is, or to have ever met Einstein in person. For most people these are all ungrounded symbols, or at least very sparsely grounded symbols. This supports a strong suspicion that Symbol Grounding is not a primary requirement to be able to create usable contexts out of Commonsense Knowledge, and to be able to reason within those contexts. It also gives pointers to possible solutions to the Symbol Grounding problem, where rich conceptual contexts might give rise to Symbol Grounding.

The Combinatorial Explosion problem: This problem cannot be solved by Commonsense Knowledge directly. Bringing Commonsense Knowledge into a system might actually trigger a Combinatorial Explosion, especially when deeper layers of inference are involved. However, building a Commonsense Knowledge solution involves the development of a Knowledge Representation model that supports the needed (deep) inferences across the Knowledge base and therefore might trigger the problem. Solving (or evading) the Combinatorial Explosion problem is therefore an integral part of solving the Commonsense Knowledge problem. Devising a structure for knowledge representation that supports the building of a Commonsense Knowledge-based world-view, should therefore include the modeling of inference systems that evade the Combinatorial Explosion problem. This also points at, again, the very high probability that implementing Commonsense Knowledge in the machine might be an AI-complete problem.

AI-complete (AI-Hard): From the previous discussed issues, we arrive at the AI-complete problem. It is accepted that there are certain problems in Artificial Intelligence that can only be solved by solving AI-complete. The Commonsense Knowledge problem seems to be a prime candidate in this category. In 1996, Push Singh wrote the following:

> "AI researchers have been trying unsuccessfully to get around the need for common sense knowledge. To solve the hard problems in AI - natural language understanding, general vision, completely trustworthy speech and handwriting recognition - we need systems with common sense knowledge and flexible ways to use it. The trouble is that building such systems amounts to "solving AI". This notion is difficult to accept, but it seems that we have no choice but to face it head on."

Almost 25 years after this was written, we can now come to some startling conclusions. It is clear that Commonsense Knowledge has been at the heart of several research efforts, either to implement it, or to get rid of the need. For several decades now, this hasn't delivered a solution. It is also becoming apparent that implementing Commonsense Knowledge in a usable way, including constant updating and deep inference, basically equates to 'solving' Artificial (General) Intelligence. The Commonsense Knowledge problem clearly shows how it is intertwined with several other identified problems in the field, which illustrates the need to solve those problems in relation to each other. It seems clear that these problems cannot be solved in isolation, and may prove to be all AI-complete problems. This point has also been argued by others (Chalmers et al. 1991).

Problems in past and current Commonsense Knowledge projects

There have been several high-profile AI-projects that were aimed at building Commonsense Knowledge databases, some of them still running today. Unfortunately, none of these projects have surmounted to anything resembling human-level Commonsense Knowledge capabilities in a way that can support higher levels of reasoning. There are obvious reasons for failure that inhibit each of these projects in this respect. The most obvious one being that these projects try to solve a (very) small subset of Commonsense Knowledge modeling, making them insufficient for real-world AGI. As will be shown, this is mostly due to a lacking Knowledge Representation model, which stems from trying to stay away from the AI-complete problem. The other obvious shortcoming is the fact that none of these systems have capabilities for Unsupervised Transfer Learning (UTL), and therefore are not capable of sustaining constant updates and additions to the knowledge base. Looking at several of these high-profile projects reveals the issues at hand.

Cyc: This project, started by Doug Lenat in 1984, was aimed at building a Commonsense Knowledge database, capable of reasoning with implicit knowledge. It consists of a large database of facts and assumptions that are coded with a complex form of predicate logic. The project didn't solve the UTL-capability problem, and knowledge was entered manually by project staff. This issue was compounded by the fact that it uses a very complex predicate notation. The project still runs today, but its 'Artificial Intelligence' designation has been abandoned a long time ago. Today it is more or less seen as an advanced search engine for logical deliberation on specific knowledge domains. Specific applications need large scale manual configuration and tuning. As knowledge is manually entered into the system, there is no need for Natural Language Processing, so the project doesn't have an NLP focus. Because of the rather involved way of entering knowledge, the system is also not capable of dynamically updating its facts and assumptions, which makes it assumable that the actual knowledge in the system consists mainly of static facts. It is clear that the project didn't aim to be an AI-complete solution, neither to try to solve any of the individual hard problems (besides Commonsense Knowledge).

Never-Ending Language Learning (NELL): Launched in 2010 at the Carnegie Mellon University, the NELL project was aimed at mining knowledge from online website pages to find facts and add those to its knowledge base. Although the project has found an impressive number of facts while running almost continuously for the last decade, its knowledge representation model is rather shallow. The project started with a limited set of predefined semantic relationships to search for, and although additional semantic structures were added over time, the system still has several limitations in regard to real AGI systems. Its biggest shortcoming is the lack of temporal and spatial semantics, making it impossible to reason in these spaces (Mitchell et al. 2015). The use of formal predicate notation to describe relationships makes it impossible to do inferences on the importance of these relationships. Referential inferences are therefore limited to (shallow) belief weights that are given to the relations.

Open Mind Common Sense (OMCS): This project, launched in 1999, uses a crowd-sourcing model to let people (via the Internet) add Commonsense Knowledge to its knowledge base (Sing et al., 2002), an approach also used by the MindPixel project (see below). The project is still running today, but its focus has been shifting over the years. The big (initial) difference from the other projects listed here, is that facts were entered in plain English sentences, including the references between concepts. Unfortunately, as it became harder over time to parse the entered sentences into a usable form, the project shifted to entering information in a more structured form, limiting entered facts to a more strict notation. It is obvious that OMCS is not capable of adding and updating its stored facts, and assumptions on those facts. It also cannot do Natural Language Processing, as the abandoning of unstructured English sentences for input illustrates clearly. Projects that are based on OMCS, like ConceptNet, inherit the shortcomings noted here.

MindPixel: The MindPixel project, running from 2000 until 2005, was another project that leveraged crowd-sourcing for its knowledge base. The conceptual knowledge representation was rather limited, with statements being marked only true or false. Using a system where twenty (human) participants had to evaluate a statement, it tried to construct the truth level of a statement. The system did not encompass information on causal or temporal relationships. In my view the biggest issue with the MindPixel approach is that its model for stating facts is limited in itself: reality consists of uncountable numbers of facts that can both be true or false, depending on their context (contextual dependability). It is impossible to model the world with only true or false statements, even if the truthfulness is weighted, without access to more profound contexts.

ThoughtTreasure: This is one of the older projects. It started in 1993 and closed in 2000. The project was released into Open Source in 2015. ThoughtTreasure consists of concepts that are linked by assertions. The contextual structure is predefined, it does not support changes to this structure over time. The assertion notation does not follow natural language models. However, the notation does not cater for additional information besides the assertion. The system leans heavily on predefined classes, definitions and structures, making the knowledge base static in nature and therefore hard to update. The ThoughtTreasure knowledge base seems to have been constructed manually, although information on this issue is unclear.

Freebase: Yet another knowledge base that was created from existing data sources and additionally crowd sourced manual contributions. The project started in 2007 and was subsequently bought by Google in 2010. In the years after, Google moved the Freebase data over to Wikidata, moved the Freebase query infrastructure over to the Knowledge Graph API, and shut down the Freebase server in 2016. Freebase made use of the Resource Description Framework (RDF), which is predominantly aimed at querying structured data (as opposed to doing inference across connected concepts) and has no provisions for UTL or dynamic updates of the data. The last available data dump of Freebase contains an impressive 1.9 billion triples. It is 250 GB uncompressed data which is quite a lot for that amount of information, which is mainly due to the somewhat convoluted RDF notation.

YAGO (Yet Another Great Ontology): The project was initiated at the Max Planck Institute, and presented in 2007 (Suchanek et al., 2007). The YAGO knowledge base is being built with extracted information from Wikipedia and WordNet. It extracts category information from Wikipedia to build its ontology and uses WordNet to enrich that ontology both qualitatively and quantitatively. This means that the knowledge base is basically built from manually entered information (on Wikipedia). As the ontology is constructed from the (manually) pre-defined categories on Wikipedia, this also implies that the system cannot infer new categories on its own. Descriptors of assertions are concepts by themselves, but they are written as formal statements (e.g. hasWonPrize, bornInYear), therefore they are not an integral part of the actual concept database.

IBM Watson: The IBM Watson project has some impressive results to show for and some serious media hype to go along with that. The Watson project is a mixed system of AI-technologies, including Machine Learning through Artificial Neural Networks. It also makes use of the above mentioned YAGO knowledge base, which points to Watson's inability to actually learn concepts and assertions by itself. It also points to the fact that Watson consists of several technologies thrown together to solve a specific problem in an arguably narrow domain: the Jeopardy game. From this, it is obvious that there is no grand AI-complete model underlying the Watson technology. The lack of real-world applications based on Watson, and IBM moving out of several markets where Watson was promised to make headway, hints at the notion that the Watson system might be lacking in expected capabilities.

As a conclusion of this section, we can say that every Commonsense Knowledge project thus far failed in respect to one or more of the known AI-problems. The common themes among these failures are insufficient knowledge representation models, inability for self-learning in any form (unsupervised, reinforcement, transfer) and a (too) strong focus on logically testable facts that leaves no room for differences in experience, fuzzy interpretation of assumptions or abstract notions like 'pain', 'timeliness' and 'desire'. Some of these projects have gathered impressive amounts of concepts and facts, but have also been running for several years (in some cases even decades) almost continuously to get to these numbers.

What we can learn from the development in children

Children's developmental psychology can give us important insights into the human knowledge ‘bootstrapping’ process. Piaget’s well-known Cognitive Developmental Theory describes four stages in child development, that appear to map closely with a symbolic bootstrapping approach to cognitive development.

Newborn babies start exploring the world through touch (the Sensorimotor stage), which seems non-symbolic at first. However, from simple observations we can come to the conclusion that a baby’s touch sense is tuned towards finding specific information about the world they are in, like the form of things and surface materials. These things translate directly to concepts that function as building blocks for describing reality. Simple experiences map directly to basic but fundamental concepts. Eventually, when speech comes into play (the Pre-operational stage), those concepts are labeled to facilitate communication. However, before a concept is labeled, it already exists in its symbolic form in the child's brain. Cognitively speaking, it is already a 'thing'. If that would not be the case, it would be impossible to label it with language when speech is added.

During the first years of child development, we educate the child by explanation. Children are wired to get ‘bootstrapped’ this way, by continuously asking questions about things and events that make up their reality. This process keeps going on, until the child has enough basic knowledge to start making conceptual connections on its own (the Concrete Operational Stage). When children experience new information, they learn to comprehend that information by integration. Their bootstrapped Commonsense Knowledge base gives them the foundation for that comprehension, and through that process, to eventually self-learn.

Finally, in later years of childhood, the bootstrapped mind has accumulated enough (symbolic) connections to have access to Rich Conceptual Contexts that can support inference with abstract concepts (the Formal Operational Stage). This gives rise to hypothetical reasoning (what-if), deduction and abduction. This accelerates integration of new information into the contexts of other already known concepts.

This makes it clear that a 'system' cannot be UTL-capable from scratch. It needs to be bootstrapped with enough existing knowledge and insight, to create contextual structures for integration of newly learned knowledge. It also hints at the innate capabilities needed to start building the contextual structures, which should be reflected in the knowledge representation model.

ASTRID project objectives

The main objective of the ASTRID project is the development of a unified knowledge model that supports knowledge bootstrapping, continuous learning from sparse data (when bootstrapped) and finally support many forms of higher reasoning: structural, causal, temporal, inductive, deductive, abductive and hypothetical. The knowledge model must support inference across references, but also 'about' references, and therefore has to use a generalized notation for its predicate structure. It also has to be able to handle 'fuzzy' abstract predicate classes like 'sometimes', 'maybe' and 'many'. The system must be able to be bootstrapped from scratch, no existing structured corpora can be used. However, use of existing corpora to augment the system in a later stage should not be inhibited.

The second important objective is the capability to handle the differentiation between the abstract concepts that describe reality, and the real-world instantiation of those conceptual abstractions. The system must also be able to manage multiple real-world instantiations of one abstract concept and differentiate between those instantiations. This capability must be supported intrinsically by the knowledge representation model. This also points at a possible (partly) solution for the AI-complete problem. Description of the implementation of this capability is beyond the scope of this paper, but mentioned for completeness of the project's objectives.

Finally, the (trained) system should be able to function at a usable level while running on fairly minimal spec hardware, with portable levels of data storage. This is important to cater for high levels of autonomy in applications of the technology. The ASTRID benchmark implementation runs on a (virtualized) Quad-core CPU running at 3 GHz, 8 Gigabyte of RAM and 512 Gigabyte of storage. For fast bootstrapping and subsequent training of the system, faster machines and clustered server stacks can be utilized. However, continuous learning from occurring events during deployment is supported on the benchmark implementation.

Knowledge representation supporting Deep Inference

When we talk about the meaning of things, we talk about the semantic values of concepts (Saumier and Chertkow, 2002). A system that can somehow support the 'understanding' of things, should therefore be semantic in nature. A symbolic system for knowledge representation is inherently semantic in nature. Predicate logic, captured in some complex coded form, although supporting a semantic description of context, is not decidedly semantic in structure or nature.

To be able to do Deep Inference across conceptual contexts, the predicates that describe the relations within those contexts must also be conceptual to be semantic. The system should be able to reason about predicates in the same way as it can reason about the concepts that are related through those predicates. To understand the value of a predicate in a specific inference, the predicate should have the same Rich Contextual Representation as every other concept in the reasoning space. The review of other Commonsense Knowledge projects, earlier in this paper, shows clearly that elaborately coded complex predicates are a weakness in the design, not a strength, as they inhibit semantic inference on the predicates.

Predicates should also not be predefined. All reviewed projects have predefined predicates that demonstrably limit the scope of these systems. It can (erroneously) be argued that this is a form of bootstrapping. However, in all these projects the predefined predicate structures also implicate that there is no innate capability to identify unknown predicate structures. Bootstrapping implies that things go their own way after the bootstrapping phase, which is clearly not the case here. Manually adding new predicates to the system (which some projects support) is clearly not what is understood to be a UTL-capable system.

The ASTRID system uses a directed multi-dimensional graph for its knowledge base. Predicates are concepts themselves and have the same contextual structure. Concepts, including predicates, are described in natural language. Relations between concepts are described with a predicate concept and several weight values to facilitate complex inferences. Within the ASTRID knowledge graph everything is a concept. Conceptual contexts are multi-dimensional while the information structure is basically single-dimensional or flat. There is only one level, being concepts. Because of this structure, the ASTRID knowledge graph can (and does) model any ontology (within the reach of its current knowledge). Any concept can be seen as the top of an ontology that can be queried through inference. Because of this, the ASTRID system does not need predefined ontologies.

Because the ASTRID knowledge graph is structured as described, it facilitates multi-level inference across contextual structures, including the predicates. This gives rise to cognitive structuring of the Rich Contextual Representations during inference. In human cognition we call this insight, or understanding. In this way inference answers the obvious question: ‘What does it mean’. If this is coupled with simple core beliefs, we get very close to solving the Symbol Grounding problem.

The ASTRID knowledge graph also solves the Word Sense Disambiguation problem. Words or concepts are only ambiguous in isolation or in very sparse contexts. As soon as a system is capable of building rich contexts though inference, the ambiguity is lifted. This is not only true for ambiguous words within one language but also for jargon, slang, and most importantly for translation to words in other languages. Word Sense Disambiguation is at the heart of language translation because even non-ambiguous words can have several translations in another language, and only (deep) contexts can solve the selection of the right word translation in any specific case.

The Semantic bootstrapping system

With enough knowledge stored in the knowledge graph, the system is capable of learning new information by itself. New information gets integrated into the already available knowledge and can be inferred through the available contexts. However, similar to how humans start to learn, inference doesn’t work on an empty knowledge graph. This is where bootstrapping comes in. Analogous to how children learn, we need a teacher (lacking a parent) that can explain the semantics of concepts that are encountered during the early learning stage, when there is not enough context yet to infer semantic meaning. We also want that teacher to be much faster than human teaching, as we want to build the Commonsense Knowledge base as fast as possible.

The ASTRID system makes use of our Semantic Bootstrapping Trainer (SBT), as the teacher for the system. Information could be entered manually, but based on how the ASTRID knowledge graph is designed this would take huge amounts of manpower (as projects like Cyc has proven). The (partial) goal of the ASTRID system is to build Rich Contextual Representations, in need of many contextual relations for each concept. The SBT works together with ASTRID’s capability of finding predicates in sentences and create large amounts of contextual relations while being trained.

Besides the manpower needed to manually enter that amount of information, the ASTRID system discovers contextual relations that get overlooked by humans. As an impressive example of this capability, after being trained with only a few short stories in the realm of human social interaction, the system inferred that 'when sick, taking medicine might be advantageous'. This knowledge was not available as a strict sentence in the trained stories.

The SBT is a semantic part-of-speech tagger that makes use of Hidden Markov Chains to distill larger contexts beyond the traditional POS-approach with bi-grams and tri-grams. It is a rule-based tagger, using only a few dozen rules. In this configuration, tagging accuracy went up to 93% after being (manually) trained with less than a thousand sentences and having compiled a contextual map of bi-grams for only two-thousand words. The SBT learns recursively: wrongly tagged words are not corrected but merely marked as fault, the system is self-correcting based on later learned words.

Zipf's law, when applied to speech, shows clearly that a fairly small set of words at the highest frequency side of the distribution, encapsulate the most used words in communication. The fifty most used words in English account for almost fifty percent of all English communication. As the SBT itself is continuously trained while in turn training the ASTRID system with natural language sentences, both the SBT and the ASTRID knowledge base expectedly reflect Zipf's law in the training results.

The SBT maps semantics to sentences, based on our proprietary semantic model, which in turn facilitates ASTRID's capability of finding the contextual relations to build the complex representation of the world (or reality). This approach makes it possible to train the ASTRID system with any available text, where the texts don't need to be formatted or manually tagged in any way. Currently, the system is being trained with short stories from American literature, as those stories represent a large part of what we call Commonsense Knowledge. So far, the ASTRID benchmark implementation has found close to fifty thousand concepts and almost half a million relations. The bulk of this was achieved in just a few days of training. This in stark contrast to the many years and sometimes decades that other (failed) projects needed to get to their current results. ASTRID finding new concepts slows down remarkably after about thirty thousand concepts, again due to Zipf's law. However, training of the ASTRID system has shown that it keeps finding many new relations between already known concepts, even when the textual domain is limited to social human interaction. This proofs that Commonsense Knowledge is heavily dependent on rich contexts.

Conclusion

I submitted an argument for the need of Commonsense Knowledge in an Artificial General Intelligent system. Subsequently, the known AI-related problems that have a connection with Commonsense Knowledge were explored and the failures in Commonsense Knowledge projects were evaluated in context of the known AI-problems. A conclusion was reached that limited knowledge representation models lay at the heart of the problems and failures.

As an introduction to the philosophies underlying the ASTRID system I've pointed at child development as a source for answers to the Commonsense Knowledge problem. It was argued that children move through a 'bootstrapping' phase before they enter the phase of UTL-capability. Additionally, it was argued that such a bootstrapping phase in Artificial General Intelligence could solve the Commonsense Knowledge problem and subsequently several related problems.

I presented the ASTRID system, which is capable of Unsupervised Transfer Learning (UTL) after being sufficiently bootstrapped by a training facility. I have shown that our Semantic Bootstrapping Trainer is capable of bootstrapping the ASTRID system and facilitates ASTRID's capability of inferring contextual relations in natural language texts without the need for predefined ontologies and/or predicates. Furthermore, I have also described the structure of the ASTRID knowledge graph, and how that structure facilitates reasoning through Deep Inference.

The ASTRID system has been designed as an AI-complete system, starting with a knowledge representation model that intrinsically facilitates higher forms of reasoning with Commonsense Knowledge and is capable of continuous UTL after being bootstrapped. The knowledge representation model, and underlying semantic model, are central to solving the AI-complete problem. The ASTRID system already solves the known AI-problems as discussed, short of the AI-complete problem. From this perspective, expectations towards solving the AI-complete problem within the ASTRID project might not be far-fetched, although this will need further research and development.

References

1. McCarthy, John. (1959). Programs with Common Sense. http://jmc.stanford.edu/articles/mcc59.html

2. Gary Marcus (2017). Deep Learning: A Critical Appraisal. Self-published. arXiv. https://arxiv.org/abs/1801.00631

3. Willems, Hans Peter. (2021). Self-learning Symbolic AI: A critical appraisal. Self-published. https://mindconstruct.com/files/selflearningsymbolicaicritcalappraisal.pdf

4. Hassabis, Demis. (2010). A systems neuroscience approach to building AGI. Singularity Summit 2010 lecture. https://www.youtube.com/watch?v=Qgd3OK5DZWI

5. Marcus, Gary. (2002). The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence. Self-published. https://arxiv.org/abs/2002.06177

6. McCarthy, John. (2007). From here to human-level AI. Artificial Intelligence 171 (2007) 1174–1182. https://www.sciencedirect.com/science/article/pii/S0004370207001476?via%3Dihub

7. Kolodner, Janet. (1992). An introduction to case-based reasoning. Artificial Intelligence Review. 6. 3-34. http://web.media.mit.edu/~jorkin/generals/papers/Kolodner_case_based_reasoning.pdf

8. Wang J. Wang M. and Fujita H. (2020). Word Sense Disambiguation: A comprehensive knowledge exploitation framework. Knowledge-Based Systems, Volume 190 , 29 February 2020, 105030. https://www.sciencedirect.com/science/article/pii/S0950705119304344?via%3Dihub

9. McCarthy, John and Hayes, Patrick J. (1969). Some Philosophical Problems from the standpoint of Artificial Intelligence. Machine Intelligence 4. 463—502. http://jmc.stanford.edu/articles/mcchay69.html

10. Hayes, Patrick J, (1981). The Frame Problem and Related Problems in Artificial Intelligence, Readings in Artificial Intelligence 1981, Pages 223-230. https://www.sciencedirect.com/science/article/pii/B9780934613033500209

11. Harnad, S. (1990) The Symbol Grounding Problem. Physica D 42: 335-346. https://arxiv.org/abs/cs/9906002v1

12. Chalmers, David & French, Robert & Hofstadter, Douglas. (1996). High-Level Perception, Representation, and Analogy: A Critique of Artificial Intelligence Methodology. Journal of Expert Theory and Artificial Intelligence. 4. https://www.researchgate.net/publication/2798485_High-Level_Perception_Representation_and_Analogy_A_Critique_of_Artificial_Intelligence_Methodology

13. T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-ending learning. In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15). AAAI Press, 2302–2310. https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/10049/9557

14. Singh, P., Lin, T., Mueller, E., Lim, G., Perkins, T., & Zhu, W.L. (2002). Open Mind Common Sense: Knowledge Acquisition from the General Public. CoopIS/DOA/ODBASE. https://www.semanticscholar.org/paper/Open-Mind-Common-Sense:-Knowledge-Acquisition-from-Singh-Lin/45a23651bcc5a6cc993d722e71b0d301a6dc9dee

15. Suchanek, Fabian & Kasneci, Gjergji & Weikum, Gerhard. (2007). Yago: A Core of Semantic Knowledge Unifying WordNet and Wikipedia. 16th international conference on World Wide Web, May 2007, Banff, Canada. pp.697 - 697. https://hal.archives-ouvertes.fr/hal-01472497

16. Saumier, D. and Chertkow, H. (2002). Semantic Memory. Current Neurology and Neuroscience Reports 2002, 2:516–522. https://link.springer.com/article/10.1007/s11910-002-0039-9