Do Machines Understand?
A Short Review of Understanding & Common Sense in Artificial Intelligence

Kristinn R. Thórisson[1,2] & David Kremelberg[1]

[1] Icelandic Institute for Intelligent Machines, Reykjavik, Iceland
[2] CADIA, School of Comp. Sci., Reykjavik University
Reykjavik, Iceland

Abstract. While numerous systems have been developed with the aim of artificial intelligence (AI), few have explicitly targeted *understanding*. While the word "understanding" can be found in the title of a few fields (scene-, image-, and language-understanding) little attention has been paid in these fields to understanding *qua* understanding. The most explicit attempt at imbuing machines with understanding may perhaps be work on common sense reasoning, which nevertheless has diverted attention away from understanding proper to human-centric notions of "common sense." Searle's "Chinese Room" thought experiment was a notable attempt at addressing the subject, yet while raising many interesting discussions, their utility for advancing progress this actual topic is debatable. Here we discuss the state of scientific knowledge on understanding, review some of the relevant literature in light of our own and others' definitions, and propose a short list of what may be necessary to create machines with a capacity to understand.

1 Introduction

Few concerted attempts have been made in the artificial intelligence (AI) literature to imbue machines with *understanding*. No special sub-field exists within the field of AI focusing on understanding *per se*. When associated with particular research tracks or fields, "understanding" only makes an explicit appearance appended to *language*, *image*, and *scene* (Biederman 1985, Song et al. 2015, Winograd 1972). Within these fields progress on the subject of understanding proper has been slow, with the focus since their inception remaining largely at the level of signal processing, parsing, and symbol manipulation. In mainstream AI, understanding *qua* understanding has generally either been ignored, discussed only in relation to particular domains (cf. Biederman 1985, Laurentini 1994, Lin et al. 2013, Marslen-Wilson and Tyler 1980, Silberman et al. 2012, Song et al. 2015, Smith and Kanade 1997, Winograd 1972), or reframed as "common sense" (cf. Lenat et al. 1990, Liu and Singh 1994, Panton et al. 2006).

In the field of philosophy discussions of understanding have been dominated to a large degree by a language-centric viewpoint (cf. Grimm 1988, 2014, Potter 1994), somewhat aligned with the symbolic approach to common sense. On the rare occasion that the concept is addressed directly in AI research literature the term is commonly put in quotes, suggesting that even among those working in the field the idea that machines may have the potential to understand remains uncertain at best, and impossible at worst.

As a result of this state of affairs, the concept of understanding seems as elusive as ever. Here we argue that the process of understanding is different and distinct from general interpretations of "common sense," as well as human common sense, attempt to provide some perspectives on the subject to clarify why we think it has not been adequately addressed to date, and briefly discuss what may be missing. A quick note on the scope of our discussion is in order: We aim to limit our discussion to work within AI, to the more common meanings of the term "understanding" when used in the vernacular, and that could arguably / hypothetically be replicated / implemented with *information structures.*[1]

The remainder of the paper is organized as follows: First we discuss how understanding has been defined and how we define it, followed by a short review of two extreme positions seemingly underlying much of the work done on the subject in AI. Next, we discuss the extent to which general understanding can be equated with "common sense," followed by a discussion of Searle's Chinese Room argument (Searle 1980). Finally, we present our approach to model building and how understanding pertains to this approach.

2 What is Understanding?

Concrete definitions of understanding are few and far between, and when they appear they tend to be rather high-level (Thórisson & Kremelberg 2017). In the philosophical literature the importance of the topic to cognition has been emphasized, with some viewing it as being distinct and different from knowledge (Franklin 1983, de Gelder 1981, Grimm 1988). The acquisition – or deepening – of understanding has been viewed by some as constituting a greater intellectual achievement then knowledge acquisition alone (Grimm 1988, 2014); however, the opposite view has also been taken (Kvanvig 2003). Overall, while more attention has been paid to the subject within philosophy than AI, it has been stated as recently as the 1980s that a proper treatment of the concept has been more or less absent from philosophical discussions (Franklin 1983). As already mentioned, philosophical treatment of the topic has tended to focus on language, with a common definition being that "true understanding" is "true belief." Platonic definitions such as this tend not to be useful in the pursuit of making machines that understand, which may in part explain its absence from the AI literature.

For the term "understanding" to have utility within the field of artificial intelligence it must refer to a something that can be measured. In our approach, an agent's understanding of a phenomenon is testable on the basis of (at least) four criteria: An agent's ability to (1) predict the phenomenon's behavior, (2) achieve goals with respect to the phenomenon, to (3) explain the phenomenon, and (4) create or re-create the phenomenon (Bieger et al. 2017, Thórisson et al. 2016).

Here we thus focus on the pragmatic aspects of understanding, and the ways in which understanding may be useful in guiding behavior (Thórisson et al. 2016). While not strictly necessary for the remaining content of this paper, we will briefly outline here

[1] Philosophical and human-specific aspects of the concept will thus not be addressed—any recourse to unknown features of physics, nature, or the natural order of things that purportedly may underly human mental experience – e.g. "consciousness," "qualia," etc. – will be strictly outside our scope here.

our own view. We see understanding as *a process* that involves cognition interacting with a rule-governed world: Understanding a certain phenomenon involves the *creation and verification of models* of that phenomenon that can be used for four kinds of tests (see paragraph above). These kinds of models will generally be more useful the more *accurately* and the more *completely* they capture *causal* and *relational* aspects of the phenomenon in question, and allow themselves to be manipulated, independently of the phenomenon they model, for the above four purposes.

In our view, "true" understanding – or *general* understanding – thus requires a mechanism that, for any subject or topic, can freely (dynamically, more or less at any point in time) identify and model relations between observed phenomena and their features, and *generates new knowledge* that rests on existing knowledge, in a way that *improves and expands* what is already known (Thórisson et al. 2016). Some systems may already exist that have a weak form of this capability, but as far as we know no system exists that is not more or less completely tied to a particular (narrow) domain.

3 Two Common Arguments For & Against Machine Understanding

While it is difficult to definitively explain the apparent lack of significant focus on understanding in the literature, two recurring and diametrically opposed forms of reasoning seem to surface, both of which serve to obviate a justification for a discrete discussion on understanding: Firstly, there is the idea that a system able to do χ implies or provides proof of understanding – i.e., the system must understand χ in order to do χ. If this is assumed, there is no point in focusing specifically on understanding, since understanding and adequate level of performance are one and the same thing, and proof of understanding directly follows from the system being able to perform the task in question. Similarly, this relates to how "understanding" has commonly been used in the AI literature: within specific domains such as language and scene understanding, the way in which this word used – "understanding" within one specific domain – implies that machines don't need to understand in the sense that we're using this word.

Secondly, and inversely, there is the idea that when a machine is able to perform a certain task it serves to prove that understanding is not required to perform the task. In this it is assumed *a priori* that machines cannot understand, and therefore, if a machine is capable of performing a certain task, it proves that the task does *not* require understanding. Through this latter form of reasoning there is no basis upon which to discuss understanding proper, as it simply does not exist within the realm of machines.

Whichever of these lines of reasoning one accepts, understanding becomes a non-issue: If understanding of χ is given in any machine that can handle or achieve χ, then all one needs to do is to get the machine to do χ; if understanding is not needed to do χ, and all one wants is a machine that can do χ, then all one needs to do is to get the machine to do χ. In either case, understanding is irrelevant.

We see both of these lines of reasoning as faulty, but barring a concrete definition of the term, a couple of comments can be made here. First, the ability to perform a task does not (necessarily) imply or prove the presence of understanding: An agent (human, machine, or other) could be provided with a list of instructions for the purposes of completing the task, and simply follow the instructions for this purpose without any understanding, as in Searle's Chinese Room argument (Searle 1980). This is probably why some have argued that a chess machine, while being capable of doing the task it was designed for, does not truly *understand* what it's doing. Inversely, the inability to do a task does not (necessarily) imply lack of understanding: The reasons for this inability can be numerous and may have nothing to do with its understanding. Essentially, we view understanding as being distinct from action; understanding can exist with no action being performed or completed, and lack of understanding can be present despite the ability of the agent to complete the task.

4 Language, Image, & Scene Understanding

Three sub-fields of AI have adopted the term "understanding" in their name: Scene understanding (cf. Lin et al. 2013, Silberman et al. 2012, Song et al. 2015), image understanding (cf. Biederman 1985, Johnson-Laird 1983, Laurentini 1994, Smith and Kanade 1997), and language understanding (cf. Marslen-Wilson and Tyler, 1980, Smith and Kanade 1997, Winograd 1972). As their titles indicate, the focus on understanding in each of these contexts is on understanding within a particular and specific domain. When used within these particular domains, "understanding" has typically been defined in as narrow a sense as "target identification" (Biederman 1985), "identifying and reconstructing objects" (Laurentini 1994), "creating a short synopsis of a video" (Smith and Kanade 1997), and "question answering" (Winograd 1972), among others. It is clear that these systems *do something* that only animals could do before, and thus fit one of the more common definitions of AI (to get machines to do what only humans could do before); however, in reference to sections 3 and 2 above, one is hard pressed to argue that the associated systems *understand* what they are doing—even the most ambitious ones (cf. Li et al. 2009) certainly do not do so on any level approaching that of a human (cf. Zelinsky 2013), performing similar tasks in similar domains.

If these systems can be said to have some sort of understanding mechanisms it is clearly confined and exclusive to their domain: There is no hope that a visual object classifier will be able to learn how to create a short synopsis of a video, or that a target-identifier will ever do question answering. So, no matter how one slices it, whatever kind or limited amount of "understanding" mechanisms they have, it could not be said to be *general*. But could it perhaps be argued that such domain-specific systems have some ability to understand *their own* domain? This comes back to what we think "understanding" really is: If we agree with the claim that "to know is to understand," then these systems "understand" at the very least the things they are able to do/achieve; their ability to understand matches in fact precisely the scope of their capabilities. If, however, we dismiss this view as semantic trickery, explaining understanding away rather than explaining it, as argued in section 3, we must be able to point out how "genuine" understanding is different from this. It boils down to this: Take two systems A and B that are in all aspects identical except that A has little or no ability to acquire new knowledge and improve its current knowledge, nor can it acquire new understanding or improve its current one autonomously. System B, however, has both of these capabilities. If we call A intelligent we might by the same token be inclined to say that it understands. However, if we see autonomous knowledge acquisition as central to general intelligence, and autonomous understanding acquisition and deepening as central to (general) understanding, then A is neither intelligent nor does it understand. Whatever we call these things, one thing is clear: These systems are *not created equal*: They differ significantly in how they interact with their operating environment; B has built-in mechanisms to autonomously improve its own operation, the ability to achieve goals, and handle exceptions, leveraging its present state of knowledge and understanding to guide future knowledge acquisition and understanding efforts.

A mechanism for acquiring knowledge autonomously and cumulatively through observation and experimentation, even when limited to a particular domain, will make a system radically different from those systems that don't have such capabilities. A mechanism that enables a system to autonomously and cumulatively generate models of causal relations between observed and experienced phenomena, and use these models to explain (to itself and others) how things really work, will make that system radically different from other systems that don't have that mechanism: Such a system is much closer to the vernacular meaning of "understanding," in our view, than any system built in AI to date, and certainly closer than other definitions of the term we have come across. The more independent such an understanding mechanism is, the more likely is the system to be one we would be inclined to say has "true" or "general understanding."

5 Is "Common Sense" General Understanding?

The most serious attempts at addressing understanding in machines date back to the early days of AI (cf. McCarthy 1959, Minsky and Papert 1970), with some work continuing through the 80's and 90's (cf. Lenat et al. 1990, Liu and Singh 2004, Panton et al. 2006). To a first approximation, this work more or less equates the concept of understanding with that of "common sense," seemingly based on an assumption that understanding and common sense are synonymous—or can be treated as such for the purposes of making intelligent machines. In the majority of prior work on common sense and "common-sense reasoning" there is the further inherent assumption that the kind of common sense that research should be aimed at is *human* common sense.

A number of systems have been developed with the specific aim of performing common-sense reasoning (cf. Cambria et al. 2012, Lenat et al. 1990, Liu and Singh 2004, Panton et al. 2006, Poria et al. 2014). Generally, within the literature, "common sense" has been defined simply as "knowledge of the world that most people have," with attempts at creating systems incorporating common sense or common-sense reasoning being focused on simply providing machines with this same human-centric knowledge (cf. Cambria et al. 2012, Carbonell and Minton 1983, Lenat et al. 1990, Lieberman and Liu 2002, Liu and Singh 2004, McCarthy 1959, Panton et al. 2006, Poria et al. 2014). Compared with humans, however, it is debatable to what extent any of these systems can be said to have what we refer to as "common sense" in people, and the extent to which they are capable of common-sense reasoning. At the very least they do not exhibit what we could generally call human "sensibility" or "rationality" as evidenced by their inability to recognize their own failures. The relationship between common sense, or common-sense reasoning, and understanding appears even in casual observation to be more nuanced than that suggested in this (and other) AI literature.

Several conclusions can be drawn when examining current common-sense and expert systems. While not serving to prove that such systems are impossible with the methods employed to date, no system so far has demonstrated automatic acquisition of common-sense knowledge. With the notable exception of Cyc (Lenat et al. 1990), hardly any have even had this as a research objective. Also with the possible exception of Cyc, very few systems incorporate more than a few thousand *axioms / rules / knowledge-nodes / facts*, while all without exception illustrate brittle behavior through relatively frequent and unexpected errors and failures. This pervasive brittleness calls into question the extent to which common sense or common-sense reasoning can be captured in these approaches. We see understanding as a process by which reliable and useful knowledge can be acquired, improved, and updated continuously and actively, the knowledge being of a form that can then be used to predict, achieve goals, and explain. In this way, understanding seems to be a foundation for common sense—a prerequisite. Few if any common-sense systems incorporate the mechanisms needed to *acquire* understanding.[2]

6 Grounding & Understanding in Searle's Chinese Room

John Searle (1980) posited in his thought experiment, now commonly referred to as "the Chinese Room," that an individual with a sufficient list of rules to translate text from English to Chinese would *appear* to understand the text, even though he in fact does not. Searle argued that a computer program that converses seemingly intelligently with humans would similarly also be doing so with no understanding of the conversation.

The thought experiment touches on a number of issues, cognitive and otherwise. Searle's frequent use the words "meaningless" and "understanding" in the elaboration and description of his argument seems to indicate that the paper's primary focus is in fact understanding, and this will be the scope of its discussion here.[3]

One of the the most popular replies to Searle from those in AI is the "systems reply," namely, that while the man in the Chinese Room does not understand Chinese, the *system* understands Chinese. In general, we agree with Searle that the Chinese Room "system" does not *really* understand (if we were to accept systems reply we would have to call it some form of *weak* understanding). However, there are some unresolved problems with Searle's argument. To really achieve what Searle ascribes to the Chinese Room – a decent translation capability, meaning a notable amount of English words, sentences, and Chinese symbols – an enormous amount of rules would be required—rules that could be used to handle all sorts of conditions described in the English text. The thought experiment assumes the English text could reference anything seen, heard, imagined—indeed, anything that can be thought of and written down by humans; the size of this set is of course enormous. So is Searle's hypothetical list of rules realistic? Could the universe be written down as a list of binary $IF - THEN$ rules and run as a simulation on a computer? In principle perhaps, but this would require orders of magnitude more rules than the number of atoms in the universe, and we come to the realization that the question is not purely about theoretical matters in the abstract, but in fact also has an important practical component: Pragmatically speaking it would

---

[2] Cyc was predicted to be able to acquire new knowledge autonomously after a certain minimum number of axioms had been provided to it by hand (Lenat et al. 1990), but it was not provided with any special mechanisms for this.

[3] As section 2 describes, our conceptualization of understanding is purely a functional one, and thus "consciousness" and other concepts in Searle's arguments that are closer to philosophy, psychology, and metaphysics than AI, have in our view little or no weight in refuting the possibility of machines that understand.

be impossible to instantiate intelligence through such lists of rules, for two main reasons. First, due to limited represenational capacity of brains, to address this problem of exceedingly large numbers of rules we would need higher-level rules in the form of generalizations, and the ability to manipulate them to produce "new rules" which were not initially explicit—in other words, we would have to be able to do *reasoning* with the rules. Even more importantly, we would need mechanisms that can generate those rules from observation and experimentation, because the world will change, and not everything can be foreseen at the outset. On this path, therefore, Searle's argument does not succeed in refuting what it was intended to refute, namely, strong AI (and "strong understanding"). However, it may be a decent refutation – albeit a roundabout one – that expert systems could possess general intelligence (of the kind observed in humans and perhaps a few other animals).

We see grounding as a process by which knowledge is verified: Knowledge whose referent and quality cannot be ascertained – like that which is given to the man in the Chinese Room – is "ungrounded" in that it has not been verified and ascertained in ways that the man in the Chinese Room is familiar with. As a result, it cannot be usefully used to guide meaningful action. So we disagree with Searle that intentionality (taken to refer to the *meaning* or *purpose* of action) cannot be imparted to machines: What needs to be present though is proper grounding of the knowledge. In this sense, in our view understanding is grounded knowledge of a particular kind (namely, the kind that captures causal and other relations in useful information structures).

In the Chinese Room the rules are strictly about translation, presumably construed as some sort of correspondence mapping, in ways which could essentially be reduced to an enormous lookup table. How this could be used to achieve decent – not to mention good – translation is perplexing: One would assume that some sort of real-world knowledge would be required in addition, to arbitrate between subtle difference in meaning, for instance. So on this count the thought experiment is based on faulty assumptions. As the rules are predefined, and thus a priori cannot be grounded in the system's mechanisms and other knowledge, they are ungrounded: The Chinese Room as a system completely lacks grounding (as also argued by Searle (1980)). Without grounding, nothing meaningful can be built on the knowledge. If, however, the giant lookup table were acquired, updated, and managed by the system itself, actively and continuously, thus allowing its contents to be verified – to be grounded in experience – this system would have the beginnings of "true" (or "strong," or "general") understanding. We would still be left with the problem that the rules presumably are only good for one purpose: translation. While not directly or elegantly shown by the Chinese Room experiment, Searle's conclusion was in the very least correct, in our view: There is little if any understanding to speak of in this scenario, whether by the man in the Chinese Room or the system as a whole.

Elsewhere in this paper, and in prior work (Thórisson et al. 2016), we have argued that understanding must rely on knowledge of causes. To this hypothesis Searle does not actually seem to disagree: "Any mechanism capable of producing intentionality must have causal powers equal to those of the brain... Any attempt literally to create intentionality artificially (strong AI) could not succeed just by designing programs but would have to duplicate the causal powers of the human brain" (Searle 1980:417). While we are perhaps reading into his statement – and it is indeed not obvious what Searle means in fact by the phrase "causal powers of the brain" – interpreting this as "knowledge of causes" makes perfect sense to us. This, however, does not preclude it from the realm of computer intelligence, as Searle seems to argue.

So, to summarize, based upon the definition of understanding that we have previously presented (Thórisson et al. 2016) and is briefly outlined in section 2 above, the type of system Searle describes does *not* understand in the strong sense—it does not implement a *general understanding* mechanisms. Whether the system as a whole may have some form of weak understanding could perhaps be argued but is certainly a rather unimportant and uninteresting proposition, having more to do with definitions than general intelligence.

7 Understanding: A Short List of Necessary Ingredients

What is called for to create a machine that understands is, at the very least, mechanisms that allow (a) identification of causal (and other) relations among observed features of things, events, and phenomena (through experimentation and reasoning), (b) separation of relevant from irrelevant factors for various goals, scenarios, and domains (through context awareness), and (c) creation of (increasingly generalized) models that can be used to better (i) predict, (ii) explain, (iii) achieve goals with respect to, and in some fashion (iv) re-create the phenomena of interest. Barring this in a control system, what we have is essentially a system controlled by "a bag of tricks": Heuristics with limited ability to handle the system's own incorrectness, changes in environment, domain shifts, etc., through growth and evolution, and a severely reduced ability to learn from its own mistakes.

Models evaluated and tested in the domain of their reference *ground them* by measuring their usefulness for modeling what they are supposed to model (using the methods a - c above). A coherent set of models that models a host of phenomena and their relations will be useful in environments that change and evolve only insofar as they can be appropriately re-evaluated and updated in a timely fashion in light of such changes.

With the proper mechanisms we see understanding achievable in machines to the same depth and breadth as is evidenced in humans. We find it likely that multiple useful approaches may exist for building models; models capable of capturing causal-relational properties are however *required* in order for goal-achievement and explanation to be possible, which are two key ways in which understanding can be evaluated: While statistical correlation between e.g. some observable variables a and b may allow limited prediction, to know whether b will be affected if a is manipulated we need to know the causal direction—does a cause b, b a or is there perhaps a third variable c that causes them both? Thus, causal information is necessary for successful and reliable goal-achievement, as well as for (proper) explanation. Other approaches than those that capture causal-relational may provide depictions of target phenomena useful for some purposes, but will fail when using them for guiding general successful goal-directed action.

8 Conclusions

This short review of the literature on understanding in the field of AI highlights a dearth of discussion on the subject. Throughout the field we see a lot that has been given minimal attention, misinterpreted, or gone missing alltogether. For the most part, understanding has simply been ignored, or discussed in a very specific sense within particular domains. If mentioned at all in the AI literature, understanding is generally discussed within the context of one specific domain, such as language or scene understanding, or equated with common sense. The most well-known argument against the possibility of understanding in machines, Searle's Chinese Room, does not address systems that can model causality or whose knowledge is in some other way properly grounded, and in our opinion fails to refute the possibility of understanding in machines.

This current state of affairs is unfortunate, as a coherent conceptualization of understanding is needed in the field of AI (and especially artificial *general* intelligence) (1) so that it can be effectively investigated, (2) to allow for a comparison of different systems with regard to their level of understanding, and (3) so that system builders can design new systems, improve current ones, and train systems with a capacity for understanding as a primary goal.

References

Biederman, I. (1985). Human Image Understanding: Recent Research and a Theory. *Computer Vision, Graphics, and Image Processing 32*(1), 29-73.

Bieger, J. & Thórisson, K. R. (2017). Evaluating Understanding. IJCAI Workshop on Evaluating General-Purpose AI, Melbourne, Australia, Sep. 20.

Cambria, E., Olsher, D., Kowk, K. (2012). Sentic Activation: A Two-Level Affective Common Sense Reasoning Framework. *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence* (186-192). Toronto, Canada: AAAI Press.

Franklin, R.L. (1983). On Understanding. *Philosophy and Phenomenological Research 43*(3), 307-328.

de Gelder, B. (1981). I Know What You Mean, But If Only I Understood You... In H. Parret & J. Bouveresse (Eds.), *Meaning and Understanding* (44-61). New York: Walter de Gruyter.

Grimm, S.R. (1988). The Value of Understanding. *Philos. Compass 7*(2), 279-299.

Grimm, S.R. (2014). Understanding as Knowledge of Causes. In A. Fairweather (Ed.), *Virtue Epistemology Naturalized* (329-345). New York: Springer.

Johson-Laird, P. (1983). *Mental Models*. Cambridge, MA: Harvard University Press.

Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Laurentini, A. (1994). The Visual Hull Concept for Silhouette-Based Image Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*(2), 150-162.

Lenat, D.B., Guha, R.V., Pittman, K., Pratt, D., et al. (1990). CYC: Toward Programs With Common Sense. *Commun. ACM 33*(8), 30-49.

Li, L., Socher, R., Fei-Fei, L. (2009). Towards Total Scene Understanding: Classification, Annotation and Segmentation in an Automatic Framework. *Proc. IIIE Conf. Computer Vision and Pattern Recognition (CVPR).*

Lin, D., Fidler, S., Urtasun, R. (2013). Holistic Scene Understanding for 3D Object Detection With RGBD Camera. *Proc. IEEE.*

Liu, H., Singh, P. (2004). ConceptNet - a Practical Commonsense Reasoning Tool-Kit. *BT Technol. J. 22*(4), 211-226.

Marslen-Wilson, W., Tyler, L.K. (1980). The Temporal Structure of Spoken Language Understanding. *Cognition 8*, 1-71.

McCarthy, J. (1959). Programs With Common Sense. *Proceedings Symposium on Mechanization of Thought Processes*. Her Majestys Stationary Office, London.

Minsky, M., Papert, S. (1970). Proposal to ARPA For Research on Artificial Intelligence at M.I.T., 1970-1971. Artificial Intelligence Memo. No. 185, Massachusetts Institute of Technology A.I. Lab.

Nivel, E., Thorisson, K.R., et al. (2012). Bounded Recursive Self-Improvement. Reykjavik University Technical Report, RUTR-13006.

Panton, K., Matuszek, C., Lenat, D., Schneider, D., Witbrock, M., Siegel, N., Shepard, B. (2006). Common Sense Reasoning - From CYC to Intelligent Assistant. In Y. Cai & J. Abascal (Eds.), *Ambient Intelligence in Everyday Life* (1-31). *LNCS 3864*. Heidelberg: Springer.

Poria, S., Gelbukh, A., Cambria, E., Hussain, A., et al. (2014). EmoSenticSense: A Novel Framework for Affective Common-Sense Reasoning. *Knowl. Based Syst.* 69, 108-123.

Potter, V.G. (1994). *On Understanding Understanding: A Philosophy of Knowledge*. New York: Fordham University Press.

Schank, R.C., Abelson, R.P. (1977). *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Press.

Searle, J.R. (1980). Minds, Brains & Programs. *The Behavioral and Brain Sciences 3*, 417-457.

Silberman, N., Hoiem, D., Kohli, P., Fergus, R. (2012). Indoor Segmentation and Support Inference From RGBD Images. *Proceedings of the 12th European Conference on Computer Vision* Volume Part V, 746-760.

Smith, M.A., Kanade, T. (1997). Video Skimming and Characterization through the Combination of Image and Language Understanding Techniques. School of Computer Science Technical Report, Carnegie Mellon University, CMU-CS-97-111.

Song, S., Lichtenberg, S.P., Xiao, J. (2015). SUN RGB-D: A RGB-D Scene Understanding Benchmark Suite. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Thórisson, K.R. & Kremelberg, D. (2017). Understanding & Common Sense: Two Sides of the Same Coin? *Proceedings of the 10th Conference on Artificial General Intelligence (AGI-17)*.

Thórisson, K.R., Kremelberg, D., Steunebrink, B.R., Nivel, E. (2016). About Understanding. In B. Steunebrink, P. Wang, & B. Goertzel, B. (Eds.), *Artificial General Intelligence: 9th International Conference, AGI 2016, New York, NY, USA, July 16-19, 2016, Proceedings*, 106-117. Cham, Switzerland: Springer.

Winograd, T. (1972). Understanding natural language. *Cognitive Psychology 3*(1), 1-191.

Zelinsky, G. J. (2013). Understanding Scene Understanding. *Front. Pscyhol. 4*, 954.