

Superintelligence Safety: A Requirements Engineering Perspective*

Hermann Kaindl

*Institute of Computer Technology
TU Wien
Vienna, Austria*

hermann.kaindl@tuwien.ac.at

Jonas Ferdigg

*Institute of Computer Technology
TU Wien
Vienna, Austria*

e1226597@tuwien.ac.at

Corresponding Author: Hermann Kaindl

Copyright © 2023 Hermann Kaindl and Jonas Ferdigg. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract

Under the headline “AI safety”, a wide-reaching issue is being discussed, whether in the future some “superhuman artificial intelligence” / “superintelligence” could pose a threat to humanity. In addition, the late Steven Hawking warned that the rise of robots may be disastrous for mankind. A major concern is that even benevolent superhuman *artificial intelligence* (AI) may become seriously harmful if its given goals are not exactly aligned with ours, or if we cannot specify precisely its objective function. Metaphorically, this is compared to king Midas in Greek mythology, who expressed the wish that everything he touched should turn to gold, but obviously this wish was not specified precisely enough. In our view, this sounds like requirements problems and the challenge of their precise formulation. (To our best knowledge, this has not been pointed out yet.) As usual in requirements engineering (RE), ambiguity or incompleteness may cause problems. In addition, the overall issue calls for a major RE endeavor, figuring out the wishes and the needs with regard to a superintelligence, which will in our opinion most likely be a very complex software-intensive system based on AI. This may even entail theoretically defining an extended *requirements problem*.

Keywords: Artificial intelligence, superintelligence, AI safety, requirements problem

1. INTRODUCTION

The idea of technology becoming sentient has been a common theme in the literature and movies for decades. One might think of the famous movie “2001: A Space Odyssey” by Stanley Kubrick,

* This is an extended and substantially revised version of the paper Towards Requirements Engineering for Superintelligence Safety previously published in Joint Proceedings of REFSQ-2020 Workshops, Doctoral Symposium, Live Studies Track, and Poster Track, Pisa, Italy, 24-03-2020, published at <http://ceur-ws.org>

the Matrix, or the Terminator movies. These stories seem to be mostly consistent in the impression that a suddenly arising uncontrolled Artificial Intelligence (AI) will not mean us well, but is more likely to stand in direct opposition to humanity's goals. Furthermore, the AIs described in books and movies are never dumb machines, but rather potent entities with cognitive superpowers far beyond the capacities of human general intelligence — they are *superintelligent*. In his book “Superintelligence” [1], Oxford philosophy professor Nick Bostrom provides good reasons to believe that we can create such an entity.

In the AI research priorities document published in [2], apart from the desirability of safety, e.g., of self-driving cars, “AI safety” is addressed with regard to “superhuman AI”. This is partly based on forecasting work on intelligence explosion and “superintelligence” [1], where the importance of given goals to be exactly aligned with those of mankind is emphasized. Even that may be insufficient, however, the system must also somehow be deliberately constructed to pursue them [1].

In terms of requests like “Make paperclips”, such a system is envisaged to take everything it can (with high “intelligence”) and to make as many paperclips as it can make, whatever it may cost. Much like the example of king Midas or the core theme of the movie with the title “Bedazzled”, we view this as a requirement that is not specified precisely enough. While other references regarding “AI safety” can be found at <https://vkrakovna.wordpress.com/ai-safety-resources/>, it is mentioned nowhere that it is actually a problem with *requirements*.

Since mankind may eventually create a “superintelligence”, it should rather sooner than later specify the requirements on such a system, which are traditionally the required functions, qualities and constraints. *Goal-oriented RE* (GORE) focusses on goals, but these are usually goals of the stakeholders rather than those of an AI system (for GORE see, e.g., [3, 4] and successor work, and the very recent mapping study [5]). Hence, we propose to extend GORE theoretically in this direction for “AI safety”, where the goals of the “superintelligence” based on AI technology are supposedly of critical importance. Note, that such goals will most likely be conditions of something to achieve, in contrast to certain problem-solving algorithms where goals are defined states, see, e.g., [6, 7]. For unidirectional search, conditions as goals can be used, in contrast to bidirectional search.

Reasoning of a superintelligence on its own goals may involve a kind of *self-awareness* [8], or at least *self-representation* [9]. It is not clear, however, what the implications are on the problem at hand.

Apart from Nick Bostrom, many AI proponents attribute a moderate likelihood of superintelligence emerging within a 100-year horizon, see below under background material. Still, it is fair to doubt that such a superintelligence could ever be built, even though AI and especially Machine Learning are delivering success stories. Also for new AI systems in the near future that are beyond the state of art, however, the core point of our paper is important, i.e., there are requirements problems and the challenge of their precise formulation.

Having a *glass-box view* on a system implementing a superintelligence may be helpful for keeping control of it, such as following an approach to (object-oriented) software engineering for AI systems [10], [11]. Also making architectural decisions upfront, e.g., for a generic architecture [12], may be useful in this regard. Another approach is to model safety *frameworks*, see, e.g., [13]. However, implementing a superintelligence may well involve different approaches such as deep neural networks,

Table 1: When will human-level machine intelligence be attained? Source: [1, p. 19]

	10%	50%	90%
PT-AI	2023	2048	2080
AGI	2022	2040	2065
EETN	2020	2050	2093
TOP100	2024	2050	2070
Combined	2022	2040	2075

which are considered opaque, or approaches yet to be developed. Hence, we take a *black-box view* here and propose to focus on requirements.

Strictly speaking, the notion “AI safety” is misleading in our opinion, since it may also include the safety of certain AI-based systems, e.g., self-driving cars. Hence, we prefer the notion of “superintelligence safety”.

Safety of software-intensive systems is certainly a related area [14, 15], but it is primarily concerned with hazards in the environment of the system and related safety risks. *Functional safety* as dealt with, e.g., in the current automotive standard ISO 26262 focuses on failures of functions and how to manage them. Hence, we do not see how previous safety research could help much with regard to superintelligence safety.

The remainder of this paper is organized in the following manner. First, we provide some background material on superintelligence safety, in order to make this paper self-contained. Then we view superintelligence safety from the perspective of RE, and envisage formulating requirements accordingly. In order to address the essence of the given issue, we argue for doing RE on the superintelligence in the first place. Finally, we motivate how the theory of the *requirements problem* should be extended to cover goals of a superintelligence.

2. BACKGROUND ON SUPERINTELLIGENCE SAFETY

Surveys revealed that AI experts attribute a moderate likelihood of superintelligence emerging within a 100-year horizon, as can be seen in TABLES 1 and 2.¹ Several AI experts were asked to estimate a year by which the likelihood of having created Human-Level Machine Intelligence (HLMI) and superintelligence, respectively, reaches a certain percentage.

Assuming the creation of a superintelligence is possible, how can we control it and avoid the doomsday scenarios so eloquently depicted by authors and film directors? While an emerging superintelligence might not be inherently malevolent, it still poses a great risk. Without any information about the internal workings of such an entity, we have to assume that its way of ‘thinking’ might be very different from the way a human thinks. A superintelligence most likely will not share human morals or have a concept of value at all, if the developers did not intentionally design it to have one.

¹ PT-AI: Participants of the conference *Philosophy and Theory of AI*, 2011. AGI: Participants of the conference *Artificial General Intelligence and Impacts and Risks of Artificial General Intelligence*, 2012. EETN: Sampled members of the Greek Association for Artificial Intelligence, 2013. TOP100: The 100 top authors in artificial intelligence, measured by citation index, May 2013

Table 2: How long from human level to superintelligence? Source: [1, p. 20]

	2 years after HLMI	30 years after HLMI
TOP100	5%	50%
Combined	10%	75%

While the most terrifying scenario is a malevolent superintelligence with the intent to eradicate humanity, another potential scenario is a superintelligence with no inherent motivations at all, simply following the instructions of its creator to the best of its abilities. Bostrom describes some of the abilities a superintelligence could have under the term “Cognitive Superpowers” [1, p. 91]. A superintelligence could excel at strategic planning, forecasting, analysis for optimizing chances of achieving distant goals, social and psychological modeling and manipulation, rhetoric persuasion, hacking into computer systems, design and modeling of advanced technologies, and the list goes on [1, p. 94].

One can see that even equipped with only a subset of these skills, a superintelligence pursuing a goal *to the best of its abilities* and without being constrained by concepts like morality or value, can cause substantial damage to its environment. Bostrom depicts this with his famous thought experiment where a superintelligence is given the simple task to make as many paper clips as possible. While the superintelligence might start off by acquiring monetary resources by predicting the stock market and building paper clip factories, it may not just stop there. Since its goal is the unconstrained maximization of the production of paperclips, it will soon discover that humans pose a hindrance to its endeavor, as we might try to stop the superintelligence from producing more paperclips or even try to shut it off. Since the superintelligence is multiple magnitudes smarter than humanity combined, humanity fails to stop the superintelligence and ultimately the whole world (the whole observable universe, in fact) will be made into paperclips.

Even if this was just a contrived example, one can see the risks of giving a superintelligence an unconstrained optimization problem. This argument is also extended in [1] to constrained optimization problems, but the details are not necessary for our paper.

3. SPECIFYING AND COMMUNICATING REQUIREMENTS TO A SUPERINTELLIGENCE

Under the (usually unrealistic) assumption that we knew the requirements already, ‘just’ telling the superintelligence about them properly is the problem here. It can be decomposed into specifying and communicating.

According to the Standard ISO/IEC 10746-2:2009 Information technology – Open Distributed Processing – Reference Model: Foundations, 7.4, a *specification* is a “concrete representation of a model in some notation”. In order to better understand this sub-problem of specifying requirements, let us follow the observation in [16] that requirements *representations* are often confused with requirements *per se*. This confusion is also widespread in practice, as exemplified in the very recent Standard ISO SE Vocabulary 24765:2017. In fact, it defines a requirement both as a “statement that

translates or expresses a need and its associated constraints and conditions” and “a condition or capability that must be present in a product ...” (dots inserted).

As it is well-known in RE, specifying requirements can be done informally, say, using some natural language, or formally, e.g., using some formal logic as the notation. Possibilities for semi-formal representations within this spectrum are many-fold.

In order to avoid *ambiguity* in the course of specifying requirements for a superintelligence, *formal* representation of the requirements should be the choice, of course, possibly in some formal logic. *Grounding the logic* used in the domain will, however, still leave loopholes for the superintelligence to ‘misunderstand’ the given specification of our requirements. For example, for some predicate on paperclips, a grounding of what paperclips are in the real world will be necessary.

In addition, *incompleteness* of the specification remains an open issue. Unfortunately, no general solution appears to be feasible at the current state of the art in RE.

Communicating a given requirements specification to a superintelligence may be investigated according to [17], based on *speech-act theory* [18], under the premise that stakeholders communicate information in the course of requirements engineering. This view was not taken for the sake of really communicating requirements in [17], but for using the semantics of various speech acts to theoretically (re-)define the *requirements problem* originally defined by Zave and Jackson [19] (see also below). Everything is assumed to be communicated by speech acts here. However, we can only communicate *representations* of requirements, and not requirements *per se*. It is clear that the text given in the examples in [17] *represents* something by describing it, the very confusion addressed in [16], where specific examples of this confusion are given.

Anyway, for communicating specifications of requirements to a superintelligence, speech-act theory could be helpful for annotating the specific kind of speech act, e.g., *Question* or *Request*. For example, the text “Can you produce paperclips?” may be interpreted either way, unless specified more precisely.

Still, we have to face that ‘perfectly’ specifying and communicating requirements to a superintelligence may not be possible. Even if we could, this would not help in case of malevolent superintelligence that would not necessarily satisfy requirements as specified and communicated. Hence, let us consider building a superintelligence in such a way that these problems can be mitigated.

4. RE ON A SUPERINTELLIGENCE

In our view, building a superintelligence should certainly include RE (much as building any non-trivial system). For specifically addressing superintelligence *safety*, we may ignore defining *functional requirements* on its superintelligent abilities. Still, some functionality may have to be defined for operationalizing certain “non-functional” safety requirements, which will in the first place be cast as *constraints* on the superintelligence.

Working on these requirements could be done as usual according to best practice of RE. This will involve major stakeholders (in particular, AI researchers), requirements elicitation may be done

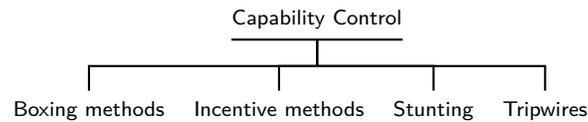


Figure 1: Overview of capability control techniques

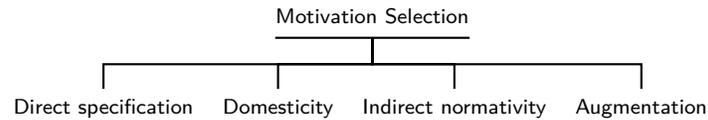


Figure 2: Overview of motivation selection techniques

including questionnaires in the Web, etc. While the core questions will be about requirements on superintelligence safety, of course, a bigger question should be asked in our opinion, what the goals and requirements are that mankind has regarding AI in the first place.

We expect that RE for *superintelligence safety* will pose even greater challenges than the usual practice of RE. After all, it is not just about conceiving such a superintelligence but about making sure that its creation will not raise uncontrollable safety risks.

Let our RE endeavor be also informed by some preliminary thought by the author who raised the issue of superintelligence safety. In [1, Chapter 9], *controlling* a superintelligence is discussed, in order to deal with this specific safety problem. Two broad classes of potential methods are distinguished — *capability control* and *motivation selection*. Within each of them, several specific techniques are examined.

We review the key ideas here, starting with *capability control*, where FIGURE 1 shows an overview of the related techniques summarized as follows [1, p. 143]:

- **Boxing methods**
“The system is confined in such a way that it can affect the external world only through some restricted pre-approved channel. Encompasses physical and informational containment methods.”
- **Incentive methods**
“The system is placed within an environment that provides appropriate incentives. This could involve social integration into a world of similarly powerful entities. Another variation is the use of (cryptographic) reward tokens. ...” (dots inserted)
- **Stunting**
“Constraints are imposed on the cognitive capabilities of the system or its ability to affect key internal processes.”
- **Tripwires**
“Diagnostic tests are performed on the system (possibly without its knowledge) and a mechanism shuts down the system if dangerous activity is detected.”

Viewed from an RE perspective, it seems as though most of what is discussed here could be elaborated as *constraint requirements*. Both current theory and practice of RE will most likely be sufficient for dealing with *capability control* as laid out here.

With regard to the ideas on *motivation selection*, FIGURE 2 shows an overview of the related techniques summarized as follows [1, p. 143]:

- Direct specification
“The system is endowed with some directly specified motivation system, which might be consequentialist or involve following a set of rules.”
- Domesticity
“A motivation system is designed to severely limit the scope of the agent’s ambitions and activities.”
- Indirect normativity
“Indirect normativity could involve rule-based or consequentialist principles, but is distinguished by its reliance on an indirect approach to specifying the rules that are to be followed or the values that are to be pursued.”
- Augmentation
“One starts with a system that already has substantially human or benevolent motivations, and enhances its cognitive capacities to make it superintelligent.”

In our opinion, these ideas will be much harder to elaborate according to the current state of the art in RE than the ideas on capability control above. When interpreting some of them in such a way as involving *dynamic goals of the superintelligence*, GORE comes to mind. However, can the current GORE approaches and especially the current theoretical formulation of the RE problem really cover that?

5. TOWARDS A NEW THEORY OF THE RE PROBLEM

In order to address this question, we raise further questions:

- What exactly is the requirements problem in this regard?
- Can we cast it in terms of the theoretical requirements problem?

In this regard, we investigate the state of the art on formulations of the (theoretical) requirements problem. Based on the seminal work of Zave and Jackson [19], Jureta et al. [17] extended the formulation of the requirements problem. The essence of this new formulation is as follows (where the notation is simplified for the purposes of this paper):

Given:

K domain assumptions

G goals

Q quality constraints and softgoals
 A attitudes (preferences)

To be found:

P plans

$$K, P \mid G, Q, A$$

An important change in this formulation is the defeasible *consequence* relation for non-monotonic satisfaction instead of (monotonic) entailment, based on the insight that the latter would be unrealistic with regard to practice. We consider this a useful variation of the previous formulation in [19], in particular with respect to what is given in terms of goals and softgoals. Hence, this formulation elaborates the RE problem towards GORE. Goals in GORE research are typically *wishes* (desires) of some agent / stakeholder.

In addition, Jureta et al. [17] introduced *attitudes* for evaluation in terms of degree of favor or disfavor of other elements in the problem formulation. While the formulation in [19] assumed that all given “requirements” are compulsory, the inclusion of attitudes leads to a more realistic formulation where some are not compulsory but optional. It also covers *preferences* for comparison between different components of the same type, which establishes orders between components of the same type. (In particular, for each preference order in A over softgoals, there is a preference order that maintains that same ordering over quality constraints that stand in the justified approximation relationship to the given softgoals defined in [17].) This makes it possible to define an optimum in terms of stakeholder attitudes.

From a practical point of view, we do not think that both quality constraints and softgoals (Q) will normally be given *a priori* (as assumed in [17]). RE should include the derivation of quality constraints from softgoals. Such derivations are actually the key task of the softgoal approach to RE.

In contrast to the original approach in [19], this theoretical formulation involves *goals*. However, these are only goals of stakeholders in RE, while goals of the system-to-be-built itself are not included.

Later, Jureta et al. [20] formally showed the fundamental differences between standard RE (as sketched above) and RE for *adaptive systems*. A system is adaptive if it can detect differences between its requirements and runtime performance and can adjust its behavior to cope with such deviations. It spans design-time and run-time; design-time, because design decisions influence the range of monitored inputs for the system, and the feedback mechanisms it will have; and run-time, because these mechanisms enable the system to react to at least some changes rather than ignore them. The solution concept, called *configurable specification*, amounts to a set of requirements configurations and evolution requirements for switching between configurations. Each configuration is shown to satisfy all properties required by [17, 19]. During adaptation, the system switches from one configuration to another and does so because awareness requirements became violated during the last period of stability.

The latter RE problem defined for adaptive systems fits an AI-based software-intensive system most likely better than the previous RE theory. However, it still takes only goals of stakeholders

into account like [17] but not goals of the system-to-be-built itself. For more conventional software and systems, this was not necessary, of course.

Rational agents like those whose decision-making has been formally verified by Dennis et al. [21] have goals assigned, but statically and at design-time. Such agents can be designed, e.g., using the knowledge level software engineering methodology for agent-oriented programming Tropos [22].

Tropos takes goals of *agents* into account for building agent-oriented software systems. In the course of GORE, the goals of stakeholders are broken down into goals assigned to agents. As long as this is done correctly, the goals of the stakeholders are aligned with the goals of the agents.

However, an AI-based system may also create new goals itself at run-time [23]. We assume that a superintelligence will (have to) be able to formulate goals on its own. This is not covered theoretically yet in terms of theoretical formulations of requirements problems.

Hence, for treating superintelligence safety as a requirements problem, we propose to develop a new theoretical formulation applicable to it, since it involves goals of both stakeholders *and* of the AI-based system-to-be-built. In particular, it involves that these goals will be aligned, and that the system-to-be-built can change them itself at run-time. It will be challenging to capture all this in new theoretical formulations of extended requirements problems covering superintelligence safety.

6. CONCLUSION

In this paper, we address the potentially very important issue of “AI safety”, in the sense of *superintelligence safety* [1], from an RE perspective. To our best knowledge, this is the first RE approach to this issue, although it may seem obvious that RE is the very discipline of choice here.

We actually distinguish two different approaches:

- Specifying and communicating requirements for specific problems to a concrete superintelligence, and
- Doing RE in the course of building a superintelligence in the sense of an AI-based software-intensive system.

Based on common wisdom of RE at the current state of the art, we tentatively conclude that ‘perfectly’ specifying and communicating requirements to a superintelligence may not be possible. And even if this became possible in the future, it would not help in case of malevolent superintelligence that would not necessarily satisfy requirements as specified and communicated.

Hence, we envision doing RE on a superintelligence to be built. Such an RE endeavor could be informed by some preliminary thought on *controlling* a superintelligence in [1]. The first approach through *capability control* may be dealt with properly through *constraint requirements*. The second approach for controlling through *motivation selection*, however, appears to go beyond the current theory of RE. In particular, we raise the challenge of extending GORE with *dynamic goals of a superintelligence*.

After having revisited the existing formulations of the (theoretical) *requirements problem*, we tentatively conclude that new and extended formulations will be needed for theoretically founded RE on superintelligence. Defining a new requirements problem for a dynamic goal model of the system-to-be-built is a challenge, but the impact on an improved understanding of problems like “AI safety” (in the sense of superintelligence safety) would be great.

7. ACKNOWLEDGMENT

We thank Tom Everitt for his comments on a previous version of this paper.

References

- [1] Bostrom N. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press. 2014.
- [2] Russell S, Dewey D, Tegmark M. Research Priorities for Robust and Beneficial Artificial Intelligence. *AI Mag*. 2015;36:105-114.
- [3] Dardenne A, van Lamsweerde A, Fickas S. Goal-Directed Requirements Acquisition. *Sci Comput Program*. 1993;20:3-50.
- [4] Mylopoulos J, Chung L, Yu E. From Object-Oriented to Goal-Oriented Requirements Analysis. *Commun ACM*. January 1999;42:31–37.
- [5] Horkoff J, Aydemir FB, Cardoso E, Li T, Maté A, et al. Goal-Oriented Requirements Engineering: An Extended Systematic Mapping Study. *Requirements Eng*. 2019;24:133-160.
- [6] Kaindl H, Kainz G, Leeb A, Smetana H. How to Use Limited Memory in Heuristic Search. In *Proc. Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*. 1995;18:226-242.
- [7] Kaindl H, Kainz G. Bidirectional Heuristic Search Reconsidered. *J Artif Intell Res*. 1997;7:283-317.
- [8] Jantsch A, Tammemäe K. A Framework of Awareness for Artificial Subjects. In *Proceedings of the 2014 International Conference on Hardware/Software Codesign and System Synthesis*. 2014; 1-3.
- [9] Kaindl H, Vallee M, Arnautovic E. Self-Representation for Self-Configuration and Monitoring in Agent-Based Flexible Automation Systems. *IEEE Trans Syst Man Cybern Syst*. January 2013;43:164-175.
- [10] Kaindl H. Object-Oriented Approaches in Software Engineering and Artificial-Intelligence. *Journal of object-oriented programming*. 1994;6:38-45.
- [11] Kaindl H. Is Object-Oriented Requirements Engineering of Interest? *Requirements Eng*. 2005;10: 81-84.

- [12] Mannion M, Lewis O, Kaindl H, Montroni G, Wheadon J. Representing Requirements on Generic Software in an Application Family Model. *Softw Reuse Adv Softw Reusability*. 6th International Conference, ICSR-6, Vienna, Austria, June 27-29, 2000. Proceedings 6 . Springer Berlin Heidelberg. 2000:153-169.
- [13] Everitt T, Kumar R, Krakovna V, Legg S. Modeling AGI Safety Frameworks With Causal Influence Diagrams. In: *Proceedings of the workshop on artificial intelligence safety 2019*. CEUR Workshop proceedings. 2019;2419.
- [14] Smith DJ, Simpson KGL. *Safety Critical Systems Handbook: A Straight Forward Guide to Functional Safety*. IEC; 2010. p. 61508 Edition) and Related Standards. Including Process IEC. 61511 and machinery IEC 62061 AND ISO 13849. Elsevier; 2011
- [15] Hulin B, Kaindl H, Rathfux T, Popp R, Arnautovic E, et al. Towards a Common Safety Ontology for Automobiles and Railway Vehicles. In: *12th European Dependable Computing Conference (EDCC)*; 2016: 189-192.
- [16] Kaindl H, Svetinovic D. On Confusion Between Requirements and Their Representations. *Requirements Eng*. 2010;15:307-311.
- [17] Jureta I, Mylopoulos J, Faulkner S. Revisiting the Core Ontology and Problem in Requirements Engineering. In: *16th IEEE International Requirements Engineering Conference*; 2008: 71-80.
- [18] Searle JR. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge, England: Cambridge University Press; 1969.
- [19] Zave P, Jackson M. Four Dark Corners of Requirements Engineering. *ACM Trans Softw Eng Methodol*. 1997;6:1-30.
- [20] Jureta IJ, Borgida A, Ernst NA, Mylopoulos J. The Requirements Problem for Adaptive Systems. *ACM Trans Manag Inf Syst*. 2015;5:1-33.
- [21] Dennis LA, Fisher M, Lincoln NK, Lisitsa A, Veres SM. Practical Verification of Decision-Making in Agent-Based Autonomous Systems. *Autom Softw Eng*. September 2016;23:305-359.
- [22] Bresciani P, Perini A, Giorgini P, Giunchiglia F, Mylopoulos J. A Knowledge Level Software Engineering Methodology for Agent Oriented Programming. In: *Proceedings of the fifth international conference on autonomous agents, agents '01*. Acad m. NY; 2001: 648-655.
- [23] Cox MT. Perpetual Self-Aware Cognitive Agents. *AI magazine*. 2007;28:32.