

Analysis of Types of Self-Improving Software

Roman V. Yampolskiy

Computer Engineering and Computer Science
Speed School of Engineering
University of Louisville, USA

roman.yampolskiy@louisville.edu

Abstract. Software capable of improving itself has been a dream of computer scientists since the inception of the field. In this work we provide definitions for Recursively Self-Improving software, survey different types of self-improving software, and provide a review of the relevant literature. Finally, we address security implications from self-improving intelligent software.

Keywords: recursive self-improvement, self-modifying code, self-modifying software, self-modifying algorithm, autogenous intelligence, bootstrap fallacy

1 Introduction

Since the early days of computer science, visionaries in the field anticipated creation of a self-improving intelligent system, frequently as an easier pathway to creation of true artificial intelligence¹. As early as 1950 Alan Turing wrote: “Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child’s? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child-brain is something like a notebook as one buys from the stationers. Rather little mechanism, and lots of blank sheets... Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child” [1].

Turing’s approach to creation of artificial (super)intelligence was echoed by I.J. Good, Marvin Minsky and John von Neumann, all three of whom published on it (interestingly in the same year, 1966): Good - “Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an ‘intelligence explosion,’ and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make” [2]. Minsky - “Once we have devised programs with a genuine capacity for self-improvement a rapid evolutionary process will begin. As the machine improves

¹This paper is based on material excerpted, with permission, from the book - Artificial Superintelligence: a Futuristic Approach © 2015 CRC Press.

both itself and its model of itself, we shall begin to see all the phenomena associated with the terms “consciousness,” “intuition” and “intelligence” itself. It is hard to say how close we are to this threshold, but once it is crossed the world will not be the same” [3]. Von Neumann - “There is thus this completely decisive property of complexity, that there exists a critical size below which the process of synthesis is degenerative, but above which the phenomenon of synthesis, if properly arranged, can become explosive, in other words, where syntheses of automata can proceed in such a manner that each automaton will produce other automata which are more complex and of higher potentialities than itself” [4]. Similar types of arguments are still being made today by modern researchers and the area of RSI research continues to grow in popularity [5-7], though some [8] have argued that recursive self-improvement process requires hyperhuman capability to “get the ball rolling”, a kind of “Catch 22” .

2 Taxonomy of Types of Self-Improvement

Self-improving software can be classified by the degree of self-modification it entails. In general we distinguish three levels of improvement – modification, improvement (weak self-improvement) and recursive improvement (strong self-improvement). However, it is easy to see that recursive improvement is a type/subset of improvement which is a subset of modification. For the purposes of this paper we will treat them as separate classes.

2.1 Self-Modification

Self-Modification does not produce improvement and is typically employed for code obfuscation to protect software from being reverse engineered or to disguise self-replicating computer viruses from detection software. While a number of obfuscation techniques are known to exist [9], ex. self-modifying code [10], polymorphic code, metamorphic code, diversion code [11], none of them are intended to modify the underlying algorithm. The sole purpose of such approaches is to modify how the source code looks to those trying to understand the software in questions and what it does [12].

2.2 Self-Improvement

Self-Improvement or Self-adaptation [13] is a desirable property of many types of software products [14] and typically allows for some optimization or customization of the product to the environment and users it is deployed with. Common examples of such software include evolutionary algorithms such as Genetic Algorithms [15-20] or Genetic Programming which optimize software parameters with respect to some well understood fitness function and perhaps work over some highly modular programming language to assure that all modifications result in software which can be compiled and evaluated. The system may try to optimize its components by creating internal tournaments between candidate solutions. Omohundro proposed the concept of

efficiency drives in self-improving software [21]. Because of one of such drives, balance drive, self-improving systems will tend to balance the allocation of resources between their different subsystems. If the system is not balanced overall performance of the system could be increased by shifting resources from subsystems with small marginal improvement to those with larger marginal increase [21]. While performance of the software as a result of such optimization may be improved the overall algorithm is unlikely to be modified to a fundamentally more capable one.

Additionally, the law of diminishing returns quickly sets in and after an initial significant improvement phase, characterized by discovery of “low-hanging fruit”, future improvements are likely to be less frequent and less significant, producing a Bell curve of valuable changes. Metareasoning, metalearning, learning to learn, and lifelong learning are terms which are often used in the machine learning literature to indicate self-modifying learning algorithms or the process of selecting an algorithm which will perform best in a particular problem domain [22]. Yudkowsky calls such process *non-recursive optimization* – a situation in which one component of the system does the optimization and another component is getting optimized [23].

In the field of complex dynamic systems, aka chaos theory, positive feedback systems are well known to always end up in what is known as an *attractor*- a region within system’s state space that the system can’t escape from [24]. A good example of such attractor convergence is the process of Metacompilation or Supercompilation [25] in which a program designed to take source code written by a human programmer and to optimize it for speed is applied to its own source code. It will likely produce a more efficient compiler on the first application perhaps by 20%, on the second application by 3%, and after a few more recursive iterations converge to a fixed point of zero improvement [24].

2.3 Recursive Self-Improvement

Recursive Self-Improvement is the only type of improvement which has potential to completely replace the original algorithm with a completely different approach and more importantly to do so multiple times. At each stage newly created software should be better at optimizing future version of the software compared to the original algorithm. As of the time of this writing it is a purely theoretical concept with no working RSI software known to exist. However, as many have predicted that such software might become a reality in the 21st century it is important to provide some analysis of properties such software would exhibit.

Self-modifying and self-improving software systems are already well understood and are quite common. Consequently, we will concentrate exclusively on RSI systems. In practice performance of almost any system can be trivially improved by allocation of additional computational resources such as more memory, higher sensor resolution, faster processor or greater network bandwidth for access to information. This linear scaling doesn’t fit the definition of recursive-improvement as the system doesn’t become better at improving itself. To fit the definition the system would have to engineer a faster type of memory not just purchase more memory units of the type it already has access to. In general hardware improvements are likely to speed up the

system, while software improvements (novel algorithms) are necessary for achievement of meta-improvements.

It is believed that AI systems will have a number of advantages over human programmers making it possible for them to succeed where we have so far failed. Such advantages include [26]: longer work spans (no breaks, sleep, vocation, etc.), omniscience (expert level knowledge in all fields of science, absorbed knowledge of all published works), superior computational resources (brain vs processor, human memory vs RAM), communication speed (neurons vs wires), increased serial depth (ability to perform sequential operations in access of about a 100 human brain can manage), duplicability (intelligent software can be instantaneously copied), editability (source code unlike DNA can be quickly modified), goal coordination (AI copies can work towards a common goal without much overhead), improved rationality (AIs are likely to be free from human cognitive biases) [27], new sensory modalities (native sensory hardware for source code), blending over of deliberative and automatic processes (management of computational resources over multiple tasks), introspective perception and manipulation (ability to analyze low level hardware, ex. individual neurons), addition of hardware (ability to add new memory, sensors, etc.), advanced communication (ability to share underlying cognitive representations for memories and skills) [28].

Chalmers [29] uses logic and mathematical induction to show that if an AI_0 system is capable of producing only slightly more capable AI_1 system generalization of that process leads to superintelligent performance in AI_n after n generations. He articulates, that his proof assumes that the *proportionality thesis*, which states that increases in intelligence lead to proportionate increases in the capacity to design future generations of AIs, is true.

Nivel et al. proposed formalization of RSI systems as autocatalytic sets – collections of entities comprised of elements, each of which can be created by other elements in the set making it possible for the set to self-maintain and update itself. They also list properties of a system which make it purposeful, goal-oriented and self-organizing, particularly: *reflectivity* – ability to analyze and rewrite its own structure; *autonomy* – being free from influence by system's original designers (*bounded autonomy* – is a property of a system with elements which are not subject to self-modification); *endogeny* – an autocatalytic ability [30]. Nivel and Thorisson also attempt to operationalize autonomy by the concept of *self-programming* which they insist has to be done in an experimental way instead of a theoretical way (via proofs of correctness) since it is the only tractable approach [31].

Yudkowsky writes prolifically about recursive self-improving processes and suggests that introduction of certain concepts might be beneficial to the discussion, specifically he proposes use of terms - Cascades, Cycles and Insight which he defines as: Cascades – when one development leads to another; Cycles – repeatable cascade in which one optimization leads to another which in turn benefits the original optimization; Insight – new information which greatly increases one's optimization ability [32]. Yudkowsky also suggests that the goodness and number of opportunities in the space of solutions be known as *Optimization Slope* while *optimization resources* and *optimization efficiency* refer to how much of computational resources an agent has access to and how efficiently the agent utilizes said resources. An agent engaging in

an *optimization process* and able to hit non-trivial targets in large search space [33] is described as having significant optimization power [23].

3 RSI Software Classification

RSI software could be classified based on the number of improvements it is capable of achieving. The most trivial case is the system capable of undergoing a single fundamental improvement. The hope is that truly RSI software will be capable of many such improvements, but the question remains open regarding the possibility of an infinite number of recursive-improvements. It is possible that some upper bound on improvements exists limiting any RSI software to a finite number of desirable and significant rewrites. Critics explain failure of scientists, to date, to achieve a sustained RSI process by saying that RSI researchers have fallen victims of the *bootstrap fallacy* [34].

3.1 How Improvements are Discovered

Another axis on which RSI systems can be classified has to do with how improvements are discovered. Two fundamentally different approaches are understood to exist. The first one is a brute force based approach [35] which utilizes Levin (Universal [36]) Search [37]. The idea is to consider all possible strings of source code up to some size limit and to select the one which can be proven to provide improvements. While theoretically optimal and guaranteed to find superior solution if one exists this method is not computationally feasible in practice. Some variants of this approach to self-improvement, known as Gödel Machines [38-43], Optimal Ordered Problem Solver (OOPS) [44] and Incremental Self-Improvers [45, 46], have been thoroughly analyzed by Schmidhuber and his co-authors. Second approach assumes that the system has a certain level of scientific competence and uses it to engineer and test its own replacement. Whether a system of any capability can intentionally invent a more capable and so a more complex system remains as the fundamental open problem of RSI research.

It is important to note that the first concrete algorithms for RSI were all by Schmidhuber. His diploma thesis from 1987 already was about an evolutionary system that learns to inspect and improve its own learning algorithm, where Genetic Programming (GP) is applied to itself, to recursively evolve better GP methods. His RSI based on the self-referential Success-Story Algorithm for self-modifying probabilistic programs was already able to solve complex tasks [47]. And finally, his self-referential recurrent neural networks run and inspect and change their own weight change algorithms [48]. In 2001, his former student Hochreiter had actually a practical implementation of such an RNN that learns an excellent learning algorithm, at least for the limited domain of quadratic functions [49, 50].

3.2 Hybrid Systems

Finally, we can consider a hybrid RSI system which includes both an artificially intelligent program and a human scientist. Mixed human-AI teams have been very successful in many domains such as chess or theorem proving. It would be surprising if having a combination of natural and artificial intelligence did not provide an advantage in designing new AI systems or enhancing biological intelligence. We are currently experiencing a limited version of this approach with human computer scientists developing progressively better versions of AI software (while utilizing continuously improving software tools), but since the scientists themselves remain unenhanced we can't really talk about self-improvement. This type of RSI can be classified as Indirect recursive improvement as opposed to Direct RSI in which the system itself is responsible for all modifications. Other types of Indirect RSI may be based on collaboration between multiple artificial systems instead of AI and human teams [51].

3.3 Other Properties

In addition to classification with respect to types of RSI we can also evaluate systems as to certain binary properties. For example: We may be interested only in systems which are guaranteed not to decrease in intelligence, even temporarily, during the improvement process. This may not be possible if the intelligence design landscape contains local maxima points.

Another property of any RSI system we are interested in understanding better is necessity of unchanging source code segments. In other words must an RSI system be able to modify any part of its source code or are certain portions of the system (encoded goals, verification module) must remain unchanged from generation to generation. Such portions would be akin to ultra-conserved elements or conserved sequences of DNA [52, 53] found among multiple related species. This question is particularly important for the goal preservation in self-improving intelligent software, as we want to make sure that future generations of the system are motivated to work on the same problem [29]. As AI goes through the RSI process and becomes smarter and more rational it is likely to engage in a de-biasing process removing any constraints we programmed into it [8]. Ideally we would want to be able to prove that even after recursive self-improvement our algorithm maintains the same goals as the original. Proofs of safety or correctness for the algorithm only apply to particular source code and would need to be rewritten and re-proven if the code is modified, which happens in RSI software many times. But we suspect that re-proving slightly modified code may be easier compared to having to prove safety of a completely novel piece of code.

We are also interested in understanding if RSI process can take place in an isolated (leakproofed [54]) system or if interaction with external environment, internet, people, other AI agents is necessary. Perhaps access to external information can be used to mediate speed of RSI process. This also has significant implications on safety mechanisms we can employ while experimenting with early RSI systems [55-63]. Finally, it needs to be investigated if the whole RSI process can be paused at any point and for any specific duration of time in order to limit any negative impact from

potential intelligence explosion. Ideally we would like to be able to program our Seed AI to RSI until it reaches certain level of intelligence, pause and wait for further instructions.

4 Conclusions

Recursively Self-Improving software is the ultimate form of artificial life and creation of life remains one of the great unsolved mysteries in science. More precisely, the problem of creating RSI software is really the challenge of creating a program capable of writing other programs [64], and so is an AI-Complete problem as has been demonstrated by Yampolskiy [65, 66]. AI-complete problems are by definition most difficult problems faced by AI researchers and it is likely that RSI source code will be so complex that it would be difficult or impossible to fully analyze [51]. Also, the problem is likely to be NP-Complete as even simple metareasoning and metalearning [67] problems have been shown by Conitzer and Sandholm to belong to that class. In particular they proved that allocation of deliberation time across anytime algorithms running on different problem instances is NP-Complete and a complimentary problem of dynamically allocating information gathering resources by an agent across multiple actions is NP-Hard, even if evaluating each particular action is computationally simple. Finally, they showed that the problem of deliberately choosing a limited number of deliberation or information gathering actions to disambiguate the state of the world is PSPACE Hard in general [68].

This paper is a part of a two paper set presented at AGI2015 with the complementary paper being: “On the Limits of Recursively Self-Improving AGI” [69].

References

1. Turing, A., *Computing Machinery and Intelligence*. Mind, 1950. **59(236)**: p. 433-460.
2. Good, I.J., *Speculations Concerning the First Ultrainelligent Machine*. Advances in Computers, 1966. **6**: p. 31-88.
3. Minsky, M., *Artificial Intelligence*. Scientific American, 1966. **215(3)**: p. 257.
4. Burks, A.W. and J. Von Neumann, *Theory of self-reproducing automata*. 1966: University of Illinois Press.
5. Pearce, D., *The biointelligence explosion*, in *Singularity Hypotheses*. 2012, Springer. p. 199-238.
6. Omohundro, S.M., *The Nature of Self-Improving Artificial Intelligence*, in *Singularity Summit*. 2007: San Francisco, CA.
7. Waser, M.R., *Bootstrapping a Structured Self-Improving & Safe Autopoietic Self*, in *Annual International Conference on Biologically Inspired Cognitive Architectures*. November 9, 2014: Boston, Massachusetts.
8. Hall, J.S., *Engineering utopia*. Frontiers in Artificial Intelligence and Applications, 2008. **171**: p. 460.
9. Mavrogiannopoulos, N., N. Kisserli, and B. Preneel, *A taxonomy of self-modifying code for obfuscation*. Computers & Security, 2011. **30(8)**: p. 679-691.
10. Anckaert, B., M. Madou, and K. De Bosschere, *A model for self-modifying code*, in *Information Hiding*. 2007, Springer. p. 232-248.

11. Petrean, L., *Polymorphic and Metamorphic Code Applications in Portable Executable Files Protection*. Acta Technica Napocensis, 2010. **51**(1).
12. Bonfante, G., J.-Y. Marion, and D. Reynaud-Plantey, *A computability perspective on self-modifying programs*, in *Seventh IEEE International Conference on Software Engineering and Formal Methods*. 2009, IEEE. p. 231-239.
13. Cheng, B.H., et al., *Software engineering for self-adaptive systems: A research roadmap*, in *Software engineering for self-adaptive systems*. 2009, Springer. p. 1-26.
14. Ailon, N., et al., *Self-improving algorithms*. SIAM Journal on Computing, 2011. **40**(2): p. 350-375.
15. Yampolskiy, R., et al., *Printer Model Integrating Genetic Algorithm for Improvement of Halftone Patterns*, in *Western New York Image Processing Workshop (WNYIPW) -IEEE Signal Processing Society*. 2004: Rochester, NY.
16. Yampolskiy, R.V., L. Ashby, and L. Hassan, *Wisdom of Artificial Crowds—A Metaheuristic Algorithm for Optimization*. Journal of Intelligent Learning Systems and Applications, 2012. **4**(2): p. 98-107.
17. Yampolskiy, R.V. and E.L.B. Ahmed, *Wisdom of artificial crowds algorithm for solving NP-hard problems*. International Journal of Bio-Inspired Computation (IJBIC). **3**(6): p. 358-369.
18. Ashby, L.H. and R.V. Yampolskiy, *Genetic Algorithm and Wisdom of Artificial Crowds Algorithm Applied to Light Up*, in *16th International Conference on Computer Games: AI, Animation, Mobile, Interactive Multimedia, Educational & Serious Games*. July 27 - 30, 2011 Louisville, KY, USA p. 27-32.
19. Khalifa, A.B. and R.V. Yampolskiy, *GA with Wisdom of Artificial Crowds for Solving Mastermind Satisfiability Problem*. International Journal of Intelligent Games & Simulation, 2011. **6**(2): p. 6.
20. Port, A.C. and R.V. Yampolskiy, *Using a GA and Wisdom of Artificial Crowds to solve solitaire battleship puzzles*, in *17th International Conference on Computer Games (CGAMES)*. 2012, IEEE: Louisville, KY. p. 25-29.
21. Omohundro, S., *Rational artificial intelligence for the greater good*, in *Singularity Hypotheses*. 2012, Springer. p. 161-179.
22. Anderson, M.L. and T. Oates, *A review of recent research in metareasoning and metalearning*. AI Magazine, 2007. **28**(1): p. 12.
23. Yudkowsky, E., *Intelligence Explosion Microeconomics*, in *MIRI Technical Report*. Available at: www.intelligence.org/files/IEM.pdf.
24. Heylighen, F., *Brain in a vat cannot break out*. Journal of Consciousness Studies, 2012. **19**(1-2): p. 1-2.
25. Turchin, V.F., *The concept of a supercompiler*. ACM Transactions on Programming Languages and Systems (TOPLAS), 1986. **8**(3): p. 292-325.
26. Sotala, K., *Advantages of artificial intelligences, uploads, and digital minds*. International Journal of Machine Consciousness, 2012. **4**(01): p. 275-291.
27. Muehlhauser, L. and A. Salamon, *Intelligence explosion: Evidence and import*, in *Singularity Hypotheses*. 2012, Springer. p. 15-42.
28. Yudkowsky, E., *Levels of organization in general intelligence*, in *Artificial general intelligence*. 2007, Springer. p. 389-501.
29. Chalmers, D., *The Singularity: A Philosophical Analysis*. Journal of Consciousness Studies, 2010. **17**: p. 7-65.
30. Nivel, E., et al., *Bounded Recursive Self-Improvement*. arXiv preprint arXiv:1312.6764, 2013.
31. Nivel, E. and K.R. Thórisson. *Self-programming: Operationalizing autonomy*. in *Proceedings of the 2nd Conf. on Artificial General Intelligence*. 2008.
32. Yudkowsky, E. and R. Hanson, *The Hanson-Yudkowsky AI-foom debate*, in *MIRI Technical Report*. 2008: Available at: <http://intelligence.org/files/AIFoomDebate.pdf>.

33. Yampolskiy, R.V., *The Universe of Minds*. arXiv preprint arXiv:1410.0369, 2014.
34. Hall, J.S., *Self-improving AI: An analysis*. Minds and Machines, 2007. **17**(3): p. 249-259.
35. Yampolskiy, R.V., *Efficiency Theory: a Unifying Theory for Information, Computation and Intelligence*. Journal of Discrete Mathematical Sciences & Cryptography, 2013. **16**(4-5): p. 259-277.
36. Gagliolo, M., *Universal search*. Scholarpedia, 2007. **2**(11):2575.
37. Levin, L., *Universal Search Problems*. Problems of Information Transmission, 1973. **9**(3): p. 265--266.
38. Steunebrink, B. and J. Schmidhuber, *A Family of Gödel Machine Implementations*, in *Fourth Conference on Artificial General Intelligence (AGI-11)*. 2011: Mountain View, California.
39. Schmidhuber, J., *Gödel machines: Fully self-referential optimal universal self-improvers*, in *Artificial general intelligence*. 2007, Springer. p. 199-226.
40. Schmidhuber, J., *Gödel machines: Towards a technical justification of consciousness*, in *Adaptive Agents and Multi-Agent Systems II*. 2005, Springer. p. 1-23.
41. Schmidhuber, J. *Gödel machines: Self-referential universal problem solvers making provably optimal self-improvements*. in *Artificial General Intelligence*. 2005.
42. Schmidhuber, J., *Ultimate cognition à la Gödel*. Cognitive Computation, 2009. **1**(2): p. 177-193.
43. Schmidhuber, J., *Completely self-referential optimal reinforcement learners*, in *Artificial Neural Networks: Formal Models and Their Applications—ICANN 2005*. 2005, Springer. p. 223-233.
44. Schmidhuber, J., *Optimal ordered problem solver*. Machine Learning, 2004. **54**(3): p. 211-254.
45. Schmidhuber, J., J. Zhao, and M. Wiering, *Shifting inductive bias with success-story algorithm, adaptive Levin search, and incremental self-improvement*. Machine Learning, 1997. **28**(1): p. 105-130.
46. Schmidhuber, J., *A general method for incremental self-improvement and multiagent learning*. Evolutionary Computation: Theory and Applications, 1999: p. 81-123.
47. Schmidhuber, J., *Metalearning with the Success-Story Algorithm*. 1997: Available at: <http://people.idsia.ch/~juergen/ssa/sld001.htm>.
48. Schmidhuber, J., *A neural network that embeds its own meta-levels*, in *IEEE International Conference on Neural Networks*. 1993, IEEE. p. 407-412.
49. Younger, A.S., S. Hochreiter, and P.R. Conwell, *Meta-learning with backpropagation*, in *International Joint Conference on Neural Networks (IJCNN'01)* 2001, IEEE.
50. Hochreiter, S., A. Younger, and P. Conwell, *Learning to learn using gradient descent*. Artificial Neural Networks—ICANN 2001, 2001: p. 87-94.
51. Leon, J. and A. Lori, *Continuous self-evaluation for the self-improvement of software*, in *Self-Adaptive Software*. 2001, Springer. p. 27-39.
52. Beck, M.B., E.C. Rouchka, and R.V. Yampolskiy, *Finding Data in DNA: Computer Forensic Investigations of Living Organisms*, in *Digital Forensics and Cyber Crime*. 2013, Springer Berlin Heidelberg. p. 204-219.
53. Beck, M. and R. Yampolskiy, *DNA as a medium for hiding data*. BMC Bioinformatics, 2012. **13**(Suppl 12): p. A23.
54. Yampolskiy, R.V., *Leakproofing Singularity - Artificial Intelligence Confinement Problem*. Journal of Consciousness Studies (JCS), 2012. **19**(1-2): p. 194–214.
55. Majot, A.M. and R.V. Yampolskiy. *AI safety engineering through introduction of self-reference into felicific calculus via artificial pain and pleasure*. in *Ethics in Science, Technology and Engineering, 2014 IEEE International Symposium on*. 2014. IEEE.
56. Yampolskiy, R. and J. Fox, *Safety Engineering for Artificial General Intelligence*. Topoi, 2012: p. 1-10.

57. Yampolskiy, R.V. and J. Fox, *Artificial General Intelligence and the Human Mental Model*. Singularity Hypotheses: A Scientific and Philosophical Assessment, 2013: p. 129.
58. Sotala, K. and R.V. Yampolskiy, *Responses to catastrophic AGI risk: A survey*. Physica Scripta. Vol 90, December 2015.
59. Yampolskiy, R.V., *What to Do with the Singularity Paradox?*, in *Philosophy and Theory of Artificial Intelligence*. 2013, Springer Berlin Heidelberg. p. 397-413.
60. Yampolskiy, R. and M. Gavrilova, *Artimetrics: Biometrics for Artificial Entities*. IEEE Robotics and Automation Magazine (RAM), 2012. **19**(4): p. 48-58.
61. Yampolskiy, R., et al., *Experiments in Artimetrics: Avatar Face Recognition*. Transactions on Computational Science XVI, 2012: p. 77-94.
62. Ali, N., D. Schaeffer, and R.V. Yampolskiy, *Linguistic Profiling and Behavioral Drift in Chat Bots*. Midwest Artificial Intelligence and Cognitive Science Conference, 2012: p. 27.
63. Gavrilova, M. and R. Yampolskiy, *State-of-the-Art in Robot Authentication [From the Guest Editors]*. Robotics & Automation Magazine, IEEE, 2010. **17**(4): p. 23-24.
64. Hall, J.S., *VARIAC: an Autogenous Cognitive Architecture*. Frontiers in Artificial Intelligence and Applications, 2008. **171**: p. 176.
65. Yampolskiy, R., *Turing Test as a Defining Feature of AI-Completeness*, in *Artificial Intelligence, Evolutionary Computing and Metaheuristics*, X.-S. Yang, Editor. 2013, Springer Berlin Heidelberg. p. 3-17.
66. Yampolskiy, R.V., *AI-Complete, AI-Hard, or AI-Easy—Classification of Problems in AI*. The 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Cincinnati, OH, USA, 2012.
67. Schaul, T. and J. Schmidhuber, *Metalearning*. Scholarpedia, 2010. **5**(6):**4650**.
68. Conitzer, V. and T. Sandholm, *Definition and complexity of some basic metareasoning problems*, in *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*. 2003: Acapulco, Mexico. p. 1099–1106.
69. Yampolskiy, R.V., *On the Limits of Recursively Self-Improving AGI*, in *The Eighth Conference on Artificial General Intelligence*. July 22-25, 2015: Berlin, Germany.