

University of Louisville

## ThinkIR: The University of Louisville's Institutional Repository

---

Faculty Scholarship

---

1-2019

### Personal Universes: A Solution to the Multi-Agent Value Alignment Problem

Roman V. Yampolskiy

University of Louisville, [roman.yampolskiy@louisville.edu](mailto:roman.yampolskiy@louisville.edu)

Follow this and additional works at: <https://ir.library.louisville.edu/faculty>



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Computer Engineering Commons](#)

---

#### Original Publication Information

Yampolskiy, R., V. (2019). Personal Universes: A Solution to the Multi-Agent Value Alignment Problem. *arXiv preprint arXiv:1901.01851*

#### ThinkIR Citation

Yampolskiy, Roman V., "Personal Universes: A Solution to the Multi-Agent Value Alignment Problem" (2019). *Faculty Scholarship*. 424.  
<https://ir.library.louisville.edu/faculty/424>

This Article is brought to you for free and open access by ThinkIR: The University of Louisville's Institutional Repository. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of ThinkIR: The University of Louisville's Institutional Repository. For more information, please contact [thinkir@louisville.edu](mailto:thinkir@louisville.edu).

# Personal Universes:

## A Solution to the Multi-Agent Value Alignment Problem

**Roman V. Yampolskiy**  
Computer Engineering and Computer Science  
Speed School of Engineering  
University of Louisville  
[roman.yampolskiy@louisville.edu](mailto:roman.yampolskiy@louisville.edu)

### Abstract

AI Safety researchers attempting to align values of highly capable intelligent systems with those of humanity face a number of challenges including personal value extraction, multi-agent value merger and finally in-silico encoding. State-of-the-art research in value alignment shows difficulties in every stage in this process, but merger of incompatible preferences is a particularly difficult challenge to overcome. In this paper we assume that the value extraction problem will be solved and propose a possible way to implement an AI solution which optimally aligns with individual preferences of each user. We conclude by analyzing benefits and limitations of the proposed approach.

**Keywords:** AI Safety, Alternate Reality, Simulation, Value Alignment Problem, VR

### 1. Introduction to the Multi-Agent Value Alignment Problem

Since the birth of the field of Artificial Intelligence (AI) researchers worked on creating ever capable machines, but with recent success in multiple subdomains of AI [1-7] safety and security of such systems and predicted future superintelligences [8, 9] has become paramount [10, 11]. While many diverse safety mechanisms are being investigated [12, 13], the ultimate goal is to align AI with goals, values and preferences of its users which is likely to include all of humanity.

Value alignment problem [14], can be decomposed into three sub-problems, namely: personal value extraction from individual persons, combination of such personal preferences in a way, which is acceptable to all, and finally production of an intelligent system, which implements combined values of humanity.

A number of approaches for extracting values [15-17] from people have been investigated, including inverse reinforcement learning [18, 19], brain scanning [20], value learning from literature [21], and understanding of human cognitive limitations [22]. Assessment of potential for success for particular techniques of value extraction is beyond the scope of this paper and we simply assume that one of the current methods, their combination, or some future approach will allow us to accurately learn values of given people. Likewise, we will not directly address how, once learned, such values can be represented/encoded in computer systems for storage and

processing. These assumptions free us from having to worry about safety problems with misaligned AIs such as perverse instantiation or wireheading [23], among many others [24].

The second step in the process requires an algorithm for value aggregation from some and perhaps even all people to assure that the developed AI is beneficial to the humanity as a whole. Some have suggested that interests of future people [25], potential people [26] and of non-human animals and other sentient beings, be likewise included in our “Coherent Extrapolated Volition” (CEV) [27], which we would like superintelligent AI to eventually implement. However, work done by moral philosophers over hundreds of years indicates that our moral preferences are not only difficult to distil in a coherent manner (anti-codifiability thesis) [28], they are also likely impossible to merge without sacrificing interests of some people [29, 30], we can say it is the Hard problem of value alignment. Results from research into multivariate optimization and voting based preference aggregation support similar conclusions [31-33].

Perhaps we should stop trying to make “one size fits all” approach to the optimization of the universe work and instead look at potential for delivering an experience customized to individual users. The superintelligent systems we are hoping to one day create, with the goal of improving lives of all, may work best if instead they strive to optimize their alignment with individual lives of each and every one of us, while giving us all freedom to be ourselves without infringing on preferences of other sentient [34, 35] beings. Such a system due to its lower overall complexity should also be easier to design, implement and safeguard.

## **2. Individual Simulated Universes**

It has been suggested that future technology will permit design [36] and instantiation of high fidelity simulated universes [37-41] for research and entertainment ([42], chapter 5) purposes as well as for testing advanced AIs [43-46]. Existing work and recent breakthroughs in virtual reality, augmented reality, inter-reality, haptics, and artificial consciousness combined with tremendous popularity of multiplayer virtual worlds such as Second Life [47-49] or Ultima Online [50] provide encouraging evidence for the plausibility of realistic simulations.

We can foresee, in a not so distant future, a point at which visual and audio fidelity of the simulations, as well as for all other senses [51] becomes so high that it will not be possible to distinguish if you are in a base reality or in a simulated world, frequently referred as hyperreality [52, 53]. In principle, it should be possible to improve local fidelity (measurable by the agent) of the simulated reality to levels beyond base reality, for example to the point of more precise measurements being possible with special instrumentation. This would effectively reverse the resolution relationship between the two realities making the base reality less believable on local scale. A variant of a Total Turing Test [54, 55], we shall call a Universal Turing Test (UTT) could be administered in which the user tries to determine if the current environment is synthetic or not [56] even if it is complex enough to include the whole universe, all other beings (as philosophical zombies [57]/Non-Playing Characters (NPCs)) and AIs. Once the UTT is consistently passed we will know, the hyperreality is upon us.

Consequently, we suggest that instead of trying to agree on convergent, universal, diverse, mutually beneficial, equalizing, representative, unbiased, timeless, acceptable to all, etc. moral/ethical norms and values, predicated on compromise [58], we look at an obvious alternative.

Specifically, we suggest that superintelligent AIs should be implemented to act as personalized simulations - *Individual Simulated Universes* (ISU) representing customized synthetically generated [7, 59] mega-environments, in the “a universe per person multi-verse framework”, which are optimally and dynamically adjusting to align their values and preferences to the Personal CEV [60] of sentient agents calling such universes “home”.

Aaronson describes the general idea as “... an infinite number of sentient beings living in simulated paradises of their own choosing, racking up an infinite amount of utility. If such a being wants challenge and adventure, then challenge and adventure is what it gets; if nonstop sex, then nonstop sex; if a proof of  $P \neq NP$ , then a proof of  $P \neq NP$ . (Or the being could choose all three: it’s utopia, after all!)” [61]. Bostrom estimates that our galactic supercluster has enough energy to support trillions of such efficiently [62] simulated universes [63]. Features of related phenomenon have been described in literature as [64]: dematerialization [65], ephemeralization [66], time-space compression [67], miniaturization [68], densification [69], virtualization [70], digitization [71], and simulation [72].

Faggella talks about opportunities presented in the virtual world over what is possible in the present reality [73]: “... ‘freedom’ could only extend so far in a real world as to border on impinging on the ‘freedom’ of others. Complete freedom would imply control over one’s environment and free choice to do what one would chose with it. It seems easy to understand how this might imply the threatening of the freedom of others in the same physical world. ... Not to mention, the physical world has many impinging qualities that would hinder any semblance of complete freedom. Matter has qualities, light has qualities, and physical bodies (no matter how enhanced) will always have limitations. If you’d like to change an aspect of our character or emotional experience, for example, we’d have to potentially tinker with brain chemicals ... . In a virtual reality, we are potentially presented not only with the freedom to extend beyond physical limitations (to transport to different times or places, to live within self-created fantasy worlds, to eliminate death and any physical risk), we would also be granted freedom from impinging or effecting others – and so allow for their full freedom an a separate virtual reality as well. ... For this reason, it seems to make sense that ... we might encounter a Bostrom-like ‘Singleton’ to rule the physical world, and a great sea of individual consciousnesses in the virtual world. The ‘Singleton’ could keep our computational substrates safe from harm and eliminate competition or danger in the physical world, while our virtual ‘selves’ would be capable of expressing and exploring the epitome of freedom on our own terms in a limitless virtual world of our own creation.”

This means that an ISU can be anything a user truly wishes it to be including dangerous, adversarial, competitive, and challenging at all levels of user competence like levels in a well-designed video game. It will let a user be anything they want to be including a malevolent actor [74, 75], a privileged person (like a king) or the exact opposite (a slave), or perhaps just a selfish user in an altruistic universe. A personalized universe doesn’t have to be fair, or just or free of perceived suffering and pain [76]. It could be just a sequence of temporary fantasies and hopefully what happens in your personalized universe stays in your personalized universe. ISU’s goal is to cater to the world’s smallest minority and its preferences, you [77, 78]! Moreover, the good news is that we know that we are not going to run out of Fun [79] even if we live much longer lives [80].

If an agent controlling the environment is not well aligning with a particular individual for whom the environment is created (during early stages of development of this technology) it may be necessary to use precise language to express what the user wants. The now defunct Open-Source Wish Project (OSWP) [81] attempted to formulate in precise and safe form such common wishes as: immortality, happiness, omniscience, being rich, having true love, omnipotence, etc [23].

For example the latest version of the properly formed request for immortality was formalized as follows: “I wish to live in the locations of my choice, in a physically healthy, uninjured, and apparently normal version of my current body containing my current mental state, a body which will heal from all injuries at a rate three sigmas faster than the average given the medical technology available to me, and which will be protected from any diseases, injuries or illnesses causing disability, pain, or degraded functionality or any sense, organ, or bodily function for more than ten days consecutively or fifteen days in any year; at any time I may rejuvenate my body to a younger age, by saying a phrase matching this pattern five times without interruption, and with conscious intent: 'I wish to be age,' followed by a number between one and two hundred, followed by 'years old,' at which point the pattern ends - after saying a phrase matching that pattern, my body will revert to an age matching the number of years I started and I will commence to age normally from that stage, with all of my memories intact; at any time I may die, by saying five times without interruption, and with conscious intent, 'I wish to be dead'; the terms 'year' and 'day' in this wish shall be interpreted as the ISO standard definitions of the Earth year and day as of 2006. [81]” Of course, this is still far from foolproof and is likely to lead to some undesirable situations, which could be avoided by development of a well-aligned system.

### **3. Benefits and Shortcomings of Personalized Universes**

ISUs can be implemented in a number of ways, either by having perfect emulations of agents reside in the simulated universe or by having current biological agents experience fully realistic simulated environments (while robotic systems take care of their bodies' biological needs), see Faggella's review of possible variants of virtual reality [82]. Both options have certain desirable properties, for example, software versions of users are much easier to modify, reset to earlier memory states [83], upgrade and backup [84, 85], while biological agents are likely to have stronger identity continuity [86]. Emulations can also be taken as snapshots from different points in the person's life and set to exist in their own independent simulations multiplying possible experiences [34] for the subset of agents derived from that particular individual. In both virtual and uploaded scenarios, it is probably desirable for the user to “forget” that they are not in the base reality via some technological means with the goal of avoiding Solipsism syndrome<sup>1</sup>.

Our proposal doesn't just allow us to bypass having to find a difficult to compute approximation to a likely impossible to solve problem of multi-agent value aggregation, but it also provides for a much better “customer experience” free of compromise on even small details which may be important to that individual. Additionally, virtual existence makes it possible to have an “undo button” for actions/experiences user might regret, something not always possible in the world of physical reality. Last, but not least any existential risks related to this particular AIs failure are limited to the simulated universe and its virtual inhabitants, not to the humanity and all life forms.

---

<sup>1</sup> [https://en.wikipedia.org/wiki/Solipsism\\_syndrome](https://en.wikipedia.org/wiki/Solipsism_syndrome)

Of course, like any AI safety mechanism ours has certain weaknesses, which will have to be explicitly addressed. Those include having to withstand agents with extreme preferences, who may wish to prevent others from exercising their self-determination and may attempt to hack and sabotage ISUs or even base reality (which should be easier to secure, with most agents and their complex preferences out of the way). Another area of concern is problems with superintelligence serving as “operating system” for the base reality and allocating non-conflicting resources for the ISUs. Finally, we should study how the philosophical questions of living in a “fake” world vs “real” world, even if it is not possible to distinguish between them by any means, impacts human psychology and well-being.

It is also important to figure out a metric to measure user-relative quality of the simulation experience not just from fidelity point of view but also from users overall satisfaction with how their values, goals and preferences are being serviced, such metrics are notoriously hard to design and easy to abuse [87]. Potential ideas may include user feedback both from within the simulation and while outside observing a recording of themselves in the simulation, feedback after trying other simulations and potentially all other simulations, and peer-review from other conscious agents both from outside and from within the same environment.

It is possible to let users “play” in other’s universes and perhaps as other characters and to allow them to discover and integrate new values to which their universe will dynamically adopt. It may also be possible for two or more agents to decide to cohabit the same universe by coming to accept a mutually satisfying set of values, but of course their individual alignment with the environment would be reduced and so it is important to provide them with a “divorce” option. We are assuming a well aligned AI, which will not attempt to directly hack the agent to game the feedback score, but out of caution, we do not recommend evolutionary competition [88-90] between ISUs as that can lead to adversarial behaviors between superintelligent agents even the base reality superintelligence would not be able to resolve.

#### **4. Conclusions**

In this exploratory paper, we advocated a solution to the hardest of the three subproblems of multi-agent value alignment, specifically value aggregation. Our “in the box” solution suggests replacing one-size-fits-all model of value satisfaction with customized and highly optimized approach which is strictly superior for all possible agents not valuing decreasing quality of value alignment for other agents. Some existing evidence from cosmology may be seen as suggesting that perhaps this approach is not so novel and in fact has already been implemented by earlier civilizations, and this universe is already a part of a multiverse [91, 92] generated by intelligence [93]. While some significant concerns with the philosophical [94], social [95] and security [96, 97] problems associated with personalized universes remain, particularly with regards to securing base reality, the proposal has a number of previously described advantages. Such advantages are likely to make it attractive to many users or to at least be integrated as a part of a more complex hybrid solution scheme. The decisions made by users of personal universes are also a goldmine of valuable data both for assessment of agents and for providing additional data to improve overall AI alignment [98]. We will leave proposals for assuring safety and security of cyberinfrastructure running personalized universes for future work. The main point of this paper is that a personal universe is a place where virtually everyone can be happy.

## Acknowledgments

The author is grateful to Elon Musk and the Future of Life Institute and to Jaan Tallinn and Effective Altruism Ventures for partially funding his work on AI Safety. Special thank you goes to all NPCs in this universe.

## References

1. Silver, D., et al., *A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play*. Science, 2018. 362(6419): p. 1140-1144.
2. Silver, D., et al., *Mastering the game of Go without human knowledge*. Nature, 2017. 550(7676): p. 354.
3. Mnih, V., et al., *Human-level control through deep reinforcement learning*. Nature, 2015. 518(7540): p. 529.
4. High, R., *The era of cognitive systems: An inside look at IBM Watson and how it works*. IBM Corporation, Redbooks, 2012.
5. Moravčik, M., et al., *Deepstack: Expert-level artificial intelligence in heads-up no-limit poker*. Science, 2017. 356(6337): p. 508-513.
6. Krizhevsky, A., I. Sutskever, and G.E. Hinton. *Imagenet classification with deep convolutional neural networks*. in *Advances in neural information processing systems*. 2012.
7. Goodfellow, I., et al. *Generative adversarial nets*. in *Advances in neural information processing systems*. 2014.
8. Bostrom, N., *Superintelligence: Paths, dangers, strategies*. 2014: Oxford University Press.
9. Yampolskiy, R.V., *AI-complete CAPTCHAs as zero knowledge proofs of access to an artificially intelligent system*. ISRN Artificial Intelligence, 2011. 2012.
10. Yampolskiy, R.V., *Artificial Superintelligence: a Futuristic Approach*. 2015: Chapman and Hall/CRC.
11. Yampolskiy, R.V., *Artificial Intelligence Safety and Security*. 2018: CRC Press.
12. Sotala, K. and R.V. Yampolskiy, *Responses to catastrophic AGI risk: a survey*. Physica Scripta, 2014. 90(1): p. 018001.
13. Everitt, T., G. Lea, and M. Hutter, *AGI Safety Literature Review*. arXiv preprint arXiv:1805.01109, 2018.
14. Soares, N. and B. Fallenstein, *Aligning superintelligence with human interests: A technical research agenda*. Machine Intelligence Research Institute (MIRI) technical report, 2014. 8.
15. Dignum, V., *Responsible Artificial Intelligence: Designing AI for Human Values*. ITU Journal: ICT Discoveries, 2017.
16. Evans, O., A. Stuhlmüller, and N.D. Goodman. *Learning the Preferences of Ignorant, Inconsistent Agents*. in *AAAI*. 2016.
17. Kim, T.W., T. Donaldson, and J. Hooker, *Mimetic vs Anchored Value Alignment in Artificial Intelligence*. arXiv preprint arXiv:1810.11116, 2018.
18. Ng, A.Y. and S.J. Russell. *Algorithms for inverse reinforcement learning*. in *Icml*. 2000.
19. Abbeel, P. and A.Y. Ng. *Apprenticeship learning via inverse reinforcement learning*. in *Proceedings of the twenty-first international conference on Machine learning*. 2004. ACM.
20. Sarma, G.P., N.J. Hay, and A. Safron. *AI Safety and Reproducibility: Establishing Robust Foundations for the Neuropsychology of Human Values*. in *International Conference on Computer Safety, Reliability, and Security*. 2018. Springer.

21. Riedl, M.O. and B. Harrison. *Using Stories to Teach Human Values to Artificial Agents*. in *AAAI Workshop: AI, Ethics, and Society*. 2016.
22. Trazzi, M. and R.V. Yampolskiy, *Building Safer AGI by introducing Artificial Stupidity*. arXiv preprint arXiv:1808.03644, 2018.
23. Yampolskiy, R.V., *Utility function security in artificially intelligent agents*. *Journal of Experimental & Theoretical Artificial Intelligence*, 2014. 26(3): p. 373-389.
24. Yampolskiy, R.V. *Taxonomy of Pathways to Dangerous Artificial Intelligence*. in *AAAI Workshop: AI, Ethics, and Society*. 2016.
25. Mulgan, T., *Future people: A moderate consequentialist account of our obligations to future generations*. OUP Catalogue, 2008.
26. Warren, M.A., *Do potential people have moral rights?* *Canadian Journal of Philosophy*, 1977. 7(2): p. 275-289.
27. Yudkowsky, E., *Coherent extrapolated volition*. Singularity Institute for Artificial Intelligence, 2004.
28. Purves, D., R. Jenkins, and B.J. Strawser, *Autonomous machines, moral judgment, and acting for the right reasons*. *Ethical Theory and Moral Practice*, 2015. 18(4): p. 851-872.
29. Yampolskiy, R.V., *Artificial intelligence safety engineering: Why machine ethics is a wrong approach*, in *Philosophy and Theory of Artificial Intelligence*. 2013, Springer Berlin Heidelberg. p. 389-396.
30. Sobel, D., *Full information accounts of well-being*. *Ethics*, 1994. 104(4): p. 784-810.
31. Arrow, K.J., *A difficulty in the concept of social welfare*. *Journal of political economy*, 1950. 58(4): p. 328-346.
32. Arrow, K.J., *Social choice and individual values*. Vol. 12. 2012: Yale university press.
33. Gehrlein, W.V., *Condorcet's paradox and the likelihood of its occurrence: different perspectives on balanced preferences*. *Theory and decision*, 2002. 52(2): p. 171-199.
34. Yampolskiy, R.V., *Detecting Qualia in Natural and Artificial Agents*. arXiv preprint arXiv:1712.04020, 2017.
35. Raoult, A. and R. Yampolskiy, *Reviewing Tests for Machine Consciousness*. Available at [https://www.researchgate.net/publication/284859013\\_DRAFT\\_Reviewing\\_Tests\\_for\\_Machine\\_Consciousness](https://www.researchgate.net/publication/284859013_DRAFT_Reviewing_Tests_for_Machine_Consciousness), 2015.
36. Knight, W., *AI software can dream up an entire digital world from a simple sketch*. December 3, 2018: Available at: <https://www.technologyreview.com/s/612503/ai-software-can-dream-up-an-entire-digital-world-from-a-simple-sketch>.
37. Bostrom, N., *Are we living in a computer simulation?* *The Philosophical Quarterly*, 2003. 53(211): p. 243-255.
38. Yampolskiy, R.V., *Future Jobs - The Universe Designer*, in *Circus Street*. 2017: Available at: <https://www.circusstreet.com/blog/future-jobs-the-universe-designer>.
39. Yampolskiy, R., *Job ad: universe designers*, in *Stories from 2045*, C. Chase, Editor. p. 50-53.
40. Chalmers, D.J., *The virtual and the real*. *Disputatio*, 2017. 9(46): p. 309-352.
41. Putnam, H., *Brain in a Vat*. *Reason, Truth and History*, 1981: p. 1-21.
42. Tegmark, M., *Life 3.0: Being human in the age of artificial intelligence*. 2017: Knopf.
43. Armstrong, S., A. Sandberg, and N. Bostrom, *Thinking inside the box: Controlling and using an oracle AI*. *Minds and Machines*, 2012. 22(4): p. 299-324.
44. Yampolskiy, R., *Leakproofing the Singularity Artificial Intelligence Confinement Problem*. *Journal of Consciousness Studies*, 2012. 19(1-2): p. 1-2.

45. Babcock, J., J. Kramár, and R.V. Yampolskiy, *Guidelines for Artificial Intelligence Containment*. arXiv preprint arXiv:1707.08476, 2017.
46. Babcock, J., J. Kramár, and R. Yampolskiy, *The AGI containment problem*, in *Artificial General Intelligence*. 2016, Springer. p. 53-63.
47. Boulos, M.N.K., L. Hetherington, and S. Wheeler, *Second Life: an overview of the potential of 3-D virtual worlds in medical and health education*. Health Information & Libraries Journal, 2007. 24(4): p. 233-245.
48. Yampolskiy, R.V. and M.L. Gavrilova, *Arithmetic: biometrics for artificial entities*. IEEE Robotics & Automation Magazine, 2012. 19(4): p. 48-58.
49. Yampolskiy, R.V., B. Klare, and A.K. Jain. *Face recognition in the virtual world: recognizing avatar faces*. in *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*. 2012. IEEE.
50. Simpson, Z.B. *The in-game economics of Ultima Online*. in *Computer Game Developer's Conference, San Jose, CA*. 2000.
51. Bushell, W.C. and M. Seaberg, *Experiments Suggest Humans Can Directly Observe the Quantum*, in *Psychology Today*. December 5, 2018: Available at: <https://www.psychologytoday.com/us/blog/sensorium/201812/experiments-suggest-humans-can-directly-observe-the-quantum>.
52. Baudrillard, J., *Simulacra and simulation*. 1994: University of Michigan press.
53. Eco, U., *Travels in hyper reality: essays*. 1990: Houghton Mifflin Harcourt.
54. Harnad, S., *The Turing Test is not a trick: Turing indistinguishability is a scientific criterion*. ACM SIGART Bulletin, 1992. 3(4): p. 9-10.
55. Schweizer, P., *The truly total Turing test*. Minds and Machines, 1998. 8(2): p. 263-272.
56. Yampolskiy, R.V., *On the origin of synthetic life: attribution of output to a particular algorithm*. Physica Scripta, 2016. 92(1): p. 013002.
57. Chalmers, D.J., *The conscious mind: In search of a fundamental theory*. 1996: Oxford university press.
58. Bostrom, N., *Moral uncertainty - towards a solution?*, in *Overcoming Bias*. 2009: Available at: <http://www.overcomingbias.com/2009/01/moral-uncertainty-towards-a-solution.html>.
59. Faggella, D., *Programmatically Generated Everything (PGE)*. August 27, 2018: Available at: <https://danfaggella.com/programmatically-generated-everything-pge/>.
60. Muehlhauser, L. and C. Williamson, *Ideal Advisor Theories and Personal CEV*. Machine Intelligence Research Institute, 2013.
61. *Visions of a Better World*, in *Scientific American*. December 19, 2018: Available at: <https://blogs.scientificamerican.com/cross-check/visions-of-a-better-world>.
62. 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.
63. Bostrom, N., *Astronomical waste: The opportunity cost of delayed technological development*. Utilitas, 2003. 15(3): p. 308-314.
64. Smart, J.M., *The transcension hypothesis: Sufficiently advanced civilizations invariably leave our universe, and implications for METI and SETI*. Acta Astronautica, 2012. 78: p. 55-68.
65. Wernick, I.K., et al., *Materialization and dematerialization: measures and trends*. Daedalus, 1996: p. 171-198.
66. Fuller, R.B., *Synergetics: explorations in the geometry of thinking*. 1982: Estate of R. Buckminster Fuller.

67. Harvey, D., *The condition of postmodernity*. Vol. 14. 1989: Blackwell Oxford.
68. Feynman, R. and D. Gilbert, *Miniaturization*. Reinhold, New York, 1961: p. 282-296.
69. Leskovec, J., J. Kleinberg, and C. Faloutsos. *Graphs over time: densification laws, shrinking diameters and possible explanations*. in *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. 2005. ACM.
70. Lévy, P. and R. Bononno, *Becoming virtual: Reality in the digital age*. 1998: Da Capo Press, Incorporated.
71. Negroponte, N., et al., *Being digital*. *Computers in Physics*, 1997. 11(3): p. 261-262.
72. Chalmers, D., *The Matrix as metaphysics*. *Science Fiction and Philosophy From Time Travel to Superintelligence*, 2003. 36.
73. Faggella, D., *Transhuman Possibilities and the "Epitome of Freedom"*. May 14, 2013: Available at: <https://danfaggella.com/transhuman-possibilities-and-the-epitome-of-freedom/>.
74. Pistono, F. and R.V. Yampolskiy, *Unethical Research: How to Create a Malevolent Artificial Intelligence*. arXiv preprint arXiv:1605.02817, 2016.
75. Brundage, M., et al., *The malicious use of artificial intelligence: Forecasting, prevention, and mitigation*. arXiv preprint arXiv:1802.07228, 2018.
76. Pearce, D., *Hedonistic Imperative*. 1995: David Pearce.
77. Rand, A., *The Ayn Rand lexicon: objectivism from A to Z*. Vol. 4. 1988: Penguin.
78. Rand, A., *The virtue of selfishness*. 1964: Penguin.
79. Ziesche, S. and R.V. Yampolskiy, *Artificial Fun: Mapping Minds to the Space of Fun*. arXiv preprint arXiv:1606.07092, 2016.
80. Kurzweil, R. and T. Grossman, *Fantastic voyage: live long enough to live forever*. 2005: Rodale.
81. Anonymous, *Wish For Immortality 1.1*, in *The Open-Source Wish Project*. 2006: Available at: <http://www.homeonthe strange.com/phpBB2/viewforum.php?f=4>.
82. Faggella, D., *The Transhuman Transition – Lotus Eaters vs World Eaters*. May 28, 2018: Available at: <https://danfaggella.com/the-transhuman-transition-lotus-eaters-vs-world-eaters/>.
83. Lebens, S. and T. Goldschmidt, *The Promise of a New Past*. 2017: Ann Arbor, MI: Michigan Publishing, University of Michigan Library.
84. Hanson, R., *The Age of Em: Work, Love, and Life when Robots Rule the Earth*. 2016: Oxford University Press.
85. Feygin, Y.B., K. Morris, and R.V. Yampolskiy, *Uploading Brain into Computer: Whom to Upload First?* arXiv preprint arXiv:1811.03009, 2018.
86. Parfit, D., *Reasons and persons*. 1984: OUP Oxford.
87. Manheim, D. and S. Garrabrant, *Categorizing Variants of Goodhart's Law*. arXiv preprint arXiv:1803.04585, 2018.
88. Lehman, J., J. Clune, and D. Misevic. *The Surprising Creativity of Digital Evolution*. in *Artificial Life Conference Proceedings*. 2018. MIT Press.
89. Lowrance, C.J., O. Abdelwahab, and R.V. Yampolskiy. *Evolution of a Metaheuristic for Aggregating Wisdom from Artificial Crowds*. in *Portuguese Conference on Artificial Intelligence*. 2015. Springer.
90. 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(02): p. 98.
91. Carr, B., *Universe or multiverse?* 2007: Cambridge University Press.

92. Vilenkin, A. and M. Tegmark, *The case for parallel universes*. Scientific American. Retrieved from: <http://www.scientificamerican.com/article/multiverse-the-case-for-parallel-universe>, 2011.
93. Gardner, J.N., *Biocosm: the new scientific theory of evolution: intelligent life is the architect of the universe*. 2003: Inner Ocean Publishing.
94. Vallentyne, P., *Robert Nozick: Anarchy, State, and Utopia*, in *Central Works of Philosophy* v5. 2014, Routledge. p. 108-125.
95. Turchin, A., *Wireheading as a Possible Contributor to Civilizational Decline*. 2018: Available at: <https://philpapers.org/rec/TURWAA>.
96. Faggella, D., *Substrate Monopoly – The Future of Power in a Virtual and Intelligent World*. August 17, 2018: Available at: <https://danfaggella.com/substrate-monopoly/>.
97. Faggella, D., *Digitized and Digested*. July 15, 2018: Available at: <https://danfaggella.com/digitized-and-digested/>.
98. Zhavoronkov, A., *Is Life A Recursive Video Game?*, in *Forbes*. December 12, 2018: Available at: <https://www.forbes.com/sites/cognitiveworld/2018/12/12/is-life-a-recursive-video-game>.