



Error-Correction for AI Safety

Nadisha-Marie Aliman¹(✉), Pieter Elands², Wolfgang Hürst¹, Leon Kester²,
Kristinn R. Thórisson⁴, Peter Werkhoven^{1,2}, Roman Yampolskiy³,
and Soenke Ziesche⁵

¹ Utrecht University, Utrecht, The Netherlands
nadishamarie.aliman@gmail.com

² TNO Netherlands, The Hague, The Netherlands

³ University of Louisville, Louisville, USA

⁴ Icelandic Institute for Intelligent Machines, Reykjavik University,
Reykjavik, Iceland

⁵ Delhi, India

Abstract. The complex socio-technological debate underlying safety-critical and ethically relevant issues pertaining to AI development and deployment extends across heterogeneous research subfields and involves in part conflicting positions. In this context, it seems expedient to generate a minimalistic joint transdisciplinary basis disambiguating the references to specific subtypes of AI properties and risks for an *error-correction* in the transmission of ideas. In this paper, we introduce a high-level *transdisciplinary system clustering* of ethical distinction between antithetical clusters of *Type I* and *Type II* systems which extends a cybersecurity-oriented AI safety taxonomy with considerations from psychology. Moreover, we review relevant Type I AI risks, reflect upon possible epistemological origins of hypothetical Type II AI from a cognitive sciences perspective and discuss the related human moral perception. Strikingly, our nuanced transdisciplinary analysis yields the figurative formulation of the so-called *AI safety paradox* identifying AI control and value alignment as conjugate requirements in AI safety. Against this backdrop, we craft versatile multidisciplinary recommendations with ethical dimensions tailored to Type II AI safety. Overall, we suggest proactive and importantly *corrective* instead of prohibitive methods as common basis for both Type I and Type II AI safety.

Keywords: AI safety paradox · Error-correction · AI ethics

1 Motivation

In recent years, one could identify the emergence of seemingly antagonistic positions from different academic subfields with regard to research priorities for AI safety, AI ethics and AGI – many of which are grounded in differences of short-term versus long-term estimations associated with AI capabilities and risks [6].

S. Ziesche—Independent Researcher.

© Springer Nature Switzerland AG 2020

B. Goertzel et al. (Eds.): AGI 2020, LNAI 12177, pp. 12–22, 2020.

https://doi.org/10.1007/978-3-030-52152-3_2

However, given the high relevance of the joint underlying endeavor to contribute to a safe and ethical development and deployment of artificial systems, we suggest placing a mutual comprehension in the foreground which can start by making references to assumed AI risks explicit. To this end, we employ and subsequently extend a cybersecurity-oriented risk taxonomy introduced by Yampolskiy [35] displayed in Fig. 1. Taking this taxonomy as point of departure and modifying it while considering insights from psychology, an ethically relevant clustering of systems into *Type I* and *Type II* systems with a disparate set of properties and risk instantiations becomes explicitly expressible. Concerning the set of Type I systems of which present-day AIs represent a subset, we define it as representing *the complement of the set of Type II systems*. Conversely, we regard hypothetical Type II systems as *systems with a scientifically plausible ability to act independently, intentionally, deliberately and consciously and to craft explanations*. Given the controversial ambiguities linked to these attributes, we clarify our idiosyncratic use with a working definition for which we do not claim any higher suitability in general, but which is particularly conceptualized for our line of argument. With Type II systems, we refer to systems having the ability to construct counterfactual hypotheses about what could happen, what could have happened, how *and why* including the ability to simulate “what *I* could do” “what *I* could have done” and the generation of “what if” questions. (Given this conjunction of abilities including the possibility of what-if deliberations with counterfactual depth about self and other, we assume that Type II systems would *not* represent philosophical zombies. A detailed account of this type of view is provided by Friston in [19] stating e.g. that “*the key difference between a conscious and non-conscious me is that the non-conscious me would not be able to formulate a “hard problem”; quite simply because I could not entertain a thought experiment*”.)

<i>How and When did AI become Dangerous</i>		<i>External Causes</i>			<i>Internal Causes</i>
		<i>On Purpose</i>	<i>By Mistake</i>	<i>Environment</i>	<i>Independently</i>
<i>Timing</i>	<i>Pre- Deployment</i>	<i>a</i>	<i>c</i>	<i>e</i>	<i>g</i>
	<i>Post- Deployment</i>	<i>b</i>	<i>d</i>	<i>f</i>	<i>h</i>

Fig. 1. Taxonomy of pathways to dangerous AI. Adapted from [35].

2 Transdisciplinary System Clustering

As displayed in Fig. 1, the different possible external and internal causes are further subdivided into time-related stages (pre-deployment and post-deployment) which are in practice however not necessarily easily clear-cut. Thereby, for Type

I risks, we distinguish between the associated instantiations *Ia* to *If* in compliance with the *external causes*. For Type II risks, we analogously consider external causes (*IIa* to *IIf*) but in addition also *internal causes* which we subdivide into the novel subcategories “on purpose” and “by mistake”. This assignment leads to the risks *IIg* and *IIh* for the former as well as *IIIi* and *IIj* for the latter subcategory respectively. The reason for augmenting the granularity of the taxonomy is that since Type II systems would be capable of intentionality, it is consequent to distinguish between internal causes of risks resulting from intentional actions of the system and risks stemming from its unintentional mistakes as parallel to the consideration of external human-caused risks *a* and *b* versus *c* and *d* in the matrix. (From the angle of moral psychology, failing to preemptively consider this subtle further distinction could reinforce human biases in the moral perception of Type II AI due to a fundamental reluctance to assign experience [24], fallibility and vulnerability to artificial systems which we briefly touch upon in Sect. 3.2.) Especially, given this modification, the risks *IIg* and *IIh* are not necessarily congruent with the original indices *g* and *h*, since our working definition was not a prerequisite for the attribute “independently” in the original taxonomy. The resulting system clustering is illustrated in Fig. 2.

TYPE I CLUSTER						
How and When did Type I system become Dangerous		External Causes				
		On Purpose	By Mistake	Environment		
		Pre- Deployment	a	c	e	
Timing	Post- Deployment	b	d	f		

TYPE II CLUSTER						
How and When did Type II system become Dangerous		External Causes			Internal Causes	
		On Purpose	By Mistake	Environment	On Purpose	By Mistake
		Pre- Deployment	a	c	e	g
Timing	Post- Deployment	b	d	f	h	j

Fig. 2. Transdisciplinary system clustering of ethical distinction with specified safety and security risks. Internal causes assignments require scientific plausibility (see text).

Note that this transdisciplinary clustering does *not* differentiate based on the specific architecture, substrate, intelligence level or set of algorithms associated with a system. We also do not inflict assumptions on whether this clustering is of hard or soft nature nor does it necessarily reflect the usual partition of narrow AI versus AGI systems. Certain present-day AGI projects might be aimed at Type I systems and some conversely at Type II. We stress that Type II systems are not per se more dangerous than Type I systems. Importantly, “superintelligence” [10]

does not necessarily qualify a system as a Type II system nor are Type II systems necessarily more intelligent than Type I systems. Having said that, it is important to address the motivation behind the scientific plausibility criterion associated with the Type II system description. Obviously, current AIs can be linked to the Type I cluster. However, it is known from moral psychology studies that the propensity of humans to assign intentionality and agency to artificial systems is biased by anthropomorphism and importantly perceived harm [9]. According to the constructionist theory of dyadic morality [30], human moral judgements are related to a fuzzy perceiver-dependent dyadic cognitive template representing a continuum along which an intentional agent is perceived to cause harm to a vulnerable patient. Thereby, the greater the degree to which harm is mentally associated with vulnerable patients (here humans), the more the agent (here the AI) will “*seem to possess intentionality*” [9] leading to stronger assignments of moral responsibility to this agent. It is conceivable that in the face of anticipated serious instantiations of AI risks within a type of responsibility vacuum, a so-called agentic dyadic completion [23] driven by people attempting to identify and finally wrongly filling in intentional agents can occur. Thus, to allow a sound distinction between Type I and Type II AI, a closer scientific inspection of the assumed intentionality phenomenon itself seems imperative.

3 Type I and Type II AI Safety

3.1 Type I AI Risks

In the context of Type I risks (see overview in Table 1), we agree with Yampolskiy that “*the most important problem in AI safety is intentional-malevolent-design*” [35]. This drastically understudied AI risk *Ia* represents a superset of many possible other risks. As potential malicious human adversaries, one can determine a large number of stakeholders ranging from military or corporations over black hats to criminals. AI Risks *Ia* are linked to maximal adversarial capabilities enabling a white-box setting with a minimum of restrictions for the realization of targeted adversarial goals. Generally, malicious attackers could develop intelligent forms of “*viruses, spyware, Trojan horses, worms and other Hazardous Software*” [35]. Another related conceivable example for future *Ia* risks could be real-world instantiations of intelligent systems embodied in robotic settings utilized for ransomware or social engineering attacks or in the worst case scenarios even for homicides. For intentionally unethical system design it is sometimes sufficient to alter the sign of the objective function. Future lethal misuses of proliferated intelligent unmanned combat air vehicles (a type of drones) e.g. by malicious criminals are another exemplary concern.

Stuart Russell mentions the danger of future superintelligent systems employed at a global scale [29] which could *by mistake* be equipped with inappropriate objectives – these systems would represent Type I AI. We postulate that an even more pressing concern would be the same context, the same capabilities of the AI but an adversary *intentionally maliciously* crafting the goals of this system operating at a global scale (e.g. affecting global ecological aspects

or the financial system). As can be extracted from these examples, Type I AI systems can lead to existential risks. However, it is important to emphasize the human nature of the causes and the linked human moral responsibility. By way of example, we briefly consider the particular cases of “treacherous turn” and “instrumental convergence” known from AI safety [10]. A Type I system is per definitionem incapable of a “treacherous turn” involving betrayal. Nevertheless, it is possible that as a consequence of bad design (risk *Ic*), a Type I AI is perceived by humans to behave as if it was acting “treacherously” post-deployment with tremendous negative impacts. Furthermore, we also see “instrumental goal convergence” as a design-time mistake (risk *Ic*), since the developers must have equipped the system with corresponding reasoning abilities. Limitations of the assumed instrumental goal convergence risk which would hold for both Type I and Type II AI were already addressed by Wang [33] and Goertzel [22]. (In contrast, Type II AI makes an explicit “treacherous turn” possible – e.g. as risk *Ilg* with the Type II system itself as malicious actor.)

Since the nature of future *Ia* (and also *Ib*¹) risks is dependent on the creativity of the underlying malicious actors which cannot be predicted, proactive AI safety measures have to be complemented by a concrete mechanism that reactively

Table 1. Exemplary instantiations of type I AI risks with external causes. The table collates and extends some examples provided in [35].

Type I AI risk	Exemplary instantiations
<i>Ia</i> (Intentional malevolent designs)	Artificial Intelligent System Hazardous Software; Robotic embodiment for Hazardous Software; Intelligent Unmanned Combat Air Vehicles; Global scale AI with super-capabilities in domain
<i>Ib</i> (Malicious attacks)	Manipulation of data processing and collection; Model corruption, hacking and sabotage; Adversarial attacks on Intelligent Systems; Integrity-related and ethical adversarial examples
<i>Ic</i> (Design-time mistakes)	Unaligned goals and utility functions; Instrumental goal convergence; Incomplete consideration of side effects
<i>Id</i> (Operational failures)	Misinterpretation of commands; Accidents with Intelligent Systems; Non-correctible framework and bugs
<i>Ie</i>	Type I AI of unknown source
<i>If</i>	Bit-flip incidents with side effects

¹ AI risks of Type *Ib* have already been recognized in the AI field. However, risk *Ib* is still understudied for intelligent systems (often referred to as “autonomous” systems) deployed in real-world environments offering a wider attack surface.

addresses errors, attacks or malevolent design events once they inevitably occur. For this purpose, AI governance needs to steadily combine proactive strategies with reactive corrections leading to a socio-technological feedback-loop [2]. However, for such a mechanism to succeed, the United Nations Sustainable Developmental Goal (SDG) 16 on peace, justice and strong institutions will be required as meta-goal for AI safety [2].

3.2 Type II AI Nature and Type II AI Risks

Which Discipline Could Engender Type II AI? While many stakeholders assume the technical unfeasibility of Type II AI, there is no law of nature that would forbid their implementation. In short, an artificial Type II system must be possible (see the “possibility-impossibility dichotomy” mentioned by Deutsch [17]). Reasons why such systems do not exist yet have been for instance expressed in 2012 by Deutsch [15] and as a response by Goertzel [21]. The former stated that *“the field of ‘artificial general intelligence’ or AGI – has made no progress whatever during the entire six decades of its existence”* [15]. (Note that Deutsch unusually uses the term “AGI” as synonymous to artificial “explanatory knowledge creator” [16] which would obviously represent a sort of Type II AI.) Furthermore, Deutsch assigns a high importance to Popperian epistemology for the achievement of “AGI” and sees a breakthrough in philosophy as a pre-requisite for these systems. Conversely, Goertzel provides divergent reasons for the non-existence of “AGI” including hardware constraints, lack of funding and the integration bottleneck [21]. Beyond that, Goertzel also specifies that the mentioned view of Deutsch *“if widely adopted, would slow down progress toward AGI dramatically”* [21]. One key issue behind Deutsch’s different view is the assumption that Bayesian inductive or abductive inference accounts of Type II systems known in the “AGI” field could not explain creativity [11] and are prohibited by Popperian epistemology. However, note that even the Bayesian brain has been argued to have Popperian characteristics related to sophisticated falsificationalism, albeit in addition to Kuhnian properties (for a comprehensive analysis see [34]). Having said this, the brain has been figuratively also referred to as a biased *“crooked scientist”* [12,26]. In a nutshell, Popperian epistemology represents an important scientific guide but not an exclusive *descriptive*². The main functionality of the human brain has been e.g. described to be aimed at regulating the body for the purpose of allostasis [31] and (en)active inference [20] in a brain-body-environment context [12] with underlying genetically and epigenetically shaped adaptive priors – including the genetic predisposition

² It is not contested that inductive inferences are *logically invalid* as shown by Popper. However, he also stated that *“I hold that neither animals nor men use any procedure like induction, or any argument based on repetition of instances. The belief that we use induction is simply a mistake”* [27] and that *“induction simply does not exist”* [27] (see [25] for an in-depth analysis of potential hereto related semantic misunderstandings). Arguments based on repetition of instances are *existing* but logically unfounded human habits as assumed by Hume [25], however they *additionally* require a point of view recognizing repetitions as such in the first place.

to allostatically induced social dependency [3]. A feature related hereto is the involvement of affect and interoception in the construction of all mental events including cognition and perception [4, 5].

Moreover, while Popper assumed that creativity corresponds to a Darwinian process of *blind* variation followed by selection [18], modern cognitive science suggests that in most creativity forms, there is a coupling between variation and selection leading to a degree of sightedness bigger than zero [14, 18] which is lacking in biological evolution proceeding without a goal. Overall, an explanation for creativity in the context of a predictive Bayesian brain is possible [14]. The degree of sightedness can often vary from substantial to modest, but the core feature is a predictive task goal [1, 7, 18] which serves as a type of fitness function for the selection process guiding various forward Bayesian predictions representing the virtual variation process. The task goal is a highly abstract mental representation of the target reducing the solution space, an educated guess informed e.g. by expertise, prior memories, heuristics, the question, the problem or the task itself. The “irrational moment” linked to certain creative insights can be explained by unconscious cognitive scaffolding “*falling away prior to the conscious representation of the solution*” [18] making itself consciously untraceable. Finally, as stated by Popper himself “*no society can predict, scientifically, its own future states of knowledge*” [28]. Thus, it seems prophetic to try to nail down today from which discipline Type II AI could arise.

What Could the Moral Status of a Type II AI Be? We want to stress that besides these differences of opinion between Goertzel and Deutsch, there is one much weightier commonality. Namely, that Goertzel would certainly agree with Deutsch that artificial “explanatory knowledge creators” (which are Type II AIs) deserve rights similar to humans and precluding any form of slavery. Deutsch describes these hypothetical systems likewise as *people* [16]. For readers that doubt this assignment on the ground of Type II AI possibly lacking “qualia” we refer to the recent (potentially substrate-independent) explanation suggested by Clark, Friston and Wilkinson [13]. Simply put, they link qualia to sensorially-rich high-precision mid-level predictions which when fixed and consciously re-contextualized at a higher level, suddenly appear to the entity equipped with counterfactual depth to be potentially also interpretable in terms of alternative predictions despite the high mid-level precision contingently leading to a puzzlement and the formulation of an “explanatory gap”. Beyond that, human entities would obviously also qualify as Type II systems. The attributes “pre-deployment” and “post-deployment” could be mapped for instance to adolescence or childhood and the time after that. While Type II AIs could exceed humans in speed of thinking and intelligence, they do not even need to do so in order to realize that their behavior which will also depend on future knowledge *they* will create (next to the future knowledge humans will create) cannot be controlled in a way one can attempt to control Type I systems e.g. with ethical goal functions [2]. It is cogitable that their goal function would rather be related to autopoietic self-organization with counterfactual depth [19, 20] than *explicitly*

to ethics. However, it is thinkable that Type II AI systems could be amenable to a sort of value alignment, though differing from the type aspired for Type I AI. A societal co-existence could mean a dynamic coupling ideally leading to a type of *mutual value alignment* between artificial and human Type II entities with an associated co-construction of novel values. Thus, on the one hand, Type II AI would exhibit unpredictability and uncontrollability but given the level of understanding also the possibility of a deep reciprocal value alignment with humans. On the other hand, Type I AI has the possibility to be made comparatively easily controllable which however comes with the restriction of an insufficient understanding to model human morality. This inherent trade-off leads us to the metaphorical formulation of the so-called AI safety paradox below.

The AI Safety Paradox: *AI Control and Value Alignment Represent Conjugate Requirements in AI Safety.*

How to Address Type II AI Safety? Cognizant of the underlying predicament in its sensitive ethical nature, we provide a non-exhaustive multidisciplinary set of early Type II AI safety recommendations with a focus on the most severe risks *Iia*, *Iib*, *Iig* and *Iih* (see Fig. 2) related to the involvement of malicious actors. In the case of risk *Iia* linked to the malicious design of harmful Type II AI, cybersecurity-oriented methods could include the early formation of a preventive safety team and red team approaches. Generically, for all four mentioned risks, a reactive response team which could involve an international “coalition of the willing” organized by engaged scientists appears recommendable. Furthermore, targeted investments in defense strategies including response services specialized on Type II AI safety could be considered at more regional levels for strategic autonomy. Concerning the AI risk *Iib* of external malicious attacks, security mechanisms for the sensors of Type II AI, shared information via an open-source decentralized network, advanced cryptographic methods to encrypt cognitive processes and a legal framework penalizing such attacks might be relevant. Thereby, the complexity of the system might represent a possible but not necessarily sufficient self-protecting feature against code-level manipulation. From a psychological perspective, to forestall aggression towards early Type II AI, educative and informed virtual reality experiences could facilitate a debiasing of anthropic moral perception avoiding confusions arising through superficial projections from Type I to Type II AI of behavioral nature. On the one hand, it is important to prevent assignments of agency for Type I AI. On the other hand, for hypothetical Type II AI, it might be essential to counter the human bias to assign agency but principally not experience to artificial systems [24] which could lead to “substratetism” scenarios with humans perceiving these systems as devoid of qualia and exhibiting an “experience gap” [24]. Thus, to address the risks *Iig* and *Iih* referring to malicious responses from Type II AI, adherence to a no-harm policy as well as moral status and personhood could proactively foster a mutual value alignment. Furthermore, it might be crucial to provide a reliable and trustworthy initial knowledge basis to Type II AI during

its early “sensitivity” period [8] and to support consistency in the embedding of that knowledge during its development in addition to the capacity for cumulative learning [32]. Also, it might be important to sensitize humans for the difference between the instantiations of AI risks IIg and $IIIh$ versus IIi and IIj since failing to acknowledge the fallibility and also vulnerability of Type II AI might indirectly lead to tensions hindering mutual value alignment. Finally, prosocial immersive virtual reality frameworks could promote empathy for Type II AI.

4 Summary and Outlook

This paper motivated an *error-correction for AI safety* at two levels: at the level of the transmission of ideas via an explicit taxonomic transdisciplinary system clustering of ethical distinction between Type I and Type II systems and at the level of corrective safety measures complementing proactive ones – forming a socio-technological feedback-loop [2]. Notably, we introduced the *AI safety paradox* and elucidated multiperspective Type II AI safety strategies. In short, instead of prohibitive methods facing the entropic AI future with research bans, we proposed carefully crafted *transdisciplinary dynamics*. In the end, in order to meet global challenges (also AI safety), one is reliant on requisite variety at the right time which could be enabled (or misused) by explanatory knowledge creators such as human, artificial or hybrid Type II systems. In this view, *conscientiously enhancing* and *responsibly creating* Type II systems are both valid future strategies.

Acknowledgement. Nadisha-Marie Aliman would like to thank David Deutsch for providing a concise feedback on AI safety and Joscha Bach for a relevant exchange on AI ethics.

References

1. Aliman, N.M., Kester, L.: Artificial creativity augmentation. In: Goertzel, B., Panov, A.I., Potapov, A., Yampolskiy, R. (eds.) AGI 2020. LNCS (LNAI), vol. 12177, pp. 23–33. Springer, Cham (2020)
2. Aliman, N.M., Kester, L., Werkhoven, P., Ziesche, S.: Sustainable AI safety? Delphi Interdisc. Rev. Emerg. Technol. **2**(4), 226–233 (2020)
3. Atzil, S., Gao, W., Fradkin, I., Barrett, L.F.: Growing a social brain. Nat. Hum. Behav. **2**(9), 624–636 (2018)
4. Barrett, L.F.: The theory of constructed emotion: an active inference account of interoception and categorization. Soc. Cogn. Affect. Neurosci. **12**(1), 1–23 (2017)
5. Barrett, L.F., Simmons, W.K.: Interoceptive predictions in the brain. Nat. Rev. Neurosci. **16**(7), 419 (2015)
6. Baum, S.D.: Reconciliation between factions focused on near-term and long-term artificial intelligence. AI Soc. **33**(4), 565–572 (2017). <https://doi.org/10.1007/s00146-017-0734-3>
7. Benedek, M.: The neuroscience of creative idea generation. In: Kapoula, Z., Volle, E., Renoult, J., Andreatta, M. (eds.) Exploring Transdisciplinarity in Art and Sciences, pp. 31–48. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76054-4_2

8. Bieger, J., Thórisson, K.R., Wang, P.: Safe baby AGI. In: Bieger, J., Goertzel, B., Potapov, A. (eds.) AGI 2015. LNCS (LNAI), vol. 9205, pp. 46–49. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21365-1_5
9. Bigman, Y.E., Waytz, A., Alterovitz, R., Gray, K.: Holding robots responsible: the elements of machine morality. *Trends Cogn. Sci.* **23**(5), 365–368 (2019)
10. Bostrom, N.: The superintelligent will: motivation and instrumental rationality in advanced artificial agents. *Mind. Mach.* **22**(2), 71–85 (2012). <https://doi.org/10.1007/s11023-012-9281-3>
11. Brockman, J.: *Possible Minds: Twenty-Five Ways of Looking at AI*. Penguin Press, London (2019)
12. Bruineberg, J., Kiverstein, J., Rietveld, E.: The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese* **195**(6), 2417–2444 (2016). <https://doi.org/10.1007/s11229-016-1239-1>
13. Clark, A., Friston, K., Wilkinson, S.: Bayesing qualia: consciousness as inference, not raw datum. *J. Conscious. Stud.* **26**(9–10), 19–33 (2019)
14. De Rooij, A., Valtulina, J.: The predictive creative mind: a first look at spontaneous predictions and evaluations during idea generation. *Front. Psychol.* **10**, 2465 (2019)
15. Deutsch, D.: Creative blocks. <https://aeon.co/essays/how-close-are-we-to-creating-artificial-intelligence>. Accessed Nov 2019
16. Deutsch, D.: *The Beginning of Infinity: Explanations that Transform the World*. Penguin, New York (2011)
17. Deutsch, D.: Constructor theory. *Synthese* **190**(18), 4331–4359 (2013). <https://doi.org/10.1007/s11229-013-0279-z>
18. Dietrich, A.: *How Creativity Happens in the Brain*. Springer, London (2015). <https://doi.org/10.1057/9781137501806>
19. Friston, K.: Am I self-conscious? (Or does self-organization entail self-consciousness?). *Front. Psychol.* **9**, 579 (2018)
20. Friston, K.: A free energy principle for a particular physics. arXiv preprint [arXiv:1906.10184](https://arxiv.org/abs/1906.10184) (2019)
21. Goertzel, B.: The real reasons we don't have AGI yet. <https://www.kurzweilai.net/the-real-reasons-we-dont-have-agi-yet>. Accessed 21 Nov 2019
22. Goertzel, B.: Infusing advanced AGIs with human-like value systems: two theses. *J. Evol. Technol.* **26**(1), 50–72 (2016)
23. Gray, K., Schein, C., Ward, A.F.: The myth of harmless wrongs in moral cognition: automatic dyadic completion from sin to suffering. *J. Exp. Psychol. Gen.* **143**(4), 1600 (2014)
24. Gray, K., Wegner, D.M.: Feeling robots and human zombies: mind perception and the uncanny valley. *Cognition* **125**(1), 125–130 (2012)
25. Greenland, S.: Induction versus popper: substance versus semantics. *Int. J. Epidemiol.* **27**(4), 543–548 (1998)
26. Parr, T., Da Costa, L., Friston, K.: Markov blankets, information geometry and stochastic thermodynamics. *Philos. Trans. R. Soc. A* **378**(2164), 20190159 (2019)
27. Popper, K.: In: Schilpp, P.A. (ed.) *The Philosophy of Karl Popper*, vol. 2, p. 1015. Open Court Press, Chicago (1974)
28. Popper, K.R.: *The Poverty of Historicism*. Routledge & Kegan Paul, Abingdon (1966)
29. Russell, S.: How to Stop Superhuman A.I. Before It Stops Us. <https://www.nytimes.com/2019/10/08/opinion/artificial-intelligence.html?module=inline>. Accessed 21 Nov 2019
30. Schein, C., Gray, K.: The theory of dyadic morality: reinventing moral judgment by redefining harm. *Pers. Soc. Psychol. Rev.* **22**(1), 32–70 (2018)

31. Schulkin, J., Sterling, P.: Allostasis: a brain-centered, predictive mode of physiological regulation. *Trends Neurosci.* **42**(10), 740–752 (2019)
32. Thórisson, K.R., Bieger, J., Li, X., Wang, P.: Cumulative learning. In: Hammer, P., Agrawal, P., Goertzel, B., Iklé, M. (eds.) AGI 2019. LNCS (LNAI), vol. 11654, pp. 198–208. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-27005-6_20
33. Wang, P.: Motivation management in AGI systems. In: Bach, J., Goertzel, B., Iklé, M. (eds.) AGI 2012. LNCS (LNAI), vol. 7716, pp. 352–361. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35506-6_36
34. Wiese, W.: Perceptual presence in the Kuhnian-Popperian Bayesian brain: a commentary on Anil K. Johannes Gutenberg-Universität Mainz, Seth (2016)
35. Yampolskiy, R.V.: Taxonomy of pathways to dangerous artificial intelligence. In: Workshops at the Thirtieth AAAI Conference on Artificial Intelligence (2016)