Let's Talk AI with Bernhard Steffen

Bernhard Steffen[1] and Barbara Steffen[2]

[1] TU Dortmund, Department of Computer Sciences, Germany,
steffen@cs.tu-dortmund.de
[2] METAFrame Technologies GmbH,
barbara.steffen@metaframe.de

"At the end, it is far beyond whatever I would have expected. LLMs will certainly have a major societal impact."

The Interviewee - Bernhard Steffen

My Personal AI Mission:
There are huge opportunities and major risks. As there is no way to stop the development of AI, my mission is helping to make the best of it.

My Takes on AI

Artificial Intelligence: It changed over time, today it's mainly machine learning with increasing focus on large language models, which, in my eyes, for the first time justify the name AI.

Trust: Concerns how humans perceive a certain entity. It should not be confused with trustworthiness, which is a property of the considered entity itself.

Explainability: Traditionally, explainability of AI focused on indications of how a certain result is achieved. Typical are so-called heat maps which highlight the most important pixel for an image classification. This approach became

obsolete with large language models. Here post-hoc rationalization seems more appropriate: Let the net provide you with reasons for the delivered result. This way, we are reaching a communication level reminiscent of dialogues between humans.

Essential Elements of Human Capabilities: I have trouble pinpointing a capability which I consider out of reach of future AI-based technologies.

The Interview

Barbara *Today I have the pleasure of interviewing Professor Bernhard Steffen from the Technical University of Dortmund. Please introduce yourself and your connection to artificial intelligence.*

Bernhard My name is Bernhard Steffen. As you mentioned, I'm a professor at the Technical University of Dortmund. Recently, I founded AISoLA, the event where this interview is taking place. I wasn't always a fan of artificial intelligence. I've observed so many hype cycles of AI throughout my life, and my appreciation only built up slowly. Some twenty years ago, I introduced active automata learning as a means to achieve model-based testing without requiring a priori models [12], which was a major hurdle in practice. Later we analyzed and explained random forests [7] by essentially transforming them into single decision trees. My interest in neural networks is, however, much younger. It started when we observed that we could treat them in a similar fashion as random forests [14]. Conceptually, this was quite elegant, but scalability was a problem. The true limits of our 'traditional' approaches, however, became apparent with the large language models (LLMs) like ChatGPT. They require a major change of mind, not only concerning their validation, but also their use as they will inevitably impact everybody's life. This insight became the driving force behind the organization of AISoLA.

Barbara *Can you name one or two specific AI-related research questions that you and your research group are currently working on?*

Bernhard We are working along two independent lines: One is that we want to verify or explain systems. We take the AI system as it is and try to understand or at least contain what it does, for example by verifying that it doesn't cause harm. The other line is that we try to build systems that exploit neural networks or other learned systems in a 'harmless' way. We have to learn how to do that properly. Of course, if the entire system is not critical, there are no problems. More challenging is the case of critical systems that embody AI components that have to be 'tamed' to guarantee the reliability of the entire system. We are currently experimenting with systems that contain AI-generated code. These systems do not contain, e.g., neural networks, but only their output, the generated code. Thus, they can be treated with the methods which have been developed in recent decades. This is a change of attitude towards AI-based systems: Do not try to verify or explain the LLMs but their outcome. In the case of code generating LLMs this sounds very natural. Nobody verifies the programmer. It is the program that is verified. I like this approach which, in a sense, puts LLMs and humans on an equal level, as I do not believe that we will ever be able to sufficiently control the LLMs themselves.

Barbara *Do you focus specifically on the developers, such as technical experts, or on end users?*

Bernhard In fact both: We have developed the concept of Language-Driven Engineering [17] to support even end users to participate in the development process of certain applications by providing them with a graphical language interface that addresses their mindsets. In fact, the aim of LDE to narrow the gap between end users and developers can be reduced even further by domain-specific natural languages: Based on, e.g., ChatGPT-based code generation, people can participate in the development process without any previous training, using natural language. I am convinced that future program development environments will increasingly take advantage of this opportunity (see e.g. [3]).

Barbara *How do you verify the program?*

Bernhard There are numerous ways to do that, based on half a century's experience in program verification [6, 9]. There are interactive methods, typically supported by theorem proving, which require already expert knowledge to handle small scale applications [8]. In contrast, model-checking techniques [5] and their statistical counterparts [1, 2] are fully automated as soon as an adequate abstraction level has been defined and they scale much better. In fact, statistical model checking, which inherently is a testing technology, applies even to larger scale neural networks where all other techniques fail. At the down-side, statistical model checking only establishes knowledge with a certain probability. This may be considered a drawback for traditional programs, but it seems quite natural for structures that have been learned with statistical methods, like neural networks. Our research focuses on the verification of, from their structure, 'traditional' programs that have been written with AI assistance. This allows us to apply the wealth of methods that have been developed in recent decades. We apply model checking and runtime verification based on active automata learning [3]. In the latter case, one looks at the entire system, which consists of a mixture of handwritten or model-driven designed code and parts that have been generated from natural language in a form of AI-assisted programming. You just look at the runtime behavior and validate the properties of the overall system. The application of active automata learning is particularly effective for web applications, as it directly works at the level of user interactions with the system and does not require us to deduce the behavior from the highly heterogeneous code structure.

"[The AIs'] so-called emergent properties may well lead to a kind of 'intelligence' far beyond our expectation and recognition, and therefore far beyond our control...[In fact, many] capabilities that experts thought were out of the scope of these systems have, in the meantime, already been accomplished."

Barbara *In your opinion, what role does trust play in the adoption of AI?*

Bernhard Trust is certainly a crucial factor if one wants people to use a new technology. What I've learned at AISoLA is that trust is a complex concept.

When we, the formal methods people, think about trust, we consider the reasons why we should trust, e.g., AI. Can I prove that it does what it promises or at least that it won't cause harm? At AISoLA I learned that these considerations about the systems' (provable) quality criteria concerns their *trustworthiness* which should clearly be distinguished from *trust*, as it is a property of the user: a user may well trust a system that is not trustworthy and vice versa. Indeed, psychologists study ways to increase trust without caring about the properties of a system, e.g., by placing quality seals or getting famous people to speak in favor. From an ethical perspective it is now important to adequately balance the two so that people's trust or fear are justified.

Orthogonal is the question of the required level of trust. Clearly, there is a big difference between using DALL-E to create a picture and sitting in a self-driving car. I learned at AISoLA that this is where the notion of *calibrated trust* comes into play. We must learn to adjust our trust level to the criticality of the considered (AI) system. This is not as easy as my example suggests. E.g., social media seem harmless at first sight, but we had to learn that they can be exploited to misinform and polarize to a level that even may endanger democracy, another topic that has been discussed at AISoLA [13], To summarize. Establishing a responsible understanding of trust is important, and it is very difficult when systems with the complexity of LLMs are involved that, inevitably, will inform our lives. We have to learn to live with these systems, and we have to find ways such that this is done in a fashion where trust and trustworthiness are in balance and calibrated to the criticality of the application.

Barbara *Are there essential measures that we should consider to ensure ethical AI use?*

Bernhard Ethics is a very different matter. As a technician I would say, whenever I have proved that a system is trustworthy, it is ethical for me to convince you to use it. But the ethical discourse is typically quite different. Famous are the discussions of dilemmas: *What is the ethical decision in a situation where one either drives straight on a road passively and kills five people, or actively steers the car to the side to kill only one person?* These kinds of questions have come up, and they are not at all technical. They belong to a different discipline which, as far as I understood, does not take a clear position here, as it is unethical to measure the value of individual lives or even compute with them. In practice, this works if these cases are rare and can be dealt with on an individual basis. However, it is not sufficient when we want to set the rules for AIs. In fact, we are exactly in this situation when talking about the acceptance of self-driving cars where similar questions arise: *Is it ethical to introduce self-driving cars if this would reduce the number of severe accidents by 90%, but with the effect that other people die?* As a technician I would say *"90% reduction is great. Let's go for it."* But this discussion seems, at least in Germany, to be far from over. At the end, the situation is similar to trust: We have to learn to adequately calibrate ethics.

Barbara *Regarding the technical capabilities of artificial intelligence in the future, on a scale of 1 to 10, where 1 stands for artificial intelligence systems like ChatGPT, and 10 for artificial general intelligence systems that surpass human capabilities. What do you think will be possible?*

Bernhard In some sense, I think what you describe as 1 is already quite advanced, because these LLMs are already quite general: They can write, translate, code, draw, play, etc., essentially based on the same machinery and with a performance that often surpasses human capabilities. This shows that we are already beyond the stage where AlphaGo demonstrated the power of AI by easily beating the top human Go players. This was amazing, but still considered harmless, because AlphaGo is just a special tool, like a crane that is considered harmless even though it can lift material far beyond human capabilities. LLMs like ChatGPT provide already a flavor of what you're referring to with artificial general intelligence. They cover a complexity beyond human reach and their 'mental capability' surpasses human expertise in many situations. Looking at the current development, I do not think that it will take long until these AIs will beat the best human experts in numerous, unrelated challenges, not only in playing Go or translating texts, but also in writing poems, drawing pictures, steering cars, and even executing surgeries or playing football when combined with adequate robotics. In fact, I cannot imagine a concrete capability that I would consider to definitely remain out of reach for AI systems in the future. The talk by Edward Lee [10] was very enlightening in this respect. He showed that capabilities that experts thought were out of the scope of these systems have, in the meantime, already been accomplished. It's not only that I think we will see that these systems can perform better on tasks we understand. Rather their so-called *emergent properties* may well lead to a kind of 'intelligence' far beyond our expectation and recognition, and therefore far beyond our control. Again, referring to Edward Lee's talk, and to experiences with Pi [22], the communication with the AI-based system may even become more attractive than with other humans: They 'listen' carefully, and answer adequately, and they are always available to provide advice. In fact, when asked for the underlying WHY, they are even great at explaining the advantages of the proposed statements or decisions. As Edward Lee says, these systems will be unbeatable when it comes to post-hoc rationalization [10]: They simply have a much wider knowledge base than any human, and it will be increasingly difficult for us to withstand this convincing power.

"[What would happen] when all technicians agree that we cannot control AI systems and should abandon them? Abandon, how? There is no such way. We need to learn to navigate the uncertainty that comes with the AI systems."

Barbara *Could you place the possible future capabilities somewhere on the scale from one to ten?*

Bernhard The power of LLMs was like a shock for me. I never had expected that this technology would carry this far. Just look at the many directions it influences today already. As I mentioned, I currently have no idea where the limits could be and I would clearly say 10, or even more: It will not be too difficult for future AIs to clearly outperform humans in most challenges. These systems may soon become emotionally attractive to humans in the way shown in *Her* [21]. In fact, the dialogues with Pi [22, 19] are already now quite fascinating in this sense.

Barbara *Looking at that, there are a lot of different scenarios being discussed about what the future will look like as AI develops. These scenarios range from dystopia to utopia. Where do you position yourself?*

Bernhard I would position myself in the middle, say 5 in a range from 1 to 10 which, admittedly, sounds boring. In the past, a lot of dystopian scenarios have been warned about, like nuclear wars, overpopulation, and global warming, to name just a few. In almost all these cases, we react(ed), according to many experts, far too hesitantly. E.g., concerning overpopulation, we were warned in the 70s that we may reach a population of up to four billion people. Now we are eight billion and we're still here. Also with global warming, I am pretty sure that we will surpass the temperature limits we set for ourselves, but I guess humanity will still survive. The strong thing about humans is that they can adapt. This does not mean that we will be able to maintain the current quality of life, but we will probably survive, hopefully not in a future as sketched in the movie Matrix, where we live in some kind of soup producing energy and the rest, even our mind, is under the control of computers.

I am not afraid of such a dystopian vision, but I believe that we, at least most of us, will increasingly follow decisions made by the authority of automated systems without questioning. In a sense, this started already with the introduction of traffic lights which let you wait even when you can clearly see that the streets are empty. The point is that we will successively forget about the notion of empty streets and simply wait – and don't care. Most people are happy with this, and they are, e.g., glad to follow a navigation system without having any idea about the surroundings they are driving through. This may work fine as long as there are some people who still know what is going on, but, perhaps, in the future, there are no such people anymore, and only the AI 'knows'. This may not be a problem if the criteria the AI uses are in line with our criteria. This may, however, easily get out of hand as nicely illustrated by The Social Dilemma [23], showing that originally 'ethical' criteria produced systems that polarize and misinform people. Please note that this did not even require particularly elaborate technology. In summary, I hope that the benefits outweigh the drawbacks of future AI-based systems, but I also believe that it will significantly change our lives in a way that forces us to adapt in ways not everybody would consider an improvement.

Barbara *Do you see the potential for a widening gap in society? For example, early adopters who are really excited about the new technical capabilities and try*

out everything until they get very good at it, while others are more hesitant and thus get left behind?

Bernhard I think this gap will definitely widen. We're living in a world which is accelerating its rate of change, and few and fewer people and organizations are able to deal with this. In this world, people or organizations that are used to and appreciate change as an opportunity have a significant advantage, while hesitant people and organizations are doomed to fail – unfortunately there are many of those in Germany! AI reinforces this trend, and it is important to be prepared for it. If an organization is what is called evolutionary fit [18] it can win, otherwise it will inevitably lose in the long term [16]. Thus, there will be many winners, but most probably many more losers in the future.

Barbara *Reflecting on the last few days and the interdisciplinary discussions, was there an insight that you found particularly interesting?*

Bernhard Perhaps most importantly, I realized that I see many things far too much from a technical perspective. For example, take the notion of trust again. I always thought achieving trust requires us to make the considered systems more robust, secure, reliable, etc. Then I listened to a talk from a psychologist where the impact of a seal on trust was considered independently of any technological quality criteria. Luckily, the study showed little impact. Looking back, I do not quite remember how people were questioned. Were they simply asked whether they value seals, or were they put in scenarios where similar systems were presented, some with and some without a seal? My intuition would be that for the latter, some impact should be measurable, at least if the seals convey some kind of seriousness, or even better, come with a known brand name. This reminds me a bit of the TÜV in Germany, the Technische Überwachungsverein, when certifying software. From the technological point of view, the TÜV does very little, but it provides certificates, which often suffice to officially operate a software system. This works quite well, because the TÜV is an authority. It is known to everybody in Germany for their biennial technical checks of cars. These checks are much more adequate for their purpose than the superficial analyses of software, which merely focus on the underlying development process and hardly concern the software as such. Is it acceptable that the TÜV authority gained for car inspections is then used to provide trust for software in the described way? As a technician I would say no.

"We have to holistically approach the situation with an interdisciplinary cooperation of experts not only from computer science, but also from, e.g., psychology, philosophy, law, education, economy, and politics."

On the other hand, we are simply not able to deal with the wealth of software systems in a way that I would consider satisfactory, and, so far, things seem to be ok. The infamous software crisis [11] never really struck, and we learned to live with the various obstacles poor software solutions provide. In the 'real world',

when you have to communicate with the public or with politicians, you have to make compromises and you have to accept the attention which some scientists get which I would rather categorize as marketing or sales people. After all, what would it help, for example, if the technicians all agree that we cannot control AI systems and should abandon them? Abandon, how? There is no such way. We need to learn to navigate the uncertainty that comes with the AI systems. Fear is not a good advisor here. My mother always told me *"when you risk to hit an obstacle, do not focus on the obstacle, but on the gap where you want to pass by".*

Barbara *Is there a specific topic or research question that you would like to see addressed from an interdisciplinary perspective, and, if so, which disciplines would you like to get involved?*

Bernhard In your thesis [16] you state that systems thinking [15] is one of the pillars for achieving evolutionary fitness, i.e., the capability to survive in a fast-changing world. I am convinced that we need the same here. We have to holistically approach the situation with an interdisciplinary cooperation of experts not only from computer science, but also from, e.g., psychology, philosophy, law, education, economics, and politics. AISoLA was founded to support exactly this cooperation. I was quite impressed by the different perspectives. I mentioned the discussions about trust and ethics already, and AI is certainly a game changer when it comes to law: how should we deal with liability caused by AI-based decisions. It certainly does not make sense to sue a system. There is no doubt that education is essential to establish an understanding of calibrated trust in the public and the impact of economical goal setting has clearly been pinpointed by the Social Dilemma [23] which calls for regulation. This underlines the importance of institutions like the Research Center on Trust and the Lamarr institute in Dortmund, which are both characterized by their interdisciplinarity.

Barbara *From your personal perspective, what should be the AI vision?*

Bernhard A year ago, I clearly would have said that we should try to control AI in a fashion we control a crane, just as a means to amplify our capabilities. However, looking at LLMs, is this possible at all? I heard people saying that it makes no sense to play Go anymore because a computer can do it better. I do not follow this argument as I still like to play Go and Chess for fun, and, concerning cranes, weightlifting is still an Olympic discipline. The same may be true for car racing: Formula One racing between humans will probably still be popular when automated cars outperform humans. But there is a difference with AI when it starts showing us the limits of our very bounded rationality [10] in real life, by, e.g., questioning each of our decisions with convincing explanations based on post-hoc rationalization. This may be fantastic at the beginning when we are still able to stay in control by revising the explanations and this way make the proposed decisions ours. We have already seen examples in science where this form of amplification led to amazing breakthroughs. But in the long run, when the AI proposals become more and more perfect, we may risk to totally delegate decision making to the AI, as many already delegated route planning to Google

Maps. Is this dystopian or utopian? We must learn to accept the AI systems as (often superior) partners rather than tools and to forget about understanding their inner workings. As mentioned before about AI-assisted programming, we should not try to verify/control the AI, but the produced output/code, just as we never verify the programmers but their work. Or, more generally, we should not try to verify the AI, but carefully inspect its responses. This way the AI may turn into a powerful partner, a big brother, if you wish, who certainly is watching us, but hopefully does not intend to harm us. I am optimistic that we can avoid dystopian emergent properties that lead us into slavery. Why should the AI aim at doing this – except, it learns by human examples which are sometimes quite scary? We have seen this with, e.g., social biases. Thus, a major challenge is to control the data sources for the learning processes.

Barbara *We need to distinguish contexts here, right? We have contexts like chess and Go, which are games with a predefined set of rules that describe the game. But the real world is less defined, much more complex, and constantly changing. So while the game or predefined settings make it easy to determine what is right and what is wrong, who won and who lost, what is a cat and what is a dog, in the real world there are many contexts or tasks for which we don't know the correct answers. So, it is more complicated to verify those answers. How should we approach these more complex scenarios?*

Bernhard You are right, learning in contexts where the outcome is clear, is quite different from learning in a context where there is no clear known answer. The dog/cat classification is a good example for the former, while the situation for chess and Go is much more involved. There is typically no known best move. The quality of a move is evaluated only many moves later when one of the players wins. This is a clear sign of success for the considered sample/play, but the impact of a certain move is still quite indirect and must be reconfirmed in many games. And there are situations where the role of a sample can only be understood in the context of the entire training set like in clustering, where huge data sets are bundled according to similarity. Also, language translation belongs to this category of what is called un/self-supervised learning, even though the output may well be checked for semantic or syntactic correctness. This additional information is typically applied in a subsequent reinforcement step. There is a wealth of methods and heuristics, but I am not the right person to tell more about them. We are focusing on the validation of the results of a learning process or, in the case of AI-assisted programming, of the resulting code.

"We must learn to accept the AI systems as (often superior) partners rather than tools and to forget about understanding their inner workings."

In most real-world applications of AI we do not know the precise answers, if they exist at all. This is where the dialogue between humans and AI comes into play, perhaps via a chain of post-hoc rationalization steps. In fact, I believe that these dialogues will boost our understanding in various disciplines. Already for playing

Go, AI helped to reveal patterns that humans would never have been able to find because of their limited rationality. AI can go quite a few steps further and reveal relationships which may be game changers in medicine, mathematics, physics, biology, etc. The recent results concerning protein folding indicate the potential, and I am convinced that this is only the beginning.

Barbara *Reflecting on the past week, how did you like AISoLA? It took place for the first time this year. Did it meet your expectations? And do you plan to continue?*

Bernhard I'm very happy with AISoLA, also from the perspective of the organizer. It was very interesting. Actually, I skipped most of the technical sessions, because I know very much what's going on there, and focused on the non-technical sessions to learn more about the other disciplines' perspectives. In connection with the LLM-shock I got earlier, it really changed my perspective. In particular, it revealed the challenges we have to deal with from the societal side. This has motivated my plan to continue organizing AISoLA yearly in the near future.

Barbara *Is there anything else you would like to add?*

Bernhard As a fan of bootstrapping and self-application [4] I first thought that I would disagree with Albert Einstein's quote *"We cannot solve our problems with the same thinking we used when we created them"* [20], because I am convinced that we have to use AI to deal with the problems caused by AI. Then I realized, Einstein did not speak about technology, but about mindset, and I agree that we need a change of mindset, not only of us, but also of the AIs. Think of the originally unintended polarization caused by social media as illustrated in The Social Dilemma [23]. It is a by-product, or as one would say today *emergent property* of profit maximization. Today's AI systems are orders of magnitude more complex than the rule-based systems mentioned in the Netflix documentary. It will therefore require a major collaborative effort to tame them.

Barbara *Then, Bernhard, thank you very much for your time and insights on AI and its further development. And especially a big thanks from me and all the interviewees for the opportunity to attend AISoLA and meet interdisciplinary experts on AI.*

Bernhard Thank you.

References

1. Alur, R., Giacobbe, M., Henzinger, T.A., Larsen, K.G., Mikučionis, M. (2019). Continuous-Time Models for System Design and Analysis. In: Steffen, B., Woeginger, G. (eds) Computing and Software Science. Lecture Notes in Computer Science, vol 10000. Springer.
2. Baier, C., Hermanns, H., Katoen, JP. (2019). The 10,000 Facets of MDP Model Checking. In: Steffen, B., Woeginger, G. (eds) Computing and Software Science. Lecture Notes in Computer Science, vol 10000. Springer.

3. Daniel Busch, Gerrit Nolte, Alexander Bainczyk, Bernhard Steffen: ChatGPT in the Loop: A Natural Language Extension for Domain-Specific Modeling Languages. AISoLA 2023: 375-390
4. Carey, S. (2004) "Bootstrapping & the origin of concepts," Daedalus. JSTOR, 133(1), pp. 59–68., https://en.wikipedia.org/wiki/Bootstrapping
5. Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem. 2018. Handbook of Model Checking (1st. ed.). Springer.
6. Floyd, Robert W. (1967). "Assigning Meanings to Programs" (PDF). In Schwartz, J.T. (ed.). Mathematical Aspects of Computer Science. Proceedings of Symposium on Applied Mathematics. Vol. 19. American Mathematical Society. pp. 19–32. ISBN 0821867288.
7. Gossen, F., Steffen, B. Algebraic aggregation of random forests: towards explainability and rapid evaluation. Int J Softw Tools Technol Transfer 25, 267–285 (2023).
8. Hähnle, R., Huisman, M. (2019). Deductive Software Verification: From Pen-and-Paper Proofs to Industrial Tools. In: Steffen, B., Woeginger, G. (eds) Computing and Software Science. Lecture Notes in Computer Science, vol 10000. Springer.
9. Hoare, C. A. R. (October 1969). "An axiomatic basis for computer programming". Communications of the ACM. 12 (10): 576–580. doi:10.1145/363235.363259. S2CID 207726175.
10. Lee, E.A. (2024). Deep Neural Networks, Explanations, and Rationality. In: Steffen, B. (eds) Bridging the Gap Between AI and Reality. AISoLA 2023. Lecture Notes in Computer Science, vol 14380. Springer.
11. Peter Naur & Brian Randell (Eds.): Software Engineering: Report of a conference sponsored by the NATO Science Committee, Garmisch, Germany, 7-11 Oct. 1968. NATO Scientific Affairs Division, Brüssel 1969
12. Raffelt, H., Merten, M., Steffen, B. et al. Dynamic testing via automata learning. Int J Softw Tools Technol Transfer 11, 307–324 (2009).
13. Schiaffonati, V., Werthner, H., Lee, E., Vardi, M., Laurus, J., Metakides, G. (2024). DigHum Panel. Beyond Chat-GPT: The Impact of AI on Academic Research. In: Steffen, B. (eds) Bridging the Gap Between AI and Reality. AISoLA 2023. Lecture Notes in Computer Science, vol 14380. Springer.
14. Maximilian Schlüter, Gerrit Nolte, Alnis Murtovi, Bernhard Steffen: Towards rigorous understanding of neural networks via semantics-preserving transformations. Int. J. Softw. Tools Technol. Transf. 25(3): 301-327 (2023)
15. System thinking: Senge, P. M. (2006) The fifth discipline: The art and practice of the learning organization. Broadway Business.
16. Barbara Steffen, Alignment-Driven Adaptation Process & Tool (ADAPT): Towards Continuous and Holistic Adaptation of Organizations, to appear in 2024
17. Bernhard Steffen, Frederik Gossen, Stefan Naujokat, Tiziana Margaria: Language-Driven Engineering: From General-Purpose to Purpose-Specific Languages. Computing and Software Science 2019: 311-344
18. Teece, D. J. and Petricevic, O. (2021) "Capability-based theories of multinational enterprise growth," in The Oxford handbook of international business strategy.
19. AISoLA, Pi-Interview - https://aisola.org/news/pi-interview/
20. Einstein - https://medium.com/@buttermanbreathes/we-cannot-solve-our-problems-with-the-same-thinking-we-used-when-we-created-them-dc3b33fdb9d4, https://www.nzherald.co.nz/world/10-famous-einstein-quotes/D46UEDPCON63F4JT3AVX4Z4J4U/
21. Her - https://en.wikipedia.org/wiki/Her_(film)

22. Pi - https://pi.ai/onboarding
23. Social Dilemma - https://en.wikipedia.org/wiki/The_Social_Dilemma