

Is Silicon Valley Diverse Enough to Resist Reduction?

Lund Pedersen, Carsten

Document Version
Final published version

Published in:
Journal of Design and Science

Publication date:
2019

License
CC BY

Citation for published version (APA):
Lund Pedersen, C. (2019). Is Silicon Valley Diverse Enough to Resist Reduction? *Journal of Design and Science*, (5). <https://jods.mitpress.mit.edu/pub/issue5-pedersen#discussions>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 02. Mar. 2024



Journal of Design and Science

Is Silicon Valley Diverse Enough to Resist Reduction?

Carsten Lund Pedersen

Published on: Jan 10, 2019

This essay takes Joichi Ito's [manifesto for resisting reduction](#) in AI as its point of departure, and asks, "Is Silicon Valley diverse enough to resist reduction in AI?"

Silicon Valley seems fascinated with the notion of 'Singularity', i.e. the moment in time when artificial intelligence (AI) becomes smarter than humans.¹ The prophecies of Singularity rest on assumptions of exponential growth in technological advances. The unilateral emphasis on Singularity in Silicon Valley is arguably problematic for the field of AI, as it means that it is being driven by a community with a rather monolithic focus. As Ito suggests in his manifesto for resisting reduction,² the collaborative model of complex adaptive systems should provide the paradigm for future work on AI. In an earlier piece, Ito similarly argued that intelligence can be perceived as being a distributed phenomenon, and therefore, humans and machines can together comprise a networked intelligence, i.e. 'extended intelligence' (EI).³ This argument was recently supported by the founding director of the Center for Collective Intelligence at MIT, Thomas Malone, who argues for the capacity of groups of people and machines to work together to obtain collective intelligence.⁴ Moreover, this claim is empirically supported.⁵ Paul J.H. Schoemaker and Philip E. Tetlock have made similar contributions where they argue for building more intelligent enterprises that combine human and artificial intelligence. That is, they provide an intricate argument for building systems that leverage the strengths of both human decision making and technology-enabled capabilities.⁶ Together these sources all point toward the validity of understanding AI ecosystems through a collective intelligence lens. But are the ecosystems that develop AI characterized by inherent collective intelligence? That is, does Silicon Valley as an ecosystem of AI developers and promoters convey collective intelligence in their combined systemic model of collaboration?

Numerous studies have established that good collaborative models, where the desired outcome is collective intelligence, are driven in part by cognitive diversity. That is, the members of the system must not only be skilled, they must understand and solve problems differently than one another, in order to obtain a system of collective intelligence. Ito argues that Silicon Valley mega companies "... are developed and run in great part by people who believe in a new religion, Singularity". Put differently, Silicon Valley seems to have homogenous notions of how to approach AI. Therefore, I argue that there is a lack of cognitive diversity among the people running Silicon Valley mega companies, and consequently, the main barrier to resisting reduction is in place. The lack of cognitive diversity results in suboptimal collective intelligence in the system, which helps explain the trend of reductionism. Ito alludes to this notion by cautioning that "In Silicon Valley, the combination of groupthink and financial success of this cult of technology has created a positive feedback system that has very little capacity for regulating through negative feedback." Thus, they lack cognitive diversity to obtain collective intelligence. By formalizing this argument using Scott E. Page's 'Diversity prediction theorem',⁷ I show that Silicon Valley's preoccupation with Singularity is a sign of a system driven by herd behavior. Here,

collective intelligence is understood as an outcome that follows the logic suggested by Malone, Laubacher and Dellarocas (2010): If we can understand the parts that make up collective intelligence, we can design powerful systems.⁸




This essay will outline different interpretations of growth that together present a diversity of models on how AI may unfold. Second, I explain how a lack of cognitive diversity is related to herd behavior – and how that may have implications for the system. One of the essential implications of the subsequent reductionism is explicated by drawing on Page’s diversity prediction theorem. Finally, I tie together the arguments to show how cognitive diversity and the resistance to reductionism are two sides of the same coin – and I emphasize that, collectively, wise decisions on AI can only be assumed if the system building it comprises a cognitively diverse set of individuals.

Different Interpretations of Growth

As Ito noted,⁹ the ‘church of Singularity’ that dominates Silicon Valley is grounded in assumptions of exponential growth. Ito similarly acknowledges that exponential growth presumes positive reinforcement – something that excites Singularitarians and scares system dynamics people, as people outside the so-called Singularity bubble may see alternative trajectories where systems self-regulate and adapt. Put differently, different interpretations exist of the potential growth trajectories of AI technology, and Singularitarians rely upon only one of them, that is, exponential growth of technological advances.

One way to resist systemic reduction of the conceptualization of a given phenomenon is to challenge the underlying tenets of the paradigmatic beliefs of the community. Put differently, networked intelligence is, in part, determined by the exploration of *new* ideas.¹⁰ If the ‘church of Singularity’ believes in exponential growth of AI, then it is relevant to question this base assumption by considering alternative trajectories. Allowing such a comprehensive and multifaceted overview of scenarios is, in essence, cognitive diversity. Moreover, technological advances have historically been described through a multitude of different developmental trajectories and diffusion curves. Consequently, it is reasonable to at least consider alternative trajectories, in order to assess different scenarios and test the verisimilitude of the varying models. It is interesting to note that this overreliance on exponential growth has similarly started to attract criticism from within the start-up community (which Silicon Valley is presumed to represent): For instance, Basecamp founder David Heinemeier Hansson has argued that the exponential curve is toxic for start-ups, and has further posited that exponential growth is “the banality of moral decline”.¹¹ Hence, a more collectively intelligent approach to growth would include alternative trajectories.

Systematically reviewing some of the most common conceptual trajectories can explicate the need for conceptual diversity – and the potential blind spots of the church of Singularity. The following will go through the models of the exponential curve, S curve and hype curve (see the figure below for an overview).

TYPE	MEANING	EXAMPLES	VISUALIZATION
<i>“The exponential curve”</i>	Growth is exponential and will continue - resulting in substantial growth	<ul style="list-style-type: none"> • Moore’s law • Singularity 	
<i>“The S curve”</i>	Growth is gradual at first, then steep, and after a while slows down and evens out	<ul style="list-style-type: none"> • Resembles an S • Diffusion of innovation 	
<i>“The hype curve”</i>	Growth is 'overestimated in the short run, but underestimated in the long run'	<ul style="list-style-type: none"> • Gartner’s hype cycle • Amara’s law 	

The exponential curve: The central tenet of a Singularitarian view of AI is the assumption of an exponential curve of technological advancement. This perspective presumes that growth happens exponentially, and therefore, what may initially be perceived as a small and insignificant development will suddenly ‘explode’ in virulent outcomes due to positive feedback mechanisms. The prime example in technological cycles is Moore’s law, that is, the observation and subsequent prediction that the number of components per integrated circuit doubles every year.¹² When analyzing AI advances from this perspective, the Singularity hypothesis makes perfect sense in its logic.

The S curve: The concept of an S curve is widely adopted among academics and practitioners, but it has similarly been utilized inconsistently, as it has been applied to explain both technological progress as well as market share. When referring to technological progress, the notion of an S curve is very closely related to the concept of ‘diffusion of technological innovation’, where the S curve can describe the technological advancement – and the diffusion of innovation describes the demand for the technology.¹³ Hence, the S curve explicates the intricate relationship between technological progress on the one hand, and the demand and diffusion among consumers of the technology on the other. Viewed from this perspective, AI technological advances may follow the shape of an S (little progress at first, then steep progress, and finally the progress evens out). However, the technological advances (and where we end up on the S) are very much dependent on the demand among AI consumers – and the extent to which AI remains among innovators and early adopters – or if it moves to the mainstream

market. Hence, the future of AI is less certain and positive from this perspective – and it is therefore inconsistent, and challenges, the Singularity movement’s logic.

The hype curve: The hype curve does not explicitly emphasize technological advances, but rather, the accompanying ‘hype’ of stakeholders (building on Gartner’s hype cycle).¹⁴ However, it is included here, as it is also very closely related to *Amara’s law*, stipulating that “We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run.”¹⁵ Hence, initial growth may be inflated, resulting in a subsequent dip – but over the long run there will be steady growth. Seen from this perspective, AI may merely be experiencing a current hype phase, which will eventually end. But in the long run it may steadily advance (although not exponentially). Taking the hype curve perspective, Singularity may be symptomatic of inflated expectations, but it may not necessarily be a wrong hypothesis in the long run. It will, however, take much longer time than anticipated to become true.

Looking at a new technological development through the various curves is useful, as (1) it provides a comprehensive overview, (2) it reduces the risk of blind spots, and (3) it can point toward the most likely outcome. Such an analysis has previously been utilized on blockchain and bitcoin.¹⁶ The present essay makes no forecasts or judgments about which curve will best describe the development of AI. However, the present essay *does* emphasize that a variety of trajectories are possible, and that a system entailing collective intelligence will maintain the diversity of models rather than lean heavily on a single model which would be symptomatic of reductionism.

Cognitive Diversity vs. Herd Behavior

A potential consequence of a monolithic view in favor of exponential growth is that it blinds alternative explanations and predictions. In the above example, two models provide alternative trajectories that highlight some of the barriers to growth. These barriers remain absent in the exponential growth model. Hence, decision makers, developers, and entrepreneurs that do not consider all of the different development trajectories, at least to some extent, may very well be taking a reductionist stance on the issue of AI development.

A reductionist tendency, caused by a lack of cognitive diversity, can have three important implications:

1. it can lead to counterfactual interpretations
2. it can lead to overconfidence and optimism bias when predicting the future of AI, and
3. it can result in the oversight of potential barriers to development and growth. Of course, these issues are somewhat related: point 2 can be symptomatic of point 1.

However, in order to fully explicate my argument, the implications are untangled into these three distinct issues. These issues will be further unfolded in the following.

The first problem with a lack of cognitive diversity, and the resultant reductionism, is that it can result in counterfactual interpretations of ongoing developments in the marketplace and technological environment. It similarly alludes to the counterfactual interpretations that may arise when a single mental model (i.e. singularity and the exponential curve) dominate:

“Whether you are on an S-curve or a bell curve, the beginning of the slope looks a lot like an exponential curve [...] Most people outside the Singularity bubble believe in S-curves, namely that nature adapts and self-regulates and that even pandemics will run their course.”

Two aspects are interesting in this description. First, it explains how Singularitarians may (wrongly) interpret early signals arising from the technological development as evidence of an exponential curve, although they may in reality be on an S-curve or bell curve. In other words, as the beginning of the slope is similar in all of the previously discussed curves, having a monolithic mental model of exponential curves may lead to oversimplified interpretations of data and flawed conclusions about reality. Second, it explains how the mental model of the exponential curve exists within a so-called ‘Singularity bubble’, although outside observers may believe in alternate explanations such as S-curves. Consequently, it illustrates the problem of a lack of cognitive diversity caused by isolated networks and echo chambers.

The second problem with a lack of cognitive diversity, and the resultant reductionism, is that it can lead to overconfidence and optimism bias when predicting the future of AI. As discussed under the previous point, it can be argued that Silicon Valley puts too much emphasis on exponential growth, which results in a Singularity bubble that filters out alternative explanations, and therefore, becomes symptomatic of a lack of cognitive diversity. A lack of cognitive diversity can lead to flawed decisions and systemic bias. That is, a reliance on a mental model favoring exponential growth may result in overconfident forecasts of technological growth and demand, leading to an optimism bias — particularly if they do not interpret data and early signals with alternative models. The present argument can be formalized mathematically through Page’s diversity prediction theorem. In Page’s work on the benefits of cognitive diversity, he untangles the elements that comprise the accuracy of collective forecasts (also referred to as ‘the wisdom of crowds’). As the concept of Singularity (and its underlying presumptions of exponential growth) is essentially a collective forecast of how AI will develop, it will be relevant to apply the lenses of the diversity prediction theorem to see what it has to say about this collective forecast. The basic explanation of the theorem is as follows:¹⁷

Collective error = Average individual error - Prediction diversity

Now, what does this theorem actually mean? It means that if we want to build a group of forecasters (say, technologists from Silicon Valley who should forecast the development of AI), then their collective accuracy in making group predictions (as measured by ‘collective error’) depends on how much the group members know about the subject (as measured by the ‘average individual error’) as well as **how differently they interpret the issue and the resultant forecast differences among the group members** (as measured by ‘prediction diversity’). Put differently, a group that makes accurate collective forecasts depends on its group members being accurate (because they know about the topic they predict about) as well as differences between the forecasts of its group members (they vary in predictions because they each have a different interpretive model). The theorem similarly states that individual ability and diversity in predictions are equally important for the collective accuracy of the group. Hence, cognitive diversity matters just as much as ability. Albeit the members of the AI community in Silicon Valley are undoubtedly skilled at a level that is arguably almost impossible to match, it is reasonable to believe that their Achilles’ heel may lie in their lack of cognitive diversity, at least if it is judged by the overreliance on exponential models and a belief in Singularity. If we were to apply the diversity prediction theorem on Silicon Valley’s predictions of AI, it could be justified that the reliance upon assumptions of exponential growth could lead to overconfidence in the form of an optimism bias. That is, a systemic error in favor of beliefs of the growth trajectories of AI advances. Compared to the other potential curves represented in table 1, the exponential curve by far has the most optimistic and dramatic outlook on the evolution of AI. Whereas a more ‘collectively intelligent’ system would take a more nuanced outlook that incorporates forecasts from all three curves, arguably conveying a more pessimistic moderation of the future where, at best, progress will be slower and on a much smaller scale. Such a system would value skilled engineers and computer scientists but balance their perspectives with varied viewpoints from outside the conventional silos that dominate areas such as academia, autonomous vehicles, and AI.

The third problem with a lack of cognitive diversity, and the resultant reductionism, is that it can overlook potential barriers to development and growth. A direct outcome of the previous points is that reductionist approaches to AI may result in oversimplified understandings of technological development. If a single model of development is presumed, and that model constitutes the most optimistic outlier among the prevalent models, then potential barriers are not taken into account by the system that relies upon such a model. The subsequent result is herd behavior – that is, companies and individuals that reinforce their beliefs and take the systemic conformity as evidence for the validity of the direction. Over the long term results will be disappointing as inflated expectations are not met. Such a systematic failure has been seen during both the dot com bubble and the financial crisis of 2008. Herd behavior in favor of the most optimistic model overlooks realistic barriers inherent in the other models, namely two: technological barriers and demand barriers. Technological barriers are relevant to consider when focusing on exponential growth, as the exemplar of

exponential advances in technology, Moore's law, has recently been declared dead due to technological limitations.¹⁸ Demand barriers are similarly relevant to consider, as (1) it is logical to believe that a certain maximum level of demand may exist in a market,¹⁹ and (2) it may be difficult to move from a technologically advanced segment to a mainstream market, or laggards.²⁰ Hence, both the technology itself, and the demands of a market, pose pressing barriers to growth – particularly exponential growth. A failure to acknowledge this will often in itself lead to a failure unless you are right in your initial assumptions of the growth trajectory being exponential.

As shown above, the reductionist tendencies, caused by a lack of cognitive diversity among Silicon Valley mega companies' conceptualization of AI development, can have three important implications: (1) it can lead to counterfactual interpretations, (2) it can lead to overconfidence and optimism bias when predicting the future of AI, and (3) it can result in the oversight of potential barriers to development and growth. When combined, these implications comprise an unfortunate intellectual environment within which AI is being developed and commercialized. Put differently, the lack of cognitive diversity – resulting in conceptual reductionism – constitutes a form of herd behavior, known from financial bubbles and collective or functional stupidity.²¹

Cognitive diversity = resisting reduction

When a system discourages active reflection and challenging the underlying assumptions that drive the system's behavior, the outcome will necessarily be reductionism, and reductionism can result in the mindless behavior that characterizes “functionally stupid” systems.²² An openly acknowledged solution to this issue is cognitive diversity and the inclusion of alternative viewpoints and theories.²³ Put differently, cognitively diverse viewpoints could improve the system by making it resistant towards systemic reduction.

Although it has previously been emphasized that one way to improve AI in Silicon Valley is to diversify the people who build it,²⁴ this argument tends to emphasize demographic diversity over cognitive diversity. Demographic diversity is undoubtedly important for a variety of valid reasons, but this essay addresses and demands another kind of diversity – that is, cognitive diversity. Cognitive diversity can help provide a better system of building and managing AI than the present model carved in Silicon Valley, namely because it allows for resistance to reduction.

The bottom line is that having a system that comprises cognitive diversity among its members will secure the needed resistance to reduction, as alternative perspectives will continuously be considered.²⁵ Only when having this cognitive diversity on board can we, realistically, expect collectively wise decisions on the development of AI. Neglecting this argument will lead to an AI paradox which entails that *a collectively unintelligent system is building artificial intelligence.*

Footnotes

1. <https://futurism.com/singularity-explain-it-to-me-like-im-5-years-old/>
↑
2. <https://pubpub.ito.com/pub/resisting-reduction>
↑
3. <https://pubpub.ito.com/pub/extended-intelligence>
↑
4. https://books.google.dk/books/about/Superminds.html?id=QeOzDwAAQBAJ&redir_esc=y
↑
5. <https://dspace.mit.edu/handle/1721.1/105080>
↑
6. <https://sloanreview.mit.edu/article/building-a-more-intelligent-enterprise/>
↑
7. <https://press.princeton.edu/titles/8757.html>
↑
8. <https://sloanreview.mit.edu/article/the-collective-intelligence-genome/>
↑
9. <https://pubpub.ito.com/pub/resisting-reduction>
↑
10. <https://www.technologyreview.com/s/525341/social-physics/>
↑
11. <http://www.wired.co.uk/article/david-heinemeier-hansson-startup-growth>
↑
12. https://en.wikipedia.org/wiki/Moore%27s_law
↑
13. https://en.wikipedia.org/wiki/Diffusion_of_innovations
↑

14. <https://www.gartner.com/technology/research/methodologies/hype-cycle.jsp>
↑
15. https://en.wikipedia.org/wiki/Roy_Amara
↑
16. <https://medium.com/@mcasey0827/speculative-bitcoin-adoption-price-theory-2eed48ecf7da>
↑
17. <https://press.princeton.edu/titles/8757.html>
↑
18. <https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/>
↑
19. https://en.wikipedia.org/wiki/Market_saturation
↑
20. https://en.wikipedia.org/wiki/Crossing_the_Chasm
↑
21. <https://phys.org/news/2016-06-discusses-functional-stupiditywhen-smart-people.html>
↑
22. <https://www.city.ac.uk/news/2016/june/the-stupidity-paradox>
↑
23. <https://www.fastcompany.com/3066345/facebooks-vp-of-engineering-cognitive-diversity-is-the-most-powerful-tool>
↑
24. <https://www.technologyreview.com/s/610637/for-better-ai-diversify-the-people-building-it>
↑
25. <https://hbr.org/2017/11/great-corporate-strategies-thrive-on-the-right-amount-of-tension>
↑