Journal of Artificial General Intelligence 4(3) 153-163, 2013
DOI: 10.2478/jagi-2013-0009

Submitted 2013-06-01
Accepted 2013-12-09

Will We Hit a Wall? Forecasting Bottlenecks to Whole Brain Emulation Development

Jeff Alstott JEFFALSTOTT@GMAIL.COM
732 Park Road NW
Washington, DC 20010, USA

Editor: Randal Koene, Diana Deca

Abstract

Whole brain emulation (WBE) is the possible replication of human brain dynamics that reproduces human behavior. If created, WBE would have significant impact on human society, and forecasts frequently place WBE as arriving within a century. However, WBE would be a complex technology with a complex network of prerequisite technologies. Most forecasts only consider a fraction of this technology network. The unconsidered portions of the network may contain bottlenecks, which are slowly-developing technologies that would impede the development of WBE. Here I describe how bottlenecks in the network can be non-obvious, and the merits of identifying them early. I show that bottlenecks may be predicted even with noisy forecasts. Accurate forecasts of WBE development must incorporate potential bottlenecks, which can be found using detailed descriptions of the WBE technology network. Bottlenecks identification can also increase the impact of WBE researchers by directing effort to those technologies that will immediately affect the timeline of WBE development.

Keywords: whole brain emulation, forecasting, bottlenecks, technology network,

1. Introduction

Whole brain emulation (WBE) is the replication of the dynamics of the human brain, at a level of description and accuracy sufficient to reproduce human behavior. Attention to WBE has increased over the past few decades. This interest is in part due to the potential ramifications of WBE, be it scientific, philosophical, or economic (Sandberg and Bostrom, 2008; Fiala, 2002; Shores, 2011; Izhikevich and Edelman, 2008; Hanson, 1994, 2009). However, some interest in WBE is driven by perceptions that it is indeed possible. Forecasts of the technology development necessary for WBE, such as Moore's law predicting increasing computational capacity, have created an expectation that WBE will arrive within the century. However, WBE is a complex endeavor with a complex network of prerequisite technologies. Most forecasts only consider a fraction of this technology network. The unexamined portions of the network may have bottlenecks that would slow the development of WBE. In order to effectively forecast WBE development, the entirety of the technology network must be rigorously described and forecasted.

Multiple analyses of how to create WBE have been put forth in recent years. Some have described the tools and implementations available today (Deca, 2012; Cattell and Parker, 2012). Others have been process oriented, describing frameworks for how to iteratively examine the brain and build software or hardware models of it (Koene, 2012; Sandberg and Bostrom, 2008; Parker, Friesz, and Pakdaman, 2006). These strategies for creating WBE have been described as roadmaps, because they mark out the steps for getting to WBE from where we are now. These roadmaps answer the question of how WBE may be developed. Another important question, however, is when WBE will be developed.

Forecasts of when WBE (and AGI) will be developed are already common, more common than the process-oriented roadmaps. These forecasts generally consider only one requirement for WBE: computational capacity. These forecasts calculate a number for the computational requirements of emulating the brain, then use Moore's law to project when the required computing power will be reached. However, WBE requires multiple prerequisite technologies to be developed. Therefore, predicting the timeline of WBE progress requires forecasting the development of each prerequisite technology. Closer analysis of the progress of WBE-related technologies can provide insights into which prerequisites will arrive last, and thus would be bottlenecks to WBE development. In order to create WBE faster, these bottlenecks could be targeted for additional attention and resources. Thus, forecasting WBE development can give strategic direction for how to proceed towards successful WBE.

In what follows, I introduce a framework for describing technology development: technology networks. I discuss how technology networks make clear the difficulties of forecasting complex technological developments such as WBE. I then argue that in order to make effective predictions of WBE development, we must make forecasts of each branch of the network, so as to identify bottlenecks. This will require a detailed technology network with close examination of each component. Lastly, I present a simple model illustrating the impact of bottlenecks on forecast accuracy. The presented understanding of bottlenecks and how to find them may improve the work of WBE forecasters and inform WBE developers.

2. Technology networks

Virtually every human technology rests on other technologies. In order to make an aluminum can, one must be able to make aluminum. In order to make a microscope, one must be able to make lenses. In order to make a WBE, one must be able to make a lot. When describing how to develop a technology, then, it is a useful framework to consider its prerequisite technologies. Each of those technologies then has its own prerequisites, and so on. This framework is a technology network (Figure 1). Technology networks bear resemblance to strategic planning maps (Dortmans, 2005), as well as maps of industry interdependency (Hausmann et al., 2011). However, technology networks were first developed and popularized in the Civilization board game and computer game series in the form of technology "trees" (Tresham, 1980; Meier and Shelley, 1991). Network frameworks have been used both to describe and model technology development in many domains (Tëmkin and Eldredge, 2007; Valverde and Solé, 2007; McNerney et al., 2011; Solé et al., 2013).

"Technology" is broadly defined in a technology network. Technologies include not only directly applied engineering capabilities but also basic scientific knowledge. It is much easier to make nuclear power if one understands nuclear physics. Technology network branches can merge and divide, with a single technology being relevant for developments in many fields. There can also be alternative routes to a technology, where different capabilities produce the same end product.

The elements of a technology network can be described at low or high resolution. That is, a technology network model could include "rockets enable space travel", or it could include "small rockets enable big rockets, and big rockets enable space travel," or it could have even more detailed components. The reality of technology progression occurs at a very fine-grained level,

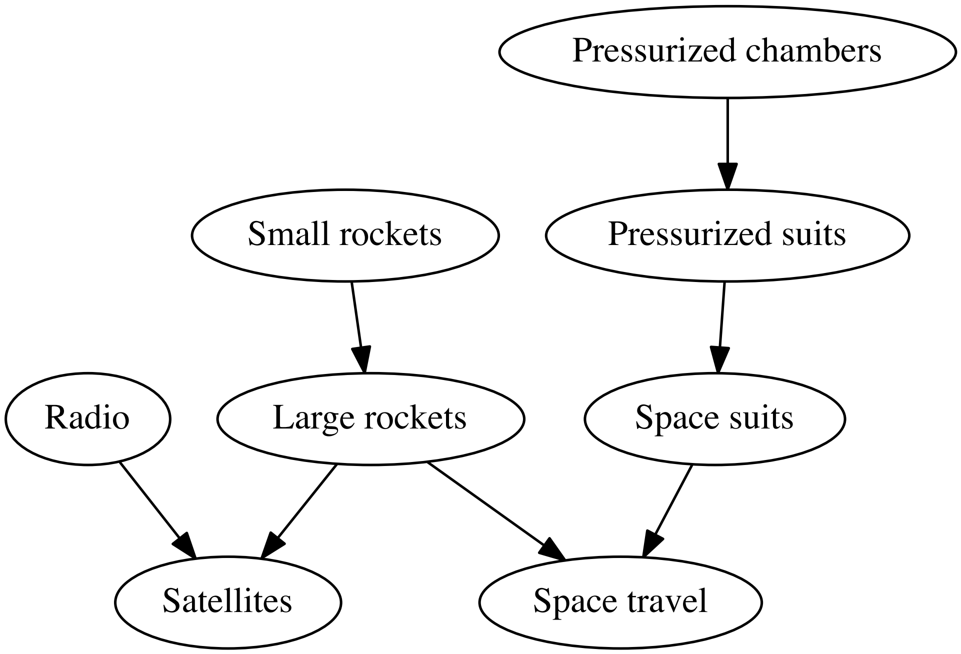

Figure 1: **A simple example of a technology network.** This network is a high-level description of some of the necessary developments for space-based technologies, with technologies enabling other technologies. It illustrates that a single technology can have multiple prerequisites, and that a single technology can enable multiple future technologies. "Technology" is broadly defined, and includes both engineering and scientific achievements. A more detailed and accurate network could be constructed by incorporating additional branches or dividing components of each branch into smaller steps.

with individual investigators making small advances on sub-projects and sub-sub-projects. The professor cannot figure out how planets form around stars until the graduate student figures out how to properly align the telescope and convert the raw data into a format they can interpret. Some may not wish to call this low-level reality "technology" development. Regardless of terminology, however, it is the process of accreting small, interdependent capabilities and knowledge that lead to large-scale technological advances.

It should now be clear that technology networks can describe a huge space. The higher the resolution of the technology network, the more individual advances unveil as necessary to get to a particular technology. Additionally, we observe more interdependencies; the astronomy graduate student was observing solar systems, but required engineering and computer science to evaluate the data. The great complexity of technology networks make them a difficult space to map out and accurately describe. And then we must remember that portions of the network which we are presently interested in – those that lead to WBE – don't yet exist.

3. Forecasting and Planning with Technology networks

"Prediction is very difficult, especially about the future."[1] Predictions of how civilization will behave at a later date have become plentiful in the last century. These predictions have historically not been accurate. When a prediction forecasts further into the future it must describe more steps in the technology network, and thus has more chances to make a mistake. Through simple compounding of individual errors, predictions of even 100 years out can become terribly inaccurate.

The problem of compounding errors can be ameliorated by focusing on one well-defined branch of the technology network, isolated from the others. By examining the behavior of just that branch, it may be possible to develop an accurate theory of how that branch will develop in the future. An example of forecasting a single branch is Moore's law. Moore's law observes a historical regularity in a few quantitatively defined elements of computational power, and projects that regularity into the future. What is unusual about Moore's law is that this precise quantitative prediction has held for the past 50 years. We will return to the relevance of Moore's law for WBE in the next section.

Forecasts of a single branch are more likely to be accurate than predicting the entire technology network. However, they are not perfect. Commercial airline cruising speeds increased at a steady rate after World War I, leading to predictions that transatlantic flights could soon take just an hour. But commercial passenger aircraft speed halted in the 1960s, just above Mach 0.8 (Lewis and Niedzwiecki, 1999). The reason for the sudden stop is that "a single technological branch" is a shaky concept. The technology network framework makes clear that technologies can be interdependent, forming interconnected branches. A branch may look like it is progressing independently, until it isn't. For commercial airlines, the primary impediment to increased cruising speeds was and is fuel efficiency; increased speeds are possible, but fuel costs are too high for commercial viability (Dunn, 2009; Peeters, Middel, and Hoolhorst, 2005). If progress were made on another technological branch to lower the cost of fuel or raise thrust fuel efficiency, then airline speeds could again increase.

The cost of thrust is a bottleneck for commercial airline speeds. Airline speeds can increase only as fast as the cost of thrust is improved, until, perhaps, another bottleneck is reached. By mapping out the technology network for airline speed, subcomponents like fuel costs can be more readily identified. By forecasting the speed of development in those subcomponents, bottlenecks could be identified before they happen. If a goal in 1950 was to see airline speeds increase to Mach 3, creating forecasts of fuel costs and fuel efficiency should have also been a goal. If those forecasts were accurate, then the fuel bottleneck could have been planned for. For example, additional research could have gone to fuel efficiency or alternative fuels to widen the bottleneck before it was reached.

Technology networks are dense and complex, and therefore statistically unlikely to be accurately predicted in their entirety. Predictive power may be increased by restricting forecasts to slim portions of a network. However, technology branches are rarely independent. Portions of the network which initially appear to lack dependencies likely rely on other branches which have not yet been described. If these other related branches are not developed, bottlenecks can occur on the primary branch (Figure 2). Thus, there could exist bottleneck technologies that would slow the

1. Original author unknown. Even accurate statements about the past can be difficult.

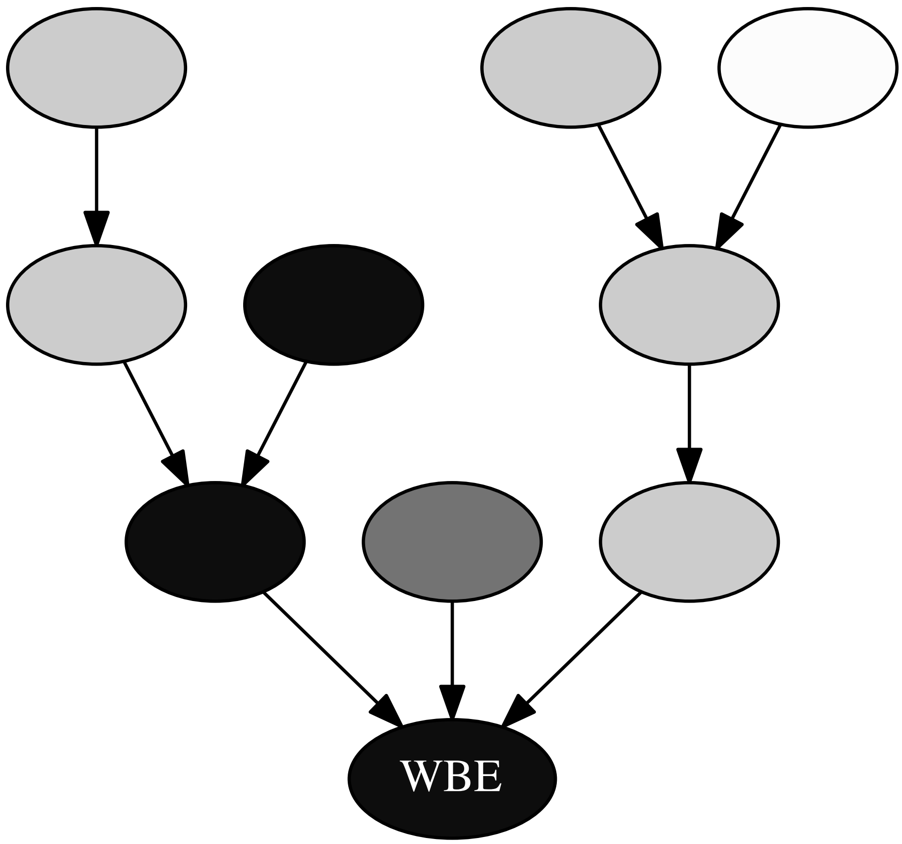

Figure 2: **Diagram of the bottlenecking effect on a WBE technology network.** Technologies (nodes) have prerequisites that are developed at different speeds (shades; darker shades take longer to develop). A technology can only develop once its slowest prerequisite is created, which may be far back in the network.

arrival of WBE until they are developed. Targeted progress in a single field may be insufficient to bring about WBE. If just one necessary section of the network is incomplete, the project will stagnate. It is possible we could identify and plan for these bottlenecks before they happen, with a sufficiently accurate WBE technology network.

4. The WBE Technology network

There are many possible ways to implement WBE, and so there are different collections of technology that would together enable WBE. Sandberg and Bostrom described possible WBE implementations and identified one method as most likely to succeed first: scan the structure of the brain at the subcellular level, then translate that scan to a software model of the brain structure, then use that model structure to simulate neuronal activity (Sandberg and Bostrom, 2008). Each of these steps is currently nontrivial, and will require developing a suite of capabilities to enable WBE.

In order to describe both the technologies for WBE and their prerequisites, Sandberg and Bostrom created a technology network (Figure 3). To date it is by far the most complete and systematic analysis of the technology needed for WBE. This network is high level, describing the broad capabilities necessary for scanning, translation, and simulation. It also identifies possible drivers of new development in each area, including Moore's law, commercial viability, and interest from the research community. This technology network is suitable for conveying the broad strokes of WBE's requirements, but it is insufficient for forecasting WBE's development. The central reason for this is that quantitative descriptions of development time only cover a portion of the network: those driven by Moore's law.

Moore's law, as illustrated by the technology network, describes a slim portion of the technology needed for WBE. However, most predictions of when WBE (or, relatedly, AGI) will be developed simply identify some number for the computational requirements for WBE, then plot out Moore's law until the computational capacity per some dollar value reaches that target. When Moore's law says the computational requirements will be met, that is the year WBE (or AGI) will be built. Such predictions do indeed put a kind of lower bound on when WBE will be created (Sandberg, 2013). However, even assuming that Moore's law will continue to hold for another century, a technology network framework lays bare that such forecasting is too simplistic. WBE will rely not only on certain computational capabilities, but also scanning, translating and modeling capabilities. Those capabilities very well could develop more slowly than Moore's law, creating a bottleneck. It is tempting to put particular emphasis on the computing power gained from Moore's law, due to the idea that it is the foundation upon which the entire WBE project rests. But this is an inaccurate model. There is no one foundation; every component technology of WBE is a foundation. Without all the pieces the project crumbles.

5. Forecasting Bottlenecks Piece by Piece

Forecasting the entirety of a large technology network with precision is virtually certain to fail. Planning exactly how all elements of WBE will be developed is unlikely to succeed. However, a rough understanding of the technology network can be sufficient to find potential bottlenecks. For each technology enabling WBE, we must quantitatively describe both the functionality necessary for WBE and the rate of progress to that goal. This has been the strategy for forecasting WBE with Moore's law. The difference is that this must be done for every piece of the network.

An example of such an analysis was performed by Sandberg and Bostrom on the discovery of neurotransmitters and neuromodulators (Sandberg and Bostrom, 2008). A complete catalogue of these ligands is probably necessary for WBE, but neuroscientists are still finding about 8 per year. However, an upper bound on the number of neurotransmitters and neuromodulators can be inferred from the genome. At historic rates the complete catalogue will be finished by the 2030s.

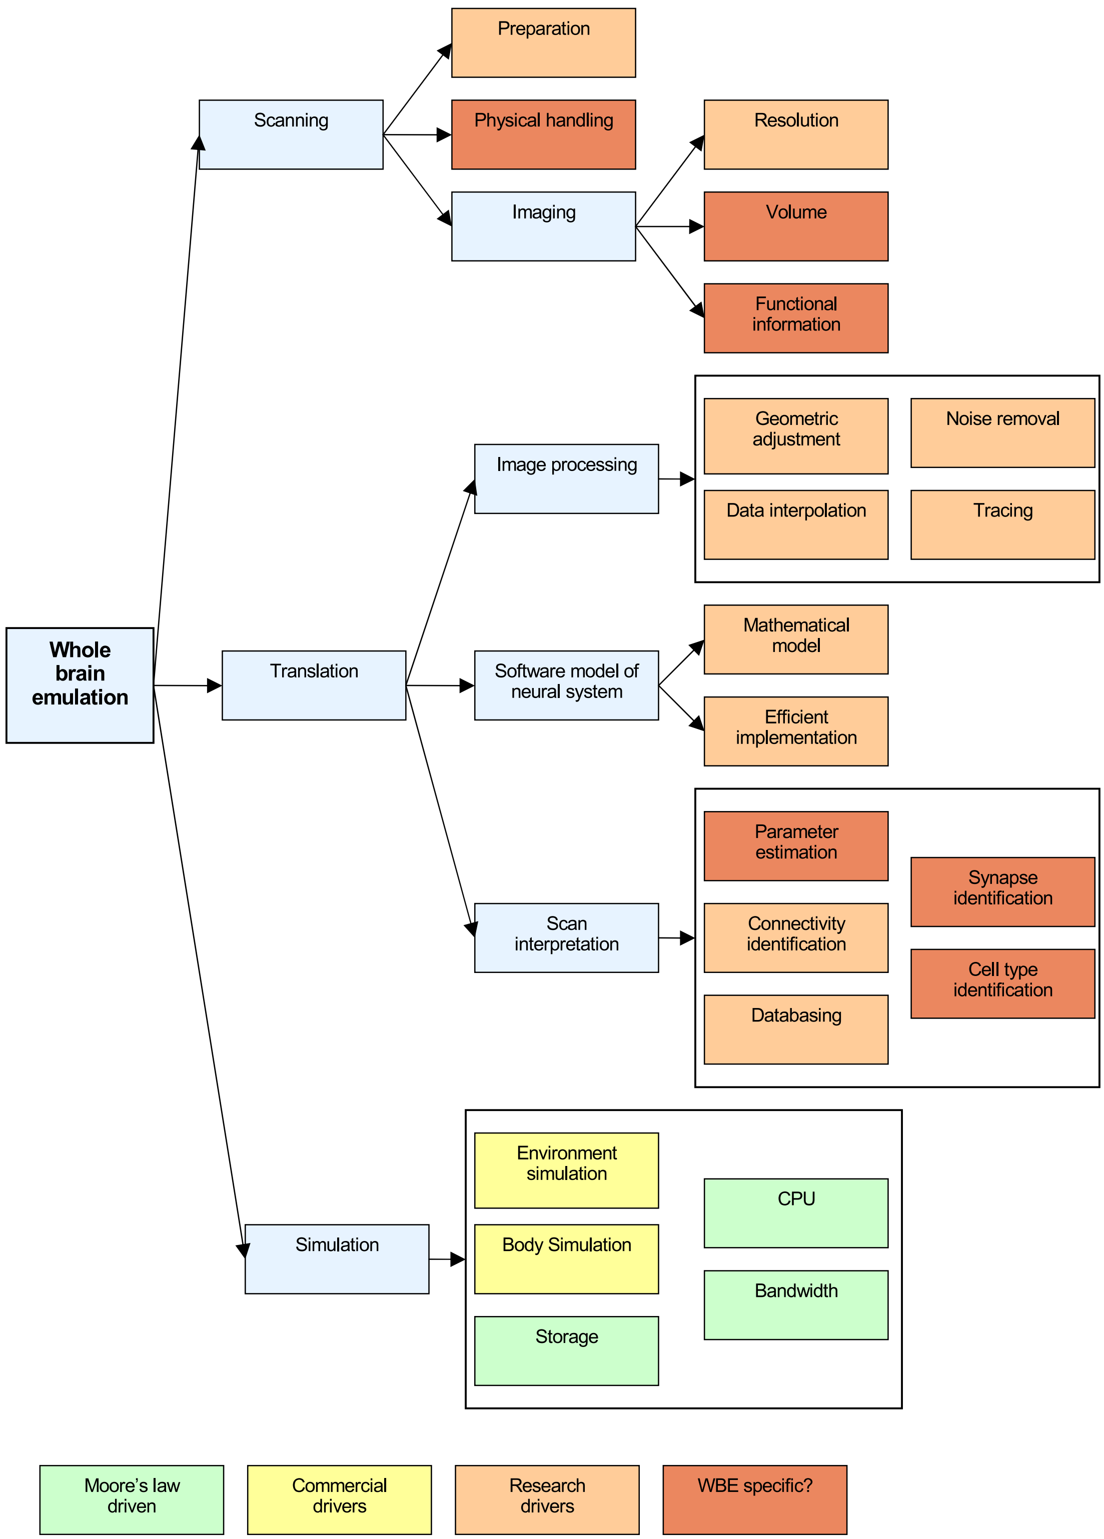

Figure 3: **A WBE technology network describes the branches required for WBE development.** Box shades correspond to forces that will likely lead to development of each technology. From Sandberg and Bostrom (2008).

The example of neurotransmitters illustrates that much of the necessary forecasting may happen on a very fine-grained level. In the WBE technology network (Figure 3) “Neurotransmitter cataloguing” would be just a subcomponent of a subcomponent of a branch leading to both “Functional information” and “Mathematical model”. It seems plausible that most accurate forecasts would be on such a low level. Moore’s law only actually predicts the behavior of the number of transistors per integrated circuit, and Kryder’s law only predicts the storage density of hard disks (Walter, 2005). Perhaps a reason these forecasts have been so successful is that they are limited in scope. Identifying bottlenecks to WBE, then, may necessitate making such low-level descriptions and predictions of scientific progress.

While Moore’s law is fine-grained, it is also high impact. This is because the particular variable of transistor density affects many other systems. Indeed, if a technology can be developed with simply more computational power, its progression can be hooked to Moore’s law and reach exponential rates (Kurzweil, 2012). It may be possible in other fields to find similarly precise-but-impactful trends. For example, a forecast from biology could be that the number of cells covered in a perfect computational model of the cellular life cycle doubles every year[2]. Such a forecast would predict an accurate model of the full life cycle of every cell in the brain by the middle of the century. Even an impactful prediction such as this could potentially be intercepted by a bottleneck from another area, however. These models might capture life cycle dynamics very well, but not any of the electrodynamics necessary for neurons in WBE. The bottleneck would then be the rate of increase in the number of cells modeled with both accurate electrodynamics and life cycle. In order to ensure that bottlenecks are found before they occur, a detailed technology network must be made, with perhaps orders of magnitude more components than in Figure 3.

Simply describing such a large network would be a feat. As previously discussed, making very accurate forecasts for each element of the network would likely be impossible. However, it is not necessary to precisely predict development rates nor the ultimate technology level need for WBE. All that is needed is a rank order of when technologies will reach the target level; those technologies that will arrive the latest are the bottlenecks. Even if the rates or target predictions have noise, the rank order can be recovered. Figure 4 shows the behavior of a simple model of 10,000 technologies, each with a rate of progress and a distance to the level needed for WBE. The rate and distance combine to create a time left to completion, which is a heavy tailed distribution (Figure 4A). Noisy estimates of both the distance and the rate are used to calculate an estimated rank order of the time remaining to develop the technology. Even with noisy estimates of distance and rate, the estimated rank order can be strongly correlated to the true rank order (Figure 4B). Thus, estimated bottlenecks strongly predict the true bottlenecks. Even with added noise of the same magnitude as the true distances and rates, correlations do not drop far below 0.4 (Figure 4C). This model is simple, but illustrates that bottlenecks can be identified even with noisy prediction.

6. Conclusions

WBE would have great impact on humanity, and accordingly WBE has received much analysis of how and when it will be developed. Forecasts of the timeline of WBE creation are difficult, however, because WBE would require a large suite of constituent technologies. This set of requisite technologies and their own prerequisites form a technology network which must be developed

2. At present, one single-celled organism has been modeled extremely well, though far from perfectly (Karr et al., 2012).

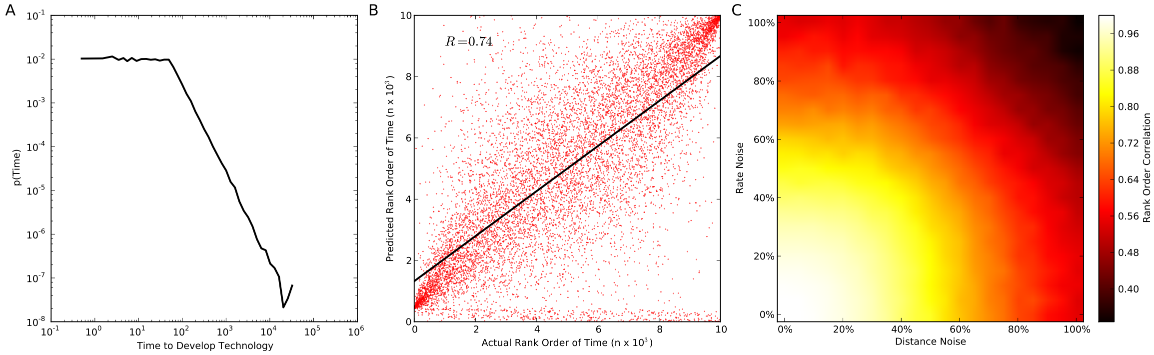

Figure 4: **A simple model of predicting technological bottlenecks shows estimates are robust to noise.** 10,000 technologies were generated with a rate of progress and a "distance" to the level needed for WBE (randomly drawn from a uniform distribution of 0 to 10 and 0 to 500 arbitrary units, respectively). Estimates of the time to completion were made from noisy estimates of the rate and distance. Gaussian noise was added to the true values, with a scale ranging from 0% to 100% of the true value. A) Time remaining to completion of the WBE requirement followed a heavy-tailed distribution, spanning many orders of magnitude. B) With noise scaled to 50% of the signal, estimates of the technologies' rank order of time to completion strongly correlated with the true values. C) The strength of the correlation between estimation and truth was a function of the noise in the rate and distance estimation. Even with noise scaled to 100% of the true values, the correlation of the time to completion rank order stayed near 0.4.

before WBE can be achieved. Accurate forecasts of WBE development must account for all the elements of this network. As WBE will only be realized when all the necessary components are in place, the slowest-progressing components of the technology network will determine when WBE arrives. These slowly-arriving technologies are bottlenecks, and can likely be identified beforehand even with noisy predictions.

For those trying to develop WBE, identifying bottlenecks before they arrive will be fruitful. Potential bottlenecks can potentially be widened by devoting more attention and resources to the slow technologies in an attempt to speed up their rate of progress. It is also possible that a bottleneck could be avoided by using a different technological strategy, taking a different path down the technology network that does not hit the bottleneck.

Creating a detailed WBE technology network and forecasting the constituent elements would require specialized knowledge of a great many areas. If WBE developers change their strategy to widen or avoid bottlenecks this will change the rates of progress, and thus the forecasts should be updated. Even if there is not deliberate intercession by WBE developers, forecasts will likely need updating over time as new data arrives. Forecasting WBE is therefore no small undertaking. But given the potential impacts of WBE, a clear understanding of when and how it may arrive would be valuable for society.

References

Cattell, R., and Parker, A. 2012. Challenges for Brain Emulation : Why is Building a Brain so Difficult ? *Natural Intelligence* 1(3).

Deca, D. 2012. Available Tools for Whole Brain Emulation. *International Journal of Machine Consciousness* 04(01):67–86.

Dortmans, P. J. 2005. Forecasting, backcasting, migration landscapes and strategic planning maps. *Futures* 37(4):273–285.

Dunn, P. 2009. Why hasn't commercial air travel gotten any faster since the 1960s? *MIT Engineering*.

Fiala, J. 2002. Three-dimensional structure of synapses in the brain and on the Web. *Proceedings of the 2002 International Joint Conference on Neural Networks. IJCNN'02 (Cat. No.02CH37290)* 1–4.

Hanson, R. 1994. If Uploads Come First: The Crack of a Future Dawn. *Extropy* 6(2).

Hanson, R. 2009. The Economics of Brain Emulations. In Healey, P., and Rayner, S., eds., *Unnatural Selection: The Challenges of Engineering Tomorrow's People*. Sterling, VA: Earthscan.

Hausmann, R.; Hidalgo, C. A.; Bustos, S.; Coscia, M.; Chung, S.; Jimenez, J.; Simoes, A.; and Yıldırım, M. A. 2011. *The Atlas of Economic Complexity*.

Izhikevich, E. M., and Edelman, G. M. 2008. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences* 105(9):3593–8.

Karr, J. R.; Sanghvi, J. C.; Macklin, D. N.; Gutschow, M. V.; Jacobs, J. M.; Bolival, B.; Assad-Garcia, N.; Glass, J. I.; and Covert, M. W. 2012. A whole-cell computational model predicts phenotype from genotype. *Cell* 150(2):389–401.

Koene, R. A. 2012. Toward tractable AGI : Challenges for System Identification in Neural Circuitry. In *Proceedings of the 5th International Conference on Artificial General Intelligence (AGI)*, 136–147.

Kurzweil, R. 2012. *How to Create a Mind: The Secret of Human Thought Revealed*. Viking Adult.

Lewis, J. S., and Niedzwiecki, R. W. 1999. Aircraft Technology and Its Relation to Emissions. In Penner, J. E.; Lister, D. H.; Griggs, D. J.; Dokken, D. J.; and McFarland, M., eds., *Aviation and the Global Atmosphere*. Cambridge, UK: Cambridge University Press.

McNerney, J.; Farmer, J. D.; Redner, S.; and Trancik, J. E. 2011. Role of design complexity in technology improvement. *Proceedings of the National Academy of Sciences* 108(22):9008–2013.

Meier, S., and Shelley, B. 1991. *Civilization*. MicroProse.

Parker, A. C.; Friesz, A. K.; and Pakdaman, A. 2006. Towards a Nanoscale Artificial Cortex. In *The 2006 International Conference on Computing in Nanotechnology*.

Peeters, P. M.; Middel, J.; and Hoolhorst, A. 2005. Fuel efficiency of commercial aircraft: An overview of historical and future trends. Technical Report November, National Aerospace Laboratory NLR.

Sandberg, A., and Bostrom, N. 2008. Whole Brain Emulation: A Roadmap. Technical report, Future of Humanity Institute, Oxford University.

Sandberg, A. 2013. Feasibility of Whole Brain Emulation. In Müller, V. C., ed., *Theory and Philosophy of Artificial Intelligence*. Berlin: Springer. 251–264.

Shores, C. 2011. Misbehaving Machines: The Emulated Brains of Transhumanist Dreams. *Journal of Evolution and Technology* 10–22.

Solé, R. V.; Valverde, S.; Casals, M. R.; Kauffman, S. A.; Farmer, D.; and Eldredge, N. 2013. The Evolutionary Ecology of Technological Innovations. 18(4):15–27.

Tëmkin, I., and Eldredge, N. 2007. Phylogenetics and Material Cultural Evolution. *Current Anthropology* 48(1):146–153.

Tresham, F. 1980. *Civilization*. Hartland Trefoil.

Valverde, S., and Solé, R. 2007. Self-organization versus hierarchy in open-source social networks. *Physical Review E* 76(4):046118.

Walter, C. 2005. Kryder's Law. *Scientific American*.