FORESIGHT

THE INTERNATIONAL JOURNAL OF APPLIED FORECASTING

Forecasting the Impact of Artificial Intelligence: Another Voice Lawrence Vanston

Response Spyros Makridakis

Forecaster in the Field feature Lawrence Vanston

Forecasting the Impact of Artificial Intelligence: Another Voice

LAWRENCE VANSTON

PREVIEW FROM AUTHOR *Over the past year I have had the pleasure to read* Foresight*'s five-part series "Forecasting the Impact of Artificial Intelligence" by Spyros Makridakis. Recently I was asked to offer my thoughts on these articles as a forecaster with a deep interest in AI. Since I have a lot to say, I agreed.*

First, Spyros has done a masterly job conveying a complex and broad-ranging topic. He has raised many of the crucial issues and made the importance of this topic very clear. There are few people with his insight and experience in forecasting who could have done this, and he has done it well. Here I will add my voice, focusing on

- *areas in which I thought the article was weak (e.g, AI performance forecasts);*
- *areas that Spyros covers well, but are crucial enough to deserve further illumination (AI's impact on employment and the dangers of AI); and*
- *areas of disagreement (brain-computer interfaces, blockchain, and AI for forecasting).*

FORECASTING AI's COMPUTATIONAL POWER

My first reaction to the AI performance forecasts cited by Spyros was "Is that all there is?" For a field as important to humanity as AI, I was expecting more. We can't really blame Spyros, because when his series appeared there weren't many compelling forecasts of AI performance that I could find either. Since then I have made some progress on forecasting AI performance, which I shared last summer at ISF2018 (Vanston, 2018) and will highlight in this section.

A key component of forecasting AI performance is forecasting computational power. Although Ray Kurzweil has argued for a future of superexponential improvement in computer performance (Kurzweil, 2005), the actual evidence supports continued exponential progress along the lines of Moore's law (see **Sidebar A** for a primer on growth curves).

Exponential progress is typical of technologies unless there is a fundamental limit that cannot be overcome, further improvement has no utility, or there is a fundamental change in technological approach. The last often results in a discontinuity or a change in the rate of improvement. The transition from discrete circuits to integrated circuits is an example. That transition increased the rate of improvement and began Moore's law, but it is not evidence of a superexponential trend, as Kurzweil argues. Another example of a discontinuity is the transition from analog modems to broadband, where we saw a dramatic bump in data rates and then a return to the previous rate of improvement. Even when two exponential trends combine, as happened for a while when we increased clock rates and wavelengths simultaneously in optical transmission systems, the combined trend is still exponential. Setting aside a major discontinuity from quantum computing, for example, or the oft-predicted but yet-to-happen death of Moore's law, the logical working forecast is exponential at historical rates.

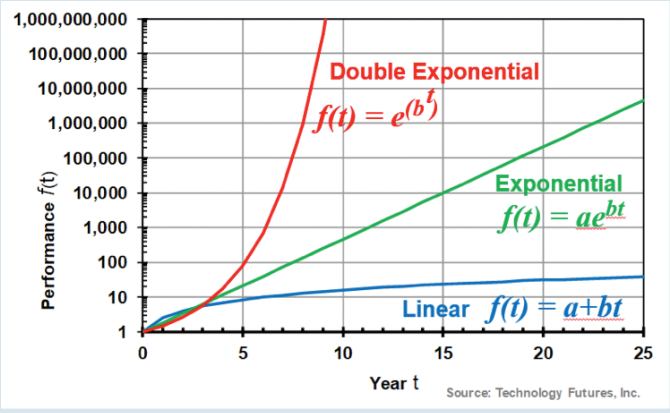

Sidebar A: Growth Rate Primer

With linear growth, performance increases by a constant ***b*** units each year. Linear growth is not typical of performance improvement in high tech, but as we discuss in the text, there is evidence that progress in many AI applications is linear.

With exponential growth, performance increases by a constant percentage r each year, where $r = e^b - 1$. Exponential growth implies a constant doubling time, $d = ln\ 2 / b$. Moore's law is the classic example of exponential growth with a doubling time of two years and an annual growth rate of 41%. Exponential growth yields a straight line when performance is plotted on a log scale as shown in the graph.

With double exponential growth, the annual growth increases exponentially, so it plots on a log scale like an exponential would on a linear scale. This is the curve to use if you want to forecast insane future performance improvement, but such improvement is unsustainable in the long run. It's easy to mistake one-time increases in the exponential growth constant ***b*** for double exponential growth, especially if you're looking for it.

One-time changes in **b** can occur when the technology approach changes (as when we went from discrete to integrated circuits), or when two exponentially improving processes are combined. The latter results in a new exponential curve since the product of two exponential curves is another exponential: $f(t) = a_1 e^{b1t} * a_2 e^{b2t} = ae^{bt}$ where $b = b_1 + b_2$.

A trove of data has been compiled by the Electronic Frontier Foundation for AI applications such as games, image recognition, speech recognition, machine translation, natural language processing, and computer programming (Electronic Frontier Foundation, 2018). These data, along with insightful analyses in blogs by Miles Brundage (2018) and Sarah Constantin (2017), reveal that, in some areas, AI has passed human performance, while in others it has a long way to go. In some cases, deep learning has caused a discontinuity in performance improvement; in others, it simply enabled continuation of the trend. Surprisingly, when measured in the customary units applied to a specific application (for example, the Elo rating for chess or BLEU score for machine translation), AI progress has typically not been exponential, but linear (Vanston 2018).

This suggests that for those AI applications still far below human performance, AI parity may be years away, barring breakthroughs. This probably explains why 92.5% of the AI fellows in Oren Etzioni's survey (Etzioni, 2016) cited by Spyros said superintelligence wouldn't be achieved for 25 years, if ever. And why almost every speech I hear from an AI expert starts with "Forget about AI and human intelligence, current AI has the IQ of a slug. Nothing to worry about here!"—before they move on to whatever they're selling.

Is there really nothing to worry about? See **Sidebar B** for my thoughts on that topic, but we are flying blind until we have better technology forecasts than those provided by Kurzweil (who is a great visionary relying on a highly questionable forecast) and the offhand opinions of experts (who know everything about AI, but may not know much about technology forecasting).

AI'S IMPACT ON EMPLOYMENT

Spyros does a balanced job of laying out the arguments regarding the impact of AI on employment. This issue, and how we handle it, may be as important and more immediate than our existential concerns

about AI. Before you continue, note that I lean optimist/realist even though I sound like a pessimist!

Optimists argue that we have had technological revolutions before, and each time we have ended up with more, but different, jobs. That's true, but each of those revolutions was unique and there have not been so many to make broad generalizations. We can indeed say based on this experience that we *could* end up with more jobs, but saying we *will* without specifics is speculative. It is argued that we just can't imagine the jobs yet. Perhaps… but if you can't imagine it, you can't count on it. Same for happy historical examples such as still having bank tellers in spite of ATMs.

Even if we do end in a happy place of full employment, there's no guarantee it will be a smooth transition. The impacts of employment dislocation can be substantial and consequential. The last American presidential election turned on a few states where the displacement of industrial jobs by technology or offshoring is either a reality or a realistic fear. What's it going to be like when the factories return home, but the parking lots are empty?

AI will doubtless generate many jobs in AI and managing AI, but realistically, how many compared to the jobs lost? And not everyone is suited by talent, disposition, or inclination for those jobs. Optimists say that creative and intellectual jobs are not threatened and that AI will supplement, not replace, them. I don't doubt that AI will help doctors be better doctors and architects be better architects. But suppose AI makes architects twice as efficient at designing buildings. Does that mean we will have twice as many buildings or half as many architects?

Optimists who concede that jobs will be lost look on the bright side and offer the prospect of more leisure time to pursue their personal interests, perhaps with the benefit of a Universal Basic Income. I will never run out of creative, productive, fulfilling things to do, but that's not true for everyone; witness people on unemployment who fall prey to opioids. Economists would say that if you must do income redistribution, it's most efficient to just give people cash. I'm not so sure. Perhaps pushing the money through programs (grants, microloans, fellowships, etc.) and institutions (schools, charities, churches, nonprofits, sports leagues, etc.) makes more sense for some individuals.

Lastly, while AI will impact employment, it likely won't be dramatic in the next 10 years. So we may have some time for forecasters to forecast, analysts to analyze, planners to plan, entrepreneurs to entrepreneur, and think tanks to develop policy positions consistent with their politics. If there are indeed problems, I believe there are solutions. The question is whether we pay enough attention to them.

Sidebar B: Is There Nothing to Worry About?

Are Elon Musk and the late Stephen Hawking, not to mention the other authors cited by Spyros, wrong to be deeply worried about AI? In my opinion, absolutely not, but perhaps for different reasons. My reasons are independent of when and whether computers achieve human intelligence or human consciousness.

1. Rather than the threats of single computers, you should worry about the interactions of millions of interconnected computers, the actions of each dictated by an AI algorithm optimized for its own purposes. Really smart humans managed to design complex derivatives that were impossible for really smart humans to unwind, almost bringing down our financial system. At least we know how those algorithms reasoned. Is the same true for AI?
2. Emerging technologies provide computers with the types of inputs and outputs that heretofore only humans and fellow animals had. This not only provides data fodder for machine learning, but assumes you can just pull the plug or take out the battery. (Try that on your iPhone!)
3. Further, AI is advancing to the point of being able to use those skills in spite of the second half of Moravec's Paradox (Moravec, 1988)—that it's "difficult or impossible to give them [computers] the skills of a one-year-old when it comes to perception and mobility." Our

machines are already approaching a one-year-old's level in many ways. Consider this list of "easy for one-year-olds, hard for computers" skills: "recognizing a face, moving around in space, judging people's motivations, catching a ball, recognizing a voice, setting appropriate goals, paying attention to things that are interesting; anything to do with perception, attention, visualization, motor skills, social skills" (Wikipedia, 2018). On some of these measures our machines seem equal to a one-year-old and working on the terrible twos.

4. This brings us to the other half of Moravec's Paradox: "It is comparatively easy to make computers exhibit adult-level performance on intelligence tests or playing checkers." As machines begin to reach fundamental levels of sensory perception, motor skill, rudimentary language skills, learning ability, and pseudo-reasoning, how far is it really to human intelligence? If what sets us apart from other vertebrates is massive amounts of pliable, general-purpose wiring, is that really so hard to duplicate? Add in the human knowledge base, already conveniently compiled, and we could be a lot closer than is now apparent.

5. Finally, AI doesn't need human intelligence or human consciousness to be dangerous. Microbes with sub-slug intelligence kill millions of us and theoretically could kill us all, so AI isn't the first intelligence to challenge humans! Once we leave programming decisions to AI we don't know what intelligence—possibly far inferior to ours, but just as dangerous—will emerge. Nor do we know whether its evolution will be on pace with vertebrate evolution or microbal evolution.

These concerns exist in the absence of human greed, malice, and bias, intentional or otherwise. In their presence, they are deeply troubling. There are indeed ways to address these problems. My fear regards the will, not the way. It's hard to say no to useful, cool technology. In the U.S. we have seen our own technology used against us to influence the fundamental course of events, including our global alliances and trade relationships, even without AI. The problems arising from AI may be harder to fix. Will we be quick enough? And, in the current winner-take-all world, what are the consequences if we aren't?

BLOCKCHAIN, IA, AND FORECASTING

Most of Part 4 of the Makridakis series is devoted to blockchain. That Spyros put blockchain on the same plane as AI surprised me, and I disagree with so many of his observations and conclusions that the margins of my copy are full. Here are my main objections:

AI doesn't need blockchain. Neither does IoT or any other cutting-edge technology, so why link them? Blockchain is an interesting character, but it is not crucial to the AI story.

Blockchain claims to have the potential to replace banks and other financial intermediaries, even government institutions. This outcome assumes that the recording of transactions is the most important aspect of the financial entity's business. However, anyone who has bought a house knows that recording the deed at the county courthouse (or the equivalent in your country) is a tiny part of the transaction. The institutions involved—agents for the buyer and seller, the title company, the lending bank, federal lending agencies, etc.—provide expertise, broad trust, fiduciary responsibility, insurance, money in the form of loans and loan guarantees, and force of law.

Similarly, the main function of Amazon and the source of its "monopolistic/oligopolistic power" is not the recording of the financial transactions. If it were, Amazon could easily be replaced, with or without blockchain. The reason it has not is that for a reasonable price Amazon provides a host of useful services and a form of trust that is far broader than blockchain's narrow definition of trust.

Blockchain is basically a narrow technology attempting to replace existing technologies that are already providing transaction documentation quickly, reliably, and cheaply. Maybe it will be widely adopted, maybe it won't, but it won't change the world. And if it does change the world, it may be for the worse; the Internet is a dangerous place.

INTELLIGENCE AUGMENTATION

Part 4 ends with a discussion of intelligence augmentation, in which Spyros views AI and humans helping each other, leveraging each other's strengths. Nothing new or controversial there—the future is an extension of the past. He then links the future of IA with direct human-machine interfaces. The latter technology may be closer than we think and may have immediate medical, military, and other benefits. However, the risks and benefits of direct human-machine interfaces need to be assessed on their own, regardless of AI. In fact, this technology may increase our susceptibility to AI domination as much as ameliorate it.

THE LONG-TERM FUTURE

Part 5 of the series starts with a discussion of the pitfalls of long-term forecasting and the top 10 emerging trends. I find a lot here to agree with. I absolutely believe in the necessity of steering between the shoals of being too conservative or too optimistic. That's why I never start a forecast without doing a drivers and constraints analysis (see **Sidebar C**) that usually includes the four factors Spyros lists, among others. Along with AI itself, half of his top 10 trends are on my own list of top information/communications technology trends, and the others are areas I know well enough to agree with their presence on the list.

Sidebar C: Drivers and Constraints

This simple method is a great tool for navigating the shoals of overoptimistic and overpessimistic technology forecasts. It's been likened to force-field analysis, but it emphasizes the dynamic instead of the static, which is the whole point of forecasting. Basically, it is a series of questions begetting more questions until you have an answer, or at least an informed opinion!

- What are the drivers for adoption?
  - How strong are they?
- What are the constraints on adoption?
  - How strong are they? Can they be overcome?
- What is the balance of drivers and constraints?
  - Will this change?
- What are the important areas of uncertainty that need to be resolved?
  - How can these be addressed to everyone's satisfaction?

Incidentally, besides preventing avoidable shipwrecks, the exercise of going through these questions provides a foundation for quantitative forecasts as well as a research agenda for both forecasting and the technology itself. For more information on drivers and constraints, see Vanston, 2008.

While we are on the topic of long-term forecasting, I will take the opportunity to plug my own field, technological forecasting (or technology forecasting)—which includes the kinds of "not so statistical forecasting" that we need when we require rigor and proven methods, but don't have much data. These methods include alternative scenarios, expert opinion, Delphi, substitution analysis, performance analysis, analogies, cross-impact analysis, and ideation tools like nominal groups and impact wheels, to name a few. *Technological Forecasting and Social Change* is the field's classic journal. Martino (1983) and Porter (2011) are a couple of the classic texts. Technological forecasting was essential in the 1960s as the digital revolution began, it was essential for the forecasting we did during the 1990s and early 2000s for the last wave of new technology (Vanston and Hodges,

2004)), and it will be essential for forecasting AI and other new-wave technologies in the future.

A few selected cautions regarding some of Spyros's specific conclusions:

- Don't assume that all computing will end up in the cloud. Computing has a long pendulum when it comes to centralization and decentralization. "Easy as using electricity" sounds a lot like the "computing utility" mantra from the late 1970s, right before the PC revolution decentralized computing. Currently, *edge computing* is as much a trend as cloud computing. With edge computing, information is stored and processed close to the user (thus, at the network's edge) to serve applications such as autonomous vehicles that require high reliability and low latency. Most likely there will continue to be local, edge, and cloud computing in some combination.
- I think Spyros sticks his neck out unnecessarily when he predicts that a wristwatch will be the interface of choice for communicating with computers. It might be, but AR/VR using glasses or headsets may reemerge as a contender as well. Or maybe a drone will fly beside us! Or maybe our current array of laptops, pads, smartphones, and big screens will survive. Choosing among alternative technologies that do the same thing is high on the list of the risky forecasting tasks (Vanston, 2008). Better to avoid it if possible!
- As I've noted above, brain-computer interfaces may have a place, but they are not necessary to survive AI. We should ask for more analyses of drivers/constraints and risks/benefits before agreeing with Spyros's optimistic conclusion regarding this technology. Otherwise, we risk having another unfortunate example of forecasts gone wrong. And very odd that he ends his opus on AI on this point!

AI FOR FORECASTING

Spyros makes one other point in Part 5 that I'd like to challenge vigorously. He cites the alleged slow progress in using AI for forecasting as an example of the weaknesses of AI. First, it's not a representative example. More importantly, it masks some fundamental changes in forecasting. The M competitions measure the ability to forecast the next few numbers in a long series of numbers with no contextual information except the frequency of the data. Years of research and prior contests have honed statistical methods so that they would be hard to be beat by any emerging technology. By not having labeled or contextual information, the strengths of AI forecasting are missed. The advantages statistical methods have in these contests are such that some AI forecasting experts have refused to participate.

Even then, the top finisher in the 2018 M4 competition, Slawek Smyl of Uber Technologies, used a hybrid of AI and conventional forecasting (Smyl, Ranganthan & Pasqua, 2018). This is more an example of teaching AI to use statistical forecasting so it doesn't have to learn it on its own than an example of happy machine/human collaboration. ISF2018 this summer was instructive. Four major presentations on AI for forecasting were authored by employees of Amazon, Microsoft, and Uber, leaders in the application of technology. All of these presentations were positive regarding AI for forecasting and optimistic that the constraints on AI (e.g., that it's data and computation intensive and that it's prone to overfitting and instability) will be overcome. Reading the article by the Amazon team, ironically in the same issue of *Foresight* as Part 5, reinforces this position (Januschowski and colleagues, 2018).

Will AI eventually replace most statistical forecasting? It's too soon to tell, but it's possible: parity has been reached in some applications, improvements in AI are likely, and constraints on AI can likely be overcome. Some of the objections to AI I've heard remind me of those raised by people engaged with prior old technologies such as traditional telephony, dial-up modems, and circuit switching. I'd mention slide rules, but that would

Will AI eventually replace most statistical forecasting? It's too soon to tell, but it's possible: parity has been reached in some applications, improvements in AI are likely, and constraints on AI can likely be overcome.

date me. And then there are the classic examples: steam locomotives and prop planes. Maybe forecasters need to sponsor their own long-range technology forecast, starting with a good drivers and constraints analysis! Even if the news is bad, I think the skill sets are highly transferable with a little training, and job security will be more like that of doctors than architects.

SUMMARY

Spyros has done a masterful job covering a complex, wide-ranging subject. I'm genuinely impressed. I have highlighted the areas of disagreement rather than the much broader areas of agreement because those are the ones that need to be talked about more. For the same reason, I have also addressed the more controversial concerns and dangers of AI more than its immense benefits. Forecasters of all stripes have a crucial role to play in helping sort these issues for the benefit of our clients and humanity. As my mentor the late Ralph Lenz once told me: "The highest value of a forecast is to raise the level of discussion." Let's talk.

REFERENCES

Brundage, M. (2018). Blog Posts, Retrieved from ***https://www.milesbrundage.com/blog-posts***

Constantin, S. (2017). Performance Trends in AI, Retrieved from Otium: ***https://srconstantin.wordpress.com/2017/01/28/performance-trends-in-ai/***

Electronic Frontier Foundation (2018). AI Progress Measurement, Retrieved from ***https://www.eff.org/ai/metrics***

Etzioni, O. (2016, September 20). No, the Experts Don't Think Superintelligent AI Is a Threat to Humanity, *MIT Technology Review*, Retrieved from ***https://www.technologyreview.com/s/602410/no-the-experts-dont-think-superintelligent-ai-is-a-threat-to-humanity/***

Januschowski, T., Gasthaus, J., Wang, Y., Rangapuram, S. S. & Callot, L. (Fall 2018). Deep Learning for Forecasting: Current Trends and Challenges, *Foresight*, 42-47.

Kurzweil, R. (2005). *The Singularity Is Near*, New York: Viking Books.

Martino, J. (1983). *Technological Forecasting for Decision Making, 2nd Ed.*, North-Holland.

Moravec, H. (1988). *Mind Children*, Harvard University Press.

Moravec, H. (1998). When Will Computer Hardware Match the Human Brain? *Journal of Evolution and Technology*, 1.

Porter, A. L., Cunningham, S.W., Banks, J., Roper, A.T., Mason, T.W. & Rossini, F.A. (2011). *Forecasting and Management of Technology*, 2nd Ed., Wiley.

Smyl, S., Ranganthan, J. & Pasqua, A. (2018, June 25). M4 Forecasting Competition: Introducing a New Hybrid ES-RNN Model, Retrieved from Uber Engineering: ***https://eng.uber.com/m4-forecasting-competition/***

Vanston, L. (2018). Forecasting AI Performance, 38th International Symposium on Forecasting, Boulder, CO, USA: International Institute of Forecasters.

Vanston, L. (2017). Forecasting Issues for the New World, 23rd IIF Workshop on Predictive Analytics and Forecasting (20-25), Munich: International Institute of Forecasting, Retrieved from ***http://predictive-analyticsandforecasting.com/wp-content/uploads/2017/10/MUNICH-WORKSHOP-VANSTON-Sepember-2017.pdf***

Vanston, L. (2008). Practical Tips for Forecasting New Technology Adoption, Telektronikk 3/4.2008, Retrieved from ***http://tfi.com/pubs/forecasting-tips.pdf***

Vanston, L. & Hodges, R. (2004). *Technology Forecasting for Telecommunications, Telektronikk*, 100(4), 32-42, Retrieved from ***http://tfi.com/pubs/w/pdf/telektronikk_peer.pdf***

Wikipedia (2018). Moravec's Paradox, Retrieved from Wikipedia: ***https://en.wikipedia.org/wiki/Moravec%27s_paradox***

Lawrence Vanston is President of Technology Futures, Inc. ***http://www.tfi.com***. See his Forecaster in the Field interview on page 48 of this issue.

lvanston@tfi.com

Response to Lawrence Vanston

SPYROS MAKRIDAKIS

I'm extremely flattered that two futurists as distinguished as Owen Davies and Lawrence Vanston have thought well to comment on my five-part *Foresight* paper "Forecasting the Impact of Artificial Intelligence." My sincere thanks to both for their insightful remarks and the opportunity they provided me to update my thinking about AI and check the rationality of my predictions. Their commentaries have considerably improved my original article, and I'm pleased that there is a fair amount of agreement among us. This response is devoted to Larry's commentary, as Owen's was a coda to my original paper with additional expert opinion gathered by TechCast's Delphi forecasts.

The Nobel laureate Paul Krugman in 1998 made some profoundly off-the-mark predictions that Owen Davies used as an epigraph to his coda in the Winter 2019 issue—including Krugman's statement that the Internet's effect on the world economy would be no greater than that of the fax machine—illuminating the shortsightedness of forecasters. Of course Krugman did not notice that a year earlier IBM's Deep Blue computer had beaten world chess champion Garry Kasparov, or that algorithms in 1998 could read handwritten characters in what were preludes to AI.

At the other extreme, forecasters should not fall into the trap of the science-fiction hype advocating that General Purpose AI (GAI) and singularity are just around the corner, competing with us and threatening our human dominance. I would like to reiterate my position that while AI is superb in games and image, speech, and text recognition, it is incapable of understanding meaning, making causal links, and exhibiting common sense. Many experts, to paraphrase Lawrence Vanston, say that "current AI has the IQ of a slug." My own prediction is that AI will take a long time to achieve the abilities of a one-year-old child to interact with its external world.

My main comments center on Larry's four major disagreements with what I have written, and present my reasoning as to why I believe I'm on the right track, although it will be many years before we know who's right and who's mistaken. Additionally, I will outline my view of the two biggest AI/IA issues facing humankind and explain why dealing with them may prove extremely difficult.

THE FOUR DISAGREEMENTS

Blockchain

Larry states that blockchain is basically a narrow technology trying to become mainstream. I disagree. I believe blockchain possesses three major advantages of great value that are not available in the traditional Internet:

- → it provides trust among unknown participants who can interact with each other with confidence;
- → it ensures enlarged safety in transactions;
- → it permits the construction of fast, trusted, and safe local networks.

These may not seem like much, but let's consider the implications of hacking houses that are connected to the Internet of Things (IoT), smart contracts, autonomous vehicles, and even the possibility of getting your thoughts stolen through a BCI (brain-to-computer interface). Unless perfect safety could be assured, the operation of IoT, smart contracts, and AVs, not to mention BCI, might be impossible. The same would go for local networks that could re-create the equivalent of the village square where neighbors can share communications and services. They must operate in a trusted and safe environment; otherwise, they would be shunned.

The Cloud

The second and third disagreements are about the cloud and its accessibility without a computer, say through a digital watch or other device. I understand

Larry's objections, but I still believe that we are at the very beginning of the utilization of cloud computing.

Two factors would contribute to moving the future toward my prediction. First, G5 and higher-level communication networks will bring Wi-Fi to all places, facilitating connections from anywhere to the cloud. Second, once voice commands become widespread and reliable, they would make keyboards obsolete: this would substantially increase the utilization of the cloud, whether this would be done by a cheap computer, a smartphone, or a digital watch. In my view, there will be no reason to own and maintain an expensive computer if the same service can be obtained more cheaply and easily from the cloud, allowing one to access and process information from any place at any time without having to carry a big device or worry about where such information is stored.

Forecast Accuracy

Will AI improve forecasting accuracy? My strong view is that improvements, if any, will be limited and will be due not to AI learning to forecast more accurately but rather to the ability of computers to optimize the parameters of the forecasting model more precisely. The top two winners of the M4 Competition (Smyl's model from Uber, and Montero-Manso from University of A Coruña, Spain) did not learn to forecast more accurately by studying past patterns or analyzing historical data. Rather, they simply found more effective ways to choose the most appropriate model and/or compute the best weights to combine the various forecasting methods used. AI forecasting models cannot work well when there are structural changes in the data or when patterns change. Our own experience has indicated that there is still a long way to go before AI forecasting models will be able to improve their accuracy through such learning.

OTHER MAJOR ISSUES

AI will give rise to two major, interrelated issues. The first will have to do with the "winner takes all" syndrome. AI applications, once developed, could be exploited at a global level at little additional expenditure, providing a huge advantage to their inventor(s).

The second will have to do with the income inequality created by such a syndrome as a single firm, or a very few firms, in some advanced countries would dominate the AI market. Furthermore, income inequality will become much worse between advanced AI nations and the rest of the world that would not have the specialized scientists or the vast research funds required to keep up.

Worse, as AI will be automating more and more repetitive labor tasks it will reduce the advantage of less-developed countries to attract firms through low labor costs, as robots and machines will be capable of operating the factories of advanced countries at similar or even lower costs. The big challenge will be to reduce income inequality within advanced AI countries through some form of guaranteed minimum income. However, it may be much harder to do so between countries. People seem unwilling to help poorer countries/regions economically while, at the same time, remaining unwilling to allow immigrants to enter their wealthier countries.

Spyros Makridakis is Professor, University of Nicosia and Director of the Future. See our Forecaster in the Field interview with Spyros in the Fall 2017 issue of *Foresight*.

makridakis.s@unic.ac.cy

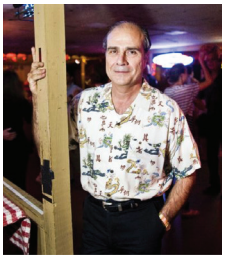

Interview with Lawrence Vanston, President, Technology Futures, Inc.

So, what do you want to do when you grow up?
I want to help people understand technology change so they can make good choices for themselves and the world.

Isn't that what you do already?
Yes, but I want to do it when I grow up, too, and better—especially the part about people understanding. I definitely have my Cassandra days.

What do people get wrong about the future?
They forget that things change; what's true today mightn't be true tomorrow. When you hear the words "never/not in my lifetime," alarms should blare. Surprisingly, experts are the worst. Then there's the opposite problem: people think change will occur overnight, are dismayed when it doesn't, then say "no one saw it coming" when it finally does. We keep making these mistakes, and simple tools for avoiding them have been around for decades.

You're a forecaster?
I call myself a "technology forecaster"…an actual thing! I deal with what'll change in the next five years plus, rather than short-term variations. Although there's plenty of overlap with short-term forecasting, I rely more on understanding the principles of technology change, less on statistical analysis. And I rarely have nice data to work with.

How did you get started in forecasting?
I was in grad school in the late 70s studying operations research. My father, John Vanston, taught a course in technology forecasting that I took and enjoyed. About that time, he started Technology Futures, Inc. (TFI). I graduated and took a dream job in network planning at Bell Labs. After the 1984 Bell System breakup, I moved back to Austin, joined TFI, and been there ever since. Best of all, I learned technological forecasting from the founders, who were more interested in right decisions than theory. Myself, I was interested in both.

Tell us more about your career to date.
Luckily, I've lived in interesting times—especially for the communications industry, home to most of my clients. The glory days were the late 80s and early 90s when we were simultaneously forecasting digital wireless, digital TV, consumer broadband, fiber optics, Internet applications, and local competition, which revolutionized the industry. Happily, most of my clients survived, I think in part due to good forecasting. The 2000s have been more evolutionary—2G-to-5G wireless; 1Mbs-to-1Gbs broadband; HD, 4K, 8K digital television—but still very interesting to a forecaster.

How do you see the future?
It's a new glory era for technology forecasting. The aforementioned technology wave changed everything over 30 years; this new wave will change things as much in the next three. AI, VR/AR, robotics, autonomous vehicles, the Internet of Things, all raise issues of timing, likelihood, path to the future, and potential impacts, the issues that technology forecasting addresses for good people to make good decisions. There will be lots to do when I grow up!

What do you do when you're not forecasting?
Besides spending time with family and friends, my passions are social dancing and art. You can spot me out dancing in Austin, the live-music capital of the world. Every November, my house becomes Art 84 where I welcome 1,200 or so visitors as part of the East Austin Studio Tour featuring art, music, and dance. I'm also on the board of Fisterra Projects, an Austin-based arts-and-sciences nonprofit. And I love to travel…a good thing, since my kids, clients, and conferences are far flung!