



Peter J. Denning

DOI:10.1145/2330667.2330677

The Profession of IT Don't Feel Bad If You Can't Predict the Future

Wise experts and powerful machines are no match for chaotic events and human declarations. Beware of their predictions and be humble in your own.

THE MACHINE THAT would predict the future." An article of that title appeared in the December 2011 issue of *Scientific American*. It suggested that advances in big data and super-computing will finally enable the old dream of an automated oracle. It set me to reflecting on what machines we already have available for forecasting and what our track record is with them. It also reminded me of a predicament I

have faced many times as a professional, on being asked to make forecasts. When can I offer forecasts that others can trust? When should I refrain?

The Work of Futurists

I began by inquiring into the work of the professionals who get paid for their forecasts.² Forecasting the future became a profession in the 1940s. Most professional futurists see their mission as investigating how social,

demographic, economic, and technological developments will shape the future. They advise on global trends, plausible scenarios, emerging market opportunities, and risk management. They are heavy users of information technology. Futurists rely on three main methods.

Revelation of current realities. Often we are oblivious or blind to what is going on around us. We operate with interpretations of the world that are unsupported by evidence. Futurists gather data and propose new interpretations grounded in that data. They then examine how policy and action might change to align with the reality. For many people, simply showing them what is already going on around them is a revelation of the future to them.

Peter Drucker was a master at this. His book *The New Realities* (Harper Business, 1989) is loaded with examples. In his chapter "When the Russian Empire Is Gone," he analyzed economic data, conversations of politicians and the media, and moods of Soviet citizens to conclude that the Soviet Union would soon fall. It did—within a year of the book's publication, even sooner than he expected.

Drucker was once asked what his method of forecasting was. He replied that he made no forecasts. He simply looked at the current realities and told people what the consequences are. When pressed to make long-term forecasts, he offered probability estimates based on past history.



Modeling. A model is a set of equations or simulations that take some observed variables (parameters) of a system and compute other values (metrics). A validated model is one whose track record shows consistently good agreement between computed and actual metrics. A validated model can be used for forecasting by declaring that its assumptions will still be valid in the future time period, and setting its parameters to the values expected in the future period. The forecast will be in error if model assumptions do not hold or if parameter estimates are incorrect. Such models have long been used in the sciences and engineering to describe natural recurrences.

Trend extrapolation is one of the simplest models. When a trend can be detected in some measure of performance, futurists can calculate future values and draw conclusions about the consequences. In 1965 Gordon Moore, a cofounder of Intel Corporation, noticed an 18-month doubling trend in the development of computer circuits (“Cramming More Components into Integrated Circuits,” in *Electronics Magazine* 38, April 1965). That is a 100-fold speedup for the same price over a decade. An industry rule of thumb is that any technology change that provides a 10-fold speedup can usher in a disruptive change. Many entrepreneurs started using the law to gauge whether their proposed disruptive technologies would be supportable by the computing power available in a few years. Moore’s Law became a guiding business model that has sustained the computer chip industry for nearly 50 years. It has started to break down as a trend because the sizes of transistors and wires are approaching a few atoms each, too small for them to function. Most trend analyses break down over longer forecast periods because eventually the trend encounters a limit.

In *The Age of Spiritual Machines* (Viking, 1999), Ray Kurzweil observed the same doubling trend in four previous generations of information technologies, and he claimed it would be present in technologies that supersede silicon. Based on that, he extrapolated Moore’s Law well into the future. He predicted a “singularity” around 2030, when he believes artificial brains will become intelligent.

Our problems with forecasts arise when we wrongly believe model assumptions or parameter forecasts will be valid.

On the other side, in *The Social Life of Information* (Harvard Business, 2000), John Seely Brown and Paul Duguid warned against overconfidence in trend extrapolation because social systems often resist and redirect changes in technology. They exposed a series of major predictions that never happened. Belief in those predictions led to the dot-com bust in 2002.

Scenarios. A scenario is a story that lays out in some detail what the future might look like under certain assumptions about trends and other factors. Futurists usually offer several scenarios under different assumptions. The method is useful to help people see how they might react to different futures, and then try to influence policies and trends so that the most attractive futures come to be. Futurists do not offer scenarios as forecasts or predictions. They sometimes give probabilities for the various futures they depict.

One thing I learned from this is that futurists actually avoid making predictions. They give you model results and scenarios and leave it to you to draw your conclusions.

Expert Predictions

Despite the caution of professional futurists, expert predictions have acquired a bad reputation. In *Future Babble* (Dutton, 2011), Dan Gardner argued that misguided trust among media forecasters in “legions of experts” has led many people down false paths. He bases his conclusions on the work of psychologist Philip Tetlock, who performed a long and careful study

of 27,450 predictions by 284 experts in many fields. Tetlock found that the performance of experts was no better than random guessing. He found that celebrity experts tend to be worse than random guessing and that “humble” experts—like the cautious futurists—tend to be slightly better than random. Consumers of these predictions tend to celebrate the successes and forget the failures.

Tetlock only evaluated predictions that were stated as definitively testable hypotheses; for example, “In five years, unemployment will be under 10%.” Many expert predictions are not so precise. Gardner says that experts are even less successful with vaguely worded long-term hypotheses than with precisely worded short-term hypotheses.

Dave Walter presented a dramatic example of failed long-term forecasts in his book *Today Then* (1992). At the 1892 Chicago Columbian Exhibition, the exhibitors speculated about how electricity, telephony, and automobiles would bring peace and prosperity in the coming century. The American Press Association invited 74 leading authors, journalists, industrialists, business leaders, engineers, social critics, lawyers, politicians, religious leaders, and other luminaries of the day to pen their forecasts of the world after 100 years.

The 1892 forecasters believed that in 1992 railroads and pneumatic tubes would be the primary means of transportation, governments would be smaller, and increased commerce would end wars. None foresaw the interstate highway system, genetic engineering, quantum physics, universal health care, mass state-sponsored education, broadcast TV and radio—or the computer. Walter concluded that many modern expert predictions are no more reliable than these.

Prediction Machines

Prediction machines are machines that forecast the future with reasonable accuracy. They are nothing mysterious. In almost every case they are validated models being applied for future conditions. How well have such machines done to date? Can they do better than experts?

Mathematical models of physi-

cal processes are the most successful examples.¹ Newtonian models of planetary motion give highly reliable predictions of the future positions of planets, asteroids, comets, and man-made vehicles. Jay Forrester's system dynamics models were very reliable for material and information flows in industrial plants. Queueing network models have been very reliable for forecasting throughputs and response times of communication networks and assembly lines. Finite element models have been very reliable for determining whether airplanes will fly or buildings will withstand earthquakes.

The common feature of these physical models is that they describe and exploit natural recurrences—laws of nature. We can assume that Newtonian physics, system feedback loops, congestion at bottleneck queues, and forces in rigid structures will continue to behave the same way in the future. We do not have to worry that the assumptions of the model will be invalid.

Our problems with forecasts arise when we wrongly believe model assumptions or parameter forecasts will be valid. In other words, we assume a recurrence that will not happen.

Many things can invalidate our assumptions of recurrences: human declarations in social systems, chaotic, or low-probably disruptive events, inherently complex systems whose rules of operation are unknown, complex adaptive systems whose rules change, environmental changes that invalidate key assumptions, and unanticipated interactions especially those never before seen. This list is hardly exhaustive.

Of these, I think the first is the most underappreciated. Human social systems are networks of commitments, and most commitments ultimately follow from human declarations. The timing and nature of declarations is unpredictable. Whether a technology is adopted or sustained in a community depends on the support of its social structure and belief systems, both of which resulted from previous declarations.³ Seely Brown and Duguid, mentioned earlier, give numerous examples of technology forecasts foiled by human declarations.

We know from experience that

We seek technology predictions in an attempt to reduce our risks, losses, and missed opportunities. We do so against great odds.

many validated models deteriorate over time. A locality principle is at work: the model assumptions are less likely to change over a short period than over a long period. Our short-range predictions are better than our long-range predictions. As a consequence, we need to frequently revalidate models to maintain our confidence that they still apply to at least the current circumstances.

What about long-term predictions? Most often, they are just flat-out wrong, as in the examples Dan Gardner and Dave Walter gave us. Occasionally they are correct but way off in the timing. Researchers at MIT predicted in the 1960s that computer utilities—forerunners of today's "cloud"—would be common by the 1980s; they were off by 30 years. Alan Kay predicted in the 1970s that personal computers would revolutionize computing; he was off by 20 years. Alan Turing in 1950 speculated that conversation machines would, by the year 2000, have a 70% chance of fooling a human for more than five minutes.⁴ He also thought that memory capacity for the machine's database would be the main obstacle. By 2012, our natural language systems are not close to this goal even though we have the memory capacity—but maybe in a few more years they will.

The few long-range predictions that do succeed late give us a forlorn hope that we can at least get the outcome right, even if the timing is off.

Nevertheless, the dream of good prediction by machine lives on. That *Scientific American* article mentioned

earlier envisions a project to build a computing system with more storage and computing power than ever before, connected globally to sensors and personal information. With new data mining methods to be developed, the system would find correlations in the data, and use them for predictions. Despite the soaring rhetoric, the system is no more likely to be successful than any other prediction machine, except when it can find and validate recurrences. It is unlikely to be successful whenever the outcome can depend on human declarations or unpredictable events.

Conclusion

We seek technology predictions in an attempt to reduce our risks, losses, and missed opportunities. We do so against great odds. Unpredictability arises not from insufficient information about a system's operating laws, from inadequate processing power, or from limited storage. It arises because the system's outcomes depend on unpredictable events and human declarations. Do not be fooled into thinking that wise experts or powerful machines can overcome such odds.

If you are called on to make forecasts, do so with great humility. Make sure your models are validated and that their assumed recurrences fit the world you are forecasting. Ground your speculations in observable data, or else label them as opinion. Be skeptical about your ability to make longer-term predictions, even with the best of models. Do not worry about the forecasts made by experts—they are no better than forecasts you can make.

Often, the most powerful and useful statement you can make when asked for a prediction is: "I don't know." **□**

References

1. Denning, P. Modeling reality. *American Scientist* 78 (Nov.–Dec. 1990); <http://denninginstitute.com/pjd/PUBS/AmSci-1990-6-modeling.pdf>.
2. Denning, P. Innovating the future: From ideas to adoption. *The Futurist*, World Future Society (Jan.–Feb. 2012), 40–45.
3. Schon, D. *Beyond the Stable State*. Norton, 1971.
4. Turing, A. M. Computing machinery and intelligence. *Mind* 59 (1950), 433–460; <http://www.loebner.net/Prizef/TuringArticle.html>.

Peter J. Denning (pjd@nps.edu) is Distinguished Professor of Computer Science and Director of the Cebrowski Institute for information innovation at the Naval Postgraduate School in Monterey, CA, is Editor of *ACM Ubiquity*, and is a past president of ACM.

Copyright held by author.