Philippe Hünenberger
ETH Zurich

Of millimeter paper and machine learning

My father used to be a gymnasium teacher. When he was not giving class, he loved to work in the cafés of our home-town, preparing his courses or correcting his copies. Sometimes, I was allowed to join him and then, I would ask him for "a function". In return, I would get a sheet of millimeter paper, his HP35 pocket calculator (with the fascinating "reverse Polish notation"), and … a mathematical expression returning a value of y for a given value of x. And off to work I was, calculating and painstakingly reporting dots on the paper to make a graph. This took hours, gratefully saved by my father for working quietly on his copies. And it also afforded me some surprises. For example, the function $y=\sqrt{(100-x^2)}$ only gave the upper half of a circle, not the lower one, and the HP calculator stubbornly refused to return a y for any x above 10 or below -10. Of course, my father would ultimately clarify these issues with me, but not before I had had time (literally hours!) to brood over them on my own ...

This was clearly drilling – deep and narrow in scope over long open-ended time stretches: a lonely, tedious and time-consuming task relying on manual exploration, intellectual processing, critical questioning and iterative experimenting. Besides a solid reputation of nerd-kid in the cafés, this approach rewarded me with a truly emotional respect for mathematical functions (they became life-long "friends") and an understanding that drilling can underpin thrilling subsequent discoveries (e.g., for the circle, complex numbers).

A few years later, I received my first computer, an Apple IIe. Now, given ten minutes or so of programming, I could plot any function on the screen, i.e. achieve almost instantaneously what had taken me a full morning a few years before. I was ready to explore a mysterious new world, that of mathematical functions: cardioids, epicycloids, astroids, Lissajous curves, Cornu spirals, … I could find new ideas in books, ask my father or other teachers, or simply try at random – I could even "play" the program as a game with friends.

I was now engaged in surfing – broad and shallow in scope over multiple short segments of time: the fast-paced and playful consideration of possibilities, relying on interactive exploration and comparatively superficial observations. This new approach rewarded me with a feeling for the richness of mathematical functions (they became a new "universe" to discover) and also with more visible achievements (even "marketable" ones, as I could earn a high-school prize for a program rotating the five Platonic solids on screen). But the mere surfing left me a bit frustrated, at least as long as I did not complement it by subsequent re-drilling on one function or the other. Just like a tour of 10 capitals of Europe in 10 days leaves you wishing you could spend afterwards 10 days in each of them separately.

When I started my PhD in Theoretical Chemistry in 1992, the research job matched my expectations: a back-and-forth oscillation (zoom in, zoom out) in Pascal's "double infinity", alternating surfing for breadth and drilling for depth, with a largely self-determined alternation schedule. Computers were useful tools, data was at the service of science, and the e-mail and internet were convenient devices for targeted and asynchronous communication. Since then, and especially over the last decade, things have changed a lot.

Undoubtedly, modern digital tools represent a fantastic extension of the human brain in terms of data access, processing throughput and communication reach. But in addition to that, they have also become overly invasive companions. Data seems to no longer be at the service of science, rather the opposite. The e-mail and internet, reinforced by an army of surveys, newsletters, mailing lists, evaluation tools and social networks, has evolved into an overflowing stream of information and an inexhaustible source of interrupts. In this noisy digital world, short-term surfing activities seem to take most of the space, while the quiet drilling activities have become a luxury. This may just be an exaggerated swing of the pendulum, triggered by the relative novelty of digital technologies. And the pendulum could certainly return to a more comfortable position provided that individual researchers and research managers both take the challenge seriously – and use their (human!) brains to control the present evolution. To this purpose, I have listed below three propositions.

Causality is stronger than correlation

Proposition one (epistemological): Causality-based models (from drilling) should be credited with a higher intrinsic value than correlation-based models (from surfing), irrespective of their relative current predictive powers for specific applications.

This judgment of value is not obvious to defend in times where big-data correlations relying on machine learning (ML) and artificial intelligence (AI) are becoming increasingly predictive, often more predictive nowadays than causal models based on elementary physical principles and numerical computations. Imagine a village where the "ancient" would predict the yield of the upcoming crop with 80% confidence based on rational thinking and a deep knowledge of the climate, plants and insects, but the "idiot" would do the same with 95% confidence by correlating intuitively a large number of more or less relevant observations in an entirely unknown fashion. Wouldn't it make sense to call the method of the "idiot" a "new paradigm" – and then maybe just kill the "ancient" to save food? Bad idea in my opinion … for at least three reasons.

First, comparing current predictive powers for specific applications is short sighted. As long as their physical Ansätze are correct, causal models have the potential of becoming fully predictive over all applications. The bottleneck is in their numerical evaluation, bounded by the current computing power. In contrast, correlation models are intrinsically limited in scope by the selection of a training set and in accuracy by the selection of input observables, irrespective of the available computing power.

Second, causal models have explicit Ansätze which are systematically improvable (e.g. Newtonian to quantum/relativistic ones), and their processing is amenable to human understanding and supervision. In contrast, the "Ansätze" of a correlation model are non-transparent, buried in the selection of training set and input observables. These may suffer from many biases, including proxy effects (if B is similar to A, then B must behave like A; punishment of the exception), assumed causality (if A correlates with B, then changing A will change B; action on a symptom), and design bias (voluntary or involuntary tuning of the training set and observables to get results matching prior expectations). In addition, the data throughput of correlation-based models is typically so high that they are beyond human supervision, and the coupling of the output of such a model to its input (e.g. funding of scientists made in proportion to their publication metrics) may create pernicious feedback loops. As a result, the uncritical use of large-scale correlation models tends to discourage serendipitous discoveries and promote self-fulfilling prophecies.

Third, only causal models represent what one should call knowledge in a humanistic perspective. In this sense, the terms ML and AI are misleading. Computers don't "learn" and they are not "intelligent". These are human characteristics, implying far more than correlation-picking (e.g. critical and orthogonal thinking, creativity, ethical accountability, emotional and social intelligence, …). Following Plato, I am convinced that knowledge is about finding what causes the shadows on the wall of the cavern, not merely about predicting patterns of motions in these shadows.

Surfing should be viewed as an extension to drilling

Proposition two (scaling-up): Surfing should be viewed as an extension to drilling, i.e. procedural understanding should precede automated application; this holds not only for scientific research, but also for learning, teaching, and education in general.

You don't give toddlers a Porsche to explore the city traffic. First, over many years, they learn how to crawl, then walk, then bike, and then drive (and then they can start saving for the Porsche!). Along the way, they progressively refine their procedural competences and feeling for danger, in parallel to scaling-up in terms of locomotion reach and speed. They also learn to distinguish between what does not need human thinking (and can thus be automated) and what definitely does (and therefore implies full brain awareness). Why should we do it differently with computers? In the context of teaching, this is what I wished to illustrate with my explorations of mathematical functions: first millimeter paper, then hacking a curve-drawing program, then surfing the function space.

You will often hear that computer-assisted techniques should be introduced as early as possible in teaching (this hype clearly extends up to the university level and beyond). The usual arguments sound like: (1) surfing is playful and interactive, thus likely to promote curiosity; (2) digital supports can be adjusted to individual learning curves; (3) this will prepare the child/student for a world where digital tools play a central role. None of the above arguments convinces me, because: (1) the type of "curiosity" induced by playful surfing is superficial and short-lived (nothing like the deep and long-lasting thirst one calls scientific curiosity); (2) creating an artificial world that adjusts miraculously to a person's needs actually impairs the development of adaptation skills (very unfortunate considering that neuroplasticity will be a key asset in the upcoming job market); (3) computer-surfing skills are relatively easy to learn if you have brain-drilling skills (but the opposite is definitely much harder).

A key pedagogical element in teaching is to trigger a curiosity-based itching for the next level of abstraction or throughput, thereby motivating the usefulness/necessity of this next level. Just as one should teach Chemistry starting from experimental observations and promoting an itching for the theoretical model explaining them, one should teach computer skills starting from step-by-step procedures and inducing a similar itching for the automation of the repetitive steps. Importantly, this scaling-up ensures that the assumptions and shortcomings of the modeling/automation procedure are evidenced explicitly, so that critical thinking is fully preserved in a subsequent faster-paced surfing phase.

Besides this scaling-up pedagogy, an open and respectful teaching atmosphere in the classroom, the promotion of critical and creative thinking, a thorough preparation, and an exemplary role of the teacher – which was already in essence the good old teaching recipe of my father – I am not sure there is so much to gain by introducing too many "innovations" in teaching, and especially not digital ones.

Similar considerations apply to research. Clearly, the modern scientific world is too complex and multi-faceted for anyone to know every technique in entire depth at any time. This is not even desirable. We all rely on a number of black boxes, i.e. procedural components (theories, models, algorithms, equations, software, data, …) for which we know the input and output, but not the inner workings down to the last details. The key question is rather about the extent of ignorance we are willing to tolerate when using a black box, without running the risk of being "fooled" by it. This limit is crossed as soon as we are no longer able to assess based on our own knowledge and thinking whether the black box is working correctly or not. An education that has involved an explicit scaling-up in the construction of a number of "standard" black boxes is definitely an asset for performing these types of assessments. The more we lazily skip to the surfing without spending effort onto the preliminary drilling (both in education and in research), the more our society will consider computers as wizards or oracles rather than tools.

Finding the balance between drilling and surfing is a major challenge nowadays

Proposition three (management of resources): Striking the appropriate balance and schedule between drilling and surfing in terms of allocated time, means and rewards is a major challenge nowadays; wise choices in this regard are of extreme importance for the long-term success of the scientific endeavor.

Individual researchers (in particular group leaders) could easily fill their agendas with surfing activities, leaving little room for drilling ones. Digital tools are not the direct cause for this, but an aggravating factor. This is because they allow a massive flow of information and requests to reach us on a quasi-instantaneous basis from all over the world, and because they represent a permanent invitation to inefficient reactive processing (ping-pong) and multitasking habits, themselves again contributing to increasing the digital flow. Yet, I am convinced that most researchers possess in principle the necessary skills and wisdom to strike the balance on their own, with "protection" tricks including: ignoring or declining most requests, delegating tasks, batching on-line periods, agreeing on communication policies, practicing temporary unreachability, and … being "sloppy" when something does not matter.

However, it is not clear how much they still have the freedom to do so in practice, considering the raise of two phenomena at the research-management level: the wish to increase the apparent productivity and immediate visibility of research, and the wish to reinforce its top-down steering. Both result in an increasing pressure on researchers to enhance what is considered to be their efficiency (the "ratio of research output to taxpayer franc", a nice expression I read recently in the NZZ) and quantifiable impact (university rankings, publication numbers and related metrics), and to work in directions that are imposed from the top based on immediate societal relevance and fashion trends (strategic goals, dedicated funding). Surfing activities tend to be more extravert, interactive, drivable, fast-paced and visible. Thus, they are more easily steered, quantified, recognized, financed and rewarded. Drilling activities, on the other hand, are typically introvert, slow, quiet and self-driven, and their effect on research quality is only visible in the long term. As a result, with a top-down management towards productivity and visibility, drilling becomes associated with a negative connotation of unproductive off time. This leads to an unhealthy tendency to minimize these activities or shift them into recovery time, as if they were no longer part of the job.

Fundamental research in an academic environment should be in first priority rigorous and creative, and only in second priority productive and visible. Historically, the production of the most efficient things (fundamental discoveries) has often been a rather inefficient process (trial-and-error, persistent work, well interpreted failures and … a bit of luck). In a society obsessed by efficiency, one should thus think carefully whether one wishes the research process to look efficient, or the research outcome to be efficient. If the latter is desired, the current management trends should be opposed, i.e. one should reinforce the trust in individual researchers.

The three above propositions are only invitations to your own thinking, a few personal suggestions for putting a new value on drilling in a world that is a bit too crazy about surfing. Maybe this thinking can help to avoid a possible future where data is the new currency and algorithms are the new priests, and in which technocrats drive the world based on curves from machine-learning, without ever having themselves put a single dot on a sheet of millimeter paper.

Prof. Dr. Philippe Hünenberger
ETH Zürich
Laboratory of Physical Chemistry
HCI G233
8093 Zurich
Switzerland
Phone +41 44 632 5503
phil@igc.phys.chem.ethz.ch
ORCID 0000-0002-9420-7998
http://www.csms.ethz.ch

Citation: Hünenberger P: Of millimeter paper and machine learning. Infozine **2018**, Special Issue 2, 6–8.
DOI 10.3929/ethz-b-000294364
Copyright: Infozine (Chemistry | Biology | Pharmacy Information Center at ETH Zurich)
Published: November 15, 2018