From Machine Learning to Human Learning: What Can Pedagogy Learn from AI Successes

Victor Timchenko[1], Yuriy Kondratenko[2,3]
Olga Kosheleva[4], Vladik Kreinovich[5],
[1]Admiral Makarov National University of Shipbuilding,
Mykolaiv, Ukraine, vl.timchenko58@gmail.com
[2]Petro Mohyla Black Sea National University, Mykolaiv, Ukraine
[3]Institute of Artificial Intelligence Problems
Kyiv, Ukraine, y_kondrat2002@yahoo.com
[4]Department of Teacher Education, University of Texas at El Paso
El Paso, TX 79968, USA, olgak@utep.edu
[5]Department of Computer Science, University of Texas at El Paso
El Paso, TX 79968, USA, vladik@utep.edu

Abstract

Many machine learning techniques – including many techniques behind the current AI-based boom in machine learning – come from the analysis of successful human learning strategies (and researchers expect that other human learning experiences can lead to even more effective AI-based systems). At this moment, so much experience have been accumulated in AI-based machine learning that it is time to start the analysis in the opposite direction – to see what can human-based pedagogy learn from AI successes. In this chapter, we provide the first results of such an analysis – some of which go somewhat against the current pedagogical wisdom.

1 Machine Learning and Human Learning – Past, Present, and Future: A Brief Introduction

Machine learning and human learning: recent past. Human learning is as old as humanity itself: we all learn from our parents, from our teachers, from our mentors. During thousands of years, humanity have accumulated a lot of experience in teaching and learning:

- we know what works and when,
- we know what does not work so well and why.

For many thousands of years, learning has been done by human teachers. When computers were invented, machines capable of what before considered intellectual functions, a natural idea was to help computers learn, i.e., to start the discipline of machine learning; see, e.g., [1].

How did researchers come up with ideas? One of the main ideas was to look for how we humans learn – just like if you want to create a flying machine, a natural ideas is to look at creatures that fly – e.g., birds. And indeed, many machine learning techniques came from how we learn:

- from neural networks that emulate our learning on the neuron level
- to techniques emulating learning on higher level.

Already on this stage:

- there were some successes in machine learning that overdid human doing similar learning tasks,
- but still, in comparison with thousands of years of successful human learning, the number of these machine learning successes was very small.

So still, the main progress in human learning was largely based only on human-learning experience.

Machine learning and human learning: present. With the development of deep learning, the situation is starting to change. Every day, newer and newer examples appear when machine learning systems provide better results that trained humans – be it:

- in many areas of medical diagnostics or
- in search for new molecules;

see, e.g., [2].

- While the number of successes of human learning continues to grow relatively slow – at approximately the same pace as in the past,
- the number of successes of machine learning is growing exponentially.

Already the number of machine learning successes has grown so fast that it is comparable to the number of human learning successes – and soon, it will overcome that number. This leads us to the next stage.

Machine learning and human learning – future: we need to use the experience of machine learning to come up with better human learning. To get better learning strategies, we learn from the past learning experiences:

- we learn the use techniques that turned out to be effective, and
- we learn not to use techniques that were not effective.

In the past, as we have mentioned, most such experience came from human learning. However, nowadays:

- as majority of learning success stories come from machine learning, and
- as the proportion of machine learning successes becomes closer and closer to 1,
- it is important to learn from successes and failures of machine learning.

This is, of course, already done when our main purpose is to develop better machine learning tools. However, it is also reasonable to use the machine learning experience to develop better human learning strategies. This brings up to the main purpose of this paper.

What we do in this paper. In this paper, we describe how the main features of successful machine learning can be used to help human learning. We present our conclusions in several sections of this paper, sections corresponding to different aspects of human learning.

It should be mentioned that somewhat unexpectedly, we come up with several conclusions that seems to contradict the usual pedagogical wisdom.

2 Who can learn?

What is the question? Before we start planning education strategies, it is important to decide whether each of the planned students is able to master this material.

What experience of human pedagogy teaches us. Experience of human pedagogy seems to indicate that some topics are not for everyone. For example:

- everyone can learn the high school physics, but
- not everyone can successfully pass PhD-level classes in theoretical physics.

Similarly:

- everyone can learn to play piano and to sing more or less in tune, but
- not everyone can successfully pass graduate-level music performance classes to become a skilled performer.

It is worth mentioning that even the idea that everyone can master high school subjects (such as basic math) is reasonably new. Even in the 19 century, many people believed that women and people of lower social status are unable to master these subjects – and many US folks still believe that, e.g., to master Calculus, one needs to have a specially organized brain. Such beliefs are not common in Europe, where dozens of years of required high school education has shown that indeed, people are capable of learning all this.

However, most people – and most researchers in pedagogy – do believe that a special brain is needed to become a theoretical physicists or a performing musician.

Comment. Not both, by the way:

- while Albert Einstein was clearly a genius in physics,
- his ability to play violin was always on an amateur level.

Several memoirs of young physicists of his time mention that before their introduction to Einstein, they were warned that:

- while Einstein loves constructive criticisms of his physical ideas,
- he is really hurt when someone criticizes his violin skills :-(

But what can we learn about it from the current AI success? A human being is born with a largely pre-organized network of neurons, a network that allows a newborn baby to cry, to feed, even to swim if placed in water.

- Some brains have more neurons and/or are better organized by the birth time.
- Some brains have fewer neurons and/or are not as well organized.

However, all these brains have the same basic structure – that later becomes more and more developed as the baby learns.

In contrast, the current AI successes:

- do not come from the specially organized neural network with pre-determined weights;
- they come from a neural network in which initial weights are assigned randomly.

From this viewpoint, the starting state of AI learning is not even the state of the newborn baby. We could call it the state of the amoeba except that the starting neural network has much more neurons that an amoeba.

And still, starting from this very primitive state, the AI systems manage to learn many skills that, for humans, require an equivalent to a PhD – such as the ability to reliably diagnose rare diseases.

What this shows is that to learn such complex topics, there is no need to have a specially structured brain – so *anyone can learn* all the courses needed to get a PhD in mathematics or in theoretical physics and become a professional in these disciplines.

Of course, like with every other topic:

- some people will learn this material faster, while
- others will require more time to learn this material,

but eventually, everyone will learn.

3 In what order should we teach?

How this is usually done. When we want students to learn a lot of material – that does not fit into a single class – then we divide this material into classes:

- first, we teach the basics,
- then, we teach more complex stuff, and
- finally, we teach the most complex technical details.

What can be learn about it from the AI successes? The current AI successes are obtained by using *deep learning*, where all the layers of the neural network learn at the same time.

- Researchers tried to train one layer at a time.
- However, training all the layers at the same time turned out to be much more effective.

In a deep neural network, roughly speaking:

- the first layer corresponds to facts,
- the next layer corresponds to more general notions,
- and, in general, each following layer corresponds to the next level of abstraction.

From this viewpoint, learning all the layers at the same time means that we should learn the whole materials in each course. For example:

- instead of studying the basics of programming in Computer Science I course,
- we should study all the levels of programming in this course, and
- then refine our knowledge of all the levels in each following course.

There is another AI-based reason why this learning scheme would work the best. Namely, in AI, it is often very effective to do what is called *transfer learning*: instead of starting from scratch, we start with a network training for one task, and re-train it for another task. This way, instead of first studying one topic from scratch and then another topic from scratch, we use whatever we learned in the first topic to start learning the second topic as well.

This is somewhat similar to how students are taught in high school, but for most higher education courses, this is a radical idea. It has a potential advantage. For example, now engineering students first study calculus – often without a clear idea of why this is needed in engineering. Because of this lack of understanding, they often do not take this course seriously enough – and thus, at the next semester, when they take an engineering course in which calculus is used, they need to re-learn calculus. In the above-described AI-motivated arrangement, the students study both calculus and its applications at about the same time. Thus, they better understand the need for calculus and hence, waste less time on re-learning it.

4 Related question: do we need rigorous foundations?

How this is done now. As we have mentioned in the previous section, many science and engineering courses – especially physics courses – are based on mathematical foundations: mostly on calculus and differential equations. These foundational courses are usually taught by mathematics faculty, who try their best to present this material is a rigorous way – the way of mathematics.

Many engineers and physicists believe that for their students, most of this rigor is a waste of time, all the students need is to remember the techniques, and there is no need for them to learn the foundations of these techniques. There have been less-rigorous textbooks written by such engineers and scientists – see, e.g., [9] – but these textbooks are an exception,

What can we learn about this from AI successes? It is known that modern AI systems are not good in reasoning and in rigor – but nevertheless, they produce much better results that systems based on rigorous derivations.

From this viewpoint, to make teaching of physicists and engineers more effective, it makes sense to decrease the amount of rigor in foundational classes to the bare minimum.

5 How should we teach: special learning techniques can be helpful, but is it possible to learn without them?

What modern pedagogy teaches us. Many pedagogical publications start with the fact that there are some problems with students learning required material. Then these papers:

- describe new teaching techniques – be it active learning, using groupwork, etc. – and
- show that the use of these techniques leads to a statistically significant improvement of the learning results.

This creates an impression that learning success is not possible if we only use traditional learning tools. This impression is rarely stated explicitly, but implicitly it is there: no one proposes a new method that would simply allow, e.g., 8th grade students to learn the material twice faster.

But what can we learn about it from AI successes? Interestingly, AI successes *do not* come from any sophisticated learning techniques: they are based on the simplest possible optimization technique – gradient descent. This technique used to be described in courses on Numerical Mathematics as an example of a technique that is not recommended due to its simplicity – but AI successes brought this technique's revival. Moreover, attempts to use more complex learning techniques led to worse results – that is why gradient descent is still actively used.

What this seems to indicate, in regard to human education, is that, in principle, we can learn even the most complex subjects by using only the traditional teaching techniques. This does not mean that more complex techniques are not helpful: they can make learning faster and more reliable – but in principle, contrary to the impression that one gets from reading the modern pedagogical papers, we can learn everything by using only these techniques.

6 How do we teach: is there a need to drill on many examples?

How this was done in the past and what pedagogy teaches us. In the past, students would have to do a lot of similar exercises on each topic. Such a drill approach is still practiced, but many pedagogical papers prefer more creative solutions.

What can we learn about this from AI successes? AI successes are largely based on training on thousands of similar examples – exactly what drill is about.

From this viewpoint, it looks like drill should remain an important part of teaching. It is definitely a good idea to make this drill less boring by adding some creativity, but it looks like drill should remain.

7 Need to balance positive and negative examples

How this is usually done. In traditional learning, students are mostly trained on correct solutions – this is what is typically explained in class, sometimes with a few examples of typical mistakes.

What can we learn about this from AI successes? The experience of AI is that learning works the best where we have approximately the same number of positive and negative examples.

With respect to human learning, this means that we need to provide students not only with many correct examples, we also need to provide them with as many incorrect examples, examples on which students will need to learn how to avoid the corresponding mistakes.

This seems to be a completely new idea in pedagogy. It is desirable to try it. Our hope is that it will work, since this is related to the known problem of partial credit. Let us illustrate this problem on the following simplified example. Suppose that we give a test with 10 problems, each worth 10 points. To solve each problem, the student needs to correctly perform 10 steps – and usually, instructors give 1 point for correctly performing each step. If for each of ten problems, the student performs correctly 9 out of 10 steps, this student gets 9 points for each of these problem. So, this student's overall grade is $10 \cdot 9 = 90$, which is still A ("excellent"). So:

- by accommodating partial credit, the student gets an excellent grade on the 10-problem test, while
- in all 10 problems, the student's answers are wrong.

Hopefully, students who are more exposed to possible mistakes will be more careful and avoid such a situation.

8 Possibility of visualization

How this is done now. Of course, everyone understands that visualization is helpful. In some topics, there are good visualization techniques, while in many others, there are none. It is not even clear whether a good visualization is possible for some abstract topics.

What can we learn about this from AI successes? One of the main reasons why some things are difficult to visualize is that:

- in many problems, we perform computations with high accuracy, but
- in visualization, we can only achieve a certain level of accuracy – e.g., we can only distinguish a small number of intensity levels, as small number of colors, etc.

At first glance, this may seem like a serious obstacles to visualization. However, in AI-based machine learning, there are empirical results – that are justified by a rigorous mathematical analysis – showing that for training, it is enough to use 8-bit numbers (see, e.g., [4]), i.e., numbers that have only $2^8 = 256$ distinguishable values.

This number is exactly how many colors we can distinguish – which shows that we can always use colors to describe all intermediate stages of the desired computations. In other words, visualization is always possible.

This can be also used to speed up computations. It is worth mentioning that not only visualization is always possible, but we can also use the color interpretation of numbers, and process color signals instead of computer-based real numbers (see, e.g., [3, 5, 7, 8]):

- this is faster – since colored light travels with the speed of light – the largest physically possible speed, and
- several light beams can be easily processed in parallel, thus further speeding up the computation process.

9 How to assess learning and how to motivate students?

How learning is assessed now, and what modern pedagogy teaches us about it. Once the teaching is over, we need to assess how well the students learned. Traditionally, the student's level of knowledge is accessed by comparing the student's answers with the correct answers:

- if all the student's answers are correct, the student gets a higher grade, while
- otherwise, if some answers are wrong, the student's grade is lower.

Many modern pedagogical publications propose to change the traditional ways of assessment. Some researchers propose not to grade students' work at all – at least for some material. Other researchers propose to replace the overall grade for a test and/or for a class with a list of grades describing the student's level of knowledge in different categories.

Some publications show that these assessment changes lead to more effective learning. On the other hand, the experience of some of our colleagues – who have tried to implement these ideas – has shown that students are often confused by these changes.

What can we learn about this from AI successes? For training, all AI systems use simple objective functions that are based on similarity between the desired answers and the answers provided by the system on the current stage of training. Attempts to introduce more complex assessment techniques did not lead to more effective AI training.

From this viewpoint, it looks like the traditional assessment techniques are good enough for training. The only change may be that:

- while originally, neural networks used the least squares optimization function – that is known to lead to the final grade being a linear combination of individual grades (see, e.g., [6]),
- modern AI systems use non-quadratic objective functions, that result in non-linear combination of individual grades.

From this viewpoint, it may be beneficial to explore nonlinear techniques for combining grades for individual assignments into a single grade.

Acknowledgments

This work was supported in part by the National Science Foundation grants 1623190 (A Model of Change for Preparing a New Generation for Professional Practice in Computer Science), HRD-1834620 and HRD-2034030 (CAHSI Includes), EAR-2225395 (Center for Collective Impact in Earthquake Science C-CIES), and by the AT&T Fellowship in Information Technology.

It was also supported by a grant from the Hungarian National Research, Development and Innovation Office (NRDI), and by the Institute for Risk and Reliability, Leibniz Universitaet Hannover, Germany.

References

[1] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, New York, 2006.

[2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, Cambridge, Massachusetts, 2016.

[3] O. Kosheleva, V. Kreinovich, V. L. Timchenko, and Y. P. Kondratenko, "From Fuzzy to Mobile Fuzzy", *Journal of Mobile Multimedia*, 2024, Vol. 20, No. 3, pp. 651–664.

[4] C. Q. Lauter and A. Volkova, "A framework for semi-automatic precision and accuracy analysis for fast and rigorous deep learning". In: M. Cornea, W. Liu, and A. Tisserand, (eds.), *Proceedings of the 27th IEEE Symposium on Computer Arithmetic ARITH 2020*, Portland, Oregon, June 7–10, 2020, pp. 103–110.

[5] A. Lupo, O. Kosheleva, V. Kreinovich, V. Timchenko, and Y. Kondratenko, "There is still plenty of room at the bottom: Feynman's vision of quantum computing 65 years later", In: Y. P. Kondratenko and A. I. Shevchenko (eds.), *Research Tendencies and Prospect Domains for AI Development and Implementation*, River Publishers, Denmark, 2024, pp. 77–86.

[6] D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*, Chapman and Hall/CRC, Boca Raton, Florida, 2011.

[7] V. Timchenko, Y. Kondratenko, and V. Kreinovich, "Logical Platforms for Mobile Application in Decision Support Systems Based on Color Information Processing", *Journal of Mobile Multimedia*, 2024, Vol. 20, No. 3, pp. 679–698, doi.org/10.13052/jmm1550-4646.2037

[8] V. Timchenko, V. Kreinovich, Y. Kondratenko, and V. Horbov, "Effectiveness Evaluations of Optical Color Fuzzy Computing", In: Y. P. Kondratenko and A. I. Shevchenko (eds.), *Research Tendencies and Prospect Domains for AI Development and Implementation*, River Publishers, Denmark, 2024, pp. 129–151.

[9] Ya. B. Zel'dovich and I. M. Yaglom, *Higher Math for Beginners*, Prentice Hall, Hoboken, New Jersey, 1988.