Memories of the Future: Systems, Human, and Cybernetic Aspects of the Emerging Post-AI World

Vladik Kreinovich[1], Miroslav Svitek[2],
Julio Urenda[3], and Olga Kosheleva[4]

[1,3,4]Departments of [1]Computer Science, [3]Matematical Sciences, and
[4]Teacher Education, University of Texas at El Paso,
500 W. University, El Paso, Texas 79968, USA
vladik@utep.edu, jcurenda@utep.edu, olgak@utep.edu
[2]Faculty of Transportation Sciences, Czech Technical University in Prague,
Konviktska 20, 110 00 Praha 1, Czech Republic, miroslav.svitek@cvut.cz

1. We are in the age of AI

- Spectacular successes of machine learning techniques have made everyone in awe of AI:
- Starting with ChatGPT, Large Language Models (LLMs) help us in many creative tasks: writing poems, coming up with syllabi, etc.
- Computers are world champions in Go – probably the most difficult human-invented game.
- Self-driving cars are a common sight in many cities.
- AI-based systems solve important partial differential equations faster than all known algorithms, etc.
- To many people:
  - AI is all we need,
  - and if the current AI is not yet perfect, all we need is to train it some more.

2. We are in the age of AI (cont-d)

- Computer Science departments and programs all over the world are renamed into AI departments and programs.
- Most faculty openings in Computer Science are now AI-related.
- All other research directions are almost frozen – to the extent that:
  - several member of the Program Committee of one of the major fuzzy conferences
  - seriously proposed to give the conference Best Paper Award to a machine learning paper
  - that has practically nothing to do with fuzzy.

3. But

- The recent success of AI was very fast and very unexpected.
- In terms of successes, it is a clear example of a steep exponential growth.
- It is natural to expect that this spectacular exponential growth will continue for some time.
- Of course, it is not possible that this exponential growth will last forever.
- There are many real-life examples where such a growth stopped.
- An economic boom leads to overheated economy and even to a crisis.
- Exponential population growth that worried people in the 1960s has drastically slowed down, etc.
- But with AI, many people expect the boom to continue for some time.
- However, our opinion is that the boom is about – if not to end, but at least to slow down.

4. But (cont-d)

- Yes, LLMs are well known to be not perfect.
- However, the hope was that:
  - just like training on more examples made machine learning results better,
  - feeding even more examples to LLMs will make them less imperfect.
- What we see now is that the current LLMs have already devoured practically everything that is there.
- Even several authors of copyrighted texts are suing the LLM-producing companies for a supposedly illegal use of their work.
- The training took practically all possible available examples, but things are still not perfect.
- Self-driving cars are still, in many situations, several times less safe than an average human driver.

5. But (cont-d)

- Even the Go championship is now in doubt:
  - while the AI indeed did easily beat Go professionals,
  - it has been shown to be defenseless against a skilled amateur.
- Maybe we can get better results if:
  - instead of the current – mostly empirically determined – aspects of neural network
  - such as activation function etc.
  - we find better ones?
- Unfortunately, there is not much hope that this will lead to significantly better results.
- Most such empirically determined aspects turned out to be provably optimal in some reasonable sense.
- In short, what we observe is a drastic slowdown of AI progress.

6. Comment

- We are *not* saying that AI progress has stopped.
- There are still many potential useful applications of modern AI, of current LLMs.
- These applications need to be developed and exploited.
- What we *are* saying is that the techniques themselves are not growing as fast as we expected.
- To continue growing, they need a boost.

7. What we do in this talk

- In this talk, we analyze where the desired boost can come from.
- Our conclusion is that it will come if we better take into account systems, human, and cybernetic aspects of the problems.
- These are largely the aspects that are largely within the scope of our IEEE Systems, Man, and Cybernetics society.
- We titled this talk "Memories of the Future" because:
  - while we always appreciate new ideas,
  - we believe that (and we argue that) many problems can be solved if we go back to the previous AI (and other) ideas.
- These ideas are currently mostly frozen.
- We need combine them with the current machine-learning-based AI techniques.

8. So what are systems, human, and cybernetic aspects?

- Let us start by describing what we mean by systems, human, and cybernetic aspects.
- *Systems* means that:
  - we consider processes as a whole,
  - including software, hardware, humans, and other objects involved in the corresponding process.
- For example, a systems approach to manufacturing means to consider all aspects:
  - from providing raw materials
  - to human satisfaction of the final products
  - to the effects on environment and health.
- *Human* means taking into account that the ultimate goal of all the human activity is to benefit the humankind.

9. What are systems, human, and cybernetic aspects (cont-d)

- So we need to take into account:
  - not only objective characteristics,
  - it is also important to take into account what humans want, and how they perceive what we offer them.
- A classical example of the difference between objective characteristics and human perception is economy.
- Objectively, we should aim at the largest and most fairly distributed Gross Domestic Product (GDP).
- However, the ultimate goal is human happiness.
- It is well-known that:
  - while, in general, GDP is correlated with happiness,
  - there are many exceptions.

10. What are systems, human, and cybernetic aspects (cont-d)

- In Mexico, people are, on average, happier that what a correlation model predicts.
- In several countries of Northern Europe, they are less happy than they are expected to be based on their GDP.
- Finally, by definition, *cybernetics* is a study of processes that are common to human, living beings, and machines.
- Studying such processes has led to many useful developments, and this usefulness makes sense.
- For example, if we want to make a machine that flies, it is useful to look at creatures that fly – birds and insects.
- As a result, we have modern airplanes with wings.
- Of course, these wings are different from the birds' ones – for example, they do not flap.
- But still the main idea is largely the same.

11. What are systems, human, and cybernetic aspects (cont-d)

- Similarly:
  - if we want to make a machine that thinks and makes decisions,
  - a natural idea to look at how we humans think and make decisions.
- This can be on the level of statements and logics, as in the historical logical approach to AI.
- This can also be on the level of neurons.
- This is, by the way, where the modern neural network-based spectacular AI models come from.

12. What do we want?

- For all our problems – be it transportation, medicine, manufacturing – we want a solution:
  - that would be satisfactory to us,
  - that would provide accurate and reliable results while maintaining privacy and security.
- We also want these solutions to be generated as fast as needed.
- This, in many cases, means as fast as possible.
- And ideally, the corresponding computations should not spend as much energy as they do now.
- Now computers take more than 5% of the world's energy consumption (and this proportion is growing).
- After all, our brains perform their computations and their reasoning at the energy level several orders of magnitude lower than the computers.

13. How do we make sure that the solution is satisfactory?

- First, we want to find out what we humans want, be it transportation, be it eating out, etc.
- This is not easy to formulate in precise computer-understandable terms.
- For example, for transportation:
  - in the first approximation, it makes sense to assume
  - that we want to get from point A to point B as fast as possible.
- But this is true only in the first approximation.
- We also need to take comfort into consideration.
- It is largely because of the comfort that:
  - in many big cities, some people prefer to drive their own cars,
  - while public transportation is often much faster.

14. How do we make sure that the solution is satisfactory (cont-d)

- How can we learn what people want?
- How can we take into account this *human* aspect of the corresponding problem?

15. Why cannot we just use current machine-learning-based AI techniques?

- At first glance, why not use machine learning (ML)?
- In many other learning tasks, they have been very successful.
- But these ML techniques are only very successful when they are fed thousands and millions of examples.
- The need for so many examples can be naturally explained by simple statistics.
- Indeed, based on a size-N sample, we can estimate the probability of a certain event with approximation accuracy $\sim 1/\sqrt{N}$.
- For example, based on 100 samples, we can estimate the desired probability with approximation accuracy $\sim 10\%$.
- So, to estimate the probability with approximation accuracy 0.1%, we need to have $N = 10^6$ – i.e., we need a million of examples.

16. Why cannot we just use current machine-learning-based AI techniques (cont-d)

- This, by the way, is the reason why self-driving cars are not yet perfect:
  - even with billions of examples $N \approx 10^9$,
  - the probability of an accident can only be estimated with approximation accuracy $1/\sqrt{N} \approx 3 \cdot 10^{-5}$.
- This means that we cannot guaranteed smaller probability of an accident.
- And this probability is still much worse than the probability for an accident for a human-driven car.
- It is easy to come up with millions and billions examples of values of the corresponding physical process.

17. Why cannot we just use current machine-learning-based AI techniques (cont-d)

- Indeed, sensors are very affordable now, we can make many measurements.
- However, it is difficult to come up with millions of survey results describing how people feel.

18. What can we do with this human-related aspect?

- Good news is that in coming up with ideas, we do not need to start from scratch.
- People have been thinking of how to describe human reasoning for quite some time.
- This was one of the main direction in traditional AI, with neural networks being a poor cousin at that time.
- It may be time to revive that research and combine machine learning with more traditional reasoning-based AI techniques.
- Also, people have been thinking of how to formalize imprecise ("fuzzy") human sentiments for quite some time.
- One of the successful techniques:
  - that led to many successful applications in the 1980s and continues to be useful
  - is *fuzzy techniques.*

19. What can we do with this human-related aspect (cont-d)

- There, we describe expert statements like "the tumor is small" by asking the expert to assign:
  - to each possible tumor size,
  - a degree to which this expert believes that such a tumor is small.
- So a reasonable idea is to also combine machine learning techniques with fuzzy.
- Such attempts are already going on – judging by many related presentations on fuzzy conferences and publications in fuzzy journals.
- We just need to take this activity more seriously:
  - not, as some people in ML community perceive it – as attempts of somewhat obscure fuzzy community to compete,
  - but as an attempt to further boost AI techniques.

20. Comments

- Why should we combine ML with anything else?
- Many researchers like the saying "if all you have is a hammer, everything starts looking like a nail".
- Hammer is just one of the tools, a very important one, no doubt.
- Similarly, ML is just one of the tools, maybe one of the most important ones.
- However, we often also need other tools to succeed.
- There is an additional advantage of combining machine learning with more traditional AI techniques, in particular, fuzzy techniques:
  - in contrast to usual machine learning – which, to a human user, is just a black box,
  - traditional AI and fuzzy techniques use natural language.
- Human users can understand how these techniques come to their conclusions.

21. Comments (cont-d)

- This understanding makes these technique more reliable to human users.
- For those who are not very familiar with modern fuzzy techniques, it is necessary to emphasize that:
  - they have gone way beyond
  - the original Zadeh's simple ideas of using values from the interval $[0, 1]$ and min and max instead of "and" and "'or".
- Modern fuzzy techniques are much more complex now.
- They do not have as many spectacular applications as in the peak of their usage, in the 1980s and 1990s.
- However, they are still successfully used.
- They are used in many cars' automatic transmissions.

22. Comments (cont-d)

- There, one of the main objectives is an imprecise objective to make the ride comfortable for the driver and for the passengers.
- They are used:
  - in many other similar human-related applications – like a rice cooker,
  - where the quality is largely subjective.

23. How do we make the solutions more accurate and reliable – or at least how to gauge their accuracy and reliability?

- Many machine learning results are way off.
- Such results are usually called *hallucinations*.
- For example, a very effective image understanding algorithm may suddenly interpret:
  - a cat picture (which, to a human, is clearly a cat picture)
  - as a dog or even as a train.
- Of course, there are cases when an algorithm makes a mistake.
- No one is perfect, and most algorithms are not perfect either.
- What many algorithms – and many humans – do in such situation is provide:
  - not only the answer itself, b
  - ut also an estimate for the accuracy and/or reliability of this answer.

24. How do we gauge the solution's accuracy and reliability (cont-d)

- For example:
  - when my doctor says that I have a hernia that needs to be operated upon,
  - all I can do is agree to this operation.
- However, suppose that my doctor says that most probably it is hernia, but an additional ultrasound test may be helpful.
- Then I would probably first undergo this test.
- From this viewpoint, it is desirable to supplement machine learning results with such estimates.

25. How do we gauge the solution's accuracy and reliability (cont-d)

- If AI's recommendations would come with estimates of their reliability,
  - and users will realize that recommendations with high reliability estimates are indeed largely reliable,
  - this will increase human trust in AI.
- In other words, we need to apply methods of Uncertainty Quantification (UQ) to such situations.
- And this is where further work is clearly needed.
- Many machine learning algorithms already produce estimates of their answer's reliability and accuracy.

26. How do we gauge the solution's accuracy and reliability (cont-d)

- However, these estimates are often way off:
  - in examples when a system recognizes a simply wavy pattern as a penguin,
  - its estimated confidence in this identification is sometimes 99%.
- What can we do about it?

27. Why cannot we just use current machine-learning-based AI techniques?

- Since ML is so good in learning, why cannot we use it to provide the desired estimates?
- At first glance, we can rather easily do it.
- We can come up with groups of images – grouped by different degree of clarity.
- For each of these groups, we can apply the trained ML tool (whose accuracy we are estimating) to each of these images.
- Based on the results, we can compute the percentage of correct answers.
- Then, we train another ML tool on the pairs (image, percentage-of-the-correct answer in the corresponding group).

28. Why cannot we just use current machine-learning-based AI techniques (cont-d)

- After this training, we expect that:
  - for every new image,
  - thus trained tool will provide a good estimate for the first tool's reliability,
  - so, we can use similar techniques to gauge the accuracy of neural network's computations.
- This seems like a natural idea, so why doesn't it work?
- When we train the neural network on pairs like (image, cat), (image, train), etc., it learns.
- However, when try to train it on pairs (image, percentage), it does not work – why?
- The answer goes back to the same statistical arguments that we used earlier.

29. Why cannot we just use current machine-learning-based AI techniques (cont-d)

- To be well trained, we need many examples.
- If we have a sufficiently large number N of examples, the system will be well trained.
- However, to get an example of a probability, we need a group of such examples.
- Suppose that we want to estimate this probability with accuracy 10%.
- Then, we need – by the same statistical argument – to have at least 100 examples in each group.
- Since each group has 100 examples, we have $N/100$ groups.
- This means that the resulting accuracy of estimating the degree of confidence is 10 times smaller than the accuracy of classification itself.

30. So what can we do?

- One of the main reasons why detection is not perfect is that we do not have a pure picture of a standard average cat.
- In this ideal case, we would have a perfect (or at least an almost perfect) recognition.
- In reality:
  - we have a picture of an individual cat that is different from the average cat, and
  - we have other objects in the picture.
- From this viewpoint, both factors act as noise:
  - the difference between an individual cat and an average cat
  - and the background.

So, the question becomes: how to gauge the effect of this noise on the system's decision?

31. So what can we do (cont-d)

- Out of three aspects that we mentioned earlier, here the most appropriate is the systems aspects.
- Indeed, after all, the trained neural network can be viewed as a system.
- For systems, noise is a usual feature, we know how to gauge the effect of noise.
- It is therefore desirable to use the general system techniques to gauge the effect of this specific noise on the result.
- Traditional methods of this estimation are methods of sensitivity analysis.
- They require that we estimate the effect of each noise component.
- In effect, they require that we estimate the partial derivative of the computation result with respect to each input.

32. So what can we do (cont-d)

- In general, this requires calling the algorithm as many times as there are inputs.
- For ML models with many inputs, this would be prohibitively long.
- Good news is that since training of a neural network is based on computing partial derivatives.
- So, we can use this built-in feature to compute the derivatives:
  - practically for free,
  - at the expense of just one additional back-propagation step.
- Hopefully, other systems-related techniques can also be designed.

33. An alternative approach: eXplainable AI (XAI)

- An alternative approach to making AI results more reliable is to make them explainable to human users.
- We mentioned this idea earlier.
- There is a lot of work in XAI, including Shapley-related approach that we discuss later.
- One of the main problems here is to understand what users perceive as an explanation.
- his is very subjective, it is clearly related to the human aspects.

34. How can we manage security and privacy?

- Security and privacy are very important.
- However, every security and privacy professional will tell you that there is no such thing as perfect security and perfect privacy.
- What we really need is to make sure that the probabilities of security and privacy violations should be small.
- For this purpose, we need to be able to estimate these probabilities.
- Similarly to the uncertainty quantification case, the traditional machine learning tools may not be very helpful here.
- Indeed, to estimate a probability, we need a group.
- This drastically decreases the number of training examples and thus, drastically decreases the estimation accuracy.
- So, for these estimations, we need to supplement machine learning with more traditional security and privacy techniques.

35. How can we manage security and privacy (cont-d)

- One way to preserve privacy – that is often used in surveys – is to replace exact numbers with ranges.
- For example, a survey:
  - instead of asking for the exact age,
  - may ask whether the age is between 20 or 30, or between 30 and 40 years old, etc.
- Of course:
  - if we only have, as inputs, ranges (and not the exact values),
  - the computation results lose some accuracy.
- In this case, an important question is how to select ranges for different inputs so that:
  - within the constraint of preserving privacy,
  - the result will be as accurate as possible.

36. How can we manage security and privacy (cont-d)

- It is known that once we have the values of all the partial derivatives, we can then reasonably easily solve this problem.
- So, we can use the above-mentioned uncertainty quantification solution for machine learning-based data processing.

37. How can we make computations faster and less energy-consuming?

- Training a neural network takes some time.
- For example, for the first version of ChatGPT, training took over a year.
- And we are talking computations on very fast high-performance computers.
- How can we make training faster?
- And how can we make it less energy-consuming?
- Three factors cause the training of machine learning models to be rather long (and rather energy-consuming).
- First, we have a large amount of examples to process.
- Each example includes many features.
- Inputting and processing all the features from all the examples requires significant time and energy;

38. How can we make computations faster and less energy-consuming (cont-d)

- Second, the training algorithms that we use require many computational steps.
- This requires a lot of time and a lot of energy.
- Finally, the hardware on which the algorithms run takes time and energy to perform each operation.
- So, to speed up computations, it is desirable to see if we can decrease the effect of these three causes.

39. How can we decrease the number of features

- A natural idea is to only use the features that affect the result the most.
- Detecting such features is a known problem in system analysis.
- There are many statistical techniques to detect the most important features.
- One of these techniques – actively use in machine learning – is the use of the Shapley value.
- This concept was originally designed to find the collaborators who were most important for the success of a project.
- However, the use of the Shapley value does not always find the truly most important features.
- There are two ways in which we can improve the situation.
- First, the original Shapley value implicitly assumes that we know the exact effect of each combination of factors.

40. How can we decrease the number of features (cont-d)

- In practice:
  - especially when the effect is measured in probabilities – as in machine learning,
  - these probabilities can only be determined approximately.
- It is therefore desirable to use modifications of the Shapley value that take this uncertainty into account.
- Second, the use of Shapley values means, in effect, that we consider a *linear* approximation to the effectiveness.
- In this approach:
  - the predicted value of using only m out of all n features is approximated by
  - the sum of the values corresponding to each feature.

41. How can we decrease the number of features (cont-d)

- Of course, most real-life systems are more complex than the linear ones.
- So, a natural idea is to use *non-linear* generalizations of the Shapley value concept.

42. How can we decrease the number of needed examples?

- We humans do not need thousands and millions of examples to learn a new concert.
- For us, dozen (or so) examples is usually enough – and a few hundreds is definitely enough.
- How do we do it?
- It is often said that:
  - the best teacher is not the one who best teaches the class material,
  - the best teacher is the one who teaches students how to learn.
- This is why we humans are so good (in comparison with the current neural networks) in learning new concepts and new ideas.
- Because we have learned how to learn.
- So, instead of training a neural network for each specific case, let us teach the neural network how to learn.

43. How can we decrease the number of needed examples (cont-d)

- Specifically, let us give, to a neural network, many examples, from different domains, of pairs in which:
  - the input consists of several examples of the desired problem-solution pairs, plus a new problem of similar type, and
  - the output is the solution to the last problem.
- When we feed a neural network pairs (image, animal name), it learns to recognize an animal from an image.
- So with this new training, we will feed it examples of a new concept:
  - it will learn how to use this concept.
  - just like we learn how to use it.

44. How can we decrease the number of needed examples (cont-d)

- For example, we give it:
  - many examples of addition triples, like $2+2=4$, $3+5=8$, etc., and
  - a new similar example, e.g., $3+6$.
- Then, the trained neural network will (hopefully) return 9.
- This idea comes from emulating how we humans do it – so it related to the cybernetic aspects.

45. Comment

- This may also be related to the fact that many of us boost their productivity when attending face-to-face conferences.
- During several days, we hear a lot of pairs (problem, solution).
- This boosts our own ability to find solutions to complex problems.
- This is just like in elementary school:
  - the process of going through many examples of addition that the teacher showed on the board
  - helped us to better perform addition (or multiplication, or whatever concept it was).

46. What can we do on the hardware level?

- In modern computers, the only thing that is moving is electrons.
- They move for two reasons:
  - when we need to communicate a bit (or a sequence of bit) from one location to another, and
  - when we need to change the state of the basic cell from 0 to 1 (or from 1 to 0).
- The time needed for each such movement is equal to the distance divided by speed.
- Electrons already move with a speed close to the speed of light.
- So the only way to decrease time is to decrease the distance.
- Similarly, the energy needed for each movement is proportional to the distance.
- So, the only way to decrease the distance is also to decrease the distance.

47. What can we do on the hardware level (cont-d)

- Already now, within a cell:
  - the distance that electrons need to travel – when the state changes from 0 to 1
  - is the size of a few thousand molecules.
- If we decrease this distance even more, we will eventually get the size of a few atoms.
- Moving electrons between atoms is what is happening during chemical reactions.
- So we arrive at the idea of *chemical computing*.
- Chemical processes is the main way processes are happening in our brain.
- Hence this falls under *cybernetic* aspects.
- So, it is desirable to pursue chemical computing.

48. What can we do on the hardware level (cont-d)

- In this pursuit, it is desirable to take into account that:
  - the larger the concentrations,
  - the faster the reaction.
- Even the simulation of high-concentration chemical reactions can indeed speed up computations.
- If we go even further in decreasing size, we get to micro-objects.
- For describing them, Newtonian physics is no longer a good approximation.
- We need to use quantum physics.
- Computing on this level will thus be, in effect, *quantum computing.*

49. Let us summarize

- Let us summarize the above analysis of:
  - how systems S, human H, and cybernetics C-related ideas and techniques
  - can help modern AI to become even more effective (and less flawed).
- We need to combine traditional AI techniques and fuzzy techniques with machine learning: H.
- We need to further develop uncertainty quantification, security, and privacy techniques to machine learning models: S.
- We need to make AI more explainable: H.
- We need to more adequately find the most important input features.
- For this, we need to further explore the use of non-linear and/or uncertainty-affected generalizations of the Shapley value: S.

50. Let us summarize (cont-d)

- We need to use the ability of a neural network to learn:
  - not only to make it learn a specific material,
  - but also to learn how to learn by itself: C.
- We need to pursue the possibility of using chemical (and quantum) computing: C.
- Let us do it.
- We cannot guarantee that all these directions would lead to successes.
- In many of these recommendations, we do not have a clear idea of what exactly to do.
- But let us try!
- As a Russian poet Mayakovsky said (it sounds better when rhymed):

  The future will not come by itself, we need to do something about it!

51. Acknowledgments

This work was supported in part:

- by the US 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),
- by the AT&T Fellowship in Information Technology, and
- by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Focus Program SPP 100+ 2388, Grant Nr. 501624329,