# What machines can learn about our complex world - and what can we learn from them?

Tomaso Aste

*Department of Computer Science, UCL, London, UK.*

*UCL Centre for Blockchain Technologies, UCL, London, UK.*

*Systemic Risk Centre, London School of Economics, London UK.*

10/02/2021

**Abstract**

*We live in a complex world where we are overwhelmed by data, but information is hard to extract and the consequences of our actions are difficult to predict. Humans have learned to navigate this complexity quite efficiently. Mathematics, statistics, and machines are still far less capable. In this short note, I provide a few general considerations, and my perspective, about modeling real, complex systems and about the automation of this process with learning machines.*

This is support material to a presentation entitled *"This world is not simple!"* given at UCL enrichment activity on February $10^{th}$ 2021.

# 1 Modeling

Humans constantly use models to interpret reality. We do it consciously and unconsciously, scientifically or instinctively but we need these instruments to negotiate with reality. We use models when we try to understand what is happening and when we try to predict what will happen next. We need them to manage and perhaps reduce the unpredictability of the world. We need them to take decisions. We need them to survive.

Models are the very essence of what we call thinking. Without models, reality will be an overwhelming amount of data with no information. Indeed, in my view, models are the instruments that transform data into information. Nowadays, scientists and engineers are developing automated modeling tools. The purpose is to give machines the instruments to select and learn models, developing their ability to interact with the real world and make autonomous decisions.

There are many different kinds of models; the ones we use, for instance, to choose our food at the market are different from the ones we use to place our satellites into the right orbit. As an example, let's think about the model we use when we cross a busy road. Let's make an effort to analyze what we normally, spontaneously, do in this situation. The process is surprisingly complex and sophisticated. Normally, we start looking if there is any vehicle and, if there isn't, then we cross. In the case there are vehicles, then we gauge their distance and if they are far enough we cross. If the vehicles are not too far, then we estimate their speed and we compute if the time they will take to arrive is sufficient for us to cross the road. There are many other variables that we evaluate. We distinguish cars from trucks and trucks from buses; they have indeed different typical speeds and behaviors. We might consider the vehicle's trajectory and, in some cases, we might look at the driver and guess his or her intent to let us pass. We also ponder, and normally dismiss, other variables that are in most cases marginal, for instance, the color of the vehicle, its brand, or the age of the driver. This operation is an example of data-driven modeling. We have learned these operations from observations and, perhaps, from imitation. We don't have a general theory; we do all this spontaneously and instinctively without noticing the amount of data that we are processing and the way we make predictions and compare and use them. We consider this *simple* but actually it is a *complex* operation that we, modern humans, can perform quite easily but that is extremely hard to encode in machines or formulate it precisely with our current mathematical tools.

# 2 Different kinds of modeling

Models are used to interpret and understand reality. Some models help us to distinguish between different scenarios such as distinguishing between friends and foes or isolating food from poison. For instance, one could have a picture of an animal as *input* and the information about the kind of animal provided by the model as *output*. Similarly, one could have a set of observations of a physical system, for instance, the positions of the planets in the sky over some time, and the model might output some mechanism to compute the motion of the planets. Models are also used to infer the internal relations between a set of variables and the problem consists of uncovering the structures of similarities, hierarchies, dependencies, causalities that are in the data-structure. Models are also used to predict outcomes that have not been observed yet. We use them constantly in our everyday life, to survive; for instance, to avoid being crushed by a bus when we cross the road. The accuracy of the model prediction in various circumstances is a measure of the goodness of the model. All these modeling tasks are not necessarily distinct and can be seen as different ways of addressing a problem and interpreting the model. Indeed, the distinction between input and output variables is in most cases just a useful convention and the laws that map the input into output coincides with the dependency structure of the dataset. Further, in a statistical setting, prediction consists in inferring the dependency between variables at different times.

Models are tools to transform data into useful information which can be used for actions. Data, observations, can be images, movies, journal articles, chats, financial prices or electric signals, or –indeed– anything that is produced by some process and carries some information. Data can always be represented as points and shapes in a space. The space might not be ordinary; it is often high-dimensional and sometimes non-euclidean. At this level of abstraction, models are maps between points in these spaces. These maps must both guide us as precisely as possible to the positions of the existing observations as well as to the likely positions of future observations. The model provides a *description* of the system by locating the observations and their interrelation in these spaces; the model can also provide *prediction* by inferring the presence and location of unobserved points.

Models can be deterministic, perhaps including some degree of uncertainty due to limited and noisy observations, or they can be probabilistic. Newton's gravitation law and the consequent modeling of the motion of celestial bodies is an example of deterministic modeling which can be written in the form $\boldsymbol{y} = f(\boldsymbol{x})$ and for which the true model can be *learned* with arbitrary precision provided a large enough number of observations.[1] The Maxwell-Boltzmann distribution describing the speed of particles in a gas is an example of probabilistic modeling where the law to be learned is the probability that a particle has a velocity smaller or equal to a given value $v$: $Probability(V \leq v)$.[2] This is a different kind of problem than the previous; however, also in this case we could formulate the problem in the form $\boldsymbol{y} = f(\boldsymbol{x})$ where, in this case, the relation represents a form of *probabilistic* dependency between the variables. Some problems cannot be represented in terms of an input and an output and the modeling task becomes the discovery of mutual dependency relations between the variables. The goal becomes to map the structure of the interactions between the system's variables and uncover their similarities, their hierarchies, and causal relations. In other contexts, one might aim instead to discover emerging behaviors from simple rules, and often for this task simulations are adopted. These are all forms of modeling, and they all serve the same purpose of helping us to navigate reality by describing what is happening and making a prediction on what will happen. In other terms, models help us to interpret existing observations and make predictions about future ones.

In some contexts, a meaningful modeling approach consists in starting from the elementary components of the system and modelling their behaviors and interactions from first principles. This microscopic modeling can produce precise explanations of the system behavior and can help to understand the origin of macroscopic phenomena from fundamental microscopic laws. For instance, this approach is sometimes used to derive the properties of materials from the constituting atoms in so-called *ab initio* simulations. However, in many real and complex systems, there is nothing comparable to the atoms. Indeed, in these systems, the elementary components are complex themselves and their laws of interactions are often unknown. In complex systems, microscopic modeling approaches often end up either being unrealistic oversimplifications or being too complex themselves to be of use to explain the underlying system.

# 3 The scientific method

Since at least the sixteenth century, science has developed its way to build and validate models that has been given the name of *the scientific method* (Popper, 1934; Gower, 1997). The scientific method is a circular approach based on observation of the system, formulation of a model, testing the goodness of such model through further observations, and comparison of the results with the ones obtained with alternative models, all under the principle of parsimony (see figure 1).

Making models that can produce predictions is at the core of scientific research. Being able to make predictions is essential to formulate hypotheses and theories

[1] Here I use the notation $\boldsymbol{x}$ and $\boldsymbol{y}$ to generically indicate two sets of input and output variables

[2] Where $V$ is the random variable indicating the velocity of a particle and $v$ is a value.

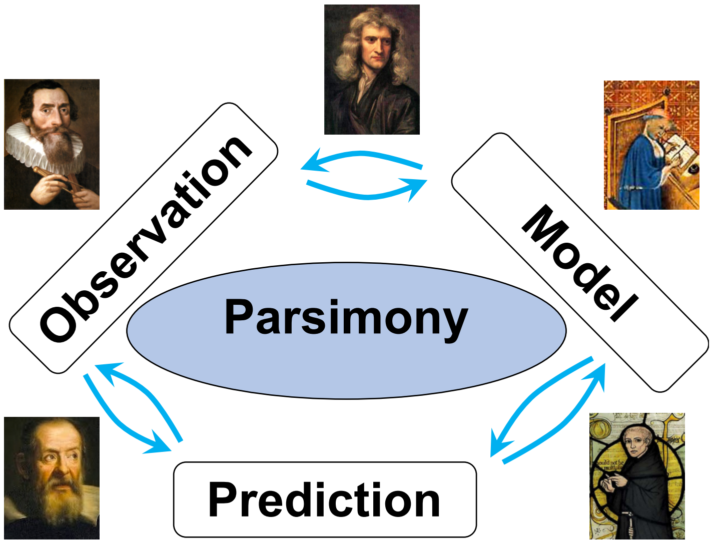


Figure 1: The scientific method is a circular process: from observations (data) we formulate models that make predictions that are tested against further observations, all all within the principle of parsimony where simpler models are preferred. The people represented in the images are, from the top right in clockwise order: Nicholas Oresme (1320-1382); William of Ockham (1285-1347); Galileo Galilei (1564-1642); Johannes Kepler (1571-1630); Isaac Newton (1643-1727). These are philosophers and natural philosophers who have greatly contributed to science and to the elaboration of the scientific method.

that can be tested with further observations and can be iteratively refined or be discarded and changed to obtain better predictions. Prediction, in this context, has a vast meaning embracing all inferences or relations or gatherings around common classifications not provided directly by the observations or not provided precisely by them. Model complexity is another key ingredient; the principle of parsimony recommends that a simpler model with fewer parameters and fewer assumptions should be preferred to a more complex model if both produce comparable predictions. This is an application of the *Occam's razor* principle introduced, in the thirteenth century, by an English Franciscan friar, William of Ockham (pictured on the bottom right of figure 1).

Epistemologists have repeatedly shown that the scientific method has not been always rigorously followed by scientists Feyerabend (1975). There are several situations where the scientific method is hard to apply, for instance, in situations where experiments cannot be repeated because the system under examination is unique and evolves with time. Nonetheless, its circular approach, that starts and ends with observations passing through modeling and prediction, is a very powerful methodology that has been at the very basis of most of the knowledge about natural and artificial systems that humankind has acquired in the last five hundred years. Furthermore, modern day machine learning and artificial intelligence approaches are rooted in this kind of circular approach and aim to implement it automatically without the assistance of human scientists. State-of-the-art machine learning and data-science protocols are expanding the observation part, digging and manipulating vast datasets; they generate vast quantities of alternative models and they place more emphasis on prediction and they tend to discard parsimony, but still, the general circular approach remain the same.

# 4 Building models (autonomously) from data

Historically, models were generated from human intuition built upon experience and then they were tested and validated, or eventually falsified, with observations and experiments.

This is a *top down* approach where mechanisms and interactions are assumed a priori and then data are used to calibrate and validate the model.

There is no need to argue or demonstrate that one of the novel elements of this epoch is the abundance and availability of data. This, combined with unprecedented access to powerful computational capabilities, is radically changing the way we are constructing and operating our models. Rather than following scientists' genius intuitions, increasingly models are constructed directly from data. The appealing aspect of this data-driven model construction process is that it can be automated. The challenge becomes to test the validity and goodness of the model, which must perform not only on the dataset from which it has been constructed but also in other circumstances, perhaps even in situations that never happened before. This is a *bottom up* approach that starts from the analytics of observation data and constructs the model directly from the data. Data-driven modeling is not only a new opportunity, but a necessity because the complexity of most of the problems we are tackling nowadays leave less space for the top down approach. It is indeed very unlikely that someone one day will come up with a genius idea to solve the problem of recurring financial crises or the prediction of earthquakes. These problems are much more complex than modeling the motion of the heavenly bodies and the top down approach based on assumptions from human intuition is less effective and of narrower applicability.

Automation of the bottom-up modeling approach expands enormously the number of models that can be proposed. The focus, therefore, moves towards their selection and validation by testing their prediction capability on data.

Data scientists and machine learners have increasingly powerful instruments that allow them to build increasingly powerful models using increasingly large data sources. However, scientists are also facing increasingly challenging problems involving increasingly complex systems, using data that are complex themselves. For instance, if we aim to construct a model to monitor the state of the economy and predict the possible unfolding of a financial crisis, we must build on past data but we must be aware that, most likely, the next crisis will be different from the ones that happened in the past. The lesson from the past experience must be 'learned' by the model; however, it is essential that the model is able to handle events that have not been observed in the past. This is clearly a hard task and we still have not developed all the instruments to achieve it.

# 5 Modeling real and complex systems

Complex systems are not abstract objects. They are very real and they are everywhere. Humans are complex systems and so are human societies. Animals –even the simplest–

are complex systems and many human-created artificial systems –such as financial markets– are complex themselves. The defining characteristic of complexity is that simplification is not possible without losing crucial properties of the system. In complex systems, important properties *emerge* from the combination of many different elements. Although their elements might be known, the emergent property is the result of their combination and prediction is hard to achieve from the analysis of the constituting elements in isolation (see Parisi (2002); Boccara (2010)). An example is a living organism, even a simple one, which is made of parts with properties and functions that might be well known and understood. However, even the deepest knowledge of all its parts, will not reveal the most important property of the organism: the fact that it is alive. Being alive is the emergent property of the whole system which is of greater importance than the sum of the parts. Understanding emergent properties is a very important part of complex systems modeling, it is –literally– a question of life and death, and it is what makes such modeling very challenging.

Henry Louis Mencken –known as the Sage of Baltimore– once noted that 'For every complex problem there is an answer that is clear, simple, and wrong'. Nonetheless, we can devise effective models to describe and predict, at some level of accuracy, the behavior of complex systems. We do it all the time. Indeed, we are complex creatures who live in a complex world. The modeling of complex systems has the same nature and scope as the modeling of any other system. In complex systems, models are used to describe the system and to predict its behavior. However, the task can be more arduous. For instance, often in complex systems not only is the system non-deterministic but also the internal rules change. Nonetheless, the modeling of these systems is still based on the scientific method's circular approach. The struggle we are facing in modeling some real, complex, systems could be showing us the horizon of applicability of the scientific method. One day, scientists, or machines, might develop alternative routes.

# 6 Non-linear modeling

Most of the scientific problems solved in the last few centuries are associated with finding a linear map $y = b_0 + b_1 x_1 + b_2 x_2 + ....$ between two subsets of variables.[3] Linear models have the great advantage that a solution can be always obtained with arbitrarily small error given a large enough number of observations. Furthermore, approximate solutions with similar error levels have coefficients $b_i$ which are at a similar distance from the coefficients of the exact solution and they tend towards the exact solution when the error goes to zero. However, a very large number of problems cannot be addressed with linear modeling. These are typically the problems associated with complex systems where the non-linear interplay between a large number of variables is a crucial feature of the systems, that cannot be simplified. In general, when non-linear models are considered, the complexity of the problem increases considerably. Indeed, there exists an almost infinite number of models (the number grows combinatorially with the number of variables) that provide approximate solutions with equivalent error levels but that are totally different from

[3] Or some function of the variables $g_0(y) = b_0 + b_1 g_1(x_1) + b_2 g_2(x_2) + ....$, with the $b_i$ constants (independent from $x_k$ and $y$).

each other.[4] For these systems, finding a good model and understanding its validity is a very delicate task. Noticeably, the instruments at the core of present-day artificial intelligence –neural networks– are highly non-linear modeling tools (Anthony and Bartlett, 2009; Mohri et al., 2018).

Models are used to describe systems and to predict their behavior. Description and prediction are two actions with no simple distinction. Often description is associated with the present and past *state* of a system whereas prediction is about the future state. If we can describe well what is the present state of a system, we might also be able to infer what will happen next. However, the two are not entirely coincident. Sometimes a good description of a system tells very little about possible future outcomes, and conversely, there are examples of good predictive models that tell very little about the system.

We might argue that the construction of models that are highly predictive but little descriptive is one of the most controversial issues that emerged in scientific research in the last decade. Indeed, neural networks produce models that sometimes can be very good for prediction, but they are often useless for description. This modeling approach is putting into question the very meaning of modeling and, I would say, the actual meaning of knowledge. Without a theoretical underpinning that humans can understand and convey to each other, on what basis can we be confident that the model will continue to provide accurate and reliable prediction in the future? What do we really know about a system in which we can predict its behavior but cannot explain why? And what does explaining really mean?

These are philosophical issues, hard to resolve. However, it is important to be conscious that we are witnessing a radical change in the way systems are analyzed and modeled and that this radical change could have profound implications for what we presently call scientific knowledge. In particular, how to interpret results becomes blurry when non-linear modeling is used. Indeed, if very different models provide similar results and similar prediction capability, the function of the model as a knowledge-building instrument that guides us in the interpretation of the system becomes questionable. Specifically, the meaning of description fades away because if the same outcome can be reached by completely different models then the details of the model and its *state* have little overall meaning. One might question if prediction alone can be considered a sufficient outcome of a model. To some extent, it must be sufficient because if we want to use the model to navigate reality we only need a precise navigation forecast independently from the mechanism that is used to reach it. From a blackbox prediction machine, we do not learn much about the system and its fundamental laws and principles, but we still could attempt to use the blackbox prediction machine to infer such laws in a similar way as we do with simulations. However, suppose for instance that we have built a blackbox machine that predicts well the motions of the major planets. Would it tell us anything about gravity? Suppose, that with an upgraded model, we would manage also to predict accurately the orbit of Mercury. Would we have improved our knowledge from Newton's dynamics to Einstein's general relativity?

[4]Also linear models can have multiple, equivalent solutions when they have more variables than constraints and therefore are underdetermined. However, non-linearity adds an higher level of complexity with solutions that are disconnected from each-other.

# 7 Model interpretability

The interpretability of artificial intelligence models is an actively debated topic that sometimes comes under the name of *explainable artificial intelligence*. The model interpretability issue touches on several fundamental issues and questions the essence of what we define knowledge. It has also several practical consequences from ethics to reliability. In this domain, most questions rest unanswered and it is not the purpose of this short note to provide solutions. What I want to stress here is that the problem of interpretability is not an exclusive problem of some specific artificial intelligence approaches, such as neural networks; it is the unavoidable consequence of the shift of present-day modeling towards data-driven non-linear tools.

To explain artificial intelligence, blackbox prediction machines can be opened up and their functioning mechanisms might be decrypted in full detail. However, the non-uniqueness of the solution, and the existence of a combinatorially large number of very different machines which provide similar prediction power, makes the interpretation of each solution meaningless. Data-driven modeling with non-linear tools provides us with millions of completely different 'theories' which might work fine. However, differently from what the top-down scientific approaches have done in the past, these competing theories are hard to select or combine in a unified paradigm that can advance what we become accostumed to calling scientific knowledge.

# 8 Conclusions

In this short note, I have highlighted a few aspects that, I believe, are relevant for present challenges regarding data-driven modeling and the automation of the modeling process for real, complex systems. I did not try to be comprehensive, I have just tried to provide my perspective and, in doing so, I have certainly ignored some parts of the current debate and several past contributions to the field.

Science is adopting these modeling instruments, that come under the names of machine learning and artificial intelligence. How scientists operate might radically change as a consequence, posing crucial questions as to the meaning of modeling, explanation, knowledge, and understanding. Science has been a collective human effort to understand the world which has been grounded on empirical observation and has been built on a shared and a continuously debated vision of the world. Will machines learn to do science autonomously? They could, but will scientists be able to understand them?

My perspective is that the fundamental paradigm shift that we are witnessing at the present time concerns the use of non-linear data-driven tools for modeling which have proved to be effective in yielding predictions but puts into question the traditional meaning associated with description and explanation. An accurate description of the system and explanation of the process used to obtain the results, might not be always necessary if the blackbox learning machine provides reliable and accurate predictions. However, there are contexts, such as ethics, law, and regulation, where the explanation of the decision-making process is a fundamental requirement. Furthermore, the lack of interpretability and absence of meaningful descriptive capability call into question the very essence of what we call knowledge.

Our society is becoming increasingly reliant on these modeling instruments, yet we still lack the fundamental understanding of these tools. I feel that it is important and urgent to contribute to the debate on the fundamental aspects of present-day machine learning modeling, and this note is my modest contribution to it.

# Acknowledgments

I wish to acknowledge the Financial Computing and Analytics group members and other colleagues at UCL for very useful discussions and for reading and commenting the draft version of this note.

# References

M. Anthony and P. L. Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.

N. Boccara. *Modeling complex systems*. Springer Science & Business Media, 2010.

P. Feyerabend. *Against Method: Outline of an Anarchistic Theory of Knowledge*. New Left Books, 1975.

B. Gower. *Scientific method: An historical and philosophical introduction*. Psychology Press, 1997.

M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2018.

G. Parisi. Complex systems: a physicist's viewpoint. *arXiv preprint cond-mat/0205297*, 2002.

K. Popper. *The logic of scientific discovery*. Julius Springer, 1934.