




POLITECNICO
MILANO 1863

DIPARTIMENTO DI MECCANICA



Artificial intelligence and statistics for quality technology: an introduction to the special issue

Bianca Maria Colosimo, Enrique del Castillo, L. Allison Jones-Farmer, Kamran Paynabar

This is an Accepted Manuscript of an article published by Taylor & Francis in journal of quality technology on 01 Nov 2021, available online:

<http://www.tandfonline.com/10.1080/00224065.2021.1987806>.

This content is provided under [CC BY-NC-ND 4.0](https://creativecommons.org/licenses/by-nc-nd/4.0/) license



Artificial Intelligence and Statistics for Quality Technology: an introduction to the special issue

Bianca Maria Colosimo^{a*}, Enrique del Castillo^b

L. Allison Jones-Farmer^c, Kamran Paynabar^d

^a Department of Mechanical Engineering, Politecnico di Milano, Milano, Italy

^b Department of Industrial and Manufacturing Engineering, The Pennsylvania State University, University Park, PA

^c Farmer School of Business, Miami University, Oxford, Ohio

^d H. Milton Stewart School of Industrial and Systems Engineering (ISyE), Georgia Institute of Technology, Atlanta

*corresponding author: biancamaria.colosimo@polimi.it

Abstract

In many applied and industrial settings, the use of Artificial Intelligence (AI) for quality technology is gaining growing attention. AI refers to the broad set of techniques which replicate human cognitive and analytical skills for problem solving, including Machine Learning, Neural Networks and Deep Learning. This paper presents a brief introduction to the special issue, where AI-based solutions are presented to solve problems that are typically faced in the area of quality technology. Limits and advantages of AI-based solutions are briefly discussed to stimulate creative attention to novel solutions and new directions for future research.

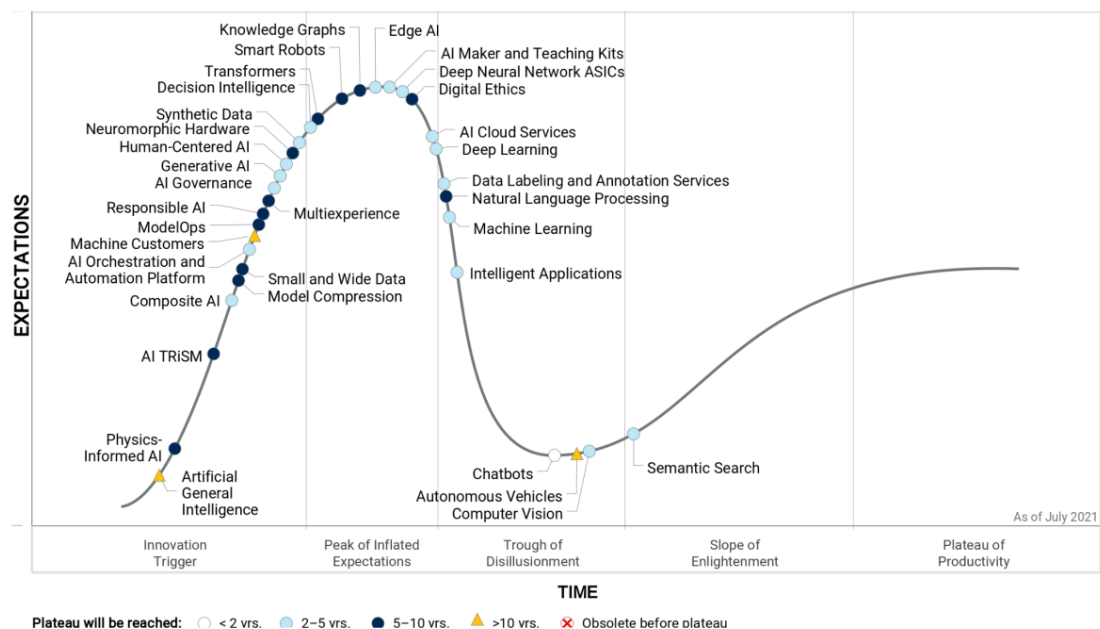
Keywords: Artificial Intelligence, Machine Learning, Quality Technology, Statistical Process Monitoring, Design of Experiments

Introduction

We are continuing to see a significant increase of approaches based on Artificial Intelligence (AI) in our daily life. Similarly, manufacturing, business and services industries are more often including AI solutions to improve the efficiency of their processes, increase the value of their solutions, and design

new products to find new businesses and markets. Nearly all the international roadmaps focusing on innovation and research include AI as a fundamental driver of future technology¹.

Despite the current hype, AI has a long history as its origin dates back to the 1950s, when the English mathematician Alan Turing published the fundamental paper “Computing Machinery and Intelligence” (Turing, 1950) and the computer scientist John McCarthy coined in 1956 the term AI to propose the first Artificial Intelligence conference, the Dartmouth Conference. AI was defined as “the science and engineering of making intelligent machines”, to emphasize a new field aimed at creating machines that can replicate human intelligence to act effectively in a wide variety of novel and possibly unexpected situations (Russell and Norvig, 2021).



Gartner.

Figure 1: Gartner’s “Hype cycle” for Artificial Intelligence, 2021 (Gartner, 2021)

Starting from this early vision, AI is currently embracing a wide spectrum of solutions and approaches, as described in Figure 1, where different AI topics are shown and located in the Gartner’s “Hype cycle of innovation” (Gartner, 2021)² where Gartner illustrates their view on when different types of AI solutions will reach a mature “Plateau of Productivity”. According to this view, few areas within AI have reached the peak of expectations, and AI remains a rapidly maturing field. Regardless of the exactness of the picture and Gartner’s view on the maturity of AI, it is true that different AI solutions have different maturity levels and different expected times to reach full maturity.

¹ <https://www.weforum.org/agenda/archive/artificial-intelligence-and-robotics/>
<https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
<https://www.state.gov/artificial-intelligence/>
<https://hbr.org/2021/02/is-china-emerging-as-the-global-leader-in-ai>

² available here <https://www.gartner.com/en/newsroom/press-releases/2021-09-07-gartner-identifies-four-trends-driving-near-term-artificial-intelligence-innovation>

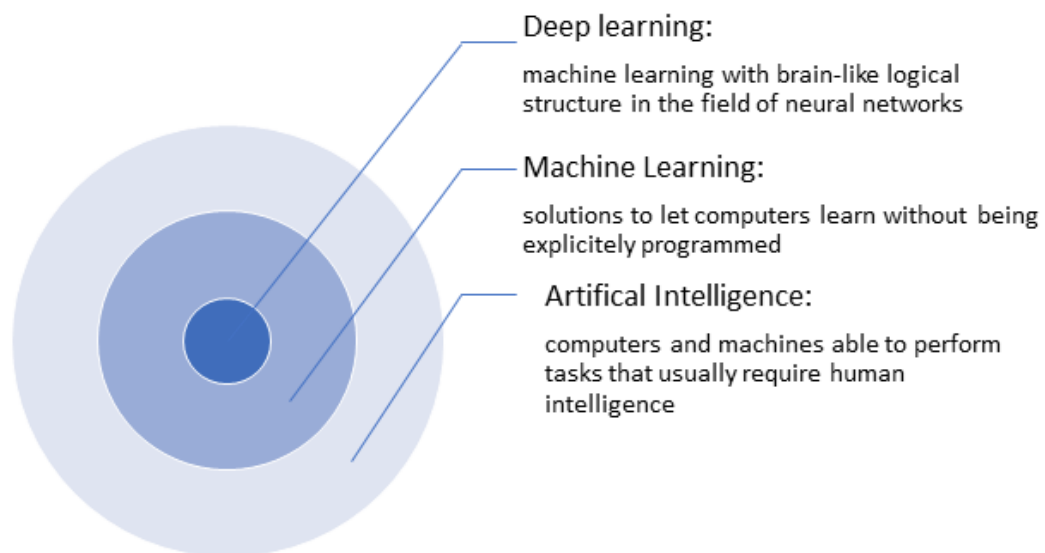


Figure 2: AI, Machine Learning and Deep Learning

It is important to clarify the terms AI and Machine Learning (ML). Figure 2 illustrates a commonly accepted illustration of these terms, showing ML and deep learning (DL) as subsets of a broader field of AI. Thompson et al. (2021) state that “AI is the broad science of mimicking human abilities” and ML “is a specific subset of AI that trains a machine how to learn.” Murphy (2012) gives a more in-depth definition of ML, defining it as “a set of methods that can automatically detect patterns in data, and then use the uncovered patterns to predict future data, or to perform other kinds of decision making under uncertainty (such as planning how to collect more data!).” With these views of the relationship between AI, ML, and DL, we see solutions developed using AI, ML, and DL as having a significant impact in changing the way in which data are used to drive decisions in the field of quality technology.

Journal of Quality Technology (JQT) views quality technology as a field that emphasizes the application of data analytics to solve real problems related to quality in industry, government, and society. Without question, the “technology” associated with the processes and products of our time are becoming increasingly complex. In addition, the “technology” that generates data from these processes and products is evolving rapidly and opening opportunities for solutions to unprecedented challenges and improvements in many domains. Accordingly, our special issue on “AI and Statistics for Quality Technology” aims to highlight some of the opportunities and challenges to be faced in the near future in the field of quality technology and to demonstrate innovative solutions to these challenges.

In light of the increasing attention to AI and its strong focus on the value of data, this special issue presents a collection of papers showing how ML or DL-based approaches can be usefully considered to solve current problems of interest in quality technology. As the field of quality technology evolves, we see ML, DL and AI solutions emerging as lively and active fields of study that embrace the discipline of Statistics with the power of computer science. The goal of these efforts is to develop automated methods of data analysis with the ultimate goal of continuously improving the quality of products and processes in industry, government and society.

The rest of this paper provides an overview of AI/ML applications in quality technology, and is organized as follows. The next three sections focus on AI solutions for process monitoring, Design of

Experiments, and predictive modeling, respectively. The last section provides a brief discussion on advantages and limits of AI-based solutions for quality problems.

AI for Quality Monitoring and Surveillance

Traditional Statistical Process Control (SPC) methods developed in the 1920's are still popular today. However, as manufacturing and data have become more complex, so have the monitoring and surveillance methods. Some degree of automation, even in traditional methods, has become inevitable given the size and dimensionality of the data sets collected thanks to advances in sensor technology. Likewise, the application areas for SPC methods have broadened significantly. Much of the research on process monitoring in the past decades focused on improving model-based methods to achieve a desired Average Run Length (ARL). As the monitoring scenarios become more complex in practice, attempts to align a model with the monitoring scenario becomes progressively more difficult.

Increasingly in practice and in research we are noticing a focus on data-driven and algorithmic monitoring methods in SPC. Some early attempts at using AI for process control date back to the 1980's (see, e.g. Thomson, 1988). Many of the early papers using ML methods for process monitoring focused on developing control charts under a conceptual framework similar to the Shewhart chart. Some of the work focused on identifying control limits in more complex data scenarios such as autocorrelation or multivariate measurements. Much of the work focused on identifying patterns of observations using Neural Networks. Psarakis (2011) and Zorriassatine and Tannock (1998) provide a review of this early application of Neural Networks to control charts. Woodall and Montgomery (2014) noted that "Despite the large number of papers on this topic [Neural Network-based control charts] we have not seen much practical impact on SPC." In some cases, the performance gains recognized by using Neural Networks in the monitoring frameworks that were similar to the Shewhart approach were marginal. In addition, we found much of this early work difficult to apply in practice due to several factors including lack of generalizability of the methods to a broader class of problems, a lack of practical guidance as to how to adopt many of the methods to different scenarios and lack of documentation and/or compute codes to implement.

Weese et al. (2016) reviewed a broader class of literature on the use of statistical learning methods in statistical process monitoring. Their review cited several main research streams, including dimension reduction methods, one-class classification methods, control chart pattern recognition, profile monitoring, and methods based on risk adjustment, neural networks, support vector machines, ensemble models, and risk-adjusted neural-network-based methods. They noted that many modern processes produce data from multiple streams and/or a hierarchical structure, often with autocorrelation and cyclical patterns. The traditional model-based SPC methods are ill suited to monitoring such data.

Recently, we have seen several important methodological challenges in process monitoring and surveillance that continue to impact SPC practice. Most of these challenges are problem driven and arise from new types of data that originate from technological advances in either the processes that generate the data or the measurement systems that record and store the data. Examples of such challenges include high-dimensionality of data, intricate temporal and spatial structure of data streams, non-stationarity and dynamicity of processes from which data are collected, missing and corrupted data,

unstructured (text) data, and manifold data, that is, cases in which the data are concentrated in a nonlinear subspace.

Traditional SPC tools fall short of addressing these process monitoring challenges as their detection power quickly decays with increased data dimensions, and they often require a stationary process and a fixed baseline, or are based on linear dimensionality reduction techniques. Additionally, before monitoring can begin, one must first establish a baseline (traditionally referred to as Phase I) sample. An initial challenge is how to establish such a sample in a complex, high-dimensional situation. Although there may be classical statistical approaches to many of these scenarios, the complexities in the data from many of these scenarios may lend themselves well to algorithmic solutions using AI tools.

Adopting an AI approach for process monitoring is not without cost. Many ML and DL models are difficult to interpret. This problem is not new to process monitoring, as signal interpretation has often been challenging even when using relatively straightforward multivariate methods. With respect to AI models, there is growing interest and rapid development in building algorithms that are transparent, interpretable and understandable. Creating interpretable AI solutions is equally important in process monitoring, where understanding the source of a process change, or the “root cause” is critical to process improvement.

Once an AI solution is built, in order to add value to a business, it must be deployed into practice. In practice, even the best developed model may not work as expected; thus it is critical to monitor the model implementation. A growing area of focus in industry is developing automated methods for monitoring the implementation of AI models in practice. Similar to process monitoring, monitoring an AI model implementation is necessary to ensure the model remains effective over time. Part of an area known as Model Operations (ModelOps), monitoring AI model performance is a key component of an AI implementation strategy. Many of the model monitoring initiatives have developed outside of the industrial statistics community, and we believe these can be enriched by further study in our field.

In this issue, two sets of authors have provided innovative solutions to challenging modern process monitoring problems using ML and DL tools. The papers address topics on complex profile monitoring with image data and monitoring with incomplete data in a sensor network.

In the paper, “Towards a better monitoring statistic for profile monitoring via variational autoencoders”, in this issue, Sergin and Yan (2021) develop a new formulation for the use of Variational Autoencoders (VAE). The VAE is a type of latent variable model that uses high-capacity deep neural networks in place of linear transformations to project high dimensional data into a “deep” latent space. This work demonstrates limitations in the current uses of VAE for process monitoring and proposes a novel and computationally feasible method that overcomes the limitations in the VAE method. Their work is supported by an extensive simulation study and is applied to image profile monitoring in a hot steel rolling process.

In “Multi-sensor based landslide monitoring via transfer learning”, Li et al. (2021) address the problem of process monitoring with incomplete data in a sensor network with applications in monitoring geographical sensors for landslides. Using a regularized parameter-based learning approach integrated with the LASSO, this research provides a method to effectively transfer information from sensors with sufficient historical data to those with a limited data history. A control chart is also developed to monitor the sensors using the transferred data.

AI/ML and the Design of Experiments

Over the years, there has been a continuous tension in the Applied Statistics literature between authors opposed to and authors in favor of using automatic or algorithmic methods (see, e.g., Breiman, 2001). With the advent of continuously bigger datasets due to the use of better and cheaper sensing equipment in industry, some degree of automation became inevitably necessary.

Even within the field of Design of Experiments (DOE) for industry, a typically “small data” area as experiments are expensive, there has been since its beginnings in the late 40's and early 50's the two aforementioned schools of thought. One school of thought advocated a carefully designed and analyzed sequence of experiments for characterizing, understanding, and eventually, optimizing a process (following the lead of G.E.P. Box and co-workers initiated in the landmark paper by Box and Wilson (1951) which translated and expanded Fisher's agricultural factorial experiments and ANOVA for the optimization of industrial processes). Another school of thought advocated algorithmic methods exclusively for process optimization that depended less on a human, e.g., the initial work of Hotelling, (1941) and Friedman and Savage (1947) on sequential experiments for the optimization of a process continued with the classic work by Robbins and Monro (1951) and Kiefer and Wolfowitz (1952). Both schools have descendants in our days, with the former school the most dominant and successful in industrial practice but the latter having a resurgence in recent decades due to the wide availability of data obtained online or as a “stream” and the emergence of computer experiments and ML methods.

One area where ML methods have moved closer to what DOE techniques are intended for is Active Learning (AL). Supervised ML models usually need to be trained on hundreds or even thousands of labeled instances to perform well. The labels are often provided by an oracle, such as a human annotator, in response to different queries about unlabeled instances. While many modern systems can easily produce a large number of unlabeled instances at low cost, the labeling process can be very difficult, expensive or time-consuming (Settles, 2009). AL has been proposed to solve these problems. The key idea of AL is to let the learning algorithm choose which data to learn from, instead of providing randomly selected labeled instances. The goal is for the algorithm to choose the most informative instance to label so that the learning model will perform better with less training data, thereby reducing the efforts of the labeling process and speeding up learning. The similitudes with DOE are clear: the instances correspond to experimental runs, learning means fitting a model, labeled instances are response values at a corresponding experimental condition, and one can even think of unlabeled data as what in optimal Design of Experiments is called a candidate list of points (for more see Li et al., 2020).

An important and open-ended question in AL is how to measure the informativeness of a data point. There have been many variations of query strategies proposed in the AL literature, such as uncertainty sampling, query-by-committee, etc. (see Settles, 2009 for details). For a binary classification problem, uncertainty sampling is equivalent to querying the instance whose posterior probability of being in any of the two classes is closest to 0.5. For multi-class classification problems, there are other uncertainty measures, such as entropy. An early reference on AL is Cohn et al. (1994), who discussed how optimal design of experiments can guide neural network exploration by selecting the training trajectories that will increase the amount of information the most in a limited number of steps.

A concept related to how to measure information in a point is the acquisition function of sequential bayesian optimization techniques, the descendants of the Kiefer-Wolfowitz stochastic optimization methods. These methods are useful, among other problems, in Computer Experiments, where automatic

algorithms are sensible, since the "experimentation" is not physical. The goal is to optimize as one performs experiments, so the sequence of design points constitutes a search for the optimum. This is in contrast with classic Response Surface Methods (Box and Draper, 2007; Myers et al., 2016; Del Castillo, 2007) where models are fit first and optimization is performed subsequently, either based on first order or second order polynomial models used as local approximations. A most popular acquisition function in Computer Experiments, using a Gaussian Process model as global approximator over a large experimental region, is the Expected Improvement method (Jones et al., 1998) but other such functions have been proposed in more recent years. For a recent review of sequential Bayesian Optimization methods, see Frazier (2018); for Gaussian Process as a main tool in ML see Rasmussen and Williams (2006) and when used as a tool in Computer Experiments see Gramacy (2020). Sequential optimization of a process can be based on convolutional neural networks (Goodfellow et al., 2016) which in turn can also be aided by DOE-Active Learning ideas.

Related to the aforementioned problems is a very large body of classic literature existing in the Optimizing Control literature, too large to discuss here. For a review of the fundamental statistical Dual Control problem, which involves simultaneously learning and controlling a process, see Pronzato (2000), who makes a connection between Dual Control and classical optimal design of experiments. Another point of contact between ML and design of experiments worth pointing out is the reliance many of these methods have on Bayesian techniques, see the work on Bayesian optimization methods based on classical response surface models by Peterson (2004), summarized by Del Castillo (2007) and in the forthcoming book by Peterson and Del Castillo (2022), and contrast with the Bayesian probabilistic view of ML popularized among others by Murphy (2012). Another very large body of literature that should be mentioned is the current emphasis in ML on finding causal effects from *observational* data, using, in particular, Bayesian networks (see e.g., Pearl, 2009, who further developed the classic work on "potential responses" by Rubin, 1978).

We cannot conclude this section without mentioning two other recent connections between DOE and ML/AI: the popular "A/B Tests" for web-based experiments (for instance, to enhance a web site operation, essentially sequences of 2-level one factor tests, see e.g. Kohavi et al., 2007) and the development of new DOE methods in networks of connected components (Eckles et al., 2017). All these inter-related past and recent developments indicate the great current potential for cross-fertilizing research between the two classical schools in design of optimization experiments in industry.

Recent work has demonstrated how modern ML approaches can contribute to the design and analysis of industrial experiments and vice versa, DOE techniques can be useful for better and faster ML. An example of the former case of approach, in this issue, is the paper by Lin et al. ("Forward stepwise Random Forest Analysis of Experimental Designs"). Here the authors show that if the goal of an experimental design is to make predictions, ML tools such as Random Forests can be used to successfully analyze experimental data under such an objective. As an example of the latter type of contribution in this special issue, Lian et al. (2021) ("Robustness with respect to class imbalance in AI classification algorithms") borrow concepts from mixture experiments to study the robustness of some AI/ML classification algorithms. As it is well known, in order for these algorithms to classify well, the quality of the training data is important, a problem not unrelated to the dual control and active learning problems. In its simplest form, that of a two-class problem, class imbalance in the training data, where one class is much more abundant leads to poor performance. These papers are instances of the two-way cross-fertilization between AI/ML and DOE, but more work is still needed in both directions.

AI in Reliability, Prognostic and Predictive Modeling of Quality Data

Predictive modeling in ML focuses on forecasting a quantity or an event by analyzing past behavior and discovering nonrandom patterns in the data. In principle, this is nothing different than the goals in classical predictive statistical inference, which has been widely used in quality and reliability modeling. Examples, of predictive modeling applications in quality and reliability include monitoring autocorrelated processes using ARMA models (Alwan and Roberts, 1988), modeling and control of variation propagation in multistage processes using state space models (Shi, 2014), offline prediction of mean time to failure of a component using various reliability models (Meeker and Escobar, 1998), online prediction of residual useful lifetime of a component by modeling its degradation signals (Gebraeel et al., 2005), process optimization by learning the relationship between the output and process variables and finding their optimal setpoints (Del Castillo, 2007). Despite the common goals, some of the methods and the language used in ML differ, creating barriers for cross-communication. An even bigger difference exists in that some of the types of data sets and volume of data used in ML are not common in traditional industrial statistics.

In ML, different types of data sets, by increasing degree of complexity, include numerical data in euclidean spaces (scalars, vectors, matrices or tensors), streaming (time based) data, network, text, and manifold data (scalars, vectors, and tensor data not on an euclidean space). Most of these data types arise in industrial quality and broad areas of work exist about their modeling. Numerical data in euclidean spaces is the basis for most classical “predictive” statistical methods including regression and their multivariable extensions (less frequently is the analysis of tensor data). Numerical data in euclidean spaces is also common in the analysis of one dimensional profile data, common not only in SPC but also in longitudinal data analysis. Streaming data is time series data in very large volumes and at rates not typically considered in traditional SPC., There is considerable work on network data originating from sociology (e.g. Wasserman and Faust, 1994) and now popularized in different areas in applied science (e.g., Newman, 2010). While there exists a large literature on categorical data analysis which is classic (e.g. see Agresti, 2013), it does not include analysis and modeling of large texts. There are a handful examples of the use of manifold data in industrial statistics, given the required different mathematical background necessary (based on differential geometry), see Del Castillo and Zhao (2020a) for a review.

Traditional statistical methods for predictive modeling have been successfully utilized in small-scale problems under some (restricting) assumptions on the data and the process structure, e.g., independence, stationarity and known distribution of data. However, these methods may not be applicable to large datasets with numerous variables or data streams collected with a high-sampling frequency, or to datasets that have functional structure such as images, point-clouds, manifolds, and networks. Additionally, if data are collected from a dynamic process (e.g. in continuous manufacturing processes), the estimated predictive models should be continually updated over time. These data/process characteristics require a new set of flexible and scalable analytical tools for predictive modeling, which are provided by ML and AI.

Low-dimensional learning methods used to reduce the dimensionality in modern process control data, for instance, can greatly benefit from ML methods. Here the subspace where the data lie is not necessarily a linear subspace, thus traditional PCA analysis is not sufficient. Methods including tensor

analysis, functional data analysis, manifold learning, dictionary learning, matrix factorization, and low-rank recovery have shown to be effective ML tools to address the foregoing challenges in creating predictive models for quality. For recent examples, Reisi et al. (2020) used tensor analysis for building a predictive model for multiple high-dimensional (HD) inputs and one HD output. Yan et al. (2019) developed a process optimization method for an HD output and LD inputs. Wang and Shi (2021) extended the tensor-based predictive modeling approach to multi-stage processes while Shan and Apley (2008) focused on using high-dimensional measurement data to understand the nature of the variability to facilitate root cause identification. For the purpose of monitoring manifold data, Del Castillo and Zhao (2020b) used the Laplace-Beltrami (LB) operator to measure and monitor the local curvature of a manifold at different points. Del Castillo and Zhao (2020a) review various methods for analyzing manifold data in the SPC and DOE domains. Del Castillo and Colosimo (2011) provide a statistical framework for shape analysis of objects while Bui and Apley (2018) developed ML solutions to detect changes in textured surfaces.

One of the main advantages of traditional predictive statistical models, especially the parametric models, is their interpretability and identifiability. The estimated model parameters can often be associated with the underlying physics/engineering principles governing the process, which make the predictive models interpretable. On the other hand, AI-based predictive models often lack interpretability, and identifying connections between the model components or parameters and the process they are modeling is difficult. Moreover, the main performance metric of this family of AI algorithms is prediction accuracy, usually assessed via multi-fold cross-validation, an idea that originated in mainstream statistics. Thus, in developing these models, little attention has been given to statistical identifiability. For example, in some tensor-based methods different sets of model parameters may result in a similar prediction accuracy, but providing different interpretations (Reisi et al., 2020). Another drawback, especially for AI predictive models, is that they require a very large sample size to achieve reasonable accuracy (Du et al., 2018). This may not always be available in manufacturing applications, where offline experimental data is expensive to collect and online streaming data are not stationary due to the process dynamics.

ML and DL methods have also been used to build predictive models for reliability analysis and prognostics. Hong et al. (2018) discussed the applications of ML and big data analytics in reliability modeling. Fang et al. (2019, 2020) proposed prognostics frameworks for HD degradation signals using tensor analysis and low-rank recovery methods. Zhang et al. (2019) provided a comprehensive review on the applications of deep learning methods in prognostics and health management (PHM). They discussed the applications of various neural networks in analysis of different types of degradation signals including time-series and image streams.

Reliability engineering research can also benefit from ML techniques and methods. Shortle and Mendel (2001) showed how the space of lifetimes of multiple component systems is not euclidean, and proposed a different space for their analysis, using notions of invariance in differential geometry common to manifold learning. There is a long stream of work on Bayesian reliability analysis, given that reliability problems usually have to deal with considerable degrees of uncertainty and prior engineering and physical knowledge can be incorporated under a Bayesian approach, which is common in ML.

In this special issue, three papers focus on development of novel predictive modeling techniques using AI tools with applications in quality and reliability. These papers address topics on multi-input- multi-output predictive models for quality in multistage manufacturing processes using DL, modeling recurrent event processes using boosted trees, and adaptive sensor selection for effective prognostics using deep neural networks (DNN), respectively.

In the first of these papers, “Deep Multistage Multi-Task Learning for Quality Prediction of Multistage Manufacturing Systems,” Yan et al. (2021) propose a deep multistage multi-task learning (DMMTL) framework to predict multivariate quality characteristics in multi-stage manufacturing processes based on multivariate process variables. The proposed model connects multiple inputs and outputs within and between stages using a small set of fully connected hidden variables. Not only can these hidden variables help with reducing the dimension of the problem, but also capture the correlation among multiple variables. They used simulations as well as a case study to validate their work and compare its performance with a few ML-based predictive models including Random Forests and Elastic Nets.

In the paper entitled “Boost-R: Gradient Boosted Trees for Recurrence Data,” Liu and Pan develop a new nonparametric method for estimation of the cumulative intensity function that characterizes recurrent event processes, which have wide application in reliability modeling and prediction. The proposed method, named Boost-R, combines the gradient boosting approach with additive regression trees. It can take the information of both static and dynamic features into account when estimating the cumulative intensity function. In addition to extensive simulation studies, the authors validated their method using a case study on modeling failure processes of oil and gas wells and provide an R package.

Finally, in “An Adaptive Sensor Selection Framework for Multisensor Prognostics,” Liu et al. propose a new framework for adaptive selection of informative sensors to estimate the residual useful life of a system based on multivariate degradation signals. Their framework can be used along a variety of neural-network-based predictive models and make them more interpretable. They extend the Relu regularization and neuron-level dropout widely used in DNNs to sensor-level dropout, which randomly removes a sensor and its corresponding nodes and connections from the NN in each iteration of training. The sensor-level dropout along with a sensor disentanglement penalty helps to determine the importance of each sensor. The proposed framework is tested and compared with benchmark methods using simulated degradation processes of turbofan aircraft engines.

Statistics vs AI: some final conclusions

In his seminal paper “Statistical Modeling: The Two Cultures” Breiman (2001), pointed out that two different cultures coexisted among statisticians. The first culture assumes data as generated by a given stochastic model while the second one uses algorithmic models treating the data mechanism as unknown (Fig 3). According to Breiman, “the statistical community has been committed to the almost exclusive use of data models” (i.e., playing in the first mainstream) leading to “irrelevant theory, questionable conclusions and keeping statisticians from working on a large range of interesting current problems.” Breiman’s experience as consultant before returning to academia allowed him to observe that algorithmic methods from the then incipient ML community could be effectively used to solve many real problems involving large complex data sets or requiring more accurate alternatives to data modeling on small data sets.

Clearly, some of the advantages of AI methods that have been also discussed in previous sections with reference to monitoring, DOE, reliability and prognostics come at the expenses of significant drawbacks in terms of interpretability, explainability and transparency, all features which refer, with different shades, to the opaqueness of AI solutions, in particular DNN and “deep” learning methods. Algorithmic

modeling in ML typically involves black-box algorithms, which can undermine trust in safety-critical domains, reduce advances in scientific knowledge, limit the transferability of approaches to domain shifts and the fairness when approaches are used for decision-making (Marcinkevičs and Vogt, 2020). To overcome all these limitations, eXplainable Artificial Intelligence (XAI) is currently embracing a set of techniques and methods to convert the so-called black-box AI algorithms to white-box algorithms, where the results achieved by the algorithms and all the parameters, variables and steps can be explained and clearly traced.

As authors working in Industrial/Engineering Statistics, we welcome further work that integrates AI, ML, and DL methods in non-opaque, innovative ways to better address old problems or to solve new, relevant problems in modern industry. The data deluge is real and in some domains new methodology is necessary, where some degree of automation is needed. Sensing equipment has advanced at a fast pace, and classical statistics due to “Student”, Fisher and Neyman, based on small samples, are not always relevant in today’s processes. Theoretical statisticians have tried over the past two decades to adapt to the great influence of this other “statistical modeling culture” by defining ML in their own territory under the name *Statistical Learning Theory* (Hastie et al. 2013). The emergence of Statistical Learning Theory rightly forms an almost mandatory course in today’s training of any theoretical or applied statistician. This theory includes computational statistical methods (Efron and Hastie, 2016), which have been incorporated into and shown value in industrial statistics over the past 20 years. However, we believe that totally automatic AI algorithms to monitor, control or optimize industrial processes, while maybe *possible* to eventually achieve, are unnecessary and possibly dangerous. The danger lies in the fact that the fuel for AI is data. A large number of AI/ML algorithms thus require a good deal of data pre-processing. The goodness, correctness, and appropriateness of all data-driven solutions depend on many features of the data, including quality, representativeness of cases, availability of the right variables, as well as ethical and privacy concerns. These elements require human judgement, and we believe this will continue in the foreseeable future. Thus, we envision solutions where AI is used to assist in decision making rather than replace the decision maker. For example, when possible, it is important to fuse pre-existing knowledge, including domain knowledge, physics and engineering laws, and first principles into models rather than rely on completely data and data-driven approaches. Additionally, when using data-driven models, one should also evaluate the models in practice with additional testing of the predictions using DOE and/or by monitoring the model performance. We believe that fusing front-end knowledge with data as well as back-end testing of model performance will help to ensure that AI is used responsibly and delivers value.

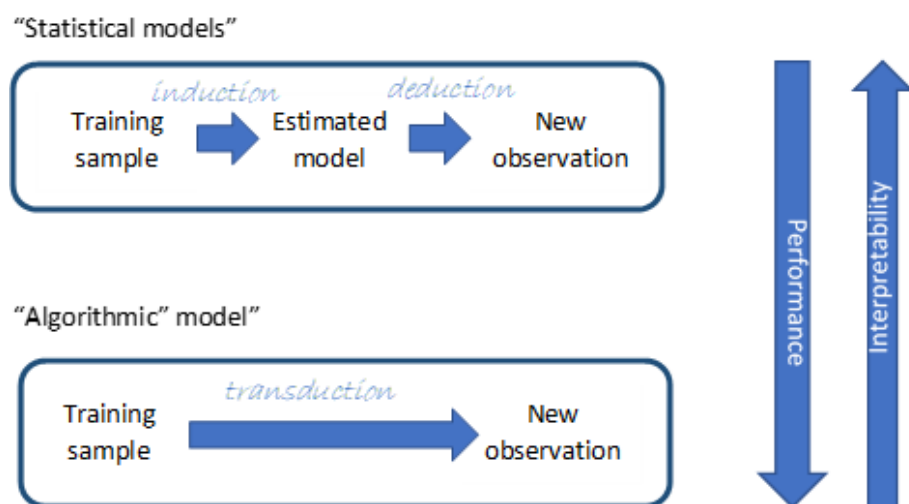


Fig 3: Traditional statistics versus ML – courtesy Simone Vantini (ENBIS Workshop, July 12-13, 2021 - *Interpretability for Industry 4.0*). The notion of transduction, or going from specific training cases to specific test cases, was introduced in the AI/ML literature by V. Vapnik (see Vapnik, 1998).

As with many difficult trade-offs, the efficiency versus transparency conflict is not going to have a clear winner. The present special issue attempts to highlight how appropriate solutions depend on the context of the problem under study, the type of research question/problem to be faced, the size and type of training and testing data, the level of detail required, and the need to transfer knowledge to different domains or to generalize the proposed method to other contexts. In this complex landscape, “the commitment to data modeling should not prevent” the community “from entering new scientific and commercial fields where the data being gathered is not suitable for analysis by data models” (Breiman, 2001). Thus, AI approaches, including ML and DL, should be further investigated and considered as potential competitors or possible solutions to real problems faced in quality technology.

References

- Agresti, A. (2003). *Categorical data analysis* (Vol. 482). NY: John Wiley & Sons.
- Alwan, L. C. and Roberts, H. V. (1988): “Time-Series Modeling for Statistical Process Control.” *Journal of Business & Economic Statistics*, 6 87-95.
- Box G. E. P. and Wilson K. (1951). On the experimental attainment of optimum conditions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 13(1):1, 45, 1951.
- Box G. E. and Draper N. R. (2007). *Response surfaces, mixtures, and ridge analyses*, volume 649. John Wiley & Sons.
- Breiman, L. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3), 199-231.
- Bui, A. and Apley, D. W., (2018) Monitoring for Changes in the Nature of Stochastic Textured Surfaces, *Journal of Quality Technology*, 50(4), pp. 363-378, 2018.
- Cohn D., Atlas L., and Ladner R.. (1994). Improving generalization with active learning. *Machine learning*, 15(2):201 {221.
- Del Castillo E. (2007). *Process optimization: a statistical approach*, volume 105. Springer Science & Business Media, 2007.
- Del Castillo, E. and Colosimo, B. M. (2011). “Statistical Shape Analysis of Experiments for Manufacturing Processes”. *Technometrics* 53(1), pp. 1–15.
- Del Castillo, E., and Zhao, X. (2020a). Industrial statistics and manifold data. *Quality Engineering*, 32(2), 155-167.
- Del Castillo, E., & Zhao, X. (2020b). Statistical process monitoring for manifold data. In *Wiley Stat-sRef: Statistics Reference Online* (pp. 1–8). John Wiley & Sons.

- Lian, J., Freeman, L., Hong, Y., and Deng, X. (2021). Robustness with respect to class imbalance in artificial intelligence classification algorithms. *Journal of Quality Technology*, xxx
- Du, S., Wang, Y., Zhai, X., Balakrishnan, S., Salakhutdinov, R., & Singh, A. (2018). How Many Samples are Needed to Estimate a Convolutional Neural Network? *NeurIPS*.
- Eckles D., Karrer B., and Ugander J. (2017). Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1), 2017.
- Efron, B., and Hastie, T. (2016). *Computer age statistical inference*. Cambridge University Press.
- Fang X., Paynabar K., and Gebraeel N. (2019). [Image-Based Prognostics Using Penalized Tensor Regression](#). *Technometrics*, 61(3), 369-384.
- Fang X., Yan H., Gebraeel N., and Paynabar K. (2020). [Multi-Sensor Prognostics Modeling for Applications with Highly Incomplete Signals](#). *IIE Transactions*, 53(5), 597-613.
- Frazier P. I. (2018). Bayesian optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*, pages 255-278. INFORMS.
- Friedman M and L. Savage. (1947). Experimental determination of the maximum of a function. In *Selected Techniques of Statistical Analysis*.
- Gartner (2021). *Hype Cycle for Artificial Intelligence, 2021*. Gartner Identifies Four Trends Driving Near-Term Artificial Intelligence Innovation. <https://www.gartner.com/en/newsroom/press-releases/2021-09-07-gartner-identifies-four-trends-driving-near-term-artificial-intelligence-innovation>
- Gebraeel, N. Z., Lawley, M. A., Li, R., & Ryan, J. K. (2005). Residual-life distributions from component degradation signals: A Bayesian approach. *IIE Transactions*, 37(6), 543-557.
- Goodfellow I., Bengio Y., and Courville A.. (2016). *Deep learning*. MIT press.
- Gramacy R. B.. (2020). *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- Hastie, T., Tibshirani, R., & Friedman, J. (2013). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction: Data Mining, Inference, and Prediction*, Second Edition, 2013, Springer Series in Statistics.
- Hao Y., Sergin N. D., Brenneman W. A., Lange S.J., and Ba S. (2021). "Deep multistage multi-task learning for quality prediction of multistage manufacturing systems." *Journal of Quality Technology* (2021): 1-27.
- Hong, Y., Zhang, M., and Meeker, W. Q., (2018), Big Data and Reliability Applications: The Complexity Dimension, *Journal of Quality Technology*, Vol. 50, pp. 135-149.
- Hotelling H.. (1941). Experimental determination of the maximum of a function. *The Annals of mathematical statistics*, 12(1):20-45, 1941.
- Hussain, F., Hussain, R., and Hossain, E. (2021). Explainable Artificial Intelligence (XAI): An Engineering Perspective. *arXiv preprint arXiv:2101.03613*.

- Jones D. R., Schonlau M., and Welch W. J.. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4): 455-492, 1998.
- Kiefer J. and Wolfowitz J. (1952). Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, pages 462-466.
- Kohavi R., Henne R. M., and Sommerfield D. (2007). Practical guide to controlled experiments on the web: listen to your customers not to the HiPPO. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 959-967, 2007.
- Li, H. Del Castillo E., and Runger G.(2020). On active learning methods for manifold data. *Test*, 29(1):1-33, 2020.
- Li, W., Tsung, F., Song, Z., Zhang, K., & Xiang, D. (2021). Multi-sensor based landslide monitoring via transfer learning. *Journal of Quality Technology*, xxxx
- Lin, C. Y. (2020). Forward stepwise random forest analysis for experimental designs. *Journal of Quality Technology*, xxx.
- Liu, X., and Pan, R. (2021). Boost-R: Gradient boosted trees for recurrence data. *Journal of Quality Technology*, xxx.
- Marcinkevičs, R., & Vogt, J. E. (2020). Interpretability and explainability: A machine learning zoo mini-tour. *arXiv preprint arXiv:2012.01805*.
- Meeker, W. Q. and Escobar, L. A. (1998). *Statistical Methods for Reliability Data*. Wiley, New York.
- Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., & Yu, B. (2019). Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44), 22071-22080.
- Murphy K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Myers R. H., Montgomery D. C., and Anderson-Cook C. M. (2016). *Response surface methodology: process and product optimization using designed experiments*. John Wiley & Sons.
- Newman, M., (2010). *Networks*. Oxford university Press.
- Peterson J. J.. (2004). A posterior predictive approach to multiple response surface optimization. *Journal of Quality Technology*, 36(2):139-153.
- Peterson, J. and del Castillo, E., (2022) "Process Optimization with Multiple Response Variables: A Predictive Distribution Approach with R and Stan", forthcoming, CRC Press, 2022.
- Pearl, J., (2009). *Casualty, Models, Reasoning and Inference*, 2nd ed. Cambridge University Press.
- Pronzato, L. (2000). Adaptive optimization and D-optimum experimental design. *Annals of Statistics*, 1743-1761.
- Psarakis, S (2011). "The use of neural networks in statistical process control charts." *Quality and Reliability Engineering International* 27, no. 5 2011: 641-650.

- Rasmussen C. E. and. Williams C. K. I (2006), Gaussian Processes for Machine Learning, the MIT Press, 2006, ISBN 026218253X. 2006 Massachusetts Institute of Technology.
- Reisi Gahrooei, M., Yan, H., Paynabar K., Shi, J. (2020), Multiple tensor-on-tensor regression: An approach for modeling processes with heterogeneous sources of data. *Technometrics*.
- Robbins, H. and Monroe, S. (1951). A stochastic approximation method. *The annals of mathematical statistics*, 400-407.
- Rubin, D.B. (1978). "Bayesian inference for causal effects: the role of randomization". *Annals of Statistics*, 7, pp. 34-58.
- Russell S. and Norvig, P. (2021) *Artificial Intelligence: A Modern Approach*, 4th Edition. 2021. Pearson.
- Sergin, N.D., and Yan, H. (2021). "Toward a better monitoring statistic for profile monitoring via variational autoencoders." *Journal of Quality Technology* 2021:xxx.
- Settles, B. (2009). Active learning literature survey. - <https://minds.wisconsin.edu/handle/1793/60660>
- Shan, X. and Apley, D. W.,(2008) Blind Identification of Manufacturing Variation Patterns by Combining Source Separation Criteria, *Technometrics*, 50(3), pp. 332—343.
- Shi J. (2014) Stream of Variations Analysis. In: Baillieul J., Samad T. (eds) *Encyclopedia of Systems and Control*. Springer, London.
- Shortle, J.F. and Mendel, M.B., (2001). Physical foundations for lifetime distributions. In *System And Bayesian Reliability: Essays in Honor of Professor Richard E Barlow on His 70th Birthday* (pp. 257-266).
- Thomson, A. C. (1988) "Real-time artificial intelligence for process monitoring and control." *IFAC Proceedings Volumes* 21, no. 13 (1988): 67-72.
- Thompson, W., Hui, L. and Bolen, A.. (2021). "Artificial Intelligence, Machine Learning, Deep Learning and More." SAS. SAS. 2021. https://www.sas.com/en_us/insights/articles/big-data/artificial-intelligence-machine-learning-deep-learning-and-beyond.html.
- Turing, A. M. and Haugeland, J. (1950). *Computing machinery and intelligence* (pp. 29-56). Cambridge, MA: MIT Press.
- Tsung, F., Song, Z., Zhang, K. (2021) "Multi-sensor based landslide monitoring via transfer learning." *Journal of Quality Technology* 2021: xx-xx
- Vapnik, V. N. (1998). *Statistical learning theory*. New York: Wiley.
- Wang, A., and Shi, J. (2021). Holistic modeling and analysis of multistage manufacturing processes with sparse effective inputs and mixed profile outputs. *IIE Transactions*, 53(5), 582-596.
- Wasserman, S., and Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press.

Weese, M., Waldyn M., Fadel M. M., and Jones-Farmer A.. (2016) "Statistical learning methods applied to process monitoring: An overview and perspective." *Journal of Quality Technology* 48, no. 1, 4-24.

Woodall, W. H., and Montgomery, D. C. (2014), "Some Current Directions in the Theory and Application of Statistical Process Monitoring," *Journal of Quality Technology*, 46, 78-94.

Yan, H., Paynabar, K., Pacella, M., (2019) "Point Cloud Data Analysis for Process Modeling and Optimization", *Technometrics*, 61, pp. 385-395.

Yan, H., Sergin, N. D., Brenneman, W. A., Lange, S. J., and Ba, S. (2021). Deep multistage multi-task learning for quality prediction of multistage manufacturing systems. *Journal of Quality Technology*, xxxx.

Zhang L.W, Lin J., Liu B., Zhang Z.C., Yan X.H., Wei M.H. (2019) "A review on deep learning applications in prognostics and health management." *IEEE Access*, 7: 162415–38.

Zorriassatine, F., and Tannock, J. D. T. (1998). A review of neural networks for statistical process control. *Journal of intelligent manufacturing*, 9(3), 209-224.