University Of AL-Qadisiyah

Available online at www.qu.edu.iq/journalcm

JOURNAL OF AL-QADISIYAH FOR COMPUTER SCIENCE AND MATHEMATICS

ISSN:2521-3504(online) ISSN:2074-0204(print)

Enhancing AI Through Statistical Methods for Improved Decision-Making

Hadeel Kamil Habeeb [a] ***, Zainb Hassan Radhy*** [b] *

[a] *Faculty of Nursing, University of Al-Qadisiya, Al- Qadisiyah, Iraq,* admin.stat21.11@qu.edu.iq

[b] *College of Computer Science and Information Technology, University of Al- Qadisiyah, Al-Qadisiya, Iraq.* zainb.hassan@qu.edu.iq

ARTICLE INFO

Article history:
Received: 28 /09/2024
Rrevised form: 14 /10/2024
Accepted : 22 /10/2024
Available online: 30 /12/2024

Keywords:

Artificial Intelligence, Machine Learning Algorithms, Bayesian Inference, Regression Analysis, AI-Statistical Hybrid Approaches

ABSTRACT.

This study examines the crucial relationship between Artificial Intelligence (AI) and statistics, which has grown in importance due to the information boom and computational advances. Modern data-intensive applications have led these domains to merge, despite their parallel history. This document explores AI's fundamentals, focusing on machine learning and deep learning, and how statistical methods help AI perform tasks like decision-making and pattern identification. We integrate AI with robust statistical modelling and prediction to improve AI transparency and effectiveness. This multidisciplinary approach emphasizes theoretical advances, practical applications, ethical considerations, and future problems at the interface of AI and statistics. We encourage AI and statistics communities to collaborate to promote innovation and responsible AI development and deployment. This collection of publications is a comprehensive resource for researchers and practitioners using AI and statistics to better decision-making and predictive analytics

MSC..

https://doi.org/ 10.29304/jqcsm.2024.16.41800

Main text

in recent years, the intersection of Artificial Intelligence (AI) and statistics has gained significant attention, driven by the explosion of data and advancements in computational capabilities. This evolving relationship is not merely a convergence of two fields; it represents a fundamental shift in how we approach problem-solving and data analysis in an increasingly complex world. At the heart of AI are concepts such as machine learning and deep learning, which rely heavily on statistical methods. Machine learning algorithms, for instance, use statistical techniques to learn from data, enabling them to make decisions and identify patterns effectively. Deep learning, a subset of machine learning, employs neural networks that mimic human brain functions, further enhancing the capacity for complex data interpretation.

*Corresponding author: Zainb Hassan Radhy

Email addresses: *zainb.hassan@qu.edu.iq*

Communicated by 'sub etitor'

1. Introduction

The history of artificial intelligence (AI) and statistics has been punctuated by long periods of mutual disinterest, but the areas are now converging. (Landgrebe & Smith, 2022) This is being driven by the enormous amounts of data that are now being collected and stored, by the increasing computational power available and by the theoretical and practical successes in building statistical learning methods. The pervasive use of data in almost all scientific disciplines as well as in industry and government has led to an increasing demand for statistical tools and hence for a renaissance in the field of statistics. (Deuse et al.2022) (Sader et al.2022) (Boyd et al., 2023) That the importance of statistics is not lost on the AI community is evidenced by the recent explosion of work on exploring connections between statistics and AI. Related to this is the emergence of a new interdisciplinary field of data mining that has grown out of the realization that the large amounts of data now being collected can yield useful knowledge. (Holmes et al.2022) Although there is debate about whether data mining is a subfield of statistics or a distinct field, the areas are united by a common concern with using automatic methods to analyze data and by the fact that much data mining work uses

statistical methods. (Harbola et al., 2022) Overall, it isn't easy to overemphasize the importance of the present relationship between AI and statistics and the prospects for the future. This is the backdrop against which the present collection of articles should be viewed. (Aggarwal et al.2022) NOMENCLATURE

1.1 Overview of Artificial Intelligence and Statistics

With the advent of computers, the field of statistics has seen enormous gains in its ability to process data and draw inferences from it. (Smolensky et al.2022) This has led to a corresponding increase in the importance of statistical ideas. At the same time, the field of artificial intelligence (AI) has sought to understand and model intelligence by building computational artefacts that can exhibit intelligent behaviour. Initially, this involved emulating human behaviour or automating simple cognitive tasks. (Lin et al.2021) However as AI has matured, there has been a movement towards building tools to automate tasks that would require intelligence if done by a human. As such, there is now increasing overlap between the goals of AI and the goals of statistics, and AI has increasingly turned to statistical inference as a means of achieving its goals. Given the historical prominence of statistical ideas in AI and the rapid increase in the complexity of the data and problems to which they are applied, we feel that it is now more important than ever to consider statistical issues when seeking to advance the goals of AI. (Enholm et al.2022) Yet statistical considerations have not always been first and foremost in AI research. We hope that this work will foster increased communication and collaboration between researchers in AI and statistics to further this important relationship.

1.2. Importance of the Relationship between AI and Statistics

Historically, the field of AI has much to learn from statistics. By considering the algorithm selection problem in the context of statistical decision theory, we can hope to understand better the situations in which different algorithms are appropriate. (Gupta et al.2022) The subfield of machine learning has been the main area of interaction between AI and statistics and has now become so large that some consider it a branch of statistics rather than a distinct field. Early work on machine learning was done largely by computer scientists, who often did not consider the principles of learning from data developed in statistics. (Dhal & Azad, 2022) This situation has led to poorly understood results and a reinvention of the statistical wheel in AI.

At the same time, many statisticians working in academia are not aware of the vast literature in AI on learning algorithms and their associated theories. The result has been a historic disconnect between the AI and statistics communities, with communication between the two communities largely limited to a few researchers. Increased cross-disciplinary research and an exchange of ideas between AI and statistics can only benefit the two fields given the current state of affairs. (Gevaert et al.2021)(Gwagwa et al., 2022)

One could argue that "AI-complete" problems are statistical problems since AI needs to have the ability to learn from data to solve problems. Alternatively, one could view statistical modelling as "soft" AI, seeking to make inferences for specific instances. In this book, we argue that there is much to be gained by considering the relationship between AI and statistics and that we can best understand the development of the two fields in the context of each other. (Kwee-Bintoro & Velez2022) (Groppe & Jain, 2024)

1. 3 Foundations of Artificial Intelligence

Machine learning, particularly supervised learning, has been AI's most successful and widely applied area. Essentially machine learning is to learn to make inferences or decisions from data. The learning is achieved by building a model from input-output pairs representing the desired solution. This model can then be used to make predictions or decisions. Sometimes the model can be a simple sequence of instructions, whilst in the best cases it will map an input to the correct output in the same way as a human would do. (Sarker, 2021)(Burkart & Huber, 2021)The model can have knowledge explicitly programmed into it, in which case machine learning is not necessary. However, a learning-based model can be more robust, accurate, and capable than a hand-designed one. Machine learning is an attractive option for an AI problem solver, but it is not always appropriate and applying it can still be an art. Sometimes it can take much effort to learn something that would have been simple to program in, and the data may inherently include many errors. Also, some learning techniques are so slow or require so much memory that they are not practically implementable. (Cai et al.2024) (Yang et al.2023)

The discipline of AI began with the conjecture that intelligent action can be considered to a large extent as an instance of "rational" decision-making. This conjecture intends to delimit the general area of AI work: The study of systems that act intelligently, while it is always possible to find a way of characterizing the behavior of any system in trivially intelligent terms, AI researchers have often been concerned to distinguish intelligent behaviour from unintelligent behaviour, that simulates an intelligent thinker rather than a random one. (Hoffmann, 2022)(Seo et al.2021) Taking this on board, the next question we should ask is "How can we automate rational decision-making?" And the very best way to answer this question is to build an automated decision maker, that is, an agent. The agent can be anything that we can define in some meaningful environment. For now, this will traditionally be a computer program, but this definition would also include robots at the state of technology of the near future. Here we assume that the robot's behavior is derived from its program and its program is derived from our work, therefore it is meaningful to talk about an intelligent thinking robot and its actions can in theory be characterizable as rational or irrational.

What then is an automated decision-maker? One that always selects the best action given what it knows. (Zhan et al.2022) This does not mean that it will always be the case that the best action is taken, rather the decision maker will not select an action that it does not think is the best one to take. Still further, to say selecting the best action really means selecting the best of a finite number of possible actions. This already sounds like something that can be implemented into a computer program. (Li & Liu, 2024)

2. Machine Learning Algorithms

Machine learning algorithms provide how a system (or agent) can adaptively construct a model of the environment and hence improve its performance. In general, the learning will be directed towards the improvement of the agent's performance on some tasks.

In the simplest case, this is just a prediction of the value of some input. This might be divided into the prediction of a class label (a binary label, or a categorization into one of several classes) or the prediction of a quantity (perhaps a real number, or a full probability distribution over some range of values). In the simplest case, the only feedback available will be the correct answer to the task; we call this supervised learning.

Later, we shall discuss semi-supervised learning, where the learner has a large amount of input for which the correct answer is not available, and reinforcement learning, where the only feedback is through the utility of the task. Unsupervised learning can be understood as a set of techniques for extracting a function to describe and understand the structure of the input.

This may be the identification of various forms of similarity and dissimilarity or the creation of a good representation of the data in a lower dimensional space. In the long term, unsupervised learning might be directed towards the development of a system that can take certain actions in an unknown environment to maximize its future rewards, using its learned understanding of the environment to simulate the action and predict the subsequent state and reward.

The types of machine learning can also be distinguished by the form of the improvement to the system's performance. For example, an agent that has availed of a single opportunity to take an action in a known state in order to assign a value to that state would like to improve its performance via that value with the least number of revisions. This is an example of learning in a static functional approximation context where the least number of revisions might be determined in a revised regression problem. An agent that has decided on a sequence of actions as well as a model for the state sequence those actions will lead to and has discovered via simulated action that on any of several identified occasions the action sequence yields too low a reward will require a more efficient editing and relearning of the sequence. This will be an example of a dynamic function approximation problem where the sequence of actions amounts to a function from the times to the states. In each case, the learning agent improves performance most efficiently by a method which is an optimization of the method used to solve the original task

2.1. Deep Learning Techniques

Deep learning is motivated by moving closer to what we believe is the ultimate learning algorithm,[1][2] one that can learn to perform any task from data. This is described in a paper by LeCun et al. (2015), where the authors ask "What is the machine learning paradigm of the future? ... Is it an algorithm that discovers a process to carry out some computation, even if that process is too complex to be formalized? Is it an unsupervised method that discovers a complex representation or abstraction in data, with much less human direction?" This is reminiscent of the quest for the Holy Grail; to discover the philosopher's stone, to uncover the mechanism behind a cause, or to learn the essence of a concept.

In pattern recognition, an interdisciplinary subject that utilizes techniques from statistics, probability, machine learning, and computer science to extract information from data, we are ultimately interested in learning a compact and exact representation of the data for decision-making and analysis. This frequently involves the automatic learning of features from the data. The idea of feature learning is not new, it is a staple topic in cognitive science and neuroscience, for which machine learning is an attempt at computational replication. Automatic learning of features is an unsupervised learning task where one seeks to learn a function that can map input x from some input space into a feature vector f(x) in a feature space, which makes the task of some supervised learning simpler.

An example of a feature might be the 'presence of a human face', where the input is an image. Deep learning and learning of features with a much greater depth is a present a good promise for the future of artificial intelligence, in the quest to replicate human cognition.

Artificial intelligence has attracted much interest in recent years. Deep learning, a theory of machine learning, can be viewed as an architecture for many different machine learning algorithms. We have developed a widespread, but somewhat superficial, understanding of the field in terms of its goals. In particular, we have focused on supervised learning involving the assignment of semantic meaning to input data. However, deep learning extends to unsupervised learning for the discovery of intricate structures in data and semi-supervised learning, where the algorithm learns a category from a partially labelled training set. Each of these domains of learning can benefit from an increase in the ability to automatically learn features from the data, for which deep learning is well suited. In this paper, we provide an overview of deep learning and the challenges it presents for the future.

2.2. Natural Language Processing

Natural Language Processing (NLP): the ability of a computer to understand and generate human language, has been a challenging but elusive goal for AI. Since the digital computer was invented, there has been a belief that one day we will be able to interact with computers using natural language. The first example of NLP was when scientists at Bell Labs in the early 1950s used an early mainframe to write a very basic Russian-to-English translation program.

As AI increased in popularity through the 1970s and 1980s, so too did work on NLP. However, much to the disappointment of funding bodies, interest in NLP waned through the 1990s as it became apparent that progress was slow. Of the many reasons for the disappointingly slow progress with NLP, one is the inherent complexity of human language. This means that simple, easy-to-use NLP tools are often linguistically naïve and are undermined by the complexity of the task they are trying to perform. Furthermore, humans are often very poor informants about the intricacies of the language they speak. This means that there is often a foundational lack of understanding about the very thing that NLP tools are trying to model.

3. Statistical Methods in Artificial Intelligence

The relation of Bayesian inference to AI is clear, given that the notion of search has been paramount in AI. It seems that the possibility of formulating a model as a probability distribution and then making probability statements about the model and data could give a unifying theory for the different types of searches and the use of statistics to guide that search.

But although the Bayesian approach to statistics is the most logical and consistent form of inference, it is not yet widely used in AI and statistics has had little impact on AI model building.

The reason for this is that the practical application of Bayesian inference is mathematically and computationally demanding. It was only the advent of modern computing power that has made Bayesian methods feasible in mainstream statistics and interest in Bayesian statistics has only grown in the last 20 years. Given that AI has been heavily dependent on the use of computers for all of its history, it is possible that Bayesian methods could yet permeate AI model building.

Statisticians have long maintained that all learning, in particular the formation of models, can be usefully thought of as a process of inference; that is the development of a probably true explanation for a set of data. This explanation will naturally be imperfect, as a variety of different data sets could lead to quite different models. The model-building process can then be thought of as a search for the best model given a data set with some criterion to judge "best". This is exactly what has been described as the optimization problem of search.

The suggestion is then that since the inference process and the optimization process are linked by the criterion, it is better to study model building within the framework of probability and decision theory. This is the tenet of the Bayesian approach. The most formal and general statement of a model comes in terms of a probability distribution. A Bayesian model consists of a prior probability distribution for the model and an explicit specification of how the model can be changed. Bayesian inference is the process of inferring about the model from the data. This is done by the formation of the given model into a probability statement for the data.

Summing over all possible models for a complex of data gives a probability distribution for the data, a marginal likelihood. Using the prior and the data specific likelihood it is possible to now obtain an optimal model with respect to some criterion, typically by finding the most probable model, that is the model with the highest posterior density.

An enduring debate in artificial intelligence research has been the relative importance of symbolic (logical) methods compared to those that are not symbolic. This issue is analogous to the debate between rule-based (deductive) and model-based (inductive) systems. Unfortunately, the terms "logical" and "symbolic" are often loosely defined and used in a variety of ways, and these distinctions can often become clouded. Instead, we will consider the nature of the statistical methods as they apply to AI, for statistics forms the basis of inductive model building. Although statistical methods are not the only inductive methods, they are the only ones with firm theoretical foundations. We will illustrate that even certain methods usually thought of as non-statistical can be reformulated to work within a statistical framework.

3.1. Bayesian Inference

We are typically interested in finding the most probable hypothesis in AI given the data. For example, to build a model or to make predictions. This is known as maximum a posteriori (MAP) estimation and can be done by maximizing the value of $P(h|D)$. We could also use the whole probability distribution to compare different models and to quantify our confidence in any inference drawn from the data. This method is called Bayesian model comparison and can be done using the probability of the data given the model:

$$P(D|h) = \int P(D|hi)P(hi)dh \quad (1)$$

Now the marginal probability P(D) may be calculated from the prior and updated probability distributions on the data:

$$P(D) = \sum P(D|hi)P(hi) \tag{2}$$

This has the effect of maximizing the probability estimate on the data given the hypothesis.

The prior probability can be updated to a posterior probability using the formula:

$$P(h|D) = P(D|h)P(h)/P(D) \tag{3}$$

so that the true probability of the hypothesis can be realized given all the evidence, by marginalizing over all values of the hypothesis using:

$$P(h|D) = \sum P(hi|D) \tag{4}$$

where hi : are all the possible values of h. This is a simple application of probability using the fact that the total probability of a set of disjoint events is the sum of the probabilities of the individual events.

We have a probability distribution P(h) on a hypothesis h. - There is some data D. - There is a likelihood function $P(D|h)$ the probability of the data given the hypothesis. The prior probability distribution and the likelihood function together determine the posterior probability distribution on the hypothesis.

Bayesian inference is a method of statistical inference in which Bayes' theorem is used to update the probability estimate for a hypothesis as more evidence or information becomes available. The essential ingredients are:

3.2. Regression Analysis

Future developments in regression analysis are expected to be in the area of nonparametric or semiparametric regression methods. These methods do not require the data to fit a specified equation and are often more flexible than current regression methods.

In multiple regression, there are several independent variables and each independent variable has a separate value of a slope. The model for multiple linear regression is:

$$Y = a + b1X1 + b2X2 + \ldots + bk + e \tag{5}$$

Where: b1 is the slope of x1 and so on. If the independent variables are not categorized and the number of interactions between variables is large, then it may be difficult to specify the model. Stepwise regression is a technique often used to build up the best model. This model will be the one for which there is no other model that has as few predictors whose value of the sum of squared errors is smaller.

Simple linear regression is a technique that predicts a dependent variable using a single independent variable. The model for simple linear regression is:

$$Y = a + bX + e \tag{6}$$

where Y is the dependent variable, X is the independent variable, a is the constant (the value of Y when X = 0), b is the slope of the line and e is the error term. The least squares technique can be used to calculate the values of a and b that provide the best-fitting line through the sample data. The line of best fit for the sample data can then be used for making predictions about the population.

Regression analysis builds on data that has already been collected. It tests a theory derived from a theoretical model that relates one or more independent variables to Y, so that predictions can be made about the value of Y based on knowledge of the independent variables.

3.3. Hypothesis Testing

The null hypothesis is the keystone of the frequentist approach to statistical testing. This approach is used to decide which of the possible competing explanations of a set of data are the most plausible. The null hypothesis is a precise formal statement that there is no effect. The purpose of the significance test is to assess the strength of evidence against the null hypothesis.

A given result is said to be statistically significant if it is very unlikely to occur when the null hypothesis is true. If the data are consistent with the null hypothesis, the test result is said to be nonsignificant. A statistically significant result provides evidence against the null hypothesis.

It is often interpreted as indicating that a real effect has been detected. Unfortunately, significance testing is one of the most misunderstood areas in statistics. Many researchers believe that a nonsignificant result means that the null hypothesis has been accepted, but this is a serious misinterpretation. The only direct way to accept a null hypothesis is to test the alternative hypothesis and find that it is insignificant.

4. Integration of AI and Statistics

In many cases, however, it is exactly this type of scenario where AI is currently being misused. Instead of building models as a tool to understand and test a specific hypothesis, it has become common to throw data at an algorithm and use the result as a general correlation or prediction without assessing its validity or its consequences[3][4].

When the stakes are high, for example in medical diagnostics, then this is much more likely to harm than good. Step 1 of the CRISP-DM data mining process is exactly what applied statisticians should still be doing; the difference is that AI has powerful methods for steps 2-4. Considering that there are situations where AI methods are appropriate and others where they are not, it should not simply be a case of AI replacing statistics. It is important to have an integrated approach where the most appropriate methodology can be chosen.

The pervasiveness and rapid progress of AI has led to scenarios where mechanistic models, which provide little direct insight, are being used as the basis of decisions with important consequences. Buckley et al. (2007) take an industrial perspective and describe the development of a chip yield optimization system, where a causal model was built to explain systematic spatial variation in yield. To build this model, it was necessary to first characterize spatial variation using an unsupervised learning method that did not assume a specific form for the spatial trend. This was conducted using a mixture model-based cluster analysis, where the number of clusters and spatial trend within each cluster was defined using a Hidden Markov Random Field. It was then possible to make inferences about the causes of low yield using a supervised learning method. In each case, the model was a quite complex AI-type model, but the use of the AI was justifiable because the system provided a unified framework for causal analysis and a specific diagnosis with an associated prediction.

4.1. *AI for Statistical Modeling and Prediction*

Automating statistics through AI methods can potentially have deep and far-reaching impacts on statistics theory. Many statistical methods and their underlying theories will likely need to be reassessed regarding their relevance to an analysis that can be conducted using intelligent computer software.

On a more practical note, it will become essential for future statisticians to be proficient in the new methods. This raises important questions about how and what statistics should be taught to students in the future.

Artificial intelligence has the potential to make significant contributions to the various processes of statistical modeling. The most evident way is through automation. Nearly all forms of statistical modelling involve a search for a "best" model according to some criterion. This involves a great deal of trial and error in manipulating the form of the model, and the use of various model selection criteria to determine the best model. All of this can be done more efficiently using search algorithms implemented in intelligent software. Another key component of statistical modelling is the interpretation of results. This can be a complex process involving the assimilation of evidence from many sources to make an inference about a particular parameter. Expert systems have the potential to automate much of this process. Simulation is an area of statistical practice which is close in spirit to AI. Use of AI methods to automate simulation model building and analysis is an area with much promise[5].

4.2. Statistical Techniques for AI Model Evaluation

In order to assess whether the model provides a good fit to the data, a test of goodness-of-fit can be employed. A well-known method is to compare the predicted outcome to the actual outcome for a random subsample of the data. If the model fits the data, then the expected outcome will be close to the actual outcome. However, it is unrealistic to expect a very close match, and so it must be judged how close is close enough. This method has given rise to prediction-specific measures of fit. Measures of fit compare the predicted outcome to the actual outcome for the whole sample. A commonly used measure of fit for binary outcome is the R2, which is calculated as the square of the correlation between the predictions and the actual results. Measures of fit are often used in conjunction with cross-validation to compare the performances of different models[6][4][7].

Another interesting validation technique for models is validating the model by bootstrapping the samples. Bootstrapping is an approach which involves the technique of sampling with replacement from the current observations. In this method, an additional dataset is generated by random selection from the original dataset, the same size as the original dataset. Data points which appear in the second dataset can be amongst those in the first dataset; a data point from the original dataset may appear multiple times in the new dataset. For instance, it may not appear at all.

Statistical evaluation is a critical aspect of the development and application of predictive models in AI, especially with the growing use of data mining methods by data scientists.

A classic use of a randomized split sample for validation of models has been supplanted in many cases by techniques that make more efficient use of the data, in particular cross-validation. Various flavor's of cross-validation can be used. For instance, in k-fold cross-validation, the data is randomly partitioned into k equal-sized subsamples. Of the k subsamples, a single subsample is retained as the validation data for testing the model, and the remaining k-1 subsamples are used to train the model. Cross-validation is then repeated k times, with each of the k subsamples used exactly once as the validation data. The k results can then be averaged to produce a single estimation.

4.3. AI and Statistical Ethics

AI is fundamentally about building an agent that can learn and make decisions. A decision is essentially choosing the best course of action based on the currently available information. In a typical research context, a decision is made to experiment to collect more information. The action is taken if the researcher believes that the experiment would provide valuable information with a consideration of the risk involved. It is important to make a very clear distinction between the AI agent and a human researcher, in which the former must make decisions by a well-defined mathematical theory and the latter may act based on intuition and experience. Here is an opportunity to use decision theory and causal inference to guide optimal data gathering and analysis for more principled and interpretable results.

AI and Statistical Jun could be represented soon. This could be represented by if we build an engine to analyze the data and give us the result, we want without the need to manipulate the data manually, or if the machine learning model has become so intelligent that it changes our information based on the result it wants rather than the information we have. These situations suggest that the younger AI would autonomously bias the data or result. This corresponds to an issue in medical research where the data analysis has already been guided by the conclusion of the study. The impact could lead to a waste of resources when the AI decides that certain data is not useful and discards them. This situation is not much different from an intelligent human researcher, hence there is a need to develop an ethical code to guide data analysis and AI in their decision-making (Davidian, 2010)[8].

5. Challenges and Future Directions

New developments in decision theory may well be a spin-off from AI research, which will also have an impact on the statistical research and potentially a positive one. All these are learning curves for modern statisticians who may be a specialist in some area of traditional statistics and who may need to evaluate the utility of AI methods for todays and future generations.

If AI methods are to be assessed in terms of their decision-making effectiveness, then it will be necessary for some of them to be compared with alternative methods using intervention experiments because assessment of the potential outcomes is a causal problem and causal inferences are made from comparative assessments of what would happen under different decision-making procedures. This points to the need to develop new methods for assessing the uncertainty in AI-based decision procedures, taking into account both the stochastic and the causal nature of the inferences that are involved.

This poses a challenge in itself for assessing the impact of AI on the statistics discipline, but the biggest challenge is to ensure that the work results in improvements in making the decisions, as opposed to just providing more complex and opaque methods. This is the fundamental issue of assessing the value of AI methods in science and society, which is a quintessentially statistical problem, and it is also an issue of assessing the opportunity costs of using AI methods as compared to the simpler and less automatic methods that are often said to be rival decision procedures.

In the last decade, the focus of the statistical community from model parameter estimation to complex decision-making process has resulted in the development of new statistical methods for AI, known as statistical learning or data mining, fueled largely by the explosion of complex data from the computer age. These developments blur the boundary between statistical AI methods and more traditional statistical methods, making it more difficult to determine what is AI and what is non-AI, and what statistical methods should properly be used in different scientific and societal decision-making problems involving data.

Next section of the book discusses the challenges and the future directions of AI and Statistics. A wide range of statistical methods are used in diverse areas of applications. The ever-growing development in AI technologies would lead to a number of opportunities to tackle the scientific questions of interest using more flexible and powerful statistical models, but it is also likely to lead to several challenges and problems.

5.1. Interpretability and Explain Ability of AI Models

This is an area that requires significant future attention to realize the potential of AI in statistics. The ability to understand, trust, and act on AI-driven inferences and decisions will often require methods that are not only essentially predictive but provide diagnostics and easily understandable justifications.

In a data analysis setting, this might mean a complex AI model providing an answer that could be verified or checked in some way, but more importantly, could be explained in simple terms relying on the most important aspects of the data. Given the rapid progress in machine learning methodology, the best models for a given task might be so complex as to be deemed untrustworthy by end users. Contrary to belief in a pure prediction approach: "If we can predict accurately, why do we need to explain?", such a situation would imply limiting the use of the best predictive methods. If the AI model is a tool to aid human decision-making, there will always be a requirement to understand why the tool is providing certain recommendations. A complex black-box medical decision support system suggesting treatment for a patient will not be trusted or implemented unless it can explain the recommendation. A high-profile example of this was the refusal of the European Parliament to consider the use of AI models for credit scoring since the reasons for a credit rating would be demanded as a legal right by the person being scored[9][10][11].

5.2. Ethical Considerations in AI-Driven Statistical Decision Making

We would like to suggest that the definition of a prescriptive model be any model that suggests an action to change the current state of affairs to a specific desired state. This broad definition ranges from classical statistical models with an assigned treatment variable to Q-learning-type models from machine learning. Below, we discuss some specific areas of prescriptive decision-making and suggest guidelines for best practice.

In the case of learning from historical data, a predictive model can be built to identify the action that would have the best outcome. Often, this will involve data mining to build a model to minimize some measure of error, without explicitly specifying the action to be chosen. In this case, it may be difficult to relate the model to the proposed action, and the model may be used to make inferences about the action in an unplanned way. The prospect of building models to directly prescribe the best action brings with it a host of issues. A medical decision support system is a classic example, and there have been many debates about the safety of using such systems.

Ethical considerations in AI-driven statistical decision making: Judging the morality of decisions is a complex task deeply rooted in philosophical and psychological analysis. Ethical considerations are pervasive throughout all stages of statistical inference and model building, from data collection to the final decisions that are made, However, the increasing role of AI in automating decision-making raises specific issues and possible hazards, especially when AI learns or advises based on previous data. This becomes a case of prescriptive decision-making concerning what actions should be taken. The distinction between descriptive and prescriptive decision-making is often overlooked. The former is decision-making based on what will happen in the future, to achieve some desired state of affairs. The latter aims to take the action that will have the best outcome compared to other possible actions.

5.3. Advancements in AI-Statistical Hybrid Approaches

At the high-dimensional end of the model spectrum, we may be making predictions using data on thousands of patients from just a few hundred genetic markers. If we are not willing to make very restrictive assumptions, most traditional statistical methods are not suitable because the data dimension exceeds the sample size. In such a setting, it is tempting to throw caution to the wind and just mine the data using a machine-learning method. An alternative compromise is to use a dimension reduction technique to summarize the genetic markers into a smaller set of variables that can then be analyzed using a traditional regression method. However, with the widespread use of machine learning and data visualization techniques, it is now possible to do this in a much more user-friendly way than in the past. High-dimensional methods and their use with data arising from modern technology are likely to remain one of the most active areas of methodological and applied research in statistics for some time.

Statistical modelling has a long-standing tradition in scientific research. Indeed, the use of S or R for data analysis and the explicit articulation of a model with a clear statement of the assumptions is what many statisticians regard as good statistical practice. This culminated in John Chambers' call, over a decade ago, for statistical computing environments to be next-generation software for statistical data analysis. He anticipated a level of automation in data analysis not provided by the current generation of statistical software. While it is arguable whether this has been achieved, there is little doubt that the automation of data analysis and the integration of machine learning and statistics, sometimes to the detriment of the latter, has been one of the major success stories of modern statistics.

Reference

[1] A. Mathew and A. Arul, "Deep Learning Techniques : An Overview Deep Learning Techniques : An Overview," no. August 2020, 2021, doi: 10.1007/978-981-15-3383-9.

[2] N. Burkart, "A Survey on the Explainability of Supervised Machine Learning A Survey on the Explainability of Supervised Machine Learning," no. January, 2021, doi: 10.1613/jair.1.12228.

[3] S. M. Liao, "Ethics of artificial intelligence," *Ethics Artif. Intell.*, no. April, pp. 1–525, 2020, doi: 10.1093/oso/9780190905033.001.0001.

[4] G. Phillips-wren, "Intelligent Decision Support Systems," no. February 2013, 2018, doi: 10.1002/9781118522516.ch2.

[5] K. Aggarwal *et al.*, "Has the Future Started? The Current Growth of Artificial Intelligence, Machine Learning, and Deep Learning," *Iraqi J. Comput. Sci. Math.*, vol. 3, no. 1, pp. 115–123, 2022, doi: 10.52866/ijcsm.2022.01.01.013.

[6] A. M. Dias, A. M. Carvalho, and P. Sampaio, "Quality 4.0: literature review analysis, definition and impacts of the digital transformation process on quality," *Int. J. Qual. Reliab. Manag.*, vol. 39, no. 6, pp. 1312–1335, 2022, doi: 10.1108/IJQRM-07-2021-0247.

[7] P. Smolensky, R. T. Mccoy, R. Fernandez, M. Goldrick, and J. Gao, "Neurocompositional computing : From the Central Paradox of Cognition to a new generation of AI systems arXiv : 2205 . 01128v1 [cs . AI] 2 May 2022," pp. 1–21.

[8] P. Wang, "On Defining Artificial Intelligence," vol. 10, no. 2, pp. 1–37, 2019, doi: 10.2478/jagi-2019-0002.

[9] J. Deuse, N. West, and M. Syberg, "Rediscovering Scientific Management - The Evolution from Industrial Engineering to Industrial Data Science," *Int. J. Prod. Manag. Eng.*, vol. 10, no. 1, pp. 1–12, 2022, doi: 10.4995/ijpme.2022.16617.

[10] S. Groppe, "Renewal of the Major Fields of New - Generation Computing," *New Gener. Comput.*, no. 0123456789, pp. 3–4, 2023, doi: 10.1007/s00354-023-00206-5.

[11] K. Seo, J. Tang, I. Roll, S. Fels, and D. Yoon, "The impact of artificial intelligence on learner – instructor interaction in online learning," *Int. J. Educ. Technol. High. Educ.*, 2021, doi: 10.1186/s41239-021-00292-9.