Artificial Intelligence is the New Astrology of Software Quality

Jack McDowell

Ying Ki Kwong

jackmcdowell@gmail.com

Abstract

Artificial Intelligence / Machine Learning are catchall terms, which often involve neural nets which are "trained" using training data and other types of input to then respond to future input. A well-trained AI can "accurately" respond to stimuli in ways that meet business needs, cutting down on human intervention or complex rule-based coding. However, just as modern science is based on fundamental first principles, software applications are supposed to be driven by clearly defined business rules and attendant business logic. Yet, AI frequently amounts to statistical correlation between "black box" models and training data with unknown scientific, legal, moral, or ethical assumptions. Therefore, concerns about biased and insufficient training data are only part of the problem.

When considering implementing AI or AI assisted technologies, management must understand the business risks of such a move, despite the seemingly unlimited enthusiasm and optimism of the technologists. The authors believe these considerations are especially relevant in the quality assurance of mission critical applications in large enterprises and government organizations.

Biography

Jack McDowell is the Statewide QA Program Manager with the State of Oregon. The Statewide QA Program provides Quality Assurance services for Oregon's Major IT Projects, and Quality Assurance consultation to Oregon State Agencies. Before this, he was a web developer and the chief editor of a community newspaper in Arlington, Virginia. Originally from Buenos Aires, Argentina where he lived before attending college in the US. He holds a master's degree in political science from the University of Oregon and a certification in ITIL.

Ying Ki Kwong is the E-Government Program Manager with the State of Oregon. Prior he was the IT Investment Oversight Coordinator in the same office and was Project Office Manager of the Medicaid Management Information System Project in the Oregon Department of Human Services. In the private sector, Dr. Kwong was CEO of a Hong Kong-based internet B2B portal and a program manager in the Video & Networking Division of Tektronix. In these roles, he has managed software-based systems/applications, products, and business process improvements. He received the doctorate from the School of Applied & Engineering Physics at Cornell University and was adjunct faculty in the School of Business Administration at Portland State University. He holds the PMP certification since 2003.

1 Introduction

Artificial Intelligence (AI) seems to be everywhere these days. The popular media often talk about AI-based tests, data processing, and even decision making in complex situations. AI can perform tasks as simple as turning on a light through voice recognition to trigger “If This Then That (IFTTT)” logic, or as complex as diagnosing medical conditions through IBM’s Watson Health (O'Leary 2022). The fact that an AI-based computing system can defeat the world champion of the board game Go is a sort of testament that AI advances in recent years have rivaled certain aspects of human intelligence (Mozur 2017).

Yet, everyone who has interacted with a voice recognition AI may have experienced a moment when an instruction, such as “turn on living room lights”, resulted in a nonsensical response. Likewise, when interacting with AI driven chatbots, frustration frequently ensues after ending up in a redirect loop. For the most part, these non-critical situations are innocuous and simply cause the user some frustrations or inconvenience. However, what happens when the business needs involve social benefits determinations (Carney 2021), shortening prison time of inmates (Rieland 2018), activating safety features of heavy machinery, or whether a self-driving car should collide with an elderly person, a child, or a tree?

AI and machine learning are often associated with “learning” neural networks trained over time to respond to different inputs. In an ideal world, a well-trained AI can “accurately” respond to stimuli in ways that meet business needs, cutting down on human intervention or complex rule-based coding. However, just as modern science is based on first principles, modern software applications are supposed to be driven by clearly defined business rules and attendant business logic that lend themselves to formal verification & validation. Yet, AI frequently amounts to statistical correlation between models and training data – however carefully done and well intentioned. At best, this statistical correlation may entail the non-existence of sound *a priori* first principles – scientific, legal, moral, ethical, or otherwise. At worst, statistical correlation may be based on biased, inaccurate, or even false premises that are embodied in models and training data that reflect real-word biases that are unscientific, illegal, immoral, or unethical.

This paper will explore how AI differs from traditional modeling and development, how AI works in practice, and the potential for inherent inaccuracy of AI Algorithms. We will discuss management considerations in addressing typical issues that may arise when applying AI techniques. These considerations include re-training of models to remove biases, better understanding of the limits of probabilistic “black box” algorithms, and potential decisions to not use AI.

2 Artificial Intelligence Software: How does it differ from Functional Software?

At its most basic level, software programs take inputs, do something with that data and produce an output. Figure 1.1 illustrates what this most basic workflow looks like. An example could be something as simple as switching a relay at a given time, where input = time, do something = compare input to value, and output = activate relay.

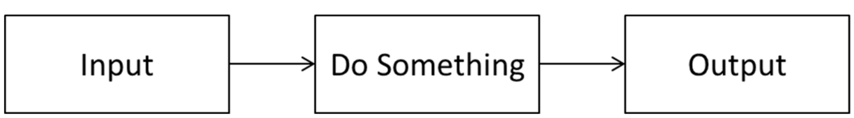

Figure 1.1: Basic Program Flow

Inputs can be user input data, data stores, big data, etc., and outputs can be returning data to a user, saving to a data store, etc.

The “Do Something” is traditionally composed of a set of operations or functions that are generated to meet a certain need. These functions can be as simple or as complex as needed to fulfill a certain task. As depicted in Figure 1.2, a simple program to determine whether x is greater than y could look as follows in C++:

```
#include <iostream>

using namespace std;

 int main() {

   int x = 5;

   int y = 2;

   if (x > y) {

      cout << "x is greater than y." << endl;

   } else {
```

Figure 1.2: Program

Which would produce the output as shown in Figure 1.3:

```
x is greater than y.

Program ended with exit code: 0
```

Figure 1.3: Output

The code above is based on the "greater than" operator >. Its underlying logic is *a priori* explicit, with no *a posteriori* (statistical) variability for acceptable input.

Whether we are thinking of procedural, functional, or object-oriented programming, specifications are translated into a finite set of operations to be conducted. These set of operations typically derive from a set of functional specifications and can number from a few to thousands of functions depending on a software's complexity. However, no matter the level of complexity, when software is developed, a person or group of people convert specifications into code.

AI programs invert this development paradigm. Instead of translating requirements into functions, AI development entails feeding training data, comprised of Inputs AND Outputs into a neural network, so that the AI can learn from the data and generate metaphorical functions.

When creating AI generated algorithms, instead of writing lines of code with explicit logic like that above, a developer may "train" a neural network using a training data set. Roughly speaking, a neural network can be thought of as layers of interconnected nodes of processing elements (neurons) as depicted in Figure 1.4, where the connection strength between interconnected nodes are adjustable parameters of the model (NeuralDesigner 2022).

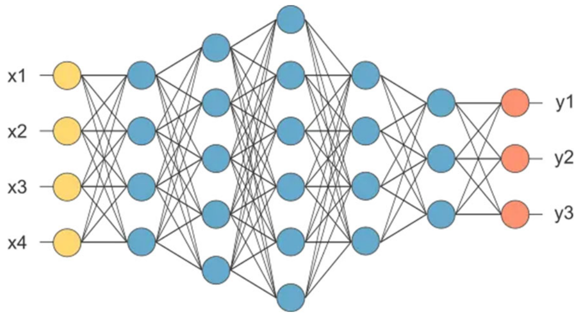

Figure 1.4: Neural Network Nodes

The process of training is to use real-world data (training data) to optimize the values of these adjustable neural network parameters. In a typical training process, a formal loss index (a measure of the different types of model errors) is minimized by iteratively adjusting neural network parameters. Iterations in the direction of lower loss index (the training direction) is done by applying an optimization algorithm (e.g., gradient descent, the Newton method, or other techniques used to numerically calculate local minimum of functions of many variables). This process of optimizing neural network parameters is as depicted in Figure 1.5 (Quesada 2022):

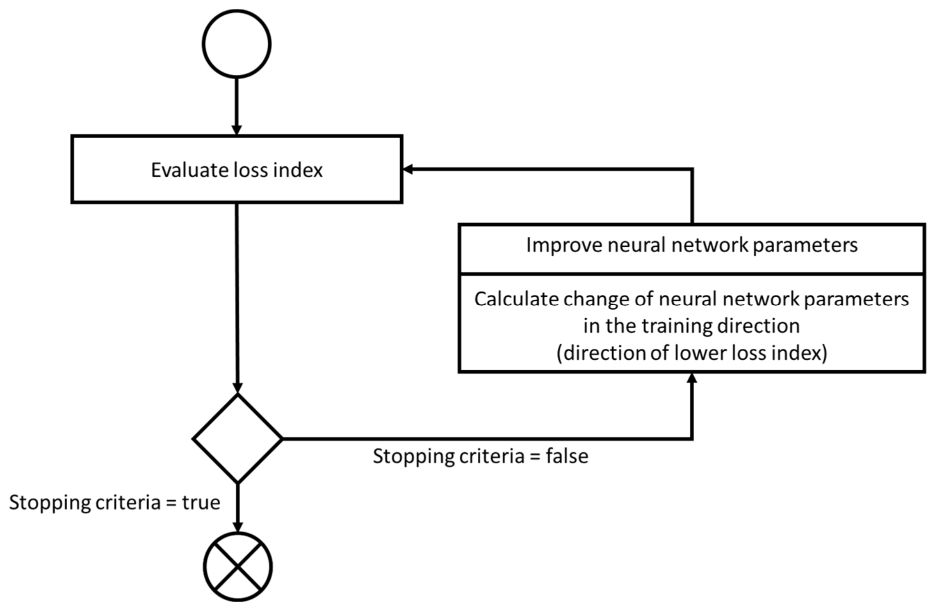

Figure 1.5: Generic Optimization Procedure for Neural Network Parameters that quantify the connection strength between neuron nodes

Generally speaking, optimization algorithms used to train neural networks that have greater computational speeds tend to require more memory and higher floating point calculation precision.

Note that the training of a neural network amounts to calculating neural network parameters that "best fit" the training data. Thus, trained neural networks have characteristics that are inherently quite different from computer codes that are based on explicit business logic. Specifically, trained neural networks contain inherent *a posteriori* (statistical) uncertainty, and they may exhibit acceptable / predictable variability only when new data to be analyzed by the trained neural network are sufficiently "like" the training data used to train the neural network in the first place.

As a model with many adjustable parameters, trained neural networks are complex emergent systems with underlying systematics that may be difficult to discern. Even with careful analysis, it is difficult to formally demonstrate or verify that a trained neural network processes information consistent with physical laws, acceptable engineering principles or practices, or in accordance with generally acceptable legal, social, or community standards. This, in a way, is the fundamental challenge of processing information using “black box” business logic that may not be explicitly traceable to specific mathematical relations, rules, decision trees, and other pre-vetted criteria.

2.1 Comparing AI in Software Systems to Traditional Development

The process of designing and building software begins with a business need or opportunity, which is then translated into specifications, which in turn result in a software product to fulfill those needs. The final product can then be tested to validate that it meets business needs.

Developing traditional applications and AI applications begins with identifying business needs, but the approach begins to differ after identifying business needs.

Typically, business needs are translated into functional and non-functional specifications, which in turn form models that perform logical or mathematical operations on data. In a way, every software process can be deconstructed to its model and specification (see Figure 2.1). In turn, software can be tested to ensure that models are acting as expected and, thus, specific requirements can be verified.

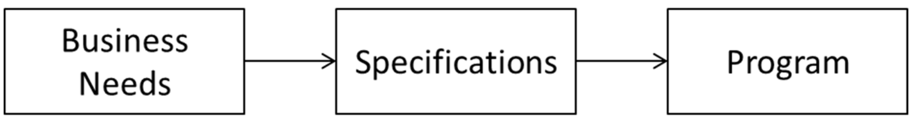

Figure 2.1. Traditional Software Development

AI software changes this paradigm by optimizing algorithms with adjustable parameters against training data based on business needs. In neural networks, this optimization process is called training; in which the element parameter (connection strengths) between processing elements (neural network nodes) are adjusted to minimize some measure of discrepancy between algorithm output and data (see Figure 2.2). As more data are collected from the real world, training is repeated to improve the quality of the algorithm. Because this process may be computationally intensive, the size of the training data sets, or frequency of re-training may be constrained in practice.

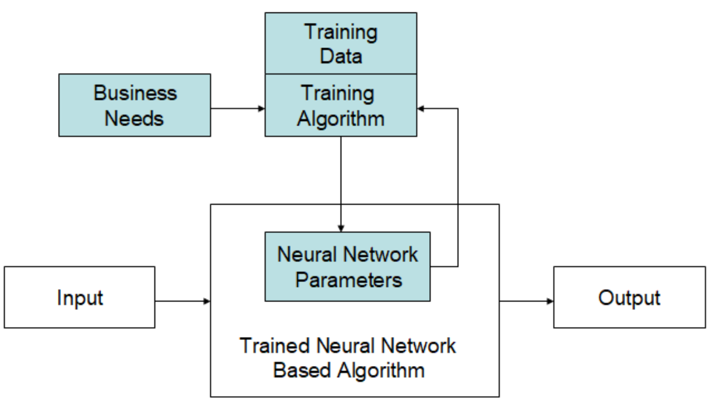

Figure 2.2: AI Training Process

3 Artificial Intelligence in Practice

3.1 AI algorithms in practice

AI and machine learning are all around us, from chatbots, to voice and phrase recognition (Amazon's Echo, Apple's Siri, Google), to cars and even SpaceX rockets and capsules (Patel 2020). Although these examples seem disparate, these applications all have the fundamental basics in common, they take an input (text, voice, telemetry data) and produce an output (answer a question, play a song, calculate a rocket's trajectory).

We may think that of the above examples, calculating a rocket's trajectory is the most complex for us to calculate. After all, if the reader were asked to play a certain song one could likely complete the task with ease; yet being given reams of telemetry data and being asked to compute a parabolic trajectory would be beyond the scope of most readers. And yet, the miscommunications we have with our smart assistants would render their use as flight computers useless. What explains the difference between these acceptable failures of AI, successes of AI in advanced computations, and when are the potential failures of AI too risky to outweigh their usefulness?

The following sections will discuss considerations regarding ensuring the quality of AI models, where AI can work well, and certain situations where the risk of utilizing AI may outweigh the benefits.

3.2 Ensuring the quality of AI models

Certain challenges in using AI driven models can be distilled by verifying and validating an application. Typically, verification and validation occur against specifications and business requirements, respectively. To conduct verification on an application, we would use manual and automated tests to ensure that specifications are met, followed by debugging of code to rectify issues. Likewise, to validate that software meets business needs, we would conduct user acceptance testing to evaluate the software, and if necessary, produce new requirements and modify the application to meet those needs.

Similarly, an AI model would be put to the test to determine if an algorithm meets business needs. However, as Gao, et al. (Gao 2019) discuss, there are unique challenges in ensuring the quality assurance of AI software. Key challenges include:

- limited data training and validation. AI algorithms are only validated with limited input data under ad-hoc contexts.
- data-driven learning features, static or dynamic, that negatively affect software outcomes, results, and actions.
- inconsistent system outputs, responses, or actions due to uncertainty inherent in statistical models.

Much of the literature on AI verification has been focused on biases introduced by incomplete or biased training data (Lee 2022). For example, research has shown that AI facial recognition algorithms trained with predominantly white faces tend to be better at identifying white faces than people of color (Najibi 2020 and Hardesty 2022). These algorithmic fails have led to contentious public discussions, such as when the IRS decided to require, and then provide alternatives, to the use of ID.me facial recognition to access tax records (Butler 2022).

Issues related to poor or biased training data can be solved by improving data sets or adding weights to datasets. However, issues related to uncertainty inherent in statistical models are harder to resolve. Whenever we think of a traditional software tool, it is possible to explore code and debug software, but when an AI is computing data, it is impossible to pinpoint what is causing an unexpected outcome or bias.

3.3 Where AI can work well

Scientific problems which can be falsified are perhaps the best ones suited to AI. This is because the AI can be given working parameters where out of bounds results are clear, and problems are repeatable experimentally. As mentioned previously, Machine Learning models are now able to compute complex rocket trajectories or drive cars autonomously. These algorithms are generally successful provided that the inputs are predictable (to match training data). When algorithms encounter data that does not fit with their existing training model, they make probabilistic assumptions to provide an answer and may not respond as expected, as in the case of autonomous vehicles confusing an overturned semi-trailer with a clear right of way (Stumpf 2020). These types of issues may be improved upon with better training data or lower thresholds for identification.

An important question to consider when deploying these systems depends on accountability for the decisions that the system made. Automations are part of many vehicles' safety features, such as airbags and automatic headlights. These features are rigorously tested to ensure that at a specific threshold, features are activated as expected. With AI automation, testing can still be conducted to ensure that safety features respond as expected in the event they are needed. Many of the so-called failures that we have seen regarding AI automation are not due to the AI itself, but due to lack of AI training and testing.

3.4 Where AI does not work well but can be improved

Most of the criticisms that we hear about AI failures are related to poor data sets or poor models. For example, facial recognition software has been found to poorly discern people of color, resulting in incorrect identifications (Najibi 2020). This has resulted in allegations that AI is biased against people of color, due to the disproportionate impacts AI has had on certain groups (Hardesty 2022). However, while the models themselves may be biased due to biased training data, the concept of utilizing AI for tasks such as facial recognition is not always inherently biased. With additional training data and better / more comprehensive data sets, these biases can be mitigated and removed.

> Because machines can treat similarly situated people and objects differently, research is starting to reveal some troubling examples in which the reality of algorithmic decision-making falls short of our expectations. Given this, some algorithms run the risk of replicating and even amplifying human biases, particularly those affecting protected groups. For example, automated risk assessments used by U.S. judges to determine bail and sentencing limits can generate incorrect conclusions, resulting in large cumulative effects on certain groups, like longer prison sentences or higher bails imposed on people of color (Lee 2022).

There is also the question of risk that needs to be addressed. For example, if AI does not recognize someone's face to unlock their phone on the first attempt, it may not be a big issue if one can try again or dial an emergency call via a fallback mechanism. However, if the same facial recognition fails to provide access to a time sensitive or mission critical service or system, the risk would be exponentially higher. "Surfacing and responding to algorithmic bias upfront can potentially avert harmful impacts to users and heavy liabilities against the operators and creators of algorithms, including computer programmers, government, and industry leaders" (Lee 2022).

4 Inherently inaccurate AI Algorithms

So far, we have discussed where algorithms can work well, and where algorithms can work well if biases and training data is comprehensive. However, area where AI may face the insurmountable limitations is when algorithms cannot be improved with additional training data due to limits of knowledge (unknown information, unknown unknowns), and where probabilistic algorithms are legally or morally problematic. The problem with AI Algorithms, is that the AI will still provide a result based on its training, and identifying an issue, and then uncovering why the result may be incorrect, can be an impossible task.

4.1 What is an inaccurate model?

Artificial neural networks are often empirical models that rely on observation rather than theory. For that reason, they only need to be consistent most of the time with empirical observations. Since empirical models are models that are not generated from scientific theories, but are generated from observations, they may not be supported by theories at all in a scientific sense.

We say that AI models are the new astrology because, much like astrology, AI models look at swaths of data and make predictions. In astrology, models are created based on correlation, not theory based causal mechanisms. According to astrology when a planet transits a house in someone's birth chart, the effects of that planet are felt by the person. For example, "Venus transiting through the 4th house brings us new comforts that make navigating difficult times smoother" (Janssens 2022). When these predictions are made, there is no causal mechanism or theoretical framework that explains why the planet Venus has brought comforts to someone navigating difficult times; rather, Astrologists claim to have observed many cases where people have found comfort as Venus was transiting their 4th house. The falsifiability of an astrological model, in this case the effects of Venus in the 4th house becomes increasingly complicated due to the additional variables (planets and houses) that may impact a person's transits in a moment in time. Therefore, while we can observe that many people with Venus transiting their 4th house may find comfort in difficult times (as may the general population), the model is not falsified by people who do not find comfort, since we can easily explain that away through additional variables, such as perhaps a Mars transit.

Scientific theories, on the other hand, are often derivable from first principles such as physical laws or accepted principles. With this said, some scientific theories are phenomenological in that they describe the empirical relationships of phenomena in a way consistent and do not contradict accepted physical laws or principles. When a model contradicts accepted physical laws and principles, we can say with good certainty that the model is not good or at least outside the regime of its validity. This is an important distinction: a model that failed to make good predictions under certain circumstances may make very good predictions within its regime of validity.

As discussed previously when AI works well, models that deal with physical laws and engineering principles are the simplest problems for AI to solve, when comprehensive training data and robust neural networks can interpret predictable data. The task of evaluating models becomes increasingly complex as we transit into the social sphere and ask AI to work with applicable laws or statutes of a specific geography, or jurisdiction, social, religious, or cultural values, and protocol and etiquette. While it is straightforward to identify a model that contradicts the laws of physics, it is much more complex to identify a model that produces outputs that contradict certain cultural values.

4.2 Limits of knowledge

Issues arise with AI when it clashes with limits of knowledge, such as when the training data does not exist or when an AI is asked to solve a new problem. For example, an AI can become an expert level GO player by being taught the game's rules and being trained with thousands of completed games. However, if an AI is tasked with the very basic task of walking, absent knowledge about walking, the results tend to be somewhat comical (TechInsider 2017). When Deep Mind taught itself to walk, it is important to remember that according to the algorithm, it was successful in completing its task of locomotion. To the observer, Deep Mind's solution appears comical because we know, from our a priori knowledge of what it is to walk, that one does not walk by running around with their arms flailing. Similarly, IBM tried to use Watson to assist with cancer screenings by training it with imaging data and found that the accuracy of Watson's predictions was no better than chance (O'Leary 2022).

4.3 Legal and moral issues with AI modeling

Inherent to probabilistic models, is the legal question as to whether an AI algorithm's conclusion is sufficient to pass legal muster, and whether its findings would hold up in court. For example, states such as Pennsylvania and Oregon used AI algorithms to assist in the determination of child placement in foster care (Burke 2022). The degree to which these determinations can be considered reliable, and these determinations can be defended is problematic in multiple ways:

1. How can we trust that the subject of an AI algorithm's decision fits within the confidence interval and probability of the model, instead of being an outlier?
2. How can we unpack the algorithm's determination, given the black box nature of the neural network?
3. Can we trust, as a society, an AI algorithm to determine such unique decisions as determining whether a child should be removed from a household and placed in foster care?

These are not merely academic musings. The U.S. Federal Government has recently published guidance related to the use of potentially biased algorithms (Jillson 2022), as well as properly trained algorithms generated from training data that may have been inappropriately obtained (Kaye 2021).

5 Conclusion

This paper considered the limitations of AI in software applications. The authors believe these considerations are especially relevant in the quality assurance of mission critical AI-based applications that support large enterprises and government organizations.

We contrasted traditional software design vs. AI-based software design as two different development paradigms:

- Traditional software design is based on algorithms whose underlying logic is *a priori* pre-determined. As such, the resultant logic can be unpacked and analyzed formally to support software verification & validation.
- AI-based software design is based on algorithm whose underlying "logic" is derived from *a posteriori* (statistical) training procedures that can be repeated in time when more training data become available. As such, the resultant logic cannot be easily unpacked and analyzed formally to support verification & validation beyond statistical correlations.

As a researcher in the Advanced Research in Cyber Systems group at Los Alamos National Labs explained, "The artificial intelligence research community doesn't necessarily have a complete understanding of what neural networks are doing; they give us good results, but we don't know how or why" (Njegomir 2022). We believe this situation has important management implications.

When considering implementing AI or AI assisted technologies, management must understand the business risks of such a move, to provide checks & balance on the seemingly unlimited enthusiasm and optimism of the technologists. As discussed in this paper, typical issues that arise with AI may include but are not limited to:

- Models may not be good or are poorly trained.
- Models may be well trained but poorly understood.
- Inherent biases are embodied in the training data.
- Lack of accountability with accuracy of models and when models should / should not work.
- Lack of formal understanding about the models beyond *a posteriori* statistical correlation.

Although astrology may provide entertainment through plausible correlations, we would not use astrology to determine social benefits or activate safety features of heavy machinery. However, we should not base business logic on plausible correlations. When we use AI in real world applications, we have the obligation to understand the regime of validity of our models and avoid underlying logic based on plausible correlations.

The authors are especially concerned about the use of AI in situations that may affect the rights, dignity, safety, and physical / mental health of individuals or minority groups. For applications in these spaces, we believe developers have the responsibilities to assure fairness and ethical development of their AI-based software. For example, if an AI-based facial recognition software is believed to be highly accurate among the general population of a country, the developers should be aware that this is a statistical statement, and this sort of statistical statement may be consistent with the same software being highly inaccurate within an ethnic minority of the same country or possibly highly inaccurate among the general population of another country (Hardesty 2022).

We would like to close this paper with the following quotes from the Brookings Institution and the United Nations:

> It is important for algorithm operators and developers to always be asking themselves: Will we leave some groups of people worse off as a result of the algorithm's design or its unintended consequences? – *Brookings Institution* (Lee 2022)

> The world needs rules for artificial intelligence to benefit humanity. The Recommendation on the ethics of AI is a major answer. It sets the first global normative framework while giving States the responsibility to apply it at their level. UNESCO will support its 193 Member States in its implementation and ask them to report regularly on their progress and practices. – *UNESCO* (Azoulay 2022)

Acknowledgement

The authors would like to acknowledge helpful discussions with Karen J. Johnson, Nicholas Betsacon, Brooke Janssens and John Cvetko during the planning of this paper.

References

Azoulay, Audrey, and Gabriela Ramos. "UNESCO Member States Adopt the First Ever Global Agreement on the Ethics of Artificial Intelligence." UNESCO, February 1, 2022. https://en.unesco.org/news/unesco-member-states-adopt-first-ever-global-agreement-ethics-artificial-intelligence.

Burke, Sally Ho and Garance. "Oregon Dropping AI Tool Used in Child Abuse Cases." AP NEWS. Associated Press, June 2, 2022. https://apnews.com/article/politics-technology-pennsylvania-child-abuse-1ea160dc5c2c203fdab456e3c2d97930.

Butler, Peter. "How to Register on the IRS Website without Taking a Video Selfie." CNET. Accessed September 30, 2022. https://www.cnet.com/personal-finance/taxes/irs-to-back-off-third-party-facial-recognition-what-happens-to-id-me/.

Carney AO, Terry, Artificial Intelligence in Welfare: Striking the Vulnerability Balance? (March 15, 2021). Carney, T., 'Artificial Intelligence in Welfare: Striking the vulnerability balance? (2020) 46(2) Monash University Law Review Advance 1-29 DOI https://doi.org/10.26180/13370369.v1, Available at SSRN: https://ssrn.com/abstract=3805329 or http://dx.doi.org/10.2139/ssrn.3805329

Gao, Jerry, Chuanqi Tao, Dou Jie, and Shengqiang Lu. "Invited Paper: What Is Ai Software Testing? and Why." *2019 IEEE International Conference on Service-Oriented System Engineering (SOSE)*, 2019. https://doi.org/10.1109/sose.2019.00015.

Hardesty, Larry. "Study Finds Gender and Skin-Type Bias in Commercial Artificial-Intelligence Systems." MIT News | Massachusetts Institute of Technology. Accessed September 30, 2022. http://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212.

Janssens, Brooke. “Venus Retrograde.” Star Oracles, February 1, 2022. https://staroracles.com/venus-retrograde/.

Jillson, Elisa, Helen Clark, and Lesley Fair. “Aiming for Truth, Fairness, and Equity in Your Company's Use of Ai.” Federal Trade Commission, May 24, 2022. https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai.

Kaye, Kate. “Why the FTC Is Forcing Tech Firms to Kill Their Algorithms along with Ill-Gotten Data.” Digiday, November 17, 2021. https://digiday.com/media/why-the-ftc-is-forcing-tech-firms-to-kill-their-algorithms-along-with-ill-gotten-data/.

Lee, Nicol Turner, Paul Resnick, and Genie Barton. “Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms.” Brookings. Brookings, March 9, 2022. https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/.

Mozur, Paul. “Google's AlphaGo Defeats Chinese Go Master in Win for A.I.” The New York Times. The New York Times, May 23, 2017. https://www.nytimes.com/2017/05/23/business/google-deepmind-alphago-go-champion-defeat.html.

Najibi, Alex. “Racial Discrimination in Face Recognition Technology.” Science in the News, October 26, 2020. https://sitn.hms.harvard.edu/flash/2020/racial-discrimination-in-face-recognition-technology.

NeuralDesigner. “3. Neural Network.” Neural networks tutorial: Neural network | Neural Designer. Accessed September 30, 2022. https://www.neuraldesigner.com/learning/tutorials/neural-network.

Njegomir, Nick. “New Method for Comparing Neural Networks Exposes How Artificial Intelligence Works.” Discover Los Alamos National Laboratory. Accessed September 30, 2022. https://discover.lanl.gov/news/0913-neural-networks.

O'Leary, Lizzie. “How IBM's Watson Went from the Future of Health Care to Sold off for Parts.” Slate Magazine. Slate, January 31, 2022. https://slate.com/technology/2022/01/ibm-watson-health-failure-artificial-intelligence.html.

Patel, Neel V. “Are We Making Spacecraft Too Autonomous?” MIT Technology Review. MIT Technology Review, July 3, 2020. https://www.technologyreview.com/2020/07/03/1004788/spacecraft-spacefight-autonomous-software-ai/.

Quesada, Alberto. 5 algorithms to train a neural network. Accessed September 30, 2022. https://www.neuraldesigner.com/blog/5_algorithms_to_train_a_neural_network.

Rieland, Randy. “Artificial Intelligence Is Now Used to Predict Crime. but Is It Biased?” Smithsonian.com. Smithsonian Institution, March 5, 2018. https://www.smithsonianmag.com/innovation/artificial-intelligence-is-now-used-predict-crime-is-it-biased-180968337/.

Stumpf, Rob. “Autopilot Blamed for Tesla's Crash into Overturned Truck.” The Drive, June 16, 2020. https://www.thedrive.com/news/33789/autopilot-blamed-for-teslas-crash-into-overturned-truck.

TechInsider. *Google's DeepMind AI Just Taught Itself To Walk*. *YouTube*. YouTube, 2017. https://www.youtube.com/watch?v=gn4nRCC9TwQ.