

Editorial

# AI and the Singularity: A Fallacy or a Great Opportunity?

Adriana Braga<sup>1</sup> and Robert K. Logan<sup>2,\*</sup> 

<sup>1</sup> Department of Social Communication, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Rio de Janeiro 22451-900, Brazil; adrianabraga@puc-rio.br

<sup>2</sup> Department of Physics, University of Toronto, Toronto, ON M5S 1A7, Canada

\* Correspondence: logan@physics.utoronto.ca; Tel.: +1-416-361-5928

Received: 18 February 2019; Accepted: 18 February 2019; Published: 21 February 2019



**Abstract:** We address the question of whether AI, and in particular the Singularity—the notion that AI-based computers can exceed human intelligence—is a fallacy or a great opportunity. We have invited a group of scholars to address this question, whose positions on the Singularity range from advocates to skeptics. No conclusion can be reached as the development of artificial intelligence is still in its infancy, and there is much wishful thinking and imagination in this issue rather than trustworthy data. The reader will find a cogent summary of the issues faced by researchers who are working to develop the field of artificial intelligence and in particular artificial general intelligence. The only conclusion that can be reached is that there exists a variety of well-argued positions as to where AI research is headed.

**Keywords:** artificial intelligence; artificial general intelligence; Singularity; cognition; emotion; computers; information; meaning; intuition; wisdom

## 1. Introduction

We made a call for papers that either support or criticize the lead paper for this Special Issue entitled *The Emperor of Strong AI Has no Clothes: Limits to Artificial Intelligence*. In this lead paper, we argued that the premise of the technological Singularity, based on the notion that computers will one day be smarter than their human creators, is false, and made use of the techniques of media ecology. We also analyzed the comments of other critics of the Singularity, as well as of those supportive of this notion. The notion of intelligence that advocates for the technological Singularity does not take into account the full dimension of human intelligence. They treat artificial intelligence as a figure without a ground. Human intelligence, as we will show, is not based solely on logical operations and computation, but also includes a long list of other characteristics, unique to humans, which is the ground that supporters of the Singularity ignore. The list includes curiosity, imagination, intuition, emotions, passion, desires, pleasure, aesthetics, joy, purpose, objectives, goals, telos, values, morality, experience, wisdom, judgment, and even humor. We asked that the contributors to this special issue either support or critique the thesis we developed in the lead paper. We received contributions concerning both sides of the argument, and have therefore put together an interesting collection of viewpoints that explores the pros and cons of the notion of the Singularity. The importance of this collection of essays is that it explores both the challenges and the opportunities of artificial general intelligence in order to clarify an important matter that has been shadowed by ideological wishful thinking, biased by marketing issues, and influenced by creative imagination about how the future would be, rather than supported by trustworthy data and grounded research.

In this sense, we want to share an article by Daniel Tukelang entitled *10 Things Everyone Should Know About Machine Learning* [1] that he has kindly given us permission to reproduce, in order to give the reader some background on the issues surrounding the notion of the Singularity.

Daniel Tukelang wrote, “As someone who often finds himself explaining machine learning to nonexperts, I offer the following list as a public service announcement”:

1. Machine learning means learning from data; AI is a buzzword. Machine learning lives up to the hype, and there is an incredible number of problems that you can solve by providing the right training data to the right learning algorithms. Call it AI if that helps you sell it, but know that AI, at least as it is used outside of academia, is often a buzzword that can mean whatever people want it to mean.
2. Machine learning is about data and algorithms, but mostly data. There is a lot of excitement about advances in machine learning algorithms, and particularly about deep learning. However, data is the key ingredient that makes machine learning possible. You can have machine learning without sophisticated algorithms, but not without good data.
3. Unless you have a lot of data, you should stick to simple models. Machine learning trains a model from patterns in your data, exploring a space of possible models defined by parameters. If your parameter space is too big, you will overfit to your training data and train a model that does not generalize beyond it. A detailed explanation requires more math, but as a rule, you should keep your models as simple as possible.
4. Machine learning can only be as good as the data you use to train it. The phrase “garbage in, garbage out” predates machine learning, but it aptly characterizes a key limitation of machine learning. Machine learning can only discover patterns that are present in your training data. For supervised machine learning tasks like classification, you will need a robust collection of correctly labeled, richly featured training data.
5. Machine learning only works if your training data is representative. Just as a fund prospectus warns that “past performance is no guarantee of future results”, machine learning should warn that it is only guaranteed to work for data generated by the same distribution that generated its training data. Be vigilant of skews between training data and production data, and retrain your models frequently, so they do not become stale.
6. Most of the hard work for machine learning is data transformation. From reading the hype about new machine learning techniques, you might think that machine learning is mostly about selecting and tuning algorithms. The reality is more prosaic: most of your time and effort goes into data cleansing and feature engineering—that is, transforming raw features into features that better represent the signal in your data.
7. Deep learning is a revolutionary advance, but it is not a magic bullet. Deep learning has earned its hype by delivering advances across a broad range of machine learning application areas. Moreover, deep learning automates some of the work traditionally performed through feature engineering, especially for image and video data. But deep learning is not a silver bullet. You cannot just use it out of the box, and you will still need to invest significant effort in data cleansing and transformation.
8. Machine learning systems are highly vulnerable to operator error. With apologies to the NRA, “Machine learning algorithms don’t kill people; people kill people.” When machine learning systems fail, it is rarely because of problems with the machine learning algorithm. More likely, you have introduced human error into the training data, creating bias or some other systematic error. Always be skeptical, and approach machine learning with the discipline you apply to software engineering.
9. Machine learning can inadvertently create a self-fulfilling prophecy. In many applications of machine learning, the decisions you make today affect the training data you collect tomorrow. Once your machine learning system embeds biases into its model, it can continue generating new

training data that reinforce those biases. And some biases can ruin people's lives. Be responsible: do not create self-fulfilling prophecies.

10. AI is not going to become self-aware, rise up, and destroy humanity. A surprising number of people seem to be getting their ideas about artificial intelligence from science fiction movies. We should be inspired by science fiction, but not so credulous that we mistake it for reality. There are enough real and present dangers to worry about, from consciously evil human beings to unconsciously biased machine learning models. So you can stop worrying about SkyNet and "superintelligence".

There is far more to machine learning than I can explain in a top 10 list. But hopefully, this serves as a useful introduction for nonexperts.

## 2. Materials and Methods

As the question of whether or not the Singularity can be achieved can only be determined after many years of research, the essays in this collection are based largely on each author's opinions of human cognition and the progress made in the field of artificial intelligence.

## 3. Results

The basic result of this collection of essays and opinions is an extensive list of the challenges in the development of artificial general intelligence and the possibility of the Singularity—the notion that an artificial general intelligence can be achieved that exceeds human intelligence.

## 4. Discussion

One of the things that readers should keep in mind when reading the articles that we have collected for this Special Issue is that AI research is still in the early stages, and AI as a data processing device is not foolproof.

The recent detection of the gravitational wave from the merger of two neutron stars as a result of scientists not trusting their AI-configured automated data-processing programs underscores the thesis that one cannot rely on artificial intelligence alone. The case study that is reviewed illustrates the point that AI, combined with human intervention, produces the most desirable results, and that AI by itself will never entirely replace human intelligence.

Ethan Siegel, an astrophysicist, science communicator, and NASA columnist, in a recent article entitled *LIGO's Greatest Discovery Almost Didn't Happen* [2], demonstrated that AI-configured computers will never replace human intelligence, but that they are nevertheless important tools that enhance human intelligence. If the scientists had relied solely on the results of their AI-configured automated data-processing programs, they would have missed a critical observation of the production of gravity waves from the merger of two neutron stars—an extremely rare event never before observed.

There are altogether three observatories for detecting gravity waves, with two LIGO (The Laser Interferometer Gravitational-Wave Observatory) detectors located at Hanford Washington and Livingston Louisiana, and the EGO (European Gravitational Observatory) detector located near Pisa, Italy. The three detectors were in agreement when they observed the first detected gravitational wave that emanated from the merger of two black holes.

A short time after the detection of the first gravitational wave at all three detectors, a signal was received at the Hanford detector consistent with the merger of two neutron stars. The problem was that no signal was registered at the other two detectors, as should have been the case if a gravitational wave had arrived at our planet according to the automated data-processing program in use at the three detectors. Without the corroborating evidence from the two other detectors, the team at Hanford would have been forced to conclude that the signal was not the detection of a gravitational wave but rather a glitch in the system. However, one of the scientists, Reed Essick, decided that it was worth examining the data from the other detectors to determine if there was a signal that had been missed as

a result of a glitch at these detectors. He went through the painstaking task of examining every signal that might have been received by the Livingston detector around the time of the event detected at Hanford. To his delight, he found that a signal had been registered at the Livingston detector but had been overlooked by the automated computer program because of a glitch at that detector. An analysis of the detector at Pisa revealed that, at the time of the event at Hanford, the Pisa detector was in a blind spot for observing the event of the two neutron stars merging. Corroborating data came from NASA's Fermis satellite, which had detected a "short period gamma-ray burst" that arrived two seconds after the arrival of the gravity wave and was consistent with the merger of two neutron stars. As a result of the due diligence and the perseverance of the LIGO team led by Essick, an astronomically important observation was rescued. In his article, Siegel drew the following conclusion from this episode:

Here's how scientists didn't let it slip away . . . If all we had done was look at the automated signals, we would have gotten just one "single-detector alert," in the Hanford detector, while the other two detectors would have registered no event. We would have thrown it away, all because the orientation was such that there was no significant signal in Virgo, and a glitch caused the Livingston signal to be vetoed. If we left the signal-finding solely to algorithms and theoretical decisions, a 1-in-10,000 coincidence would have stopped us from finding this first-of-its-kind event. But we had scientists on the job: real, live, human scientists, and now we've confidently seen a multimessenger signal, in gravitational waves and electromagnetic light, for the very first time.

Consequently, the conclusion that we reached as a result of Siegel's story and his conclusion is as follows:

1. AI combined with human intervention produces the most desirable results.
2. AI by itself will never entirely replace human intelligence.
3. One cannot rely on AI, no matter how sophisticated, to always get the right answer or reach the correct conclusion.

The variety of opinions presented in this Special Issue will give the readers much to think about vis-à-vis AI, artificial general intelligence, and the Singularity. Readers are invited to correspond to the coeditors with their reactions to, and opinions of, these essays. Thank you for your attention.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Tukelang, D. 10 Things Every Small Business Should Know About Machine Learning. Available online: <https://www.inc.com/quora/10-things-every-small-business-should-know-about-machine-learning.html> (accessed on 29 November 2017).
2. Siegel, E. LIGO's Greatest Discovery Almost Didn't Happen. Available online: <https://medium.com/starts-with-a-bang/ligos-greatest-discovery-almost-didnt-happen-a315e328ca8> (accessed on 24 April 2018).



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).