

THE AI THOUGHT BOOK

Inspirational
Thoughts & Quotes
on
Artificial Intelligence

Murat Durmus

*Quotations are drops of thoughts that
sometimes reach deep into our soul
and give us insights that a hundred
books cannot.*

Contents

Preface - 7

Artificial Intelligence - 8

AI-Ethics - 19

Explainable AI (XAI) - 28

Philosophy - 32

Data & Business - 35

Education & Future of Work - 40

Society & Humanity - 44

Mixed - 49

**Three essays for the fundamental
understanding of AI - 55**

Epilog - 69

Appendix 1: Glossary - 74

**Appendix 2: A brief History of Artificial
Intelligence - 89**

Appendix 3: The criminal Potential of AI - 96

**Appendix 4: Some significant achievements
in the field of AI since 2010 - 109**

ARTIFICIAL INTELLIGENCE


Murat Durmus
(CEO AISOMA)

when you talk about AI, you usually don't mean ..

this



but rather this



Limitations:

- Extreme sensitivity to **adversarial perturbations**
- Extreme sensitivity to **any input change** not seen in the training data
- It can only make sense of **what it has seen before**

I have done a terrible thing. I have demystified Artificial Intelligence.



*Recognizing that two points of data are connected is not enough. The System must ask **why** one point affects another.*


Robustness

The foremost question we should always ask ourselves when we get serious about AI:


“How can we make artificial intelligence systems that are robust in the face of lack of knowledge about the world?”



*Artificial Intelligence is highly **Interdisciplinary**. Therefore, let's approach it in a **Multidisciplinary** & **Holistic** way.*



AI-ETHICS



Discussions about AI Ethics are still mostly conducted in academic circles. But one can already see that many companies are seriously dealing with it. One thing that seems clear to me:

*There is no time to lose. The coming generation will have to deal more with ethical issues than with technological ones. Therefore: **Solving the Tech Industry's Ethics problem must start in the Classroom.***



Ethical guidelines, frameworks, and tool-kits are not enough.

Ethical guidelines, frameworks, and toolkits are not enough; they can only address AI ethical problems selectively but not holistically. As AI becomes more advanced and more widely used, we need to create more social and institutional structures and spaces that encourage, guide, support, and sufficiently reward ethical behavior. These structures would greatly facilitate and motivate people to live and work more ethically.




No Women, no trustworthy AI!
#diversitymatters


DATA & BUSINESS

Companies and authorities that continue to ignore Artificial Intelligence or neglect to integrate it into their organizations and processes will pay a high price for it. They will no longer be competitive. They will drown in Data and Complexity because the world is becoming *more and more Data-Driven*. ~ Murat Durmus

Amazing! So many Insights and Opportunities ...



Still Discussing ...



Help! Help!
Help!



Created by Murat Durmus (CEO AISOMA)



Data is not the new Oil, nor the new Electricity. Data (Information) is everything because we are literally drowning in it and only use a fraction of it so far. Imagine that the entire available treasure of data is lifted and made accessible. The possible applications that are offering exceed my imagination.

Duplication versus Simulation

There is still much confusion about this point in the AI community. With this article, I want to present my view on the relationship between duplication and simulation because it is of great importance that there is clarity here.

The philosopher John Searle has attached great importance to this point by explaining that a simulation is not duplicated. A machine cannot duplicate human thought, but at best, simulate it. On the fact that simulation and duplication are two pairs of boots, I fully agree with him.

Suppose we have two kinds of objects in front of us, say, an Audi A4 (neither my favorite car nor do I drive it) and a second object that someone claims to be a “duplicate” or a “model” of the Audi A4. What exactly does that mean? What is a model of the A4? It means exactly what a ten-year-old who is interested in car models understands by it. Namely, there is a direct correspondence between the external stimuli, internal states, and behavior of the A4 and the inputs, internal states, and outputs of the model. The correspondence does not necessarily have to be one hundred percent. Thus, some external stimuli, states, and behaviors of Model A4 may not be present in the

APPENDIX 2: A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) is a growing discipline of sixty years that encompasses a range of sciences, theories, and techniques (including mathematical logic, statistics, probabilities, computational neurobiology, computer science and philosophy) that aim to mimic the cognitive abilities of humans. Its developments are closely related to those in computer science. They have resulted in computers being able to perform increasingly complex tasks that previously could only be assigned to a human.

However, this automation is still a long way from human intelligence in the strict sense, which has brought the term into criticism among some experts. The final stage of their research (a “strong” AI, i.e., the ability to contextualize very different specialized problems completely autonomously) is not comparable to current achievements (“weak” or “moderate” AI, extremely efficient in its training domain). “Strong” AI, which so far exists only in science fiction, would require advances in basic research (not just perfor-

APPENDIX 3: THE CRIMINAL POTENTIAL OF AI

Note: I did not include this appendix to denigrate AI or to stir up fears. I believe that AI will bring more benefits to humanity than any other technology to date. I want to point out the potential dangers and how it can be abused.

AI can be implicated in crime in several ways. Most obviously, AI could be used as a tool for crime, using its capabilities to facilitate actions against real-world targets: predicting the behavior of people or institutions to discover and exploit vulnerabilities; generating fake content for extortion or to damage reputations; performing acts that human perpetrators cannot or will not perform themselves for reasons of danger, physical size, speed of response, etc. Although the methods are new, the crimes themselves may be traditional in nature – theft, extortion, intimidation, terror.

Alternatively, AI systems themselves may be the target of criminal activity: Circumventing protective systems that stand in the way of a crime; evading detection or prosecution of crimes already committed; causing trusted or critical systems to fail or misbehave cause harm or undermine public trust.

AI could also provide context for a crime. The fraudulent activity could depend on the victim believing that a certain AI functionality is possible when it is not – or that it is possible but not used for the fraud.

Of course, these categories are not mutually exclusive. As in the adage about catching a thief, an AI system attack may itself require an AI system to be carried out. The fraudulent simulation of non-existent AI capabilities could be executed using other AI methods that exist.

Crimes vary enormously. They may be directed against individuals or institutions, businesses or customers, property, government, the social fabric, or public discourse. They may be motivated by financial gain, acquisition of power, or change in status relative to others. They may enhance or damage reputations or relationships, change policy, or sow discord; such effects may be an end in themselves or a stepping stone to a broader goal. They may be committed to mitigate or avoid punishment for other crimes. They may be driven by a desire for revenge or sexual gratification or to further religious or political goals. They may express nothing more than a nihilistic urge to destroy, vandalize, or commit violence for its own sake.

The extent to which AI can amplify this variety of criminal acts depends mostly on how much they are embedded in a computational environment: Robotics is advancing rapidly, but AI is better suited to participate in a bank fraud than in a bar fight. This preference for the digital over the physical world is a weak defense. However, because today's society is deeply dependent on complex computer networks, not only

for finance and commerce but also for all forms of communication, politics, news, work, and social relationships. People now conduct large parts of their lives online, get most of their information there, and their online activities can make or break their reputations. This trend is likely to continue for the foreseeable future. Such an online environment, where data is property and information power, is ideally suited for exploitation by AI-based criminal activities that can have significant real-world consequences. Moreover, unlike many traditional crimes, crimes in the digital domain are often highly reproducible: once developed, techniques can be shared, repeated, and even sold, opening up the potential for commercializing criminal techniques or providing “crime as a service.” This can lead to a lowering of technological barriers as criminals are able to outsource the more challenging aspects of their AI-based crimes.

Listed below are some potential hazards.

Audio and video imitation

People have a strong tendency to believe their own eyes and ears, so audio and video evidence has traditionally been given a lot of credibilities (and often legal force), despite the long history of photo trickery. But recent developments in Deep Learning, mainly using GANs (see above), have greatly expanded the scope for generating fake content. Persuasive impersonations of targets following a fixed script can already be produced, and interactive impersonations are expected to follow. Delegates saw multiple criminal applications for such “deepfake” technologies to exploit people’s implicit trust in these media, including Impersonation of children to elderly parents via vid-

APPENDIX 4: SOME SIGNIFICANT ACHIEVEMENTS IN THE FIELD OF AI SINCE 2010

2010

DeepMind Technologies is founded

A British AI company acquired by Google in 2014 and is part of Alphabet Inc. DeepMind Technologies' most amazing products are the Neural Turing Machine, AlphaFold, Wavenet and WaveRNN, and AlphaGO. In 2014, DeepMind received the "Company of the Year" award from Cambridge Computer Laboratory.



IBM's Watson computer beats human champions on game show Jeopardy

Watson is an interrogative computer system capable of answering questions posed in natural language. It was developed as part of IBM's DeepQA project by a research team led by study director David Ferrucci. Watson was named after IBM's founder and first CEO, industrialist Thomas J. Watson. The computer system was originally developed to answer questions on

the quiz show Jeopardy! In 2011, the Watson computer system competed against champions Brad Rutter and Ken Jennings on Jeopardy! and won the first prize of \$1 million.

2011

The Google Brain Project

The project was first launched in 2011 as a part-time research project by Google employees Jeff Dean, Greg Corrado, and Stanford University professor Andrew Ng. The project first received significant attention in June 2012, when a computer cluster of 16 thousand computers designed to replicate the human brain early recognized a cat based on YouTube images.



Apple introduced SIRI on the iPhone 4s

Apple launched SIRI as the first speech assistance program with the iPhone 4S. Speech Interpretation and Recognition Interface (SIRI) uses voice commands to perform specific user tasks. These voice commands include calling a person, setting the alarm, sending an email, opening text messages, answering questions, asking for recommendations, and using multiple.