

Ethical Dilemmas in AI and Financial Services

Ethical Conondrums with AI

Dr. Philippe J.S. De Brouwer
Honorary Consul of Belgium in Kraków
guest professor at the UJ, AGH, UEK and UW
board member of AGH
SVP at HSBC in Kraków



3rd Annual Global Future of BPO, Shared Services & GBS Conference

2026-04-12

- ▶ Academic
 - ▶ Theoretical physicist and Business Engineer; solved the “fallacy of large numbers” puzzle (Samuelson).
 - ▶ Challenged Markowitz’s Mean Variance Theory, developed Maslowian Portfolio Theory.
- ▶ Career
 - ▶ Fortis/BNP (capital-guaranteed funds, director),
 - ▶ KBC (merged companies, CEO, market growth).
 - ▶ CIO of KBC Ireland (€32B, 1000 funds).
 - ▶ Head of Analytics RBS (financial risk, analytics, data)
 - ▶ SVP at HSBC (Model Risk Management CoE).
- ▶ Author: “The Big R-Book” and numerous scientific papers.



Table of Contents

Defining Ethics

Artificial Intelligence

Example: Insurance Onboarding

Ethics in Data Science

Conclusions

Table of Contents I

Defining Ethics

Artificial Intelligence

Example: Insurance Onboarding

Ethics in Data Science

Conclusions

Defining Ethics

Ethics

the study of what is morally right and wrong, or a set of beliefs about what is morally right and wrong.

Dilemma 1: Kill the Baby or Old Person?




Figure: You can only correct the course of the car a little, so you can reduce the deaths from two to one. Do you choose track A or B? — source: MIT Technology Review — Hao, 2018

Dilemma 1: Who to Kill - take 2




Figure: source: Science — Bonnefon, Shariff, and Rahwan, 2016

Dilemma 1: Who to Kill - take 2




Figure: source: Science — Bonnefon, Shariff, and Rahwan, 2016

Should we veer and kill the person who is on the pavement?

Dilemma 1: Who to Kill - take 2




Figure: source: Science — Bonnefon, Shariff, and Rahwan, 2016

Should we veer and kill the person who is on the pavement? What if there is only one

Dilemma 1: Who to Kill - take 2




Figure: source: Science — Bonnefon, Shariff, and Rahwan, 2016

Should we veer and kill the person who is on the pavement? What if there is only one

The viewpoint of the squirrel



Figure: Is removing a tree that houses a squirrel ethically correct when building your home? What if you take the viewpoint of the squirrel? What if an advanced society that regards us as squirrels wants to make Earth their home?

What is ethically acceptable depends on your viewpoint



Figure: Ethics is complex and it depends on your point of reference and you frame.

Artificial Intelligence

Artificial Intelligence

A cross-disciplinary approach to understanding, modelling, and replicating intelligence and cognitive processes by invoking various computational, mathematical, logical, mechanical, and even biological principles and devices. —Brundage, 2015

or

AI in Britannica.com

the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings

First Some Simple Questions

- ▶ What is exactly Intelligence and how to measure it?

First Some Simple Questions

- ▶ What is exactly Intelligence and how to measure it?
the ability to acquire and apply knowledge and skills (Oxford Dictionary)

First Some Simple Questions

- ▶ What is exactly Intelligence and how to measure it?
the ability to acquire and apply knowledge and skills (Oxford Dictionary)
- ▶ What is exactly consciousness?

First Some Simple Questions

- ▶ What is exactly Intelligence and how to measure it?
the ability to acquire and apply knowledge and skills (Oxford Dictionary)
- ▶ What is exactly consciousness?
the state of being aware of and responsive to one's surroundings - Oxford Dictionary

First Some Simple Questions

- ▶ What is exactly Intelligence and how to measure it?
the ability to acquire and apply knowledge and skills (Oxford Dictionary)
- ▶ What is exactly consciousness?
the state of being aware of and responsive to one's surroundings - Oxford Dictionary
- ▶ What is Self-consciousness? Sentience? Life?

- ▶ What if AI becomes more intelligent than people? Does it get moral status?

Issues with Morality and AI

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?

Issues with Morality and AI

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)

Issues with Morality and AI

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)
- ▶ What when we can “upload” a human brain? Who owns it? Can we switch it off?

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)
- ▶ What when we can “upload” a human brain? Who owns it? Can we switch it off?
- ▶ Does a conscience and sentient AI deserve rights?

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)
- ▶ What when we can “upload” a human brain? Who owns it? Can we switch it off?
- ▶ Does a conscience and sentient AI deserve rights?
- ▶ Who are we to decide this?

Issues with Morality and AI

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)
- ▶ What when we can “upload” a human brain? Who owns it? Can we switch it off?
- ▶ Does a conscience and sentient AI deserve rights?
- ▶ Who are we to decide this?
- ▶ Can one delegate agency to an AI? Who has the moral agency?

- ▶ What if AI becomes more intelligent than people? Does it get moral status?
- ▶ What about the singularity and benevolence of machines?
- ▶ What about enhanced humans (e.g. math implant in brain?)
- ▶ What when we can “upload” a human brain? Who owns it? Can we switch it off?
- ▶ Does a conscience and sentient AI deserve rights?
- ▶ Who are we to decide this?
- ▶ Can one delegate agency to an AI? Who has the moral agency?
- ▶ Should we manipulate data for “good”? (and counter-act bias?)

Example: Insurance Onboarding

The data contains

- ▶ driving related:
 - ▶ average daily travel
 - ▶ professional or private use
 - ▶ living in urban or city area
- ▶ car related:
 - ▶ type of car (eg. minivan, sports car, sedan, etc.)
 - ▶ age of the car
- ▶ behaviour on our books:
 - ▶ time with us
 - ▶ old claims
- ▶ driver related:
 - ▶ has driver's licence been revoked
 - ▶ penalty points on driver's licence
 - ▶ years on job (work related)

The data contains

- ▶ driving related:
 - ▶ average daily travel
 - ▶ professional or private use
 - ▶ living in urban or city area
- ▶ car related:
 - ▶ type of car (eg. minivan, sports car, sedan, etc.)
 - ▶ age of the car
- ▶ behaviour on our books:
 - ▶ time with us
 - ▶ old claims
- ▶ driver related:
 - ▶ has driver's licence been revoked
 - ▶ penalty points on driver's licence
 - ▶ years on job (work related)

... but also contains

- ▶ related to gender and family:
 - ▶ number of kids at home
 - ▶ single parent or not
 - ▶ marital status
 - ▶ gender
- ▶ education and profession
 - ▶ education
 - ▶ occupation
- ▶ wealth related:
 - ▶ value of the car
 - ▶ value of the home

We call these the **“protected features”**

The cost of being politically correct

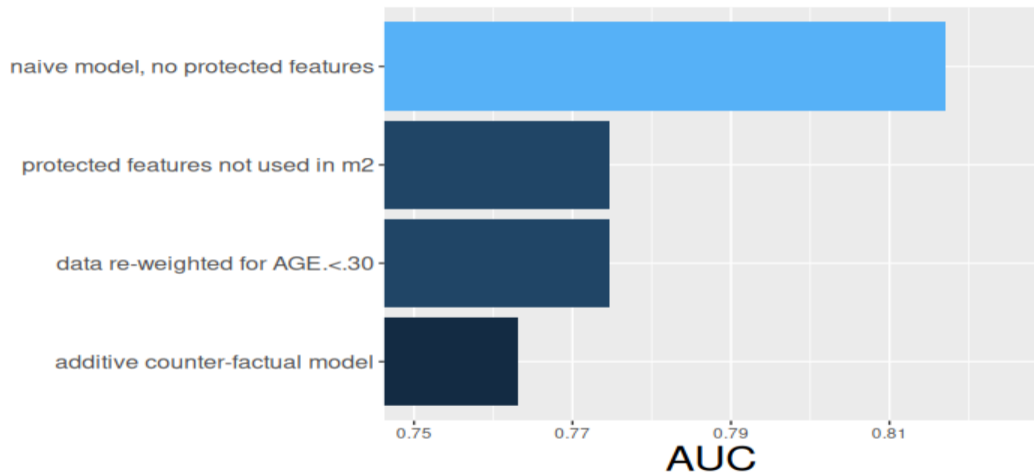




Figure: The discriminating power of different models. The higher the AUC, the stronger the model, the lower the losses on the portfolio, the higher profit (or lower prices) can be.

How to avoid bias? Two alternatives



How to avoid bias? Two alternatives



Ethics in Data Science

- ▶ user consent

Issues

- ▶ user consent
- ▶ privacy, data protection

Issues

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition
- ▶ discrimination:

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition
- ▶ discrimination:
 - ▶ bias


- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition
- ▶ discrimination:
 - ▶ bias
 - ▶ fairness

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition
- ▶ discrimination:
 - ▶ bias
 - ▶ fairness
 - ▶ accountability

- ▶ user consent
- ▶ privacy, data protection
- ▶ intellectual property rights, data ownership
- ▶ safekeeping: avoiding data theft, cyber security
- ▶ trust, surveillance, free will
- ▶ competition
- ▶ discrimination:
 - ▶ bias
 - ▶ fairness
 - ▶ accountability
 - ▶ transparency

Conclusions

Visual Biases are systematic miss-interpretations



Fairness is a matter of perspective




Figure: Fairness is a matter of perspective

Fairness is a matter of perspective



Figure: Fairness is a matter of perspective

Bias is learnt



Figure: Bias is assuming a solution based on what you have learnt before. It can be useful, but usually is limiting.

1. Ethics are the moral principles guiding human behavior.

Conclusions

1. Ethics are the moral principles guiding human behavior.
2. Some ethical intuitions are innate, but...

Conclusions

1. Ethics are the moral principles guiding human behavior.
2. Some ethical intuitions are innate, but...
3. Complex ethical challenges often provoke disagreement; perspective shapes moral judgments.

Conclusions

1. Ethics are the moral principles guiding human behavior.
2. Some ethical intuitions are innate, but...
3. Complex ethical challenges often provoke disagreement; perspective shapes moral judgments.
4. Morality cannot be reduced to fixed rules but requires context-sensitive judgment.

Conclusions

1. Ethics are the moral principles guiding human behavior.
2. Some ethical intuitions are innate, but...
3. Complex ethical challenges often provoke disagreement; perspective shapes moral judgments.
4. Morality cannot be reduced to fixed rules but requires context-sensitive judgment.
5. Data science presents unique ethical dilemmas involving fairness, bias, and accountability.

Conclusions

1. Ethics are the moral principles guiding human behavior.
2. Some ethical intuitions are innate, but...
3. Complex ethical challenges often provoke disagreement; perspective shapes moral judgments.
4. Morality cannot be reduced to fixed rules but requires context-sensitive judgment.
5. Data science presents unique ethical dilemmas involving fairness, bias, and accountability.
6. Ethical decision-making involves navigating trade-offs with no perfect solutions; responsibility matters.