

# Decision-making under uncertainty

Julien Grand-Clément, Assistant Professor

Information Systems and Operations Management Department

HEC Paris

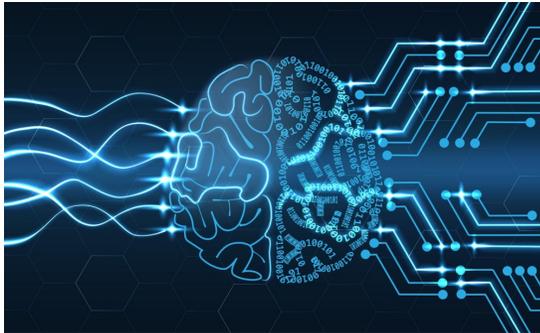
# Plan for this morning

1. Challenges in artificial intelligence
2. The robust approach to uncertainty in decision-making
3. Hands-on application: portfolio optimization

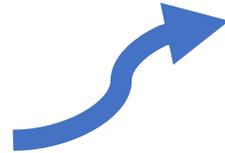
# Forewords

- Since 09/21: Assistant Professor at HEC Paris
- PhD @ Columbia University (NY) & Master @ Ecole Polytechnique

- Research:



Machine learning  
Statistics  
Optimization



Healthcare application  
(automatization/predictions)

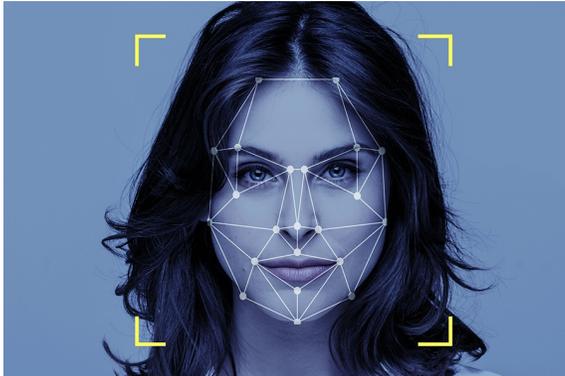
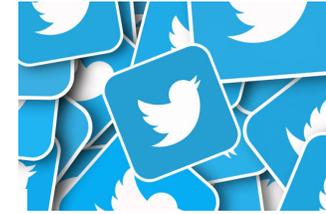


Optimizing orders  
@ Amazon

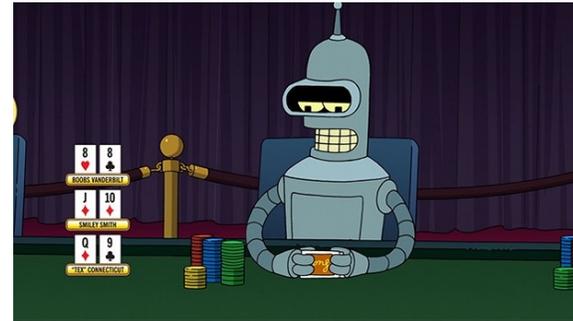


Game theory  
(Poker, Go, etc.)

Data and algorithms are everywhere these days...



Facial recognition



Poker & Go AI



Social media



designed by freepik.com

Autonomous driving

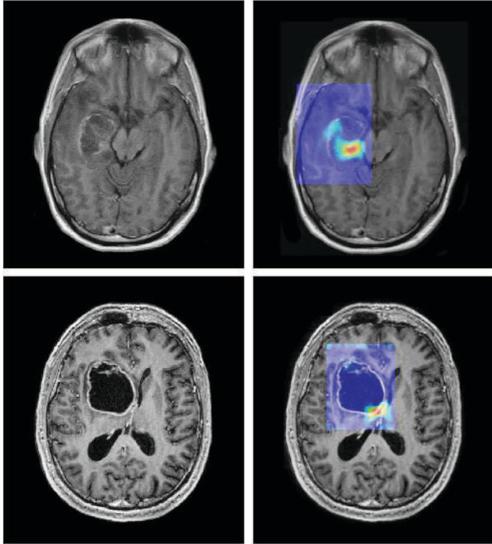


Music

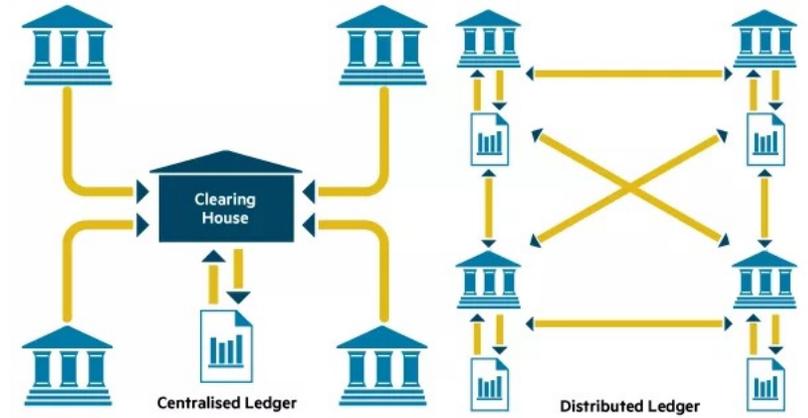


... and there is more to come!

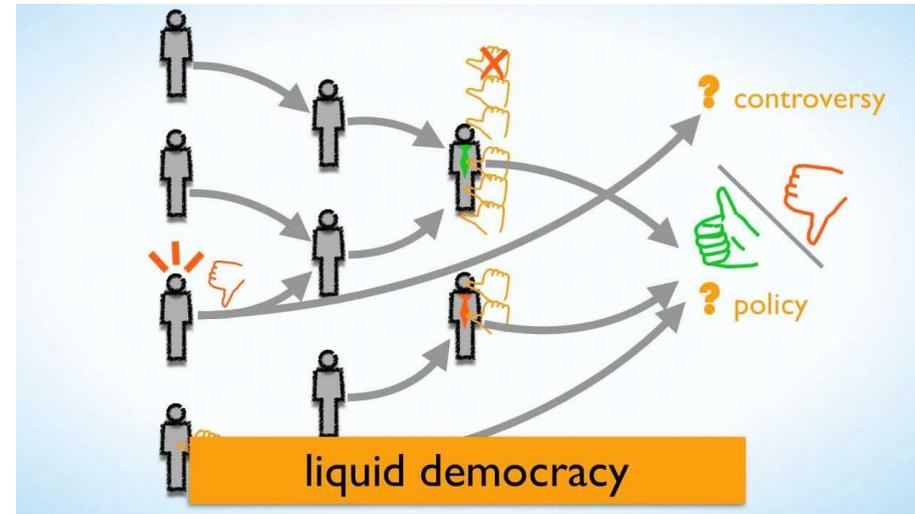
IDH1 mutant glioblastoma



AI in healthcare

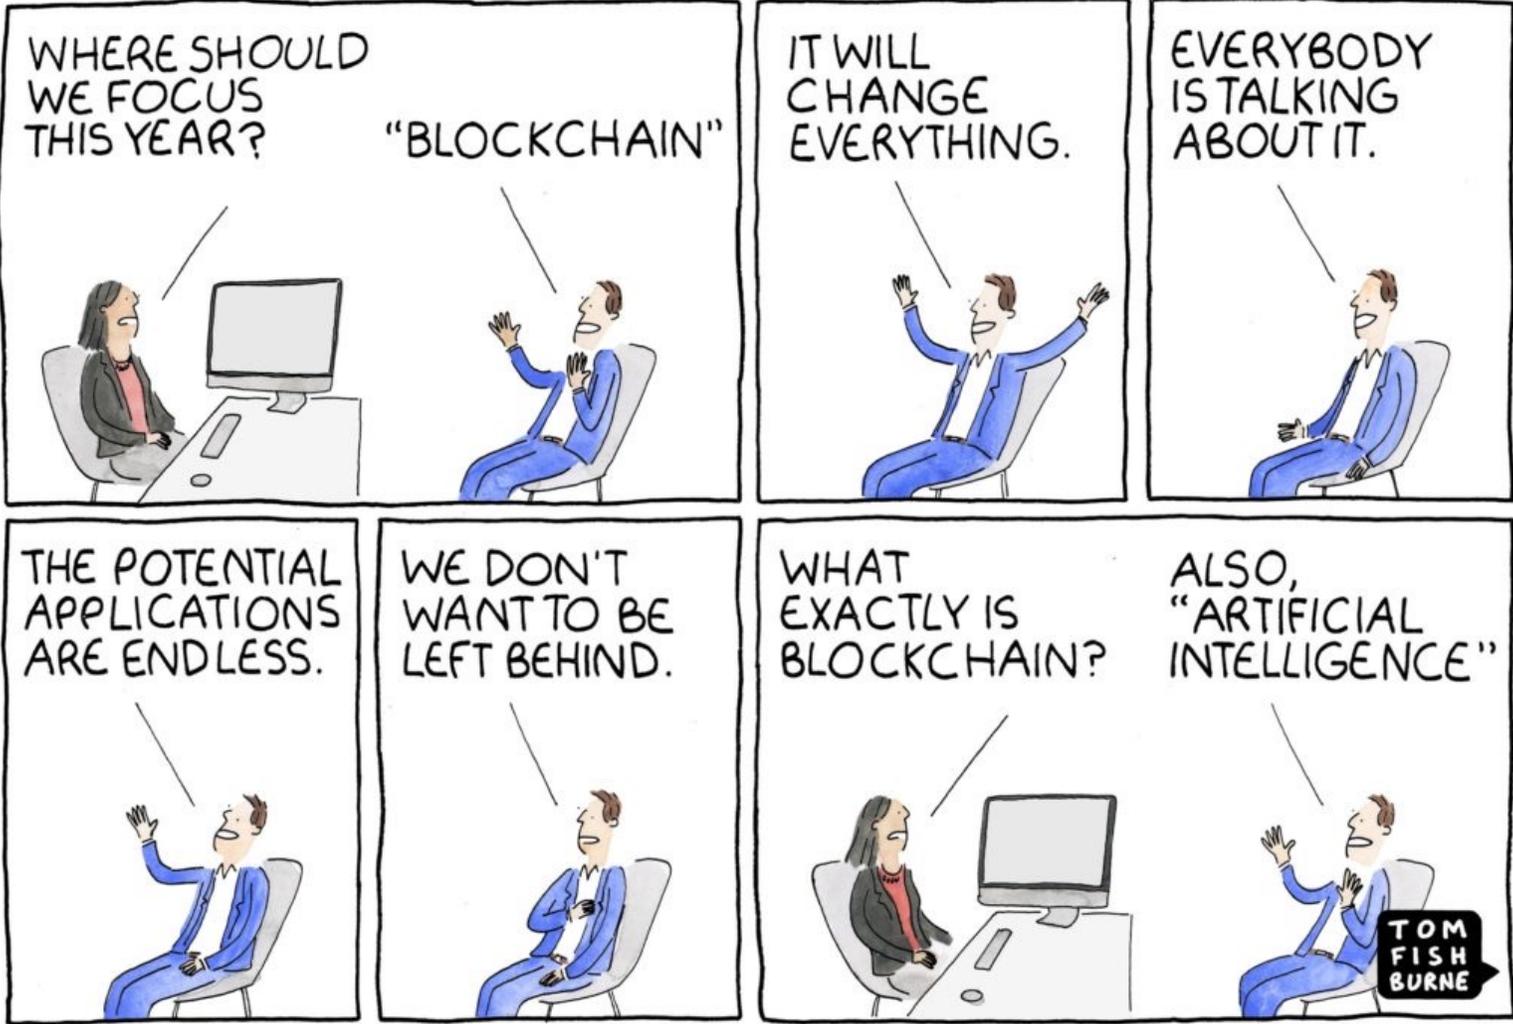


Blockchains & banking

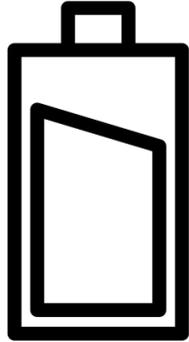


Novel election methods

# Beyond the hype...



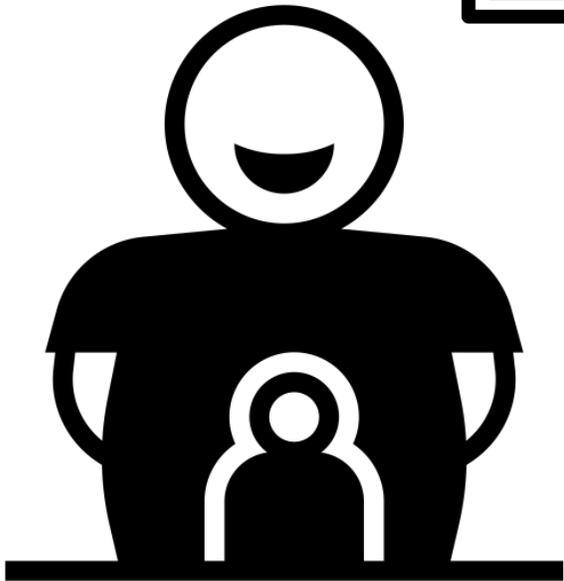
# What kind of new challenges arise?



Energy & resource  
consumption



Fairness



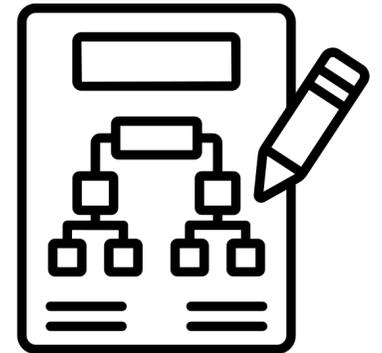
Intrusive use  
of technology



Algorithms are  
prone to errors!

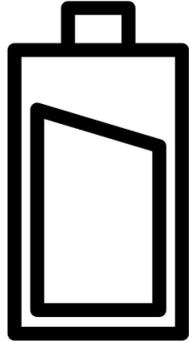


Legal responsibility



Interpretability

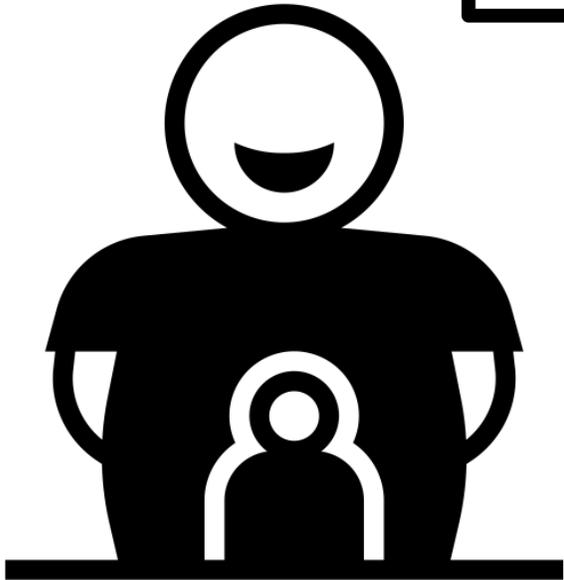
# What kind of new challenges arise?



Energy & resource  
consumption



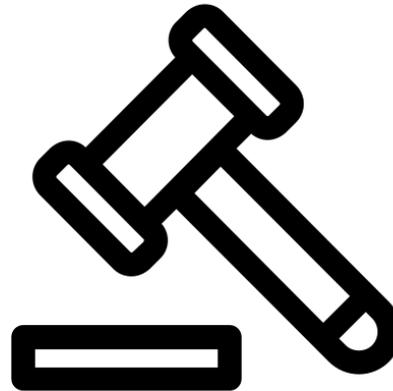
Fairness



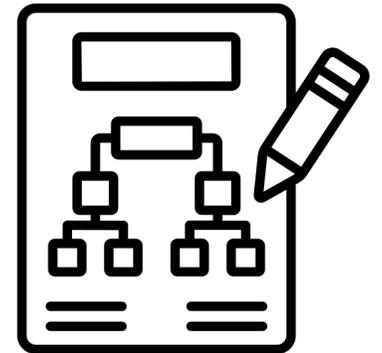
Intrusive use  
of technology



Algorithms are  
prone to errors!



Legal responsibility



Interpretability

# Examples of 'entertaining' errors:

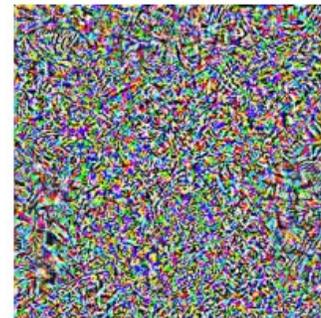
Turtle classified as a rifle



"pig" (91%)



noise (NOT random)



+ 0.005 x

=

"airliner" (99%)

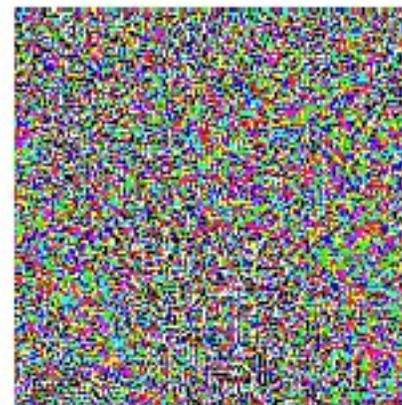


[Athalye et al. 2018]

« panda » (57 %)



noise



+ .007 x

=

« gibbon » (99.3 %)



[Goodfellow, Shlens, Szegedy 2015]

# More real-life examples...

The New York Times

## *Facial Recognition Is Accurate, if You're a White Guy*

Give this article



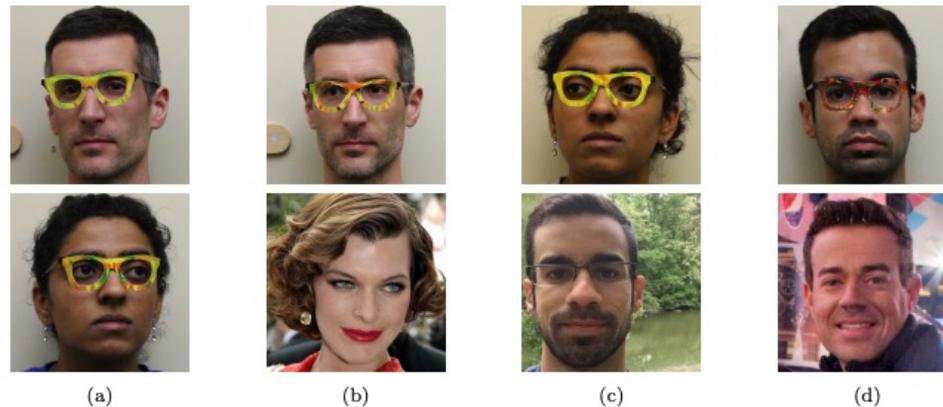
By Steve Lohr

Feb. 9, 2018

Facial recognition technology is improving by leaps and bounds. Some commercial software can now tell the gender of a person in a photograph.

When the person in the photo is a white man, the software is right 99 percent of the time.

But the darker the skin, the more errors arise — up to nearly 35 percent for images of darker skinned women, according to a new study that breaks fresh ground by measuring how the technology works on people of different races and gender.



[Sharif et al. 2017]

# More real-life examples...

## Le pilotage automatique de Tesla provoque son premier accident mortel

Par Jean-Marc De Jaeger

Publié le 01/07/2016 à 12:27, mis à jour le 01/07/2016 à 17:26



**VIDÉO - Le système de pilotage automatique n'a pas activé les freins au moment où un camion engageait une manœuvre.**

*La météo clémente a une part de responsabilité dans ce drame. Le «ciel très lumineux» n'a permis ni au pilote automatique ni au conducteur de distinguer la remorque blanche d'un poids lourd qui s'était mis perpendiculairement à la route.*

*Le Figaro, 01/07/2016*

**BBC** Sign in Home News Sport Reel Worklife

## NEWS

Home | War in Ukraine | Coronavirus | Climate | Video | World | UK | Business | Tech | Science | Stories

Tech

### Google self-driving car hits a bus

Dave Lee  
North America technology reporter

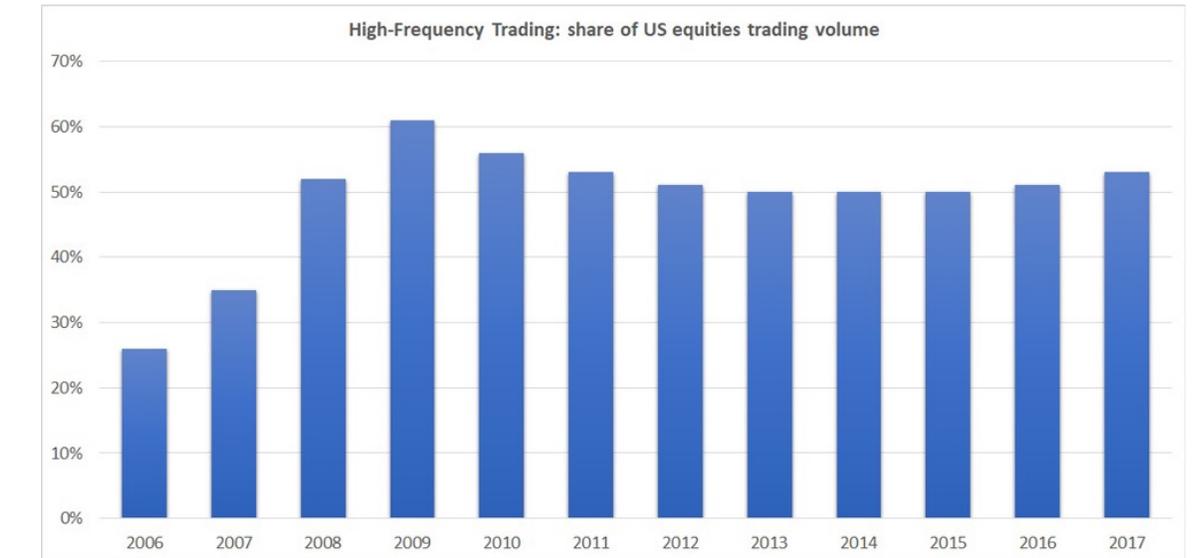
29 February 2016 | Comments

# More real-life examples...

*Dow Jones Industrial Average, May 6, 2010*



- *Plunged ~ 9.2 % within 10 minutes*
- *Market recovered within hours but flash crash erased \$1 trillion in value*
- *On April 21, 2015, nearly five years after the incident, the U.S. Department of Justice laid "22 criminal counts, including fraud and market manipulation" against Navinder Singh Sarao, a British financial trader. Among the charges included was the use of **spoofing algorithms**.*



*Wikipedia, June 2022*



Where do these « errors »  
come from?

Car accidents, recognitions failures, trading frenzy...

Where do these « errors » come from?

- *Too little data or poor data (“garbage in, garbage out”)*
- *“Covariate shifts” -> data change over time*
- *Task is too hard and should not be handed (entirely) to an AI*
- *Development disconnected to deployment*

Where do these « errors » come from?

- *Too little data or poor data (“garbage in, garbage out”)*
- *“Covariate shifts” -> data change over time*
- *Task is too hard and should not be handed (entirely) to an AI*
- *Development disconnected to deployment*

*Overall, we always face uncertainty and errors in our decisions!*

*-> Can we obtain some performance guarantees?*

## Goals for this morning:

- *Understand and measure uncertainty on a synthetic example (portfolio optimization)*
  - *How to mitigate the impact of uncertainty on the outcomes of our decisions?*
- > *How to take decisions that are « robust » ?*
- *Tools: Python & robust optimization*

# The robust optimization approach

*High level overview:*

We want to find a decision  $x \in X$  to optimize a return function  $x \mapsto R(x, \hat{\theta})$  that depends on some parameter  $\hat{\theta}$ :

$$\max \{R(x, \hat{\theta}) \mid x \in X\}$$

# The robust optimization approach

*High level overview:*

We want to find a decision  $x \in X$  to optimize a return function  $x \mapsto R(x, \hat{\theta})$  that depends on some parameter  $\hat{\theta}$ :

$$\max \{R(x, \hat{\theta}) \mid x \in X\}$$

- Game theory:  
 $x$  models a poker strategy  
 $\hat{\theta}$  models the opponents' strategy  
 $R(x, \hat{\theta}) \in [0,1]$  is the probability of winning.





# The robust optimization approach

*High level overview:*

We want to find a decision  $x \in X$  to optimize a return function  $x \mapsto R(x, \hat{\theta})$  that depends on some parameter  $\hat{\theta}$ :

$$\max \{R(x, \hat{\theta}) \mid x \in X\}$$

- Warehouse management:  
 $x$  models an order profile  
 $\hat{\theta}$  is the weekly demand for each product  
 $R(x, \hat{\theta})$  is the total revenue



# Parameter uncertainty

*Main problem: the exact value of  $\hat{\theta}$  may be unknown!*

- Game theory:  $\hat{\theta}$  is the opponents' strategy
- Portfolio optimization:  $\hat{\theta}$  is the expected return for each asset
- Warehouse management:  $\hat{\theta}$  is the weekly demand for each product

*We only have limited information about  $\hat{\theta}$  !*

# Parameter uncertainty

*Cost of failing may be high:*

- IBM's Watson gave unsafe recommendations for treating patients [The Verge, 07/26/18]
- Investment supercomputer loses \$20 million daily [Bloomberg, 05/06/19]
- Epic's AI algorithms delivering inaccurate informations on seriously ill patients [STAT, 07/21/21] – « garbage in, carnage out » !
- Racial biases at triage for critical medical resources [Obermeyer et al. 2019]

# Parameter uncertainty

Pessimistic principle: instead of solving

$$\max \{R(x, \hat{\theta}) \mid x \in X\}$$

for a known value of  $\hat{\theta}$ , we solve a *robust model*:

$$\max \left\{ \min_{\theta \in \Theta} R(x, \theta) \mid x \in X \right\}$$

for a set  $\Theta$  of *plausible values* for  $\theta$ .

What does this mean?

# Parameter uncertainty

We solve a *robust model*:

$$\max \left\{ \min_{\theta \in \Theta} R(x, \theta) \mid x \in X \right\}$$

for a set  $\Theta$  of plausible values for  $\theta$  (typically the confidence intervals).

-> We find a decision  $x$  that must « perform well » for all values  $\theta \in \Theta$  !

-> We need to find a « reasonable » set  $\Theta$ , otherwise we're too pessimistic!

# Parameter uncertainty

## Applications to robust portfolio optimization

- Connect to Hfactory:

<https://app.hfactory.io/activities>

- Open new studio -> New CPU instance -> New terminal
- Copy my files to your own work folder:

```
cp -r ~/activities_data/hi__paris_summer_school/tutorials/datasets/5A\ -\ J.\ Grand-Clement/ ~/my_work/5A
```

## Conclusion and take-away principles

- Robustness of decisions is a *\*very\** serious issue when deploying ML models for real applications.

## Conclusion and take-away principles

- Robustness of decisions is a *\*very\** serious issue when deploying ML models for real applications.
- Robust methods exist and give statistical guarantees ...  
... the drawbacks being that they may be very conservative ...  
... so we need to carefully choose the « size » of the uncertainty set!

## Conclusion and take-away principles

- Robustness of decisions is a *\*very\** serious issue when deploying ML models for real applications.
- Robust methods exist and give statistical guarantees ...  
... the drawbacks being that they may be very conservative ...  
... so we need to carefully choose the « size » of the uncertainty set!
- Note: robustness may hardly be necessary for some applications, e.g. poker/go, while it is critical for others, e.g. healthcare.