The Unbearable Lightness of AI

R. Guerraoui

2024


What is AI?


Why Now?


What Now?

What is AI?

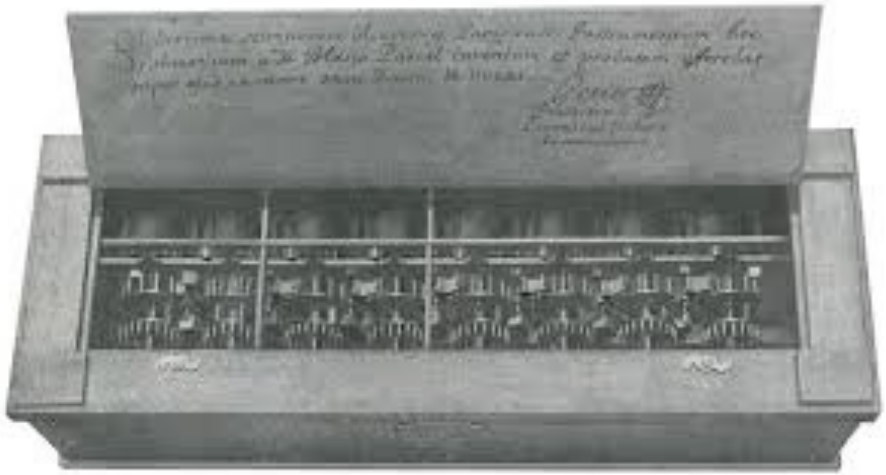

AI is the ability of a machine to solve a problem that only humans thought they could solve

AI is the ability of a human to solve a problem that only mathematicians thought they could solve

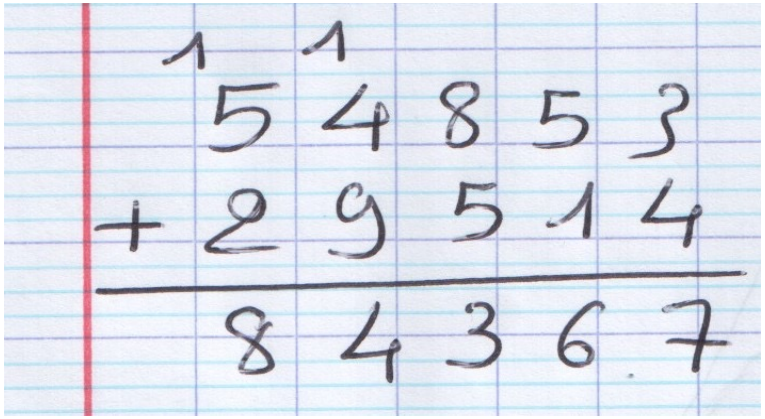

Algorithmi

The universal machine

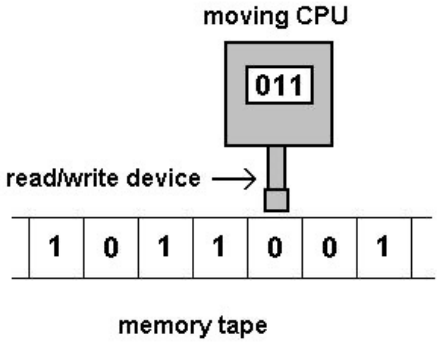

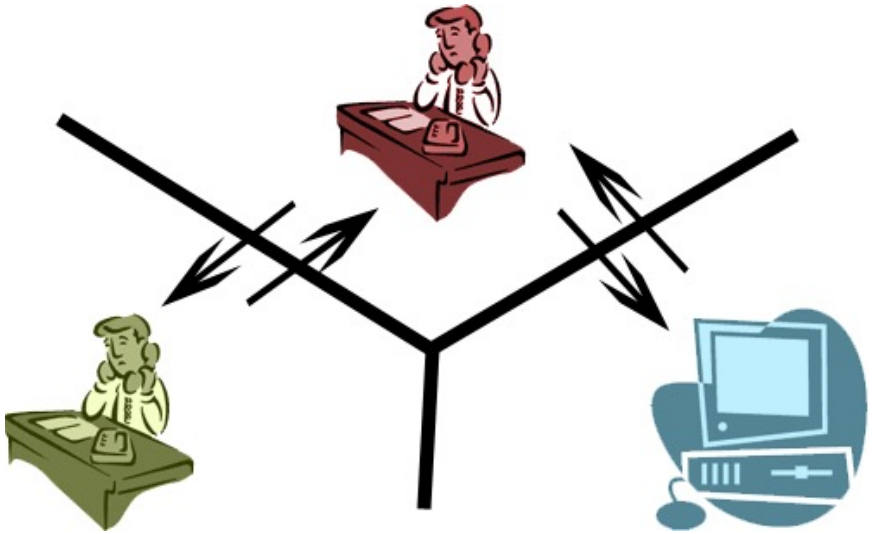

Can a machine think?

AI is (a Prowess of) Computer Science

Deep Blue: 1997

Jeopardy: 2011

Rembrand: 2016

Go: 2017


What is AI?


Why Now?


What Now?

Why Now?

(1) Algorithms that learn

(2) Machines that network

Algorithms that learn

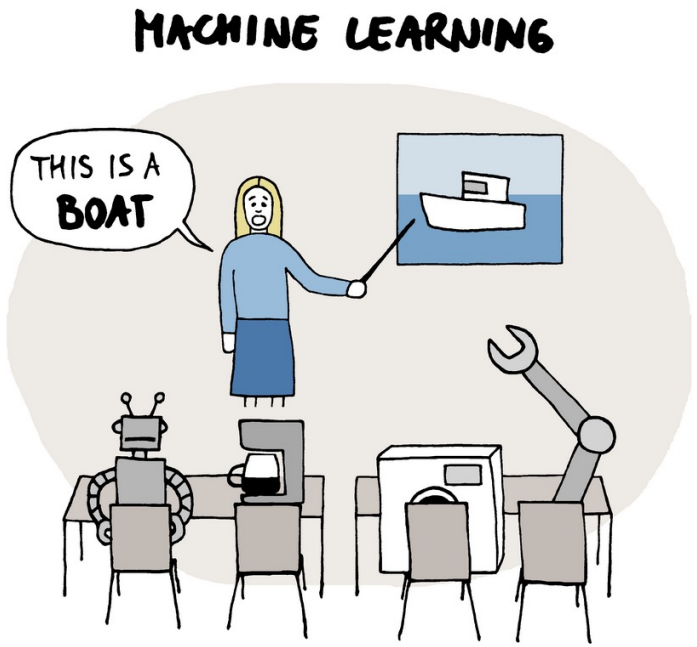

From data and mistakes


Mediego.


ouest
france


GROUPE
PUBLI
HEBDOS


LA
VOIX
DU
NORD


Le Monde


LA LIBERTÉ


LE FIGARO


Courrier
picard


LA DÉPÊCHE
DU MIDI


Swissquote


PARIS
NORMANDIE.fr


investir
LE JOURNAL DES FINANCES


Les Echos


L'union
L'Ardennais

Collaborative filtering

Each user has a profile

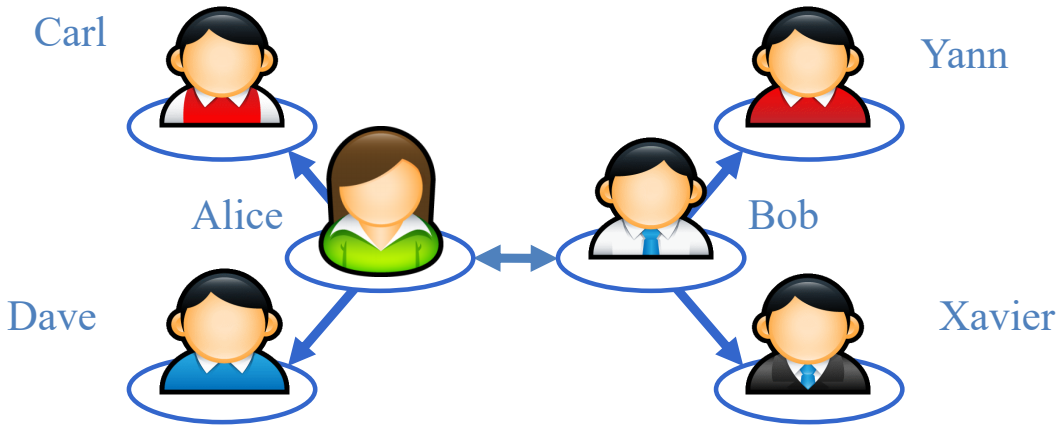

Learning from neighborgs

Algorithms that network

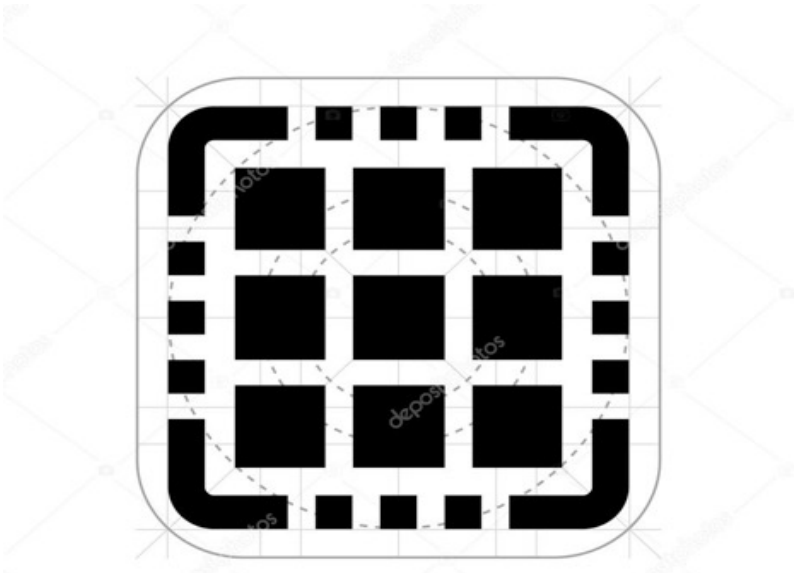

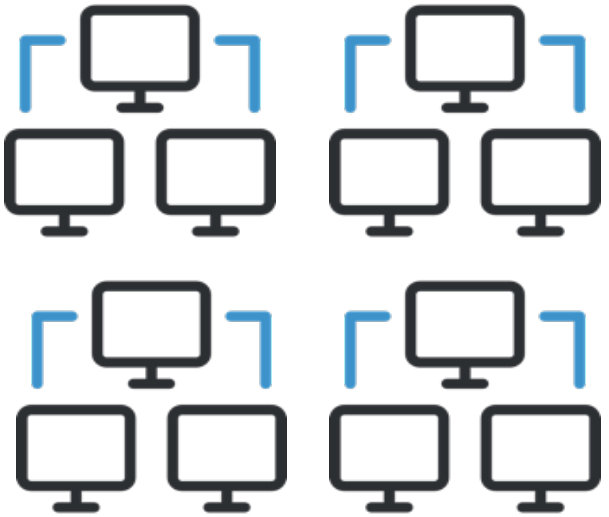

Parallel and distributed

The power of distributed computing

Federated Learning

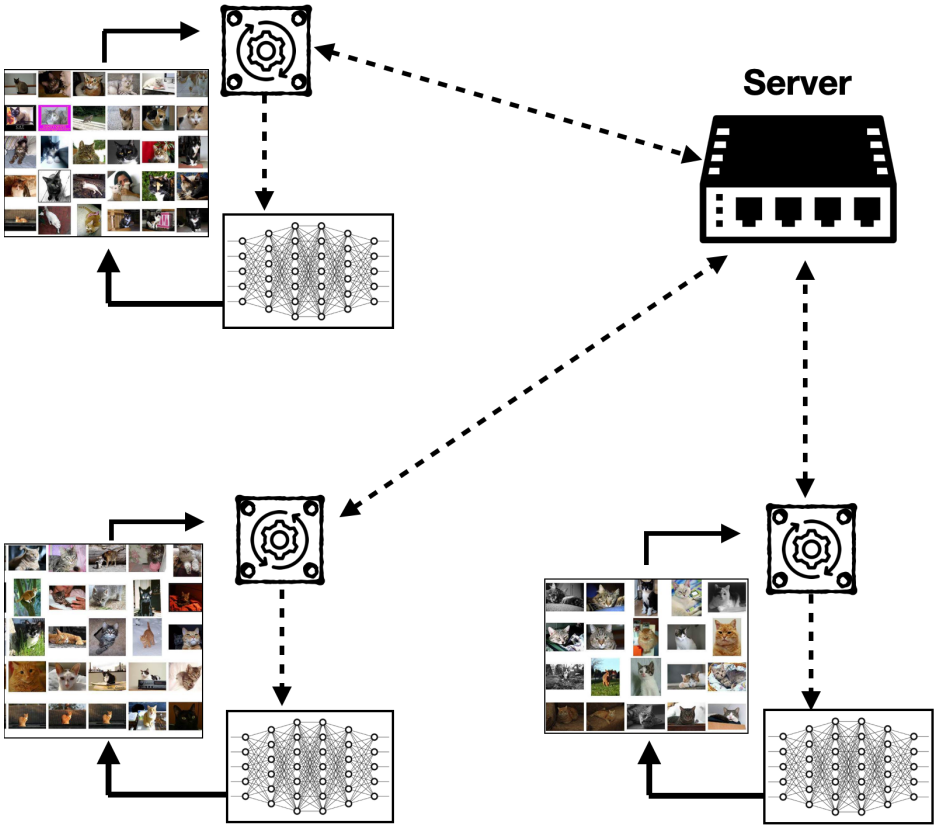

Parameter space $\subseteq \mathbb{R}^d$ — Data space — Local samples

$$\text{loss} : \Theta \times \mathcal{Z} \to \mathbb{R} , \qquad S_i \subset \mathcal{Z}^m , \; i \in [n]$$

$$\text{Loss}^{(i)}(\theta) := \frac{1}{m} \sum_{z \in S_i} \text{loss}(\theta , z)$$

Local loss function

Goal: $\min_{\theta \in \Theta} \frac{1}{n} \sum \text{Loss}^{(i)}(\theta)$

Global loss function

Benefits

Scalability
Preserves data ownership
Data diversity

X

FROM CIRCUS ANIMALS TO PETS

AI is faster and makes fewer mistakes than humans

But AI makes faster and bigger mistakes

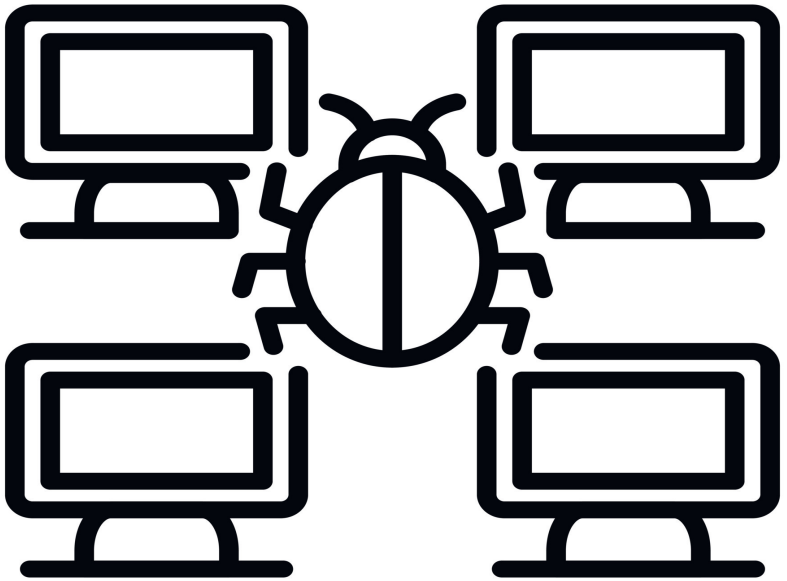

Bad Network = Bad Everything


What is AI?


Why Now?


What Now?

What Now?

Master the principles

Protect but beware

Principles

Gradient Descent, Jaccard Similarity, Central Limit Theorem, …

Distributed Computing, Networking, Consensus, Privacy….

Robustness

Theorem (BEGS17): no linear combination of DSGD updates is robust to a single adversarial worker

Distributed Stochastic Gradient Descent (DSGD)

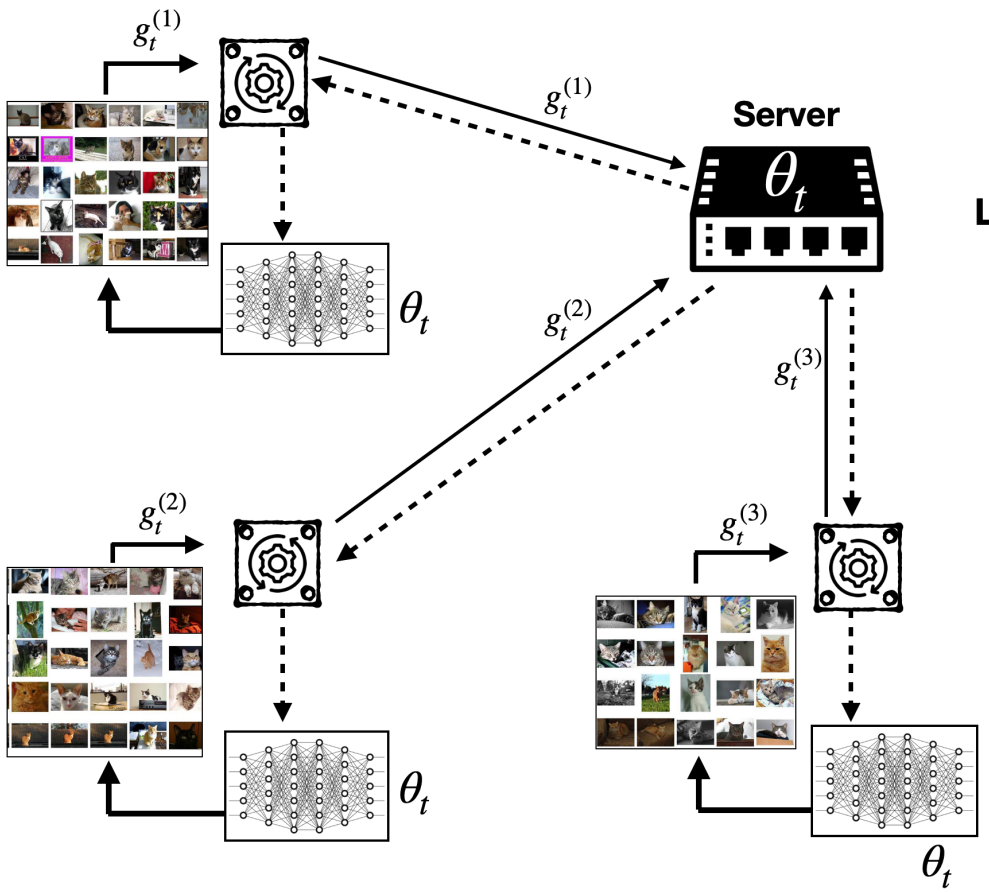

Local stochastic gradient

Local Phase: Each *node* i computes $g_t^{(i)} := \nabla \text{loss}\left(\theta_t\ ,\ z_t^{(i)}\right)$

$z_t^{(i)} \sim \mathcal{U}\left(S_i\right)$

Global Phase: The server updates the model:

$$\theta_{t+1} = \theta_t - \gamma_t\ \text{Avg}\left(g_t^{(1)}\ ,\ \ldots\ ,\ g_t^{(n)}\right)$$

Average of gradients

$$\theta_t \longrightarrow \min_{\theta \in \Theta}\ \frac{1}{n} \sum \text{Loss}^{(i)}\left(\theta\right)$$

X

Problem with Standard Learning Schemes

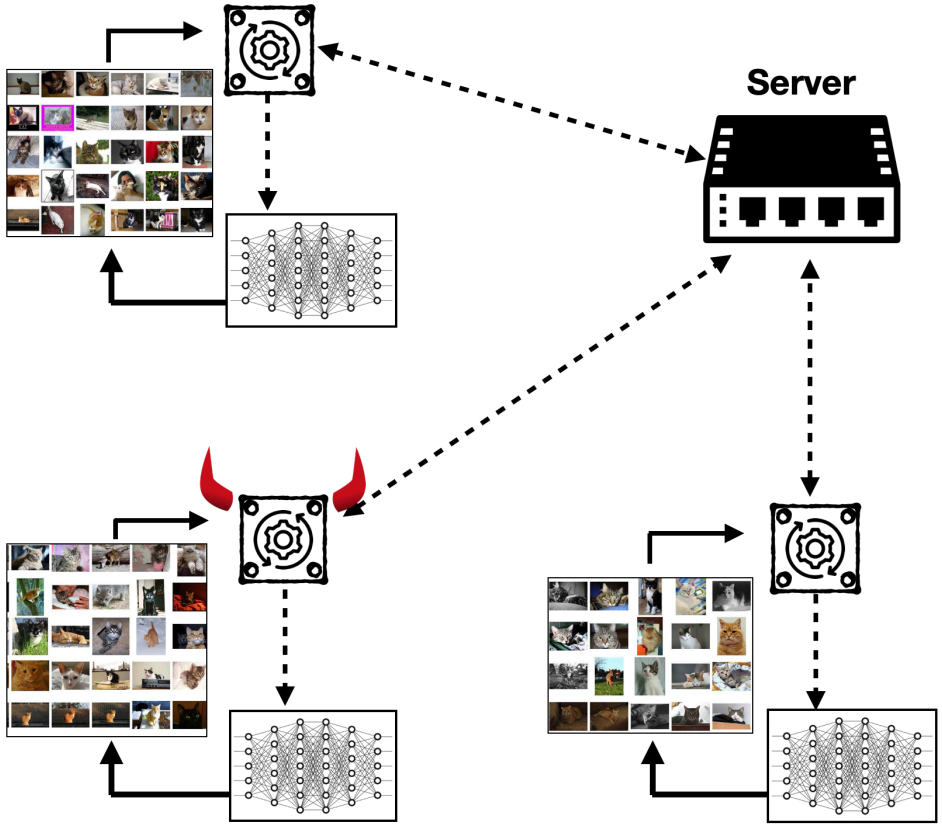

Trust **all data and machines**

$$\theta_{t+1} = \theta_t - \gamma_t \, \text{Avg}\left(g_t^{(1)}, \ \ldots \ , g_t^{(n)}\right)$$

Ineffective when some machines *misbehave*

Consensus Impossibility

Theorem (FLP85): no algorithm can solve consensus among an asynchronous network of machines

State Machine Replication

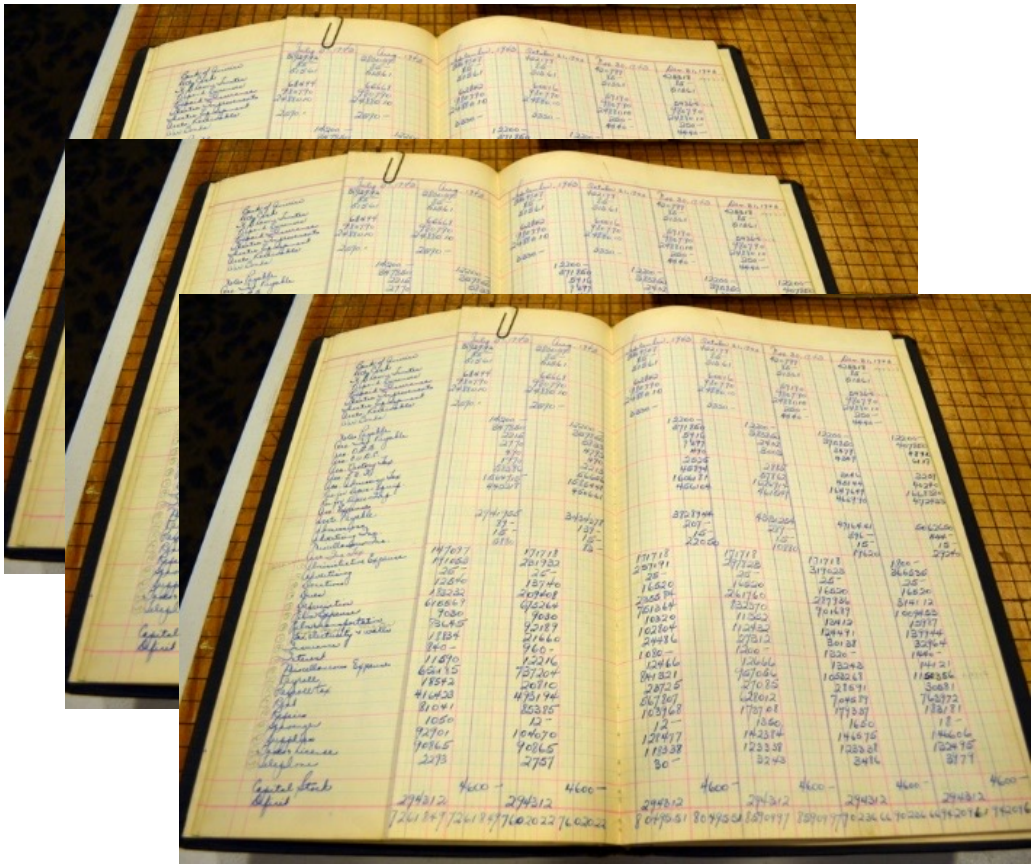

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

Robust Machine-Learning

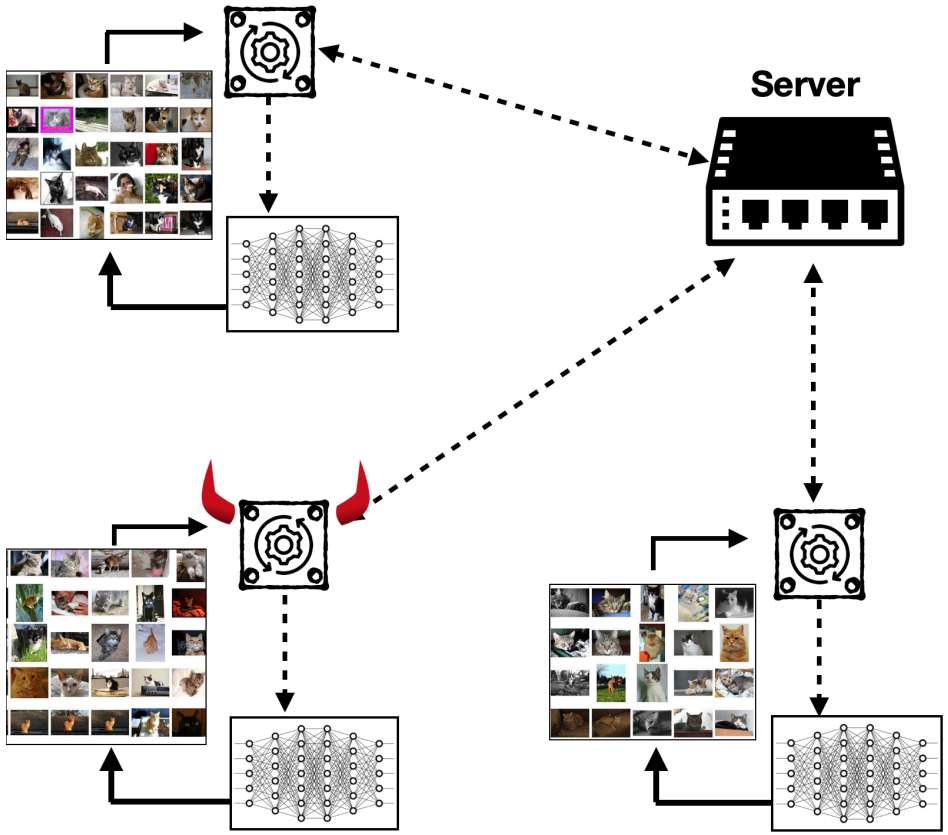

f nodes *of unknown identity* are **adversarial**

In minority

Need not follow the algorithm

$$\text{Loss}^{(i)}(\theta) := \frac{1}{m} \sum_{z \in S_i} \text{loss}(\theta, z)$$

Robust ML Goal: $\min_{\theta \in \Theta} \frac{1}{n-f} \sum_{i \in H} \text{Loss}^{(i)}(\theta)$

Set of *honest* nodes

?

We do not know this set!

"Fault-Tolerance in Distributed Optimization: The of Redundancy". **Gupta** and Vaidya. *PODC 2020*

X

Privacy vs Robustness vs Efficiency

Theorem (AGS23): private, robust and efficient machine learning is impossible

Decidability

Theorem (Turing36): algorithms cannot solve all problems (halting, printing, satisfactoriness…)

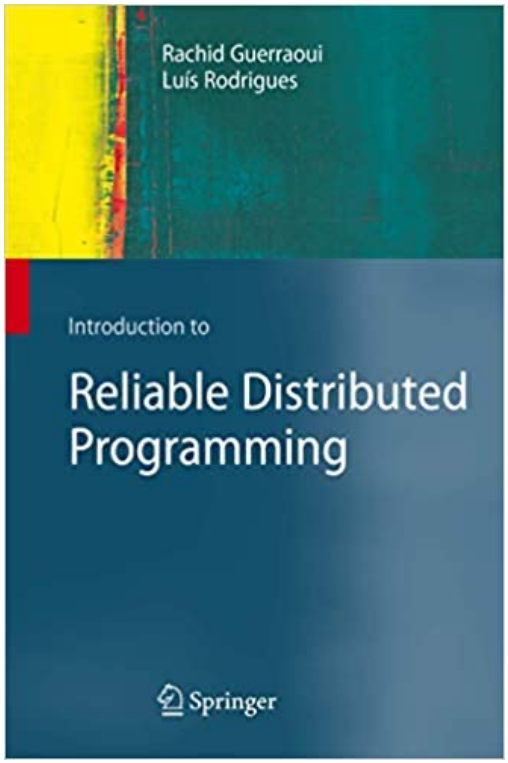

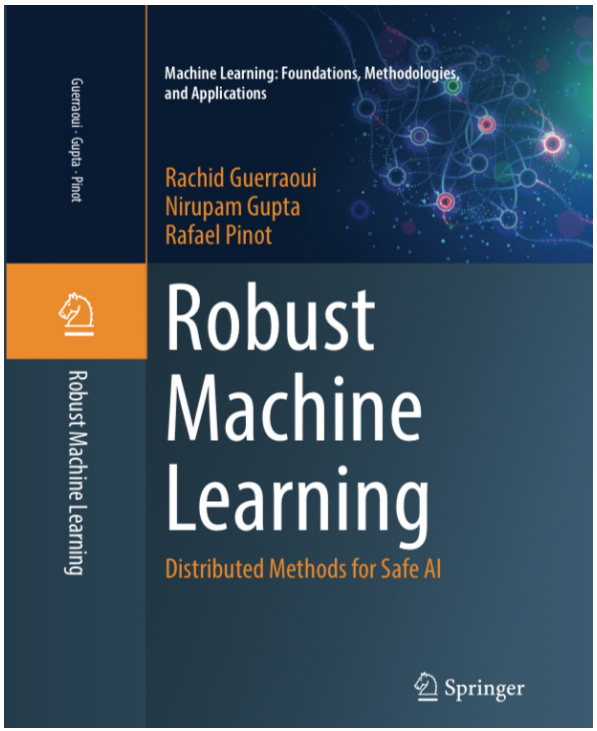

« Here are my principles, if you do not like them, I have others » G. Marx