

# La Statistica al CNR al servizio del Paese

Società Italiana di Statistica e Consiglio Nazionale delle Ricerche

5 Luglio 2022

## Artificial Intelligence Based on *Statistical Learning* Ontology, Effectiveness and Limitations

**Giovanni Cerulli**

IRCrES-CNR

Research Institute on Sustainable Economic Growth  
National Research Council of Italy

# OUTLINE

## **Part 1.**

AI definition, paradigms, application and limitations

## **Part 2.**

AI research at CNR-IRCRES: Optimal Policy Learning (OPL)

# PART 1

AI definition, paradigms,  
application and  
limitations

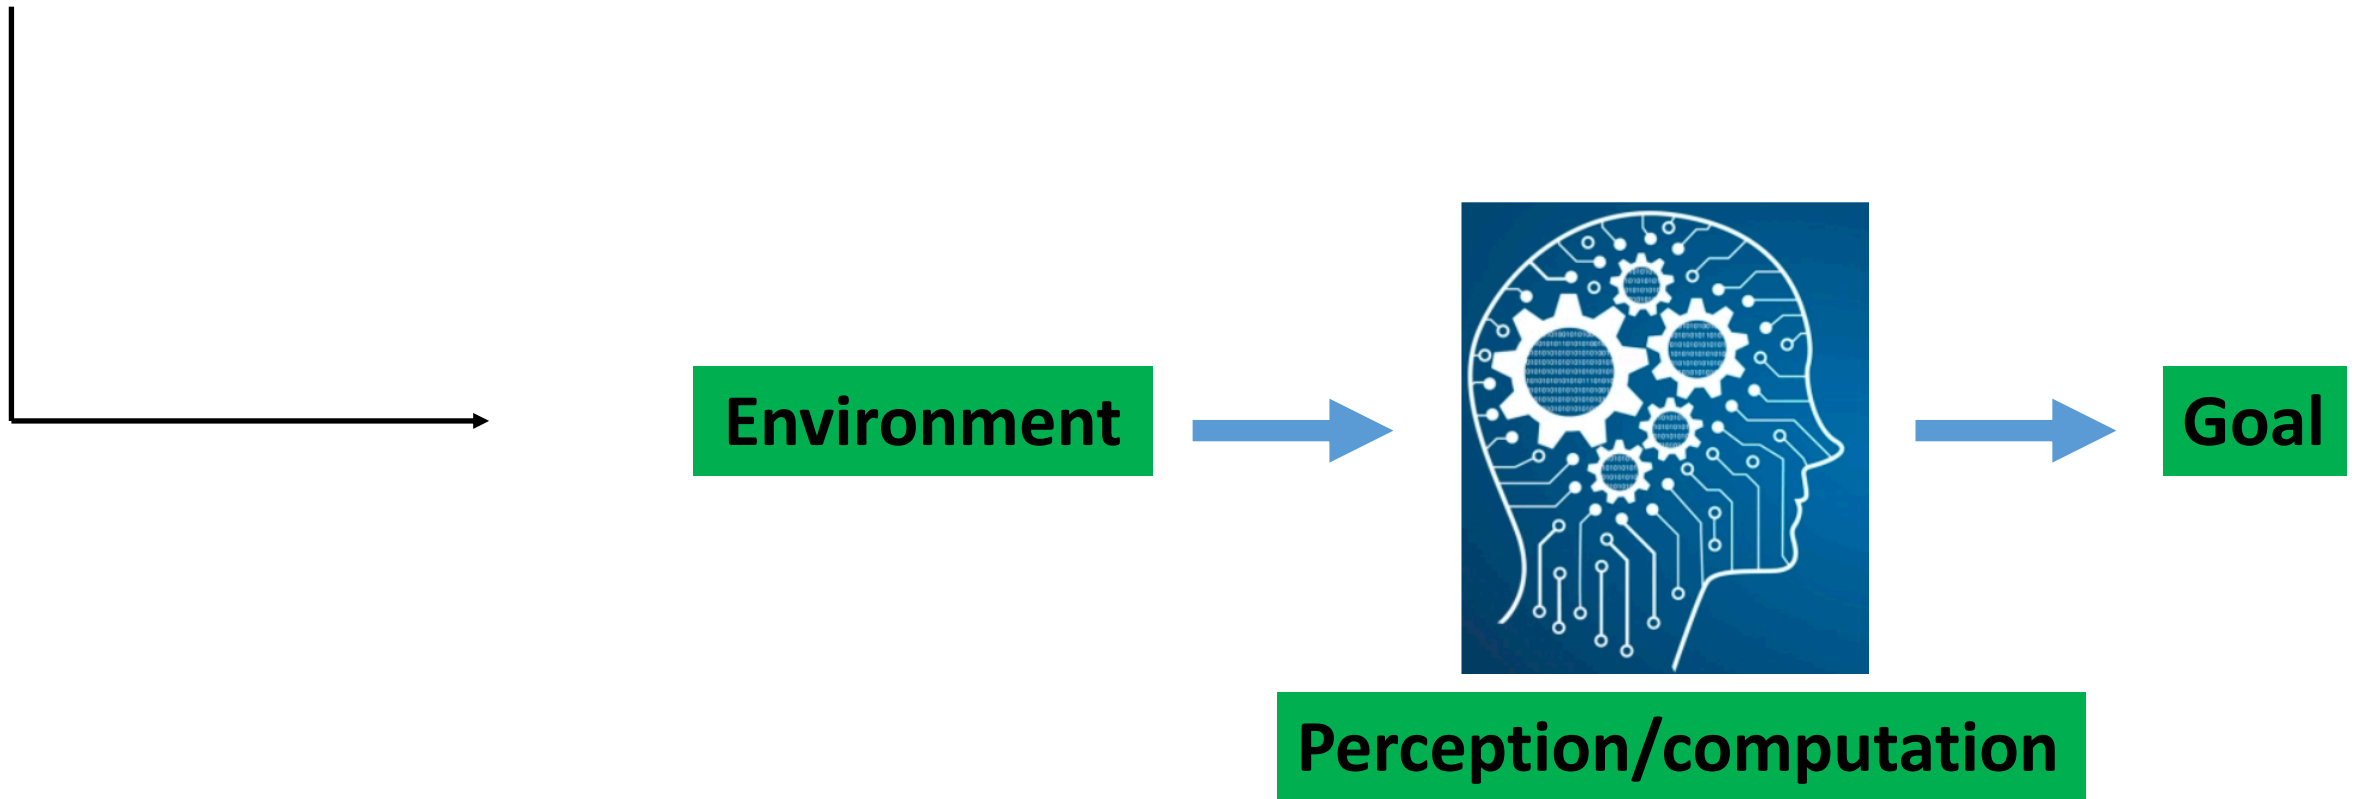
# What is **Artificial Intelligence**?

- **Artificial intelligence (AI)** is intelligence demonstrated by artifacts (e.g., machines)
- AI is opposed to **natural intelligence (NI)** displayed by animals, including humans
- AI endows machines of human-like **cognitive** functions, such as **learning** and **problem solving**



# What is **intelligence**?

**Ability** to perceive an environment, elaborates on it, and take **actions/decisions** that *maximize the chance of achieving a given goal*



# The ingredients of the *cake*: $AI = C + B + A$

**Computation**

+

**Big data**

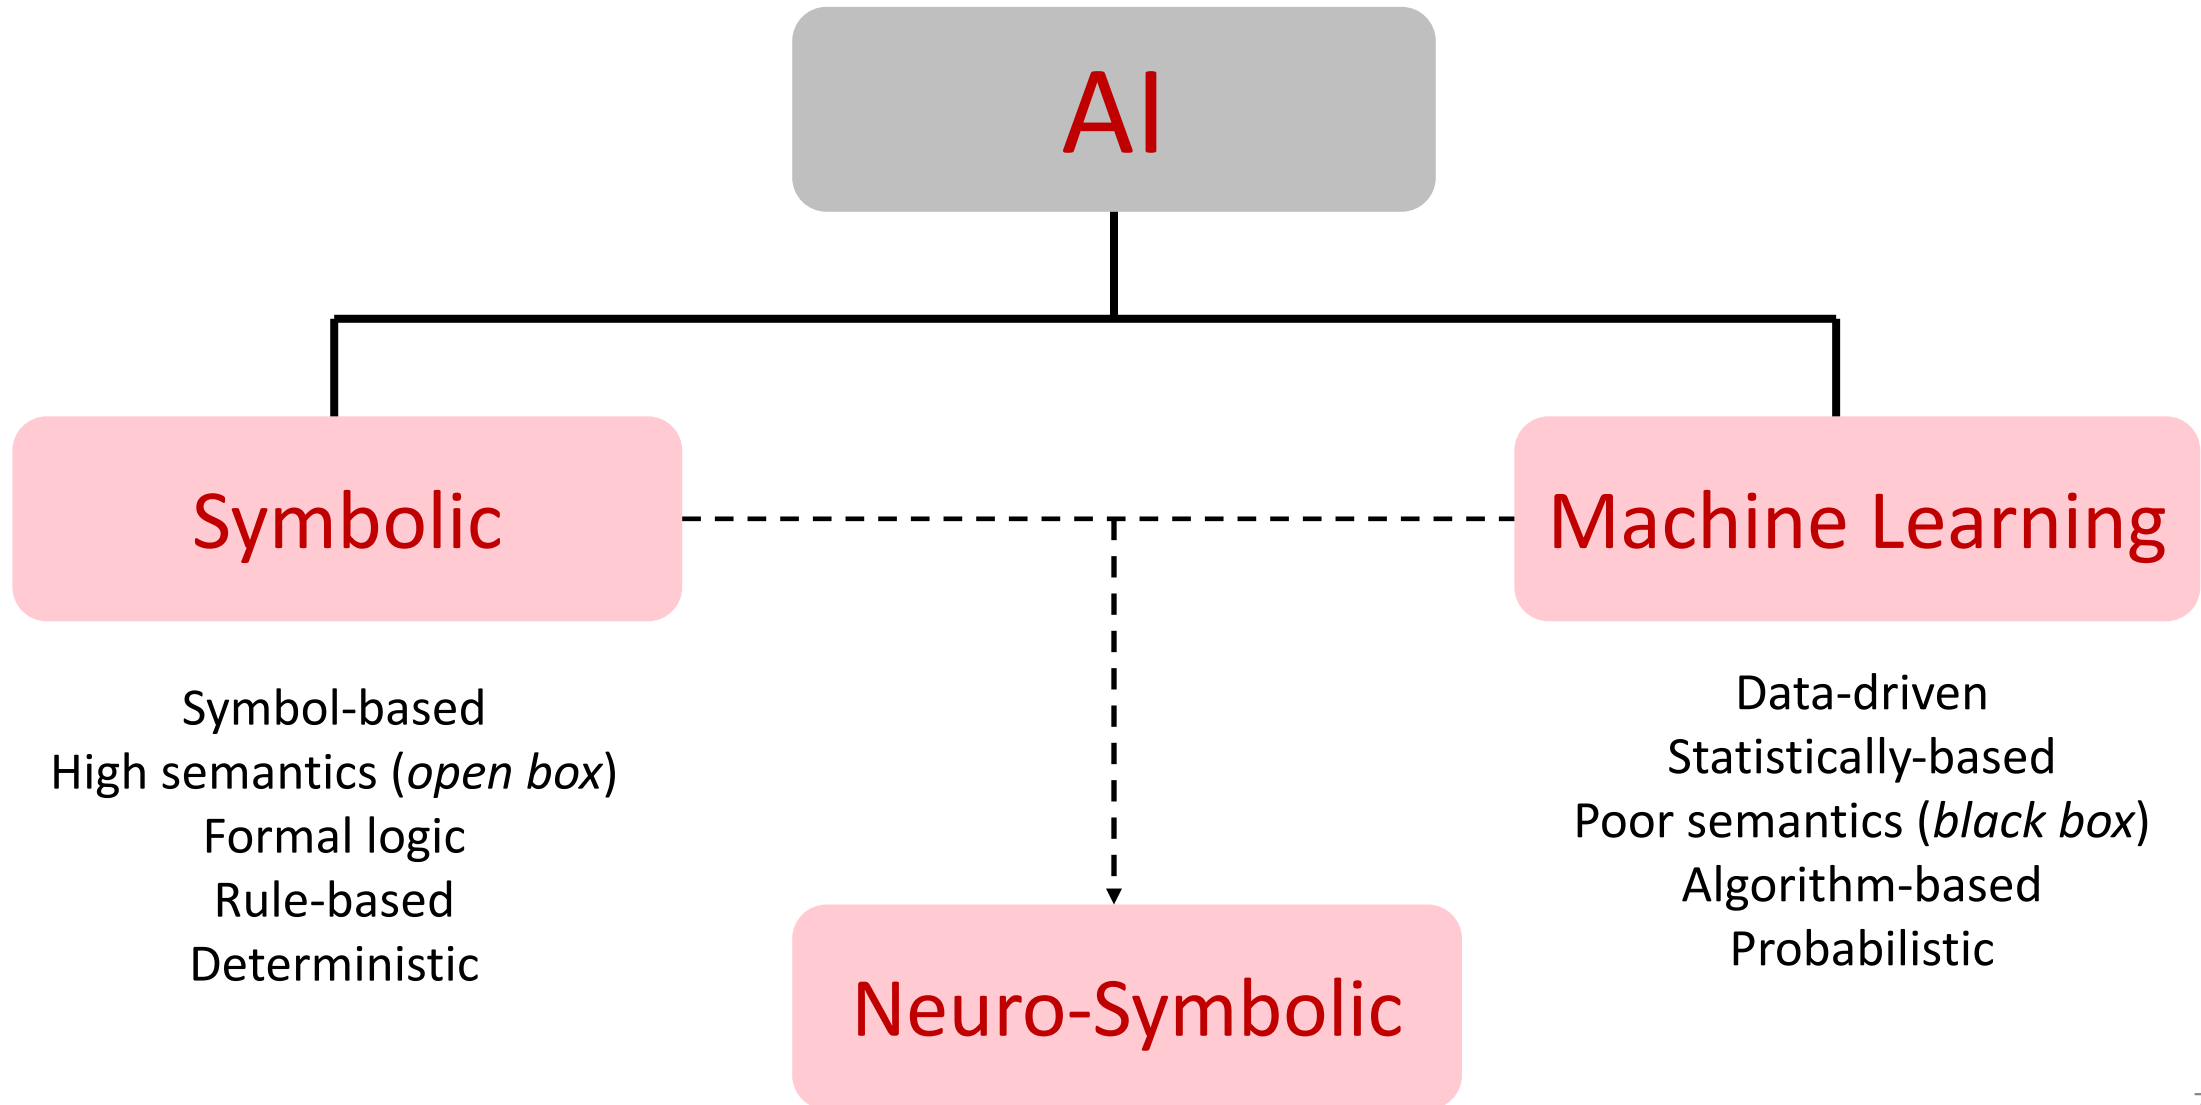
+

**Algorithms**

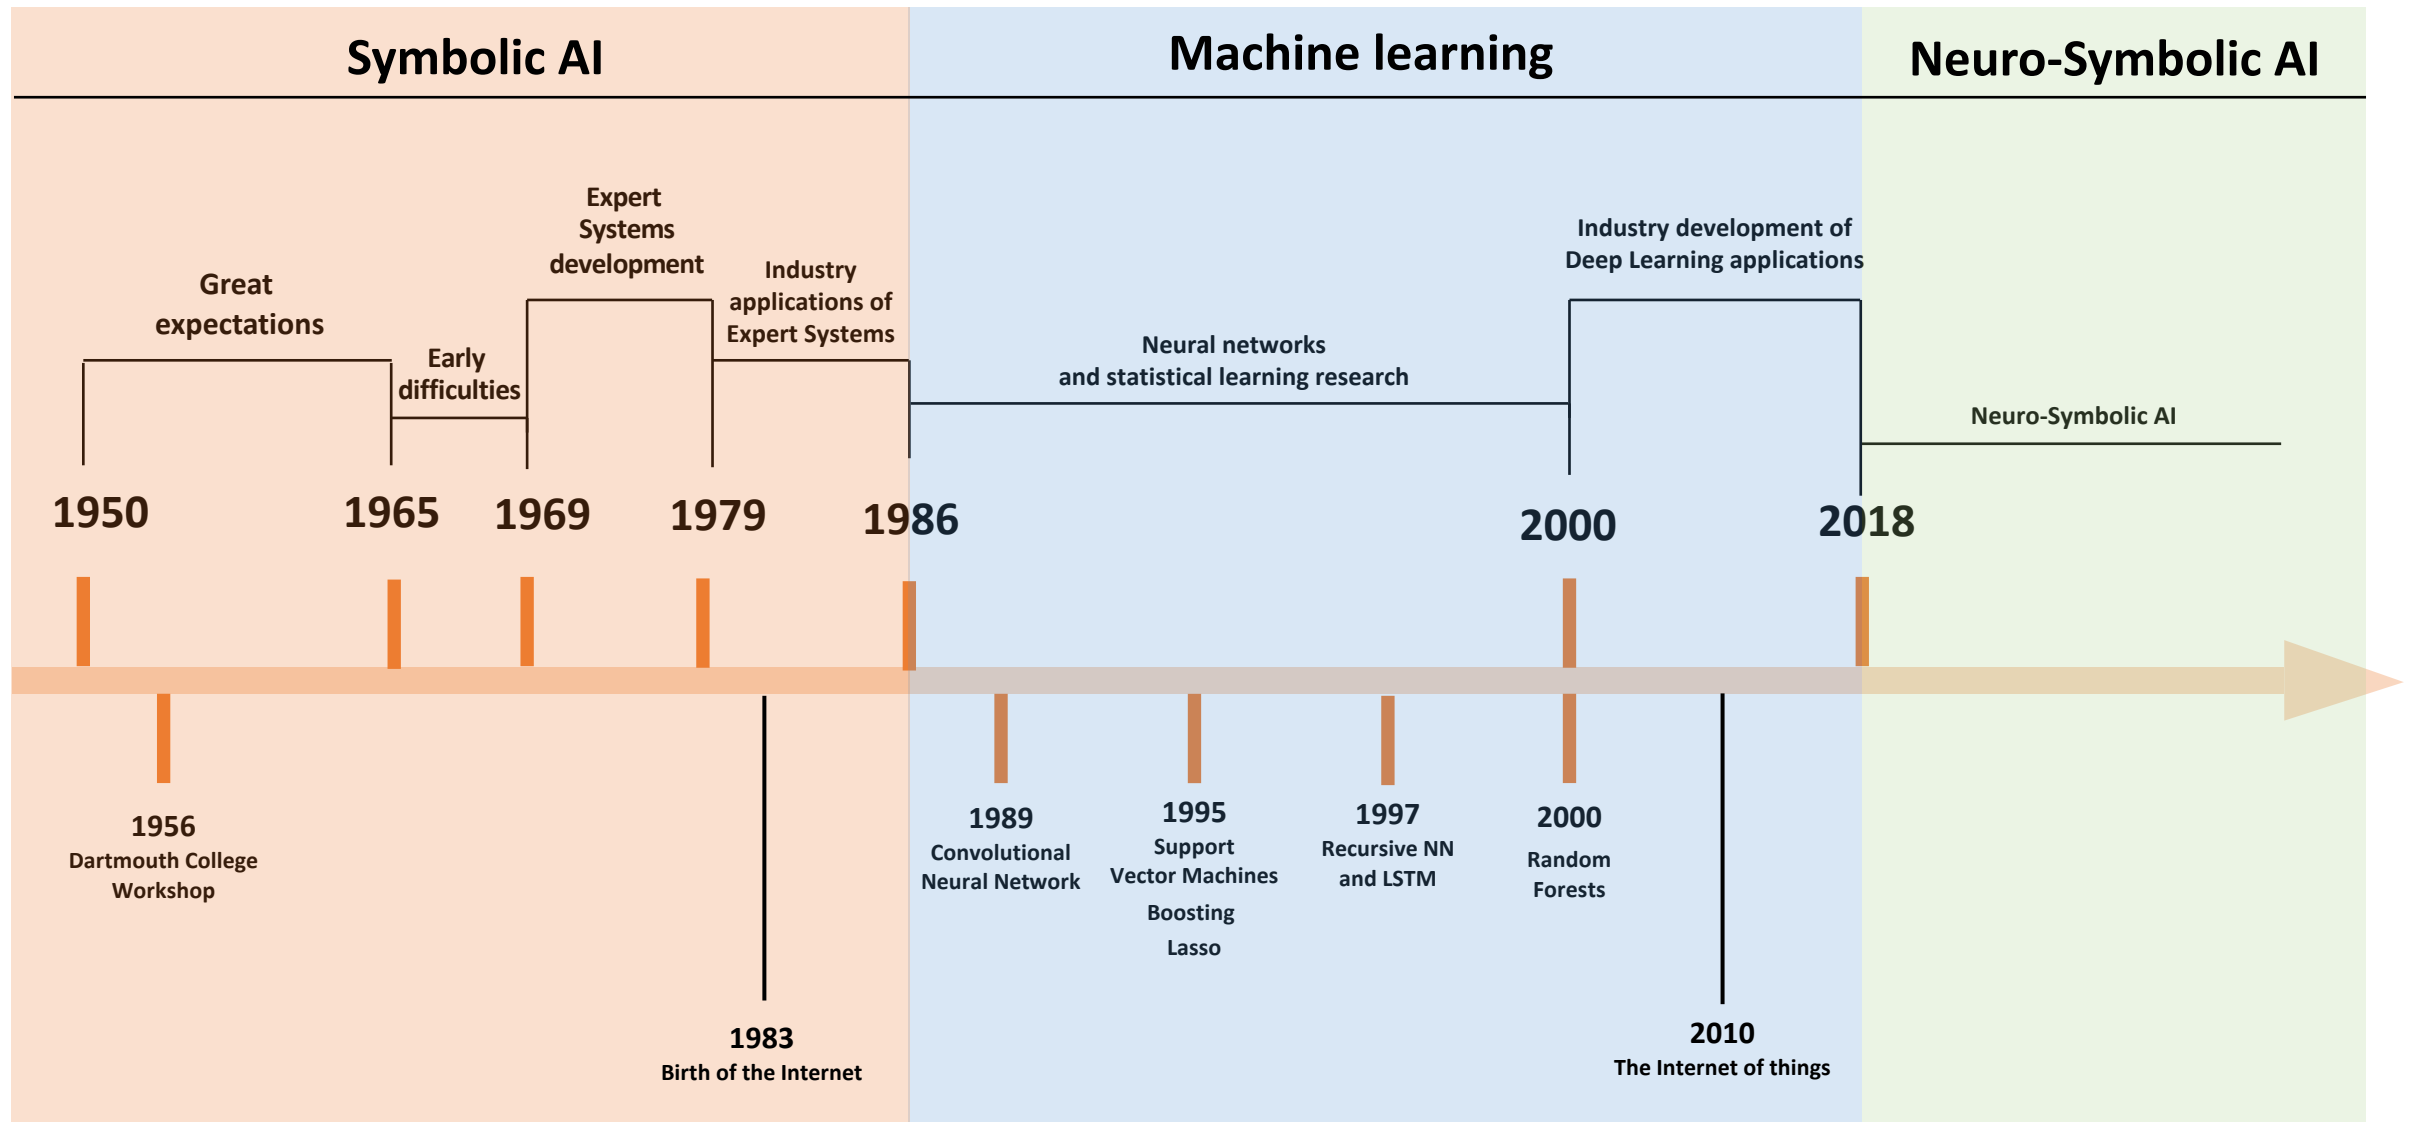


**AI** is the **product** of a tremendous increase in C, B, and A

# AI paradigms?



# Timeline of AI development?



**Symbolic AI**

**vs.**

**Statistical learning**

# Rise and fall of **Symbolic AI** - 1

What is **Symbolic AI** ?

Branch of artificial intelligence attempting to explicitly represent **human learning** in a **declarative form** (i.e. facts and rules)

Based on Thomas Hobbes statement:

*«Thinking is manipulation of symbols and Reasoning is computation»*



Allen Newell, Herbert A. Simon  
Pioneers in Symbolic AI



Building “General Problem Solvers”

Hard-wired rule-based reasoning systems like Expert Systems became the foundation for almost 40 years of AI research

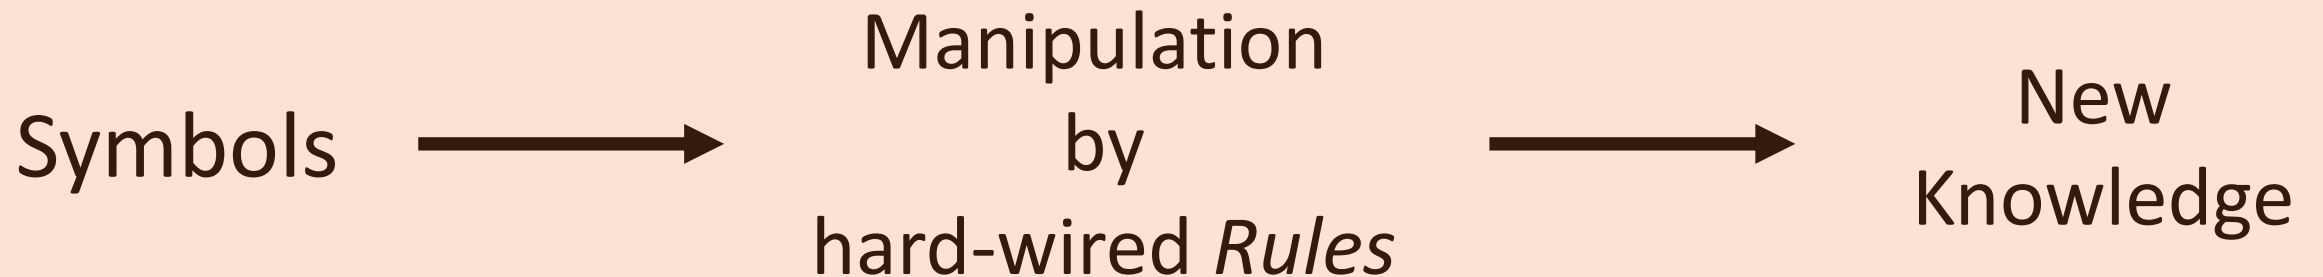
# Rise and fall of **Symbolic AI** - 2

## What is a symbol?

*“A perceptible something that stands for something else”*

- Alphabet symbols, numerals, road signs, music symbols, etc.
- A symbol such as ‘apple’ it symbolizes something which is edible, red in color. In some other language, we might have some other symbol which symbolizes the same edible object

New knowledge comes up from manipulating symbols via *logical rules*



# Rise and fall of **Symbolic AI** - 3

## The fall of Symbolic AI

The main problem with Symbolic AI was that it aimed at

**TEACHING A MACHINE LOGICAL RULES**

instead of

**LETTING THE MACHINE TO LEARN FROM EXPERIENCE**

Complex intelligence tasks, general patterns' recognition and unspecialized procedural tasks cannot successfully carried out using the Symbolic AI paradigm



# Rise and fall of **Symbolic AI** - 4

The Symbolic AI cannot solve the “*learning-by-experience paradox*”



**Emma** was here a six year old girl. She spoke a fluent Italian. She was able to understand it and make herself understood by other Italian people

Emma was *never taught a single Italian grammar rule*. She did not even know about the existence of grammar rules

But she was able to recognize words, understand their meaning, and process them accordingly

She has learnt Italian *by experience*

# The rise of *Statistical Learning* - 1

In 1959, Arthur Samuel, pioneered the field of Machine Learning (ML), by defining it as the *“field of study that gives computers the ability to learn without being explicitly programmed”*

ML can be understood as a set of computational methods that use *“**experience**”* (i.e. *past event frequencies*) to improve performance by generating **accurate predictions**

**IF HUMANS LEARN FROM EXPERIENCE**



**LET MACHINES LEARN FROM DATA**

# The rise of *Statistical Learning* - 2

Data ( $y, X$ )

Collection of many single events occurred in the past (frequencies = experiences) linking an event  $y$  and many events  $X$

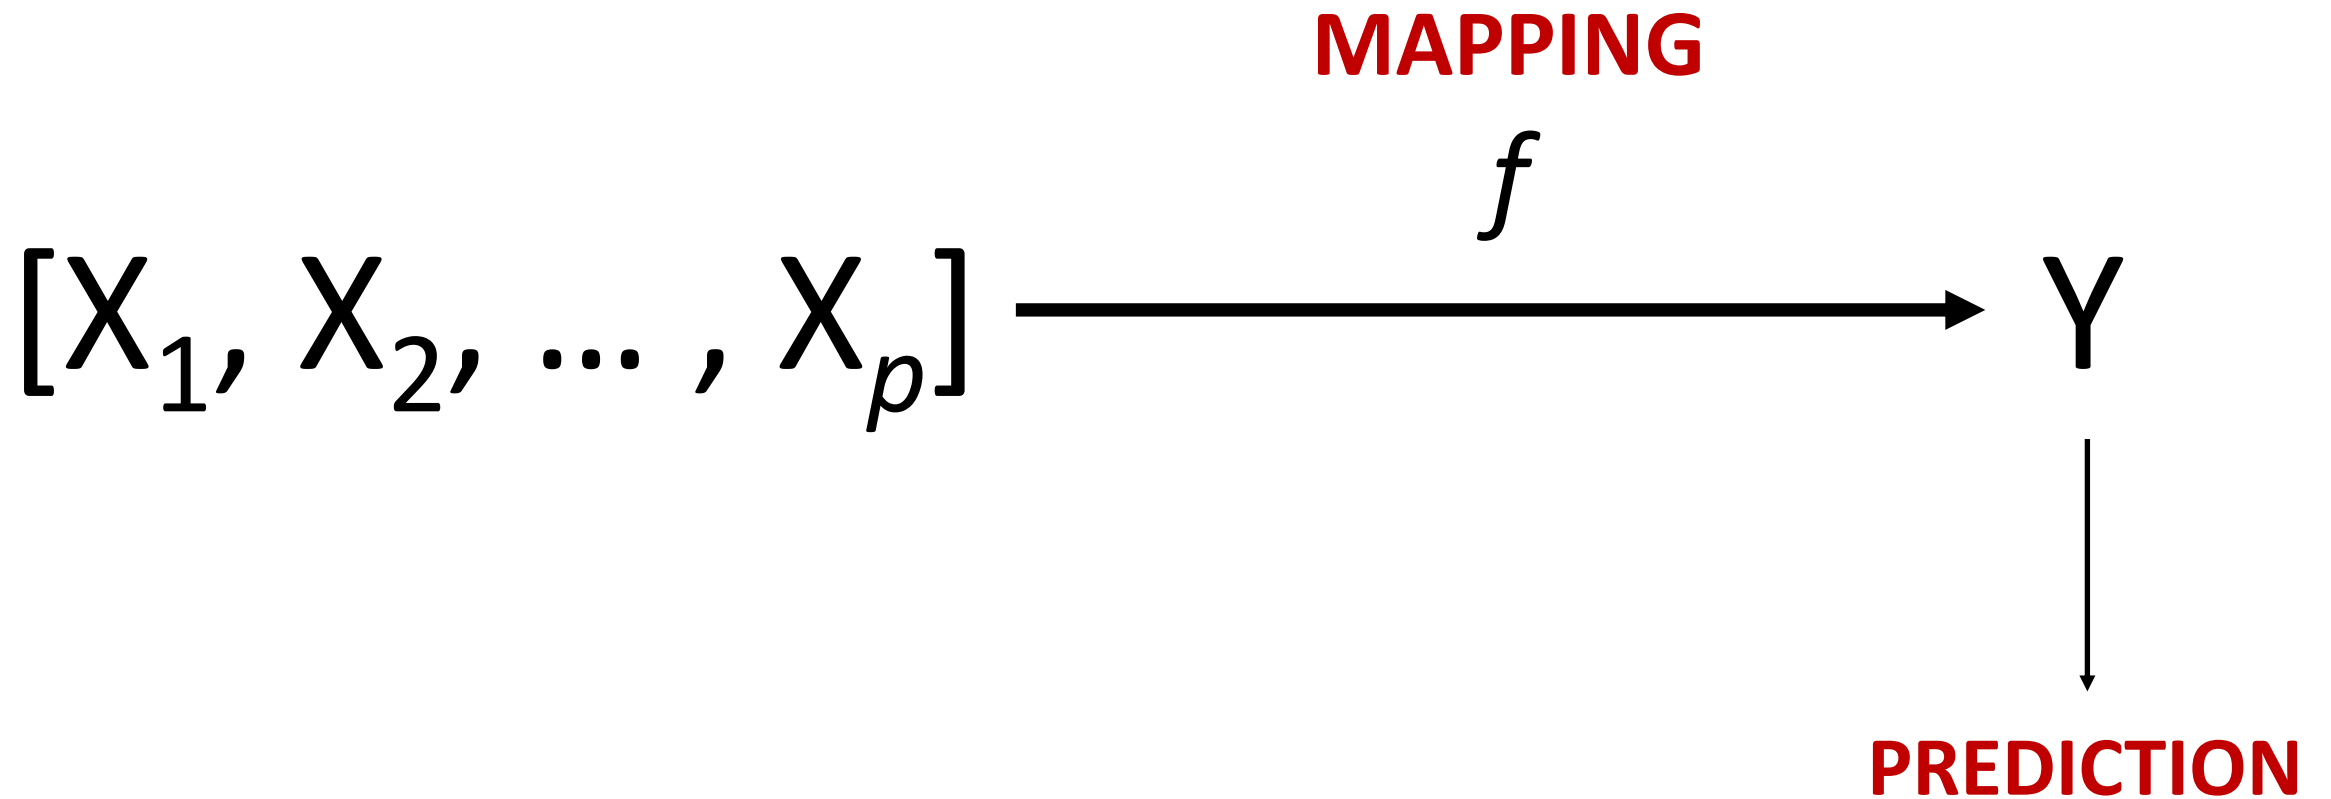


Learning algorithm: a mapping of  $X$  into  $y$



Learning = higher predictive power

# The rise of *Statistical Learning* - 3



# The rise of *Statistical Learning* - 4

The fundamental **statistical object** designing the **mapping** is the **conditional expectation** of  $y$  given  $\mathbf{X} = [X_1, X_2, \dots, X_p]$ :

$$E(Y \mid X_1, X_2, \dots, X_p) = f(\mathbf{X})$$

NOTE: we do not focus on  $P(y \mid \mathbf{X})$  as it is too demanding

# EXAMPLE

Numeral **recognition**


*Symbolic AI*

**VS.**

*Machine Learning*

# Numerical recognition

**Task:** Making a machine able to recognize whether this numeral is a 1 or a 7




This number is a 1

This number is a 7

# Numerical recognition: Symbolic AI - 1

**Possible strategy.** Put the numbers into a rectangle. Then, consider the angle A as follows:




Come up with this possible **rule**:

- If ( $A \leq 45^\circ$ ): it is a 7
- Else if ( $A > 45^\circ$ ): it is a 1



# Numerical recognition: Symbolic AI - 2

**Complication.** Number 7 is commonly written using a steeper bar (we can then exploit angle B):




Come up with this possible **rule**:

If  $(A \leq 45^\circ) \ \& \ (B > 45^\circ)$ : it is a 7

If  $(A > 45^\circ) \ \& \ (B \leq 45^\circ)$ : it is a 1

# Numerical recognition: Symbolic AI - 3

**Complication.** Number 7 is commonly written using a steeper bar (we can then exploit angle B):



Come up with this possible **rule**:

If  $(A \leq 45^\circ) \ \& \ (B > 45^\circ)$ : it is a 7

If  $(A > 45^\circ) \ \& \ (B \leq 45^\circ)$ : it is a 1

# Numerical recognition: limits of **Symbolic AI** - 1

Can we design **rule-based symbolic algorithms** able to let a machine recognize “hand-written” numerals?

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

0 1 2 3 4 5 6 7 8 9

# Image recognition: Limits of **Symbolic AI** - 2

Can we design **rule-based symbolic algorithms** able to let a machine recognize patterns like these ones?

Puppy  
or  
Muffin ?



# Numerical recognition: limits of **Symbolic AI** - 3

- The previous **rule-based recognition algorithm** is simple, and it is likely to assure a recognition accuracy around 90% or even more
- However, when considering other numerals, the problem gets more complicated, and we need a larger set of rules to come up with a good recognition accuracy
- But what about recognition of fuzzier images, as "hand-written" numerals? In this case things get much more complicated and rule-based recognition may dramatically fail or require extravagant ruling patterns, that most of the times are "local rules" not generalizable to other recognition settings
- **Machine Learning** – i.e. learning from experience – is a workable more effective solution

# Numerical recognition: Machine Learning

ML uses past recorded experience on how people connected specific numerals to specific configurations (patterns) of handwritten numerals

0 1 2 3 4 5 6 7 8 9

*0 1 2 3 4 5 6 7 8 9*

*0 1 2 3 4 5 6 7 8 9*

*0 1 2 3 4 5 6 7 8 9*

*0 1 2 3 4 5 6 7 8 9*


*0 1 2 3 4 5 6 7 8 9*

*0 1 2 3 4 5 6 7 8 9*

# Numerical recognition: Machine learning

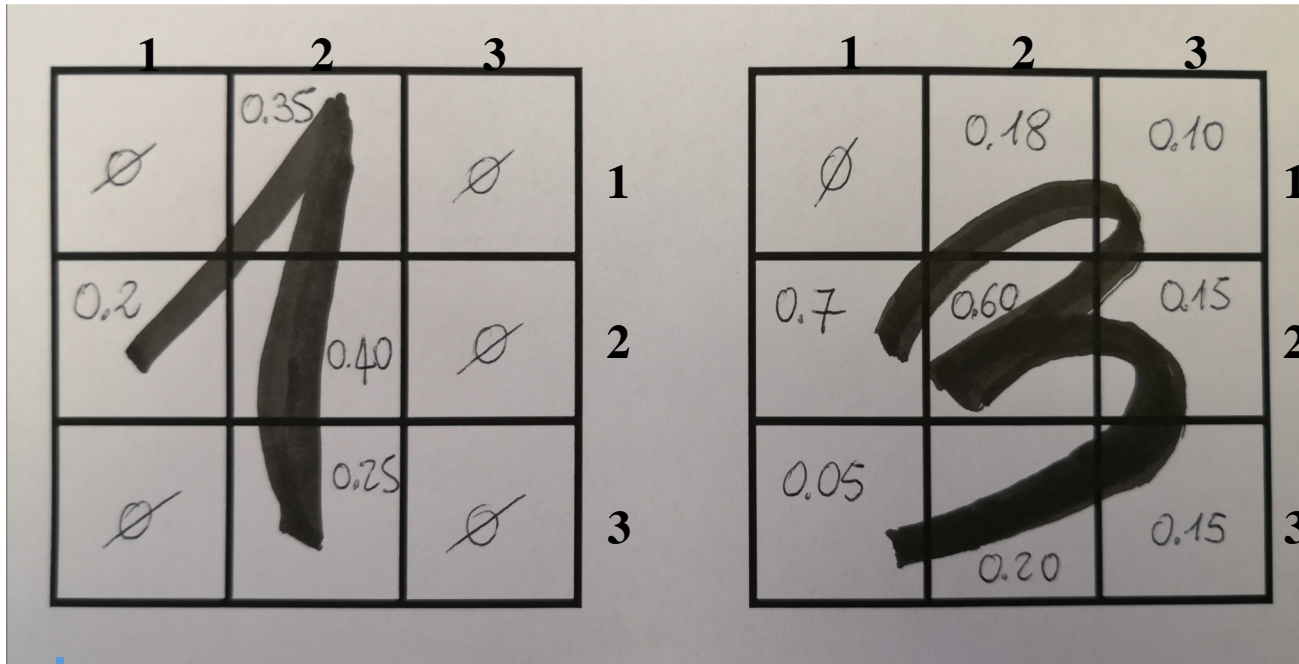
**Task:** Making a machine able to recognize whether the numeral is a 1 or a 3

- ML re-organizes the problem as a  $(y,X)$  **mapping** problem where  $(y,X)$  is a **(training) dataset**
- So, ML needs to generate a dataset able to map  $y$  (the true numeral) into  $X$  (the usual ways people write numerals 1 and 3 by-hand)
- The trick is to decompose the numeral into adjacent equally spaced pieces (pixels) and use them as predictors ( $X$ )



next slide





- Each single square is one **pixel**
- The **number inside each pixel** is the **percentage** of the pixel area covered by the ink

**Data representation of the two images**

Y	x1	x2	x3	x4	x5	x6	x7	x8	x9	
Image of "1" →	1	0	0.35	0	0.2	0.4	0	0	0.25	0
Image of "3" →	3	0	0.18	0.1	0.7	0.6	0.15	0.05	0.2	0.15



# Based on many images we aim to **predict new images**



	Y	x1	x2	x3	x4	x5	x6	x7	x8	x9
Image of "1" →	1	0	0.35	0	0.2	0.4	0	0	0.25	0
Image of "3" →	3	0	0.18	0.1	0.7	0.6	0.15	0.05	0.2	0.15
Image of "1" →	1	0	0.30	0	0.1	0.4	0	0.01	0.22	0
Image of "1" →	1	0	0.35	0.01	0.3	0.4	0	0	0.23	0
Image of "3" →	3	0	0.16	0.1	0.65	0.6	0.16	0.09	0.3	0.14
Image of "3" →	3	0	0.15	0.1	0.75	0.7	0.14	0.06	0.2	0.15
Image to predict →	?	0	0.30	0	0.1	0.4	0	0.01	0.22	0
Image to predict →	?	0	0.35	0	0.2	0.4	0	0	0.25	0


# Protocol of imaging **recognition**

1. Build a **training dataset** made of all the images
2. Use a **learner**, i.e. a mapping between  $\{Y\}$  and  $X=\{x_1, \dots, x_9\}$
3. As long as certain **configurations** of  $X$  are likelier to be associated to a specific  $Y$ , a proper learner **can learn how to link a new configuration** (i.e., a new image) to a **specific numeral**



More complex images recognition (faces, objects) follow a similar logic

# Face completion: predict the **lower** half of a face by knowing the **upper** half



*Olivetti faces dataset*

# **Limitations and failure of statistical learning**

# ML limitations

- **Ethics**

When decisions are taken by a machine, who is **responsible** for their effect?

- **Causality**

The mapping found out by a ML algorithm is based on **correlation, not on causality**

- **Scope**

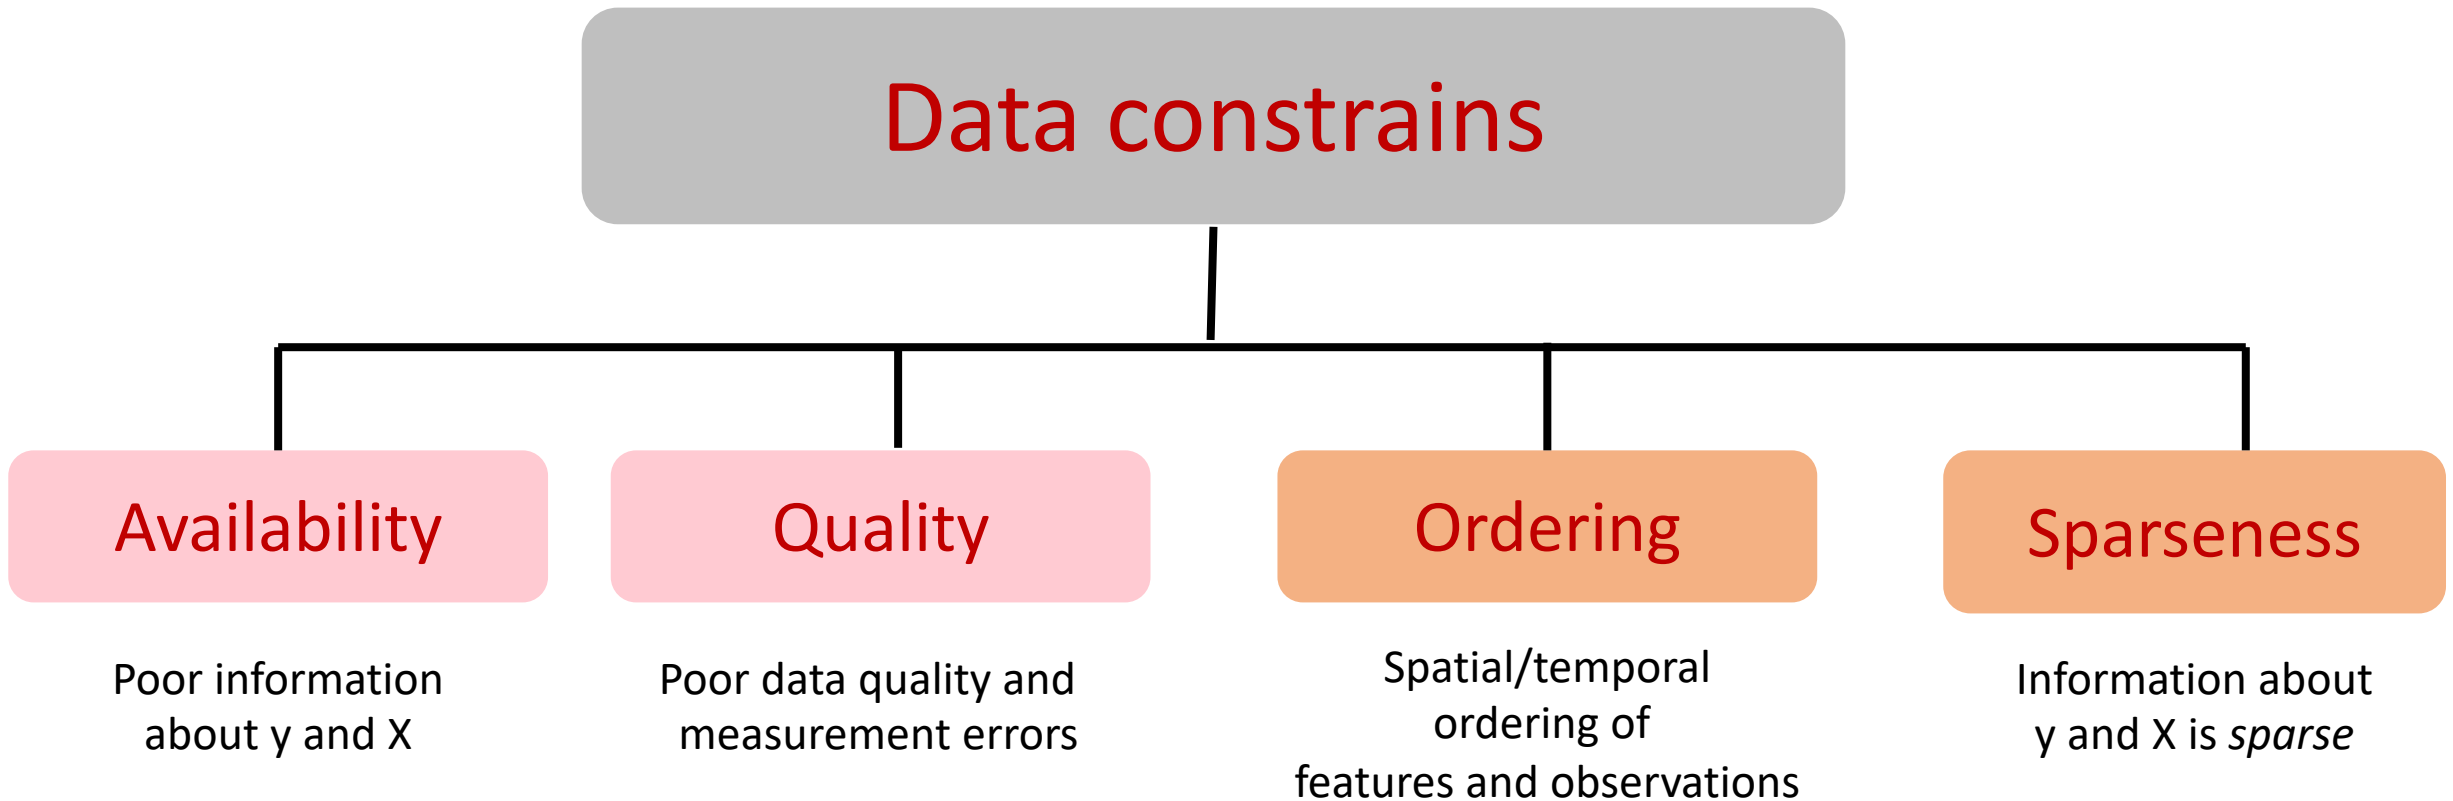
Should we use ML in whatever analytical context? Even when the aim is pure prediction, ML algorithms can skip **fundamental theory**: *“should we make weather forecasts without relying on Newton's laws of motion”?*

- **Interpretability**

There exists a fundamental **trade-off between predictability and interpretability**. In some research contexts interpretability is more important than predictability

# ML failure

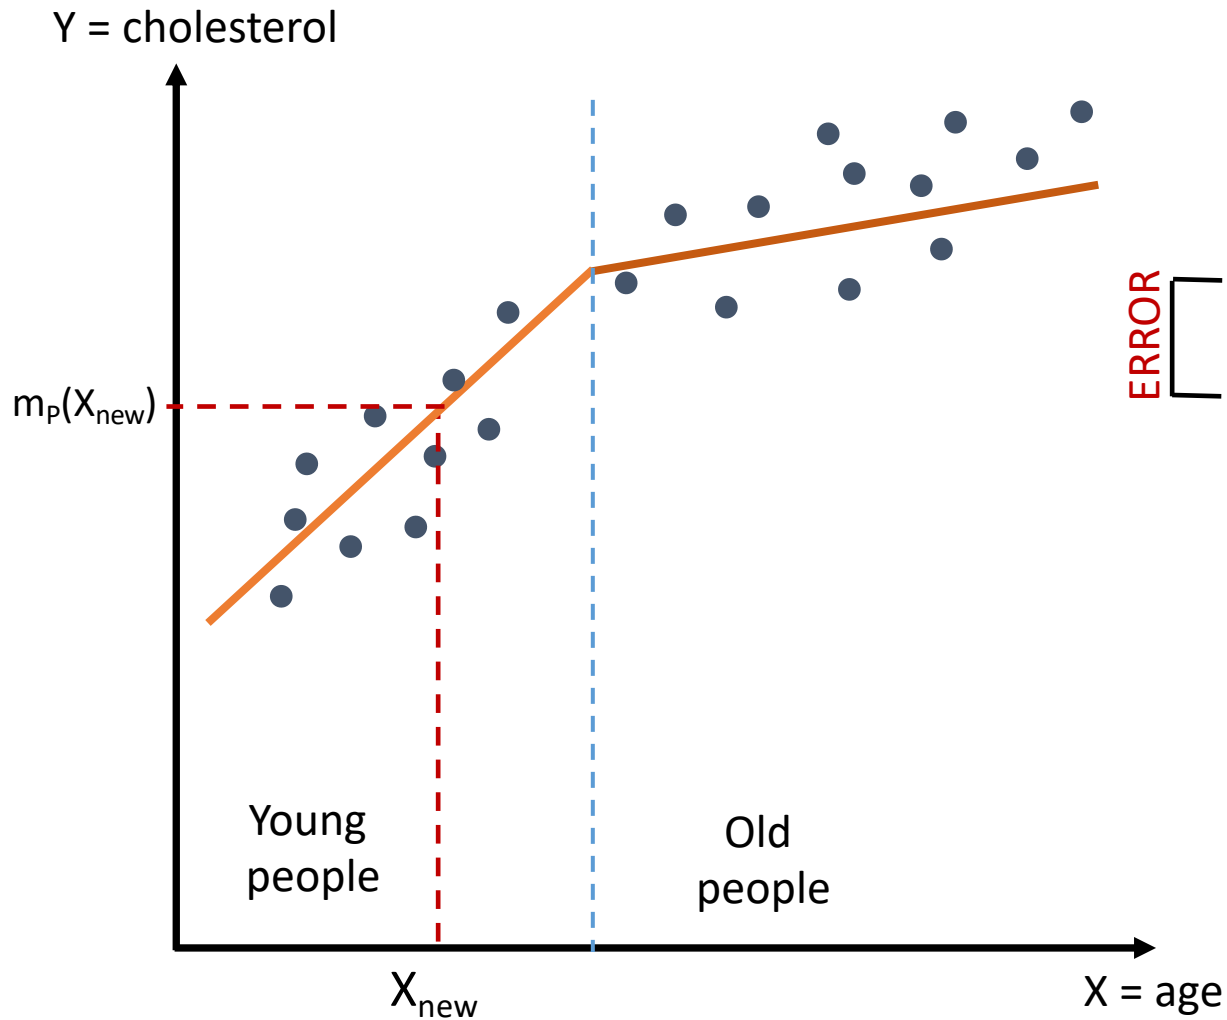
Can ML prediction fail? **Yes!** == > It may depend on **data constrains**



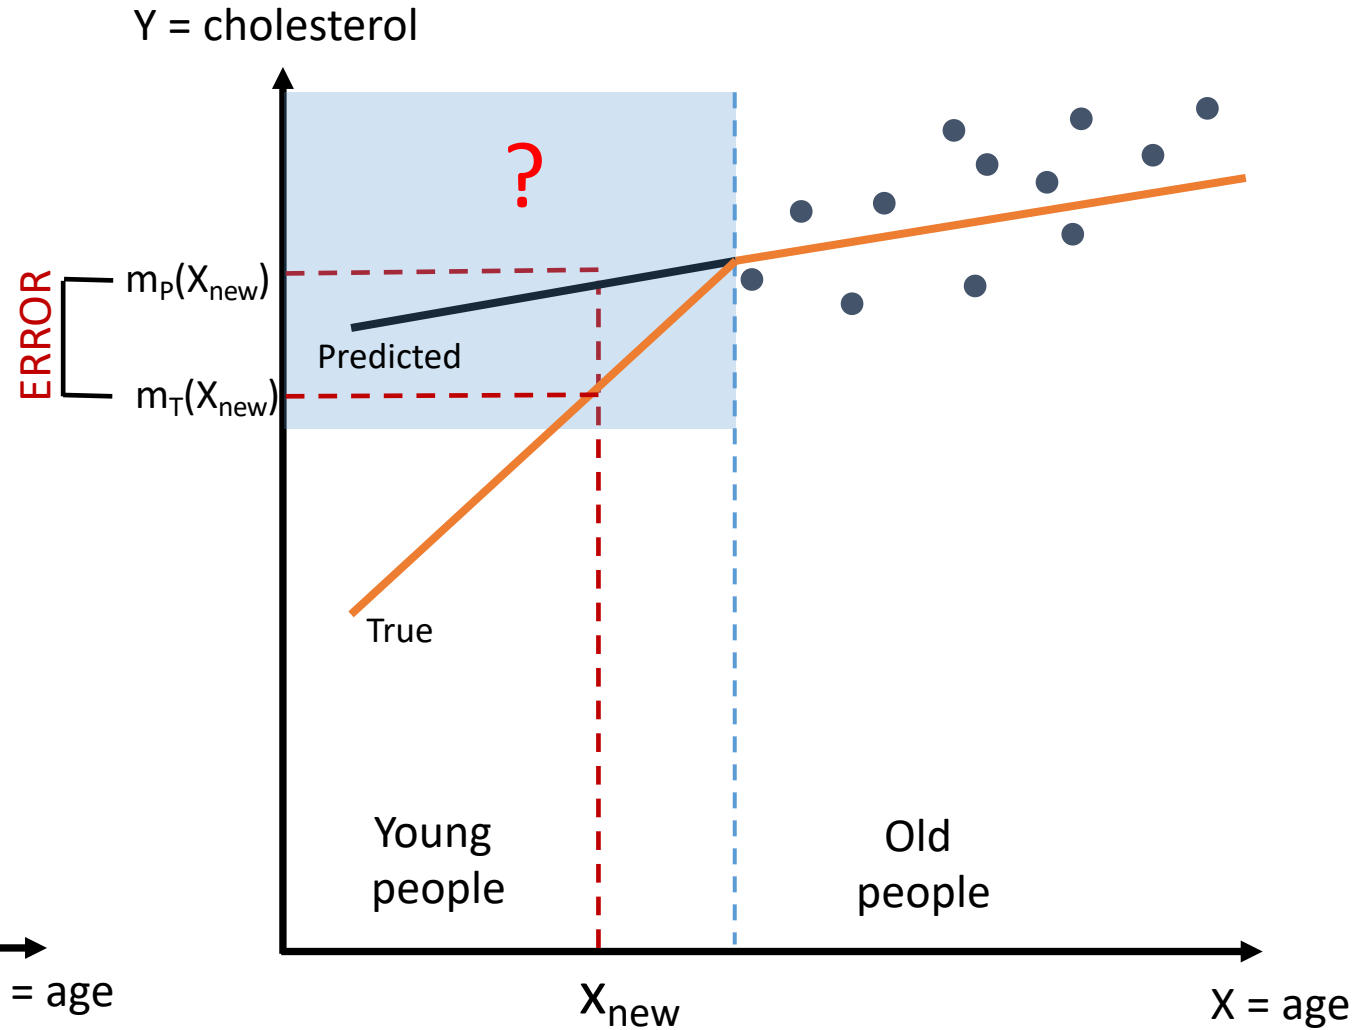
# The consequences of **data sparseness**

# Data sparseness weakens prediction

Low sparseness



High sparseness






# The consequences of **data unordering**

# Higher data ordering implies higher prediction

## Imaging recognition

### Face recognition



1. Each pixel is a variable in the dataset
2. Adjacent pixels are highly correlated

*This depends on the fact that, by nature, the nose is at the center of the face, with eyes on the left and right and forehead at the top, ....*

Similar to **spatial correlation** in statistics

## Natural Language processing

### One-token-ahead prediction



1. Find the more likely word after a **sequence of words**
2. Train over a frequency-context (e.g., *whatsup chats*)

*This depends on human beings  
**conventionality** / **Homophily***

Similar to **serial correlation** in statistics

# Lower data ordering implies lower prediction

- In the **social-sciences**, information is strongly **unordered**
- This explains why **deep learning** prediction is **poorly successful** for micro socio-economic data (data on workers, businesses, farms, etc.)
- Generally, **other ML methods** than neural networks do better with socio-economic data (*Random forests, Boosting, SVM*)

# Conclusions

- **AI matters:** revolutionizing human activities and society
- **Machine Learning:** today's leading AI paradigm
- **ML limits:** ethics, interpretability, and scope
- **ML failure:** *sparseness, poor ordering*



Can **symbolic AI** help ML in **failing** situations?

The **neuro-symbolic** approach seems a promising way to go  
(TBA in a next seminar .... 😊)

# Bibliography

- ❑ Bateson, G. 1979. *Mind and nature: a necessary unity*. London: Wildwood House.
- ❑ Cerulli, G. 2020. Improving econometric prediction by machine learning, *Applied Economics Letters*, Forthcoming.
- ❑ Dawkins, R. 1976. *The Selfish Gene*. Oxford University Press, Oxford.
- ❑ Gareth, J., Witten, D., Hastie, D.T., Tibshirani, R. 2013. *An Introduction to Statistical Learning : with Applications in R*. New York, Springer.
- ❑ Harari, Y.N. 2015. *Sapiens: a Brief History of Humankind*. New York:Harper.
- ❑ Hastie, T. , R. Tibshirani , and J. Friedman. 2001. *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. Berlin: Springer-Verlag.
- ❑ Kaplan, J. 2016. *Artificial Intelligence: What Everyone Needs to Know*. Oxford University Press: Oxford.
- ❑ Kleinberg, et al. 2015. Prediction Policy Problems. *American Economic Review*, 105 (5): 491-95.
- ❑ Magrini M. 2017. *Cervello. Manuale dell'utente. Guida semplificata alla macchina più complessa del mondo*. Giunti.
- ❑ Mullainathan, S. , and J. Spiess. 2017. Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, 31 (2): 87–106.
- ❑ Noyan, M.A. 2018. *Plato, Aristotle and Machine Learning*. <https://towardsdatascience.com/plato-aristotle-and-machine-learning-521c3252cbd1>
- ❑ Raschka, S., Mirjalili, V. 2019. *Python Machine Learning*. 3rd Edition, Packt Publishing.
- ❑ Simon, H.A. 1962. The Architecture of Complexity. *Proceedings of the American Philosophical Society*, Vol. 106, No. 6 ( , pp. 467-482.
- ❑ Singh R. 2019. *The Rise and Fall of Symbolic AI*. <https://towardsdatascience.com/rise-and-fall-of-symbolic-ai-6b7abd2420f2>
- ❑ Van der Laan, M. J. , and S. Rose. 2011. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Heidelberg: Springer.
- ❑ Varian, H. R. 2014. Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, 28 (2): 3–28.
- ❑ Zhou, Z. H. 2012. *Ensemble Methods: Foundations and Algorithms*. London, UK: CRC Press.

## PART 2

# AI research at CNR-IRCREES: Optimal Policy Learning (OPL)

# What AI research at CNR-IRCRES?

**Statistical learning for optimal  
socio-economic policy design**



**Optimal Policy Learning  
(OPL)**

# OPL definition

- **What is policy learning?**

Process of improving program **welfare** achievements by re-iterating similar policies over time

- **Optimal treatment assignment**


Policymakers can **optimally fine-tune the treatment assignment** of a prospective policy using the results from an RCT or observational study. Assignment rules depends on the **class of policies** considered (here we focus on threshold-based and linear-combination policies)

- **Maximizing constrained welfare**

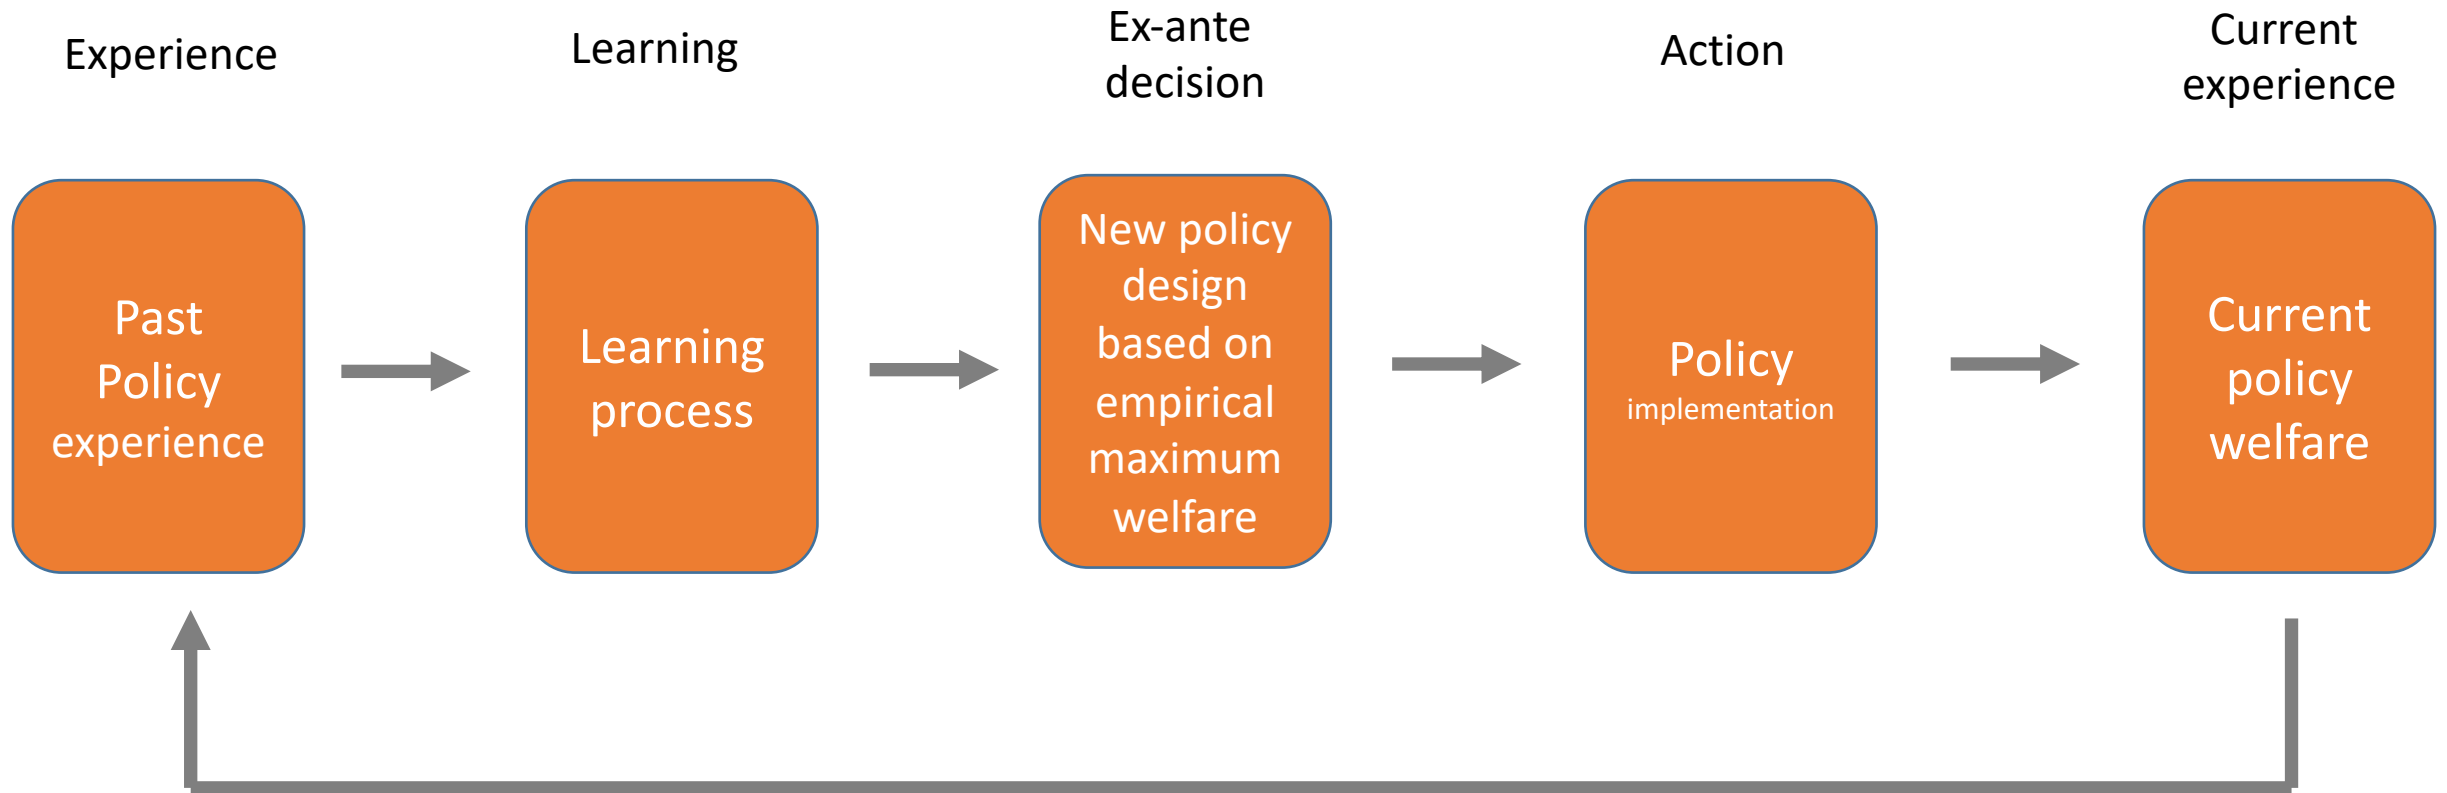
The policymaker hardly manage to reach the best solution (**unconstrained maximum welfare**) because of institutional/economic contains of various sort



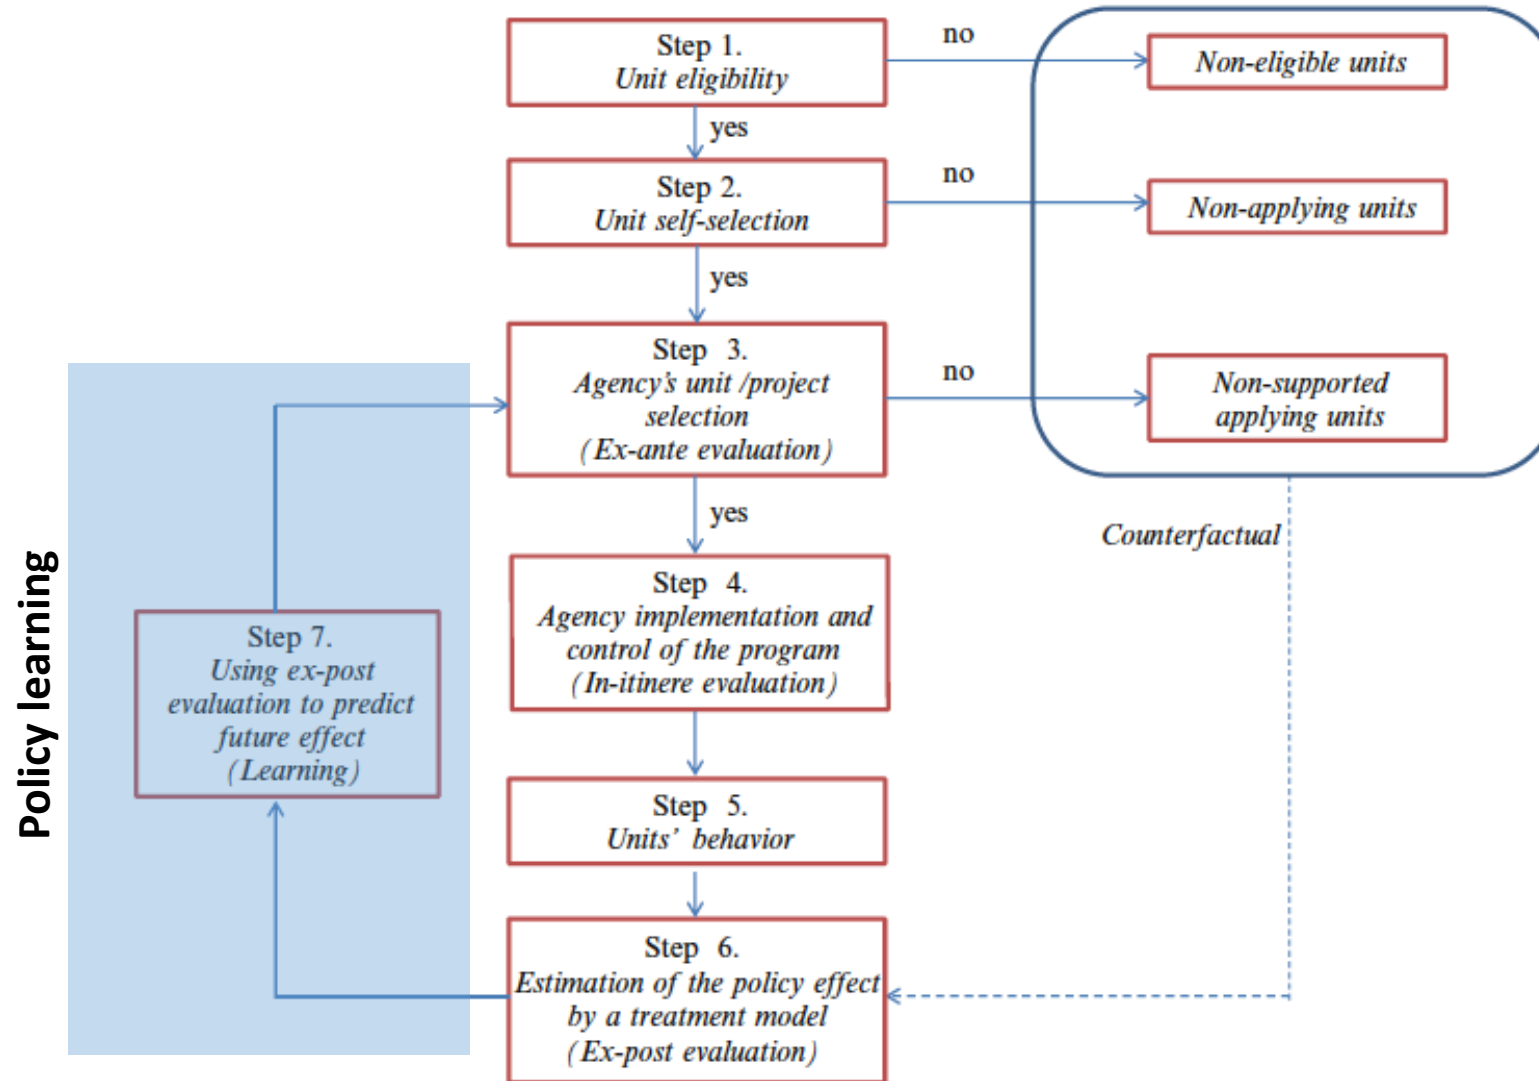
# Policy as selection problem



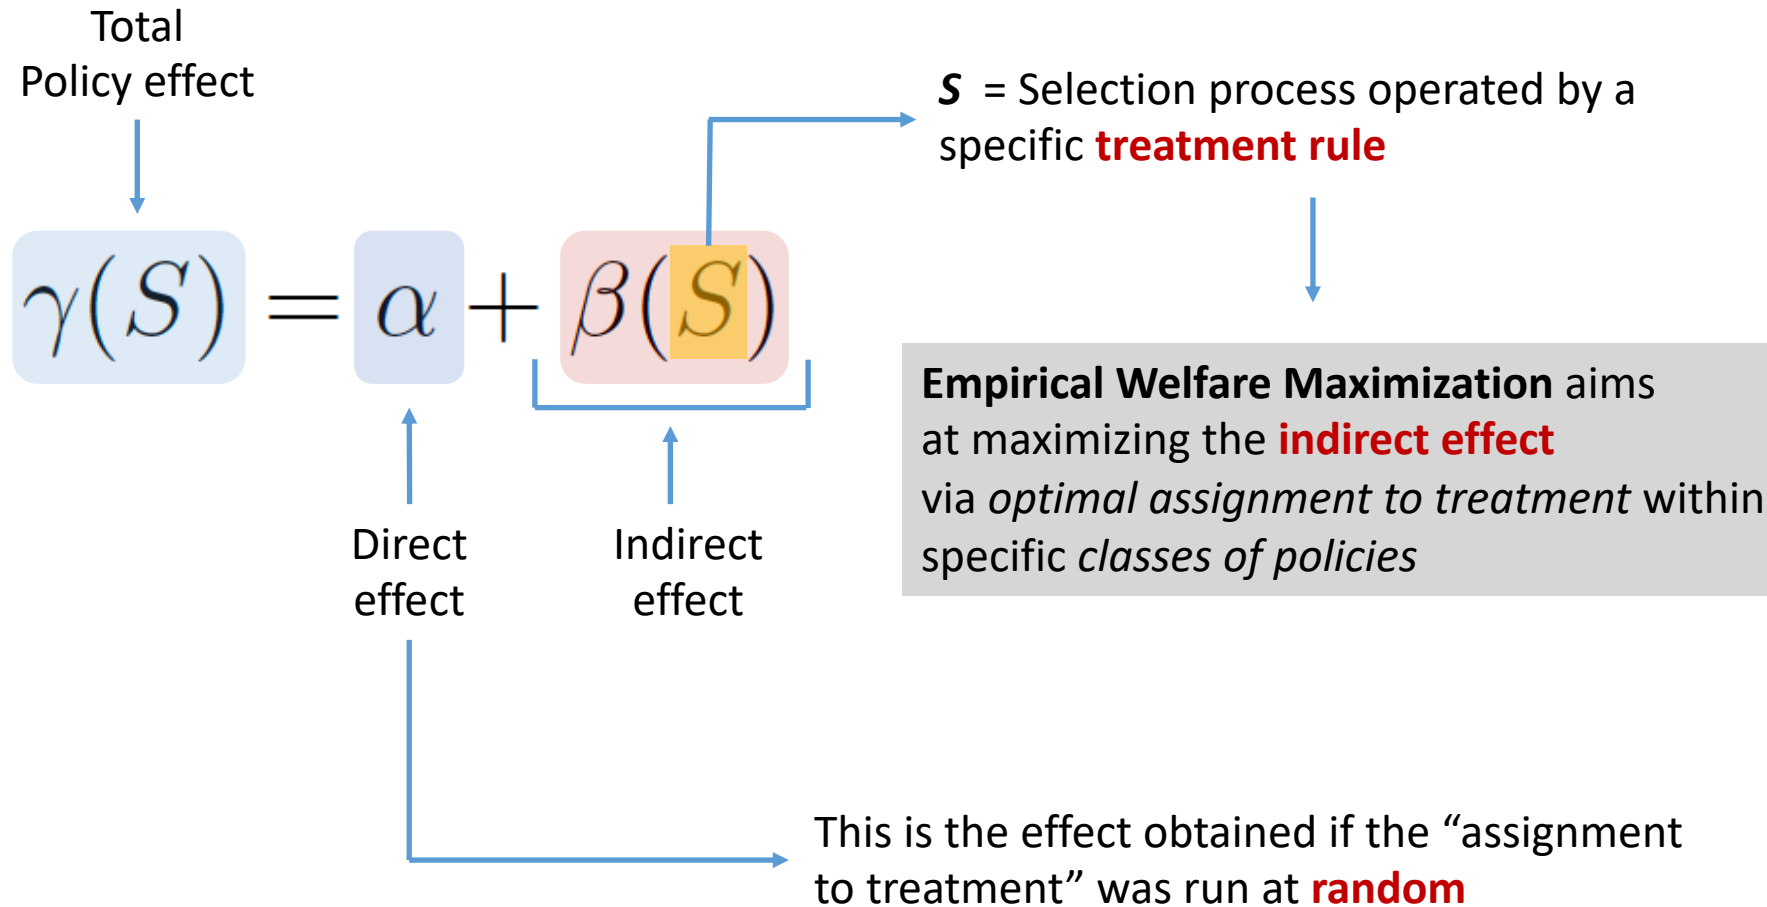
# Learning policy from **experience**



# Policy learning within policy EVALUATION cycle



# Policy direct and indirect effect



# Optimal treatment assignment

Let  $X$  be an individual's vector of characteristics,  $Y$  an outcome of interest,  $T = \{0, 1\}$  a binary treatment. A policy assignment rule  $\mathcal{G}$  is a function mapping  $X$  to  $T$ , specifying which individuals are or are not to be treated:

$$\mathcal{G} : X \rightarrow T$$

Define the (population) policy conditional average treatment effect as:

$$\tau(X) = E(Y_1|X) - E(Y_0|X)$$

where  $Y_1$  and  $Y_0$  represent the two potential outcomes of the policy, and  $E_X[\tau(X)] = \tau$  the average treatment effect.

# Optimal treatment assignment

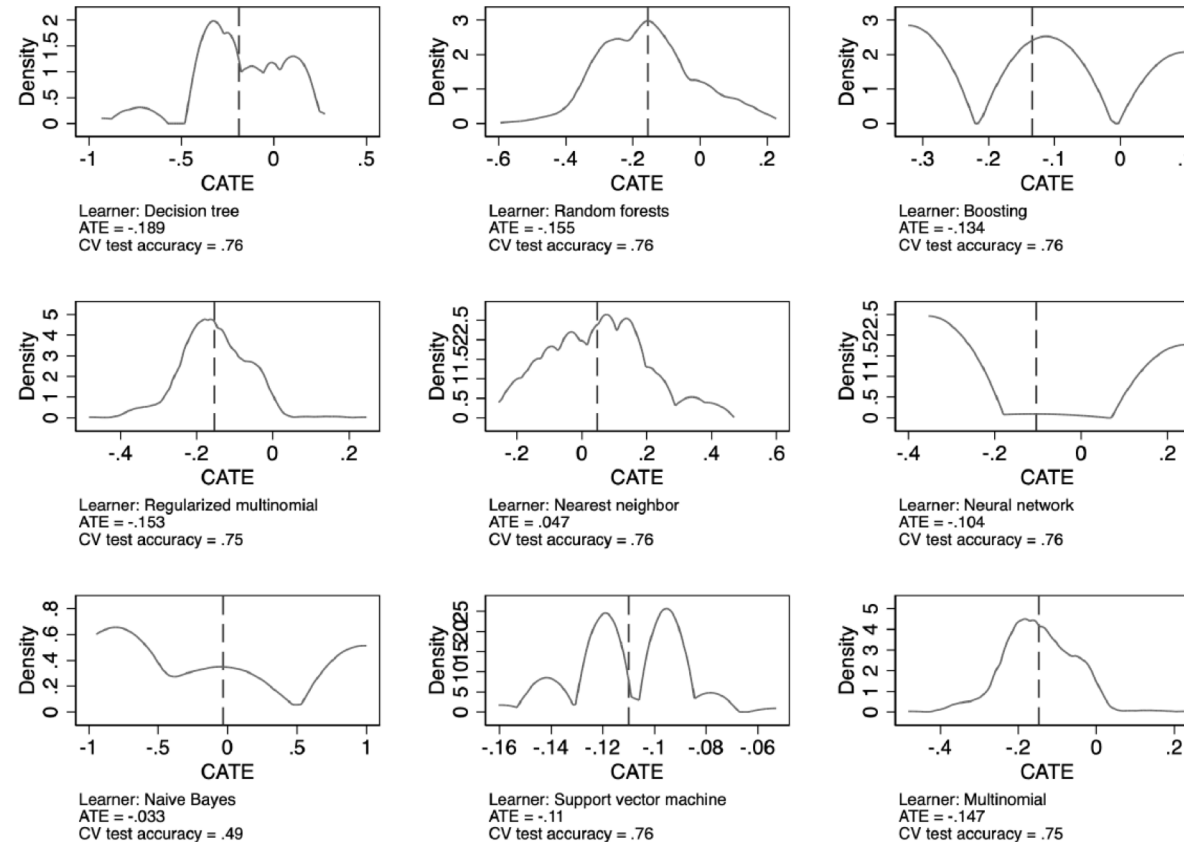
Under **selection-on-observables**, we know that:

$$\tau(X) = E(Y|X, T = 1) - E(Y|X, T = 0)$$

These two **conditional expectations** are **identified** by data. Whatever **ML algorithm** can be used for estimation (Boosting, Random forests, Neural networks, Nearest neighbor, etc.)

Extension to **selection-on-unobservables** straightforward

# ML estimation of $\tau(X)$



Estimation of the **distribution** of the **conditional average treatment effects (CATE)** using the ML methods implemented via **c\_ml\_stata\_cv** (Cerulli, 2022). Note: dashed vertical line indicates the **average treatment effect (ATE)**.

# Optimal unconditional welfare

The estimated policy actual total effect (or *welfare*)

$$\widehat{W} = \sum_{i=1}^N T_i \cdot \hat{\tau}(X_i)$$

and the estimated policy *unconstrained* optimal total effect (or *unconstrained maximum welfare*) as:

$$\widehat{W}^* = \sum_{i=1}^N \hat{T}_i^* \cdot \hat{\tau}(X_i)$$

where:

$$\hat{T}_i^* = \mathbf{1}[\hat{\tau}(X_i) > 0]$$

is the estimated optimal unconstrained policy assignment.

The difference between the estimated (unconstrained) maximum achievable welfare and the estimated welfare associated to the policy actually run is called *regret*, and it is defined as:

$$\widehat{regret} = \widehat{W}^* - \widehat{W}$$



# NAÏVE OPTIMAL SELECTION

1. Given  $\{X, Y, T\}$  from an already-implemented policy: estimate the **idiosyncratic effect**  $\tau(X)$ . This means we have learnt the mapping:

$$X \rightarrow \tau(X) \quad (\textit{learning from experience})$$

2. Consider a prospective second policy round with a new eligible set  $\{X'\}$ , and compute the learnt  $\{\tau(X')\}$  over  $X'$ .
3. Rank individuals so that:  $\tau(X_1') > \tau(X_2') > \tau(X_3') > \dots > 0$ .
4. Given a monetary budget  $C$  and a unit cost  $c_i$ , find  $N_1^*$ :

$$\sum_{i=1}^{N_1^*} c_i = C$$

# Constrained welfare maximization - 1

- ❑ Eligibility, budget, ethical, or institutional constraints make policymakers unable to implement the *optimal unconstrained policy assignment*
- ❑ They are obliged to rely on a constrained assignment rule selecting treated units according to their characteristics
- ❑ The welfare thus obtained may **drop down**
- ❑ Policymakers can however produce the **largest feasible constrained welfare**

# Constrained welfare maximization - 2

- The policymaker wants to treat only “young” people
- In theory, he can continue to use the naïve approach, by excluding from treatment all the individuals with age smaller than a certain age  $A^*$
- The problem is that different  $A^*$  can induce different level of welfare
- The problem becomes that of **choosing  $A^*$  to maximize the effect/welfare**

# Example: constrained univariate threshold-based policy

- The policymaker wants to treat only “young” people
- In theory, he can continue to use the naïve approach, by excluding from treatment all the individuals with age smaller than a certain age  $A^*$
- The problem is that different  $A^*$  can induce different level of welfare
- The problem becomes that of **choosing  $A^*$  to maximize the effect/welfare**

# Policy classes

There exist however several **classes of policies** used by policymakers to select in a constrained decision context. The most popular are:


- Threshold-based

- Linear combination


- Fixed-depth decision trees

# Policy classes


Threshold-based



Linear combination



Fixed-depth tree




Legend:

 Decision boundary

 Selection area

2-depth tree 



# Optimal **constrained** welfare

→ The corresponding **welfare** is a function of  $c_x$ :

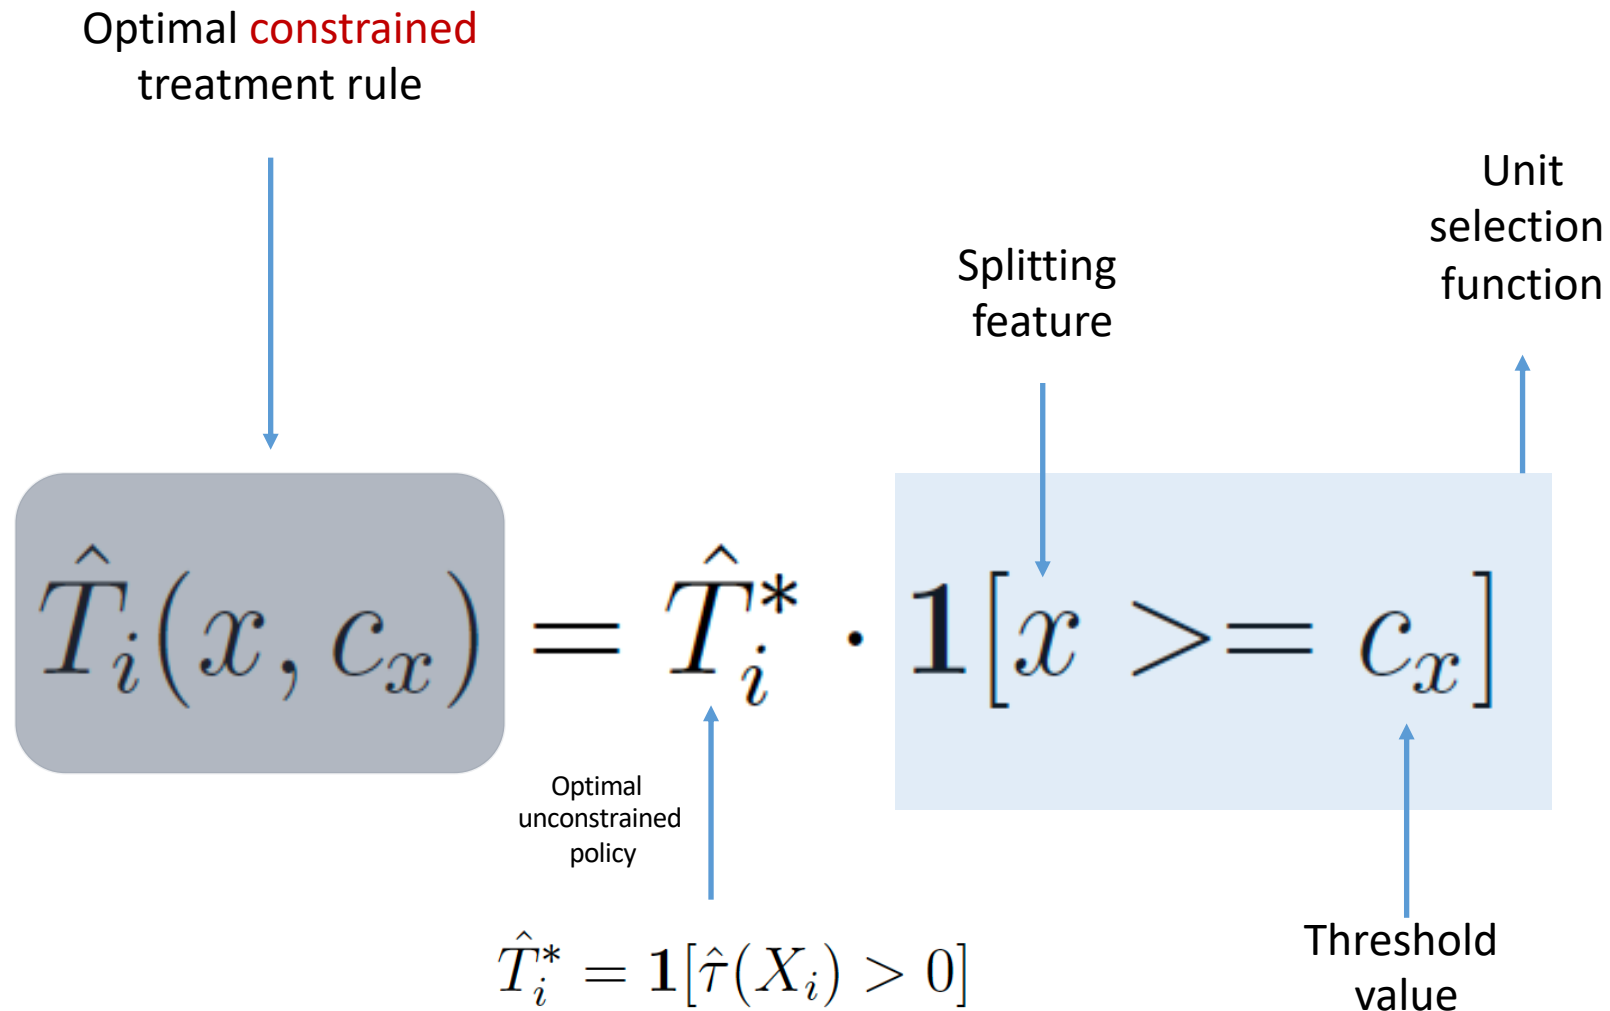
$$\widehat{W}(x, c_x) = \sum_{i=1}^N \widehat{T}_i(x, c_x) \cdot \tau(X_i)$$

We define the optimal choice of the threshold  $c_x$  as the one maximizing  $\widehat{W}(x, c_x)$  over  $c_x$ :

$$c_x^* = \operatorname{argmax}_{c_x} [\widehat{W}(x, c_x)]$$

If  $c_x^*$  exists, the estimated optimal constrained welfare will thus be equal to  $\widehat{W}(c_x^*)$ .

# Threshold-based policy





# Optimal **constrained** treatment rule (*multivariate case*)

Policymakers rely on two or more selection indicators

$$\hat{T}_i(c_x, c_z) = \hat{T}_i^* \cdot \mathbf{1}[x \geq c_x] \cdot \mathbf{1}[z \geq c_z]$$

Optimal unconstrained policy

Splitting feature x

Threshold Value for x

Splitting feature z

Threshold Value for z

# Implementation algorithm

---

**Procedure.** Threshold-based optimal policy assignment

---

1. Suppose to have data from an RCT or from an observational study consisting of the information triple  $(Y, X, T)$  available for every unit involved in the program.
  2. Run a quasi-experimental method with observable heterogeneity, estimate  $\tau(X)$ , and compute the (estimated) actual total welfare of the policy  $\widehat{W}$ .
  3. Identify the estimated optimal unconstrained policy  $\widehat{T}^*$ , and compute  $\widehat{W}^*$ , i.e. the estimated maximum total welfare achievable by the policy, and estimate the regret as  $\widehat{W}^* - \widehat{W}$ .
  4. Consider an estimated constrained selection rule  $\widehat{T}(x, c)$  based on a given set of selection variables,  $x$ , and related thresholds,  $c$ , and define the estimated maximum constrained welfare as  $\widehat{W}(x, c)$ .
  5. Build a greed of  $K$  possible values for  $c \in \{c_1, \dots, c_K\}$ , compute the optimal vector of thresholds  $c_{k^*}$  and the corresponding maximum estimated welfare  $\widehat{W}(x, c_{k^*})$  thus achieved.
-

# Linear combination policy (bivariate case)

Generates a **score** to compare with a threshold

$$\hat{T}_i(c_1, c_2, c_3) = \hat{T}_i^* \cdot \mathbf{1}[c_1 x_1 + c_2 x_2 \geq c_3]$$

Optimal unconstrained policy

score

threshold

The diagram illustrates the components of the linear combination policy equation. A blue arrow points from the title 'Linear combination policy' to the equation. Another blue arrow points from the text 'Generates a score to compare with a threshold' to the indicator function in the equation. A third blue arrow points from the text 'Optimal unconstrained policy' to the term  $\hat{T}_i^*$ . A blue bracket under the expression  $c_1 x_1 + c_2 x_2$  is labeled 'score', and another blue bracket under  $c_3$  is labeled 'threshold'.

# — Application

**DATA:** National Supported Work Demonstration (NSWD), an RCT by LaLonde (1986).

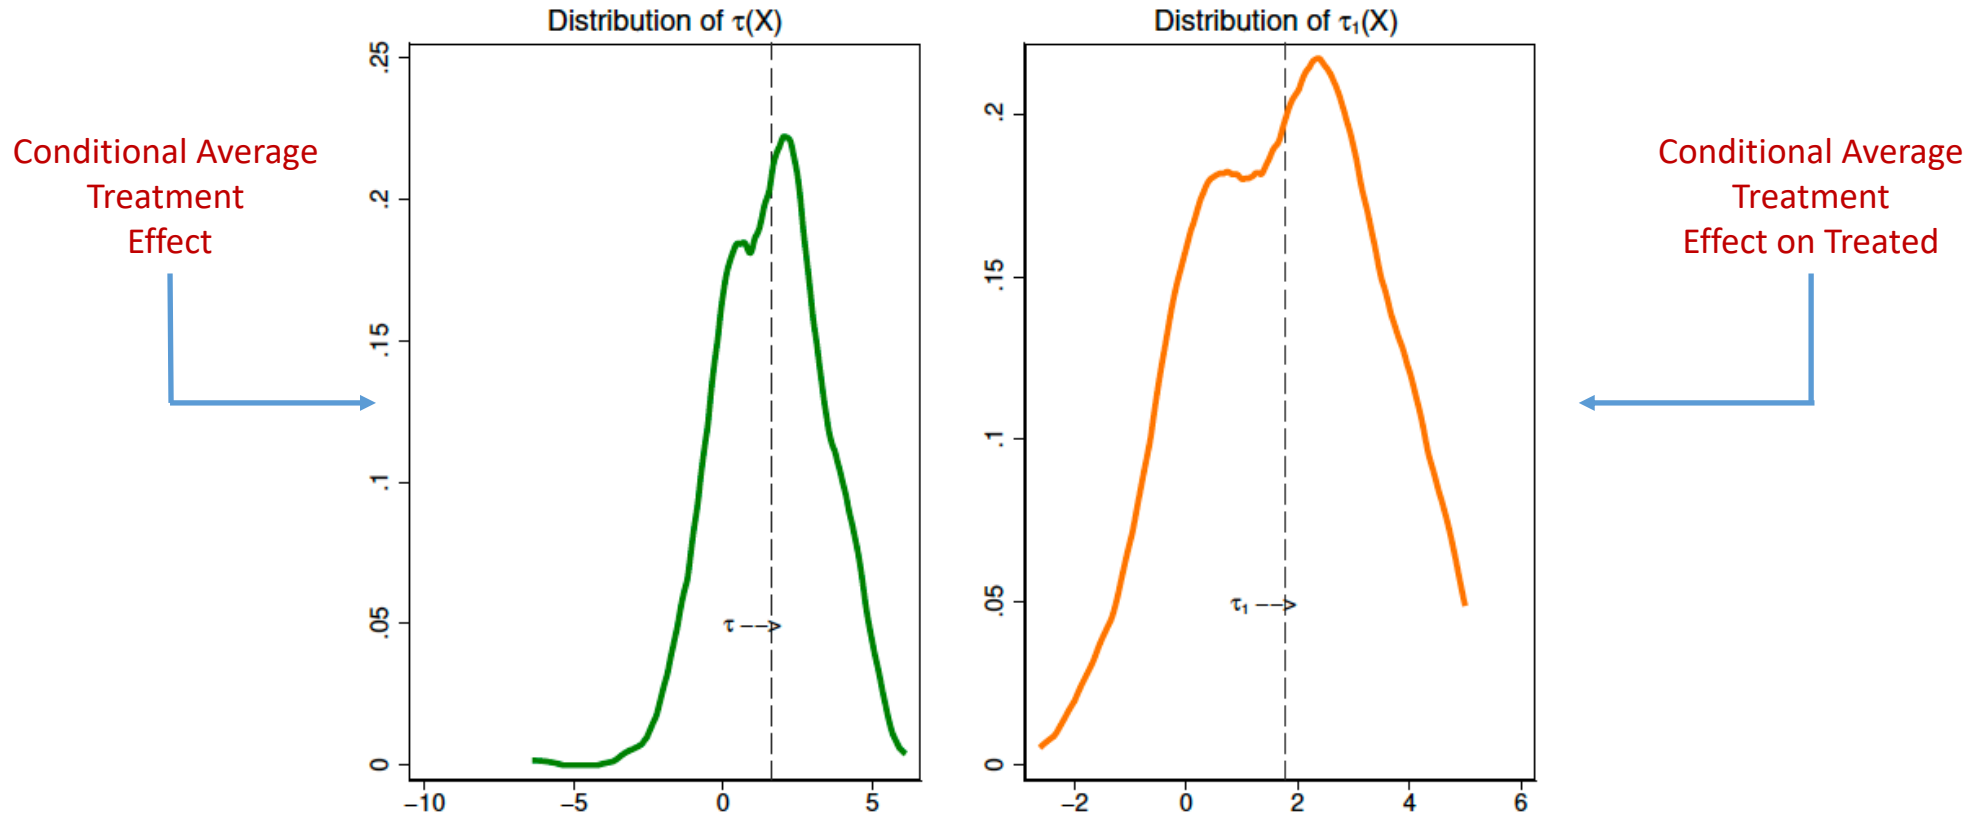
**TARGET:** Effect of a 1976 job training program on people real earnings in 1978

**CONTROLS:** age, race, educational attainment, previous employment condition, real earnings in 74 and 75



# Estimation of $ATE(x)$ and $ATET(X)$

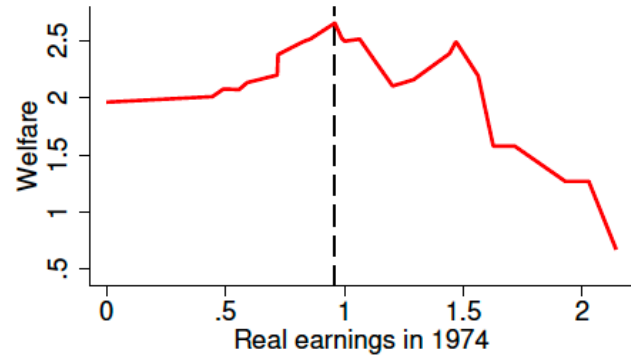
Figure 1: Distribution of  $\hat{\tau}(X)$  and  $\hat{\tau}_1(X)$ . Program: National Supported Work Demonstration (NSWD). Data: LaLonde (1986). Target variable: Real earnings in 1978. Estimation technique: Regression-adjustment (with observable heterogeneity).




# Constrained welfare maximization (univariate)

Figure 2: Computation of the policy optimal selection threshold in univariate cases. Program: National Supported Work Demonstration (NSWD). Data: LaLonde (1986). Target variable: real earnings in 1978. Univariate selection variables: real earnings in 1974, age, and educational attainment.

$$\text{AWG} = 2.65 - 1.76 = 0.89$$



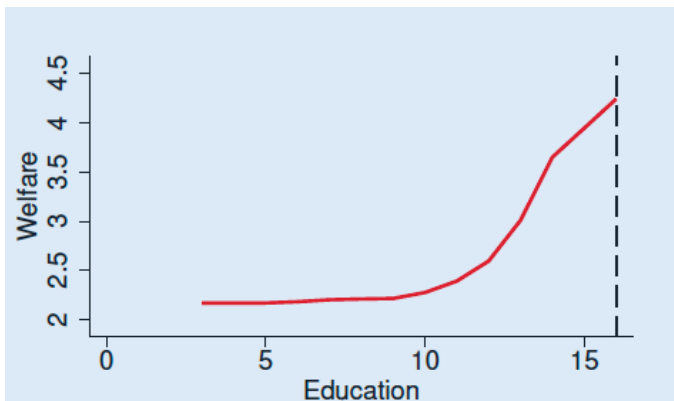
Optimal threshold = .96  
Optimal average welfare = 2.65  
Number of treated units = 108 out of 443



Optimal threshold = 42  
Optimal average welfare = 2.85  
Number of treated units = 16 out of 443

$$\text{AWG} = 2.85 - 1.76 = 1.09$$

$$\text{AWG} = 4.24 - 1.76 = 2.48$$




Optimal threshold = 16  
Optimal average welfare = 4.24  
Number of treated units = 0 out of 443

Monotonicity of welfare on educational attainment

Reference ATET = 1.76  
AWG = Average Welfare Gain

# Constrained welfare maximization (Bivariate)

Figure 3: Computation of the policy optimal decision boundary in the bivariate case. Program: National Supported Work Demonstration (NSWD). Data: LaLonde (1986). Target variable: real earnings in 1978. Bivariate selection variables: real earnings in 1975 and age.



# Empirical welfare maximization: relevant issues

## 1. Monotonicity

Welfare increases monotonically with a feature  
=> *too few to treat or too many to treat*

## 2. Sparseness

$X'$  comes from a *different joint distribution* than  $X$


**Trade-offs** arising in this case, so the best to do is offering the policymaker a “**menu**” of possible treatment choices given, for example, a pre-fixed budget




# SPARSENESS

the distribution of  $X$  and  $X'$  have **low overlap**

High sparseness



Low sparseness



# A SOLUTION TO MONOTONICITY

Trade-offs and the “menu-strategy”

## EXAMPLE

Computation of policy optimal decision boundaries in the bivariate case, when one of the two selection variables (age) is fixed at its optimal threshold, and the threshold of the other variable (education) is varying. Program: National Supported Work Demonstration (NSWD). Data: LaLonde (1986). Target variable: real earnings in 1978. Bivariate selection variables: age and educational attainment.

**AGE** —————> set at its optimal level

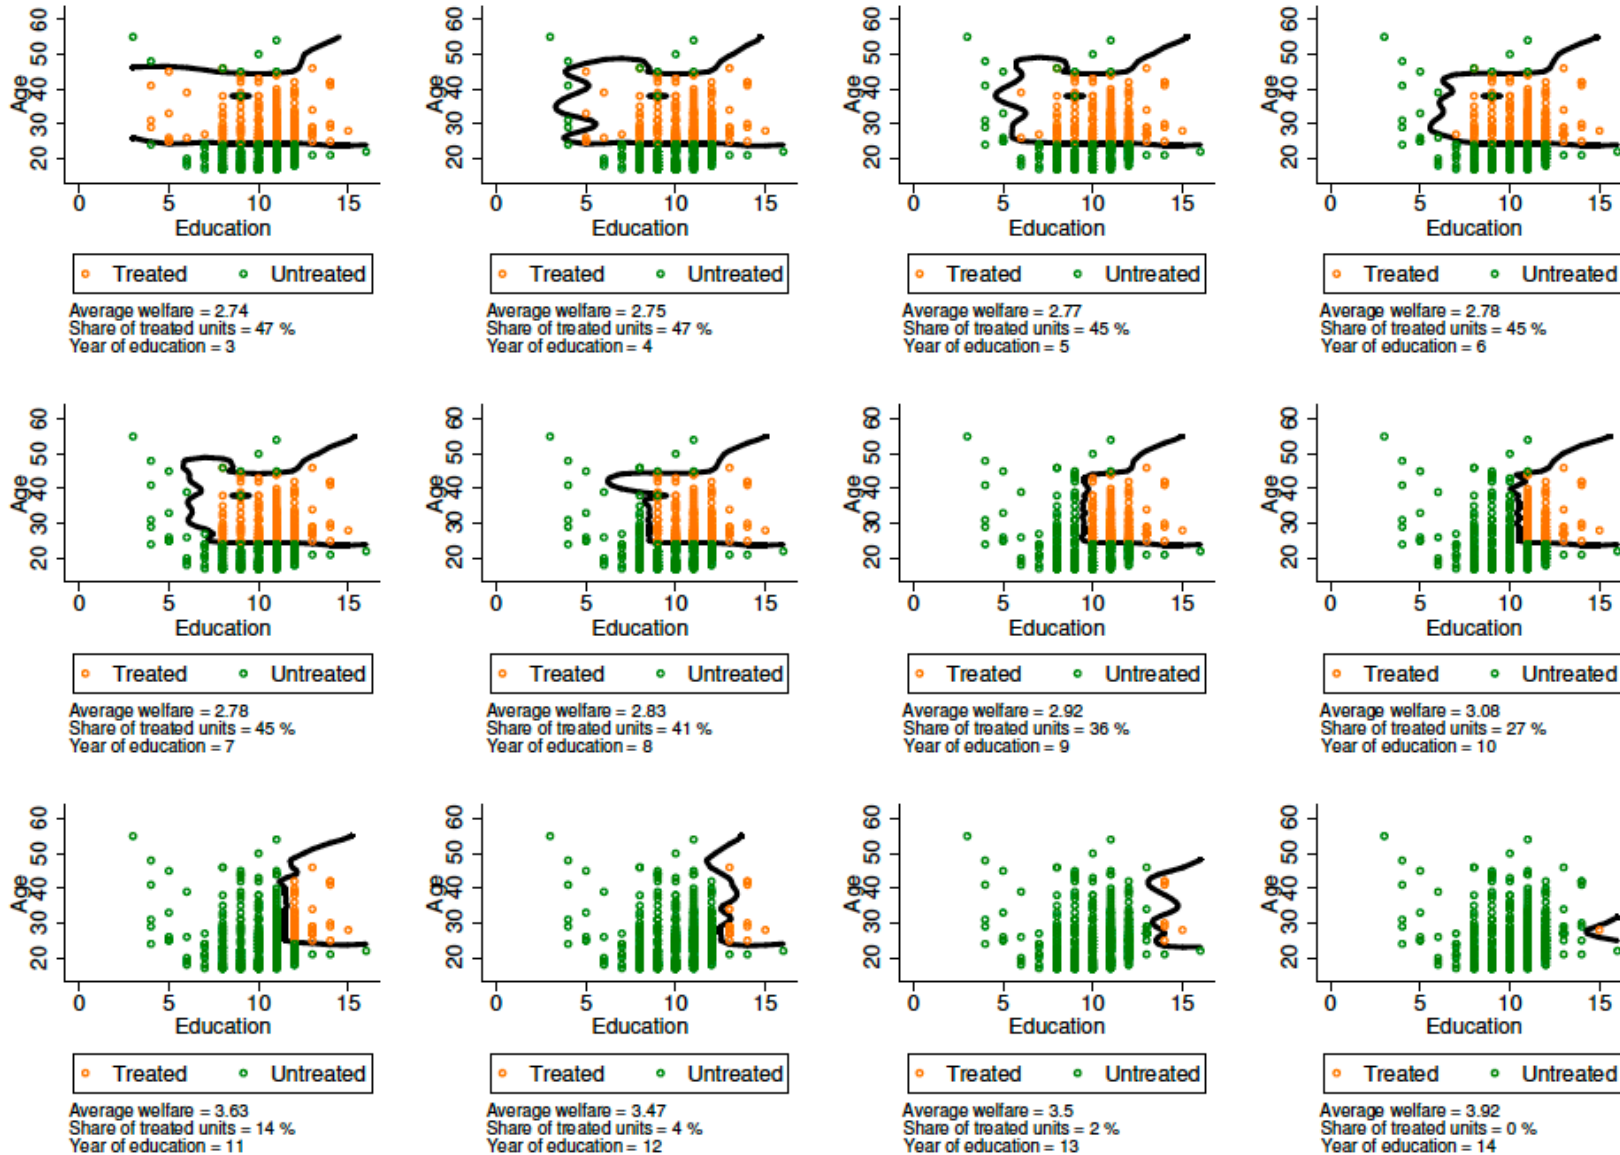
**EDUCATION** —————> free to vary



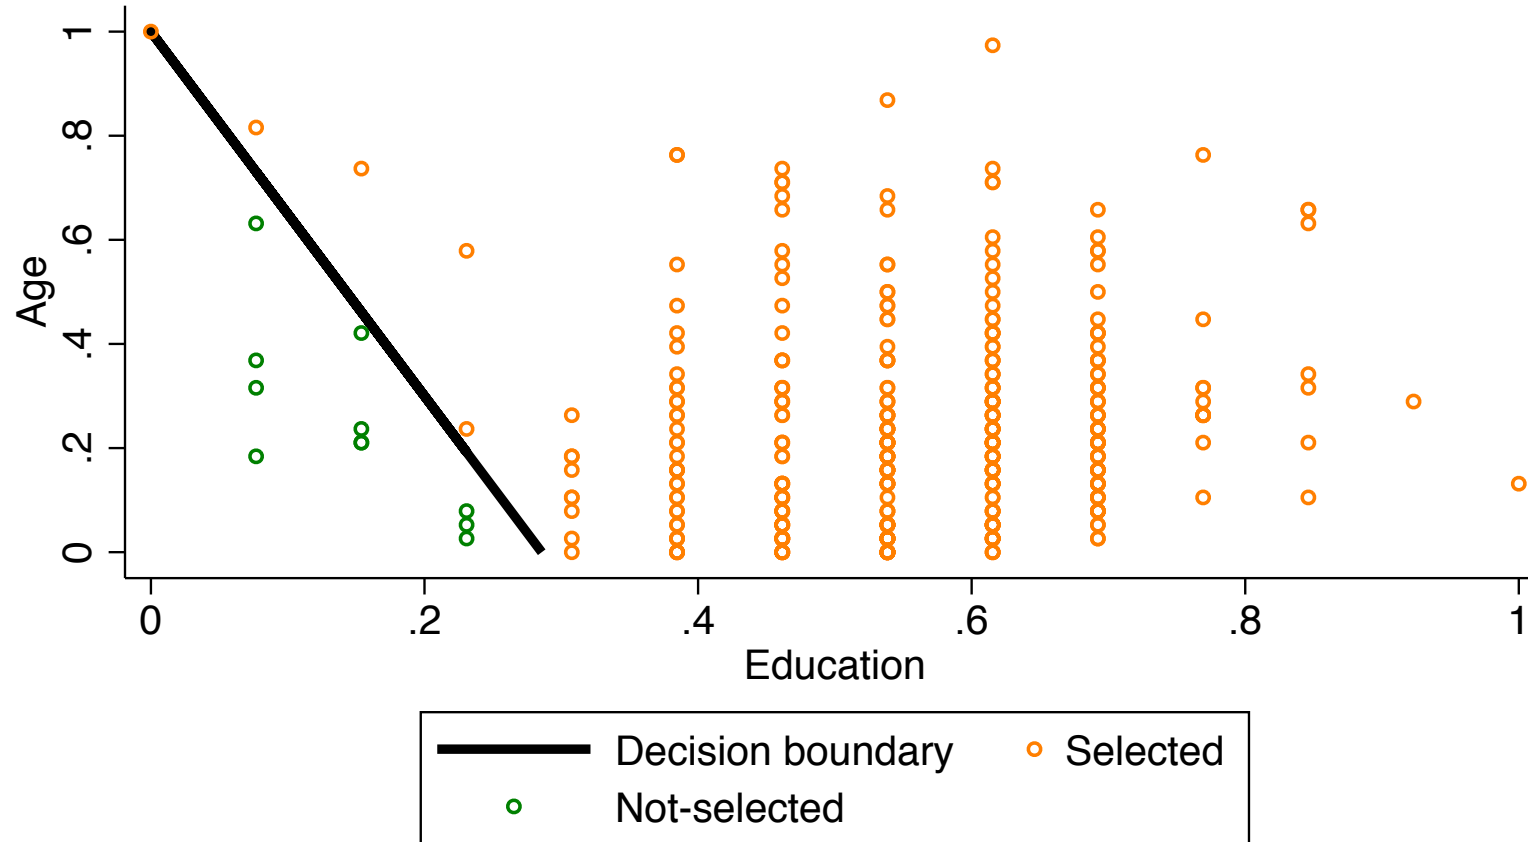
Feature plagued by monotonicity



# Trade-offs and the “menu-strategy”



# OPTIMAL SELECTION WITH A Linear combination policy



Total optimal welfare = 748  
Total oracle welfare = 764  
Regret (absolute) = 15.53  
Regret (%) = 2.03  
Average welfare = 2.24  
Average oracle welfare = 2.23  
Share of treated units = 75 %



# — SOFTWARE

We formed a research group for **OPL software implementation**:

## **Stata**

Cerulli (CNR), **opl** command

## **R**

Guardabascio (Perugia University) and Brogi (Istat)

## **Python**

De Fausti (Istat)



# Conclusion

- **Optimal Policy Learning**: AI/ML new frontier of policy design
- **Machine Learning** algorithms for estimating policy effects
- Generalization to many **policy classes**
- Producing **Stata/R/Python** software platforms for OPL
- CNR-IRCRES leading OPL development
- Future development: **DATA-DRIVEN DECISION MAKING** (within which OPL is a subset of model)