Computational topology for safe, reliable, explainable and green Artificial Intelligence

Javier Perera-Lago

14th May 2025

REXASI-PRO

REXASI PRO

REliable & eXplainable Swarm Intelligence for People with Reduced mObility

REXASI-PRO partners

REXASI-PRO | Partners

Participant No. *	Participant organisation name
1 (Coordinator)	Spindox Labs
2	Italian National Council of Research
3	Deutsches Forschungszentrum für Künstliche Intelligenz
4	Dalle Molle Institute for Artificial Intelligence
5	ROYAL HOLLOWAY AND BEDFORD NEW COLLEGE
6	V-Research
7	AITEK
8	UNIVERSIDAD DE SEVILLA
9	Hovering Solution
10	EURONET
11(Subcontracting)	Scuola di Robotica (Ethics)

REXASI-PRO objectives

Use Cases

1. Navigation in crowded environments

2. Flying robot mapping

3. Collaborative navigation

REXASI-PRO tasks

The partners were divided into 8 Work Packages (WPs). The Cimagroup research team was mainly involved in

WP6: Decision Science and Topology-based methods for Greener AI

Specifically in the tasks:

T6.2: Topology-based energy consumption optimization of Pedestrian Detection algorithm

T6.3: Topology-based optimization of robot fleet behavior

T6.2

Topology-based energy consumption optimization of Pedestrian Detection algorithm

Artificial Intelligence: the training problem

Machine Learning models depend on a set of parameters that need to be adjusted. The setting or *learning* of the optimal parameters requires a lot of real-world data.

Nowadays, we have more and more sophisticated models and more massive data sets. Because of this, the costs derived from developing new AI are growing continually.

Increasing computations in AI

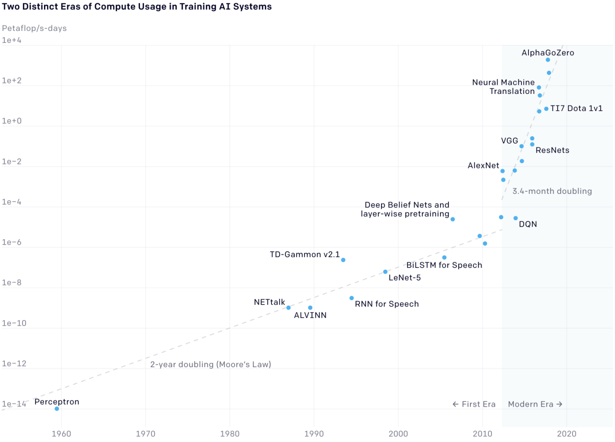

Increasing computations in AI

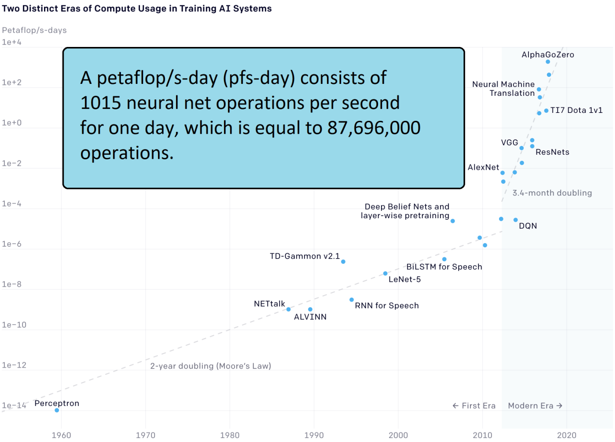

A petaflop/s-day (pfs-day) consists of 1015 neural net operations per second for one day, which is equal to 87,696,000 operations.

*Chart taken from the OpenAI blog: AI and compute

Red AI vs Green AI

- Red AI: AI research that seeks to improve the performance of models through the use of massive computational power without taking costs into account.
- Green AI: AI research that, in addition to seeking good results, seeks to reduce the consumption of resources.

Green AI: 4 approaches

According to the literature, there are four main ways to reduce the costs in Machine Learning:

- Compact Architecture Design
- Energy-efficient Training Strategies
- Energy-efficient Inference
- Efficient Data Usage

Green AI: 4 approaches

According to the literature, there are four main ways to reduce the costs in Machine Learning:

- Compact Architecture Design
- Energy-efficient Training Strategies
- Energy-efficient Inference
- Efficient Data Usage $\longleftarrow$ We will focus on this approach

Efficient data usage: Data Reduction

We want to reduce the size of the dataset, trying that the reduced dataset gives us a good representation of the full dataset.

$$\boxed{X_F\text{: Full Dataset}} \xrightarrow{\text{Reduction}} \boxed{X_R\text{: Reduced Dataset}}$$

$$\text{Properties}(X_F) \approx \text{Properties}(X_R)$$

Efficient data usage: Data Reduction

The idea is to use the reduced dataset for model training instead of the full dataset, making the process less expensive and giving similar results.

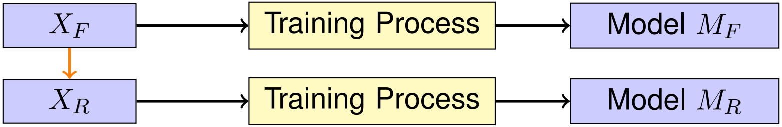

$$\text{Properties}(X_F) \approx \text{Properties}(X_R) \Rightarrow \text{Model } M_F \approx \text{Model} M_R$$

Ways to reduce a dataset

There are two main ways of reducing the size of a dataset:

- **Reducing feature size**: eliminating irrelevant or redundant features diminishes the dataset size and mitigates the risk of overfitting.

$$X_{N\times D} \longrightarrow Y_{N\times d}\ (d << D)$$

- **Reducing sample size**: discarding redundant or noisy examples and alleviating imbalances between classes can improve the training process.

$$X_{N\times D} \longrightarrow Z_{n\times D}\ (n << N)$$

Ways to reduce a dataset

There are two main ways of reducing the size of a dataset:

- **Reducing feature size**: eliminating irrelevant or redundant features diminishes the dataset size and mitigates the risk of overfitting.

$$X_{N\times D} \longrightarrow Y_{N\times d}\ (d << D)$$

- **Reducing sample size**: discarding redundant or noisy examples and alleviating imbalances between classes can improve the training process.

$$X_{N\times D} \longrightarrow Z_{n\times D}\ (n << N)$$

Size reduction

There are many reduction methods, which we classified into four categories:

- **Statistic-based methods,** which extract a subset either at random or using concepts from statistics and probability.
- **Geometry-based methods,** which use the distance matrix of the dataset to perform the reduction.
- **Ranking-based methods,** which order the items by some criterion and select the best ones.
- **Wrapper methods,** which perform the data reduction during the training process itself.

Data Reduction

Consider for example this classification dataset:

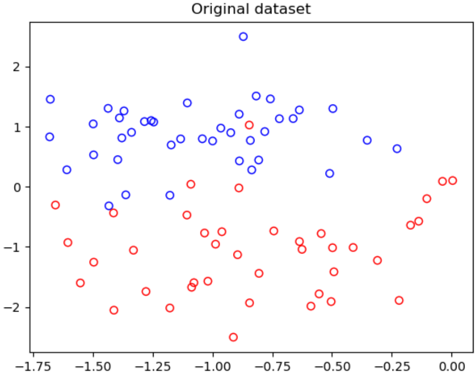

Data Reduction

We can reduce it using many methods:

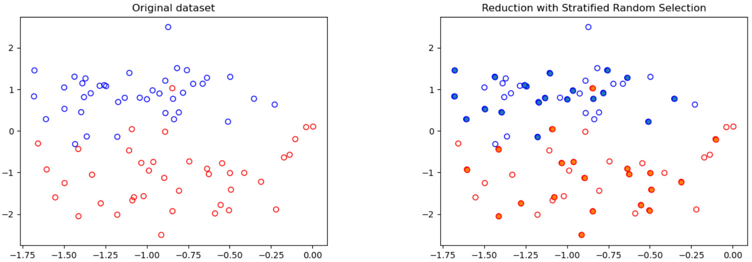

Verdecchia R, Cruz L, Sallou J, et al.: Data-centric green AI an exploratory empirical study. In: 2022 International Conference on ICT for Sustainability (ICT4S). IEEE, 2022; 35–45.

Data Reduction

We can reduce it using many methods:

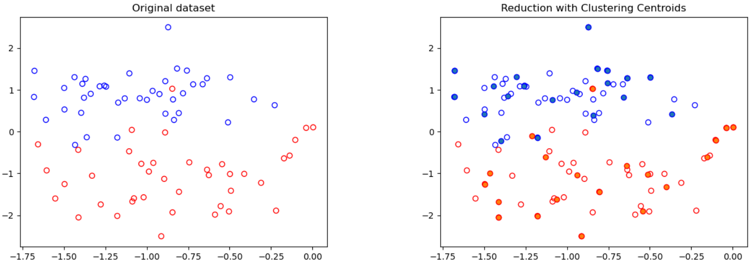

Olvera-López JA, Carrasco-Ochoa JA, Martínez-Trinidad JF, et al.: A review of instance selection methods. Artif Intell Rev. 2010; 34: 133–143.

Data Reduction

We can reduce it using many methods:

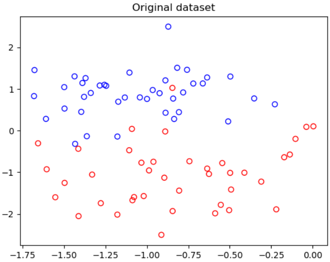

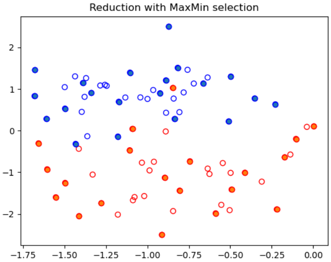

Lacombe C, Hammoud I, Messud J, et al.: Data-driven method for training data selection for deep learning. In: 82nd EAGE Annual Conference & Exhibition. European Association of Geoscientists & Engineers, 2021; 2021. : 1–5.

Data Reduction

We can reduce it using many methods:

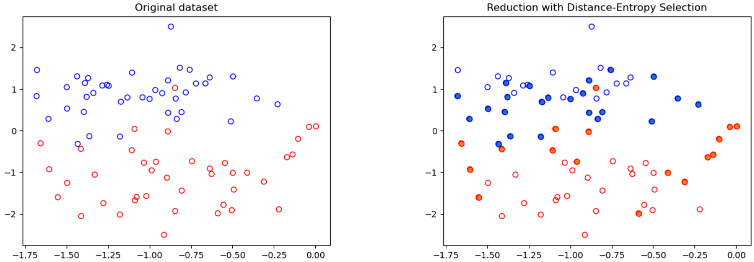

Li Y, Chao X: Distance-entropy: an effective indicator for selecting informative data. Front Plant Sci. 2022; 12: 818895.

Data Reduction

We can reduce it using many methods:

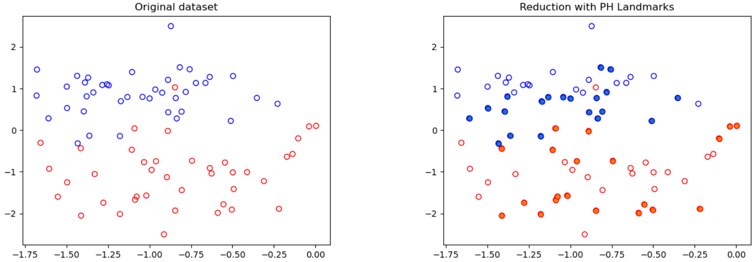

Stolz BJ: Outlier-robust subsampling techniques for persistent homology. J Mach Learn Res. 2023.

Data Reduction

We can reduce it using many methods:

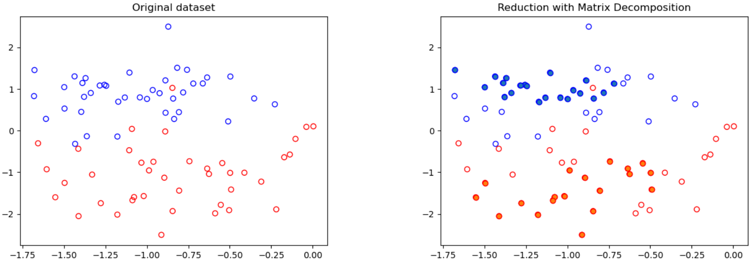

Ghojogh B, Crowley M: Instance ranking and numerosity reduction using matrix decomposition and subspace learning. In: Canadian Conference on Artificial Intelligence. 2019; 160–172.

Data Reduction

We can reduce it using many methods:

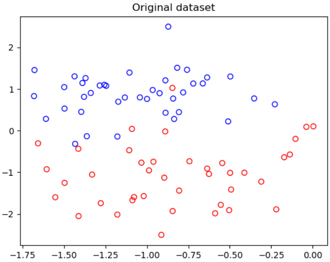

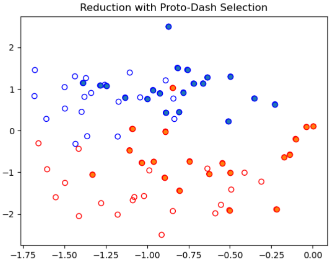

Gurumoorthy KS, Dhurandhar A, Cecchi G, et al.: Efficient data representation by selecting prototypes with importance weights. In: 2019 IEEE International Conference on Data Mining (ICDM). IEEE, 2019; 260–269.

Data Reduction

There are many reduction methods, and we created a Python module to apply and compare them.

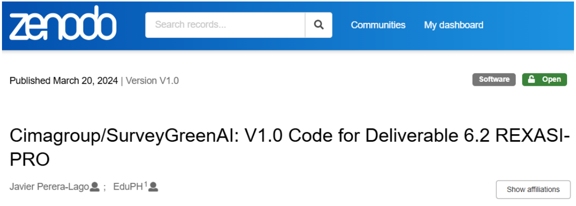

Perera-Lago, J., Toscano-Duran, V., Paluzo-Hidalgo, E., Gonzalez-Diaz, R., Gutiérrez-Naranjo, M. A., & Rucco, M. (2024). An in-depth analysis of data reduction methods for sustainable deep learning. Open Research Europe, 4(101), 101.

ε-representativeness

We ask ourselves:

How can we measure if a reduced dataset gives a good representation of the full dataset?

We will use the concept of ε-**representativeness**, which uses pairwise distances to measure the similarity between the full dataset and a reduced version of it.

Gonzalez-Diaz, R., Gutiérrez-Naranjo, M. A., & Paluzo-Hidalgo, E. (2022). Topology-based representative datasets to reduce neural network training resources. Neural Computing and Applications, 34(17), 14397-14413.

ε-representativeness

Let's assume we are trying to solve a classification task, and our dataset $\mathcal{D}$ is defined:

$$\mathcal{D} = \{(x, c_x) | x \in X \subset \mathbb{R}^n, c_x \in [[0, k]]\}$$

where $[[0, k]] = \{0, 1, 2, \cdots, k\}$. For each point $x \in X$, there is a label c_x that tells us its class. Each point belongs to one and only one class.

ε-representativeness

Definition: ε-representative point

Given a real number $\varepsilon > 0$ which we call the representation error, a labelled point (x, c_x) is ε-representative of $(\tilde{x}, c_{\tilde{x}})$ if $c_x = c_{\tilde{x}}$ and $||x - \tilde{x}|| \leq \varepsilon$. We denote $x \approx_\varepsilon \tilde{x}$.

ε-representativeness

Example of ε-representative points.

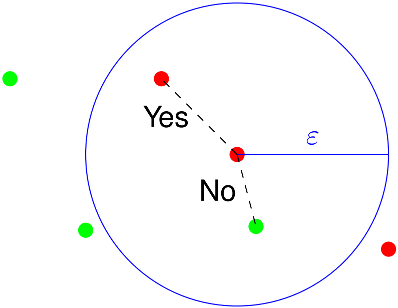

ε-representativeness

We extend ε-representativeness between pair of points to define the ε-representativeness between datasets:

Definition: ε-representative dataset

A dataset $\tilde{\mathcal{D}} = \{(\tilde{x}, c_{\tilde{x}}) | \tilde{x} \in \tilde{X} \subset \mathbb{R}^n, c_{\tilde{x}} \in [[0, k]]\}$ is ε-representative of $\mathcal{D} = \{(x, c_x) | x \in X \subset \mathbb{R}^n, c_x \in [[0, k]]\}$ if there exists an isometric transformation $f : \tilde{X} \to \mathbb{R}^n$, such that for any $(x, c_x) \in \mathcal{D}$ there exists $(\tilde{x}, c_{\tilde{x}}) \in \tilde{\mathcal{D}}$ satisfying that $f(\tilde{x}) \approx_\varepsilon x$.

ε-representativeness

ε-representative datasets preserve persistent homology:

ε-representativeness

ε-representative datasets preserve persistent homology:

Theorem 1 [1]

If the dataset $\tilde{\mathcal{D}}$ is ε-representative of $\mathcal{D}$, then

$$d_B(\mathrm{Dgm}_q(X), \mathrm{Dgm}_q(\tilde{X})) \leq 2\varepsilon$$

where $q \leq n$, $\mathrm{Dgm}_q(X)$ and $\mathrm{Dgm}_q(\tilde{X})$ are the persistence diagrams of the Vietoris-Rips filtrations computed from X and $\tilde{X}$, and d_B denotes the bottleneck distance between their persistence diagrams.

ε-representativeness

Given a dataset $\mathcal{D}$, a reduction $\mathcal{D}_R$ and an isometry $i : \mathcal{D}_R \to \mathbb{R}^d$, the minimum ε such that $\mathcal{D}_R$ is ε-representative dataset of $\mathcal{D}$ is:

$$\varepsilon^* = \max_{k=1,\dots,c} \max_{x : c_x = k} \min_{x' : c_{x'} = k} ||x - i(x')||$$

Applying data reduction

We applied some experiments about data reduction on the **Collision Dataset.**

It consists on a set of simulations where a platoon of vehicles navigates an environment. The classification task consists in deciding whether the platoon will collide based on features such as the number of cars and their speed.

Mongelli, M., Ferrari, E., Muselli, M., & Fermi, A. (2019). Performance validation of vehicle platooning through intelligible analytics. IET Cyber-Physical Systems: Theory & Applications, 4(2), 120-127.

Applying data reduction

We trained a fixed Multi-Layer Perceptron with the full dataset and with many reduced dataset given by six different methods and we got the following results:

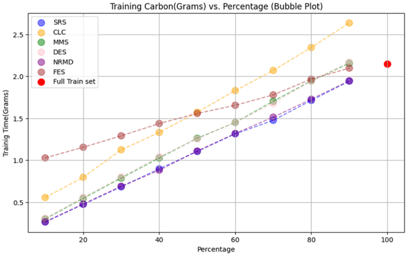

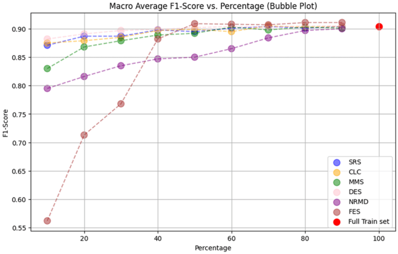

Applying data reduction

There is a significant correlation between ε-representativeness of the subset and the F1-score of the trained network.

Perera-Lago, J., Toscano-Duran, V., Paluzo-Hidalgo, E., Gonzalez-Diaz, R., Gutiérrez-Naranjo, M. A., & Rucco, M. (2024). An in-depth analysis of data reduction methods for sustainable deep learning. Open Research Europe, 4(101), 101.

	Spearman's ρ	p-value
10%	-0.38	0.0
20%	-0.43	0.0
30%	-0.42	0.0
40%	-0.39	0.0
50%	-0.22	0.1
60%	-0.15	0.24
70%	-0.19	0.14
80%	-0.07	0.58
90%	-0.14	0.3

Applying data reduction

We also performed some experiments reducing the Collision Dataset in another family of models more interpretable by construction: Decision Trees.

Perera-Lago, J., Toscano-Durán, V., Paluzo-Hidalgo, E., Narteni, S., & Rucco, M. (2024, July). Application of the representative measure approach to assess the reliability of decision trees in dealing with unseen vehicle collision data. In World Conference on Explainable Artificial Intelligence (pp. 384-395). Cham: Springer Nature Switzerland.

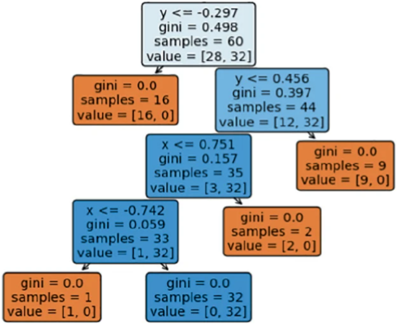

Applying data reduction

In this case, we also found that:

- Subsets with better ε-representativeness train decision trees with higher accuracy
- Subsets with better ε-representativeness train decision trees more similar to the tree train with the full dataset in terms of feature importance

T6.3

Topology-based optimization of robot fleet behavior

Navigation behaviors

A **behavior** is a local navigation algorithm that acts on each autonomous agent, trying to reach a target point or direction as fast as possible while avoiding obstacles.

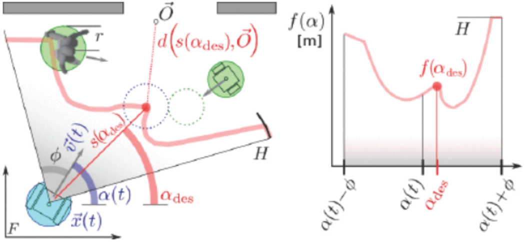

Figure from *Guzzi, J., Giusti, A., Gambardella, L. M., Theraulaz, G., & Di Caro, G. A. (2013, May). Human-friendly robot navigation in dynamic environments. In 2013 IEEE international conference on robotics and automation (pp. 423-430). IEEE.*

Navigation behaviors

Human-Like

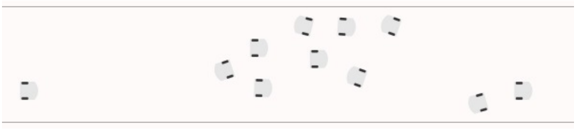

ORCA

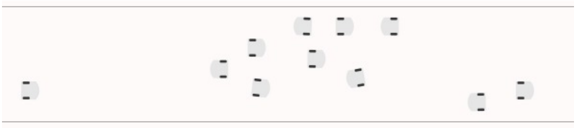

Social Force

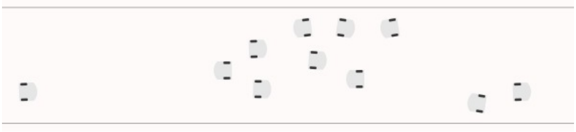

Navigation behaviors

In this task we had the following objectives:

1. To use Persistent Homology to define a measure for order and stability within a fleet of autonomous agents.
2. To use this measure to compare the performance of the three navigation behaviors shown before.

PH to distinguish behaviors

To tackle this objectives, we define the *induced matching distance*.

Consider two sets of points $X = \{x_1, x_2, \dots, x_n\}$ and $Z = \{z_1, z_2, \dots, z_n\}$ with a bijection

$$\begin{array}{rccc} f_\bullet \colon & X & \to & Z \\ & x_i & \mapsto & z_i \end{array}$$

and two symmetric non-negative functions $d_X \colon X \times X \to \mathbb{R}^+$ and $d_Z \colon Z \times Z \to \mathbb{R}^+$.

We want a distance between the barcodes $\mathrm{B}(X)$ and $\mathrm{B}(Z)$.

Comparing barcodes

A classical method is the q-Wasserstein distance:

$$W_q(\mathrm{B}(X), \mathrm{B}(Z)) = \inf_{\mu \in M} \left(\sum_{\substack{(a,\ell) \in \operatorname{Rep} B(X) \\ \mu((a,\ell))=(b,\ell')}} |a-b|^q \right)^{1/q},$$

M is the set of all partial matchings $\mu \colon \operatorname{Rep} \mathrm{B}(X) \nrightarrow \operatorname{Rep} \mathrm{B}(Z)$.

Comparing barcodes

W_q compares all the possible partial matchings in M and uses the optimal one.

However, the bijection $f_\bullet \colon X \to Z$ induces an isomorphism:

$$f_0 \colon \, \mathrm{H}_0(\mathrm{VR}_0(X)) \to \mathrm{H}_0(\mathrm{VR}_0(Z))$$

and therefore a specific partial matching $\sigma_f^0 \in M$.

Induced matching distance

Then, we propose the q-induced matching distance:

$$d^q_{f_0}(\mathrm{B}(X), \mathrm{B}(Z)) = \left(\sum_{\substack{(a,\ell)\in \mathrm{Rep}\, B(X) \\ \sigma^0_f((a,\ell))=(b,\ell')}} |a-b|^q \right)^{1/q}$$

Clearly, $W_q(\mathrm{B}(X), \mathrm{B}(Z)) \leq d^q_{f_0}(\mathrm{B}(X), \mathrm{B}(Z))$

Induced matching distance

Let X_0 be the set of points $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3, \mathsf{P}_4$ at time 0.
Let X_1 be the set of points $\mathsf{P}_1, \mathsf{P}_2, \mathsf{P}_3, \mathsf{P}_4$ at time 1.
The bijection $f_\bullet \colon X_0 \to X_1$ is the trivial one, $f_\bullet(\mathsf{P}_i) = \mathsf{P}_i$.

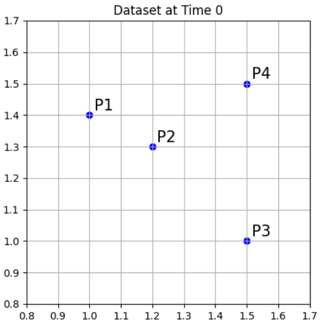

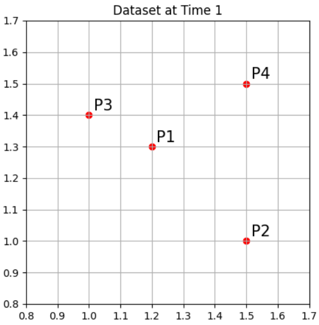

Induced matching distance

These are the partial matchings that define the distances $W_q(\mathrm{B}(X_0), \mathrm{B}(X_1))$ and $d^q_{f_0}(\mathrm{B}(X_0), \mathrm{B}(X_1))$

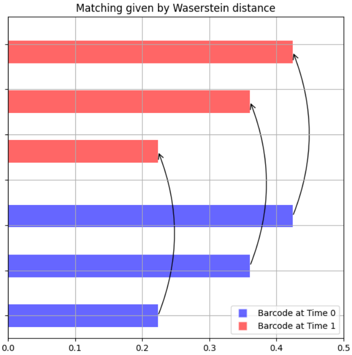

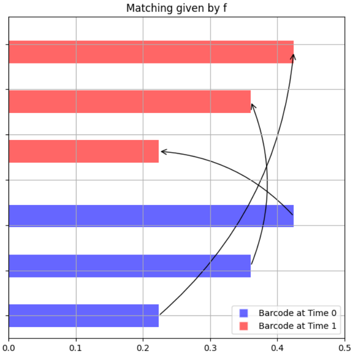

Induced matching signal

Application to a group of navigating agents.

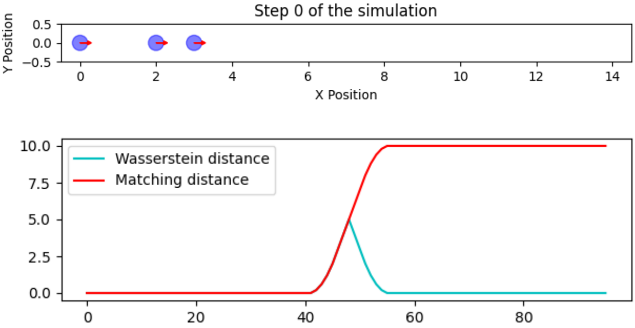

Induced matching signal

We want to use the induced matching signal as a measure of order and stability within the fleet.

Also, we want to know if it is helpful to distinguish between the three different behaviors we have shown (Human-Like, ORCA and Social Force)

Navground

We use Navground, a Python simulator for robots navigation.

Corridor scenario:

- 15m long, 3.5m wide, both ends connected.
- 10 agents with 0.8m of diameter and 1.2m/s of optimal speed.
- 5 agents driving left, 5 agents driving right.

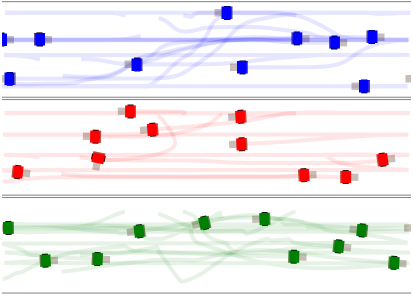

We run 200 simulations with 900 steps for each behavior type.

Induced matching signal

Given a simulation, we apply the following steps:

1. For $i = 1, \ldots, 10$, Agent $i \longrightarrow a^i = \left\{a_t^i = (x_t^i, y_t^i, \alpha_t^i)\right\}_{t=1}^{900}$
2. For $t = 1, 2, \ldots, 850$, $Z_t = \left\{z_t^i = \{a_t^i, a_{t+10}^i, \ldots, a_{t+50}^i\}\right\}_{i=1}^{10}$
3. DTW as distance $\longrightarrow \left\{\, \mathrm{VR}_0(Z_t)\right\}_{t=1}^{850} \longrightarrow \left\{\, \mathrm{B}(Z_t)\right\}_{t=1}^{850}$
4. For $t = 1, \ldots, 800$,
   $f_\bullet^t \colon Z_t \to Z_{t+50} \longrightarrow m = \left\{d^1_{f_0^t}(B(Z_t), B(Z_{t+50}))\right\}_{t=1}^{800}$

m is called the induced matching signal of the simulation

Induced matching signal

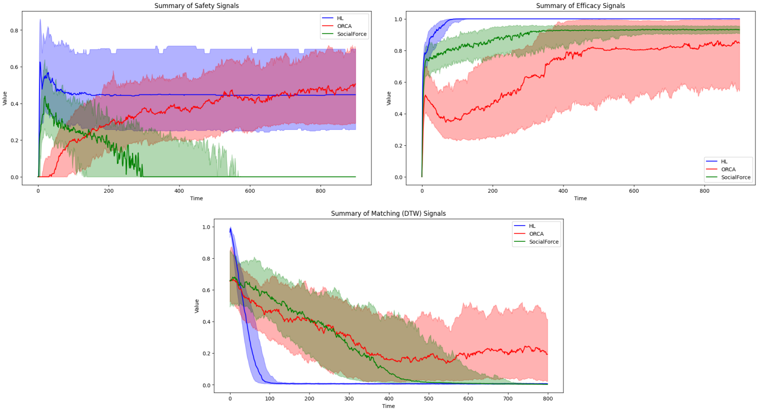

Induced matching signal

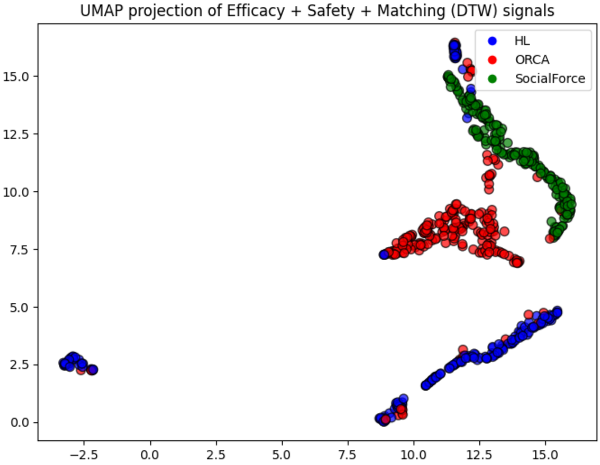

CrossTorus scenario

We also performed the same experiment on another scenario called CrossTorus:

CrossTorus

In this case we found more difficult to distinguish between behaviors.

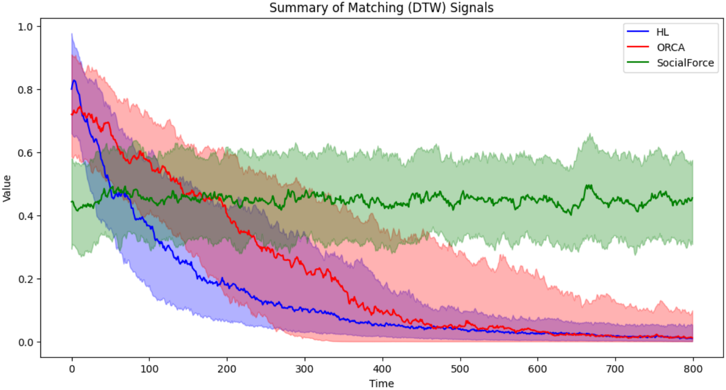

CrossTorus

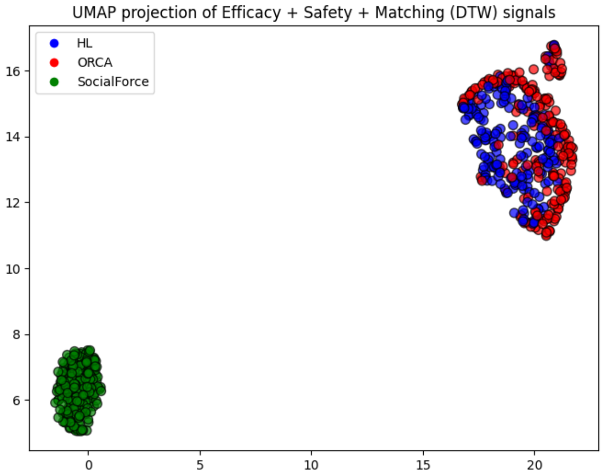
UMAP projection of Efficacy + Safety + Matching (DTW) signals
HL
ORCA
SocialForce
16
14
12
10
8
6
0
5
10
15
20

Thanks for your attention.