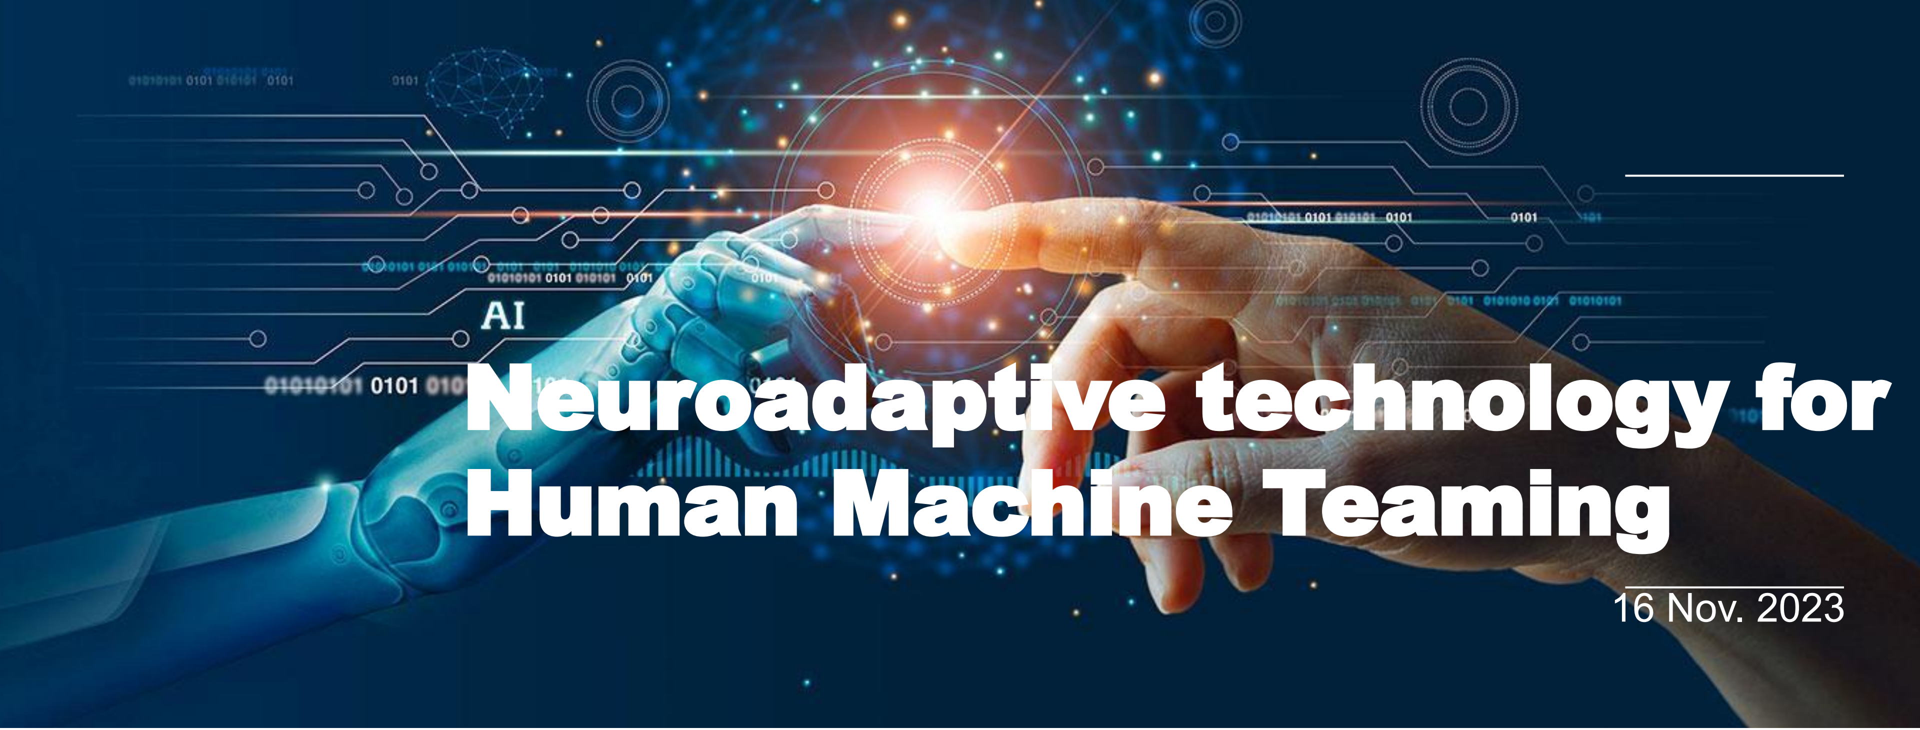

Neuroadaptive technology for Human Machine Teaming

16 Nov. 2023

ANITI

Université de Toulouse

Contributions from PhD theses of X. XU & G. ANGELOTTI

- *“In search of invariants in varying electroencephalography signals for brain-computer interfaces”*
  defended by **Xiaoqi XU** on January 27th 2023.

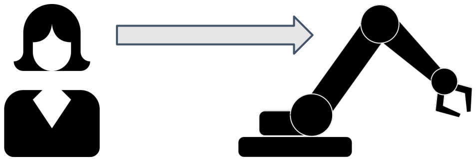

- *“Advances in Risk-Aware Offline Reinforcement Learning: A Study of Data Augmentation, Explainability, and Policy Selection”*
  defended by **Giorgio ANGELOTTI** on June 12th 2023.

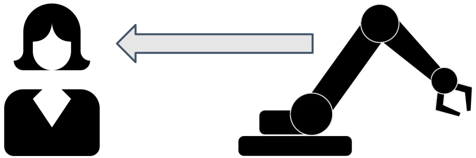

ANITI
Université de Toulouse

Giorgio ANGELOTTI – some PhD Thesis contributions

Context and Contributions

Context
How to improve Human-Machine Interaction (HMI) offline?
(i.e. from previously collected data?)

Data collection with Human in the Loop is expensive
→ Lack of HMI data
→ Data (inter-subject) variability *(cf. presentation X.Xu)*

Contributions in Offline Reinforcement Learning (ORL):

- Data augmentation
- Risk-aware Policy selection
- Application to HMI

ANITI

Université de Toulouse

Offline Reinforcement Learning

Levine et al.
Offline Reinforcement Learning: Tutorial, Review and Perspectives on Open Problems
(2020) arXiv preprint

- No interaction to explore (no trial and error)
- Learns from a data set of experiences
- Possibly **become better** than the recorded agent

Challenges

- *Vanilla Temporal Difference algorithms not good without exploration*

$$\frac{\delta Q_i^*(s_t,a_t)}{\delta i} \approx \frac{Q_{i+1}^*(s_t,a_t) - Q_i^*(s_t,a_t)}{\alpha_i} = r_t + \gamma \boxed{\max_a Q_i^*(s_{t+1},a)} - Q_i^*(s_t,a_t).$$

Optimistic in front of Uncertainty

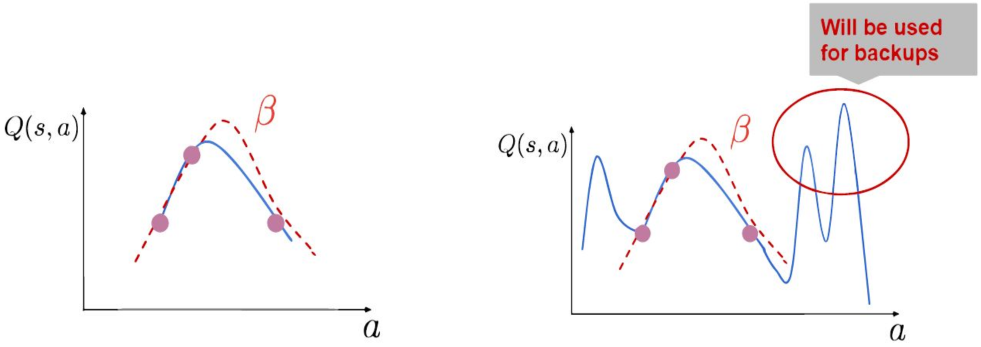

Figure 2.4: Taken from the presentation of Kumar, Fu et al. (2019) at NeurIPS 2019. Q-function

Works

Angelotti, Drougard, Chanel
Offline Learning for Planning: A Summary
(2020) Proceedings of the 1st Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning at the ICAPS

Review of the literature of offline learning (missing at the time)

Angelotti, Drougard, Chanel
Expert-guided Symmetry Detection in MDPs
(2022) Proceedings of the 14th International Conference on Agents and Artificial Intelligence

Angelotti, Drougard, Chanel
Data Augmentation Through Expert-Guided Symmetry Detection to Improve Performance in Offline Reinforcement Learning
(2023) Proceedings of the 15th International Conference on Agents and Artificial Intelligence

Can we improve the data efficiency of algorithms in the offline context?
Validation of expert proposed symmetries

Angelotti, Drougard, Chanel
An Offline Risk-aware Policy Selection Method for Bayesian MDPs
Submitted to an International Journal

How to select a deterministic risk-sensitive policy between many?
Robust offline policy selection with the Bayesian formalism

Angelotti, et al.
Our method applied to HMI
Submitted to an International Conference

Can we apply these methods to Human-Computer Interaction with Physiological computing?
Robust application to Firefighter Robot use case

Angelotti, Diaz Rodriguez
Towards a more efficient computation of individual attribute and policy contribution for post-hoc explanation of cooperative multi-agent systems using Myerson values
(2023) Elsevier's Knowledge-Based Systems

Explainability for Multi-Agent Systems

Contribution:

*Check if an **alleged symmetry** k **is valid*** and use the resulting knowledge for data augmenting.

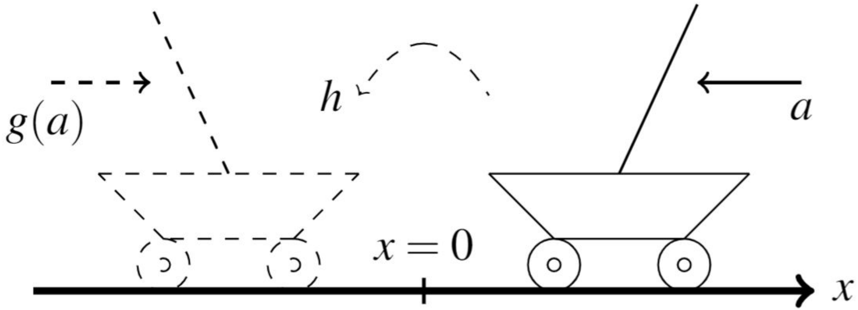

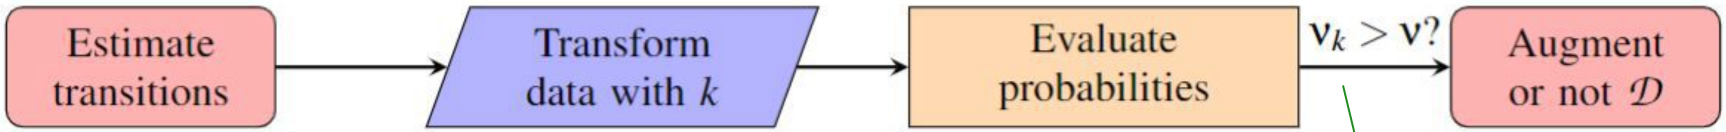

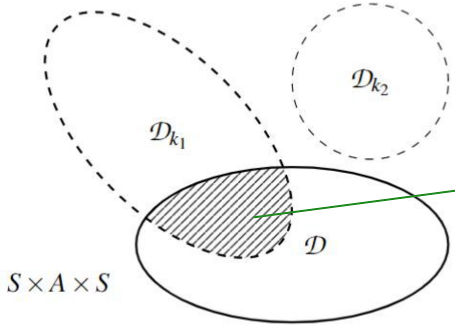

proportional to the size of the intersection!

I Data augmentation/efficiency

$$U^{\pi} = \mathbb{E}_{s\sim\rho}[V^{\pi}(s)] \qquad \Delta U = U^{\hat{\pi}_k} - U^{\hat{\pi}}$$

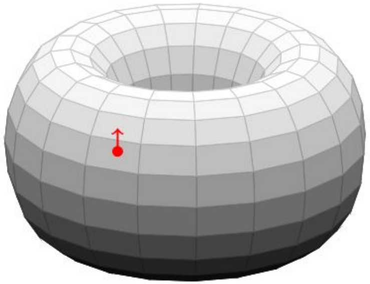

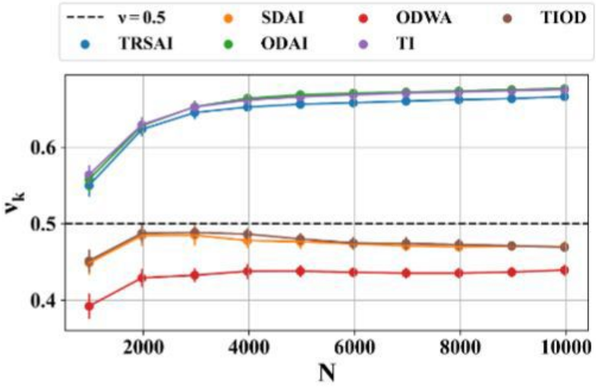

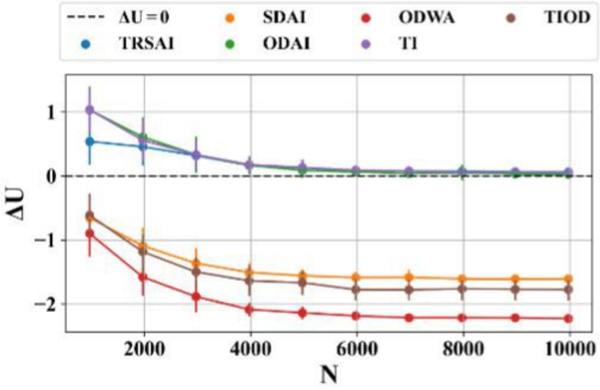

II Risk-aware Policy selection

Contribution:

*A method to offline evaluate and select **deterministic** policies in a **risk-sensitive** way*

technique to do so for *small finite states and actions* MDPs.

Intuition behind risk-sensitive metrics

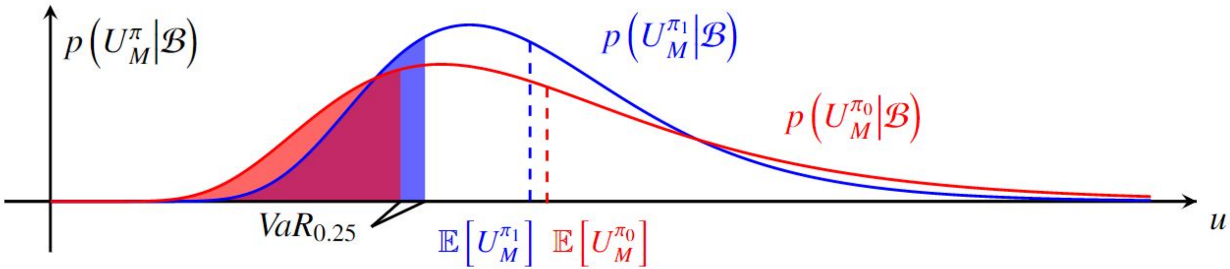

Exploitation vs Caution – step 1

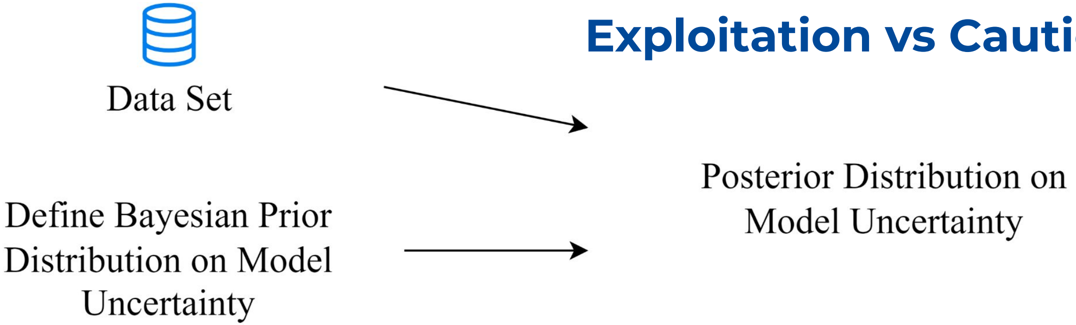

Exploitation vs Caution – step 2

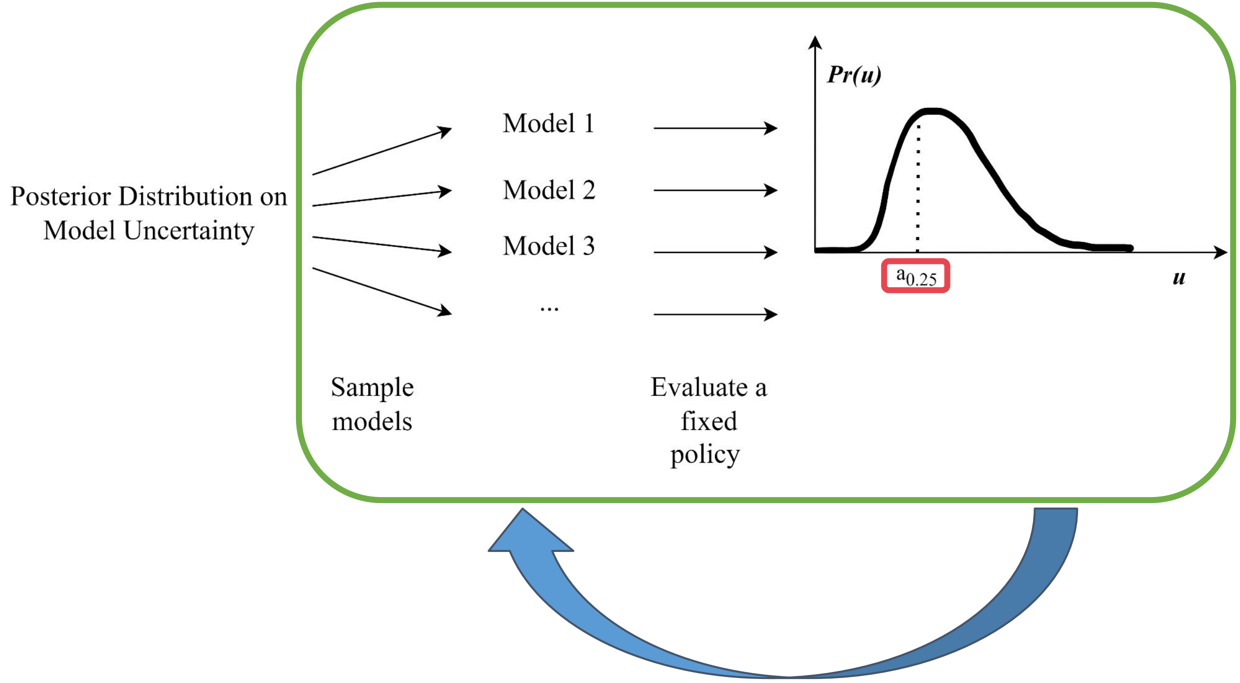

Policy_1

Estimate the same value for every different policy in the candidate set,

{Policy_1, ..., Policy_n}

then select the one that maximizes the Risk-sensitive metric!

Reiterate until the estimate has the wanted statistical significance

II Risk-aware Policy selection

Environments

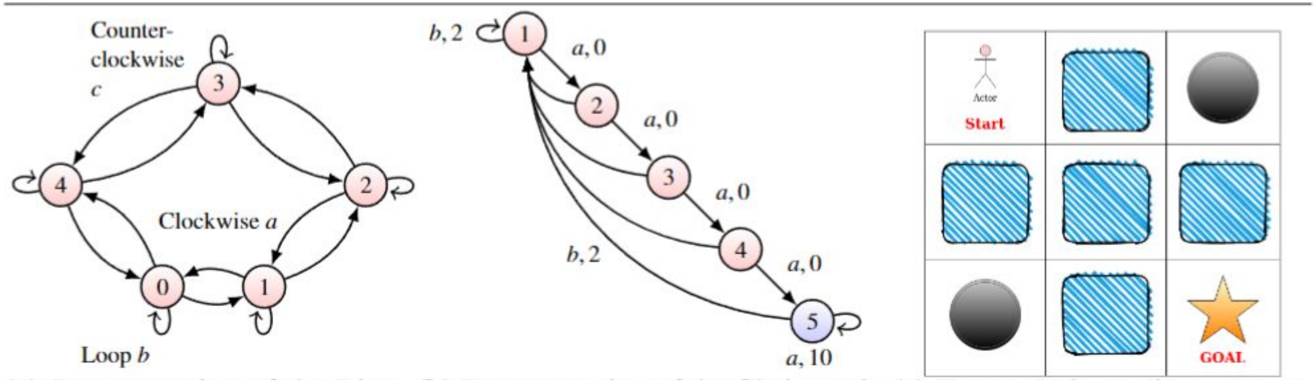

II Risk-aware Policy selection

Results

Environment	Metrics	Baseline SPIBB	BOPAH	BCR	NORBU	Selection Method $EvC_{VaR_{0.25}}$	$EvC_{CVaR_{0.25}}$	$UnO_{VaR_{0.25}}$	$UnO_{CVaR_{0.25}}$
Ring	*Max*	0.61	0.48	0.74	**0.84**	**0.82**	0.71	**0.82**	0.72
	Mean	-0.29	-0.28	-0.01	**0.03**	**0.01**	-0.04	-0.26	-0.27
	Median	-0.31	-0.34	**0.0**	**0.0**	**0.0**	**0.0**	-0.27	-0.33
	Min	-0.78	**-0.68**	-0.82	-0.71	**-0.82**	**-0.82**	-0.96	-0.96
Chain	*Max*	**0.55**	0.54	**0.55**	**0.55**	**0.55**	**0.55**	0.54	0.54
	Mean	0.0	0.01	0.01	**0.02**	**0.01**	**0.01**	**0.01**	**0.01**
	Median	**-0.01**	**-0.01**	**-0.01**	**-0.01**	**-0.01**	**-0.01**	**-0.01**	**-0.01**
	Min	-0.38	-0.16	**-0.15**	**-0.15**	**-0.16**	**-0.16**	**-0.16**	**-0.16**

(< 10 states, < 10 actions)

Ring: NORBU ☑️ , EvC ☑️

Chain: No best method

Usage of trivial policy increases with batch size

Policy selection rate by EvC $VaR_{0.25}$ and EvC $CVaR_{0.25}$ in Ring for different batch sizes.

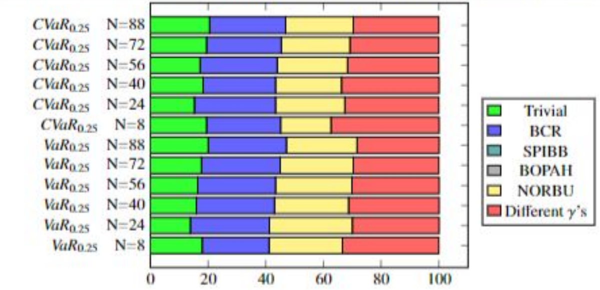

Policy selection rate by EvC $VaR_{0.25}$ and EvC $CVaR_{0.25}$ in Chain for different batch

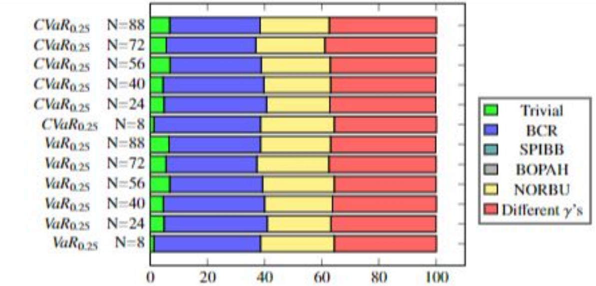

ANITI
Université de Toulouse

Application to Mixed-Initiative Human-Robot Interaction

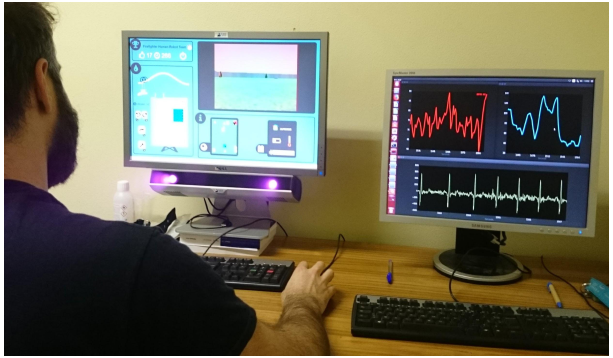

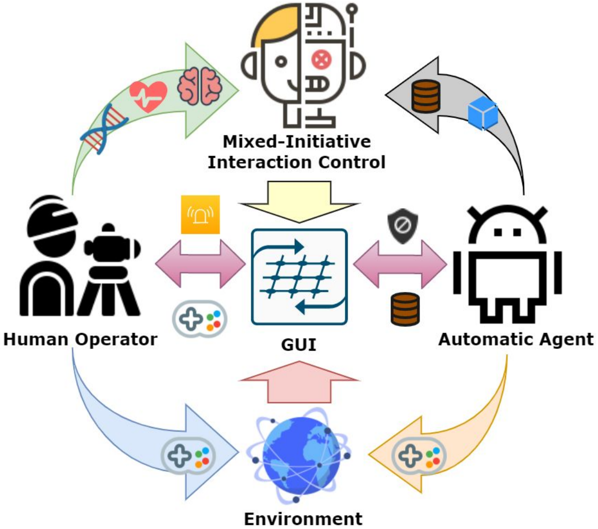

- **Human Supervision**
- **Dangerous** consequences of bad policies
- **Limited previously collected** data set
- **Partial observability**

III Application to HMI

Application: Offline RL with Human-in-The-Loop and Physiological Computing

III Application to HMI

What is missing? *A method to include model uncertainty for learned POMDPs and to compute a robust policy*

Contribution:

technique to do so, *specific to the application to* _our use case_

Batch of experiences
Random Interaction Policy

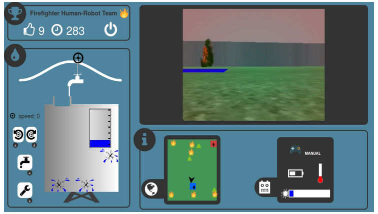

Eye tracking (ET)
Electrocardiography (ECG)
Human actions on the interface (HAI)

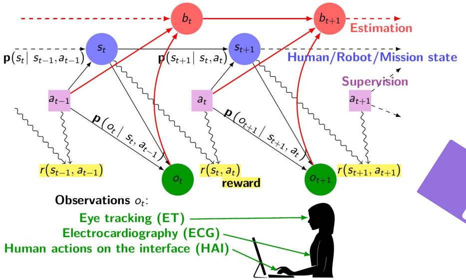

Uncertainty in observation function

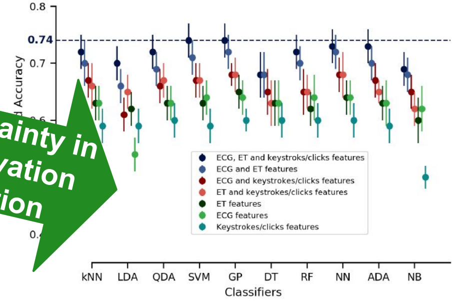

Observation function p(o|s',a)

Uncertainty in transition function

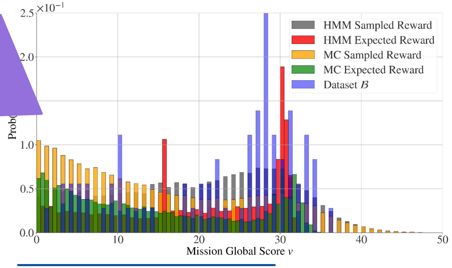

Transition function p(s'|s,a)

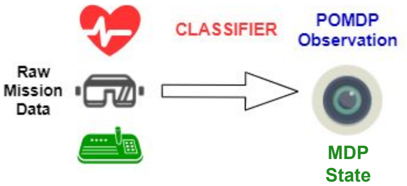

Classifier confusion matrix:
Dirichlet Prior for POMDP Observation

SOLVE ROBUSTLY THE POMDP AND THE MDP TAKING INTO ACCOUNT _MODEL UNCERTAINTY_

ANITI
Université de Toulouse

III Application to HMI

Evaluate the mixed-initiative interaction policy

compare following policies:

1. ***Data collector policy (Random)***
2. *Full **automatic** policy*
3. ***MDP adaptive strategy** with physiological and behavioural data*
4. ***POMDP adaptive strategy** with physiological and behavioural data*

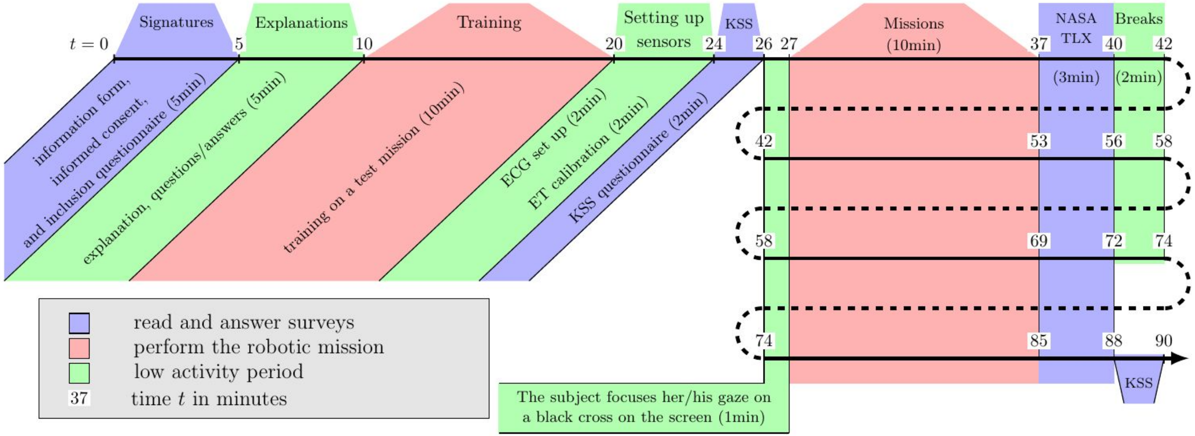

CER-2018-070

ANITI

Université de Toulouse

III Application to HMI

Policy	mean	std	min	25%	Median	75%	max
Random ($\mathcal{D}$)	22.1	10.0	1.0	12.8	26.5	29.3	36.0
Random ($\mathcal{D}_{val}$)	17.9	9.6	4.0	8.3	20.0	26.8	32.0
MDP ($\mathcal{D}_{val}$)	22.7	7.6	6.0	18.0	22.5	**28.8**	**39.0**
POMDP ($\mathcal{D}_{val}$)	**23.4**	**5.7**	**14.0**	**18.5**	**24.0**	27.5	37.0
FA ($\mathcal{D}_{val}$)	*25.6*	*5.0*	*9.0*	*25.3*	*27.0*	*28.0*	*32.0*

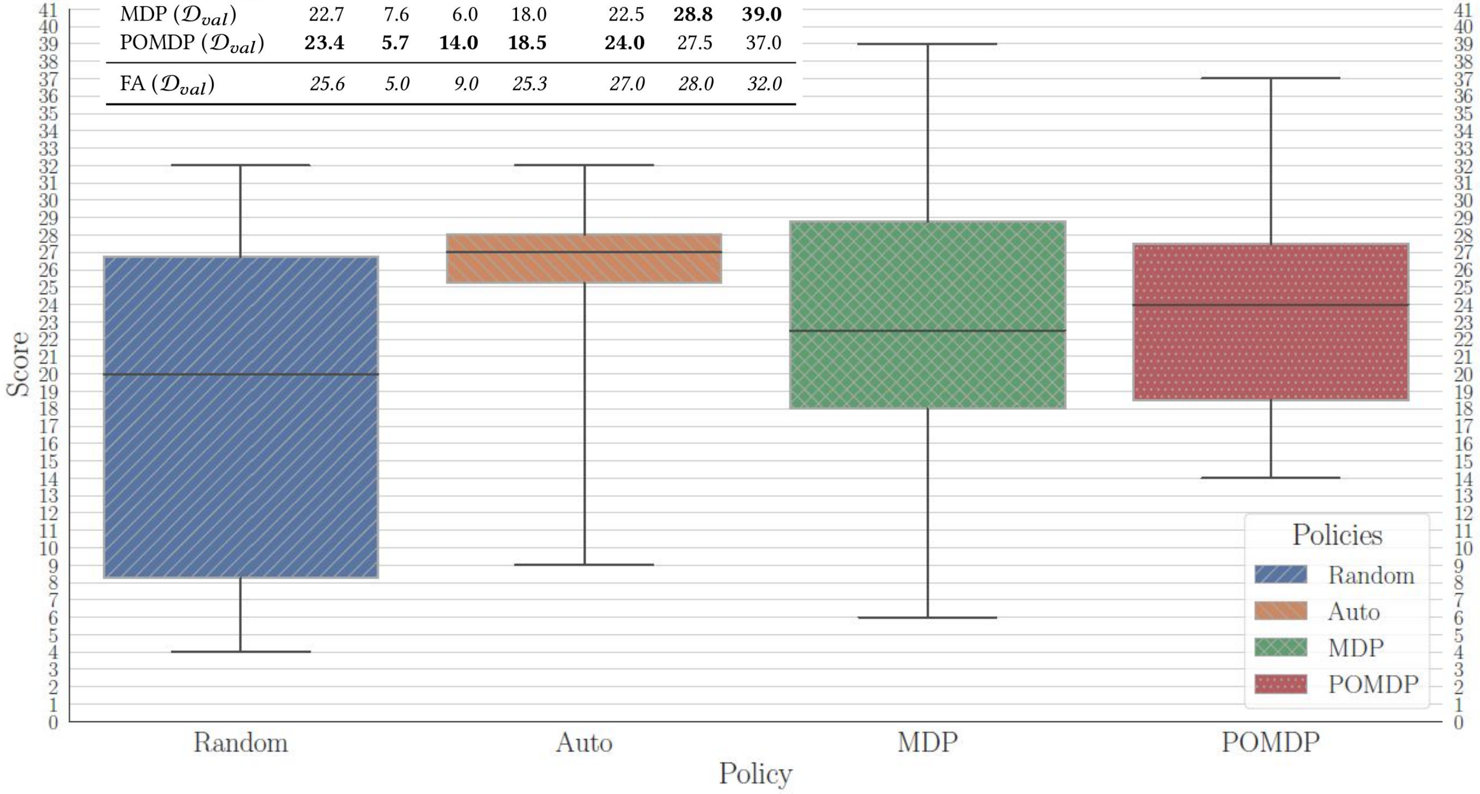

Belief of performance

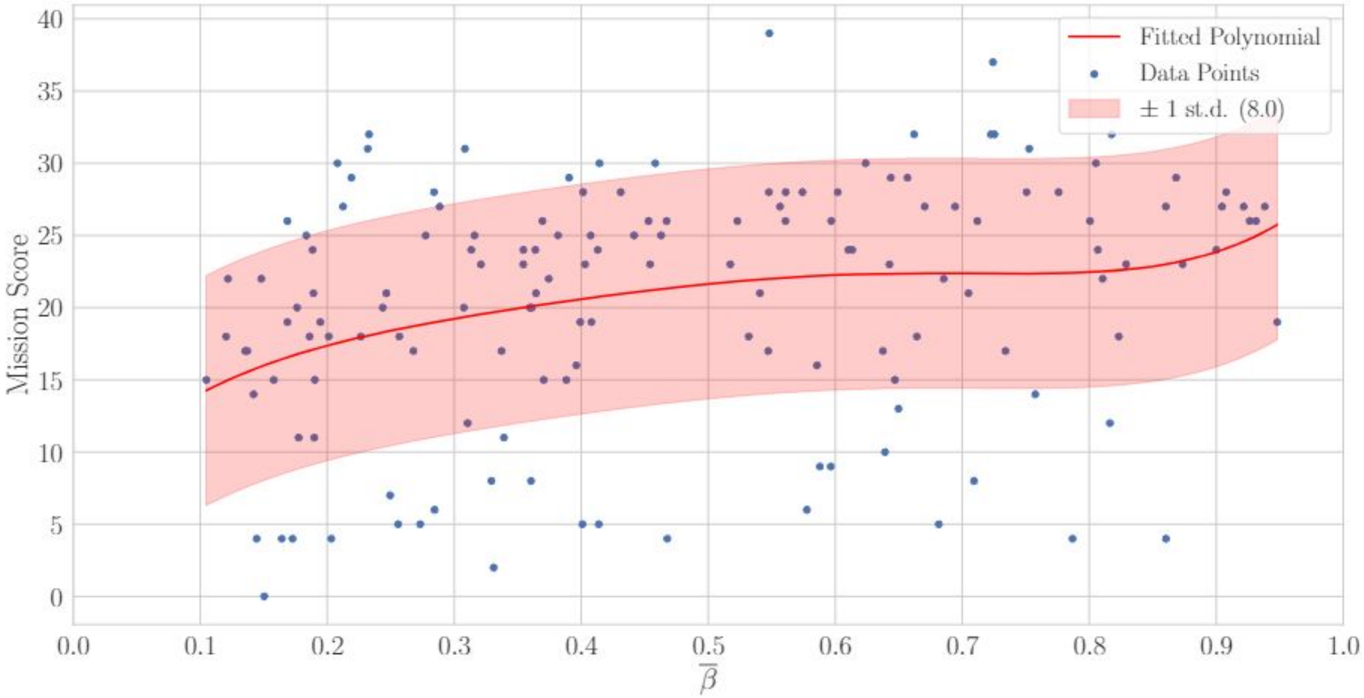

Spearman's $\rho = 0.325$

$p(\overline{\beta})$-value < 0.001

Brain Computer Interface for Human Machine Teaming

Offline Reinforcement Learning could improve Human Machine Interaction.

Teaming up better with people,

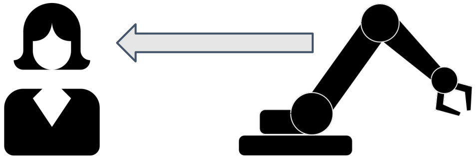

by better understanding their state.

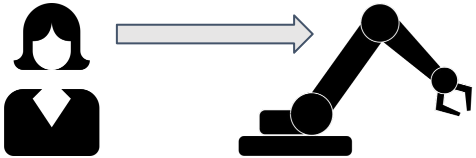

→ EEG-based Brain Computer Interface!

ANITI

Université de Toulouse

Xiaoqi XU – some PhD Thesis contributions

Context and Contributions

Context:

How to design a classifier for BCI that tackles inter-subject EEG variability?

Contributions:

- Laplacian (spatial)
- Path signature (temporal)
- Topological data analysis (spatio-temporal)

Context

EEG varies between sessions and subjects

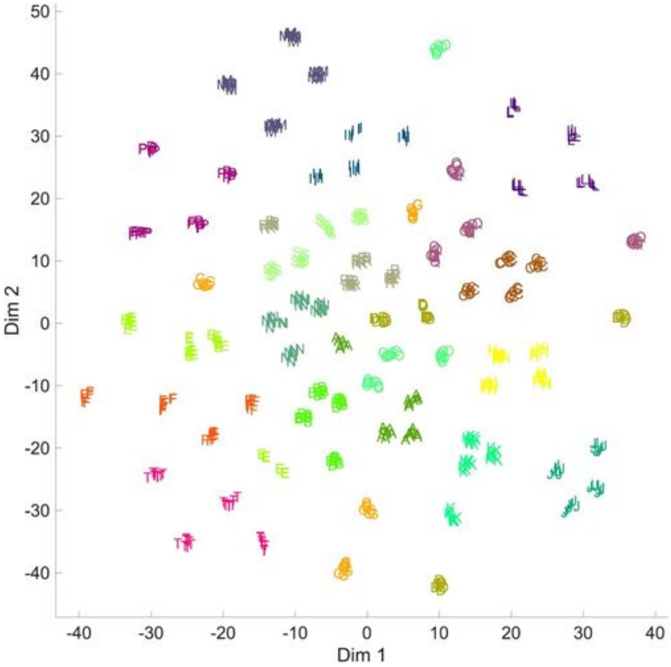

Nishimoto, T., Higashi, H., Morioka, H., & Ishii, S. (2020). Eeg-based personal identification method using unsupervised feature extraction and its robustness against intra-subject variability. Journal of Neural Engineering, 17(2), 026007.

Context

EEG variability

Internal

- cellular noise
- electrical noise
- synaptic noise

- neural plasticity
- brain topography
- circadian rhythms

Microscopic

Macroscopic

- electrodes' placement

- environmental noise
- task

External

Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008)
Masquelier, T. (2013)
Scrivener, C. L. & Reader, A. T. (2022)

Croce, P., Quercia, A., Costa, S., & Zappasodi, F. (2018)
McIntosh, A. R., Kovacevic, N., & Itier, R. J. (2008)
Gibson, E., Lobaugh, N. J., Joordens, S., & McIntosh, A. R. (2022)

Solutions

Transfer learning

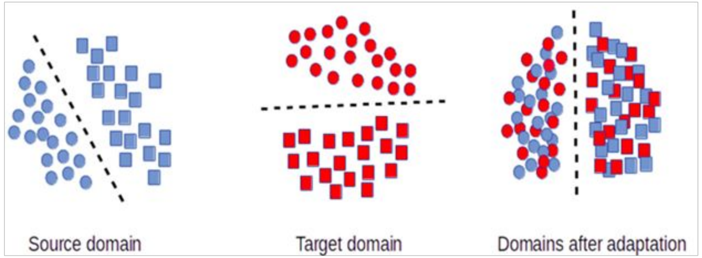

Lotte, F., Bougrain, L., Cichocki, A., Clerc, M., Congedo, M., Rakotomamonjy, A., & Yger, F. (2018). A review of classification algorithms for eeg-based brain–computer interfaces: a 10 year update. Journal of Neural Engineering

Jayaram, V., Alamgir, M., Altun, Y., Scholkopf, B., & Grosse-Wentrup, M. (2016). Transfer learning in brain-computer interfaces. IEEE Computational Intelligence Magazine, 11(1), 20–31.

Adaptive waveform learning

Model the neural events through adaptive kernels

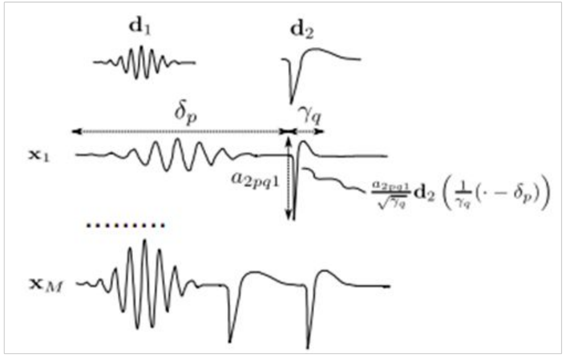

Hitziger, S., Clerc, M., Saillet, S., Bénar, C., & Papadopoulo, T. (2017). Adaptive waveform learning: A framework for modeling variability in neurophysiological signals. IEEE Transactions on Signal Processing, 65(16), 4324–4338.

Contributions

Objectif: robust features against inter-subject variability

Look at the big common picture

- Laplacian: hierarchical representation that encodes the intrinsic geometry

Time is elastic

- Path signature: invariant under time reparametrization and captures order

Rules of change are unchanged

- TDA: extracts topological properties of the attractor of EEG dynamics

ANITI
Université de Toulouse

Publications

- **X. Xu**, N. Drougard, R. N. Roy (2021). Dimensionality Reduction via the Laplace-Beltrami Operator: Application to EEG-based BCI, *IEEE EMBS Conf. NER*
- **X. Xu**, N. Drougard, R. N. Roy (2021). Topological Data Analysis as a New Tool for EEG Processing. Front. Neurosci. 15:761703.
- E. Jahanpour *, **X. Xu** *, M. F. Hinss, N. Drougard, R. N. Roy (2021). A neuroergonomic approach to performance estimation in a psychomotor vigilance task. *Neuroergonomics Conference*
- Roy R., Hinss M., Darmet L., Ladouce S., Jahanpour E., Somon B., **Xu X.** , Drougard N., Dehais F., Lotte F. (2022). Retrospective on the First Passive Brain-Computer Interface Competition on Cross-Session Workload Estimation. *Front. Neuroergon.*
- **X. Xu**, D. Lee, N. Drougard, R. N. Roy (accepted for publication in Scientific Reports). Signature methods for brain-computer interfaces. Preprint available at Research Square [https://doi.org/10.21203/rs.3.rs-2476159/v1]
- **X. Xu**, N. Drougard, R. N. Roy (to be submitted). Tackling inter-subject variability in brain-computer interface via topological data analysis.

ANITI

Université de Toulouse

Data availability:

All of the datasets analysed during the thesis are publicly available. The links for download:

- *BCI competition IV 2a dataset*: https://www.bbci.de/competition/iv/
- *Physionet motor imagery dataset*: https://physionet.org/content/eegmmidb/1.0.0/
- *Passive BCI competition dataset*: https://zenodo.org/record/4917218#.Y8pgFKfMI5k

Code availability:

Code of these methods can be found on github:

- *Path signature*: https://github.com/XiaoqiXu77/Signature_BCI
- *TDA*: https://github.com/XiaoqiXu77/TDA_BCI

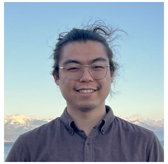

Path signature

- Originated in pure math to solve stochastic differential equations
- Borrowed by machine learning community as a feature map for time series (e.g. handwritten character recognition, diagnosis of bipolar disorder etc.)
- Never used for BCI

Chen, K.-T. (1958). Integration of paths – a faithful representation of paths by noncommutative formal power series. Transactions of the American Mathematical Society, 89, 395–407
Lyons, T. J., Caruana, M., & Lévy, T. (2007). Differential equations driven by rough paths. Springer.
Yang, W., Jin, L., & Liu, M. (2016). Deepwriterid: An end-to-end online text-independent writer identification system. IEEE Intelligent Systems, 31(2), 45–53.
Perez Arribas, I., Goodwin, G. M., Geddes, J. R., Lyons, T., & Saunders, K. E. A. (2018). A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder. Translational psychiatry, 8(1), 274–274.

Path signature

Path $X_t : [a,b] \to \mathbb{R}^d$ $\longrightarrow$ Signature $S(X)_{a,b} = (1, \underbrace{S(X)^1_{a,b}, \ldots, S(X)^n_{a,b}}_{\text{level 1}}, \underbrace{S(X)^{1,1}_{a,b}, S(X)^{1,2}_{a,b}}_{\text{level 2}}, \ldots)$

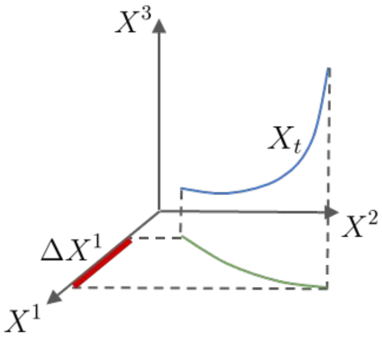

$$S(X)^i_{a,b} = \int_a^b dX^i_t = X^i_b - X^i_a = \Delta X^i$$

$$S(X)^{i_1,\ldots,i_k}_{a,b} := \int_{a<t<b} S(X)^{i_1,\ldots,i_{k-1}}_{a,t}\, dX^{i_k}_t$$

$$= \int_{a<t_k<b} \int_{a<t_{k-1}<t_k} \cdots \int_{a<t_1<t_2} dX^{i_1}_{t_1}\, dX^{i_2}_{t_2} \ldots dX^{i_k}_{t_k}$$

ANITI

Université de Toulouse

Path signature

Path $X_t : [a,b] \to \mathbb{R}^d$ $\longrightarrow$ Signature $S(X)_{a,b} = (1, \underbrace{S(X)^1_{a,b}, \ldots, S(X)^n_{a,b}}_{\text{level 1}}, \underbrace{S(X)^{1,1}_{a,b}, S(X)^{1,2}_{a,b}}_{\text{level 2}}, \ldots)$

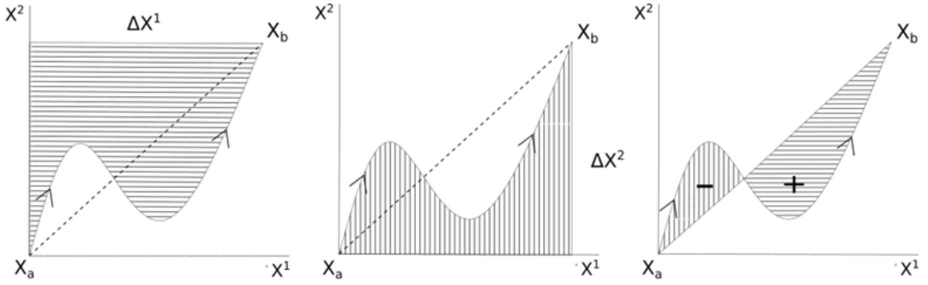

(a) $S(X)^{1,2}_{a,b}$

(b) $S(X)^{2,1}_{a,b}$

(c) $\frac{1}{2}(S(X)^{1,2}_{a,b} - S(X)^{2,1}_{a,b})$ $\longrightarrow$ lead matrix

Properties

- It fully **characterizes** paths up to tree-like equivalence (paths which retrace themselves along some subsection)
- Any continuous classification boundary in the path space can be approximated by a **linear** boundary in the signature space
- Can be **efficiently** computed, and used in an online setting
- Invariant under translation and time reparametrization

ANITI

Université de Toulouse

Application on BCI

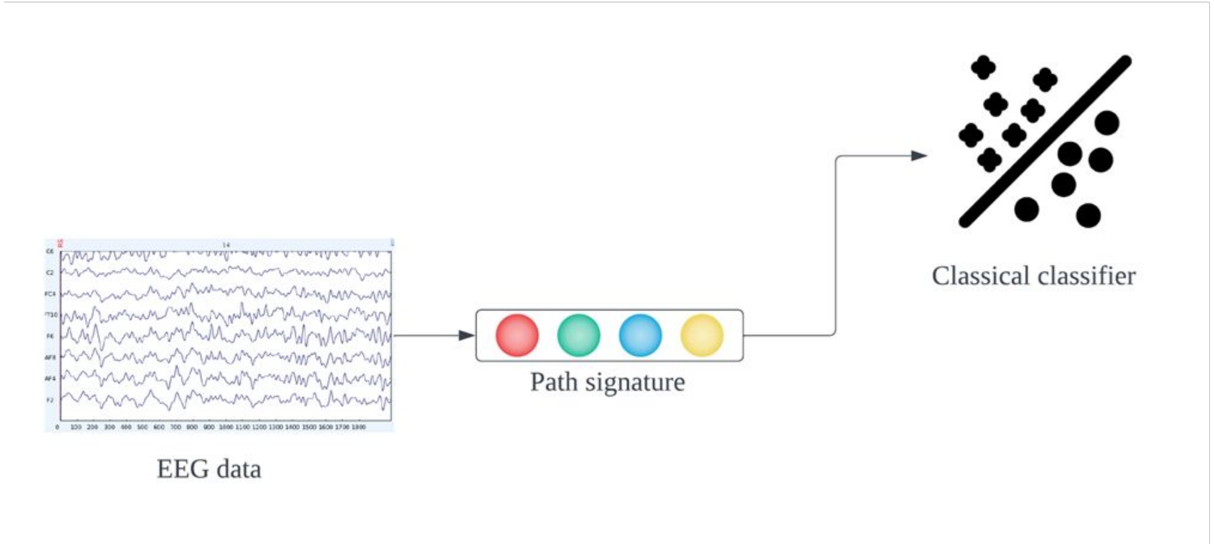

Results

Features: study of truncation levels

	level	1	2	3	4
intra	SVM	50.5(12.9)	66.1(15.2)	55.7(13.3)	56.9(13.9)
	LDA	51.0(12.6)	63.4(14.0)	54.7(14.2)	54.6(14.2)
	LR	50.9(11.3)	**67.1**(13.5)	56.6(13.3)	56.6(14.0)
	RF	49.7(11.6)	59.7(16.1)	58.0(16.8)	56.6(18.0)
	MLP	52.3(11.2)	61.6(16.4)	54.8(13.3)	56.8(15.1)
inter	SVM	52.6(3.6)	53.9(6.1)	54.2(6.2)	52.8(4.9)
	LDA	52.5(4.9)	53.5(6.2)	53.2(4.9)	53.9(6.0)
	LR	53.5(4.4)	54.7(5.8)	54.4(6.4)	53.6(4.9)
	RF	51.0(4.5)	54.6(6.0)	54.2(7.8)	52.0(4.6)
	MLP	53.1(5.6)	**58.7**(8.3)	56.5(5.2)	52.2(5.7)

- The lead-lag relationship captured by the level 2 signature seems to be relevant with the underlying neural mechanisms

Cyclicity analysis

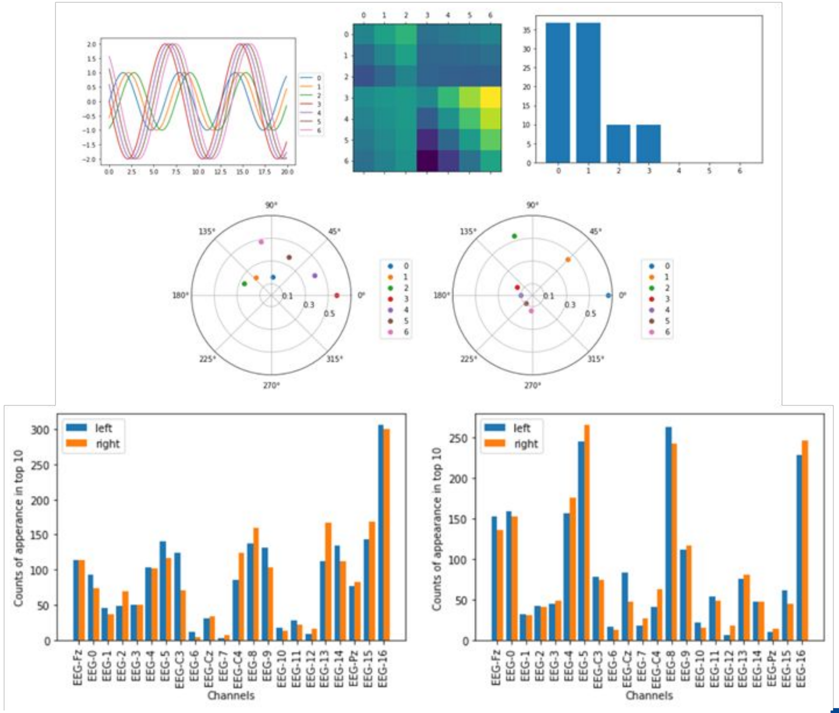

ANITI
Université de Toulouse

Application on BCI

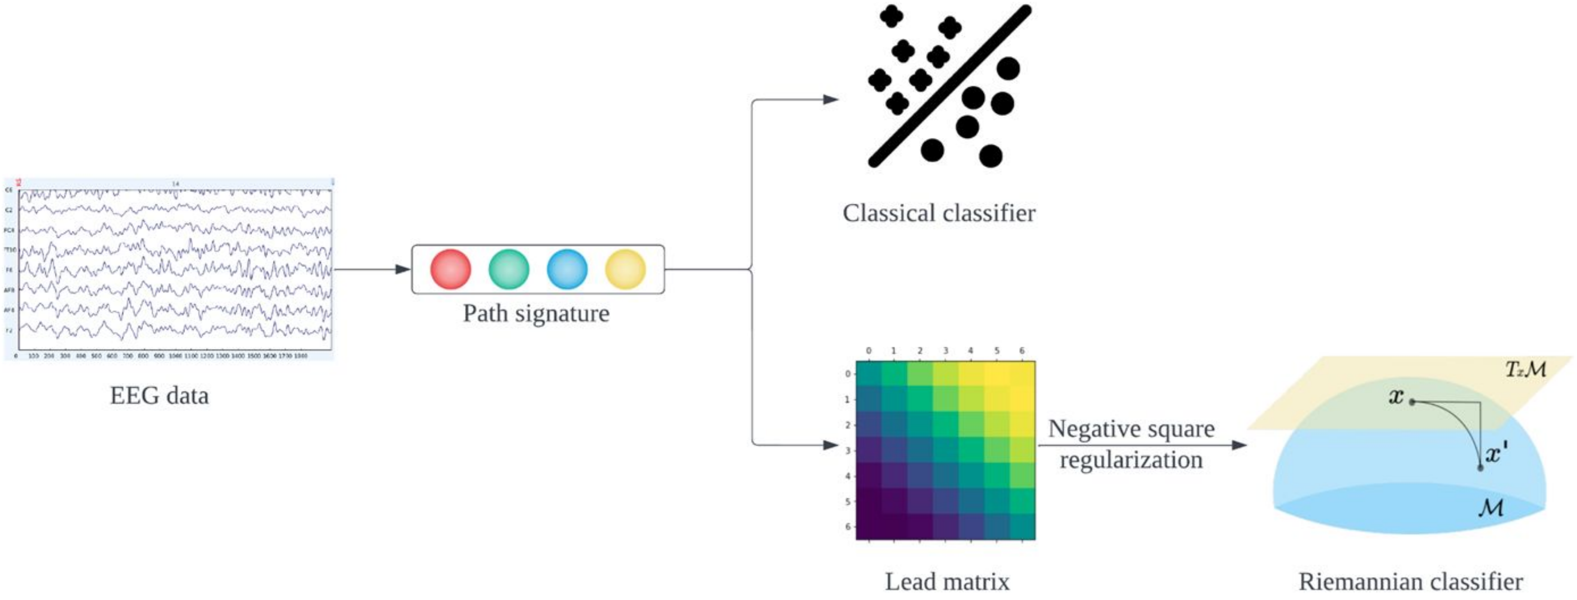

Results

	#channels	#classes	Signature intra	Signature inter	Covariance intra	Covariance inter
BCI Competition IV 2a	22	2	71.4(18.1)	66.1(11.8)	81.1(16.6)	69.2(15.9)
Physionet MI-BCI	64	2	60.1(23.9)	47.0(11.0)	63.8(24.2)	46.2(14.8)
Passive BCI competition	61	3	88.9(10.5)	41.4(7.4)	90.9(9.3)	42.0(5.6)
BCI Competition IV 2a	22	2	71.4(18.1)	66.1(11.8)	81.1(16.6)	69.2(15.9)
Physionet MI-BCI	22	2	58.2(24.5)	51.6(12.4)	64.1(24.7)	51.5(16.0)
Passive BCI competition	21	3	58.0(17.7)	39.7(6.1)	71.2(15.9)	36.9(5.4)

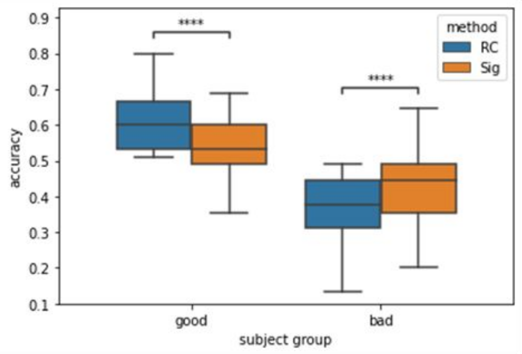

- The signature-based matrices are more robust to inter-subject variability than covariance matrices, especially on noisy and low-quality data

Xiaoqi Xu, Darrick Lee, Nicolas Drougard, Raphaëlle N. Roy (2023). Signature methods for brain-computer interfaces. Accepted for publication in Scientific Reports

ANITI
Université de Toulouse