

SOPHIA

SPRINGBOARD FOR ARTIFICIAL INTELLIGENCE

EDITION 2022

BOOKLET ABSTRACTS

23 > 25 NOVEMBER 2022

📍 SOPHIA ANTIPOLIS

[SOPHIA-SUMMIT.COM](https://sophia-summit.com)
#SOPHIASUMMIT



Table of contents

3D electromechanical modelling for cardiac resynchronisation therapy planning, Gaëtan Desrues [et al.]	5
A Flood of Problems Drained – Fixing Valves the Smart Way, Christoph Kögler	7
A General Theory for Client Sampling in Federated Learning, Yann Fraboni . . .	8
A comparative cost assessment of coalescing epidemic control strategies in heterogeneous artificial social-contact networks., Jan Broekaert [et al.]	10
A new dense hybrid visual odometry approach with multi-masks, Ziming Liu [et al.]	12
A new dense hybrid visual odometry approach with multi-masks, Ziming Liu [et al.]	13
ABEILLE: a novel method for ABerrant Ex-pression Identification empLoying machine Learning from RNA-sequencing data, Justine Labory	14
AI for mobility assistance or reeducation with quantified evaluation of the subject’s motor functions, Romain Tissot [et al.]	16
Artificial intelligence technologies and firms’ innovation: evidence from hospital suppliers, Sofia Patsali [et al.]	18
Automatic Argumentation Quality Assessment with Neural Graph Embeddings, Santiago Marro [et al.]	20
Causal based fairness : a methodology to use counterfactual fairness, Christèle Tarneç [et al.]	22
ChemOddity : Machine Learning to predict Odorant Detection Threshold., Maxence Lalis [et al.]	24

Data collection methodology for Responsible Artificial Intelligence in Health, Andreea Oprea [et al.]	25
Decentralized Federated Learning (DFL) platform, Frederic Guyard [et al.]	27
Deep Learning Framework for Model Correction in Cardiac Electrophysiological Imaging, Victoriya Kashtanova [et al.]	29
Deep learning on tabular data for the travel industry, Simon Nanty [et al.]	31
Don't fear the unlabelled: safe semi-supervised learning via simple debiasing, Hugo Schmutz [et al.]	33
Explainable Electrocardiogram Analysis with Wave Decomposition: Application to Myocardial Infarction Detection, Yingyu Yang [et al.]	35
Fallacious Argument Classification in Political Debates, Pierpaolo Goffredo [et al.]	36
Federated Learning for Data Streams, Othmane Marfoq [et al.]	37
Federated Learning under Heterogeneous and Correlated Client Availability, Angelo Rodio [et al.]	39
Forward Modelling of M/EEG: Towards a New Automatic Head and Brain Tissue Segmentation System, Ludovic Corcos [et al.]	41
Foveated Neural Networks, Matteo Tiezzi [et al.]	43
Human decision-making in the presence of algorithm adviser: An Experimental emphasis., Mathieu Chevrier [et al.]	44
Human vs AI: The role of human capital in forecast automation, Artyom Yepremyan [et al.]	45
Informative labels in Semi-Supervised Learning, Aude Sportisse [et al.]	47
Integrating machine learning methods to single-cell experimental workflow increases throughput and accuracy for target identification in immuno-oncology, Marielle Péré [et al.]	49
Learning to act under uncertainty: application to bio-inspired olfactory search, Aurore Loisy [et al.]	51
Learning to act under uncertainty: application to bio-inspired olfactory search, Aurore Loisy [et al.]	53

Local Model Reconstruction Attacks in Federated Learning and their Uses, Ilias Driouich [et al.]	55
Localization of Brain Microstructure Changes Associated with Alzheimer’s Disease using Class Activation Maps, Aymene Mohammed Bouayed [et al.]	56
Lucky Luke Maintainer: Detect camera’s defects before seeing them, Antoine Amalric [et al.]	58
NERD (Nematode EffectoR Discovery) : a tool to predict proteins involved in nematodes’ plant parasitism., Djampa Kozłowski	59
On the (Non-)Reliance on Algorithms - A decision-theoretic account, Bernard Sinclair-Desgagné	61
On the Two-fold Role of Logic Constraints in Deep Learning, Gabriele Ciravegna [et al.]	62
Optimal Transport for Graph Representation Learning, Cédric Vincent-Cuaz [et al.]	63
Predicting odorant-receptor activation with protein language and graph neural networks, Matej Hladiš [et al.]	65
Prospects of AI-augmented medical treatment acceptance, Zakaria Babutsidze [et al.]	66
Quantum Decision Modeling for the Travel Industry, Federico Tiblias [et al.]	68
Quantum-based Sentiment Analysis for the Travel Industry, Massimiliano Pronesti [et al.]	70
Real-Time High-Resolution Traffic Monitoring with Distributed Acoustic Sensing and AI, Yacine Khacef [et al.]	72
Refrigerant leak detection in industrial refrigeration systems with vapor compressors, Amal Mtibaa [et al.]	73
Shallow convolutional neural network with rank-1 Fourier domain weights for brain signal classification, Sara Sedlar [et al.]	75
Speeding up convergence in decentralized multi-agent learning through smart neighbor picking, Marina Costantini [et al.]	77
The Zero-Inflated Poisson Dynamic Latent Block Model with an Application to Pharmacovigilance Data, Giulia Marchello [et al.]	79

The graph embedded topic model, Dingge Liang [et al.]	80
Unsupervised Echocardiography Registration through Patch-based MLPs and Transformers, Yingyu Yang [et al.]	81
” Why do I have chickenpox? ” Template-based explanations for correct and incorrect diagnosis., Benjamin Molinet [et al.]	82
Author Index	82

3D electromechanical modelling for cardiac resynchronisation therapy planning

Gaëtan Desrues ^{*† 1,2}, Serge Cazeau ³, Maxime Sermesant ⁴

¹ Inria Sophia Antipolis - Méditerranée – Institut National de Recherche en Informatique et en Automatique – France

² Microport CRM – Institut National de Recherche en Informatique et en Automatique : EpioneTeam, MicroPort CRM – France

³ Microport CRM – MicroPort CRM – France

⁴ Inria Sophia Antipolis - Méditerranée – Institut National de Recherche en Informatique et en Automatique : EpioneTeam – France

Despite important medical advances over the last few decades, cardiovascular diseases remain a leading cause of death globally and the number one cause of death in the EU. Cardiovascular diseases accounts for 45% of all deaths in Europe according to the European Commission (2017). Half of this mortality is due to heart failure (HF), which can be caused or further aggravated by electrical dyssynchrony. Cardiac resynchronization therapy (CRT) is a procedure used to correct such dyssynchrony by inserting electrodes into the cardiac muscle to artificially deliver electrical signals and ensure a coordinated contraction. However, 30% of CRT candidates are non responders. In the last decades, personalized electromechanical models of the heart have been developed and turn to be great assets for clinicians. They allow for better patient and therapy selection, and help in CRT response analysis.

However, personalisation of 3D models is still challenging due to the complexity of the system itself (coupling of electrophysiology, haemodynamics and mechanics), large number of parameters to be tuned and high computational demand.

We present a fully automatic pipeline for cardiac simulation, including personalisation of anatomy, electrophysiology and mechanics. First, we introduce the anatomical model, including anisotropic meshing and His-Purkinje network generation. Then we show how the complex electrophysiological activity can be modelled for fast simulations with the Eikonal model, and introduce a pacing method compatible with the fast Purkinje conduction network. 12-lead Electrocardiogram (ECG) are generated from the activation map. Finally, we use the Bestel-Clement-Sorine model to describe the cardiac electromechanical activity. The heart is described as a passive isotropic hyperelastic material accounting for elasticity and friction in the cardiac extracellular matrix surrounding the fibres. The activation map computed beforehand is coupled to the active orthotropic contraction part, which accounts for the active stress along cardiac fibres and elasticity between sarcomeres and Z-discs.

To personalize the model, we couple machine learning methods with optimisation algorithms and minimize simulations' output with routine patient data: ECG, echocardiography data and general patient information. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES), a stochastic method for real-parameter optimisation of non linear functions is used, together with supervised learning methods. We compare parametric methods (linear and RBF kernel regression) and non-parametric ones (k-nearest neighbors, multilayer perceptron) for regression

*Speaker

†Corresponding author: gaetan.desrues@inria.fr

problems.

Personalisation of anatomical data shows great results with error below 5% on matched indices such as ventricles volume and wall thickness, allowing to generate normal, hypertrophic and dilated geometries. In this study, we use a cohort of 90 patients undergoing CRT implantation, with different pacing configurations, to validate the personalisation of both electrophysiological and mechanical models, as well as the prediction to CRT response. We also obtained very low ($< 1\%$) for the electrophysiological personalisation and 19% error on predicted CRT outcome, on the available data (QRS axis and duration). Mechanical personalisation error is still high for some indices and prediction to CRT response is ongoing work, but already shows good tendency on important indices such a left pre-ejection interval, reported to be a good descriptor of resynchrony. Further work will include prediction methods of the optimal pacing location for each patient and hopefully avoid long and repeated trials in operation room during implantation of CRT devices.

We have successfully personalised cardiac models on patients undergoing CRT implantation and show how these models can be a great tool for clinicians to reduce the implantation time and the non-response rate of CRT candidates. Machine learning helps in building patient-specific simulations that behaves very closely to the patient's cardiac function. Endorsement of automatic patient-specific cardiac modelling methods into the clinical word is more and more feasible with the development of machine learning algorithms. They improve the analysis of our current knowledge of the cardiac function and make it possible to obtain physiological results, that can have direct impact on patients' life, in ever shorter times.

Keywords: cardiac modelling, CRT, ECG, patient, specific simulation, stimulation

A Flood of Problems Drained – Fixing Valves the Smart Way

Christoph Kögler * ¹

¹ Infineon Technologies AG – Königsbrücker Str. 180 01099 Dresden, Germany

What began as two co-innovation projects between Infineon, GlobalFoundries, T-Systems and other partners, has developed into a blue-print for providing integrated AI systems combining hardware and software. By combining MEMS sensor technology, machine learning and edge computing, the team brought essential stability to two extremely important components of highly automated semiconductor production: the ultra-pure water supply and the overhead hoist transport (OHT) system. The first test runs were carried out in the GlobalFoundries fab in Dresden and the one at Infineon will follow next. Christoph Kögler, former head of IoT & Cloud Solutions at T-Systems and now Director Innovation, Digitalization and AI at Infineon, gives a first-hand insight into these two applications.

The goal of the projects was to maintain the valves and transport vehicles in time and, above all, without production downtime. For example, if an OHT vehicle failed during operation, it would block parts of the rail system in the clean room and thus slow down production. Accordingly, the complex mechanics of the vehicles have so far been monitored manually and maintained as a precaution. During the two three months projects, the team developed a novel, easy-to-use and highly scalable solution for the complex task of monitoring pumps and transport vehicles without adding anything to the moving machinery itself. The results are impressive: Based on customer feedback, the OHT Health Predictor alone can reduce unscheduled downtime in the affected areas by 25 to 35 percent. The implemented IoT solution includes everything from sensors to data pre-processing in edge computing hardware and display on dashboards. The team has selected suitable hardware, implemented machine learning algorithms for early detection of failures – e.g., ball bearing and brake damage – and realized a flexible cloud platform for evaluating the pump valves and vehicles. A configurable dashboard shows high-level KPI's e.g., a damage estimation indicator for every machine creating a dynamic priority list for maintenance scheduling and planning.

We will show not only the business rationale for the customer to invest into these two solutions, but also why a specific Infineon MEMS microphone was chosen and how the projects were organized as a collaboration between different highly specialized partners. We will present the technical architecture consisting of several stages of smart sensors, edge computing hardware and cloud-native applications and derive important insights for other customer-facing cloud projects, e.g., the need for scalability-by-design. And, finally, we will provide important lessons learned from the implementation of these projects – and some fun stories that inevitably happen when hardware, software, client expectations and stark reality meet each other. The presentation will conclude with an outlook on current developments using the experience and data gained from the collaboration between the partners.

Keywords: P2S, Condition Monitoring, Predictive Maintenance, Cloud, Edge, IoT, AI

*Speaker

A General Theory for Client Sampling in Federated Learning

Yann Fraboni * ¹

¹ Inria – L’Institut National de Recherche en Informatique et en Automatique (INRIA) – France

Federated Learning (FL) has gained popularity in the last years as it enables different clients to jointly learn a global model without sharing their respective data. Among the different FL approaches, federated averaging (FedAvg) has emerged as the most popular optimization scheme. An optimization round of FedAvg requires data owners, also called clients, to receive from the server the current global model which they update on a fixed amount of Stochastic Gradient Descent (SGD) steps before sending it back to the server. The new global model is then created as the weighted average of the client updates, according to their data ratio.

A drawback of naive implementations of FedAvg consists in requiring the participation of all the clients to every optimization round. As a consequence, the efficiency of the optimization is limited by the communication speed of the slowest client, as well as by the server communication capabilities. To mitigate this issue, the original FedAvg algorithm already contemplated the possibility of considering a random subset of m clients at each FL round such that in expectation the resulting global model is identical to the one obtained when considering all the clients. The current default unbiased sampling scheme consists in sampling m clients according to a Multinomial Distribution (MD), where the sampling probability depends on a client data ratio. Nevertheless, when clients have identical amount of data, clients can also be sampled uniformly without replacement. In this case, Uniform sampling has been experimentally shown to yield better results than MD sampling.

Previous works proposed alternative to MD and Uniform sampling with the aim of improving FL convergence. In (2), MD sampling was extended to account for clusters of clients with similar data characteristics, while in (3), clients sampling probabilities are defined depending on the Euclidean norm of the clients local work. While these works are based on the definition and analysis of specific sampling procedures, there is currently a need for a general theoretical framework to elucidate the impact of client sampling on FL convergence.

The main contribution of this work consists in deriving a general theoretical framework for FL optimization allowing to clearly quantify the impact of client sampling on the global model update at any FL round. This contribution has important theoretical and practical implications. First, we demonstrate the dependence of FL convergence on the variance of the aggregation weights. Second, we prove for the first time that the convergence speed is also impacted through sampling by the resulting covariance between aggregation weights. From a practical point of view, we establish both theoretically and experimentally that client sampling schemes based on aggregation weights with sum different than 1 are less efficient. We also prove that MD sampling is outperformed by Uniform sampling only when clients have identical data ratio. Finally, we show that the comparison between different client sampling schemes is appropriate only when considering a small number of clients. Our theory ultimately shows that MD sampling should be used as default sampling scheme, due to the favorable statistical properties and to the resilience

*Speaker

to FL applications with varying data ratio and heterogeneity. In this work, we also showed that our theory encompasses advanced FL sampling schemes, such as the ones proposed in (2) and (3). Finally, while the contribution of this work is in the study of the impact of a client sampling on the global optimization objective, further extensions may focus on the analysis of the impact of clients selection method on individual users' performance, especially in presence of heterogeneity.

This work was accepted at FL-IJCAI'22 and is available at <https://arxiv.org/abs/2107.12211>.

- (1) Chen et al. (2020), Optimal Client Sampling for Federated Learning.
- (2) Fraboni et al. (2021), Clustered sampling: Low-variance and improved representativity for clients selection in federated learning.
- (3) Khaled et al. (2020), Tighter theory for local sgd on identical and heterogeneous data.
- (4) Li et al. (2020), On the convergence of fedavg on non-iid data.
- (5) Wang et al. (2020), Tackling the objective inconsistency problem in heterogeneous federated optimization.

Keywords: federated learning, client sampling, data heterogeneity, convergence rate, SGD

A comparative cost assessment of coalescing epidemic control strategies in heterogeneous artificial social-contact networks.

Jan Broekaert ^{*† 1}, Davide La Torre ¹, Faizal Hafiz ¹, Marco Repetto ²

¹ SKEMA Business School – Skema Business School – Campus Sophia Antipolis 60 rue Dostoïevski
06902 Sophia Antipolis, France

² Digital Industries, Siemens Italy – 4 Via Privata Vipiteno, Milan, 20128, Italy, Italy

The outcome of an epidemic depends on the proper disease characteristics, the heterogeneity of the receptive agents, the spatial distribution and connectivity of the agents, and the variety of deployed control strategies aimed at containing the epidemic burden. In practice, the deployment of the control strategies is constrained by economic, social, technological, and scientific factors, which often depend on regional contingencies. Our model allows a comparative cost assessment for two connected subpopulations of a network that deploy coalescing control strategies, coinciding or divergent, depending on their respective asymmetric resources; a contact confinement strategy (‘strict & short’ vs ‘weak & long’), a vaccination strategy (‘early & lower-efficiency’ vs ‘later & higher-efficiency’), or any possible combination of both strategies. The population-specific *infection load*, *death load*, and the deployed control strategies incur different costs in each subpopulation. The multi-objective goal of a minimized epidemic burden and a minimized economic burden from *policy cost* and *productivity impact* is compounded using a scaled scalarization method. (Charpentier *et al* 2020, Gros *et al* 2021). The optimal outcome is obtained from the Nash equilibrium in the ‘2-player’ compounded pay-off matrix of the total cost and is provided for a scale range of deployment cost and GDP difference between the two players. Our network-based epidemic model implements a probabilistic Susceptible-Exposed-Infectious-Removed-Dead (SEIRSD) dynamics and also includes effects of *ego-network* support and an effect of loss of immunity by lack of exposure. The repeated simulations ($n = 120$) are run on heterogeneous artificial social-contact networks ($N=1000$) with high clustering coefficients and high cycle sub-graphs occurrence to realistically mimic observed person-person contact networks (Pastor-Satorras *et al*, 2015). The model parameters for the epidemic diffusion over the network are fitted to reported observations of the COVID-19 pandemic. (Godio *et al* 2020, Bjornstad *et al* 2020, korolev.2020}. The effects of the adopted Nash Equilibrium control strategies on the oscillatory and endemic phases of the epidemic are demonstrated on the SEIRD phase-diagrams (Fig. 1) and discussed.

References:

Charpentier, A., Elie, R. and Laurière, M. and Tran, V.C., (2020), COVID-19 pandemic control: balancing detection policy and lockdown intervention under ICU sustainability. *Math. Model. Nat. Phenom.*, 15, 57.

Gros, C., Valenti, R., Schneider, L., Valenti, K. and Gros, D., (2021), Containment efficiency and control strategies for the corona pandemic costs. *Scientific Reports*, 11,1, 6848.

*Speaker

†Corresponding author: jan.broekaert@skema.edu

Pastor-Satorras, R., Castellano, C., Van Mieghem, P. and Vespignani, A., (2015), Epidemic processes in complex networks, *Rev. Mod. Phys.*, 87, 3, 925-979.

Godio, A. , Pace, F. and Vergnano, A., (2020), SEIR-Modeling of the Italian Epidemic of SARS-CoV-2 Using Computational Swarm Intelligence, *International journal of environmental research and public health*, 17,10,3535

Bjornstad, O.N., Shea, K., Krzywinski, M. and Altman, N. (2020), The SEIRS model for infectious disease dynamics, *Nature Methods*, 17, 6, 557-558.

Korolev, I. (2021), Identification and estimation of the SEIRD epidemic model for COVID-19, *Journal of econometrics*, 220,1, 63-85.

Fig. 1 SEIRSD phase-diagrams for the baseline epidemic on one rendition of the non-confined artificial social-contact network (N=1000).

Keywords: Cost Optimization, Health Policy, Network dynamics, Nash equilibrium, SEIRSD epidemic, COVID, 19

A new dense hybrid visual odometry approach with multi-masks

Ziming Liu *¹, Philippe Martinet¹, Ezio Malis¹

¹ Inria – L’Institut National de Recherche en Informatique et en Automatique (INRIA) – France

Automated vehicles have been mature with new technologies in sensors and algorithms. For an autonomous driving vehicle, the basic requirement is to know where it is, then make the plan and control the environment. When we want an autonomous vehicle to localize itself in real-time, it is required to enrich the extracted information, including geometry information and abstract semantic information (1). The ability to detect and recognize all encountered situations and then decide the appropriate behavior is essential. Visual odometry is an important part of the localization module of autonomous vehicles. Recent advances in deep learning approaches have given rise to hybrid visual odometry approaches that combine both deep networks and traditional pose estimation methods (2). One limitation of deep learning approaches is the availability of ground truth data needed to train the neural networks. For example, it is extremely difficult, if not impossible, to obtain a ground truth dense depth map of the environment to be used for stereo visual odometry. Even if unsupervised training of networks has been investigated, supervised training remains more reliable and robust. In this paper, we propose a new hybrid dense stereo visual odometry approach in which a dense depth map is obtained with a network that is supervised using ground truth poses that can be more easily obtained than ground truth depths maps (3). However, machine learning methods generate hallucinated depths even in areas where it is impossible to estimate the depth due to several reasons, like occlusions, homogeneous areas, etc. Generally, this produces wrong depth estimation that leads to errors in odometry estimation. To avoid this problem, we propose a new approach to generate multiple masks that will be combined to discard wrong pixels and therefore increase the accuracy of visual odometry (4).

(1) A. I. Comport, E. Malis, and P. Rives, "Real-time quadrifocal visual odometry," *IJRR*, vol. 29, no. 2-3, pp. 245–266, 2010.

(2) Zhan H, Weerasekera C S, Bian J W, et al. Visual odometry revisited: What should be learnt? in *ICRA*. IEEE, 2020: 4203-4210.

(3) Z. Liu, E. Malis, and P. Martinet, "A new dense hybrid stereo visual odometry approach," in *IROS*, IEEE, 2022

(4) Z. Liu, E. Malis, and P. Martinet, "Multi-masks Generation for Increasing Robustness of Dense Direct Methods," in *ICRA reviewed*, IEEE, 2023

Keywords: autonomous driving, depth estimation, visual odometry, localization, hybrid methods

*Speaker

A new dense hybrid visual odometry approach with multi-masks

Ziming Liu *¹, Ezio Malis¹, Philippe Martinet¹

¹ Inria – L’Institut National de Recherche en Informatique et en Automatique (INRIA) – France

Automated vehicles have been mature with new technologies in sensors and algorithms. For an autonomous driving vehicle, the basic requirement is to know where it is, then make the plan and control the environment. When we want an autonomous vehicle to localize itself in real-time, it is required to enrich the extracted information, including geometry information and abstract semantic information (1). The ability to detect and recognize all encountered situations and then decide the appropriate behavior is essential. Visual odometry is an important part of the localization module of autonomous vehicles. Recent advances in deep learning approaches have given rise to hybrid visual odometry approaches that combine both deep networks and traditional pose estimation methods (2). One limitation of deep learning approaches is the availability of ground truth data needed to train the neural networks. For example, it is extremely difficult, if not impossible, to obtain a ground truth dense depth map of the environment to be used for stereo visual odometry. Even if unsupervised training of networks has been investigated, supervised training remains more reliable and robust. In this paper, we propose a new hybrid dense stereo visual odometry approach in which a dense depth map is obtained with a network that is supervised using ground truth poses that can be more easily obtained than ground truth depths maps (3). However, machine learning methods generate hallucinated depths even in areas where it is impossible to estimate the depth due to several reasons, like occlusions, homogeneous areas, etc. Generally, this produces wrong depth estimation that leads to errors in odometry estimation. To avoid this problem, we propose a new approach to generate multiple masks that will be combined to discard wrong pixels and therefore increase the accuracy of visual odometry (4).

(1) A. I. Comport, E. Malis, and P. Rives, "Real-time quadrifocal visual odometry," *IJRR*, vol. 29, no. 2-3, pp. 245–266, 2010.

(2) Zhan H, Weerasekera C S, Bian J W, et al. Visual odometry revisited: What should be learnt? in *ICRA*. IEEE, 2020: 4203-4210.

(3) Z. Liu, E. Malis, and P. Martinet, "A new dense hybrid stereo visual odometry approach," in *IROS*, IEEE, 2022

(4) Z. Liu, E. Malis, and P. Martinet, "Multi-masks Generation for Increasing Robustness of Dense Direct Methods," in *ICRA reviewed*, IEEE, 2023

Keywords: autonomous driving, depth estimation, visual odometry, localization, hybrid methods

*Speaker

ABEILLE: a novel method for ABerrant Ex-pression Identification empLoying machine Learning from RNA-sequencing data

Justine Labory * ^{1,2}

¹ Maison de la Modélisation, de la Simulation et des Interactions [Sophia-Antipolis] – Université Côte d’Azur – 1361 Route des Lucioles 06560 Valbonne, France, France

² Medical Data Laboratory – MDLab - MSI - Université Cote d’Azur – France

Omics technologies have revolutionized the world of biology and medicine tailoring the way to next-generation healthcare. The advent of whole-exome sequencing (WES) and whole-genome sequencing (WGS) has greatly accelerated the identification of variants in previously unknown rare disease genes. Although these technologies are mainstays in Mendelian disease diagnosis, their success rate for detecting causal variants is far from complete, ranging from 25 to 50%. Several variants remain as variants of unknown significance (VUS) or they are missed due to the inability to prioritize them. Recently the employ of RNA sequencing (RNA-seq) has been proposed. This technology provides a direct probing of RNA abundance and sequence of both coding and non-coding genome, including allele-specific expression and splice isoforms. Despite the very promising premises of RNA-seq to detect new responsible genes, the pioneering works employing this technique improved the diagnostic power of only 10% (Labory et al., 2020). Transcriptome analysis facilitates genome-wide interpretation of DNA variants, specifically three aberrant events can be analyzed: aberrant expression, aberrant splicing, allelic imbalance or allele-specific ex-pression. Currently, some improvements have been achieved in the identification of aberrant splicing events and allele-specific expression thanks to the development of novel tools designed for rare disease (Mertes et al., 2021). On the other hand, the identification of aberrant gene expression (AGE) in this context requires a paradigm-shift toward a novel way to analyze gene expression data. Usually, to identify AGE, a comparison of gene expression level between two groups of individuals with different conditions such as: healthy/diseased, exposed/unexposed to treatment, or others, is carried out. The significance of the results is measured by statistical tests whose statistical power decreases if the two groups are made up of too few individuals. This approach is not adapted in the context of rare diseases. By definition, rare diseases concern only a very small number of subjects, thus the availability of large cohorts is uncommon. Most importantly, very often, replicates for the same individual are not available in such clinical context. Furthermore, the disposal of a control group, such as healthy individuals, is often limited. Moreover, the heterogeneity of these pathologies is very high, which means that the same disease present in several patients will not be due to the same responsible genes. Consequently, these individuals cannot be combined to constitute the diseased group. Therefore, classical statistical methods cannot be used for the detection of AGEs in the context of rare diseases. Machine learning methods via neural networks including autoencoders (AEs) or variational autoencoders (VAEs) have shown promising performances in medical fields (Pratella et al., 2021).

*Speaker

Here, we describe ABEILLE, (ABerrant Expression Identification empLoying machine LEarning from sequencing data), a novel method for the identification of AGE from RNA-seq data without the need of replicates or a control group, using a flexible model obtained after testing several parameters (Labory et al., 2022).

ABEILLE combines the use of a VAE (variational autoencoder), able to model any data without specific assumptions on their distribution, and a decision tree to classify genes as AGE or non-AGE. An anomaly score is associated to each gene in order to stratify AGE by severity of aberration.

We compare ABEILLE performances to the state-of-the-art alternatives by using semi-synthetic data and a real dataset from the study by Kremer et al. (Kremer et al., 2017), demonstrating the importance of the flexibility of the VAE configuration to identify potential pathogenic candidates.

References:

Labory J, Fierville M, Ait-El-Mkadem S, Bannwarth S, Paquis-Flucklinger V, Bottini S. Multi-omics approaches to improve mitochondrial disease diagnosis: challenges, advances and perspectives. *Frontiers*. 2020

Mertes, C., Scheller, I. F., Yépez, V. A., Çelik, M. H., Liang, Y., Kremer, L. S., ... & Gagneur, J. (2021). Detection of aberrant splicing events in RNA-seq data using FRASER. *Nature communications*, 12(1), 1-13.

Pratella D, Ait-El-Mkadem S, Bannwarth S, Paquis-Flucklinger V, Bottini S. A Survey of Autoencoder Algorithms To Pave the Diagnosis of Rare Diseases. *International journal of Molecular Sciences*. 2021

Labory J, Le Bideau G, Pratella D, Yao JE, , Ait-El-Mkadem S, Bannwarth S, El-Hami L, , Paquis-Flucklinger V, Bottini S (2022). ABEILLE: a novel method for ABerrant Expression Identification empLoying machine Learning from RNA-sequencing data. *Bioinformatics*.
Kremer,L.S. et al. (2017) Genetic diagnosis of Mendelian disorders via RNA sequencing. *Nature Communications*, 8, 15824.

Keywords: mitochondrial diseases, diagnostic stalemate, RNAsequencing, autoencoder

AI for mobility assistance or reeducation with quantified evaluation of the subject's motor functions

Romain Tissot ^{*} ¹, Jean-Pierre Merlet ²

¹ HExapode, PHysiologie, AssISTance et Objets de Service – Inria Sophia Antipolis - Méditerranée, 3iA Côte d'Azur, Université de Nice Sophia-Antipolis – 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis cedex, France

² equipe HEPHAISTOS – L'Institut National de Recherche en Informatique et en Automatique (INRIA) – France

Cable-Driven Parallel Robots (CDPR) are constituted of cables that can be independently coiled and uncoiled which can induce movement of the attached end-effector in the workspace. Solving the positioning of this robots requires to take into account the sagging of the cable. In other words the model involves mass and elasticity of the cables which generates very complex equations that tend to be demanding in computational time without having the certainty to have all existing solutions.

Indeed, in the recent studies about CDPR, it has been shown that solving the Direct and Inverse Kinematics (DK/IK) may lead to find multiple solutions. With interval analysis (IA), it is possible to look for all solutions within intervals for the unknowns. Computation time will then depend on the size of the search space and while all solutions in this space will be found the algorithm will miss solutions outside this space.

Another option to generate solutions for the two kinematics analysis is to use Continuation. A first use is to discover solutions for the kinematic problems starting from the solutions that can be obtained if the cables were rigid. The principle, then, is to move the two physical parameters that characterize elasticity and sagging toward their real value by using small incremental steps and using Newton to determine the solution at each step. However this approach may also miss solutions because of singularities occurring during the algorithm. A second use of continuation is interesting for AI : starting from a small number of solutions obtained if IA or continuation, we determine new kinematics solutions by changing by small step the problem inputs. This allows one to obtain very quickly a large number of solutions, thereby providing a training set for AI.

Still IA and continuation are extremely computer intensive and this is why we recently wanted to see if artificial intelligence tools could help us to solve this issues. Using already known sets of equations and solution as training sets has shown that IA, although very fast, is globally bad at solving the system with huge errors (up to 300%). Better results are obtained with hybridation where the neural network (NN) outputs is used as an initial guess for the Newton method but still we get at most about 30 % of the solutions. We have identified the reason of this poor result : classical cost functions do not take well into account the structure of the equations and the respective influence of the unknowns on the value of the equations. We are currently investigating an hybridation of unsupervised learning where the cost function will be

^{*}Speaker

based on an indicator whose value should be below a fixed threshold to ensure the convergence of the Newton scheme. Still such NN will output at best a single solution and therefore it will be needed to build a large number of Nns, each one devoted to a specific branch among the possible solutions.

Keywords: cable, driven parallel robot, cable sagging, kinematics, neural network, unsupervised learning.

Artificial intelligence technologies and firms' innovation: evidence from hospital suppliers

Sofia Patsali ^{*† 1,2}, Jackie Krafft ¹, Michele Pezzoni ¹

¹ Groupe de Recherche en Droit, Economie et Gestion – Université Nice Sophia Antipolis (1965 - 2019), Centre National de la Recherche Scientifique : UMR7321, Université Côte d'Azur : UMR7321 – GREDEG - Bâtiment 2 - Campus Azur du CNRS - 250 rue Albert Einstein - CS 10269 - F06905 SOPHIA ANTIPOLIS Cedex, France

² Bureau d'économie Théorique et Appliquée – AgroParisTech, université de Strasbourg, Université de Lorraine, Centre National de la Recherche Scientifique : UMR7522, Institut National de Recherche pour l'Agriculture, l'Alimentation et l'Environnement : UMR1443 – Université de Lorraine, UFR Droit Sciences Economiques et Gestion, 13 place Carnot CO 70026, 54035 Nancy Cedex Université de Strasbourg, Faculté des Sciences Economiques et de Gestion, 61 avenue de la Forêt Noire 67085 Strasbourg Cedex, France

This work analyses the impact of supplying AI-driven medical equipment to hospitals on the innovative performance of suppliers.

Public engagement in developing AI technologies for health care has significantly expanded in France in recent years (Villani 2018, OPECST 2019). France is set to become a European "champion of AI" (OCDE 2020). In this respect, the French government has established many initiatives to enhance the economic impact of AI-driven technologies, such as the network of interdisciplinary institutes "3IA Institutes". In light of the above trends, policymakers, health professionals and academics are looking for insights into the economic and societal impacts of health-related AI technologies. So far, economists have studied the impact of AI on labour market and earnings inequality (Brynjolfsson et al. 2018, Autor and Salomons 2017, Acemoglu and Restrepo 2020, Acemoglu et al. 2020), neglecting the impact of AI on firms' innovative activities.

This paper aims to fill the gap by building an original dataset merging data on the procurement of AI-powered products to hospitals with firms' patent applications. Specifically, we draw upon three data sources. *First*, we use public procurement administrative and financial data provided by the Unicancer network of hospitals for 2010 - 2017. Unicancer hospitals are private non-profit health establishments exclusively devoted to treatment, research, and teaching in oncology. Unicancer suppliers are among the top 30 patent applicants worldwide in the domain of AI applications in oncology (WIPO 2019). Compared to existing studies on procurement and innovation, Unicancer provided us with confidential information on the full list of selected and non-selected firms applying to become hospitals' suppliers, the criteria used to evaluate firms, and the final evaluation of firms on which the selection was based. *Second*, we retrieve the description for each supplied product from the Global Medical Device Nomenclature (GMDN). The GMDN is a system of internationally agreed descriptors used to identify medical device products and is managed by the Food and Drug Administration. *Third*, we retrieve selected and non-selected firms' patents applications from the European Patent Office statistical database (Patstat).

To assess the impact of procurement of AI-powered products on firms' innovative performance, a particular challenge is linking products to a set of technological domains that are potentially

*Speaker

†Corresponding author: sofia.patsali@univ-cotedazur.fr

affected by becoming a hospital supplier. Indeed, a large part of the firms in our sample are multinational companies innovating in a broad set of technological domains. Some of these domains are not relevant for the products supplied to hospitals. To identify the technological domains relevant for the supplied product, we assess the similarity of the GMDN product description with the abstracts of the patents attributed to the supplier firm. To do so, we use a neural network algorithm for text analysis that transforms documents into vectors according to the semantic meaning of the words appearing in their texts (Mikolov et al. 2013). Once identified the patents that are more similar to the product description, we extract their IPC (International Patent Classification) technology classes, and we consider these classes as the ones that are potentially influenced by becoming a supplier of the hospital.

Using a diff-in-diffs approach, we assess the impact of becoming a supplier on firms' propensity to patent in technology classes relevant to the product supplied. Our results show that firms selected as Unicancer suppliers have a significantly higher propensity to innovate if compared to firms that are not selected. Specifically, Unicancer suppliers file one patent application more per year in the three years after the procured contract than companies that candidated for being suppliers but were not selected.

References

Acemoglu, D., Lelarge, C., & Restrepo, P. (2020, May). Competing with robots: Firm-level evidence from France. In *AEA Papers and Proceedings* (Vol. 110, pp. 383-88).

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188-2244.

Autor, D., & Salomons, A. (2017, June). Robocalypse now: Does productivity growth threaten employment. In *Proceedings of the ECB Forum on Central Banking: Investment and Growth in Advanced Economie* (pp. 45-118).

Brynjolfsson, E., Rock, D., & Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The economics of artificial intelligence: An agenda* (pp. 23-57). University of Chicago Press.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Keywords: Innovation in health, artificial intelligence, medical equipment, neural network text analysis, Unicancer.

Automatic Argumentation Quality Assessment with Neural Graph Embeddings

Santiago Marro ^{*† 1}, Elena Cabrio ², Serena Villata ³

¹ Université Côte d’Azur, Inria, CNRS, I3S, France – Université Côte d’Azur, Inria, CNRS, I3S, France
– France

² Université Côte d’Azur, Inria, CNRS, I3S, France – Université Côte d’Azur, Inria, CNRS, I3S, France
– France

³ Université Côte d’Azur, Inria, CNRS, I3S – Université Côte d’Azur, Inria, CNRS, I3S – France

Argumentation is used by people both internally, by evaluating arguments and counterarguments to make a decision, and externally, e.g., by exchanging arguments to reach an agreement or to promote a position. A major component of the argumentation process concerns the assessment of a set of arguments and of their conclusions in order to establish their justification status, and therefore compute their acceptability degree. The assessment of the justification status of the statements supported by arguments allows the agent to decide what to believe and what to do. Argumentation semantics provide formal criteria to determine which sets of arguments (i.e., extensions) can be regarded as collectively acceptable (Baroni, Caminada, and Giacomin 2011). However, the assessment of the arguments acceptability is only a (basic) part of the complex assessment tasks required in argumentative processes in many everyday life applications, e.g., in medicine and education.

The issue of assessing an argumentation is particularly critical when considering the different aspects of artificial argumentation, from the identification of real natural language arguments and their relations in text, to the computation of the justification status of abstract arguments, to the gradual assessment of arguments. Despite some approaches addressing the automatic assessment of natural language arguments (Wachsmuth et al. 2017, 2020; Saveleva et al. 2021), this issue remains largely unsolved.

In this talk, we address this open issue and we answer the following research question: what are the basic quality dimensions to characterize natural language argumentation and how to automatically assess them?

More precisely, we propose an Argument Mining (AM) approach to identify and classify natural language arguments along with quality dimensions. In artificial argumentation, Argument(at)ion Mining aims at extracting arguments from text and at analyzing them (Cabrio and Villata 2018).

In this paper, we decide to characterize argument quality along with three quality dimensions for natural language argumentation, i.e., cogency, rhetoric and reasonableness. Cogency estimates the acceptability of the premises that are relevant to the argument’s conclusion and their sufficiency to draw the conclusion, rhetoric determines the rhetorical strategy employed in the argument’s conclusion (if any) from the three options of ethos, logos and pathos, and reasonableness rates if the argument adequately rebuts its counterarguments, assessing, therefore, the dialectical quality dimension of the argumentation.

*Speaker

†Corresponding author: santiago.marro@inria.fr

Our interest focuses on the education scenario, where students are asked to interact with our AM system to assess the quality of their persuasive essays with respect to these three quality dimensions. To train our AM model, we annotated an existing dataset of 402 student persuasive essays (Stab and Gurevych 2017) with these quality dimensions.

We then propose a new deep learning AM method based on a transformer architecture, exploiting the structure of the argumentation graph through graph embeddings. Our approach addresses in an automatic way the evaluation process proposed in social science by Stapleton and Wu (2015), through a scoring rubric for persuasive writing that integrates the assessment of both argumentative structural elements and reasoning quality. The obtained results are satisfactory and outperform standard baselines and similar approaches in the literature.

References

- Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *Knowl. Eng. Rev.* 26(4):365–410.
- Cabrio, E., and Villata, S. 2018. Five years of argument mining: a data-driven analysis. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 5427–5433. ijcai.org.
- Saveleva, E.; Petukhova, V.; Mosbach, M.; and Klakow, D. 2021. Graph-based argument quality assessment. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, 1268–1280.
- Stab, C., and Gurevych, I. 2017. Parsing argumentation structures in persuasive essays. *Computational Linguistics* 43(3):619–659.
- Stapleton, P., and Wu, Y. A. 2015. Assessing the quality of arguments in students’ persuasive writing: A case study analyzing the relationship between surface structure and substance. *Journal of English for Academic Purposes* 17:12–23.
- Wachsmuth, H., and Werner, T. 2020. Intrinsic quality assessment of arguments. In *Proceedings of the 28th International Conference on Computational Linguistics*, 6739–6745.
- Wachsmuth, H.; Naderi, N.; Hou, Y.; Bilu, Y.; Prabhakaran, V.; Thijm, T. A.; Hirst, G.; and Stein, B. 2017. Computational argumentation quality assessment in natural language. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, 176–187. Valencia, Spain: Association for Computational Linguistics.

Keywords: argument mining, argumentation, ai, transformers, nlp, argument quality

Causal based fairness : a methodology to use counterfactual fairness

Christèle Tarnec ^{*† 1}, Frederic Guyard ^{* ‡ 1}

¹ ORANGE – Orange – Orange Labs Sophia, 905 rue Albert Einstein, 06921 Sophia-Antipolis Cedex, France

AI enabled predictions are based on correlations between the **observed selected variables** and the predicted outcome, which make the prediction model not sufficiently robust but also **potentially unfair** as the mechanism by which social and historical direct and indirect discrimination influence the training data is not captured.

To develop fair AI-based services by design, causal reasoning seems to be valuable. Several fairness metrics based on causal reasoning are proposed by the literature, among which the counterfactual fairness.

Causal reasoning, in addition to fit intuition, can help to prevent misuse of correlation and improve explainability. Causal reasoning is based on reverse reasoning. on assumptions. For example: If I were living in another neighborhood "other things being equal" would I be selected for the job? In order to be able to perform such reasoning, a **causal graph** is needed. However, **discovering causal graph is in its infancy and a pertinent causal graph is hard to obtain and often not reliable.** Several algorithms exist to discover causal graph but the discovered graph may differ depending on the algorithm used and on its parameters choice. The causal graphs discovered by these algorithms typically present several issues: wrongly directed edges, un-existing discovered edge, partial discovery, etc.

We focused on **counterfactual fairness** and **demonstrated empirically the difficulty** to discover a unique causal graph and **to assert that a model is counterfactually fairer than another.** Counterfactual Fairness is graph dependent: when creating a model that is counterfactually fair, we use a discovered graph (i.e., a graph revealed by one of the existing algorithms to discover a graph) and not the real graph. If several different graphs are discovered, it is impossible to say which one is the more counterfactually fair. How can one assess if a model is fairer than another one?

Measuring with statistical fairness metrics seems not to be relevant as most of them compare the input distribution and the output distribution (all of them except Statistical Parity and Conditional Statistical Parity) and input distribution is potentially biased, as based on a discovered graph that may not match with the real graph.

Even when using Statistical Parity or Conditional Statistical Parity to say which model is the more counterfactually fair, there is no proof or theoretical results demonstrating that statistical notions can help answering this question.

We propose an approach to overcome this problem: **we setup a methodology** that can be used to evaluate fairness of different models, measure the sensibility and the stability of the fairness

*Speaker

†Corresponding author: christele.tarnec@orange.com

‡Corresponding author: frederic.guyard@orange.com

with respect to the graph discovered (generating the Real-World data distribution, computing causal strength, simulating an observation process etc.) This methodology offers the possibility to have various experimentations (sensitivity to unmeasured confounding for instance, other causal metrics, etc.), to use it upside down in practical cases (Discover graph then generate data according to the graph discovered and evaluate if the dataset is a representative sample of the generated data to strengthen causal discovery). We also show experimentally how to build counterfactual models when several causal graphs have been proposed by different algorithms or by different human experts.

Keywords: fairness, causal reasoning, causality, counterfactual, causal discovery graph, ethical AI

ChemOddity : Machine Learning to predict Odorant Detection Threshold.

Maxence Lalis *¹, Hladis Matej¹, Nicolas Baldovini¹, Jérémie Topin^{†1}

¹ Institut de Chimie de Nice UMR 7272, Université Côte d’Azur, CNRS, Nice, France – Université Côte d’Azur (UCA) – France

Odorant Detection Threshold (ODT) is the minimum concentration of an odorant that can be detected by human. Assessing ODT is complex, time consuming and requires both equipment and manpower. At this moment only a small subset of odorants has an ODT characterized. To get an estimate of this value, we propose a machine learning model that aims to determine the detection threshold for any molecules. A database of more than 1700 compounds has been gathered to train and evaluate a predictive model. We assessed several combinations of models’ architecture (Random Forest, k-Nearest Neighbors, Support Vector Machine, Graph Neural Network, Ensemble models) and molecular descriptors (3D molecular descriptors, fingerprints, graphs). We selected a Graph Neural Network, after optimizing parameters for each combination, whose performances are RMSE= 1.14, MAE=0.8 log (ODT) and an R2 of 0.61. Our model predicts detection limits ranging from parts-per-million concentration (ppm) to parts-per-trillion (ppt). This allows us to custom design chemicals with extremely low detection threshold. Such compounds can be used in various domains such as perfume and food industry or safety and pollution control. The predictive model and the entirety of the data will be available.

Keywords: Chemistry, Olfaction, Odorant Detection Threshold, Predictive Model, Graph Neural Network

*Speaker

[†]Corresponding author: jeremie.topin@univ-cotedazur.fr

Data collection methodology for Responsible Artificial Intelligence in Health

Andreea Oprescu , Gloria Miró-Amarante , Lutgardo García-Díaz , Ángel Chimenea-Toscano , Victoria E. Rey , Ricard Martínez-Martínez , María Del Carmen Romero-Terneró * ¹

¹ Universidad de Sevilla – C/ S. Fernando, 4, C.P. 41004-Sevilla, Spain

A medical field that is increasingly benefiting from Artificial Intelligence applications is Gynecology and Obstetrics. In previous work, we showed that artificial intelligence (AI) technology and obstetric control by physicians can improve pregnancy health, leading to better pregnancy outcomes and overall better experience, also reducing any possible long-term effects that can be produced by complications. This work presents a data collection methodology for responsible AI in health and a case study in the pregnancy domain. It is a qualitative descriptive study on the preferences and expectations expressed by pregnant women about responsible AI and affective computing. A 41 item structured interview was distributed among 150 pregnant patients attending prenatal care at Hospital Virgen del Rocío and Clinic CAREMUJER (Seville, Spain) during the months of October and November 2020. In this study, a substantial interest in intelligent pregnancy solutions has been revealed among pregnant women. Participants with a lower level of interest reported privacy concerns and a lack of trust in AI solutions. Regarding affective computing-based intelligent solutions specifically, most participants reported positively, and no significant difference was found between women having a healthy or a high-risk pregnancy on this matter. Our findings also suggest that there is a high demand for intelligent personalized solutions among participants. On the topic of sharing pregnancy data with the healthcare provider in favor of scientific research, pregnant women assisting public health care services were found to be more likely to share their data when the provider was a public health system rather than a private entity. Pregnant women who are interested in using an AI pregnancy application share a strong idea that it should be responsible, trustworthy, useful, and safe. Similarly, we found that pregnant women would change their mind about their concerns and they would feel more confident if the intelligent solution gives explanations about the system decisions and recommendations, as the XAI approach promotes.

Our study on intelligent healthcare systems focuses on aspects such as Human-Centered Design, eXplainable AI, Privacy, and Information Security by Design (3).

Obtaining data is a crucial step in an AI workflow. Traditionally, data for AI applications was obtained from medical records and medical instruments. Data fusion can be used to improve the quality of AI models, since it allows more data types, from wearable devices, environmental sensors, and smartphones, among others. However, as the type of data and the number of interconnected devices increase, the privacy and information security of the solutions must be reinforced. In this context of multimodal data, we propose a preliminary data collection methodology considering the seven key requirements AI systems must meet to be deemed trustworthy by the EU guidelines.

Our study shows that, in general, pregnant women showed interest in using an intelligent health-

*Speaker

care solution to assist them throughout their pregnancy. However, various obstacles to the adoption of technological solutions have been identified, mainly due to concerns about privacy and security, lack of digital skills, and skepticism about technology. Addressing privacy and information security requirements is vital for successful implementations of AI solutions in the healthcare sector.

The development of regulatory compliance processes for a technology cannot be static. The very nature of AI makes this principle a rule. We must understand that we are in an early phase of this technology, in which we have learned that data-analysis processes are very sensitive to bias. On the other hand, we begin to entrust basic decision-making processes to this technology, with legal or material consequences for its recipients. Proper compliance requires a state of ongoing monitoring and updating that is deployed at several levels: 1) learn from the operation of the technology itself. The results obtained, the operation errors, incidents, and ultimately any verified or verifiable element should also be indexed and studied by legal support. And not only to prevent potential conflicts and responsibilities, but above all to improve compliance conditions; 2) deepen the design of compliance by proposing improvements when necessary; 3) accompany each phase or evolution of the product.

Keywords: responsible artificial intelligence (RAI), explainable artificial intelligence (XAI), emotional computing affective computing, user, centered design, human, centered design, privacy, security, pregnancy

Decentralized Federated Learning (DFL) platform

Frederic Guyard ^{*† 1}, Hugo Miralles ², Tamara Tomic ^{* ‡ 3}

¹ ORANGE – Orange – Orange Labs Sophia, 905 rue Albert Einstein, 06921 Sophia-Antipolis Cedex, France

² Université Côte d’Azur, Orange, CNRS, Inria, I3S, France – Université Côte d’Azur, Orange Group, Centre National de la Recherche Scientifique - CNRS, L’Institut National de Recherche en Informatique et en Automatique (INRIA), Université Côte d’Azur, Inria, CNRS, I3S, France – France, France

³ Orange Labs [Sophia Antipolis] – Orange Labs – 905 rue Albert Einstein 06921 Sophia Antipolis France, France

The volumes of sensitive data collected on user equipment with computation capacities sufficient to train models tend to increase with 5G deployment and network virtualization. Orange as a trustful operator embraces development of services free of collecting user data, for example based on Federated Learning (FL) (1). Classical centralized FL collaboratively trains a unique global model over iterations on local data. It consists of the workers (W) that train models locally, the parameter manager (P) that configures and controls the system and the aggregator (A) unit that accumulates the learning knowledge of workers over iterations (2).

In this demonstration, we present a Decentralized Federated Learning (DFL) platform developed in Python with a flexible functional architecture that supports multi-thread and multi-session (several global models may be trained at the same time). The platform has a communication interface for passing messages between its components. The Federated Learning algorithms (W, A, and P) are coded and deployed by users, and then they are distributed by the platform among edge devices. We note that user equipment stores data and codes locally. Compared to FL, DFL platform has an additional component, the scheduler (S). It deals with a fleet of workers that potentially change over model training iterations (W registration/detachment). A single user equipment can host several training sessions, having different functionalities (W/S/P) per session. Moreover, multiple sessions can implement different communication topologies (e.g., star, fully connected) and various aggregation strategies (e.g., FedAvg, ClusteredFL).

We deploy DFL on the RaspberryPi (R) cluster, each R representing one user device. Apart the network router and graphical demonstrator units, the rest of the Rs have one of the functionalities of W, S and P per session, hosting several sessions on a same R. We demonstrate execution of 3 different learning sessions in parallel (different learning algorithms with different Python libraries and datasets).

By adding generic microservices (S and A), new Ws can register to the FL process over learning iteration and several data scientists can launch FL processes in parallel on the same fleet of devices. The DFL platform is distributed, scalable and portable to docker or virtual machine environments. Furthermore, at time the platform is only an implementation of the specifications, the platform could be developed using other tools (edge-specific libraries such as Tensorflow Lite) or other programming languages.

*Speaker

†Corresponding author: frederic.guyard@orange.com

‡Corresponding author: tamara.tomic@orange.com

This flexibility allows us to perform tests for use cases for business as well as research (heterogeneous data, devices, usages) on different complexity algorithms (ranging from MLP to LSTM) by simulating user equipment characteristics.

(1) P. Kairouz, et al., "Advances and open problems in Federated Learning," *Foundations and Trends in Machine Learning*, vol 14, issue 1-2, pp 1-210, <https://arxiv.org/abs/1912.04977v3>, 2021.

(2) H. Ludwig et al., IBM Federated Learning: an Enterprise Framework White Paper, <https://arxiv.org/pdf/2020>.

Keywords: Federated Learning

Deep Learning Framework for Model Correction in Cardiac Electrophysiological Imaging

Victoriya Kashtanova ^{*† 1,2}, Maxime Sermesant ^{1,2}

¹ Inria Sophia Antipolis - Méditerranée – Institut National de Recherche en Informatique et en Automatique – 2004 route des Lucioles BP 93 06902 Sophia Antipolis, France

² 3iA Côte d’Azur – 3iA Côte d’Azur – Sophia Antipolis, France

Imaging the electrical activity of the heart can be achieved with invasive catheterisation. However, the resulting data are sparse and noisy. Mathematical modelling of cardiac electrophysiology can help the analysis but solving the associated mathematical systems can become unfeasible. It is often computationally demanding, for instance when solving for different patient conditions. Furthermore, it is still difficult to reduce the discrepancy between the output of idealized mathematical models and clinical measurements, which are usually noisy.

To alleviate this limitation, we present a new framework to model the dynamics of cardiac electrophysiology at low cost. This framework is based on the integration of a low-fidelity (or incomplete) physical model and a learning component implemented here via neural networks. The latter acts as a complement to the physical part and handles all quantities and dynamics that the simplified physical model neglects. This construction allows the framework to learn from data of different complexity.

Using data of action potential of different origin (i.e in-silico data simulated via the Ten Tusscher - Panfilov ionic model, as well as real ex-vivo optical mapping data), we demonstrate that this framework allows us to reproduce the complex dynamics of the transmembrane potential and to correctly identify the relevant physical parameters, even when only partial and noisy measurements are available.

Overall, our results suggest that automated learning of cardiac electrophysiology dynamics is feasible and has great potential. This may be useful for applications concerning fast parameterization of computational heart models. Additionally, this combined model-based and data-driven approach could improve cardiac electrophysiological imaging and provide predictive tools.

(1) Aliev, R.R., Panfilov, A.V.: A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals* 7(3), 293–301 (1996)

(2) He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE conf. CVPR*. pp. 770–778 (2016)

(3) Kashtanova, V., Ayed, I., Arriuela, A., Potse, M., patrick gallinari, Sermesant, M.: Deep learning for model correction in cardiac electrophysiological imaging. In: *Int. Conf. MIDL* (2022)

*Speaker

†Corresponding author: victoriya.kashtanova@inria.fr

- (4) Pop, M., Sermesant, M., Lepiller, D., Truong, M.V., McVeigh, E.R., Crystal, E., Dick, A., Delingette, H., Ayache, N., Wright, G.A.: Fusion of optical imaging and mri for the evaluation and adjustment of macroscopic models of cardiac electro(1)physiology: A feasibility study. *Medical Image Analysis* 13(2), 370–380 (2009)
- (5) Ten Tusscher, K.H.W.J., Panfilov, A.V.: Alternans and spiral breakup in a human ventricular tissue model. *Am. J. Physiol. - Heart Circ. Physiol.* 291(3), H1088– H1100 (2006)

Keywords: AI for Healthcare, Deep Learning, Physics based learning, Cardiology, Electrophysiology, Simulations

Deep learning on tabular data for the travel industry

Simon Nanty * ¹, Michael Defoin-Platel *

1

¹ Amadeus SAS – Amadeus SAS – France

Machine learning for computer vision or natural language processing has seen major advances over the last decades with the emergence of deep learning. Successful deep neural network architectures and methods have been proposed on numerous homogeneous but unstructured data types. For example, depending on the data particularities, convolutional networks became the reference for image or speech data, while recurrent or transformer-based networks are now state of the art for textual data. However, many other high-impact applications of machine learning are supported by data that are structured in tabular format. These applications are very common in industrial sectors such as health, finance, or travel to name a few.

Unlike image or text data, tabular data are heterogeneous in the sense that they are composed of multiple types of attributes, such as continuous, ordinal, or discrete categorical features. Those features often have a different meaning, representing quantities of different units for example, and can also be very diverse in terms of distribution or cardinality (for categorical features). Moreover, the dependencies between them can be hard to model compared to spatial or temporal dependencies. Another aspect, found in many use cases and less prevalent in vision or language datasets, is the availability of pre-existing domain knowledge about these data. We can think of constraints on the target space or of monotonic relationships between some of the features and the target variable that we want to enforce. However, this domain knowledge can be quite difficult to integrate in a machine learning model.

For tabular data, the application of deep learning methods is still very challenging and has seen very limited successes in the literature so far. Classical machine learning models for tabular datasets, such as gradient boosted trees or random forests, remain the most widely used algorithms in industry or in machine learning competitions and still outperform most of the deep learning approaches proposed to address the issues linked to data heterogeneity and diversity (1, 4).

According to (1), these new approaches can be categorized into three main groups. The first category focuses on data transformation and representation such as feature encoding (2) or self-supervised reduction of the input features. A second group investigates advances on the architecture of neural network itself. In particular, some authors adapt to the tabular case the successful architectures or elements from other domains, such as convolutions, Transformers, Batch Normalization, Skip-Connection, etc. (2) Finally, in a third group, emphasis is put on regularization of network parameters with a wide variety of methods (3).

In this work, we review the most promising tabular deep learning advances and apply them to a large travel industry dataset. We then investigate different ways to integrate domain knowledge

*Speaker

into the model. Finally, we conclude with a discussion on what makes tabular data the next frontier of deep learning.

(1) Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., & Kasneci, G. (2021). Deep neural networks and tabular data: A survey. arXiv preprint arXiv:2110.01889.

(2) Gorishniy, Y., Rubachev, I., & Babenko, A. (2022). On Embeddings for Numerical Features in Tabular Deep Learning. arXiv preprint arXiv:2203.05556.

(3) Kadra, A., Lindauer, M., Hutter, F., & Grabocka, J. (2021). Well-tuned simple nets excel on tabular datasets. *Advances in neural information processing systems*, 34, 23928-23941.

(4) Shwartz-Ziv, R., & Armon, A. (2022). Tabular data: Deep learning is not all you need. *Information Fusion*, 81, 84-90.

Keywords: Deep learning, Tabular data

Don't fear the unlabelled: safe semi-supervised learning via simple debiasing

Hugo Schmutz *^{1,2,3,4}, Olivier Humbert⁵, Pierre-Alexandre Mattei^{4,6}

¹ Laboratoire Jean Alexandre Dieudonné – Université Nice Sophia Antipolis, Centre National de la Recherche Scientifique – France

² UMR E4320 – Université Côte d'Azur, Université Nice Sophia Antipolis (... - 2019), COMUE Université Côte d'Azur (2015 - 2019), Commissariat à l'énergie atomique et aux énergies alternatives : DRF/JOLIOT, COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019) – Faculté de Médecine, 28 avenue de Valombrose 06107 Nice Cedex 2, France

³ Modèles et algorithmes pour l'intelligence artificielle – Inria Sophia Antipolis - Méditerranée, Université Nice Sophia Antipolis (1965 - 2019), Laboratoire Jean Alexandre Dieudonné, Scalable and Pervasive softwARe and Knowledge Systems – 2004 route des Lucioles BP 93 06902 Sophia Antipolis, France

⁴ 3iA Côte d'Azur – 3iA Côte d'Azur – France

⁵ UMR E4320 – Université Côte d'Azur, Université Nice Sophia Antipolis (... - 2019), COMUE Université Côte d'Azur (2015 - 2019), Commissariat à l'énergie atomique et aux énergies alternatives : DRF/JOLIOT, COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019), COMUE Université Côte d'Azur (2015 - 2019) – Faculté de Médecine, 28 avenue de Valombrose 06107 Nice Cedex 2, France

⁶ Modèles et algorithmes pour l'intelligence artificielle – Inria Sophia Antipolis - Méditerranée, Scalable and Pervasive softwARe and Knowledge Systems, Université Nice Sophia Antipolis (... - 2019), Laboratoire Jean Alexandre Dieudonné – France

Semi-supervised learning (SSL) provides an effective means of leveraging unlabelled data to improve a model's performance. Even though the domain has received a considerable amount of attention in the past years, most methods present the common drawback of lacking theoretical guarantees. Our starting point is to notice that the estimate of the risk that most discriminative SSL methods minimise is biased, even asymptotically. This bias impedes the use of standard statistical learning theory and can hurt empirical performance. We propose a simple way of removing the bias. Our debiasing approach is straightforward to implement and applicable to most deep SSL methods. We provide simple theoretical guarantees on the trustworthiness of these modified methods, without having to rely on the strong assumptions on the data distribution that SSL theory usually requires. In particular, we provide generalisation error bounds for the proposed methods. We evaluate debiased versions of different existing SSL methods, such as the Pseudo-label method and Fixmatch, and show that debiasing can compete with classic deep SSL techniques in various settings by providing better calibrated models. Additionally, we provide a theoretical explanation of the intuition of the popular SSL methods.

*Speaker

Keywords: Semi, supervised learning, Unsupervised leaning, safe learning, asymptotic statistics, variance reduction, controle variatres, consistency

Explainable Electrocardiogram Analysis with Wave Decomposition: Application to Myocardial Infarction Detection

Yingyu Yang ^{*† 1,2}, Marie Rocher ³, Pamela Mocerì ³, Maxime Sermesant ^{1,2}

¹ Inria – L’Institut National de Recherche en Informatique et en Automatique (INRIA) – France

² UCA – University Côte d’Azur – France

³ CHU – CHU Nice – France

Automatic analysis of electrocardiograms with adequate explain-ability is a challenging task. Many deep learning based methods have been proposed for automatic classification of electrocardiograms. However, very few of them provide detailed explainable classification evidence. In our study, we explore explainable ECG classification through explicit decomposition of single-beat (median-beat) ECG signal. In particular, every single-beat ECG sample is decomposed into five subwaves and each subwave is parameterised by a Frequency Modulated Moebius. Those parameters have explicit meanings for ECG interpretation. In stead of solving the optimisation problem iteratively which is time-consuming, we make use of an Cascaded CNN network to estimate the parameters for each single-beat ECG signal. Our preliminary results show that with appropriate position regularisation strategy, our neural network is able to estimate the subwave for P, Q, R, S, T events and maintain a good reconstruction accuracy (with R2 score 0.94 on test dataset of PTB-XL) in a unsupervised manner. Using the estimated parameters, we achieve very good classification and generalisation performance on myocardial infarction detection on four different datasets. The features of high importance are in accordance with clinical interpretations.

Keywords: ECG analysis, Reconstruction, Explainable ML, Myocardial infarction classification

*Speaker

†Corresponding author: yingyu.yang@inria.fr

Fallacious Argument Classification in Political Debates

Pierpaolo Goffredo ^{*} ¹, Shohreh Haddadan , Vorakit Vorakitphan , Serena Villata[†] ², Elena Cabrio ³

¹ Laboratoire d'Informatique, Signaux, et Systèmes de Sophia Antipolis – Centre National de la Recherche Scientifique : UMR7271, Université Côte d'Azur, 3iA Côte d'Azur, L'Institut National de Recherche en Informatique et en Automatique (INRIA) – 2000, route des Lucioles - Les Algorithmes - bât. Euclide B 06900 Sophia Antipolis, France

² Université Côte d'Azur, Inria, CNRS, I3S – Université Côte d'Azur, Inria, CNRS, I3S – France

³ Université Côte d'Azur, Inria, CNRS, I3S, France – Université Côte d'Azur, Inria, CNRS, I3S, France – France

Because of their significance in critical thinking, fallacies play a key role in arguing since antiquity. Their function is even more important now, since modern argumentation technologies confront difficult problems such as detecting deceptive and manipulating arguments in news stories and political debates, as well as generating counter-narratives.

Despite significant progress in this regard, classifying arguments as being fallacious remains a difficult and unresolved process.

We make two contributions: first, we present a novel annotated resource of 39 political debates from the U.S. Presidential Campaigns, where we annotated six main categories of fallacious arguments (e.g., ad hominem, appeal to authority, appeal to emotion, false cause, slogan, slippery slope), leading to 1628 annotated fallacious arguments; second, we take on this novel task of classifying fallacious arguments and define a neural architecture based on transformers outperforming state-of-the-art results and standard baselines. Our findings demonstrate the significance of argument relations and components in this task.

Keywords: NLP, Fallacy, Classification, Contest, Political Debates, Corpora

*Speaker

†Corresponding author: villata@i3s.unice.fr

Federated Learning for Data Streams

Othmane Marfoq *^{1,2}, Giovanni Neglia², Richard Vidal³, Laetitia Kameni³

¹ Accenture Technology Labs – Accenture – Sophia Antipolis, France

² Inria Sophia Antipolis – Inria Sophia Antipolis - Méditerranée, Inria Sophia Antipolis - Méditerranée – France

³ Accenture – Accenture Labs [Sophia Antipolis] – France

Federated learning (FL) (1) is an effective solution to train machine learning models on the increasing amount of data generated by IoT devices and smartphones while keeping such data localized. Heterogeneity remains a key and fundamental challenge in federated learning, and several recent works explored different sources of heterogeneity. In particular personalized federated learning (2), is proposed to handle the statistical heterogeneity resulting from the non-i.i.d.-ness of the data across clients in FL. Ensemble distillation (3), sub-model training (4), and local memorization(5) handle system heterogeneity resulting from varying computational and memory capabilities across clients.

Federated learning usually involves the minimization of an objective function, which is only available through unbiased estimates of its gradients. This problem, also referred to as stochastic approximation, encapsulates both cases when the objective function is the expected or empirical risk. The first case corresponds to the scenario where one can sample new data points at every iteration, while the latter one refers to the case where a fixed dataset is available before the beginning of the learning process and never updated afterward.

Most previous works on federated learning, e.g., (1, 2, 3), focus on the second case, i.e., the minimization of the empirical risk. They assume that every client collects and stores first all the samples it will use during the training process. Waiting for clients to collect reasonably enough data before starting the learning can be sub-optimal (or even impossible) in many cases, because (1) this approach ignores new samples collected by the clients after the start of the training, (2) clients may have limited memory capacities, and can not store a large number of data samples. In extreme settings (e.g., sensors network), some clients are only able to use the last few samples they collected, due to hardware restrictions. It is therefore necessary to design federated algorithms able to learn from a data stream under limited memory capacities.

In this work, we first formulate and study the problem of streaming federated learning, and highlight a new source of heterogeneity, which we name temporal data collection heterogeneity. To the best of our knowledge streaming FL and temporal data collection heterogeneity were not previously studied in the context of federated learning. We characterize the effect of such heterogeneity through a new bias-variance trade-off controlled by the relative importance of older samples in comparison to newer ones. We propose a general FL algorithm (see Appendix) to learn from data streams through an opportune weighted empirical risk minimization. Our theoretical analysis provides insights to configure such algorithm, and we evaluate its performance on a wide range of machine learning tasks and different data arrival patterns (see Appendix).

*Speaker

Our work highlights the challenges brought by learning from data streams in the context of federated learning. It provides the first formulation of this problem and sets an initial foundation to study it. We believe that this work opens the door for further research and study of federated learning from data streams.

References:

- (1) McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. "Communication-efficient learning of deep networks from decentralized data." In *Artificial intelligence and statistics*, pp. 1273-1282. PMLR, 2017.
- (2) Smith, Virginia, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. "Federated multi-task learning." In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 4427-4437. 2017.
- (3) Lin, Tao, Lingjing Kong, Sebastian U. Stich, and Martin Jaggi. "Ensemble distillation for robust model fusion in federated learning." *Advances in Neural Information Processing Systems* 33 (2020): 2351-2363.
- (4) Horvath, Samuel, Stefanos Laskaridis, Mario Almeida, Ilias Leontiadis, Stylianos Venieris, and Nicholas Lane. "Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout." *Advances in Neural Information Processing Systems* 34 (2021): 12876-12889.
- (5) Marfoq, Othmane, Giovanni Neglia, Richard Vidal, and Laetitia Kamani. "Personalized Federated Learning through Local Memorization." In *International Conference on Machine Learning*, pp. 15070-15092. PMLR, 2022.

Keywords: Federated Learning, Data Stream, Stochastic Approximation

Federated Learning under Heterogeneous and Correlated Client Availability

Angelo Rodio *^{1,2}, Francescomaria Faticanti², Othmane Marfoq^{2,3},
Giovanni Neglia^{1,2}, Emilio Leonardi⁴

¹ 3iA Côte d’Azur – 3iA Côte d’Azur – France

² Inria Sophia Antipolis – Inria Sophia Antipolis - Méditerranée, Inria Sophia Antipolis - Méditerranée – France

³ Accenture Technology Labs – Accenture – Sophia Antipolis, France

⁴ Politecnico di Torino – Italy

The enormous amount of data generated by mobile and IoT devices motivated the emergence of distributed machine learning training paradigms. Federated Learning (FL) is an emerging framework where geographically distributed devices (or clients) participate in the training of a shared Machine Learning (ML) model without sharing their local data. FL was proposed to reduce the overall cost of collecting a large amount of data as well as to protect potentially sensitive users’ private information. In the original Federated Averaging algorithm ({FedAvg}) (1), a central server selects a random subset of clients from the set of available clients and broadcasts them the shared model. The sampled clients perform a number of independent Stochastic Gradient Descent (SGD) steps over their local datasets and send their local model updates back to the server. Then, the server aggregates all the received client updates to produce a new global model, and a new training round begins. In each iteration of {FedAvg}, typically a few hundred devices are chosen randomly by the server to participate.

In real-world scenarios, the availability/activity of clients is dictated by exogenous factors that are beyond the control of the orchestrating server and hard to predict. For example, only smartphones that are idle, under charge, and connected to broadband networks are commonly allowed to participate in the training process. These eligibility requirements can make the availability of devices correlated over time and space. For example, {temporal correlation} may origin from a smartphone being under charge for a few consecutive hours and then ineligible for the rest of the day. Similarly, the activity of a sensor powered by renewable energy may depend on natural phenomena intrinsically correlated over time (e.g., solar light, wind circulation). {Spatial correlation} refers instead to correlation across different clients, which often emerges as consequence of users’ geographical distribution. For example, clients in the same time zone often exhibit similar availability patterns, e.g., due to time-of-day effects.

Temporal correlation in the data sampling procedure is known to negatively affect the performance of ML training even in the centralized setting (2) (3) and can potentially lead to {catastrophic forgetting}: the data used during the final training phases can have a disproportionate effect on the final model, "erasing" the memory of previously learned information. Catastrophic forgetting has also been observed in FL, where clients in the same geographical area have more similar local data distributions and clients’ participation follows a cyclic daily pattern (leading to spatial correlation) (4). Despite this evidence, a theoretical study of the convergence of FL algorithms under temporally and spatially correlated client participation is still missing.

*Speaker

This work provides the first convergence analysis of {FedAvg} under heterogeneous and correlated client availability. We assume that clients' temporal and spatial availability follows an arbitrary finite-state Markov process: this assumption models a realistic scenario in which the activity of clients is correlated and, at the same time, still allows the analytical tractability of the system. Our theoretical analysis (i) quantifies the negative effect of correlation on the algorithm's convergence rate through an additional term depending on the spectral properties of the Markov chain; (ii) points out a trade-off between two conflicting objectives: slow convergence to the optimal model, or fast convergence to a biased model, i.e., a model that minimizes an objective function different from the initial target. Guided by insights from the theoretical analysis, we propose {CA-Fed}, a federated learning algorithm which dynamically assigns weights to clients and balances the trade-off between maximizing convergence speed and minimizing model bias. Interesting that {CA-Fed} can decide to ignore clients with low availability and large correlation and outperforms other state-of-the-art algorithms (5), with the incidental but appreciated consequence of also reducing the overall computation or communication costs.

(1) McMahan, Brendan, et al. "Communication-efficient learning of deep networks from decentralized data." *Artificial intelligence and statistics*. PMLR, 2017.

(2) Sun, Tao, et al. "On markov chain gradient descent." *Advances in neural information processing systems*, 2018.

(3) Doan, Thinh T., et al. "Finite-time analysis of stochastic gradient descent under Markov randomness." *arXiv preprint arXiv:2003.10973*, 2020.

(4) Eichner, Hubert, et al. "Semi-cyclic stochastic gradient descent." *International Conference on Machine Learning*. PMLR, 2019.

(5) Ribero, Monica, et al. "Federated Learning Under Intermittent Client Availability and Time-Varying Communication Constraints." *arXiv preprint arXiv:2205.06730*, 2022.

Keywords: Federated learning, distributed optimization, heterogeneous device participation

Forward Modelling of M/EEG: Towards a New Automatic Head and Brain Tissue Segmentation System

Ludovic Corcos ¹, Theodore Papadopoulo ¹, Samuel Deslauriers-Gauthier
* ¹

¹ Centre Inria d'Université Côte d'Azur – L'Institut National de Recherche en Informatique et en Automatique (INRIA) – France

Magnetoencephalography and electroencephalography (together M/EEG) are imaging modalities that allow the non-invasive measurement of the magnetic field and the electric potential generated by cortical activity. Inferring which brain areas generated the observed M/EEG measurements is not a trivial task and is referred to as the inverse problem. A common way to solve the problem is to assume that brain sources act like current dipoles in a volume conductor, in this case the head whose geometry can be obtained from magnetic resonance imaging (MRI). The relationship between brain sources and M/EEG measurements can therefore be modeled, a process called the solving forward problem. This process can be seen as injecting anatomical priors into the inverse problem (1). However, extracting the anatomical information from MRI needed to solve the forward problem is lengthy and tedious with existing tools. In this work, we present the first step in the creation of an automated pipeline to generate a volume conductor model from T1 and T2 images.

Our objective is to obtain a 7-tissue segmentation of the MRI of each individual, consisting of skin, bone, muscle, white matter, gray matter, cerebrospinal fluid, and background. Note that existing segmentation systems do not consider the skin, bone, and muscle layers, instead focus only on the brain (2). Given this segmentation, a 3D model can be built where a specific electrical conductivity is assigned to each label. To segment the MRI volumes, we trained our own deep learning (DL) architecture, which is an adaptation of the U-NET model, under Tensorflow. The training dataset was generated from 2226 T1 and T2 MRI volumes provided by the Human Connectome Project. These data are not segmented and must therefore be annotated to generate suitable training data for our DL architecture. First, we developed a tool to remove the background from images based on the watershed technique. Next, we non-linearly registered all subjects to a reference (subject 0) in order to obtain the deformation field for each individual. We also performed an automatic segmentation of the brain itself using the FSL FAST (3). In parallel, we manually segmented the skin, bone, muscle, and cerebrospinal fluid in the MRI of patient 0 (only 4 tissues). By back projecting these two segmentations (fast + manual) into native subject space, we obtained the 7-tissue segmentation in subject specific space for all subjects which was used to train our U-net.

Figure 1 illustrates the segmentation obtained with our system in a few seconds. We can observe the accurate segmentation of the different tissues which will be assigned specific conduction values for the construction of the volume conductor. This is an important first step in the creation of an automated pipeline to solve the forward problem in M/EEG.

*Speaker

Keywords: MEG, EEG, Segmentation, Forward Modelling, U, net

Foveated Neural Networks

Matteo Tiezzi , Simone Marullo , Alessandro Betti * ¹, Enrico Meloni ,
Lapo Faggi , Marco Gori ², Stefano Melacci

¹ Université Côte d'Azur, Inria, CNRS, Laboratoire I3S, Maasai team, Nice, France – 3iA Côte d'Azur
– France

² Université Côte d'Azur – University Côte d'Azur, 3iA Côte d'Azur – France

The classic computational scheme of convolutional layers leverages filter banks that are shared over all the spatial coordinates of the input, independently on external information on what is specifically under observation and without any distinctions between what is closer to the observed area and what is peripheral. In this talk we propose to go beyond such a scheme, introducing the notion of Foveated Convolutional Layer (FCL), that formalizes the idea of location-dependent convolutions with foveated processing, i.e., fine-grained processing in a given-focused area and coarser processing in the peripheral regions. We show how the idea of foveated computations can be exploited not only as a filtering mechanism, but also as a mean to speed-up inference with respect to classic convolutional layers, allowing the user to select the appropriate trade-off between level of detail and computational burden. FCLs can be stacked into neural architectures and we evaluate them in several tasks, showing how they efficiently handle the information in the peripheral regions, eventually avoiding the development of misleading biases. When integrated with a model of human attention, FCL-based networks naturally implement a foveated visual system that guides the attention toward the locations of interest, as we experimentally analyze on a stream of visual stimuli.

*Speaker

Human decision-making in the presence of algorithm adviser: An Experimental emphasis.

Mathieu Chevrier * ¹, Eric Guerci , Brice Corgnet , Julie Rosaz

¹ Mathieu CHEVRIER – Groupe de Recherche en Droit, Economie et Gestion (GREDEG) – 250 Rue Albert Einstein, Sophia Antipolis 06560 Valbonne, France

The objective of this paper is first to propose an experimental design that allows us to measure the appreciation and the aversion that people have towards an algorithm in a risky and repeated task.

Second, we observe the impact of the different strategies proposed by the algorithm advisor on the participants. We show in this paper that people like algorithms that use a biased strategy more than those that use the optimal strategy in a risky context.

Keywords: Algorithm aversion/appreciation, Experimental economics, Human preferences, risk and uncertainty, Cognitive bias, Strategy

*Speaker

Human vs AI: The role of human capital in forecast automation

Artyom Yepremyan * ¹, Francesco Castellaneta , Bruno Cirillo

¹ SKEMA Business School – SKEMA Business School-UCA – France

Forecasts about future developments in firms' internal and external environments are a critical source of firms' competitive advantage (Barney, 1986). To this purpose, firms have traditionally relied on the forecasting ability of their human capital. However, individuals' limitations to the access and interpretation of information about past and future trends can undermine the firm's forecasting accuracy (Durand, 2003). Scholars have thus documented how the use of artificial intelligence has progressively substituted individuals in firms' forecasting tasks. Despite the evidence that machine learning algorithms can fix human biases (Choudhury et al., 2020), there are also increasing concerns that algorithms further perpetuate human bias (Cowgill & Tucker, 2019). Therefore, recent research has started enquiring human-machine complementarities (Choudhury et al., 2020) as a solution for improving machine learning predictions accuracy.

Nevertheless, we still know little about the nature of human-machine interaction as source of a firm's forecasting ability, especially in contexts characterized by lack of information and high uncertainty. Under these circumstances, the algorithmic bias requires close attention. Within the scope of this study, we aim to shed light on the trade-off between substitution and complementarity of human capital and artificial intelligence and we ask: "How and under what conditions does human capital help improve machine learning's predictive performance?".

To tackle this question, we used macroeconomic and financial data on a sample of U.S. publicly listed firms and matched it with data on historical analyst forecasts to train a Random Forest (RF) algorithm that predicts the earnings per share of firms in our sample. We perform two sets of analyses. First, we compare the accuracy of human predictions with those of the RF. To identify the likely mechanisms explaining the findings, we explore how this comparison varies with the level of uncertainty and information asymmetry, and with various characteristics of the analysts' human capital. Our results indicate that AI is more accurate than humans, thus corroborating the substitution effect. However, when the analyst has low experience on a firm and there is a high level of information asymmetry the analyst tends to outperform the RF. These results highlight that individuals have a competitive edge over machine learning when exploring new contexts characterized by the lack of hard, codified data. These results provide robust support to the idea that the nature of the human capital and artificial intelligence interaction is complex, and it can vary with different types and dimensions of human capital.

With these considerations in mind, we performed a second set of analyses where we modelled such interaction by integrating analysts forecasts in the training set of another RF model. Thus, we compared the estimated of this model with those of our first RF trained only on historical codified data. The results indicate that the RF trained with human inputs produces, on average, lower predictive error, thus in line with prior studies. However, in contrast with prior

*Speaker

studies, they also show that the RF trained with human inputs performs worse when it builds on human capital associated with high experience on the focal firm and in cases characterized by high uncertainty. Nevertheless, such experience turns to be useful to the performance of the RF when there is high uncertainty on the focal firm.

Our study contributes to the emerging literature on human-machine interaction (Choudhury et al., 2020; Krakowski et al., 2022) by exploring the conditions under which human capital and machine learning solution should substitute or complement each other. Thus, we add to the literature on strategic human capital by shedding light on what type of human capital has competitive edge over AI and can complement it.

Bibliography

Barney, J. B. (1986). Strategic Factor Markets: Expectations, Luck, and Business Strategy. *Management Science*, 32(10), 1231–1241. <https://doi.org/10.1287/mnsc.32.10.1231>

Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, 41(8), 1381–1411. <https://doi.org/10.1002/smj.3152>

Cowgill, B., & Tucker, C. E. (2019). Economics, fairness and algorithmic bias. *Preparation for: Journal of Economic Perspectives*.

Durand, R. (2003). Predicting a firm's forecasting ability: The roles of organizational illusion of control and organizational attention. *Strategic Management Journal*, 24(9), 821–838. <https://doi.org/10.1002/smj.339>

Krakowski, S., Luger, J., & Raisch, S. (2022). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, smj.3387. <https://doi.org/10.1002/smj.3387>

Keywords: AI, machine learning, human capital, human machine complementarities, strategic decision making

Informative labels in Semi-Supervised Learning

Aude Sportisse *^{1,2,3}, Pierre-Alexandre Mattei^{1,2,3}, Charles Bouveyron^{1,2,3}, Hugo Schmutz^{1,2,3}, Oliver Humbert^{1,3,4}

¹ 3iA Côte d’Azur – 3iA Côte d’Azur – France

² Inria Sophia Antipolis – Inria Sophia Antipolis (MAASAI) – France

³ Université Côte d’Azur – Université Côte d’Azur – Parc Valrose, 28, avenue Valrose 06108 Nice Cedex 2, France

⁴ Centre de Lutte contre le Cancer Antoine Lacassagne [Nice] – Unicancer, Université Côte d’Azur – 33, avenue de Valombrose, 06189 Nice, France

Missing values can be considered as *unavailable* entries in a data set, and are often coded as NA (for Not Available) in code languages. They can occur for many reasons: unanswered questions in a survey, lost data, sensing machines that fail, aggregation of multiple sources, etc. Classical statistical methods cannot be directly applied on the data sets which contain missing values (you can just have in mind that computing $NA+1$ is impossible) (1). In data science, most people delete all the rows (or all the columns) which contain missing values. Most of the time, this naive strategy is unfortunately not suitable: (i) there can be a huge loss of information, by deleting entire rows or columns, (ii) it is rare that a sub-population of the data is representative of the general population. This last situation raises the problem of the information contained in a missing value. If the process that causes the data to be missing, called the missing-data mechanism, depends on the data values themselves, the missing values are said informative. It is for example the case when rich people are less inclined to reveal their income. This work focuses on this case of missing values, which is the most realistic case but also the most challenging one. In semi-supervised learning (2), we have access to features but the outcome variable is missing for a part of the data. In real life, although the amount of data available is often huge, labeling the data is costly and time-consuming. It is particularly true for image data sets: images are available in large quantities on image banks but they are most of the time unlabeled. It is therefore necessary to ask experts (doctors if they are medical images) to label them (assign them a class, an output variable). In this context, people are more inclined to label images of some classes which are easy to recognize. The unlabeled data are thus informative missing values, because the unavailability of the labels depends on their values themselves. Typically, the goal of semi-supervised learning is to learn predictive models using all the data (labeled and unlabeled ones). Again, classical methods do not consider the missing data mechanism and lead to biased estimates if the missing values are informative (3). We aim at designing new semi-supervised algorithms that benefit from theoretical guarantees and than handle missing labels possibly informative.

(1) Little, R. J., & Rubin, D. B. (2019). *Statistical analysis with missing data* (Vol. 793). John Wiley & Sons.

(2) Chapelle, O., Scholkopf, B., & Zien, A. (2009). Semi-supervised learning (Chapelle, O. et al., eds.; 2006)(book reviews). *IEEE Transactions on Neural Networks*, 20(3), 542-542.

*Speaker

(3) Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., & Goodfellow, I. (2018). Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems*, 31.

Keywords: semi supervised learning, informative labels, missing not at random values

Integrating machine learning methods to single-cell experimental workflow increases throughput and accuracy for target identification in immuno-oncology

Marielle Péré *^{1,2}, Asma Chalabi[†], Diego Oyarzun³, Madalena Chaves[‡]
¹, Jérémie Roux[§]^{1,2}

¹ Inria Sophia Antipolis - Méditerranée – Institut National de Recherche en Informatique et en Automatique, Équipe Biocore – France

² Institut de Recherche sur le Cancer et le Vieillissement – Institut National de la Santé et de la Recherche Médicale : U1081, Centre National de la Recherche Scientifique : UMR7284, Université Côte d’Azur – France

³ University of Edinburgh – United Kingdom

Cell response heterogeneity upon treatment is a main obstacle in preclinical development of efficacious cancer drugs, due to the emergence of drug-tolerant cells. We have previously developed a **single-cell workflow, Fate-Seq** (Meyer *et al.* 2020), to **profile drug-tolerant persisters**, based on predictions of their drug response. To achieve this goal, *Fate-Seq* couples 3 single-cell techniques (Bian *et al.* 2022): first the prediction of the cell response phenotype (resistant or sensitive) for clonal cancer cells treated with a chosen drug, then the separation of the predicted resistant cells from the predicted sensitive ones and finally the RNA sequencing of the cells.

To **automatize** and **increase the prediction throughput**, we present 3 major improvements in our workflow using **machine learning models to classify cell drug response** and **determine the molecular factors of non-genetic resistance to a drug**. These molecular factors represent good candidates to be **targeted during a co-treatment**, in combination with the first drug analyzed with our pipeline.

First, we demonstrate how we combine **image processes and machine learning classification models** to automatically **track cells overtime, and detects important cellular events** like division or death. The output of this first technique are short and sparse fluorescent time-trajectories, that represents the transcriptomic activity in response to the drug, with a unique signal for each cell.

We then introduce our **eDRUGs (early Drug Response UpGraded) classifier**, that combines **mechanistic modeling of apoptosis** (cell death) and **machine learning classification models to predict cell drug response**, within an hour, for a maximum number of cells, using the fluorescent time-trajectories as input. This new method is twice as accurate as our previous prediction method (Péré *et al.* 2022).

Finally, we will also propose a **novel analysis method of sc-RNA-seq data** obtained with

*Speaker

[†]Corresponding author: asma.chalabi@inria.fr

[‡]Corresponding author:

[§]Corresponding author:

Fate-Seq. This method consists in training binary classifiers on the scRNAseq expression data obtained from the pipeline, using a range of models and explainable AI techniques such as DeepLift (Shrikumar *et al.* 2017), in addition to clustering techniques, to obtain attribution scores for each gene. These scores are expected to reveal a reduced gene set, possibly containing only tens of genes, that are predictive of drug resistance.

References

Bian, et al. Coupling live-cell imaging and in situ isolation of the same single cell to profile the transient states of predicted drug-tolerant cells. STAR protocols, 2022.

Meyer, et al. Profiling the non-genetic origins of cancer drug resistance with a single-cell functional genomics approach using predictive cell dynamics. Cell Systems, 2020.

Péré, et al Modeling isogenic cancer cell response upon varying trail stimulations to decipher the kinetic determinants of cell fate decision. Elsevier, 2022.

Shrikumar, et al. Learning important features through propagating activation differences. In International conference on machine learning, 2017.

Keywords: Single, cell, Machine learning, ODE models, Cancer, Time, series classification, Pharmacodynamics, Drug, tolerant persisters.

Learning to act under uncertainty: application to bio-inspired olfactory search

Aurore Loisy ^{*† 1}, Christophe Eloy ¹

¹ Institut de Recherche sur les Phénomènes Hors Equilibre – Aix Marseille Université : UMR7342 / UMR6594 / UMR138, Ecole Centrale de Marseille, Centre National de la Recherche Scientifique : UMR7342 / UMR6594 / UMR138 – France

How to track down a source of odor when there is no track to follow? Many insects, such as mosquitoes, are capable of reaching a source of odor from dozens of meters away. While terrestrial animals can follow odor trails on the ground, such trails do not exist in the air. Indeed, any smooth aerial trail is rapidly destroyed by random wind fluctuations. As a result, insects have to navigate an odor landscape made of few randomly scattered patches of odor. The same challenge is faced by marine crustaceans: to find food or mates, they use olfactory cues to navigate plumes of odor that are significantly altered by oceanic turbulence.

To decipher the search algorithms used in Nature (1), a model POMDP (partially observable Markov decision process) was devised by (2) to mimick the searching conditions under which these animals operate. Far from a "toy" problem, this POMDP encompasses a physical model of odor dispersion in turbulence and its solutions are used for the design of olfactory robots, known as sniffer robots, that track chemicals emitted by explosives and mines.

In this POMDP, the agent must find a stationary target (the source of odor) hidden in a grid world using stochastic partial observations (odor detection events). Good policies must ensure that the source is always found, but to be optimal they must also minimize the duration of the search. This POMDP belongs to a narrower class of problems, called partially observable stochastic shortest path problems, for which a few formal mathematical results exist. Yet, the large size of the problem prevents the use of current POMDP solvers (3). The best solution known to date is a simple, yet remarkably efficient, heuristic based on greedy exploration (2). Here we propose a model-based deep reinforcement learning algorithm able to find near-optimal search strategies (4).

We first reformulate the POMDP as a belief-MDP – a Markov decision process where the states are replaced by belief-states to account for uncertainty – and define its optimal cost function – a mapping from belief states (probability distributions over possible source locations) to their optimal cost (minimum expected time to reach the source when starting from that belief state). Solving the POMDP is then equivalent to computing the optimal cost function for all the (infinitely many) belief-states. We use a deep neural network to approximate the optimal cost function, and train it using a novel algorithm that combines value iteration with model-free deep reinforcement learning techniques (5). The trained agent is found to consistently beat the best heuristic agent in a variety of searching conditions.

While the computational cost of training neural networks will necessarily limit the size of the POMDPs that can be addressed with our method, it is worth mentioning that these results were

*Speaker

†Corresponding author: aurore.loisy@gmail.com

obtained using a CPU-only standard desktop computer. In addition, the method is generic and robust: it does not require any prior human knowledge nor rely on the fine-tuning of hyperparameters. Therefore we expect it to be applicable to a wide range of large POMDPs that are not solvable, even approximately, by standard methods.

From a biological standpoint, we hypothesize that artificial intelligence can help us rationalizing animal behaviour. The learned search strategy, which can be interpreted as a perfect compromise between gathering additional information and exploiting current knowledge, generate trajectories which features can be compared to experimental data. This will be the topic of future work.

- (1) Hein et al., "Natural search algorithms as a bridge between organisms, evolution, and ecology", *Proc. Natl. Acad. Sci. U. S. A.* 113, 9413-9420 (2016).
- (2) Vergassola et al., "Infotaxis as a strategy for searching without gradients", *Nature* 445, 7126 (2007).
- (3) Kurniawati, "Partially observable Markov decision processes and robotics", *Annu. Rev. Control Robot. Auton. Syst.* 5, 253-277 (2022).
- (4) Loisy and Eloy, "Searching for a source without gradients: how good is infotaxis and how to beat it", *Proc. R. Soc. A* 478, 20220118 (2022).
- (5) Mnih et al., "Human-level control through deep reinforcement learning", *Nature* 518, 529–533 (2015).

Keywords: Decision, making under uncertainty, POMDP, deep reinforcement learning, olfactory search, bio, inspired robotics

Learning to act under uncertainty: application to bio-inspired olfactory search

Aurore Loisy ^{*† 1}, Christophe Eloy ¹

¹ Institut de Recherche sur les Phénomènes Hors Equilibre – Aix Marseille Université : UMR7342 / UMR6594 / UMR138, Ecole Centrale de Marseille, Centre National de la Recherche Scientifique : UMR7342 / UMR6594 / UMR138 – France

How to track down a source of odor when there is no track to follow? Many insects, such as mosquitoes, are capable of reaching a source of odor from dozens of meters away. While terrestrial animals can follow odor trails on the ground, such trails do not exist in the air. Indeed, any smooth aerial trail is rapidly destroyed by random wind fluctuations. As a result, insects have to navigate an odor landscape made of few randomly scattered patches of odor. The same challenge is faced by marine crustaceans: to find food or mates, they use olfactory cues to navigate plumes of odor that are significantly altered by oceanic turbulence.

To decipher the search algorithms used in Nature (1), a model POMDP (partially observable Markov decision process) was devised by (2) to mimick the searching conditions under which these animals operate. Far from a "toy" problem, this POMDP encompasses a physical model of odor dispersion in turbulence and its solutions are used for the design of olfactory robots, known as sniffer robots, that track chemicals emitted by explosives and mines.

In this POMDP, the agent must find a stationary target (the source of odor) hidden in a grid world using stochastic partial observations (odor detection events). Good policies must ensure that the source is always found, but to be optimal they must also minimize the duration of the search. This POMDP belongs to a narrower class of problems, called partially observable stochastic shortest path problems, for which a few formal mathematical results exist. Yet, the large size of the problem prevents the use of current POMDP solvers (3). The best solution known to date is a simple, yet remarkably efficient, heuristic based on greedy exploration (2). Here we propose a model-based deep reinforcement learning algorithm able to find near-optimal search strategies (4).

We first reformulate the POMDP as a belief-MDP – a Markov decision process where the states are replaced by belief-states to account for uncertainty – and define its optimal cost function – a mapping from belief states (probability distributions over possible source locations) to their optimal cost (minimum expected time to reach the source when starting from that belief state). Solving the POMDP is then equivalent to computing the optimal cost function for all the (infinitely many) belief-states. We use a deep neural network to approximate the optimal cost function, and train it using a novel algorithm that combines value iteration with model-free deep reinforcement learning techniques (5). The trained agent is found to consistently beat the best heuristic agent in a variety of searching conditions.

While the computational cost of training neural networks will necessarily limit the size of the POMDPs that can be addressed with our method, it is worth mentioning that these results were

*Speaker

†Corresponding author: aurore.loisy@gmail.com

obtained using a CPU-only standard desktop computer. In addition, the method is generic and robust: it does not require any prior human knowledge nor rely on the fine-tuning of hyperparameters. Therefore we expect it to be applicable to a wide range of large POMDPs that are not solvable, even approximately, by standard methods.

From a biological standpoint, we hypothesize that artificial intelligence can help us rationalizing animal behaviour. The learned search strategy, which can be interpreted as a perfect compromise between gathering additional information and exploiting current knowledge, generate trajectories which features can be compared to experimental data. This will be the topic of future work.

- (1) Hein et al., "Natural search algorithms as a bridge between organisms, evolution, and ecology", Proc. Natl. Acad. Sci. U. S. A. 113, 9413-9420 (2016).
- (2) Vergassola et al., "Infotaxis as a strategy for searching without gradients", Nature 445, 7126 (2007).
- (3) Kurniawati, "Partially observable Markov decision processes and robotics", Annu. Rev. Control Robot. Auton. Syst. 5, 253-277 (2022).
- (4) Loisy and Eloy, "Searching for a source without gradients: how good is infotaxis and how to beat it", Proc. R. Soc. A 478, 20220118 (2022).
- (5) Mnih et al., "Human-level control through deep reinforcement learning", Nature 518, 529-533 (2015).

Keywords: Decision, making under uncertainty, POMDP, deep reinforcement learning, olfactory search, bio, inspired robotics

Local Model Reconstruction Attacks in Federated Learning and their Uses

Ilias Driouich * ¹, Chuan Xu ², Giovanni Neglia ³, Frederic Giroire ⁴, Eoin Thomas

¹ Inria-SIC [Sophia Antipolis] – Inria Sophia Antipolis - Méditerranée – 2004, Route des Lucioles BP 93 06902 Sophia Antipolis cedex, France

² Inria Sophia Antipolis – Inria Sophia Antipolis - Méditerranée, Inria Sophia Antipolis - Méditerranée – France

³ MAESTRO (INRIA Sophia Antipolis) – INRIA, Université Montpellier II - Sciences et Techniques du Languedoc – France

⁴ COATI (INRIA Sophia Antipolis / Laboratoire I3S) – CNRS : UMR6070, Université de Nice Sophia-Antipolis, INRIA – France

Federated learning (FL) naturally offers a certain level of privacy, as clients' data is not collected by a third party. Nevertheless, recent works have demonstrated that maintaining the data locally does not provide itself formal privacy guarantees. Especially, an (honest-but-curious) adversary can infer some sensitive client information just by eavesdropping the exchanged messages. In this paper, we initiate the study of local model reconstruction attacks for federated learning, where an honest-but-curious adversary eavesdrops the messages exchanged between a targeted client and the server, and then reconstructs the local/personalized model of the victim. The local model reconstruction attack allows the adversary to trigger other classical attacks in a more effective way, since the local model only depends on the client's data and can leak more private information than the global model learned by the server. Additionally, we propose a novel model-based attribute inference attack in federated learning leveraging the local model reconstruction attack. We provide an analytical lower-bound for this attribute inference attack. Empirical results using real world datasets from airline industry confirm that our local reconstruction attack works well for both regression and classification tasks. Moreover, we benchmark our novel attribute inference attack against the state-of-the-art attacks in federated learning. Our attack results in higher reconstruction accuracy especially when the clients' datasets are heterogeneous (as it is common in federated learning)

Keywords: Federated learning, privacy, Attacks.

*Speaker

Localization of Brain Microstructure Changes Associated with Alzheimer’s Disease using Class Activation Maps

Aymene Mohammed Bouayed ¹, Rachid Deriche ¹, Samuel Deslauriers-Gauthier ^{* 1}

¹ Centre Inria d’Université Côte d’Azur – L’Institut National de Recherche en Informatique et e n Automatique (INRIA) – France

Recently, a general analytical formula to extract all of the 4th order Rotation Invariant Features (RIFs) from diffusion Magnetic Resonance Imaging (dMRI) data has been proposed (1). These invariants have been shown to be linked to the underlying brain microstructure (1) and were recently used to identify Alzheimer Disease (AD) patients (2). While these features indeed contain information that is useful to the identification of AD, the classification is based on the analysis of the whole brain volume and does not pinpoint local changes associated with AD. In this work, we propose to use explainable AI tools, namely Class Activation Maps (CAMs) (3), to localize brain microstructure changes associated with AD.

Class Activation Maps (CAMs) were first introduced in (3). They can be viewed as a heat map which shows for a particular output neuron what region of the input it is most sensitive to conditioned on an input image. CAMs are a weighted sum of the output feature maps of the last convolutional layer by the weights that connect them to the output classification neuron. The resulting CAM is then interpolated to the same dimension of the input image so that it can be overlaid. We extend this method to the case of 3D scans by using 2D convolution on slices of the scans and creating a CAM for each slice by following the same steps as before. This gives us a 3D CAM for the full brain scan.

We calculated 3D CAMs on the same data set used in (2), namely the ADNI - SIEMENS and the ADNI - GE medical data sets, and their best multi-RIF neural network model. The data set ADNI - SIEMENS contains 46 AD and 352 Normal Connectivity (NC) subjects respectively, whereas the ADNI - GE medical contains 191 AD and 419 NC subjects respectively. The obtained CAMs are illustrated in Figure 1. For both AD and NC subjects, the model assigns a high importance to similar voxels of the white matter in the vicinity of the ventricles. This reflects that there are white matter changes captured by the RIFs in the highlighted region that are significant enough to discriminate between the classes. Moreover, from Figure 1 we can see that the regions of interest for the model trained on the ADNI - SIEMENS data set expressed as CAMs encompass the CAMs of the model trained on ADNI - GE medical data set.

In summary, our results indicate that AD may cause microstructure changes in the white matter that can be quantified using rotation invariant features of dMRI and localized using class activation maps.

*Speaker

Keywords: dMRI, Rotation invariant features, Convolutional neural network, Class activation maps

Lucky Luke Maintainer: Detect camera's defects before seeing them

Antoine Amalric ^{*† 1}, Hedi Aloui ^{* ‡ 1}

¹ Videtics – Videtics – 535 route des Lucioles, Bâtiment 3, Les Aqueducs, 06560 Sophia-Antipolis, France

Globally there are more than 770 million cameras already installed which represent a colossal investment made by cities and private structures. That enormous scaling is mostly due to the rise of artificial intelligence to analyze streams that were otherwise limited to human teams monitoring capabilities. However, it left behind a technological debt as it has become more and more overwhelming, if not impossible, to maintain such gigantic infrastructures efficiently. The maintenance cycle can be described as the ordered sequence of camera flaw detection, notification, planning and, finally, manual intervention. Currently, none of those steps are automatic and the cost of maintaining such installations keeps increasing due to human error and time consuming verification tasks. As a direct consequence, the technology we are supposed to leverage to face today's problems and tomorrow's challenges, becomes limited in potential and use. The camera's defects take longer to be detected, therefore crippling the technology used, at best making the investment in the system obsolete and at worst, putting human lives at risk. The figures are clear, too many cameras and not enough operators to process their streams. That's where AI, for the second time, helps us close the gap between the ever growing size of those installations and the need to maintain them. Our work directly tackles what we've identified as the bottleneck of the pipeline maintenance namely the verification step. Our deep learning model, Lucky Luke Maintainer, can adapt to any environment may it be indoor or outdoor. It allows to do, in real time, a complete analysis of all cameras of the installation and can detect defaults that range from camera tampering to extreme illumination settings that all preclude a smooth analytics pipeline. With such a tool, installations could now be maintained more efficiently and unlock an unlimited scalability capability. The new choke point now lies in the planification and repair stage that may one day be solved by robot intervention.

Keywords: intelligent video analytics, maintenance, deep learning, real time, smart cites, camera tampering

*Speaker

†Corresponding author: antoine.amalric@videtics.com

‡Corresponding author: hedi.aloui@videtics.com

NERD (Nematode EffectoR Discovery) : a tool to predict proteins involved in nematodes' plant parasitism.

Djampa Kozłowski ^{*† 1}

¹ Maison de la Modélisation, de la Simulation et des Interactions (MSI) – Université Côte d'Azur (UCA) – 1361 Route des Lucioles, 06560 Valbonne, France, France

Feeding safely over 9 billion people by 2050 with minimal impact on the planet is one of the main challenges humanity is facing (fao.org). Plant parasitic nematodes (PPNs) cause billions of dollars of crop loss per year and thus represent a serious threat to the already fragile food security (Jones et al., 2013). PPNs are microscopic roundworms that secrete effector proteins inside their host plant to manipulate their development, defence systems, metabolism, and physiology. Understanding how effectors operate is thus necessary to drive innovations in crop protection against PPNs. However, because of their diversity, accurate detection of effector proteins in nematode genomes is challenging (Viera et al., 2019). Indeed, most of the known effectors are not conserved across PPNs genera and have specific functions in the host-parasite interaction (Cotton et al., 2013). Therefore, simple criterion such as the presence of a signal peptide for secretion and the absence of a transmembrane region in the protein sequence, also known as 'SP no TM' method, is not sufficient for a reliable effector proteins prediction (Viera et al., 2019). In this context, we developed the NERD (Nematode EffectoR Discovery) tool suite. NERD aims to automatically create, tune, and select supervised Machine Learning (ML) pipelines which best discriminate sets of proteins. Briefly, NERD is subdivided into four modules. NERD-features computes 38 features from protein sequence including physicochemical properties, predicted sub-cellular localization, and presence of specific motifs/domains. Using these data, NERD-build creates pipelines combining 4 possible scaling methods, 3 possible sampling methods and 11 different supervised ML approaches (275 pipelines in total). Pipelines performances are evaluated by cross-validation using user-defined metrics. Pipelines are then ranked according to their performances and the n best are selected for the next step (with n set by the user). NERD-tuning explores hyper-parameter values of each brick of the previously selected pipelines and returns the best parameters combination. NERD-selection evaluates pipeline performances using optimized hyperparameters values. It also combines the pipelines using ensemble approaches to create HARD and SOFT voting classifiers and compare their performances with the individual models. Finally, it elects the m bests individual and ensemble pipelines using a user-defined metric (with m set by the user). The trained models can then be used to predict the label (effector vs. non-effector) of a protein sequence.

We used the NERD tool suite to train PPNs effector protein predictors from a positive dataset composed of known and experimentally validated effectors (330 sequences), and a negative dataset composed of 'widely' conserved proteins among nematodes, more likely to represent housekeeping genes than effectors (1548 sequences). We optimized and selected the models by maximizing the following custom metric: precision * 0.8*recall. This choice was made to minimize the amount of false positive prediction, which is crucial when identifying candidates for

*Speaker

†Corresponding author: djampa.kozlowski@univ-cotedazur.fr

experimental validation. All the generated models outperform the reference ‘SPnoTM’ method . The best model, a HARD voting classifier combining a random forest and a KNN pipeline, shows excellent performances on the validation set with a precision of 98.55% vs 85.91% for ‘SPnoTM’. This model also has good generalization performances since it has a precision of 80.27% vs 68.81% for ‘SPnoTM’ on a set of proteins from PPN species distant from the one used for training. This result suggests despite poor conservation at the sequence level, conserved/similar features exist between distant parasites’ effector proteins, which will be further investigated. We finally used the model to identify effector proteins in the sets of predicted proteins of various PPN species and established a list of potential candidates for experimental validation.

1 fao.org

2 Jones, et al. Top 10 plant-parasitic nematodes in molecular plant pathology: Top 10 plant-parasitic nematodes. *Mol. Plant Pathol.* *14* :946–961, 2013.

3 Vieira, et al. Plant-Parasitic Nematode Effectors - Insights into Their Diversity and New Tools for Their Identification. *Biotic Interactions* 50:37–43, 2019.

4 Cotton, et al. The genome and life-stage specific transcriptomes of *Globodera pallida* elucidate key aspects of plant parasitism by a cyst nematode. *Genome Biol.* 15:R43, 2014.

Keywords: Plant health, Plant parasite nematodes, Machine learning, Supervised model, Effector proteins, Genomic.

On the (Non-)Reliance on Algorithms - A decision-theoretic account

Bernard Sinclair-Desgagné * ¹

¹ Bernard Sinclair-Desgagné – Skema Business School and Université Côte d’Azur (GREDEG) – France

A wealth of empirical evidence shows that people display opposite behaviors when deciding whether to rely on an algorithm, even if it is inexpensive to do so and using the algorithm should enhance their own performance. This talk will bring theoretical support for these conflicting facts. For ‘advisory’ algorithms (like recommender systems), which provide information only, I will look at *the value of algorithmic information* to explain human behavior, refine current policy recommendations, and make further predictions. For ‘performative’ algorithms, which can carry out complex actions like managerial or medical decisions, I will invoke *the value of algorithmic control* to characterize the circumstances where a rational decision maker would be favorable/unfavorable to using such algorithms. The latter analysis will build on an intuitive formal model of *AI as a black box*.

Keywords: AI Economics and Management, Algorithm aversion/appreciation, Value of information, Value of control, Black, box AI

*Speaker

On the Two-fold Role of Logic Constraints in Deep Learning

Gabriele Ciravegna ^{*† 1}, Frédéric Precioso ¹, Marco Gori ¹

¹ Inria – MAASAI team - INRIA Sophia Antipolis – France

In the last few years, Deep Learning (DL) has achieved impressive results in a variety of problems ranging from computer vision to natural language processing. Nonetheless, the excitement around the field may remain disappointed since there are still many open issues. To mitigate some of these problems, we consider the Learning from Constraints framework. In this setting learning is conceived as the problem of finding task functions while respecting the constraints representing the available knowledge. In the Active Learning scenario, First-Order Logic knowledge is converted into constraints and their violation is checked as a guide for sample selection. In the Adversarial Defence scenario, we employ domain knowledge to defend from Adversarial Attacks since it provides a natural way to detect adversarial examples. While some relationships are known properties of the considered environments, DNNs can also autonomously develop new relation patterns. Therefore, we also propose a novel Learning of Constraints formulation which aims at understanding which logic constraints are satisfied by the task functions. This allows explaining DNNs, otherwise commonly considered black-box classifiers. We propose an end-to-end differentiable approach, extracting logic explanations from the same classifier. The method relies on an entropy-based layer which automatically identifies the most relevant concepts. It enables the distillation of concise logic explanations in several safety-critical domains, outperforming state-of-the-art white-box models.

Keywords: Deep Learning, AI, Logic, Domain Knowledge, Hybrid Systems

*Speaker

†Corresponding author: gabriele.ciravegna@inria.fr

Optimal Transport for Graph Representation Learning

Cédric Vincent-Cuaz ^{*† 1,2}, Rémi Flamary ^{3,4}, Corneli Marco ⁵, Titouan Vayer ⁶, Nicolas Courty ⁷

¹ Modèles et algorithmes pour l'intelligence artificielle – Inria Sophia Antipolis - Méditerranée, Scalable and Pervasive softwARe and Knowledge Systems, Université Nice Sophia Antipolis (... - 2019),
Laboratoire Jean Alexandre Dieudonné – France

² Université Côte d'Azur – Université Côte d'Azur (UCA) – France

³ Université Côte d'Azur – Univerité Côte d'Azur – Université Côte d'Azur 06000 Nice, France

⁴ Centre de Mathématiques Appliquées - Ecole Polytechnique – Ecole Polytechnique, Centre National de la Recherche Scientifique : UMR7641 – France

⁵ Université Côte d'Azur, Inria, CNRS, Laboratoire J.A.Dieudonné, Maasai team, Nice, France. –
Université Côte d'Azur (UCA) – France

⁶ Laboratoire de l'Informatique du Parallélisme – Centre National de la Recherche Scientifique :
UMR5668 / URA1398, Université de Lyon, Institut National de Recherche en Informatique et en
Automatique, Université Claude Bernard Lyon 1, École Normale Supérieure - Lyon – France

⁷ UBS/IRISA – Université de Bretagne Sud, IRISA/OBELI, Université de Bretagne Sud – France

Graph representation learning (GRL) for graph-level tasks aims to design a vector encoding information from both node features and graph topology. To this end, we propose to leverage Optimal Transport theory to add dictionary learning (DL) to the arsenal of unsupervised GRL, and to improve the topology sensitivity of graph neural networks (GNN) dedicated to supervised GRL.

First, DL is a key tool for representation learning, that explains the data as linear combination of few basic elements. Yet, this analysis is not amenable in the context of graph learning, as graphs usually belong to different metric spaces. We fill this gap by proposing a new Graph DL approach (1), which uses as data fitting term, the Fused Gromov-Wasserstein (FGW), (3)) distance which encodes simultaneously feature and structure dissimilarities by solving a soft graph-matching problem. In our work, graphs are jointly encoded through their nodes' pairwise relations and their node features, and modeled as convex combination of graph atoms, i.e. dictionary elements, estimated thanks to a stochastic algorithm operating on a dataset of unregistered graphs. GDL is completed by a novel upper bound that can be used as a fast approximation of FGW in the embedding space. We show the relevance of our approach for unsupervised GRL of graph datasets and for online graph subspace estimation and tracking.

Second, we focus on GNN which rely on two main components: node features embedding through message passing, and aggregation with a specialized form of pooling. The structural information is solely taken into account implicitly in these two steps. We propose in this work ((2)) a novel point of view, which places FGW distances to some learnable graph templates at the core of the graph representation. We postulate that the vector of FGW distances to a set of template graphs has a strong discriminative power, which is then fed to a non-linear classifier for final predictions. Distance embedding can be seen as a new layer, and can leverage on existing message passing techniques to promote sensible feature representations. The optimal set of template graphs is also learnt in an end-to-end fashion by differentiating through this layer. We empirically validate

*Speaker

†Corresponding author: cedric.vincent-cuaz@inria.fr

our claim on several graph classification datasets, known to stress the expressiveness and generalization abilities of GNN (4), and show that our method surpasses state-of-the-art approaches.

(1) Vincent-Cuaz, C., Vayer, T., Flamary, R., Corneli, M. & Courty, N. Online graph dictionary learning. ICML 2021.

(2) Vincent-Cuaz, C., Flamary, R., Corneli, M., Vayer, T. & Courty, N. Template based Graph Neural Network. NeurIPS 2022.

(3) Vayer, T., Chapel, L., Flamary, R., Tavenard, R., & Courty, N. Optimal Transport for structured data with ICML 2019.

(4) Xu, L., Hu, W, Leskovec, J., Jegelka, S. How Powerful are Graph Neural Networks?. ICLR 2019.

Keywords: Graph Representation Learning, Optimal Transport, Dictionary Learning, Graph Neural Networks

Predicting odorant-receptor activation with protein language and graph neural networks

Matej Hladiš ^{*† 1}, Maxence Lalis ¹, Sébastien Fiorucci ¹, Jérémie Topin^{‡ 1}

¹ Université Côte d’Azur – Université Côte d’Azur (UCA) – Parc Valrose, 28, avenue Valrose 06108 Nice Cedex 2, France

Odour perception in mammals is triggered by interactions between volatile organic compounds and a subset of hundreds of proteins called olfactory receptors (ORs). Molecules activate these receptors in a complex combinatorial coding allowing mammals to discriminate a vast number of chemical stimuli. One odorant can activate many different ORs, and an OR can accept several classes of molecules, called ligands. To date, predicting molecule-induced activation for ORs is highly challenging since 43% of ORs have no identified ligands and only 205 out of the 385 human receptors have a known active compound. To tackle the OR-activation prediction, we combine (CLS) token embedding from protBERT1 protein language model with a molecular graph and propose a tailored GNN architecture incorporating inductive biases from the protein-molecule binding. We abstract the biological process of protein-molecule activation as the injection of a molecule into a protein specific environment. Moreover, only a limited amount of curated data linking a molecule to a set of ORs is available. To fill this gap, we gathered and curated a new dataset of 46 650 OR-molecule pairs from the literature, tripling the size of the currently largest dataset of this kind. On a high confidence dose-response data, we show that our approach outperforms previous work by more than 30%. The model correctly identifies 70% of ligands with 69% precision in an i.i.d. case and keeps the precision above 65% when tested on entirely unseen molecules and ORs. Finally, we analyse the biological relevance of the model predictions and its agreement with the theory of combinatorial coding in olfaction. By linking olfactory perception of molecules with their OR activations we confirm that odorants are recognized by a unique subset of ORs and that these subsets are specific to an olfactory characteristic. This is in full agreement with previous laboratory observations.

Keywords: Olfaction, Protein molecule interaction, Olfactory receptors, Graph neural networks, Protein language modelling

*Speaker

†Corresponding author: matej.hladis@univ-cotedazur.fr

‡Corresponding author: jeremie.topin@univ-cotedazur.fr

Prospects of AI-augmented medical treatment acceptance

Zakaria Babutsidze *^{1,2}, Ewa Wlodarczyk-Boudenot, Ibrahim Abdel Malak

¹ Observatoire français des conjonctures économiques – Sciences Po – France

² SKEMA Business School – SKEMA Business School-UCA – France

Artificial Intelligence (AI) can today be found everywhere from spam filters, fraud prevention, weather forecasting, face recognition to self-driving cars and many more. AI has particularly high potential for creating value in healthcare. It has already been successfully employed in diagnosing various diseases at early stages of development (Gulshan et al. 2016) and predicting heart attacks (Hutson 2017) that could potentially save millions of lives. Recent scrutiny of possibility and desirability of automation of healthcare tasks finds surprisingly large overlap between the two categories, which gives ground for optimism in terms of incorporating AI technologies in healthcare systems (Fruehwirt and Duckworth 2021).

Bulk of medical AI applications happen at the background, without the involvement of a patient, or a proxy decision-maker (someone with medical power of attorney). However, many medical applications are patient-facing. For such patient-facing medical AI applications it is important that the patient, or a decision-maker on behalf of the patient, is willing to accept AI's part in the procedure, as patient's consent is usually necessary in the process. This brings up questions of acceptance of AI technologies by individuals who do not necessarily understand advantages and risks presented by these new technologies, but are being put in spot to make a consequential decision. Even though technology acceptance studies have been around for long period (e.g., Venkatesh et al. 2003), given specificities of AI and speed of its injection into medical system, re-examination of established concepts in this specific setting is desirable. It is important to study whether the introduction of AI elements in established medical procedures can undermine patient's trust toward these processes. We examine this very question in medical setting and study drivers of AI acceptance in high-stake medical situations.

We conduct two experimental survey studies. Both studies put forward a hypothetical situation where subject's family member requires surgical intervention and where the experimental subject has the medical power of attorney. They need to decide whether to accept the surgical intervention. In both studies we randomly split subjects in two groups – control and treatment. While in the control situation the whole medical procedure is performed by human doctors, in the treatment situation we have AI technology participating. The difference between the two studies is the phase in which AI technology is inserted in the medical procedure. In study 1 AI technology is used in non-invasive phase (i.e., predicting the likelihood of success of the surgery), in study 2 AI is inserted in the invasive phase of the procedure (i.e., AI guided robot is physically assisting human doctors during the surgery). We find that while total effect of AI involvement in study 1 is negative, the same effect in study 2 is, in fact, positive. We also find that an important driver of AI-augmented healthcare acceptance is subject's trust in artificial intelligence. Trust in AI, in its turn, is affected by AI knowledge that subject possesses. References Fruehwirt, W., Duckworth, P. (2021) Towards better healthcare: What could and should be automated? *Technological Forecasting and Social Change*, 172, 120967.

*Speaker

Gulshan, V., Peng, L., Coram, M., et al. (2016) Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *Journal of the American Medical Association*, 316(22), 2402–2410.

Hutson, M. (2017) Self-taught artificial intelligence beats doctors at predicting heart attacks. *Science*, 14(04).

Venkatesh, V., Morris, M. G., Davis, G. B., Davis, F. D. (2003) User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478.

Quantum Decision Modeling for the Travel Industry

Federico Tiblias ^{*† 1,2,3}, Alix Lhéritier ^{*}

¹, Nicolas Bondoux ¹, Rodrigo Acuna Agost ¹, Marios Kountouris ²

¹ Amadeus – Amadeus SAS – France

² EURECOM – EURECOM – France

³ Politecnico di Torino – Italy

Travelers face many choice situations whose outcomes are very hard to predict by service providers such as airlines or travel agencies. In particular, compared to traditional applications in retail, travelers face a very dynamic set of alternatives: prices change constantly, different flight combinations can serve a single destination and their availability change as tickets are sold.

Traditional choice models are based on different assumptions about human decision making. For example, the widely used Multinomial Logit (MNL) model satisfies Luce’s axiom (1) also known as independence of irrelevant alternatives, which states that the probability of selecting one alternative over another from a set of many alternatives is not affected by the presence or absence of other alternatives in the set.

Pairwise Choice Markov Chains (2) and its feature-based extension PCMC-Net (3) have recently been introduced to overcome limitations of choice models based on traditional axioms unable to express empirical observations from modern behavior economics like context effects occurring when the preference between two options is altered by adding a third alternative.

However, the suitability of Classical Probability Theory has been challenged by empirical experiments like in (4), where participants ranked a conjunction of events X and Y as being more likely than the single event X, violating the axioms of classical probability.

Quantum Cognitive models (see, e.g., (5)) use the formalism of Quantum Probability Theory as a set of principles for inference, encompassing fallacies in decision making such as the conjunction fallacy of the aforementioned experiment, question order effects and others.

We investigate the suitability of decision models based on Quantum Probability Theory to represent human behaviors observed in the Travel Industry.

References

(1) R. D. Luce, Individual Choice Behavior: A Theoretical analysis, New York, NY, USA: Wiley, 1959.

(2) S. Ragain and J. Ugander, "Pairwise choice Markov chains," *Advances in neural infor-*

*Speaker

†Corresponding author: federico.tiblias@amadeus.com

mation processing systems, vol. 29, 2016.

(3) A. Lhéritier, "PCMC-Net: Feature-based Pairwise Choice Markov Chains," in *International Conference on Learning Representations*, 2020.

(4) A. Tversky and D. Kahneman, "Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment.," *Psychological review*, vol. 90, p. 293, 1983.

(5) E. M. Pothos and J. R. Busemeyer, "Quantum cognition," *Annual review of psychology*, vol. 73, p. 749–778, 2022.

Keywords: Choice Modeling, Decision Modeling, Cognition, Quantum Probability Theory

Quantum-based Sentiment Analysis for the Travel Industry

Massimiliano Pronesti ^{*† 1,2,3}, Alix Lhéritier ^{*}

¹, Nicolas Bondoux ¹, Mourad Boudia ¹, Marios Kountouris ²

¹ Amadeus – Amadeus SAS – France

² EURECOM – EURECOM – France

³ Politecnico di Torino – Italy

Analyzing and understanding customer feedback on their travel experiences, like flights or stays at hotels, is key for the Travel Industry. Feedback can be given by a score but also by a free text describing many aspects of the experience. Understanding the sentiment behind each sentence of a feedback is key to understand the nuances of the experience and provide better and customized offers for future travelers.

Recently, Natural Language Processing (NLP) has made giant leaps in understanding language, by means of huge deep learning models like Transformers (1). These models are based on the distributional theory—a word’s meaning is defined by the context in which it appears—and learn a representation of words in a vector space. These architectures are not tied to a specific language or grammar but require large amounts of data to be trained.

Another line of research, sparked by linguists such as N. Chomsky and J. Lambek, is based on the compositional theory of grammatical types and provides an algebraic treatment of grammar (2). This allows, for example, to mathematically prove that a given sentence is coherent and makes sense.

The categorical compositional distributional (DisCoCat) model (3) combines the distributional theory with the compositional one. In a nutshell, words are represented in tensor spaces and are combined by means of linear maps, allowing to compute the meaning of a whole sentence as a tensor value.

On classical computers, these operations are exponential in the number of words. However, it was shown (4) that thanks to the formalism of Categorical Quantum Mechanics, these operations naturally map to quantum operations as those performed by currently available Noisy Intermediate-Scale Quantum (NISQ) devices, making their complexity only linear in the number of words (see e.g. (5)).

Assuming the grammatical structure is given, the learning problem reduces to learning a proper representation of the individual words such that the final representation of the sentences is useful for the given task—in our case, sentiment classification. By leveraging the grammar, this approach has the potential of learning with a significantly reduced amount of data.

In this work, we explore the suitability of such Quantum Natural Language Processing models on

*Speaker

†Corresponding author: massimiliano.pronesti@amadeus.com

two sentiment analysis datasets. The first one consists of 14K tweets on six US Airlines, which lacks a good English grammar and therefore limits the benefits of grammar-based models. The second one is made of hotel reviews that enjoy a good grammar. We show that grammar-based models outperform Recurrent Neural Network based architectures on sentence understanding tasks and potentially compete with pre-trained state-of-the-art NLP models, despite the comparison not being fair at this stage of advancement of the quantum technology.

References

- (1) A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- (2) J. Lambek, "The mathematics of sentence structure," *The American Mathematical Monthly*, vol. 65, p. 154–170, 1958.
- (3) B. Coecke, M. Sadrzadeh and S. J. Clark, "Mathematical foundations for a compositional distributional model of meaning," *Linguistic Analysis*, vol. 36, p. 345–384, 2010.
- (4) S. Abramsky and B. Coecke, "A categorical semantics of quantum protocols," in *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.*, 2004.
- (5) B. Coecke, G. de Felice, K. Meichanetzidis and A. Toumi, "Foundations for near-term quantum natural language processing," *arXiv preprint arXiv:2012.03755*, 2020.

Keywords: Natural Language Processing, Quantum Computing, Category Theory

Real-Time High-Resolution Traffic Monitoring with Distributed Acoustic Sensing and AI

Yacine Khacef * ¹, Martijn Van Den Ende , Cédric Richard , André Ferrari , Anthony Sladen

¹ Yacine Khacef – Observatoire de la Cote d’Azur, 3iA Côte d’Azur, Université Côte d’Azur (UCA) – France

Traffic management is one of the main challenges that modern cities are facing. Real-time traffic data are a first-order requirement for a range of critical tasks, like traffic jam detection, travel time estimation and the improvement of infrastructures for safer and more environmentally-friendly transportation. Existing techniques for traffic monitoring are mainly based on roadside cameras and road-embedded loop detectors, which have several limitations and drawbacks such as high deployment and maintenance costs, and low spatial resolution. As a low-cost alternative, we are currently developing a real-time traffic monitoring solution based on Distributed Acoustic Sensing (DAS). DAS turns existing fibre-optic (telecommunication) cables into an array of vibration sensors with metric spatial resolution and a range of over one hundred kilometer. Given that telecommunication fibres are often deployed along existing traffic infrastructures, DAS holds great potential for recording vehicles traffic flows with unprecedented spatial resolution. We have devised Deep-Learning based algorithms to reliably detect and extract traffic information from DAS data, which are designed to operate faster than real-time. The output of the algorithms (i.e., the traffic information) is visually presented to the user via a dashboard that includes various data analysis tools, accessible directly from a web browser. Underlying this dashboard is a dedicated time-series database where the DAS data and derived products (vehicle counts, traffic speeds, among others) are stored as they come out of the processing pipeline. We showcase a prototype that focuses on the city of Nice, highlighting the potential of DAS for low-cost, high-resolution, and real-time traffic monitoring in urban environments.

Keywords: Traffic Monitoring, Deep Learning, Smart Cities, Distributed Acoustic Sensing

*Speaker

Refrigerant leak detection in industrial refrigeration systems with vapor compressors

Amal Mtibaa * ^{1,2}, Valentina Sessa ¹, Gilles Guerassimoff ¹, Stéphane Alajarin ²

¹ Centre de Mathématiques Appliquées (CMA) – MINES ParisTech - École nationale supérieure des mines de Paris – Sophia Antipolis, France

² Clauger 3E – Clauger – Brignais, Lyon, France

Over the years, the industrial refrigeration market has continued to amply evolve following its use in various applications such as food and beverage processing, cold storage, and others. An efficient operation of a refrigeration system can improve its environmental impact by decreasing energy consumption and greenhouse gas emissions. For this reason, fault detection and diagnosis (FDD) (1) has accrued the attention of many researchers in this field. However, the number of those dealing with the industrial use case is small. Loss of refrigerant is one of the most crucial faults that can affect the efficiency and safety of an industrial refrigeration system. Leakages can remain undetected until sufficient refrigerant has been lost resulting in severe environmental and financial damages.

This work presents a refrigerant leak detection solution for industrial vapor compression systems. A data-driven approach is proposed where a system collects temperature and pressure sensors' data as well as ambient temperature and general electrical power. The collected data are used to feed a Machine Learning model which predicts the fault-free liquid reservoir level. Afterward, the predicted level is compared with the actual one to verify whether a leak is taking place, (2), (3).

The main contribution of our work is presented in what follows:

- Differently from what was proposed in the literature, we consider the reservoir's liquid level as our target. This choice was suggested by our numerical experiments and technical experience. Indeed, we observed that the reservoir's liquid level provides a good indication for detecting gradual loss at an early stage.
- Our leak detection system is able to detect gradual leaks regardless of the physical installation aspects (e.g., number of evaporators, type of refrigeration fluid).
- We work on actual industrial installations data, allowing us to present a proof of concept, as industrial data are noisier than laboratory data and usually not supervised.

The steps performed by the proposed refrigerant leak detection system are described in the following. First, the collected sensors' data go through a preprocessing phase for cleaning and anomaly detection. Then, a regression model is implemented for the prediction of the fault-free liquid reservoir level. In this work, we compare the performance of tree-based regression methods, such as Random Forest, Extreme Forest, Gradient Boosting, XGboost, and Light GBM,

*Speaker

and Facebook’s forecasting model Prophet. The comparison was implemented by using common time series error metrics as well as dedicated ones. We performed our numerical experiments on four actual industrial cases, characterized by different refrigeration installations and with data collected in different seasons.

Finally, we present an example of leak detection and highlight the different challenges that come with it.

References

- (1) Jianli Chen et al. "A review of computing-based automated fault detection and diagnosis of heating, ventilation and air conditioning systems". In: *Renewable and Sustainable Energy Reviews* 161 (2022), pp. 112–395.
- (2) S. A. Tassou and I. N. Grace. "Fault diagnosis and refrigerant leak detection in vapour compression refrigeration systems". en. In: *International Journal of Refrigeration* 28.5 (2005), pp. 680–688.
- (3) Dinh Anh Tuan Tran et al. "An enhanced chiller FDD strategy based on the combination of the LSSVRDE model and EWMA control charts". In: *International journal of refrigeration* 72 (2016), pp. 81–96.

Keywords: Leak detection, industrial refrigeration systems, fault detection and diagnosis, machine learning

Shallow convolutional neural network with rank-1 Fourier domain weights for brain signal classification

Sara Sedlar ^{*† 1}, Samuel Deslauriers-Gauthier ², Rachid Deriche ²,
Theodore Papadopoulo ^{‡ 2}

¹ l'Equipe-Projet ATHENA – CRISAM - Inria Sophia Antipolis - Méditerranée, Université Côte d'Azur (UCA) – 2004 Rte des Lucioles, 06902 Valbonne, France

² l'Equipe-Projet ATHENA – CRISAM - Inria Sophia Antipolis - Méditerranée, Université Côte d'Azur (UCA) – 2004 Rte des Lucioles, 06902 Valbonne, France

A brain activity can be described by a cortical region where it is located and its temporal course. Magnetic field strength and electric potential recorded at scalp (or in its proximity) by magneto- and electro-encephalography (M/EEG) devices are direct measures of these activities. They can be explained by Maxwell's equations with quasi-static approximations (1). Consequently, we can assume that cortical brain activities spread instantaneously and linearly over the measuring sensors, thus a multivariate M/EEG signal can be represented as a sum of rank-1 multivariate signals corresponding to individual activities and noise. A rank-1 multivariate signal corresponds to an outer product of a brain activity temporal course and its spatial map to the sensors. As the human head can be approximated by a sphere, spatial maps can be seen as spherical signals and thus be represented in terms of spherical harmonic basis. This representation is less sensitive to the spatial distribution of the sensors which varies between subjects and sessions. Recent studies have shown that brain waveforms are often of a transient and recurrent nature (2). This is also the case in brain-computer interface (BCI) if the waveforms are evoked by external sensory stimuli. Assuming transience and recurrence in the temporal courses, they can be modeled as convolution of Dirac impulses and characteristic temporal patterns.

Given the rank-1 assumption, spherical head approximation, and transience and recurrence in the temporal courses, we propose a convolutional neural network (CNN) with rank-1 spectral domain trainable kernels for M/EEG signal classification. As convolutional layers in CNNs perform correlation, the objective of the model is to learn spatial and temporal kernels which resemble the characteristic temporal patterns and spatial maps. To constrain temporal kernels to extract features within relevant spectral ranges, they are represented in terms of discrete cosine basis.

As illustrated in Figure 1, the model is composed of a single layer feature extraction module realized with rank-1 trainable spectral domain kernels, a feature selection module realized with a simple thresholding and max-pooling operator, and a feature classification module composed of one fully connected layer. The model is compared with the state-of-the-art CNN models on the BCI problem of mental workload classification from EEG signals (3) and on the motor-task MEG signal (4) classification problem. We have shown that our model can achieve state-of-the-art performance with a significantly lower number of parameters. Given this and its speed both during training and test phase, it is well suited for portable devices in BCI.

(1) Jukka Sarvas. Basic mathematical and electromagnetic concepts of the biomagnetic inverse

*Speaker

†Corresponding author: sara.sedlar@inria.fr

‡Corresponding author: theodore.papadopoulo@inria.fr

problem. *Physics in Medicine
Biology*, vol. 32, no. 1, page 11, 1987.

(2) Freek van Ede, Andrew J Quinn, Mark W Woolrich and Anna C Nobre. Neural oscillations: sustained rhythms or transient burst-events? *Trends in Neurosciences*, vol. 41, no. 7, pages 415–417, 2018.

(3) Marcel F Hinss, Bertille Somon, Frédéric Dehais and Raphaëlle N Roy. Open EEG Datasets for Passive Brain-Computer Interface Applications: Lacks and Perspectives. In *2021 10th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 686–689. IEEE, 2021.

(4) David C Van Essen, Kamil Ugurbil, Edward Auerbach, Deanna Barch, Timothy EJ Behrens, Richard Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtisset al. The Human Connectome Project: a data acquisition perspective. *Neuroimage*, vol. 62, no. 4, pages 2222–2231, 2012.

Keywords: M/EEG classification, rank 1 CNN, brain computer interfaces

Speeding up convergence in decentralized multi-agent learning through smart neighbor picking

Marina Costantini * ¹, Thrasyvoulos Spyropoulos ^{1,2}

¹ EURECOM – EURECOM – Campus SophiaTech, 450 Route des Chappes, CS 50193 - 06904 Biot
Sophia Antipolis cedex, France

² Technical University of Crete [Chania] – Plateia Agiou Titou 73132 Chania - Crete, Greece

As the excitement for the breakthroughs of machine learning (ML) grows, so do concerns for how these data-based approaches threaten privacy. This has led to the emergence of a new ML paradigm: decentralized learning. In contrast to traditional ML, where data must be gathered at a central location, in the decentralized setting many interconnected agents cooperatively train an ML model exploiting the data of all nodes, but keeping the data private. This paradigm has allowed e.g. to perform clinical diagnosis using the data of multiple hospitals without compromising confidentiality (1).

In this work, we propose a new asynchronous decentralized algorithm where agents can activate anytime and choose a neighbor to make an update together (2). Methods with such minimal coordination avoid incurring extra synchronization costs that may slow down convergence (see Fig. 1). However, most asynchronous approaches assume that the neighbor selection is done at random, overlooking the possibility of taking into account the optimization landscape at the time of making the choice. Instead, here we propose to exploit the degree of freedom of the neighbor selection to speed up convergence.

In particular, we formulate the decentralized learning problem as a constrained optimization problem, and show that by solving this problem through its dual we move the control variables from the network nodes to the edges, thus allowing for an asynchronous algorithm that updates one edge variable at a time. We then remark that this algorithm can be cast into the (centralized) Coordinate Descent method (CD) (3), with one crucial difference: in CD any coordinate may be updated, but in our setting only the variables of the edges connected to the activated node are available.

We thus define the new class of Set-wise CD algorithms, and propose two neighbor selection rules: random uniform (SU-CD) and Gauss-Southwell (SGS-CD), which selects the neighbor whose gradient has the largest magnitude (4). Since previous analyses in CD do not apply, we develop a new methodology based on norms uniquely defined for each algorithm considered that allows us to obtain linear convergence rates for smooth and strongly convex functions. Furthermore, we prove that the speedup of SGS-CD with respect to SU-CD can match the maximum degree in the network, and confirm in simulations that this speedup increases abruptly with network connectivity. Moreover, our results apply also to the distributed parallel setting.

Keywords: decentralized learning, optimization algorithms, multiagent learning, asynchronous al-

*Speaker

gorithms, coordinate descent

The Zero-Inflated Poisson Dynamic Latent Block Model with an Application to Pharmacovigilance Data

Giulia Marchello * ¹, Corneli Marco ², Charles Bouveyron ^{3,4}

¹ Université Côte d'Azur, Inria, CNRS, Laboratoire J.A.Dieudonné, Maasai team, Nice, France – Université Cote d'Azur, Institut National de Recherche en Informatique et en Automatique, Centre Sophia-Antipolis Méditerranée – France

² Université Côte d'Azur, Inria, CNRS, Laboratoire J.A.Dieudonné, Maasai team, Nice, France. – Université Côte d'Azur (UCA) – France

³ Laboratoire J.-A. Dieudonne, UMR CNRS 7531, Université Cote d'Azur – CNRS : UMR7531 – France

⁴ Equipe Maasai, Inria Sophia Antipolis – Institut National de Recherche en Informatique et en Automatique – France

The simultaneous clustering of observations and features of data sets (known as co-clustering) has recently emerged as a central machine learning application to summarize massive data sets. However, most existing models focus on continuous and dense data in stationary scenarios, where cluster assignments do not evolve over time. This work introduces a novel latent block model for the dynamic co-clustering of data matrices with high sparsity. To properly model this type of data, we assume that the observations follow a time and block dependent mixture of zero-inflated distributions, thus combining stochastic processes with the time-varying sparsity modeling. To detect abrupt changes in the dynamics of both cluster memberships and data sparsity, the mixing and sparsity proportions are modeled through systems of ordinary differential equations. The inference relies on an original variational procedure whose maximization step trains fully connected neural networks in order to solve the dynamical systems. An application to adverse drug reaction in pharmacovigilance is also proposed, where the proposed model recognize clusters in a meaningful way by identifying safety events that were consistent with retrospective knowledge. Hence, our aim is to propose this dynamic co-clustering method as a tool for automatic safety signal detection, to support medical authorities.

Keywords: Co clustering, zero inflated distributions, dynamic systems, variational inference, pharmacovigilance.

*Speaker

The graph embedded topic model

Dingge Liang ^{*† 1}, Marco Corneli ², Charles Bouveyron ³, Pierre Latouche ⁴

¹ Université Côte d'Azur – L'Institut National de Recherche en Informatique et en Automatique (INRIA) – Parc Valrose, 28, avenue Valrose 06108 Nice Cedex 2, France

² Université Côte d'Azur, Inria, Maison de la Modélisation des Simulations et des Interactions (MSI), Maasai team, Nice, France. – Université Côte d'Azur, Institut National de Recherche en Informatique et en Automatique, Centre Sophia-Antipolis Méditerranée – France

³ Université Côte d'Azur, Inria, CNRS, Laboratoire J.A.Dieudonné, Maasai team, Nice, France. – Université Côte d'Azur (UCA) – France

⁴ Mathématiques Appliquées Paris 5 – Institut National des Sciences Mathématiques et de leurs Interactions : UMR8145, Centre National de la Recherche Scientifique : UMR8145, Université de Paris : UMR8145 – France

Most of current graph neural networks (GNNs) developed for the prevalent text-rich networks typically treat texts as node attributes. This kind of approaches unavoidably results in the loss of important semantic structures and restricts the representational power of GNNs. In this chapter, we introduce a document-similarity based graph convolutional network (DS-GCN) encoder to combine graph convolutional networks and embedded topic models for text-rich network representation. Then, a latent-position based decoder is used to reconstruct the graph while maintaining its topology. Similarly, the document matrix is rebuilt using a decoder that takes both topic and word embeddings into account. By including a cluster membership variable, we thus develop an end-to-end clustering technique relying on a new deep probabilistic model called graph embedded topic model (GETM). The effectiveness of GETM in fusing the graph topology structure and the topic embeddings is illustrated with numerical experiments performed on three simulated scenarios, which additionally emphasize the performance in terms of node clustering of the proposed methodology.

Keywords: Deep generative models, Graph convolutional networks, Topic modeling, Clustering

*Speaker

†Corresponding author: dingge.liang@inria.fr

Unsupervised Echocardiography Registration through Patch-based MLPs and Transformers

Yingyu Yang ^{*† 1,2}, Zihao Wang ^{1,2}, Herve Delingette ^{1,2}, Maxime Sermesant ^{1,2}

¹ Inria – L’Institut National de Recherche en Informatique et e n Automatique (INRIA) – France

² UCA – University Côte d’Azur – France

Image registration is an essential but challenging task in medical image computing, especially for echocardiography, where the anatomical structures are relatively noisy compared to other imaging modalities. Traditional (non-learning) registration approaches rely on the iterative optimization of a similarity metric which is usually costly in time complexity. In recent years, convolutional neural network (CNN) based image registration methods have shown good effectiveness. In the meantime, recent studies show that the attention-based model (e.g., Transformer) can bring superior performance in pattern recognition tasks. In contrast, whether the superior performance of the Transformer comes from the long-winded architecture or is attributed to the use of patches for dividing the inputs is unclear yet. This work introduces three patch-based frameworks for image registration using MLPs and transformers. We provide experiments on 2D-echocardiography registration to answer the former question partially and provide a benchmark solution. Our results on a large public 2D-echocardiography dataset show that the patch-based MLP/Transformer model can be effectively used for unsupervised echocardiography registration. They demonstrate comparable and even better registration performance than a popular CNN registration model. In particular, patch-based models better preserve volume changes in terms of Jacobian determinants, thus generating robust registration fields with less unrealistic deformation. Our results demonstrate that patch-based learning methods, whether with attention or not, can perform high-performance unsupervised registration tasks with adequate time and space complexity.

Keywords: Unsupervised Registration, MLP, Transformer, Echocardiography

*Speaker

†Corresponding author: yingyu.yang@inria.fr

” Why do I have chickenpox? ”

Template-based explanations for correct and incorrect diagnosis.

Benjamin Molinet ^{*†}, Serena Villata ¹, Elena Cabrio, Maya Medjad,
Santiago Marro

¹ Université Côte d’Azur, Inria, CNRS, I3S – Université Côte d’Azur, Inria, CNRS, I3S – France

In current Artificial Intelligence (AI) applications, it is crucial to provide methods to explain the decisions of these tools in a way that can be understood by humans. The automatic generation of explanations based on arguments is important to support automatic decision making in critical domains. This work focuses on the medical domain and aims at generating explanations for a given diagnosis. Given a clinical case proposed to students in medicine, we automatically generate explanations about why a certain diagnosis is correct and why the alternative answers proposed to the students are not. We base our analysis on the 1843 symptoms we retrieved from the Human Phenotype Ontology (HPO) for each diagnosis in our dataset. To enrich this dataset with clinical knowledge, we also annotated 314 clinical cases with labels defined by UMLS (the Unified Medical Language System). We thus provide a complete pipeline that takes as input a clinical case as well as the correct and incorrect answers, and generates natural language explanations. The pipeline includes a symptom detection system for the analysed clinical case and a matching module between the representations in the clinical case and in the HPO ontology. We obtain good results in both symptom detection and alignment, outperforming competitive baselines.

Keywords: Explainable AI, Natural Language Explanations, Symptoms detection

*Speaker

†Corresponding author: benjamin.molinet@inria.fr

Author Index

- Abdel Malak, Ibrahim, 63
Acuna Agost, Rodrigo, 65
Alajarin, Stéphane, 70
Aloui, Hedi, 55
Amalric, Antoine, 55
- Babutsidze, Zakaria, 63
Baldovini, Nicolas, 21
Betti, Alessandro, 40
Bondoux, Nicolas, 65, 67
Bouayed, Aymene Mohammed, 53
Boudia, Mourad, 67
Bouveyron, Charles, 44, 76, 77
Broekaert, Jan, 7
- Cabrio, Elena, 17, 33, 79
Castellaneta, Francesco, 42
Cazeau, Serge, 2
Chalabi, Asma, 46
Chaves, Madalena, 46
Chevrier, Mathieu, 41
Chimenea-Toscana, Ángel, 22
Ciravegna, Gabriele, 59
Cirillo, Bruno, 42
Corcos, Ludovic, 38
Corgnet, Brice, 41
Corneli, Marco, 77
Costantini, Marina, 74
Courty, Nicolas, 60
- Defoin-Platel, Michael, 28
delingette, herve, 78
Deriche, Rachid, 53, 72
Deslauriers-Gauthier, Samuel, 38, 53, 72
Desrués, Gaëtan, 2
DRIOUICH, Ilias, 52
- Eloy, Christophe, 48, 50
- Faggi, Lapo, 40
Faticanti, Francescomaria, 36
Ferrari, André, 69
Fiorucci, Sébastien, 62
Flamary, Rémi, 60
Fraboni, Yann, 5
- García-Díaz, Lutgardo, 22
Giroire, Frederic, 52
Goffredo, Pierpaolo, 33
Gori, Marco, 40, 59
Guerassimoff, Gilles, 70
Guerci, Eric, 41
Guyard, Frederic, 19, 24
- Haddadan, Shohreh, 33
Hafiz, Faizal, 7
Hladiš, Matej, 62
Humbert, Oliver, 44
HUMBERT, Olivier, 30
- Kameni, Laetitia, 34
Kashtanova, Victoriya, 26
Khacef, Yacine, 69
Kountouris, Marios, 65, 67
kozlowski, djampa, 56
Krafft, Jackie, 15
Kögler, Christoph, 4
- La Torre, Davide, 7
LABORY, Justine, 11
Lalis, Maxence, 21, 62
Latouche, Pierre, 77
Leonardi, Emilio, 36
Lhéritier, Alix, 65, 67
LIANG, Dingge, 77
liu, ziming, 9, 10
Loisy, Aurore, 48, 50
- malis, ezio, 9, 10
Marchello, Giulia, 76
Marco, Corneli, 60, 76
MARFOQ, Othmane, 34, 36
Marro, Santiago, 17, 79
martinet, philippe, 9, 10
Martínez-Martínez, Ricard, 22
Marullo, Simone, 40
Matej, Hladis, 21
Mattei, Pierre-Alexandre, 30, 44
Medjad, Maya, 79
Melacci, Stefano, 40
Meloni, Enrico, 40

Merlet, Jean-Pierre, 13
 Miralles, Hugo, 24
 Miró-Amarante, Gloria, 22
 Mocerri, Pamela, 32
 Molinet, Benjamin, 79
 Mtibaa, Amal, 70

Nanty, Simon, 28
 Neglia, Giovanni, 34, 36, 52

Oprescu, Andreea, 22
 Oyarzun, Diego, 46

Papadopoulo, Theodore, 38, 72
 Patsali, Sofia, 15
 Pezzoni, Michele, 15
 Precioso, Frédéric, 59
 Pronesti, Massimiliano, 67
 Péré, Marielle, 46

Repetto, Marco, 7
 Rey, Victoria E., 22
 Richard, Cédric, 69
 Rocher, Marie, 32
 Rodio, Angelo, 36
 Romero-Ternero, María del Carmen, 22
 Rosaz, Julie, 41
 Roux, Jérémie, 46

Schmutz, Hugo, 30, 44
 Sedlar, Sara, 72
 Sermesant, Maxime, 2, 26, 32, 78
 Sessa, Valentina, 70
 Sinclair-Desgagné, Bernard, 58
 Sladen, Anthony, 69
 Sportisse, Aude, 44
 Spyropoulos, Thrasyvoulos, 74

Tarnec, christèle, 19
 thomas, eoin, 52
 Tiblias, Federico, 65
 Tiezzi, Matteo, 40
 Tissot, Romain, 13
 Topin, Jérémie, 21, 62
 Totic, Tamara, 24

van den Ende, Martijn, 69
 Vayer, Titouan, 60
 Vidal, Richard, 34
 Villata, Serena, 17, 33, 79
 Vincent-Cuaz, Cédric, 60
 Vorakitphan, Vorakit, 33

Wang, Zihao, 78

Wlodarczyk-Boudenot, Ewa, 63

Xu, Chuan, 52

YANG, Yingyu, 32, 78
 Yepremyan, Artyom, 42