L’IA pour le monitoring : à la recherche de l’inconnu

Détection d’anomalies sur séries temporelles

Présentation par Marielle Malfante

Travaux effectués au CEA / LIST / DSCIN / LIIM, in collaboration with …

Health / Medical	Environnemental	Other

• Epilepsy (EEG data, Clinatec)
• Cellular biologie (dry mass of cells, CEA/DTBS, CNRS)
	• Seismic data
	  ○ Volcano-seismic (Ubinas, Merapi, IGP, BPPTKG, IPGP, GIPSA-Lab)
	  ○ Mars (Insight mission, IPGP)
	  ○ French territory monitoring (CEA/DASE)
	  ○ LFE (ISTerre)
	• Bioacoustic (Biophonia)
	• Oceans vitality through scalops behavior monitoring (BeBest, GIPSA-Lab)
		• Ship sounds (predictive maintenance)
		• Industrial data
		• Synthetic data sometimes …

« Small AI » ?

L'ACTIVITÉ « IA ALGORITHMES INNOVANTS » EN QUELQUES MOTS (1/2)

- **Création de l'équipe fin 2019:**
  - 2019 : 1ère arrivée
  - 2022 : 4 ingénieurs-chercheurs permanents, 4 thèses, 2 postdoc.
- **Point de départ et fil conducteur scientifique : les « petites IA »**

« FAIRE TRANSITER L'INFORMATION PLUTÔT QUE LES DONNÉES »

Contexte smart sensor, CPS, EdgeIA

(Très) basse consommation, objets de taille réduite (peu de mémoire, puissance de calcul limitée)
Haute efficacité énergétique (algorithmes « optimisés » : compromis précision vs ressources disponibles)

« AMENER L'IA SUR DES APPLICATIONS MONDE RÉEL »

Problématiques théoriques à adresser :
l'IA pour le monde réel pose des questions qui requièrent des innovations algorithmiques en comparaison de l'IA de laboratoire.
Collaborations « monde réel » nécessaires pour valider la pertinence des approches sur divers contextes applicatifs,
par exemple pour le monitoring environnemental ou de la personne.

« Small AI » ?

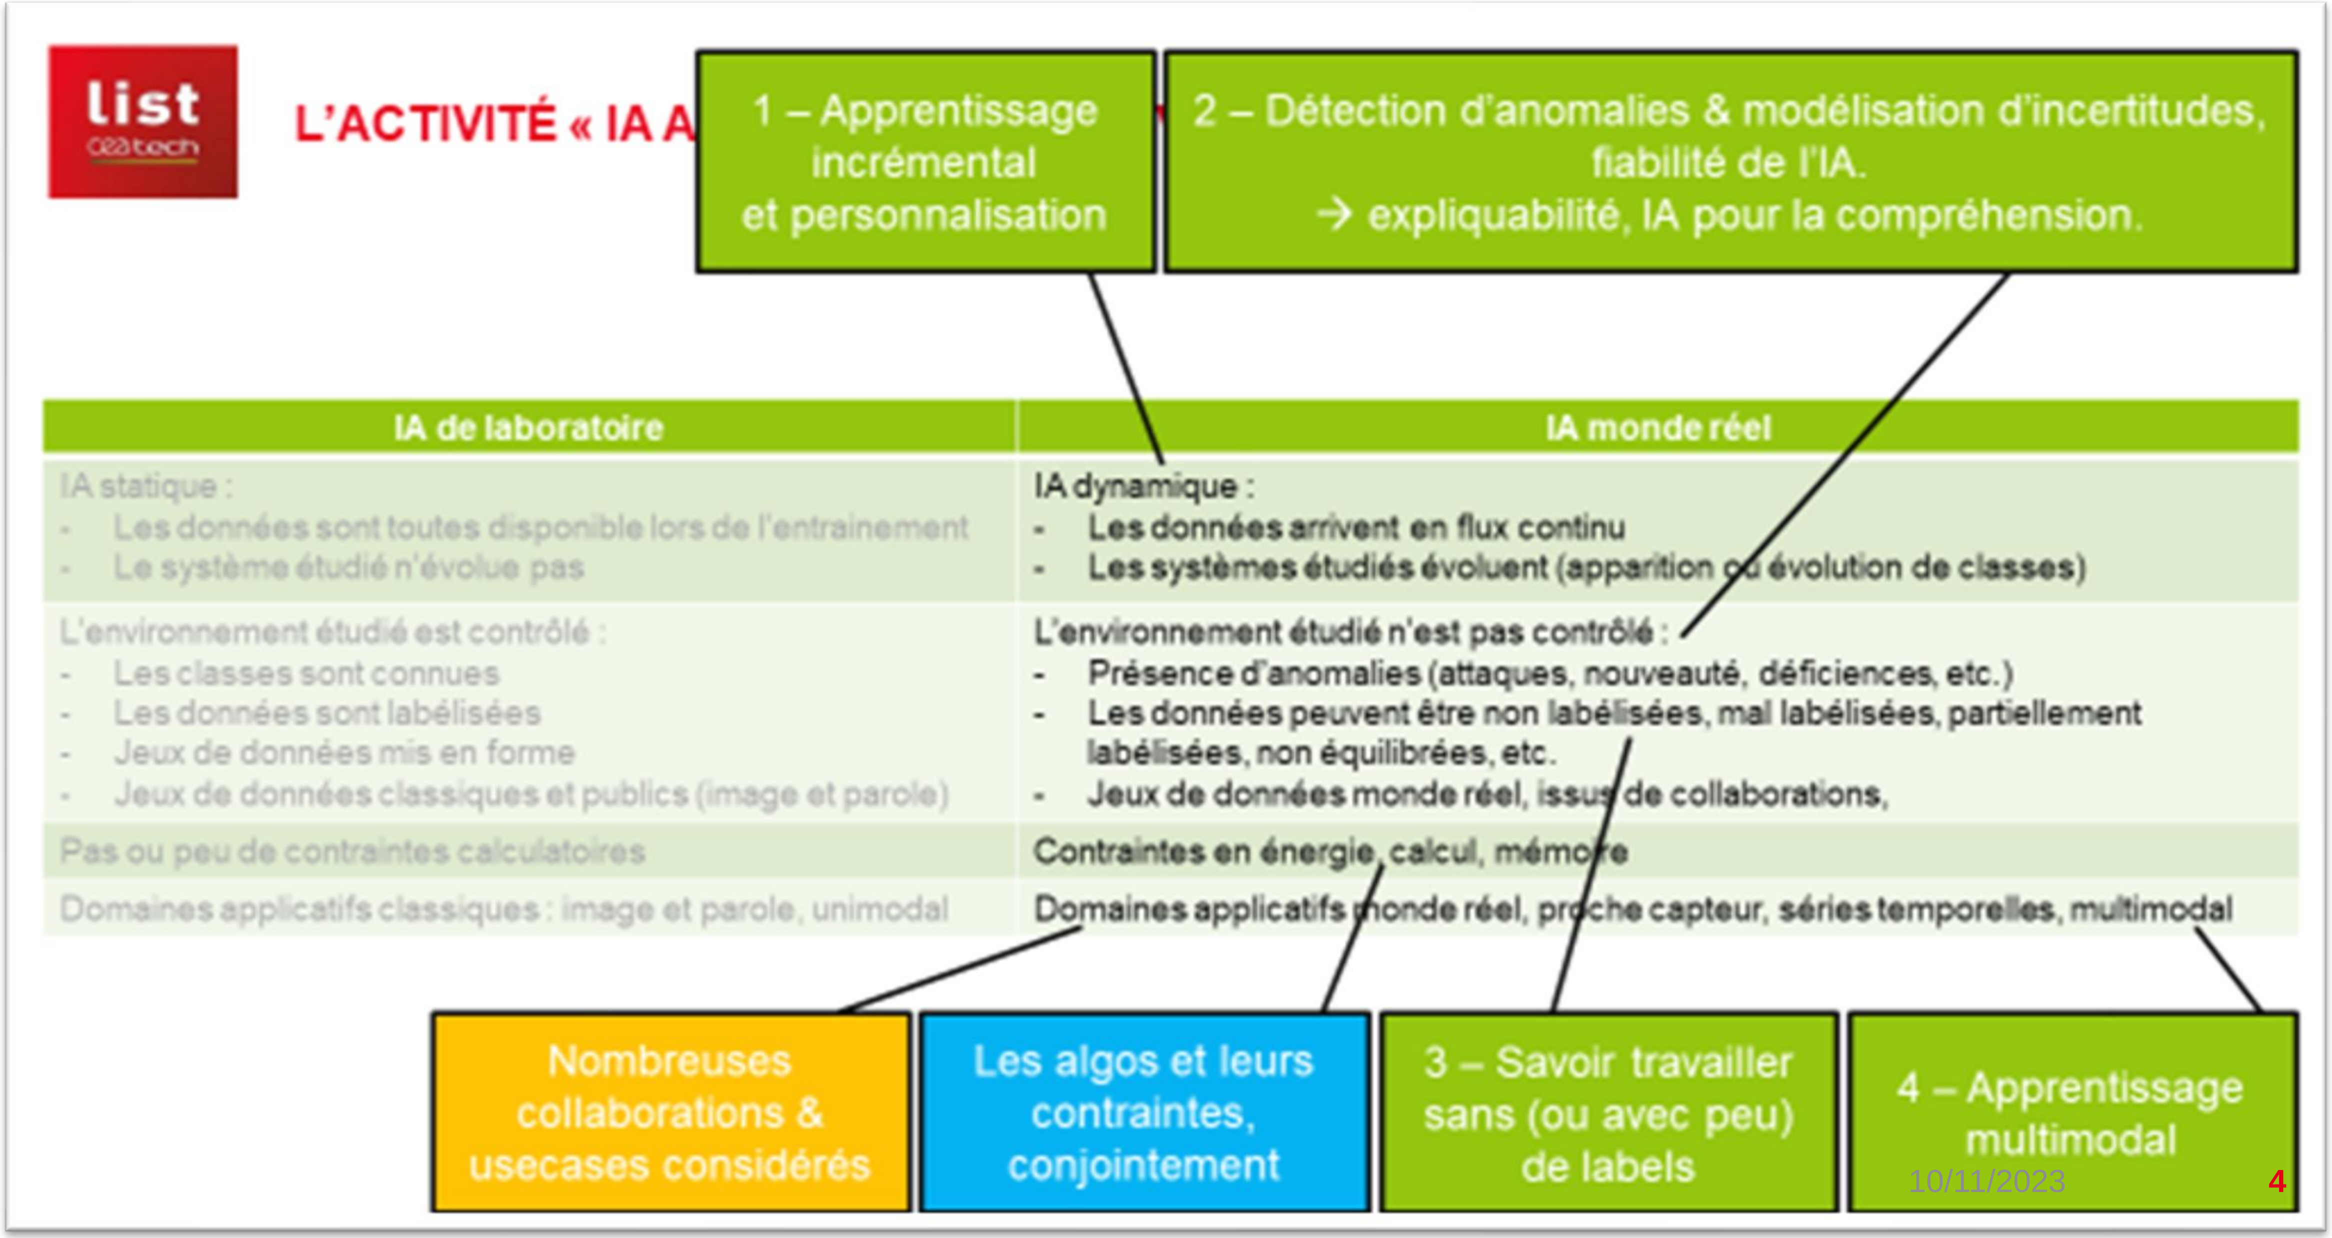

list ceatech

L'ACTIVITÉ « IA A

1 – Apprentissage incrémental et personnalisation

2 – Détection d'anomalies & modélisation d'incertitudes, fiabilité de l'IA.
→ expliquabilité, IA pour la compréhension.

IA de laboratoire	IA monde réel
IA statique : - Les données sont toutes disponible lors de l'entrainement - Le système étudié n'évolue pas	IA dynamique : - Les données arrivent en flux continu - Les systèmes étudiés évoluent (apparition ou évolution de classes)
L'environnement étudié est contrôlé : - Les classes sont connues - Les données sont labélisées - Jeux de données mis en forme - Jeux de données classiques et publics (image et parole)	L'environnement étudié n'est pas contrôlé : - Présence d'anomalies (attaques, nouveauté, déficiences, etc.) - Les données peuvent être non labélisées, mal labélisées, partiellement labélisées, non équilibrées, etc. - Jeux de données monde réel, issus de collaborations,
Pas ou peu de contraintes calculatoires	Contraintes en énergie, calcul, mémoire
Domaines applicatifs classiques : image et parole, unimodal	Domaines applicatifs monde réel, proche capteur, séries temporelles, multimodal

Nombreuses collaborations & usecases considérés

Les algos et leurs contraintes, conjointement

3 – Savoir travailler sans (ou avec peu) de labels

4 – Apprentissage multimodal

1. Why do we need anomaly detection?

- Focus on Open Set Recognition

The problem

Most of the time* when training a ML or AI model, we start from a training dataset.

Problem:

This training set will *never ever* fully represent the 'real world'

Open Set Recognition domains formalizes the distinction between:

- The data that we have to build model,
- And the fact that the real world will always be wilder and contain unexpected patterns

Two questions to consider:

- Do we have data?
  Yes / No
  (distinction from ML point of view)
- Is this class « of interest », is it expected? (distinction from human, interpretation point of view)

Unknown Known Classes
= Classes of interest, expected classes, but we do not have data

Known Known Classes
= Classes of interest (positive classes), we have data

Unknown Unknown Classes
= Surprises !

Known Unknown Classes
= Classes of little or no interest (negative classes), we have data

Geng, C., Huang, S. J., & Chen, S. (2020). Recent advances in open set recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, *43*(10), 3614-3631.

So …what should we do?

Unknown Known Classes = Classes of interest, expected classes, but we do not have data	Known Known Classes = Classes of interest (positive classes), we have data
Unknown Unknown Classes = Surprises !	Known Unknown Classes = Classes of little or no interest (negative classes), we have data

Use appropriate methods for each of the quadrant,
of the data type.

- Supervised learning to recognize what you know
- Semi-supervised learning* to detect the unexpected (= say I don't know)
- Among the unexpected, unsupervised learning to try and identify new patterns, new classes

Ok, but if it is that simple, why don't we all do it?
Spoiler: It's actually not that simple ☺

- Supervised
  - → Often, imbalanced dataset, especially between KKC and KUC --- > Focus #2
  - → Reliability issue is not widely acknowledged: we trust the output probabilities we have …. And we should not.
- Semi-supervised --- > Focus #1
  - → Everything in between supervised and unsupervised learning is still a bit fuzzy:
    - For historical reasons (one word, several concept: one-class classifiers VS not fully labeled dataset)
    - *Can be mixed with other types of learning, for instance self-supervised (which is also often mixed with unsupervised : no need of manual labelling does not mean that it is unsupervised)
  - → *Is semi-supervised really semi-supervised ? Can also be unsupervised : hypothesis of large numbers.
- How can do we check model performances when we do not know what we are looking for?
- Unsupervised
  - → is hard to interpret

2. How to AD?

- Focus #1: Finding the unknown with StArDusTS
- Focus #2: Recovering known anomalies by managing inbalanced datasets

StArDusTS Self-supervised Anomaly Detection on Time Series

- Observation #1: We are looking for the unknown: the dataset is not labelled for this anomaly detection task.
- Observation #2: Neural Networks need a form of supervision to be trained.

→ Proposition with StArDusTS approach:

– To keep in mind –

- One of the strength of neural network is their abilities to learn representation of data. This is the idea behind self-supervised learning.
- Conditions for training:
  - Large dataset of continuous time series
  - The dataset reflects the **nominal** time series (anomalies are not represented in the training set, or are in minority). Anomaly = What is not normal.

- Bailly, R., Malfante, M., Allier, C., Ghenim, L., & Mars, J. I. (2021, November). Deep anomaly detection using self-supervised learning: application to time series of cellular data. In ASPAI 2021-3rd International Conference on Advances in Signal Processing and Artificial Intelligence.
- Bailly, R., Malfante, M., Allier, C., Ghenim, L., & Mars, J. (2023, August). Comparaison des capacités prédictives de réseaux de neurones, application à la masse sèche de cellules. In GRETSI 2023-29ème Colloque Francophone de Traitement du Signal et des Images.
- Bailly, R., Malfante, M., Allier, C., Ghenim, L., & Mars, J. I. (2021, June). Self-supervised learning for anomaly detection on time series: application to cellular data. In Conférence sur L'apprentissage Automatique.


1000 µm

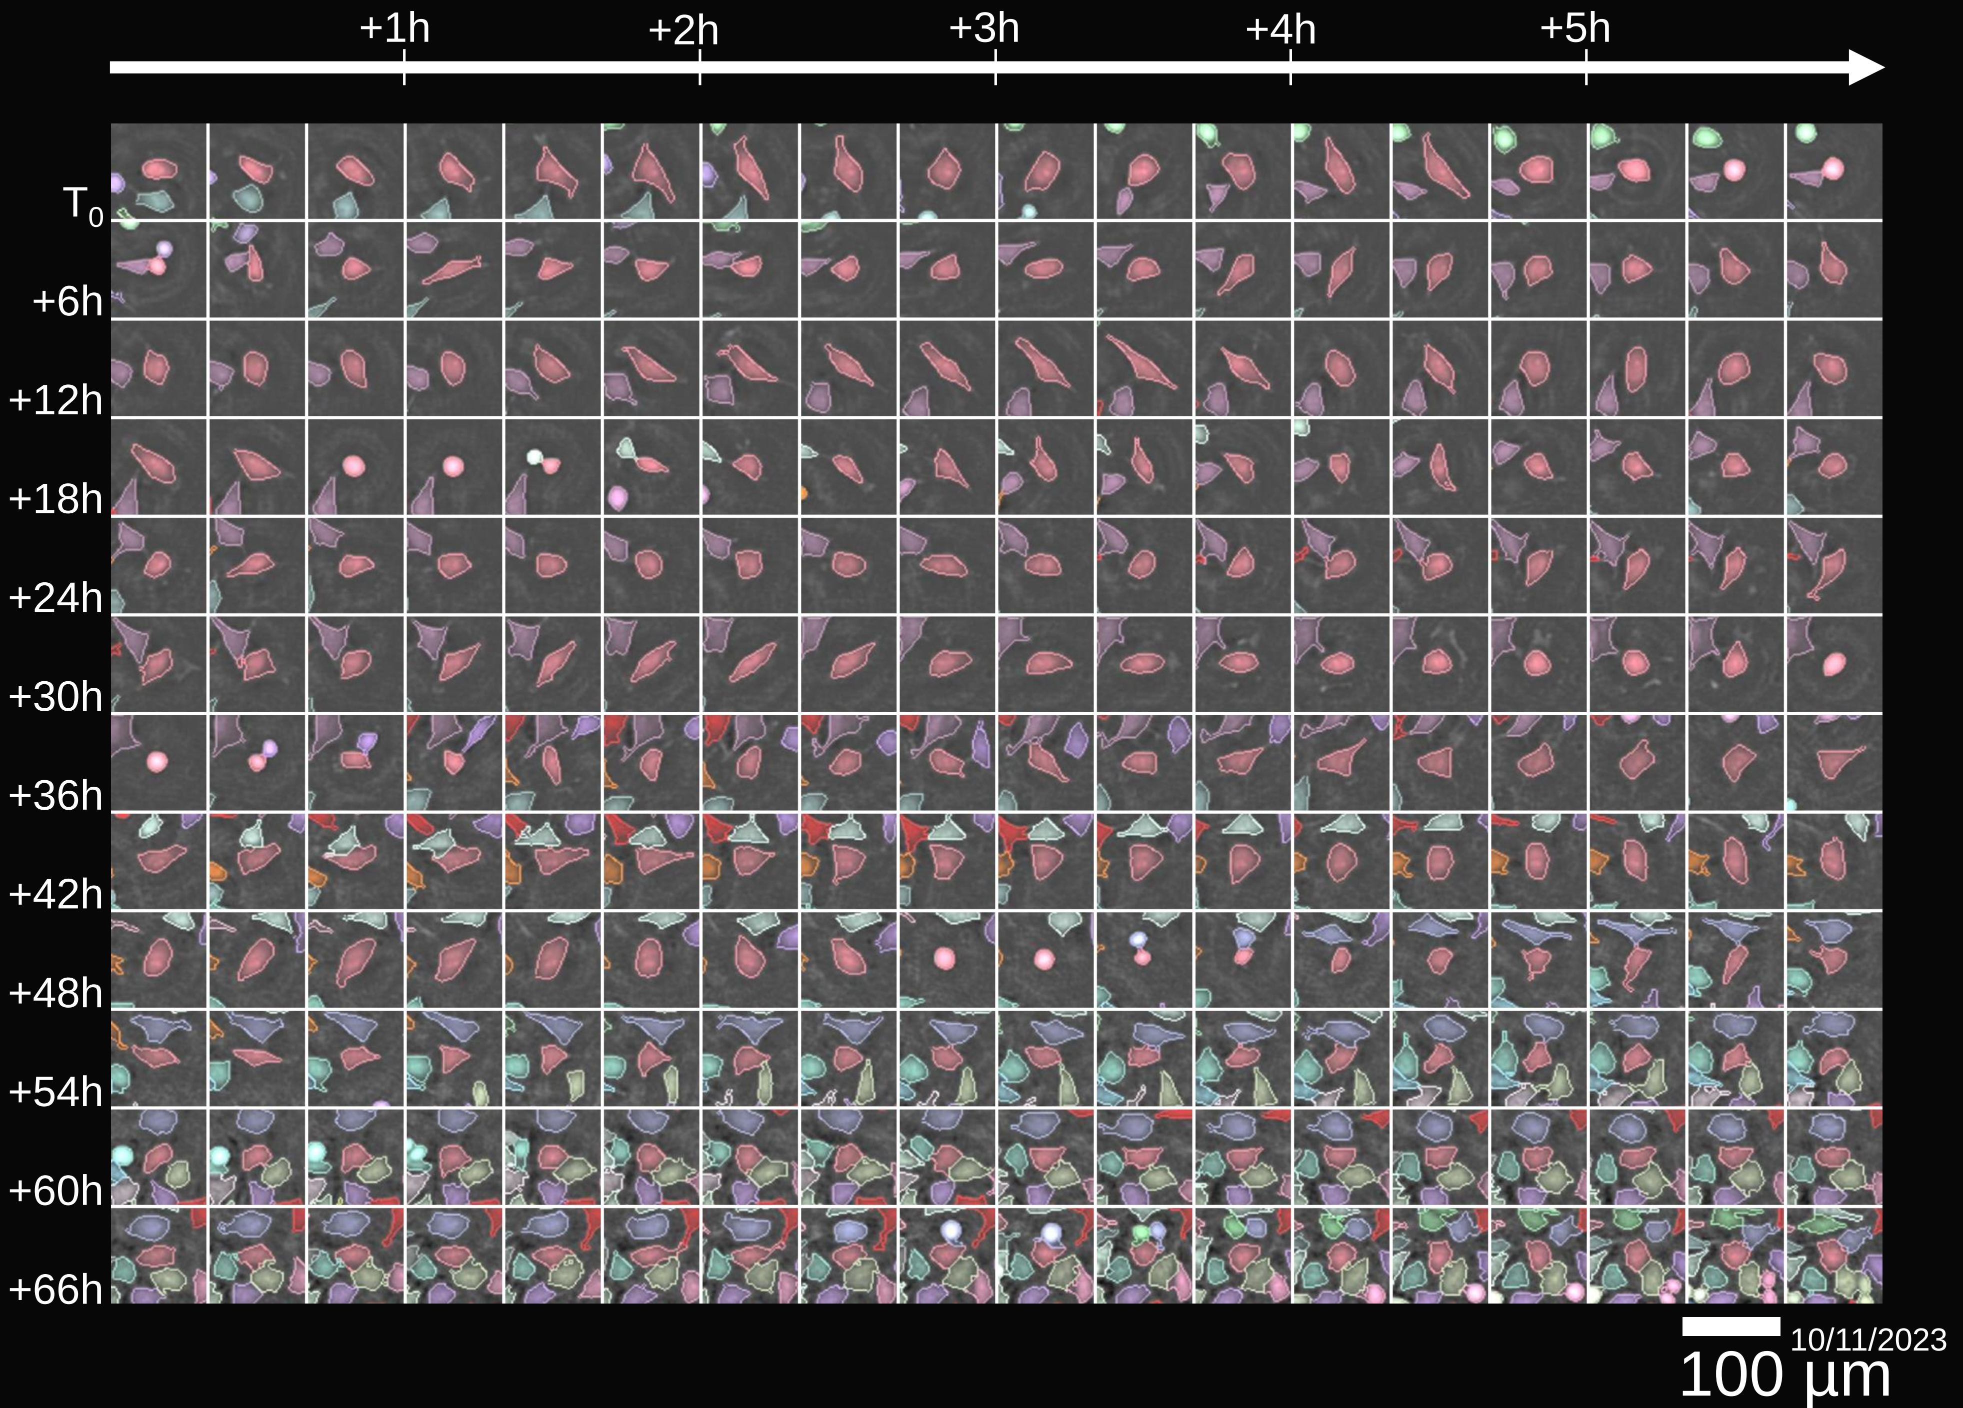
+1h
+2h
+3h
+4h
+5h
T_0
+6h
+12h
+18h
+24h
+30h
+36h
+42h
+48h
+54h
+60h
+66h
100 µm

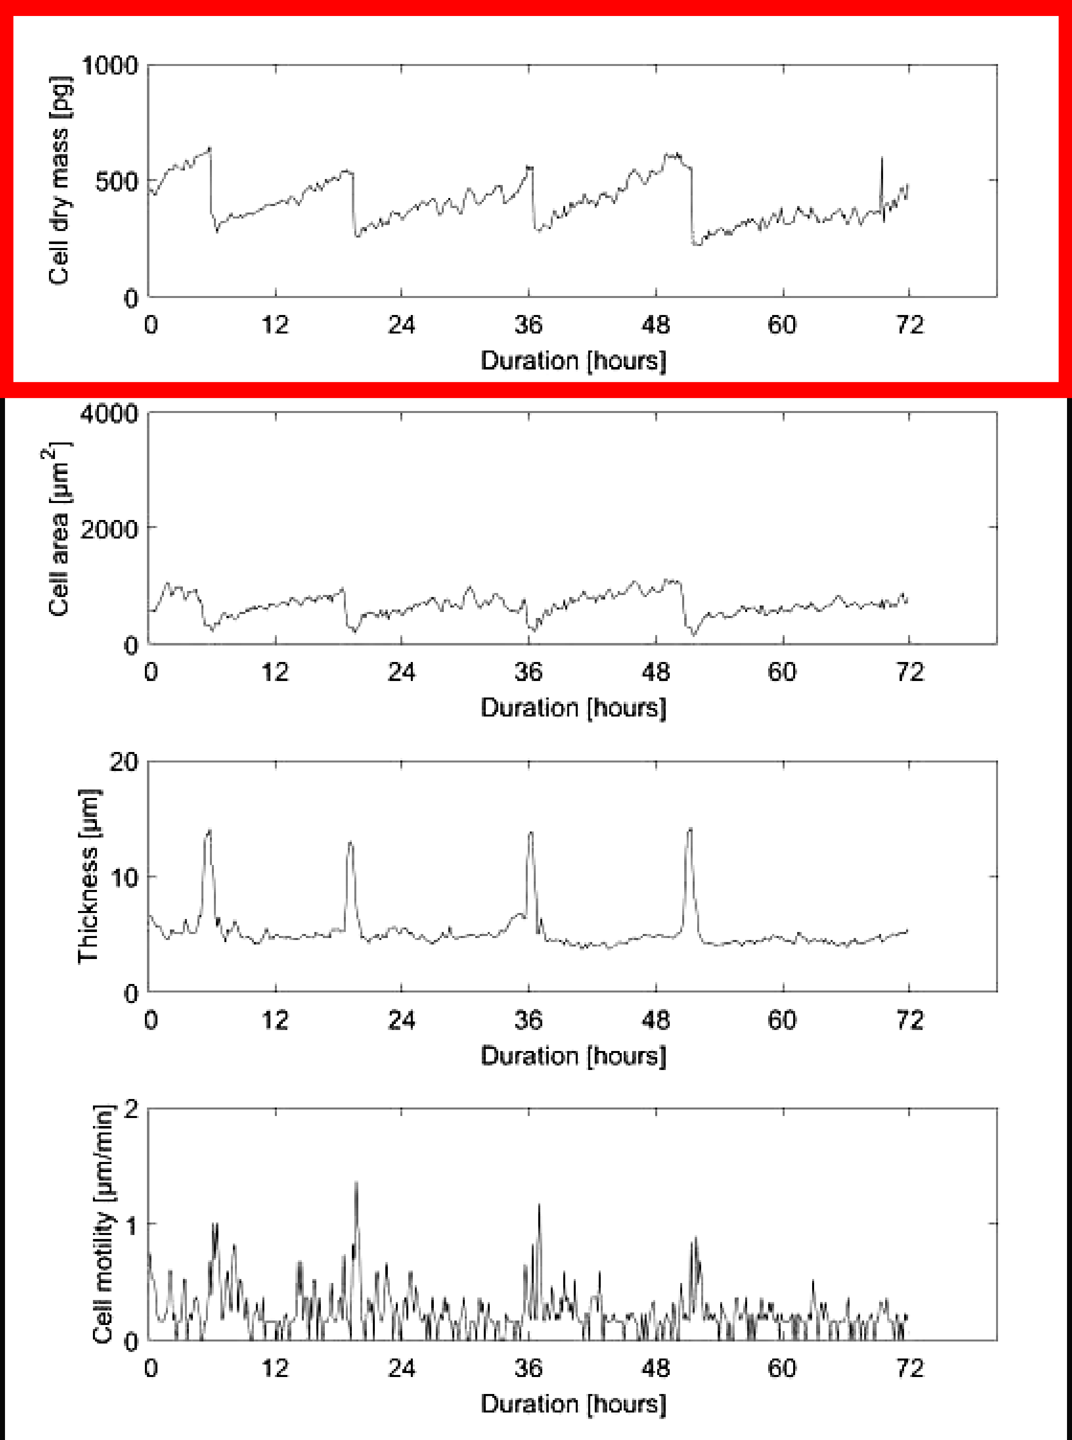
Cell dry mass [pg]
1000
500
0
0 12 24 36 48 60 72
Duration [hours]
Cell area [µm²]
4000
2000
0
0 12 24 36 48 60 72
Duration [hours]
Thickness [µm]
20
10
0
0 12 24 36 48 60 72
Duration [hours]
Cell motility [µm/min]
2
1
0
0 12 24 36 48 60 72
Duration [hours]

Managing inbalanced datasets: when looking for known events is not that simple

- Project CIME (Confiance de l'Intelligence artificielle pour le Monitoring Environnemental): classification task on continuous seismic data (RESIF + STEAD), 32 months of data, 6 classes,
- How to detect the bias?

From Accuracy~0.75

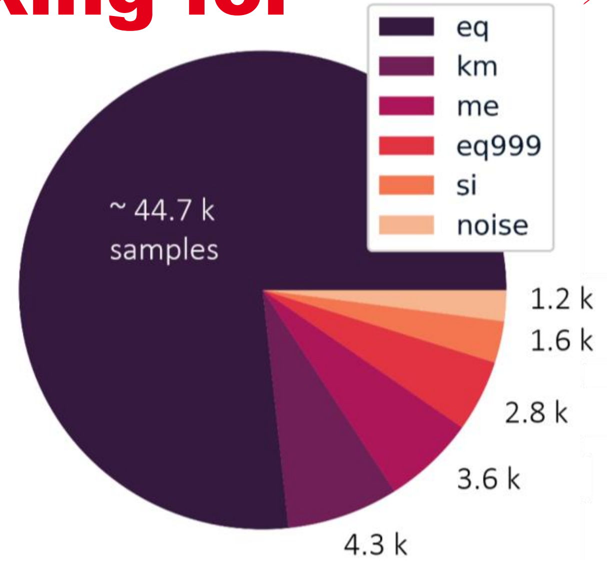

Post-doc: Chantal Van Dinther

Collaboration with CEA / DAM / DASE: Pierre Gaillard, Yoann Cano

Van Dinther, C., Malfante, M., Gaillard, P., & Cano, Y. (2023, August). Réduction du biais dans la classification de données sismique: méthodes de gestion des jeux de données asymétriques. In *GRETSI'2023*.
Van Dinther, C., Malfante, M., Gaillard, P., & Cano, Y. (2023). *Increasing the reliability of seismic classification: A comparison of strategies to deal with class size imbalanced datasets* (No. EGU23-16410). Copernicus Meetings.

3. In a nutshell

Anomaly Detection: Take home messages

- Why considering anomaly detection?
  - To improve the reliability of AI system when used in Open Set Recognition context
  - Also to develop models when you do not have labels.
  - Also to understand: anomaly = source of knowledge
- Careful to biased models, one metric is often not enough (also consider energy consumption, memory storage, environmental impact, etc)
- Alternative to anomaly detection? → Uncertainty modeling

- What can we do once we have detected unusual patterns?
  - Study them! & incorporate them in future classification models → AI to understand
  - Also, act to improve the reliability of future models:
    - Explainabililty
    - Multimodality

Just to keep in mind ...

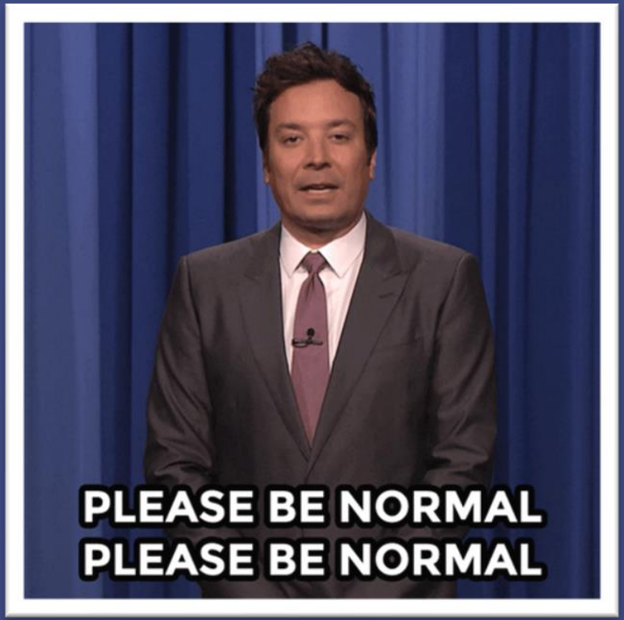

Thank you for your attention ☺

Questions? Reactions? Comments?

Contact info: Marielle.malfante@cea.fr

Introduction on Lens free imaging

- Collaboration with CEA / LETI / DTBS, Carnot DAV Micro3DN project: Cédric Allier, Chiara Paviolo, Lamya Ghenim, Sophie Morales, Caroline Paulus, etc…

- Idea
  - Microscopy, but optical lens are replaced by reconstruction algorithms.
  - Lens-free imaging allows data collection of thousands of cell growth, in particular of time series of their mass.
  - Also smaller than classic microscops

- Dataset
  - Unlabelled
  - 30hours long time series, sampled at f_s = 10 min
  - Training : 189 565 time series,
  - Validation : 53 056 time series,
  - Testing 1 & 2 : 59 719 & 827 733 time series

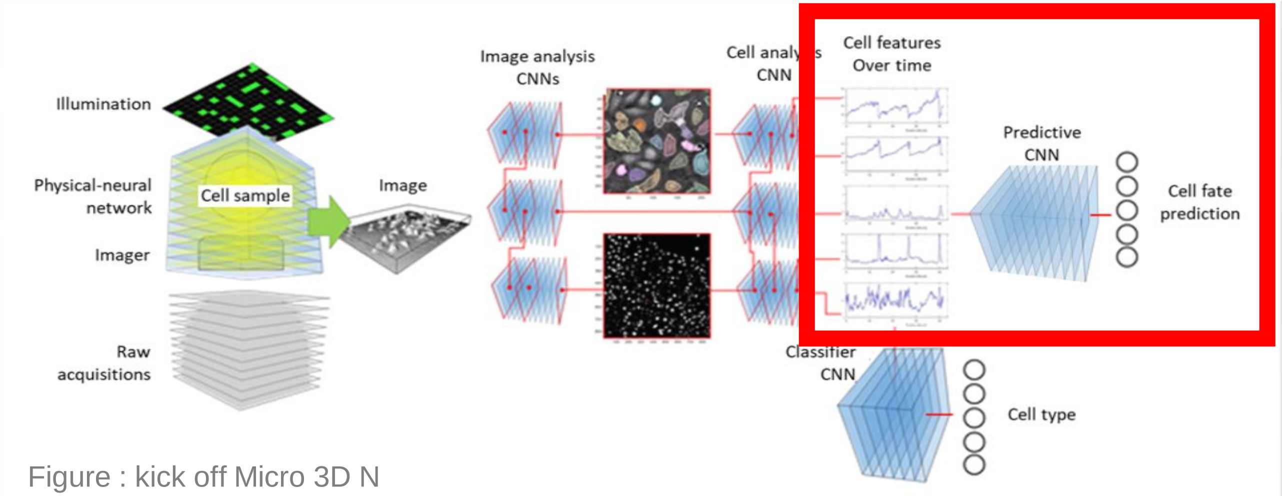

Figure : kick off Micro 3D N

■ More results on lens free imaging

#1 Validation of predictive capabilities

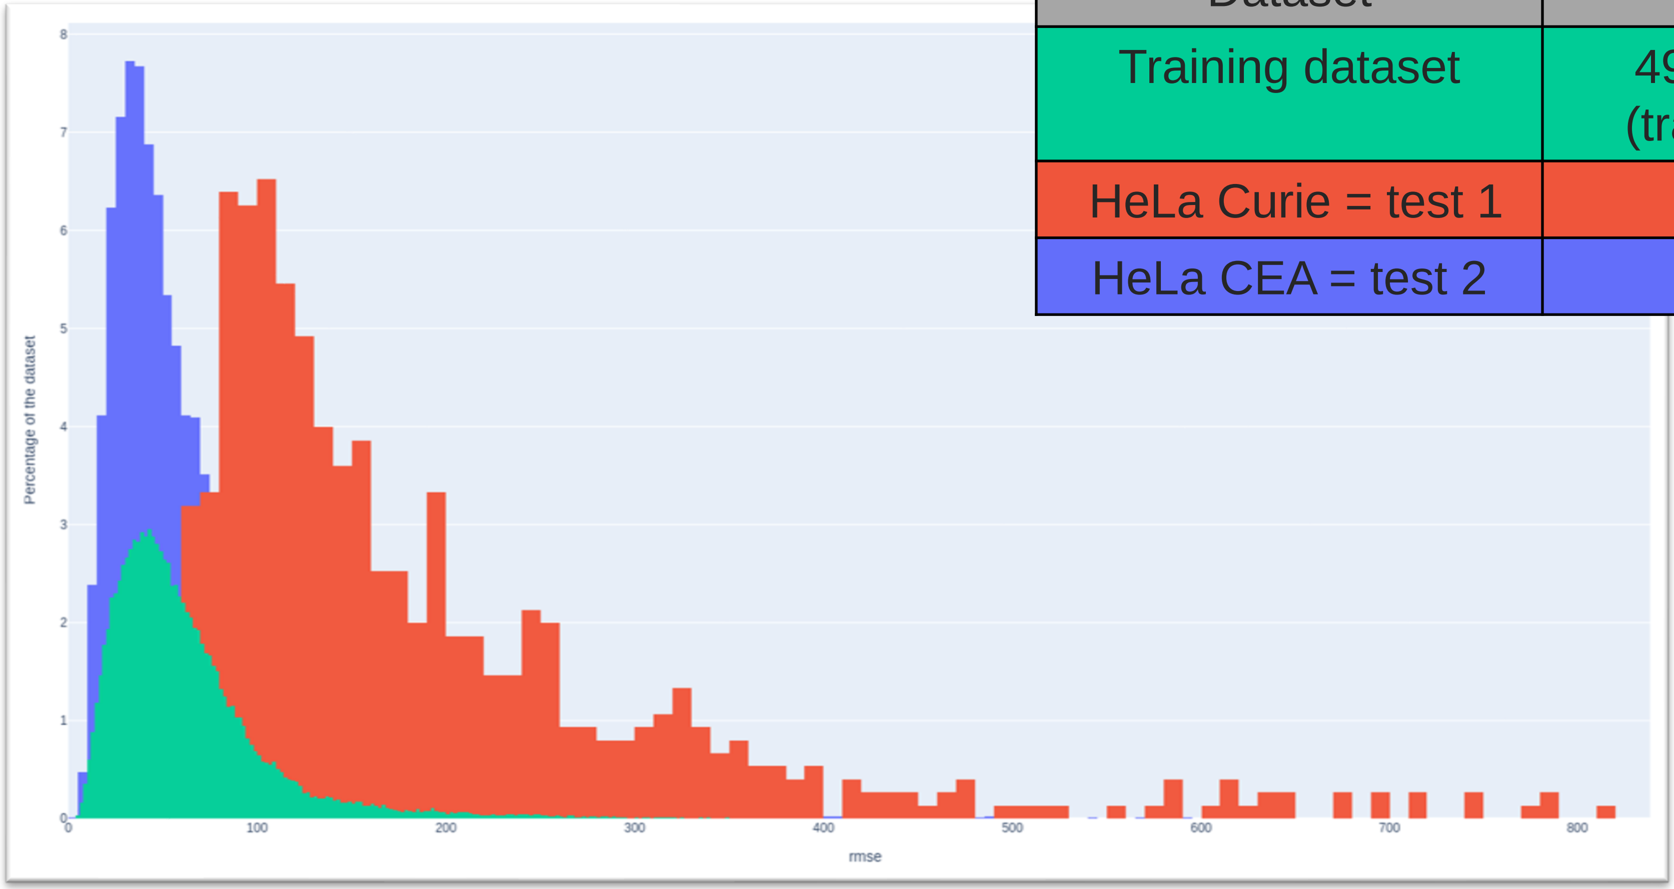

Dataset	Number of data	
Training dataset	496 311 (training)	57 310 (val.)
HeLa Curie = test 1	59 719 (exp 1a)	
HeLa CEA = test 2	827 733 (exp 2)	

#1 Validation of predictive capabilities

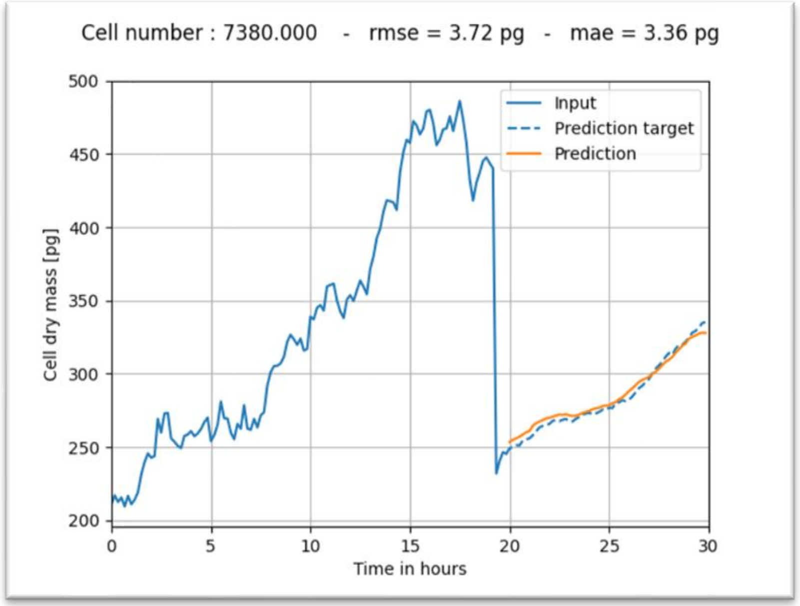

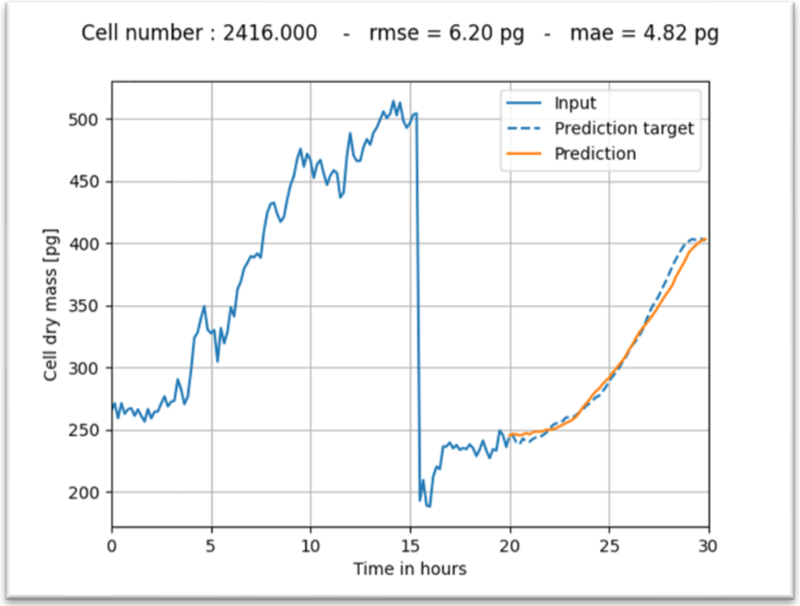

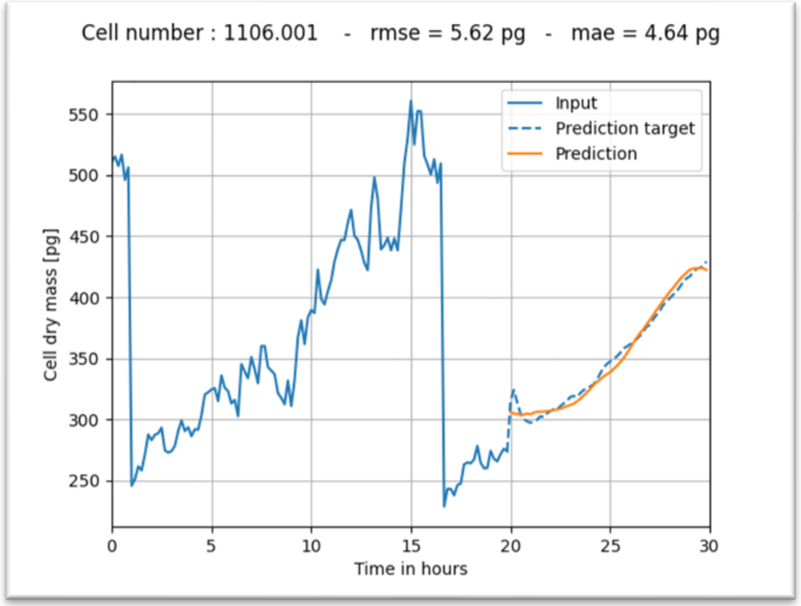

#2 Bad predictions… are anomalies

After manual review of the raised anomalies, 4 causes are identified:

1.	Anomaly on the cell	40%
2.	Anomaly on the cell, causing an anomaly on segmentation or tracking algorithms	31%
3.	Problem on the segmentation or tracking algorithm	26%
4.	False detection: bad prediction made by the network	3%

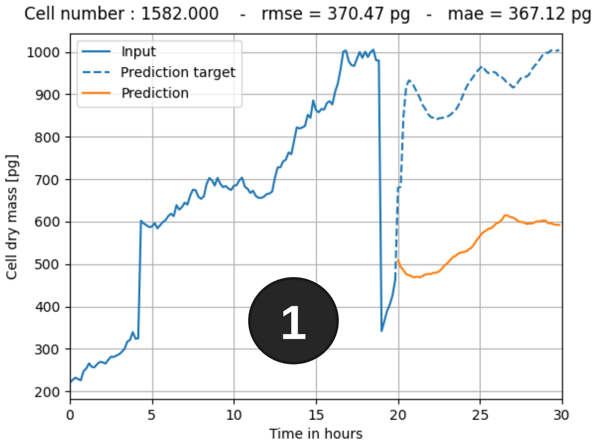

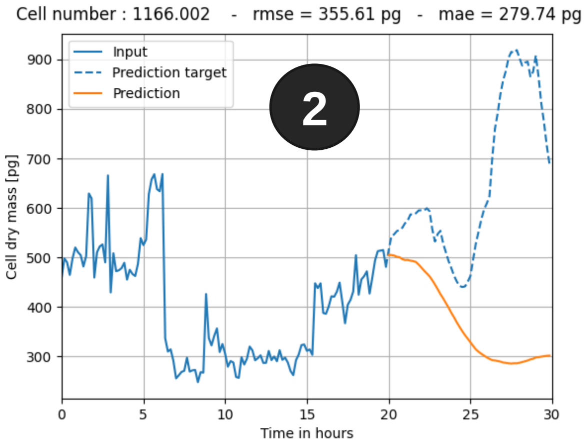

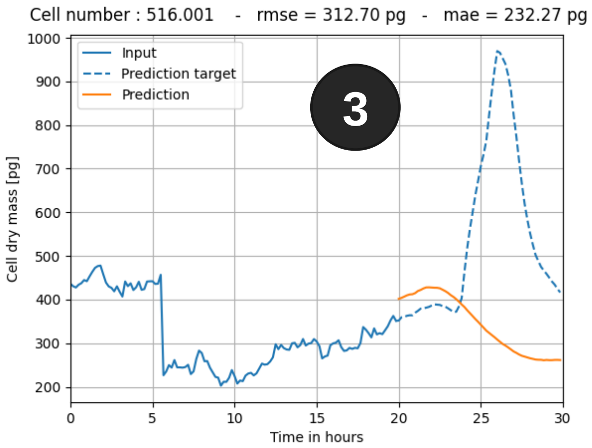

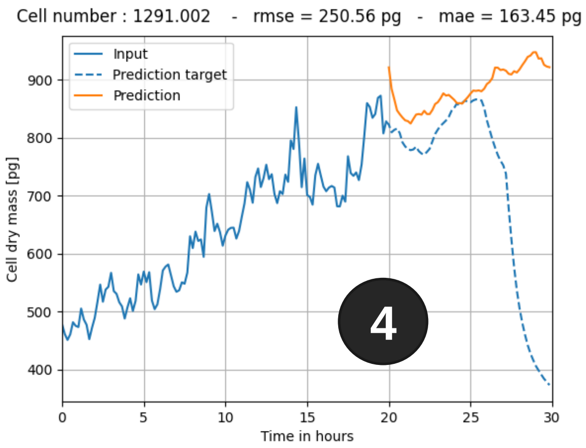

#3 What about the network architecture?

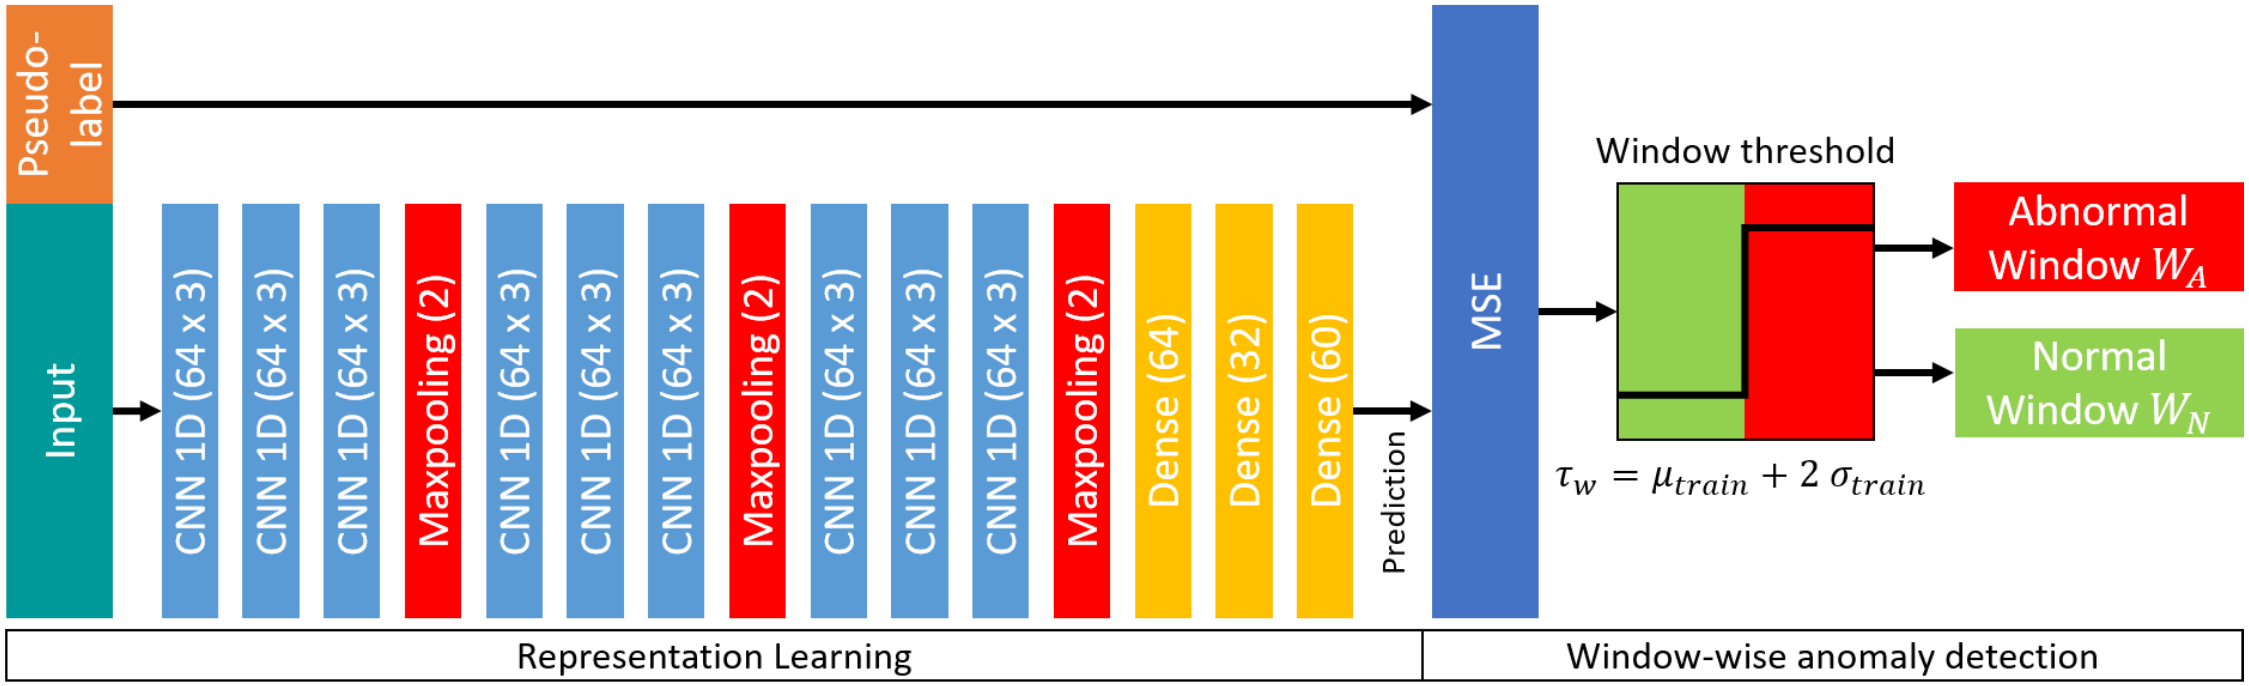

#3 What about the network architecture?

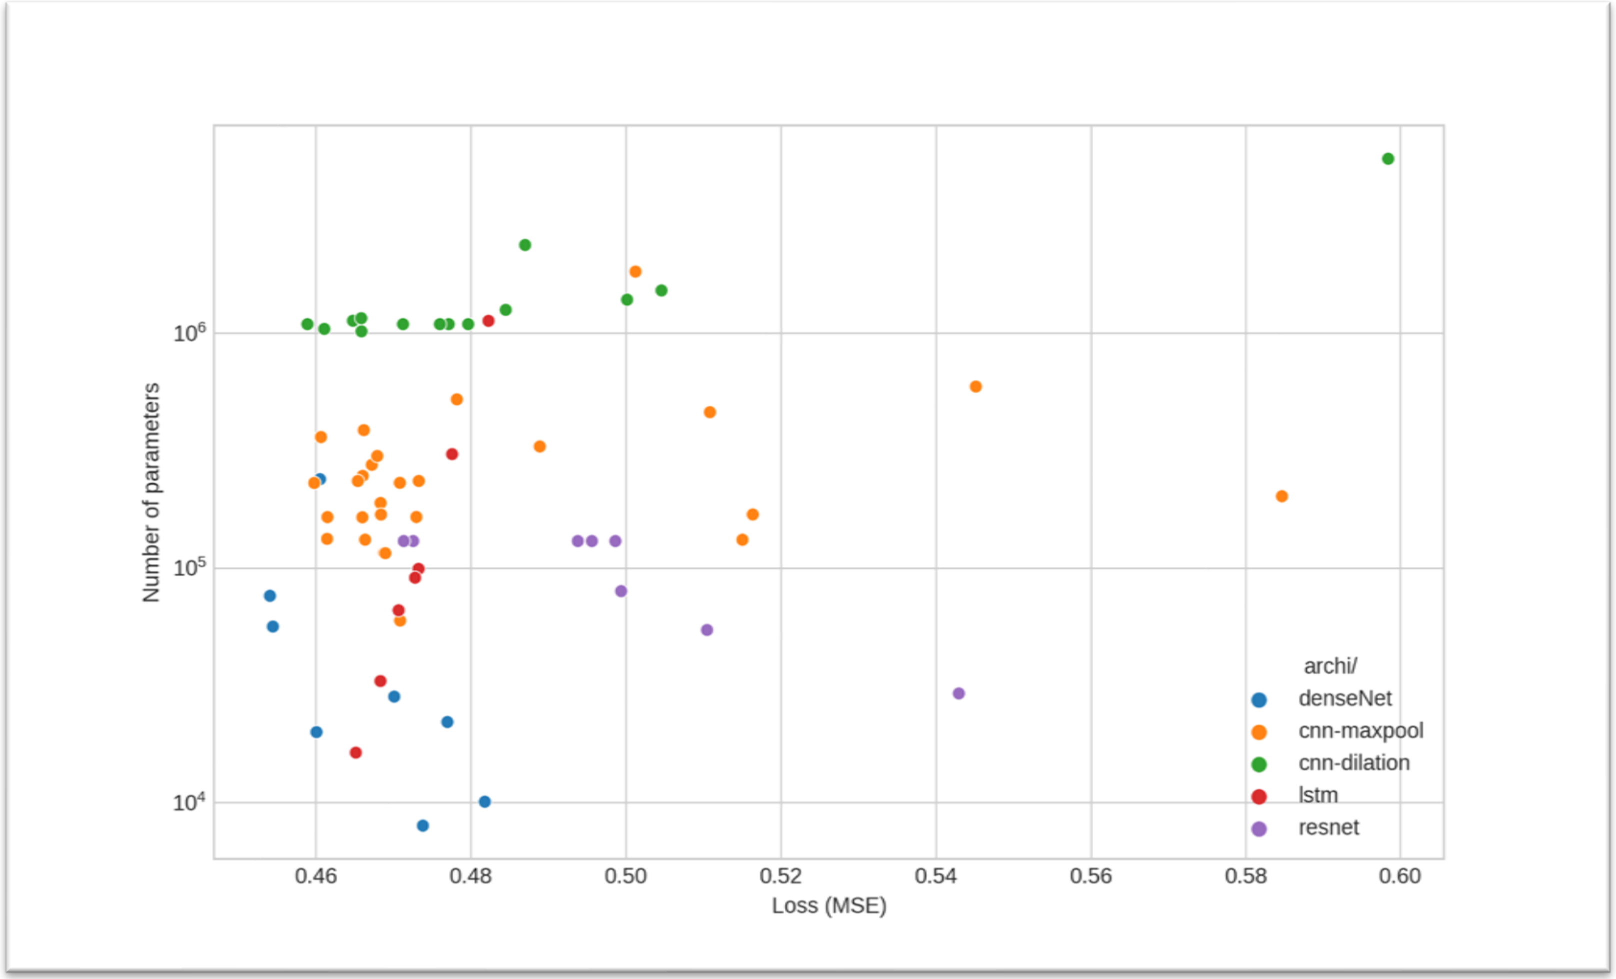

Benchmark on networks architectures:

- Réseaux CNN1d
- Réseaux denses
- Réseaux type ResNet
- Réseaux LSTM

With the following metrics:

- MSE,
- Number of parameters
- Training time
- Inference time

#4 On continuous time series

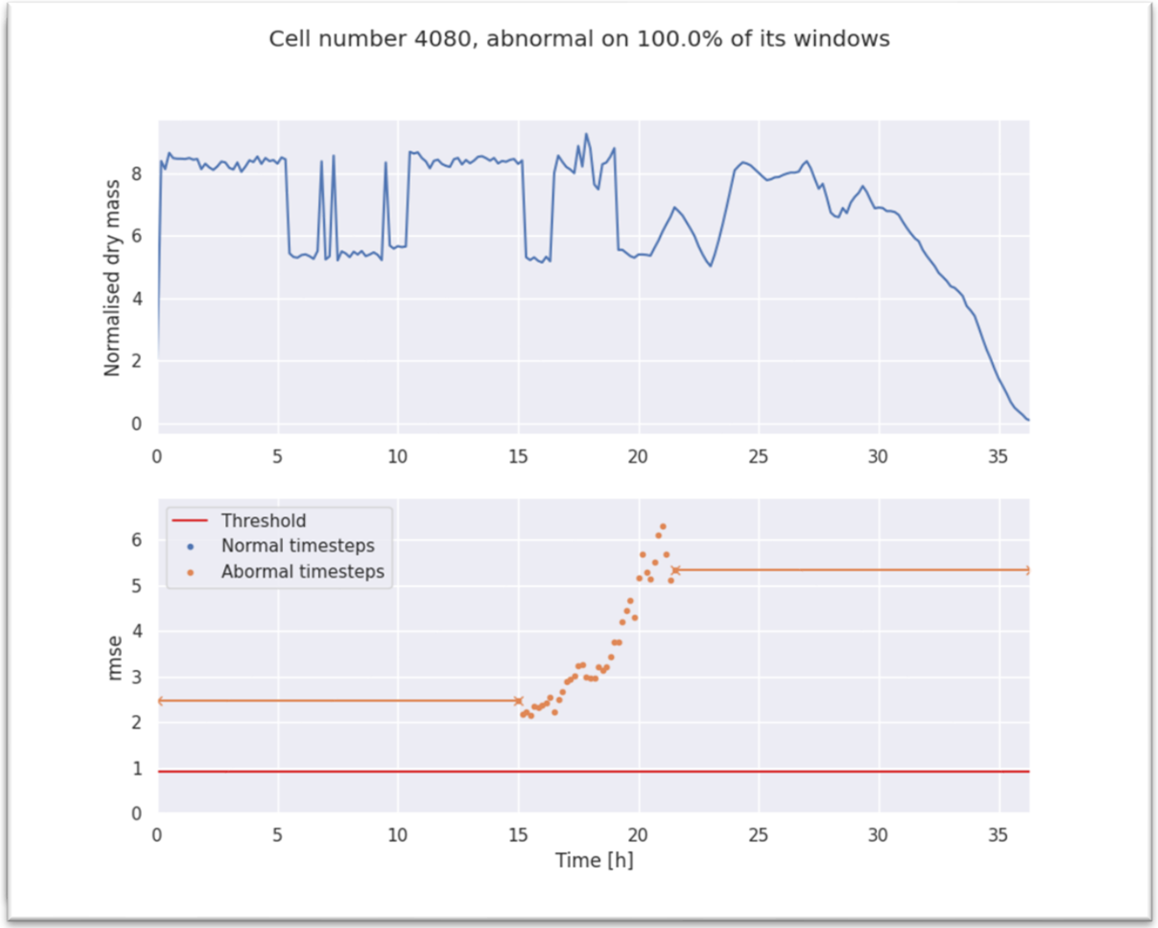

Use of the model on longer time series, inclusion on temporal coherence

Is a cell normal or not, can it evolve?