(Real-time) Image processing using machine learning and artificial intelligence for the protection of fusion reactors

13th ITER International School

Erwan Grelier

Based on works from Marie-Hélène Aumeunier, Léo Dubus, Erwan Grelier, Raphaël Mitteau, Victor Moncada, Valentin Gorse, Alexis Juven, Julie Bonnail, Christian Staron, and the WEST team

Fusion machine monitoring and protection

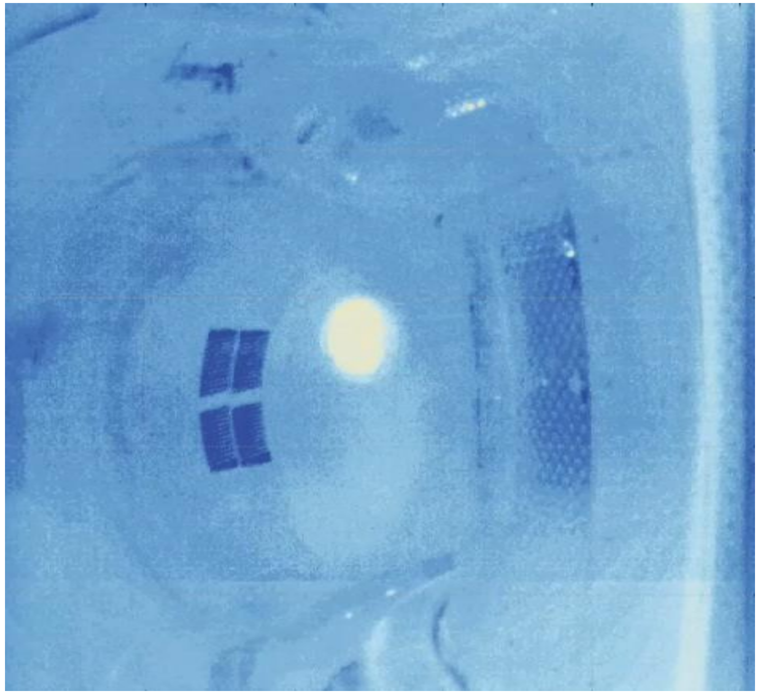

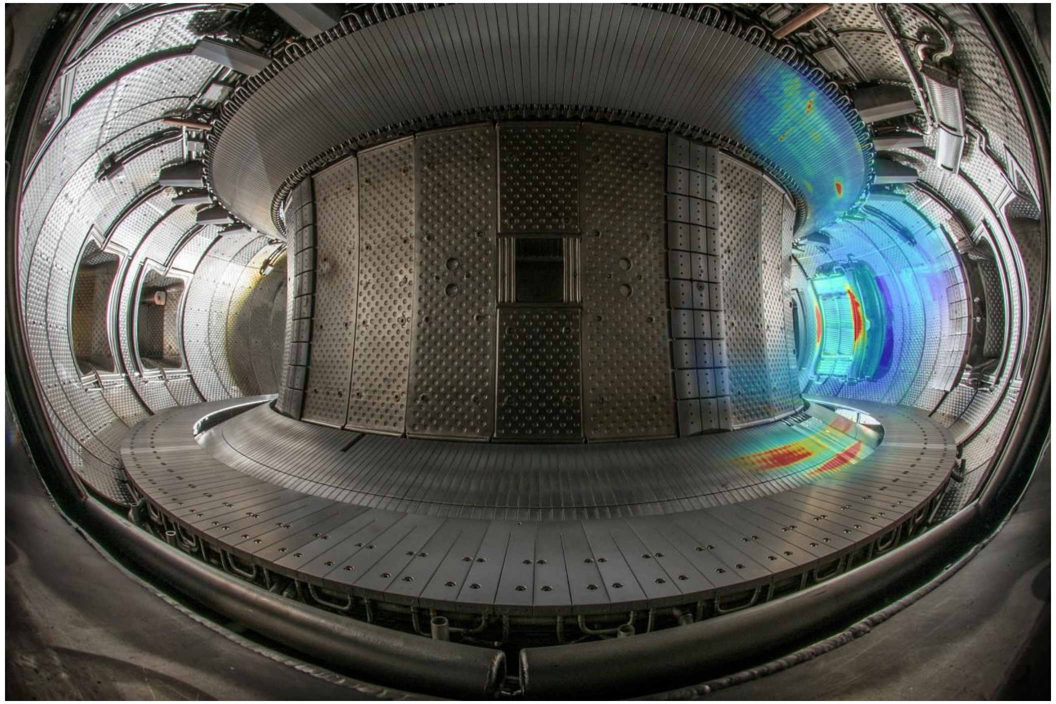

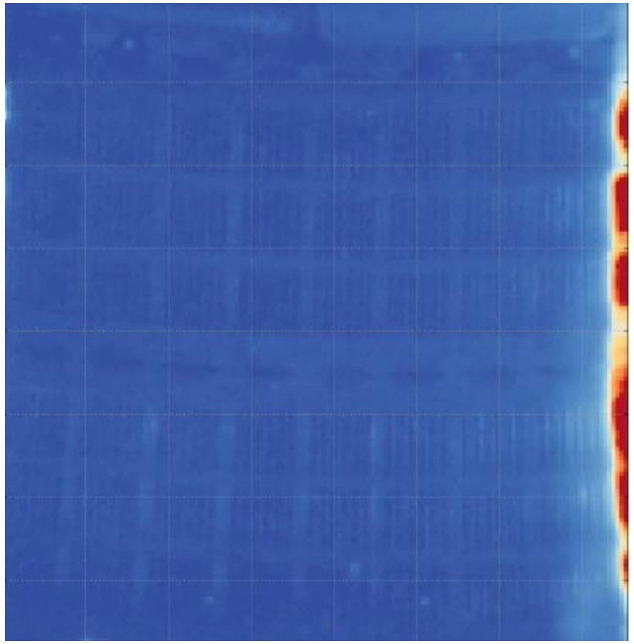

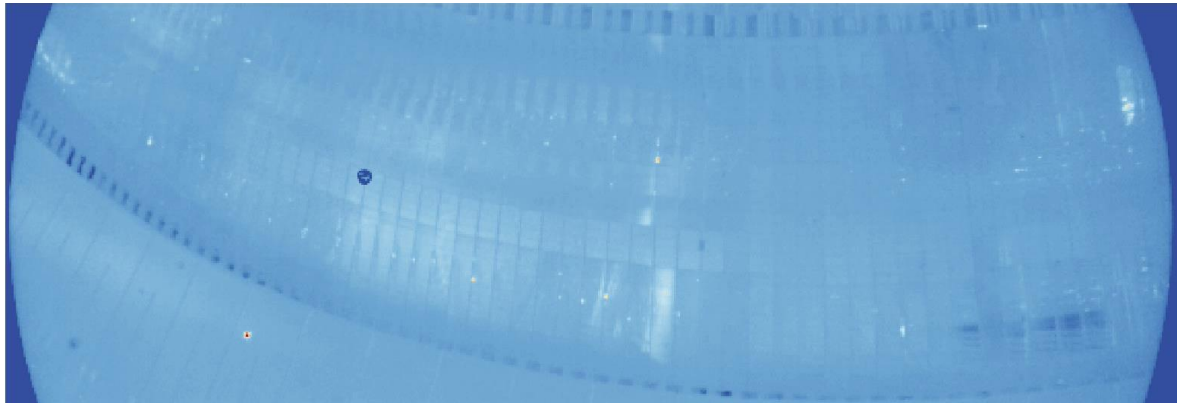

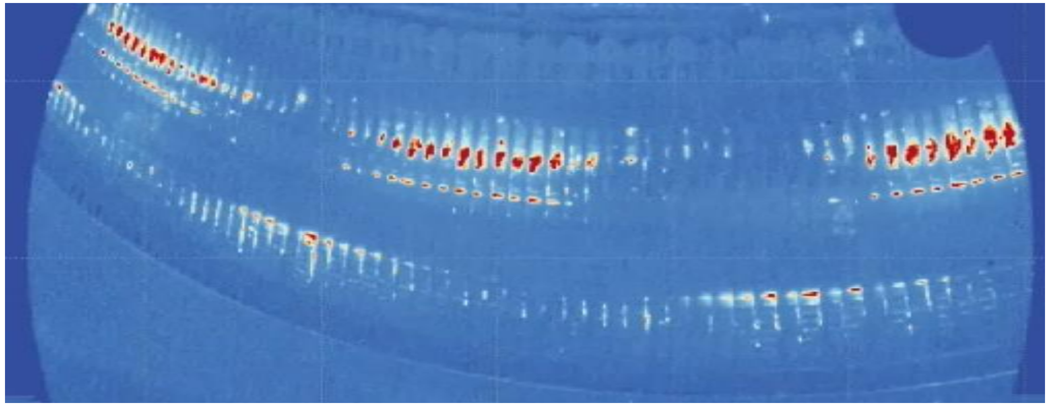

Outline

1. Infrared thermography diagnostic
2. Machine protection using infrared thermography
3. Artificial intelligence for machine protection
4. ML-based methods implemented at WEST
5. Extension to other fusion machines

1. Infrared thermography diagnostic

Infrared thermography diagnostic

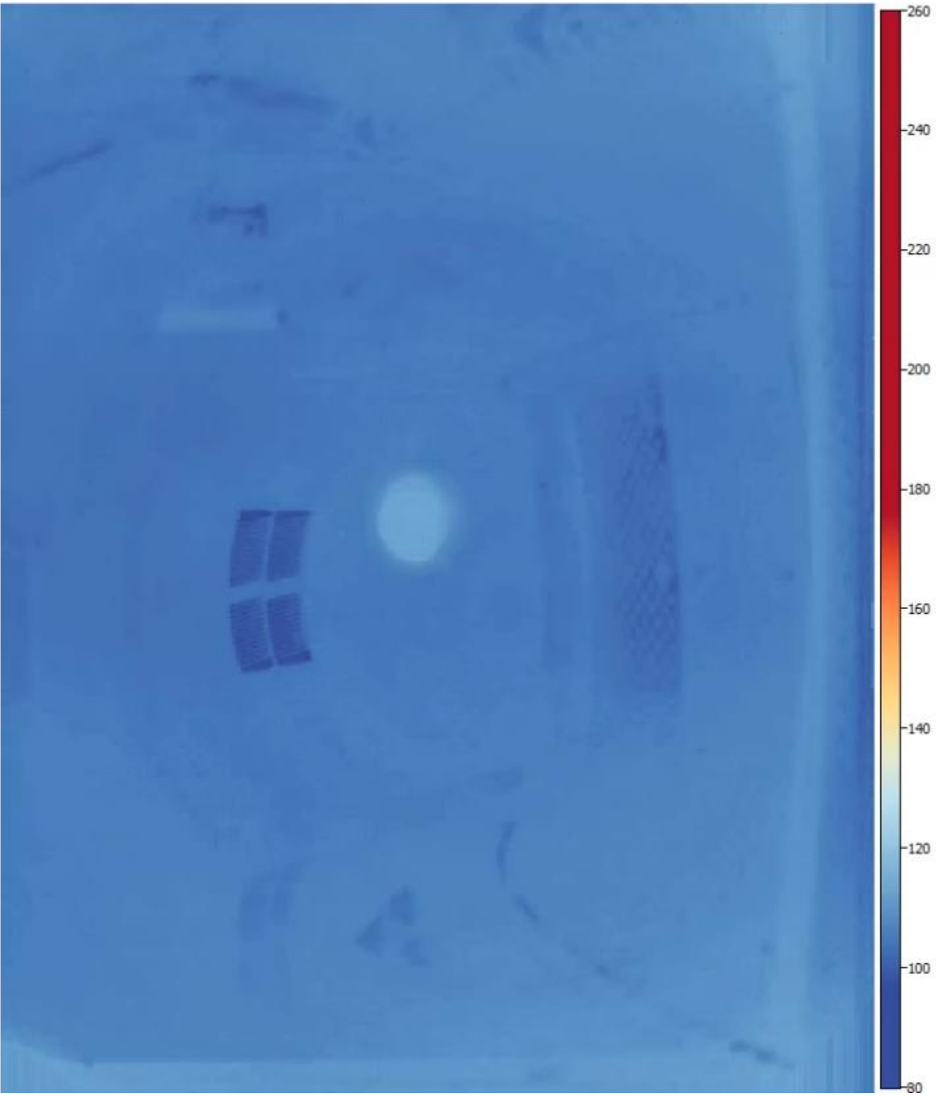

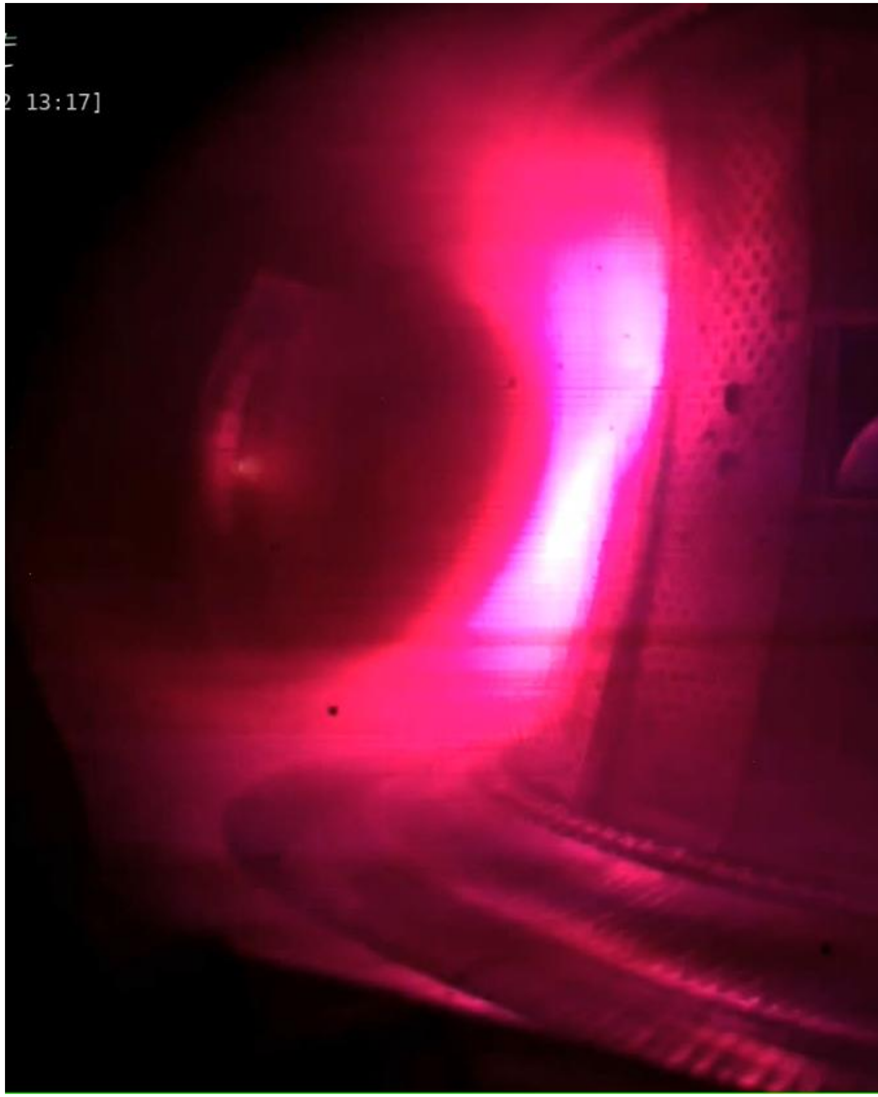

Infrared thermography diagnostic

- Infrared thermography can be used to measure the apparent temperature of the Plasma Facing Components (PFC)
  → Black body temperature / emissivity of 1
  → Actual temperature of the PFCs requires knowing their emissivity
- Furthermore, artifacts in the infrared (IR) images complicate the temperature measures (*e.g.* reflections, UFOs)

→ Reliable monitoring for machine protection using thermography requires complex systems, simulations and machine learning

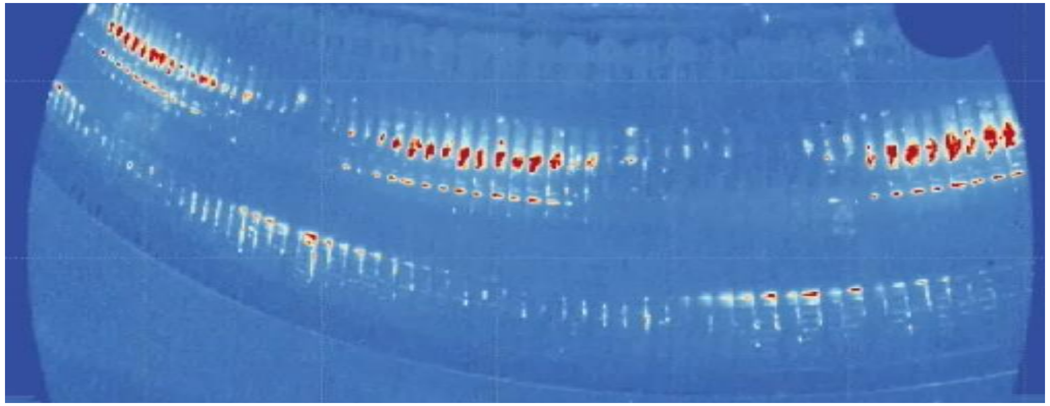

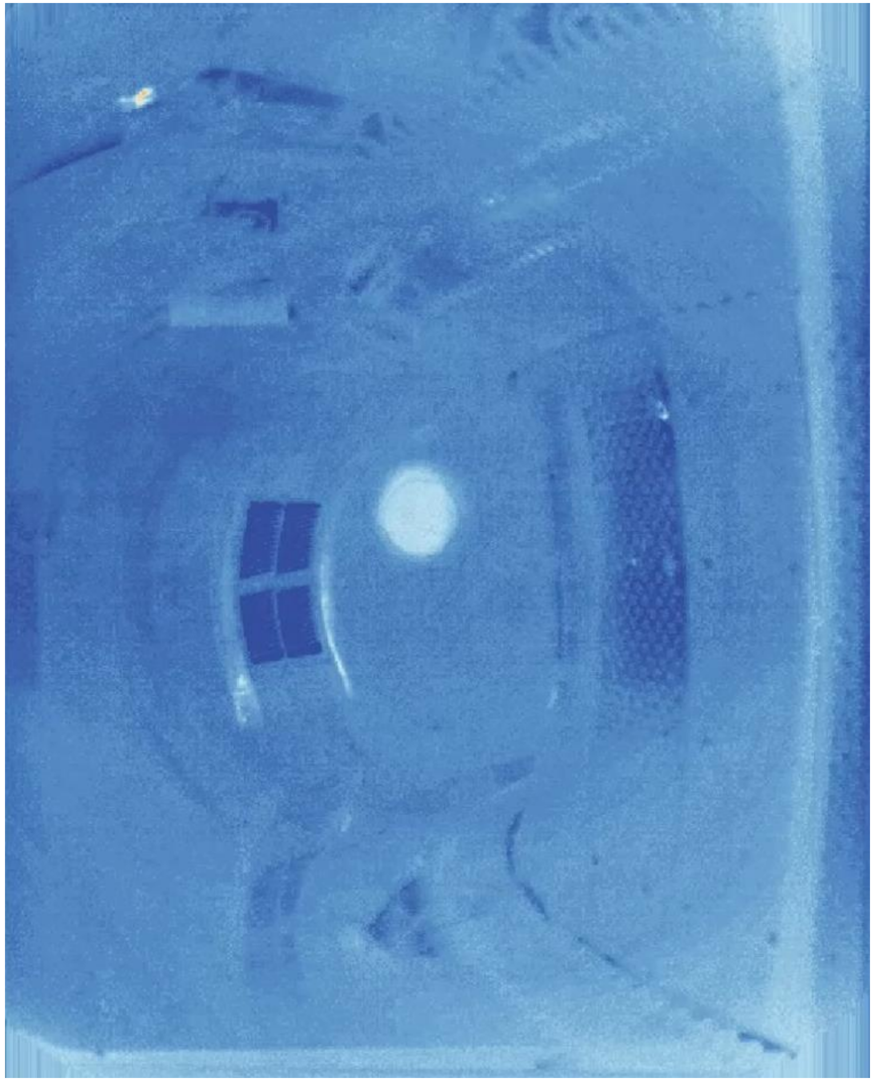

Infrared thermography diagnostic at WEST

- 12 purpose-built thermography cameras and 7 endoscopes
  - covering about 52% (45m2) of the inner wall
  - SWIR and MWIR bands (~1.7 and 4µm)
  - 16bit 512x640 infrared images
  - at 50Hz (20ms between frames)
  - Apparent temperature measured (emissivity of 1)
- First wall coverage and spatial resolution
  - Antennas fully covered (3 ICRH + 2 LHCD) - 4 mm/px
  - 82% of the lower divertor - 2.5 mm/px
  - One wide angle view
  - Very high resolution endoscope for the divertor - 0.1mm/px

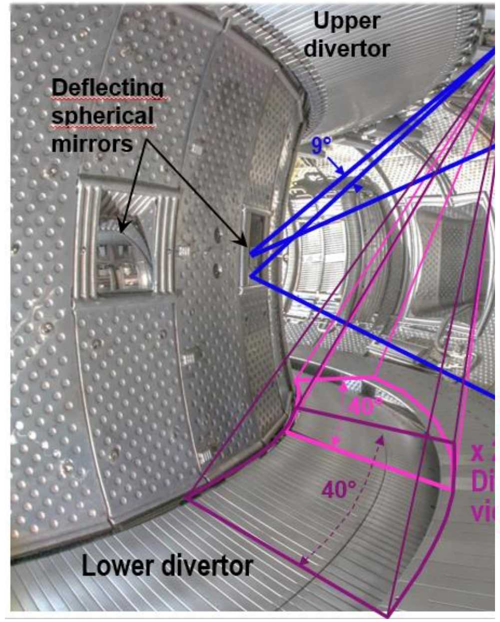

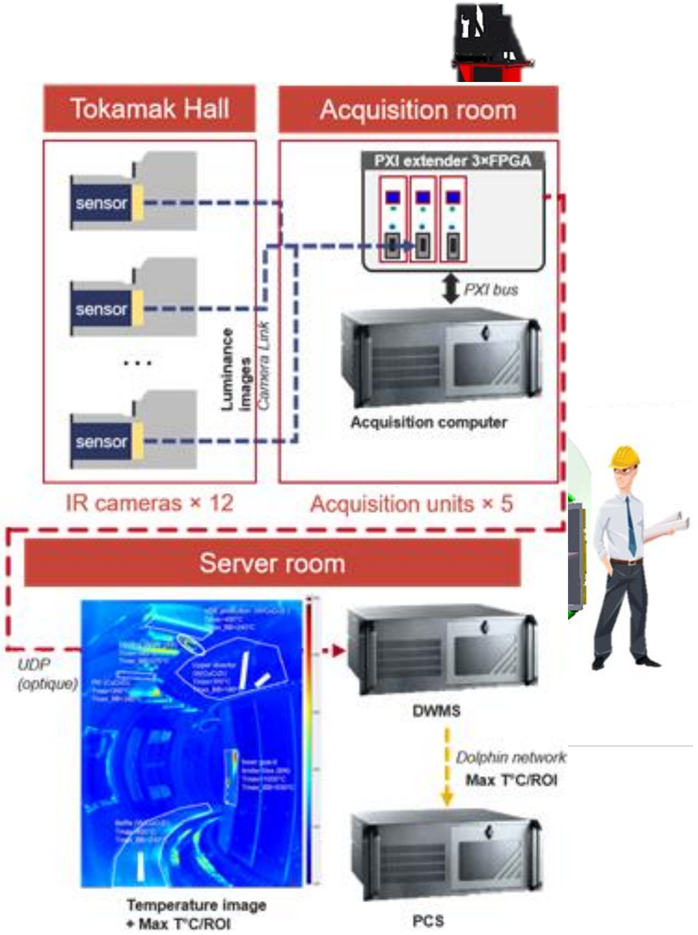

- Current monitoring: temperature thresholds on fixed regions of interest (ROI)
- DWMS (Diagnostic: Wall Monitoring System) equipped with a GPU for real-time inference using AI models

Infrared thermography diagnostic at WEST

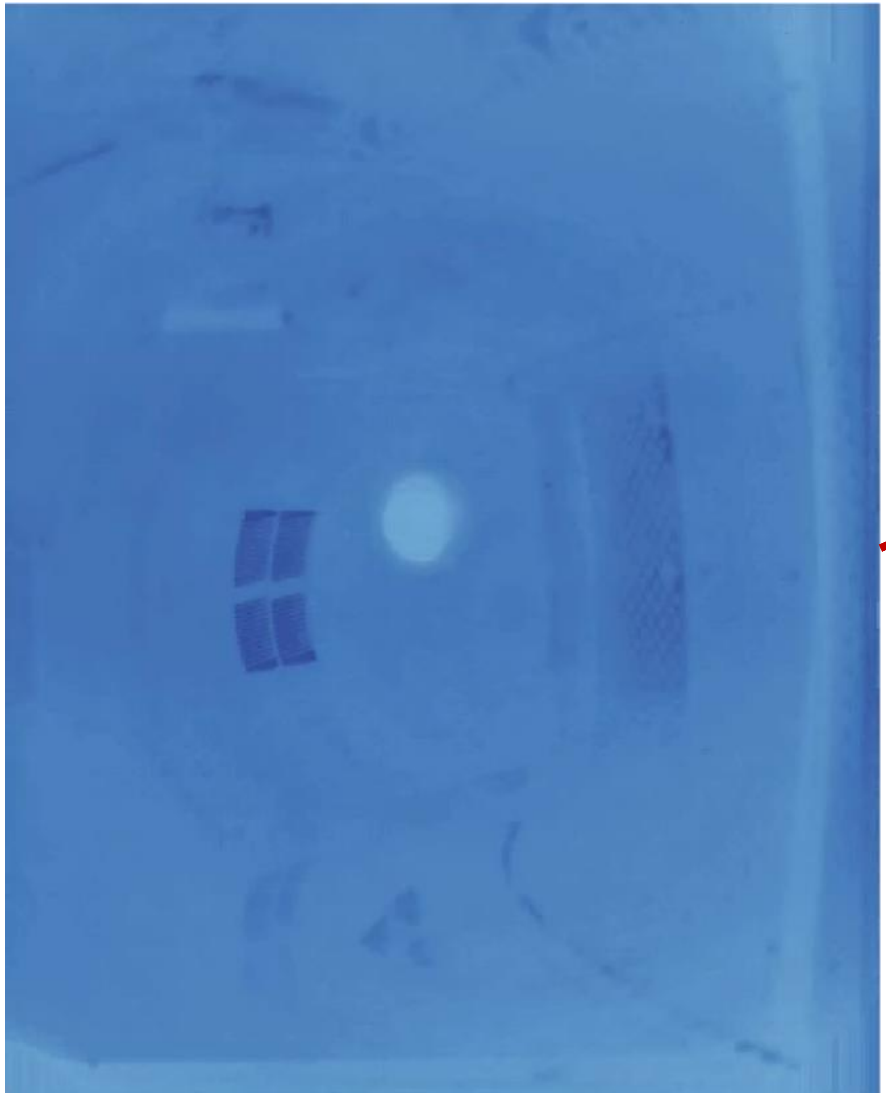

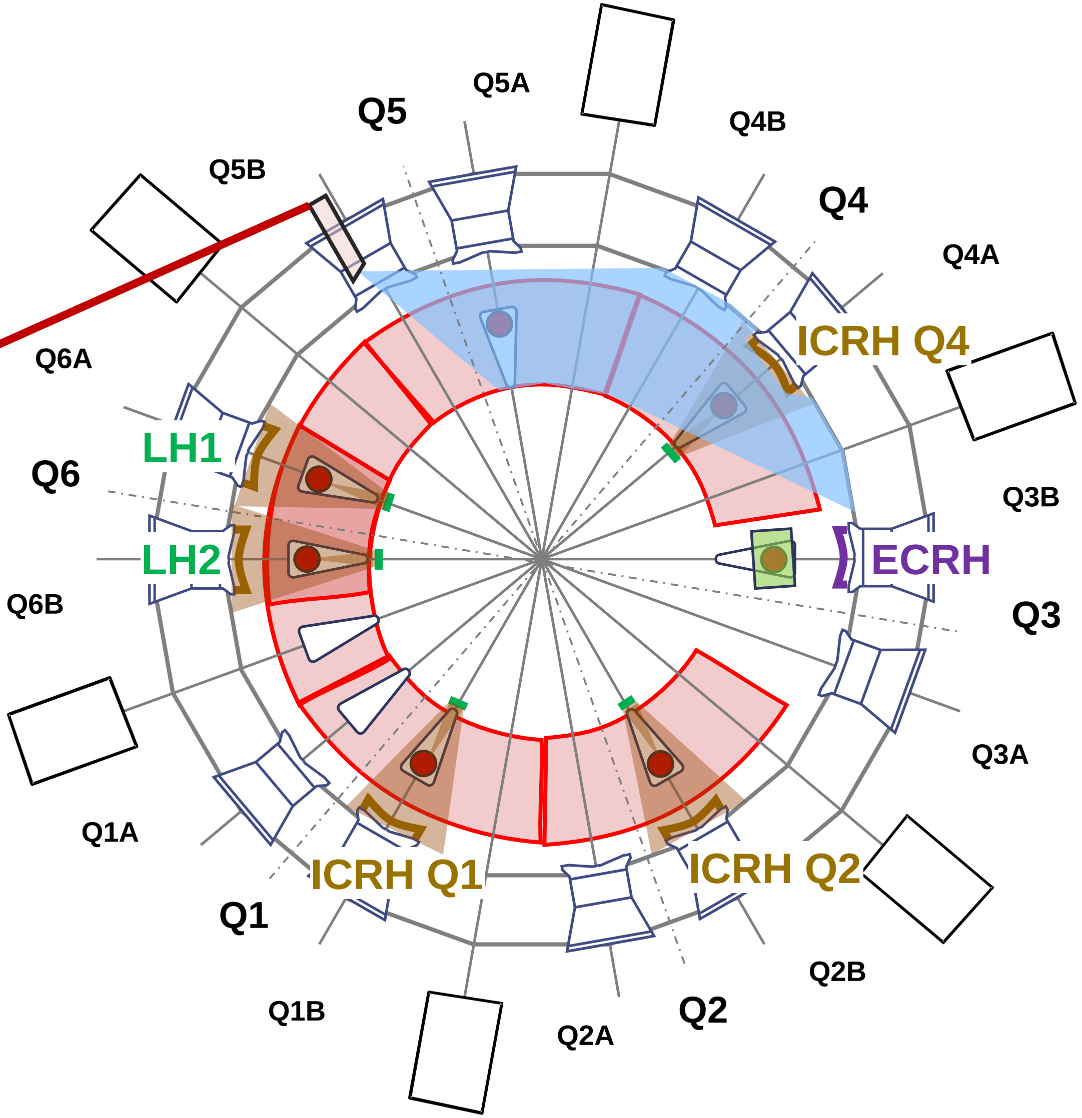

Infrared thermography diagnostic at ITER

- Wide Angle Viewing System (WAVS)
- 20 lines of sight in total (15 Equatorial and 5 Upper)
  - 77% of the chamber (80% of the divertor and 77% of the first wall)
  - MWIR (~3-5µm)
  - at 100Hz (10ms between frames) in full frame
- Spatial resolution
  - 8mm/px for the divertor view
  - 12mm/px for the ICRH antenna
  - 24mm/px for the remaining tangential views

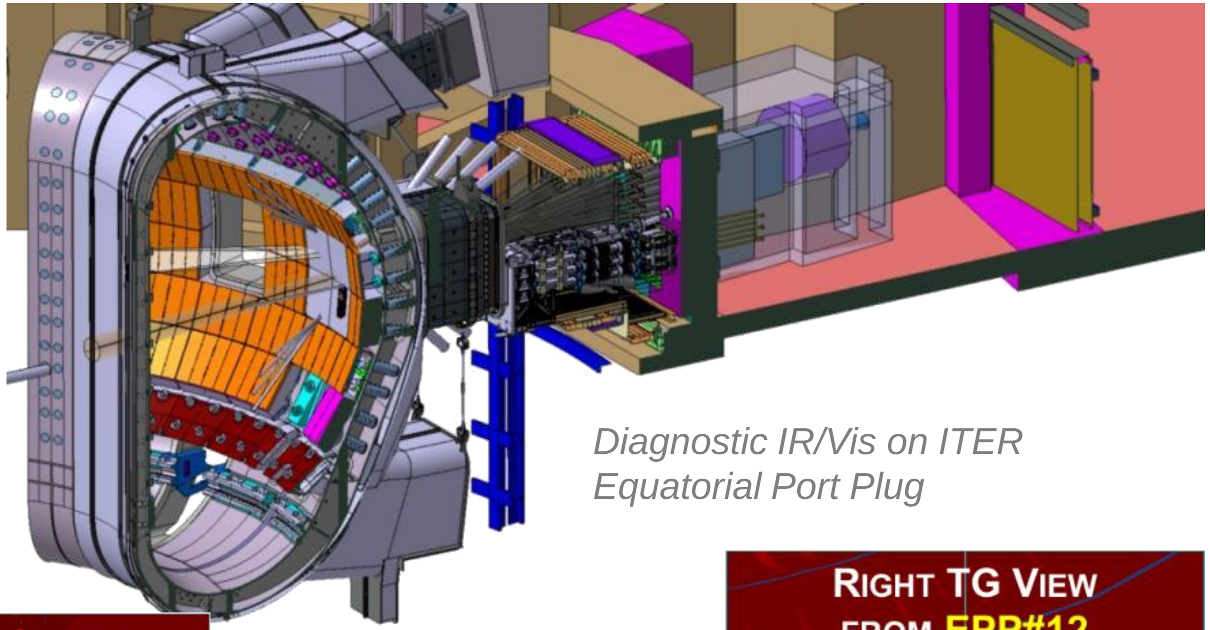

Diagnostic IR/Vis on ITER Equatorial Port Plug

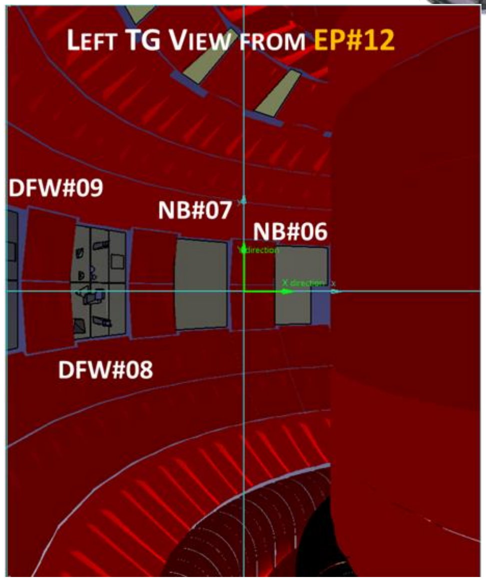

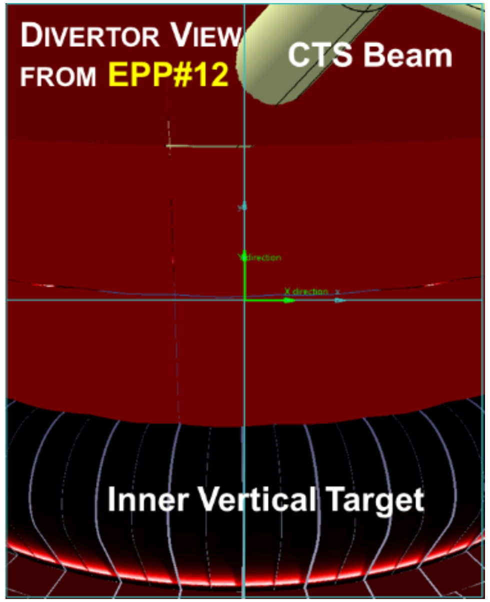

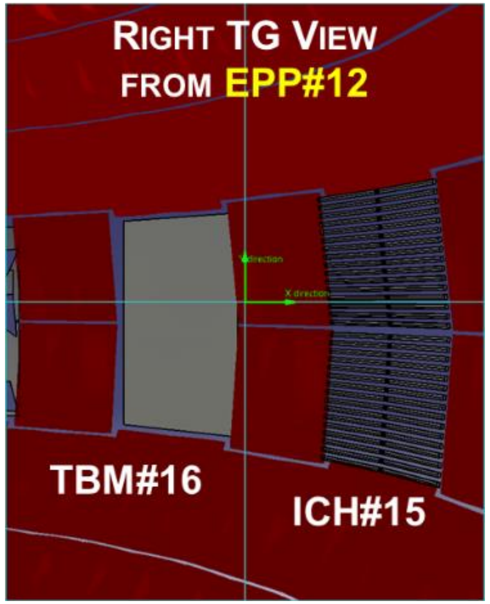

[Visuals from WAVS team]

Fusion machine monitoring and protection

Investment protection + performance optimization + large quantities of data during real-time operation

→ need for the development of an intelligent automated process, part of the Wall Monitoring System (WMS), for the real-time protection of fusion machines based on infrared (IR) imagery

→ at WEST, we are working on the real temperature estimation from IR data (emissivity evaluation) and on methods to clean the IR images from non relevant hot spots and detect thermal events of interest for machine monitoring and protection

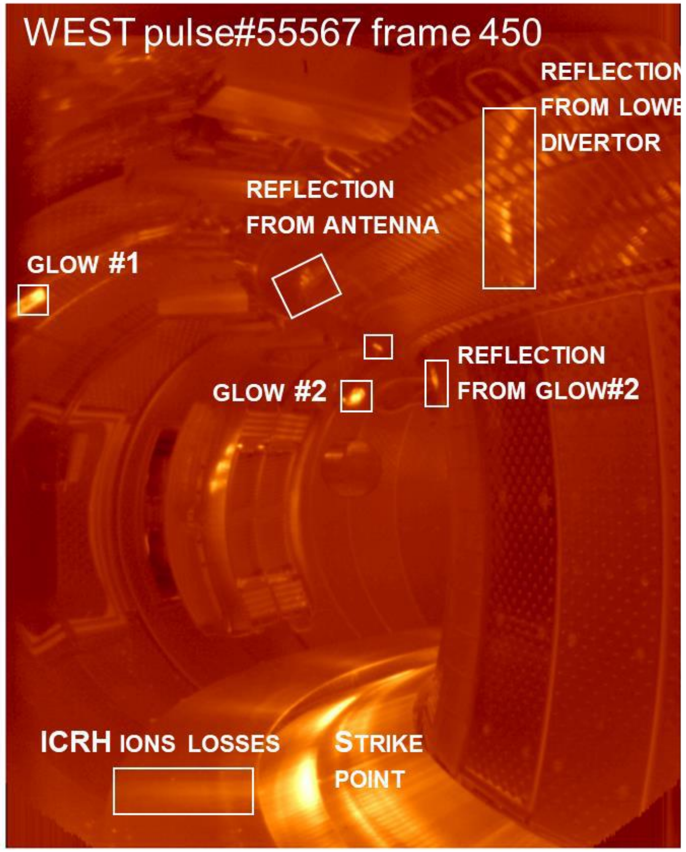

M-H. Aumeunier *et al.*, Nucl. Fusion 64 (2024) 086044

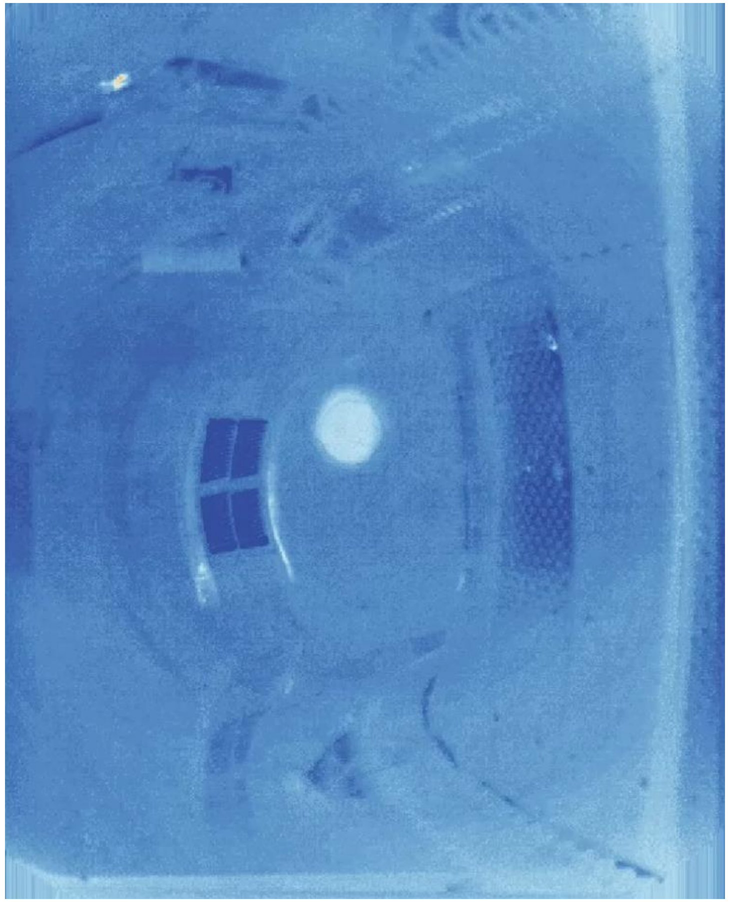

2. Machine protection using infrared thermography

2.1. Real-time protection and feedback

2.2. Post-experiment machine protection

Real-time protection and feedback

DWMS
(Diagnostic: Wall
Monitoring System)

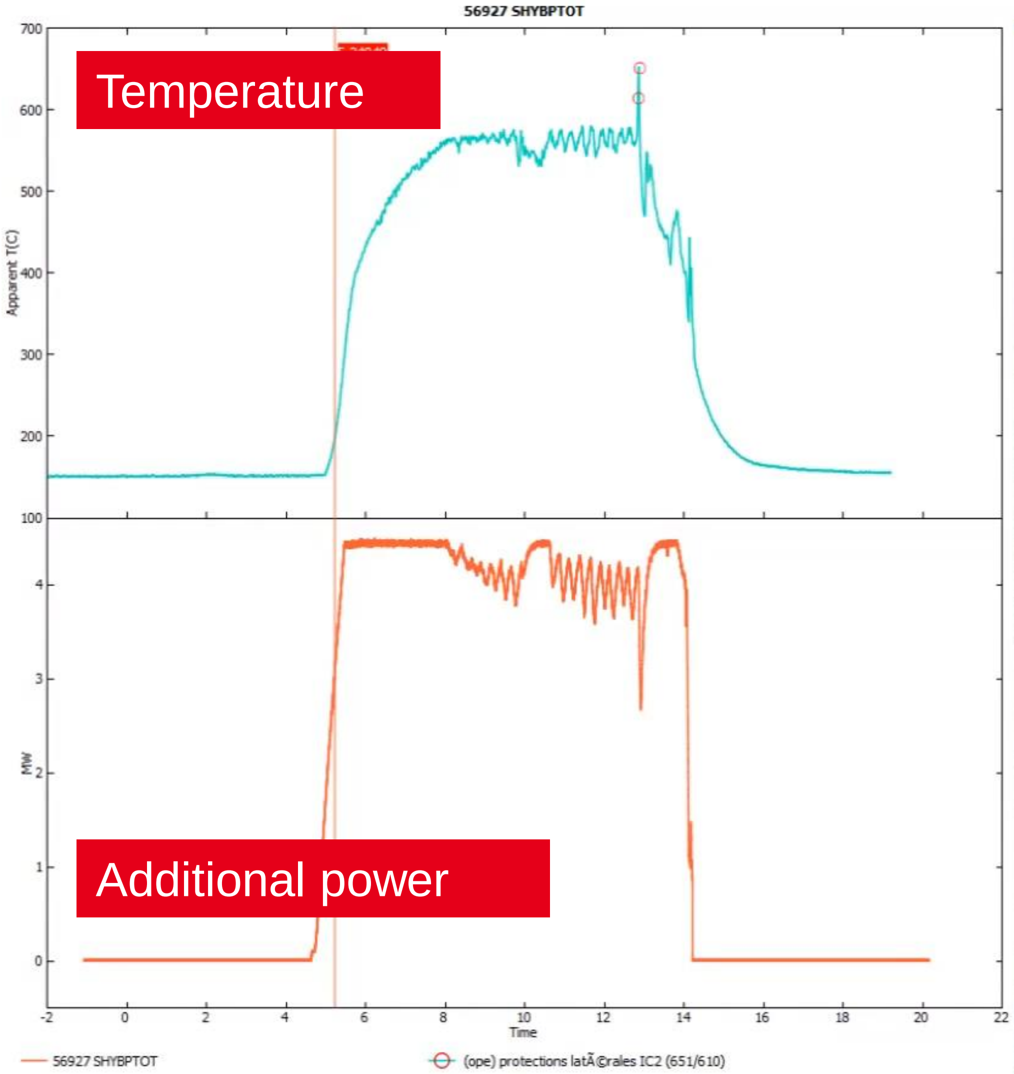

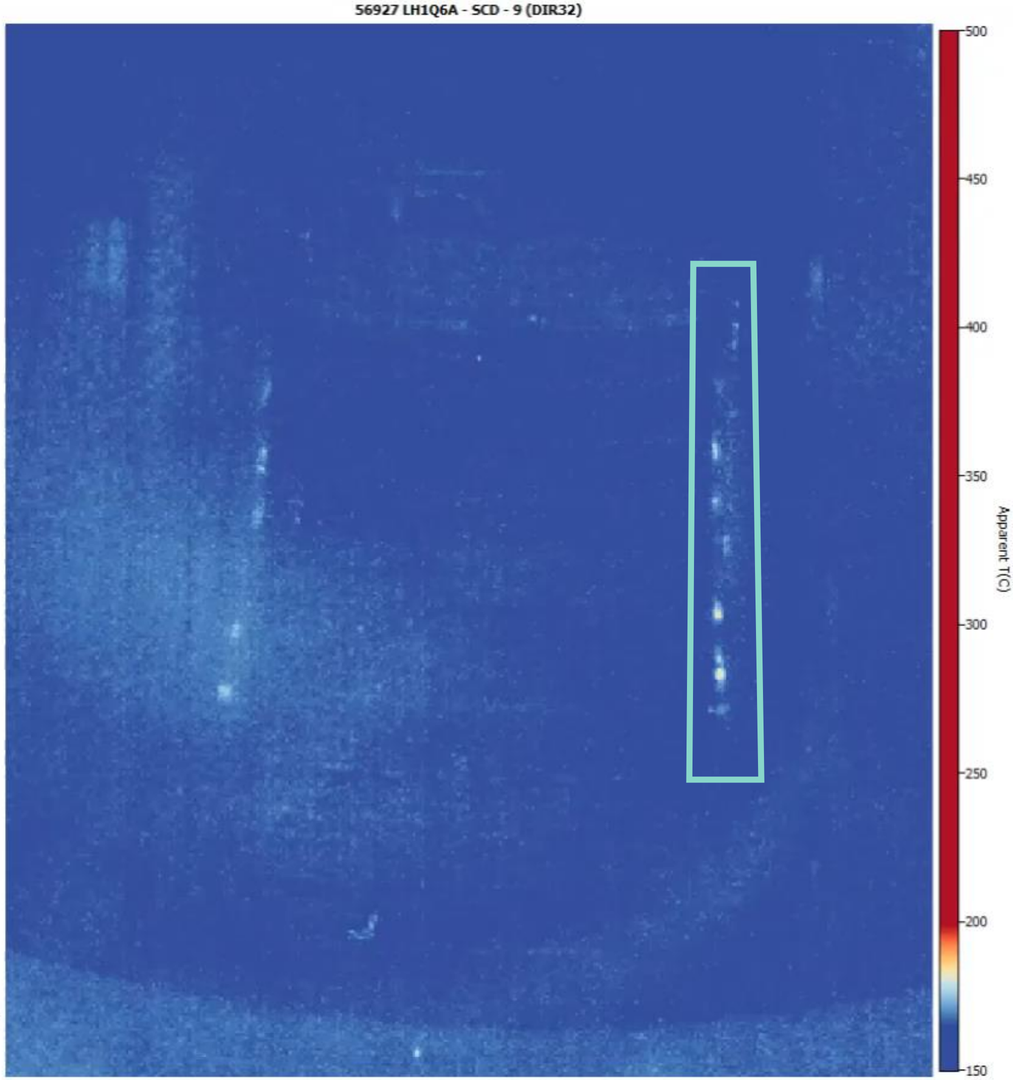

Intelligent real-time protection and feedback

- Role: "Ensure that we stay in the operational domain of the machine"
- Framework that receives and processes the infrared images, and writes the result on the real-time network for feedback control
- All in real-time (less than 20ms)
- With good performance to
  - protect the machine
  - while limiting as much as possible the impact on the operation

Challenges

- Occurrences of false positives (UFOs, reflections, fast electrons…)
- Computational resources

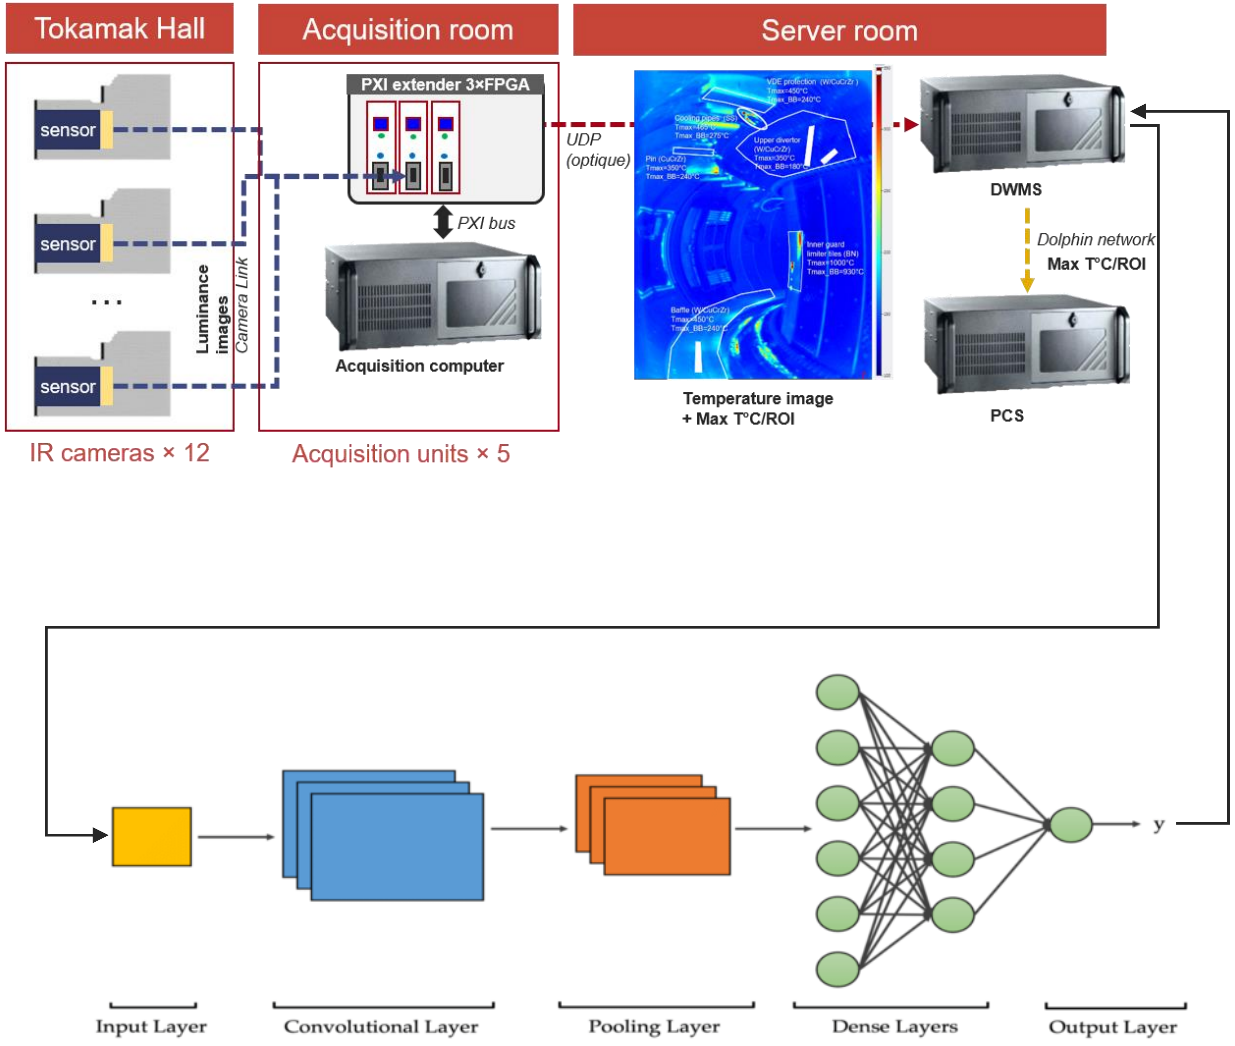

Real-time protection and feedback

What does it mean to protect the reactor while limiting as much as possible the impact on the operation?
How does it impact the choices we make on the model and its parameters?

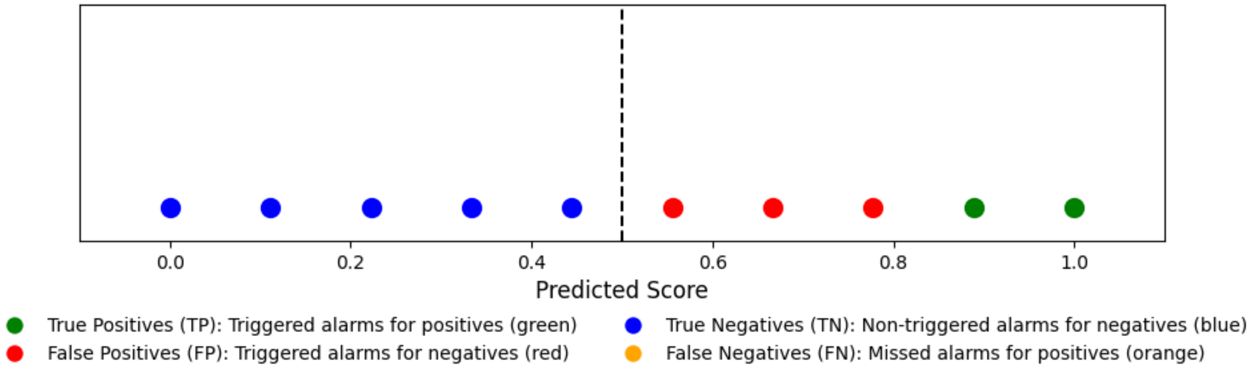

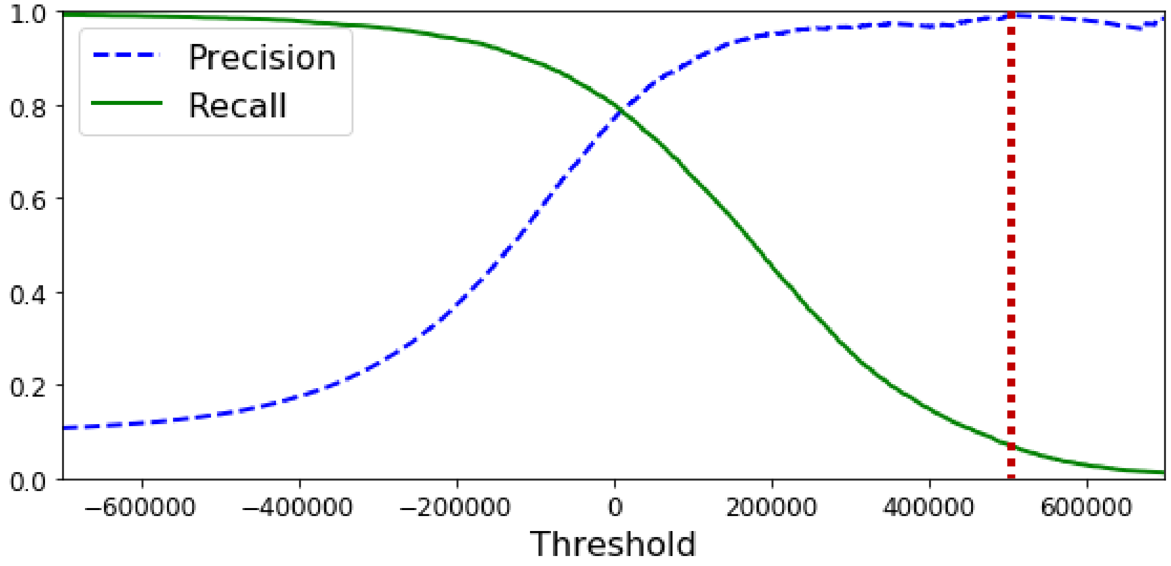

[https://datascience-george.medium.com/the-precision-recall-trade-off-aa295faba140]

Favor recall
“Do not miss important events”

Favor precision
“Do not produce false alarms”

→ Need to find a compromise

Post-experiment machine protection

“Human expertise” branch

- Between pulses (~10min between two pulses)
- Analysis performed by the PFC Protection Officer (PPO), actions taken by the operating team
- Currently: the PPO looks at all the infrared (IR) movies, and performs a post-pulse analysis of the thermal events
- Need for a framework that can access the movies automatically after an experiment, analyze them and save the results in a dedicated database
- Contrary to the real-time branch, recall can be favored compared to precision, to be sure not to “miss” an important detection

Post-experiment machine protection

Challenges

- The analysis can be difficult when there is not much time between pulses (~10min) to analyze the videos of the 10 cameras
- The analysis of a same thermal event can differ between experts
- The results of the automatic detections (from ML models) need to be presented in a clear and understandable way

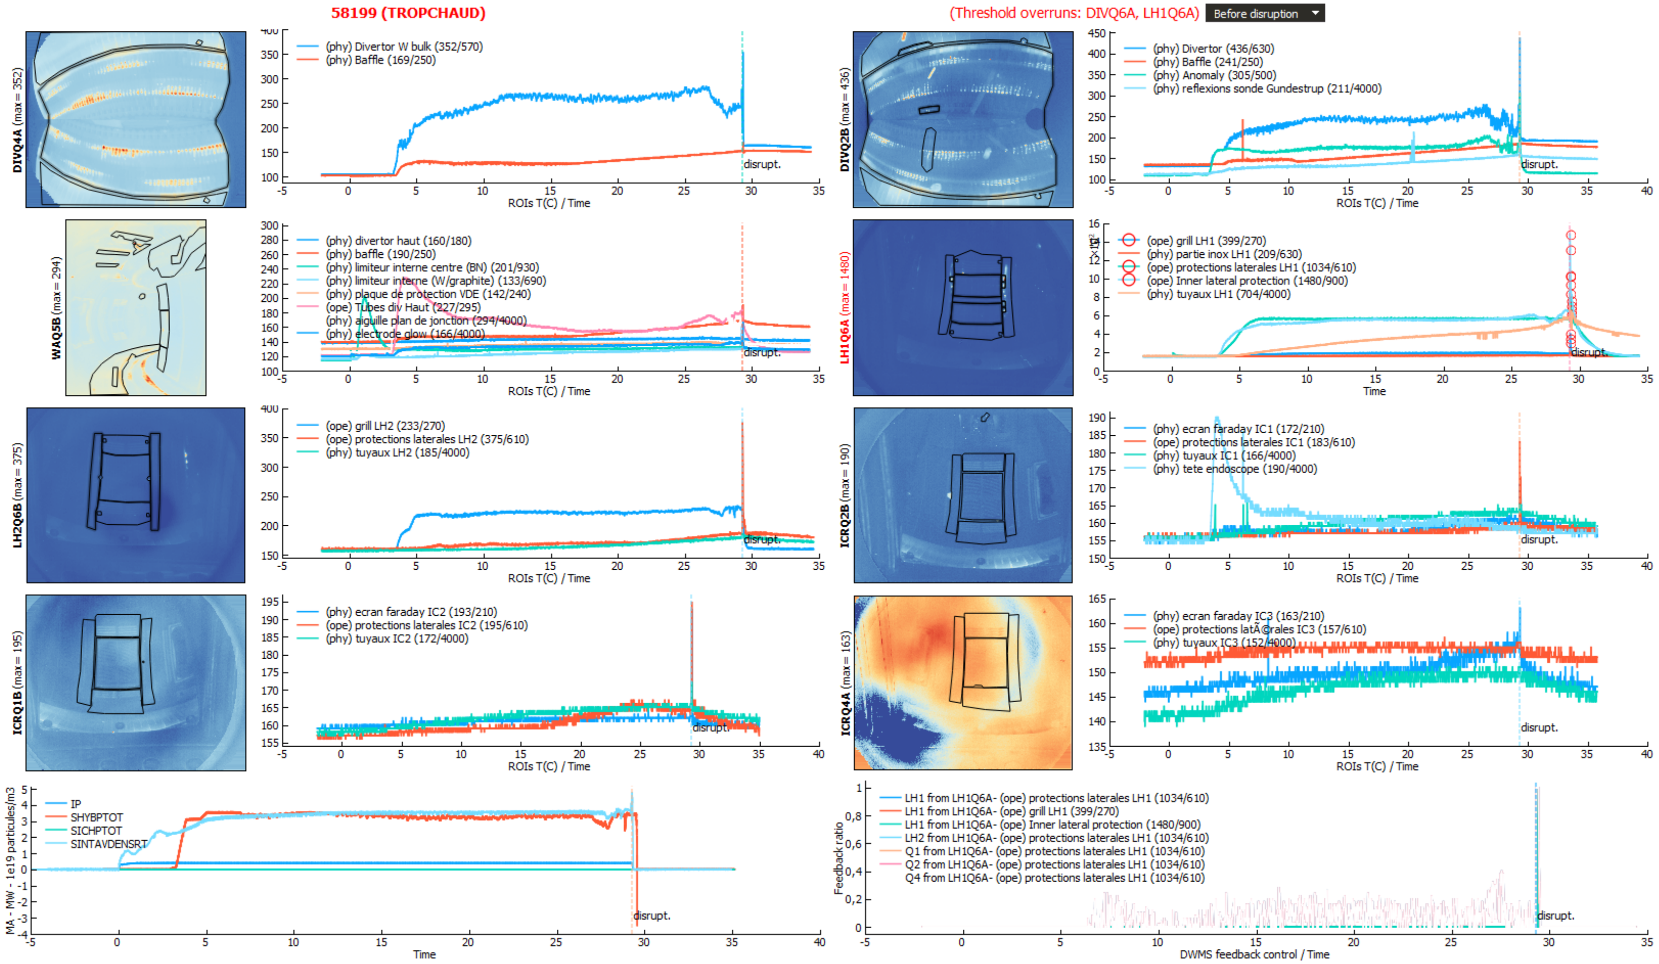

3. Artificial intelligence for machine protection

or the importance of data

Different types of machine learning tasks

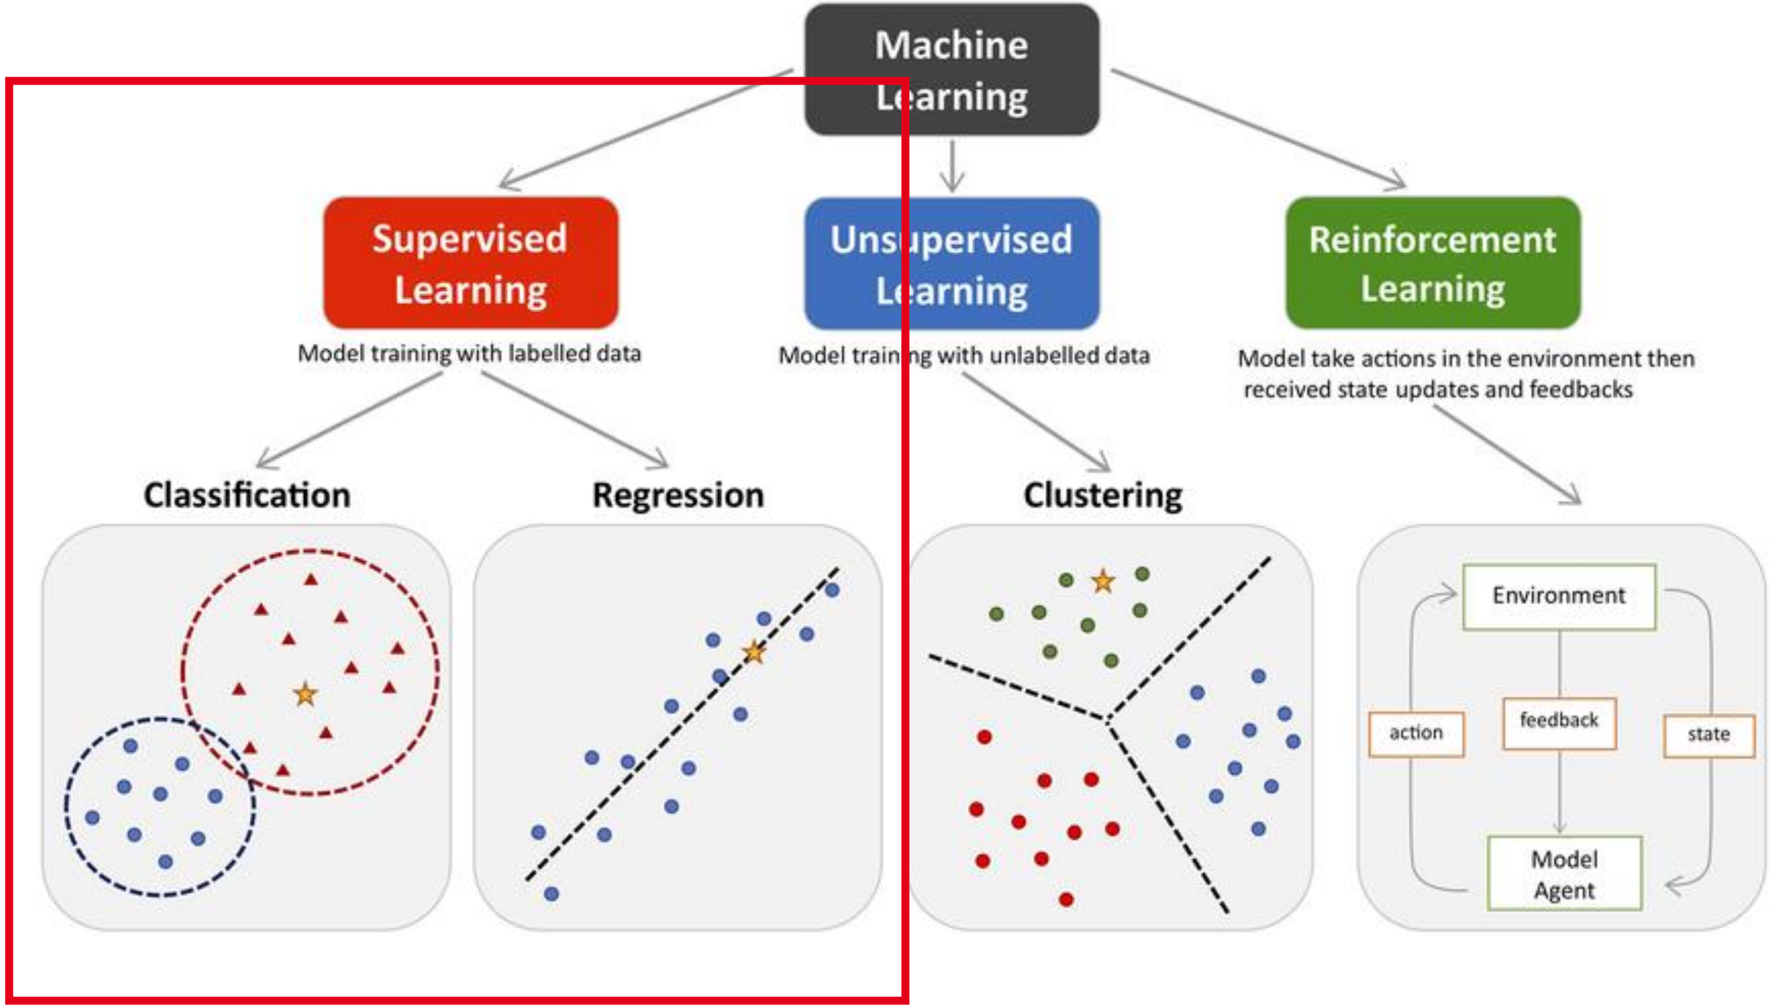

Image from [Peng, Junjie & Jury, Elizabeth & Dönnes, Pierre & Ciurtin, Coziana. (2021). Machine Learning Techniques for Personalised Medicine Approaches in Immune-Mediated Chronic Inflammatory Diseases: Applications and Challenges. Frontiers in Pharmacology. 12. 10.3389/fphar.2021.720694.]

Workflow of ML applied to a given problem

- Define the problem: what do the “clients” want our model to predict? Which data is available as inputs? Is ML the right approach? What are the outputs of the model? What does it imply on the dataset creation step? Which speed should the model attain?
- Create the database: create the datasets required to train, validate and evaluate the model, using annotation tools dedicated to the task type (classification, detection, segmentation, tracking…)

- Choose a model architecture adapted to the task and data
- Train the model on the training set
- Use the validation set to select the proper hyperparameters (model architecture, learning rate, stopping iteration) → requires the definition of metrics understandable by the clients
- Evaluate the model’s performance with the test set → estimated level of performance on new, similar experiments; ensure safety, compliance, and clients trust
- Integrate and deploy the model to use it during or between experiments
- Optionally, use the results of inferences to improve the model (active learning or human-in-the-loop learning)

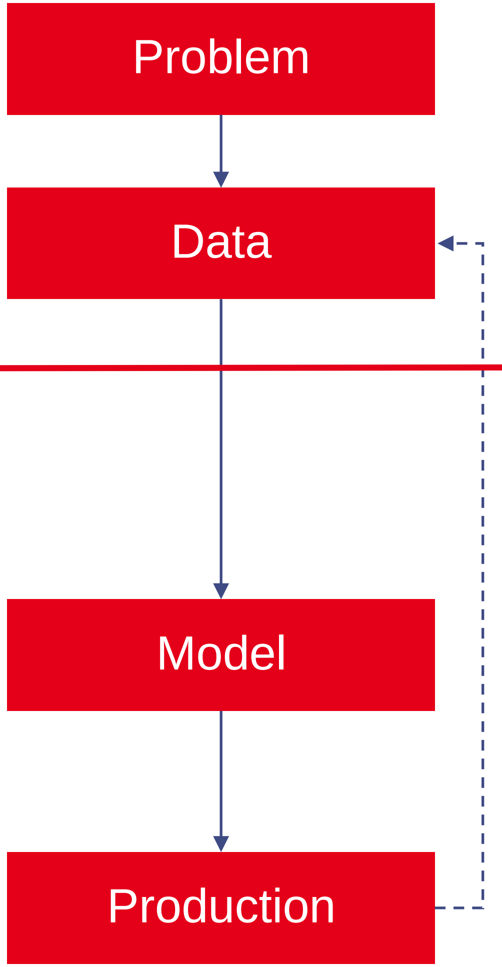

Data: natural vs infrared images

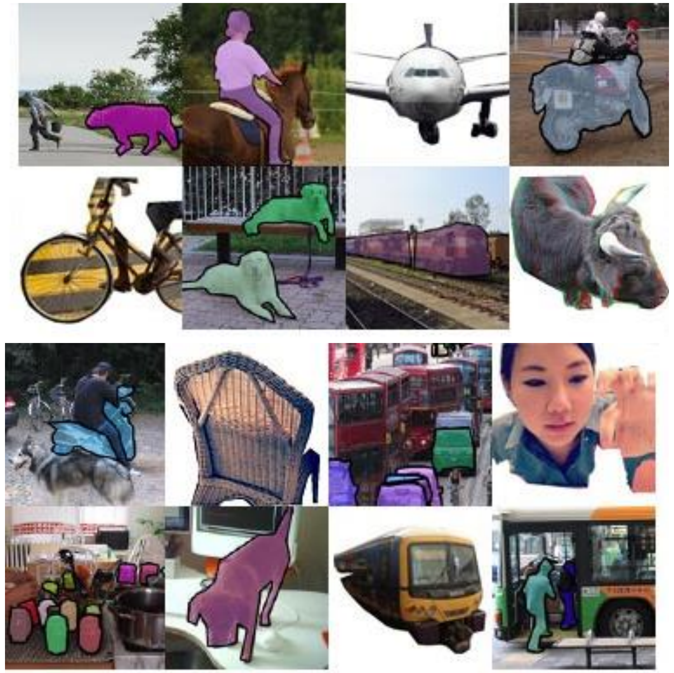

[https://cocodataset.org]

- Usually images
- Color (3 channels)
- Pixel value may not be important (luminance)
- Easy to have diversity (many diverse images on the internet)
- Easy to annotate (scenes or objects from everyday life)

- Videos
- Gray scale (1 channel)
- Pixel value is important (apparent temperature)
- Diversity hard to attain (requires several machines, configurations and scenarios)
- Hard to annotate (requires knowledge of physics and of the machine)

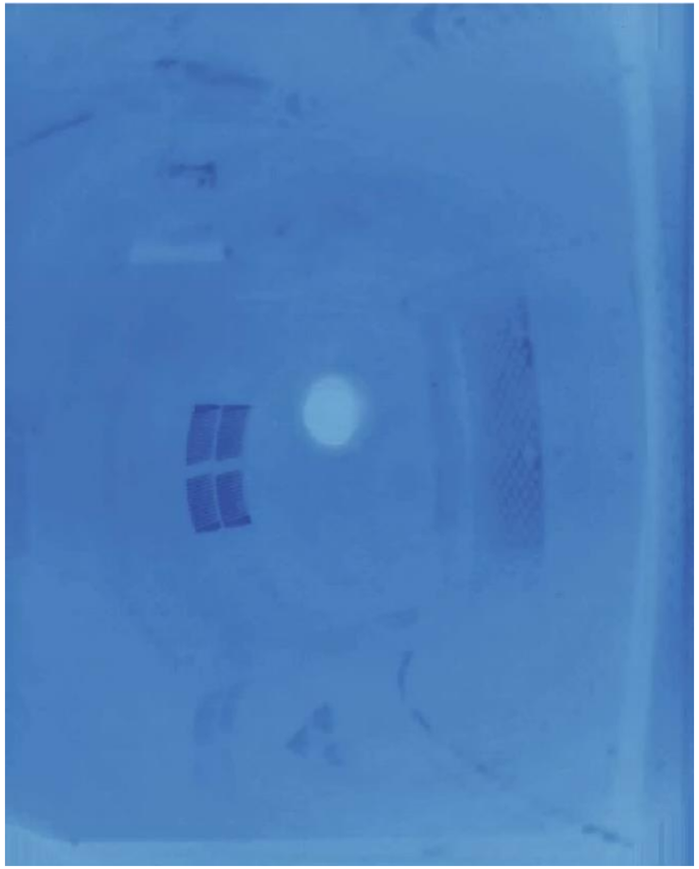

→ Existing datasets relying on natural images may not help for machine protection with infrared data
→ Existing methods and models must be adapted to receive infrared images

Optimizing IR data diversity in datasets

1. **Annotation budget constraint**
   The available budget often limits the total number of images that can be annotated. This restriction may lead to over-representation of certain patterns and under-representation of edge cases

   Consequences:

   - Poor generalization to unseen data
   - Bias in model predictions, impacting robustness
   - Reduced ability to handle edge cases or rare events

2. **How to maximize diversity?**
   Ensuring diverse coverage of movie features (machine configuration, PFC state, plasma scenario, *etc.*) can be modeled as filling a high-dimensional hypercube space. Each dimension represents a characteristic

   → Selecting representative data points within the annotation budget requires balancing uniform coverage and emphasis on rare or critical features

Optimizing IR data diversity in datasets

3. **Shift from whole movies to sequences**

- **Initial approach**:
  Annotating entire movies resulted in redundancy, as many frames are visually similar and add almost no information
- **Current approach**:
  Focus on sequences or individual frames that maximize feature diversity

  → Reduced annotation redundancy, increased diversity

4. **Strategies for diversity optimization**

- **Active learning**:
  Prioritize labeling images that maximize diversity in feature space
- **Cluster-based sampling**:
  Identify and sample from underrepresented clusters in the feature space

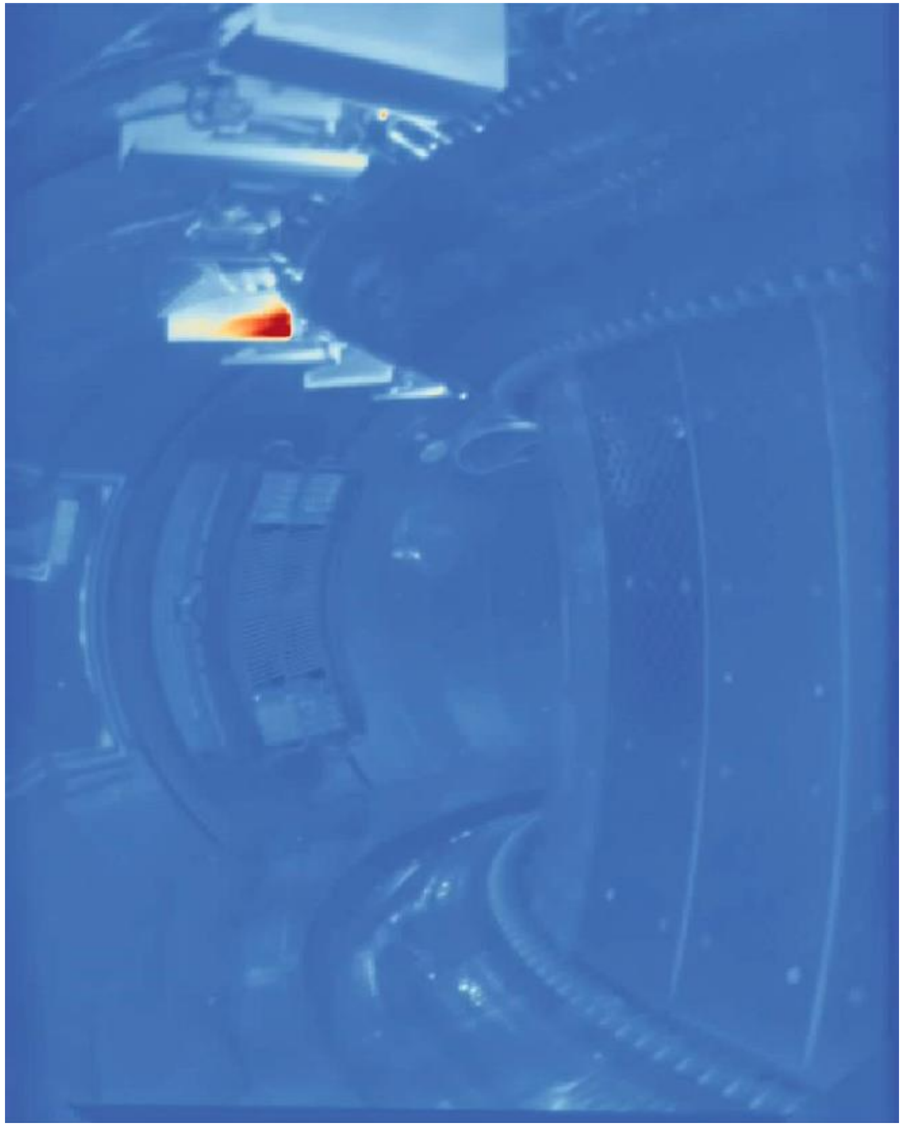

Difficulty of performance quantification

You (the experts)

- Classical ML metrics like
  - Precision
  - Recall
  - F1-score
  - Accuracy
- Task-specific metrics like
  - Mean Average Precision (detection/segmentation)
  - Tversky loss (imbalanced segmentation)
- Domain-specific metrics like
  - Number of disruptions prevented
  - Number of temperature overrun prevented

…

The stakeholders

- Performance
- Risk
- Money
- Time

It is crucial to be able to work with both experts metrics, and metrics that can be understood by the clients and the stakeholders, to make the proper choices during conception and ensure trust in the provided solution

4. ML-based methods implemented at WEST

A quick overview

Reminder: ML workflow

- Define the problem, assess if it is solvable and decide if ML is the way to go
- Create the datasets
- Choose a model architecture
- Train the model on the training set
- Use the validation set to select the proper hyperparameters
- Evaluate the model's performance with the test set
- Deploy the model

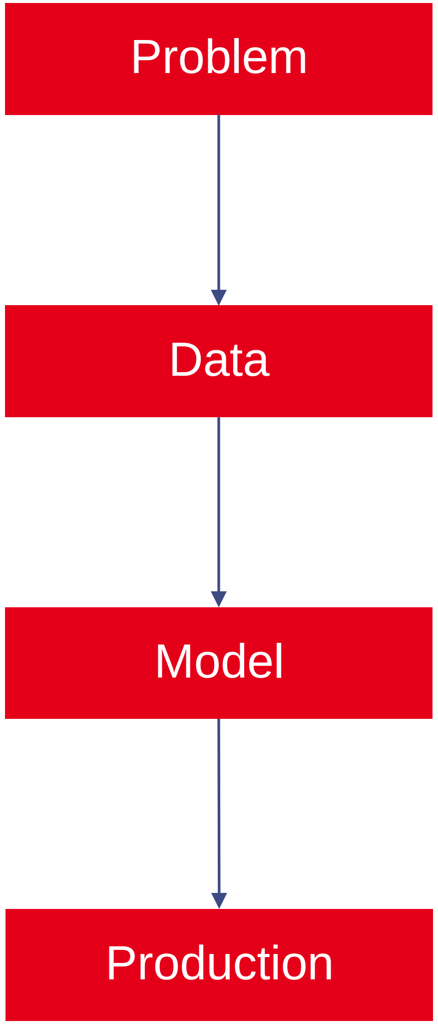

ML-based methods and their use on WEST

- Various ML-based methods are developed on WEST data:
  - Thermal event detection
  - Electric arc detection
  - Strike line description
  - UFO detection
  - Runaway electrons detection
- Deployed on DWMS for (near) real-time use or on a server for between pulse computation
- Tested during the C9 campaign (Dec. 23 – Apr. 24, ~1300 pulses) and ongoing C10 campaign
- Other methods are being/will be developed:
  - Temperature inversion
  - Synthetic infrared image generation

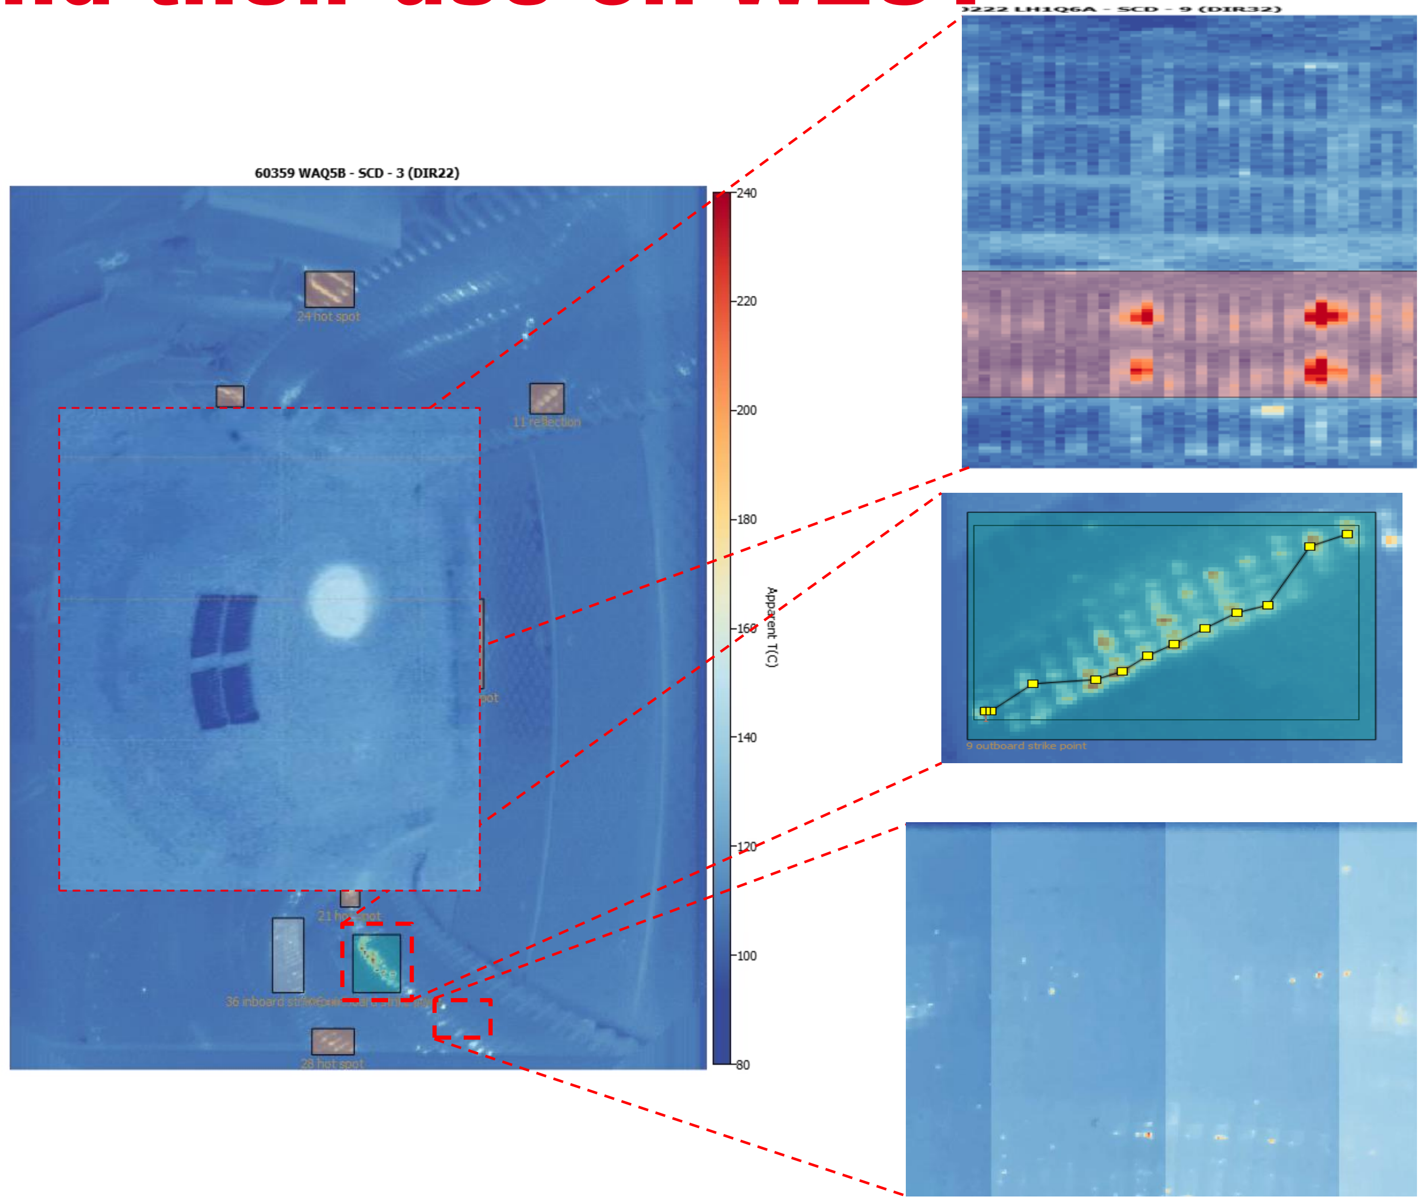

Electric arc detection

- **Problem:** detect electric arcs on the grill of lower hybrid heating antennae in real-time during experiments
  - Classification problem (electric arc yes/no)
  - Real-time constraint (50Hz), with interfacing with the feedback control loop
  - IR images as inputs
  - Imbalanced problem (very low occurrence of electric arcs)
  - Available resources: computer with a CPU, connected to the acquisition system

- **Datasets:**
  - Identification of experiments with and without electric arcs
  - Manual annotations of IR images from identified movies
  - 12595 images from 115 movies annotated using a functionality specifically developed in the ThermaVIP software

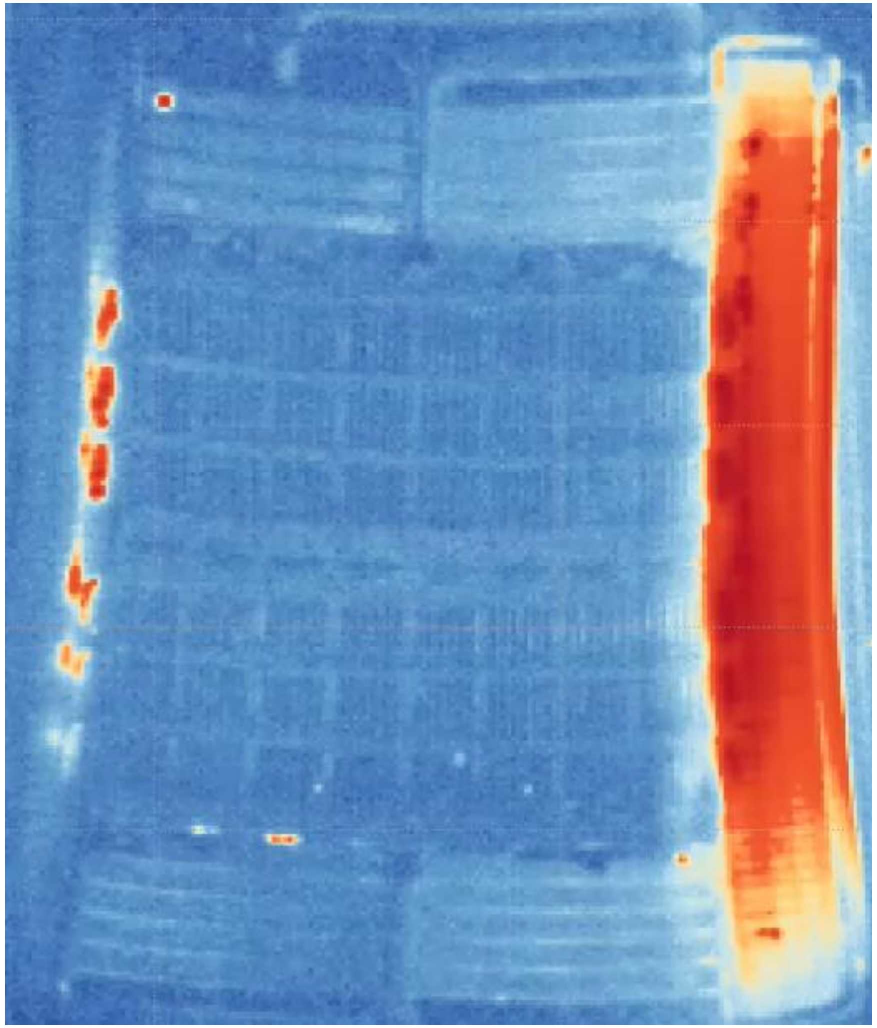

Electric arc detection

- **Model:** convolutional neural network (~104k parameters, small enough to run on CPU at > 50Hz)
- **Performance:**
  - Mean inference time on DWMS (pure CPU): ~8ms
  - Metrics on test set (2226 images, detection threshold 0.91)
    - Balanced accuracy: 98.86%
    - Precision: 100.00%
    - Recall : 97.73%
    - F1-score: 98.85%

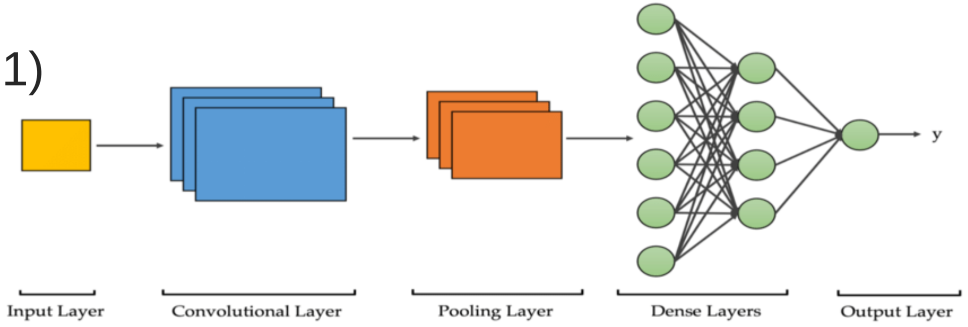

- After discussion with the heating system operators, precision has been favored over recall, to minimize the risk of false positives, by adopting a very high detection threshold of 0.998
- **Deployment:**
  - on DWMS for real-time use during operation, with a feedback control on the additional power injected
  - dedicated dashboard for post-experiment analysis

V. Gorse et al. “AI-based Techniques for Real-Time Detection and Characterization of Thermal Events: First Results on WEST Tokamak”. *To be published*

Thermal events detection

- **Problem:** on all the lines of sight, detect various thermal events, classify and track them
  - Object detection and multiple object tracking problems
  - To be run right after a pulse, but with the goal of real-time detection → model must be real-time compatible
  - Infrared movies as inputs
  - Available resources: server equipped with multiple GPUs
- **Datasets:**
  - Identification of infrared movies as diverse as possible: different lines of sight, plasma scenarios and machine configuration
  - Manual annotation (bounding box + classification + tracking) of the thermal events in the movies
  - 123740 images from 118 movies annotated using functionalities specifically developed in the ThermaVIP software

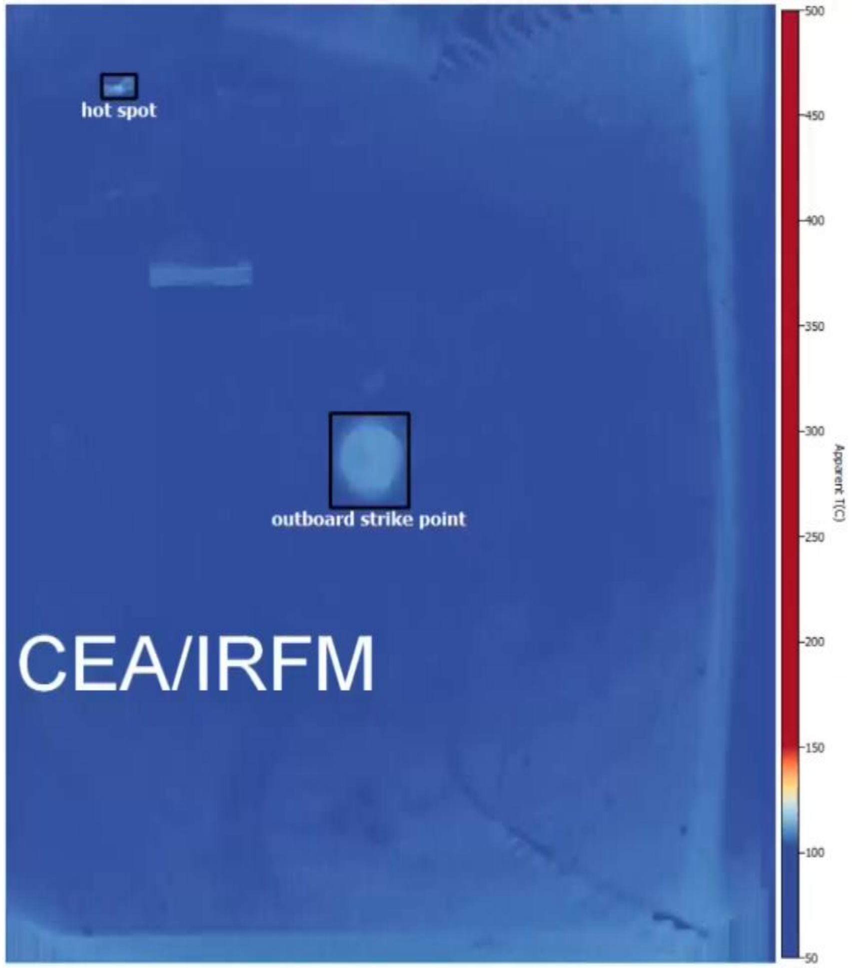

Thermal events detection

- **Models:**
  - Faster R-CNN for the detection (~28.3M parameters)
  - SORT for the tracking (no parameters)
- **Performance on test set:**

Camera type	Recall@0.5	Precision@0.5	AP@0.5
Divertor	40.46%	37.71%	32.68%
LH antenna	52.33%	75.66%	50.14%
Wide-angle	75.85%	47.38%	59.98%
All	55.87%	60.75%	50.82%

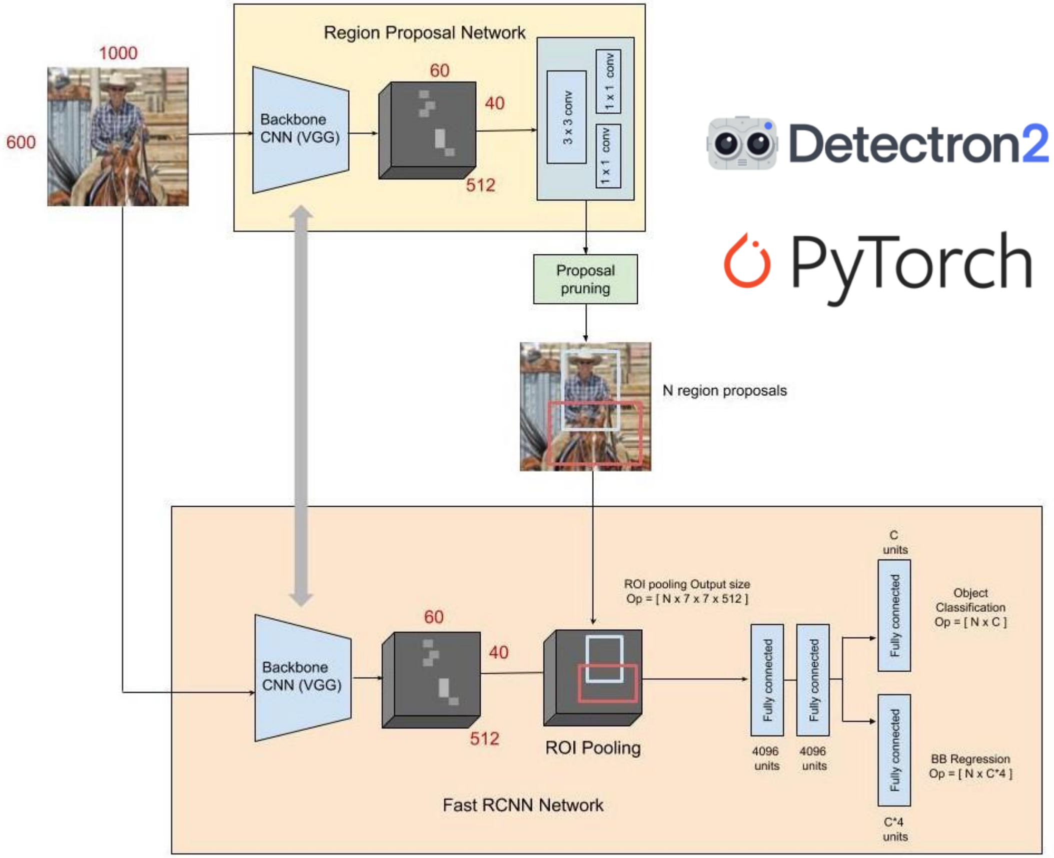

Detectron2

PyTorch

- **Deployment:**
  - Automatic post-experiment inference on all the lines of sight
  - Demonstrator of near real-time inference (~30fps vs 50Hz) deployed on one line of sight

S. Ren et al. "Faster r-cnn : Towards real-time object detection with region proposal networks". *arXiv:1506.01497 [cs]*. http://arxiv.org/abs/1506.01497

E. Grelier et al. "Deep Learning-Based Process for the Automatic Detection, Tracking, and Classification of Thermal Events on the in-Vessel Components of Fusion Reactors". *Fusion Engineering and Design* 192: 113636. https://doi.org/10.1016/j.fusengdes.2023.113636

Strike line automatic description

- **Task:** automatic strike line description (angle + curvature + skeleton)
  - Regression + segmentation problem
  - No real-time constraint, but should be quick since it uses the results of the previous method
  - IR images cropped on strike lines as inputs
  - Available resources: server equipped with multiple GPUs
- **Datasets:**
  - Synthetic images of strike lines generated using CAST3M
  - Angle, curvature and skeleton automatically generated
  - Enabled the creation of 32400 images of diverse strike lines without any manual annotation

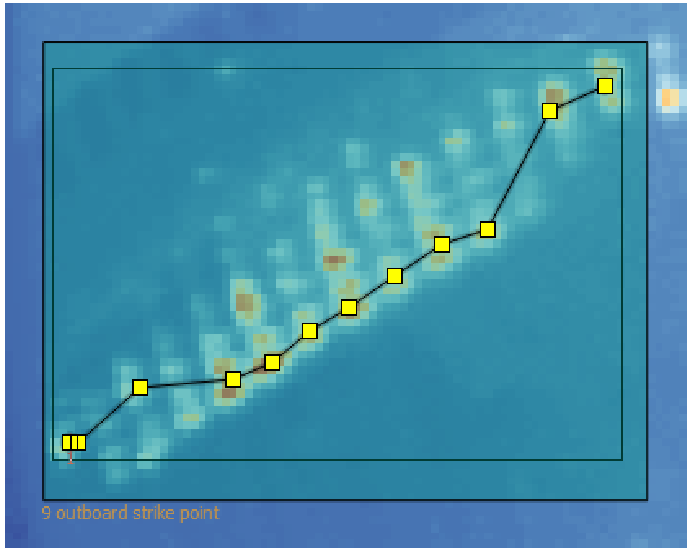

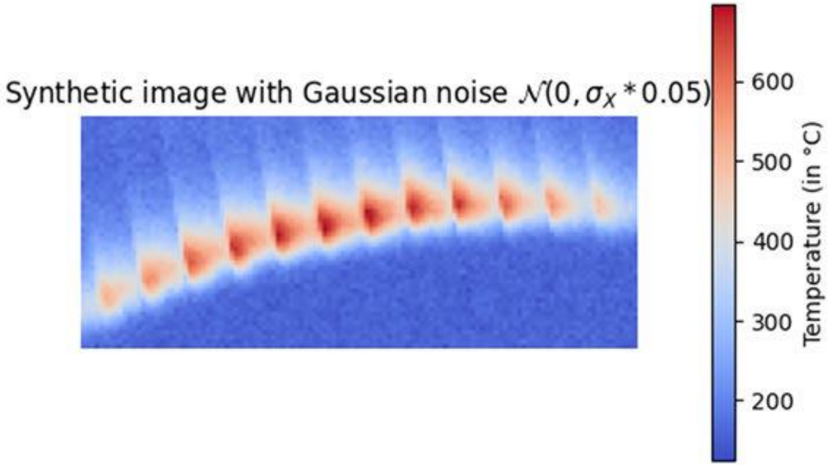

Strike line automatic description

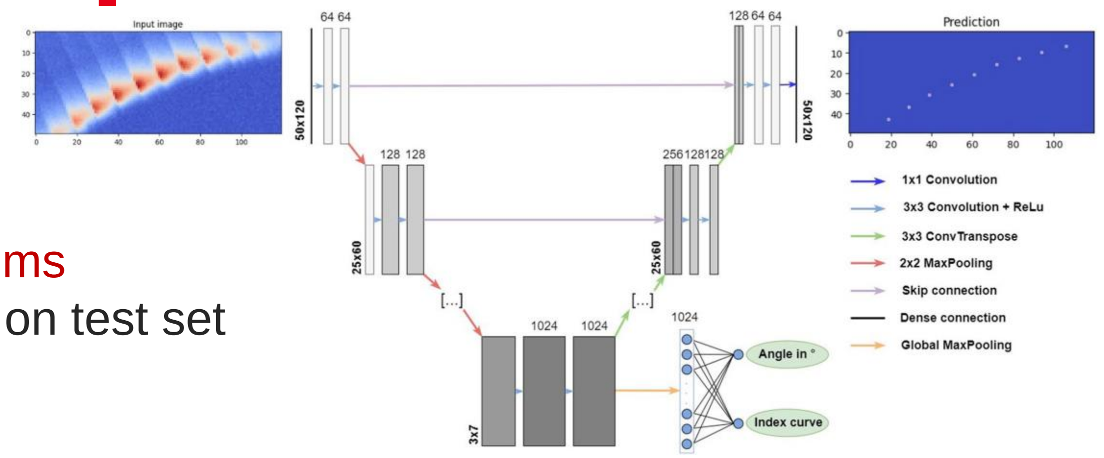

- **Model:** physics constrained U-net (~34M parameters)

- **Performance:**
  - Mean inference time on DWMS (GPU + CPU): ~3ms
  - Metrics (regression and segmentation) computed on test set

- **Deployment:**
  - Can be run on demand either during or between experiments
  - Example of use during C9 campaign (early 2024)

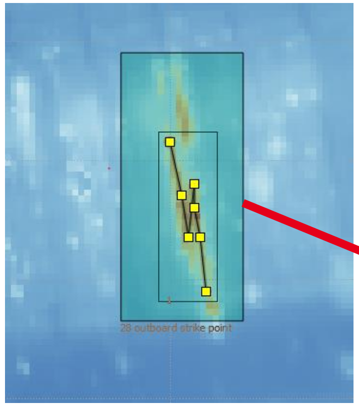

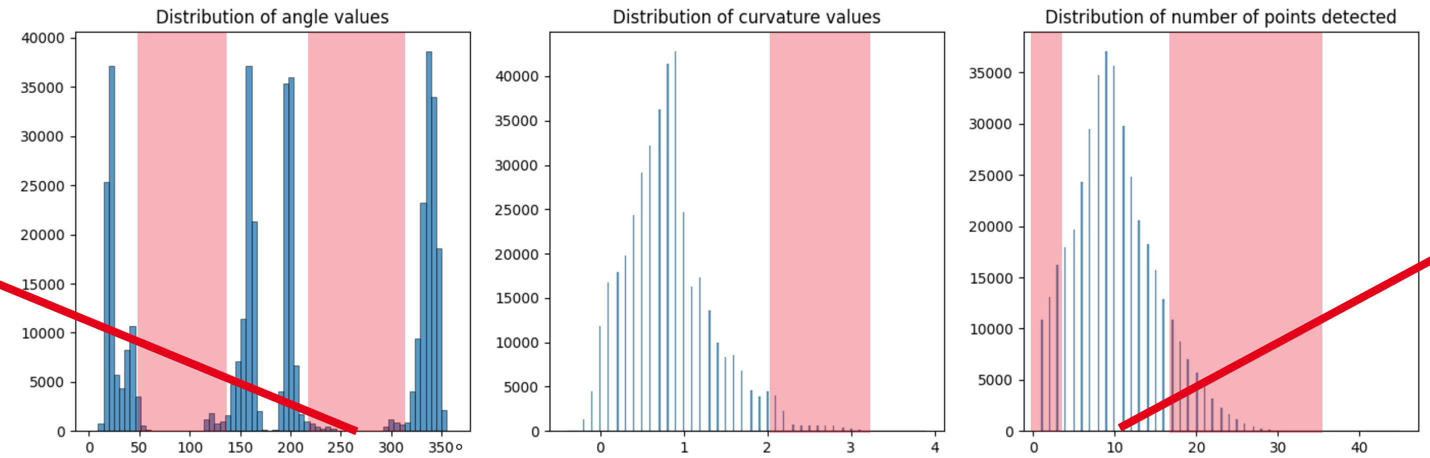

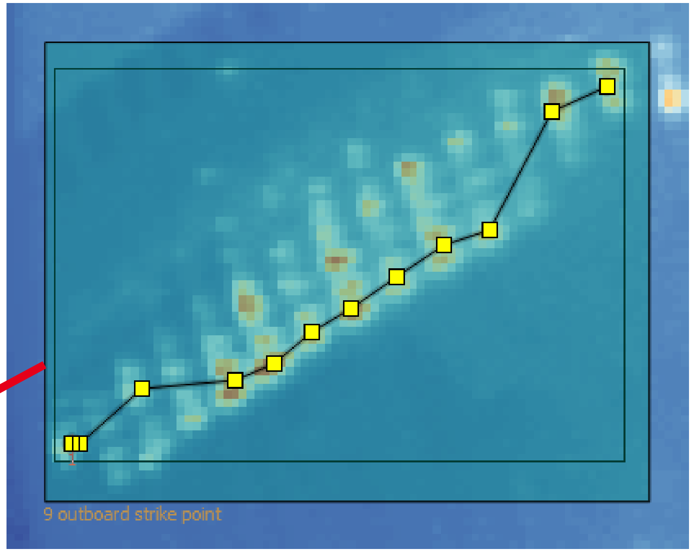

V. Gorse et al. “Using a Physics Constrained U-Net for Real-Time Compatible Extraction of Physical Features from WEST Divertor Hot-Spots”. *Journal of Fusion Energy* 43, nº 1: 13. https://doi.org/10.1007/s10894-024-00405-y

Runaway electrons detection

- **Problem:** detect the synchrotron radiation characteristic of the presence of runaway electrons
- **Model:** same architecture used for electric arc detection
- **Deployment:**
  - Inter-experiment automatic analysis
  - Dedicated web interface for visualization

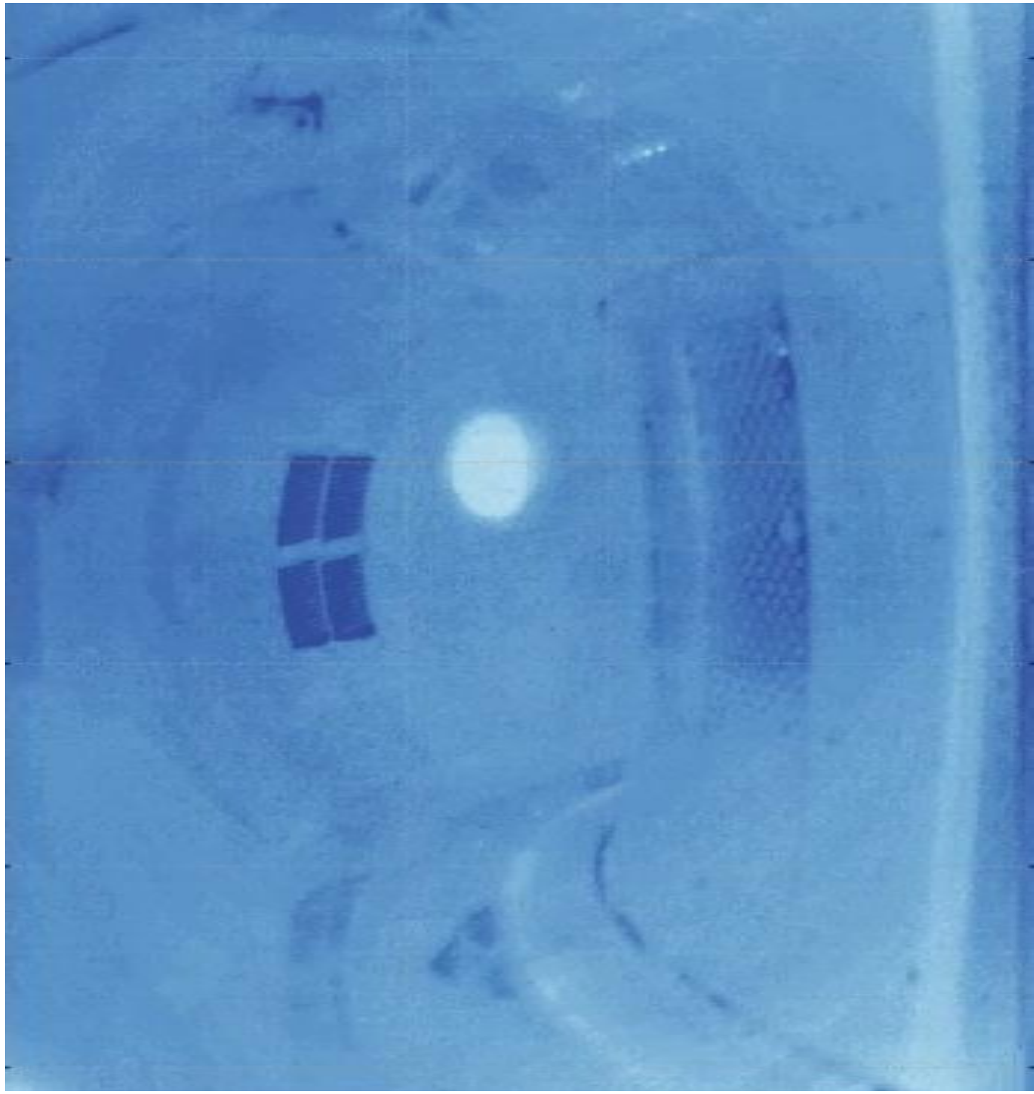

Building blocks developed for other problems:
- dataset creation,
- model choice, training and evaluation,
- deployment for automatic inter-experiment analysis with a dedicated dashboard

5. Extension to other fusion machines

W7-X, JT60-SA, ITER…

Extension to ITER: requirements

"
1. Plasma Control based on quantitative image analysis:
    - Basic image processing for e.g. maximum signal extraction and localization
    - Fast system response (~1 ms)
2. Plasma Control based on qualitative image analysis:
    - Advanced image processing (i.e. phenomenological analysis)
    - Slower system response (>10 ms) that can be used for e.g. false alarms mitigation raised by fast – but basic – algorithms.

Ideally, the two strategies should be mixed.
"

System Design Description (DDD) 55.G1 Vis/IR Equatorial Port Wide-Angle Viewing System

→ ITER will require an advanced protection system based on infrared data from day one of operation

→ It is necessary to develop cross-machines models now, using data from existing fusion machines

Diversity between fusion machines

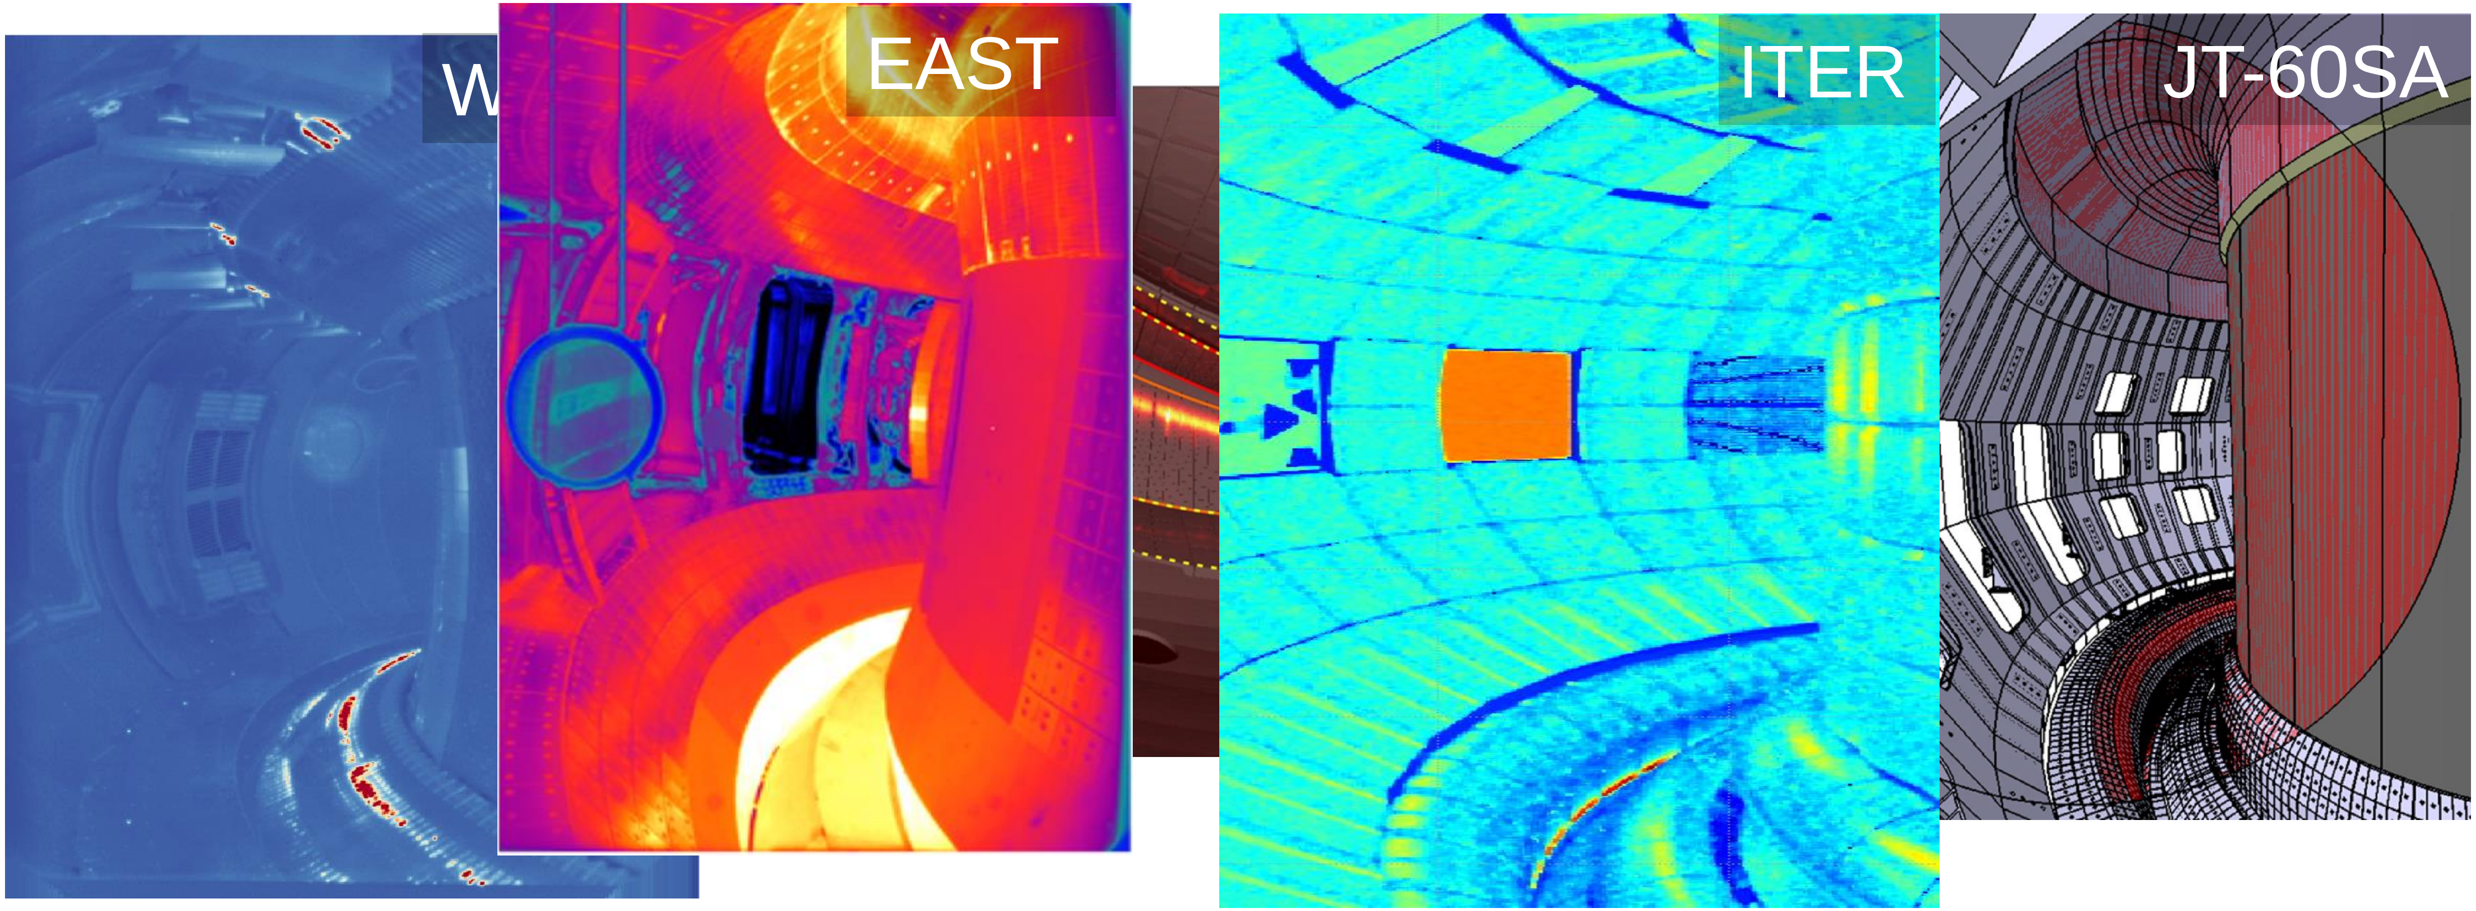

Sources for the images:

- [Jakubowski, M. *et al*. (2018). Infrared imaging systems for wall protection in the W7-X stellarator (Invited). Review of Scientific Instruments, 89(10), 10E116. https://doi.org/10.1063/1.5038634]
- [Aumeunier, M.-H. *et al.* (2017). Impact of reflections on the divertor and first wall temperature measurements from the ITER infrared imaging system. *Nuclear Materials and Energy*, *12*, 1265-1269. https://doi.org/10.1016/j.nme.2017.02.014]
- WEST movie database
- [Yang, Z., Wang, Y., Zhang, C. et al. Techniques Used for Registration and Reconditioning of Infrared Images of the Experimental Advanced Superconducting Tokamak (EAST) Divertor. J Fusion Energ 39, 184–191 (2020). https://doi.org/10.1007/s10894-020-00251-8]
- [Kamiya, K., Itami, K., Takeuchi, M., & Enokuchi, A. (2014). Design study of a wide-angle infrared thermography and visible observation diagnostic on JT-60SA. In Fusion Engineering and Design (Vol. 89, Issue 12, pp. 3089–3094). Elsevier BV. https://doi.org/10.1016/j.fusengdes.2014.09.011]

Need for common standards and datasets

- IRFM developed several open-source packages and standards:
  - librir (https://github.com/IRFM/librir): infrared video storage and exchange
  - thermal-events (https://github.com/IRFM/thermal-events): infrared video annotation creation and exchange
  - thermavip (https://github.com/IRFM/thermavip): offline multi-sensor data analysis, firm real-time processing, and online visualization of sensor data; currently used at WEST and W7-X

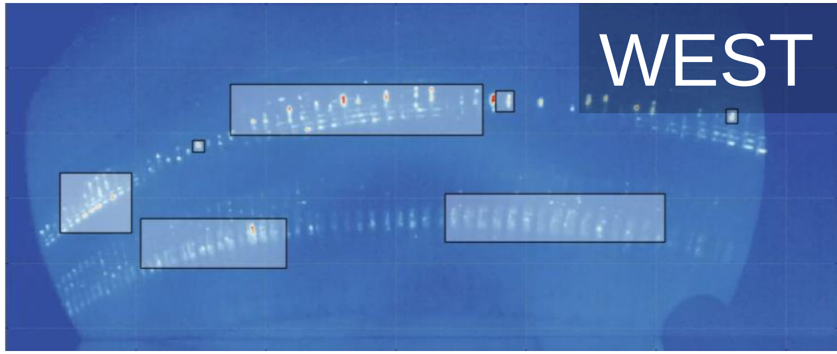

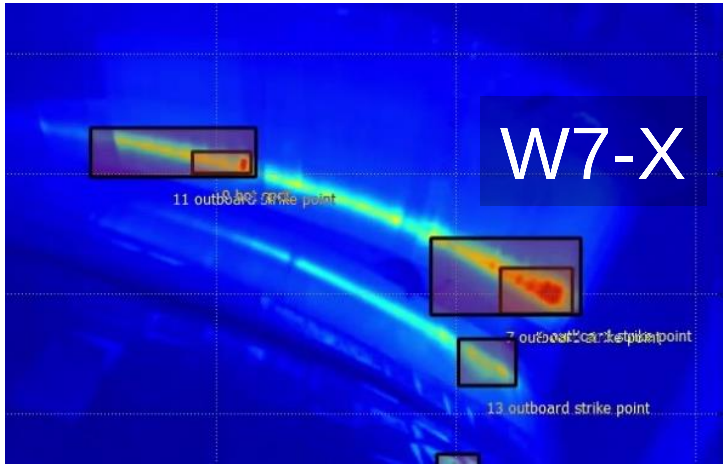

- All the annotations and detections related to the WEST tokamak are stored in a centralized SQL database, using the standards created at IRFM

A demonstrator of a common database of infrared videos and annotations between WEST and W7-X has been deployed on the EUROfusion gateway

E. Grelier et al. "An Open Source Fusion Machine Agnostic Standard for the Exchange and Processing of Infrared Videos and Video Annotations“. *To be published*

Core techniques for model adaptation

- Transfer learning:
  - Use pre-trained models as feature extractors
  - Fine-tune higher layers while freezing lower ones
- Domain adaptation:
  - Align features from known and new machines using adversarial training
- Self-supervised learning:
  - Train on unlabeled data by predicting or simulating missing information
- Simulated data:
  - Use synthetic data from new machines to pre-train models before real-world deployment

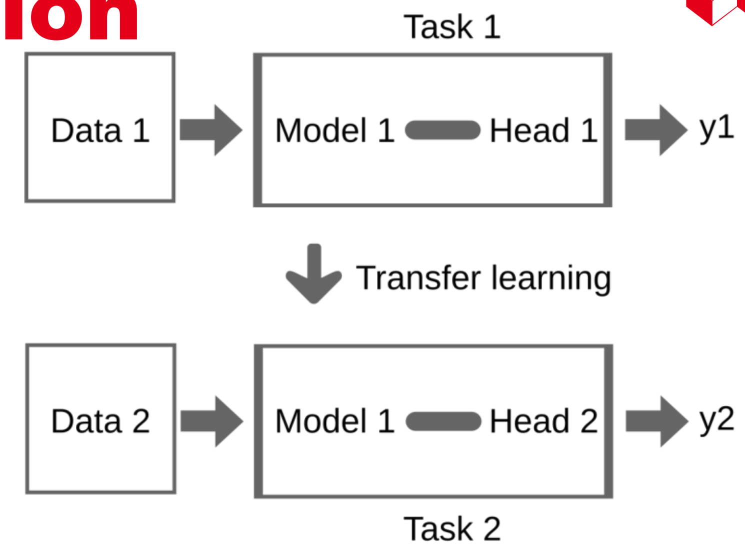

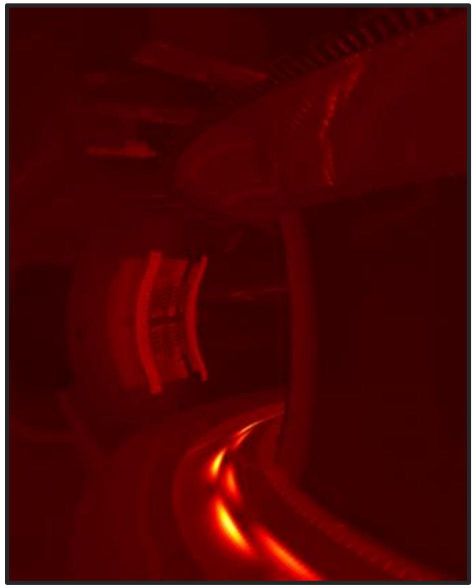

Images: [**A. Juven *et al.*** "U-Net for temperature estimation from simulated infrared images in tokamaks". *Nuclear Materials and Energy (Vol. 38, p. 101562).* . https://doi.org/10.1016/j.nme.2023.101562] and [By Biggerj1 - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=143857678]

Deployment and continuous improvement

- Active learning:
  - Query uncertain predictions for human labeling to iteratively improve the model
- Online learning:
  - Continuously adapt the model with streaming data from new machines
- Cross-machine validation:
  - Test models across diverse machine types to ensure robustness

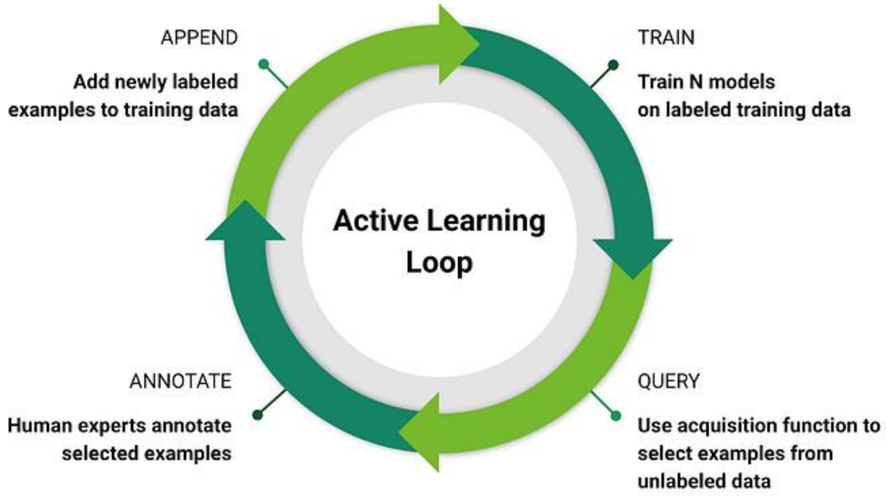

[https://medium.com/nvidia-ai/scalable-active-learning-for-autonomous-driving-a-practical-implementation-and-a-b-test-4d315ed04b5f]

Workflow for adaptation to other machines

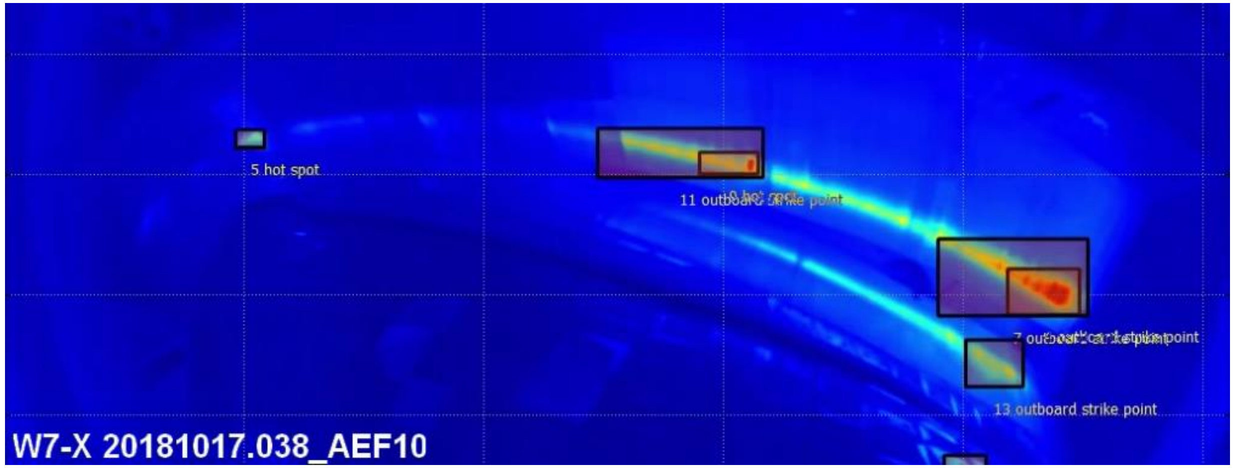

- Train models on existing machine data using transfer learning
- Incorporate simulated data from new machines for pre-training and validation
- Deploy models with active and online learning for continuous improvement with new data
- Use cross-validation across machines to maintain robustness

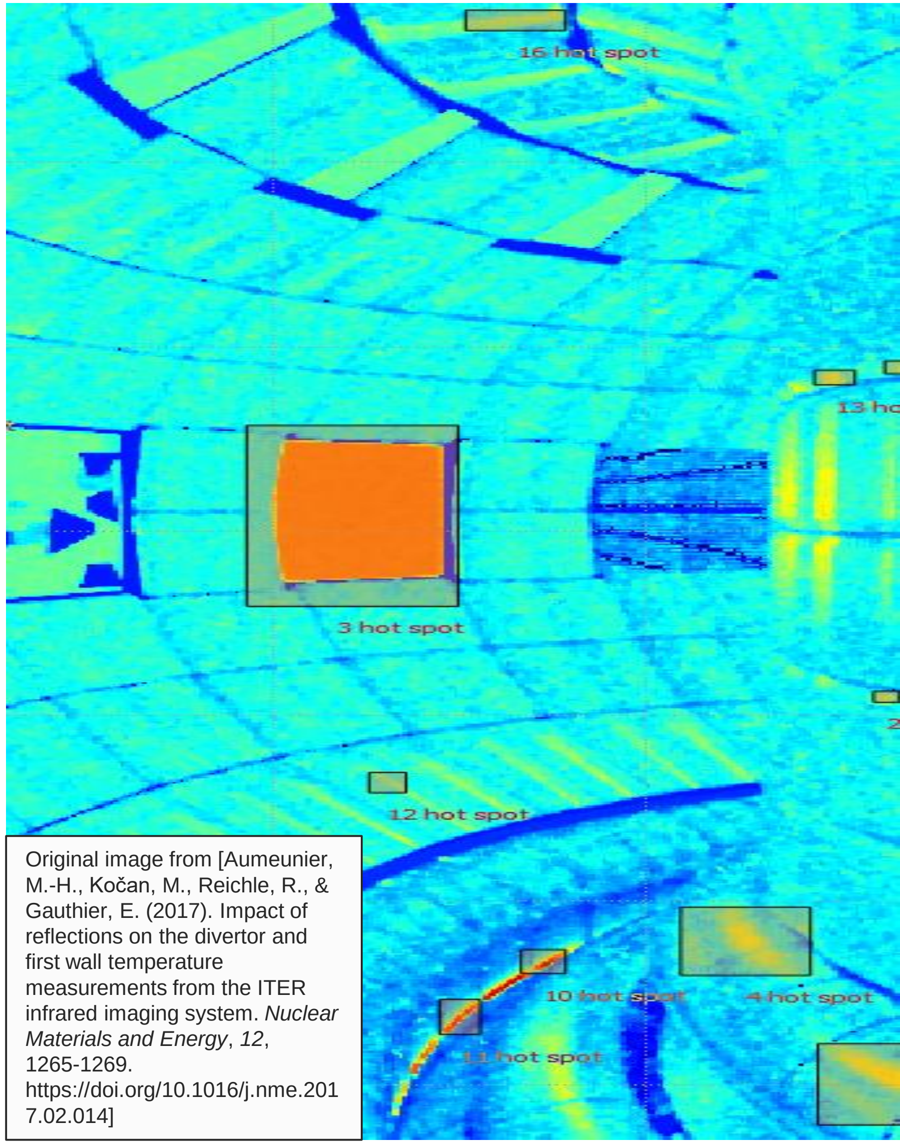

Original image from [Aumeunier, M.-H., Kočan, M., Reichle, R., & Gauthier, E. (2017). Impact of reflections on the divertor and first wall temperature measurements from the ITER infrared imaging system. *Nuclear Materials and Energy*, *12*, 1265-1269. https://doi.org/10.1016/j.nme.2017.02.014]

6. Conclusion and key takeaways

Conclusion

- Infrared thermography is a crucial tool for machine monitoring and protection
- ITER will require an advanced protection system based on infrared data from day one of operation
- Machine learning is a powerful tool for the development of such systems on existing machines and simulations from future machines
- Diverse data from several fusion machines should be put together to train generalizable models

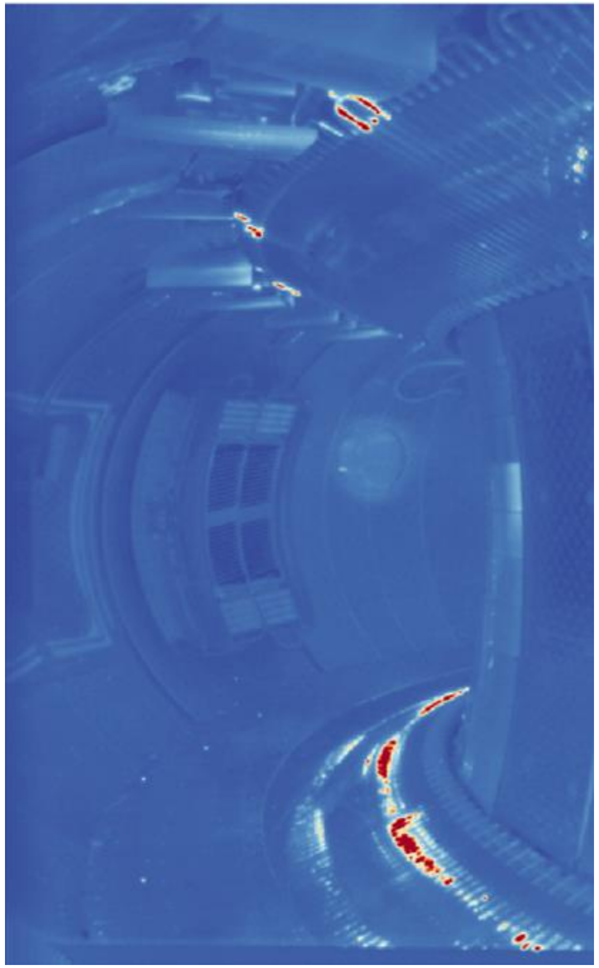

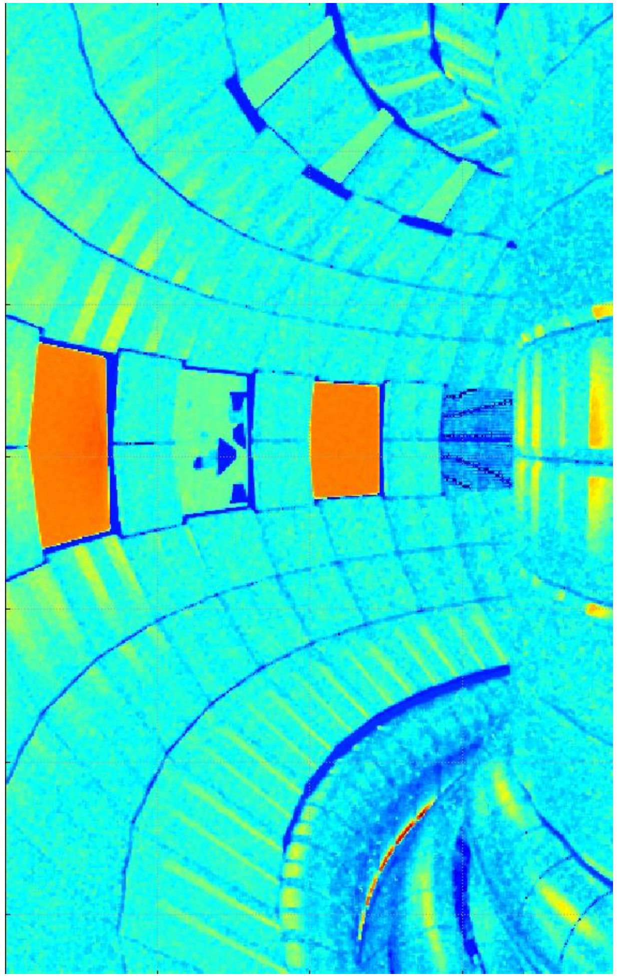

WEST ITER

Key takeaways

- Data is key! "Garbage in, garbage out"
- A good model cannot fix bad data, but good data may make up for a less-than-optimal model
- Modularity is key to quickly be able to tackle new problems, upcycling existing codes and architectures
- Appropriate metrics and model explainability are crucial for the public trust and acceptability of the model, especially for nuclear machines
- Good ways of displaying the results is paramount to the adoption of the model during operation

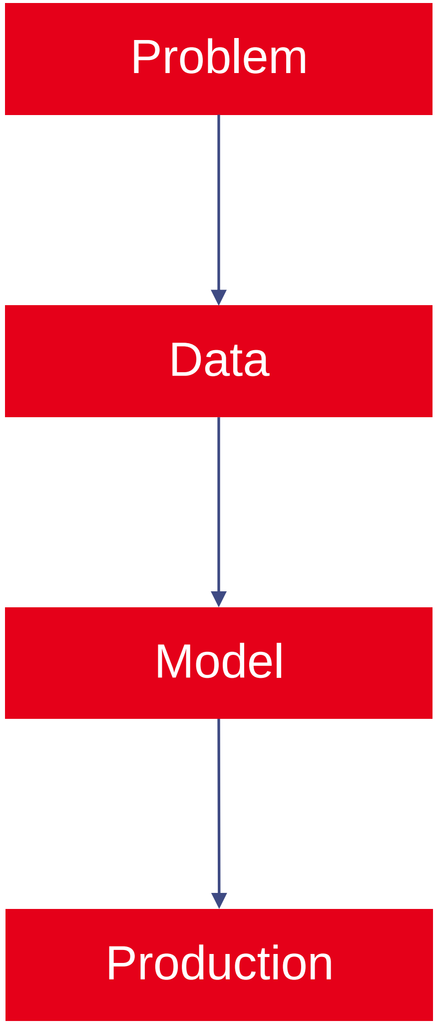

Thank you for your attention

Erwan Grelier (erwan.grelier@cea.fr)

CEA, IRFM

Saint-Paul-Lez-Durance, F-13108, France

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.