Computer Vision for Machine Protection in Nuclear Fusion

8th EIROforum School on Instrumentation
13th – 17th May 2024

Bartłomiej Jabłoński, Dariusz Makowski
Lodz University of Technology (TUL), Lodz, Poland

Aleix Puig Sitjes, Marcin Jakubowski
Max Planck Institute for Plasma Physics (IPP), Greifswald, Germany

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), and in part by the Engineering Grant EEG21-17, and in part by the Polish Ministry of Science and Higher Education within the International Co-financed Projects (PMW). 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.

Agenda

- ❑ Brief introduction to the Wendelstein 7-X (W7-X) stellarator and infrared diagnostics
- ❑ Computer vision applications:
  - Thermal overload detection/anticipation (machine protection, real-time, image processing)
  - Anomalous image detection (unsupervised learning, anomaly detection, autoencoder)
  - Dataset preparation (image annotation, image processing)
  - Thermal event instance segmentation (supervised learning, deep learning, detection, segmentation)
- ❑ Summary

The Wendelstein 7-X Stellarator

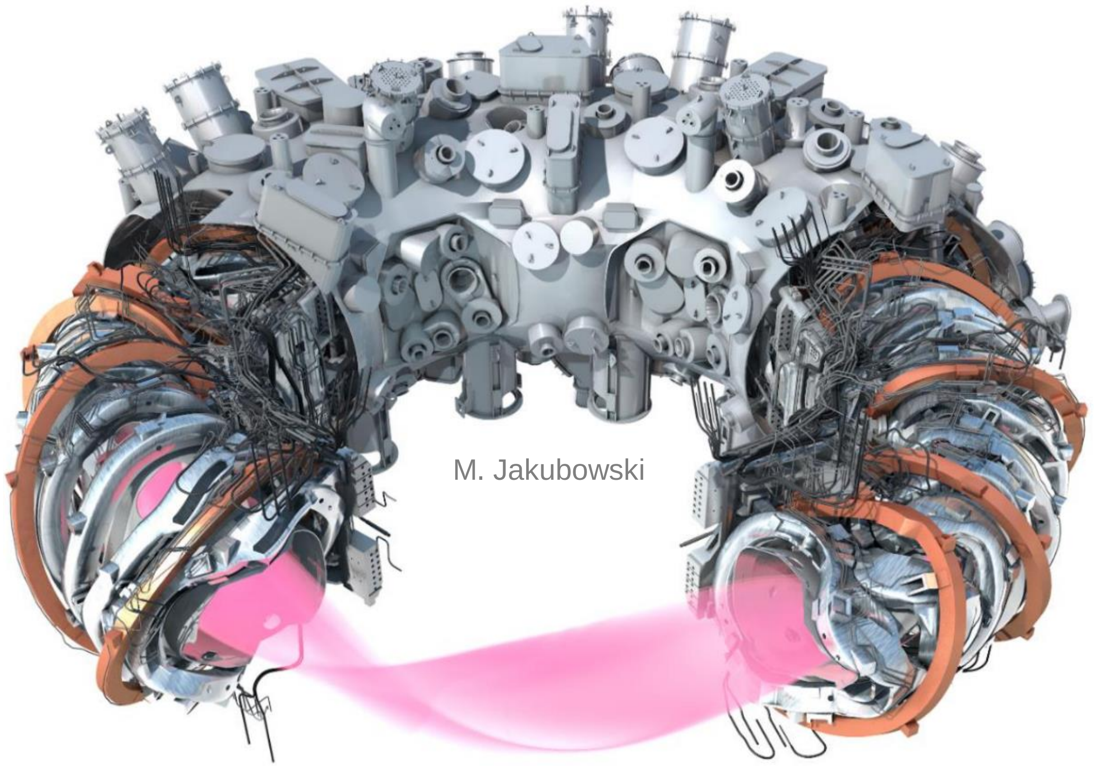

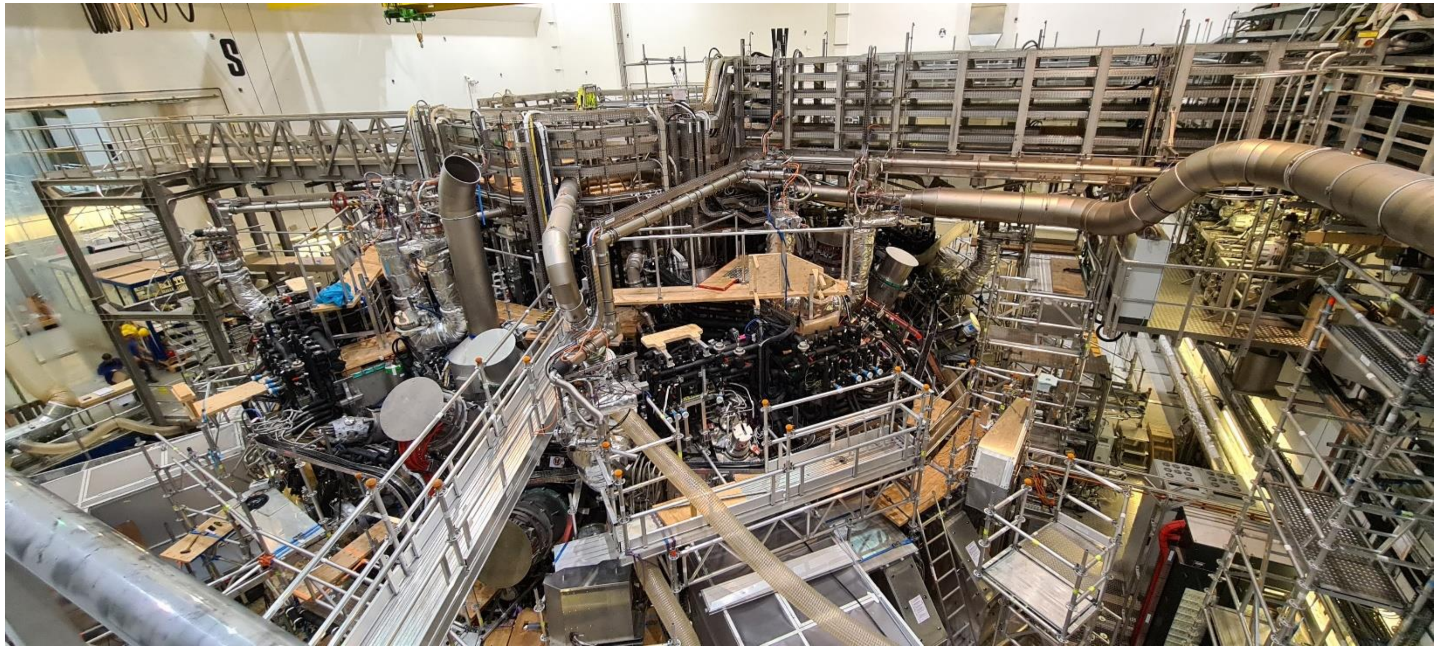

- The largest and most advanced stellarator in the world
- Located at the Max Planck Institute for Plasma Physics (IPP) in Greifswald, Germany
- Imaging systems are used for plasma diagnostics

- Expected to sustain plasma for **30 min** in the upcoming Operational Phases (OPs) 2.N (2.2 from September 2024)
- Reached **480 s** (8 min) discharge of 1.3 GJ energy turnover with on average 2.7 MW of heating power in OP2.1 (2023)

Plasma Image Diagnostics

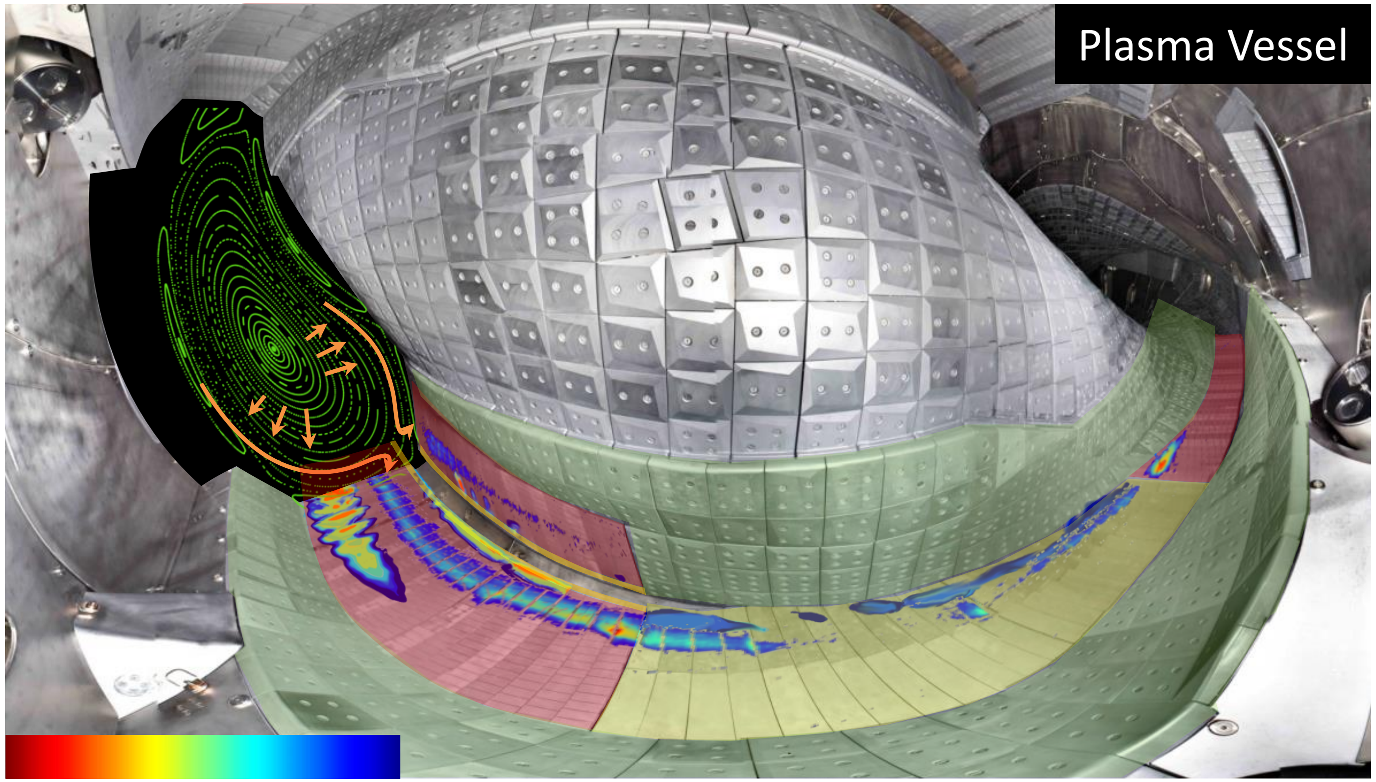

Divertor heat load [a.u.]

J. Fellinger *et al.*, *PFMC-19*, 2023

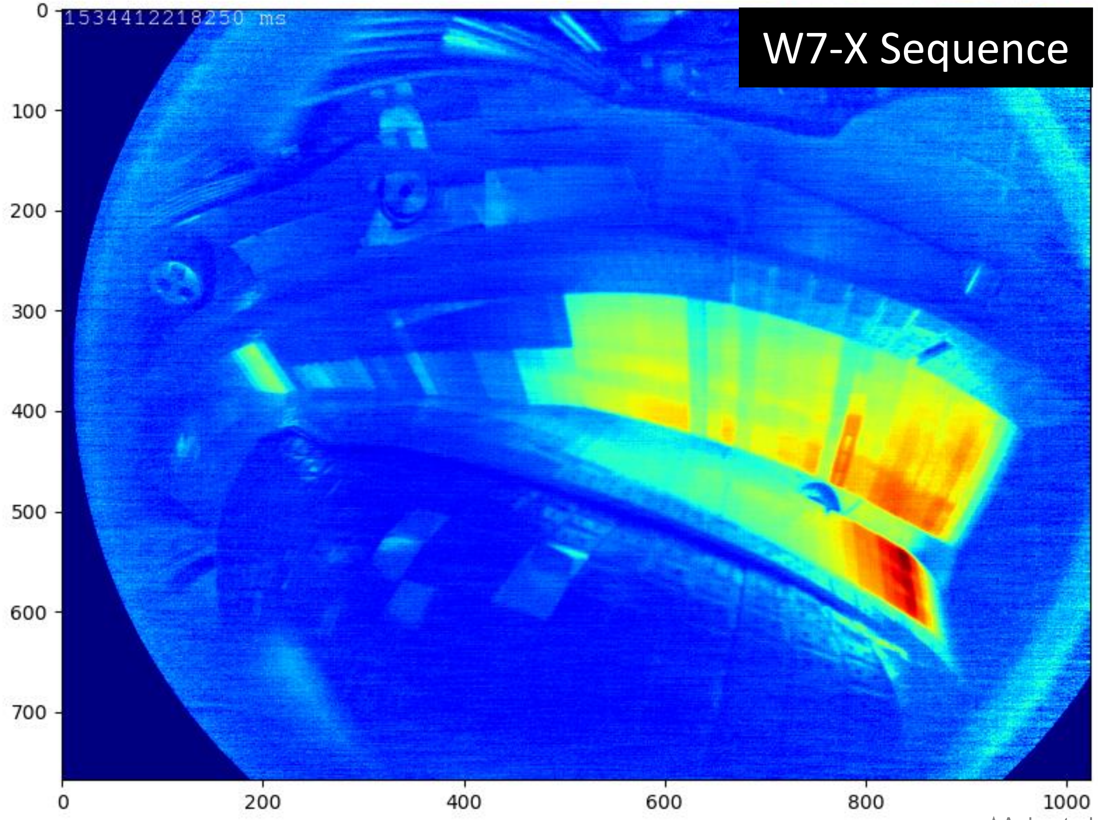

*Animated

Basic Protection

- Protect water-cooled Plasma Facing Components (PFCs) from thermal overloads with infrared (IR) cameras
- Trigger the Fast Interlock System (FIS) to terminate a discharge when a thermal overload is anticipated

Advanced Protection & Control (facilitated by deep learning)

- Thermal event detection
- Feedback control

W7-X has 12 IR cameras, and 10 divertor units are monitored

Thermal Overload Detection (TOD)

Process **entire FoV** without pre-defined RoIs and assess risk to **dynamically** adjust temperature thresholds.

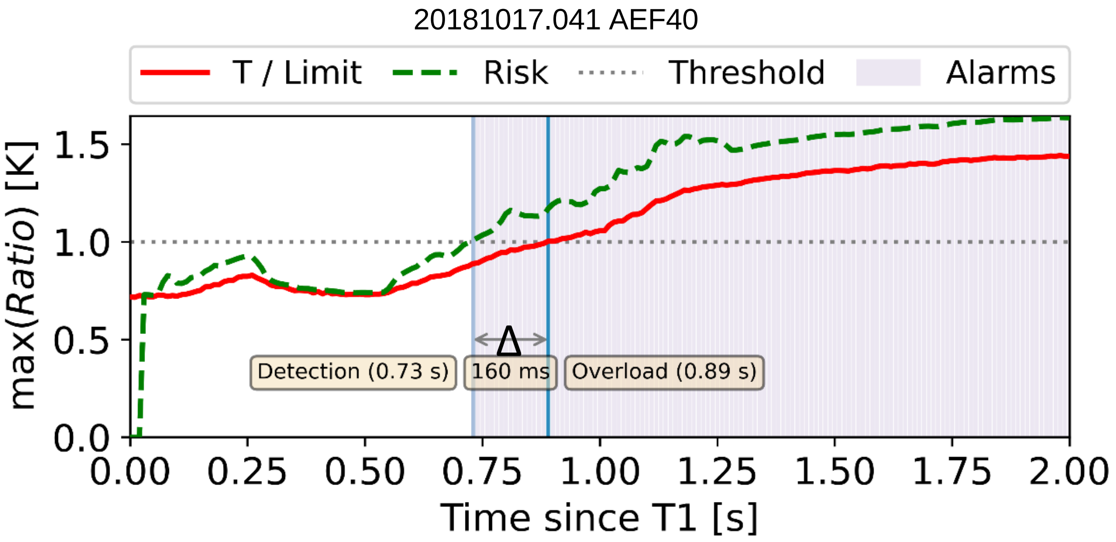

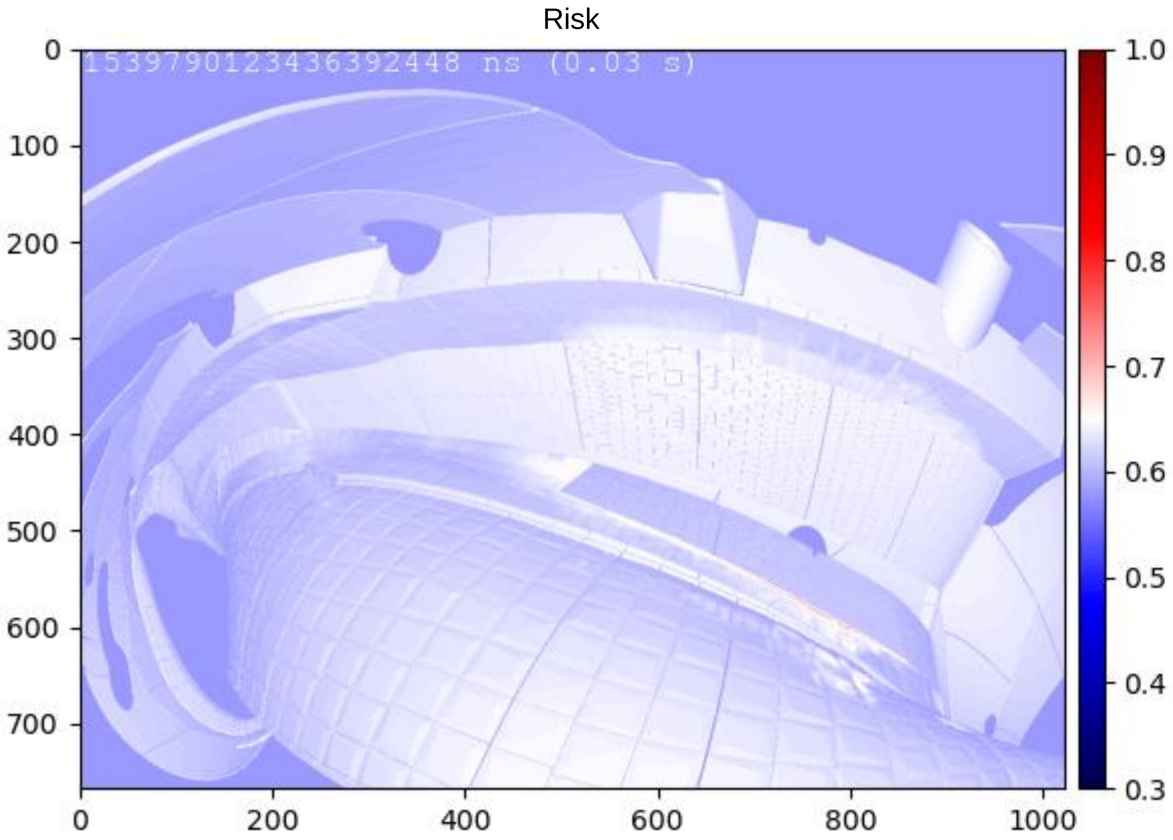

Requirement: **Δ ≥ system delay**, i.e., $\begin{cases} 50\,ms, & CPU \\ 20\,ms, & GPU \end{cases}$

*Animated

Thermal Overload Detection (TOD) – Evaluation

Evaluated the entire Operational Phase (OP) 1.2, totalling 1419 discharges corresponding to 12`074 discharge sequences (**19`447`678** images).

Conclusion: Full-frame infrared image processing with dynamic temperature thresholding improves the versatility and effectiveness of real-time machine protection in thermonuclear fusion devices compared to the protection of predefined regions of interest with fixed temperature thresholding.

	Delay = 50 ms (CPU)		Delay = 20 ms (GPU)	
Threshold	Sensitivity	FPR	Sensitivity	FPR
Dynamic	0.970	0.052	0.997	0.042

	Delay = 30 ms (CPU)		Delay = 20 ms (GPU)	
Threshold	Sensitivity	FPR	Sensitivity	FPR
80% of the PFC limit	0.990	0.263	0.997	0.263
85% of the PFC limit	0.984	0.166	0.995	0.166
90% of the PFC limit	0.963	0.095	0.994	0.095
95% of the PFC limit	0.927	0.043	0.958	0.043

$$\uparrow \text{senstivity} = \frac{TP}{TP + FN} \qquad \downarrow \text{FPR} = 1 - \text{specificity} = \frac{FP}{FP + TN}$$

False positives: Heating was turned off after the alarm [dynamic CPU]

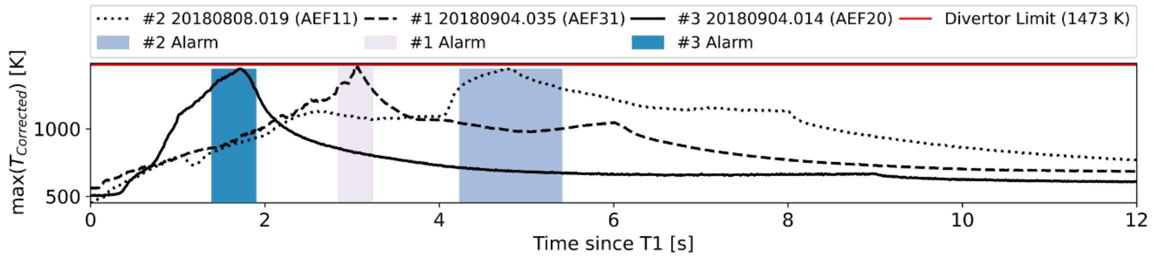

False positives: Operation close to the temperature limits [dynamic CPU]

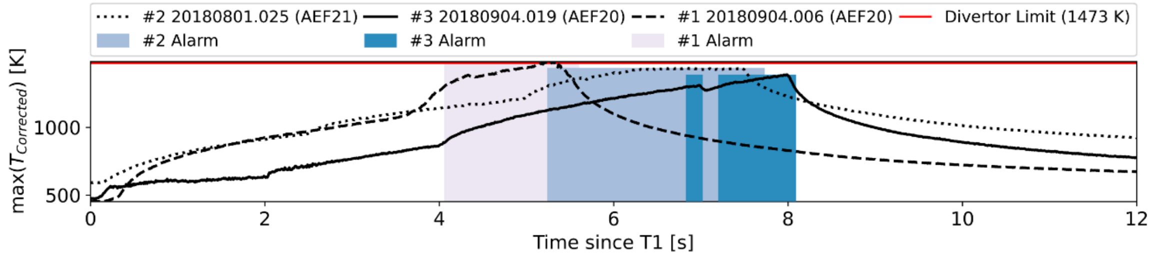

Wendelstein 7-X

Transmission interference

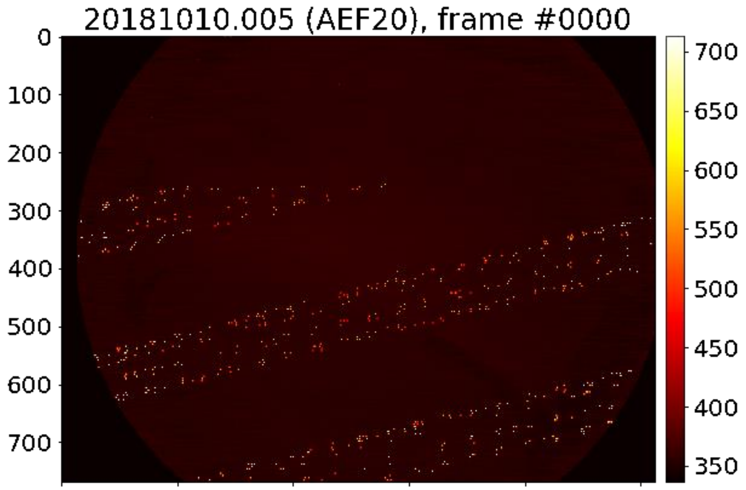

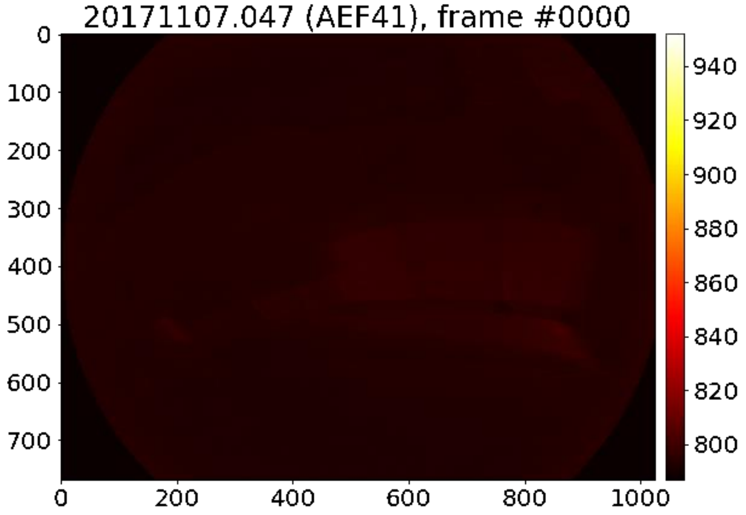

Narrow dynamic range

Transmission interference and bottom artifacts

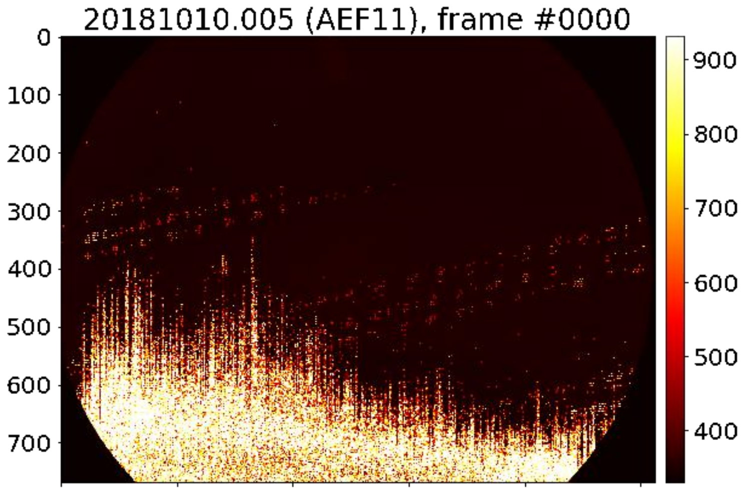

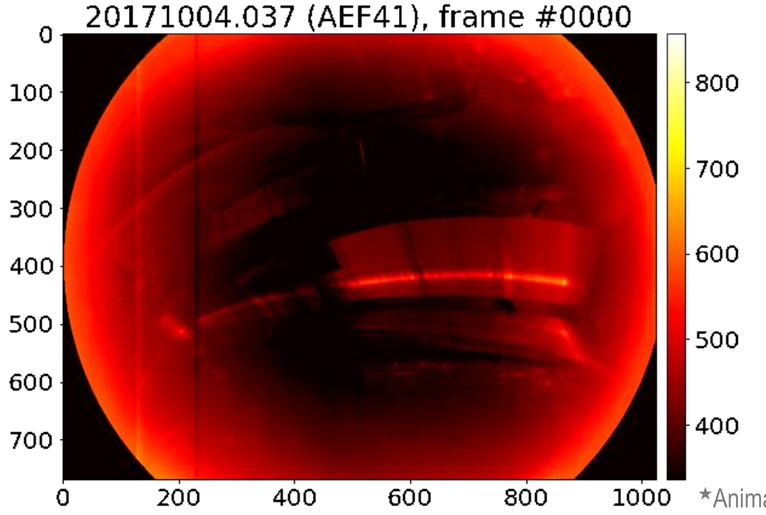

*Animated

Vertical strap and strong vignetting

✓ Correct discharge sequence data:

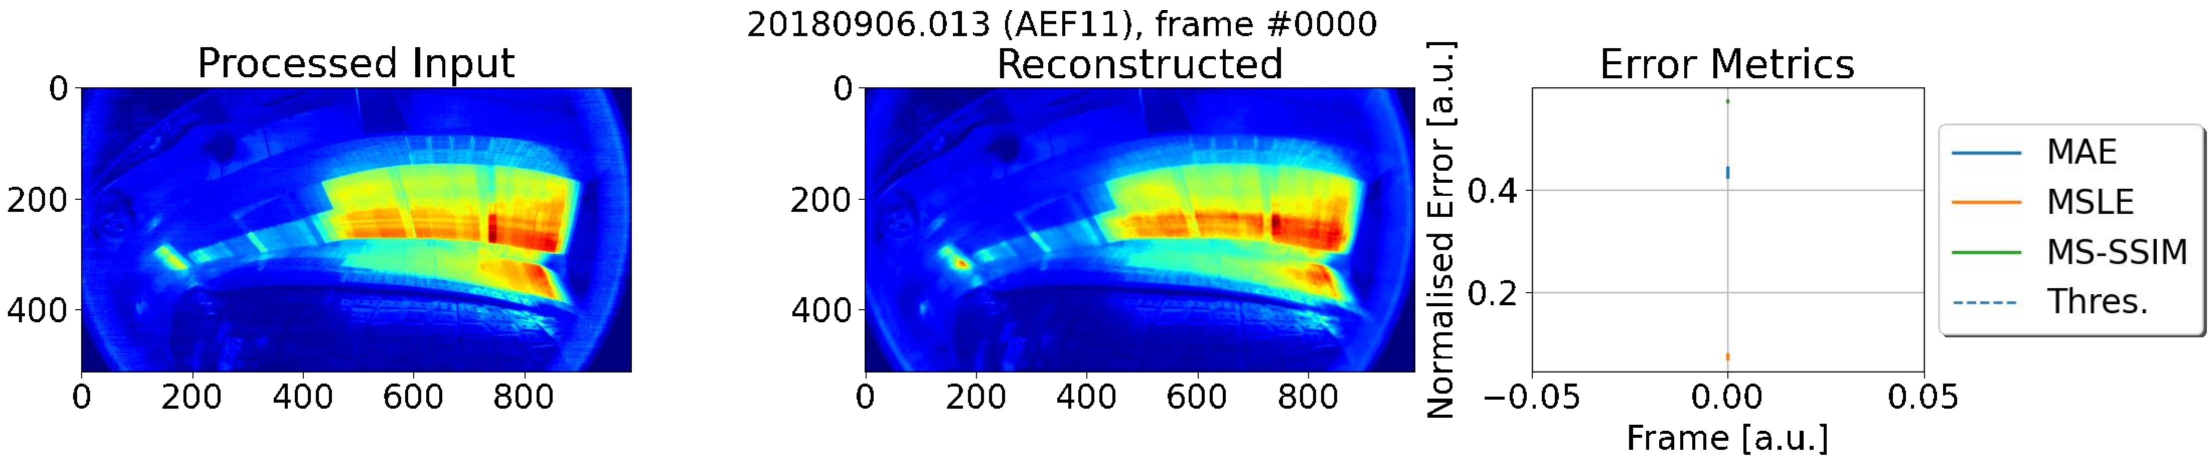

✗ Incorrect discharge sequence data, i.e., vertical strap and strong vignetting:

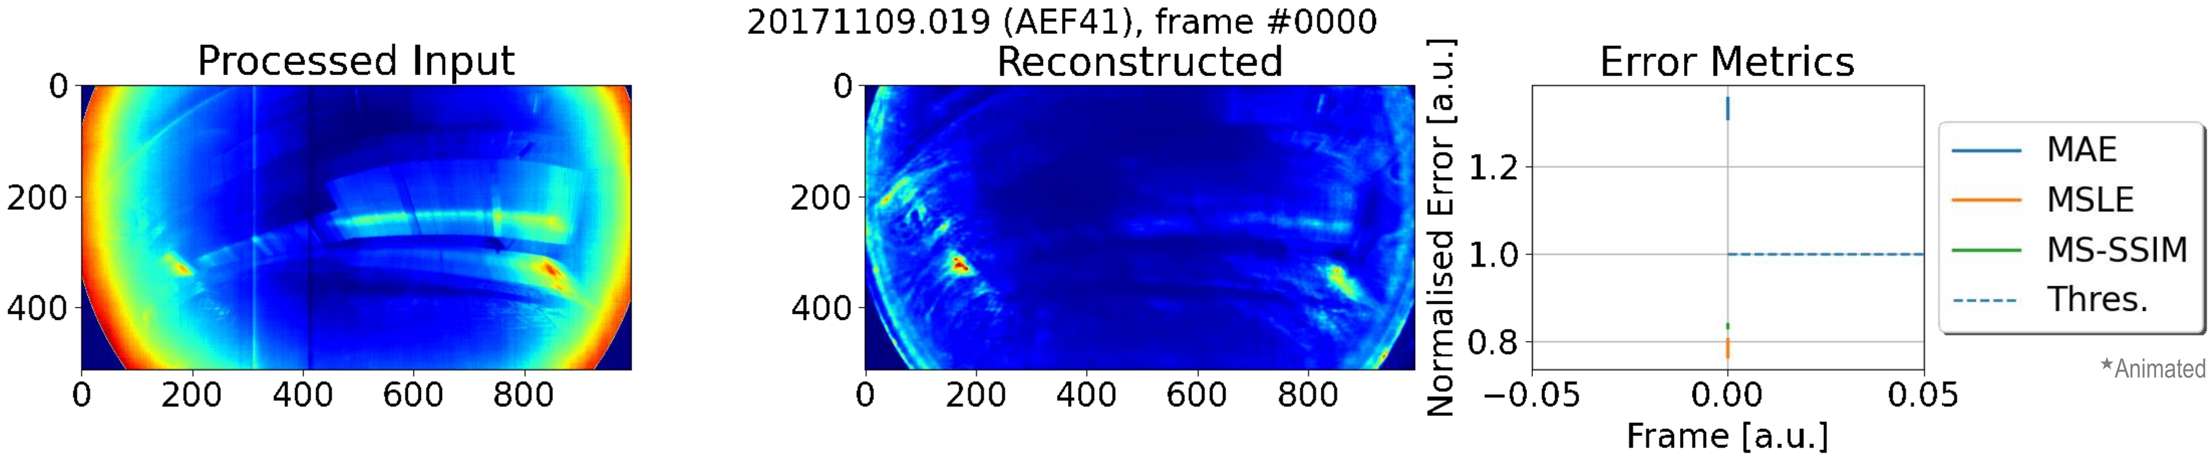

Anomaly Detection – Evaluation

Manually identified **96** invalid discharges (not counting sequences with partially incorrect images) out of **12`998** OP1.2 discharge sequences excluding noise-only or blank cases (0's)

(1) AUC = **1.0000** discrimination between good vs bad images (balanced test dataset of 8108 randomly sampled images)

(2) AUC = **0.9938** discrimination between noise-only vs 125 – 175 °C plasma build-up images after T1 signal (balanced dataset of 414 randomly sampled images)

❖ Total latency **5.5 ms** (Tesla T4 – FP16: transfer, pre-processing, inference, errors)

Distribution of Reconstruction Error (1)

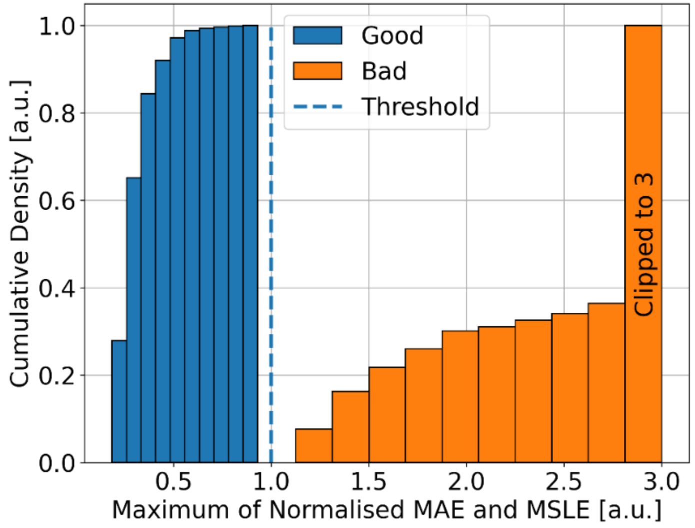

Receiver Operating Characteristic (1,2)

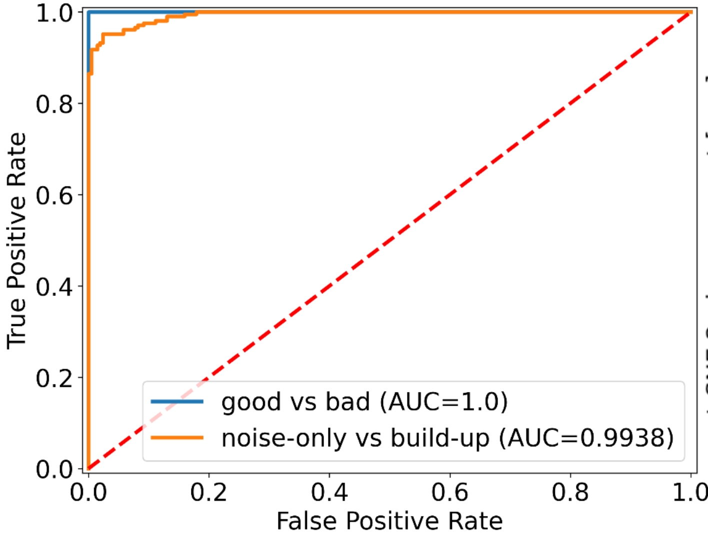

Latent Space Projection (1)

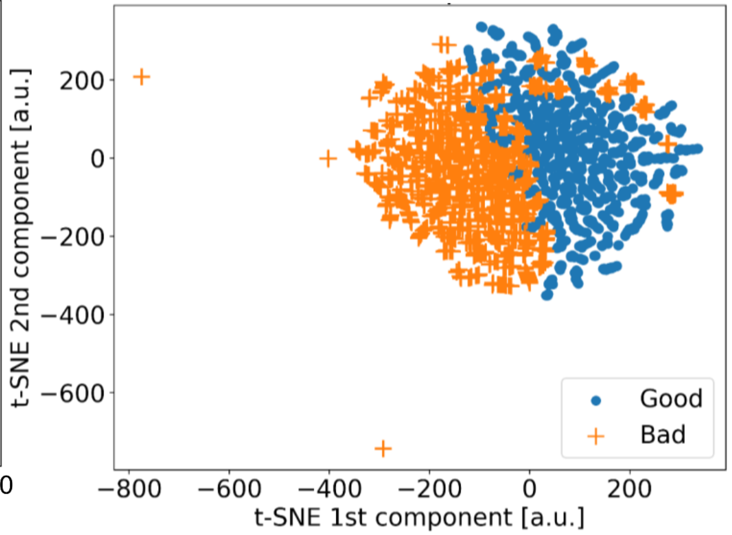

Thermal Events in Infrared Images

[1] Strike-line (SL)

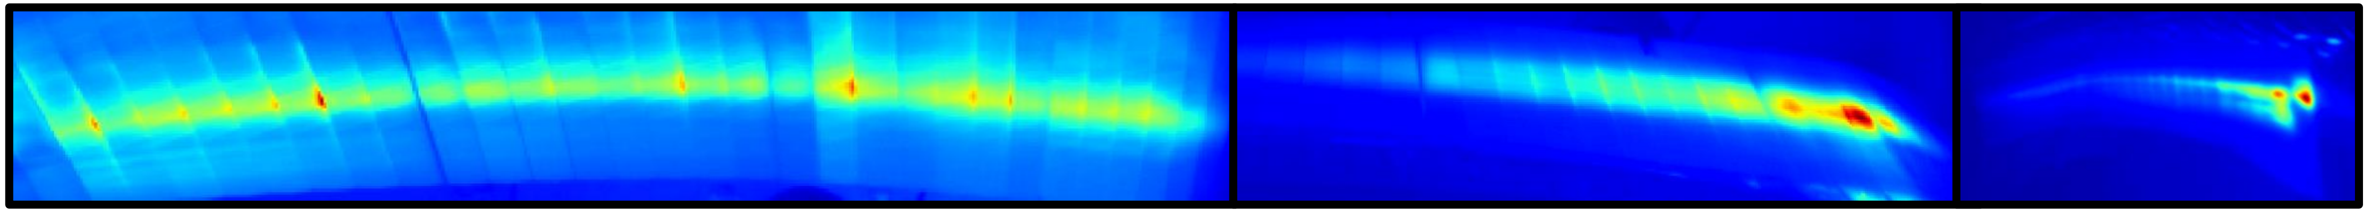

[2] Reflection (R)

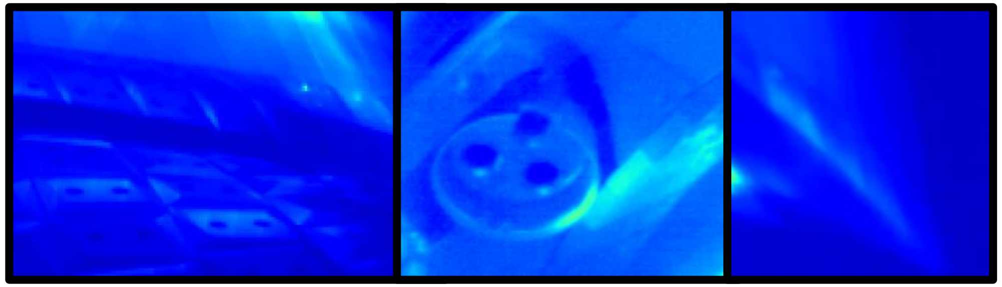

[3] Hot-spot (HS)

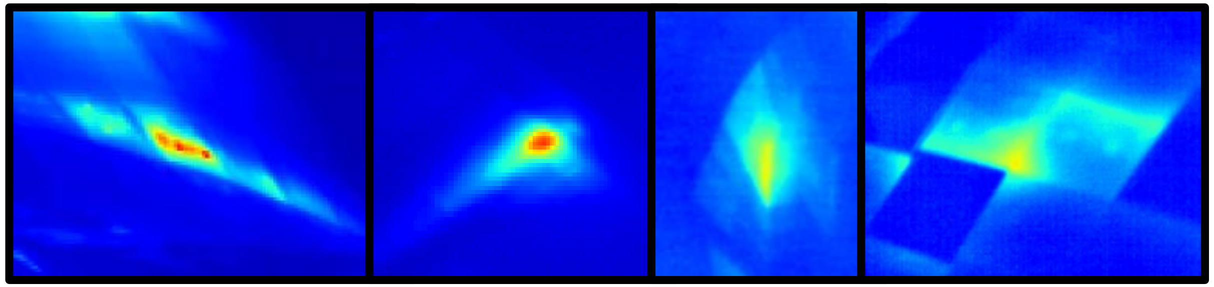

[4] UFO

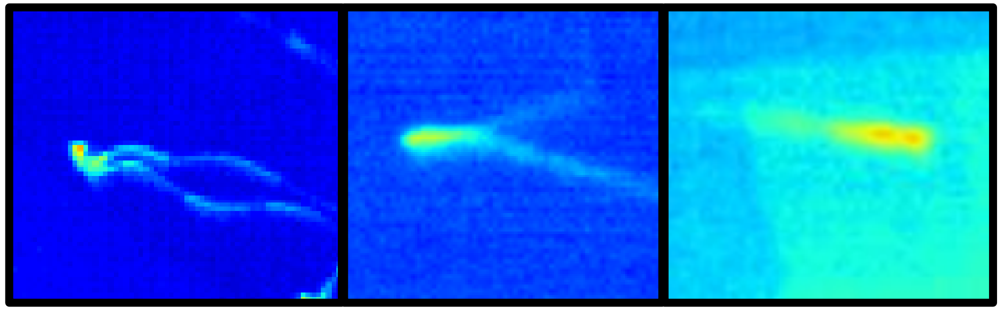

[5] Leading Edge (LE)

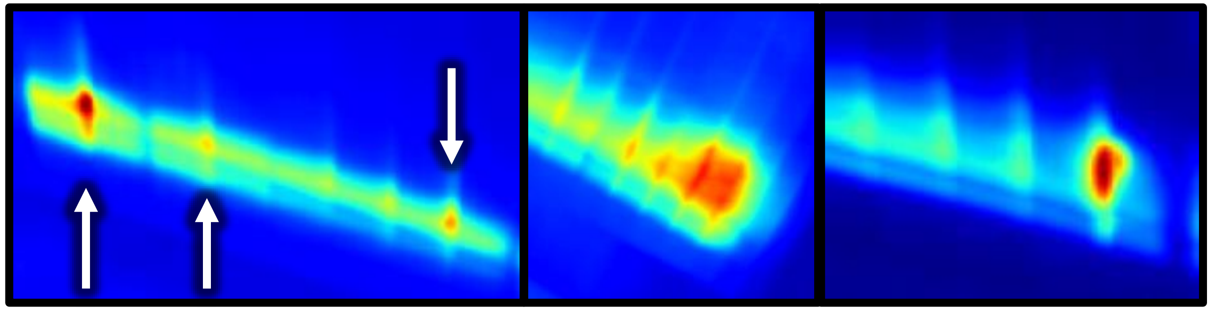

Thermal Events – Literature Review

Tungsten Environment in Steady-state Tokamak (WEST) Cadarache, France

„Deep learning and image processing for the automated analysis of thermal events on the first wall and divertor of fusion reactors” 09.2022

„Deep learning-based process for the automatic detection, tracking, and classification of thermal events on the in-vessel components of fusion reactors” 03.2023

Issues with a deep learning for computer vision dataset:

- Manually annotated by a group of experts
- Highly time-consuming (complex, noisy)
- Inconsistencies between annotators (fuzzy event borders)
- Bounding-boxes, not segmentation masks
- Device specific issues, higher image resolutions, framerates

Propose a **semi-automatic** approach for infrared image annotation for thermonuclear fusion devices to **accelerate the annotation process** and **impose consistency** while integrating expert knowledge.

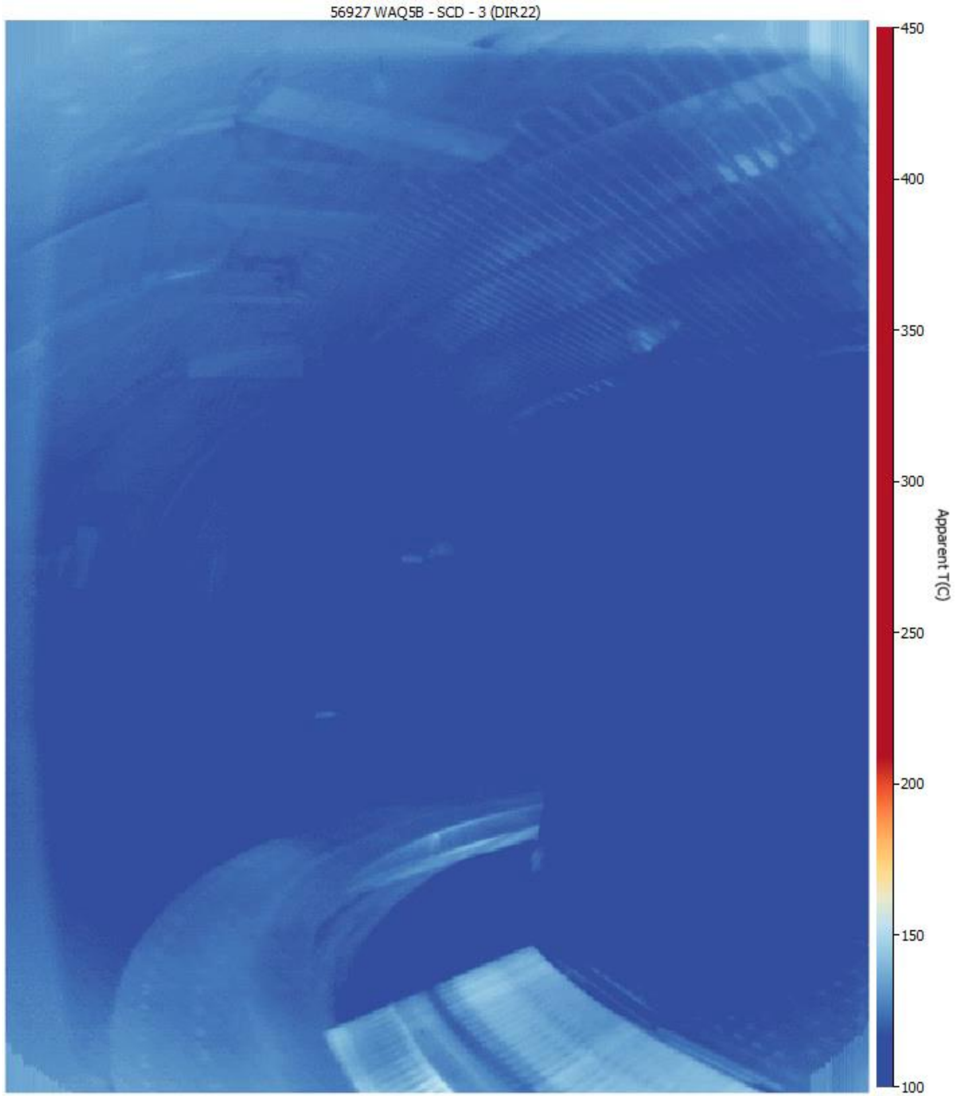

X. Courtois *et al.*, *IR Calibration Workshop*, 2023

*Animated

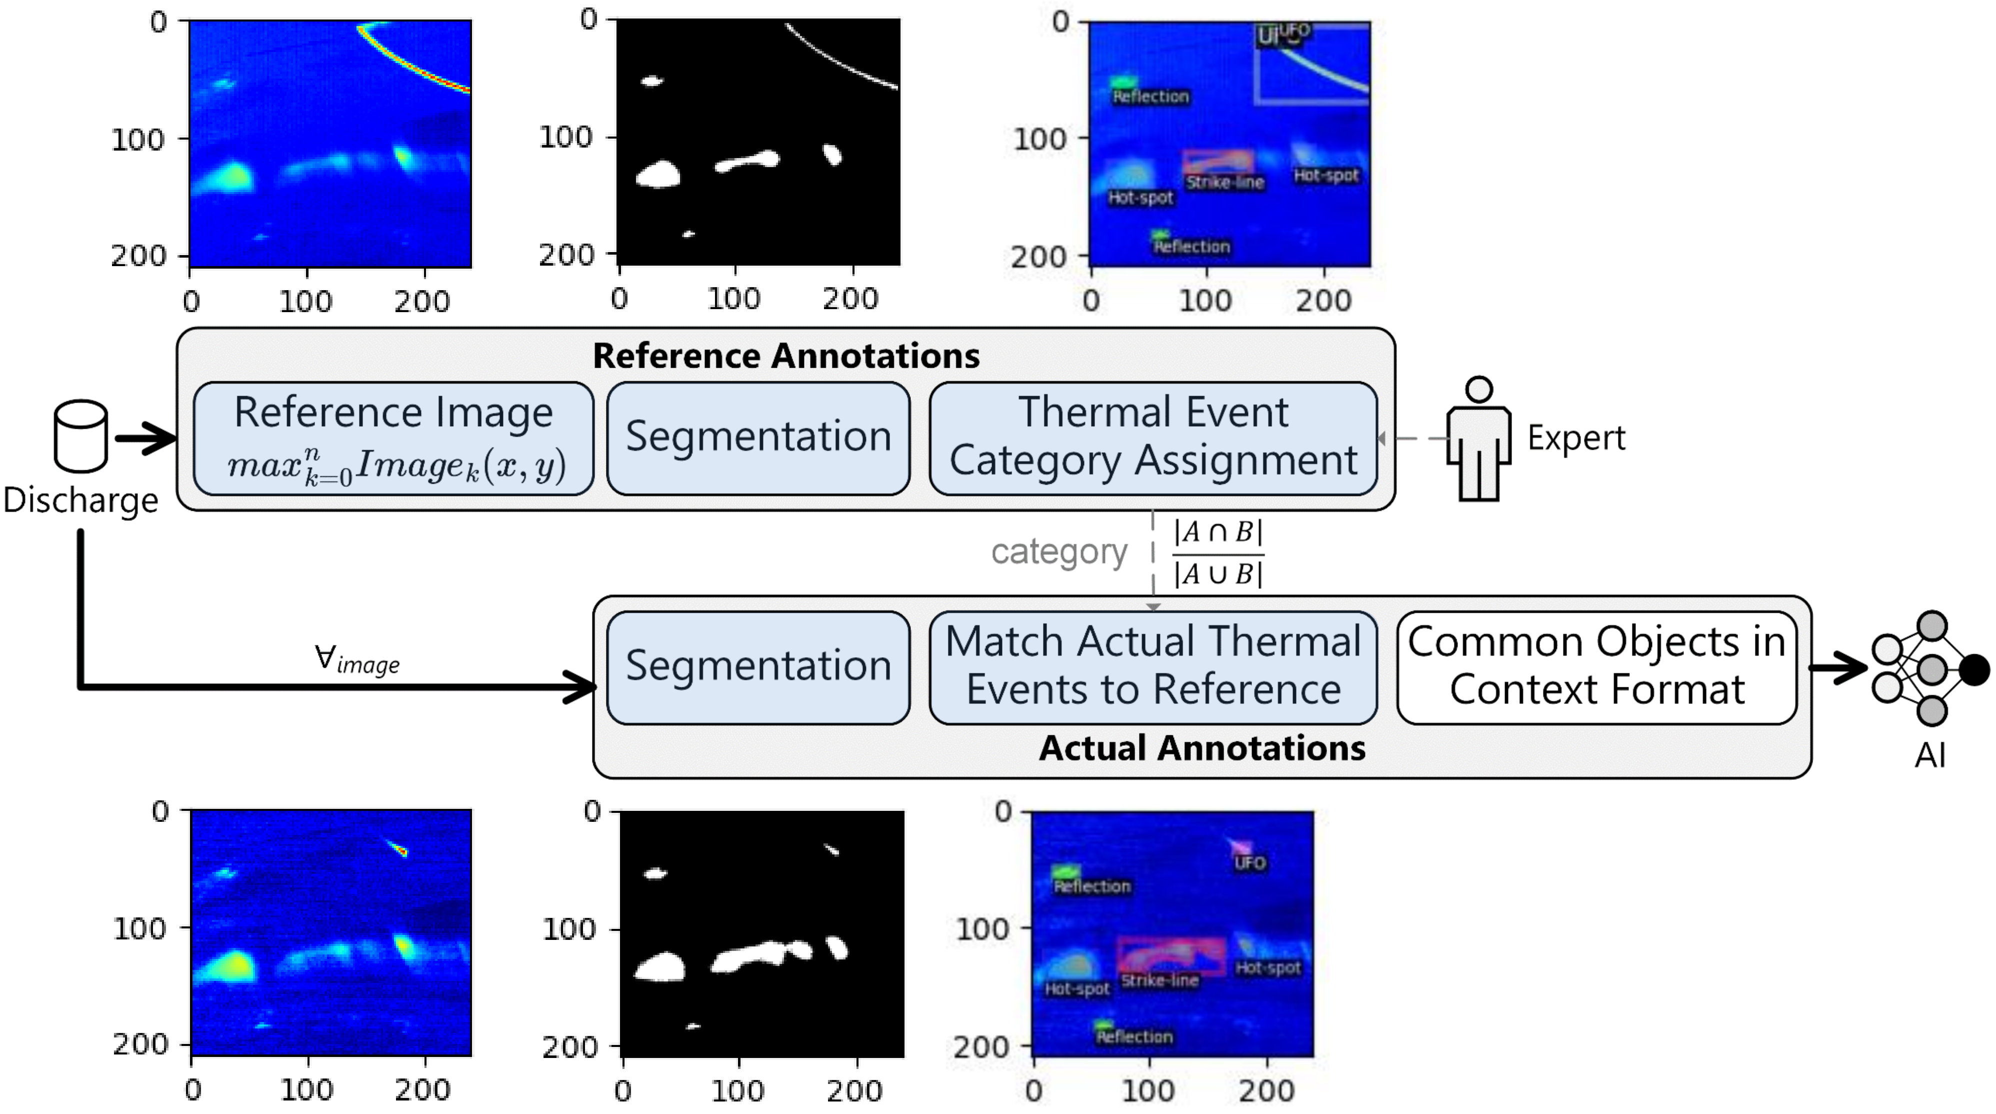
Reference Annotations
Reference Image
$max_{k=0}^{n} Image_k(x,y)$
Segmentation
Thermal Event Category Assignment
Expert
Discharge
category
$\frac{|A \cap B|}{|A \cup B|}$
$\forall_{image}$
Segmentation
Match Actual Thermal Events to Reference
Common Objects in Context Format
Actual Annotations
AI
UFO
Reflection
Hot-spot
Strike-line
0
100
200

Qualitative Results – Difference

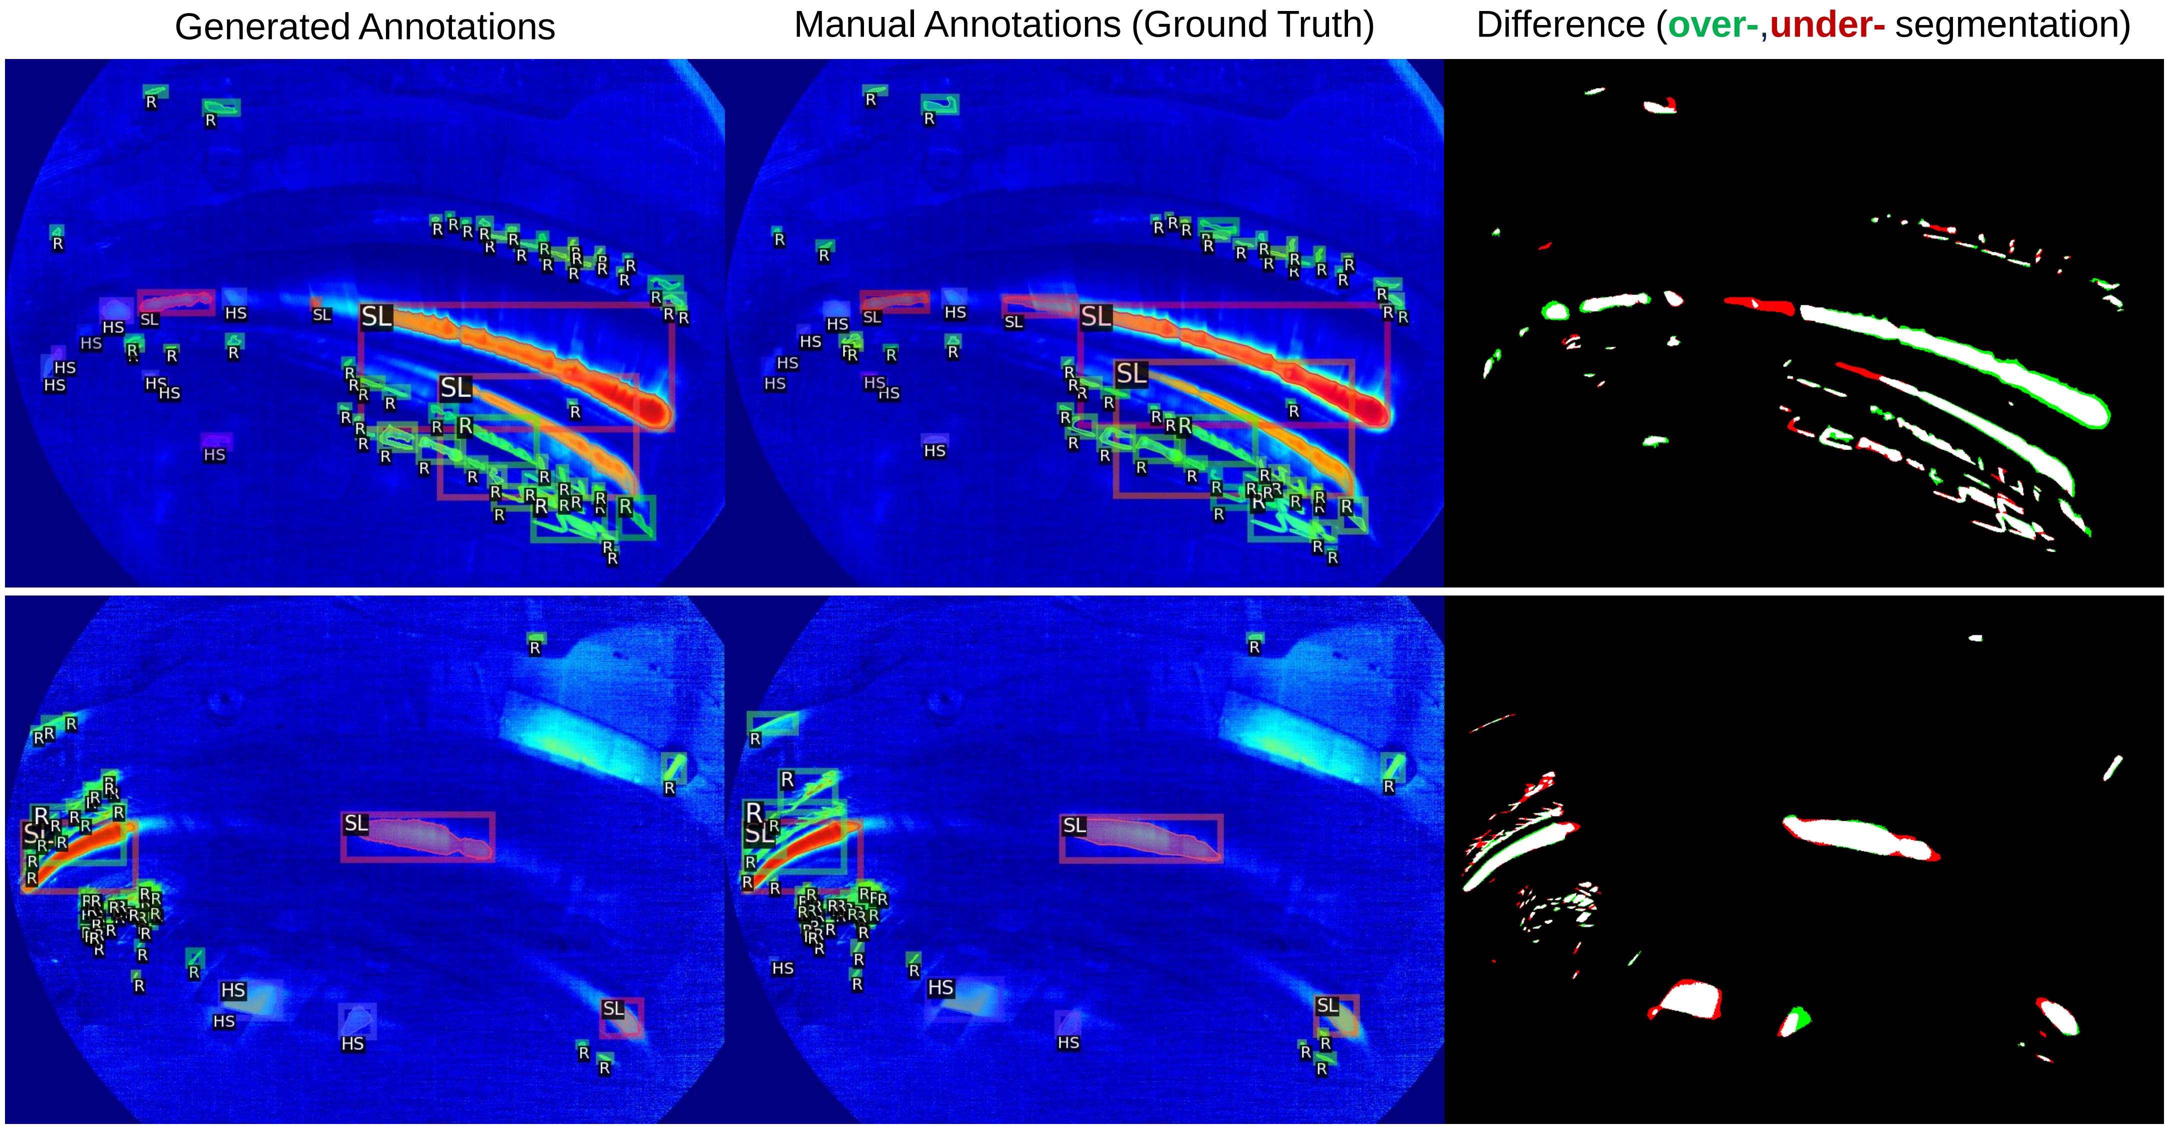

Qualitative Results – Generated Sequence

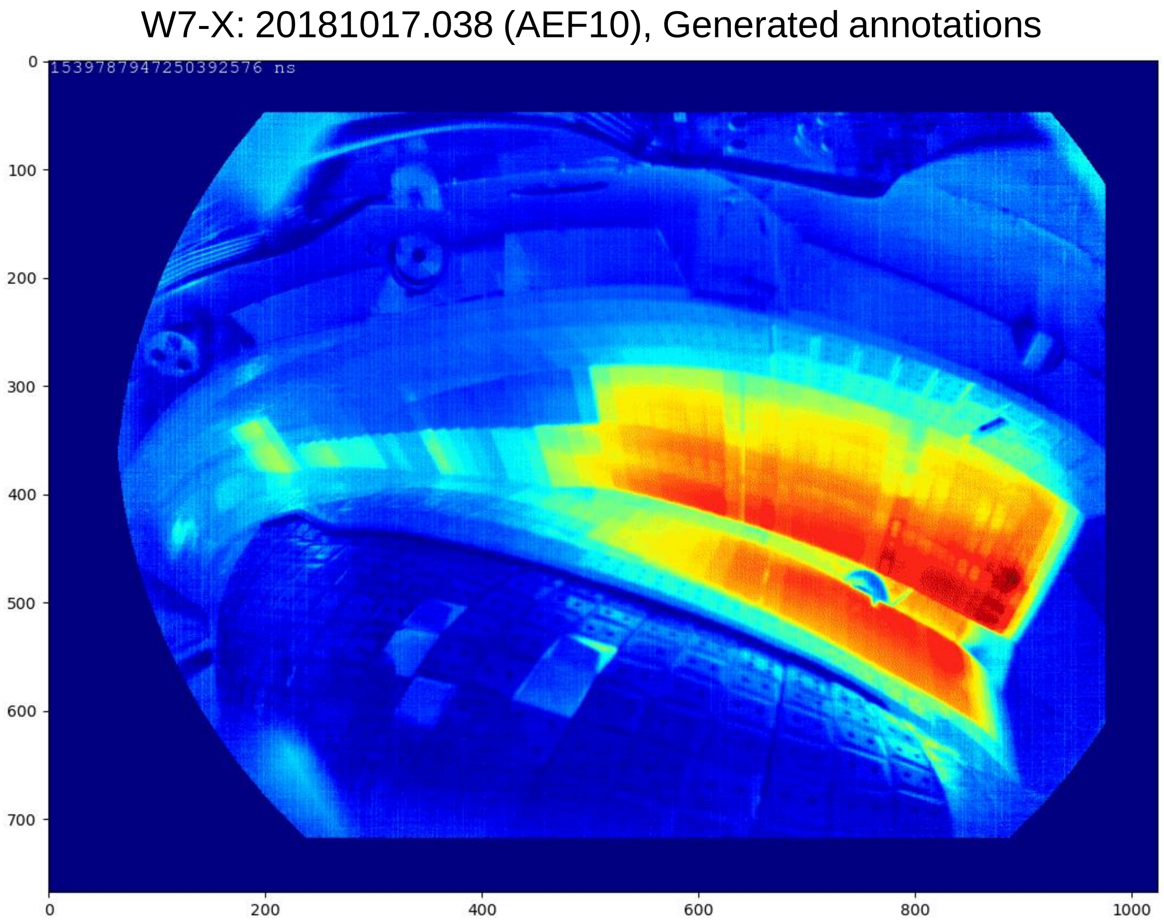

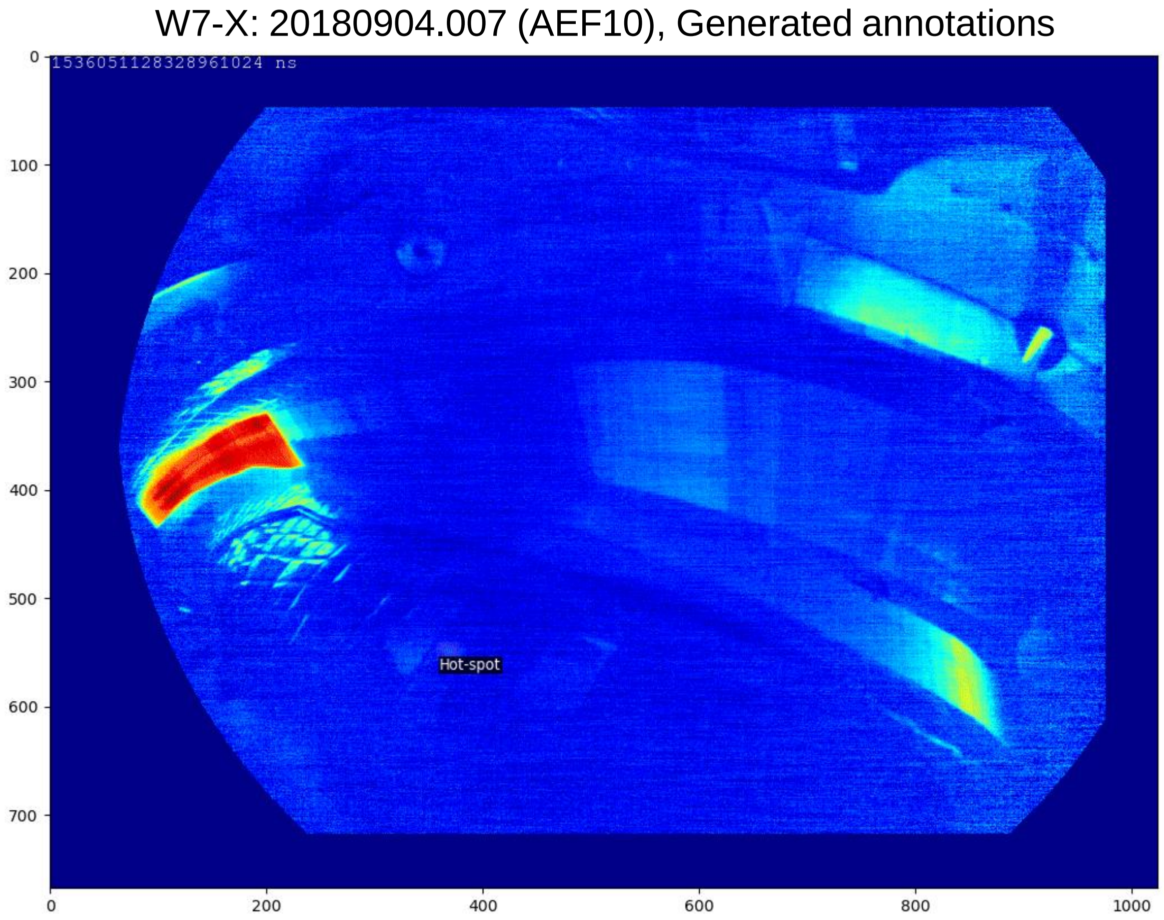

*Animated

Quantitative Results

Metric

Sørensen–Dice Coefficient (SDC)

$$SDC(a,b) = \frac{2\sum_i^N a_i b_i}{\sum_i^N a_i + \sum_i^N b_i}$$

temperature over limit weighted SDC (tlwSDC)

$$tlwSDC(a,b,w) = \frac{2\sum_i^N a_i b_i w_i}{\sum_i^N a_i w_i + \sum_i^N b_i w_i}$$

$$w_i = \frac{T_i - T_i^{Background}}{T_i^{Limit}}$$

Evaluation dataset of 21 manually annotated Ground Truth (GT) images prepared with: *„ClickSEG: A Codebase for Click-Based Interactive Segmentation”*

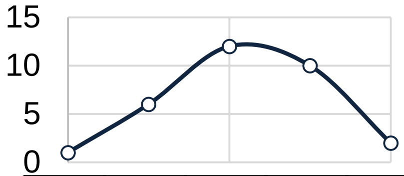

$\vec{w}$	1	6	12	10	2		
$\vec{a}$ (GT)	0	1	1	1	0	SDC	$tlwSDC$
$\vec{b}_1$	0	0	1	1	0	0.8	0.88
$\vec{b}_2$	0	1	1	0	0	0.8	0.78

	Otsu's thresholding (baseline)		Adaptive Gaussian thresholding		Proposed method	
Mask	$\overline{\text{SDC}} \pm s_{\text{SDC}}$	$\overline{\text{tlwSDC}} \pm s_{\text{tlwSDC}}$	$\overline{\text{SDC}} \pm s_{\text{SDC}}$	$\overline{\text{tlwSDC}} \pm s_{\text{tlwSDC}}$	$\overline{\text{SDC}} \pm s_{\text{SDC}}$	$\overline{\text{tlwSDC}} \pm s_{\text{tlwSDC}}$
Any	0.586 ± 0.173	0.777 ± 0.085	0.780 ± 0.050	0.881 ± 0.035	**0.825 ± 0.030**	**0.904 ± 0.018**
Strike-line	0.721 ± 0.230	0.860 ± 0.198	0.828 ± 0.068	0.917 ± 0.033	**0.868 ± 0.048**	**0.935 ± 0.024**
Reflection	0.108 ± 0.126	0.194 ± 0.189	0.679 ± 0.072	0.761 ± 0.073	**0.748 ± 0.044**	**0.818 ± 0.044**
Hot-spot	0.433 ± 0.316	0.564 ± 0.346	0.725 ± 0.119	0.829 ± 0.094	**0.787 ± 0.097**	**0.877 ± 0.069**

The same proposed annotation method but different segmentation methods!

Examples on WEST Images

#56927 Tangential

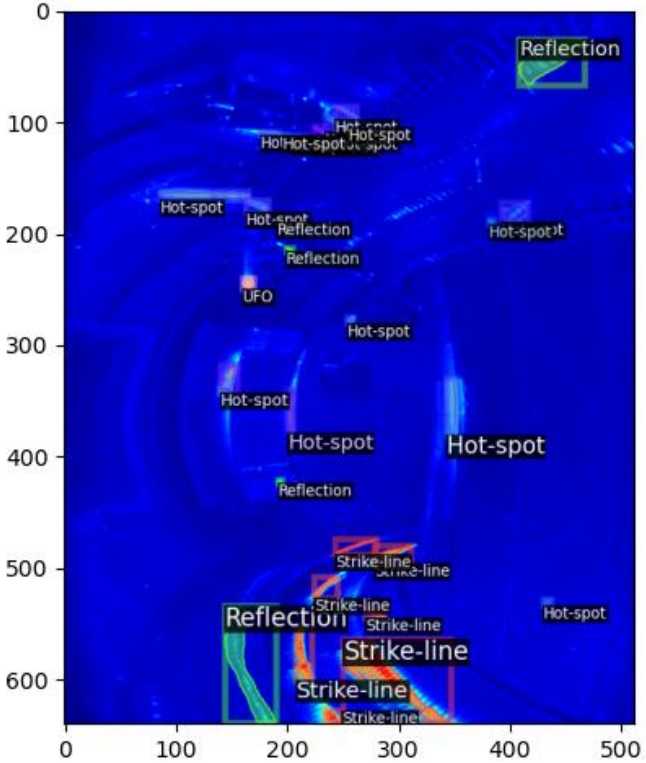

Generated in this work

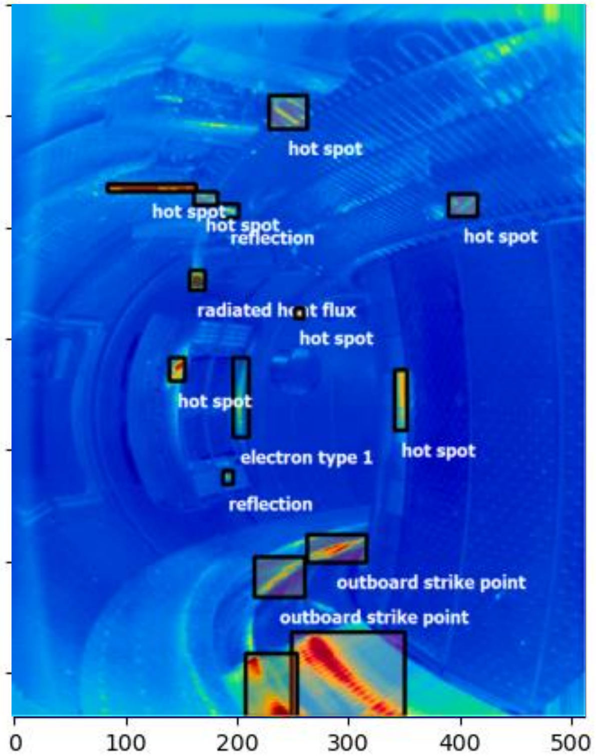

E. Grelier *et al.*, *Fusion Engineering and Design*, 113636, 2023

#55210 Tangential

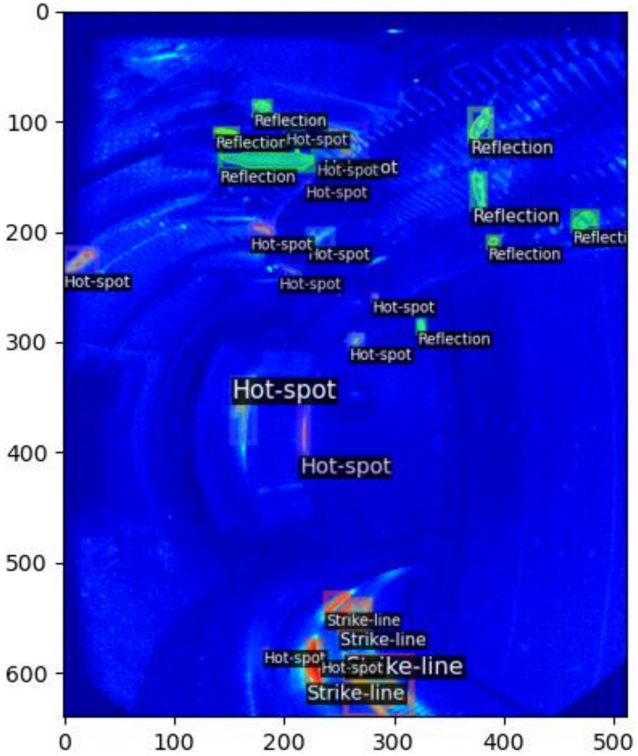

Generated in this work

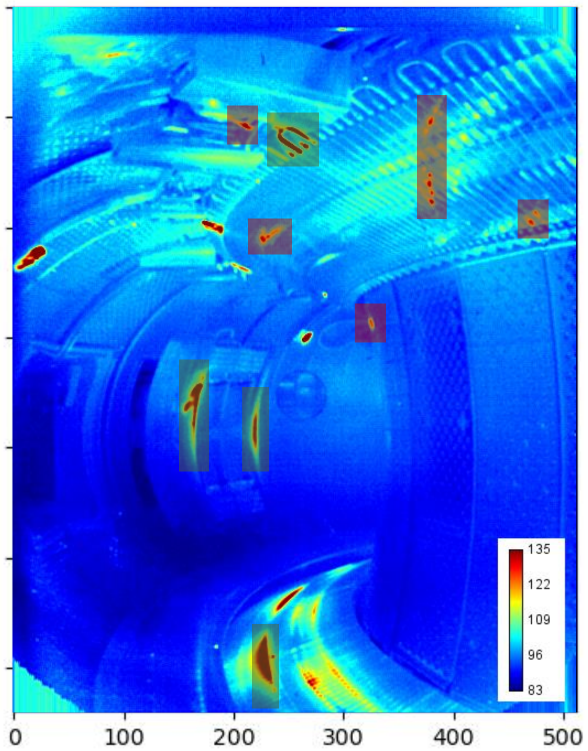

M-H. Aumeunier *et al.*, *IR Calibration Workshop*, 2023

Conclusions: Deterministic image processing method based on a reference image and max-tree representation significantly accelerates the thermal event annotation process by consistently generating annotations of a high degree of similarity to manual annotations in infrared images from nuclear fusion devices.

Annotated Dataset

- A dataset generated with the semi-automatic annotation method from OP1.2 discharge sequences acquired with IRCAM Caleo
- **134** training discharge sequences: **213`883** images, and 3`968 images without any thermal event
- **21** test discharge sequences (sampled every 10^{th} image to reduce the test time): **3`657** images
- Supported annotation formats: COCO and YOLO

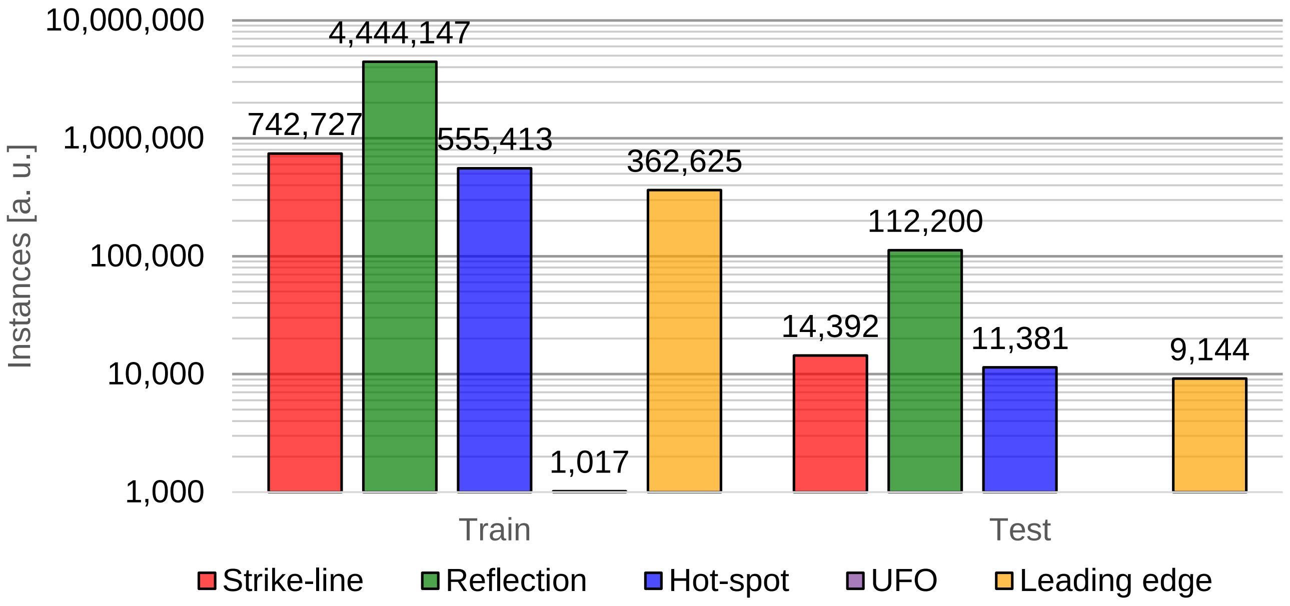

Metallic devices (ITER), compared to graphite devices (W7-X), experience even more reflections

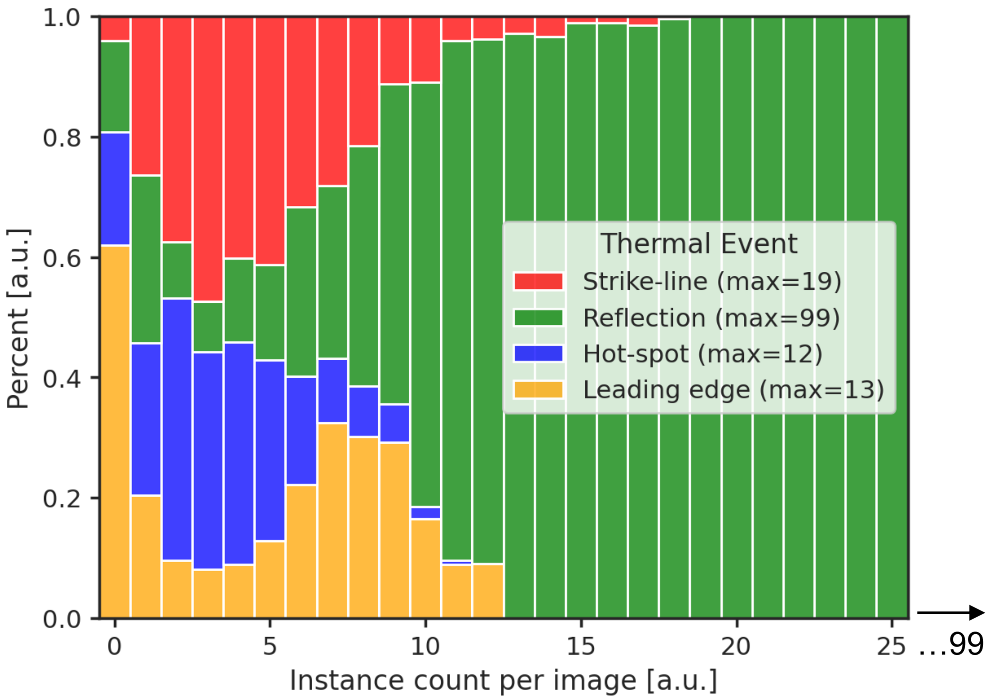

The maximum number of thermal events in an image = 126

Annotated Dataset – Train Split Exploration

Bounding-box centroids | Bounding-box dimensions

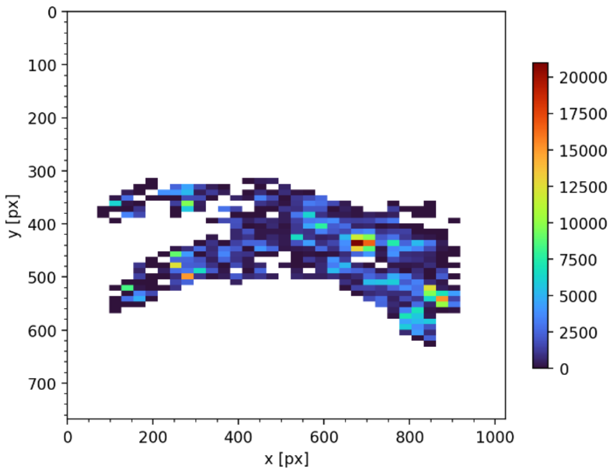

Strike-line ↑

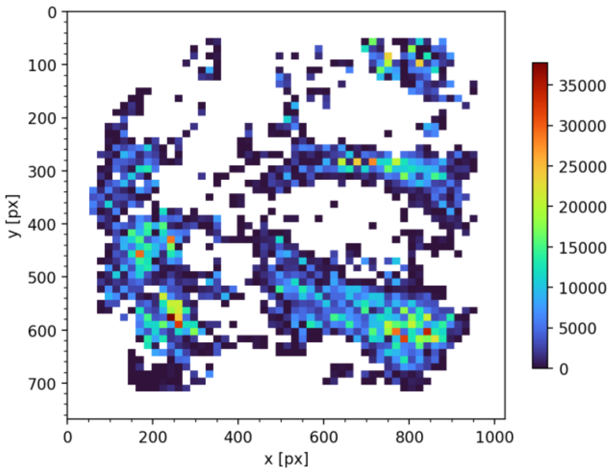

Reflection ↑

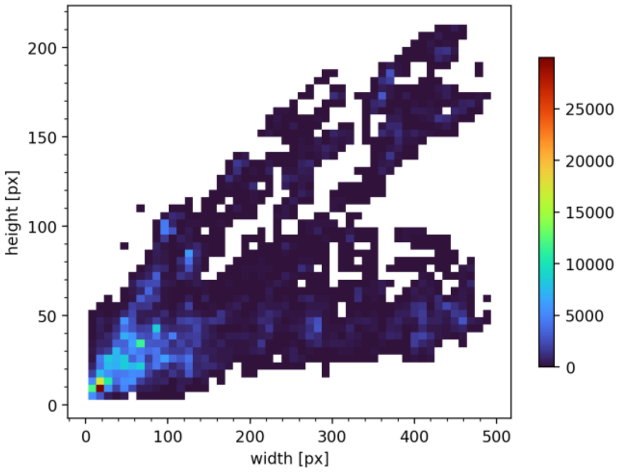

Strike-line ↑

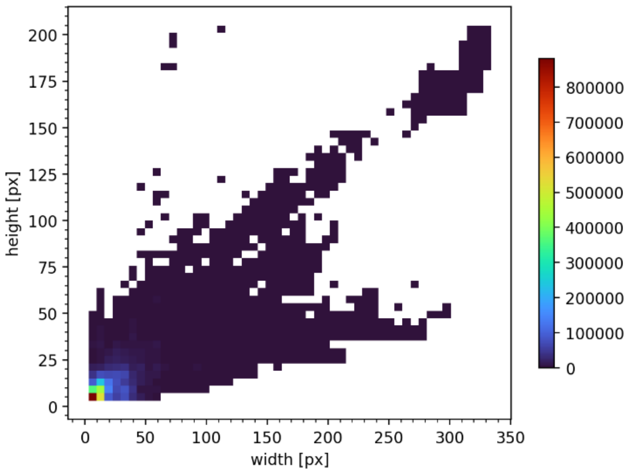

Reflection ↑

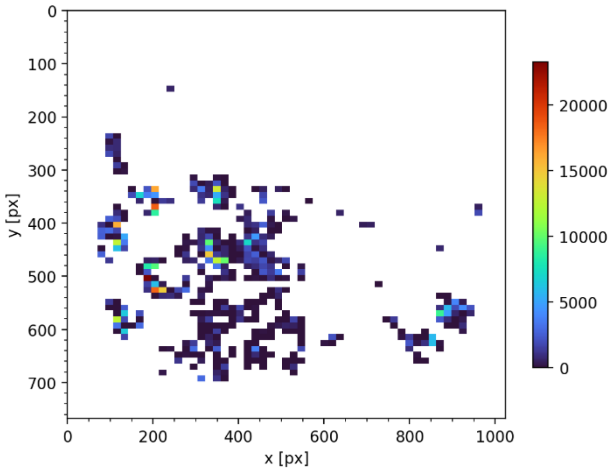

Hot-spot ↑

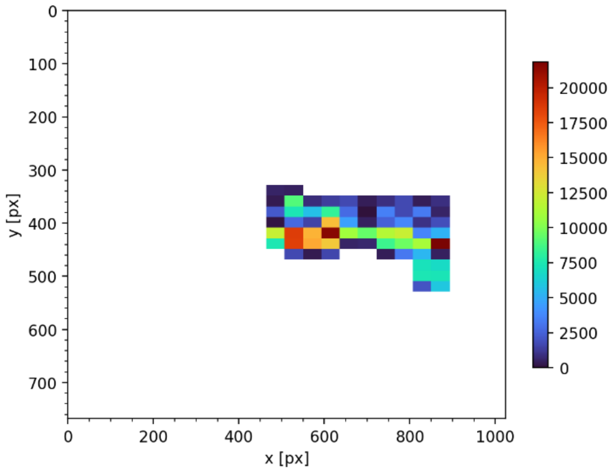

Leading edge ↑

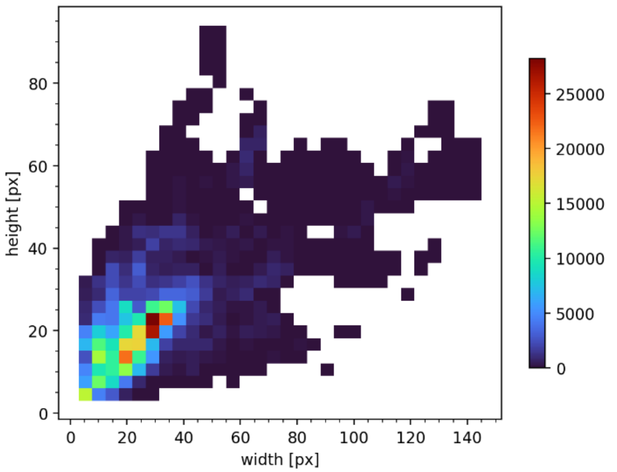

Hot-spot ↑

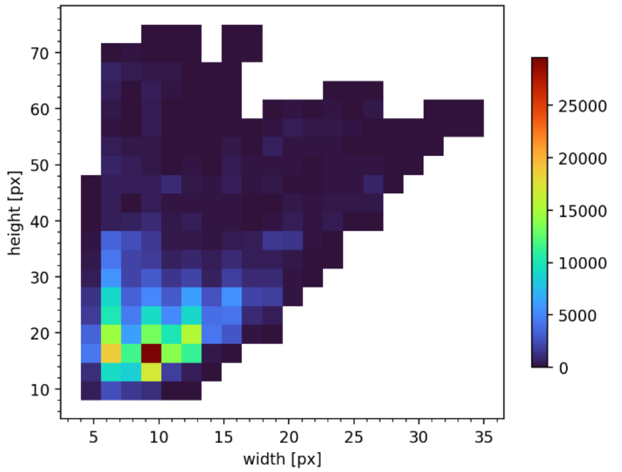

Leading edge ↑

Instance Segmentation – Qualitative Results

20180904.007 (AEF10), high-iota (FTM) configuration

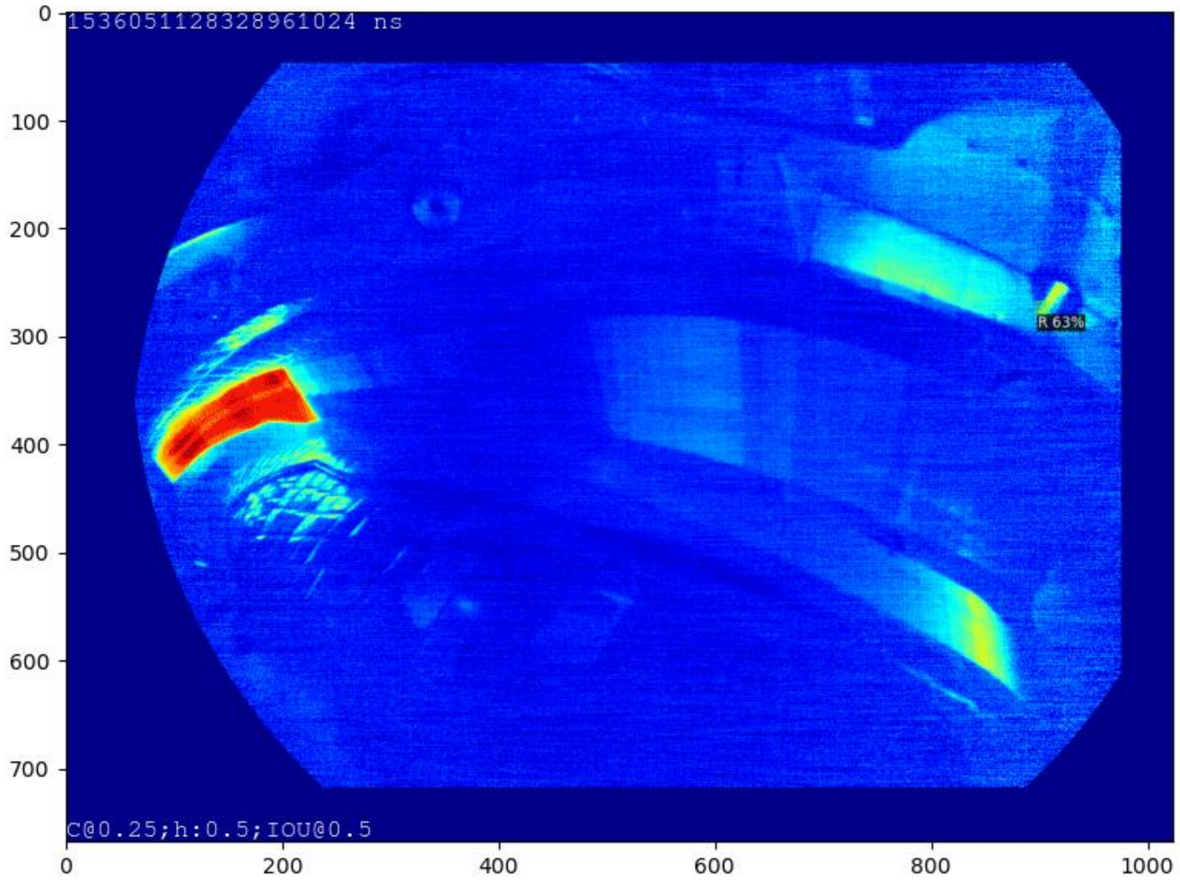

20181017.038 (AEF10), standard (EJM) configuration

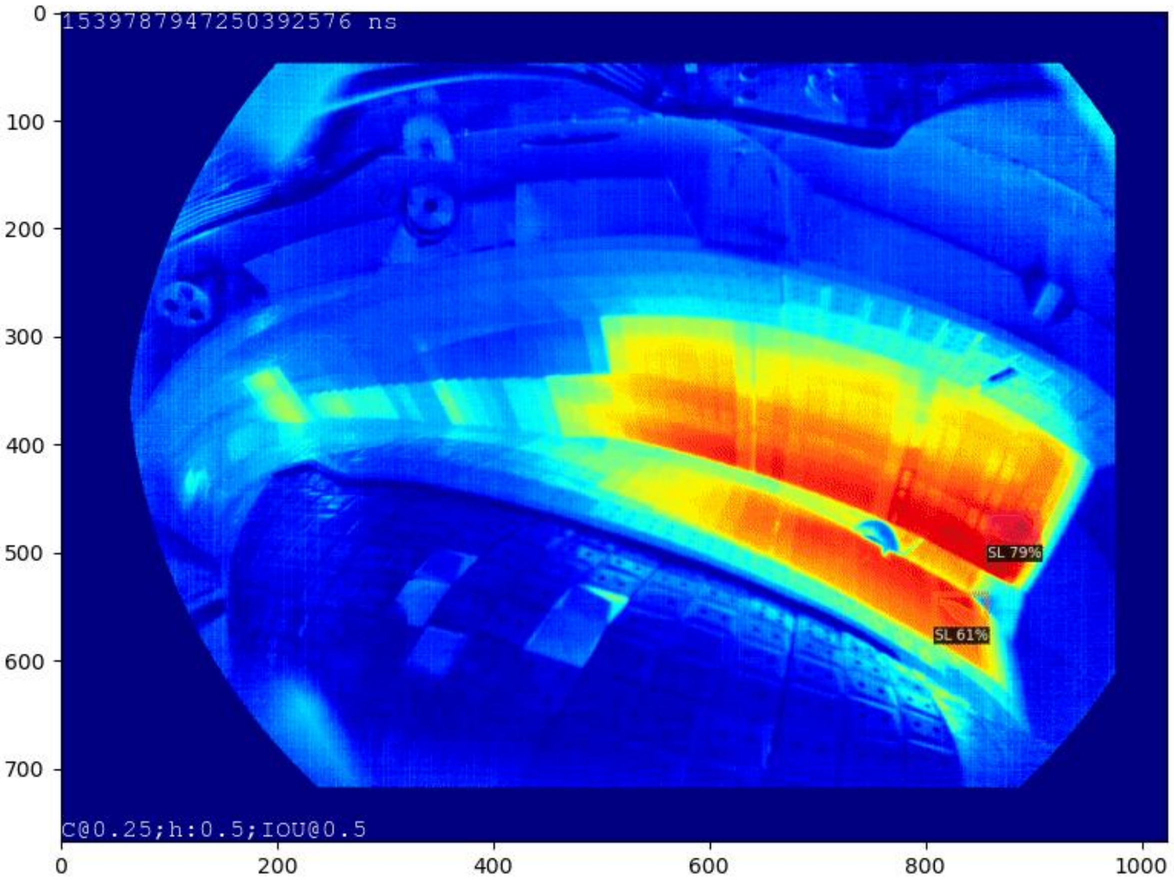

*Animated

Scratch Mask R-CNN (R50, GN): T1 (heating start) → T4 (heating termination), visualize every 5th image

Instance Segmentation – Qualitative Results

20180829.040 (AEF51), low-iota (DBM) configuration

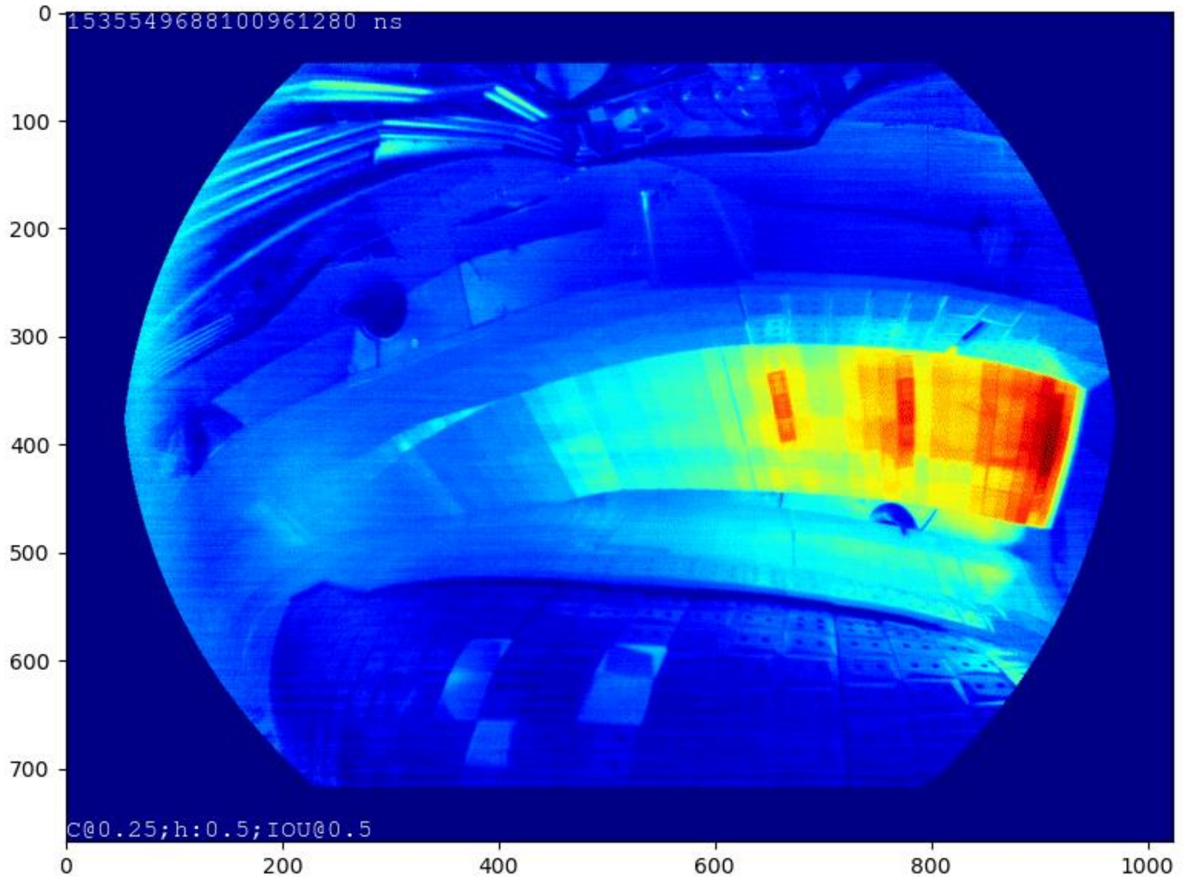

20181004.032 (AEF51), low-iota (DBM) configuration

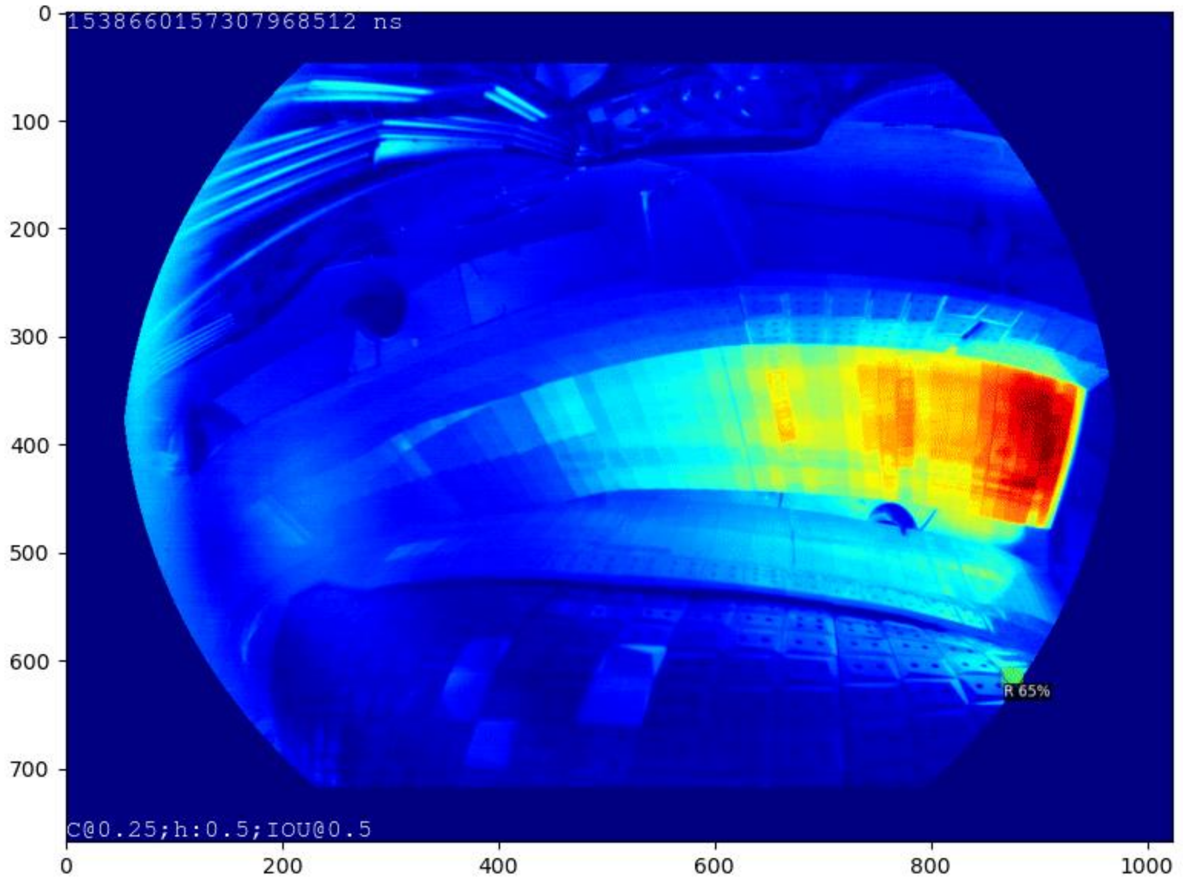

*Animated

Pretrained YOLOv8 (large): T1 (heating start) → T4 (heating termination), visualize every 5th image

Real-Time Instance Segmentation – Quantitative Results

#	Model	Pretrained	Inference shape [C×H×W]	#Params	Bounding-Box		Mask		TensorRT [ms] end-to-end
					mAP	AP@50	mAP↓	AP@50	
1	Cascade Mask R-CNN (R18, BN)	✗	1×768×1024	58.7 M	**45.97**	74.32	**43.67**	**75.69**	-
2	Mask R-CNN (R50, GN)	✗	1×768×1024	45.3 M	44.05	**74.41**	43.38	75.29	93.13
3	Mask R-CNN (R18, GN)	✗	1×768×1024	32.3 M	43.79	74.21	42.88	75.56	69.31
4	Mask R-CNN (R18, GN→BN)	✗	1×768×1024	32.3 M	43.67	73.57	42.86	75.30	44.07
5	Mask R-CNN (R50, BN)	✗	1×768×1024	44.0 M	43.37	73.37	41.87	74.30	40.53
6	Mask R-CNN (R50, BN)	✓	3×768×1024	44.0 M	42.85	72.41	40.70	73.31	40.53
7	YOLOv8 (medium)	✓	3×1024×1376	27.2 M	43.58	70.26	34.75	70.46	28.82
8	YOLOv8 (small)	✓	3×1024×1376	11.8 M	43.27	70.65	34.59	71.15	16.25
9	YOLOv8 (nano)	✓	3×1024×1376	**3.3 M**	42.42	70.63	34.25	70.69	11.12
10	YOLOv8 (nano)	✓	3×768×1024	**3.3 M**	40.14	67.68	29.72	66.94	**8.76**
11	MaskDINO (R50, DETR)	✓	3×1024×1024	43.8 M	38.05	73.43	33.50	73.69	-

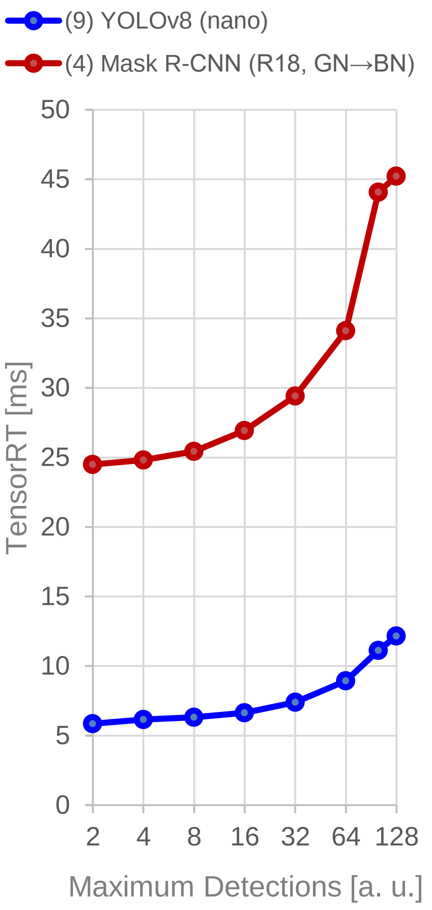

Evaluation criteria are maximum detections = 100, minimum confidence = 5% and the inference times include pre- and post-processing on NVIDIA Tesla T4.
YOLOv8 models offer faster inference in exchange for a lower segmentation performance (especially localization of small events).

R50/18: ResNet 50/18 backbone
GN: Group Normalization (PyTorch → ONNX supports *opset*<18; therefore, GN = Reshape + Instance Normalization + Reshape + Mul + Add or TRT plugin)
BN: Batch Normalization (BN + Conv → Conv (folding))
DETR: Detection with Transformers

Annotated Dataset – Thermal Event Temperature

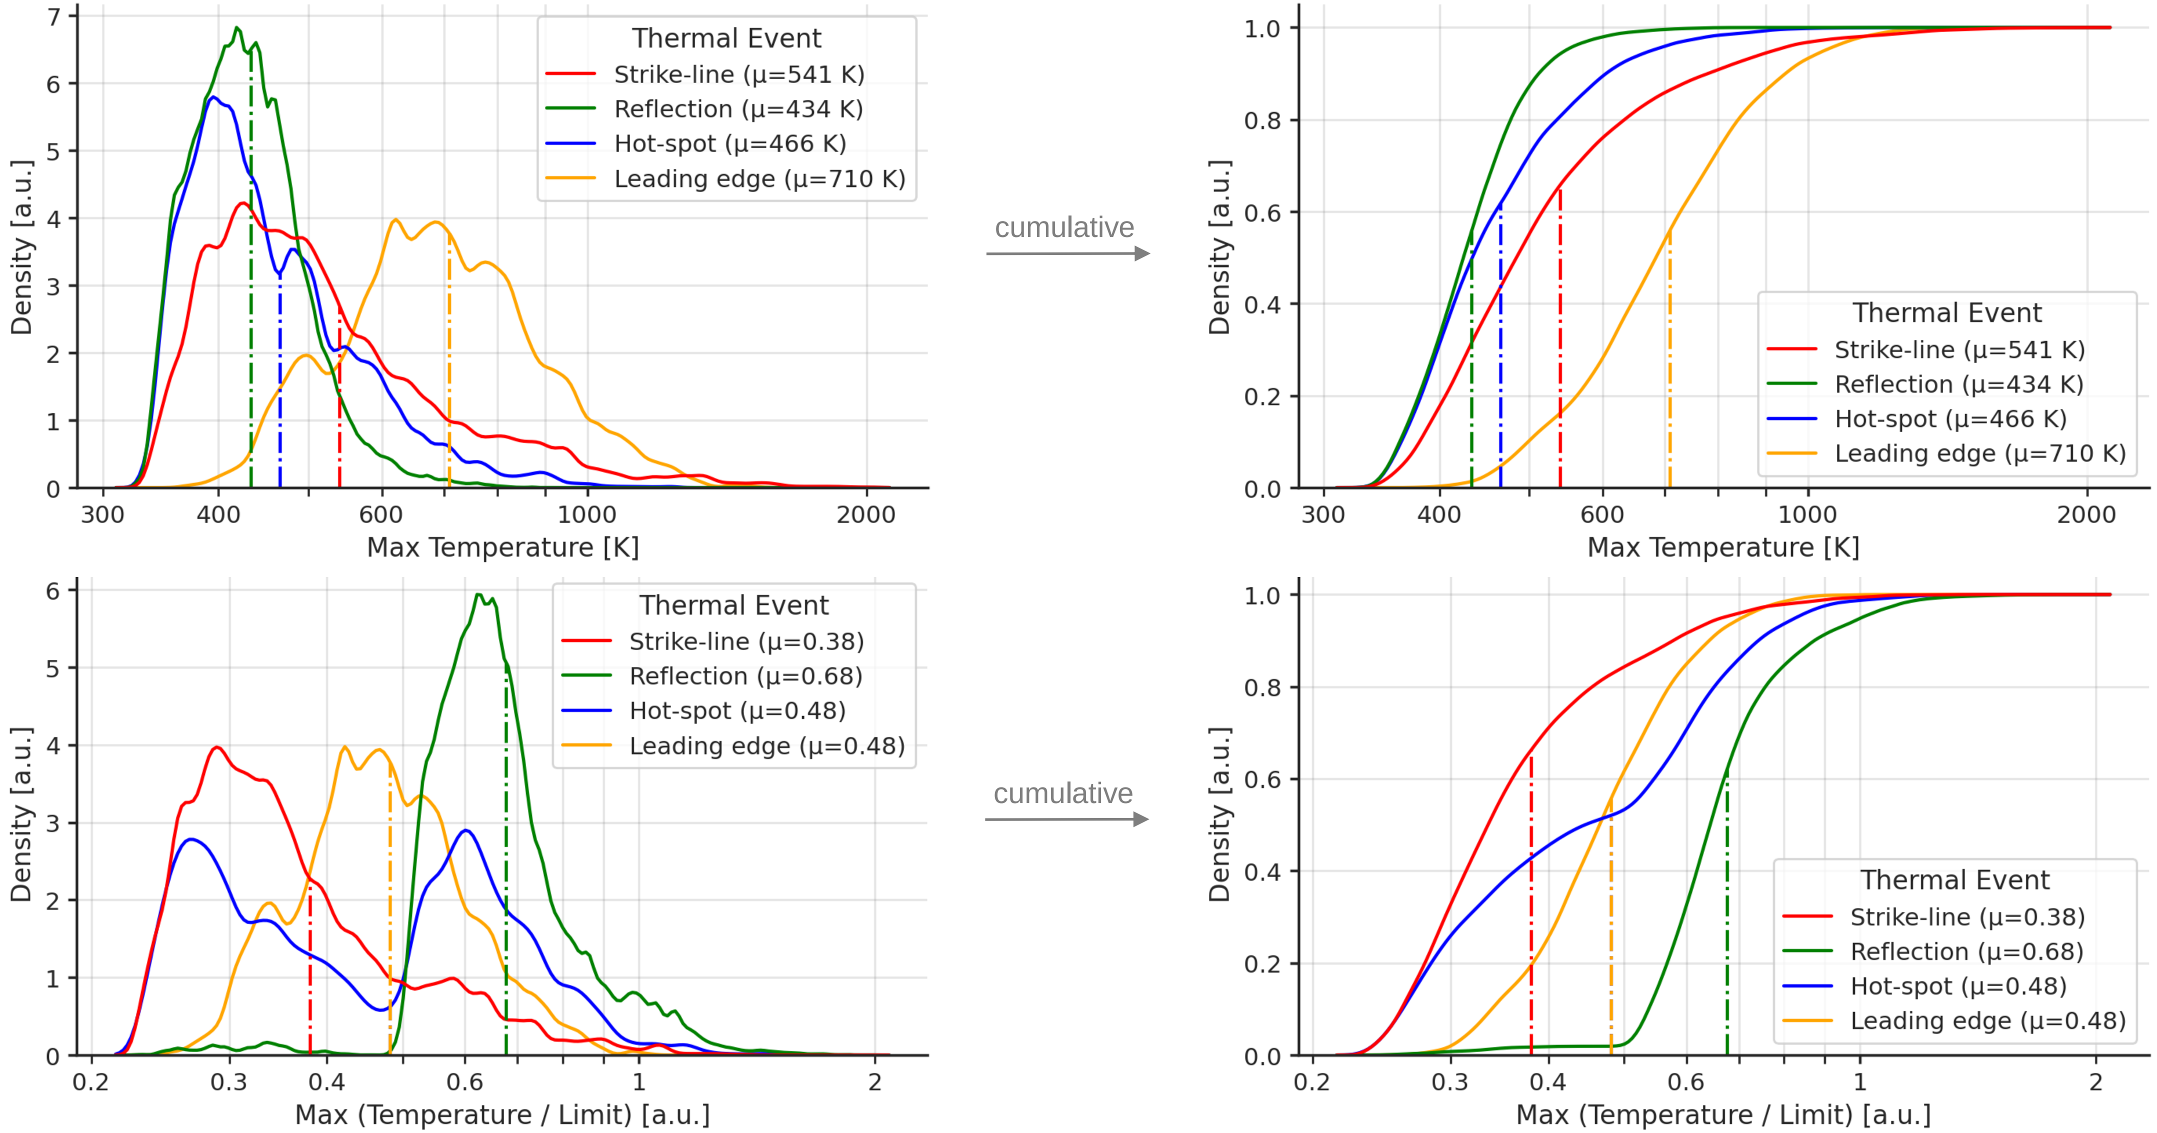

Machine Protection – Non-Dangerous Overload Suppression

Overload (o) = region of pixels exceeding the corresponding temperature limit in image (S)
ND = non-dangerous (positive class), D = dangerous (negative class), Ø = no matching detection

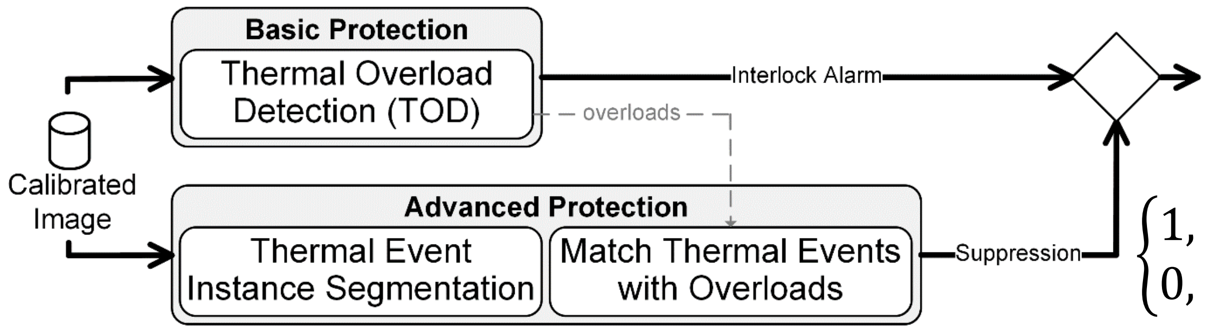

$$\begin{cases} 1, & \text{if } \forall o \in S, o \in ND \quad \text{(suppress)} \\ 0, & \text{otherwise.} \end{cases}$$

ND	reflection, leading edge	D	strike-line, hot spot, UFO

Ground Truth	Prediction	Outcome
ND	ND	TP: good, supressed unnecessary alarm
D	ND	FP: bad, supressed dangerous alarm
ND	D	FN, not supressed unnecessary alarm
D	D	TN, not suppressed dangerous alarm
ND	Ø→D	FN, not supressed unnecessary alarm
D	Ø→D	TN, not supressed unnecessary alarm

Statistics	
Overload OP1.2 sequences	640
Excl. training sequences	(-32) 608
Only NDs in image	(**20.23%**) 123
Only Ds in image	303
Total reflections (ND)	(62%) 1005
Total strike-lines (D)	327
Total hot spots (D)	261
Total leading edges (ND)	(1%) 17

Machine Protection – Non-Dangerous Overload Suppression Evaluation

↑ Precision (P) $= \dfrac{TP}{TP + FP}$ (TP + FP = number of predictions) ↑ Recall (R) $= \dfrac{TP}{TP + FN}$ (TP + FN = number of ground truth annotations)

↑ Average Precision (AP) = $\sum_n (R_n - R_{n-1}) P_n$, for n confidence thresholds

↑ R@P(1) = recall at precision of 1 (no FPs)

Model	Per Image			Per Overload		
	AP	R@P(1)	Suppressions ↓	AP	R@P(1)	Suppressions
afbb8a1 (1)	1.000	1.000	123	1.000	0.988	1010
662ca1e (4)	0.974	0.967	119	0.998	0.976	997
35b16da (5)	0.986	0.911	112	0.999	0.985	1007
94f2e90 (3)	0.947	0.894	110	0.996	0.761	778
c633384 (6)	0.971	0.805	99	0.998	0.894	914
efef13d (2)	0.983	0.724	89	0.999	0.801	819
e1412b9 (7)	0.994	0.992	122	0.999	0.850	869
c2ba23e (10)	0.981	0.976	120	0.996	0.408	417
683d08f (9)	0.998	0.870	107	0.999	0.748	764
e5a3a55 (8)	0.994	0.715	88	0.996	0.633	647

Exemplary overloads

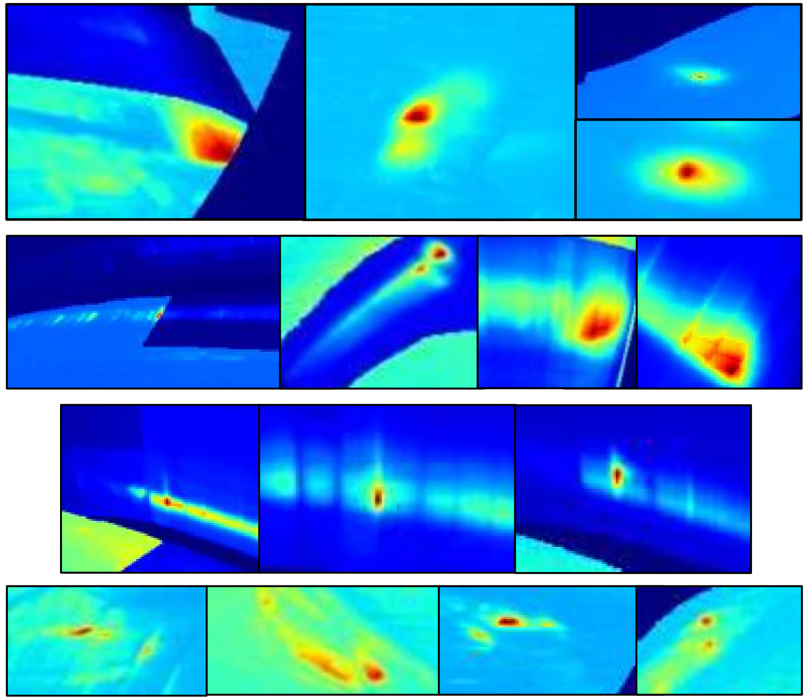

Mask R-CNN

YOLOv8

Exemplary false positive

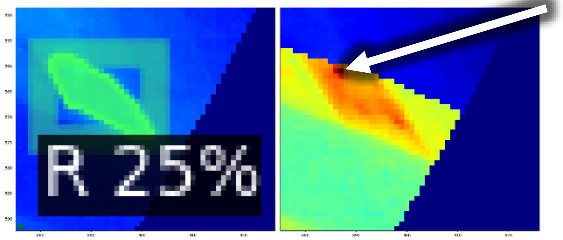

Summary

Presented varied applications of computer vision for machine protection in the nuclear fusion device – Wendelstein 7-X:

I. Real-time **overload** anticipation in an entire FoV with dynamic thresholds…
→ protect PFCs from overheating

II. Real-time **anomalous image detection**…
→ verify if images are correct so that PFCs can be protected with [I]

III. Offline semi-automatic infrared discharge **annotation method**…
→ facilitate supervised deep learning in [IV]

IV. Real-time **thermal event instance segmentation**…
→ scene understanding and non-dangerous overload suppression to reduce alarms in [I]

Computer Vision for Machine Protection in Nuclear Fusion

8th EIROforum School on Instrumentation
13th – 17th May 2024

Bartłomiej Jabłoński, Dariusz Makowski
Lodz University of Technology (TUL), Lodz, Poland

Aleix Puig Sitjes, Marcin Jakubowski
Max Planck Institute for Plasma Physics (IPP), Greifswald, Germany

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), and in part by the Engineering Grant EEG21-17, and in part by the Polish Ministry of Science and Higher Education within the International Co-financed Projects (PMW). 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.

Risk Model

(1) $\mu_{Temperature}$

$$\boxed{\overline{T'_{xy}}(t)} = \frac{T_{xy}(t-[n-1]t_{frame}) + T_{xy}(t-[n-2]t_{frame}) + \cdots + T_{xy}(t)}{n}$$

(2) $\Delta_{Temperature}$

$$\boxed{\Delta\overline{T'_{xy}}(t)} = (1-\alpha)\boxed{\Delta\overline{T'_{xy}}(t}-t_f) + \alpha\left(\boxed{\overline{T'_{xy}}(t)} - \boxed{\overline{T'_{xy}}(t}-t_f)\right)$$

(3) q (heat-flux) estimate

$$\boxed{q_{xy}(t)} = \frac{\boxed{\Delta\overline{T'_{xy}}(t)}}{\cancel{\left(\frac{t_f}{C^m_{xy}}\right)^{1}}}$$

(4) Threshold

$$\boxed{T^{th}_{xy}(t)} = \min\left\{T^{limit}_{xy} - \boxed{q_{xy}(t)}\cancel{\left(\frac{t_r}{C^m_{xy}}\right)^{1}}, T^{th\_max}_{xy}\right\}$$

(5) Risk

$$\hat{r}_{xy}(t) = \frac{\boxed{\overline{T'_{xy}}(t)}}{\boxed{T^{th}_{xy}(t)}}(1-S_{xy})$$

1. Full-frame processing (1280×1024) instead of Regions of Interest (RoIs)
2. Dynamic temperature threshold instead of a fixed threshold

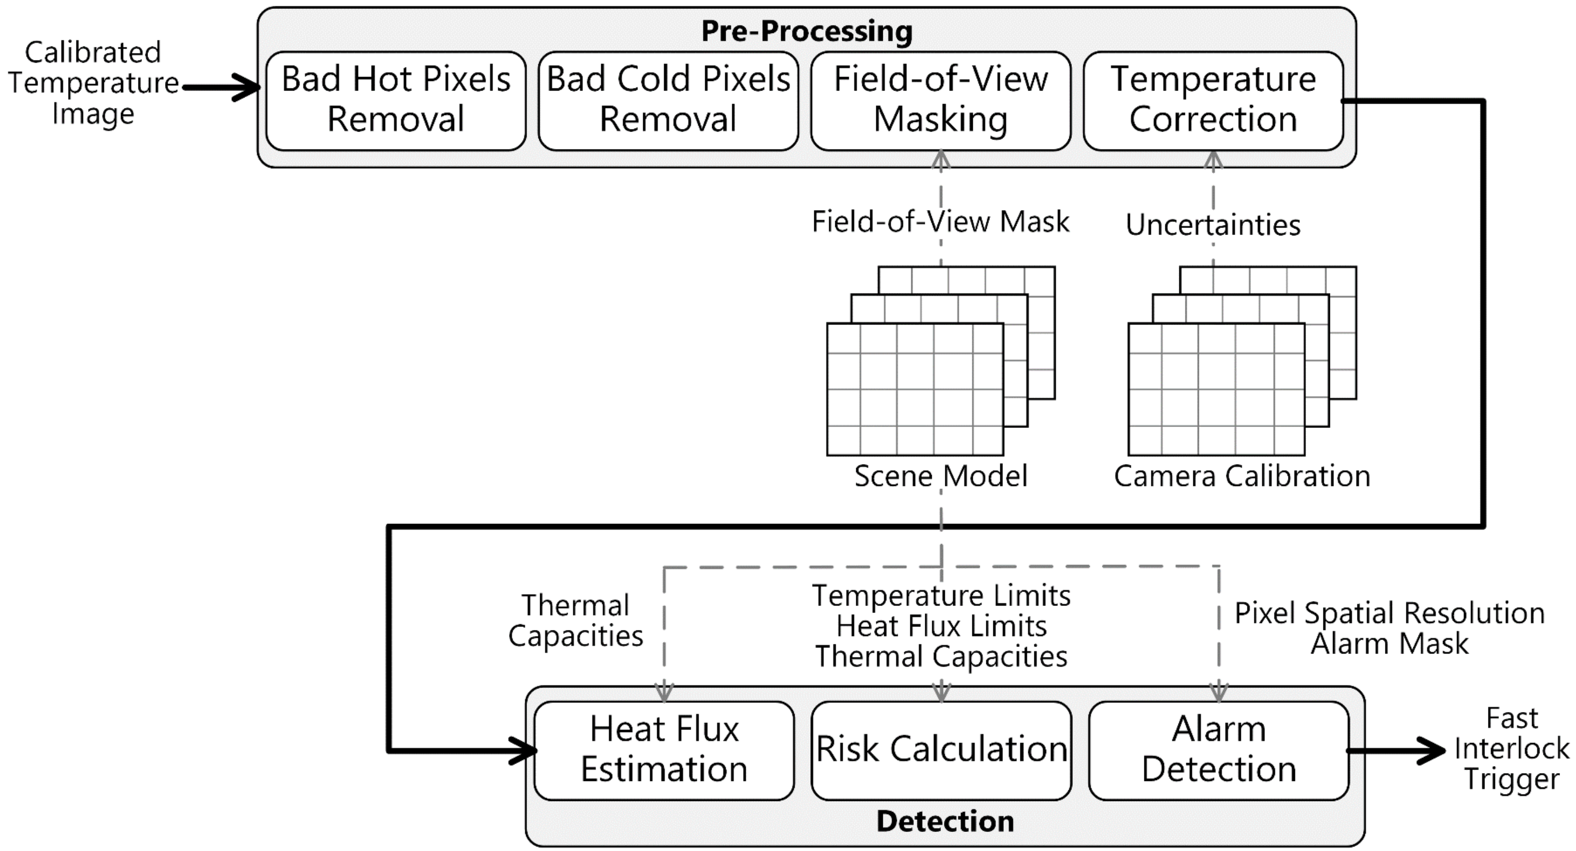

Processing Performance

- Real-time constraint is **110 ms** = 160 ms – 50 ms (safety margin)
- Acquisition rate is **100 Hz**, a new image every 10 ms
- CoDaS Fast Control Station (FCS) triggers the next processing step every 10 ms
- AMD EPYC 7402P 24-core **CPU** and NVIDIA Tesla T4 **GPU**
- 1-channel **16-bit 1280×1024** IR images

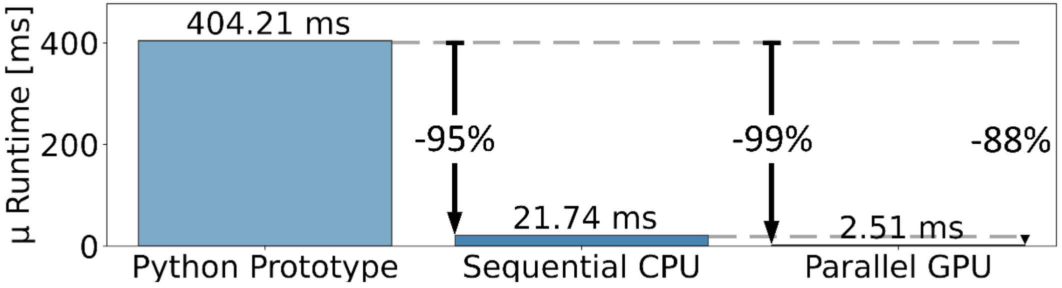

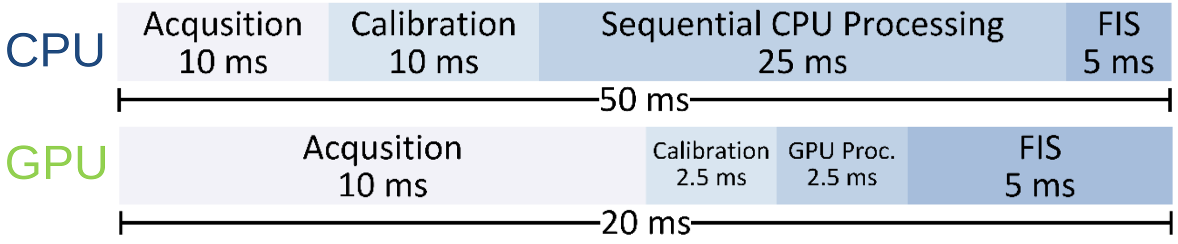

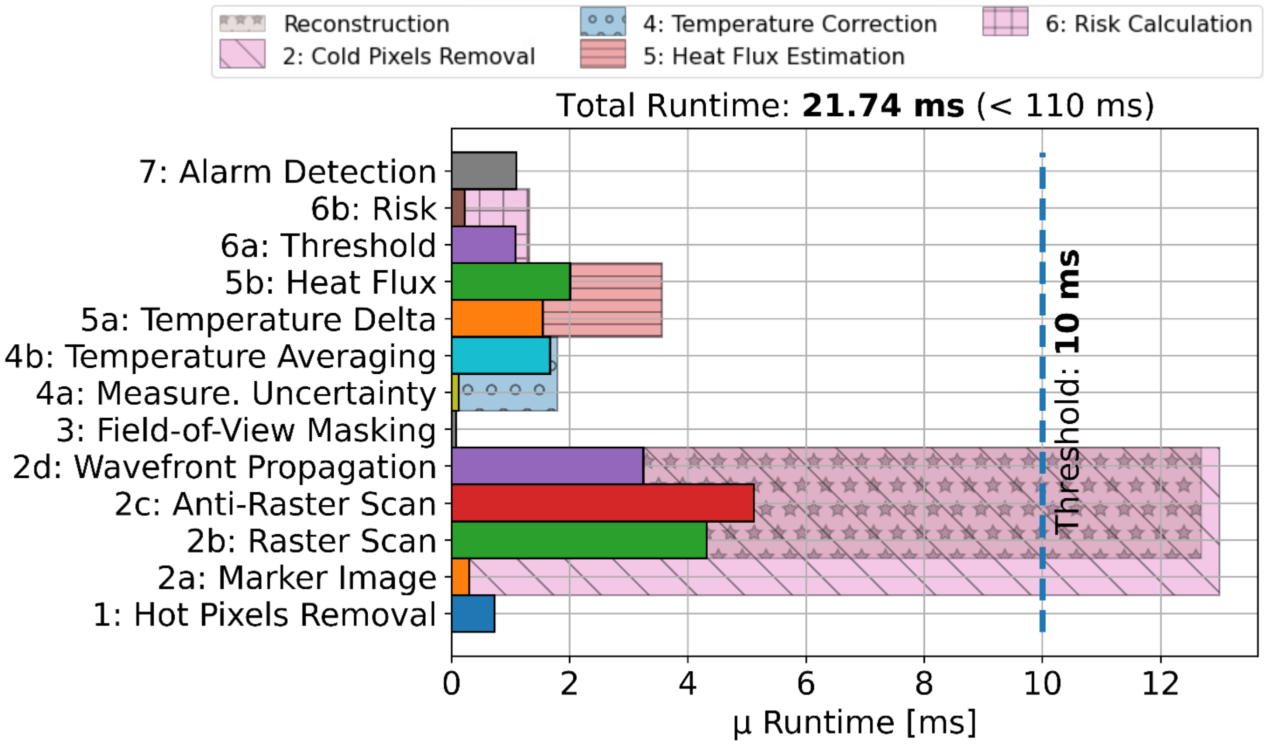

Evaluation Criteria for Binary Output

- During the previous experimental campaign **OP1.2** (2017-2018), no autonomous protection system was in operation, and the inertially-cooled (**uncooled**) test divertors were affected by overloads
- It enables **backtesting**
- During **OP2.1** (2022-2023), PFCs will be **water-cooled** and become **vulnerable** to overloads

Detection Result	**Definition**
True Positive (TP)	The alarm is triggered within $[-1000, -50]$ ms before the overload
True Negative (TN)	NEITHER the alarm NOR the overload occurs
False Positive (FP)	EITHER the alarm is triggered within $(-\infty, -1000)$ ms before the overload OR the alarm is triggered when no overload occurs
False Negative (FN)	The alarm is triggered within $(-50, \infty)$ ms before the overload, i.e., it is too late to compensate for the system delay

Thermal Overload Detection – Evaluation

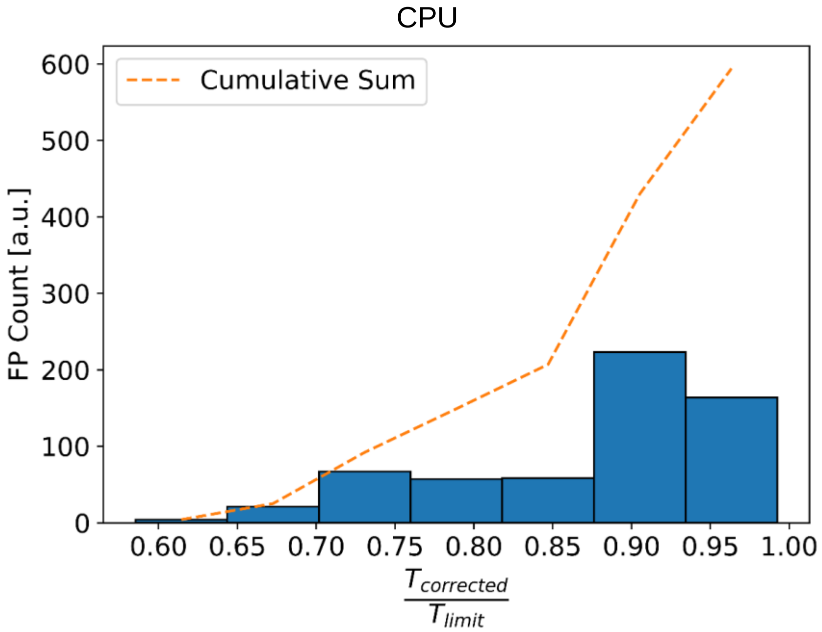

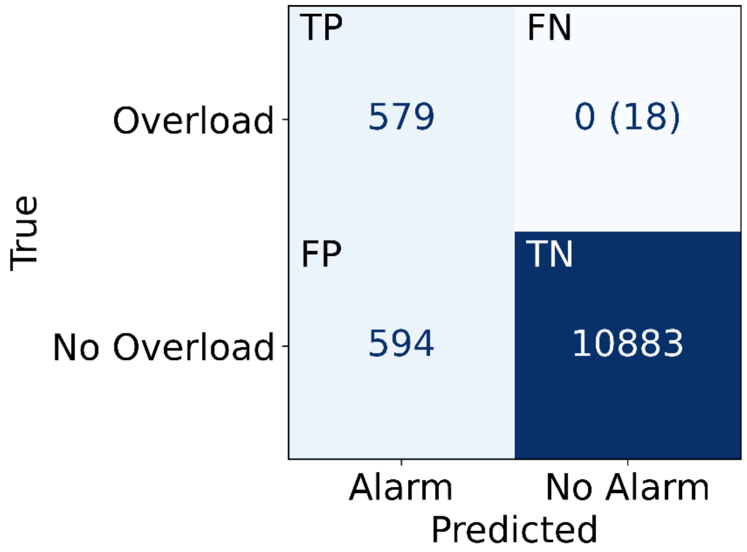

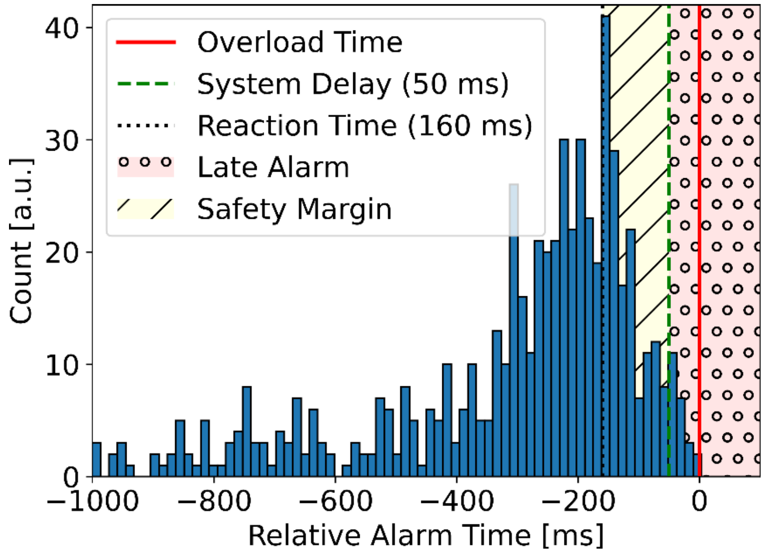

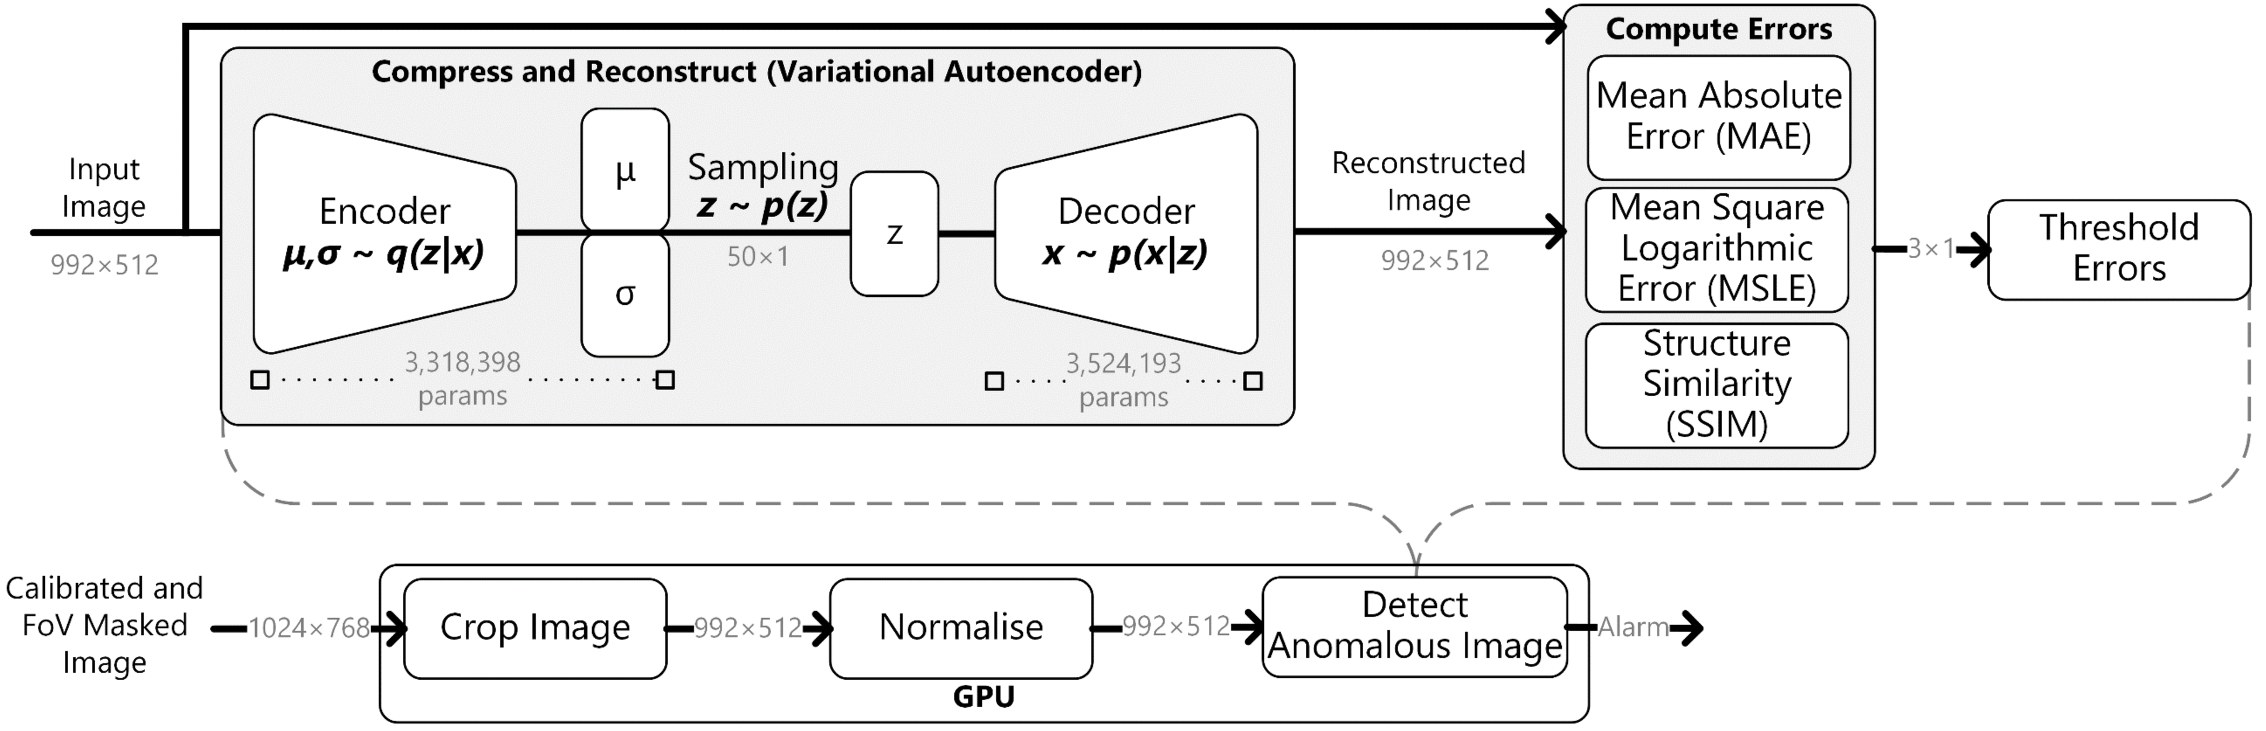

Trained only on correct images!

✗ Incorrect discharge sequence data, i.e., only-noise:

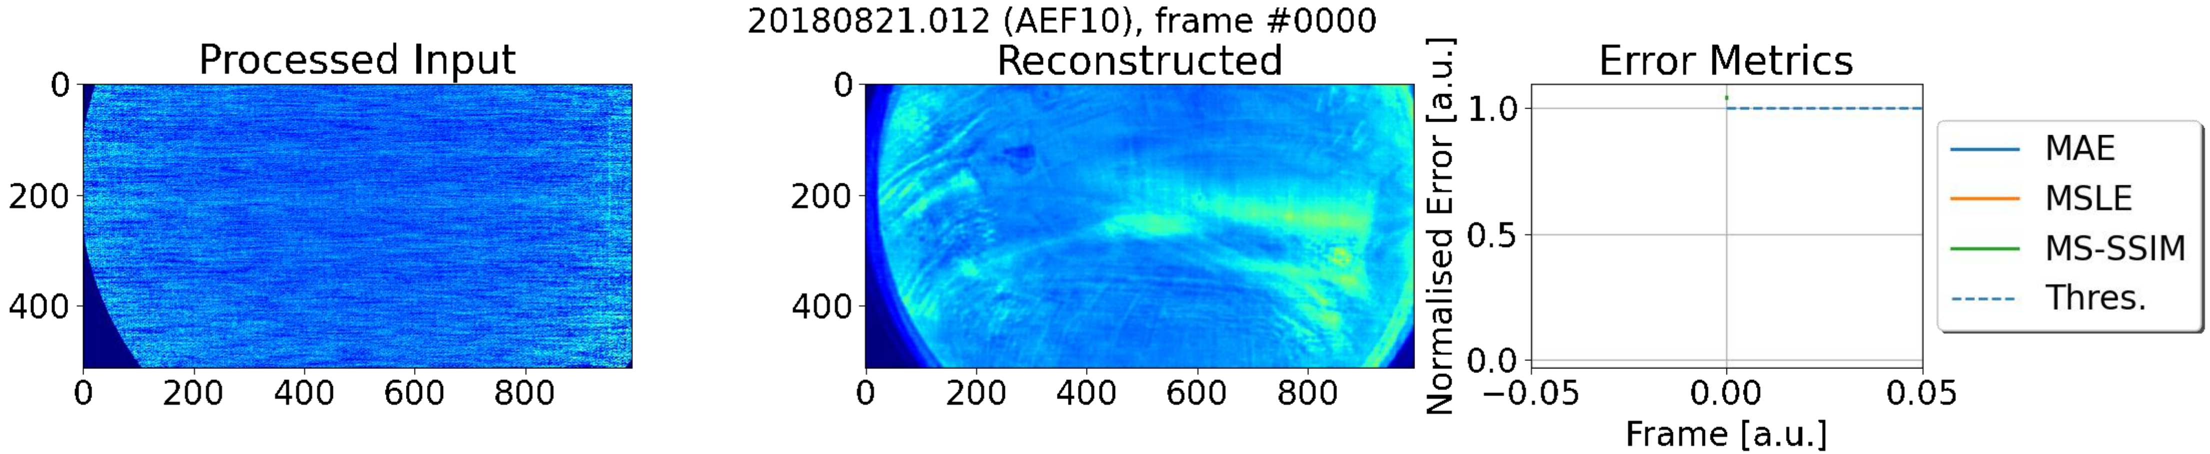

✗→✓ Incorrect to correct discharge sequence data, i.e., plasma build-up transition from noise-only to observable heat loads:

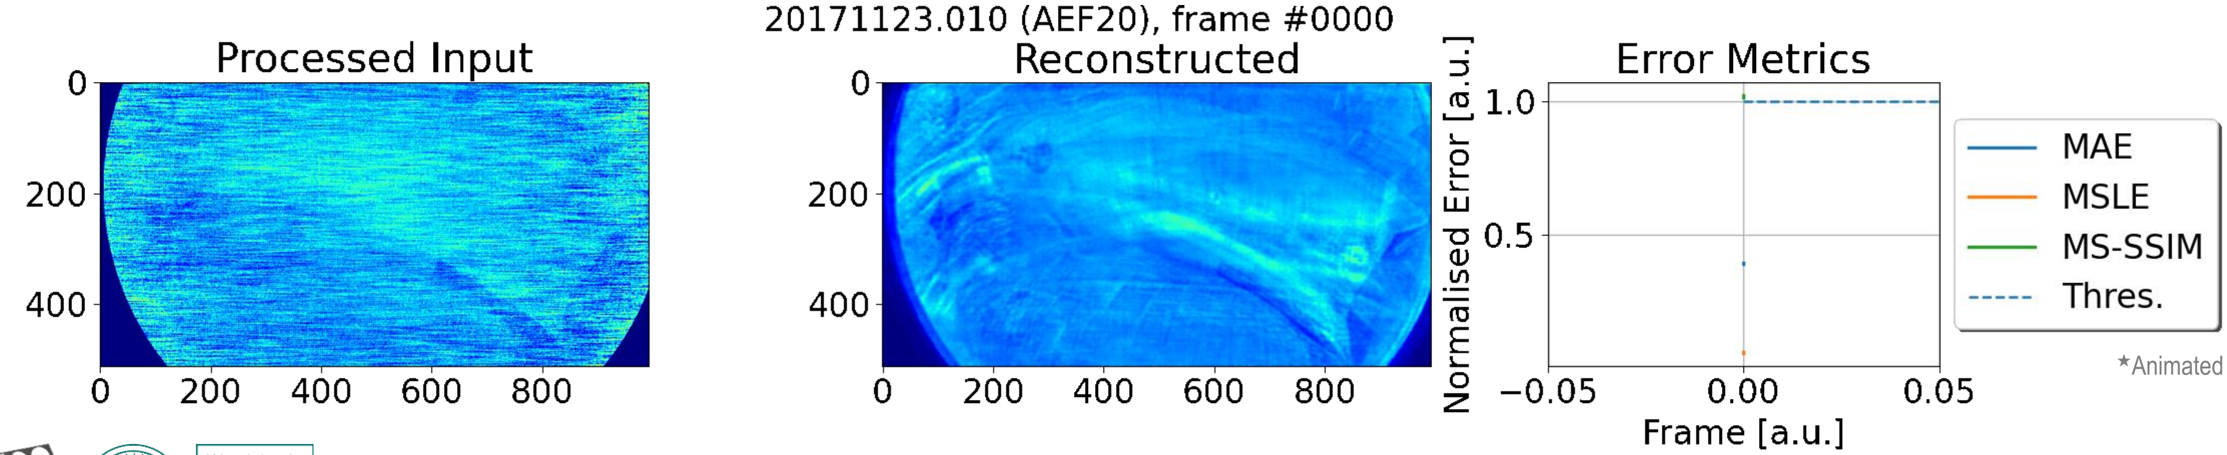

*Animated

Variational Autoencoder Formulas

Prior $p_\Theta(z)$
Posterior $p_\Phi(z|x)$
Likelihood $p_\Theta(z|x)$

$$p_\Phi(z|x) \approx p_\Theta(z)$$

$$-D_{KL}[P_\Phi(z|x) \| P_\Theta(z)]$$

$$z = \mu(z) + \Sigma(z) * \epsilon$$

$$L(\Phi, \Theta, x) = -D_{KL}[P_\Phi(z|x) \| P_\Theta(z)] + E_{q(z|x)}[log(p_\Theta(x|z))]$$

$$D_{KL}(P\|Q) = -\frac{1}{2}\sum_{i=0}^{n} 1 + \log\sigma_{qi}^2 - \mu_{qi}^2 - \sigma_{qi}^2$$

$$MAE(y_{true}, y_{pred}) = \frac{\sum_{i=1}^{n} \left|y_{true}^i - y_{pred}^i\right|}{n}$$

$$MSLE(y_{true}, y_{pred}) = \frac{\sum_{i=1}^{n} \left(\log\left(y_{true}^i + 1\right) - \log\left(y_{pred}^i + 1\right)\right)^2}{n}$$

$$BCE(y_{true}, y_{pred}) = -\frac{\sum_{i=1}^{n} y_{true}^i \log y_{pred}^i + \left(1 - y_{true}^i\right)\log\left(1 - y_{pred}^i\right)}{n}$$

$$SSIM(y_{true}, y_{pred}) = \frac{1}{n}\sum_{i=1}^{n} \frac{\left(2\mu_{true_i}^2\mu_{pred_i}^2 + c_1\right)\left(2\sigma_{true-pred_i} + c_2\right)}{\left(\mu_{true_i}^2 + \mu_{pred_i}^2 + c_1\right)\left(\sigma_{true_i}^2 + \sigma_{pred_i}^2 + c_2\right)}$$

- t-distributed stochastic neighbour embedding

Segmentation Method

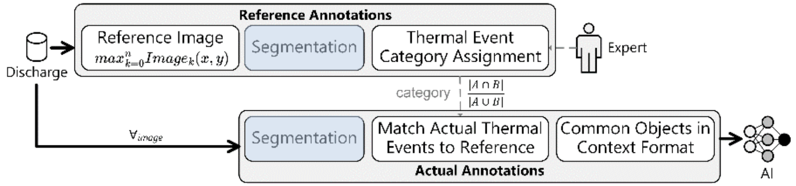

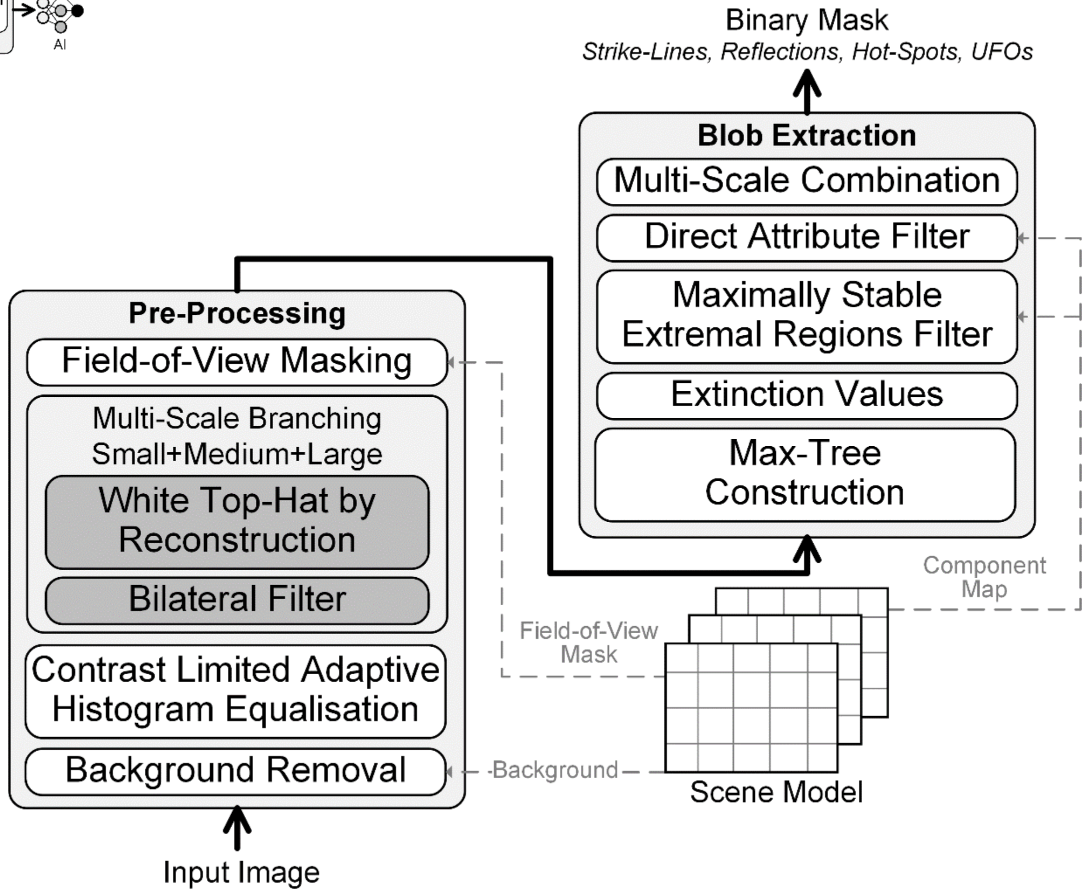

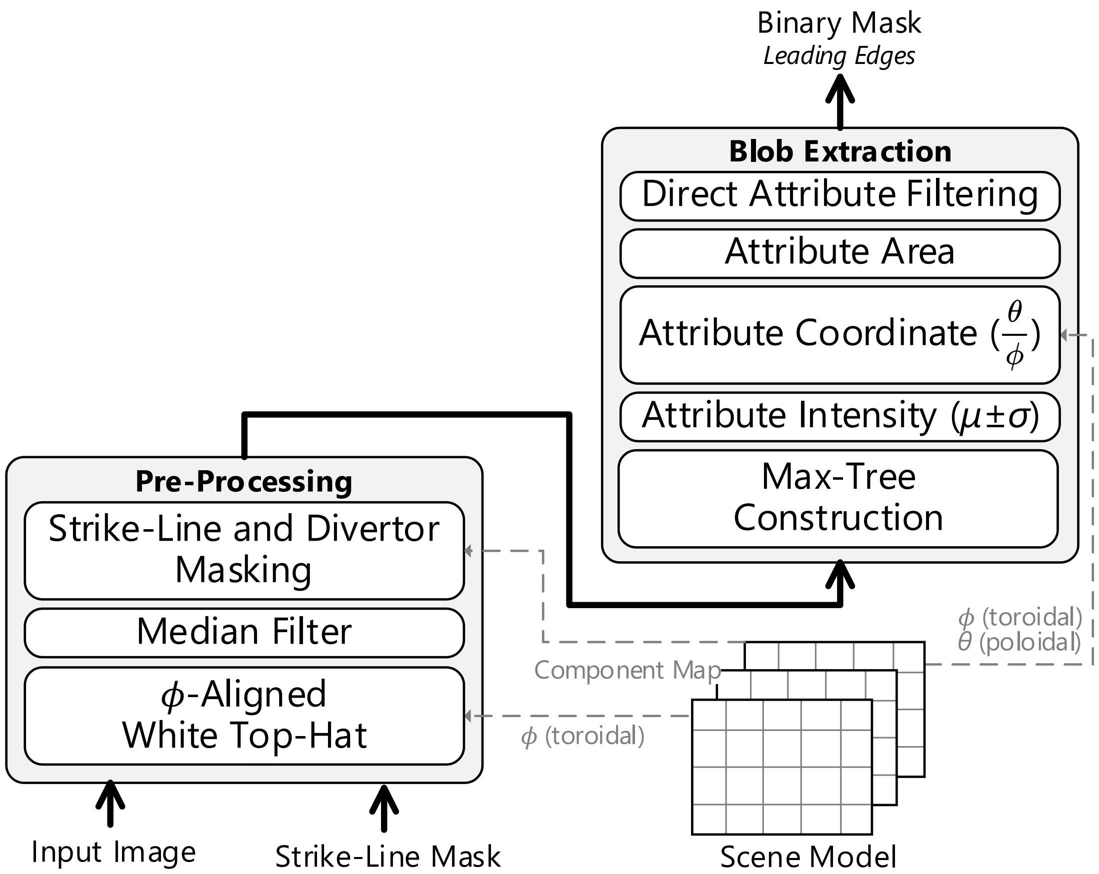

Φ-Aligned White Top-Hat

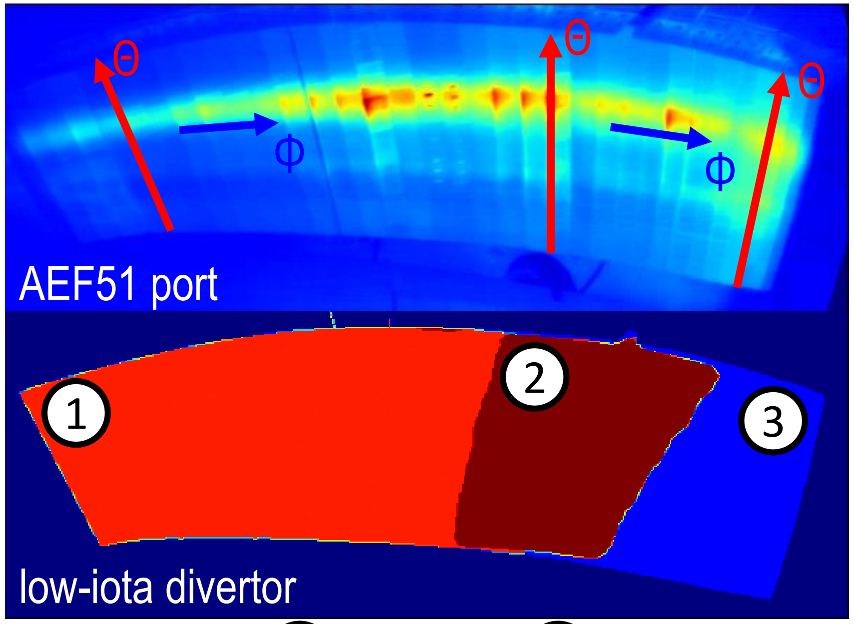

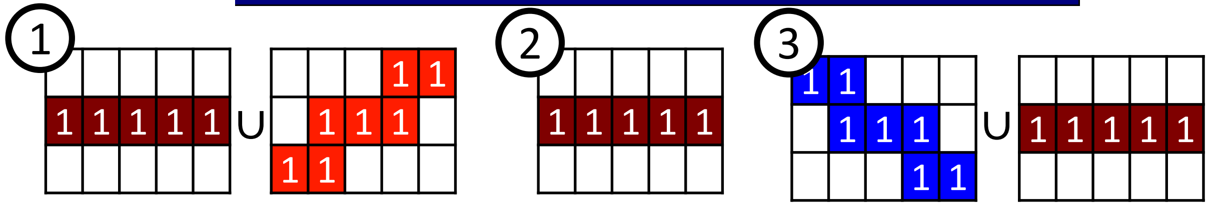

toroidal (**Φ**, phi)

poloidal (**Θ,** theta)

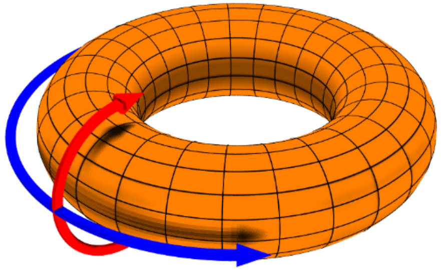

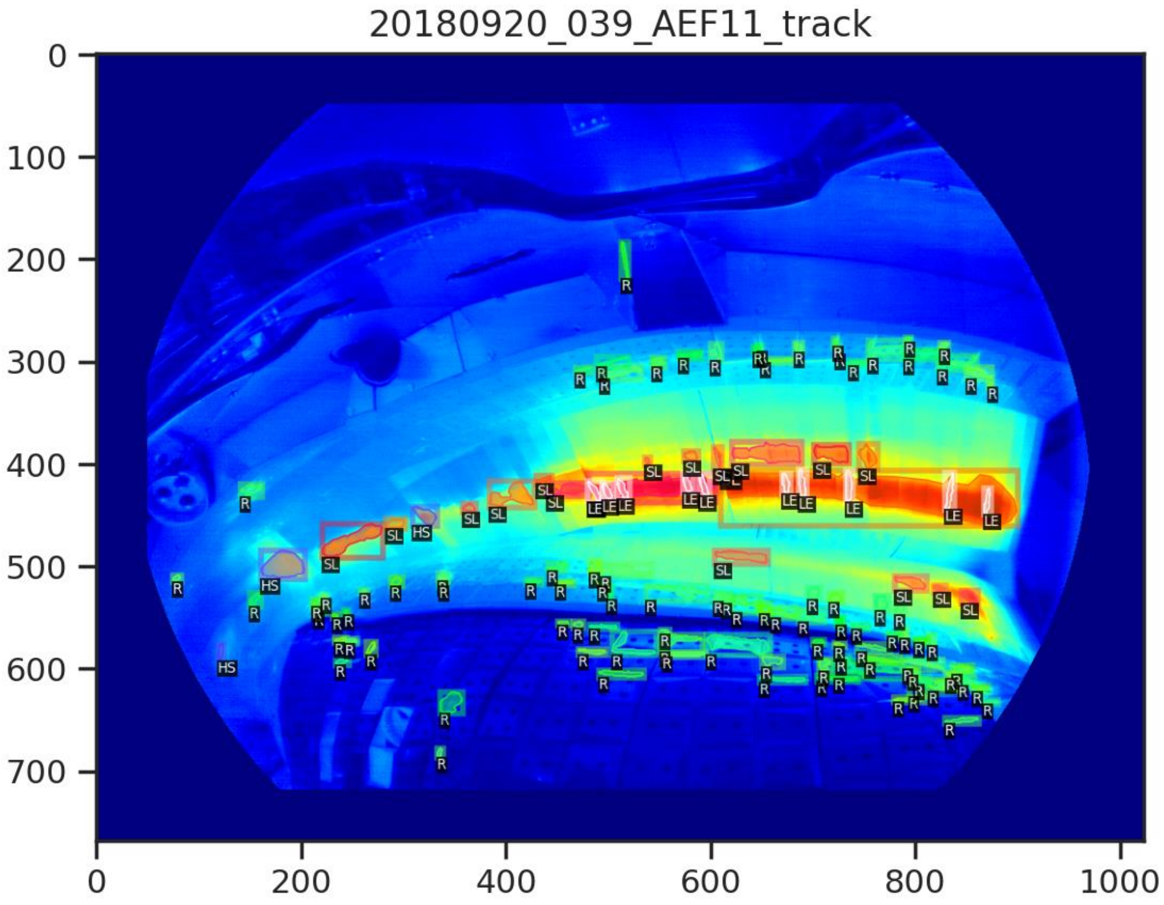

Discharge sequence containing most SL, R, LE and total thermal events.

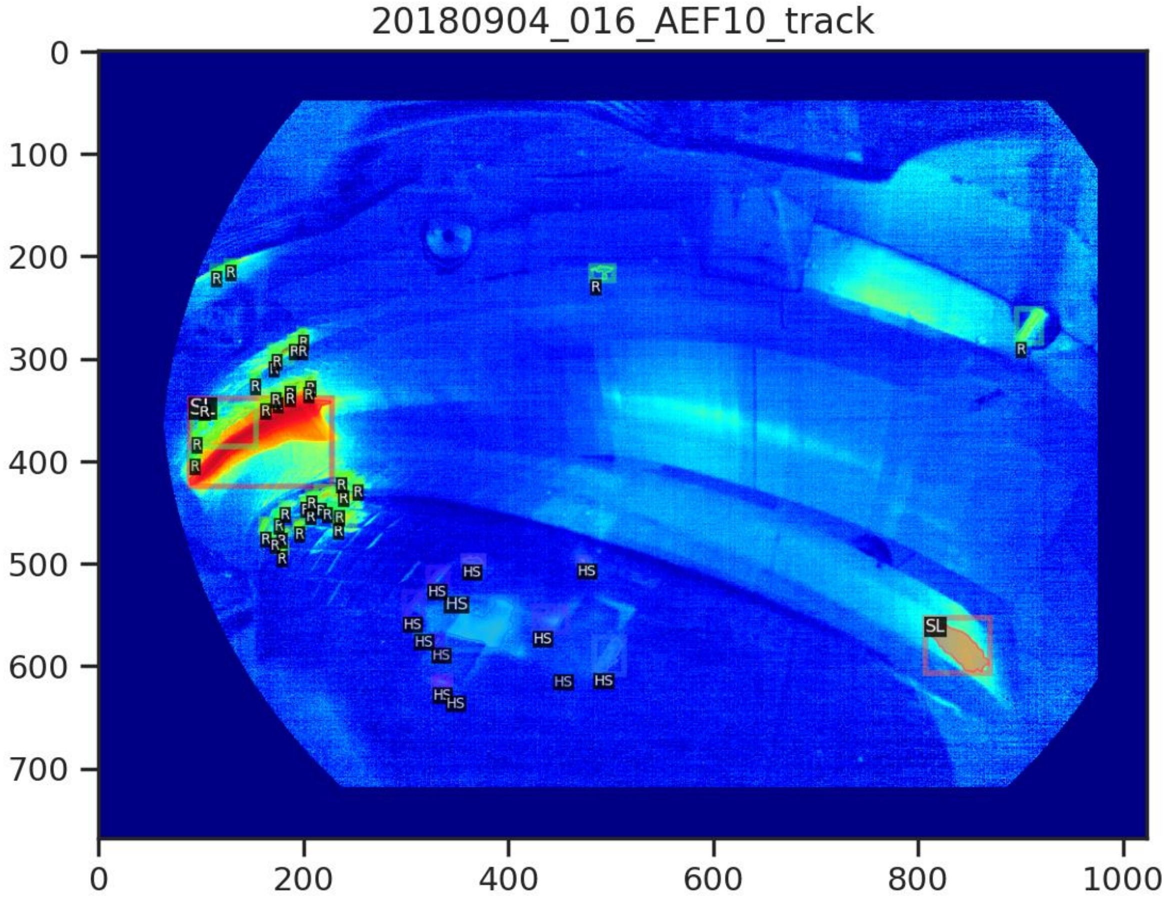

Discharge sequence containing most HS.

Mask R-CNN

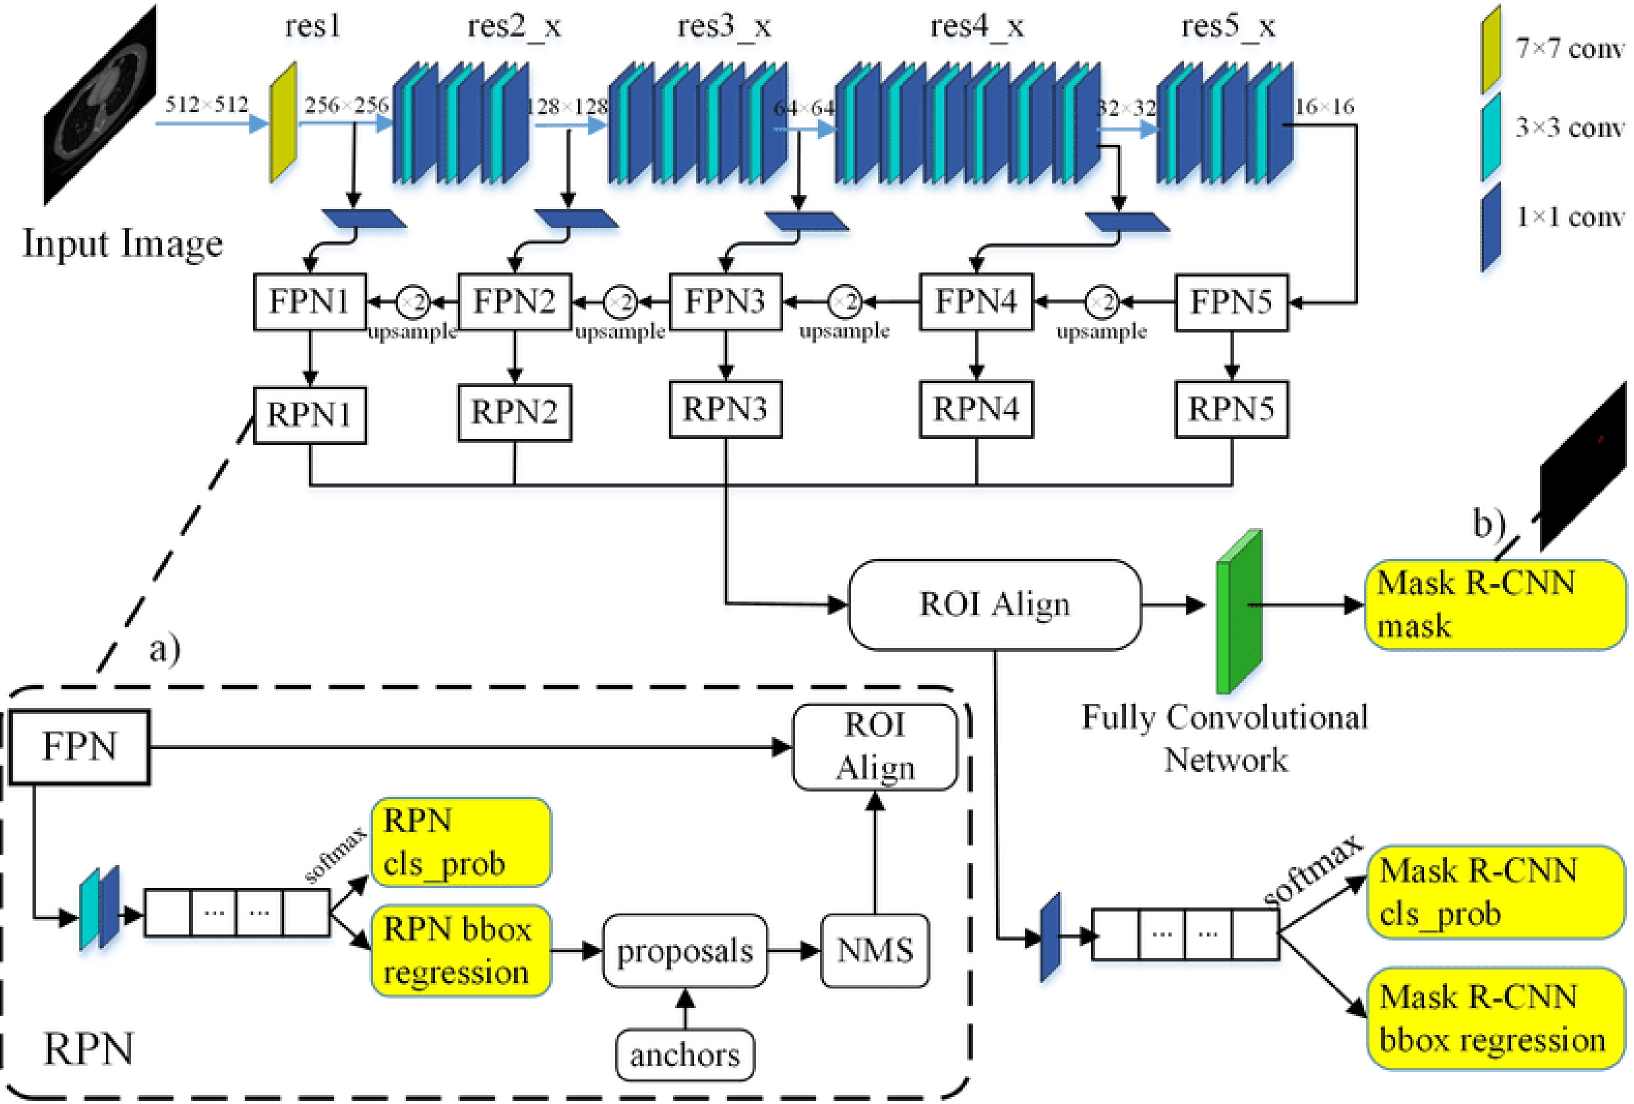

res1
res2_x
res3_x
res4_x
res5_x
512×512
256×256
128×128
64×64
32×32
16×16
Input Image
7×7 conv
3×3 conv
1×1 conv
FPN1
FPN2
FPN3
FPN4
FPN5
×2
upsample
RPN1
RPN2
RPN3
RPN4
RPN5
ROI Align
Fully Convolutional Network
b)
Mask R-CNN mask
a)
FPN
ROI Align
RPN cls_prob
softmax
RPN bbox regression
proposals
NMS
anchors
RPN
softmax
Mask R-CNN cls_prob
Mask R-CNN bbox regression

YOLOv8

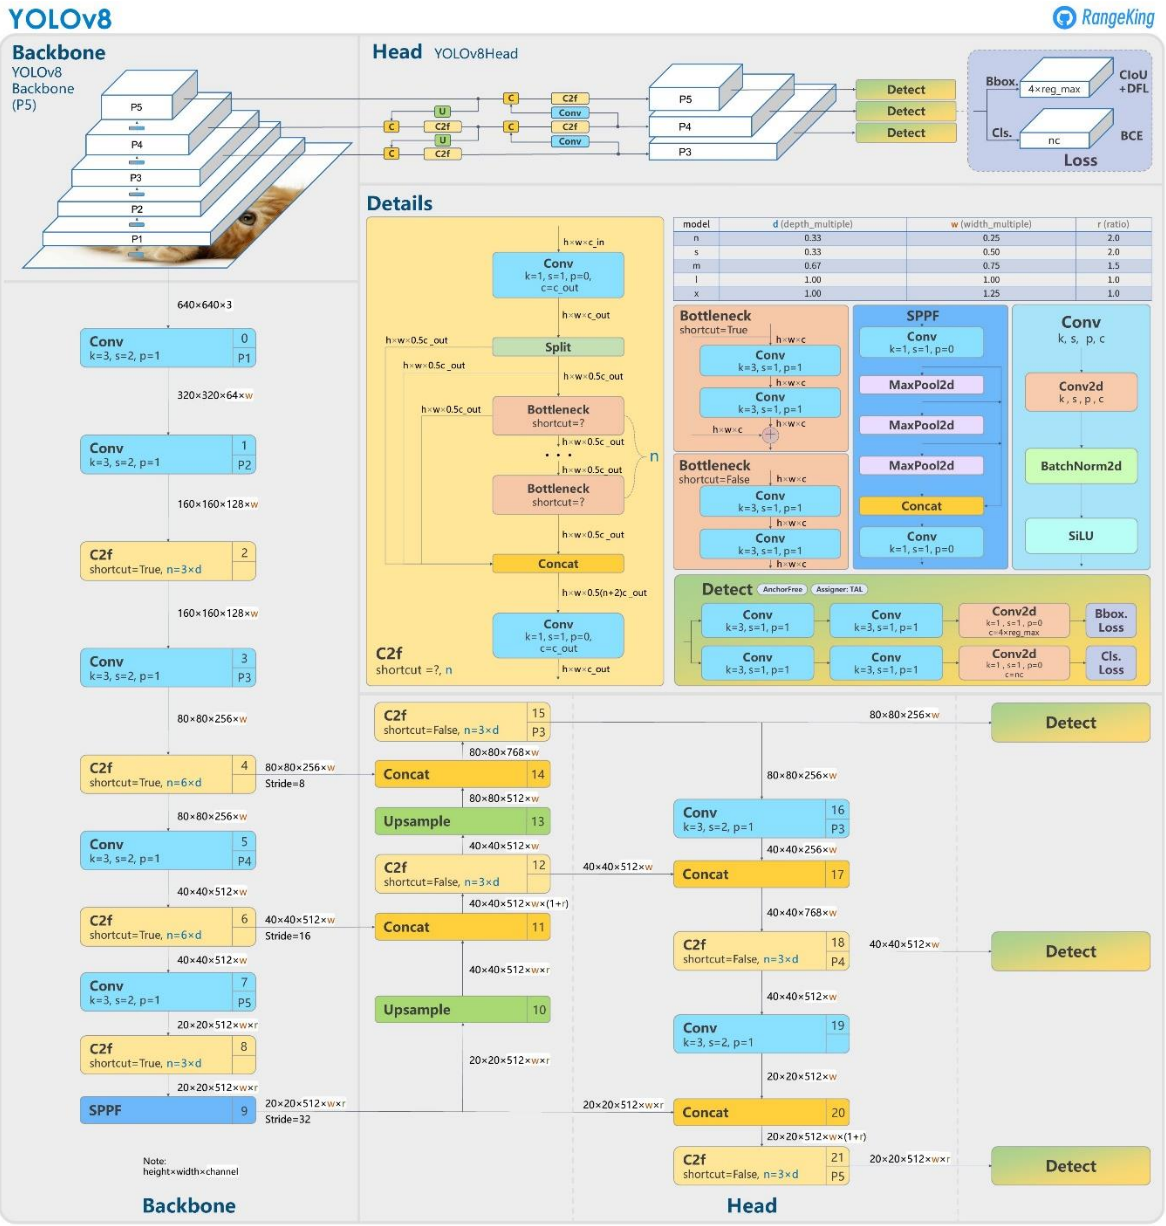

YOLOv8
RangeKing
Backbone
Head
Details
Note: height×width×channel

mAP – mean Average Precision

1. Model predictions.
2. Assign each predicted instance to the ground truth instance with the highest IoU above the threshold (IoU ≥@).

$$\text{IoU}(A,B) = \frac{|A \cap B|}{|A \cup B|}$$

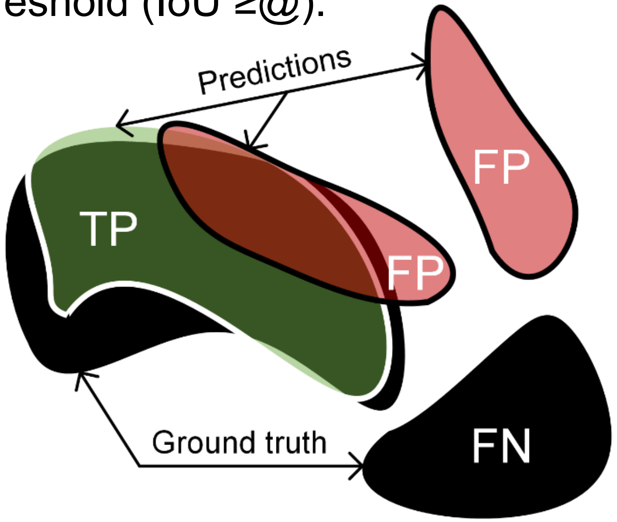

3. Sort the predicted instances based on confidence scores in descending order.
4. **PR**: Calculate precision and recall at each step as you iterate through the sorted predictions:
   a. **TP**: Prediction has a correct category and sufficient IoU;
   b. **FP**: Prediction has an incorrect category or insufficient IoU;
   c. **FN**: If ground truth is not matched with any prediction;

$$\text{Precision (P)} = \frac{TP}{TP + FP}$$ (TP + FP = number of predictions)

$$\text{Recall (R)} = \frac{TP}{TP + FN}$$ (TP + FN = number of ground truth annotations)

5. **AP**: Interpolate and compute the area under the precision-recall curve.
6. **mAP**: Calculate the mean of the individual class average precisions.

$$\text{mAP} = \frac{\sum_{category} AP_{category}}{n_{categories}}$$

7. **mAP@0.5:0.05:0.95**: Repeat for different IoU thresholds @ and calculate the mean.

$$\text{mAP@}0.5{:}\,0.95 = \frac{\sum_{IoU} mAP@IoU}{n_{IoUs}}$$

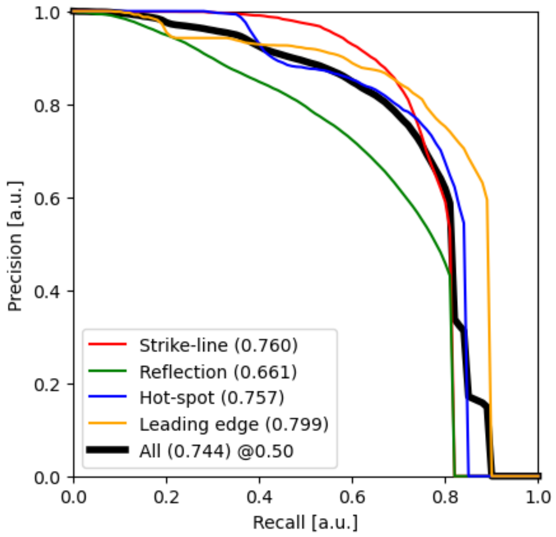

Bad Cold Pixel Removal

Morphological Closing by Reconstruction

841	815	**730**	830	815
841	815	830	815	815
830	830	451	830	841
830	815	451	841	841
815	433	451	451	433
433	433	**320**	433	418
418	433	451	433	418

841	815	**815**	830	815
841	815	830	815	815
830	830	451	830	841
830	815	451	841	841
815	433	451	451	433
433	433	**433**	433	433
433	433	451	433	433

841	830	**830**	830	830
841	830	830	830	830
830	830	**830**	830	841
830	830	**830**	841	841
815	451	451	451	451
433	433	**451**	433	433
433	433	451	433	433

Morphological Closing

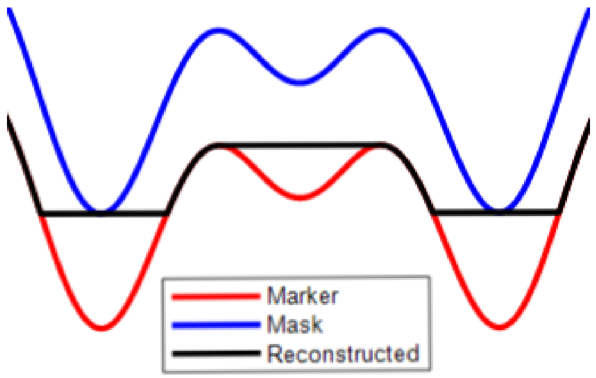

Max-Tree

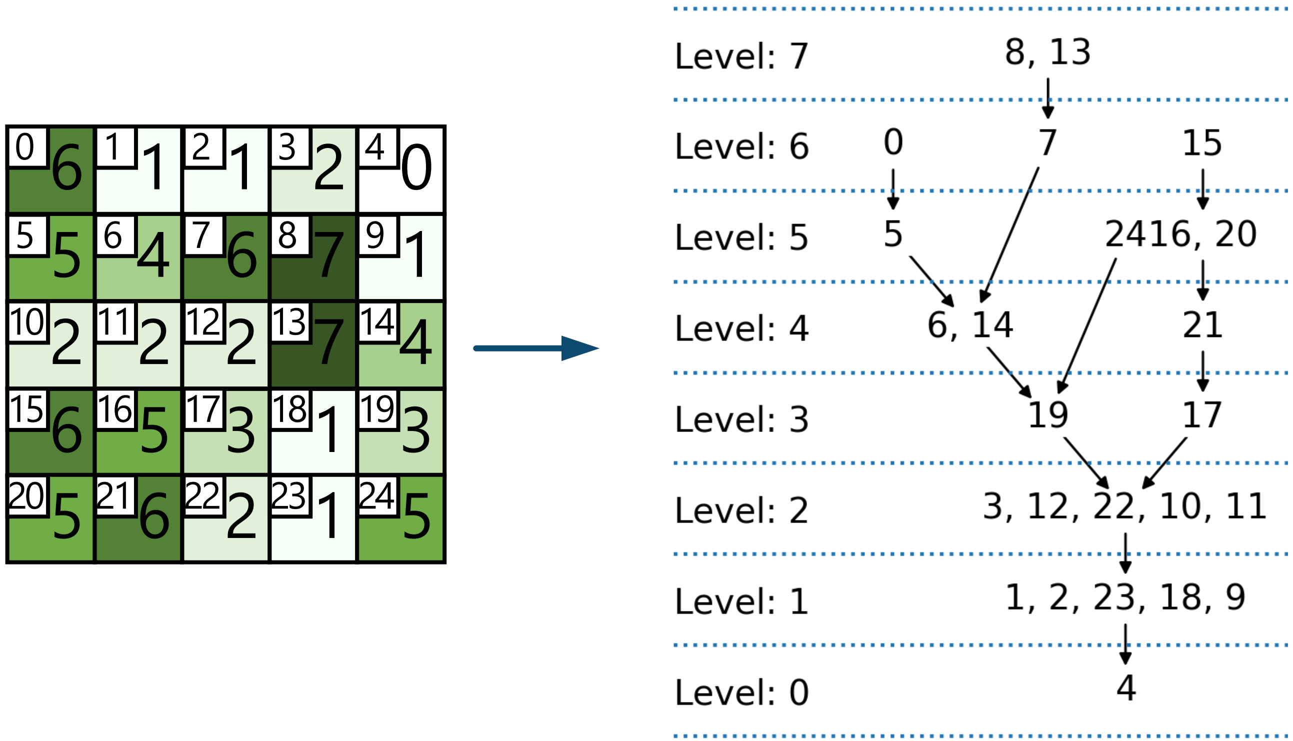
Level: 7
8, 13
Level: 6
0
7
15
Level: 5
5
2416, 20
Level: 4
6, 14
21
Level: 3
19
17
Level: 2
3, 12, 22, 10, 11
Level: 1
1, 2, 23, 18, 9
Level: 0
4

Real-Time Constraint

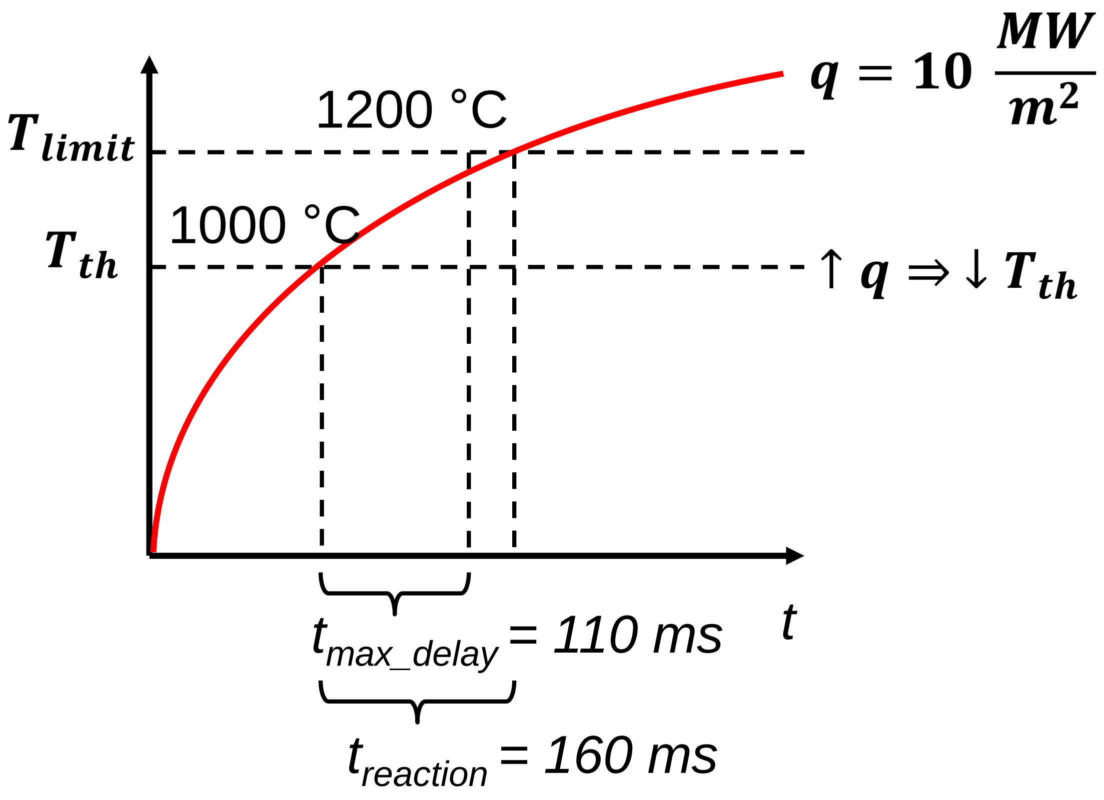

VAE Formulas

Prior	$p_\Theta(z)$
Posterior	$p_\Phi(z\|x)$
Likelihood	$p_\Theta(z\|x)$

$$p_\Phi(z|x) \approx p_\Theta(z)$$

$$-D_{KL}[P_\Phi(z|x) \| P_\Theta(z)]$$

$$z = \mu(z) + \Sigma(z) * \epsilon$$

$$L(\Phi, \Theta, x) = -D_{KL}[P_\Phi(z|x) \| P_\Theta(z)] + E_{q(z|x)}[log(p_\Theta(x|z))]$$

$$D_{KL}(P\|Q) = -\frac{1}{2}\sum_{i=0}^{n} 1 + \log\sigma_{qi}^2 - \mu_{qi}^2 - \sigma_{qi}^2$$

$$MAE(y_{true}, y_{pred}) = \frac{\sum_{i=1}^{n}\left|y_{true}^{i} - y_{pred}^{i}\right|}{n}$$

$$MSLE(y_{true}, y_{pred}) = \frac{\sum_{i=1}^{n}\left(\log\left(y_{true}^{i} + 1\right) - \log\left(y_{pred}^{i} + 1\right)\right)^2}{n}$$

$$BCE(y_{true}, y_{pred}) = -\frac{\sum_{i=1}^{n} y_{true}^{i}\log y_{pred}^{i} + \left(1 - y_{true}^{i}\right)\log\left(1 - y_{pred}^{i}\right)}{n}$$

$$SSIM(y_{true}, y_{pred}) = \frac{1}{n}\sum_{i=1}^{n}\frac{\left(2\mu_{true_i}^2\mu_{pred_i}^2 + c_1\right)\left(2\sigma_{true-pred_i} + c_2\right)}{\left(\mu_{true_i}^2 + \mu_{pred_i}^2 + c_1\right)\left(\sigma_{true_i}^2 + \sigma_{pred_i}^2 + c_2\right)}$$

- t-distributed stochastic neighbour embedding

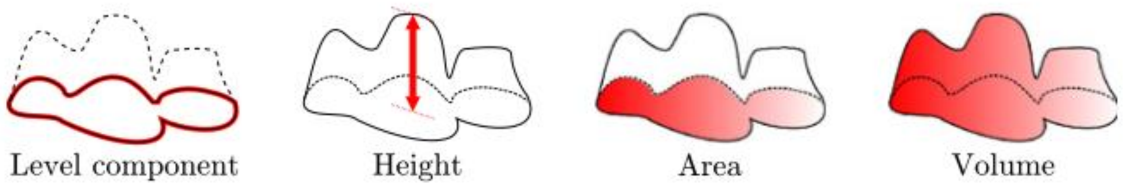

Fig. 4. Attributes from the relief top of a region of the image.

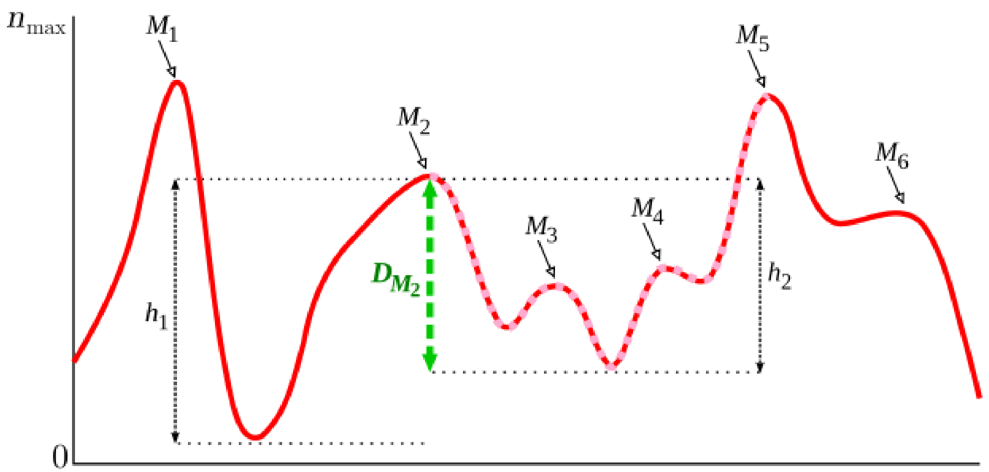

Fig. 3. Dynamics D_{M_2} of regional maximum M_2. Minimum height climb down from M_2 to reach another higher maximum (M_1 or M_5).

WEST

#55210 Divertor