



HAL
open science

Detection of Thermal Events by Semi-Supervised Learning for Tokamak First Wall Safety

Christian Staron, Hervé Le Borgne, Raphaël Mitteau, Erwan Grelier, Nicolas Allezard

► **To cite this version:**

Christian Staron, Hervé Le Borgne, Raphaël Mitteau, Erwan Grelier, Nicolas Allezard. Detection of Thermal Events by Semi-Supervised Learning for Tokamak First Wall Safety. 2023. cea-04099065v2

HAL Id: cea-04099065

<https://cea.hal.science/cea-04099065v2>

Preprint submitted on 18 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Detection of Thermal Events by Semi-Supervised Learning for Tokamak First Wall Safety

Christian Staron, Hervé Le Borgne, Raphaël Mitteau, Erwan Grelier, Nicolas Allezard

Abstract—This paper explores a semi-supervised object detection approach to detect thermal events on the internal wall of tokamaks. A huge amount of data is produced during an experimental campaign by the infrared (IR) viewing systems used to monitor the inner thermal shields during machine operation. The amount of data to be processed and analyzed is such that protecting the first wall is an overwhelming job. Automating this job with artificial intelligence (AI) is an attractive solution, but AI requires large labelled datasets that are not readily available for tokamak walls. Semi-supervised learning (SSL) is a possible solution to being able to train deep learning models with a small amount of labelled data and a large amount of unlabelled data. SSL is explored as a possible tool to rapidly adapt a model trained on an experimental campaign *A* of tokamak WEST to a new experimental campaign *B* by using labelled data from campaign *A*, a little labelled data from campaign *B* and a lot of unlabelled data from campaign *B*. Model performance is evaluated on two labelled datasets and two methods including semi-supervised learning. Semi-supervised learning increased the mAP metric by over six percentage points on the first smaller-scale database and over four percentage points on the second larger-scale dataset depending on the method employed.

Index Terms—Fusion reactors protection, infrared thermography, semi-supervised learning, domain adaptation, object detection

I. INTRODUCTION

THE search for abundant, cheap, low-carbon energy is driving research into nuclear fusion. This reaction, in which two atomic nuclei assemble to form a heavier nucleus while releasing a large amount of energy, is naturally at work in the sun and can be carried out on Earth within tokamaks [1]–[4] or stellarators. Such devices confine a plasma in a limited volume at very high temperatures thanks to powerful magnetic fields.

However, some particles escape inevitably from confinement and reach the Plasma Facing Components (PFC), leading to intense though localized heating. Thermal events appear on the first wall and can damage these in-vessel components when exceeding armour material allowable. Among numerous

This paper was produced by the Commissariat à l’Énergie Atomique et aux énergies alternatives (CEA – French Alternative Energies and Atomic Energy Commission), France. We acknowledge the financial support of the Cross-Disciplinary Program on Numerical Simulation of CEA. This work was granted access to the HPC resources of IDRIS under the allocation 2022-AD010513390R1 made by GENCI. We also thank Laurent Letellier and Patrick Sayd who were the source of this collaborative work.

Christian Staron, Raphaël Mitteau and Erwan Grelier are with the IRFM, an institute part of the CEA, F-13108 St Paul Lez Durance, France

Hervé Le Borgne and Nicolas Allezard are with the List, an institute part of the CEA and Université Paris-Saclay, F-91120, Palaiseau, France.

sensors to monitor the first wall, infrared (IR) viewing systems provide the most relevant information about the surface temperature of the inner walls of tokamaks (Figure 1).

The live image feeds are used for real-time feedback control during operation. These data are also used for scientific post-processing after the experiments, to localize and characterize the different classes of thermal events that occur during machine operation. This task requires long and tedious work by high-level experts of plasma physics who would benefit from automation through computer vision approaches.

First attempts in this direction were based on simple thresholds [5] but were tedious to tune manually and suffered from high sensitivity to false detection. It has been slightly improved with the modelling of the background by clustering [6], but had still limited performance in practice. More recent approaches rely on large neural networks learned from visual data. Grelier, Mitteau and Moncada [7] trained a Cascade R-CNN model [8] from scratch using a dataset of 325 thermal events distributed in seven classes, manually annotated from 20 infrared movies from the WEST tokamak. They improved their approach more recently [9] with a Faster R-CNN model [10]. Szucs *et al.* [11] used the YOLOv5 [12] model, with a backbone pre-trained on ImageNet and fine-tuning with 471 real images from the W7-X stellarator and 250 synthetic images of hot spots. Such approaches nevertheless still require tedious manual work to annotate the IR video films.

Using a dedicated tool (described in [7]), labelling one IR film takes between 20 and 30 minutes. Hence, labelling the whole database with these tools manually would take over three years as the database contains more than 13 000 films.

Even by taking multiple precautions, the process of annotation is known to be subject to multiple potential biases [13]. The main potential subject for thermal event labelling is the *label bias* because “different annotators may assign different labels to the same type of [thermal event]” [14]. Since the knowledge of plasma fusion progresses with time, even the same experts can make subtle changes in their assessments to categorize thermal events, between two annotation sessions several months apart.

This article addresses the two aforementioned limitations of the current approaches and proposes an automated process to detect and classify thermal events on the internal wall of tokamaks, which requires a significantly lesser amount of manual work while maintaining efficiency. We rely on the recent semi-supervised approach SoftTeacher [15] to take advantage of the large amount of unlabelled data produced

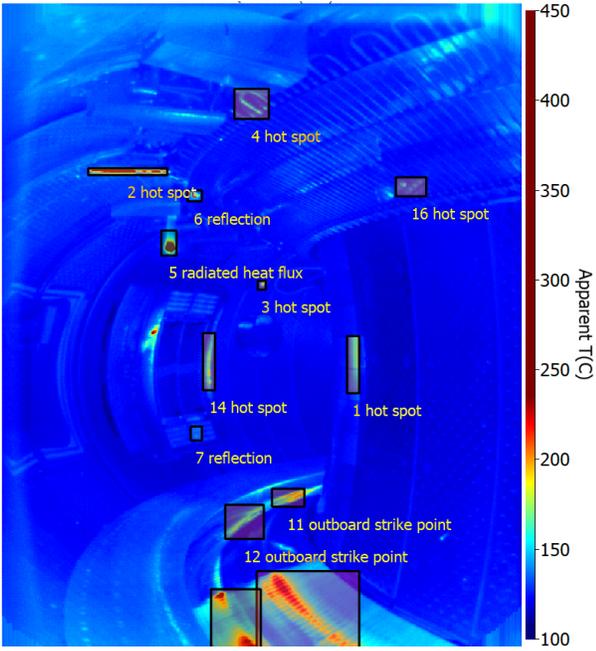


Figure 1. Annotation examples of several classes of hot spots on the wide-angle line of sight of the tokamak WEST [4]

by the IR cameras, in the vein of the semi-supervised principles [16], [17]. While still requiring a limited quantity of manual annotations, this approach reduces their amount by one to two orders of magnitude, with a controlled and limited drop in performance.

The study also shows a phenomenon that was previously unidentified. From an experimental campaign to the following one, small changes in the tokamak’s configuration make the images different without being clearly visible to human experts. However, although difficult to see “humanly”, these changes cause a collapse of the performance with a supervised approach, learned during an acquisition campaign and applied as is to a subsequent campaign. The method proposed here overcomes this issue with a minimal manual annotation effort on a novel campaign and adapts the model to the (humanly imperceptible) new visual domain.

To summarize, the contributions of this paper are:

- a novel method for visual monitoring of the inner wall of tokamaks, based on semi-supervised learning. Through this paradigm, we show it addresses two major limits of the previous fully-supervised approaches, namely the need for massive annotated data and the subjectivity of the human annotator that introduces noise to the labels;
- the identification of a novel phenomenon that arises between two acquisition campaigns in a tokamak. While difficult to see, even for human experts, it has a major impact on fully-supervised approaches. We show that our SSL-based approach is much more robust to this issue.

In the following, the method based on semi-supervised learning is described in section II. The data handled in the study, which come from two different acquisition campaigns, as well as several experimental results are reported in section III, showing the interest of the semi-supervised

approaches with regards to the compromise between manual annotation effort and detection performance on the one hand, and the ability to adapt to a new domain resulting from a new acquisition campaign on the other hand. Finally, we discuss the limitations of the method and possible future impact in section IV.

II. PROPOSED APPROACH

Detecting and classifying thermal events on the internal wall of tokamaks is processed here as an object detection problem. Given an input image, an object detection model predicts a collection of bounding boxes that localize and classify the identified hot spots in the image. Each bounding box consists of 5 predictions: (x, y, w, h) and the confidence for the class label. The (x, y) coordinates represent the upper-left corner of the box, w its width and h its height.

Most recent work proposed to estimate such a model in a supervised learning setting [7], [9], [11]. They rely on a large labelled dataset $L_{\text{train}}^{\text{sup}}$ and minimize a loss, whose exact form depends on the outputs design of the model. In that case, the loss is the sum of a cross-entropy loss for classification and a smooth L_1 loss for regression. The regression allows us to estimate the values (x, y, w, h) for each bounding box, while the classification loss results in an estimation of the confidence over the class label. The final loss \mathcal{L}^{sup} is a weighted sum of all these losses, that is minimized with stochastic gradient descent.

A. Semi-supervised learning to detect thermal events

The main limit of the supervised approach is that it requires a large dataset $L_{\text{train}}^{\text{sup}}$, which is tedious to manually annotate. The Semi-Supervised Learning (SSL) paradigm is adopted as it needs a much smaller labelled training set L_{train} but takes advantage of a large unlabelled training set U_{train} that is obtained at marginal cost during each acquisition campaign of plasma fusion within the tokamak. The general process of SSL on infrared images encompasses two steps:

- 1.) Step 1: (burn-in phase) Train the model on annotated images belonging to L_{train} . The burn-in phase is a fully supervised learning phase, also used as a reference later in the paper
- 2.) Step 2: Apply the model to the images of U_{train} to estimate a pseudo-label for each of them. Then train the model with both the labelled images from L_{train} and the pseudo labelled images from U_{train}

Since the size of L_{train} is typically 1% of that of $L_{\text{train}}^{\text{sup}}$, it reduces significantly the manual effort of annotation required. On the other hand, the size of U_{train} is of the same order as that of $L_{\text{train}}^{\text{sup}}$ but since the pseudo-labels are obtained automatically, it does not require manual effort.

In the vein of the recent literature in SSL for object detection [15], [18], [19] this work relies on a student-teacher architecture that comprises two branches. The teacher model receives the input of unlabelled images weakly augmented and annotates them with pseudo labels. On the other branch, the input images are strongly augmented and the student model uses these images and the estimated pseudo labels to train the

hot spot detector itself. During training, the student model is trained by minimizing the weighted sum of a supervised loss \mathcal{L}^{sup} estimated with labelled images and an unsupervised loss $\mathcal{L}^{\text{unsup}}$ that results from unlabelled images only. In practice, $\mathcal{L} = \mathcal{L}^{\text{sup}} + \lambda \mathcal{L}^{\text{unsup}}$, where λ is a hyperparameter that is 0 during the burn-in phase and is non-null during Step 2. The teacher model is an exponential mean average (EMA) [20] of the student model. Its parameters θ^{teacher} are updated according to $\theta^{\text{teacher}} \leftarrow \alpha \theta^{\text{teacher}} + (1 - \alpha) \theta^{\text{student}}$ at each iteration, where α is a hyperparameter slightly less than one. Note that to detect an object in an image, we face a class imbalance between the boxes that contain a potential object (foreground) and those that do not (background). This imbalance is managed by object detectors such as Faster-RCNN, which is used at the core of our approach, thanks to a binary detector of *objectness*, that acts on the last convolutional layer of the backbone, which is further used to detect (through regression) and recognize the class of the hot spot (through classification).

Applying data augmentations to input images is a well-known efficient strategy in SSL [21]. They enable a diversification particularly of the unlabelled data so that the teacher model does not fall into the trap of always predicting the same boxes at each iteration. Different augmentations are applied depending on the set from which the data is from.

- Weak augmentations are applied to unlabelled images that are given as input to the teacher model. Weak augmentations include resizing the image or flipping it in relation to the vertical axis that separates the image in half.
- Supervised augmentations are applied to the labelled images given as input to the student model. Supervised augmentations add to weak augmentations one extra augmentation that is picked from a list of augmentations with RandAugment [22]. The image can be for instance sharpened, equalized, brightened, contrasted or even not augmented at all.
- Strong augmentations are applied to the unlabelled images given as input to the student model. Strong augmentations add to the already applied supervised augmentations transformations like translations, shearing, rotating or erasing parts of the image by Cutout [23].

A choice was made on which augmentations could be applied to WEST’s infrared images because the pixel values in the IR images have a physical interpretation, namely the surface temperature measurements. All augmentations listed in [15] that strongly affect the pixel values are thus not applied. This includes augmentations that equalize the image, solarize it or mix the color channels. The weak augmentations are unchanged. They are the same augmentations used when training a supervised model in Step 1. The kept supervised augmentations comprise of augmentations that contrast the images, brighten or sharpen them. Strong augmentations are linked to more physical augmentations. The preserved augmentations consist of translating, shearing and erasing part of the images at each iteration.

B. Adaptation to a novel acquisition campaign

As mentioned above, there may be a difference between images generated by IR cameras for subsequent acquisition campaigns. These differences are either imperceptible to the human eye or blinded by the human brain which tends to correct the images unconsciously. Differences can be due to the installation of new components or various plasma configurations that affect the location of the power deposition.

In this paper, the images are collected from two such campaigns, denoted by A and B , described in subsection III-A. As shown experimentally in Section III-C, when a detection/classification model is trained with data from campaign A , its performance drops dramatically when applied directly on data from campaign B . For a new campaign, the parameters θ_A previously learned need to be adapted to new parameters θ_B that provide acceptable performance. Such an issue can be seen as a domain adaptation problem [24] that can be addressed with transfer learning approaches such as using universal backbones [25], [26].

Fine-tuning is a common method that enables the quick adaptation of model weights. Hence, one can therefore adjust a model’s weights when trained on campaign A with images from campaign B to quickly gain good performance on a test set containing exclusively images from campaign B . However, it still requires a significant amount of labelled data to avoid overfitting and reach the performance of a model that would be learned from scratch on images from both A and B .

This article aims at demonstrating that an approach based on SSL is more efficient than fine-tuning for a given budget of labelled images. The two following scenarios are considered in the following:

- method 1: the initial burn-in is performed from scratch with a small amount of labelled images, either from campaign A or B . Then, Step 2 is performed with these images and unlabelled images from B only.
- method 2: the initial burn-in is performed from scratch with a small amount of labelled images from campaign A then fine-tuned with images from B . Then Step 2 is performed similarly to method 1, with these images and unlabelled images from B only.

Method 1 reflects an approach that requires retraining the model at each new campaign, with a small amount of labelled data. Method 2 is closer to an ideal realistic use case for WEST’s operating team, in which the model learned previously in campaign A is directly adapted to campaign B without training from scratch.

C. Implementation details

The implementation of our model (Figure 2) is built upon the code of SoftTeacher [15], that itself relies on MMDetection [27]. The (student) detection model is a Faster R-CNN with a ResNet-50 [28] backbone. Such a model is thus similar to that used by recent works in a supervised setting [9]. Hence, the supervised loss \mathcal{L}^{sup} is the sum of a cross entropy loss for classification and a smooth L_1 loss for regression, computed with labelled data only. The unsupervised loss $\mathcal{L}^{\text{unsup}}$ is also the sum of a classification and a regression losses, but the

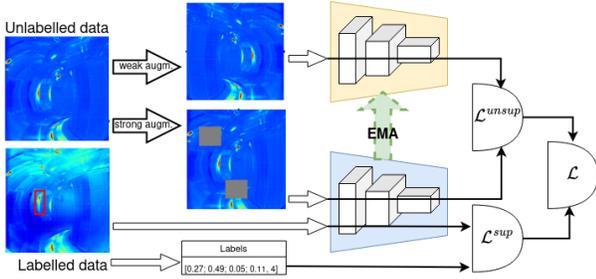


Figure 2. The architecture of the model used in this study. The student model (bottom, in blue) learns a supervised loss \mathcal{L}^{sup} with labelled data and an unsupervised loss \mathcal{L}^{unsup} with unlabelled data, which the pseudo-label is provided by the teacher (top, yellow). At each epoch, the teacher is updated from the weights of the student with an exponential moving average (EMA). More details can be found in [15] which code was used at the base of our implementation

Table I
MAIN HYPERPARAMETERS ADOPTED TO TRAIN FASTER R-CNN AND
SOFTTEACHER MODELS

	Supervised training (burn-in)	SSL training
Relative weight λ	–	4.0
Learning rate	5×10^{-3}	5×10^{-4}
Batch size	16	20
Total iterations	80 000	
Scheduler	40 000 and 53 333 iterations	
Weight Decay	0.0001	
Momentum	0.9	
EMA update α	0.999	

annotation is provided through a pseudo-label computed by the soft-teacher. The unsupervised classification loss is modified such that the background boxes are weighted according to their reliability score, while the unsupervised regression loss is modified to filter the 2D boxes that the teacher model does not regress with enough stability after a small jittering [15]. Table I shows the main hyperparameters used to train both the supervised models and the semi-supervised models. The setup of the hyperparameters results either from the values proposed by [15] or thanks to a grid search conducted on a $2/3 / 1/3$ split of the training dataset of \mathcal{D}^1 (described in subsection III-A). Hence, the relative weight λ of the unsupervised loss is fixed at 4.0 during Step 2. The batch sizes are reduced to retain the number of images per GPU while taking into account the number of GPU used. The learning rates are first adapted proportionally to the changes in batch sizes to keep their ratio constant, then further reduced if needed to ensure convergence. The models are trained on 4 NVIDIA Tesla V100 16 GB GPU, each computing a quarter of the batch size before it is aggregated to estimate the gradient.

The supervised models of the burn-in stage are trained on a number of images from campaign *A* equivalent to 1% of the fully labelled dataset and a number of images from campaign *B* corresponding to 0.5% of the labelled dataset. Models are often trained on 10%, 5% and 1% of labelled data in literature [18], [19], usually on the RGB images of MSCOCO [29]. The decision to train models on 1% relies on the fact that it is the

Table II
THE DIFFERENT THERMAL EVENT CLASSES IN THE LABELLED DATASETS

Class number	Class name	Number of hot spots per class	
		\mathcal{D}^1	\mathcal{D}^2
1	electron type 1	2 103	0
2	inboard strike point	13 078	26 132
3	outboard strike point	19 569	32 434
4	reflection	856	724
5	radiated heat flux	631	0
6	UFO	30	149
7	hot spot	61 237	119 751

hardest case because the least amount of labelled data is dealt with during training but also it is the most realistic use case for WEST, as a lot of unlabelled data is at disposal.

III. EXPERIMENTAL RESULTS ON WEST'S DATASETS

A. WEST Dataset

The work of this article is based on images from the tangential line of sight of the WEST tokamak [4], [30]. The line is equipped with an infrared camera, and provides an overall wide-angle (WA) view of the WEST tokamak. It has thus a key role in the protection of the machine because it has the largest variety of thermal events (Figure 1).

A WEST fusion experiment lasts typically between 15 and 60 seconds, and the longest one up to several minutes. Acquisition campaigns are organized on a regular basis, separated by several months to service the tokamak and its utilities. The images resulting from the videos are of size 512×640 px and encoded over 16 bits.

The thermal events are labelled according to seven classes for the purpose of this work (reported in Table II, illustrated in Figure 3, and presented in [7]), while a full data processing pipeline would have many more classes. Distinguishing the different categories and assigning a label to a bounding box is a difficult task, that requires a good knowledge of the machine and the use of both the temporal behavior of the thermal event and data from other diagnostics. Each label consists of four integers corresponding to the surrounding box of the spot, and one to its class. The six first classes correspond to well-defined wall events and the seventh is the catchall class, which is appropriate when the exact definition of the thermal event is undefined. In practice, with progress in plasma fusion concerning the identification of these thermal events, new classes can be defined in the future and include a part of the annotations in this catchall class. Note that the distinction is made between “thermal events” and “hot spots”: a hot spot is related to one single timestamp. By contrast, a thermal event is a sequence of hot spots. This distinction is important because some classes of thermal events can only be determined using spatio-temporal information. For instance, UFOs can be classified as such only by looking at the movement between subsequent hot spots.

Acquisition Campaigns. WEST’s experimental campaigns last over several months. WEST goes through many different plasma configurations during these campaigns and evolves

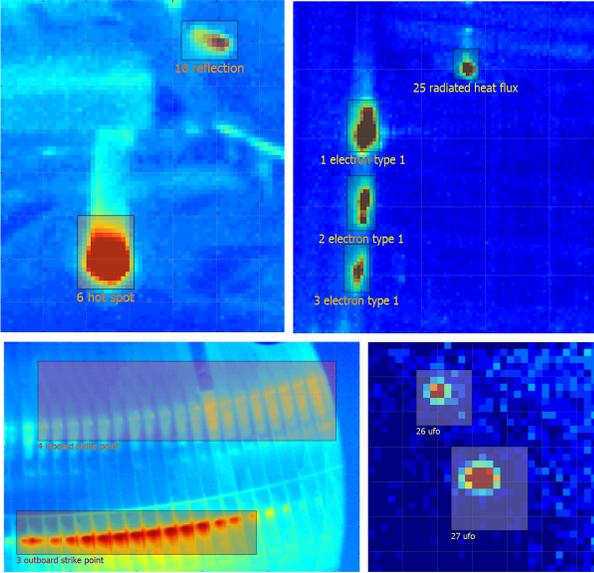


Figure 3. Examples of different manual annotations (upper left: reflection and hot spot; upper right: radiated heat flux and electrons type 1; lower left: inboard and outboard strike point; lower right: UFOs)

from one campaign to another (adding new components, new diagnostics...).

Such changes make the images acquired by the IR cameras different from one campaign to another, as illustrated in Figure 4. Some changes may not be detected by human experts because human tend to correct the images unconsciously, but those changes have a significant impact on the low-level statistics of the images and, as a result, on the performance of the automatic vision systems.

The images we consider were extracted from video films acquired during two different acquisition campaigns to be able to study this effect, namely the campaign *A* which took place between July and November 2019, and the campaign *B* (November 2020 – January 2021).

Annotation Campaigns. Two annotation sessions were organized to study potential differences between annotation campaigns. \mathcal{D}^1 is made of eight video films annotated in September 2021, corresponding to 12 156 images [7]. \mathcal{D}^2 was annotated by the same expert in August 2022, on 21 other video films, corresponding to 18 336 images [9]. Note that both \mathcal{D}^1 and \mathcal{D}^2 contain images from acquisition campaigns *A* and *B*. A unique expert annotated both \mathcal{D}^1 and \mathcal{D}^2 to mitigate the risk of inter-annotator variation. The experiments of \mathcal{D}^1 are picked by experts because they contain interesting events from a machine protection standpoint, whereas the ones of \mathcal{D}^2 are selected algorithmically to maximize the diversity in machine configurations (plasma configuration, duration, injected power and energy). Table II shows the number of occurrences of each class of hot spot for both datasets. A value of 0 indicates that the dataset does not contain any image that has a given class of thermal event.

Splitting for Experiments.

Each of the two annotated datasets is divided into a train and a test set. The splits are made at the video film level to avoid data leakage. If a detector were trained with images

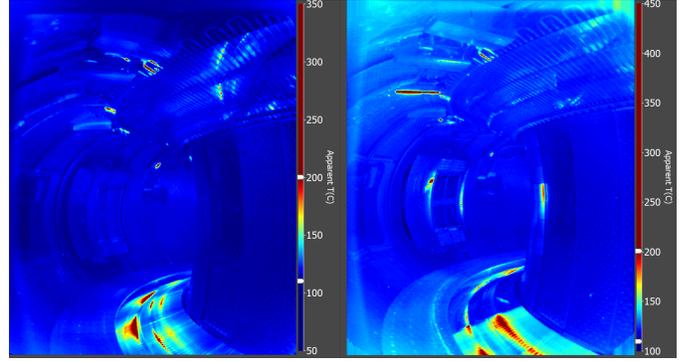


Figure 4. Image from campaign *A* (left) during an experiment with similar experimental conditions as the image from campaign *B* (right)

from a given video film, it would be much easier to recognize hot spots from images of the same video film (even if they are not seen during training), due to the strong temporal coherency within each video film. Table III gives the number of images in the training and test sets. In the following, each subset is referred by the notation X_β^α , where $X \in \{A, B\}$ is the acquisition campaign, $\alpha \in \{1, 2\}$ refers to the annotation session datasets (\mathcal{D}^1 or \mathcal{D}^2) and $\beta \in \{\text{train}, \text{test}\}$ is the split. \bar{X} denotes the complement of X with respect to the set of labelled data from the campaign X . The considered subsets of X contain images that are equally distributed between the video films.

The same number of images is extracted from each labelled film to have a well-defined representation of each film in the training set. During experiments, if one needs to train on a number of images equivalent to *e.g.* 1% of the dataset, 96 labelled images are considered for \mathcal{D}^1 and 144 for the larger dataset \mathcal{D}^2 .

B. Evaluation metric

The mean Average Precision (mAP) is used to measure the performance. It is a reference metric in object detection that both reflects the localization and classification ability of the model, and aggregates the performance for all the classes considered. The mAP values reported for SSL experiments in the following are always those of the student model. The mAP implementation is the same as the one adopted for the COCO challenges, usually noted $\text{mAP}_{0.5:0.95}$, that is the mAP averaged over an “intersection over union” (IoU) between the detected hot spot and the ground truth at 0.5, 0.55, 0.6, ..., 0.95. Following [31], the mAP is only computed for boxes with a confidence higher than 10^{-3} .

During the preliminary tests of our work, we considered all the standardized COCO metrics, but did not notice any remarkable results both globally or by class. In other words, the relative order of the performances was always the same regardless of the metric. For this reason, we only report the (integrated) mAP as described above.

C. Inter-campaigns and Annotations Domain Shift

Subsection III-A explains that the datasets have been created to study the transfer of a model learned on the data of

Table III

CONTENT OF THE LABELLED DATASETS \mathcal{D}^1 AND \mathcal{D}^2 USED IN EXPERIMENTS. IMAGE DATA COME FROM TWO ACQUISITION CAMPAIGNS (A AND B) AND LABELS FROM TWO ANNOTATIONS CAMPAIGNS (CORRESPONDING TO THE DATASETS \mathcal{D}^1 AND \mathcal{D}^2).

	\mathcal{D}^1		\mathcal{D}^2	
	Train	Test	Train	Test
Number of films from A	2	1	8	1
Number of films from B	3	2	9	3
Number of images	7 879	4 277	15 130	3 206
Notation	$A^1_{\text{train}} + B^1_{\text{train}}$	$A^1_{\text{test}} + B^1_{\text{test}}$	$A^2_{\text{train}} + B^2_{\text{train}}$	$A^2_{\text{test}} + B^2_{\text{test}}$

campaign A to the test data of a campaign B. Firstly, a series of experiments in a fully labelled setting is conducted. The reporting of the results is in Table IV.

The results are quite good (lines (i), (iv) and (v) of Table IV) when the model is trained and tested on data from the same campaign (and annotated homogeneously in \mathcal{D}^1 : see below). The results on B^1 are better than on A^1 but the inclusion of data from B boosts performance. This may be due to a larger training dataset or more relevant images in B^1 . However, the performance drops dramatically to near zero (lines (ii) and (iii) of Table IV) when the training and test data are from different campaigns. These experiments show that the changes in tokamak configurations from one campaign to another produce subtle changes in the low-level statistics of the images that are not visible to experts (see Figure 4) but have a significant impact on performance, comparable to that of a domain shift.

The datasets enable the study of the influence of a potential light change in the annotation process. Indeed, while both datasets \mathcal{D}^1 and \mathcal{D}^2 contain data from campaigns A and B, \mathcal{D}^2 is annotated almost one year after \mathcal{D}^1 . Even though the annotators are the same, the criterion to classify the thermal events by experts changed slightly, introducing minor differences between the two. For instance, following discussions with experts, for \mathcal{D}^2 , the reciprocating Langmuir probe saw its classification change from “radiated heat flux” to “hot spot”, and some thermal events on the upper divertor were annotated and classified as “reflection”. Furthermore, the differences in diversity of the configuration of the machine, described in subsection III-A, also introduce differences between the two datasets.

Good performance is still attained when the training and test data is homogeneous (lines (vii), (x) and (xi) of Table IV) but a significant drop is noted when the training set and the test one come from different annotation sessions (lines (vi) and (viii) of Table IV). The fact that the model trained on \mathcal{D}^2 has better performance than that trained on \mathcal{D}^1 shows that later data encapsulates earlier ones. The domain shift identified in \mathcal{D}^1 is nevertheless still present in \mathcal{D}^2 (lines (ix) to (xii) of Table IV).

The change in model performance between datasets and campaigns demonstrates that there is a need to quickly adjust model weights when faced with a new experimental campaign. The following subsections explicit how semi-supervised learning can address this need.

Table IV
mAP FOR DIFFERENT CAMPAIGNS AND DATASETS

Experiment number	Initial training		Evaluation	
	Number of images	set	Test set	mAP
(i)	5 575	A^1_{train}	A^1_{test}	0.381
(ii)	5 575	A^1_{train}	B^1_{test}	0.030
(iii)	2 304	B^1_{train}	A^1_{test}	0.037
(iv)	2 304	B^1_{train}	B^1_{test}	0.435
(v)	7 879	$A^1_{\text{train}} + B^1_{\text{train}}$	$A^1_{\text{test}} + B^1_{\text{test}}$	0.434
(vi)	7 879	$A^1_{\text{train}} + B^1_{\text{train}}$	$A^2_{\text{test}} + B^2_{\text{test}}$	0.087
(vii)	15 130	$A^2_{\text{train}} + B^2_{\text{train}}$	$A^2_{\text{test}} + B^2_{\text{test}}$	0.385
(viii)	15 130	$A^2_{\text{train}} + B^2_{\text{train}}$	$A^1_{\text{test}} + B^1_{\text{test}}$	0.168
(ix)	4 606	B^2_{train}	A^2_{test}	0.030
(x)	4 606	B^2_{train}	B^2_{test}	0.330
(xi)	10 524	A^2_{train}	A^2_{test}	0.396
(xii)	10 524	A^2_{train}	B^2_{test}	0.017

D. SSL Gets Near Supervised Performance at Lower Cost

This section compares the approach based on semi-supervised learning and the classical supervised setting [7], [9]. Experiments are conducted both on \mathcal{D}^1 and \mathcal{D}^2 . The results are reported in respectively Table V and Table VI with the same experiments for both datasets, numbered from (i) to (vii). The supervised settings is reproduced with the same code as our approach for a fair comparison.

The best possible performance is achieved with the Faster R-CNN model in a supervised setting with a ResNet-50 backbone, when the model is trained with all data (from campaigns A and B). This performance varies from 0.32 to 0.477 (line (i) in both tables) depending on the test dataset. The performance logically decreases between 0.267 and 0.435 when only a part of the training data is used (line (ii) to (iv)), depending on the exact settings either at training or testing.

The performance drops with regard to the best fully supervised model with the proposed approach based on SSL (line (vi) of both tables). The drop is at most 4 mAP percentage points on \mathcal{D}^1 with $\frac{7879}{96+48} \approx 54$ times less data and 3.7 percentage points on \mathcal{D}^2 with $\frac{15130}{144+72} \approx 70$ times less data. If one uses the same number of annotated data (line (iii) versus (vi)) the SSL approach surpasses the supervised approach. On \mathcal{D}^1 , the SSL approach even surpasses the model trained on all the data from B^1_{train} (0.456 vs 0.435) when both are tested on

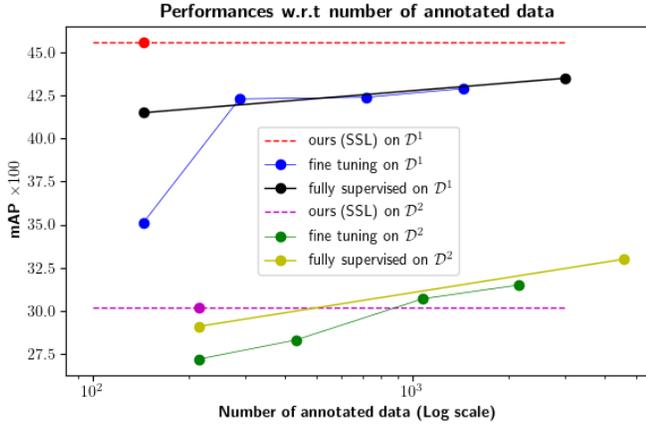


Figure 5. Performance according to the number of annotated data (Log scale in x) for our approach (SSL) compared to fine-tuning and supervised learning, both on \mathcal{D}^1 and \mathcal{D}^2 . The performance of our approach is given for a unique number of annotated data (points). The dashed line is merely drawn as an aid for figure readability : the dashed line recalls the reference level for the continuous lines

B_{test}^1 (line (vi) versus (ii)), although it gets lower ones (0.303 vs 0.33) on \mathcal{D}^2 . Semi-supervised learning is therefore a way to get near fully supervised performance at a much smaller labelling cost.

E. SSL Better Adapts than Fine-Tuning

The SSL approach is compared to the classical fine-tuning as a method to adapt the model to a novel campaign. On line (v) of Table V and Table VI, a model is first trained during 80 000 iterations with 1% of data from campaign A (96 images for \mathcal{D}^1 , 144 for \mathcal{D}^2), then fine-tuned during 10 000 iterations with 0.5% data from campaign B (48 images for \mathcal{D}^1 , 72 for \mathcal{D}^2). Fine-tuning allows the expected adaptation, since the performance on B_{test} is no longer almost null (line (v) versus (iv) of Table V and Table VI, respectively). There is an expected decrease of performance on A_{test} as fine-tuning specializes the model’s parameters to the new campaign B (line (v) of Table V and Table VI). Fine-tuning remains however inferior to the model trained with the same amount of annotated data from A and B (line (iii) versus (v) of Table V and Table VI, respectively).

For both datasets, the same amount of labelled data is handled but additional unlabeled data is used for the SSL model training. The SSL approach surpasses fine-tuning by 3 to 10 percentage points (line (vi) versus (v) of Table V and Table VI, respectively). With a small amount of labelled data representing less than 2% of all available data¹, SSL even surpasses the supervised approach on the later campaign B_{test} (line (vi) versus (iii) of Table V and Table VI, respectively), in accordance with the results of the previous section. A study including an increased number of annotated images used for fine-tuning to a factor of 2, 5, and 10 times the number of annotated data used with SSL is carried out. The results

¹ 144 annotated images from \mathcal{D}^1 used in line (vi) of Table V, from a total of 7879 images; 216 from \mathcal{D}^2 used in line (iii) of Table VI, from a total of 15130.

are reported in Figure 5. The fine-tuning on \mathcal{D}^2 requires 4 times more labelled data to reach the same performance as SSL, while on \mathcal{D}^1 , the performance of fine-tuning is still significantly below, even with 10 times more annotated data. We also report the results for the supervised learning approaches with 2304 annotated data on \mathcal{D}_1 and 4606 ones on \mathcal{D}_2 , that are presented in detail in section subsection III-D.

These experiments on \mathcal{D}^1 and \mathcal{D}^2 show that it is in the tokamak operating team’s best interest to apply semi-supervised learning rather than fine-tuning. It should however be noted that the SSL approach includes a burn-in stage that consists in 80 000 iterations with data from both campaigns, while the fine-tuning approach can directly use the model already available from a previous campaign and compute the adaptation during only 10 000 iterations. The need for the burn-in phase can make the SSL less attractive than fine-tuning for practical usage.

F. SSL Enables Rapid Campaign Adaptation

The combination of fine-tuning with SSL during the second step of learning (method 2 in Section II) is applied to address the limits of the previous section. It consists in starting from a model available from a previous campaign (campaign A), fine-tuning it on the novel campaign (campaign B) then continuing to learn with unlabelled data from the novel campaign. Note that the step of fine-tuning is required since the model from the previous campaign is not reliable enough to create good pseudo-labels. The results of this approach are reported on line (vii) of Table V and Table VI.

The performance is better (on \mathcal{D}^1) or on-par (on \mathcal{D}^2) with a supervised approach with the same amount of labelled data (line (iii) of Table V and Table VI) but do not need to retrain the model from scratch. In comparison to the SSL approach only (method 1, line (vi) of Table V and Table VI) the results are globally equivalent, depending on the considered test set. As a consequence, this second version of the approach that relies on SSL still surpasses the fine-tuning only approach (line (v) of Table V and Table VI).

Applying semi-supervised learning is therefore a labelling cost-effective way to adapt a model trained on a certain experimental campaign to a new one and rapidly get good detection performance on images from the new campaign. The pairing of fine-tuning and semi-supervised learning even achieves the best mAP value after an experiment including semi-supervised learning on \mathcal{D}^2 , the larger and more realistic dataset.

IV. CONCLUSION AND PERSPECTIVES

This article explored the application of semi-supervised learning to infrared cameras used to monitor the inside of fusion reactors during experimental campaigns. Semi-supervised learning can help quickly adapt a model trained on a certain acquisition campaign to a new one with a lesser data labelling cost. This work establishes that semi-supervised learning outperforms fine-tuning when faced with a new campaign, and that semi-supervised learning can achieve near fully supervised performance but with extensively fewer annotations.

Table V
EXPERIMENTAL RESULTS ON \mathcal{D}^1

Experiment number	Initial training		Fine-tuning		Semi-supervised learning		Test set	mAP
	Number of images	Set	Number of images	Set	Number of images	Unlabelled set		
(i)	7 879	$A_{\text{train}}^1 + B_{\text{train}}^1$	–	–	–	–	A_{test}^1 B_{test}^1 both	0.397 0.477 0.434
(ii)	2 304	B_{train}^1	–	–	–	–	B_{test}^1	0.435
(iii)	96/48	$A_{\text{train}}^1 / B_{\text{train}}^1$	–	–	–	–	A_{test}^1 B_{test}^1	0.282 0.415
(iv)	96	A_{train}^1	–	–	–	–	A_{test}^1 B_{test}^1	0.267 0.020
(v)	96	A_{train}^1	48	B_{train}^1	–	–	A_{test}^1 B_{test}^1	0.202 0.351
(vi)	96/48	$A_{\text{train}}^1 / B_{\text{train}}^1$	–	–	2 256	\bar{B}_{train}^1	A_{test}^1 B_{test}^1	0.357 0.456
(vii)	96	A_{train}^1	48	B_{train}^1	2 256	\bar{B}_{train}^1	A_{test}^1 B_{test}^1	0.355 0.416

Table VI
EXPERIMENTAL RESULTS ON \mathcal{D}^2

Experiment number	Initial training		Fine-tuning		Semi-supervised learning		Test set	mAP
	Number of images	Set	Number of images	Set	Number of images	Unlabelled set		
(i)	15 130	$A_{\text{train}}^2 + B_{\text{train}}^2$	–	–	–	–	A_{test}^2 B_{test}^2 both	0.425 0.320 0.385
(ii)	4 606	B_{train}^2	–	–	–	–	B_{test}^2	0.330
(iii)	144/72	$A_{\text{train}}^2 / B_{\text{train}}^2$	–	–	–	–	A_{test}^2 B_{test}^2	0.389 0.291
(iv)	144	A_{train}^2	–	–	–	–	A_{test}^2 B_{test}^2	0.374 0.015
(v)	144	A_{train}^2	72	B_{train}^2	–	–	A_{test}^2 B_{test}^2	0.12 0.272
(vi)	144/72	$A_{\text{train}}^2 / B_{\text{train}}^2$	–	–	4 534	\bar{B}_{train}^2	A_{test}^2 B_{test}^2	0.388 0.302
(vii)	144	A_{train}^2	72	B_{train}^2	4 534	\bar{B}_{train}^2	A_{test}^2 B_{test}^2	0.372 0.317

A little over six mAP percentage points separate at most the performance of the trained semi-supervised model with the performance of the model trained on all data on the labelled datasets of tokamak WEST but this performance is accomplished with up to 70 times fewer annotations.

Semi-supervised learning is therefore an attractive tool for future scientific experimental campaigns. This article focuses on the wide-angle viewing line of sight of the WEST tokamak. Experiments including semi-supervised learning could be carried out on other lines of sight available among the 12 cameras installed on WEST as a continuation of this work. Other possibilities include adapting a model trained on an experimental campaign of WEST to other fusion devices or in the near future to the International Thermonuclear Experimental Reactor (ITER) [2].

DATA AVAILABILITY

The data that supports the findings of this study are available from the corresponding author upon reasonable request, according to the CEA general data-sharing framework, or through formalized collaboration agreement.

REFERENCES

- [1] F. Hofmann *et al.*, “Creation and control of variably shaped plasmas in TCv,” *Plasma Physics and Controlled Fusion*, vol. 36, p. B277, Dec. 1994.
- [2] Y. Shimomura, R. Aymar, V. Chuyanov, M. Huguet, R. Parker, and ITER Joint Central Team, “ITER overview,” *Nuclear Fusion*, vol. 39, p. 1295, Sept. 1999.
- [3] S. Wu, “An overview of the EAST project,” *Fusion Engineering and Design*, vol. 82, no. 5, pp. 463–471, 2007. Proceedings of the 24th Symposium on Fusion Technology.
- [4] J. Bucalossi *et al.*, “Operating a full tungsten actively cooled tokamak: overview of WEST first phase of operation,” *Nuclear Fusion*, vol. 62, p. 042007, Feb. 2022.

- [5] J. Traversé *et al.*, “Image analysis of infrared quantitative data applied to tokamak survey and safety: present and future,” in *IEEE International Symposium on Intelligent Signal Processing (WISP)*, pp. 1–5, Oct. 2007.
- [6] V. Martin, F. Bremond, J.-M. Traversé, V. Moncada, and G. Dunand, “Thermal event recognition applied to tokamak protection during plasma operation,” in *2009 IEEE Instrumentation and Measurement Technology Conference*, pp. 1690–1694, 2009.
- [7] E. Grelier, R. Mitteau, and V. Moncada, “Deep learning and image processing for the automated analysis of thermal events on the first wall and divertor of fusion reactors,” *Plasma Physics and Controlled Fusion*, vol. 64, p. 104010, Sept. 2022.
- [8] Z. Cai and N. Vasconcelos, “Cascade r-cnn: Delving into high quality object detection,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6154–6162, 2018.
- [9] E. Grelier, R. Mitteau, and V. Moncada, “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*, vol. 192, p. 113636, 2023.
- [10] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems* (C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, eds.), vol. 28, Curran Associates, Inc., 2015.
- [11] M. Szűcs *et al.*, “A deep learning-based method to detect hot-spots in the visible video diagnostics of Wendelstein 7-X,” *Journal of Nuclear Engineering*, vol. 3, no. 4, pp. 473–479, 2022.
- [12] G. Jocher *et al.*, “ultralytics/yolov5: v3.1 - Bug Fixes and Performance Improvements,” Oct. 2020.
- [13] S. Fabbrizzi, S. Papadopoulos, E. Ntoutsis, and I. Kompatsiaris, “A survey on bias in visual datasets,” *Computer Vision and Image Understanding*, vol. 223, p. 103552, 2022.
- [14] A. Torralba and A. A. Efros, “Unbiased look at dataset bias,” in *CVPR 2011*, pp. 1521–1528, 2011.
- [15] M. Xu *et al.*, “End-to-end semi-supervised object detection with Soft Teacher,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3040–3049, 2021.
- [16] Y. Ouali, C. Hudelot, and M. Tami, “An overview of deep semi-supervised learning,” 2020.
- [17] N. Karaliolios, F. Chabot, C. Dupont, H. L. Borgne, R. Audigier, and Q.-C. Pham, “Generalized pseudo-labeling in consistency regularization for semi-supervised learning,” in *Int. Conf. on Image Processing*, 2023.
- [18] K. Sohn, Z. Zhang, C.-L. Li, H. Zhang, C.-Y. Lee, and T. Pfister, “A simple semi-supervised learning framework for object detection,” 2020.
- [19] Y.-C. Liu *et al.*, “Unbiased teacher for semi-supervised object detection,” 2021.
- [20] A. Tarvainen and H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” in *Advances in Neural Information Processing Systems* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), vol. 30, Curran Associates, Inc., 2017.
- [21] Q. Xie, Z. Dai, E. Hovy, T. Luong, and Q. Le, “Unsupervised data augmentation for consistency training,” in *Advances in Neural Information Processing Systems* (H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, eds.), vol. 33, pp. 6256–6268, Curran Associates, Inc., 2020.
- [22] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, “RandAugment: Practical automated data augmentation with a reduced search space,” 2019.
- [23] T. DeVries and G. W. Taylor, “Improved regularization of convolutional neural networks with Cutout,” 2017.
- [24] A. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, “A brief review of domain adaptation,” in *Advances in Data Science and Information Engineering*, (Cham), pp. 877–894, Springer International Publishing, 2021.
- [25] Y. Tamaazousti, H. Le Borgne, and C. Hudelot, “Mucale-net: Multi categorical-level networks to generate more discriminating features,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (Los Alamitos, CA, USA), pp. 5282–5291, IEEE Computer Society, July 2017.
- [26] Y. Tamaazousti, H. Le Borgne, C. Hudelot, M.-E.-A. Seddik, and M. Tamaazousti, “Learning more universal representations for transfer-learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 9, pp. 2212–2224, 2020.
- [27] K. Chen *et al.*, “MMDetection: Open MMLab detection toolbox and benchmark,” *arXiv preprint arXiv:1906.07155*, 2019.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [29] T.-Y. Lin *et al.*, “Microsoft COCO: Common Objects in Context,” in *Computer Vision – ECCV 2014* (D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, eds.), (Cham), pp. 740–755, Springer International Publishing, 2014.
- [30] X. Courtois *et al.*, “Full coverage infrared thermography diagnostic for west machine protection,” *Fusion Engineering and Design*, vol. 146, pp. 2015–2020, 2019. SI:SOFT-30.
- [31] Q. Zhou, C. Yu, Z. Wang, Q. Qian, and H. Li, “Instant-teaching: An end-to-end semi-supervised object detection framework,” 2021.