



HAL
open science

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

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

► **To cite this version:**

Christian Staron, Hervé Le Borgne, Raphaël Mitteau, Nicolas Allezard, Erwan Grelier. Detection of Thermal Events by Semi-Supervised Learning for Tokamak First Wall Safety. 2023. cea-04099065v1

HAL Id: cea-04099065

<https://cea.hal.science/cea-04099065v1>

Preprint submitted on 16 May 2023 (v1), last revised 18 Jan 2024 (v2)

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⁽¹⁾, Hervé Le Borgne⁽²⁾, Raphaël Mitteau⁽¹⁾, Nicolas Allezard⁽²⁾, Erwan Grelier⁽¹⁾

(1) CEA, IRFM, F-13108 St Paul Lez Durance, France

(2) Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

Abstract—This paper explores a semi-supervised object detection approach to detect hot spots 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 analysed 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 databases which 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 performances are 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 employed method.

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]. 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 localised heating. Hot spots appear on the first wall, and can damage these in-vessel components when exceeding armour material allowables. Among numerous 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.

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.

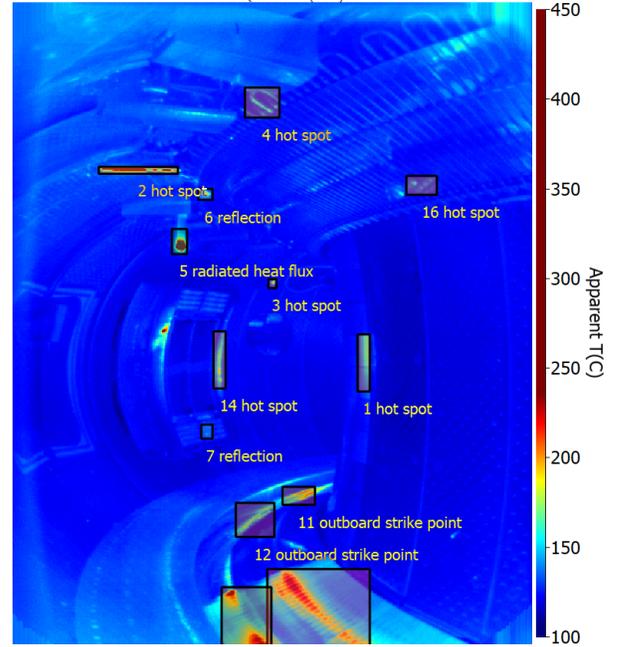


Figure 1. Annotation examples of several types of hot spots on the wide-angle line of sight of the Tokamak WEST [4]

The resulting live image feeds are used for real-time feedback control during operation, and also for science after the experiments. In that last case, the experts aim at localizing and characterizing the different types of hot spots 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 thresholding and clustering [5]. More recent approaches rely on large neural networks learned from visual data [6], [7], such as Faster R-CNN [8]. The results are much more informative, they nevertheless still require tedious manual work to annotate the IR video films (Figure 1).

This article proposes an automated process that detects and classifies thermal events on the internal wall of Tokamaks that requires a significant lesser amount of manual work while maintaining efficiency. The large amount of unlabelled data produced by the IR cameras is taken advantage of, in the vein of the semi-supervised principles [9]. 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 performances.

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 performances 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.

In the following, Section II presents the data handled in the study, that result from two different acquisition campaigns. The method based on semi-supervised learning is described in III. Several experimental results are reported in Section IV, showing the interest of the semi-supervised approaches with regards to the compromise between manual annotation effort and detection performances on the one hand, and the ability to adapt to a new domain resulting from a new acquisition campaign on the other hand. Last, the main results are presented, including potential limitations to the method and possible future works in Section V.

II. WEST’S LABELLED IMAGE DATASETS

The work of this article is based on images from the infrared cameras of Tokamak WEST [4]. The monitoring of WEST’s PFC and more generally its first wall is provided by a total of 12 infrared viewing systems [10]. These viewing systems are used for science and real-time feedback control during operation [11]. Each one is installed at different locations to capture multiple fields of view. The images for this study come from the wide-angle (WA) line of sight of WEST. They are illustrated in Figure 1. The WA line of sight provides an overall view of the Tokamak and thus has a key role for the protection of the machine. Moreover, this view has the largest variety of hot spots, while other views may contain only a restricted number of the types of hot spots further considered.

A WEST fusion experiment typically lasts between 30 and 60 seconds, during which each camera produces a video film. Acquisition campaigns are organized on a regular basis separated by several months to maintain the Tokamak. The images resulting from the videos are of size 512×640 px and encoded over 16 bits.

The hot spots are labelled according to seven classes for the purpose of this work, while a full system would have many more classes. Each label consists of four integers corresponding to the surrounding box of the spot, and one to the class that characterizes its type. The six first classes correspond to well defined wall events and the seventh is the catchall class, that is appropriate when the exact definition of the hot spot is undefined. In practice, with progresses in plasma fusion with regard to the identification of these hot spots, new classes can be defined in the future and include a part of the annotations in this catchall class. The different types of labelled hot spots, illustrated in Figure 2, are the following:

Table I
THE DIFFERENT HOT SPOT CLASSES IN THE LABELLED DATASETS

Class number	Class names	Number of labels 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

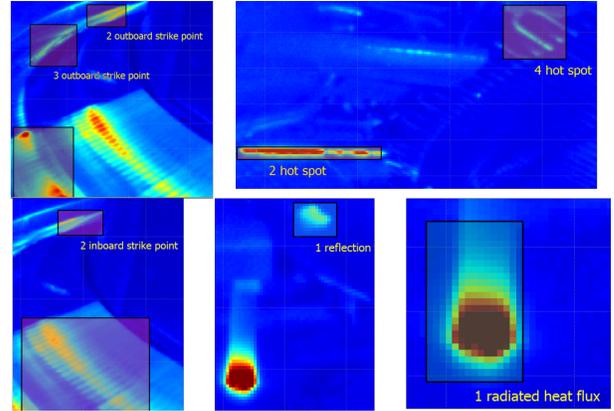


Figure 2. Examples of different hot spot annotations (from left to right: outboard strike point, hot spot, inboard strike point, reflection, radiated heat flux)

- ‘electron type 1’: a hot spot that appears on the lateral protections of antennas. Its appearance is correlated with the power injected by the antenna.
- ‘inboard strike point’: a strike line located on the inboard part of the divertor.
- ‘outboard strike point’: a strike line located on the outboard part of the divertor.
- ‘reflection’: the photonic reflection of a thermal event on the metallic wall of the vessel.
- ‘radiated heat flux’: a hot component, such as a Langmuir probe, heated by plasma radiation.
- ‘UFO’: a hot particle or speck of dust that moves from frame to frame.
- ‘hot spot’: a catchall class for events that do not belong to any of the other categories

Image labelling is performed with a user-friendly annotation tool detailed in [6]. This tool assists experts during the annotation task, allowing to reduce annotation time and inconsistencies between annotators. Labelling one IR film takes between 20 and 30 minutes. Hence, labelling the whole dataset 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 [12]. The main potential subject for hot spot labelling is the *label bias* due to the fact that “different annotators may assign different labels to the same type of [hot spot]” [13]. Since the knowledge in plasma fusion progresses with time, even the

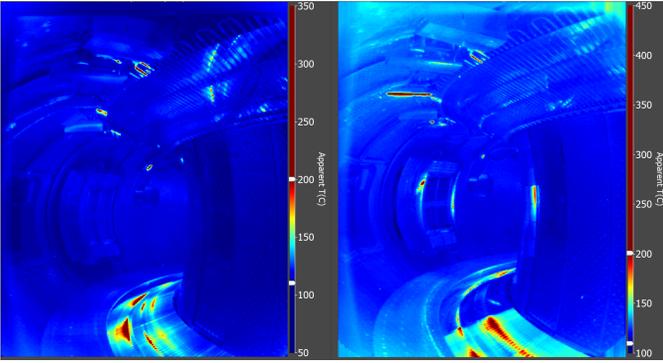


Figure 3. Image from campaign *A* (left) during an experiment with similar experimental conditions as the image for campaign *B* (right)

same experts can make subtle changes in their assessments to categorize hot spots, between two annotation sessions several months apart. Two annotation sessions were organised to study such potential biases. \mathcal{D}^1 is made of eight video films annotated in September 2021, corresponding to 12 156 images. \mathcal{D}^2 was annotated by the same expert in August 2022, on 21 other video films, corresponding to 18 336 images. A unique expert annotated both \mathcal{D}^1 and \mathcal{D}^2 to mitigate the risk of inter-annotator variation.

WEST's experimental campaigns last over several months. Hundreds of experiments are conducted during these experimental campaigns for the development of multiple plasma scenarios. WEST goes through many different plasma configurations during these campaigns and evolves from one campaign to another (adding new components, new diagnostics...). Such changes make the images acquired by the IR cameras quite different from one campaign to another, as illustrated in Figure 3. Such changes may not be detected by human experts but have a significant impact on the low-level statistics of the images and, as a result, on the performances of the automatic vision systems. Each labelled dataset was made 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* which was a shorter campaign between November 2020 and January 2021.

Each of the two annotated datasets are 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 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 II gives the number of images in the training and test sets, for both datasets \mathcal{D}^1 and \mathcal{D}^2 . 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 and $\beta \in \{\text{train}, \text{test}\}$ is the split in the following. \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.

III. PROPOSED APPROACH

Detecting and classifying hot spots 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 localise 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.

Recent works proposed to estimate such a model with supervised learning [6], [14]. It relies on a large labelled dataset $L_{\text{train}}^{\text{sup}}$ and minimises a loss, which the exact form depends on the outputs design of the model. Grelier *et al.* trained a Faster R-CNN [8]. The same base model is selected in the following for fair comparison. In that case, the loss is the sum of log losses for classification and smooth L_1 losses for regression. The regression allows to estimate the t-uplet (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 (SGD).

A. Semi-supervised learning to detect hot spots

The main limit of the supervised approach is that it requires a large dataset $L_{\text{train}}^{\text{sup}}$, that 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} .
- 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]–[17] this work relies on a student-master architecture that comprises two branches. The master 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 master model is an exponential mean average (EMA) [18] of the student model. Its parameters θ^{master} are updated

Table II
COMPOSITION OF THE LABELLED DATASETS CONSIDERED IN THIS STUDY

	\mathcal{D}^1		\mathcal{D}^2	
	Train	Test	Train	Test
Nb. films from A	2	1	8	1
Nb. films from B	3	2	9	3
Nb. images (100%)	7 879	4 277	15 130	3 206
Notation	$A_{\text{train}}^1 + B_{\text{train}}^1$	$A_{\text{test}}^1 + B_{\text{test}}^1$	$A_{\text{train}}^2 + B_{\text{train}}^2$	$A_{\text{test}}^2 + B_{\text{test}}^2$

according to $\theta^{\text{master}} \leftarrow \alpha\theta^{\text{master}} + (1 - \alpha)\theta^{\text{student}}$ at each iteration, where α is a hyperparameter slightly less than one.

Applying data augmentations to input images is a well-known efficient strategy in SSL [19]. They enable a diversification particularly of the unlabelled data so that the master 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 master 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 [20]. The image can be for instance sharpened, equalised, 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 [21].

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 [17] that strongly affect the pixel values are thus not applied. This includes augmentations that equalise the image, solarise it or mixes the colour channels. The weak augmentations are unchanged. They are the same augmentations used when training a supervised model at 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 in translating, shearing and erase 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 which affect the location of the power deposition.

An explanation on how the images are collected from two such campaigns, namely A and B is in Section II. As shown experimentally in Section IV-B, when a detection/classification model is trained with data from campaign A , its performances drop dramatically when applied directly on data of campaign B . For a new campaign, the parameters θ_A previously learnt need to be adapted to new parameters θ_B that provide acceptable performances. Such an issue can be seen as a domain adaptation problem [22] that can be addressed with transfer learning approaches such as using universal backbones [23], [24].

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 performances 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 performances of a model that would be learnt from scratch on images from both A and B .

This article claims 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 to retrain 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 learnt previously in campaign A is directly adapted to campaign B without training from scratch.

C. Implementation details

The implementation of the model in this work is built upon the code of SoftTeacher [17], that itself relies on MMDDetection [25]. The (student) detection model is a Faster R-CNN with a ResNet-50 [26] backbone. The setup of the hyperparameters results either from the values proposed by [17] or thanks to the grid search conducted on a $2/3 / 1/3$ split of the training dataset of \mathcal{D}^1 . Hence, the relative weight λ of the unsupervised loss is fixed at 4.0 during Step 2. Table III exposes the other main hyperparameters used to train both

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

	Supervised training (burn-in)	SSL training
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	

the supervised models and the semi-supervised models. The models are trained on 4 NVIDIA Tesla V100 16 GB GPU.

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 [15], [16], usually on the RGB images of MSCOCO [27]. The decision to train models on 1% relies on the fact that it is the hardest case because the least amount of 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.

It is chosen to use a number of images that can be equally distributed amongst the labelled video films during training. The same number of images are extracted from each labelled film to have a well defined representation of each film in the training set. In dataset \mathcal{D}^1 , for an experiment involving training on a number of images equivalent to 1% of the dataset, 96 labelled images are exploited. The larger dataset \mathcal{D}^2 handles 144 labelled images when supervised models are trained on a number of images comparable to 1% of \mathcal{D}^2 . The number of images taken for \mathcal{D}^1 corresponds in reality to 1.02% and the number of images for \mathcal{D}^2 coincides with 0.95%.

IV. EXPERIMENTAL RESULTS ON WEST’S DATASETS

A. Evaluation metric

The mean Average Precision (mAP) is used to measure the performances on the infrared images. It is a reference metric in object detection that both reflects the localisation and classification ability of the model. 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 [28], the mAP is only computed for boxes with a confidence higher than 10^{-3} .

B. Inter-campaigns and Annotations Domain Shift

Section II explains that the datasets have been made to study the transfer of a model learned on the data of campaign *A* to the test data of a campaign *B*. Firstly, a series of experiments

Table IV
mAP FOR DIFFERENT CAMPAIGNS AND DATASETS

Experiment number	Initial training		Test set	mAP
	Number of images	set		
(a)	5 575	A^1_{train}	A^1_{test}	0.381
(b)	5 575	A^1_{train}	B^1_{test}	0.030
(c)	2 304	B^1_{train}	A^1_{test}	0.037
(d)	2 304	B^1_{train}	B^1_{test}	0.435
(e)	7 879	A^1_{train} $+B^1_{\text{train}}$	A^1_{test} $+B^1_{\text{test}}$	0.434
(f)	7 879	A^1_{train} $+B^1_{\text{train}}$	A^2_{test} $+B^2_{\text{test}}$	0.087
(g)	15 130	A^2_{train} $+B^2_{\text{train}}$	A^2_{test} $+B^2_{\text{test}}$	0.385
(h)	15 130	A^2_{train} $+B^2_{\text{train}}$	A^1_{test} $+B^1_{\text{test}}$	0.168
(j)	4 606	B^2_{train}	A^2_{test}	0.030
(k)	4 606	B^2_{train}	B^2_{test}	0.330
(l)	10 524	A^2_{train}	A^2_{test}	0.396
(m)	10 524	A^2_{train}	B^2_{test}	0.017

in a fully labelled setting are conducted. The reporting of the results is in Table IV.

The results are quite good (lines (a), (d) and (e) 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 performances. This may be due to a larger training dataset or more relevant images in B^1 . However, the performances drop dramatically to near zero (lines (b) and (c)) when the training and test data are from different campaigns. These experiments show that the changes in plasma configurations from one campaign to another produces subtle changes on the low-level statistics of the images that are not visible to experts (see Figure 3) but have a significant impact on performances, 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 if the annotators are the same, the criterion to classify the hot spots by experts changed slightly, introducing minor differences between the two. Good performances are still attained when the training and test data is homogeneous (line (g), (k) and (l)) but a significant drop is noted when the training set and the test one come from different annotation sessions (lines (f) and (h)). The fact that the model learned on \mathcal{D}^2 has better performances than that trained on \mathcal{D}^1 shows that later data encapsulates more earlier ones. The domain shift identified in \mathcal{D}^1 is nevertheless still present in \mathcal{D}^2 (lines (j) to (m)).

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

C. SSL Gets Near Supervised Performances at Lower Cost

This section compares the approach based on semi-supervised learning (SSL) and the classical supervised setting [6], [14]. Experiments are conducted both on \mathcal{D}^1 and \mathcal{D}^2 . The results are reported in resp. Table V and Table VI) with the same experiments for both datasets, numbered from (i) to (vii).

The better possible performances are 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*). These performances vary from 0.32 to 0.477 (line (i) in both tables) depending on the test dataset. The performances logically decrease 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 performances drop 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 4.2 mAP percentage points on \mathcal{D}^1 with $\frac{7879}{96+48} = 54$ times less data and 6.1 percentage points on \mathcal{D}^2 with $\frac{15130}{144+72} = 70$ 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_{train}^1 (0.456 vs 0.435) when both are tested on 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 performances at a much smaller labelling cost.

D. Semi-Supervised 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*, then fine tuned during 10 000 iterations with 0.5% data from campaign *B*. Fine tuning allows the expected adaptation, since the performance on B_{test} are no longer almost null (line (v) versus (iv)). There is an expected decrease of performances on A_{test} as fine-tuning specialises the model's parameters to the new campaign *B* (line (v)). Fine tuning remains however inferior to the model trained with the same amount of annotated data from *A* and *B* (line (iii) versus (v)).

For both datasets, the same amount of labelled data is handled but additive unlabeled data is used for the SSL model trainings. The SSL approach surpasses fine-tuning by 3 to 10 percentage points (line (vi) versus (v)). With a small amount of labelled data (representing less than 2% of all available data), SSL even surpasses the supervised approach on the later campaign B_{test} (line (vi) versus (iii)), 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 $\times 2 / \times 5 / \times 10$ the number of annotated data used with SSL is carried out. The results are reported in Figure 4. The fine-tuning on \mathcal{D}^2 requires 4 times more labelled data to reach the same performances as SSL, while on \mathcal{D}^1 , the performances of fine-tuning are still significantly below, even with 10 times more annotated data.

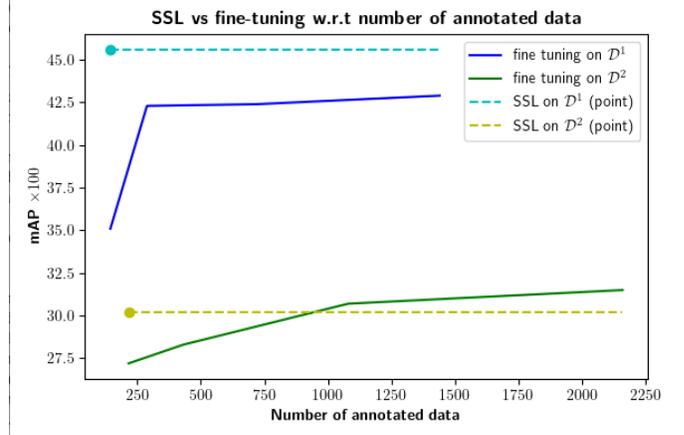


Figure 4. Performances of fine-tuning w.r.t the number of annotated data used in \mathcal{D}^1 (blue) and \mathcal{D}^2 (green). The performances of SSL are given for a unique number of annotated data (cyan/yellow point)

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. However in practice, 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.

E. Semi-supervised learning enables rapid campaign adaptation

The combination of fine-tuning with SSL during the second step of learning (method 2 in Section III) 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 continue 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 is reported on line (vii) of Table V and Table VI.

The performances are 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)) but do not need to retrain the model from scratch. In comparison to the SSL approach only (method 1, line (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)).

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 performances 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.

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

V. CONCLUSION AND PERSPECTIVES

This article explores 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 also establishes that semi-supervised learning outperforms fine-tuning when faced with a new campaign, and that semi-supervised learning can achieve near fully supervised performances but with extensively less annotations. A little over six mAP percentage points separate at most the performances of the trained semi-supervised model with the performances of the model trained on all data on the labelled datasets of Tokamak WEST but these performances are accomplished with up to 70 times less annotations.

Semi-supervised learning is therefore an attractive tool for future scientific experimental campaigns. This article focused

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 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 datasharing framework, or through formalized collaboration agreement.

REFERENCES

- [1] F. Hofmann, J. B. Lister, *et al.*, “Creation and control of variably shaped plasmas in tcv,” *PPCF*, vol. 36, p. B277, dec 1994.

- [2] Y. Shimomura, R. Aymar, V. Chuyanov, *et al.*, "ITER overview," *Nuc. Fus.*, vol. 39, p. 1295, sep 1999.
- [3] S. Wu, "An overview of the east project," *Fus. Eng. and Des.*, 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," *Nuc. Fus.*, vol. 62, p. 042007, Apr 2022.
- [5] V. Martin, F. Bremond, J. Travers, *et al.*, "Thermal event recognition applied to tokamak protection during plasma operation," in *2009 IEEE Instrumentation and Measurement Technology Conference*, pp. 1690–1694, 2009.
- [6] 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," *PPCF*, vol. 64, p. 104010, oct 2022.
- [7] M. Szűcs, T. Szepesi, C. Biedermann, *et al.*, "A deep learning-based method to detect hot-spots in the visible video diagnostics of wendelstein 7-x," *Jour. of Nuc. Eng.*, vol. 3, no. 4, pp. 473–479, 2022.
- [8] 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, *et al.*, eds.), vol. 28, Curran Associates, Inc., 2015.
- [9] Y. Ouali, C. Hudelot, and M. Tami, "An Overview of Deep Semi-Supervised Learning," Tech. Rep. arXiv:2006.05278, arXiv, July 2020. arXiv:2006.05278 [cs, stat].
- [10] X. Courtois, M. Aumeunier, C. Balorin, *et al.*, "Full coverage infrared thermography diagnostic for WEST machine protection," *Fus. Eng. and Des.*, vol. 146, pp. 2015–2020, sep 2019.
- [11] R. Mitteau, C. Belafdil, C. Balorin, *et al.*, "WEST operation with real time feed back control based on wall component temperature toward machine protection in a steady state tungsten environment," *Fus. Eng. and Des.*, vol. 165, p. 112223, Apr. 2021.
- [12] S. Fabbrizzi, S. Papadopoulos, E. Ntoutsis, *et al.*, "A survey on bias in visual datasets," *Computer Vision and Image Understanding*, vol. 223, p. 103552, 2022.
- [13] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *CVPR 2011*, pp. 1521–1528, 2011.
- [14] 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," *Fus. Eng. and Des.*, vol. 192, p. 113636, Jul 2023.
- [15] K. Sohn, Z. Zhang, C. Li, *et al.*, "A simple semi-supervised learning framework for object detection," 2020.
- [16] Y. Liu, C. Ma, Z. He, *et al.*, "Unbiased teacher for semi-supervised object detection," *CoRR*, vol. abs/2102.09480, 2021.
- [17] M. Xu, Z. Zhang, H. Hu, *et al.*, "End-to-end semi-supervised object detection with soft teacher," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [18] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. of the 31st Int. Conf. on Neural Information Processing Systems, NIPS'17*, (Red Hook, NY, USA), p. 1195–1204, Curran Associates Inc., 2017.
- [19] Q. Xie, Z. Dai, E. H. Hovy, *et al.*, "Unsupervised data augmentation for consistency training," in *NeurIPS*, 2020.
- [20] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "RandAugment: Practical automated data augmentation with a reduced search space," 2019.
- [21] T. DeVries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," 2017.
- [22] A. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, "A brief review of domain adaptation," in *Advances in Data Science and Information Engineering*, pp. 877–894, Springer International Publishing, 2021.
- [23] Y. Tamaazousti, H. Le Borgne, and C. Hudelot, "Mucal-net: Multi categorical-level networks to generate more discriminating features," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (Honolulu), 2017.
- [24] Y. Tamaazousti, H. Le Borgne, C. Hudelot, *et al.*, "Learning more universal representations for transfer-learning," *IEEE T. Pattern Analysis and Machine Intelligence*, vol. 42, no. 9, pp. 2212–2224, 2020.
- [25] K. Chen, J. Wang, J. Pang, *et al.*, "MMDetection: Open MMLab Detection Toolbox and Benchmark," *arXiv:1906.07155 [cs, eess]*, June 2019.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [27] T.-Y. "Lin, M. Maire, S. Belongie, and others", ""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".
- [28] Q. Zhou, C. Yu, Z. Wang, *et al.*, "Instant-teaching: An end-to-end semi-supervised object detection framework," *CoRR*, vol. abs/2103.11402, 2021.