CONFERENCE PRE-PRINT

WEST ADVANCED WALL PROTECTION ACHIEVEMENTS TOWARD LONG PULSE OPERATION

R. MITTEAU, M.H. AUMEUNIER, L. DUBUS, J. GERARDIN, V. GORSE, E. GRELIER, V. MONCADA, S.VIVES, X. LITAUDON and the WEST Team

CEA, IRFM
Saint-Paul-lez-Durance, France
Email: raphael.mitteau@cea.fr
For the WEST team, see http://west.cea.fr/WESTteam

M. JAKUBOWSKI, J. FELLINGER
Max-Planck-Institut für Plasmaphysik
Greifswald, Germany

Abstract

Long pulse operation in magnetic fusion devices requires well controlled plasma power exhaust to the divertor & wall, and avoidance of wall hot spots that could evolve in wall damage. At WEST, 10 major plasma facing components are monitored using 10 series of temperature/power indicators, based on multiple diagnostic systems, among which the infrared viewing system is especially relevant. These indicators span from the most basic ones (temperatures, power and energy from deterministic models) to advanced processes using artificial intelligence acquired through machine learning. Some advanced processes do operate in real time, and feedback on power actuators through the plasma control system, providing active control toward remaining within the safe operational domain. Some intervene as forensic tools post discharge to identify possible dangerous situation regarding the power loading to the wall, so that the discharge plan is adjusted to avoid running into potential wall events. No critical wall power event happened during the campaigns C9 to C11 (2024-2025), totalling about 13h of plasma, that would have affected the campaigns operational plan. While it cannot be demonstrated that the active & intelligent wall protection enabled the new plasma duration record of 1337 seconds, it is bound that the wall protection system as a whole helped significantly obtaining this record by preventing wall hot spots to become critical during the campaigns.

1. INTRODUCTION

Exploiting and extending WEST capability to long pulse operation represents a large part of WEST's experimental activity during the campaigns C9 to C11 (2024-2025), totalling about 13h of plasma [1]. Long Pulse Operation (LPO) in magnetic fusion plasma devices requires well-controlled plasma power exhaust to the In Vessel Components (IVCs), essentially divertor & main chamber wall. Any deviation to the carefully engineered power exhaust scheme results in wall hot spots and possible damage to the IVCs. This effect is exacerbated by plasma duration, because LPO is enabled by employing actively cooled components, with a specific operation mode. Their operation domain is primarily limited by the surface heat flux (q, [W/m²]) and power (P, [W]), in addition to the standard & customary temperature limits commonly used for non-actively cooled IVCs. Flux quantities govern armour delamination and coolant critical heat flux (for water-cooled components), while these flux quantities are of much less concern for non-actively cooled IVCs. LPO requires a control over the engineering heat fluxes at the component surface, but no diagnostic produces actual measures of engineering "q" or "P". Hence having good estimates of the fluxes is essential. Moreover, actively cooled components have severe damage modes because of the cooling circuit, making the detection of abnormal thermal events essential.

At WEST, 10 major IVCs (lower and upper divertor, baffle, inner guard limiters, five Lower Hybrid (LH) and ion cyclotron heating antennas, dump plates, Fig. 1) are monitored using 10 series of power / temperature indicators, largely based on the image streams from the InfraRed (IR) viewing system. They span from the most basic ones (heat fluxes, temperatures, power and energy from deterministic models) to advanced processes using neural networks and machine learning techniques. Some processes operate in Real Time (RT) and feedback on actuators through the Plasma Control System (PCS), providing active strategies towards remaining within the operational domain [2]. Other processes intervene post discharge as forensic tools to identify possible deteriorations of the wall and divertor under power loading. They are used for the preparation of the following discharge.

The infrared viewing system is a key contributor to IVCs protection. This is because of its unique combination of being a quantitative measure, being produced as a 2D + time data stream. The infrared viewing systems provide the radiance of the hot thermal scene of the IVCs. WEST is equipped with 6 IR viewing systems which main

objective is wall safety, while other IR viewing systems are used only for science and are not further mentioned here. The WEST infrared system is thoroughly described in [3].

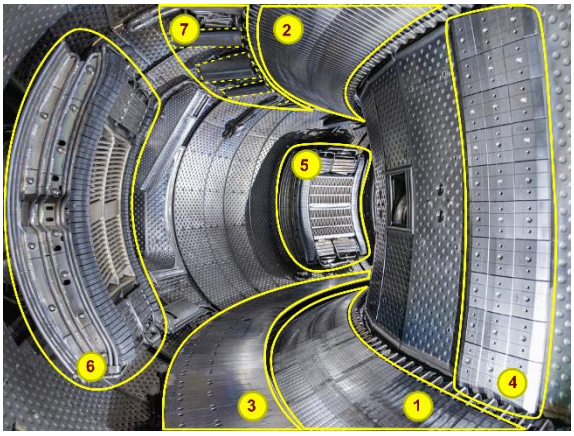

Fig. 1: WEST inner vessel (2024) with 7 of the 10 major internal components being monitored in infrared: 1) the lower divertor, 2) the upper divertor, 3) the lower divertor baffle, 4) the inner guard limiter, 5) the LH2 lower hybrid antenna, 6) the Q1 ICRH antenna, 7) the VDE and ripple dump plates.

While unrivalled as a wall health monitoring tool by the human wall monitoring experts (Plasma facing component Protection Officers, PPOs), the infrared images are complex to interpret. Obtaining the surface temperature from the scene radiance requires having sufficient knowledge of the thermal scene surface properties (emissivity ε, reflectivity r, possible presence of thermally resistive layers). However, all these properties are variable with location and time, and photonic properties such as ε and r depend on the light incidence angle. Hot emission sources outside of the Field of View (FoV) would affect the scene radiance, and should normally be taken into account. The conversion of the thermal scene to a surface temperature map is also affected by the predominantly diffusive or reflective nature of the wall. Obtaining a good enough estimation of IVC surface temperature gets tricky when ε becomes very low, which is the case in devices with metallic walls [4]. Moreover, converting the temperature history into the surface heat flux requires temperature flux inversion. This inversion is made uncertain by many uncertainties and complexities (non-linearities in material and heat transfer properties, complex geometries including armour castellations and segmentation, possible presence of thermally resistive slag layer at the surface of the IVCs, in addition to the uncertainties associated to ε and r already cited above). A comprehensive effort is underway at WEST, to untangle all these effects, though the use of inverse methods and synthetic diagnostic simulation [5]. This innovative method is due to be integrated in the WEST safety chain through a vessel thermal map, once mature enough. For the present purpose of machine safety, the estimation of the true component temperature is obtained by assuming a constant emissivity for each component, and neglecting the diffuse light reflexion. The emissivity is adjusted regularly, either by parameter estimation from cross diagnostic analysis or by in-vessel direct measures. The basic wall safety system relies merely on monitoring the surface temperatures in prescribed Regions Of Interest (ROI). They are compared to temperature limits, and, before reaching the limit, a controlled temperature band is set which feeds back on the power of the additional heating system. These temperature traces are shown in Fig. 2 for discharge #60738, along with the temperature control bands.

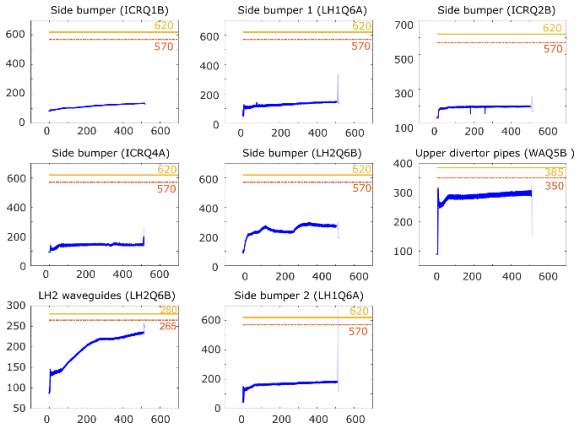

Fig. 2: Temperature plots of monitored regions of interest that define the operational domain of the LH antenna LH2 during the 514 seconds discharge #60738.

Above the maximum prescribed temperature, a hard-wired connection triggers a soft plasma termination. Such a security system based on temperature in ROIs has many shortcomings:

— It is a temperature-based rather than a flux-based control, with the above-cited drawback.
— Frequent and manual adjustment of ROIs are needed, because of changes in the field of view (positional drift of the line of sight, or component thermal response change).
— False alarms do happen, when an unexpected photonic reflexion or a hot chip of material passes through a ROIs.
— Certain safety actions should rather be triggered by specific shapes or patterns rather than a temperature.

Better systems should: a) take into account the shape and/or the time sequence of the event, as visible on the IR frames; b) use other diagnostic measures relevant to IVCs, such as local temperatures measures (thermocouples, fibber Bragg grating) [6], coolant calorimetry [7], and finally all machine data. These various data should then be merged. These processes are best done using engineering or physical models and/or Machine Learning (ML).

This paper focusses on the 7 advanced wall protection processes that either rely on ML techniques, or cross-diagnostics data merging: the thermal event detector, the divertor strike-lines descriptor, the Lower Hybrid (LH) heating system arc detector, the Runaway Electrons (REs) detector, the Unidentified Flying Objects (UFO) detector, the post discharge automatic expert assessment of the wall thermal response, and the RT estimation of heat flux to the divertor from RT power balance. They are detailed hereafter.

2. ADVANCED WALL PROTECTION TECHNIQUES AND METHODS

2.1. Model based estimation of the incident heat flux on selected internal components

Having solid estimates of the IVC thermal flux density is key to IVCs monitoring. The divertor engineering heat flux is obtained without the IR system using a model based estimator. The standard power distribution expression is used:

$$q_{div}^{eng} = r_{in-out} \cdot r_{ripple} \cdot \frac{\boldsymbol{B_{OSP}} \cdot \boldsymbol{n_{Div}}}{\|\boldsymbol{B_{OSP}}\|} \cdot \frac{p_{Cond}}{4 \cdot \pi \cdot r_{OSP} \cdot \lambda_q^{OMP}} \cdot \frac{\|\boldsymbol{B_{OMP}}\|}{\boldsymbol{B_{OMP}} \cdot \boldsymbol{e_z}}$$

$\boldsymbol{B}$ is the total magnetic field. The subscripts OMP and OSP stand for Outer MidPlane and the Outer Strike Point, respectively. p_{Cond} is the conducto-convective power exiting the plasma ($p_{Cond} = p_{ohmic} + p_{add} - p_{rad}$), r_{omp} is the major radius of the OMP , and λ_q^{OMP} is the heat flux decay length of the scrape-off layer at that same location. $\boldsymbol{n_{div}}$ and $\boldsymbol{e_z}$ are the unit vectors normal to the outer divertor target and to the horizontal plane, respectively. r_{in-out} is a scalar parameter expressing the in-out power deposition asymmetry (r_{in-out} = 4/3), and r_{ripple} is the peaking factor associated to the ripple of the toroidal magnetic field (r_{ripple} = 2 for WEST divertor). **Bold** quantities are vectors. p_{ohmic}, $\boldsymbol{B_{OSP}}$ and $\boldsymbol{B_{OMP}}$ are produced by the magnetic diagnostic cubicle, p_{add} by the additional heating systems cubicles, and p_{rad} by the bolometry diagnostic. The heat flux decay length is prescribed at 7.5 mm, according to [8]. It has little variation over the customary WEST operation range, and the divertor heat flux for another decay length can readily be obtained by proportional scaling. All significantly variable quantities are read by the wall monitoring system cubicle through the shared memory. The divertor heat flux is written in another section of that shared memory, being made available to other PCS processes. This basic model produces a RT estimation of the divertor incident heat flux, independently from the IR diagnostic. The RT calculated heat flux is shown at 4 MW/m² in Fig. 3, for discharge #61376. Its time variation follows closely the divertor power time variation. This engineering heat flux q_{div}^{eng} is directly comparable to the divertor allowable of 10 MW/m² at WEST, and it stays well below that limit. This process operated flawlessly during C11 campaign.

Other power indicators are calculated from power models. One noticeable indicator is the cumulated energy density to WEST vacuum vessel bellows, that heat-up as a result of the radiation. The bellow incident heat flux is derived from the separatrix radiated heat flux, using a form factor calculated from a geometric model. The maximum cumulated energy density on the bellows is obtained by combining the time-integrated the RT measured bolometry heat flux combined with the radiation form factor. The bellow cumulated energy stays well below the maximum allowable energy density of 2.1 MJ/m² during the long discharges. This functionality is currently operating as a proof of principle.

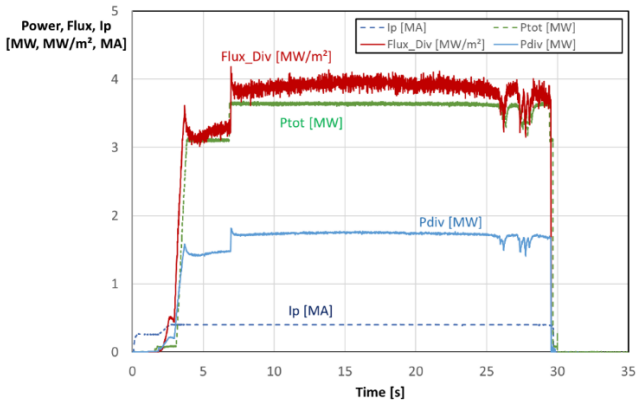

Fig. 3: Time traces of the RT calculated divertor heat flux and power from a RT power balance calculation, along with the measured plasma current.

2.2. Learning-based identification of thermal events

Introducing intelligence into the detectors acting on the 2D image stream of the IR viewing diagnostic shall give access to better detectors. The objective is to mimic the capability of human operators interpreting the IR movies and pointing at thermal events, primarily according to their shape, location, and their time evolution, rather than to the temperature value. This section addresses theses processes based on machine learning: the thermal event, the strike line, the LH arc, the REs and the UFOs detectors.

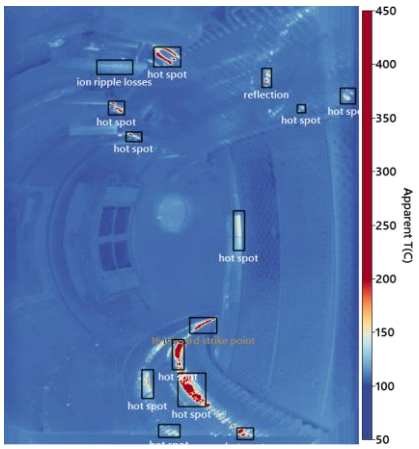

Fig. 4: Bounding boxes surrounding thermal events being detected during WEST discharge #60223.

The thermal event detector aims at detecting thermal events based on their shape or texture. The detector is built from a database of manually annotated events having occurred previously and available as a series of images. It is based on the concept of bounding boxes, a rectangle that surrounds the hot spot, being made available as metadata associated to the image. A time tracking module acting on the movies links the individual hot spots (relative to a single frame) to thermal events (which relates to a time sequence, characterised by starting and ending times). Several annotation supporting tools have been built on purpose to generate these annotations. These tools facilitate greatly the generation of annotated databases, which are the cornerstone of Artificial Intelligence (AI) developments. These developments of IR image-based advanced wall protection processes are done in an international framework, mainly within the EUFOfusion consortium. Thermal events formats (instances, categories, annotations) are developed jointly with IPP/W7-X [9,10]. The ability to share data and knowledge is viewed as crucial, because cross-machine diversity is key to obtain high-performance models. Most image data workflows used in both CEA/WEST and IPP/W7-X institutes use common libraries. The formats and frameworks are accessible to other possible users, through a policy of open-source software developments [11,12]. Data privacy and security are enabled by proprietary front-end layers. Further compatibility with ITER and the IMAS data storage architecture is pursued with EUROfusion [13], including the previously mentioned merging with synthetic data.

Various detection algorithms have been employed over time, namely mask-RCNN, then faster RCNN, and finally YOLO [14]. Bounding boxes of detected events are given in Fig. 4 for #60223. The “hot-spot” label is the catch-all class which fetches thermal events that could not be classified in a known category. These detectors operate in RT, and offer a smart alternative to rigid ROIs. After the discharge, the thermal events are stored in the WEST database. They are ranked according to severity index, which is basically the ratio of peak temperature to allowable temperature. The ranked thermal events are accessible through a web-based interface, and also by the THERMAVIP software suite [15].

The next tool is the strikeline characterisation tool. It aims at providing the strikeline curvature and angle within the image, for determining whether a strikeline is considered normal or pathologic. Curvature and angle are simple indicators that would be used unconsciously by the PPOs to decide whether a given strikeline is normal or not. Normal strikeline are merely processed for temperature and flux, while pathologic strikelines are further expertised for deeper understanding. The curvature is rated on an integer scale (from 0 to 3), and the curvature is given in degree, relative to the image main axis. The strikeline tool operates downstream the thermal event detector tool, on the bounding boxes having been labelled as “strikeline”. It is a constrained U-net [16], which both produces the strikeline skeletons, as illustrated by the connected yellow squares in Fig. 5 below, and also the two quantities of the angle and curvature. Both numbers are then compared to ranged standards, and the strikelines classified according on whether they are normal or not.

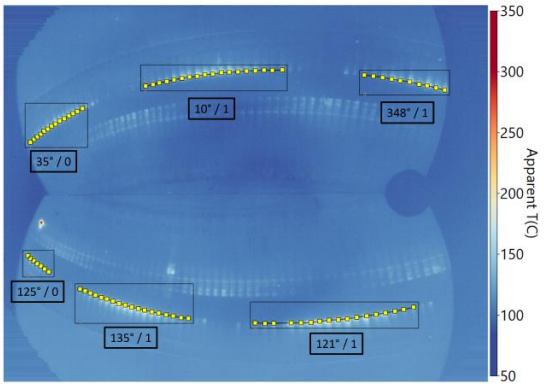

Fig. 5: Extraction of strikelines properties during #57396: segmentation of peak thermal regions, angle and curvature. (Angle in ° / Curvature index from 0 to 3).

Another AI-based detector is the LH arc detector. It detects electrical arcs on the FoV looking at the waveguides of the LH launcher. The IR images provide the most relevant signature of electrical arcs, better than the copper signal or the antenna reflection coefficient. The arc detector is a lightweight neural network acting on the image frame. It flags the images where an arc is detected, without even producing the bounding box (Fig. 6). The training is done on 47 films with arcs, and achieves a exactness of 87.3% and an F1-score of 85.4%. Despite its relatively long lag time of 20 ms (resulting from the image acquisition frequency), the image-based detector is viewed by the heating experts as the most relevant signature of the presence of an arc. The detector is operating in RT on the image stream. It communicates the possible presence of arcs to the PCS through the shared memory.

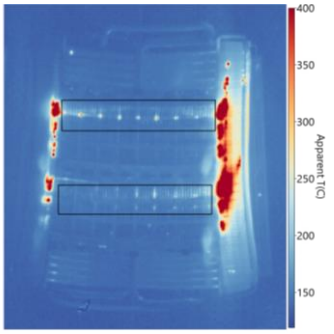

Fig. 6: Electrical arcs detected by a compact arc detector on the LH1 lower-hybrid heating antenna. The bounding boxes surrounding the arcs have been manually drawn, as the detector does not aim at obtaining bounding boxes.

The data processing framework being established, it becomes easier to build further specialized AI-based image detectors. The UFO detector (named UFOund) is developed as a result of the need to get an accounting of UFOs, following the 2022 high fluence campaign [17, 18]. UFOund is made of three convolutional blocks followed by an average pooling layer and a linear feed-forward model. The total number of trainable parameters of the UFOund model is 151442. This small size enables it to run on small GPUs in near RT during operation. The dataset is generated from 60 pulses of the C9 campaign, leading to the manual annotation of 295 infrared movies. The total number of annotated frames is 134712, among which 7603 contained UFOs (5.64% of the total number). The model has been implemented in Python using PyTorch, and trained and evaluated on a dedicated server equipped with three NVIDIA A30 GPUs (each with 24Gb of VRAM).

Following the same principle, a RE detector is developed toward building a image database of RE events. The RE detector is based on the synchrotron light emission of the RE beam which spectra includes some emission in the waveband of the IR detector (3.9 µm), hence being visible on the tangential FoV.

2.3. The post discharge thermal event diagnostic tool

Further ability to integrate the elaborate knowledge of the WEST thermal scenes is researched, in relation with the plasma session general parameters. This knowledge is virtually defined as the collective memory of the PPOs. It is too complex to be formulated as a list of rules. There are too many dependencies, qualitative judgements, and uncertainties. Large Language Models (LLM) emerge as a possible candidate technique to integrate such complex knowledge. Practically at WEST, LLM developments are based on the discharge summary. The summary is the IR viewing system dashboard. It is the first tool used by the PPOs for analysing the thermal scenes of a plasma experiment once finished. The PPO dashboard is a one-screen summary view of all Fields of View (FoV) of a given discharge. There are two sub-tables for each FoV: the “maxed” view of the FoV, which contain the maximum intensity of each pixel over the time of the discharge, and the time traces of the ROIs. The “maxed” view is an artificially reconstructed picture allowing to give a fast and synthetic representation of the most relevant hot spots having occurred in a FoV throughout the discharge. Two additional views are provided at the bottom of the dashboard, namely the discharge key parameters (plasma current, additional power for each heating system, plasma density) allowing to relate possible hot spots to notable events of the discharge, and a view with the expected feedback control from the IR wall monitoring system on the PCS. An LLM training database is created from these dashboards, associated to textual comments from the PPOs. The multimodal LLM Llava 1.6 [19] is used, which is an appropriate compromise in terms of performance and confidentiality, while still being an open-source model. The interfacing of the dashboard data to the LLM uses a dedicated transfer tool, since the LLM only works with suites of “tokens in input” (lexical or data units). An original projection module for tokenisation interface is trained, allowing to associate both IR images, temperature traces and textual content in a suite of tokens, while staying within the limit of the maximum number of tokens (4096). The model has 13 billions parameters. It is fine-tuned on CEA/CNRS national supercomputers. Once fine-tuned, the model is repatriated to WEST, and takes the name LLM4PPO [20]. The inference is calculated on WEST data servers a few minutes after the discharge. A typical result of LLM4PPO is provided in Fig. 7. The pulse number and FoV are selected on the left hand side table, and the analysis results is displayed on the right hand side after a few seconds in a textual formal. The sentences point at either a normal or abnormal thermal scene and sequence, providing further qualifications directly understandable by PPOs.

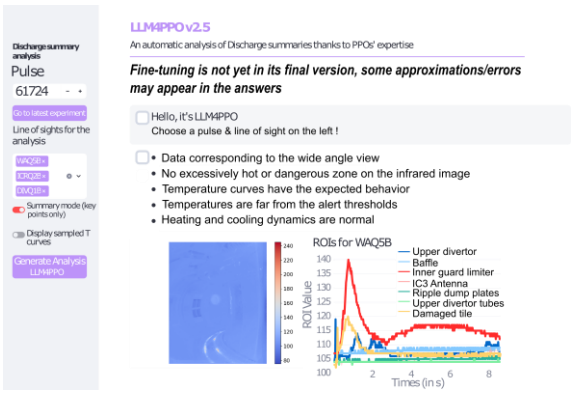

Fig 7: LLM4PPO result for #61724 (The original screen output has been edited with larger fonts, and the prompt has been translated from French to English).

3. RESULTS

3.1. Performance of machine learning based processes

Like workers health and safety monitoring, quantifying performance for matters of operation safety and machine protection is a broad field in itself. Poor performance is obvious when mishaps and accidents occur. By contrast, performance quantification attracts little attention when everything runs well. The performance is also tuneable with numerical parameters, according to whether the users favour catching all events, at the price of risking some false positives, or if the users prefer to avoid false positives, accepting the risk of missing some detections. Finally, performance quantification connects with data science performance evaluation, which is a deeply specialized topic [21]. The appropriate metrics depends on the usage or client:

— Loss functions are used for the training phase of the model. Loss function metrics are basically numerical distances, taken between two data points.
— Metrics such as precision, recall, and F1 score are used when it comes about comparing the performance of various possible candidate classification processes. There are known shortcomings to these indicators, for example for imbalanced classification processes. More elaborated indicators are possible (precision versus recall curve, receiver operating characteristic analysis). These tools are specialist's tools and address experts' communities.
— And finally, user metrics may be defined, in agreement with clients and users. For example, a user metric can be derived from a polling satisfaction test. User metrics are simple and easily explainable, but may be not as rigorous as the previous ones.

For the strikeline characterisation tool described in section 2.1, quantitative results obtained on the test set (synthetic images) demonstrate near perfect segmentation metrics (F1-score > 98%) and low regression errors (mean absolute error < 1° for angle and < 0.05 for curvature). The tool was then used satisfactorily by the users, although user defined metrics are yet to be developed.

UFOund achieves a balanced accuracy of 0.78 and a F1-score of 0.67 on an unseen test set with a detection threshold of 0.95, and gives very good qualitative results during operation at WEST.

The quantification of the LLM4PPO performance is even more complex to assess. There is no widely accepted "ground truth" to compare with. The quantification of the expert's system performance relies on evaluating four successive versions of LLM4PPO, combined with independent PPOs human notations from polls. The latest version of LLM4PPO (v2.5) has a 30% rate of incorrect responses, which is still a 20% improvement compared to the initial version v1.0. LLM4PPO is currently used as a "proof-of-principle" tool in the control room, with care. It is still being improved by training, i.e. by augmenting the training database using PPOs expert analyses.

3.2. System performance

During the experimental campaign C9 (Jan-Apr 2024), the availability of the wall protection is of 94.1%. The availability index reflects a state where the whole chain of wall protection is fully operational, which is an outstanding availability. The partial availability, defined as the state where one or several safety components do not operate normally, is even larger. Partial availability still allows plasma operation, while certain wall safety functions might not be active. 15.6 % of the discharges entered the active control region, meaning that automatic wall protection measures were activated. Only 1 discharge out of 1389 is lost as a result of the hard-wired wall protection system, because the automatic feedback control failed. This shows a >99.9% reliability of the wall protection system.

4. CONCLUSION

The WEST Wall monitoring system is developed by incorporating a-priori knowledge either from scientific models or phenomenological know-how recorded as annotated databases. This allows making advanced usage of dense information being hardly accessible in the extensive 2D+time infrared recordings of in-vessel IR viewing systems. The IR data is supplemented with signals coming for other diagnostic through data merging, giving access to flux quantities which contribute greatly to the safe operation of actively cooled components. The relevant thermal event and flux indicators are used both in real time through machine actuators, and also post discharge for expert's analyses toward adjusting the discharge sequence during experimental sessions. No critical wall

power event happened during C9 to C11 (2024-2025, about 13 hours of plasma cumulated). While it cannot be demonstrated that active wall protection enabled smooth operation, and especially the new plasma duration record of 1337 s, it is bound that the active wall protection as a whole helped significantly obtaining these records by preventing wall hot spots to become critical during the campaigns.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] BUCALOSSI, J., Overview of WEST contributions to the new ITER baseline and fusion power plants, this conference.
[2] MITTEAU, R., BELAFDIL, C., BALORIN, C., 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. Des. 165 (2021) 112223
[3] COURTOIS, X., AUMEUNIER, M.H., BALORIN, C. et al., Design and status of the new WEST IR thermography system, Fus. Eng. Des. 136 (2018) 1499-1504
[4] GASPAR, J., CORRE Y., RIGOLLET, F., et al., Overview of the emissivity measurements performed in WEST: in situ and post-mortem observations, Nucl. Fus vol. 62(2022)096023
[5] AUMEUNIER, M.-H., JUVEN, A., GERARDIN, J. et al., Surface temperature measurement from infrared synthetic diagnostic in preparation for ITER operations Nucl. Fusion Vol. 64 (2024) 086044
[6] CORRE, Y., CHANET, N., COTILLARD, R., et Al., First temperature database achieved with Fiber Bragg Grating sensors in uncooled plasma facing components of the WEST lower divertor, Fus. Eng. Des., Vol. 170(2021)112528
[7] GASPAR, J., GERARDIN, J. CORRE, Y., et al., Calorimetry measurement for energy balance and energy distribution in WEST for L-mode plasmas, Nucl. Fus.;Vol.64, Nb 3, (2024) 036018
[8] FEDORCZAK, N., GASPAR J., CORRE, Y., Cross diagnostics measurements of heat load profiles on the lower tungsten divertor of WEST in L-mode experiments, Nucl. Mat. Ener. 27 (2021) 100961
[9] JAKUBOWSKI, M, Infrared imaging systems for wall protection in the W7-X stellerator, Rev. Sci. Instrum. 2018 89 (10) 10E116
[10] PUIG SITJES, A.; JAKUBOWSKI, M.; NAUJOKS, D.; et al., Real-Time Detection of Overloads on the Plasma-Facing Components of Wendelstein 7-X. Appl. Sci. 2021, 11, 11969
[11] https://github.com/IRFM/thermavip
[12] GRELIER, E., MONCADA, V., MITTEAU R., An open source fusion machine agnostic standard for the exchange and processing of infrared videos and video annotations, Fus. Eng. Des. Vol. 215 (2025) 114934
[13] LITAUDON, X., Eurofusion contributions to ITER nuclear operation, Nucl. Fusion 64 (2024) 112006
[14] GRELIER, E, MITTEAU, R., MONCADA, V., Learning-Based Process for the Automatic Detection, Tracking, and Classification of Thermal Events on the In-Vessel Components of Fusion Reactors Fus. Eng. Des. 192 (2023) 113636
[15] MONCADA, V., COURTOIS, X., DUBUS, L., Software platform for imaging diagnostic exploitation applied to edge plasma physics and real-time PFC monitoring, Fus. Eng. Des. Vol. 190(2023)113528
[16] GORSE, V., GRELIER, E., MONCADA,V., et al., Real-time monitoring system for detection and characterization of thermal events on WEST Tokamak: Implementation and first results, Fus. Eng. Des. 215 (2025) 114960
[17] GASPAR, J., ANQUETIN, Y., CORRE, Y., et Al., Thermal and statistical analysis of the high-Z tungsten-based UFOs observed during the first deuterium high fluence campaign of the WEST tokamak, Nuc. Mat. Ener. Vol. 41 (2024) 101745
[18] GRELIER E. ; BONNAIL, J., COURTOIS, X., Automated UFO detection via infrared diagnostics in fusion reactors: Application to the WEST tokamak Fus. Eng. and Des., Vol. 221 (2025) 115401/UFOund/
[19] HAOTIAN L., et al., LLaVA-NeXT: Improved reasoning, OCR, and world knowledge. Jan. 2024.
[20] GORSE, V., Decision Support for In-Operation Monitoring of the WEST Tokamak First Wall Using Multimodal Large Language Model (LLM) on Infrared Imaging, submitted to Knowledge based systems
[21] JAPKOWICZ, N., Evaluating Learning Algorithms, A Classification Perspective, Cambridge University Press, August 2011