

CONFERENCE PRE-PRINT**DISRUPTION PREDICTION FOR FUTURE TOKAMAK REACTORS FROM DIFFERENT PERSPECTIVES AND WITH DIFFERENT METHODS**

¹Wei Zheng, ¹Xinkun Ai, ¹Fengming Xue, ¹Zhong Yu, ¹Chengshuo Shen, ²Dalong Chen, ²Bihao Guo, ¹Nengchao Wang, ¹Ming Zhang, ¹Yonghua Ding, ¹Zhongyong Chen, ²Bingjia Xiao, ²Biao Shen

¹ International Joint Research Laboratory of Magnetic Confinement Fusion and Plasma Physics, State Key Laboratory of Advanced Electromagnetic Technology , Huazhong University of Science and Technology

² Institute of Plasma Physics, HFIPS, Chinese Academy of Sciences, Hefei, China

Email: yhding@hust.edu.cn, zhengwei@hust.edu.cn

1. INTRODUCTION

Accurate disruption prediction is one of the keys for tokamaks to be commercially viable reactors. Machine learning disruption predictors exhibit good performance on the machines they are trained on. However, future reactors cannot provide enough disruption data before they damage themselves [1-3]. Current research has explored cross-tokamak disruption prediction by using data from multiple existing tokamaks, showing that the physical information contained in such data can reduce the model's dependency on data from a new device [4-6]. Nevertheless, this approach still requires accumulating a certain number of discharge shots on the new device for model training and cannot achieve disruption prediction starting from the very first discharge. Moreover, for future large scale experimental tokamak reactors, at different operation phases, they will have different goals, generate different data with different risks. The requirements on disruption prediction system are different. Studies have indicated that traditional machine learning models lack sufficient robustness against such variations [7]. Consequently, research has been conducted to enable disruption models to automatically adapt to changing operation scenarios. For example, one strategy involves periodically adding the latest discharge samples to the training set and retraining the model, allowing model to learn discharge characteristics in new operation scenarios [8,9]. Auxiliary strategies, such as discarding outdated samples from the training set, help the model focus more on recent data, thereby enhancing its adaptability to new operation scenarios. However, the adaptive capability of existing models still requires further improvement. To address these issues, we designed different methods and strategies tailored to different operation phases. In the initial phase, when no disruption data is available, we present an anomaly detection disruption predictor which can be deployed for the very first shot of a new machine. And based on the anomaly detection model, a cross-machine adaptive transfer method is proposed, enabling the model to rapidly adapt to the changing operation scenarios of new machine. As operation continues and a limited amount of data becomes available, we employ physics-guided transfer learning to adapt models trained on existing machines to the new machine, minimizing the demand for new-machine data while maintaining predictive performance. Additionally, we incorporate a disruption budget consumption (DBC) framework to quantify the impact of disruptions across different operational phases. This allows for risk-aware operation and model evaluation, ensuring that disruption prediction systems not only maximize accuracy but also minimize machine damage throughout the tokamak's lifecycle. Together, these methods form a comprehensive, adaptive disruption prediction strategy suitable for future large-scale tokamaks.

2. ADAPTIVE ANOMALY DETECTION DISRUPTION FROM SCRATCH

In future high-performance tokamaks such as ITER, the plasma energy storage is extremely high. The disruption could cause severe damage to the device, making it impossible to collect sufficient disruption data to train supervised learning models. To address the challenge of scarce disruption samples in training supervised models, the J-TEXT team has studied the anomaly detection-based disruption prediction method that requires only non-disruptive samples for training[10]. This approach eliminates the need for disruption data, overcoming the issues of data imbalance and labeling difficulties inherent in supervised methods. Four anomaly detection predictors are developed based on autoencoder (AE), one-class support vector machine (OCSVM), k-nearest neighbor (KNN), and angle-based outlier detection (ABOD). These models are tested on both J-TEXT [11] and EAST machines, achieving approximately 90% true positive rate (TPR) and below 10% false positive rate (FPR).

However, the scarcity of available data during the initial operation phase of future new device, especially with no data available at the first discharge. This makes it difficult to train and fine-tune models, thereby hindering

effective disruption prediction early on. Traditional anomaly detection models cannot use disruption precursor samples during training, resulting in significant wastage of data. To reduce the dependency on new machine data during cross-tokamak transfer, the J-TEXT team designed an Enhanced Convolutional Autoencoder Anomaly Detection (E-CAAD) model (shown in Figure 1)[12]. Equation 1-1 presents the loss function of the E-CAAD model, which can minimize the RCE-1 for disruption precursor samples during training. Therefore, this model can be trained using only non-disruptive samples, but also uses disruption precursor samples when they become available. This approach not only overcomes the class imbalance problem in supervised models, but also improves data efficiency compared to traditional anomaly detection methods, making it more suitable for data-scarce scenarios in early-stage operation of new machines.

$$loss = \frac{1}{n} \sum_{i=1}^n [y \cdot RCE_i^y + \alpha \cdot \beta \cdot (1 - y) \cdot RCE_i^{y-1}], \quad \beta = \begin{cases} \frac{n_N}{n_D}, & n_D \geq 1 \\ 0, & n_D = 0 \end{cases} \quad (1)$$

Where n_N and n_D are the numbers of non-disruption and disruption precursor samples, respectively, in the training set. β is the dynamic balance parameter of positive and negative samples, automatically calculated by the model before each retraining step. α is the model's attention parameter to the class of disruption precursor samples, set as a hyperparameter before model deployment and experimentally tuned on existing machines to search for the optimal hyperparameter.




Figure 1 Network structure of convolution autoencoder

Moreover, differences in machine structure, operational parameters, and control systems among different tokamaks make it difficult to directly transfer a high-performance data-driven disruption prediction model trained on one machine to another. In future high-parameter tokamak machines, the damage caused by disruptions will be extremely severe, necessitating disruption prediction from the first discharge after the main systems are operational. Moreover, activities such as wall conditioning, diagnostic upgrades, and operation range exploration on new machines will lead to changes in plasma operation scenarios, requiring the prediction model to possess adaptive abilities to continuously track these variations. To address these challenges, the J-TEXT team proposed a comprehensive strategy based on the E-CAAD predictor[13]. This strategy integrates three key approaches: the existing-machine data assistance strategy, the adaptive learning from scratch strategy, and the adaptive alarm threshold adjustment strategy. Together, these enable disruption prediction from the first discharge on a new machine and allow the predictor to automatically adapt to changes in plasma operation scenarios.

Existing-machine data assistance strategy: Using J-TEXT as the existing machine and EAST as a simulated future device, the J-TEXT team conducted cross-machine transfer experiments with the E-CAAD model. The results demonstrate that the E-CAAD model trained on the existing machine can effectively distinguish between disruption precursors and non-disruption samples on the new device, confirming the feasibility of disruption prediction from the first discharge on a new device.

Adaptive learning from scratch strategy: Building on the existing-machine data assistance strategy, the team further introduced an adaptive learning strategy and an adaptive alarm threshold adjustment strategy based on the E-CAAD model. The adaptive learning strategy involves periodically adding the most recent samples to the training set and retraining the model, enabling it to learn features of new operation scenarios. This approach allows the model to make full use of scarce early-phase data. And the weight function incorporated into the model's loss

function gradually reduces the influence of outdated samples and emphasizes learning from the latest discharges, facilitating quick adaptation to changing operation scenarios.

Adaptive alarm threshold adjustment strategy: This strategy generates a reference threshold based on the reconstruction error (RCE) characteristics of historical discharges. It then dynamically scales this threshold according to the real-time estimated RCE distribution under new operation scenarios, addressing the challenge of threshold selection during data-scarce periods.

Using the above strategies, cross-machine transfer experiments of the E-CAAD model from J-TEXT to EAST are conducted. The results, shown in Figure 2, demonstrate successful disruption prediction from the first discharge on the new machine while maintaining the model’s adaptability to dynamic changes in operation scenarios. With a 20 ms response time reserved for the Disruption Mitigation System (DMS), the transferred model achieved a TPR of 85.88% and an FPR of 6.15%. This performance is comparable to that of models trained with large amounts of EAST historical data, proving the effectiveness of the proposed cross-machine transfer method.




Figure 2 Adaptive adjustment of disruption alarm threshold during E-CAAD predictor's adaptive learning process from scratch on EAST

3. PHYSICS GUIDED DEEP TRANSFERABLE DISRUPTION PREDICTOR

Though the previously introduced method can provide basic protection to the machine, a higher false alarm rate than acceptable limits its performance at the beginning of the experiments. Thus we introduce a physics guided deep transferable disruption predictor, which is trained with past experiment data with little data from the new machine[14]. We focus on finding a balance in the exploration-exploitation dilemma. On the one hand, physics-based models exploit decades of accumulated experimental knowledge, offering reliable and interpretable rules such as density limits or safety factor limits. Yet, this knowledge is incomplete and cannot capture all device-specific disruption pathways. On the other hand, data-driven approaches can explore hidden correlations and fine-grained features, but they require sufficient labelled discharges from the target machine, which is an unrealistic condition for newly built or early-phase machines. The key question is therefore how to balance the exploitation of known physics with the exploration enabled by machine learning, under different data availability regimes.

To address this, we propose the Physics-guided Domain Adaptation Network (PhyDANet), which integrates physics both explicitly and implicitly into the transfer learning process, and combines data exploration with domain adaptation techniques. Explicitly, we embed known operational limits and instability proxies into the loss function as hard constraints, ensuring the model’s predictions remain physically meaningful even when trained with very limited data. Implicitly, we introduce a feature reconstruction task that guides the representation learning towards disruption-relevant physical precursors, such as tearing modes or radiation profile anomalies, without restricting the model’s ability to discover additional predictive patterns, so that the model generalizes better by learning known disruption precursor physics patterns that have been proven on multiple tokamaks. As for domain adaptation, both Domain Adversarial Neural Network (DANN) and supervised deep Maximum Mean Discrepancy (MMD) are used, for marginal and conditional domain shifts alignment. A diagram of the PhyDANet is illustrated in Figure 3.

The model is validated in J-TEXT as the existing, source tokamak and EAST as the newly-built, target tokamak. Two data regimes are set up to simulate the realistic deployment stages, an extremely scarce regime (with one

disruptive and one non-disruptive discharge from EAST) and a low-resource regime (with ten disruptive and ten non-disruptive discharges from EAST). Numerical experiments are conducted and results are summarized in Table 1.

Table 1 Performance of PhyDANet under different data regimes and physics-guided strategies.

	Data	+ Physical Loss	+Reconstruction	No physics guidance
Only EAST	2 shots	0.8357	0.8039	0.7992
	20 shots	0.8902	0.9053	0.8756
J-TEXT+EAST	J-TEXT 1021 shots	0.8026	0.7762	0.8094
	EAST 2 shots			
Domain Generalization	J-TEXT 1021 shots	0.9449	0.9475	0.9212
	EAST 20 shots			

In the extremely scarce regime (only two EAST discharges), direct training without physics guidance quickly overfits to noise, achieving only marginal predictive capability. By contrast, introducing physical loss constraints provides a substantial boost in performance, anchoring predictions to known physical limits such as density and safety factor thresholds. Reconstruction is less effective here, as the lack of data limits its ability to guide feature learning. This confirms that when exploration through target data is not possible, exploitation of well-established physics remains the most reliable path. However, combining with J-TEXT data through domain generalization introduces mixed effects. On one hand, the additional data enriches the representation space and provides exposure to more disruption patterns. On the other hand, the domain gap between J-TEXT and EAST can mislead the model when alignment is insufficient, leading to negative transfer. Physical loss helps mitigate this risk by steering the model away from spurious alignments, but performance gains remain modest compared to physics-only guidance. This highlights the challenge of leveraging external data under extreme scarcity, where robust physics priors remain indispensable.

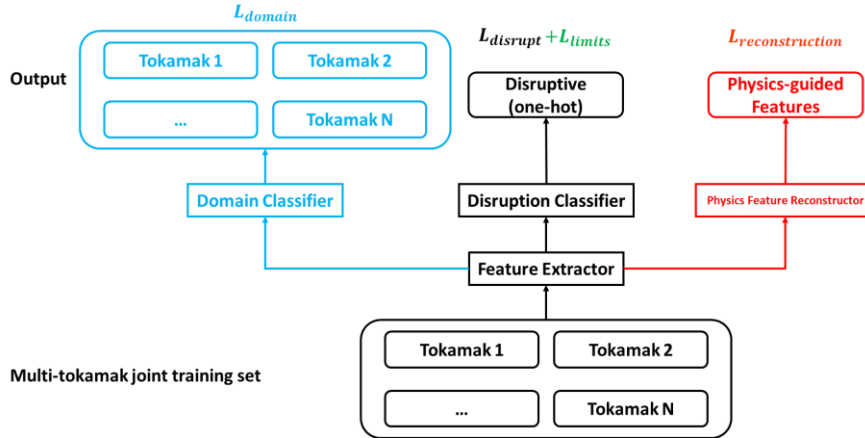


Figure 3 Diagram of the Physics-guided Domain Adaptation Network (PhyDANet)

In the low-resource regime (twenty EAST discharges), the balance begins to shift. With more target data available, domain adaptation becomes more effective: aligning marginal and conditional distributions across J-TEXT and EAST allows the model to explore cross-machine knowledge while adapting to target-specific characteristics. Here, reconstruction also begins to play a stronger role, as the increased data variation supports learning detailed, disruption-related representations. When combined, representation alignment and reconstruction yield the best overall performance, demonstrating that as exploration through target data becomes feasible, fine-grained feature exploitation overtakes physics-only constraints as the main driver of predictive gains.

The results illustrate a dynamic balance: in the absence of sufficient target data, physics priors dominate and anchor the model to general physical mechanisms; as more target data become available, domain adaptation and representation learning exploit cross-machine knowledge; and with abundant data, the predictor shifts towards device-specific exploration of fine-grained features. This layered integration of physics and data makes the predictor both more transferable and more robust, offering a pathway for reliable disruption protection from the first discharges of future reactors.

4. DISRUPTION BUDGET GUIDED OPERATION

Plasma disruption has a significant impact on the life cycle of tokamak machines. Therefore, taking plasma disruption as a management objective is an important part of the operation and experimental research of tokamaks.

To address the issue of plasma disruption prediction across different operation phase on the tokamak, first of all, we will explore the operating range according to the ITER operation plan[15]. Then the core lies in evaluating the disruption in different operating phase, studying the impact on the tokamak, and providing the disruption budget consumption (DBC) values; Next, we will conduct cross-phase disruption prediction research, incorporate the concept of DBC, construct different datasets, and evaluate the performance of the model in different shots from different operation phase. However, the performance in the high-parameter operation phase is more important. Meanwhile, the total DBC of the dataset also represents the impact of disruption on the machine that the training model needs to bear. Finally, we combined DBC with the model prediction results and proposed a new performance evaluation method for disruption prediction. Compared with traditional evaluation methods, we can evaluate the model from the perspective of the impact of disruption on the machine or the life cycle of the tokamak device. The concept of DBC brings a new perspective to rupture prediction.

DBC is used to quantify the impact of disruptions in different operational phases of a tokamak plasma discharge. Since each operational phase has different plasma current (magnetic energy) and thermal energy, the severity of a disruption is not the same across phases. DBC measures how much of the “disruption budget” is consumed in each operational phase. It’s calculated by plasma current and toroidal field. For disruption shots, the thermal energy and CQ rate should be taken into consideration.

Then plasma discharges are divided into five classes (class1–class5) based on plasma current ranges. Class 1 and 2 (low-current phases), all available shots are used in the training set. For class 3 and class 4 (medium-current phases), 80% of the shots are used for training and 20% for testing. A unique DBC-based filtering is applied at this stage: shots with excessively high DBC values are removed from the training set. The rationale is that high-DBC shots often represent extreme cases that can reduce the generalization ability of the model if overrepresented in training. For class 5, only a limited number of shots (10) are included in the training set, while the majority are reserved for testing. This ensures that the evaluation phase reflects the model’s ability to predict disruptions under the most demanding plasma conditions. The DBC value is defined as formula (1). Potential DBC is planned budget consumption for dataset. For shots disrupted in training set and miss classified in test set, the disruption DBC will be used. For shots undisrupted, classified disruption in test set and all class 1 and class2 shots, the DBC is 0.

$$DBC = \begin{cases} \frac{Ip - Ip_{min}}{Ip_{std}} + \frac{Bt - Bt_{min}}{Bt_{std}} \text{ potential DBC;} \\ \frac{Ip - Ip_{min}}{Ip_{std}} + \frac{Bt - Bt_{min}}{Bt_{std}} + \frac{P_{tot} - P_{tot_{min}}}{P_{tot_{std}}} + \frac{CQ_{rate} - CQ_{rate_{min}}}{CQ_{rate_{std}}} \text{ potential DBC;} \\ 0 \text{ undisruption;} \end{cases} \quad (2)$$




Figure 4 Cost curve for model performance on test 3, 4 and 5; model trained without 10 shots from class 5 (blue line) and with 10 shots from class 5 (yellow line).

In the Figure 4, the model performance is presented. The cost curve integrates the probability of disruption in a given class (operational phase). In Figure 4, the probability of positive is the probability of disruption. Different operational phases have different disruption rates, and therefore the misclassification costs vary accordingly. To enable fair comparison, the cost is normalized across datasets with different disruption rates. This allows the model’s performance to be directly compared under varying plasma conditions. The normalized cost is the DBC value used by model in test set. For example, model trained by adding limit class 5 shots have 4 false negative in the test 5, then the DBC value of these 4 shots will be used and normalized by the whole DBC value of test 5. The normalized cost on test 3,4 and 5 are plotted in the cost space along with their probability of positive (probability

of disruption). They are connected as the cost curve in cost space. By plotting the normalized cost against the disruption probability, the generalization ability of the model across operational ranges can be assessed. A lower cost indicates that the model not only predicts disruptions accurately but also balances the trade-off between false alarms and missed disruptions in phases with different risks. There are 2 models trained tested on test 3,4 and 5. The first one is trained with shots in class 2, 3 and 4. The other one uses the same shots and adding 10 shots from class 5 with low DBC value. Including some class 5 disruption data in the training set improves the generalization of the model across operational phases, reducing the overall prediction cost and enhancing performance in high-risk plasma regimes. The cost curve does not only show that the model with class 5 data have better performance on test 5 and test 3, but also demonstrate how many percent of DBC it can save comparing with the model without class 5 data.

Compared with the model generalization on 3 test set, the model performance on test 5—measured by both AUC and DBC cost—is more critical. A DBC-based filtering strategy is applied: shots with excessively high DBC values are removed from the training set. The DBC-based filtering is applied on class 3 and class 4 shots in training set. With this strategy, the DBC of train set can be modified depending on removed shots belonging to which operation phase. By removing the excessively high DBC shots in class 4 and class 3 separately, evident results are shown in Table 2. The shots removed from class 3 are less than shots removed from class 4. But the DBC strategy applied on class 4 yields better performance in terms of both AUC and normalized DBC on the test set. This means with less training set DBC, the DBC strategy can perform better on test 5. It can be found that when 15 shots removed from class 3, model has higher AUC than model removing 20 shots from class 4, but the DBC cost on test set is higher too. So, DBC cost on test set is as important as the AUC for disruption prediction.

Table 2 DBC strategy on class 3 and class 4, model performance with AUC, normalize DBC and DBC on train set.

Train set	AUC	Normalized DBC(test)	DBC(train)
c1,c2,c3; c5*10 ;c4-20	0.94	0.023	1414.58
c1,c2,c3; c5*10 ;c4-35	0.97	0.035	1196.4
c1,c2; c5*10 ;c4,c3-15	0.95	0.025	1583
c1,c2; c5*10 ;c4,c3-25	0.91	0.043	1502.1

Table 3 DBC strategy on class 3 and class 4, with similar train set DBC, more test set DBC and lower AUC on class3.

Train set	AUC	Normalized DBC(test)	DBC(train)
c1,c2,c3; c5*10 ;c4-20	0.94	0.023	1414.58
c1,c2; c5*10 ;c4,c3-30	0.92	0.051	1469.3

Some results are presented in Table 3, when similar train set DBC are used, the model performance differs. The strategy that removes 20 shots from class 4 achieves both a higher AUC and a lower DBC cost on the test set compared to the strategy applied to class 3. The DBC strategy highlights that more training data does not always yield better performance or less test set DBC cost. Instead, carefully selecting training sets based on DBC and operation phase can significantly improve prediction in higher operational phases.

5. CONCLUSION

This work has systematically addressed the critical challenge of disruption prediction for future tokamak reactors by proposing a phased and multi-faceted strategy tailored to the distinct operational phases and data-availability scenarios of next-generation machines like ITER. Unlike conventional one-model-fits-all approaches, our framework is built on the premise that disruption prediction requirements evolve throughout a tokamak's lifecycle, necessitating adaptive, context-aware, and risk-informed solutions.

In the initial operational phase, where no disruption data exists, the E-CAAD model, combined with a cross-machine adaptive transfer method, provides a foundational safety net. This transfer method integrates three key approaches: the existing-machine data assistance strategy, the adaptive learning from scratch strategy, and the adaptive alarm threshold adjustment strategy. Together, these enable disruption prediction from the first discharge on a new machine and allow the model to automatically adapt to changes in plasma operation scenarios.. Experimental validation on EAST has demonstrated its ability to achieve a high TPR (85.88%) with a manageable FPR.15%), proving effective even during the data-scarce initial phase of a new device.

As operations progress and a limited amount of device-specific data becomes available, PhyDANet offers a sophisticated solution to the exploration-exploitation dilemma. By explicitly embedding physical constraints into the loss function and implicitly guiding feature learning towards known disruption precursors, PhyDANet effectively balances the reliability of established physics with the adaptability of data-driven discovery. The results show that physical guidance is crucial under extreme data scarcity, while domain adaptation and feature

reconstruction become increasingly powerful as more target machine data is accumulated, ensuring robust performance across different data regimes.

Furthermore, the introduction of the DBC framework provides an essential risk-aware perspective, quantifying the impact of disruptions on machine lifetime across different operation phases. This approach allows for more meaningful, operationally-grounded model evaluation and training set curation, ensuring that prediction models are optimized not just for accuracy, but for minimizing cumulative machine damage—particularly in high-risk, high-current operational regimes.

Collectively, these methods form a cohesive and scalable disruption prediction ecosystem. The anomaly detection method ensures safety from the first discharge; physics-guided transfer learning enables knowledge leverage and adaptation with minimal new data; and the DBC framework aligns technical performance with long-term engineering constraints. This layered strategy provides a structured pathway for the gradual and safe integration of machine learning into the operation of future fusion reactors.

While the methods presented show significant promise in tests conducted on existing machines like J-TEXT and EAST, their ultimate validation will come from deployment on future large-scale tokamaks. Future work will focus on the real-time implementation of these integrated strategies, their seamless coupling with disruption mitigation systems, and extension to a broader range of machine classes and operational scenarios. This research lays a necessary methodological foundation for the safe, efficient, and economically viable operation of future fusion energy systems.

ACKNOWLEDGEMENTS

The authors would like to acknowledge helpful input from the J-TEXT and the EAST team. This work was supported by National Key R&D Program of China under Grant (No. 2024YFE03240200, and No. 2022YFE03040004), National Natural Science Foundation of China (NSFC) under Project Numbers Grant (No. T2422009 and No. 12375219), Natural Science Foundation of Wuhan under Project Numbers Grant (No. 2024040701010040), and Interdisciplinary Research Program of HUST (No. 2024JCYJ017), Hubei International Science and Technology Cooperation Projects (No. 2022EHB003).

REFERENCES

- [1] J.X. Zhu, C. Rea, K. Montes, et al. Hybrid deep-learning architecture for general disruption prediction across multiple tokamaks, *Nucl. Fusion*, 61 (2020) 026007.
- [2] A. Loarte, Required R&D in Existing Fusion Facilities to Support the ITER Research Plan, in, 2020.
- [3] W. Zheng, F.M. Xue, C.S. Shen, et al., Overview of machine learning applications in fusion plasma experiments on J-TEXT tokamak, *Plasma Sci. Technol.*, 24 (2022) 12.
- [4] C.S. Shen, W. Zheng, B.H. Guo, Y.H. Ding, D.L. Chen, X.K. Ai, et al. Cross-tokamak disruption prediction based on domain adaptation[J]. *Nuclear Fusion*, 2024, 64(6): 066036.
- [5] J.X. Zhu, C. Rea, K. Montes, R.S. Granetz, R. Sweeney, R.A. Tinguely. Hybrid deep-learning architecture for general disruption prediction across multiple tokamaks[J]. *Nuclear Fusion*, 2021, 61(2): 026007.
- [6] J. Zhu, C. Rea, R.S. Granetz, E.S. Marmar, K.J. Montes, R. Sweeney, et al. Scenario adaptive disruption prediction study for next generation burning-plasma tokamaks[J]. *Nuclear Fusion*, 2021, 61(11): 114005.
- [7] A. Murari, R. Rossi, E. Peluso, M. Lungaroni, P. Gaudio, M. Gelfusa, et al. On the transfer of adaptive predictors between different devices for both mitigation and prevention of disruptions[J]. *Nuclear Fusion*, 2020, 60(5): 056003.
- [8] A. Murari, M. Lungaroni, E. Peluso, P. Gaudio, J. Vega, S. Dormido-Canto, et al. Adaptive predictors based on probabilistic SVM for real time disruption mitigation on JET[J]. *Nuclear Fusion*, 2018, 58(5): 056002.
- [9] A. Murari, M. Lungaroni, M. Gelfusa, E. Peluso, J. Vega. Adaptive learning for disruption prediction in non-stationary conditions[J]. *Nuclear Fusion*, 2019, 59(8): 086037.
- [10] X.K. Ai, W. Zheng, M. Zhang, et al., Tokamak plasma disruption precursor onset time study based on semi-supervised anomaly detection, *Nuclear Engineering and Technology*, (2023).
- [11] Y. Ding et al., "Overview of the recent experimental research on the J-TEXT tokamak," *Nuclear Fusion*, , 64 (2024) 112005.
- [12] X.K. Ai, W. Zheng, M. Zhang, et al., Cross-tokamak deployment study of plasma disruption predictors based on convolutional autoencoder, *Plasma Physics and Controlled Fusion*, (2024).

- [13] X.K. Ai, W. Zheng, M. Zhang, et al., Adaptive anomaly detection disruption prediction starting from first discharge on tokamak, Nucl. Fusion, 65 (2025) 036011.
- [14] W. Zheng, F. Xue, Z. Chen, et al., Disruption prediction for future tokamaks using parameter-based transfer learning, Communications Physics, 6 (2023) 181.
- [15] Y. Zhong, W. Zheng, Z.Y. Chen, et al., High-beta disruption prediction study on HL-2A with instance-based transfer learning, Nucl. Fusion, 64 (2024) 096012.