9th Asia-Pacific Conference on Plasma Physics, 21-26 Sep, 2025 at Fukuoka

FusionMAE: large-scale pretrained model to optimize and simplify diagnostic and control of fusion plasma

Zongyu Yang[1], Zhenghao Yang[1&2], Wenjing Tian[1&2], Jiyuan Li[1], Xiang Sun[1], Guohui Zheng[1&3], Songfen Liu[3], Niannian Wu[1&4], Rongpeng Li[4], Zhaohe Xu[1], Bo Li[1], Zhongbing Shi[1], Zhe Gao[2], Wei Chen[1], Xiaoquan Ji[1], Min Xu[1] & Wulyu Zhong[1]

[1] Southwestern Institute of Physics, Chengdu, China, [2] Tsinghua University, Beijing, China,
[3] Nankai University, Tianjin, China, [4] Zhejiang University, Hangzhou, China

e-mail (speaker): zy-yang@swip.ac.cn

Accurate and robust fusion of heterogeneous diagnostic signals is essential for real-time control and scientific understanding in magnetic confinement fusion (MCF) devices. However, the presence of incomplete measurements, sensor dropouts, and secondary-derived quantities pose significant challenges for traditional methods. In this work, we propose FusionMAE, a novel masked autoencoder architecture tailored for the fusion of multichannel time-series diagnostics in tokamak plasmas. By leveraging a Transformer-based encoder–decoder design inspired by masked auto-encoder in computer vision, FusionMAE is capable of reconstructing missing diagnostic signals and extracting a compact, high-fidelity embedding of plasma states.

Our dataset comprises over 2,400 discharges from the HL-3 tokamak, spanning four experimental campaigns. Each discharge contains 88 synchronized diagnostic channels, resampled to 1 kHz, covering magnetic, kinetic, radiative, and shape-related measurements. We implement a structured masking strategy, where 25% of diagnostic channels are randomly masked during training, in addition to inherent invalid signals. The encoder processes unmasked channels through per-channel MLP projection and a stack of multi-head self-attention layers. The compressed latent representation is then decoded to reconstruct the full set of signals.

FusionMAE demonstrates strong denoising and imputation capabilities, achieving over 96% Pearson correlation in reconstructing masked channels. Notably, it also recovers secondarily-derived diagnostics such as plasma boundary parameters (R, Z, a, κ), total beta, and loop voltage, without explicit physics constraints. The learned latent embedding proves to be semantically aligned with key plasma parameters and remains stable even under 20% random channel dropout.

In downstream tasks, we integrate FusionMAE outputs into a suite of representative models: a deep disruption predictor (DPR), an equilibrium reconstruction network (EFIT-NN), and a next-step plasma evolution predictor (PPR). In all three cases, models trained on FusionMAE-reconstructed or embedded features match or outperform those trained on full raw diagnostics, especially in scenarios with noisy or partially missing input. Additionally, UMAP visualization of embeddings reveals natural clustering aligned with plasma current, magnetic field strength, and core density, suggesting physical interpretability of the latent space.

FusionMAE thus serves as a general-purpose, diagnosis-agnostic framework that unifies data fusion, noise suppression, and downstream task acceleration. Its capacity to handle corrupted, incomplete, or derived signals makes it particularly suited for long-pulse and reactor-scale operation scenarios, where diagnostic redundancy and reliability vary dynamically. Future work will focus on real-time deployment and cross-device generalization to facilities.

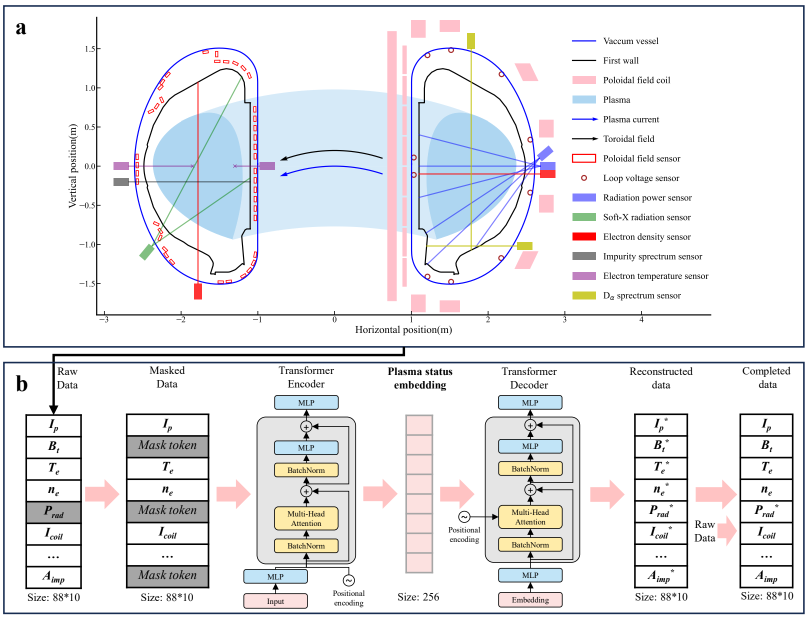

Fig.1: Overall architecture of FusionMAE.(a) Diagram of the diagnostic systems employed in this study.(b) Workflow and neural network structure of FusionMAE, which compresses data from multiple diagnostic systems into a unified plasma status embedding.