

# DATA-DRIVEN PLASMA EQUILIBRIUM FORECASTING IN MAGNETIC FUSION TOKAMAK

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The most promising approach to achieving nuclear fusion is through tokamaks, which confine plasma using magnetic fields. Understanding the current plasma equilibrium state in tokamaks is critical for effective plasma control. Unlike previous studies, which reconstruct equilibrium from magnetic field information, our work forecasts future equilibrium based on past equilibrium states. Specifically, we formulate the plasma equilibrium prediction task as a video prediction task, a well-explored problem in the machine learning community. This formulation allows us to capture the spatio-temporal dynamics of plasma states and provides a foundation for multimodal modeling of data streams from tokamak operations. Our methodology, incorporating a physics-inspired learning technique for physically reliable predictions, achieved plausible results in forecasting future plasma equilibrium up to 200 ms ahead compared to baselines. This approach holds promise for predicting plasma instabilities and preventing disruptions, marking a significant step towards developing stable fusion reactors.

## 1 INTRODUCTION

One way to categorize human history is by examining the main energy sources utilized in each era. From the discovery of fire to the exploitation of fossil fuels, and more recently, to the adoption of nuclear energy and solar power, humanity has continually evolved its primary energy sources. To address climate change and global energy shortages facing humanity, and to enter a new era, developing new clean energy sources has become an urgent mission. Nuclear fusion technology emerges as the most promising candidate for next-generation clean energy and has been studied for decades (Sheffield, 1994; Harms et al., 2000; Freidberg, 2007). It offers the potential for nearly limitless, clean energy production by replicating the processes that power the sun.

Currently, the most viable approach for achieving nuclear fusion is through tokamak devices such as the International Thermonuclear Experimental Reactor (ITER), the largest international collaborative scientific project in the world (Rebut, 1995). Tokamaks use powerful electromagnets to confine high-temperature plasma within a vacuum vessel to induce nuclear fusion reactions. However, numerous scientific and engineering challenges hinder the commercialization of nuclear fusion. In detail, one of the primary obstacles is plasma instability and the resulting plasma disruptions (Schuller, 1995). To predict and prevent these issues, it is essential to understand plasma equilibrium, a state where the magnetic field’s confining force on plasma particles balances the force of the particles attempting to diffuse outward. While plasma equilibrium information can be computed by analyzing magnetic field data using an algorithm called equilibrium fitting (EFIT) (Lao et al., 1985), its application is limited due to the difficulty of performing real-time calculations.

To overcome these challenges, the nuclear fusion community has recently begun to actively introduce machine learning (ML) and data-driven approaches (Seo et al., 2021; Degraeve et al., 2022; Lao et al., 2022; Char et al., 2023; Joung et al., 2023; Seo et al., 2024; Kim et al., 2024). ML techniques are being applied to tasks such as controlling tokamaks, predicting disruptions, and inferring plasma equilibrium. Along with the aforementioned studies, this work aims to address challenges in the nuclear fusion domain using data-driven methods, specifically by forecasting plasma equilibrium in tokamaks. Unlike the previous study that predicted plasma equilibrium using magnetic field data, we propose predicting future plasma equilibrium based on past plasma equilibrium data. This ap-

proach enables us to address the temporal dynamics of plasma, introducing a novel problem setting in fusion research that goes beyond existing studies focusing on current-state reconstruction.

More precisely, we formulated the problem as a video prediction task by interpreting plasma equilibrium data as a sequence of temporal snapshots. We applied widely known video prediction models to plasma data and introduced methods to incorporate plasma physics knowledge into the model training process. As a result, our model was able to predict plasma equilibrium up to approximately 200 ms into the future without explicit tokamak control information. This is considered a sufficiently long-term prediction, as plasma physics typically deals with phenomena occurring on timescales of a few to tens of milliseconds inside a tokamak.

Our research introduces a novel data-driven approach to nuclear fusion, an under-explored field within the ML community. By tackling complex spatio-temporal plasma equilibrium data rather than the scalar features used in previous studies (Wai et al., 2022; Degrave et al., 2022; Wan et al., 2023; Char et al., 2023), we demonstrate the potential of ML in fusion science. In this manner, our work can serve as a foundation for developing multimodal models in fusion research, paving the way for the simultaneous analysis of diverse fusion data types. Our findings, based on data from Korea Superconducting Tokamak Advanced Research (KSTAR) (Lee et al., 2001), are applicable to ITER, given the structural similarities between these tokamaks. Finally, current research efforts are focused on developing a digital twin of KSTAR and conducting virtual fusion experiments (Kwon et al., 2022). Our study has the potential to significantly advance the ultimate goal of comprehensive virtual fusion experimentation.

The contributions of our work can be summarized as follows:

- We formulated the plasma equilibrium forecasting problem as a video prediction task, introducing this under-explored problem to both nuclear fusion and ML communities.
- We introduced a physics-based loss function inspired by plasma equilibrium equations and demonstrated its effectiveness.
- We evaluated our proposed model against widely-used baseline methods across various domains and interpreted the results from the plasma physics perspective.

## 2 PRELIMINARY

This section provides a brief introduction to the background knowledge necessary to understand our work. For a deeper understanding of nuclear fusion or how ML can be applied in the field of fusion, readers may refer to Hutchinson (2002), Ikeda (2007), Freidberg (2007), Wesson & Campbell (2011), McCracken & Stott (2013), and Spangher et al. (2024).

### 2.1 NUCLEAR FUSION AND TOKAMAK

Nuclear fusion is a physical process in which two light atomic nuclei combine into a single heavier nucleus under high temperature and pressure, simultaneously releasing energy. This process results in plasma, known as the fourth state of matter, where electrons are separated from atomic nuclei and can move freely. Einstein’s famous equation  $E = mc^2$  can be used to calculate the amount of energy produced during nuclear fusion, showing that even a small amount of fuel can generate an enormous amount of energy. For example, to produce the same energy as 125 kg of deuterium and tritium fuel, approximately 2.7 million tons of coal are needed <sup>1</sup>.

In fact, nuclear fusion reactions occur in the sun, and its brightness is due to these reactions. While nuclear fusion in the sun occurs at approximately 15 million Kelvin due to immense gravitational pressure, achieving fusion on Earth requires temperatures of 100–150 million Kelvin because we cannot replicate the sun’s gravitational conditions. However, building a vessel or material that can directly contact plasma at such high temperatures is practically infeasible. Therefore, researchers have devised a method using magnetic fields to confine the plasma in a vacuum state, preventing it from contacting the vessel walls. This concept is implemented in devices known as tokamaks, which are donut-shaped machines. The tokamak consists of various components, namely multiple electromagnets for confining and controlling plasma, heating systems such as neutral beam injection

<sup>1</sup><https://www.iter.org/sci/FusionFuels>

(NBI) and electron cyclotron resonance (ECR) heating systems, and multiple types of sensors for monitoring the tokamak’s operational status.

Representative tokamak devices, including KSTAR (Lee et al., 2001) in South Korea, DIII-D (Luxon, 2002) in the United States, EAST (the EAST et al., 2009) in China, and JET (Jacquinot et al., 1999) in Europe, are contributing crucial insights into plasma physics as well as the development and operation of ITER (Rebut, 1995). Among these, KSTAR and EAST, which leverage superconducting electromagnets, can operate for significantly longer periods compared to other tokamaks that do not use superconductors. This key difference provides a distinct advantage, enabling long-term plasma experiments and detailed investigations. Our work utilizes real experimental data from KSTAR; consequently, our findings have potential applications in the prediction and analysis of long-duration tokamak operations.

## 2.2 PLASMA EQUILIBRIUM

The most crucial information for controlling and analyzing plasma inside a tokamak is plasma equilibrium. Plasma equilibrium refers to the balance between the magnetic forces confining the plasma and the outward pressure of the plasma itself. The Grad-Shafranov (GS) equation (Grad & Rubin, 1958; Shafranov, 1966), a partial differential equation describing this equilibrium state, illustrates the spatial distribution of magnetic fields under conditions satisfying plasma equilibrium:

$$\Delta^* \psi = -\mu_0 R J_\phi = -\mu_0 R^2 \frac{dp(\psi)}{d\psi} - F(\psi) \frac{dF(\psi)}{d\psi}. \quad (1)$$

In this equation,  $\psi = \psi(R, Z)$  represents the poloidal magnetic flux function dependent on coordinates  $R$  and  $Z$ , defined within the cylindrical coordinate system of KSTAR. The term  $\mu_0$  denotes the permeability of free space, while  $p(\psi)$  represents the plasma pressure and  $F(\psi)$  is related to the toroidal magnetic field, both being functions of  $\psi$ .  $\Delta^* \psi$  is a two-dimensional non-linear partial derivative of  $\psi$ , defined as  $\Delta^* = R \frac{\partial}{\partial R} \frac{1}{R} \frac{\partial}{\partial R} + \frac{\partial^2}{\partial Z^2}$ .

Based on the  $R$ ,  $Z$ , and  $J_\phi$ , which is the toroidal component of the plasma current density, we can obtain the solution  $\psi$  (and  $\Delta^* \psi$ ) of the GS equation. This information serves as an essential key to understanding the state of the plasma for controlling and analyzing the tokamak. However, it is challenging to determine the solution of the GS equation, and the EFIT algorithm (Lao et al., 1985) was developed to effectively find the solution. In brief, EFIT employs an iterative calculation process to find the solution based on given initial conditions. Therefore, it has limitations in scenarios requiring real-time application such as tokamak control due to its heavy computational time. Consequently, the real-time EFIT (rt-EFIT) algorithm (Ferron et al., 1998; Moret et al., 2015; Huang et al., 2017) has been proposed to reduce computational time by using the solutions at the previous time step as initial values, but this approach is also limited due to lower accuracy.

The reconstruction of plasma equilibrium through EFIT calculations provides abundant information for both controlling and understanding plasma behavior. Primarily, the key output  $\psi$  is utilized to determine the last closed flux surface (LCFS). The LCFS, which delineates the boundary of the confined plasma, is a crucial parameter in tokamak operations. By accurately identifying the LCFS, researchers can enhance plasma stability, optimize tokamak operational methods, and gain deeper insights into plasma dynamics. Furthermore, the results from the EFIT calculations can be combined with various experimental data collected from tokamak diagnostics. This allows for the derivation of important plasma parameters, including temperature profiles, density distributions, and magnetic pressure, which are crucial for optimizing plasma stability, confinement efficiency, and fusion reaction rates.

## 2.3 DATA-DRIVEN APPROACHES IN FUSION DOMAIN

Due to the computational cost of the EFIT algorithm, data-driven approaches have been recently proposed as alternatives to predict plasma equilibrium faster and with comparable accuracy to the EFIT algorithm (Joung et al., 2019; Wai et al., 2022; Lao et al., 2022; Joung et al., 2023; Lu et al., 2023). GS-DeepNet serves as a notable example (Joung et al., 2023). This model predicts the current plasma equilibrium in real time with the current magnetic field information given as input. Notably,

162 GS-DeepNet introduced an unsupervised learning mechanism without requiring ground truth EFIT  
 163 data. Our method, in contrast, predicts the current or future equilibrium based on past equilibrium  
 164 information **without magnetic field data, and can potentially collaborate with the aforementioned**  
 165 **methods by bootstrapping from the equilibria they reconstruct.**

166 Beyond the plasma equilibrium prediction, data-driven techniques are increasingly being applied  
 167 to various challenges in the field of nuclear fusion. For example, Seo et al. (2021), Degraeve et al.  
 168 (2022), and Char et al. (2023) demonstrated the potential of ML in tokamak control; moreover, Seo  
 169 et al. (2024) and Kim et al. (2024) showed that ML-based control can help avoid plasma instabilities,  
 170 which are a major challenge in tokamak operation. While most studies have utilized a number of  
 171 simple and raw features from the tokamak as inputs, our method distinguishes itself by leveraging  
 172 spatio-temporal information, such as  $\psi$  and  $\Delta^*\psi$  from the EFIT calculation. This approach allows  
 173 us to capture complex plasma dynamics and their evolution over time, providing a framework that  
 174 can be extended to support diverse types of data with complex structures generated from tokamaks.  
 175 Our study, along with existing research, establishes a foundation for the development of multimodal  
 176 models capable of integrating and understanding the enormous volume of data from tokamaks.

### 177 178 179 3 METHOD

180  
181 In this section, we define our problem setup and introduce our proposed approach. Our key idea is  
 182 to interpret the plasma equilibrium prediction task as a video prediction task in the ML domain. We  
 183 present the advantages of this approach, as well as propose a technique inspired by plasma physics.  
 184

#### 185 186 3.1 FORMULATING PLASMA EQUILIBRIUM PREDICTION AS VIDEO PREDICTION

187  
188 Video prediction is a task of predicting future frames based on the given past frames. Formally, given  
 189  $T$  past frames  $X_{t-T:t-1} = \{x_i\}_{t-T}^{t-1}$  at time  $t$ , the objective of this task is to obtain  $T'$  future frames  
 190  $Y_{t:t+T'-1} = \{x_i\}_t^{t+T'-1}$ , where  $x_i$  represents an image of dimensions  $C \times H \times W$ , with  $C$  channels  
 191 and spatial resolution of  $H \times W$ . The EFIT data from fusion experiments have similar features to  
 192 video data, allowing us to approach fusion plasma prediction as a video prediction problem. In  
 193 this context, each EFIT frame  $x_i$  can be considered as a two-dimensional image where channels  
 194 correspond to  $\psi$  and  $\Delta^*\psi$ .  
 195

196 This video-based approach offers several advantages over the traditional pixel-level approach, which  
 197 treated  $\psi$  and  $\Delta^*\psi$  values at each point as separate data. By simultaneously utilizing  $\psi$  and  $\Delta^*\psi$   
 198 information from multiple points, we can expect more accurate predictions. For instance, this com-  
 199 prehensive approach enables the model to identify complex spatial patterns and relationships that  
 200 traditional pixel-level methods could not capture. Furthermore, this framework has the potential to  
 201 be extended to a multimodal model in the future by employing additional spatio-temporal data col-  
 202 lected from tokamaks. Moreover, by processing spatial and temporal information concurrently, this  
 203 method offers an opportunity to analyze plasma dynamics from a spatio-temporal perspective. This  
 204 is particularly significant because the aforementioned GS equation, which is currently the primary  
 205 method for reconstructing plasma equilibrium, describes only spatial changes but not temporal evo-  
 206 lution. If the model based on video prediction approach provides us with significant results regarding  
 207 plasma’s temporal dynamics, it could lead to novel scientific insights into plasma behavior.

208 **This intersection of physics and ML is promising (Chen et al., 2022), as ML researchers’ expertise**  
 209 **in video prediction can offer fresh perspectives unconstrained by traditional physical knowledge. To**  
 210 **this end, we explore the application of SimVP (Gao et al., 2022), a widely adopted and versatile**  
 211 **video prediction model architecture. SimVP processes input frames by employing spatial down-**  
 212 **sampling, transforming spatio-temporal features through inception modules, and decoding them to**  
 213 **generate predictions at the original resolution.** Despite the existence of various video prediction  
 214 models, we selected this one for its straightforward and versatile nature in video prediction tasks.  
 215 Specifically, SimVP does not make assumptions about the data type, making it an ideal candidate  
 for the plasma equilibrium prediction task. This flexibility is crucial for our future plans to develop  
 multimodal models capable of integrating diverse data from tokamak operations.

### 3.2 GRAD-SHAFRANOV CONSTRAINT LOSS

The  $\psi$  and  $\Delta^*\psi$  data used as input and output in our study are derived from the EFIT algorithm, which is based on the GS equation (Equation 1 in Section 2.2). Essentially, these data inherently satisfy the GS equation, implying that any accurate prediction from our model should also follow this fundamental constraint. To ensure this, we introduce a GS-constraint loss to train the model in addition to the mean absolute error (MAE) on  $\psi$  and  $\Delta^*\psi$ . The GS-constraint loss helps to maintain the physical consistency of our predictions within the theoretical foundation of plasma equilibrium.

In detail, we apply  $\Delta^*$  operator, defined as  $\Delta^* = R \frac{\partial}{\partial R} \frac{1}{R} \frac{\partial}{\partial R} + \frac{\partial^2}{\partial Z^2}$ , to the model output  $\psi_{pred}$  to calculate  $\Delta^*\psi_{est}$ . Subsequently, we compare it with the model output  $\Delta^*\psi_{pred}$  to compute the MAE, which serves as our GS-constraint loss as follows:

$$\mathcal{L}_{GS} = \mathbb{E} \|\Delta^*\psi_{est} - \Delta^*\psi_{pred}\|_1 = \mathbb{E} \|D_{\Delta^*}(\psi_{pred}) - \Delta^*\psi_{pred}\|_1. \quad (2)$$

Here,  $D_{\Delta^*}(\psi_{pred})$  represents the application of the  $\Delta^*$  operator to  $\psi_{pred}$  using finite difference approximations. Since  $\psi$  and  $\Delta^*\psi$  are defined on a fixed grid over  $R$  and  $Z$ , and since the coordinates  $R$  and  $Z$  are fixed in our data representation, we cannot use automatic differentiation to compute spatial derivatives in this case. Alternatively, following Ren et al. (2022), we approximated the  $\Delta^*$  operator by employing a finite difference method, shifting  $\psi$  horizontally and vertically. Although this approximation technique is numerically less accurate than actual differentiation, it presents the most feasible solution for our problem. The approximation error (reported in Table 1 using  $\dagger$  symbol) in the GS-constraint loss term could lead to inaccurate predictions. Therefore, we scaled the GS-constraint loss by a coefficient less than one and added it to the pixel MAE loss in order to mitigate numerical errors:  $\mathcal{L}_{Total} = \mathcal{L}_{MAE} + \alpha \mathcal{L}_{GS}$ , where  $\alpha = 0.01$ .

Meanwhile, one might propose using the leftmost and rightmost terms of Equation 1 for constraining the model predictions. This approach could allow for the natural calculation of derivative values through automatic differentiation. However, this approach faces practical limitations due to the ambiguity in the GS equation, which does not provide information about the exact forms of  $p(\psi)$  and  $F(\psi)$ . Specifically, although we understand the physical meaning of  $p$  and  $F$ , we cannot precisely determine their functional forms or characteristics as functions of  $\psi$ . In practice, the spatial distribution of current ( $FF'$  term) in the GS equation is typically approximated using low-order functions, such as first to third-order polynomials. Consequently, we opted for the previously described loss based on finite differences in a conservative manner. This conservative choice was made to ensure both numerical stability during training and physical fidelity of the model, as the alternative method could potentially lead to ambiguous physical interpretations and unstable training results. While this limitation could be addressed if additional conditions or data for  $p$  and  $F$  were available, such considerations remain a topic for future research <sup>2</sup>.

## 4 EXPERIMENTS

This section presents the experiments and results that demonstrate the effectiveness of our approach. It begins by introducing the dataset used for our experiments, along with the model training details. We then introduce the evaluation metrics, which are adaptations of commonly used metrics in ML problems, tailored for the plasma equilibrium prediction context. Finally, we conduct a comprehensive analysis through both quantitative and qualitative assessments of the experimental results.

### 4.1 KSTAR DATASET

We evaluated our method using experimental data from KSTAR. Our dataset comprises 789 KSTAR discharges collected from the 2017 and 2018 campaigns, split into training and test sets with an 8:2 ratio. A single data instance includes offline EFIT calculation results, plasma current ( $I_p$ ) information, and heating information (NBI and ECR). The EFIT data, which are the primary data used in our study, consist of a two-channel video ( $\psi$  and  $\Delta^*\psi$ ) for each shot, calculated at 50 ms intervals.

<sup>2</sup>In our problem setup,  $p(\psi)$  and  $F(\psi)$  can be implemented by stacking  $1 \times 1$  convolutional layers.




Figure 1: Visualization of example KSTAR EFIT data used in this work. The left plot shows  $\psi$  and  $\Delta^*\psi$ . The horizontal and vertical axes represent the  $R$  and  $Z$  coordinates of KSTAR, respectively, measured in meters. The original  $\psi$  and  $\Delta^*\psi$  data are displayed as heatmaps, while the researchers usually present them as contour lines. The LCFS, highlighted in lime, indicates the current confinement state and configuration of the plasma. The right plot illustrates plasma current information along with heating information. The plasma current, shown in red on the left y-axis, represents the overall operational status of the shot. Heating information is displayed on the right y-axis, with a blue line for NBI and a green line for ECR. The gray-shaded region represents the time interval for which the EFIT data are depicted in the upper plots.

Figure 1 illustrates the EFIT data from KSTAR discharge shot #18948 at 2.60 seconds. Each frame consists of a two-channel heatmap representing  $\psi$  and  $\Delta^*\psi$ , which satisfy the GS equation. The original EFIT data was calculated at a  $65 \times 65$  resolution, which is specifically calibrated resolution for KSTAR operations, but we cropped one pixel from the right and bottom, resulting in a  $64 \times 64$  resolution for convenience. In addition, we applied min-max normalization to  $\psi$  and  $\Delta^*\psi$  based on their statistics. The preprocessed EFIT data are visualized using contour lines, as shown in Figure 1. The most important feature for fusion researchers is the LCFS, drawn in a lime-colored line, which represents the plasma boundary. The LCFS is one of the most essential elements in plasma control and analysis, requiring precise estimation. To determine the LCFS, we generated 80 contour lines and selected the outermost closed line as the LCFS.

In addition to the EFIT data, we employed plasma current and heating information for each shot. Although plasma current information was not used in training, it was utilized in the analysis process because it provides insight into the overall operational status of the shot. The heating information contains both NBI and ECR data. For analysis purposes, we used the raw heating data, but downsampled it to match the EFIT sampling rate during training. This plasma current and heating information shows the operational status of the tokamak for each shot. For instance, examining the right plot in Figure 1, which illustrates discharge # 18948, we can observe the overall operational phases: a ramp-up stage until around 2 seconds, followed by a flat-top, and a ramp-down starting near 13 seconds. Furthermore, this experiment did not utilize ECR heating, but employed three NBI devices for plasma heating. Preprocessing of the NBI and ECR data was minimal due to their less significant numerical deviations. We only replaced negative values in the NBI and ECR with zeros, and there was no further normalization. When using the heating information in training, we replicated the scalar values of heating information to match the shape of  $\psi$  and  $\Delta^*\psi$ . Then, these replicated values were concatenated with  $\psi$  and  $\Delta^*\psi$  to form a four-channel input data.

## 4.2 BASELINE AND EVALUATION METRIC

**Evaluation metric** The basic evaluation was conducted using the pixel-level MAE, a common metric in video prediction tasks. In our work, we computed MAE for specific regions of interest in the  $\psi$  and  $\Delta^*\psi$ , rather than considering the entire region. The first area is the central region. This area was prioritized because the prediction accuracy in the outer regions, where plasma does not exist, is relatively less important. Second, we assessed the divertor region. The divertor is a crucial component of the tokamak, responsible for extracting high-temperature exhaust and residues from the plasma. Accurate prediction of plasma boundary in the divertor is essential for sustaining long-period operations and preventing damage to the tokamak. In detail, the plasma boundary al-

324 lows us to predict and evaluate the accuracy of key features such as the X-point (the intersection of  
 325 magnetic flux surfaces) and the divertor strike points where plasma and residues make contact with  
 326 the divertor. However, predicting these features accurately requires centimeter to millimeter-scale  
 327 precision (Eldon et al., 2020), which presents a challenge for the spatial resolution of our current  
 328 EFIT data. In this context, we extended our evaluation to assess the accuracy of the LCFS deter-  
 329 mined from the predicted  $\psi$ . As a potential alternative evaluation metric, we experimented with  
 330 intersection over union (IoU) and mean IoU (mIoU) to compare the similarity between the ground  
 331 truth and predicted LCFS shapes. The detailed analysis of these metrics and their limitations are  
 332 discussed in Section B.

333 **Baseline** Since our proposed method differs from existing research in both approach and results,  
 334 it is challenging in making direct and fair comparisons. To address this, we selected several base-  
 335 line methods that are widely adopted across different domains. First, we employed a time-lagged  
 336 baseline, widely recognized as a simple yet robust benchmark in time series forecasting problems.  
 337 This baseline outputs ground truth values from our EFIT dataset, delayed by a single time step.  
 338 In detail, we applied a 50 ms time lag, which represents a significantly long period in the context  
 339 of plasma dynamics. For example, plasma dynamics is typically analyzed with a time resolution  
 340 of less than 10 ms, and tokamak control systems require a response time on the order of a few  
 341 ms. We also presented a constant prediction baseline that assumes that the most recently observed  
 342 plasma state persists unchanged. Next, we conducted experiments using ConvLSTM (Shi et al.,  
 343 2015), PredRNN (Wang et al., 2017), and PredRNN-v2 (Wang et al., 2023), which are fundamen-  
 344 tal video prediction models. Finally, we compared our results with dynamic mode decomposition  
 345 (DMD) (Schmid, 2010), a widely-used linear model for predicting spatio-temporal dynamics, and  
 346 Chronos (Ansari et al., 2024), a general-purpose pretrained time series forecasting model.

### 347 4.3 TRAINING AND IMPLEMENTATION DETAILS

348 In our experiments, we used the implementation of model architectures from the OpenSTL li-  
 349 brary (Tan et al., 2022; 2023), while developing training and evaluation code from scratch to fit  
 350 our problem setup. For hyperparameters, SimVP was configured with  $N_S = 4$  and  $N_T = 8$  layers,  
 351 with hidden dimensions of  $d_S = 64$  and  $d_T = 256$ , respectively. The ConvLSTM model consisted  
 352 of four layers, each with a hidden dimension of 96. PredRNN and PredRNN-v2 were also config-  
 353 ured with four layers, but each utilized a hidden dimension of 128. To simplify our experiments, we  
 354 set the input and output frame lengths to be equal, in other words,  $T = T'$ . These sequences contain  
 355 one, four, or eight frames, representing time spans of 50 ms, 200 ms, and 400 ms, respectively.

356 The models shared the following training hyperparameters: We utilized the AdamW opti-  
 357 mizer (Loshchilov & Hutter, 2019) with a learning rate of  $10^{-4}$ , weight decay of 0.05, and gradient  
 358 clipping of 1. The dataset was sampled every four timesteps, and the batch size was set to 128.  
 359 The number of training epochs was 1, 500, and we applied an exponential moving average update  
 360 with  $\beta = 0.9999$ . To ensure statistical robustness, all experiments were repeated with three random  
 361 seeds, and we reported the average values. Each experiment was conducted on a single NVIDIA  
 362 H100 GPU and completed within 48 hours.

363 Predictions with DMD were generated autoregressively from the last observation after fitting it with  
 364  $T$  past frames. Since DMD inherently operates on 1D vectors, we tested two approaches: (1) flat-  
 365 tening the 2D EFIT data into 1D vectors and (2) performing predictions on a per-pixel basis. For  
 366 the flattening approach, we considered both joint prediction of  $\psi$  and  $\Delta^*\psi$  (i.e., predicting a vector  
 367 of dimension  $64 \times 64 \times 2 = 8,192$ ) and separate prediction (i.e., predicting two vectors, each of di-  
 368 mension  $64 \times 64 = 4,096$ ). All DMD experiments were conducted using the PyDMD library (Demo  
 369 et al., 2018). As for the Chronos model, we employed Chronos-t5-base, which has a com-  
 370 parable number of parameters to SimVP, and assessed its zero-shot accuracy without finetuning.  
 371 To enhance reliability, the final predictions of the Chronos model were obtained by averaging the  
 372 outcomes of 15 repeated samplings for each data point.

### 373 4.4 RESULTS AND ANALYSIS

374 We evaluated the efficacy of the GS-constraint loss and conducted a comprehensive analysis of the  
 375 prediction accuracy under various scenarios, varying both the types of input data and the prediction  
 376 length. Table 1 summarizes our experimental results. Among the baselines presented in the first and  
 377

Table 1: Qualitative evaluation of plasma equilibrium prediction.

Method	$\psi$ MAE ( $\downarrow$ )		$\Delta^*\psi$ MAE ( $\downarrow$ )		GS
	Center	Divertor	Center	Divertor	const. ( $\downarrow$ )
Baseline (50 ms delay)	0.00165	0.00161	0.01493	0.00742	0.04679 <sup>†</sup>
Baseline (Constant)	0.00676	0.00687	0.02946	0.01533	0.04679 <sup>†</sup>
DMD (200 ms)	0.93483	0.89436	14.4225	9.37053	5.09277
DMD (200 ms) $\psi$ only	0.98264	1.26488	-	-	-
DMD (200 ms) pixel-level	0.01053	0.01028	0.03441	0.01965	0.03163
Chronos-base (200 ms)	0.01758	0.01756	0.06306	0.04728	47.0440
ConvLSTM (200 ms)	0.00385	0.00463	0.04160	0.02701	0.34480
PredRNN (200 ms)	0.00287	0.00355	0.05799	0.03089	1.54382
PredRNNv2 (200 ms)	0.00603	0.00720	0.06159	0.03474	2.01905
SimVP (400 ms)	0.00175	0.00178	0.02434	0.01365	0.38020
SimVP (400 ms) + $\mathcal{L}_{GS}$	0.00165	0.00172	0.02396	0.01347	0.21892
SimVP (200 ms)	0.00128	0.00124	0.01825	0.00996	0.36242
SimVP (200 ms) + $\mathcal{L}_{GS}$	0.00123	0.00125	0.02035	0.01108	0.21476
SimVP (200 ms) $\psi$ only	0.00122	0.00119	-	-	-
SimVP (50 ms)	0.00116	0.00109	0.01850	0.01007	0.47197
SimVP (50 ms) + $\mathcal{L}_{GS}$	0.00106	0.00103	0.01945	0.01050	0.22314
SimVP (200 ms) + $\mathcal{L}_{GS}$ + Heat.	0.00181	0.00194	0.03291	0.01823	0.23188
SimVP (200 ms) + $\mathcal{L}_{GS}$ + Rev.	0.00132	0.00136	0.02289	0.01229	0.22324
SimVP (50 ms) + $\mathcal{L}_{GS}$ + Heat.	0.00172	0.00176	0.02330	0.01296	0.22133

second groups of the table, the time-lagged baseline demonstrated the lowest MAEs. This result is primarily due to the averaged metrics being reported across all timesteps. Notably, as illustrated in Figure 4 in Section A, the 50 ms ahead prediction accuracy of the pixel-level DMD model is comparable to that of SimVP. Further details can be found in A.

**GS-constraint loss** The third group in Table 1 demonstrates the changes in MAE when the GS-constraint loss was introduced. The models using the GS-constraint loss generally showed lower MAEs on  $\psi$  and the predicted  $\psi$  and  $\Delta^*\psi$  more closely satisfied the GS equation. However, we observed performance degradation in  $\Delta^*\psi$  prediction in the case of the 50 ms and 200 ms models. This trade-off is acceptable considering that  $\psi$  prediction is relatively more important than  $\Delta^*\psi$  prediction in the nuclear fusion domain. We also conducted experiments predicting only  $\psi$ , which resulted in improved prediction accuracy. Nevertheless, it should be noted that while the accuracy of  $\psi$  predictions may increase, this does not necessarily guarantee their physical validity; thus, further investigation is required to determine whether using only  $\psi$  is practically acceptable.

**Input configuration and Prediction length** Prior to our experiments, we hypothesized that incorporating heating information, which partially represents control signals of the tokamak, would enhance the prediction accuracy of the models. However, contrary to our expectations, Table 1 and Figure 2 (a) demonstrate that the inclusion of heating data did not bring improvement, or led to a decrease in accuracy in some cases. These unexpected results suggest that developing a multi-modal model capable of integrating multiple data types from a tokamak requires more than simply adding extra input channels. It indicates a need for architectural improvements in the model design to properly leverage these additional data sources.

Alternatively, we can hypothesize that the current heating information may be insufficient for accurate forecasting of  $\psi$  and  $\Delta^*\psi$ . In fact, heating in a tokamak is not limited to external devices such as ECR and NBI. There is also Ohmic heating, where the plasma is heated by its own electrical resistance. If Ohmic heating contributes more significantly to the future plasma equilibrium than the ECR and NBI, this could explain the current performance degradation when including heating data.

Subsequently, we examined the impact of prediction length on the models. In typical time series forecasting tasks, shorter prediction lengths generally yield lower error rates. As evidenced in Fig-



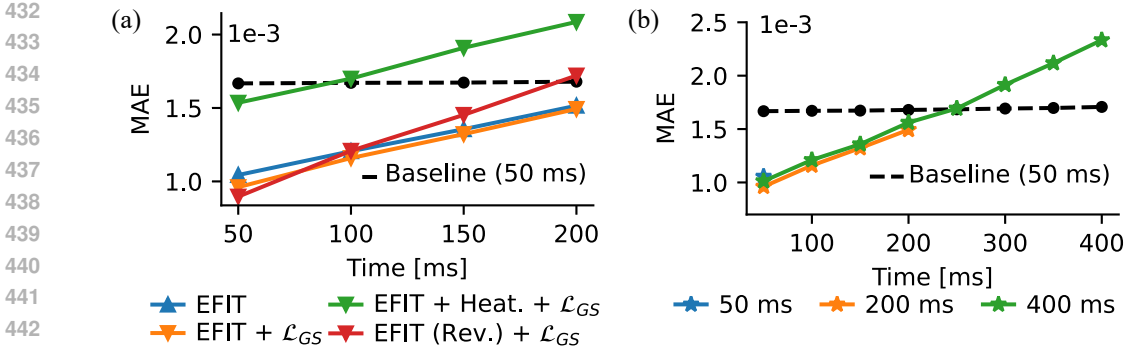


Figure 2: MAEs of  $\psi$  predictions across various model and input configurations. These two plots show MAE values for  $\psi$  predictions in the center region at different future time points. (a) All models, except the one using heating information, more accurately predict  $\psi$  up to 200 ms into the future compared to the 50 ms lagged baseline. (b) Models with input and output lengths of 200 ms and 400 ms exhibit higher prediction accuracy at 50 ms compared to the model with 50 ms input and output lengths.

Figure 2 (b), our results align with this trend, showing that models predicting 200 ms into the future achieve lower MAEs compared to those predicting 400 ms. Interestingly, we observed that the model predicting only 50 ms ahead demonstrated higher errors than the other two models, contradicting our initial intuition.

This unexpected finding raises the possibility that plasma states from the more distant past (beyond 50 ms) might have a potential influence on future plasma equilibrium, and this hypothesis requires further rigorous investigation to confirm. If substantiated through future study, this observation might offer a new perspective on plasma dynamics from a temporal viewpoint. Such a perspective could potentially complement the spatial description provided by the GS equation, which illustrates plasma equilibrium in space. However, it is crucial to emphasize that these are preliminary interpretations. Extensive experimental and theoretical investigations would be necessary to validate these hypotheses for understanding plasma dynamics.

In addition, we recognize that our current problem setting may not always require long-term equilibrium predictions. In fact, for many practical applications, short-term predictions are often sufficient and more relevant. To enhance prediction accuracy for the nearest future (50 ms ahead), we experimented with a counterintuitive technique of reversing the input sequence. The rationale behind this approach stems from the architecture of SimVP, where convolutional layers update the temporal information of input features. The features of the earliest input frame contribute most significantly to calculating the features of the nearest future frame in the output. By reversing the input sequence, we enable the most recent historical information to be utilized in computing the nearest future information. Figure 2 (a) shows that the model with reversed input performs best at 50 ms, validating the effectiveness of this heuristic idea. This technique could potentially be further refined by applying loss only to the nearest future frames.

**Qualitative analysis** Lastly, we analyzed the actual prediction results by visualizing the model outputs. Figure 3 shows the prediction results for discharge shot #20279 from 9.35 s to 9.50 s. This corresponds to the ramp-down phase of the experiment. The ramp-down phase involves reducing energy after maintaining plasma in a stable confinement state (*i.e.*, the flat-top phase). In other words, it is the process of terminating the fusion reaction and experiment. We focused our analysis on this phase because, during ramp-down, confinement gradually weakens, leading to plasma instabilities and an increased risk of disruptions, requiring careful control.

The results show that the predicted plasma boundary closely matches the ground truth. However, prediction accuracy decreases near the X-point region, suggesting that further refinement is necessary to achieve the level of precision required for plasma shaping and divertor control. Despite this limitation, it is noteworthy that we achieved this level of prediction using only past equilibrium states. This suggests that by enhancing forecasting accuracy through integrating these data and architectural improvements, the model could be practically used in various scenarios.

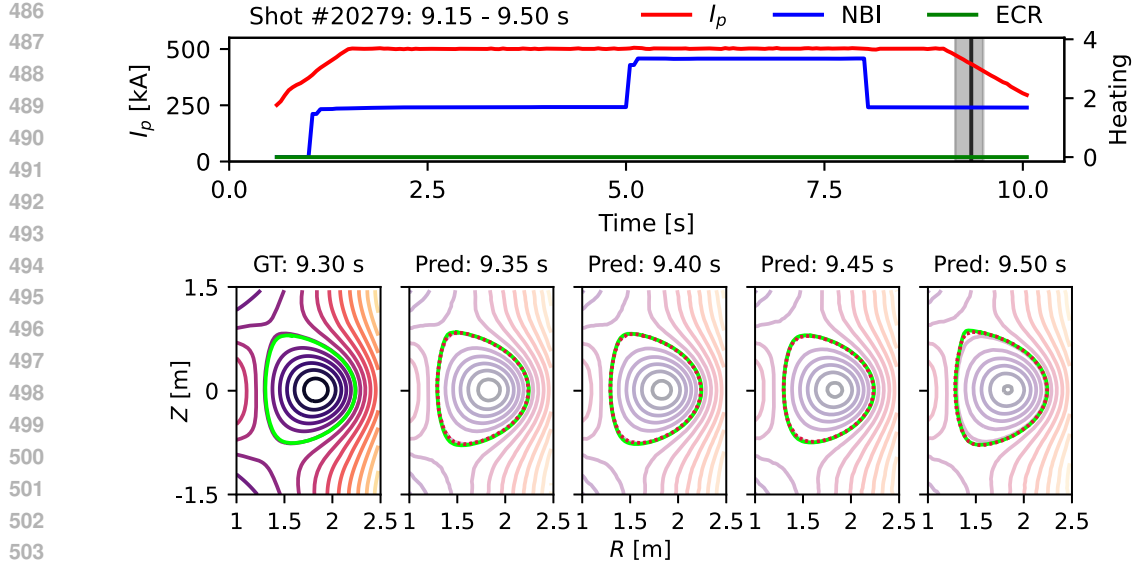


Figure 3: Visualization of model predictions at Shot #20279. The model uses  $\psi$  and  $\Delta^*\psi$  values from 9.15 s to 9.30 s as input to predict the corresponding values from 9.35 s to 9.50 s. The top plot illustrates the overall operational status of the shot, with input and output intervals shaded in gray and separated by a black vertical line. In the bottom row, the leftmost plot displays the last frame of the input  $\psi$ , while the subsequent plots show the predicted results in chronological order. In each  $\psi$  plot, the ground truth LCFS is depicted by a lime-colored line, and the LCFS determined from the model output is shown as a red dotted line.

## 5 DISCUSSION AND CONCLUSION

In this paper, we introduced a novel problem of forecasting future plasma equilibrium in tokamaks based on past states. We formulated this as a video prediction problem and conducted experiments across various methods and scenarios. Our results showed that the models can predict the poloidal magnetic flux ( $\psi$ ) up to 200 ms into the future more accurately than the baselines. In addition, we found that introducing a loss function based on the GS equation improved the predictions of  $\psi$  and resulted in physically reliable forecasting. However, we observed an unexpected accuracy degradation when integrating partial tokamak control signals, *i.e.*, NBI and ECR heating data, suggesting the need for further investigation to effectively leverage additional data from the tokamak. We conclude by discussing the potential applications of our results in the fusion domain and future directions.

**Outlook** In recent years, the nuclear fusion community has shown increasing interest in developing digital twins of tokamaks (Kwon et al., 2022; Iglesias et al., 2017). Beyond simple 3D models, these digital twins serve as tools for virtual fusion experiments that integrate various experimental data and simulation results. This approach offers two significant advantages: it reduces the time and energy costs required for each discharge, and it can minimize the risk of damage in the tokamak that may occur during experiments. A notable example is the Virtual KSTAR (Kwon et al., 2022), a digital twin of KSTAR. The techniques developed in this project are expected to be applied in the future development of digital twins for ITER and K-DEMO (Cho et al., 2022).

Typically, fusion experiments are conducted based on detailed control scenarios, and virtual experiments follow a similar approach; The system should be capable of generating simulated plasma discharge results based on input scenarios describing the operational plan. To achieve this, a multi-modal model is required to efficiently utilize not only current EFIT data but also various data from the tokamak, including heating information and magnetic field data (*i.e.*, control information). In this context, our proposed method for predicting plasma equilibrium can be extended to serve as a foundation for fusion multimodal models. We believe that our results can potentially contribute to predicting plasma instabilities, preventing disruptions, and gaining novel insights into plasma dynamics, thereby taking a significant step towards realizing data-driven stable fusion reactors.

## REFERENCES

- 540  
541  
542 Abdul Fatir Ansari, Lorenzo Stella, Ali Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin  
543 Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham  
544 Kapoor, Jasper Zschiegner, Danielle C. Maddix, Hao Wang, Michael W. Mahoney, Kari Torkkola,  
545 Andrew Gordon Wilson, Michael Bohlke-Schneider, and Bernie Wang. Chronos: Learning the  
546 language of time series. *Transactions on Machine Learning Research*, 2024.
- 547 Ian Char, Joseph Abbate, Laszlo Bardoczi, Mark Boyer, Youngseog Chung, Rory Conlin, Keith  
548 Erickson, Viraj Mehta, Nathan Richner, Egemen Kolemen, and Jeff Schneider. Offline model-  
549 based reinforcement learning for tokamak control. In *Proc. the Annual Learning for Dynamics  
550 and Control Conference (LADC)*, pp. 1357–1372, 2023.
- 551 Boyuan Chen, Kuang Huang, Sunand Raghupathi, Ishaan Chandratreya, Qiang Du, and Hod Lip-  
552 son. Automated discovery of fundamental variables hidden in experimental data. *Nature Compu-  
553 tational Science*, 2(7):433–442, 2022.
- 554 Ara Cho, Jae-Min Kwon, Hyun-Kyung Chung, Juhung Kim, Jisung Kang, Won-Jae Choi, and  
555 Eun Sang Lee. A planning study for virtual demo development in korea. *Fusion Engineering and  
556 Design*, 176:113026, 2022.
- 557 Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco  
558 Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de las Casas, Craig Don-  
559 ner, Leslie Fritz, Cristian Galperti, Andrea Huber, James Keeling, Maria Tsimpoukelli, Jackie  
560 Kay, Antoine Merle, Jean-Marc Moret, Seb Noury, Federico Pesamosca, David Pfau, Olivier  
561 Sauter, Cristian Sommariva, Stefano Coda, Basil Duval, Ambrogio Fasoli, Pushmeet Kohli, Ko-  
562 ray Kavukcuoglu, Demis Hassabis, and Martin Riedmiller. Magnetic control of tokamak plasmas  
563 through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- 564 Nicola Demo, Marco Tezzele, and Gianluigi Rozza. Pydmd: Python dynamic mode decomposition.  
565 *Journal of Open Source Software*, 3(22):530, 2018. doi: 10.21105/joss.00530. URL <https://doi.org/10.21105/joss.00530>.
- 566  
567 D. Eldon, A.W. Hyatt, B. Covele, N. Eidietis, H.Y. Guo, D.A. Humphreys, A.L. Moser, B. Sammuli,  
568 and M.L. Walker. High precision strike point control to support experiments in the diii-d small  
569 angle slot divertor. *Fusion Engineering and Design*, 160:111797, 2020.
- 570  
571 J.R. Ferron, M.L. Walker, L.L. Lao, H.E. St. John, D.A. Humphreys, and J.A. Leuer. Real time  
572 equilibrium reconstruction for tokamak discharge control. *Nuclear Fusion*, 38(7):1055, 1998.
- 573  
574 Jeffrey P. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007.
- 575  
576 Zhangyang Gao, Cheng Tan, Lirong Wu, and Stan Z. Li. Simvp: Simpler yet better video prediction.  
577 In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 3170–  
578 3180, 2022.
- 579  
580 Harold Grad and Hanan Rubin. Hydromagnetic equilibria and force-free fields. *Journal of Nuclear  
581 Energy*, 7(3-4):284–285, 1958.
- 582  
583 A.A. Harms, K.F. Schoepf, and D.R. Kingdon. *Principles of Fusion Energy: An Introduction to  
584 Fusion Energy for Students of Science and Engineering*. World Scientific, 2000.
- 585  
586 Yao Huang, Bing-Jia Xiao, and Zheng-Ping Luo. Fast parallel grad-shafranov solver for real-time  
587 equilibrium reconstruction in east tokamak using graphic processing unit\*. *Chinese Physics B*,  
588 26(8):085204, 2017.
- 589  
590 J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):  
591 90–95, 2007.
- 592  
593 Ian H Hutchinson. Principles of plasma diagnostics. *Plasma Physics and Controlled Fusion*, 44  
(12):2603–2603, 2002.
- D. Iglesias, P. Bunting, S. Esquembri, J. Hollocombe, S. Silburn, L. Vitton-Mea, I. Balboa, A. Huber,  
G.F. Matthews, V. Riccardo, F. Rimini, and D. Valcarcel. Digital twin applications for the jet  
divertor. *Fusion Engineering and Design*, 125:71–76, 2017.

- 594 K. Ikeda. Progress in the iter physics basis. *Nuclear Fusion*, 47(6):E01, 2007.  
595
- 596 J. Jacquino, V.P. Bhatnagar, J.G. Cordey, L.D. Horton, D.F.H. Start, R. Barnsley, P. Breger, J.P.  
597 Christiansen, S. Clement, S.J. Davies, J.K. Ehrenberg, L.-G. Eriksson, G.M. Fishpool, M. Gade-  
598 berg, P.J. Harbour, H.J. Jäckel, K. Lawson, J. Lingertat, C.G. Lowry, C.F. Maggi, G.F. Matthews,  
599 R.D. Monk, D.P. O'Brien, E. Righi, G. Saibene, R. Sartori, B. Schunke, A.C.C. Sips, M.F.  
600 Stamp, D. Stork, J.D. Strachan, A. Tanga, K. Thomsen, and JET Team. Overview of iter physics  
601 deuterium-tritium experiments in jet. *Nuclear Fusion*, 39(2):235, 1999.
- 602 Semin Joung, Jaewook Kim, Sehyun Kwak, J.G. Bak, S.G. Lee, H.S. Han, H.S. Kim, Geunho Lee,  
603 Daeho Kwon, and Y.-C. Ghim. Deep neural network grad-shafranov solver constrained with  
604 measured magnetic signals. *Nuclear Fusion*, 60(1):016034, 2019.
- 605  
606 Semin Joung, Y.-C. Ghim, Jaewook Kim, Sehyun Kwak, Daeho Kwon, C. Sung, D. Kim, Hyun-  
607 Seok Kim, J. G. Bak, and S. W. Yoon. Gs-deepnet: mastering tokamak plasma equilibria with  
608 deep neural networks and the grad-shafranov equation. *Scientific Reports*, 13(1):15799, 2023.
- 609  
610 Jinsu Kim, Jeongwon Lee, Jaemin Seo, Yeongsun Lee, and Yong-Su Na. Disruption prediction and  
611 analysis through multimodal deep learning in kstar. *Fusion Engineering and Design*, 200:114204,  
612 2024.
- 613  
614 Jae-Min Kwon, Hosaeng Choi, JaeSeok Ki, Sang Yun Park, Sang Hyeon Park, Young Jin Kim,  
615 Hanju Cho, Soojin Kim, Hyun Sic Chae, Kun-Su Lee, Minhoo Woo, Taegu Lee, Dosub Lee,  
616 Jinseop Park, Taehyun Tak, Tongnyeol Rhee, Kimin Kim, and Hyunsun Han. Development of a  
617 virtual tokamak platform. *Fusion Engineering and Design*, 184:113281, 2022.
- 618  
619 L L Lao, S Kruger, C Akcay, P Balaprakash, T A Bechtel, E Howell, J Koo, J Leddy, M Leinhauser,  
620 Y Q Liu, S Madireddy, J McClenaghan, D Orozco, A Pankin, D Schissel, S Smith, X Sun, and  
621 S Williams. Application of machine learning and artificial intelligence to extend efit equilibrium  
622 reconstruction. *Plasma Physics and Controlled Fusion*, 64(7):074001, 2022.
- 623  
624 L.L. Lao, H. St. John, R.D. Stambaugh, A.G. Kellman, and W. Pfeiffer. Reconstruction of current  
625 profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611, 1985.
- 626  
627 G.S. Lee, M. Kwon, C.J. Doh, B.G. Hong, K. Kim, M.H. Cho, W. Namkung, C.S. Chang, Y.C.  
628 Kim, J.Y. Kim, H.G. Jhang, D.K. Lee, K.I. You, J.H. Han, M.C. Kyum, J.W. Choi, J. Hong, W.C.  
629 Kim, B.C. Kim, J.H. Choi, S.H. Seo, H.K. Na, H.G. Lee, S.G. Lee, S.J. Yoo, B.J. Lee, Y.S. Jung,  
630 J.G. Bak, H.L. Yang, S.Y. Cho, K.H. Im, N.I. Hur, I.K. Yoo, J.W. Sa, K.H. Hong, G.H. Kim, B.J.  
631 Yoo, H.C. Ri, Y.K. Oh, Y.S. Kim, C.H. Choi, D.L. Kim, Y.M. Park, K.W. Cho, T.H. Ha, S.M.  
632 Hwang, Y.J. Kim, S. Baang, S.I. Lee, H.Y. Chang, W. Choe, S.G. Jeong, S.S. Oh, H.J. Lee, B.H.  
633 Oh, B.H. Choi, C.K. Hwang, S.R. In, S.H. Jeong, I.S. Ko, Y.S. Bae, H.S. Kang, J.B. Kim, H.J.  
634 Ahn, D.S. Kim, C.H. Choi, J.H. Lee, Y.W. Lee, Y.S. Hwang, S.H. Hong, K.-H. Chung, D.-I. Choi,  
635 and KSTAR Team. Design and construction of the kstar tokamak. *Nuclear Fusion*, 41(10):1515,  
636 2001.
- 637  
638 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. the International  
639 Conference on Learning Representations (ICLR)*, 2019.
- 640  
641 Jingjing Lu, Youjun Hu, Nong Xiang, and Youwen Sun. Fast equilibrium reconstruction by deep  
642 learning on east tokamak. *AIP Advances*, 13(7), 2023.
- 643  
644 J.L. Luxon. A design retrospective of the diiii-d tokamak. *Nuclear Fusion*, 42(5):614, 2002.
- 645  
646 Garry McCracken and Peter Stott. *Fusion: the energy of the universe*. Academic Press, 2013.
- 647  
648 J.-M. Moret, B.P. Duval, H.B. Le, S. Coda, F. Felici, and H. Reimerdes. Tokamak equilibrium  
649 reconstruction code liuqe and its real time implementation. *Fusion Engineering and Design*, 91:  
650 1–15, 2015.
- 651  
652 P.-H. Rebut. Iter: the first experimental fusion reactor. *Fusion Engineering and Design*, 30(1):  
653 85–118, 1995.

- 648 Pu Ren, Chengping Rao, Yang Liu, Jian-Xun Wang, and Hao Sun. Phycrnet: Physics-informed  
649 convolutional-recurrent network for solving spatiotemporal pdes. *Computer Methods in Applied  
650 Mechanics and Engineering*, 389:114399, 2022.
- 651 Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of  
652 Fluid Mechanics*, 656:5–28, 2010. doi: 10.1017/S0022112010001217.
- 653 F C Schuller. Disruptions in tokamaks. *Plasma Physics and Controlled Fusion*, 37(11A):A135,  
654 1995.
- 655 Jaemin Seo, Y.-S. Na, B. Kim, C.Y. Lee, M.S. Park, S.J. Park, and Y.H. Lee. Feedforward beta  
656 control in the kstar tokamak by deep reinforcement learning. *Nuclear Fusion*, 61(10):106010,  
657 2021.
- 658 Jaemin Seo, SangKyeun Kim, Azarakhsh Jalalvand, Rory Conlin, Andrew Rothstein, Joseph Ab-  
659 bate, Keith Erickson, Josiah Wai, Ricardo Shousha, and Egemen Kolemen. Avoiding fusion  
660 plasma tearing instability with deep reinforcement learning. *Nature*, 626(8000):746–751, 2024.
- 661 VD Shafranov. Plasma equilibrium in a magnetic field. *Reviews of plasma physics*, 2:103, 1966.
- 662 John Sheffield. The physics of magnetic fusion reactors. *Rev. Mod. Phys.*, 66:1015–1103, 1994.
- 663 Xingjian Shi, Zhoung Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun WOO.  
664 Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Proc.  
665 the Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- 666 Lucas Spangher, Allen M. Wang, Andrew Maris, Myles Stapelberg, Viraj Mehta, Alex Saperstein,  
667 Stephen Lane-Walsh, Akshata Kishore Moharir, Alessandro Pau, and Cristina Rea. Position:  
668 Opportunities exist for machine learning in magnetic fusion energy. In *Proc. the International  
669 Conference on Machine Learning (ICML)*, pp. 46303–46322, 2024.
- 670 Cheng Tan, Zhangyang Gao, Siyuan Li, and Stan Z Li. Simvp: Towards simple yet powerful  
671 spatiotemporal predictive learning. *arXiv preprint arXiv:2211.12509*, 2022.
- 672 Cheng Tan, Siyuan Li, Zhangyang Gao, Wenfei Guan, Zedong Wang, Zicheng Liu, Lirong Wu, and  
673 Stan Z. Li. Openstl: A comprehensive benchmark of spatio-temporal predictive learning. In *Proc.  
674 the Advances in Neural Information Processing Systems (NeurIPS)*, pp. 69819–69831, 2023.
- 675 Baonian Wanfor the EAST, HT-7 Teams, and International Collaborators. Recent experiments in the  
676 east and ht-7 superconducting tokamaks. *Nuclear Fusion*, 49(10):104011, 2009.
- 677 J.T. Wai, M.D. Boyer, and E. Kolemen. Neural net modeling of equilibria in nstx-u. *Nuclear Fusion*,  
678 62(8):086042, 2022.
- 679 Chenguang Wan, Zhi Yu, Alessandro Pau, Olivier Sauter, Xiaojuan Liu, Qiping Yuan, and Jian-  
680 gang Li. A machine-learning-based tool for last closed-flux surface reconstruction on tokamaks.  
681 *Nuclear Fusion*, 63(5):056019, 2023.
- 682 Yunbo Wang, Mingsheng Long, Jianmin Wang, Zhifeng Gao, and Philip S Yu. Predrnn: Recurrent  
683 neural networks for predictive learning using spatiotemporal lstms. In *Proc. the Advances in  
684 Neural Information Processing Systems (NeurIPS)*, 2017.
- 685 Yunbo Wang, Haixu Wu, Jianjin Zhang, Zhifeng Gao, Jianmin Wang, Philip S. Yu, and Mingsheng  
686 Long. Predrnn: A recurrent neural network for spatiotemporal predictive learning. *The IEEE  
687 Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45(2):2208–2225, 2023.
- 688 J. Wesson and D.J. Campbell. *Tokamaks*. OUP Oxford, 2011.
- 689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

A ADDITIONAL QUANTITATIVE RESULTS

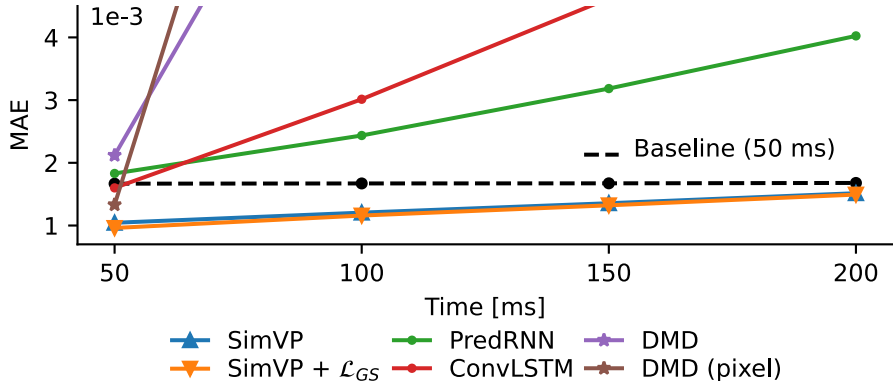


Figure 4: MAEs of  $\psi$  predictions across various methods. This plot presents MAE values for  $\psi$  predictions in the center region at different future time points. All methods, except SimVP, had lower accuracy than the 50 ms lagged baseline. At 50 ms ahead, DMD with pixel-level predictions performed comparably to SimVP. The results of Chronos and PredRNNv2, which showed relatively larger errors compared to other models, were omitted.

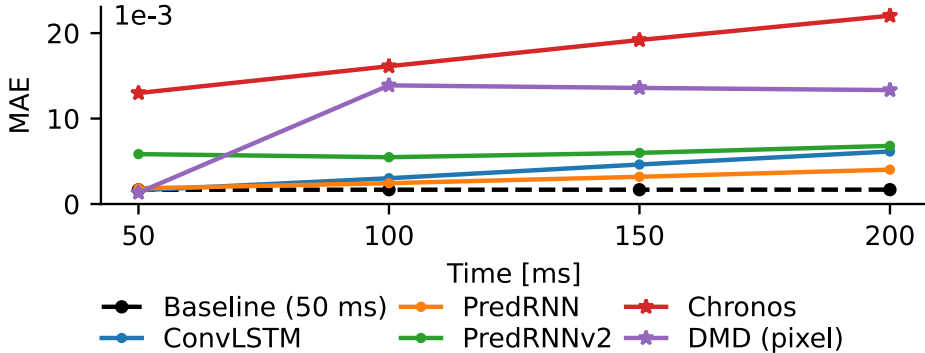


Figure 5: MAEs of  $\psi$  predictions in the center region at different future time points. This plot shows the models (Chronos and PredRNNv2) and y-axis values omitted in Figure 4. The MAE values of the DMD model, which reach up to 3.59, are excluded due to their excessively high magnitude. Unlike other methods, the pixel-level DMD model exhibits a rapid increase in error at 100 ms.

Table 1 illustrates that the models including DMD, Chronos, ConvLSTM, PredRNN, and PredRNNv2 demonstrated higher errors compared to the time-lagged baseline. However, Figure 4 demonstrates that these models did not fail entirely in predictions across all timesteps. For predictions 50 ms ahead, all methods except PredRNNv2 and Chronos achieved accuracy comparable to the baseline. Notably, DMD with pixel-level predictions performed similarly to SimVP.

Beyond 100 ms, the prediction errors of the DMD models increased sharply, exhibiting the steepest rise among the comparative models. This may be attributed to numerical instability during the fitting process of the DMD models, particularly when simultaneously predicting  $\psi$  and  $\Delta^* \psi$ . Evidence includes warning messages from the `PyDMD` library and the observation that the MAE at 200 ms reaches 3.59, which is approximately 1,700 times larger than the error at 50 ms.

Interestingly, even the pixel-level DMD model, which involves far fewer variables, exhibits a significant increase in error as depicted in Figure 5. This suggests that additional factors contribute to the observed performance drop. One possible explanation is that accurately predicting the temporal dynamics of plasma requires accounting for spatial information. Furthermore, the non-linear nature of plasma dynamics may limit the ability of linear models such as DMD to predict long-term behavior beyond 100 ms. However, these hypotheses cannot be conclusively verified with the current experimental results, highlighting the need for further investigation.

## B DETAILS AND LIMITATIONS OF IOU-BASED LCFS EVALUATION




Figure 6: IoU measurements of the LCFS obtained from two consecutive frames in Shot #19993, varying the number of contour lines. The left y-axis represents the IoU, while the right y-axis shows the computation time for both the LCFS and IoU.

The nuclear fusion community has typically measured accuracy by comparing parametric information extractable from the LCFS, such as poloidal angle (Wan et al., 2023). However, these methods are limited to examining only partial aspects of the LCFS shape. This limitation highlights the need for quantitative evaluation metrics that can measure the similarity of the overall LCFS shape.

To address this issue, we experimented with the IoU metric, which is familiar to ML researchers, as a potential method to compare the similarity between the ground truth LCFS and the predicted LCFS. The process begins by drawing  $N$  contour lines using `Matplotlib` (Hunter, 2007). These lines are stored as lists of  $(R, Z)$  coordinates, allowing for the identification of closed surfaces by comparing the start and end coordinates of each contour line. The outermost closed surface is selected as the LCFS and converted into a `Polygon` object of `Shapely` library to calculate the IoU. We calculated the IoU for each timestep and then reported the mIoU across the entire test data.

While this method of estimating the LCFS and IoU is straightforward, it has limitations that hinder its use as a primary evaluation metric. As illustrated in Figure 6, the IoU calculation is sensitive to the number of contour lines used. Based on this result, we concluded that IoU should be used only as a supplementary evaluation metric to MAE.