

On learning latent dynamics of the AUG plasma state

A. Kit,^{1,2, a)} A.E. Järvinen,² Y.R.J. Poels,^{3,4} S. Wiesen,⁵ V. Menkovski,³ R. Fischer,⁶ M. Dunne,⁶ and ASDEX-Upgrade Team

¹⁾ *University of Helsinki, FI-00014 Helsinki, Finland*

²⁾ *VTT Technical Research Centre of Finland, FI-02044 VTT, Finland*

³⁾ *Eindhoven University of Technology, Mathematics and Computer Science, NL-5600MB Eindhoven, The Netherlands*

⁴⁾ *École Polytechnique Fédérale de Lausanne, Swiss Plasma Center, CH-1015 Lausanne, Switzerland*

⁵⁾ *Forschungszentrum Jülich GmbH, Institut für Energie- und Klimaforschung - Plasmaphysik, DE-52425 Jülich, Germany*

⁶⁾ *Max-Planck-Institut für Plasmaphysik, D-85748 Garching, Germany*

(Dated: 19 March 2024)

In this work, we demonstrate the utility of state representation learning applied to modeling the time evolution of electron density and temperature profiles at ASDEX-Upgrade (AUG). The proposed model is a deep neural network which learns to map the high dimensional profile observations to a lower dimensional state. The mapped states, alongside the original profile’s corresponding machine parameters are used to learn a forward model to propagate the state in time. We show that this approach is able to predict AUG discharges using only a selected set of machine parameters. The state is then further conditioned to encode information about the confinement regime, which yields a simple baseline linear classifier, while still retaining the information needed to predict the evolution of profiles. We then discuss the potential use cases and limitations of state representation learning algorithms applied to fusion devices.

I. INTRODUCTION

We investigate the use of state representation learning to model the time evolution of plasmas at ASDEX-Upgrade (AUG). As reviewed in¹, state representation learning (SRL) focuses on learning low dimensional features of an environment that evolve in time and are influenced by actions. An SRL model posits a system’s state at a given time, s_t , with observations, o_t , which are noisy measurements of the state. The state evolves under the influence of actions, leading to future states s_{t+1} , which can again be measured, o_{t+1} . In this work, we consider AUG to be the environment in which actions, $a_t \in \mathcal{A}$, are made at a time step t , where \mathcal{A} is the action space. At AUG, actions are ‘machine control parameters’, such as plasma current, magnetic field strength, gas puffing rate, etc. The change in machine parameters induces a change in the plasma state, s_t to s_{t+1} . Full information of the true plasma state is not accessible, but diagnostic systems, such as Thomson scattering and reflectometry, provide partial and noisy observations, $o_t \in \mathcal{O}$, of the plasma state. Examples of machine parameters and observations are depicted in Figure 1.

The goal of this work is to investigate methods to learn a useful representation of the AUG plasma state, $s_t \in \mathcal{S}$, with which we can then learn a forward model, $p(s_{t+1}|s_t, a_t)$, to predict the evolution and dynamics of the plasma state. In this work, a useful state representation is defined as one which represents the high dimensional observations in a state that conforms to the actual

degrees of freedom of the system, i.e., a state that retains the information content of the observations.

Compressed or lower dimensional state representations are desirable for control^{2,3}, predictive inference^{4,5}, and interpretation⁶, among others. Ideally, such representations contain the essential aspects of the system. In practice, it is hard to enforce this, thus, the overarching question we seek to address is: What constitutes the objective function that learns a useful state?

In this work, a variational autoencoder⁷ (VAE) is used to learn a state representation of electron density and temperature profile measurements of AUG plasmas. The state representation and machine parameters are used to train a forward model to predict the dynamical evolution of the state. The end result is a model that can predict dynamical evolution of the electron density and temperature profiles directly from a sequence of machine parameters (Fig. 1).

II. DATASET

The dataset used in this analysis consists of 1000 high confinement mode (H-mode) discharges that are non-disruptive, deuterium fuelled, and without impurities. For each pulse, observations are outer-mid-plane electron profiles of density, n_e and temperature, T_e , and actions are machine control parameters.

The electron profiles are obtained from the IDA⁸ at AUG, which applies Bayesian probability theory to fit a spline to core-edge measurements originating from lithium beams, electron cyclotron emissions, Thomson scattering and interferometry.

The following machine control parameters were se-

^{a)}Electronic mail: adam.kit@helsinki.fi

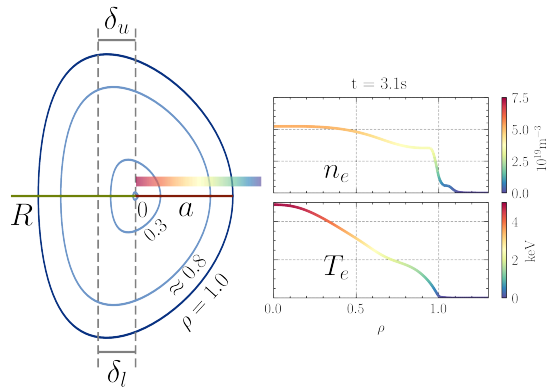


FIG. 1. A representation of the plasma state at ASDEX-Upgrade. On the left, a 2D cross section of the plasma with various flux surfaces labeled by their flux surface coordinate ρ . The confined region of the plasma spans from the core ($\rho = 0.0$) to the separatrix ($\rho = 1.0$). Machine parameters related to the shape of the plasma are labeled; the upper/lower triangularity, δ_u/l , the major and minor radius, R, a . On the right, observations of the electron density and temperature are visualized. The main plasma kinetic profiles are typically remapped to the outer-mid-plane (location corresponding to the colorbar in the left). Flux surfaces move during the course of the discharge, as does the magnitude of the electron profiles; we seek to model the dynamics of the kinetic profiles in this work.

lected: total plasma current, I_P , safety factor magnitude at 95% flux surface, q_{95} , total deuterium injection rate, D_{TOT} , plasma major radius, R , plasma elongation κ , upper triangularity δ_u , lower triangularity, δ_l , aspect ratio, $A = R/a$ (where R and a are the major and minor radii of the plasma), and total heating power normalized by the Martin LH-threshold scaling⁹, P_{TOT}/P_{LH} . The motivation for using the Martin scaling instead of just P_{TOT} is that P_{TOT}/P_{LH} is a normalized parameter with respect to plasma scenarios and can be applied for any device size that was used in the scaling. The chosen machine parameters are considered to be ‘controllable’ even though not all are strictly knobs on the tokamak, i.e., I_P is not the current through the central solenoid but is a quantity that is achieved through attenuation of controllable parameters (the current through the central solenoid). The same holds for the parameters related to the plasma shape even though they are all reconstructed values of the plasma state (δ_l, δ_u, A, R and κ). The machine parameters are linearly interpolated in time with respect to the IDA measurement frequency (1kHz) in order to homogenize the sampling frequencies of both actions and observations.

Then, the combined set of observations and actions are time-wise downsampled, which transformed the time step frequency from 1kHz to 200Hz, i.e., observations and actions were selected every 5ms as opposed to the 1ms true sampling interval. This is done in part because the observations and actions chosen in this work tend to have

low variability within 5ms intervals.

The data is split into training, validating and testing subsets, consisting of 853 (~ 5500 real-time seconds), 137, 137 discharges respectively. Additionally, discharges coming from the same shot request are binned into the same subset, which helps ensure that the training set does not include similar discharges as the validating and testing sets. The observations and actions are z-score normalized via the mean and standard deviation of the training set.

III. MODEL

The model follows ‘World Modeling’ as first proposed in¹⁰; here an **observational** model is used to compress measurements at a given time to a latent distribution, s_i , and a **forward** model to evolve this distribution into the future, s_j where $j > i$. Following more recent advances in World Modeling^{2,11}, we additionally train the observational and forward models in one computational graph. Finally, a physics prior is introduced to guide the representation to be more physically informative.

The goal of the observational model is to learn a function that can reconstruct observations from a given state, i.e., learn $\phi(\theta_\phi) : o_t \rightarrow s_t$ and $\pi(\theta_\pi) : s_t \rightarrow \hat{o}_t$, where ϕ has parameters θ_ϕ . We assume that these relations are not deterministic therefore we argue to treat these function as probability distributions. To do so, we employ a VAE, a probabilistic generative model, consisting of an encoder (ϕ) and decoder (π) distribution. The encoder and decoder distribution are parameterized by neural networks. The distributions are learned by minimizing the reconstruction error (L^2 norm) between observations, o_t , and their reconstructions, \hat{o}_t , in combination with a regularizing term (Kullback-Leibler divergence) on a prior belief of the encoder distribution, $\phi_{\theta_\phi}(s_t | o_t)$, resulting in our implementation of the VAE objective function:

$$\mathcal{L}_{\text{obs}} = \mathbb{E}_{\phi, \pi} [||\hat{o}_t - o_t||_2] + \text{KL}(\phi_{\theta_\phi}(s_t | o_t) || p(s_t))$$

where $p(s_t) = \mathcal{N}(0, 1)$ is our prior belief about the state distribution, and the expectation $\mathbb{E}_{\phi, \pi}$ is estimated using the *re-parameterization trick*⁷. The prior belief about the state distribution is strictly a design choice to allow for an unconstrained model; future work would include identifying physics-based distribution as priors. The effect of different norms for the reconstruction loss was not explored in this work due to the efficiency of the L^2 norm. The architecture of the observational model is given in Table I.

The goal of the forward model is to predict a future state s_{t+1} from the previous state s_t and machine control parameters a_t , i.e., to learn the mapping $f(\theta_f) : s_t, a_t \rightarrow \hat{s}_{t+1}$. Since the state s_t is a distribution, the output of the forward model, \hat{s}_{t+1} , is the parameters of a probabilistic distribution, also parameterized by a neural network,

TABLE I. Observational model architecture. The observations are 1D profiles with two channels, the density and temperature. The 1D convolution and 1D transposed convolution layers are denoted as Conv. and Transp. Conv. respectively. The parameters of the convolution are denoted as (in channels= i , out channels= o , kernel width= k , stride= s), i.e., a Convolution layer with 2 in channels, 4 out channels, a kernel width of 4 and stride of 2 would be denoted as Conv. (2, 4, 4, 2). The Encoder to State has two components, denoting the mean and standard deviation of the latent variable.

Model Component	Layer(s)	Activation
Encoder	Conv. (2, 4, 4, 2)	ReLU
	Conv. (4, 8, 4, 2)	ReLU
	Conv. (8, 16, 4, 2)	ReLU
	Conv. (16, 32, 4, 2)	None
Encoder to State	Dense (320, 8), Dense (320, 8)	None
Decoder	Dense (8, 128)	None
	Transp. Conv. (128, 16, 5, 3)	ReLU
	Transp. Conv. (16, 8, 6, 3)	ReLU
	Transp. Conv. (8, 4, 6, 3)	ReLU
Output	Transp. Conv. (4, 2, 6, 3)	ReLU
	Dense (165, 200)	None

TABLE II. Forward Architecture. The initial layer of the forward model has size 8 (state size) + 9 (action size) = 17. Similar to the observational model, the forward model outputs the parameters of a distribution, therefore the last layers parameterize the mean and standard deviation.

Model Component	Layer(s)	Activation
State to State	Dense(8+9, 20)	None
	Dense (20, 8), Dense (20, 8)	None

from which \hat{s}_{t+1} is sampled. To match s_{t+1} with \hat{s}_{t+1} the following objective function is used:

$$\mathcal{L}_f = \mathbb{E}_{\phi, f} [\text{KL}(\phi(s_{t+1} | o_{t+1}) || f(\hat{s}_{t+1} | s_t, a_t))]$$

where once again, the Kullback-Liebler divergence is used. The architecture of the forward model is given in Table II.

Together, a prediction of a future state, s_{t+1} , is made by first encoding o_t to s_t , then transitioning from s_t to s_{t+1} (Fig 2). In this fashion, the forward model and observational model are trained simultaneously by minimizing the following objective:

$$\mathcal{L} = \mathcal{L}_{\text{obs}} + \mathcal{L}_f$$

It is worth noting here that by allowing gradients to flow from \mathcal{L}_f back through to the encoder of the observational model, the learned state representation is expected to retain properties that facilitate state dynamics prediction by the forward model.

Additionally, two forms of regularization are added to the model: i) a penalty on violation of static pressure conservation in reconstruction $\propto ||n_t T_t - \hat{n}_t \hat{T}_t||_1$ and ii) the *pushforward* trick from¹², where the forward model predicts \hat{s}_{t+i} , with $i > 1$, using the previous forward model prediction of s_{t+i-1} . Ideally, the pressure penalty

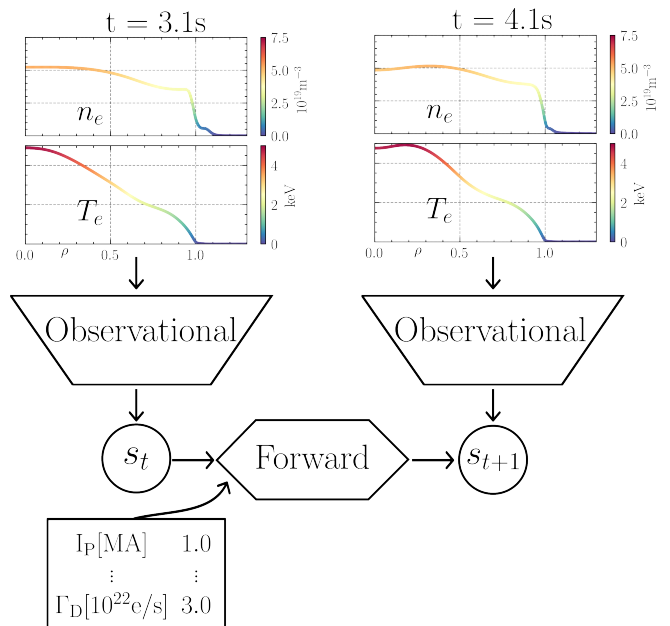


FIG. 2. Graphical representation of the full model. Electron profiles are encoded via the *observational* model to the state s_t . To predict s_{t+1} , the *forward* model takes s_t and actions a_t , here the plasma current I_p and gas puff rate Γ_D . The observational model can be used to decode the state to retrieve kinetic profiles.

encourages the observational model to encode physically consistent electron density and temperature reconstructions and was first explored in⁴. The pressure penalty is very similar to the reconstruction error, however we believe the pressure penalty helps to regularize the predictions of the density and temperature with respect to fluctuations around a pressure value. For example, if the reconstruction error is 0, then the pressure error is also 0. Yet, if the reconstruction error is non-zero, then the pressure error could be ≥ 0 , i.e., the fluctuations in temperature and density may even out to yield zero pressure error. If the model is wrong, we would rather the model learn to be wrong in this way. The *push-forward* trick¹² aims to stabilize auto-regressive models in long-range planning. During training, the number of time steps to rollout, i , is determined per mini-batch by sampling from a uniform distribution $U[0, N]$ where N is the number of epochs trained thus far. The loss is only calculated between the final rollout state, \hat{s}_{t+i} , and corresponding state s_{t+i} . In other words, we cut the gradients in the unrolling stage.

Since the space of observations \mathcal{O} is constrained to \mathcal{R}^+ , the output of the observation model is clamped to output only positive real values during training. This is done by clamping the output of the observational decoder.

Due to the competing objectives and to normalize the reconstructed pressure penalty with respect to the remaining objectives, we found it useful to weight individual components with scaling factors. Training hyperpa-

parameters and objective penalty weights are given in Table III. A rigorous search for optimal hyperparameters, including state size, was not conducted. However, the final configurations III were selected among others by obtaining the lowest error on the validation dataset.

TABLE III. Objective weights and training parameters used for the SRL model. All weights are scalar multiplied by their corresponding objective value per mini-batch update. KL_{obs} is applied to the KL term in \mathcal{L}_{obs} . KL_{f} is applied to the KL term of in \mathcal{L}_{f} . L_{ot}^1 is applied to the L^1 term of \mathcal{L}_{obs} . L_{pt}^1 is applied to the L^1 pressure penalty.

Objective Weights				
Name	KL_{obs}	KL_{f}	L_{ot}^1	L_{pt}^1
Value	0.01	1.0	100	0.0001
Training Hyperparameters				
Name	Batch Size	Optimizer	Learning Rate	
Value	8	Adam ¹³	0.02	

IV. RESULTS

A. Observational Model

The quality of the observational model can be determined by comparing the observation with the reconstruction (Figure 3). For the test-set discharges, a mean absolute error (MAE) between reconstruction and observation has a mean of $0.28 \pm 0.13 \cdot 10^{19} \text{m}^{-3}$ and 0.11 ± 0.07 (keV) for density and temperature, respectively. These lossy compression results are expected, as the state space is only 8-D while the observational space is 400 points for each time step. Increasing the state space dimensionality would likely yield lower reconstruction error. The average reconstruction error (MAE) of the test set discharges for $\rho = 0.0, 0.5, 0.9, 1.0$ is given in Table IV.

We find the reconstruction quality of the observational model to be sufficiently accurate to proceed with the forward model.

TABLE IV. The MAE of the observational model’s reconstructions for various ρ . The MAE for AUG #36150 (Fig. 3) is provided to compare with the average over the test set discharges. Standard deviations for all values are calculated over 100 sample reconstructions, given the injection of noise in the VAE.

Average MAE over test set discharges				
ρ	0.0	0.5	0.9	1.0
n_e (10^{19}m^{-3})	0.46 ± 0.6	0.31 ± 0.13	0.33 ± 0.17	0.41 ± 0.28
T_e (keV)	0.3 ± 0.12	0.11 ± 0.09	0.06 ± 0.03	0.03 ± 0.06
AUG #36150 (Fig. 3)				
n_e (10^{19}m^{-3})	0.23 ± 0.18	0.43 ± 0.27	0.45 ± 0.27	0.78 ± 0.6
T_e (keV)	0.26 ± 0.17	0.16 ± 0.12	0.1 ± 0.06	0.06 ± 0.06

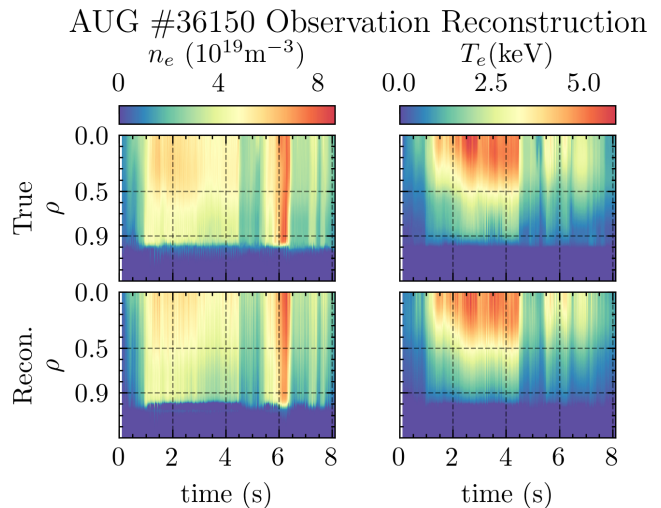


FIG. 3. Observational model’s reconstruction of AUG discharge #36150. The left and right figures show the density and temperature profile evolutions, respectively. The top and bottom plots show the true profiles and model reconstruction, respectively. The x and y-axis on all figures are the same. The reconstruction error is similar to that of the average of the test set discharges, as the MAE of the density and temperature profiles averaged over this discharge are 0.32 ± 0.29 (10^{19}m^{-3}) and 0.12 ± 0.12 (keV), respectively. Mean and standard deviation of the errors are calculated over 100 sample reconstructions. Respective errors at $\rho = 0.0, 0.5, 0.9, 1.0$ are given in Table IV.

B. Forward Model

To determine the predictive quality of the forward model on AUG discharges, we first encode the observations at $t = 0$ to a state s_0 via the observational model, then the forward model rolls s_0 out to the final time step using true actions and its own predictions of $\hat{s}_{t>0}$. Each \hat{s}_t is then decoded into profiles via the observational decoder. The MAE of all time steps over all discharges for the forward model is $0.97 \pm 0.55 \cdot 10^{19} \text{m}^{-3}$ and 0.35 ± 0.24 (keV) for density and temperature, respectively. The average reconstruction error (MAE) of the test set discharges for $\rho = 0.0, 0.5, 0.9, 1.0$ is given in Table V. The average percentage reconstruction error (MAPE) is given in Table VI. Both the MAE and MAPE are reported due to large variations radially of the magnitude of density and temperature, i.e., at $\rho > 1.0$ density and temperature are relatively low compared to $\rho = 0.0$. The mean accumulation of error does not rapidly increase over time for test-set discharges (Fig. 4).

The forward model is able to capture the profile evolution for various discharges, for example, a plasma scenario with feedback density control (Fig. 5), as well as a power step-wise ramp up (Fig. 6).

TABLE V. The MAE of the forward model’s reconstructions for various ρ . The MAE of the AUG discharges visualized in this work are provided for comparison with the average over the test set discharges. Standard deviations for all values are calculated over 100 sample reconstructions.

Average MAE over test set discharges				
	$\rho = 0.0$	0.5	0.9	1.0
n_e (10^{19} m^{-3})	1.51 ± 0.83	1.13 ± 0.76	1.07 ± 0.61	0.68 ± 0.43
T_e (keV)	0.96 ± 0.72	0.36 ± 0.24	0.15 ± 0.11	0.05 ± 0.09
AUG #34828 (Fig. 5)				
n_e	0.597 ± 0.57	0.39 ± 0.36	0.38 ± 0.28	0.25 ± 0.21
T_e	0.87 ± 0.33	0.11 ± 0.1	0.07 ± 0.05	0.04 ± 0.02
AUG #36022 (Fig. 6)				
n_e	0.67 ± 0.57	0.83 ± 0.26	0.74 ± 0.31	0.5 ± 0.37
T_e	0.54 ± 0.41	0.17 ± 0.1	0.09 ± 0.06	0.02 ± 0.02
AUG #36150 (Fig. 10)				
n_e	2.57 ± 1.02	1.86 ± 0.93	1.76 ± 0.79	2.99 ± 0.84
T_e	1.22 ± 1.05	0.55 ± 0.39	0.32 ± 0.28	0.54 ± 0.21
AUG #36669 (Fig. 11)				
n_e	1.21 ± 1.23	1.01 ± 0.86	0.88 ± 0.79	0.74 ± 0.5
T_e	0.87 ± 0.51	0.40 ± 0.45	0.15 ± 0.09	0.1 ± 0.02

TABLE VI. The mean-absolute percentage error (MAPE) of the forward model’s reconstructions for various ρ . The MAPE is calculated as the L^1 difference between the predicted and true value, divided by the true value. The MAPE of the AUG discharges visualized in this work are provided for comparison with the average over the test set discharges. Standard deviations for all values are calculated over 100 sample reconstructions. Large deviations (MAPE > 100%) in the edge are expected, as the temperature and density tend to be relatively low ($n_e < 10^{18} \text{ m}^{-3}$, $T_e < 140 \text{ eV}$).

Average MAPE over test set discharges				
ρ	0.0	0.5	0.9	1.0
n_e (%)	28.89 ± 20.37	28.55 ± 19.97	34.31 ± 29.76	376.6 ± 326.6
T_e (%)	65.82 ± 101.6	26.53 ± 17.71	73.60 ± 167.2	208.7 ± 353.4
AUG #34828 (Fig. 5)				
n_e	10.73 ± 14.63	12.47 ± 15.57	12.05 ± 10.23	17.73 ± 15.60
T_e	31.33 ± 15.38	13.67 ± 15.33	24.29 ± 27.46	52.74 ± 79.44
AUG #36022 (Fig. 6)				
n_e	12.97 ± 10.18	15.775 ± 10.15	15.837 ± 15.82	22.33 ± 54.52
T_e	21.70 ± 16.77	13.788 ± 10.32	31.516 ± 97.79	43.27 ± 129.8
AUG #36150 (Fig. 10)				
n_e	53.9 ± 53.54	44.58 ± 47.45	54.44 ± 56.12	263.6 ± 147.1
T_e	42.76 ± 46.58	27.17 ± 21.95	51.54 ± 57.96	521.5 ± 177.8
AUG #36669 (Fig. 11)				
n_e	33.14 ± 46.30	32.66 ± 46.56	40.44 ± 84.46	104.4 ± 311.2
T_e	25.77 ± 17.39	17.74 ± 16.20	18.59 ± 20.06	22.31 ± 20.98

C. Forward Model with Auxiliary Regressor

Expanding on previous work⁴, we further regularize the state representation by learning a mapping $r(\theta_r) : s_t \rightarrow P_{\text{TOT}}/P_{\text{LH}}, \tau_E$, where $P_{\text{TOT}}/P_{\text{LH}}$ is the aforementioned normalized power action, and τ_E the global confinement time. Inspired by⁶, we split s_t into two subspaces $s_{c,t}$ and $s_{/c,t}$, and apply the mapping only on $s_{c,t}$. The dimensionality of $s_{c,t}$ is 2, one dimension for each regressed variable, and the dimensionality of $s_{/c,t}$ is kept to 6 so that the total dimensionality of s_t is preserved from the previous experiment. Then, the mapping $r(\theta_r)$ is made to be linear with respect to the regressed variables $P_{\text{TOT}}/P_{\text{LH}}$ and τ_E , i.e., a diagonal matrix that maps one dimension of $s_{c,t}$ to $P_{\text{TOT}}/P_{\text{LH}}$ and the other remaining dimension to τ_E . The original loss function \mathcal{L} then gains

the following additional $\mathcal{L}1$ loss term:

$$\mathbb{E}_{r,\phi}[||a_{c,t} - \hat{a}_{c,t}||_1]$$

where $a_{c,t}$ and $\hat{a}_{c,t}$ are the true regressed variables and their reconstructions, respectively.

With the additional regressor, it was seen that the model can infer the confinement time and power variable via an observed a state s_t . The observed state can be encoded by either the observational model or predicted as before using the forward model (Fig. 7).

While the reconstructed value of τ_E tends to be higher than the true value, the reconstructed $P_{\text{TOT}}/P_{\text{LH}}$ values are quite close to true observations. The elevated τ_E predictions might be caused by biases originating from the beginning and end of the plasma discharges and will be investigated in futures studies. The main message in this




FIG. 4. The test set forward model error (MAPE) as a function of time. The error per step is calculated as the average over the density and temperature profiles up to $\rho \leq 1.0$ for that step. The reason for radial cutoff is the very low values of temperature and density at $\rho > 1.0$. The spikes at the beginning and end are likely due the discharge entering and exiting H-mode, where the density and temperature rapidly change and are therefore difficult to precisely match.

proof-of-principle work is to demonstrate the attachment of semantically meaningful information from the plasma state to the trained state representation with the auxiliary regression modules.

Since one dimension of $s_{c,t}$ is linear with respect to $P_{\text{TOT}}/P_{\text{LH}}$, we receive a simplified H-mode classifier without additional training. Assuming that $P_{\text{TOT}}/P_{\text{LH}}$ is sufficiently accurate in quantifying the presence of H-mode, i.e, a plasma is in H-mode when $P_{\text{TOT}}/P_{\text{LH}} \geq 1$ and sufficient predictive quality of the auxiliary regressor, then via the linear mapping r , there exists an equivalent threshold, $s_{c,t} > H_{\text{thresh}}$ within the relevant power dimension of $s_{c,t}$, (Fig. 8).

However, the additional objective comes at cost, as we observed small penalties on the reconstructive quality in the forward and observational models (Fig. 9). This is likely due to the two dimensions no longer free to compress information *only* pertaining to profile reconstruction information. It is likely that increasing the size of the state would resolve this.

D. Model Limitations

The forward model is limited by the type and distribution of machine parameters that is provided. For example, as only the total heating power is provided, any discharges where input power mixtures is varied midshot are subject to relatively large prediction errors (Fig. 10). Another example is strong tungsten accumulation in the plasma, such as observed in some of the actively cooled divertor experiments (Fig. 11). Since the actions selected in this work do not show the corresponding variations seen on these discharges, the forward model will mispredict the resulting profiles. However, the experimental observations generally fall within the 90% confidence in-

AUG #34828 Forward Model Reconstruction

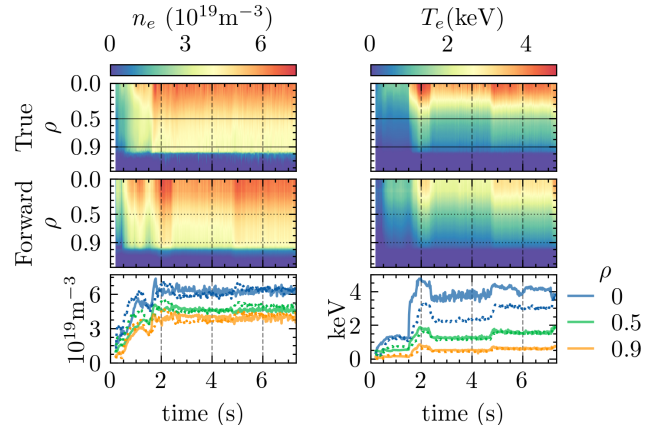


FIG. 5. AUG #34828 comparison of true (top) and forward model predicted (middle) electron density and temperature profiles are plotted along with time traces of the radial values at $\rho = 0.0, 0.5$ and 0.9 (bottom). All rows in the left column are associated with density, the right with temperature. The solid/dotted grid lines in the top/middle plot correspond to the solid/dotted traces in the bottom plot. The initial state s_0 is sampled from the encoder of the observational model given the profiles at t_0 . The initial state is then propagated in time alongside actions and previous state prediction via the forward model. The density prediction is worse at the beginning of the pulse, likely due to the fluctuation feedback of the gas flow rate, but eventually stabilizes to the true value. The resulting reconstructions of the predicted states from the forward model demonstrate the capability to handle sufficiently complex plasma scenarios. Respective errors of the density and temperature at $\rho = 0.0, 0.5, 0.9, 1.0$ are given in Table V & VI.

terval (Fig. 12).

V. DISCUSSION

In this work, we have demonstrated the utility of state representation learning towards learning the plasma state at ASDEX-Upgrade. Our proposed model is able to predict the electron density and temperature profiles from machine parameters only. Additionally, we demonstrate the functionality of learning a state representation by incorporating a simplified H-mode classifier into the model while retaining the ability to predict the time evolution of plasma profiles.

The MAPE/MAE errors are reported in Table V, VI. While these points estimates can have significant deviations, we note that the correct profile generally falls within the 90% confidence interval, as depicted in Figure 12. Future work will explore more rigorous evaluations of the predictive distribution.

The forward model developed in this work has a very limited capacity (Table II). It might be useful to improve the forward model's capacity, e.g., via a recurrent neural

AUG #36022 Forward Model Reconstruction

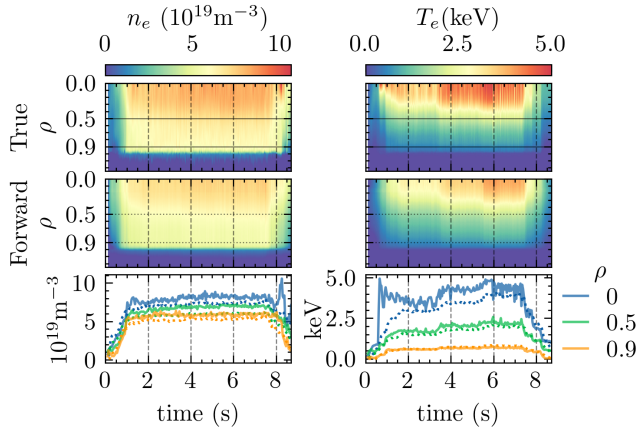


FIG. 6. AUG #36022 comparison of true (top) and forward model predicted (middle) electron density and temperature profiles are plotted along with time traces of the radial values at $\rho = 0.0, 0.5$ and 0.9 (bottom). All rows in the left column are associated with density, the right with temperature. The solid/dotted grid lines in the top/middle plot correspond to the solid/dotted traces in the bottom plot. The initial state s_0 is sampled from the encoder of the observational model given the profiles at t_0 . The initial state is and then propagated in time alongside actions and previous state prediction via the forward model. In this discharge, the total power input is increased step-wise, leading to a step-wise increase in core electron temperature, which match reconstructions of the predicted state. Respective errors of the density and temperature at $\rho = 0.0, 0.5, 0.9, 1.0$ are given in Table V & VI.

network as demonstrated in². An alternative approach was demonstrated in^{16,17}, where a deep learning variation on Kalman filters is learned by sampling the *transition* between states with a VAE. Also, our approach to using a linear transform for time-stepping latent representations has some commonalities with the work in¹⁸, where Koopman operator theory is used to guide auto-encoders into learning Koopman eigenfunctions from data, i.e., the latent space has globally linear dynamics. A key difference being that we predict the mean and standard deviation of a distribution, which we sample at each time step. Here, perhaps, there is a connection to latent neural stochastic differential equation models¹⁹ (NSDE), where our autoregressive formulation can be considered a crude discretization of such an NSDE.

The cumulative error plot (Fig. 4) is in some sense troubling. We expected to see a step wise accumulation of error, as with traditional forward model predictors. Interestingly, it appears that the observational model encodes information about the machine parameters, even though this information is only propagated through the forward model. As a result, when the plasma reaches flat top, the forward model likely pulls the state prediction to a previously seen steady state representation. We believe this has to do with *verifiability*²⁰, and appears to be both a feature and a bug. An open question is then

AUG #34814




FIG. 7. The predicted and true time traces of τ_E (top) and P_{TOT}/P_{LH} (bottom) from AUG #34814. The predictions of the observational model (Obs.) are obtained by encoding observations to s_t and applying the auxiliary mapping. The predictions of the forward model (Forw.) are obtained by encoding the first observation to an initial state, i.e., $o_0 \rightarrow s_0$, then rolling out with the forward model until the last action.

to what extent the latent dynamic model simply learns to propagate the state through a series of steady states rather than (ideally) predicting the temporal evolution of the plasma. Along this line, we hypothesise that one could learn a steady state model within a DIVA/CCVAE-like framework^{6,21}, as in⁴, treat all time slices as steady state, and ultimately forgo the forward model. Future studies will investigate steady-state and dynamical models within the context of verifiability.

We believe an important question is what information, and at what frequency, is needed to predict future plasma states. Additional questions arise; assuming the true plasma state is Markovian, as we suspect, then what observables are necessary to capture that? Also, if we can approximate the true state sufficiently well in a low dimensional representation, then a) what observations are used to learn such a state and b) what actions are needed to (accurately) propagate that state in time?

Take, for example, at a given time, observations of the the wall and plasma facing components in comparison to observations of the core. If we observe tungsten in the edge via spectroscopic measurements²², this accumulation is not immediately seen in the core. Thus an open question is how to reconcile the differing time scales of differing order phenomena in tokamaks?

A limitation to our model is that it is data-driven and generative. Outside of constraining the output of the decoder to \mathcal{R}^+ , we do not enforce the model to predict ‘physically valid’ plasmas. It is of interest, then, to look into how we may constrain the representation to be physically valid.

Future work would explore including MHD stability and instability information into the state. Such a model could provide time-of-flight information on whether a plasma crosses a stability threshold, and if so, possibly what instability may be triggered.




FIG. 8. **Top:** The time trace of state dimension $s_{c1,t}$ is encoded by the forward model using the actions of AUG # 34814. The horizontal line (colorbar vertical) marks the value of $s_{c1,t}$ which corresponds to an inferred $P_{TOT}/P_{LH} = 1$. The coloring is found by applying the auxiliary regressor to the range of $s_{c1,t} \in [-10, 0]$. Due to the linear capacity of the auxiliary mapping, the output of the mapping on $s_{c1,t}$ on the interval $[-10, 0]$ does not change in time, nor does it change with respect to any other state variable. Like¹⁴, we arrive at a model that can predict different regimes, albeit in very different fashion. **Bottom:** The time trace of P_{TOT}/P_{LH} for AUG #34814, with horizontal line marking where $P_{TOT}/P_{LH} = 1$.

ACKNOWLEDGMENTS

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

This work made use of the Finnish CSC computing infrastructure under project #2005083.

The work of A.E.J and A.K. was partially supported by the Research Council of Finland grant no. 355460.

The authors would like extend thanks to Professor Satoshi Hamaguchi for hosting the ICDDPS-4 conference, where this work was presented as a oral contribution.

A.K. would like to thank Ms. Green for fruitful discussions and artistic inspiration throughout the progress of this work.

A.K. would like to thank Ivan Zaitsev and Kostantinos Papadakis for energetic conversations surrounding the development of this work.




FIG. 9. Even with the auxiliary regressor, the state representation and forward model are still able to capture complex time evolution of AUG discharges. The true density and temperature time traces at various ρ are opaque to show the slight differences with predicted time traces.

DATA AVAILABILITY STATEMENT

The experimental data used for training the deep learning models in this work is stored at the data storage facilities of ASDEX-Upgrade and the authors do not have permission to make this data publicly available. However, the training data preparation routines will be provided on request such that anyone with access to the data can regenerate the training dataset. The Python codes encompassing the deep learning algorithms are available in GitHub at <https://github.com/DIGIfusion/latent-state-modeling>.

¹T. Lesort, N. Díaz-Rodríguez, J. F. Goudou, and D. Filliat, “State representation learning for control: An overview,” *Neural Networks* **108**, 379–392 (2018).

²D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, “Mastering diverse domains through world models,” *arXiv preprint arXiv:2301.04104* (2023).

³J. Abbate, R. Conlin, R. Shousha, K. Erickson, and E. Kolemen, “A general infrastructure for data-driven control design and implementation in tokamaks,” *Journal of Plasma Physics* **89**, 895890102 (2023).

⁴A. Kit, A. Järvinen, S. Wiesen, Y. Poels, and L. Frassinetti, “Developing deep learning algorithms for inferring upstream separatrix density at JET,” *Nuclear Materials and Energy* **34** (2023), 10.1016/j.nme.2022.101347.

⁵B. Zhu, M. Zhao, H. Bhatia, X.-q. Xu, P.-T. Bremer, W. Meyer, N. Li, and T. Rognlien, “Data-driven model for divertor plasma detachment prediction,” *Journal of Plasma Physics* (2022).

⁶T. Joy, S. Schmon, P. Torr, S. N., and T. Rainforth, “Capturing label characteristics in {vae}s,” in *International Conference on Learning Representations* (2021).

⁷D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” 2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings, 8–9 (2014).

⁸R. Fischer, C. J. Fuchs, B. Kurzan, W. Suttrop, and E. Wolfgram, “Integrated data analysis of profile diagnostics at ASDEX upgrade,” *Fusion Science and Technology* **58**, 675–684 (2010).

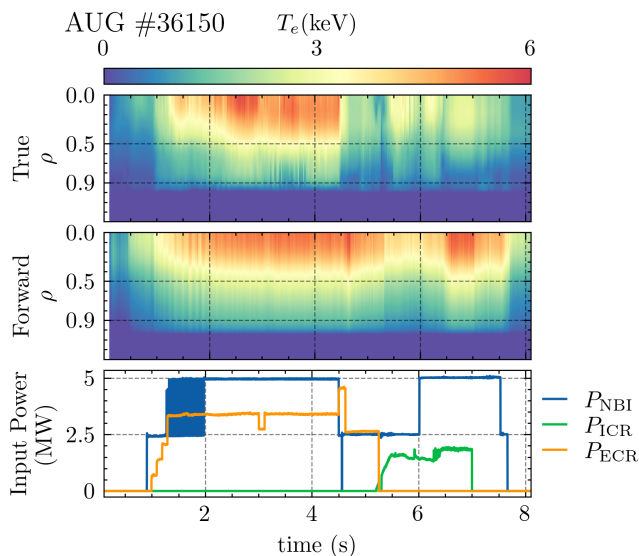


FIG. 10. In AUG #36150, the power starts as mainly NBI and ECR driven, however at 4.5 seconds, NBI is rapidly cut off and supplemented with ICR. Since $P_{\text{TOT}}/P_{\text{LH}}$ is the only available power variable to the forward model, it observes a steady stream of power, which does not induce major changes in the inferred plasma state. It is likely that including separate variables for each power parameter would increase the resilience of the model to similar discharges. Respective errors of the density and temperature at $\rho = 0.0, 0.5, 0.9, 1.0$ are given in Table V. Predictions obtained from the model without the auxiliary regressor.

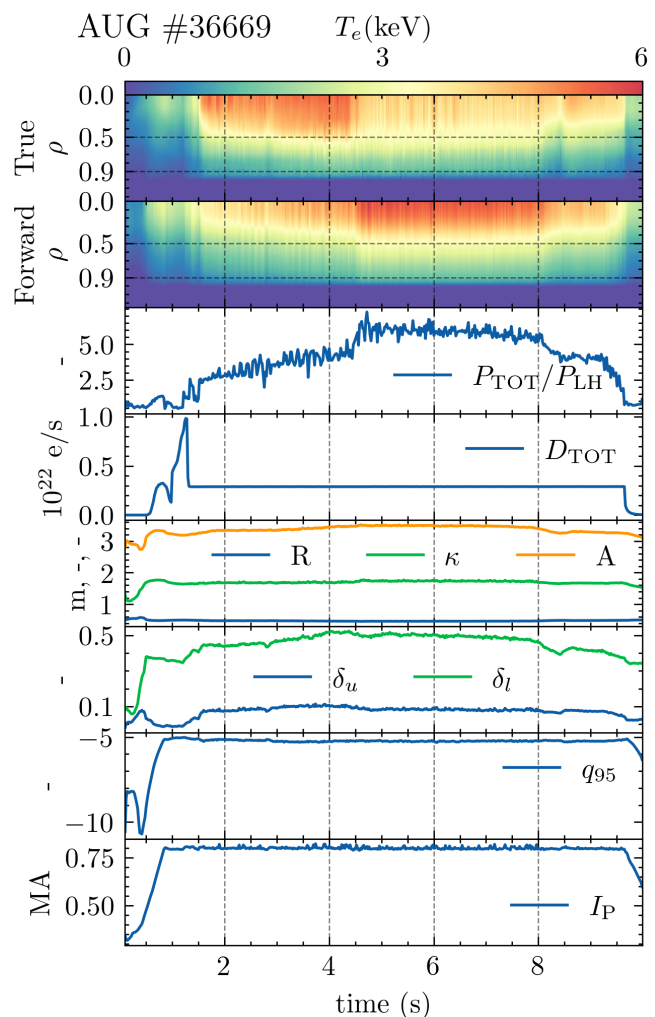


FIG. 11. AUG #36669, the top two plots are the true and forward model predicted temperature profiles and remaining plots are the corresponding actions for the pulse. The y-axis for the actions figures are the units, and their corresponding name is given within the plot. In AUG #36669, the device is configured to test actively cooled divertor plates, and accumulation of impurities lead to an increase in core radiated power¹⁵, dropping the core temperature. The actions supplied to the model do not sufficiently encode this information and the model does not predict the temperature decrease. It is likely that including additional actions that are more correlated with detached/attached plasmas would increase the resilience of the model, such as those used in⁵. Respective errors of the density and temperature at $\rho = 0.0, 0.5, 0.9, 1.0$ are given in Table V. Predictions obtained from the model without the auxiliary regressor.

- ⁹T. Takizuka and Y. R. Martin, “Power requirement for accessing the H-mode in ITER,” *Journal of Physics: Conference Series OPEN ACCESS* **123**, 12033 (2008).
- ¹⁰D. Ha and J. Schmidhuber, “World models,” (2018), 10.5281/ZENODO.1207631.
- ¹¹D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, “Learning latent dynamics for planning from pixels,” in *International Conference on Machine Learning* (2019) pp. 2555–2565.
- ¹²J. Brandstetter, D. Worrall, and M. Welling, “Message passing neural pde solvers,” (2023), arXiv:2202.03376 [cs.LG].
- ¹³D. P. Kingma and J. L. Ba, “Adam: A Method for Stochastic Optimization,” 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (2014).
- ¹⁴F. Matos, V. Menkovski, A. Pau, G. Marceca, F. Jenko, and the TCV Team, “Plasma confinement mode classification using a sequence-to-sequence neural network with attention,” *Nuclear Fusion* **61**, 046019 (2021).
- ¹⁵R. Neu, A. Kallenbach, M. Balden, V. Bobkov, J. W. Coenen, R. Drube, R. Dux, H. Greuner, A. Herrmann, J. Hobirk, H. Höhnle, K. Krieger, M. Kočan, P. Lang, T. Lunt, H. Maier, M. Mayer, H. W. Müller, S. Potzel, T. Pütterich, J. Rapp, V. Rohde, F. Rytter, P. A. Schneider, J. Schweinzer, M. Sertoli, J. Stober, W. Suttrop, K. Sugiyama, G. Van Rooij, and M. Wischmeier, “Overview on plasma operation with a full tungsten wall in ASDEX Upgrade,” *Journal of Nuclear Materials* **438**, S34–S41 (2013).
- ¹⁶M. Karl, M. Soelch, J. Bayer, and P. van der Smagt, “Deep variational bayes filters: Unsupervised learning of state space models from raw data,” (2017), arXiv:1605.06432 [stat.ML].
- ¹⁷P. Becker-Ehmck, J. Peters, and P. van der Smagt, “Switching linear dynamics for variational bayes filtering,” (2019),

arXiv:1905.12434 [stat.ML].

- ¹⁸B. Lusch, J. N. Kutz, and S. L. Brunton, “Deep learning for universal linear embeddings of nonlinear dynamics,” *Nature Communications* 2018 9:1 **9**, 1–10 (2018).

- ¹⁹A. Ryzhikov, M. Hushchyn, and D. Derkach, “Latent neural stochastic differential equations for change point detection,” (2022), arXiv:2208.10317 [cs.LG].

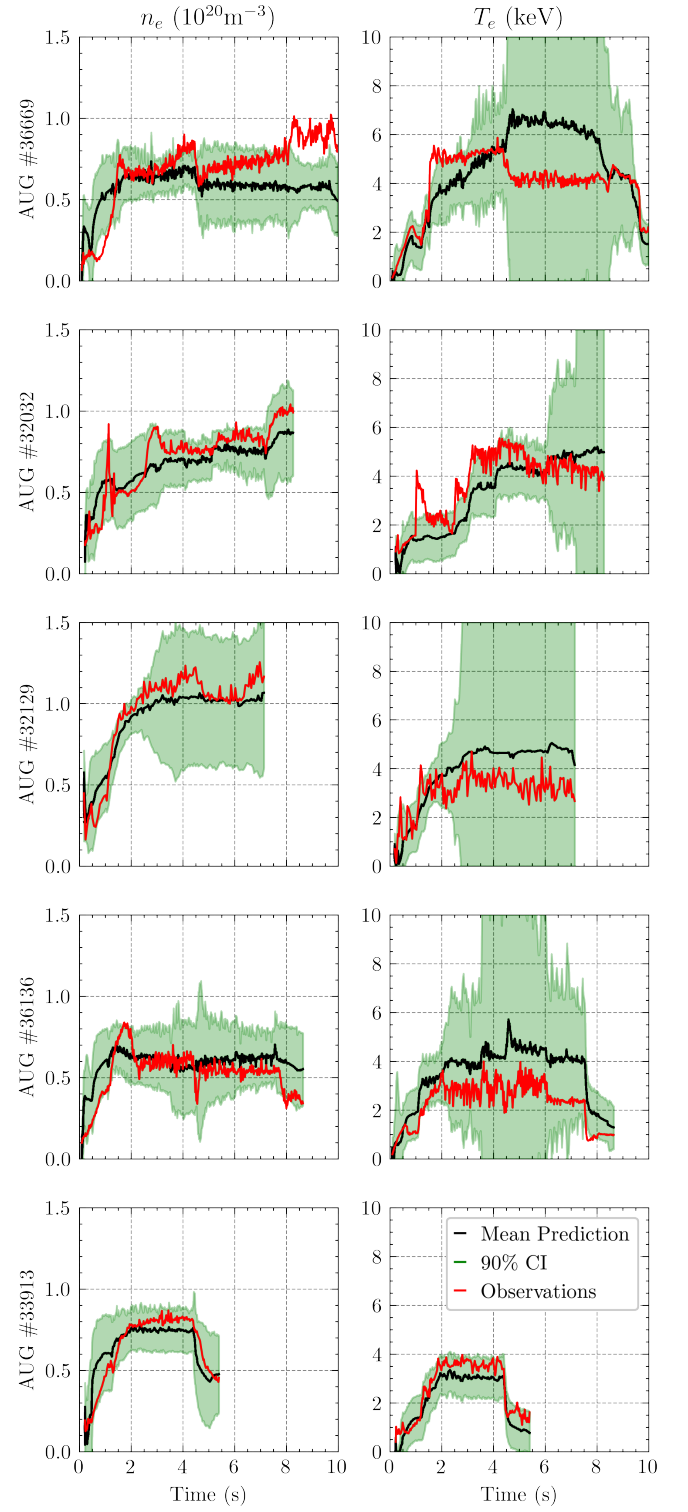


FIG. 12. Forward model predictions for core density and temperature ($\rho = 0.0$), with the density on the left and temperature on the right. The black lines depict mean predictions of 1000 samples generated using the forward model, whereas the green shaded regions represent the 90% confidence interval of the forward model distribution $p(s_{t+1}|s_t, a_t)$.

²⁰C. Balsells-Rodas, Y. Wang, and Y. Li, “On the identifiability of markov switching models,” 2305.15925 [cs, stat].

²¹M. Ilse, J. M. Tomczak, C. Louizos, and M. Welling, “Diva: Domain invariant variational autoencoders,” (2019), arXiv:1905.10427 [stat.ML].

²²R. Neu, K. B. Fournier, D. Schlögl, and J. Rice, “Observations of x-ray spectra from highly charged tungsten ions in tokamak plasmas,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **30**, 5057 (1997).