Data augmentation for disruption prediction via robust surrogate models

Katharina Rath[1,2]†, David Rügamer[1,5], Bernd Bischl[1], Udo von Toussaint[2], Cristina Rea[4], Andrew Maris[4], Robert Granetz[4] and Christopher G. Albert[3]

[1]Department of Statistics, Ludwig-Maximilians-Universität München

[2]Max-Planck-Institut für Plasmaphysik, Garching

[3]Fusion@OEAW, Institute of Theoretical and Computational Physics, Graz University of Technology

[4]Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, Massachusetts

[5]Institute of Statistics, RWTH Aachen University

(Received xx; revised xx; accepted xx)

The goal of this work is to generate large statistically representative datasets to train machine learning models for disruption prediction provided by data from few existing discharges. Such a comprehensive training database is important to achieve satisfying and reliable prediction results in artificial neural network classifiers. Here, we aim for a robust augmentation of the training database for multivariate time series data using Student-t process regression. We apply Student-t process regression in a state space formulation via Bayesian filtering to tackle challenges imposed by outliers and noise in the training data set and to reduce the computational complexity. Thus, the method can also be used if the time resolution is high. We use an uncorrelated model for each dimension and impose correlations afterwards via coloring transformations. We demonstrate the efficacy of our approach on plasma diagnostics data of three different disruption classes from the DIII-D tokamak. To evaluate if the distribution of the generated data is similar to the training data, we additionally perform statistical analyses using methods from time series analysis, descriptive statistics, and classic machine learning clustering algorithms.

1. Introduction

Disruptions pose serious challenges to the operation and design of tokamaks. Due to rapidly growing instabilities, thermal energy is rapidly lost during a disruption, the magnetic confinement of the plasma is destroyed, and energy is deposited into the confining vessel, potentially causing serious damages. Hence, to maintain a reliable fusion operation, disruption mitigation mechanisms should be triggered with sufficient warning time prior to the disruption. Recent advances on real-time disruption prediction have been made using machine learning (Rea & Granetz 2018; Kates-Harbeck *et al.* 2019; Rea *et al.* 2020, 2019; Aymerich *et al.* 2022; Pau *et al.* 2019). Disruption prediction is a challenging task for various reasons. One of them is the imbalanced data situation; for some disruption classes, only a few measurements are available, making it difficult to obtain robust results. This is challenging, especially when working with neural networks, as they require a large training data set in order to give satisfying results and to avoid overfitting (see e.g. Aggarwal (2018)). However, generating such an amount of training

† Email address for correspondence: katharina.rath@ipp.mpg.de

data from additional discharges is expensive and also potentially harmful for the reactor. Particularly with regard to future reactors such as ITER or SPARC, a sufficient data set won't be available at the time these reactors will start operating.

Data augmentation is one possibility to balance the training data set by creating rare disruption events and thereby improve the prediction performance of machine learning models. The aim of data augmentation is to produce an arbitrarily large amount of artificial samples that have the same statistical properties as the original small data set. Especially in the context of image classification, data augmentation is a widely used technique to improve the prediction accuracy and avoid overfitting (Shorten & Khoshgoftaar 2019). Commonly used methods are random transformation-based approaches, such as cropping or flipping. However, these methods are not expedient for the task at hand, as time-dependencies and the causal structure of physical signals are destroyed by such transformations (Iwana & Uchida 2021; Wen *et al.* 2021). More elaborate methods for multivariate time series generation using neural networks (Yoon *et al.* 2019) require substantially more samples per class than usually available for disruption prediction. Other advanced data augmentation methods are based on decomposition into trend, seasonal/periodic signal, and noise (Cleveland *et al.* 1990; Wen *et al.* 2019) or involve statistical modelling of the dynamics using, e.g., mixture autoregressive models (Kang *et al.* 2020).

Here, we tackle the above-mentioned challenges by relying on a non-parametric Bayesian approach to design the multivariate surrogate model based on Student-t process regression (Roth *et al.* 2017; Shah *et al.* 2014) to generate additional data. This model is closely related to the more commonly used Gaussian process regression (Williams & Rasmussen 1996). One drawback of Gaussian processes regression is the assumption of Gaussian noise, which is inaccurate due to outliers in the present application case. This results in unreliable uncertainty estimates. Therefore, our approach rather builds on the related Student-t processes that allow a heavy tailed noise distribution and give robust results even for noisy data corrupted by outliers.

Another challenge imposed by high-resolution time series data is the computational complexity of multivariate Gaussian or Student-t process regression of $O(N^3)$, where $N = DT$ is the number of training data points given by the product of dimensions D and time steps T of the multivariate time series. For typical values of $N > 1000$, traditional regression requires unbearable computing time. We instead use the state space formulation of a Student-t process as a linear time invariant stochastic differential equation, which can be solved using a corresponding filter and smoother (Solin & Särkkä 2015). In the case of a Gaussian process, the analogous approach is the well-known Kalman filter and RTS smoother (Särkkä 2013; Särkkä & Solin 2019). This ansatz reduces the computational complexity to $O(N)$, making it also suitable for high-resolution time series.

Here, we are working with a multi-output state space model to generate multivariate time series. We first assume that dimensions of the multivariate time series are not correlated. This is done to avoid the requirement of optimizing all hyperparameters at the same time, which is practically unfeasible due to the limited amount of available data. To still account for signal interdependencies, we then induce correlations and cross-correlations via coloring transformations in a post-processing step.

To balance the training data set, we use several local surrogate models to generate data coming from different disruption classes. From a small set – usually less than 10 discharges – of multivariate time series with D measurement signals coming from one disruption class with similar operating conditions, we estimate the posterior distribution. We then sample from the trained model in order to generate similar data that enlarges the

training database. To evaluate if the generated samples are from the same distribution as the training data, we use several methods from time series analysis, descriptive statistics, and clustering algorithms to show that generated and training samples are almost indistinguishable.

2. Methods

2.1. *Student-t processes*

Student-t processes (TPs) are a generalization of the widely used Gaussian processes (GPs) (Shah *et al.* 2014; Williams & Rasmussen 1996). TPs allow for a heavy tailed noise distribution (estimated by an additional hyperparameter $\nu > 2$) and therefore put less weight on outliers compared to GPs (Shah *et al.* 2014; Roth *et al.* 2017). This is illustrated in figure 1 for a test case of synthetic data corrupted by outliers. As in GP regression, we consider a set of N training observations $\mathcal{D} = \{(t_i, y_i)\}_{i=1}^{T}$ of scalar function values $y_i = f(t_i)$ plus measurement noise at training points t_i with $i = 0, 1, ..., T$ (in our case: time). We model these data points using a Student-t process with zero mean and covariance function $k(t, t')$,

$$f(t) \sim \mathcal{TP}(0, k(t, t'), \nu). \tag{2.1}$$

Similar to the GP, a kernel function $k(t, t')$ quantifies the covariance between values of f at times (t, t') and yields an $N \times N$ covariance matrix $K_{ij} = k(t_i, t_j)$ for the random vector of all observed y_i. Kernel hyperparameters determine further details, e.g., a length scale l quantifies how fast correlations vanish with increasing distance in t. The additional hyperparameter $\nu > 2$ corresponds to the degrees of freedom that specify the noise distribution. The predicted distribution of a scalar output $f(t_*)$ at test point t_* is given in closed form by

$$\mathbb{E}[f(t_*)] = \boldsymbol{k}_*^\top K_y^{-1} \boldsymbol{y}, \tag{2.2}$$

$$\mathbb{V}[f(t_*)] = \frac{\nu - 2 + \boldsymbol{y}^\top K_y^{-1} \boldsymbol{y}}{\nu - 2 + N} (k_{**} - \boldsymbol{k}_*^\top K_y^{-1} \boldsymbol{k}_*), \tag{2.3}$$

where $K_y = K + \sigma_n^2 I$ is the measurement noise parametrized by the noise variance σ_n^2. $\boldsymbol{k}_*$ is an $N-$dimensional vector with the ith entry being $k(t_*, t_i)$. $k_{**} = k(t_*, t_*)$ describes the covariance between training and test data and the variance at the test point t_*. In contrast to GP regression, the posterior variance $\mathbb{V}[f(t_*)]$ of the prediction explicitly depends on training observations by taking data variability into account and results in more reliable uncertainty estimates. An analogous expression to (2.3) is obtained for the covariance matrix between predictions at multiple t_* (Shah *et al.* 2014).

2.2. *State space formulation*

As in GP regression, the computational complexity increases with $O(N^3)$, as an inversion of the covariance matrix via Cholesky factorization is necessary to train TPs (Williams & Rasmussen 1996). This makes GP and also TP regression unfavorable for high-resolution time series data. However, as shown by (Solin & Särkkä 2015), the TP regression problem can be reformulated as an m-th order linear time invariant stochastic differential equation (SDE):

$$\frac{d\hat{\boldsymbol{f}}(t)}{dt} = F\hat{\boldsymbol{f}}(t) + L\boldsymbol{w}(t), \tag{2.4}$$

$$f(t_i) = H\hat{\boldsymbol{f}}(t_i), \tag{2.5}$$

where $\hat{\boldsymbol{f}}(t) = (f(t), \frac{df(t)}{dt}, ..., \frac{d^{m-1}f(t)}{dt^{m-1}})^\top$, the feedback matrix F and noise effect matrix L are derived from the underlying TP, $H = (1, 0, ..., 0)$ is the measurement or observation matrix, and $\boldsymbol{w}(t)$ is a vector of white noise processes with spectral density γQ, where γ is a scaling factor (Solin & Särkkä 2015).

To solve this SDE for discrete points in time by estimating the posterior distribution $p(\hat{\boldsymbol{y}}_{0:T}|\boldsymbol{y}_{1:T})$ of the latent state $\hat{\boldsymbol{y}}_{0:T}$ given noisy observations $\boldsymbol{y}_{1:T}$, we use the corresponding Student-t filter and smoother as outlined in Solin & Särkkä (2015). Here, the posterior is estimated by using marginal distributions: (1) filtering distribution $p(\hat{\boldsymbol{y}}_t|\boldsymbol{y}_{1:t})$ given by the update step in Algorithm 1, (2) prediction distribution $p(\hat{\boldsymbol{y}}_{t+k}|\boldsymbol{y}_{1:t})$ given by the prediction step in Algorithm 1 for k steps after the current time step t, and (3) smoothing distributions $p(\hat{\boldsymbol{y}}_t|\boldsymbol{y}_{1:T})$ for $t < T$ given by Algorithm 2 (Särkkä 2013). The initial distribution is determined by the prior state mean given by the measurements at $t = 0$ and prior state covariance P_0 given by the stationary covariance (Solin & Särkkä 2015). The augmented states df/dt that are not measured and noise are initialized with 0.

For example, the state space formulation of the Matérn 3/2 kernel is given by the following expressions for feedback, noise effect matrix, and spectral density (Särkkä & Solin 2019):

$$F = \begin{pmatrix} 0 & 1 & 0 \\ -\lambda^2 & -2\lambda & 0 \\ 0 & 0 & -\infty \end{pmatrix}, P_0 = \begin{pmatrix} \sigma^2 & 0 & 0 \\ 0 & \sigma^2\lambda^2 & 0 \\ 0 & 0 & \sigma_n^2 \end{pmatrix}, H = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}, L = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \tag{2.6}$$

where $\lambda = \sqrt{3}/l$. Hyperparameters l, σ^2, σ_n^2 and ν needed in the Student-t filter algorithm are estimated by minimizing the negative log-likelihood (Solin & Särkkä 2015). The log-likelihood is sequentially calculated using the Student-t filter (Algorithm 1). When the hyperparameters are optimized, the predictive distribution is first calculated via Algorithm 1 and then smoothed using Algorithm 2. In order to include the noise model with σ_n^2 corresponding to $K_y = K + \sigma_n^2 I$ in traditional TP regression, the SDE is directly augmented by the entangled noise model. As the model is not only augmented with the noise model, but also with the first derivative of the target function we want to predict, we can immediately infer $df(t)/dt$ from the given observations y.

Here, the task at hand concerns multivariate time series with multiple measurements n with D dimensions Y where the ith row is $\boldsymbol{y}_i = \boldsymbol{f}(t_i)$ at every time step t_i. To facilitate the training of the model, we consider an uncorrelated model, such that the random processes associated are not correlated. In traditional GP/TP regression, this corresponds to a multi-output model with a block-diagonal covariance matrix. The multi-output state space model to estimate $p(\hat{Y}_{0:T}|Y_{1:T})$ is built by stacking the univariate SDE models resulting in a block-diagonal structure for feedback and covariance matrices. Then, the dynamics of $\boldsymbol{y}_i$ are independent. We sample uncorrelated multivariate time series from this model and apply coloring transformations in a following post-processing step to account for correlations (section 2.4). Each dimension has its own set of hyperparameters in order to grasp dynamics that happen on different time scales. The measurement

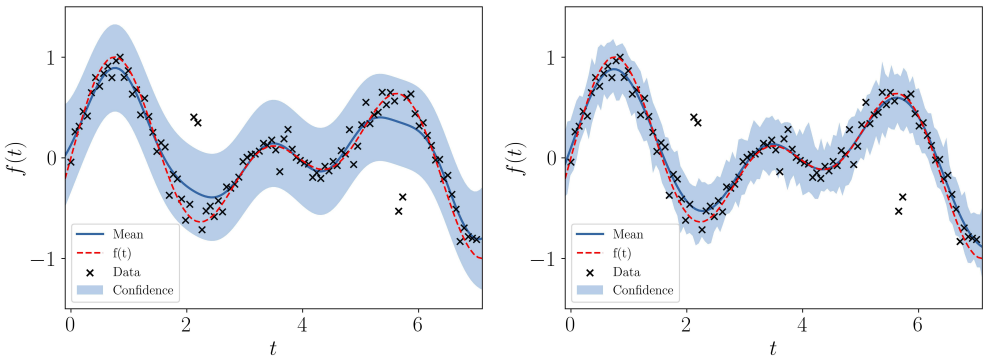

FIGURE 1. Predicted mean and 95% confidence band with (*a*) Gaussian process and (*b*) Student-t process trained on $N = 100$ training data points following $f(t) = \sin(2t)\cos(0.4t)$ corrupted by Gaussian noise $0.1\mathcal{N}(0, 1)$, with several outliers.

covariance matrix R (Algorithm 1) is estimated using the covariance of n measurements for each dimension at every time step.

2.3. *Student-t sampler*

To sample from the estimated posterior distribution, we employ a Student-t sampler, which is a modified version of the sampling technique presented by Durbin & Koopman (2002). First, we draw a t-distributed random sequence $\hat{X}_{0:T} = \hat{\boldsymbol{x}}_{i,0:T}$ from the prior estimated by the trained Student-t model. These sequences are initialized by $\mathcal{T}(0, P_0)$ and then filtered using Algorithm 1 and smoothed via Algorithm 2, which yields $\mathbb{E}(\hat{X}_{0:T}|Y^+_{1:T})$ where $Y^+_{1:T} = H\hat{X}_{0:T}$, with the stacked measurement matrix $H = (I, 0, 0)$ that extracts only the first component of $\hat{\boldsymbol{x}}_t$ in every time-step t. $Y^+_{1:T}$ are data associated with the filtered and smoothed sequence $\hat{X}_{0:T}$ given by (A 2). Finally, to obtain a random sequence $\bar{Y}_{0:T} = \bar{\boldsymbol{y}}_{i,0:T} \sim p(\hat{Y}_{0:T}|Y_{1:T})$, we combine

$$\bar{Y}_{1:T} = H(\mathbb{E}(\hat{Y}_{0:T}|Y_{1:T}) + \hat{X}_{0:T} - \mathbb{E}(\hat{X}_{0:T}|Y^+_{1:T})), \tag{2.7}$$

where H extracts the first component of $\hat{\boldsymbol{y}}_t$ in every time-step t. This procedure gives a D-dimensional multivariate time series for T time steps.

2.4. *Post-processing*

Given the trained model, we sample data $\bar{Y}_{1:T}$ from the estimated posterior, where rows are dimensions $\bar{y}_i$ and columns are time steps. $\bar{Y}_{1:T}$ can be split into a mean given by the smoothing distribution and deviations due to the sampling. Correlations between dimensions D of the generated data are not reproduced correctly with the uncorrelated model. To inscribe the average covariance Σ over all samples that we empirically observe in the training data $Y_{1:T}$ into the generated data $\bar{Y}_{1:T}$ with covariance matrix $\bar{\Sigma}$ with small but larger than zero off-diagonal elements, we first perform a ZCA whitening (also known as Mahalanobis) transformation (see e.g. Kessy *et al.* (2018)):

$$Z = \bar{\Sigma}^{-1/2}\bar{Y}. \tag{2.8}$$

The transformed data Z have a diagonal covariance matrix Λ_Z, with unit variances on the diagonal. We then color the generated data via a coloring transformation (Kessy *et al.* 2018):

$$\tilde{Y} = \Sigma^{1/2} Z = \Sigma^{1/2} \bar{\Sigma}^{-1/2} \bar{Y}, \tag{2.9}$$

obtaining data $\tilde{Y}$, which now have the same (temporally local) covariance as the training data Y.

Another possibility is to directly take the distribution of the training data covariance matrix Σ over samples into account by using samples from a corresponding multivariate Gaussian distribution as data covariance matrices. This generates variation in the covariance of the generated data, especially if there are local differences between the samples. However, on average for a large enough sample size, we recover the training data covariance matrix Σ.

To also take time-lagged correlations into account, we must adjust not only covariances but also cross-covariances in our generated data. Therefore, we use the cross-covariance matrix given by

$$\bar{\Sigma}_{c,rs}(t_1, t_2) = \mathbb{E}[(\bar{y}_{r,t_1} - \mu_{r,t_1})(\bar{y}_{s,t_2} - \mu_{s,t_2})], \tag{2.10}$$

where the expected value $\mathbb{E}[\cdot]$ is estimated by averaging over all combinations of lags $t_1 - t_2$ in addition to the sample mean. Here, $\mu_{i,t}$ is the expected value of $\bar{y}_{i,t}$. To decorrelate and color the data in the way described above, we formally use a global covariance matrix Σ_g of size $DT \times DT$ involving correlations both over time and across dimensions of the multivariate time-series. The global covariance matrix is a periodic block matrix given by

$$\Sigma_{g,(t_1 D+r)(t_2 D+s)} = \Sigma_{c,rs}(t_1, t_2) \tag{2.11}$$

for the cross-covariance Σ_c with lag $t_1 - t_2$.

$$\tilde{Y} = \Sigma_g^{1/2} Z = \Sigma_g^{1/2} \bar{\Sigma_g}^{-1/2} \bar{Y}. \tag{2.12}$$

This incorporates the empirical cross-covariance for all time lags and between all dimensions D of the generated data.

3. Evaluation of generated data

As the generated data serve as augmented training data for later analyses, statistical properties of the original training data should be reflected in the generated data. Therefore, we perform statistical tests to check if training and generated share key statistical properties.

3.1. *Distribution and Wasserstein distance*

To measure the distance between the distribution of the training and the generated data, we use the Wasserstein-1 metric (Villani 2008):

$$W_1(P, V) = \inf_{\gamma \in \Gamma(P,V)} \int_{\mathbb{R}\times\mathbb{R}} |x - y| d\gamma(x, y), \tag{3.1}$$

where $\Gamma(P, V)$ denotes the set of all probability distributions on $\mathbb{R} \times \mathbb{R}$, with P, V being its marginals. The minimizer γ of (3.1) denotes the optimal transport plan to transport P to V. We compare each signal separately and average the corresponding Wasserstein distances. Although the problem concerns time series data, we discard all time information and only consider the global distribution of the data due to the small amount of available training data samples.

3.2. *Maximum mean discrepancy two-sample test*

In addition to the Wasserstein distance, we perform the kernel two-sample test (Gretton *et al.* 2012) for each signal (again discarding time information). The null hypothesis we want to test is that both n training data $y_{i,1:N}$ and m generated data samples $\tilde{y}_{i,1:T}$ follow the same distribution P. We use the maximum mean discrepancy (MMD) test statistic via a kernel g:

$$\mathrm{MMD}^2 = \frac{1}{n(n-1)} \sum_{i,j=1}^{n} g(y_{i,1:T}, y_{j,1:T}) + \frac{1}{m(m-1)} \sum_{i,j=1}^{m} g(\tilde{y}_{i,1:T}, \tilde{y}_{i,1:T}) - \frac{2}{nm} \sum_{i=1}^{n} \sum_{j=1}^{m} g(y_{i,1:T}, \tilde{y}_{i,1:T}), \tag{3.2}$$

where $g(x, y) = \exp(-||x - y||^2/(2\sigma^2))$ with $\sigma = \mathrm{Median}(|\Upsilon_i - \Upsilon_j|)/2$ and Υ is the combined sample of $y_{i,1:T}$ and $\tilde{y}_{i,1:T}$. To estimate a threshold for the acceptance of the null hypothesis for a given confidence level, bootstrapping is performed via mixing samples $y_{i,1:T}$ and $\tilde{y}_{i,1:T}$, which generates a distribution with 10000 samples that satisfies the null hypothesis. Finally, we can estimate a p-value for the MMD of the generated data distributions.

3.3. *Auto- and cross-correlation*

To evaluate if the generated data reflect the temporal dependencies of the training data, we calculate auto- and cross-correlations for training and generated data by normalizing the cross-covariance Σ_c in (2.10) by $1/(\sigma_{r,t_1}\sigma_{s,t_2})$. Here, $\sigma_{s,t}$ is the standard deviation of $\tilde{y}_{s,t}$. If $r = s$, this diagnostic becomes the auto-correlation – see, e.g., Park (2017). For $t_1 = t_2$, the local correlation matrix follows. We evaluate the mean squared error to the auto- and cross-correlation of the training data. Evidently, the global coloring transformation (2.10) produces a perfect match in this diagnostic.

3.4. *Power spectral density*

All frequencies that are present in the training data set should also appear in the generated data. This can be evaluated using power spectral density (PSD), which provides an estimate of power distribution across the frequency of a signal. We evaluate the mean squared error between the PSD of the training data and generated data.

3.5. *Embedding via principal component analysis*

We apply 2-D principal component analysis (PCA) on training data with flattened temporal dimension and project the generated data onto the first two principal components of the training data to evaluate the embedding and visualize if both training and generated data lie on the same submanifold.

The distance between the embedded distributions of training and generated data is measured by using the sliced Wasserstein distance that takes advantage of the very efficient calculation of 1-D Wasserstein distances (Bonneel *et al.* 2015; Flamary *et al.* 2021). The multivariate distribution is sliced and randomly projected on a one-dimensional subspace, and the corresponding 1-D Wasserstein distances are averaged to obtain an estimation for the multivariate distribution. With an increasing number of projections, the sliced Wasserstein distance converges. Here, we use 10^3 projections to estimate the distance W_{emb} between the embedded distributions.

3.6. *Multivariate functional principal component analysis*

For the evaluation of the correctly represented temporal evolution of the generated data, we apply multivariate functional principal component analysis (mfPCA) on the training data and project the generated data onto the eigenbasis of the training data (Happ & Greven 2018). Then, we reconstruct both training and generated data with the same eigenbasis and evaluate the variance of the residuals.

3.7. *Dynamic time warping*

For time series comparison, dynamic time warping (DTW) is widely used to measure the similarity between two temporal sequences $y_{i,1:T}$ and $\tilde{y}_{j,1:T}$ (Berndt & Clifford 1994). This metric is formulated as an optimization problem:

$$\mathrm{DTW}(y_{i,1:T}, \tilde{y}_{j,1:T}) = \min_{\gamma} \sqrt{\sum_{(i,j)\in\gamma} d(y_i, \tilde{y}_j)^2}, \tag{3.3}$$

where γ is the alignment path such that the Euclidean distance between $y_{i,1:T}$ and $\tilde{y}_{j,1:T}$ is minimal. Hence, DTW gives the distance between two time series with the best temporal alignment. We compare each training data sample with each generated data sample and use the mean to compare different post-processing methods.

3.8. *Self-organizing maps on time series*

Finally, we apply time series clustering based on DTW self-organizing maps (SOMs) on both the training and generated data (Vettigli 2018). If the generated data are a potentially useful extension of the training data, the clustering should show similar results. Therefore, we compute a clustering model on the training data and use the trained model to predict cluster labels of both the training and generated data. From the predicted labels, we evaluate the F1 score (Murphy 2022) with the ground truth.

4. Numerical experiments

We evaluate the performance of the proposed model using disruption data from several discharges from the DIII-D tokamak taken from the 2016 experimental campaign. These disruptions were already included in previously published papers on data-driven applications in fusion (Montes *et al.* 2021).

We cluster the available data sets depending on the similarity of the conditions and on the occurring instability. Here, we use the model to augment five signals of the training data set (referred to as β_n, normalized internal inductance li, plasma elongation κ, safety factor q_{95}, Greenwald fraction n/n_G) for different disruptions: (1) disruption due to magnetohydrodynamic (MHD) instability (Shots 166463, 166464, 166465, 166466, 166468, 166469), (2) disruption due to MHD instability induced by resonant magnetic perturbations (RMP) applied to control Edge Localized Modes (ELMs) (Shots 166452, 166454, 166456, 166457, 166460), and (3) density accumulation due to detachment (Shots 166933, 166934, 166937). For each disruption class, the model is trained on these few available training samples. The choice of signals is influenced by the use case of augmenting the training database for a neural network for disruption prediction, but in general, the method is extendable to any number and any kind of signals.

Following the flow shown in figure 2, preprocessing is performed on the training data. As we are primarily interested in the behavior close to a disruption, we align the samples according to their end time and only consider the stable flat-top phase. Additionally, all

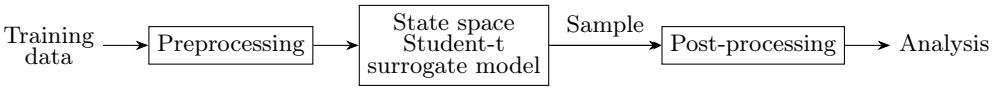

FIGURE 2. Data processing flow.

data are rescaled via min-max scaling to a range between $[-0.5, 0.5]$. This stabilizes the optimization of the hyperparameters in the Student-t filter algorithm, as the input to the optimizer is of order 1. Missing data points are interpolated linearly. All discharges are sampled every 25 ms. Then, we set up the state space Student-t surrogate model. In all experiments, a Matérn 3/2 kernel as in (2.6) is used. We train the surrogate model by optimizing its hyperparameters by minimizing the negative log-likelihood using the Scipy implementation of L-BFGS-B (Virtanen *et al.* 2020), and resulting values for all experiments can be found in the appendix B, table 3. Each signal has its own set of hyperparameters in order to be able to handle dynamics that happen on different time scales. Subsequently, we apply the Student-t filter and smoother (Algorithm 1 and 2) with optimized hyperparameters to our data. From the estimated distribution, we draw 1000 samples from the posterior using the Student-t sampler and perform the coloring transformations in the post-processing. Finally, we evaluate the generated data sets by using the defined metrics. In general, the generation of the time series samples is of $O(N)$, but some of the used metrics to evaluate the generate data are not. Therefore, we limited the number of samples in the given analysis to 1000.

5. Results and analysis

For each disruption class, we draw 1000 samples from the posterior estimated by the trained model and compare four available post-processing methods: (1) uncorrelated model (here, no post-processing is performed), (2) coloring transformation with empirical covariance matrix, (3) coloring transformation with sampled covariance matrix, and (4) coloring transformation with empirical cross-covariance matrix to account for lagged correlations. The results for test cases 2 and 3 are shown in the appendix in C.1 and C.2.

In figure 3, a visual comparison is given between training data and generated data for the coloring transformation with empirical cross-covariance matrix, together with the estimated mean and 95% confidence intervals for the disruption data from DIII-D with disruption due to MHD instability (test case 1). The model is able to capture the general trend given by the training data and can also reproduce outliers. In general, the generated data fit the distribution of the training data.

We continue with a thorough statistical analysis, which allows a ranking of the different post-processing methods following the metrics outlined in section 3. The results are given in table 1 for disruption data from DIII-D with disruption due to MHD instability. Other experiments give similar results, as indicated in the appendix in C.1, table 4, and C.2, table 6 for test cases (2) and (3), respectively.

To put the calculated numbers into context, we evaluate the Wasserstein distance between the different training data sets originating from different disruption classes. We found $W_1 = 0.16 \pm 0.02$ between test case (1) and (2); $W_1 = 0.20 \pm 0.08$ between test case (1) and (3); and $W_1 = 0.25 \pm 0.15$ between test case (2) and (3). Additionally, the 2-D PCA embedding of the training data together with the estimation of the 2-D sliced Wasserstein distance is estimated. We observe $W_{\text{emb}} = 0.277 \pm 0.003$ between test case (1) and (2); $W_{\text{emb}} = 0.402 \pm 0.005$ between test case (1) and (3); and $W_{\text{emb}} = 0.445 \pm 0.005$ between test case (2) and (3). The achieved Wasserstein distances between training and generated data for one disruption class are typically more than five times smaller in all

post-processing methods as given in tables 1, 4 and 6. This is promising, as it means that the augmented data are much more similar to their proper class than disruption classes are to each other in this measure. Besides the Wasserstein distance, DTW is also difficult to interpret without context. Again, we calculate the metric between the training sets and obtain DTW $= 1.3 \pm 0.7$ between test case (1) and (2); DTW $= 2.7 \pm 1.5$ between test case (1) and (3); and DTW $= 2.4 \pm 1.3$ between test case (2) and (3). The distances between the generated and training data for each class lie sufficiently below the distances between different training data sets. The training data were also reconstructed using the multivariate functional PCA with 5 components. We observe the following reconstruction mean squared errors for test case (1) 0.006, (2) 0.003, and (3) 0.008. We use the first five eigenfunctions of the training data as a basis to project the generated data of each test case. However, the reconstruction error of the generated data with included correlations in the post-processing is still in the same order.

Figure 4 displays the kernel density of the 2-D PCA embedding of the generated data in the eigenspace of the training data. All four methods generate data that lie on the same submanifold as the training data. However, when cross-covariances are included, the shape of the training data is better reproduced. The number of available training data samples is very limited, as we are working with manually labeled disruptive data from DIII-D. Therefore, the results here only give an idea if the features apparent in the training data are also apparent in the generated data.

The superiority of the post-processing with the empirical cross-covariance is also apparent in figure 5, where the auto- (on the diagonal) and cross-covariance for all estimated signals is shown. As we are inscribing the empirical cross-covariance into the uncorrelated generated data from the model, the cross-covariance fits exactly, and the cross-covariances lie on top of each other. When using either the empirical covariance or the sample covariance, only the cross-covariance at lag 0 matches the cross-covariance of the training data. Both post-processing methods give on average the same cross-covariance for 1000 generated samples. Additionally, the difference in covariance at lag 0 is shown in figure 6.

Finally, we use SOMs for time series clustering to evaluate if the label prediction works similarly well for the generated data. Here, we only use three classes, as the training data look quite similar for different signals. The results for three different experiments are given in table 2. Between the four post-processing methods, no significant difference is evident. The clustering algorithm performs as well on all methods as on the original training data.

6. Conclusion and Outlook

We applied Student-t process regression in a state space formulation to introduce robust data augmentation for multivariate time series. The state space formulation reduces the computational complexity and is thus suitable for high-resolution time series. We used the model to learn the distribution of time series coming from a given disruption class. From the estimated posterior, time series were generated to augment the training database. To evaluate if the original and generated data share key statistical properties, multiple statistical analyses and classic machine learning clustering algorithms have been carried out. We found that within the scope of the used metrics, the generated time series resemble the training data to a sufficient extent.

When the method is applied to augment the training database for the neural network disruption predictor, a thorough analysis of the existing (labeled) training database is necessary to decide which disruption classes are not available in sufficient quantity. For

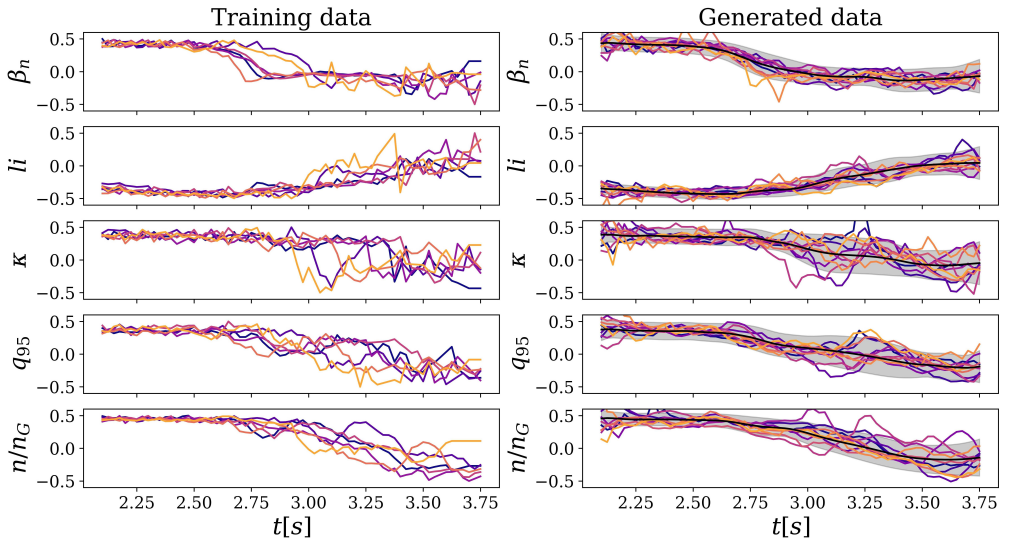

FIGURE 3. (*a*) Training data and (*b*) 10 generated data sets from the state space Student-t surrogate model together with the estimated mean (black solid line) and 95% confidence (grey shaded region) for the disruption data from DIII-D, with disruption due to MHD instability (test case 1). Different colors correspond to different shots of training data and different samples of the generated data, respectively.

Metric	uncorrelated	emp. cov	emp. crosscov	sample cov
W_1	0.035 ± 0.013	0.035 ± 0.011	0.038 ± 0.012	$\mathbf{0.036 \pm 0.01}$
MMD p-value	0.68 ± 0.35	0.82 ± 0.18	$\mathbf{0.92 \pm 0.06}$	0.85 ± 0.13
MSE ρ_{rs}	0.019 ± 0.009	0.018 ± 0.009	$\mathbf{0.0011 \pm 0.0004}$	0.017 ± 0.009
W_{emb}	0.078 ± 0.004	0.084 ± 0.0006	$\mathbf{0.08 \pm 0.0007}$	0.083 ± 0.006
MSE PSD $[10^{-6}]$	5 ± 4	4 ± 3	$\mathbf{3 \pm 3}$	4 ± 3
DTW	0.8 ± 0.5	0.8 ± 0.4	$\mathbf{0.7 \pm 0.3}$	0.85 ± 0.4
MSE mfPCA	0.137	0.015	**0.011**	0.022

TABLE 1. Post-processing method comparison for disruption data from DIII-D, with disruption due to MHD instability (test case 1). Mean and standard deviation over five dimensions and $N = 1000$ samples generated from the trained model for statistical metrics described in section 3. Best values are highlighted in bold.

Train	Test	training	uncorrelated	emp. cov	emp. crosscov	sample cov
original	generated	0.75	0.74	0.75	0.74	0.78
generated	original	0.88	0.86	0.90	0.90	0.89
mix	mix	0.89	0.89	0.89	0.89	0.89

TABLE 2. $F1$ score for DTW SOM clustering of different post-processing methods for disruption data from DIII-D, with disruption due to MHD instability (test case 1).

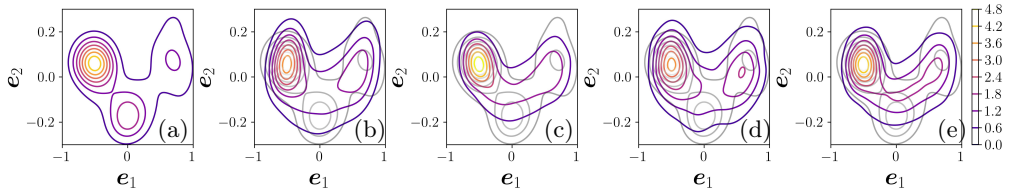

FIGURE 4. Kernel density of the 2-D PCA embedding of the (*a*) training data and generated data via (*b*) cross-covariance, (*c*) empirical covariance, (*d*) sampled covariance and (*e*) uncorrelated model for disruption data from DIII-D, with disruption due to MHD instability (test case 1). The embedded training data are shown in gray in all plots. The color scale representing the density is the same in all plots.

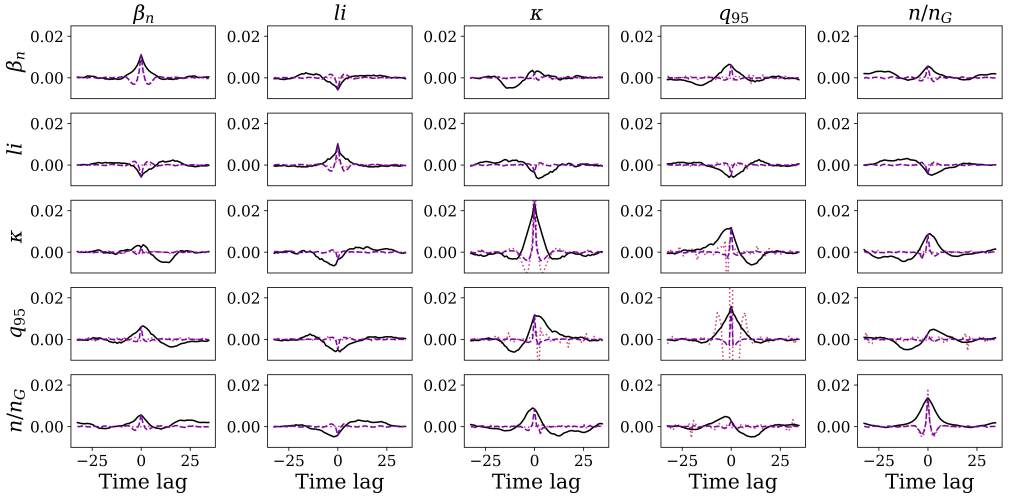

FIGURE 5. Comparison of the cross-covariance in the training and generated data with cross-covariance (solid lines on top of each other, numerical error of order 10^{-16}), covariance or sampled covariance post-processing (dashed lines), and uncorrelated model (dotted line) for disruption data from DIII-D, with disruption due to MHD instability (test case 1)

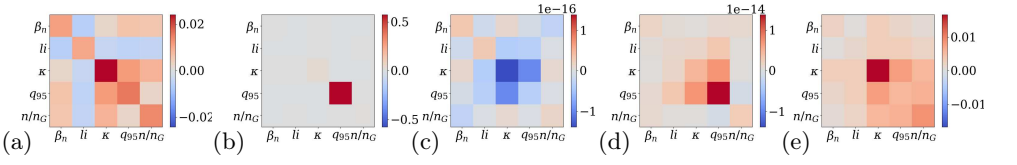

FIGURE 6. Comparison of the covariance of training data *(a)* and the difference from the generated data *(b)* with uncorrelated model, *(c)* empirical covariance, *(d)* cross-covariance, and *(d)* sampled covariance post-processing for disruption data from DIII-D, with disruption due to MHD instability (test case 1). Note the different scaling in the color scale.

each of those classes, we will train the surrogate model and then be able to generate data to balance the data set. Subsequently, the performance of the neural network trained with the augmented training database will be evaluated. Due to the broad range of evaluation metrics, we are optimistic that the generated data will improve and robustify the performance.

Another perspective regards disruption prediction of future devices, where little data will be available to train machine learning-based approaches. In this case, the surrogate

model could be used and updated, as more data are being collected and can therefore update machine learning-driven models.

To improve the proposed method, the integration of correlations and cross-correlations on the level of a multivariate surrogate model instead of the coloring in post-processing will be investigated in future work. Another possible extension of the current method could also take spatial information of profiles into account (Wilkinson *et al.* 2020).

However, the approach developed here is sufficiently generic to be used for data augmentation in a broad range of applications, e.g., time series in climate research.

Acknowledgements

The present contribution is supported by the Helmholtz Association of German Research Centers under the joint research school HIDSS-0006 "Munich School for Data Science - MUDS" (KR) and the MIT-Germany Lockheed Martin Seed Fund (KR, CR, AM, UvT). 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 (CA). This work was supported by the Federal Ministry of Education and Research (BMBF) of Germany by Grant No. 01IS18036A (DR, BB). This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE- SC0014264, and DE-FC02-04ER54698.

Disclaimer. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Appendix A. Algorithm

Algorithm 1 Multivariate Student-t filter (Solin & Särkkä 2015)

Init:

$$\hat{\boldsymbol{y}}_{0|0} = \boldsymbol{y}_0, \quad P_{0|0} = P_0, \quad \boldsymbol{\nu}_0 = \boldsymbol{\nu}, \quad \gamma_0 = I_D \tag{A 1}$$

for $t = 1, 2, ..., T$ **do**

Filter prediction:

$$\hat{\boldsymbol{y}}_{t|t-1} = A_{t-1}\hat{\boldsymbol{y}}_{t-1} \tag{A 2}$$

$$P_{t|t-1} = A_{t-1}P_{t-1}A_{t-1}^\top + \gamma_{t-1}Q_{t-1}, \tag{A 3}$$

where $A_t = \exp(F\Delta t)$ and $Q_t = P_0 - A_t P_0 A_t^\top$.

Filter update (if measurement $\boldsymbol{y}_t$ with mean $\bar{\boldsymbol{y}}_t$ is available):

$$\boldsymbol{v}_t = \bar{\boldsymbol{y}}_t - H_t\hat{\boldsymbol{y}}_t \tag{A 4}$$

$$S_t = H_t P_{t|t-1} H_t^\top + R \tag{A 5}$$

$$\gamma_t = \frac{\gamma_{t-1}}{\nu_t - 2}(\nu_{t-1} - 2 + \boldsymbol{v}_t S_t^{-1} \boldsymbol{v}_t) \tag{A 6}$$

$$K_t = P_{t|t-1} H_t^\top S_t^{-1} \tag{A 7}$$

$$\hat{\boldsymbol{y}}_{t|t} = \hat{\boldsymbol{y}}_{t|t-1} + K_t \boldsymbol{v}_t \tag{A 8}$$

$$P_{t|t} = \frac{\gamma_t}{\gamma_{t-1}}(P_{t|t-1} - K_t S_t K_t^\top) \tag{A 9}$$

end for

Algorithm 2 Multivariate Student-t smoother (Solin & Särkkä 2015)

Init:

$$\hat{\boldsymbol{y}}_T = \hat{\boldsymbol{y}}_{T|t}, \quad P_T = P_{T|t} \tag{A 10}$$

for $t = T-1, T-2, ..., 1$ **do**

Smoother prediction:

$$\hat{\boldsymbol{y}}_{t+1|t} = A_t \hat{\boldsymbol{y}}_{t|t} \tag{A 11}$$

$$P_{t+1|t} = A_t P_{t|t} A_t^\top + \gamma_t Q_t \tag{A 12}$$

Smoother update:

$$G_t = P_{t|t} A_t^\top P_{t+1|t}^{-1} \tag{A 13}$$

$$\hat{\boldsymbol{y}}_{t|T} = \hat{\boldsymbol{y}}_{t|t} + G_t(\hat{\boldsymbol{y}}_{t+1|T} - \hat{\boldsymbol{y}}_{t+1|t}) \tag{A 14}$$

$$P_{t|T} = \frac{\gamma_T}{\gamma_t}(P_{t|t} - G_t P_{t+1|T} G_t^\top) + G_t P_{t+1|T} G_t^\top \tag{A 15}$$

end for

Appendix B. Hyperparameters

For the different test cases, we used the hyperparameters given in table 3.

Test case	hyp	β_n	li	κ	q_{95}	n/n_G
(1)	ν	2.19	2.58	2.15	2.21	2.1
	σ_n^2	0.024	0.01	0.056	0.029	0.02
	σ^2	1.74	1.76	1.96	1.60	1.60
	l	19.6	28.3	20.3	20.1	15.7
(2)	ν	3.4	2.57	2.36	2.49	2.7
	σ_n^2	0.023	0.032	0.036	0.033	0.022
	σ^2	1.65	0.53	1.36	1.87	1.91
	l	17.7	19.8	9.63	19.1	16.8
(3)	ν	2.14	2.12	2.01	2.71	2.55
	σ_n^2	0.163	0.044	0.493	0.016	0.011
	σ^2	0.62	1.73	1.24	1.21	0.58
	l	17.6	11.5	4.5	12.8	10.0

TABLE 3. Optimized hyperparameters for the state space Student-t surrogate model for all test cases.

Metric	uncorrelated	emp. cov	emp. crosscov	sample cov
W_1	0.027 ± 0.013	0.02 ± 0.006	0.022 ± 0.007	$\mathbf{0.02 \pm 0.006}$
MMD p-value	0.617 ± 0.335	0.869 ± 0.153	$\mathbf{0.885 \pm 0.107}$	0.876 ± 0.15
MSE ρ_{rs}	0.013 ± 0.017	0.013 ± 0.017	$\mathbf{0.005 \pm 0.005}$	0.014 ± 0.017
W_{emb}	0.0458 ± 0.0003	0.0496 ± 0.0004	$\mathbf{0.0466 \pm 0.0004}$	0.0539 ± 0.0004
MSE PSD $[10^{-6}]$	7 ± 9	3 ± 4	$\mathbf{1 \pm 2}$	3 ± 4
DTW	0.86 ± 0.37	0.78 ± 0.32	$\mathbf{0.65 \pm 0.31}$	0.83 ± 0.39
MSE mfPCA	0.011	0.009	**0.007**	0.011

TABLE 4. Post-processing method comparison for disruption data from DIII-D with disruption due to MHD instability (test case 2). Mean and standard deviation over five dimensions and $N = 1000$ samples generated from the trained model for statistical metrics described in section 3. Best values are highlighted in bold.

Appendix C. Results for other test cases

In the following sections, the results for test cases 2 and 3 are presented.

C.1. *Test case 2: disruption due to MHD instability during RMP ELM control*

A visual comparison of the training and the generated data for test case 2 is shown in figure 7. Here, the disruption occurs due to MHD instability induced by resonant magnetic perturbations (RMP) applied to control Edge Localized Modes (ELMs) (Shots 166452, 166454, 166456, 166457, 166460). The results of the statistical analysis are given in table 4 and are in the same order as for test case 1. Figure 8 displays the kernel density of 2-D PCA embedding of the generated data. Again, the results show that the generated data lives on the same submanifold for all four post-processing methods. Figures 9 and 10 show the cross-covariance and the covariance of the training and generated data. In table 5, the $F1$ score for dynamic time warping SOM clustering is given.

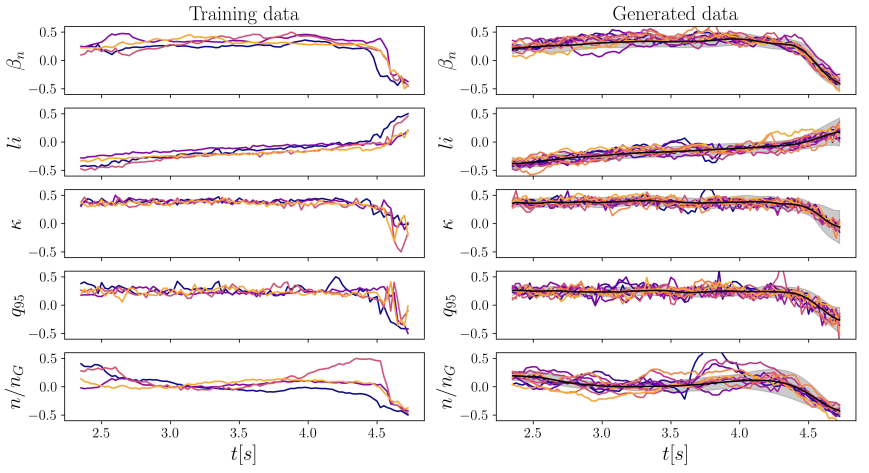

FIGURE 7. (a) Training data and (b) 10 generated data sets from the state space Student-t surrogate model together with the estimated mean (black solid line) and 95% confidence (grey shaded region) for the disruption data from DIII-D with disruption due to MHD instability (test case 2). Different colors correspond to different shots of training data and different samples of the generated data, respectively.

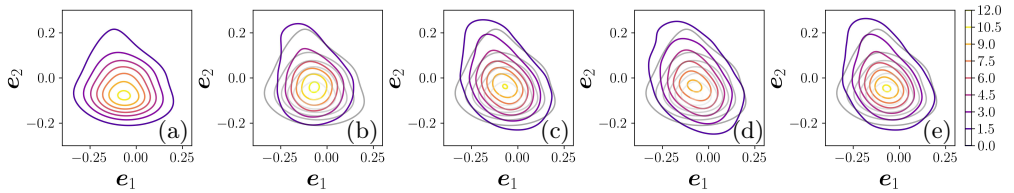

FIGURE 8. Kernel density of 2-D PCA embedding of (a) training data and generated data via (b) cross-covariance, (c) empirical covariance, (d) sampled covariance and (e) uncorrelated model for disruption data from DIII-D with disruption due to MHD instability (test case 2). The embedded training data are shown in gray in all plots. The colorscale representing the density is the same in all plots.

Train	Test	training	uncorrelated	emp. cov	emp. crosscov	sample cov
original	generated	1.0	1.0	1.0	0.94	1.0
generated	original	1.0	0.91	1.0	1.0	1.0
mix	mix	1.0	1.0	1.0	0.96	1.0

TABLE 5. $F1$ score for dynamic time warping SOM clustering of different post-processing methods for disruption data from DIII-D with disruption due to MHD instability (test case 2).

C.2. *Test case 3: Density accumulation*

For the third test case with a disruption occurring due to density accumulation (Shots 166933, 166934, 166937), the visual comparison is given in figure 11 followed by the results of the statistical analysis in table 6. The embedding is shown in figure 12, and the cross-covariance and covariance is displayed in figures 13 and 14, respectively. The

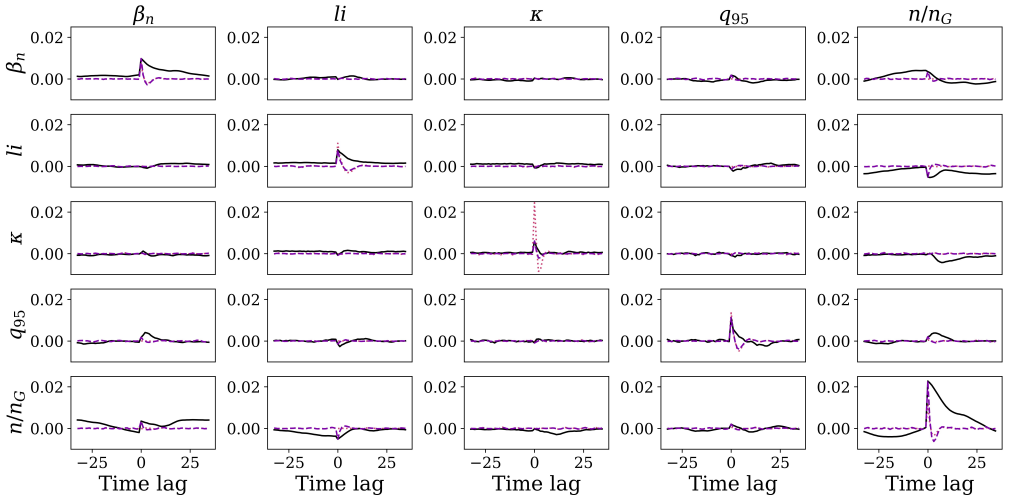

FIGURE 9. Comparison of cross-covariance of training data and generated data with cross-covariance (solid lines on top of each other, numerical error of order 10^{-16}), covariance or sampled covariance (dashed lines) post-processing and uncorrelated model (dotted line) for disruption data from DIII-D with disruption due to MHD instability (test case 2).

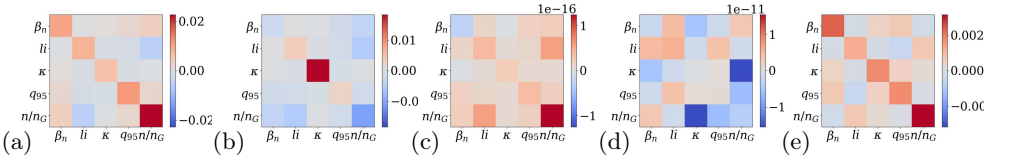

FIGURE 10. Comparison of covariance of training data *(a)* and difference of generated data *(b)* with uncorrelated model, *(c)* empirical covariance, *(d)* cross-covariance and *(d)* sampled covariance post-processing for disruption data from DIII-D with disruption due to MHD instability (test case 2). Note the different scaling in the color scale.

results obtained for the $F1$ score for dynamic time warping SOM clustering are given in table 7.

REFERENCES

AGGARWAL, CHARU C. 2018 *Neural Networks and Deep Learning*. Cham: Springer.

AYMERICH, E., SIAS, G., PISANO, F., CANNAS, B., CARCANGIU, S., SOZZI, C., STUART, C., CARVALHO, P.J., FANNI, A. & CONTRIBUTORS, JET 2022 Disruption prediction at JET through deep convolutional neural networks using spatiotemporal information from plasma profiles. *Nuclear Fusion* **62** (6), 066005.

BERNDT, DONALD J. & CLIFFORD, JAMES 1994 Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining*, *AAAIWS'94* , p. 359–370. AAAI Press.

BONNEEL, NICOLAS, RABIN, JULIEN, PEYRÉ, GABRIEL & PFISTER, HANSPETER 2015 Sliced and Radon Wasserstein Barycenters of Measures. *Journal of Mathematical Imaging and Vision* **1** (51), 22–45.

CLEVELAND, ROBERT B., CLEVELAND, WILLIAM S., MCRAE, JEAN E. & TERPENNING, IRMA 1990 Stl: A seasonal-trend decomposition procedure based on loess (with discussion). *Journal of Official Statistics* **6**, 3–73.

DURBIN, J. & KOOPMAN, S. J. 2002 A simple and efficient simulation smoother for state space time series analysis. *Biometrika* **89** (3), 603–615.

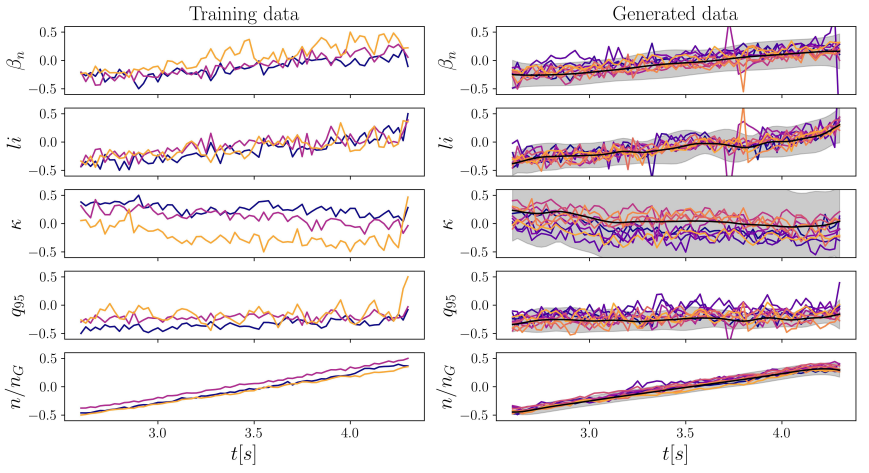

FIGURE 11. (a) Training data and (b) 10 generated data sets from the state space Student-t surrogate model together with the estimated mean (black solid line) and 95% confidence (grey shaded region) for disruption data from DIII-D with density accumulation (test case 3). Different colors correspond to different shots of training data and different samples of the generated data, respectively.

Metric	uncorrelated	emp. cov	emp. crosscov	sample cov
W_1	0.071 ± 0.088	0.03 ± 0.026	$\mathbf{0.025 \pm 0.019}$	0.03 ± 0.028
MMD p-value	0.43 ± 0.32	0.86 ± 0.17	0.84 ± 0.09	$\mathbf{0.87 \pm 0.13}$
MSE ρ_{rs}	0.0083 ± 0.0075	0.0076 ± 0.007	$\mathbf{0.0019 \pm 0.0019}$	0.007 ± 0.007
W_{emb}	0.334 ± 0.004	0.091 ± 0.001	$\mathbf{0.0598 \pm 0.0008}$	0.106 ± 0.001
MSE PSD $[10^{-6}]$	240 ± 33	11 ± 19	$\mathbf{0.8 \pm 0.5}$	79 ± 13
DTW	1.7 ± 2.2	0.9 ± 0.6	$\mathbf{0.8 \pm 0.6}$	0.9 ± 0.6
MSE mfPCA	0.138	0.021	**0.010**	0.025

TABLE 6. Post-processing method comparison for disruption data from DIII-D with density accumulation (test case 3). Mean and standard deviation over five dimensions and $N = 1000$ samples generated from the trained model for statistical metrics described in section 3. Best values are highlighted in bold.

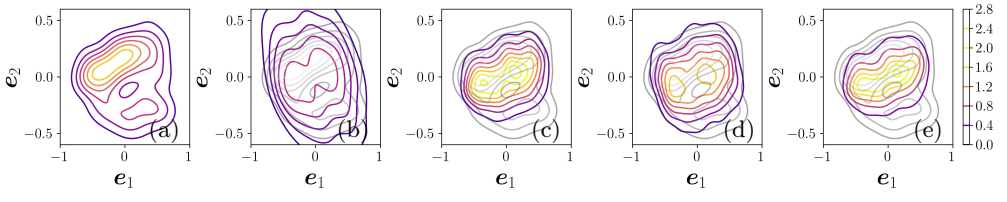

FIGURE 12. Kernel density of 2-D PCA embedding of (a) training data and generated data via (b) cross-covariance, (c) empirical covariance, (d) sampled covariance and (e) uncorrelated model for disruption data from DIII-D with density accumulation (test case 3). The embedded training data are shown in gray in all plots. The colorscale representing the density is the same in all plots.

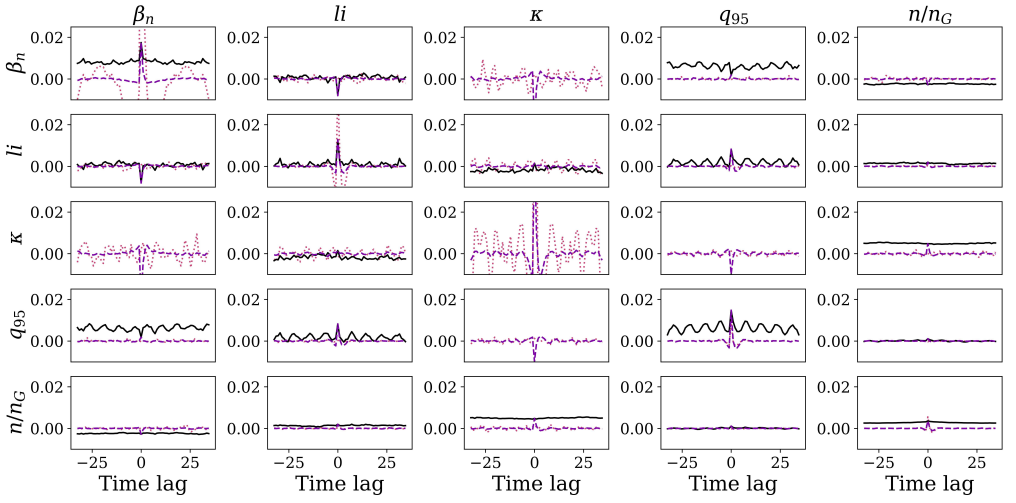

FIGURE 13. Comparison of cross-covariance of training data and generated data with cross-covariance (solid lines on top of each other, numerical error of order 10^{-16}), covariance or sampled covariance (dashed lines) post-processing and uncorrelated model (dotted line) for disruption data from DIII-D with density accumulation (test case 3).

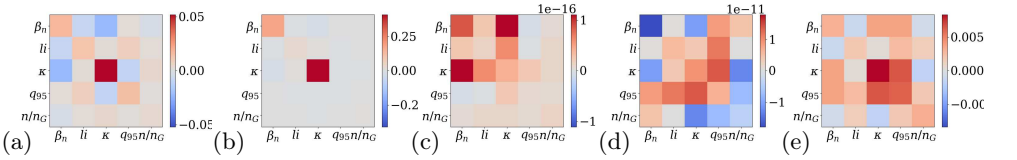

FIGURE 14. Comparison of covariance of training data *(a)* and difference of generated data *(b)* with uncorrelated model, *(c)* empirical covariance, *(d)* cross-covariance and *(d)* sampled covariance post-processing for disruption data from DIII-D with density accumulation (test case 3). Note the different scaling in the color scale.

Train	Test	training	uncorrelated	emp. cov	emp. crosscov	sample cov
original	generated	0.81	0.98	1.0	0.85	0.99
generated	original	0.99	0.92	0.93	0.93	0.93
mix	mix	0.96	0.96	0.96	0.94	0.97

TABLE 7. $F1$ score for dynamic time warping SOM clustering of different post-processing methods for disruption data from DIII-D with density accumulation (test case 3).

FLAMARY, RÉMI, COURTY, NICOLAS, GRAMFORT, ALEXANDRE, ALAYA, MOKHTAR Z., BOISBUNON, AURÉLIE, CHAMBON, STANISLAS, CHAPEL, LAETITIA, CORENFLOS, ADRIEN, FATRAS, KILIAN, FOURNIER, NEMO, GAUTHERON, LÉO, GAYRAUD, NATHALIE T.H., JANATI, HICHAM, RAKOTOMAMONJY, ALAIN, REDKO, IEVGEN, ROLET, ANTOINE, SCHUTZ, ANTONY, SEGUY, VIVIEN, SUTHERLAND, DANICA J., TAVENARD, ROMAIN, TONG, ALEXANDER & VAYER, TITOUAN 2021 Pot: Python optimal transport. *Journal of Machine Learning Research* **22** (78), 1–8.

GRETTON, ARTHUR, BORGWARDT, KARSTEN M., RASCH, MALTE J., SCHÖLKOPF, BERNHARD & SMOLA, ALEXANDER 2012 A kernel two-sample test. *J. Mach. Learn. Res.* **13** (null), 723–773.

Happ, Clara & Greven, Sonja 2018 Multivariate functional principal component analysis for data observed on different (dimensional) domains. *Journal of the American Statistical Association* **113** (522), 649–659, arXiv: https://doi.org/10.1080/01621459.2016.1273115.

Iwana, Brian Kenji & Uchida, Seiichi 2021 An empirical survey of data augmentation for time series classification with neural networks. *PLOS ONE* **16** (7), 1–32.

Kang, Yanfei, Hyndman, Rob J. & Li, Feng 2020 Gratis: Generating time series with diverse and controllable characteristics. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **13** (4), 354–376, arXiv: https://onlinelibrary.wiley.com/doi/pdf/10.1002/sam.11461.

Kates-Harbeck, Julian, Svyatkovskiy, Alexey & Tang, William 2019 Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature* **568** (7753), 526–531.

Kessy, Agnan, Lewin, Alex & Strimmer, Korbinian 2018 Optimal whitening and decorrelation. *The American Statistician* **72** (4), 309–314, arXiv: https://doi.org/10.1080/00031305.2016.1277159.

Montes, K.J., Rea, C., Tinguely, R.A., Sweeney, R., Zhu, J. & Granetz, R.S. 2021 A semi-supervised machine learning detector for physics events in tokamak discharges. *Nuclear Fusion* **61** (2), 026022.

Murphy, Kevin P. 2022 *Probabilistic Machine Learning: An introduction*. MIT Press.

Park, K.I. 2017 *Fundamentals of Probability and Stochastic Processes with Applications to Communications*. Springer International Publishing.

Pau, A., Fanni, A., Carcangiu, S., Cannas, B., Sias, G., Murari, A. & and, F. Rimini 2019 A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at JET. *Nuclear Fusion* **59** (10), 106017.

Rea, Cristina & Granetz, Robert S. 2018 Exploratory machine learning studies for disruption prediction using large databases on diii-d. *Fusion Science and Technology* **74** (1-2), 89–100, arXiv: https://doi.org/10.1080/15361055.2017.1407206.

Rea, C., Montes, K.J., Erickson, K.G., Granetz, R.S. & Tinguely, R.A. 2019 A real-time machine learning-based disruption predictor in DIII-d. *Nuclear Fusion* **59** (9), 096016.

Rea, C., Montes, K. J., Pau, A., Granetz, R. S. & Sauter, O. 2020 Progress toward interpretable machine learning–based disruption predictors across tokamaks. *Fusion Science and Technology* **76** (8), 912–924, arXiv: https://doi.org/10.1080/15361055.2020.1798589.

Roth, Michael, Ardeshiri, Tohid, Özkan, Emre & Gustafsson, Fredrik 2017 Robust bayesian filtering and smoothing using student's t distribution. *CoRR* **abs/1703.02428**, arXiv: 1703.02428.

Shah, Amar, Wilson, Andrew Gordon & Ghahramani, Zoubin 2014 Student-t processes as alternatives to gaussian processes. *Artificial Intelligence and Statistics* .

Shorten, Connor & Khoshgoftaar, Taghi M. 2019 A survey on image data augmentation for deep learning. *Journal of Big Data* **6**, 1–48.

Solin, Arno & Särkkä, Simo 2015 State Space Methods for Efficient Inference in Student-t Process Regression. *Tech. Rep.*.

Särkkä, Simo 2013 *Bayesian Filtering and Smoothing*. *Institute of Mathematical Statistics Textbooks* . Cambridge University Press.

Särkkä, Simo & Solin, Arno 2019 *Applied Stochastic Differential Equations*. *Institute of Mathematical Statistics Textbooks* . Cambridge University Press.

Vettigli, Giuseppe 2018 Minisom: minimalistic and numpy-based implementation of the self organizing map.

Villani, C. 2008 *Optimal Transport: Old and New*. *Grundlehren der mathematischen Wissenschaften* . Springer Berlin Heidelberg.

Virtanen, Pauli, Gommers, Ralf, Oliphant, Travis E., Haberland, Matt, Reddy, Tyler, Cournapeau, David, Burovski, Evgeni, Peterson, Pearu, Weckesser, Warren, Bright, Jonathan, van der Walt, Stéfan J., Brett, Matthew, Wilson, Joshua, Millman, K. Jarrod, Mayorov, Nikolay, Nelson, Andrew R. J., Jones, Eric, Kern, Robert, Larson, Eric, Carey, C J, Polat, İlhan, Feng, Yu, Moore, Eric W., VanderPlas, Jake, Laxalde, Denis, Perktold, Josef, Cimrman, Robert, Henriksen, Ian, Quintero, E. A., Harris, Charles R.,

Archibald, Anne M., Ribeiro, Antônio H., Pedregosa, Fabian, van Mulbregt, Paul & SciPy 1.0 Contributors 2020 SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272.

Wen, Qingsong, Gao, Jingkun, Song, Xiaomin, Sun, Liang, Xu, Huan & Zhu, Shenghuo 2019 Robuststl: A robust seasonal-trend decomposition algorithm for long time series. *Proceedings of the AAAI Conference on Artificial Intelligence* **33** (01), 5409–5416.

Wen, Qingsong, Sun, Liang, Yang, Fan, Song, Xiaomin, Gao, Jingkun, Wang, Xue & Xu, Huan 2021 Time series data augmentation for deep learning: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization.

Wilkinson, William J, Chang, Paul E, Andersen, Michael Riis & Solin, Arno 2020 State Space Expectation Propagation: Efficient Inference Schemes for Temporal Gaussian Processes. *Tech. Rep.*.

Williams, Christopher K. I. & Rasmussen, Carl Edward 1996 Gaussian processes for regression. In *Advances in Neural Information Processing Systems 8*, pp. 514–520. MIT press.

Yoon, Jinsung, Jarrett, Daniel & van der Schaar, Mihaela 2019 Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems* (ed. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox & R. Garnett), , vol. 32. Curran Associates, Inc.