Plasma State discovery using Bayesian methods

I. Kharitonov[1], G.F. Subbotin[1], D. Sorokin[1], A. Granovskiy[1], E. Adishchev[1], M.S. Stokolesov, E.N. Khairutdinov[1], M.R. Nurgaliev[1]

[1] *Next Step Fusion, Luxembourg, Mondorf-les-Bains, 56A Avenue François Clément, L-5612*
e-mail: ikh@nextfusion.org

A novel state discovery algorithm has been developed to optimize plasma states in tokamak scenarios. Optimizing plasma states, i.e., finding optimal configurations based on a specific functional, is a very challenging task due to the high dimensionality of the problem and the slow response of the simulator. This could become one of the decisive factors in fusion power plant scenarios for tokamak optimization. The process is envisioned as an iterative cycle, where the device is optimized first, followed by the plasma state, and then the device again. By embedding magnetic and kinetic parameters, along with sensor data, into a latent space using a generative model, we achieved scalable optimization and reliable state reconstruction. Additionally, the generative model helps to manage the uncertainty of the simulation process, ensuring safer and more interpretable operations, as it addresses the inherent lack of interpretability common in machine learning methods.

The problem of optimizing plasma state parameters is formulated in a probabilistic framework. The optimization algorithm uses a Gaussian Process (GP) as the surrogate model. However, GPs perform poorly in high-dimensional spaces. For D3D, the state vector includes E-coil currents, F-coil currents, ion and electron temperatures at the core and boundary, the pressure gradient (PPRIME), and toroidal flux function (FFPRIME) vectors (see Fig. 1). This results in a dimensionality exceeding 100. Another challenge is that the environment is very slow, limiting the number of plasma states we can sample. To address this, we introduce a second surrogate model

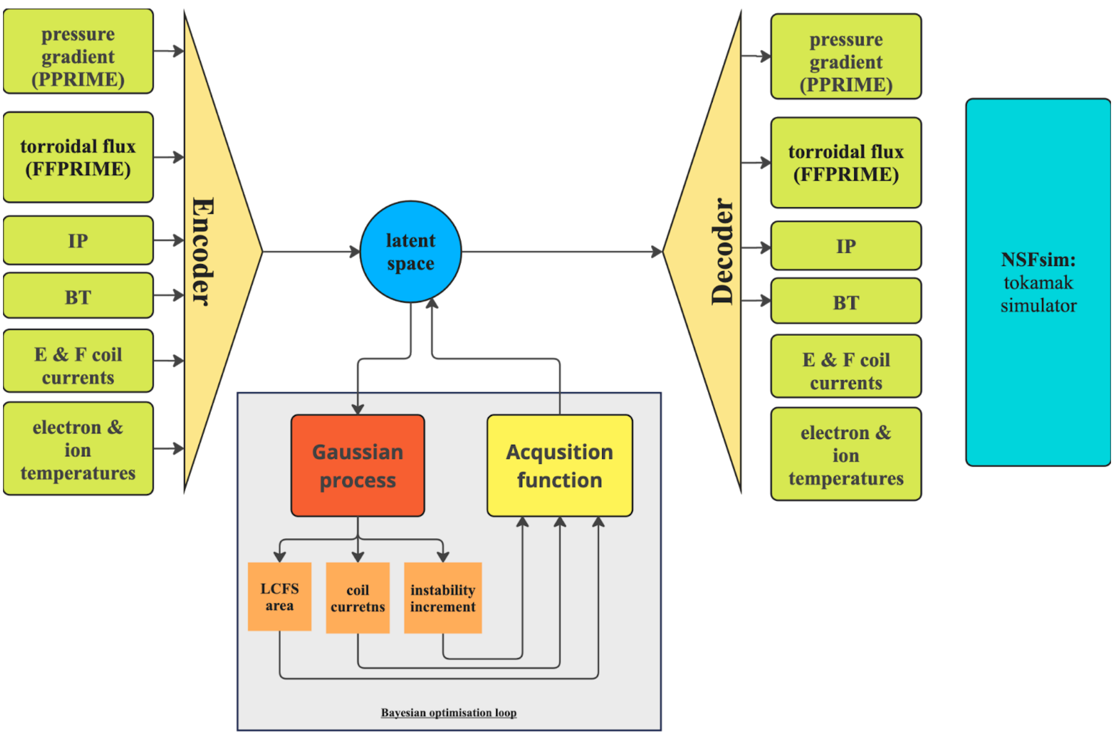

Fig. 1. Integration of Bayesian optimization loop into variational autoencoder-decoder

based on a generative model – a variational autoencoder (VAE) – to learn a compact representation of plasma states. This enables the algorithm to operate in a much smaller latent space (with a dimensionality of 8 in our case) instead of the original input space. The strength of this model lies in its ability to compress high-dimensional vectors, identify patterns in the data, and encode them implicitly. Visualization of states in latent spaces is given in Fig. 2.

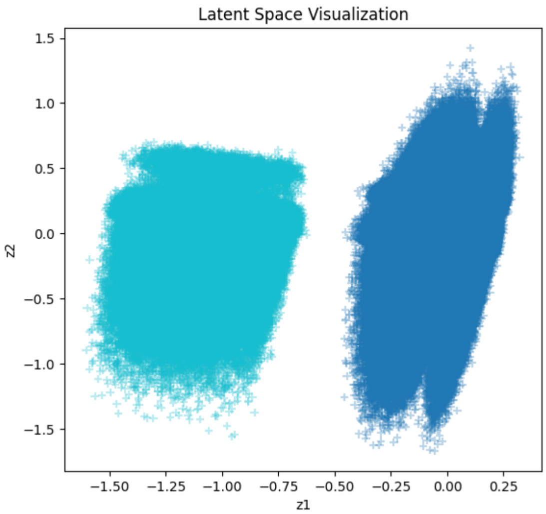

Fig. 2. Example of state separation in latent space

Another advantage of the generative model is its ability to generate plasma states for specific conditions and sample states from the model. For instance, we know that the same plasma state can be achieved using different combinations of coil currents. This model provides a probabilistic distribution, allowing the generation of feasible plasma states. Additionally, this representation model can be utilized in control agents as a state reconstructor. Moreover, since the VAE is a neural network, sampling new data from it is significantly faster than querying the environment. Moreover, the generative model predicts a distribution of parameters instead of fixed values, providing an estimate of the model's uncertainty. This allows us to determine when the model's predictions should be trusted or avoided, based on its confidence level. This is crucial for safety, as machine learning methods often lack interpretability, which can lead to unsafe outcomes if predictions are blindly followed without understanding their reliability. The pipeline operates as follows:

1. Use an environment capable of modeling a specific tokamak.
2. Initially, random plasma states are sampled from the environment to train the generative model (VAE).
3. Plasma state optimization begins using the surrogate model (Gaussian Process). This model describes the relationship between a quality function (e.g., maximum plasma volume, minimum instability increment, or minimum coil currents) and the vector in the latent space. The GP provides both the value of the function and the uncertainty of its estimate.
4. The optimization seeks the maximum of this function using an acquisition function that considers both the value and the variance of the estimate.
5. The Gaussian Process model generates new samples, which are decoded into plasma states by the VAE and fed back into the environment. The environment solves equilibrium and kinetic problems and augments the plasma state with additional data, such as plasma shape. This updated dataset is added to the training set of the generative model, and the process is repeated iteratively.

We performed ablation studies comparing direct Bayesian optimization with and without the variational autoencoder. We analyzed the latent representations and provided examples of conditional generation of plasma states. We also presented the results of plasma state optimization as a set, which was analyzed from both a control perspective and a tokamak engineering perspective to demonstrate its applicability in a tokamak optimization pipeline.