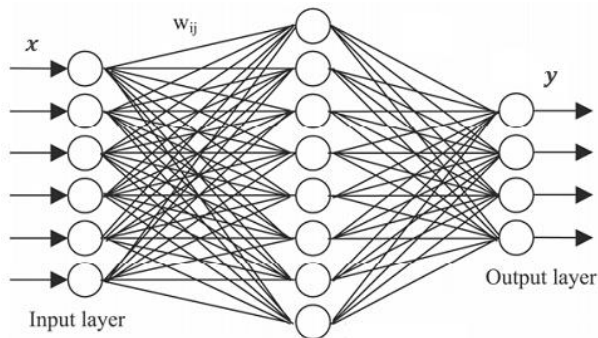


Control and Operation of Tokamaks

Machine Learning for plasma control



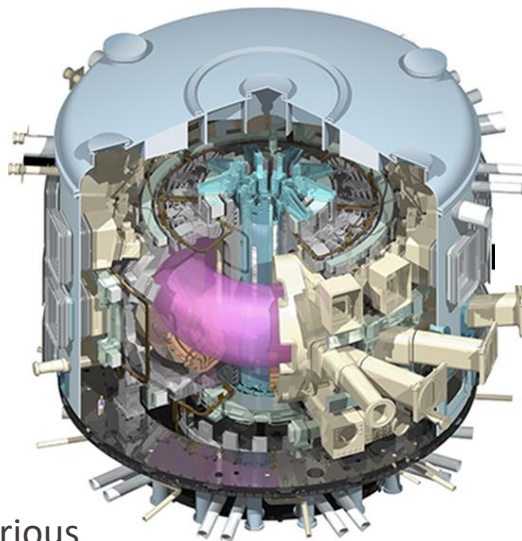
Alessandro Pau
alessandro.pau@epfl.ch

16/02/2023, Lausanne

Data in fusion: a challenge in itself

■ **Tokamaks** are one of the most complex system in nature...

■ **Massive amount of data** (Big data – 2PB/day at ITER, high bandwidth diagnostics)

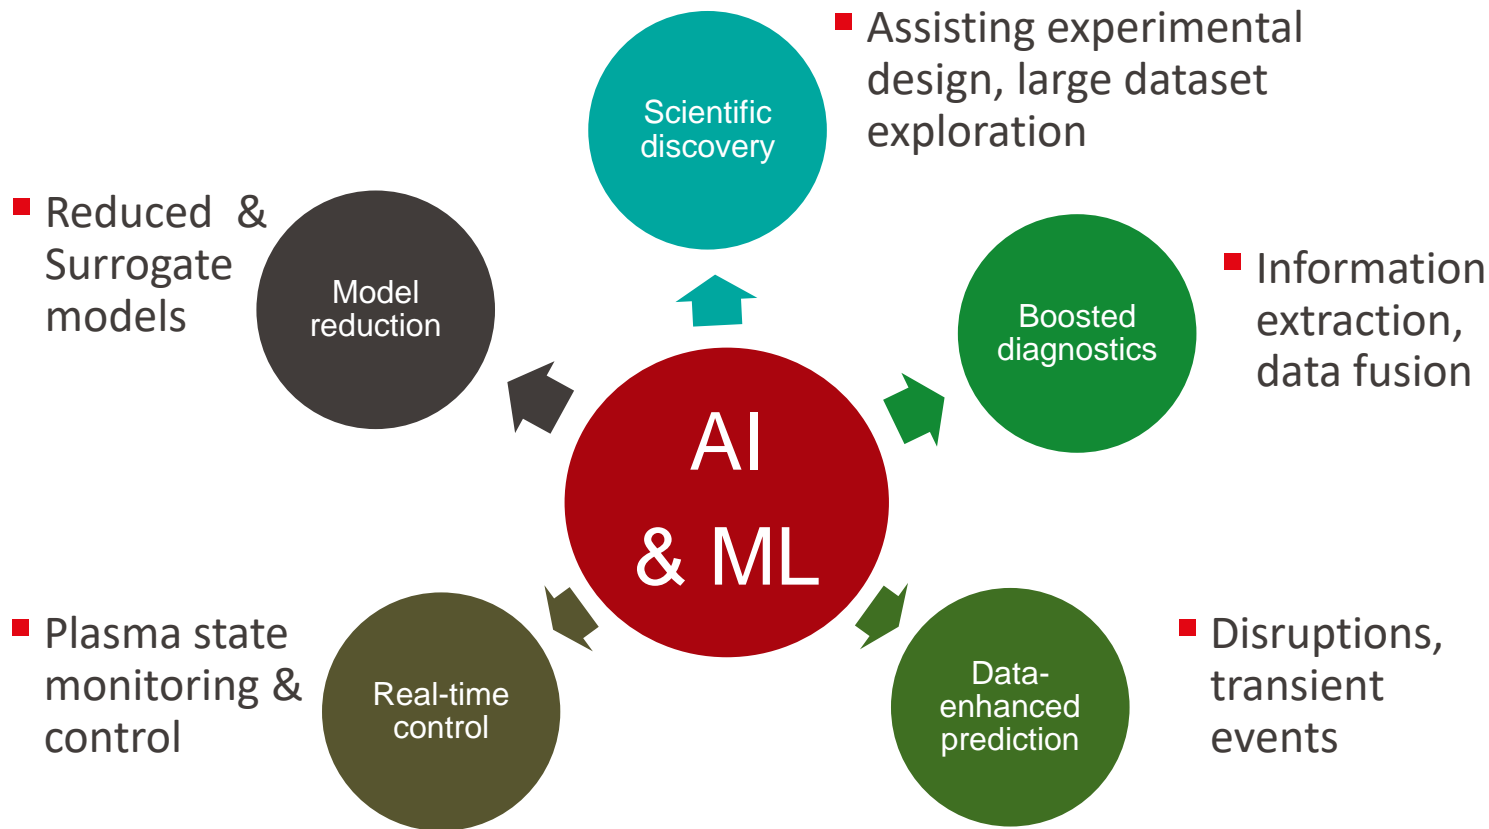


■ **High-dimensionality** (many diagnostics measuring various plasma properties)

■ **Multi-scales, multi-physics** (integrated tokamak modeling)

■ **Heterogenous** (various formats, facility dependent data ecosystem)

Advancing fusion by leveraging AI and ML



EPFL Control augmentation in modern Plasma Control Systems

- Magnetic control via DRL

REF: [F.J. Degraeve, F. Felici et al. *Nature* 2022]

- plasma state **monitoring** and **forecasting** for control augmentation

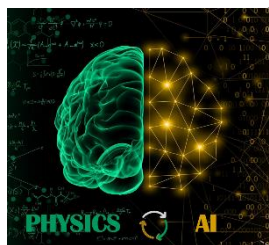
- Detection of **off-normal events** to react with specific control tasks in **real-time**

- **Proximity to operational limits**

REF: [Pau et al *IEEE-TPS* 2018]

REF: [Pau et al *NF* 2019]

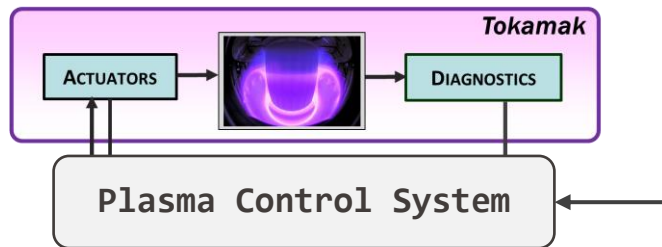
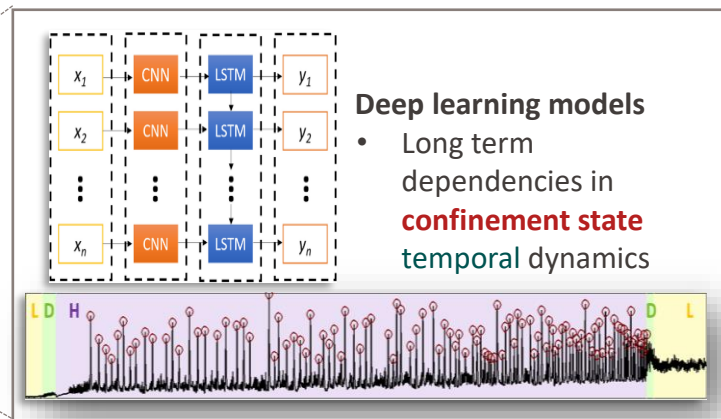
EVENT DETECTION



REF: [F. Matos et al. *NF* 2020]

REF: [F. Matos et al. *NF* 2021]

REF: [G. Marceca et al. *NeurIPS* 2021]

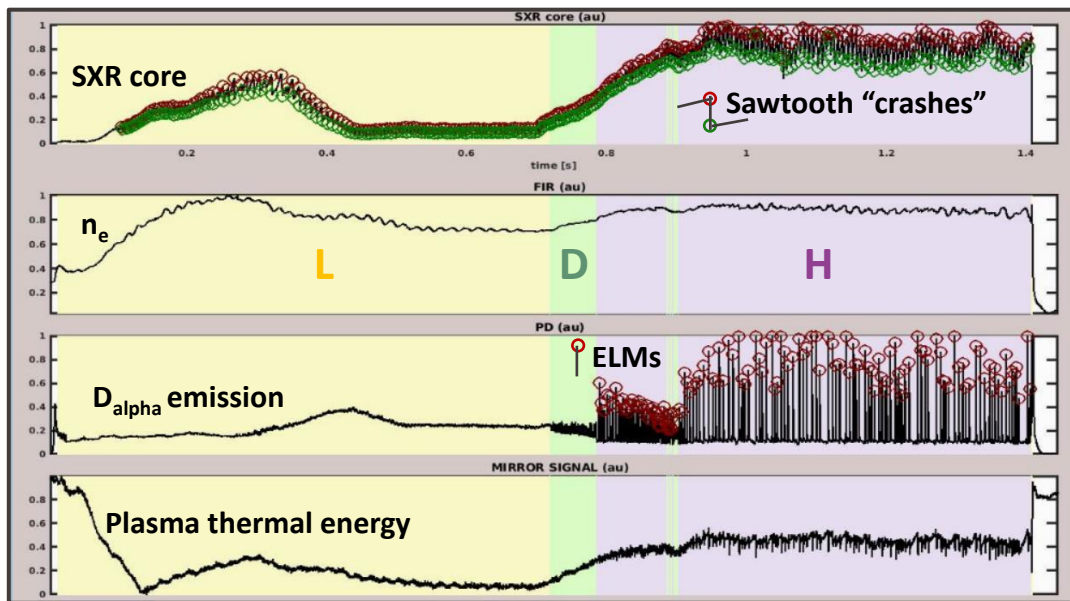


...combination & integration of:

- **Physics-, model- & ML-based** approaches

EPFL Event detection and plasma state classification

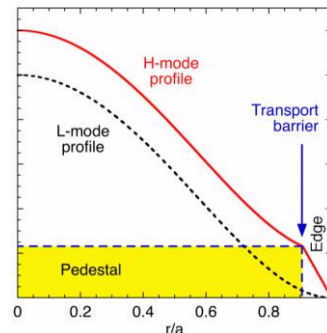
- Plasma confinement states [L=Low; D=Dithering; H=High]



LHD phases from validated-labeled file
 A total of 138 ELMs were found in the loaded data
 A total of 249 ST_MDs were found in the loaded data

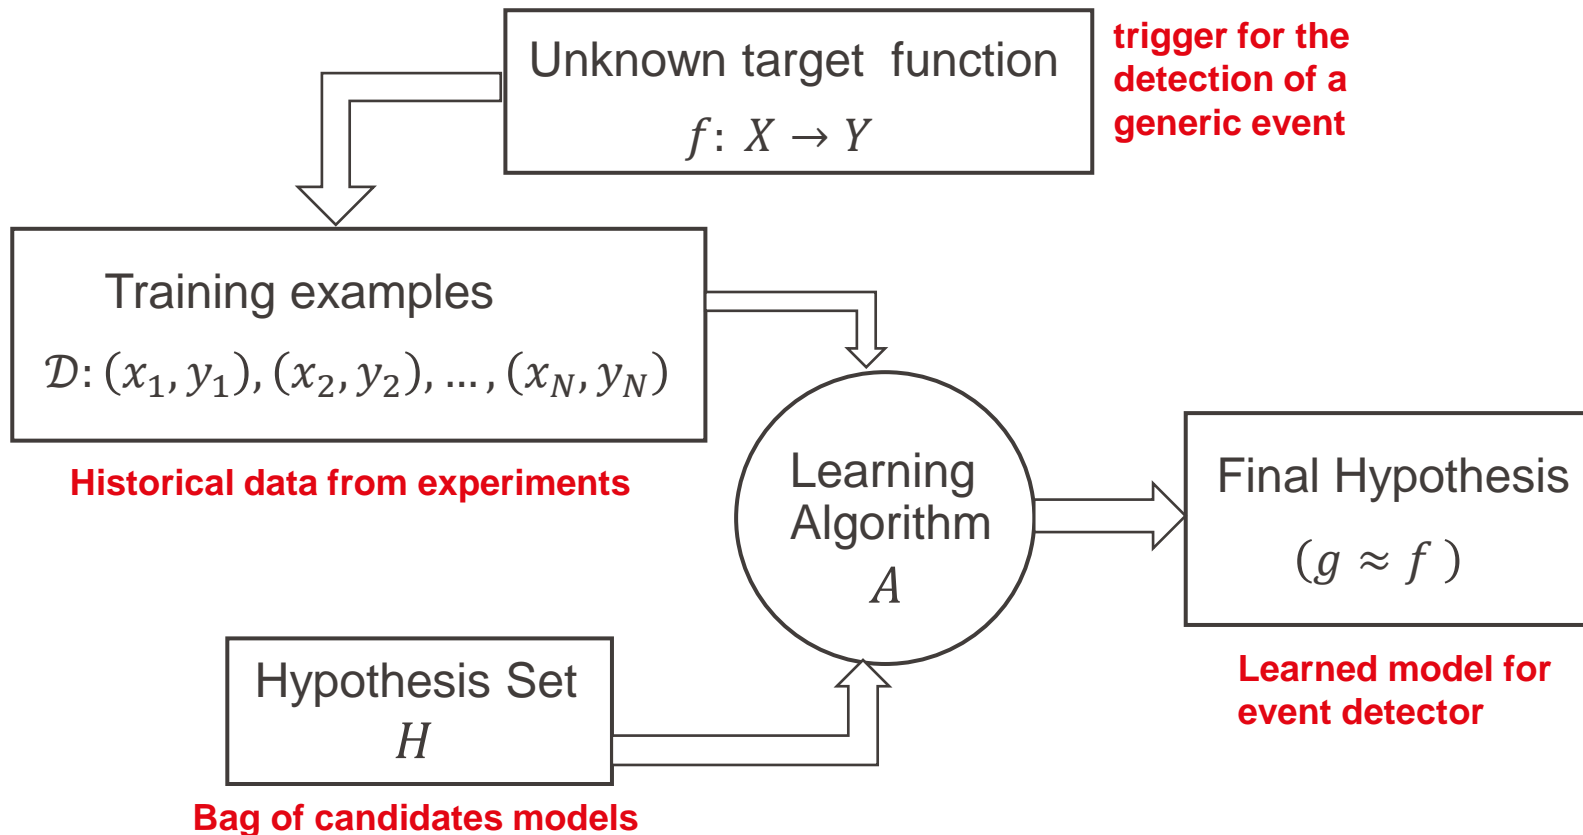
↳ An experiments have potentially **hundreds of events**....

- Plasma can evolve in one of several possible **confinement states** (typical categorization in **L=Low; D=Dithering; H=High**).
- By applying **sufficient heating power**, the plasma spontaneously transitions from a low to a high confinement state
- **H-mode**: improved energy confinement state with **reduced particles and energy transport** outwards formation of an edge transport barrier (ETB) and a cyclic MHD instability called Edge Localized Modes (ELMs).



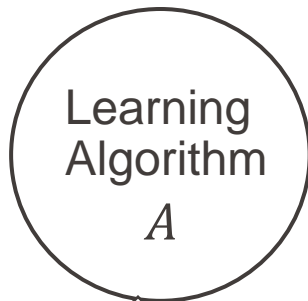
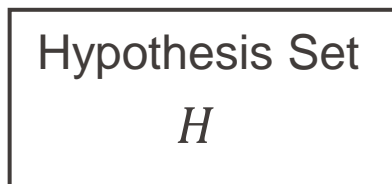
**How can we formulate
a control problem
leveraging ML?**

Formulation of the learning problem



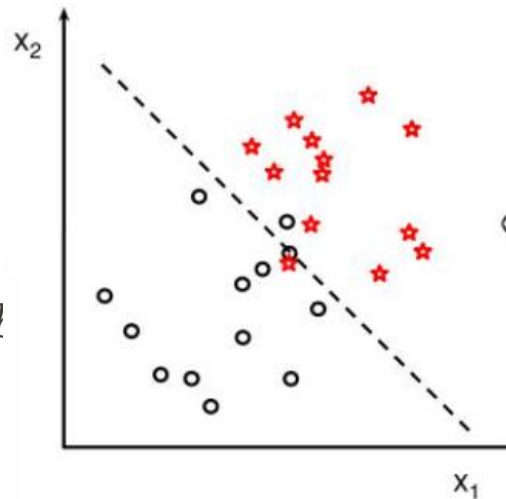
Formulation of the learning problem

Bag of candidates models

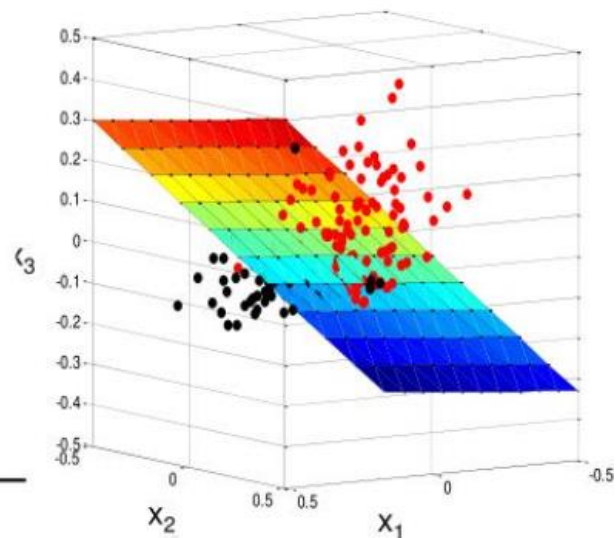


Simple generic linear hypothesis:
Functional form $h(x)$

$$h(x) = \text{sign} \left(\left(\sum_{i=1}^d w_i \cdot x_i \right) + l \right)$$

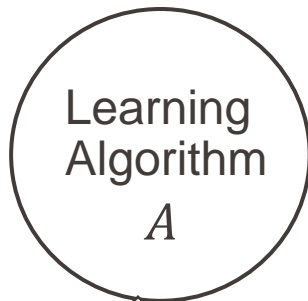
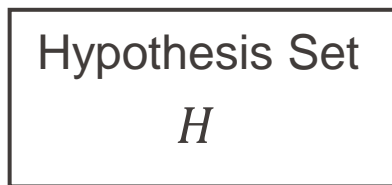


Linear decision boundary



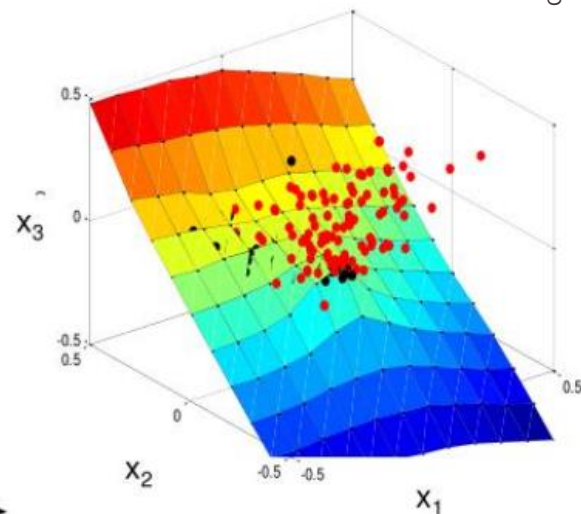
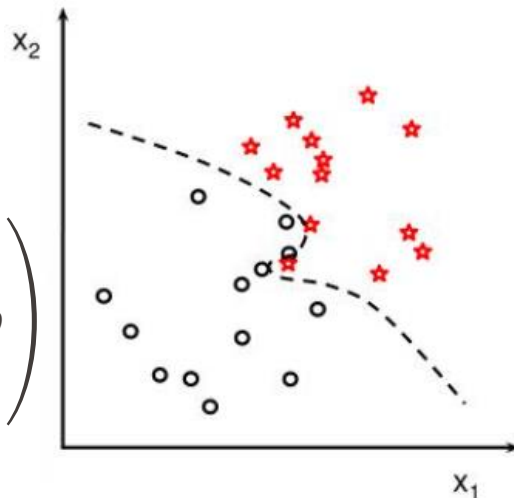
Formulation of the learning problem

Bag of candidates models



Simple generic linear hypothesis:
Functional form $h(x)$

$$h(x) = \text{sign} \left(\left(\sum_{i=1}^d w_i \cdot x_i \right) + b \right)$$



Functional form $h(x) \rightarrow$ *perceptron*

$$h(x) = \text{sign} \left(\left(\sum_{i=1}^d w_i \cdot x_i \right) + b \right)$$

If the problem is linearly separable:

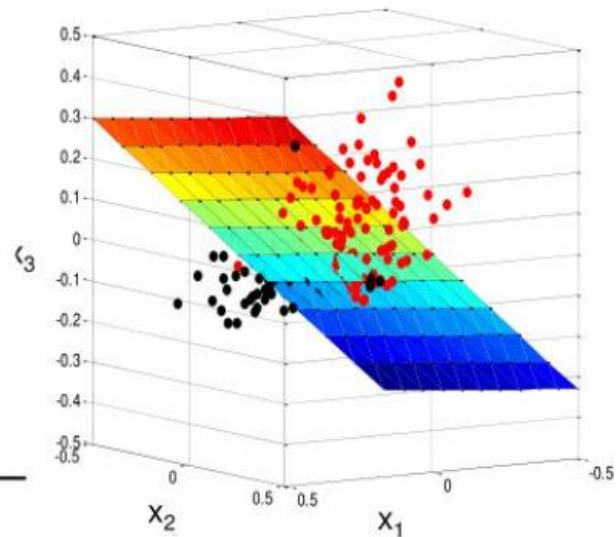
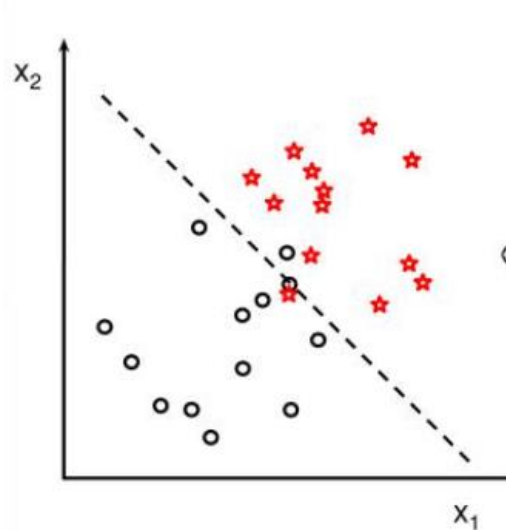
- Set of $w = (w_1, w_2, \dots, w_d)$, and b such that $h(x) = y$ is satisfied for all training samples.

Iterative procedure:

$$w(t+1) = w(t) + y(t) \cdot x(t)$$

Training examples
 $\mathcal{D}: (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$

Linear decision boundary



Training data

$$\mathcal{D}: (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$$

Batch Learning:

- training on a fixed **dataset entirely available** to the learning algorithm, with model's parameters being updated after each iteration through the data.
 - typically **more computationally efficient**, but less flexible to adapt to new data distributions.

Active learning:

- the learning algorithm is able to **interactively query** an information source to obtain the **desired outputs on new data points** (most informative data points to learn from)
 - often used when there is a limited amount of labelled data available: selecting which data points to learn from, the model can learn more effectively and efficiently.

Online Learning:

- the algorithm receives **one example at a time**, with model's parameters being updated incrementally as new data comes in.
 - Useful in case of limitations on computing and storage

...Type of learning: Reinforcement Learning

Training data

(input, output, reward/penalty)

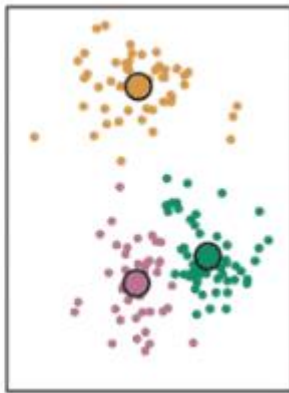
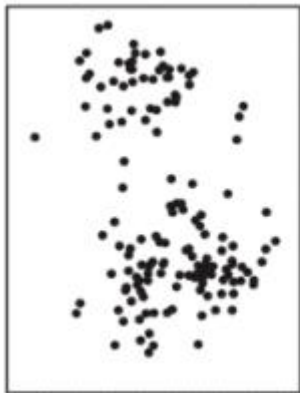
Reinforcement learning:

- an “**agent**” learns to make decisions by continuously interacting with an environment and receiving **feedback** in the form of **rewards** or **penalties**.
 - The goal of reinforcement learning is to learn a “**policy**”, which is a **mapping from states to actions**, that maximizes the cumulative reward the agent receives over time.
-
- **Training data** consists of **sequences of states, actions, and rewards**.
 - Learning by **trial-and-error**, where the agent takes actions, receives rewards, and updates its policy based on the observed rewards until convergence to an optimal solution

...Type of learning: Unsupervised Learning

Training data

$(x_1, \dots), (x_2, \dots), \dots, (x_N, \dots)$

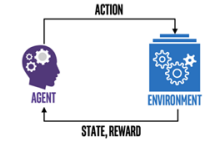
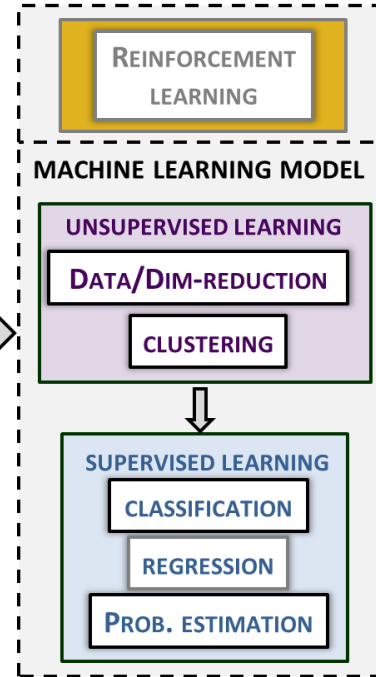
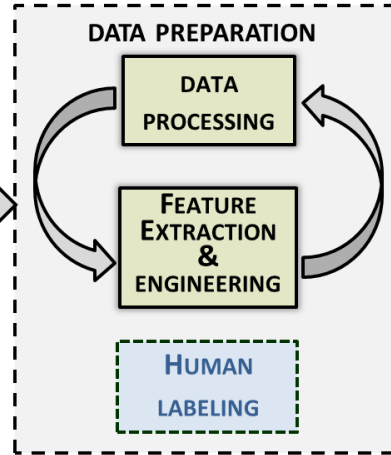
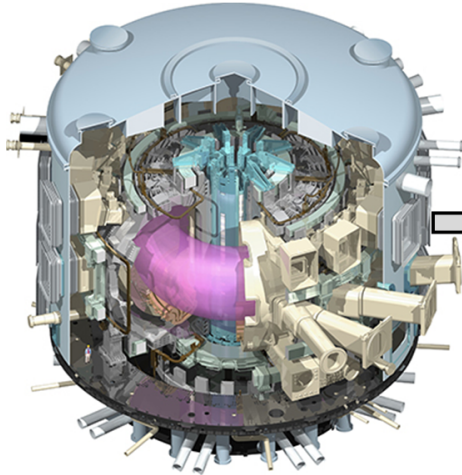


Unsupervised Learning:

- useful to discover **patterns** or **structure** in the data, with **no labelled data**. The learning algorithm task is to identify structure in the data, such as **grouping** similar examples according to a well defined **metric**.
- Some common unsupervised learning techniques:
 - **Clustering**: grouping of similar examples into clusters,
 - **dimensionality_reduction**: projection of the data into a lower-dimensional space while preserving as much of the structure of the data as possible
 - **anomaly detection**: identification of examples that are significantly different from the majority of the data (...novelty detection).

...Typical Machine learning workflow

- Massive volume of data
- High-dimensional;



PLASMA CONTROL & SIMULATORS

- DATA VISUALIZATION
- PATTERN RECOGNITION
- GENERATIVE MODELS
- ANOMALY DETECTION
- EVENT/OBJECT DETECTION
- IMAGE PROCESSING
- TIME SERIES FORECASTING

- Heterogenous
- multiple timescales

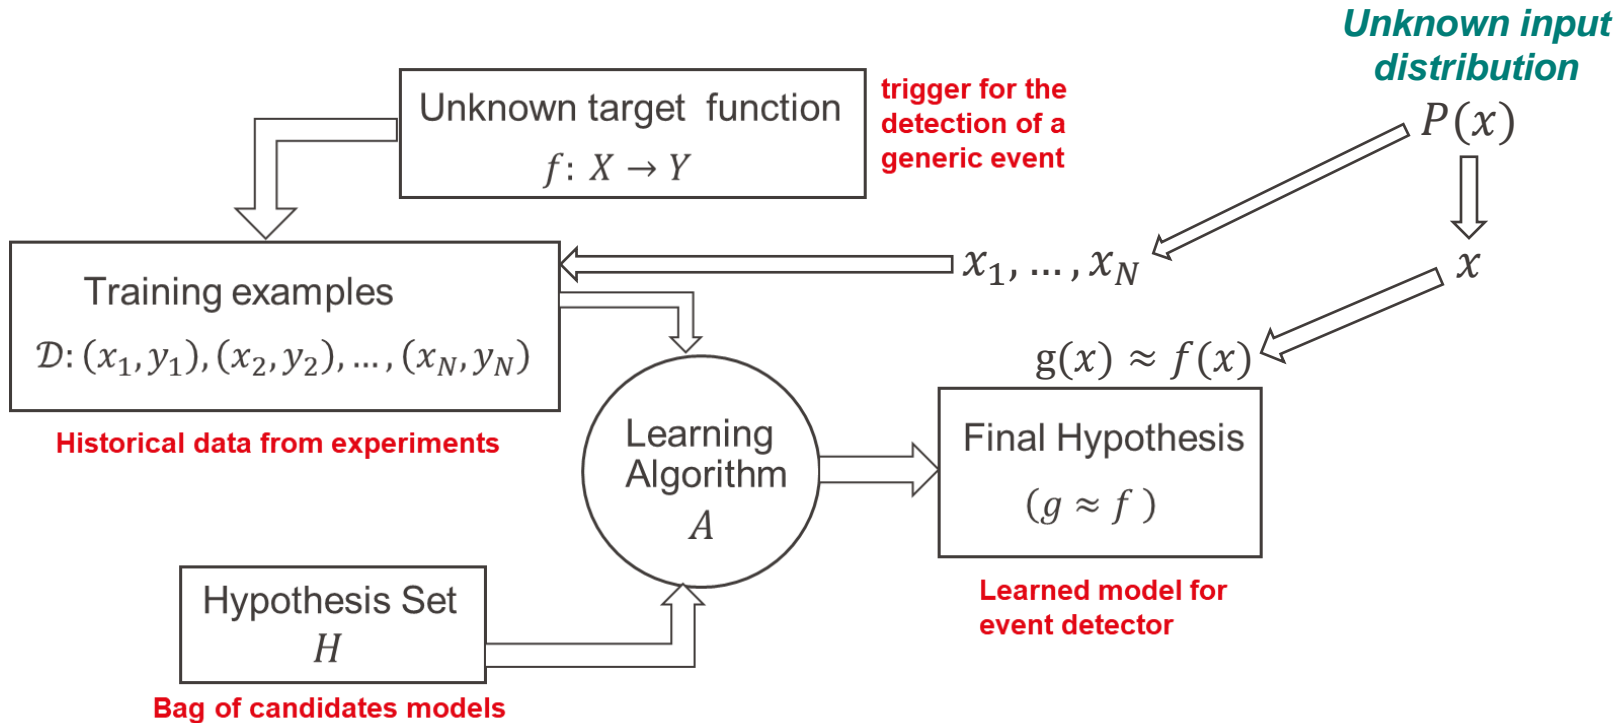
Unknown target function

$$f: X \rightarrow Y$$

Model fitting, or training:

- Learn the **unknown target function** describing the mapping $f(x, \theta) \rightarrow y$
- find the **set of parameters** θ that best describe the mapping between the input and output variables in the data.
- A good fitting enables the model to **generalize** from the training data to new, unseen data.
- The **quality of the model fit** depends on a number of things:
 - choice of model;
 - quality of the input data (and labelling in supervised settings);
 - choice of optimization algorithm used to fit the model.

When learning is feasible...



Hoeffding Inequality

upper bound

$$P([E_{in}h(x) - E_{out}h(x)] \geq \epsilon) \leq 2e^{-2\epsilon^2N} \text{ for any } \epsilon > 0$$

*Unknown input
distribution*

$P(x)$



x_1, \dots, x_N



$g(x) \approx f(x)$

Hoeffding Inequality

upper bound

$$P([E_{in}h(x) - E_{out}h(x)] \geq \epsilon) \leq 2e^{-2\epsilon^2N} \quad \text{for any } \epsilon > 0$$



$\mathcal{H} = \{h_1, \dots, h_M\} \rightarrow g$ *h is fixed before generating
the dataset, g is our final
hypothesis given \mathcal{D}*



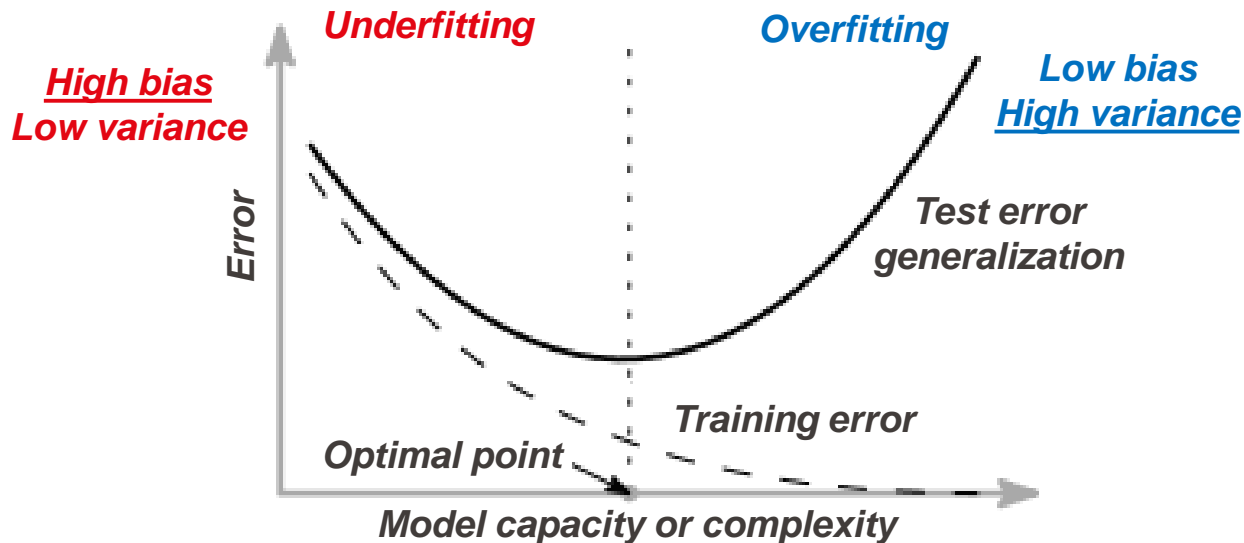
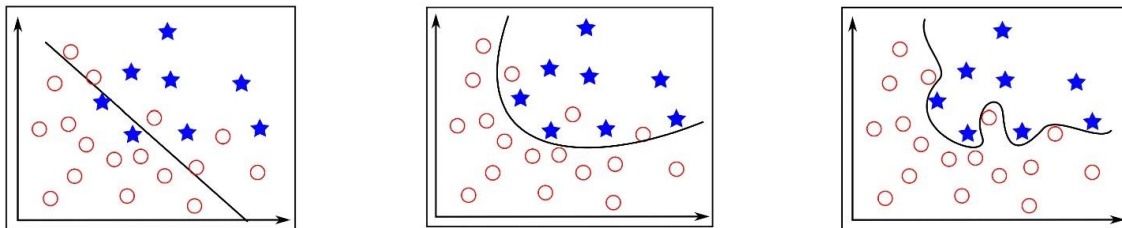
$$P([E_{in}h(x) - E_{out}h(x)] \geq \epsilon) \leq 2Me^{-2\epsilon^2N} \quad \text{for any } \epsilon > 0$$



Vapnik-Chervonenkis (VC) dimension: measure of the **complexity** of a learning problem.

- used to derive **bounds** on the **generalization** error of learning algorithms, and to compare the **expressive power** of different families of functions or models.

Underfitting and Overfitting: bias/variance trade-off



$\hat{\theta}$ is a “point estimator”
(predicted “target”)

$$\text{Bias}[\hat{\theta}] = E[\hat{\theta}] - \theta$$

$E[\hat{\theta}]$ is evaluated over
the training sets

$$\text{Var}[\hat{\theta}] = E[\hat{\theta}^2] - (E[\hat{\theta}])^2$$

$$\text{Var}[\hat{\theta}] = [(E[\hat{\theta}] - \hat{\theta})^2]$$

$$\begin{aligned} \text{Bias}[\hat{f}(x)] &= E[\hat{f}(x)] - f(x) \\ \text{Var}[\hat{f}(x)] &= [E[\hat{f}(x)] - f(x)]^2 \end{aligned}$$



$$y = f(x) - \begin{cases} \text{true function } f \\ \text{noise error } \varepsilon \text{ with mean } \mathbf{0} \text{ and variance } \sigma^2 \\ \text{true target (label) } y \end{cases}$$



Error on unseen training data x

$$\text{Err}(x) = E[(y - \hat{f}(x))^2]$$



Bias/variance decomposition of the squared error [derivation]

$$\text{Err}(x) = (\text{Bias}[\hat{f}(x)])^2 + \text{Var}[\hat{f}(x)] + \sigma^2$$

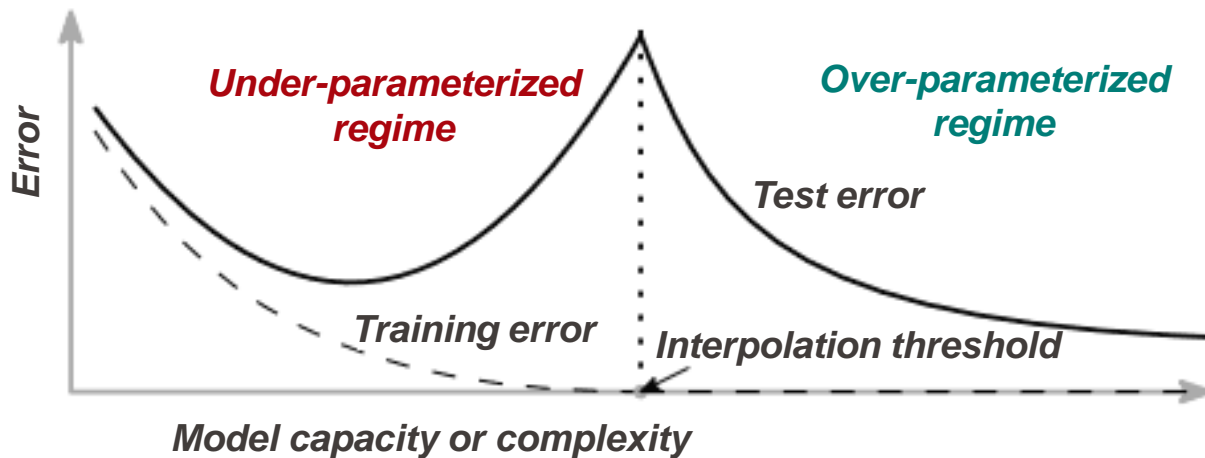
Irreducible error

How far the learned model is from the true function

Changes when the model is trained on different data samples

EPFL Underfitting and Overfitting: bias/variance trade-off

- **Classical machine learning** operates in the under-parameterized regime
- **Modern machine learning** operates in the over-parameterized regime



Phenomenon known as: “**double descend**”

- Model-wise double descend
- Epoch-wise double descend

Model fitting, or training:

- Learn the unknown target function describing the relation $f(x, \theta) \rightarrow y$
- find the set of parameters θ that best describe the mapping between the input and output variables in the data.
- Given the input data \mathcal{D} , solve a optimization problem in terms of minimization of a **objective or loss function**

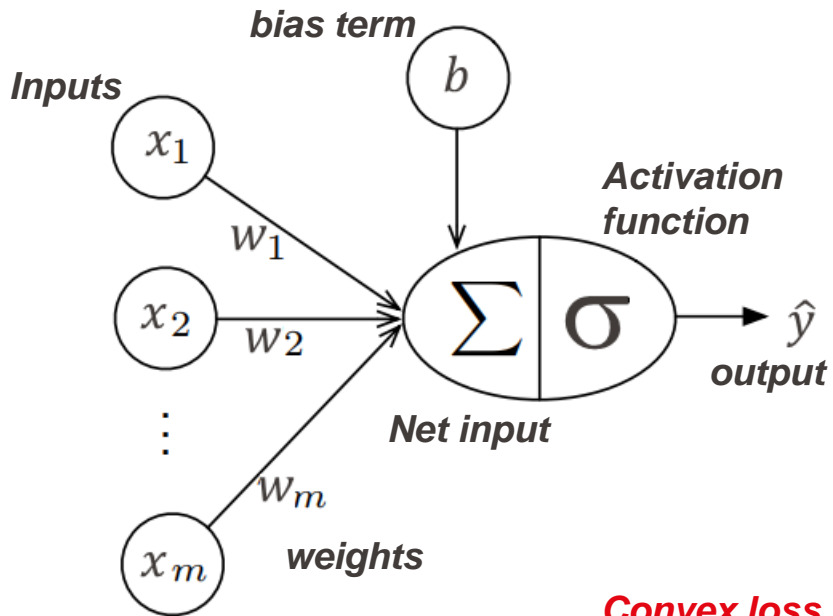
Training examples

$$\mathcal{D}: x_1, x_2, \dots, x_N$$

Loss function

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \mathcal{L}(\mathcal{D}|\theta)$$

- We call **inference** the procedure with which we quantify the uncertainty or confidence in the estimate $\hat{\theta}$

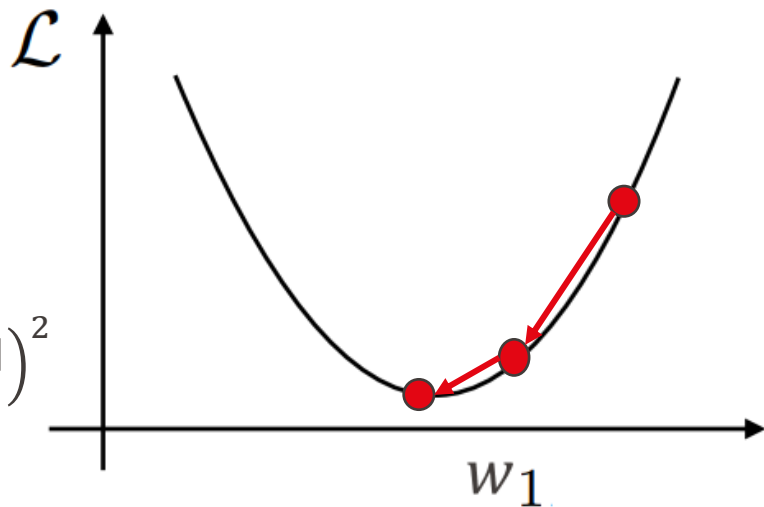


$$\hat{y} = \sigma \left(\left(\sum_{i=1}^m w_i \cdot x_i \right) + b \right) = \sigma(\mathbf{x}^T \mathbf{w} + b)$$

$$\mathcal{D}: (\mathbf{x}^{[1]}, y^{[1]}), (\mathbf{x}^{[2]}, y^{[2]}), \dots, (\mathbf{x}^{[N]}, y^{[N]}) \in \mathfrak{R}^m$$

Convex loss function

$$\mathcal{L}(\mathbf{w}, b) = \sum_j \left(\hat{y}^{[j]} - y^{[j]} \right)^2$$



On-line mode:

- Learning faster but more noisy (shuffling each epoch)

$$\hat{y} = \sigma \left(\left(\sum_{i=1}^m w_i \cdot x_i \right) + b \right) = \sigma(\mathbf{x}^T \mathbf{w} + b)$$

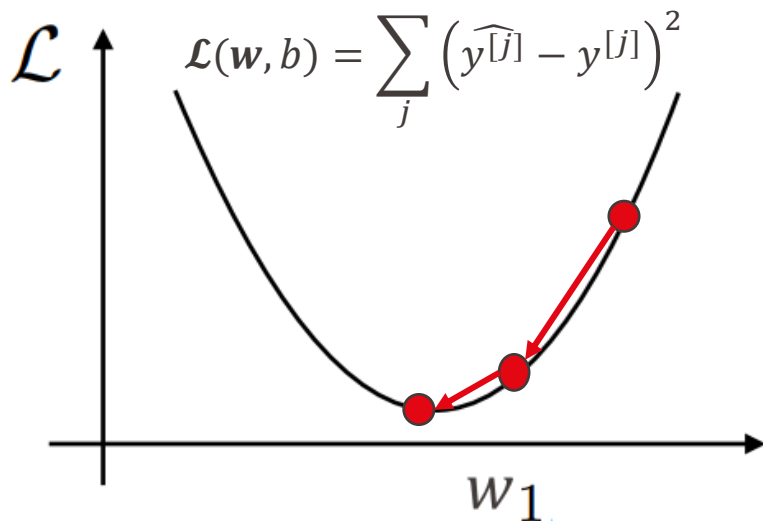
Batch mode:

- Slower but less sensitive to noise

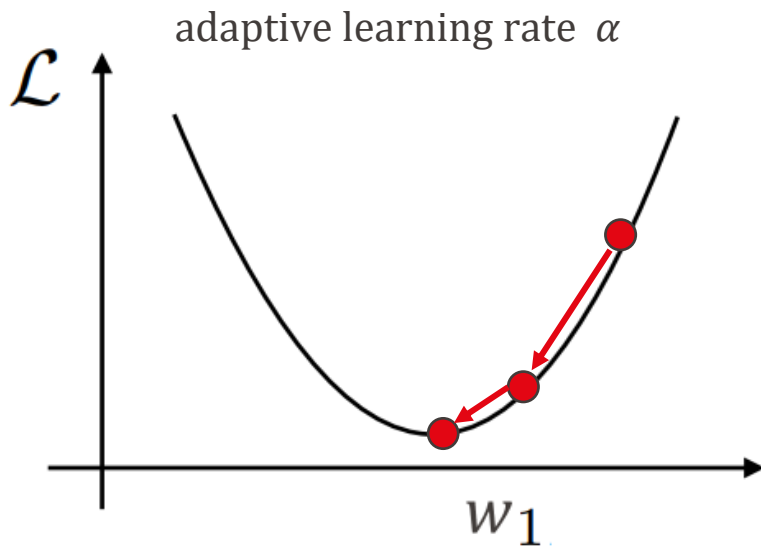
$$\mathcal{D}: (\mathbf{x}^{[1]}, y^{[1]}), (\mathbf{x}^{[2]}, y^{[2]}), \dots, (\mathbf{x}^{[N]}, y^{[N]}) \in \mathfrak{R}^m$$

Mini-batch mode (typically used in DL)

- In between the previous two: with respect to batch settings, the update is done for each “mini-batches”.
- Vectorization
- Less noisy than online-mode & learning faster than batch

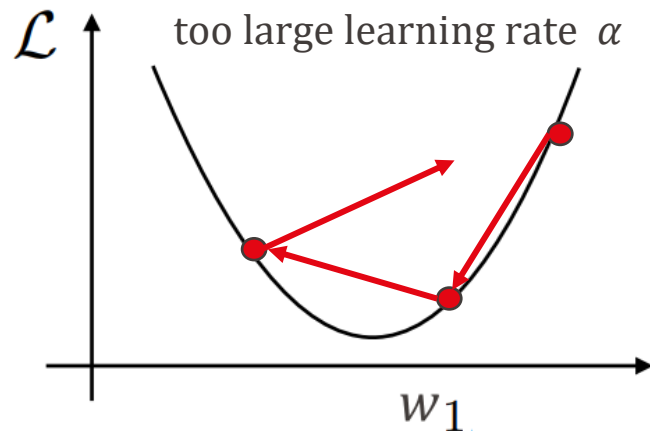
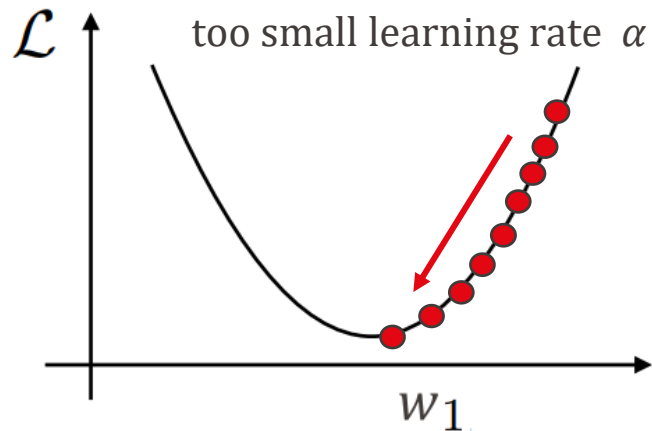
**Other training paradigm:**

- **Stochastic Gradient Descent (SGD)**
- **Batch Normalization (BN)**

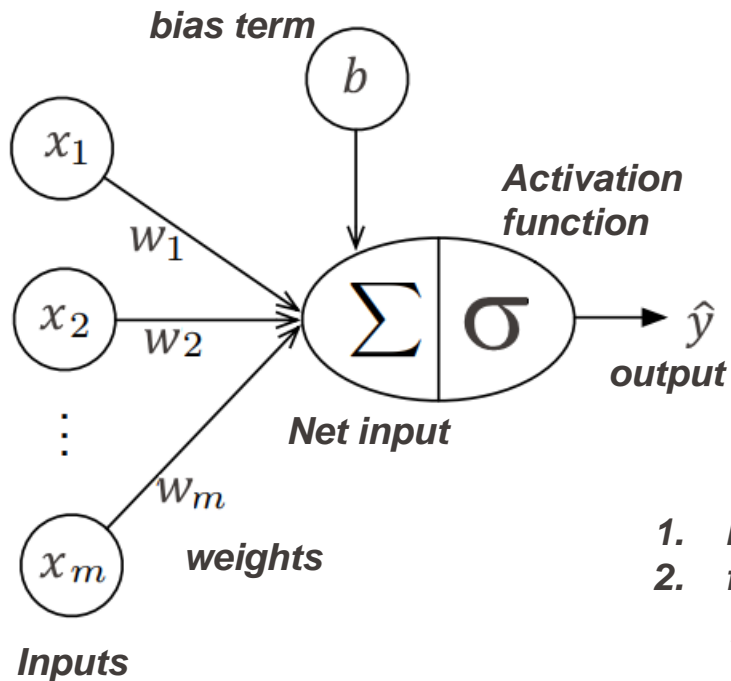


Convex Loss function

$$\mathcal{L}(w, b) = \sum_j \left(\widehat{y}^{[j]} - y^{[j]} \right)^2$$



Fit/Train of ML models...



$$\hat{y} = \sigma \left(\left(\sum_{i=1}^m w_i \cdot x_i \right) + b \right) = \sigma(\mathbf{x}^T \mathbf{w} + b)$$

Given a training set:

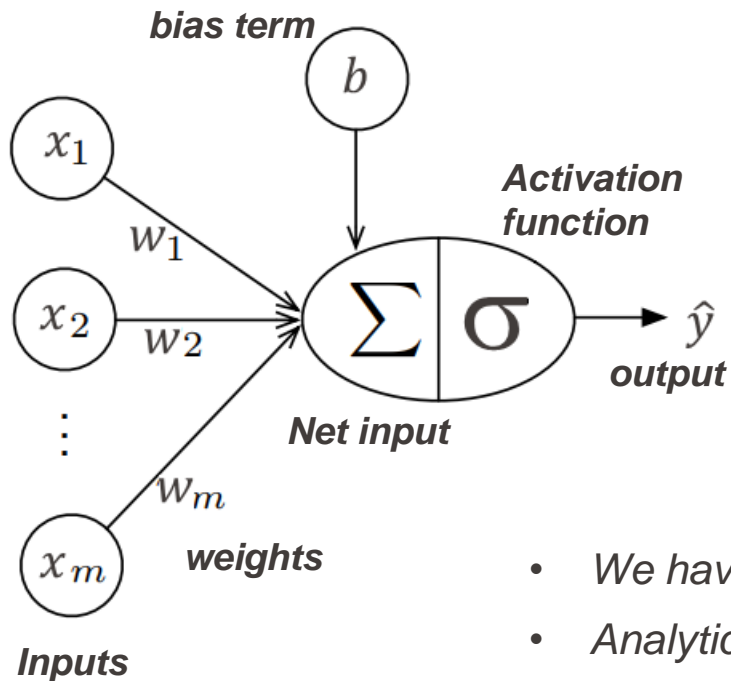
$$\mathcal{D}: (\mathbf{x}^{[1]}, y^{[1]}), (\mathbf{x}^{[2]}, y^{[2]}), \dots, (\mathbf{x}^{[N]}, y^{[N]}) \in \mathbb{R}^m$$

On-line mode with Stochastic Gradient Descent

1. Initialize w, b
2. for every training epoch:
 1. for every $(x^{[j]}, y^{[j]})$ in \mathcal{D} : (or over mini-batches)
 1. $\widehat{y}^{[j]} = \sigma(\mathbf{w}^T \mathbf{x}^{[j]} + b)$ compute prediction
 2. $\nabla_{\mathbf{w}, b} \mathcal{L} = (y^{[j]} - \widehat{y}^{[j]}) \cdot \mathbf{x}^{[j]}$ (with $x^{[0]} = \mathbf{1}$ for b)
 3. $\mathbf{w}, b = \mathbf{w}, b + \alpha \cdot (-\nabla_{\mathbf{w}, b} \mathcal{L})$ (α is the learning rate)

Convex loss function

$$\mathcal{L}(\mathbf{w}, b) = \sum_j (y^{[j]} - \widehat{y}^{[j]})^2$$



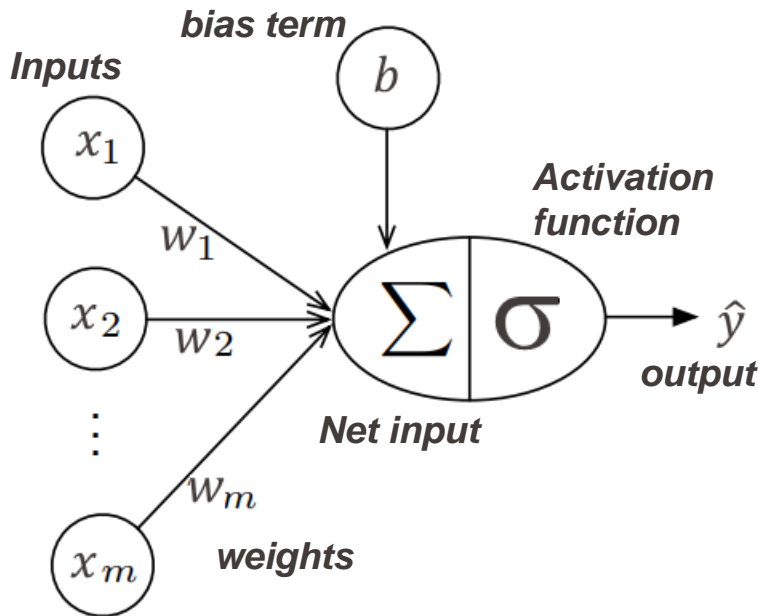
$$\hat{y} = \sigma \left(\left(\sum_{i=1}^m w_i \cdot x_i \right) + b \right) = \sigma(\mathbf{x}^T \mathbf{w} + b)$$

Given a training set:

$$\mathcal{D}: (x^{[1]}, y^{[1]}), (x^{[2]}, y^{[2]}), \dots, (x^{[N]}, y^{[N]}) \in \mathbb{R}^m$$

Optimization problems with Least-Squares

- We have to fit basically a **linear regression model**
- Analytical solution (**normal equation**): $\mathbf{w} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$
- Sometimes **matrix inversion computationally expensive** (large datasets)
- We can **learn** this parameters **iteratively**, fitting neural networks...



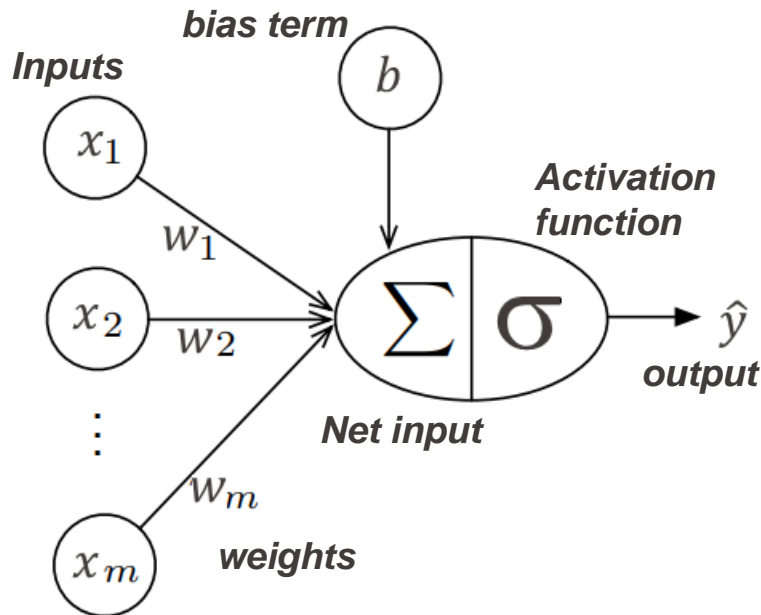
On-line mode with Stochastic Gradient Descent

$$1. \frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial \left(\frac{1}{2N} \sum_j (y^{[j]} - \hat{y}^{[j]})^2 \right)}{\partial w_i}$$

Convex Loss function

$$\mathcal{L}(w, b) = \frac{1}{2N} \sum_j (y^{[j]} - \hat{y}^{[j]})^2 \text{ (MSE)}$$

$$\hat{y} = \sigma(x^T w + b) \text{ (prediction)}$$



Convex Loss function

$$\mathcal{L}(\mathbf{w}, b) = \frac{1}{2N} \sum_j (\hat{y}^{[j]} - y^{[j]})^2 \text{ (MSE)}$$

$$\hat{y} = \sigma(\mathbf{x}^T \mathbf{w} + b) \quad \text{(prediction)}$$

On-line mode with Stochastic Gradient Descent

$$1. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial \left(\frac{1}{2N} \sum_j (\hat{y}^{[j]} - y^{[j]})^2 \right)}{\partial w_i}$$

$$2. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial \left(\frac{1}{n} \sum_{j^2} \frac{1}{2} (\sigma(\mathbf{w}^T \mathbf{x}^{[j]}) - y^{[j]})^2 \right)}{\partial w_i}$$

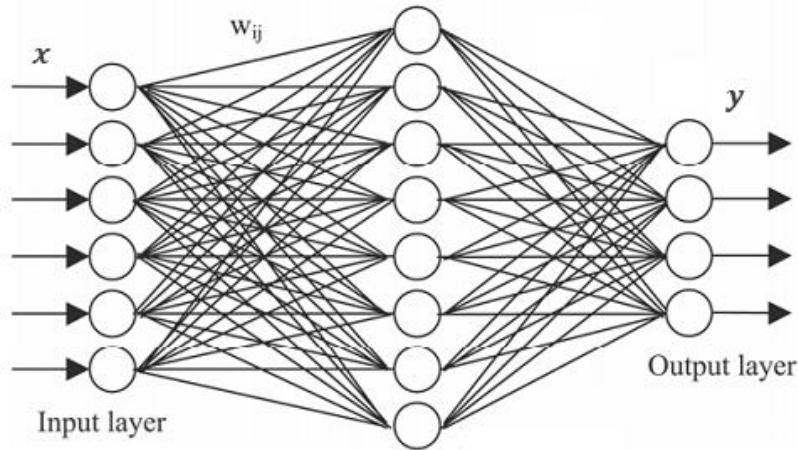
(chain rule for $f(\sigma(h(\mathbf{w})))$)

$$\frac{\partial f(\sigma(h(\mathbf{w})))}{\partial w_i} = \frac{\partial f}{\partial \sigma} \cdot \frac{\partial \sigma}{\partial h} \cdot \frac{\partial h}{\partial w_i} \quad \text{where} \quad \begin{cases} f = (\sigma - y_i)^2 \\ \sigma = I(h) \\ h = \mathbf{w}^T \mathbf{x}^{[j]} \end{cases}$$

$$3. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{1}{n} \sum_j (\sigma(h) - y^{[j]}) \cdot \frac{d\sigma}{dh} \cdot \frac{\partial (\mathbf{w}^T \mathbf{x}^{[j]})}{\partial w_i}$$

$$4. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{1}{n} \sum_j (\sigma(\mathbf{w}^T \mathbf{x}^{[j]}) - y^{[j]}) \cdot x_i^{[j]}$$

Back-propagation (Jacobians)



Convex Loss function

$$\mathcal{L}(\mathbf{w}, b) = \frac{1}{2N} \sum_j \left(\widehat{y}^{[j]} - y^{[j]} \right)^2 \text{ (MSE)}$$

$$\widehat{y} = \sigma(\mathbf{w}^T \mathbf{x} + b) \quad \text{(prediction)}$$

On-line mode with Stochastic Gradient Descent

$$1. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial \left(\frac{1}{2N} \sum_j (\widehat{y}^{[j]} - y^{[j]})^2 \right)}{\partial w_i}$$

$$2. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{\partial \left(\frac{1}{n} \sum_j \frac{1}{2} (\sigma(\mathbf{w}^T \mathbf{x}^{[j]}) - y^{[j]})^2 \right)}{\partial w_i}$$

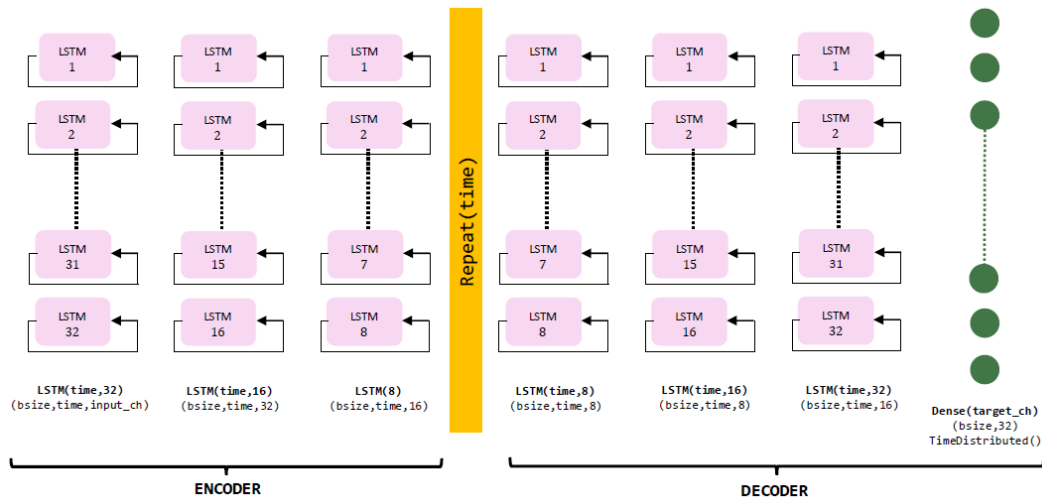
(chain rule for $f(\sigma(h(\mathbf{w})))$)

$$\frac{\partial f(\sigma(h(\mathbf{w})))}{\partial w_i} = \frac{\partial f}{\partial \sigma} \cdot \frac{\partial \sigma}{\partial h} \cdot \frac{\partial h}{\partial w_i} \quad \text{where} \quad \begin{cases} f = (\sigma - y_i)^2 \\ \sigma = I(h) \\ h = \mathbf{w}^T \mathbf{x}^{[j]} \end{cases}$$

$$3. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{1}{n} \sum_j (\sigma(h) - y^{[j]}) \cdot \frac{d\sigma}{dh} \cdot \frac{\partial (\mathbf{w}^T \mathbf{x}^{[j]})}{\partial w_i}$$

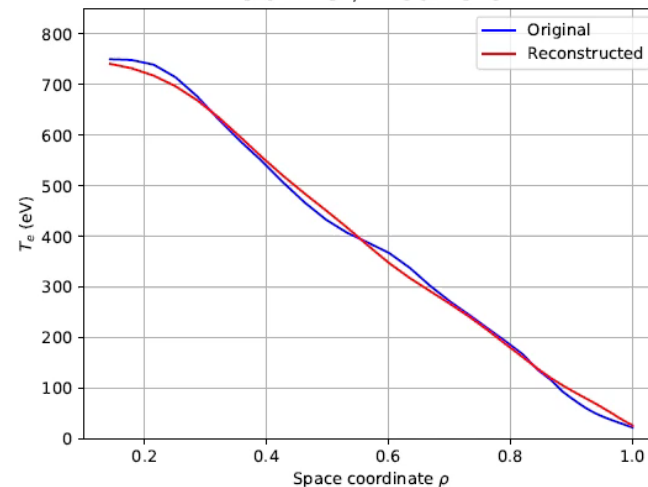
$$4. \quad \frac{\partial \mathcal{L}}{\partial w_i} = \frac{1}{n} \sum_j (\sigma(\mathbf{w}^T \mathbf{x}^{[j]}) - y^{[j]}) \cdot \mathbf{x}_i^{[j]}$$

Reconstruction of T_e profile from TS



REF:[Cristina Venturini: Master Thesis]

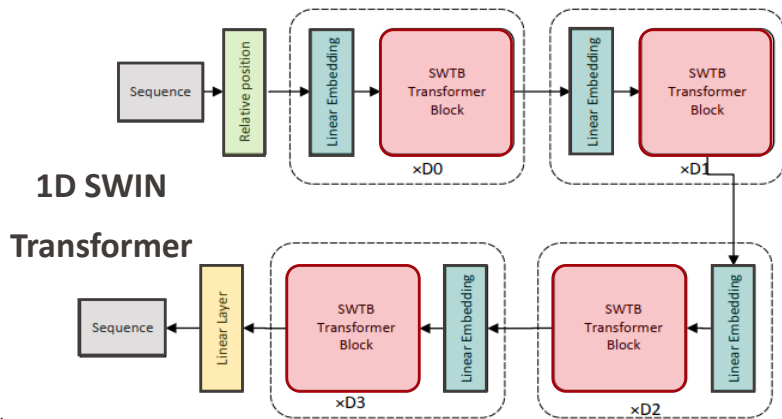
Reconstruction of T_e profile - Model 1da
Shot 71254, Time 0.4232 s



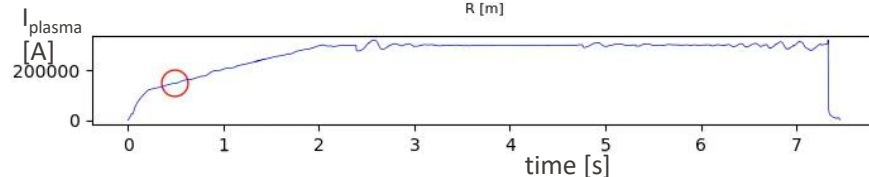
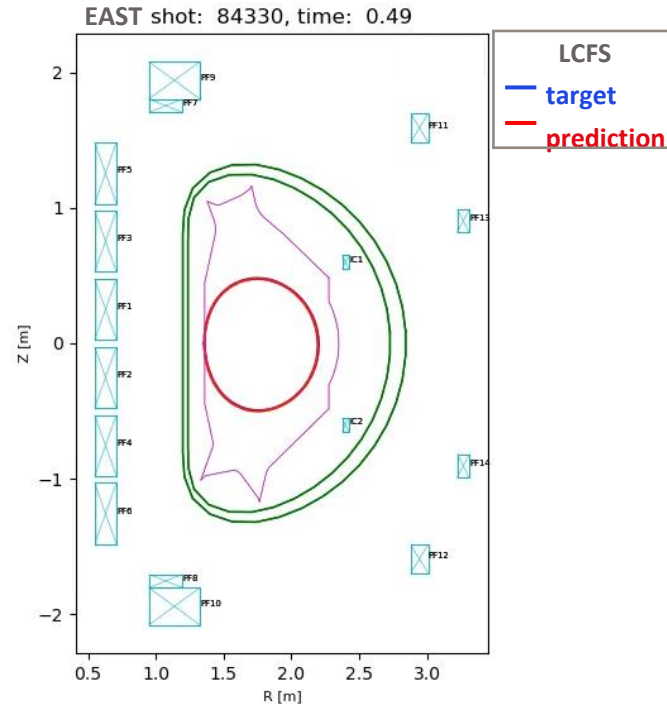
- **Encoder-Decoder**: embedding of high-dimensional space and reconstruction in the original input space
- Learning of **short & long-range dependencies** of the temporal dynamics of the profile evolution;

Magnetic equilibrium reconstruction (LCFS)

- magnetic equilibrium reconstruction:
 - complex time-varying, non-linear, **multi-scale**...
 - Modeling **sequences** with large variations in the time scales... **“attention is all you need”**!
 - ...**one-step ahead** prediction of the magnetic field evolution in time (Last Closed Flux Surface)



REF: [C. Wan, A. Pau, O Sauter et al 2022 (in review)]



A probabilistic perspective

- We call **inference**^(*) the procedure with which we quantify of the uncertainty or confidence in the estimate $\hat{\theta}$.
$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \mathcal{L}(\mathcal{D}|\theta)$$
- On a probabilistic perspective we reason in terms of **Probability Density Estimation** for the joint probability distribution of our dataset \mathcal{D} (a sample from the population)
- Under **i.i.d assumption** (training examples sampled independently and identically from the population representing the input domain \mathcal{D}):

$$p(\mathcal{D}|\theta) = \prod_{n=1}^N p(y_n|x_n, \theta) \quad LL(\mathcal{D}|\theta) \triangleq \log p(\mathcal{D}|\theta) = \sum_{n=1}^N \log p(y_n|x_n, \theta)$$

- Therefore the optimization problem can be seen as maximizing a probability

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} p(\mathcal{D}|\theta)$$

- **inference**^(*) in the deep learning community refers to predicting $p(y_n|x_n, \hat{\theta})$

- Therefore the optimization problem translates in maximizing the **Log-Likelihood (LL)**,

$$LL(\mathcal{D}|\boldsymbol{\theta}) \triangleq \log p(\mathcal{D}|\boldsymbol{\theta}) = \sum_{n=1}^N p(y_n|x_n, \boldsymbol{\theta}) \quad \hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} LL(\mathcal{D}|\boldsymbol{\theta})$$

- Which can be also seen as minimizing the **Negative Log-Likelihood (NLL)**:

$$NLL(\mathcal{D}|\boldsymbol{\theta}) \triangleq -\log p(\mathcal{D}|\boldsymbol{\theta}) = -\sum_{n=1}^N p(y_n|x_n, \boldsymbol{\theta}) \quad \hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} NLL(\mathcal{D}|\boldsymbol{\theta})$$

- Estimating the probability density function (high uncertainty if the sampling distribution is small) is usually done with two common approaches:
 - **Maximum Likelihood Estimation (MLE)**
 - **Maximum a Posteriori (MAP):**

- **Maximum Likelihood Estimation (MLE):** frequentist approach for estimating the set of parameters $\hat{\theta}$ of a model by finding the values that maximize the log-likelihood $LL(\mathcal{D}|\theta)$.
- **Interpretation:** $LL(\mathcal{D}|\theta)$ describes the probability of observing the data given the model parameters $\hat{\theta}$. The likelihood function is known if data are **i.i.d.** $\hat{\theta}$ resulting from MLE are the most probable values given the data.
- **Maximum a Posteriori (MAP):** Bayesian approach for estimating the values of the parameters $\hat{\theta}$ that maximize the posterior probability,
- **Interpretation:** MAP describes probability of the parameters given the data, and allows incorporating **prior knowledge** about the parameters into the estimation process. This prior knowledge is specified as a probability distribution and allows us to account for uncertainty in the data.

- **Bayes' Theorem** that describes how to update the probability of a event (or hypothesis) based on new evidence or information.
- What do we mean with **Bayesian Inference**?

Given a dataset:

$$\mathcal{D}: (\mathbf{x}^{[1]}, y^{[1]}), (\mathbf{x}^{[2]}, y^{[2]}), \dots, (\mathbf{x}^{[N]}, y^{[N]})$$

Posterior

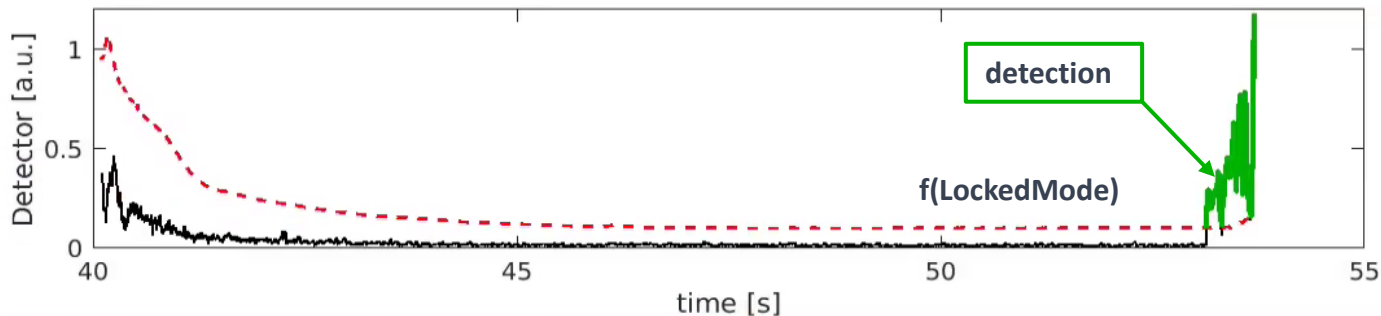
$$P(\boldsymbol{\theta}|\mathcal{D}) = \frac{\overset{\text{Likelihood}}{P(\mathcal{D}|\boldsymbol{\theta})} \cdot \overset{\text{Prior}}{P(\boldsymbol{\theta})}}{\underset{\text{Evidence or Marginal}}{P(\mathcal{D})}}$$

$\boldsymbol{\theta}$ is an unknown random variable

$$= \frac{P(\mathcal{D}|\boldsymbol{\theta}) \cdot P(\boldsymbol{\theta})}{\int P(\mathcal{D}, \boldsymbol{\theta}') p(\boldsymbol{\theta}') d\boldsymbol{\theta}'}$$

- The **posterior** gives an indication of the **uncertainty** about our fitting parameter $\boldsymbol{\theta}$ given the data \mathcal{D} , according to the **prior knowledge** we have.
- Extremely powerful for **online learning**:
 - $P(\boldsymbol{\theta}|\mathcal{D}_{1:t}) \propto P(\mathcal{D}_{1:t}|\boldsymbol{\theta}) \cdot P(\boldsymbol{\theta}|\mathcal{D}_{1:t-1})$

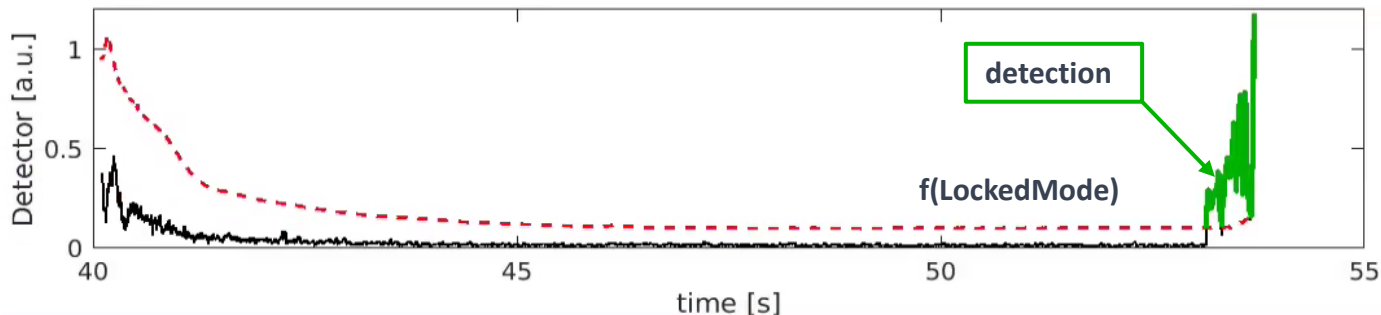
Bayesian inference: event-detection problem



- We have a RT-detector for **Locked Modes** (LM - common disruption precursor):
- The detector works very well:
 - It has an **accuracy** of **99%** (correct detection when there actually is a LM)
 - It has a very low **false positive** rate **0.1%**
 - In our sampling distribution **2%** of the discharges exhibits a Locked Mode

$$P(\theta|\mathcal{D}) = \frac{P(\mathcal{D}|\theta) \cdot P(\theta)}{P(\mathcal{D})}$$

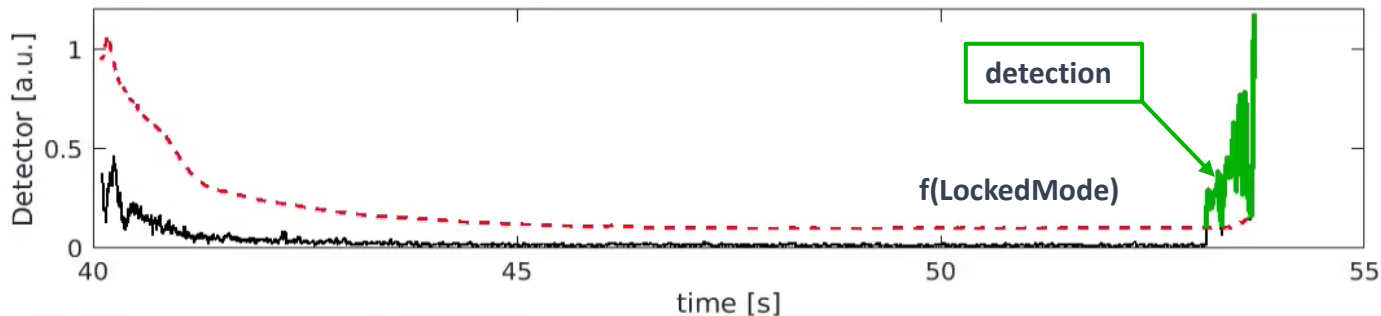
- Consider the case where we run a discharge and the locked mode detector triggers an alarm.
 - **What is the probability that there was actually a locked mode?**



- We have a RT-detector for **Locked Modes** (LM - common disruption precursor:
- The detector works very well:
 - It has an **accuracy** of **99%** (correct detection when there actually is a LM)
 - It has a very low **false positive** rate **0.1%**
 - In our sampling distribution **2%** of the discharges exhibits a Locked Mode

$$P(LM|detect) = \frac{0.99 \quad 0.02}{P(detect)} = \frac{P(detect|LM) \cdot P(LM)}{P(detect)} \quad ?\%$$

- What is the probability that there was actually a locked mode?



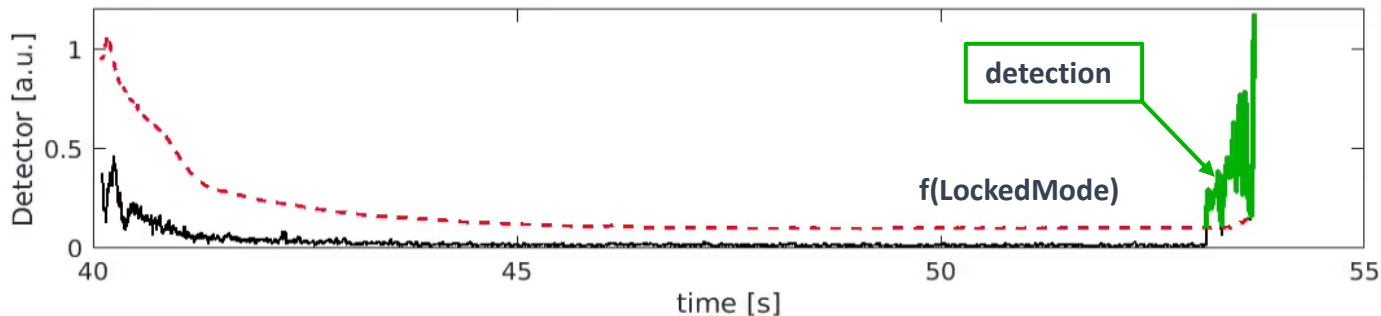
- We have a RT-detector for **Locked Modes** (LM - common disruption precursor:
- The detector works very well:
 - It has an **accuracy** of **99%** (correct detection when there actually is a LM)
 - It has a very low **false positive** rate **0.1%**
 - In our sampling distribution **2%** of the discharges exhibits a Locked Mode

$$P(LM|detect) = \frac{0.99 \quad 0.02}{P(detect)} = \frac{P(detect|LM) \cdot P(LM)}{P(detect)} \quad ?\%$$

$$P(detect) = P(detect | LM) \cdot P(LM) + P(detect | \sim LM) \cdot P(\sim LM)$$

- What is the probability that there was actually a locked mode?

Bayesian inference: event-detection problem



- We have a RT-detector for **Locked Modes** (LM - common disruption precursor:
- The detector works very well:
 - It has an **accuracy** of **99%** (correct detection when there actually is a LM)
 - It has a very low **false positive** rate **0.1%**
 - In our sampling distribution **2%** of the discharges exhibits a Locked Mode

$$P(LM|detect) = \frac{0.99 \quad 0.02}{P(detect)} = \frac{P(detect|LM) \cdot P(LM)}{P(detect)} \quad ?\%$$

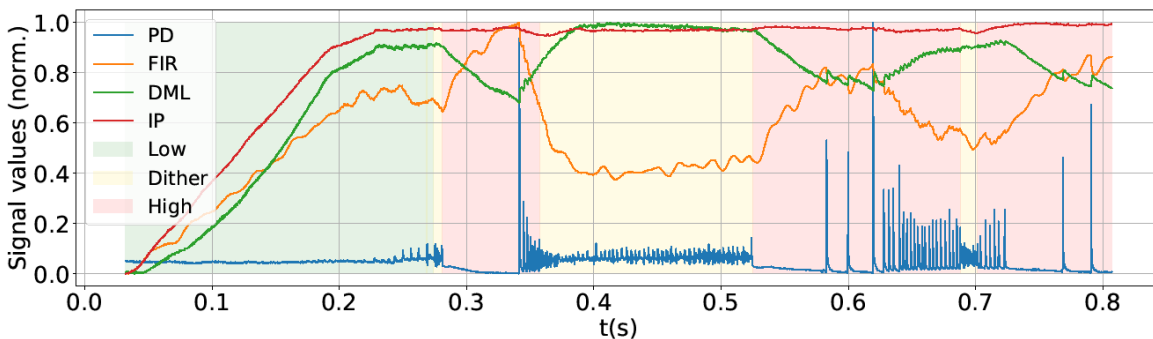
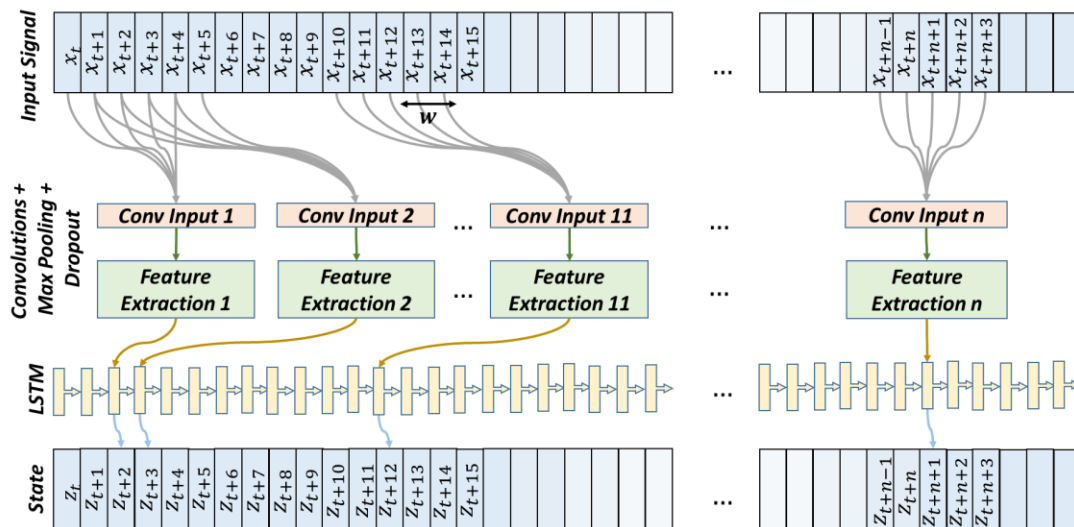
$$P(detect) = P(detect|LM) \cdot P(LM) + P(detect|\sim LM) \cdot P(\sim LM)$$

$$P(LM|detect) = \frac{0.99 \quad 0.02}{0.0208} = \sim 0.953$$

- What is the probability that there was actually a locked mode? **95%**

Event Detection & Disruption prevention...

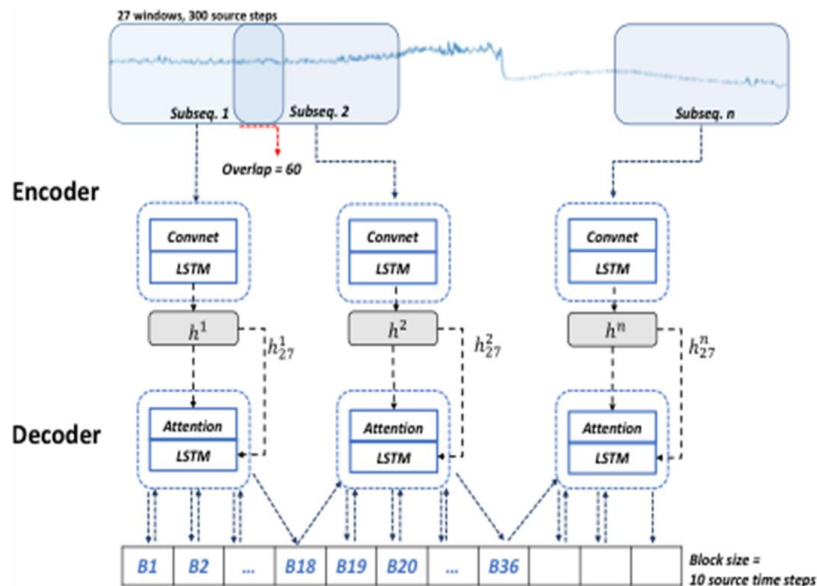
- **Deep Learning** model based on a **convolutional-RNN (LSTM)**
- **Probability** of the plasma of being in a given **confinement state** (accounting for temporal evolution)
- **RT implementation** (nice example of integration with physics-based models in the framework of **off-normal events handling & disruption avoidance**)



REF: [Matos et al NF 2020]

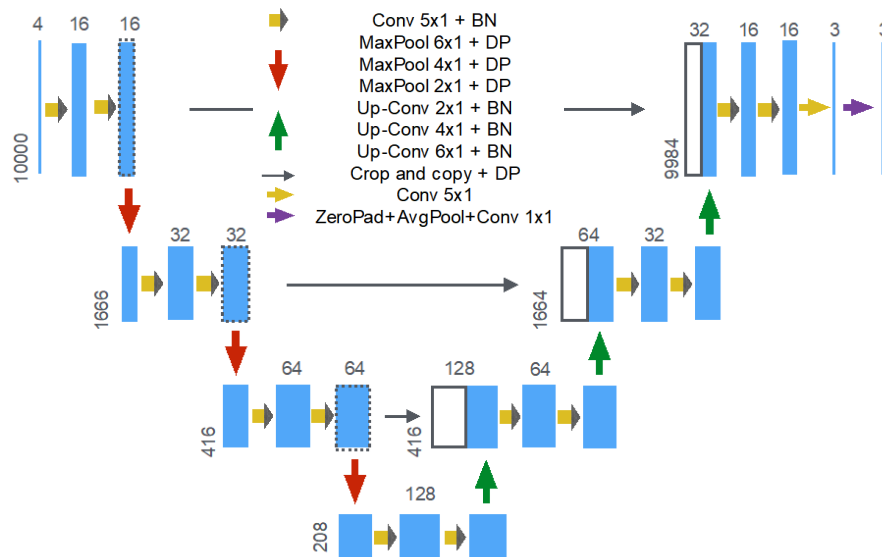
SEQUENCE 2 SEQUENCE MODEL:

- Model not constrained to have same **source/target resolutions**.
- Decoder was extended with an **attention** layer to capture **larger context** of long input sequences.



UTIME MODEL:

- Multi-scale convolutional** structure allows to capture **patterns at different scales** present in the plasma.
- processing the whole signal at once (offline) with the ability to see at a **wider contextual information**.



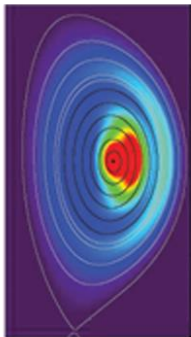
REF: [Marceca et al NEURIPS 2021]

REF: [Matos et al NF 2021]

Disruptions: a complex physics picture

REF: [P. de Vries et al. *NF* 2011]

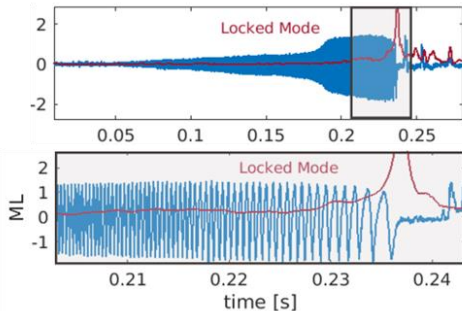
Impurity
Accumulation
(radiation)



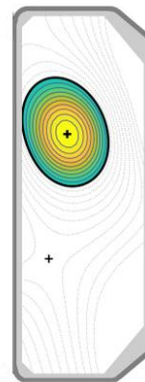
Tearing Modes
Beta Limit
Radiative Collapse
Auxiliary Shutdown Low q
Edge Cooling
NTM MARFE
Locked Mode
Impurity Accumulation
UFO Error Fields Locked Mode
Density limit
Impurity Influx
Technical Failure
VDE

- Wide plethora of **physics** and **technical** cause (systems failure, etc.)
- ... **loss of control** and potential **damages** to the machine.

Locked
modes
(MHD)



Vertical
Displacement
Events

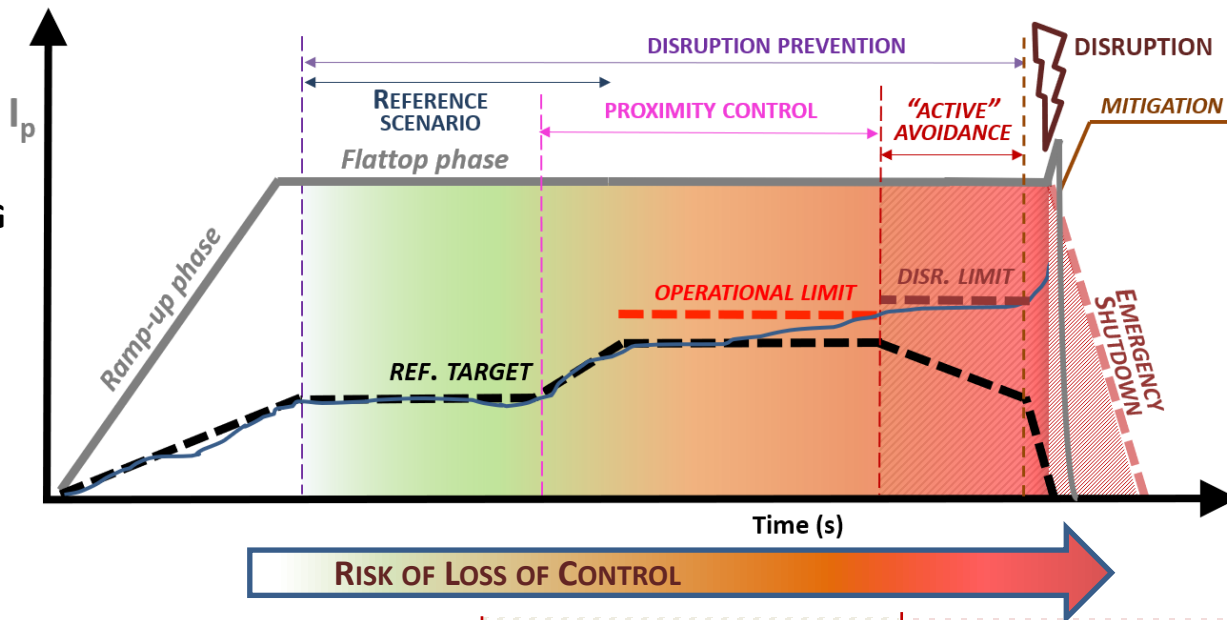


Disruption prevention in a nutshell

PLASMA STATE MONITORING & DISRUPTION PREVENTION

a key requirement for plasma control throughout the entire evolution of the discharge

REF: [A. Pau et al EPS 2022]



REF. SCENARIO

- keep the target scenario stable against disturbances (ST, ELM, MHD modes, etc.)

PROXIMITY CONTROL

- keep stability while pushing performance by regulating proximity to stability & controllability boundaries

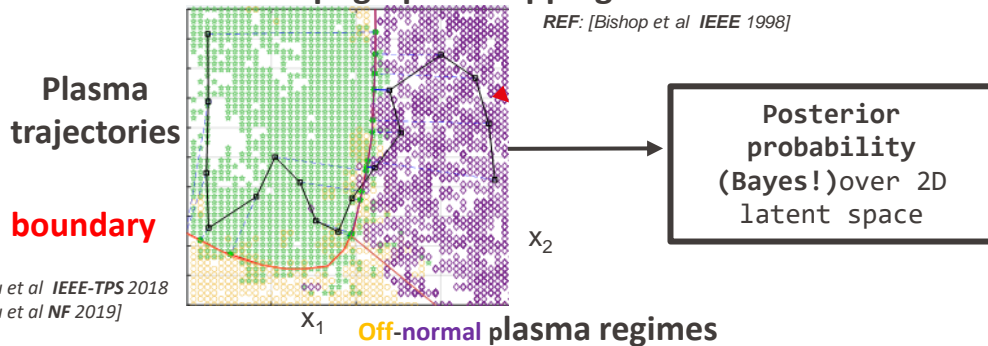
ACTIVE AVOIDANCE

- asynchronous response when crossing operational boundaries (danger levels)

EMERGENCY SHUTDOWN

- Fast controlled shutdown
- mitigation

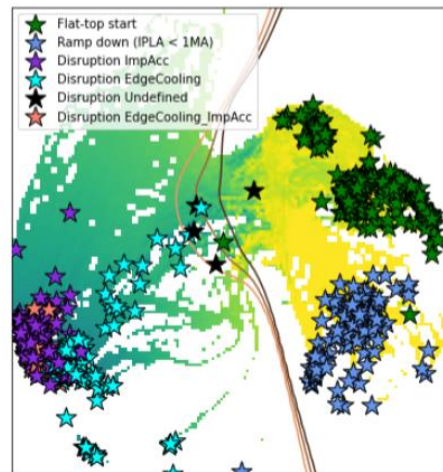
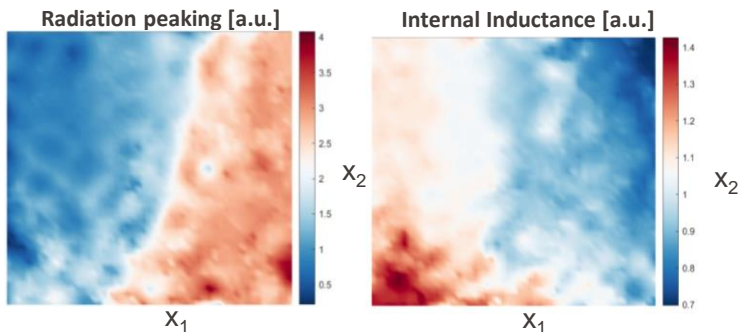
Generative Topographic Mapping



(towards model transparency and interpretability)

- Trajectories over a low-dimensional latent space encode different disruption paths
- Clustering of plasma regions with different risk of disruptions to unveil boundaries

2D Manifold key physics quantities distributions



REF: [Andrin Bürli: Master Thesis]

- **Learning From Data** by Yaser Abu-Mostafa
- **Pattern Recognition and Machine Learning** by Christopher M Bishop
- **Bayesian Reasoning and Machine Learning** by David Barber
- **Machine Learning: A Probabilistic Perspective** by Kevin P. Murphy
- **The Elements of Statistical Learning** by T. Hastie, R. Tibshirani, J. Friedman