GPYTORCH: BLACKBOX MATRIX-MATRIX GAUSSIAN PROCESS INFERENCE WITH GPU ACCELERATION

Jacob R. Gardner, **Geoff Pleiss**,
David Bindel, Kilian Q. Weinberger, Andrew Gordon Wilson

Cornell University

DEEP LEARNING & RAPID PROTOTYPING

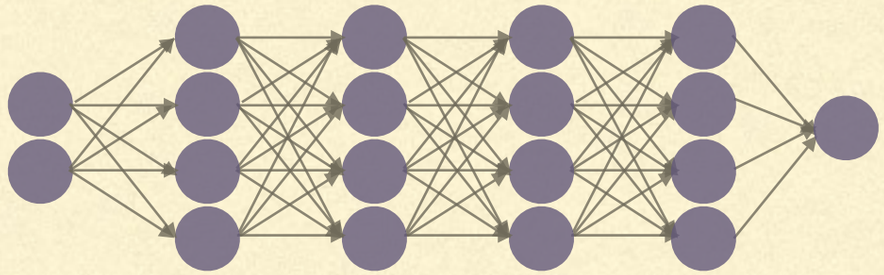

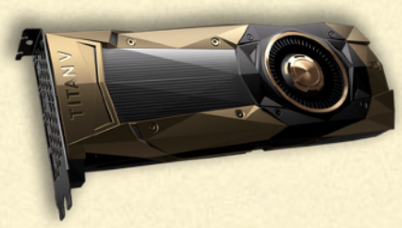

1. GPU acceleration

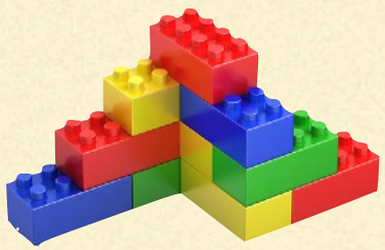

2. Modular design

3. Open source software

DEEP LEARNING & RAPID PROTOTYPING

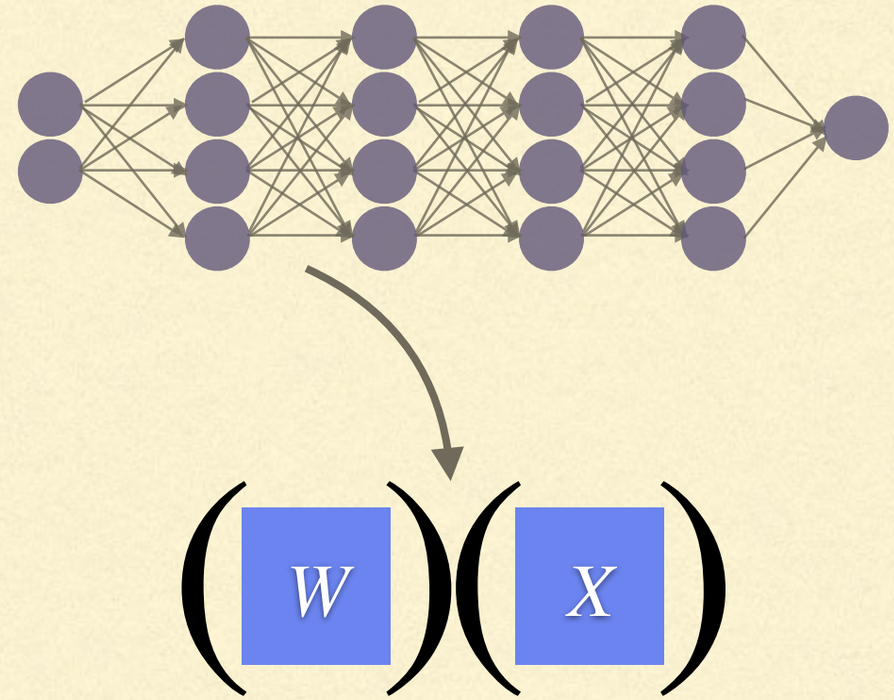

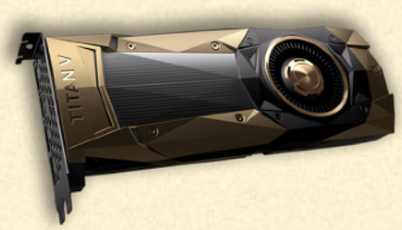

1. GPU acceleration

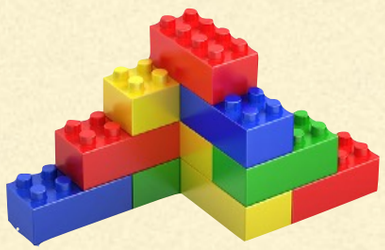

2. Modular design

3. Open source software

DEEP LEARNING & RAPID PROTOTYPING

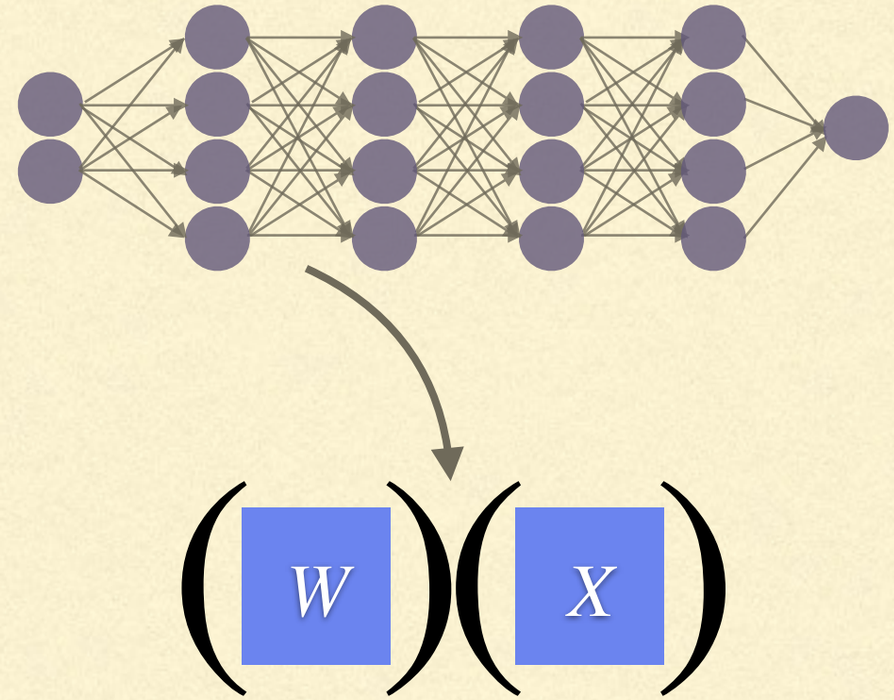

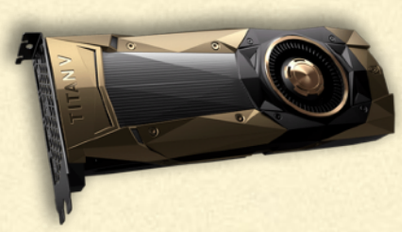

1. GPU acceleration

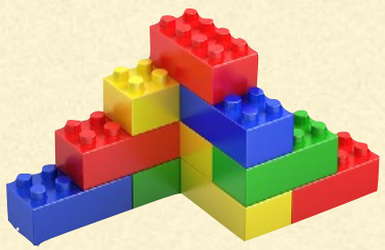

2. Modular design

```
model.add_module(
    nn.Linear(64, 128)
)
```

3. Open source software

DEEP LEARNING & RAPID PROTOTYPING

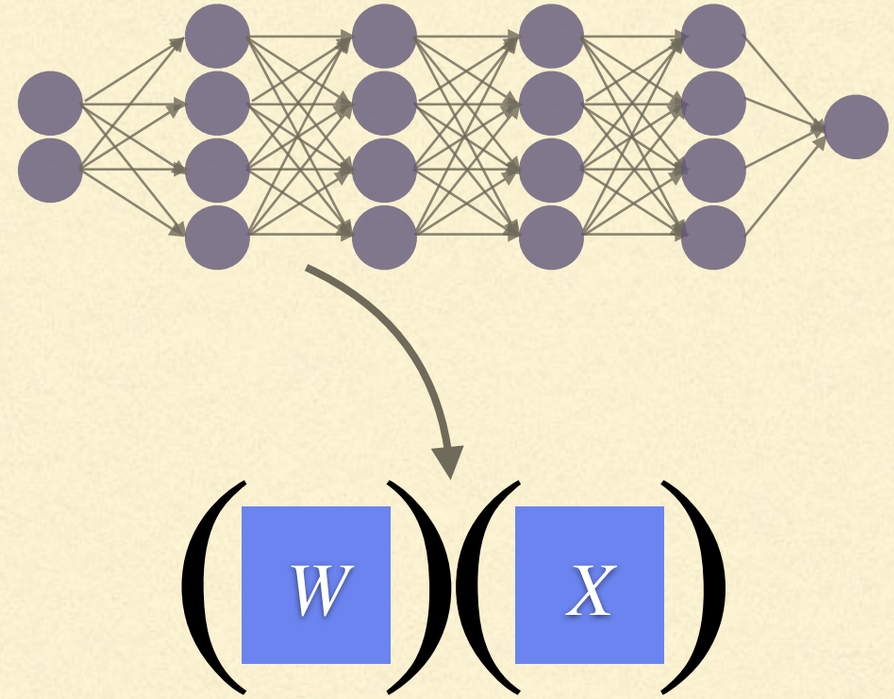

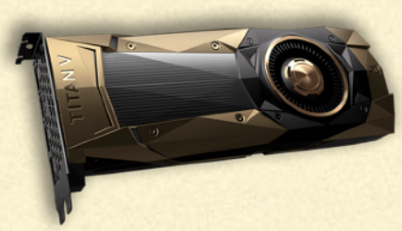

1. GPU acceleration

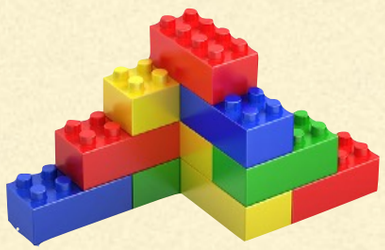

2. Modular design

```
model.add_module(
    nn.Linear(64, 128)
)
```

3. Open source software

GAUSSIAN PROCESSES (GPS) & RAPID PROTOTPYING?

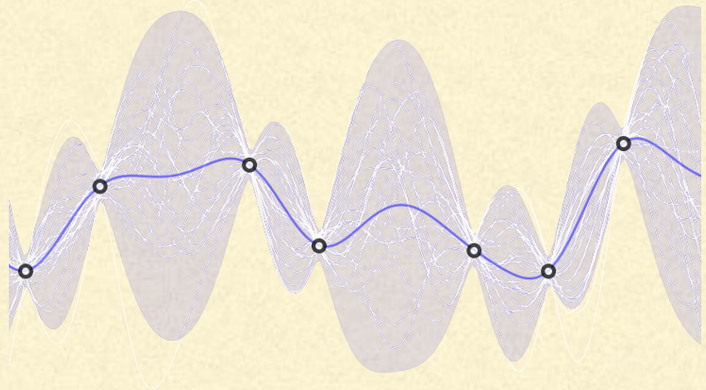

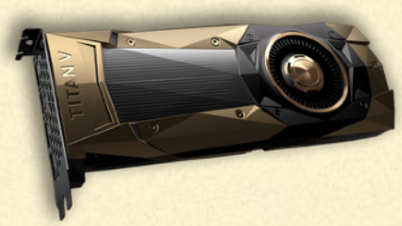

1. GPU acceleration?

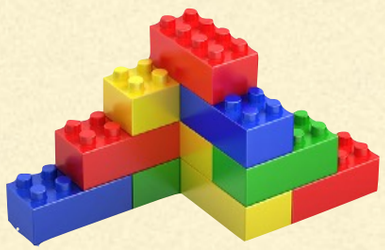

2. Modular design?

GAUSSIAN PROCESSES (GPS) & RAPID PROTOTPYING?

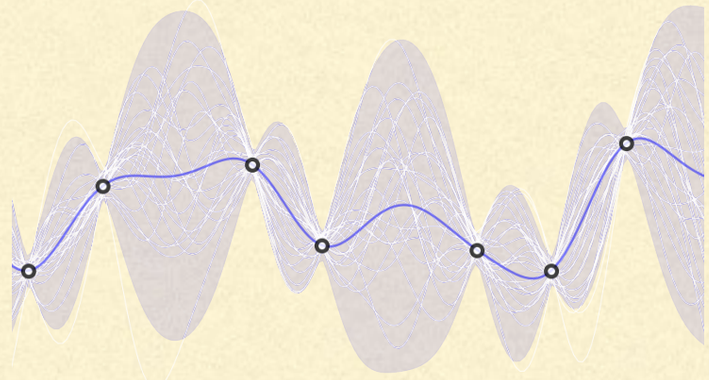

data covariance matrix

$$\log\det(K) \qquad K^{-1}y$$

1. GPU acceleration?

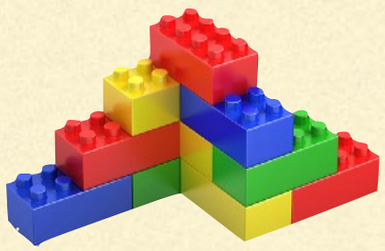

2. Modular design?

GAUSSIAN PROCESSES (GPS) & RAPID PROTOTPYING?

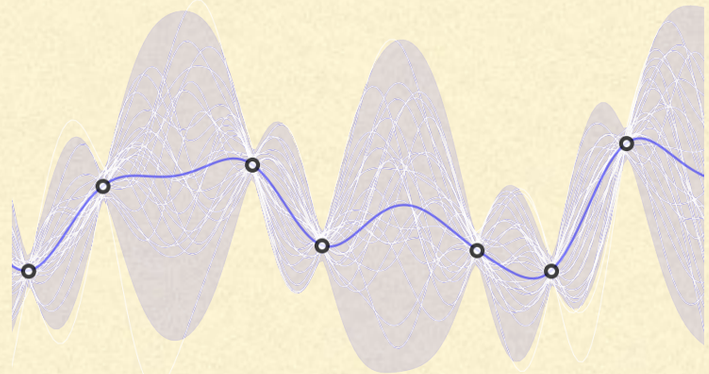

data covariance matrix

$$\log\det(K) \qquad K^{-1}y$$

1. GPU acceleration?

2. Modular design?

~~$(K + \sigma^2 I)^{-1}$~~

$$\sigma^{-2}I - \sigma^{-4}K_{XU}L^{-\top}\left(I + L^{-1}K_{XU}K_{XU}L^{-\top}\sigma^{-2}\right)^{-1}L^{-1}K_{UX}$$

BLACKBOX MATRIX-MATRIX (BBMM)INFERENCE

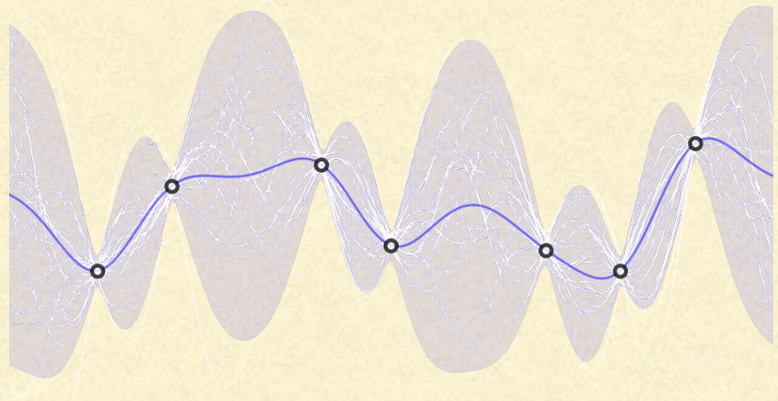

BLACKBOX MATRIX-MATRIX (BBMM)INFERENCE

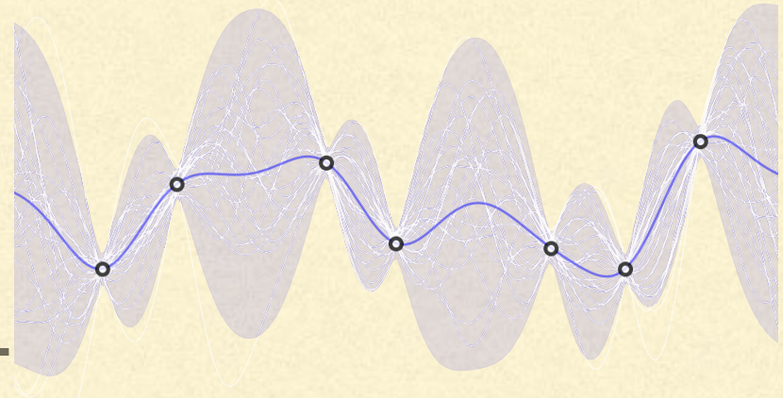

data covariance matrix

$$\left(\quad K \quad\right)\left(\quad \cdot \quad\right)$$

(Modified) batch conjugate gradients

$$\begin{bmatrix} \log\det(K) \\ K^{-1}y \end{bmatrix}$$

BLACKBOX MATRIX-MATRIX (BBMM)INFERENCE

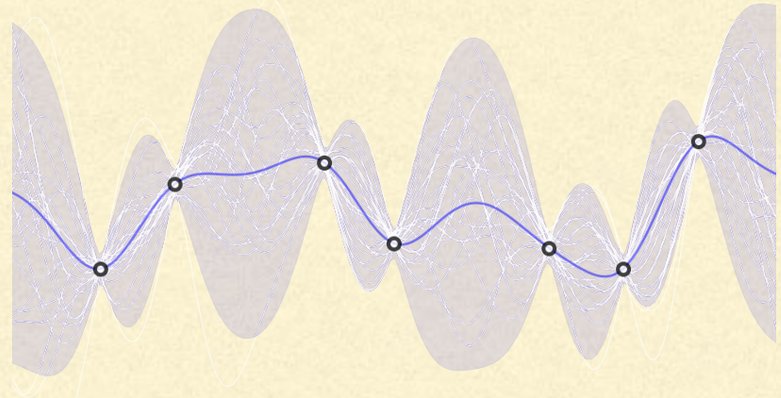

Matrix-matrix multiplication w/ K

(Modified) batch conjugate gradients

$$\begin{bmatrix} \log\det(K) \\ K^{-1}y \end{bmatrix}$$

BLACKBOX MATRIX-MATRIX (BBMM)INFERENCE

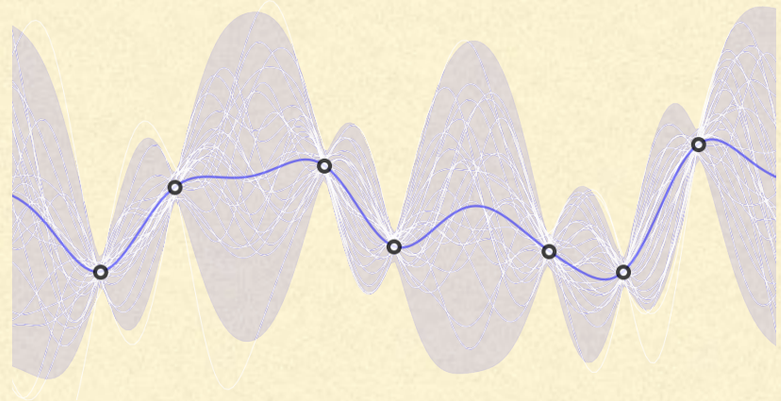

Matrix-matrix multiplication w/ K

(Modified) batch conjugate gradients

$$\begin{bmatrix} \log\det(K) \\ K^{-1}y \end{bmatrix}$$

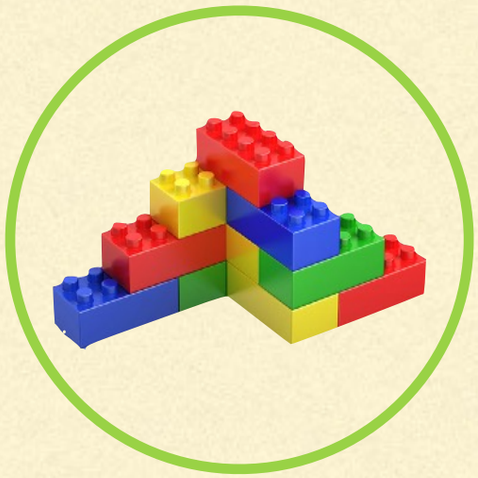

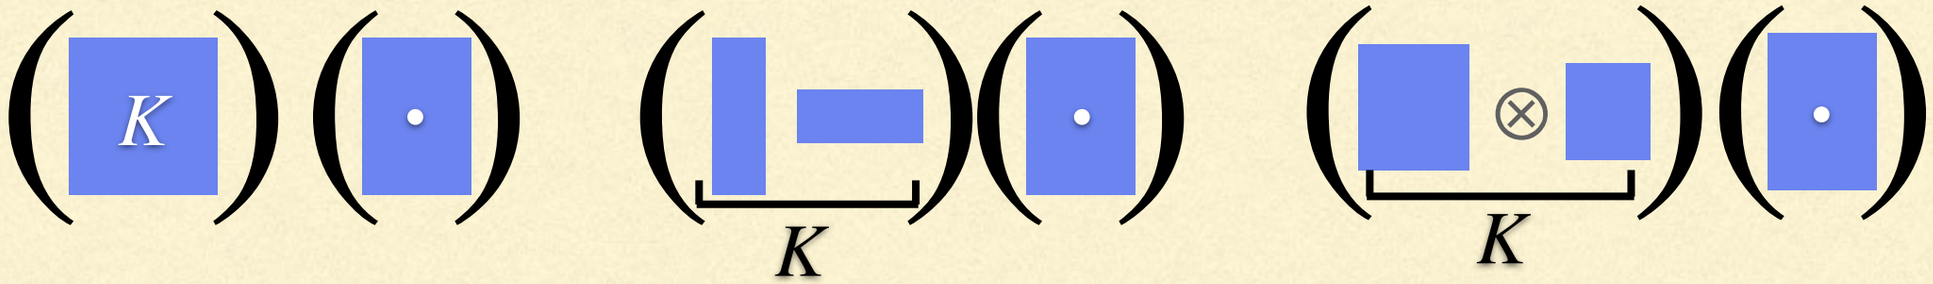

Std. GP

Scalable GP

Multi-output GP

~~$O(n^3)$~~ $O(n^2)$

GPYTORCH

A BBMM LIBRARY

GPYTORCH
A BBMM LIBRARY

BBMM (Modified CG)

```
In [ ]: train()
```

Standard (Cholesky)

```
In [ ]: train()
```

GPYTORCH
A BBMM LIBRARY

GPS WITH BBMM INFERENCE

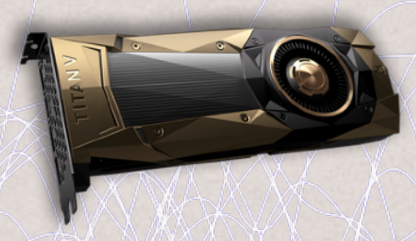

1. GPU acceleration

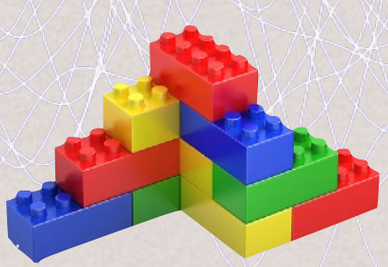

2. Modular design

3. Open source

GPyTorch

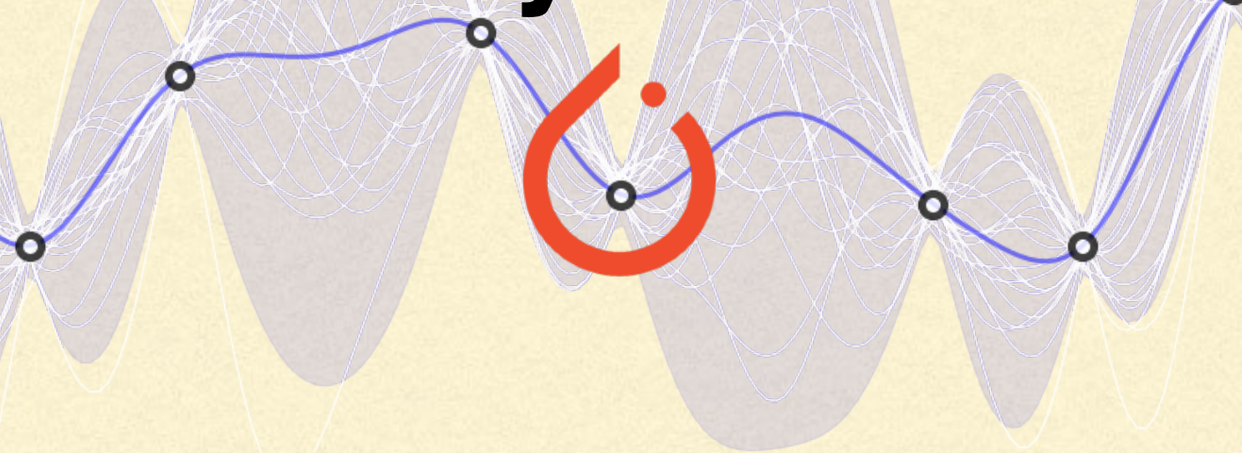

Poster #13

http://gpytorch.ai