Trust Your Robots!
Predictive Uncertainty Estimation of Neural Networks with Sparse Gaussian Processes

A. Research Questions

A method for Neural Networks:

- Fast uncertainty estimation – running on a real robotic system.
- Reliable uncertainty estimation as good as Gaussian Processes.

This will enable robotic introspection for improving robustness and safety!

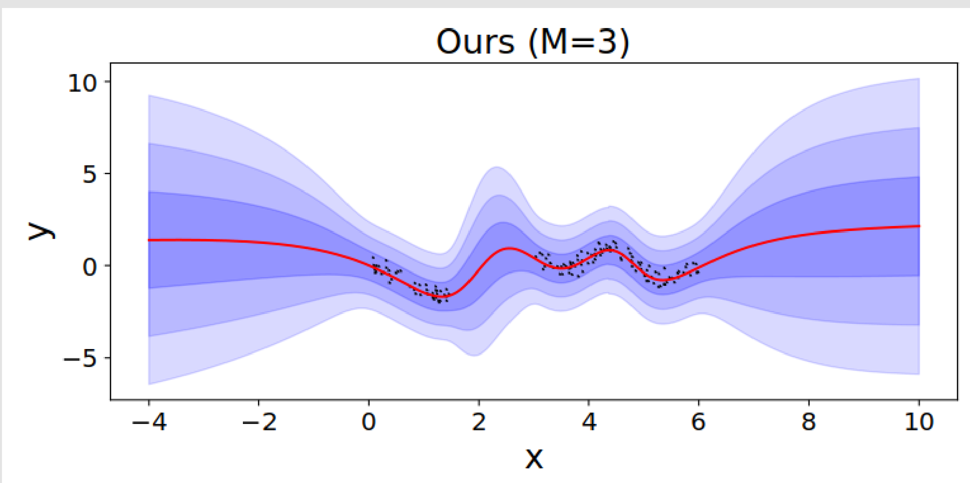

B. Our Main Idea

We propose a predictive model:

- Neural Networks for accurate most-likely predictions.
- Sparse Gaussian Processes for reliable uncertainty estimates.

See figure below for illustrations

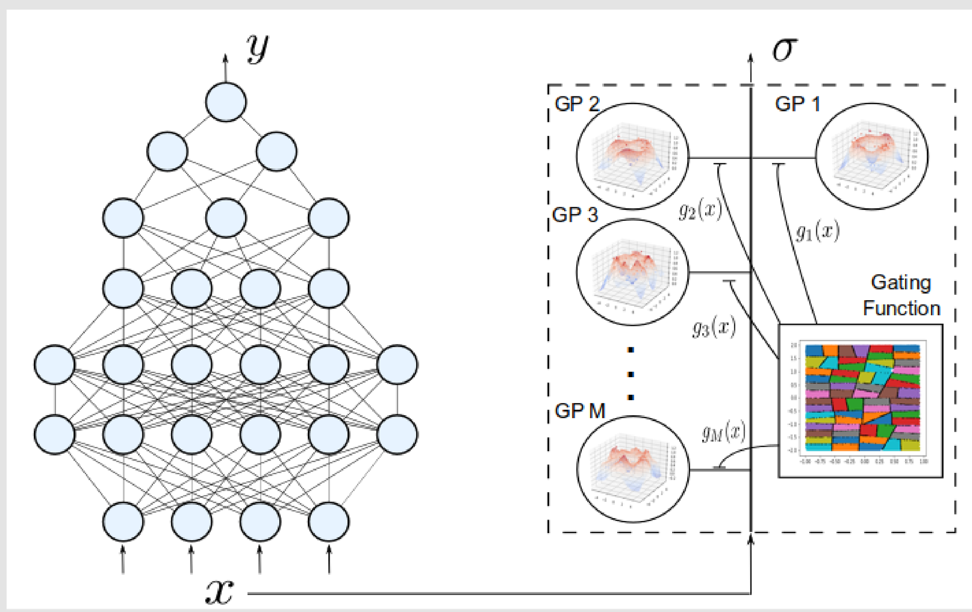

C. The Proposed Method

- Linearize Neural Networks:

$$\tilde{y} = J_f(x)\theta + \epsilon$$

In a function space view, we can view this as a GP with the Neural Tangent Kernel.

- Casting into sparse GPs:

$$\tilde{y} = \sum_{m=1}^{M} g_m(x)\widetilde{f_m}(x) + \epsilon_m$$

$$\tilde{f}_m(\mathrm{x}) \sim \mathrm{GP}(0, \frac{1}{\delta_m} \mathrm{J}^{\mathrm{T}}_{\mathrm{f_m}}(\mathrm{x})\mathrm{J}_{\mathrm{f_m}}(\mathrm{x}))$$

Divide and conquer principle to improve scalability of GPs.

In this mixtures of experts model:

- Generative model - scalability:

$$\begin{bmatrix}\tilde{y}_m \\ \tilde{f}_m\end{bmatrix} \sim N\left(0, \begin{bmatrix} K_m + \sigma_0 I & k_{*,m} \\ k_{m,*} & k_{m,**}\end{bmatrix}\right)$$

- Predictive in Regression – closed form:

$$\Sigma_m = k_{m,**} - k_{m,*}^T (K_m + \sigma_0 I)^{-1} k_{m,*} + \sigma_0$$

- Predictive Classification – closed form:

$$p(c|z_m) = \left(\frac{z_m}{\sqrt{1 + \lambda_{m,0}\Sigma_m(x^*)}}\right)$$

E. Main Results

The selected highlights are:

- Scalability tested upto approx. two million data points, and support for distributed training.

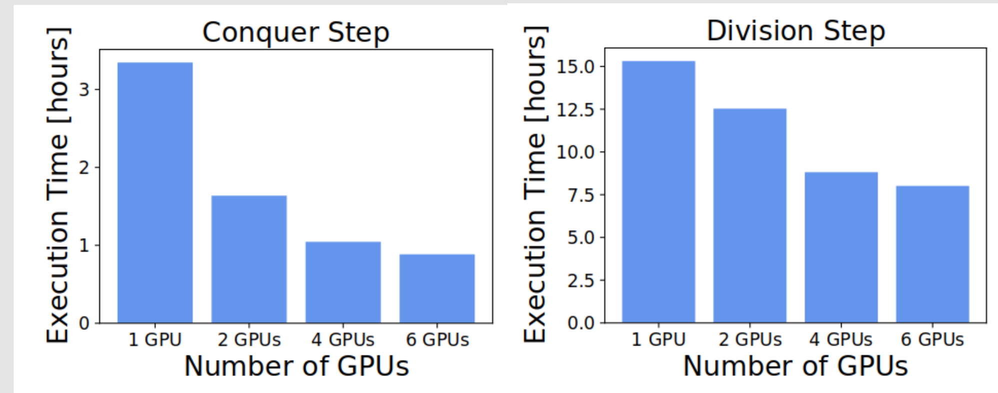

- Real robot test: performance and run-time.

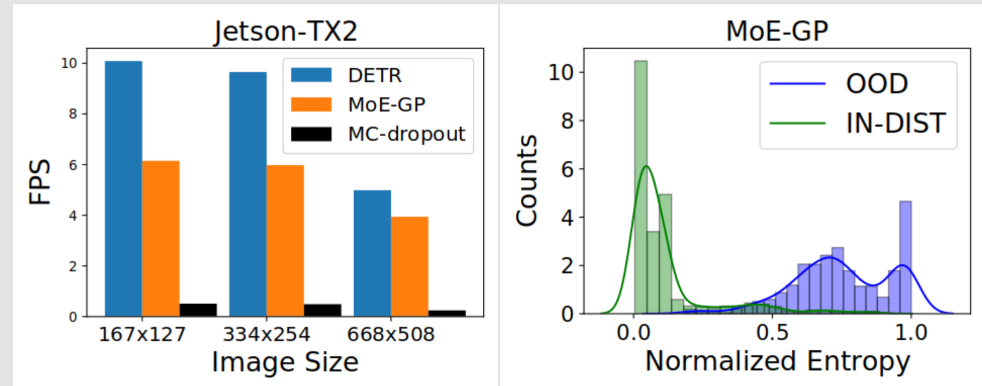

- Theory that connects neural networks to mixtures of GP experts. Ablation studies and comparisons in the paper!

F. Key Take-Aways

Summary of our contributions are:

- A method to estimate predictive uncertainty of neural networks with sparse Gaussian Processes.
- Theory that connects neural networks to sparse Gaussian Processes -> practical algorithm.

paper video code

Deutsches Zentrum für Luft- und Raumfahrt DLR
Karlsruhe Institute of Technology
High Performance Humanoid Technologies

By Jongseok Lee, Jianxiang Feng, Matthias Humt, Marcus G Müller and Rudolph Triebel