Maximum Likelihood Uncertainty Estimation: Robustness to Outliers

Safe AI 2022 (The AAAI's Workshop on Artificial Intelligence Safety)

Deebul S. Nair[1,3], Nico Hochgeschwender[1,2], Miguel A. Olivares-Mendez[3]

1 Dept. of Computer Science, Bonn-Rhein-Sieg University of Applied Sciences, Germany.
2 German Aerospace Center (DLR), Institute for Software Technology, Germany.
3 SpaceR Research Group, University of Luxembourg, Luxembourg.

Hochschule Bonn-Rhein-Sieg University of Applied Sciences

b-it

uni.lu UNIVERSITÉ DU LUXEMBOURG

Introduction

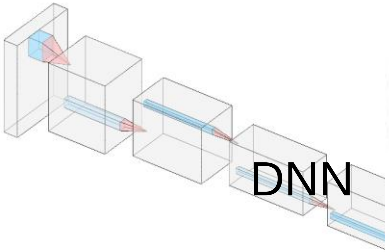

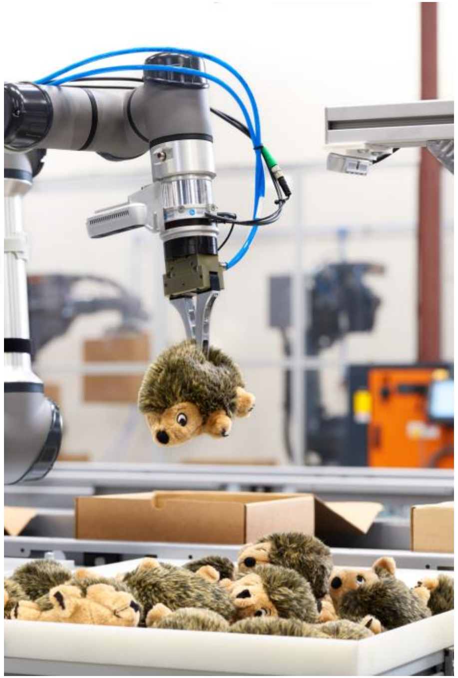

Uncertainty Estimation

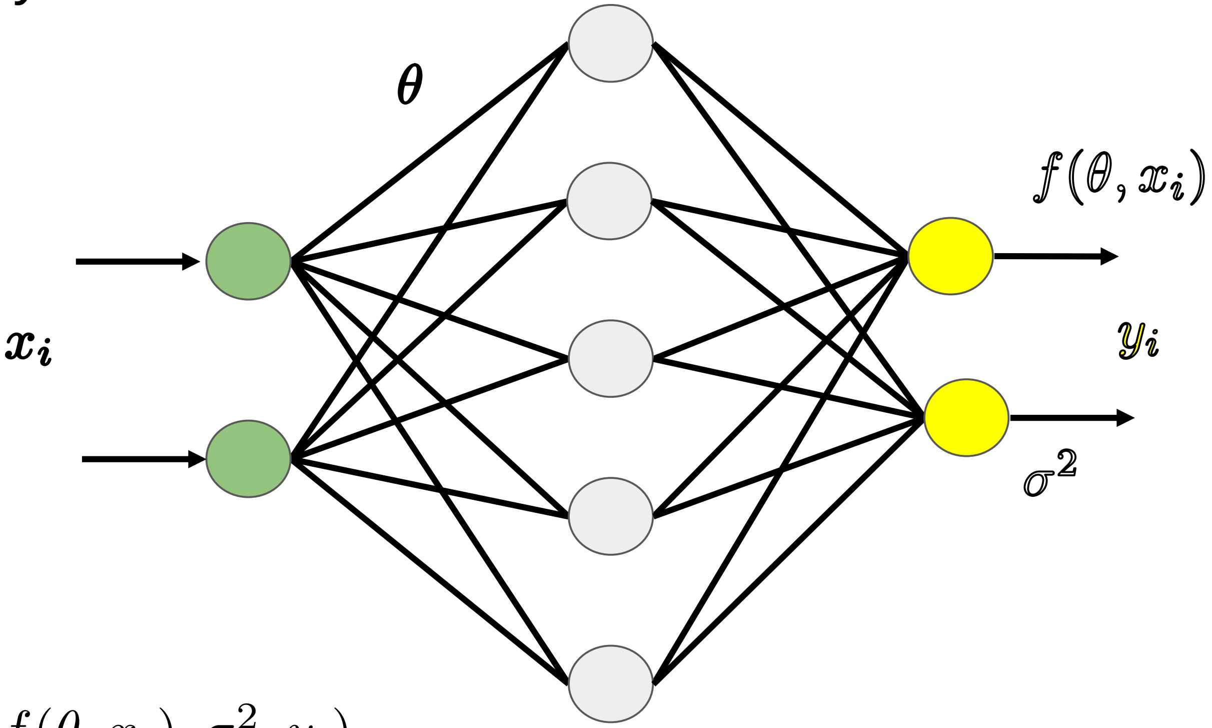

$$\underset{\theta}{\operatorname{argmin}} \quad \mathcal{L}(f(\theta, x_i), \sigma^2, y_i)$$

Robustness to Label Outliers

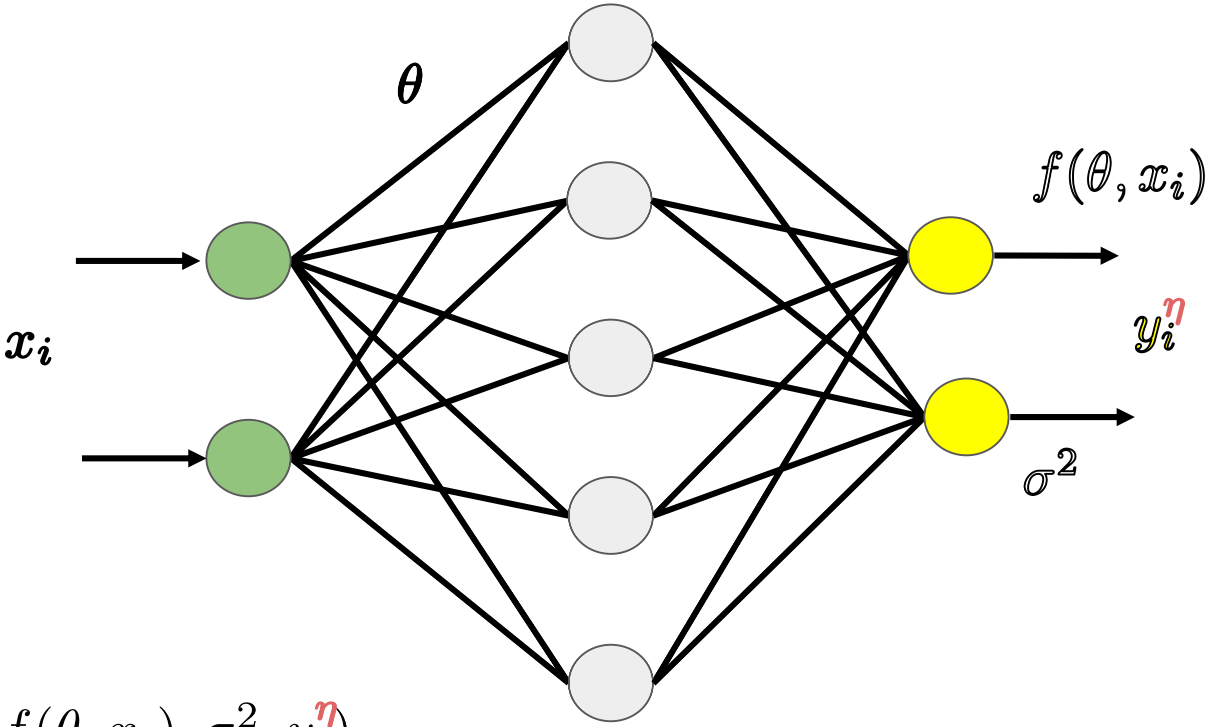

$$\underset{\theta}{\operatorname{argmin}} \quad \mathcal{L}(f(\theta, x_i), \sigma^2, y_i^\eta)$$

Robustness to Label Outliers

Robustness of uncertainty estimation in the presence of **label outliers**

Contributions:

1. Current uncertatinty estimation methods are **not robust** to label outliers

2. A robust loss function

$$\underset{\theta}{\operatorname{argmin}} \quad \mathcal{L}(f(\theta, x_i), \sigma^2, y_i^{\eta})$$

Motivation: Monocular Depth Estimation

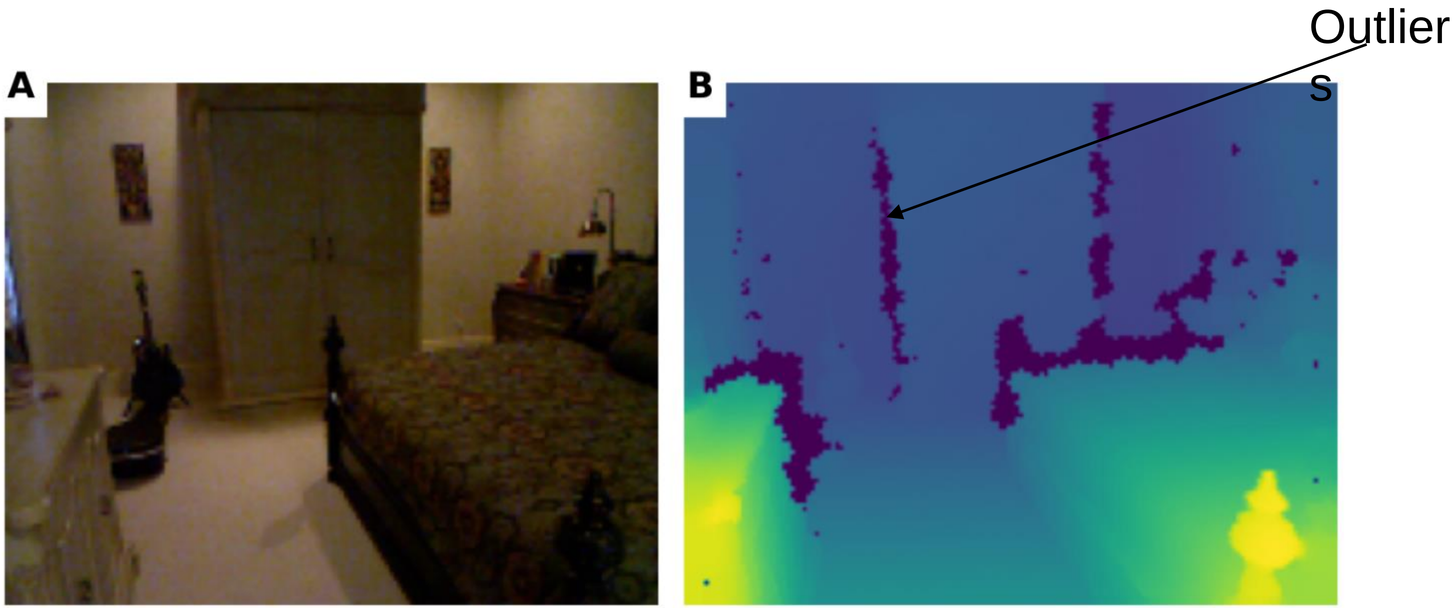

x_i y_i^{η}

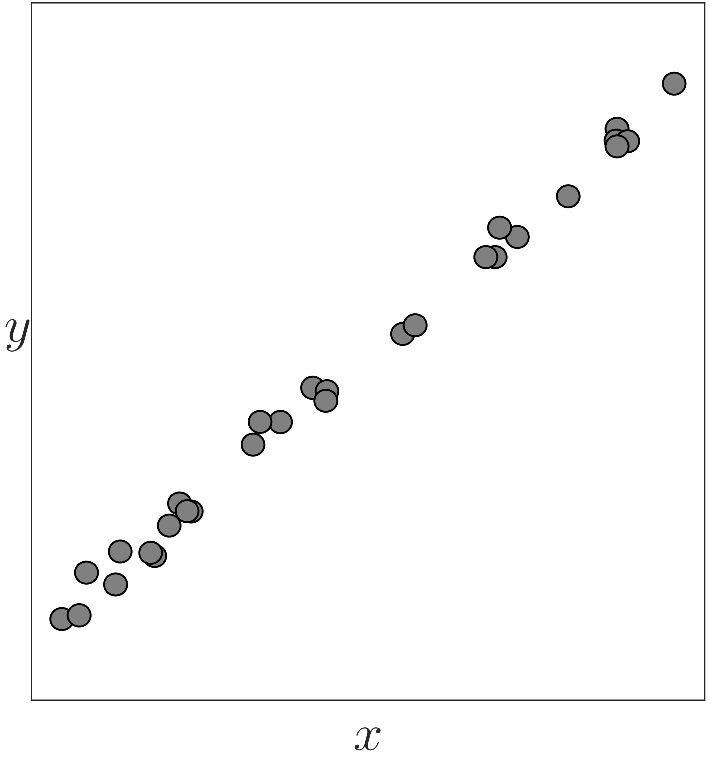

y
x

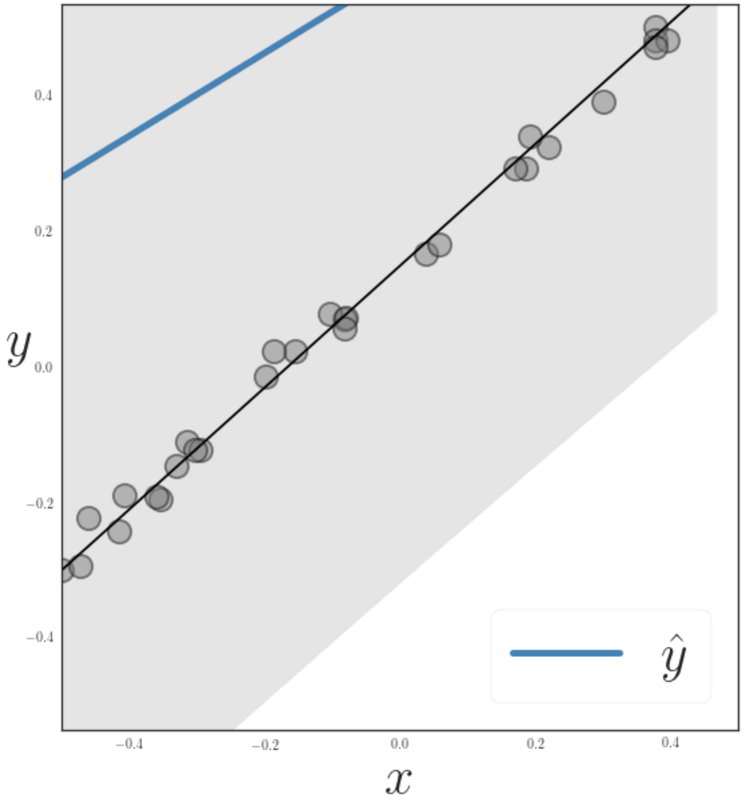

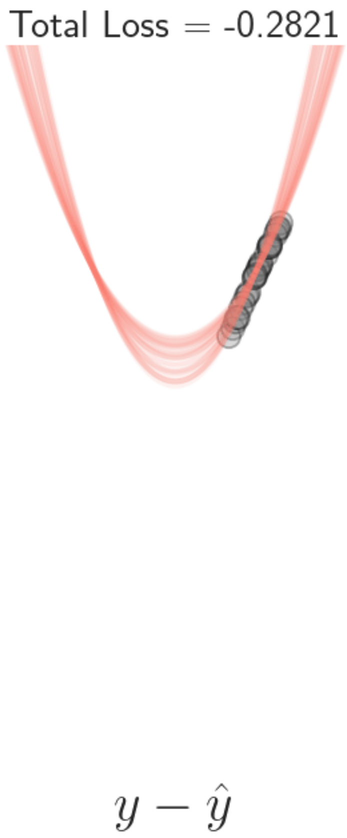

Nix et al, 1994

Uncertainty estimation : Gaussian NLL Loss

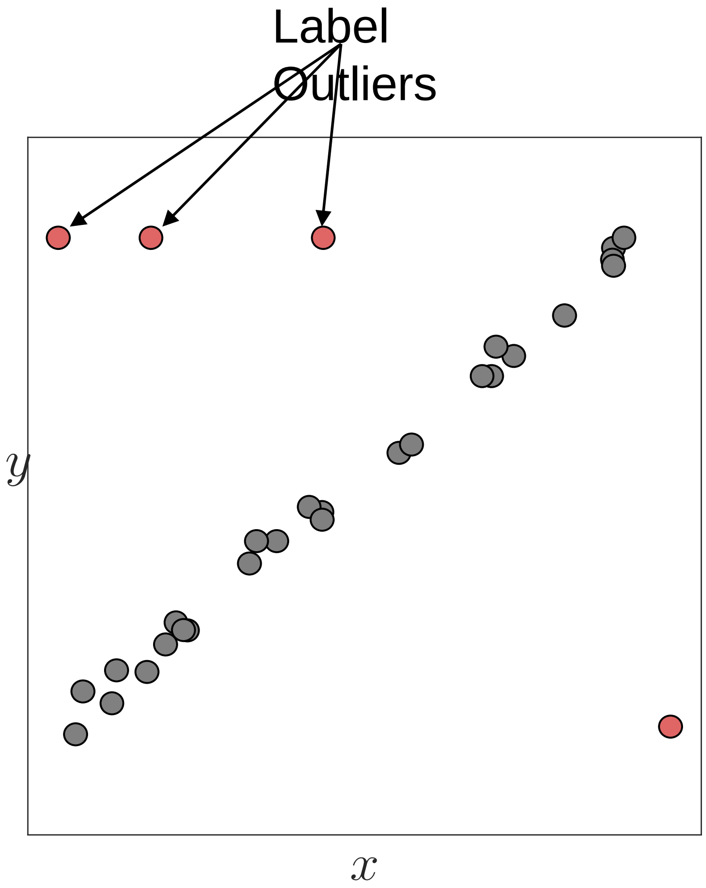
Label
Outliers
y
x

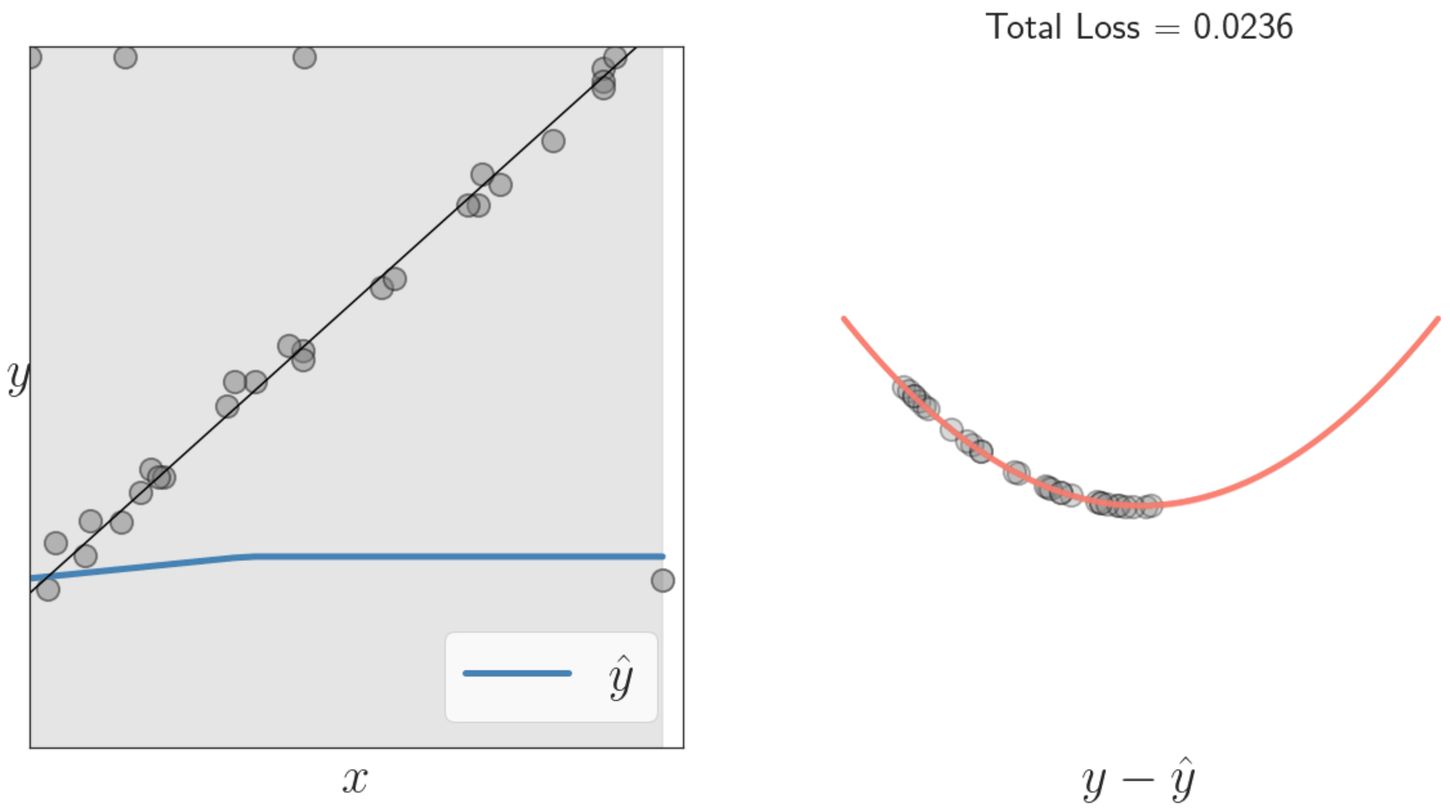

Nix et al, 1994

Uncertainty estimation : Gaussian NLL Loss

Gaussian

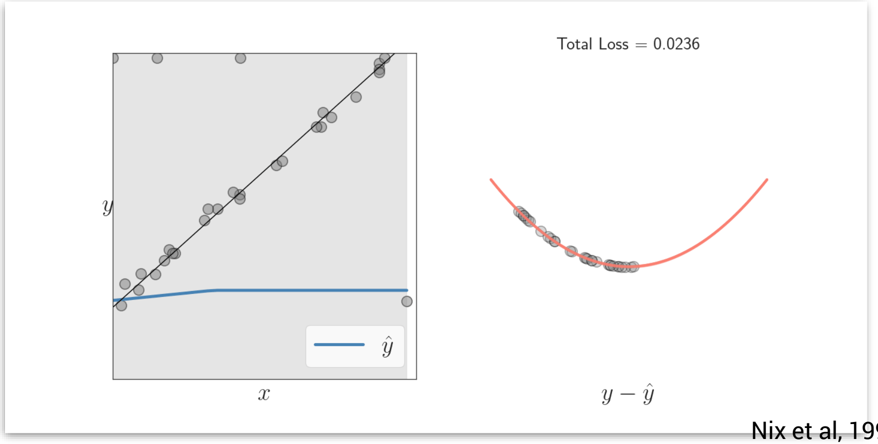

Nix et al, 1994

Dropout

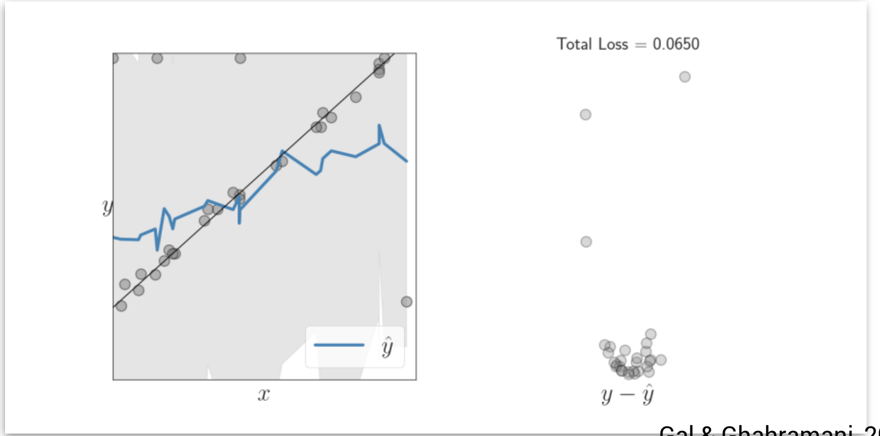

Gal & Ghahramani, 2016

Ensemble

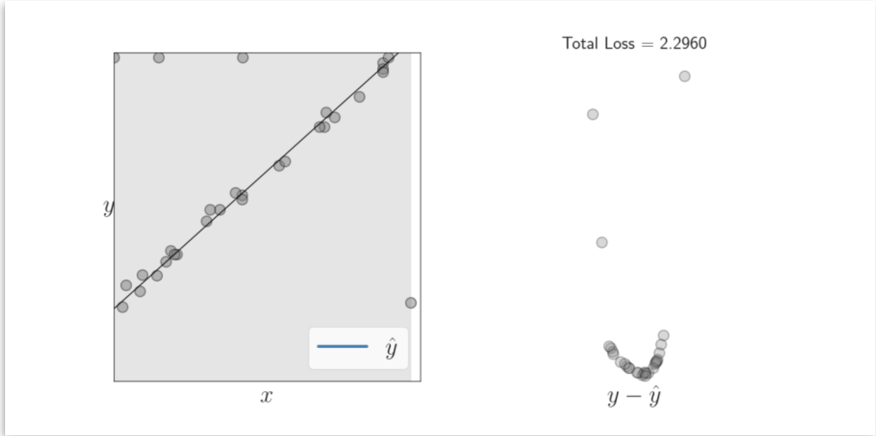

B Lakshminarayanan et al, 2017

Evidential

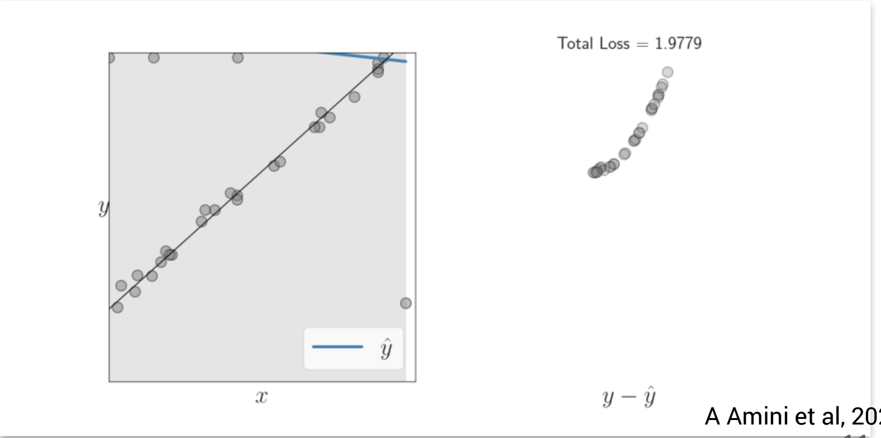

A Amini et al, 2020

Proposed Approach

Heavy Tailed Distribution

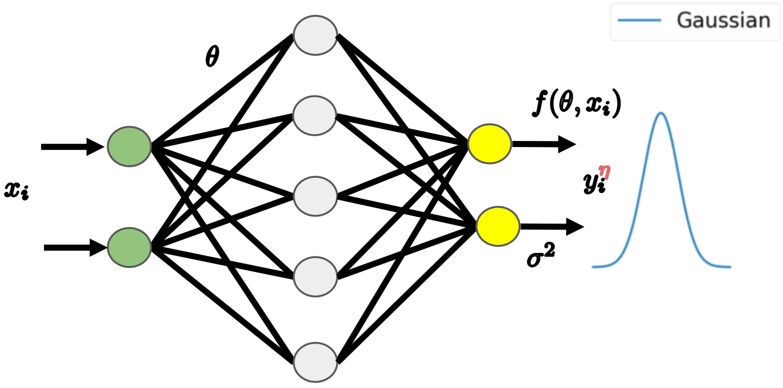

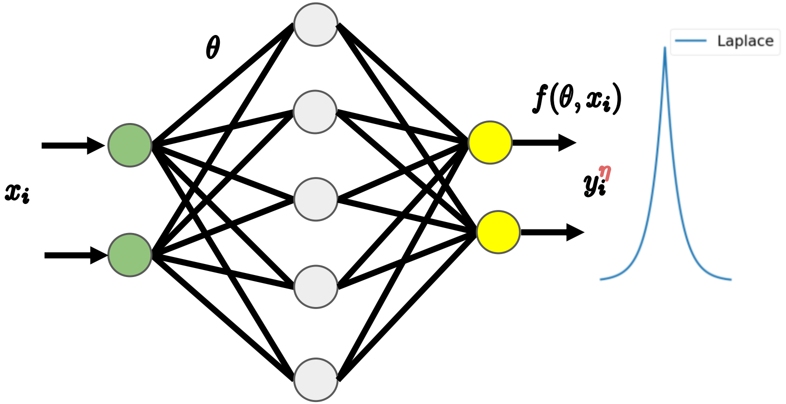

Gaussian NLL Loss

Laplace NLL Loss

Laplace Negative Log Likelihood

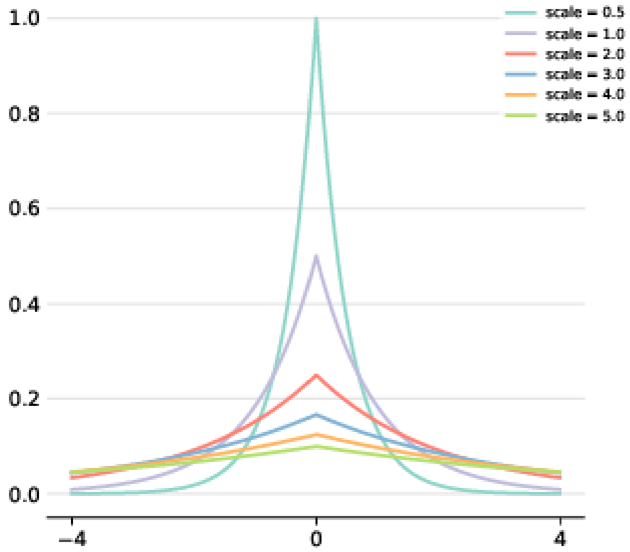

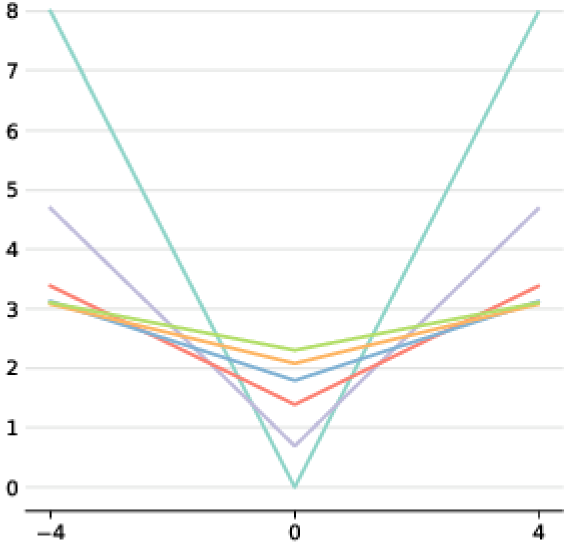

$$p(y|\mu, s) = \frac{1}{2s}\exp(-\frac{|y-\mu|}{s})$$

$$NLL(x, y, \theta) = \log(2f_\theta^s) + \frac{|y-f_\theta^\mu|}{f_\theta^s}$$

Laplace Negative Log Likelihood

Monotonic w.r.t inputs |x| and scale s
(useful for graduated non-convexity)

Smooth w.r.t. x and scale

Bounded first and second derivatives
(no exploding gradients)

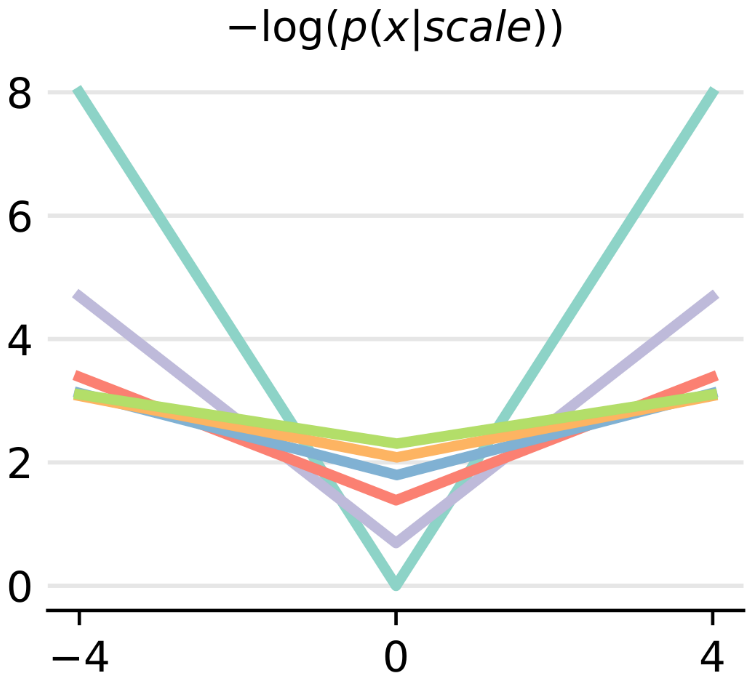

$$NLL(x, y, \theta) = \log(2 f^s_\theta) + \frac{|y - f^\mu_\theta|}{f^s_\theta}$$

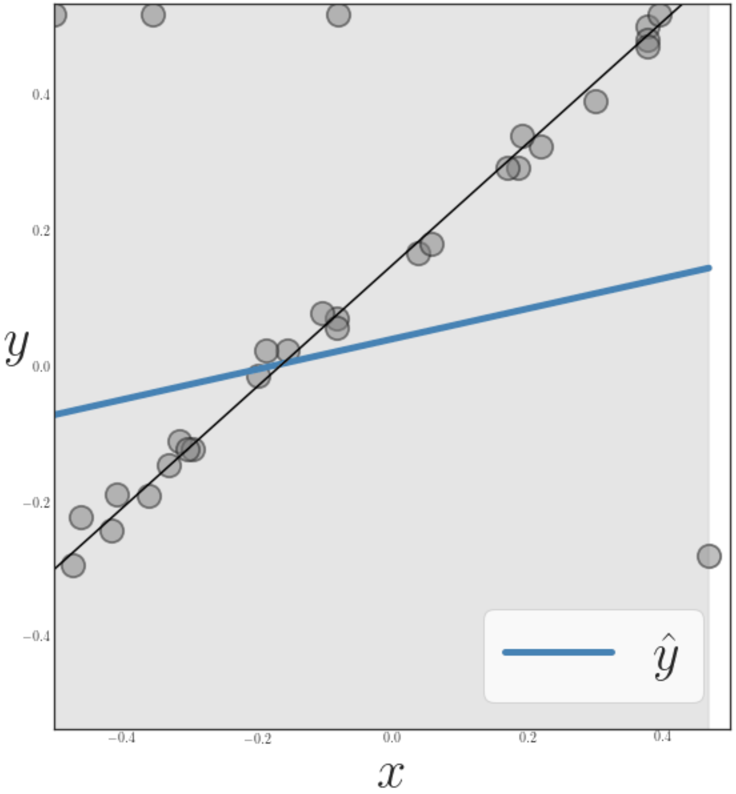

Total Loss = 0.3988

Results

Monocular Depth Estimation

- NYU-Depth-v2-dataset
  - with outliers
  - cleaned dataset
- Architecture:
  - U-Net (Ronneberger, Fischer, and Brox 2015) with spatial dropout.
- Uncertainty Metrics :
  - Interval Score
  - Entropy


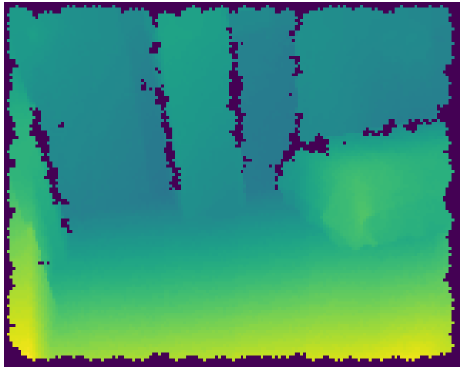
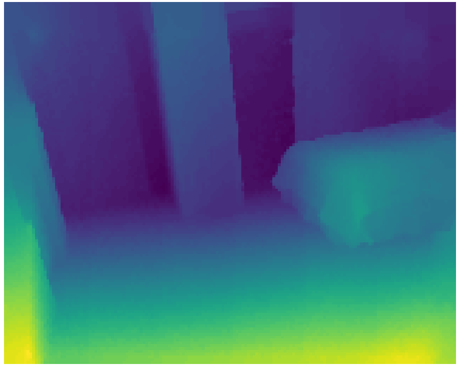

Clean vs Outlier Labels

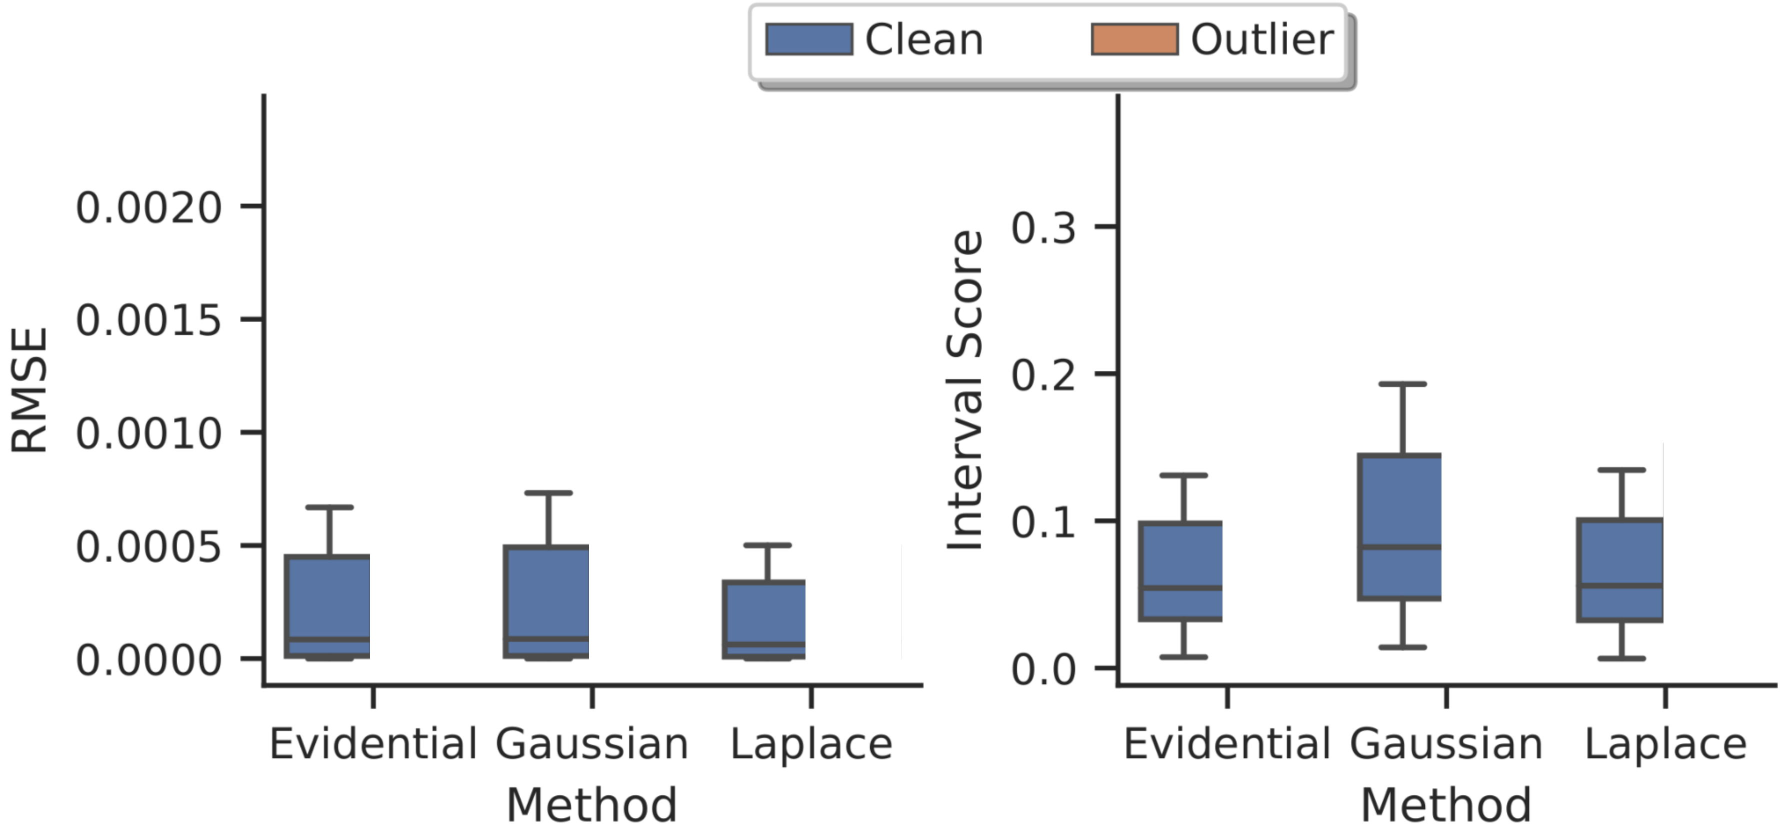

Clean vs Outlier Labels

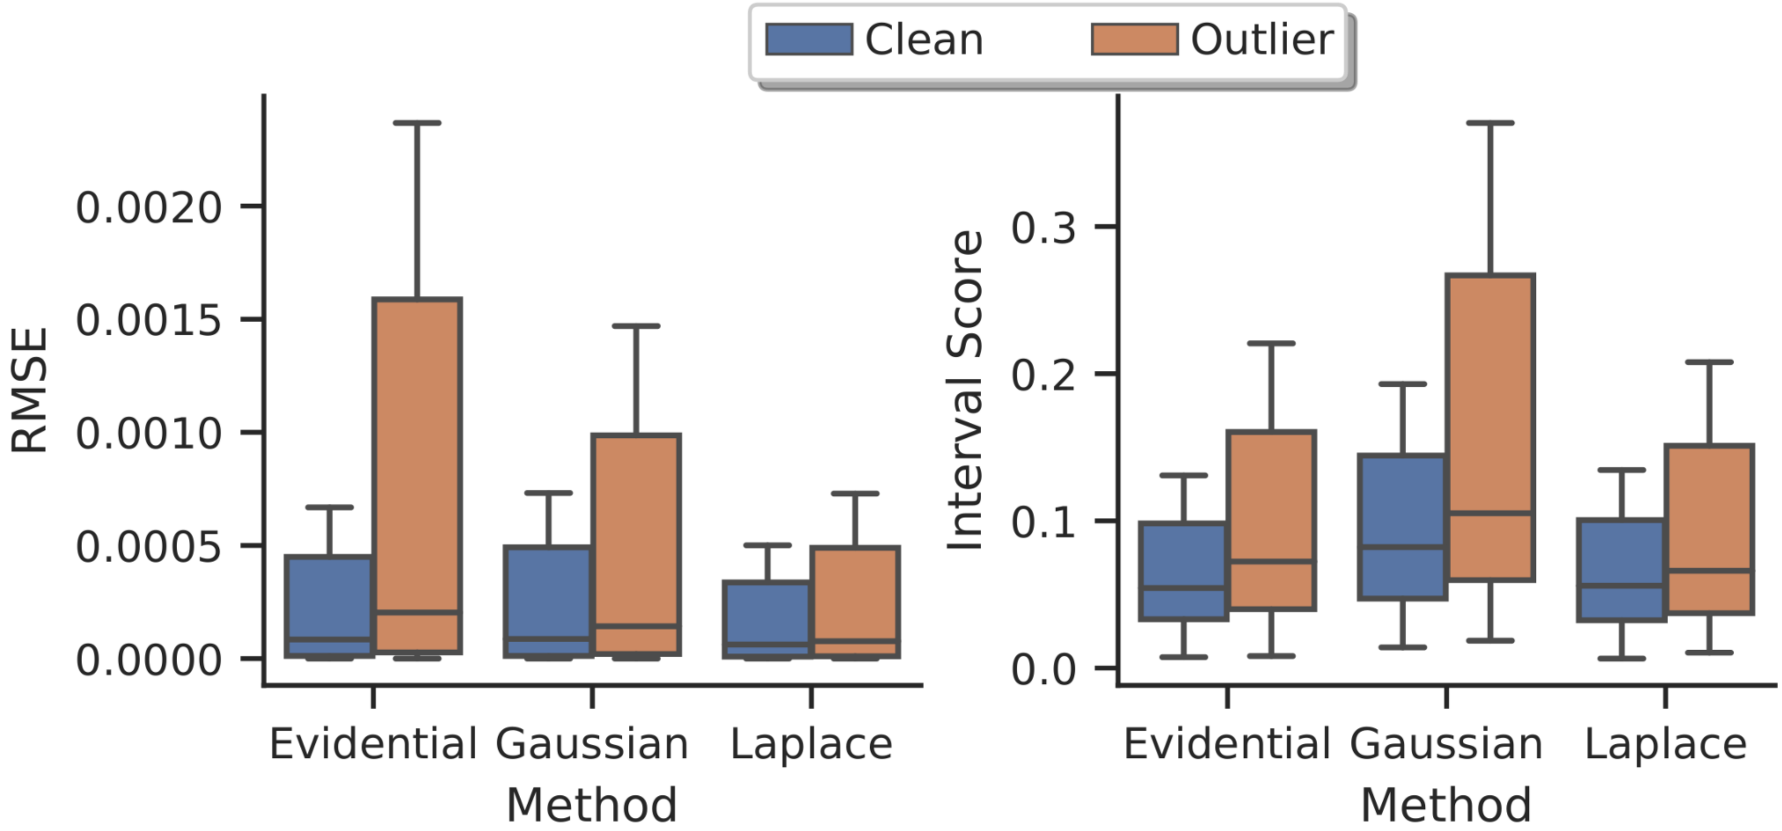

Out-of-Distribution

NYUv2 Dataset
In-Distribution

Appollo Dataset
Out-of-
Distribution

Out-of-Distribution

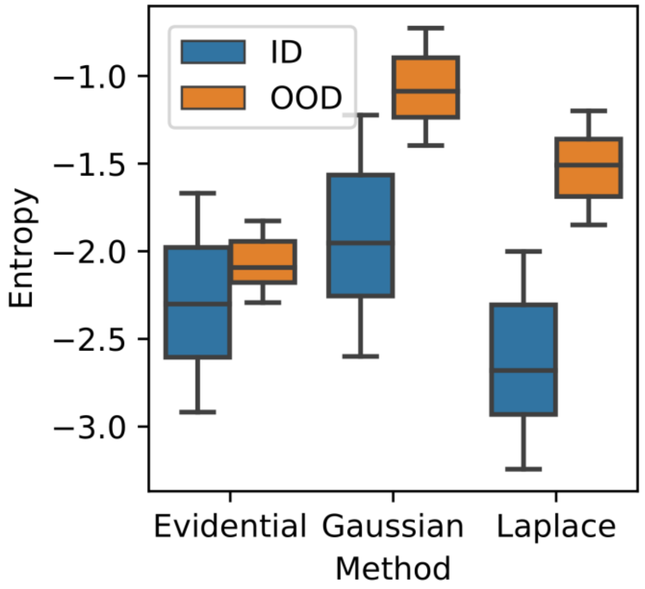

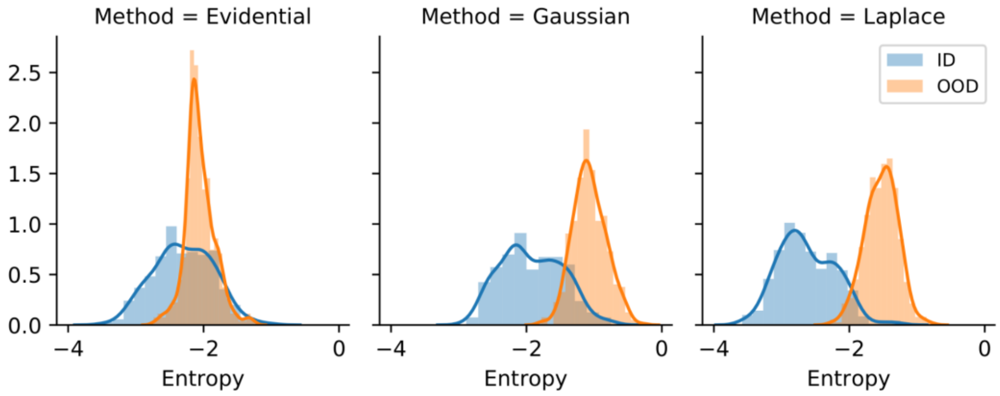

Conclusion

- We improve robustness of uncertainty estimation using a **heavy-tailed distribution based loss function**.

- When applied to high dimensional datasets containing outliers, such as depth estimation datasets, the Laplace loss function is able to **better estimate the uncertainties.**

- The proposed robust loss function could benefit in building software which uses uncertainty from the neural network for safe deployment of deep neural networks in autonomous systems.

Thank you

Real World Datasets

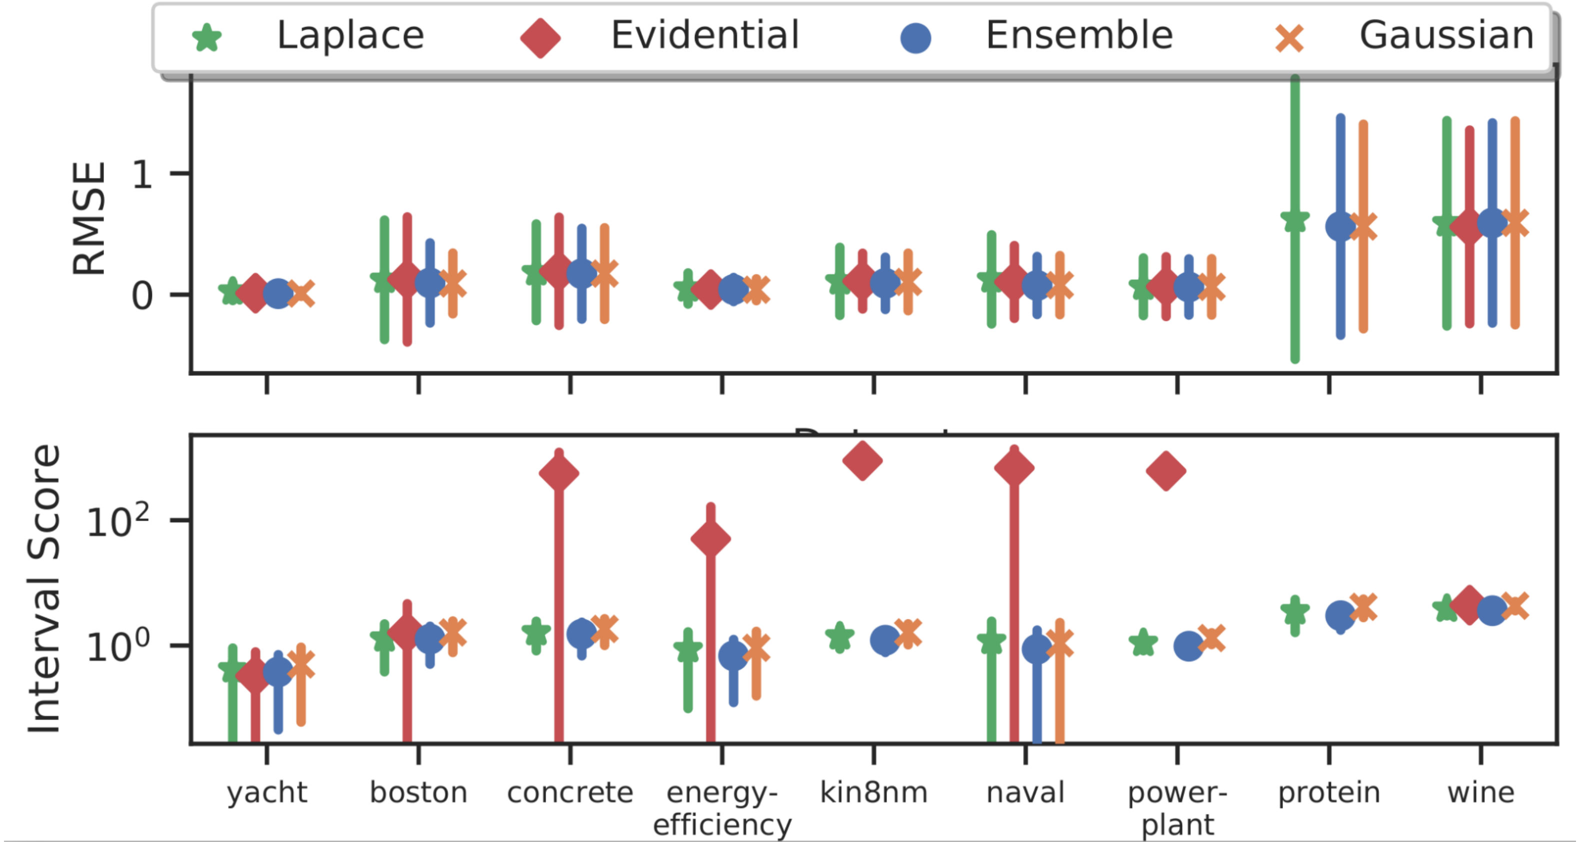

Breakaway point Benchmark: Toy dataset

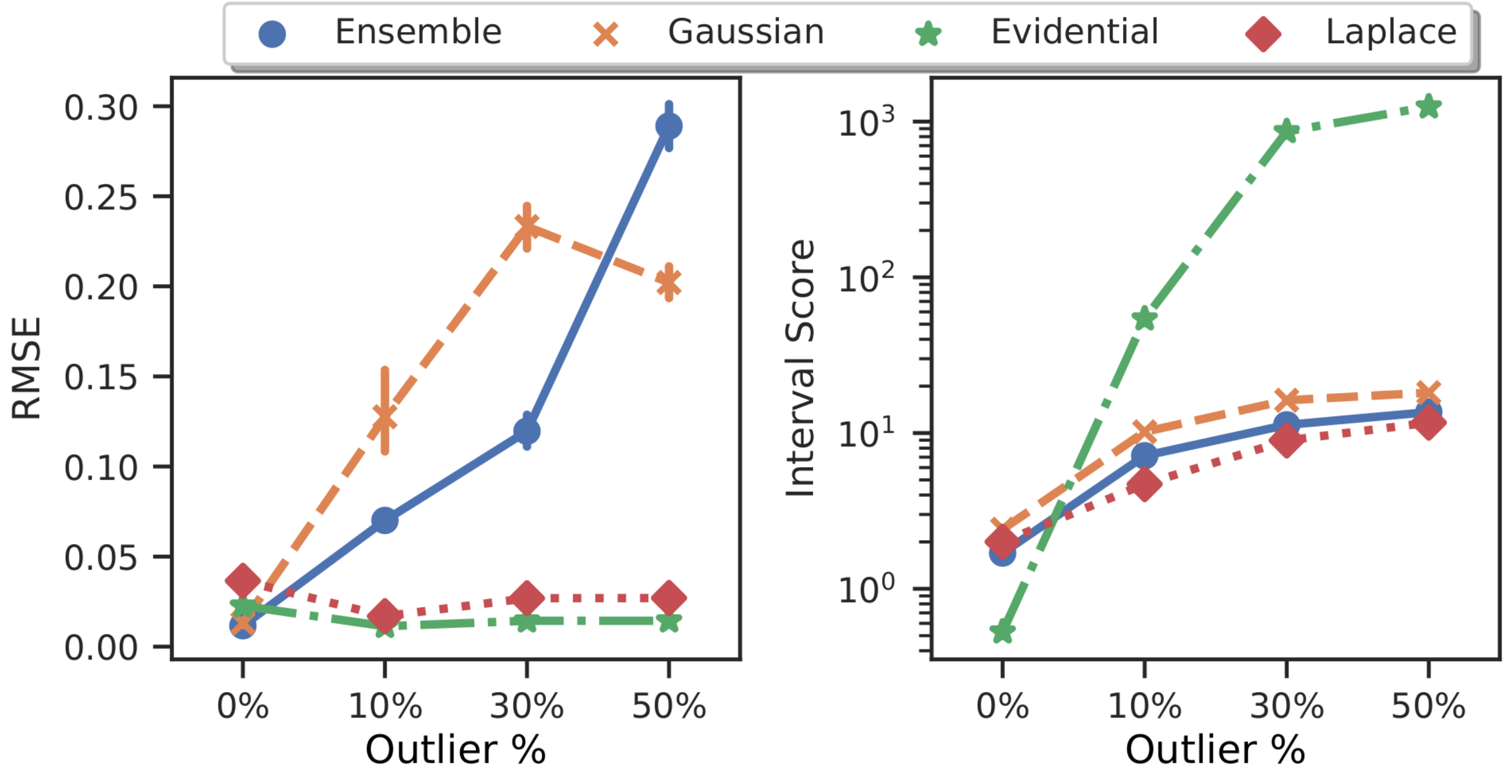

Adversarial Attack

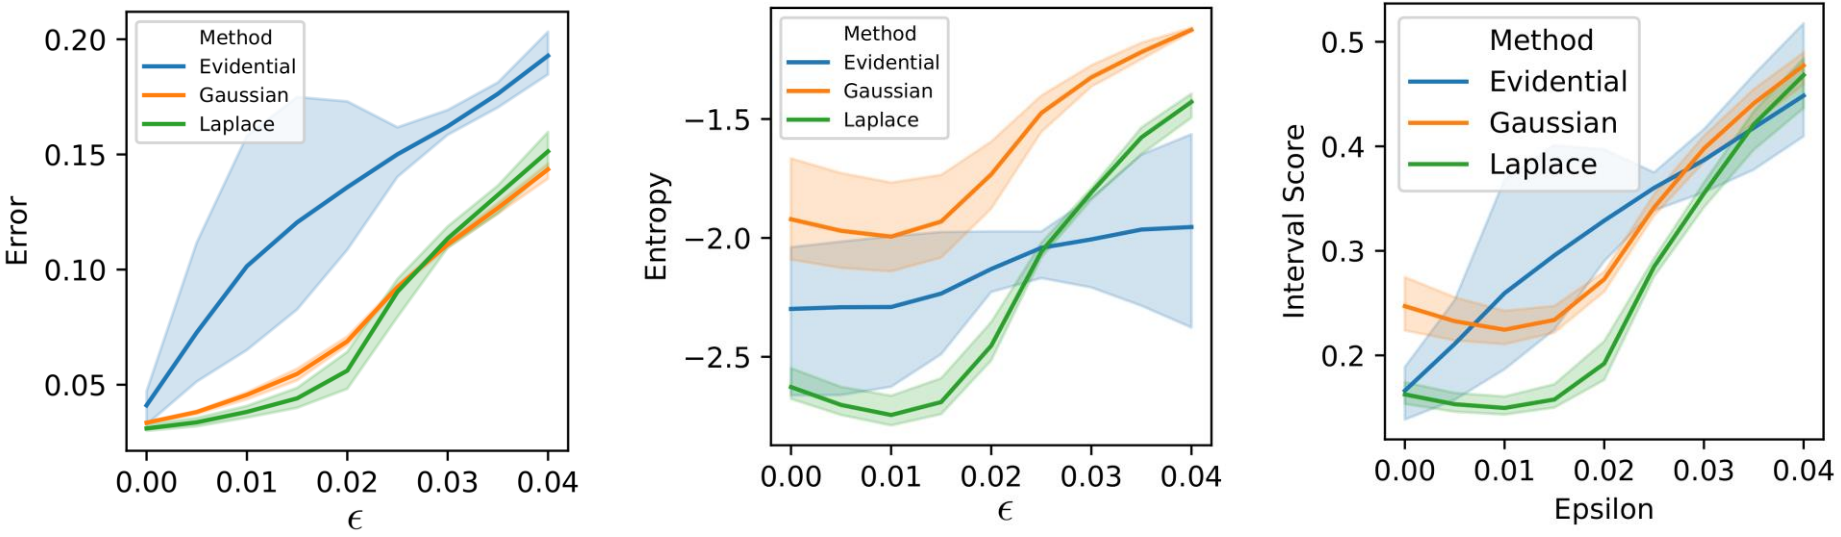

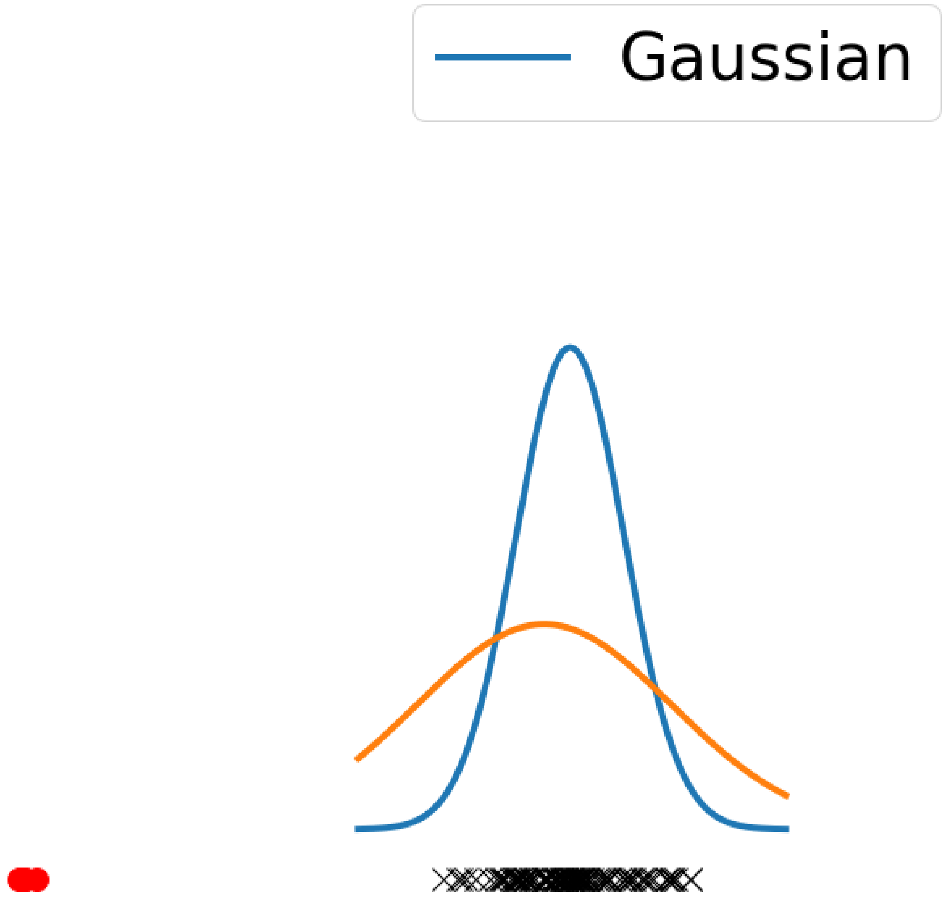
Gaussian

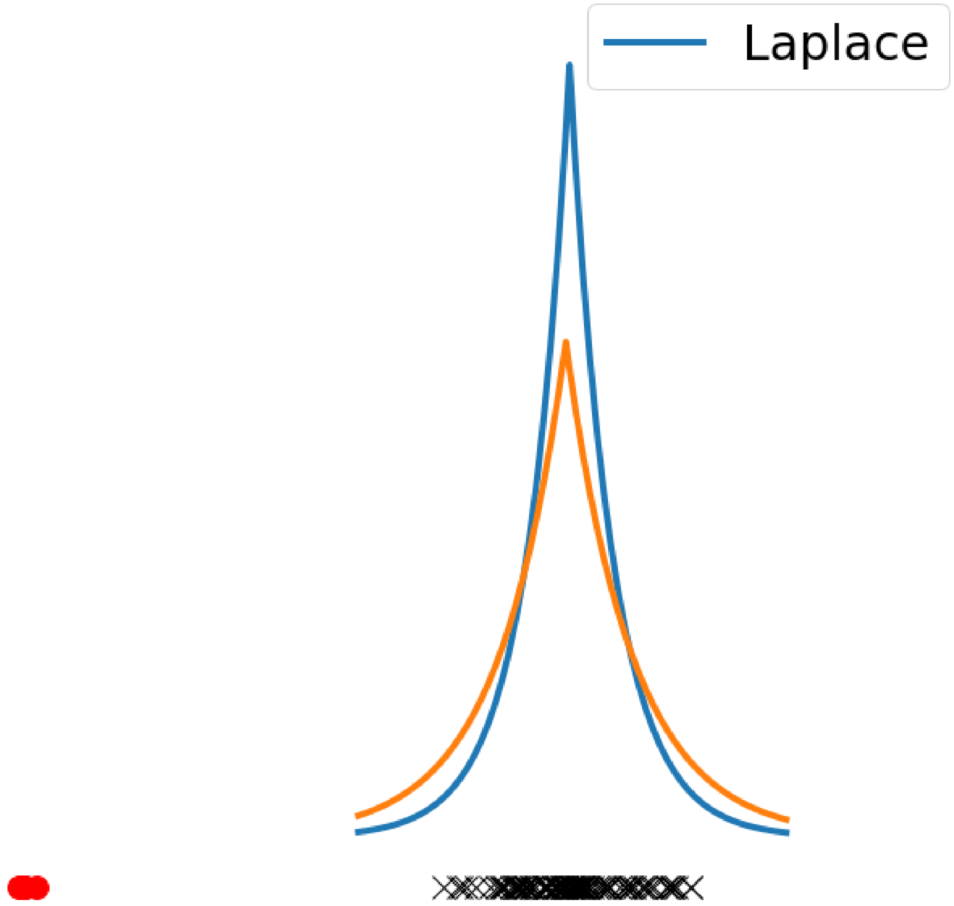
Laplace

Confidence Calibration

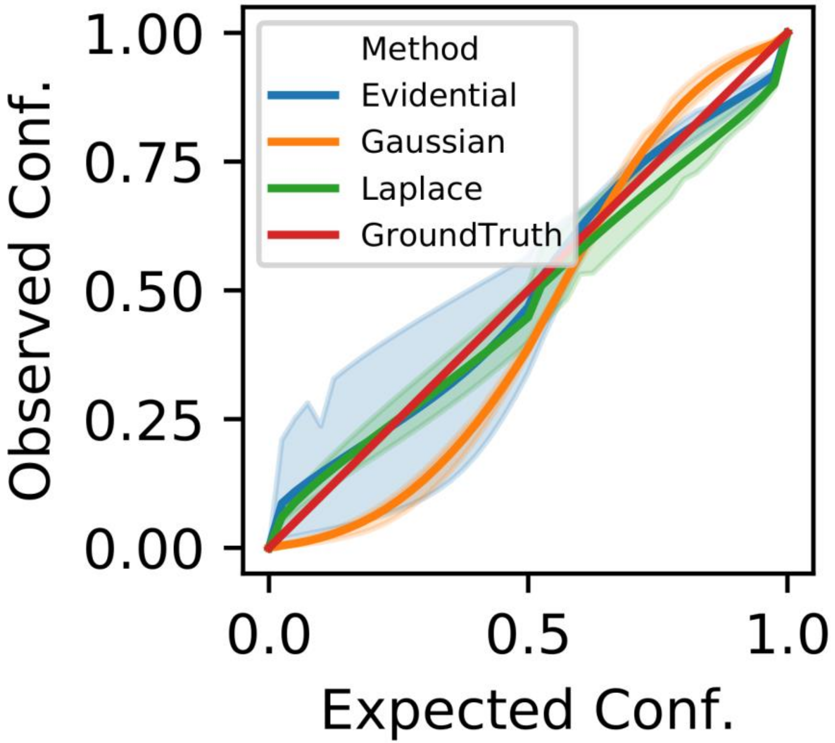

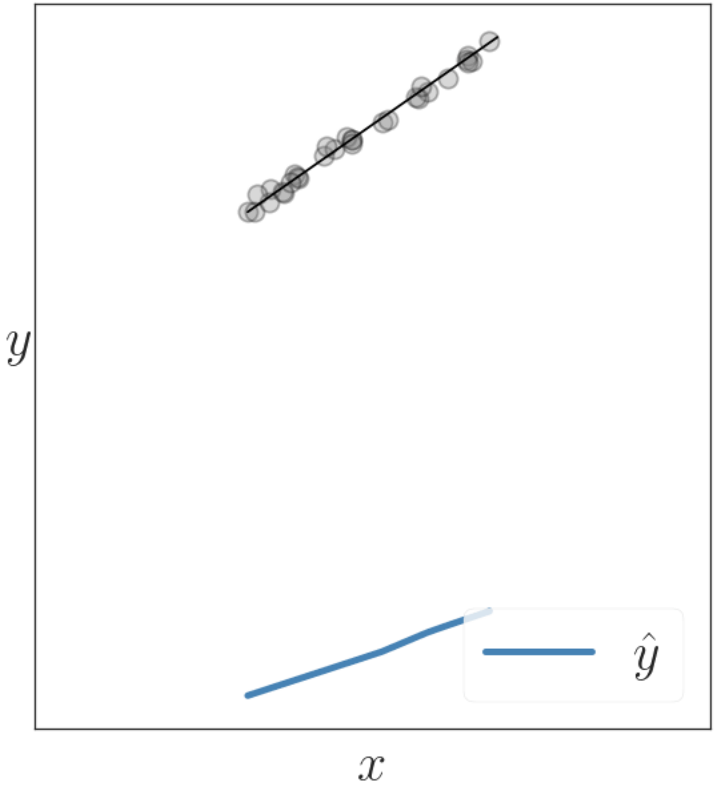

L2 loss

Monocular Depth Estimation

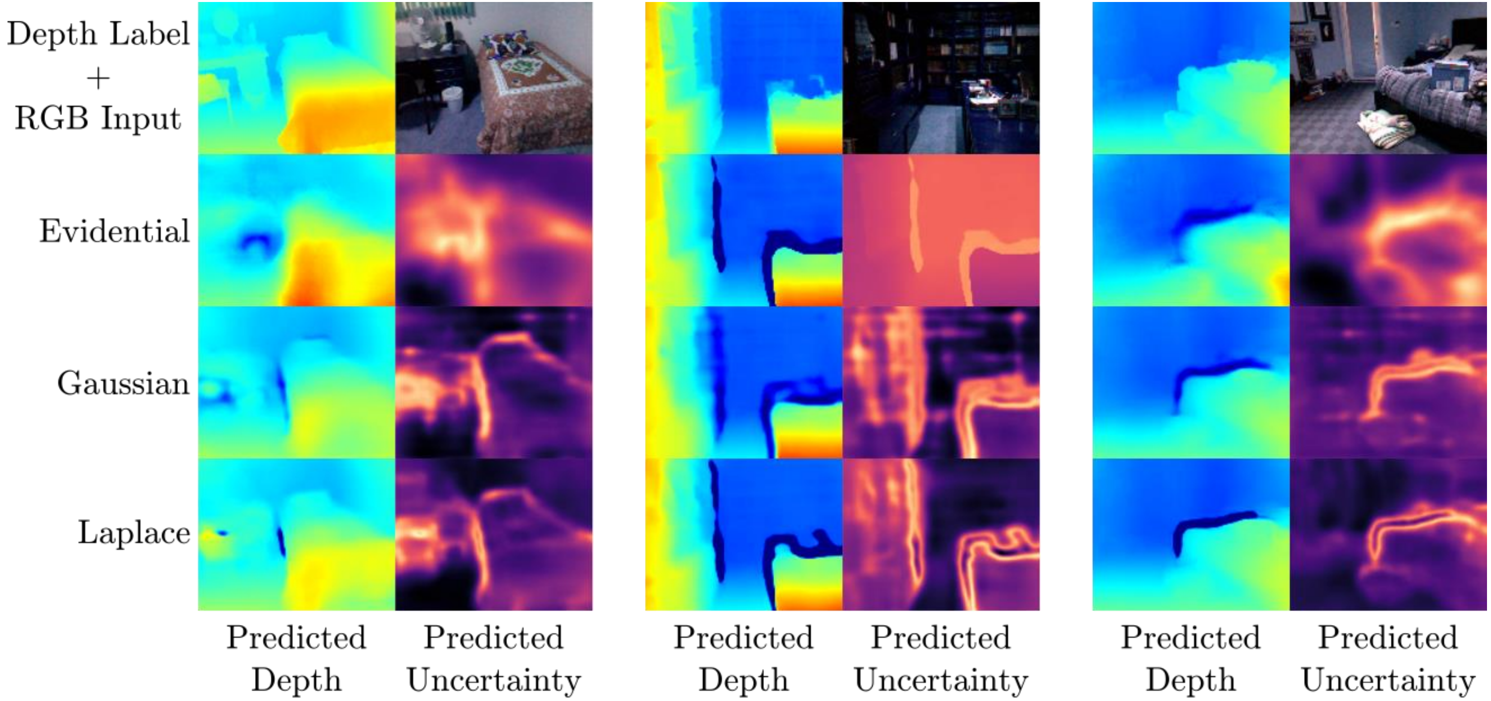

Regression

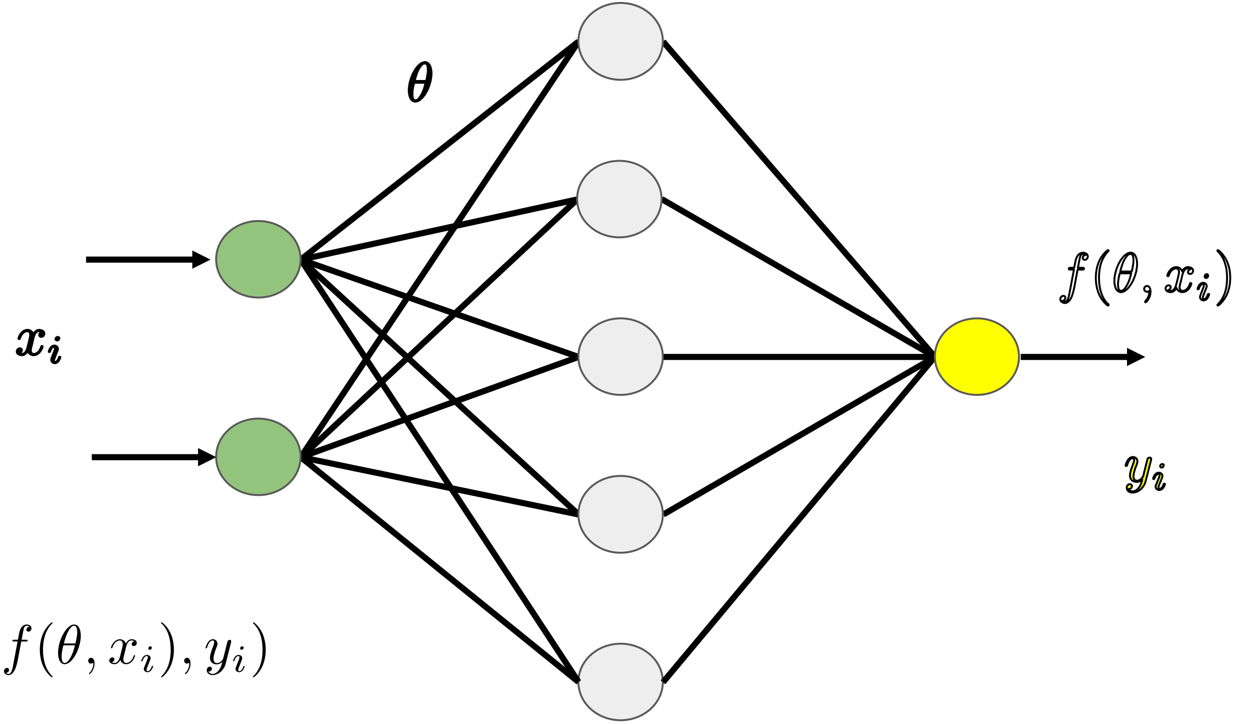

$$\underset{\theta}{\operatorname{argmin}} \quad \mathcal{L}(f(\theta, x_i), y_i)$$