NOISE-AWARE DIFFERENTIALLY PRIVATE REGRESSION VIA META-LEARNING

Ossi Räisä*, Stratis Markou*, Matthew Ashman, Wessel P. Bruinsma,
Marlon Tobaben, Antti Honkela, Richard E. Turner
University of Helsinki, University of Cambridge, Microsoft Research

*Equal Contribution

INTRODUCTION

- Goal: probabilistic regression, output predicted Gaussian mean and variance that
  - adapts to a new dataset in one forward pass, and
  - is differentially private with regards to that dataset
- We meta-train a convolutional neural process with simulated data
- We add noise with an improved functional mechanism

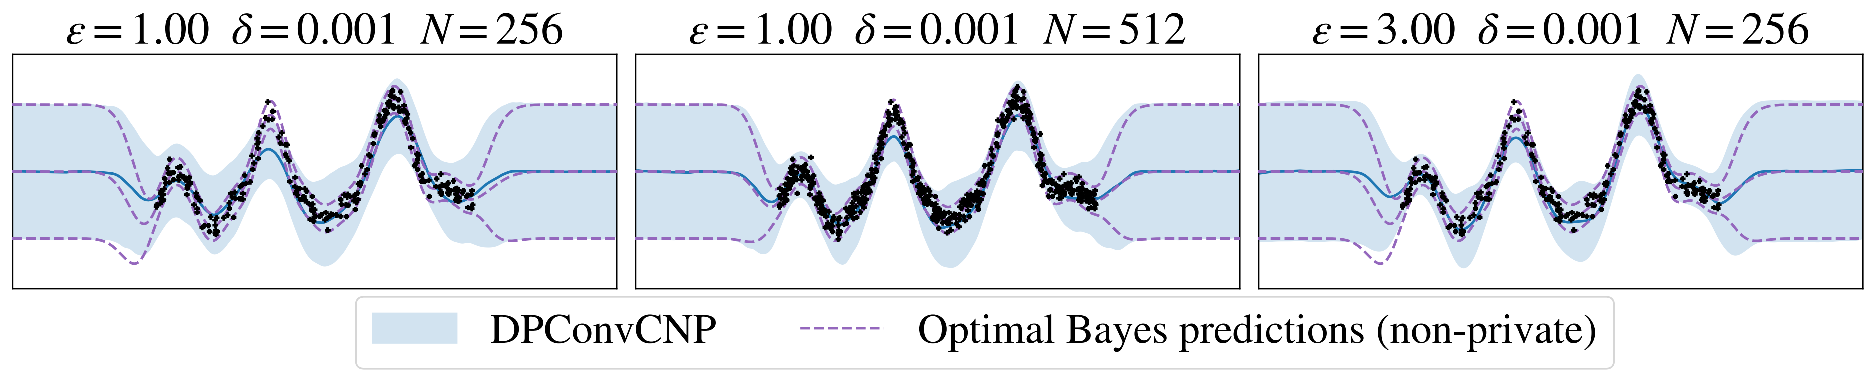

META-LEARNING

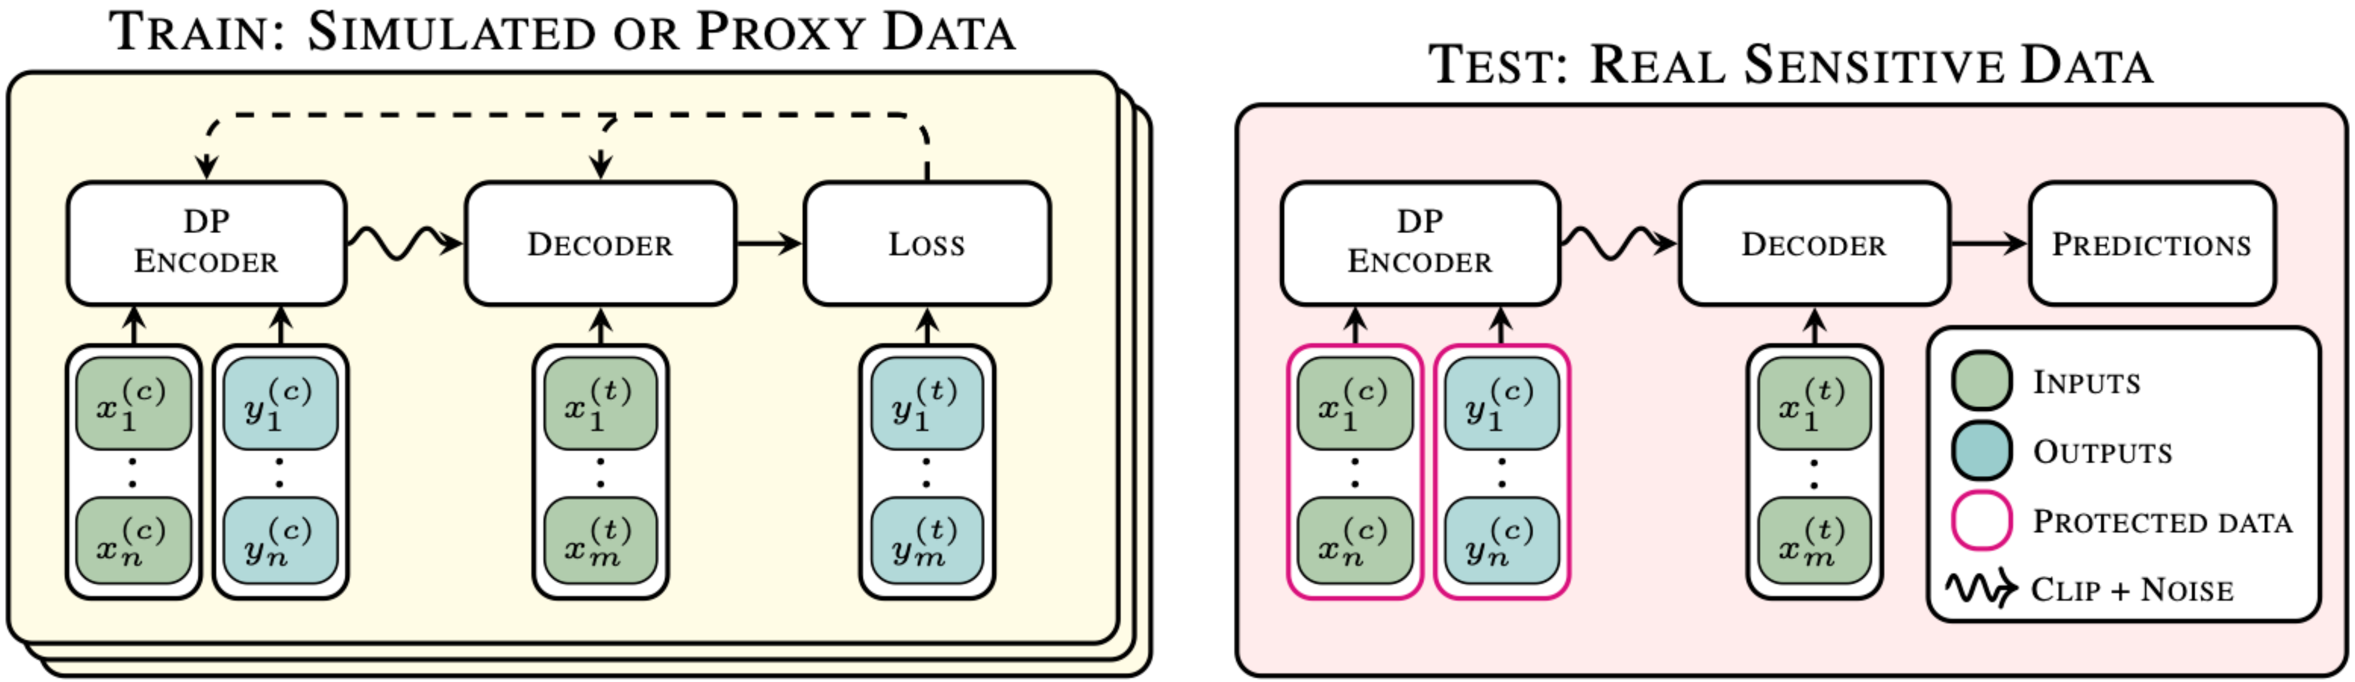

CONVOLUTIONAL CONDITIONAL NEURAL PROCESS

- Our base model: ConvCNP
- Encode context set to two channels:
  - Density $r^{(d)}$, encodes location
  - Signal $r^{(s)}$, encodes location · value
- Learns decoder CNN

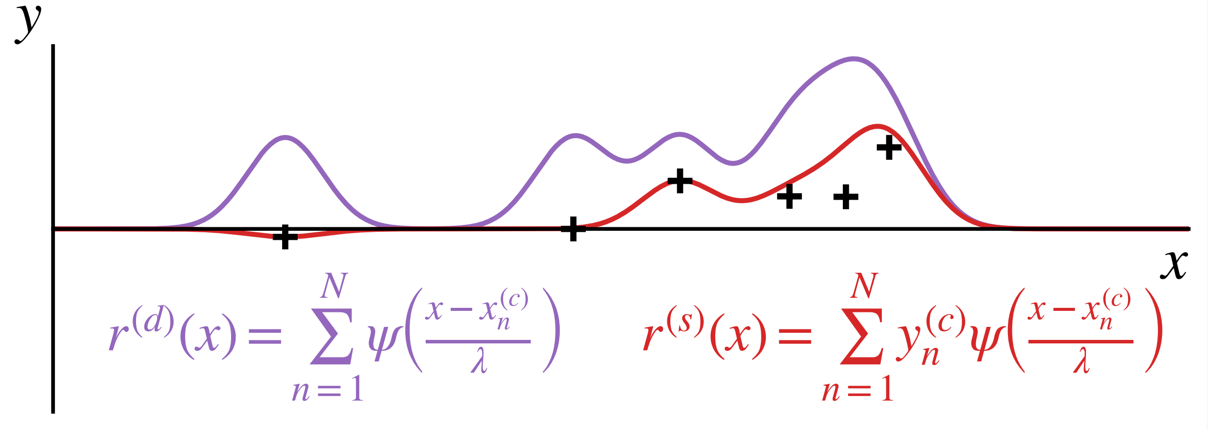

$$r^{(d)}(x) = \sum_{n=1}^{N} \psi\left(\frac{x - x_n^{(c)}}{\lambda}\right) \qquad r^{(s)}(x) = \sum_{n=1}^{N} y_n^{(c)} \psi\left(\frac{x - x_n^{(c)}}{\lambda}\right)$$

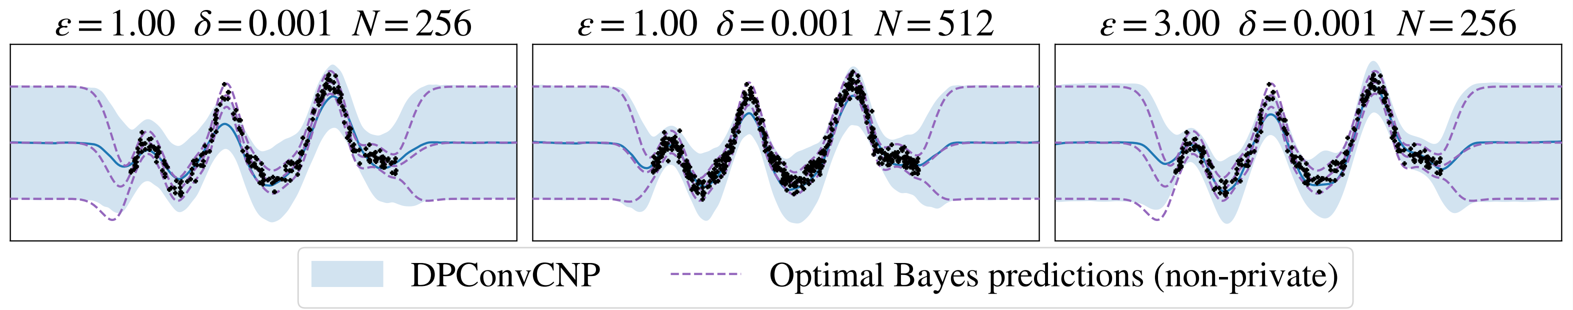

IMPROVED FUNCTIONAL MECHANISM

- We prove a Gaussian DP bound for the functional mechanism:
  - $c = {}^{\Delta}/_{\mu}$
- Convert μ to (ϵ, δ) → lower noise variance for same privacy bound

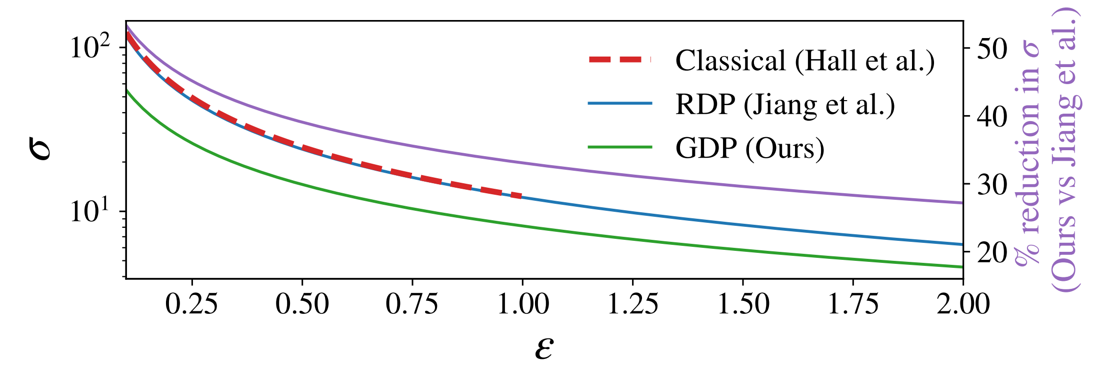

$\Delta^2 = 10, \delta = 0.001$

Hall et al. “Differential privacy for functions and functional data” JMLR 2013
Jiang et al. “Functional Rényi Differential Privacy for Generative Modeling” NeurIPS 2023

DP-CONVCNP

- Recall ConvCNP encoding:
  - Density $r^{(d)}$, encodes location
  - Signal $r^{(s)}$, encodes location · value
- We use the functional mechanism to privatise them
- Density channel has finite sensitivity
- Clipping y is needed for signal channel

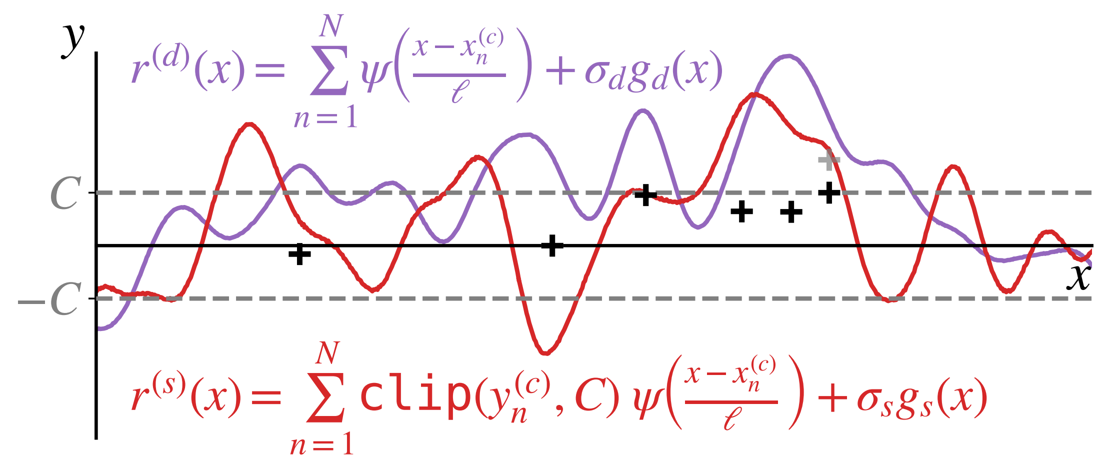

$$r^{(d)}(x) = \sum_{n=1}^{N} \psi\left(\frac{x - x_n^{(c)}}{\ell}\right) + \sigma_d g_d(x)$$

$$r^{(s)}(x) = \sum_{n=1}^{N} \mathtt{clip}(y_n^{(c)}, C)\, \psi\left(\frac{x - x_n^{(c)}}{\ell}\right) + \sigma_s g_s(x)$$

RESULTS: DOBE !KUNG DATASET

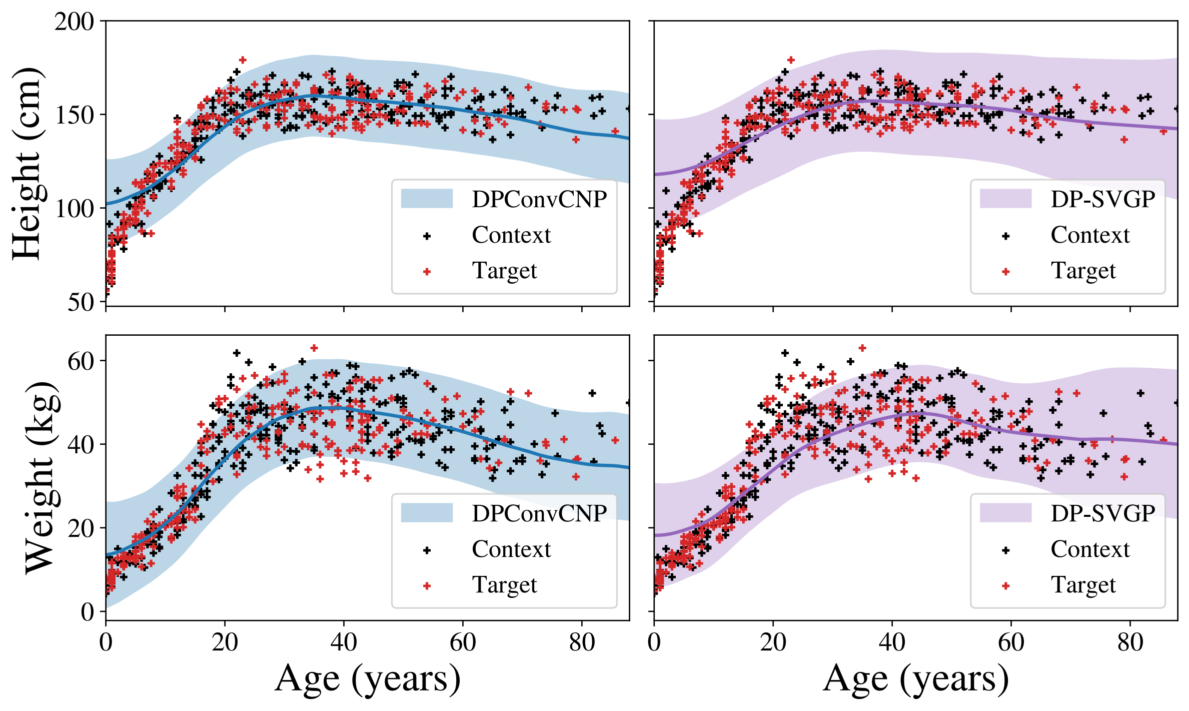

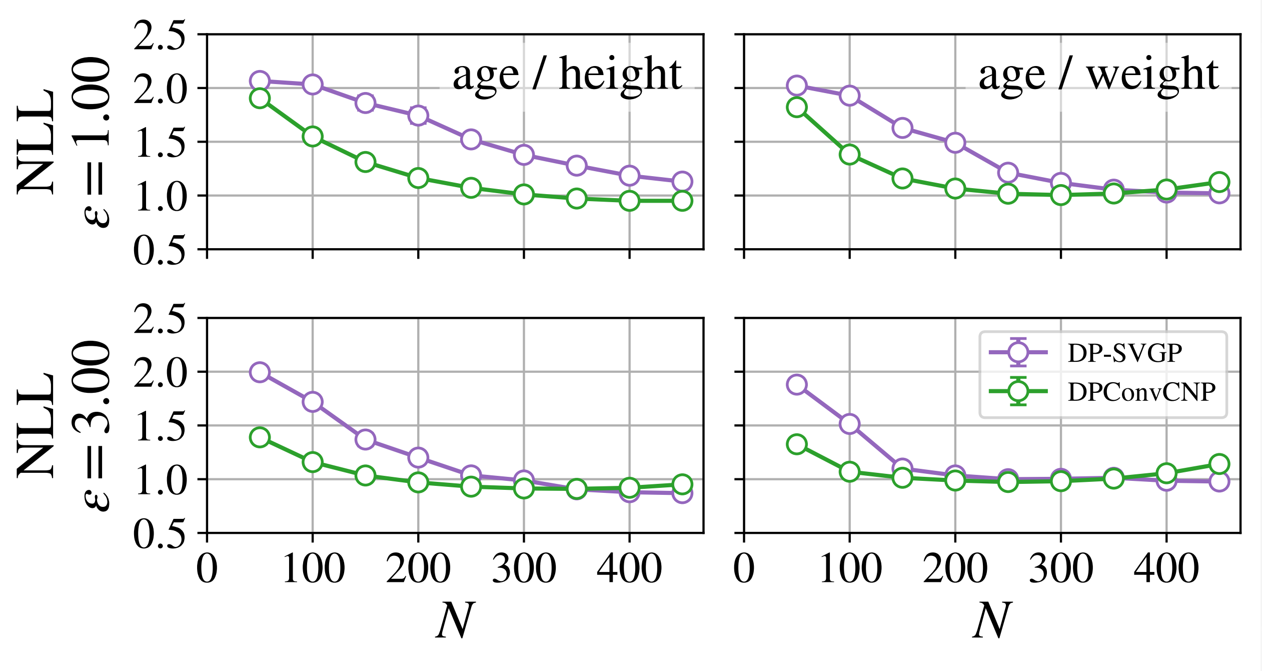

CONCLUSION

- We prove an improved privacy bound for the functional mechanism
- We develop DPConvCNP using the improved functional mechanism
- DPConvCNP provides:
  - Noise-aware predictions
  - Privacy for the context set
  - Meta-training with simulated data, no need for public data
  - Improved accuracy over DP Gaussian process baseline

THANK YOU

Visit our poster for more information