Conditional Neural Processes

Weisz S., Buonomo, A, Yuang, L, Fathullah, Y.

UNIVERSITY OF CAMBRIDGE

Introduction

Conditional Neural Processes (CNPs) *meta-learn* a mapping from context sets D_c to *predictive distributions* at target locations x_T, $p_\theta(y_T|x_T; D_C)$, using *neural networks.*

Desirable Properties

1. Data-efficient (using meta-learning)
2. Fast predictions at test time: $\mathcal{O}(n+m)$ for predicting at m target locations with n context observations
3. Good uncertainty representation (by modelling stochastic processes)
4. Data-driven expressivity (using deep learning)

Model Architecture

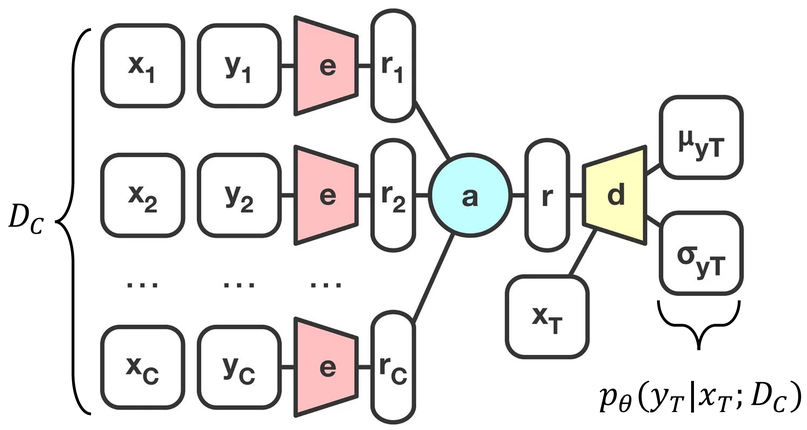

Maximum Likelihood Training

Minimize using gradient descent:

$$\mathcal{L}(\theta) = -\mathbb{E}_{D_C, D_T \sim P}\left[\sum_{(x_t, y_t) \in D_T} \log p_\theta(y_t|x_t; D_C)\right]$$

1D Function Regression

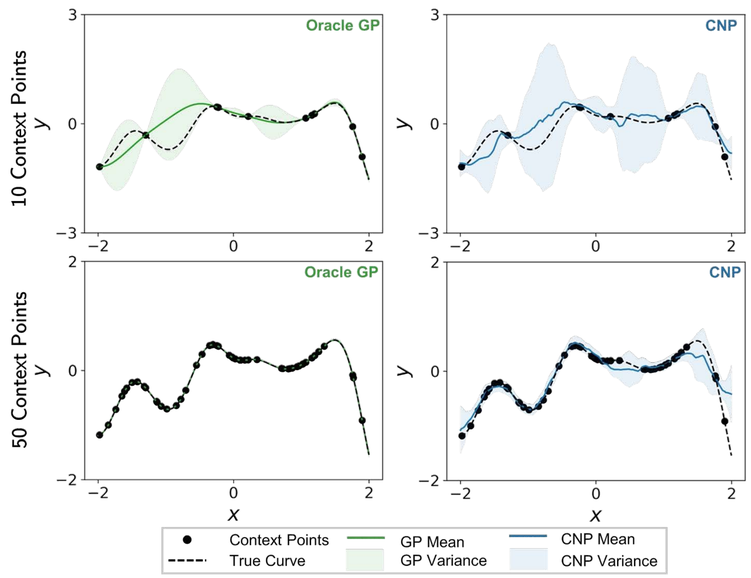

MNIST Image Completion

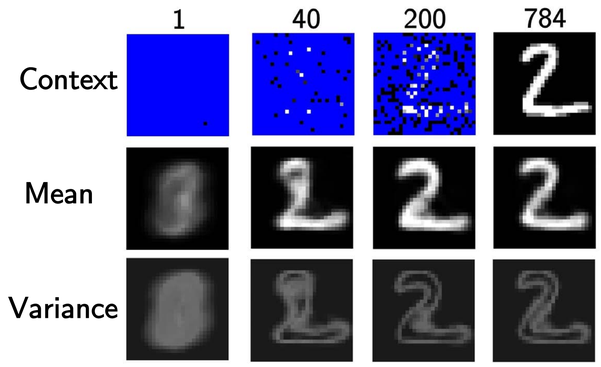

CelebA Image Completion

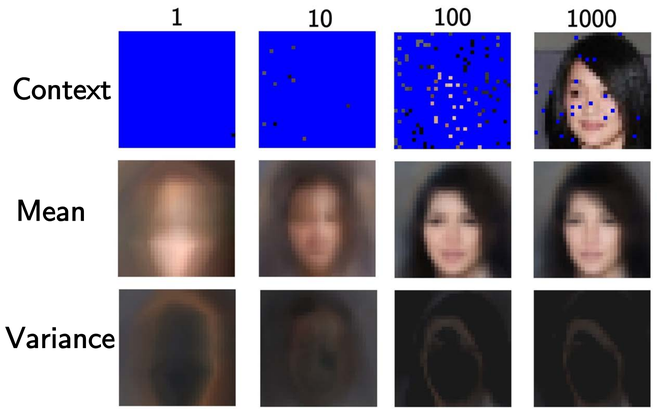

Image Completion Flexibility

Handle arbitrary context set patterns

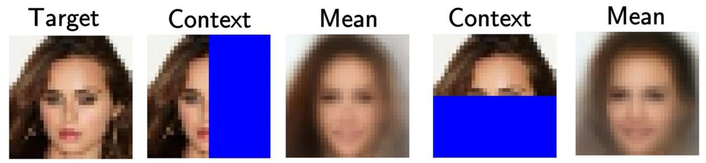

Increase image resolution

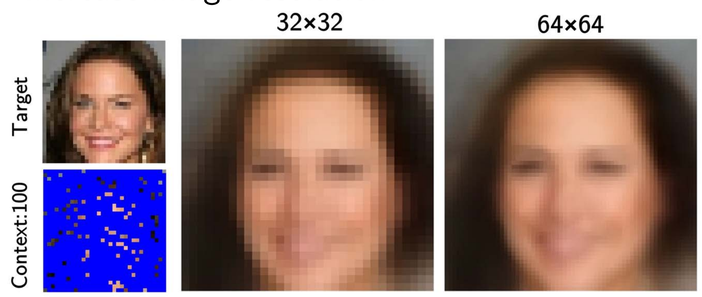

Extension: ConvCNPs

Add inductive bias: Translation Equivariance

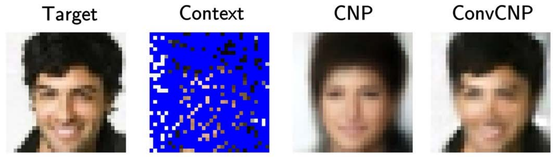

Avoid underfitting

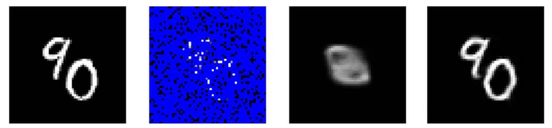

Generalise & extrapolate

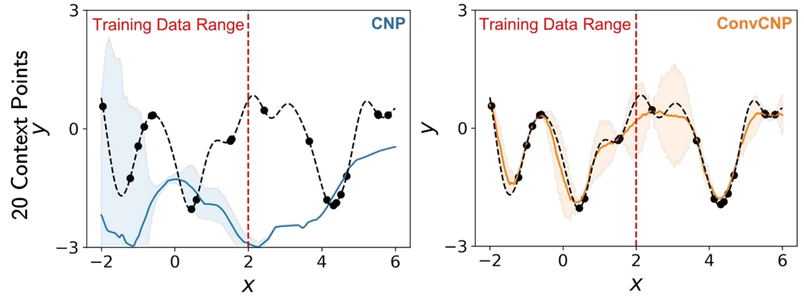

References

[1] Garnelo, Marta, et al. "Conditional neural processes." *International Conference on Machine Learning*. PMLR, 2018.
[2] Gordon, Jonathan, et al. "Convolutional conditional neural processes." *arXiv preprint arXiv:1910.13556* (2019).
[3] Dubois, Yann, et al. Neural Process Family. http://yanndubs.github.io/Neural-Process-Family/. Sept.2020