Real-time Uncertainty Decomposition for Online Learning Control Supplementary Material

Anonymous authors
Paper under double-blind review

This document provides supplementary information for the ICRL 2021 submission.

Table 1: Evaluation time in seconds

	1D_centered	1D_split	2D_square	2D_gaussian	pmsm_temperature	sarcos
GPmodel	0.004984	0.004733	0.005563	0.083960	1.499310	6.770838
BNN	4.229852	4.257875	11.048753	20.556395	60.889223	121.837305
Dropout	2.660853	2.649146	7.388201	13.246666	38.791683	77.995343
EpiOut	0.121921	0.094260	0.107400	0.215205	0.491637	0.894823

Table 2: Training time seconds

	1D_centered	1D_split	2D_square	2D_gaussian	pmsm_temperature	sarcos
GPmodel	0.221481	0.231211	0.234307	11.773749	545.812856	13845.742551
BNN	102.226074	103.123688	47.040921	456.007891	2222.078883	4415.747424
Dropout	3.485833	3.519852	1.646240	15.074952	72.355323	144.469012
EpiOut	20.289957	18.967542	13.004467	146.365084	2483.204981	12475.670199

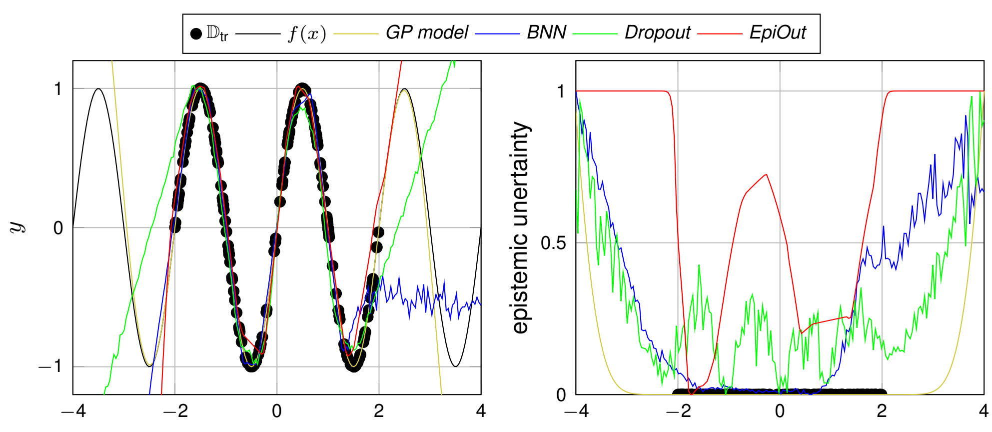

Figure 1: The point estimate (mean prediction) for the different models along with the training data *1D centered* and the true underlying function $f(x) = \sin(\pi x)$ are shown on the left. The right plot shows the epistemic uncertainty estimate.

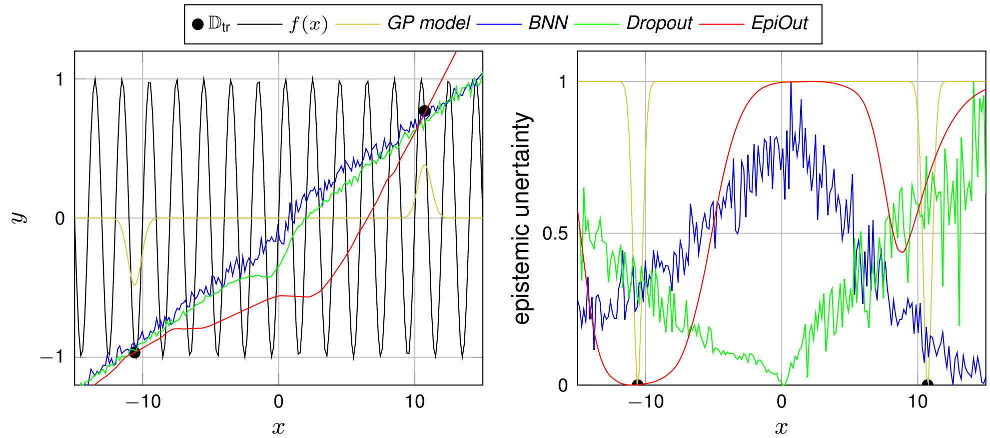

Figure 2: The point estimate (mean prediction) for the different models along with the training data *1D sparse* and the true underlying function $f(x) = \sin(\pi x)$ are shown on the left. The right plot shows the epistemic uncertainty estimate.

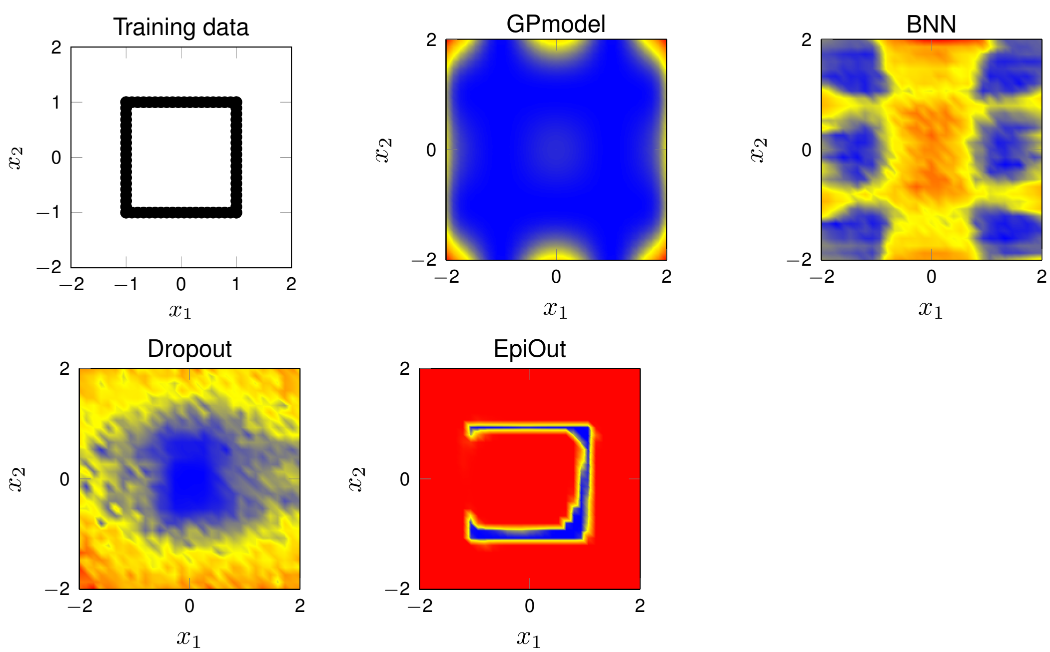

Figure 3: The training data and the predicted epistemic uncertainty by the considered models on the 2D square data set.

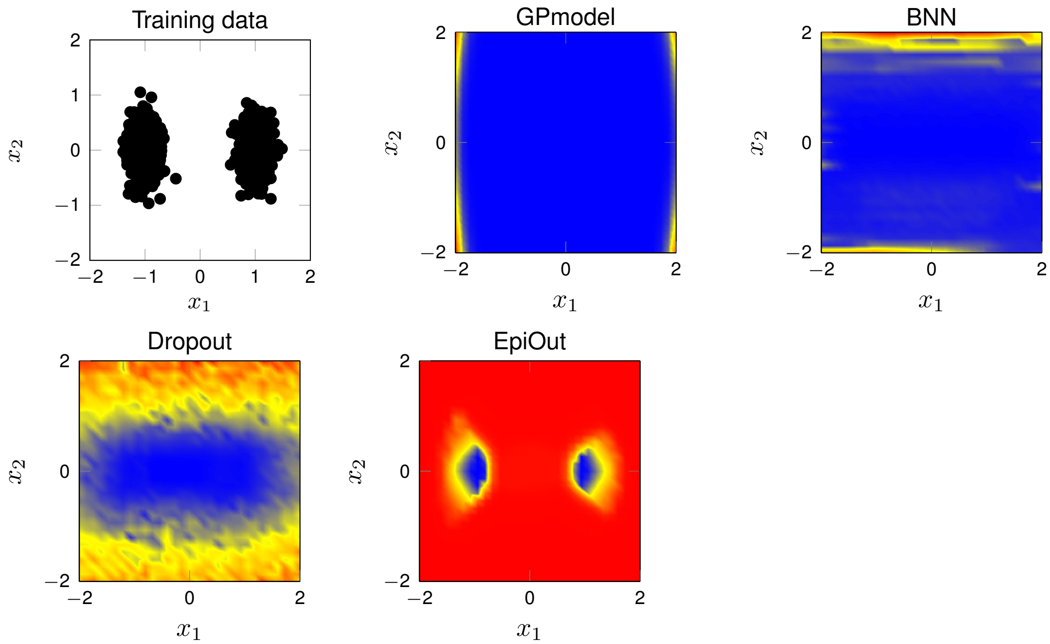

Figure 4: The training data and the predicted epistemic uncertainty by the considered models on the 2D Gaussian data set.

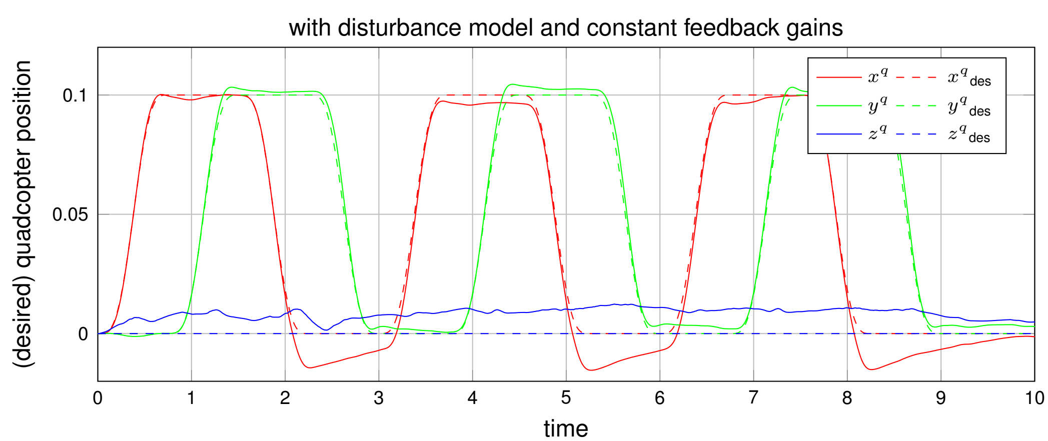

Figure 5: The tracking performance of the quadcopter with disturbance model but constant gains $\beta = 0$ shows an increased tracking error of 0.00944 (vs 0.00733).

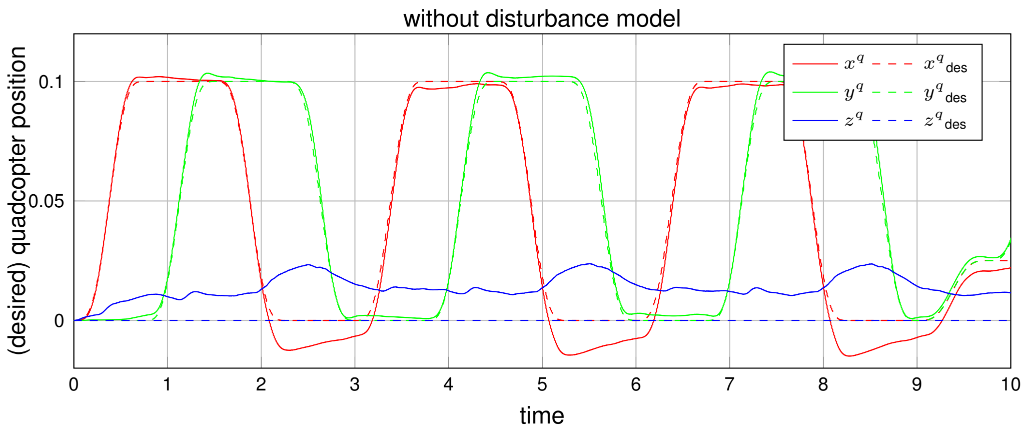

Figure 6: The tracking performance of the quadcopter without disturbance model. It is clearly visible, that the tracking error is larger (RMSE is 0.0153) than with the disturbance model.

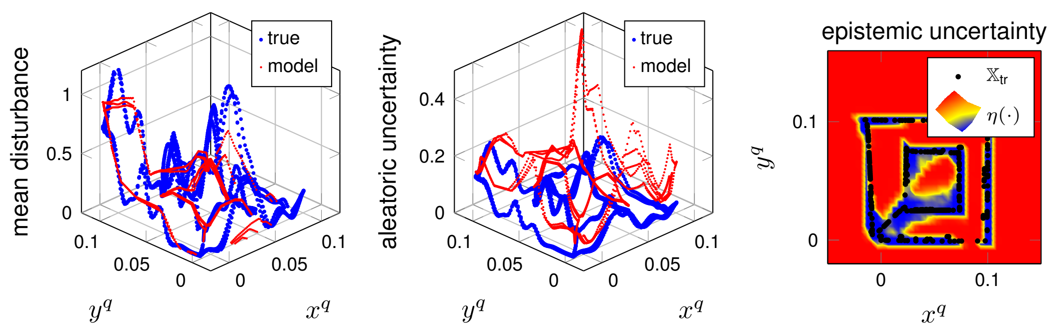

Figure 7: The mean disturbance model (left) captures most of the thermals, and the aleatoric uncertainty (middle) is slightly overestimated by the model. The "true" values are obtained from 1000 samples of the disturbance. The epistemic uncertainty (right) shows that the model is only confident close to the desired trajectory.

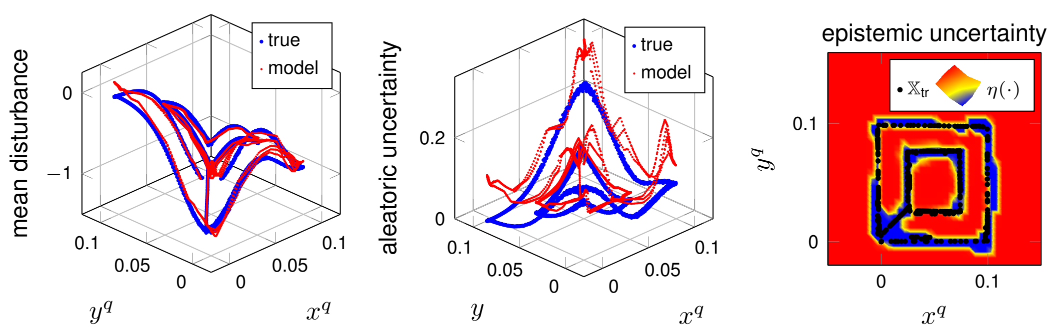

Figure 8: Results for a synthetic disturbance composed of Gaussians.