Concerns regarding Theorem 3.3 Thank you for raising the interesting question on the conditions for asymptotic convergence to which an answer is provided in [1]. There it read as follows: Let $\mathbb{B}_\rho(\boldsymbol{x})$ denote a set of training points around $\boldsymbol{x}$ with radius $\rho > 0$, then the posterior variance converges to zero if there exists a function $\rho(N)$ for which $\rho(N) \leq k(\boldsymbol{x}, \boldsymbol{x})/L_k \ \forall N$, $\lim_{N\to\infty} \rho(N) = 0$, $\lim_{N\to\infty} \left|\mathbb{B}_{\rho(N)}(\boldsymbol{x})\right| = \infty$ holds (L_k: Lipschitz constant of $k(\cdot,\cdot)$). This is achieved e.g. if a constant fraction of all samples lies on the point $\boldsymbol{x}$. We will add the reference [1] and a discussion on the conditions in the paper. In order to address the concerns of Reviewer #2, we will improve the clarity of Theorem 3.3 by reformulating lines 190-191 as follows: "Furthermore, consider an infinite data stream of observations $(\boldsymbol{x}_i, y_i)$ of an unknown function $f : \mathbb{X} \to \mathbb{R}$ with ...". Making Theorem 3.3 quantitative as suggested by Reviewer #2 follows directly from its proof and we will substitute (11) by $P(\sup_{\boldsymbol{x}\in\mathbb{X}} \|\nu_N(\boldsymbol{x}) - f(\boldsymbol{x})\| \in \mathcal{O}(\log(N)^{-\epsilon})) \geq 1 - \delta$.

Boundedness of (3) and (4) As expected by Reviewer #2, (3) and (4) grow with the order of $N^{\frac{1}{2}}$ (see supplementary material (42), (43)). Although unbounded, they grow slow enough to allow the proof of Theorem 3.3 such that the main contribution of these bounds lies in the asymptotic analysis. However, in practical applications there are various ways to estimate tighter constants such as e.g. global optimization. We will add a brief discussion on this in the updated paper.

Noise on input data Reviewer #1 pointed out, that Assumption 3.1. might be violated in the control example due to noise on input data . However, in the presented setup, there is no noise on the input because $f(\cdot)$ does not map from current state to next state, but from the state $\boldsymbol{x}$ to the time derivative of state x_2. Thus input data $\boldsymbol{x}$ and output data $\dot{x}_2$ are measured with two different sensors. Here we made the assumption, that observations of $\dot{x}_2$ are corrupted by noise, while $\boldsymbol{x}$ is measured noise free, which is of course debatable. But in practice, measuring the time derivative is usually realized with finite difference approximations, which injects significantly more noise than a direct measurement. Therefore, Assumption 3.1 is valid for our experimental setup. We will include the given reasoning in the updated paper.

Relation to existing approaches We disagree with Reviewer #2 regarding the originality and significance of our contribution. Even though the bounds in Theorem 3.1 and [2, Theorem 6] look similar, their practical applicability is very different. Once the prior is fixed, all parameters for (7) can be easily computed such that a reliable error bound can be determined. In contrast, [2, Theorem 6] requires the information gain and a bound on the RKHS norm, which is assumed to be known (belonging to an RKHS does not suffice to compute the uniform error bound). In practice, we have observed that these parameters pose a high hurdle which has prevented the rigorous application of this theorem in control applications and typically heuristic constants without theoretical foundations are applied, see e.g. [3]. Therefore, even though both approaches have limits regarding their assumptions, Theorem 3.1 can be rigorously applied in practice, whereas this has been an issue with [2, Theorem 6]. We thank Reviewer #3 for pointing out the previous work [4] which derives a Lipschitz bound approximation for GPs. Although we think this work suggests a valuable estimator for the Lipschitz constant, it does not provide any theoretical guarantees. We will discuss this difference in the updated paper.

Previous work require bounded observation noise Reviewer # 2 argues, that previous work, e.g. [2] are capable of dealing with unbounded noise. Even though [2] generally uses Gaussian noise, the (for this work) most relevant result in [2, Theorem 6] mentions the condition that "the noise variables ϵ_t are uniformly bounded by σ".

Minor comments Thanks for pointing out various typos, we will fix all of them. As suggested by Reviewer #1, we will add a definition of a uniform error bound, extend the proof sketch for Theorem 3.2 and add sketches in the same style for Theorem 3.1, 3.3 and 4.1. Furthermore, reviewer #3 asked to consider a more complex control example: Generally, this is possible within this framework, where (12) becomes $\dot{x}_1 = x_2,\ \dot{x}_2 = \dot{x}_3,\ \cdots\ \dot{x}_d = f(\boldsymbol{x}) + u$, with $\boldsymbol{x} = [x_1\ x_2\ \cdots\ \boldsymbol{x}_d]^\intercal$ using a definition $r = [\boldsymbol{\lambda}^\intercal\ 1]\boldsymbol{e}$ where the coefficients in $\boldsymbol{\lambda} \in \mathbb{R}^{d-1}$ are Hurwitz. The robotic example can also directly be extended to arbitrary degrees of freedoms, however, for the sake of focus on the main results on the error bounds, we would keep the current control examples.

References

[1] A. Lederer, J. Umlauft, and S. Hirche, "Posterior variance analysis of Gaussian processes with application to average learning curves," *arXiv preprint: arXiv:1906.01404*, 2019.

[2] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[3] F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Advances in Neural Information Processing Systems*, 2017, pp. 908–918.

[4] J. González, Z. Dai, P. Hennig, and N. Lawrence, "Batch bayesian optimization via local penalization," in *Artificial intelligence and statistics*, May 2016, pp. 648–657.