

Incremental Sensorimotor Learning with Constant Update Complexity

Arjan Gijsberts and Giorgio Metta

Robotics, Brain, and Cognitive Sciences

Italian Institute of Technology, Genoa, Italy

{arjan.gijsberts, giorgio.metta}@iit.it

Introduction

The field of robotics is increasingly moving toward applications that involve unstructured human environments. This domain is challenging from a learning perspective, since subsequent observations are dependent and the environment is typically non-stationary. This non-stationarity is not limited to the external environment, as internal sensorimotor relationships may be subject to change as well. Specific causes include environmental influences (e.g., temperature) and wear-and-tear, as well as self-initiated changes (e.g., changing dynamics due to tool-use).

Successful modeling of sensorimotor relationships therefore necessitates an open-ended learning process that continuously updates existing models when novel observations become available. Adaptation and prediction have to take place while the robot is operating in the environment and should be time-efficient to ensure responsive behavior. The implied requirements of incremental updates and $\mathcal{O}(1)$ per-sample time complexity, however, are not well supported by current machine learning methods. For instance, kernel methods are “impaired” by formulating the solution in terms of a kernel expansion that grows linearly with the number of training samples. Even algorithms targeted specifically for real-time, incremental robotic learning, such as LWPR (Vijayakumar, D’souza, and Schaal 2005), are characterized by steadily increasing computational requirements as the number of training samples grows ad infinitum. Furthermore, this latter method generally requires a large amount of training samples and manual hyperparameter tuning to attain satisfactory performance.

Incremental Learning with Random Features

Several approaches have been proposed to combine the desired learning performance of kernel methods with incremental updates and bounded computational complexity. Typically, these involve an additional procedure that actively removes samples from the kernel expansion once a predefined budget has been reached (Nguyen-Tuong, Seeger, and Peters 2009, among others). Though effective at bounding the computational complexity, this removal procedure may require costly bookkeeping and invalidates existing theoret-

ical results on learning with kernel methods. More recently, Gijsberts and Metta presented promising results using ridge regression combined with a numerically stable update routine (Gijsberts and Metta 2011). Moreover, non-linearity is achieved by approximating (shift invariant) kernel functions using a finite dimensional feature mapping composed of random projections (Rahimi and Recht 2008). The accuracy of this kernel approximation can be improved arbitrarily by increasing the number of projections, albeit at increased computational cost. Furthermore, since this feature mapping can be considered a kernel in its own right, the method is supported by existing theoretical results on learning with (kernel) ridge regression (Pan and Xiao 2009; Zhdanov and Kalnishkan 2010, for example).

Here we extend this approach by applying the kernel approximation and update routine in the context of Bayesian linear regression. In effect, the resulting method can be considered an incremental variant of the Sparse Spectrum GPR for batch learning (Lázaro-Gredilla et al. 2010). Advantages of this alternative method over the earlier ridge regression variant are the availability of a predictive distribution and convenient hyperparameter optimization using log likelihood maximization. These advantages are acquired at the cost of slightly increased computational requirements and a Gaussianity assumption on the signal noise.

Results

The incremental version of SSGPR is evaluated on the task of learning inverse manipulator dynamics, using nearly half an hour of samples collected from the iCub humanoid (Gijsberts and Metta 2011). For this experiment, the feature mapping is initialized to approximate an anisotropic Gaussian kernel using 50 random projections. The resulting method is compared with incremental LWPR, Gaussian Process Regression (GPR) trained in batch (Rasmussen and Williams 2005), and an analytical Rigid Body Dynamics (RBD) model (Siciliano and Khatib 2008). The initial 15000 samples are used for batch training in case of GPR, and for hyperparameter optimization and calibration in case of incremental SSGPR, LWPR, and RBD. Both incremental learning methods therefore start training during evaluation on the test samples.

Fig. 1 demonstrates how the average prediction error develops as the number of test samples increases. Interestingly,

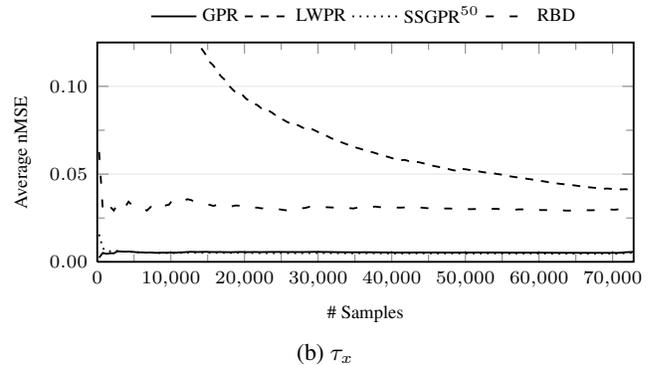
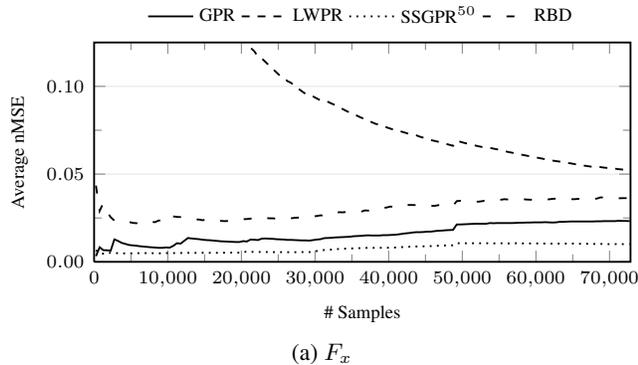


Figure 1: Convergence of the average one-step-ahead prediction error of the considered methods when predicting the force F and torque τ in the x -axis of the sensor reference frame. The results for SSGPR⁵⁰ are the average error over 25 randomized runs. The standard deviation over the various runs is negligible and error bars are therefore omitted for clarity.

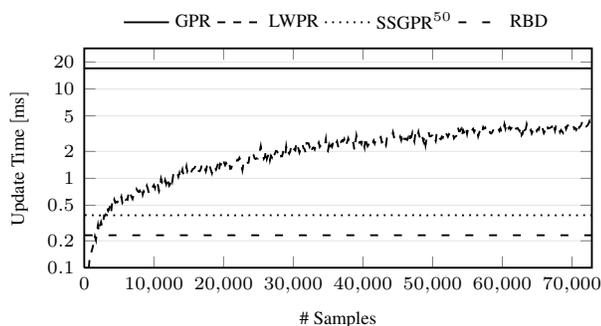


Figure 2: Sample time with respect to the number of training samples for the force components of the iCub dataset. The results for GPR and RBD only include the prediction time, whereas the timing for LWPR and SSGPR also includes the model update. Note the log-scale on the y -axis.

GPR and RBD exhibit decreasing accuracy when predicting forces, which can be contributed to temperature induced sensor drift. Incremental SSGPR is notably less affected by this drift, since continuous updates to its model help to compensate for these gradual changes. Furthermore, the results demonstrate that 50 random features are sufficient to outperform the competing methods and to match the full GPR when predicting torques. Although LWPR seemingly handles the sensor drift well, its performance is significantly inferior and characterized by slow convergence. In addition, the corresponding timing results (cf. Fig. 2) confirm its increasing computational requirements. This renders LWPR unsuitable for applications requiring open-ended learning or involving hard real-time constraints. The per-sample update time of incremental SSGPR remains constant over time, as per design, and even approaches the efficient analytical RBD model. This demonstrates that the proposed method achieves excellent generalization performance while requiring a bounded and predictable time to process a sample.

Acknowledgments

This work was supported by European Commission project ITALK (ICT-214668).

References

- Gijsberts, A., and Metta, G. 2011. Incremental learning of robot dynamics using random features (accepted). In *2011 IEEE International Conference on Robotics and Automation*.
- Lázaro-Gredilla, M.; Quiñonero-Candela, J.; Rasmussen, C. E.; and Figueiras-Vidal, A. R. 2010. Sparse spectrum gaussian process regression. *Journal of Machine Learning Research* 11:1865–1881.
- Nguyen-Tuong, D.; Seeger, M. W.; and Peters, J. 2009. Model learning with local gaussian process regression. *Advanced Robotics* 23(15):2015–2034.
- Pan, Z.-W., and Xiao, Q.-W. 2009. Least-square regularized regression with non-iid sampling. *Journal of Statistical Planning and Inference* 139(10):3579–3587.
- Rahimi, A., and Recht, B. 2008. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20*. MIT Press. 1177–1184.
- Rasmussen, C. E., and Williams, C. K. I. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. MIT Press.
- Siciliano, B., and Khatib, O. 2008. *Springer Handbook of Robotics*. Springer.
- Vijayakumar, S.; D’souza, A.; and Schaal, S. 2005. Incremental online learning in high dimensions. *Neural Computation* 17(12):2602–2634.
- Zhdanov, F., and Kalnishkan, Y. 2010. An identity for kernel ridge regression. In *Proceedings of the 21st International Conference on Algorithmic Learning Theory, ALT’10*, 405–419. Springer-Verlag.