A Bayesian Regression Approach to Terrain Modeling

C. Plagemann[1], S. Mischke[1], K. Kersting[2], S. Prentice[3], N. Roy[3], W. Burgard[1]

[1]Univ. of Freiburg, Dept. of Computer Science, Germany
[2]Fraunhofer Institute IAIS, Sankt Augustin, Germany
[3]Massachusetts Institute of Technology, CSAIL, Cambridge, USA

Surface models that represent the distribution of elevations at given points in the plane are important in many areas such as outdoor robotics or the geo-sciences. We introduce a nonparametric, probabilistic regression framework based on locally-adaptive Gaussian processes for learning such models from sets of elevation samples. Our approach does not require a discretization of the space, it can be learned efficiently by making justified model approximations, and it is fully predictive in the sense that, for every point in the plane, a predictive distribution of elevation values can be estimated.

We discuss several ways of representing and learning local smoothness in our model, which allow the user to balance model accuracy and computational efficiency depending on the application scenario.

We demonstrate the usefulness of our approach in difficult test settings, including a quadruped robot equipped with a light-weight laser range finder, which observes rocky terrain, plans a path as well as a sequence of foot placements, and then executes its plan (see Figure).

Related Publications

[1] C. Plagemann, S. Mischke, S. Prentice, K. Kersting, N. Roy, and W. Burgard. Learning Predictive Terrain Models for Legged Robot Locomotion. In Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nice, France, 2008. To appear.

[2] C. Plagemann, K. Kersting, and W. Burgard. Nonstationary Gaussian Process Regression using Point Estimates of Local Smoothness. In European Conference on Machine Learning (ECML), Antwerp, Belgium, 2008. To appear.

[3] T. Lang, C. Plagemann, and W. Burgard. Adaptive Nonstationary Kernel Regression for Terrain Modeling. In Robotics: Science and Systems (RSS), Atlanta, Georgia, USA, June, 2007.

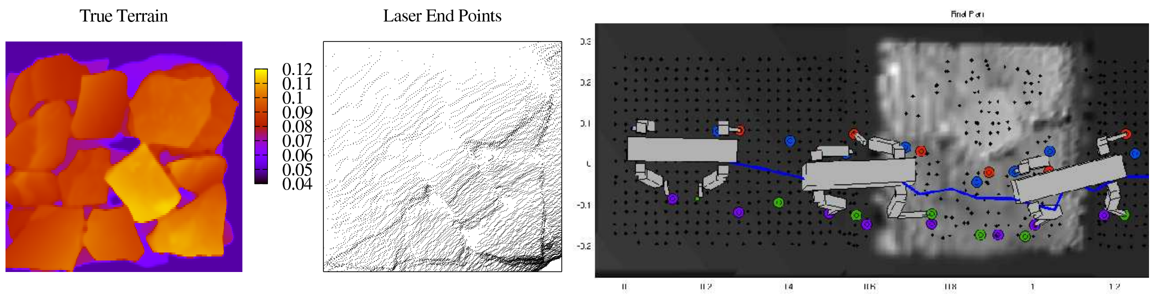

Figure 1: Our goal is to recover the true elevation function (left) from a set of elevation samples (middle). As an application example, we discuss a quadruped robot walking over rocky terrain (right).