Active Filtering for robotic tactile learning

Hannes P. Saal
School of Informatics
University of Edinburgh
Edinburgh, UK
h.saal@sms.ed.ac.uk

Jo-Anne Ting
School of Informatics
University of Edinburgh
Edinburgh, UK
jting@ed.ac.uk

Sethu Vijayakumar
School of Informatics
University of Edinburgh
Edinburgh, UK
sethu.vijayakumar@ed.ac.uk

1 Motivation

The problem in active learning is generally to determine controls $\mathbf{x}$ such that a parameter $\boldsymbol{\theta}$ can be estimated as accurately as possible given observations $\mathbf{y}$ and an observation model $\mathbf{y} = f(\boldsymbol{\theta}, \mathbf{x})$. In the sequential setting, a distribution over $\boldsymbol{\theta}$ is updated after each observation, and new controls $\mathbf{x}$ are selected to maximize the mutual information:

$$argmax_x I(\boldsymbol{\theta}; \mathbf{y}|\mathbf{x}) = argmax_x \int \int p(\boldsymbol{\theta}, \mathbf{y}|\mathbf{x}) \, log \frac{p(\boldsymbol{\theta}, \mathbf{y}|\mathbf{x})}{p(\boldsymbol{\theta}) p(\mathbf{y}|\mathbf{x})} \, d\boldsymbol{\theta} \, d\mathbf{y}$$

We focus on problems with robotics applications, where we are faced with some additional constraints: decisions have to be taken quickly, sometimes in real-time, and observations and controls are likely to be high-dimensional. Moreover, in our setting, observation functions are not available in analytical form, are likely to be nonlinear and need to be learned from training data. The evaluation of mutual information is, in these cases, analytically intractable, and sampling methods become untenable in high dimensions.

While we are mainly interested in the problem where $\boldsymbol{\theta}$ is constant, our methods can easily applied to problems with time-varying theta (i.e. state dynamics).

In the following, we present our methods, and then go on to describe experimental evaluation on a real robotic setup.

2 Methods and Results

We use Gaussian process regression for learning the observation function. This method models the predictive distribution as a Gaussian, where mean and variance are determined from the training data. This leads to the effect that the predictive variance will be high in regions with no training data, leading to low information there.

In the following, we present two different approaches for active filtering, based on Gaussian process observation functions.

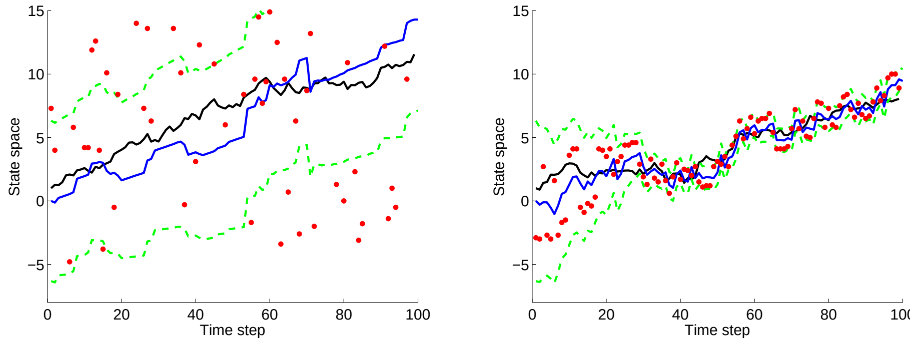

Figure 1: Comparison of standard filter with random actions (left) and active filter (right). The black line denotes the target which moves randomly with a bias term in each time step. The blue line denotes the mean of the state belief in each time step. The green lines denote the 95%-confidence interval of the belief. As can be seen, the active filter converges much faster as only informative controls are chosen. Red dots correspond to controls (plotted in the state space). The observation function in this example is a Gaussian bump function which is only high when the controls are close to the target position.

2.1 Assumed Density Filter

Recently, an assumed density filter using Gaussian processes has been proposed [1]. We extend this filter to the active setting by noting that the mutual information can be calculated analytically, as it reduces to a function of the variance of the posterior, i.e. the new criteron is:

$$argmin_x|\mathbf{\Sigma} - \mathbf{Z}(\mathbf{x})^T\mathbf{Q}(\mathbf{x})\mathbf{Z}(\mathbf{x})|$$

where $\mathbf{\Sigma}$ is the current state belief distribution, $\mathbf{Z}$ is the cross-covariance between the current state distribution and the observations, and $\mathbf{Q}$ is the marginal observation variance.

Using gradient ascent allows us to optimize actions quickly at each time step.

This filter converges considerably faster than a filter with random controls (see Figure 1).

2.2 Particle filter

Filtering problems where the observation function is highly non-linear can pose problems for filters using linear updates and cause them to diverge. A popular alternative are particle filters. When using these, we are faced with the problem that the optimization step is very slow as all integrals involved have to be solved using Monte-Carlo integration. Additionally, there is no analytic expression for the gradient of the mutual information. This problem becomes more pronounced in higher dimensions, where the number of samples required goes up exponentially.

In our approach, we use a quadratic divergence measure [2] instead of standard mutual information:

$$I_q(\boldsymbol{\theta};\mathbf{y}|\mathbf{x}) = \int\int (p(\boldsymbol{\theta},\mathbf{y}|\mathbf{x}) - p(\boldsymbol{\theta})p(\mathbf{y})|\mathbf{x})^2 \, d\boldsymbol{\theta} \, d\mathbf{y}$$

This measure allows us to solve integrals over $\mathbf{y}$ analytically, and therefore speed up the computation of the information criterion considerably—as opposed to sampling the full mutual information.

3 Experimental evaluation

We apply this framework to the problem of determining dynamic object properties from tactile sensor data gathered by touch sensors mounted on a robotic hand-arm system (see Figure 2). More specifically, the robot's task is to determine the viscosity of liquids by shaking bottles containing

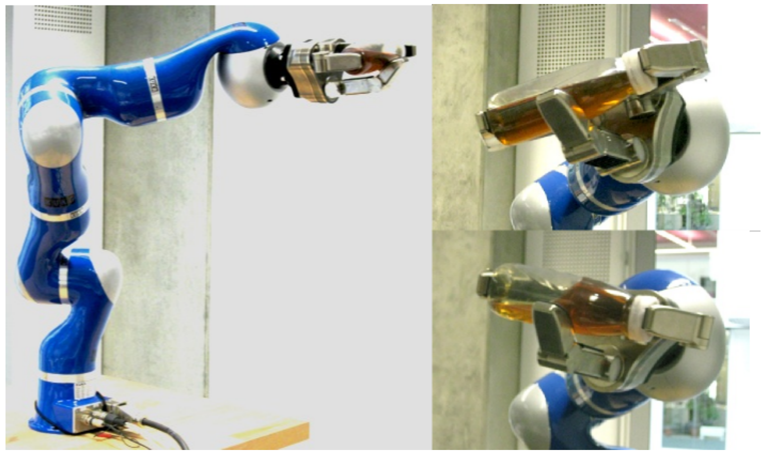

Figure 2: Experimental setup of robotic viscosity discrimination. Left: DLR light-weight arm with Schunk dexterous hand attached. Right: Close-up of the robot hand while shaking a liquid-filled bottle.

those liquids. The robot has to actively set the shaking frequency and shaking direction so as to maximize information gain. In this setting, all quantities involved are continuous, and the observation space is high-dimensional (23D). We find that convergence to the correct viscosity is significantly faster using active filtering than when using random controls to govern the shaking movements.

Acknowledgments

HPS is supported through the EPSRC and MRC by a Neuroinformatics Doctoral Training Centre scholarship. This study was supported by the EU-funded SENSOPAC project.

References

[1] Deisenroth, M.P., Huber, M.F., and Hanebeck, U.D. (2009) Analytic moment-based Gaussian Process filtering. *Proceedings of the 26th International Conference on Machine Learning*, Montreal, Canada.

[2] Torkkola, K. (2003) Feature extraction by non parametric information maximization. *Journal of Machine Learning Research*, 3: 1415–1438.