Active Sequential Learning with Tactile Feedback

Jo-Anne Ting

University of British Columbia

jating@cs.ubc.ca

RSS 2010 Workshop: Active Learning for Robotics

June 27, 2010

Estimating dynamical properties of objects

- Weight
- Centre of mass
- Compressibility
- Viscosity
- ...

Joint work with Hannes Saal & Sethu Vijayakumar (EDI)

THE UNIVERSITY of EDINBURGH
informatics

Relevant papers: Saal, Ting & Vijayakumar (AISTATS 2010); Saal, Vijayakumar & Johansson (J. of Neuro. 2009)

Problem formulation

Task: To determine dynamical variables quickly and reliably

Challenges:

- Learning problem: Need mapping between actions and sensory observations
- Control problem: To find actions that are most informative

How do humans do it?

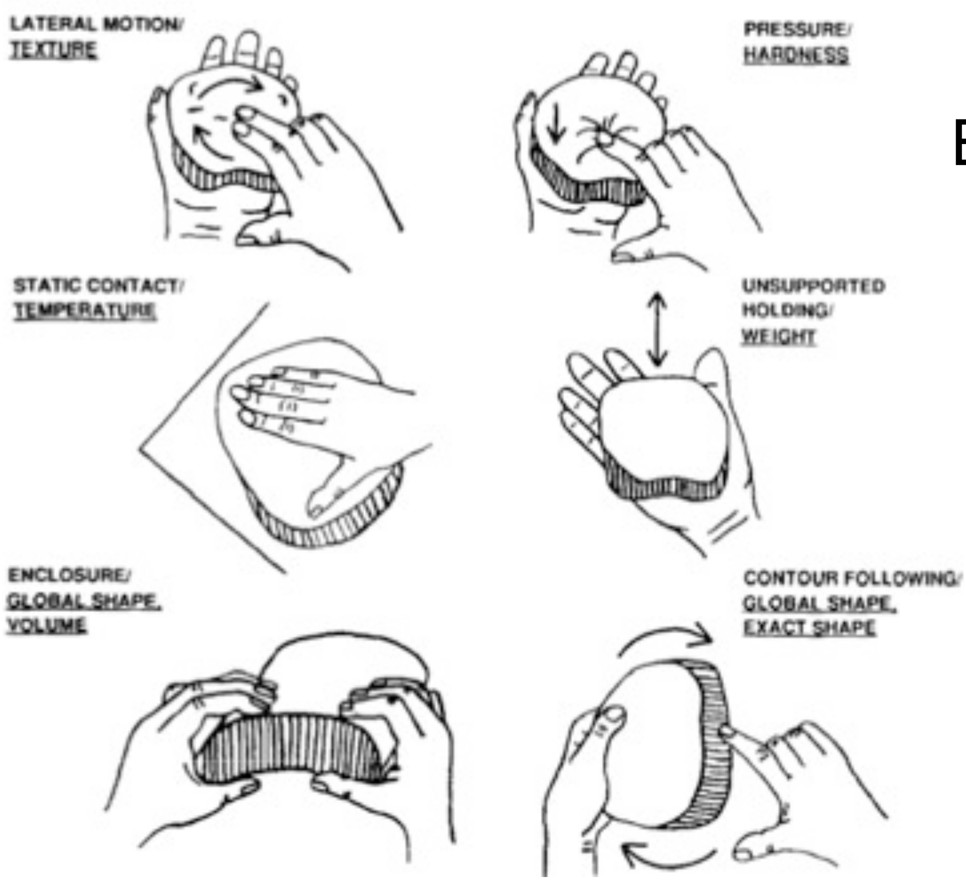

Lederman & Klatzky, 1993

Exploratory actions to start

↓

Highly informative actions once know more about object

Example approach: tactile feature discrimination

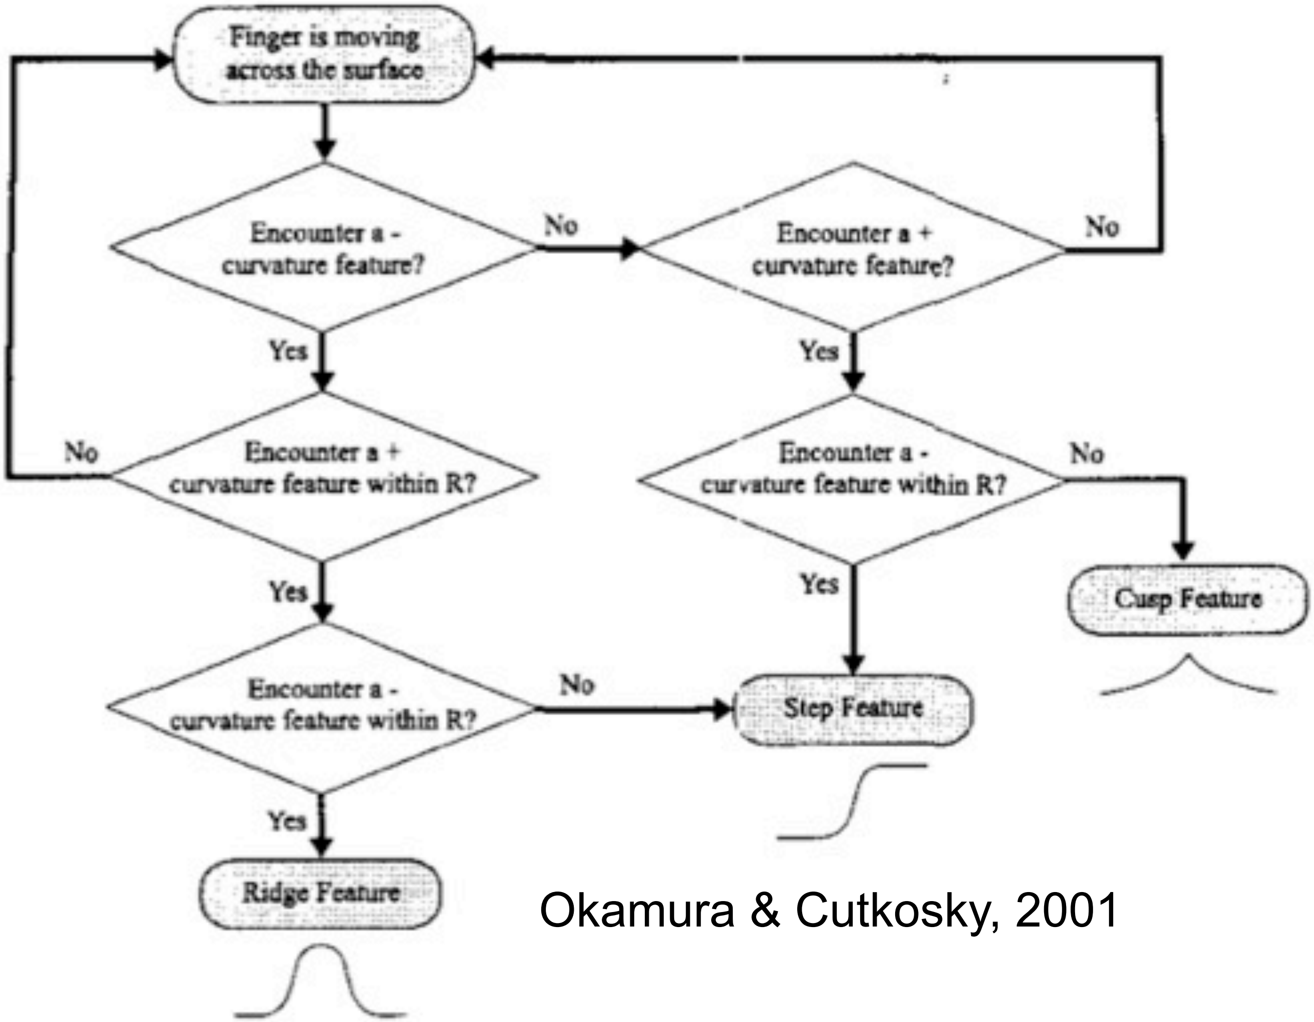

Okamura & Cutkosky, 2001

Example application: estimating viscosity

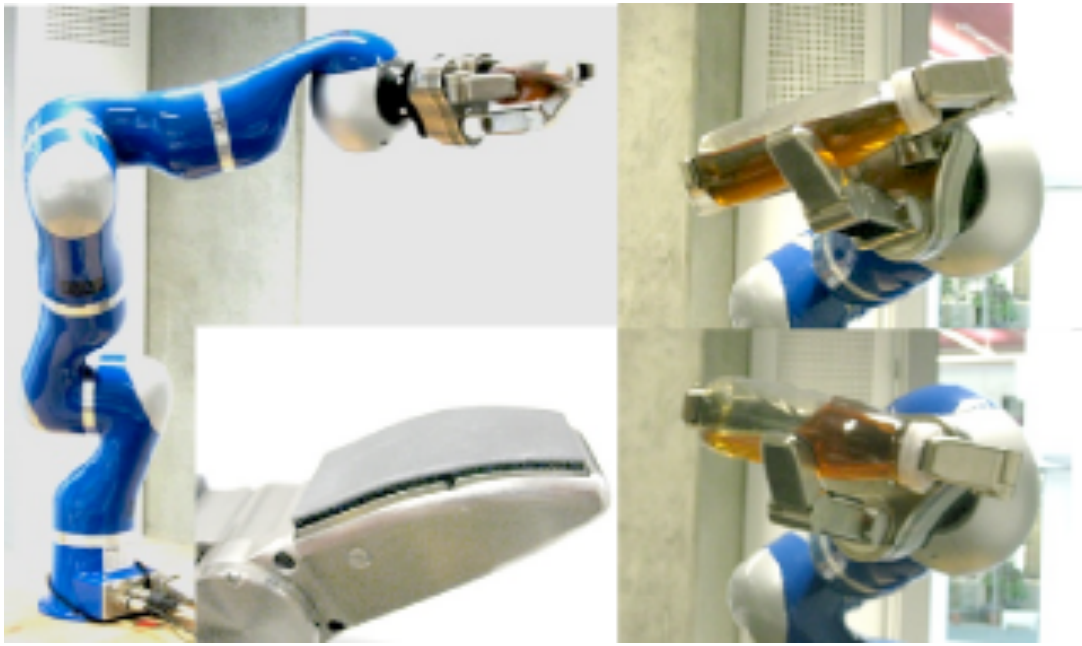

- Parameter: viscosity θ
- Observations: tactile feedback $\mathbf{y}$
- Actions: shaking frequency, angle, ... $\mathbf{x}$

Estimating viscosity

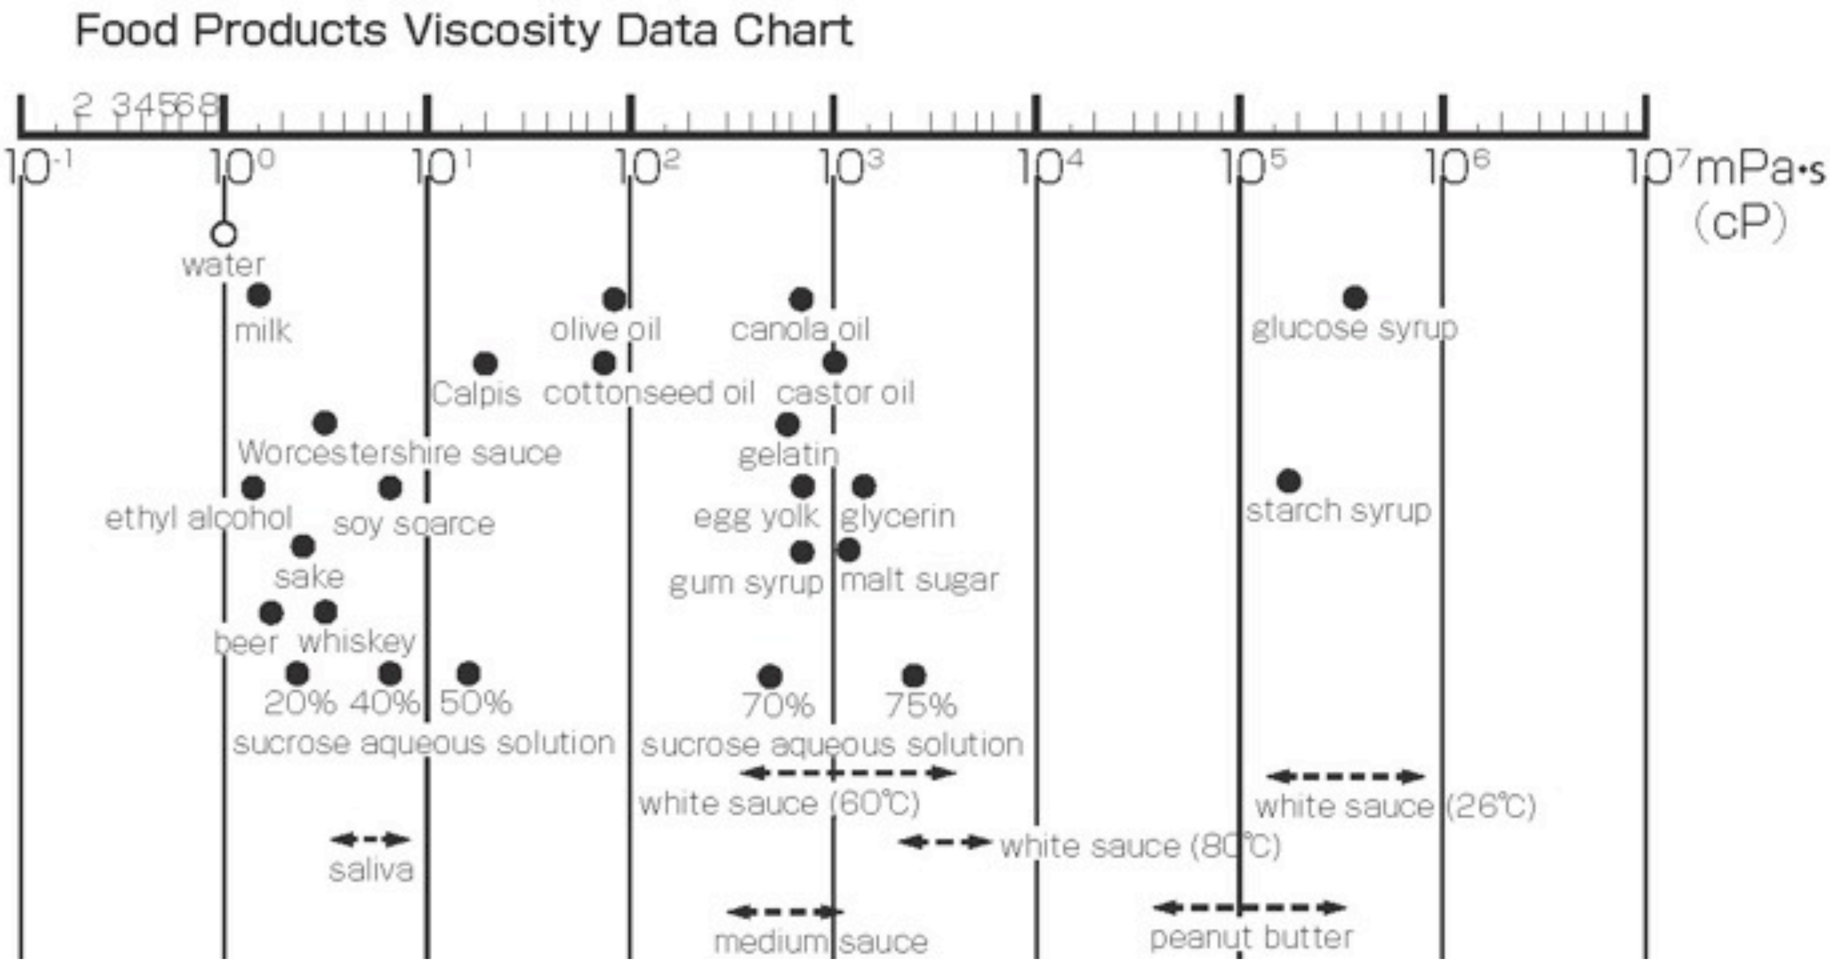

“Active” sensing

- Approach: Choose actions $\mathbf{x}$ such that mutual information between $\mathbf{y}$ and θ is maximized, e.g.,

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathbf{X}} \mathbf{I}(\theta; \mathbf{y} | \mathbf{x})$$ where

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

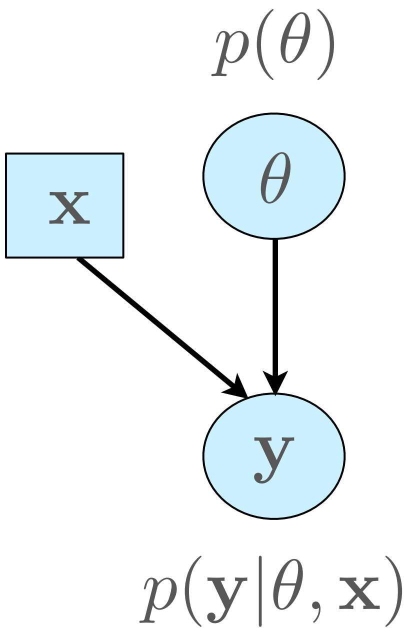

Active sensing: observation model

- Assuming $\mathbf{y} = f(\mathbf{x}, \theta) + \epsilon_{\mathbf{y}}$,

  we can place a Gaussian Process prior (Williams & Rasmussen, 1995) over f:

$$y_m(\theta, \mathbf{x}) \sim \mathrm{GP}(0, \mathbf{K}_m)$$

- Predictive distribution over $\mathbf{y}$ is Gaussian.

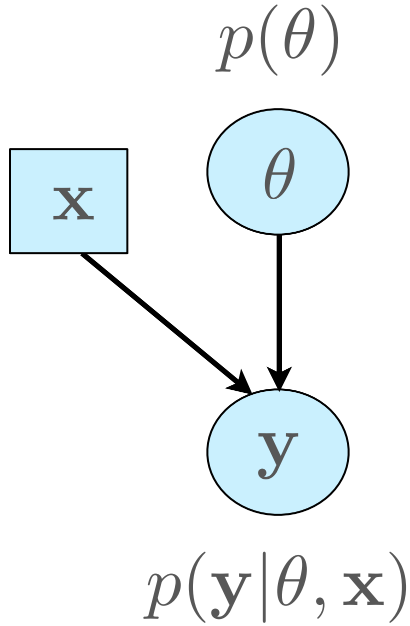

Active sensing: calculating the posterior

- Given prior $p_{t-1}(\theta)$, action $\mathbf{x}_t$, and observation $\mathbf{y}_t$, posterior is:

$$p_t(\theta|\mathbf{y}_t, \mathbf{x}_t) = \frac{p(\mathbf{y}_t|\theta, \mathbf{x}_t)p_{t-1}(\theta)}{p(\mathbf{y}_t|\mathbf{x}_t)}$$

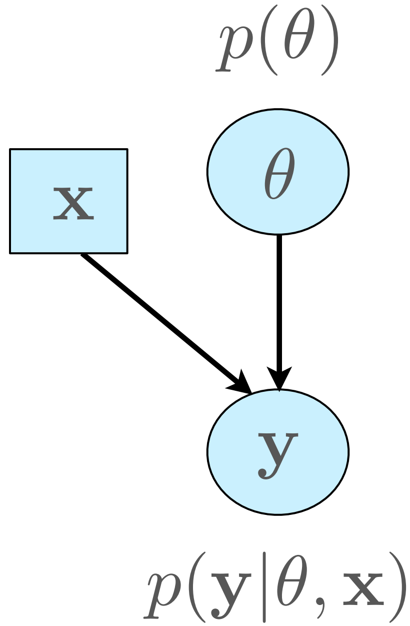

- Evaluate with MC sampling or Gaussian approximation.

Active sensing: calculating the posterior

- If $p_t(\theta|\mathbf{y}_t, \mathbf{x}_t) = N(\mu_t, \mathbf{\Sigma}_t)$, then we get:

$$\mu_t = \mu_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} (\mathbf{y}_t^{obs} - \mathbf{m}_t)$$
$$\mathbf{\Sigma}_t = \mathbf{\Sigma}_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} \mathbf{C}_t$$

Active sensing: calculating the posterior

- If $p_t(\theta|\mathbf{y}_t, \mathbf{x}_t) = N(\mu_t, \mathbf{\Sigma}_t)$, then we get:

$$\mu_t = \mu_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1}(\mathbf{y}_t^{obs} - \mathbf{m}_t)$$
$$\mathbf{\Sigma}_t = \mathbf{\Sigma}_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} \mathbf{C}_t$$

cross-covariance

marginal mean

marginal variance

Active sensing: calculating the posterior

- If $p_t(\theta|\mathbf{y}_t, \mathbf{x}_t) = N(\mu_t, \mathbf{\Sigma}_t)$, then we get:

$$\mu_t = \mu_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} (\mathbf{y}_t^{obs} - \mathbf{m}_t)$$
$$\mathbf{\Sigma}_t = \mathbf{\Sigma}_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} \mathbf{C}_t$$

Active sensing: calculating the posterior

- If $p_t(\theta|\mathbf{y}_t, \mathbf{x}_t) = N(\mu_t, \mathbf{\Sigma}_t)$, then we get:

$$\mu_t = \mu_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1}(\mathbf{y}_t^{obs} - \mathbf{m}_t)$$
$$\mathbf{\Sigma}_t = \mathbf{\Sigma}_{t-1} + \mathbf{C}_t^T \mathbf{S}_t^{-1} \mathbf{C}_t$$

where $\mathbf{m}$, $\mathbf{S}$, $\mathbf{C}$ can be evaluated analytically (e.g., Girard et al. 2003). For example:

$$\mathbf{m} = \int p(\mathbf{y})\mathbf{y}d\mathbf{y} = \int p(\theta)p(\mathbf{y}|\theta)\mathbf{y}d\theta d\mathbf{y}$$

Active sensing: entire loop

- Start with a prior $p_{t-1}(\theta)$
- Take an action $\mathbf{x}_t^*$

$$\begin{aligned}\mathbf{x}_t^* &= \arg\max_{\mathbf{x}_t} \mathbf{I}(\theta_{t-1}; \mathbf{y}_t | \mathbf{x}_t) \\ &= \arg\max_{\mathbf{x}_t} |\mathbf{C}_t(\mathbf{x})\mathbf{S}_t(\mathbf{x})\mathbf{C}_t(\mathbf{x})^T|\end{aligned}$$

- Calculate posterior $p_t(\theta)$
- Repeat loop until convergence

Challenges

- Observations, actions & parameters are continuous
- Sensor model has to be **learnt** from data
- Observations are high-dimensional
- Decisions have to be taken quickly

Viscosity example: learnt model

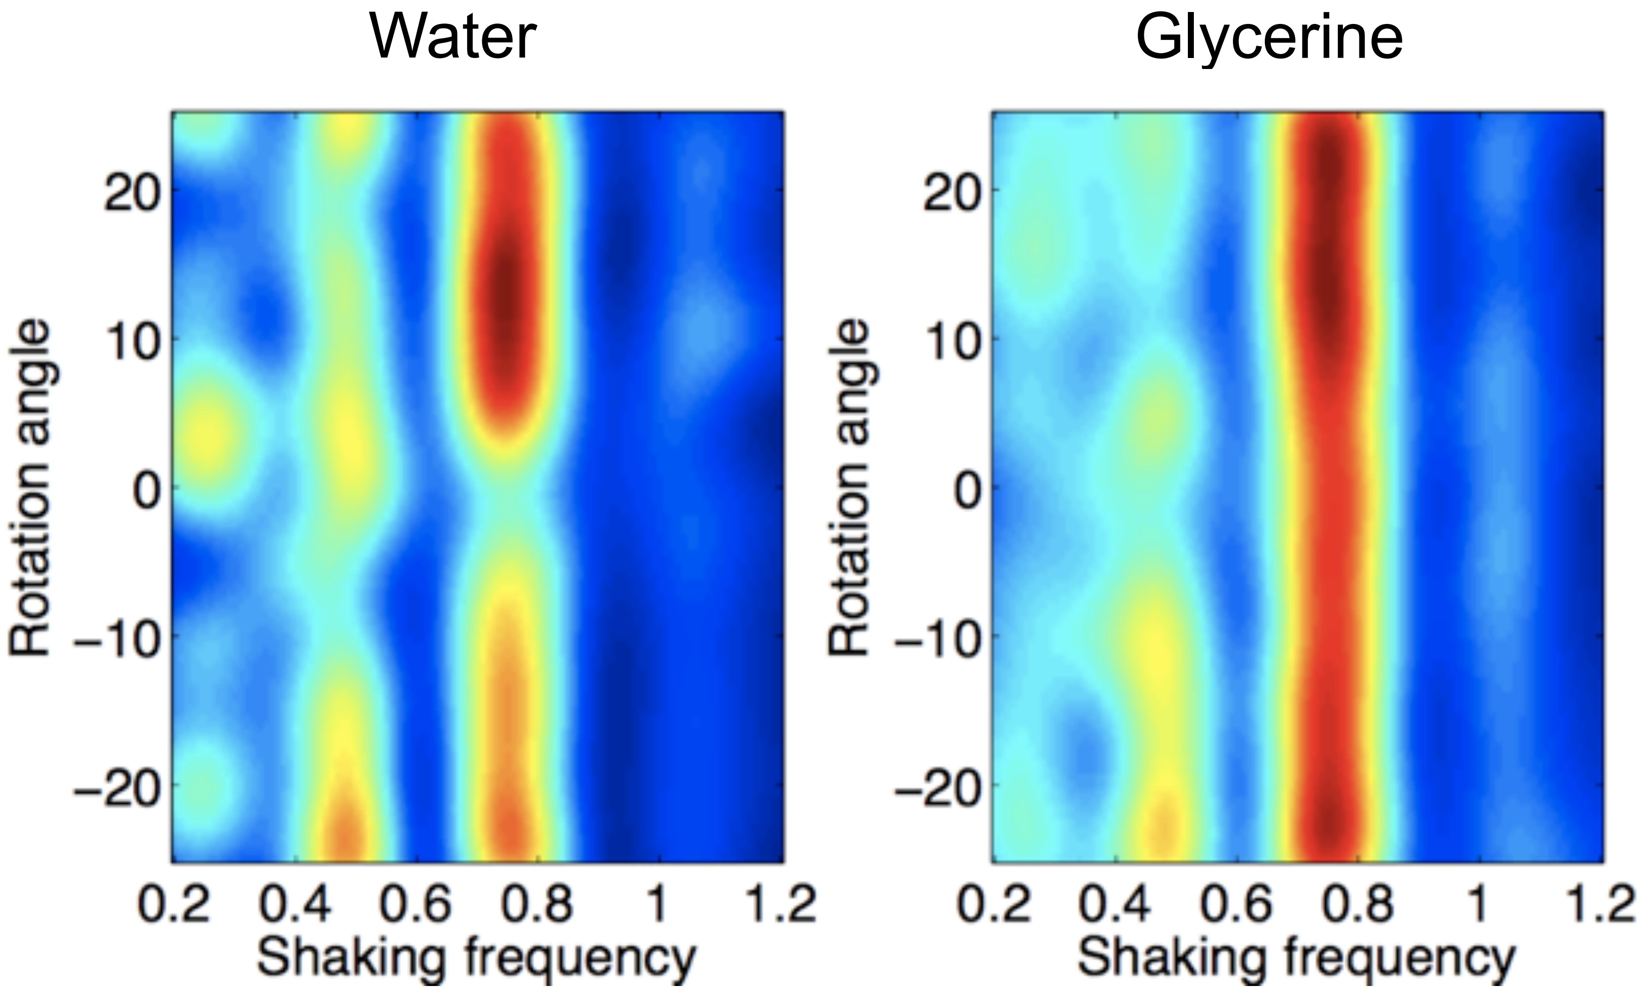

Viscosity example: video

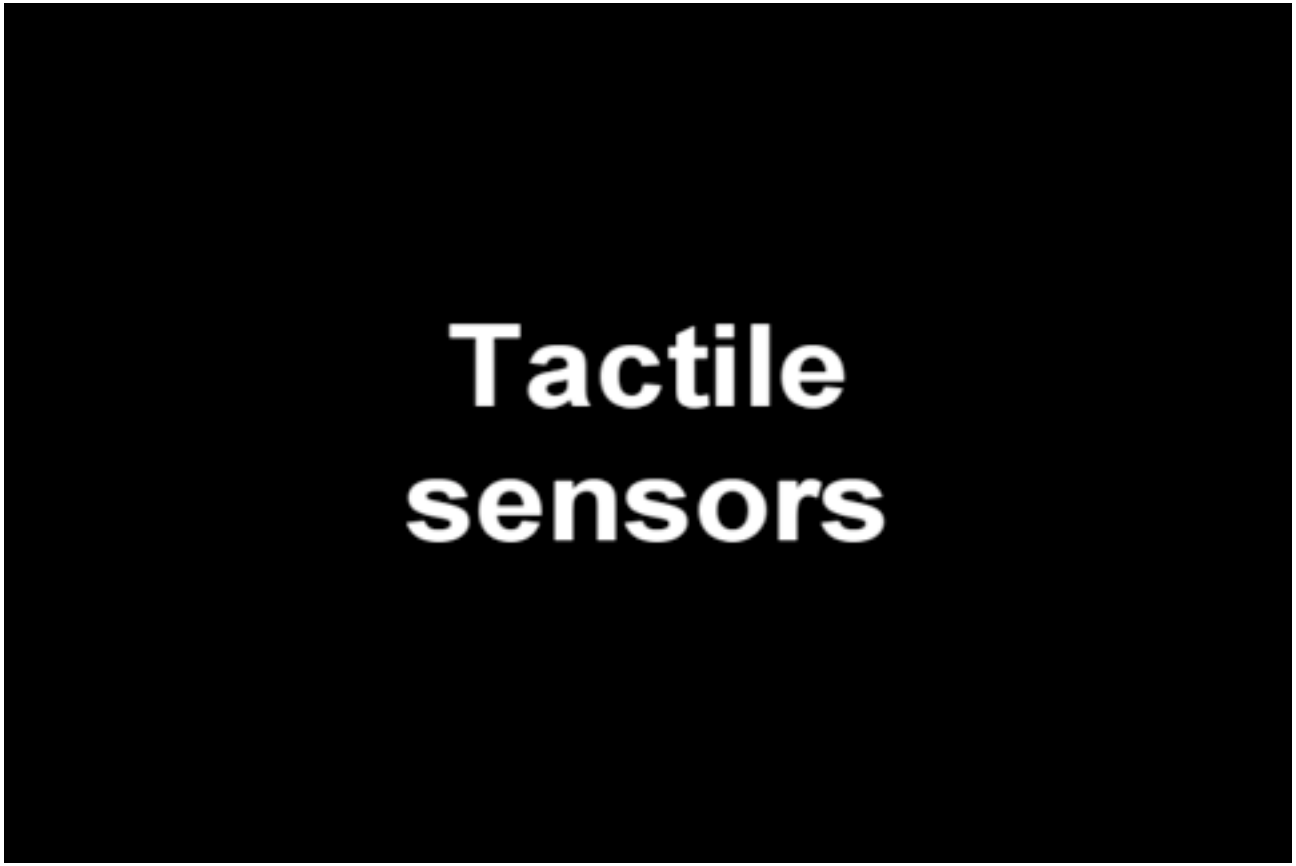

Information landscape

Broad prior $p(\theta)$

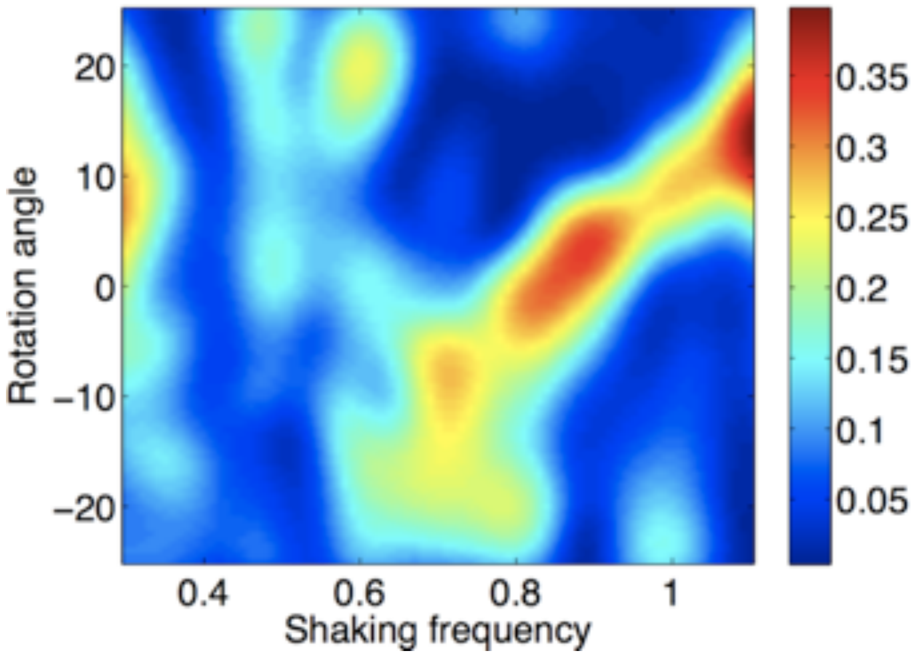

Narrow prior $p(\theta)$

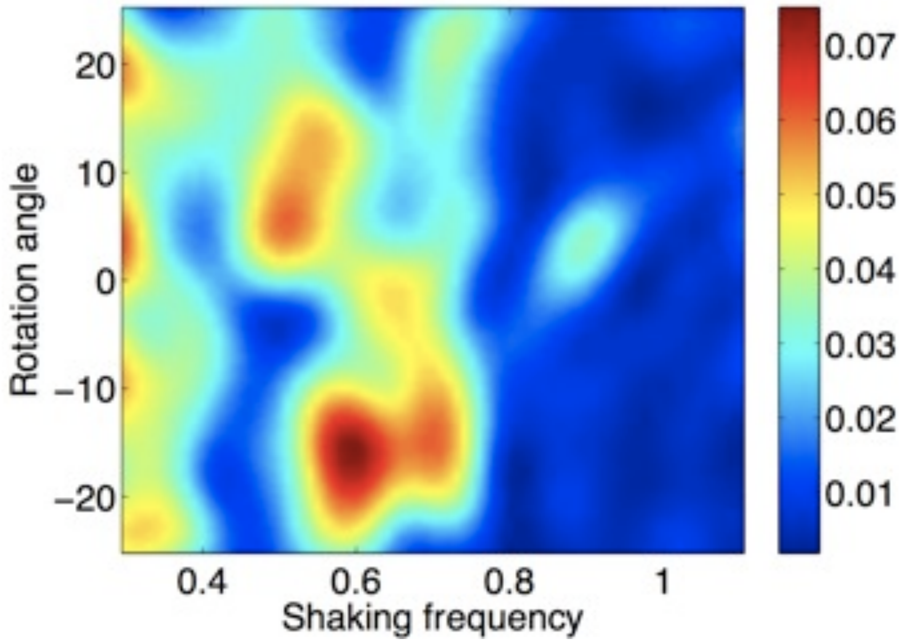

Sequentially estimating viscosity

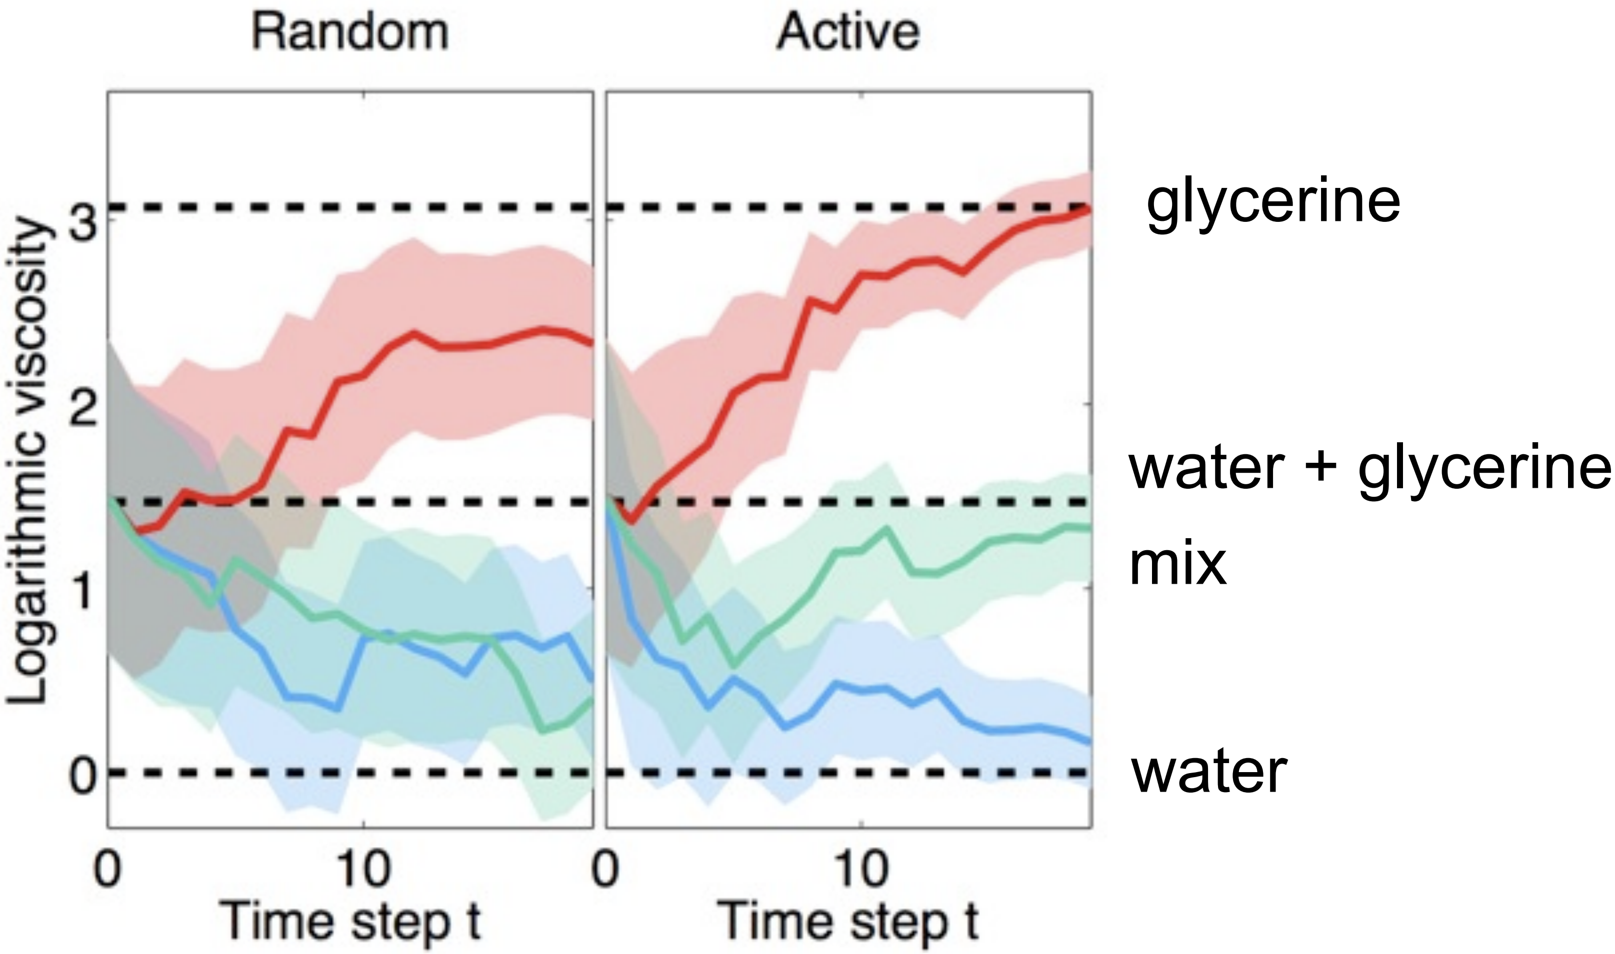

Comparing different action strategies

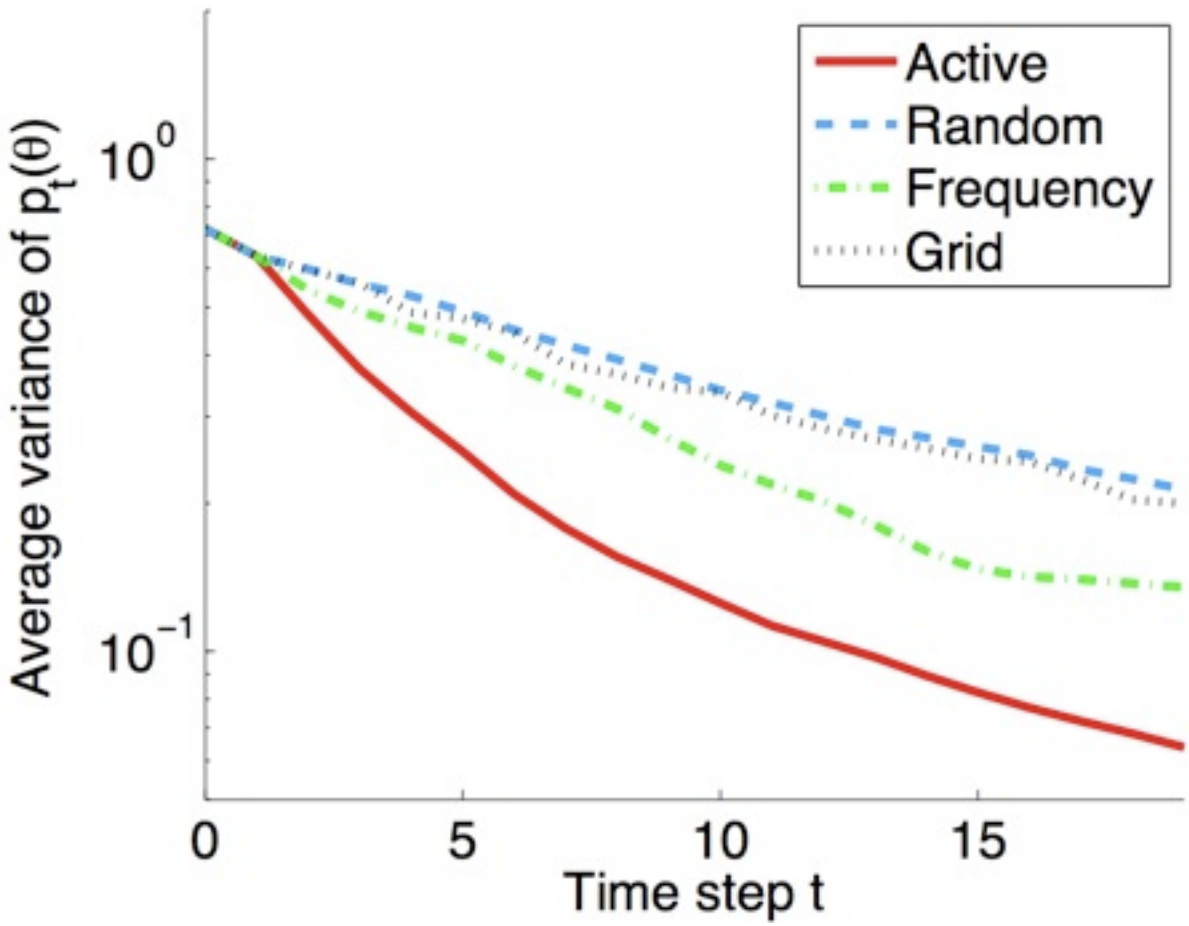

On-going and future work

- Scale up to higher-dimensional action space
- Include sparse versions of GPs for large data sets
- Extend to multi-task GPs (for multivariate regression)
- Find invariances in actions
- ...