# Preference Galleries for Material Design

Eric Brochu, Abijeet Ghosh and Nando de Freitas

THE UNIVERSITY OF BRITISH COLUMBIA

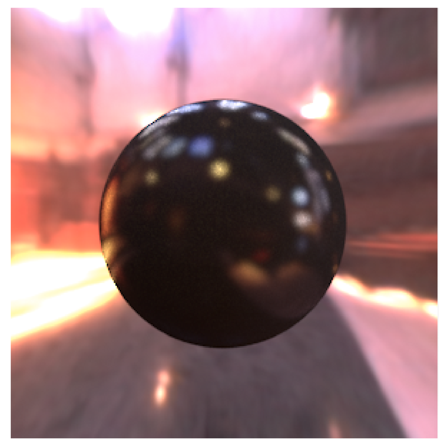

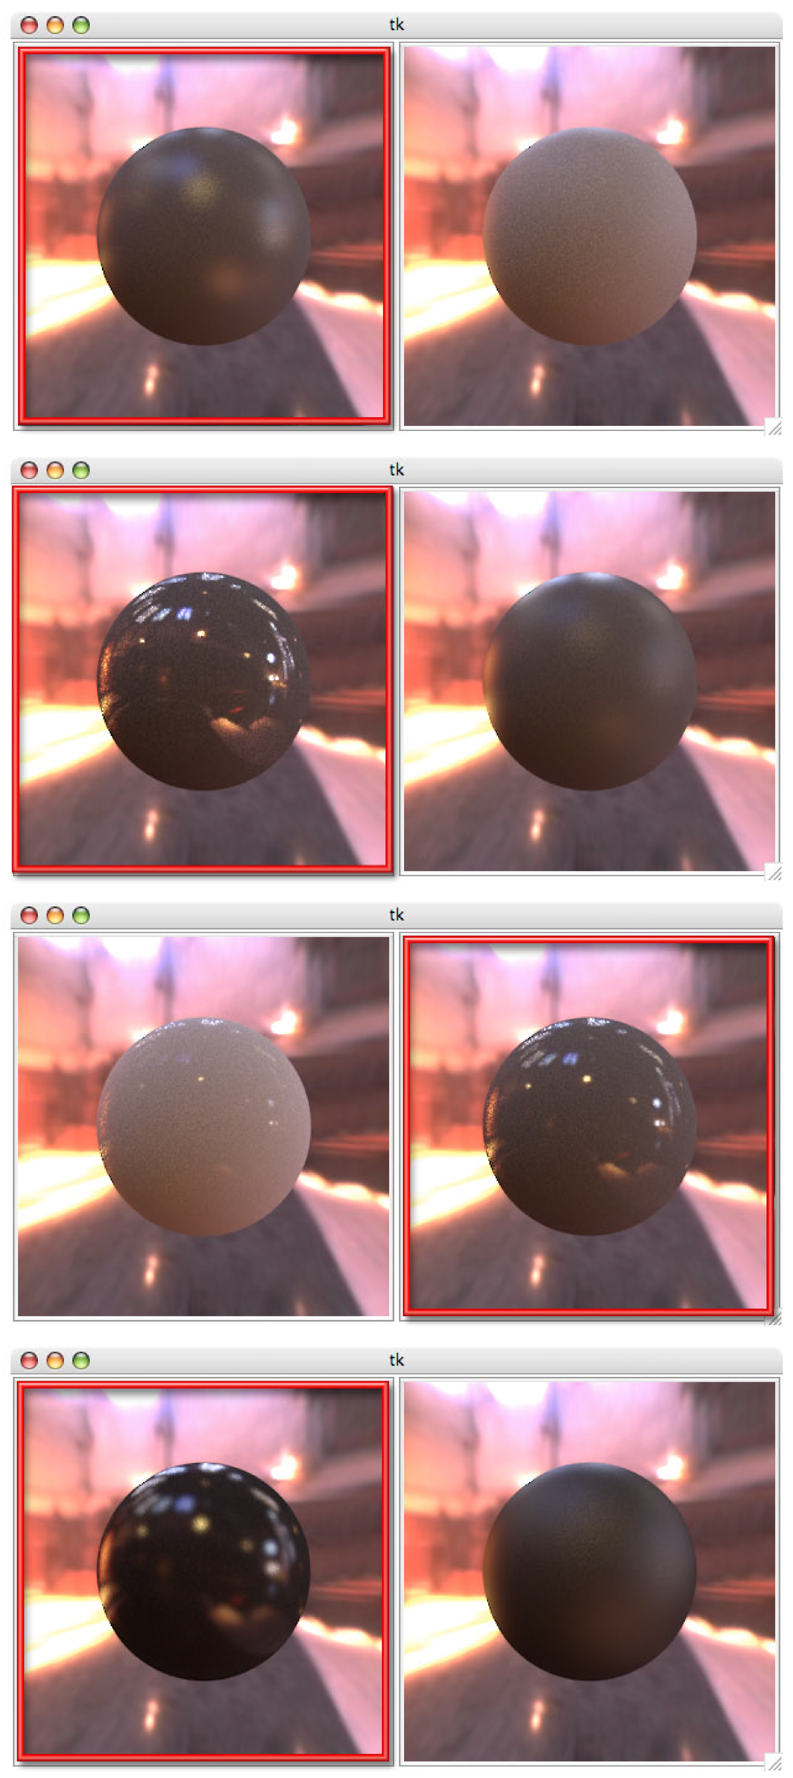

## gallery examples

Examples of the use of our material design gallery.

Subjects are assigned "target" image (top). At each iteration, they select image that they think is most like the target (shown with red box).

Average number of iterations was less than half the alternative methods compared against.

## motivation

Many problem in modelling for graphics, animation, etc. involve setting a number of dependent, unintuitive parameters.

People often know what they are looking for, but there can be a lot of "parameter twiddling" involved to get the right result. This is hard, time-consuming, and frustrating.

However, people can look at a set of examples and easily identify the ones most like the target!

So let's use Machine Learning to find examples for the user!

## background

We use ideas from Machine Learning, Active Learning and Experimental Design for a new modelling framework:

MACHINE LEARNING: Gaussian Processes are a way of fitting a model to data that not only permits predictions, but tells you the variance on those predictions.

ACTIVE LEARNING: How do we compute the utility of showing asking a human questions?

EXPERIMENTAL DESIGN: How do we select samples to optimize the response and maximize information gain?

## material design

We apply our techniques to designing materials using a gallery of images.

4D-parameter space for BRDF modelling. Render images by sampling this (continuous) space.

Use 2-image gallery design, comparing image of highest predicted value and image of highest expected improvement.

At each step, user indicates preference, model is updated and a new gallery is selected, until user finds what they are looking for.

### (1) Show user gallery, ask for preferences

We elicit preferences from the user by showing a gallery of examples.

Each example is generated from a set of parameters (modelled as a vector in a Hilbert space).

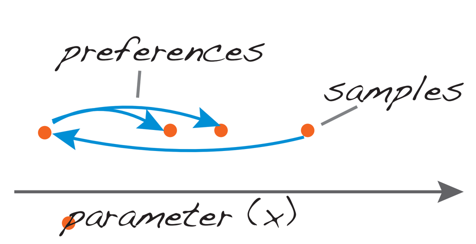


These samples and preferences become the training data.

### (2) Model valuation of user preferences

Using the preferences as input, and a Gaussian Process prior, assign values to the data that both make the observed data and the (posterior) GP likely.

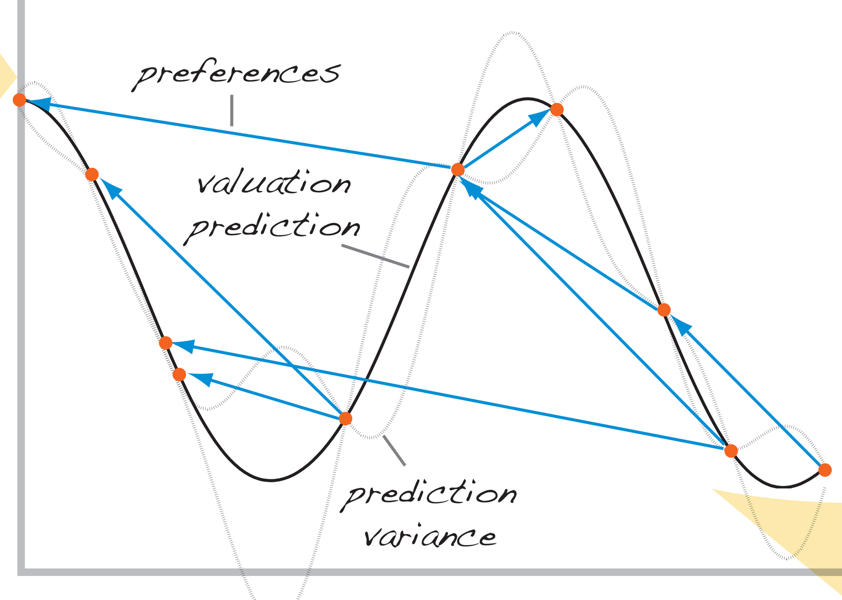


GP allows us to find prediction and prediction variance ("certainty") of any point in parameter space.

### (3) Measure exploration-exploitation

We define an Expected Improvement Function which combines prediction mean and variance to get a measure of the utility of sampling anywhere in the parameter space.

$$EI(\mathbf{x}) = (\mu_{\max} - \mu(\mathbf{x}))\Phi(d) + s(\mathbf{x})\phi(d)$$

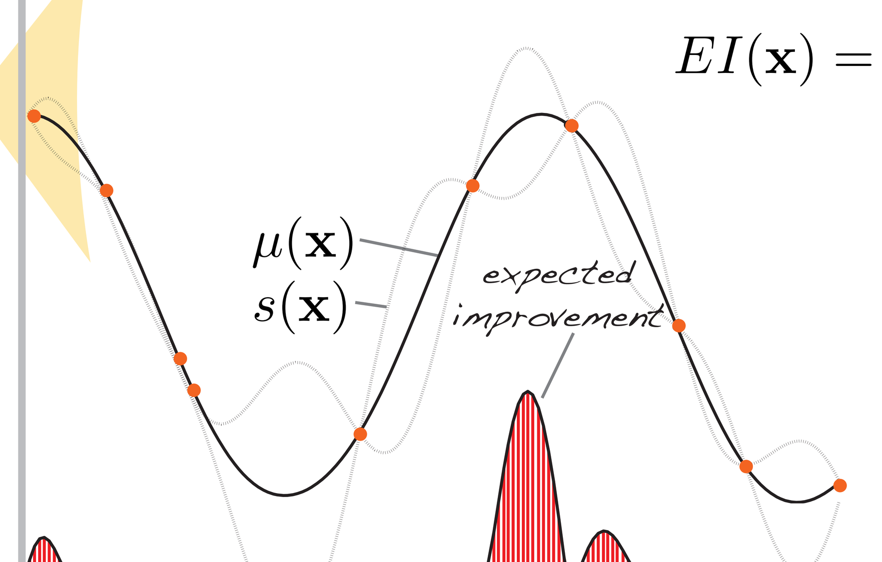


where

$$d = \frac{\mu_{\max} - \mu(\mathbf{x})}{s(\mathbf{x})}$$

$\Phi$ = normal CDF

$\phi$ = normal PDF

### (4) Generate new gallery

Now the problem of selecting the gallery becomes an experimental design.

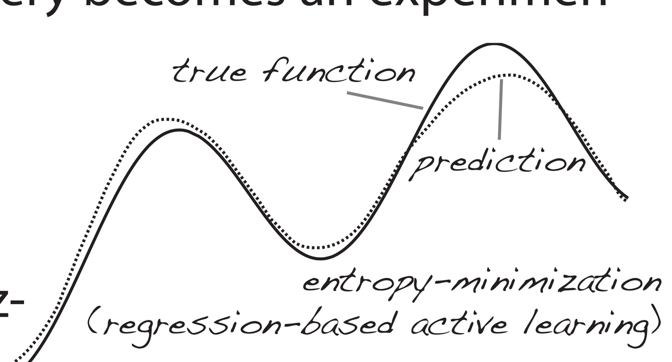


We cannot directly optimize the latent valuation, but we can (approximately) optimize the EIF to pick utility-maximizing points.

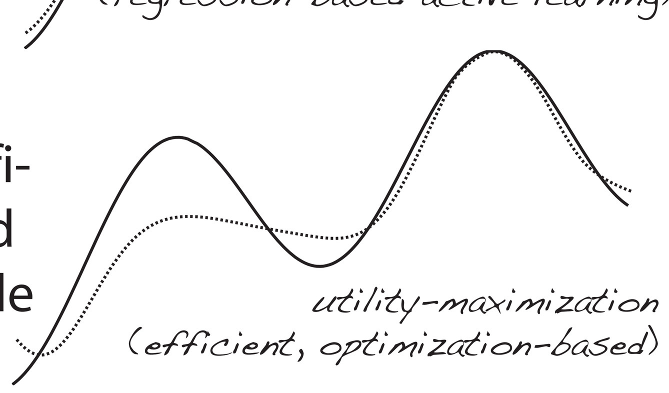


The result is that we query the user efficiently: number of queries is minimzed because we don't try to learn the whole function, just the "interesting" part.

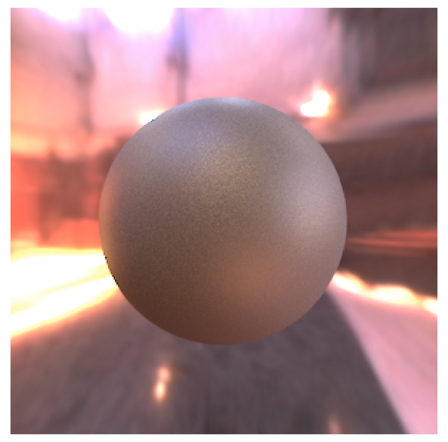

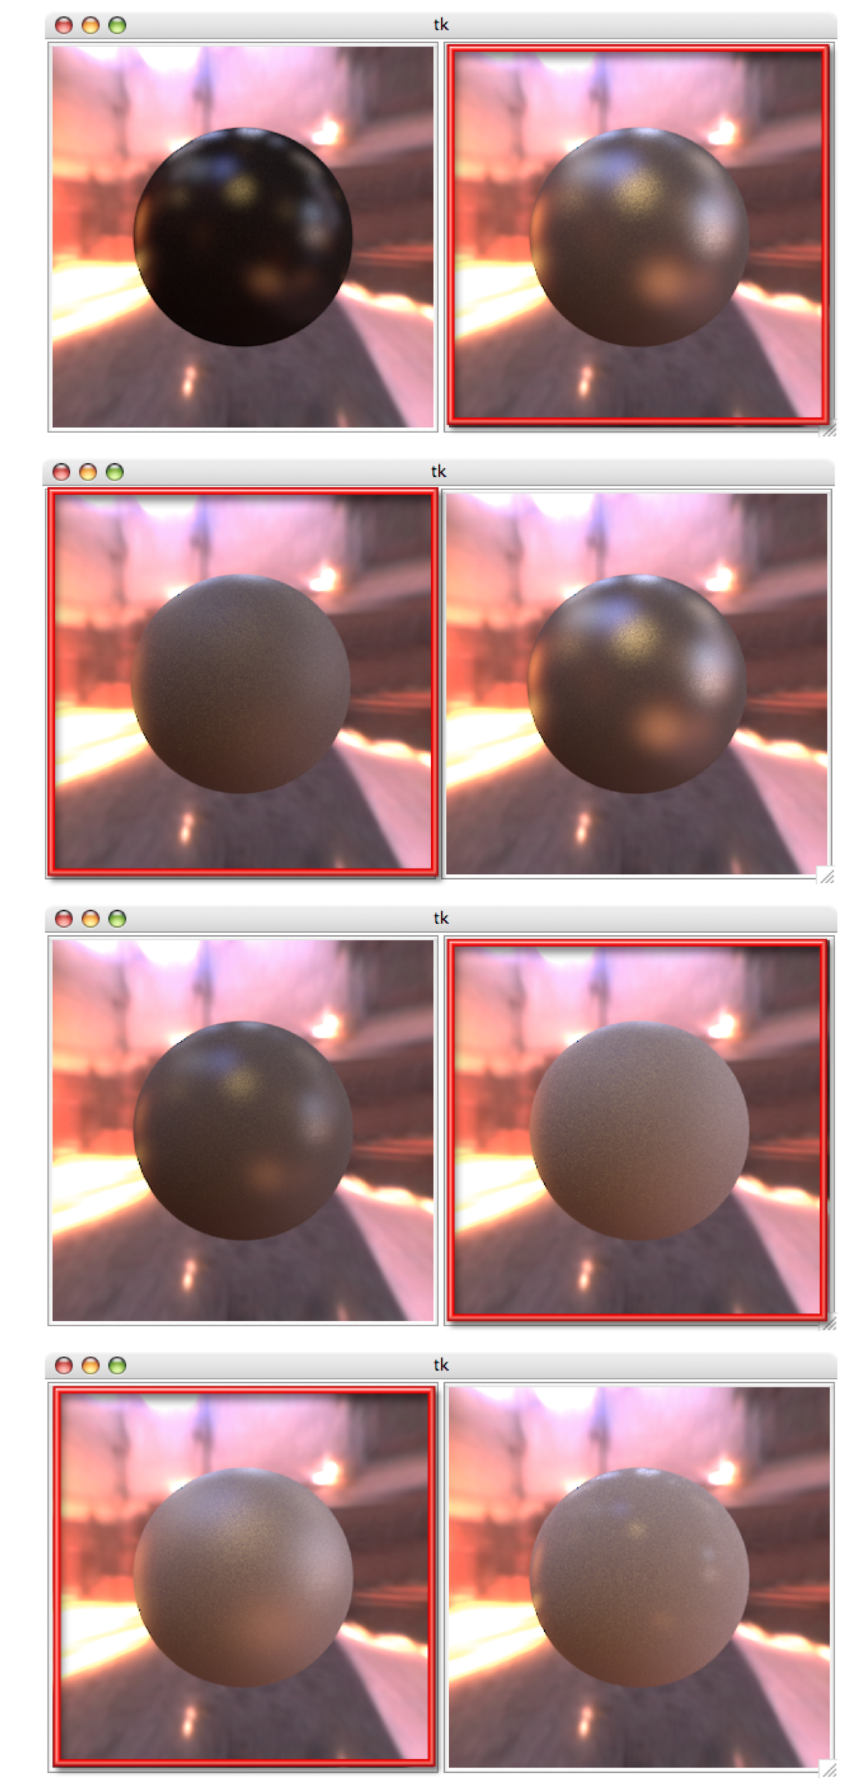

## theoretic example

The examples shown here are simplifiedto provide an intuition, and use only one dimension (parameter).

In reality, the model is designed for multi-dimensional parameter spaces (4D for the material gallery)

SIGGRAPH2007