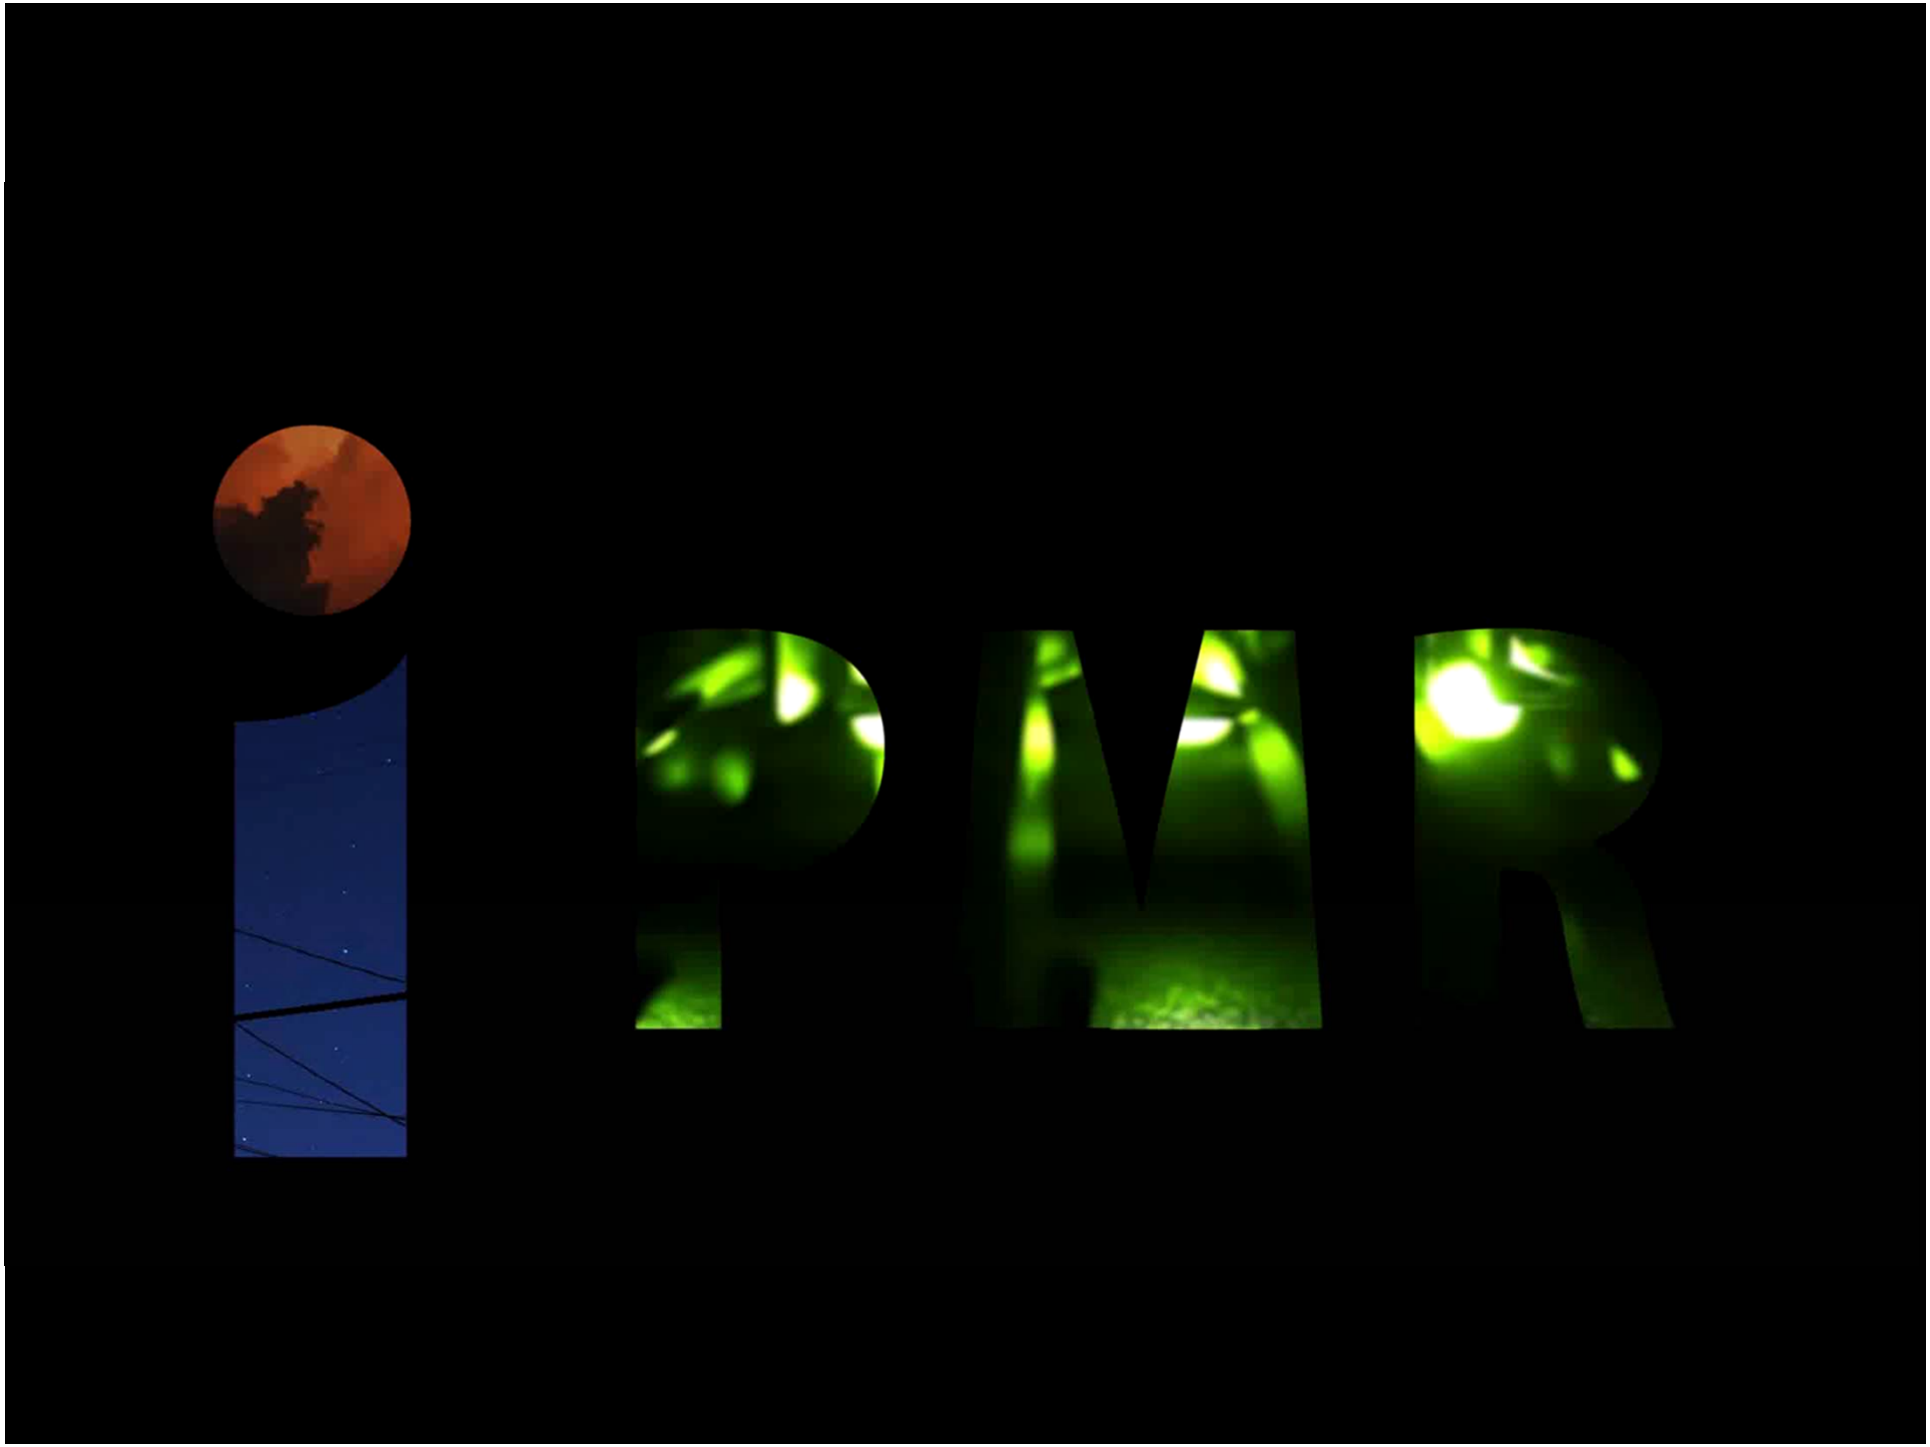
iPMR

Welcome

What is the point of this course?

Compare with MLPR

Image courtesy of daveseven, used under creative commons. http://www.flickr.com/photos/daveseven/512897762/

What is the object label for this image?

Car.

1976 Panther J72.

Compare with MLPR

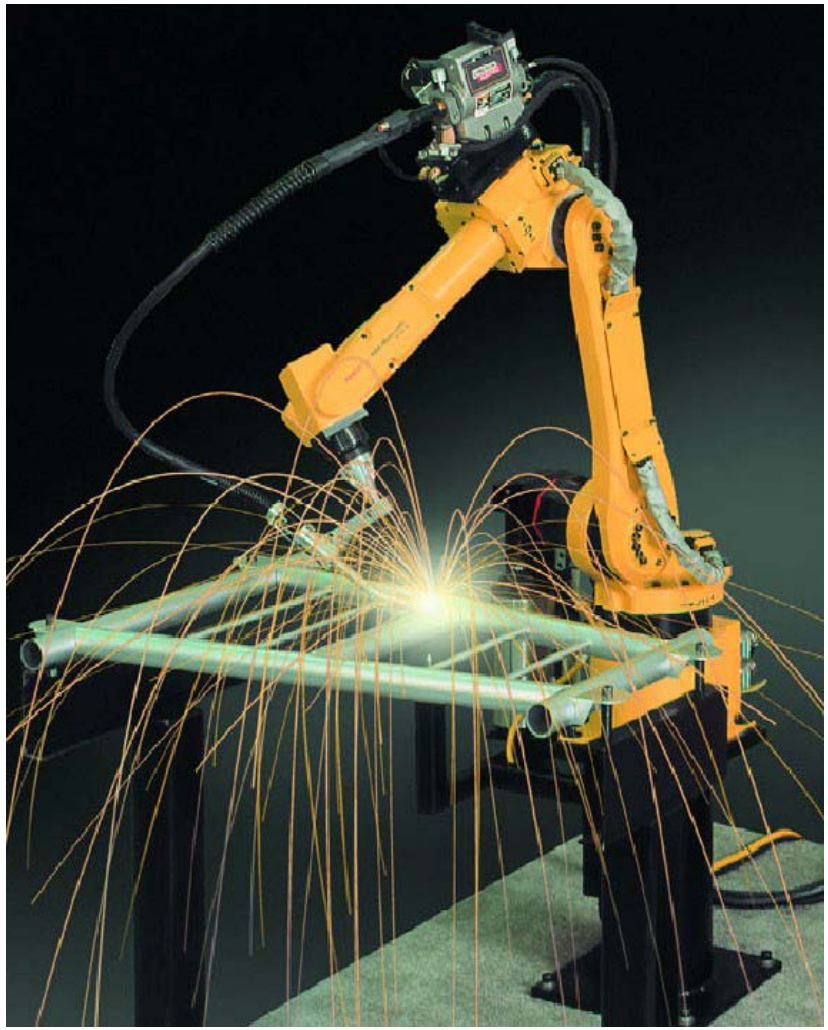

Image © Orange Indus. Creative commons licence

What position will a robot arm be in if we apply a specific set of torques?

Compare with MLPR

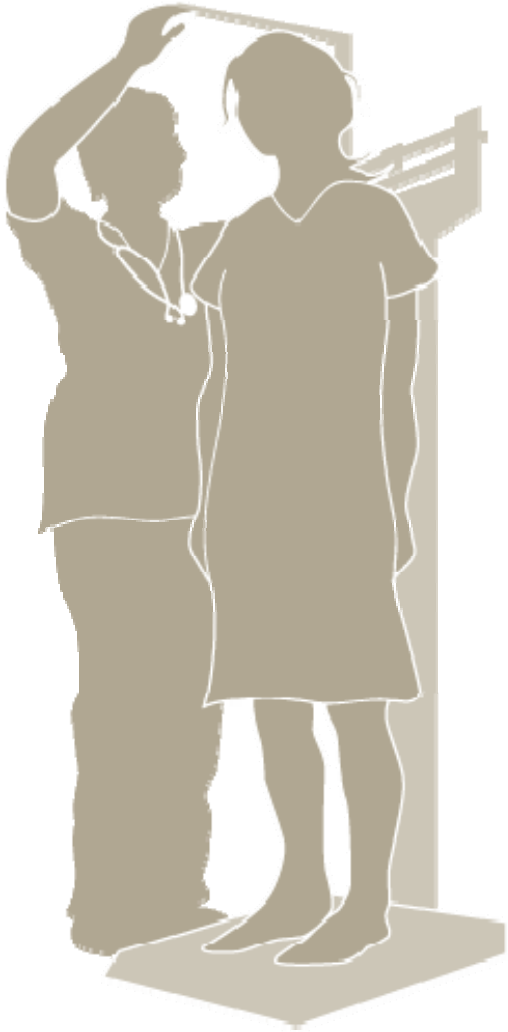

Might even involve predicting a number of things at once.

The Difference

- Predict a value
- Or provide error bars
  - i.e. simple distribution.
- Predict each part individually.
- Always condition on the same things (supervised learning task)

- May want to say how things covary with one another.
- May want to be precise about distributions
  - Need to evaluate risks and probabilities.
  - Need to make good decisions based on those risks.
- May want to condition on different things (unsupervised learning task)

What is the point of this course?

A Comment on Randomness

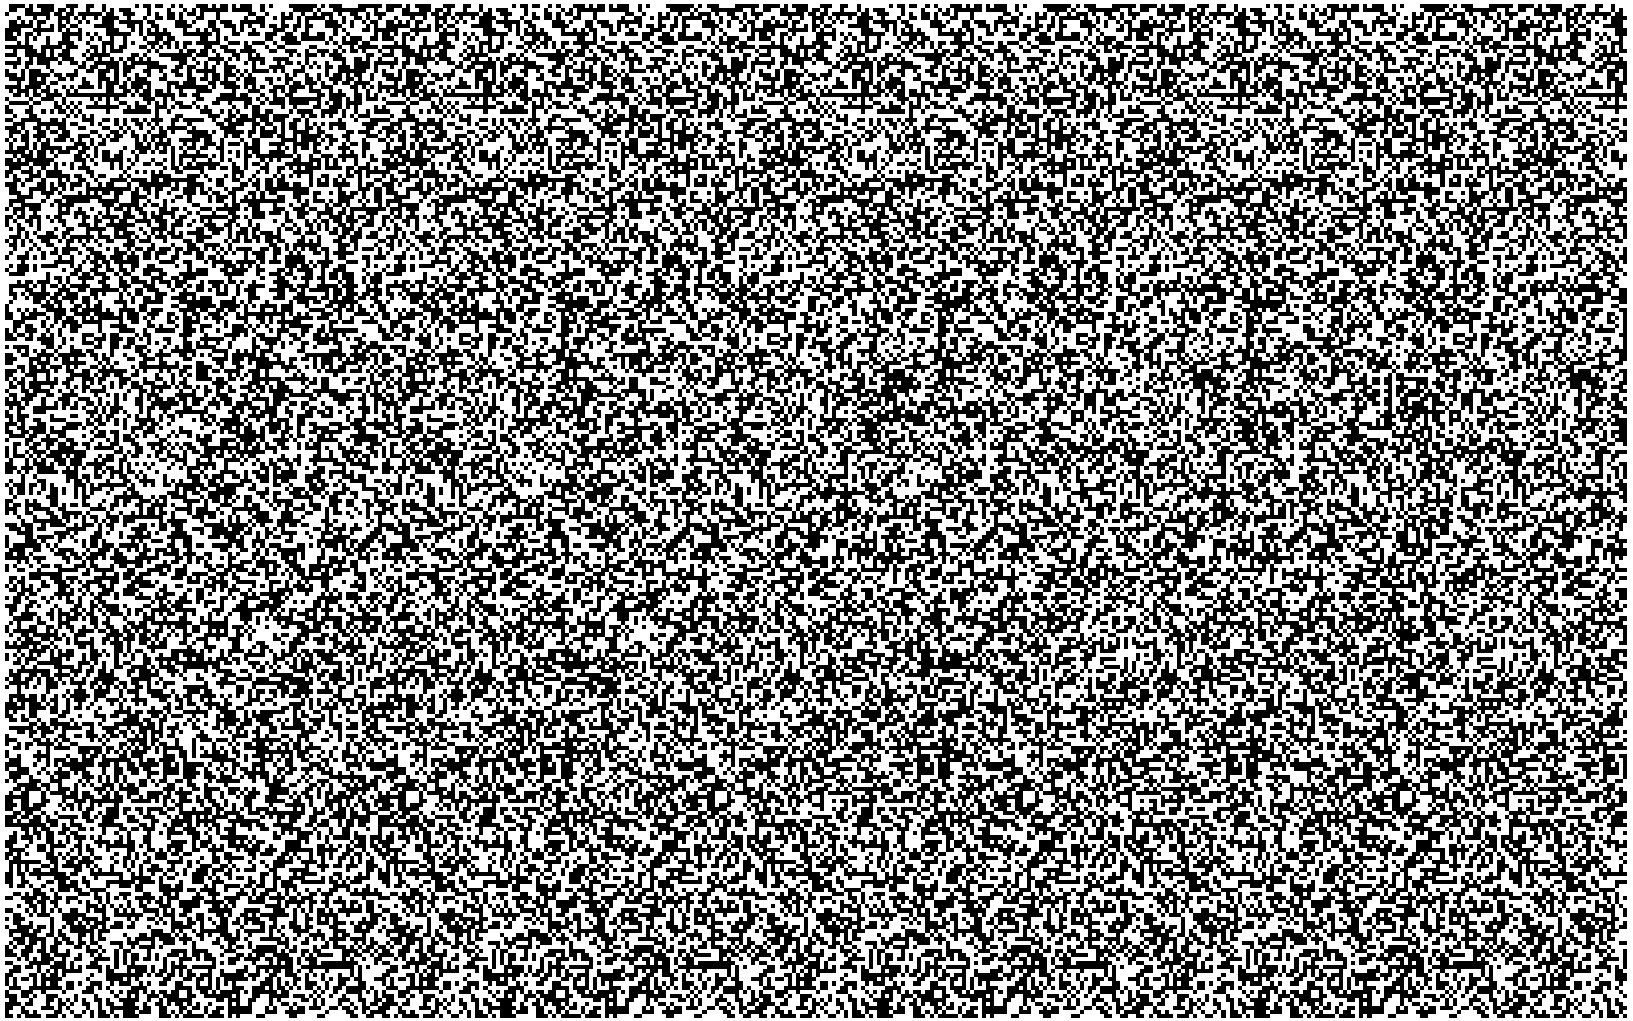

Decorating Example

- Decorating can be done like this

Image Creative Commons tumblr: wkn-source-code

Images Creative Commons tumblr: brindledog

Decorating Example

- Or like this

Image Creative Commons tumblr: fryeggs

Images Creative Commons tumblr: fryeggs

Decorating Example

- Or like this

Image Creative Commons tumblr: patternforplunder

Decorating Example

But it is not a great idea.

Hints on Decorating

- When choosing a colour and texture for a sofa…

$$P(\text{sofa colour}|\text{got taste})$$

$$P(\text{sofa colour} = \text{putrid orange}|\text{got taste}) \approx 0$$

Explain: , Marginal Distributions, Probability, Conditioning, Line=Conditioned On, Interpret.

Hints on Decorating

- But not enough…

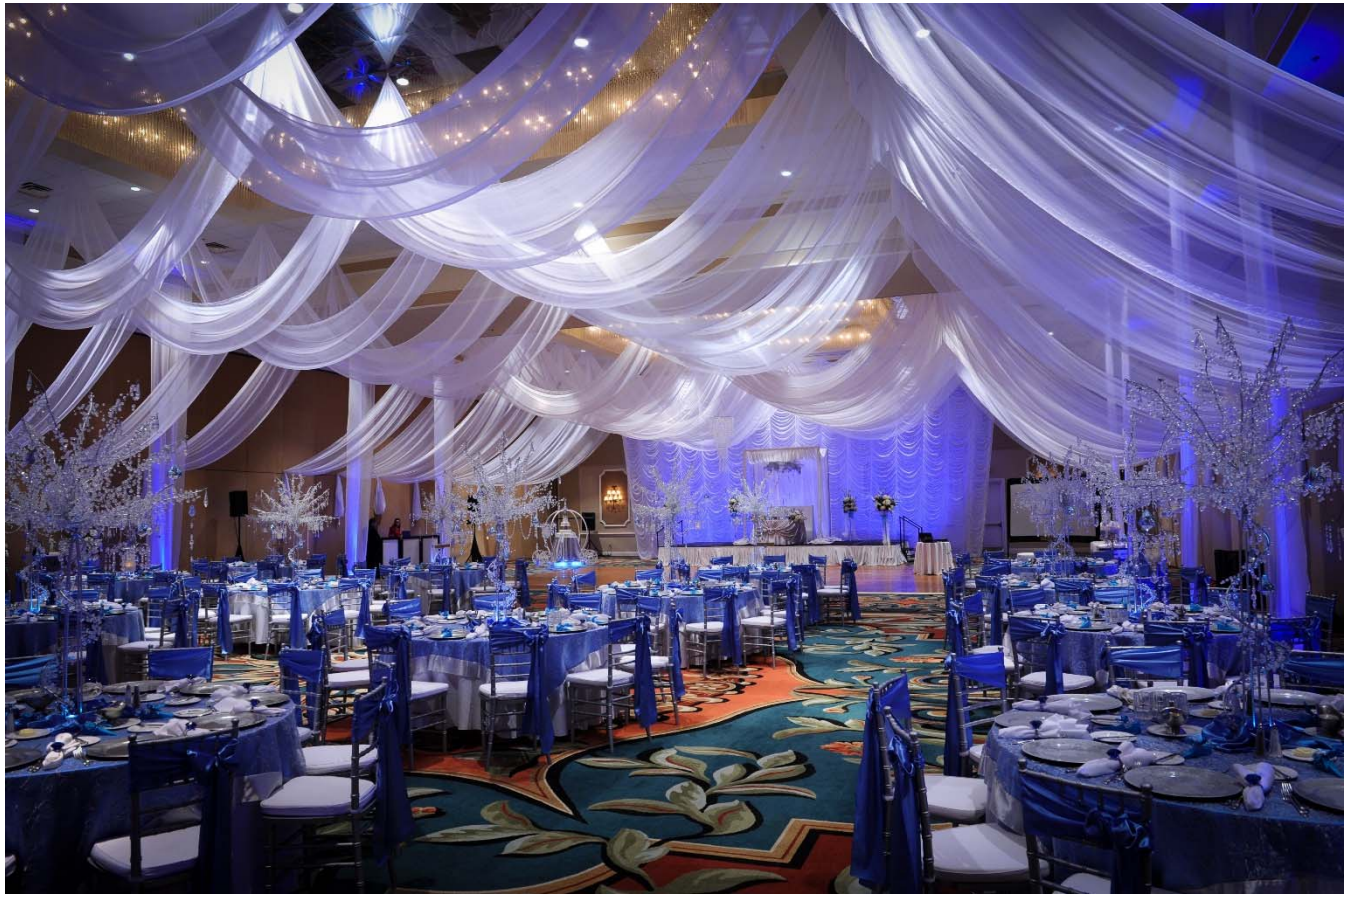

Hints on Decorating

- This is also a probability distribution

$$P(\text{lighting colour}|\text{carpet colour}, \text{carpet pattern}, \text{got taste})$$

Explain: Conditional distribution, Joint distribution. Mention rules of probability.

What is the point of this course?

This course is about giving the distribution of colours for everything in a tasteful room…

$$P(\text{carpet colour}, \text{sofa colour}, \text{lighting colour}, \text{wall colour}, \text{painting colour}, \ldots | \text{got taste})$$

Luckily it is not just about rooms…

Explain: Choosing colours jointly. Generalise. Lots of tasteful rooms: sampling.

Break

Where? How?

- Where do we get these distributions from? And how do we use them?
- How do we choose the structure and type of distributions?

= Modelling

- Where do we get these distributions from?

= Learning

- How do we use them?

= Inference

Inference

- Suppose we know (precisely)

$$P(\text{carpet colour}, \text{sofa colour}, \text{lighting colour}, \\ \text{wall colour}, \text{painting colour}, \ldots | \text{got taste})$$

what might we ask?

Real life questions:

What is the employability of a person given their details?

What is the distribution of future stock prices given the past prices?

Learning and Modelling

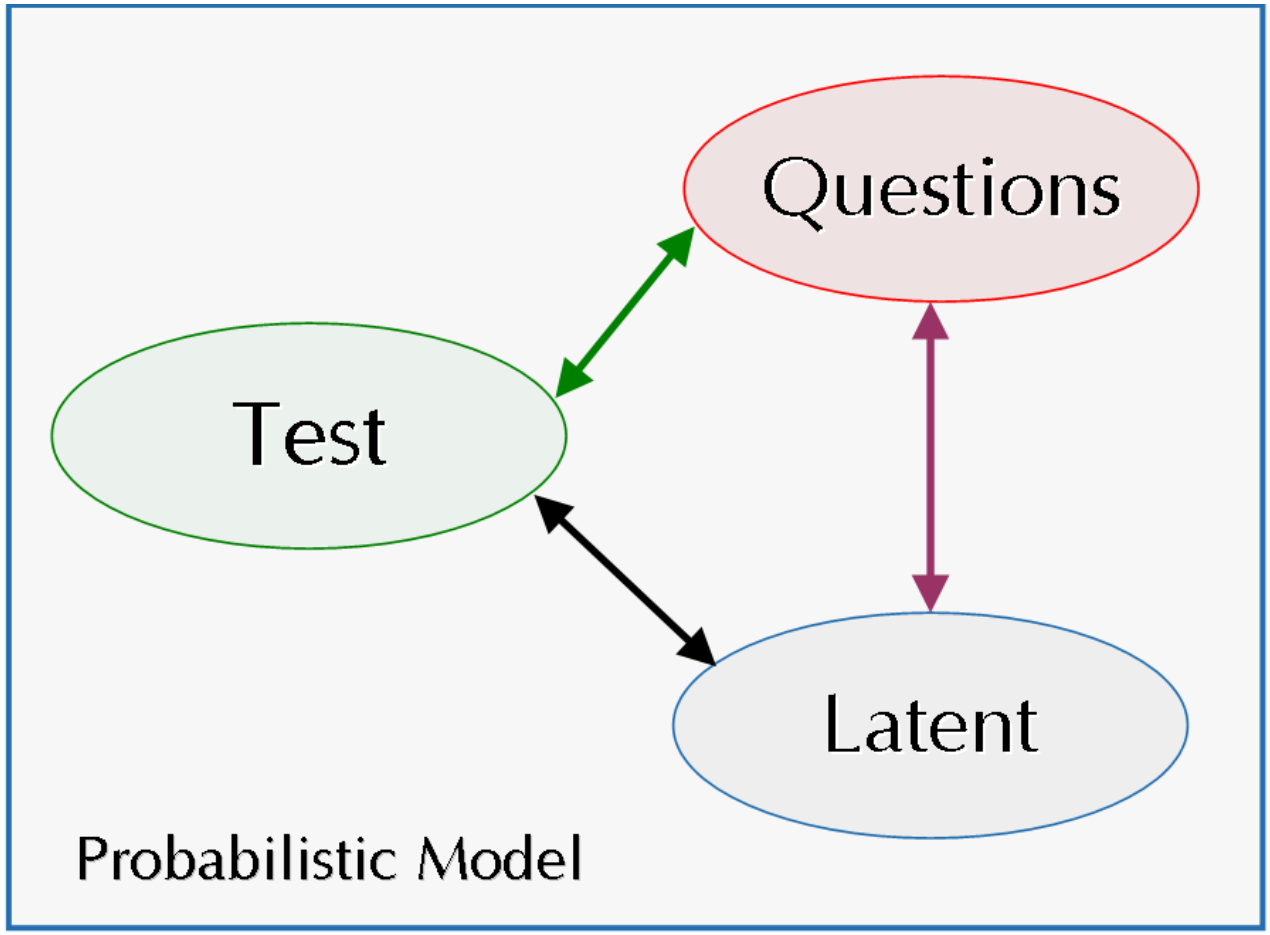

Learning

- Prior beliefs →Model and Prior distribution.
- Suppose we have

$$P(\text{carpet colour}, \text{sofa colour}, \text{lighting colour}, \\ \text{wall colour}, \text{painting colour}, \ldots | \text{got taste}, \theta)$$

but not some parameters θ.

- What should we do?

Example

H,T,H,T,H,H,H

- What a priori assumptions need we make?
  - Each coin toss is independent of other coin tosses.
  - Each toss is with the same coin.
  - Coins can be biased.
    - Most coins are not biased.
    - Those that are, are only slightly biased.
    - Occasional double headed, or double tailed coins.
  - Encode this as $P(\theta)$ where $\theta = P(\text{head})$.
- Once this is done we can uses data to learn θ.

An aside: Priors in the Brain

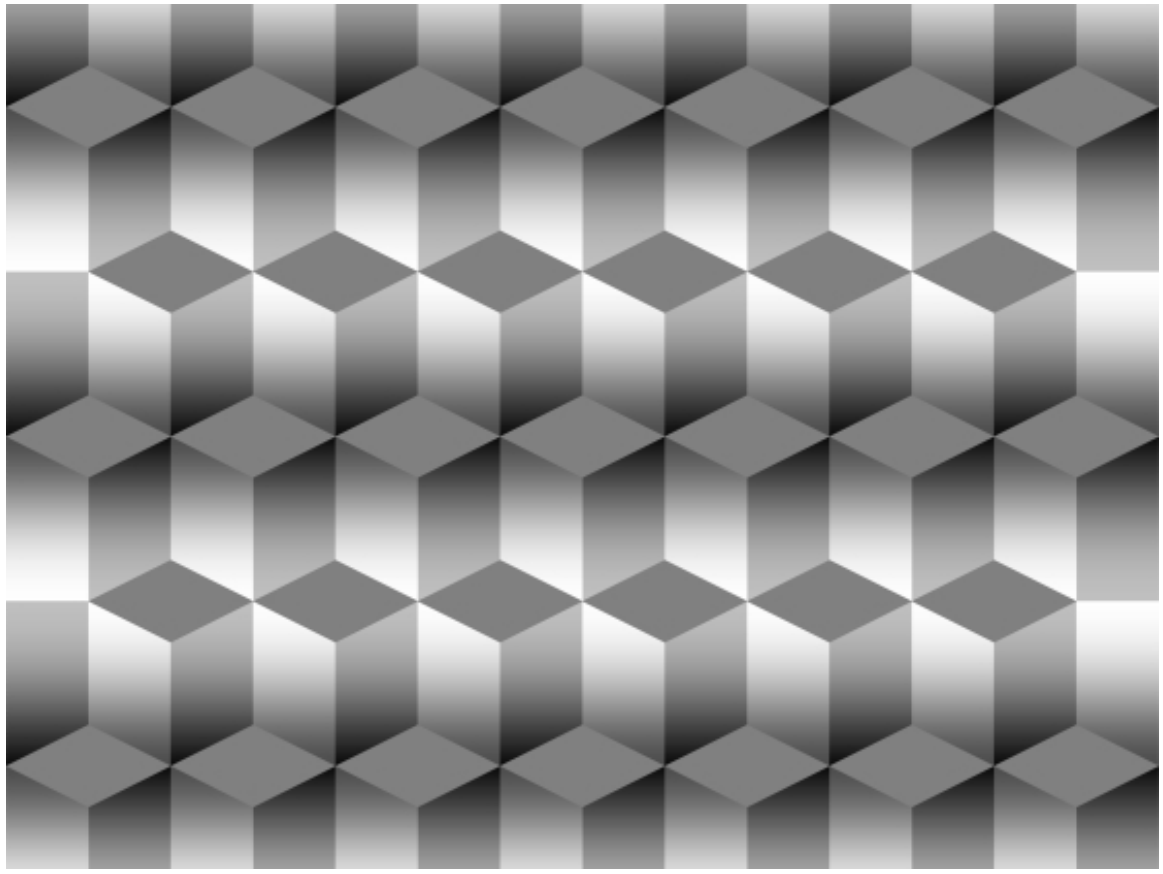

Logvinenko Illusion


GARDNER

How?

- Learning is actually the same as inference but at a different level.
- Decorating example.

- Summary. The whole course relies on the basic rules of probability, and some basic information theory.

Typical Problems

- Relating multiple variables: diagnostics, credit risk, genetics, language, vision,
- Time series modelling: epidemiology, finance,
- Latent discovery: clustering, phylogeny, parsing.
- Making decisions: planning, robotics, multi-agent systems, games.

Take home

- We need to move beyond simple models that predict one variable given others.
  - We need to handle many variables together.
  - We need to handle different distributions.
  - We need to use appropriate distributions.
- We can use these models to answer more varied questions via inference.
- We can discover appropriate models from data via learning.
- Inference and Learning are applications of the rules of probability.

Summary

- PMR is about modelling joint distributions and unsupervised learning
- Now you have a feel for what that means
- Things that will be important
  - Understanding the rules of probability.
  - Understanding some key probability distributions and densities.
- See website.