Calcul bayesien dans le systeme sensorimoteur

Paris, 2005

Sensorimotor integration

Decisions: Probabilities and Utilities

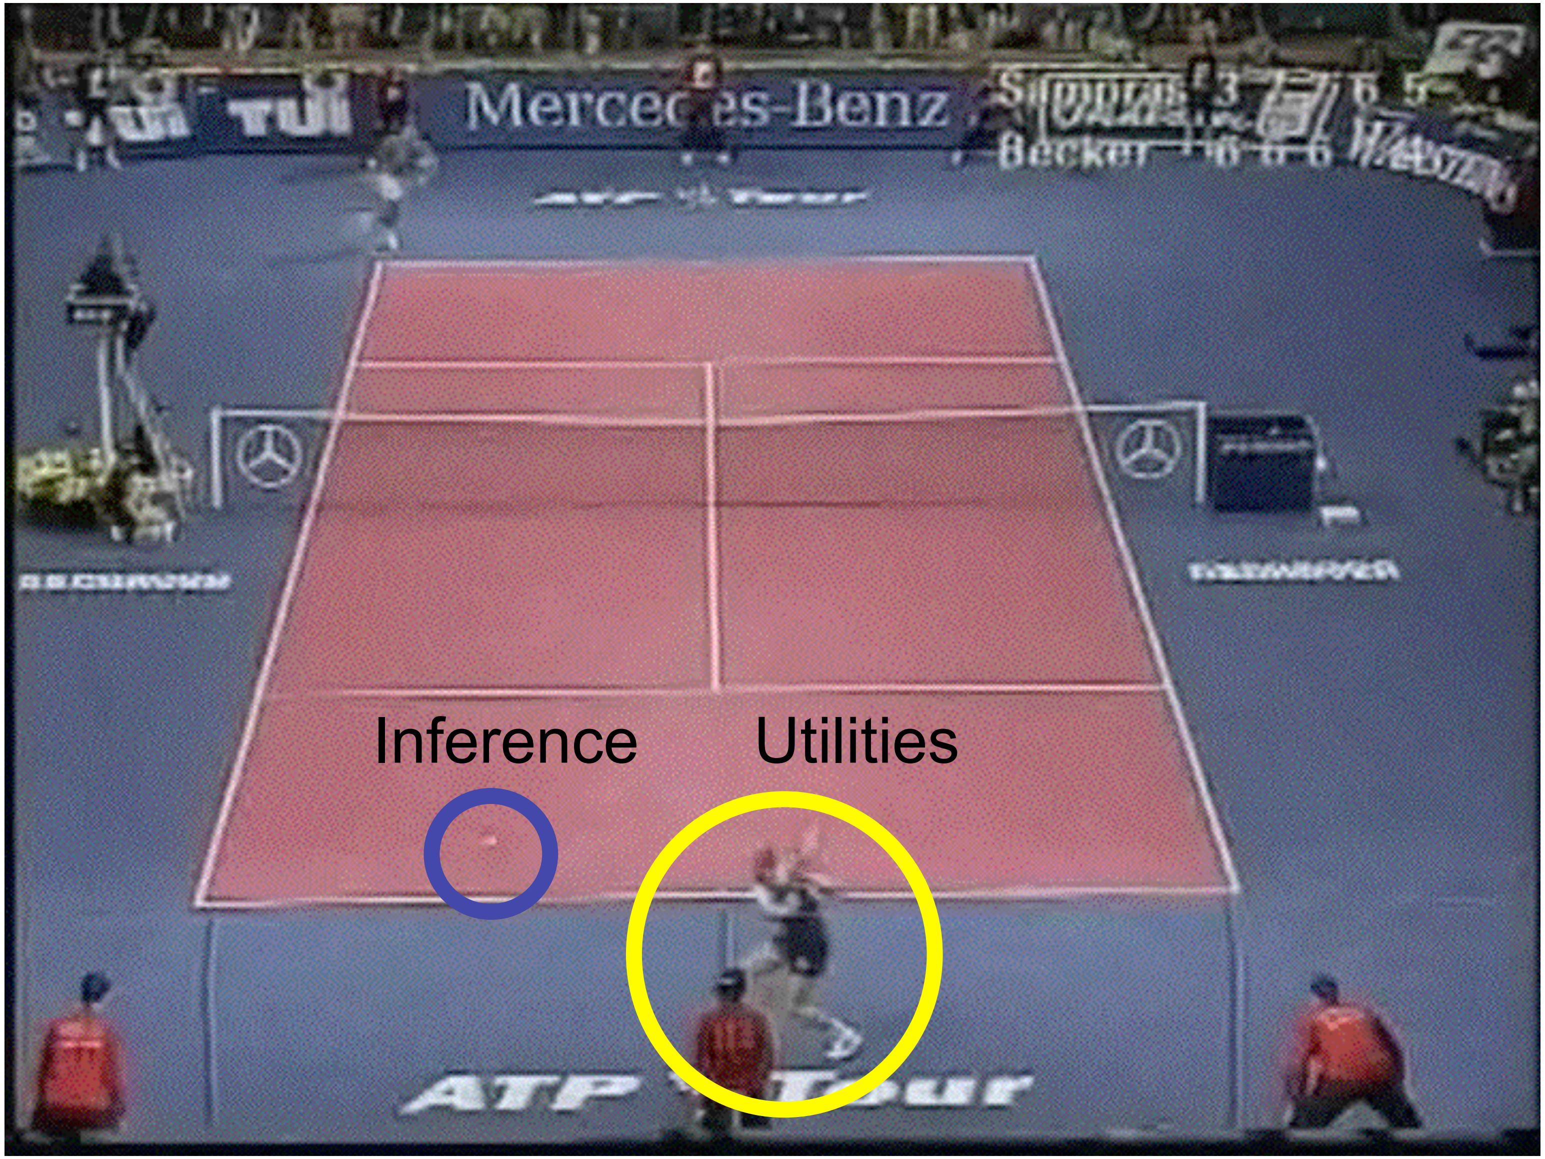

Decisions are important

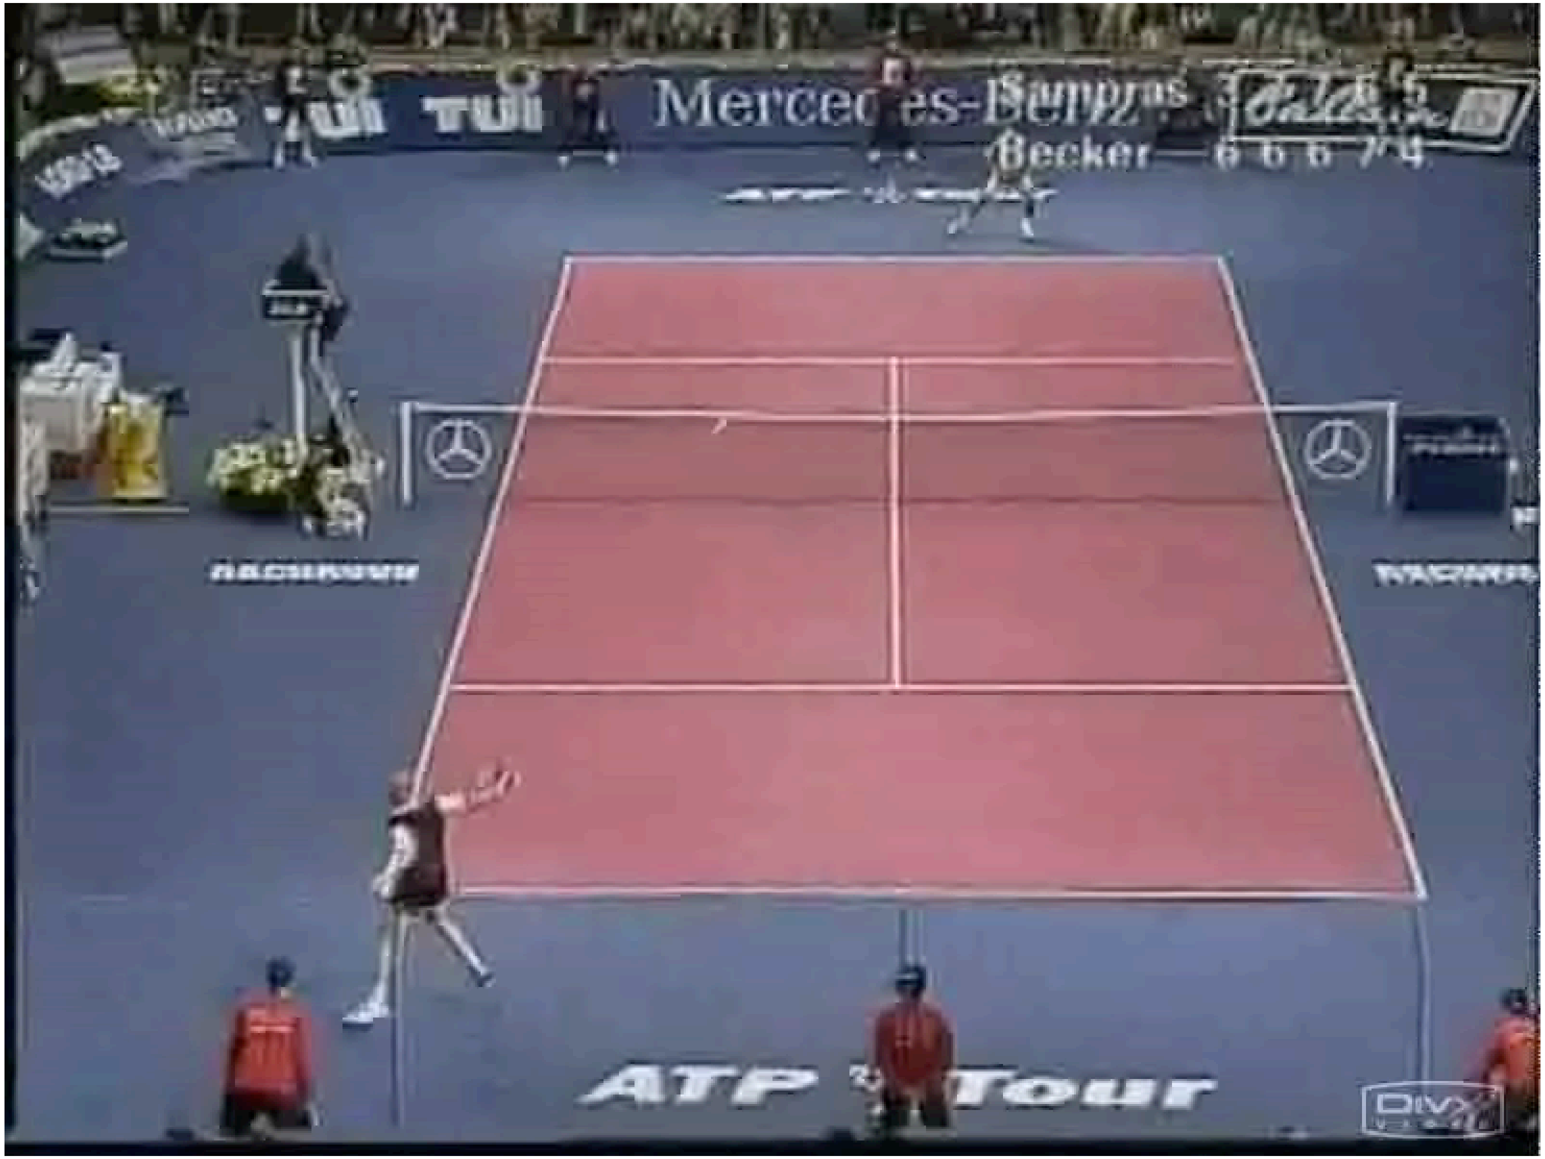

Where will this ball land?

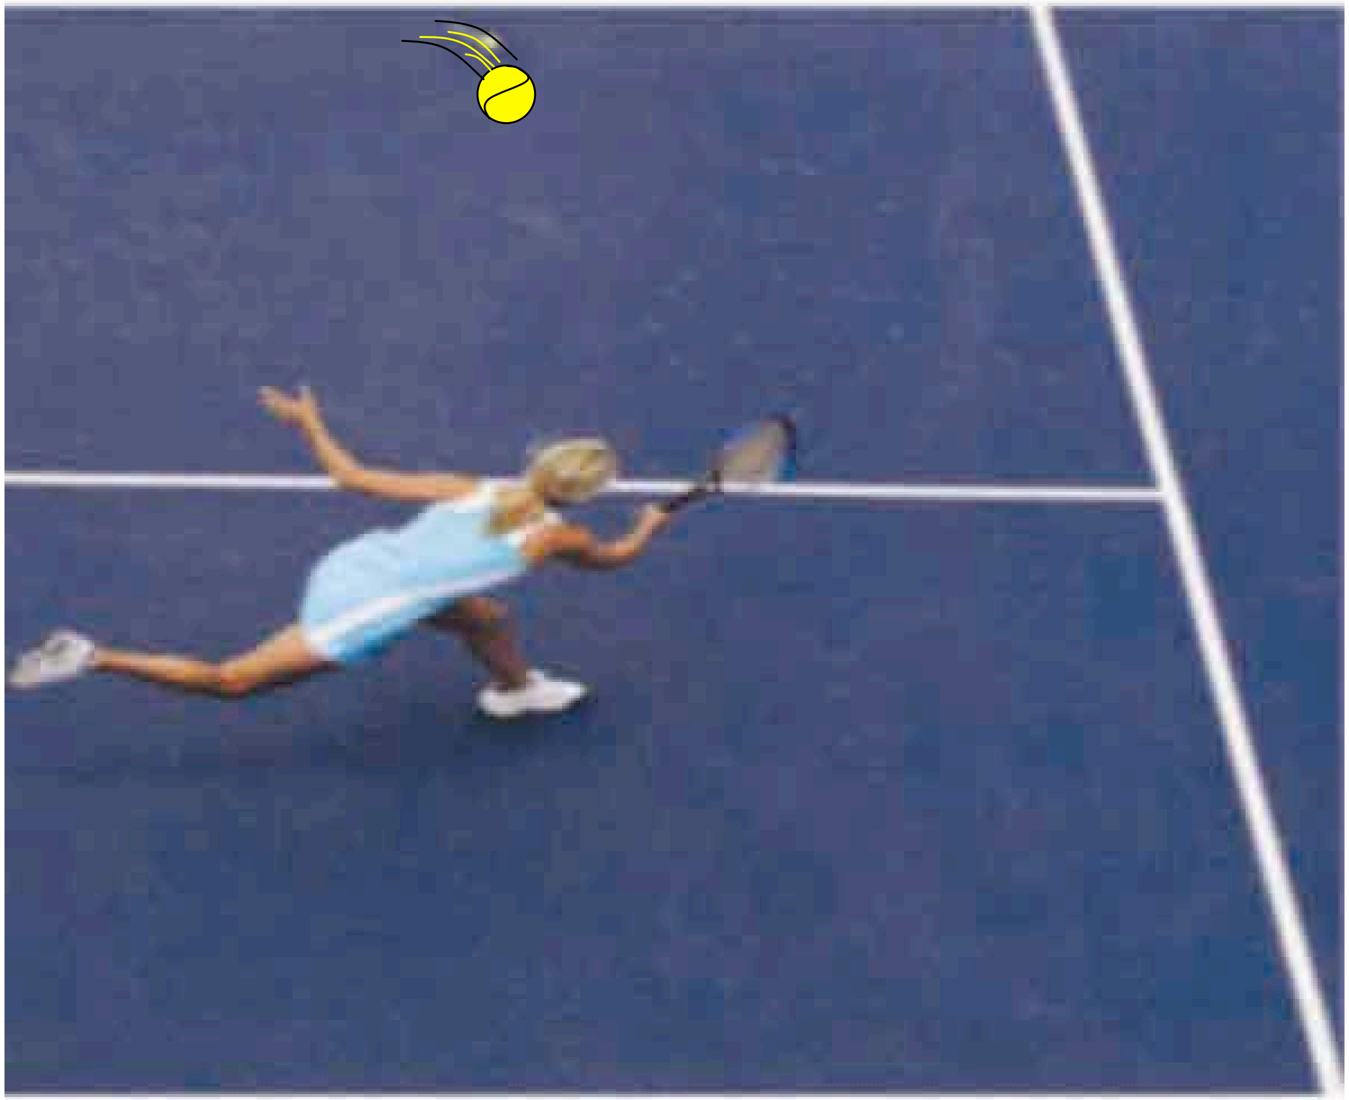

Looking at the ball: Likelihood

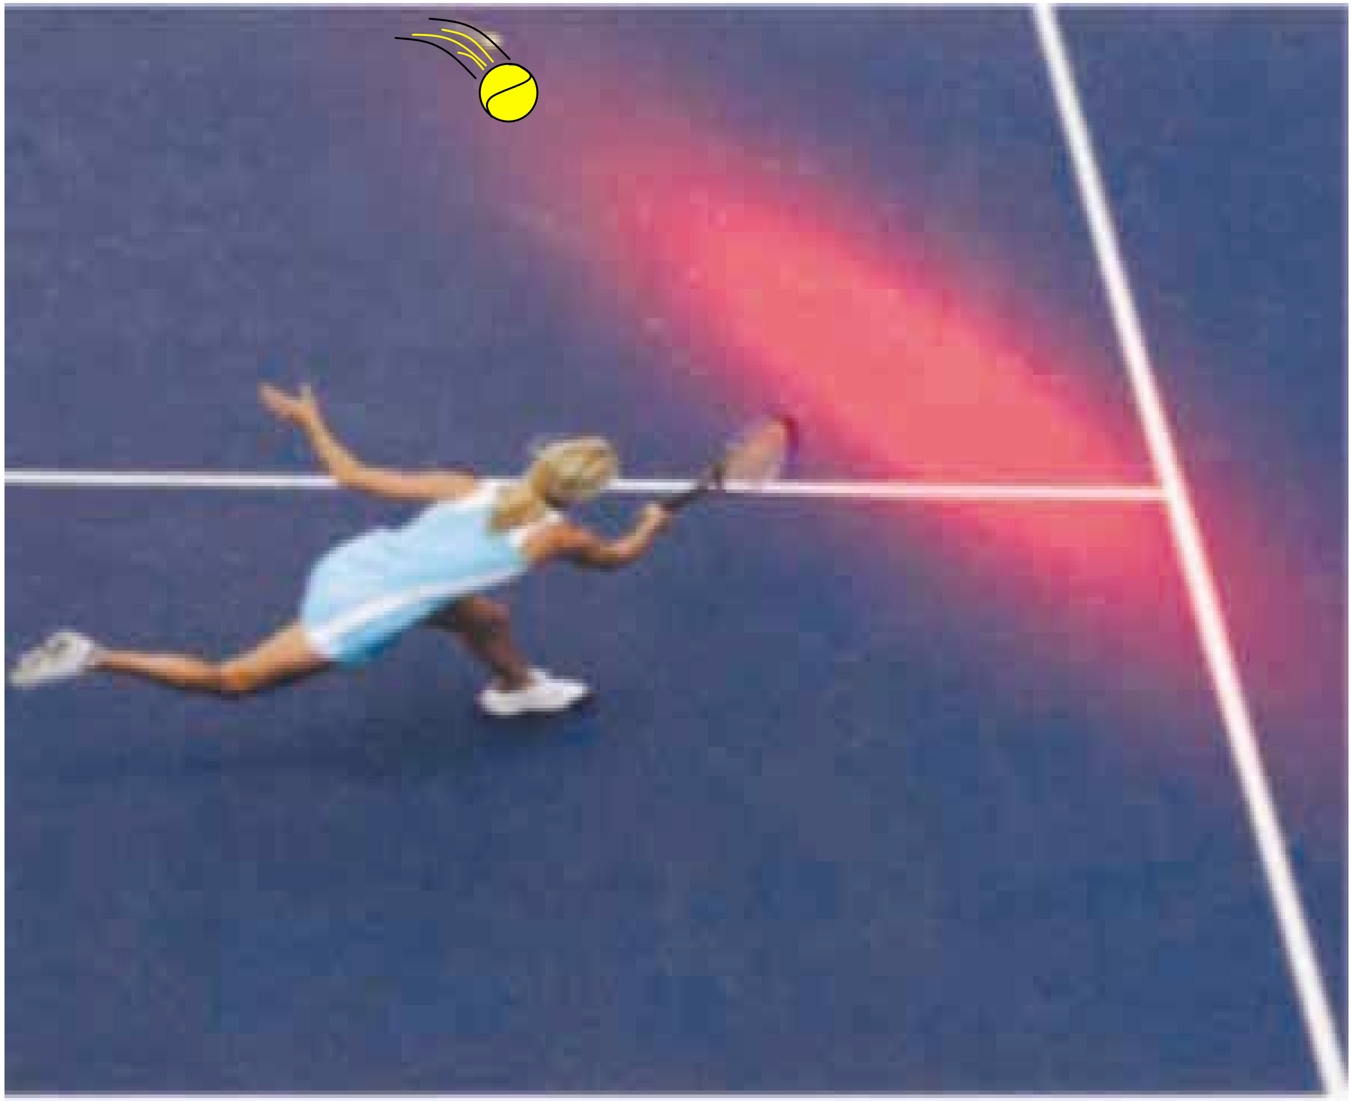

Prior knowledge

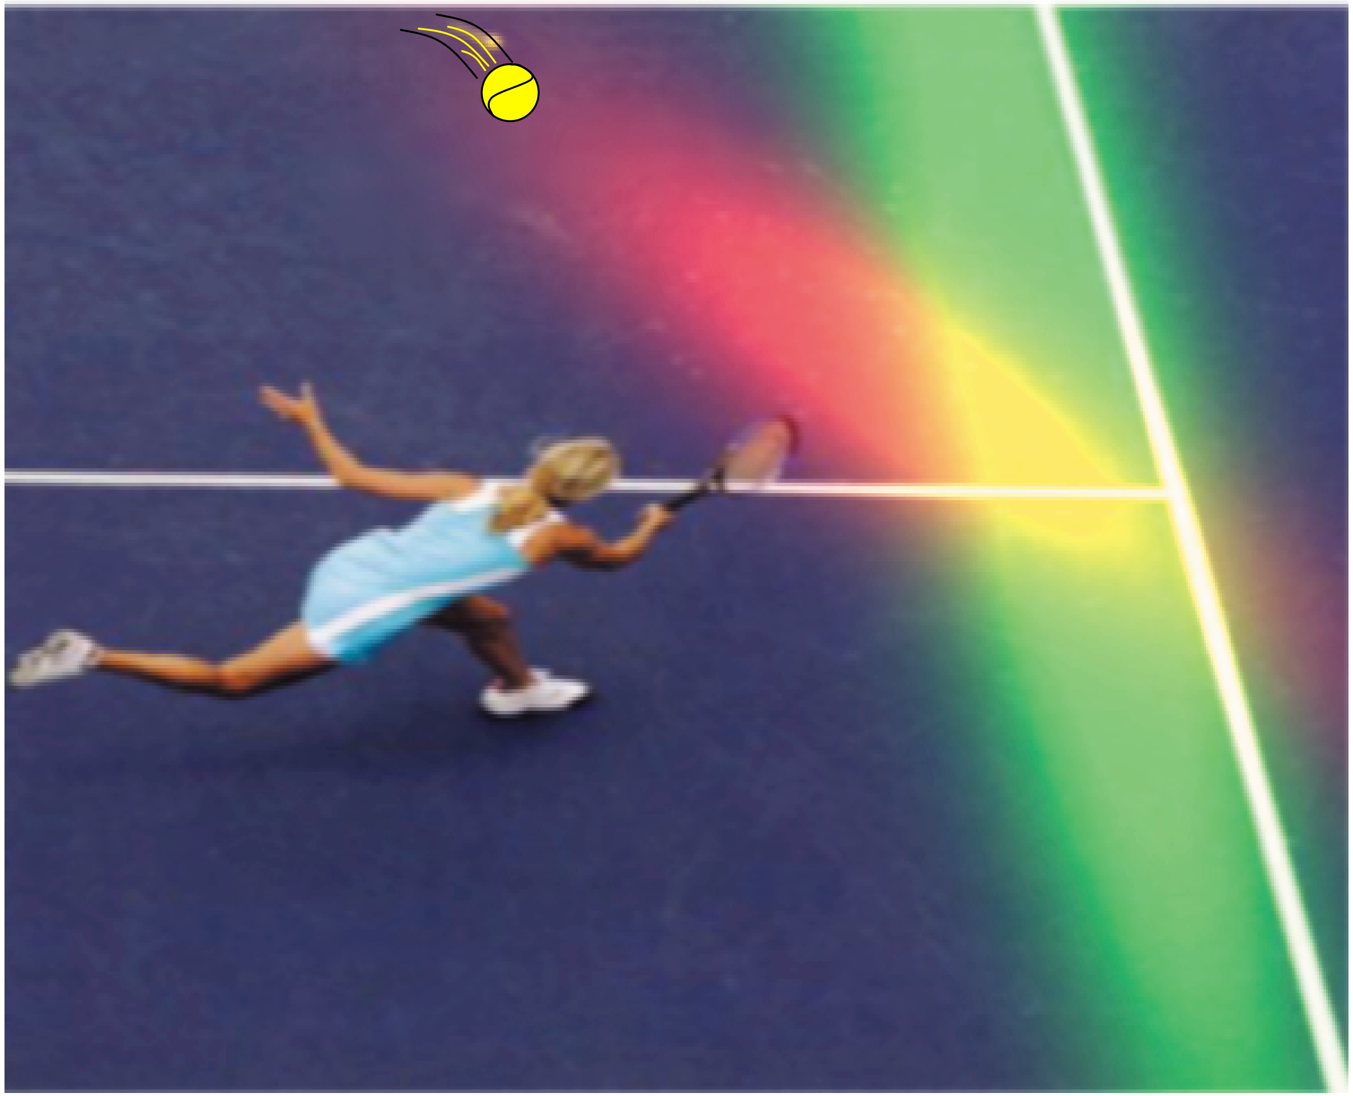

Combined

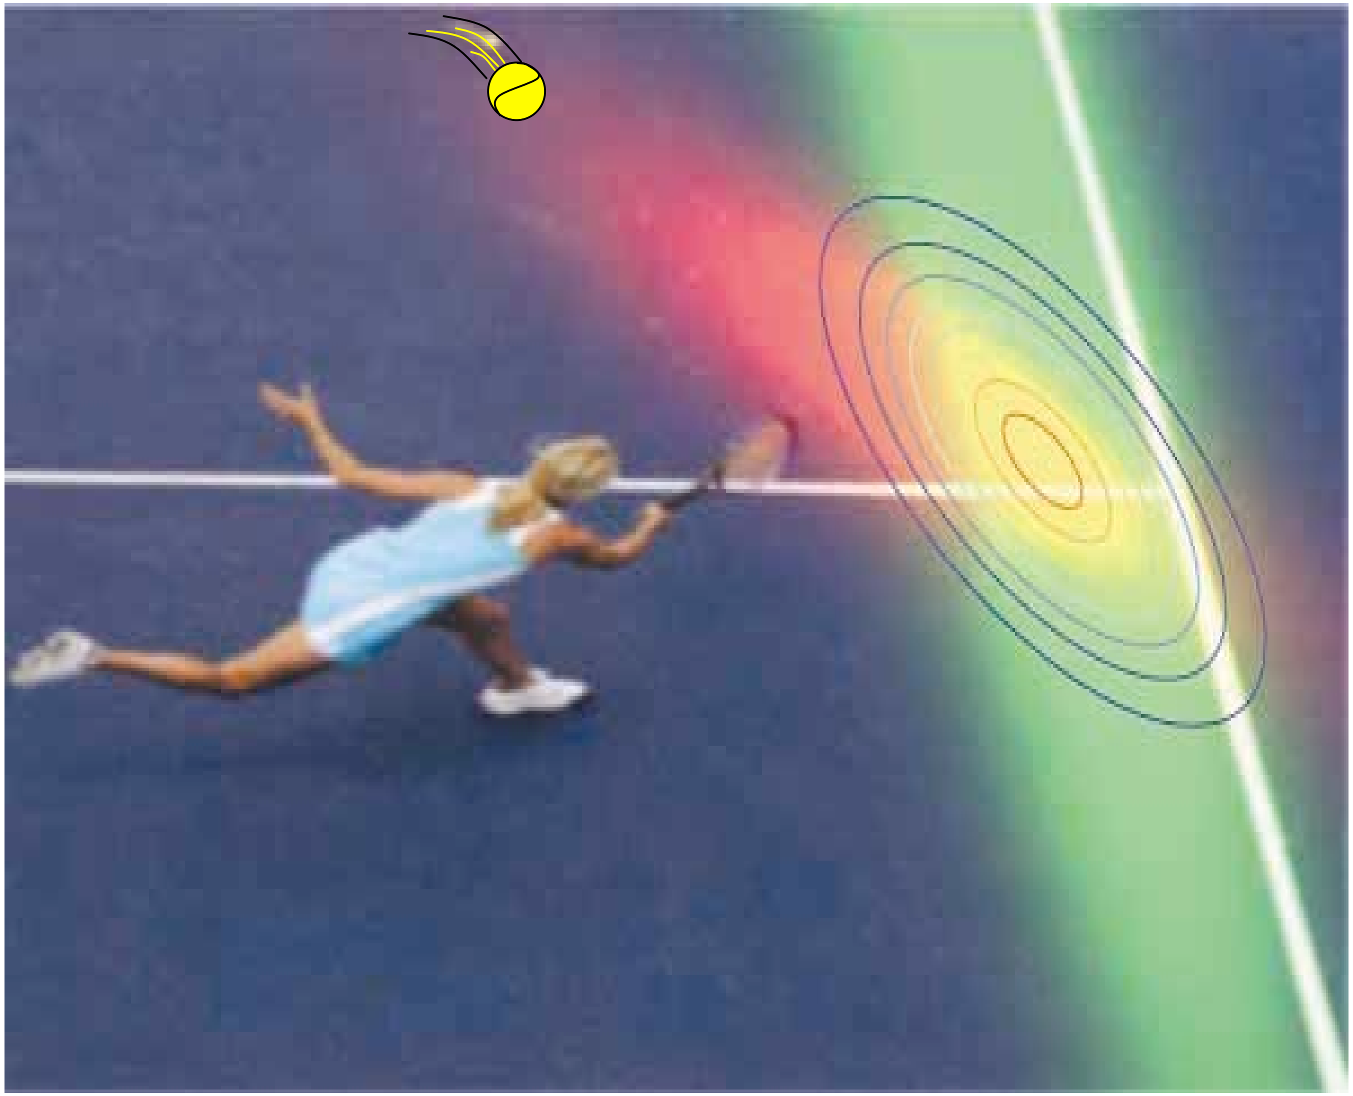

Bayesian statistics

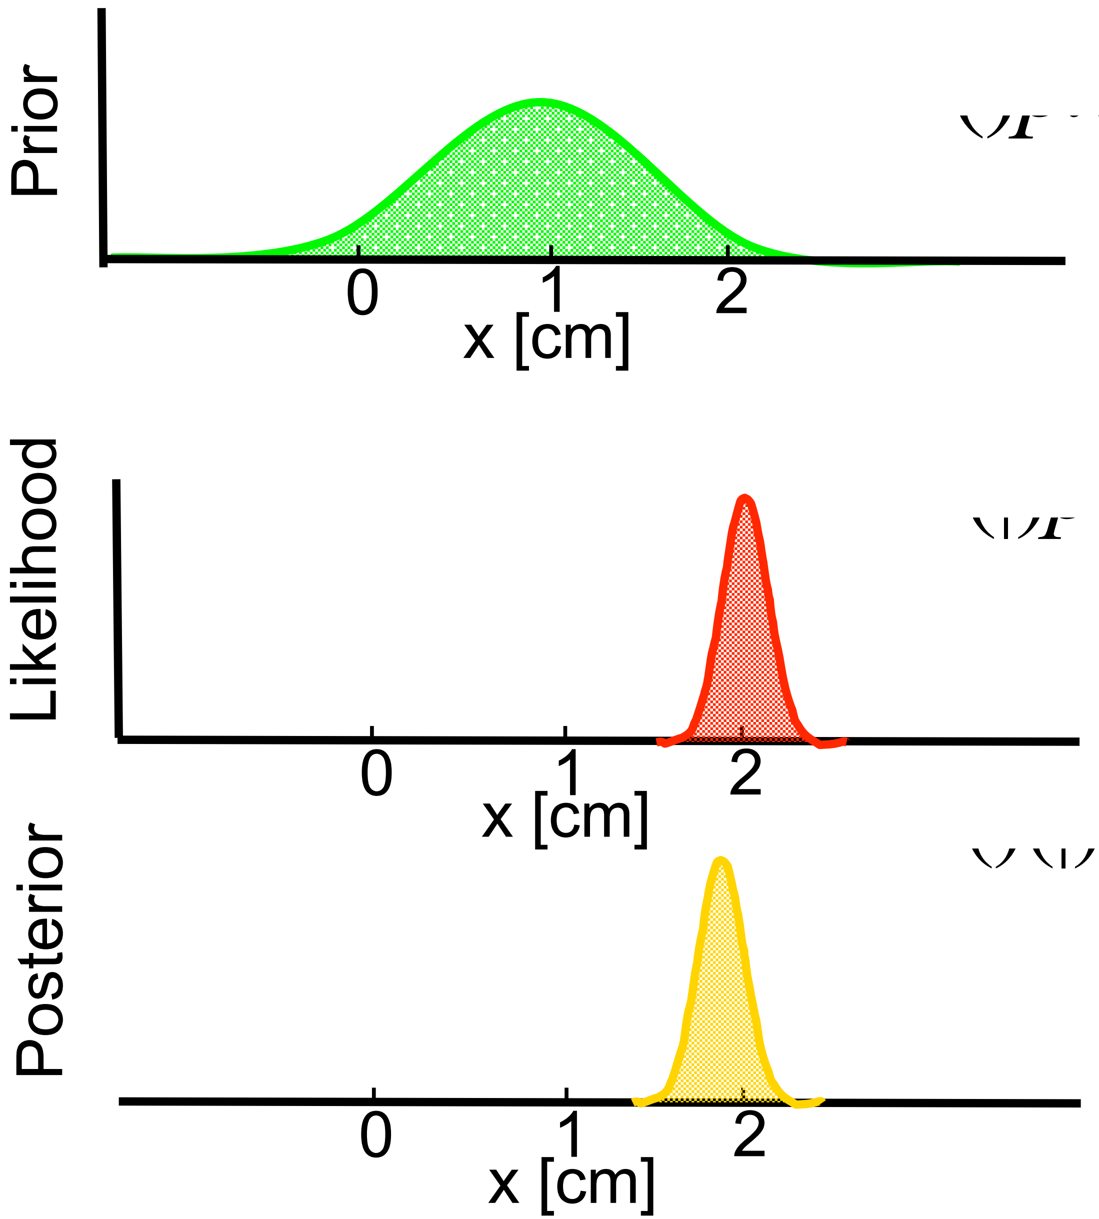

Bayesian statistics

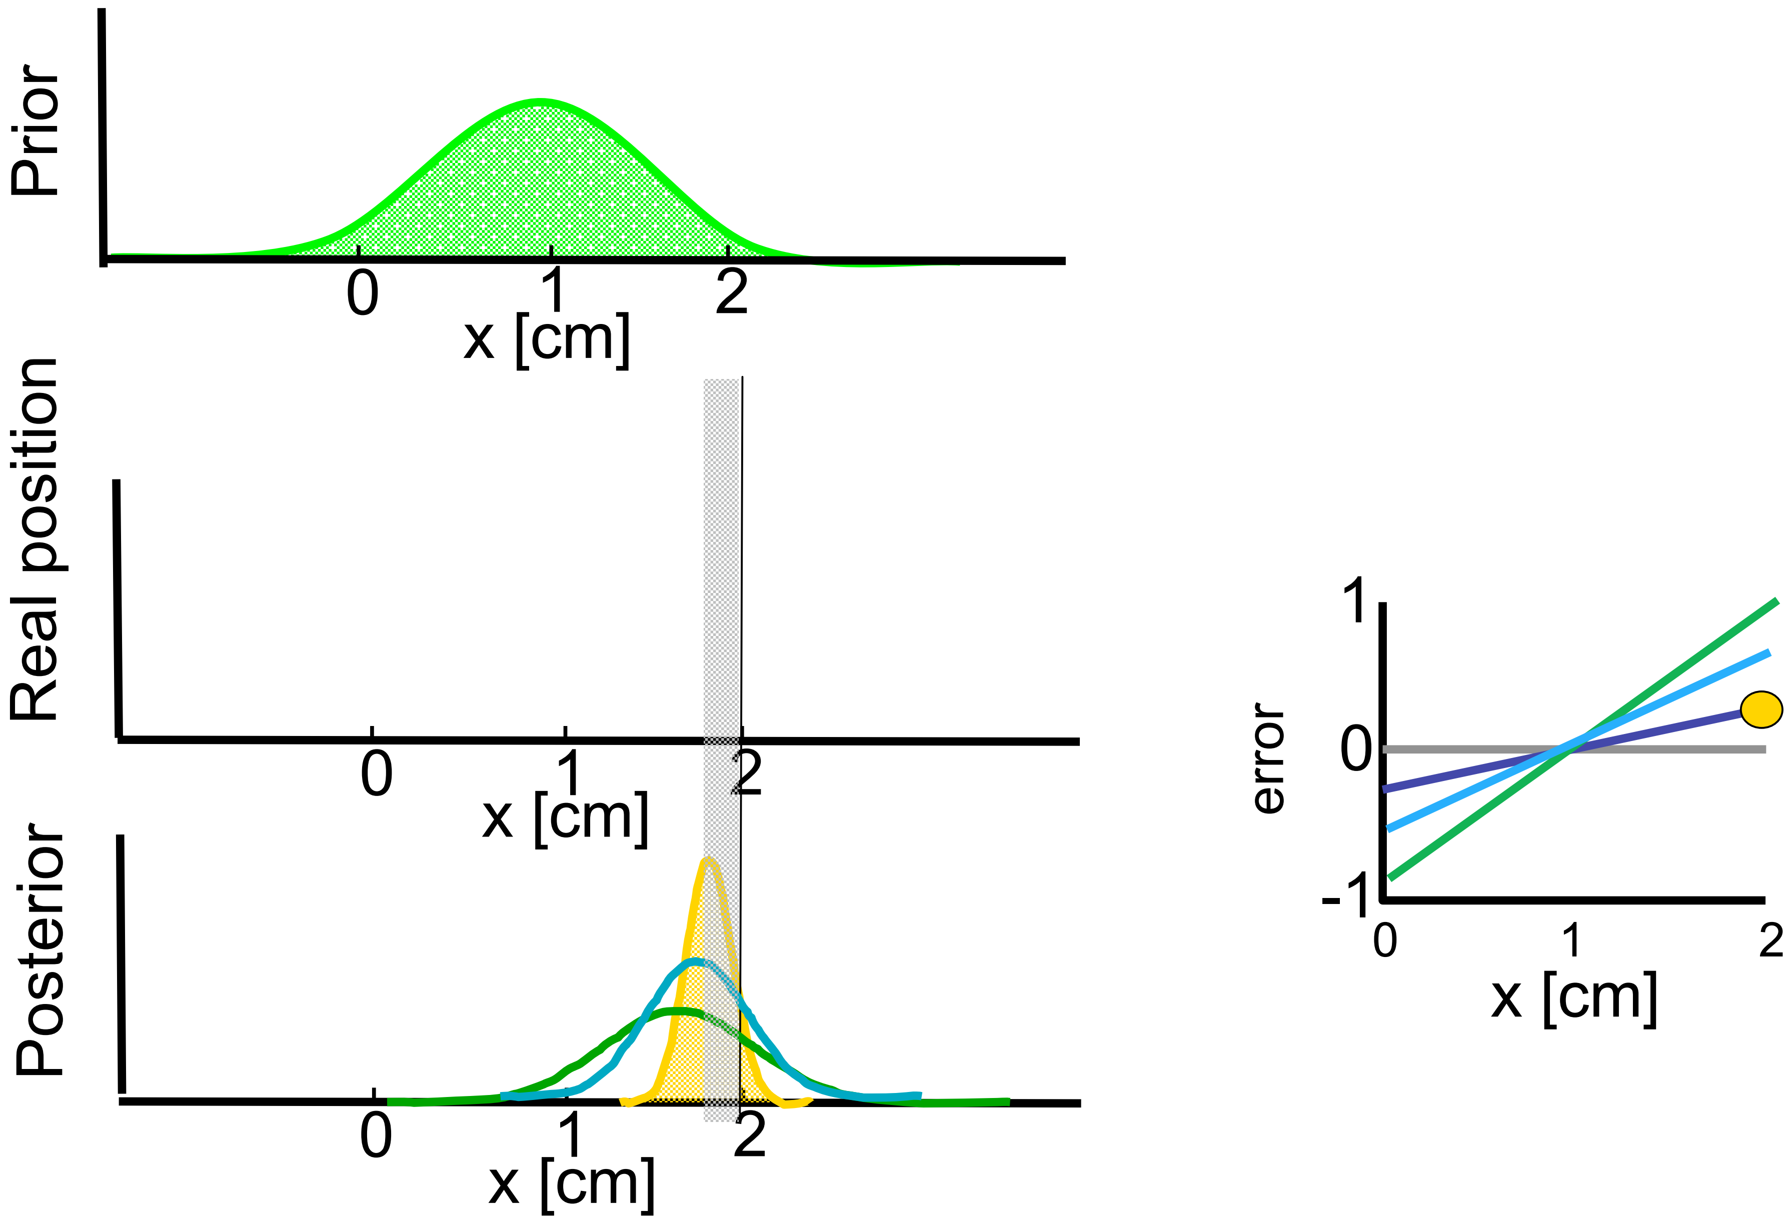

Experimentally creating uncertainty

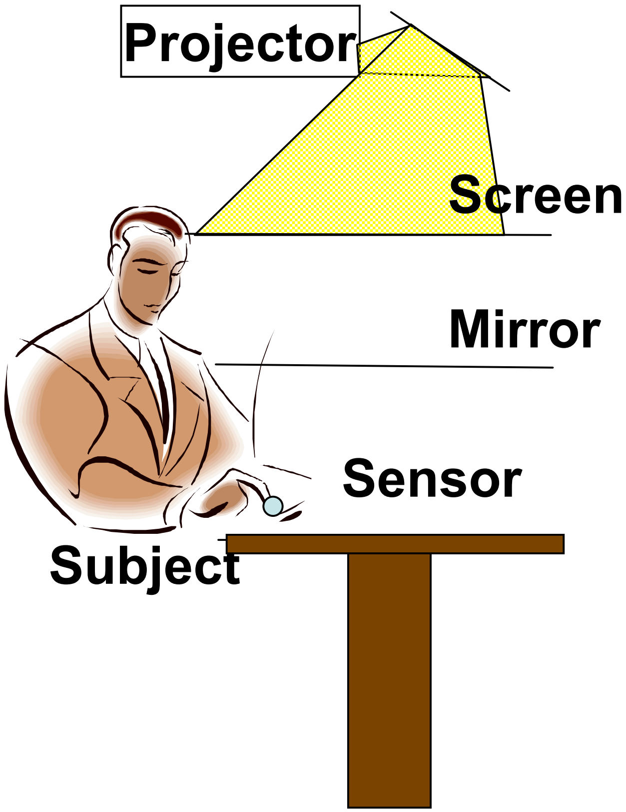

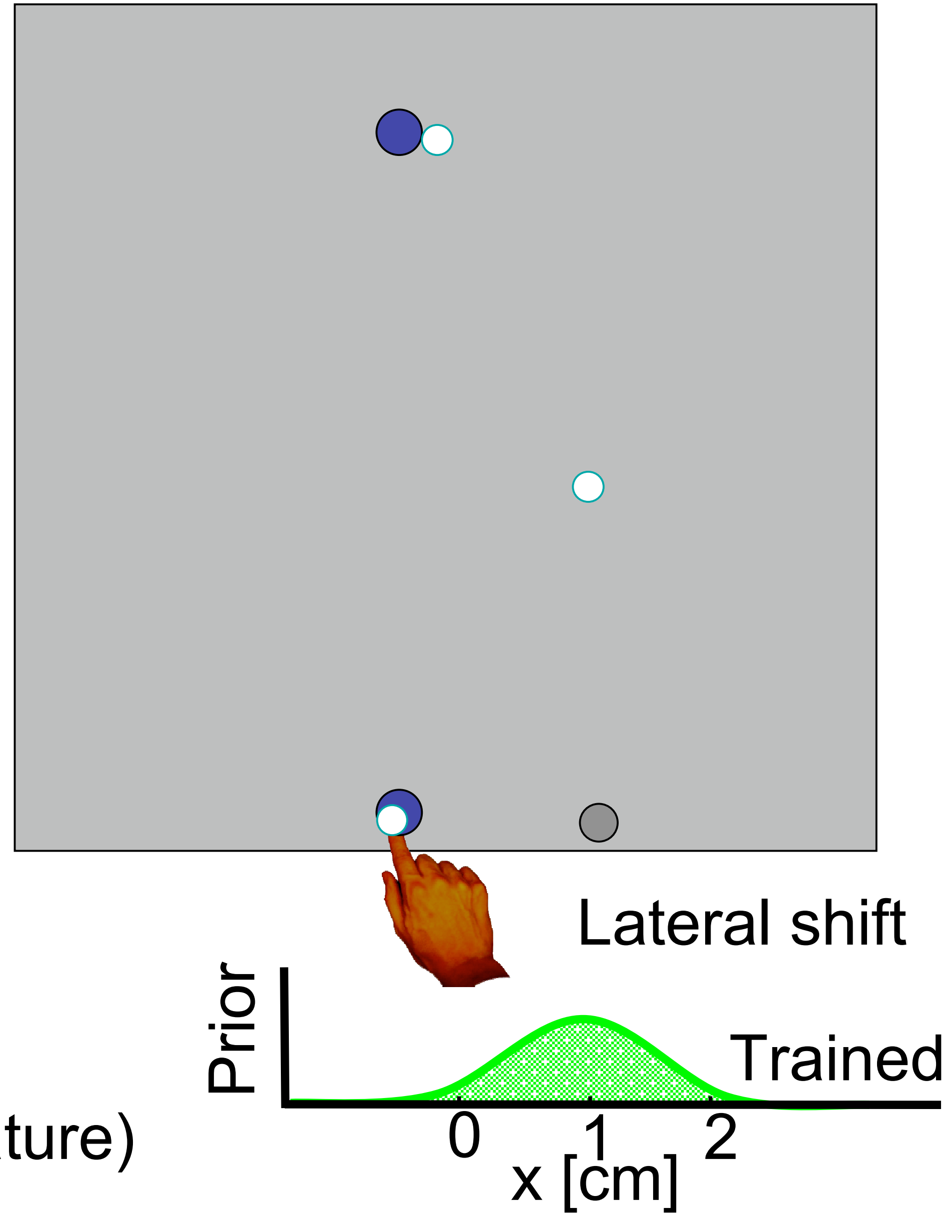

(Körding & Wolpert 2004, Nature)

Evidence: Uncertainty in the feedback

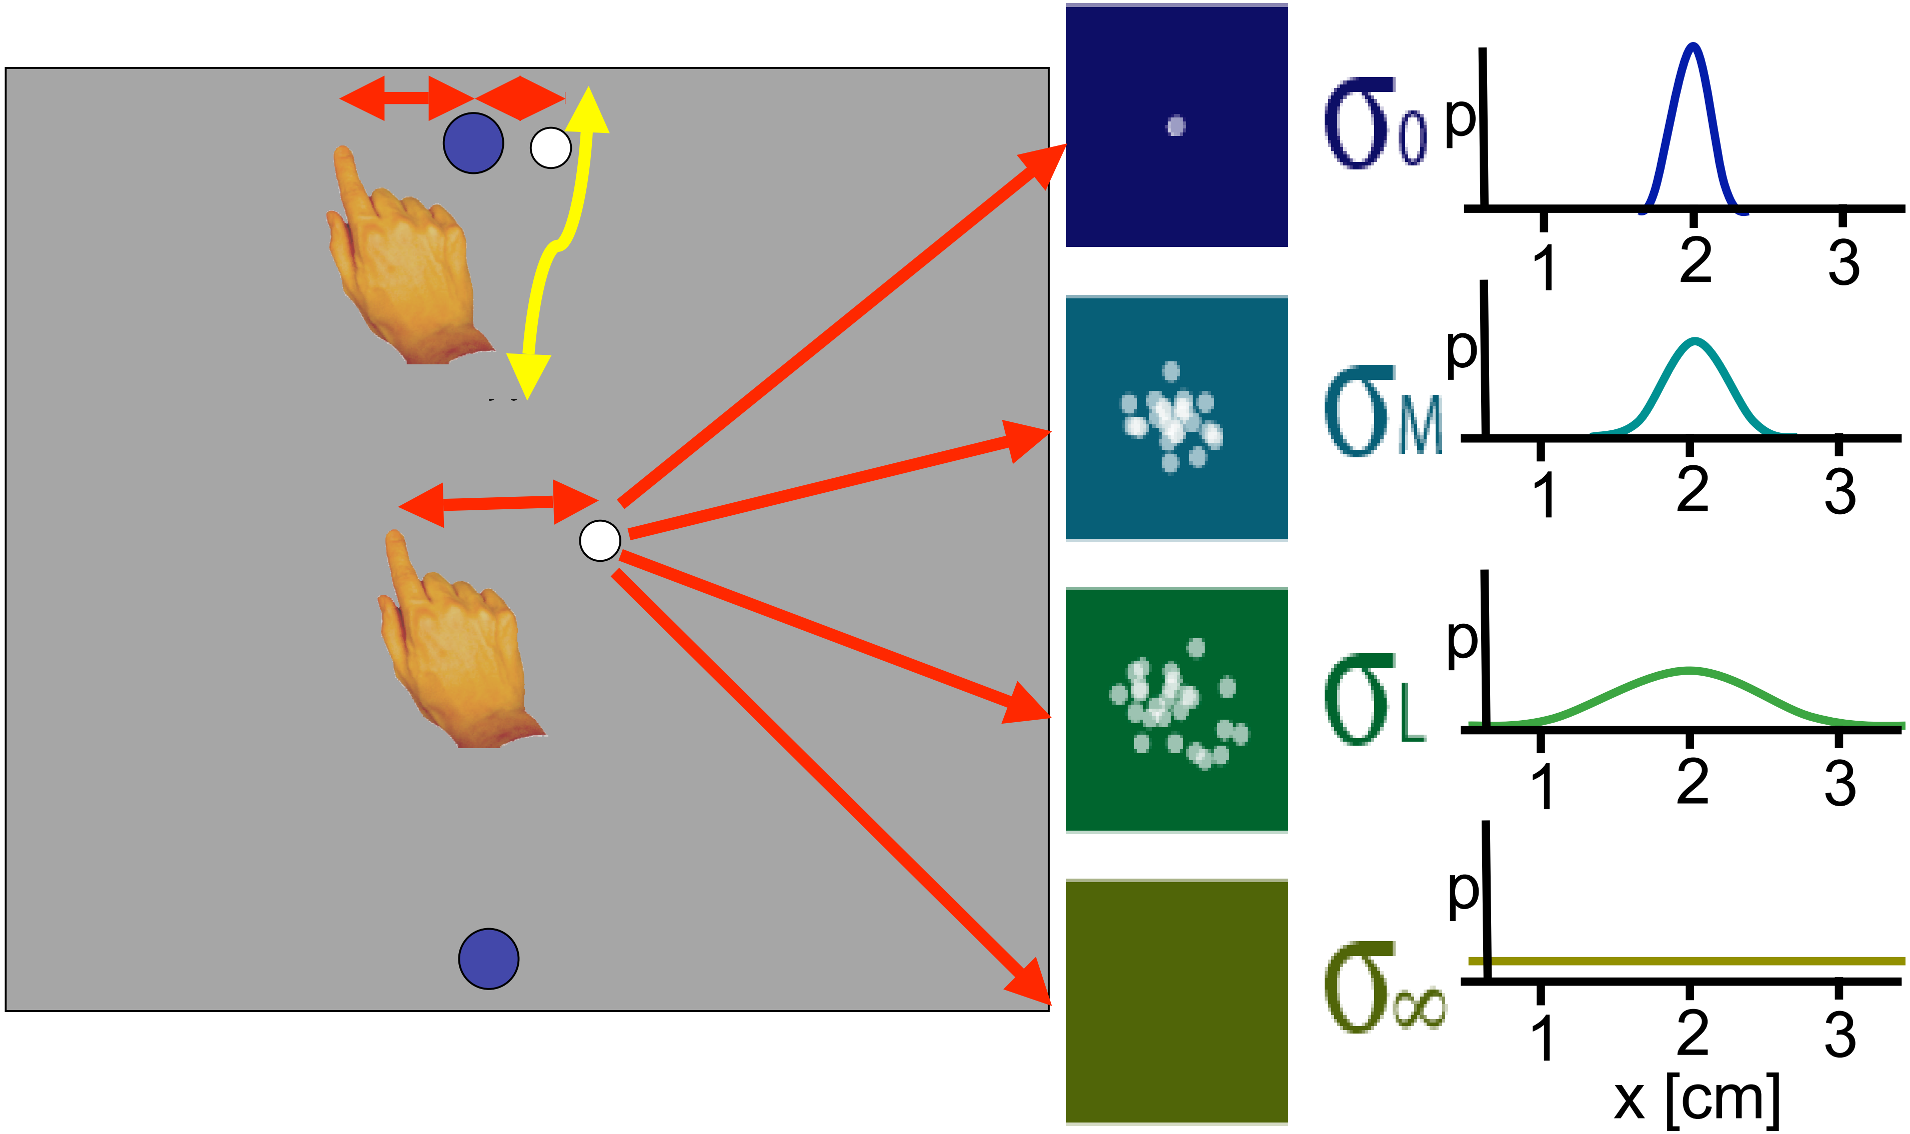

Model 1: Naïve Compensation

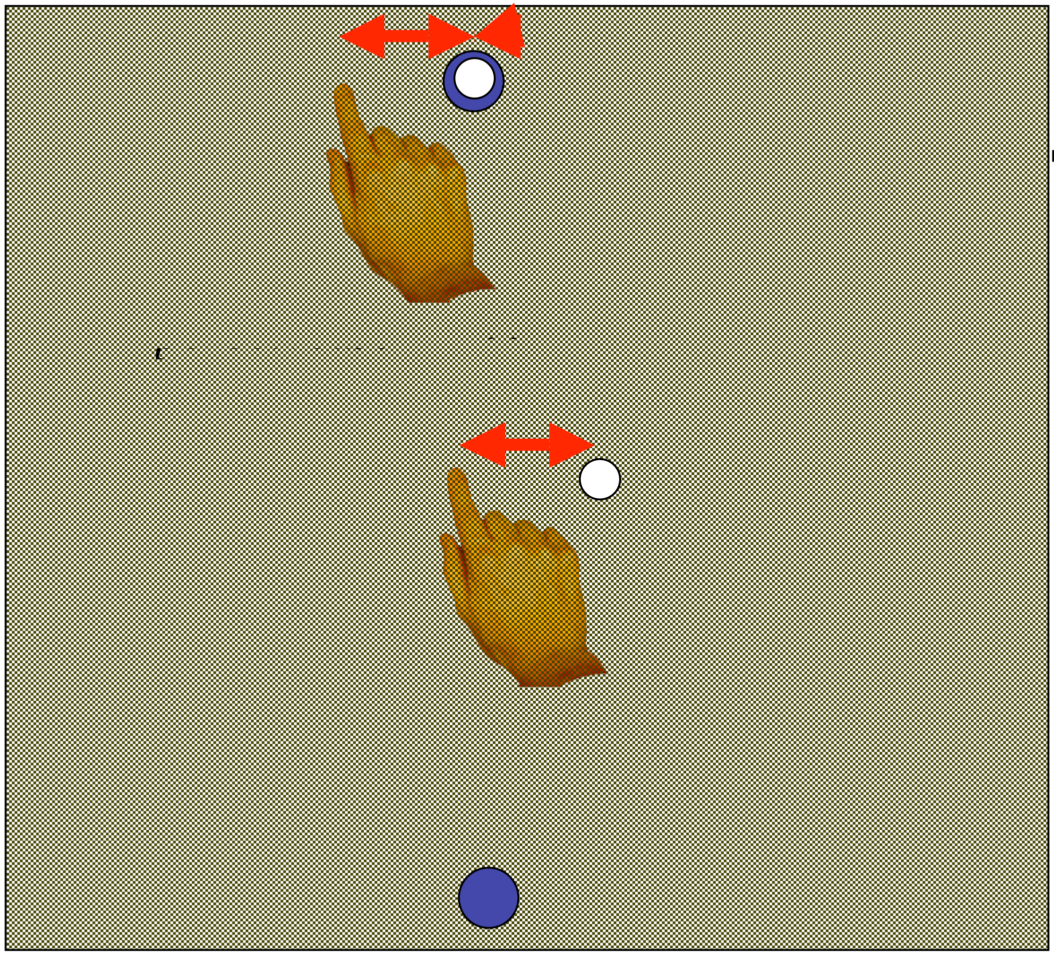

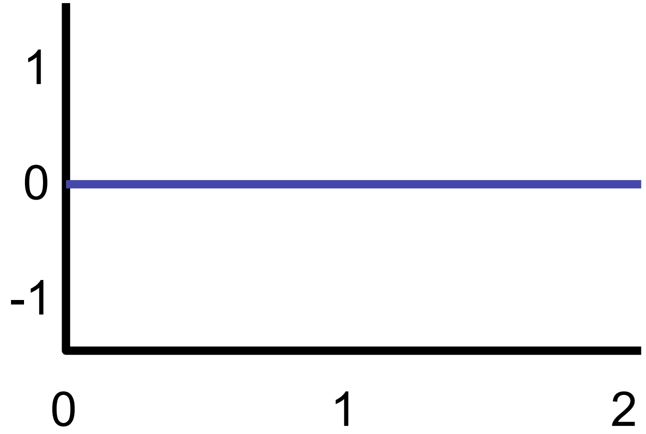

Mean and deviation

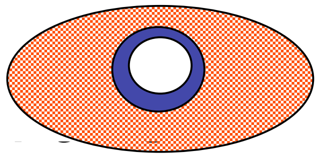

Model 2: Optimal Bayesian Compensation

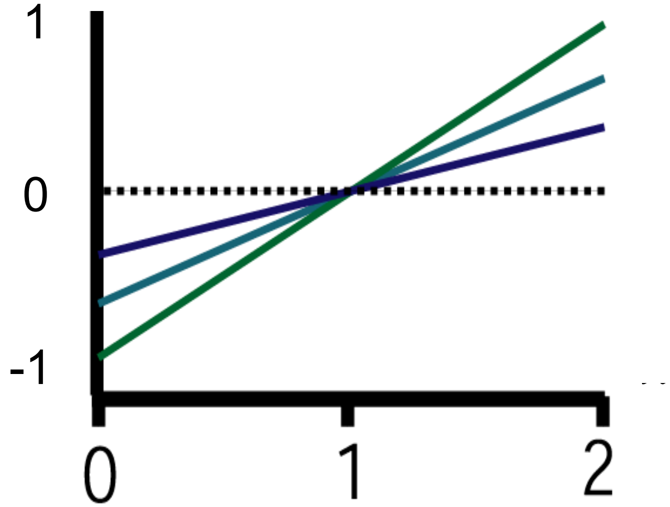

Mean and deviation

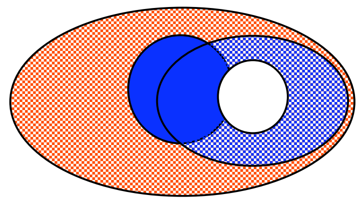

Model 1: full compensation

Model 3: Supervised Learning

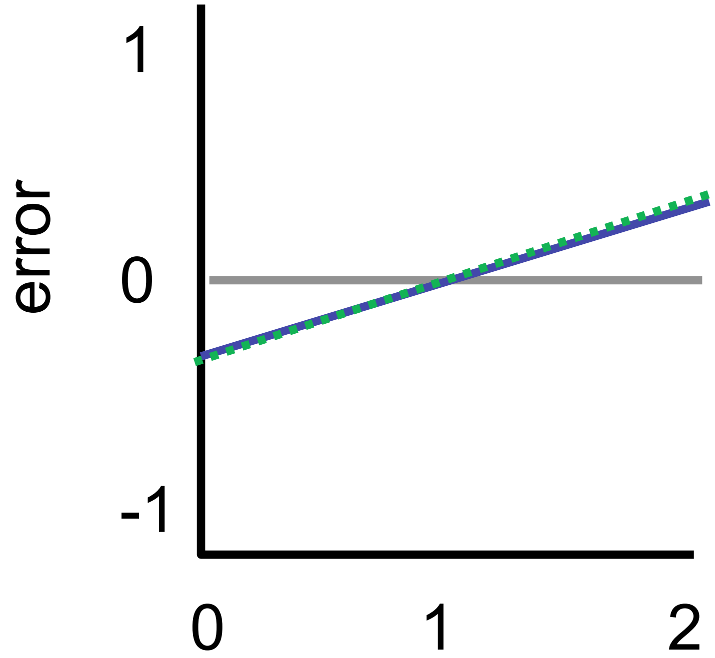

Feedback only in one case -> only one strategy can be learnt

Results: single subject

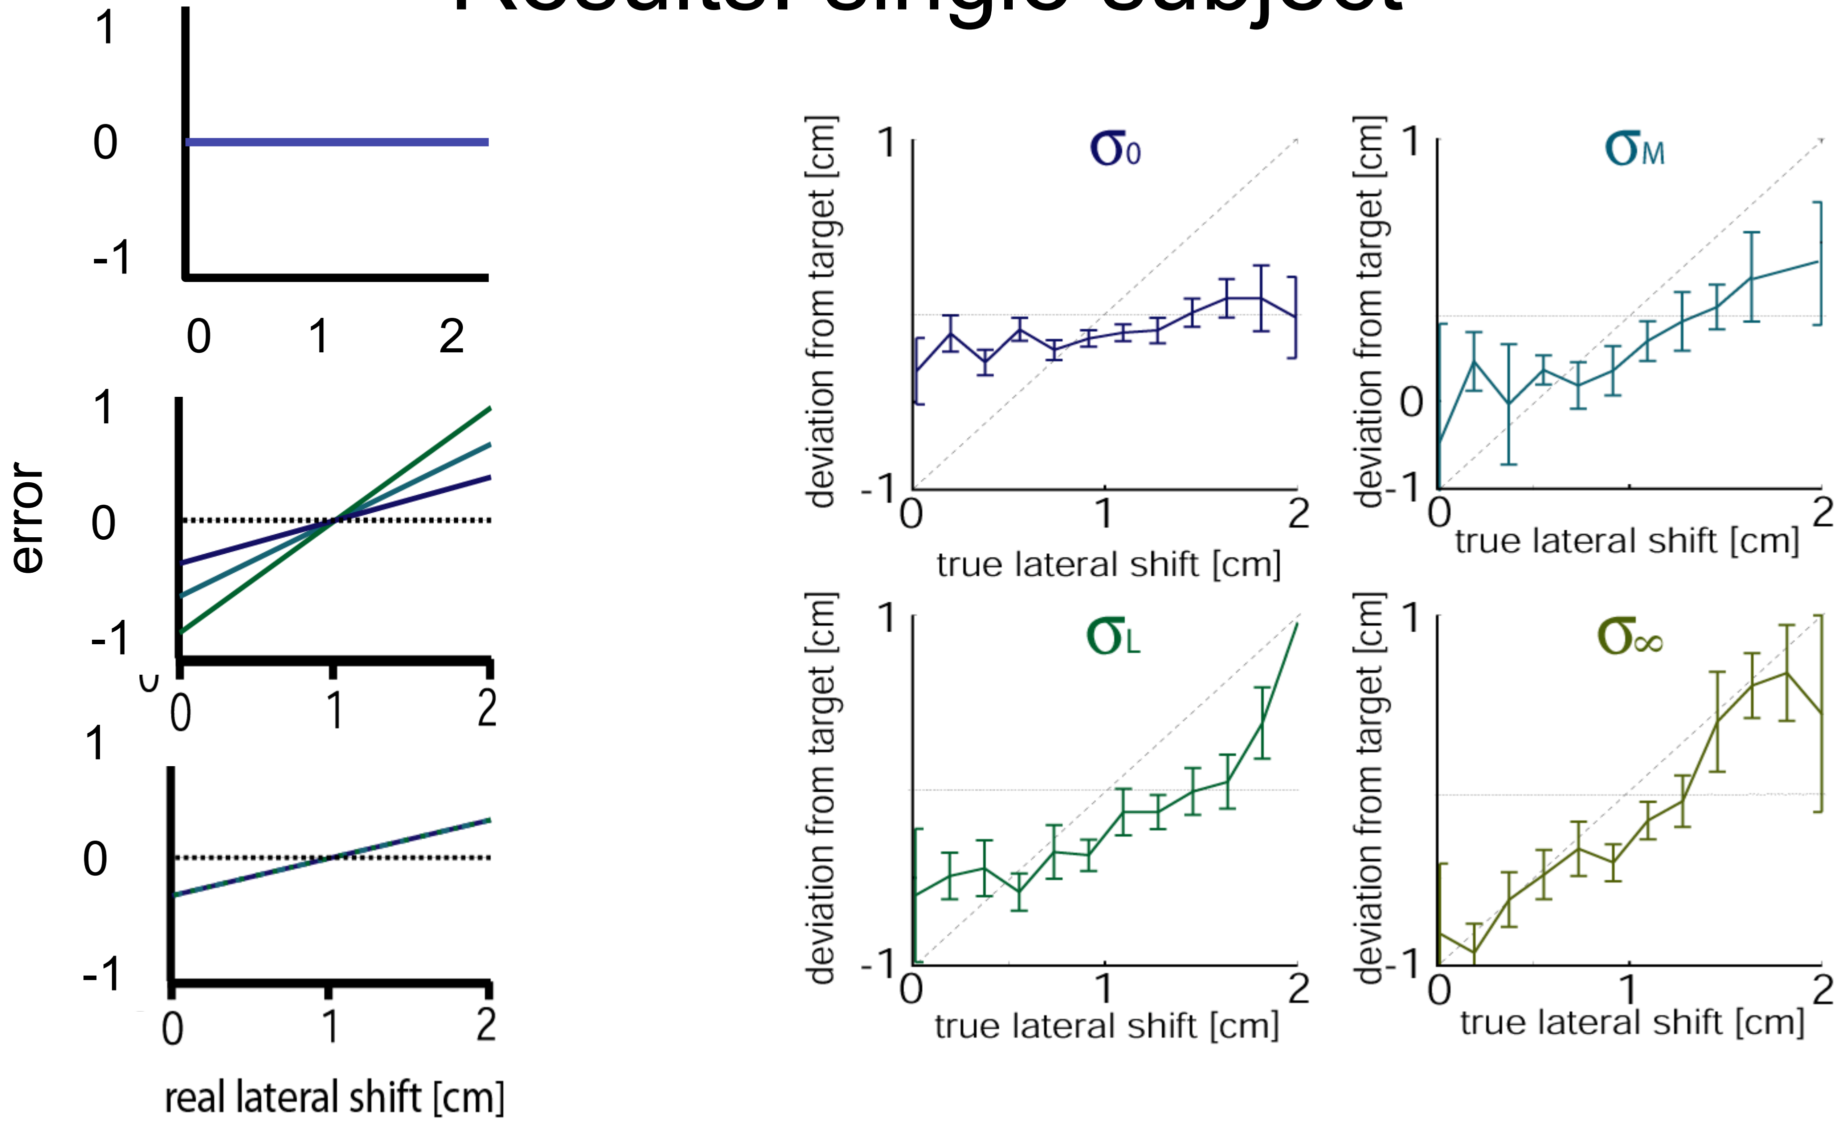

Supports model 2: Bayesian

Slope over 10 subjects

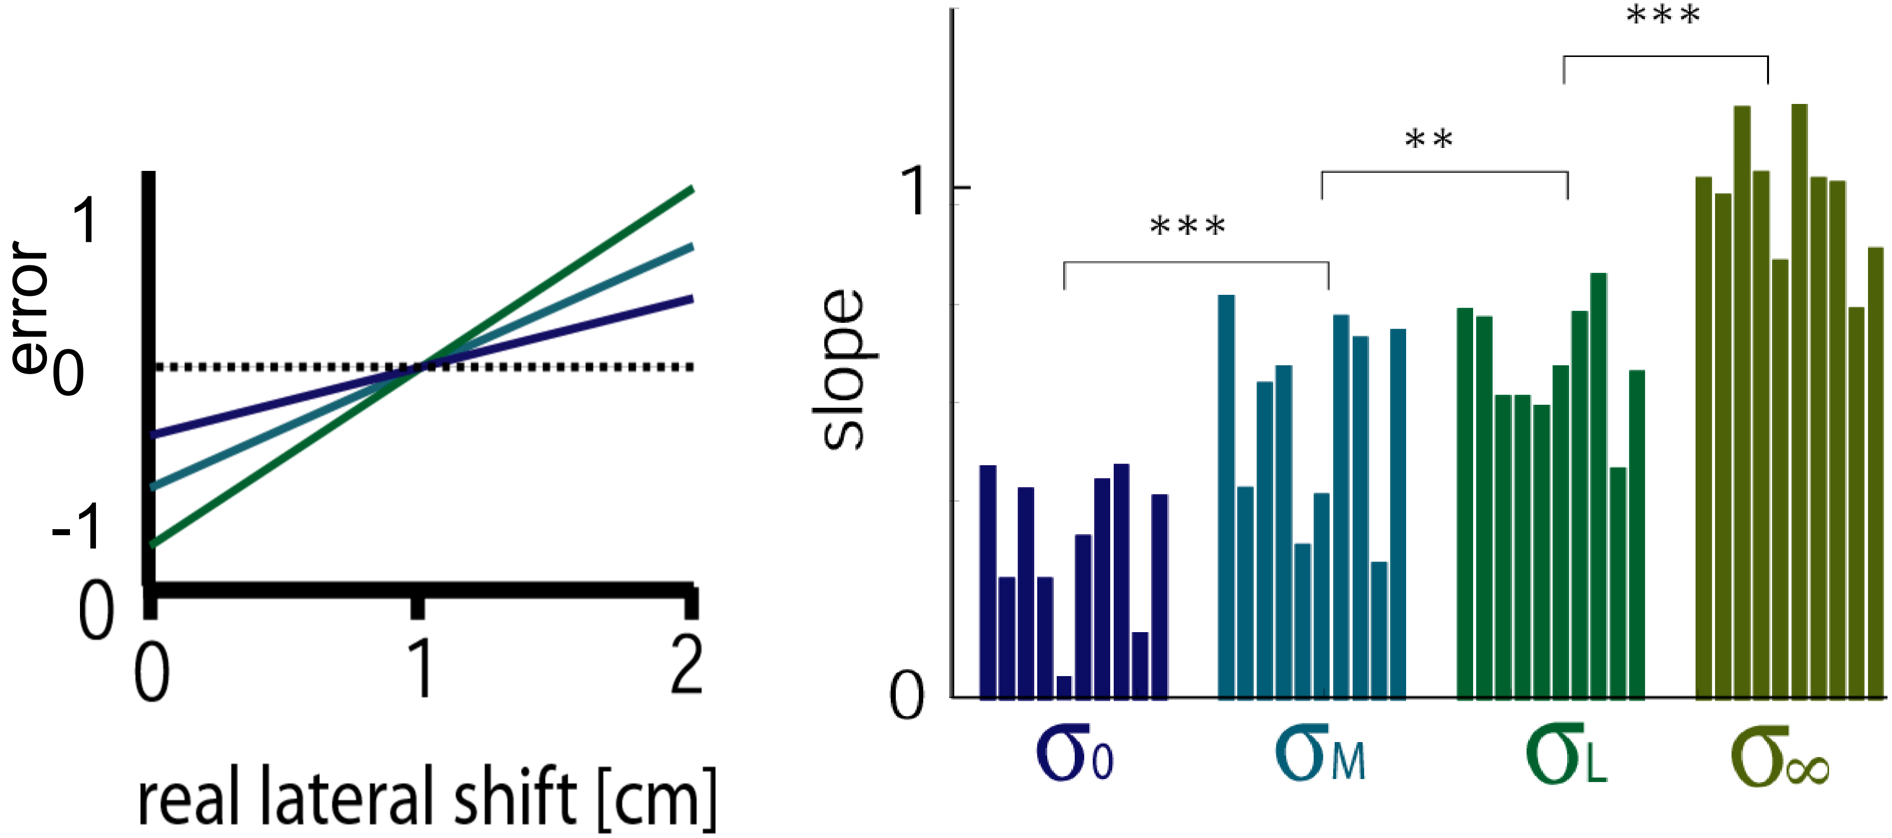

Measuring the prior

Assume subjects use maximum a posteriori strategy

See also Paninski 2004, NIPS

Estimating subjects' priors

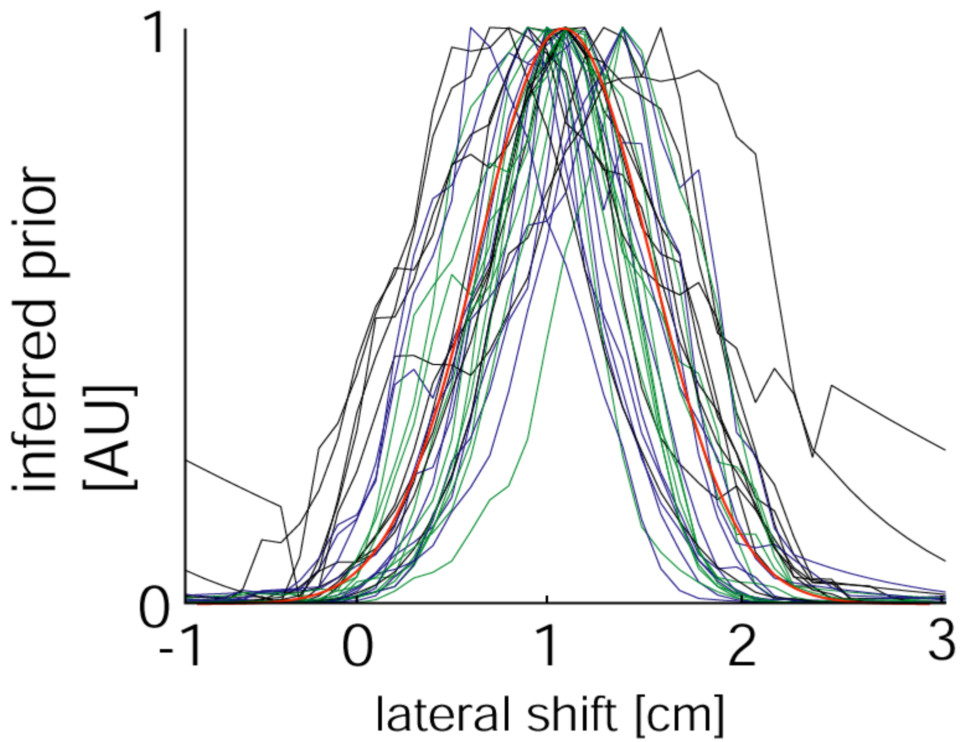

Non gaussian Distributions

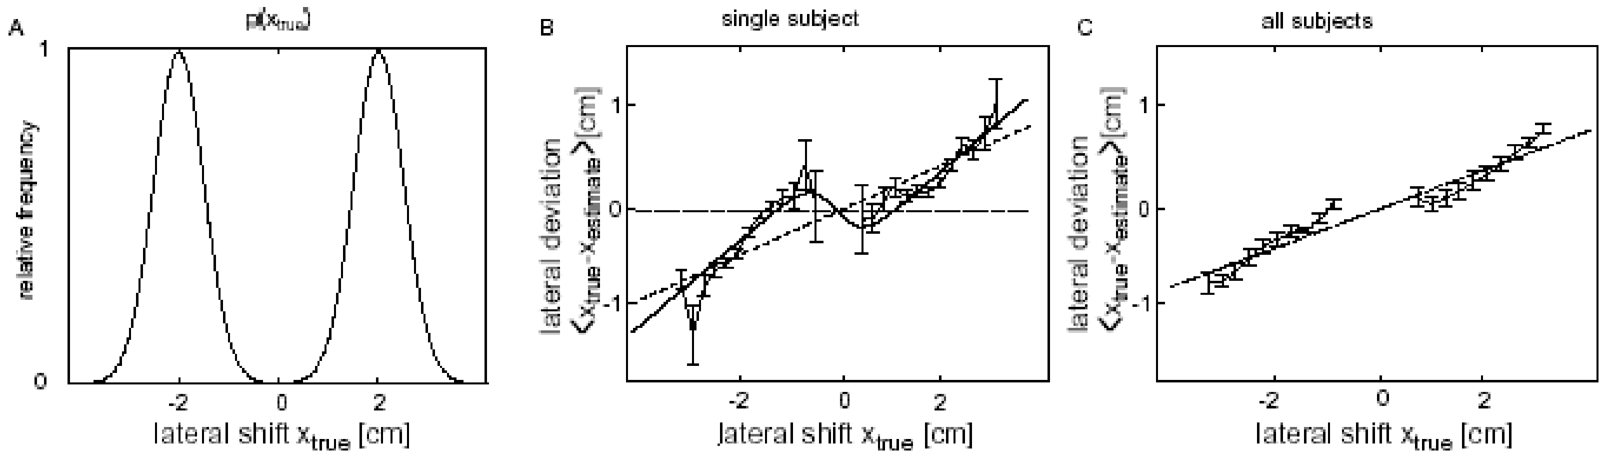

(Körding& Wolpert NIPS 2003)
(compare Miyazaki et al, J Neurophys 2005)

2) Reward and Loss

Motor outcome | Loss = -Objective=-Utility=-Reward

 = 

 = - 

 = - 

Utility functions for forces

1 hour

10 minutes

Typical example

Economics / Indifference Curves

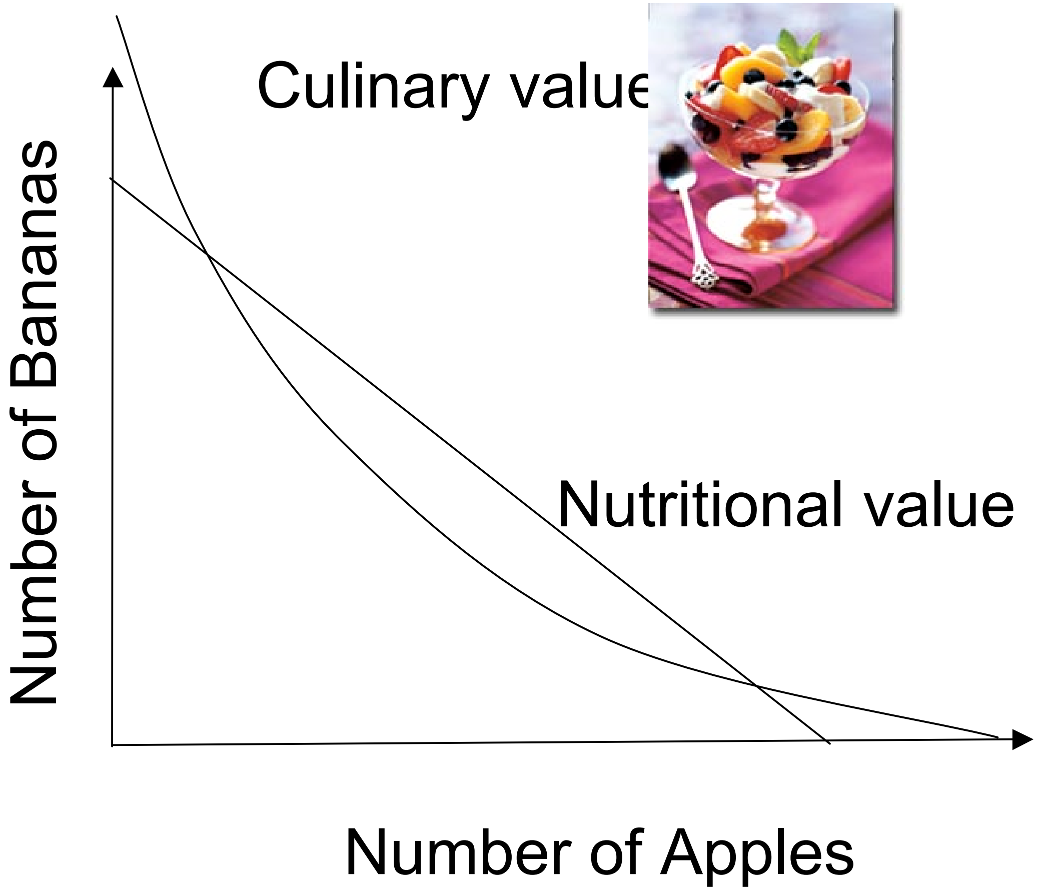

The problem space for forces

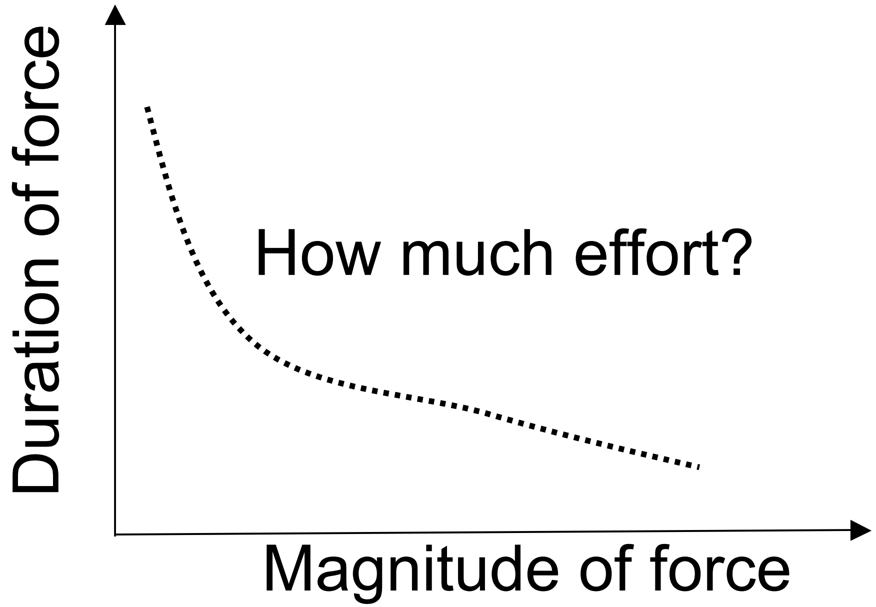

Experimental setup

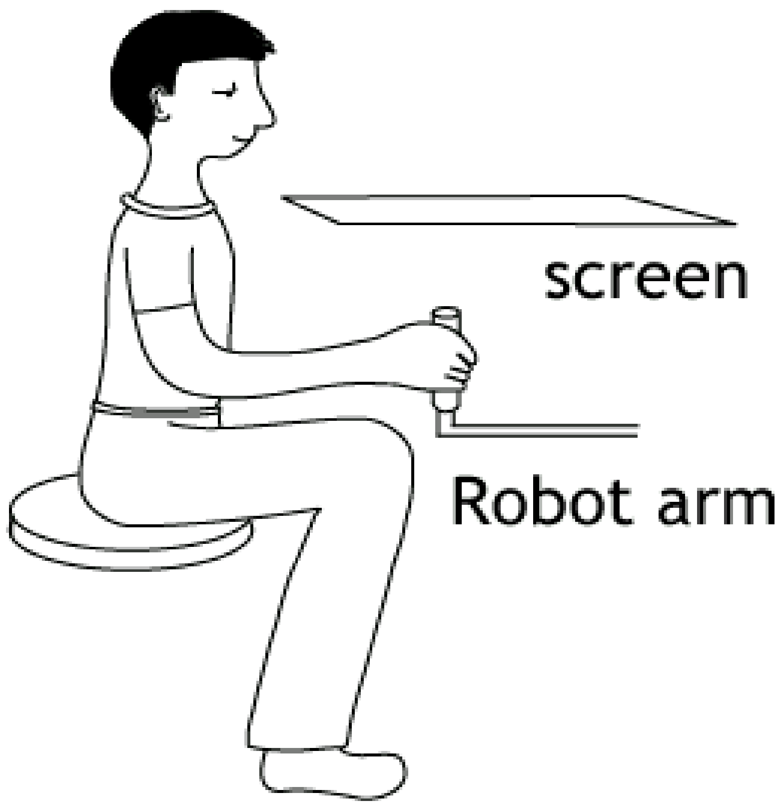

Force preference experiment

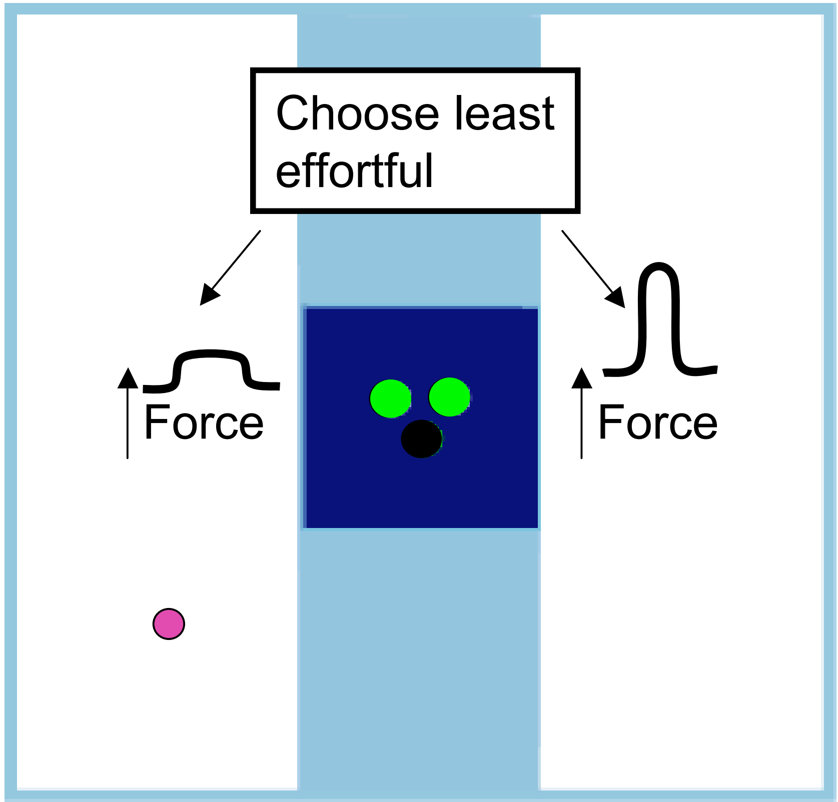

Staircase yields Effort( )=Effort( )

Compare walking experiments of Jean-Paul Laumond and his students

Results (subject 1)

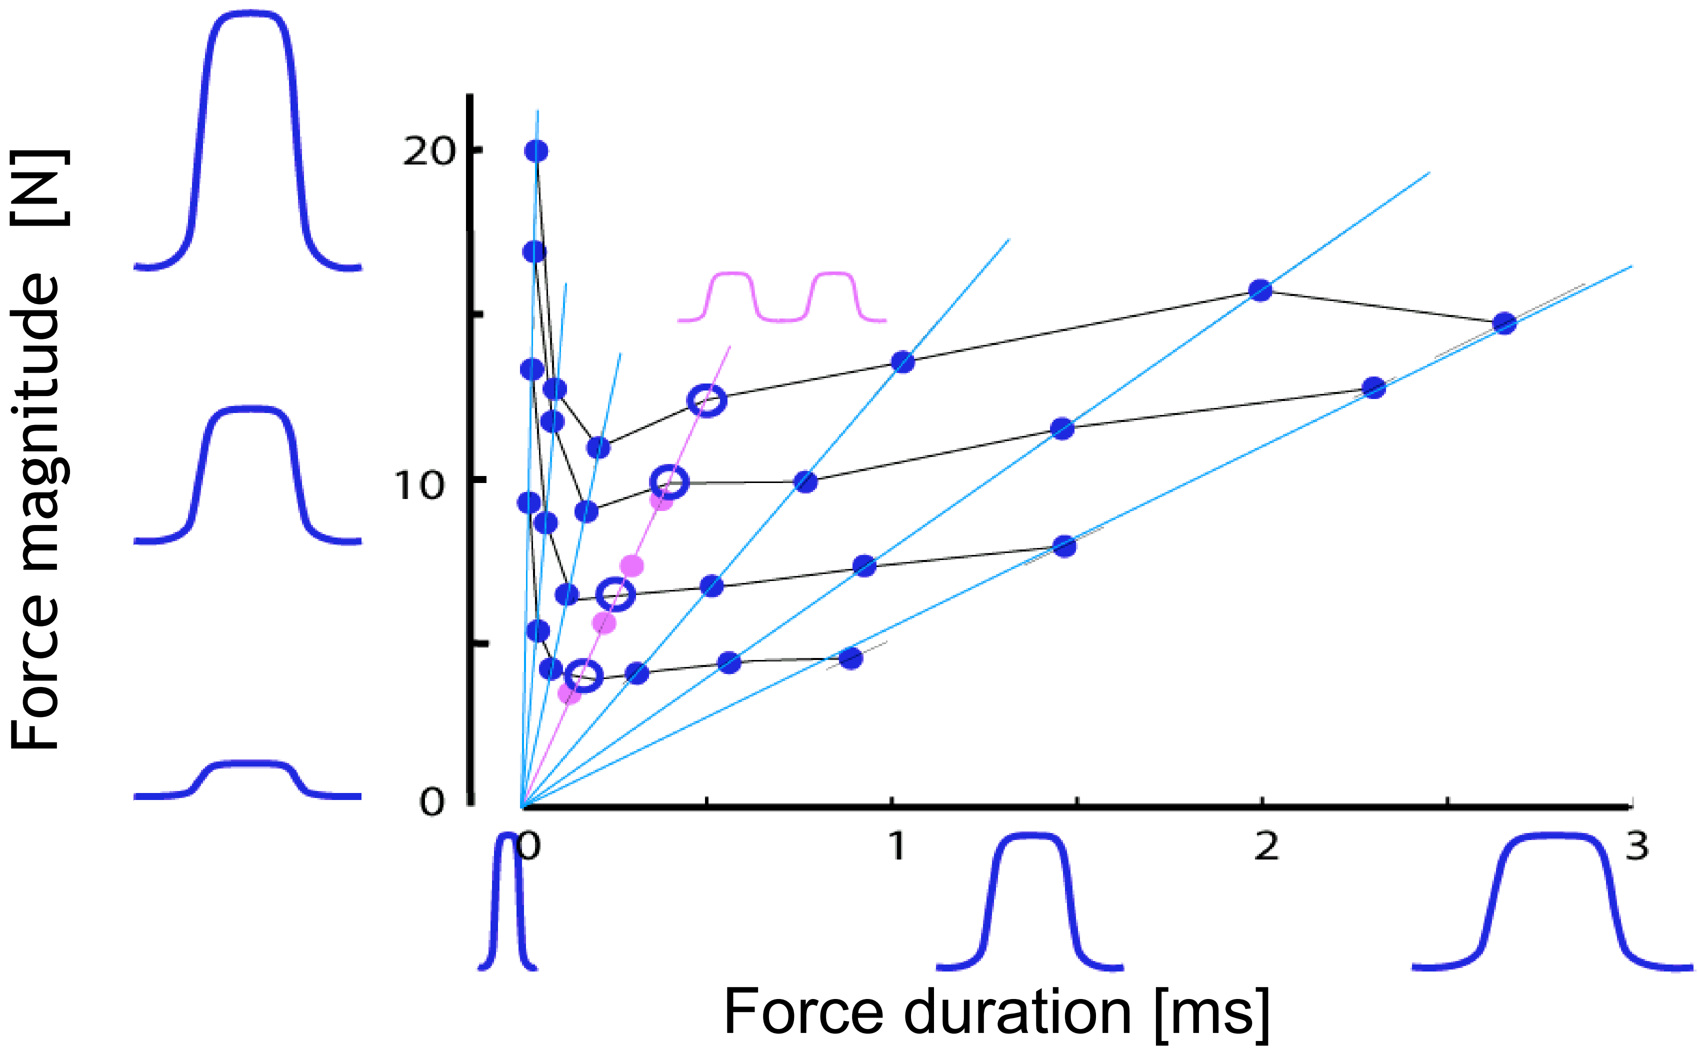

Results

(Population data n=5)

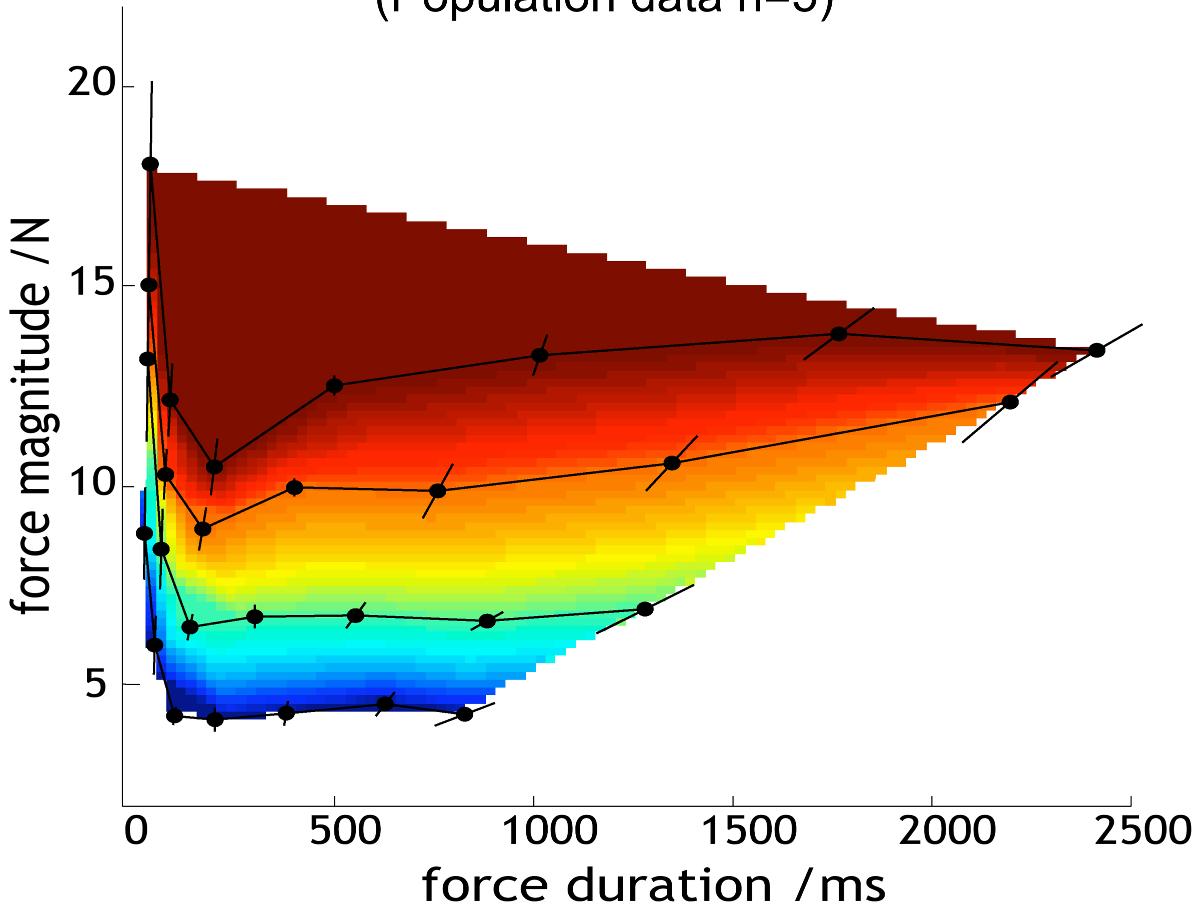

Körding, Fukunaga & Wolpert, PLOS Biology

Link to muscle properties is still missing

Conclusions

- I) People use Bayesian statistics to optimally estimate positions for movements
- II) Utility functions are useful to describe human movement decisions
- Conceptual framework for understanding movement decisions

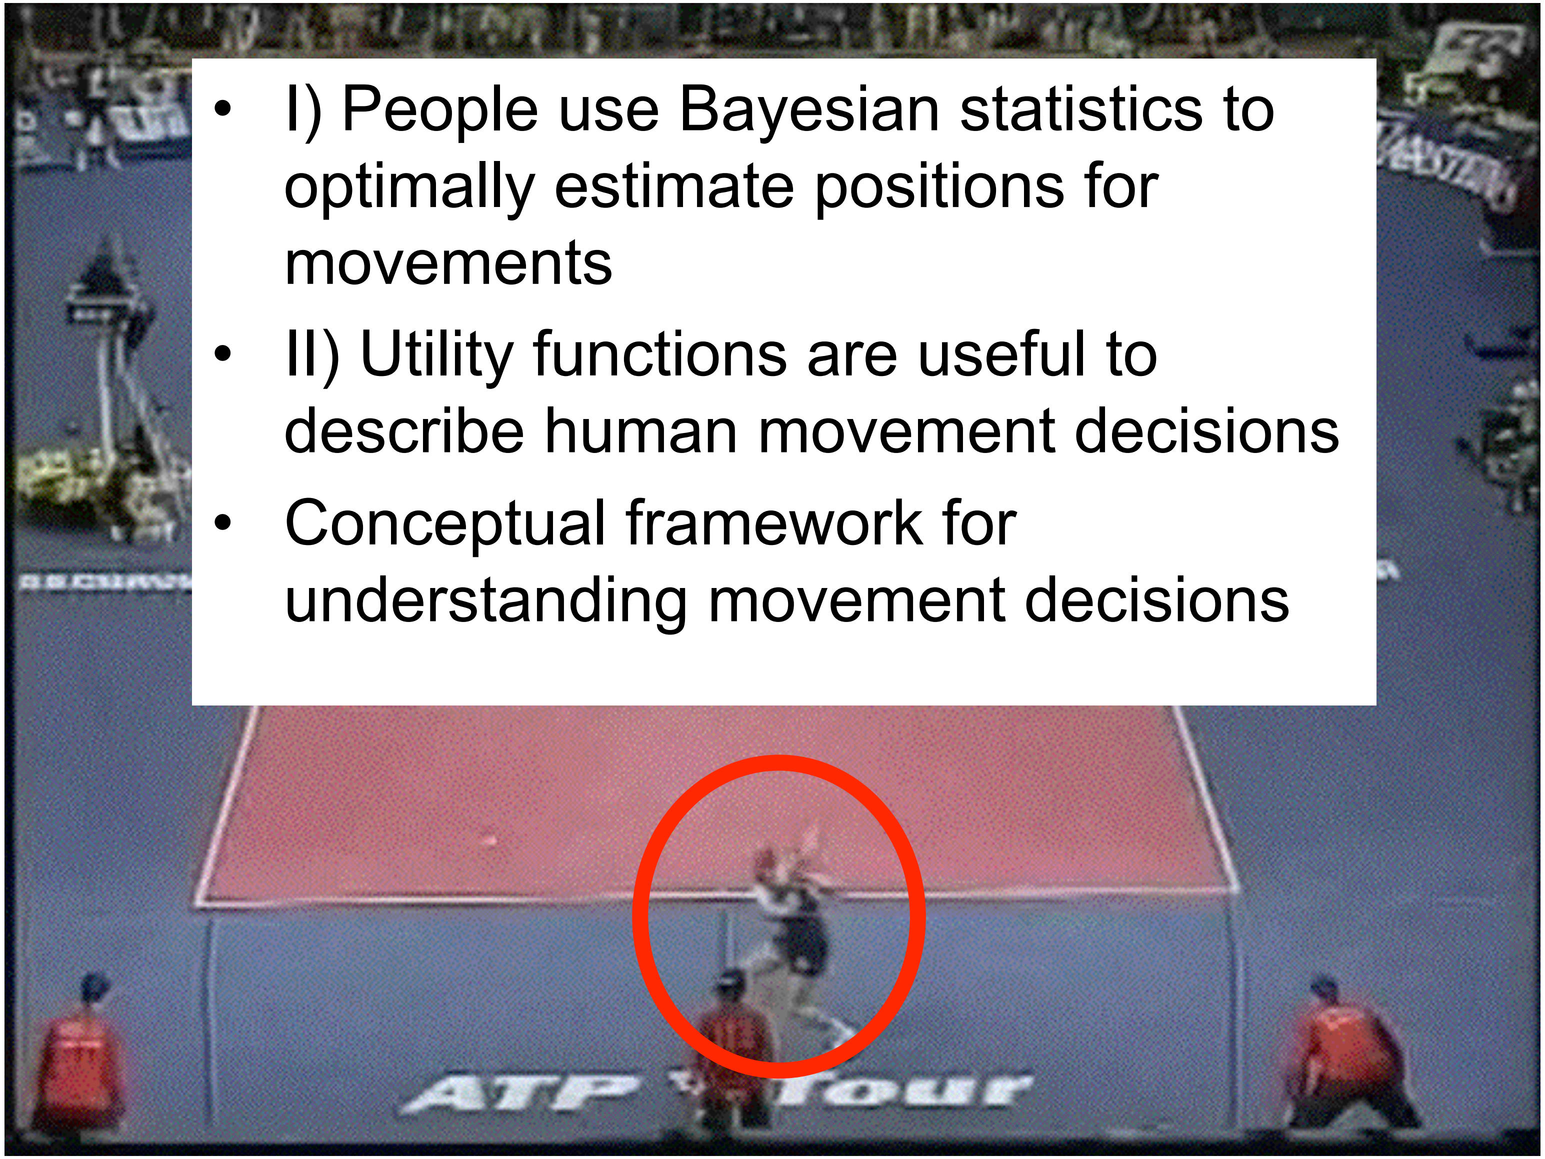

Acknowledgements

Daniel Wolpert
Izumi Fukunaga
Wolpert lab

Peter Dayan
Gatsby Computational Unit

3) Choosing a sequence (work in progress)

Reinforcement learning

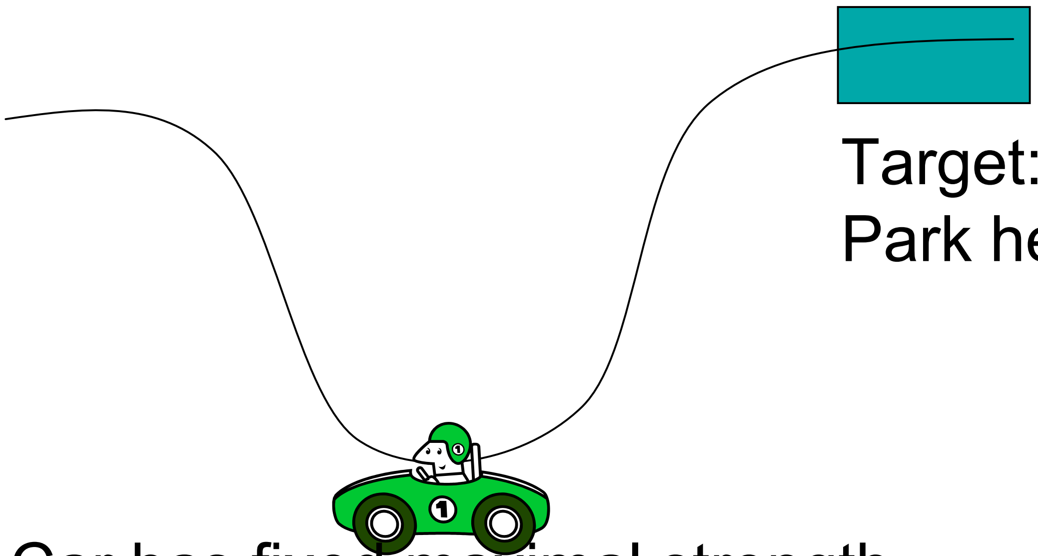

Sutton & Barto 1998

My brothers new toy

After Learning

Experimental setup

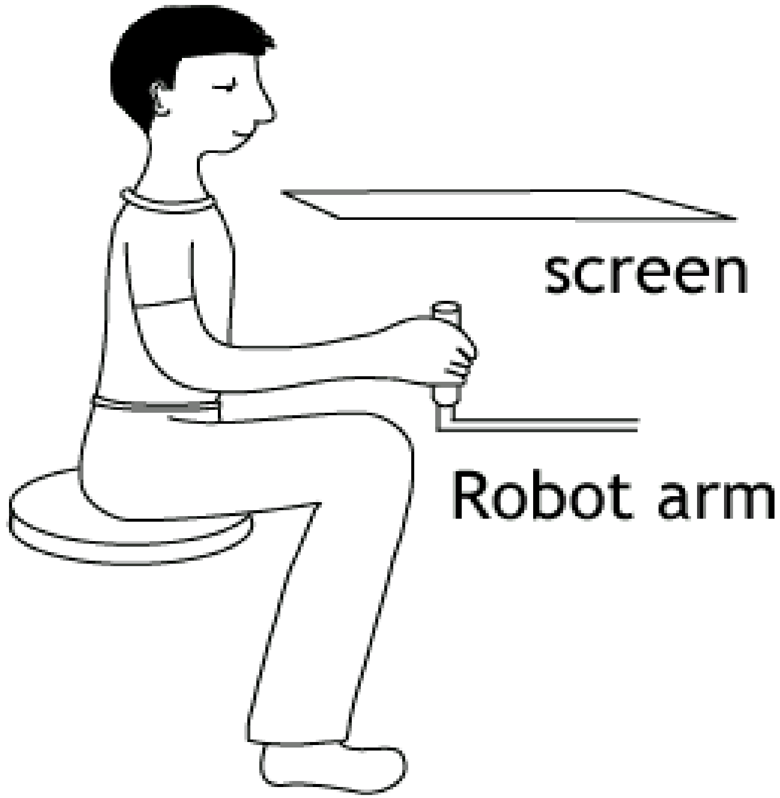

Robot simulates mass in force field

The state space

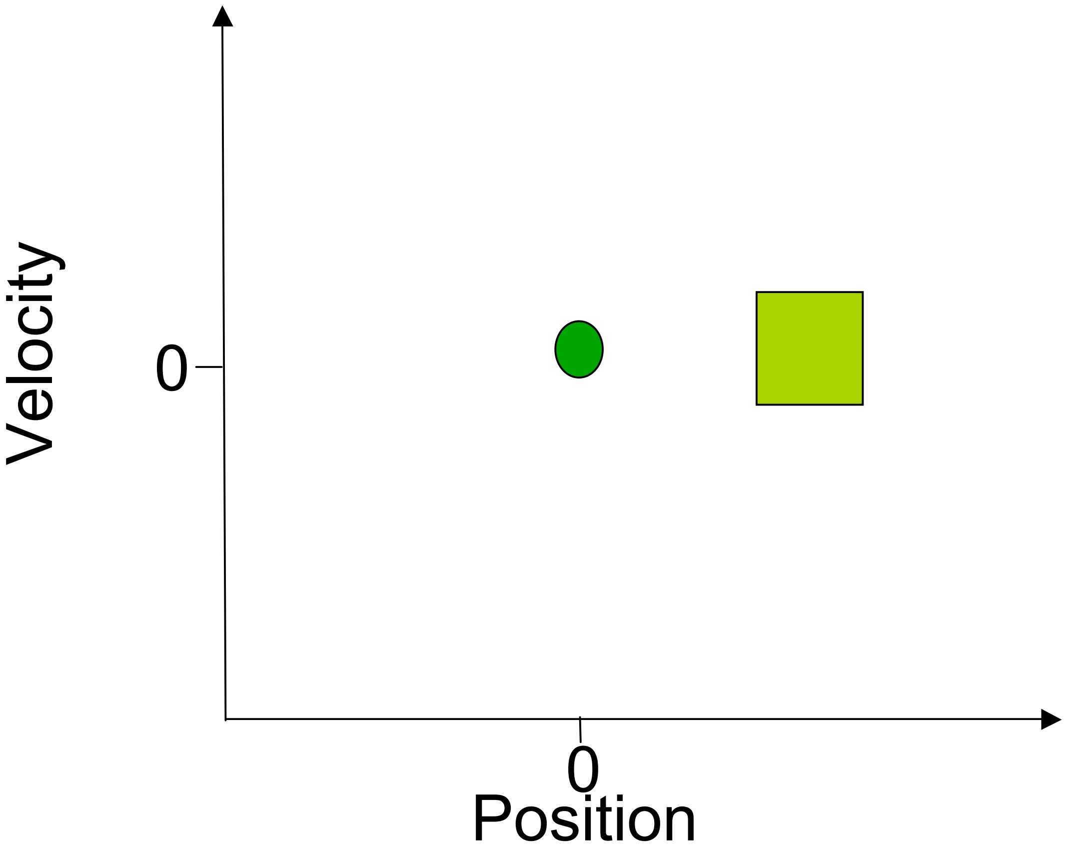

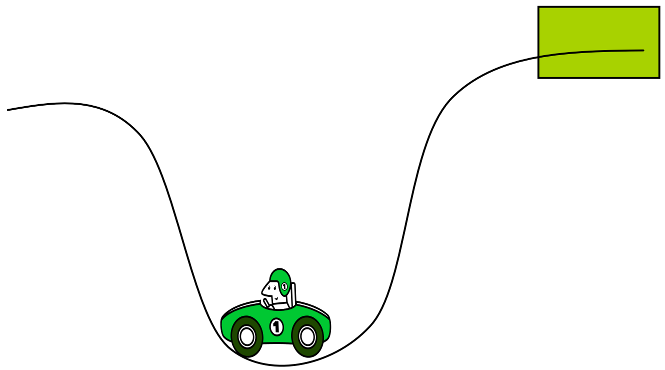

Policy= Force as function of state

Measured policy and policy predicted by theoretical reinforcement learning

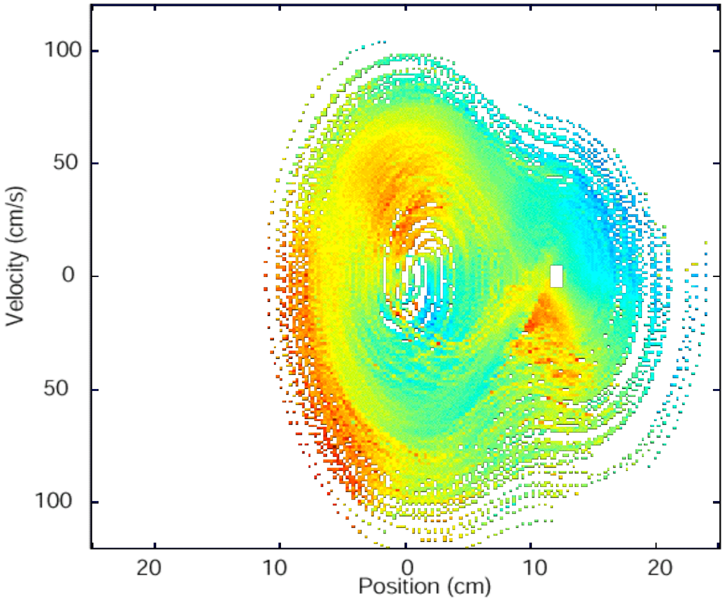

Using loss function

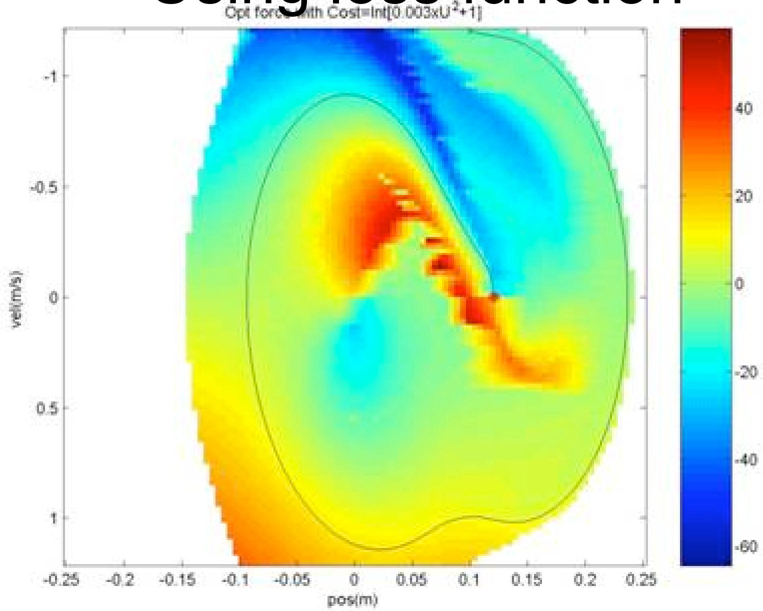

Force

Assume loss function is simple
infer Loss function _ predict behaviour
uncertainty in velocity of 20cm/s _Partially observed
delay ~250ms _ Non markovian

Goal: understanding combination probabilities, utilities