University of Antwerp

Bayesian Inference and the Integration of Multisensory Cues

Glen McLachlan

Department of Engineering Management, University of Antwerp, 2000 Antwerp, Belgium

Table of contents

Predator example

Predator example

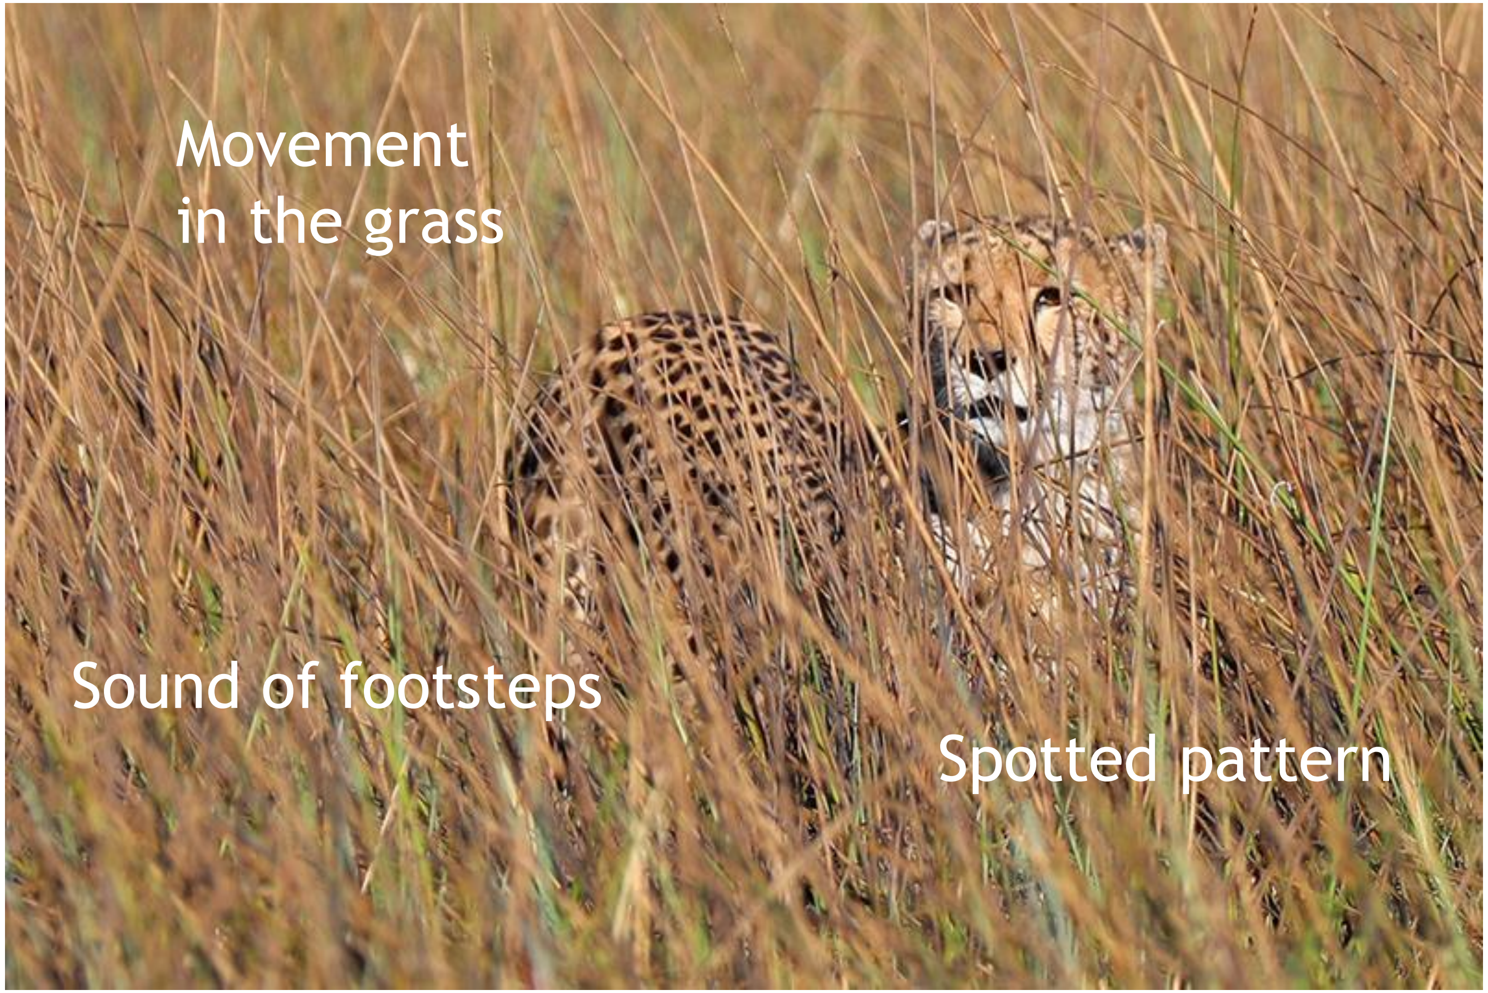

Predator example

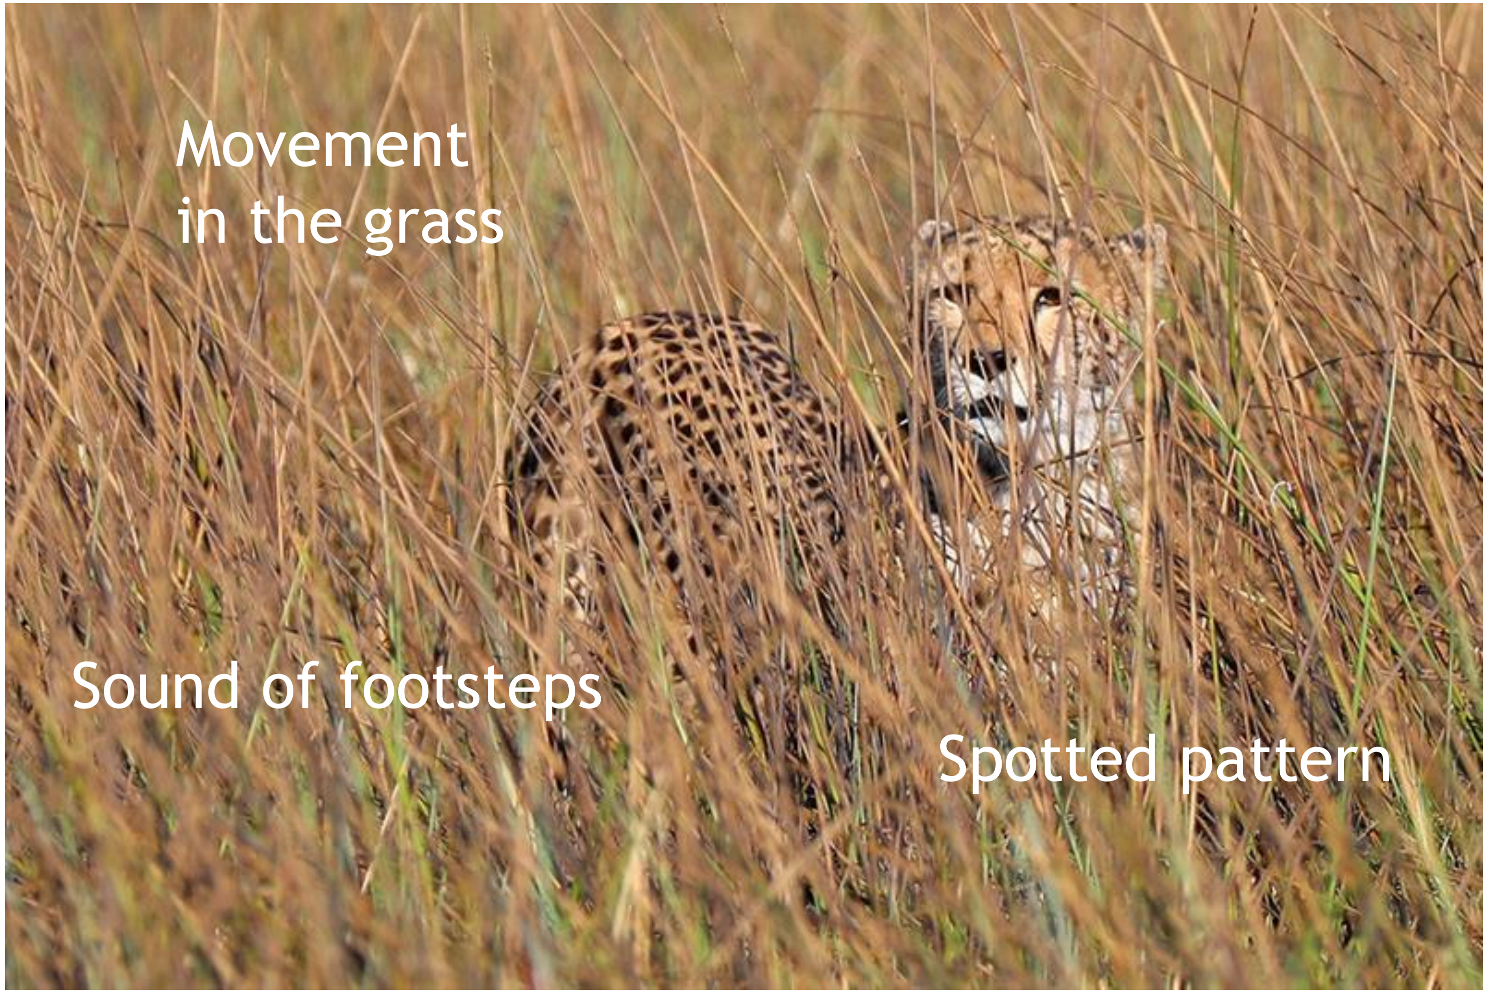

- We update our knowledge about the world by combining information across senses.

Predator example

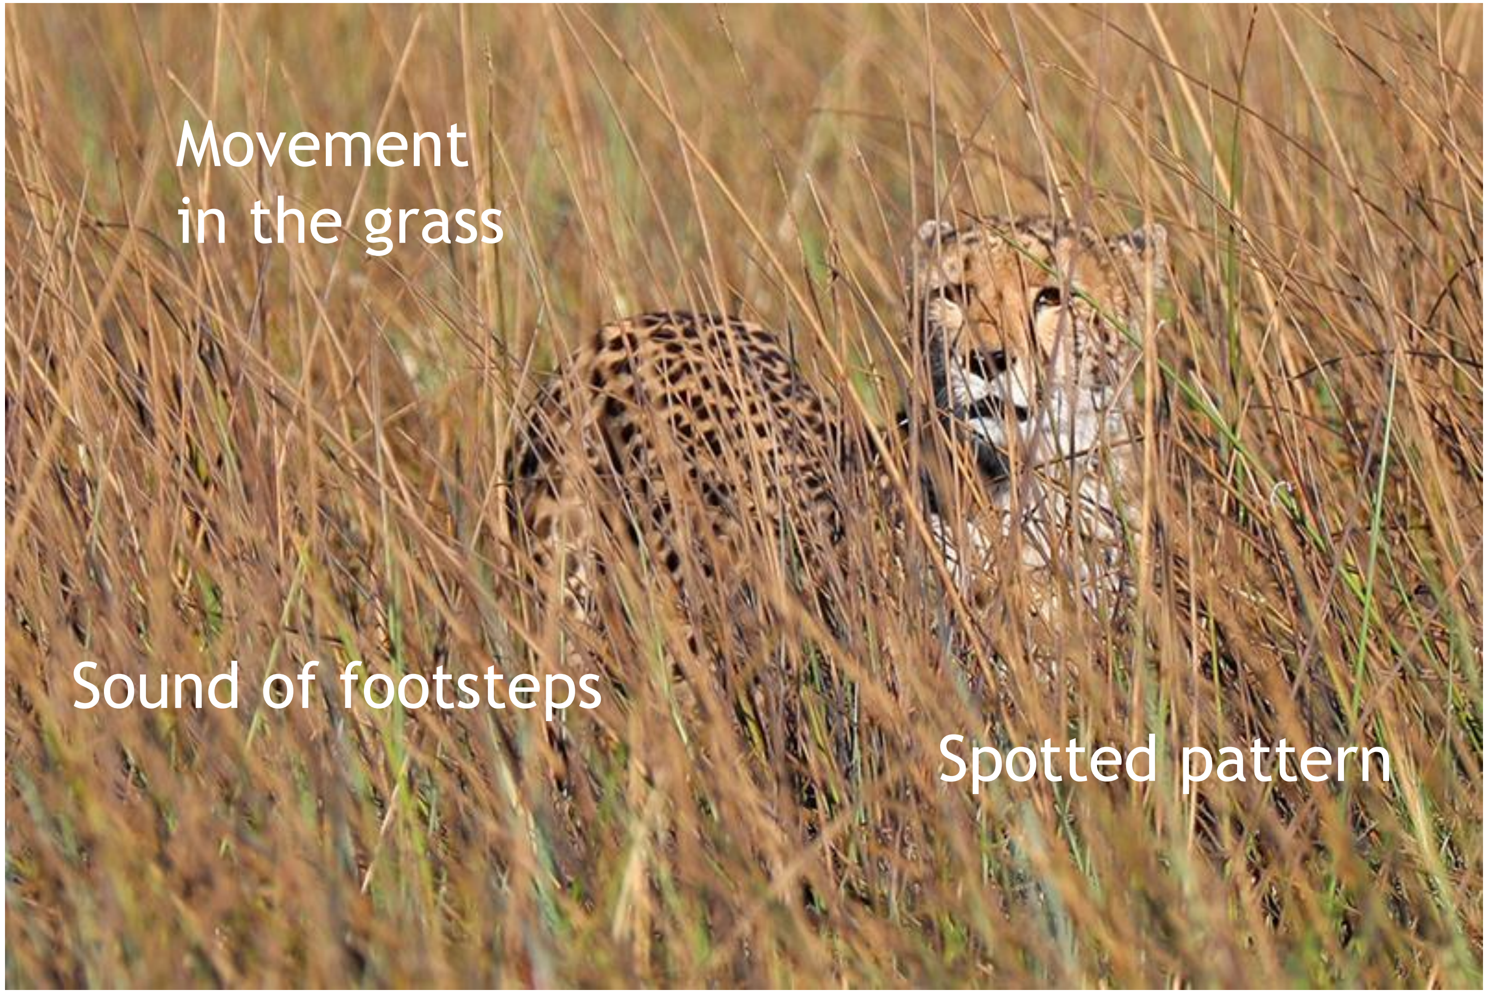

- We update our knowledge about the world by combining information across senses.
- But we also combine cues within one modality! E.g., ITD + ILD + spectral cues.

Finding confidence in uncertainty

- Information from our senses contains noise
  - External uncertainty in the signal, e.g., background noise or fog
  - Internal uncertainty in the model, e.g., variability in neuronal responses
- Still perception is generally accurate…

- We can explain this through **Bayesian inference**
  - **Inference:** forming beliefs about the environment through observations
  - **Bayesian:** by applying Bayes' Theorem

Prior, likelihood & posterior

Bayes' Theorem: $$p(H|X) = \frac{p(X|H)p(H)}{p(X)} \propto p(X|H)p(H)$$

- p(H) **prior**: probability of hypothesis H prior to measurements X
- p(X|H) **likelihood**: probability of measurements X given hypothesis H
- p(H|X) **posterior**: probability of hypothesis H given measurements X
- p(X): probability of measurements (functions as a normalisation constant)

Probability of having a rare disease H given a positive test

$$\textbf{Bayes' Theorem:} \quad p(H|X) = \frac{0.99 \cdot 0.001}{0.01098}$$

- p(H) **prior**: probability of having the disease
- p(X|H) **likelihood**: probability of positive test if we have the disease
- p(H|X) **posterior**: probability of having the disease if we test positive
- P(X): probability of testing positive

Probability of having a rare disease H given a positive test

$$\textbf{Bayes' Theorem:} \quad 0.09016 = \frac{0.99 \cdot 0.001}{0.01098}$$

- p(H) **prior**: probability of having the disease
- p(X|H) **likelihood**: probability of positive test if we have the disease
- p(H|X) **posterior**: probability of having the disease if we test positive
- P(X): probability of testing positive

Prior, likelihood & posterior

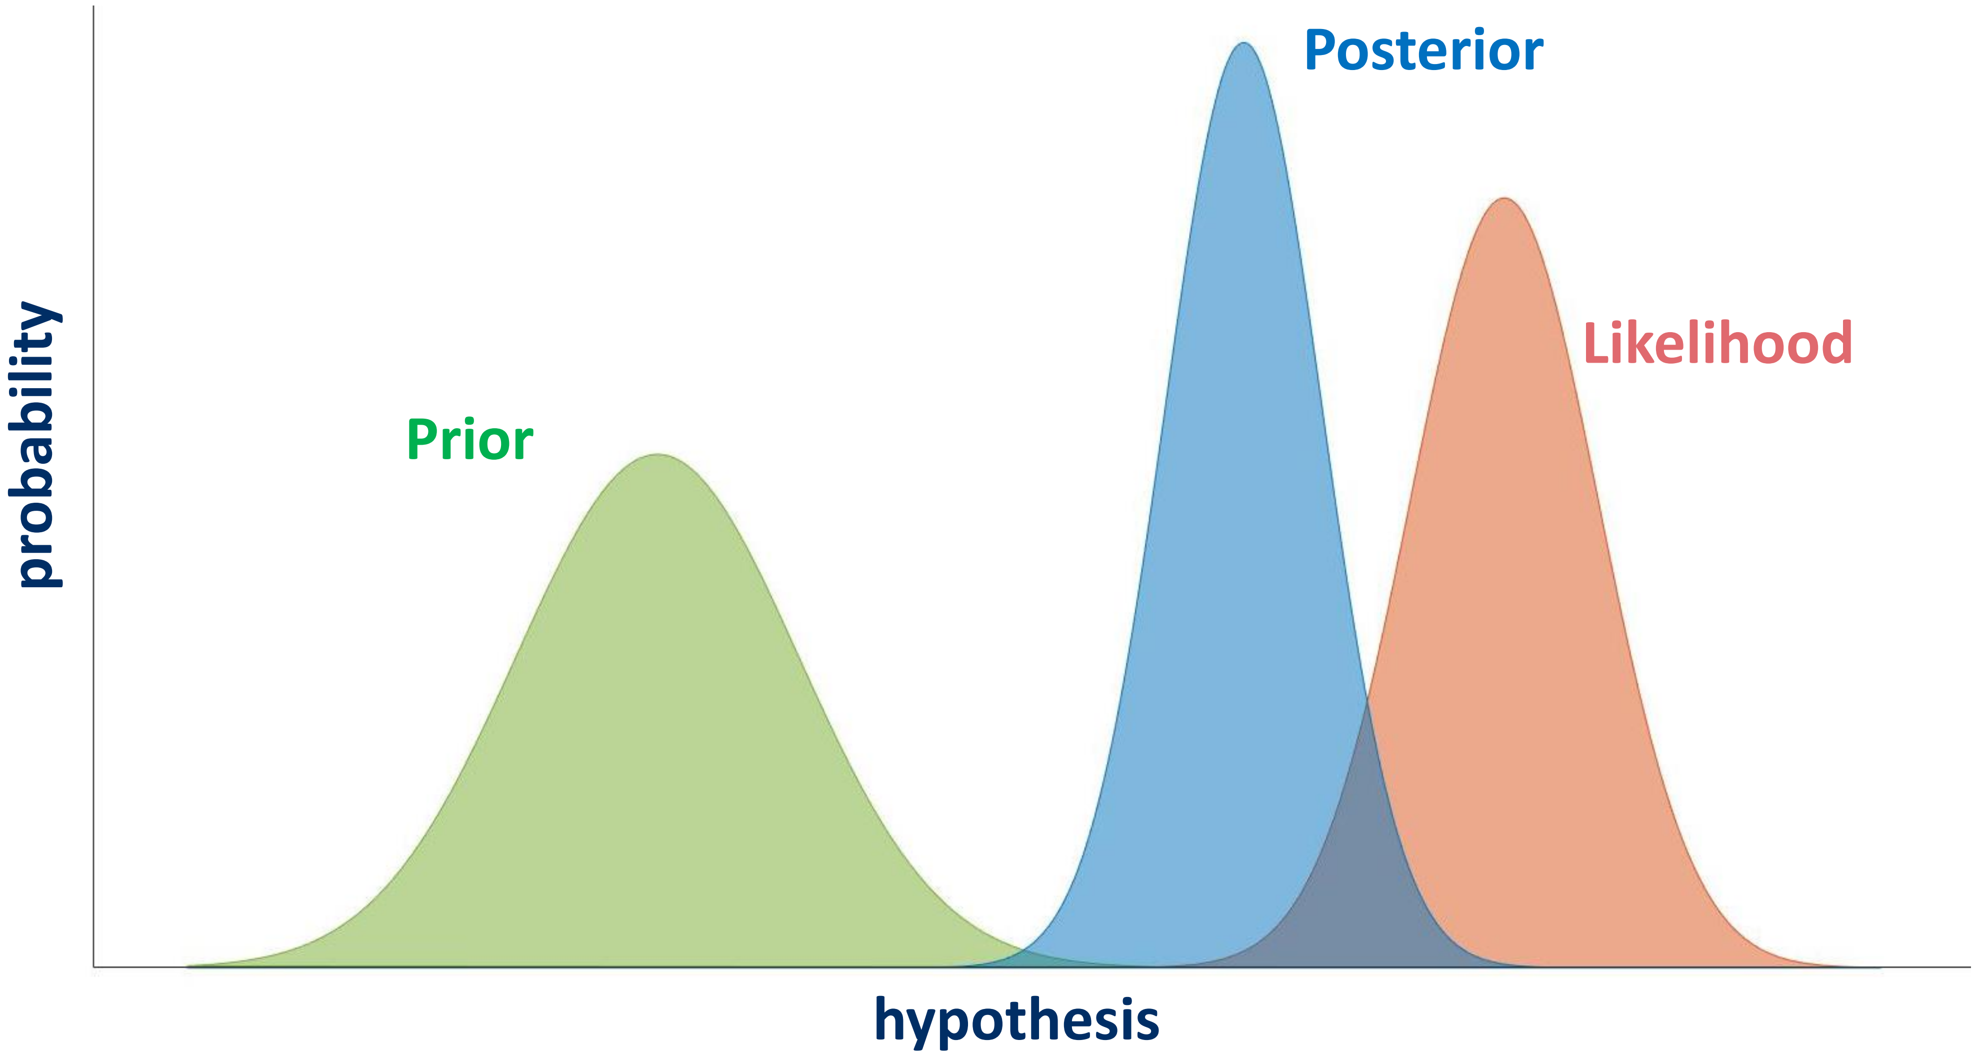

Prior, likelihood & posterior

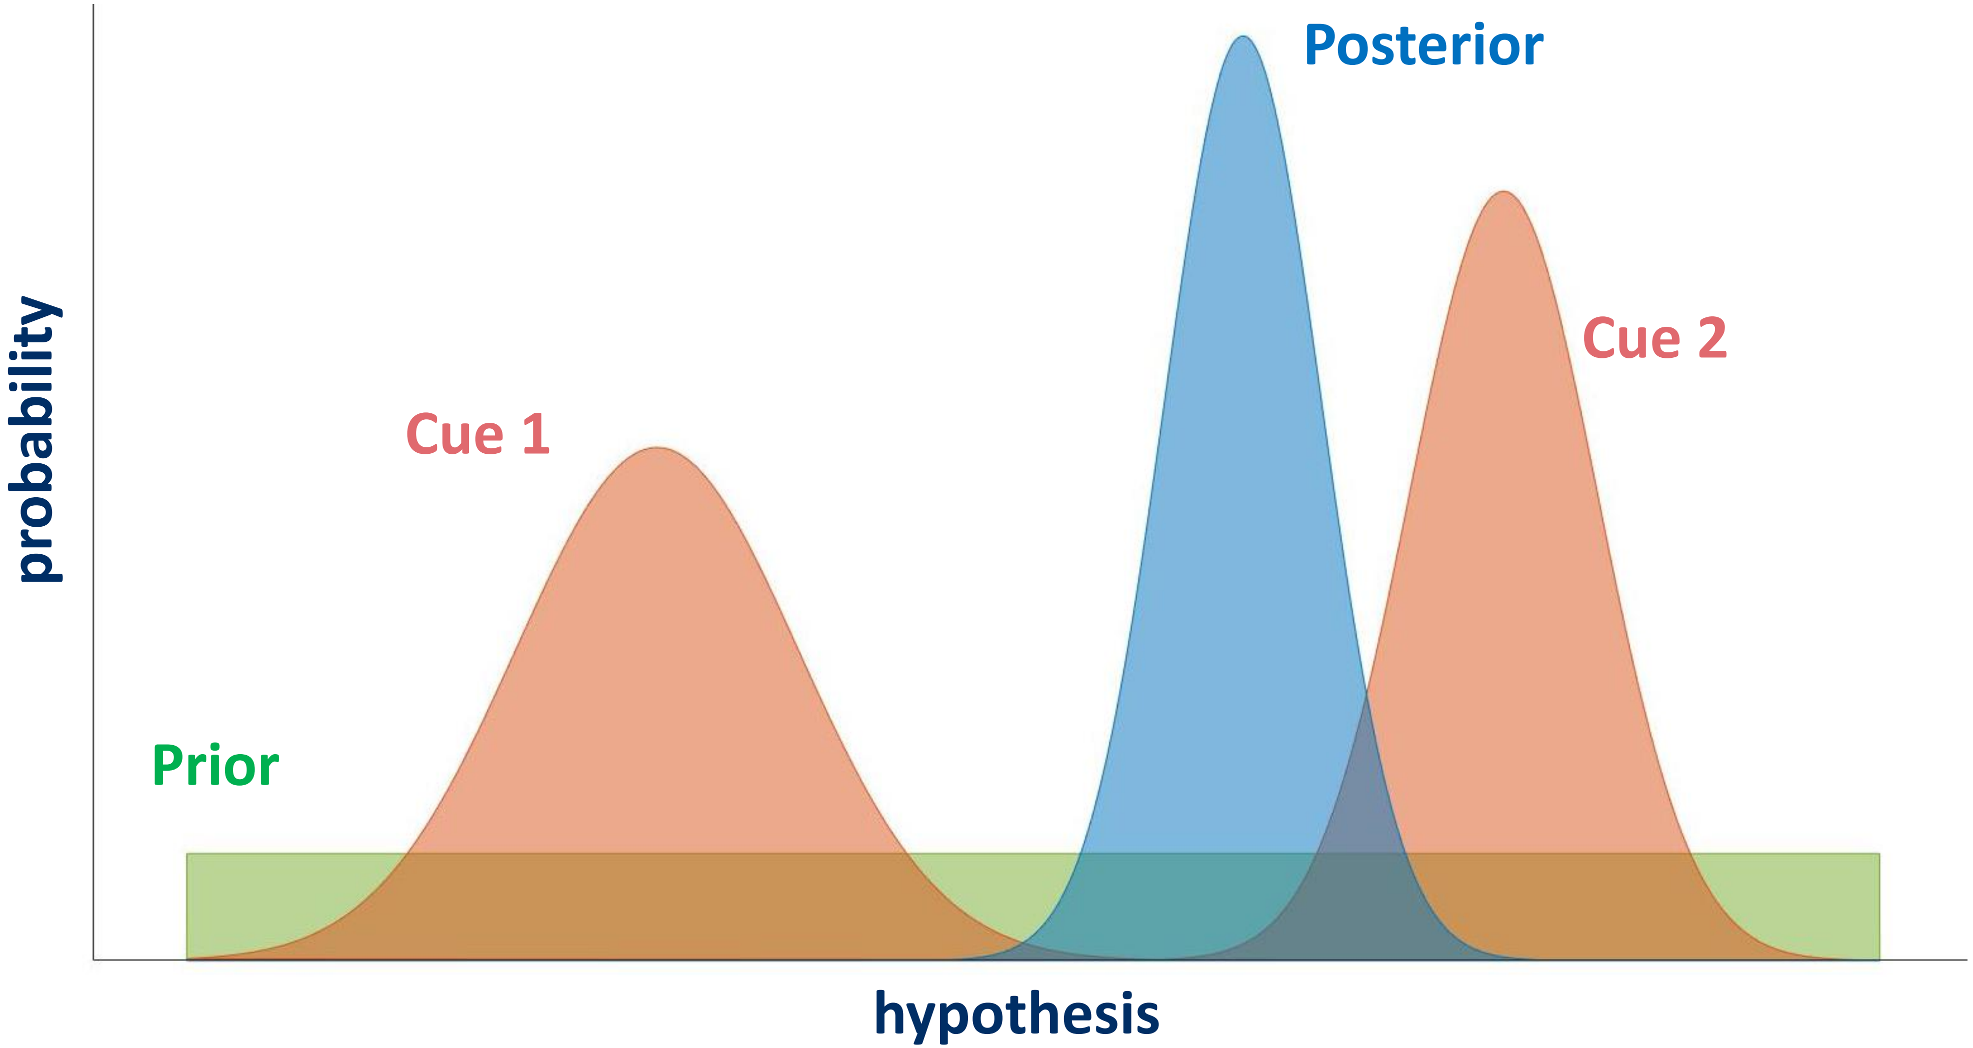

Loss functions and decision making

- We have a posterior distribution, now we need to make a point estimate
- Loss function: defines the consequences of the action
  - E.g., for sound localisation the loss is a function of estimation error
- Minimising the estimation error is equivalent with maximum a-posteriori (MAP) estimation, i.e., the mode of the posterior.

Loss Functions and Decision Making

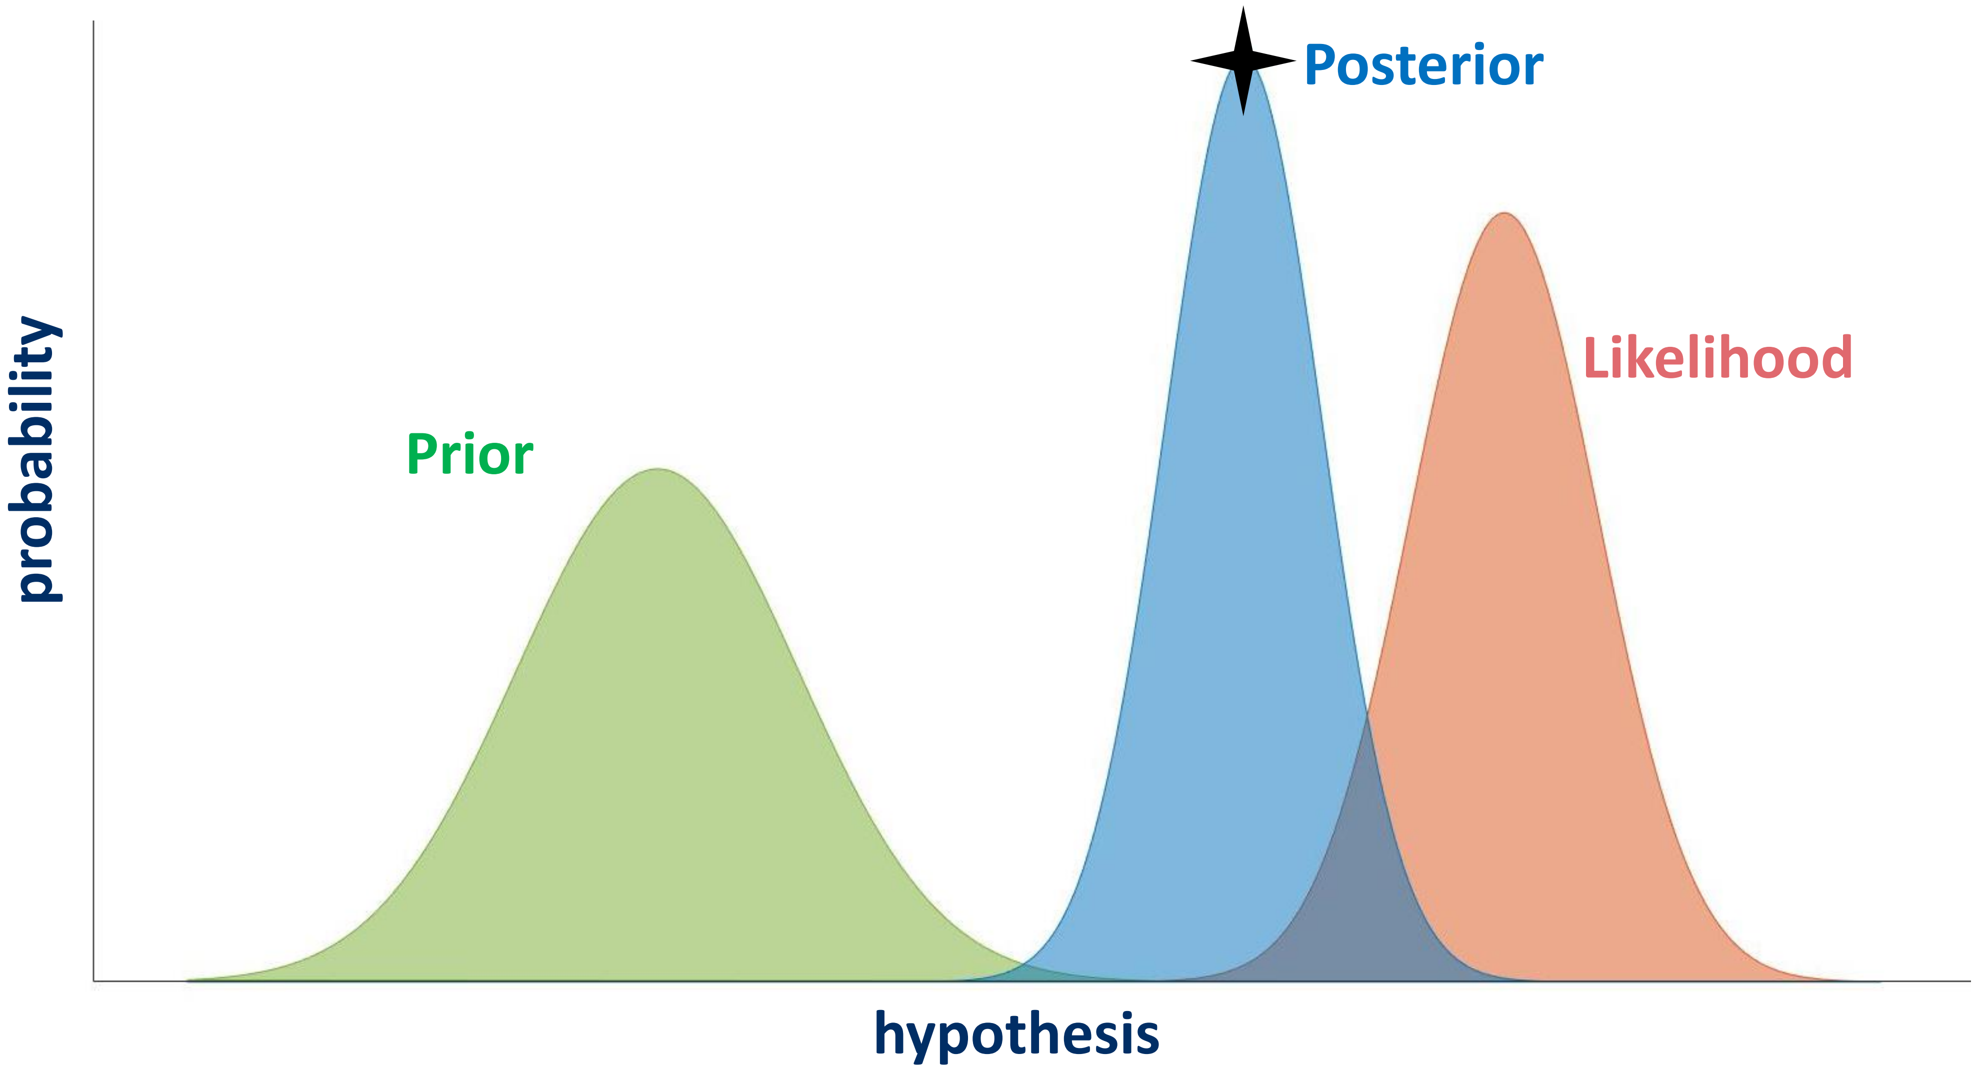

Benefits of Bayesian inference

- Encodes probability distributions, not just single value estimates
- Incorporates prior knowledge
- Modular: across and within sensory modalities
  - Combine multisensory cues, e.g., auditory, visual and sensorimotor
- Integrate over time -> iterative updates as more data comes in

Sound Localisation Example

Sound Localisation Example

Sound source localisation: what is the sound source direction?

- Step 1: Formulate a prior distribution
- Step 2: Derive likelihood from the data (e.g. through a sensor model)
- Step 3: Form posterior by combining prior and likelihood
- Step 4: Posterior becomes new prior

$$P(\psi|y) \propto p(y|\psi)p(\psi)$$

ψ: source direction, y: sensory data (ITD)

Sound Localisation Example

Sound source localisation: what is the sound source direction?

- Step 1: Formulate a prior distribution $p(\psi)$
  - Listener assumes sources to come from the front

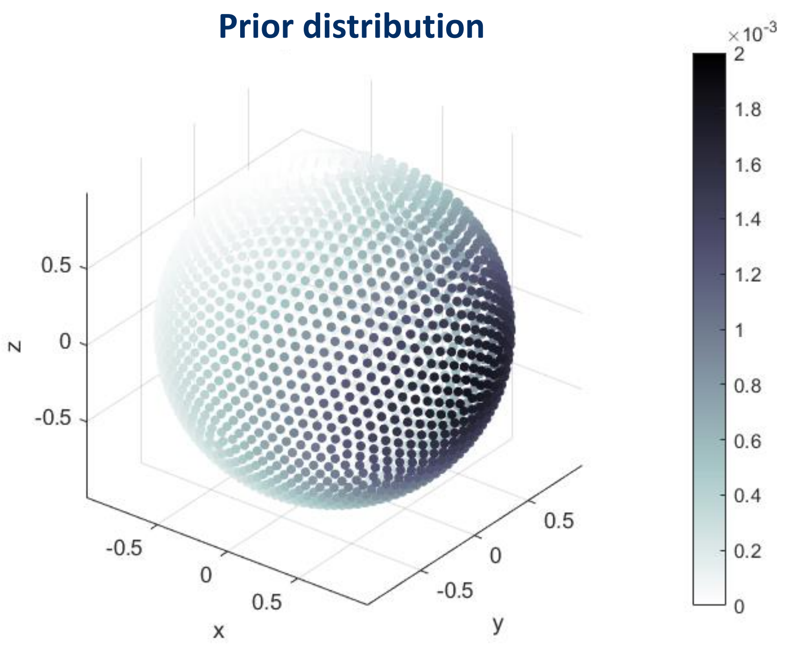

University of Antwerp

Sound Localisation Example

Sound source localisation: what is the sound source direction?

- Step 2: Form likelihood $p(y|\psi)$ from a sensor model and acoustic cues
  - Binaural cues give us a cone of confusion

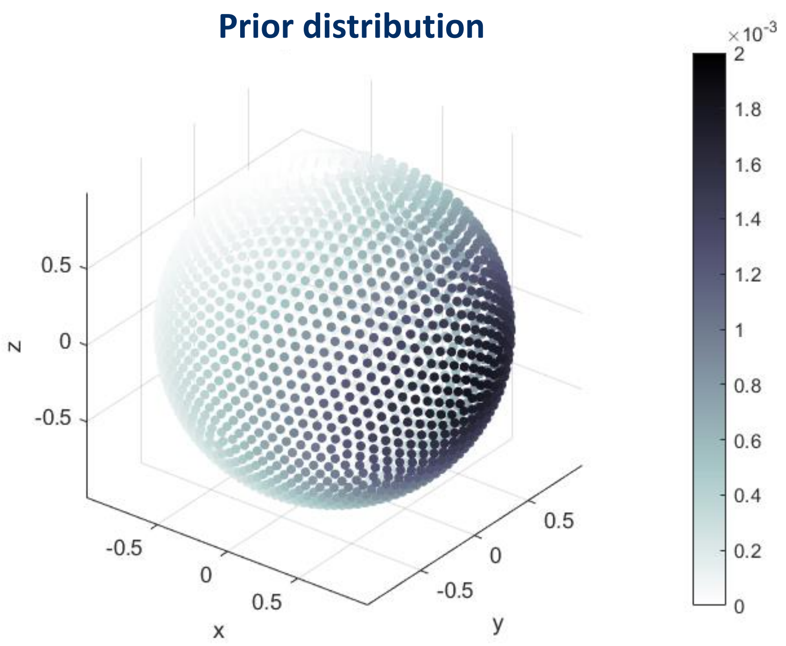

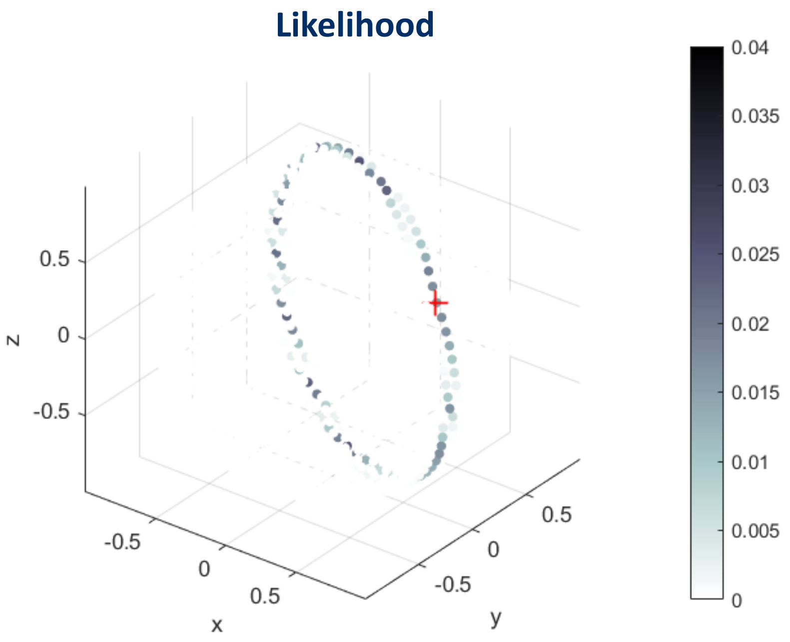

University of Antwerp

Sound Localisation Example

Sound source localisation: what is the sound source direction?

- Step 3: Form posterior by combining prior and likelihood
  - Decreased uncertainty

Prior distribution

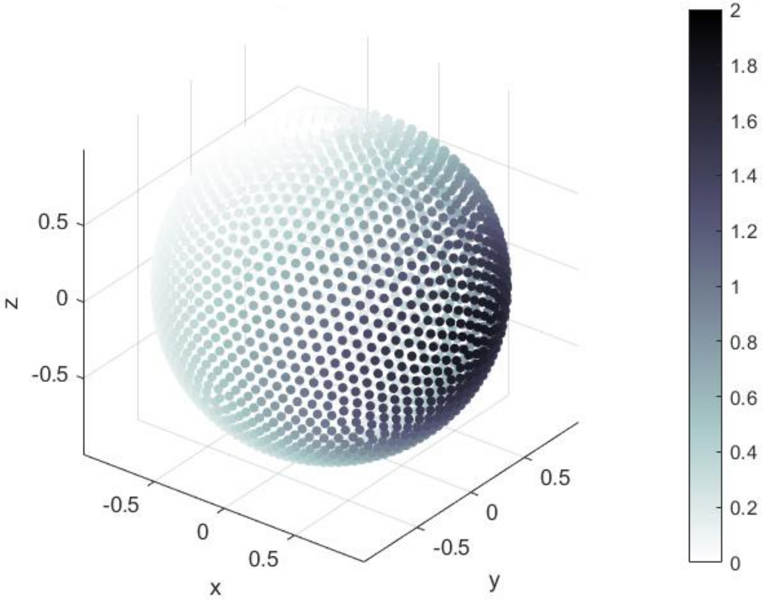

Likelihood

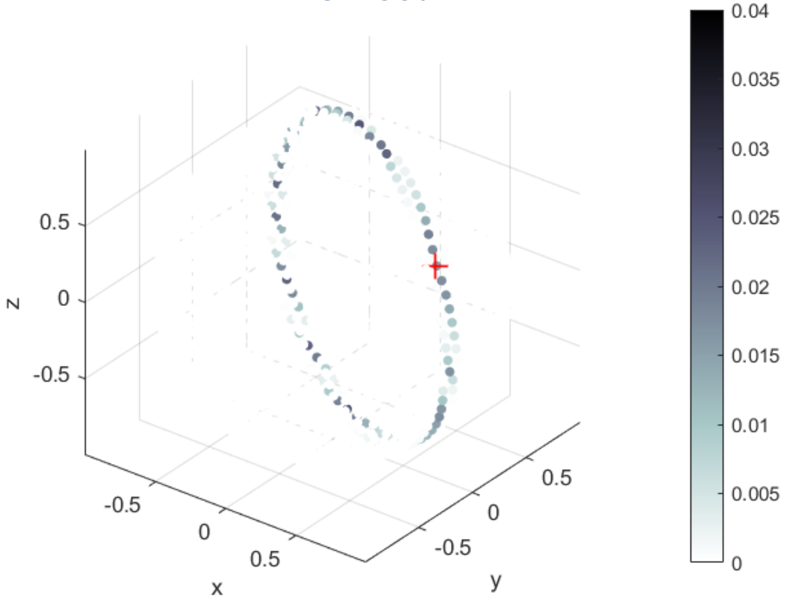

Posterior PDF

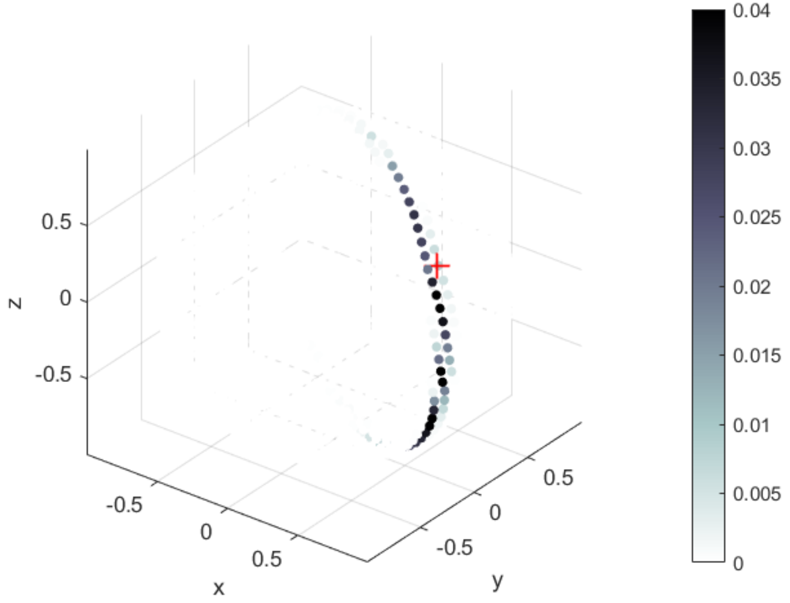

University of Antwerp

Sound Localisation Example

Sound source localisation: what is the sound source direction?

- Step 4: Posterior becomes new prior -> repeat!
  - As new information becomes available we can keep updating our posterior

Prior distribution

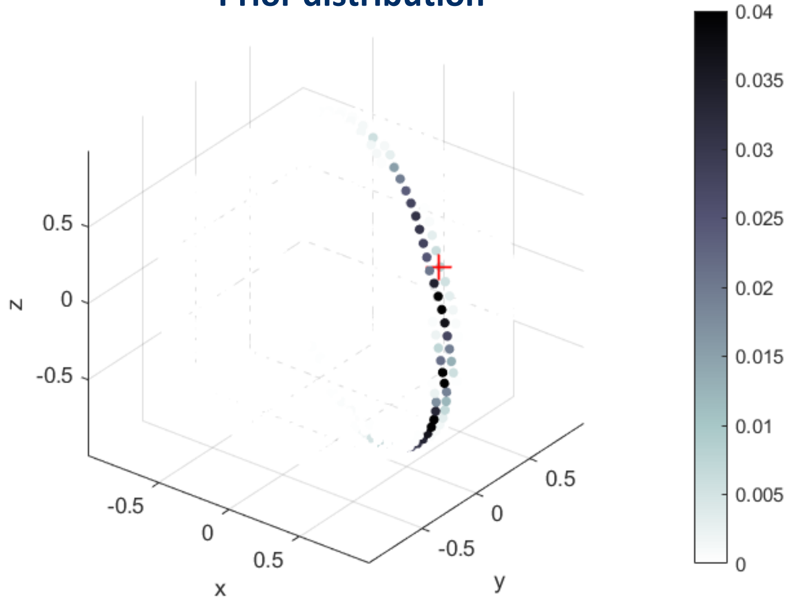

. . .

Model for Active Sound Localisation

Model for Active Sound Localisation

$$p_{t_i} = C \cdot p(y_A(t_i)|\theta_H(t_i), \psi) \cdot p_{t_{i-1}}$$

y_A = observer acoustic information
θ_H = true head orientation

Bayes' Theorem: $$p(\psi|X) = \frac{p(X|\psi)p(\psi)}{p(X)} \propto p(X|\psi)p(\psi)$$

p_{t_i} = posterior
$p_{t_{i-1}}$ = prior
$p(y_A(t_i)|\theta_H(t_i), \psi)$ = sensor model, i.e., likelihood
C = normalisation constant

Generative model (likelihood)

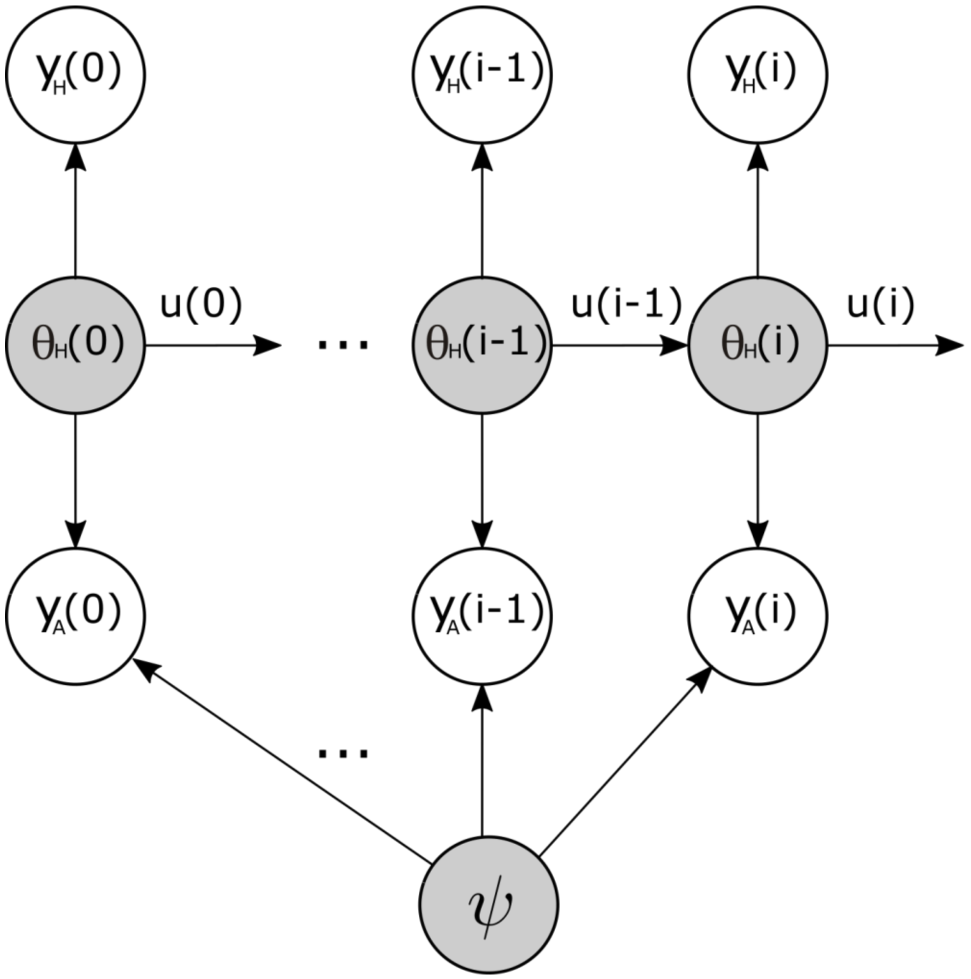

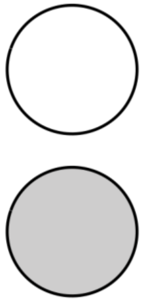

= observed variables

= true variables

$y_H(t_i)$ = obs. head orientation
$y_A(t_i)$ = obs. acoustic information
$u(t_i)$ = motor input
θ = true head orientation
ψ = true sound source direction

Model for Active Sound Localisation

$$p_{t_i} = C \cdot p_{t_{i-1}} \cdot \int_{\theta_H} p(y_A(t_i)|\theta_H(t_i),\ \psi) \cdot p(\theta_H(t_i)|y_H(t_0\!:t_i),\ u(t_0\!:t_{i-1}))d\theta_H$$

p_{t_i}	= posterior
$p_{t_{i-1}}$	= prior
$p(y_A(t_i)\|\theta_H(t_i), \psi)$	= acoustic sensor model
$p(\theta_H(t_i)\|y_H(t_0\!:t_i),\ u(t_0\!:t_{i-1}))$	= motor sensor model
C	= normalisation constant

Now let’s have a look at the code