Inferring Cognitive Models from Data using Approximate Bayesian Computation

CHI 2017

Antti Kangasrääsiö[1], Kumaripaba Athukorala[1], Andrew Howes[2], Jukka Corander[3], Samuel Kaski[1], Antti Oulasvirta[4]

[1]Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University, Finland
[2]School of Computer Science, University of Birmingham, UK
[3]Department of Biostatistics, University of Oslo, Norway
[4]Helsinki Institute for Information Technology HIIT, Department of Communications and Networking, Aalto University, Finland

Tuesday 9.5.2017

Main Contribution

Parameters of cognitive models can be tuned
based on commonly available observation data
using Approximate Bayesian Computation (ABC)

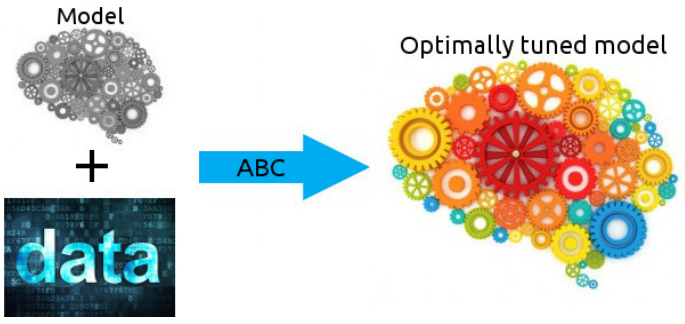

Examples of Cognitive Models

Examples of Cognitive Models

Early work:
Fitt's law (1954)
Speed/accuracy trade-off
in pointing tasks

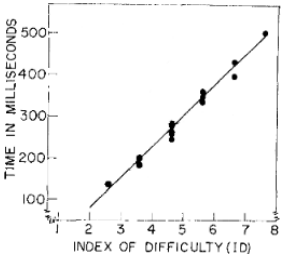

Examples of Cognitive Models

Early work:
Fitt's law (1954)
Speed/accuracy trade-off
in pointing tasks

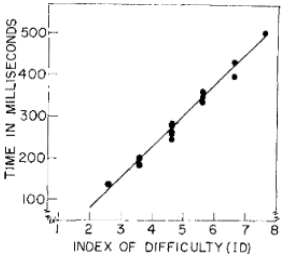

Modern work:
ACT-R (1997)
Cognitive tasks with
visual elements

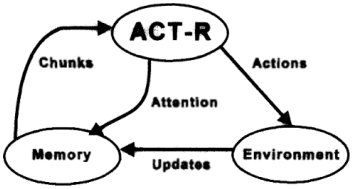

Examples of Cognitive Models

Early work:
Fitt's law (1954)
Speed/accuracy trade-off
in pointing tasks

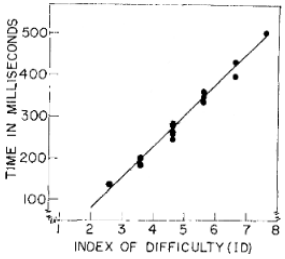

Modern work:
ACT-R (1997)
Cognitive tasks with
visual elements

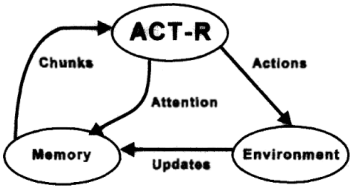

Recent work:
Menu search (CHI 2015)
Visual search in menus

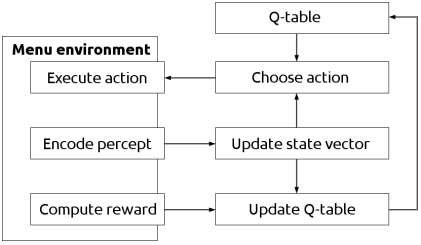

Examples of Cognitive Models

Early work:
Fitt's law (1954)
Speed/accuracy trade-off
in pointing tasks

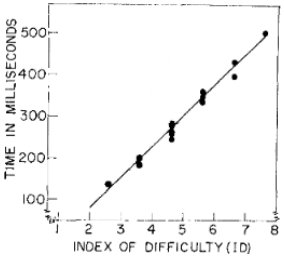

Modern work:
ACT-R (1997)
Cognitive tasks with
visual elements

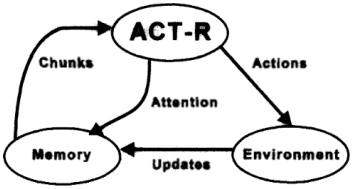

Recent work:
Menu search (CHI 2015)
Visual search in menus

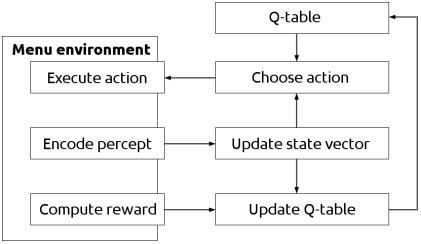

Yesterday: A model of visual decision-making by Chen, Starke, Baber & Howes

Why Model Human Cognition

Why Model Human Cognition

Predicting user behavior

Why Model Human Cognition

Predicting user behavior

Explaining cognitive processes

Why Model Human Cognition

Predicting user behavior

Explaining cognitive processes

Allowing computer systems to adapt better to the user

Model Tuning

Model Tuning

Goal: select parameters that lead to accurate predictions

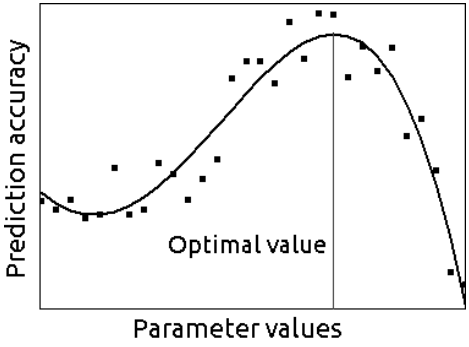

Model Tuning

Goal: select parameters that lead to accurate predictions

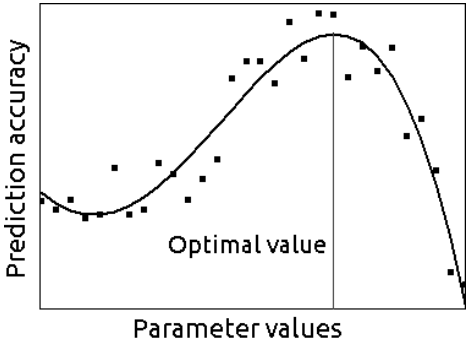

Examples of parameters

Model Tuning

Goal: select parameters that lead to accurate predictions

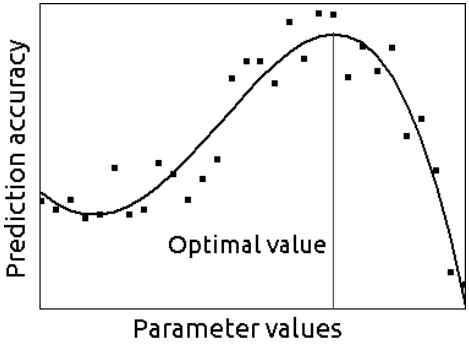

Examples of parameters

- Duration of a visual fixation

Model Tuning

Goal: select parameters that lead to accurate predictions

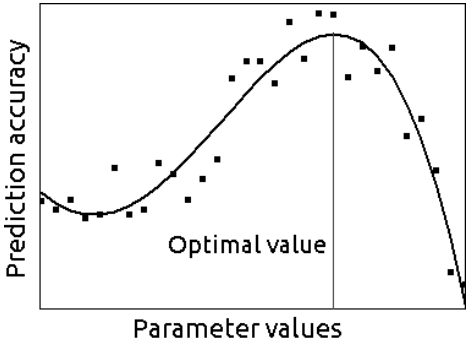

Examples of parameters

- Duration of a visual fixation
- Acuity and size of peripheral visual field

Model Tuning

Goal: select parameters that lead to accurate predictions

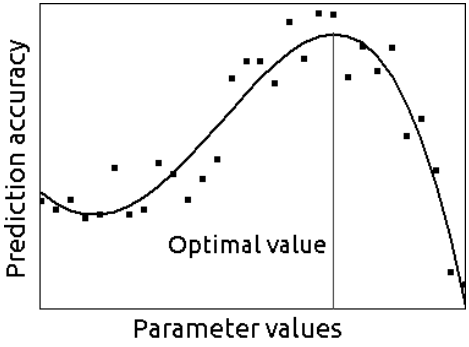

Examples of parameters

- Duration of a visual fixation
- Acuity and size of peripheral visual field
- Reward for completing the task

Model Tuning

Goal: select parameters that lead to accurate predictions

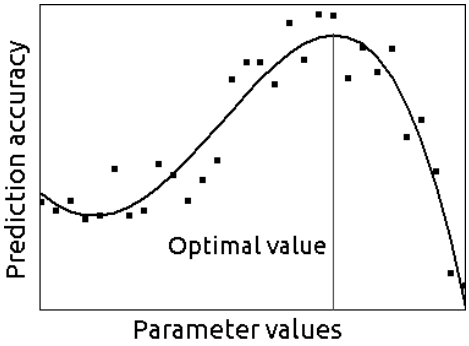

Examples of parameters

- Duration of a visual fixation
- Acuity and size of peripheral visual field
- Reward for completing the task
- Penalty for wasting one second of time

Model Tuning

Goal: select parameters that lead to accurate predictions

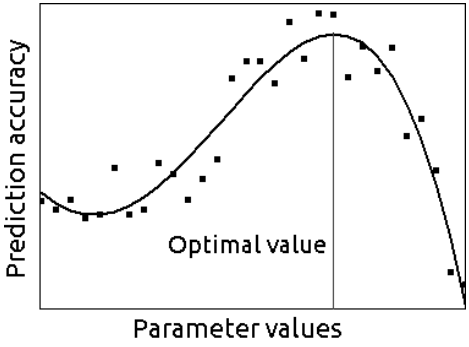

Examples of parameters

- Duration of a visual fixation
- Acuity and size of peripheral visual field
- Reward for completing the task
- Penalty for wasting one second of time
- Duration of short-term memory

Model Tuning

Goal: select parameters that lead to accurate predictions

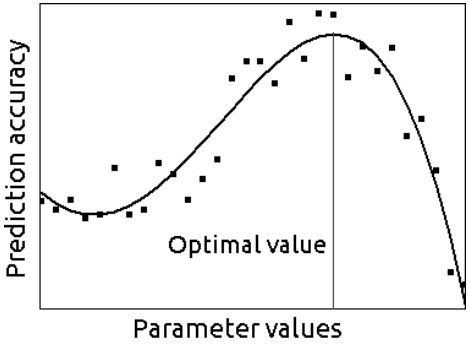

Examples of parameters

- Duration of a visual fixation
- Acuity and size of peripheral visual field
- Reward for completing the task
- Penalty for wasting one second of time
- Duration of short-term memory
- And so on ..

Characteristics of Modern Cognitive Models

Characteristics of Modern Cognitive Models

No simple way to tune parameters

Highly non-linear model structure, not possible to directly compute best parameter values

Characteristics of Modern Cognitive Models

No simple way to tune parameters

Highly non-linear model structure, not possible to directly compute best parameter values

Computing predictions takes time

Once we fix the parameter values, the model needs to first learn a behavior strategy

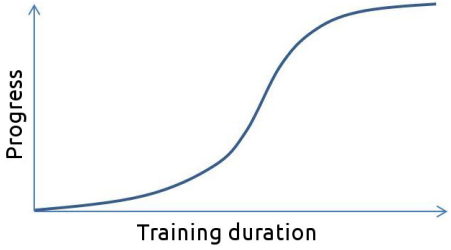

Traditional Tuning Methods

Traditional Tuning Methods

Manual Parameter Tuning

Slow: needs human labour
No guarantees of optimality

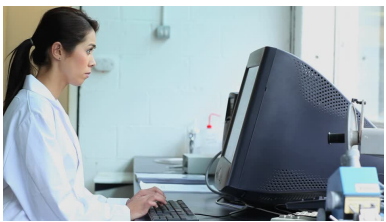

Traditional Tuning Methods

Manual Parameter Tuning

Slow: needs human labour
No guarantees of optimality

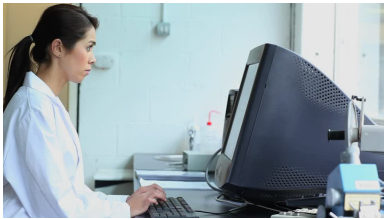

Grid search

Might take a very long time
No uncertainty estimates

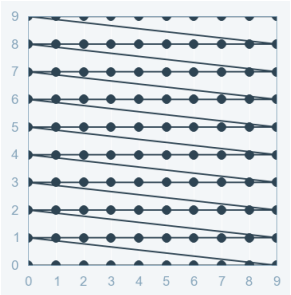

Approximate Bayesian Computation (ABC)

ABC Inference Process

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space
- (Repeat until converged)

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space
- (Repeat until converged)

Inference Results

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space
- (Repeat until converged)

Inference Results

- Based on the probabilistic model, we can compute an approximate probability distribution over the parameter space

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space
- (Repeat until converged)

Inference Results

- Based on the probabilistic model, we can compute an approximate probability distribution over the parameter space
- Maximum of the distribution gives us the Maximum Likelihood (ML) estimate

Approximate Bayesian Computation (ABC)

ABC Inference Process

- Choose parameter values for the model
- Simulate predictions
- Evaluate discrepancy between predictions and observations
- Use a probabilistic model to estimate discrepancy in different regions of parameter space
- (Repeat until converged)

Inference Results

- Based on the probabilistic model, we can compute an approximate probability distribution over the parameter space
- Maximum of the distribution gives us the Maximum Likelihood (ML) estimate
- Width of distribution corresponds to uncertainty

Approximate Bayesian Computation (ABC)

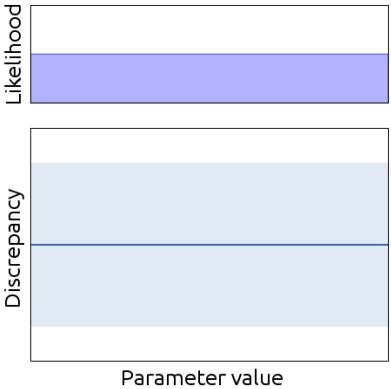

Model Predictions

Discrepancy

Observation Data

Approximate Bayesian Computation (ABC)

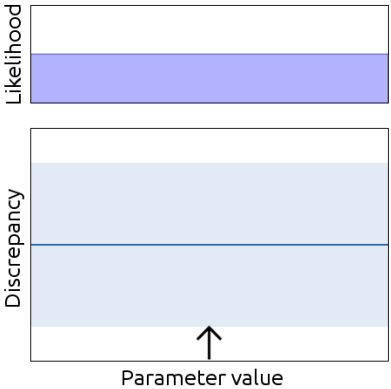

Model Predictions

Discrepancy

Observation Data

Approximate Bayesian Computation (ABC)

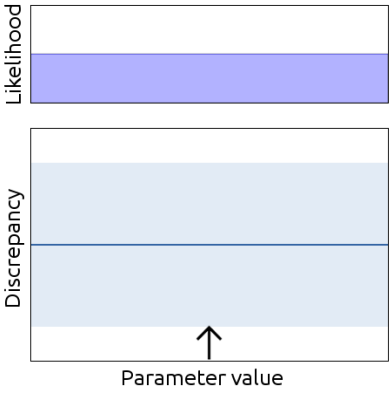

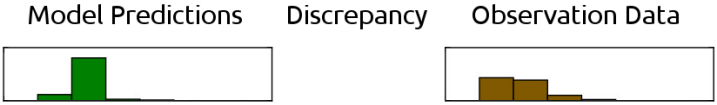

Approximate Bayesian Computation (ABC)

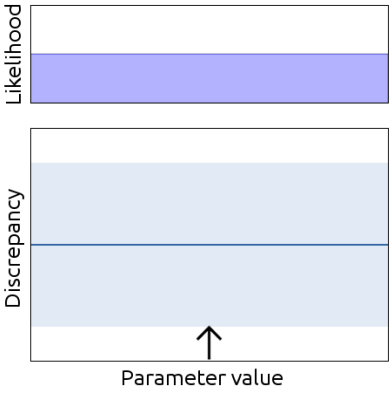

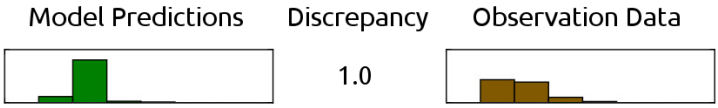

Approximate Bayesian Computation (ABC)

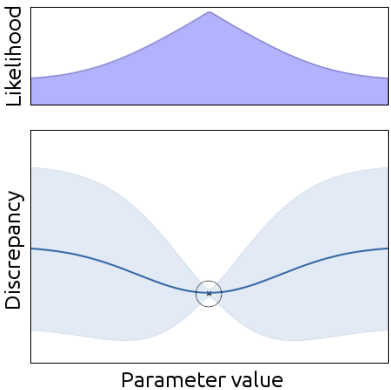

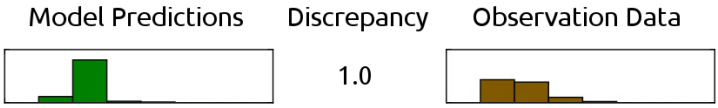

Approximate Bayesian Computation (ABC)

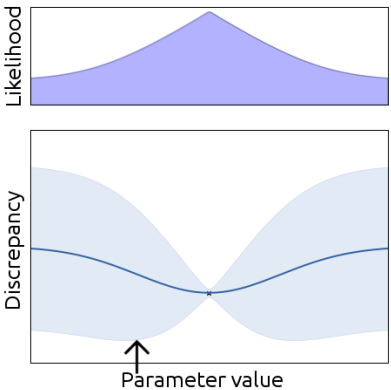

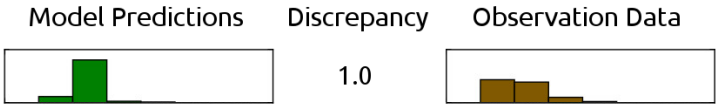

Approximate Bayesian Computation (ABC)

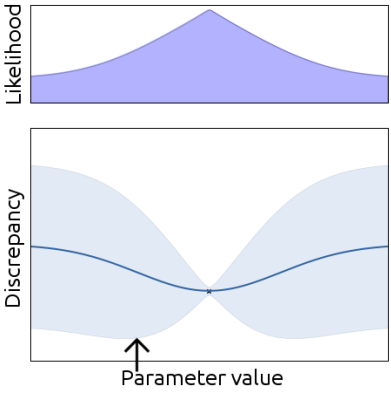

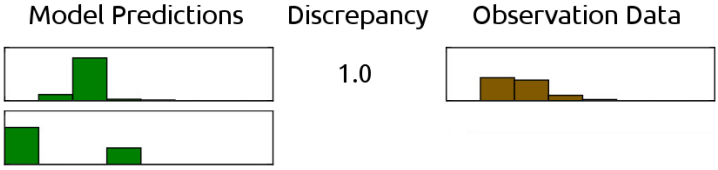

Approximate Bayesian Computation (ABC)

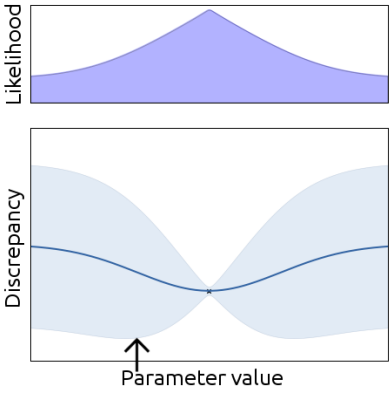

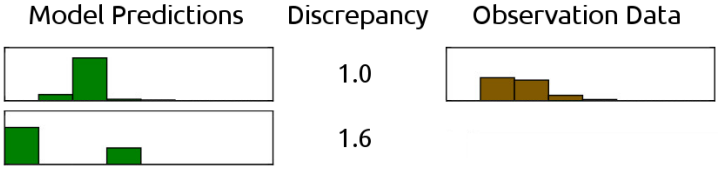

Approximate Bayesian Computation (ABC)

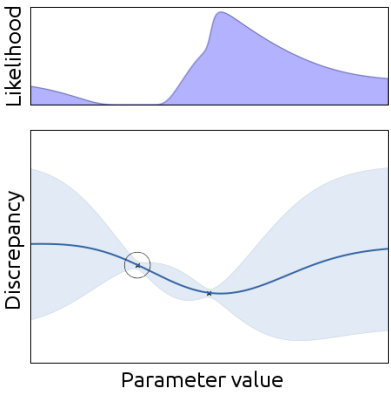

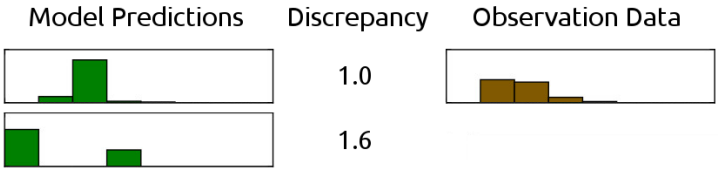

Approximate Bayesian Computation (ABC)

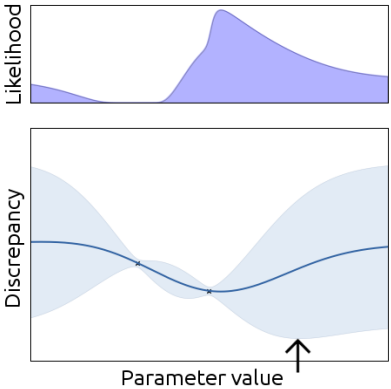

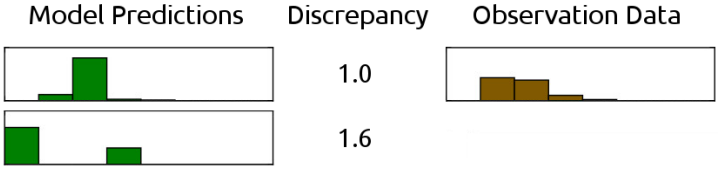

Approximate Bayesian Computation (ABC)

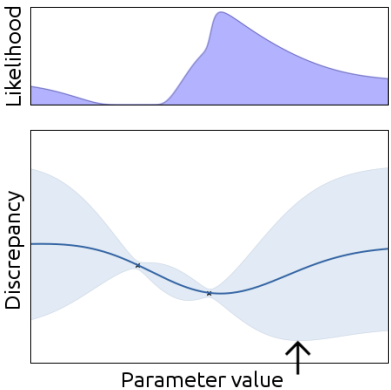

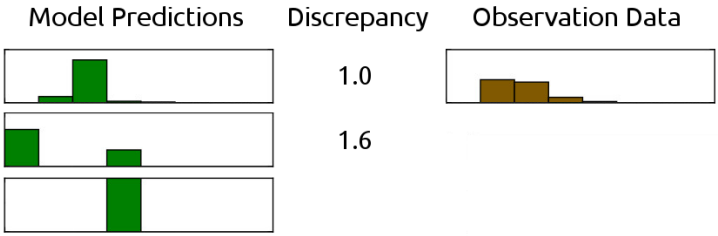

Approximate Bayesian Computation (ABC)

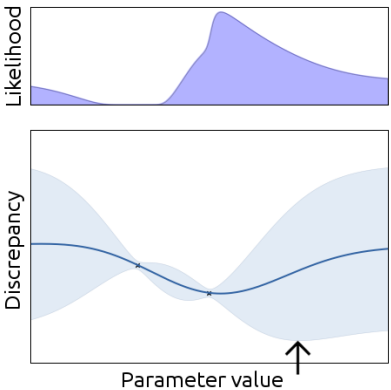

Approximate Bayesian Computation (ABC)

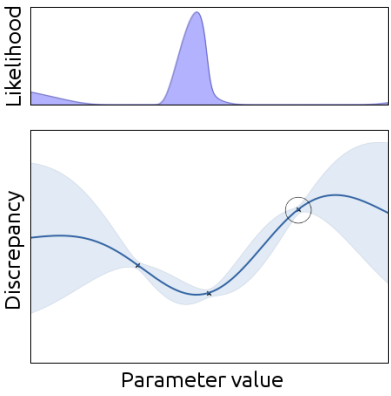

Approximate Bayesian Computation (ABC)

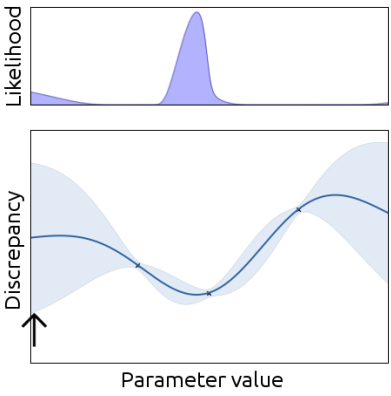

Approximate Bayesian Computation (ABC)

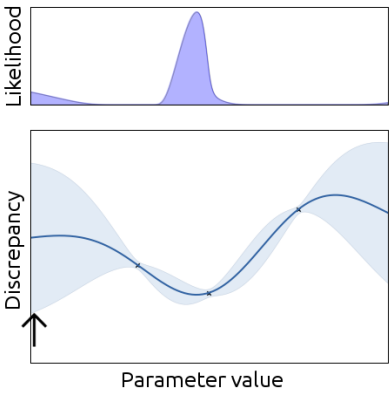

Approximate Bayesian Computation (ABC)

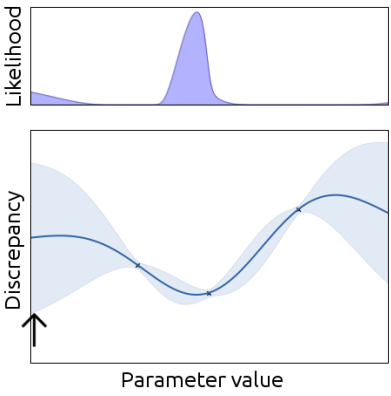

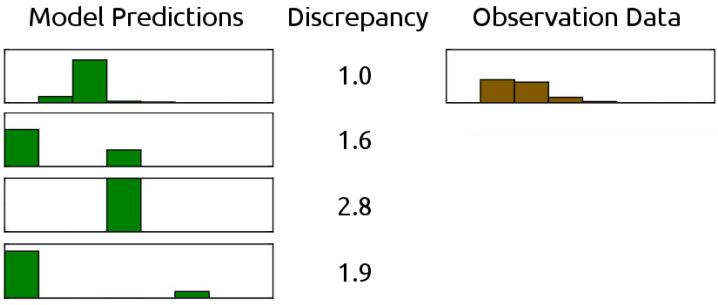

Approximate Bayesian Computation (ABC)

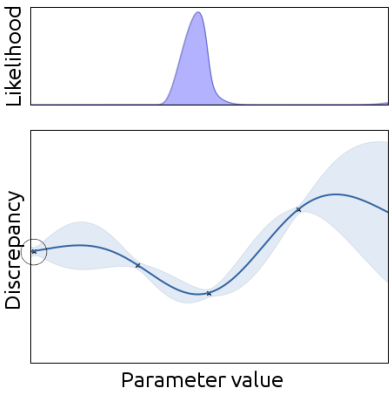

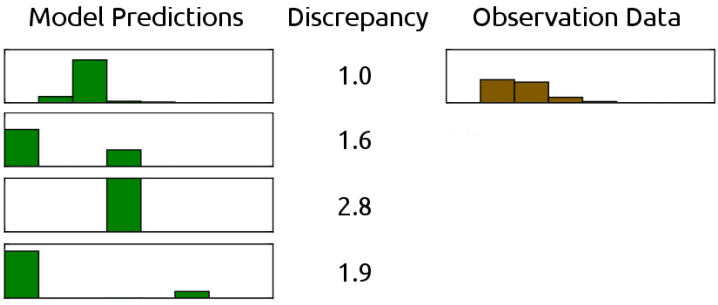

Approximate Bayesian Computation (ABC)

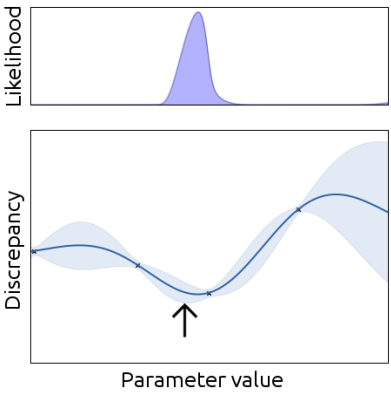

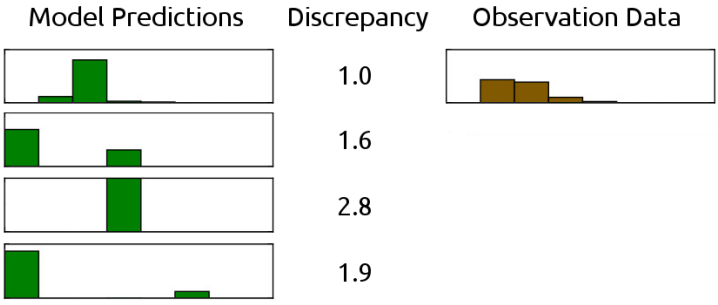

Approximate Bayesian Computation (ABC)

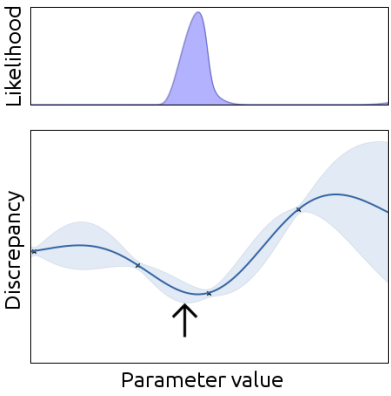

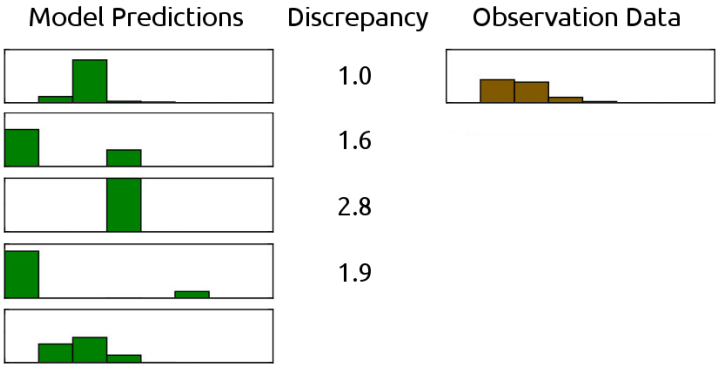

Approximate Bayesian Computation (ABC)

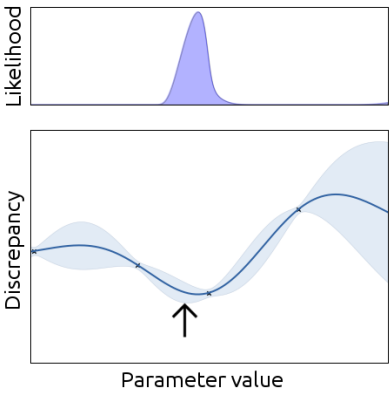

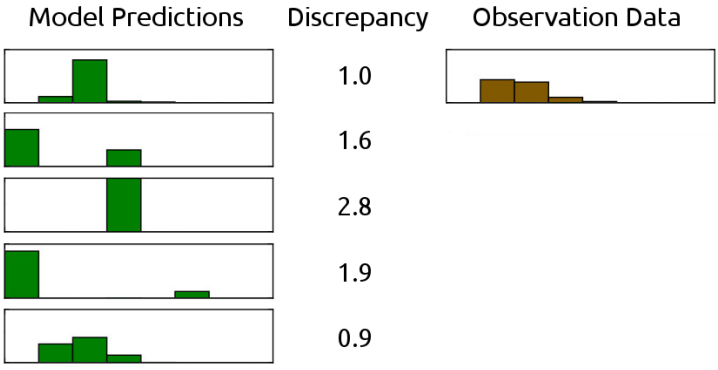

Approximate Bayesian Computation (ABC)

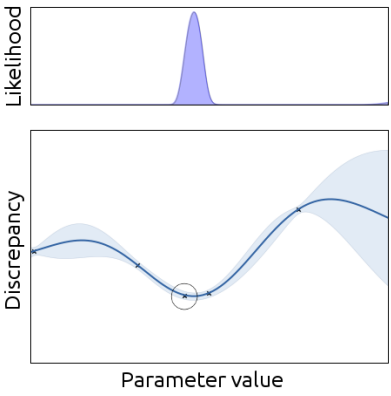

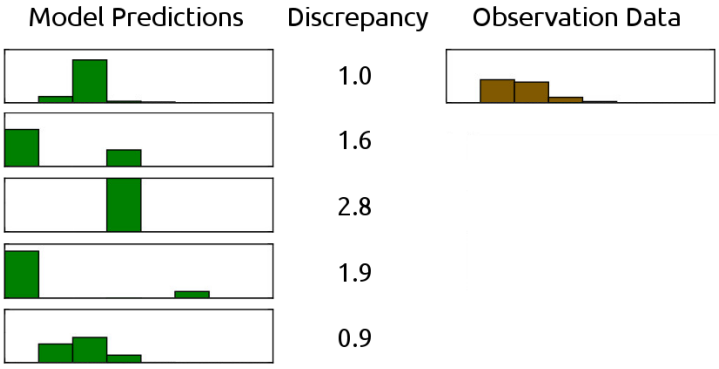

Approximate Bayesian Computation (ABC)

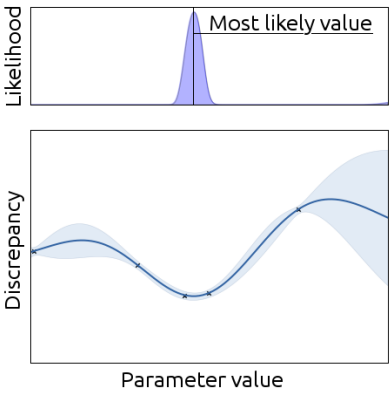

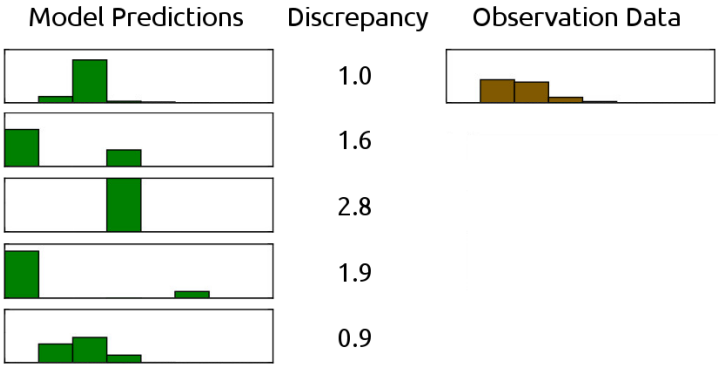

Approximate Bayesian Computation (ABC)

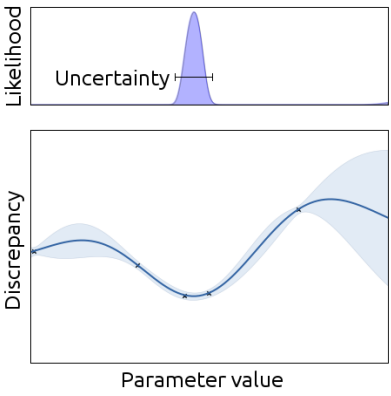

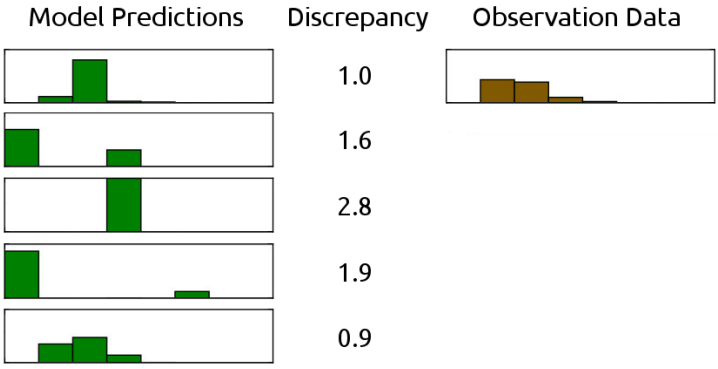

Approximate Bayesian Computation (ABC)

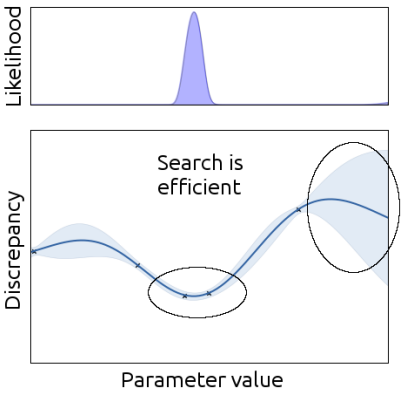

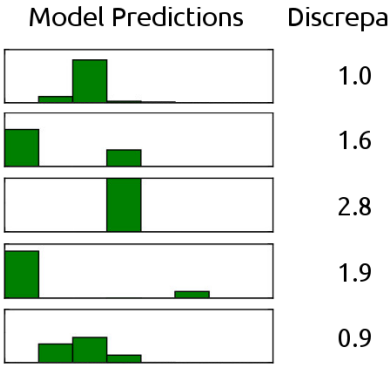

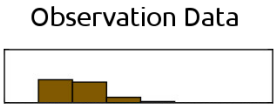

Approximate Bayesian Computation (ABC)

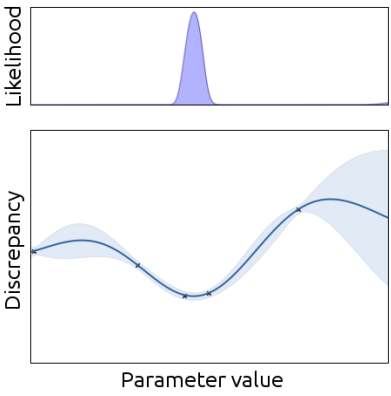

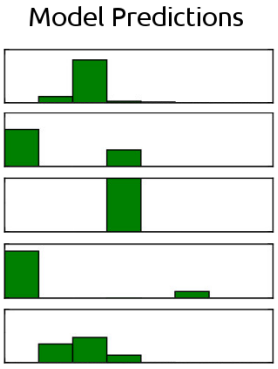

Discrepancy

1.0

1.6

2.8

1.9

0.9

Observation Data

No special requirements for observation data

Case Study: Visual Search

Case Study: Visual Search

A recent model for visual search in drop-down menus (Chen *et al.* CHI 2015)

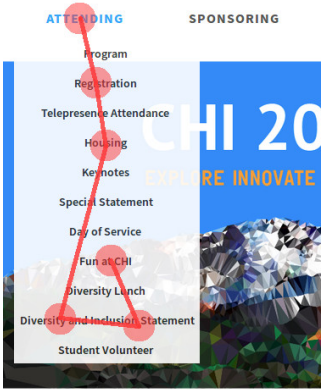

Case Study: Visual Search

A recent model for visual search in drop-down menus (Chen *et al.* CHI 2015)

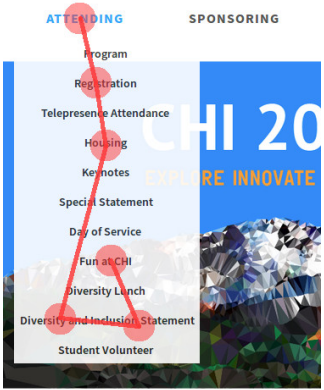

With ABC, the predictions match better to observations

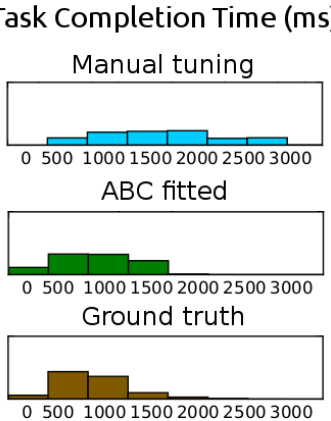

Model Development and Individual Modelling with ABC

Fit different models to same dataset

This allowed us to estimate how different changes to the model affected the predictions

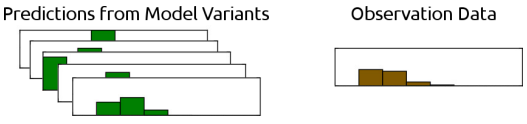

Model Development and Individual Modelling with ABC

Fit different models to same dataset

This allowed us to estimate how different changes to the model affected the predictions

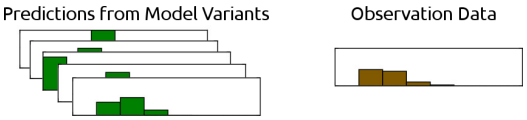

Fit same model to different datasets

By fitting models to observations from individual users we were able to improve the predictions over a population level model

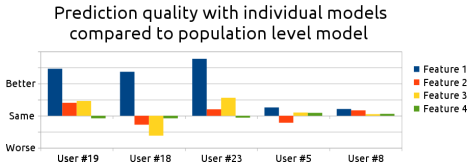

Resources for Practical Use of ABC

Resources for Practical Use of ABC

Requirements for ABC

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model
  (with tunable parameters)

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model
  (with tunable parameters)
- Prior knowledge of
  reasonable parameter values

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model
  (with tunable parameters)
- Prior knowledge of
  reasonable parameter values
- Observation dataset

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model
  (with tunable parameters)
- Prior knowledge of
  reasonable parameter values
- Observation dataset
- Discrepancy measure

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model (with tunable parameters)
- Prior knowledge of reasonable parameter values
- Observation dataset
- Discrepancy measure

Technical Considerations

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model (with tunable parameters)
- Prior knowledge of reasonable parameter values
- Observation dataset
- Discrepancy measure

Technical Considerations

- There exists programming frameworks for ABC (e.g. ELFI: github.com/elfi-dev/elfi)

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model (with tunable parameters)
- Prior knowledge of reasonable parameter values
- Observation dataset
- Discrepancy measure

Technical Considerations

- There exists programming frameworks for ABC (e.g. ELFI: github.com/elfi-dev/elfi)
- The more parameters you infer, the more challenging inference is

Resources for Practical Use of ABC

Requirements for ABC

- A simulator model (with tunable parameters)
- Prior knowledge of reasonable parameter values
- Observation dataset
- Discrepancy measure

Technical Considerations

- There exists programming frameworks for ABC (e.g. ELFI: github.com/elfi-dev/elfi)
- The more parameters you infer, the more challenging inference is
- Cluster computers offer speedups

Take-home Messages

Tuning the parameters of cognitive models based on real observations is important for making good predictions

Approximate Bayesian Computation (ABC) allows you to do this with only very mild technical requirements

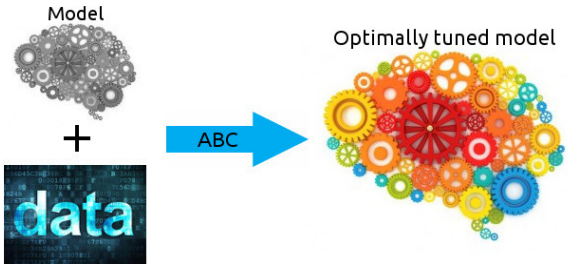

ELFI, a Python library for ABC: github.com/elfi-dev/elfi

Tech paper: Inverse Reinforcement Learning from Summary Data: arxiv.org/abs/1703.09700