Bridging Human and Machine Learning: Using discrete Markov Chain Monte Carlo with People to explore human categories

Anne Hsu

Computer Science

University College London

Mind’s challenge: to infer complex structure from limited, noisy, ambiguous input

e.g.

-Infer 3D representations from 2D retina projection

-Recognize complex sounds from pressure waves

-Learn language from hearing limited sentences allowed in language

-Learn categories from limited examples

Mind's challenge: to infer complex structure from limited, noisy, ambiguous input

Uses of Bayesian probability theory and machine learning methods in cognitive science

- **Benchmark**, posing ideal rational solution which is then compared with psychological behaviour to find underlying principles of how the mind works

- **Model** of psychological processes

- **Tool** for measuring mental representations

Mind's challenge: to infer complex structure from limited, noisy, ambiguous input

Uses of Bayesian probability theory and machine learning methods in cognitive science

- **Benchmark**, posing ideal rational solution which is then compared with psychological behaviour to find underlying principles of how the mind works
- **Model** of psychological processes
- **Tool** for measuring mental representations

Mind's challenge: to infer complex structure from limited, noisy, ambiguous input

Uses of Bayesian probability theory and machine learning methods in cognitive science

- **Benchmark**, posing ideal rational solution which is then compared with psychological behaviour to find underlying principles of how the mind works

- **Model** of psychological processes

- **Tool** for measuring mental representations

  structure of human categories

Human concepts and categories

- We form categories to make sense and find structure in overwhelmingly complex world

- Concept is a mental representation that picks out a set of entities, or a category.

- Every day, we are constantly referring to our categories, updating them, and creating new ones

- Categories can span over any concept in life
  - Healthy foods, activities
  - Family and friends vs. foes
  - Preferable activities, choices
  - Efficient travel modes

Human concepts and categories

- Background on category representation research

- Challenge: Difficult to measuring human categories in real life (high-dimensional, complex) domains

- New method of measuring people's representation of categories: Markov Chain Monte Carlo (MCMC)
  - sampling algorithm from computer science adapted to measure human representations of complex real life environments

Background

- How we learn and represent concepts & categories is major topic of psychological research

Background

- How we learn and represent concepts & categories is major topic of psychological research

- One assumes category membership depends on features of an item *x*. Membership is not rigid. P(*C*=*dog*|*x*)

x

x

Background

- How we learn and represent concepts & categories is major topic of psychological research

- One assumes category membership depends on features of an item *x*. Membership is not rigid. P(*C*=*dog*|*x*)

x

x

- Two popular models of category representation: *prototypes* and *exemplars*

Background

Models of category representation:

Prototype

(Medin & Schaffer, 1978; Nosofsky, 1986)

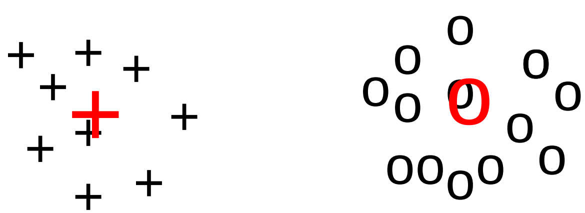

Exemplar

(Reed, 1972)

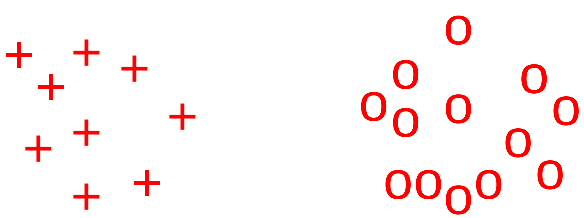

Both models can also be combined

(Anderson, 1990; Griffiths et al., 2008).

Background

Stimuli used in categorization research

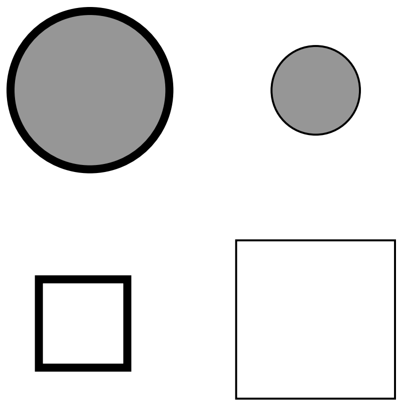

Shepard, Hovland, & Jenkins,1963

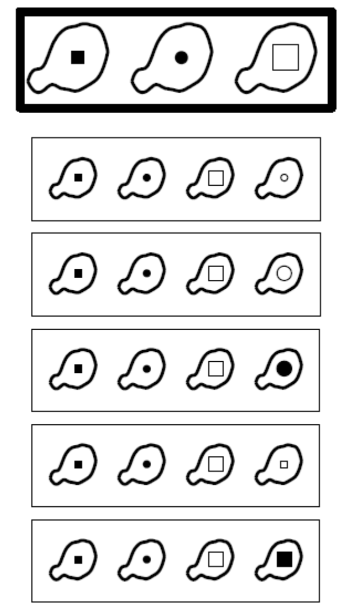

Griffiths et al., 2008
Feldman 2000

Challenge of measuring human categories

But what about real life concepts and categories?

- It is very difficult to measure the structure of human categories and concepts
- Real life items are complex and high dimensional

Challenge of measuring human categories

Features often cannot be interpreted independently

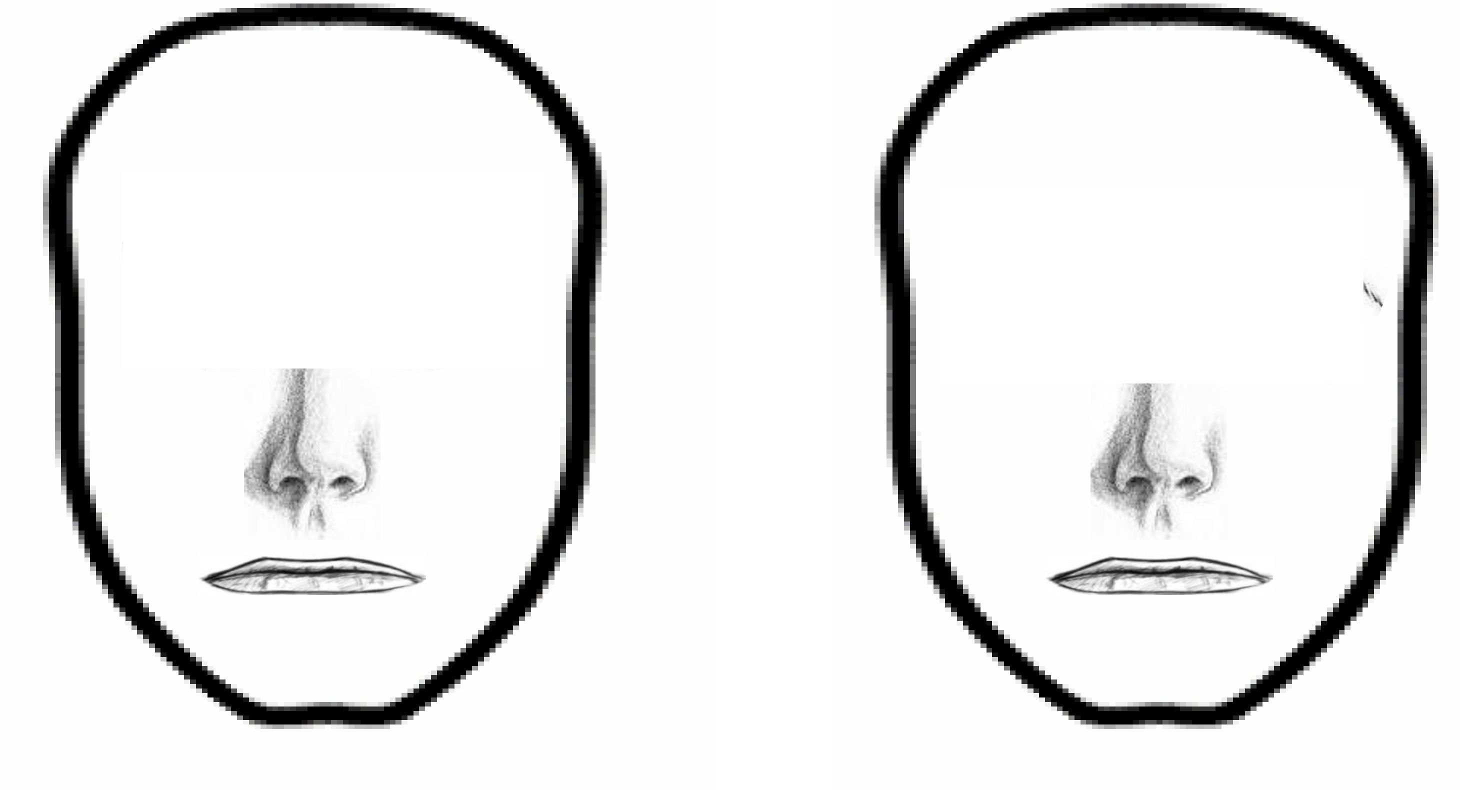

Challenge of measuring human categories

Features often cannot be interpreted independently

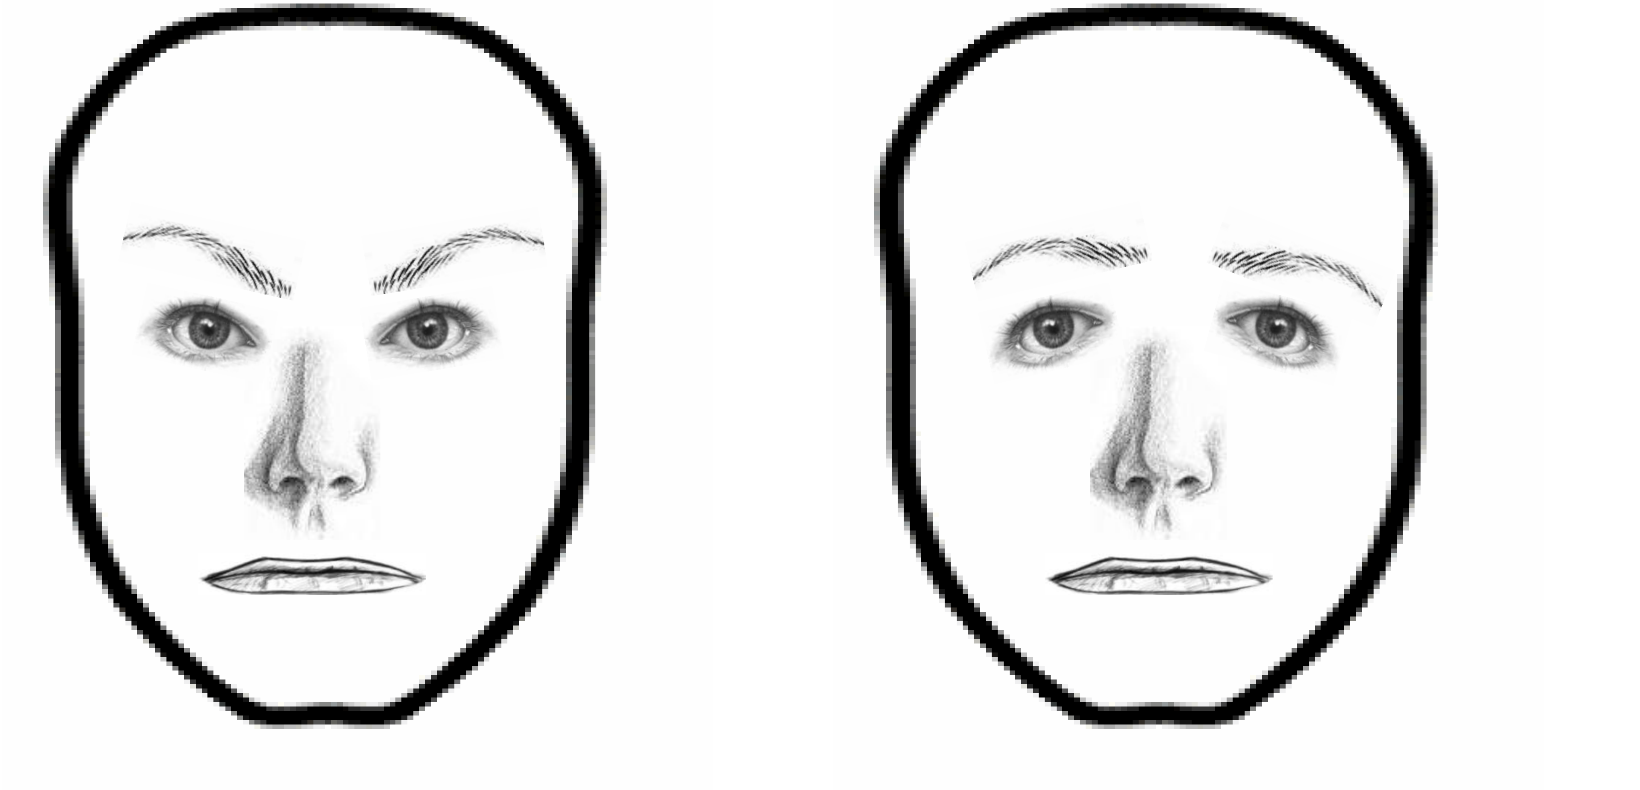

Challenge of measuring human categories

Features often cannot be interpreted independently

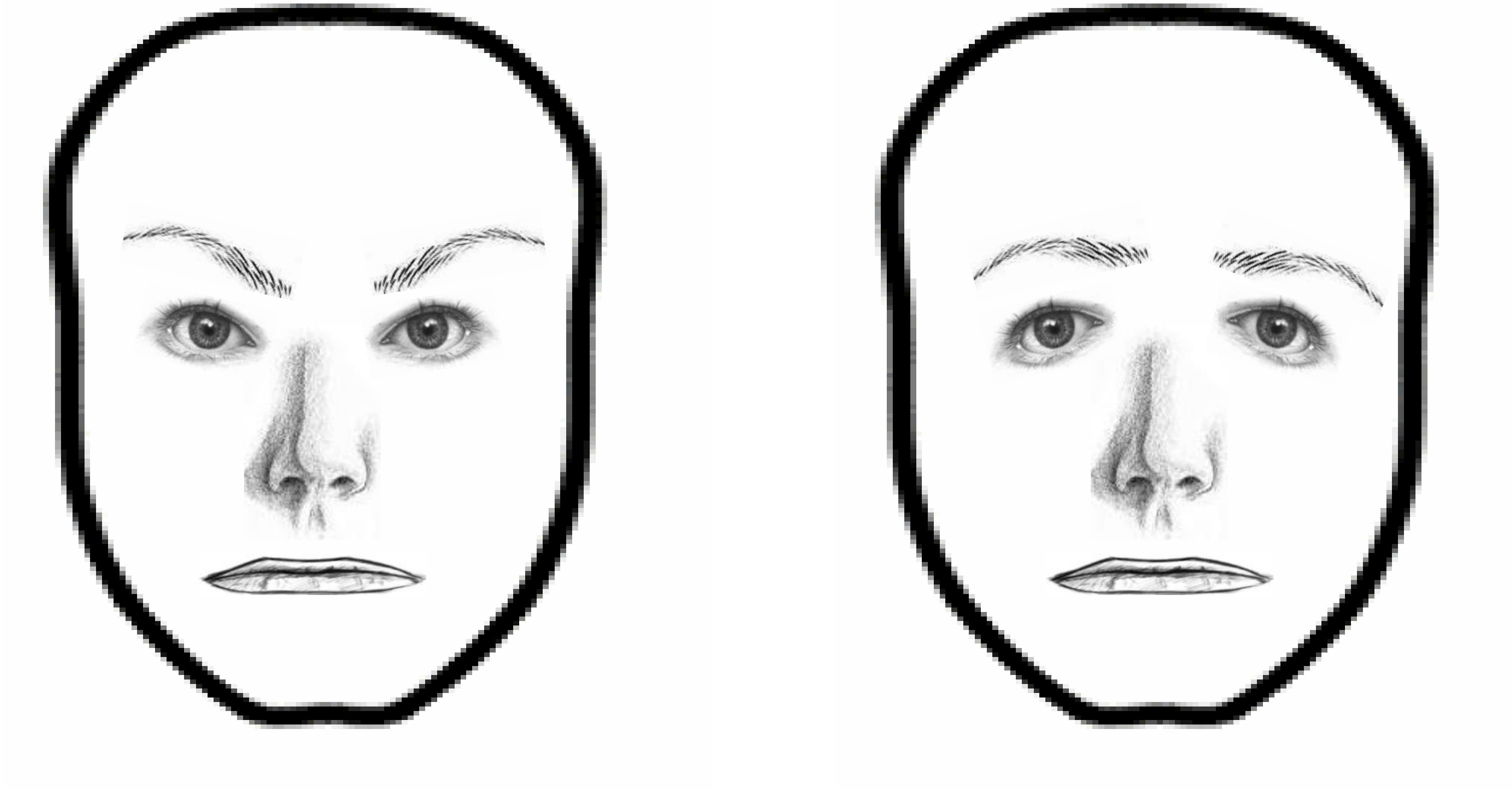

Many interdependent features-> extremely high dimensional stimuli

Challenge of measuring human categories

Problem: Human categories often have complex structure, span high dimensional spaces. How do we measure complex, real-life categories?

A start: Assume all items can be assigned a weight associated with its “representativeness” of the category

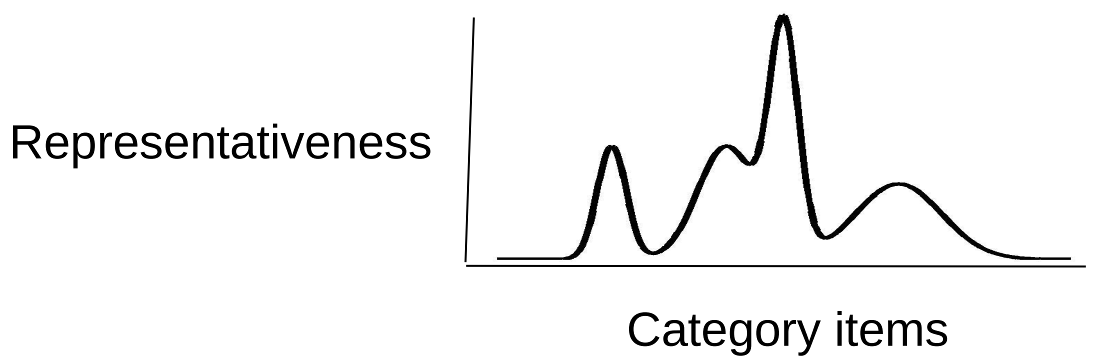

A start: Assume all items can be assigned a weight associated with its "representativeness" of the category

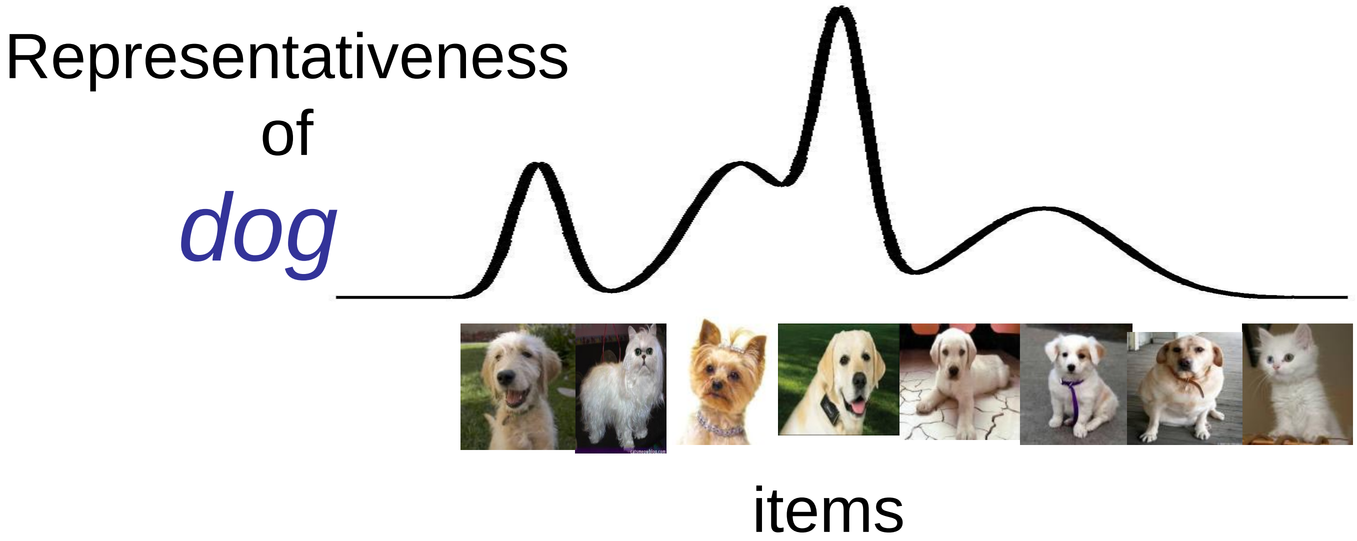

Challenge of measuring human categories

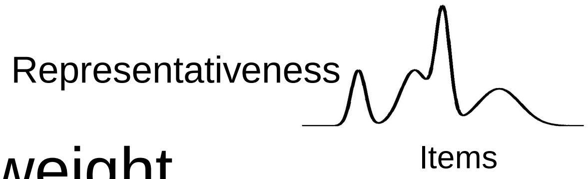

A start: Assume all items can be assigned a weight associated with its "representativeness" of the category

Which is a giraffe?

Probability $p(x|c)$

Which phone?

Preferences $f(x)$

Which looks more happy?

Strength $s(x)$

Challenge of measuring human categories

Example: Happy faces

-One possibility is to just ask:

How representative is this of a happy face?

Problems:

-Hard for people answer this question (no idea of scale)

-Will require prohibitively large number of ratings to explore the category (1000 ratings for 1000 items)

-Another possibility is to pick random pairs and ask:

Which is more representative of a happy face?

Advantage:

-Binary choice task is easier to answer

Problems:

- Will require prohibitively large number of ratings to sample all 499500 possible pairs

Solution: Markov Chain Monte Carlo (MCMC) for exploring human categories

- uses easy binary choice task
- efficiently explores category space by finding and focusing on highly representative regions
- Adapted from widely and long-used algorithm from computer science (Metropolis et al. 1953)

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

-States are examples from set of stimuli you are exploring

-We'll need to evaluate what is "nearby", will explain later.

MCMC method for exploring human categories:

1) Start at current state S_C

2) Choose at random a nearby proposal state S_P

3) Let participant choose whether to accept proposal as new state

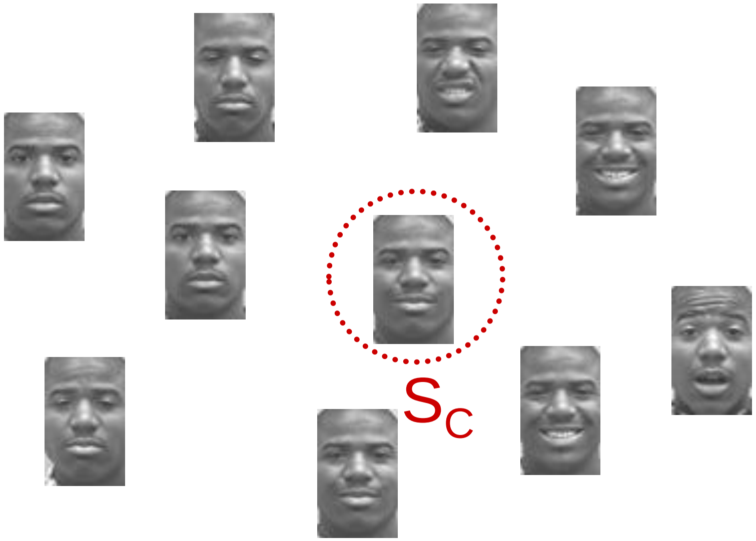

Chain:

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

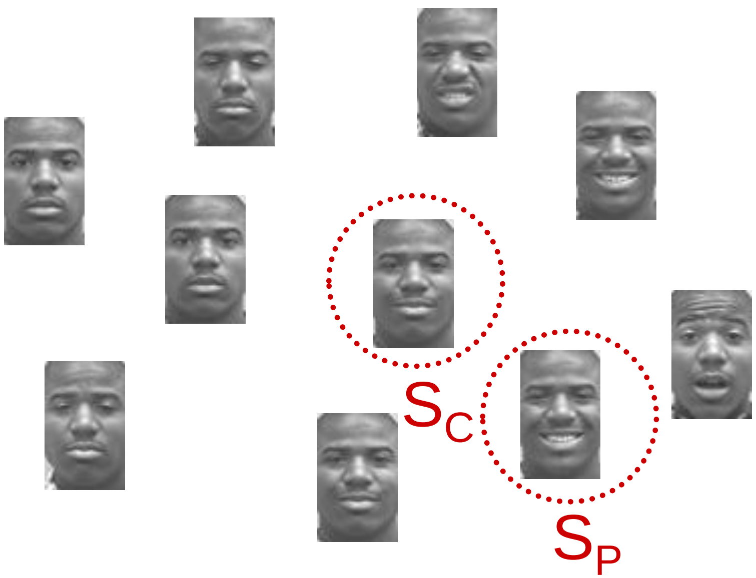

Chain:

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

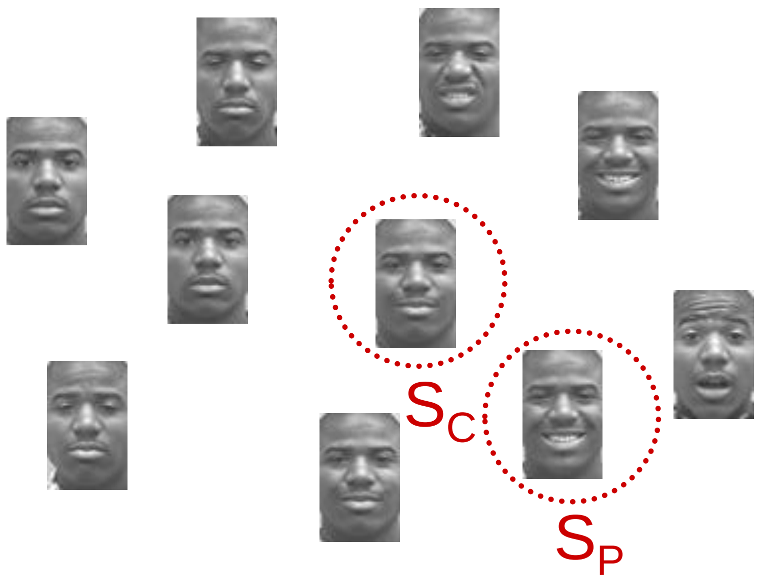

Experiment screen:

Which face looks more happy?

Chain:

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a at random nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

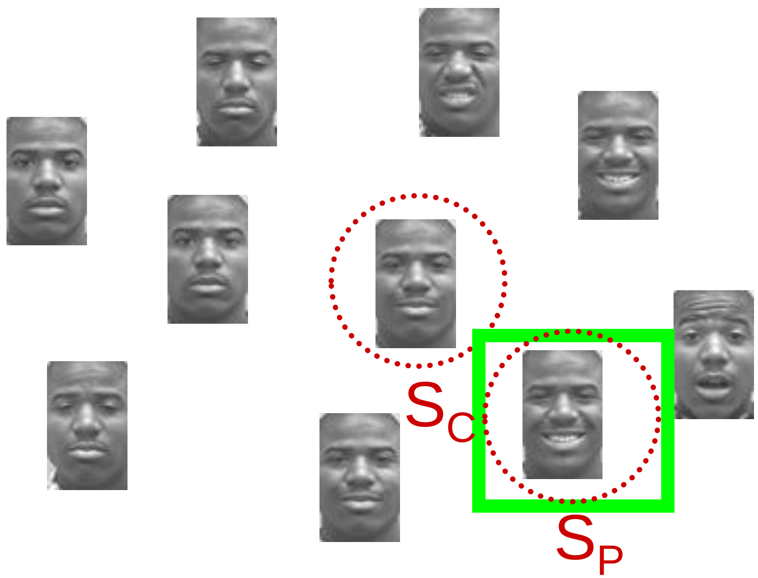

Experiment screen:

Which face looks more happy?

Chain:

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

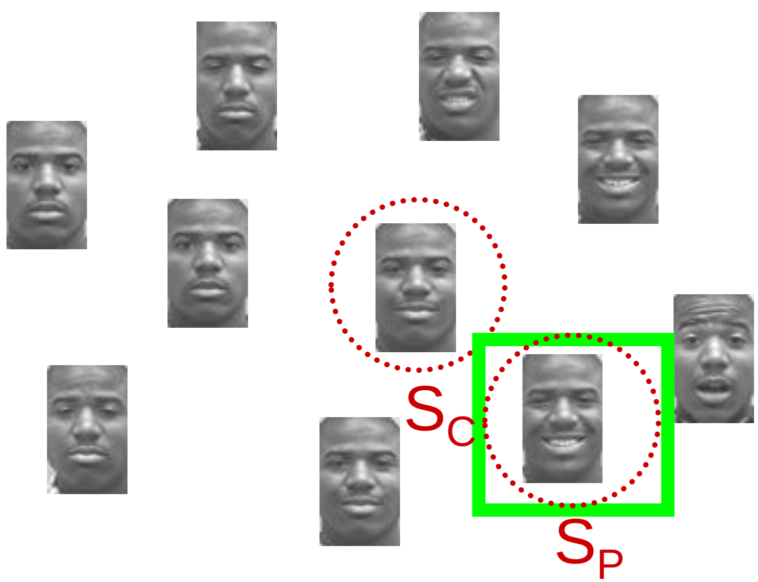

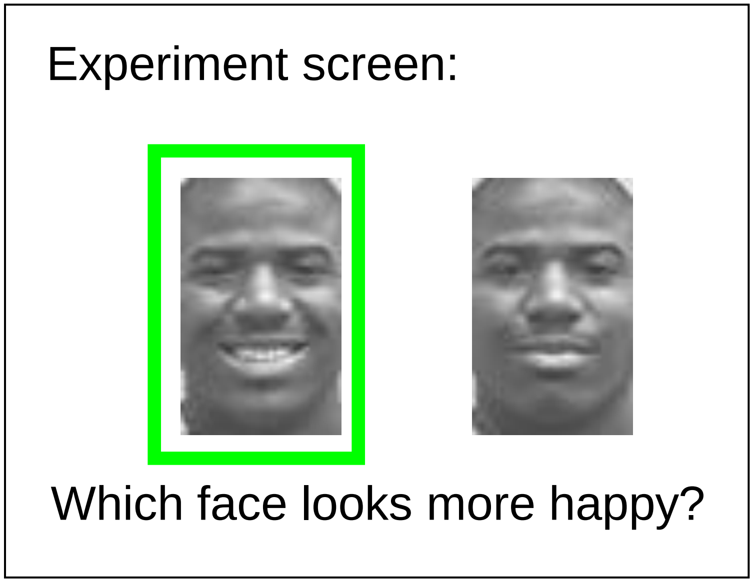

Add selected state to chain

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

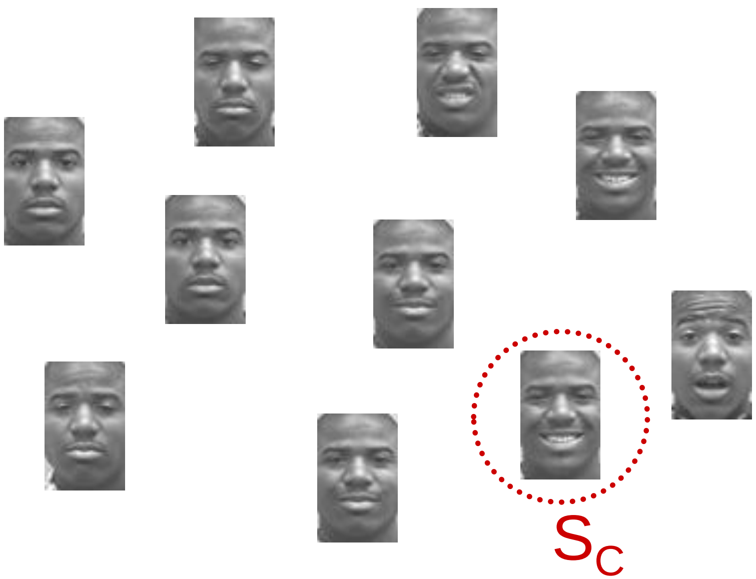

Chosen state is the new current state

(Here the proposed state became the current state)

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

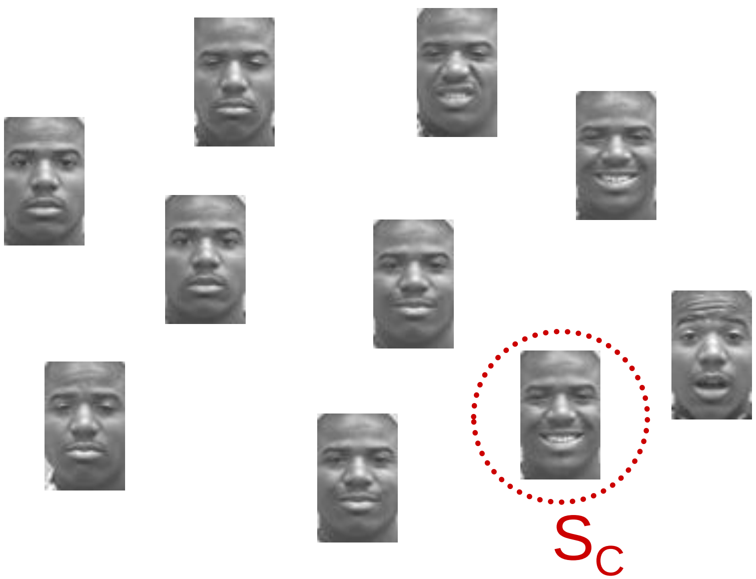

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

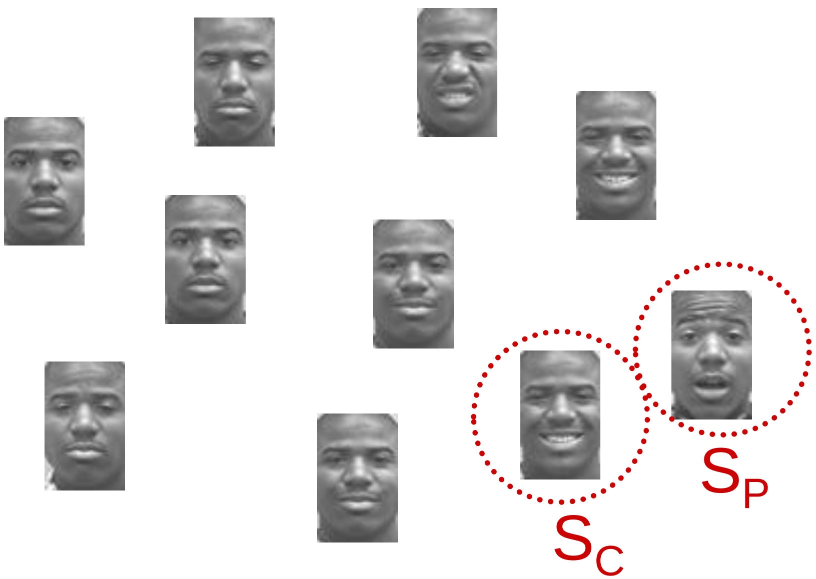

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

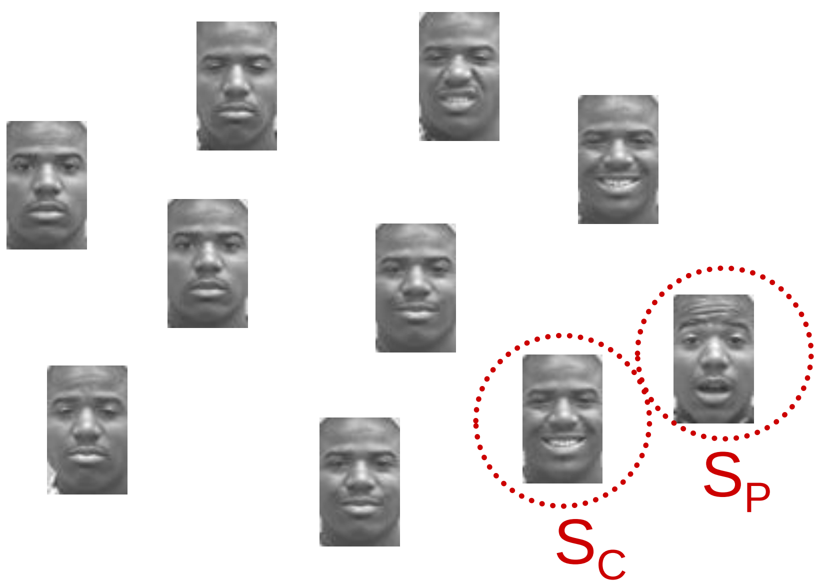

Experiment screen:

Which face looks more happy?

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

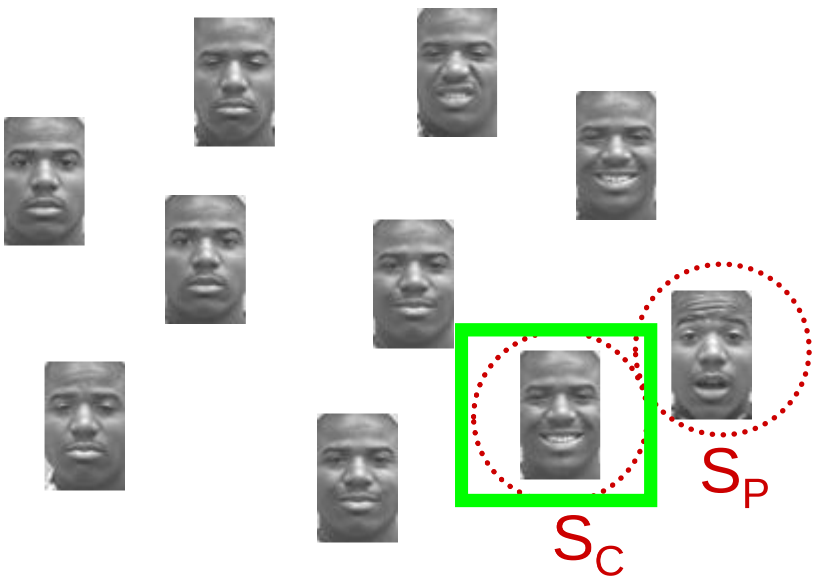

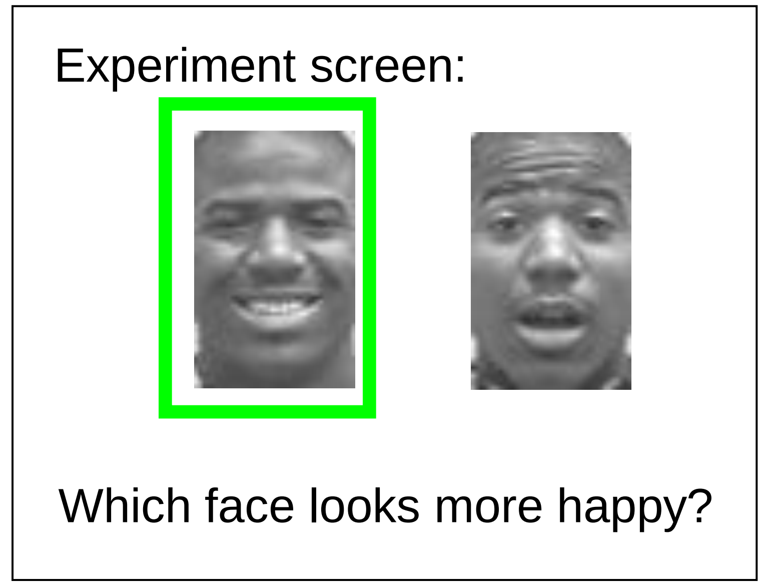

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

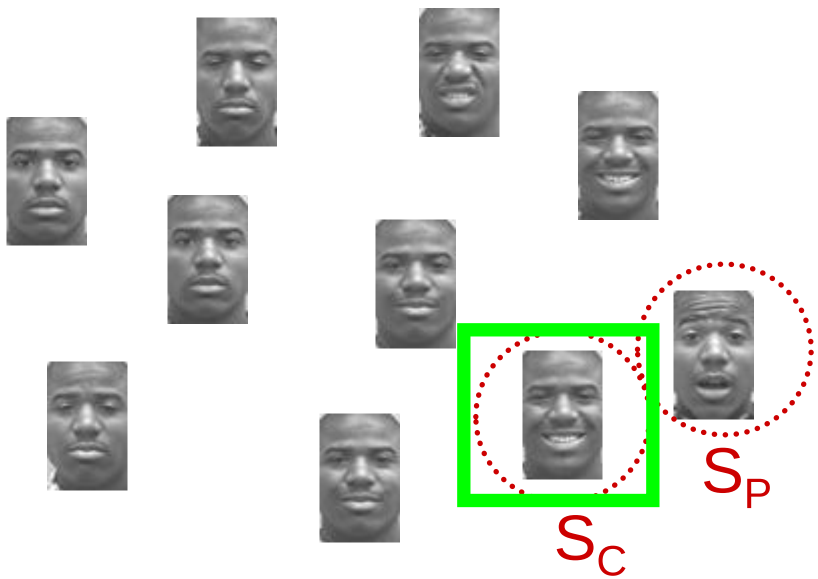

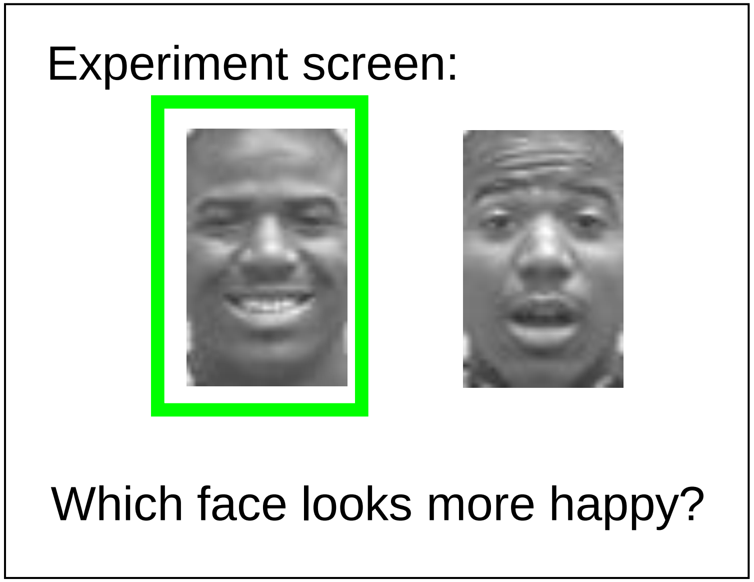

Add selected state to chain

Chain: 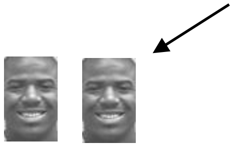

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

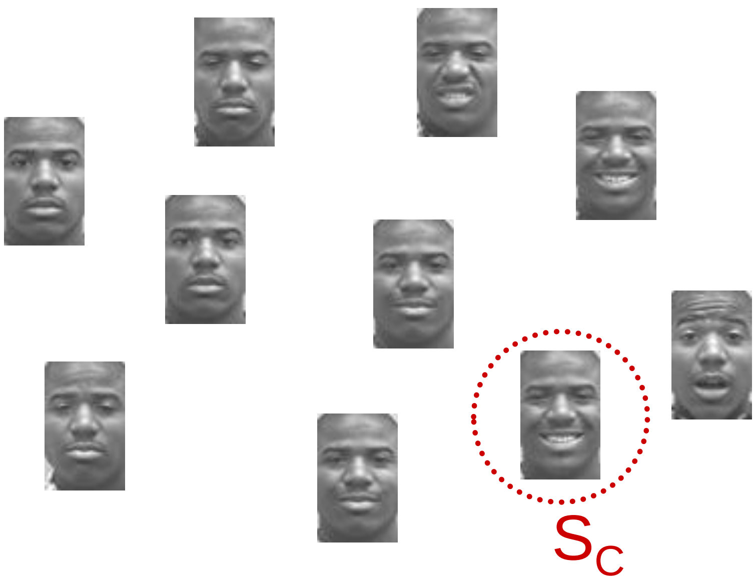

Chosen state is the new current state

(*Here current state stayed the same*)

Add selected state to chain

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

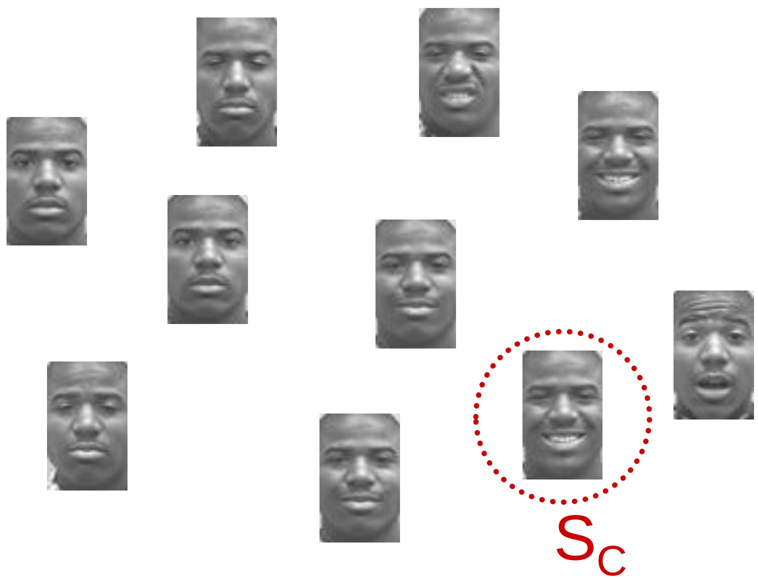

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

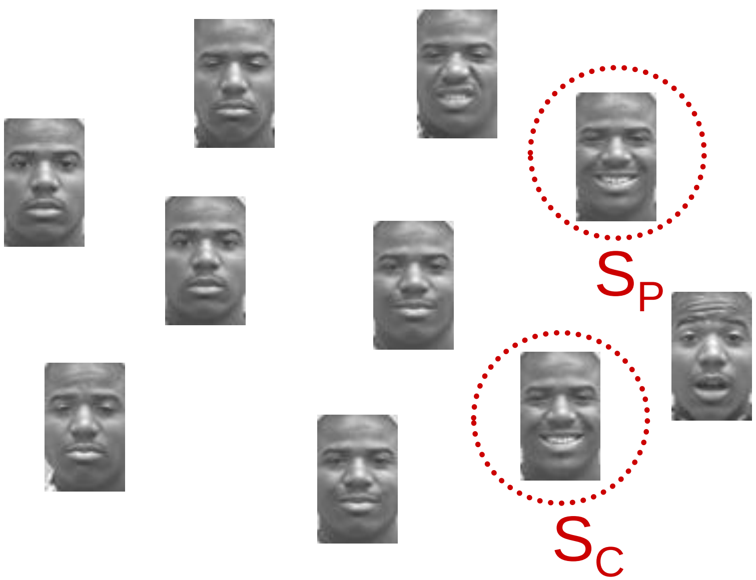

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

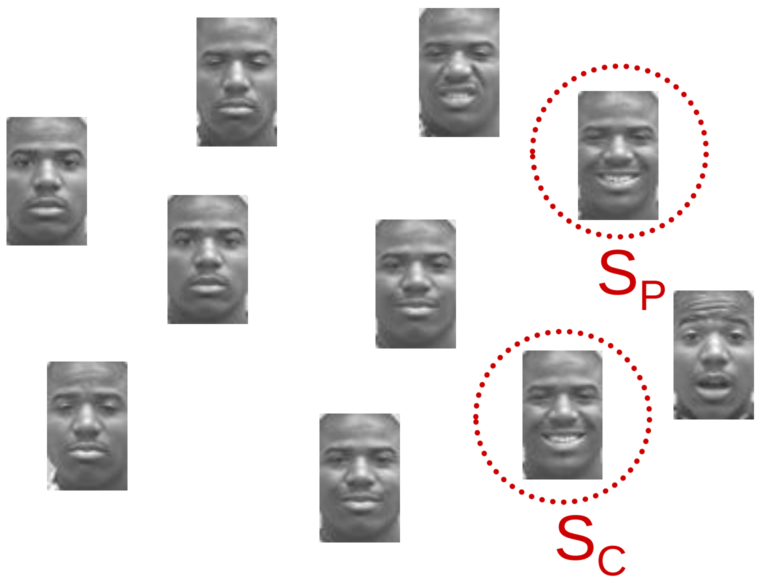

Experiment screen:

Which face looks more happy?

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose at random a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

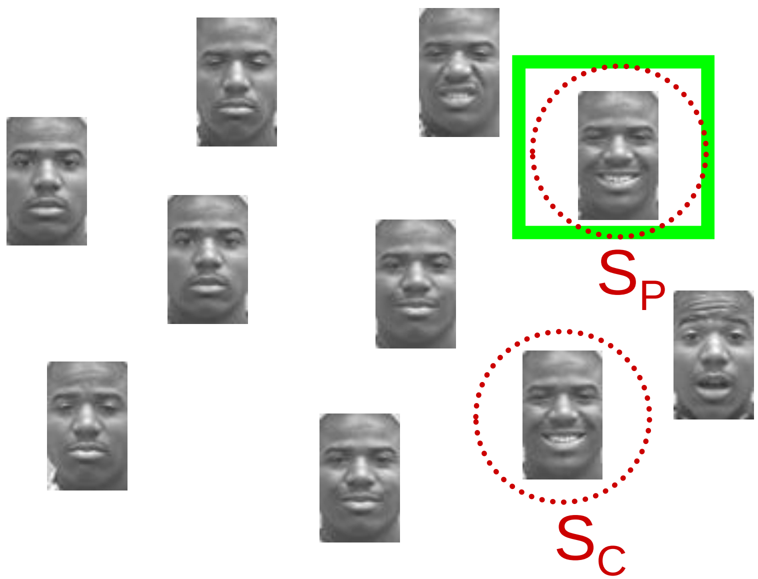

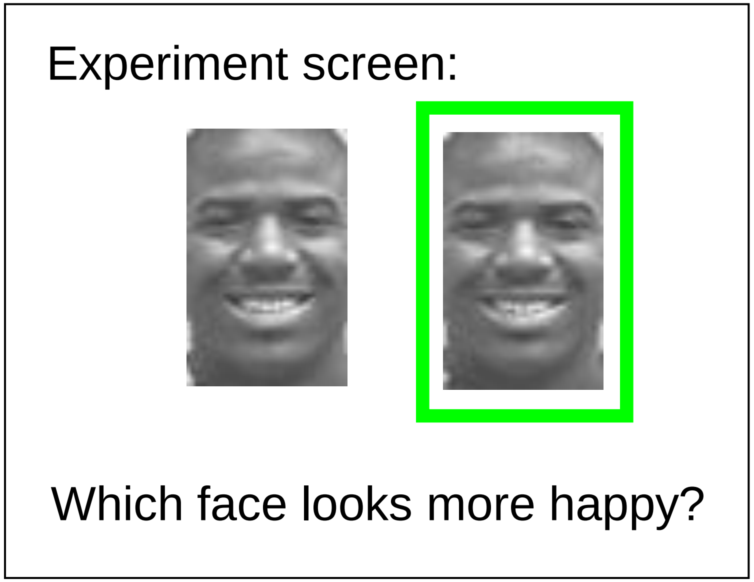

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

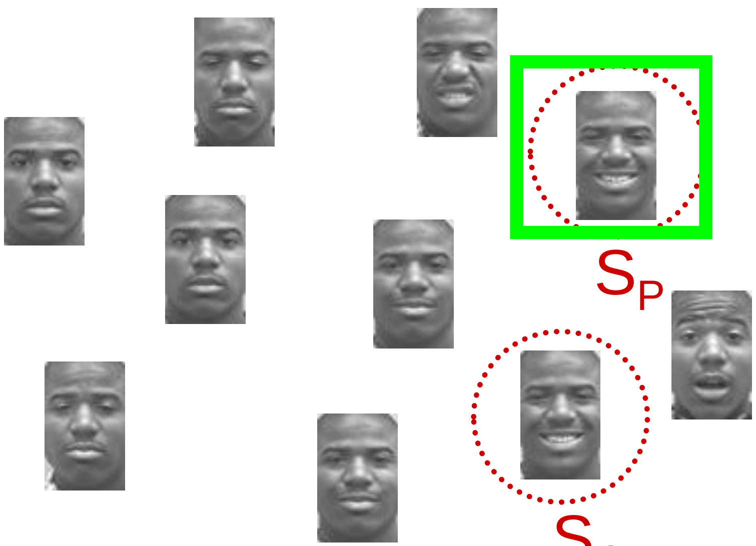

Experiment screen:

Which face looks more happy?

Add selected state to chain

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

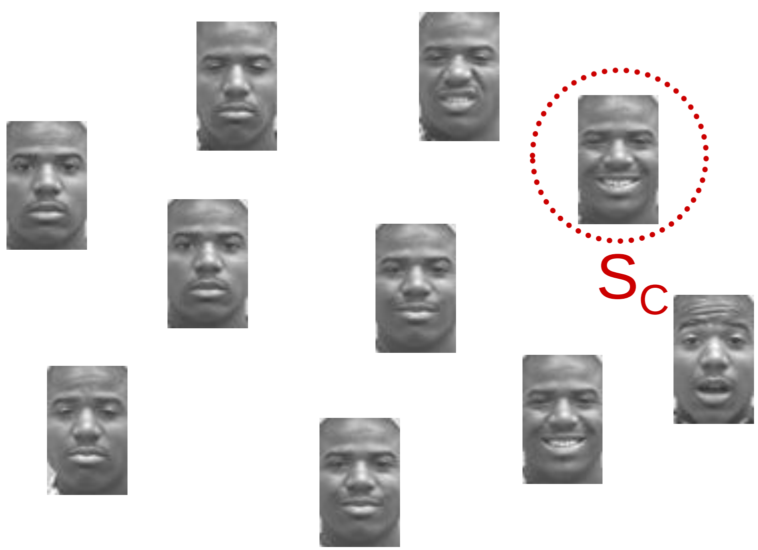

Chosen state is the new current state

(*Here the proposed state became the current state*)

Chain: 

MCMC method for exploring human categories:

1) Start at current state S_C
2) Choose a nearby proposal state S_P
3) Let participant choose whether to accept proposal as new state

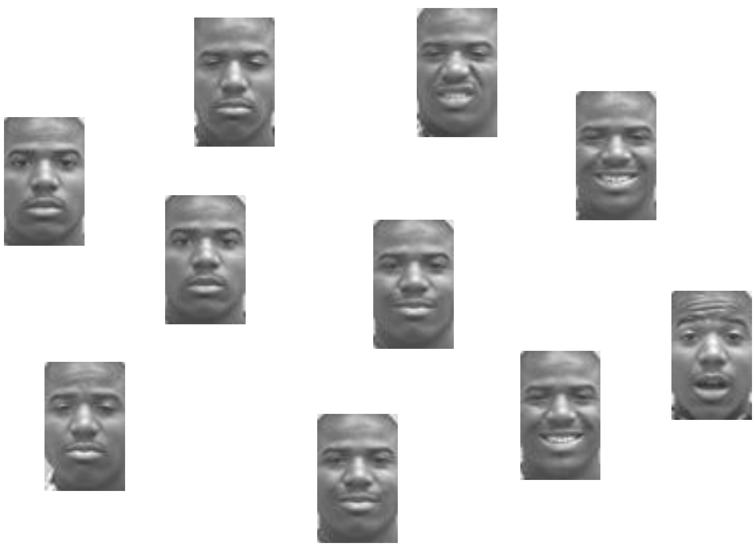

Eventually you end up with a long chain of states

Chain:  ...

MCMC method for exploring human categories:

. . .

Chain:  . . .

Build a histogram over states

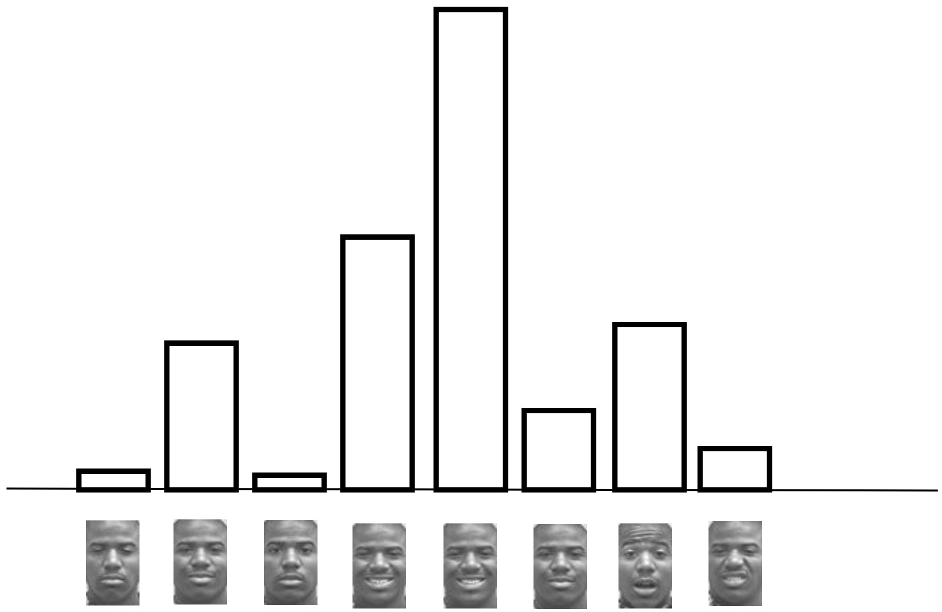

MCMC method for exploring human categories:

Chain:  . . .

. . .

Build a histogram over states

Representativeness
of a happy face

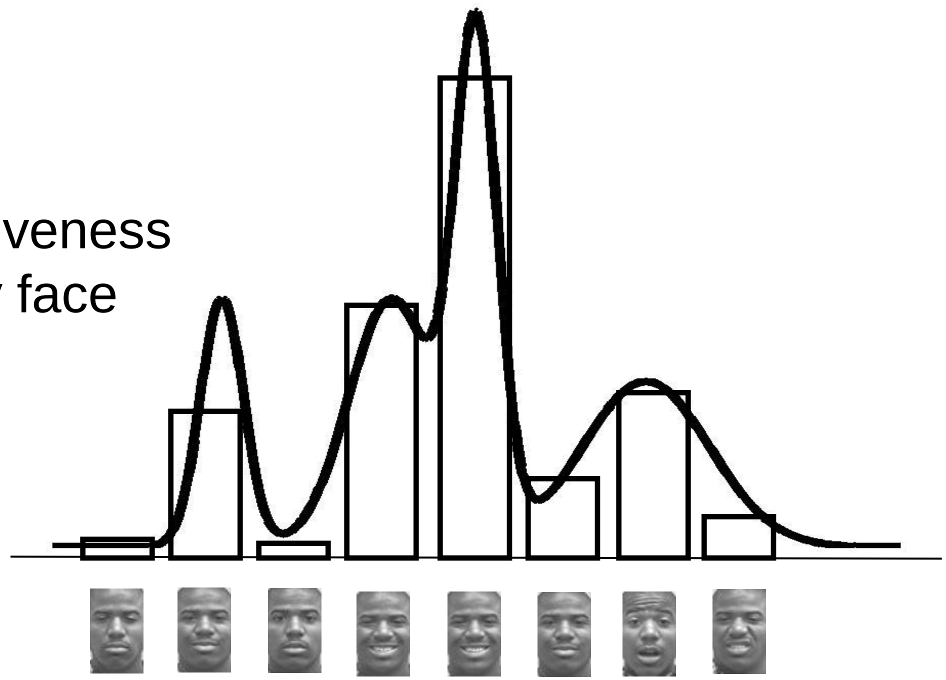

MCMC algorithm from computer science

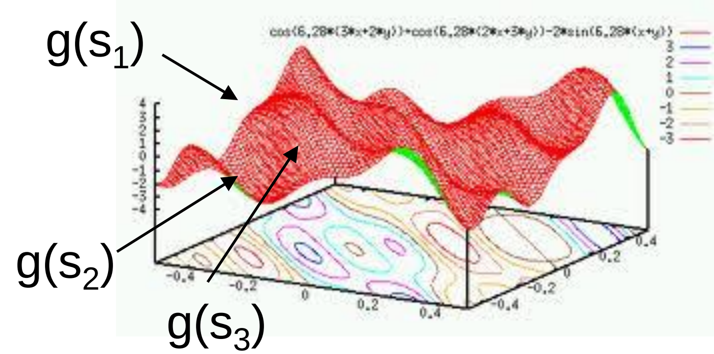

Goal: to sample from complex high-dimensional probability distribution

Assume: know relative weights $g(S)$, *the shape of distribution*

Algorithm:
1) Start at current state S_C
2) Choose a nearby proposal state S_p
3) Accept proposal as new state with probability $A(S_p;S_C)$

MCMC acceptance probability

$$A(S_p;S_C)=\frac{g(S_p)}{g(S_C)+g(S_p)}$$

MCMC algorithm from computer science

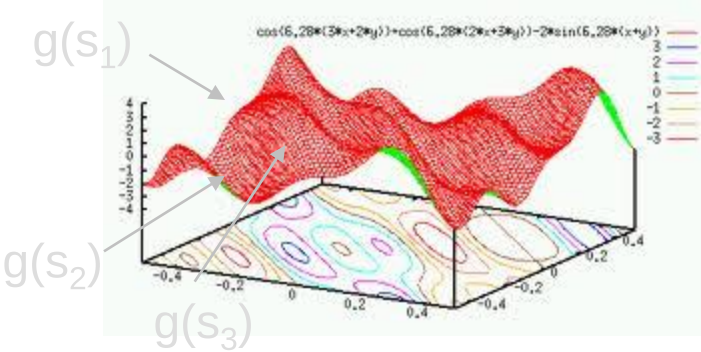

Goal: to sample from complex high-dimensional probability distribution

Assume: know relative weights $g(S)$, *the shape of distribution*

Computer Science Algorithm:
1) Start at current state S_C
2) Choose a nearby proposal state S_p
3) Accept proposal as new state with probability $A(S_p;S_C)$

MCMC acceptance probability

$$A(S_p;S_C)=\frac{g(S_p)}{g(S_C)+g(S_p)}$$

Human MCMC algorithm

3) Let participant choose whether to accept proposal as new state. Assume they accept proposal with $L(S_p;S_C)$

Luce's choice rule for binary choice

$$L(S_p;S_C) = \frac{W(S_p)}{W(S_C)+W(S_p)}$$

Here we don't need the shape of distribution

MCMC algorithm from computer science

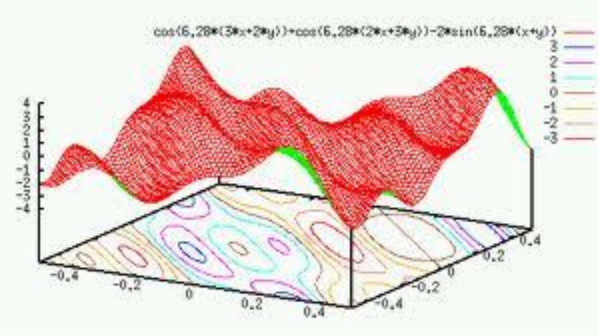

Goal: to sample from complex high-dimensional probability distribution

Assume: know relative weights g(S), *the shape of the complex function*

Computer Science Algorithm:

1) Start at current state S_C
2) Choose a nearby proposal state S_p
3) Accept proposal as new state with probability $A(S_p;S_C)$

Human MCMC algorithm

1) Start at current state S_C
2) Choose a nearby proposal state S_p
3) Let participant choose whether to accept proposal as new state. Assume they accept proposal with $L(S_p;S_C)$

MCMC acceptance probability

$$A(S_p;S_C)=\frac{g(S_p)}{g(S_C)+g(S_p)}$$

Luce's choice rule for binary choice

$$L(S_p;S_C) = \frac{W(S_p)}{W(S_C)+W(S_p)}$$

Key observation: similar structure

Convergence in task with stick figure animals

Example experiment screen:

Results:

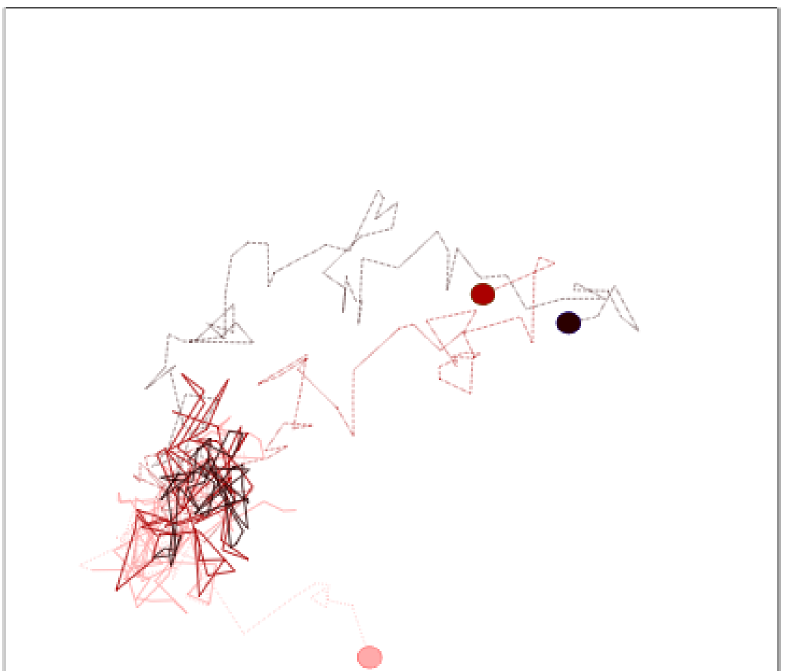

Sanborn, Griffiths, Shriffin 2010

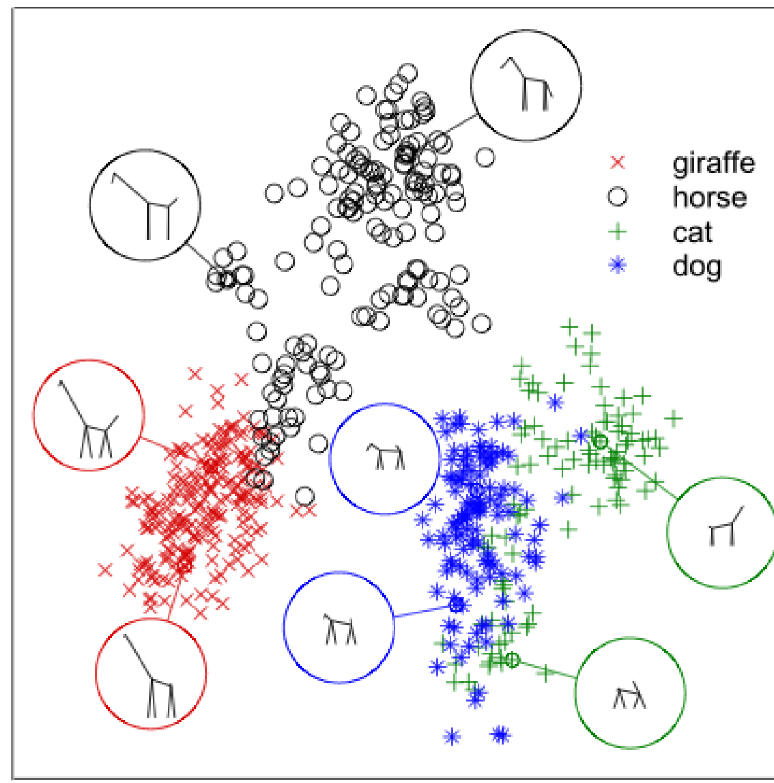

How to make a proposal state?

- Previous work used easily parameterized stimuli

- We want to explore large variety of real life categories and concepts

- How do we propose a nearby state for arbitrary stimulus spaces for which definitions of “nearby” is not obvious (e.g. sets of images)?

How to make a proposal state?

Define similarity metric and
quantify similarity over all pairs

					…
	1	0.7	0.5	0.6	
	0.7	1	0.6	0.1	
	0.5	0.6	1	0.4	
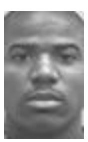	0.6	0.1	0.4	1	
⋮					

How to make a proposal state?

Define similarity metric and quantify similarity over all pairs

1	0.7	0.5	0.6
0.7	1	0.6	0.1
0.5	0.6	1	0.4
0.6	0.1	0.4	1

Run 'B-matching' to obtain graph of N nearest neighbours (e.g. N=3)

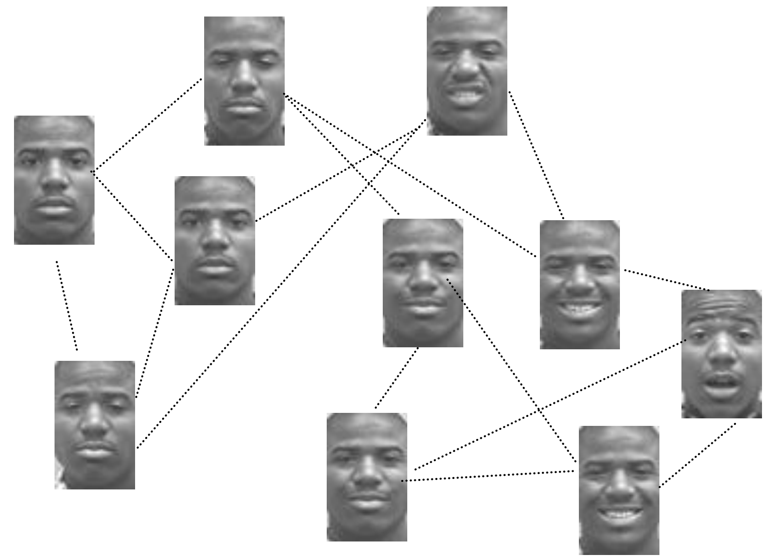

Make sure your set is fully connected

Discrete MCMC makes proposal based on similarity measure between items.

Define similarity metric and quantify similarity over all pairs

1	0.7	0.5	0.6
0.7	1	0.6	0.1
0.5	0.6	1	0.4
0.6	0.1	0.4	1

Run 'B-matching' to obtain graph of N nearest neighbours (e.g. N=3)

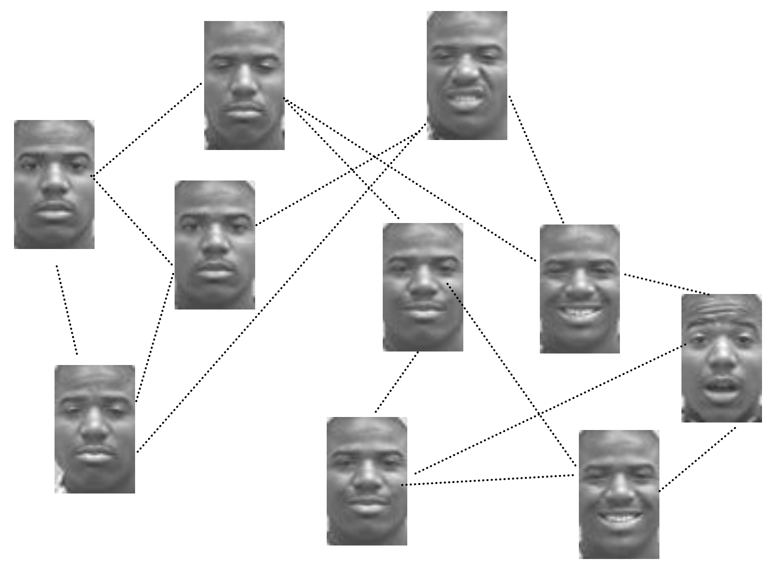

Set proposals to be uniform over neighbours, with small chance of uniform over all stimuli

Discrete MCMC makes proposal based on similarity measure between items.

Define similarity metric and quantify similarity over all pairs

				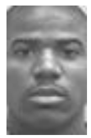	…
	1	0.7	0.5	0.6	
	0.7	1	0.6	0.1	
	0.5	0.6	1	0.4	
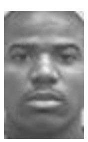	0.6	0.1	0.4	1	
⋮					

Run ‘B-matching’ to obtain graph of N nearest neighbours (e.g. N=3)

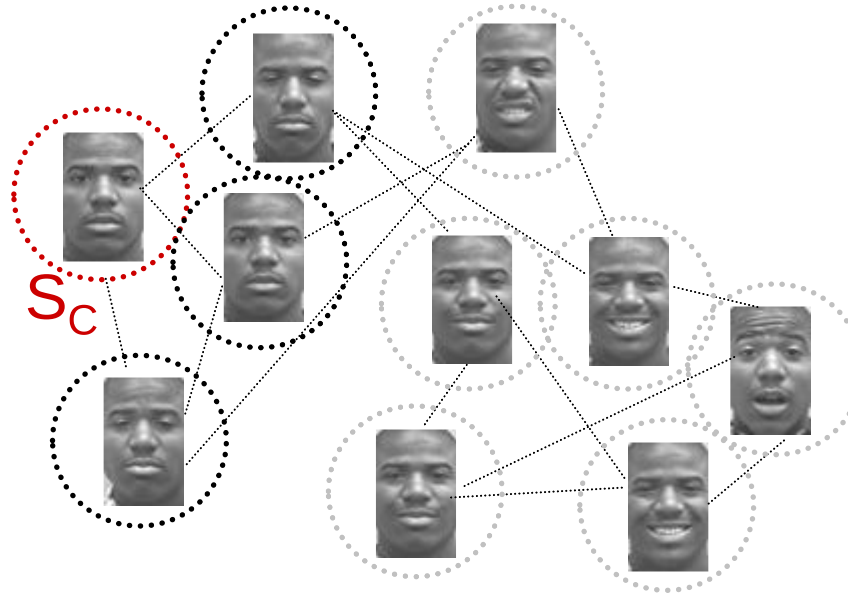

Set proposals to be uniform over neighbours, with small chance of uniform over all stimuli

Discrete MCMC makes proposal based on similarity measure between items.

Define similarity metric and quantify similarity over all pairs

	1	0.7	0.5	0.6
	0.7	1	0.6	0.1
	0.5	0.6	1	0.4
	0.6	0.1	0.4	1

. . .

⋮

Run 'B-matching' to obtain graph of N nearest neighbours (e.g. N=3)

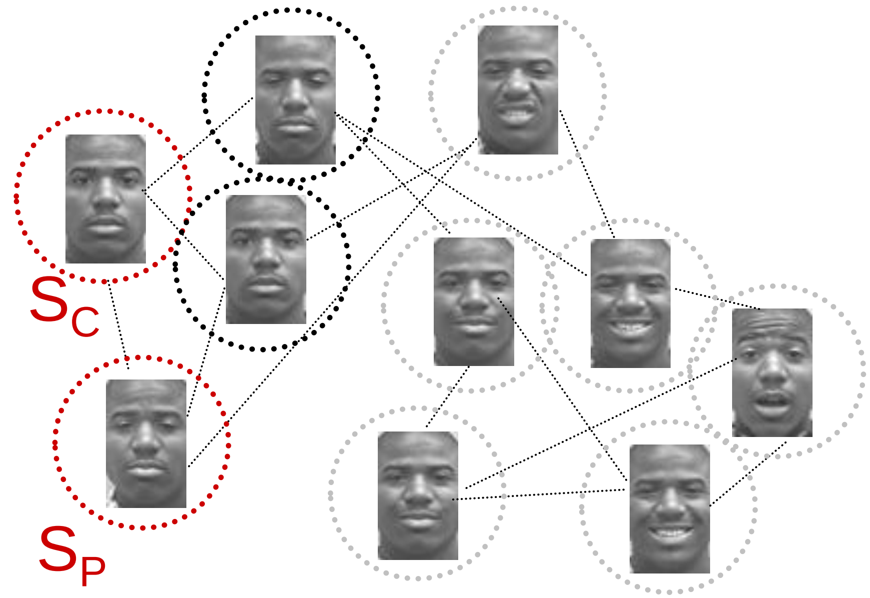

Set proposals to be uniform over neighbours, with small chance of uniform over all stimuli

Discrete MCMC makes proposal based on similarity measure between items.

Define similarity metric and quantify similarity over all pairs

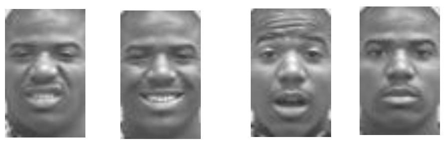

1	0.7	0.5	0.6	…
0.7	1	0.6	0.1	
0.5	0.6	1	0.4	
0.6	0.1	0.4	1	
⋮				

Run ‘B-matching’ to obtain graph of N nearest neighbours (e.g. N=3)

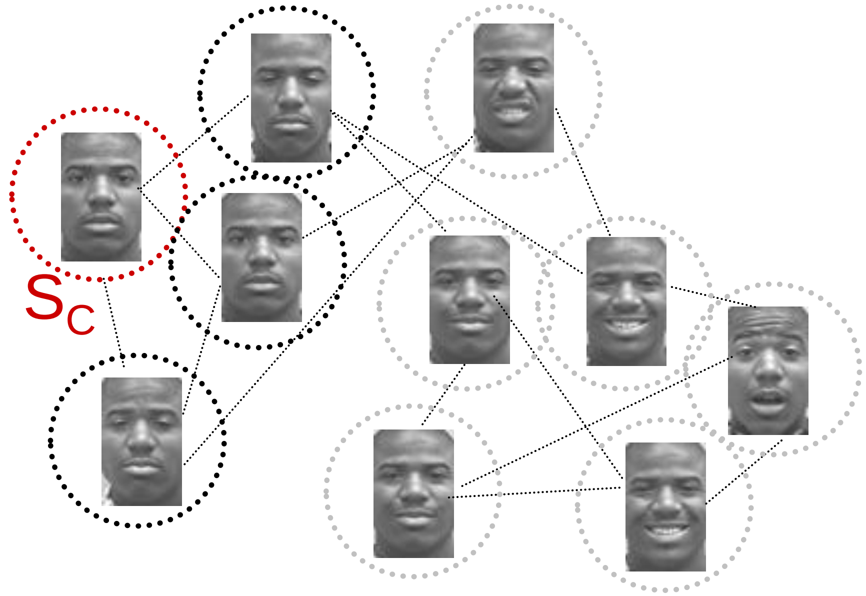

Set proposals to be uniform over neighbours, with small chance of uniform over all stimuli

Discrete MCMC makes proposal based on similarity measure between items.

Define similarity metric and quantify similarity over all pairs

…

1	0.7	0.5	0.6
0.7	1	0.6	0.1
0.5	0.6	1	0.4
0.6	0.1	0.4	1

⋮

Run 'B-matching' to obtain graph of N nearest neighbours (e.g. N=3)

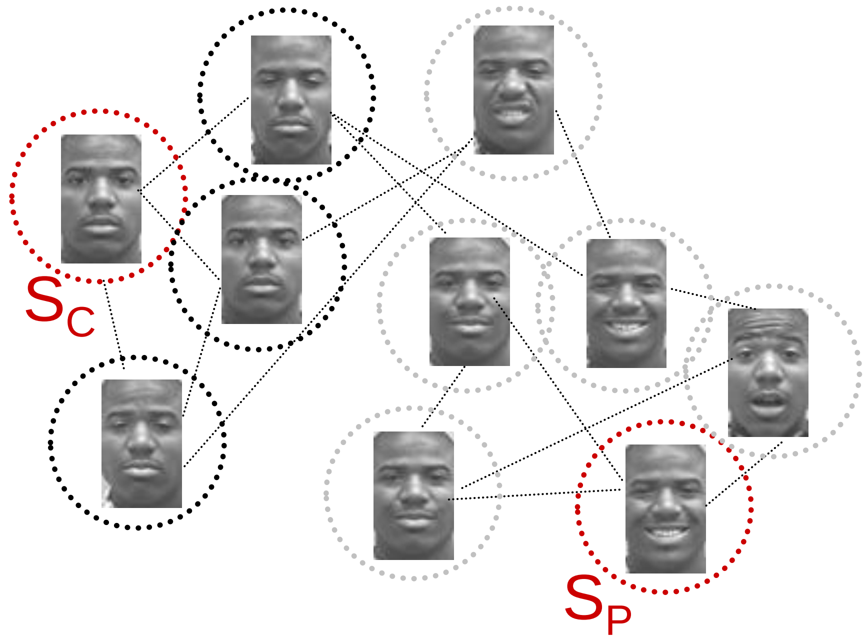

Set proposals to be uniform over neighbours, with small chance of uniform over all stimuli

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

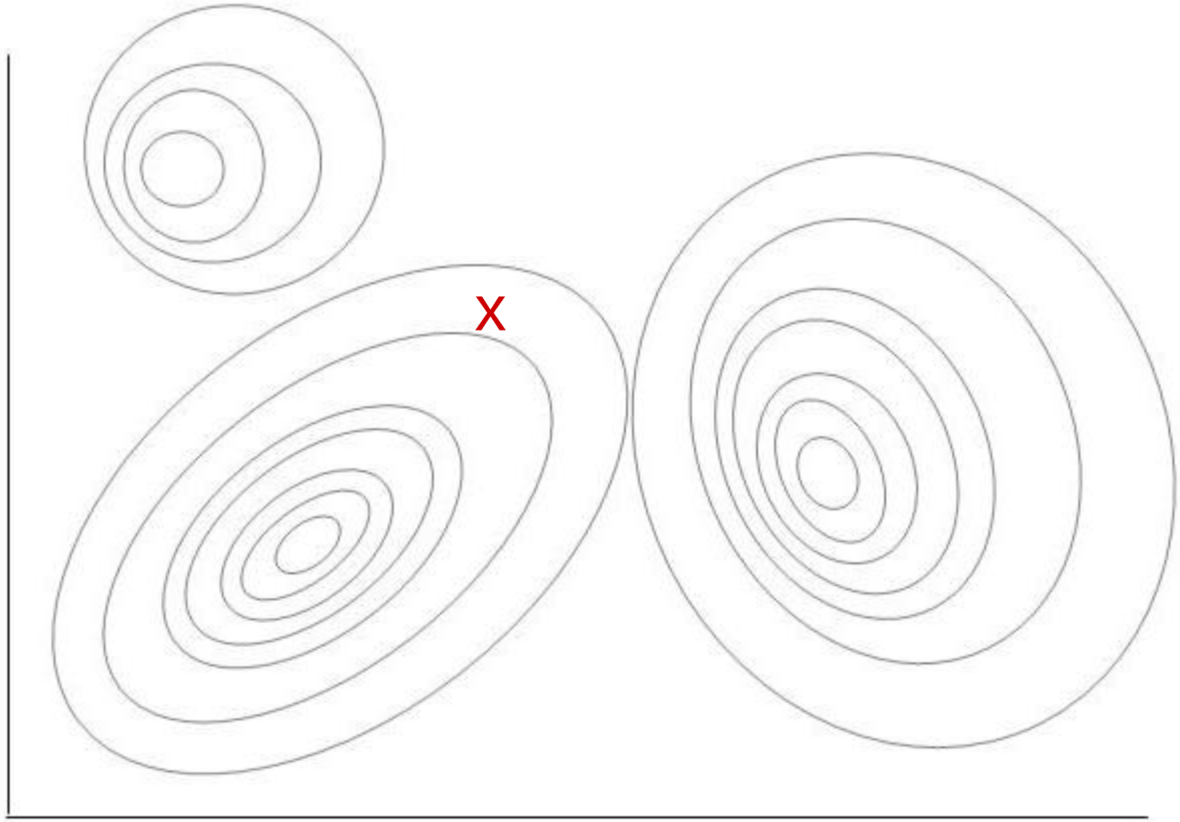

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

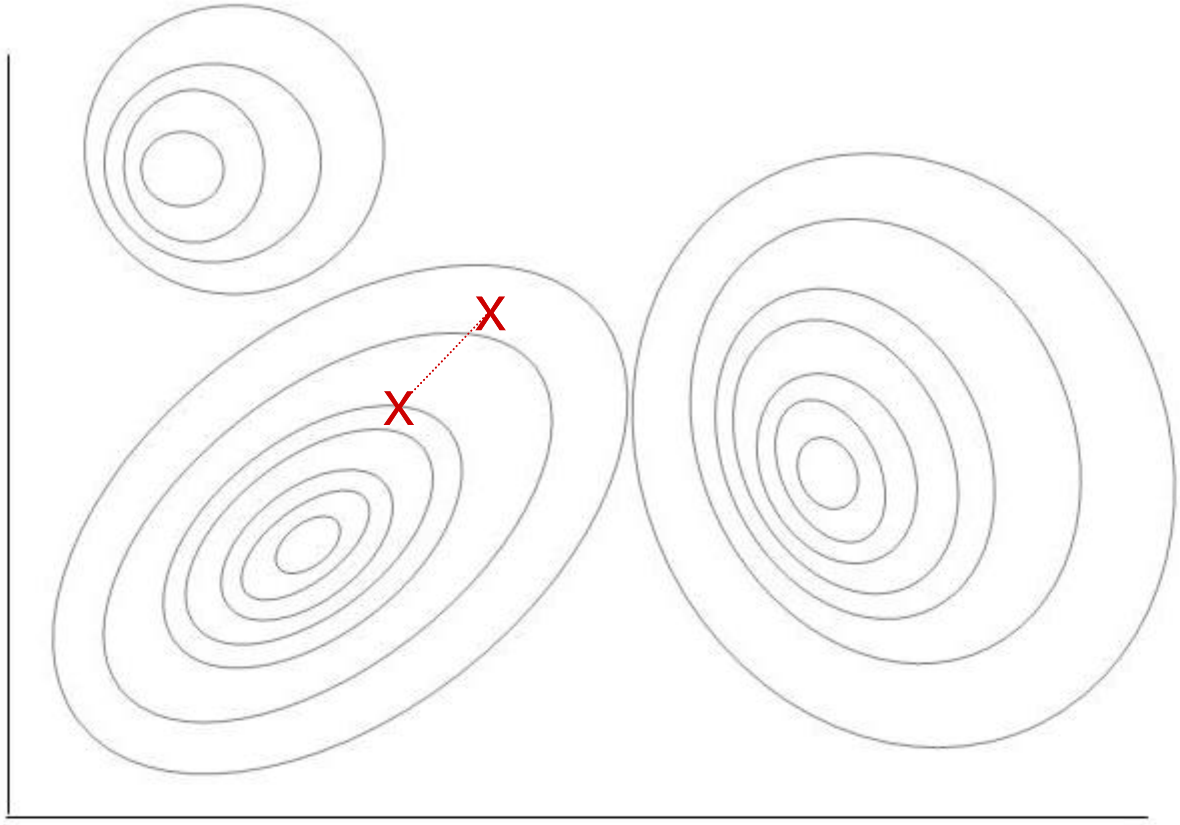

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

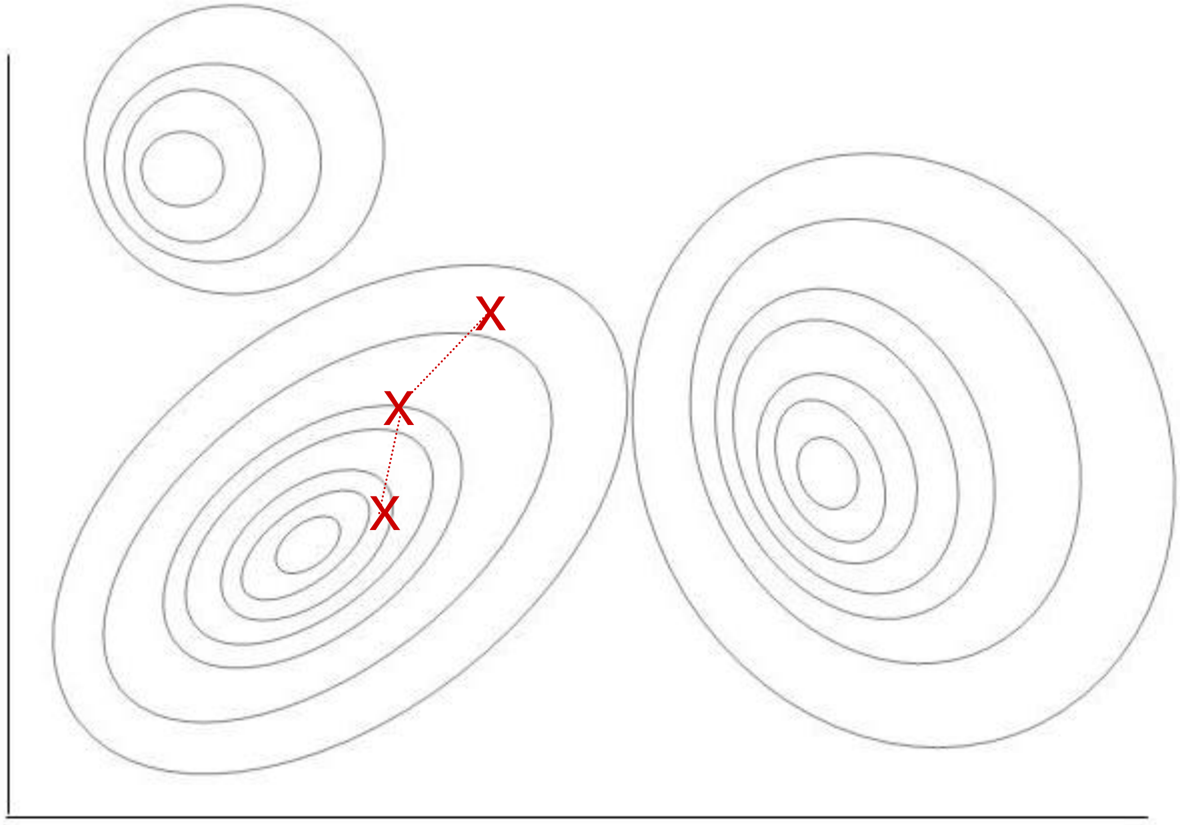

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

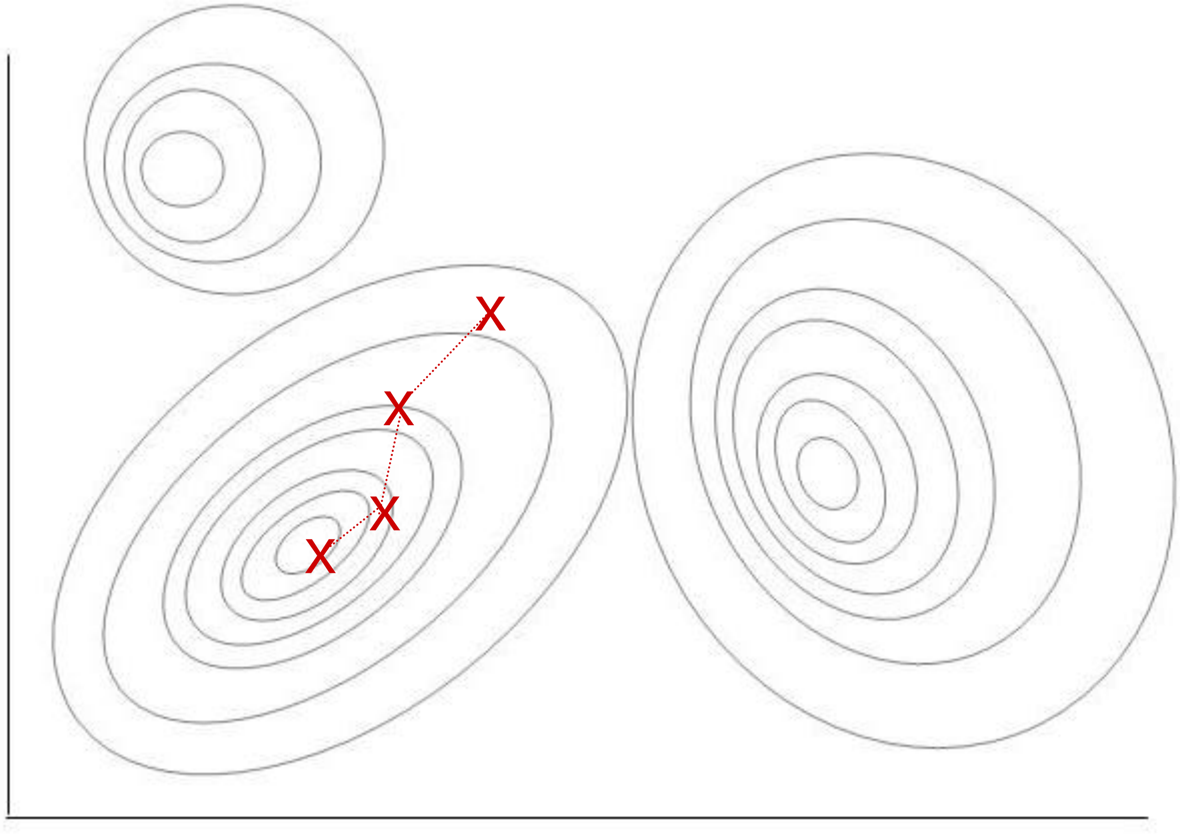

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

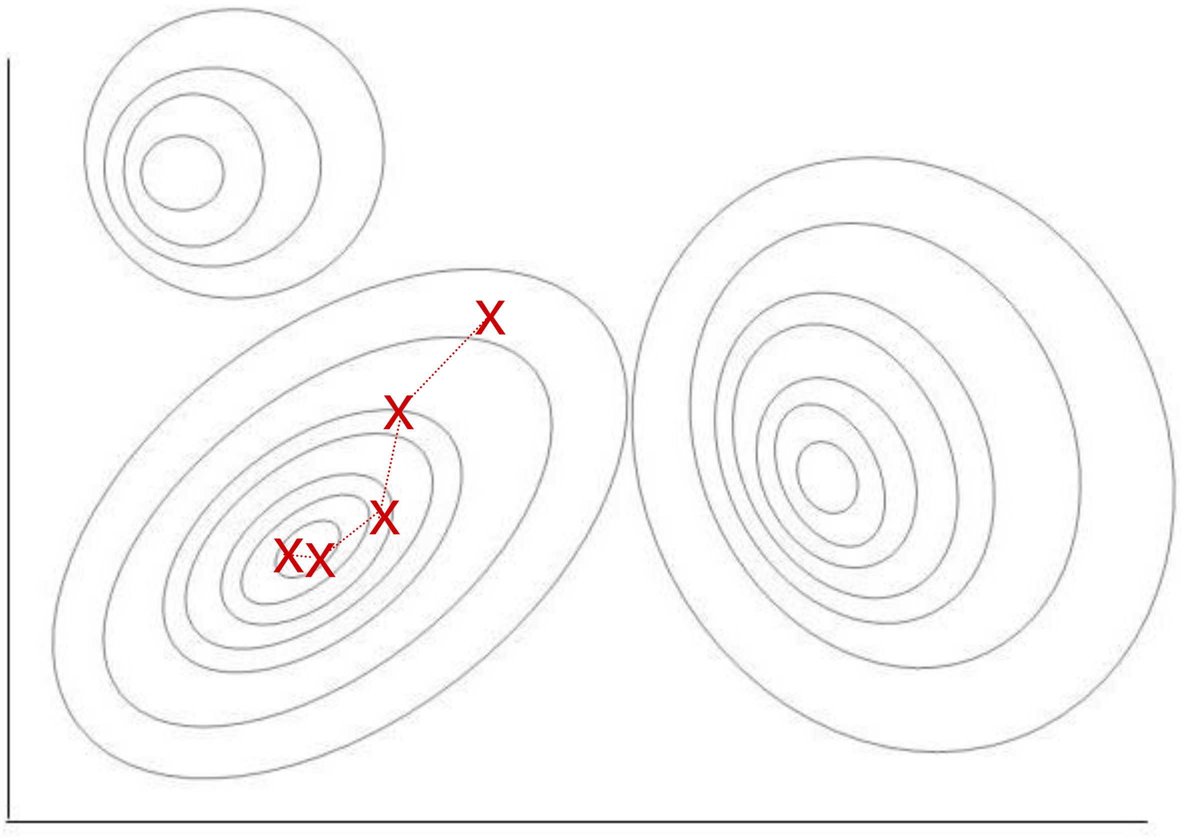

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

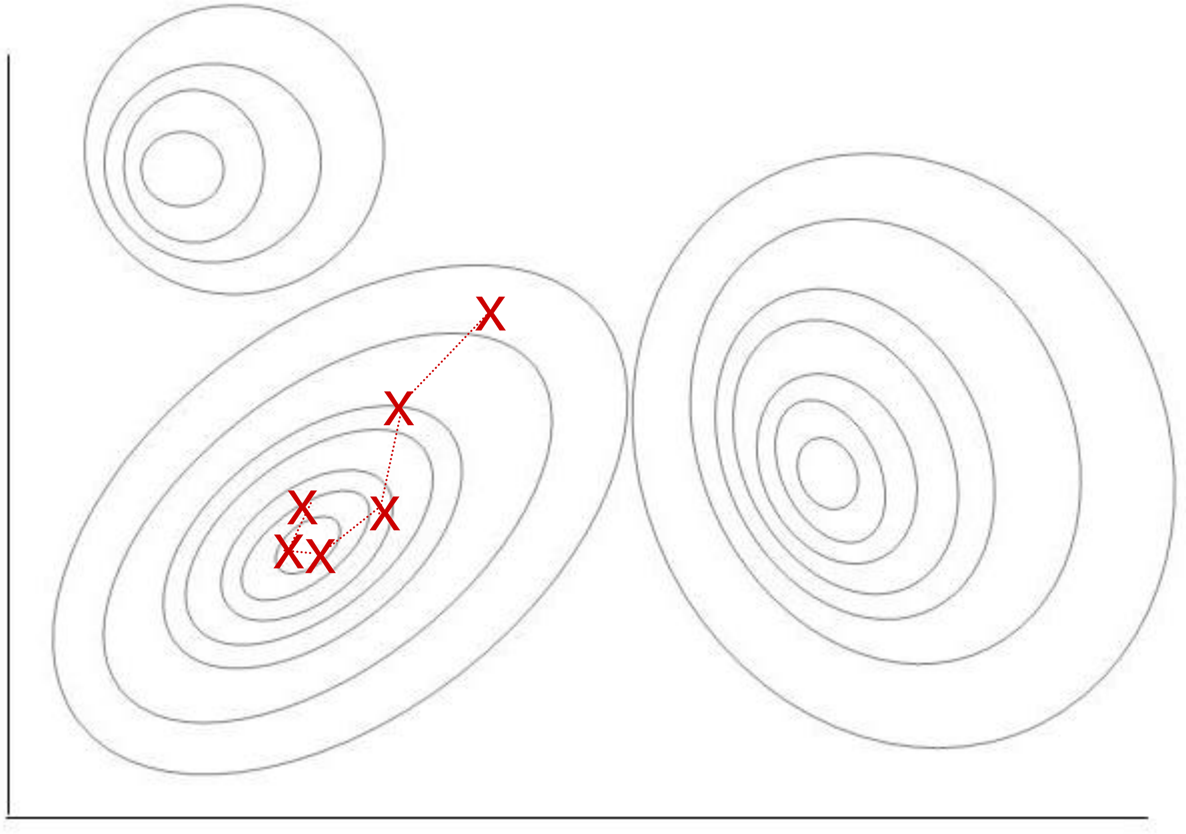

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

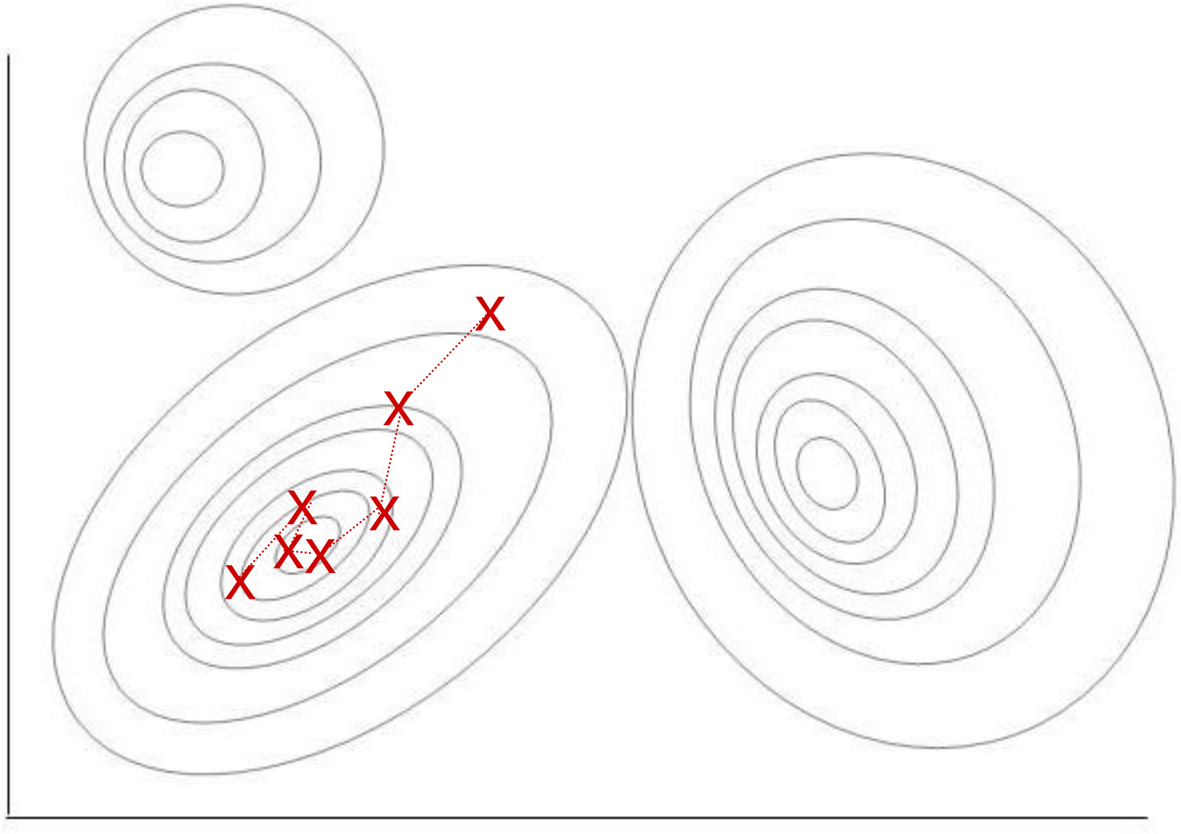

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

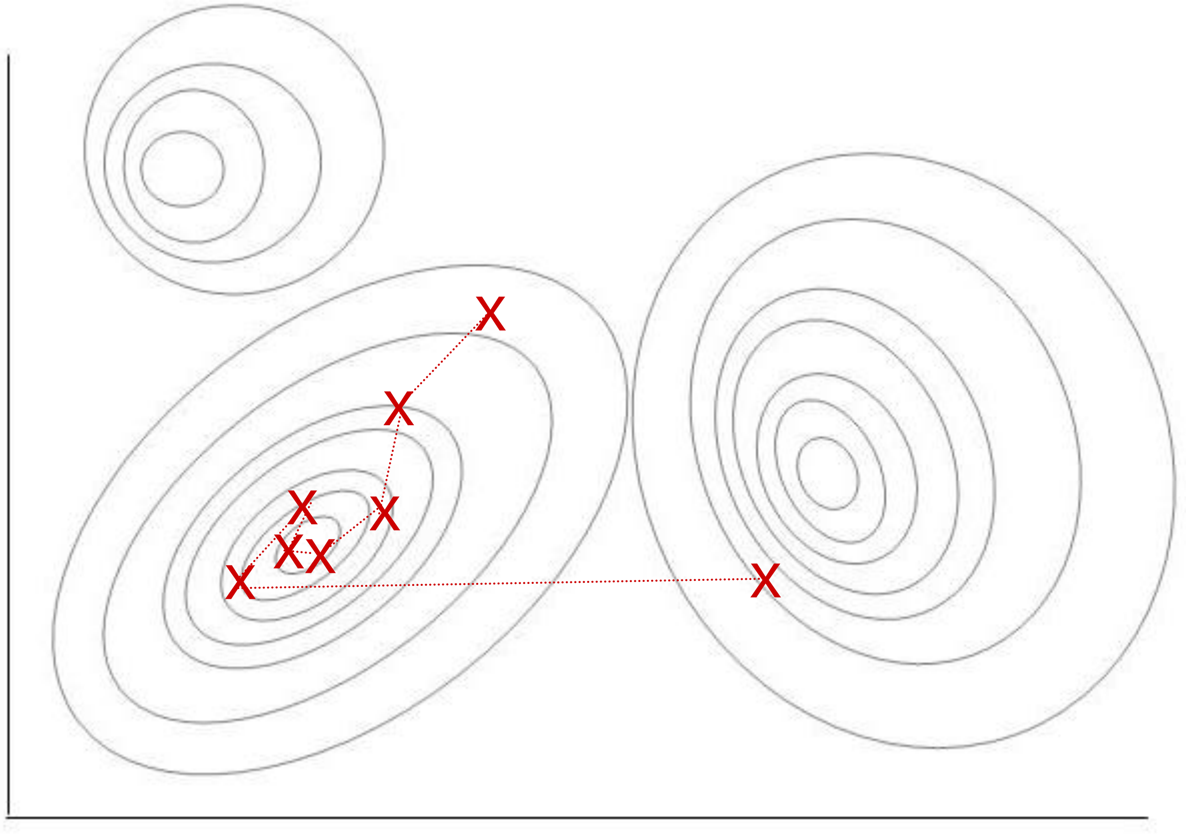

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

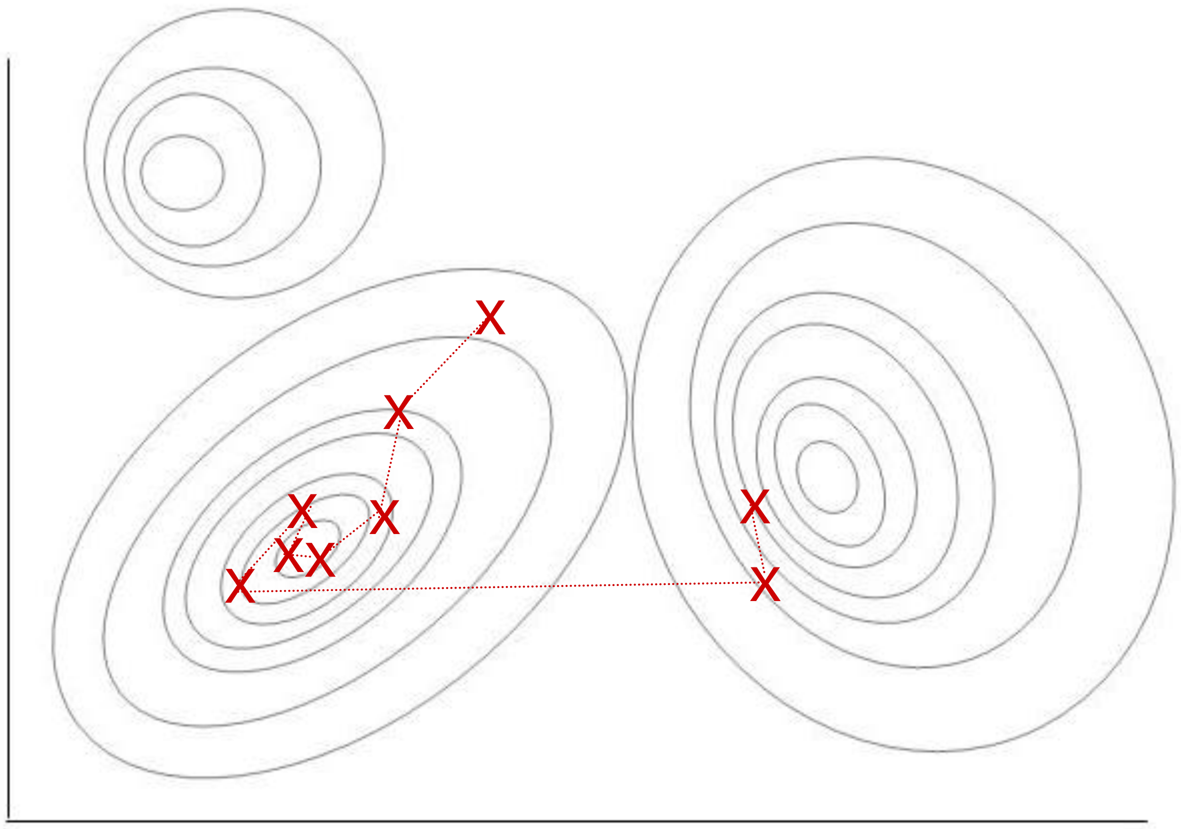

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

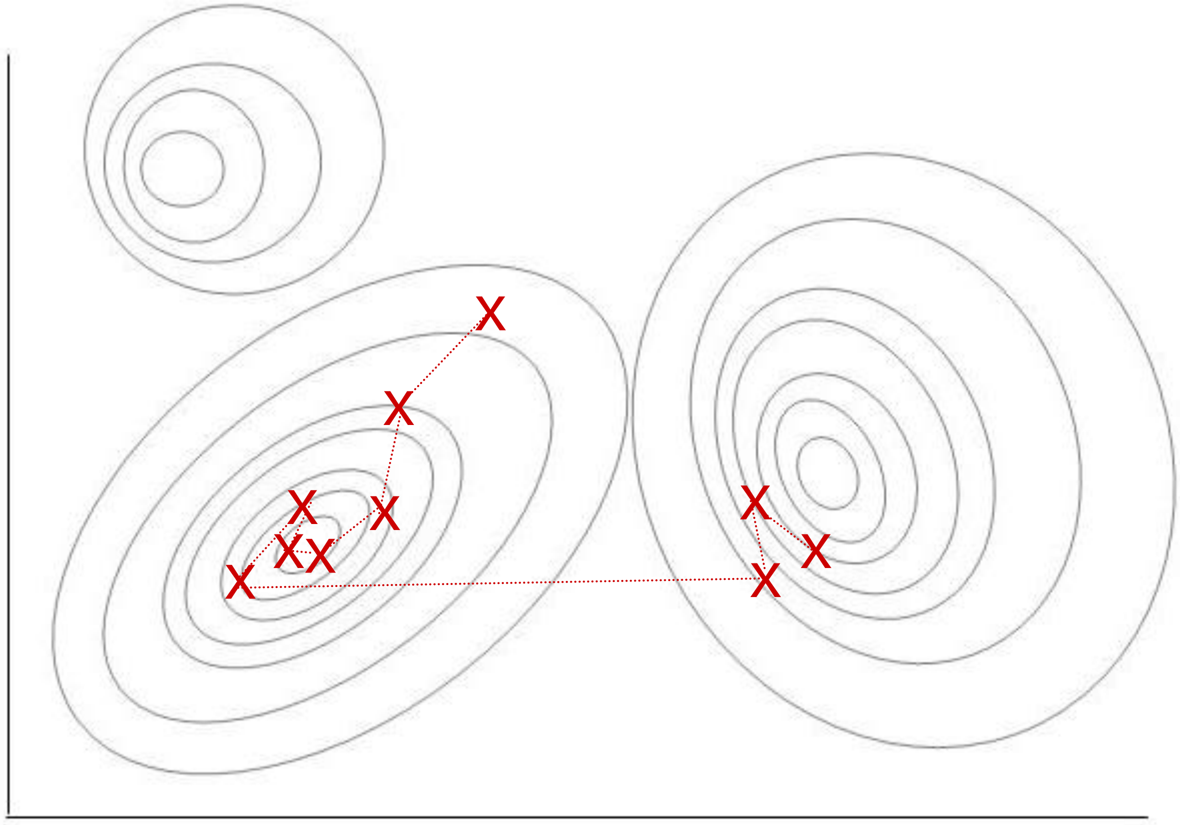

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

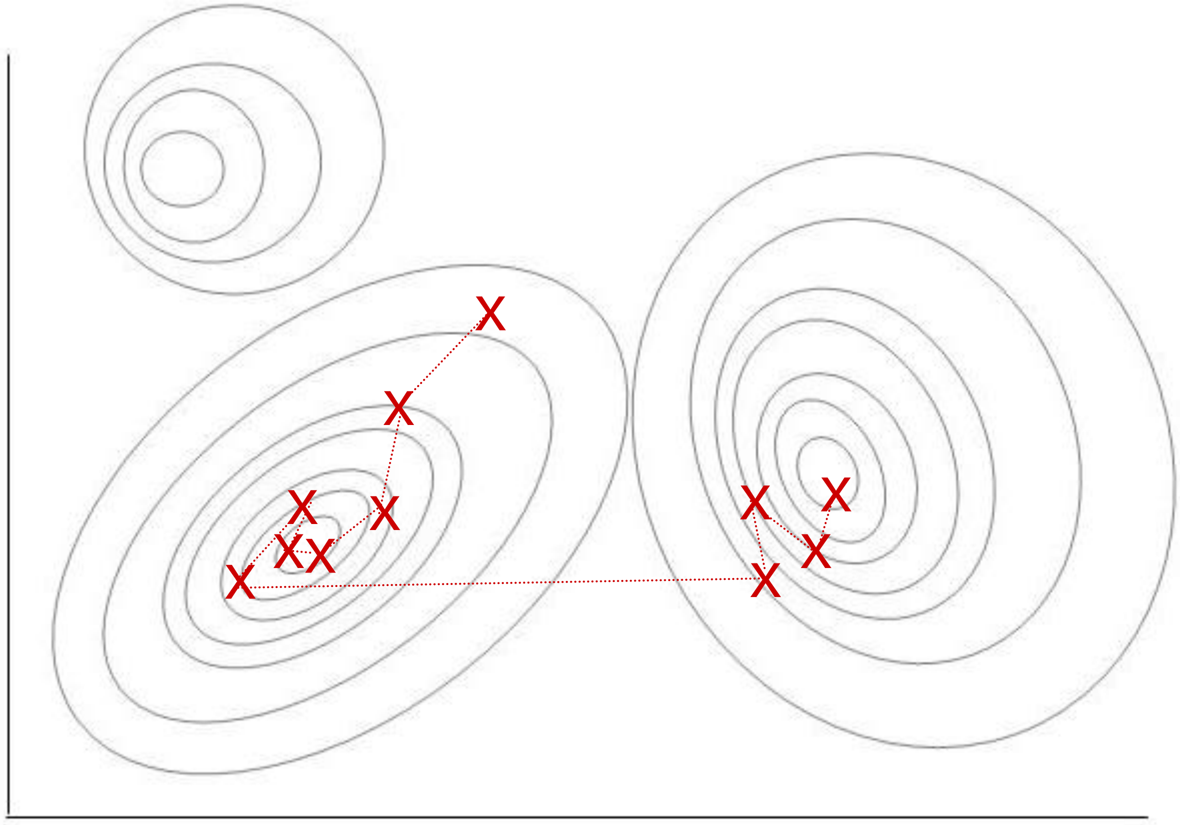

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

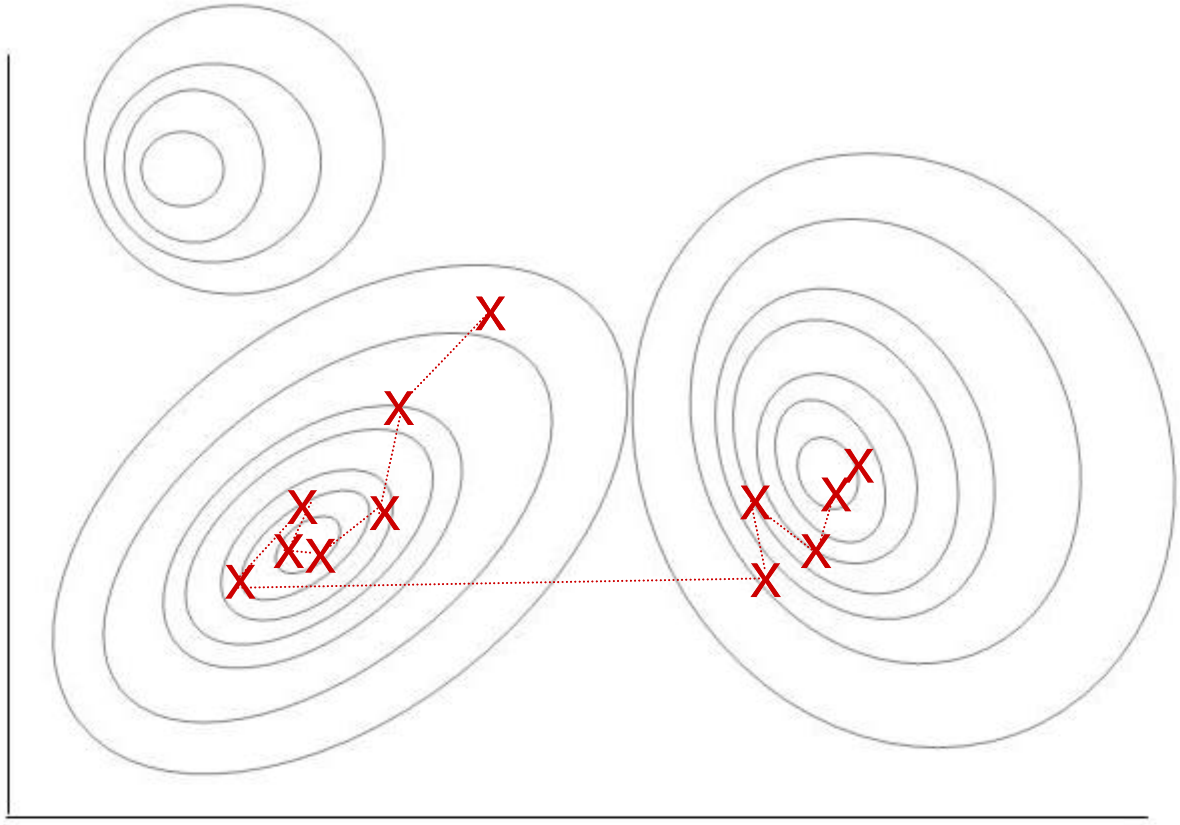

MCMC method for exploring human categories:

Algorithm efficiently explores regions of high representativeness

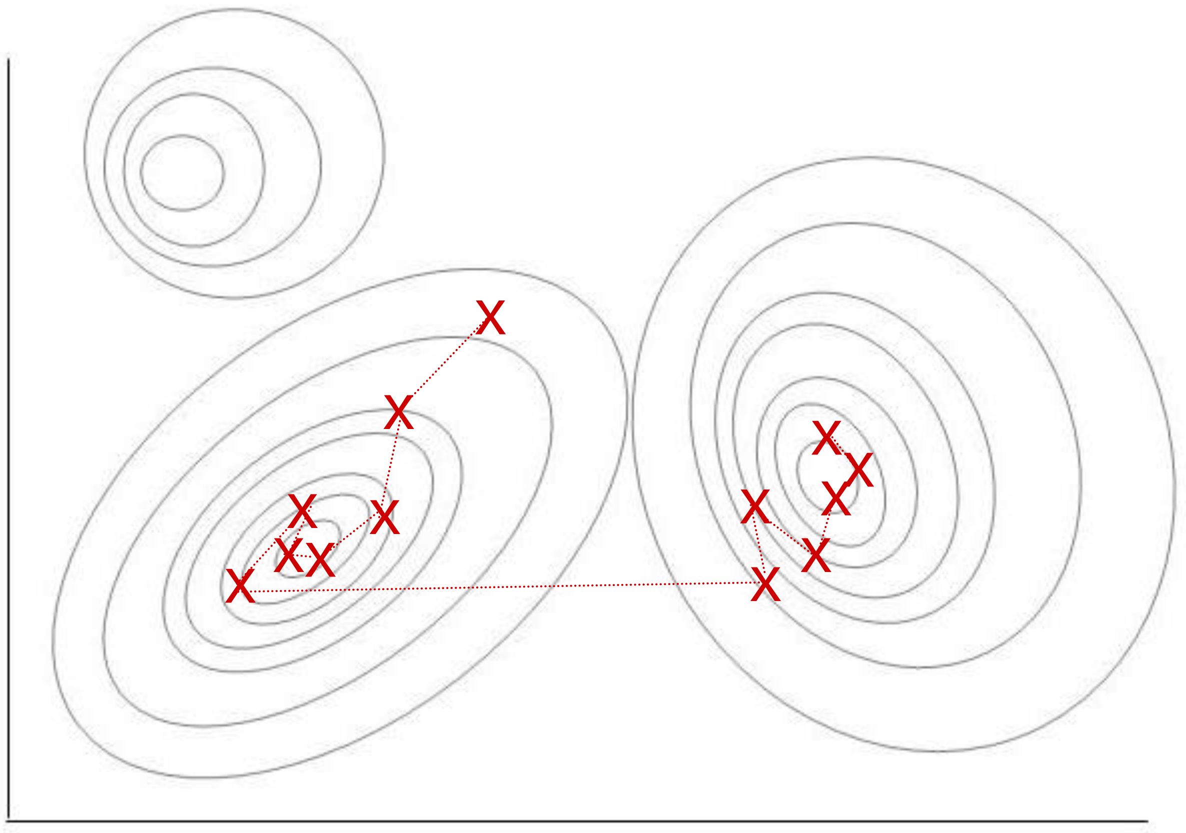

Using MCMC we can explore a wide variety of real life categories

Emotions: Which face is more Happy/Sad?

(Hsu, A., Martin, J., Sanborn, A., Griffiths, T. *submitted*)

. . .
set of 1000 faces

Happy

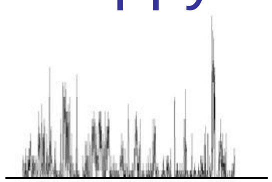

1000 faces

Ave. face

Sad

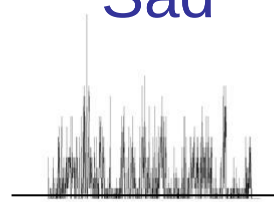

1000 faces

Ave. face

Using MCMC we can explore a wide variety of real life categories

Seasons: Which picture is more iconic of (e.g. spring)?

(Hsu, A., Martin, J., Sanborn, A., Griffiths, T. *submitted*)

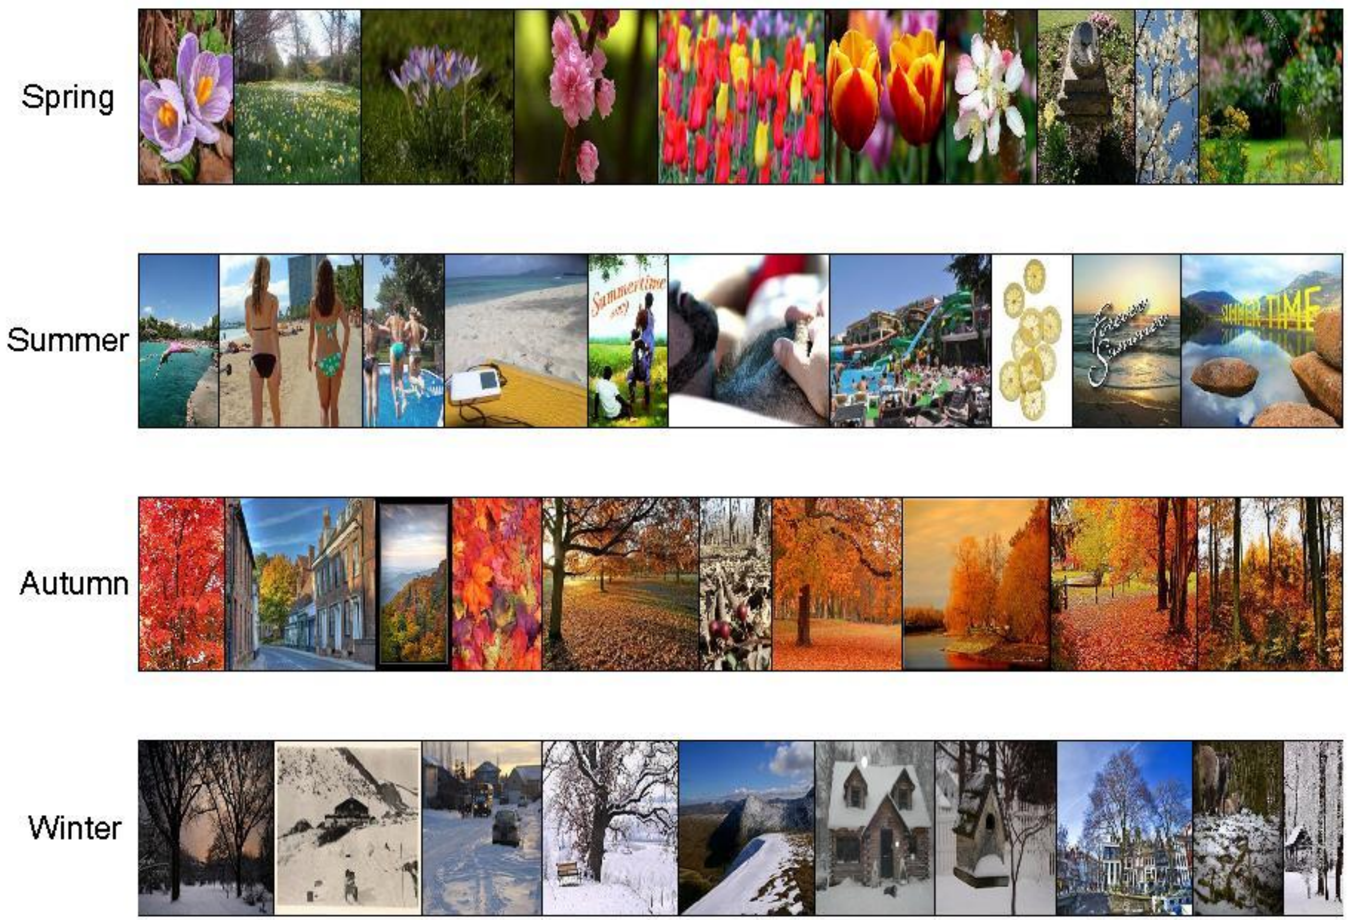

Using MCMC we can explore a wide variety of real life categories

Liberal vs. Conservative Morality: Which word is more relevant to morality?

(Hsu, A., Martin, J., Sanborn, A., Griffiths, T. *In preparation*)

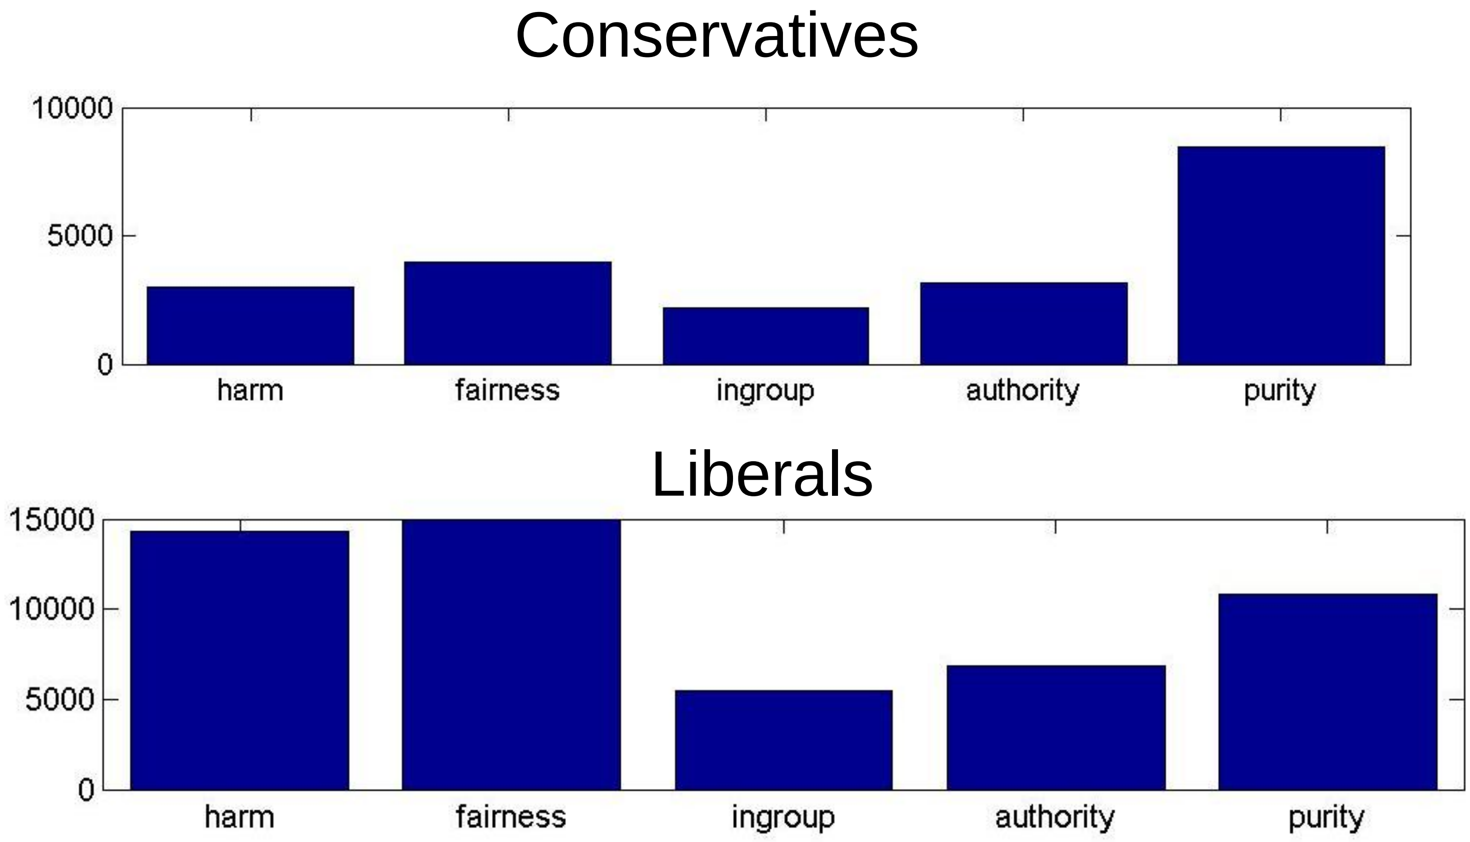

Results collapsed across 160 words pre-chosen to be relevant for morality

Using MCMC we can explore a wide variety of real life categories

Measuring Consumer preferences: Which wine is preferred?

(Hsu, A., Coenen, A., Lewis, R., Cheung, B. *submitted*)

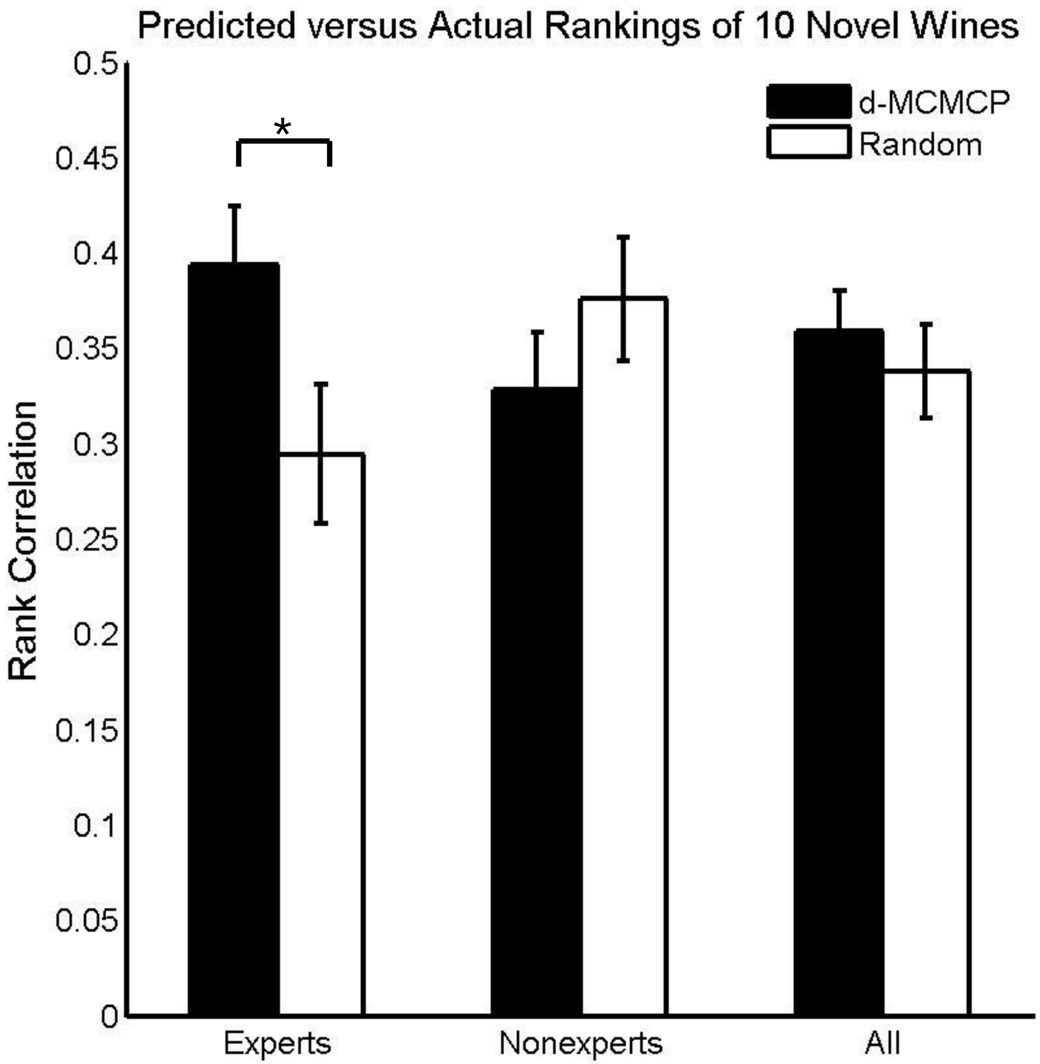

MCMC is able to measure preferences of experts significantly better when tested on predictions for novel wine preferences

* $p<0.05$

Using MCMC we can explore a wide variety of life categories.

Faces
- Emotions (Prof. Tom Griffiths; UC Berkeley)
- Trustworthiness (Costi Rezlescu; UCL)

Images
- Iconic images of cities, seasons, holidays (Prof. Tom Griffiths; UC Berkeley)

Legal reasoning
- Stereotypes of criminals (Dr D. Lagnado; UCL)

Consumer Preferences
- Preferences for product features (Anna Coenen, Dr. R. Lewis, Dr. B. Cheung; Decision Technology)

Medical decision making
- Hospital choice (Dr I. Vlaev, Dr D. King,Dr H. Lee; Imperial College)
- Surgical fracture diagnosis (Dr R. Emery, Dr I. Vlaeg; Imperial)

Thank you to:

Tom Griffiths

Jay Martin

Adam Sanborn