Model-agnostic Fits for Understanding Information Seeking Patterns in Humans

Soumya Chatterjee, Pradeep Shenoy

IIT Bombay

Google Research India

Modeling human decision-making under uncertainty

- A fundamental problem in cognitive science:
  - How do humans *seek and integrate information* while making decisions?
  - How do we capture deviations from optimality? individual variation?

- Applications in human-computer or human-AI interactions
  - Avoiding “echo chambers” in social media
  - Encouraging diversity in recommendations & media consumption
  - Building effective cognitive assistants, e.g., health & wellness tracking

Previous Modelling Approaches

- Propose model from *first principles* of learning/optimization, e.g., reinforcement learning [1]
  - tweak model to cover idiosyncrasies of specific task

- Propose *mechanistic models* that fit data, e.g., exponential weighting, or drift-diffusion models [2]
  - Describe the *process* without explaining the *goal*

- Challenges in these approaches
  - Explanatory power limited by *inductive bias* proposed by modeler
  - Need sufficient data for fitting models

[1] Vision: a computational investigation into the human representation and processing of visual information. Marr, D. 1982. ISBN 0716712849
[2] The diffusion decision model: Theory and data for two-choice decision tasks. Ratcliff, R.; and McKoon, G. 2008, Neural Computation 20(4)

Our work

- Contribution 1: avoid modeler-specified inductive biases
  - Deep learning approach for fitting human behavior
  - DNN architecture that reflects the structure of the task, but not the goals or rewards
  - Recover subjects' action policies purely by replicating behavior

Model agnostic

Captures behavioral biases

Our work

- Contribution 1: avoid modeler-specified inductive biases
    - Deep learning approach for fitting human behavior — Model agnostic
    - DNN architecture that reflects the structure of the task, but not the goals or rewards
    - Recover subjects' action policies purely by replicating behavior — Captures behavioral biases
- Contribution 2: handle extreme data paucity, for subject-specific models
    - Leverage large subject populations and shared parameters
    - Simultaneously learn population-level and individual-level model fits — Subject-specific fits with just 6 trials per subject !

Our work

- Contribution 1: avoid modeler-specified inductive biases
  - Deep learning approach for fitting human behavior

    Model agnostic
  - DNN architecture that reflects the structure of the task, but not the goals or rewards
  - Recover subjects' action policies purely by replicating behavior

    Captures behavioral biases
- Contribution 2: handle extreme data paucity, for subject-specific models
  - Leverage large subject populations and shared parameters
  - Simultaneously learn population-level and individual-level model fits

    Subject-specific fits with just 6 trials per subject !
- Non-goal for this talk: *explaining implicit policies* contained in DNNs
  - Significantly raise performance bar for alternate predictive/explanatory models
  - Future work: interpretable DNN policy learning

Behavioural Task: Sequential sampling & choice

Multi-step sampling task [1]:

- Goal: guess row with max product
  - Alternatively, row with min total
- At each step, subject chooses
  - *Sample* from allowed row OR
  - Guess final answer
- Cost for sampling; reward/penalty for guess correctness
- 32445 subjects with 1.2m trials (!)
- Open source

Figure adapted from [1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Behavioural Task: Sequential sampling & choice

Multi-step sampling task [1]:

- Goal: guess row with max product
  - Alternatively, row with min total
- At each step, subject chooses
  - *Sample* from allowed row OR
  - Guess final answer
- Cost for sampling; reward/penalty for guess correctness
- 32445 subjects with 1.2m trials (!)
- Open source

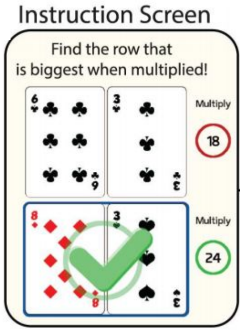

Figure adapted from [1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Behavioural Task: Sequential sampling & choice

Multi-step sampling task [1]:

- Goal: guess row with max product
  - Alternatively, row with min total
- At each step, subject chooses
  - *Sample* from allowed row OR
  - Guess final answer
- Cost for sampling; reward/penalty for guess correctness
- 32445 subjects with 1.2m trials (!)
- Open source

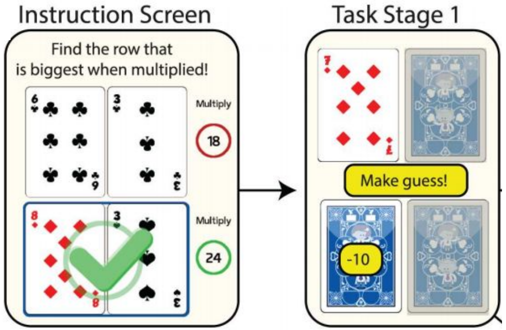

Figure adapted from [1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Behavioural Task: Sequential sampling & choice

Multi-step sampling task [1]:

- Goal: guess row with max product
  - Alternatively, row with min total
- At each step, subject chooses
  - *Sample* from allowed row OR
  - Guess final answer
- Cost for sampling; reward/penalty for guess correctness
- 32445 subjects with 1.2m trials (!)
- Open source

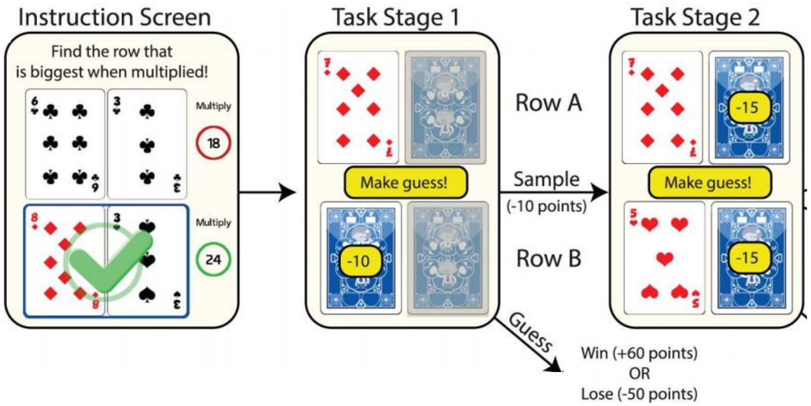

Figure adapted from [1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Behavioural Task: Sequential sampling & choice

Multi-step sampling task [1]:

- Goal: guess row with max product
  - Alternatively, row with min total
- At each step, subject chooses
  - *Sample* from allowed row OR
  - Guess final answer
- Cost for sampling; reward/penalty for guess correctness
- 32445 subjects with 1.2m trials (!)
- Open source

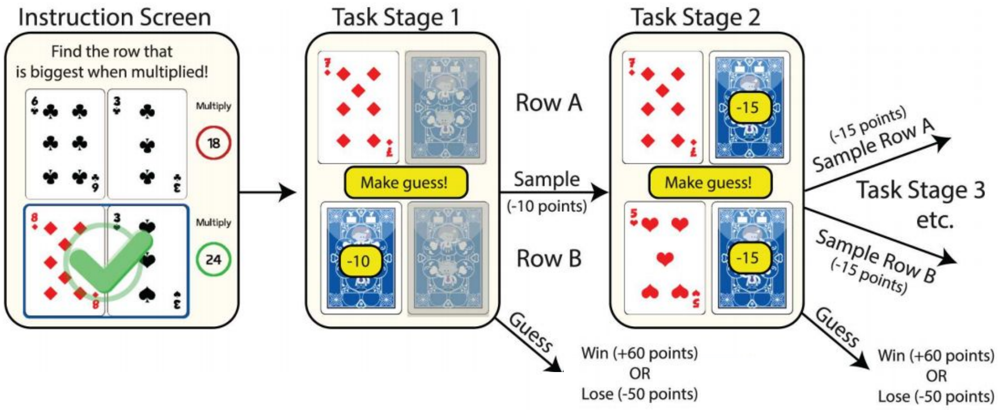

Figure adapted from [1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Proposed Models

- Baseline: Card Value Based Model [1]
  - Simple softmax heuristic model with hand crafted parameters (population level)
- DNN Population Model (Pop-DNN)
  - Cascaded DNNs fitted to behaviour from entire population
- DNN Subject Specific Model (Subj-DNN)
  - Subject specific embeddings
  - Other parameters shared across subjects
- Multiple tasks Model (Multi-DNN)
  - Multiple nets one for each task
  - Shared subject embeddings

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Proposed Models

- Baseline: Card Value Based Model [1]
  - Simple softmax heuristic model with hand crafted parameters (population level)
- DNN Population Model (Pop-DNN)
  - Cascaded DNNs fitted to behaviour from entire population
- DNN Subject Specific Model (Subj-DNN)
  - Subject specific embeddings
  - Other parameters shared across subjects
- Multiple tasks Model (Multi-DNN)
  - Multiple nets one for each task
  - Shared subject embeddings

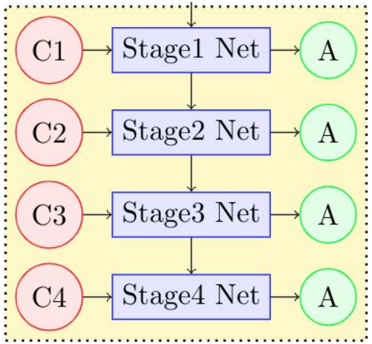

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Proposed Models

- Baseline: Card Value Based Model [1]
  - Simple softmax heuristic model with hand crafted parameters (population level)
- DNN Population Model (Pop-DNN)
  - Cascaded DNNs fitted to behaviour from entire population
- DNN Subject Specific Model (Subj-DNN)
  - Subject specific embeddings
  - Other parameters shared across subjects
- Multiple tasks Model (Multi-DNN)
  - Multiple nets one for each task
  - Shared subject embeddings

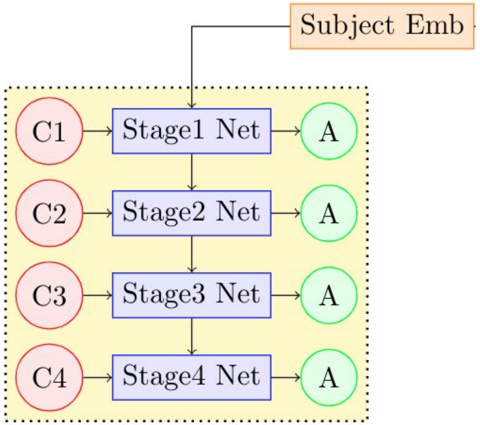

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Proposed Models

- Baseline: Card Value Based Model [1]
    - Simple softmax heuristic model with hand crafted parameters (population level)
- DNN Population Model (Pop-DNN)
    - Cascaded DNNs fitted to behaviour from entire population
- DNN Subject Specific Model (Subj-DNN)
    - Subject specific embeddings
    - Other parameters shared across subjects
- Multiple tasks Model (Multi-DNN)
    - Multiple nets one for each task
    - Shared subject embeddings

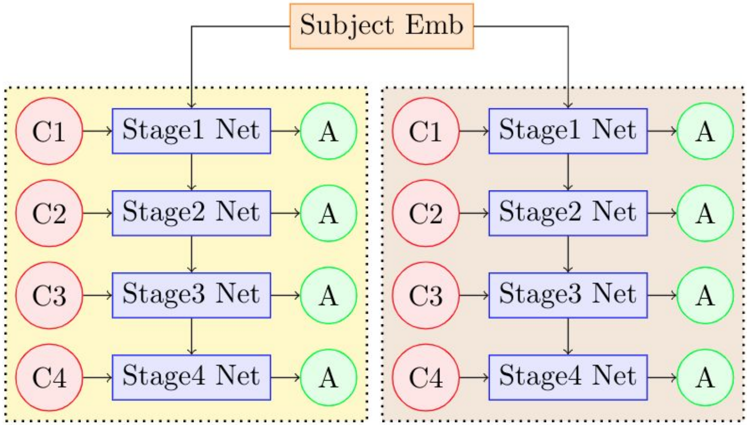

*some details omitted - see paper for details

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Evaluation of proposed approach

1. Does our model fit data better?

2. Does the model capture known biases at a population level?

3. Do subject embeddings capture individual policy variations?

4. Does pooling data across subjects really help?

5. Can learned embeddings *generalize* beyond task?

Results 1a - Decision-making at population level

Better fits to data

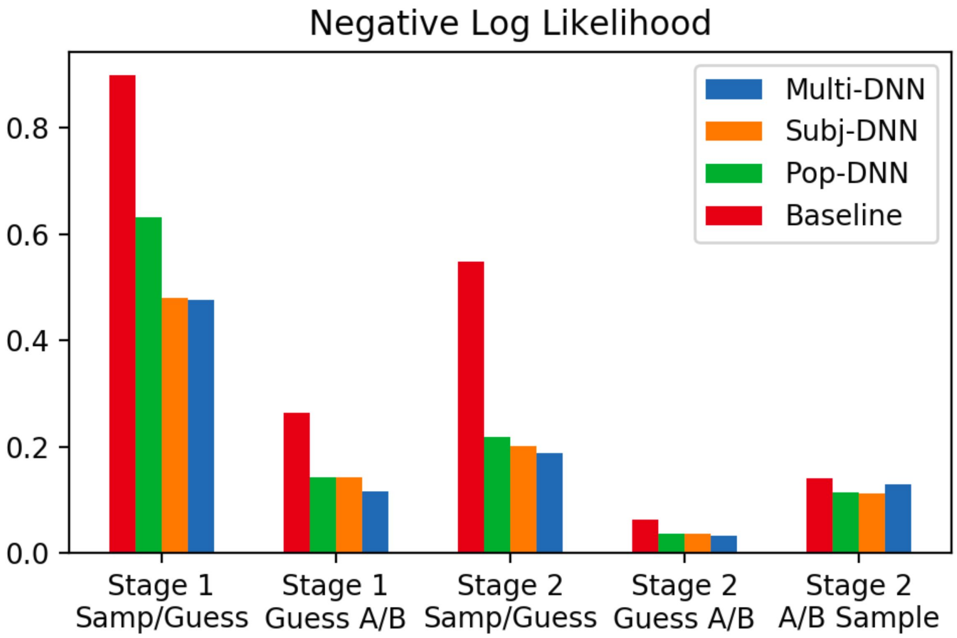

Both DNN & subject embeddings *improve fit significantly*

Results 1b - Model captures behavior variation

Model behavior correlates strongly with human behavior

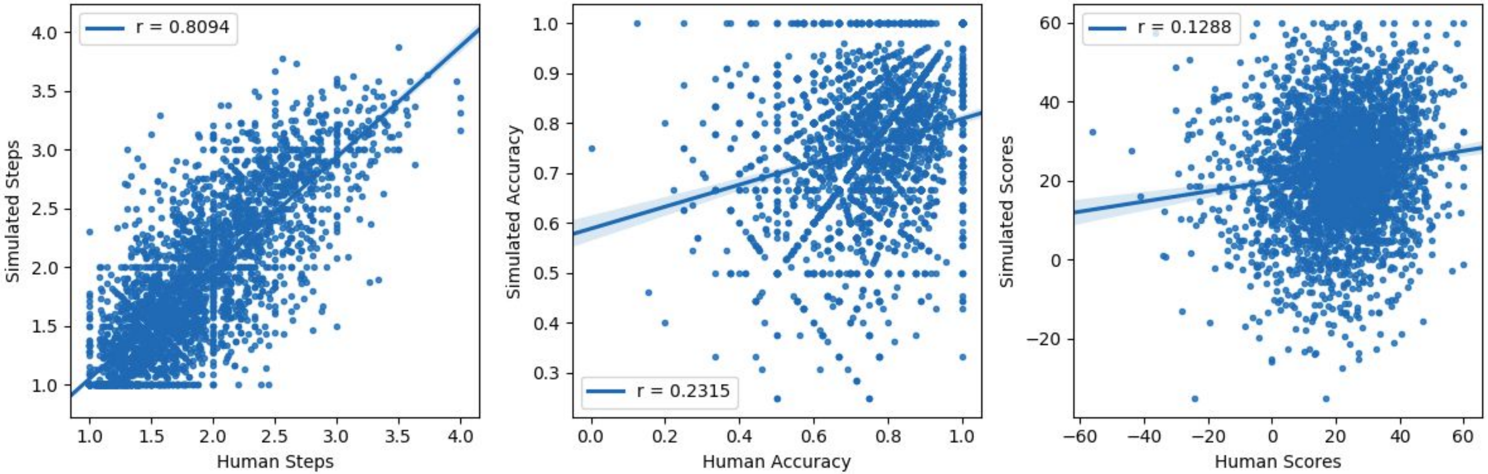

Multi-DNN simulation *significantly correlates with* human behavior ($p < 10^{-10}$)

- Pop-DNN (no subject embeddings) is *not correlated* with behavior

Evaluation of proposed approach

1. Does our model fit data better?
2. Does the model capture known biases at a population level?
3. Do subject embeddings capture individual policy variations?
4. Does pooling data across subjects really help?
5. Can learned embeddings *generalize* beyond task?

Results 2a - Model captures biases in behaviour

Approaching the positive bias [1]

Framing (MaxProd vs MinProd) influences whether to sample or guess

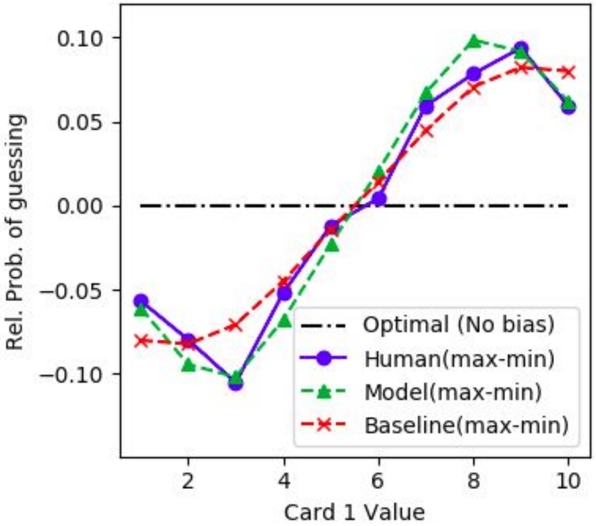

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Results 2a - Model captures biases in behaviour

Approaching the positive bias [1]

Framing (MaxProd vs MinProd) influences whether to sample or guess

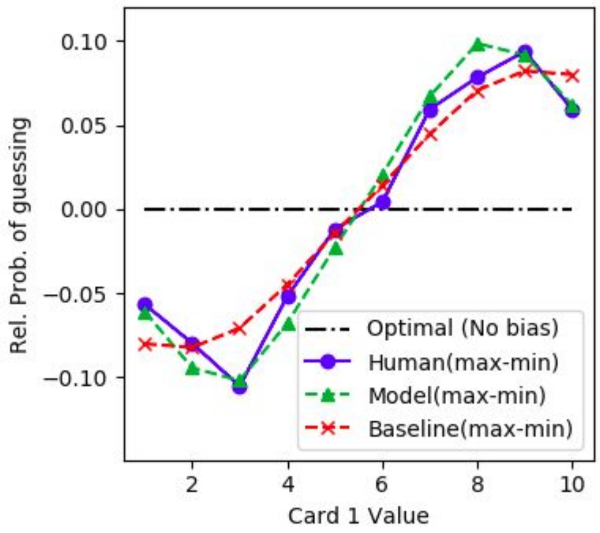

Rejecting the unsampled bias [1]

Subjects less likely to choose a row as answer if they had chosen not to sample from it

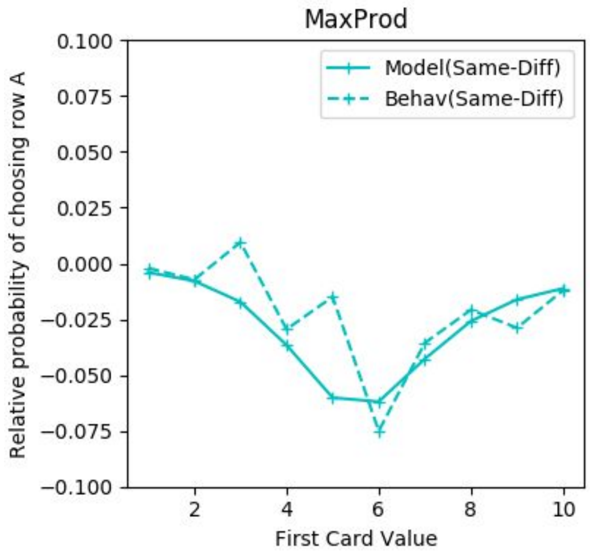

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Results 2b - Model captures biases in behaviour

Sampling the favourite bias [1]

- Humans choose to sample if offered from current favorite (confirmation bias?)

- Creates suboptimal asymmetry between “find MaxProd” & “find MinProd”

MaxProd

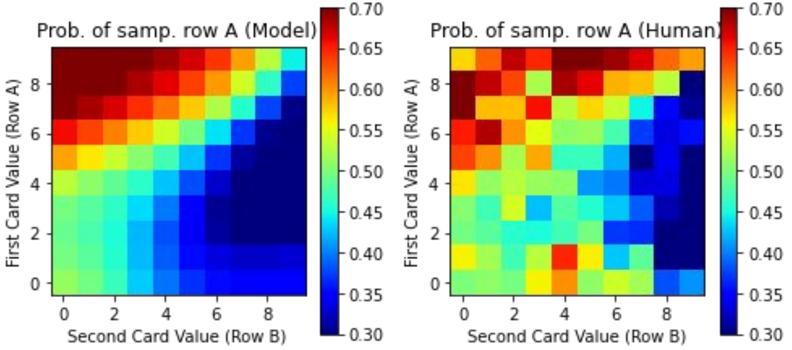

MinProd

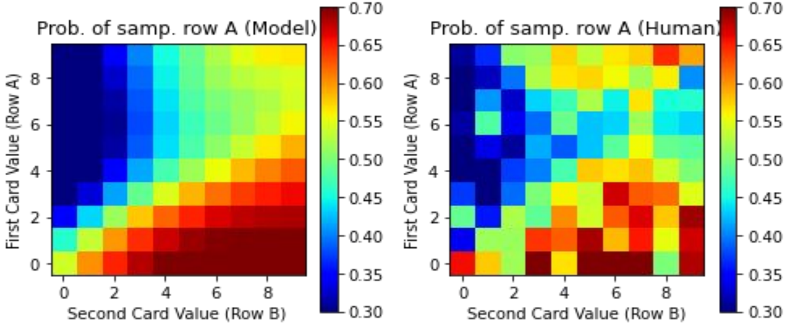

[1] Approach-Induced Biases in Human Information Sampling. Hunt et, al, 2016, PLOS Biology, 14(11)

Evaluation of proposed approach

1. Does our model fit data better?
2. Does the model capture known biases at a population level?
3. Do subject embeddings capture individual policy variations?
4. Does pooling data across subjects really help?
5. Can learned embeddings *generalize* beyond task?

Results 3a - Embeddings capture individual variation

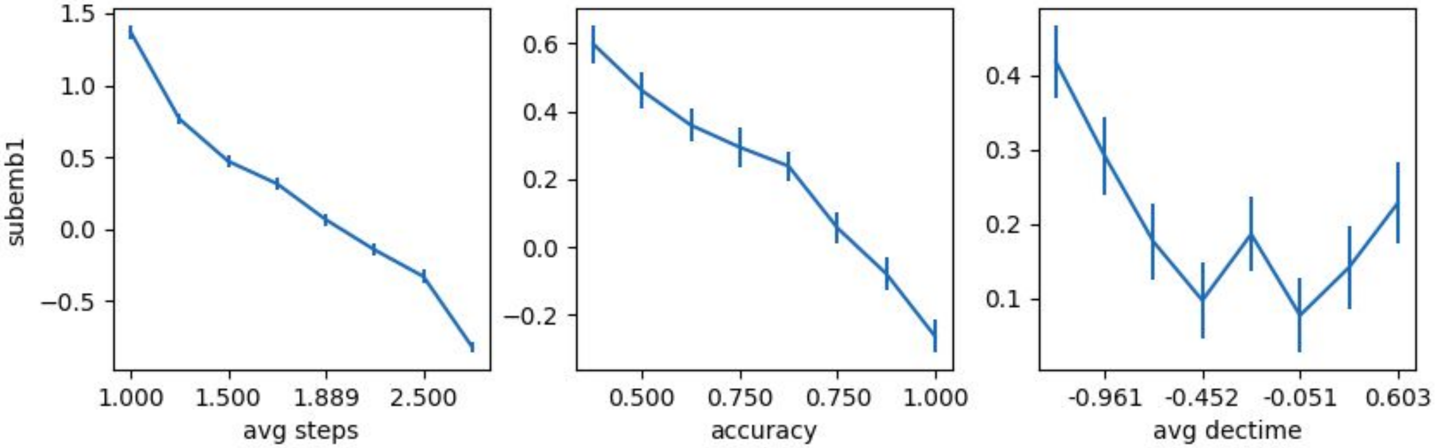

- We correlated the embedding dimensions with behavioral measures in the task
- Learned embeddings do contain information about subjects' performance
    - average embedding values are statistically different across buckets
- Decision time, typically related to subjective uncertainty about choice, is also captured
    - no access to this data during training

Results 3b - Discovering demographics from data

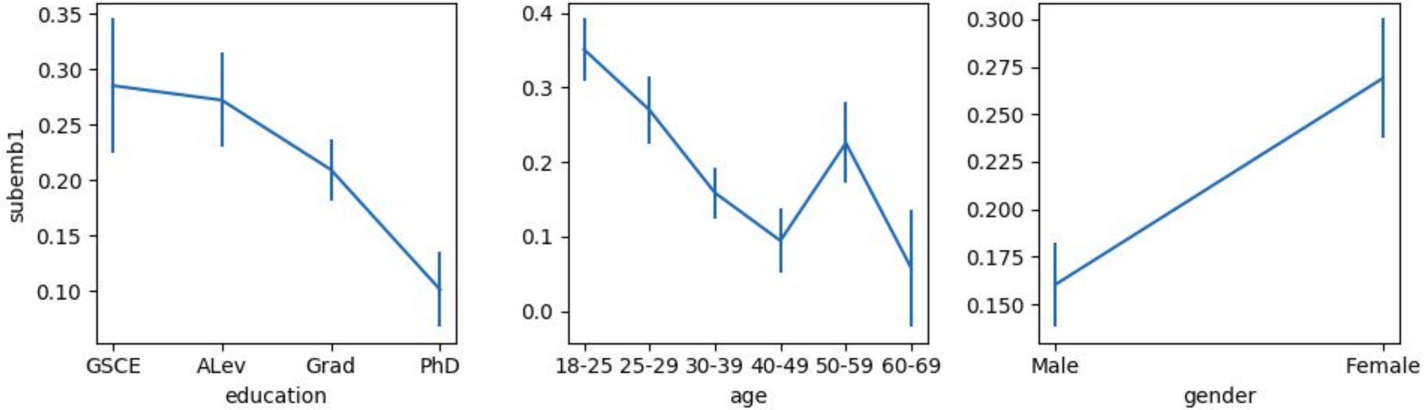

- We compare subject embeddings across education, age and gender
- Subject embeddings covary with education, age and gender
- Model never had access to any of this information
- Subject embeddings learn meaningful things

Evaluation of proposed approach

1. Does our model fit data better?
2. Does the model capture known biases at a population level?
3. Do subject embeddings capture individual policy variations?
4. Does pooling data across subjects really help?
5. Can learned embeddings *generalize* beyond task?

Results 4 - Sample Complexity

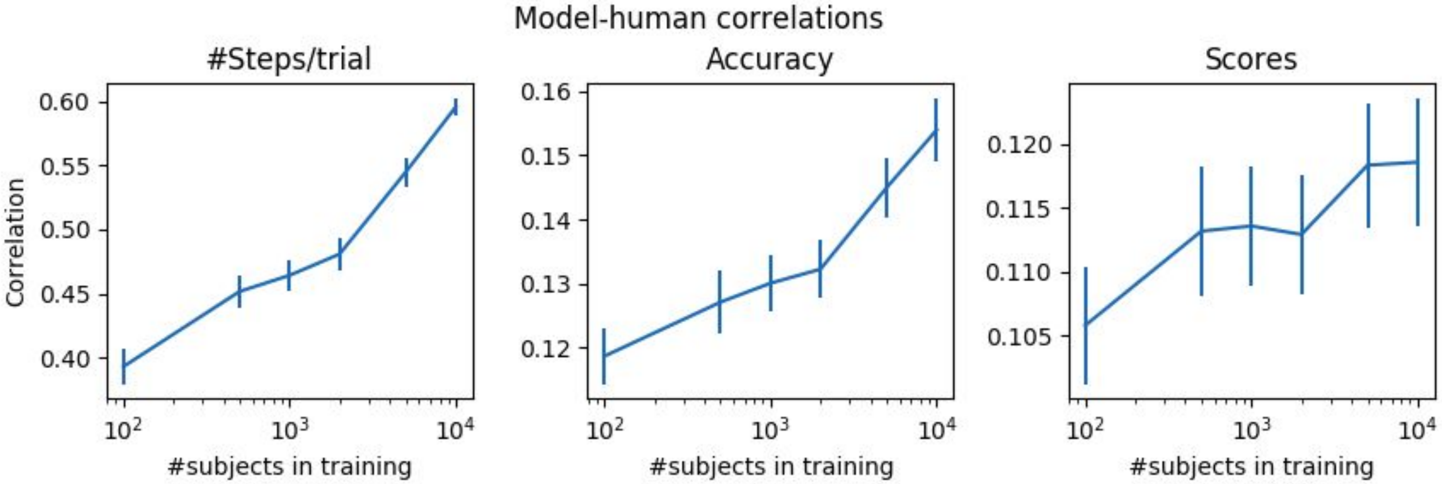

- Fix a group of test subjects (A). Add data from other subjects (additional subjects - B)
- Train on 6-7 trials per subject from A and B. Evaluate on remaining 4 trials for subjects in A
- *Increasing #subjects in B improves performance*
    - No additional data from group A
- Lack of per subject data compensated for by pooled training

Evaluation of proposed approach

1. Does our model fit data better?
2. Does the model capture known biases at a population level?
3. Do subject embeddings capture individual policy variations?
4. Does pooling data across subjects really help?
5. Can learned embeddings *generalize* beyond task?

Results 5 - Generalization to other tasks

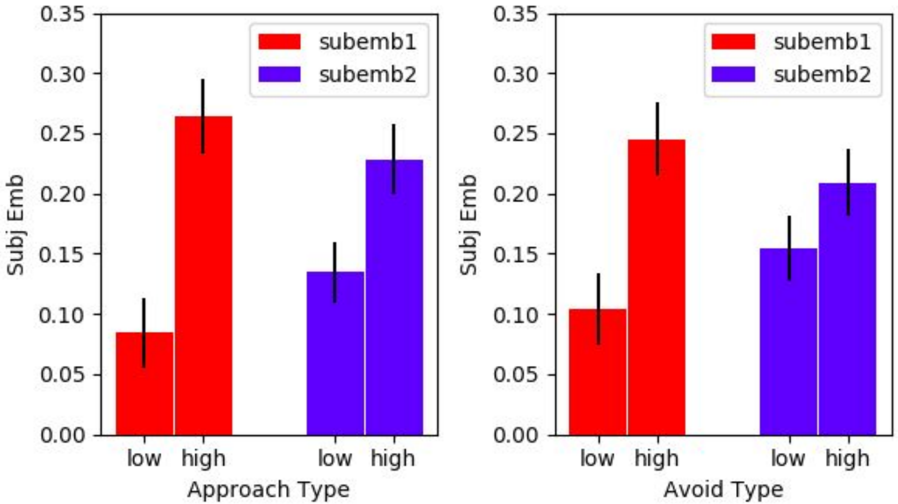

- Subject embeddings correlate with measures on secondary task
    - Approach-avoidance parameter measured on separate gambling-related task [1]
- Mean embedding value significantly different on low/high buckets of approach parameters

[1] Dopaminergic Modulation of Decision Making and Subjective Well-Being. Rutledge et, al, 2015, Journal of Neuroscience, 35 (27)

Conclusion

We presented a model-agnostic, multi-task approach for modeling human behavior in an information-seeking task.

Key contributions:

- High accuracy fits with sparse data, via pooled learning
- No assumptions about task goals or inductive biases
- Capture individual variation in the task, including biases
- Simple, low-dimensional *representation* of subjective parameters that generalize beyond current task

Thank You

Model-agnostic Fits for Understanding Information Seeking Patterns in Humans

Soumya Chatterjee and Pradeep Shenoy