

# Sequential Sampling Models of Adaptive Human Decision-Making (FA9550-11-1-0181)

PI: Michael Lee (UC Irvine)

Joachim Vandekerckhove (UC Irvine)

Senior Personnel: Shunan Zhang, James Pooley

## **AFOSR Program Review:**

**Mathematical and Computational Cognition Program**

**Computational and Machine Intelligence Program**

**Robust Decision Making in Human-System Interface Program**

**(Jan 28 – Feb 1, 2013, Washington, DC)**



# Adaptive Sequential Sampling (Lee)

## Objective:

Develop and evaluate new sequential sampling models that are adaptive

- within a single trial, optimizing decisions sensitive to time pressure
- change boundaries over sequences of trials
- allow for structured search

## Technical Approach:

Empirical data collection in series of experiments

New model development based on statistical and theoretical development

Evaluation of new and existing models using data

## DoD Benefit:

Better models of human and optimal decision-making in dynamic environment

- understanding, predicting and classifying human decision-makers
- automated adaptive and time-sensitive machine decision-making to emulate or support humans in tactical and strategic systems

## Budget:

Actual/  
Planned \$K

FY 11/12	FY 12/13
113	113

Annual Progress  
Report Submitted?

Y

N

Project End Date: 2013

## List of Project Goals


1. Collect empirical data to measure people's cue search and decision-making behavior in changing environments
2. Develop a self-regulating accumulator (SRA) model of decision-making suited to cue-based environments
3. Evaluate the SRA model, and traditional reinforcement learning models, against the human data
4. Develop sequential sampling models that optimize under time pressure and deadlines
5. Relate accumulator (race) and diffusion (random walk) sequential sampling models
6. Implement model inference using Approximate Bayesian computation, Synthetic Likelihoods

## Progress Towards Goals (or New Goals)

1. Collect empirical data to measure people's cue search and decision-making behavior in changing environments
2. Develop a self-regulating accumulator (SRA) model of decision-making suited to cue-based environments
3. Evaluate the SRA model, and traditional reinforcement learning models, against the human data
4. Develop sequential sampling models that optimize under time pressure and deadlines
5. Relate accumulator (race) and diffusion (random walk) sequential sampling models
6. Implement model inference using Approximate Bayesian Computation, Synthetic Likelihoods


# Sequential Sampling Models

- Gather evidence by drawing samples from an evidence distribution until a fixed critical level is reached for one decision or the other
  - Fixed threshold, consistent with Type I error optimality results
  - Samples are drawn iid, so there is no notion of search or environmental structure or change



# Sequential Sampling Models

- In most applications, trials are independent
  - Boundaries are not just constant with a decision, but over a sequence of decisions



- One over-arching goal of grant is to move beyond fixed boundaries in sequential sampling models of human decision-making
  - Within trials, optimize with respect to time-pressured optimization criteria
  - Between trials, adapt or regulate boundaries as environment changes
- Other (related) goal is to consider non-stationary evidence sampling


# Process Rationality of Heuristic Decision-Making



Lee, M.D., & Zhang, S. (2012). Evaluating the process coherence of take-the-best in structured environments. *Judgment and Decision Making*, 7, 360-372.

# Cue Based Decision-Making

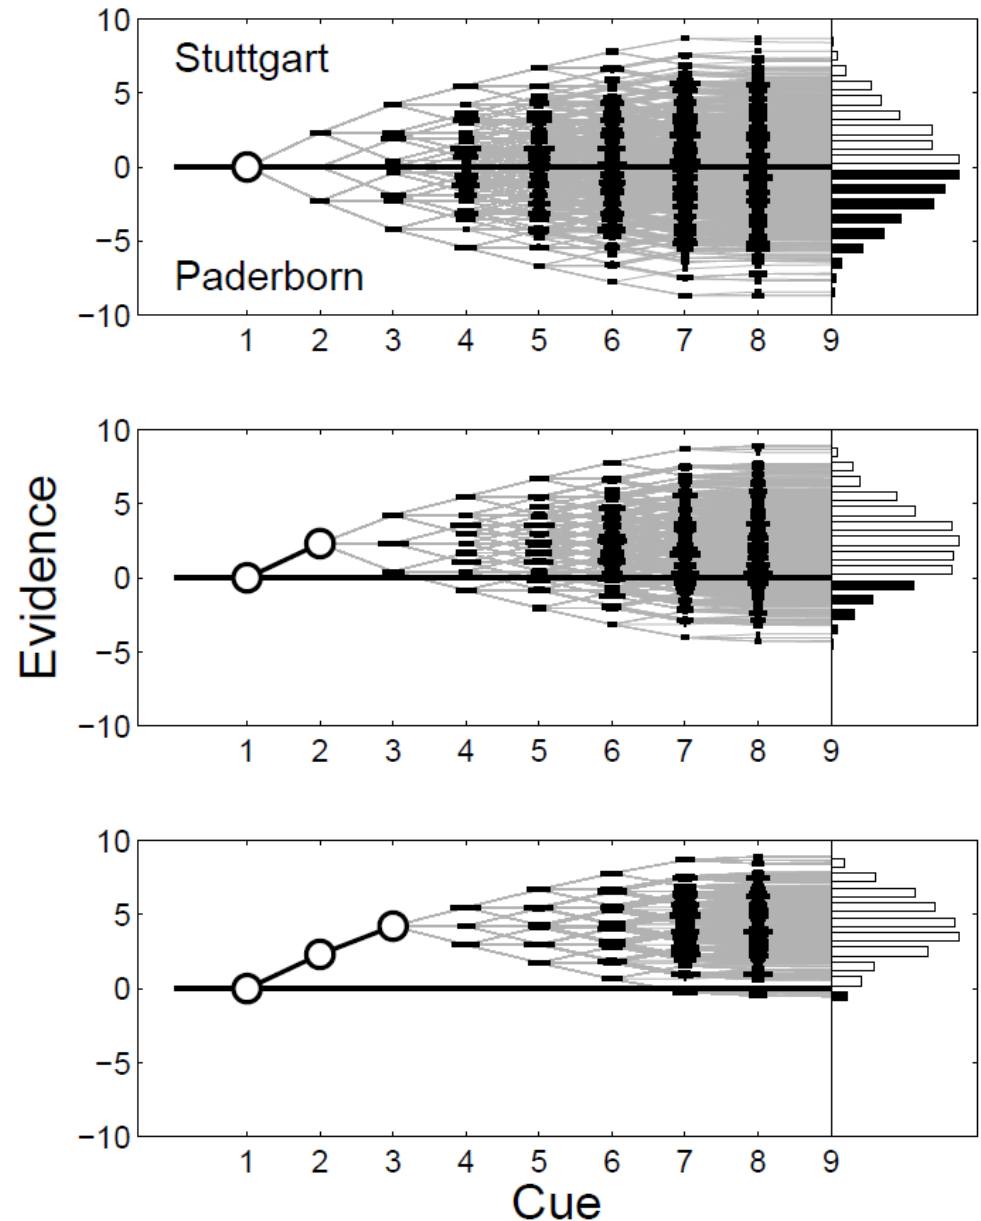
- One domain to study non-homogenous evidence samples and search is in cue-based decision-making
  - Which is Stuttgart or Paderborn is larger, based on cues like whether or not they have a soccer team in the Bundesliga
- Cues have different discriminabilities and validities, and so provide different evidence with different probabilities






# Process Coherence of Limited Search

- Gigerenzer et al study heuristic models like take-the-best with limited search, relying on environmental structure
  - Search in validity order, and stop once a discriminating cue is found
- We present a sequential sampling analysis, giving a rationale for limited search in terms of process coherence
  - Stop searching when answer cannot change



# When Limited Search Works

- We showed that limited search works when
  - search has diminishing returns, so later information is less important
  - the environment has a correlated structure, so that the first evidence is predictive of the rest




## Converging Boundaries Within A Trial



Zhang, S., Lee, M.D., Vandekerckhove, J., Maris, G., and Wagenmakers, E.-J. (submitted). On the relationship between diffusion and accumulator sequential sampling models.


# Diffusion and Accumulator Evidence Accrual

- Two extremes of evidence gathering
  - Diffusion models combine evidence in a single tally
  - Accumulation models gather evidence for each alternative in their own tally



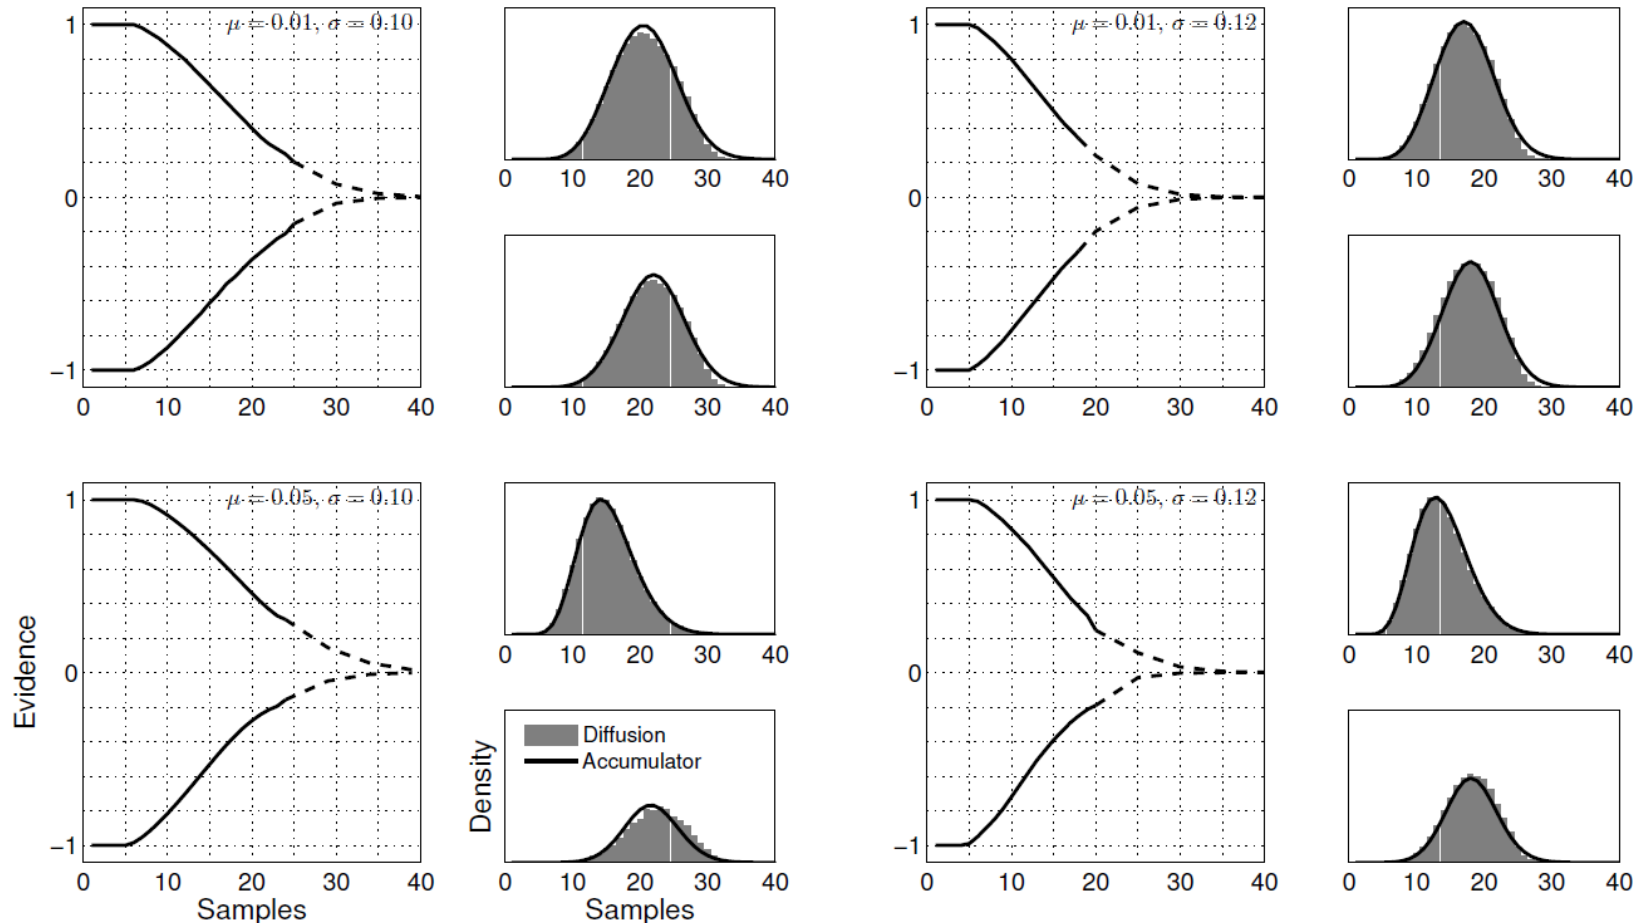
# Equating Diffusion and Accumulator Processes

- Accumulator distributions are matched by diffusion distributions with converging boundaries
  - Generate decision and response time data from a standard accumulator model
  - Find the boundaries that lead a diffusion process to match this behavior



# General Equivalence

- We have a proof that the equivalence can always be found, and an algorithm for finding the converging boundaries
  - Theoretical challenge in the finding that boundaries are often asymmetric



# Optimality Under Stochastic Deadlines



Zhang, S., Lee, M.D., & Wagenmakers, E.-J.  
(in preparation). Optimal diffusion boundaries  
under a class of stochastic deadlines.


# Optimality Under Stochastic Deadlines

- Constant boundaries in sequential sampling models are not consistent with the psychological constraint that most decisions must be made in a limited time
- Assume a loss function in which the goal is to gather as much information as possible before the deadline
  - Deadline is draw from a Gamma distribution
  - Penalty for exceeding the deadline is  $d$




# Optimal Boundaries

- Solve via dynamic programming
  - Similar to Frazier and Yu (2006), except we measure utility not in terms of accuracy, but information gathered
  - Narrow evidence distribution but broad deadline distribution, for different penalties  $\mathbf{d}$



# Interpreting Accumulator Models

- The resulting boundaries converge in a way that qualitatively matches the accumulator equivalence result
  - Gives one interpretation of what an accumulator model is optimizing in terms time-pressured decision making




# Adapting Search in Changing Environments



Lee, M.D., Newell, B.R., &  
Vandekerckhove, J. (in preparation).  
Reinforcement learning and self-regulating  
accumulator accounts of search in  
dynamic environments.

# Simple Search Task

- Must choose between two soil samples, on the basis of 9 binary cues, searched in decreasing validity order

	Sample A	Sample B		Choice
Actinium	Yes	Yes		<input type="button" value="A"/> <input type="button" value="B"/>
Radiation	No	No		
Promethium	Yes	No		<input type="button" value="Correct ?"/>
Carbon	No	Yes		
Gravimetric	No	Yes		
Seismic	?	?	<input type="button" value="Find Out"/>	
Europium	?	?		
Underground	?	?		
Microscopic	?	?		<input type="button" value="Trial 3 of 200"/>


# Simple Search Task

- Measure decision accuracy, and the 'Proportion of Extra Cues' (PEC) searched beyond the first discriminating cue


	Sample A	Sample B		Choice
Actinium	Yes	Yes		A B
Radiation	No	No		
Promethium	Yes	No	0/7	Correct ?
Carbon	No	Yes	1/7	
Gravimetric	No	Yes		
Seismic	?	?	Find Out	
Europium	?	?		
Underground	?	?		
Microscopic	?	?	7/7	Trial 3 of 200

# Non-Stationary Environment Task

- Subjects do 200 trials like this, but (without them being told) the environment changes twice




# Non-Stationary Environment Task



# Non-Stationary Environment Task


First discriminating cue gives no information, but full search gives answer







# Non-Stationary Environment Task

**Search to first discriminating cue gives answer, as does searching all cues**



# Individual and Overall Stopping Behavior


- Three blocks show, with individual differences
  - Limited search
  - Errors triggering more extended search
  - Return to more limited search, not triggered by errors




# Additional Experiments

- In two additional experiments, we encouraged limited search where possible
  - Experiment 2: Short time penalty (3s) for searching cues
  - Experiment 3: Monetary cost to searching cues

Experiment 2



Experiment 3




# Self-Regulating Accumulator

- Regulating a boundary means making **covert** decisions about whether to move it up or down,
  - Based on success of current decision-making
- But this is the same problem we already solved by using the sequential sampling model for **overt** decisions
- This is the basis for Vickers' (1979) self-regulating accumulator (SRA) model, which adapts boundaries to maintain a target level of confidence
  - Natural, elegant, parsimonious hierarchical structure, with four parameters
    - Target confidence
    - Twitchiness to adapt
    - Size of adaptation
    - Starting caution

# Model Evaluation

- Currently, control complexity by fitting parameters to first 100 trials, and generate predictions for remaining 100 trials
  - Tests generalization to a new environment change




# Simple Comparison Models




- Other attempts to address regulation have largely relied on reinforcement learning ideas
  - Myung and Busemeyer (1989), Busemeyer and Myung (1992), Erev (1998), Maddox and Bohill (1998, 2001), Rieskamp and Otto (2006), Simen (2006)
- Two comparison models, based on continual reinforcement driven by error, or error-and-effort signals
  - When an error is made, search some additional proportion of cues
  - When an error is made, search some additional proportion, and when no error is made, search some smaller proportion
- A final comparison model, a reduced form of the SRA model that adapts on every trial
  - Model-based test for the presence of lagged adaptation in human decision-making

# (Preliminary) Model Evaluation

- Order of generalization-only weighted mean-square-error of search and decisions









# Subjects Who Need SRA Explanations





# Subjects Who Need SRA Explanations



# Main Findings

- Adaptation of search is sensitive to environment change, and is not always error driven
- Some people are well described by simple RL, others need an account like SRA with latent monitoring and lagged and punctate regulation
  - Especially as search or monetary costs are introduced, and behavior becomes more careful and consistent, and the full models are more readily distinguished from their simple counterparts

# Interaction with Other Groups and Organizations

- Ben Newell's lab, University of New South Wales
  - Experimental design to measure search behavior
  - Links between sequential sampling models and fast and frugal heuristic models
- Eric-Jan Wagenmakers' lab, University of Amsterdam
  - Relationship between diffusion and accumulation processes
  - Stochastic deadlines in sequential sampling
- Scott Brown (Newcastle), Angela Yu (UCSD)

# List of Publications Attributed to the Grant

## ■ Accepted

- Lee, M.D., & Newell, B.R. (2011). Using hierarchical Bayesian methods to examine the tools of decision-making. *Judgment and Decision Making*, 6, 832-842.
- Lee, M.D., & Zhang, S. (2012). Evaluating the process coherence of take-the-best in structured environments. *Judgment and Decision Making*, 7, 360-372.
- Lee, M.D., & Pooley, J.P. (in press). Correcting the SIMPLE model of free recall. *Psychological Review*. Accepted 28-Aug-2012.

## ■ Submitted

- Zhang, S., Lee, M.D., Vandekerckhove, J., Maris, G., and Wagenmakers, E.-J. (submitted). On the relationship between diffusion and accumulator sequential sampling models.
- van Ravenzwaaij, D., Moore, C.P., Lee, M.D., & Newell, B.R. (submitted). Is take the best for the best? A hierarchical Bayesian modeling approach to searching and stopping.

## ■ In preparation


- Lee, M.D., Newell, B.R., & Vandekerckhove, J. (in preparation). Reinforcement learning and self-regulating accumulator accounts of search in dynamic environments.
- Zhang, S., Lee, M.D., & Wagenmakers, E.-J. (in preparation). Optimal diffusion boundaries under a class of stochastic deadlines.

**Thanks! Questions?**




# Accumulator Sequential Sampling

- Sample evidence for each alternative independently, until one or other reaches threshold
- Provides a good model of confidence, as 'balance of evidence' difference between totals
- Adapted to cue-based information environments







# Self-Regulating Accumulator

- Each boundary is now regulated by an ‘internal’ accumulator process (Vickers 1979, Vickers & Lee 1998)
- Gathers evidence of under- or over-confidence
- Enough evidence leads to an internal regulatory decision to make the threshold for the basic decision-making process






# Self-Regulating Accumulator





# Subjects Who Don't Need SRA Explanations



# Subjects Who Don't Need SRA Explanations

