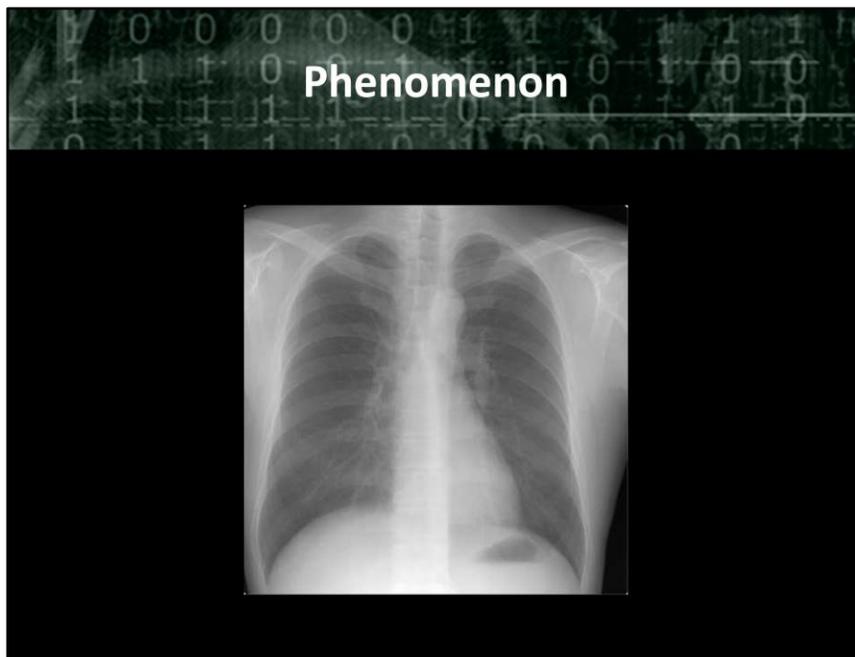




Feedback-Driven Perceptual Learning

Granger Danger
Josh Moskowitz
Daniel Lametti
Xiang Mou

Decision Making is an intimate part of our every day life. Making accurate decisions in a timely manner can be the difference between life and death for an animal that has to decide whether a movement in the bushes is just the wind or a predator about to have his next meal



How does experience relate to perceptual abilities?

When examining an X-ray for disease a cardiologist must quickly and accurately recognize whether there is disease in the patient. A study done in the late 80s found that more senior cardiologists were better at recognizing differences in abnormal xrays when compared to junior doctors and non-doctors, suggesting they gained the ability to recognize subtle differences in abnormal scans that less experienced doctors could not see. This learning occurred over a long period of time where these doctors were given feedback as to whether the features they saw in the scans indicated disease. Doctors are rewarded for correctly recognizing disease in patient and face serious repercussions for not doing so.

Phenomenom



We are not talking about the kind of these kinds of higher-level decisions made with prolonged deliberation. Slower decisions with even a mean RT greater than 1.5 s can induce multiple evidence accumulation processes, which are difficult to impossible to model given standard decision theory models.

Question

How does performance in a 2AFC perceptual judgement task change over time with feedback?

Two-Alternative Forced Choice

What changes occur in perceptual discrimination over time due to learning the task demands?

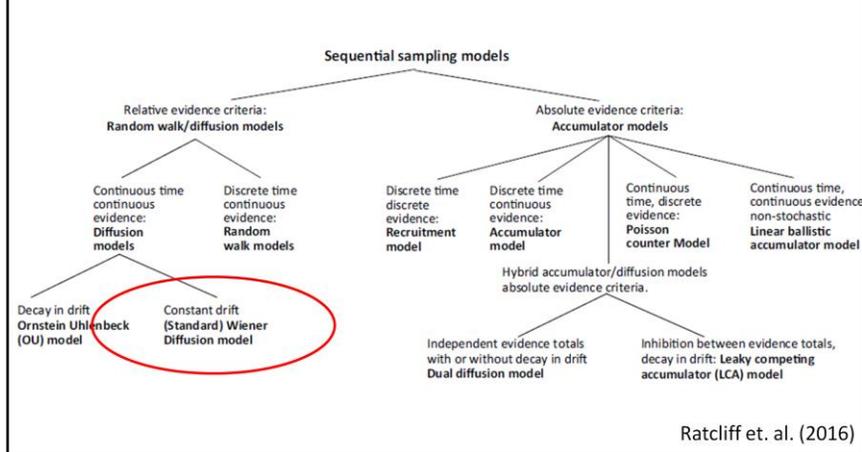
- changes in accuracy and reaction time

How are these changes in learning moderated by feedback

- how does feedback influence the accuracy and reaction time

Background

Evidence accumulation models



These models are based on the premise that the representation of stimuli in the central nervous system is inherently variable or noisy and to make a decision about a stimulus, one must accumulate successive samples of this noisy stimulus representation until a criterion quantity of evidence is obtained.

2 General Classes

Relative Stopping Rule Models (Random Walk)

evidence in favor of one response is evidence against the other response

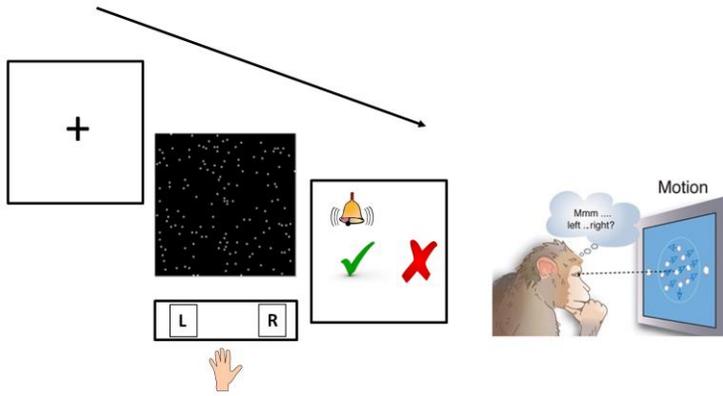
Absolute stopping Rule (Accumulator)

Evidence in favor of one response is accumulated by one counter at the same time as evidence for another response is accumulated in a different counter

When a counter reaches some threshold that response is selected

Called absolute because an increase in evidence for 1 response doesn't change the evidence for the other

Task

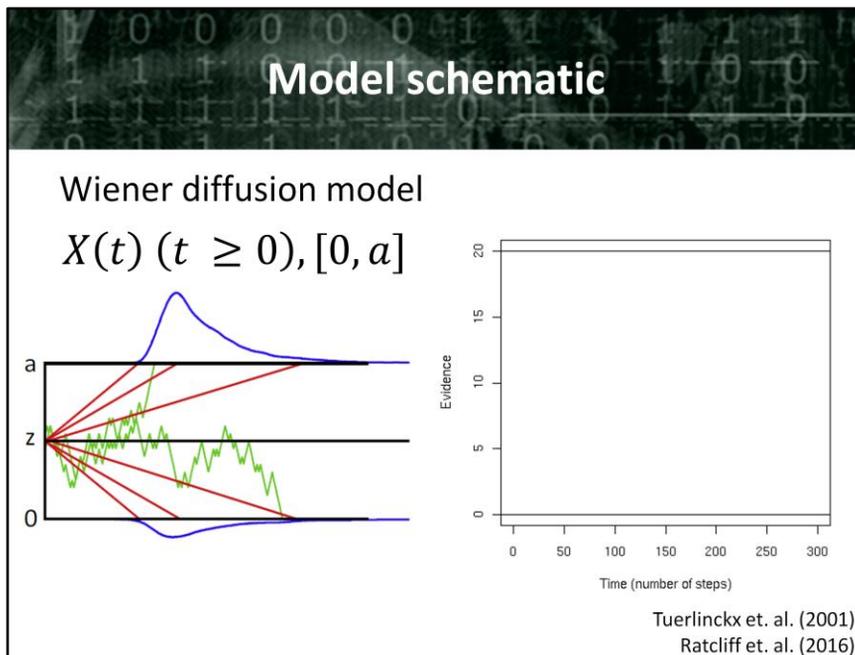


Hypotheses

1. Feedback will change the rate of learning on the task
 - A. Negative feedback should lead to higher RTs but higher accuracy
 - B. Positive feedback should lead to lower RTs and cause a temporary dip in accuracy
2. Learning should result in a general improvement in the speed-accuracy trade-off

If a subject encounters negative feedback (incorrect trial) then they will take more time on the subsequent trial to accumulate evidence

If a subject encounters positive feedback (correct trial) they will take less time on the subsequent trial to accumulate evidence



It is a stochastic process that can be represented by a continuous random variable, $X(t)$ denoting the position of the process in state space at some time t , restricted to the interval $0, a$, starting at z

Here in green are 2 simulated paths in the diffusion model

The blue curves represent the RT distributions for correct response (top) and incorrect response (bottom).

Red lines represent fastest, medium, and slowest responses from the RT distribution

Parameters and variables

Independent variables (input):

1. Number of Participants
2. Number of Trials
3. Learning Rate (gain factor)
4. Coherence Level

Dependent variables (output):

1. Reaction time
2. Correct/Incorrect choice

Model parameters

- Used DMAT Toolbox to generate artificial data (single trial RTs and accuracy)

	Symbol	Parameter	Interpretation
Decision process	a	Boundary separation	Speed-accuracy trade-off (high a means high accuracy)
	z	Starting point	Bias for either response ($z = a/2$ is neutral)
	v	Drift rate	Amount of input information; Quality of the stimulus
Nondecision	T_{cr}	Nondecision time	Sum of all other processes involved (motor RT, encoding . . .)
Intertrial variability	s_z	Intertrial range of z	Participant's variability in bias
	s_t	Intertrial SD of T_{cr}	Participant's variability in nondecision time
	η	Intertrial SD of v	Variability in stimulus quality, or variability in attention or motivation

Model
Constants

Vandekerckhove &
Tuerlinckx (2008)

Used a 3rd party MATLAB toolbox to establish the parameters we'll use to generate our data set.

Model equations

- Drift rate from distribution $N(\mu, \sigma)$
 - Mean will change on trial-by-trial basis
- 5 levels of Coherence represented by initial drift rates of 0, 0.05, 0.1, 0.15, 0.20
- Reaction time for a given trial:

$$RT = T_D + T_{er}$$

Reaction time is a combination of decision time and non-decision time

Model equations

Represent learning as a gain on the drift rate based on the history of correct/incorrect responses:

$$\text{Drift rate} = \text{Gain} * \text{baseline drift rate}_{20}$$

If response was correct:

$$\begin{aligned} \text{Success}_{\text{count}} &= \text{Success}_{\text{count}} + 1 \\ \text{Gain} &= -1 * e^{-\tau * \text{Success}_{\text{count}}} + 2 \end{aligned}$$

Else if response was incorrect:

$$\begin{aligned} \text{Success}_{\text{count}} &= \text{Success}_{\text{count}} - 1 \\ \text{Gain} &= -1 * e^{-\tau * \text{Success}_{\text{count}}} + 2 \end{aligned}$$

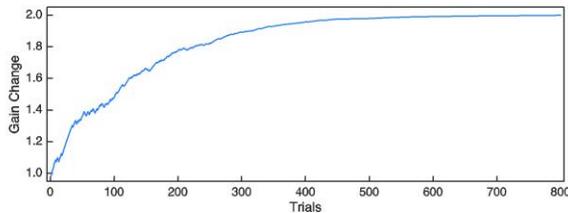
We don't ever change the baseline drift rate, because that relates to the level of perceptual discriminability (coherence level)

Instead we'll adjust the drift rate up/down along this exponential function, if participant was incorrect, the drift rate will go down, and RT should increase

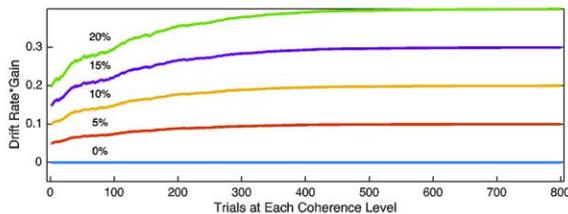
Tau is our learning rate, equal to 0.015 for our simulation, greater gain for larger values of this

Simulations / results

Change in drift rate
over the experiment
 $w/\tau = 0.015$



Greater learning
(increase in drift rate
relative to baseline) in
high coherence trials

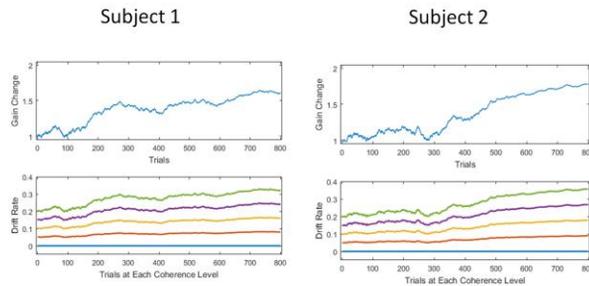


The figure on top shows the change in the gain formula G over time, across all coherence levels,

The figure on the bottom breaks down the change in drift rate, for each coherence level. For higher levels of coherence learning has more of an impact, for the 0% coherence, no learning could occur.

Simulations / results

We can sometimes see larger variations in drift rate, depending on how many incorrect/correct feedback trials occur in a row

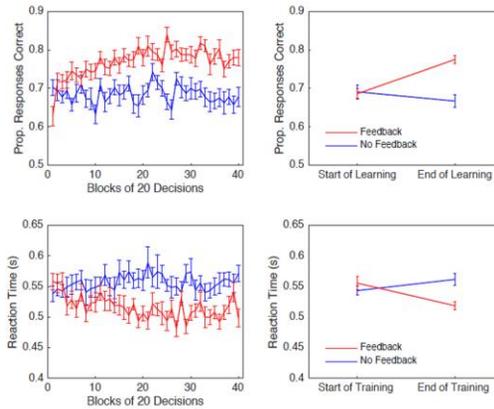


We ran additional simulations that show more inter-subject variability in changes in gain of the drift rate.

Also tried making the change in gain for negative feedback greater than for positive, and generally just increasing the changes in gain, but that resulted in RTs quickly reaching a floor or ceiling level

Simulations / results

- Time-Courses of Learning vs. No Learning



Start of learning is first 50 decision, end of learning is last 50 (trials 750-800)

Simulations / results

- Feedback results in improved accuracy and lower RTs compared to constant drift rate model

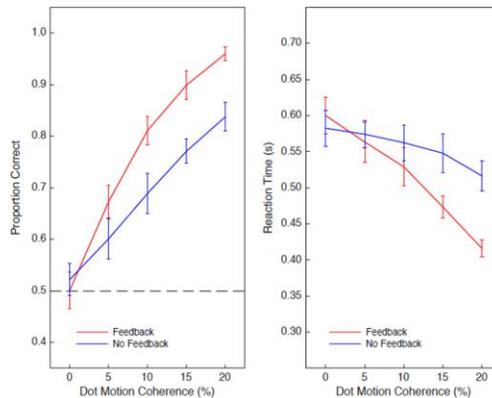
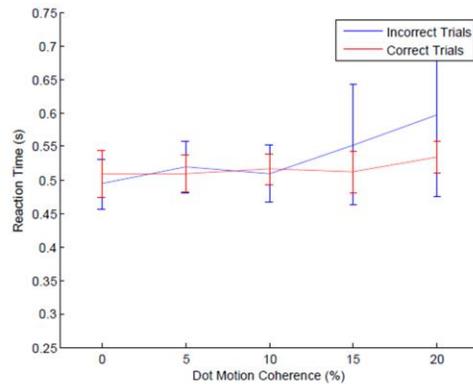


Figure on left shows in red the model we incorporated a gain on the drift rate, which was upgraded or downgraded depending on feedback.

RTs on the right. Data showcase improved accuracy and reduced RTs given feedback.

Model testing re. hypotheses

Performance on trial immediately after feedback



Making an incorrect response leads to higher RTs on the trial immediately following that incorrect trial

Making a correct response leads to lower RT on the trial immediately following

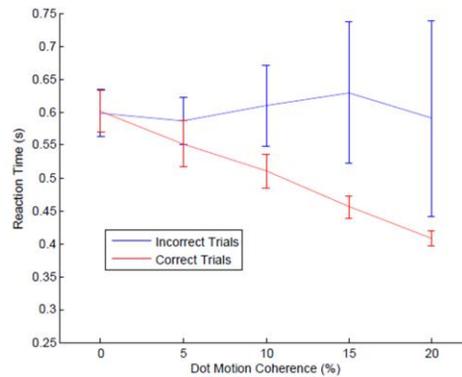
This is consistent with our hypothesis:

If a subject encounters negative feedback (incorrect trial) then they will take more time on the subsequent trial to accumulate evidence

If a subject encounters positive feedback (correct trial) they will take less time on the subsequent trial to accumulate evidence

Model testing re. hypotheses

RT for Correct Trials
consistent with
existing literature



Graph shows reaction time as a function of dot motion coherence for correct and incorrect trials.

Generally, as the difficulty of the stimuli increases, we'd expect RTs for correct responses to increase. And we see this relationship by the red line

Critical model evaluation

- Can model various learning rates associated with task difficulty
- Difficult to separate task difficulty from learning
 - Might need to modify alternative parameters such as threshold, and non-decision time
- Wanted to factor in the potential for reward to optimize a speed-accuracy trade-off

Consider alternatives such as nonstationary models of evidence accumulation, where there is a collapsing threshold, or the drift rate increases within a trial.

Summary & conclusions

- Feedback can have a beneficial effect on task performance
- Experimental data is crucial for setting initial parameters of a model and verifying model performance
- Might be worth sacrificing model accuracy for understandability

Feedback decreased task RTs and increased accuracy

Couldn't track changes in the state of the stochastic process as the algorithm used in the toolbox did not model the stochastic $X(t)$ path directly, rather it directly simulated the time it would take to reach a decision boundary