Learning Category Instances and Feature Utilities in a Feature-Selection Model

Michael Martin and Christian Lebiere
Carnegie Mellon University, Pittsburgh, PA, USA

MaryAnne Fields and Craig Lennon
Army Research Laboratory, Aberdeen, MD, USA

Presented at the 2016 ACT-R Post-Graduate Summer School, Lancaster, PA

Autonomous Systems

- Perceptual systems tend to feed forward to cognitive systems that provide little feedback
- Establish a feedback loop between perceptual and cognitive systems
- Exploit *cognitive context* to augment bottom-up perceptual approaches

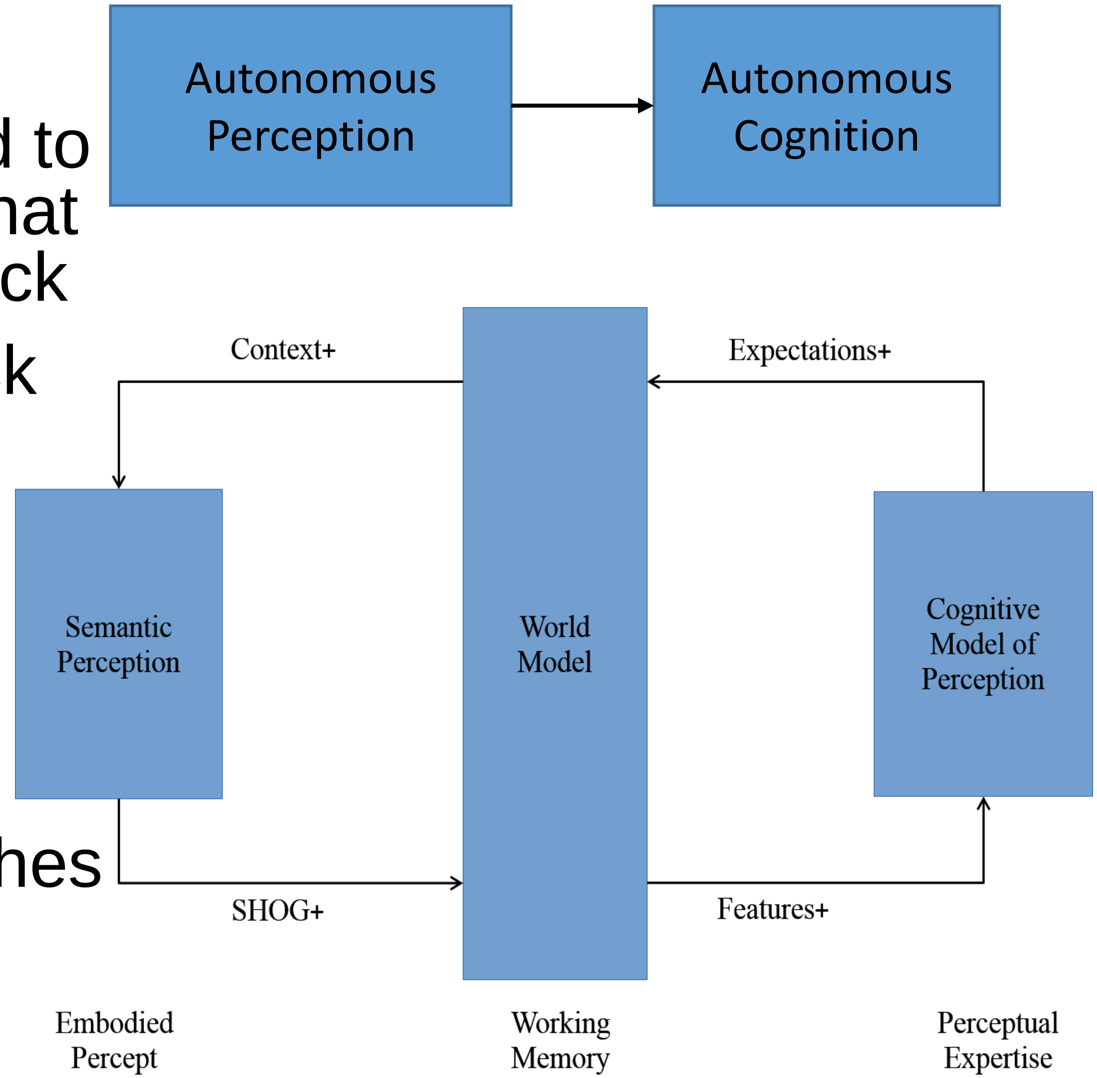

Intent

- ACT-R = a priori architecture of ML methods
  - Traditionally, ML methods tailored to problem ad-hoc
  - Classification & Feature Selection usually treated separately
- Create generic cognitive module for classifying objects, scenes or situations
  - Feature Selection Model (YACM)
  - FS Model = Prediction problem (which class) + Inference problem (important features)
- Design goals
  - Take input from any perceptual module
  - Learn to classify (Instance-Based Learning)
  - Learn which features to use (Feature Selection)
  - Stay close to architecture for generalizability
    - Respect default parameter values
    - Observe strict harvesting
  - Fixed representational scheme will not work
    - Flexible chunks in v7 make this easy

EGCM-RT (Lamberts, 2000) & EBRW (Nosofsky & Palmeri, 1997)

- Encode features one at a time until a category decision is made
  - Decision depends on similarity b/t S & instances in memory
  - S representation gradually constructed via info accumulation process
  - Accumulation stops as function of evidence for category membership
- Retrieved exemplars during info accumulation drive a random walk process
  - Implies a random walk pointer
  - Exemplars retrieved until pointer exceeds criterion value for 1 of classes

Feature Selection Model

- Combine IBL w/FS to anchor perceptual labels
- S information accumulated in imaginal buffer
- Make class inference for each feature encoded
- Random walk pointer & criterion value represented by relative utility of & competition among encoding & decision productions

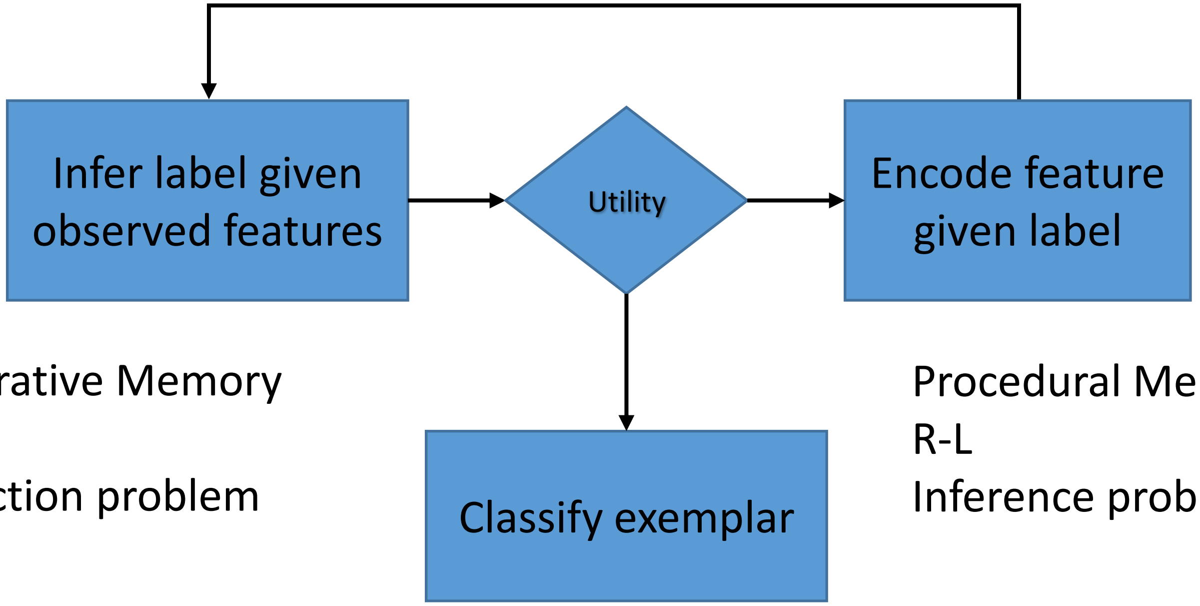

Declarative Memory
IBL
Prediction problem

Procedural Memory
R-L
Inference problem

Cohn-Kinade Facial Expressions

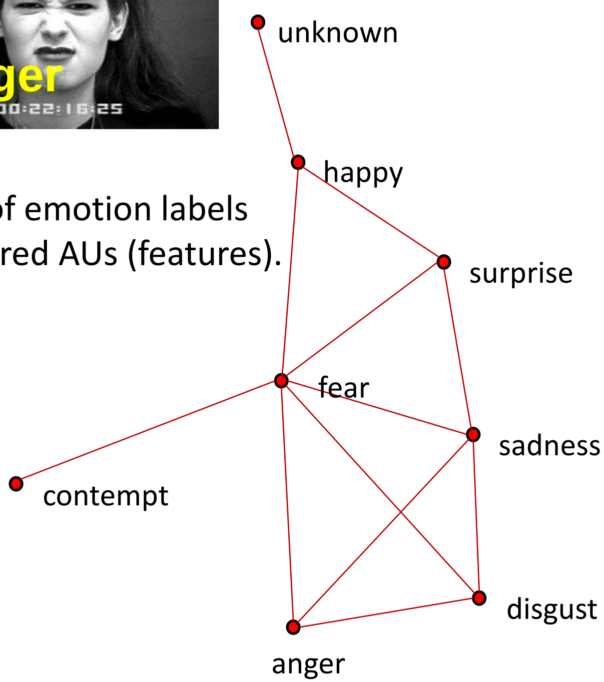

Similarity of emotion labels based on shared AUs (features).

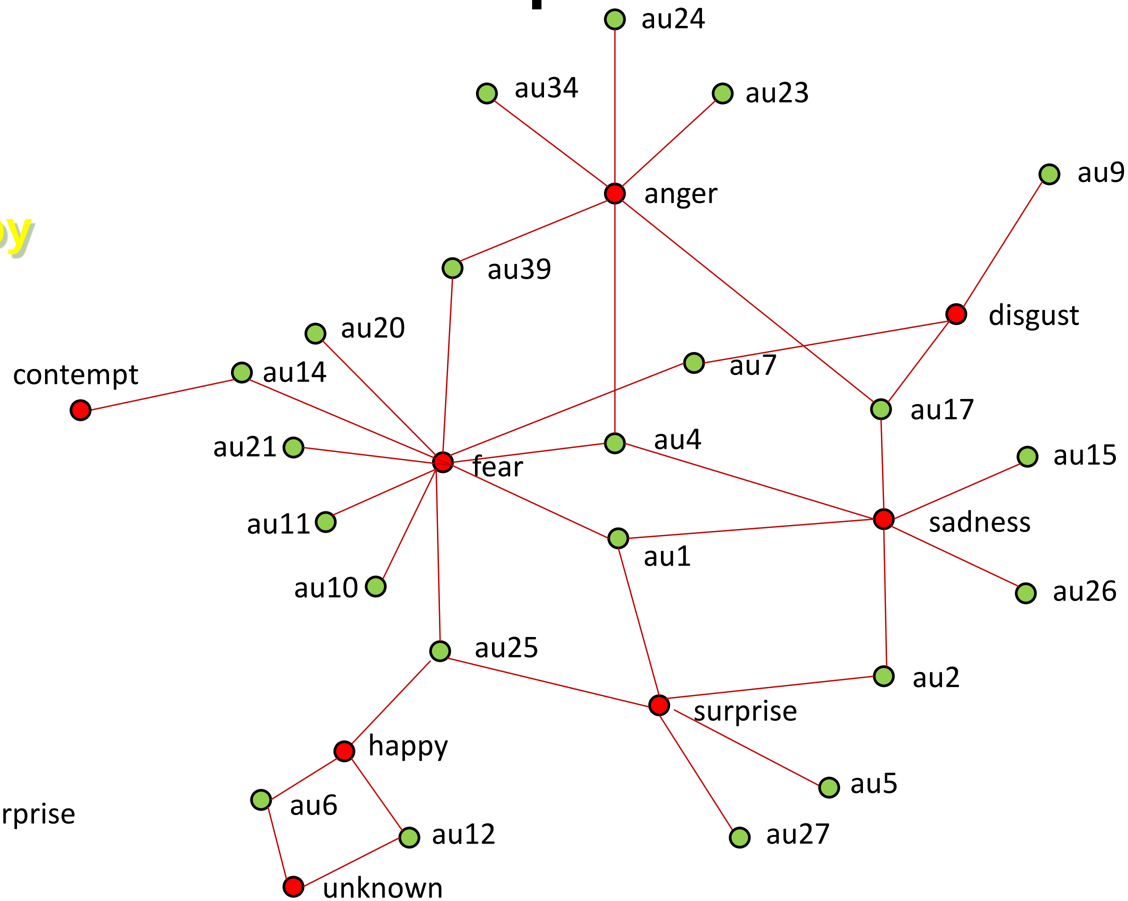

Emotion x Feature network indicates that emotion labels are associated with both shared- and distinct facial Action Units. Categorization depends on different subsets of overlapping and non-overlapping features.

Baseline: Retrieve Exemplars

That's 20% correct

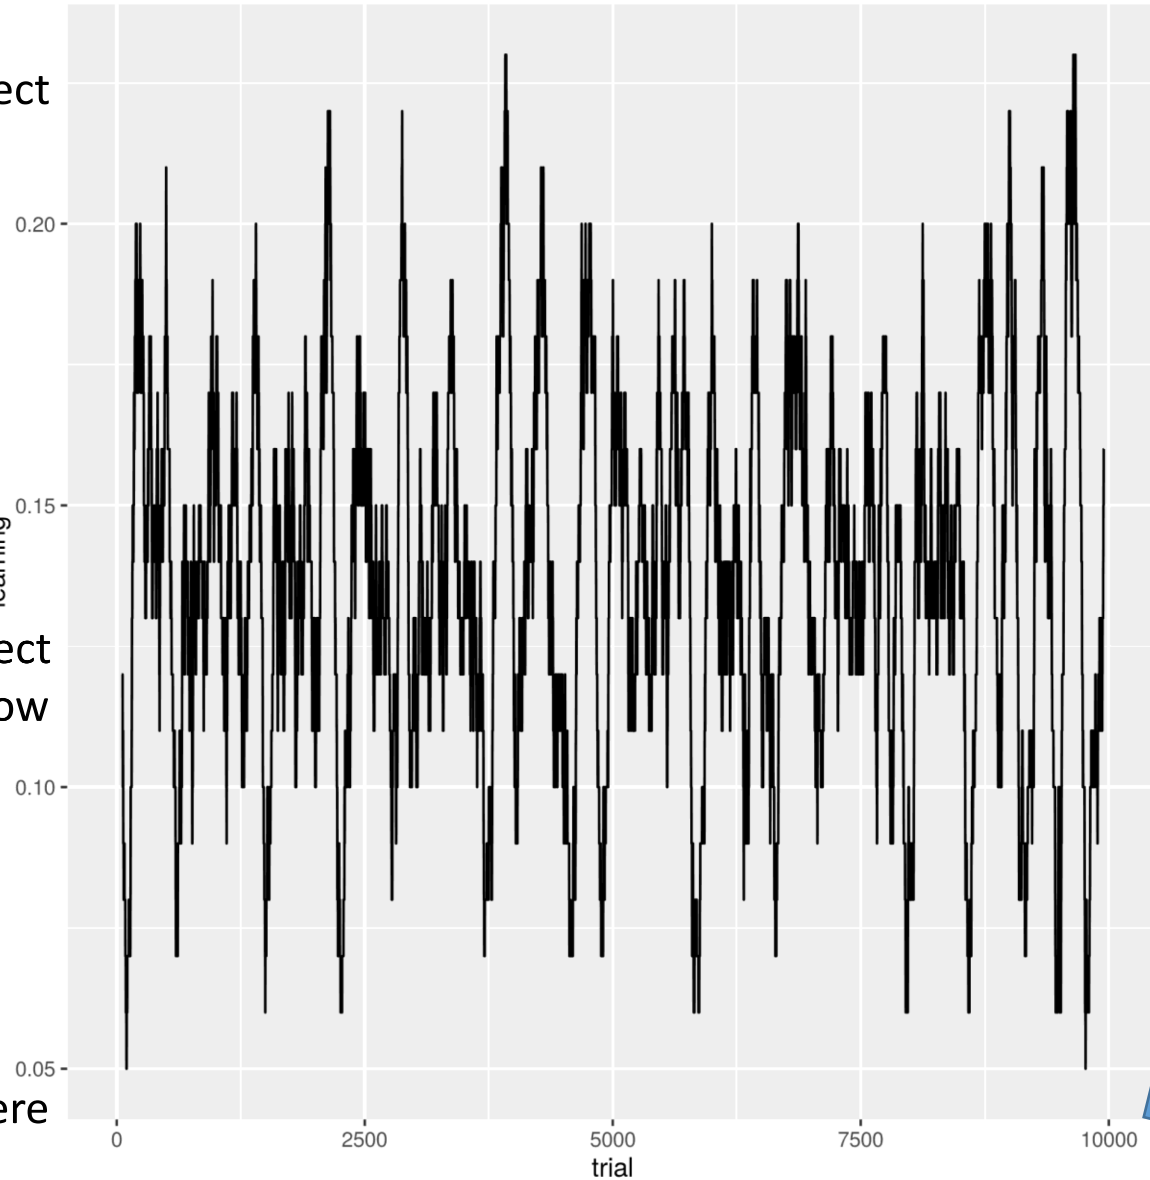

Proportion correct for sliding window 100 trials wide

Anger everywhere

<u>Activation</u>

:bll nil

:mas 10

:ans nil

<u>Utility</u>

Reward 10, -1

:alpha 0.2

:egs 0

10,000 trials

188 chunks

Retrieve Exemplars w/Noise

That's 25% correct

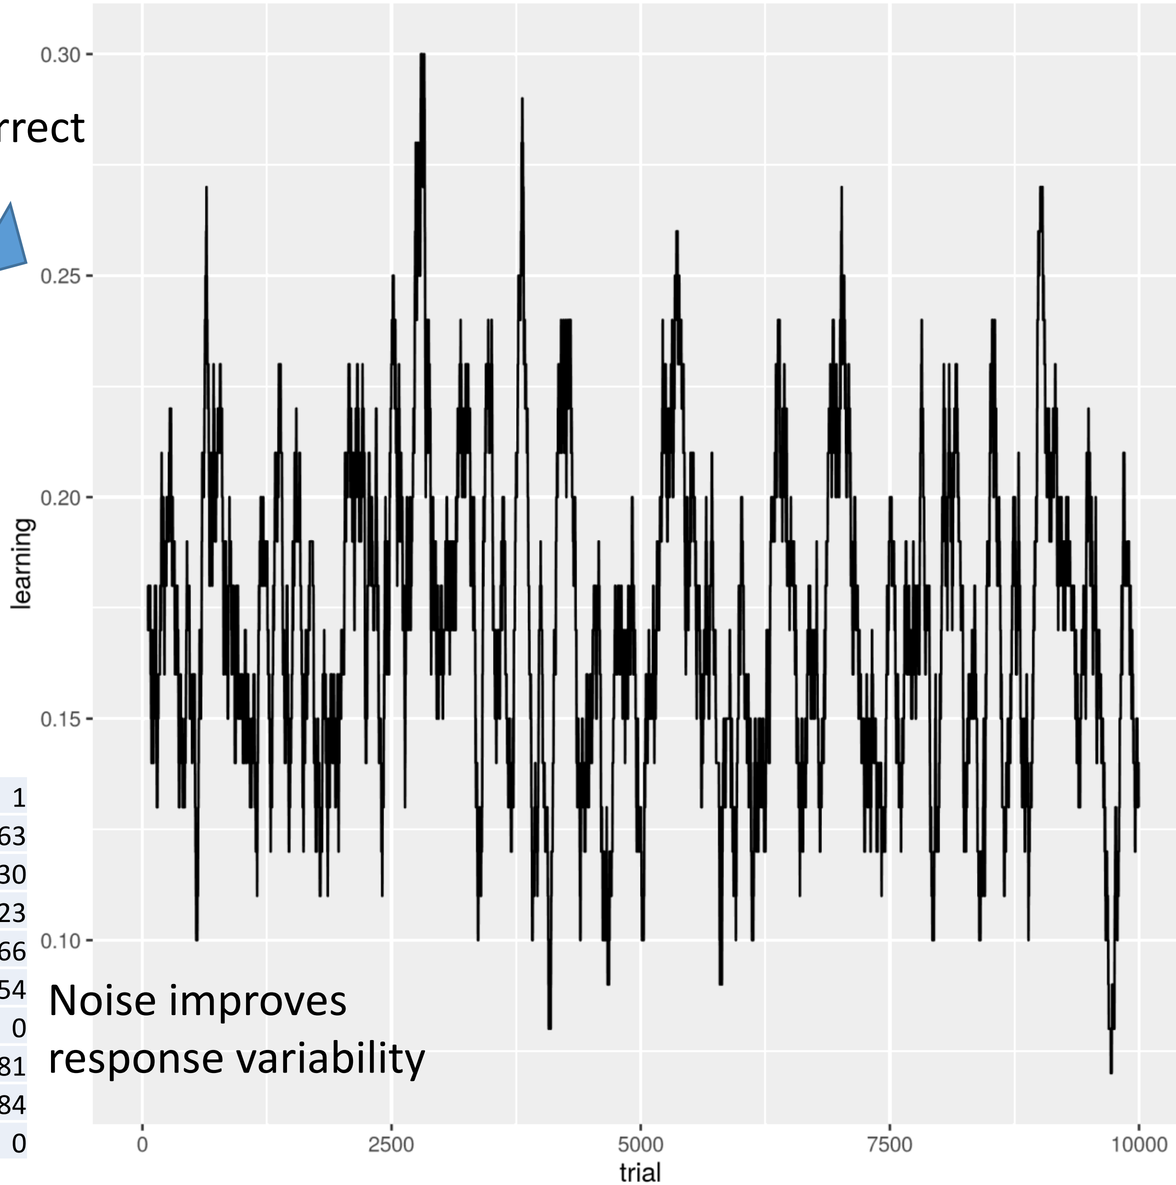

	0	1
anger	4266	663
contempt	646	30
disgust	1457	323
fear	795	66
happy	1533	454
neutral	10	0
sadness	892	81
surprise	1684	584
unknown	16	0

Noise improves response variability

<u>Activation</u>

:bll nil

:mas 10

:ans .25

<u>Utility</u>

Reward 10, -1

:alpha 0.2

:egs 1.0

3695 chunks

Retrieve Situated Prototype

That's 50% correct

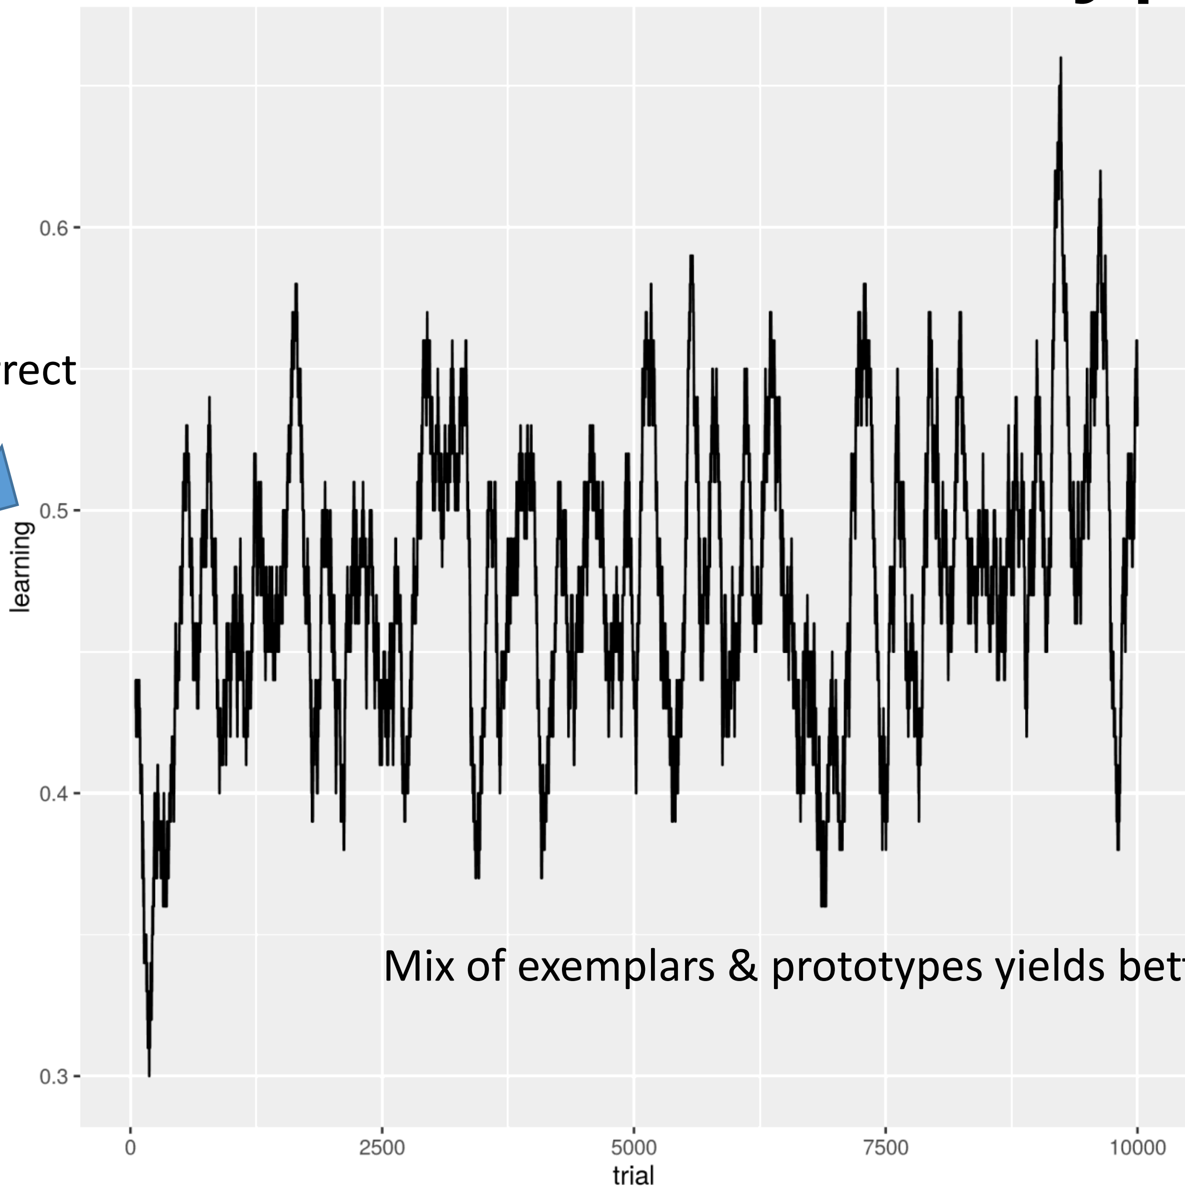

Mix of exemplars & prototypes yields better performance

<u>Activation</u>

:bll nil

:mas 10

:ans .25

:tmp nil ($\sqrt{2}$ * ans)

<u>Utility</u>

Reward 10, -1

:alpha 0.2

:egs 1.0

11,562 chunks

Augment Context Effect

That's 60% correct

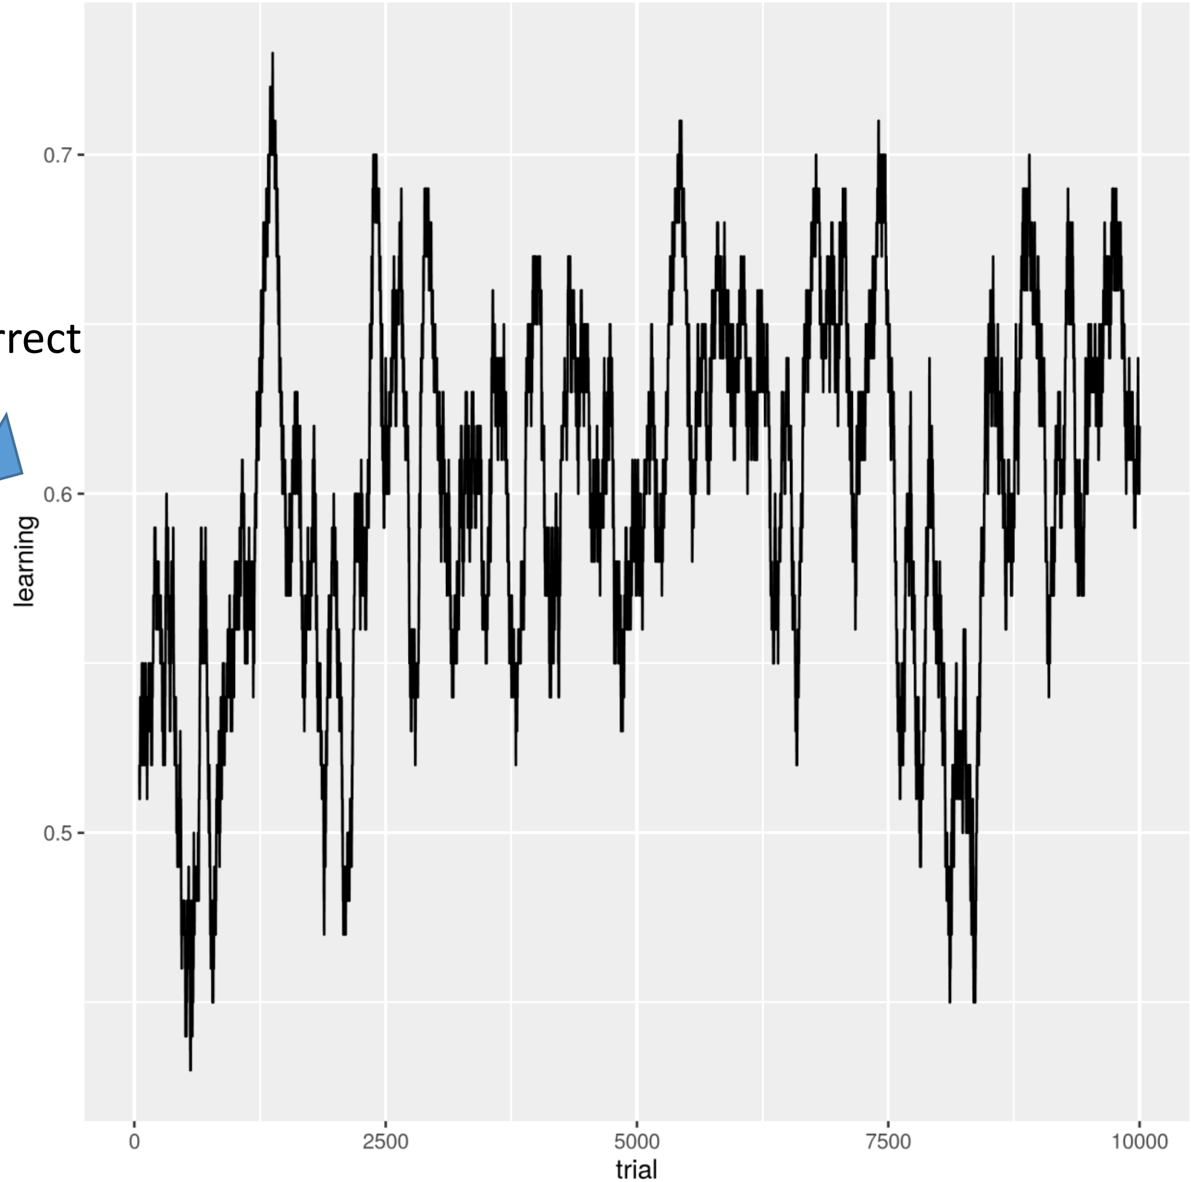

Activation
:bll nil
:mas 10
:ans .25
:tmp 10

Utility
Reward 10, -1
:alpha 0.2
:egs 1.0

11,591 chunks

Confusions After 9000 Trials

Overgeneralization
Many false alarms

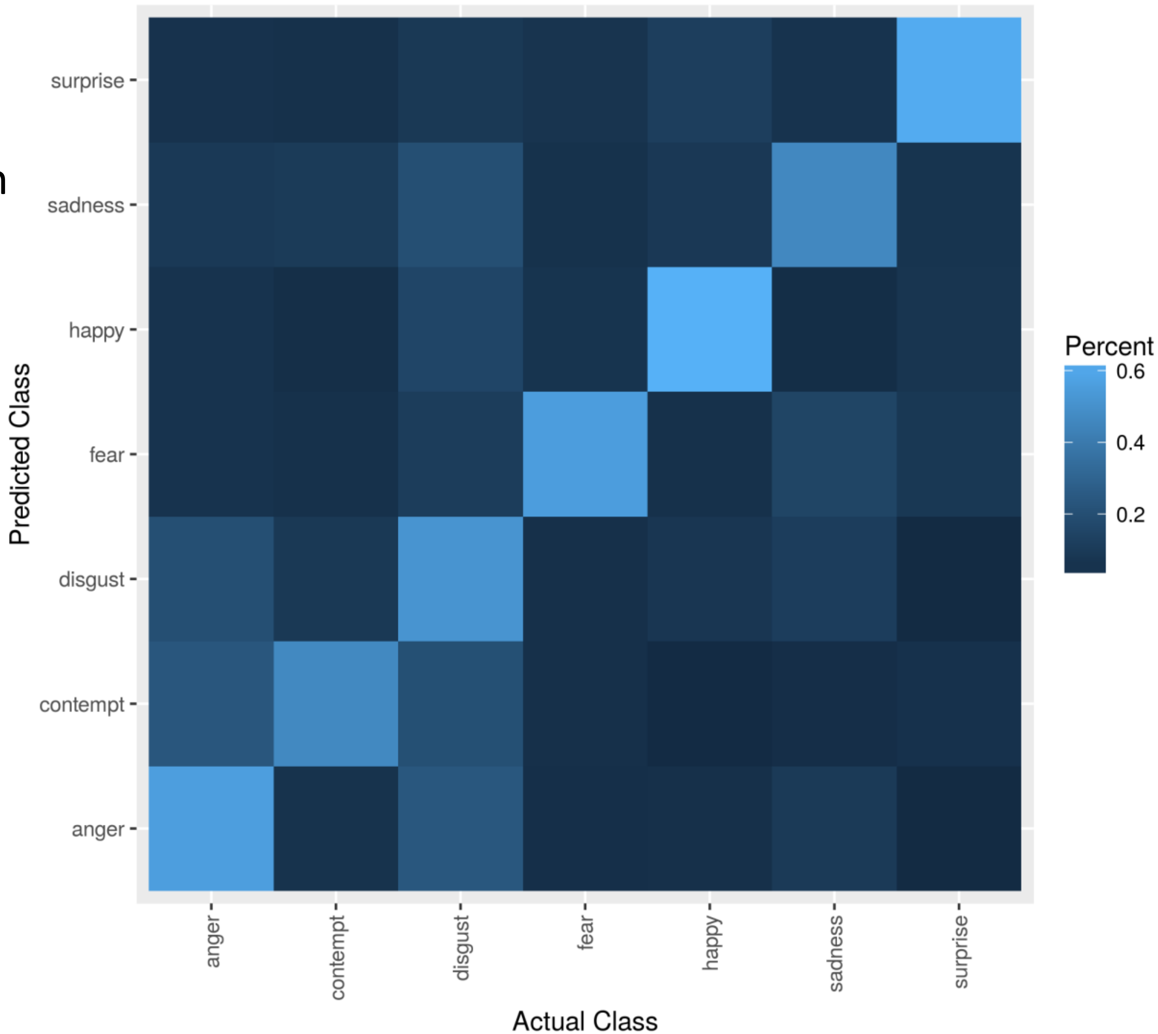

Control Reward Distribution

That's 80% correct

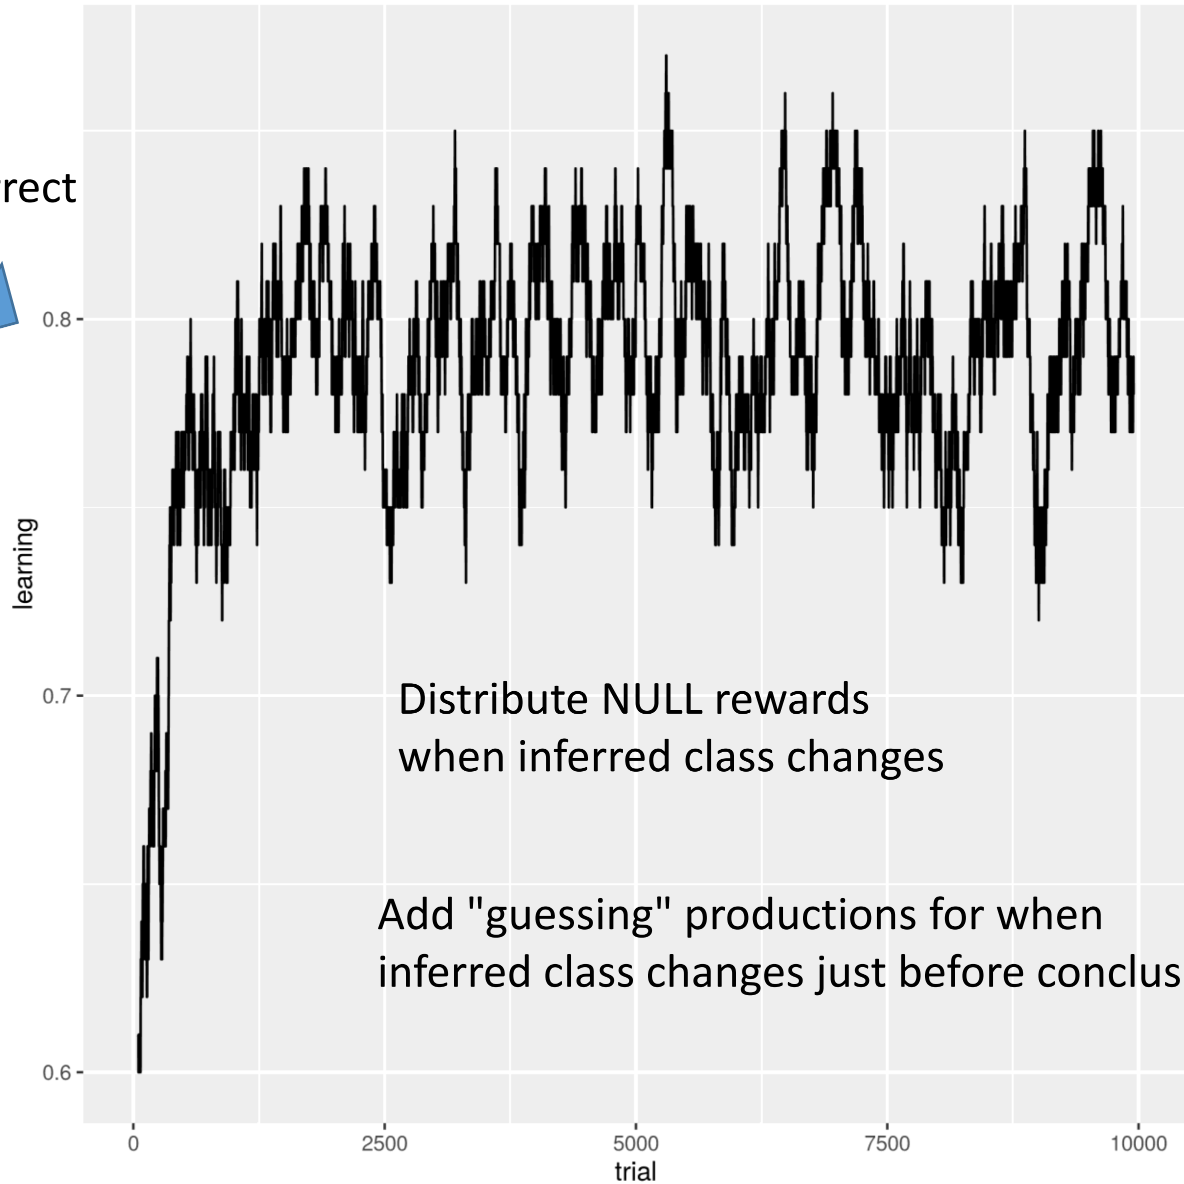

Distribute NULL rewards
when inferred class changes

Add "guessing" productions for when
inferred class changes just before conclusion is drawn.

<u>Activation</u>
:bll nil
:mas 10
:ans .25
:tmp 10

<u>Utility</u>
Reward 10, -1
:alpha 0.2
:egs 1.0

9721 chunks

Confusions After 9000 Trials

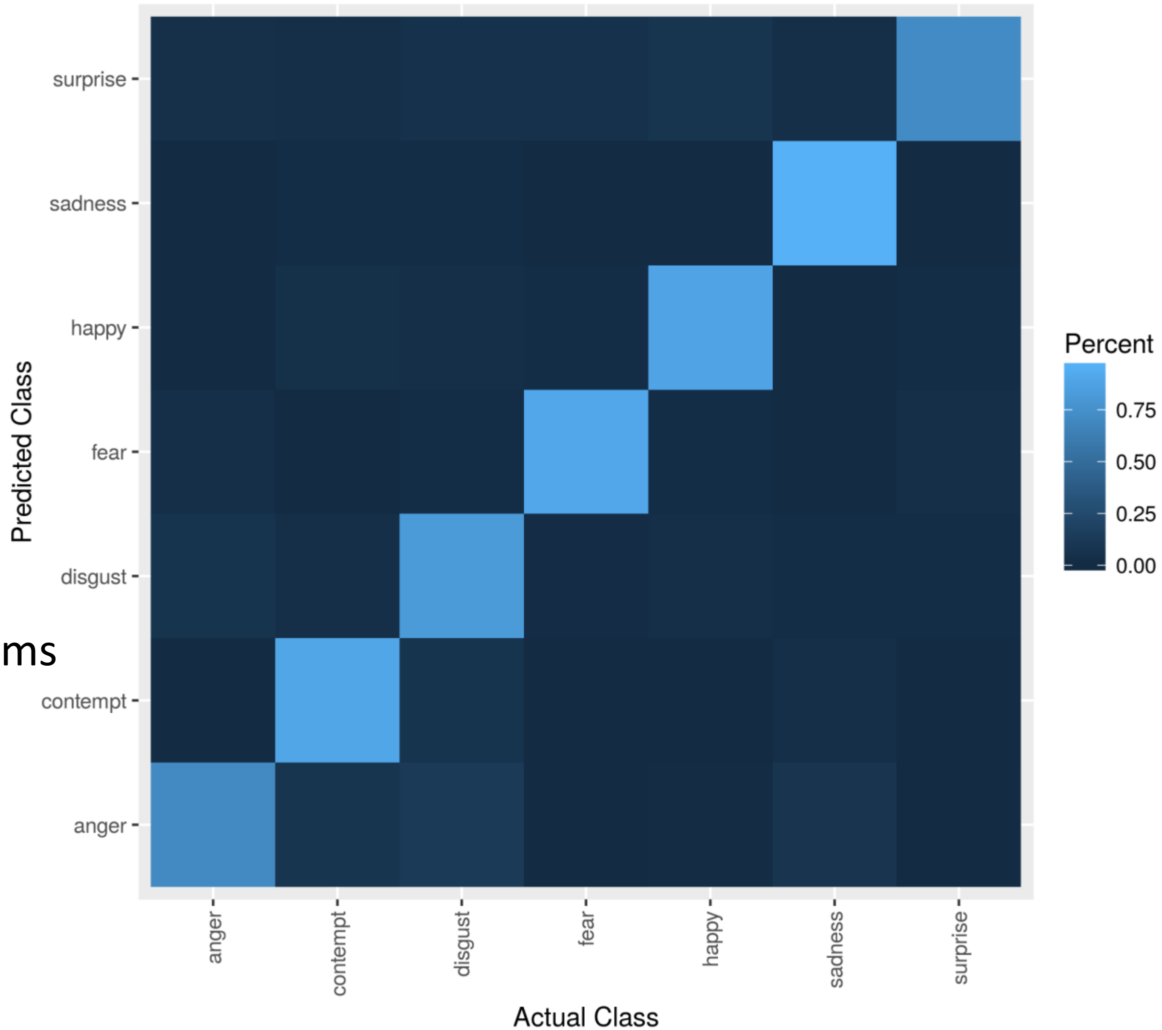

Fewer false alarms

Class-specific Decisions

That's 80% correct

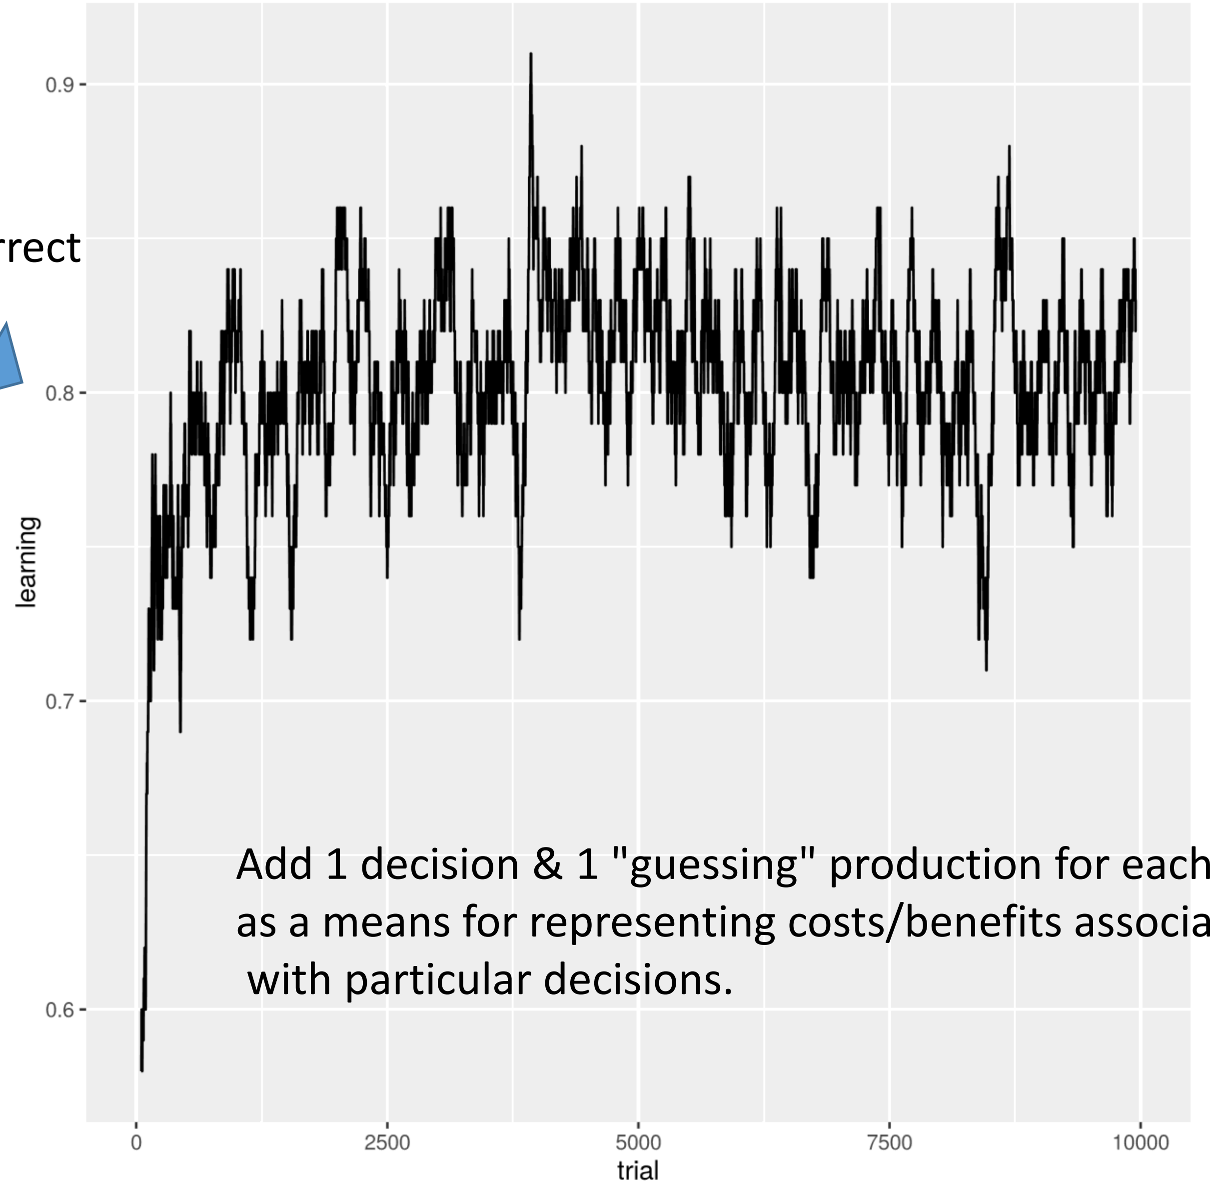

Add 1 decision & 1 "guessing" production for each class as a means for representing costs/benefits associated with particular decisions.

Activation
:bll nil
:mas 10
:ans .25
:tmp 10

Utility
Reward 10, -1
:alpha 0.2
:egs 1.0

10,161 chunks

Confusions After 9000 Trials

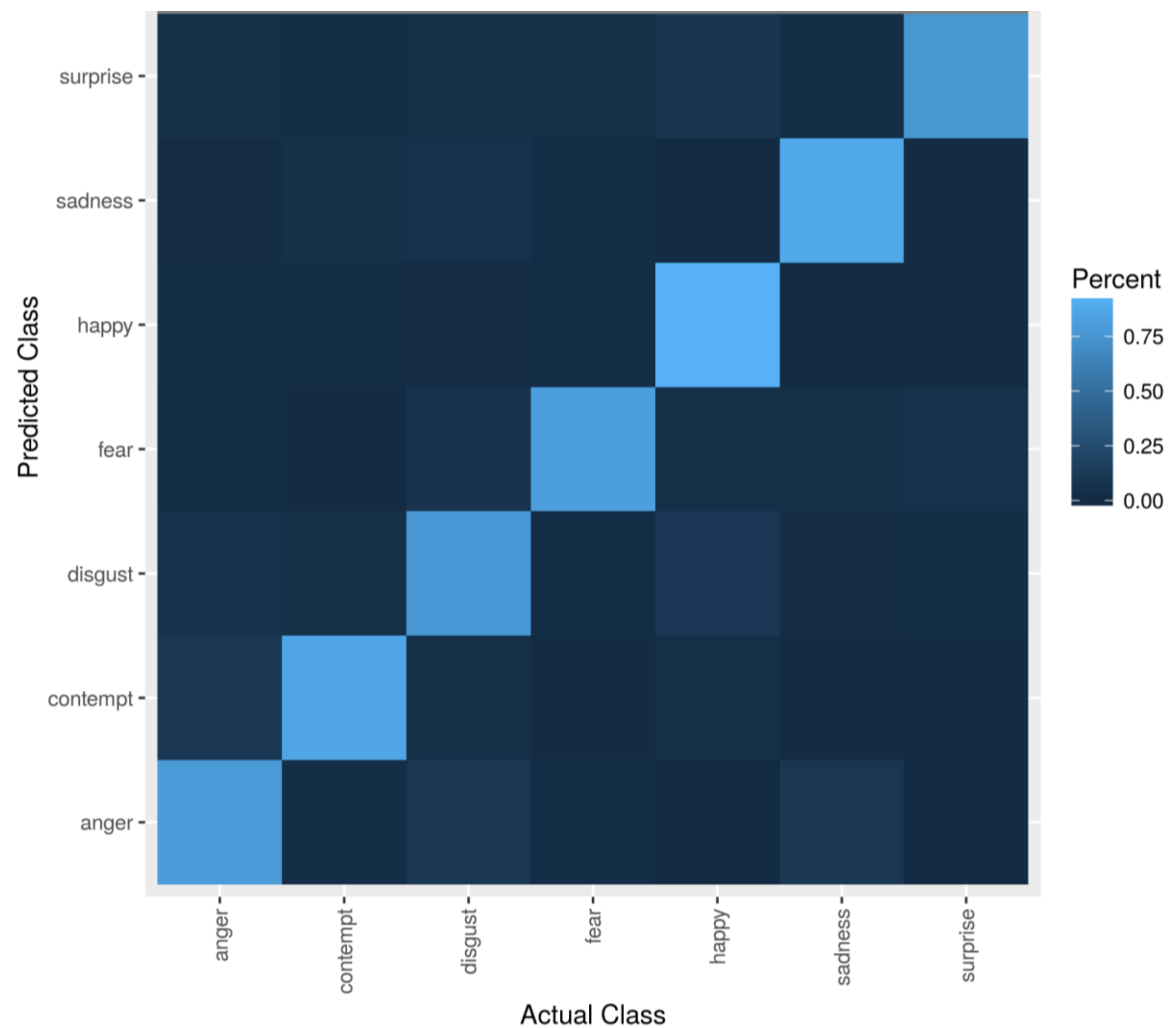

Feature Counts by Class

Model: Feature present or absent

Data: Feature presence only

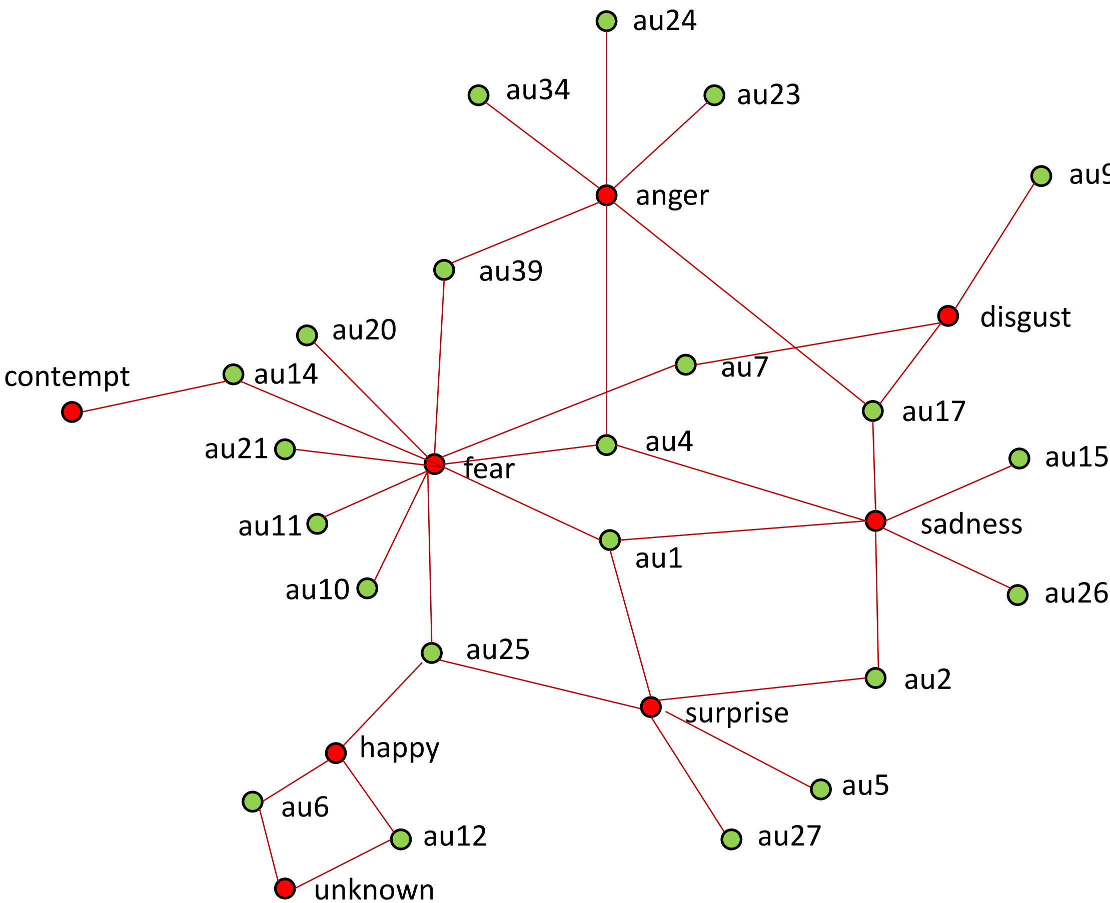

df$response: anger

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
2.00	9.00	16.50	18.69	29.00	39.00

---

df$response: contempt

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
6.00	22.75	27.50	28.82	38.00	39.00

---

df$response: disgust

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
1.00	8.00	21.00	21.17	33.00	39.00

---

df$response: fear

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
6.00	18.00	28.00	25.89	35.00	39.00

---

df$response: happy

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
1.0	13.0	24.0	23.1	34.5	39.0

---

df$response: sadness

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
4.00	13.75	23.50	22.84	30.50	39.00

---

df$response: surprise

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.00	2.00	11.00	14.64	26.00	39.00

Gold & Coal

- Gold Nugget
  - Flexible chunks
  - Link b/t symbolic & subsymbolic representation
  - Subsymbolic learning that is a mix of traditional ML methods
- Lump of Coal
  - Lack of a "simulated annealing" process in R-L
  - Rewards tied to time scale
  - Lack of production-specific exploitation/exploration