Summary
What we've learned, and what we haven't

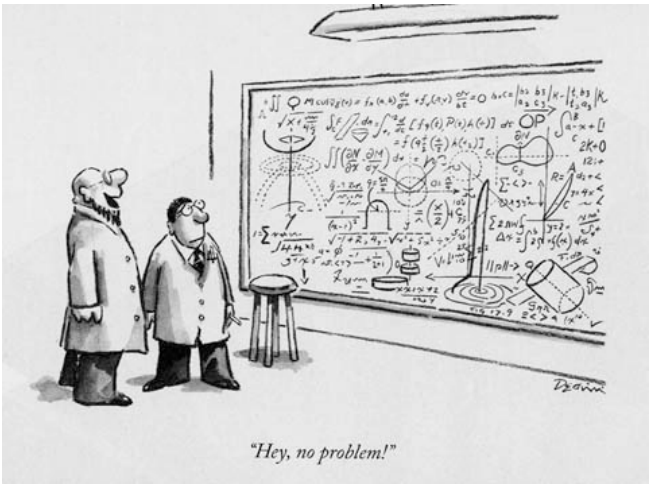

PSY/NEU338: Animal learning and decision making:
Psychological, computational and neural perspectives

What did you take home from this course?

If you had to summarize 1-3 take-home messages
- things you did not know (facts)...
- things you did not know how to do...
- things you had not thought about...

...before taking this course,
what would they be?
(ie. in what way has this course impacted on you?)

a few words about the exam

- 3 hours
- 25 short answer questions (~5 minutes each, 4-6 lines to answer) - similar to questions from the assignments
- I want to see that you've understood the material, and that you can reason about it intelligently (not that you have memorized the terms)
- open material (books, notes, no laptops)
- think simple, direct - not trick questions
- composition:
  ½ "easy", ¼ "challenging", ¼ "difficult"
  ½ behavior, ½ modeling, ⅓ brain
- review session prior to the exam (Nathan will schedule)
- review questions: you write them! you get points!

Review questions: How it works

- launch Jan 1st on piazza
- you have to write a question *and* an answer (4-5 lines max)
- technicalities: write the question as a question and the answer as a reply; note what topic your question pertains to (from the title of the slide show) - see below
- discussion by others allowed; corrections for answers allowed
- good questions get 1 point
  good answers (by others) get 1 point
  the first good question on a topic gets 2 points
- these points are *real*. for the exam. you can get up to 5.

overview

- a few words about the exam
- the big picture (+ caveats)
- a few paradoxes (and fewer resolutions)
- advice for a young investigator

what did we have here?

- classical conditioning (error driven learning, R-W, TD, dopamine & basal ganglia, fMRI of prediction errors, opponent processes)
- instrumental conditioning (free operant schedules, reinforcement learning, Actor/Critic, Q values, comparing models using neural data)
- actions and habits (S-R vs R-O, devaluation, model based and model free RL, uncertainty based arbitration)
- comparison between Pavlovian and instrumental conditioning (omission schedules)
- learning as a model of disorders (drug abuse as learning gone awry, latent inhibition and schizophrenia)
- extinction (latent cause models, Bayesian inference)
- generalization and discrimination (configural versus elemental models, Pearce & similarity)

the BIG picture

brain

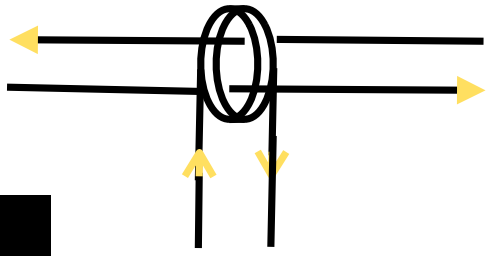

behavior

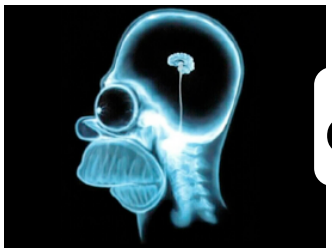

computation

$\delta(t)=r(t)+V(t+1)-V(t)$

the BIG picture
normative models

how do animals behave?

how should animals behave if they were optimal?

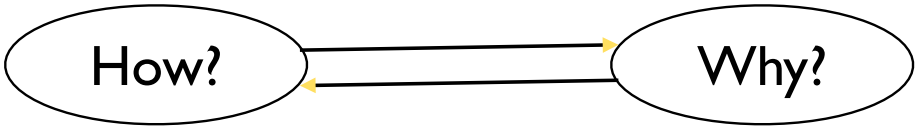

Marr's three levels:
the problem, the algorithm, the implementation

The problem: maximize reward, minimize punishment

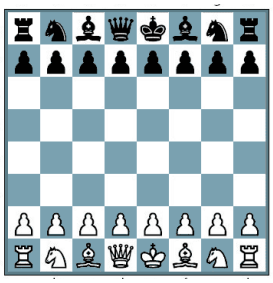

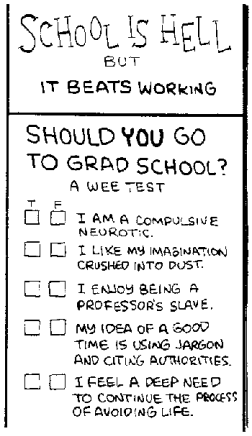

Decision making is hard because:

- Reward/punishment may be delayed
- Outcomes may depend on a series of actions

⇒ "credit assignment problem" (Sutton, 1978)

Marr's three levels:
the problem, the algorithm, the implementation

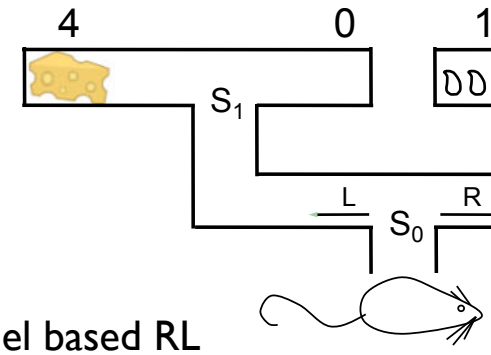

Model based RL

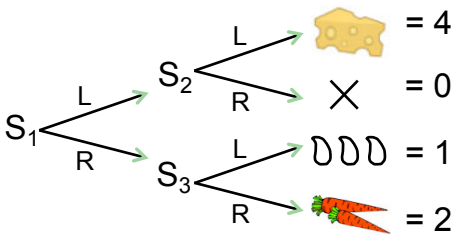

Model free RL

$Q(S_0,L) = 4$
$Q(S_0,R) = 2$

$Q(S_1,L) = 4$
$Q(S_1,R) = 0$

$Q(S_2,L) = 1$
$Q(S_2,R) = 2$

Marr's three levels:
the problem, the algorithm, the implementation

learning a model of the environment ("cognitive map") + planning in the map

error-correcting learning rule to learn values (R-W/TD)

Model based RL

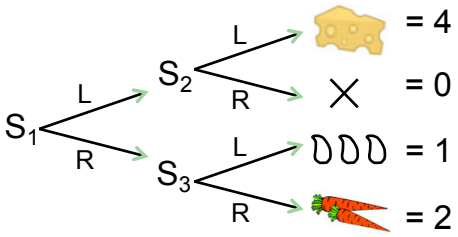

Model free RL

$Q(S_0,L) = 4$
$Q(S_0,R) = 2$

$Q(S_1,L) = 4$
$Q(S_1,R) = 0$

$Q(S_2,L) = 1$
$Q(S_2,R) = 2$

Marr's three levels:
the problem, the algorithm, the implementation

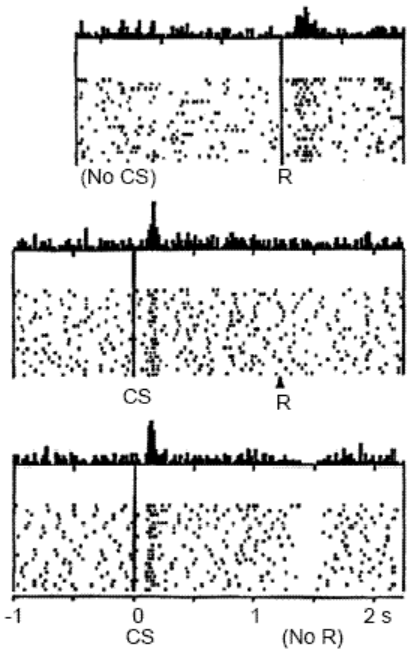

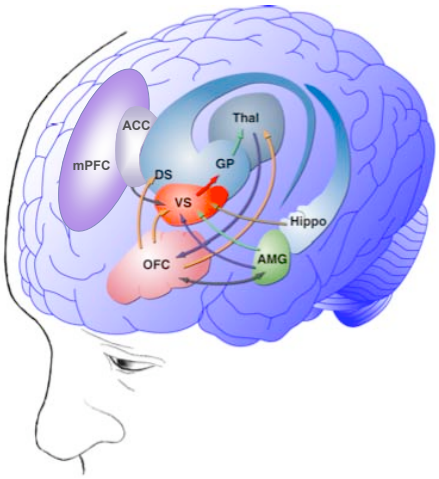

Marr's three levels:

the problem, the algorithm, the implementation

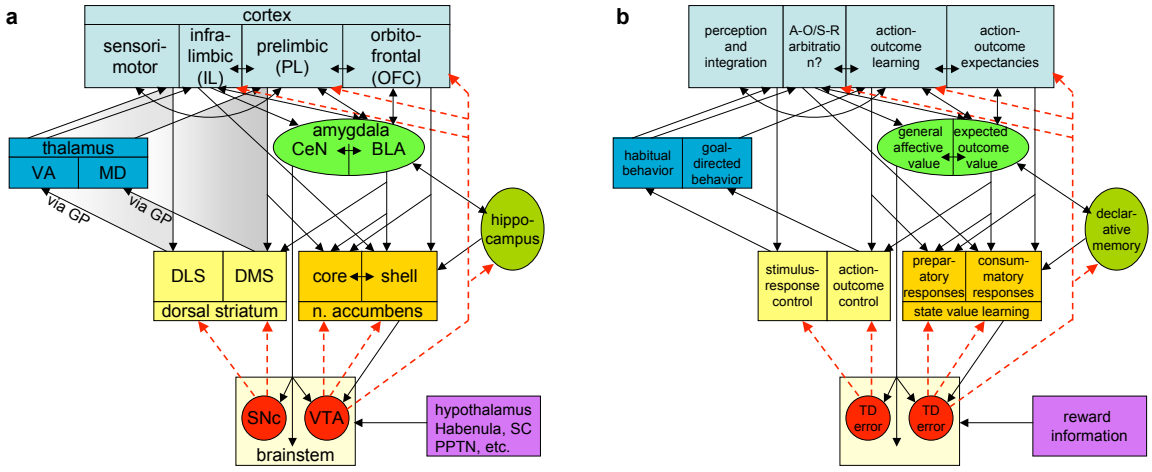

caveats: this was all through a very specific lens...

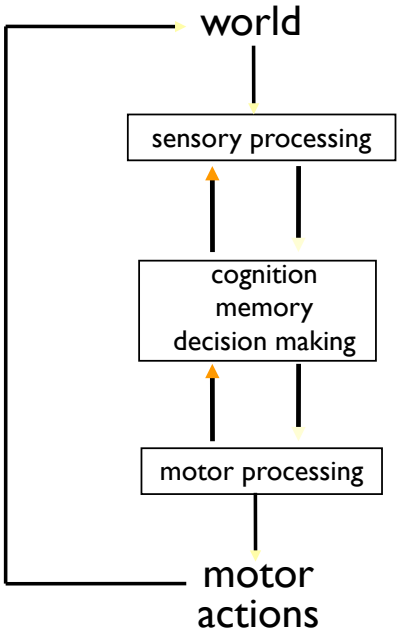

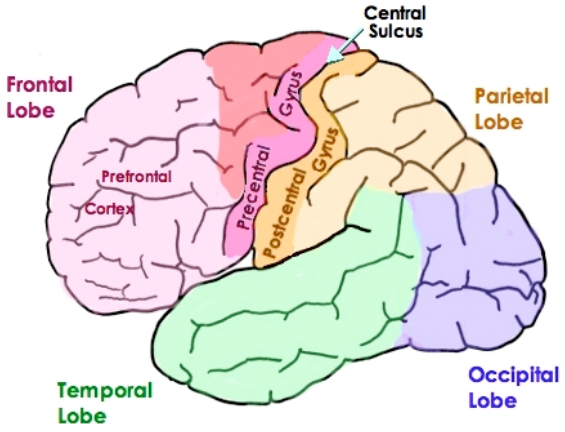

caveats: this was all through a very specific lens...

Hazy, Frank & O'Reilly, 2007

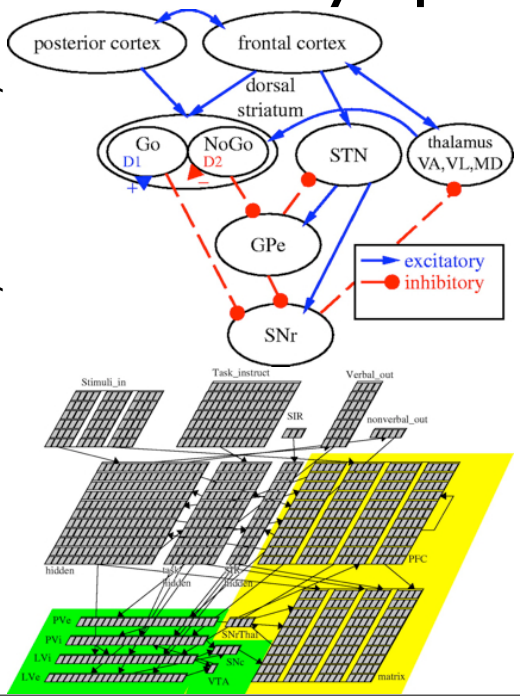

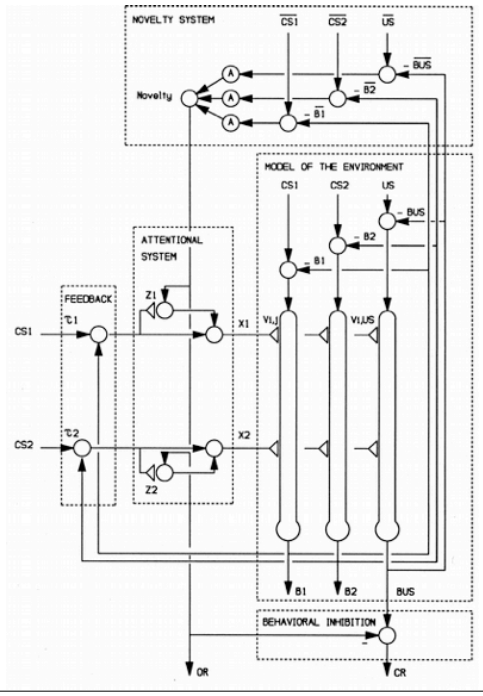

Schmajuk, Lam & Gray, 1996

keep in mind...

- *Behavior is complicated* (many mechanisms interactions) ➪ design experiments to investigate each mechanism in isolation, minimizing effect of others
  ⇒ beware of false dichotomies
- *All models are wrong, some models are useful (G. Box)*
  ⇒ be aware of the scope of the model (what does it focus on? What does it ignore intentionally? unintentionally?)
- *More questions than answers...*
- *This is a great time to be studying behavior and the brain!*
  (we are getting absolutely incredible data, computers are getting very fast, synergies between fields are leading to a deep and integrated understanding…)

overview

- a few words about the exam
- the big picture (+ caveats)
- a few paradoxes (and fewer resolutions)
- advice for a young investigator

I. asymmetries galore

- which is easier to learn:
  - A- / AB+
  - A+ / AB-

- what would you predict (behaviorally, for the tone) for:
  - 4min noise → 10sec tone → reward
    1min noise → 10sec tone → nothing
  - 4min noise → 10sec tone → nothing
    1min noise → 10sec tone → reward

1. asymmetries galore

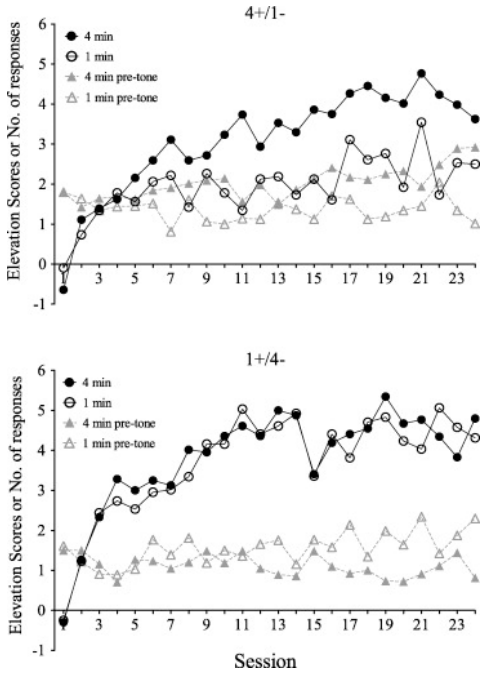

Todd, Winterbauer & Bouton 2010

2. attention and associability

- William James's famous (non-)definition of attention: "Everyone knows what attention is. It is the taking possession by the mind in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought...It implies withdrawal from some things in order to deal effectively with others." (Principles of Psychology, 1890)
- huge and interesting topic of research
- emphasis on: where is the bottleneck of attention (cognitive psychology); how is attention realized in the brain (neuroscience)

2. attention and associability

'downward unblocking'

stage 1: A → 2 food pellets

stage 2: AB → 1 food pellet.

What do you predict would be learned for B?

(Dickinson, Hall & Mackintosh, 1976)

2. attention and associability

- A normative perspective:
  What would you attend to more, a less surprising cue (you know what it predicts), or a more surprising one (you don't know what it predicts)?
- Pearce & Hall: attention for learning vs. attention for performance
  - attend to predictive stimuli for performance
  - attend to surprising stimuli for learning
  - surprise (=prediction error) enhances associability (=learning rate on next trial). Prediction: giving a larger US just prior to extinction accelerates extinction (validated)
  - learning not through error correction: $V^{new} = V^{old} + \eta R$ (explains paradoxical results of downward unblocking)

2. attention and associability

Holland & Gallagher: Amygdala (CeN) and attention

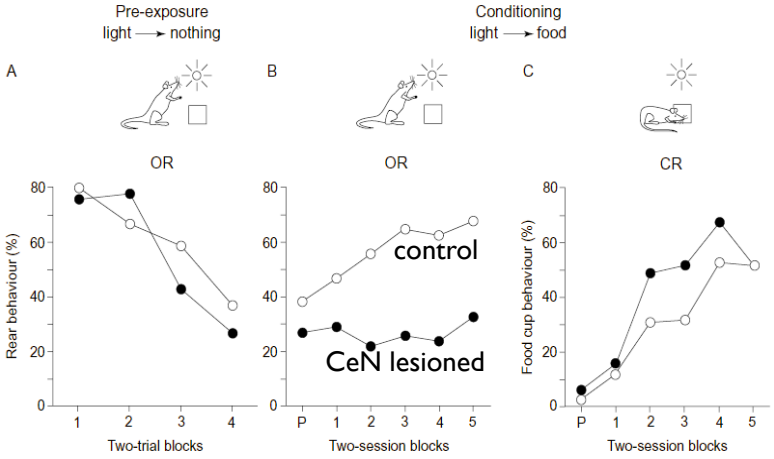

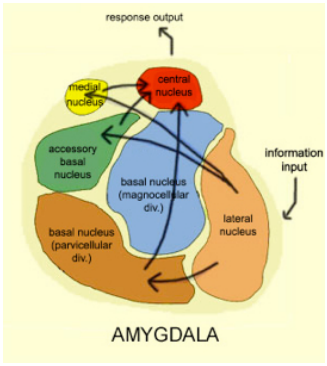

unconditional OR intact

conditional OR disrupted

CRs intact

2. attention and associability

Holland & Gallagher: Amygdala (CeN) and attention

Training condition	Phase 1	Phase 2	Test
Consistent	light → tone → food light → tone → nothing	light → tone → food light → tone → nothing	light → food
Shift	light → tone → food light → tone → nothing	~~light → tone → food~~ light → nothing	light → food

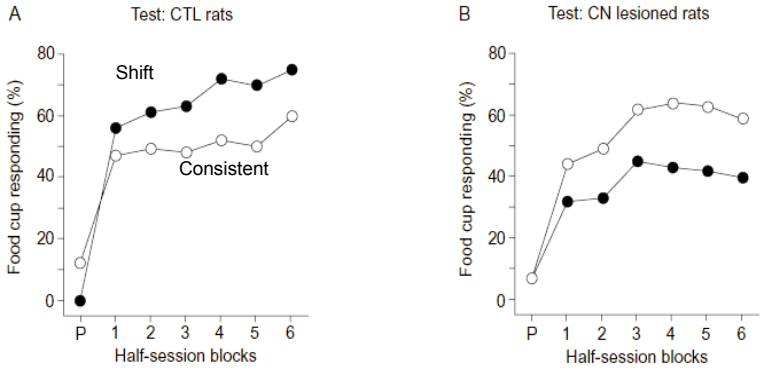

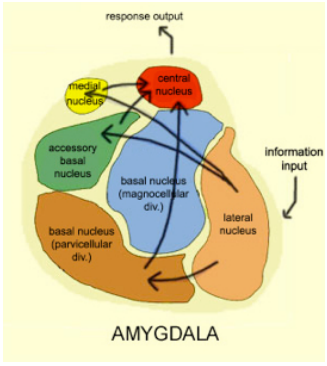

surprise enhances associability (CeN dependent) (also ACh dependent)

3. Serotonin and depressive realism

- serotonin (5HT) is implicated in signaling aversive events and their prediction (opposite of dopamine)
- serotonin is implicated in inhibiting behavior (as a result of the prediction of aversive outcomes?)
- BUT: the main treatment for depression is SSRIs that *elevate* levels of serotonin

3. Serotonin and depressive realism

"Markov models of trains of thought"

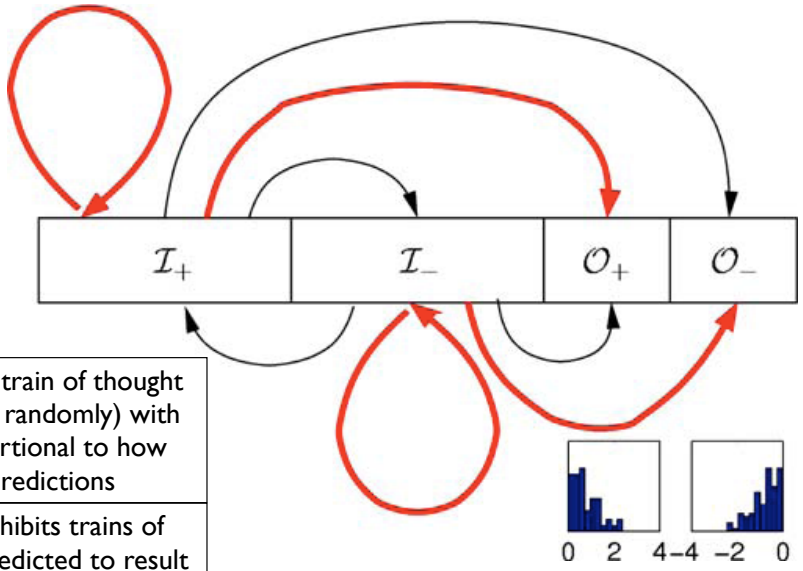

Policy: terminate a train of thought (and start another randomly) with probability proportional to how negative its predictions

Idea: serotonin inhibits trains of thought that are predicted to result in negative outcomes → bias toward optimistic valuation

Dayan & Huys, 2008

Studying learning is like solving a huge puzzle!

constraints/paradoxes of behavior

computational models

neural data

What I wanted you to take home from this course...

- facts you did not know:
  - prediction errors as a way to think about learning
- things you did not know how to do:
  - psychological tools for thought (clever experimental design)
  - computational tools for thought (think about what the brain does through the lens of what it computes)
- things you had not thought about:
  - notice conditioning in everyday life - prediction errors, habits etc.
  - relationship between complex behaviors and simple paradigms/models

advice for a graduating senior

- the doors don't close; they just become a bit harder to shove open
- as long as you are doing what you want to do *now* you will be happy! (and isn't that the whole point?)
- decision making trick: what do I want to do *first*?
- remember: it is the journey that matters, not the end point

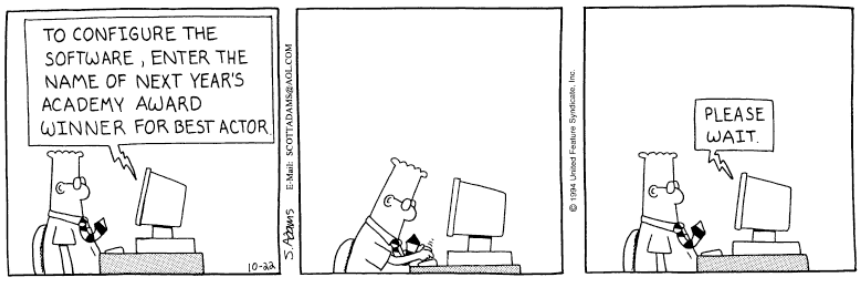

THANK YOU!