CCN Lecture

Depression & Reinforcement Learning (cont.)

26 / 10 / 2017

Sam Rupprechter

s1520312@sms.ed.ac.uk

Outline

- Depression
- Reinforcement Learning (RL)
- RL Impairments in Depression
- → Modelling Theory

Major Depressive Disorder (MDD)

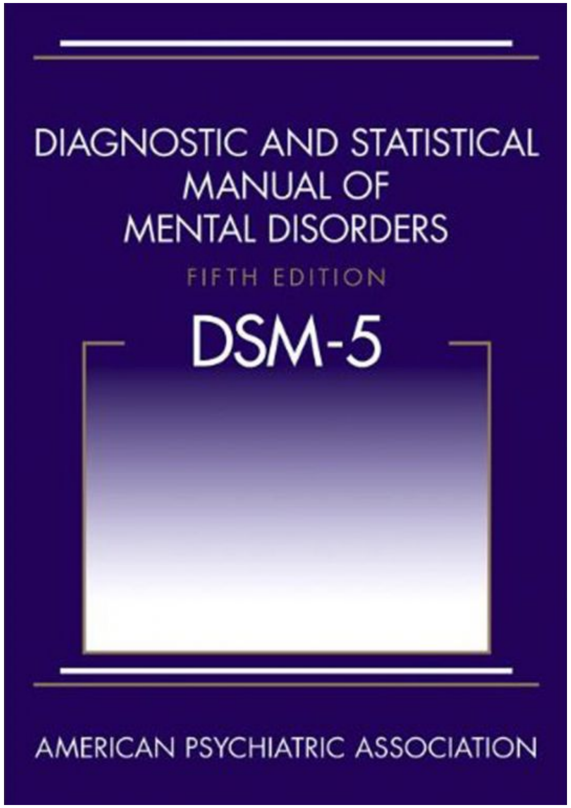

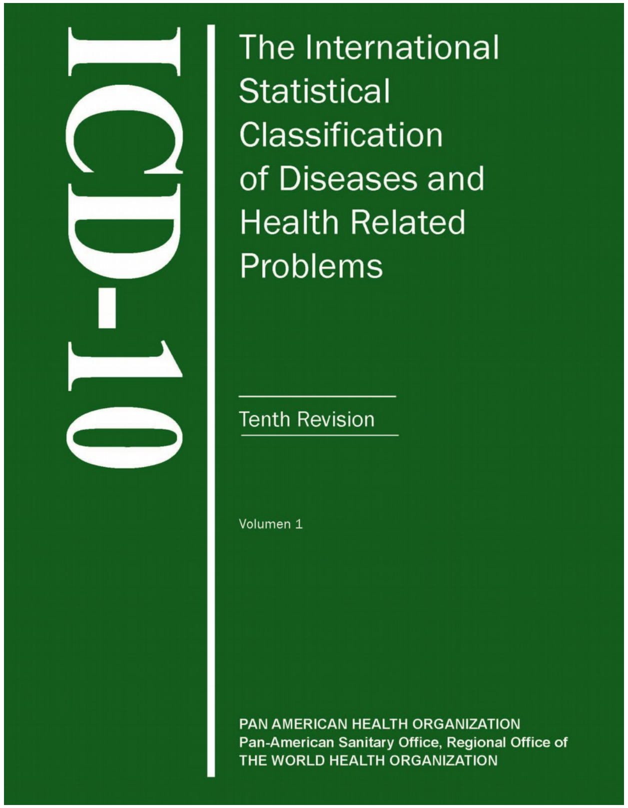

Major Depressive Disorder (MDD)

- Symptoms:
  - **Depressed mood**
  - **Anhedonia** (inability to experience pleasure)
  - Loss of energy, fatigue
  - Change in weight or appetite
  - Insomnia / Hypersomnia
  - Psychomotor agitation / retardation
  - Feelings of worthlessness or excessive or inappropriate guilt
  - Concentration difficulties
  - Suicidal thoughts / ideation

Major Depressive Disorder (MDD)

- Categorical view has little basis in biology?
  - Research moves towards dimensional view
- RDoC framework
  - Multiple levels of analysis
    - Neural circuitry, genes, behaviour
- Endophenotypes
  - Anhedonia
  - Neuroticism

Treatment

- Cognitive Behavioural Therapy (CBT)
- Antidepressant medication
  - Selective Serotonin Reuptake Inhibitors (SSRIs)
    - Primary first line treatment
  - Serotonin-Norepinephrine Reuptake Inhibitor (SNRIs)
  - Tricyclic Antidepressants (TCAs)
- Electroconvulsive therapy (ECT), Surgery
  - Very severe, treatment-resistant cases

MDD Theories

- Cognitive Theory (Beck, 2008)
  - Negative cognitive schemas (CBT targets those)
  - e.g. biased recalling of negative events
- Learned Helplessness (Seligman, 1972)
- Stress → deficits in reinforcement / reward processing (learning) → anhedonia (Pizzagalli, 2014)
  - 70-80% of Major Depressive Episodes preceded by major life event

Learned Helplessness

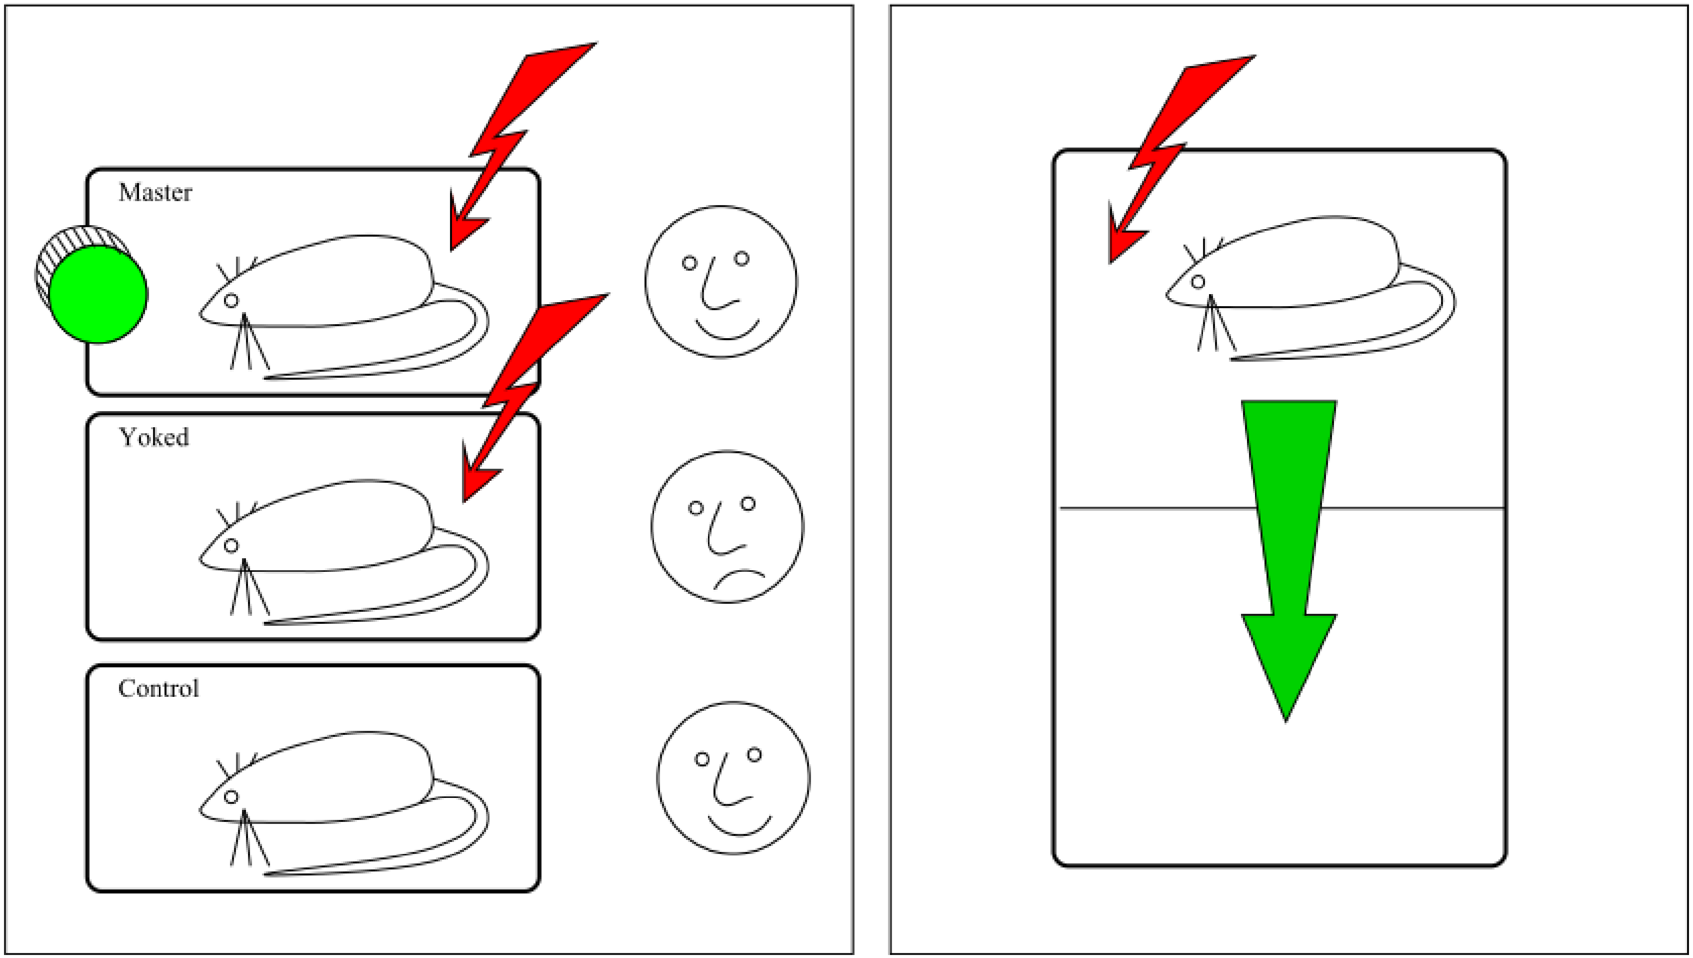

Huys et al., 2008; NIPS

Impact

- High (lifetime) prevalence (esp. in developed countries)
  - USA: 16.2% (Kessler et al., 2003)
  - UK / Europe: 7-10% (Ayuso-Mateos et al., 2001)
  - Depression rates are rising (e.g. Mojtabai et al., 2016)

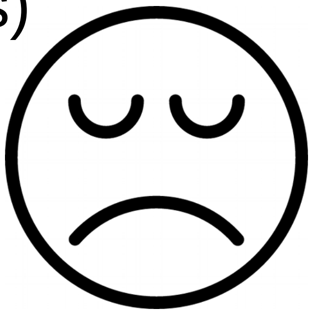

- High economic impact (Europe: €92 billion in 2010) (Olesen et al., 2012)
- People are suffering
  - Risk factor for suicide (Olfson et al., 2017)
    - And suicide rates are increasing
  - Cognitive Impairments (e.g. Snyder, 2013)
    - Attention, concentration, executive functioning, working memory, …
  - Impairments in Reinforcement Learning (Chen et al., 2015)

RL Impairments in MDD

- Signal Detection Task
- fMRI studies
- Computational Modelling

- Iowa Gambling Task (?)
- Reversal Learning (?)

Signal Detection Task (e.g. Pizzagalli et al., 2005)

- One stimulus rewarded more often
  - (healthy) participants become biased towards it

Reinforcement Learning (RL)

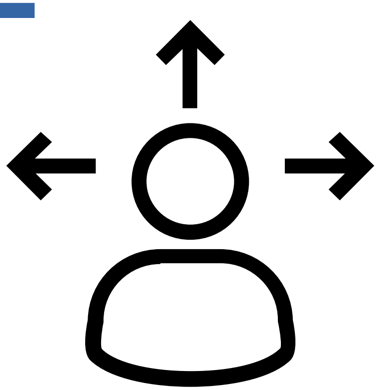

- Update values

$$V(t+1) = m \times V(t) + \varepsilon \times (\rho \times r(t) - V(t))$$

- Decide between two options

$$p(a \mid V, \theta) = \frac{1}{1 + \exp(-\beta \times (V_a - V_b))}$$

MDD Modelling Studies (behavioural)

- Chase et al., 2010
  - Lower learning rates
- Kunisato et al., 2012
  - Lower temperature parameter
- Huys et al., 2013
  - Lower reward sensitivity
- Beevers et al., 2013
  - Higher temperature parameter
- Dombrovski et al., 2010
  - Lower memory [in suicide attempters]

Brain Activity

- Model-based fMRI (e.g. Kumar et al., 2008; Gradin et al., 2011)
  - No real behavioural differences
  - Abnormal reward prediction errors
  - Abnormal expected reward values

(Behavioural) Modelling

- How do I actually “fit” a model to data?
  - Try to find “optimal” values for the parameters of the model that our data “most likely”

    (maximize the probability of observed choices)

Maximize the Likelihood

$$L = p(A \mid V, \theta) = \prod_{a \in A} p(a \mid V, \theta)$$

- Multiplying lots of small numbers is a bad idea… take the log instead!
- Instead of maximizing log likelihood → we usually minimize negative log likelihood

$$NLL = -\sum_{a \in A} \log p(a \mid V, \theta)$$

Example Experiment

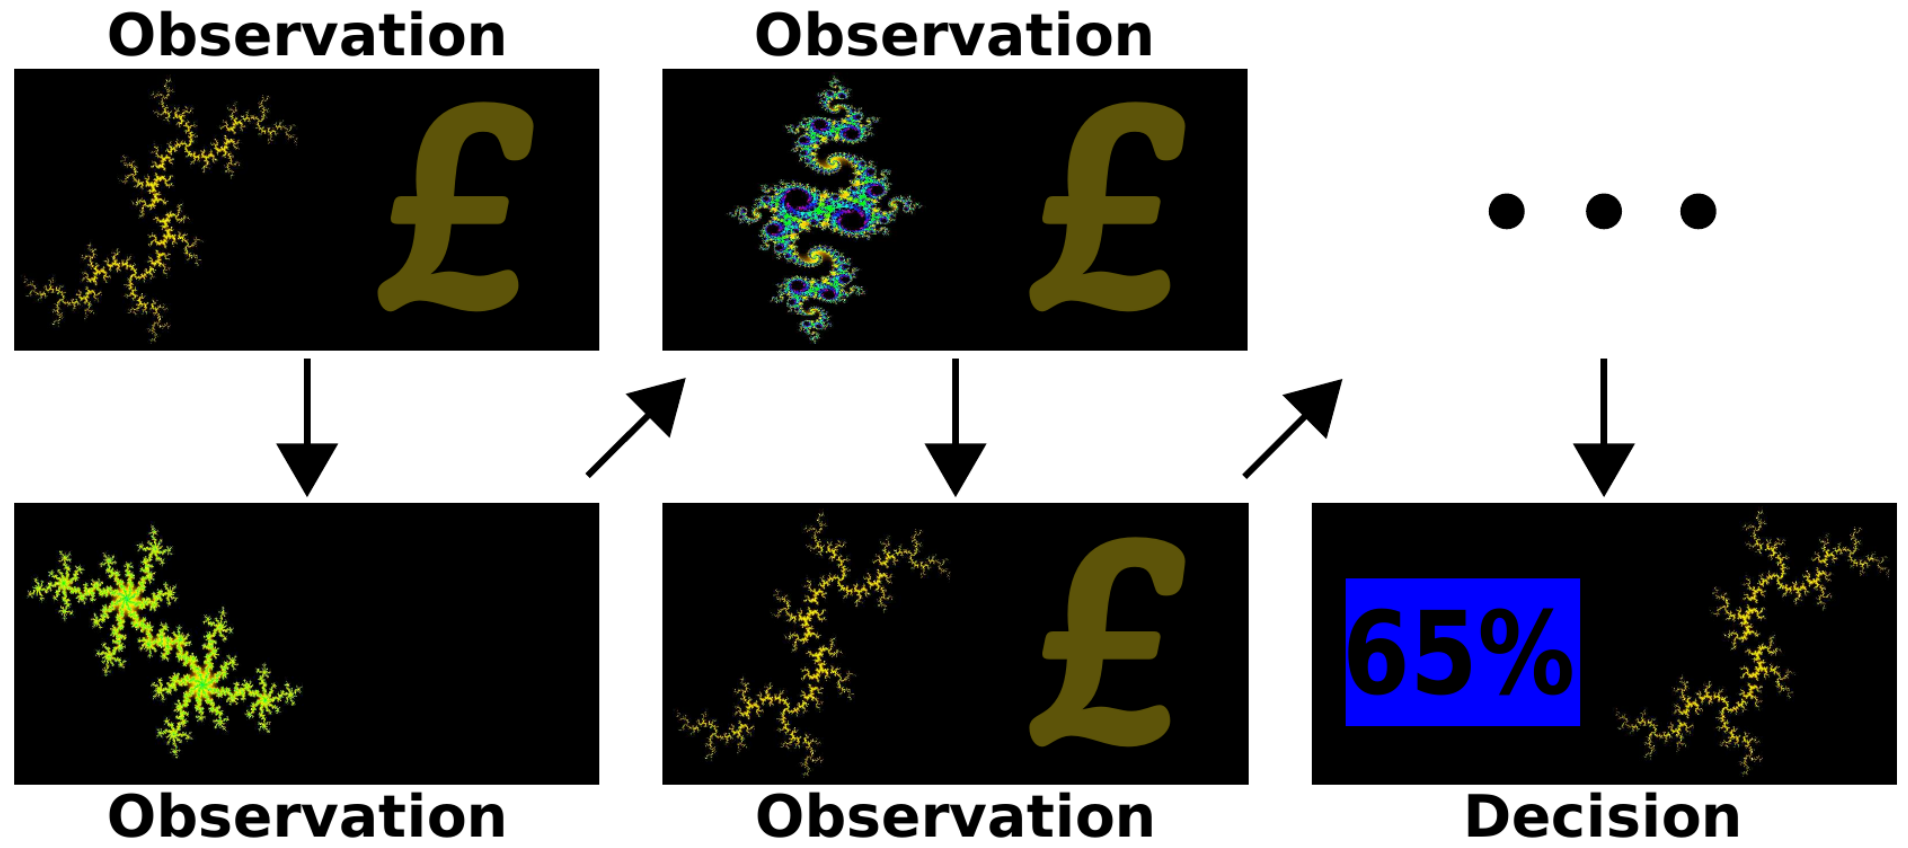

Stankevicius et al., 2014; Further work in progress

Our Model

- Value Update

$$V_i^{t+1} = A \times V_i^t + r_i^t$$

- Decision

$$p(\text{choose fractal i}) = \frac{1}{1 + \exp(-\beta(f(V_i) - \phi_i))}$$

NLL in MATLAB

```
function nll = neg_log_likelihood(data, theta)
    A = theta(1);
    beta = theta(2);
    X = data.decisions;
    T = data.num_trials;
    r = data.obs_rewards;
    p = data.phis;
    V = zeros(T, 1);
    for i = 1:size(r, 2)
        V = A*V + r(:, i);
    end
    probs = logsig(X .* beta .* (V/4 - p));
    nll = -sum(log(probs));
end
```

Estimate parameters

- Different options (e.g. gradient descent)
  - We will simply use one of the built-in functions

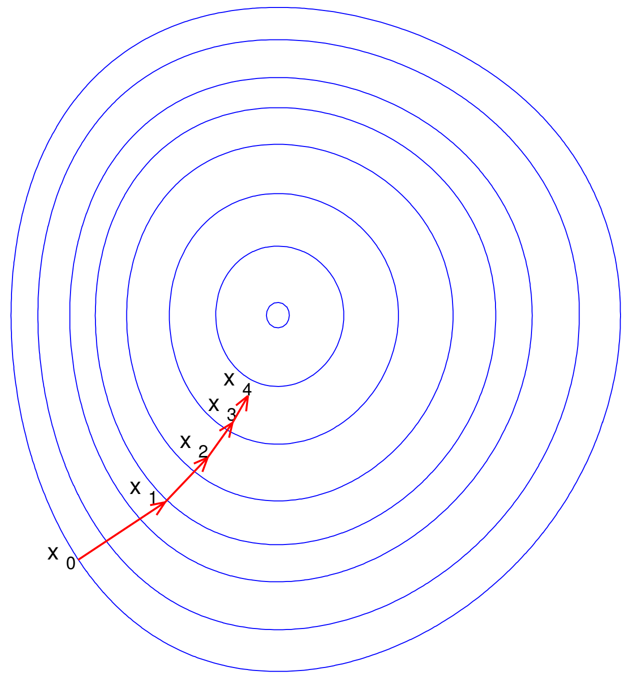

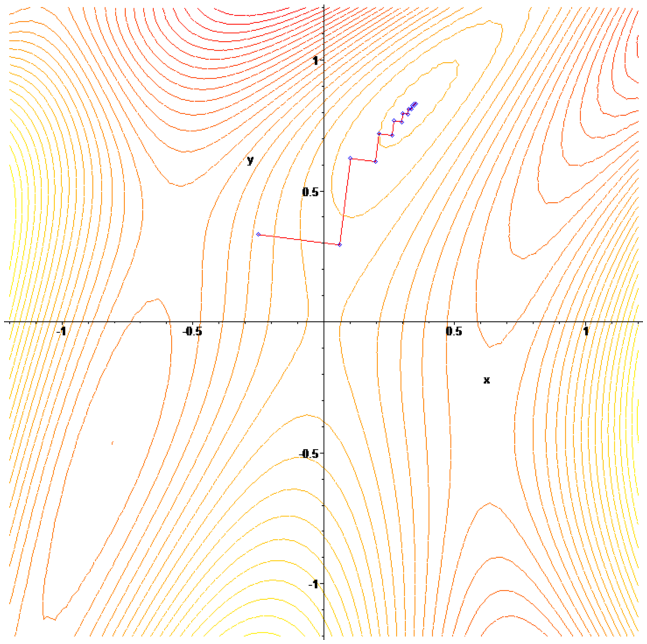

Possible Issues

- Lots of local minima
- Surface around minima is very flat

- Initialise with different starting points
  - Randomly
  - Grid

fminunc

```
f = @(x)(neg_log_likelihood(data, x));
thetas = fminunc(f, [0;0]);
```

How good is our estimation?

- If we are making inferences based on specific parameter values (e.g. look at group differences), we better make sure that those estimates are reliable
- **Simulate data** from estimated parameters
  - Does generated data “look like” the original data? (similar summary statistics, evolution of values, …?)
    - How much does the generated data vary?
  - Re-fit parameters to simulated data and compare parameters (e.g. look at the correlation: hopefully close to 1)
    - How much do the simulated parameters vary?

How good is our estimation?

- Look at the curvature (Hessian / 2nd order derivative) at the estimated point (Hessian returned by `fminunc`)
  - Take inverse to get covariance matrix

Correlated Estimated Parameters

- Might cause issues during inference
  - e.g. if two parameters are (highly) negatively correlated
    - We can arbitrarily change one of the parameters and then adjust the second parameter so as to keep the previous “maximum” likelihood (extreme example)
    - What does that mean if we are interested in the actual values of these parameters (e.g. for group comparisons)?

Correlated Estimated Parameters

- As long as the abs(correlation) is < 1 both parameters will explain “something”
  - Unclear what value between 0 and 1 would be “too high”; will depend on the problem; use simulations
- Parameters might actually be correlated
  - People who learn faster (higher learning rate) might be better at “remembering” what they learned (lower discounting)
- Make sure parameters are distinguishable in the mathematical formulation
  - c.f. reward sensitivity and inverse temperature

Constraining Parameters

- Often we know what range of values is sensible for specific parameters (0 < learning rate < 1)
  - Want to make sure estimated parameters lie within that range
  - E.g. force parameters to be positive by exponentiating them at the beginning of the likelihood function
    - Optimisation function (fminunc) can search whole space (-inf to +inf)

Model Comparison

- How do we choose a model (hypothesis)?
- We want a Trade-off
  - Which model fits our data best? (accuracy)
    - Likelihood
  - Which model is the simplest? (complexity)
    - Number of parameters
- Turn to Bayesian model comparison…

Occam's razor

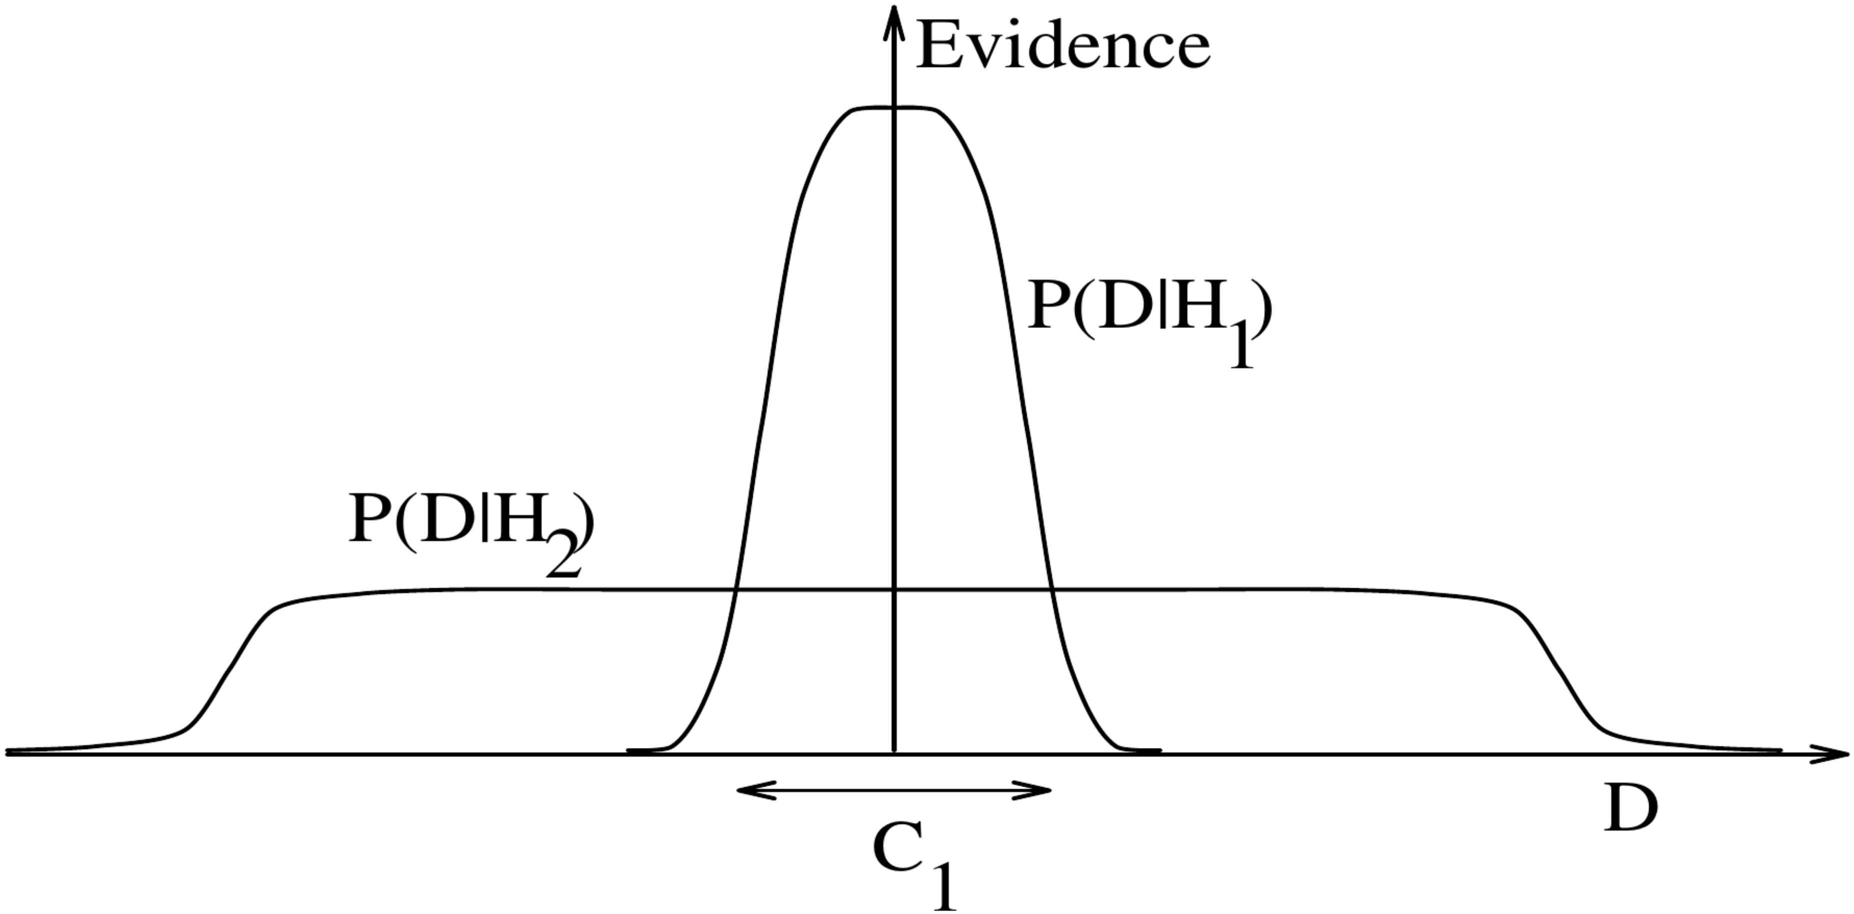

MacKay, 2003

Bayesian Model Comparison

- … or rather approximations

$$AIC = 2 \times NLL + 2 \times d$$

$$BIC = 2 \times NLL + d \times \log(n)$$

- Calculate for each model
- Choose model with **lowest** value (if difference > 10)
- Note that adding “redundant” parameters might affect the comparisons

Model Recovery Simulations

- **Do we have the data we need to answer the questions we are asking?**
- Confusion matrix
  - For each model m
    - Generate data from m
    - Fit all models to this data
    - Does model comparison choose m?
    - (repeat steps inside loop multiple times)

	H1	H2	H3
H1	20	0	0
H2	0	19	1
H3	0	1	19

References

- Ayuso-Mateos, J. L.; Vázquez-Barquero, J. L.; Dowrick, C.; Lehtinen, V.; Dalgard, O. S.; Casey, P.; Wilkinson, C.; Lasa, L.; Page, H.; Dunn, G. & others; Depressive disorders in Europe: prevalence figures from the ODIN study; The British Journal of Psychiatry, RCP, 2001, 179, 308-316
- Bakic, J.; Pourtois, G.; Jepma, M.; Duprat, R.; Raedt, R. & Baeken, C.; Spared internal but impaired external reward prediction error signals in major depressive disorder during reinforcement learning; Depression and anxiety, Wiley Online Library, 2017, 34, 89-96
- Beck, A. T.; The evolution of the cognitive model of depression and its neurobiological correlates; American Journal of Psychiatry, Am Psychiatric Assoc, 2008, 165, 969-977
- Beevers, C. G.; Worthy, D. A.; Gorlick, M. A.; Nix, B.; Chotibut, T. & Maddox, W. T.; Influence of depression symptoms on history-independent reward and punishment processing; Psychiatry research, Elsevier, 2013, 207, 53-60
- Chase, H.; Frank, M.; Michael, A.; Bullmore, E.; Sahakian, B. & Robbins, T.; Approach and avoidance learning in patients with major depression and healthy controls: relation to anhedonia; Psychological medicine, Cambridge University Press, 2010, 40, 433-440
- Chen, C.; Intelligence moderates reinforcement learning: a mini-review of the neural evidence; Journal of neurophysiology, Am Physiological Soc, 2015, 113, 3459-3461
- Dombrovski, A. Y.; Clark, L.; Siegle, G. J.; Butters, M. A.; Ichikawa, N.; Sahakian, B. J. & Szanto, K.; Reward/punishment reversal learning in older suicide attempters; American Journal of Psychiatry, Am Psychiatric Assoc, 2010, 167, 699-707
- Huys, Q. J.; Pizzagalli, D. A.; Bogdan, R. & Dayan, P.; Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis; Biology of mood & anxiety disorders, BioMed Central, 2013, 3, 12
- Kessler, R. C.; Berglund, P.; Demler, O.; Jin, R.; Koretz, D.; Merikangas, K. R.; Rush, A. J.; Walters, E. E. & Wang, P. S.; The epidemiology of major depressive disorder: results from the National Comorbidity Survey Replication (NCS-R); Jama, American Medical Association, 2003, 289, 3095-3105
- Kunisato, Y.; Okamoto, Y.; Ueda, K.; Onoda, K.; Okada, G.; Yoshimura, S.; Suzuki, S.-i.; Samejima, K. & Yamawaki, S.; Effects of depression on reward-based decision making and variability of action in probabilistic learning; Journal of behavior therapy and experimental psychiatry, Elsevier, 2012, 43, 1088-1094

References

- MacKay, D. J.; Information theory, inference and learning algorithms; Cambridge university press, 2003
- Mojtabai, R.; Olfson, M. & Han, B.; National trends in the prevalence and treatment of depression in adolescents and young adults; Pediatrics, Am Acad Pediatrics, 2016, e20161878
- O'Doherty, J. P.; Hampton, A. & Kim, H.; Model-based fMRI and its application to reward learning and decision making; Annals of the New York Academy of sciences, Wiley Online Library, 2007, 1104, 35-53
- Olesen, J.; Gustavsson, A.; Svensson, M.; Wittchen, H.-U. & Jönsson, B.; The economic cost of brain disorders in Europe; European journal of neurology, Wiley Online Library, 2012, 19, 155-162
- Olfson, M; Blanco, C; Wall, M & others; National trends in suicide attempts among adults in the united states; JAMA Psychiatry, 2017
- Pizzagalli, D. A.; Jahn, A. L. & O'Shea, J. P.; Toward an objective characterization of an anhedonic phenotype: a signal-detection approach; Biological psychiatry, Elsevier, 2005, 57, 319-327
- Pizzagalli, D. A.; Depression, stress, and anhedonia: toward a synthesis and integrated model; Annual review of clinical psychology, NIH Public Access, 2014, 10, 393
- Seligman, M. E.; Learned helplessness; Annual review of medicine, Annual Reviews 4139 El Camino Way, PO Box 10139, Palo Alto, CA 94303-0139, USA, 1972, 23, 407-412
- Snyder, H. R.; Major depressive disorder is associated with broad impairments on neuropsychological measures of executive function: A meta-analysis and review; American Psychological Association, 2013
- Stankevicius, A.; Huys, Q. J.; Kalra, A. & Seriès, P.; Optimism as a prior belief about the probability of future reward; PLoS Comput Biol, Public Library of Science, 2014, 10, e1003605
- Webb, C. A.; Dillon, D. G.; Pechtel, P.; Goer, F. K.; Murray, L.; Huys, Q. J.; Fava, M.; McGrath, P. J.; Weissman, M.; Parsey, R. & others; Neural correlates of three promising endophenotypes of depression: evidence from the EMBARC study; Neuropsychopharmacology, Nature Publishing Group, 2016, 41, 454

Images

- http://img.zanda.com/item/07091220000003/1024x768/Diagnostic_and_Statistical_Manual_of_Mental_Disorders.jpg
- http://new.paho.org/bulletins/images/stories/Bulletins/Epidemiological_Bulletin/Abril_2009/10_eng.jpg
- https://www.nimh.nih.gov/images/rdoc/rdoc_logo_nobg_153390_1.png
- https://image.flaticon.com/icons/svg/492/492338.svg
- https://image.flaticon.com/icons/svg/116/116321.svg
- https://image.flaticon.com/icons/svg/554/554420.svg
- https://image.flaticon.com/icons/svg/164/164996.svg
- https://upload.wikimedia.org/wikipedia/commons/2/2f/Haemodynamic_response_function.svg
- https://upload.wikimedia.org/wikipedia/commons/f/ff/Gradient_descent.svg
- https://upload.wikimedia.org/wikipedia/commons/d/db/Gradient_ascent_%28contour%29.png