# Modelling behavioural data 

## Quentin Huys <br> MA PhD MBBS MBPsS

Translational Neuromodeling Unit, ETH Zürich
Psychiatrische Universitätsklinik Zürich

## Outline

- An example task
- Why build models? What is a model
- Fitting models
- Validating \& comparing models
- Model comparison issues in psychiatry


## Example task



## Example task



Guitart-Masip, Huys et al. Submitted

## Example task



Guitart-Masip, Huys et al. Submitted

## Example task



## Analysing behaviour

- Standard approach:
- Decide which feature of the data you care about
- Run descriptive statistical tests, e.g.ANOVA
- Many strengths
- Weakness
- Piecemeal, not holistic / global
- Descriptive, not generative
- No internal variables


## Analysing behaviour

- Standard approach:
- Decide which feature of the data you care about
- Run descriptive statistical tests, e.g.ANOVA

- Many strengths
- Weakness
- Piecemeal, not holistic / global
- Descriptive, not generative
- No internal variables


## Models

- Holistic
- Aim to model the process by which the data came about in its "entirety"
- Generative
- They can be run on the task to generate data as if a subject had done the task
- Inference process
- Capture the inference process subjects have to make to perform the task.
- Do this in sufficient detail to replicate the data.
- Parameters
- replace test statistics
- their meaning is explicit in the model
- their contribution to the data is assessed in a holistic manner


## A simple Rescorla-Wagner model

- Q values

$$
\mathcal{Q}_{t}\left(a_{t}, s_{t}\right)=\mathcal{Q}_{t-1}\left(a_{t}, s_{t}\right)+\epsilon\left(r_{t}-\mathcal{Q}_{t-1}\left(a_{t}, s_{t}\right)\right)
$$

$a_{t}$ action on trial $t$; can be either 'go' or 'logo'
$s_{t} \quad$ stimulus presented on trial $t$
$\epsilon \quad$ learning rate

- Key points:


- Q is the key part of the hypothesis
- formally states the learning process in quantitative detail
- formalizes internal quantities that are used in the task


## Actions

- Q values

$$
\mathcal{Q}_{t}\left(a_{t}, s_{t}\right)=\mathcal{Q}_{t-1}\left(a_{t}, s_{t}\right)+\epsilon\left(r_{t}-\mathcal{Q}_{t-1}\left(a_{t}, s_{t}\right)\right)
$$

- Action probabilities:"softmax" of Q value

$$
\begin{aligned}
p\left(a_{t} \mid s_{t}, h_{t}, \beta\right) & =p\left(a_{t} \mid \mathcal{Q}\left(a_{t}, s_{t}\right), \beta\right) \\
& =\frac{e^{\beta \mathcal{Q}\left(a_{t}, s_{t}\right)}}{\sum_{a^{\prime}} e^{\beta \mathcal{Q}\left(a^{\prime}, s_{t}\right)}}
\end{aligned}
$$

- Features:

$$
\begin{aligned}
p\left(a_{t} \mid s_{t}\right) & \propto \mathcal{Q}\left(a_{t}, s_{t}\right) \\
0 & \leq p(a) \leq 1
\end{aligned}
$$

- links learning process and observations
- choices, RTs, or any other data
- link function in GLMs
- man other forms


## Fitting models I

- Maximum likelihood (ML) parameters

$$
\hat{\theta}=\underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta)
$$

- where the likelihood of all choices is:

$$
\begin{aligned}
\mathcal{L}(\theta) & =\log p(\left\{a_{t}\right\}_{t=1}^{T} \mid\left\{s_{t}\right\}_{t=1}^{T},\left\{r_{t}\right\}_{t=1}^{T}, \underbrace{\theta}_{\beta, \epsilon}) \\
& =\log p\left(\left\{a_{t}\right\}_{t=1}^{T} \mid\left\{\mathcal{Q}\left(s_{t}, a_{t} ; \epsilon\right)\right\}_{t=1}^{T}, \beta\right) \\
& =\log \prod_{t=1}^{\beta} p\left(a_{t} \mid \mathcal{Q}\left(s_{t}, a_{t} ; \epsilon\right), \beta\right) \\
& =\sum_{t=1}^{T} \log p\left(a_{t} \mid \mathcal{Q}\left(s_{t}, a_{t} ; \epsilon\right), \beta\right)
\end{aligned}
$$

## Fitting models II

- No closed form
- Use your favourite method
- gradients
- fminunc / fmincon...
- Gradients for RW model

$$
\begin{aligned}
\frac{d \mathcal{L}(\theta)}{d \theta} & =\frac{d}{d \theta} \sum_{t} \log p\left(a_{t} \mid \mathcal{Q}_{t}\left(a_{t}, s_{t} ; \epsilon\right), \beta\right) \\
& =\sum_{t} \frac{d}{d \theta} \beta \mathcal{Q}_{t}\left(a_{t}, s_{t} ; \epsilon\right)-\sum_{a^{\prime}} p\left(a^{\prime} \mid \mathcal{Q}_{t}\left(a^{\prime}, s_{t} ; \epsilon\right), \beta\right) \frac{d}{d \theta} \beta \mathcal{Q}_{t}\left(a^{\prime}, s_{t} ; \epsilon\right) \\
\frac{d \mathcal{Q}_{t}\left(a_{t}, s_{t} ; \epsilon\right)}{d \epsilon} & =(1-\epsilon) \frac{d \mathcal{Q}_{t-1}\left(a_{t}, s_{t} ; \epsilon\right)}{d \epsilon}+\left(r_{t}-\mathcal{Q}_{t-1}\left(a_{t}, s_{t} ; \epsilon\right)\right)
\end{aligned}
$$

## Little tricks

- Transform your variables

$$
\begin{aligned}
\beta & =e^{\beta^{\prime}} \\
& \Rightarrow \beta^{\prime}=\log (\beta) \\
\epsilon & =\frac{1}{1+e^{-\epsilon^{\prime}}} \\
& \Rightarrow \epsilon^{\prime}=\log \left(\frac{\epsilon}{1-\epsilon}\right)
\end{aligned}
$$

$$
\frac{d \log \mathcal{L}\left(\theta^{\prime}\right)}{d \theta^{\prime}}
$$

- Avoid over/underflow

$$
\begin{aligned}
y(a) & =\beta \mathcal{Q}(a) \\
y_{m} & =\max _{a} y(a) \\
p & =\frac{e^{y(a)}}{\sum_{b} e^{y(b)}}=\frac{e^{y(a)-y_{m}}}{\sum_{b} e^{y(b)-y_{m}}}
\end{aligned}
$$

## ML characteristics



## ML characteristics

- ML is asymptotically consistent, but variance high
- I0-armed bandit, infer beta and epsilon

- Hessian $\frac{d^{2}}{d \theta d d_{s}} \mathcal{L}(\theta)$ can be used to derive confidence intervals and identify poorly constrained estimates


## ML characteristics

- ML is asymptotically consistent, but variance high
- I0-armed bandit, infer beta and epsilon

200 trials, I stimulus, 10 actions, learning rate $=.05$, beta $=2$


- Hessian $\frac{d^{2}}{d \theta_{i d} \theta_{j}} \mathcal{L}(\theta)$ can be used to derive confidence intervals and identify poorly constrained estimates
- ML can overfit... more later


## ML characteristics

- ML is asymptotically consistent, but variance high
- I0-armed bandit, infer beta and epsilon

200 trials, I stimulus, 10 actions, learning rate $=.05$, beta $=2$


- Hessian $\frac{d^{2}}{d \theta_{i d} \theta_{j}} \mathcal{L}(\theta)$ can be used to derive confidence intervals and identify poorly constrained estimates
- ML can overfit... more later


## ML characteristics

- ML is asymptotically consistent, but variance high
- I0-armed bandit, infer beta and epsilon
beta and epsilon can
200 trials, I stimulus, 10 actions, learning rate $=.05$, beta $=2$
 trade off

- Hessian $\frac{d^{2}}{d \theta_{i d} d_{j}} \mathcal{L}(\theta)$ can be used to derive confidence intervals and identify poorly constrained estimates
- ML can overfit... more later


## Priors

Not so smooth


Smooth


## Priors

Not so smooth


Smooth


## Priors

Not so smooth


Smooth


## Maximum a posteriori estimate

$$
\begin{aligned}
& \mathcal{P}(\theta)=p\left(\theta \mid a_{1 \ldots T}\right)=\frac{p\left(a_{1 \ldots T} \mid \theta\right) p(\theta)}{\int d \theta p\left(\theta \mid a_{1 \ldots T}\right) p(\theta)} \\
& \log \mathcal{P}(\theta)=\sum_{t=1}^{T} \log p\left(a_{t} \mid \theta\right)+\log p(\theta)+\text { const. } \\
& \frac{\log \mathcal{P}(\theta)}{d \alpha}=\frac{\log \mathcal{L}(\theta)}{d \alpha}+\frac{d p(\theta)}{d \theta}
\end{aligned}
$$

- If likelihood is strong, prior will have little effect
- mainly has influence on poorly constrained parameters
- if a parameter is strongly constrained to be outside the typical range of the prior, then it will win over the prior


## Maximum a posteriori estimate



200 trials, I stimulus, 10 actions, learning rate $=.05$, beta $=2$

$$
m_{\text {beta }}=0, m_{\text {eps }}=-3, n=1
$$

## But

What prior parameters should I use?

## ML characteristics: group data



## ML characteristics: group data



## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability


## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model


## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model
- Summary statistic
- treat parameters as random variable, one for each subject
- overestimates group variance as ML estimates noisy


## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model
- Summary statistic
- treat parameters as random variable, one for each subject
- overestimates group variance as ML estimates noisy
- Random effects
- prior mean = group mean


## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model
- Summary statistic
- treat parameters as random variable, one for each subject
- overestimates group variance as ML estimates noisy
- Random effects
- prior mean = group mean

$$
p\left(\mathcal{A}_{i} \mid \mu_{\theta}, \sigma_{\theta}\right)=\int d \theta_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p\left(\theta_{i} \mid \mu_{\theta}, \sigma_{\theta}\right)
$$

## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model
- Summary statistic
- treat parameters as random variable, one for each subject
- overestimates group variance as ML estimates noisy
- Random effects
- prior mean = group mean

$$
p\left(\mathcal{A}_{i} \mid \mu_{\theta}, \sigma_{\theta}\right)=\int d \theta_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p(\theta_{i} \mid \underbrace{\mu_{\theta}, \sigma_{\theta}}_{\zeta})
$$

## ML characteristics: group data



- Fixed effect
- conflates within- and between- subject variability
- Average behaviour
- disregards between-subject variability
- need to adapt model
- Summary statistic
- treat parameters as random variable, one for each subject
- overestimates group variance as ML estimates noisy
- Random effects
- prior mean = group mean

$$
p\left(\mathcal{A}_{i} \mid \mu_{\theta}, \sigma_{\theta}\right)=\int d \theta_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p(\theta_{i} \mid \underbrace{\mu_{\theta}, \sigma_{\theta}}_{\zeta})
$$

## Estimating the hyperparameters

- MAP

$$
\log \mathcal{P}(\theta)=\mathcal{L}(\theta)+\log \underbrace{p(\theta)}_{=p(\theta \mid \zeta)}+\text { const. }
$$

- Empirical Bayes: set them to ML estimate

$$
\hat{\zeta}=\underset{\zeta}{\operatorname{argmax}} p(\mathcal{A} \mid \zeta)
$$

- where we use all the actions by all the $k$ subjects

$$
\mathcal{A}=\left\{a_{1 \ldots T}^{k}\right\}_{k=1}^{K}
$$

## ML estimate of top-level parameters




$$
\hat{\zeta}=\underset{\zeta}{\operatorname{argmax}} p(\mathcal{A} \mid \zeta)
$$

## Estimating the hyperparameters

- Effectively we now want to do gradient ascent on:

$$
\frac{d}{d \zeta} p(\mathcal{A} \mid \zeta)
$$

- But this contains an integral over individual parameters:

$$
p(\mathcal{A} \mid \zeta)=\int d \theta p(\mathcal{A} \mid \theta) p(\theta \mid \zeta)
$$

- So we need to:

$$
\begin{aligned}
\hat{\zeta} & =\underset{\zeta}{\operatorname{argmax}} p(\mathcal{A} \mid \zeta) \\
& =\underset{\zeta}{\operatorname{argmax}} \int d \theta p(\mathcal{A} \mid \theta) p(\theta \mid \zeta)
\end{aligned}
$$

## Expectation Maximisation

$$
\begin{aligned}
\log p(\mathcal{A} \mid \zeta) & =\log \int d \theta p(\mathcal{A}, \theta \mid \zeta) \\
& =\log \int d \theta q(\theta) \frac{p(\mathcal{A}, \theta \mid \zeta)}{q(\theta)} \\
& \geq \int d \theta q(\theta) \log \frac{p(\mathcal{A}, \theta \mid \zeta)}{q(\theta)} \\
k^{\text {th }} \mathrm{E} \text { step: } q^{(k+1)}(\theta) & \leftarrow p\left(\theta \mid \mathcal{A}, \zeta^{(k)}\right) \\
k^{\text {th }} \mathrm{M} \text { step: } \zeta^{(k+1)} & \leftarrow \underset{\zeta}{\operatorname{argmax}} \int d \theta q(\theta) \log p(\mathcal{A}, \theta \mid \zeta)
\end{aligned}
$$

- There are other approaches
- Monte Carlo
- Analytical conjugate priors
- Variational Bayes
- Iterate between
- Estimating MAP parameters given prior parameters
- Estimating prior parameters from MAP parameters


## EM with Laplace approximation

- E step: $q^{(k+1)}(\theta) \leftarrow p\left(\theta \mid \mathcal{A}, \zeta^{(k)}\right)$
- only need sufficient statistics to perform M step
- Approximate $p\left(\theta \mid \mathcal{A}, \zeta^{(k)}\right) \sim \mathcal{N}\left(\mathbf{m}_{k}, \mathbf{S}_{k}\right)$
- and hence:

E step: $\quad q_{k}(\theta)=\mathcal{N}\left(\mathbf{m}_{k}, \mathbf{S}_{k}\right)$

$$
\begin{aligned}
\mathbf{m}_{k} & \leftarrow \underset{\theta}{\operatorname{argmax}} p\left(\mathbf{a}_{k} \mid \theta\right) p\left(\theta \mid \zeta^{(i)}\right) \\
\mathbf{S}_{k}^{-1} & \left.\leftarrow \frac{\partial^{2} p\left(\mathbf{a}^{k} \mid \theta\right) p\left(\theta \mid \zeta^{(i)}\right)}{\partial \theta^{2}}\right|_{\theta=\mathbf{m}_{k}}
\end{aligned}
$$

## EM with Laplace approximation

- Next update the prior

$$
\text { Prior mean }=\text { mean of MAP estimates }
$$

M step: $\quad \zeta_{\mu}^{(i+1)}=\frac{1}{K} \sum_{k} \mathbf{m}_{k}$

$$
\begin{aligned}
& \qquad \zeta_{\nu^{2}}^{(i+1)}=\frac{1}{N} \sum_{i}\left[\left(\mathbf{m}_{k}\right)^{2}+\mathbf{S}_{k}\right]-\left(\zeta_{\mu}^{(i+1)}\right)^{2} \\
& \text { Prior variance depends on inverse Hessian } S \text { and variance } \quad \text { Take uncertainty of estimates } \\
& \text { into account }
\end{aligned}
$$

- And now iterate until convergence


## Hierarchical / random effects models

- Advantages
- Accurate group-level mean and variance
- Outliers due to weak likelihood are regularized
- Strong outliers are not
- Useful for model selection
- Disadvantages
- Individual estimates $\theta_{i}$ depend on other data, i.e. on $\mathcal{A}_{j \neq i}$ and therefore need to be careful in interpreting these as summary statistics
- Error bars on group parameters (especially group variance) are difficult to obtain
- More involved; less transparent


## Link functions

- Sigmoid $p(a \mid s)=\frac{e^{\beta \mathcal{Q}(a, s)}}{\sum_{a^{\prime}} e^{\beta \mathcal{Q}\left(a^{\prime}, s\right)}}$
- $\mathcal{E}$ - greedy $p(a \mid s)= \begin{cases}c & \text { if } a=\operatorname{argmax}_{a} \mathcal{Q}(a, s) \\ \frac{1-c}{|a|-1} & \text { else }\end{cases}$

- irreducible noise $p(a \mid s)=\frac{1-g}{2}+g \frac{e^{\beta \mathcal{Q}(a, s)}}{\sum_{a^{\prime}} e^{\beta \mathcal{Q}\left(a^{\prime}, s\right)}}$
- critical sanity check I: reasonable link function?


- other link functions for other observations


## Model comparison

- A fit by itself is not meaningful
- Generative test
- qualitative
- Comparisons
- vs random
- vs other model -> test specific hypotheses and isolate particular effects in a generative setting


## Model fit: likelihood

- How well does the model do?
- choice probabilities:

$$
\begin{aligned}
\mathbb{E} p(\text { correct }) & =e^{\mathcal{L}(\hat{\theta}) / K / T} \\
& =e^{\log p(\mathcal{A} \mid \theta) / K / T} \\
& =\left(\prod_{k, t=1}^{K, T} p\left(a_{k, t} \mid \theta_{k}\right)\right)^{\frac{1}{K T}}
\end{aligned}
$$

- typically around 0.65-0.75 for 2-way choice
- for I0-armed bandit example
- pseudo-r²:I-L/R
- better than chance?

$$
\begin{aligned}
\mathbb{E}\left[N_{k}(\text { correct })\right] & =\mathbb{E}\left[p_{k}(\text { correct })\right] T \\
p_{\text {bin }}\left(\mathbb{E}\left[N_{k}(\text { correct })\right] \mid N_{k} d, p_{0}=0.5\right) & <1-\alpha
\end{aligned}
$$



## Generative test

- Model: probability(actions)
- simply draw from this distribution, and see what happens


- Critical sanity test: is the model meaningful?
- Caveat: overfitting


## Overfitting



## Model comparison



## Model comparison

- Averaged over its parameter settings, how well does the model fit the data?

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta) p(\theta \mid \mathcal{M})
$$

- Model comparison: Bayes factors

$$
B F=\frac{p\left(\mathcal{M}_{1} \mid \mathcal{A}\right)}{p\left(\mathcal{M}_{0} \mid \mathcal{A}\right)}=\frac{p\left(\mathcal{A} \mid \mathcal{M}_{1}\right) p\left(\mathcal{M}_{1}\right)}{p\left(\mathcal{A} \mid \mathcal{M}_{2}\right) p\left(\mathcal{M}_{2}\right)}
$$

- Problem:
- integral rarely solvable
- approximation: Laplace, sampling, variational...


## Why integrals? The God Almighty test




## Why integrals? The God Almighty test




## Why integrals? The God Almighty test





These two factors fight it out Model complexity vs model fit

## Bayesian Information Criterion

- Laplace's approximation (saddle-point method)




## Bayesian Information Criterion

- Laplace's approximation (saddle-point method)




## Bayesian Information Criterion

- Laplace's approximation (saddle-point method)




## Bayesian Information Criterion

- Laplace's approximation (saddle-point method)



## Bayesian Information Criterion: one subject

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M})
$$

## Bayesian Information Criterion: one subject

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \overbrace{-}^{\substack{p(\mathcal{A} \mid \theta) p(\theta \mid \mathcal{M}) \\ \text { is propto Gaussian }}}
$$

## Bayesian Information Criterion: one subject

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \overbrace{2}^{\substack{p(\mathcal{A} \mid \theta)) p(\theta \mid \mathcal{M}) \\
\text { is propto Guussian }}} \\
& \approx p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right) p\left(\theta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
\end{aligned}
$$

## Bayesian Information Criterion: one subject

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \underbrace{\substack{p(\mathcal{A} \mid \theta)(\theta(\theta \mid \mathcal{M}) \\
\text { is propto Gaussian }}}_{\substack{p(\theta \mid \mathcal{M})=\text { const. } \\
\text { Model doest prefer } \\
\text { particular }}} \\
& \approx p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right) p\left(\theta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
\end{aligned}
$$

## Bayesian Information Criterion: one subject

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \overbrace{}^{\substack{p(\mathcal{A} \mid \theta) \\
\text { is propto Gaussian }}} \\
& \approx p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right) p\left(\theta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|} \\
\log p(\mathcal{A} \mid \mathcal{M}) & \approx \log p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi)
\end{aligned}
$$

## Bayesian Information Criterion: one subject



## Bayesian Information Criterion: one subject

$$
\begin{aligned}
& p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \\
& \approx p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right) p\left(\theta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|} \\
& \log p(\mathcal{A} \mid \mathcal{M}) \approx \log p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi) \\
& \Sigma_{i i} \propto \frac{1}{T} \Rightarrow \frac{1}{2} \log (|\Sigma|) \approx-\frac{N}{2} \log (T) p(\theta \mid \mathcal{M}) \\
& \text { is propto Gaussian } \\
& \text { Moded doestit prefer } \\
& \text { particular }
\end{aligned}
$$

## Bayesian Information Criterion: one subject

$$
\begin{aligned}
& p(\mathcal{A} \mid \theta) p(\theta \mid \mathcal{M}) \\
& \text { is propto Gaussian } \\
& p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) p(\theta \mid \mathcal{M}) \\
& \approx p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right) p\left(\theta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|} \\
& \log p(\mathcal{A} \mid \mathcal{M}) \approx \log p\left(\mathcal{A} \mid \theta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi) \\
& \Sigma_{i i} \propto \frac{1}{T} \Rightarrow \frac{1}{2} \log (|\Sigma|) \approx-\frac{N}{2} \log (T) \quad \text { Bayesian Information Criterion (BIC) } \\
& \approx \quad-N \quad \text { Akaike Information Criterion (AIC) }
\end{aligned}
$$

## Group data

- Multiple subjects
- Multiple models
- do they use the same model? If not parameters are not comparable
- which model best accounts for all of them?
- Multiple groups
- difference in models?
- difference in parameters?
- $2^{k}$ possible model comparisons
- Multiple parameters
- $2^{\mathrm{k}}$ possible correlations with any one psychometric measure


## Group data - approaches

- Summary statistic
- Treat individual model comparison measure as summary statistics, do ANOVA or t-test
- Fixed effect analysis
- Subject data independent

$$
\begin{aligned}
\log p(\mathcal{A} \mid \mathcal{M}) & =\sum_{i} \log p\left(\mathcal{A}_{i} \mid \mathcal{M}\right) \\
& =\sum_{i} \log \int d \theta_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p\left(\theta_{i} \mid \mathcal{M}_{i}\right) \approx-\frac{1}{2} \sum_{i} \mathrm{BIC}_{i}
\end{aligned}
$$

- Random effects analyses
- Hierarchical prior on group parameters

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \zeta \int d \theta p(\mathcal{A} \mid \theta) p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Hierarchical prior on models

$$
p\left(\mathcal{A}, \mathcal{M}_{k}, r \mid \alpha\right)=p\left(\mathcal{A} \mid \mathcal{M}_{k}\right) p\left(\mathcal{M}_{k} \mid r\right) p(r \mid \alpha)
$$

## Group-level likelihood

- Contains two integrals:
- subject parameters
- prior parameters

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$



## Evaluating p(A|M)

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$



## Evaluating p(A|M)

- Two integrals

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$



## Evaluating $p(A \mid M)$

- Two integrals

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Top level first:



## Evaluating p(A|M)

- Two integrals

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately


## - Top level first:

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})
$$



## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})
$$

## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})
$$

$p(\mathcal{A} \mid \mathcal{M})=\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})$


## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M}) \overbrace{2}^{\substack{p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M}) \\
\text { is propto Gaussian }}} \\
& \approx p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) p\left(\zeta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
\end{aligned}
$$

## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M}) \overbrace{2}^{\substack{p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M}) \\
\text { is propto Gaussian }}} \\
& \approx p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) p\left(\zeta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
\end{aligned}
$$

Model doesn't prefer particular $\zeta$

## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})
$$

$$
\approx p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) p\left(\zeta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
$$

$\log p(\mathcal{A} \mid \mathcal{M}) \approx \log p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi)$

Model doesn't prefer particular $\zeta$

## Evaluating $p(A \mid M)$

- Two integrals
- tricky

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M})
$$

$$
\approx p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) p\left(\zeta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|}
$$

$\log p(\mathcal{A} \mid \mathcal{M}) \approx \log p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi)$

## Evaluating p(A|M)

- Two integrals

$$
p(\mathcal{A} \mid \mathcal{M})=\int d \theta p(\mathcal{A} \mid \theta, \mathcal{M}) \int d \zeta p(\theta \mid \zeta) p(\zeta \mid \mathcal{M})
$$

- Step by step: approximating levels separately
- Approximate at the top level
- less action

$$
\begin{aligned}
p(\mathcal{A} \mid \mathcal{M}) & =\int d \zeta p(\mathcal{A} \mid \zeta, \mathcal{M}) p(\zeta \mid \mathcal{M}) \\
& \approx p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) p\left(\zeta^{M L} \mid \mathcal{M}\right) \times \sqrt{(2 \pi)^{N}|\Sigma|} \\
\log p(\mathcal{A} \mid \mathcal{M}) & \approx \log p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right)+\frac{1}{2} \log (|\Sigma|)+\frac{N}{2} \log (2 \pi)
\end{aligned}
$$

## Approximating level I

- Still leaves the first level:
- Approximate integral by sampling, e.g. importance sampling for few dimensions $(<10)$

$$
\begin{aligned}
\log p\left(\mathcal{A} \mid \zeta^{M L}, \mathcal{M}\right) & =\log \int d \theta p(\mathcal{A} \mid \theta) p\left(\theta \mid \zeta^{M L}\right) \\
& \approx \log \frac{1}{B} \sum_{b=1}^{B} p\left(\mathcal{A} \mid \theta^{b}\right) \\
\theta^{b} & \sim p\left(\theta \mid \zeta^{M L}\right)
\end{aligned}
$$

## Group-level BIC

$$
\begin{aligned}
\log p(\mathcal{A} \mid \mathcal{M}) & =\int d \boldsymbol{\zeta} p(\mathcal{A} \mid \boldsymbol{\zeta}) p(\boldsymbol{\zeta} \mid \mathcal{M}) \\
& \approx-\frac{1}{2} \mathrm{BIC}_{\mathrm{int}} \\
& =\log \hat{p}\left(\mathcal{A} \mid \hat{\boldsymbol{\zeta}}^{M L}\right)-\frac{1}{2}|\mathcal{M}| \log (|\mathcal{A}|)
\end{aligned}
$$

## Example task



## Example task



Guitart-Masip, Huys et al. Submitted

## Model validation: generating data



Guitart-Masip et al. 20I I, Guitart-Masip, Huys et al. Submitted

## Model validation: generating data



$$
p\left(\text { go } \mid s_{t}\right) \propto \mathcal{Q}_{t}\left(\text { go } \mid s_{t}\right)+\operatorname{bias}(\text { go })
$$




Guitart-Masip et al. 20I I, Guitart-Masip, Huys et al. Submitted

## Model validation: generating data



$$
\begin{aligned}
p\left(\mathrm{go} \mid s_{t}\right) & \propto \mathcal{Q}_{t}\left(\mathrm{go} \mid s_{t}\right)+\operatorname{bias}(\mathrm{go})+\mathcal{V}_{t}\left(s_{t}\right) \\
\mathcal{V}_{t}\left(s_{t}\right) & =\mathcal{V}_{t-1}\left(s_{t}\right)+\epsilon\left(r_{t}-\mathcal{V}_{t-1}\left(s_{t}\right)\right)
\end{aligned}
$$




Guitart-Masip et al. 20I I, Guitart-Masip, Huys et al. Submitted

## Model validation: generating data

$\mathrm{P}(\mathrm{go}) \propto$ value of stimulus



$$
\begin{aligned}
p\left(\mathrm{go} \mid s_{t}\right) & \propto \mathcal{Q}_{t}\left(\mathrm{go} \mid s_{t}\right)+\operatorname{bias}(\mathrm{go})+\mathcal{V}_{t}\left(s_{t}\right) \\
\mathcal{V}_{t}\left(s_{t}\right) & =\mathcal{V}_{t-1}\left(s_{t}\right)+\epsilon\left(r_{t}-\mathcal{V}_{t-1}\left(s_{t}\right)\right)
\end{aligned}
$$




Guitart-Masip et al. 20I I, Guitart-Masip, Huys et al. Submitted

## Model comparison: overfitting?



## Model comparison: overfitting?



## Model comparison: overfitting?



## Model comparison: overfitting?





## How does it do?


₹ $\mathrm{EBC}_{\mathrm{L}}$


Int


## How does it do?



Fitted by EM... too nice?
$\Sigma_{\mathrm{i}} \mathrm{BIC}_{\mathrm{LL}}^{\mathrm{i}}$


Int


## Top-level Laplacian approximation

- Estimating the top-level determinant
- using 2nd order finite differences

$$
\begin{aligned}
\left.\frac{d^{2}}{d h_{i j}^{2}} p(\mathcal{A} \mid \boldsymbol{\zeta})\right|_{\boldsymbol{\zeta}=\hat{\boldsymbol{\zeta}}^{M L}} \approx & \frac{1}{\delta^{2}}\left[p\left(\mathcal{A} \mid \hat{\boldsymbol{\zeta}}^{M L}+\delta \mathbf{e}_{i}\right)-\right. \\
& \left.2 p\left(\mathcal{A} \mid \hat{\boldsymbol{\zeta}}^{M L}\right)+p\left(\mathcal{A} \mid \hat{\boldsymbol{\zeta}}^{M L}-\delta \mathbf{e}_{j}\right)\right]
\end{aligned}
$$

- the shifted likelihoods can be evaluated by shifting the samples.


## Group level errors



| - Increasing Subject \# N |
| :--- |
| - Increasing Trials T |



Total \# of observations ( $\mathrm{N} \times \mathrm{T}$ )


Total \# of observations ( $\mathrm{N} \times \mathrm{T}$ )

## Posterior distribution on models

- Generative model for models



## Bayesian model selection - equations

- Write down joint distribution of generative model
- Variational approximations lead to set of very simple update equations
- start with flat prior over model probabilities

$$
\alpha=\alpha_{0}
$$

- then update

$$
\begin{aligned}
u_{k}^{i} & =\left(\int d \theta_{i} p\left(\mathcal{A}_{i}, \theta_{i} \mid \mathcal{M}_{k}\right)\right) \exp \left(\Psi\left(\alpha_{k}\right)-\Psi\left(\sum_{k} \alpha_{k}\right)\right) \\
\alpha_{k} & \leftarrow \alpha_{0, k}+\sum_{i} \frac{u_{k}^{i}}{\sum_{k} u_{k}^{i}}
\end{aligned}
$$

## Group Model selection

Integrate out your parameters

## Questions in psychiatry l: regression

- Parametric relationship with other variables $\psi$
- do standard second level analyses
- can use Hessians to determine weights

$$
\begin{aligned}
\text { E step: } \quad \begin{aligned}
q_{k}(\theta) & =\mathcal{N}\left(\mathbf{m}_{k}, \mathbf{S}_{k}\right) \\
\mathbf{m}_{k} & \leftarrow \underset{\theta}{\operatorname{argmax}} p\left(\mathbf{a}_{k} \mid \theta\right) p\left(\theta \mid \zeta^{(i)}\right) \\
\mathbf{S}_{k}^{-1} & \left.\leftarrow \frac{\partial^{2} p\left(\mathbf{a}^{k} \mid \theta\right) p\left(\theta \mid \zeta^{(i)}\right)}{\partial \theta^{2}}\right|_{\theta=\mathbf{m}_{k}}
\end{aligned} .
\end{aligned}
$$

- better: compare two models

$$
\begin{array}{rc}
\text { Model 1: } & \prod_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p\left(\theta_{i} \mid \mu_{0}, \sigma\right) \\
\text { i.e. } & \theta_{i} \sim \mathcal{N}\left(\mu_{0}, \sigma\right) \\
\text { Model 2: } & \prod_{i} p\left(\mathcal{A}_{i} \mid \theta_{i}\right) p\left(\theta_{i} \mid \mu_{0}, c, \sigma, \psi_{i}\right) \\
\text { i.e. } & \theta_{i} \sim \mathcal{N}\left(\mu_{0}+c \psi_{i}, \sigma\right)
\end{array}
$$

## Regression

- Standard regression analysis:

$$
\mathbf{m}_{i}=\mathbf{C r}_{i}+\Sigma^{1 / 2} \boldsymbol{\eta} \quad \forall i
$$

- Including uncertainty about each subject's inferred parameters

$$
\mathbf{m}_{i}=\mathbf{C r}_{i}+\left(\Sigma^{1 / 2}+\mathbf{S}_{i}^{1 / 2}\right) \boldsymbol{\eta} \quad \forall i
$$

- Careful: Finite difference estimates S can be noisy!
- regularize...


## Questions in psychiatry II: group differences

- Do groups differ in terms of parameter(s)?
- Cannot compare parameters across different models
- even very similar parameters can account for different effects




- For models with k parameters, there are $2^{\mathrm{k}}$ possible comparisons
- multiple comparisons?
- posterior over models (Stephan et al. 2009)


## Group differences in parameters

- Are two groups similar in parameter x?
- ANOVA: compare likelihood of two means to likelihood of one global mean. Take degrees of freedom into account.
- But: this tries to account for the parameters with one or two groups, not for the data
- Compare models with separate or joint parameter \& prior:

| Model 1 | $\varepsilon$ | $\beta_{1}, \beta_{2}$ |
| :---: | :---: | :---: |
| Model 2 | $\varepsilon$ | $\beta$ |

## Questions in psychiatry III: Classification

- Who belongs to which of two groups?
- How many groups are there?


## Model comparison again

- What is 'significant'?

$$
\begin{aligned}
& B F= \frac{p\left(\mathcal{A} \mid \mathcal{M}_{1}\right)}{p\left(\mathcal{A} \mid \mathcal{M}_{2}\right)} \\
& p(\Lambda<\eta)
\end{aligned}
$$

| $\log _{10}\left(B_{10}\right)$ | $B_{10}$ |
| :--- | :--- |
| 0 to $1 / 2$ | 1 to 3.2 |
|  |  |
| $1 / 2$ to 1 | 3.2 to 10 |
| 1 to 2 | 10 to 100 |
| $>2$ | $>100$ |

Evidence against $H_{0}$
Not worth more than a bare
mention
Substantial
Strong
Decisive
Kaas and Raftery 95

- "Spread of effect" in group comparisons
- Better model does not mean a behavioural effect is concentrated in one parameter
- Obvious raw differences spread between parameters


## Behavioural data modelling

- Are no panacea
- statistics about specific aspects of decision machinery
- only account for part of the variance
- Model needs to match experiment
- ensure subjects actually do the task the way you wrote it in the model
- model comparison
- Model = Quantitative hypothesis
- strong test
- need to compare models, not parameters
- includes all consequences of a hypothesis for choice


## Modelling in psychiatry

- Hypothesis testing
- otherwise untestable hypotheses
- internal processes
- Limited by data quality
- Look for strong behaviours, not noisy
- "Holistic" testing of hypotheses
- Marr's levels
- physical
- algorithm
- computational

