Automated Adaptive Design in Real Time

Desi R. Ivanova[1] (desi.ivanova@stats.ox.ac.uk)

Adam Foster[1], Steven Kleinegesse[2], Michael U. Gutmann[2], Tom Rainforth[1]

[1] Department of Statistics, University of Oxford [2] School of Informatics, University of Edinburgh

Designing good experiments

Fundamental challenge in science

Data enables
- Learning about complex systems
- Better decision making
- Scientific progress

But data collection is expensive

Good experiments help us minimise costs and maximise quality of data

Next best question?

Smart surveys

Most promising compound?

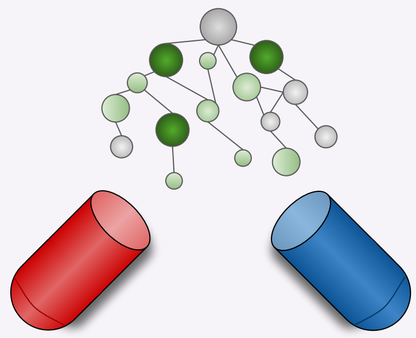

Drug discovery

Next document to annotate?

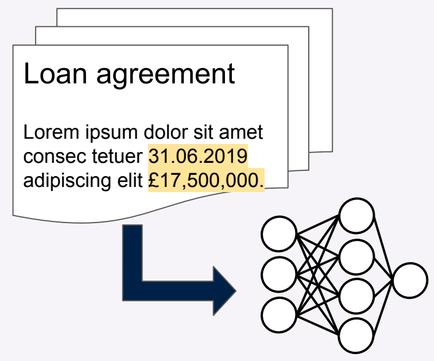

Active learning

...

...

...

What is a good design?

Bayesian Experimental Design

Principled mathematical framework

BED: Optimal design maximises expected information gain, **EIG**

Example: delay discounting
- Questions of the form "Do you prefer **£R now** or £100 in **D days**?"
- Design = {£R, D}
- Outcome = {Reject, Accept}

Live experiment
Traditional BED framework

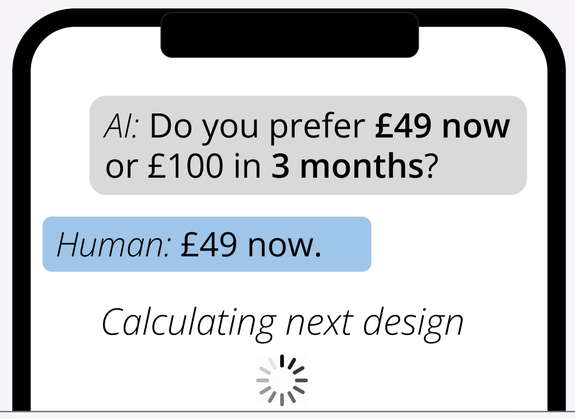

Experiment lifecycle: Traditional BED

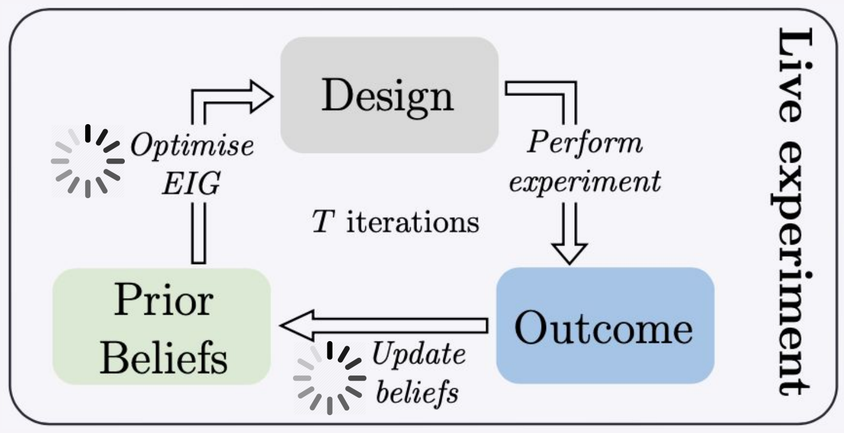

Prohibitively expensive to run in real time

Our solution: Policy-based BED

Critical change from past methods

ICML International Conference On Machine Learning

Key idea: optimise a design policy, instead of individual designs

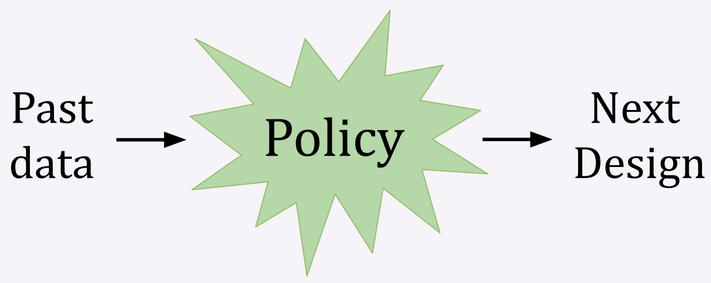

✓ Policy can be **optimised upfront**

✓ Make design decisions with a **single function evaluation** at deployment

Theorem: Total Information

Total EIG of a T-step experiment is the sum of the EIGs for each experiment iteration.

$$\sum_{t=1}^{T} \mathrm{EIG}_t(\mathrm{design}_t) = \mathrm{EIG}_{1\to T}(\mathrm{policy})$$

✓ Generalises traditional BED

✓ Single expression for EIG, allows end-to-end training

Experiment lifecycle: Policy-based BED

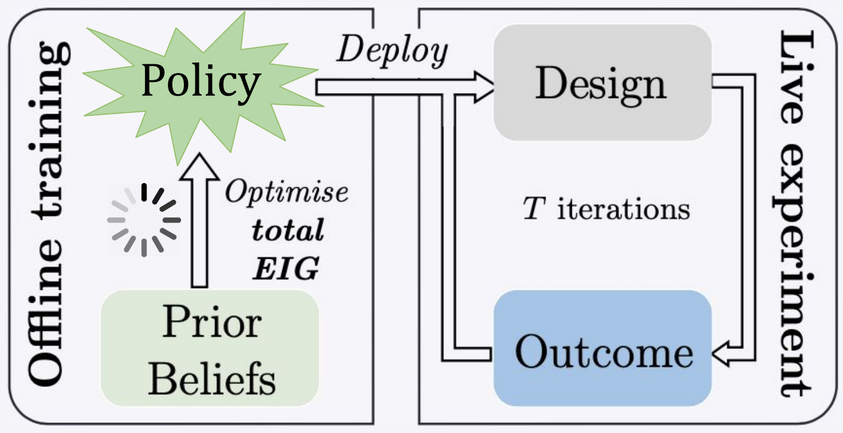

What is the policy and how do we train it?

Deep Adaptive Design (DAD)

A framework for policy-based BED

ICML International Conference On Machine Learning

NEURAL INFORMATION PROCESSING SYSTEMS

Key idea: parametrise the policy with a neural network—the DAD network

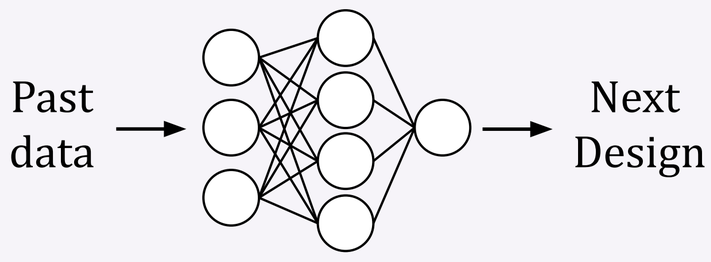

✓ Flexible and customisable

✓ Introduced a **novel architecture**, ensuring efficient training

Live experiment
Policy-based BED with DAD

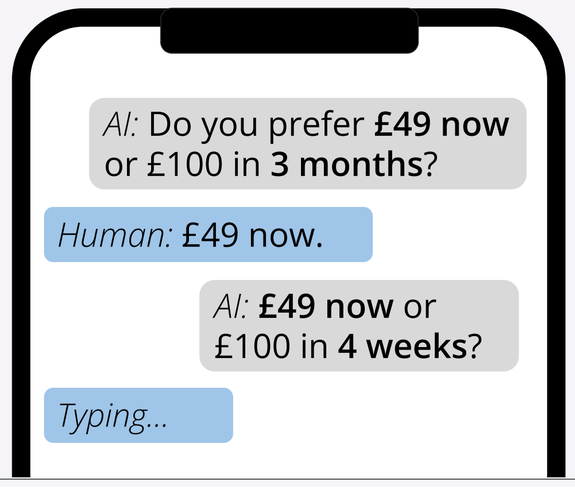

Benefits: DAD vs Traditional BED

✓ Over a **million times faster**
- Design decisions in milliseconds (vs ~3m)

✓ **~20% higher** information gain
- Higher quality data, lower costs

Impact: DAD has potential to revolutionise the way experiments are conducted

➔ Highly general method (& open source!)

➔ Contribution to scientific knowledge

➔ Wider economic and societal benefits

Next: extensions and more applications