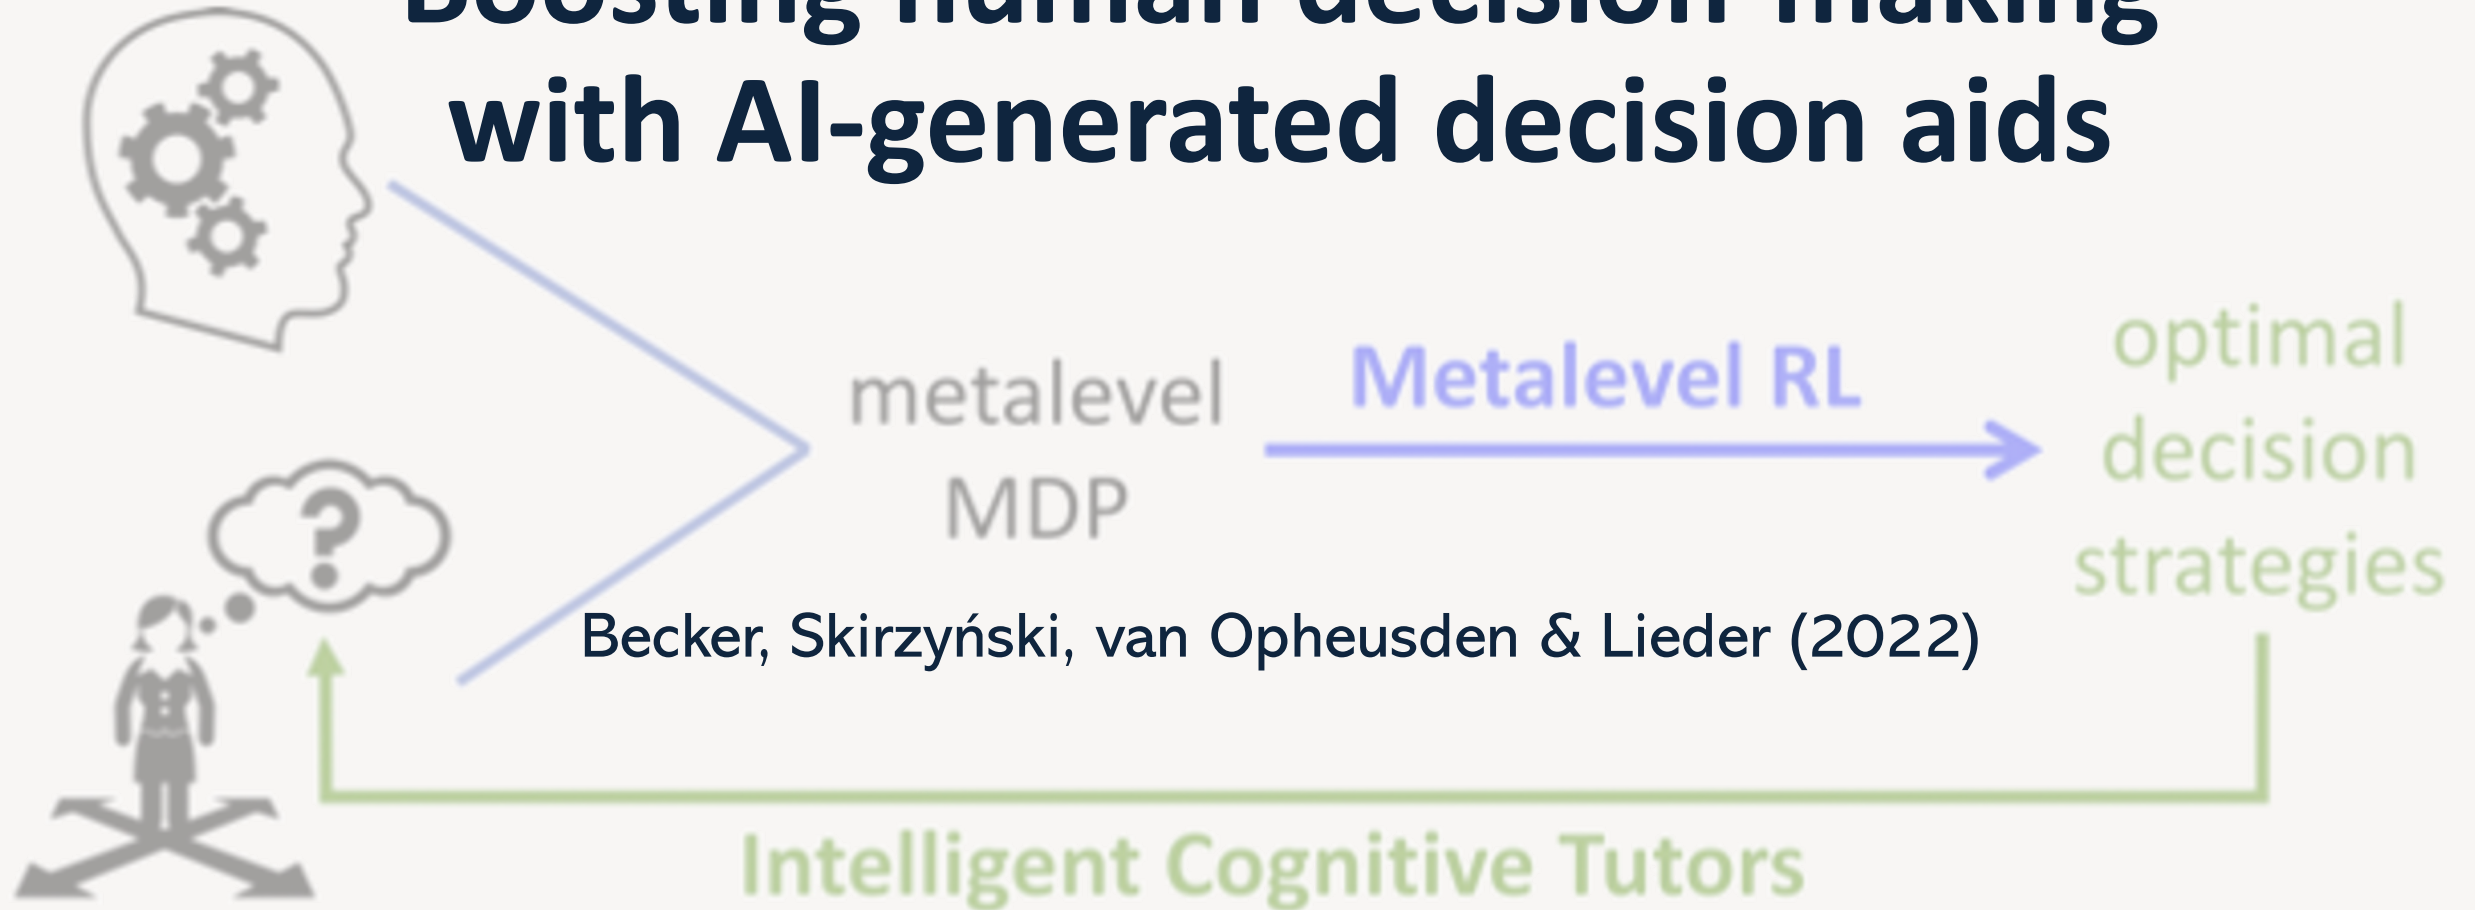


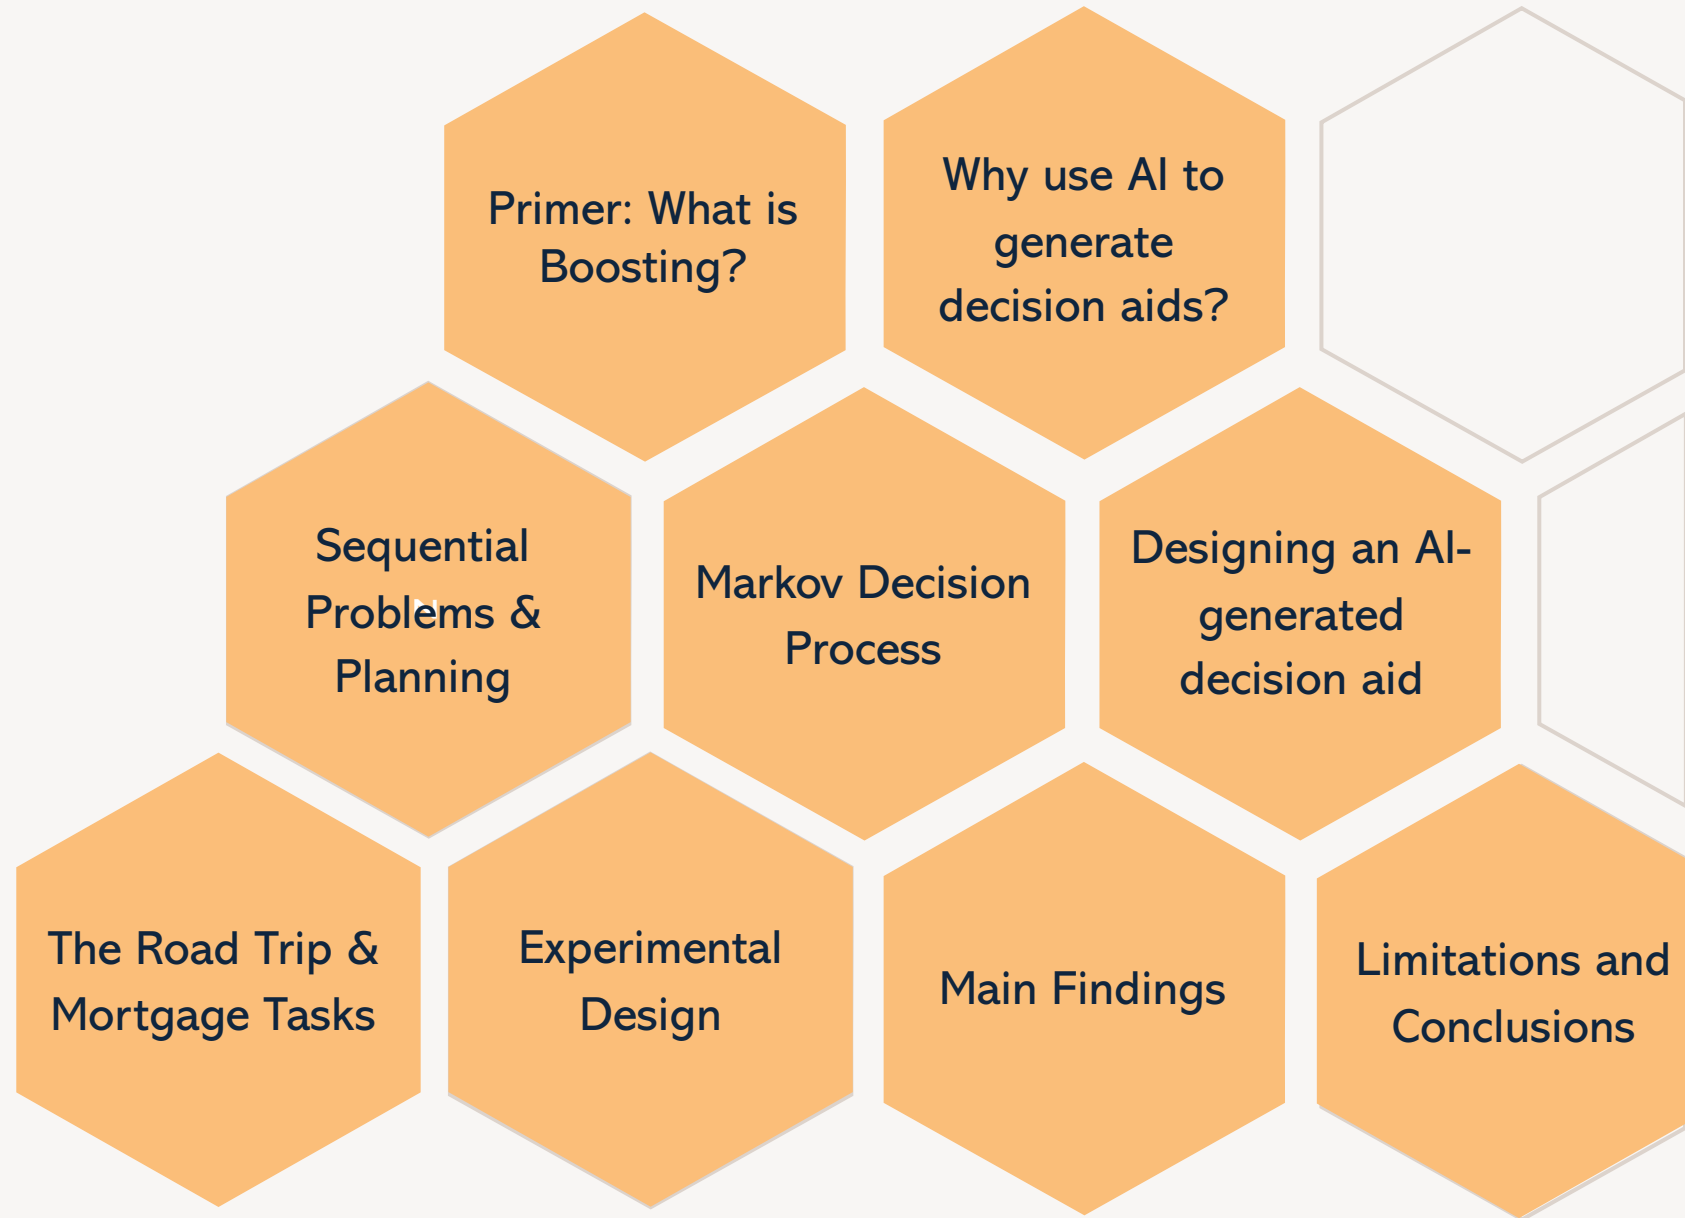
Boosting human decision-making with AI-generated decision aids



“Our work provides a proof of concept that it is possible to discover and describe [resource-] rational decision strategies automatically.”

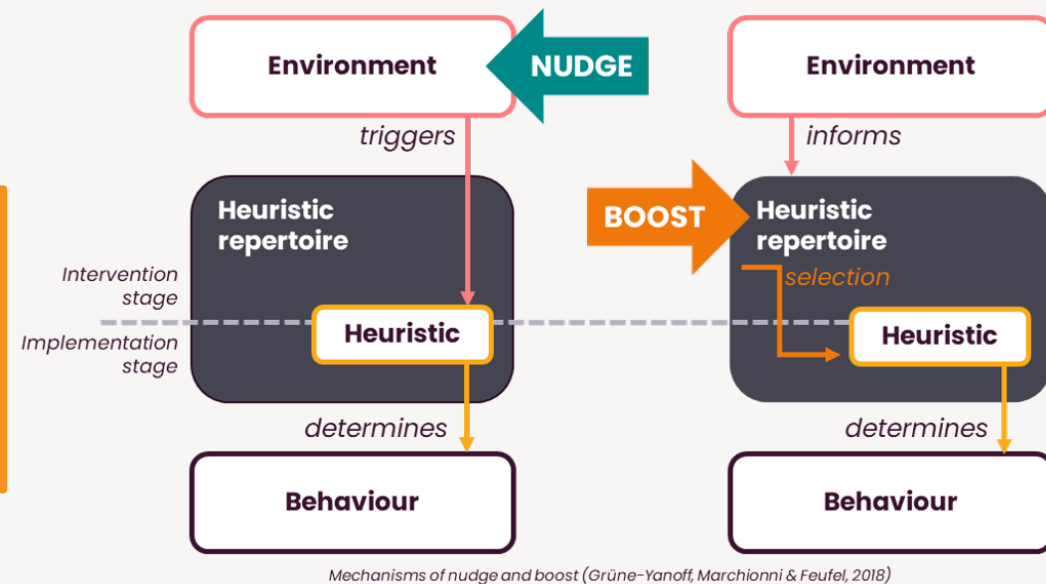


Agenda



Primer: What is Boosting?

“The objective of boosts is to foster people’s competence to make their own choices - that is, to exercise their own agency.” – Hertwig and Grüne-Yanoff (2017)



- Boosting differs from Nudges in many ways...
 - Nudges are often external interventions that result in one-off behavioural changes.
 - Boosting uses training or decision aids to induce lasting behavioural changes.
- *Example of a Nudge:* set organ donation as a default option and people have to opt-out
- *Example of a Boost:* tell people to choose credit cards with low interest rates (rule of thumb)

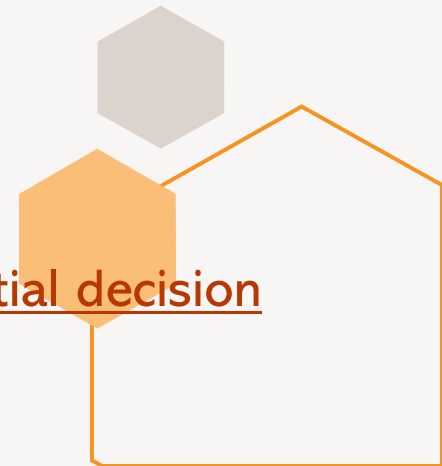
Why use AI to generate decision aids?

- **Motivation:**

- People make sub-optimal decisions. Boosting is shown to improve people's decisions but decision aids are tedious to design manually, especially for complex decision problems.
- Finding an optimal strategy may be more difficult than computing the optimal decision.
Why not use AI to discover efficient heuristics to (approximately) solve problems?


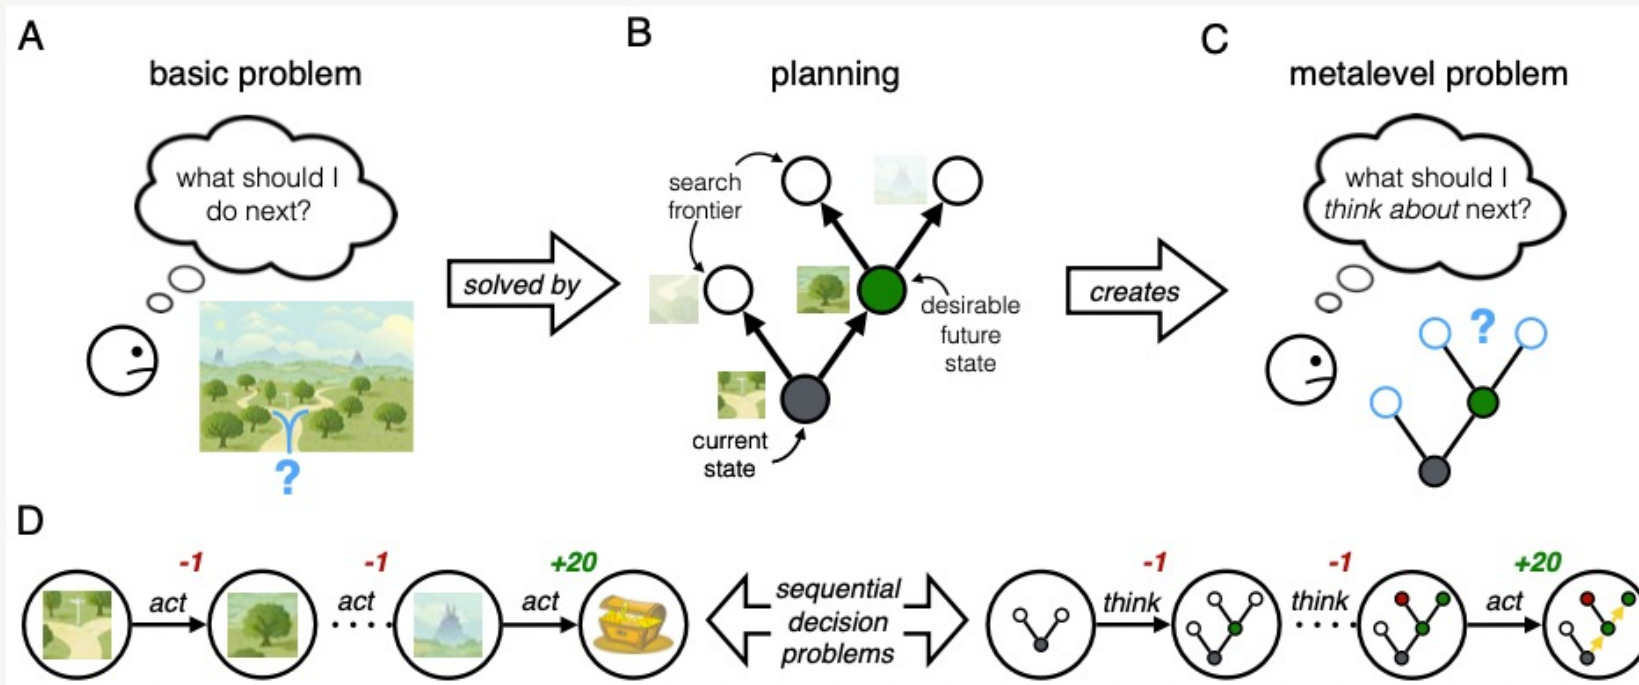
- **Research Questions:**

1. Can we use AI (specifically, Reinforcement Learning) to generate decision aids?
 - How can we interpret RL policies? (Explainable RL)
 - Can people understand the AI-generated decision aids?
2. Do AI-generated decision aids improve people's decisions when facing sequential decision problems?



Sequential Problems & Planning (Metalevel)


- Based on **Resource-Rational Theory** of optimal human planning (Callaway et al., 2021)
- When faced with a sequential decision problem, people **plan ahead** to find the best strategy
 - Planning itself can be modelled into a sequential decision problem where the decision-maker needs to decide what information to gather and when to stop planning (metalevel)



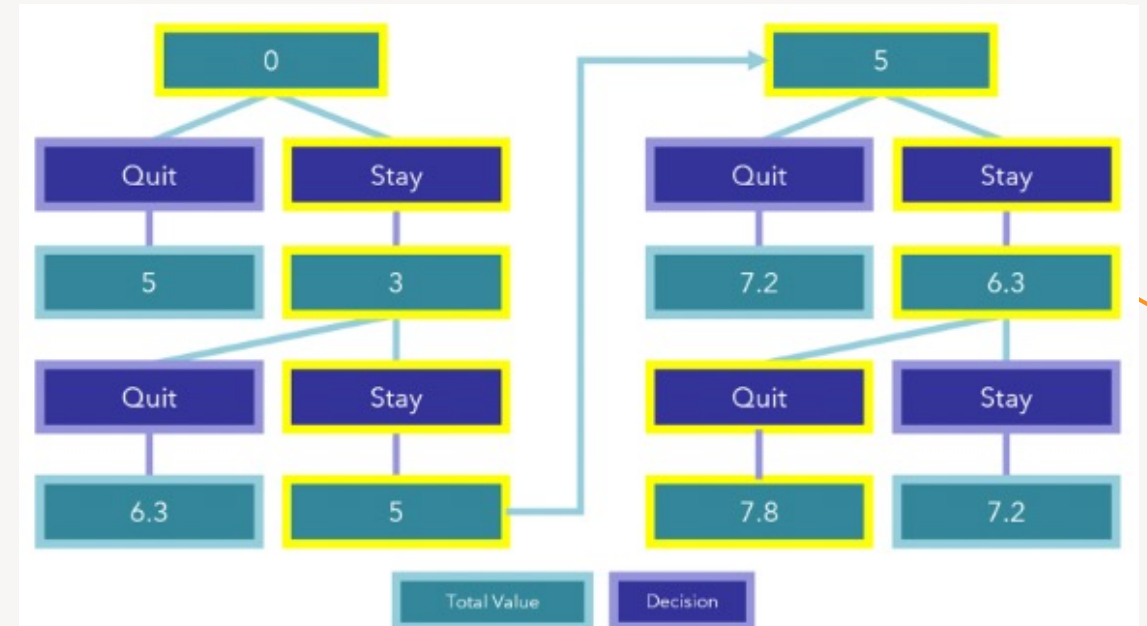

Caveat: structure of the problem is pre-defined (sequential), optimal strategy and solution are known

Source: <https://towardsdatascience.com/introduction-to-reinforcement-learning-markov-decision-process-44c533ebf8da>

Markov Decision Process



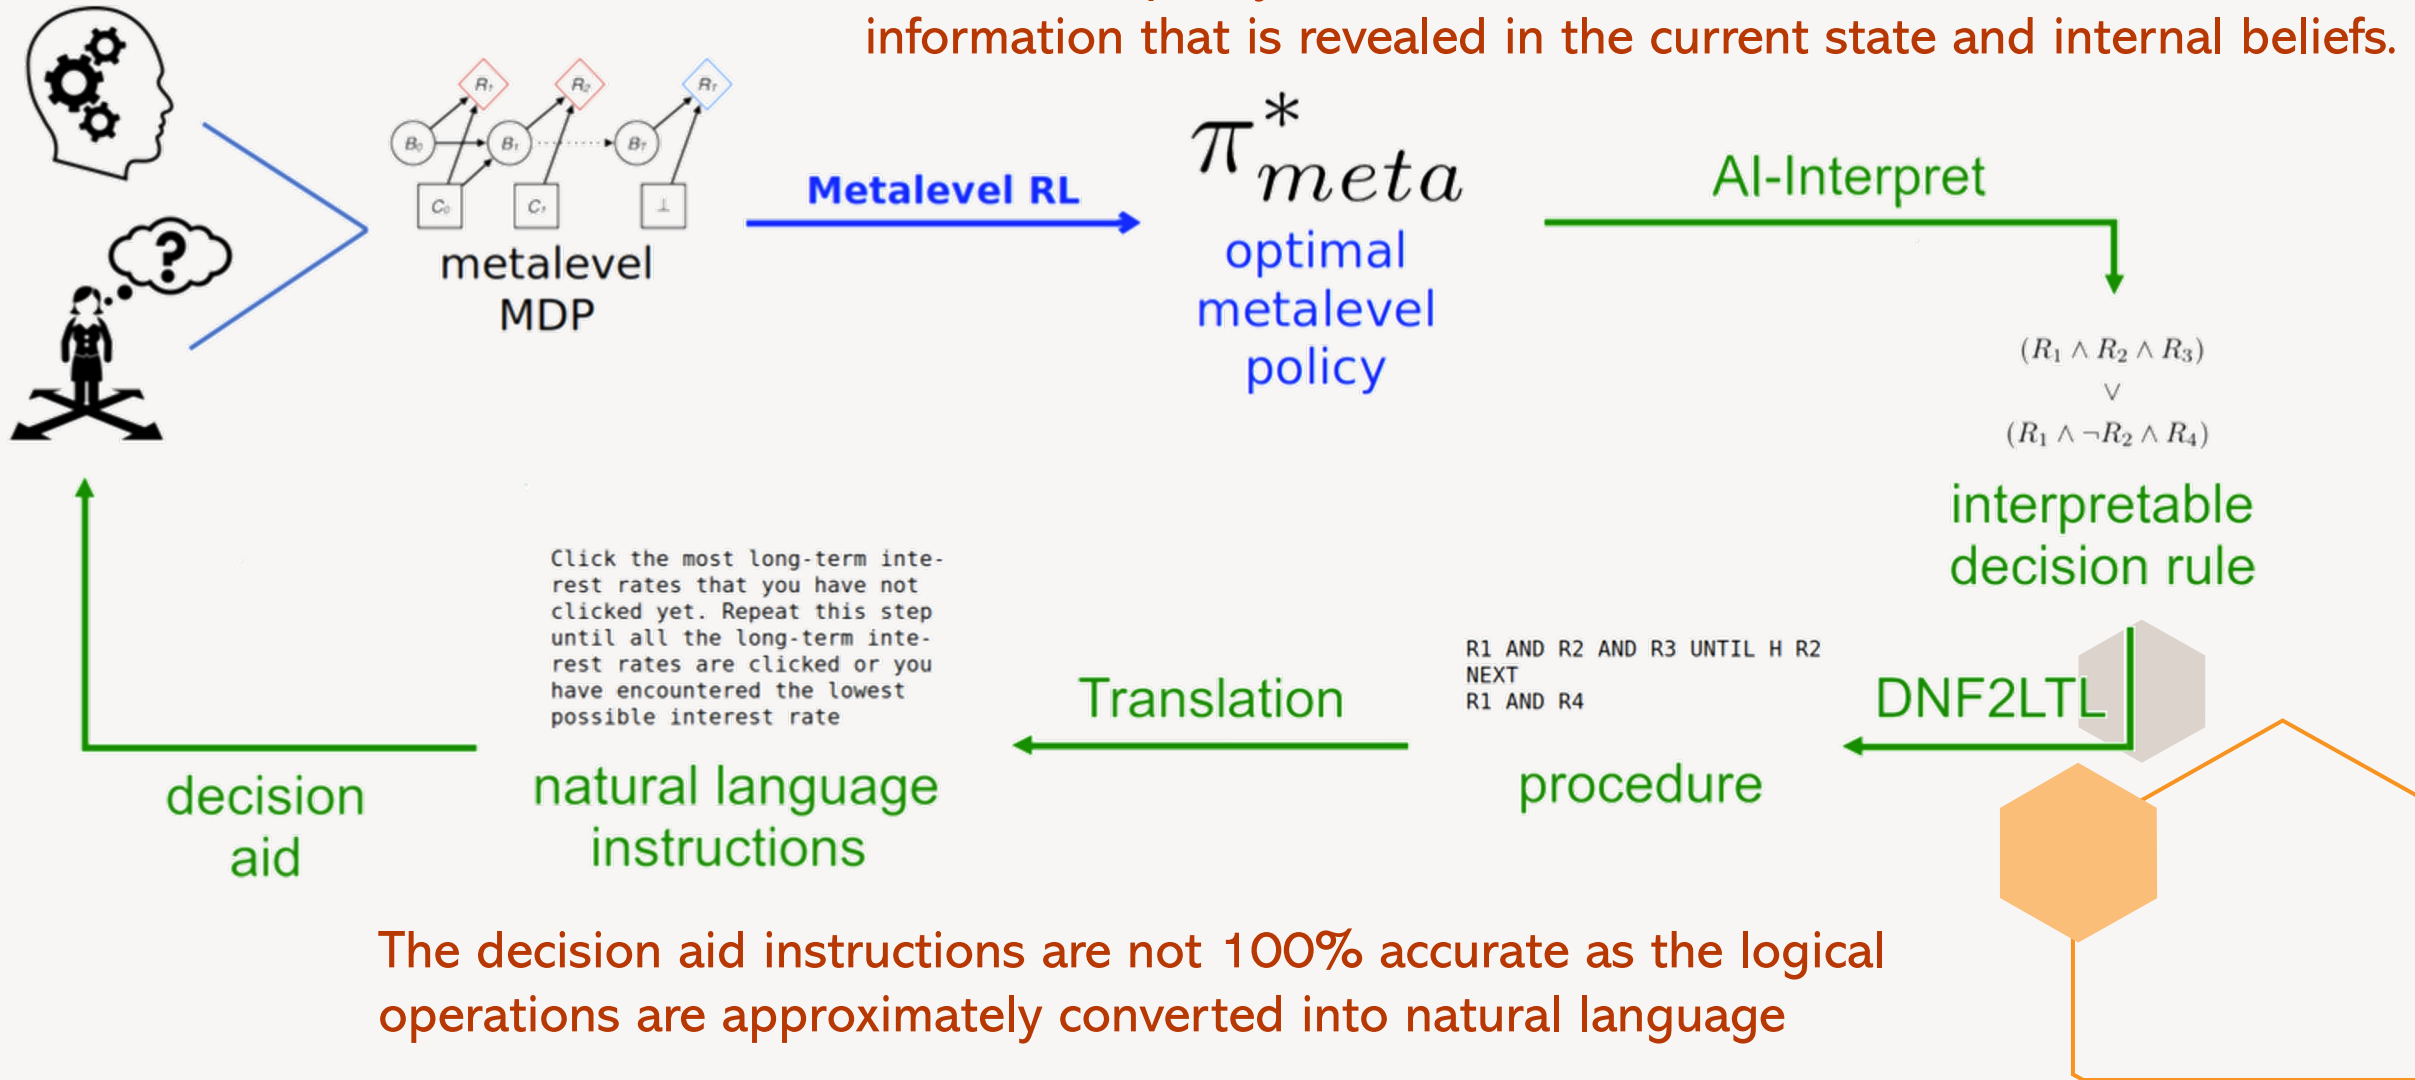
- MDP models the “sequential interaction between an agent and a stochastic environment” (Callaway et al., 2021)
- Use RL: agent maps states to actions to maximise the value of all states at the same time
 - This optimal policy is then converted into a decision aid (using a decision tree)
- Example: Play the following game over 4 rounds



Source: <https://neptune.ai/blog/markov-decision-process-in-reinforcement-learning>

Designing an AI-generated decision aid

A metalevel policy: choose the next action based on the information that is revealed in the current state and internal beliefs.



The Road Trip & Mortgage Tasks

For both tasks, participants pay a price for every click that reveals information

Road Trip Task (Travel Salesman Problem)

Mortgage Task

- Find the cheapest route to any one of the two
- Find the cheapest mortgage

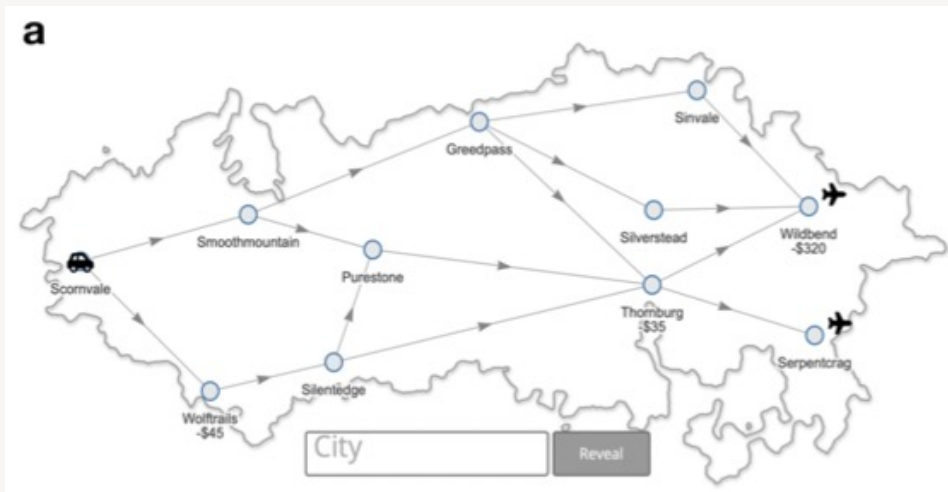
Airports

2 mortgage options with rates increasing,

- One a
- other

What do you think is the best strategy (maximises rewards & minimises clicking cost)?

- Each node has a cost ("hotel cost")



b

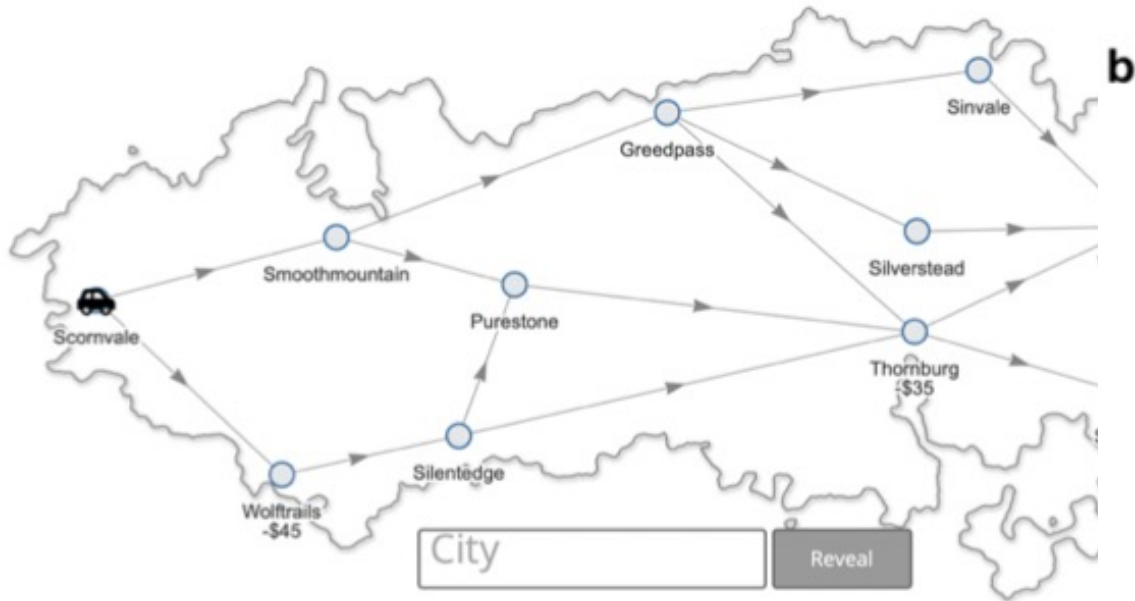
Mortgage plan A	2022 ?	2023 - 2027 ?	2028 - 2052 ?
Mortgage plan B	2022 ?	2023 - 2027 ?	2028 - 2052 ?
Mortgage plan C	2022 ?	2023 - 2027 ?	2028 - 2052 ?

Tasks with Decision Aids

a

Advice for scoring a high bonus:

Look up the prices of the most distant hotels that you have not looked up yet. Repeat this step until all the distant hotels' prices are looked up or you have encountered the lowest possible hotel price.



b

Advice for scoring a high bonus:

Click the most long-term interest rates that you have not clicked yet. Repeat this step until all the long-term interest rates are clicked or you have encountered the lowest possible interest rate

Mortgage plan A	2022 ?	2023 - 2027 ?	2028 - 2052 ?
Mortgage plan B	2022 ?	2023 - 2027 ?	2028 - 2052 ?
Mortgage plan C	2022 ?	2023 - 2027 ?	2028 - 2052 ?

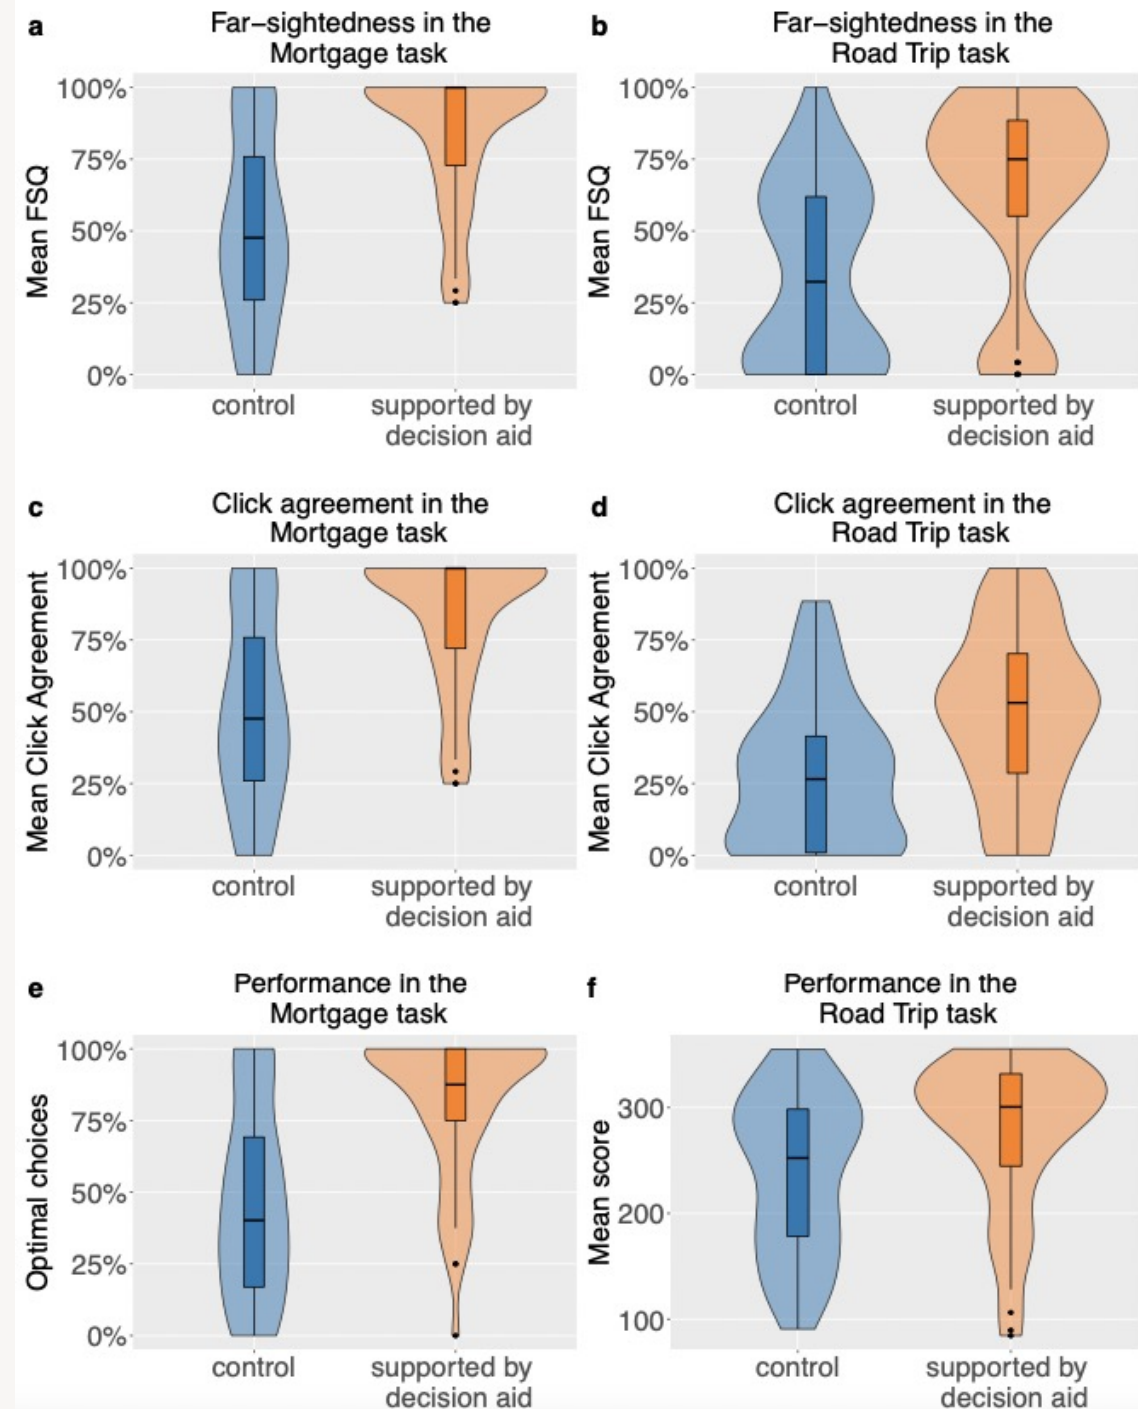
Experimental Design

- Between-subjects design: control vs. decision aid treatment
 - Participants are freely to use the decision aid in the decision-aid treatment
- HOWEVER, participants who are in the decision aid treatment are **primed**
 - '**Advice for scoring a high bonus**' is visible on the task interface
 - The instructions tell them that the decision aid is a near-optimal strategy for gathering information.
- Use clicking activity to track how people plan, gather information, and make their final decision.
- No feedback




Main Findings

- *All differences are significant*
- [a & b] When the decision aid is present, most participants clicked on the final nodes first
 - Road Trip: airport nodes
 - Mortgage: 30yr interest rates
- [c & d] Most participants followed the decision aid strategies
- [e & f] Greater improvement in performance in Mortgage Task compared to Road Trip Task
 - Perhaps participants lacked financial literacy (no test to confirm)




Limitations & Conclusion


- Participants might have followed the instructions without understanding why they make any sense.
 - For the Mortgage Task, were the participants following the decision aid strategy because they don't know better or they actually understand how interest or compounding works (domain knowledge)
 - Not differentiated in the study
 - For the Road Trip Task: it is a more intuitive task, no domain knowledge required -> but need to check if participants know that the cheapest airport hotel is \$20.



Limitations & Conclusion

- Decision aids are automatically designed based on the *resource-rational theory*
 - There may be other (approximating) algorithms that discover better strategies
 - Carsten mentioned the “alignment problem”: the optimal strategy may not account for people’s preferences and objectives
 - Do the decision aids “economically” improve people’s choice? *Not really the case in the Road Trip Task...*
- Need to follow the general metalevel MDP framework to use this method
- How can we translate more complex RL policies into decision aids?





Thank you!

**Comments &
Questions?**

