

Real-time Planning as Data-driven Decision-making

Maximilian Fickert¹ and Tianyi Gu² and Leonhard Staut¹ and
Sai Lekyang² and Wheeler Ruml² and
Jörg Hoffmann¹ and Marek Petrik²



The authors

Introduction

■ The Authors

■ Real-time Search

■ An Example

■ This work

Nancy

Results

Conclusions

Maximilian Fickert



Tianyi Gu



Leonhard Staut



Sai Lekyang



Wheeler Ruml



Jorg Hoffmann



Marek Petrik



Problem Setting: Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

■ This work

Nancy

Results

Conclusions

Environments:

single agent

discrete state, discrete action

fully observability

online planning: interleaving planning and execution

deterministic state transition

Real-time heuristic search:

can't wait until have full plan

return the next action **within a given time bound**

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

■ This work

Nancy

Results

Conclusions

An example: path finding



agent performs search for a bounded time

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

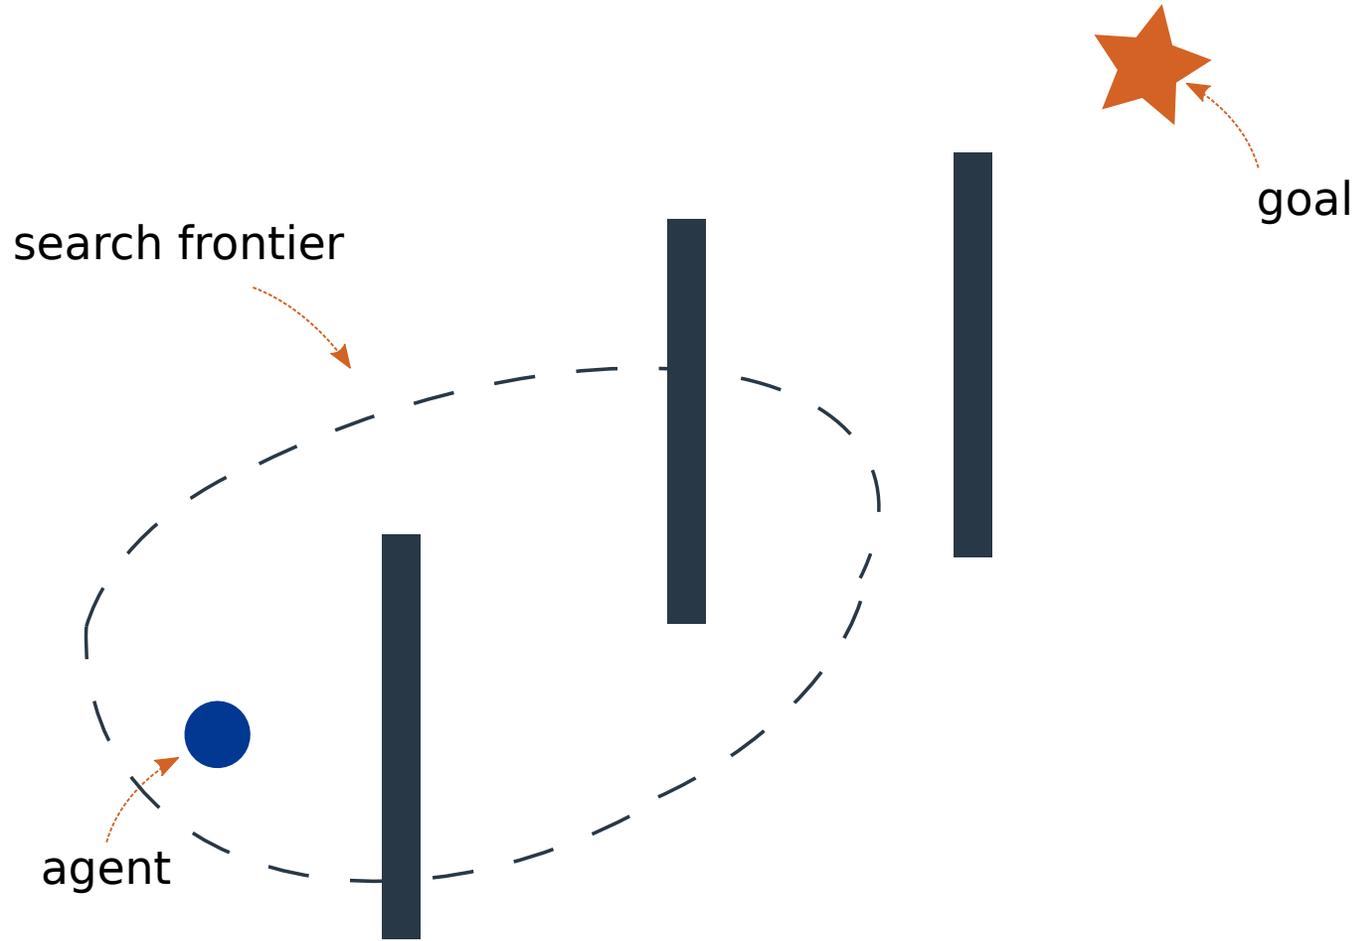
■ This work

Nancy

Results

Conclusions

An example: path finding



agent performs search for a bounded time

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

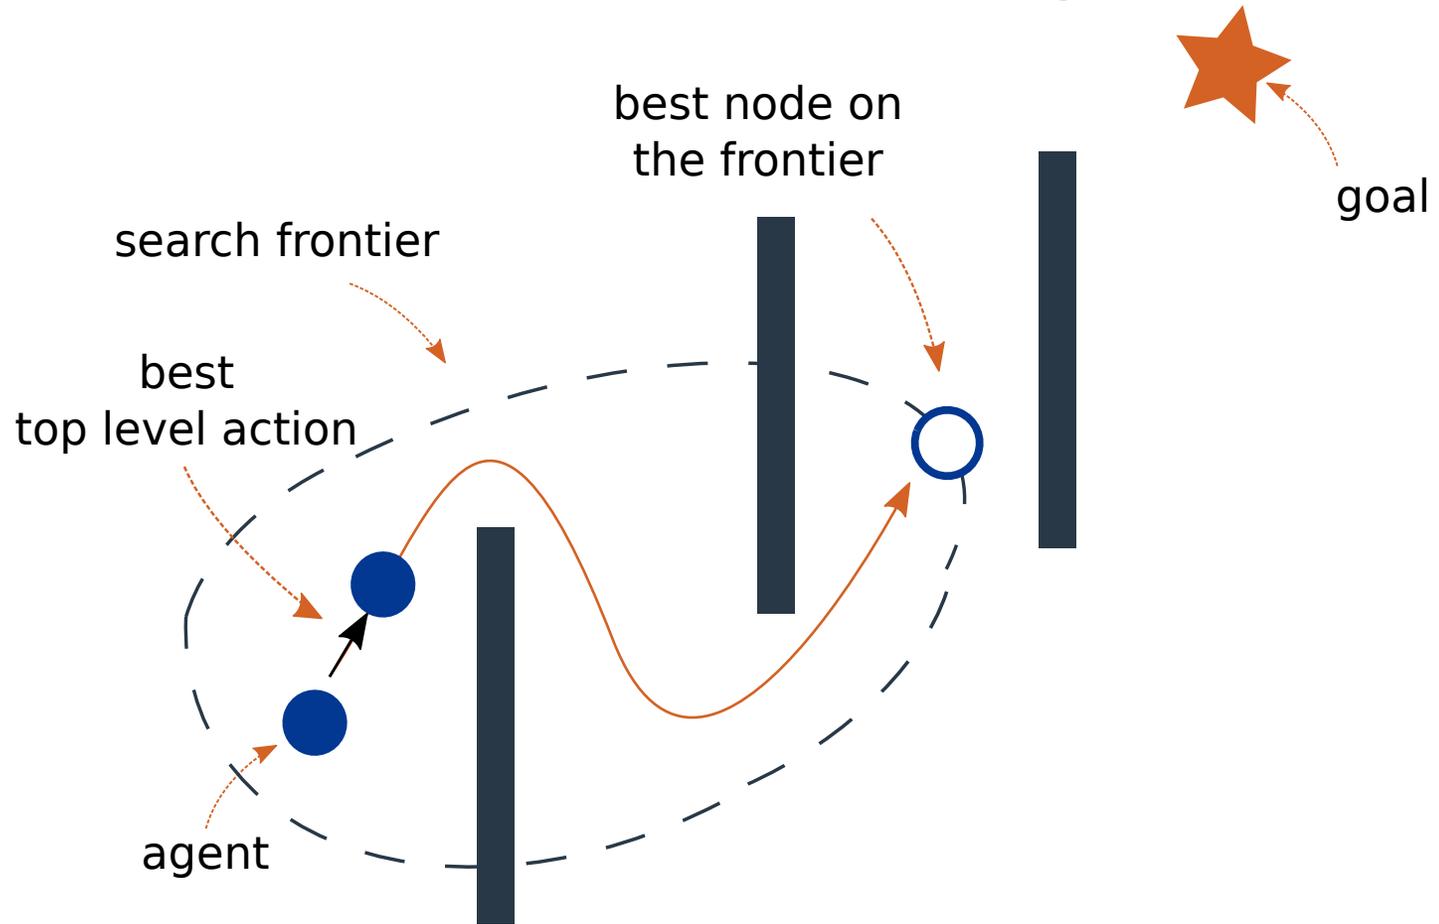
■ This work

Nancy

Results

Conclusions

An example: path finding



agent commits to best action and executes

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

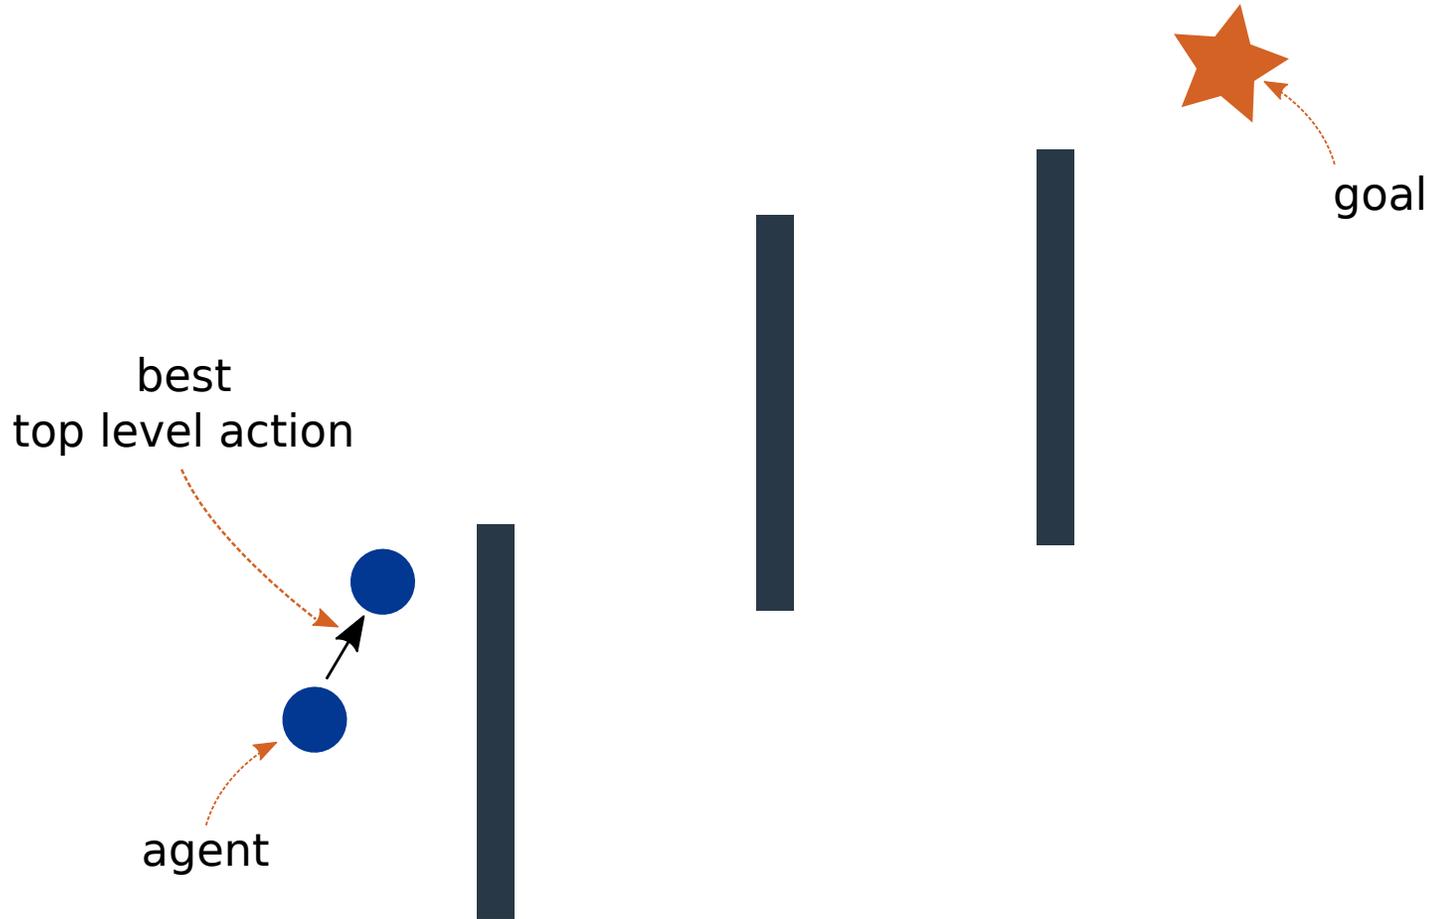
■ This work

Nancy

Results

Conclusions

An example: path finding



agent commits to best action and executes

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

■ This work

Nancy

Results

Conclusions

An example: path finding



agent commits to best action and executes

An Example of Real-time Heuristic Search

Introduction

■ The Authors

■ Real-time Search

■ An Example

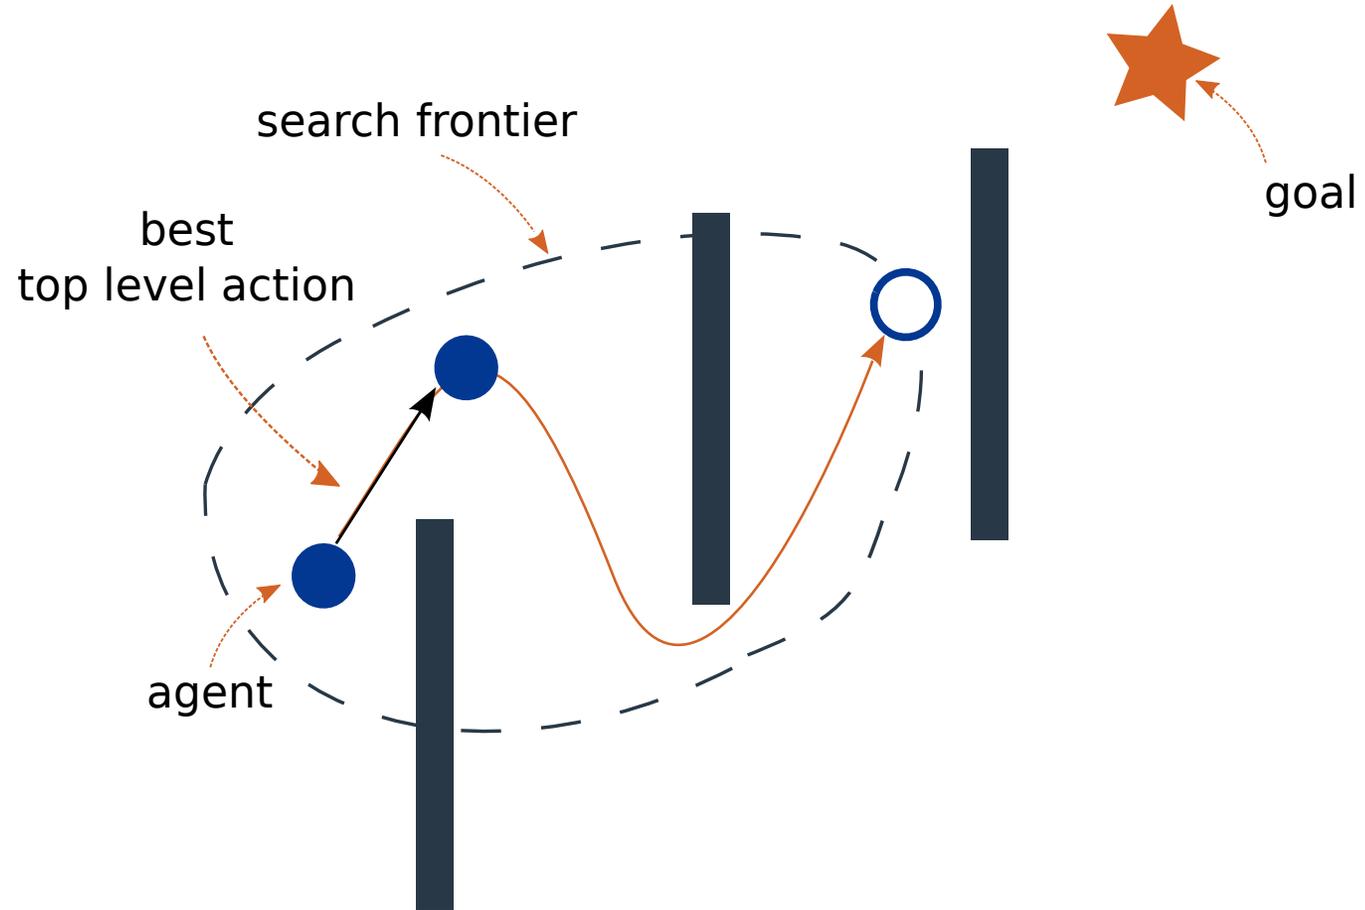
■ This work

Nancy

Results

Conclusions

An example: path finding



online planning: interleaving search and action execution
“receding horizon control”

Our Focus: Rational Real-time Planning

Introduction

- The Authors
- Real-time Search
- An Example
- This work

Nancy

Results

Conclusions

Challenge:

limited time → limited lookahead → value uncertainty

This work:

A beginning of rational approach for real-time planning

Borrowing idea from Bayesian RL

Simplified setting

Put focus on bounded rationality

Introduction

Nancy

■ Lookahead

■ Our Approach

Results

Conclusions

The Nancy Framework

Lookahead Strategy: A Troublesome Example

Introduction

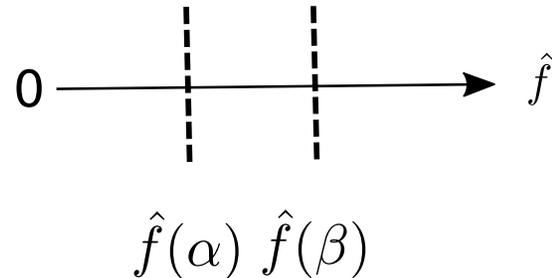
Nancy

Lookahead

Our Approach

Results

Conclusions



\hat{f} is expected value

Should an agent expand nodes under α or β ?

Lookahead Strategy: A Troublesome Example

Introduction

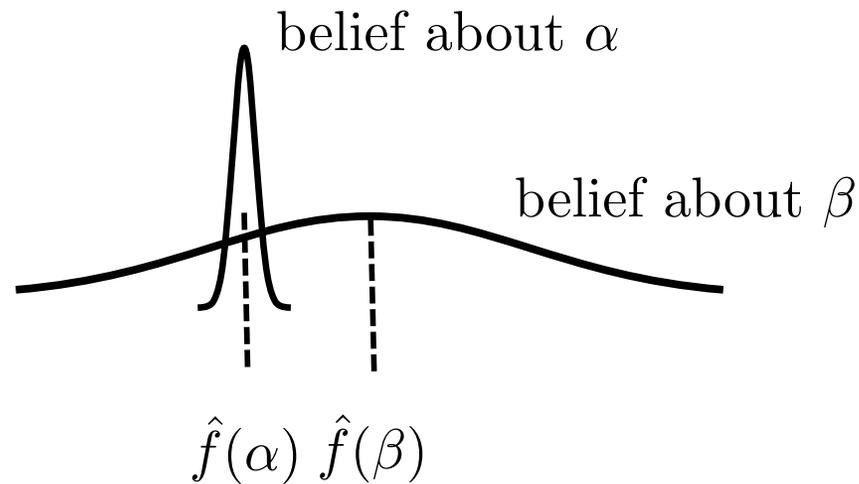
Nancy

Lookahead

Our Approach

Results

Conclusions



\hat{f} is expected value

Should an agent expand nodes under α or β ?

\hat{f} is not the answer: what to do?
want to maximize value of information
need to consider uncertainty of estimates

Lookahead Strategy: A Troublesome Example

Introduction

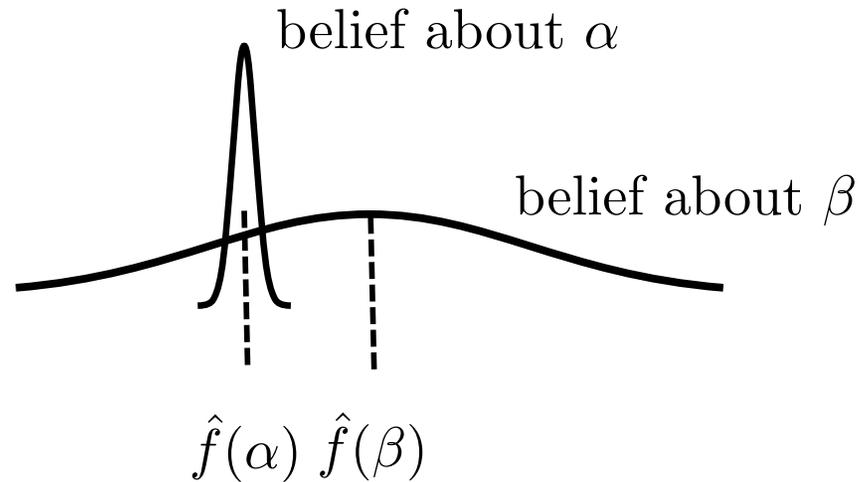
Nancy

Lookahead

Our Approach

Results

Conclusions



\hat{f} is expected value

Should an agent expand nodes under α or β ?

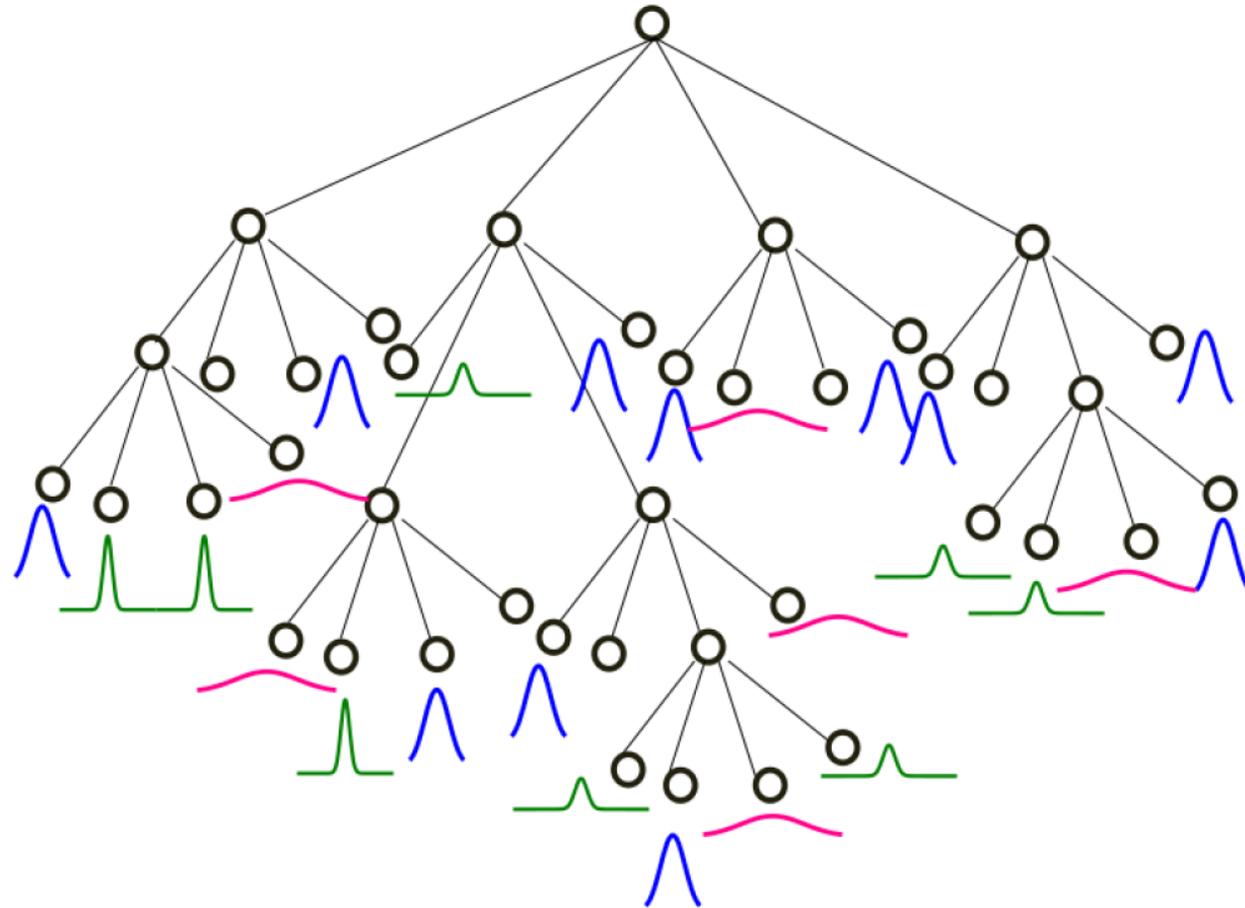
Our approach: expand nodes which **minimize expected regret**

$$\mathbb{E} \left[\underbrace{f^*(\alpha) - f^*(\beta)}_{\text{our regret}} \mid \underbrace{f^*(\beta) < f^*(\alpha)}_{\text{in cases when } \alpha \text{ not best}} \right]$$

Our Approach: The Nancy Framework

How to best use limited number of expansions?

- Introduction
- Nancy
 - Lookahead
 - Our Approach**
- Results
- Conclusions



Our Approach: The Nancy Framework

Introduction

Nancy

■ Lookahead

■ Our Approach

Results

Conclusions

How to best use limited number of expansions?

estimate value uncertainty

one option: uncertainty correlate distance-to-go
see details in Mitchell et al., AAI 2019

use the estimate to guide search

see details in Mitchell et al., AAI 2019

can use data to estimate distributions

see details in Fickert et al., AAI 2020

This PRL-20 paper:

comparison against other methods that try to exploit
value uncertainty
visualization of distributions estimated from data

Introduction

Nancy

Results

■ Data Viz

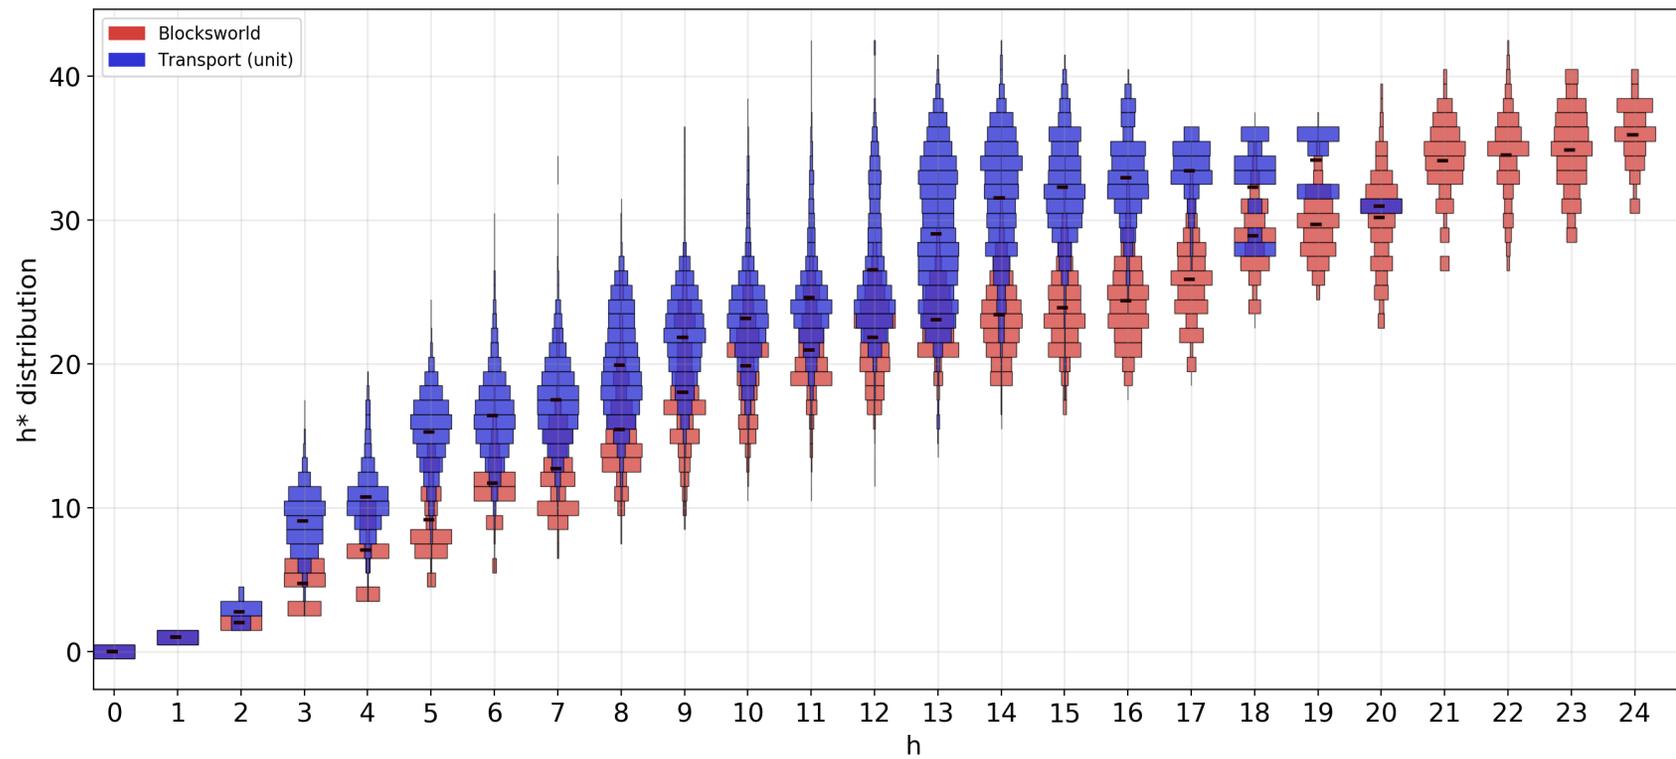
■ Results

Conclusions

Results

Example h^* distribution: Transport vs Blocks World

What does the actual cost-to-go value uncertainty distribution look like?



Beliefs are different from domain to domain

Comparison to IE and MCTS on Classic Search Domains

Introduction

Nancy

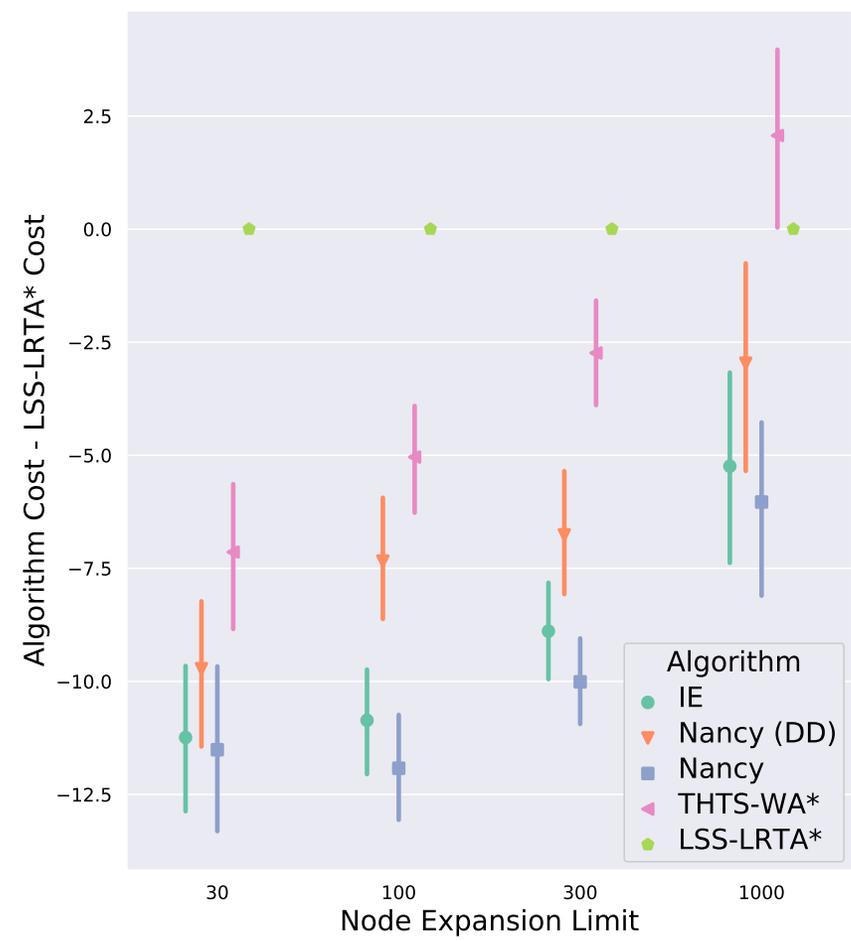
Results

■ Data Viz

■ Results

Conclusions

40 Pancake



Nancy outperforms conventional approaches

Introduction

Nancy

Results

Conclusions

■ Summary

Conclusions

Summary

Introduction

Nancy

Results

Conclusions

■ Summary

- Nancy starts to explore an optimal way of doing online heuristic search
- Nancy is competitive or better than conventional approaches

More broadly:

- Setting isolates the issue: unlike in MDPs or RL, all uncertainty is due to **bounded rationality**
- **Metareasoning** about uncertainty pays off, even for deterministic domains!

Questions?

Introduction

Nancy

Results

Conclusions

Questions

■ Questions?

