Metareasoning for Heuristic Search Using Uncertainty

Tianyi Gu

University of New Hampshire

Committee members:
Momotaz Begum
Laura Dietz
Levi Lelis
Marek Petrik
Wheeler Ruml (Advisor)

The Thesis of My Dissertation

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search can benefit from representing uncertainty

scalar heuristic $\rightarrow$ belief distribution that represents uncertainty

The Thesis of My Dissertation

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search can benefit from representing uncertainty

scalar heuristic $\rightarrow$ belief distribution that represents uncertainty

show in three problem settings:

- real-time heuristic search
- concurrent planning and execution
- bounded-cost search

What is Heuristic Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search: a planning approach

What is Heuristic Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search: a planning approach

planning is a model-based AI method, it models the environment as a state space and finds a sequence of actions that accomplishes some objective

What is Heuristic Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search: a planning approach

planning is a model-based AI method, it models the environment as a state space and finds a sequence of actions that accomplishes some objective

heuristic search:

{states, actions} $\rightarrow$ {V, E}

planning problem $\rightarrow$ find a path from s_{init} to $\{s_{goal}\}$

guide graph search by a heuristic estimate of cost-to-goal

What is Heuristic Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search associates costs with states, used to guide search

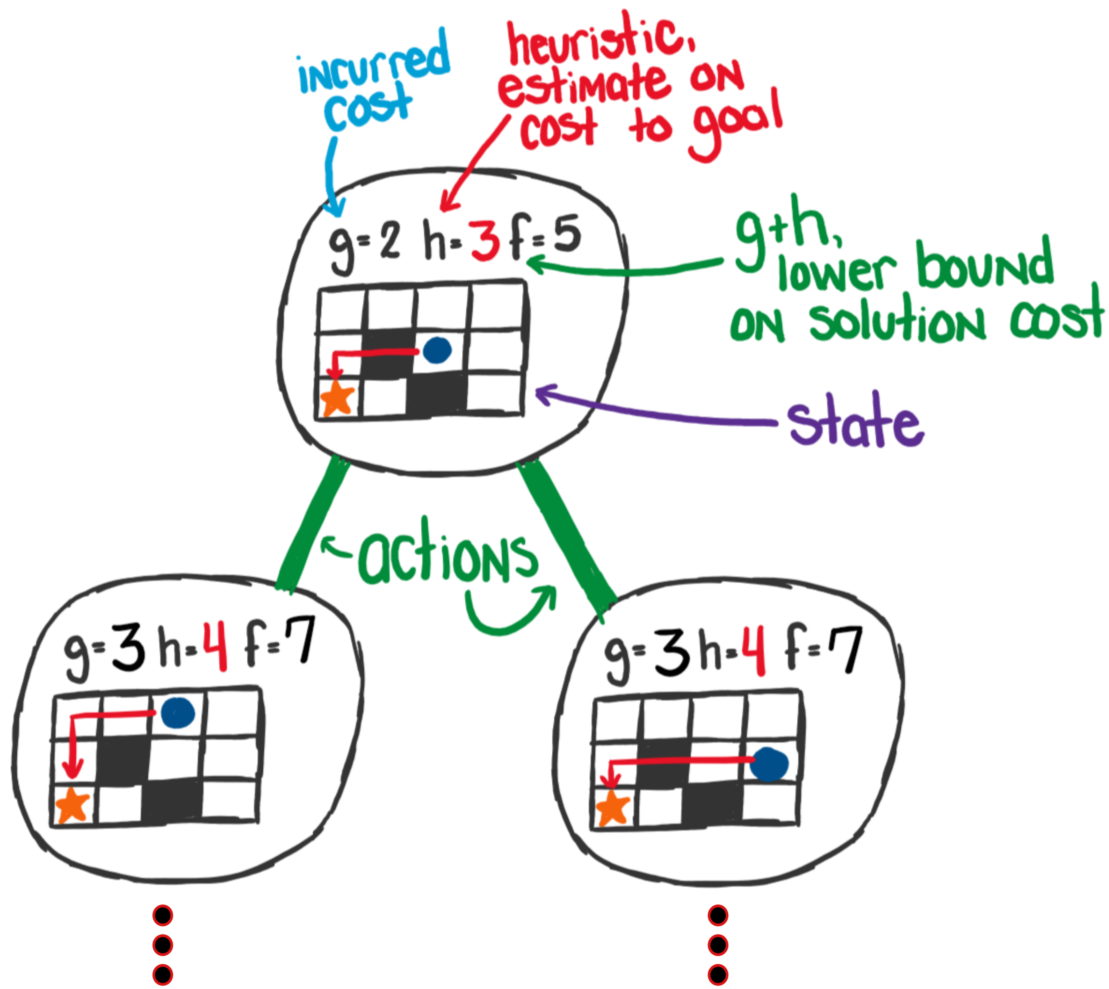

(Image credit: Andrew Mitchell)

Introduction
■ Thesis
□ Heuristic Search
■ Other Search
■ Metareasoning
■ Distributional Method
■ Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

A*: expands the node with minimal f value
returns optimal path
optimal search can take too long!
because it must expand every node with $f < C^*$[1]

[1] How Good is Almost Perfect, Malte Helmert and Gabriele Roger, AAAI, 2008.

Heuristic Search

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

A*: expands the node with minimal f value
returns optimal path
optimal search can take too long!
because it must expand every node with $f < C^*$[1]

What if we don't have time?

[1] How Good is Almost Perfect, Malte Helmert and Gabriele Roger, AAAI, 2008.

What If We Are Under Time Pressure?

What If We Are Under Time Pressure?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

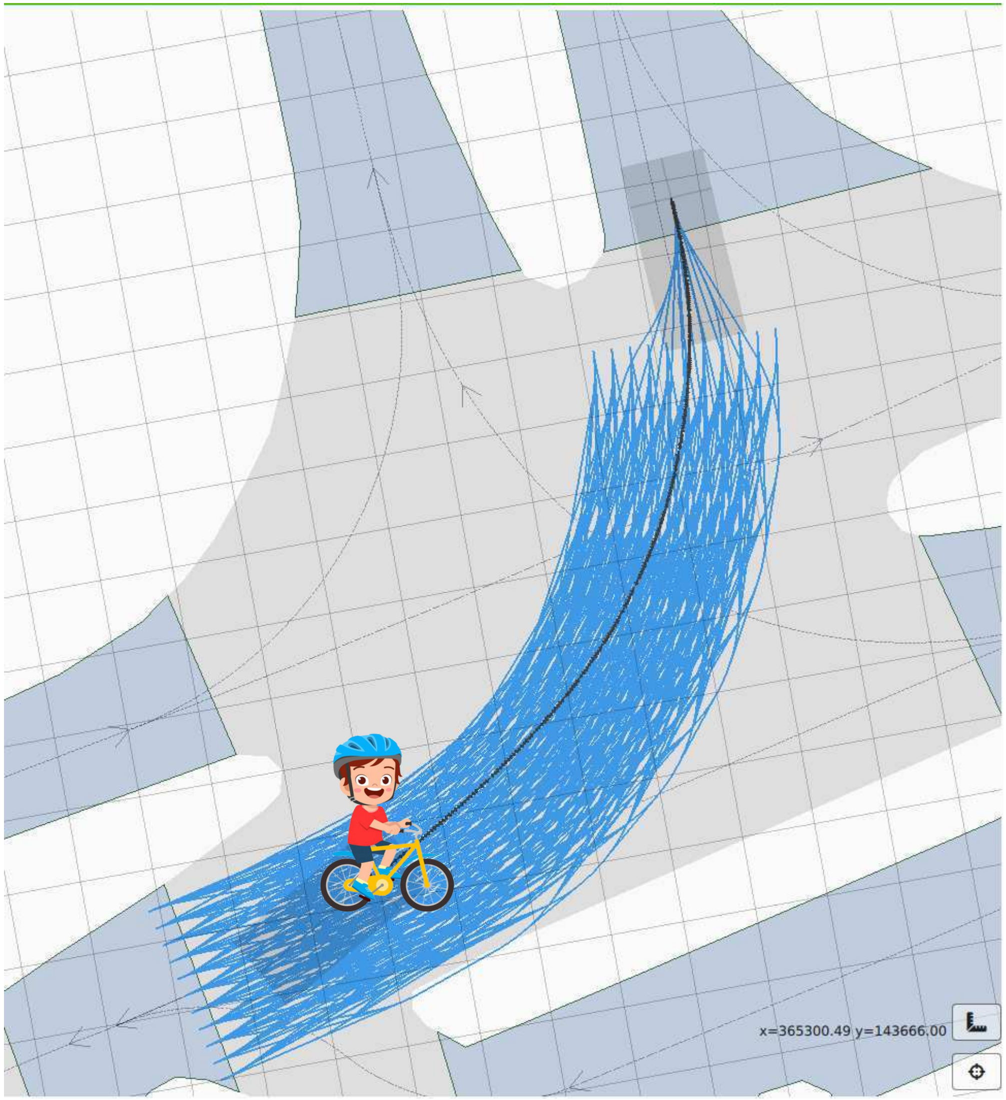

- large state space
- limited resource
- hard time bound

Alternatives to Optimal Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

1. real-time heuristic search
   time bound
2. concurrent planning and execution
   system must be stay under control
3. bounded-cost search
   cost bound

Alternatives to Optimal Search?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

1. real-time heuristic search
   time bound
2. concurrent planning and execution
   system must be stay under control
3. bounded-cost search
   cost bound

metareasoning:
reasoning about which reasoning to do

Metareasoning: Reasoning About Which Reasoning To Do

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

planning

Metareasoning: Reasoning About Which Reasoning To Do

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

metareasoning

Metareasoning: Reasoning About Which Reasoning To Do

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

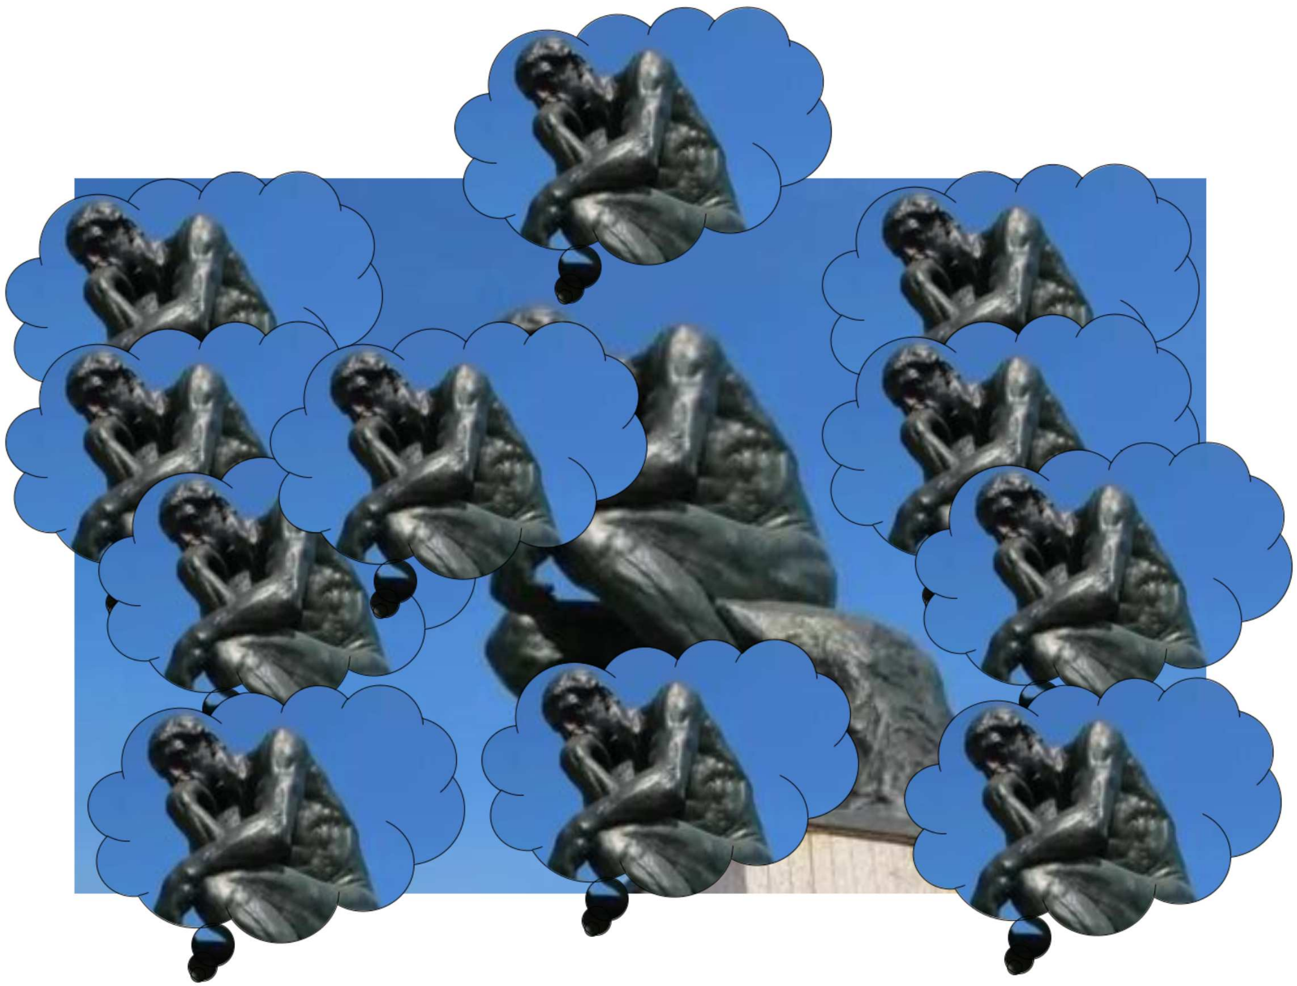

Make sure we are not dying from overthink and never act!

Why Use Distributional Methods?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

bounded rationality $\rightarrow$ uncertainty

Intuitively, distribution can be better than scalar-value based methods because it quantifies uncertainty, which is what search resolves.

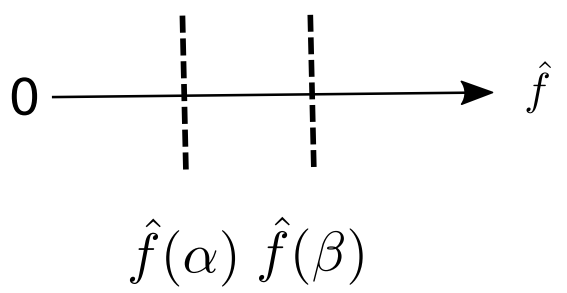

$\hat{f}$ is expected value

Should an agent expand nodes under α or β?

Why Use Distributional Methods?

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

bounded rationality $\rightarrow$ uncertainty

Intuitively, distribution can be better than scalar-value based methods because it quantifies uncertainty, which is what search resolves.

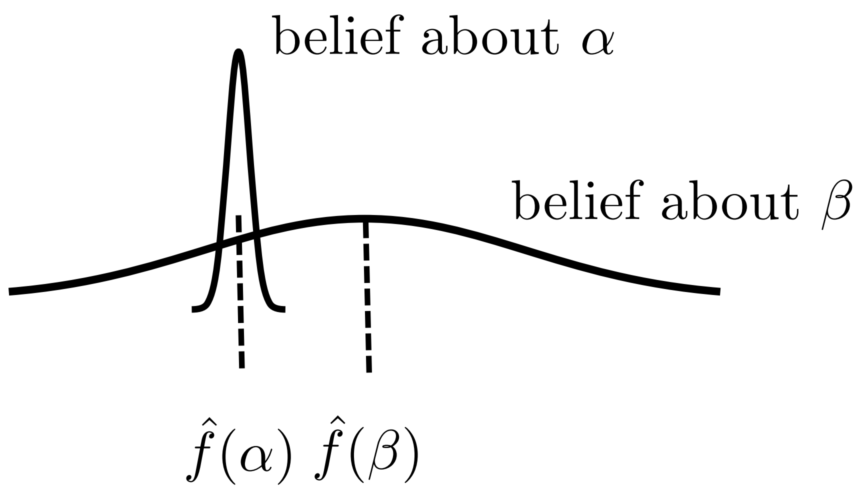

$\hat{f}$ is expected value

Should an agent expand nodes under α or β?

Overview

Introduction
- Thesis
- Heuristic Search
- Other Search
- Metareasoning
- Distributional Method
- Overview

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

heuristic search can benefit from representing uncertainty

- improving real-time search by representing uncertainty
  DDNancy: AAAI-20 *
- improving concurrent planing and execution by representing uncertainty
  FACS: IntEx-21
- improving bounded-cost search by representing uncertainty
  XES: IJCAI-21

Data-driven Real-time Search as Decision-making Under Uncertainty: Data-driven Nancy

Joint work with Maximilian Fickert, Leonhard Staut, Sai Lekyang, Wheeler Ruml, Joerg Hoffmann, and Marek Petrik

What is Real-time Heuristic Search?

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

An example: path finding

agent performs search for a bounded time

What is Real-time Heuristic Search?

Introduction

Real-time Search

- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

An example: path finding

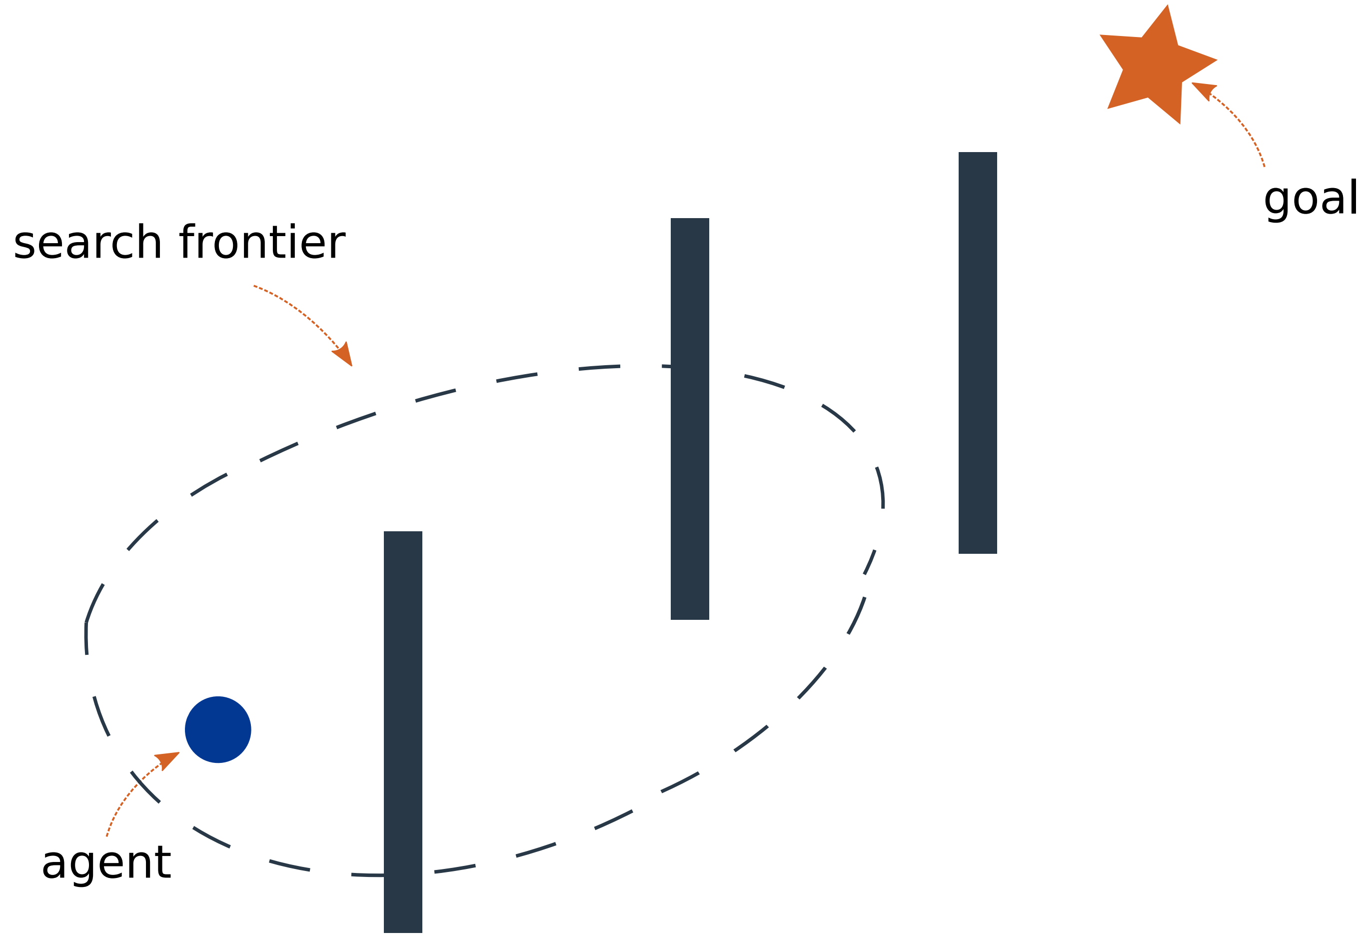

agent performs search for a bounded time

What is Real-time Heuristic Search?

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

An example: path finding

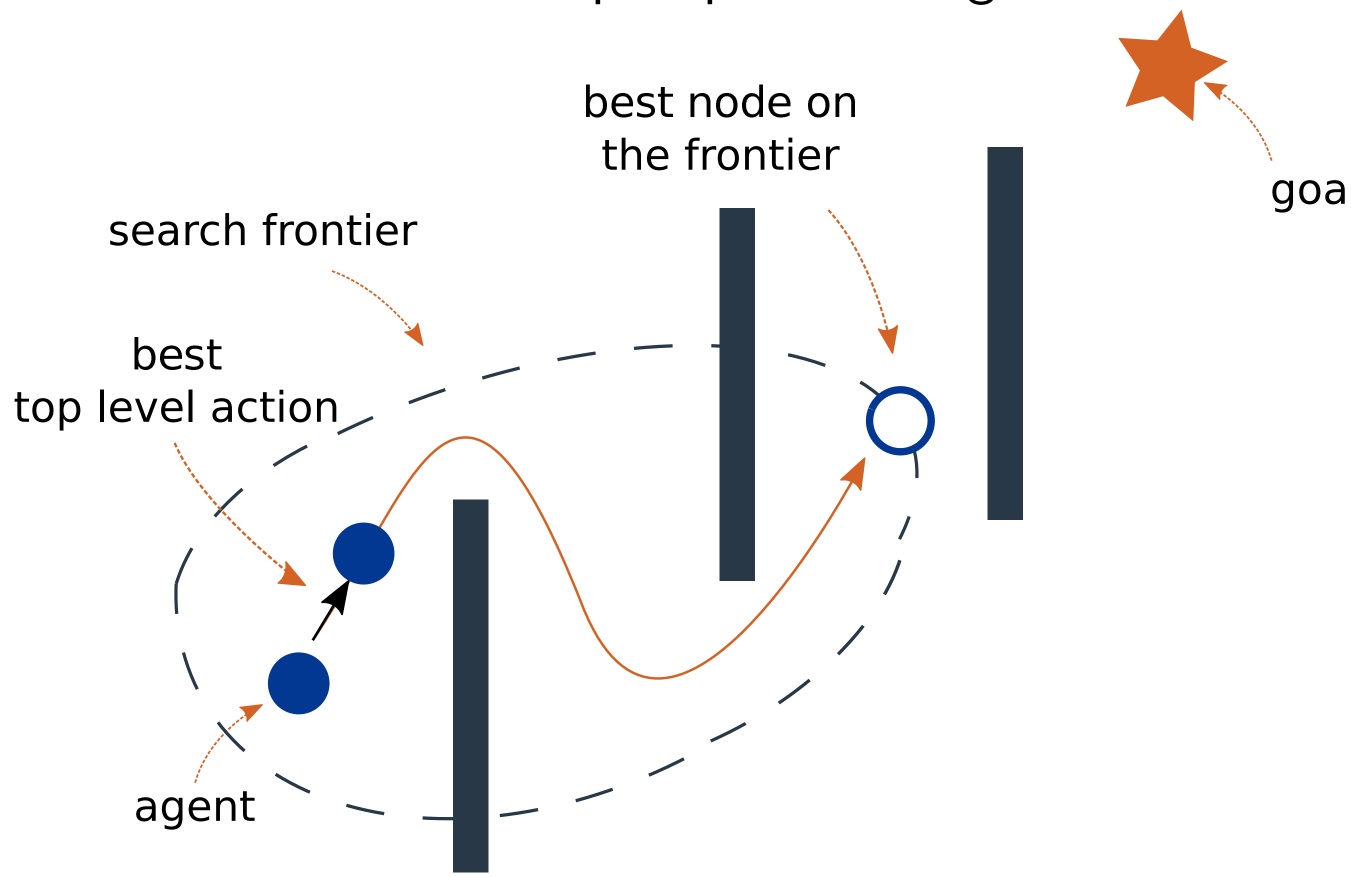

agent commits to best action and executes

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

An example: path finding

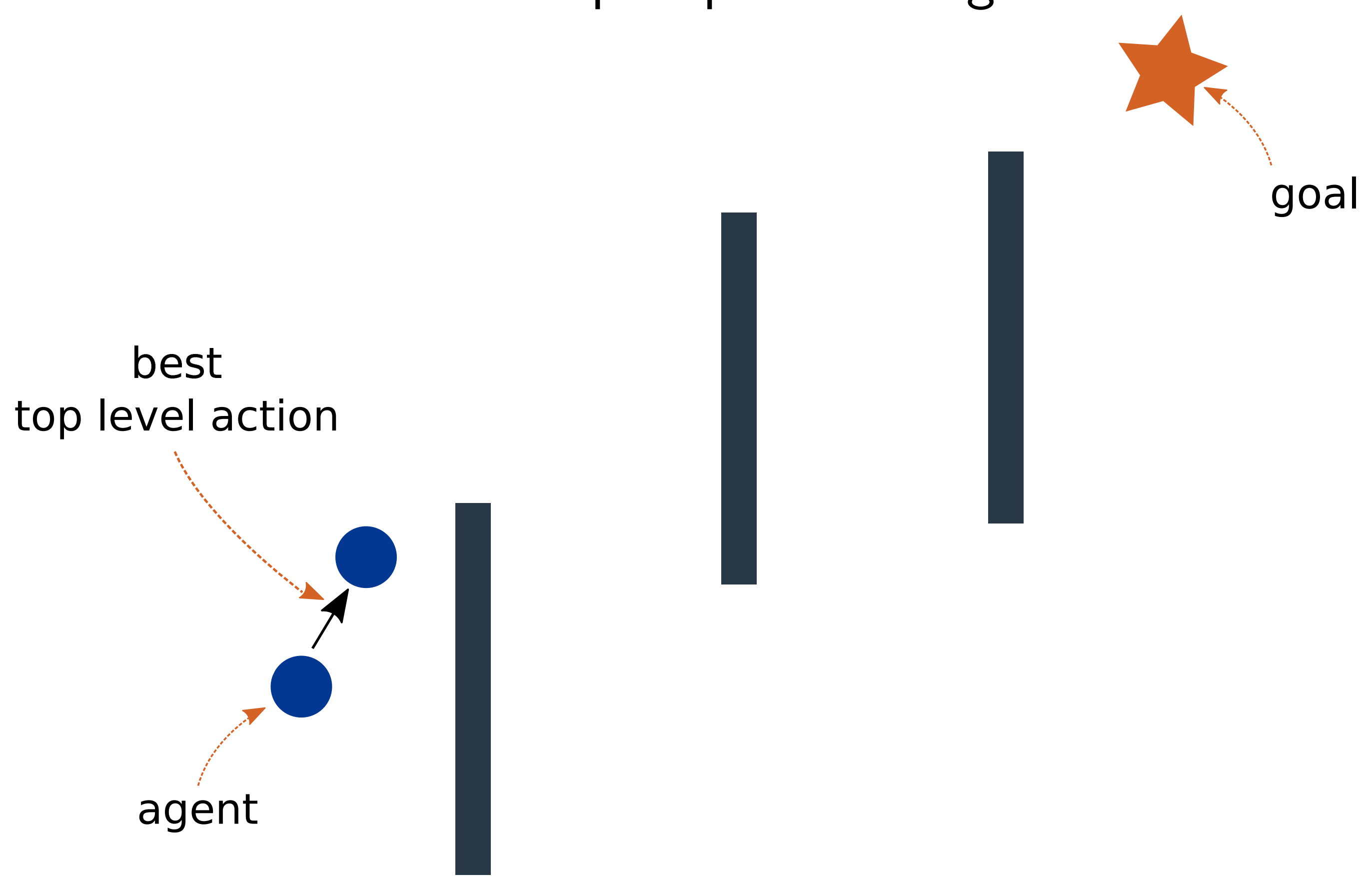

agent commits to best action and executes

What is Real-time Heuristic Search?

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

An example: path finding

agent commits to best action and executes

What is Real-time Heuristic Search?

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

An example: path finding

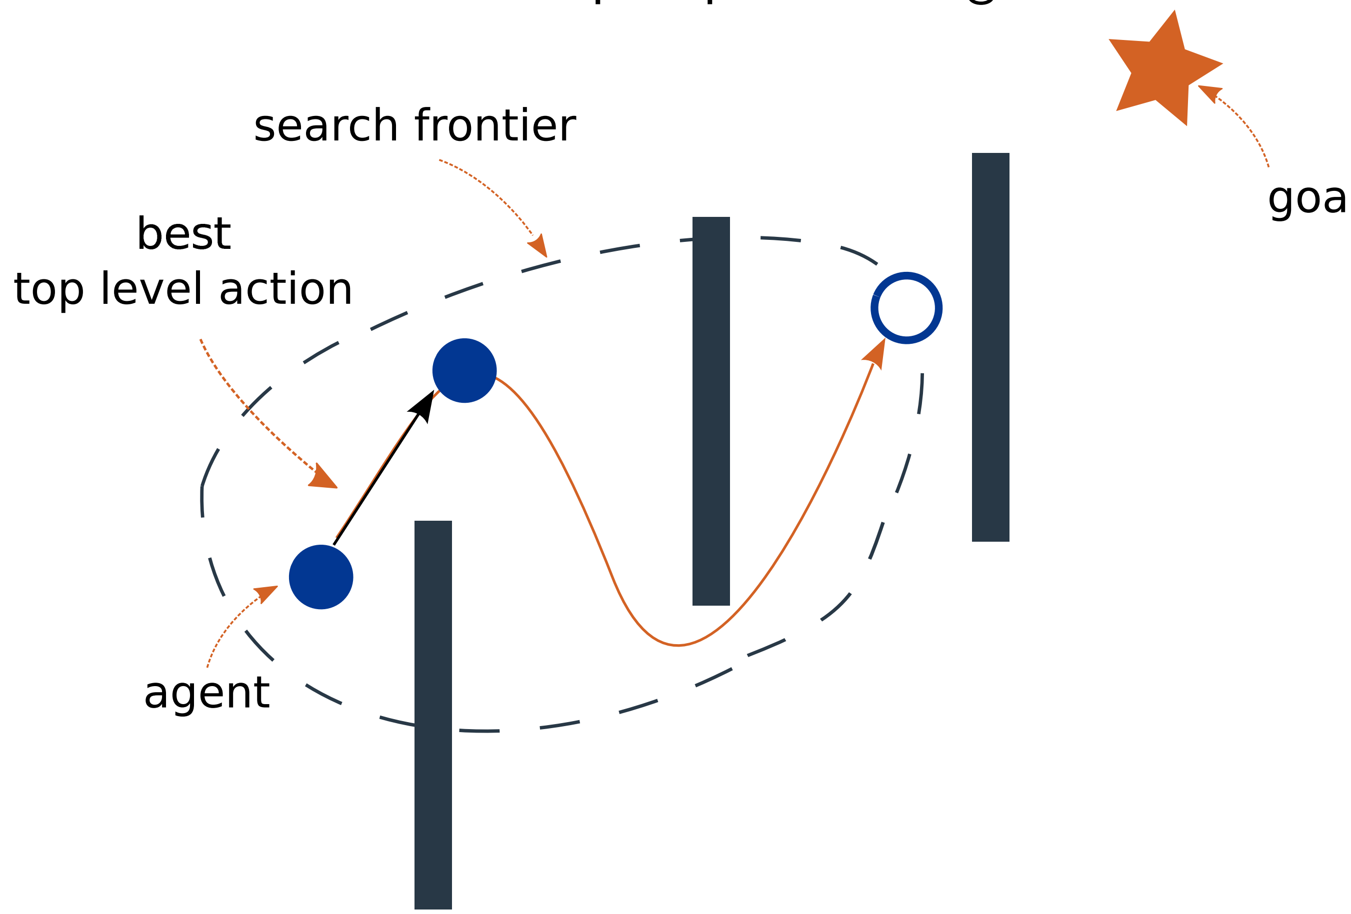

online planning: interleaving search and action execution

"receding horizon control"

A Classic Approach: LSS-LRTA* (Koenig&Sun 2008)

Introduction

Real-time Search

- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

three phases:

1. Lookahead Phase:

   expands nodes with minimum f
   to explore the search space

A Classic Approach: LSS-LRTA* (Koenig&Sun 2008)

Introduction

Real-time Search

- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

three phases:

1. Lookahead Phase:
   expands nodes with minimum f
   to explore the search space
2. Decision-making Phase:
   backup the minimum f from search frontier ('minimin')
   select top level action with minimum f to execute

A Classic Approach: LSS-LRTA* (Koenig&Sun 2008)

Introduction
Real-time Search
■ Real-time Search
■ Lookahead
■ The Beliefs
■ Data
■ Data Viz
■ Result
■ Summary
Online Planning
Bounded-Cost Search
Conclusion

three phases:

1. Lookahead Phase:
   expands nodes with minimum f
   to explore the search space
2. Decision-making Phase:
   backup the minimum f from search frontier ('minimin')
   select top level action with minimum f to execute
3. Learning Phase:
   update heuristic values
   (to escape local minima and avoid infinite loops)

A Classic Approach: LSS-LRTA* (Koenig&Sun 2008)

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

three phases:

1. Lookahead Phase:
   expands nodes with minimum f
   to explore the search space
2. Decision-making Phase:
   backup the minimum f from search frontier ('minimin')
   select top level action with minimum f to execute
3. Learning Phase:
   update heuristic values
   (to escape local minima and avoid infinite loops)

repeat until at a goal

proved to be complete for consistent heuristic

A Classic Approach: LSS-LRTA* (Koenig&Sun 2008)

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

three phases:

1. Lookahead Phase:
   expands nodes with minimum f
   to explore the search space
2. Decision-making Phase:
   backup the minimum f from search frontier ('minimin')
   select top level action with minimum f to execute
3. Learning Phase:
   update heuristic values
   (to escape local minima and avoid infinite loops)

repeat until at a goal

proved to be complete for consistent heuristic

derived from offline search, but optimal for online?

Lookahead Phase: A Troublesome Example

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

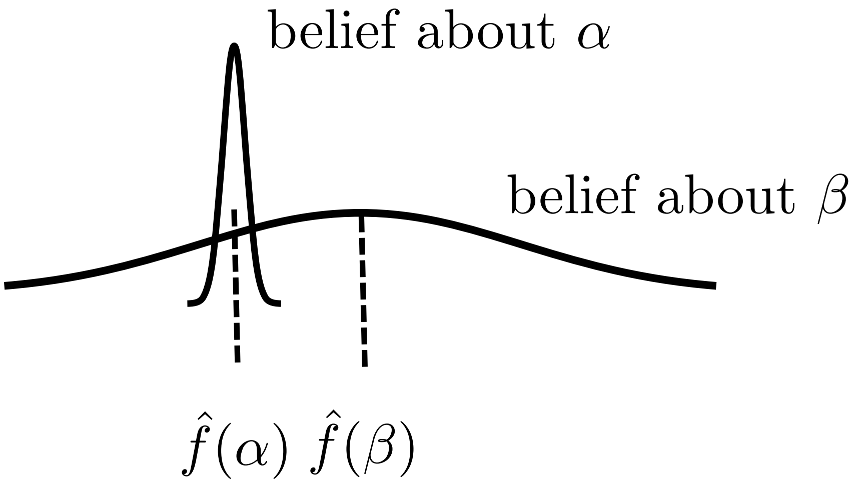

$\hat{f}$ is expected value

Should an agent expand nodes under α or β?

Lookahead Phase: A Troublesome Example

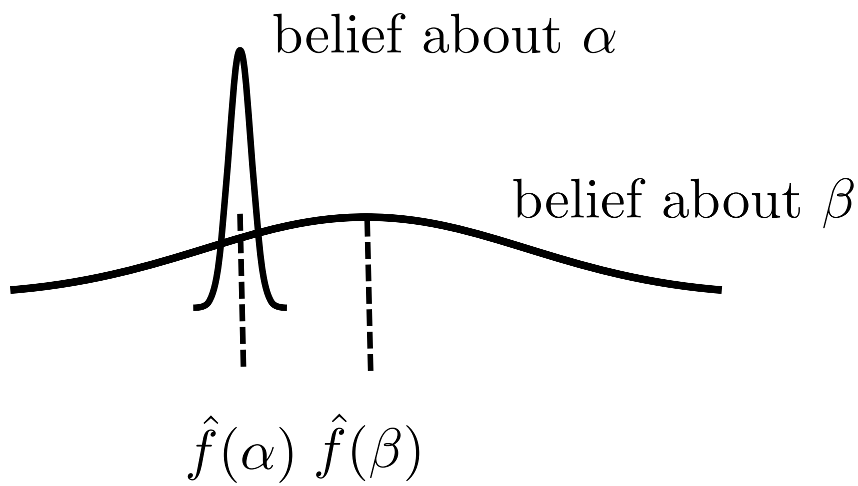

$\hat{f}$ is expected value

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
Risk-based evaluation: minimize expected regret

Risk-based Lookahead Example

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

expand under α or β?

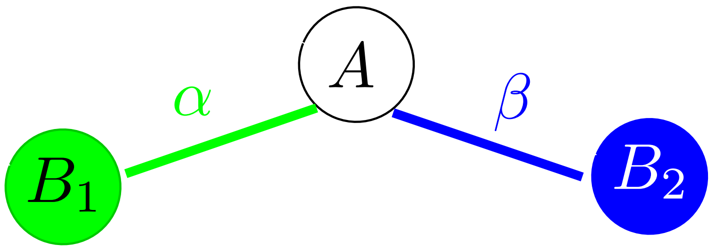

Risk-based Lookahead Example

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

expand under α or β?

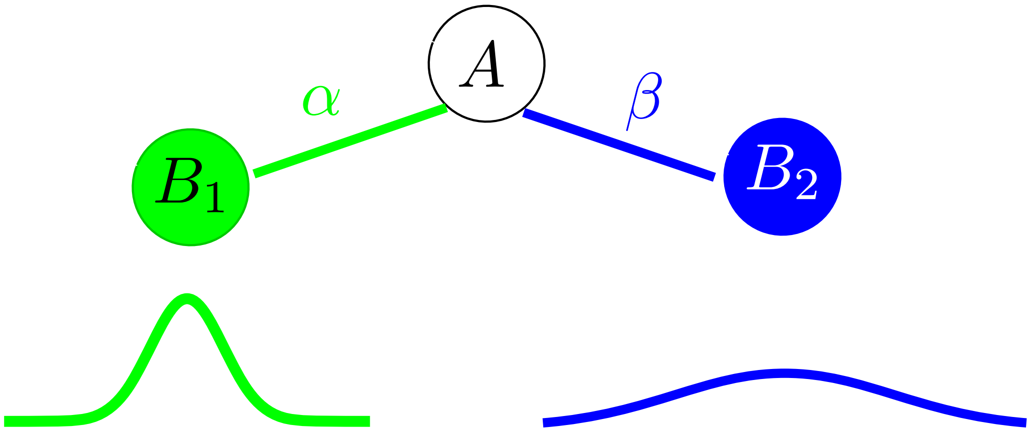

need 2 things:

1) current beliefs

2) estimate of how beliefs might change with search

Risk-based Lookahead Example

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

expand under α or β?

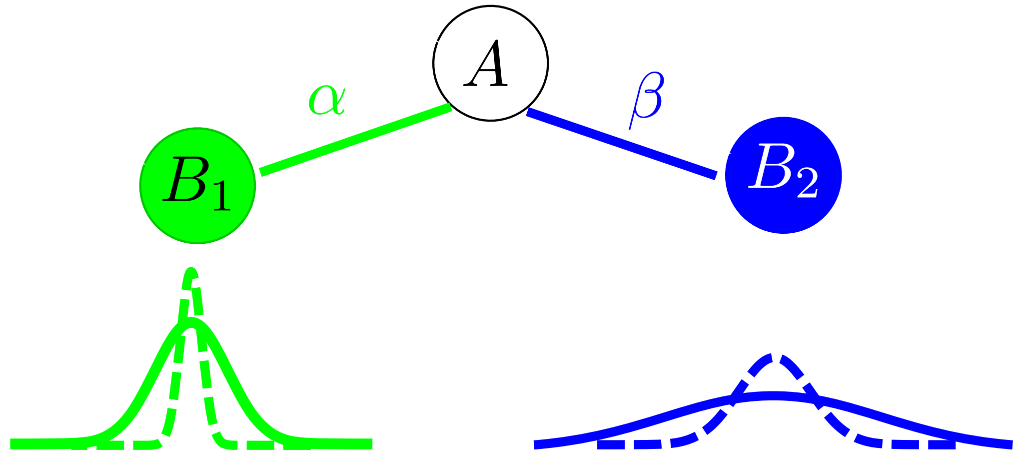

need 2 things:

1) current beliefs
2) estimate of how beliefs might change with search

Risk-based Lookahead Example

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

expand under α or β?

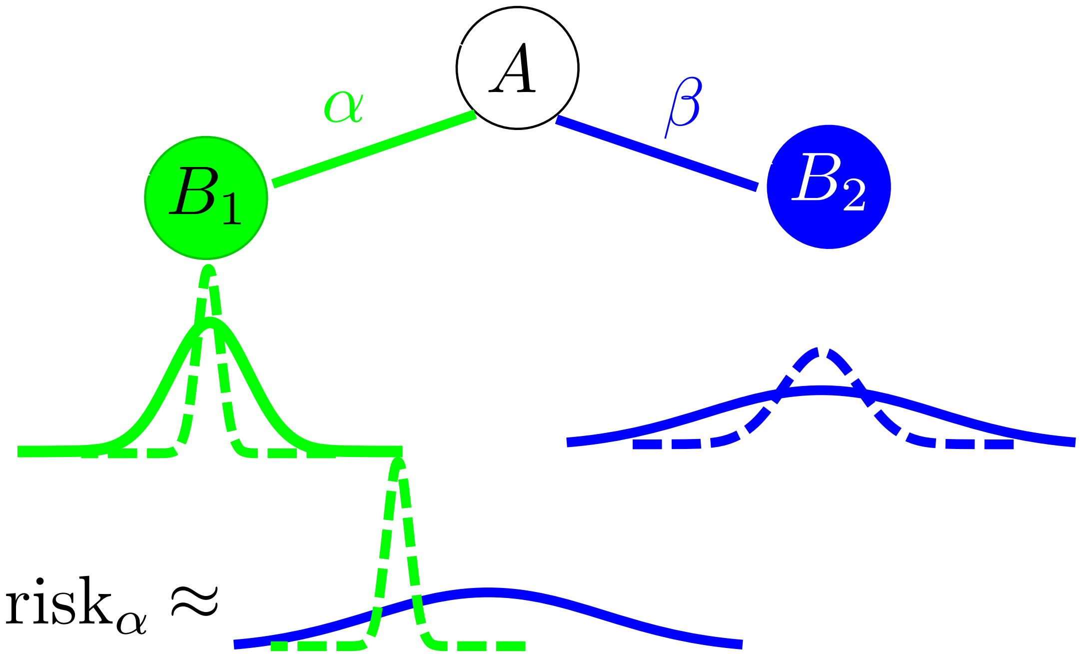

Risk: expected regret if a suboptimal action is selected

α is TLA with lowest expected value, other is β

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

Risk-based Lookahead Example

Introduction
Real-time Search
Real-time Search
Lookahead
The Beliefs
Data
Data Viz
Result
Summary
Online Planning
Bounded-Cost Search
Conclusion

expand under α or β?

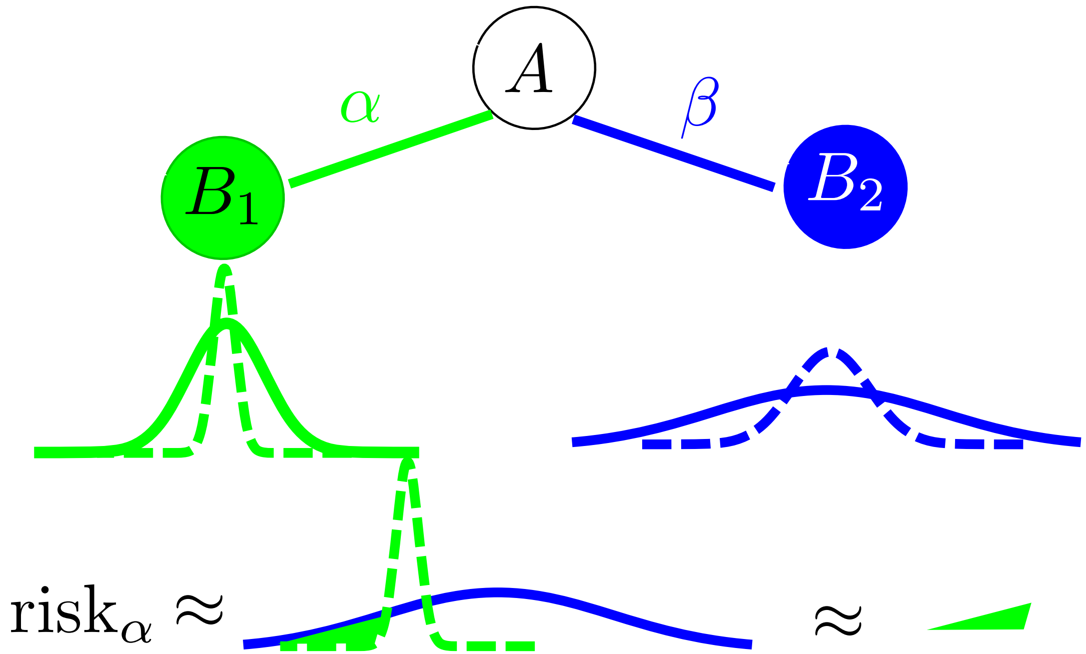

Risk: expected regret if a suboptimal action is selected

α is TLA with lowest expected value, other is β

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

Risk-based Lookahead Example

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

expand under α or β?

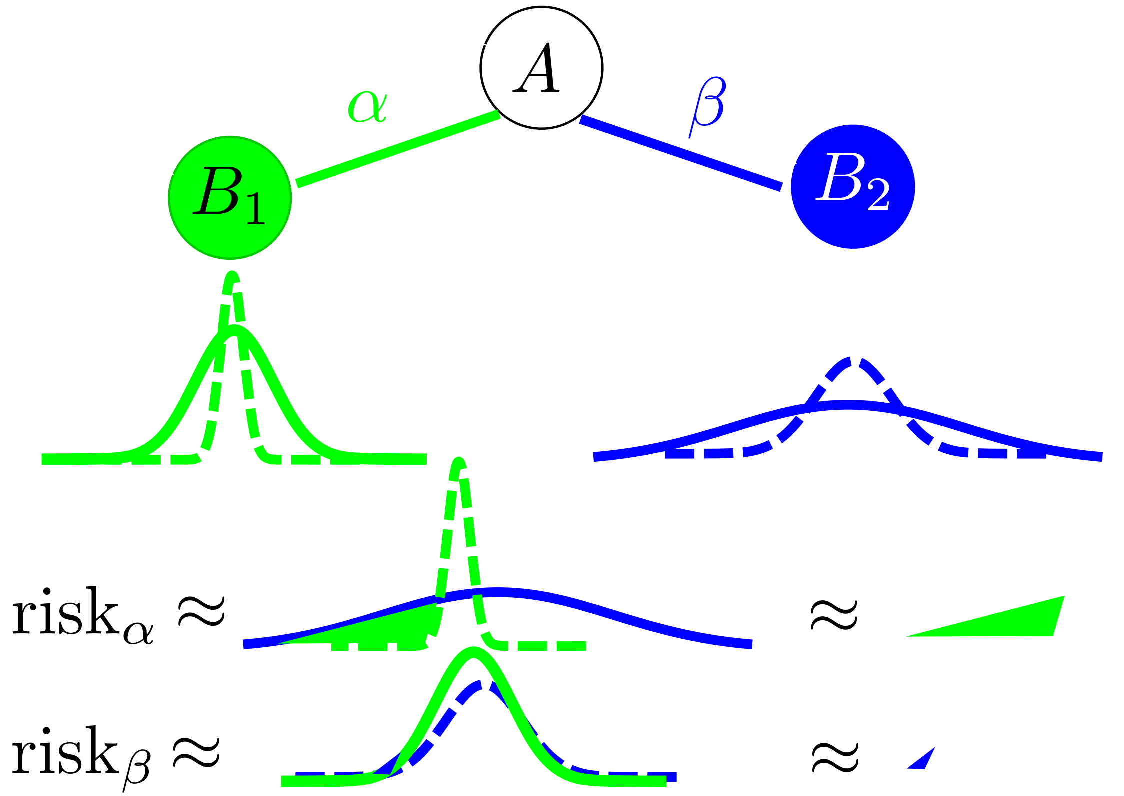

expand under the TLA that minimizes risk!

expand under β!

How to Form The Belief Distribution?

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary
Online Planning
Bounded-Cost Search
Conclusion

Heuristic values: scalar $\rightarrow$ probability distribution (belief)

But where do beliefs come from?

How to Form The Belief Distribution?

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

Heuristic values: scalar $\rightarrow$ probability distribution (belief)

But where do beliefs come from?

Nancy (Mitchell et al 2019):

truncated Gaussian based on $\hat{f}$ and f,
few parameters allows online learning

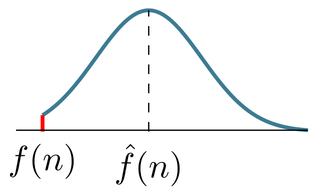

My work: Data-Driven Nancy:

expressive histogram,
many parameters requires offline learning

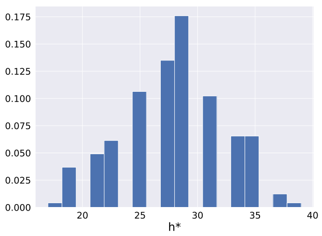

Learning a Model of Heuristic Error

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

belief: distribution of h* given features of state (h)

Gathering data:

run weighted-A* on random problems and collect all states

for each observed h value:

pick most common 200 states from the collection,

compute h*

Learning a Model of Heuristic Error

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

belief: distribution of h* given features of state (h)

Gathering data:

run weighted-A* on random problems and collect all states

for each observed h value:

pick most common 200 states from the collection,

compute h*

compute h*: need powerful optimal solver (eg. IDA*$_{CR}$[2] with pattern database heuristic)

[2] Reducing reexpansions in iterative-deepening search by controlling cutoff bounds, U.K. Sarkar, P.P. Chakrabarti, S. Ghose, S.C. De Sarkar, Artificial Intelligence, 1991.

Example h* distribution: Sliding Puzzle

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

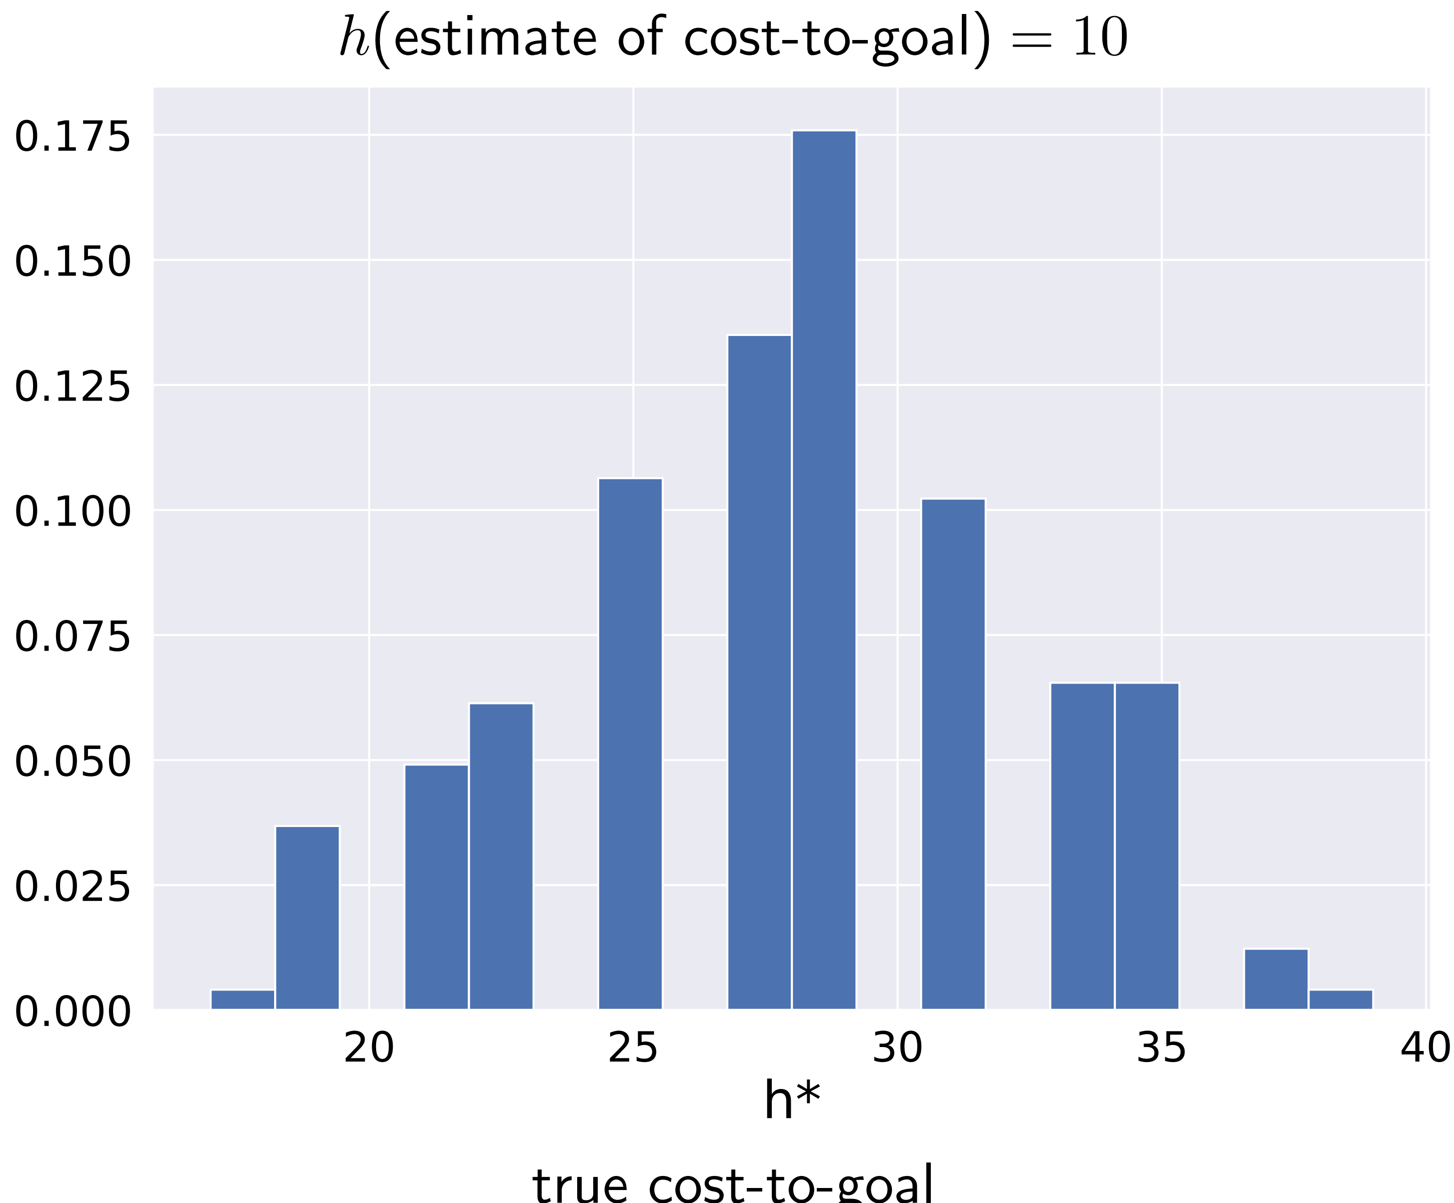

Example h^* distribution: Transport vs Blocks World

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

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

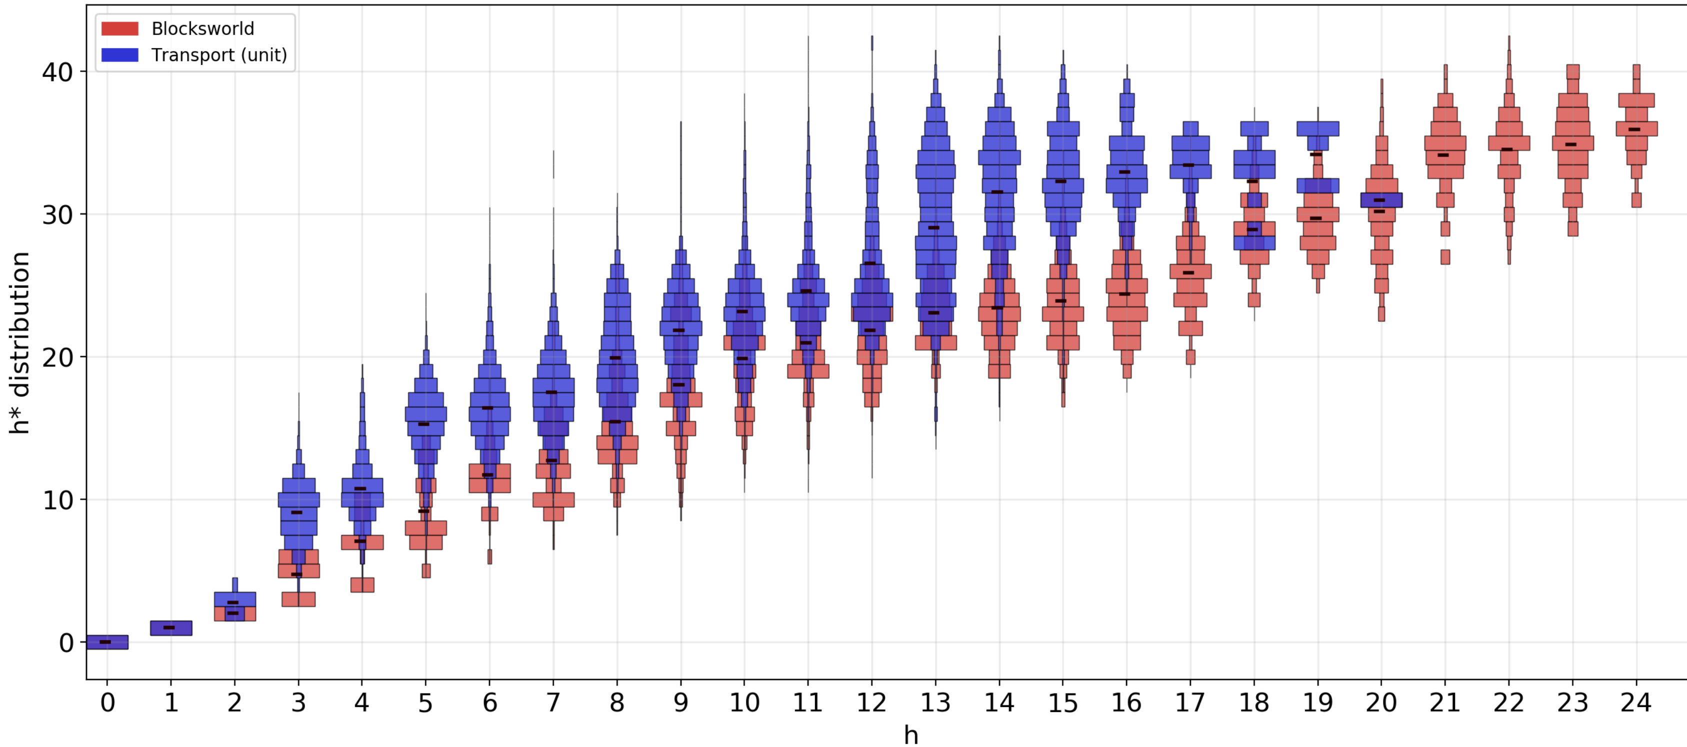

Beliefs are different from domain to domain

(Image credit: Leonhard Staut)

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

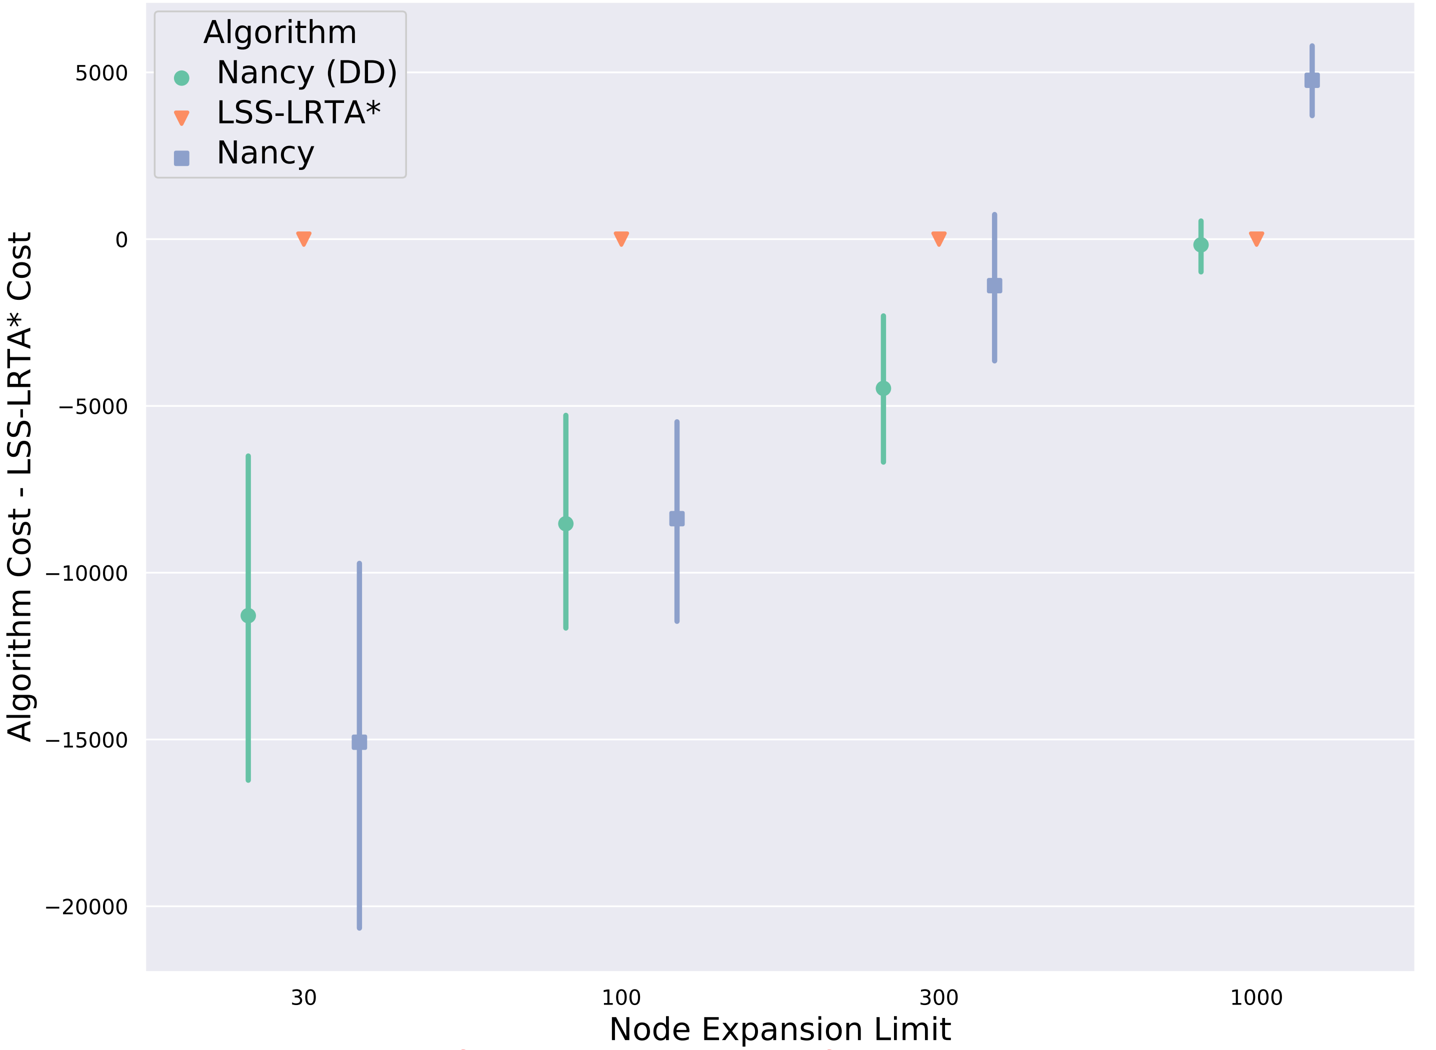

Both version of Nancy outperform conventional approach!

Results on Racetrack Problem

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

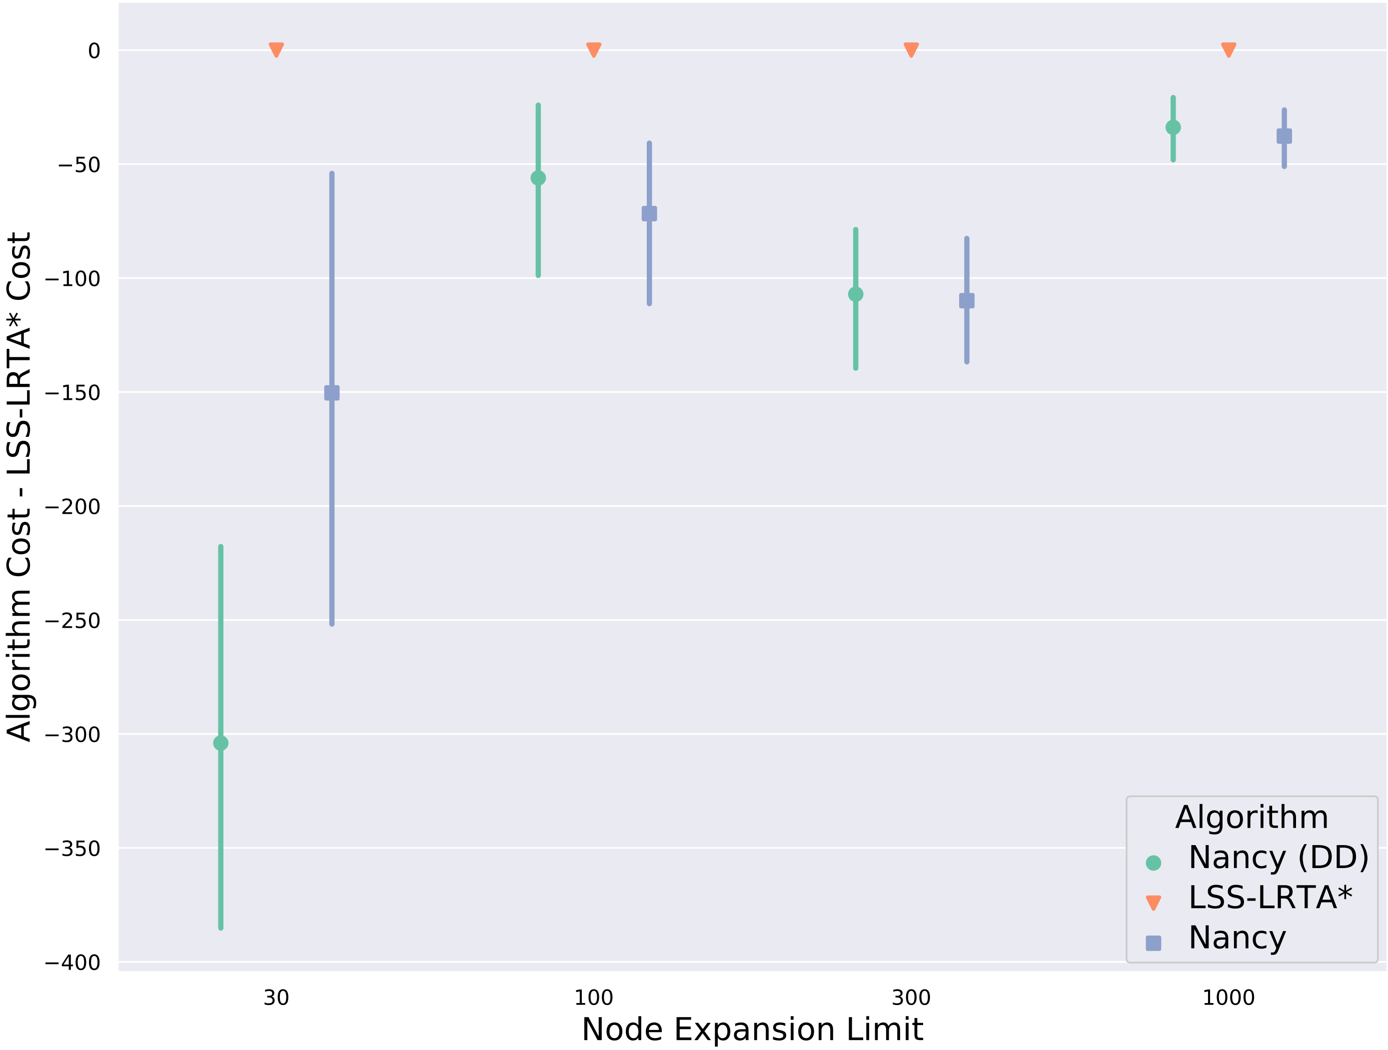

Data-Driven Nancy is more robust!

Mean Solution Cost on Planning Domains

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

Domain	Lookahead	LSS-LRTA*	Nancy	Nancy (DD)
Blocksw.	100	46	**33**	38
	300	36	**30**	34
	1000	30	32	**27**
Transport	100	631	615	**496**
	300	519	559	**485**
	1000	499	567	**422**
Transport (unit-cost)	100	48	40	**31**
	300	47	**30**	34
	1000	35	29	**27**
Elevators (unit-cost)	100	50	**35**	39
	300	32	**29**	30
	1000	34	27	**26**

Data-Driven Nancy is more robust!

(Table credit: Maximilian Fickert, Leonhard Staut)

CPU Time on Sliding Puzzle Problem

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

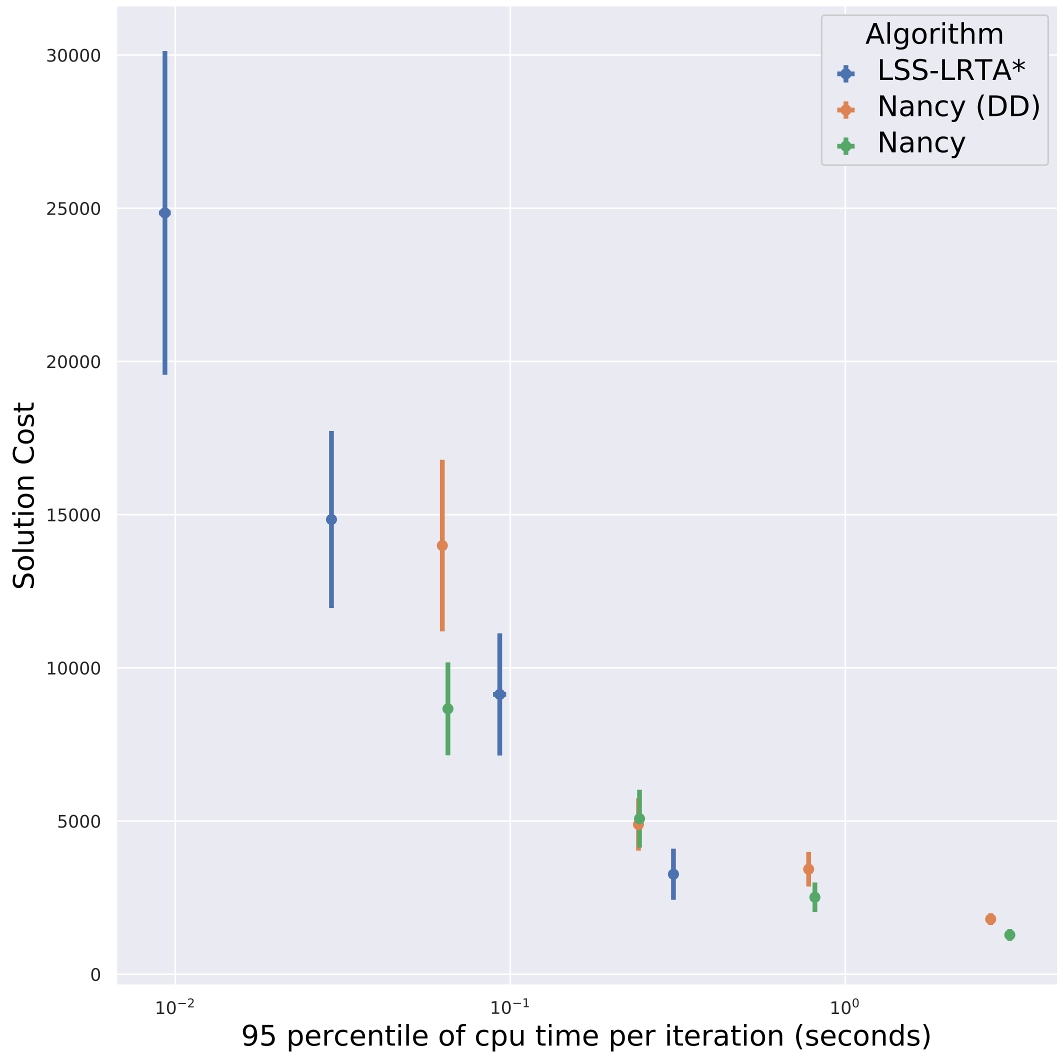

Nancy incurs overhead but worth it!

Comparison to IE and MCTS

Introduction

Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

Interval Estimation:

IE chooses the TLA with the lowest lower bound on the 95% confidence interval of the backed-up belief

Comparison to IE and MCTS

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

Interval Estimation:

IE chooses the TLA with the lowest lower bound on the 95% confidence interval of the backed-up belief

MCTS

use THTS-WA* in the expansion phrase

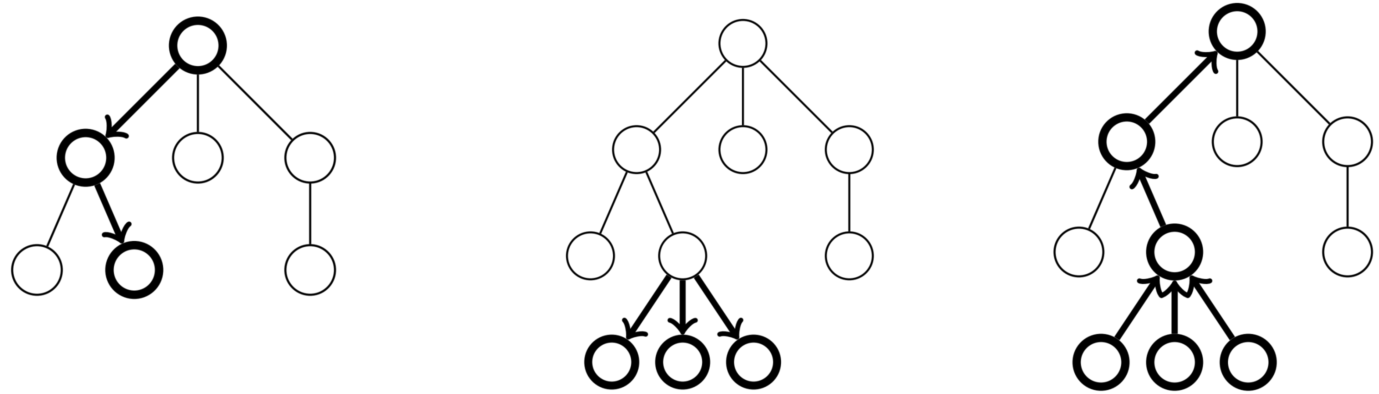

Action Selection | Initialization | Backup function

(Image credit: Tim Schulte and Thomas Keller)

Comparison to IE and MCTS on 40 Pancake

Introduction

Real-time Search

- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning

Bounded-Cost Search

Conclusion

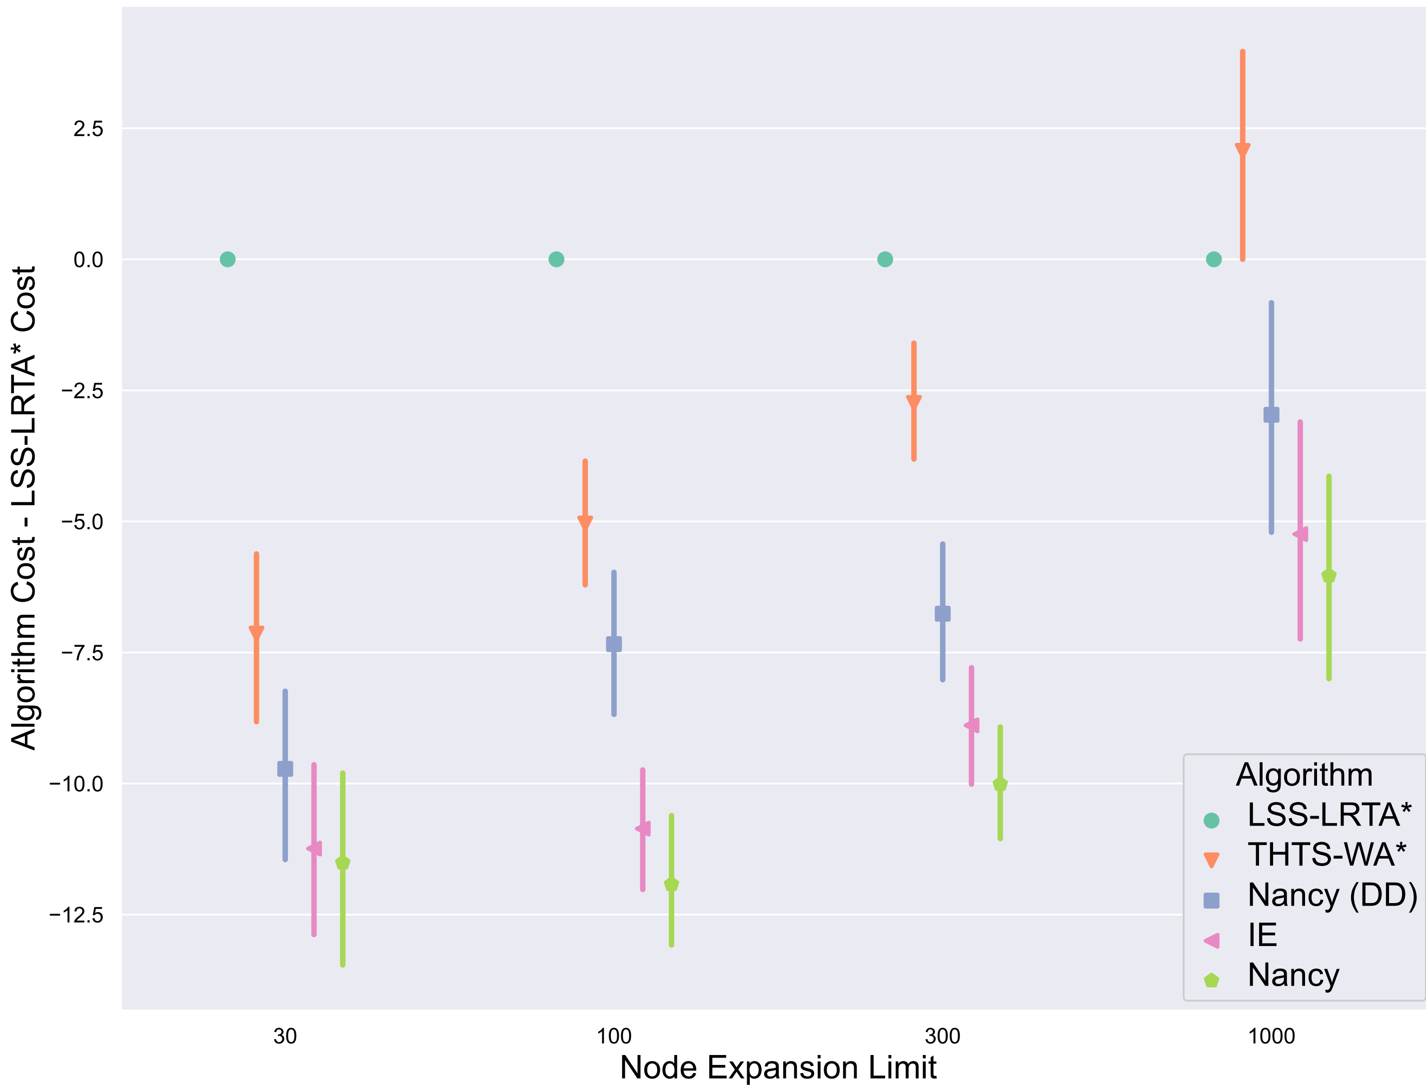

Reasoning about uncertainty helps!

Summary

Introduction
Real-time Search
- Real-time Search
- Lookahead
- The Beliefs
- Data
- Data Viz
- Result
- Summary

Online Planning
Bounded-Cost Search
Conclusion

- Distributional methods provide significant improvements compared to conventional LSS-LRTA* in real-time search
- Nancy starts to explore an optimal way of doing online heuristic search
- Data-driven approach provides an alternative way to implement the Nancy framework, it performs better when assumption fails

Status of the work:

1. *Beliefs We Can Believe In: Replacing Assumptions with Data in Real-Time Search.* AAAI Conference on Artificial Intelligence (AAAI), 2020.
2. *Real-time Planning as Data-driven Decision-making.* ICAPS Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL), 2020.
3. In preparation: *Real-time Planning as Decision-making Under Heuristic Value Uncertainty*, Journal of Artificial Intelligence Research (JAIR)

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

When to Commit to an Action in Online Planning?

Joint work with Shahaf Shperberg, Eyal Shlomo Shimony, Wheeler Ruml and Erez Karpas

Introduction
Real-time Search
Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary
Bounded-Cost Search
Conclusion

An example: highway navigation

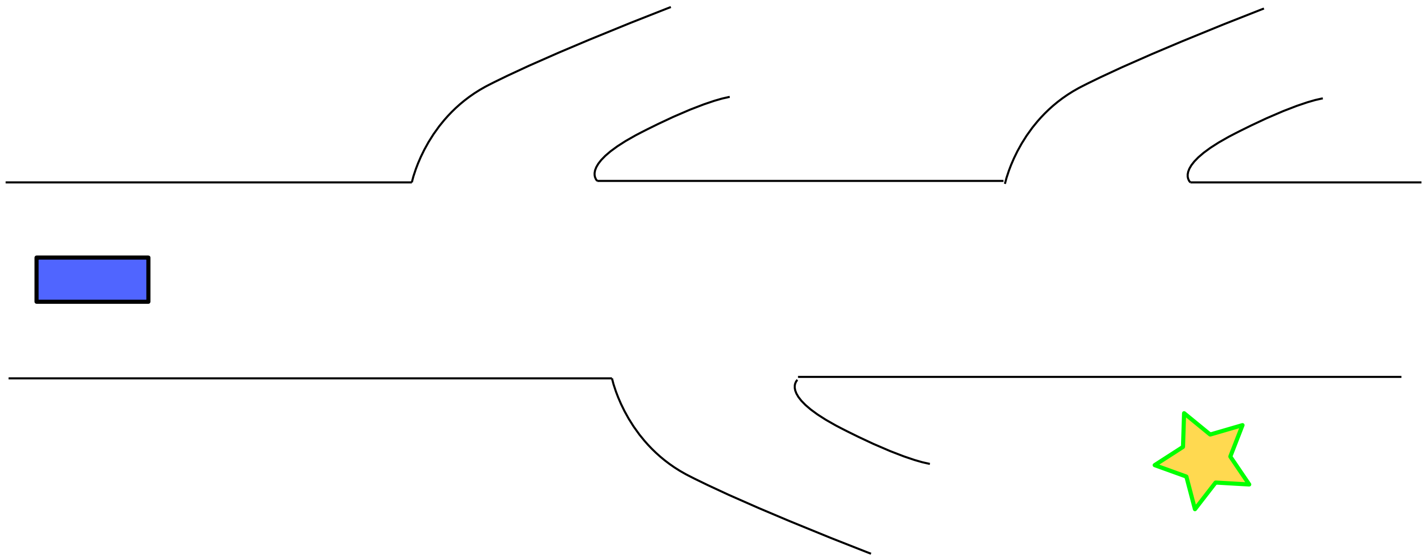

agent performs search for a bounded time

Introduction
Real-time Search
Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search
Conclusion

An example: highway navigation

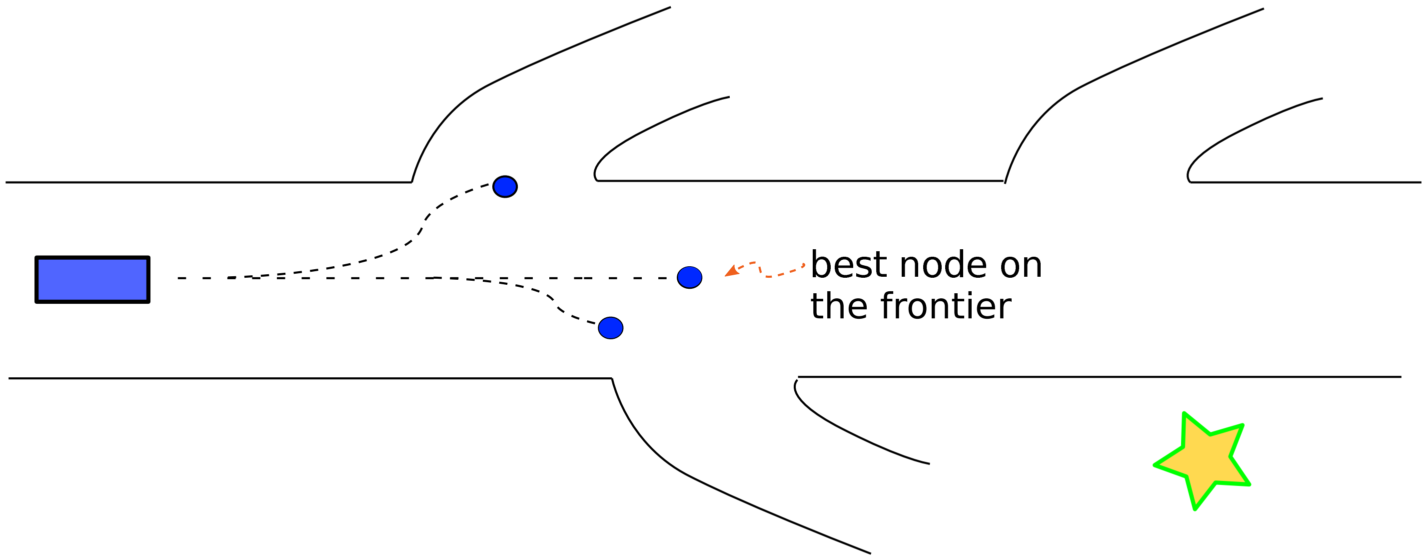

agent performs search for a bounded time

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

An example: highway navigation

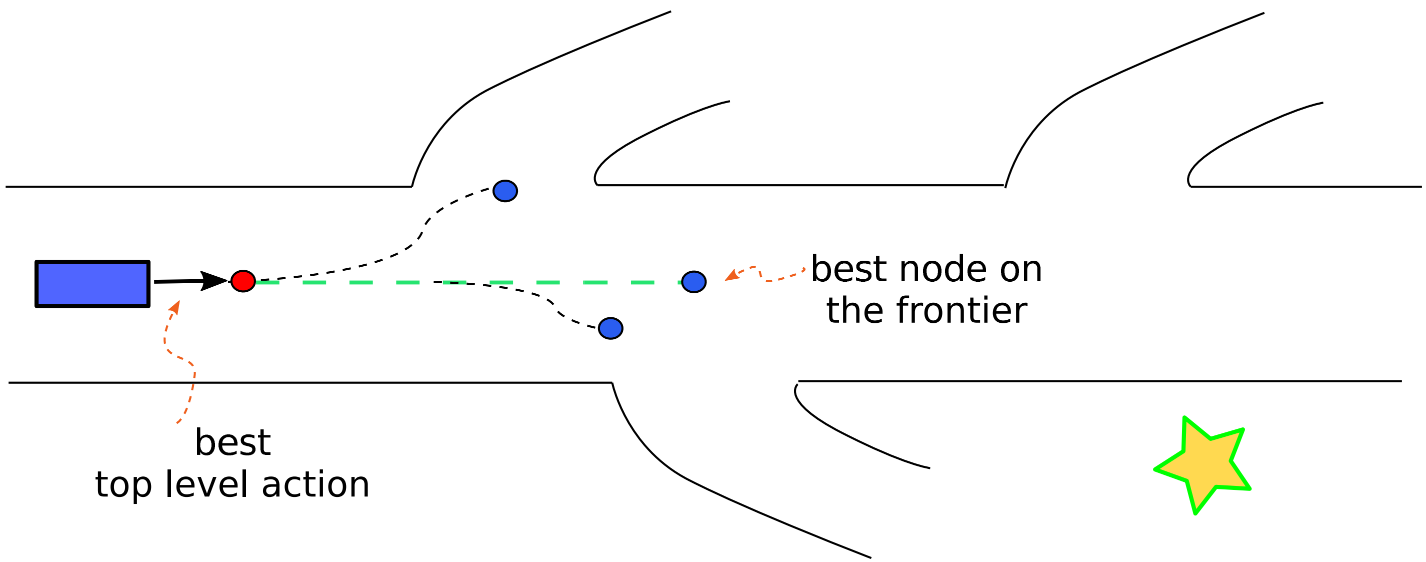

agent commits to best action and executes

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

An example: highway navigation

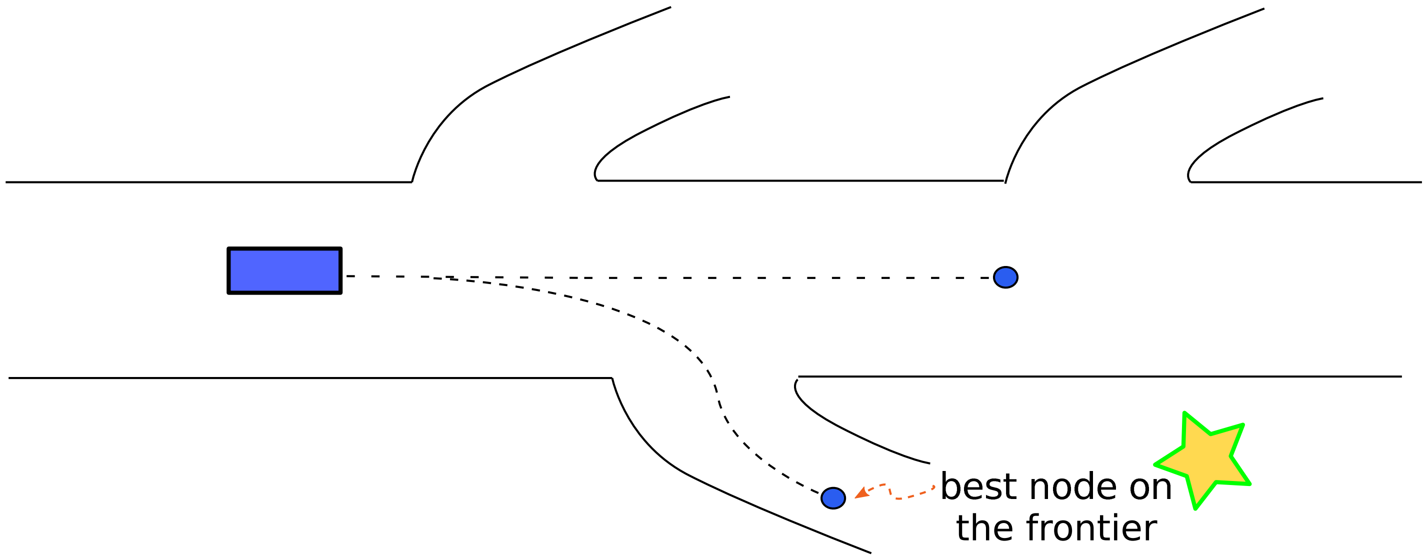

agent commits to best action and executes

An Example of Online Planning Using Heuristic Search

Introduction

Real-time Search

Online Planning

- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

An example: highway navigation

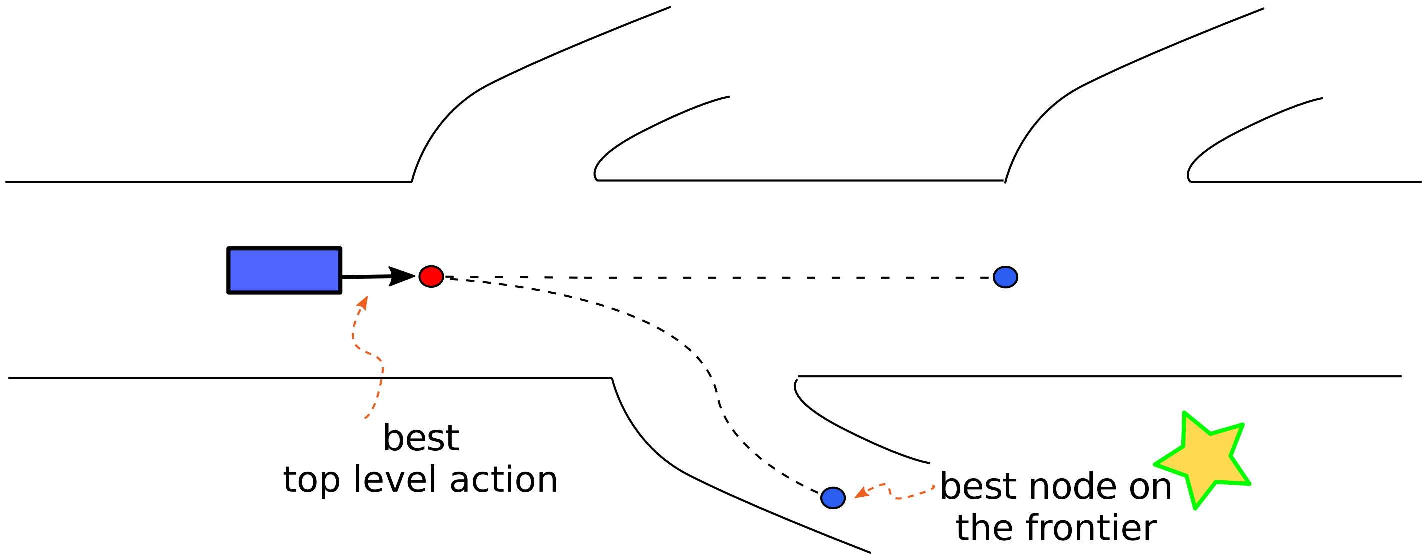

online planning: interleaving search and action execution

"receding horizon control"

The Meta-level Problem: Commit or Not Commit

Introduction
Real-time Search
Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search
Conclusion

For each node along the best prefix path:
should we commit?

The Meta-level Problem: Commit or Not Commit

For each node along the best prefix path:

should we commit?

fixed strategies:

always commit one (Korf 1990)

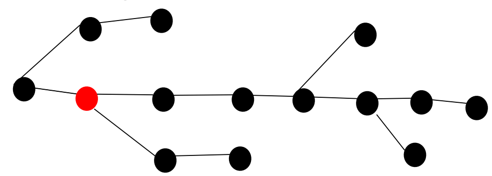

always commit all (Koenig&Sun 2008, Burns et al 2013)

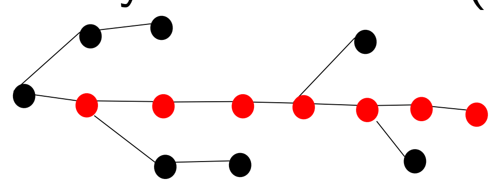

Can we do better?

Introduction
Real-time Search
Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search
Conclusion

always commit all is too risky

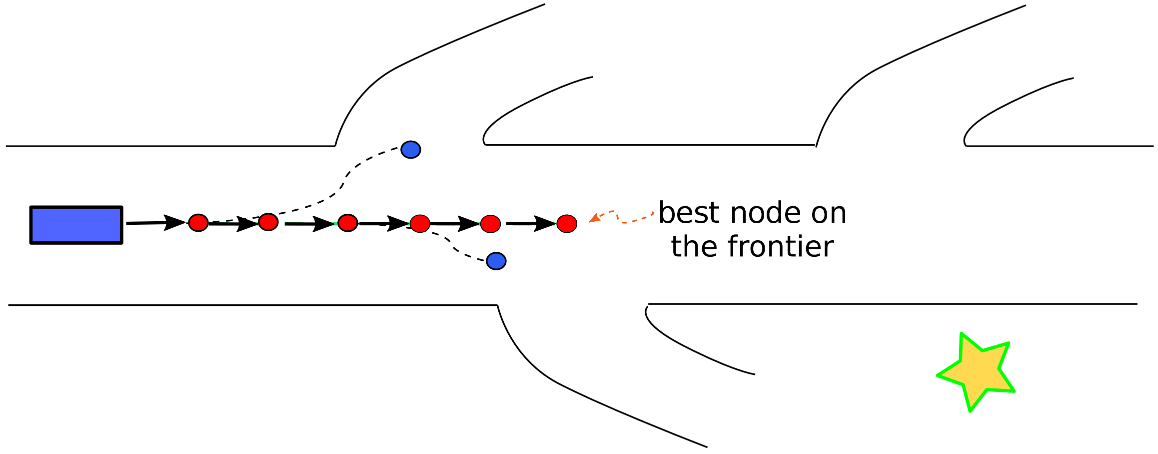

The Meta-level Problem: Commit or Not Commit

Introduction

Real-time Search

Online Planning

- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

always commit all is too risky

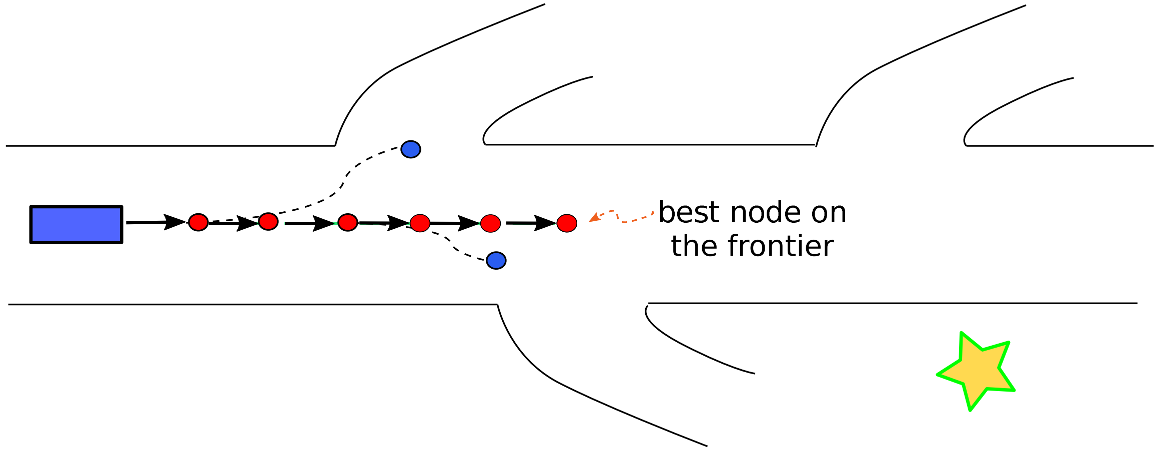

always commit one is too conservative

The Meta-level Problem: Commit or Not Commit

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

ideal:

commit if an action in prefix is certainly the best

to gain more planning time for next iteration

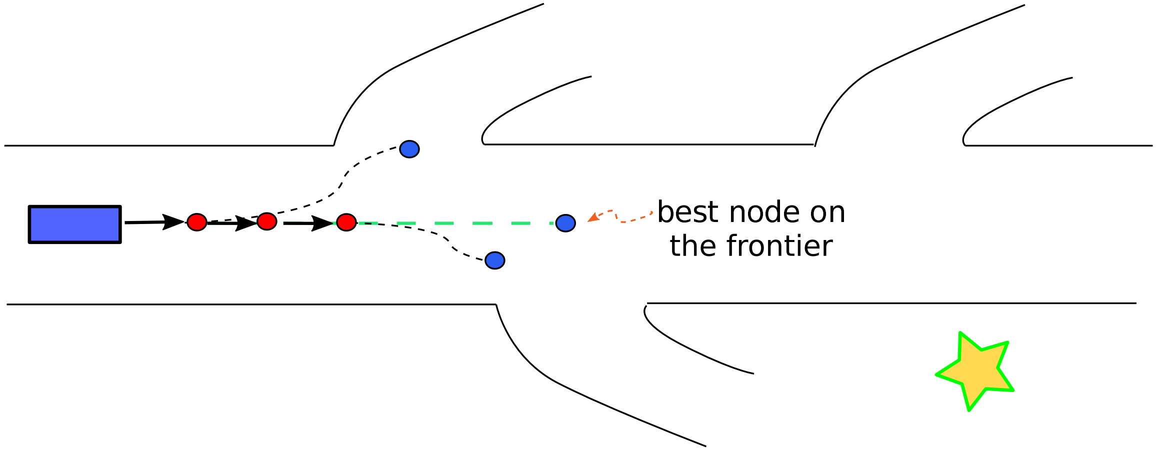

Our Approach: Flexible Action Commitment Search (FACS)

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

we propose a principled way to make meta-level decision

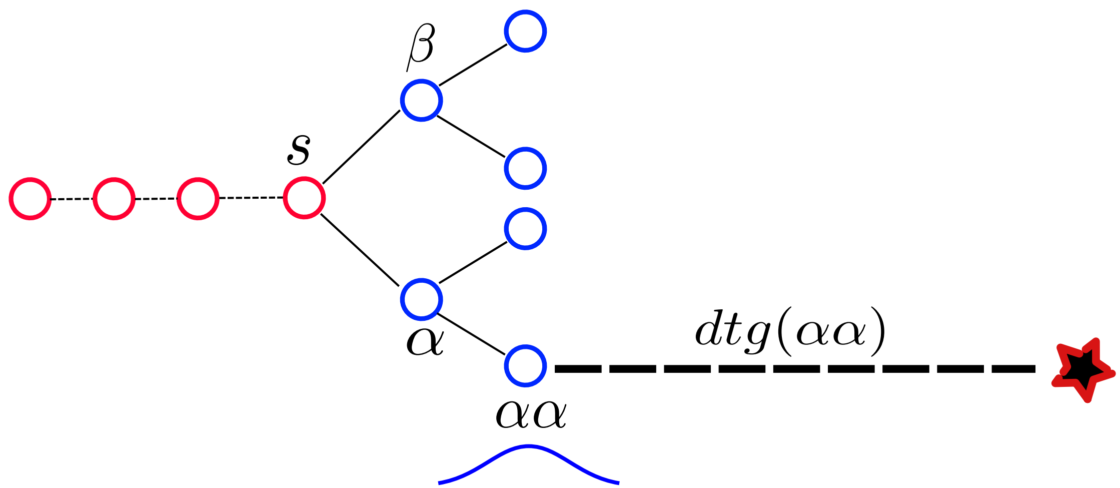

$$U_{\text{commit}} = \mathbb{E}\left[\min(X^d_{\alpha\alpha}, X^d_{\alpha\beta})\right]$$

$$U_{\text{don't commit}} = P_{\text{choose } \alpha} \cdot U_\alpha + (1 - P_{\text{choose } \alpha}) \cdot U_\beta$$

Results

Introduction

Real-time Search

Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary

Bounded-Cost Search

Conclusion

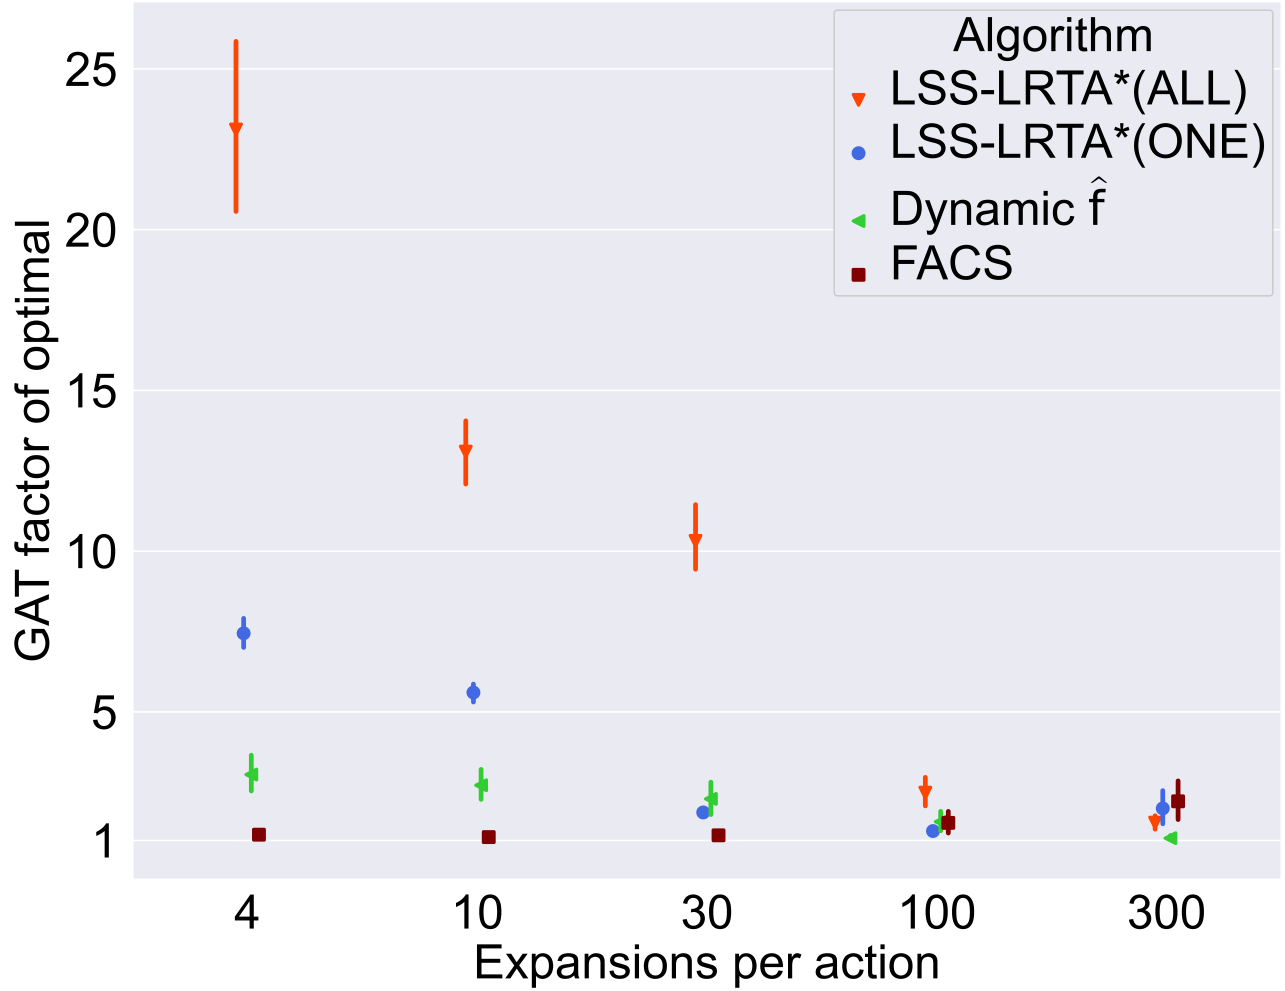

FACS consistently performs the best

Summary

Introduction
Real-time Search
Online Planning
- An Example
- Action Commmitment
- Our Approach
- Results
- Summary
Bounded-Cost Search
Conclusion

- FACS starts to explore a principled way of doing online action commitment by reasoning uncertainty
- FACS is better than fixed baseline strategies in synthetic grid pathfinding scenarios.
- Deliberation on how to allocate search effort can benefit online planning

Status of the work:

1. *When to Commit to an Action in Online Planning.* ICAPS Workshop on Integrating, Planning, Acting, and Execution (IntEx-21), 2021
2. In preparation: *When to Commit to an Action in Online Planning.*, International Symposium on Combinatorial Search (SoCS 2022)
3. In preparation: *Situated Safe Interval Path Planning for Dynamic Environments.*, International Joint Conference on Artificial Intelligence (IJCAI 2022)

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
- State of The Art
- XES
- Result
Conclusion

Bounded-Cost Search Using Estimates of Uncertainty

Joint work with Maximilian Fickert and Wheeler Ruml

What is Bounded-Cost Search?

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
- State of The Art
- XES
- Result
Conclusion

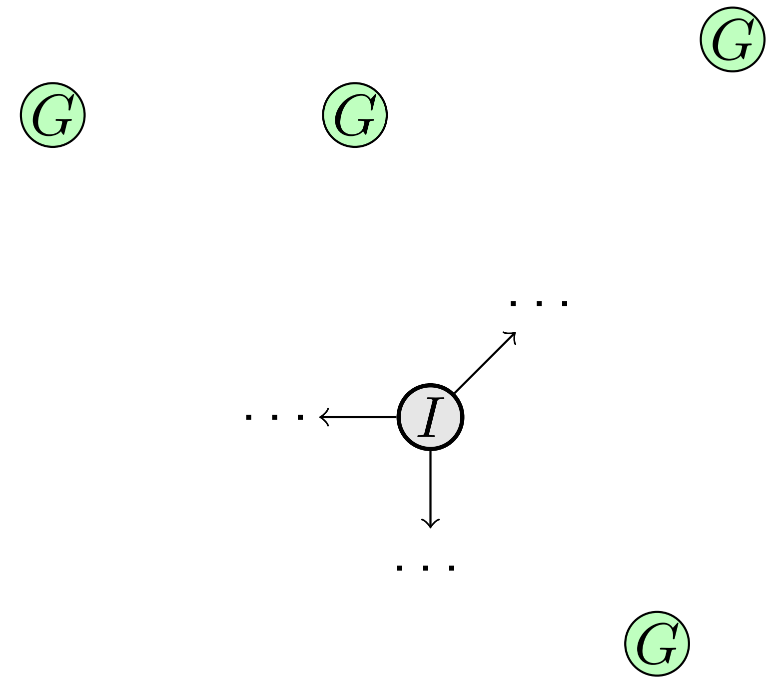

What is Bounded-Cost Search?

Introduction

Real-time Search

Online Planning

Bounded-Cost Search

- State of The Art
- XES
- Result

Conclusion

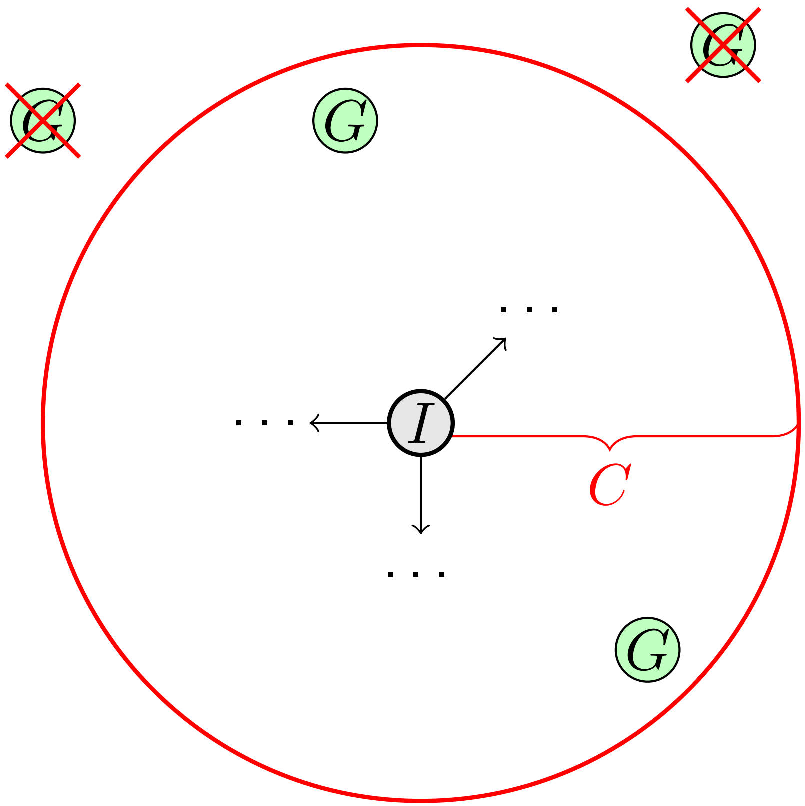

Objective: Find a plan with cost at most C as fast as possible.

State of The Art

Potential Search:

best first search on $f_{lnr}(n) = \frac{h(n)}{C - g(n)}$

does not consider search effort

BEES:

expand the node, among those estimated

to be within the bound, that is closest to a goal

does not consider the uncertainty of its estimate (brittle)

Our Approach: Expected Effort Search

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
- State of The Art
- XES
- Result

Conclusion

1. Explicitly estimate the probability of finding a solution within bound $p(n)$

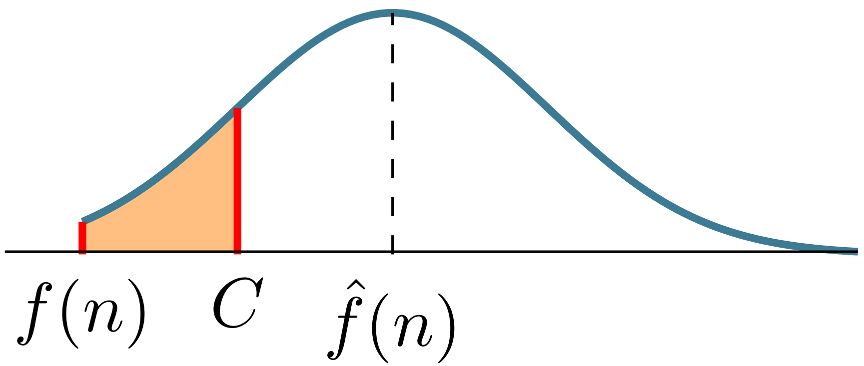

2. estimate total search effort by $d(n)$
3. best first search on expected search effort $d(n)/p(n)$

optimal efficiency proof under three unrealistic assumptions (see dissertation for detail)

Our Approach: Expected Effort Search

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
- State of The Art
- XES
- Result

Conclusion

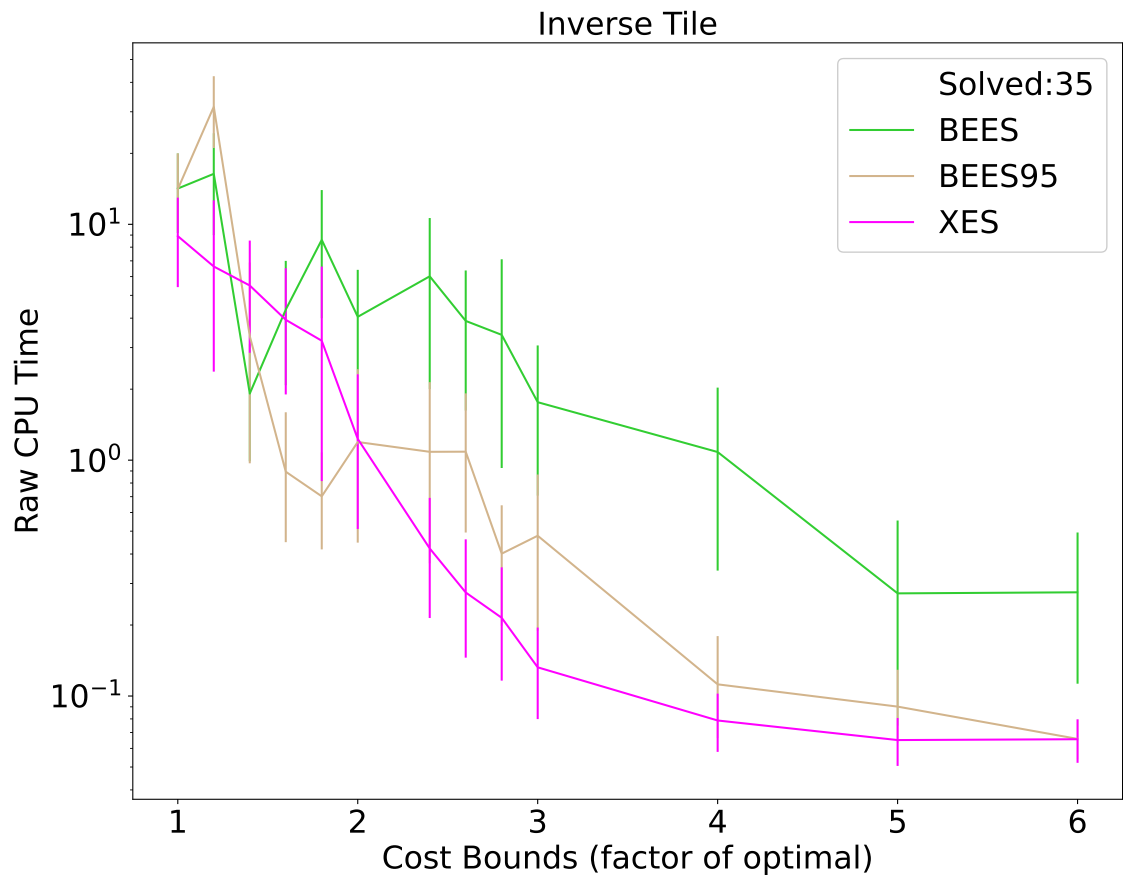

XES finds a valid solution faster!

Our Approach: Expected Effort Search

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
- State of The Art
- XES
- Result

Conclusion

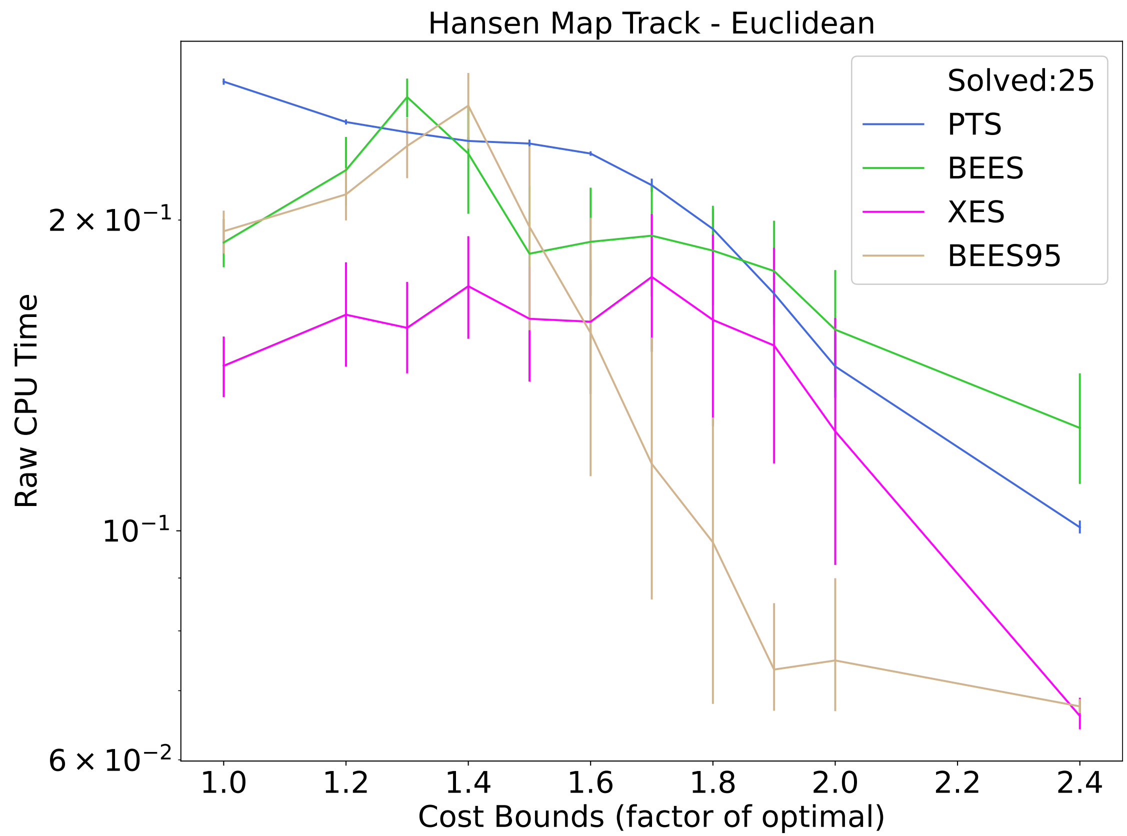

Status of the work:

1. *Bounded-Cost Search Using Estimates of Uncertainty.* International Joint Conference on Artificial Intelligence (IJCAI), 2021.
2. under review: *New Results in Bounded-Suboptimal Search.* AAAI Conference on Artificial Intelligence (AAAI), 2022.

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
■ Conclusions

Conclusion

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
■ Conclusions

The thesis of my dissertation:
heuristic search can benefit from representing uncertainty

- real-time search
- concurrent planning and execution
- bounded-cost search

Questions?

- Introduction
- Real-time Search
- Online Planning
- Bounded-Cost Search
- Conclusion
- Questions
  - Questions?

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

Back-up Slides

How To Compute F-Hat

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

$$\hat{f} = g + \hat{h} = g + h + \epsilon d$$

$\hat{f}$ is the expected value for optimal plan cost

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

Best-first search on the expected effort: $\frac{T}{p}$

- $T(n)$: search effort to find a solution under n
- $p(n)$: probability that n leads to a solution within C

Expected Effort Search (XES)

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

Best-first search on the expected effort: $\frac{T}{p}$

- $T(n)$: search effort to find a solution under n
- $p(n)$: probability that n leads to a solution within C

n_1: $T = 10$, $p = 0.5$ $\rightsquigarrow 20$

n_2: $T = 6$, $p = 0.25$ $\rightsquigarrow 24$

Our Approach: Flexible Action Commitment Search (FACS)

Introduction

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

Back-up Slides

- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

we propose a principled way to make meta-level decision

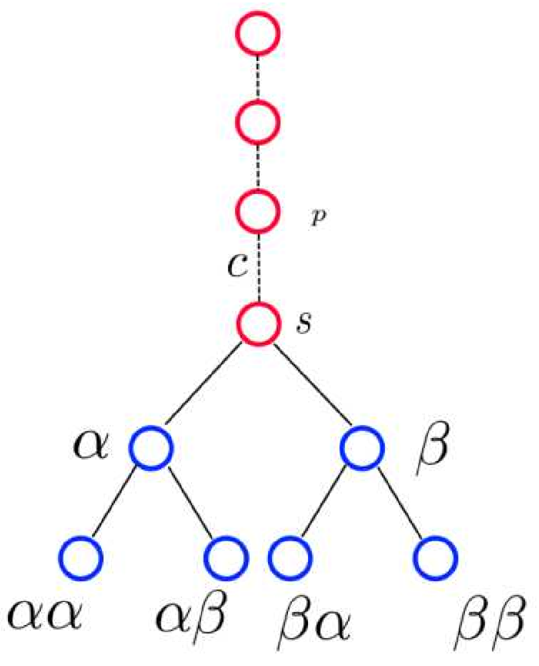

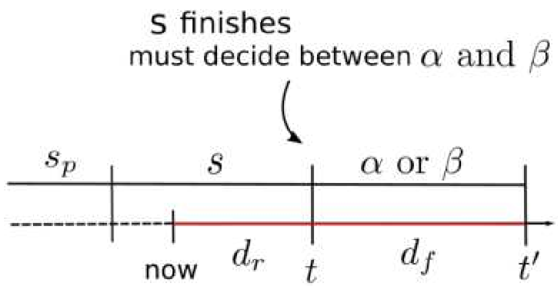

Our Approach: Flexible Action Commitment Search (FACS)

we propose a principled way to make meta-level decision

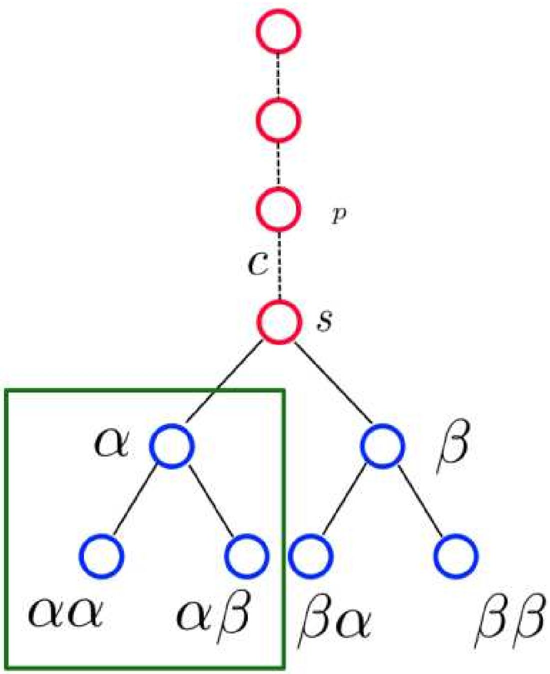

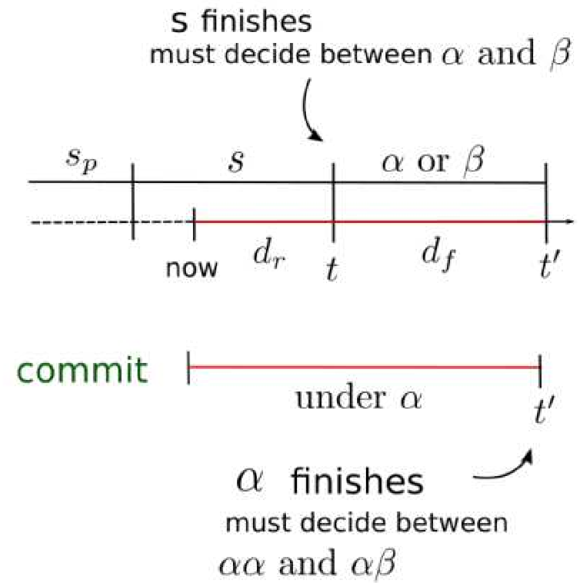

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

we propose a principled way to make meta-level decision

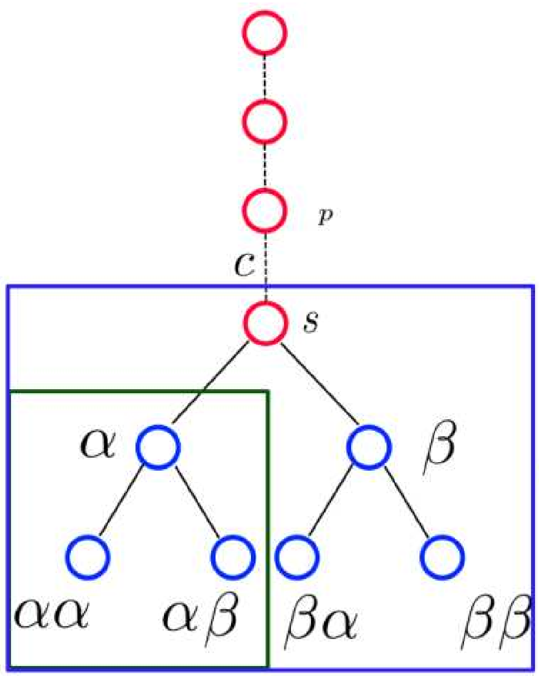

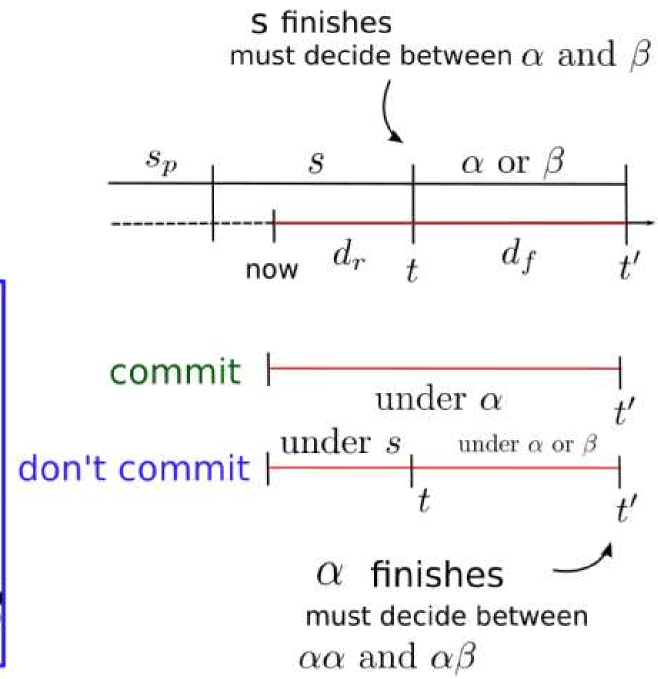

FACS: The Effect of Search

Introduction

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

Back-up Slides

- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

belief of where $\hat{f}$ will be after search:

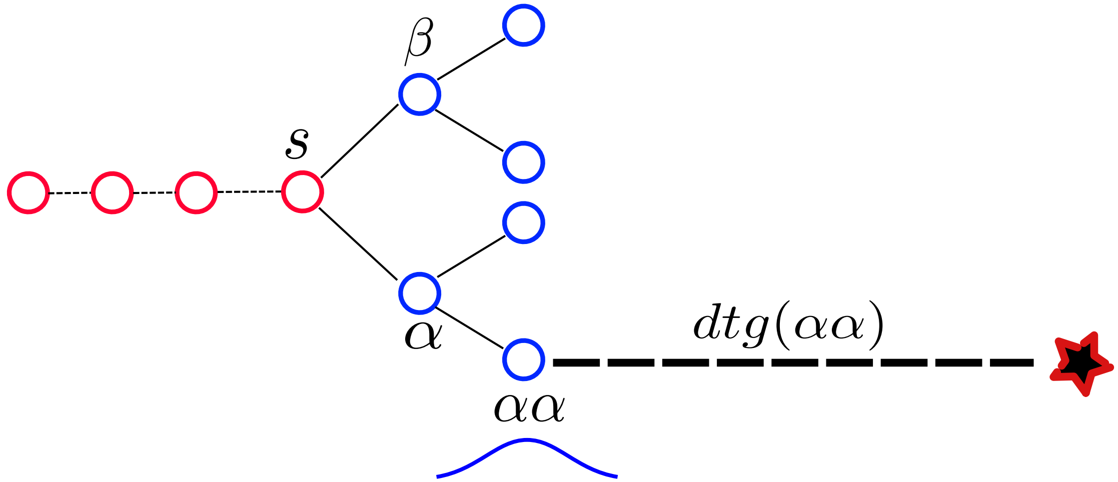

FACS: The Effect of Search

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

belief of where $\hat{f}$ will be after search:

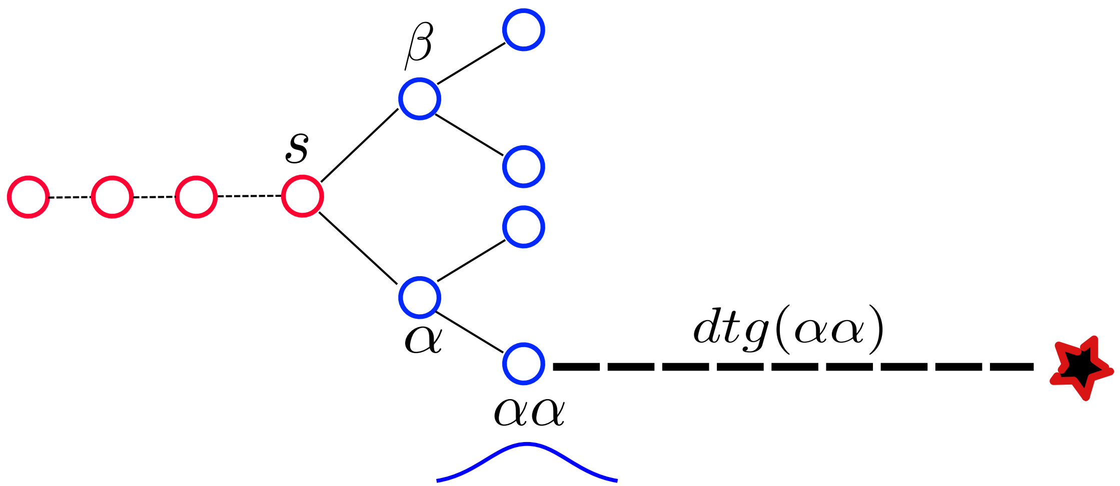

$$X_{\alpha\alpha}^{d} \sim \mathcal{N}(\hat{f}, var(dtg, d))$$

FACS: Compute Utility

Introduction

Real-time Search

Online Planning

Bounded-Cost Search

Conclusion

Back-up Slides

- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

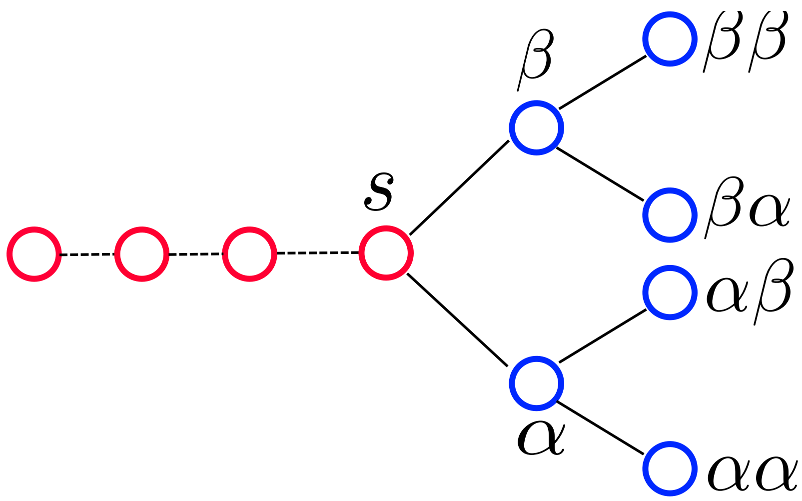

$$U_{\text{commit}} = \mathbb{E}\left[\min(X^d_{\alpha\alpha}, X^d_{\alpha\beta})\right]$$

where $d = (d_r + d_f)/2$

$$U_{\text{don't commit}} = P_{\text{choose } \alpha} \cdot U_\alpha + (1 - P_{\text{choose } \alpha}) \cdot U_\beta$$

commit when $U^{t'}_{\text{commit}} > U^{t'}_{\text{don't commit}}$

Synthetic Grid Pathfinding

Introduction
Real-time Search
Online Planning
Bounded-Cost Search
Conclusion
Back-up Slides
- f-hat
- Expected Effort Search (XES)
- FACS Detail
- FACS Belief
- FACS Decision
- FACS Domain

- Left: tar pit area $\rightarrow$ high cost for reckless committing
- Right: corridor area $\rightarrow$ need long lookahead to observe the local minima
- Middle: empty area $\rightarrow$ gain lookahead, no harm to commit