Planning Under Time Pressure as Decision-making Under Uncertainty

Tianyi Gu

Advisor: Wheeler Ruml

University of New Hampshire

What is Planning?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

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 Planning?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

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

one method of planning: heuristic search!

What is Planning?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

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

one method of planning: heuristic search!

heuristic search:

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

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

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

heuristic search associates costs with states, used to guide search

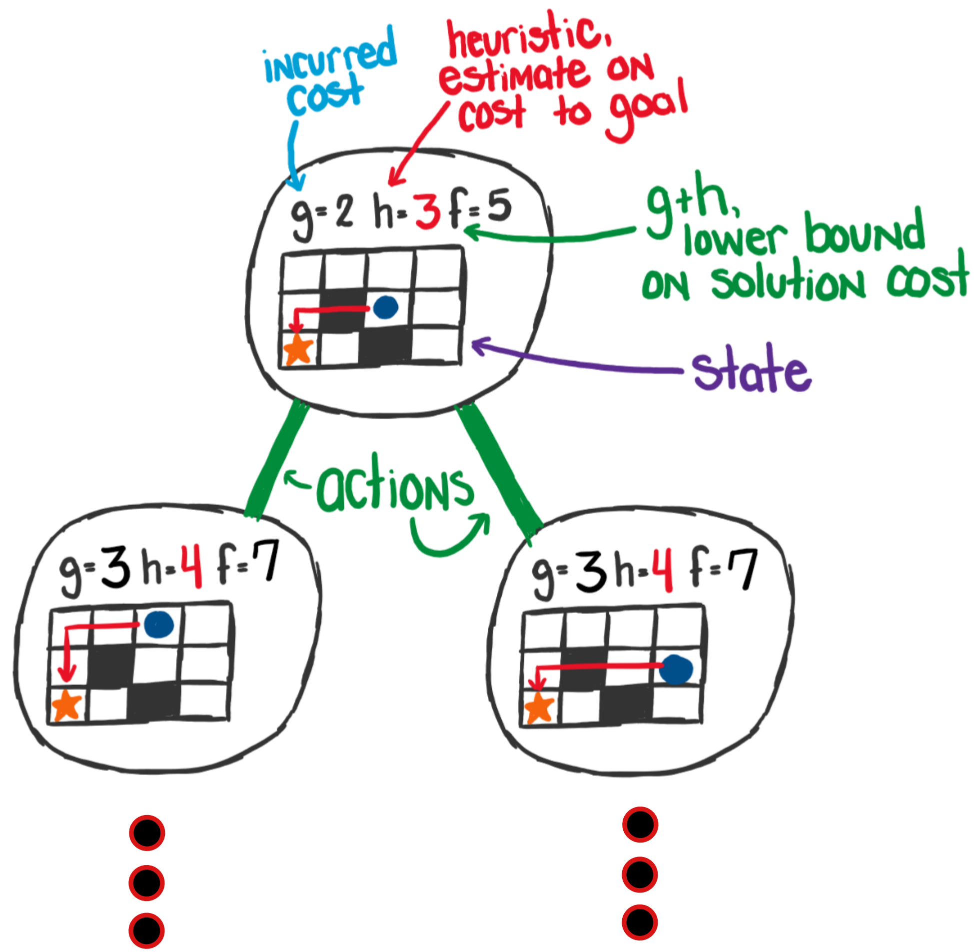

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

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
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

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]

other alternatives to optimal search:
anytime search, greedy search
no guarantee faster than A*[2]

[1] How Good is Almost Perfect, Malte Helmert and Gabriele Roger, AAAI, 2008.
[2] When does Weighted A* Fail, Christopher Wilt and Wheeler Ruml, SoCS, 2012.

Heuristic Search

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]

other alternatives to optimal search:
anytime search, greedy search
no guarantee faster than A*[2]

What if we need strong guarantee on responsiveness?

real-time heuristic search!

[1] How Good is Almost Perfect, Malte Helmert and Gabriele Roger, AAAI, 2008.
[2] When does Weighted A* Fail, Christopher Wilt and Wheeler Ruml, SoCS, 2012.

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

agent performs search for a bounded time

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

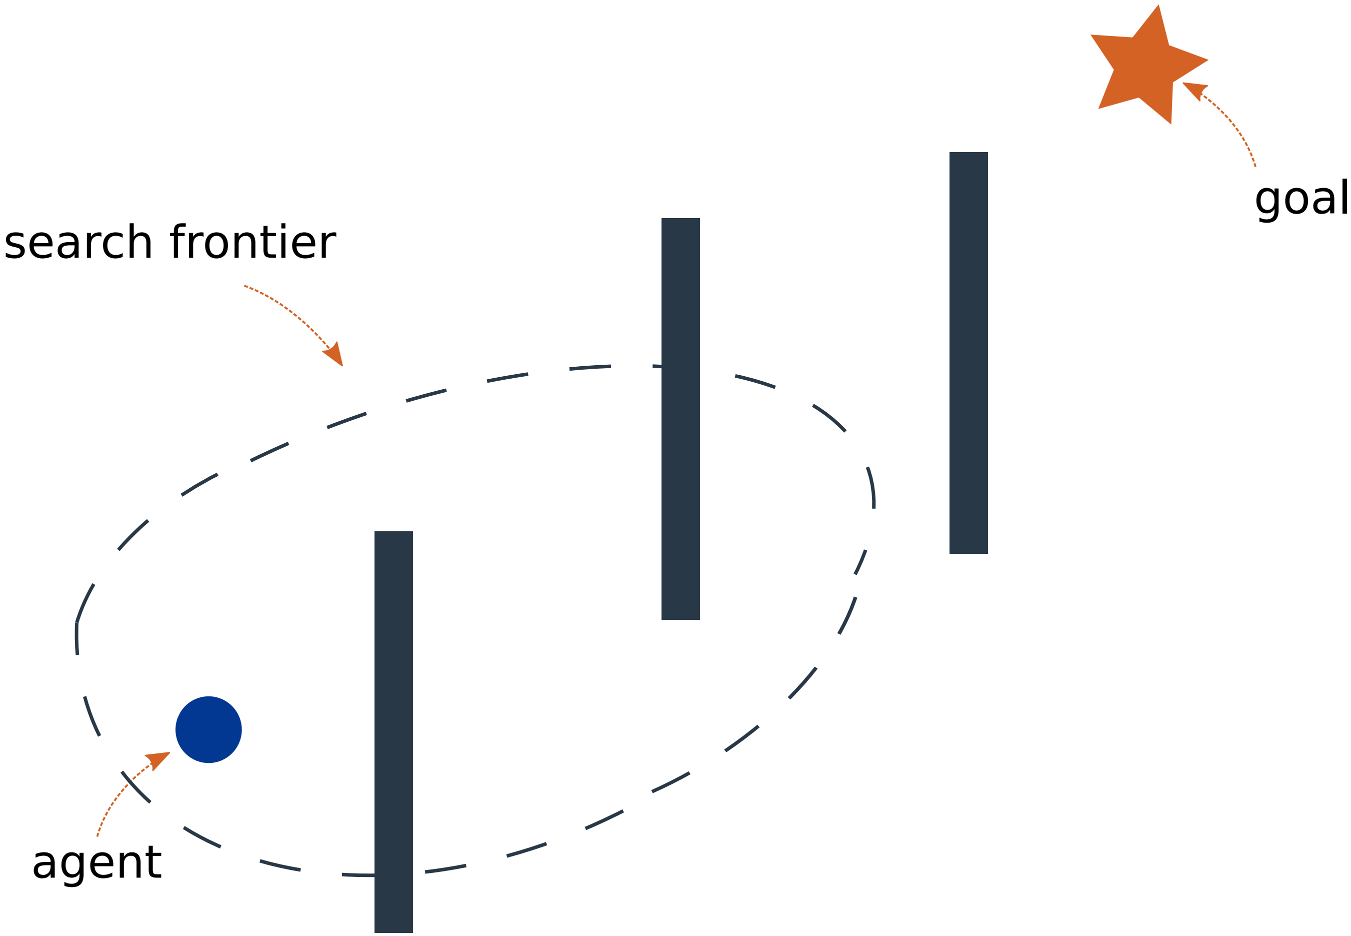

agent performs search for a bounded time

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

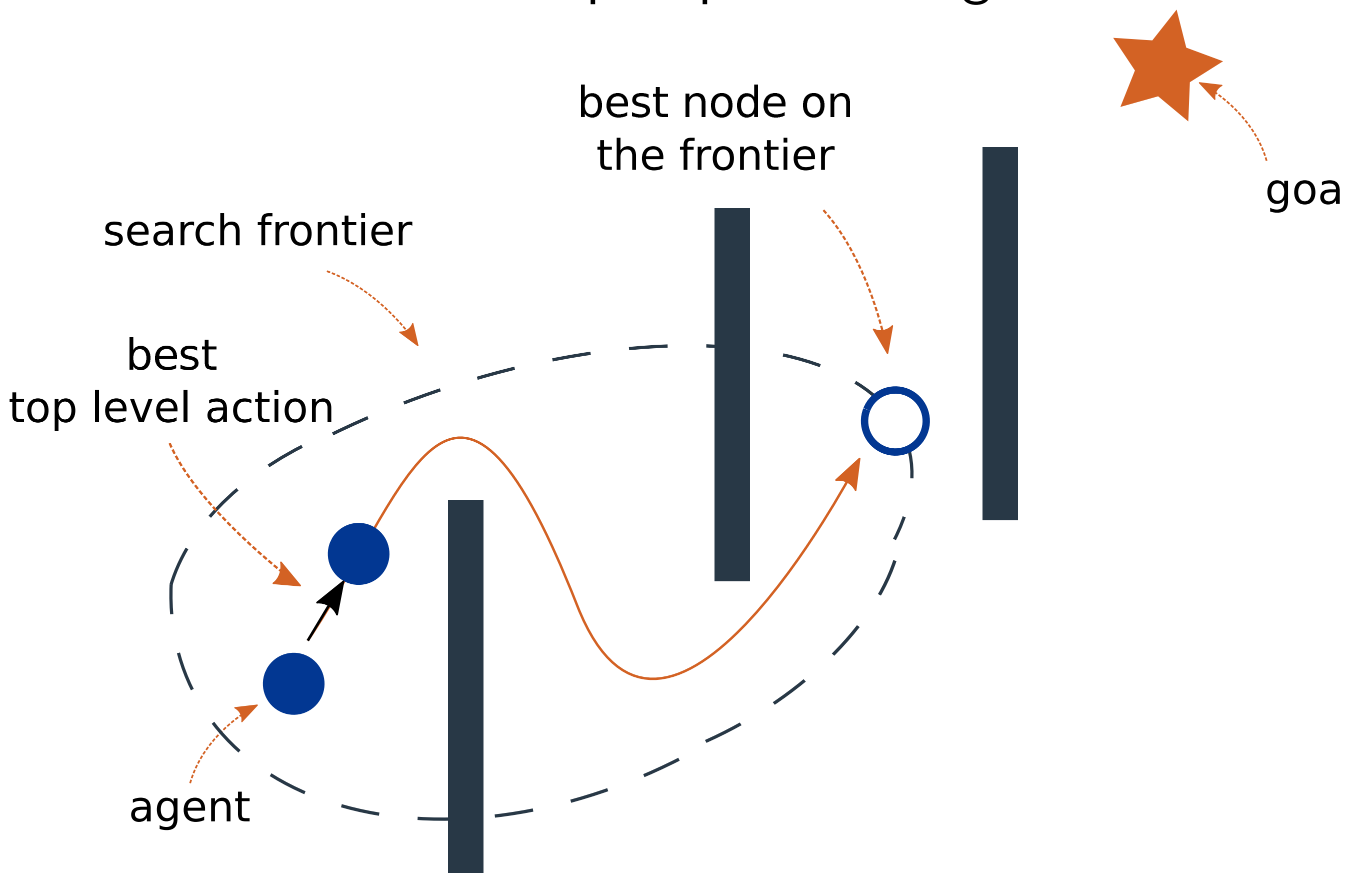

agent commits to best action and executes

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

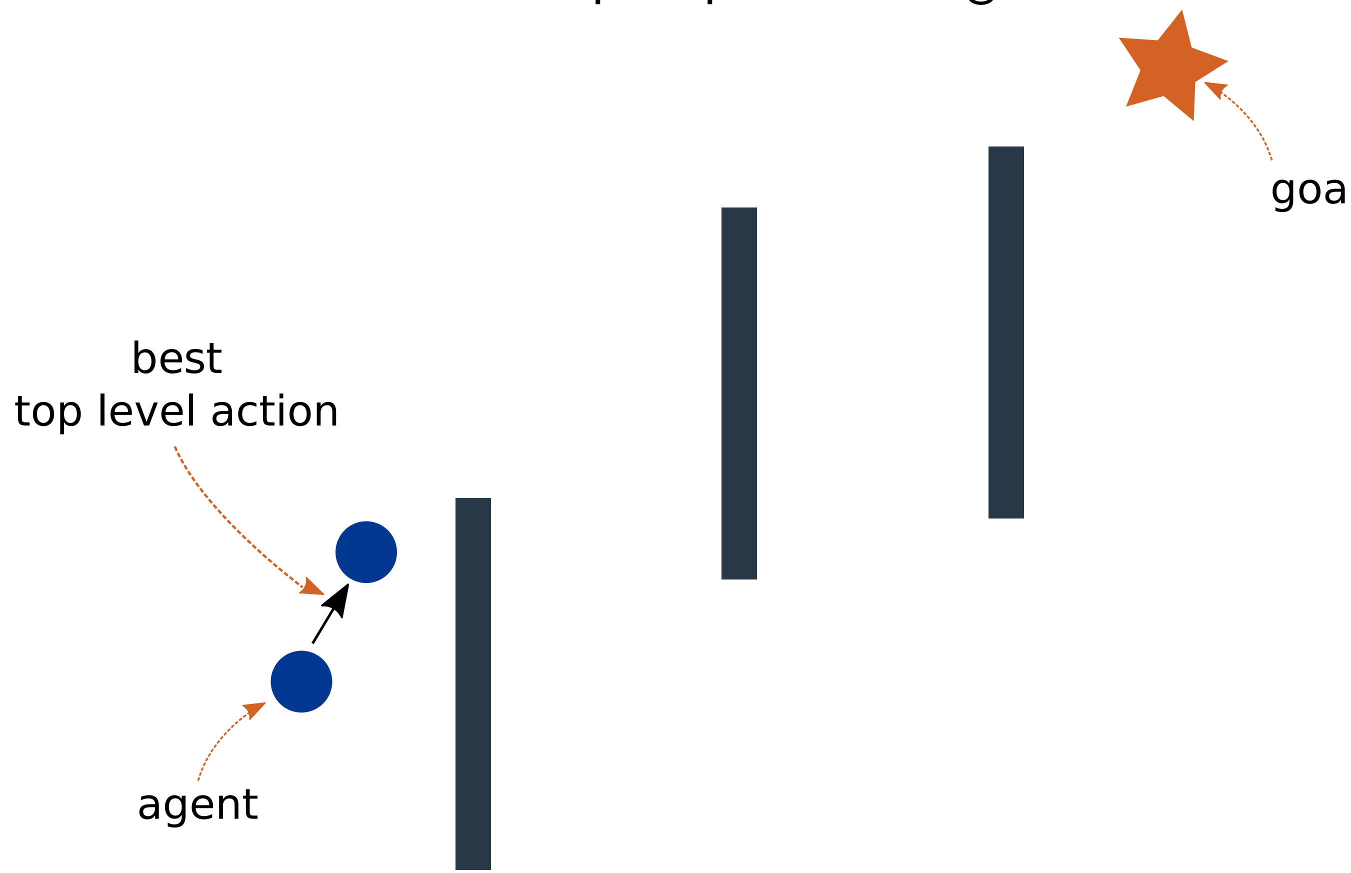

agent commits to best action and executes

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

agent commits to best action and executes

What is Real-time Heuristic Search?

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

An example: path finding

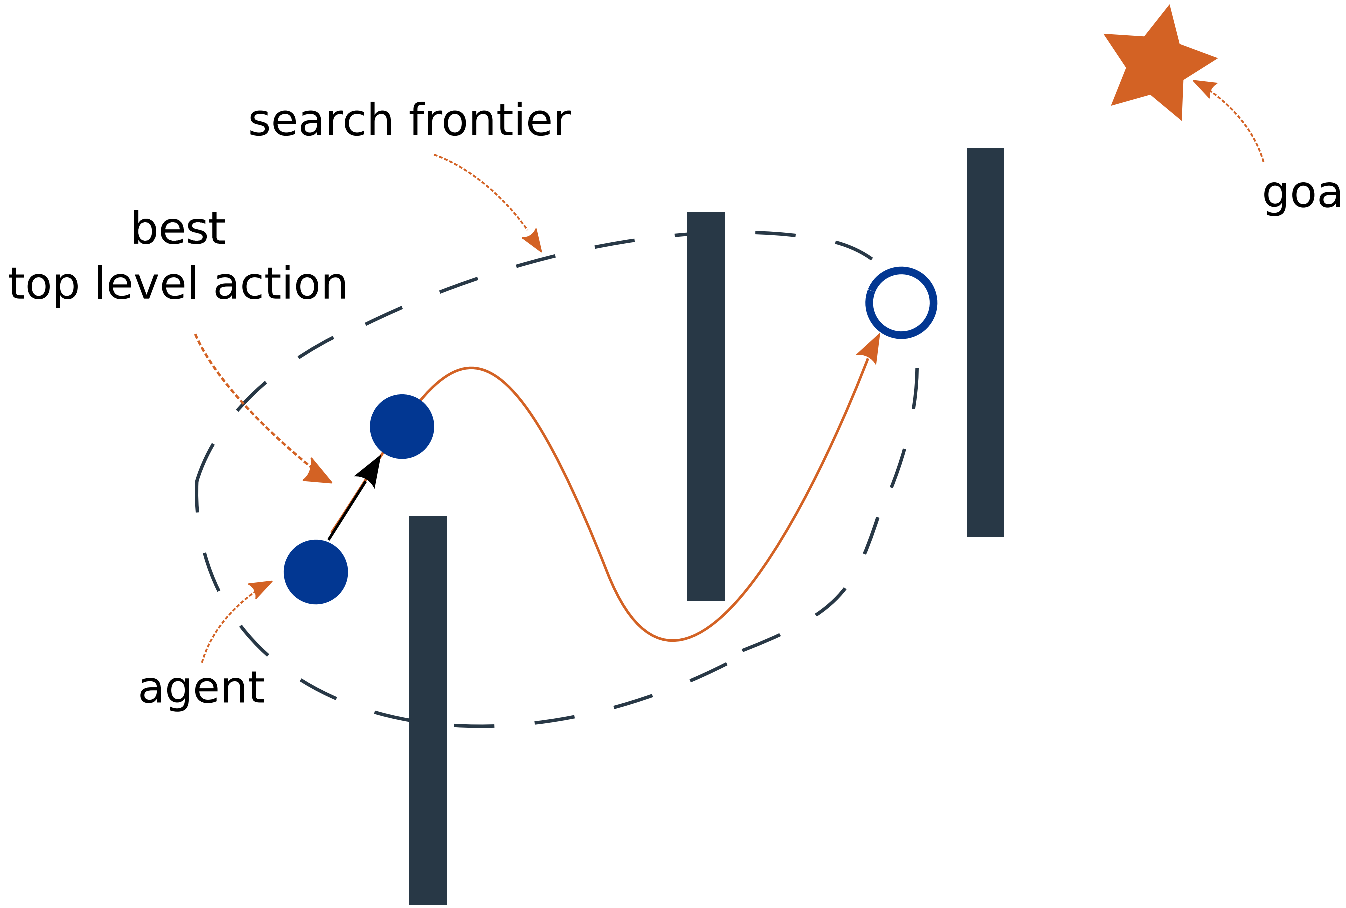

online planning: interleaving search and action execution

"receding horizon control"

Motivation for Real-time Heuristic Search

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

Real-time heuristic search:

return the next action within a time bound

Applications:

interacting with humans
- vacuum robot
- siri agent

dynamic environment
- autonomous vehicle
  inaccurate sensor
  update model online

Motivation for Real-time Heuristic Search

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

Real-time heuristic search:

return the next action within a time bound

Applications:

interacting with humans
- vacuum robot
- siri agent

dynamic environment
- autonomous vehicle
  inaccurate sensor
  update model online

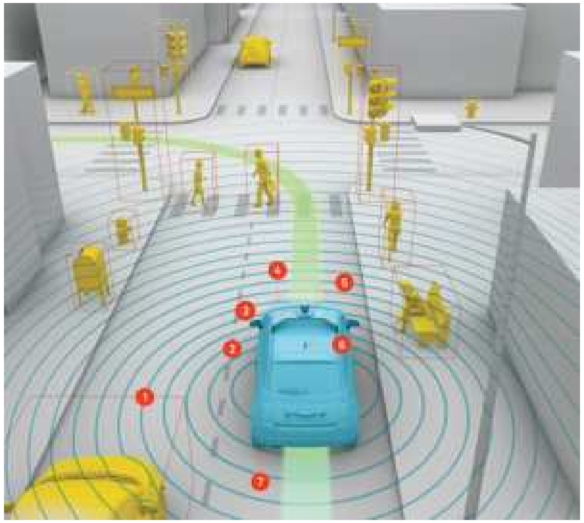

Overview

Introduction
- Planning
- Heuristic Search
- Real-time Search
- Overview

Nancy

Data-Driven Nancy

Other Research

Conclusions

Guide search by belief distribution:
a better way to plan under time pressure

- The Nancy Framework
  reconsider real-time search
- Data-Driven Nancy
  a more flexible model
- Other Research
  suboptimal search and robotics

Introduction

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

Real-time Search as Decision-making Under Uncertainty: The Nancy Framework

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

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy
Other Research
Conclusions

three phases:

1. Lookahead Phase:
   expands nodes with minimum f
   to explore the search space

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

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs
Data-Driven Nancy
Other Research
Conclusions

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

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

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
Nancy
- Decision-making
- Lookahead
- The Beliefs
Data-Driven Nancy
Other Research
Conclusions

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)

proved to be complete for consistent heuristic

repeat until at a goal

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

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy
Other Research
Conclusions

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)

proved to be complete for consistent heuristic

repeat until at a goal

derived from offline search, but optimal for online?

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs
Data-Driven Nancy
Other Research
Conclusions

Should an agent at A move to B_1 or B_2?
(x_i are unknown but i.i.d. uniform 0-1)

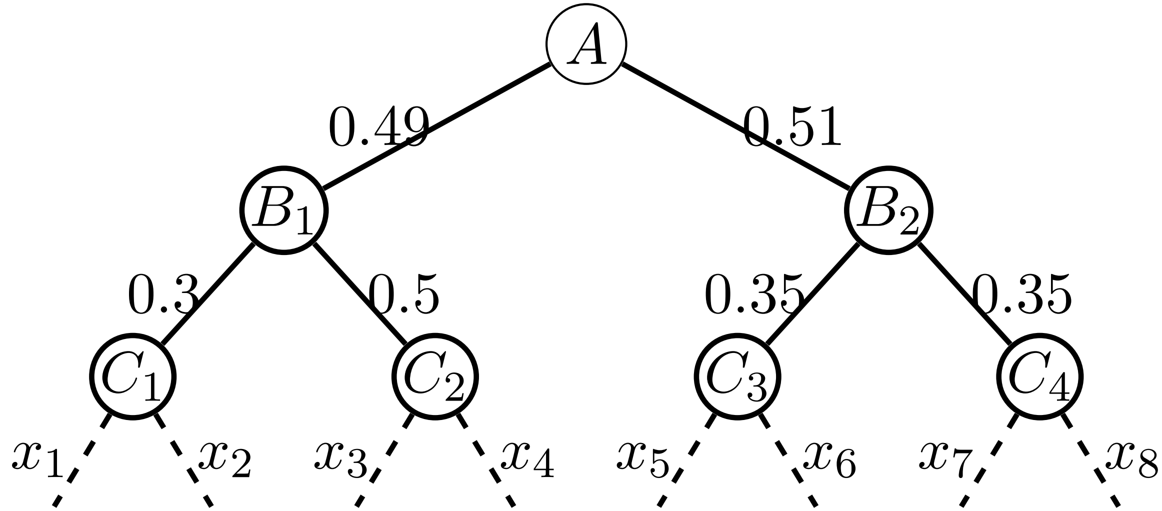

random tree domain (Pemberton & Korf 1995)

$f = g + h = g + 0$ is lower bound on optimal plan cost

Decision-making Phase: A Troublesome Example

Introduction

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

Should an agent at A move to B_1 or B_2?
(x_i are unknown but i.i.d. uniform 0-1)

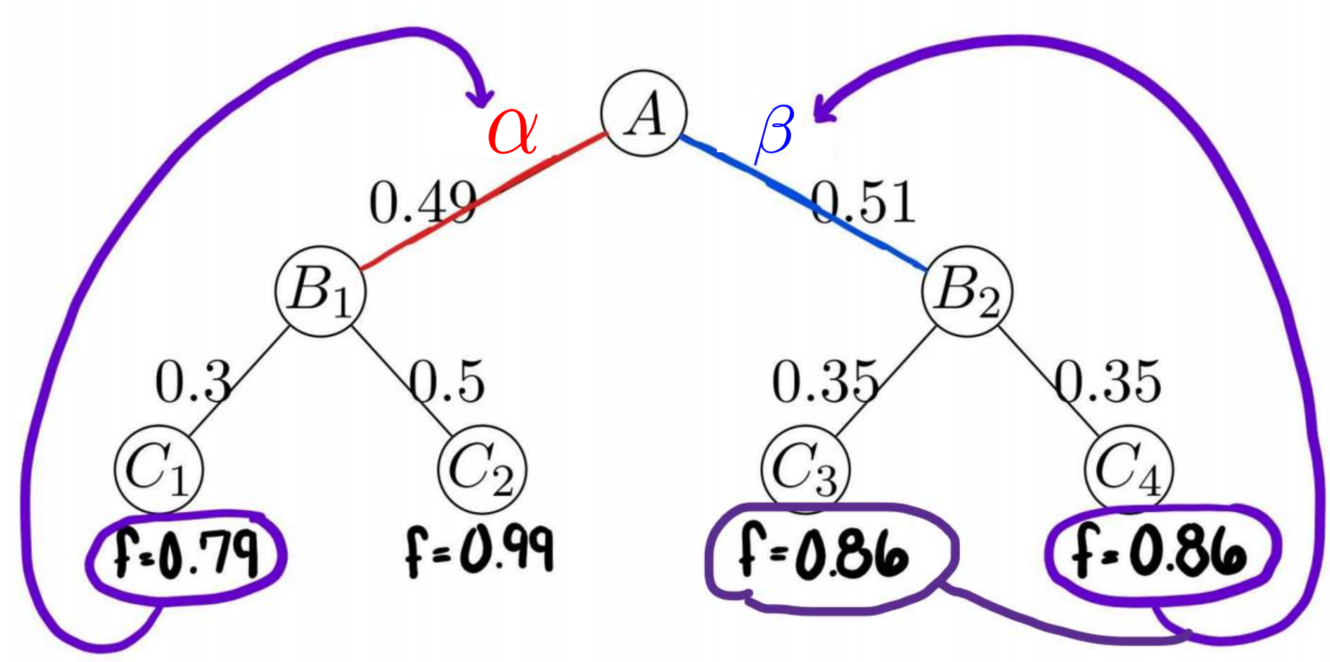

decision theory says minimize expected value
lower bound: not suitable for rational action selection

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy
Other Research
Conclusions

Should an agent at A move to B_1 or B_2?
(x_i are unknown but i.i.d. uniform 0-1)

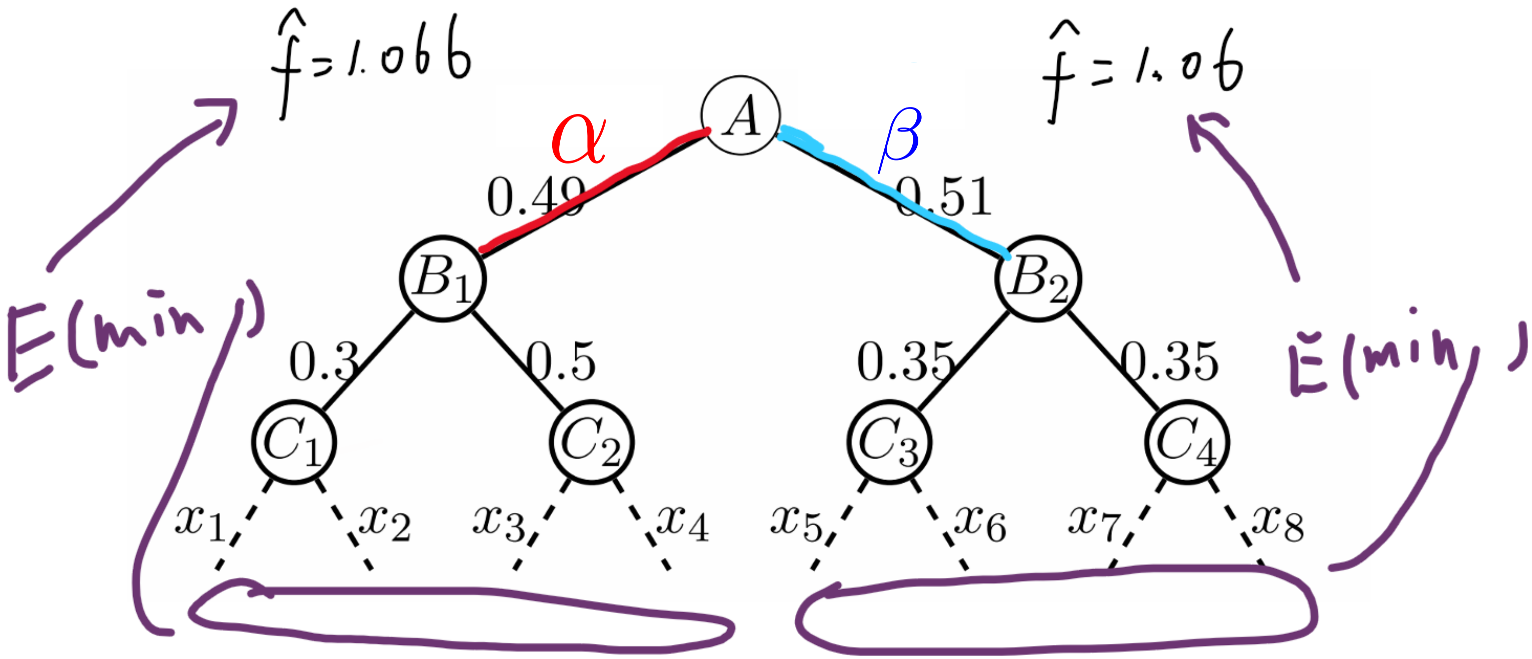

$\hat{f}$ is expected plan cost

Decision-making Phase: A Troublesome Example

Should an agent at A move to B_1 or B_2?
(x_i are unknown but i.i.d. uniform 0-1)

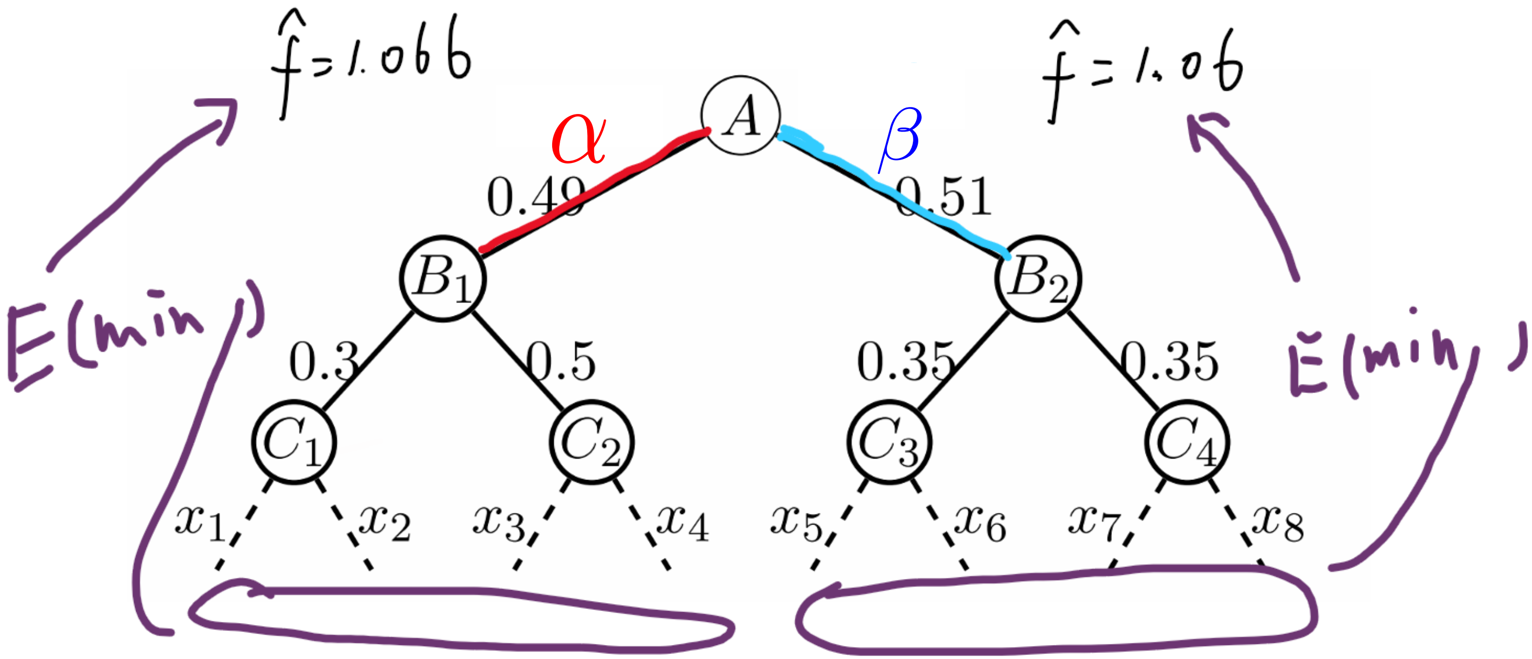

$\hat{f}$ is expected plan cost

f is not the answer: should minimize expected value!

plan under time pressure $\rightarrow$ bounded rationality

Lookahead Phase: A Troublesome Example

Introduction

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

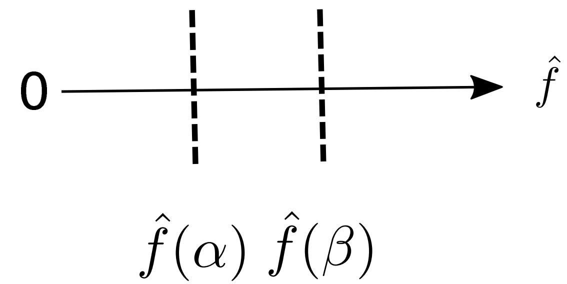

$\hat{f}$ is expected value

Should an agent expand nodes under α or β?

Lookahead Phase: A Troublesome Example

Introduction

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

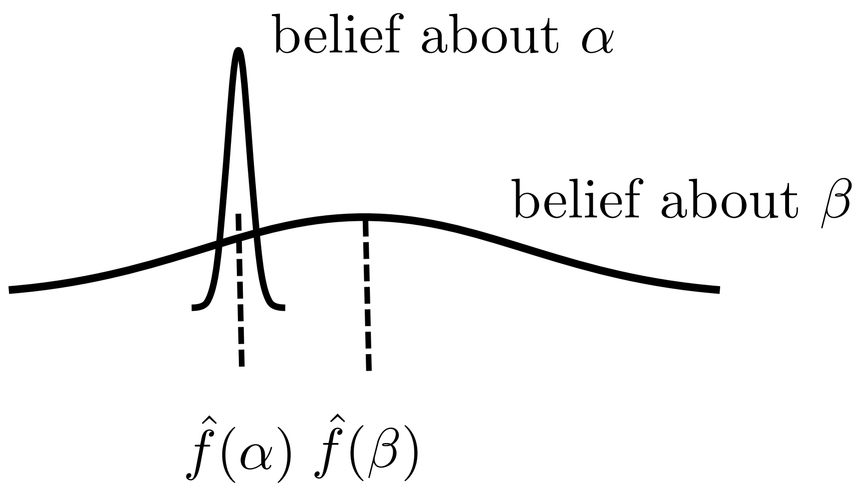

$\hat{f}$ is expected value

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

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy
Other Research
Conclusions

Risk-based lookahead [3]**:**

want to maximize value of information
expand nodes which minimize expected regret
relies on belief of values
choose expansions that decrease uncertainty in beliefs

[3] Real-time Planning as Decision-making Under Uncertainty, Andrew Mitchell, Wheeler Ruml, Fabian Spaniol, Joerg Hoffmann, and Marek Petrik, AAAI, 2019.

Risk-based Lookahead Example

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs
Data-Driven Nancy
Other Research
Conclusions

expand under α or β?

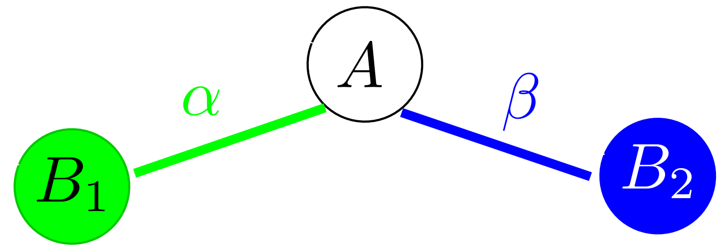

Risk-based Lookahead Example

Introduction
Nancy
Decision-making
Lookahead
The Beliefs
Data-Driven Nancy
Other Research
Conclusions

expand under α or β?

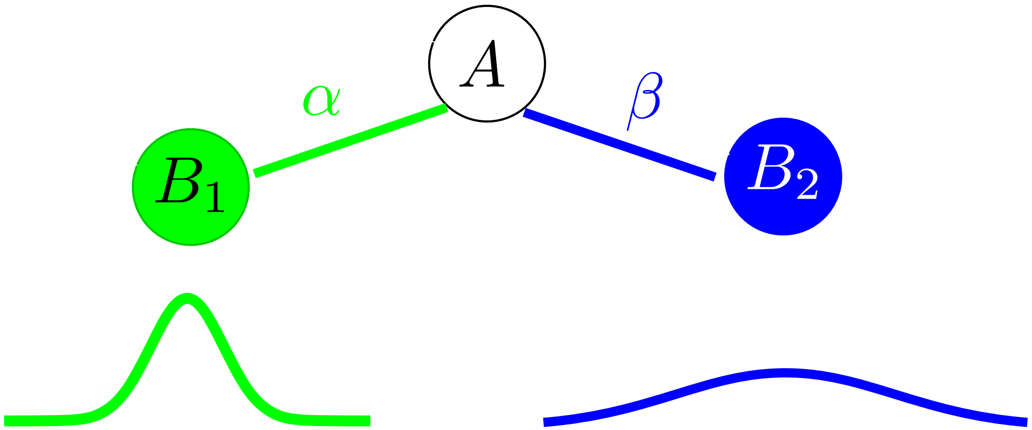

need 2 things:

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

Risk-based Lookahead Example

Introduction
Nancy
Decision-making
Lookahead
The Beliefs
Data-Driven Nancy
Other Research
Conclusions

expand under α or β?

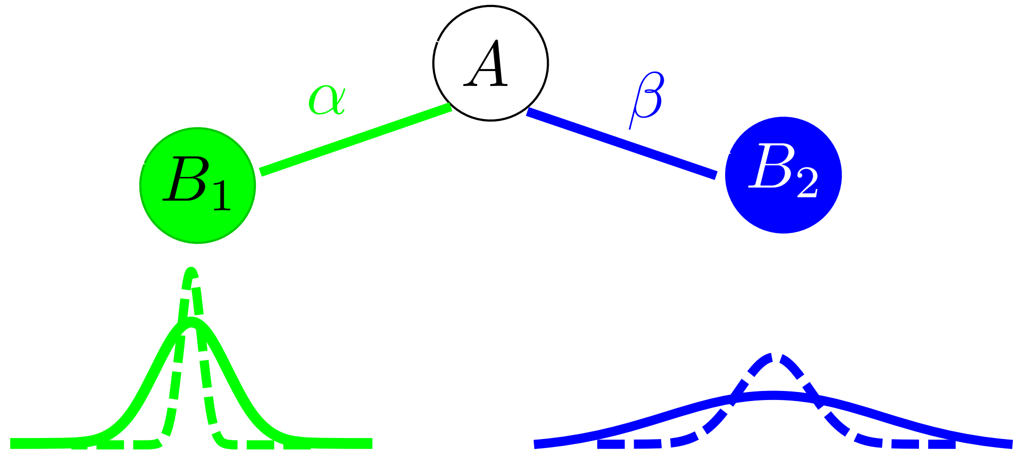

need 2 things:

1) current beliefs

2) estimate of how beliefs might change with search

Risk-based Lookahead Example

expand under α or β?

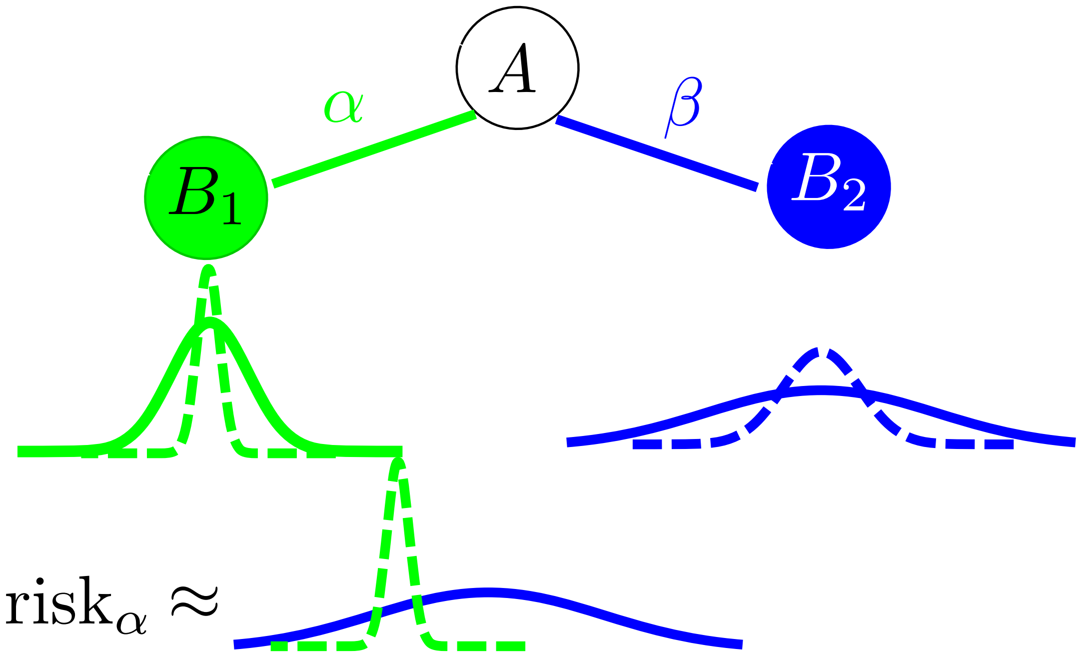

Risk: expected regret if a suboptimal action is selected

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
Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy
Other Research
Conclusions

expand under α or β?

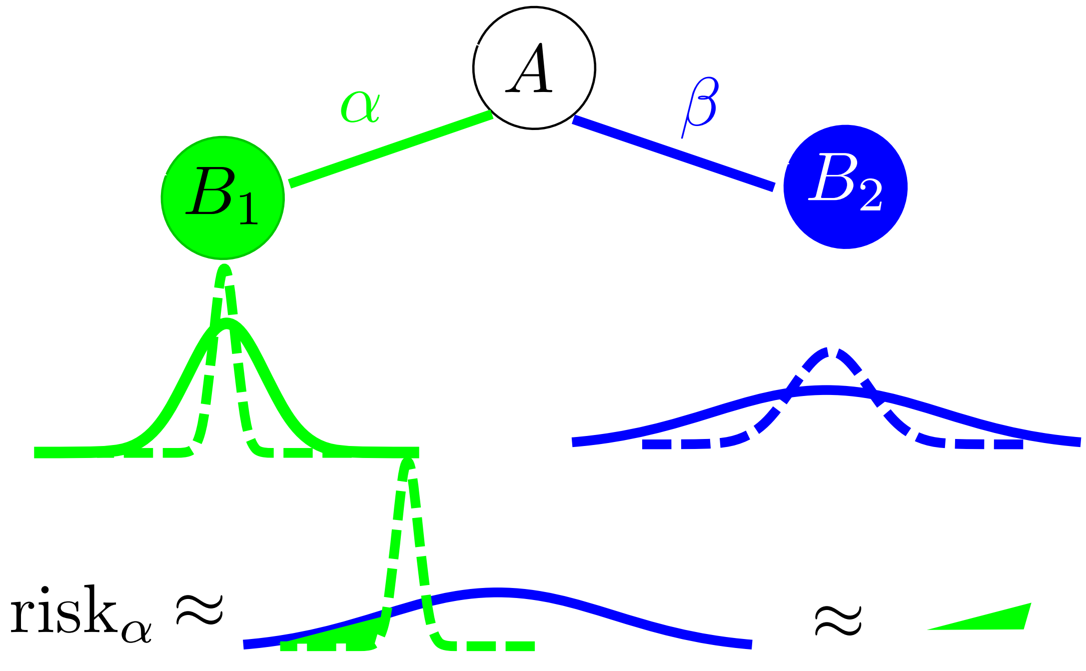

Risk: expected regret if a suboptimal action is selected

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

Nancy
- Decision-making
- Lookahead
- The Beliefs

Data-Driven Nancy

Other Research

Conclusions

expand under α or β?

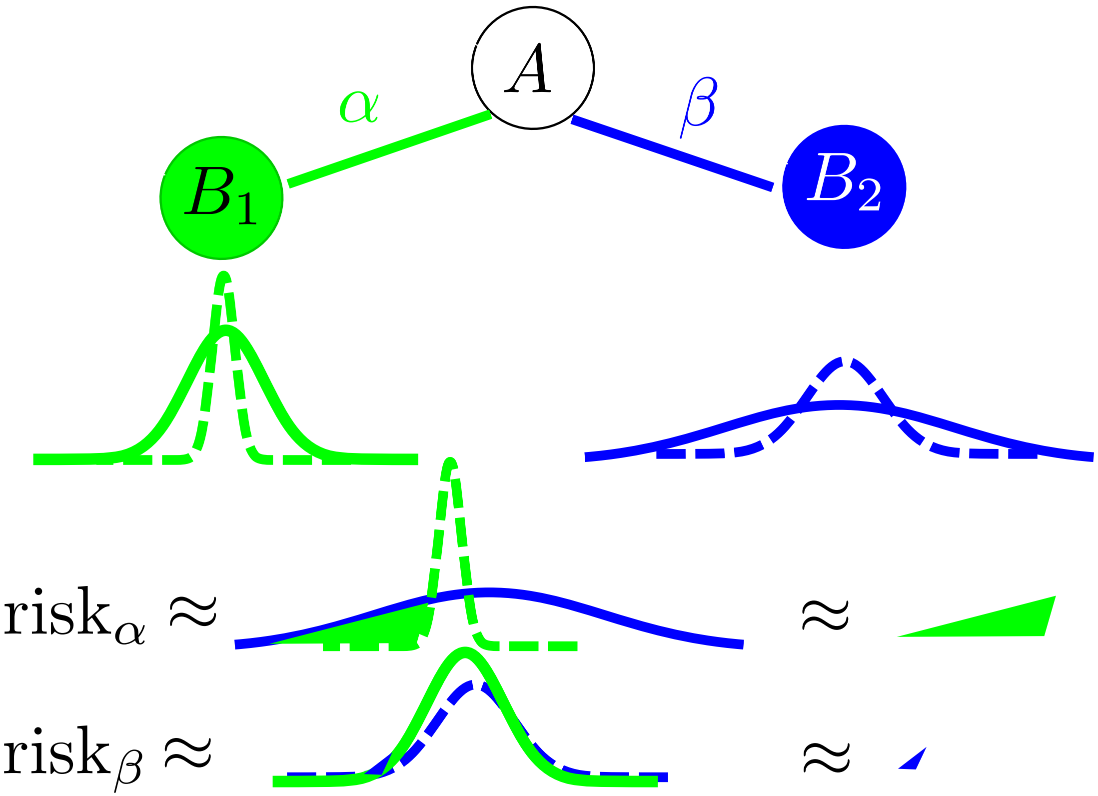

$\text{risk}_\alpha \approx$ $\approx$

$\text{risk}_\beta \approx$ $\approx$

expand under the TLA that minimizes risk!

expand under β!

Introduction
Nancy
Decision-making
Lookahead
The Beliefs
Data-Driven Nancy
Other Research
Conclusions

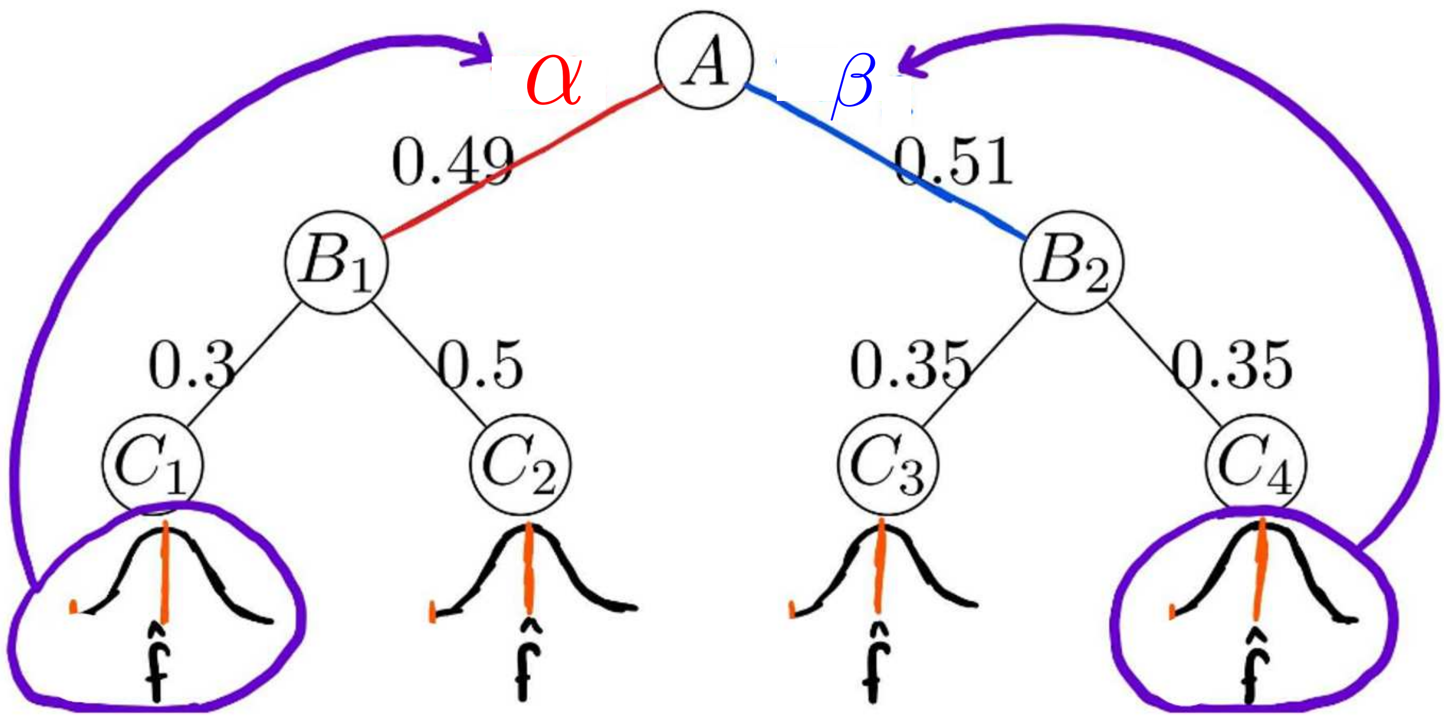

Nancy:

parent $\leftarrow$ belief with minimum $\hat{f}$ among successors

conveys an entire belief distribution

How to Form The Belief Distribution?

Introduction
Nancy
- Decision-making
- Lookahead
- The Beliefs
Data-Driven Nancy
Other Research
Conclusions

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

But where do beliefs come from?

How to Form The Belief Distribution?

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

But where do beliefs come from?

Nancy:

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

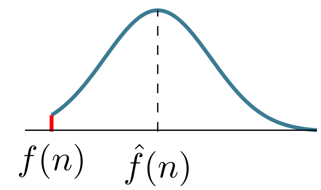

My work: Data-Driven Nancy[4]:

expressive histogram,
many parameters requires offline learning

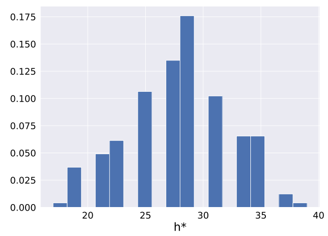

[4] Beliefs We Can Believe In: Replacing Assumptions with Data in Real-Time Search, Maximilian Fickert, Tianyi Gu, Leonhard Staut, Wheeler Ruml, Joerg Hoffmann, and Marek Petrik, AAAI, 2020.

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary
Other Research
Conclusions

Data-Driven Nancy

Learning a Model of Heuristic Error

Introduction

Nancy

Data-Driven Nancy

- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research

Conclusions

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

Gathering data:

- run offline suboptimal search on random problems
- collect all visited states
- for each observed h value:

  pick most common 200 states from the collection, compute h*

Example h* distribution: Sliding Puzzle

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research
Conclusions

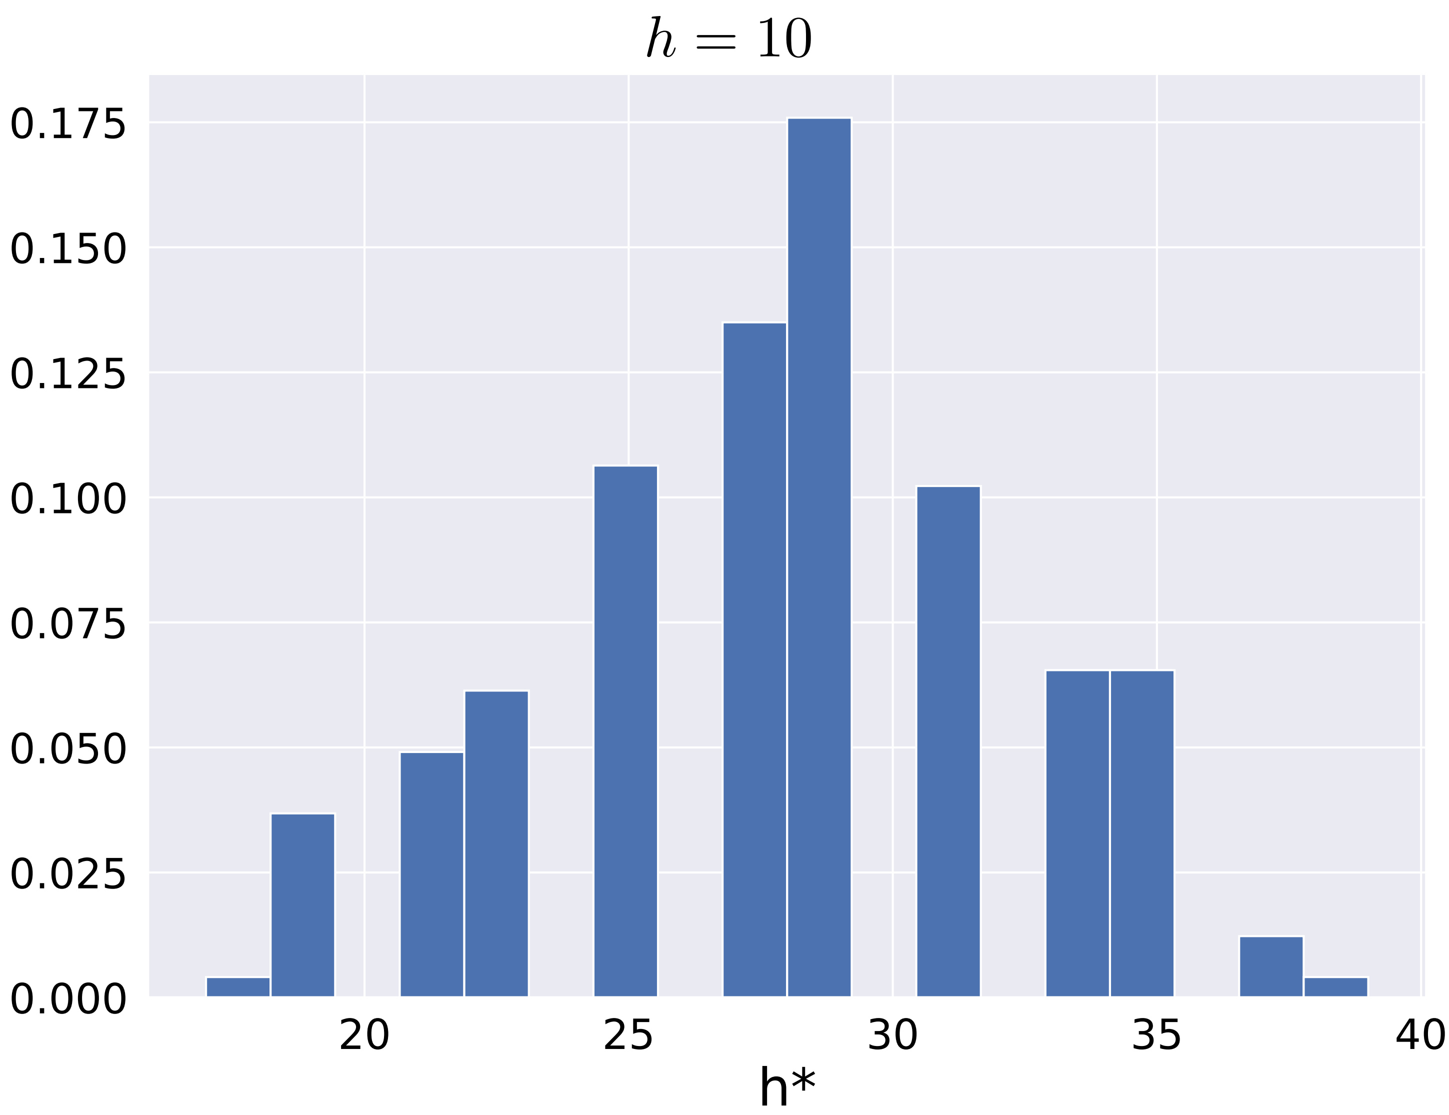

Example h^* distribution: Transport vs Blocks World

Introduction

Nancy

Data-Driven Nancy

- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research

Conclusions

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

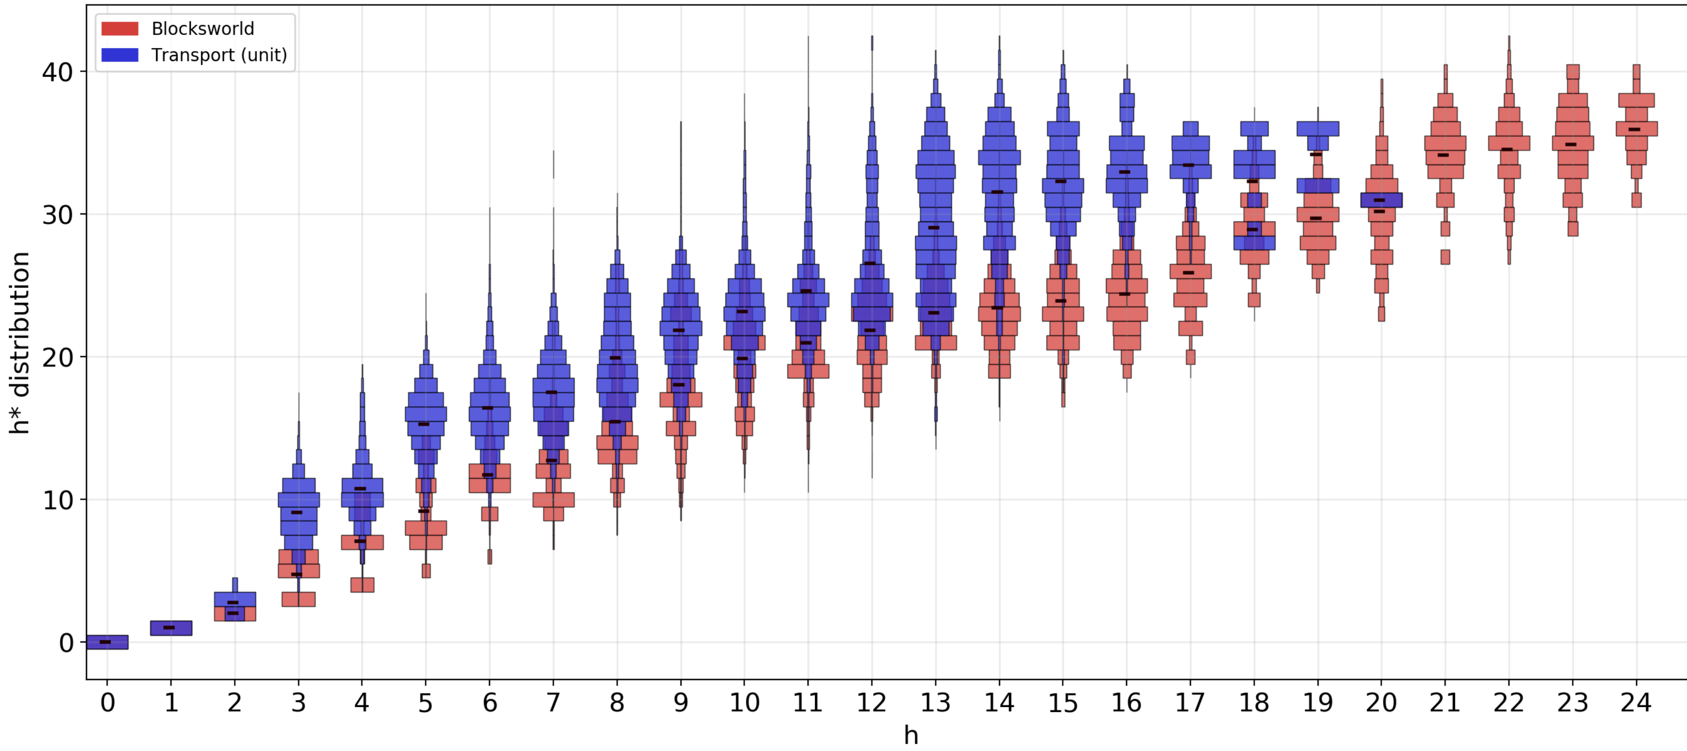

Beliefs are different from domain to domain

Completeness proof

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research
Conclusions

Lemma 1 *Under assumptions of goal-awareness and finite state space, if a real-time search algorithm is incomplete, it must have a circulating set* $S_\circ$*.*

Lemma 5 *Under our assumptions, a reasonable real-time search algorithm cannot have a circulating set.*

Theorem 1 *Under our assumptions, a reasonable real-time search algorithm will eventually reach a goal.*

Lemma 7 *Nancy is a reasonable real-time search algorithm.*

Lemma 8 *LSS-LRTA* is a reasonable real-time search algorithm.*

This proof applies to any LSS-LRTA*-style algorithm: no longer need heuristic consistency!

Mean Solution Cost on Planning Domains

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research
Conclusions

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**

Both version of Nancy outperform conventional approach!

Search Domains

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary

Other Research
Conclusions

sliding tile puzzle

    uniform, heavy ,inverse

pancake puzzle

    different size

racetrack

    reminiscent of autonomous driving

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary
Other Research
Conclusions

pancake puzzle

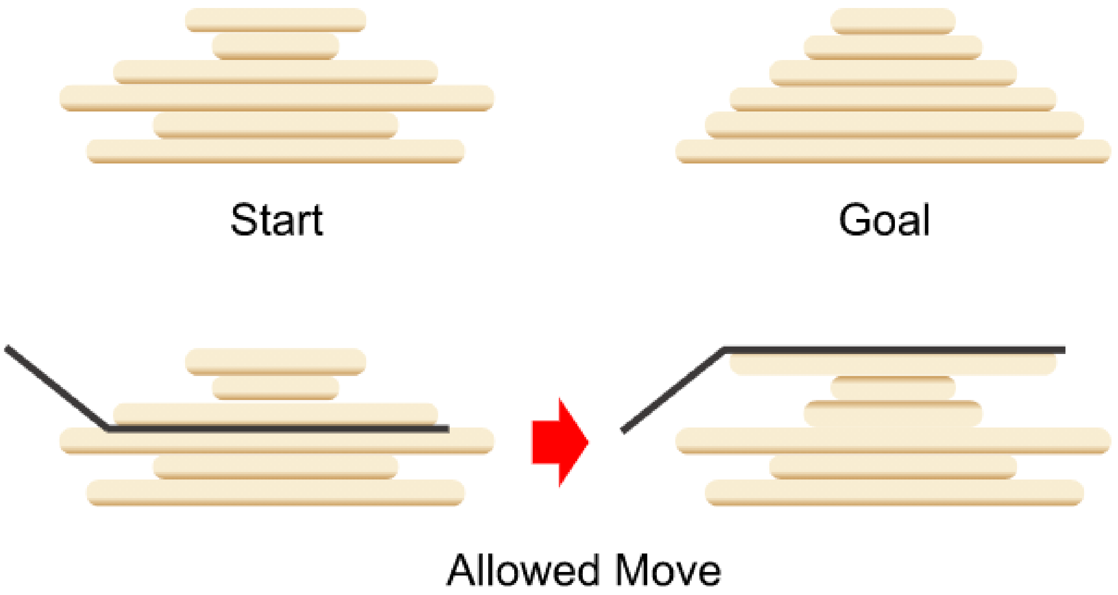

Comparison to IE and MCTS on Classic Search Domains

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary
Other Research
Conclusions

Nancy outperforms conventional approaches and MCTS[5]

[5] Real-time Planning as Data-driven Decision-making, Maximilian Fickert, Tianyi Gu, Leonhard Staut, Sai Lekyang, Wheeler Ruml, Joerg Hoffmann, and Marek Petrik, Bridging the Gap Between AI Planning and Reinforcement Learning (PRL), 2020.

Summary

Introduction
Nancy
Data-Driven Nancy
- Data
- Data Viz
- Completeness
- Planning
- Search
- Summary
Other Research
Conclusions

- Nancy starts to explore an optimal way of doing online heuristic search
- Nancy outperforms conventional LSS-LRTA* in cost and run time
- Data-driven approach increases robustness
- General completeness proof

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!

Introduction
Nancy
Data-Driven Nancy
Other Research
■ Bounded-Cost Search
■ Bounded-Suboptimal Search
■ Motion Planning
■ Robotics
Conclusions

Other Research

Introduction

Nancy

Data-Driven Nancy

Other Research

- Bounded-Cost Search
- Bounded-Suboptimal Search
- Motion Planning
- Robotics

Conclusions

distributional methods can also benefit other types of search

bounded-cost search: problem, cost bound $\rightarrow$ find a solution within bound as quickly as possible

Our Approach: Expected Effort Search (XES) [6]

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

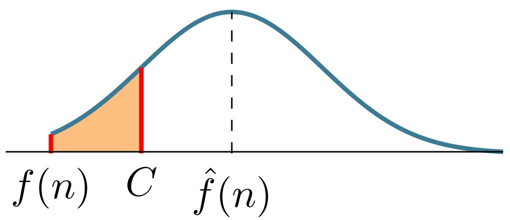

2. estimate total search effort by $d(n)$
3. best first search on expected search effort $d(n)/p(n)$

[6] Bounded-cost Search Using Estimates of Uncertainty, Maximilian Fickert, Tianyi Gu, and Wheeler Ruml, IJCAI, 2021.

Other Research: 1/3 Bounded-Cost Search

Introduction
Nancy
Data-Driven Nancy
Other Research
■ Bounded-Cost Search
■ Bounded-Suboptimal Search
■ Motion Planning
■ Robotics
Conclusions

distributional methods can also benefit other types of search

bounded-cost search: problem, cost bound → find a solution within bound as quickly as possible

Our Approach: Expected Effort Search (XES) [6]

previous algorithms are brittle

XES is now new state-of-the-art!

[6] Bounded-cost Search Using Estimates of Uncertainty, Maximilian Fickert, Tianyi Gu, and Wheeler Ruml, IJCAI, 2021.

Other Research: 2/3 Bounded-Suboptimal Search

Introduction

Nancy

Data-Driven Nancy

Other Research
- Bounded-Cost Search
- Bounded-Suboptimal Search
- Motion Planning
- Robotics

Conclusions

distributional methods can also benefit other types of search

bounded-suboptimal search: problem, suboptimal bound $\rightarrow$ find a solution within bound as quickly as possible

Our Approach: Dynamic Expected Effort Search (DXES)

estimates on cost and bound!

Introduction

Nancy

Data-Driven Nancy

Other Research
- Bounded-Cost Search
- Bounded-Suboptimal Search
- Motion Planning
- Robotics

Conclusions

distributional methods can also benefit motion planning

Motion Planning: find collision-free trajectory for robot

Effort-guided planning: Bayesian Effort-Aided Search Tree[7]

- abstract graph
- edge = binomial distribution of online estimate on planning effort

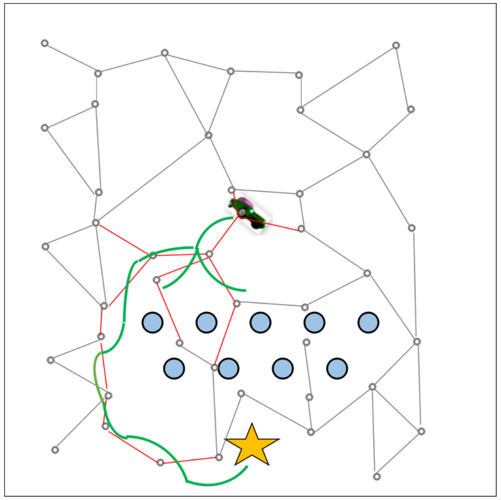

Estimate effort → Guide motion tree growth toward easy way

[7] An Effort Bias for Sampling-based Motion Planning, Scott Kiesel, Tianyi Gu, and Wheeler Ruml, IROS, 2017.

Other Research: 3/3 Motion Planning

Introduction
Nancy
Data-Driven Nancy
Other Research
■ Bounded-Cost Search
■ Bounded-Suboptimal Search
■ Motion Planning
■ Robotics
Conclusions

distributional methods can also benefit motion planning

Motion Planning: find collision-free trajectory for robot

Effort-guided planning: Bayesian Effort-Aided Search Tree[7]

- abstract graph
- edge = binomial distribution of online estimate on planning effort

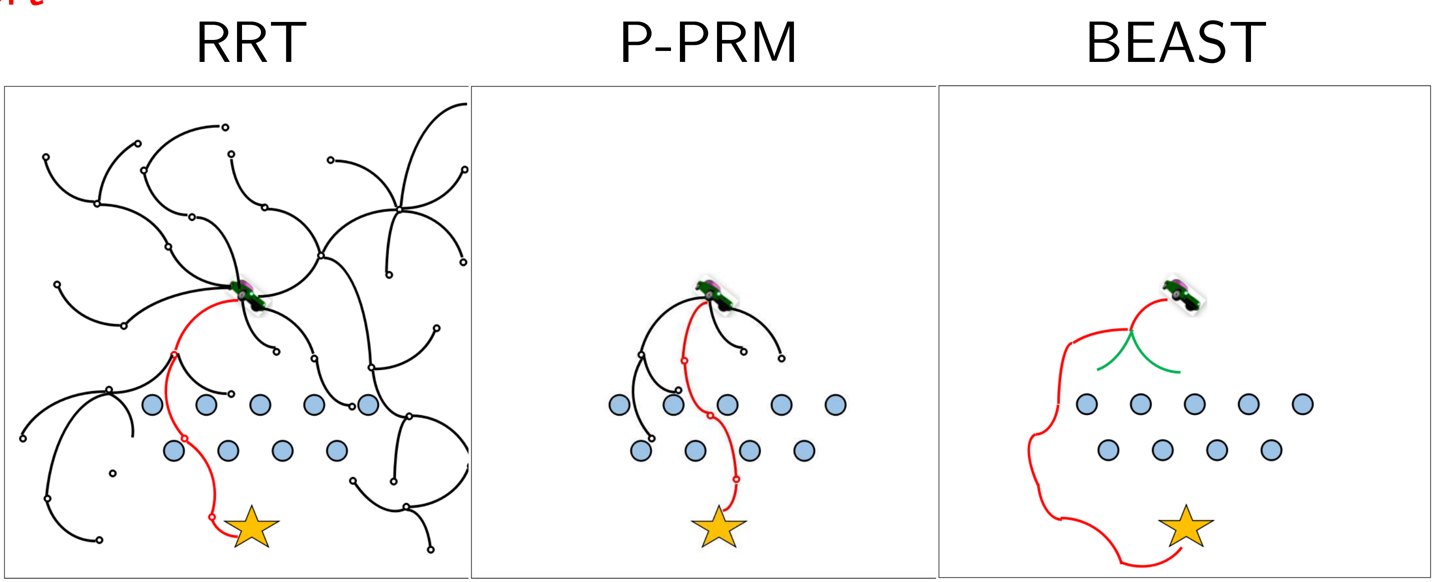

Beast find solution faster than P-PRM and RRT

[7] An Effort Bias for Sampling-based Motion Planning, Scott Kiesel, Tianyi Gu, and Wheeler Ruml, IROS, 2017.

Other Research: Dementia-care Robot

Introduction
Nancy
Data-Driven Nancy
Other Research
- Bounded-Cost Search
- Bounded-Suboptimal Search
- Motion Planning
- Robotics

Conclusions

Can robot and current available technologies help?

proof-of-concept demo[8] [9] : smart home & lay-user friendly robot

[8] An Adaptive Software Framework for Dementia-care Robots, T Gu, M Begum, N Zhang, D Xu, S Arthanat, and D LaRoche, PlanRob, 2020.

[9] Caregiver Perspectives on A Smart Home-based Socially Assistive Robot for Individuals with Alzheimer's Disease and Related Dementia, S Arthanat, M Begum, T Gu, N Zhang, D Xu, and D LaRoche, Disability and Rehabilitation: Assistive Technology, 2020.

Other Research: Dementia-care Robot

Can robot and current available technologies help?

proof-of-concept demo[8] [9] : smart home & lay-user friendly robot

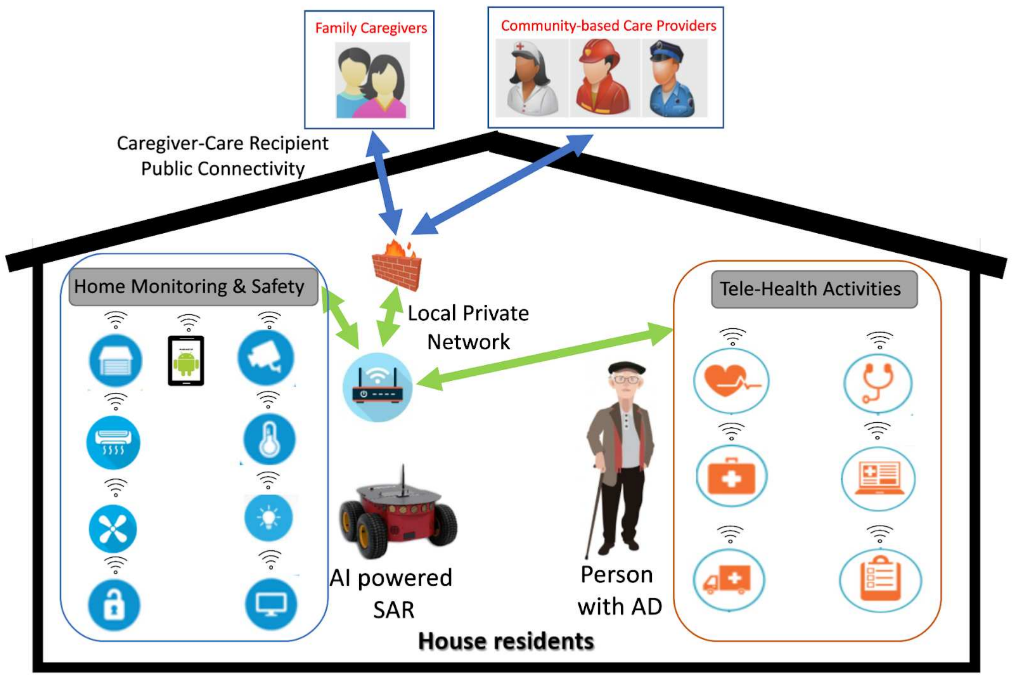

[8] An Adaptive Software Framework for Dementia-care Robots, T Gu, M Begum, N Zhang, D Xu, S Arthanat, and D LaRoche, PlanRob, 2020.

[9] Caregiver Perspectives on A Smart Home-based Socially Assistive Robot for Individuals with Alzheimer's Disease and Related Dementia, S Arthanat, M Begum, T Gu, N Zhang, D Xu, and D LaRoche, Disability and Rehabilitation: Assistive Technology, 2020.

Other Research: Dementia-care Robot

Can robot and current available technologies help?

proof-of-concept demo[8] [9] : smart home & lay-user friendly robot

Questionnaire 3: Programming an Alerting Protocol

To Prevent Wandering

This form is to demonstrate how you can set up an alerting protocol for the robot to prevent your family member from wandering outside.

Please fill in the information below

To prevent your family member from stepping out

1. What time duration should your family member not go out?

   From: _____ To: _____

2. Who is the person I should call if your family member does not come back after the reminder?

   Name: ___________

   Phone: ___________

3. Should the robot call emergency personnel too?
   - ☐ Yes
   - ☐ No

4. If yes, how soon after the family member does not come back? ______minutes

5. If your family member is not back, what is the likely place the emergency personnel need to look for?

6. Is there anyone else you want the robot to call? What is the phone number?

   Name: ____________________

   Phone: ____________________

Questionnaire 2: Programming a Reminder Protocol

Medication Intake

This form is to demonstrate how you can set up a reminder protocol for the robot to help manage your family member's medication.

Please fill in the information below

For medication intake

1. What time do you want your family member to take his or her medications?

   ____

2. Where is the medication bottle kept? e.g. kitchen table

   ____

3. Will the medication bottle get moved from where it is kept usually?
   - ☐ Yes
   - ☐ No

4. What should the robot do if your family member cannot find the medication?

   ☐ Locate the medication in the house and
   - ☐ Remind your family member or
   - ☐ Call you

   OR

   ☐ Call you

5. How many times you want the robot to remind your family member before calling you and asking you to communicate with the family member?

   _____ times every _____ minutes

---

[8] An Adaptive Software Framework for Dementia-care Robots, T Gu, M Begum, N Zhang, D Xu, S Arthanat, and D LaRoche, PlanRob, 2020.

[9] Caregiver Perspectives on A Smart Home-based Socially Assistive Robot for Individuals with Alzheimer's Disease and Related Dementia, S Arthanat, M Begum, T Gu, N Zhang, D Xu, and D LaRoche, Disability and Rehabilitation: Assistive Technology, 2020.

Introduction
Nancy
Data-Driven Nancy
Other Research
■ Bounded-Cost Search
■ Bounded-Suboptimal Search
■ Motion Planning
■ Robotics
Conclusions

Can robot and current available technologies help?

proof-of-concept demo[8] [9] : smart home & lay-user friendly robot

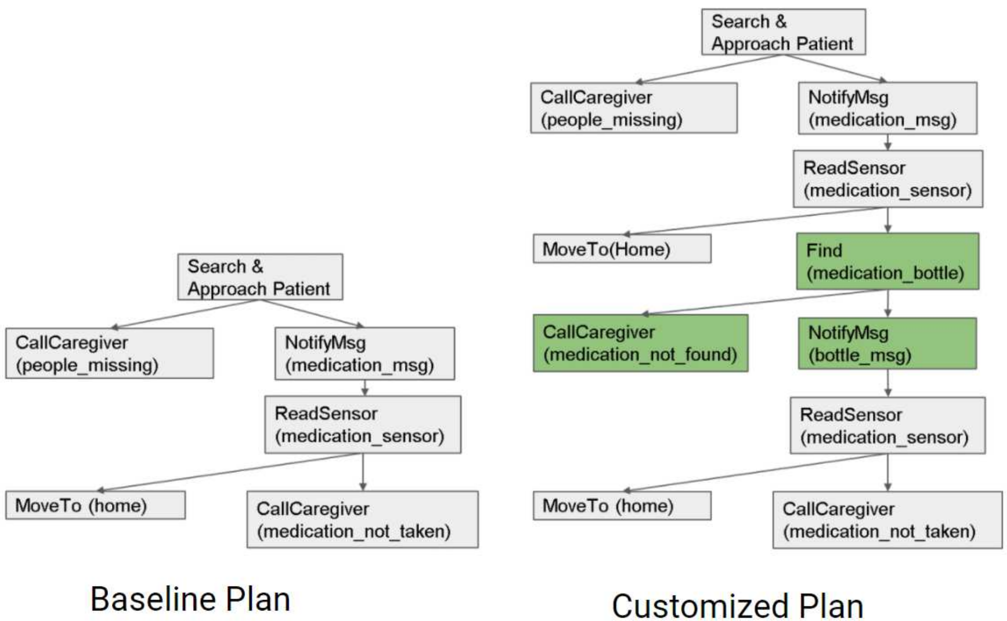

[8] An Adaptive Software Framework for Dementia-care Robots, T Gu, M Begum, N Zhang, D Xu, S Arthanat, and D LaRoche, PlanRob, 2020.

[9] Caregiver Perspectives on A Smart Home-based Socially Assistive Robot for Individuals with Alzheimer's Disease and Related Dementia, S Arthanat, M Begum, T Gu, N Zhang, D Xu, and D LaRoche, Disability and Rehabilitation: Assistive Technology, 2020.

Introduction
Nancy
Data-Driven Nancy
Other Research
Conclusions
■ Summary

Conclusions

Summary

Introduction
Nancy
Data-Driven Nancy
Other Research
Conclusions
Summary

Examples of using distribution to guide search:

- real-time planning: Nancy, DDNancy
- suboptimal search: XES, DXES
- robotics: BEAST, dementia-care robot

Summary

Introduction
Nancy
Data-Driven Nancy
Other Research
Conclusions
Summary

Examples of using distribution to guide search:

- real-time planning: Nancy, DDNancy
- suboptimal search: XES, DXES
- robotics: BEAST, dementia-care robot

Exciting time in AI!

- Planning, RL, ML, Robotics

Summary

Introduction
Nancy
Data-Driven Nancy
Other Research
Conclusions
Summary

Examples of using distribution to guide search:
- real-time planning: Nancy, DDNancy
- suboptimal search: XES, DXES
- robotics: BEAST, dementia-care robot

Exciting time in AI!
- Planning, RL, ML, Robotics

Much work needs to be done!
- data-driven + planning
- statistics + model-based approach

Questions?

Introduction
Nancy
Data-Driven Nancy
Other Research
Conclusions
Questions
- Questions?