

Search Algorithms as Agents
— or —
There's More to Life Than $h(n)$

Wheeler Ruml
and the UNH AI Group, esp. Jordan Thayer



UNIVERSITY *of* NEW HAMPSHIRE

(thanks to the NSF RI and DARPA CSSG programs for support)

The AI Vision

Heuristic Search

■ The AI Vision

■ Alg as Agent

■ What is Search?

■ Related Work

■ Problem Settings


Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion



Search Algorithms as Agents

Heuristic Search

■ The AI Vision

■ Alg as Agent

■ What is Search?

■ Related Work

■ Problem Settings


Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion



What is Search?

Heuristic Search

■ The AI Vision

■ Alg as Agent

■ **What is Search?**

■ Related Work

■ Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

acting under uncertainty to maximize utility

What is Search?

Heuristic Search

■ The AI Vision

■ Alg as Agent

■ **What is Search?**

■ Related Work

■ Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

acting under uncertainty to maximize utility
= **all of AI**

What is Search?

Heuristic Search

■ The AI Vision

■ Alg as Agent

■ **What is Search?**

■ Related Work

■ Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

acting under uncertainty to maximize utility
= **all of AI**

possible sources of information, 'sensors'
how to exploit information, 'aggregation', 'filtering'

Related Work

Heuristic Search

- The AI Vision
- Alg as Agent
- What is Search?
- Related Work
- Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

Rational search:

- Othar Hansson
- Andy Mayer

Metareasoning, 'bounded optimality':

- Eric Wefald
- Shlomo Zilberstein
- Eric Horvitz
- Stuart Russell

Anytime search:

- Tom Dean
- Mark Boddy

Problem Settings

Heuristic Search

- The AI Vision
- Alg as Agent
- What is Search?
- Related Work
- Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

optimal: minimize solution cost
must expand all with $f(n) < f^*(opt)$

Problem Settings

Heuristic Search

- The AI Vision
- Alg as Agent
- What is Search?
- Related Work
- Problem Settings

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

optimal: minimize solution cost
must expand all with $f(n) < f^*(opt)$

greedy: minimize solving time

bounded suboptimal: minimize time subject to relative cost bound (factor of optimal)

bounded cost: minimize time subject to absolute cost bound

contract: minimize cost subject to absolute time bound

utility function: maximize utility function of cost and time
eg, goal achievement time =
plan makespan + search time

Heuristic Search

Greedy Search

- Direct Approach
- Distance-to-Go
- Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

Greedy Search

A Direct Approach

how to minimize solving time?

Heuristic Search

Greedy Search

■ **Direct Approach**

■ Distance-to-Go

■ Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

A Direct Approach

Heuristic Search

Greedy Search

Direct Approach

Distance-to-Go

Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

how to minimize solving time?

how to minimize number of expansions?

A Direct Approach

Heuristic Search

Greedy Search

Direct Approach

Distance-to-Go

Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

how to minimize solving time?

how to minimize number of expansions?

take the shortest path to a goal

A Direct Approach

Heuristic Search

Greedy Search

Direct Approach

Distance-to-Go

Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

how to minimize solving time?

how to minimize number of expansions?

take the shortest path to a goal

for domains with costs, this is **not** $h(n)$

A Direct Approach

Heuristic Search

Greedy Search

■ **Direct Approach**

■ Distance-to-Go

■ Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion

how to minimize solving time?

how to minimize number of expansions?

take the shortest path to a goal

for domains with costs, this is **not** $h(n)$

source #1 of 5: distance-to-go

Distance-to-Go Estimates

Heuristic Search

Greedy Search

■ Direct Approach

■ Distance-to-Go


■ Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion



Distance-to-Go Estimates

Heuristic Search

Greedy Search

■ Direct Approach

■ Distance-to-Go


■ Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion



$$d_{cheapest}^*(n) = 3$$

$$d_{nearest}^*(n) = 2$$

$d_{nearest}$ is potentially independent of h ,
can be computed as h with unit costs

$d_{cheapest}$ often computable alongside h

d can be inadmissible

Distance-to-Go Estimates

Heuristic Search

Greedy Search

■ Direct Approach

■ Distance-to-Go


■ Speedy Results

Bounded Search

Contract Search

Utility Functions

Conclusion



$$d_{cheapest}^*(n) = 3$$

$$d_{nearest}^*(n) = 2$$

$d_{nearest}$ is potentially independent of h ,
can be computed as h with unit costs

$d_{cheapest}$ often computable alongside h

d can be inadmissible

Speedy Search: best-first search on $d(n)$ (Thayer et al, SoCS-09)

Speedy Results

Heuristic Search

Greedy Search

■ Direct Approach

■ Distance-to-Go

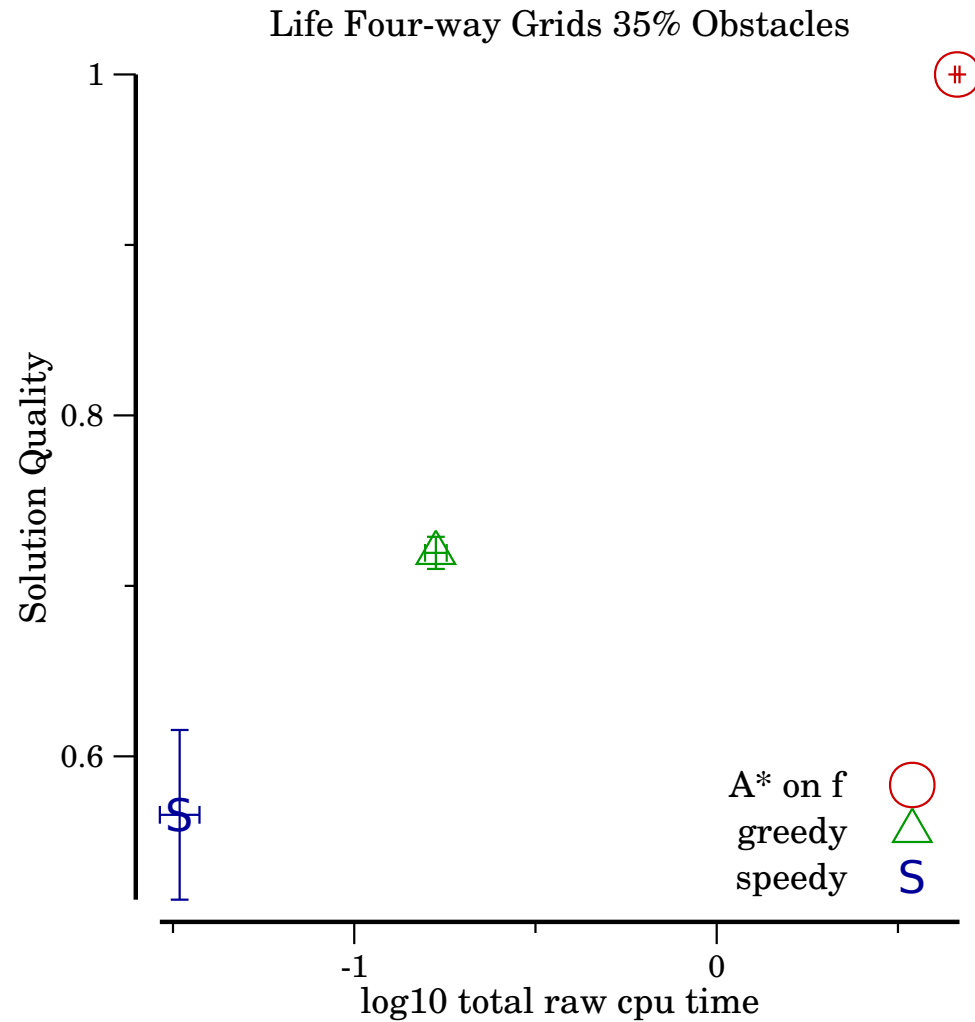
■ Speedy Results

Bounded Search

Contract Search

Utility Functions

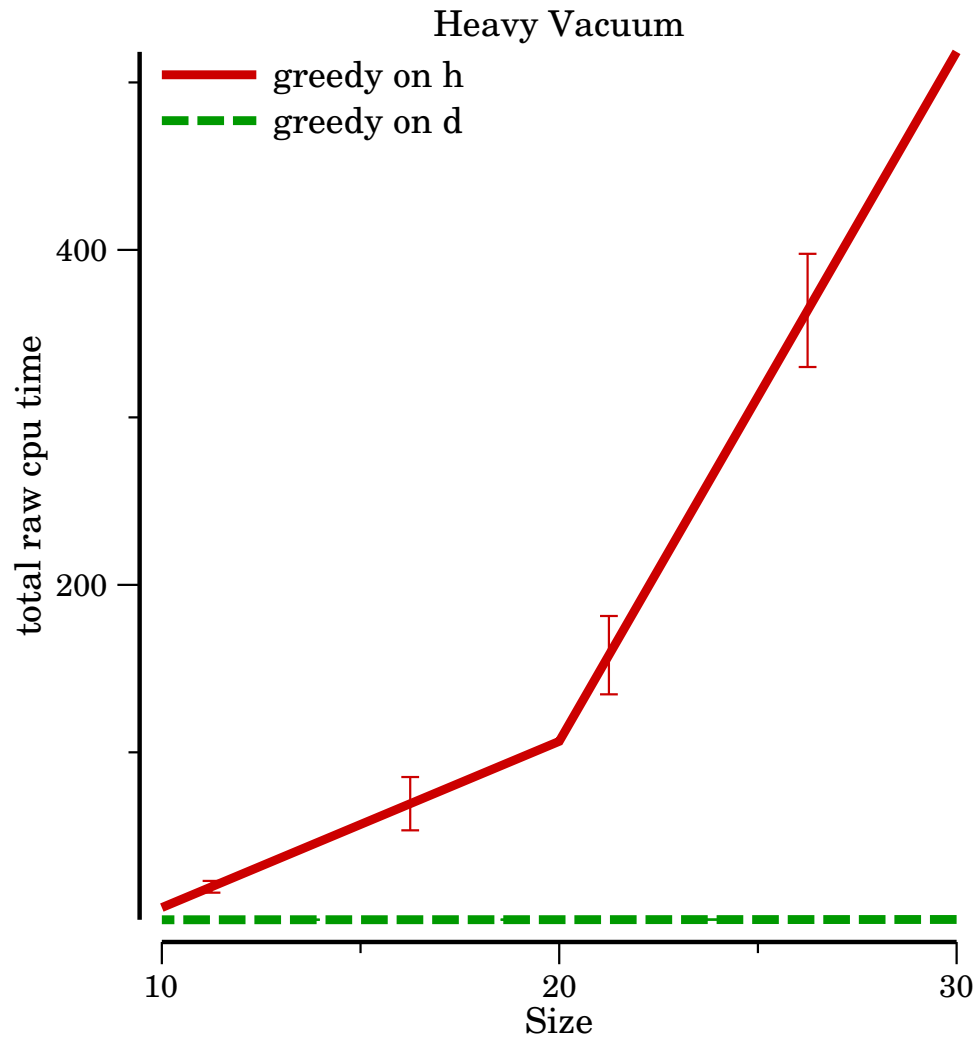
Conclusion



speedy finds solutions faster

Speedy Results

- Heuristic Search
- Greedy Search
- Direct Approach
- Distance-to-Go
- Speedy Results**
- Bounded Search
- Contract Search
- Utility Functions
- Conclusion



speedy scales better

Heuristic Search

Greedy Search

Bounded Search

- Direct Approach
- A_ϵ^*
- A_ϵ^* 's Flaw
- Estimated Cost
- Learning Cost
- EES
- EES Order
- EES Results
- The Story So Far

Contract Search

Utility Functions

Conclusion

Bounded-Suboptimal Search

A Direct Approach

minimize solving time subject to relative cost bound (factor of optimal)

Heuristic Search

Greedy Search

Bounded Search

■ **Direct Approach**

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

A Direct Approach

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

minimize solving time subject to relative cost bound (factor of optimal)

weighted A^* ($f'(n) = g(n) + w \cdot h(n)$) is **simple but ad hoc**

A Direct Approach

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

minimize solving time subject to relative cost bound (factor of optimal)

weighted A^* ($f'(n) = g(n) + w \cdot h(n)$) is **simple but ad hoc**

expand the node closest to a solution within the bound

A Direct Approach

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

minimize solving time subject to relative cost bound (factor of optimal)

weighted A^* ($f'(n) = g(n) + w \cdot h(n)$) is **simple but ad hoc**

expand the node closest to a solution within the bound known to not work!

A_ϵ^* (Pearl and Kim, IEEE PAMI 1982)

intuition: of all solutions within the bound, the nearest should be the fastest to find

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

A_ε* (Pearl and Kim, IEEE PAMI 1982)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ε*

■ A_ε*'s Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

intuition: of all solutions within the bound, the nearest should be the fastest to find

$$f(n) = g(n) + h(n)$$

best_f: open node with minimum f

two lists:

open: as usual, sorted on $f(n)$

focal: subset of open with $f(n) \leq w \cdot f(\text{best}_f)$, sorted on $d(n)$

A_ε*: best-first search using *focal*

A_ε* (Pearl and Kim, IEEE PAMI 1982)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ε*

■ A_ε*'s Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

intuition: of all solutions within the bound, the nearest should be the fastest to find

$$f(n) = g(n) + h(n)$$

best_f: open node with minimum f

two lists:

open: as usual, sorted on $f(n)$

focal: subset of open with $f(n) \leq w \cdot f(\text{best}_f)$, sorted on $d(n)$

A_ε*: best-first search using *focal*

Why doesn't it work?

A_ϵ^* 's Flaw (Thayer et al, SoCS-09)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far


Contract Search

Utility Functions

Conclusion

open: as usual, sorted on $f(n)$

focal: subset of open with $f(n) \leq w \cdot f(best_f)$, sorted on $d(n)$



A_ϵ^* 's Flaw (Thayer et al, SoCS-09)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

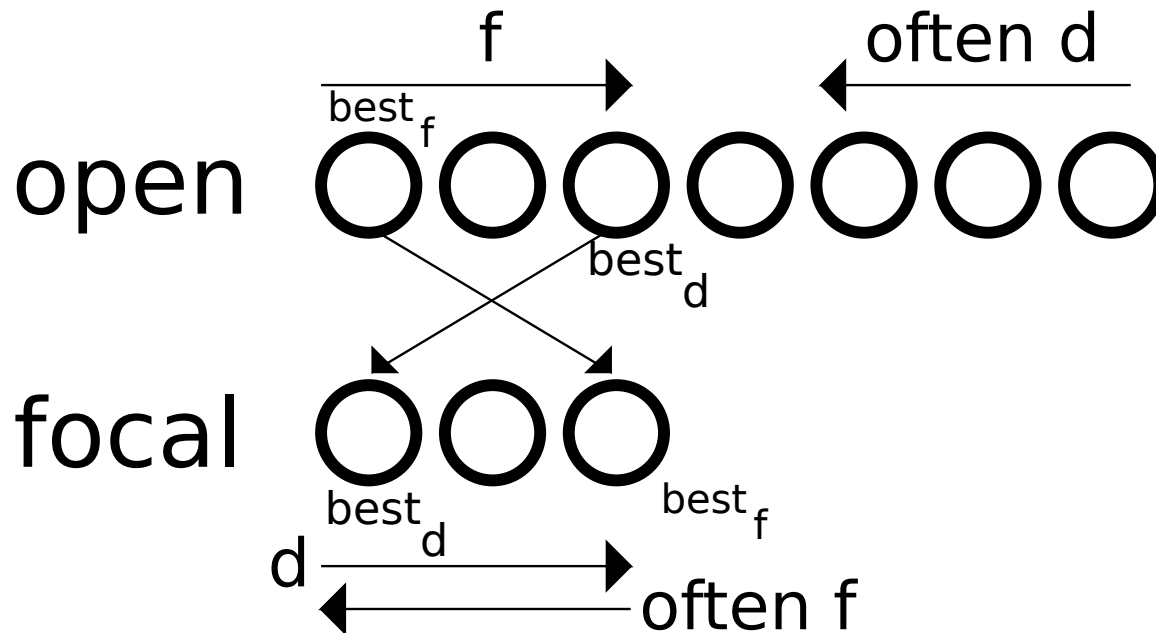
Contract Search

Utility Functions

Conclusion

open: as usual, sorted on $f(n)$

focal: subset of open with $f(n) \leq w \cdot f(best_f)$, sorted on $d(n)$



A_ϵ^* 's Flaw (Thayer et al, SoCS-09)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

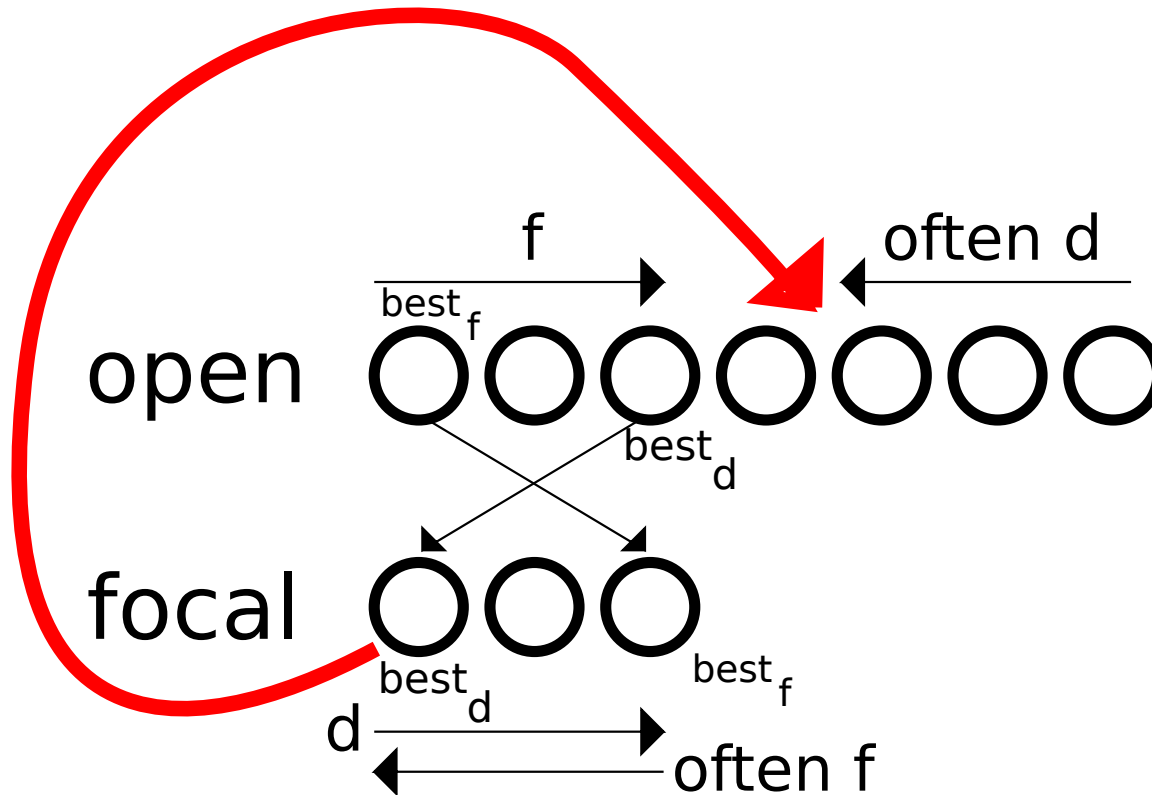
Contract Search

Utility Functions

Conclusion

open: as usual, sorted on $f(n)$

focal: subset of open with $f(n) \leq w \cdot f(best_f)$, sorted on $d(n)$



f rises as search progresses (h is admissible)

$best_d$'s children don't qualify for focal

Estimated Cost

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_{ϵ}^*

■ A_{ϵ}^* 's Flaw

■ **Estimated Cost**

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

intuition: pursuing the shortest solution within the bound
should be fast

intuition': an unbiased estimates of cost won't always rise

Estimated Cost

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ **Estimated Cost**

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion


intuition: pursuing the shortest solution within the bound
should be fast

intuition': an unbiased estimates of cost won't always rise

source #2 of 5: unbiased cost estimates $\hat{f}(n)$

Learning Unbiased Cost Estimates

Every expansion gives evidence for heuristic's error!



$$f^*(p) = f^*(bc)$$

$$h^*(p) = h^*(bc) + c(p, bc)$$

$$\epsilon_h = (h(bc) + c(p, bc)) - h(p)$$

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far


Contract Search

Utility Functions

Conclusion

Learning Unbiased Cost Estimates

Every expansion gives evidence for heuristic's error!



$$f^*(p) = f^*(bc)$$

$$h^*(p) = h^*(bc) + c(p, bc)$$

$$\epsilon_h = (h(bc) + c(p, bc)) - h(p)$$

$$\hat{h}(n) = h(n) + \overline{\epsilon_h} \cdot d(n)$$

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far


Contract Search

Utility Functions

Conclusion

Learning Unbiased Cost Estimates

Every expansion gives evidence for heuristic's error!



$$f^*(p) = f^*(bc)$$

$$h^*(p) = h^*(bc) + c(p, bc)$$

$$\epsilon_h = (h(bc) + c(p, bc)) - h(p)$$

$$\hat{h}(n) = h(n) + \bar{\epsilon}_h \cdot d(n)$$

can do this for $d(n)$ too...

$$\hat{h}(n) = h(n) + \bar{\epsilon}_h \cdot \hat{d}(n)$$

see Thayer et al (ICAPS-11) for details

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

Explicit Estimation Search (EES, IJCAI-11)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

$best_f$: open node with minimum f

$best_{\hat{f}}$: open node with minimum \hat{f}

Explicit Estimation Search (EES, IJCAI-11)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

$best_f$: open node with minimum f

$best_{\hat{f}}$: open node with minimum \hat{f}

three lists (!):

$open$: as usual, but sorted on $\hat{f}(n)$

$focal$: subset of open with $\hat{f}(n) \leq w \cdot \hat{f}(best_{\hat{f}})$, sorted on $\hat{d}(n)$

$cleanup$: same as $open$, but sorted on $f(n)$

Explicit Estimation Search (EES, IJCAI-11)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

$best_f$: open node with minimum f

$best_{\hat{f}}$: open node with minimum \hat{f}

three lists (!):

$open$: as usual, but sorted on $\hat{f}(n)$

$focal$: subset of open with $\hat{f}(n) \leq w \cdot \hat{f}(best_{\hat{f}})$, sorted on $\hat{d}(n)$

$cleanup$: same as $open$, but sorted on $f(n)$

$best_{\hat{d}}$: first node on $focal$

Explicit Estimation Search (EES, IJCAI-11)

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

$best_f$: open node with minimum f

$best_{\hat{f}}$: open node with minimum \hat{f}

three lists (!):

$open$: as usual, but sorted on $\hat{f}(n)$

$focal$: subset of open with $\hat{f}(n) \leq w \cdot \hat{f}(best_{\hat{f}})$, sorted on $\hat{d}(n)$

$cleanup$: same as $open$, but sorted on $f(n)$

$best_{\hat{d}}$: first node on $focal$

estimated w -admissible node with minimum \hat{d}

EES Expansion Order

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ **EES Order**

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

$best_f$: open node with minimum f

$best_{\hat{f}}$: open node with minimum \hat{f}

$best_{\hat{d}}$: estimated w -admissible node with minimum \hat{d}

node to expand next:

1. pursue the shortest solution within the bound
2. pursue the estimated cheapest solution
3. raise the lower bound on optimal cost

in other words:

1. **if** $\hat{f}(best_{\hat{d}}) \leq w \cdot f(best_f)$ **then** $best_{\hat{d}}$
2. **else if** $\hat{f}(best_{\hat{f}}) \leq w \cdot f(best_f)$ **then** $best_{\hat{f}}$
3. **else** $best_f$

find a solution provably within the bound as quickly as possible

EES Results

Heuristic Search

Greedy Search

Bounded Search

Direct Approach

A^*_ϵ

A^*_ϵ 's Flaw

Estimated Cost

Learning Cost

EES

EES Order

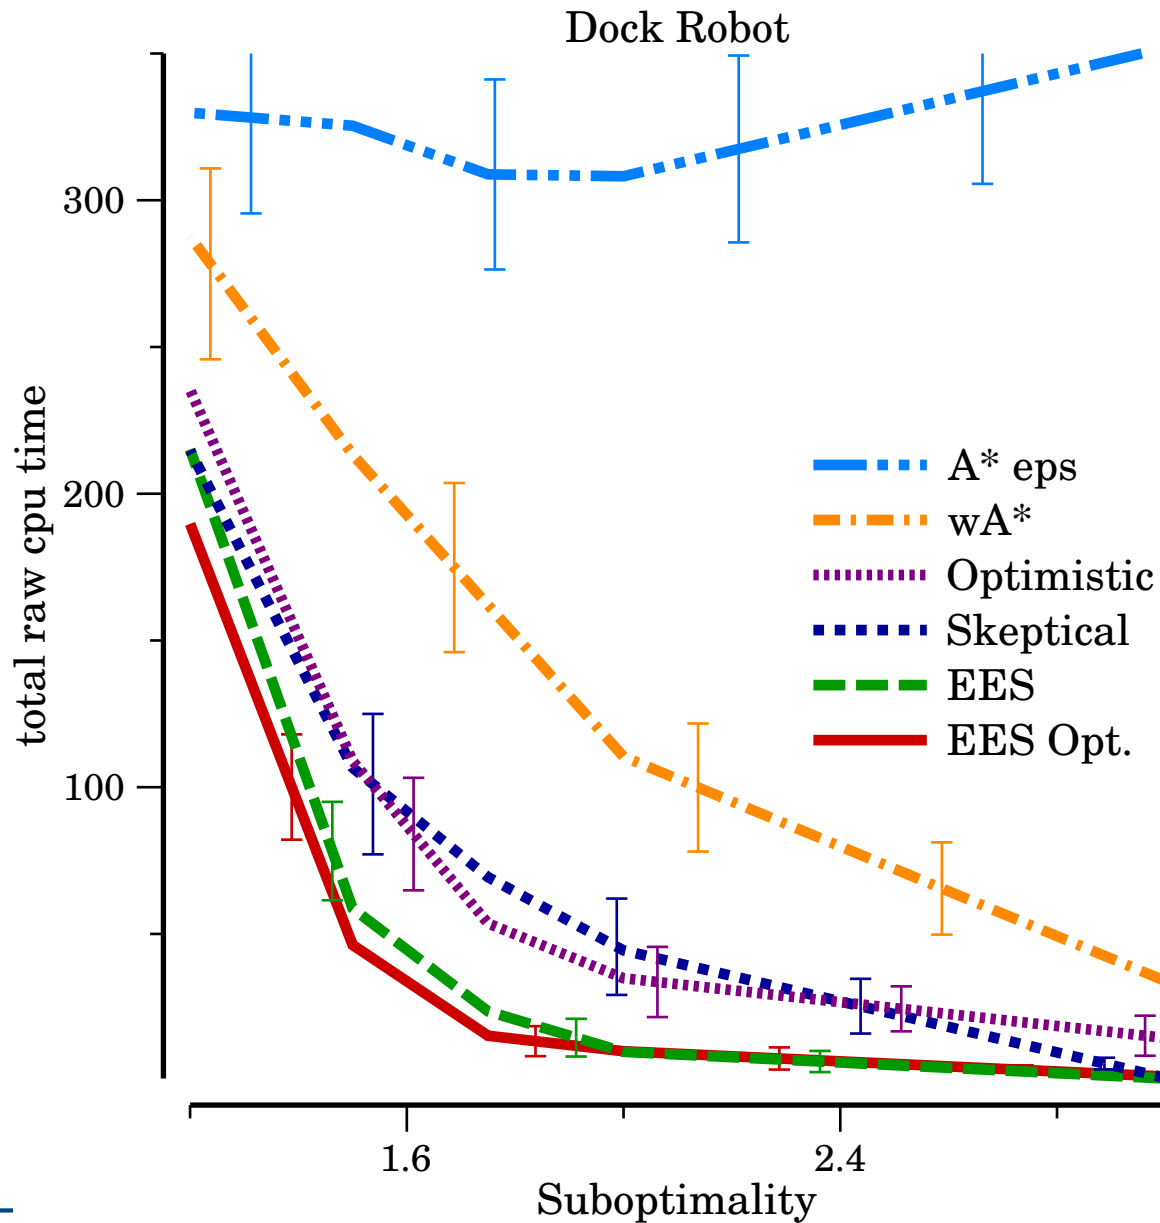
EES Results

The Story So Far

Contract Search

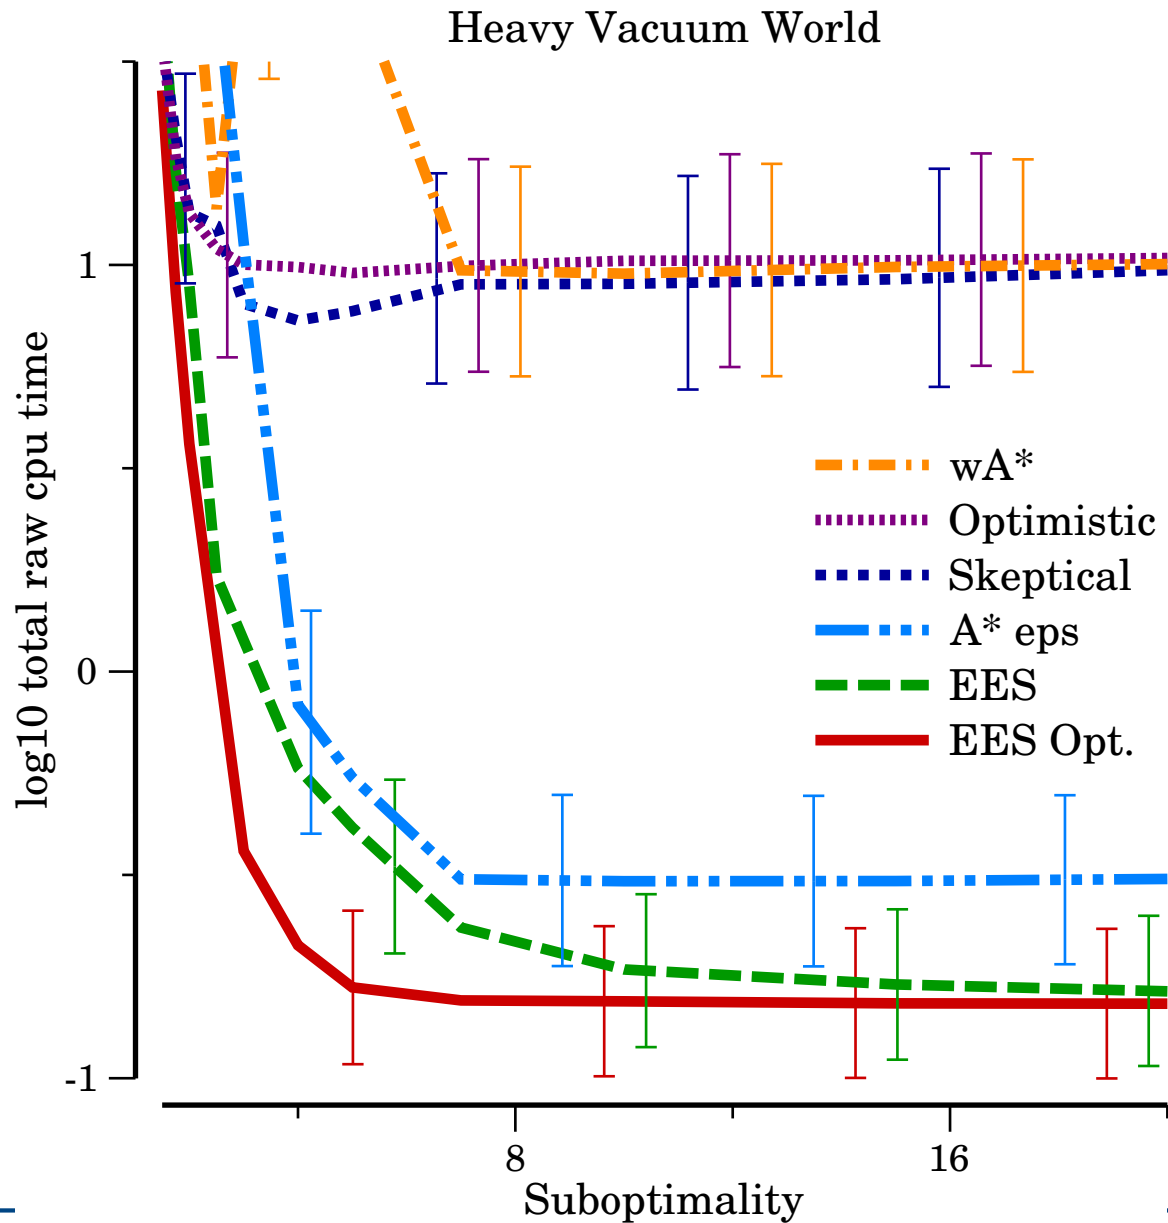
Utility Functions

Conclusion




EES Results

- Heuristic Search
- Greedy Search
- Bounded Search
- Direct Approach
- A^*_ϵ
- A^*_ϵ 's Flaw
- Estimated Cost
- Learning Cost
- EES
- EES Order
- EES Results**
- The Story So Far
- Contract Search
- Utility Functions
- Conclusion



EES Results



The Story So Far

Heuristic Search

Greedy Search

Bounded Search

■ Direct Approach

■ A_ϵ^*

■ A_ϵ^* 's Flaw

■ Estimated Cost

■ Learning Cost

■ EES

■ EES Order

■ EES Results

■ The Story So Far

Contract Search

Utility Functions

Conclusion

- search algorithms as agents
- more information sources
 1. $d(n)$: distance-to-go
 2. $\hat{f}(n)$: expected cost (expansion experience)
- more problem settings
 - ◆ greedy search: Speedy
 - ◆ bounded-suboptimal search: EES

Heuristic Search

Greedy Search

Bounded Search

Contract Search

■ Direct Approach

■ Vacillation

■ DAS

■ DAS Results

Utility Functions

Conclusion

Contract Search

A Direct Approach

minimize cost subject to absolute time bound

Heuristic Search

Greedy Search

Bounded Search

Contract Search

■ **Direct Approach**

■ Vacillation

■ DAS

■ DAS Results

Utility Functions

Conclusion

A Direct Approach

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ [Direct Approach](#)

■ [Vacillation](#)

■ [DAS](#)

■ [DAS Results](#)

[Utility Functions](#)

[Conclusion](#)

minimize cost subject to absolute time bound

anytime algorithms: series of solutions, eg, **unknown** deadline

A Direct Approach

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ [Direct Approach](#)

■ [Vacillation](#)

■ [DAS](#)

■ [DAS Results](#)

[Utility Functions](#)

[Conclusion](#)

minimize cost subject to absolute time bound

anytime algorithms: series of solutions, eg, **unknown** deadline

while time remains

expand node with best reachable solution

A Direct Approach

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ [Direct Approach](#)

■ [Vacillation](#)

■ [DAS](#)

■ [DAS Results](#)

[Utility Functions](#)

[Conclusion](#)

minimize cost subject to absolute time bound

anytime algorithms: series of solutions, eg, **unknown** deadline

while time remains

expand node with best reachable solution

best: **expected** cost

reachable: expected **search time to goal**

Vacillation

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ Direct Approach

■ **Vacillation**

■ DAS

■ DAS Results

[Utility Functions](#)

[Conclusion](#)

how long to reach goal $\hat{d}(n)$ steps away?

Vacillation

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ Direct Approach

■ **Vacillation**

■ DAS

■ DAS Results

[Utility Functions](#)

[Conclusion](#)

how long to reach goal $\hat{d}(n)$ steps away?

how long to go one step toward a goal?

estimate $\overline{\Delta e}$, time between a node's generation and expansion

Vacillation

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

■ Direct Approach

■ **Vacillation**

■ DAS

■ DAS Results

[Utility Functions](#)

[Conclusion](#)

how long to reach goal $\hat{d}(n)$ steps away?

how long to go one step toward a goal?

estimate $\overline{\Delta e}$, time between a node's generation and expansion

just record current time in node when generating

source #3 of 5: search vacillation

$$d_{max} = \text{time remaining} / \overline{\Delta e}$$

Deadline-Aware Search (Dionne et al, SoCS-11)

Heuristic Search

Greedy Search

Bounded Search

Contract Search

■ Direct Approach

■ Vacillation

■ DAS

■ DAS Results

Utility Functions

Conclusion

1. while time remains and *open* is not empty
2. $d_{max} \leftarrow$ estimate bound
3. $s \leftarrow$ pop $best_f$ from *open*
4. if *s* is a goal
5. save if best so far
6. else if $\hat{d}(s) < d_{max}$
7. expand *s*
8. else
9. add *s* to *pruned*
10. if *open* is empty, recover using *pruned*

DAS Results

Heuristic Search

Greedy Search

Bounded Search

Contract Search

■ Direct Approach

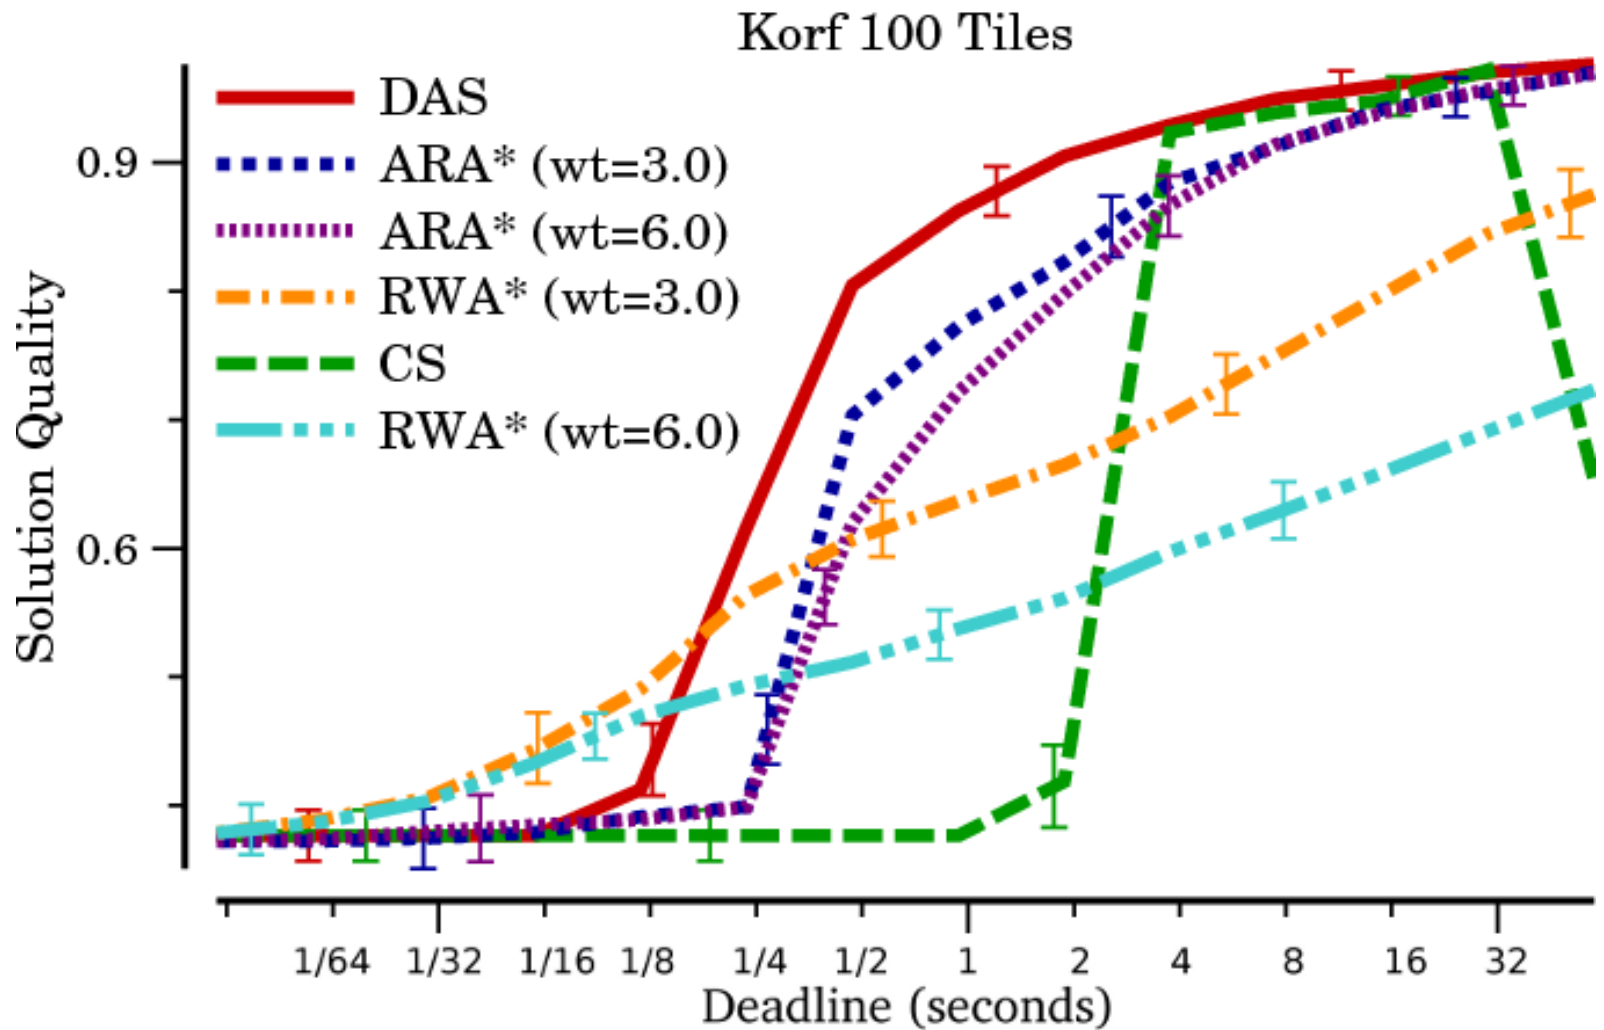
■ Vacillation

■ DAS

■ DAS Results

Utility Functions

Conclusion



DAS Results

Heuristic Search

Greedy Search

Bounded Search

Contract Search

■ Direct Approach

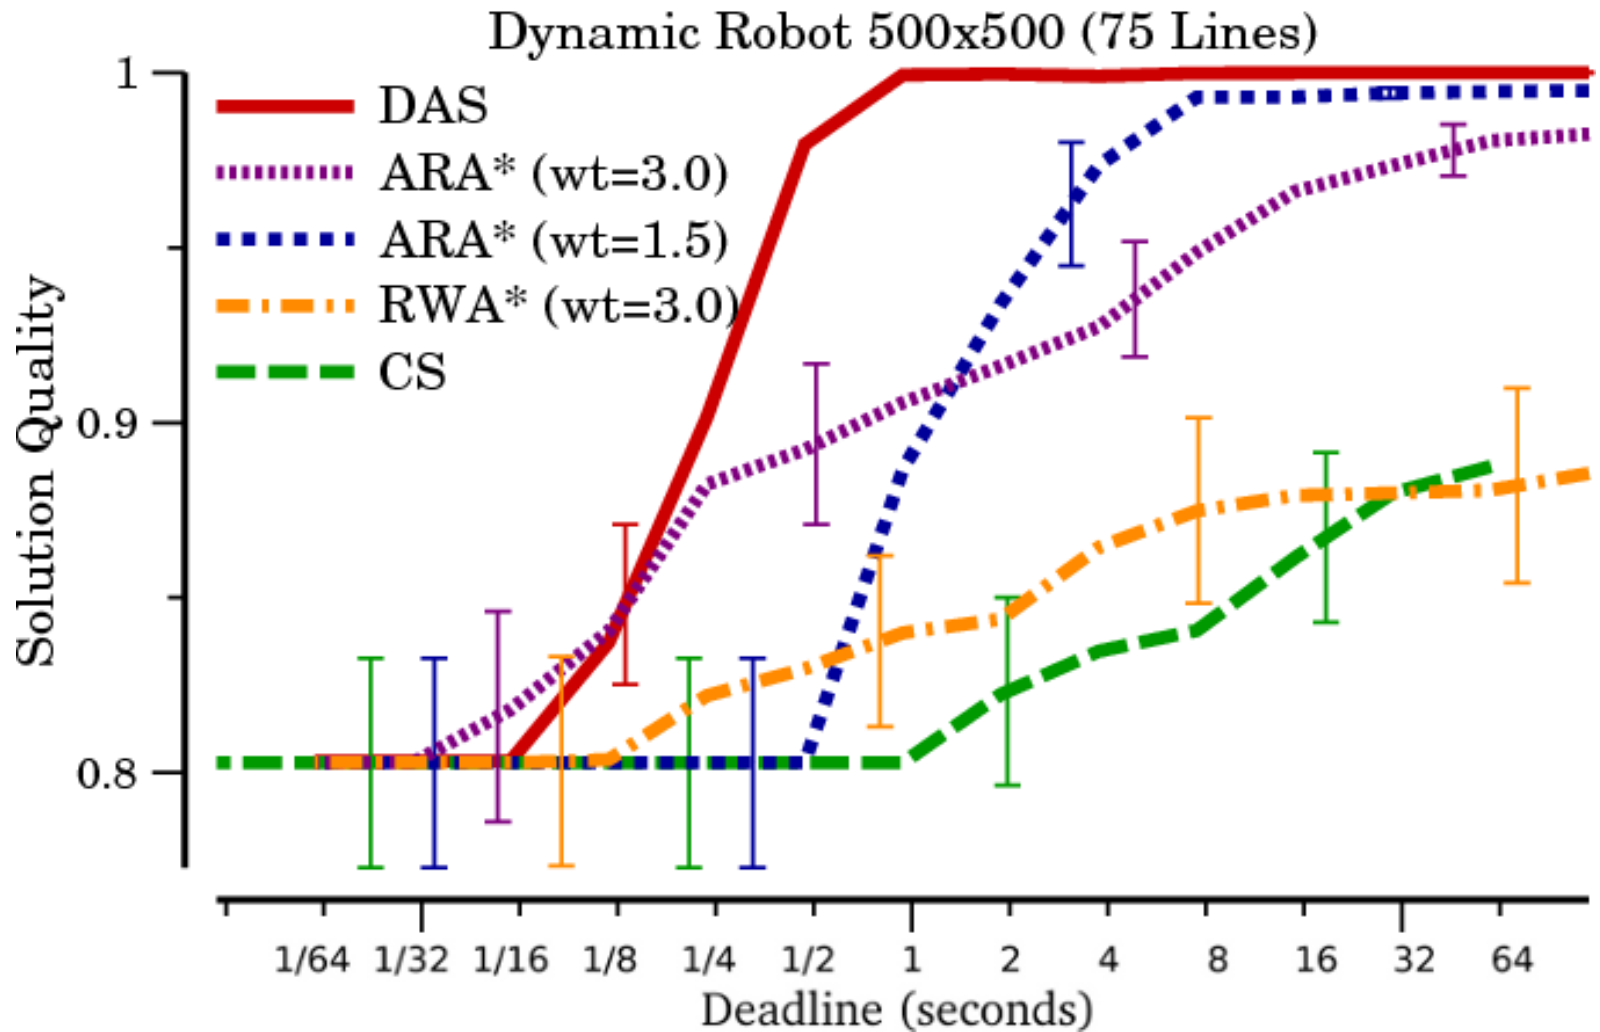
■ Vacillation

■ DAS

■ DAS Results

Utility Functions

Conclusion



Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

- Utility Fns
- Anytime Algs
- BUGSY
- Properties
- BUGSY Results

Conclusion

Search with a Utility Function

Utility-based Search

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

[Utility Functions](#)

■ **Utility Fns**

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

[Conclusion](#)

balance time and cost according to user's utility function

Example: linear utility function

for solution of cost f produced after time t :

$$U(f, t) = -w_f \cdot f - w_t \cdot t$$

Utility-based Search

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

[Utility Functions](#)

■ **Utility Fns**

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

[Conclusion](#)

balance time and cost according to user's utility function

Example: linear utility function

for solution of cost f produced after time t :

$$U(f, t) = -w_f \cdot f - w_t \cdot t$$

Example: minimize goal achievement time

if cost = plan makespan,

$$U(f, t) = -f - t$$

Utility-based Search

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

Conclusion

balance time and cost according to user's utility function

Example: linear utility function

for solution of cost f produced after time t :

$$U(f, t) = -w_f \cdot f - w_t \cdot t$$

Example: minimize goal achievement time

if cost = plan makespan,

$$U(f, t) = -f - t$$

anytime algorithms?

Problems with Anytime Algorithms

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

Conclusion

Requires a termination policy, assuming:

1. relevant solver features for predicting progress are known
2. training data available
3. new instance is similar in relevant aspects to training
4. relevant instance aspects are known

Problems with Anytime Algorithms

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY

■ Properties


■ BUGSY Results

Conclusion

Requires a termination policy, assuming:

1. relevant solver features for predicting progress are known
2. training data available
3. new instance is similar in relevant aspects to training
4. relevant instance aspects are known

Impossible to design optimally:



Must know the user's trade-off!

Best-first Utility-guided Search, Yes! (BUGSY, IJCAI-07)

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

Conclusion

source #4 of 5: user's true objective/utility function

best-first search according to utility:

$$U(n) = \max_{s \text{ under } n} (-w_f \cdot f(s) - w_t \cdot t(s))$$

Best-first Utility-guided Search, Yes! (BUGSY, IJCAI-07)

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY

■ Properties

■ BUGSY Results

Conclusion

source #4 of 5: user's true objective/utility function

best-first search according to utility:

$$U(n) = \max_{s \text{ under } n} (-w_f \cdot f(s) - w_t \cdot t(s))$$

convert $\hat{d}(n)$ to $t(n)$

approximate s under n by cheapest and nearest

- need $d_{cheapest}$, $d_{nearest}$, $h_{cheapest}$, $h_{nearest}$
- source #5 of 5: cost-to-go to nearest, $h_{nearest}$
- seems straightforward in many domains

Properties

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY


■ Properties

■ BUGSY Results

Conclusion

Different from anytime algorithms

- no need for termination policy (training data, precomputation)
- can spend all effort pursuing one solution
- no fixed trade-off



Properties

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs

■ BUGSY


■ Properties

■ BUGSY Results

Conclusion

Different from anytime algorithms

- no need for termination policy (training data, precomputation)
- can spend all effort pursuing one solution
- no fixed trade-off



BUGSY Results

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

■ Utility Fns

■ Anytime Algs


■ BUGSY

■ Properties

■ BUGSY Results

Conclusion

grid pathfinding



B — BUGSY

R — ARA* with termination policy learned off-line

A* — A*

S — Speedy

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

■ Rational Search

■ Search

■ Summary

Conclusion

Rational Search

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)


[Utility Functions](#)

[Conclusion](#)

Rational Search

■ Search

■ Summary



Rational Search

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)


[Utility Functions](#)

[Conclusion](#)

■ Rational Search

■ Search

■ Summary



The Search Problem

Heuristic Search

Greedy Search

Bounded Search

Contract Search

Utility Functions

Conclusion

■ Rational Search

■ Search

■ Summary

Traditional:

1. initial state
2. goal predicate
3. expand — yields $g(n)$
4. $h(n)$

The Search Problem

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

[Utility Functions](#)

[Conclusion](#)

■ Rational Search

■ Search

■ Summary

Traditional:

1. initial state
2. goal predicate
3. expand — yields $g(n)$
4. $h(n)$

Possible:

5. $d(n)$: distance-to-go
6. $\hat{f}(n)$: expected cost, inadmissible h (expansion experience)
7. $\overline{\Delta e}$: expansion delay, time to goal (expansion experience)
8. $U(f, t)$: user's utility function, true objective
9. $h_{nearest}(n)$: cost-to-go to nearest

Summary

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

[Utility Functions](#)

[Conclusion](#)

■ Rational Search

■ Search

■ Summary

- search algorithms as agents
 - ◆ expansion as sensor reading instead of proof step!
 - ◆ connections:
 - reinforcement learning
 - metareasoning
 - decision-making
- more information sources
 - what else can we exploit? and how?
- more problem settings
 - ◆ greedy search: Speedy
 - ◆ bounded-suboptimal search: EES
 - ◆ contract search: DAS
 - ◆ utility-based search: BUGSY

The University of New Hampshire

tell your students to apply to grad school in CS at UNH!

[Heuristic Search](#)

[Greedy Search](#)

[Bounded Search](#)

[Contract Search](#)

[Utility Functions](#)

[Conclusion](#)

■ Rational Search

■ Search

■ Summary



- friendly faculty
- funding
- individual attention
- beautiful campus
- low cost of living
- easy access to Boston, White Mountains
- strong in AI, infoviz, networking, bioinformatics