# **Machine Learning for Classical Planning:** Neural Network Heuristics, Online Portfolios, and State Space Topologies

Patrick Ferber

November 17, 2022

UNIVERSITÄT
DES
SAARLANDES

University
of Basel

# Planning

# Planning

# Planning

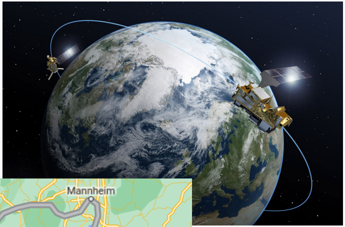

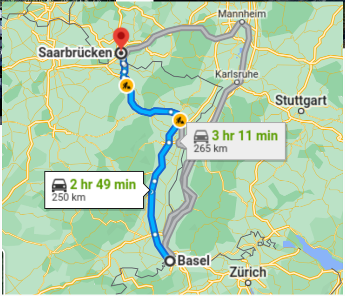
Mannheim
Saarbrücken
Karlsruhe
Stuttgart
3 hr 11 min
265 km
2 hr 49 min
250 km
Basel
Zürich

# Planning

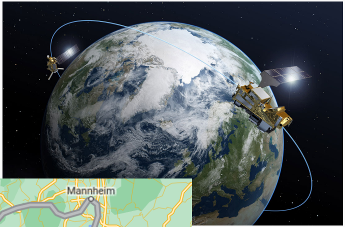

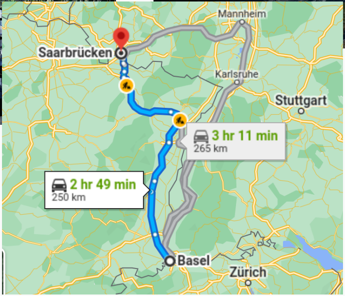


sequence of actions to satisfy a goal

# Classical Planning

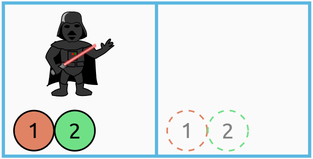

# Classical Planning

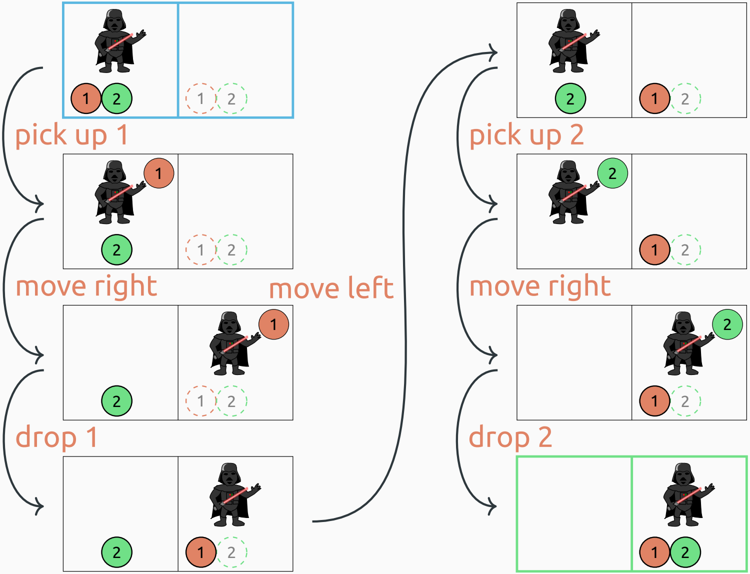

# Greedy Best-First Search

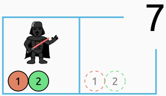

# Greedy Best-First Search

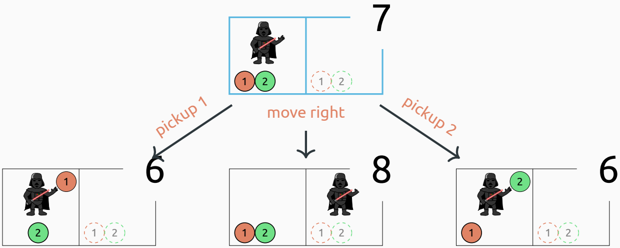

# Greedy Best-First Search

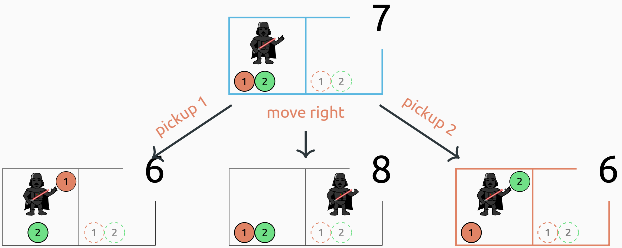

# Greedy Best-First Search

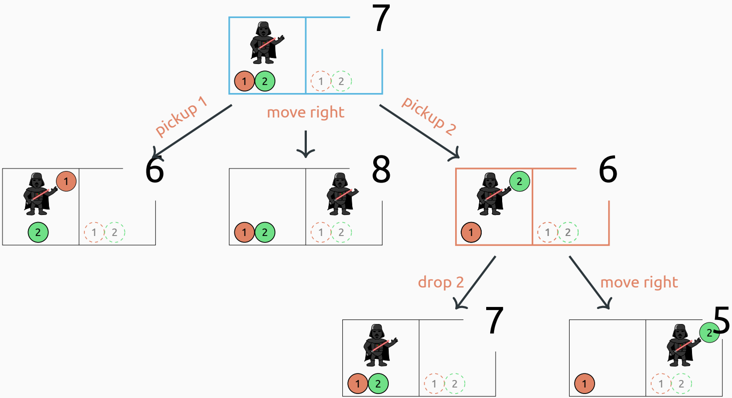

# Greedy Best-First Search

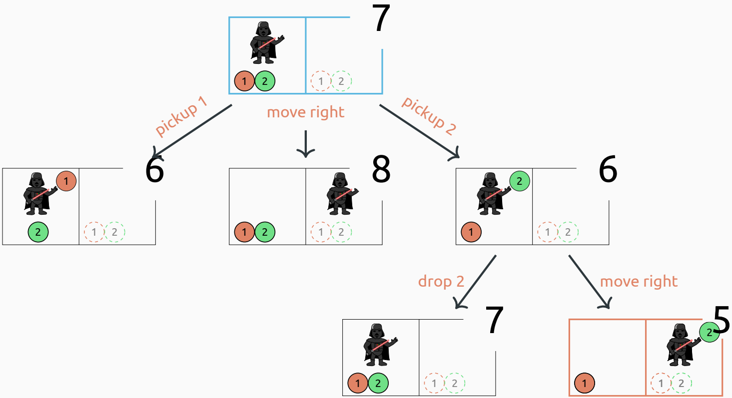

# Greedy Best-First Search

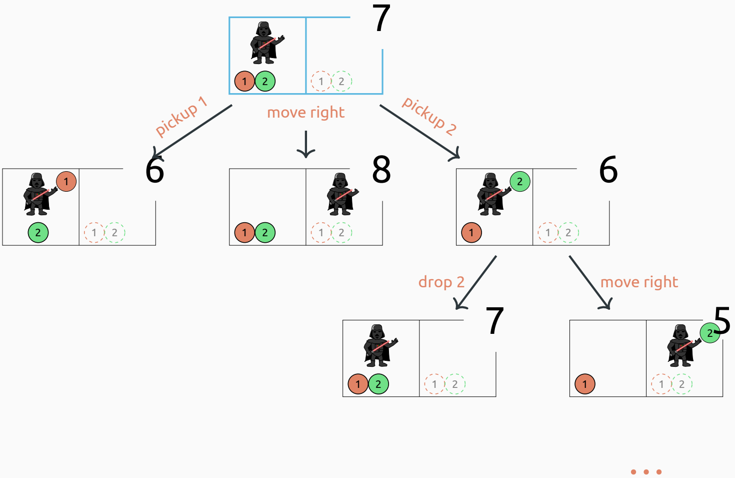

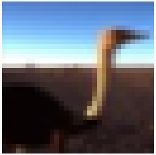

→ ostrich

# Machine Learning

 → car

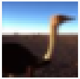 → ostrich

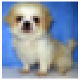 → dog

# Machine Learning

sample  $\longrightarrow$ car

sample  $\longrightarrow$ ostrich

sample 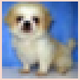 $\longrightarrow$ dog

# Machine Learning

| | Features | Labels |
|---|---|---|
| sample |  | → car |
| sample |  | → ostrich |
| sample | 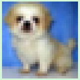 | → dog |

Heuristics

Online Portfolios

State Space Topologies

Heuristics

Online Portfolios

State Space Topologies

# Heuristic

h( 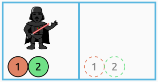 ) = 7

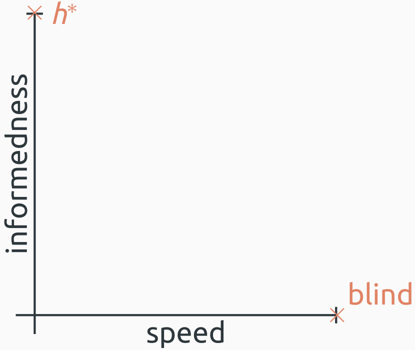

# Heuristic

h( 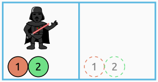 ) = 7

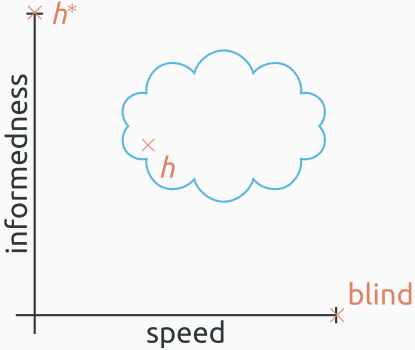

# Heuristic

h( 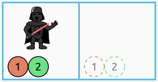 ) = 7


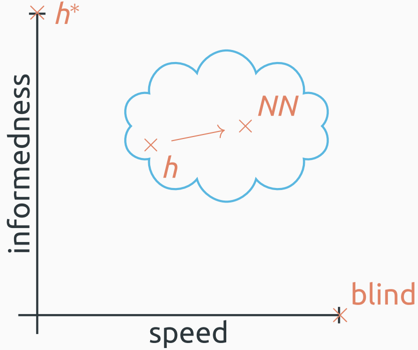

# Learning Heuristics: Take I

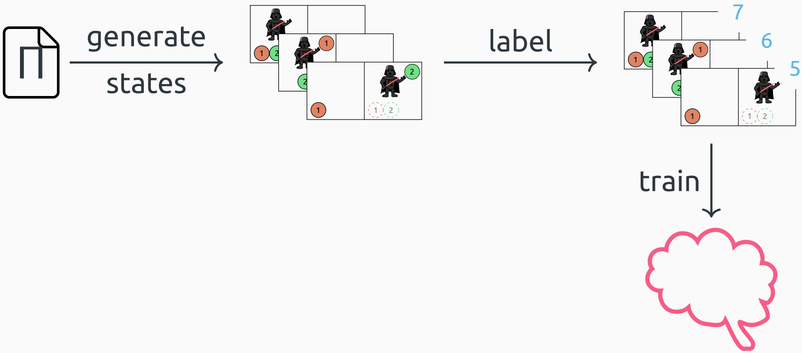


Workflow: Ferber *et al.* [2020]

# Sampling States

## Progression Random Walks

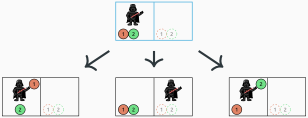

## Sampling States

### Progression Random Walks

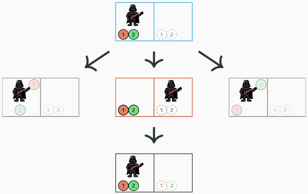

## Progression Random Walks

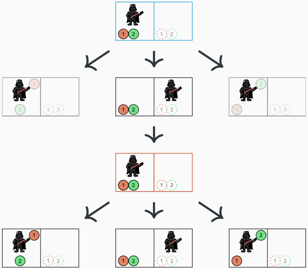

# Sampling States

## Progression Random Walks

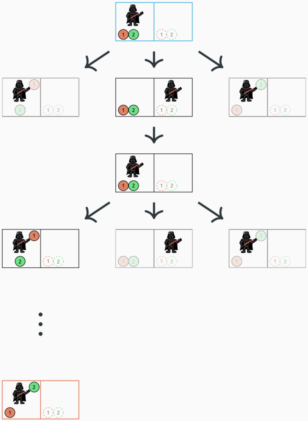

# Labeling States

$h^*($ 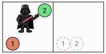 $)$

$h^{FF}($ 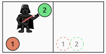 $)$

$GBFS($ 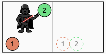 $, h^{FF})$

$h^{FF}$: Hoffmann and Nebel [2001]

# Labeling States

$\cancel{h^*(}$ 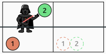 $\cancel{)}$

$h^{FF}($ 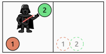 $)$

$GBFS($ 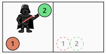 $, h^{FF})$

$h^{FF}$: Hoffmann and Nebel [2001]

# Labeling States

$h^{*}($ 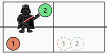 $)$

$h^{FF}($ 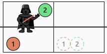 $)$

$GBFS($ 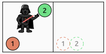 $, h^{FF})$

$h^{FF}$: Hoffmann and Nebel [2001]

# Neural Network Heuristic

# Neural Network Heuristic

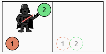


at-vader(left) at-vader(right) ⋯ carry(ball2)

1 0 1

↓


↓

h

# Neural Network Heuristic

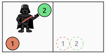


| at-vader(left) | at-vader(right) | ⋯ | carry(ball2) |
|---|---|---|---|
| 1 | 0 | | 1 |

FFN

2

| | | | | |
|---|---|---|---|---|
| 0.8 | 0.7 | 0.4 | 0.2 | 0.1 |

# Limitation Teacher Search

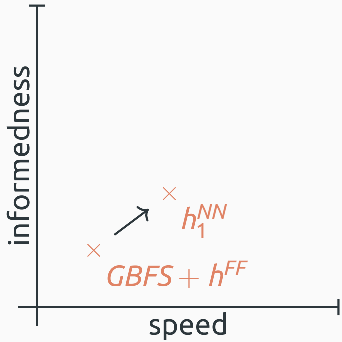

# Limitation Teacher Search

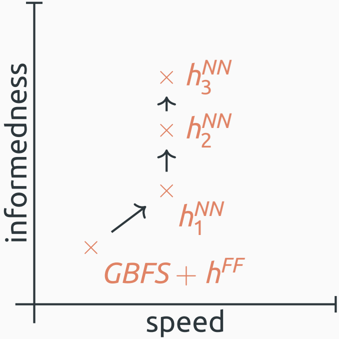

# Learning Heuristics: Take II

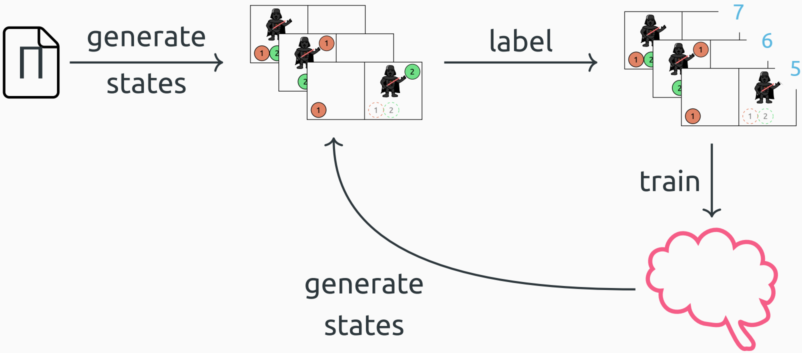


Ferber *et al.* [2022b]

# Labeling States

$GBFS($ 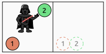

$, h^{FF})$

$GBFS($ 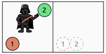

$, h^{NN})$

$exp(GBFS($ 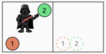

$, h^{NN}))$

$AVI($ 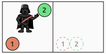

$, h^{NN})$

# Sampling States

## Progression Random Walks

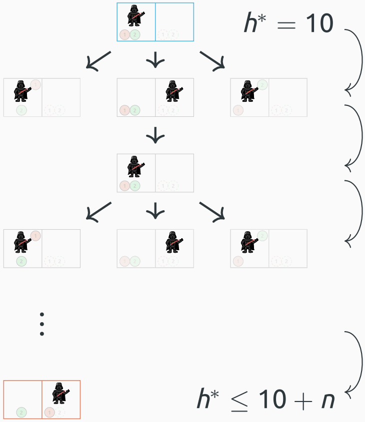

$h^* = 10$

$h^* \leq 10 + n$

# Sampling States

Progression Random Walks

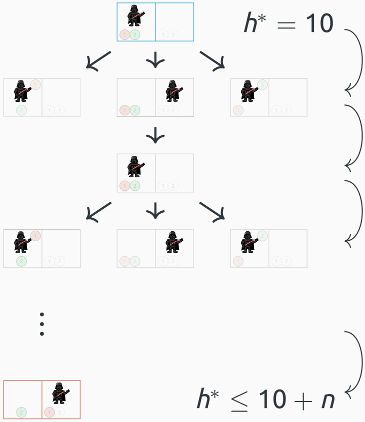


Regression Random Walks

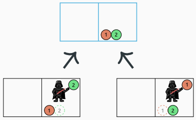

# Sampling States

## Progression Random Walks

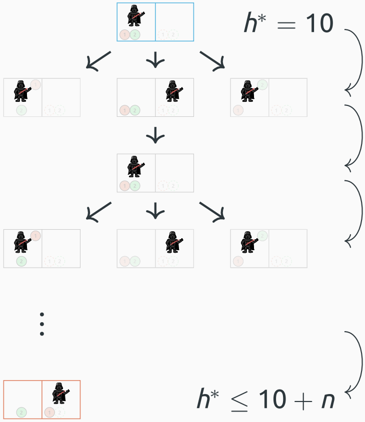


## Regression Random Walks

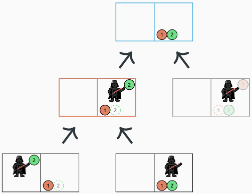

# Sampling States

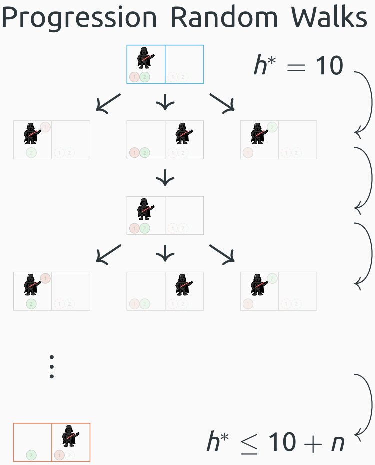


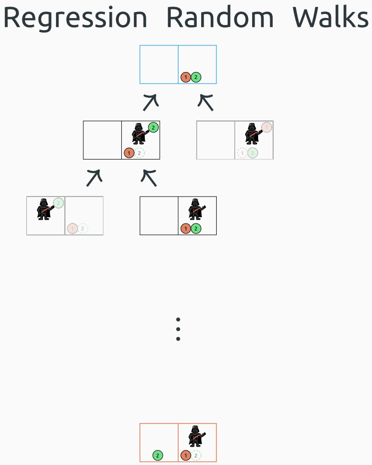

# Sampling States

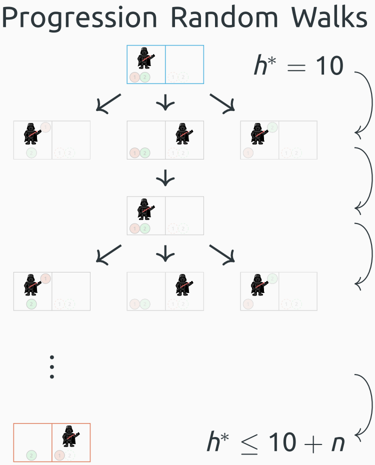


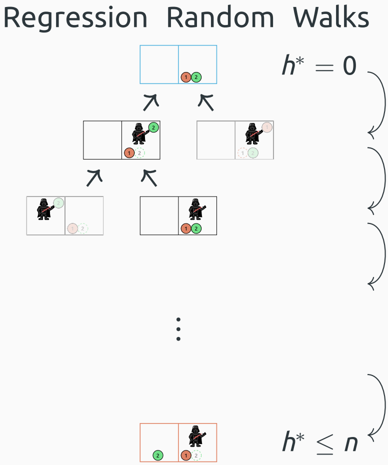

# Heuristic Coverages

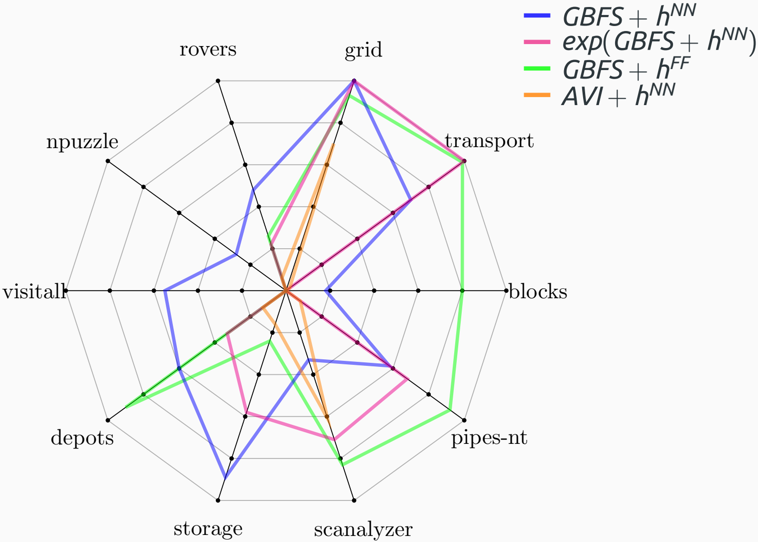




Ferber *et al.* [2022b]

# Heuristic Coverages

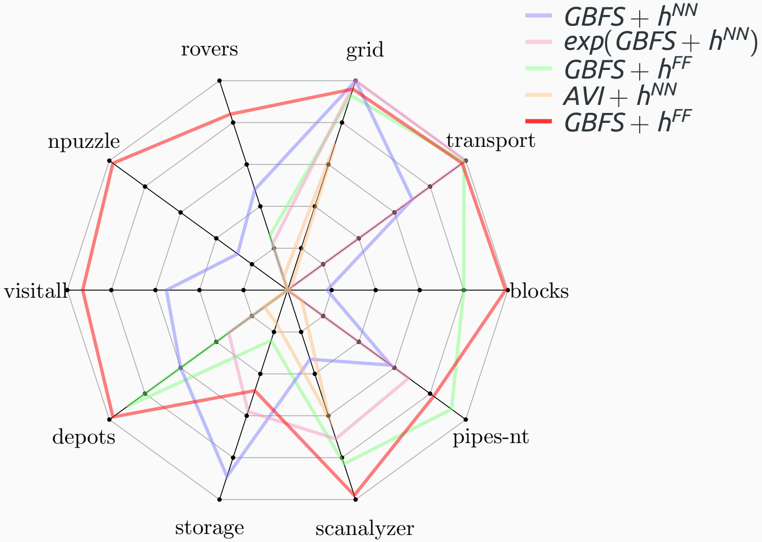




Ferber *et al.* [2022b]

Heuristics

Online Portfolios

State Space Topologies

# Planner Coverages

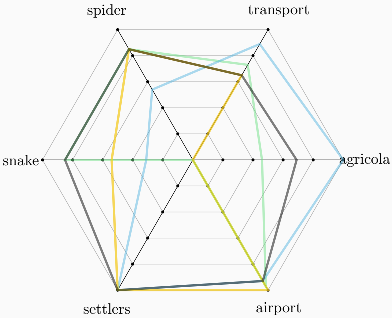




Sievers *et al.* [2019a]

# Planner Coverages

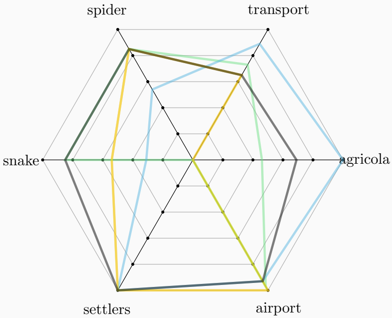


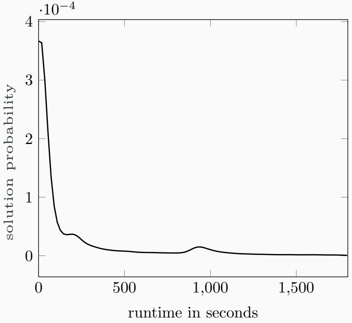




Sievers *et al.* [2019a]

# Planner Coverages

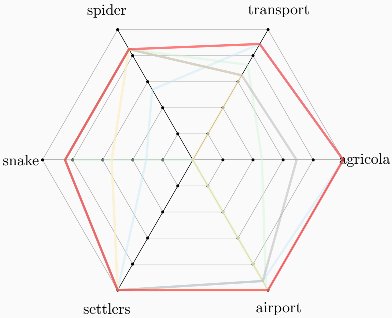


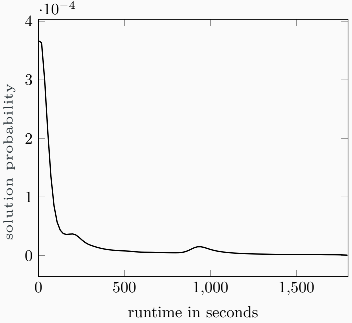




Sievers *et al.* [2019a]

# Planner Coverages

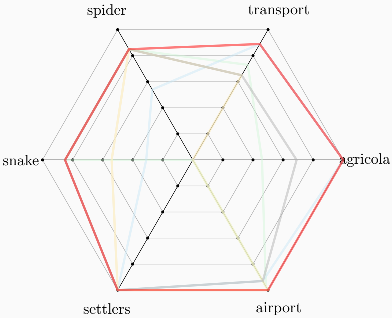


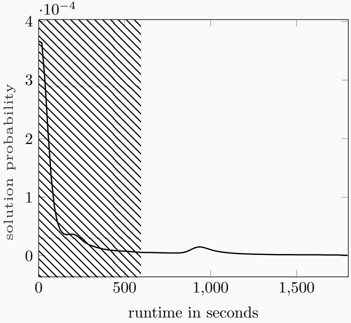




Sievers *et al.* [2019a]

# Online Portfolio

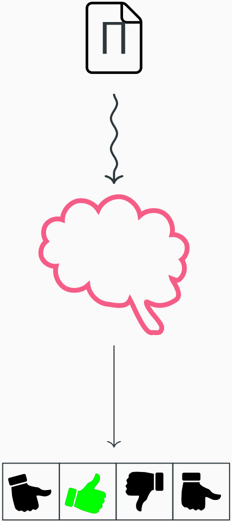




Sievers *et al.* [2019a]; Ma *et al.* [2020]

# Online Portfolio

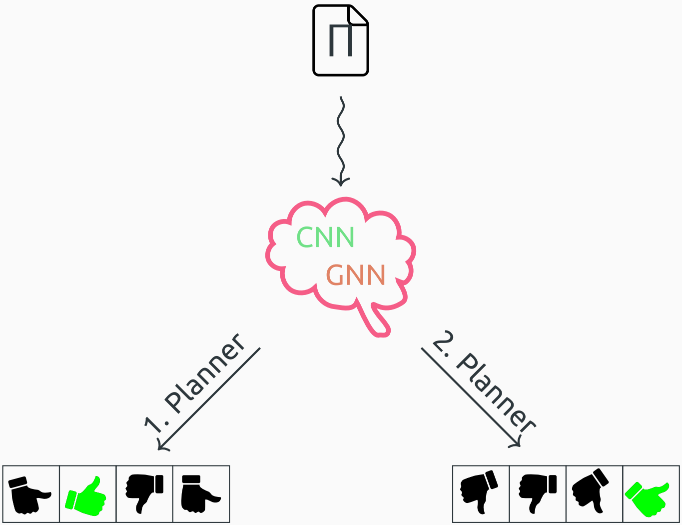




Sievers *et al.* [2019a]; Ma *et al.* [2020]

# Learning Online Portfolios

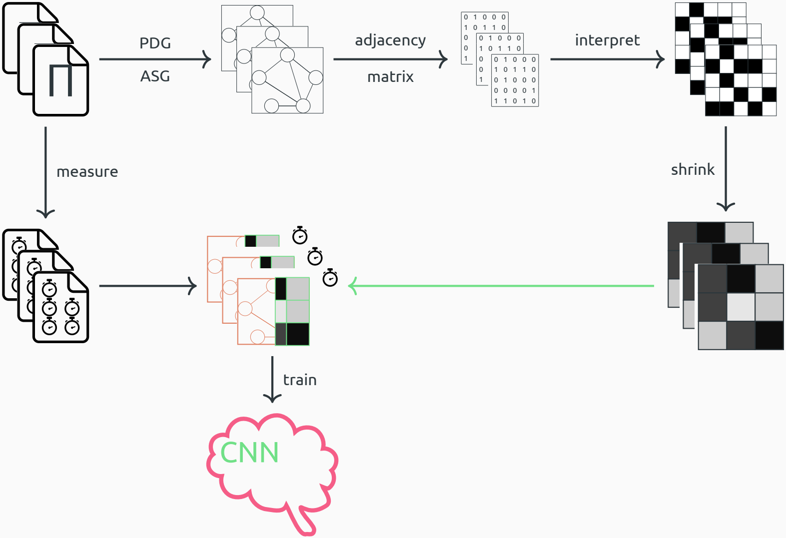




PDG: Pochter *et al.* [2011]    ASG: Sievers *et al.* [2019b]    Framework: Sievers *et al.* [2019a]; Ma *et al.* [2020]

# Learning Online Portfolios

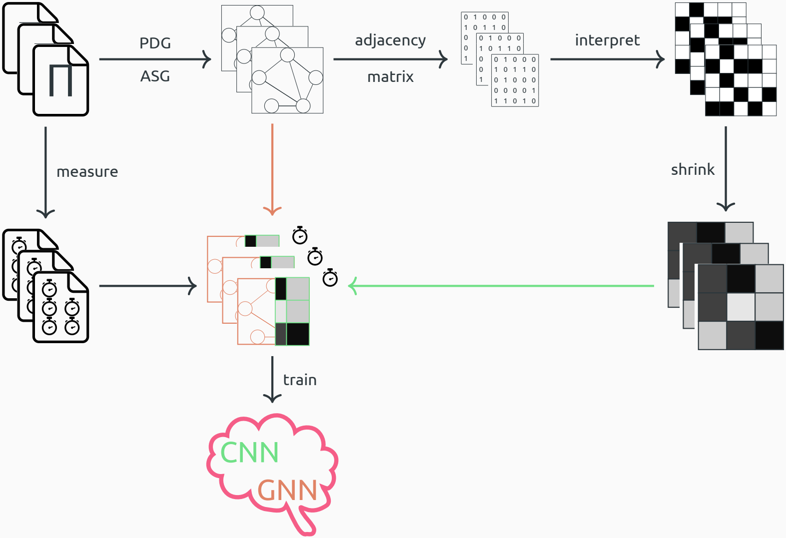




PDG: Pochter *et al.* [2011]   ASG: Sievers *et al.* [2019b]   Framework: Sievers *et al.* [2019a]; Ma *et al.* [2020]

# Portfolio Coverages

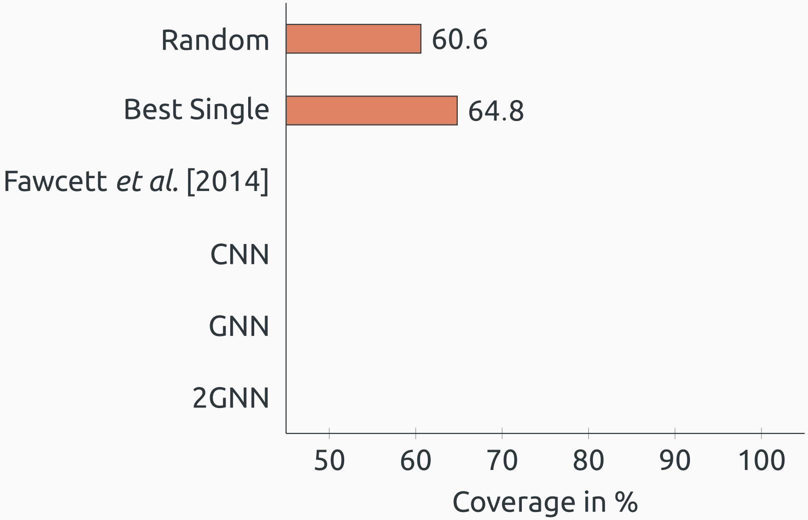




Ma *et al.* [2020]

# Portfolio Coverages

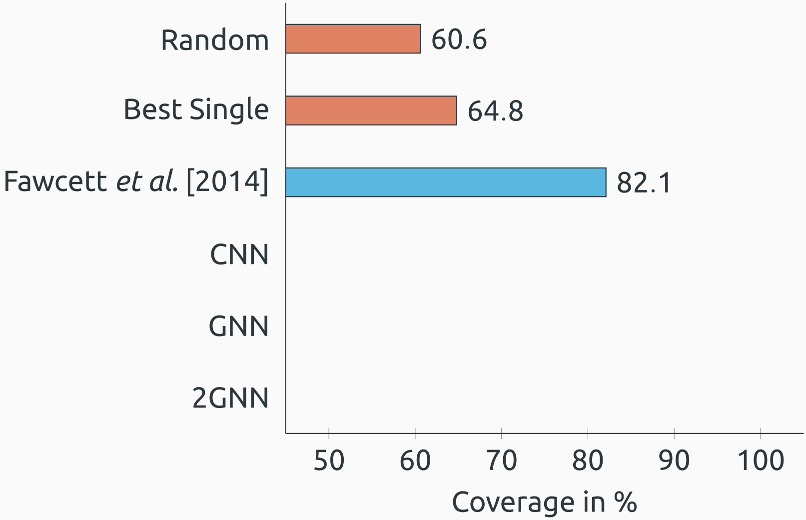




Ma *et al.* [2020]

# Portfolio Coverages

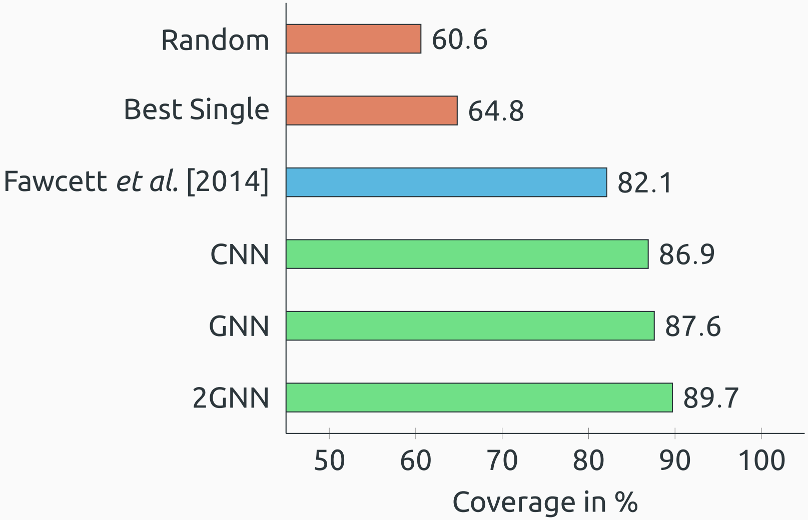




Ma *et al.* [2020]

Heuristics

Online Portfolios

State Space Topologies

# State Space Topologies

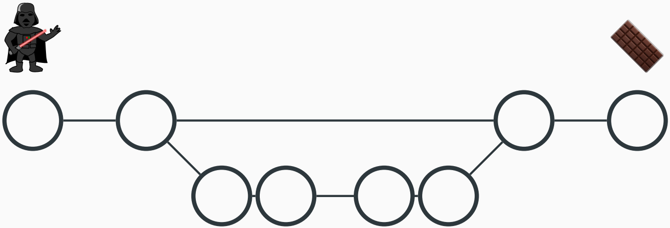



Hoffmann [2005]

# State Space Topologies

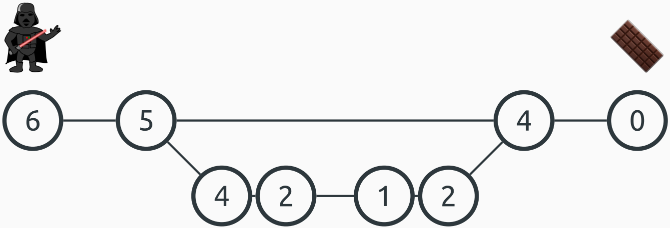




Hoffmann [2005]

# State Space Topologies

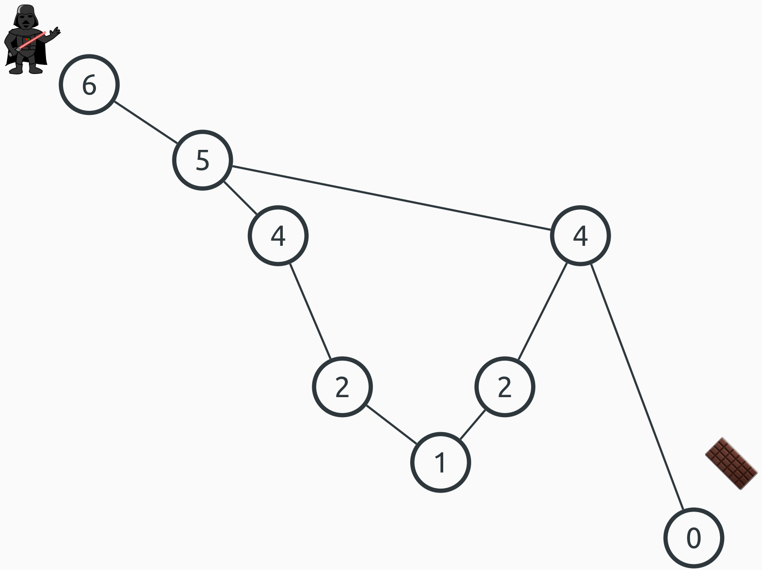




Hoffmann [2005]

# State Space Topologies

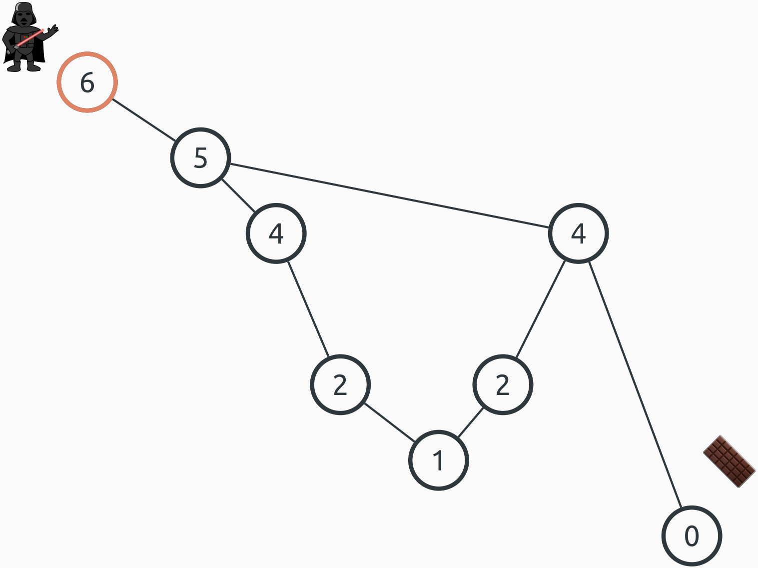




Hoffmann [2005]

# State Space Topologies

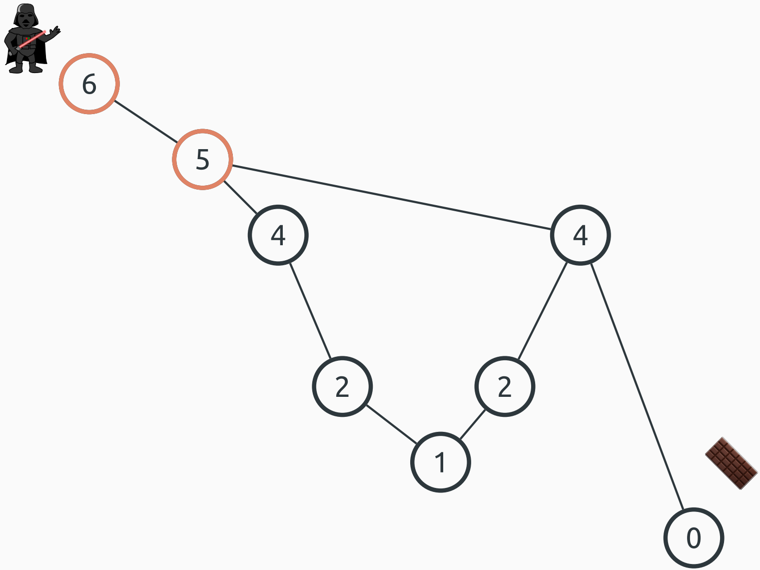




Hoffmann [2005]

# State Space Topologies

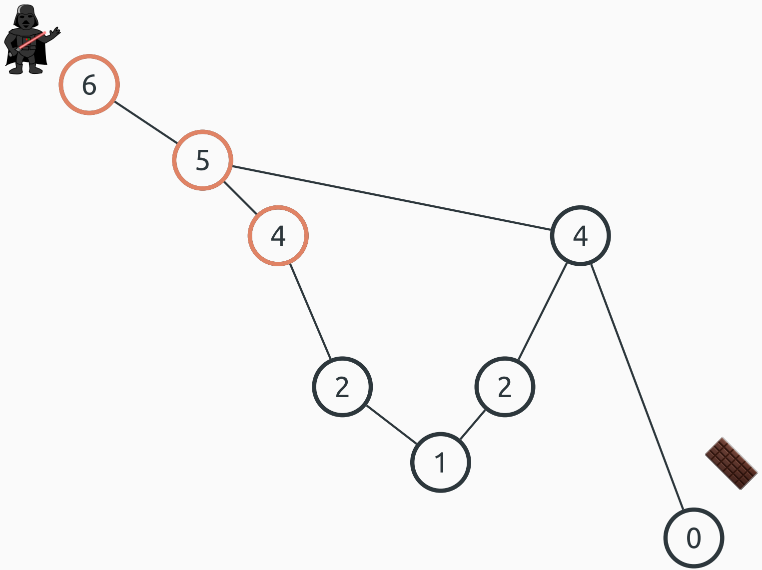




Hoffmann [2005]

# State Space Topologies

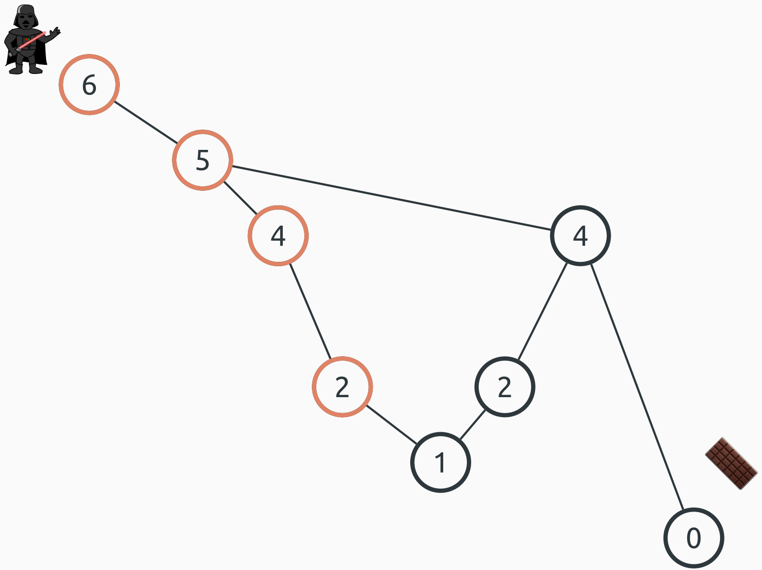




Hoffmann [2005]

# State Space Topologies

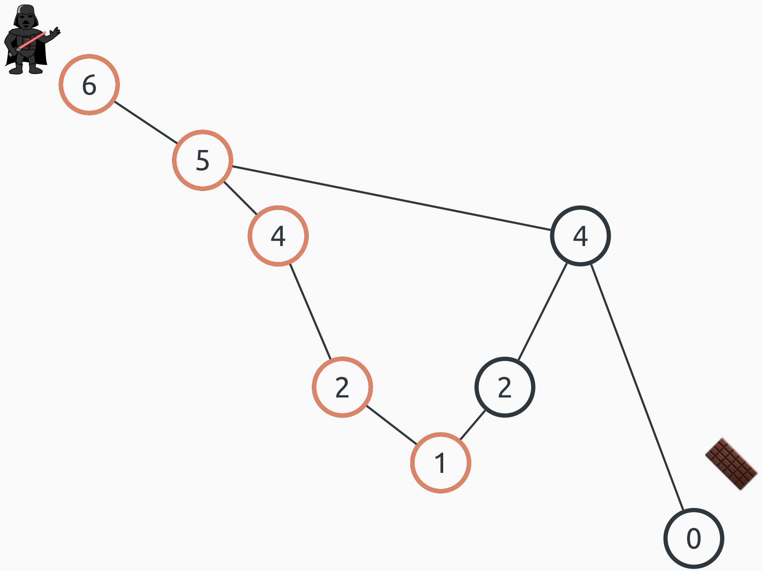




Hoffmann [2005]

# State Space Topologies

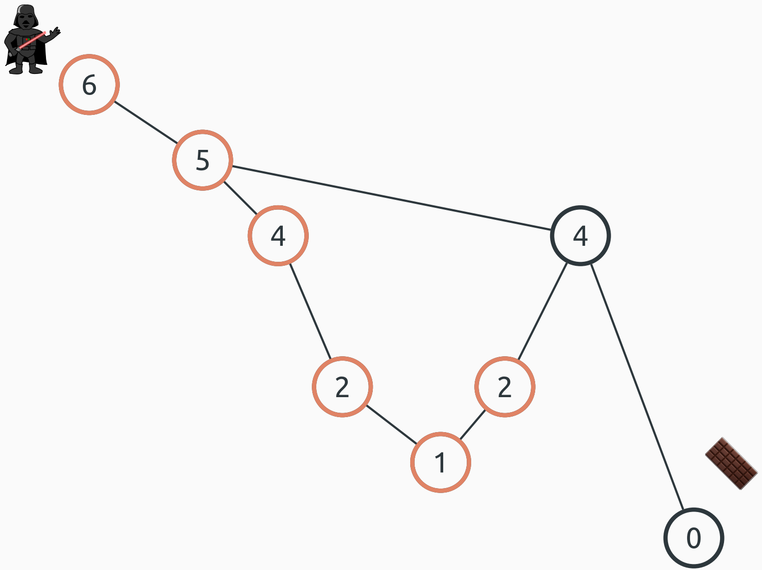




Hoffmann [2005]

# State Space Topologies

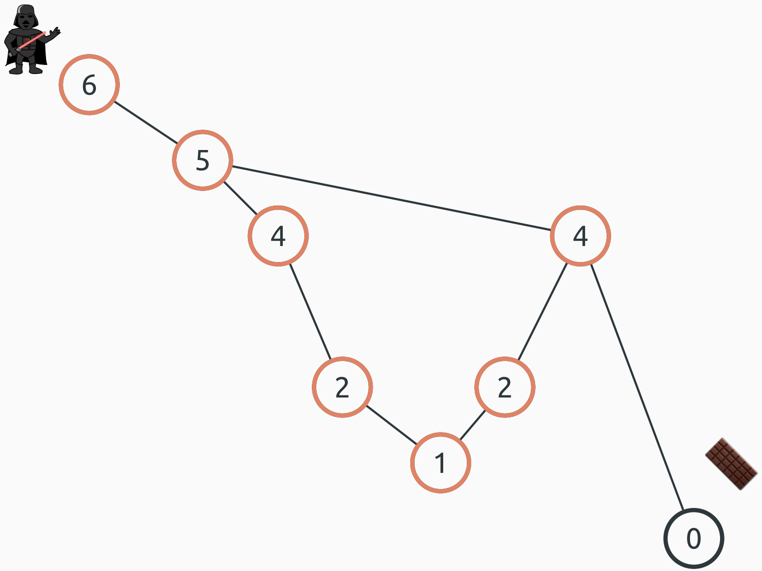


Hoffmann [2005]

# State Space Topologies

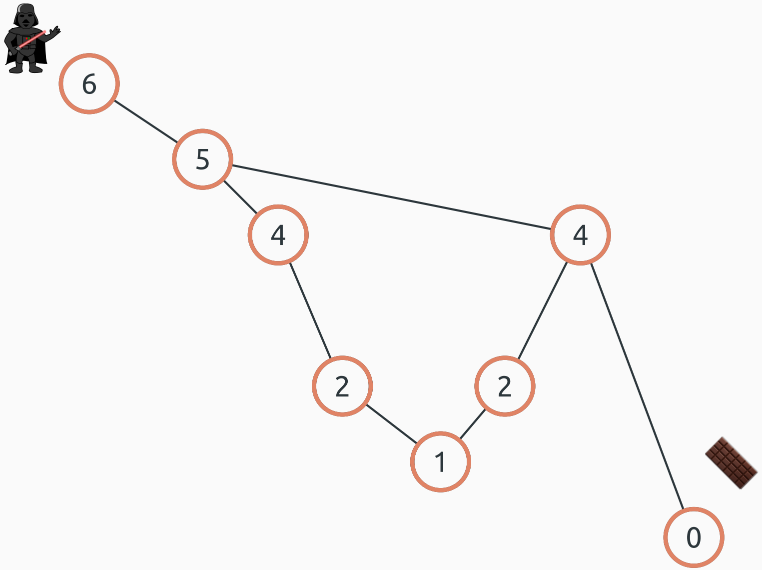




Hoffmann [2005]

# State Space Topologies

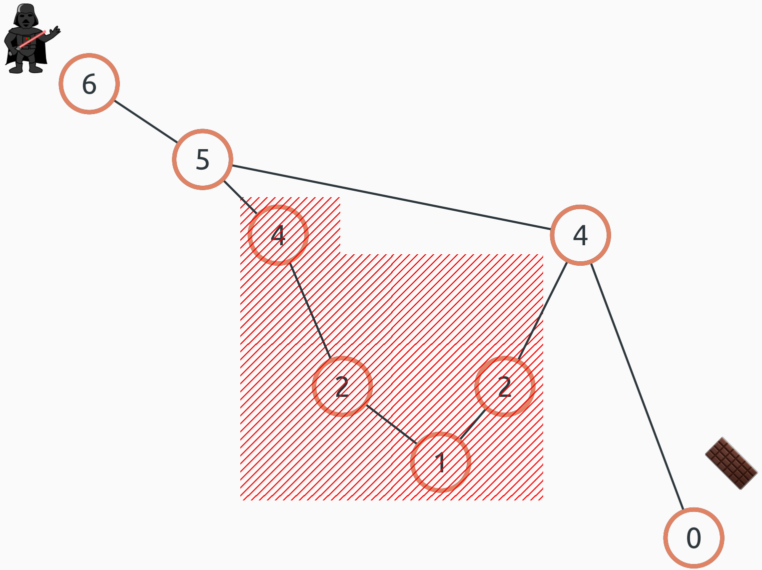


Hoffmann [2005]

# Progress States

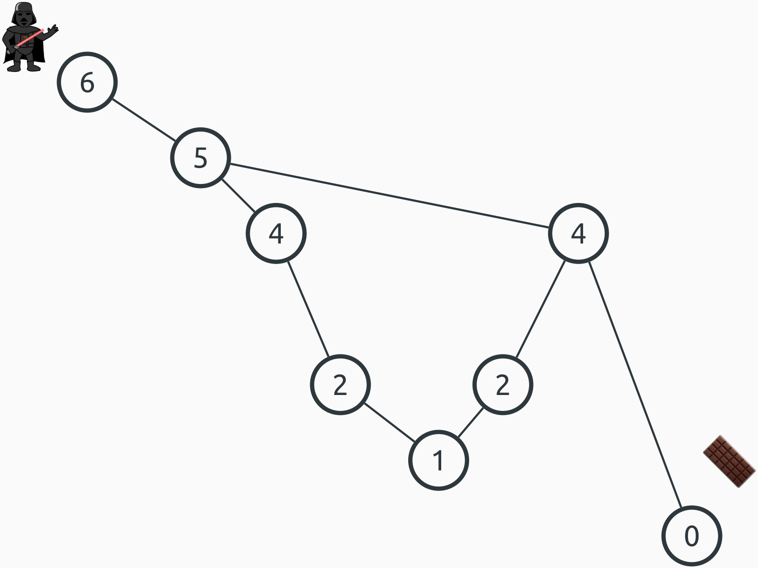




HWM: Wilt and Ruml [2014]    Progress States: Heusner *et al.* [2018]

# Progress States

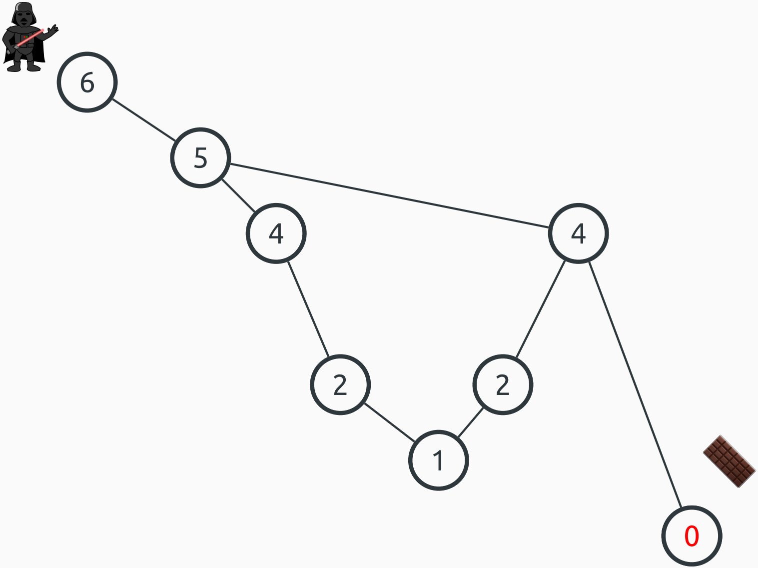




HWM: Wilt and Ruml [2014]    Progress States: Heusner *et al.* [2018]

# Progress States

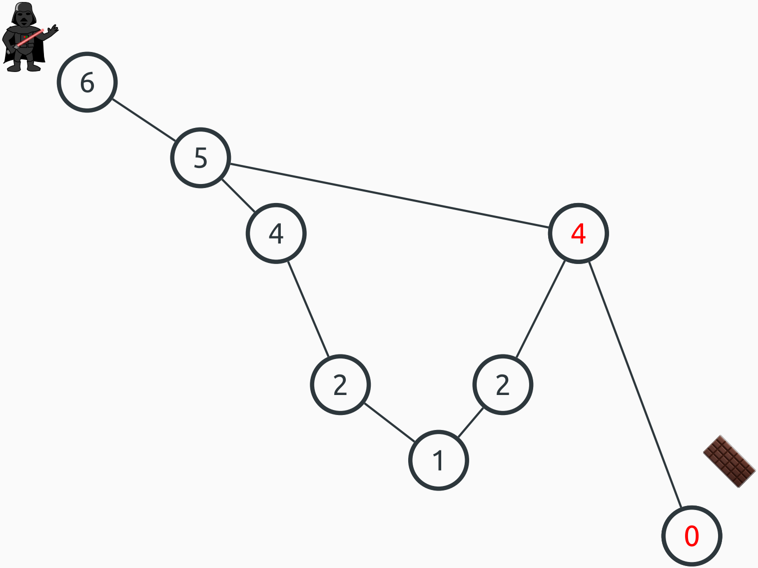




HWM: Wilt and Ruml [2014]  Progress States: Heusner *et al.* [2018]

# Progress States

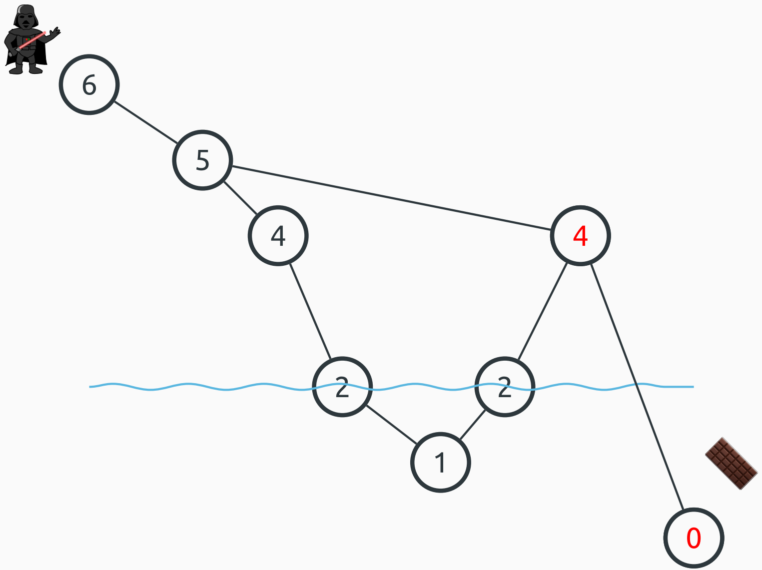




HWM: Wilt and Ruml [2014] Progress States: Heusner *et al.* [2018]

# Progress States

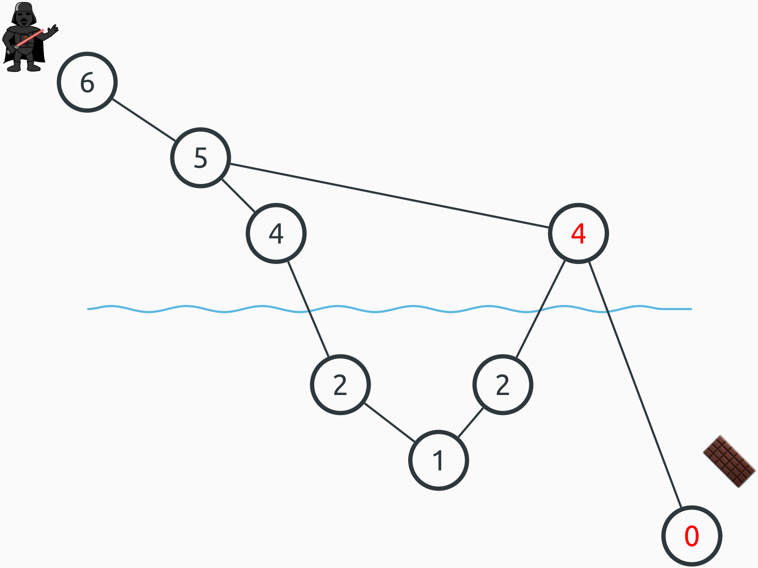




HWM: Wilt and Ruml [2014]  Progress States: Heusner *et al.* [2018]

# Progress States

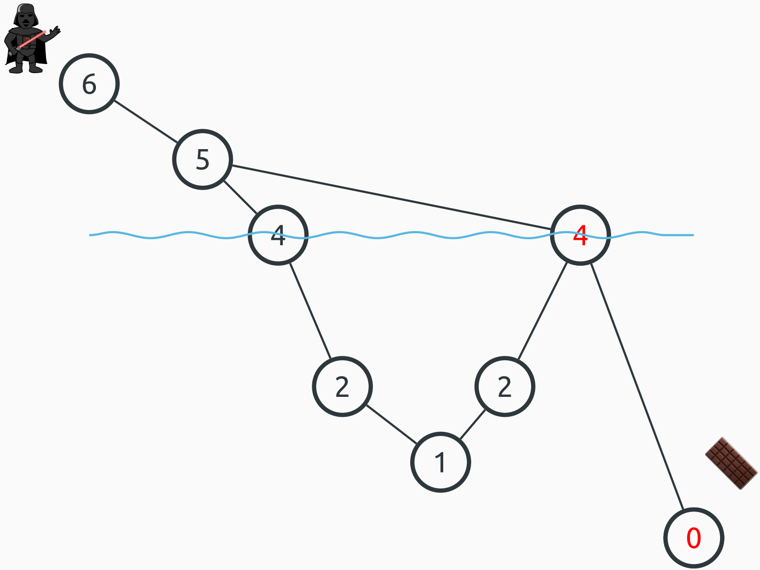




HWM: Wilt and Ruml [2014]    Progress States: Heusner *et al.* [2018]

# Progress States

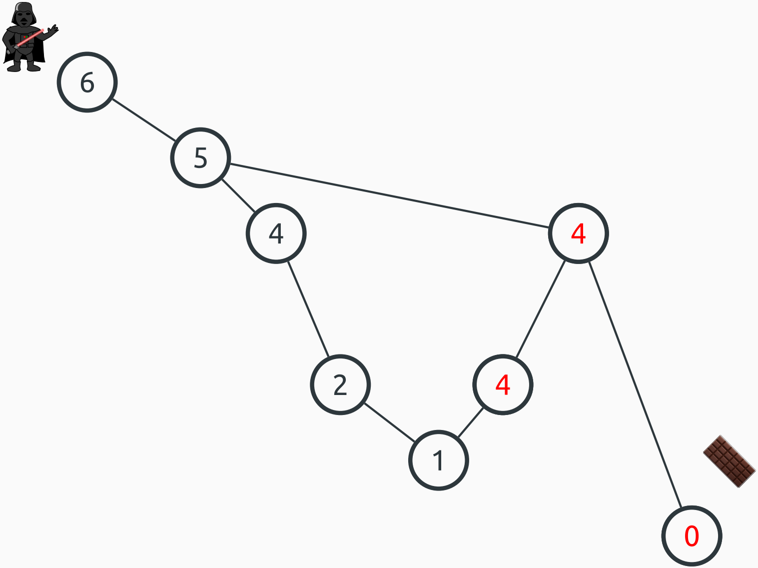




HWM: Wilt and Ruml [2014]    Progress States: Heusner *et al.* [2018]

# Progress States

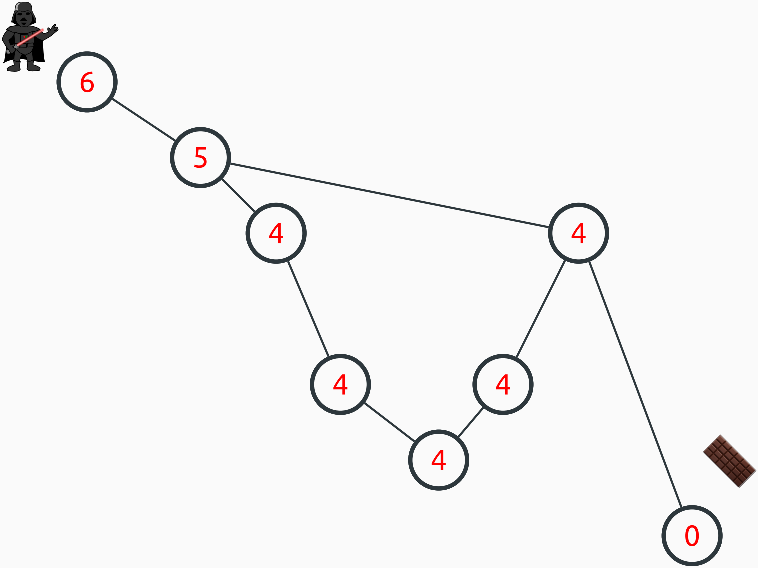




HWM: Wilt and Ruml [2014]     Progress States: Heusner *et al.* [2018]

# Progress States

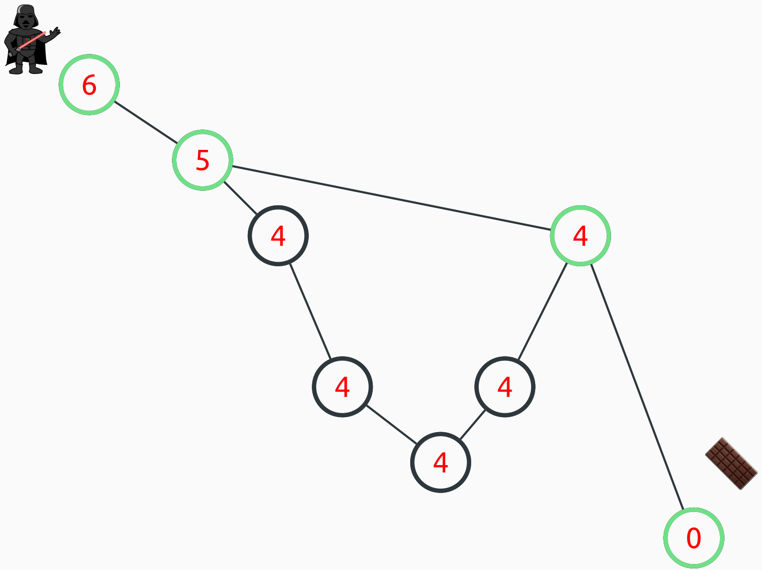




HWM: Wilt and Ruml [2014]    Progress States: Heusner *et al.* [2018]

# Description Logic Formulas

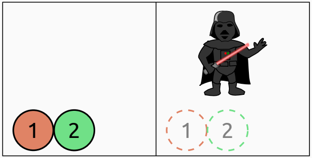


$$
\begin{aligned}
&|\text{at-ball}(?\text{ball}, \text{left})| = 0 \lor \\
&|\text{at-vader}(?\text{room}) \cap \{\text{left}\}| > 0 \land |\text{carry}(?\text{ball})| = 0 \lor \\
&|\text{at-vader}(?\text{room}) \cap \{\text{right}\}| > 0 \land |\text{carry}(?\text{ball})| > 0 \\
&\equiv \text{is-progress-state}
\end{aligned}
$$



Description Logic: Baader *et al.* [2003]  Formula: Ferber *et al.* [2022a]

# Identifying Progress States

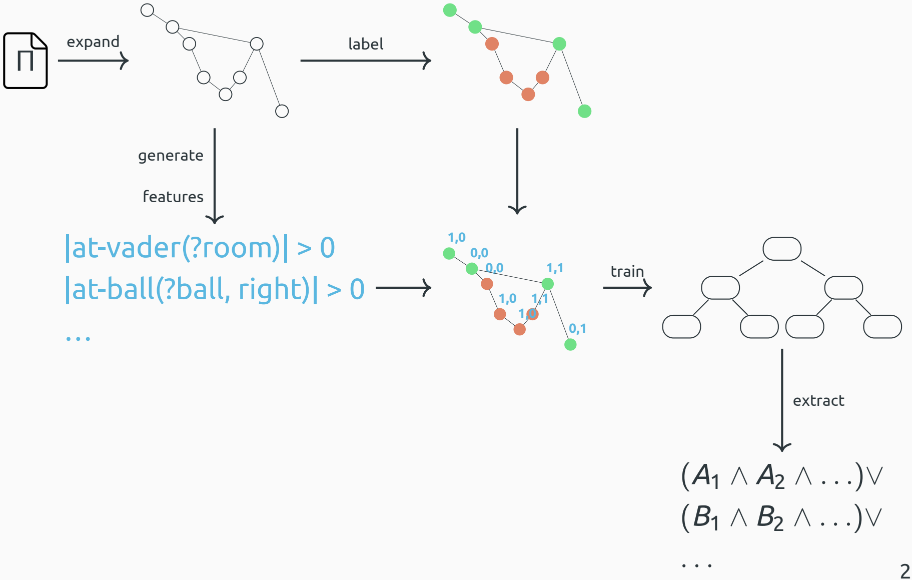


Ferber *et al.* [2022a]

# Expansions GBFS+$h^{FF}$

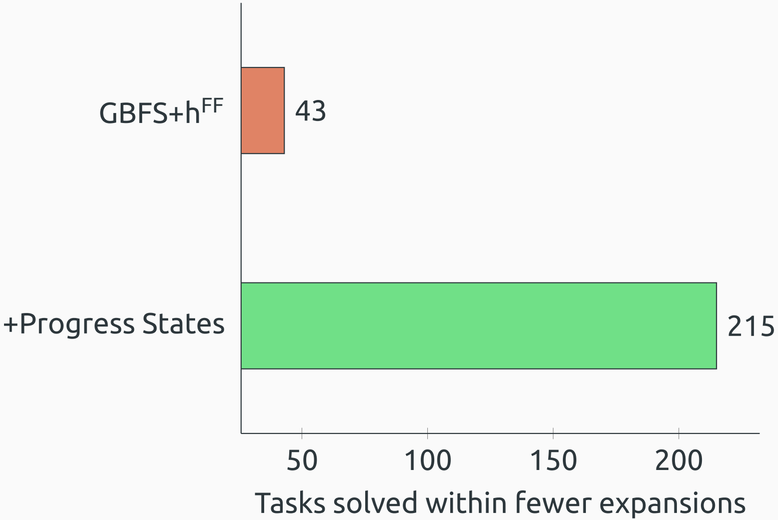


Ferber *et al.* [2022a]

# Generalized Bench Transition System

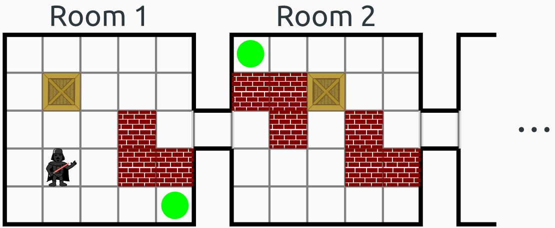

# Generalized Bench Transition System

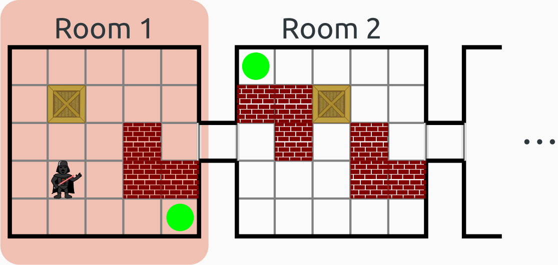

# Generalized Bench Transition System

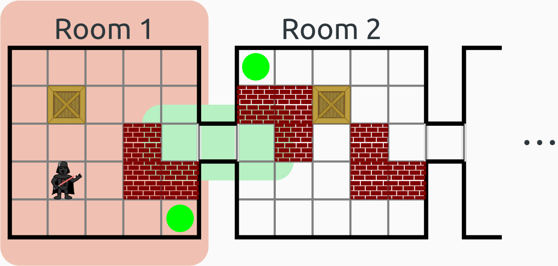

# Generalized Bench Transition System

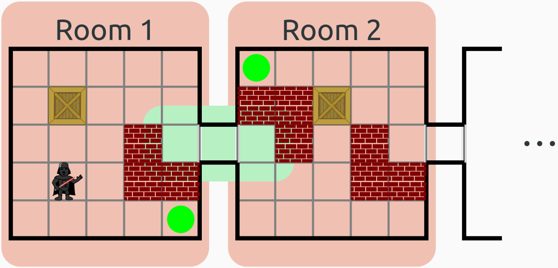

# Generalized Bench Transition System

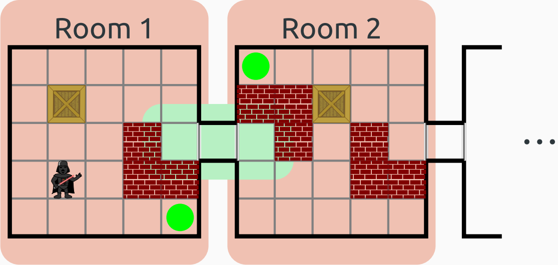


M: current room tidy &
other room untidy
G: in untidy room
A: crate moved

M: untidy crate in current room
G: only tidied crates in current room
A: change room

M: all rooms tidy
G: True
A: False

# Co-Authors & Colleagues

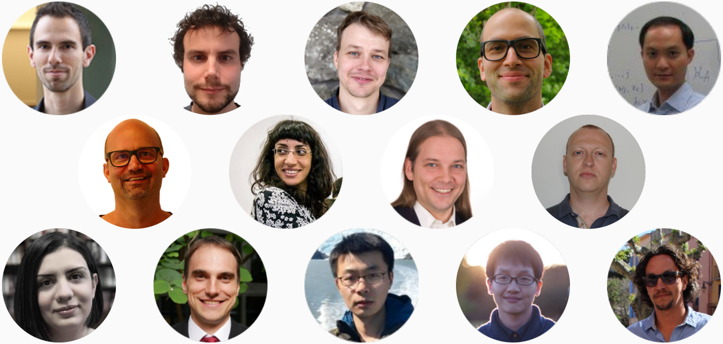

# Conclusion

**Heuristics**

highly complementary

can exceed model-based heuristics

**Online Portfolios**

state of the art

**State Space Topologies**

automatically generated

improves search

split into subproblems

# Construct GBTS

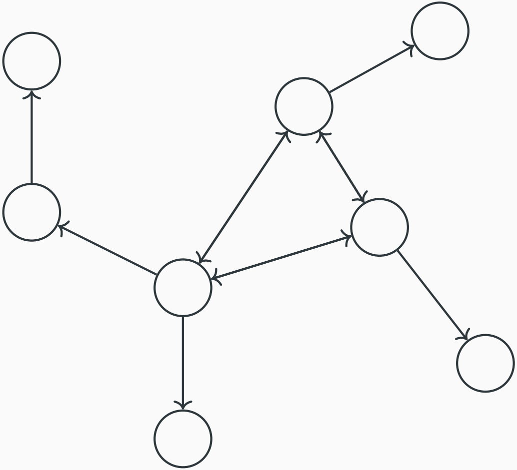

# Construct GBTS

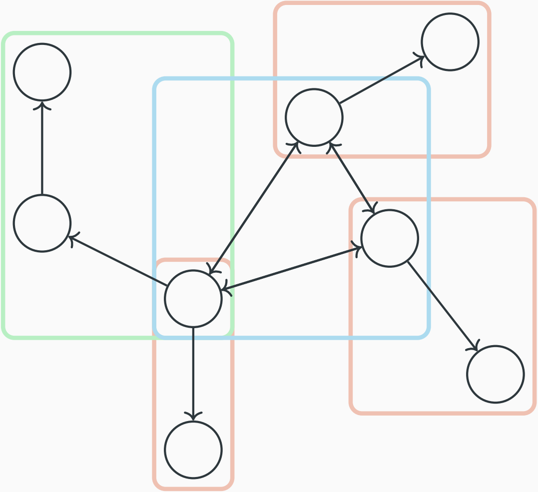

# Construct GBTS

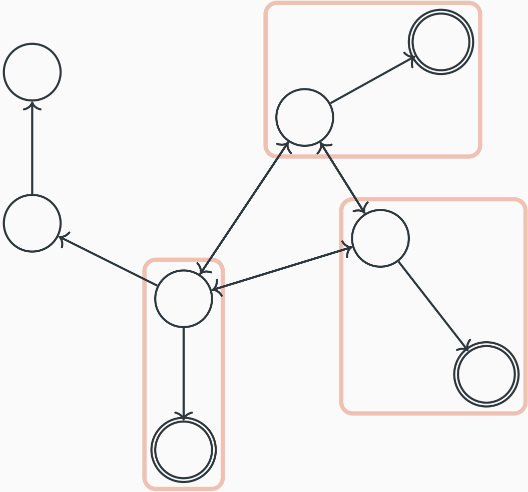

# Construct GBTS

Membership:

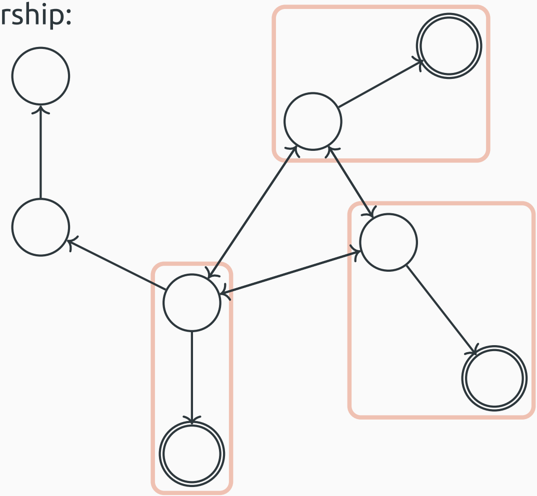

# Construct GBTS

Membership:

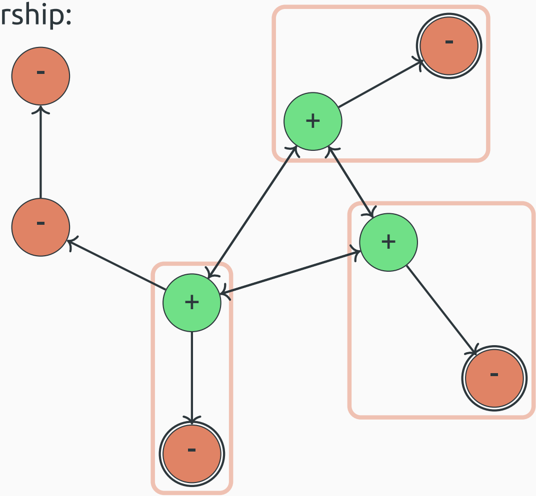

# Expansions with Progress States

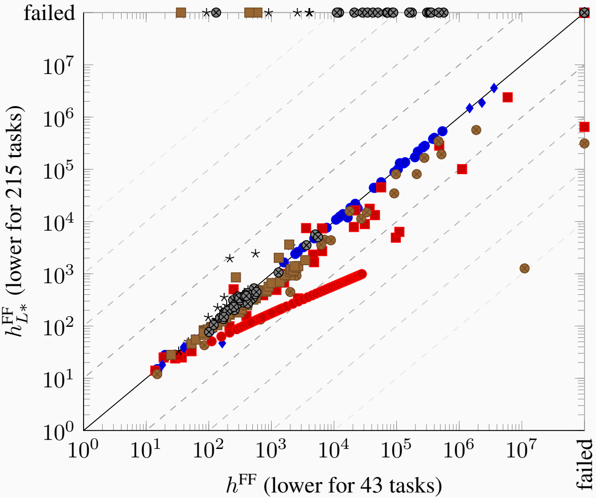


Ferber *et al.* [2022a]

# References

Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications.* Cambridge University Press, 2003.

Chris Fawcett, Mauro Vallati, Frank Hutter, Jörg Hoffmann, Holger Hoos, and Kevin Leyton-Brown. Improved features for runtime prediction of domain-independent planners. In Steve Chien, Alan Fern, Wheeler Ruml, and Minh Do, editors, *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling (ICAPS 2014)*, pages 355–359. AAAI Press, 2014.

# References ii


Patrick Ferber, Malte Helmert, and Jörg Hoffmann. Neural network heuristics for classical planning: A study of hyperparameter space. In Giuseppe De Giacomo, editor, *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, pages 2346–2353. IOS Press, 2020.

Patrick Ferber, Liat Cohen, Jendrik Seipp, and Thomas Keller. Learning and exploiting progress states in greedy best-first search. In Luc De Raedt, editor, *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI 2022)*, pages 4740–4746. IJCAI, 2022.

Patrick Ferber, Florian Geißer, Felipe Trevizan, Malte Helmert, and Jörg Hoffmann. Neural network heuristic functions for classical planning: Bootstrapping and comparison to other methods. In Sylvie Thiébaux and William Yeoh, editors, *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)*, pages 583–587. AAAI Press, 2022.

# References iii


Manuel Heusner, Thomas Keller, and Malte Helmert. Best-case and worst-case behavior of greedy best-first search. In Jérôme Lang, editor, *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pages 1463–1470. IJCAI, 2018.

Jörg Hoffmann and Bernhard Nebel. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.

Jörg Hoffmann. Where 'ignoring delete lists' works: Local search topology in planning benchmarks. *Journal of Artificial Intelligence Research*, 24:685–758, 2005.

Tengfei Ma, Patrick Ferber, Siyu Huo, Jie Chen, and Michael Katz. Online planner selection with graph neural networks and adaptive scheduling. In Vincent Conitzer and Fei Sha, editors, *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, pages 5077–5084. AAAI Press, 2020.

# References iv


Nir Pochter, Aviv Zohar, and Jeffrey S. Rosenschein. Exploiting problem symmetries in state-based planners. In Wolfram Burgard and Dan Roth, editors, *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*, pages 1004–1009. AAAI Press, 2011.

Silvan Sievers, Michael Katz, Shirin Sohrabi, Horst Samulowitz, and Patrick Ferber. Deep learning for cost-optimal planning: Task-dependent planner selection. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, pages 7715–7723. AAAI Press, 2019.

Silvan Sievers, Gabriele Röger, Martin Wehrle, and Michael Katz. Theoretical foundations for structural symmetries of lifted PDDL tasks. In Nir Lipovetzky, Eva Onaindia, and David E. Smith, editors, *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*, pages 446–454. AAAI Press, 2019.

# References v

Christopher Wilt and Wheeler Ruml. Speedy versus greedy search. In Stefan Edelkamp and Roman Barták, editors, *Proceedings of the Seventh Annual Symposium on Combinatorial Search (SoCS 2014)*, pages 184–192. AAAI Press, 2014.