

AI Planning

17. Comparing Heuristic Functions

h^{foo} vs. h^{bar} : What's the Difference Anyway?

Álvaro Torralba, Cosmina Croitoru



Winter Term 2018/2019

Thanks to Prof. Jörg Hoffmann for slide sources

Agenda

- 1 Introduction
- 2 The Compilability Framework
- 3 Example Proofs
- 4 A Walk Through the Zoo [for Reference]
- 5 Conclusion

“The Zoo”

All the wild (and admissible) animals we’ve learned about:

- h^1, h^2, h^3, \dots :
- h^+ :
- h^{\max} :
- $\{h^\alpha\}$:
- $\{h_L^{\text{LM}}\}$:
- Elementary delete relaxation landmark heuristics $\{h_{L+}^{\text{LM}}\}$:

What Do We Know About “The Zoo”?

Ignoring Deletes

 h^{\max} h^+

Abstractions

PDB

M&S

Critical Paths

 h^1 h^2 h^3

...

Landmarks

 $h_{L^+}^{\text{LM}}$ h_L^{LM}

What Do We *Want* to Know About “The Zoo”?

Motivation for comparing heuristics: It's a big zoo: $h^1, h^2, h^3, \dots, h^+, h^{\max}, \{h^\alpha \mid \alpha \text{ PDB}\}, \{h^\alpha \mid \alpha \text{ M\&S}\}, \{h_L^{\text{LM}}\}, \{h_{L+}^{\text{LM}}\}$.

- Given any one planning task, which one should we use? Are perhaps some heuristics dominated by others?
- Are all these differences meaningful anyway? Or can we “simulate” some framework using another framework?

Restriction: We consider **admissible heuristics only**.

- Of lower bounds h and h' , the bigger one is better.
- There is no similarly clear method to decide which one of two inadmissible heuristics (i.e., heuristics that are neither lower- nor upper-bounding) is “better”.

Our Agenda for This Chapter

- ② **The Compilability Framework:** We introduce, and discuss in detail, what it means for one family of heuristics to be able to “simulate” another family of heuristics.
- ③ **Example Proofs:** We give some concrete compilability and uncomplability proofs between particular families of heuristics as introduced in the previous chapters. In particular, we detail the aforementioned LM-cut heuristic, which was originally conceived as part of such a proof argument.
- ④ **A Walk Through the Zoo:** We briefly summarize the entire set of compilability results known about our “Zoo” of admissible heuristics, at this time.

Domination Between Heuristics

→ The only traditional means of comparing heuristic functions is by dominance:

Reminder:

→ **Chapter 7**

Definition (Domination). Let Π be a planning task, and let h and h' be admissible heuristics for Π . We say that h' *dominates* h if $h \leq h'$, i.e., for all states s in Π we have $h(s) \leq h'(s)$.

→ This is a very limited framework because it only allows us to compare **single heuristics**, as opposed to **classes of heuristics**. See next slide.

What About “The Zoo”?

(Red: Conclusions from standard “domination” concept.)

Ignoring Deletes

 h^{\max} h^+

Abstractions

PDB

M&S

Critical Paths

 h^1 h^2 h^3

...

Landmarks

 $h_{L^+}^{\text{LM}}$ h_L^{LM}

Domination Between *Classes* of Heuristics?

→ To compare classes of heuristics, which question should we ask?

Ⓐ **For every $h \in H$, do all $h' \in H'$ dominate h ?**

→ In “The Zoo”:

Ⓑ **For every $h \in H$, does there exist $h' \in H'$ that dominates h ?**

→ In “The Zoo”:

Ⓒ **For every partitioned sum $\sum_{i=1}^n h_i[c_i]$ of $h_i \in H$, does there exist a partitioned sum $\sum_{i=1}^n h'_i[c_i]$ of $h'_i \in H'$ that dominates $\sum_{i=1}^n h_i[c_i]$?**

→ Way to go! After all, no reason to ignore additivity here.

Some Simple Terminology

Start/finish: We will be considering compilations from a “start class” H into a “finish class” H' .

Reminder:

→ **Chapter 15**

$h[c_i]$ is h computed on Π with cost function replaced by c_i .

Closedness:

- A class H of heuristics is *closed* if $h \in H$ implies $h[c_i] \in H$ for any cost function c_i .
- We consider only closed H for the rest of this Chapter, allowing us to do partitionings *within* classes of heuristics.
- All of our families (Critical Paths, PDB, etc.) are closed: All our heuristics can be applied to any cost function.

The Compilability Framework

Definition (Compilability Between Heuristics). Let H and H' be closed classes of admissible heuristic functions. We say that H can be *compiled* into H' , written $H \preceq H'$, if there exists an algorithm with:

- (i) **Input:** Task Π , *state* s , heuristic $h \in H$ for Π .
- (ii) **Output:** Heuristics $h'_1, \dots, h'_k \in H'$ for Π , and cost partitioning c'_1, \dots, c'_k for Π , such that $h(s) \leq \sum_{i=1}^k h'_i[c'_i](s)$.
- (iii) **Runtime:** *Polynomial in the size of the input.*

If we can fix $k = 1$, we say that H can be compiled into *individual* H' , and write \preceq^i . If we can fix h'_1, \dots, h'_k and c'_1, \dots, c'_k for all states s in Π , we say that H can be *universally* compiled into H' , and write \preceq_u .

→ H can be compiled into H' if, given a *state* s and $h \in H$, we can *efficiently* obtain an at least as good lower bound *for* s by a partitioned sum of $h'_i \in H'$.

→ If we need only one h' , then the compilation is *individual*. If we can fix one combination of h' for all states of Π , then the compilation is *universal*.

“Heuristic h ” as an Algorithm Input/Output?

“**Data Structure h** ”: We overload h to denote (a) the heuristic function itself as usual, and (b) **some structure allowing to compute h in time polynomial in the size of that structure.**

- For **abstraction heuristics** with abstraction mapping α , the “data structure h ” is the **abstract state space Θ^α** .
- For h^m and h^{\max} , the “data structure h ” is the **task Π** itself.
- For $h_{L^+}^{\text{LM}}$, the “data structure h ” is the **action set L^+ and the task Π** itself.
- For h^* and h_L^{LM} , “data structure h ” is the original **state space**, for h^+ it is the **state space** of the delete-relaxed problem.
→ **Exponentially large in $\|\Pi\|$ as these h are hard to compute.**

→ Computational efficiency is important when examining how to compile heuristics into each other. Our compilation framework inputs, and outputs, representations allowing to compute h efficiently.

Questionnaire

Question!

Say $H = \{h^*\}$. Does there exist a family H' of heuristic functions into which H can be compiled?

(A): Yes

(B): No

The Compilability Framework: Remarks

→ H can be compiled into H' if, given a **state s** and $h \in H$, we can **efficiently** obtain an at least as good lower bound **for s** by a partitioned sum of $h'_i \in H'$.

→ If we need only one h' , then the compilation is **individual**. If we can fix one combination of h' for all states of Π , then the compilation is **universal**.

- The **per-state compilation** is needed because good cost partitionings are state-dependent (cf. **Chapter 15**).
- “Efficient” means that **we can find the “structure h'_i ” in polynomial time**. Without this restriction, the compilation would be useless (e.g., everything could be compiled into h^*).
- The special case \preceq_u^i is (B) on slide 11. Examples:

Compiling *Combinations* of $h \in H$

Theorem. Let H and H' be closed classes of admissible heuristic functions, where $H \preceq H'$. Then there exists an algorithm with:

- i) **Input:** Task Π , state s , $h_1, \dots, h_n \in H$, cost partitioning c_1, \dots, c_n .
- ii) **Output:** Heuristics $h'_1, \dots, h'_k \in H'$ for Π , and cost partitioning c'_1, \dots, c'_k , such that $\sum_{i=1}^n h_i[c_i](s) \leq \sum_{j=1}^k h'_j[c'_j](s)$.
- iii) **Runtime:** Polynomial in the size of the input.

Transitivity and Reflexivity

Corollary (Compilability is Transitive). *Let H , H' , and H'' be closed classes of admissible heuristic functions, where $H \preceq H'$ and $H' \preceq H''$. Then $H \preceq H''$.*

Proposition (Compilability is Reflexive). *Let H be a closed class of admissible heuristic functions. Then $H \preceq H$.*

Proof.

→ We can think of “ \preceq ” like “ \leq for classes of heuristics”.

Notation: If $H \preceq H'$ and $H' \preceq H$, we write $H \equiv H'$.

Questionnaire

Question!

Consider STRIPS planning, and $H = \{h^+\}$. Does there exist a family H' of heuristic functions into which H can be compiled?

(A): Yes

(B): No

Before We Begin

- In what follows, we'll look at a few example compilability and uncompileability proofs.
- I've basically selected some simple examples for illustration, and one example because it is interesting and important in practice.
- All these examples are from [Helmert and Domshlak (2009)].
- We'll ignore from now on the heuristics that cannot be computed in time polynomial in $\|\Pi\|$: h^+ and h_L^{LM} .
- By definition, no h can be compiled into such heuristics, unless h itself already uses an exponentially large "data structure h ".
- As for compiling such heuristics into other h , cf. slide 15 and slide 19.

Learning Things About “The Zoo”, Part I

Example Proof: h_{L+}^{LM} to PDB

Theorem ($h_{L+}^{LM} \not\leq \text{PDB}$). Let H be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} , and let H' be the family of abstraction heuristics h^α where α is a projection. Then $H \not\leq H'$.

Example Proof: h_{L+}^{LM} to M&S

Theorem ($h_{L+}^{LM} \preceq^i \text{M\&S}$). Let H be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} , and let H' be the family of abstraction heuristics h^α where α is constructed by merge-and-shrink abstraction. Then $H \preceq^i H'$.

Learning Things About “The Zoo”, Part II

Example Proof: h_{L+}^{LM} to h^1

Theorem ($h_{L+}^{LM} \preceq_u^i h^1$). Let H be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} , and let $H' = \{h^1\}$ be the family that contains only h^1 . Then $H \preceq_u^i H'$.

Learning Things About “The Zoo”, Part III

Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



Example Proof: h^1 to h_{L+}^{LM}


Theorem ($h^1 \preceq h_{L+}^{LM}$). Let $H = \{h^1\}$ be the family that contains only h^1 , and let H' be the family of elementary delete relaxation landmark heuristics h_{L+}^{LM} . Then $H \preceq H'$.

Proof. LM-cut!

Reminder: (Rough Intuition!)

→ **Chapter 14**

Get L as a *cut* between the initial state and the “0-cost goal zone”;
reduce the cost of each action in L by $\min_{a \in L} c(a)$; iterate.



LM-cut Algorithm in Detail

Assume: $I = \{i\}$; $G = \{g\}$; for all a we have $|pre_a| \geq 1$. (Without loss of generality: this can be achieved with simple transformations.)

Assume: $h^1(s) \neq \infty$; else, return $L_1^+ = \emptyset$, so that $h_{L_1^+}^{LM} = \infty$ (cf. **Chapter 14**).

Input: Planning task Π , state s .

Output: Delete relaxation landmarks L_i^+ for s along with a cost partitioning c_i .

$i := 1$

loop do

 Compute $h^1[c](s)$; **if** $h^1(s) = 0$ **then stop**

 For each a , select $p_a \in pre_a$ with maximal $h^1(s, \{p\})$

 Build a graph L whose nodes are the facts, and

 with a labeled arc $p_a \xrightarrow{a} q$ whenever $q \in eff_a$

$L_i^+ :=$ the labels of a cut in L between i and

 the part of L from which g can be reached with 0 cost

$c_i(a) := \min_{a \in L_i^+} c(a)$ for the actions in L_i^+ , and $c_i(a) := 0$ elsewhere

 For $a \in L_i^+$, subtract $\min_{a \in L_i^+} c(a)$ from $c(a)$

$i := i + 1$

LM-cut Algorithm: Detailed Example

LM-cut Algorithm: Why Does This Work?

→ We still need to prove the slide 26 Theorem:

Proof. Assume that $h^1(s) \neq \infty$ (else we return the empty landmark).

Say LM-cut stops after k iterations. We need to prove that (a)

$\sum_{i=1}^k h_{L_i^+}^{\text{LM}}[c_i](s) \geq h^1(s)$, and (b) LM-cut runs in time polynomial in $\|\Pi\|$.

If $h^1(s) = 0$, both (a) and (b) hold trivially. Else, consider the first iteration of the LM-cut algorithm, characterized by the sum $h(s) + h'(s)$, where $h(s) = h_{L_1^+}^{\text{LM}}(s)$ with cost function c_1 , and $h'(s)$ is $h^1(s)$ with cost function $c - c_1$. We prove that $h^1(s) \leq h(s) + h'(s)$. Applying this recursively to the second iteration and so forth proves (a); as the number of 0-cost actions increases strictly in each iteration, and as each iteration takes time polynomial in $\|\Pi\|$, we have (b).

As $h(s) = h_{L_1^+}^{\text{LM}}(s) = \min_{a \in L_1^+} c(a)$, proving $h^1(s) \leq h(s) + h'(s)$ comes down to proving that **reducing action costs c by c_1 does not decrease $h^1(s)$ by more than $\min_{a \in L_1^+} c(a)$** . Observe that $h^1(s)$ corresponds exactly to cheapest paths in the graph G . Such paths enter the “0-cost goal zone” (blue on previous slide) exactly once, hence subtracting c_1 (the step into the “0-cost goal zone”) reduces their cost by at most $\min_{a \in L_1^+} c(a)$ (the cost of that step). This concludes the proof.

Questionnaire

Observe: We designed LM-cut to find a partitioned-sum LM heuristic that dominates h^1 . The very first thing that LM-cut does is “Compute $h^1(s)$ ”.

Question!

Does this mean that LM-cut is completely useless?

(A): Yes

(B): No

What We Learned About “The Zoo”

Further Known Results

Corollaries of what we've seen:

- Since $h^{\max} = h^1$ and $h^1 \equiv h_{L+}^{\text{LM}}$: $h^{\max} \equiv h_{L+}^{\text{LM}}$.
- Since $h^1 \equiv h_{L+}^{\text{LM}}$, and $h_{L+}^{\text{LM}} \not\leq \text{PDB}$: $h^1 \not\leq \text{PDB}$, and since $h^{\max} = h^1$: $h^{\max} \not\leq \text{PDB}$. Further, since $h^1 \leq h^m$: $h^m \not\leq \text{PDB}$.
- Since $h^1 \equiv h_{L+}^{\text{LM}}$, and $h_{L+}^{\text{LM}} \preceq^i \text{M\&S}$: $h^1 \preceq^i \text{M\&S}$, and since $h^{\max} = h^1$: $h^{\max} \preceq^i \text{M\&S}$.

Additional results:

- $\text{PDB} \not\leq h^m$ [Helmert and Domshlak (2009)].
- Since $\text{PDB} \preceq \text{M\&S}$: $\text{M\&S} \not\leq h^m$, and since $h^1 \leq h^m$ and $h^1 \equiv h_{L+}^{\text{LM}}$: $\text{M\&S} \not\leq h_{L+}^{\text{LM}}$.
- $h^m \not\leq \text{M\&S}$ [Helmert *et al.* (2014)].

The Full “Zoo”, As We Know It

Summary

- Algorithms for computing admissible heuristics can be compared by the strength of the lower bounds they can produce efficiently.
- A family H of admissible heuristics can be **compiled** into family H' , written $H \preceq H'$, if for every state s and $h \in H$ we can efficiently construct a partitioned sum of heuristics $h' \in H'$ so that $h \leq \sum h'[c]$.
- If $H \preceq H'$, then partitioned sums over H can be compiled into partitioned sums over H' as well. Compilability is transitive and reflexive.
- Delete relaxation landmarks cannot be compiled into PDBs, but can be compiled into merge-and-shrink; they are equivalent to h^1 (and thus to h^{\max} as well). All other results are negative.

Remarks

- Compilability of delete relaxation landmarks into merge-and-shrink but not PDBs is yet another indication of the superiority of merge-and-shrink.
- The most surprising result was that we can efficiently find delete relaxation landmarks compiling h^1 . The implementation, LM-cut, is currently the strongest admissible heuristic available. In particular, it yields *much* better heuristic values than h^1 (i.e., it dominates h^1 in theory, and vastly outperforms it in practice).
- All other results being negative show non-domination, i.e., the respective techniques, for example critical path heuristics and merge-and-shrink, are incomparable (none dominates the other).
- (Practice is another story anyway.)

Reading

- *Landmarks, Critical Paths and Abstractions: What's the Difference Anyway?* [Helmert and Domshlak (2009)].

Available at:

<http://ai.cs.unibas.ch/papers/helmert-domshlak-icaps2009.pdf>

Content: The only publication on the compilability framework, to date. Briefly introduces the framework against the (fixed) background of admissible sums via cost partitioning, and delete relaxation landmark heuristics. Provides detailed proofs for the compilations/lack thereof between landmarks and PDBs/M&S, and between landmarks and h^1 . Briefly mentions non-compilability of PDBs into h^m . Runs experiments showing the excellent performance of the LM-cut heuristic, that implements the compilation of h^1 into landmarks.

References I

Malte Helmert and Carmel Domshlak. Landmarks, critical paths and abstractions: What's the difference anyway? In Alfonso Gerevini, Adele Howe, Amedeo Cesta, and Ioannis Refanidis, editors, *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*, pages 162–169. AAAI Press, 2009.

Malte Helmert, Patrik Haslum, Jörg Hoffmann, and Raz Nissim. Merge & shrink abstraction: A method for generating lower bounds in factored state spaces. *Journal of the Association for Computing Machinery*, 61(3), 2014.