AI Planning

14. Landmark Heuristics

It's a Long Way to the Goal, But How Long Exactly?
Part IV: *Ticking Off the Items On a To-Do List*

Álvaro Torralba, Cosmina Croitoru

SAARLAND UNIVERSITY
COMPUTER SCIENCE

Winter Term 2018/2019

Thanks to Prof. Jörg Hoffmann for slide sources

Agenda

We Need Heuristic Functions!

→ Landmarks (LMs) are a method to relax planning tasks, and thus automatically compute heuristic functions h.

We cover the 4 different methods currently known:

- Critical path heuristics: Done. → **Chapter 8**
- Delete relaxation: Basically done. → **Chapters 9 and 10**
- Abstractions: Done. → **Chapters 11–13**
- Landmarks. → **This Chapter**

→ Each of these have advantages and disadvantages. (We will do a formal comparison in **Chapter 17**.)

→ LM heuristics research yielded lots of exciting results since 2009. They boost the performance of satisficing planning when combined with delete relaxation heuristics, *and* they are among the most successful methods for computing lower-bound estimators.

Landmarks in a Nutshell

Problem: Bring key B to position 1.

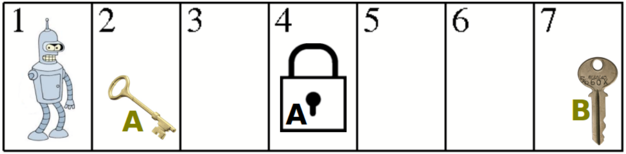

Landmarks:

- robot-at-2, robot-at-3, robot-at-4, robot-at-5, robot-at-6, robot-at-7.
- lock-open, have-key-A, have-key-B, . . .

→ A landmark is something that every plan for the task must satisfy at some point.

- Find landmarks in a pre-process to planning.
- Heuristic value(state) := number of yet un-achieved landmarks. ("Number of open items on the to-do list")

Before We Begin

- Landmarks were originally introduced as a method for problem decomposition [Hoffmann *et al.* (2004)].
- They traditionally come with a colorful variety of concepts defining orderings between them.
- Here we only discuss the generation of heuristic functions.
- We consider only the two most canonical forms of landmarks, and we do not cover LM orderings at all.
- Traditionally, LMs are mostly formulated in STRIPS; we'll do FDR (it doesn't really make a difference here). Remember that "facts" p in FDR are variable/value pairs.

Our Agenda for This Chapter

2. **Landmarks:** We start by defining the two forms of landmarks we will consider, and we discuss their connections and differences.
3. **Landmark Heuristics:** We specify how to turn landmarks (assuming they are provided as input) into heuristic functions. We introduce a notion of orthogonality which implies additivity.
4. **Detecting Landmarks:** We state that, in general, detecting landmarks is computationally hard, and we introduce and discuss the most commonly used approximation methods.

Fact Landmarks

"Something that every plan must satisfy at some point." **Take 1:**

Definition (Fact Landmark). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let s be a state. A fact p is a fact landmark for s if $p \notin s$, and for every plan $\langle a_1, \dots, a_n \rangle$ for s, there exists t so that $p \in s[\![\langle a_1, \dots, a_t \rangle]\!]$.*

→ A fact landmark is a variable value that is currently false, but that must become true at some point along every plan.

→ We'll often use "LM" for "Landmark".

→ Any spontaneous ideas for facts that will always be landmarks?

Where Fact Landmarks Fail

FindPath example: Actions $move(X,Y)$ pre X eff Y; init A, goal E.

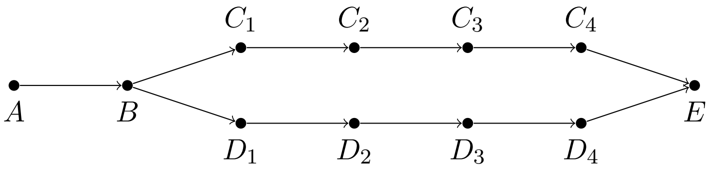

$\rightarrow$ Fact LMs for I?

To the rescue: disjunctive landmarks!

Disjunctive Action Landmarks

"Something that every plan must satisfy at some point." **Take 2:**

Definition (Disjunctive Action Landmark). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let s be a state. A set $L \subseteq A$ is a disjunctive action landmark for s if every plan for s contains an action L is minimal if there exists no $L' \subsetneq L$ that is a disjunctive action landmark for s.*

→ A disjunctive action LM is a set of actions at least one of which must occur in every plan. The LM is minimal if it contains no unnecessary actions.

Terminology: The action set induced by a fact p is $L(p) := \{a \in A \mid p \in \mathit{eff}_a\}$.

Proposition (Fact LMs Induce Disjunctive Action LMs). *Let Π be an FDR planning task, let s be a state, and let p be a fact landmark for s. Then $L(p)$ is a disjunctive action landmark for s.*

Proof. Since p must become true at some point, it must be in an action effect.

→ Is $L(p)$ always minimal?

Questionnaire

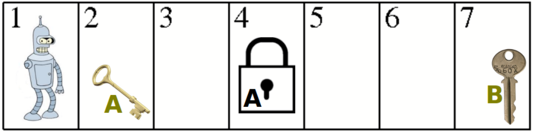

Fact landmarks p: robot-at-2, robot-at-3, robot-at-4, robot-at-5, robot-at-6, robot-at-7, lock-open, have-key-A, have-key-B.

Actions: MoveXY (*pre* robot-at-X[, lock-open for $Y = 4$]; *eff* robot-at-Y); PickXY (*pre* robot-at-X, key-Y-at-X; *eff* have-key-Y); DropXY (*pre* robot-at-X, have-key-Y; *eff* key-Y-at-X); OpenLockX for $X \in \{3, 5\}$ (*pre* robot-at-X, have-key-A; *eff* lock-open).

Question!

How many of the 9 fact landmarks p induce disjunctive action LMs $L(p)$ of size $|L(p)| > 1$? (And how many of the $L(p)$ with $|L(p)| > 1$ are minimal?)

(A): 0

(B): 7

(C): 8

(D): 9

Induced vs. All Disjunctive Action Landmarks

FindPath example: Actions $move(X,Y)$ pre X eff Y; init A, goal E.

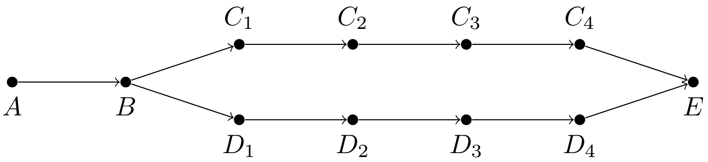

→ Fact LMs for I:

→ Disjunctive action LMs for I induced by these:

→ Minimal disjunctive action LMs for I **not** induced by these?

→ *Some* disjunctive action LMs are induced by fact LMs; most of them aren't.

→ Note the difference in the possible numbers of fact/disjunctive action LMs.

Elementary Landmark Heuristics

Definition (Elementary Landmark Heuristic). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task with state space $\Theta_\Pi = (S, A, T, I, G)$, and let $L \subseteq A$. The elementary landmark heuristic h_L^{LM} for Π given L is the function $h_L^{\mathsf{LM}} : S \mapsto \mathbb{R}_0^+$ where $h_L^{\mathsf{LM}}(s) = \min\{c(a) \mid a \in L\}$ if L is a disjunctive action landmark for s, and $h_L^{\mathsf{LM}}(s) = 0$ otherwise.*

→ If L is indeed a landmark, the elementary landmark heuristic given L returns the cost of the cheapest action in L; otherwise, it returns 0.

Remarks:

- h_L^{LM} is just a formal vehicle to elegantly express the goal distance estimates derived from LMs in terms of the heuristic functions framework.
- It has to be "min" over L, not "max" or "sum": intended meaning of L is that the planner may choose which action to use. Neither sum'ing nor max'ing would be admissible.
- If L is induced by a fact landmark p, this just means to "account for the cheapest action that achieves p".

Elementary Landmark Heuristics are Admissible

Theorem (h^{LM} is Admissible). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let $L \subseteq A$. Then h_L^{LM} is consistent and goal-aware, and thus also admissible and safe.*

Proof.

Questionnaire

Question!

Say s is a dead-end state. What are the (a) fact landmarks and (b) disjunctive action landmarks for s?

Question!

Say s is a dead-end state. Can $h^{\mathsf{LM}}_L(s)$ return ∞?

And Now?

Question!

Is h_L^{LM} a high-quality heuristic function?

(A): Yes. (B): No.

Orthogonal Landmarks

Terminology. $L_1, \dots, L_k \subseteq A$ are orthogonal if $L_i \cap L_j = \emptyset$ for $i \neq j$.

Theorem (The Sum of Orthogonal h^{LM} is Admissible). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let $L_1, \dots, L_k \subseteq A$ be orthogonal. Then $\sum_{i=1}^{k} h^{\mathsf{LM}}_{L_i}$ is consistent and goal-aware, and thus also admissible and safe.*

Proof.

The Canonical Landmark Heuristic

Terminology. The compatibility graph for $\mathcal{C} = \{L_1, \dots, L_n\}$ has vertices L_i and an arc (L_i, L_j) iff $L_i \cap L_j = \emptyset$.

Definition (Canonical Heuristic). *Let Π be an FDR planning task, let $\mathcal{C} = \{L_1, \dots, L_n\}$ be a collection of action subsets, and let $cliques(\mathcal{C})$ be the set of all maximal cliques in the compatibility graph for $\mathcal{C}$. Then the canonical heuristic $h^{\mathcal{C}}$ for $\mathcal{C}$ is defined as $h^{\mathcal{C}}(s) = \max_{\mathcal{D} \in cliques(\mathcal{C})} \sum_{L_i \in \mathcal{D}} h^{\mathsf{LM}}_{L_i}(s)$.*

→ The canonical heuristic maximizes over all largest orthogonal subsets of our landmarks collection.

Remarks:

- To reduce overlaps, minimal disjunctive action LMs are desirable.
- $h^{\mathcal{C}}$ is the best possible admissible heuristic we can derive from $\mathcal{C}$ *using the orthogonality criterion*. Despite this, on slide 22, we get $h^{\mathcal{C}} = 1$.
- Better heuristics can be obtained using cost partitioning or hitting sets (→ Chapter 15).

Questionnaire

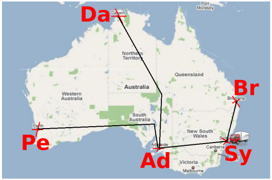

- Variables: $at : \{Sy, Ad, Br, Pe, Da\}$; $v(x) : \{T, F\}$ for $x \in \{Sy, Ad, Br, Pe, Da\}$.
- Actions: $drive(x, y)$ where x, y have a road.
- Costs: $Sy \leftrightarrow Br : 1$, $Sy \leftrightarrow Ad : 1.5$, $Ad \leftrightarrow Pe : 3.5$, $Ad \leftrightarrow Da : 4$.
- Initial state: $at = Sy, v(Sy) = T, v(x) = F$ for $x \neq Sy$.
- Goal: $at = Sy, v(x) = T$ for all x.

Induced by fact LMs: $\{drive(Ad, Pe)\}$, $\{drive(Ad, Da)\}$, $\{drive(Sy, Br)\}$, $\{drive(Pe, Ad),\ drive(Da, Ad),\ drive(Sy, Ad)\}$.

Additional disjunctive action LMs: $\{drive(Ad, Sy),\ drive(Br, Sy)\}$; $\{drive(Pe, Ad)\}$, $\{drive(Da, Ad)\}$, $\{drive(Sy, Ad)\}$.

Question!

Canonical heuristic $h^{\mathcal{C}}(I)$ from these?

Questionnaire, ctd.

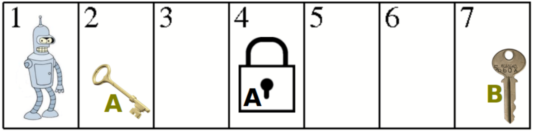

Fact landmarks p**:** robot-at-2, robot-at-3, robot-at-4, robot-at-5, robot-at-6, robot-at-7, lock-open, have-key-A, have-key-B. (Unit-cost actions)

Actions: MoveXY (*pre* robot-at-X[, lock-open for $Y = 4$]; *eff* robot-at-Y); PickXY (*pre* robot-at-X, key-Y-at-X; *eff* have-key-Y); DropXY (*pre* robot-at-X, have-key-Y; *eff* key-Y-at-X); OpenLockX for $X \in \{3, 5\}$ (*pre* robot-at-X, have-key-A; *eff* lock-open).

Question!

Considering the collection of disjunctive action LMs $L(p)$ **induced by these** p**, what is the value of the canonical heuristic** $h^{\mathcal{C}}$**?**

(A): 6	(B): 7
(C): 8	(D): 9

Elementary Landmark Heuristics in Practice (Up Next!)

$$“h^{\mathsf{LM}}_L(s) = \min\{c(a) \mid a \in L\} \text{ if } L \text{ is a disjunctive action landmark for } s, \text{ and } h^{\mathsf{LM}}_L(s) = 0 \text{ otherwise.”}$$

→ So will we keep L fixed, and check for every search state s whether or not it's a LM? No, because checking LMs is expensive. Instead, we design “landmark generation” algorithms, which guarantee to produce *only* LMs, but which do not guarantee to produce *all* LMs.

And then:

- (A) Offline generation, online update: Generate LMs $L_1, \dots, L_n$ for the initial state once before planning begins. Maintain flags throughout search to remember which ones have not been achieved yet.
- (B) Online generation: Generate LMs $L_1, \dots, L_n$ individually for each s.

But How to *Detect* those Landmarks in the First Place?

→ How to obtain a collection of disjunctive action landmarks?

Theorem (Checking Landmarks is Hard). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let s be a state. It is* **PSPACE**-*complete to decide whether or not a fact p is a fact landmark for s, and it is* **PSPACE**-*complete to decide whether or not an action set $L \subseteq A$ is a disjunctive action landmark for s.*

Proof. By a reduction from PlanEx. Given the task $\Pi = (V, A, c, I, G)$ for which we need to decide PlanEx, we construct $\Pi' := (V \cup \{x\}, A \cup \{a_1, a_2\}, c', I \cup \{x = 0\}, G)$ by introducing a new variable x with domain $\{0, 1\}$ as well as two new actions a_1, a_2 of which a_1 sets x from 0 to 1, and a_2 has precondition $x = 1$ and effect G. (We obtain c' from c by assigning arbitrary costs to a_1, a_2.)

Then

So is all lost?

→ How to obtain a collection of disjunctive action landmarks?

Answer: "It is **PSPACE**-complete to decide whether or not a fact p is a fact landmark for s, and it is **PSPACE**-complete to decide whether or not an action set $L \subseteq A$ is a disjunctive action landmark for s."

Question!

So is all lost?

(A): Yes. (B): No.

Detecting *Some* LMs, Take 1: Necessary Subgoals

Definition (Necessary Subgoals). *Let Π be an FDR planning task, and let s be a state. A fact p is a necessary subgoal in Π for s if $p \notin s$ and either:*

1. *$p \in G$; or*
2. *there exists a necessary subgoal q in Π for s so that $p \in \bigcap_{a \in A, q \in eff_a} pre_a$.*

→ Necessary subgoals are top-level goals plus shared preconditions.

("subgoal" here=singleton fact, not fact subset as for critical path heuristics.)

Proposition (Necessary Subgoals are Landmarks). *Let Π be an FDR planning task, and let s be a state. If p is a necessary subgoal in Π for s, then p is a fact landmark for s.*

Proof. By structural induction. The claim holds trivially for necessary subgoals of kind (i). For (ii), if q is a fact landmark for s, then q must be achieved at some point which by construction involves achieving p first.

Strategy: Given state s, detect necessary subgoals p_i for s by simple backchaining from the goal: start at $p \in G \setminus s$, then iteratively apply (ii) until no more new necessary subgoals are found.

Necessary Subgoals vs. Fact Landmarks

FindPath example: Actions $move(X,Y)$ pre X eff Y; init A, goal E.

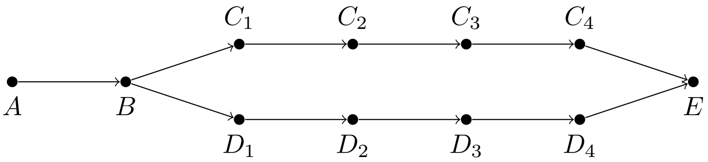

→ Fact landmarks for I?

→ Necessary subgoals for I?

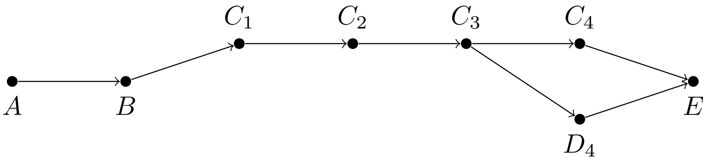

→ Fact landmarks for I?

→ Necessary subgoals for I?

Questionnaire

Problem: Bring key B to position 1.

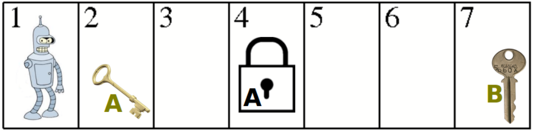

Actions: MoveXY (*pre* robot-at-X[, lock-open for $Y = 4$]; *eff* robot-at-Y); PickXY (*pre* robot-at-X, key-Y-at-X; *eff* have-key-Y); DropXY (*pre* robot-at-X, have-key-Y; *eff* key-Y-at-X); OpenLockX for $X \in \{3, 5\}$ (*pre* robot-at-X, have-key-A; *eff* lock-open).

Question!

What are the necessary subgoals for I in this planning task?

Detecting *Some* LMs, Take 2: Delete Relaxation LMs

Definition (Delete Relaxation LM). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let s be a state. A fact p [respectively an action set $L \subseteq A$] is a delete relaxation landmark for s if $p \notin s$, and for every relaxed plan $\langle a_1^+, \dots, a_n^+ \rangle$ for s, there exists t so that $p \in s[\![\langle a_1^+, \dots, a_t^+ \rangle]\!]$ [respectively so that $a_t \in L$].*

Proposition (Checking Delete Relaxation LMs is Easy). *Let $\Pi = (V, A, c, I, G)$ be an FDR planning task, and let s be a state. It can be decided in polynomial time whether or not a fact p, respectively an action set L, is a delete relaxation landmark for s.*

Proof.

Detecting Delete Relaxation LMs

How to detect delete relaxation fact LMs?

- How to find all? For every fact, run test on previous slide.
- Not such a good idea in practice: Relaxed planning is polynomial time but not dirt-cheap, and there may be 100s–1000s of facts.
- A direct method computes all "causal" delete relaxation fact landmarks by a fixed point computation [Keyder *et al.* (2010)].

How to detect delete relaxation disjunctive action LMs?

- How to find all? For every $L \subseteq A$, run test on previous slide.
- Completely useless idea in practice: Exponentially many L.
- **Vanilla solution:** Use $L(p)$ induced by delete relaxation fact LM p.
- **Advanced solution LM-cut:** [Helmert and Domshlak (2009)] 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. We'll give details in **Chapter 17**; illustration see next slide.

Detecting Action LMs: Fact-Induced vs. LM-cut

Propagating Landmarks

Remember? *"Heuristic value(state) := number of yet un-achieved landmarks (number of open items on the to-do list)."*

→ Here's how to "maintain the to-do list":

Proposition (Propagating Landmarks). *Let* Π *be an FDR planning task, let* L *be a disjunctive action LM for* I*, and let* s *be a state. If* $s = I[\![\vec{a}]\!]$ *where* $\vec{a}$ *does not use any action from* L*, then* L *is a disjunctive action LM for* s*.*

Strategy: Before search, detect disjunctive action landmarks for I. During forward search, maintain a flag for each L saying whether or not it was used yet. (For fact LMs p, the flag says whether p has already been true at some point.)

→ This is option (A) on slide 28. Re-computation for each s is option (B) on slide 28.

Delete Relaxation LMs: Properties

→ Necessary subgoals ⊆ delete relaxation landmarks ⊆ real landmarks.
→ Delete relaxation landmarks lower-bound h^+.

Precisely:

Proposition (Delete Relaxation LM Properties). *Let Π be an FDR planning task, and let s be a state. Then all of the following hold:*

1. *If p is a necessary subgoal for s, then p is a delete relaxation LM for s.*
2. *If p respectively L is a delete relaxation LM for s, then it is a LM for s.*
3. *If L is a delete relaxation LM for s, then $h^{\mathrm{LM}}_L(s) \leq h^+(s)$.*

Proof. (i): Same argument as in the proof that p is a LM. (ii): Every real plan for s is also a relaxed plan for s, so must use p respectively L. (iii): Trivial.

Necessary Subgoals vs. Delete Relaxation LMs vs. LMs

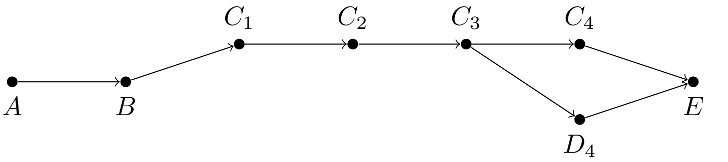

Fact LMs for I: B, C_1, C_2, C_3, E. **Necessary subgoals for I:** E.

→ Delete relaxation fact LMs for I?

And now: Say init $A, (x = 1)$; goal $E, (x = 1)$; $move(D_4, E)$ sets $x := 0$.

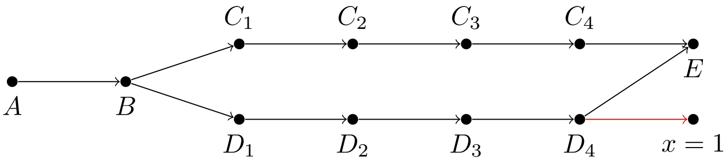

→ Fact LMs for I?

→ Delete relaxation fact LMs for I?

Questionnaire

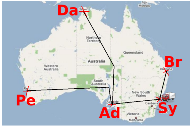

- Actions: $drive(x, y)$ where x, y have a road.
- Costs: $Sy \leftrightarrow Br : 1$, $Sy \leftrightarrow Ad : 1.5$, $Ad \leftrightarrow Pe : 3.5$, $Ad \leftrightarrow Da : 4$.
- Initial state: $at = Sy, v(Sy) = T, v(x) = F$ for $x \neq Sy$.
- Goal: $at = Sy, v(x) = T$ for all x.

Question!

Minimal disjunctive action LMs for I, and $h^{\mathcal{C}}(I)$?

Question!

Minimal delete relaxation disjunctive action LMs for I, and $h^{\mathcal{C}}(I)$?

Summary

- A landmark (LM) is something that every plan must satisfy. A fact LM must hold at some point on every plan, a disjunctive action LM is a set of actions one of which must be used by every plan.
- Fact LMs induce disjunctive action LMs; however, most disjunctive action LMs are not induced in this way.
- The elementary LM heuristic returns the cost of the cheapest action in a disjunctive action LM.
- Disjunctive action LMs are orthogonal if they are disjoint. Orthogonal elementary LM heuristics are summed admissibly in the canonical heuristic.

  → Stronger methods are cost partitioning and (even stronger) hitting sets, to be considered in the **Next Chapter**.
- Checking LMs is hard. Practical methods are sound but incomplete, detecting some LMs, namely necessary subgoals or delete relaxation LMs.
- Vanilla method: Detect (some) fact LMs and use the induced disjunctive action LMs. Much stronger method LM-cut: Iteratively cut between the initial state and the "0-cost goal zone".

Remarks: LM Definitions

Historical:

- Landmarks were originally just fact landmarks, and were introduced as a means to *decompose* the task: Find LMs for I in a pre-process, feed them one-by-one to the planner [Hoffmann *et al.* (2004)].

Technical:

- Various kinds of *orderings* between landmarks are in use: "A must be achieved (directly) before B", "A should be achieved before B or else we would need to delete B and re-achieve it after A", ...
- Instead of just facts, we can use arbitrary propositional formulas ϕ over the facts (or even quantification over PDDL objects).
- If ϕ is a disjunction of facts, then that corresponds very closely to disjunctive action landmarks.
- I've chosen the two particular notions as presented because the "vanilla method" to compute landmark heuristics is by considering the disjunctive action landmarks induced by the fact landmarks.

Remarks: LM Heuristics

Historical:

- The idea to generate heuristics based on landmarks was first conceived by [Zhu and Givan (2003)], never properly published and forgotten all about.
- The (basic) idea was re-discovered by the authors of LAMA [Richter *et al.* (2008); Richter and Westphal (2010)]. Which subsequently won two IPCs.
- Both the initial attempt and LAMA use non-admissible landmarks heuristics, basically counting the number of non-achieved fact landmarks (= summing up elementary landmark heuristics induced by fact landmarks, without ensuring independence).

Technical: (We will consider this in detail in the **Next Chapter**)

- The best admissible landmark heuristics in practice use cost partitioning [Karpas and Domshlak (2009); Helmert and Domshlak (2009)].
- One can use hitting sets over landmarks to obtain even better heuristics, but these tend to be too costly computationally [Bonet and Helmert (2010)].

Remarks: Detecting LMs

- The original LMs detection method found delete relaxation fact LMs, mostly the necessary subgoals [Hoffmann *et al.* (2004)].
- LAMA does that, plus additional methods based on domain transition graphs (cf. **Chapter 5**); it propagates LMs for I to avoid having to re-detect [Richter and Westphal (2010)].

- The first admissible LM heuristic uses the disjunctive action LMs induced by LAMA's fact LMs [Karpas and Domshlak (2009)].
- The first technique using disjunctive action LMs *not* induced by fact LMs was LM-cut [Helmert and Domshlak (2009)]. The iterated cut algorithm is done anew for every search state. Despite this, LM-cut is the most successful admissible LM heuristic in practice, to date.

Remarks: Planning Tools and Performance Using LMs

- Original use for probem decomposition gave reasonable speed-ups for FF and another satisficing heuristic search planner [Hoffmann *et al.* (2004)].
- LAMA [Richter and Westphal (2010)] introduced the idea to use both, a delete relaxation heuristic and a LM heuristic, in Fast Downward's dual-queue greedy best-first search framework. The LM heuristic improves performance significantly in some domains. LAMA won the 1st prizes for satisficing planners at IPC'08 and IPC'11.
- BJOLP [Karpas and Domshlak (2009); Domshlak *et al.* (2011)] uses admissible combination of disjunctive action LMs induced by fact LMs. It was part of the 1st-prize winning portfolio in the optimal track of IPC'11.
- LM-cut [Helmert and Domshlak (2009)] also uses admissible combination of disjunctive action LMs, but of more general such LMs not induced by fact LMs (cf. slide 38). It was part of the 1st-prize winning portfolio, and of the 2nd-prize winning portfolio, in the optimal track of IPC'11. It was the strongest single-heuristic optimal planner in IPC'11.

Reading

- *Ordered Landmarks in Planning* [Hoffmann *et al.* (2004)].

  Available at:

  http://fai.cs.uni-saarland.de/hoffmann/papers/jair04.pdf

  Content: The first paper on landmarks. Focusses mostly on ordering relations and problem decomposition; not directly relevant to the content of this chapter, but useful as a background read.

- *Sound and Complete Landmarks for And/Or Graphs* [Keyder *et al.* (2010)].

  Available at:

  http://www.dtic.upf.edu/~ekeyder/ECAI10_Landmarks.pdf

  Content: A nice and clean "modern" paper on landmarks. Contains, among other things, the fixed point algorithm computing all (causal) delete relaxation fact landmarks.

Reading, ctd.

- *Cost-Optimal Planning with Landmarks* [Karpas and Domshlak (2009)].

  Available at:

  `http://iew3.technion.ac.il/~dcarmel/Papers/Sources/ijcai09a.pdf`

  Content: The "alarm clock" waking LMs up to the modern age of cost-optimal planning! Admissible combination by going from fact LMs to disjunctive action LMs, optimal cost partitioning by compilation to linear programming **(→ Chapters 15–16)**, LM-A^* to handle this path-dependent heuristic.

  Recapitulates LAMA's heuristic along the way so may be used to get an idea of that one as well.

Reading, ctd.

- *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 LM-cut paper. As if LM-cut wasn't enough, it also introduces the comparison framework for admissible heuristics (→ **Chapter 17**).

Reading, ctd.

- *Strengthening Landmark Heuristics via Hitting Sets* [Bonet and Helmert (2010)].

  Available at:

  http://ai.cs.unibas.ch/papers/bonet-helmert-ecai2010.pdf

  Content: Introduces the idea to use minimum-cost hitting sets over disjunctive action landmarks, instead of combining elementary landmark heuristics. Shows that the minimum-cost hitting set over sufficiently large collections of delete relaxation disjunctive action landmarks is equal to h^+.

References I

Blai Bonet and Malte Helmert. Strengthening landmark heuristics via hitting sets. In Helder Coelho, Rudi Studer, and Michael Wooldridge, editors, *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI'10)*, pages 329–334, Lisbon, Portugal, August 2010. IOS Press.

Carmel Domshlak, Malte Helmert, Erez Karpas, Emil Keyder, Silvia Richter, Gabriele Röger, Jendrik Seipp, and Matthias Westphal. BJOLP: The big joint optimal landmarks planner. In *IPC 2011 planner abstracts*, pages 91–95, 2011.

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.

Jörg Hoffmann, Julie Porteous, and Laura Sebastia. Ordered landmarks in planning. *Journal of Artificial Intelligence Research*, 22:215–278, 2004.

References II

Erez Karpas and Carmel Domshlak. Cost-optimal planning with landmarks. In Craig Boutilier, editor, *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, pages 1728–1733, Pasadena, California, USA, July 2009. Morgan Kaufmann.

Emil Keyder, Silvia Richter, and Malte Helmert. Sound and complete landmarks for and/or graphs. In Helder Coelho, Rudi Studer, and Michael Wooldridge, editors, *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI'10)*, pages 335–340, Lisbon, Portugal, August 2010. IOS Press.

Silvia Richter and Matthias Westphal. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research*, 39:127–177, 2010.

Silvia Richter, Malte Helmert, and Matthias Westphal. Landmarks revisited. In Dieter Fox and Carla Gomes, editors, *Proceedings of the 23rd National Conference of the American Association for Artificial Intelligence (AAAI'08)*, pages 975–982, Chicago, Illinois, USA, July 2008. AAAI Press.

Lin Zhu and Robert Givan. Landmark extraction via planning graph propagation. In *ICAPS 2003 Doctoral Consortium*, pages 156–160, 2003.