

LAMA 2008 and 2011

Silvia Richter

NICTA
silvia.richter@nicta.com.au

Matthias Westphal

Albert-Ludwigs-Universität Freiburg
westpham@informatik.uni-freiburg.de

Malte Helmert

Albert-Ludwigs-Universität Freiburg
helmert@informatik.uni-freiburg.de

Abstract

LAMA is a propositional planning system based on heuristic search with landmarks. This paper describes two versions of LAMA that were entered into the 2011 International Planning Competition: the original LAMA as developed for the 2008 competition and a new re-implementation of LAMA that uses the latest version of the Fast Downward Planning Framework.

Landmarks are propositions that must be true in every solution of a planning task. LAMA uses a heuristic derived from landmarks in conjunction with the well-known FF heuristic. LAMA builds on the Fast Downward Planning System using non-binary (but finite domain) state variables and multi-heuristic search. A weighted A* search is used with iteratively decreasing weights, so that the planner continues to search for plans of better quality until the search is terminated. LAMA combines cost-to-goal and distance-to-goal estimates with the aim of finding good solutions using reasonable runtime.

Introduction

LAMA is a planning system based on heuristic state space search, in the spirit of FF (Hoffmann and Nebel 2001) and Fast Downward (Helmert 2006). It won the sequential satisficing track of the International Planning Competition in 2008 and is entered into this year's competition (2011) as a reference point.

A detailed description of LAMA can be found in a recent JAIR article by Richter and Westphal (2010). This paper gives a brief overview of LAMA, focusing mainly on its architectural structure, landmarks, and the differences between the 2008 and 2011 versions. This paper overlaps significantly with the IPC 2008 planner description of LAMA (Richter and Westphal 2008). Readers familiar with that paper may skip straight to Section *LAMA 2008 versus LAMA 2011*, where we describe LAMA 2011.

LAMA builds on the Fast Downward System, inheriting the general structure of Fast Downward, the translation of PDDL tasks with binary state variables to representations with finite-domain variables, and a search architecture that is able to exploit several heuristics simultaneously. One core feature of LAMA is the use of *landmarks* as a heuristic and for generating preferred operators. The landmark heuristic was introduced in a AAAI 2008 article (Richter, Helmert, and Westphal 2008). Other core features of LAMA are an

iterated search using restarts (Richter, Thayer, and Ruml 2010), and the combination of cost and distance estimates.

Structure of the Planner

LAMA consists of three separate programs:

1. the *translator* (written in Python),
2. the *knowledge compilation* module (written in C++), and
3. the *search engine* (also written in C++).

To solve a planning task, the three programs are called in sequence; they communicate via text files.

Translator

The purpose of the *translator* is to transform the planner input, specified in the propositional fragment of PDDL (including ADL features and derived predicates, but not the preferences and constraints introduced for IPC 2006), into a finite-domain state representation similar to the SAS⁺ formalism (Bäckström and Nebel 1995).

The main components of the translator are an efficient grounding algorithm for instantiating schematic operators and axioms and an invariant synthesis algorithm for determining groups of mutually exclusive facts. Such fact groups are consequently replaced by a single finite-domain state variable encoding *which* fact (if any) from the group is satisfied in a given world state.

The groups of mutually exclusive facts found during translation serve an important purpose for determining orders between landmarks. For more details on the translator component, see an article by Helmert (2009). We have modified this component only in some minor ways, including the extraction of all mutually exclusive facts (as mentioned above), the handling of action costs for IPC 2008, and some small enhancements.

Knowledge Compilation

Using the finite-domain task representation generated by the translator, the *knowledge compilation* module is responsible for building a number of data structures which play a central role in the subsequent landmark generation and search.

For example, *domain transition graphs* are produced which encode for each state variable the ways in which it

may change its value through operator applications. Furthermore, the knowledge compilation module constructs *successor generators* and *axiom evaluators*, data structures for efficiently determining the set of applicable actions in a given state of the planning task and for evaluating the values of derived state variables. We refer to Helmert (2006) for more detail on the knowledge compilation component.

Search Engine

Using the data structures generated by the knowledge compilation module, the *search engine* attempts to find a plan using heuristic search with some enhancements, such as the use of *preferred operators* (similar to helpful actions in FF) and *deferred heuristic evaluation*, which mitigates the impact of large branching factors in planning tasks with fairly accurate heuristic estimates (Richter and Helmert 2009). Deferred heuristic evaluation means that states are not evaluated upon generation, but upon expansion. States are thus not selected for expansion according to their own heuristic value, but according to that of their parent. If many more states are generated than expanded, this leads to a substantial reduction in the number of heuristic estimates computed, if at a loss of heuristic accuracy.

The rules of the 6th International Planning Competition (IPC 2008), for which LAMA was designed, suggest a type of search that takes plan quality into account. LAMA first runs a greedy best-first search, aimed at finding a solution as quickly as possible. Once a plan is found, it searches for progressively better solutions by running a series of weighted A* searches with decreasing weight. The cost of the best known solution is used for pruning the search, while decreasing the weight makes the search progressively less greedy, trading speed for solution quality (Richter, Thayer, and Ruml 2010).

The search engine is configured to use several heuristic estimators (namely, the FF heuristic and the landmark heuristic) within an approach called *multi-heuristic search* (Helmert 2006; Röger and Helmert 2010). This technique attempts to exploit strengths of the utilised heuristics in different parts of the search space in an orthogonal way. To this end, it uses separate open lists for each of the different heuristics as well as separate open lists for the preferred operators of each heuristic. Newly generated states are evaluated with respect to all heuristics, and added to all open lists (with the value estimate corresponding to the heuristic of that open list). When choosing which state to expand next, the search engine alternates between the different heuristics, and expands states from preferred-operator queues with higher priority than states from other queues.

Landmarks

Landmarks are variable assignments that must occur at some point in every solution plan. They were first introduced by Porteous, Sebastia and Hoffmann (2001) and later studied in more depth by the same authors (Hoffmann, Porteous, and Sebastia 2004). Consider the logistics example task in Fig. 1, where the goal is to transport the packet from location *B* to location *F*. In order to achieve the goal, the packet




Figure 1: A logistics task: transport packet *x* from *B* to *F*.

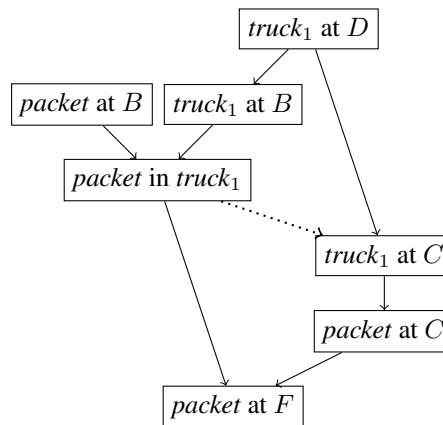


Figure 2: Partial landmark graph for the example task in Fig. 1, showing simple landmarks.




Figure 3: Domain transition graph for the packet from Fig. 1.

must be loaded onto $truck_1$ and unloaded at the airport C , in order to then be loaded into one of the planes p_1 or p_2 . Hence, the facts “*packet is on truck₁*”, and “*packet is at C*” are landmarks for this task. It is also possible to infer orders between landmarks, e. g. in this case, that “*packet is on truck₁*” must be true *before* “*packet is at C*”. The landmarks and orders can be stored in a directed graph called the landmark graph. For our example, a partial landmark graph is depicted in Fig. 2.

Our algorithm for finding landmarks and orderings between them is similar to the one by Porteous and Cresswell (2002), which is in turn based on the one by Hoffmann et al. Like Porteous and Cresswell we admit *disjunctive* landmarks (sets of propositions of which one needs to be true at some point), but we adapted the algorithm to the SAS⁺ setting and use domain transition graphs to derive further landmarks.

We find landmarks by backchaining from already known landmarks, starting with the goals (which are landmarks by definition, as they have to be true in every solution plan). For any given landmark L that is not true in the initial state, we consider the *shared preconditions* of its *possible first achievers*. Possible first achievers are those operators that a) have L as an effect, and b) can be possibly applied at the end of a partial plan (starting in the initial state) which has never made L true. Their shared preconditions are those propositions (if any) that are a precondition for each of the operators. Every such shared precondition must be true in order to reach L and is thus a landmark, which can be ordered before L .

Since it is PSPACE-hard to determine the set of *actual* first achievers of a landmark L , we use an over-approximation containing every operators that can *possibly* be a first achiever. By intersecting over the preconditions of more operators we do not lose correctness, though we may of course miss out on some landmarks. The approximation of first achievers of L is done with the help of a relaxed planning graph (RPG) (Hoffmann and Nebel 2001). During construction of the RPG we leave out any operator that would add L . When the relaxed planning graph levels out, its last set of facts is an over-approximation of the set of facts that can be achieved *before* L in the planning task; we denote it by $pb(L)$ (for *possibly before*). Any operator that is *applicable* given $pb(L)$ and achieves L is a possible first achiever of L .

We also create disjunctive sets of facts from the first achievers’ shared preconditions, such that a set contains one precondition fact from each first achiever. Hence, these sets form disjunctive landmarks. All facts in a disjunctive landmark must stem from the same predicate symbol, and we discard any sets of size greater than 4 in order to limit the number of possible sets.

We find further landmarks by exploiting the domain transition graphs (DTGs) generated by the knowledge compilation module. For each variable, a corresponding DTG has a node for each value that can be assigned to the variable, and arcs for possible transitions between them (where a transition can be achieved through operator application). For example, assume that the location of the packet in our example

is encoded with a state variable v . The DTG of v is depicted in Fig. 3. From its initial value B , the location of the packet can change to “in $truck_1$ ” (denoted as t_1 for short), from there to any of the locations A, B, C and D and so on.

Given a *simple* (i. e., non-disjunctive) landmark $L = \{v \mapsto l\}$ that is not part of the initial state, we consider the DTG of the landmark’s variable v . If there is a node that occurs on *every* path from the *initial state value* $s_0(v)$ of the variable to the *landmark value* l , then that node corresponds to a landmark value l' of the variable: We know that every plan achieving L requires that v take on the value l' , hence the fact $L' = \{v \mapsto l'\}$ can be introduced as a new landmark and ordered before L . To find these kinds of landmarks, we iteratively remove one node from the DTG and test with a simple graph algorithm whether $s_0(v)$ and l are still connected – if not, the removed node corresponds to a landmark. Nodes corresponding to assignments of v which are not in $pb(L)$ are removed from the DTG prior to this test, as they can only occur *after* B and do not have to be tested. However, we remember these nodes and if such a node is later found to be a landmark (e. g. by the backchaining procedure), we can introduce an ordering between B and the node.

Consider again the landmark graph of our example in Fig. 2. Most of the landmarks and orders in it can be found by the backchaining procedure even when restricting it to simple, i. e., non-disjunctive, landmarks, because the propositions are direct preconditions of their successor nodes in the graph. There are two exceptions: “*packet in truck₁*” and “*packet at C*”. These two landmarks are however found with the DTG method. The DTG in Fig. 3 shows immediately, that the package location must be both t_1 and C on any path from the initial state (where it has value B) to the goal F .

If we introduced another truck in the left city, the fact “*packet in truck₁*” would no longer be a landmark. However, using disjunctive landmarks we would get a landmark for the packet being inside one of the two trucks.

Inconsistencies found in the translating phase are exploited to determine further orders between the landmarks, using the definition of *reasonable orders* and the conditions proposed by Hoffmann et al. (2004). For example, the order depicted by a dotted line in Fig. 2 is such a reasonable order. For more details on how landmarks and their orders are derived, see the JAIR 2010 article (Richter and Westphal 2010).

The Landmark Heuristic

The LAMA planning system uses landmarks as a pseudo-heuristic. We estimate the goal distance of a state s by the number of landmarks l that still need to be achieved from s onwards. We estimate this number as $\hat{l} := n - m + k$, where n is the total number of landmarks, m is the number of landmarks that are *accepted*, and k is the number of accepted landmarks that are *required again*. A landmark B is accepted in a state s if it is true in that state and all landmarks ordered before B are accepted in the predecessor state from which s was generated. An accepted landmark remains accepted in all successor states. An accepted landmark is *required again* if it is not true in s and it is a direct precondition

of some landmark which is not accepted. Note that \hat{l} is not a proper state heuristic in the usual sense, as its definition depends on the way s was reached during search. Nevertheless, it can be used like a heuristic.

We also generate preferred operators along with the landmark heuristic. An operator is preferred in a state if applying it achieves an *acceptable* landmark in the next step, i. e., a landmark whose predecessors have already been accepted. If no acceptable landmark can be achieved within one step, the preferred operators are those which occur in a relaxed plan to the nearest acceptable landmark.

Action Costs

The landmark heuristic as outlined above estimates the goal distance of states, i. e., the number of operator applications needed to reach the goal state from a given state. Due to the inclusion of action costs in IPC 2008, however, we are interested in generating least-cost plans rather than short plans. Hence, the heuristics used during search should also estimate the *cost* of reaching the goal from a state, not just its goal distance.

The FF heuristic that is also used in our framework can easily be adapted to action costs, as we can use action costs in the underlying *additive heuristic* (Bonet and Geffner 2001). When generating a relaxed plan from the additive heuristic estimates, we simply use the cost rather than the length of that relaxed plan as our estimate for the cost-to-go of a given state. See Keyder and Geffner (2008) for a detailed description of this cost-sensitive version of the FF heuristic which they call $FF(h_a)$.

For the landmark heuristic this method is not directly applicable, since no actual plan is formed by the heuristic. Instead, we *weight* landmarks with an estimate on their minimum cost. Rather than counting the number of landmarks that still need to be achieved from a state, the heuristic value is then the sum of all minimum costs of those landmarks. The cost estimate for each landmark is the minimum cost that is required to make the landmark true for the first time, i. e., the minimum of all action costs of its first achievers.

Zero-cost actions can lead to problems in a standard cost-sensitive search like weighted A*. Since zero-cost actions can always be added to a search path “for free”, i. e., without negative side effects, the search may explore very long search paths without getting closer to a goal. In the worst case, this can prevent it from finding a solution within the given time limit. To avoid this problem, LAMA combines cost and distance estimates in a simple fashion, by counting for each action its cost *plus 1 for its distance*. This method has largely overcome the problems of zero-cost actions in our experiments, and offers a simple trade-off between the aim of finding cheap solutions and the need to find solutions within reasonable time. Of course this means that a plan consisting of five originally zero-cost actions is deemed worse by LAMA than a plan consisting of two actions that originally cost one, whereas the opposite is true. A smaller value for the additive constant would lessen the problem, though not solve it completely.

LAMA 2008 versus LAMA 2011

We have entered two versions of LAMA into the 2011 planning competition. LAMA 2008 is largely the same system as the one that won the sequential satisficing track in 2008. However, bug fixes have been incorporated since the 2008 competition, in particular to the translator component, and some improvements to the invariant synthesis in the translator have taken place. The translator component in LAMA 2008 is identical to the translator used in LAMA 2011 and other recent Fast Downward offshoots.

LAMA 2011 is a reimplementations of LAMA in the latest version of the Fast Downward planning framework. In the past two years, the Fast Downward framework has undergone significant changes aimed at back-integrating various offshoots of the original Fast Downward code and making the framework more modular. LAMA 2011 is entered into the competition in particular for comparison with LAMA 2008 and for comparison with other planners based on the latest Fast Downward code.

Compared to the 2008 version, LAMA 2011 uses a more space-efficient way of storing landmark information and a more efficient implementation of the FF heuristic in cases where the values of this heuristic are large. LAMA 2011 furthermore uses a different configuration sequence for its search algorithms. In line with our observations that focus on solution quality may harm coverage (Richter and Westphal 2010), LAMA 2011 runs one iteration of greedy best-first search *ignoring costs* before it starts the sequence of search iterations used by LAMA 2008 (cost-sensitive greedy best-first search followed by weighted A* searches).

References

- Bäckström, C., and Nebel, B. 1995. Complexity results for SAS⁺ planning. *Computational Intelligence* 11(4):625–655.
- Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *AIJ* 129(1):5–33.
- Helmert, M. 2006. The Fast Downward planning system. *JAIR* 26:191–246.
- Helmert, M. 2009. Concise finite-domain representations for PDDL planning tasks. *AIJ* 173:503–535.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *JAIR* 14:253–302.
- Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered landmarks in planning. *JAIR* 22:215–278.
- Keyder, E., and Geffner, H. 2008. Heuristics for planning with action costs revisited. In *Proc. ECAI 2008*, 588–592.
- Porteous, J., and Cresswell, S. 2002. Extending landmarks analysis to reason about resources and repetition. In *Proceedings of the 21st Workshop of the UK Planning and Scheduling Special Interest Group (PLANSIG '02)*, 45–54.
- Porteous, J.; Sebastia, L.; and Hoffmann, J. 2001. On the extraction, ordering, and usage of landmarks in planning. In *Proc. ECP 2001*, 37–48.

Richter, S., and Helmert, M. 2009. Preferred operators and deferred evaluation in satisficing planning. In *Proc. ICAPS 2009*, 273–280.

Richter, S., and Westphal, M. 2008. The LAMA planner — Using landmark counting in heuristic search. IPC 2008 short papers, <http://ipc.informatik.uni-freiburg.de/Planners>.

Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR* 39:127–177.

Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In *Proc. AAAI 2008*, 975–982.

Richter, S.; Thayer, J. T.; and Ruml, W. 2010. The joy of forgetting: Faster anytime search via restarting. In *Proc. ICAPS 2010*, 137–144.

Röger, G., and Helmert, M. 2010. The more, the merrier: Combining heuristic estimators for satisficing planning. In *Proc. ICAPS 2010*, 246–249.