The TorchLight Tool: Analyzing Search Topology Without Running Any Search

JÖRG HOFFMANN
INRIA, Nancy, FRANCE

ABSTRACT

The ignoring delete lists relaxation is of paramount importance for both satisficing and optimal planning. In earlier work [Hoffmann, JAIR'05], it was observed that the optimal relaxation heuristic h^+ has amazing qualities in many classical planning benchmarks, in particular pertaining to the complete absence of local minima. The proofs of this are hand-made, raising the question whether such proofs can be lead automatically by domain analysis techniques. The TorchLight tool answers this question in the affirmative.

TorchLight is based on a connection between causal graph structure and h^+ topology. It distinguishes between **global analysis** and **local analysis**. Global analysis shows the absence of local minima once and for all, for the entire state space of a given planning task. Local analysis computes what we call the *success rate*, which estimates the percentage of individual sample states not on local minima and thus allows to make finer distinctions. Finally, **diagnosis** summarizes structural reasons for analysis failure, thus indicating domain aspects that may cause local minima.

TorchLight Results Overview

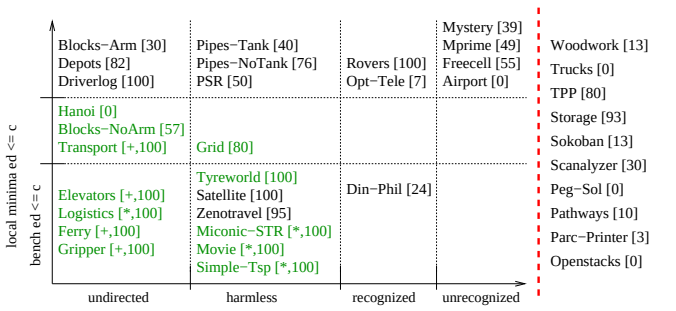

Taxonomy of Hoffmann [JAIR'05]. Green: no local minima under h^+.

"*": global analysis always succeeds. "+": local analysis always succeeds if run on optimal relaxed plans. Numbers: average success rate per domain, for local analysis (run on h^{FF}'s relaxed plans) when sampling a single state per domain instance.

Local Analysis (simplified)

Optimal rplan dependency graph oDG^+: *Assume (X, s_I, s_G, O), $s \in S$, optimal relaxed plan $P^+(s)$, $x_0 \in X$, $o_0 \in P^+(s)$ taking $t_0 = (s(x_0), c)$; denote $P^+_{\leq 0}(s) := P^+(s)$ up to o_0.*

- *Unique leaf x_0; arc (x, x') iff an operator in $P^+_{\leq 0}(s)$ takes a transition on x' preconditioned on x*
- *Non-leaf x: $oDTG^+_x$ is DTG sub-graph traversed by $P^+_{\leq 0}(s)$*

Successful oDG^+:

- *oDG^+ is acyclic*
- *If delete p of t_0 is relevant for "rest of $P^+(s)$", then $P^+(s)$ can be re-arranged so that all such p are re-achieved up front*
  [boarding passenger in Miconic deletes "not-boarded()"; picking ball in Gripper deletes "free-gripper()", re-achieved by dropping ball]
- *Non-leaf x: $oDTG^+_x$ transitions invertible and no harmful side effects*
  [moving vehicle along road-map]

Theorem. *$\exists$ successful $oDG^+ \implies s$ is not a local minimum under h^+.*
Proof. By moving only non-leaf vars x within $oDTG^+_x$, we can reach a state s_0 where t_0 can be applied. h^+ remains constant on the path, by virtue of inverting the executed operators in $P^+(s)$. After applying t_0, h^+ decreases because we can remove o_0 from the relaxed plan.

Diagnosis

If o_0 fails due to t_0 delete of p, collect (PDDL action name O, predicate name P) where o_0 instantiates O and p instantiates P; weight by frequency.

Global Analysis (simplified)

Global dependency graph gDG: *Assume (X, s_I, s_G, O), $s \in S$, goal variable $x_0 \in X$, o_0 taking $t_0 = (s(x_0), c)$.*

- *Unique leaf x_0; (x, x') precondition-effect arcs in causal graph*

Successful gDG:

- *gDG is acyclic*
- *Side-effect deletes of t_0 do not occur anywhere except (perhaps) in o_0*
  [boarding passenger in Miconic deletes "not-boarded()"]
- *Non-leaf x: all DTG_x transitions invertible and no harmful side effects*
  [moving vehicle along road-map]

Theorem. *$\forall gDG$ successful $\implies$ no local minima under h^+.*
Proof. In every non-goal s, every optimal relaxed plan $P^+(s)$ will move one goal var x_0 "for its own sake only". The oDG^+ for x_0 and its first move o_0 is contained in the respective gDG. Thus oDG^+ is acyclic, and all $oDTG^+_x$ transitions invertible/no harmful side effects. Side-effect deletes of t_0 irrelevant by prerequisite; "own" delete $s(x_0)$ irrelevant because x_0 moves for its own sake only. Altogether, oDG^+ is successful.

Illustrative Example: No Local Minima

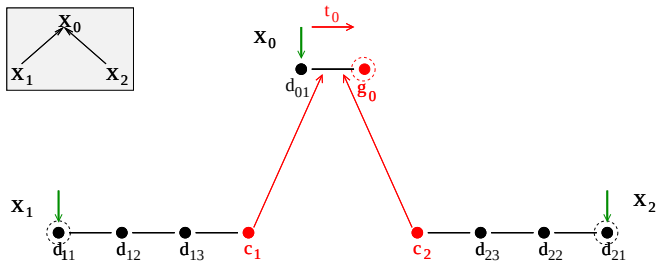

Three variables x_0, x_1, x_2. Top left: causal graph. All transitions invertible and no side effects. Green: where we are. Red: what we need to do.

TorchLight Analysis of NoLM Example

Details: see TorchLight verbose demo.

Local analysis:

Relaxed plan: $\langle x_1 12, x_1 23, x_1 3c, x_2 12, x_2 23, x_2 3c, o_0 \rangle$

Leaf-var x_1 with $x_1 12$ doesn't work because delete $x_1 = d_{11}$ is relevant (goal). Same for $x_2, x_2 12$. Other moves of x_1, x_2: start value $\neq s$.

Leaf-var x_0 with o_0: oDG^+ = causal graph is acyclic; delete $x_0 = d_{01}$ is irrelevant; $oDTG^+_x$ transitions for x_1, x_2 invertible and no side effects.

$\implies$ oDG^+ for leaf-var x_0 with o_0 is successful!

Global analysis:

Any transition of x_1, x_2: no causal graph predecessors hence no non-leaf vars (and acyclic). No side effects at all.

Any transition of x_0: causal graph predecessors x_1, x_2 with invertible/no side effects transitions; acyclic. No side effects at all.

$\implies$ all gDG successful!

Illustrative Example: Local Minima

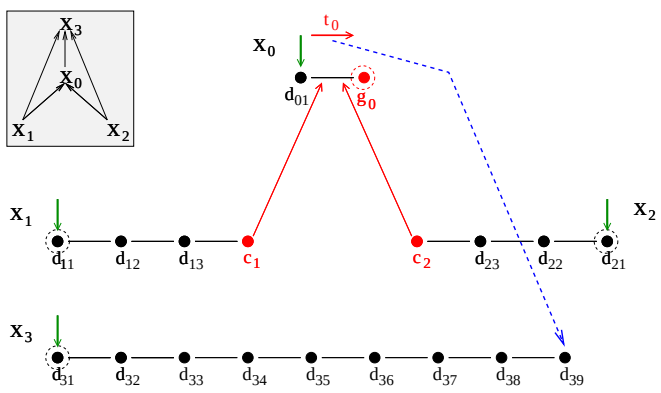

As above, but fourth variable x_3 that is already in its goal d_{31}; side effect of t_0 setting x_3 to d_{39} far away from its goal.

TorchLight Analysis of LM Example

Details: see TorchLight verbose demo.

Local analysis:

Relaxed plan: $\langle x_1 12, x_1 23, x_1 3c, x_2 12, x_2 23, x_2 3c, o_0 \rangle$

Leaf-var x_0 with o_0: as before, oDG^+ = causal graph is acyclic, $oDTG^+_x$ transitions for x_1, x_2 invertible and no side effects, delete $x_0 = d_{01}$ is irrelevant. However, delete $x_3 = d_{31}$ is relevant and not re-achieved inside relaxed plan.

$\implies$ this oDG^+ not successful! (others neither, as before)

Diagnosis:

o_0 failed due to t_0 delete of $x_3 = d_{31}$.

Global analysis:

Transition t_0 of x_0: as before, causal graph predecessors x_1, x_2 with invertible/no side effects transitions, acyclic. However, side effect on x_3!

$\implies$ this gDG is not successful!

Improving TorchLight: Strengthening Global Analysis?

Two major weaknesses of global analysis vs. local analysis:

(1) "(x, x') precondition-effect arc in causal graph" vs. "(x, x') iff an operator in $P^+_{\leq 0}(s)$ takes a transition on x' preconditioned on x"
[("carry-ball-b", "free-gripper") in causal graph due to dropping ball b; ("free-gripper", "carry-ball-b") in causal graph due to picking up the same ball b]

(2) "Side-effect deletes of t_0 irrelevant" vs. "$P^+(s)$ can be re-arranged so that all relevant deletes of t_0 re-achieved up front"
[picking ball in Gripper deletes "free-gripper()", re-achieved by dropping ball]

Hence local analysis, but not global analysis, succeeds in Elevators, Ferry, Gripper, Transport.

Addressing (1): sufficient conditions for "operator o never precedes operator o_0 in an optimal relaxed plan". Adapt [Hoffmann&Nebel, ECP'01]?

Addressing (2): sufficient conditions for "if o is in optimal relaxed plan then so is o'". Variant of landmarks analysis?

Improving TorchLight: Characterizing "Good Cases"?

Extrapolate "reasons" for local analysis success?
(Thanks to anonymous reviewer for suggesting.)

Using TorchLight: Targeted Macro-Actions?

Local analysis succeeds $\implies$ path to state with strictly smaller h^+ value!
NoLM Example: move x_1 to c_1, move x_2 to c_2, apply t_0.

Similar to relaxed-plan-execution macros [Vidal, ICAPS'04]?
Stronger if (and only if?) to-and-fro moves of non-leaf vars are needed.
(Macro can be exponentially long in depth of oDG^+ ...)

Using TorchLight: Performance Prediction?

Highly informative search space features!
("Enforced Hill-Climbing succeeds iff success rate $\geq T$"
$\implies$ 71.9% correct, vs. baseline 60.7%)

Use for automatic planner configuration!
Even online! Analyze $P^+(s)$, search more/less greedily if "yes"/"no"

Using TorchLight: Targeted Abstraction?

Global analysis succeeds $\implies$ problem tractable by chaining "macros"!

Remove diagnosed "harmful" effects until global analysis succeeds?
[transportation domains: remove fuel usage]

Option: stop anytime; run planner inside heuristic!
[Grid: allow to carry several keys at same time]

Using TorchLight: PDDL Modeling Guidance?

This whole work happened because "planning end-users" (Carlos & Luciana) complained about not having such guidance!

Diagnosis points out "critical" aspects of model
$\implies$ user may omit these aspects!
$\implies$ versioning for trade-off precision vs. costs!

[(a) end-users might not know that fuel consumption hurts, and (b) removing it might still yield useable plans ...]

References

Hoffmann, JAIR'05. Where 'ignoring delete lists' works: Local search topology in planning benchmarks. *JAIR* 24:685–758.
Hoffmann&Nebel, ECP'01. RIFO revisited: Detecting relaxed irrelevance. *6th European Conference on Planning*, 325–336.
Vidal, ICAPS'04. A lookahead strategy for heuristic search planning. *14th International Conference on Automated Planning and Scheduling*, 150–160.