A* Wars: The Fight for Improving A* Search for Troubleshooting with Dependent Actions

Thorsten J. Ottosen and Finn Verner Jensen
Department of Computer Science, Aalborg University, Denmark

Troubleshooting with actions

A Troubleshooting model consists of

- A set of *faults* $\mathcal{F}$ ($\mathrm{f}_i \in \mathcal{F}$) that is potentially causing the problem.
- A set of *actions* $\mathcal{A}$ ($\mathrm{A}_i \in \mathcal{A}$) that can fix the problem.
- A dynamic set of *evidence* $\boldsymbol{\varepsilon} = \{\mathrm{A} \in \mathcal{A} \,|\, \mathrm{A}$ failed to fix the problem (written $\mathrm{A} = \neg\mathrm{a}$) $\}$.
- A *cost* $\mathrm{C}_{\mathrm{A}}(\boldsymbol{\varepsilon})$ for each action A, possibly depending on evidence $\boldsymbol{\varepsilon}$.
- A *Bayesian Network* that provides $\mathrm{P}(\mathrm{A}\,|\,\boldsymbol{\varepsilon})$, $\mathrm{P}(\mathrm{A}\,|\,\mathrm{f},\boldsymbol{\varepsilon})$ and $\mathrm{P}(\mathrm{f}\,|\,\boldsymbol{\varepsilon})$.

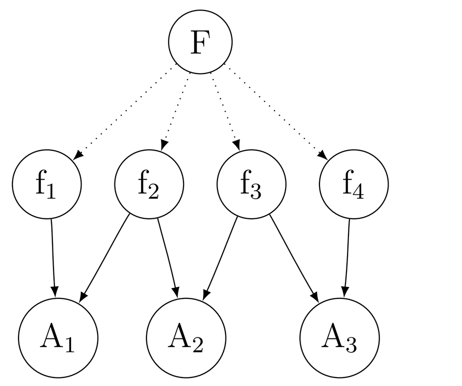

	f_1	f_2	f_3	f_4
$\mathrm{P}(\mathrm{a}_1\,|\,\mathrm{F})$	1	1	0	0
$\mathrm{P}(\mathrm{a}_2\,|\,\mathrm{F})$	0	1	1	0
$\mathrm{P}(\mathrm{a}_3\,|\,\mathrm{F})$	0	0	1	1
$\mathrm{P}(\mathrm{F})$	0.20	0.25	0.40	0.15
$\mathrm{C}_{\mathrm{A}_1} = \mathrm{C}_{\mathrm{A}_2} = \mathrm{C}_{\mathrm{A}_3} = 1$				

Figure 1: Left: a simple model for a troubleshooting scenario with dependent actions. The dotted lines indicate that the faults f_1 to f_4 are states in a single fault node F. A_1, A_2 and A_3 represent actions, and parents of an action node A are faults which may be fixed by A. Right: the quantitative part of the model.

Definition 1. The *expected cost of repair* (ECR) of a troubleshooting sequence $\mathrm{s} = \langle \mathrm{A}_1, \ldots, \mathrm{A}_n \rangle$ with costs $\mathrm{C}_{\mathrm{A}_i}$ is the mean of the costs until an action succeeds or all actions have been performed:

$$\mathrm{ECR}\,(\mathrm{s}) = \sum_{i=1}^{n} \mathrm{C}_{\mathrm{A}_i}(\boldsymbol{\varepsilon}^{i-1}) \cdot \mathrm{P}\left(\boldsymbol{\varepsilon}^{i-1}\right) \,.$$

The goal is to determine a sequence with the lowest ECR.

Example (ECR calculation)

Consider a sequence for the model in Figure 1:

$$\begin{aligned}\mathrm{ECR}\,(\langle \mathrm{A}_2, \mathrm{A}_3, \mathrm{A}_1 \rangle) &= \mathrm{C}_{\mathrm{A}_2} + \mathrm{P}(\neg\mathrm{a}_2) \cdot \mathrm{C}_{\mathrm{A}_3} + \mathrm{P}(\neg\mathrm{a}_2, \neg\mathrm{a}_3) \cdot \mathrm{C}_{\mathrm{A}_1} \\ &= \mathrm{C}_{\mathrm{A}_2} + \mathrm{P}(\neg\mathrm{a}_2) \cdot \mathrm{C}_{\mathrm{A}_3} + \mathrm{P}(\neg\mathrm{a}_2) \cdot \mathrm{P}(\neg\mathrm{a}_3\,|\,\neg\mathrm{a}_2) \cdot \mathrm{C}_{\mathrm{A}_1} \\ &= 1 + \frac{7}{20} \cdot 1 + \frac{7}{20} \cdot \frac{4}{7} \cdot 1 = 1.55 \,.\end{aligned}$$

The set of faults that can be repaired by an action A is denoted $\mathrm{fa}\,(\mathrm{A})$. For example, in Figure 1 we have $\mathrm{fa}\,(\mathrm{A}_2) = \{\mathrm{f}_2, \mathrm{f}_3\}$. In models where actions can have $\mathrm{P}(\mathrm{a}\,|\,\boldsymbol{\varepsilon}) = 1$, $\mathrm{fa}(\cdot)$ is a dynamic entity which we indicate by writing $\mathrm{fa}(\cdot\,|\,\boldsymbol{\varepsilon})$.

Definition 2. The *efficiency* of an action A given evidence $\boldsymbol{\varepsilon}$ is the probability that the actions solves the problem divided by its cost, that is

$$\mathrm{ef}(\mathrm{A}\,|\,\boldsymbol{\varepsilon}) = \frac{\mathrm{P}(\mathrm{A} = \mathrm{a})}{\mathrm{C}_{\mathrm{A}}(\boldsymbol{\varepsilon})} \,.$$

A* and monotonicity of the function $\underline{\mathrm{ECR}}$

A* is a best-first search algorithm that works by continuously expanding a frontier node n for which the value of the *evaluation function*

$$f(n) = g(n) + h(n),$$

is minimal until finally a goal node t is expanded (Hart et al., 1968). If node m is reachable from node n, $c(n,m)$ is the cost from n to m. Then $g(n) = c(s,n)$ where s is the start node, and $h(n)$ is the *heuristic function* that guides (or misguides) the search by estimating the cost $c(n,t)$. For Troubleshooting we have

$$f(n) = \underbrace{\mathrm{ECR}\,(\boldsymbol{\varepsilon}^{\mathrm{n}})}_{\mathrm{g(n)}} + \underbrace{\underline{\mathrm{ECR}}(\boldsymbol{\varepsilon}^{n})}_{\mathrm{h(n)}} \,,$$

where $\mathrm{ECR}\,(\boldsymbol{\varepsilon}^{\mathrm{n}})$ is the ECR of the sequence defined by the path from s to n.

Definition 3 (Vomlelová and Vomlel, 2003)**.** Let $\mathcal{E}$ denote the set containing all possible evidence. The function $\underline{\mathrm{ECR}} : \mathcal{E} \mapsto \mathcal{R}^+$ is defined for each $\boldsymbol{\varepsilon}^n \in \mathcal{E}$ as

$$\underline{\mathrm{ECR}}(\boldsymbol{\varepsilon}^n) = \mathrm{P}(\boldsymbol{\varepsilon}^n) \cdot \sum_{\mathrm{f} \in \mathcal{F}} \mathrm{P}(\mathrm{f}\,|\,\boldsymbol{\varepsilon}^n) \cdot \mathrm{ECR}^*(\boldsymbol{\varepsilon}^n \cup \mathrm{f}) \,.$$

where $\mathrm{ECR}^*(\boldsymbol{\varepsilon}^n \cup \mathrm{f})$ is the optimal cost when a fault f is known.

Example (ECR^* calculation)

Assume the fault f can be repaired by two actions A_1 and A_2 and that $\mathrm{P}(\mathrm{a}_1\,|\,\mathrm{f}) = 0.9$ and $\mathrm{P}(\mathrm{a}_2\,|\,\mathrm{f}) = 0.8$. Furthermore, let both actions have cost 1. Since instantiating the fault node renders the actions conditionally independent, $\mathrm{P}(\mathrm{a}\,|\,\boldsymbol{\varepsilon} \cup \mathrm{f}) = \mathrm{P}(\mathrm{a}\,|\,\mathrm{f})$ and the efficiencies of the two actions are 0.9 and 0.8, respectively. We get

$$\mathrm{ECR}^*(\boldsymbol{\varepsilon} \cup \mathrm{f}) = \mathrm{ECR}\,(\langle \mathrm{A}_1, \mathrm{A}_2 \rangle) = \mathrm{C}_{\mathrm{A}_1} + \mathrm{P}(\neg\mathrm{a}_1\,|\,\mathrm{f}) \cdot \mathrm{C}_{\mathrm{A}_2} = 1 + 0.1 \cdot 1 = 1.1 \,.$$

because the optimal sequence with independent actions is found by ordering the actions w.r.t. descending initial efficiency (Kadane and Simon, 1977).

Definition 4. A heuristic function $h(\cdot)$ is *monotone* if

$$h(n) \leq c(n,m) + h(m),$$

whenever m is a successor node of n.

For monotone heuristic functions A* is *guaranteed* to have found the optimal path to a node when the node is expanded (Hart et al., 1968).

Theorem 1. *Under the assumption of no questions, constant costs, a single initial fault, and conditional independence of actions given that the fault is known, then the heuristic function* $\underline{\mathrm{ECR}}(\cdot)$ *is monotone.*

Hybrid-A* Algorithm

Definition 5. A *dependency graph* for a troubleshooting model given evidence $\boldsymbol{\varepsilon}$ is the undirected graph with a vertex for each action $\mathrm{A} \in \mathcal{A}(\boldsymbol{\varepsilon})$ and an edge between two vertices A_1 and A_2 if $\mathrm{fa}(\mathrm{A}_1\,|\,\boldsymbol{\varepsilon}) \cap \mathrm{fa}(\mathrm{A}_2\,|\,\boldsymbol{\varepsilon}) \neq \emptyset$.

Definition 6. A *dependency set leader* for a troubleshooting model given evidence $\boldsymbol{\varepsilon}$ is the first action of an optimal sequence in a connectivity component in the dependency graph given $\boldsymbol{\varepsilon}$ (a *dependency set*).

Theorem 2 (Koca and Bilgiç, 2004)**.** *The globally optimal sequence is given by the following algorithm:*

I. Construct the dependency sets and retrieve the set leaders.
II. Calculate $\mathrm{ef}(\cdot)$ *for all set leaders.*
III. Select the set leader with the highest $\mathrm{ef}(\cdot)$ *and perform it.*
IV. If it fails, update the probabilities, and continue in step (2).

Hybrid-A*: *We exploit Theorem 2 and avoid branching whenever the most efficient action belongs to a small dependency set (which is solved by brute-force).*

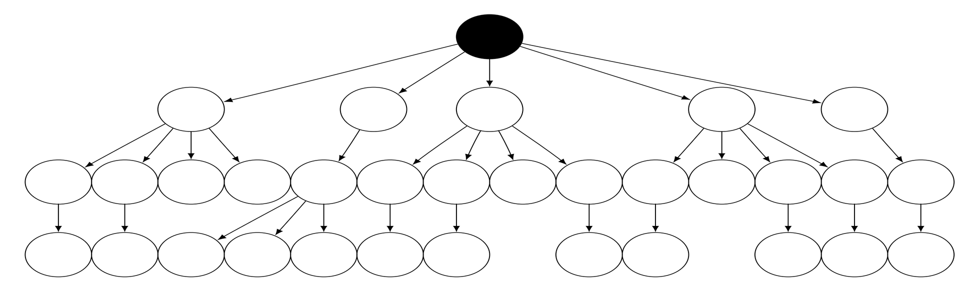

Figure 2: An example of what the search tree looks like in the hybrid approach. For some nodes, the normal A* branching is avoided, and near goal nodes this branching is almost avoided for all nodes.

Experimental results

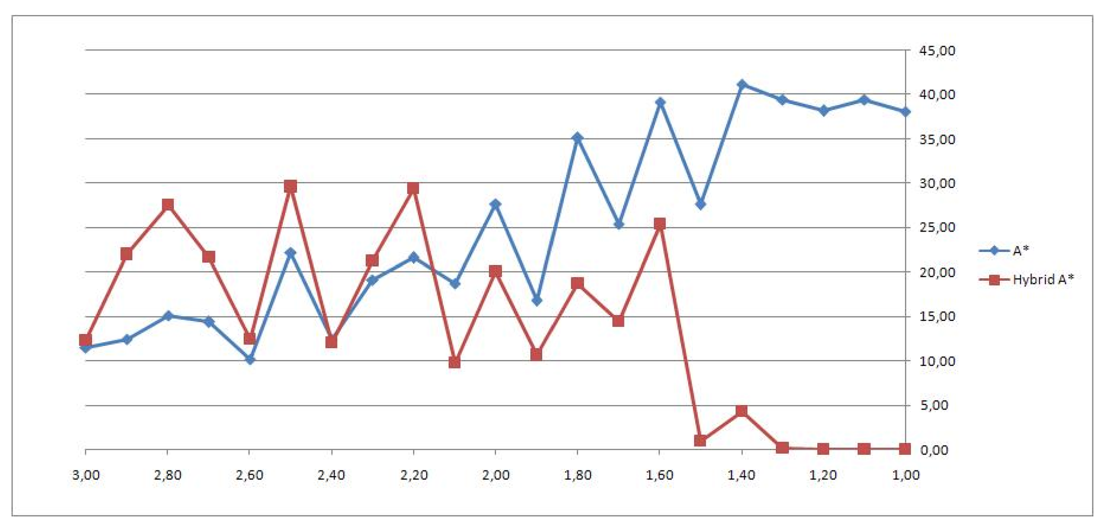

Figure 3: Comparison of normal A* (Ottosen and Jensen, 2008) with the hybrid approach. The X-axis indicates average dependency of the model (that is, the average size of $\mathrm{fa}(\cdot)$ over all actions), and the Y-axis represents time in seconds. All models had 20 actions and 20 faults.

References

P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Systems Science and Cybernetics*, SSC-4(2):100–7, 1968.

J. Kadane and H. Simon. Optimal strategies for a class of constrained sequential problems. *The Annals of Statistics*, 5:237–255, 1977.

E. Koca and T. Bilgiç. A troubleshooting approach with dependent actions. In R. L. de Mántaras and L. Saitta, editors, *ECAI 2004: 16th European Conference on Artificial Intelligence*, pages 1043–1044. IOS Press, 2004. ISBN 1-58603-452-9.

T. J. Ottosen and F. V. Jensen. Better safe than sorry—optimal troubleshooting through A* search with efficiency-based pruning. In *Proceedings of the Tenth Scandinavian Conference on Artificial Intelligence*, pages 92–97. IOS Press, 2008. ISBN 978-1-58603-867-0.

M. Vomlelová and J. Vomlel. Troubleshooting: Np-hardness and solution methods. *Soft Computing Journal, Volume 7, Number 5*, pages 357–368, 2003.