Distance Estimates for Planning in the Discrete Belief Space

Jussi Rintanen
Albert-Ludwigs-Universität Freiburg
Institut für Informatik
Georges-Köhler-Allee, 79110 Freiburg im Breisgau
Germany

Abstract

We present a general framework for studying heuristics for planning in the belief space. Earlier work has focused on giving implementations of heuristics that work well on benchmarks, without studying them at a more analytical level. Existing heuristics have evaluated belief states in terms of their cardinality or have used distance heuristics directly based on the distances in the underlying state space. Neither of these types of heuristics is very widely applicable: often goal belief state is not approached through a sequence of belief states with a decreasing cardinality, and distances in the state space ignore the main implications of partial observability.

To remedy these problems we present a family of admissible, increasingly accurate distance heuristics for planning in the belief space, parameterized by an integer n. We show that the family of heuristics is theoretically robust: it includes the simplest heuristic based on the state space as a special case and as a limit the exact distances in the belief space.

Introduction

Planning under partial observability is one of the most difficult planning problems because algorithms for finding plans have to – at least implicitly – address the incompleteness of the knowledge about the environment in which the plan execution will take place. The incompleteness of knowledge leads to the notion of belief states: depending on the type planning, belief states are either probability distributions on the state space, or, in the non-probabilistic versions of the partially observable planning problem, sets of states. In either case, the exponentially bigger size of the set of belief states – the belief space – in comparison to the state space is a major obstacle in achieving efficient planning.

Following the lead in classical planning (Bonet & Geffner 2001), also restricted types of planning in the belief space, most notably the planning problem without any observability at all (sometimes known as conformant planning), has also been represented as a heuristic search problem (Bonet & Geffner 2000). However, the implementation of heuristic search in the belief space is more complex than in classical planning because of the difficulty of deriving good heuristics. First works on the topic have used distances in the state space (Bonet & Geffner 2000) and cardinalities of the belief states (Bertoli, Cimatti, & Roveri 2001). On some types of problems these heuristics work well, but not on all, and the two proposed approaches have orthogonal strengths.

Many problems cannot be solved by blindly taking actions that reduce the cardinality of the current belief state: the cardinality of the belief state may stay the same or increase during plan execution, and hence the decrease in cardinality is not characteristic to belief space planning in general.

Similarly, distances in the state space completely ignore the most distinctive aspect of planning with partial observability: the same action must be used in two states if the states are not observationally distinguishable. A given (optimal) plan for an unobservable problem may increase the actual current state-space distance to the goal states (on a given execution) when the distance in the belief-space monotonically decreases, and vice versa. Hence, the state space distances may yield wildly misleading estimates of the distances in the corresponding belief space.

To achieve more efficient planning it is necessary to develop belief space heuristics that combine the strengths of existing heuristics based on cardinalities and distances in the state space. In this work we present such a family of heuristics, parameterized by a natural number n. The accuracy of the distance estimates improves as n grows. As the special case $n = 1$ we have a heuristic based on distances in the state space, similar to ones used in earlier work. When the cardinality of the state space equals n, the distance estimates equal the actual distances in the belief space.

Planning in the Belief Space

Let S be a set of states. A planning problem $\langle I, O, G\rangle$ without observability is abstractly defined in terms of

1. an initial belief state $I \subseteq S$,
2. a set of operators O, each operator $o \in O$ associated with a transition relation $R_o \subseteq S \times S$ describing the possible transitions when o is applied (the operators are nondeterministic and hence the transition relations are not restricted to partial functions), and
3. a set G of goal states (the goal belief state).

In the special case of deterministic (classical) planning the set I consists of one state and all the relations $R_o, o \in O$ are partial functions. The objective in deterministic planning is to find a sequence $o_1, \ldots, o_n$ of operators so that $s_0 R_{o_1} s_1 R_{o_2} s_2 \cdots s_{n-1} R_n s_n$ where $I = \{s_0\}$ and $s_n \in G$.

The state space is viewed as a graph $\langle S, R_1, R_2, \ldots, R_n \rangle$ with nodes $s \in S$ and sets R_i of directed edges labeled with operators $o_i \in O$. Classical planning is finding a path from the initial state to one of the goal states.

In planning without observability (conformant planning), the objective is similarly to find a sequence of operators reaching a goal state, but because of the unknown initial state (equivalently, an initial state nondeterministically chosen from I) and the nondeterministic transitions, the intermediate stages during plan execution are not states, but sets of states, describing the sets of possible states at every stage of plan execution.

Belief states in the belief space are sets of states, that is, the belief space is the powerset of the state space. An operator o maps a belief state B to the belief state

$$\{s' \in S | s \in B, sR_o s'\}$$

if[1] for every $s \in B$ there is $s' \in S$ such that $sR_o s'$.

Similarly a labeled graph $\langle 2^S, R^b_1, R^b_2, \ldots, R^b_n \rangle$ can be constructed for the belief space. Here the set of directed edges R^b_i are defined so that belief state B is related to B' by R^b_i (that is $BR^b_i B'$) if B is mapped to B' by o_i. Like in classical planning, these relations R^b_i are partial functions, and similarly to classical planning, planning without observability (conformant planning) is finding a path from the initial belief state I to a belief state G' such that $G' \subseteq G$.

Algorithms

Planning in the belief space can be solved like a state-space search problem: start from the initial belief state and repeatedly follow the directed edges in the belief space to reach new belief states, until a belief state that is a subset of the goal states is reached.

When a heuristic search algorithm is used for avoiding the enumeration of all the belief states, then it is important to use some informative heuristic for guiding the search.

The most obvious heuristic would be an estimate for the distance from the current belief state B to the set of goal states G. In the next section we discuss some admissible heuristics that are derived *from distances in the state space* that underlies the belief space. These are the most obvious heuristics one could use. The rest of the paper is dedicated to investigating more informative heuristics that are not directly derived from the distances in the state space.

Distance Heuristics from the State Space

Research on conformant planning so far has concentrated on distance heuristics derived from the distances of individual states in the state space. The distance of a belief state is for example the maximum length of a shortest path from a constituent state to a goal state.

[1]This is the applicability condition, corresponding to the requirement that the preconditions of the operator is satisfied in every possible state of the belief state (precondition is part of the commonly used syntactic definitions of operators.)

Bonet and Geffner (2000) have used a related distance measure, the optimal expected number of steps of reaching the goal in the corresponding probabilistic problem. Bryce and Kambhampati (2003) have considered efficient approximations of state space distances.

The most obvious distance heuristics are based on the weak and strong distances in the state space. The weak distances of states are based on the following inductive definition. Sets D_i consist of those states from which a goal state is reachable in i steps or less.

$$\begin{aligned} D_0 &= G \\ D_{i+1} &= D_i \cup \{s \in S | o \in O, s' \in D_i, sR_o s'\} \end{aligned}$$

A state s has *weak distance* $d \geq 1$ if $s \in D_d \backslash D_{d-1}$ and distance 0 if $s \in G$. This means that it is possible to reach one of the goal states starting from s by a sequence of d operators if the nondeterministic alternatives play out favorably. Of course, nondeterministic actions may come out unfavorably and a higher number of actions may be needed, or the goal may even become unreachable, but if it is possible that the goals are reached in d steps then the weak distance is d.

Strong distances are based on a slightly different inductive definition. Now D_i consists of those states for which there is a guarantee of reaching a goal state in i steps or less.

$$\begin{aligned} D_0 &= G \\ D_{i+1} &= D_i \cup \{s \in S | o \in O, s' \in D_i, sR_o s', \\ &\qquad sR_o s'' \text{ implies } s'' \in D_i \text{ for all } s''\} \end{aligned}$$

A state s has *strong distance* $d \geq 1$ if $s \in D_d \backslash D_{d-1}$ and strong distance 0 if $s \in G$.

Next we derive distance heuristics for the belief space based on state space distances. Both weak and strong distances yield an admissible distance heuristic for belief states, but strong distances are (not always properly) higher than weak distances and therefore a more accurate estimate for plan length.

Definition 1 (State space distance) *The* state space distance *of a belief state* B *is* $d \geq 1$ *when* $B \subseteq D_d$ *and* $B \not\subseteq D_{d-1}$*, and it is* 0 *when* $B \subseteq D_0 = G$.

Even though computing the exact distances for a typical succinct representation of transition systems, like STRIPS operators, is PSPACE-hard, the much higher complexity of planning problems with partial observability still often justifies it: this computation would in many cases be an inexpensive preprocessing step, preceding the much more expensive solution of the partially observable planning problem. Otherwise cheaper approximations can be used.

n-Distances

The essence of planning without observability is that the same sequence of actions has to lead to a goal state for every state in the belief state. The distance heuristics from the strong distances in the state space may assign distance 1 to both state s_1 and state s_2, but the distance from $\{s_1, s_2\}$ may be anywhere between 1 and infinite. The inaccuracy in estimating the distance of $\{s_1, s_2\}$ based on the distances of s_1

and s_2 is that the distances of s_1 and s_2 in the state space may be along paths that have nothing to do with the path for $\{s_1, s_2\}$ in the belief space. That is, the actions on these three paths may be completely different.

This leads to a powerful idea. Instead of estimating the distance of B in terms of distances of states $s \in B$ in the state space S, let us estimate it in terms of distances of n-tuples of states in the product state space S^n, in which there is a transition from $\langle s_1, \ldots, s_n \rangle$ to $\langle s'_1, \ldots, s'_n \rangle$ if and only if there is an operator o that allows a transition from s_i to s'_i for every $i \in \{1, \ldots, n\}$. Here the important point is that the transition between the tuples is by using the same operator for every component state. This corresponds to the necessity of using the same operator for every state, because observations cannot distinguish between them.

This leads to a generalization of strong distances. We define the distances of n-tuples of states as follows.

$$\begin{aligned} D_0 &= G^n \\ D_{i+1} &= D_i \cup \{\sigma \in S^n | o \in O, \sigma' \in D_i, \sigma R^n_o \sigma', \\ &\qquad \sigma R^n_o \sigma'' \text{ implies } \sigma'' \in D_i \text{ for all } \sigma''\} \end{aligned}$$

Here R^n_o is defined by

$$\langle s_1, \ldots, s_n \rangle R^n_o \langle s'_1, \ldots, s'_n \rangle \text{ if } s_i R_o s'_i \text{ for all } i \in \{1, \ldots, n\}.$$

Now we can define the n-distance of a belief state as follows.

Definition 2 (n-distance) *Belief state B has n-distance $d = 0$ if for all $\{s_1, \ldots, s_n\} \subseteq B$, $\langle s_1, \ldots, s_n \rangle \in D_0$ (this is equivalent to $B \subseteq G$.) Belief state B has n-distance $d \geq 1$ if for all $\{s_1, \ldots, s_n\} \subseteq B$, $\langle s_1, \ldots, s_n \rangle \in D_d$ and for some $\{s_1, \ldots, s_n\} \subseteq B$, $\langle s_1, \ldots, s_n \rangle \notin D_{d-1}$. If the distance d is not any natural number, then $d = \infty$.*

So, we look at all the n-element subsets of B, see what their distance to goals is in the product state space S^n, and take the maximum of those distances. Notice that when we pick the elements $s_1, \ldots, s_n$ from B, we do not and cannot assume that the elements $s_1, \ldots, s_n$ are distinct, for example because B may have less than n states. Of course, the definition assumes that B has at least one state. The 1-distance of a belief state coincides with the state space distance defined in the previous section.

The motivation behind n-distances is that computing the actual distance of belief states is very expensive (as complex as the planning problem itself) but we can use an informative notion of distances for "small" belief states of size n.

Next we investigate the properties of n-distances. The first result shows that n-distances are at least as good an estimate as m-distances when $n > m$. This result is based on a technical lemma that shows that m-tuples from the definition of m-distances are included in the n-tuples of the definition of n-distances.

Lemma 3 (Embedding) *Let $n > m$ and let $D_0, D_1, \ldots$ be the sets in the definition of m-distances, and $D'_0, D'_1, \ldots$ the sets in the definition of n-distances.*

Then for all $i \geq 0$, all belief states B and all $\{s_1, \ldots, s_m\} \subseteq B$, if $\langle s_1, \ldots, s_m, s'_{m+1}, \ldots, s'_n \rangle \in D'_i$ where $s'_k = s_m$ for all $k \in \{m+1, \ldots, n\}$, then $\langle s_1, \ldots, s_m \rangle \in D_i$.

Proof: By induction on i. Base $i = 0$: If $\{s_1, \ldots, s_m\} \subseteq B$ and $\langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle \in D'_0$ (state s_m is repeated so that the number of components in the tuple is n), then $\langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle \in G^n$. Consequently, $\langle s_1, \ldots, s_m \rangle \in D_0 = G^m$.

Inductive case $i \geq 1$: We show that if $\langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle \in D'_i$, then there is an operator $o \in O$ such that for any states $s'_1, \ldots, s'_m$ such that $s_i R_o s'_i$ for all $i \in \{1, \ldots, m\}$, $\langle s'_1, \ldots, s'_m \rangle \in D_{i-1}$, and hence $\langle s_1, \ldots, s_m \rangle \in D_i$.

So assume $\langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle \in D'_i$. Hence there is an operator $o \in O$ such that for all $s'_1, \ldots, s'_n$ such that $s_i R_o s'_i$ for all $i \in \{1, \ldots, m\}$ and $s_m R_o s'_j$ for all $j \in \{m+1, \ldots, k\}$, the n-tuple $\langle s'_1, \ldots, s'_n \rangle$ is in D'_{i-1}. Now $\langle s'_1, \ldots, s'_m, s'_m, \ldots, s'_m \rangle \in D'_{i-1}$ because this tuple is one of those reachable from $\langle s_1, \ldots, s_m, s_m \ldots, s_m \rangle$, and hence by the induction hypothesis $\langle s'_1, \ldots, s'_m \rangle \in D_{i-1}$. Because this holds for all s'_i reachable by R_o from the corresponding s_i, we have $\langle s_1, \ldots, s_m \rangle \in D_i$. $\square$

The embedding of m-distances in n-distances as provided by the lemma easily yields the result that n-distances are more accurate than m-distances when $n > m$.

Theorem 4 *Let d_n be the n-distance for a belief state B and d_m the m-distance for B. If $n > m$, then $d_n \geq d_m$.*

Proof: So assume that the m-distance of B is d_m. This implies that there is $\{s_1, \ldots, s_m\} \subseteq B$ such that $\langle s_1, \ldots, s_m \rangle \notin D_{d_m - 1}$. Let $\sigma = \langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle$ where s_m is repeated $n - m + 1$ times. By Lemma 3 $\sigma \notin D'_{d_m - 1}$. Hence there is $\{s_1, \ldots, s_m, s_m, \ldots, s_m\rangle \subseteq B$ such that $\langle s_1, \ldots, s_m, s_m, \ldots, s_m \rangle \notin D'_{d_m}$. Hence the n-distance d_n of B is greater than or equal to d_m. $\square$

So, 2-distances are a better estimate for belief states than the estimate given by state space distances, and 3-distances are a better estimate than 2-distances, and so on.

For the last result we need another lemma, which we give without a proof. The lemma states that the components of the tuples in the sets D_i may be reordered and replaced by existing components.

Lemma 5 *Let $D_0, D_1, \ldots,$ be the sets in the definition of n-distances. Then if $\langle s_1, \ldots, s, s', \ldots, s_k \rangle$ is in D_i, then so is $\langle s_1, \ldots, s', s, \ldots, s_k \rangle$, and if $\langle s_1, s_2, s_3, \ldots, s_k \rangle$ is in D_i, then so is $\langle s_1, s_1, s_3, \ldots, s_k \rangle$.*

Theorem 6 *Let the belief space be 2^S, where the cardinality of the state space is $n = |S|$. Then the n-distance of B equals the distance of B in the belief space.*

Proof: Omitted because of lack of space. $\square$

In summary, the accuracy of the n-distances grows as n grows, and asymptotically when n equals the number of states it is perfectly accurate.

In addition to including state space distances as a special case, the family of n-distances also takes into account the cardinalities of belief states, although only in a restricted manner as determined by the magnitude of n. Consider the belief state $B = \{s_1, s_2\}$. Its 2-distance is determined by the membership of the tuples $\sigma_1 = \langle s_1, s_2 \rangle$ (and symmetrically $\langle s_2, s_1 \rangle$), $\sigma_2 = \langle s_1, s_1 \rangle$, and $\sigma_3 = \langle s_2, s_2 \rangle$ in the sets D_i. The distance of σ_1 is at least as high as that of σ_2 and σ_3, because any sequence of actions leading to goals that is applicable for $\{s_1, s_2\}$ is also applicable for s_1 alone and for s_2 alone, and there might be shorter action sequences applicable for s_1 and s_2 but not for $\{s_1, s_2\}$. Therefore, any reduction in the size of a belief state, like from $\{s_1, s_2\}$ to $\{s_1\}$ would appropriately improve the n-distance estimate.

Accuracy of the Heuristics

Preceding results show that the accuracy of n-distances increases as n grows, reaching perfect accuracy when n equals the cardinality of the state space. Can we demonstrate advantages of n-distances on concrete planning problems? We shed some light on this issue next, but the reader should note that the benchmarks used by us and others are rather simple and unlikely to reflect properties of more challenging problems. Consequently, this section just illustrates the impact the improved heuristics can have on different problems.

We have implemented a planner that does heuristic search in the belief space, and three heuristics for guiding the search algorithms implemented in this planner: the 1-distances, the 2-distances, and the cardinality of belief states. The first two heuristics are admissible, and can be used in connection with optimal heuristic search algorithms like A$*$. The third heuristic, size of the belief states, does not directly yield an admissible heuristic, and we use it only in connection with a search algorithm that does not rely on admissibility.

The planner is implemented in C and represents belief states as BDDs with the CUDD system from Colorado University. CUDD provides functions for computing the cardinality of a belief state. Our current implementation does not support n-distances for n other than 1 and 2.

The search algorithms implemented in our planner include the optimal heuristic search algorithm A$*$, the suboptimal family of algorithms WA$*$[2], and suboptimal best-first search which first expands those nodes that have the lowest estimated remaining distance to the goal states.

The main topic to be investigated is the relative accuracy of n-distances, and as a secondary topic we briefly evaluate the effectiveness of different types of search algorithms and heuristics. We use the following benchmarks. Regrettably there are few meaningful benchmarks; all the interesting ones are for the more general problem of partially observable planning.

- Bonet and Geffner (2000) proposed one of the most interesting benchmarks for conformant planning so far, sorting networks (Knuth 1998). A sorting network consists of an ordered (or a partially ordered) set of gates acting on a number of input lines. Each gate combines a comparator and a swapper: if first input is greater than the second, then swap the values. The goal is to sort the input sequence. The sorting network always has to perform the same operations irrespective of the input, and hence exactly corresponds to planning without observability.

  Our size parameter is the number of inputs.

- In the empty room benchmark a robot without any sensors moves in a room to north, south, east and west and its goal is to get to the middle of the room. This is possible by going to north or south and then west or east until the robot knows that it is in one of the corners. Then it is easy to go to the goal position. The robot does not know its initial location.

  Size n characterizes room size $2^n \times 2^n$.

- Our blocks world benchmark is the standard blocks world, but with several initial states and modified to be solvable without observability even when the initial state is not known. Operators are always applicable, but nothing happens if the relevant blocks are not accessible. The initial belief state consists of all the possible configurations of the n blocks, and the goal is to build a stack consisting of all the blocks in a fixed order.

- The ring of rooms benchmark involves a round building with a cycle of n rooms with a window in each that can be closed and locked. Initially the state of the windows and the location of the robot is unknown. The robot can move to the next room either in clockwise or counterclockwise direction, and then close and lock the windows. Locking is possible only if the window is closed. Locking an already locked window and closing an already closed window does not have any effect.

  The size parameter is the number of rooms.

There are other benchmarks considered in the literature, but their flavor is close to some of the above, and many can be easily reformulated as planning with full observability.

Table 1 makes a comparison on the accuracy of 1-distances and 2-distances on a number of problem instances. For each heuristic we first give the distance estimate for the initial belief state, followed by the percentage of the actual distance. The actual distance (= length of the shortest plan) was determined by A$*$ and is given in the last column.

As expected, on most of the problems 2-distances are strictly better estimates than 1-distances, and surprisingly, on one of the problems, the empty room navigation problem, the 2-distances equal the lengths of the shortest plans.

For the ring of room problems 1-distances and 2-distances are the same, and coincide with the actual shortest plan length. This is because of the simple structure of the problem and its belief space. It seems that 2-distances would also not provide an advantage over 1-distances on many other problems in which there are no dependencies between state variables with unknown values.

The sorting network problem is the most difficult of the benchmarks in terms of the relation between difficulty and number of state variables. Every initial state (combination of input values) in this benchmark can be solved by a sorting network with a small number of gates (more precisely $\lfloor \frac{n}{2} \rfloor$),

[2] We parameterize WA$*$ with $W = 5$, giving a 5 times higher value to the estimated remaining distance than to the distance so far, yielding solutions having cost at most 5 times the optimal.

	1-distance		2-distance		exact
	len	%	len	%	len
sort02	1	1.00	1	1.00	1
sort03	1	0.33	2	0.67	3
sort04	2	0.40	3	0.60	5
sort05	2	0.22	3	0.33	9
sort06	3	0.25	4	0.33	12
sort07	3	0.19	5	0.31	16
sort08	4	0.16	6	0.32	19
ring03	8	1.00	8	1.00	8
ring04	11	1.00	11	1.00	11
ring05	14	1.00	14	1.00	14
ring06	17	1.00	17	1.00	17
ring07	20	1.00	20	1.00	20
BW02	2	0.67	3	1.00	3
BW03	4	0.57	5	0.71	7
BW04	6	0.46	8	0.62	13
BW05	7	0.41	9	0.53	17
emptyroom01	2	1.00	2	1.00	2
emptyroom02	4	0.50	8	1.00	8
emptyroom03	8	0.40	20	1.00	20
emptyroom04	16	0.36	44	1.00	44
emptyroom05	32	0.35	92	1.00	92

Table 1: Accuracy of 1-distances and 2-distances on a number of problem instances

which makes the 1-distances small. Increasing n monotonically increases n-distances, but the increase is slow because for a small number of input combinations the smallest network sorting them all is still rather small, and as the number of input value combinations is exponential in the number of inputs, only a tiny fraction of all combinations is covered.

Table 2 gives runtimes on all combinations of search algorithm and heuristic. We only report the time spent in the search algorithm, ignoring a preprocessing phase during which BDDs representing 1-distances and 2-distances are computed. Computing 2-distances is more expensive than computing 1-distances because there are twice as many variables in the BDDs and the efficiency of BDDs decreases as BDDs grow. The higher accuracy of 2-distances is often reflected in the runtimes.

On the empty room problems, performance of A$*$ and 1-distances quickly deteriorates as room size grows, while with 2-distances A$*$ immediately constructs optimal plans even for bigger rooms. On sorting networks and WA$*$ 2-distances lead to a better performance because of its advantage over 1-distances in accuracy, but finding bigger optimal networks is still very much out of reach.

On all of the problems, best-first search is the fastest to find a plan, but plans were much longer on the empty room and blocks world problems, and slightly worse on sorting networks. For bigger sorting networks the cardinality heuristic combined with best-first is the best combination, as runtimes with the other heuristics and with A$*$ and WA$*$ grow much faster. We believe that our collection of benchmarks is too small to say conclusively anything general about the relative merits of the heuristics.

Interestingly, our planner with best-first search and the cardinality heuristic produces optimal sorting networks up to size 8 (Bonet and Geffner (2000) report producing networks until size 6 with an optimal algorithm), and for bigger networks the difference to known best networks is first relatively small, but later grows; see Table 3.

Related Work

Bonet and Geffner (2000) were one of the first to apply heuristic state-space search to planning in the belief space. They used a variant of the state space distance heuristic considered by us, with the difference that they were addressing probabilistic problems and considered expected distances of states under the optimal probabilistic plan, instead of the non-probabilistic weak or strong distances.

Bryce and Kambhampati (2003) compute distances with Graphplan's (Blum & Furst 1997) planning graphs, and recognize that the smallest of the weak distances of states in a belief state – as is trivially obtained from planning graphs – is not very informative, and propose improvements based on multiple planning graphs: for a formula $\chi_1 \vee \chi_2 \vee \cdots \vee \chi_n$ describing a belief state, compute the estimate for each χ_i separately. Then an admissible estimate for the whole belief state is bounded from above by $\max_{i=1}^{n} \min_{s \in \sigma(\chi_i)} \delta(s)$ where $\sigma(\chi_i)$ is the set of states described by χ_i, and $\delta(s)$ is the distance from state s to a goal state. This is below the state space distances (1-distances) because minimization is used, not maximization. It may be difficult to do distance maximization with planning graphs as they do not represent most dependencies between state variables.

Smith and Weld's (1998) multiple planning graphs and especially their induced mutexes are related to our n-distances. They compute a kind of approximation of our n-distances, but as this computation is based on distances of state variable values as in the work by Bryce and Kambhampati, the approximation is not very good. With the multiple planning graphs there do not appear to be useful ways of controlling the accuracy parameter n, and Smith and Weld then essentially consider n that equals the number of initial states for deterministic problems.

Haslum and Geffner (2000) have defined a family of increasingly accurate heuristics for classical deterministic planning. Their accuracy parameter n is the number of state variables analogously to our parameter n of states. However, they give approximations of distances in the state space (our 1-distances), not in the belief space, and many of the phenomena important for conditional planning, like nondeterminism, do not show up in their framework.

Conclusions and Future Work

In this paper, we have presented a family of distance heuristics for planning in the belief space, and shown that this family behaves in a robust way, generalizes distance estimates based on distances in the state space, and asymptotically yields perfectly accurate estimates. Our experimental study on a number of benchmark problems indicates that on all problems considered, the estimates are a proper improvement over estimates based on state space distances, except when both are perfectly accurate.

instance	A∗				WA∗				best first					
	1-distance		2-distance		1-distance		2-distance		1-distance		2-distance		cardinality	
	time	len	time	len	time	len	time	len	time	len	time	len	time	len
sort02	0.00	1	0.00	1	0.00	1	0.00	1	0.00	1	0.00	1	0.00	1
sort03	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3
sort04	0.00	5	0.01	5	0.00	5	0.00	5	0.00	6	0.00	6	0.00	5
sort05	0.12	9	0.15	9	0.13	9	0.07	9	0.00	10	0.01	10	0.00	9
sort06	139.41	12	154.64	12	251.87	12	25.81	12	0.01	15	0.01	15	0.00	12
sort07	$> 2h$		$> 2h$		$> 2h$		$> 2h$		0.01	21	0.01	20	0.01	16
sort08	$> 2h$		$> 2h$		$> 2h$		$> 2h$		0.02	28	0.05	28	0.02	19
ring03	0.01	8	0.01	8	0.00	8	0.00	8	0.00	8	0.00	8	0.01	8
ring04	0.00	11	0.01	11	0.00	11	0.00	11	0.00	11	0.01	11	0.00	11
ring05	0.01	14	0.03	14	0.01	14	0.04	14	0.01	14	0.03	14	0.01	14
ring06	0.03	17	0.12	17	0.03	17	0.14	17	0.03	17	0.14	17	0.03	17
BW02	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3	0.00	3
BW03	0.01	7	0.00	7	0.00	7	0.01	7	0.00	7	0.01	7	0.00	7
BW04	0.71	13	0.93	13	0.04	13	0.06	14	0.02	14	0.04	14	0.03	14
BW05	180.47	17	307.62	17	1.26	17	2.87	17	0.40	22	1.41	21	0.36	34
emptyroom01	0.00	2	0.00	2	0.00	2	0.00	2	0.00	2	0.00	2	0.00	4
emptyroom02	0.00	8	0.00	8	0.00	8	0.00	8	0.00	12	0.00	8	0.00	12
emptyroom03	0.16	20	0.01	20	0.03	24	0.00	20	0.00	50	0.00	20	0.00	36
emptyroom04	37.28	44	0.07	44	10.59	52	0.06	44	0.09	222	0.06	44	0.01	106
emptyroom05	$> 2h$		0.92	92	$> 2h$		0.89	92	2.53	950	0.89	92	0.03	342

Table 2: Runtimes and plan sizes of a number of problem instances

inputs	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
gates (best known)	1	3	5	9	12	16	19	25	29	35	39	45	51	56	60
gates (our planner)	1	3	5	9	12	16	19	26	31	39	46	56	64	74	81

Table 3: Sizes of sorting networks found by best-first search and cardinality heuristic. Networks up to 8 inputs are optimal. From 9 on optimal network sizes are not known. Total runtime for 16 inputs is 5.81 seconds on a 800 MHz Pentium.

To make the use of 2-distances feasible on bigger problem instances a further investment in implementation techniques is needed. Most direct approach would be to give up exact computation of 2-distances. First, one could use generic approximation techniques for making BDDs smaller: compute a BDD of small size given upper and lower bounds for the Boolean function. Second, 2-distances can be approximated by abstracting half of the state variables away: for $\langle s_1, s_2\rangle$ ignore the even ones for s_1 and the odd ones for s_2.

So far we presented heuristics for the unobservable planning, but of course the main interest is in conditional planning with partial observability. The n-distances are not admissible for partially observable planning in general: increased observability decreases belief state distances. This will be addressed in future research.

Acknowledgements

This work was partially funded by DFG grant RI 1177/2-1.

References

Bertoli, P.; Cimatti, A.; and Roveri, M. 2001. Heuristic search + symbolic model checking = efficient conformant planning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, 467–472.

Blum, A. L., and Furst, M. L. 1997. Fast planning through planning graph analysis. *Artificial Intelligence* 90(1-2):281–300.

Bonet, B., and Geffner, H. 2000. Planning with incomplete information as heuristic search in belief space. In Chien, S.; Kambhampati, S.; and Knoblock, C. A., eds., *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems*, 52–61. AAAI Press.

Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129(1-2):5–33.

Bryce, D., and Kambhampati, S. 2003. Heuristic guidance measure for conformant planning. ICAPS'03 workshop on Planning under Uncertainty.

Haslum, P., and Geffner, H. 2000. Admissible heuristics for optimal planning. In Chien, S.; Kambhampati, S.; and Knoblock, C. A., eds., *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems*, 140–149. AAAI Press.

Knuth, D. E. 1998. *Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley Publishing Company.

Smith, D. E., and Weld, D. S. 1998. Conformant Graphplan. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98) and the Tenth Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, 889–896. AAAI Press.