
Seminar Thesis: Efficient Planning under Uncertainty with Macro-actions

Ragnar Mogk
Department of Computer Science
Technische Universität Darmstadt
ragnar.mogk@stud.tu-darmstadt.de

1 Introduction

This thesis is a summary, analysis and discussion of a paper by Ruijie He, Emma Brunskill and Nicholas Roy called Efficient Planning under Uncertainty with Macro-actions [2].

For humans it is normal to live in a world we do not fully understand nor have complete knowledge of. Even if we just consider the room we are currently in, or the near surroundings if we are in an open area, we do not know a lot about our proximity. The area we have a sharp vision of is actually very small, and for a lot of positional tasks our other senses are not very useful. This is what is meant by partially observable or, as the title of the paper calls it, uncertainty. This is a problem not only humans have, but any machine, any robot, we build has to deal with as well. We want our robots to solve problems for us in the real world so they have to deal with real world problems. Even with all the things we can not observe we can still solve a lot of problems. When crossing a road it is impossible to check both sides for incoming cars at the same time, but it is sufficient to check them one after the other because we know, or at least we believe, that in a short enough time frame the state of street will not change considerably. We can use this belief to plan our future actions – if the street is empty make a first step, if it is not wait for it to become empty. What seems easy to a human is a hard problem for a robot. There exist solutions that work well for small problems like walking across the street, because it has essentially only two states, the street is empty or not and two actions, walk or wait. The algorithm proposed in the paper claims to solve much larger real world problems. It does so by not considering every state and action by itself, but instead considers a sequence of actions and their combined effect on the future state.

The rest of this thesis tries to explain how to transfer these ideas to our robot in a slightly more formal way. Section 1.1 introduces a formal mathematical model for the problem, sections 1.2 and 1.3 illustrate ways how to tackle such problems. The main part of the paper, the Posterior Belief Distribution (PDB) algorithm, is explained in section 2 and is discussed together with some evaluation results in section 3.

1.1 Markov Decision Problems and Partially Observable Markov Decision Problems

The first thing when solving the problem is to clearly describe and at least slightly formalise it. The type of problem considered here is known as a Partially Observable Markov Decision Problem (POMDP). To understand the partially observable case, a simple Markov Decision Problem (MDP) is explained first and later extended.

The world is described as a state $s \in \mathcal{S}$ where our actor, the robot, can decide which action $a \in \mathcal{A}$ it wants to take. There are many ways to encode states s and actions a , for a simple light bulb the only states might be on and off and the only possible action is to push the switch. For a flying helicopter the state could encode xyz-coordinates, rotation, velocity, acceleration, distance to the ground, position of obstacles and so on. All these are continuous values which make the size of the state space $|\mathcal{S}|$ large. Actions for this example could be commands to the motors of the rotors of the

helicopter, which are also continuous. Generally the problem gets harder when the size of the state space and the action space $|\mathcal{S}| + |\mathcal{A}|$ gets larger. When the robot takes an action a in some state s this has normally an influence on the world, so the current state changes to a new state s' . This change is encoded by the transition function \mathcal{T} , which calculates the new state s' from the current state s and the selected action a . We have $s' = \mathcal{T}(s, a)$. The fact that the transition function only depends on the current state and not on any previous state is called the Markov Property, hence the name Markov Decision Process. Interestingly, this is not much of a limitation. If some value of a previous state is needed, it can be encoded as part of the current state, which obviously has the disadvantage of increasing the size of the state space $|\mathcal{S}|$. Often the system is non deterministic, which means that the same action a in the same state s does not always lead to the same next state s' . This means the transition function \mathcal{T} is no longer a function but a probability $\mathcal{T}(s, a, s') = p(s'|s, a)$. Here \mathcal{T} no longer returns the next state s' , but the probability of s' to be the next state. To be solvable a Markov Decision Problem also needs a reward function $r(s, a)$ which calculates the reward of taking action a in state s . To solve the problem the robot needs to take the actions which will result in the maximum expected reward. Which action the robot should take in a state is encoded in a policy π , so that $\pi(s) = a$ where a is the action to be taken in state s . The policy which results in the maximum expected reward is called the optimal policy π^* and is the solution to the MDP.

A lot of real world situations can be modelled as a Markov Decision Problem, but to solve them the robot needs complete knowledge of the state, which is in almost all cases impossible to get because of limited sensor capabilities. Instead the robot receives an observation $z \in \mathcal{Z}$ after each action, that encodes its sensor data. This observation depends on the current state, but is also often not deterministic because of noisy sensors. The observation of some part of the system gives the Partially Observable Markov Decision Problem its name. Based on the actions the robot takes and the observations it receives, the robot forms an assumption about the state it is currently in. This assumption is called the belief $b \in \mathcal{B}$ and is a probability distribution over states, so that $b(s)$ is the probability that s is the current state. A very problematic attribute of the belief space is its size, if the state space is discrete a belief encodes one probability for each of the possible states, it can be imagined as a vector containing one entry for each state. This implies that the belief space has one dimension for each possible state. Because even small problems often have a lot of states, the dimension of the belief space is large as well. This is known as the curse of dimensionality and one of the problems with POMDBs.

But even hard problems are solvable, if not exact then maybe good enough for practical purposes. The next sections will have a look at some ideas how to solve them.

1.2 Forward Search

One possible solution for a POMDP is to compute a value function V which will assign each belief a value $V(b)$. This value encodes all the future rewards from that belief onwards following an optimal policy. Once one has the Value function it is easy to get the optimal policy π^* , the robot just has to take the action that changes the belief to the one with the highest value. But computing the value function for every belief is hard, because of the size of the belief state. Instead we perform a forward search from the current belief onward to estimate the value of the most interesting beliefs, the ones in the near future.

To compute the estimated value of an initial belief b_0 a tree like structure is constructed. See figure 1 for a simple visualisation. The root of the tree is the initial belief b_0 which branches out to all possible actions. Each of these action branches again, resulting in all possible new beliefs which in turn result from the observation received after taking that action. There are multiple possible resulting beliefs after one action because neither the state transition nor the received observation are deterministic. This branching continues until a pre set depth is reached. Then the value of each belief can be computed by traversing the tree from the leaves to the root. At each action a the average of the value of the child nodes is the expected future reward. This average is weighted by the probability of the action resulting in that belief. This expectation is added to the reward of taking the action a when in the parent belief b to obtain the expected value of taking action a in belief b . The value of the belief is the maximum value obtainable by taking any of the actions, because actions can be freely chosen, unlike received observations.




Figure 1: Figure 1: And-Or tree used by forward search. The value of an action a results from the expectation of the value of the possible child beliefs. The Value of a belief results from taking the maximum value of all its child actions.

This method has the advantage of only considering beliefs actually reachable from the current state. However, the number of states to consider still grows exponentially with the size of the tree, and calculating the beliefs resulting from an action and any of the possible observations is expensive. But if the size of the tree is small only very little steps into the future are considered, the result is a robot that can not plan ahead and will take actions that are bad in the long run.

Even with this shortcomings, the underlying principle of this method is very similar to the other algorithms described in this thesis. The Posterior Belief Distribution Algorithm itself is based on this idea, but greatly reduces the computations needed to approximate the expected value of a sequence of actions. But first the next section introduces slightly different approaches to reduce the number of beliefs to consider, so called point based methods.

1.3 Point Based Methods

The paper itself refers to Successive Approximations of the Reachable Space under Optimal Policies (SARSOP) [4] and Point-Based POMDP Algorithms: Improved Analysis And Implementation (HSVI) [5] as often considered state-of-the-art. As such, this needs to be improved upon so it is useful understand them for comparisons. This section will explain these two methods briefly, before the Posterior belief distribution algorithm is explained, to have a basis to refer to. As HSVI is the older algorithm and SARSOP is a direct improvement, HSVI is explained first directly followed by the improvements in SARSOP.

The idea behind point based methods is to sample only certain points in the belief state instead of using every belief. The consequence is a reduced complexity but also a reduced quality of the results. The method with which the points to sample are chosen also has a huge influence on the quality of the solution. Prior methods used some kind of equidistant grid on the belief state, but HSVI proposes a better method. The belief state contains a lot of beliefs that can never be reached by any sequence of actions from the starting belief b_0 and the size of the space of beliefs reachable from the initial belief, $|\mathcal{R}(b_0)|$ is much smaller than the size of the belief space $|\mathcal{B}|$. The beliefs not reachable from the initial belief, $\mathcal{B} \setminus \mathcal{R}(b_0)$ have no positive influence on the result, so they can be ignored. Another observation used in HSVI that the more actions are needed to reach a certain belief b from the starting belief b_0 the less important it becomes. This allows them to build a value function starting from the initial belief and sampling from there going down the tree of beliefs reachable from the current one. The sampling stops if, because of the discounted influence of far away beliefs, the value function is not expected to improve further. Then the sampling might continue with another path, until the bounds for the calculated value function are exact enough. Which path to sample is dictated by heuristics which guess possible improvements of the value function.

SARSOP takes the idea of sampling only from reachable beliefs $\mathcal{R}(b_0)$ one step further and tries to sample only from beliefs reachable with an optimal sequence of actions $\mathcal{R}^*(b_0)$. But the optimal sequence of actions is already the solution to the problem. SARSOP uses this fact and only tries to estimate $\mathcal{R}^*(b_0)$ from which the optimal policy can be extracted. Estimating $\mathcal{R}^*(b_0)$ also uses

heuristics and prior samples to select beliefs which are likely in $\mathcal{R}^*(b_0)$ and reject beliefs which are not. This improves the quality of the solution in every step and as such has an interesting property a lot of iterative algorithms have, if the calculation takes too long one can simply stop the execution and take the current results. These results might be less than optimal, but in most cases they are still better than nothing at all.

Contrary to these very algorithmic and somewhat heuristic methods, the Posterior Belief Distribution Algorithm described in the next chapter, takes a more analytical approach.

2 The Posterior Belief Distribution Algorithm

The Posterior Belief Distribution Algorithm (PBD) in its core is very similar to forward search, but it reduces the computational complexity by removing the exponential growth of the search space when increasing the search horizon, this allows to look further into the future. The first way to do this, is with macro actions. Macro actions are not specific to PBD and as such are explained in a more general way. There are basically two reasons the search tree grows exponentially with its height, first there are multiple possible actions at each time step, second each action can result in one of multiple observations. Macro actions define a sequence of actions which is always completely executed before considering any new action. This means that branching inside one macro action only occurs because of different observations, but no longer because of different actions. Branching naturally still does occur when one macro action has finished and the next one is considered, but if macro actions are long enough it might be sufficient to only consider one of them to get good results.

This reduces the number of branches but introduces a new problem, the generation of macro actions. As macro actions only have to be constructed once for each problem, we will just consider them to be given to us together with the other problem parameters. Still, there is research to generate them automatically, see [1].

There might be more macro actions for the whole system than for each individual state. Some actions might make no sense for some state so they will never be considered in that state.

But even with macro actions, the exponential growth of the search tree remains, because of the branching of the beliefs. Sampling of the remaining tree as done in HSVI and SARSOP is a solution, but is still quite expensive and depends on the quality of the chosen samples. There is a third solution, computing the effect of one macro action analytically. But to do this, we need to introduce a new concept belief distributions.

2.1 Belief Distributions

There are two steps to understand belief distributions, the first is to understand what a belief distribution is, and the second how to calculate them. First let us reconsider what a belief is. For discrete states a belief essentially gives each state a probability. So state s_n has probability p_n the sum of these probabilities has to be 1 to make the belief a proper distribution. For continuous state the belief encodes the distribution of the state variables. In a lot of cases the belief is just a Gaussian (normal) distribution which can be easily described by its mean and covariance. If so, the belief distribution is a distribution over the mean and the covariance. This is very useful, because a single belief distribution can encode the outcome of a complete macro action and it is not much more complicated than a normal belief. In fact if the system is modelled in such a way, that all the beliefs are Gaussian distributions, then the belief distribution itself is also Gaussian. This makes it comparatively easy to handle.

We now assume that the state transition and the observation are both linear functions with some Gaussian noise added. This assumption causes the beliefs to have the wanted property of being Gaussian distributions themselves, so that $b_t = \mathcal{N}(\mu_t, \Sigma_t)$. With this model [3] describes a method, the Kalman filter, which is expanded for this problem in [2], which allows to directly compute the mean μ_{t+1} of the next belief from the mean μ_t of the current belief, depending on the received observation z . The covariance can also be computed from the prior belief and is independent of any received observation. As the observation z itself is drawn from a Gaussian distribution, if z is substituted in the calculation for the next mean μ_{t+1} it follows, that μ_{t+1} itself becomes a

random variable. This is what a belief distribution has been introduced as earlier in this section, a distribution of state distributions. To illustrate this point further, if one samples a value from the distribution of means, that is equivalent to receiving one specific observation, as it is just the noise of the observation that is encoded in the belief distribution. But a macro action consists of more than one action and for the second action we no longer have a current belief, but a belief distribution. To compute the next belief distribution we have to integrate over the current belief distribution which can be illustrated as looking at every possible belief and combine all the possible outcomes the action, dictated by the macro action, can have in that belief and then combining them back into one belief distribution. Fortunately, we do not need to calculate the integral explicitly but can instead combine both distributions analytically, as they are both Gaussian. The result is also Gaussian, our new belief distribution. This can be done iteratively for each action in the macro action and will result in one final belief distribution, which is, again, Gaussian distributed.

This is quite useful, as compared to plain forward search or point based methods, we do not need to compute whole trees of beliefs or have multiple paths to sample values from, but can just iterate this computation of belief distributions. This means that the complexity of the calculation is linear in the depth of the search horizon, which allows to look further into the future than the other methods would allow. But we still need to actually make use of our new knowledge of the future, we need to calculate which of the possible macro actions, which determine the resulting belief distribution, has the highest expected reward.

2.2 Calculating Expected Reward

One of the assumptions we make, is that the best solution to the task also yields the highest reward. This is mostly a practical concern, because the best solution is defined by giving at the highest reward. So the important thing to know is which of the possible actions has the highest expected reward.

The reward of a macro action $\tilde{a}_{1:L}$ consisting of the actions a_1, a_2, \dots, a_L can be written as

$$r(b_0, \tilde{a}_{1:L}) = r(b_0, a_1) + \gamma \int_{z_1} p(z_1 | b_0, a) Q(b^{a_1, z_1}, \tilde{a}_{2:L}) \quad (1)$$

where $b^{a,z}$ is the updated belief after taking an action a and receiving an observation z . The reward of a macro action is the sum of the reward for taking the first action a_1 and the expected reward for taking the rest of the actions $\tilde{a}_{2:L}$. This expected reward is discounted by a factor γ which has to be less than 1 if a macro action had infinite length, to guarantee a convergence of the expected reward. In this case it could just be 1. More interesting is the integral which is like a weighted average over all possible future rewards. The Q function calculates the future reward for a belief and a macro action and is very similar to r , the difference being that it also handles the case of just getting an action instead of a macro action. To illustrate this formula, in the case of two discrete observations z_1 and z_2 we just have to possible outcomes after taking action a_1 in belief b_0 . We then just sum up the expected reward of both of this new beliefs b^{a_1, z_1} and b^{a_1, z_2} weighted by the probability of the respective observation. So if one possible outcome is good but unlikely and the other is bad and likely, the current belief will also be bad, but not as much as there is a small chance to end in a good belief.

It is possible to rewrite (1) after expanding the recursive definition of Q to

$$r(b_0, \tilde{a}_{1:L}) = r(b_0, a_1) + \sum_{i=2}^L \gamma^{i-1} r(b_{dist}^{i-1}, a_i) \quad (2)$$

where b_{dist}^i is the belief distribution after i steps of the macro action. Some explanation is needed why these definitions are equivalent, for that this definition is needed:

$$r(b_{dist}, a_i) = \int_b \int_s b(s) b_{dist}(b) r(s, a_i) ds db \quad (3)$$

This is the expected reward distribution and it is calculated as one would expect. We integrate over all possible beliefs and states because that is what a belief distribution contains, and we sum up every reward we get in every possible state, weighted by the probability of the state given the belief distribution. To put it even more simply, if a state has a reward of 4, and this state has a probability

of 0.2 in a belief that itself has a probability of 0.5 in a belief distribution, this state contributes a reward of $0.2 \cdot 0.5 \cdot 4 = 0.4$ to the total reward of the belief distribution. So what equation (2) does is to just sum the expected reward of each belief distribution on the way from the initial belief, along the macro action, to the last belief distribution in which an action is executed. But calculating the nested integrals over all states and beliefs is not something we want to do explicitly, as the state space might be very large and the belief space is even larger. But as in section 2.1 when the belief b and the belief distribution b_{dist} are Gaussian, then the two distributions in $r(b_{dist}, a)$ can be combined analytically and if the reward model is a sum of Gaussians or a polynomial function of the state, then $r(b_{dist}, a)$ itself can also be computed analytically. Arbitrary reward models may be approximated by sampling, but that might lose some of the efficiency won by doing direct computations.

We can now evaluate a macro action, it just remains to combine them and we can start making decisions.

2.3 Making Decisions

If we just evaluate each macro action at a time step and execute the best one, we run into similar problems with only evaluating a single action. The algorithm loses far-sightedness, it makes decisions that are good in the short term, but may never find a better solution that requires some worse actions first. This is why we want to plan multiple macro actions ahead. To do this, we can use the methods we have already covered in the introduction, forward search and sampling. We can not directly build the forward search tree out of macro actions and belief distributions, as there might not exist a single optimal macro action for a whole belief distribution. What we do instead is to sample some beliefs from the resulting belief distribution. This is like sampling a sequence of observations, but as the belief distribution is just a Gaussian we do not need as many samples to get good results, as sampling observation sequences would need. For each of the sampled beliefs we can then evaluate the macro actions again. So instead of a tree consisting of belief, action, belief, action, the new forward search tree consists of belief, macro action, belief distribution, sampled beliefs and then repeat. We can then again take the maximum reward over all actions and the expectation over the sampled beliefs, as in the normal forward search tree.

Now with a possible lookahead into the far future we can evaluate which macro action sequence will likely have the greatest reward. With that knowledge the first action of the first macro action is executed and a new belief is calculated from the received observation. As this information might improve the knowledge of the world, the selected macro action is abandoned and a new best path is calculated, using the newly obtained knowledge.

This concludes the introduction to the Posterior Belief Distribution Algorithm, the next section tries to fit this algorithm into its place for useful applications.

2.4 Where it Fits

The Posterior Belief Distribution Algorithm is designed to solve large Partially Observable Markov Decisions Problems. Large in this case means that it has a lot of states. To do this we need some things first. The most obvious is an actual problem. Typical examples include robots trying to navigate in a room with obstacles and find something with a position that is unknown, or a helicopter flying over some target area observing multiple vehicles. One important part of these problems is that they have some parts that are either unknown from the start and have to be discovered first, or that parts of the state is changing while the robot is unable to observe it. A lot of real world situations come with these problems. But PBD can not solve arbitrary of these problems. The state transition and observation model have to be representable as linear functions with Gaussian noise. While this is often possible, especially if one just approximate these functions, this might not work in all cases. Besides the state transition and observation model, we also need a reward function and a set of macro actions which can both be calculated prior to the algorithm. PBD is also a method that only works online, it needs the actual observations it gets from the sensors to update its current belief of the state.

Even though there are all this preconditions PBD can be used in a lot of cases as seen in the next section.

3 Evaluation

For the performance of an algorithm there are two important aspects, one is how fast a solution is found in terms of calculation time. The other is how good the solution is in terms of the expected reward. What we want is to find the optimal solution instantly, but what we get is often a trade-off between time and quality, sacrificing one for the other. The PDB algorithm is comparatively fast and gives good solutions. Fast in this case means, it is not exponential in the search horizon, but only in the search horizon of the macro actions. For long macro actions this greatly reduces the computational complexity. It does so by employ some preconditions that hold for a lot of the target problems. But as it uses approximations, one example is sampling from the belief distributions, it is very unlikely to find an optimal policy.

There is actually no lower bound given for the quality of the generated policy, but it is proven, that the calculated value of a belief is not much bigger than the optimal value of the belief. This does not give any guarantees about the other direction, the calculated value might be too small. Still, some encouraging experimental results are presented in the next section. Keep in mind that these evaluations assume that good macro actions are already known beforehand and do not evaluate the cost of constructing them, or the influence of different sets of macro actions.

3.1 Experimental Results

This sections serves two purposes, the primary one is to illustrate the performance of the algorithm, the secondary is to give an overview over the kinds of problems which satisfy the conditions needed by the algorithm, a linear Gaussian state transition and observation model.

The first experiment is a modified variant of the Rocksample problem described in [5]. In the basic version of the Rocksample problem there exists an area with rocks with known locations. However some of the rocks are interesting and some are not. This information is unknown at the beginning and can be obtained with a noisy sensor, that works better when near a rock. As sampling is expensive, only the interesting rocks should be sampled. This problem does not require planning into the far future as the information is gathered when near the rocks. The modified version of this problem, called Information Search Rocksample (ISRS), introduces so called information beacons. The robot only gets information about the rocks when near an information beacon, this forces some planning ahead, especially if the information beacons are far away from the rocks. A possible set of macro actions for this experiment is one that allows the robot to move to each of the rocks or beacons. This means that the forward search graph only encodes the decisions in which order the interesting points are visited, and not how to move from one to the other, macro actions provide a useful abstraction over less relevant parts of the problem. PBD was compared to the point-based algorithm SARSOP [4] presented in section 1.2. SARSOP was not able, even with substantial computation time to look far enough into the future to make efficient use of the information beacons. Instead it generates a policy that directly moves to the rocks using only the unreliable information it receives. PBD on the other hand does generate a policy that visits the beacons and so generates higher average rewards. If SARSOP is modified to also use macro actions instead of single steps, it actually generates a better policy than PBD but also needs about two orders of magnitude longer. However, if one increases the size of the ISRS problem, more rocks and a bigger area, PBD still generates good results quite fast, while SARSOP reaches its limits.

The second experiment is target monitoring. Here a helicopter monitors some targets moving on the plane below him with a downward facing camera. The higher above ground the helicopter is the bigger the area it can monitor, but the measurement of the target position becomes less accurate. There are some areas of interest and the helicopter gets a reward if it reports a target inside one of these areas. However if it reports a target as inside while it is not it will get a penalty (a negative reward). The macro actions in this case were to move to one of the assumed positions of the targets in certain heights and one for just hovering. The policy generated by PBD is actually quite intuitive, it hovers in place to conserve energy when nothing happens, moving also has a small negative reward, if one of the targets might enter one of the areas of interest it flies near that target to exactly monitor the point where it enters the area. In tests with multiple targets the policy generated by PBD was the only one which did not miss most of the entries into areas of interest.

Even though these experimental results show that PBD works as intended there are still limitations, some of which are discussed in the next section.

3.2 Discussion

As seen in the previous section the precondition of a linear Gaussian system is not a limitation for a lot of examples, as they can be modelled in a suitable way. PBD also generates quite good policies in a comparatively short time. But a lot of the quality seems to be from the good selection of macro actions. As the macro actions are reflected in the policy, selecting macro actions that do what is expected will also cause the policy to be limited to these good actions. This is one of the big advantages of PBD but if there are no good macro actions, policy generation will also be worse. It is also important to limit the number of macro actions, as many of them can cause performance problems.

Performance is another interesting point, both of the evaluated experiments are quite simple compared to what one would expect in a real world scenario. To actually monitor a coastline of some kilometers length with a reasonable resolution of some meters, one would need a state space that is probably too big to be handled in a reasonable time. On the other hand, the authors did not yet optimize the algorithm specifically for performance, so there might be still some big improvements.

So even PBD is an improvement, there is still work to be done, especially on integrating automatic macro action generation and performance enhancements.

4 Conclusion

In this thesis a summary of the Posterior Belief Distribution Algorithm which is a forward search algorithm, that uses macro actions and belief distributions to analytically compute expected rewards in a Partially Observable Markov Decision Problem. It was demonstrated that PBD can be used to solve comparatively large problems with good results and can also be applied to a variety of different problems.

References

- [1] Ruijie He, Emma Brunskill, and Nicholas Roy. PUMA: Planning under uncertainty with macro-actions. In Maria Fox and David Poole, editors, *AAAI*. AAAI Press, 2010.
- [2] Ruijie He, Emma Brunskill, and Nicholas Roy. Efficient planning under uncertainty with macro-actions. *J. Artif. Int. Res.*, 40(1):523–570, January 2011.
- [3] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Basic Eng.*, 82:35–45, 1960.
- [4] Hanna Kurniawati, David Hsu, and Wee Sun Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces, 2008.
- [5] Trey Smith and et al. Point-based pomdp algorithms: Improved analysis and implementation.