



Partially Observable Markov Decision Processes

[Tutorial](#) | [Papers](#) | [Talks](#) | [Code](#) | [Repository](#)
[Back](#) | [POMDP Tutorial](#) | [Next](#)

Littman, et al's Witness Algorithm

The next approach, Littman, et al (1994), attacked the problem slightly differently. This algorithm, called *witness*, uses the same basic structure as Sondik and Cheng; i.e., it defines regions for a vector and looks for a point where that vector is not dominant.

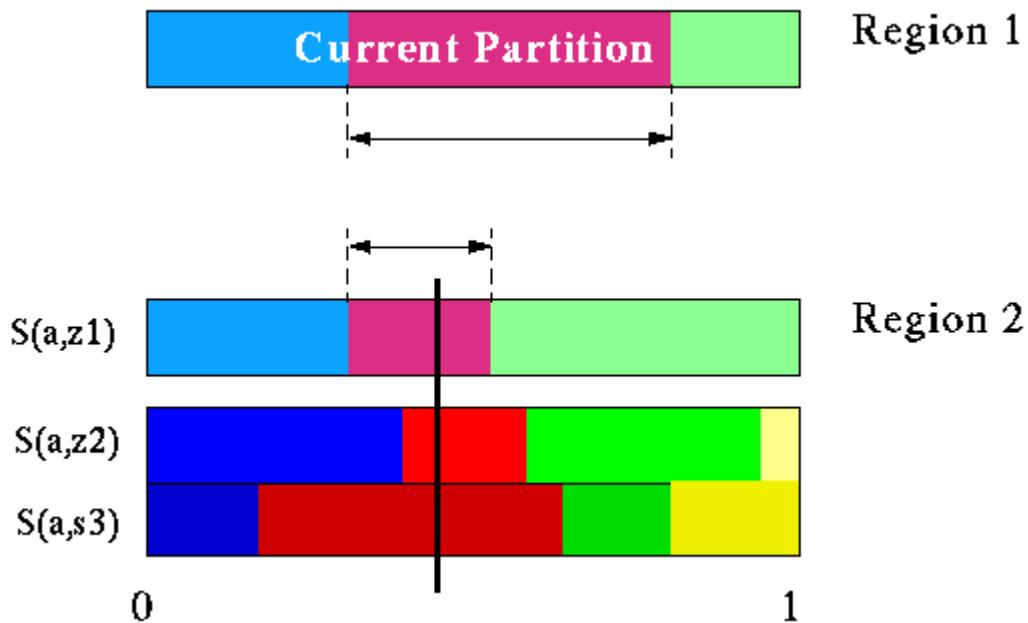
Unlike Sondik's algorithm it doesn't worry about all the actions all the time. It concentrate on finding the best value function for each of the actions separately. Once it find these it will combine them into the final v' value function.

Another way it simplifies the problem, aside from doing each action in isolation, is that it also only worries about one observation at a time. Remember that a vector at a belief point is constructed by selecting a transformed vector from v for each observation. The witness algorithm, like Cheng's algorithm, starts with an arbitrary belief point and generates its vector. It adds this vector to a set and assumes that this set is v' it then goes about trying to prove or disprove whether or not this set truly is v' .

In constructing a vector we make a choice for each observation; the choice being the selection of one of the v vectors representing a particular future strategy. The algorithm then looks at the individual choices made, observation by observation to see where a different choice would yield a better value. Like the other algorithms, it defines the region where it is assured that the particular choice is best. If it can find a belief point where a different strategy would be better, then this serves as a witness to the fact that our current set of vectors is not yet the real v' we are after.

The regions defined by this algorithm are also best describe by the intersection of two simpler regions. We have actually seen these two regions before. One is exactly the same as the regions that Cheng's algorithm defines; namely the region over which the current vector dominates all the other vectors we have generated up to this point. This region is shown in the top of the next figure. Note that at this point in this figure, we are at a point in the algorithm where we have generated three vectors.

Stated another way: this region simply reduces the search for a point to the region over the current approximation we have (i.e., all the vectors we have generated to this point). This prevents us from looking for better points in a region where we already know there are better points.



Witness regions

The second region defines the region over which the current future strategy will be best for the observation z_1 . If we wander out of this region, then we could substitute a different course of action for this observation and get a better value. When considering the strategy for one observation, we ignore the other observations. This second region, is actually a small subset of one of Sondik's regions. However, if we find a belief point where we could substitute another future strategy, we cannot simply plug in this new strategy. Because we are considering this in isolation from the other observations, it could be that in addition to changing the choice for this observation, we could also change the choice for another observation. Thus, finding a belief point where the current observation's choice could be changed, just gives us a witness to the fact that there is a point where we can do better. We can then take this point and generate the real best vector for it (taking into account all the observation choices).

We repeat this for every observation to get a set of vectors for a particular action and then we do this for all actions. We can then combine and eliminate dominated vectors to get the real set we were after, v' .

Continue

[Back](#) | [POMDP Tutorial](#) | [Next](#)