A* Variants

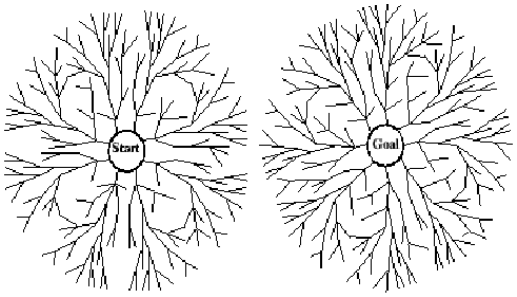

Bidirectional search

Search from the start to the goal AND from the goal to the start

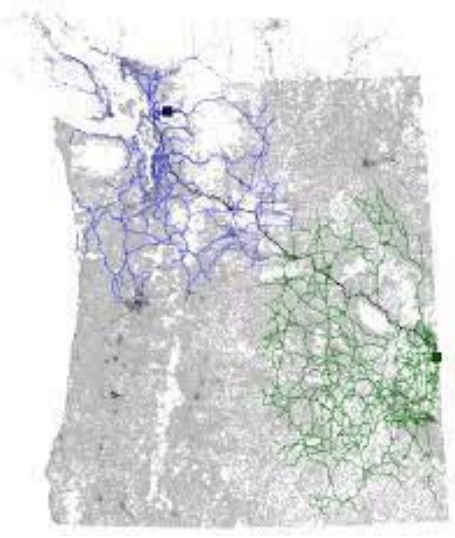

~O($b^{d/2}$) - significantly less than unidirectional search

If heuristic is *consistent*, you are finished when both sides expand the same node

A* Variants

Beam search

Limit the size of the frontier to the nodes with the lowest cost

Not complete

Not optimal

Used anyway, especially in very open environments

Beam Search for text generation (NLP)

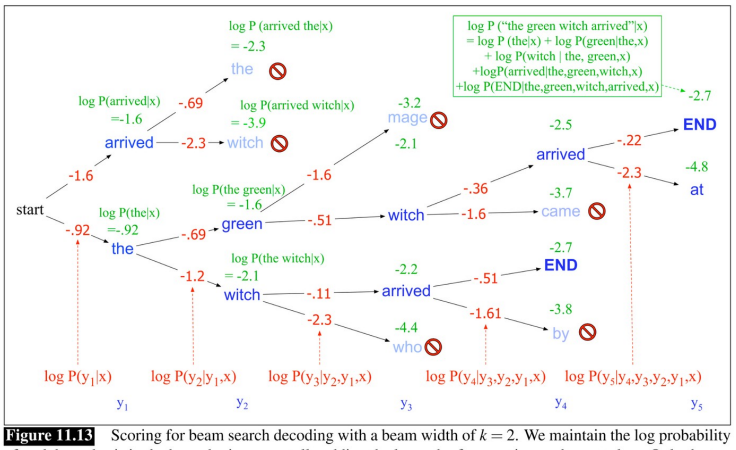

Figure 11.13 Scoring for beam search decoding with a beam width of $k = 2$. We maintain the log probability of each hypothesis in the beam by incrementally adding the logprob of generating each next token. Only the top k paths are extended to the next step.

A* Variants

Weighted A*

Modify the emphasis on the heuristic function

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

Sometimes we change λ as we get closer to the goal

Be careful modifying $\lambda > 1$

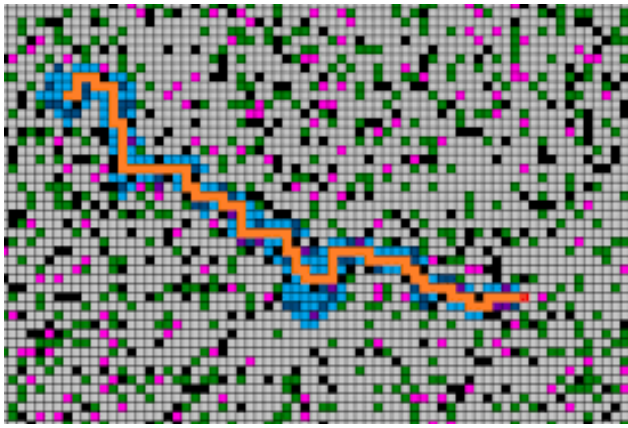

A* Variants

Iterative Deepening A* (IDA*)

Set an **f(n)** limit, see if a solution arises

- If so, return that path
- Otherwise, increase limit

Lower spatial requirements – **O(d)** vs **O(b^d)**