Informed (Heuristic) Search Algorithms

The generic informed/heuristic search algorithm is called best-first search. Best-first search works just like uniform-cost search (UCS) except we store our frontier as a priority queue sorted by an *evaluation function* known as f(n). Just like UCS, best-first search chooses to examine nodes with the lowest values of f(n) first. Where this algorithm differs from UCS is that f(n) is an estimate of the lowest-cost of path from the initial state, to the state at node n, to any goal state. Recall that UCS's frontier is sorted by g(n), which is the lowest cost of the path from the initial state to the state at node n, so UCS does not try to estimate the cost of the path *after* n towards the goal; it only takes into account the cost of the path *before* n.

Best-first search typically uses a heuristic function, denoted h(n), as an estimate of the cost from node n to a goal state. By changing the definition of f(n) to various functions involving g(n) and/or h(n), we obtain three different algorithms:

If we define $f(n) = g(n)$, best-first search degrades to uniform cost search.
If we define $f(n) = h(n)$, best-first search becomes an algorithm called greedy best-first search.
If we define $f(n) = g(n) + h(n)$, best-first search becomes an algorithm called A* search (or just A*, pronounced "A-star").

Admissibility and consistency of the heuristic function

An admissible heuristic is one that **never overestimates the cost to reach the goal**. Because h(n) estimates the cost from node n to any goal state, h(n) must never be greater than the true cost from n along the cheapest path to a goal state.

A consistent heuristic is, informally, one that always decreases in a "consistent" manner as one moves along a path from the initial state to the goal state(s). Formally, a heuristic h(n) is consistent if, for every node n and every successor n' of n, the estimated cost of reaching the goal from n is no greater than the step cost from n to n' plus the estimated cost of reaching the goal from n': $h(n) \le COST(n, a, n') + h(n')$ or equivalently, $h(n) - h(n') \le COST(n, a, n')$.

Another way to interpret this is a consistent heuristic **never overestimates the cost of a single step** from n to n'.