A Project on Fast Trajectory Replanning for Computer Games for "Introduction to Artificial Intelligence" Classes

Sven Koenig and William Yeoh
Department of Computer Science
University of Southern California
300 Henry Salvatori Computer Science Center (SAL)
941 W 37th Street
Los Angeles, CA 90089-0781
{skoenig, wyeoh}@usc.edu

Abstract

This stand-alone path-planning project for an undergraduate or graduate artificial intelligence class relates to video game technologies and is part of our effort to use computer games as a motivator in projects without the students having to use game engines. Heuristic search and, in particular, A* are among the important single-agent search techniques and thus good candidates for a project in artificial intelligence. In this project, the students need to code A* and then extend it to Adaptive A*, a fast trajectory replanning algorithm, to move game characters in initially unknown gridworlds to a given target. Adaptive A* is an incremental version of A* that often searches faster than A* since it updates the heuristics between searches to find solutions to series of similar search tasks potentially faster than is possible by solving each search task from scratch. This project requires students to develop a deep understanding of A* and heuristics to answer questions that are not yet covered in textbooks. The project is versatile since it allows for theoretical questions and implementations. We list a variety of possible project choices, including easy and difficult questions.

Introduction

The Department of Computer Science at the University of Southern California created a Bachelor's Program in Computer Science (Games) and a Master's Program in Computer Science (Game Development), which not only provide students with all the necessary computer science knowledge and skills for working anywhere in industry or pursuing advanced degrees but also enable them to be immediately productive in the game-development industry. They consist of regular computer science classes, game-engineering classes, game-design classes, game cross-disciplinary classes and a final game project. See Zyda and Koenig, Teaching Artificial Intelligence Playfully, Proceedings of the AAAI-08 Education Colloquium, 2008 for more information. The undergraduate and graduate versions of the "Introduction to Artificial Intelligence" class at the University of Southern California are regular computer science classes that are part of this curriculum. We are now slowly converting these classes to use games as the domain for projects because games motivate students, which we believe increases enrollment and retention and helps us to educate better computer scientists.

Artificial intelligence becomes more and more important for computer games, now that games use graphics libraries that produce stunning graphics and thus no longer gain much of a competitive

advantage via their graphics capabilities. Many games need trajectory-planning capabilities and thus use search algorithms. Some games already use machine learning or planning algorithms. For example, "Black and White" uses a combination of inductive learning of decision trees and reinforcement learning with neural networks, and "F.E.A.R" uses goal-oriented action planning. Thus, games can be used to illustrate many areas of artificial intelligence.

It was tempting for us to use a publicly available game engine throughout the class and then ask the students to perform several projects in it, such as a search project to plan the trajectories of the game characters and a machine-learning project to make them adapt to the behavior of their opponents. However, students familiar with game development already have plenty of exposure to game engines while students not interested in game development do not need the additional overhead. To put all students on the same footing and allow them to concentrate on the material taught in class, we decided to go with several small projects that do not need large code bases.

This technical report describes one particular project from the undergraduate and graduate versions of the "Introduction to Artificial Intelligence" class. Heuristic search and, in particular, A* are among the important single-agent search techniques and thus good candidates for a project in artificial intelligence. We therefore developed a project where the students use a generalization of A*, Adaptive A*, to move game characters in initially unknown gridworlds to a given target. The students need to code A* and then develop a deep understanding of A* and heuristics to answer questions that are not yet covered in textbooks and require computational thinking when they extend A* to Adaptive A*. Some of the questions, such as Question 3, discuss phenomena that were initially surprising even to researchers in the area of heuristic search. Overall, the search project is versatile since it allows for theoretical questions and for implementations. We list a variety of possible project choices, including easy and difficult questions. Teachers need to select among them since we list too many of them for a project of a reasonable size and some of them are difficult research questions. (Questions tagged with asterisks are the ones that we tend to use in our own projects.) Teachers also need to provide information on which programming language(s) the students are allowed to use, how they can determine the runtime of their programs, what they need to submit for their projects (for example, whether they need to submit their programs or traces of their programs on example search problems), how they should submit their projects (for example, on paper or electronically) and by when they need to submit their projects.

If you are using this assignment in your class or have any questions, comments or corrections, please send us an email at skoenig@usc.edu. If you need to cite this assignment (in your own publications or to give us credit), please cite this technical report as Sven Koenig and William Yeoh, A Project on Fast Trajectory Replanning for Computer Games for "Introduction to Artificial Intelligence" Classes, Technical Report, Department of Computer Science, University of Southern California, Los Angeles (California, USA). Updates, errata and supporting material for the project (including the latex version of this document to allow teaches to customize this assignment) can be found on our webpages idm-lab.org/gameai, where we will also release additional projects in the future. If there is sufficient interest, we will create sample solutions for teachers at accredited colleges and universities.

Acknowledgments

1   Project: Fast Trajectory Replanning for Computer Games

Figure 1: Total Annihilation by Cavedog

Consider characters in real-time computer games, such as Total Annihilation shown in Figure 1. To make them easy to control, the player can click on known or unknown terrain, and the game characters then move autonomously to the location that the player clicked on. They always observe the terrain within their limited field of view and then remember it for future use but do not know the terrain initially (due to "fog of war").

We study a variant of this search problem in a gridworld where an agent has to move from its current cell to the given cell of a non-moving target. Gridworlds are commonly used in real-time computer games. They are discretizations of terrain into square cells that are either blocked or unblocked. The initial cell of the agent is unblocked. The agent can move from its current cell in the four main compass directions (east, south, west and north) to any adjacent cell, as long as that cell is unblocked and still part of the gridworld. All moves take one time step for the agent and thus have cost one. The agent always knows which (unblocked) cell it is in and which (unblocked) cell the target is in. The agent knows that blocked cells remain blocked and unblocked cells remain unblocked but does not know initially which cells are blocked. However, it can always observe the blockage status of its four adjacent cells (which is its field of view) and remember this information for future use. The objective of the agent is to reach the target.

A common-sense and tractable movement strategy for the agent is the following one: The agent assumes that cells are unblocked unless it has already observed them to be blocked and uses search with the freespace assumption, that is, always moves along a path that satisfies three properties: 1) It is a path from the current cell of the agent to the target. 2) It is a path that the agent does not know to be blocked and thus assumes to be unblocked (= a presumed unblocked path). 3) It is a shortest such path. Whenever the agent observes additional blocked cells while it follows its current path, it remembers this information for future use. If such cells block its current path, then its current path might no longer be a shortest presumed unblocked path from the current cell of the agent to the target. Then, the agent stops moving along its current path, searches for another shortest presumed unblocked path from its current cell to the target, taking into account the blocked cells that it knows about, and then moves along this path. The cycle stops when the agent either reaches the target or determines that it cannot reach the target because there is no presumed unblocked path from its current cell to the target and it is thus separated from the target by blocked cells. In the former case, the agent reports that it reached the target. In the latter case, it reports that it cannot reach the target. This movement strategy has two desirable properties: 1) The agent is guaranteed to reach the target if it is not separated from it by blocked cells. 2) The trajectory is provably short (but not necessarily optimal). 3)
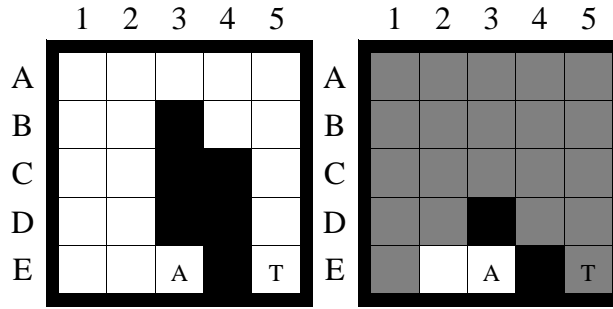
Figure 2: First Example Search Problem (left) and Initial Knowledge of the Agent (right)

The trajectory is believable since the movement of the agent is directed toward the target and takes the blockage status of all observed cells into account but not the blockage status of any unobserved cell.

As an example, consider the gridworld of size 5 × 5 shown in Figure 2 (left). Black cells are blocked, and white cells are unblocked. The initial cell of the agent is marked A, and the target is marked T. The initial knowledge of the agent about blocked cells is shown in Figure 2 (right). The agent knows black cells to be blocked and white cells to be unblocked. It does not know whether grey cells are blocked or unblocked. The trajectory of the agent is shown in Figure 3. The left figures show the actual gridworld. The center figures show the knowledge of the agent about blocked cells. The right figures again show the knowledge of the agent about blocked cells, except that all cells for which it does not know whether they are blocked or unblocked are now shown in white since the agent assumes that they are unblocked. The arrows show the shortest presumed unblocked paths that the agent attempts to follow. The agent needs to find another shortest presumed unblocked path from its current cell to the target whenever it observes its current path to be blocked. The agent finds such a path by finding a shortest path from its current cell to the target in the right figure. The resulting paths are shown in bold directly after they were computed. For example, at time step 1, the agent searches for a shortest presumed unblocked path and then moves along it for three moves (first search). At time step 4, the agent searches for another shortest presumed unblocked path since it observed its current path to be blocked and then moves along it for one move (second search). At time step 5, the agent searches for another shortest presumed unblocked path (third search), and so on. When the agent reaches the target it has observed the blockage status of every cell although this is not the case in general.

1.1   Modeling and Solving the Problem

The state spaces of the search problems are simple: The states correspond to the cells, and the actions allow the agent to move from cell to cell. Initially, all action costs are one. When the agent observes a blocked cell for the first time, it increases the action costs of all actions that enter or leave the corresponding state from one to infinity or, alternatively, removes the actions. A shortest path in this state space then is a shortest presumed unblocked path in the gridworld.

Thus, the agent needs to search in state spaces in which action costs can increase or, alternatively, actions can be removed. The agent searches for a shortest path in the state space whenever the length of its current path increases (to infinity). Thus, the agent (of the typically many agents in real-time computer games) has to search repeatedly until it reaches the target. It is therefore

Figure 3: Trajectory of the Agent for the First Example Search Problem

important for the searches to be as fast as possible to ensure that the agent responds quickly and moves smoothly even on computers with slow processors in situations where the other components of real-time computer games (such as the graphics and user interface) use most of the available processor cycles. Moore's law does not solve this problem since the number of game characters of real-time computer games will likely grow at least as quickly as the processor speed. In the following, we use A* to determine the shortest paths, resulting in Repeated A*.

```
1   procedure ComputePath()
2   while g(s_goal) > min_{s'∈OPEN}(g(s') + h(s'))
3      remove a state s with the smallest f-value g(s) + h(s) from OPEN;
4      CLOSED := CLOSED ∪ {s};
5      for all actions a ∈ A(s)
6         if search(succ(s, a)) < counter
7            g(succ(s, a)) := ∞;
8            search(succ(s, a)) := counter;
9         if g(succ(s, a)) > g(s) + c(s, a)
10           g(succ(s, a)) := g(s) + c(s, a);
11           tree(succ(s, a)) := s;
12           if succ(s, a) is in OPEN then remove it from OPEN;
13           insert succ(s, a) into OPEN with f-value g(succ(s, a)) + h(succ(s, a));

14  procedure Main()
15  counter := 0;
16  for all states s ∈ S
17     search(s) := 0;
18  while s_start ≠ s_goal
19     counter := counter + 1;
20     g(s_start) := 0;
21     search(s_start) := counter;
22     g(s_goal) := ∞;
23     search(s_goal) := counter;
24     OPEN := CLOSED := ∅;
25     insert s_start into OPEN with f-value g(s_start) + h(s_start);
26     ComputePath();
27     if OPEN = ∅
28        print "I cannot reach the target.";
29        stop;
30     follow the tree-pointers from s_goal to s_start and then move the agent along the resulting path
           from s_start to s_goal until it reaches s_goal or one or more action costs on the path increase;
31     set s_start to the current state of the agent (if it moved);
32     update the increased action costs (if any);
33  print "I reached the target.";
34  stop;
```

Figure 4: Pseudocode of Repeated Forward A*

A* can search either from the current cell of the agent toward the target (= forward), resulting in Repeated Forward A*, or from the target toward the current cell of the agent (= backward), resulting in Repeated Backward A*.

1.2 Repeated Forward A*

The pseudocode of Repeated Forward A* is shown in Figure 4. (The minimum over an empty set is infinity on Line 2.) It performs the A* searches in ComputePath(). A* is described in your textbook and therefore only briefly discussed in the following, using the following notation that can be used to describe general search problems rather than only search problems in gridworlds: S denotes the finite set of states. $s_{start} \in S$ denotes the start state of the A* search (which is the current state of the agent), and $s_{goal} \in S$ denotes the goal state of the A* search (which is the state of the target). $A(s)$ denotes the finite set of actions that can be executed in state $s \in S$. $c(s, a) > 0$ denotes the action cost of executing action $a \in A(s)$ in state $s \in S$, and $succ(s, a) \in S$ denotes the resulting successor state.

A* maintains five values for all states s that it encounters: a g-value $g(s)$ (which is infinity initially), which is the length of the shortest path from the start state to state s found by the A* search and thus an upper bound on the distance from the start state to state s; an h-value (= heuristic) $h(s)$ (which is user-supplied and does not change), which estimates the goal distance of state s (= the distance from state s to the goal state); an f-value $f(s) := g(s) + h(s)$, which estimates the distance from the start state via state s to the goal state; a tree-pointer $tree(s)$

(which is undefined initially), which is necessary to identify a shortest path after the A* search; and a search-value $search(s)$, which is described below.

A* maintains an open list (a priority queue which contains only the start state initially). A* identifies a state s with the smallest f-value in the open list [Line 2]. If the f-value of state s is no smaller than the g-value of the goal state, then the A* search is over. Otherwise, A* removes state s from the open list [Line 3] and expands it. We say that it expands state s when it inserts state s into the closed list (a set which is empty initially) [Line 4] and then performs the following operations for all actions that can be executed in state s and result in a successor state whose g-value is larger than the g-value of state s plus the action cost [Lines 5-13]: First, it sets the g-value of the successor state to the g-value of state s plus the action cost [Line 10]. Second, it sets the tree-pointer of the successor state to (point to) state s [Line 11]. Finally, it inserts the successor state into the open list or, if it was there already, changes its priority [Line 12-13]. (We say that it generates a state when it inserts the state for the first time into the open list.) It then repeats the procedure.

Remember that h-values $h(s)$ are consistent (= monotone) iff they satisfy the triangle inequalities $h(s_{\text{goal}}) = 0$ and $h(s) \leq c(s, a) + h(succ(s, a))$ for all states s with $s \neq s_{\text{goal}}$ and all actions a that can be executed in state s. Consistent h-values are admissible (= do not overestimate the goal distances). A* searches with consistent h-values have the following properties. Let $g(s)$ and $f(s)$ denote the g-values and f-values, respectively, after the A* search: First, A* searches expand all states at most once each. Second, the g-values of all expanded states and the goal state after the A* search are equal to the distances from start state to these states. Following the tree-pointers from these states to the start state identifies shortest paths from the start state to these states in reverse. Third, the f-values of the series of expanded states over time are monotonically nondecreasing. Thus, it holds that $f(s) \leq f(s_{\text{goal}}) = g(s_{\text{goal}})$ for all states s that were expanded by the A* search (that is, all states in the closed list) and $g(s_{\text{goal}}) = f(s_{\text{goal}}) \leq f(s)$ for all states s that were generated by the A* search but remained unexpanded (that is, all states in the open list). Fourth, an A* search with consistent h-values $h_1(s)$ expands no more states than an otherwise identical A* search with consistent h-values $h_2(s)$ for the same search problem (except possibly for some states whose f-values are identical to the f-value of the goal state) if $h_1(s) \geq h_2(s)$ for all states s.

Repeated Forward A* itself executes ComputePath() to perform an A* search. Afterwards, it follows the tree-pointers from the goal state to the start state to identify a shortest path from the start state to the goal state in reverse. Repeated Forward A* then makes the agent move along this path until it reaches the target or action costs on the path increase [Line 30]. In the first case, the agent has reached the target. In the second case, the current path might no longer be a shortest path from the current state of the agent to the state of the target. Repeated Forward A* then updates the current state of the agent and repeats the procedure.

Repeated Forward A* does not initialize all g-values up front but uses the variables *counter* and $search(s)$ to decide when to initialize them. The value of *counter* is x during the xth A* search, that is, the xth execution of ComputePath(). The value of $search(s)$ is x if state s was generated last by the xth A* search (or is the goal state). The g-value of the goal state is initialized at the beginning of an A* search [Line 22] since it is needed to test whether the A* search should terminate [Line 2]. The g-values of all other states are initialized directly before they might be inserted into the open list [Lines 7 and 20] provided that state s has not yet been generated by the current A* search ($search(s) < counter$). The only initialization that Repeated Forward A* performs up front is to initialize $search(s)$ to zero for all states s, which is typically automatically done when the memory is allocated [Lines 16-17].
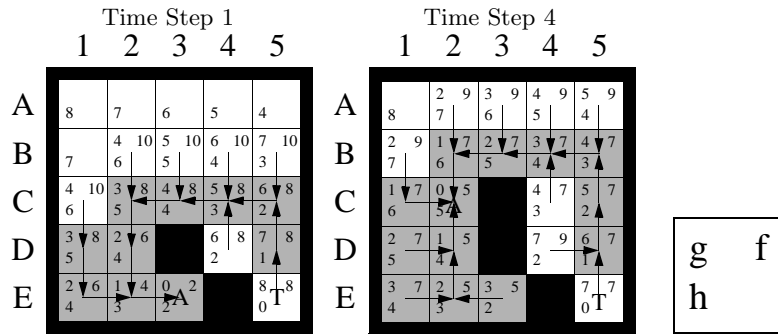
Figure 5: Repeated Forward A*

1.3 Implementation Details

Your version of Repeated A* should use a binary heap to implement the open list. The reason for using a binary heap is that it is often provided as part of standard libraries and, if not, that it is easy to implement. At the same time, it is also reasonably efficient in terms of processor cycles and memory usage. You will get extra credit if you implement the binary heap from scratch, that is, if your implementation does not use existing libraries to implement the binary heap or parts of it. We give you this extra credit to allow you to make connections to other classes and experience first hand how helpful the algorithms and data structure class that you once took can be. You can read up on binary heaps, for example, in Cormen, Leiserson and Rivest, Introduction to Algorithms, MIT Press, 2001.

Your version of Repeated A* should use the Manhattan distances as h-values. The Manhattan distance of a cell is the sum of the absolute difference of the x coordinates and the absolute difference of the y coordinates of the cell and the cell of the target. The reason for using the Manhattan distances is that they are consistent in gridworlds in which the agent can move only in the four main compass directions.

Your implementation of Repeated A* needs to be efficient in terms of processor cycles and memory usage since game companies place limitations on the resources that trajectory planning has available. Thus, it is important that you think carefully about your implementation rather than use the pseudocode from Figure 4 blindly since it is not optimized. (For example, the closed list in the pseudocode is shown only to allow us to refer to it later when explaining Adaptive A*.) Make sure that you never iterate over all cells except to initialize them once before the first A* search since your program might be used in large gridworlds. Do not determine which cells are in the closed list by iterating over all cells (represent the closed list explicitly instead, for example in form of a linked list). This is also the reason why the pseudocode of Repeated Forward A* does not initialize the g-values of all cells at the beginning of each A* search but initializes the g-value of a cell only when it is encountered by an A* search.

Do not use code written by others but test your implementations carefully. For example, make sure that the agent indeed always follows a shortest presumed unblocked path if one exists and that it reports that it cannot reach the target otherwise. Make sure that each A* search never expands a cell that it has already expanded (= is in the closed list).

We now discuss the example search problem from Figure 2 (left) to give you data that you can use to test your implementations. Figure 5 shows the first two searches of Repeated Forward A* for the example search problem from Figure 2 (left). Figure 5 (left) shows the first A* search,
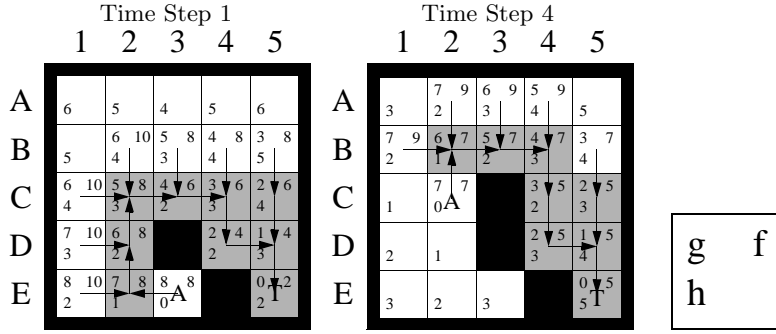
Figure 6: Repeated Backward A*

and Figure 5 (right) shows the second A* search. In our example problems, A* breaks ties among cells with the same f-value in favor of cells with larger g-values and remaining ties in an identical way. All cells have their user-supplied h-values in the lower left corner, namely the Manhattan distances. Generated cells (= cells that are or were in the open list) also have their g-values in the upper left corner and their f-values in the upper right corner. Expanded cells (= cells that are in the closed list) are shown in grey. (The cell of the target does not count as expanded since the A* search stops immediately before expanding it.) The arrows represent the tree-pointers, which are necessary to identify a shortest path after the A* search. Similarly, Figure 6 shows the first two searches of Repeated Backward A*, which searches from the target to the current cell of the agent.

1.4 Improving Repeated Forward A*

Adaptive A* uses A* searches to repeatedly find shortest paths in state spaces with possibly different start states but the same goal state where action costs can increase (but not decrease) by arbitrary amounts between A* searches.[1] It uses its experience with earlier searches in the sequence to speed up the current A* search and run faster than Repeated Forward A*. It first finds the shortest path from the current start state to the goal state according to the current action costs. It then updates the h-values of the states that were expanded by this search to make them larger and thus future A* searches more focused. Adaptive A* searches from the current state of the agent to the target since the h-values estimate the goal distances with respect to a given goal state. Thus, the goal state needs to remain unchanged, and the state of the target remains unchanged while the current state of the agent changes. Adaptive A* can handle action costs that increase over time.

To understand the principle behind Adaptive A*, assume that the action costs remain unchanged to make the description simple. Assume that the h-values are consistent. Let $g(s)$ and $f(s)$ denote the g-values and f-values, respectively, after an A* search from the current state of the agent to the target. Let s denote any state expanded by the A* search. Then, $g(s)$ is the distance from the start state to state s since state s was expanded by the A* search. Similarly, $g(s_{\text{goal}})$ is the distance from the start state to the goal state. Thus, it holds that $g(s_{\text{goal}}) = gd(s_{\text{start}})$, where $gd(s)$ is the goal distance of state s. Distances satisfy the triangle inequality:

---

[1]Start state refers to the start of the search, and goal state refers to the goal of the search. The start state of the A* searches of Repeated Forward A*, for example, is the current state of the agent. The start state of the A* searches of Repeated Backward A*, on the other hand, is the state of the target.

10

$$\begin{aligned} gd(s_{\text{start}}) &\leq g(s) + gd(s) \\ gd(s_{\text{start}}) - g(s) &\leq gd(s) \\ g(s_{\text{goal}}) - g(s) &\leq gd(s). \end{aligned}$$

Thus, $g(s_{\text{goal}}) - g(s)$ is an admissible estimate of the goal distance of state s that can be calculated quickly. It can thus be used as a new admissible h-value of state s (which was probably first noticed by Robert Holte). Adaptive A* therefore updates the h-values by assigning

$$h(s) := g(s_{\text{goal}}) - g(s)$$

for all states s expanded by the A* search. Let $h_{\text{new}}(s)$ denote the h-values after the updates.

The h-values $h_{\text{new}}(s)$ have several advantages. They are not only admissible but also consistent. The next A* search with the h-values $h_{\text{new}}(s)$ thus continues to find shortest paths without expanding states that have already been expanded by the current A* search. Furthermore, it holds that

$$\begin{aligned} f(s) &\leq gd(s_{\text{start}}) \\ g(s) + h(s) &\leq g(s_{\text{goal}}) \\ h(s) &\leq g(s_{\text{goal}}) - g(s) \\ h(s) &\leq h_{\text{new}}(s) \end{aligned}$$

since state s was expanded by the current A* search. Thus, the h-values $h_{\text{new}}(s)$ of all expanded states s are no smaller than the immediately preceeding h-values $h(s)$ and thus, by induction, also all previous h-values, including the user-supplied h-values. An A* search with consistent h-values $h_1(s)$ expands no more states than an otherwise identical A* search with consistent h-values $h_2(s)$ for the same search problem (except possibly for some states whose f-values are identical to the f-value of the goal state, a fact that we will ignore in the following) if $h_1(s) \geq h_2(s)$ for all states s. Consequently, the next A* search with the h-values $h_{\text{new}}(s)$ cannot expand more states than with any of the previous h-values, including the user-supplied h-values. It therefore cannot be slower (except possibly for the small amount of runtime needed by the bookkeeping and h-value update operations), but will often expand fewer states and thus be faster. You can read up on Adaptive A* in Koenig and Likhachev, Adaptive A* [Poster Abstract], Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 1311-1312, 2005.

Figure 7 shows the first two searches of Adaptive A* for the example search problem from Figure 2 (left). The number of cell expansions is smaller for Adaptive A* (20) than for Repeated Forward A* (23), demonstrating the advantage of Adaptive A* over Repeated Forward A*. Figure 7 (left) shows the first A* search of Adaptive A*. All cells have their h-values in the lower left corner. The goal distance of the current cell of the agent is eight. The lower right corners show the updated h-values after the h-values of all grey cells have been updated to eight minus their g-values, which makes it easy to compare them to the h-values before the h-value update in the lower left corners. Cells D2, E1, E2 and E3 have larger h-values than their Manhattan distances to the target, that is, the user-supplied h-values. Figure 7 (right) shows the second A* search of Adaptive A*, where Adaptive A* expands three cells (namely, cells E1, E2 and E3) fewer than Repeated Forward A*, as shown in Figure 5 (right).
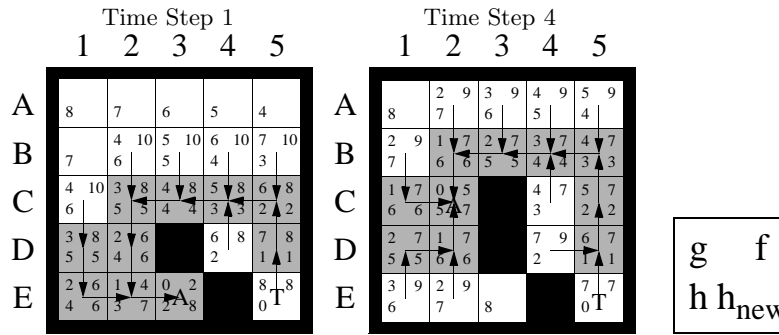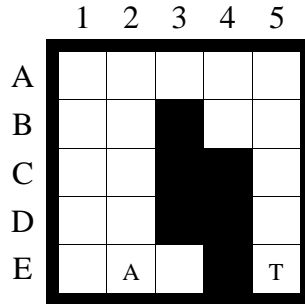
Figure 7: Adaptive A*

Figure 8: Second Example Search Problem

1.5 Experiments

Perform all experiments in the same 50 gridworlds of size 101×101. Generate their corridor structure with a depth-first search (with random tie breaking) and then make 100 randomly chosen blocked cells unblocked. See the program at idm-lab.org/gameai for details, which you are free to use. Note that sometimes more than one shortest presumed unblocked path exists. Different search algorithms might then move the agent along different shortest presumed unblocked paths (as seen in Figures 5 and 6), and the agent might observe different cells to be blocked. Then, the shortest presumed unblocked paths and the trajectories of the agent can start to diverge. This is okay since it is difficult to make all search algorithms find the same shortest presumed unblocked path in case there are several of them.

2 Questions

Answer the following questions.

2.1 Question 0: Very Easy

Read the chapter in your textbook on uninformed and informed (heuristic) search and then read the project description again. Make sure that you understand A* and the concepts of admissible and consistent h-values. Explain why the first move of the agent for the example search problem from Figure 8 is to the east rather than the north given that the agent does not know initially which cells are blocked.

Figure 9: Third Example Search Problem

2.2 Question 1: Easy

The text in the introduction of the project states that "The agent is guaranteed to reach the target if it is not separated from it by blocked cells." Give a convincing argument that the agent in finite gridworlds indeed either reaches the target or discovers that this is impossible in finite time. Prove that the number of moves of the agent until it reaches the target or discovers that this is impossible is bounded from above by the number of unblocked cells squared.

2.3 * Question 2: Easy

Repeated Forward A* needs to break ties to decide which cell to expand next if several cells have the same smallest f-value. It can either break ties in favor of cells with smaller g-values or in favor of cells with larger g-values. Compare both versions of Repeated Forward A* with respect to their runtime or, equivalently, number of expanded cells. Explain your observations in detail, that is, explain what you observed and give a reason for the observation. Hint: For the implementation part, priorities can be integers rather than pairs of integers. For example, you can use $c \times f(s) - g(s)$ as priorities to break ties in favor of cells with larger g-values, where c is a constant larger than the largest g-value of any generated cell. For the explanation part, consider which cells both versions of Repeated Forward A* expand for the example search problem from Figure 9.

2.4 * Question 3: Moderately Difficult

Compare Repeated Forward A* and Repeated Backward A* with respect to their runtime or, equivalently, number of expanded cells. Explain your observations in detail, that is, explain what you observed and give a reason for the observation. Both versions of Repeated A* should break ties among cells with the same f-value in favor of cells with larger g-values and remaining ties in an identical way, for example randomly.

2.5 * Question 4: Moderately Difficult

The text in Section "Implementation Details" states without a proof that "the Manhattan distances are consistent in gridworlds in which the agent can move only in the four main compass directions." Prove that this is indeed the case. Furthermore, the text in Section "Improving A*" assumes that the action costs remain unchanged and states without a proof that "The h-values

$h_{\text{new}}(s)$ ... are not only admissible but also consistent." Prove that Adaptive A* leaves initially consistent h-values consistent even if action costs can increase.

2.6   * Question 5: Easy

Compare Repeated Forward A* and Adaptive A* with respect to their runtime. Explain your observations in detail, that is, explain what you observed and give a reason for the observation. Both search algorithms should break ties among cells with the same f-value in favor of cells with larger g-values and remaining ties in an identical way, for example randomly.

2.7   Question 6: Very Difficult

Come up with ideas for how to generalize Adaptive A* to the case where the target moves, for example, because the agent is trying to catch a hostile agent or because the agent is trying to catch up with a friendly agent. Assume that the agent always knows the current cell of the target. Make sure that the agent always follows a shortest presumed unblocked path from its current cell to the target. Explain why it is necessary to change Adaptive A* to implement this movement strategy. Implement and test your generalization of Adaptive A*. Note that the agent needs to find another shortest presumed unblocked path from its current cell to the target whenever its current path is blocked or the target moves off the current path. The agent is guaranteed to reach the target if it is not separated from it by blocked cells and the target either moves more slowly than the agent (for example, does not move from time to time) or moves suboptimally from time to time.

2.8   Question 7: Very Difficult

In real-time computer games, blockages can disappear. Come up with ideas for how to generalize Adaptive A* to the case where the action costs can decrease as well as increase over time. Assume that the agent always knows all current action costs. Make sure that the agent always follows a shortest presumed unblocked path from its current cell to the target. Explain why it is necessary to change Adaptive A* to implement this movement strategy. Implement and test your generalization of Adaptive A*. (D* Lite mentioned in the next question handles these cases without changes.)

2.9   Question 8: Difficult

There are search algorithms other than Adaptive A* that also use their experience with previous similar searches to speed up the current A* search and run faster than Repeated Forward A*. Read up on D* Lite in Koenig and Likhachev, Fast Replanning for Navigation in Unknown Terrain, Transactions on Robotics, 21, (3), 354-363, 2005. You can find an implementation of D* Lite at idm-lab.org/gameai, which you are free to use. Compare Adaptive A* and D* Lite with respect to their runtime. Adaptive A* should break ties among cells with the same f-value in favor of cells with larger g-values. (D* Lite is actually the trajectory-planning algorithm of choice if the target does not move, both because it can handle the case where the action costs can decrease as well as increase and because it is faster than Adaptive A*. Unfortunately, it is too complex to result in a class project that helps you understand A* better. However, keep in

mind that you might want to use D* Lite instead of Adaptive A* when you implement actual real-time computer games.)

2.10 Question 9: Very Difficult

Weighted A* uses the f-values $f(s) = g(s) + c \times h(s)$, where $c \geq 1$ is a constant. Increasing c often decreases the runtime of Weighted A* but increases the lengths of the paths it finds. ARA* is a search algorithm that uses its experience with previous searches to speed up the current Weighted A* search and run faster than Repeated Forward Weighted A*. Read up on Weighted A* and ARA* in Likhachev, Gordon and Thrun, ARA*: Anytime A* with Provable Bounds on Sub-Optimality, Advances in Neural Information Processing Systems 16, 2004. Implement ARA* and then compare Adaptive A* and ARA* with respect to their trajectory length and runtime. Adaptive A* should break ties among cells with the same f-value in favor of cells with larger g-values.

2.11 Extra Credit 1: Moderately Difficult

You performed all experiments in gridworlds of size 101×101 but some real-time computer games use maps whose number of cells is up to two orders of magnitude larger than that. It is then especially important to limit the amount of information that is stored per cell. For example, the tree-pointers can be implemented with only two bits per cell. Suggest additional ways to reduce the memory consumption of your implementations further. Then, calculate the amount of memory that they need to operate on gridworlds of size 1001×1001 and the largest gridworld that they can operate on within a memory limit of 4 MBytes.

2.12 Extra Credit 2: Moderately Difficult

We used A* as baseline with its open list implemented as binary heap. However, there are alternatives.

First, one can use breadth-first search instead of A*. Compare Repeated Forward Breadth-First Search and Repeated Forward A* with respect to their runtime. Repeated Forward A* should break ties among cells with the same f-value in favor of cells with larger g-values. Do not implement breadth-first search as A* with zero h-values since it can run substantially faster than that. For example, it can use a first-in first-out queue instead of a priority queue to implement the open list and terminate once it is about to generate the goal cell rather than expand it. Describe how you implemented the breadth-first searches of Repeated Forward Breadth-First Search and then characterize situations when the runtime of Repeated Forward Breadth-First Search is smaller than that of Repeated Forward A* and vice versa (for example, by varying the sizes of the gridworlds).

Second, the open list of Repeated Forward A* can be implemented with data structures other than a binary heap, for example, with buckets (an array of sets of cells that is indexed with the f-values of the cells). Compare both versions of Repeated Forward A* with respect to their runtime or, equivalently, number of expanded cells. The version of Repeated Forward A* that uses a binary heap should break ties among cells with the same f-value in favor of cells with larger g-values. The version of Repeated Forward A* that uses buckets can break ties in a similar (but not completely identical) way by implementing the sets as last-in first-out queues.

2.13 Extra Credit 3: Difficult

Performance differences between two search algorithms can be systematic in nature or only due to sampling noise (= bias exhibited by the selected test cases since the number of test cases is always limited). One can use statistical hypothesis tests to determine whether they are systematic in nature. Read up on statistical hypothesis tests (for example, in Cohen, Empirical Methods for Artificial Intelligence, MIT Press, 1995) and then describe for one of the experimental questions above exactly how a statistical hypothesis test could be performed and what its outcome is for your experimental data.