

Incorporating Human Relationships Into Path Planning

Nathan R. Sturtevant

Department of Computer Science
University of Denver
Denver, CO, USA
sturtevant@cs.du.edu

Abstract

While pathfinding is ubiquitous in games, it is most often abstracted away from everything else that is going on in the game and viewed as a simple point-to-point shortest path problem. The notable exception to this is work on tactical pathfinding. Tactical pathfinding incorporates tactical elements, most often for battle simulations, into the pathfinding engine in order to improve high-level behavior. This paper looks at how the relationship between characters in a game should influence the paths that are taken through the game. We suggest a simple model of character relationships and show how this can be used to influence the paths that are taken through the world. Finally, we discuss design-related issues of how this could be incorporated into games.

Introduction and Overview

The problem of finding point-to-point paths in a world is a common task in many games, and there are a wide variety of planning techniques and pathfinding representations (Board and Ducker 2002; Nieuwenhuisen et al. 2004; Demyen and Buro 2006; Kallmann 2010; Sturtevant 2011) which have been developed for solving this problem within the constraints of modern video games. These approaches form a technical solution to the problem of finding paths, and as such, some might argue that this isn't really Artificial Intelligence (AI), or that it isn't an particularly interesting aspect of AI in video games.

We argue, however, that if agents cannot move around the world in a competent manner, then they will not appear intelligent. It doesn't matter what techniques are applied to give this illusion of intelligence, only that the illusion remains relatively intact throughout a playable experience. Because the reasoning process behind pathfinding can be expensive, research into representations and other technical problem details is an important part of maintaining this competency.

But, the sentiment behind the critique also has some merit. The best path between two points is not necessarily the shortest one as measured by path length. The best path should take into account everything that is going on within the world being simulated. Work in the area of tactical pathfinding (Millington and Funge 2009) is the primary research that has considered richer world and cost rep-

resentations, although a majority of this work has been in the context of first-person shooters (Straatman, van der Sterren, and Beij 2005; Kondeti et al. 2005) or military simulations (Darken, McCue, and Guerrero 2010). Example applications of this work in games has been to help agents avoid moving through corridors of fire, or to coordinate movement for tactical purposes.

In this paper we suggest that there are many situations where a pathfinding engine should be taking into account the relationships between characters when planning paths. That is, the paths planned by characters in games should take into account their relationships and emotions, reflecting the state of mind of the agents. Properly integrating this into the pathfinding process has the potential to enhance the ability for game designers to create compelling characters that engage in unique interactions and performances within a game.

We propose a simple model of character intent when moving and look at how this model can be incorporated into a pathfinding engine, and then discuss the practical application of this work. Note that this work focuses primarily on the path being planned for movement. The selection of body and facial animations when following the path are a separate area of work beyond the scope of this paper.

Background and Related Work

There has been a significant amount of research on human relationships, but very little of this work is focused directly on the modeling needed for this paper. We point out several particular fields of research and some related work in crowd simulation, but there are a large number of publications outside of AI which have some level of relevance to this work.

Studies of Human Space and Behavior

The field of proxemics looks at how humans move and position themselves in relationship to each other. This is influenced by the emotional state of an individual as well as cultural norms (Hall and Hall 1969). The sense of personal space depends on who is being interacted with and the cultural norms of those who are interacting. This work divides space around an individual into four categories, which we describe below. These were further sub-divided into close and far phases, although we will not make that distinction here.




Figure 1: The relationship distances, as defined by Edward Hall.

- **Intimate distance.** This is described as “the distance of love-making and wrestling, comforting and protecting.” In other words, this is the distance of the most tender and most harsh words between individuals.
- **Personal distance.** This is the distance that most non-contact species maintain; it is close enough to touch, but not close enough that vision is distorted.
- **Social distance.** This distance is close enough to allow for social interactions, with most physical details still clear, but far enough that in a louder environment voices might need to be raised to communicate clearly.
- **Public distance.** This is the limit at which humans can clearly communicate, but vocal nuances begin to be lost, and non-verbal communication becomes more dominant.

Visually, as shown in Figure 1, these distance form concentric circles or ellipses around an individual, from the furthest distances where an individual is safest, to the intimate distances where an individual is most vulnerable. We will use this formation as one axis of a simplified model for planning preferences.

Related to work on proxemics is work on describing human movement. This originated from work by Rudolf Laban (Laban and Ullmann 1966), and has been widely applied to dance and other artistic endeavors. This work provides insights not only to high-level motion planning, but also to the individual animations used for characters. Although we will not use this work directly here, it has been used previously for analyzing motion capture data (Boucharde and Badler 2007).

Motion Planning

There are several areas in which high-level concepts or relationships have been used to influence character movement; we highlight two here.

In the area of motion planning in AI, Guy et. al. (2011) looked at how personality models influence character behavior as parts of crowds. This work is notable in that it has significant validation studies from humans observing the behaviors of simulated agents. An analysis of the behaviors resulted in a simple two-dimensional model based on “extraversion” and “carefulness”, with a spectrum of behaviors in between. These behaviors are applied generically across

a crowd, as opposed to directly for goal-directed movement, which is the focus of this paper.

Tactical pathfinding The term ‘tactical pathfinding’ has been used in a variety of contexts, but is not always well-defined. A popular game AI textbook (Millington and Funge 2009) informally defines tactical pathfinding as being the same as regular pathfinding except that the cost function for search must be modified to take account of tactical information. Other work (Kamphuis, Rook, and Overmars 2005) gives the definition of tactics is as “the employment and arrangement of forces in battle.” While our work is not directly concerned with battle, it is concerned with similar influences on path planning behavior.

Detailed information has been provided on how the game Killzone has implemented tactics (Straatman, van der Sterren, and Beij 2005). At a high level, these approaches simply weight the graph in areas that the character should not travel, causing the basic A* planner to minimize visits to these areas.

Essentially, the goal of tactical pathfinding is to build a richer cost model of movement based on environmental factors which can positively or negatively influence the character that is moving through the world. We prefer this broader definition as it is not directly linked to battle or military usage, but can be more broadly applied to any cost model or environment. In this view, this research enhances tactical pathfinding techniques to provide a basis for using human relationships during planning.

A Model of Movement Planning Preferences

Consider a character planning through an environment that is physically safe (e.g. there are no military threats), but in which the relationships are deep and rich. How should the character move through this environment? What is a space of interesting behaviors that characters might exhibit when faced with this challenge? These are questions we hope to address with a simple relational model of movement.

We will define the actual search problem formally later in the paper, but for now we distinguish between the *character* who is planning and *target* individuals which are either along a route being planned, or at the destination of the route. Note that these are not fixed roles; a character could easily be a target during a different planning process.

While there are many animations, expressions, and variations in gait and posture which are important for building a character, these do not directly influence the planning process or the path that is taken. We suggest that two main characteristics should influence the planning process. The first characteristic has to do with how the targets will perceive the character when the character moves around the world. While there may be different motivations behind it, ultimately a character either wants to be seen by targets, or does not want the targets to see it. This forms the first axis of our model. The second characteristic has to do with how the character will arrive at the destination, assuming that the destination is a target. We borrow Hall’s model of proxemics for this measure, the second axis in our model.

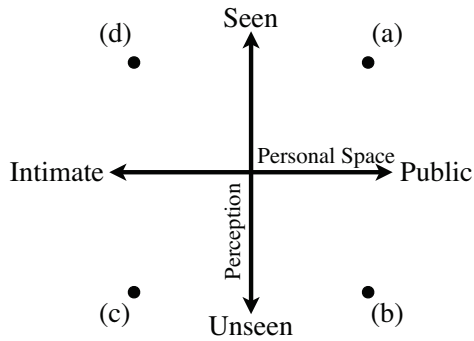


Figure 2: A two-dimensional model of how characters will seek or avoid other characters when planning paths.

The model is shown in Figure 2, with the horizontal axis for personal space and the vertical axis for path perception. This scale is measured relative to a single target character in the environment, whether a shopkeeper, a crime boss, or a friend of the opposite sex. We now look at these axes and points in the model space in more detail.

Assuming that we have a model of what the targets can see and perceive in the environment, the character can either avoid or seek these areas. A character with a large gambling debt, for instance, would probably want to avoid his bookie. A character who just went through a break-up might want to avoid her ex-boyfriend. Another character might choose a path that passes by the target towards which he feels affection.

The second dimension of the classifier is where the interactions with the target will occur on the spectrum between public and intimate space. A lover would approach the object of her affection and arrive in the intimate space around the target, while a heckler might arrive in the public space, close enough to heckle, but far enough to stay safe from any reaction that the heckling induces. A spy, which has collected information for a target, but wishes to avoid direct association with the target, would also interact with the target in the public space, where body language or other signals could be sufficient communication to help set up a future meeting.

In this classification scheme, we attempt to avoid the use of social norms, such as terms like ‘polite’, as this interpretation of behaviors can differ between cultures. Instead, we seek to rely on primitives primarily based on difference, which is the final manifestation of the cultural norms driving the behavior.

For illustrative purposes, we describe the behavior of four outlying points on the scale, labeled (a), (b), (c) and (d) in Figure 2. Note that default behaviors of most characters would fall near the origin of this scale. These points represent the extremes in behavior, however these are also probably the most interesting characters to have in a game. At these extremes, integrating this model into path planning has the potential to increase the realism of the characters involved. Sample behavior at each of the points follows:

Point (a) describes an approach into the public space of the target where the character has the intention of being seen. A character modeled by point (a) might push through a crowd while calling out threats to the target, without fully approaching the character. This behavior might be seen as showing off to the crowd that the character in question has power over the target.

Point (b) describes an approach in the public space where the character wishes to remain unseen. This approach could describe an assassin which wishes to assassinate a target from a distance while remaining unknown, or a spy that is observing the behavior of the target. Note that we could distinguish who the targets are to get very different behavior here. Point (b) might describe the behavior around police in a crowd when the character is trying to meet up with another agent, but the behavior towards the agent/target might be described by point (c).

Point (c) describes an intimate and unseen approach. In this example, consider a target that is loosely surrounded by body guards. This target has done something to upset the character, but the character is unable to approach the target directly without engaging the body guards first. Thus, the character must approach the target in a manner that helps it remain unseen. But, once a confrontation begins, it is direct and occurs in the intimate space. (Perhaps the character grabs the target from behind.)

Finally, point (d) describes a seen approach into intimate space. This could either be two lovers embracing in a public space, or two rivals directly approaching and arguing. An example from Hall is a baseball manager yelling in the face of an umpire.

Technical Implementation

There are two high-level details that need to be part of a pathfinding engine in order to support the ideas described here. The first is the support for dynamically weighted regions. The second is the support for maintaining information about heading during planning. The argument for dynamically weighted regions is relatively simple – if there are locations we want to avoid, we need to increase the cost of traveling through those regions, or, in the extreme case, make them completely inaccessible. The argument for maintaining heading information is more subtle, but it supports the perception axis and the goal of being seen or unseen. In particular, a character’s face is usually its most identifying feature. Thus, if a character wants to be seen by a target, the character will ideally face the target so that the character and the target can see each other. One possibility that we do not use here is that a “cry out” action is added to the search space, which would encourage targets within a certain range to turn and face the character.

Formally speaking, we can now define a search task for a character as a directed graph $G = \{V, E\}$, a start state $s \in V$ and a goal state $g \in V$. Additionally, we define n targets $T = t_1, t_2 \dots, t_n$ that should influence any path found between s and g . Each edge $e \in E$ has some cost of travel. This cost can be decomposed into several quantities: (1) The underlying distance traveled along the edge and (2) the cost associated with traversing the edge relative to each

of the n targets. We divide this into two costs associated with the axes of our model, a perception cost, c_p , and a personal space weight, c_{ps} .

Let P be the perception weight, ranging from $-1 \dots 1$, where 1 means the character wants the target to perceive the character. Let Δ be the relative angle from where the character to the target. Let d be the distance from the character to the target. Let k_1 and k_2 be constants. Then, the perception cost of traversing an edge is $c_p = k_1 \cdot \max(-P * \cos(\Delta \cdot k_2), 0) / d$. If the character is behind the target, then the cosine term will be -1.0, and c_p will be positive, meaning that there is additional cost to traversing that edge. If the character is directly in front of the target, the cosine term will be 1.0, and c_p will be 0, meaning that there is no additional cost. It is possible to use a minimum cost of less than 0, but negative or even very low edge weights can cause an agent to take strange paths. k_1 influences how strongly the movement constraints should be followed. k_2 widens or narrows the effective field of view of the target.

Let CP be the personal space weight. We treat CP as the desired distance when passing by targets, although it could also be treated as a preferred minimum distance. Then, $c_{ps} = \min(\text{fabs}(CP - d) / k_3, k_4)$, where k_3 and k_4 influence the tolerance and cost of being outside the preferred personal space.

Putting everything together, given an edge $e = \{s_1, s_2\}$, the total cost of traversing the edge is $c(e) = c_{ps} + c_p + d(s_1, s_2)$, where $d(x, y)$ is the distance between x and y .

While we do not consider it here, it is not difficult to incorporate the length of visibility into this model (van der Sterren 2002). This is done by adding a variable to the search space tracking how long a character has been visible. As such, it should be clear that there are many variants on this model that might be more suitable for a particular domain. This model was chosen because it works, as is demonstrated in the next section.

Example Scenarios

The primary results of using the techniques described here are visual, so we chose a few parameters in our model and demonstrate the paths that are followed by the character. Our experiments are performed on grid worlds. Characters in the grid world can either move forward or change their heading, a more realistic model of character movement than simply using the grid as an 8-connected world.

While navigation meshes (Board and Ducker 2002) have become quite popular in games, adding weighted regions to such maps can result in an expensive pathfinding process (Rowe and Alexander 2000). Navigation meshes are an efficient representation of free space. As such, we recommend that a discretization, such as a grid world, be used for at least some stages of the planning, particularly in the final stages when approaching a target.

In our first experiment, the character will approach the personal space of a target in the presence of a second target. We will then vary whether the character is avoiding or seeking the second target. In this experiment $k_1 = 10$, $k_2 = 1.0$, $k_3 = 2$ and $k_4 = 1.0$. P for the first target is 1.0; we will vary P for the second target. CP is set to 5.0.





Figure 3: Varying whether we want to be seen or not seen by the target in the middle.




Figure 4: Varying the preferred personal space to 2.0.

In Figure 3 we show the difference between when the second target has $P = 0.2$ and $P = -0.2$. The start state is in the bottom right corner. The goal state is in the upper left corner. The character is approaching the first target next to the goal, and must either avoid or seek the attention of the second target in the middle, which is facing upwards. When $P = -0.2$ (top) the second target is avoided while moving towards the goal. When $P = 0.2$ (bottom), the chosen path passes in front of the second target.

It is important to note that the path taken is exactly distance 5 from the second target, whether above or below, as influence by the CP parameter. When we set $CP = 2$ and keep $P = -0.2$, we get the behavior shown in Figure 4.

Our second experiment, shown in Figure 5, uses the same parameters as the first, except that this time there is only one target. We vary the desired perception from $P = -1.0$ (top) to $P = -0.5$ (middle) and $P = 1.0$ (bottom). When the character strongly desires to avoid being seen by the target, it approaches the target from below, and then moves in front of the target at the last possible step. When the character has only a partial desire to avoid the target, it chooses a path which does a better job of avoiding the target (middle), but

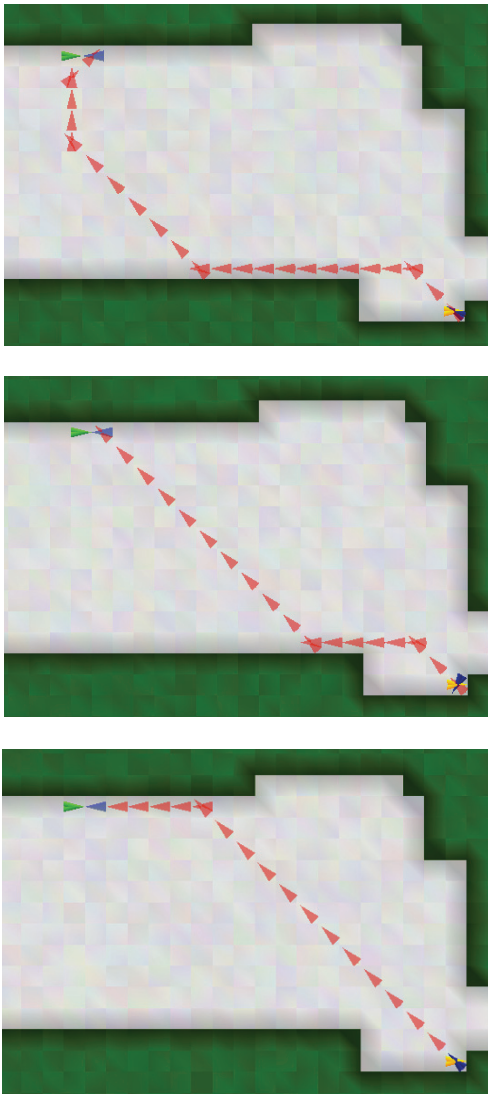


Figure 5: Varying whether we want to be seen or not seen by the target being approached.

still stays generally in front of the target. Changing the value of P does not smoothly change the behavior. Instead, the behavior jumps as changes in the weights and parameters increase the cost of travel.

A last experiment, shown in Figure 6, illustrates how behavior can change with small changes in the problem being solved. The character starts above the target, and the target is facing to the right. On the left hand side of the figure, the character takes the shortest path (distance-wise) to the target. But, when the initial distance is increased by one, an alternate path is found which is longer, but does a better job avoiding being seen by the target. (Note that the obstacles here do not influence the visibility of the agent, only the possible paths which can be taken.)

These examples demonstrate how the parameters of the model influence the behavior of the agents. We note two

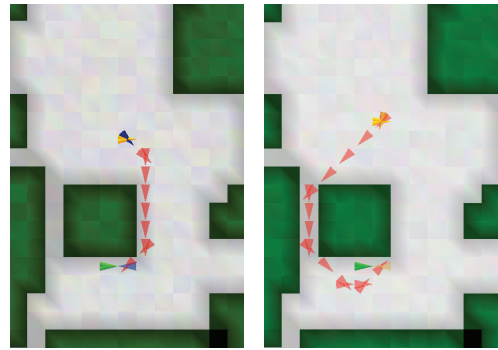


Figure 6: The change in avoidance behavior as the distance to the target changes.

things about the performance. First, there is some tuning required to find suitable parameters for achieving the desired motion. Similar, there is also tuning in the cost functions that can be used. While we chose cost functions that generated reasonable behavior, more complicated cost functions are also possible using the same parameters. Again, this is a matter of tuning for the performance desired.

Second, and more importantly, the cost of search using the model suggested here is more expensive than planning without the model. The exact cost depends on the weights used. Larger weights essentially make the default heuristic less accurate, and thus require more search to find a solution. On relatively small problems, this isn't an issue, but long-term planning will be prohibitively expensive without combining the approach with some form of abstraction (Board and Ducker 2002; Sturtevant 2011). We deliberately do not measure the cost planning in our results as our focus is on the behavior produced instead of the technical means to minimize the planning cost.

Limitations

While we are advocating for incorporating relationship models into pathfinding, we do see a number of limitations to the approach. In general, these limitations just suggest that the ideas should be applied in a limited and focused way, which is ideal, as they increase the cost of planning in the world.

When playing a game, one of the primary human limits is attention. Humans cannot process and understand everything that is going on in the world. If every character in a game is moving around the world in accordance to the relationships established by the game, most of the movement choices will be completely lost on the human player, or worse, will confuse the player into thinking that the characters are unintelligent. Thus, we suggest the following principles for the application of this work:

- **Limited application:** These techniques should be used at points when the user will notice and appreciate them most.
- **Integration with design:** The techniques will be most useful when applied jointly with the design of levels and user experiences. There are a limited number of ways of

planning to approach a target at the end of an alley; these sorts of areas would not benefit from this work.

- **Train the user:** A game using these ideas should train the user to take advantage of the capabilities of the pathfinding system. If the user is unaware that the pathfinding engine has the capacity to plan in a unique way, the capability may be wasted.
- **Integration with animation and sound:** The use of animations can be used to convey the state of mind of the character for which planning is being performed. These might be key for allowing the user to understand why the character is taking the given path. If the planning is happening for your character, they might tell you exactly why they are taking the given route.
- **Appropriate use:** We do not argue that every aspect of this work should be used in every game. Sometimes it is the conflicts that this work might avoid that make the game more interesting. But, social norms can also be exploited. If a user opts to avoid using the pathfinding system and approaches a target in an awkward or impolite manner, the target might comment on this, or in other ways react appropriately.

Open Research Issues

There has been a significant amount of recent work on methods for improving the speed of planning on grid worlds (Harabor and Grastien 2012; Botea 2012; Antsfeld et al. 2012), these techniques do not work well in dynamic and weighted worlds. This suggests that further work must be done to measure and evaluate pathfinding techniques not only in the scenarios suggested in this paper, but for tactical pathfinding approaches in general. The development of fast and robust methods which are tolerant of significant changes to the map is a significant open technical challenge in the area of pathfinding for video games. But, such techniques could encourage further work and applications of tactical pathfinding.

While we have presented a model in this paper that accounts for relationships between characters in a game and how these relationships should influence the paths taken, the work has not been deployed. In his invited talk at AIIDE 2010, Chris Hecker made the connection between game AI techniques and game design. This paper has presented new techniques, but it has not incorporated them into a compelling game design. This is a non-trivial task, but we are working on how even a small playable experience could be created which makes compelling use of this work. Perhaps the more difficult work, however, is pairing suitable animations with the paths that are created, as in the end, the player of a game primarily perceives the animations used, not the reasons why a particular path was found.

Conclusions

This paper presents a two-dimensional model of personal relationships that can be used to influence the paths planned by a character in a game. The model is built on the idea of personal space and perception of the character by a set of

targets in the world. A number of small scenarios demonstrate how the parameters influence the movement of characters through the world, and provide an initial concept of how this work could be applied in practice.

References

- Antsfeld, L.; Harabor, D. D.; Kilby, P.; and Walsh, T. 2012. Transit routing on video game maps. In Riedl, M., and Sukthankar, G., eds., *AIIDE*. The AAAI Press.
- Board, B., and Ducker, M. 2002. Area navigation: Expanding the path-finding paradigm. *Game programming gems* 3:240–255.
- Borrajo, D.; Felner, A.; Korf, R. E.; Likhachev, M.; López, C. L.; Ruml, W.; and Sturtevant, N. R., eds. 2012. *Proceedings of the Fifth Annual Symposium on Combinatorial Search, SOCS 2012, Niagara Falls, Ontario, Canada, July 19-21, 2012*. AAAI Press.
- Botea, A. 2012. Fast, optimal pathfinding with compressed path databases. In Borrajo et al. (2012).
- Bouchard, D., and Badler, N. I. 2007. Semantic segmentation of motion capture using laban movement analysis. In Pelachaud, C.; Martin, J.-C.; André, E.; Chollet, G.; Karpouzis, K.; and Pelé, D., eds., *IVA*, volume 4722 of *Lecture Notes in Computer Science*, 37–44. Springer.
- Darken, C. J.; McCue, D.; and Guerrero, M. 2010. Realistic fireteam movement in urban environments. In Youngblood, G. M., and Bulitko, V., eds., *AIIDE*. The AAAI Press.
- Demyen, D., and Buro, M. 2006. Efficient triangulation-based pathfinding. In *AAAI*, 942–947. AAAI Press.
- Guy, S. J.; Kim, S.; Lin, M. C.; and Manocha, D. 2011. Simulating heterogeneous crowd behaviors using personality trait theory. In Bargteil, A. W., and van de Panne, M., eds., *Symposium on Computer Animation*, 43–52. Eurographics Association.
- Hall, E. T., and Hall, E. T. 1969. *The hidden dimension*. Anchor Books New York.
- Harabor, D. D., and Grastien, A. 2012. The jps pathfinding system. In Borrajo et al. (2012).
- Kallmann, M. 2010. Navigation queries from triangular meshes. In Boulic, R.; Chrysanthou, Y.; and Komura, T., eds., *MIG*, volume 6459 of *Lecture Notes in Computer Science*, 230–241. Springer.
- Kamphuis, A.; Rook, M.; and Overmars, M. H. 2005. Tactical path finding in urban environments. In *First International Workshop on Crowd Simulation*.
- Kondeti, B.; Nallacharu, M.; Youngblood, M.; and Holder, L. 2005. Interfacing the dargagnan cognitive architecture to the urban terror first-person shooter game. In *IJCAI-05 Workshop on Reasoning, Representation, and Learning in Computer Games*, 55–60.
- Laban, R. v., and Ullmann, L. 1966. *Choreutics*. London: Macdonald & Evans.
- Millington, I., and Funge, J. 2009. *Artificial Intelligence for Games, Second Edition*. Morgan Kaufmann.

Nieuwenhuisen, D.; Kamphuis, A.; Mooijekind, M.; and Overmars, M. H. 2004. Automatic construction of roadmaps for path planning in games. In *International Conference on Computer Games: Artificial Intelligence, Design and Education*, 285–292.

Rowe, N. C., and Alexander, R. S. 2000. Finding optimal-path maps for path planning across weighted regions. *The International Journal of Robotics Research* 19(2):83–95.

Straatman, R.; van der Sterren, W.; and Beij, A. 2005. Killzones ai: dynamic procedural combat tactics. In *Game Developers Conference*. Citeseer.

Sturtevant, N. 2011. A sparse grid representation for dynamic three-dimensional worlds. In *Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 73–78.

van der Sterren, W. 2002. Tactical path-finding with A*. *Game Programming Gems* 3:294–306.