AGENT NAVIGATION USING POTENTIAL FIELDS AND FORWARD CHAINING

Graeme Bell & Michael Weir {gbb@dcs.st-and.ac.uk, mkw@dcs.st-and.ac.uk}.
School of Computer Science, University of St. Andrews, Scotland. KY16 9SS.

Key words to describe the work: Navigation, Potential Fields, Robotics.

Key Results: Increasingly competent heuristics produce the novel Forward Chaining heuristic, allowing navigation past obstacle configurations that are impossible for traditional potential field methods.

How does the work advance the state-of-the-art?: Gradient based potential field navigation heuristics cannot overcome the *Local Minimum Problem*. Forward Chaining overcomes the LMP and allows successful navigation in most 2D environments.

Motivation (problems addressed): To reopen research into potential field based navigation, by making potential field based navigation feasible for a much wider range of environments than was possible with traditional gradient based approaches.

1. Introduction and Problem Description

Many physical and virtual agents are faced with the task of navigation. One problem that might be faced by a real-world robot or computer game agent is the task of navigating from some starting position **S** to a goal position **G,** in a two-dimensional environment containing obstacles which may not be travelled through. This is represented in Figure 1.

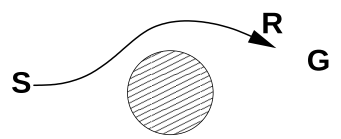

Fig. 1 An environment with start and goal positions and a single obstacle. An example navigation path for the agent R is shown.

Potential fields can be used to address this problem by using the intuitive metaphor that obstacles should 'push' the agent away and the goal should 'pull' the agent towards it **[1,2]**. This metaphor is typically realised by having the gradient at a position on a potential field direct the agent according to the agent's position relative to the goal and relative to obstacles. The goal is modelled using an attracting potential function whose single minimum of zero height lies at the goal's position, and an obstacle is modelled as a function with extremely high potential at positions inside the obstacle's boundary. Outside of the obstacle's boundary, the potential gradually falls off to zero as the distance from the obstacle increases. The goal and obstacle functions are then added together, and gradient descent is used to direct travel downwards on the potential field using the derivative of this combined goal-obstacle potential function. In very simple environments, this causes successful navigation to the global minimum of the field and thus the goal. Unfortunately, potential fields associated with more complex environments often possess minima that do not correspond to the goal's position. An example of a 'trap' obstacle configuration that produces such a potential field is given in Figure 2. In these cases, when gradient descent directs the agent into the local minimum basin, the agent becomes trapped and unable to make further progress as the potential surrounding the agent is higher in all directions locally. Further progress cannot be made and navigation is unsuccessful. This is known as the *local minimum problem* **[2,3]**.

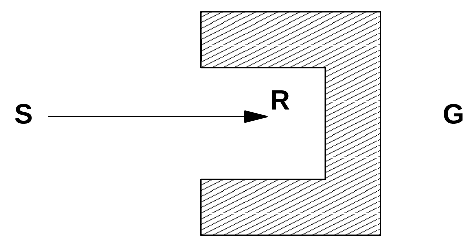

Fig. 2 A concave obstacle configuration that generates a local minimum at R, preventing navigation from S to G.

If a potential field based technique was able to escape from minimum basins such as that created by the concave obstacle configuration shown in Figure 2, then successful navigation would be possible. This would be important, as the local minimum problem is a significant and generally intractable problem for the entire field of heuristic search.

2. Forward Chaining

Several observations can be made, leading to a novel approach that surmounts the problematic obstacle configurations. Firstly, a key problem is that gradient descent is taking place on a fixed potential surface. By allowing the potential surface to vary, we could provide an opportunity to let 'downhill' become a different direction that leads gradient descent out of the trap. To provide this flexibility, we allow a goal to be temporarily replaced by a more local *subgoal* **[4,5]**. This is similar to humans choosing to go towards a place that is not the final destination of their route - knowing that they will be able to steer better to their destination by breaking a journey into small steps. By *chaining* the subgoals into a sequence that connects to the goal, the agent can be led gradually to its goal location.

Secondly, we recognise that failure has been induced by travel being brought to a standstill. We keep the agent on the move through selecting the next subgoal target state to be the point of lowest potential value (according to the goal and obstacles) on a circle whose centre is the current state. This transforms the *Local Minimum Problem* into a *Local Oscillation Problem* of various step lengths.

Thirdly, we make the most important move by eliminating back and forth motion, i.e. the Local Oscillation Problem of step length 1, through restricting subgoal selection to the *forward* semi-circle relative to the last direction in which movement took place. For shallow C-shaped boundaries, this enables the agent to move forward along the obstacle boundary even when the gradient points away from the goal. Further (but more minor) modifications to subgoal placement are made to overcome other shapes of concave curvature and increasing difficulty. The heuristics are straightforward to implement and are not computationally expensive. A sample path using the final FWDS4 subgoal placement heuristic is shown in Figure 3, together with the oscillatory paths of less evolved ancestors.

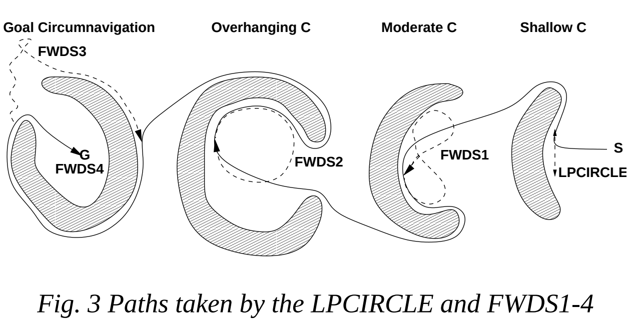

Fig. 3 Paths taken by the LPCIRCLE and FWDS1-4 heuristics when faced with the obstacle course shown. FWDS4 succeeds in navigating in the presence of all of the types of obstacle curvature represented above and generates an efficient path from the start S to the goal G.

4. Summary of Experimental Results

Experiments were carried out to evaluate the success of the heuristic on a range of shapes, under a number of different initial conditions. The results confirmed that the heuristic works as expected, and makes the potential field approach feasible and successful for 2D map navigation on non-trivial maps with obstacles that contain concave curvature.

5. Further Work

The concepts used by Forward Chaining appear to be generalisable to higher dimensions. Presently, Forward Chaining is being extended for use in 3D environments. It is hoped that it will be possible to extend the heuristic to allow navigation through n-dimensional configuration spaces with unrealisable regions.

6. Conclusion

The Forward Chaining technique for obstacle navigation has been presented, along with an outline of the development of the heuristics that produce Forward Chaining. The problems presented to potential field based navigation by concave obstacles have been dealt with by introducing a novel travel heuristic that allows success by turning the intractable local minimum problem into a tractable oscillation problem.

Our experimental results indicate that Forward Chaining is a significant improvement over traditional gradient-based approaches.

7. References

[1] Khatib, O., (1986), 'Real-Time Obstacle Avoidance for Manipulators and Mobile Robots', *The International Journal of Robotics Research,* 5(1).

[2] Dudek, G. (2000), *Computational Principles of Mobile Robotics*, Cambridge University Press.

[3] *Russell, S. & Norvig, P. (2002), Artificial Intelligence: A Modern Approach*, Prentice Hall.

[4] Weir, M. & Fernandes, A., (1994), Tangent Hyperplanes and subgoals as a means of controlling direction in goal finding, *in* 'Proceedings of the World Conference on Neural Networks', Volume 3, pp 438-443.

[5] Lewis, J. & Weir, M. (2000), Using Subgoal Chaining to Address the Local Minimum Problem, *in* 'Proceedings of the International ICSC Symposium on Neural Computation', 2000.