ECAI 2016
G.A. Kaminka et al. (Eds.)

doi:10.3233/978-1-61499-672-9-1762

Hole in One: Using Qualitative Reasoning for Solving Hard Physical Puzzle Problems

Xiaoyu Ge[1] and **Jae Hee Lee**[2] and **Jochen Renz**[1] and **Peng Zhang**[1]

Abstract. The capability of determining the right sequence of physical actions to achieve a given task is essential for AI that interacts with the physical world. The great difficulty in developing this capability has two main causes: (1) the world is continuous and therefore the action space is infinite, (2) due to noisy perception, we do not know the exact physical properties of our environment and therefore cannot precisely simulate the consequences of a physical action.

In this paper we define a realistic physical action selection problem that has many features common to these kind of problems, the minigolf hole-in-one problem: given a two-dimensional minigolf-like obstacle course, a ball and a hole, determine a single shot that hits the ball into the hole. We assume gravity as well as noisy perception of the environment. We present a method that solves this problem similar to how humans are approaching these problems, by using qualitative reasoning and mental simulation, combined with sampling of actions in the real environment and adjusting the internal knowledge based on observing the actual outcome of sampled actions. We evaluate our method using difficult minigolf levels that require the ball to bounce at several objects in order to hit the hole and compare with existing methods.

1 INTRODUCTION

One of the grand visions of Artificial Intelligence is to build robots with similar everyday capabilities as humans, who can live among us and assist us with many of our daily tasks. To progress towards more capable and more human-like robots, we need to develop methods and technology that allow robots to successfully interact with their environment. When we humans are faced with a "*physical action selection problem*", i.e., a problem that requires selecting a physical action that achieves the desired goal, we are very good at coming up with a qualitative solution and with a qualitative prediction of the consequences of an action. We have selected one particular physical action selection problem that is an actual real-world problem and that covers many common aspects of physical action selection problems. We call our problem the "*Hole-in-One*" problem in reference to the problem in mini golf of identifying and executing a shot that sinks the ball with this single shot. Variants of the hole-in-one problem occur frequently, not just in mini golf, in Pachinko, in pool billiard, curling or in a multitude of physics-based video games such as Angry Birds, but also in many everyday situations.

What these problems have in common is that the selected action can be one of infinitely many possible force vectors. Once a force vector is given and the physical properties of all objects and the environment are exactly known, it is possible to compute the exact trajectory of the ball and to see if that force vector solves the problem. However, we have to identify a force vector out of infinitely many possibilities that solves the problem. While a geometrical or analytical solution of these problems is typically not possible if the obstacle course is non-trivial, humans are very successful in solving these kind of problems. These problems become even harder to solve when we do not know the exact physical setting. We often only know what we can see and our perception is thus the limiting factor in what we know about the physical setting. Because of the uncertainty about the physical environment, potential solutions to the problem need to be executed in the actual environment before we can be sure that it is a solution. If it is no solution, we need to find ways of adjusting it so that it will eventually lead to a solution.

There are two key research streams in reasoning about physical systems, namely qualitative physical reasoning [2] and simulation-based reasoning [1]. The main advantage of qualitative reasoning is that it can rapidly draw useful conclusions from incomplete information [3]. [5] proposed a framework for reasoning about the motion of a 2D ball by qualitative simulation. While most of the previous work focuses on representing physical systems and describing (or predicting) consequences of actions, our method is solving a considerably harder problem as it has to find applicable actions from the infinite action space. In robotics, there has been extensive research on motion planning [8] or manipulation planning [4]. [9] developed a framework aiming at combining learning and planing and employing qualitative reasoning and linear temporal logic. There has been some work [7] on teaching robots to play mini golf.

In this paper we propose a solution to this problem: by a combination of qualitative reasoning and mental simulation as well as through a repeated process of limited sampling in the actual environment, observation of the consequences and adjusting our mental simulation. Using our proposed method, we are able to solve even very complicated instances of the hole-in-one problem. An extended version of this paper is available [6].

2 MODELING AND SOLVING HOLE-IN-ONE

We choose the following idealisation of the physical environment, which is often used in physics puzzle games: 1. The environment is a restricted 2D plane. 2. Objects are 2D rigid bodies with polygonal or circular shape. 3. There is a uniform downward gravitational force. 4. Object movements and interactions obey Newtonian physics. 5. Physics parameters of objects and the environment remain constant. We call this environment PHYS2D. An instance of PHYS2D is a tuple $\langle E, \mathbf{O}, \mathbf{P}, \mathbf{T} \rangle$, where E is the restricted plane where the objects are located, $\mathbf{O}$ a finite set of static rigid objects O, each of which has a shape, a location, an angle and a type, $\mathbf{P}$ is a set of physics

[1] ANU, Australia, email:{xiaoyu.ge,jochen.renz,p.zhang}@anu.edu.au
[2] QCIS, FEIT, UTS, Australia, email:jaehee.lee@uts.edu.au

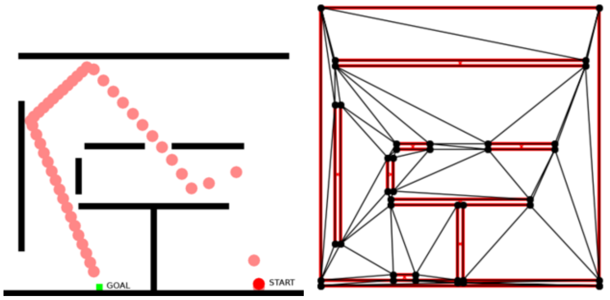

(a) (b)

Figure 1: (a) Illustration of the problem domain in this paper. The green region is the target location H and the solid red circle is the ball B. An identified solution is shown. (b) Triangulation of the scenario

parameters that hold in the environment, such as gravity, and $\mathbf{T}$ is a set of object types and their respective physics properties such as mass and friction, or whether the object can move after being hit or remains static. We assume that all objects are initially static and stable under gravity. A physical action can be applied to an object O by exerting an impulse at a certain position p of the exterior boundary of O. An *impulse* is defined as a pair (θ, μ) where $\theta \in [0, 2\pi)$ is the direction and $\mu \in [0, \mu_{\max}]$ the magnitude of the impact. $\mu_{\max}$ is the maximal magnitude allowed in the environment, both θ and μ are continuous. In this paper we assume there is only *one* start object and *one* goal region. We call this physical action selection problem the Hole-in-One problem.

Definition 1. (Hole-in-One) **Instance:** An instance of Hole-in-One (see Fig. 1a for an example) is a tuple $\langle E, \mathbf{O}, \mathbf{P}, \mathbf{T}, B, H \rangle$, where we use a scenario of PHYS2D and determine a ball $B \in \mathbf{O}$ as the start object and H as the target hole, a polygon in E with a given location. **Solution:** A *solution* is a physical action applied to an object B such that B is delivered to the hole H as a consequence of the putt. To simplify the problem, we fix p to be the centroid of B.

The Hole-in-One problem distinguishes itself from common AI planning problems in that its search space is *infinite* and in particular the action space is *continuous*: Infinitely many different actions can allow an object to take infinitely many different paths. We propose the following method to solve this problem:

- As input scenario, we take the information about the physical environment that we obtained through potentially noisy perception.
- We first partition the free space of the given scenario into finitely many triangular zones (Fig. 1b).
- We defined qualitative rules that describe the physics that govern the transition of moving objects between the triangular zones. We use these rules to generate sequences of qualitative transitions between zones that coincide with potential real paths a moving object can take to achieve the goal. We call such a sequence a *qualitative path*.
- Once all qualitative paths are determined, we rank them by their likelihood of being realized, before we try to realize them.
- We now use a physics simulator that approximates the environment based on our input scenario to search for physical actions that realize the qualitative paths in their ranked order, i.e., actions that allow objects to follow the qualitative paths.
- The solutions we obtain here are not necessarily solutions in the real environment. Therefore, whenever we obtain a solution in our simulator, we immediately apply the solution to the real environment and see if it works. If it does not lead to a real solution, we adjust the object information in our simulator before we continue with the previous step. we will not adjust the triangulation or the qualitative paths when we adjust objects in our simulator.

3 EVALUATION

We evaluated our method in a virtual environment simulated by the Box2D (www.box2d.org) physics engine. The method also uses Box2D for its internal simulation with an incomplete knowledge of the environment. We perturb the visual input to the internal simulation by rotating each object at an angle sampled from a zero mean Gaussian. We created 72 mini-golf scenarios. The scenario designs are inspired by the game levels of a popular video game of mini-golf [3]. We compare our method with a solver (S_G) which uses a goal-oriented sampling strategy. The sampling strategy of S_G is similar to the one used in [9] that adjusts actions according to the distance between the final position of the ball and the target destination. Our method outperforms S_G in all the scenarios. S_G is less efficient because there could be many local optima in a problem. By contrast, our method can detect more different types of solutions (if there are any), which helps to avoid these local optima. Qualitative reasoning and triangulation can be achieved efficiently; it takes on average 4 seconds to generate qualitative paths based on a triangulation with around 60 zones. As the noise level increases, our method can still detect and realize qualitative paths that lead to the goal. Such qualitative paths have similar bounce sequences as their counterparts derived from perfect triangulation. Further details can be found in [6].

4 CONCLUSION

We studied a realistic problem that contains some of the essential components AI needs to successfully interact with the real world: being able to predict the consequences of physical actions and to select a physical action out of an infinite number of actions that achieves a specified goal. The proposed method involves a combination of qualitative reasoning and internal simulation together with testing proposed actions in the real environment and, if necessary, adjusting our internal knowledge based on the new observations.

REFERENCES

[1] P.W. Battaglia, J.B. Hamrick, and J.B. Tenenbaum, ‘Simulation as an engine of physical scene understanding’, *Proceedings of the National Academy of Sciences*, **110**(45), 18327–18332, (2013).

[2] E. Davis, ‘Physical reasoning’, *Handbook of Knowledge Representation*, 597–620, (2008).

[3] E. Davis, G. Marcus, and A. Chen, ‘Reasoning from radically incomplete information: The case of containers’, in *ACS*, volume 273, p. 288, (2013).

[4] M.R. Dogar, *Physics-Based Manipulation Planning in Cluttered Human Environments*, Ph.D. dissertation, MIT, 2013.

[5] K.D. Forbus, ‘A study of qualitative and geometric knowledge in reasoning about motion’, Technical report, Cambridge, USA, (1981).

[6] X. Ge, J.H. Lee, J. Renz, and P. Zhang, ‘Hole in one: Using qualitative reasoning for solving hard physical puzzle problems’, in *Proc. 29th International Workshop on Qualitative Reasoning*, (2016).

[7] S.M. Khansari-Zadeh, K. Kronander, and A. Billard, ‘Learning to play minigolf: A dynamical system-based approach’, *Advanced Robotics*, **26**(17), 1967–1993, (2012).

[8] S. Kumar and S. Chakravorty, ‘Adaptive sampling for generalized probabilistic roadmaps’, *J. Control Theory Appl*, **10**(1), 1–10, (2012).

[9] D. Wolter and A. Kirsch, ‘Leveraging qualitative reasoning to learning manipulation tasks’, *Robotics*, **4**(3), 253–283, (2015).

[3] http://www.eivaagames.com/games/mini-golf-game-3d/