Applying AI Algorithms to Physical Games: Proposing an Autonomous AI Robotic Foosball Player

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Patrick Zhang

Spring, 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Rosanne Vrugtman, Department of Computer Science

Applying AI Algorithms to Physical Games: Proposing an Autonomous AI Robotic Foosball Player

CS4991 Capstone Report, 2023

Patrick Zhang
Computer Science
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia USA
pgz3nx@virginia.edu

ABSTRACT

Although Artificial Intelligence (AI) has been used to achieve super-human performance in games such as chess, poker, go, and StarCraft, these programs lack a physical component which prevents these algorithms from being able to play physical games, such as foosball, ping pong, or pool. A colleague and I have successfully designed an autonomous foosball playing robot using rudimentary algorithms and strategies. However, it could be improved with recently developed technologies like YOLO object detection and self-play reinforcement learning. These would allow the system to defeat professional human players and achieve superhuman level. Robots and AI have a large potential for social impact expanding the types of games the algorithms can be applied to, reducing the need for multiple people to play physical multiplayer games, and allowing for faster skill development though playing against more skilled opponents.

1. INTRODUCTION

One of the most exciting applications of AI is the ability to achieve superhuman ability in games such as chess, poker, and StarCraft. The appeal of these programs is multifaceted. It is interesting to see these games being played as perfectly as possible. Superhuman AI programs have introduced new ideas and strategies that humans have never considered, completely changing the way humans practice and prepare. The impact of these programs on the experience of amateur and professional human players cannot be overstated.

Currently, these superhuman AIs have primarily been developed on board games such as chess and poker because one can easily play the games with a computer by programming the rules of the game. Thus, the program just needs to output what move to do or what action to take. For physical games, in addition to determining the most optimal action, a physical component is needed to perform the action. Oftentimes, the physical component requires sophisticated movements, as in ping pong. This is why physical games do not have superhuman AIs. For foosball, there are fewer actions than ping pong (move or spin a lever), which is why it is reasonable to create an autonomous robot for.

2. RELATED WORKS

Silver, et. al. (2018) advanced reinforcement learning and deep learning algorithms used to train the strategies utilized by the superhuman AI programs to play chess. It is a general algorithm, which has allowed the program to be applied to other board games such as go and shogi. Many other superhuman chess programs have been developed using similar techniques. If applied to foosball, novel and potentially superhuman strategies could be discovered and developed.

Robots playing physical games is nothing new. Popescu, et. al. (2021) researched how to use drones to play foosball. Ping pong is a more physically demanding game from a robotics perspective than foosball. As a result, much of the research is focused on the robot's physical ability to hit the ping pong ball, rather than the AI techniques to advance ping pong strategy. Regardless, this is an exciting example of robots playing physical games.

There have also been attempts to make robotic foosball players, such as the one created by Krause, et. al. (2018). In addition to robotics, it uses neural networks to "learn" a strategy to coordinate the levers. This differentiates it from similar attempts. However, there is still room to build on this approach. They do not mention the performance of this robot against humans, but they note that it will take years for the robot to train, meaning it would take an immense effort by humans to play against it. This could be improved using more sophisticated training methods such as the self-play algorithms developed by Silver et. al.

3. PROJECT DESIGN

This section introduces a rudimentary design of the autonomous robotic foosball player and proposes several improvements.

3.1 Overview

Adding a physical component to an AI program significantly influences the design. As previously mentioned, using AI to play board games such as chess and poker is heavily researched. Most of the algorithms used to develop these programs are better suited for games that can be played online. As a result, physical games such as foosball have not had superhuman programs because they require complex robotic systems to play. However, other than that, the same algorithms used to develop superhuman chess and StarCraft programs can be used to develop strategies to play physical games such as foosball. Games such as foosball require strategy and coordination in addition to raw physical capabilities, so there is a lot of room for computers to improve human knowledge and abilities.

The game of foosball presents an interesting set of challenges. In other games such as ping pong, you only hit the ball once before the other player hits it, so there is less strategy, and more physical ability (such as accuracy and power) involved. In foosball, there are multiple levers (rows of players), so you can pass the ball between levers, resulting in the possibility of complex coordination. Because each lever can only do two things, move laterally and spin, the robotics are relatively simple, allowing more effort to be spent on the strategy element. This makes foosball an ideal game to introduce superhuman level strategy to physical games.

3.2 First Iteration

A first iteration of an autonomous AI robotic foosball table has been designed and built by me and a colleague. It contains 3 components: vision, strategy, and robotics.

3.2.1 Vision

The foosball table needs to know where the ball is at all times. This is accomplished using a camera mounted above the table, which records the table. Then, a computer vision program runs an object detection algorithm to calculate the coordinates of the ball.

The camera used is a high-speed HD webcam. The images need to be high definition so that there is enough detail to locate the ball, and it needs to be sufficiently high speed to allow the program to react in time to the ball's movements. The object detection algorithm uses an open-source library called OpenCV. The first step is a threshold, where light pixels are set to white and dark pixels are set to black. Because the ball is white, the resulting image is black, except for the pixels that make up the ball. Then, contour detection is run to identify groups of pixels that belong to the same object. The biggest such group is likely the ball. Lastly, a “minimum enclosing circle” algorithm is run to get the coordinates of the circle that are the ball.

3.2.2 Strategy

The coordinates of the ball are fed into the next component, which is the strategy component. The path of the ball is calculated using the past coordinates. Assuming the ball is going in a straight line, it is relatively simple to predict the future coordinates of the ball, while accounting for bounces off of the edge of the table.

The main part of the strategy is deciding what the levers should do (where to move them and when to spin) based on the calculated future points of the ball. Currently, a rudimentary strategy is implemented where the lever simply moves such that a player is where the path of the ball intersects the lever. When the ball gets close, spin the lever to kick the ball as fast as possible.

3.2.3 Robotics

The robotic component needs to accomplish two things: move the levers back and forth and spin the levers to kick the ball. Linear actuators are used to move the levers and servo motors to drive the linear actuators and rotate the levers. One motor per lever drives the linear actuators and one motor per lever rotates the levers, resulting in four linear actuators and eight servo motors for a standard foosball table with four levers per side.

3.3 Proposed Improvements

The high-speed object detection and advanced robotics can consistently hit the ball back towards the opponent’s goal, which is enough to defeat novice human players. However, it is nowhere near the level of a professional foosball player. Unlike novice games, professional players slowly and methodically score goals by passing the ball from one lever to another, while avoiding the opponents’ players. To match the success of AI in games like chess and poker, there are several recent technologies allow an autonomous robotic foosball player to defeat a professional player. In particular, these improvements are in the vision and strategy components.

3.3.1 Vision Improvements

The current object detection algorithm is heuristics-based and has many shortcomings. It relies on the color of the ball being very different from the rest of the objects in the field of view. When the ball is partially obstructed, for example by a player or a lever, the algorithm is unable to detect the ball. Thus, the ball is only detected in ideal conditions, not all conditions.

The solution is using a machine learning-based computer vision algorithm such as YOLO (You Only Look Once), which works by using a neural network to find potential bounding boxes of objects and simultaneously assigning probabilities that those bounding boxes are accurate. The best bounding boxes are selected and merged if necessary. These steps allow the algorithm to consider the entire context of the image in to make accurate detections. YOLO has consistently outperformed other object detection methods in both accuracy and speed. This allows the ball to be detected even when partially obstructed because the algorithm can infer that the obstructed part is also a part of the ball. Also, the algorithm does not rely on the ball being a different color than the rest of the objects in the field of view. These factors allow for a much more robust object detection system.

One of the latest versions of YOLO is YOLOv7, which was released in 2022. YOLOv7 can achieve speeds of up to 160 frames per second, which is fast enough to react during foosball games.

3.3.2 Strategy Improvements

Currently, the robotics just try to hit the ball back towards the opponent. To defeat professional human players, however, the robotic foosball table needs to employ much more sophisticated strategies. Given the predicted path of the ball, the strategy component needs to calculate the most optimal actions - where to move the levers and strike the ball. The reinforcement learning algorithms used on chess and poker can be applied to this problem.

This has never been attempted with foosball before. Advanced chess AIs are trained through self-play, which is where the AI plays against itself and slowly learns better and better strategies. When applied to foosball, this would require controlling both sides by the AI, then making both sides play against each other.

Such self-play training processes are better suited for games that can be simulated such as board games because the games can be sped up. Physical games such as foosball do not have this luxury. As a result, this kind of training could take a very long time. To ameliorate this issue, an online simulation of foosball could be created for training the strategy component. Once the strategy learned from this approach is integrated into the rest of the physical system, the robotics will be able to take the optimal actions.

4. ANTICIPATED RESULTS

After applying recently developed technologies, the improved robotic foosball player should be able to employ sophisticated strategies to defeat professional human players, achieving superhuman level. The object detection component should be able to detect the coordinates of the ball quickly and accurately in all conditions, including when the ball is obstructed or colored similarly to the table/other objects. The strategies employed should be methodical and at least as effective as those employed by professional human players. The strategies would likely involve moving the ball within players of the same lever and passing the ball to players of other levers. The AI should be able to pass the ball to other levers or shoot the ball accurately at the goal (rather than in the general direction of the goal), while avoiding opposing players. When defending, the AI should be able to detect the ball and calculate its path fast enough to move the defenders into position in time. Such a robot should strike the ball with more accuracy, speed, and sophistication to execute more complicated scoring attacks, and would show the power of AI combined with physical components.

5. CONCLUSION

This project extended existing superhuman AIs in board games such as chess and poker to foosball, a physical game. This enables players of foosball to experience the same benefits that superhuman AIs brought to chess and poker, such as enabling faster improvement and more sophisticated analysis. A fully autonomous robotic foosball player with a rudimentary strategy was implemented and described. Recent technologies that would further improve the performance of the player such as reinforcement learning and YOLO object detection were introduced. These improvements should allow the autonomous player to go from defeating novice players to defeating professional foosball players.

6. FUTURE WORK

The next step of this project is to implement the proposed sophisticated techniques. This would likely involve integrating more powerful computing to run the more advanced YOLO object detection models. Also, further research into more efficient ways of simulating or training the reinforcement learning models is required. Finally, the design of this project can be expanded to other physical games such as ping pong, pool, shuffleboard, and more.

REFERENCES

[1]. Kundu, R. (2023, January 17). *Yolo algorithm for object detection explained [+examples]*. V7. https://www.v7labs.com/blog/yolo-object-detection

[2]. Popescu, R. I., Raison, M., Popescu, G. M., Saussié, D., & Achiche, S. (2021). *Design and development of a novel type of table tennis aerial robot player with tilting propellers.* Mechatronics, 74. https://doi.org/10.1016/j.mechatronics.2021.102483

[3]. Silver, D., Kasparov, G., & Habu, Y. (2018, December 6). *Alphazero: Shedding new light on chess, Shogi, and go*. Google DeepMind. https://www.deepmind.com/blog/alphazero-shedding-new-light-on-chess-shogi-and-go

[4]. *Table soccer robot with Artificial Intelligence*. Bosch Global. (n.d.). https://www.bosch.com/stories/bend-it-like-bosch/

[5]. Wang, C.-Y., Bochkovskiy, A., & Liao, H.-Y. M. (2022). *YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors.* ArXiv (Cornell University). https://doi.org/10.48550/arxiv.2207.02696