

Human-Robot real world play: Hide-and-Seek

Many similarities of the hide-and-seek game can be found in the interactions between humans, or robots and humans, in urban spaces, similar to looking for a person in a crowded urban environment (Goldhoorn et al. 2014). It is suggested that the game of hide-and-seek is an ideal domain for studying cognitive functions in robots and it is a basic mechanism for human robot interaction in mobile robotics, because hide-and-seek requires the robot to navigate, search, interact on, and predict actions of the opponent (Johansson and Balkenius 2005).

In this project, we develop a hide and seek game (the seeker is the robot and the hider is the person or object) to investigate how the robot navigates, interacts on, and predicts action of the opponent. The project implements machine learning techniques such as reinforcement learning, which rewards to robot everytime it successfully finds a hider. This information is used to train the robot to find a hider, based on purely visual inputs.

Through the artistic novelty of human-robot playing we explore the real-world implications in supporting human disaster response teams. For example, one of the many practical applications of this experiment is to find persons; as hide-and-seek could be seen as a simplification of search-and-rescue.

Our version of the hide and seek is as follows: there are two players, hider and seeker, which play the game on a grid square of a fixed amount of cells. The grid contains: a free cell called the base, other free cells on which the players can move, and obstacle cells that are not accessible by the players and also limit their mutual visibility.

The base cell, where the seeker will start each game and where the hider will attempt to reach the base before the seeker "catches it." The game will run the maximum amount of time steps. At each time step, the hider and the seeker take turns to move to a free neighbouring cell (8-connectivity neighborhood (i.e. a maximum of 9 actions for each player)). The hider wins if it successfully reaches the base before being caught by the seeker, and the seeker wins if it captures the hider within one neighboring block. The result will be a tie if neither the seeker catches the hider, nor the hider manages to reach the base.

In our game, we believe that the seeker's position is always fully observable and the hider's position is not always observable, (the hider may be partially or fully covered by obstacle cells). The seeker will rotate 360 degrees to capture images of the seeker's surroundings. The images can be stitched in an image-stitching programme to construct an environment which the seeker can identify the hider.

The process of identifying a hider is done by an image classifier using a model-free reinforcement learning approach. To classify an object from the image as 'the hider' is first given a training dataset of images of a hider, which, through trial and error, the seeker will identify which area of the screen the hider is in. The size and location of the hider will be used as inputs for the model-based reinforcement learning, which will determine the cell the seeker should navigate to.

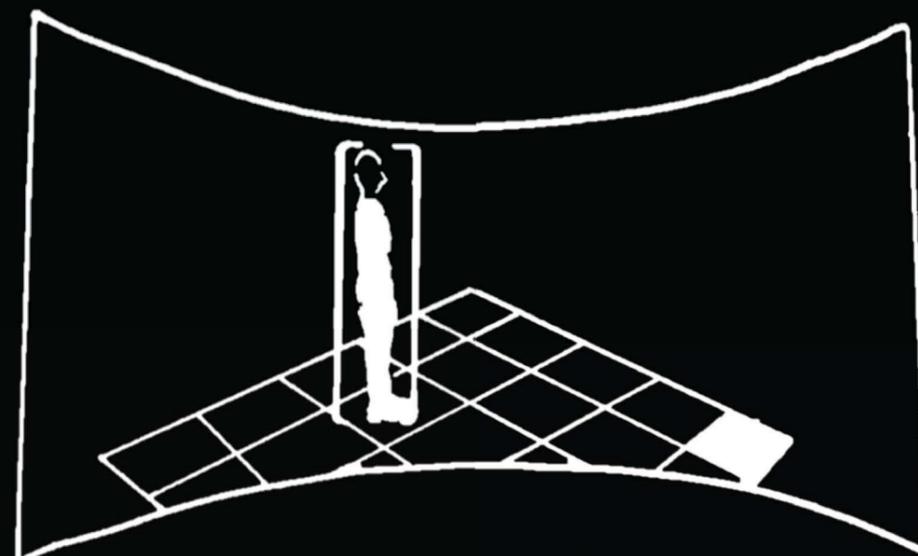


Figure 1: this sketch shows the environment created from the image-stitching programme. The seeker identifies the hider and the base from the constructed environment.

The human and surrounding audience will be able to view the robots camera stitching at every move on the projector with text explaining the status of the search.

The following text will be displayed to animate the seeker's search for the hider:

- "Where did you go?" (if hider is obstructed by object.)
- "I see you!" (if hider is in camera viewing.)
- "HAHA got you now" (when seeker is close to the hider.)

In order for the seeker to take an action in response to the hider's movement, a model-based reinforcement learning algorithm will make use of a reward function to successfully catch the hider. Reinforcement learning (RL) is an effective class of decision making algorithms (Luo, 2017), which will be used to train a robot to find a hider based purely on visual inputs.

The seeker will explicitly gather statistics from the constructed environment, use a reward function to take an action. The statistics that will be computed are:

- the distance of the seeker from the hider (d_{sh})
- the hider from the base (d_{hb})
- the seeker from the base (d_{sb})

These will be obtained from analysing the environment constructed from the 'stitched images'.

The reward function is based on a seeker move that:

- Minimises the distance of the seeker from the hider (d_{sh})
- Maximises the hider from the base (d_{hb})
- The distance of the seeker from the base is less than the distance of the hider from the base ($d_{sb} < d_{hb}$).

If all of these parameters are met then a reward is given. Both the robot and human have to move one cell per turn, and may not maintain their current position.

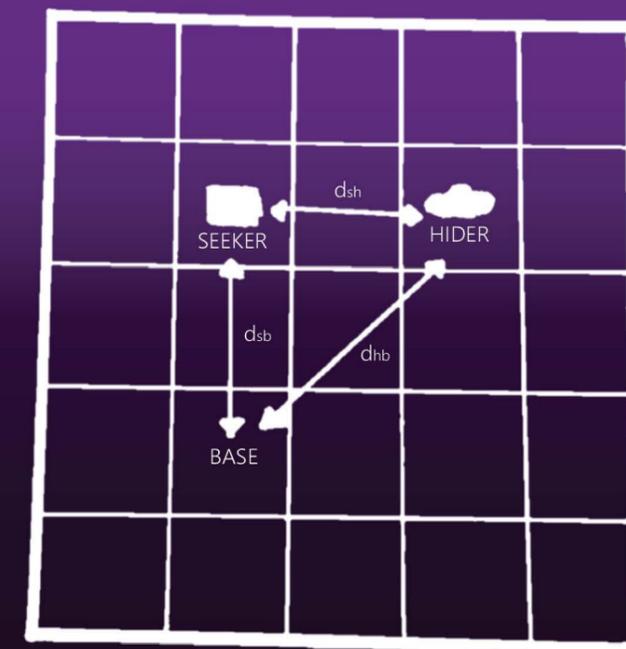


Figure 2: This sketch shows a birds eye image of the distances used to compute the reward function.

This experiment will give us important insights into the functionality of automated seeker methods used by a robot in real-world playing (interacting and predicting) against humans. Although the novelty of human-robot playing may evoke purely artistic expression, this experiment explored the first steps toward real world search-and-rescue simulation.

References:

- Armada, M. A. et al., (2014) ROBOT2013: First Iberian Robotics Conference, 505 Advances in Intelligent Systems and Computing, 253.
- Goldhoorn, A., Sanfeliu, A. and Alquezar, R. (2014) Analysis of methods for playing human robot hide-and-seek in a simple real world urban environment. In ROBOT2013: First Iberian Robotics Conference (pp. 505-520). Springer, Cham.
- Johansson, E., Balkenius, C. (2005) It's a child's game: Investigating cognitive development with playing robots. In: International Conference on Development and Learning, vol. 164.
- Luo, R. (2017) Collecting boxes in the Unreal. (available at: http://www.rodgerluo.com/projects/rl_drones.html).