

Learning to Move Autonomously in a Hostile World (sketches_0161)

Supplemental Figures

Leslie Ikemoto, Okan Arıkan, David Forsyth
University of California, Berkeley




Figure 1: **Avoiding an Obstacle.** This sequence was recorded from a live, interactive demo, in which the user controls the crate and can move it anywhere at anytime. A virtual agent is tasked with traveling from the left of the scene to the target on the right. While the agent is running towards the target (**left**), the user suddenly moves the crate into the agent's path. The local motion planner successfully replans the motion on the fly. The agent not only dodges the crate and meets the target position (**right**), but does so seamlessly, since transitions are derived from a pre-computed motion graph. This sequence further illustrates that the local motion planner has learned to balance competing goals of (a) reaching the target position quickly, and (b) avoiding obstacles.




Figure 2: **Hiding from an Enemy.** In this scenario, the agent must hide from the enemy skull who scares him. On the **left**, the agent begins walking towards the goal, but notices said skull around the corner. He hides just behind the crate (**inset**) until the enemy disappears. The system then leads the agent to his goal position (**right**). This sequence demonstrates that the local motion planner *independently* learned the emergent behavior of using obstacles to hide from enemies; behaviors are not explicitly encoded.



Figure 3: **Choosing Alternatives.** When the agent reaches a branching point in the motion graph, the local motion planner evaluates the available alternatives. In this example, there are three: walking backwards to the left (**blue**), walking to the right then turning left (**green**), and walking to the right (**red**). The local motion planner estimates the long-term expected goodness of each path by predicting where each obstacle and visible enemy will be in one second, assuming they continue moving at their current velocities. For each alternative, the local planner calculates the position of the agent after one second. Here, the planner predicts that a crate will land in the path of the green alternative.




Figure 4: In this final scene, the agent waits behind a crate to avoid a patrolling enemy (**left**). Once the enemy disappears, the agent runs through the opening in the obstacles, then times his exit through the sliding crates (**right**). The sliding crates are shown as translucent, and the arrows on the left side indicate their range of motion.