

# Communication protocol for the Mountain-Car Simulator

May 16, 2006

The first message sent by the agent is a `GETTASK` message, which returns the number of state variables (2 in our case) and the number of action variables for this environment.

At the start of each episode, the agent sends the message `INIT`. The environment sends back a message starting with the word `START`, followed by the initial position and velocity of the car. Note that the environment can be started with either a fixed position (at the bottom of the hill) or with a random position.

Then, the agent uses the message `STEP` to make a step. This message contains the dimensionality of the action vector (in this case, 1), and the action. The environment responds with a message containing the reward, a termination flag (1 if the episode has ended, 0 otherwise), the number of state variables sent (in our case, always 2) and the state variables (position and velocity).

If anything unexpected happens, the environment closes the socket and terminates the session.