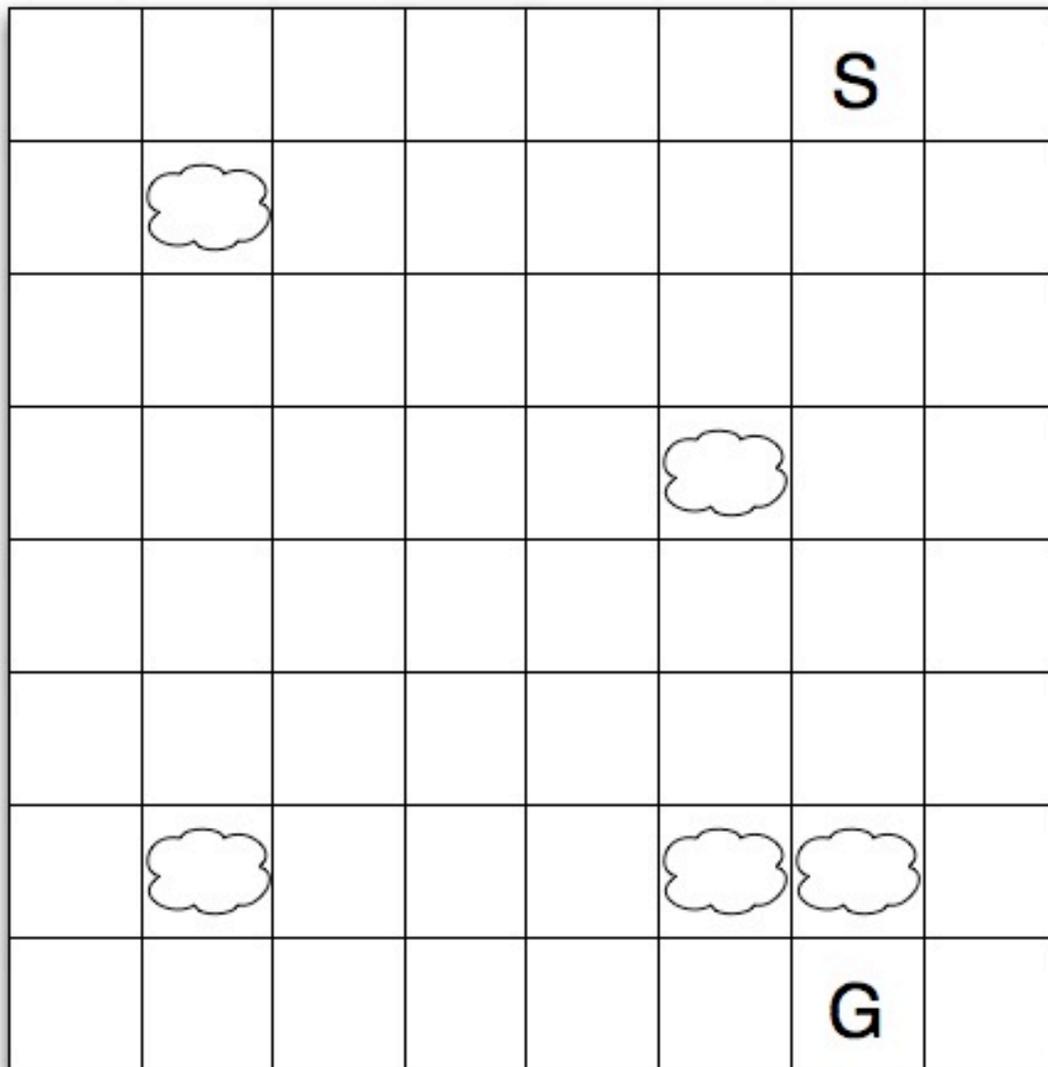


Gridworld with Mines

The figure below shows a standard gridworld, with start and goal states, but with one difference: there are mines located at various positions on the grid. These mines are stationary and cause the agent to be destroyed if touched. The actions are the standard four, *up*, *down*, *right*, and *left*. The state is represented by a integer value of the state label, comprised of the row index (*x*) multiplied by the number of columns plus the column index (*y*), resulting in 215 unique states. Let us treat this as a undiscounted episodic task, with constant rewards of +10 for reaching the goal, -10 for hitting a mine and -1 otherwise. Agent is started in a new random starting state at the beginning of each episode.



 = MINE