CSE 573: Artificial Intelligence
Winter 2019

Hanna Hajishirzi

Markov Decision Processes

slides from
Dan Klein, Stuart Russell, Andrew Moore, Dan Weld, Pieter Abbeel, Luke Zettelmoyer

Recap: Defining MDPs

- Markov decision processes:
  - Set of states S
  - Start state s_0
  - Set of actions A
  - Transitions P(s'|s,a) (or T(s,a,s'))
  - Rewards R(s,a,s') (and discount γ)

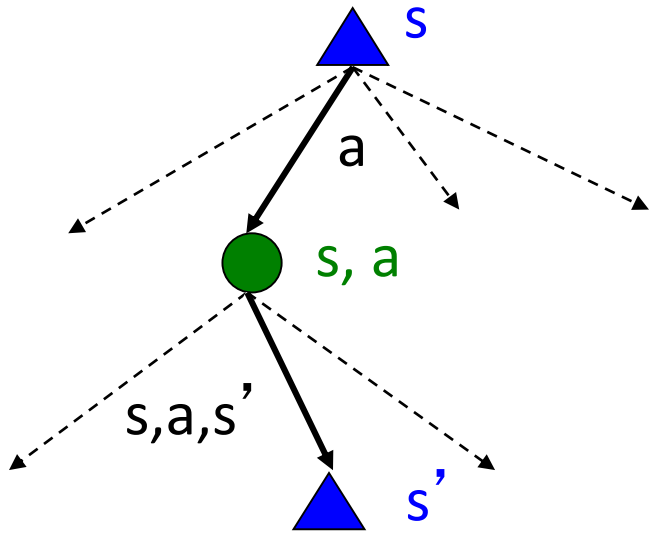

- MDP quantities so far:
  - Policy = Choice of action for each state
  - Utility = sum of (discounted) rewards

Values of States (Bellman Equations)

- Fundamental operation: compute the (expectimax) value of a state
  - Expected utility under optimal action
  - Average sum of (discounted) rewards
  - This is just what expectimax computed!

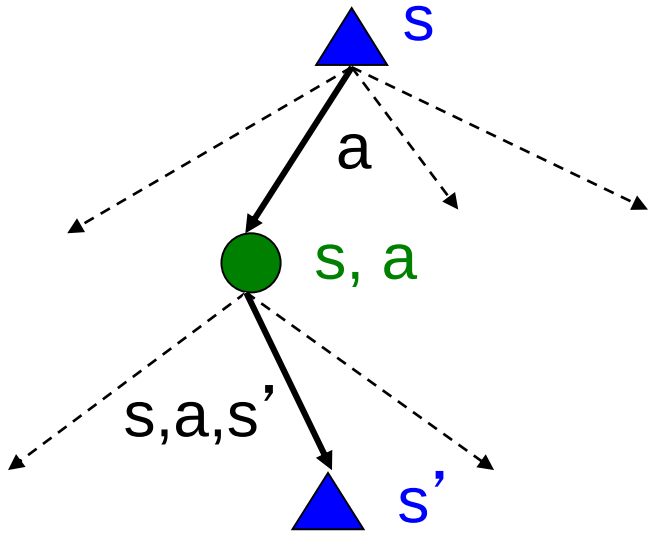

- Recursive definition of value:

$$V^*(s) = \max_a Q^*(s,a)$$

$$Q^*(s,a) = \sum_{s'} T(s,a,s') \left[R(s,a,s') + \gamma V^*(s') \right]$$

$$V^*(s) = \max_a \sum_{s'} T(s,a,s') \left[R(s,a,s') + \gamma V^*(s') \right]$$

Racing Search Tree

- We're doing way too much work with expectimax!
- Problem: States are repeated
  - Idea: Only compute needed quantities once
- Problem: Tree goes on forever
  - Idea: Do a depth-limited computation, but with increasing depths until change is small
  - Note: deep parts of the tree eventually don't matter if $\gamma < 1$

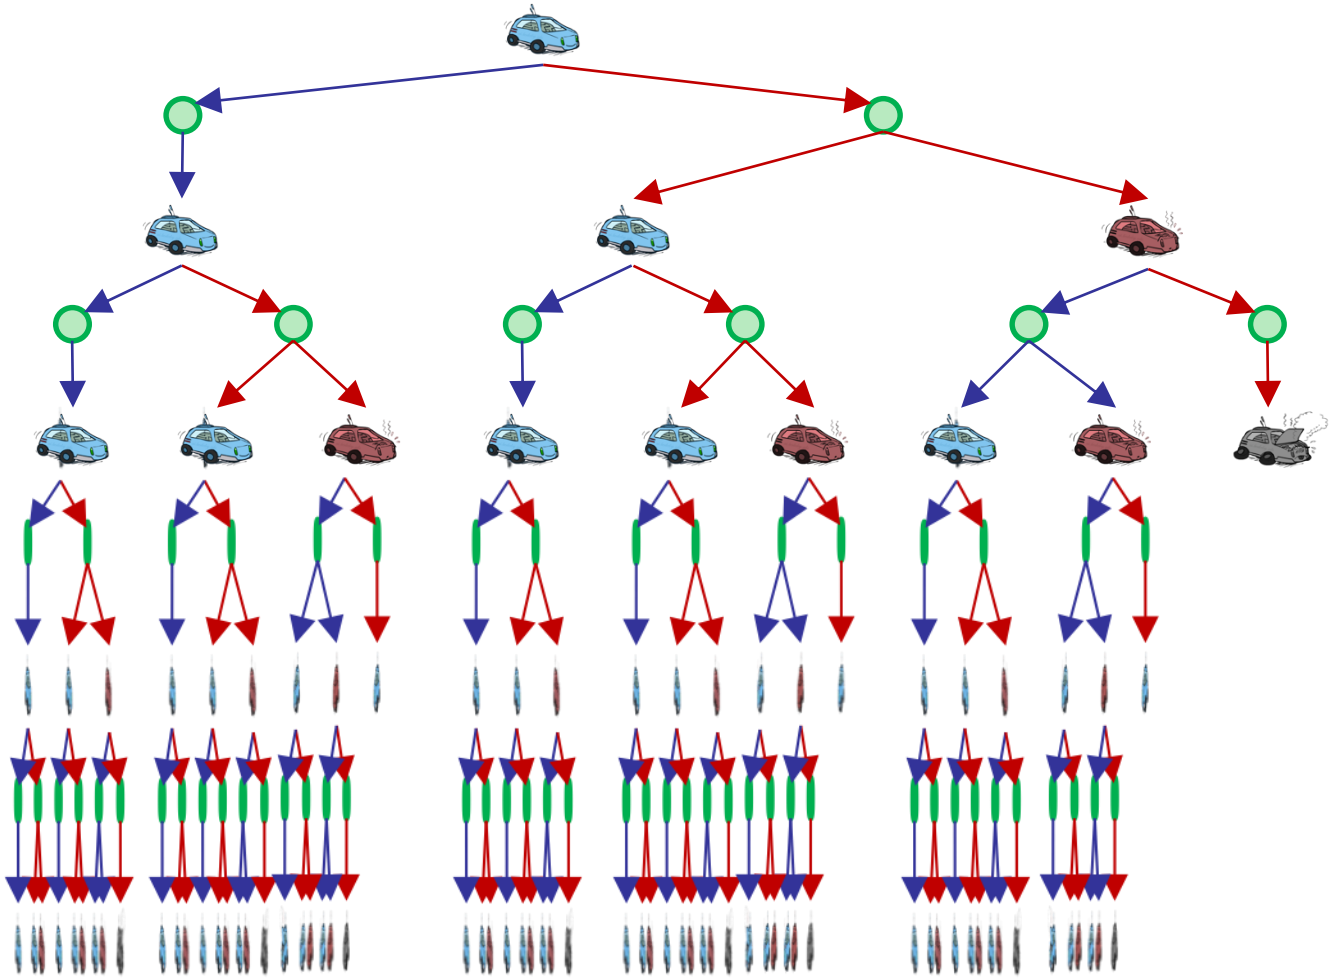

Time-Limited Values

- Key idea: time-limited values
- Define $V_k(s)$ to be the optimal value of s if the game ends in k more time steps
  - Equivalently, it's what a depth-k expectimax would give from s

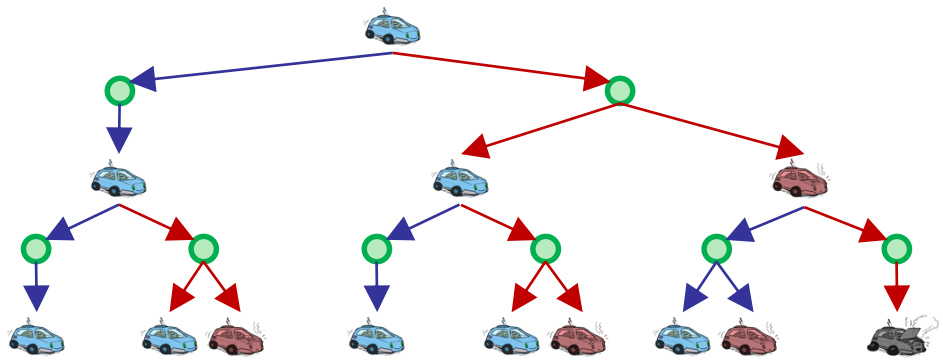

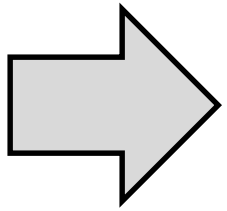

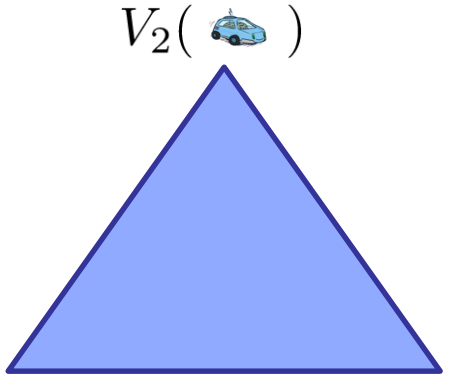

Computing Time-Limited Values

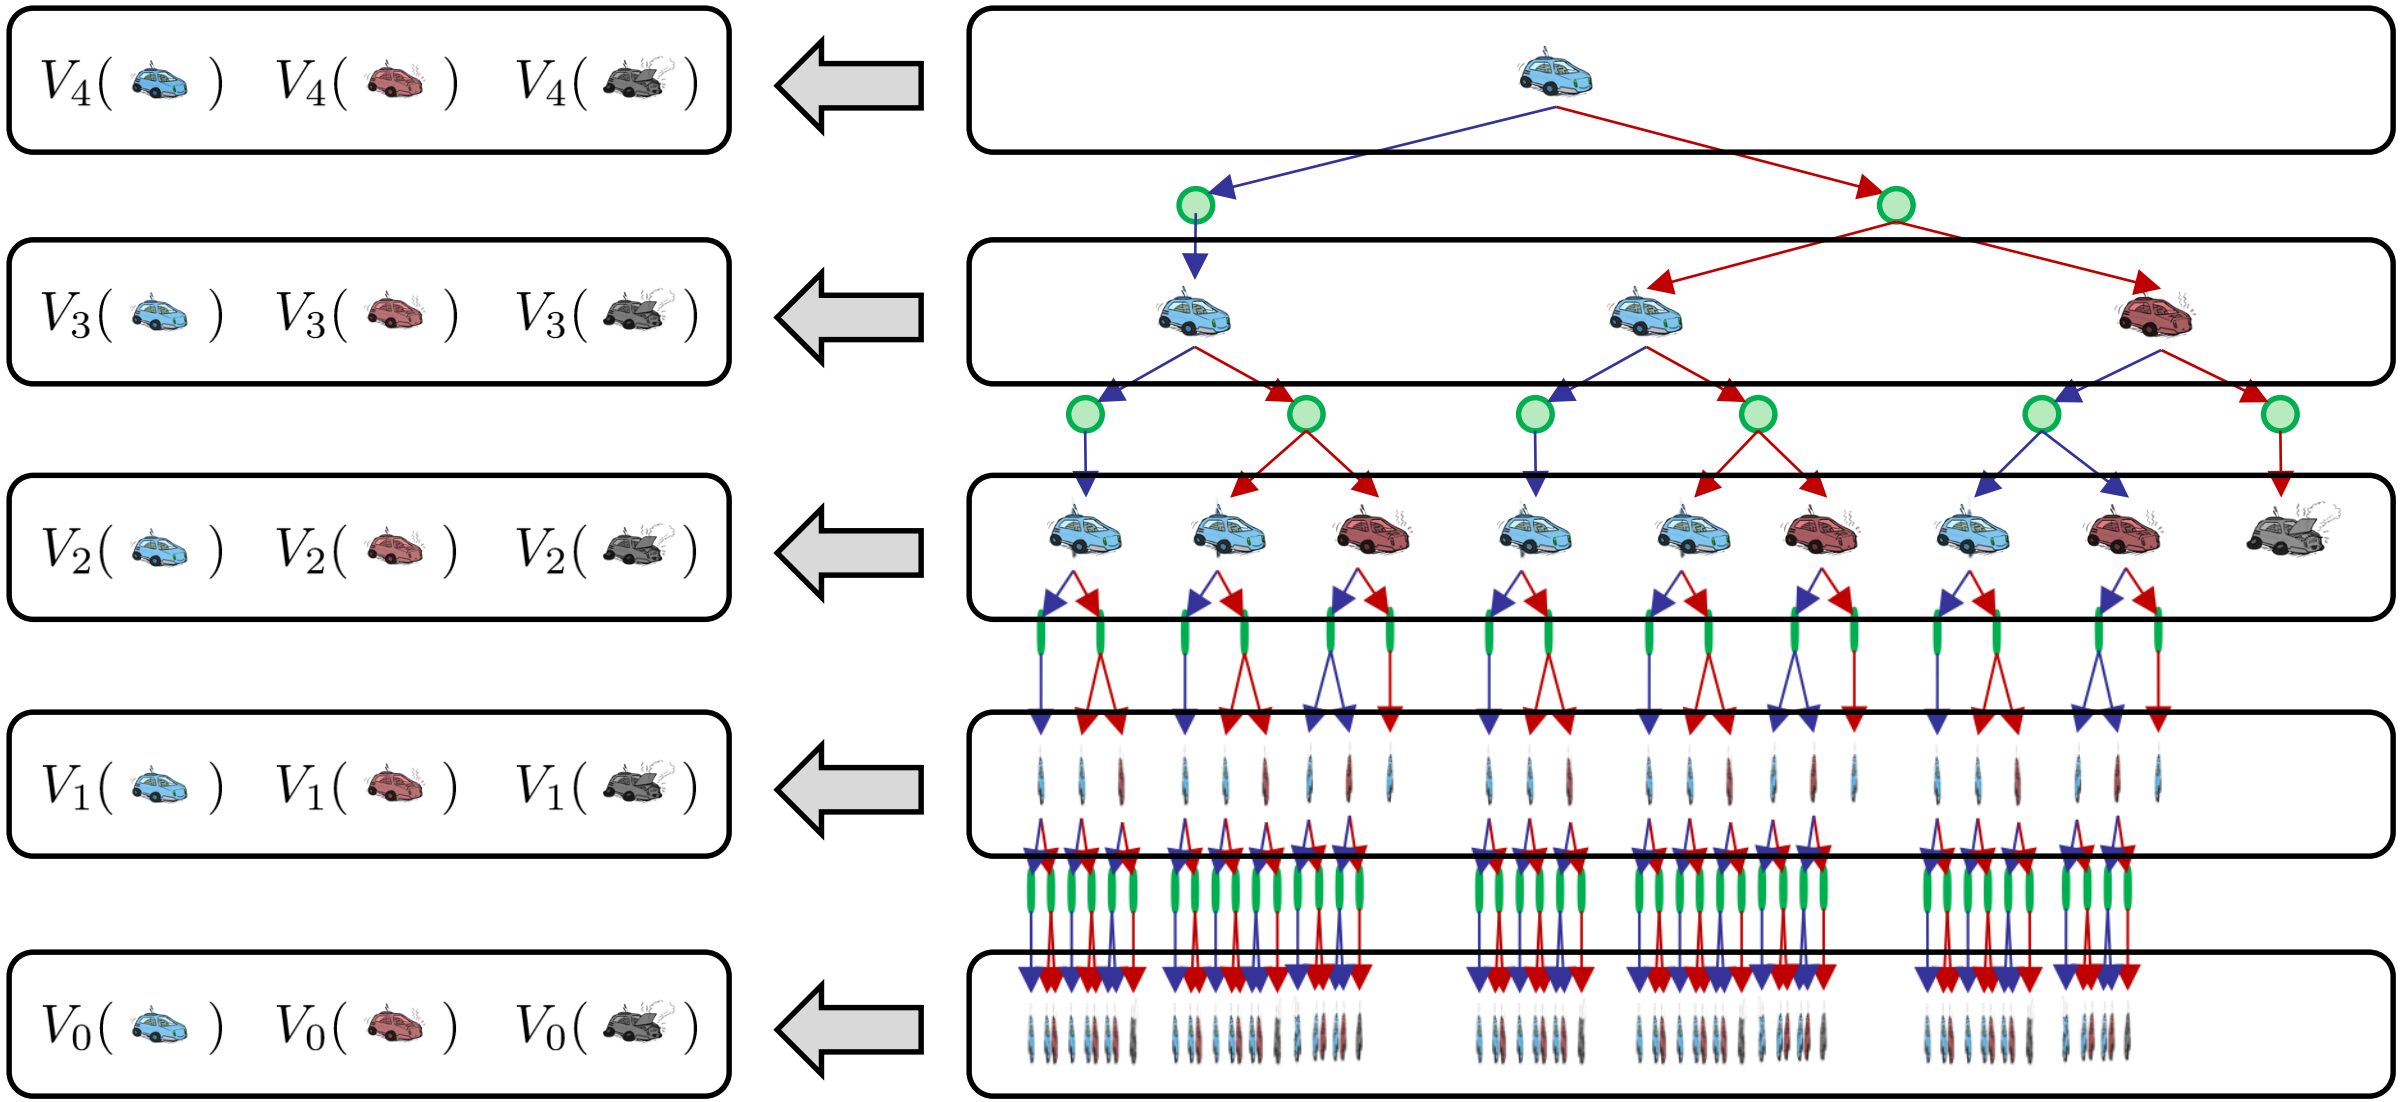

k=0

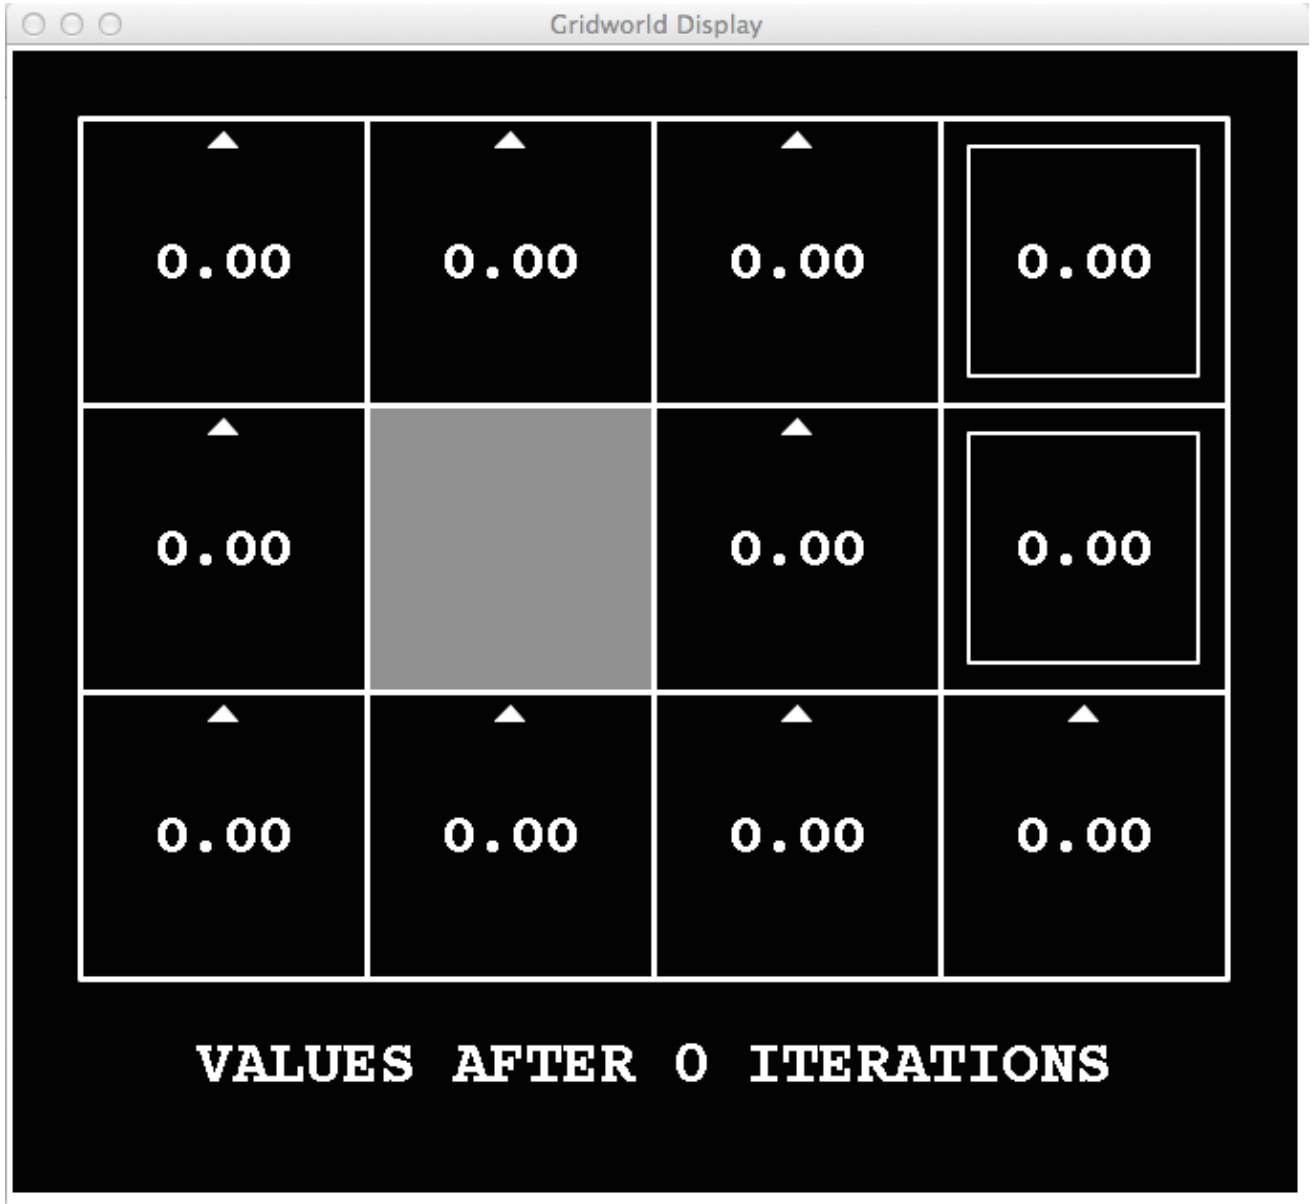

Noise = 0.2
Discount = 0.9
Living reward = 0

k=1

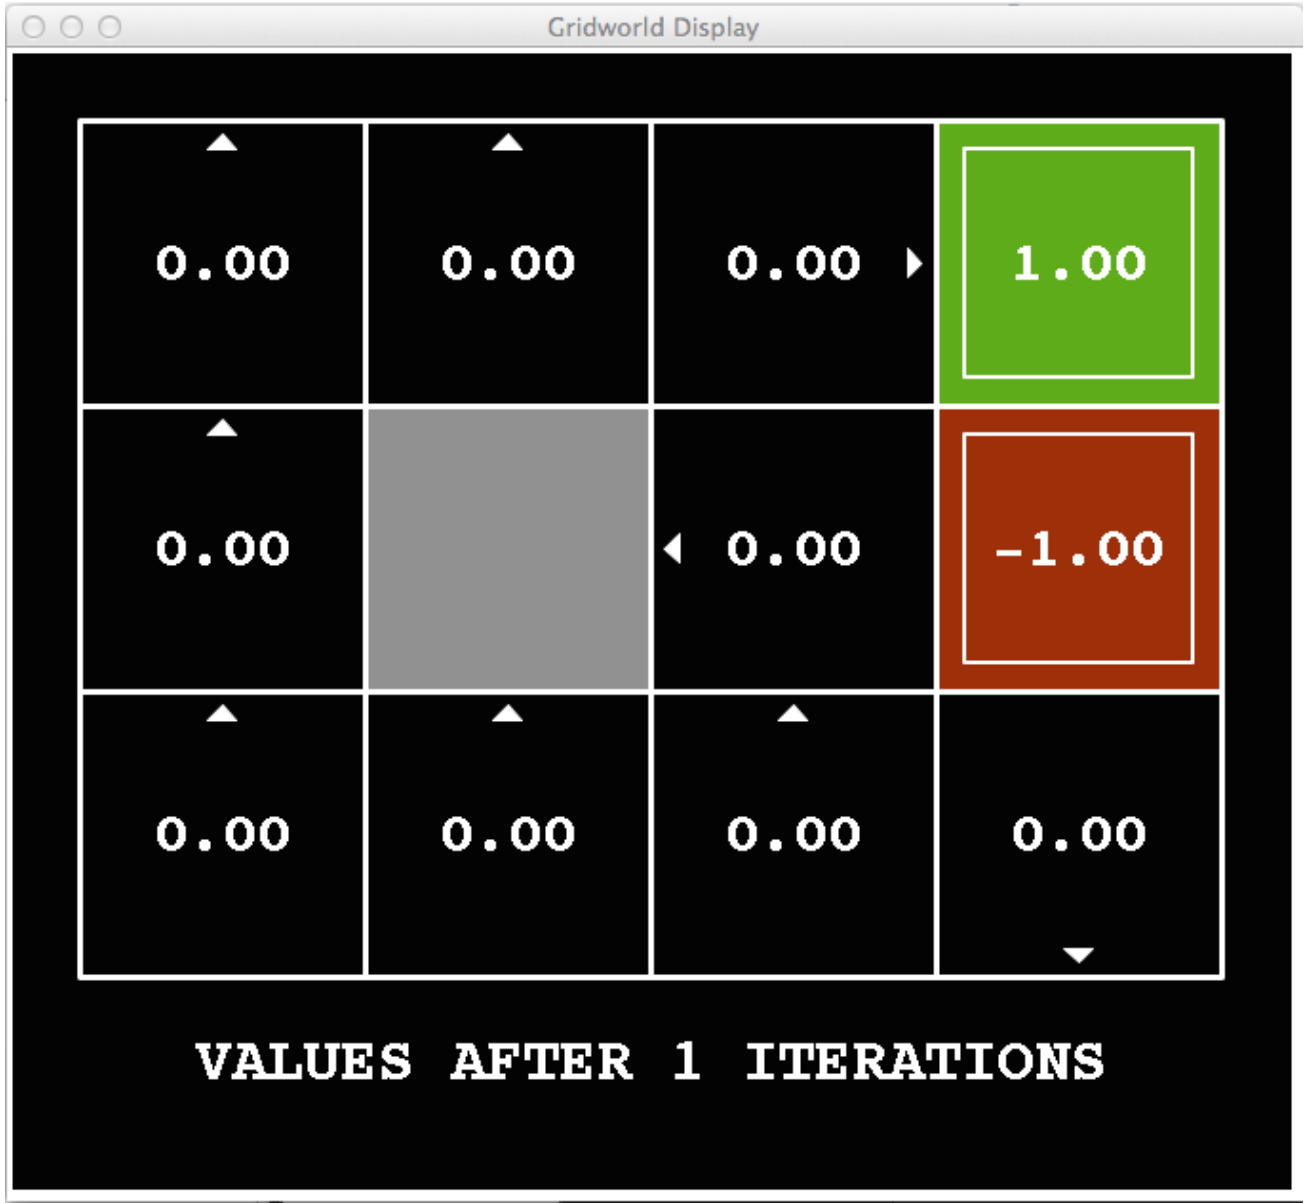

Noise = 0.2
Discount = 0.9
Living reward = 0

k=2

Noise = 0.2
Discount = 0.9
Living reward = 0

Bellman Updates

Example: γ=0.9, living reward=0, noise=0.2

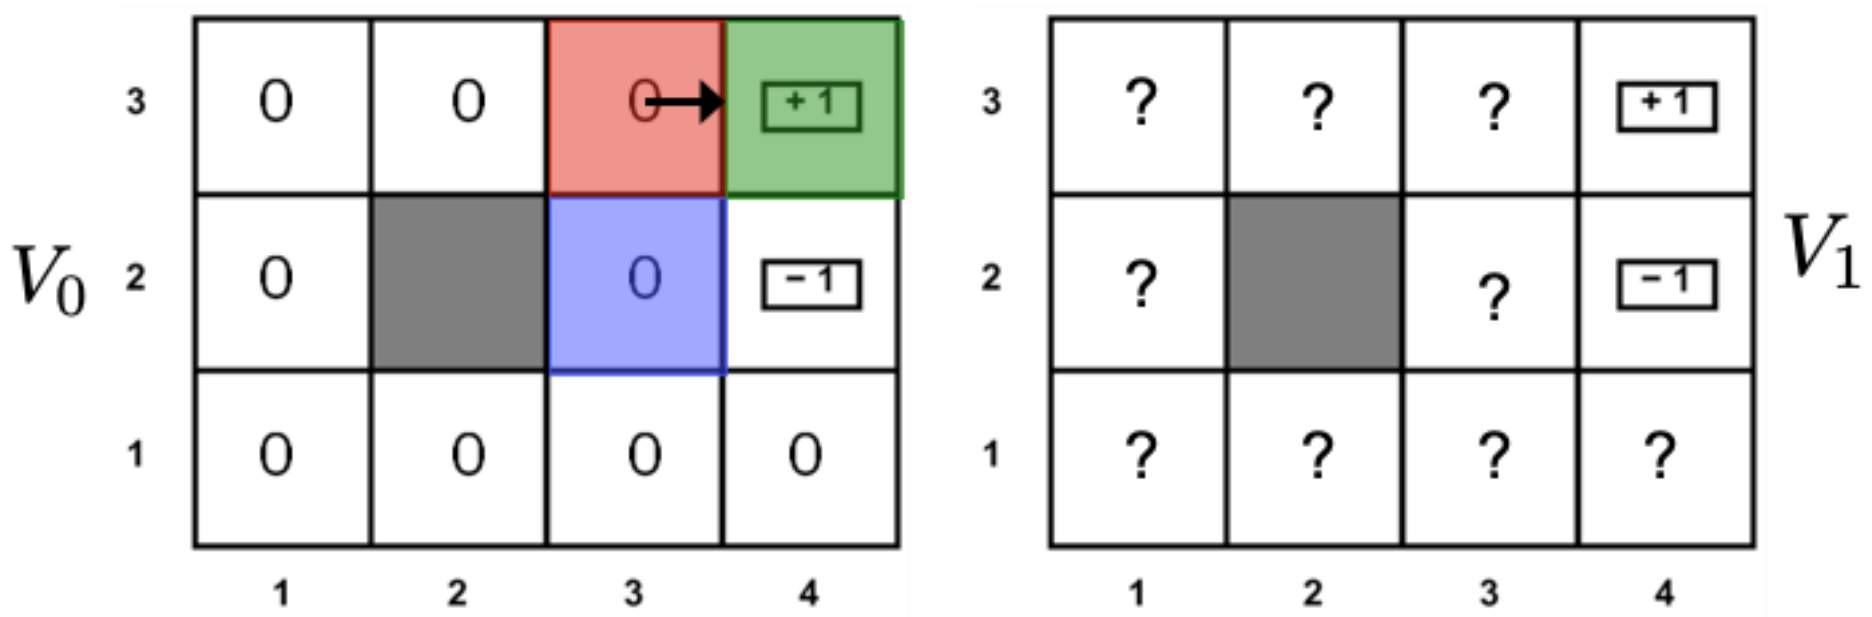

$$V_{i+1}(s) = \max_a \sum_{s'} T(s,a,s') \left[R(s,a,s') + \gamma V_i(s') \right] = \max_a Q_{i+1}(s,a)$$

$$Q_1(\langle 3,3 \rangle, \text{right}) = \sum_{s'} T(\langle 3,3 \rangle, \text{right}, s') \left[R(\langle 3,3 \rangle, \text{right}, s') + \gamma V_i(s') \right]$$

$$= 0.8 * [0.0 + 0.9 * 1.0] + 0.1 * [0.0 + 0.9 * 0.0] + 0.1 * [0.0 + 0.9 * 0.0]$$

Example: Value Iteration

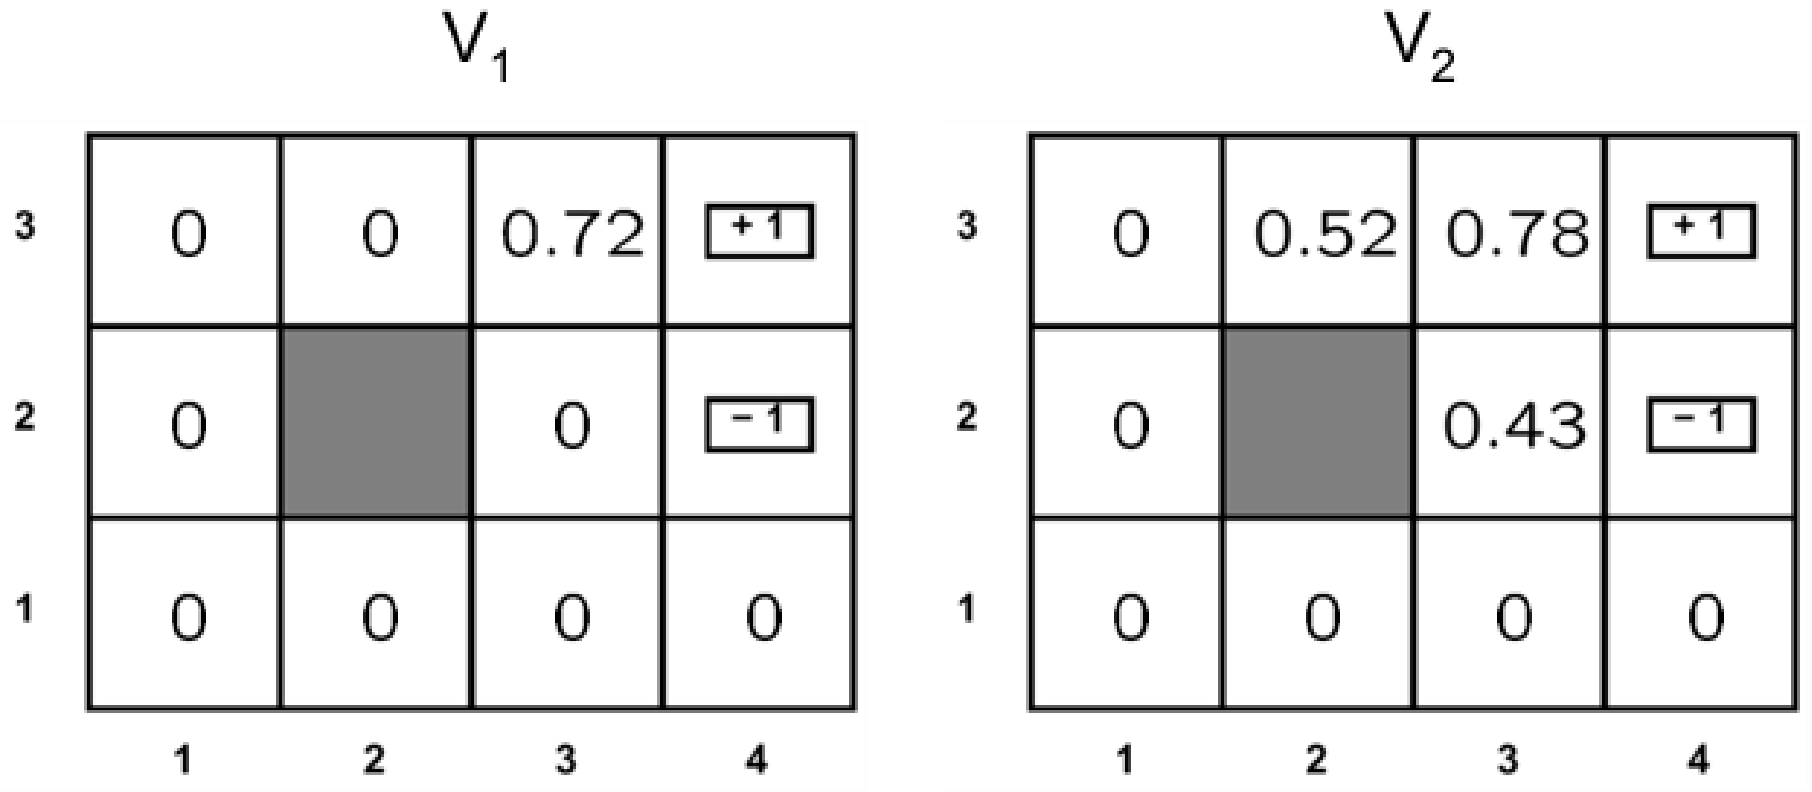

- Information propagates outward from terminal states and eventually all states have correct value estimates

k=3

Noise = 0.2
Discount = 0.9
Living reward = 0

k=4

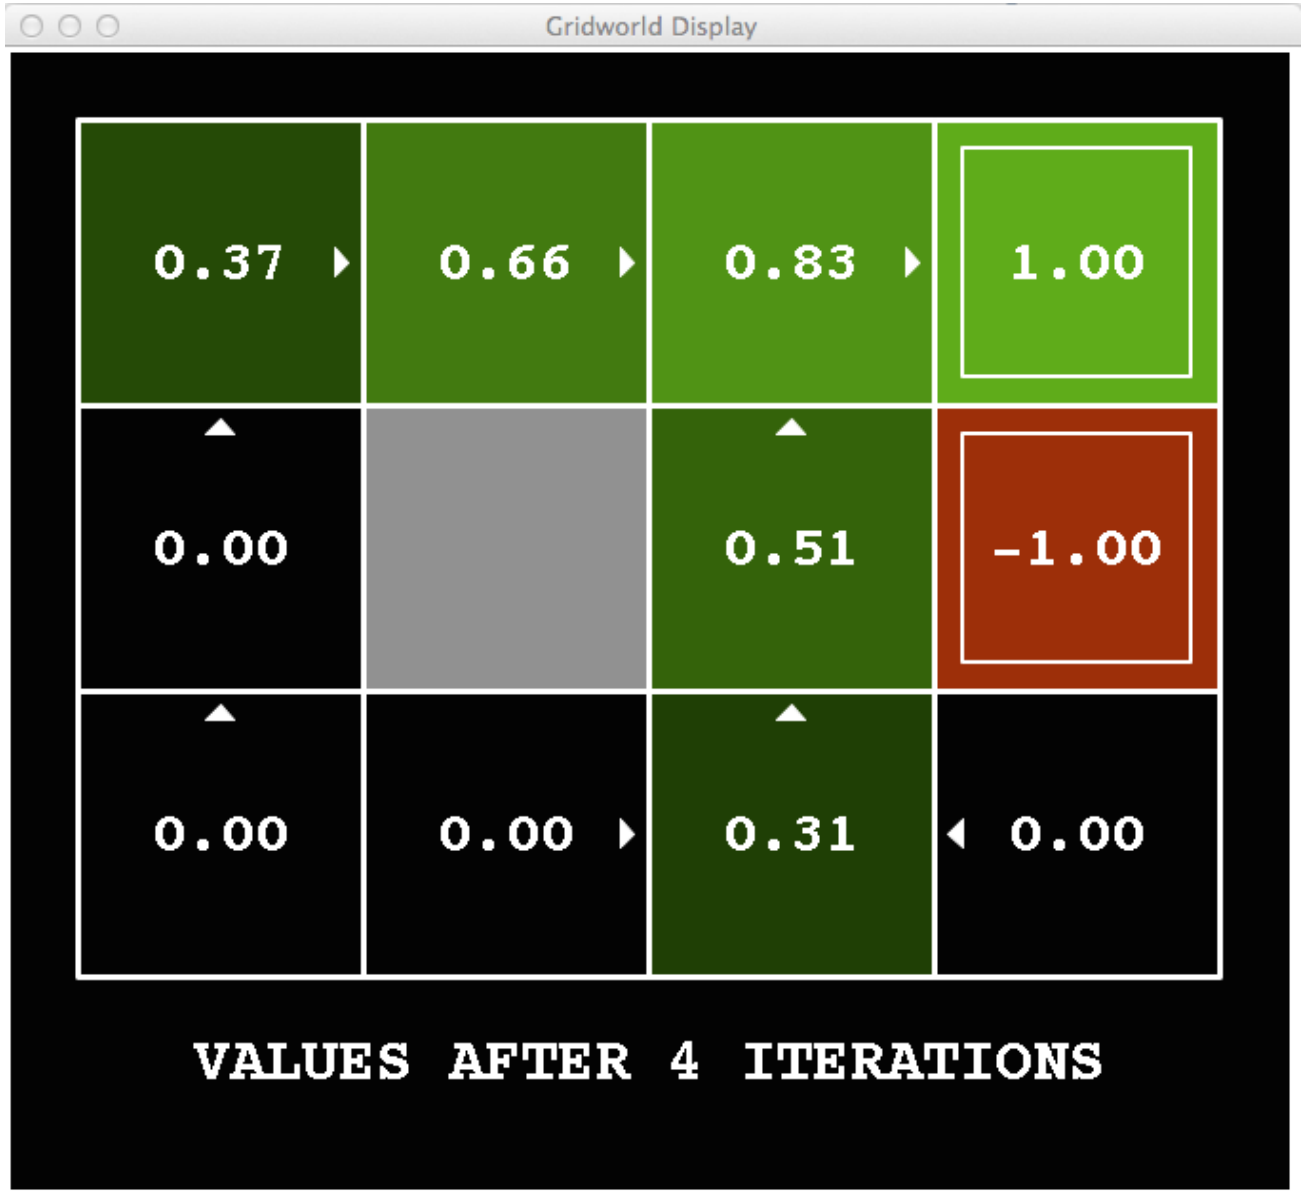

Noise = 0.2
Discount = 0.9
Living reward = 0

k=5

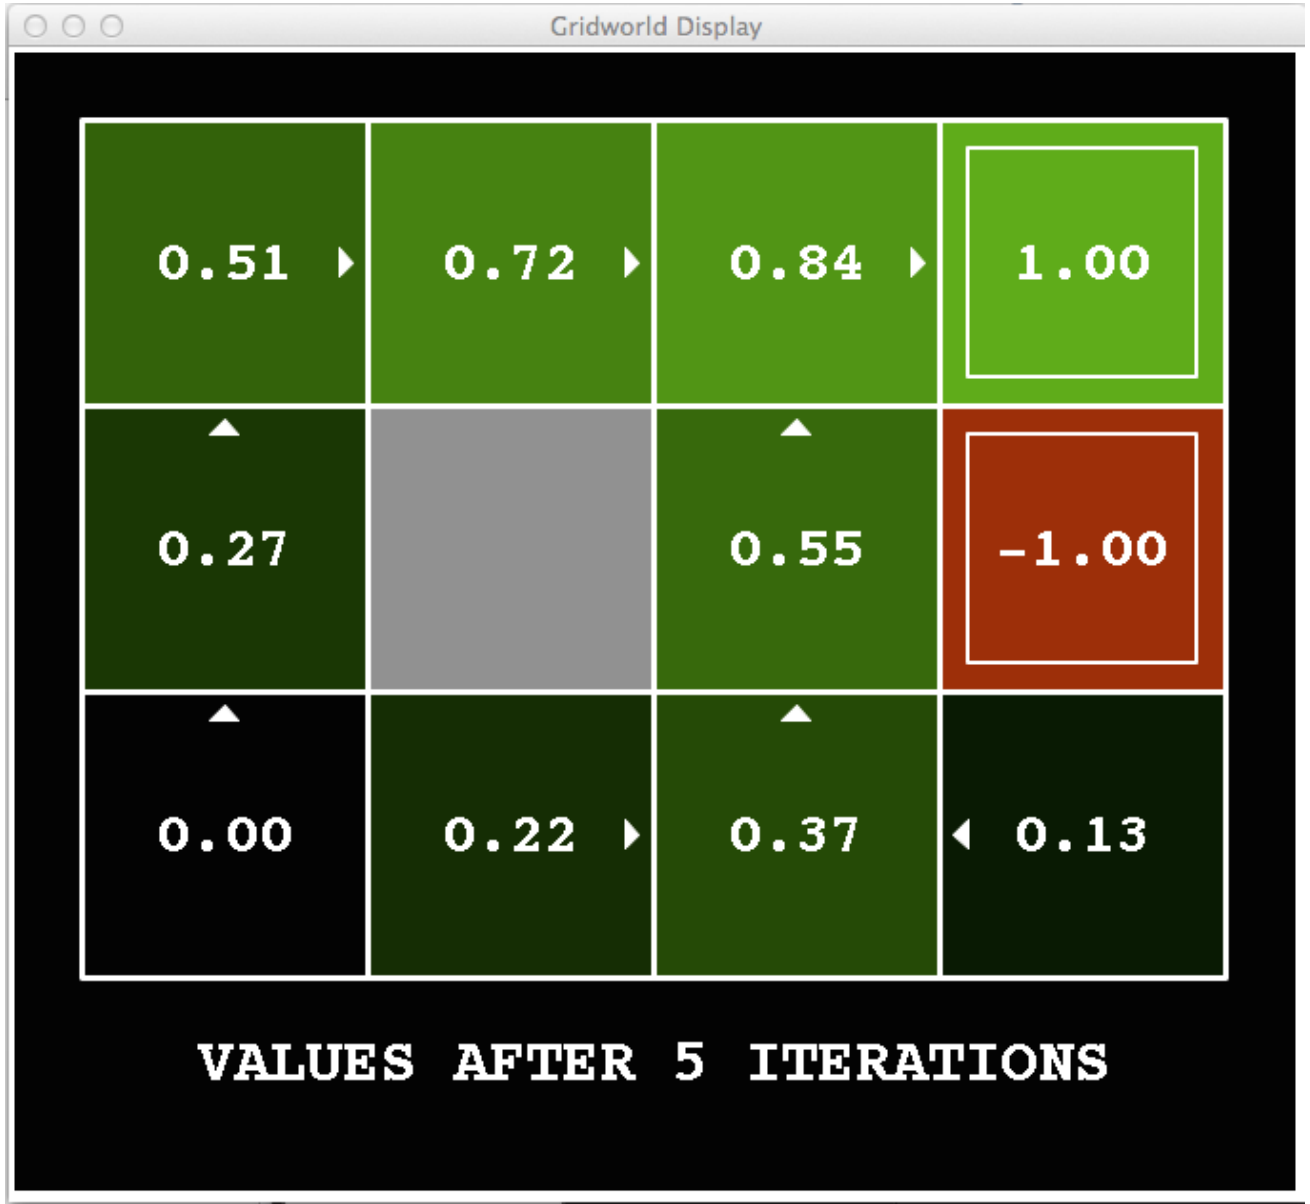

Noise = 0.2
Discount = 0.9
Living reward = 0

k=6

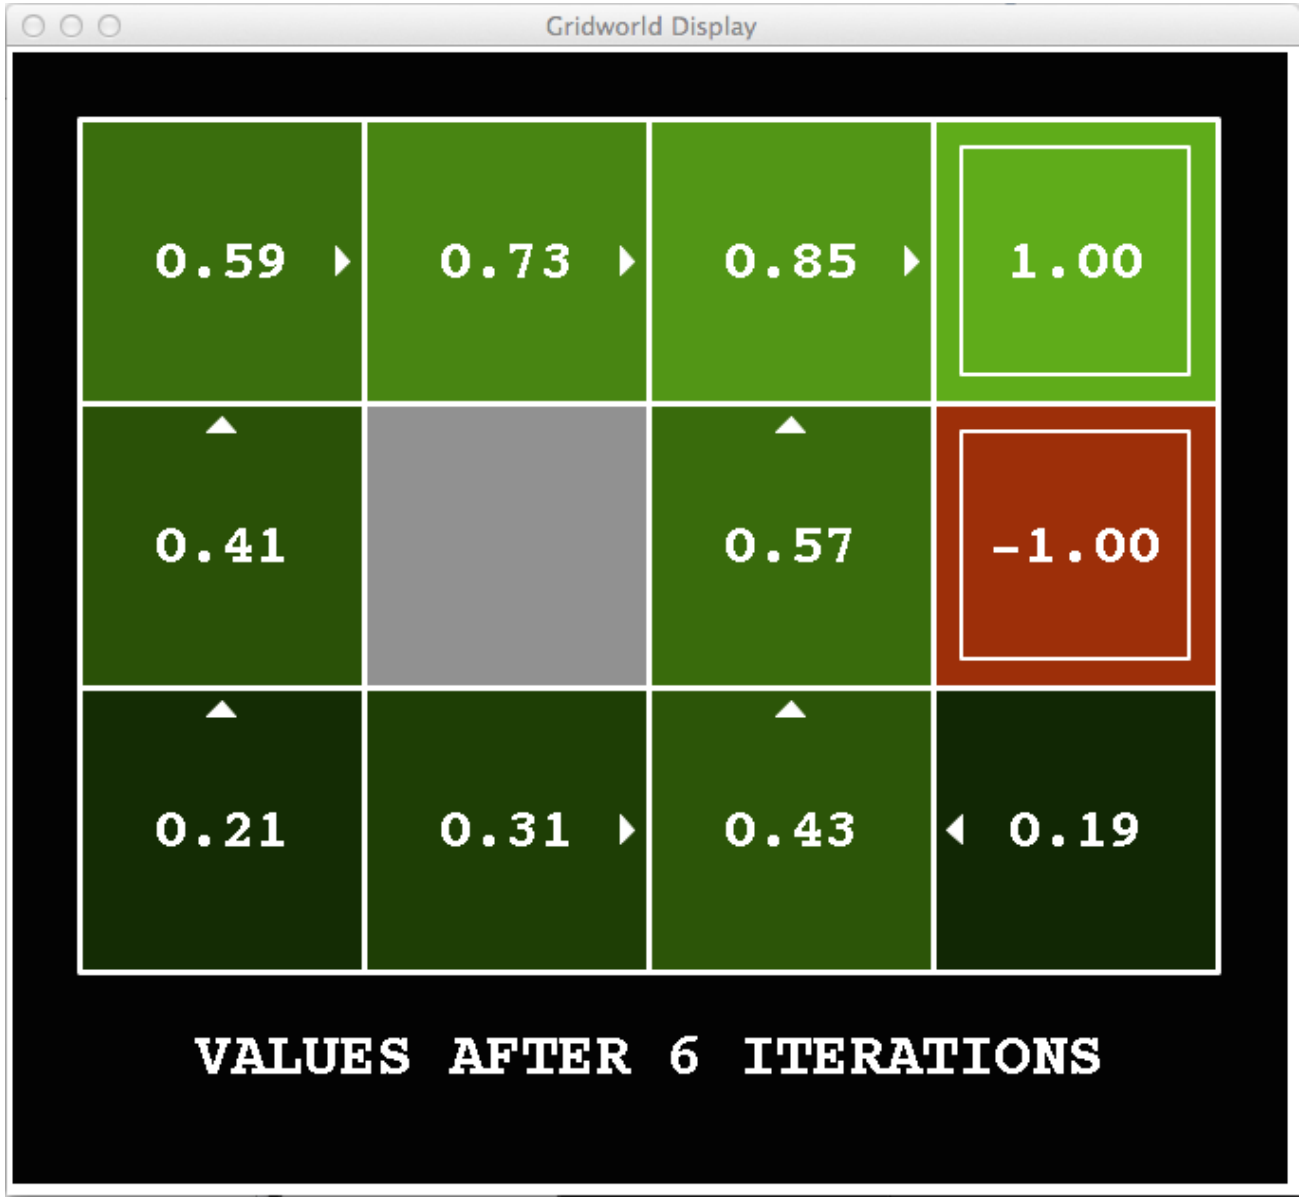

Noise = 0.2
Discount = 0.9
Living reward = 0

k=7

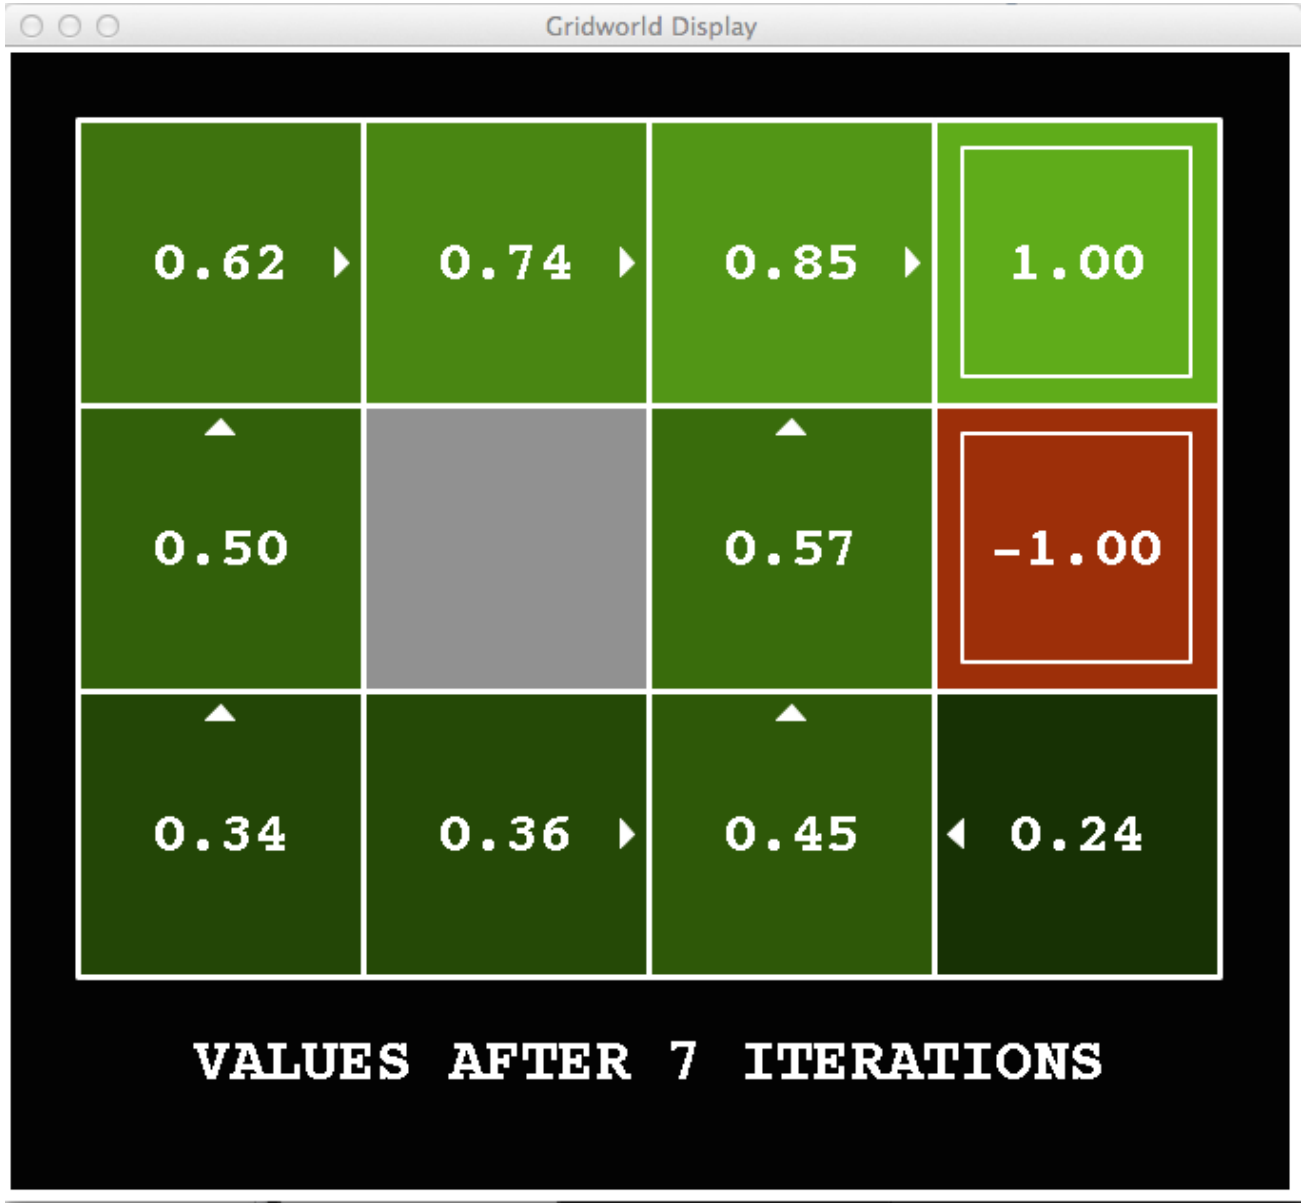

Noise = 0.2
Discount = 0.9
Living reward = 0

k=8

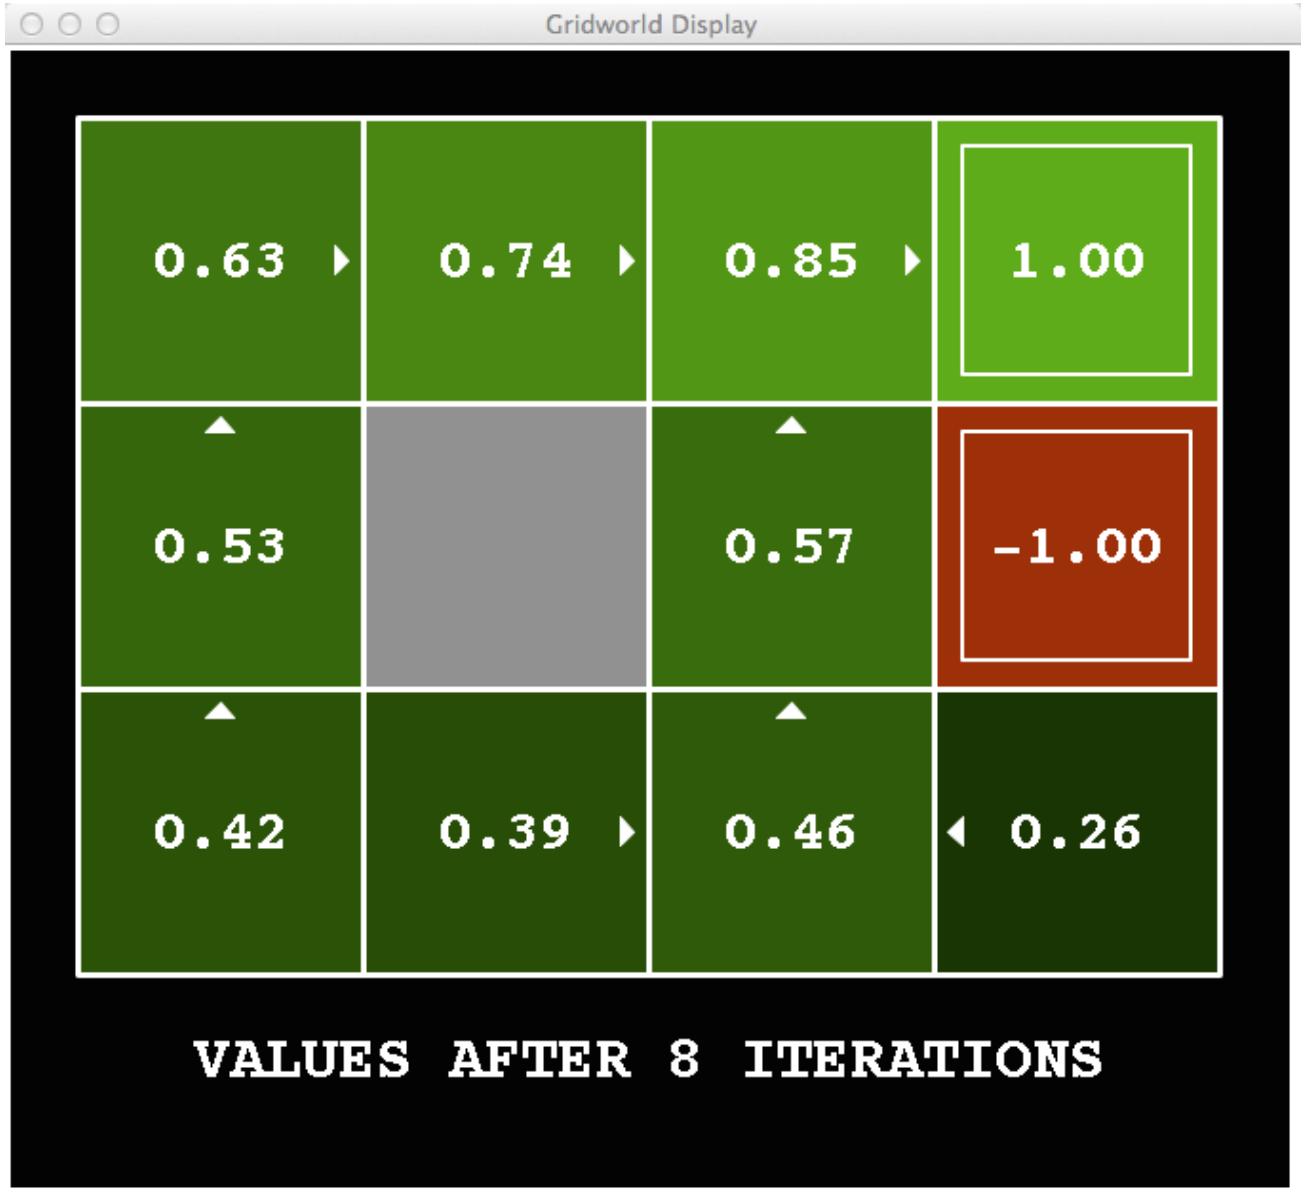

Noise = 0.2
Discount = 0.9
Living reward = 0

k=9

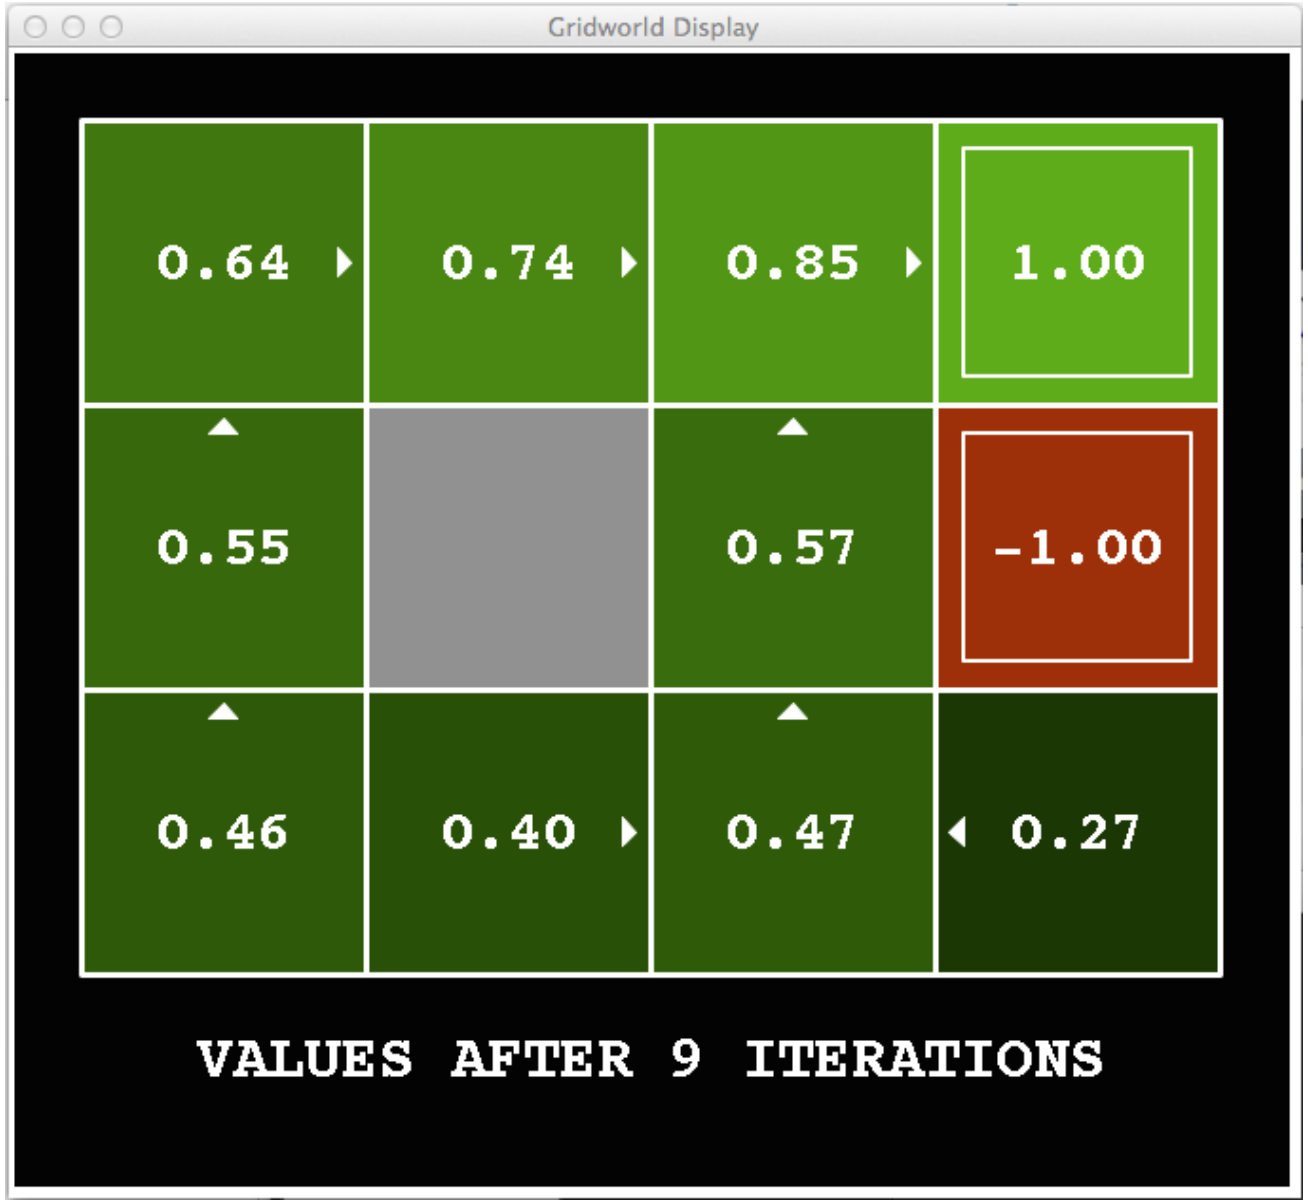

Noise = 0.2
Discount = 0.9
Living reward = 0

k=10

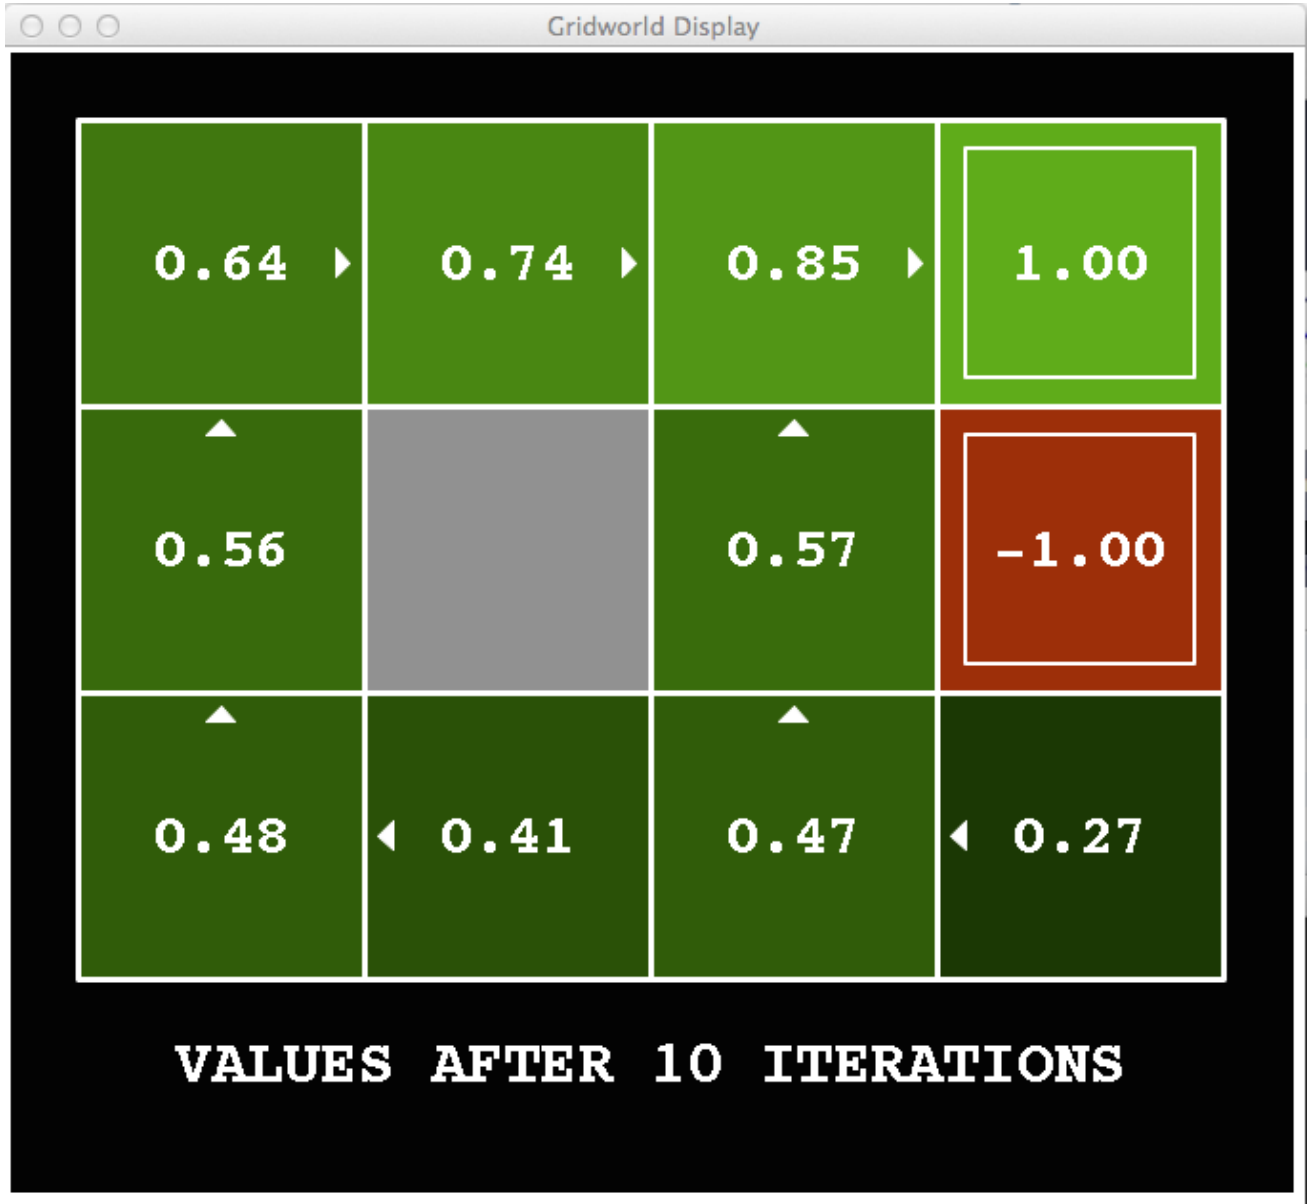

Noise = 0.2
Discount = 0.9
Living reward = 0

k=11

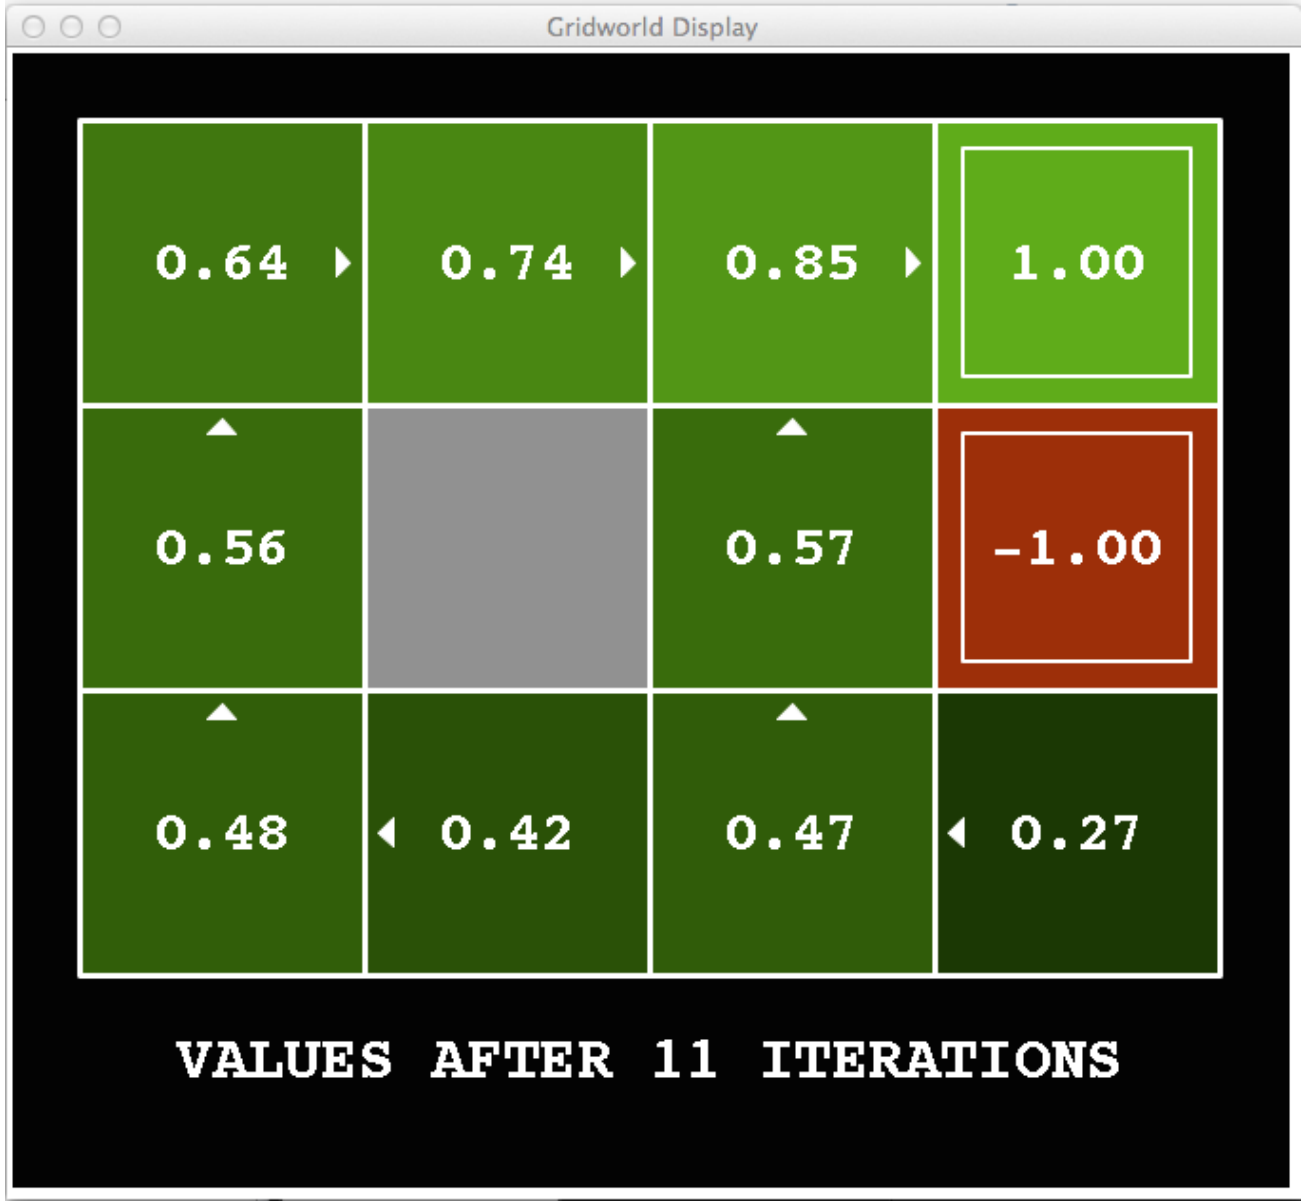

Noise = 0.2
Discount = 0.9
Living reward = 0

k=12

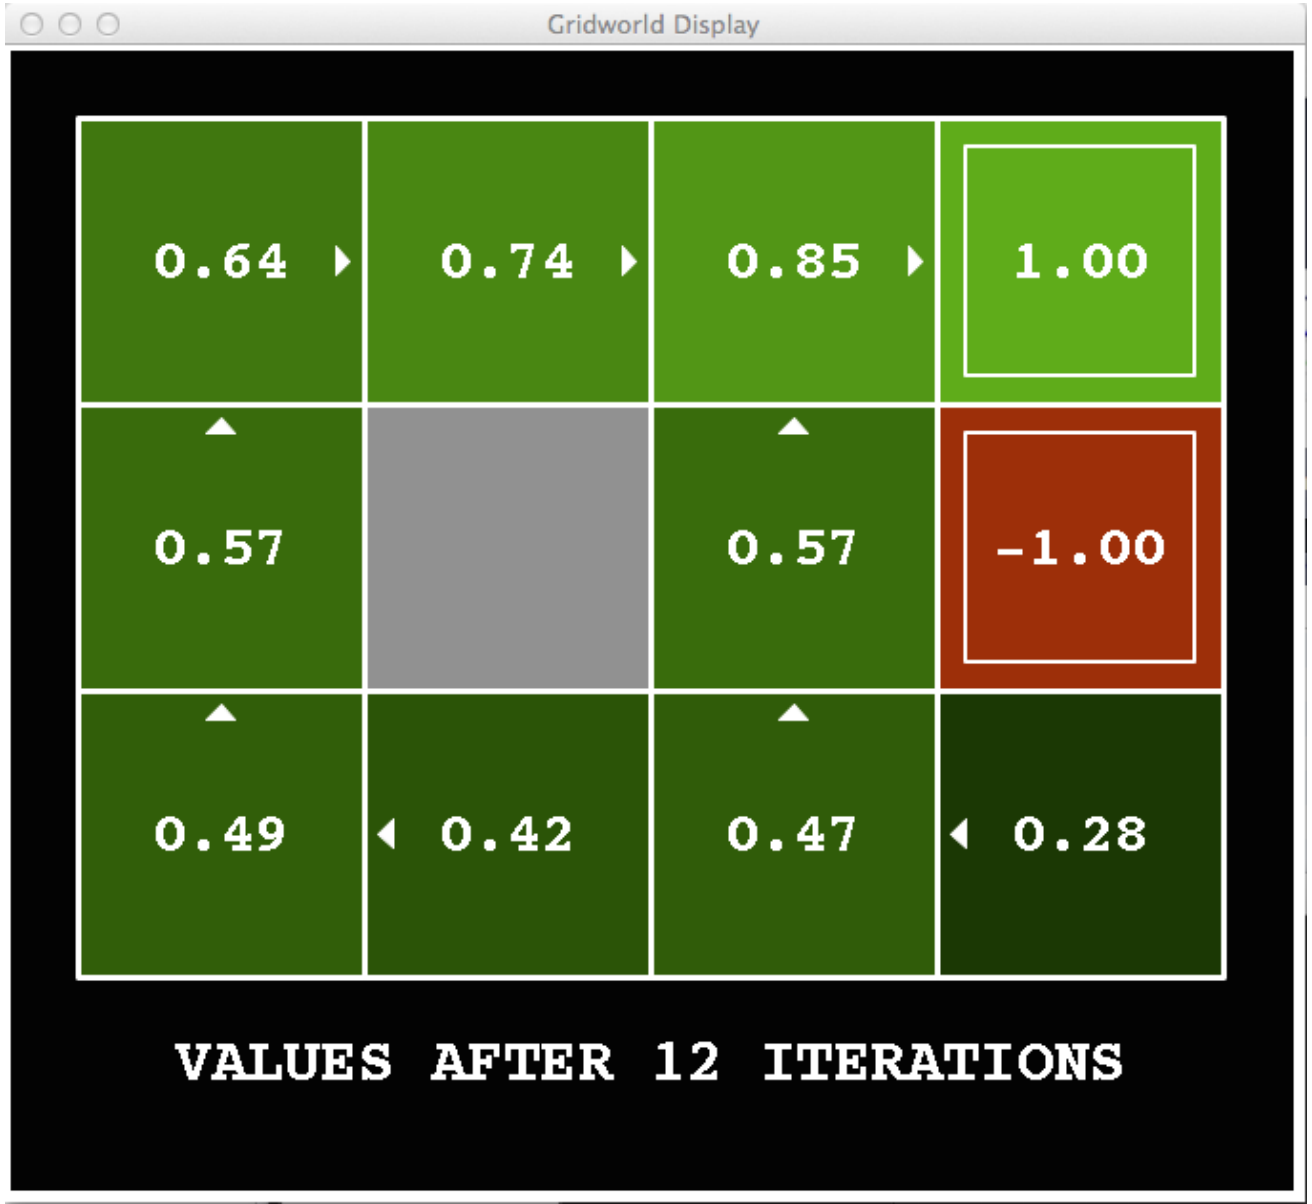

Noise = 0.2
Discount = 0.9
Living reward = 0

k=100

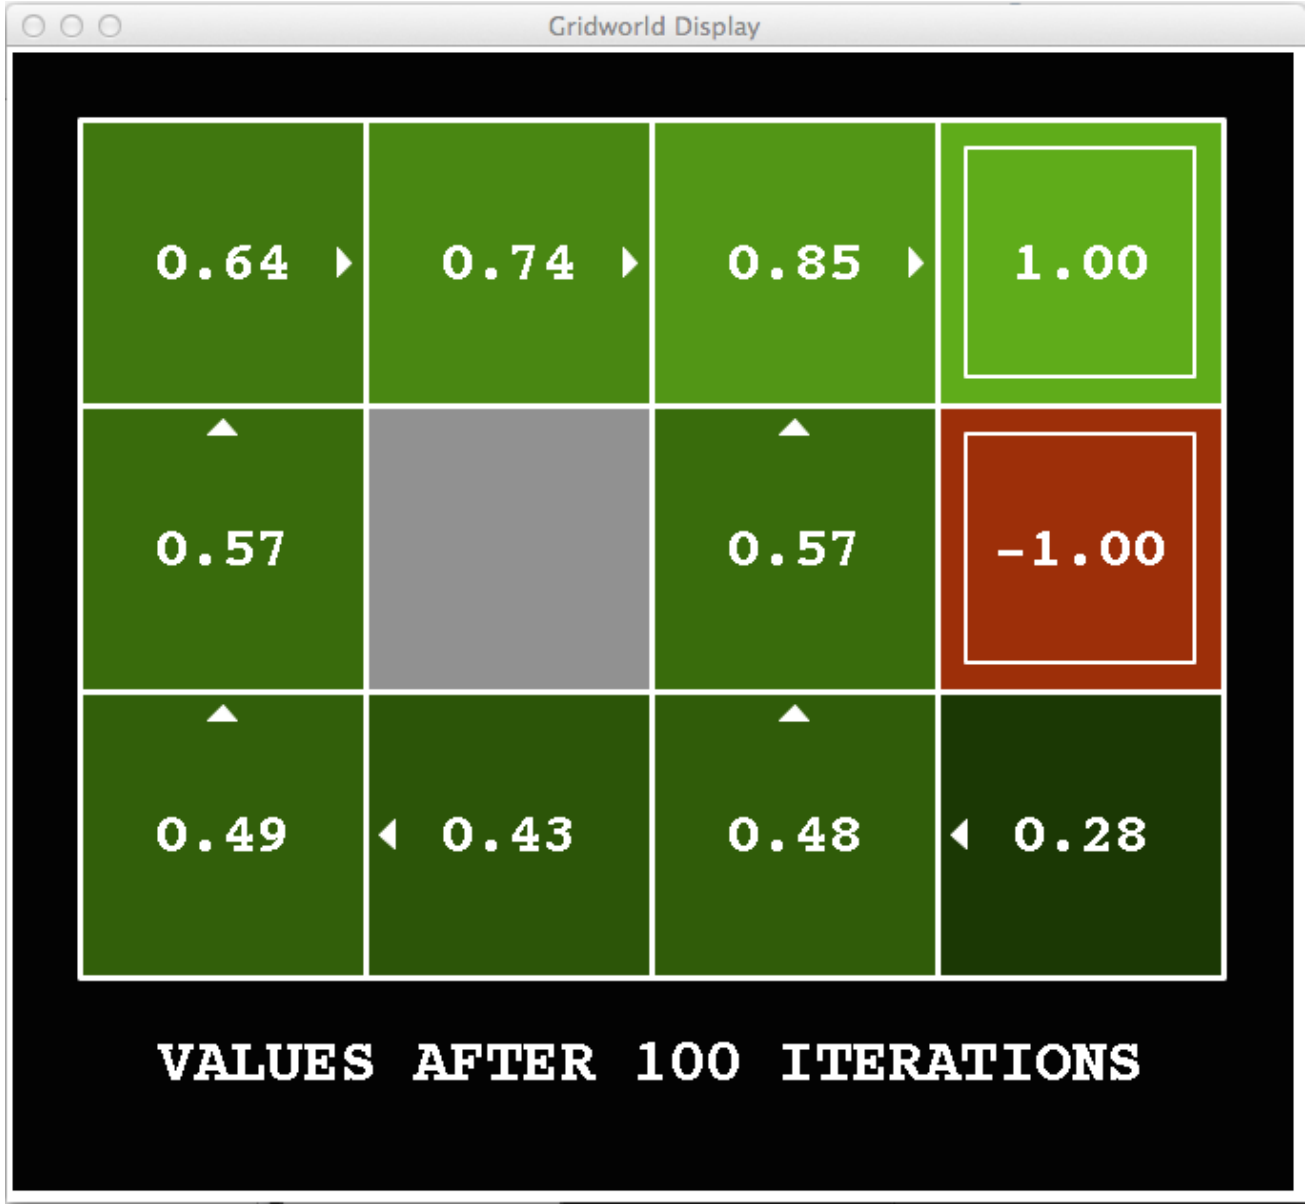

Noise = 0.2
Discount = 0.9
Living reward = 0