

Topics in Reinforcement Learning:  
AlphaZero, ChatGPT, Neuro-Dynamic Programming,  
Model Predictive Control, Discrete Optimization  
Arizona State University  
Course CSE 691, Spring 2024

Links to Class Notes, Videolectures, and Slides at  
<http://web.mit.edu/dimitrib/www/RLbook.html>


Dimitri P. Bertsekas  
dbertsek@asu.edu

Lecture 9

Combined Estimation/Control: Sequential Estimation, Bayesian Optimization, and  
Adaptive Control with a POMDP Approach  
Application to the Wordle Puzzle

- 1 Sequential Estimation of a Parameter Vector
- 2 Bayesian Optimization of Functions with Hard-to-Compute Values
- 3 Combined Estimation and Control - Adaptive Control
- 4 On-Line Solution of the Wordle Puzzle by Rollout

# Sequential Estimation of a Parameter Vector $\theta$



## Use of costly observations to estimate a parameter vector $\theta$

- The number and type of observations are subject to choice
- Instead, the outcomes of the observations obtained are evaluated on-line with a view towards stopping or modifying the observation process
- This involves sequential decision making, thus bringing DP to bear

## Example: Select one of two hypotheses using costly sequential observations

Given a new observation, we can accept one of the hypotheses or obtain a new observation at cost  $C$  (cf. quality control, the sequential probability ratio test, 1940s).

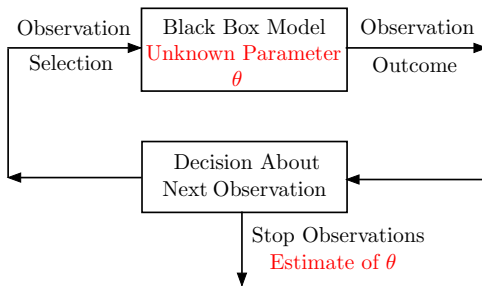
# Applications of Sequential Estimation

- Classical sequential **experiment design** problems or **sequential sampling strategies** in statistics.
  - ▶ Select one of multiple hypotheses.
  - ▶ Design of clinical trials or tests for medical diagnosis.
- Classical **sequential search problems** (e.g., search and rescue).
- **Route planning through a sensor network** for sequential information collection.
- **Sequential decoding** problems (e.g., the Mastermind and Wordle puzzles, to be discussed later).
- **Surrogate and Bayesian optimization** for minimizing “black box” functions (to be discussed first).

**An important distinction:** Does the current choice of observation affect the availability, the quality, or the cost of future observations?

- If no, we call this a **simple sequential estimation problem** (we will discuss it first in the context of Bayesian optimization).
- If yes, this can be viewed as a **combined estimation and control problem**, and can be viewed within the context of **adaptive control**.

# Surrogate Optimization of “Black Box” Functions




Minimize approximately a function whose values at given points are obtained only through time-consuming calculation, simulation, or experimentation

- Introduce a **parametric model of the cost function with parameter  $\theta$** .
- Observe sequentially the true cost function at a few observation points.
- Construct a model of the cost function (the **surrogate**) by estimating  $\theta$ .
- **Minimize the surrogate** to obtain a suboptimal solution.
- How to select observation points based on results of previous observations?
- **Exploration-exploitation tradeoff**: Observing at points likely to have near-optimal value vs observing at points in relatively unexplored areas of the search space.

- **Geostatistical interpolation** (“kriging” pioneered by the South African engineers Matheron and Krige in a goldmining context): Identify locations of high gold distribution based on samples from a few boreholes.
- **Design optimization**, e.g., aerodynamic design using hardware prototypes, materials design, drug development, etc.
- **Hyperparameter selection** of machine-learning models, including the architectural parameters of the deep neural network of AlphaZero.

# Bayesian Optimization of a Black Box Function $f$




- Introduce a parameter vector  $\theta = (\theta_1, \dots, \theta_m) \in \mathbb{R}^m$  where  $\theta_u = f(u)$ , i.e.,  $\theta$  is  $f$
- Observations are of the form  $z = f(u) + w$  (important special case is  $w = 0$ )
- Estimate  $\theta$  with  $N \ll m$  noisy measurements at chosen points  $u_1, \dots, u_N$
- We assume that  $\theta$  has a given a priori distribution  $b_0 = (b_{0,1}, \dots, b_{0,m})$  over  $\mathbb{R}^m$  (values of  $f$  at “neighboring” points should be correlated)
- After observations at points  $u_1, \dots, u_k$  of the form  $z_{u_i} = \theta_{u_i} + w_{u_i}$ , we choose the next point  $u_{k+1}$  at which to observe the value of  $f$ .
- Update the posterior distribution  $b_k$  with an estimator  $b_{k+1} = B_k(b_k, u_{k+1}, z_{u_{k+1}})$  ( $b_k$  is essentially the surrogate cost function after the  $k$ th observation)
- Gaussian case: If  $b_0$  and the noises  $w_u$  are Gaussian,  $b_k$  can be updated using closed form Gaussian process regression formulas.

# Illustration of the True Cost Function $f$ and its Surrogate

Black is the true cost function

Purple is the surrogate cost function



After 7 noise-free observations

The surrogate is specified by the posterior distribution  $b_k$  (mean and standard deviation at the different points are shown in the figure)



**Key Question:** How to select sequentially the observation point  $u_{k+1}$  given the observation results  $z_{u_1}, \dots, z_{u_k}$  from previously selected points  $u_1, \dots, u_k$

## A DP view

- **Introduce a POMDP model:** The posterior  $b_k$  (given the observations up to time  $k$ ) is the belief state,  $u_k$  is the control, the belief estimator  $b_{k+1} = B_k(b_k, u_{k+1}, z_{u_{k+1}})$  is the system. The cost function is based on the cost of the observations, and the “quality” of the surrogate obtained at the end.
- The dominant method in practice: **Use a greedy/myopic policy, based on an acquisition function.**
- The acquisition function  $A_k(b_k, u_{k+1})$  is a **heuristic measure of “benefit”** for selecting point  $u_{k+1}$  for observation when the belief state is  $b_k$ .
- Myopic policy: Selects the next point at which to observe,  $\hat{u}_{k+1}$ , as

$$\hat{u}_{k+1} \in \arg \max_{u_{k+1} \in \{1, \dots, m\}} A_k(b_k, u_{k+1})$$

- An alternative method: **Use rollout with a myopic base policy;** it has been advocated in several research works since 2016, with promising results.

# Examples of Acquisition Functions for Myopic Bayesian Optimization

The myopic policy maximizes over  $u_{k+1}$  the acquisition function  $A_k(b_k, u_{k+1})$ :

$$\hat{u}_{k+1} \in \arg \max_{u_{k+1} \in \{1, \dots, m\}} A_k(b_k, u_{k+1})$$

A common example of acquisition function: **Upper confidence bound**


$$A_k(b_k, u) = T_k(b_k, u) + \beta R_k(b_k, u), \quad \beta > 0 \text{ is a tunable parameter}$$

- Here  $T_k(b_k, u) = -\text{Mean of } f(u)$ , and  $R_k(b_k, u) = \text{Standard deviation of } f(u)$  (under the posterior distribution  $b_k$ ).
- $T_k(b_k, u)$  can be viewed as an **exploitation index** (encoding our desire to search within parts of the space where  $f$  takes low value), while  $R_k(b_k, u)$  can be viewed as an **exploration index** (encoding our desire to search within parts of the space that are relatively unexplored).


Another example of acquisition function: **Expected improvement**

$A_k(b_k, u)$  is the **expected value of the reduction of  $f(u)$  relative to the minimal value of  $f$  obtained up to time  $k$**  (under the posterior distribution  $b_k$ ).

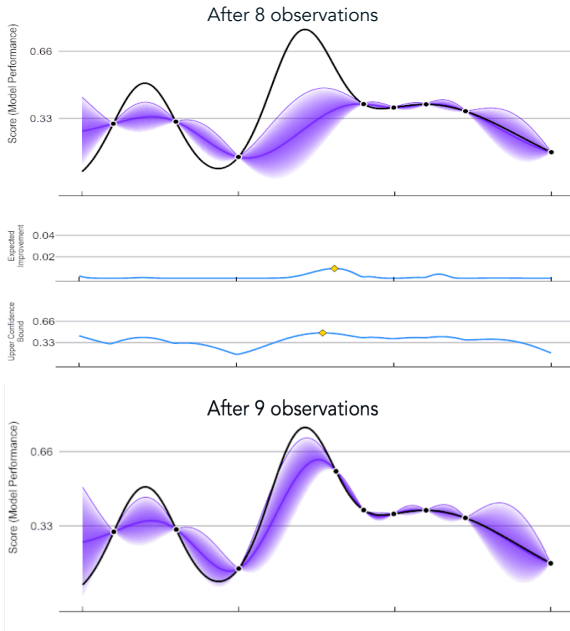
# Maximization Example (From Wikipedia Article on BO): True Function is Black, Surrogate Function is Purple; Observations are Noise-Free




# Maximization Example II



# Maximization Example III



# Maximization Example IV



# DP Algorithm for POMDP Formulation of Bayesian Optimization

$$J_k^*(b_k) = \min_{u_{k+1} \in \{1, \dots, m\}} \left[ c(u_{k+1}) + E_{z_{u_{k+1}}} \left\{ J_{k+1}^*(B_k(b_k, u_{k+1}, z_{u_{k+1}})) \mid b_k, u_{k+1} \right\} \right]$$

where  $c(u)$  is the cost of observation at  $u$ . Proceeds backwards from a terminal cost

$$J_N^*(b_N) = G(b_N) \quad (\text{measures the quality of the surrogate obtained at the end})$$

Approximation in value space (replace  $J_{k+1}^*$  with  $\tilde{J}_{k+1}$ )

$$\tilde{u}_{k+1} \in \arg \min_{u_{k+1} \in \{1, \dots, m\}} Q_k(b_k, u_{k+1})$$


where  $Q_k(b_k, u_{k+1})$  is the (approximate) Q-factor corresponding to the pair  $(b_k, u_{k+1})$ :

$$Q_k(b_k, u_{k+1}) = c(u_{k+1}) + E_{z_{u_{k+1}}} \left\{ \tilde{J}_{k+1}(B_k(b_k, u_{k+1}, z_{u_{k+1}})) \mid b_k, u_{k+1} \right\}$$

## Rollout


Use as  $\tilde{J}_{k+1}$  the cost function of a myopic base heuristic based on an acquisition function (or approximation thereof); first proposed by Lam, Wilcox, and Wolpert (2016), and followed up by others (promising, but relatively untested at present).

# Truncated Rollout with a Myopic Base Heuristic





# Adaptive Control with a POMDP Formulation and Rollout



Deterministic system  $x_{k+1} = f(x_k, \theta, u_k)$ ,  $\theta \in \{\theta^1, \dots, \theta^m\}$ : unknown parameter


- $\theta$  has known initial distribution  $b_0$  and stays constant. It is observed indirectly through perfect observation of  $x_k$
- View  $\theta$  as part of an augmented state  $(x_k, \theta)$  that is partially observed
- Bellman equation for optimal cost function  $J_k^*$ :

$$J_k^*(I_k) = \min_{u_k} \sum_{i=1}^m b_{k,i} (g(x_k, \theta^i, u_k) + J_{k+1}^*(I_k, u_k, f(x_k, \theta^i, u_k)))$$

where  $I_k = (x_0, \dots, x_k, u_0, \dots, u_{k-1})$  is the information state at time  $k$ , and  $b_{k,i} = P\{\theta = \theta^i \mid I_k\}$ ,  $i = 1, \dots, m$ , is the belief state (estimated on-line)

- Approximation in value space: Use approximation  $\tilde{J}^i(f(x_k, \theta^i, u_k))$  in place of  $J_{k+1}^*(I_k, u_k, f(x_k, \theta^i, u_k))$ . Minimize over  $u_k$  to obtain a one-step lookahead policy
- Example 1:  $\tilde{J}^i$  is the cost function of the optimal policy corresponding to  $\theta^i$
- Example 2:  $\tilde{J}^i$  is the cost function of a known policy assuming  $\theta = \theta^i$  (this is rollout)

# Rollout for Adaptive Control with a POMDP Formulation



At  $x_k$ , we minimize  $\hat{Q}_k(x_k, u_k)$ , the average Q-factor of  $u_k$ , defined by

$$\hat{Q}_k(x_k, u_k) = \sum_{i=1}^m b_{k,i} Q_k(x_k, u_k, \theta^i),$$

where  $Q_k(x_k, u_k, \theta^i)$  is the Q-factor computed assuming that  $\theta = \theta^i$

$$Q_k(x_k, u_k, \theta^i) = g_k(x_k, \theta^i, u_k) + J_{k+1, \pi^i}(f_k(x_k, \theta^i, u_k))$$

If  $\pi^i \equiv \pi$ , cost improvement over  $\pi$  can be proved

# A Rollout Approach for Solving On-Line the Wordle Puzzle (Joint Work with Siddhant Bhambri and Amrita Bhattacharjee)

## Overview

- There is a hidden mystery word/code word  $\theta$  drawn from an initial mystery list according to a known distribution. **In the standard version of the puzzle this distribution is uniform.**
- The mystery list shrinks as a result of guesses/observations.
- The guesses are chosen based on feedback about the mystery word provided by the preceding guesses.
- The puzzle is solved when the mystery list shrinks to a single element.
- We want to minimize the expected number of guesses to solve the puzzle.
- Important fact: **The belief distribution over the current mystery list remains uniform through the solution process.**
- This makes possible the solution by exact DP, with days of computation (Selby 2022).
- **Without the uniform initial belief distribution assumption (and/or small variations in the problem structure), the exact solution would be impossible.**
- **Rollout can solve near optimally the puzzle (and its variations) on-line much faster.**

# Playing Wordle Using an Online Rollout Algorithm for Deterministic POMDPs



Siddhant Bhambri






Amrita Bhattacharjee






Dimitri Bertsekas

# The Wordle Puzzle



Easy mode



Hard mode

## How To Play

Guess the Wordle in 6 tries.

- Each guess must be a valid 5-letter word.
- The color of the tiles will change to show how close your guess was to the word.

### Examples

WEARY

W is in the word and in the correct spot.

PILLS

I is in the word but in the wrong spot.

VAGUE

U is not in the word in any spot.



[Log in or create a free NYT account](#) to link your stats.

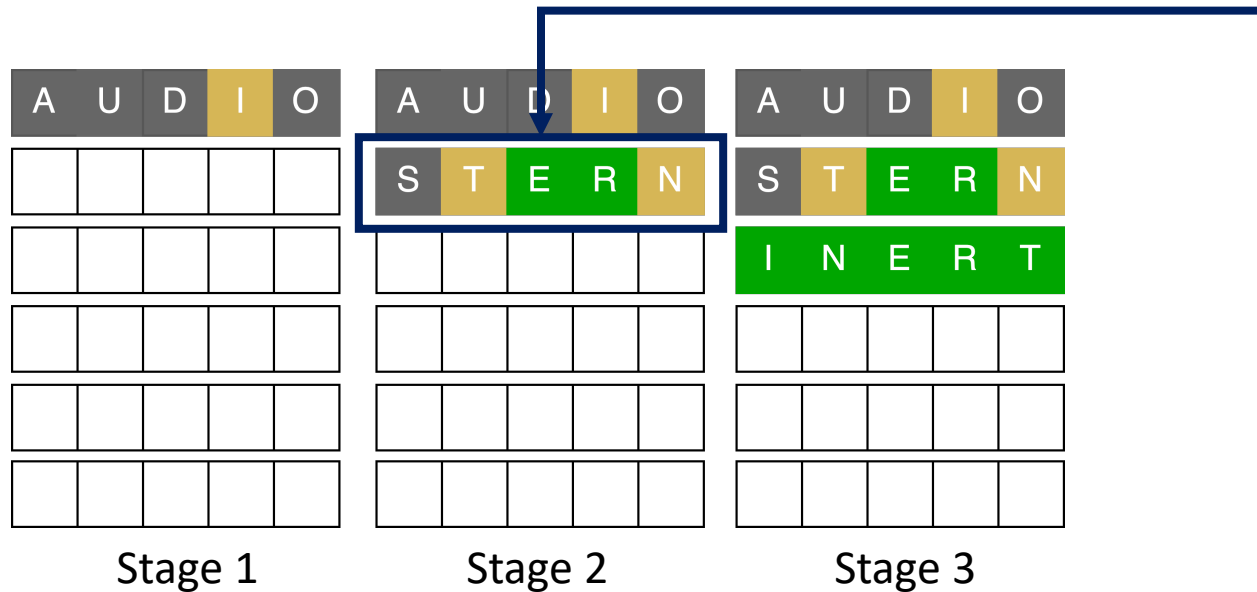
A new puzzle is released daily at midnight. If you haven't already, you can [sign up](#) for our daily reminder email.

Have feedback? Email us at [nytgames@nytimes.com](mailto:nytgames@nytimes.com).

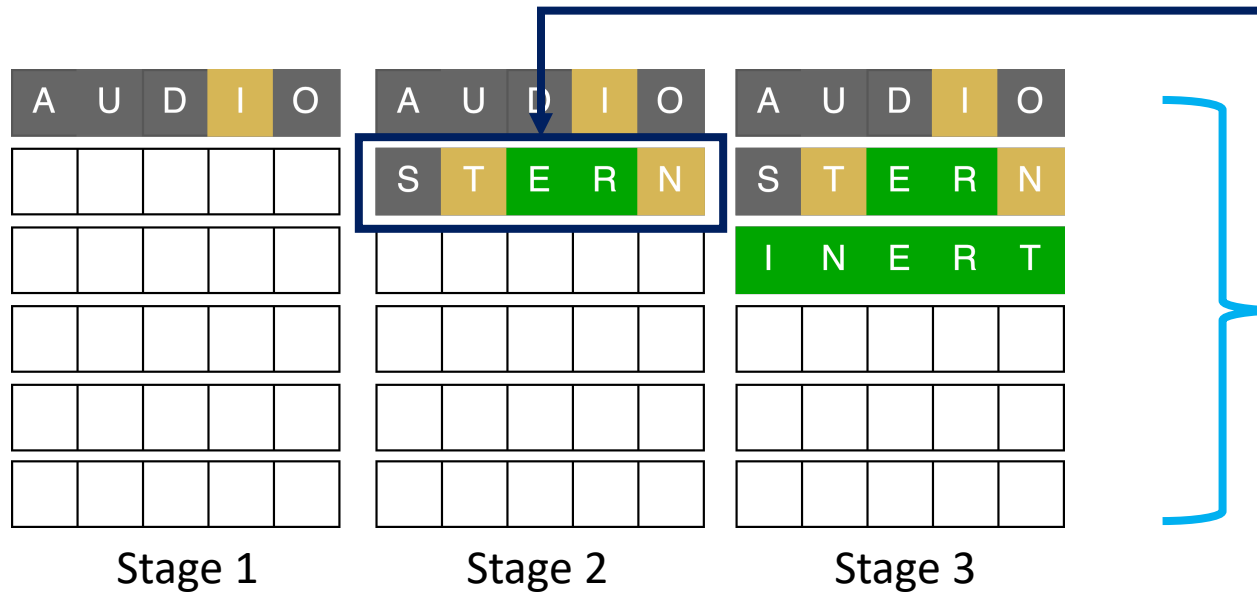
# Wordle as a POMDP

**States (S):** Subset of the initial mystery list of 2,315 words

**Actions (A):** Set of 12,972 guess words



# Wordle as a POMDP

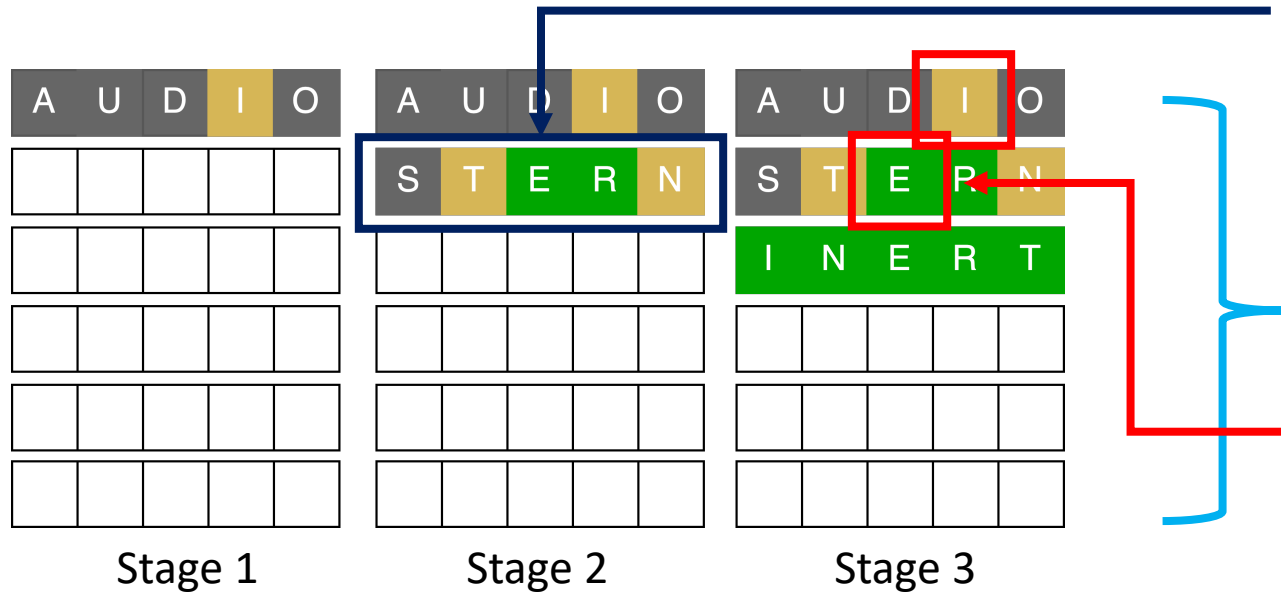


**States (S):** Subset of the initial mystery list of 2,315 words

**Actions (A):** Set of 12,972 guess words

**Transitions (T):** probability of going from one mystery word list to the next.

# Wordle as a POMDP



**States (S):** Subset of the initial mystery list of 2,315 words

**Actions (A):** Set of 12,972 guess words

**Transitions (T):** probability of going from one mystery word list to the next.

**Cost (C):** cost of utilizing a guess word (=1)

Observations from the game: colored observations for each letter



# Optimal Solution using Dynamic Programming

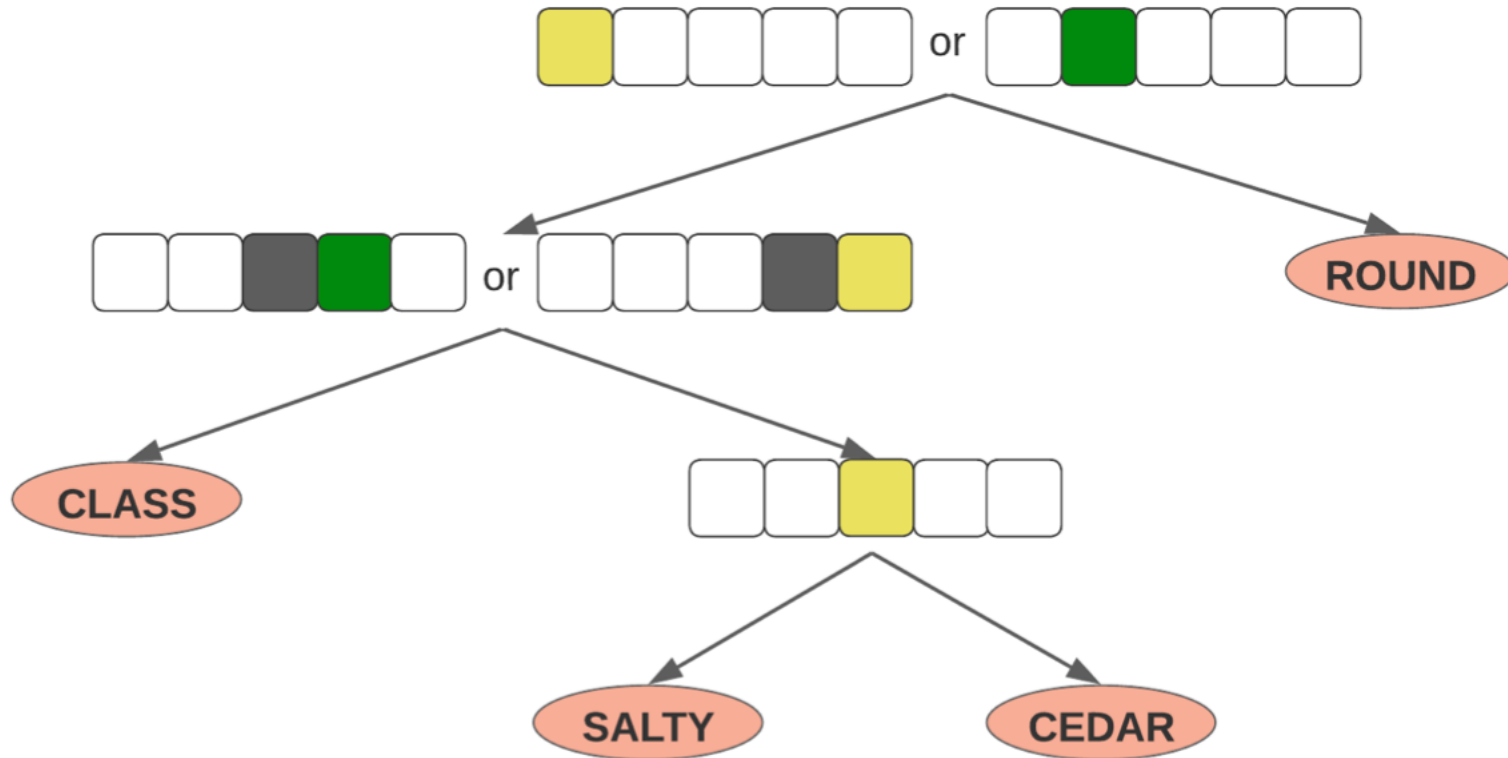
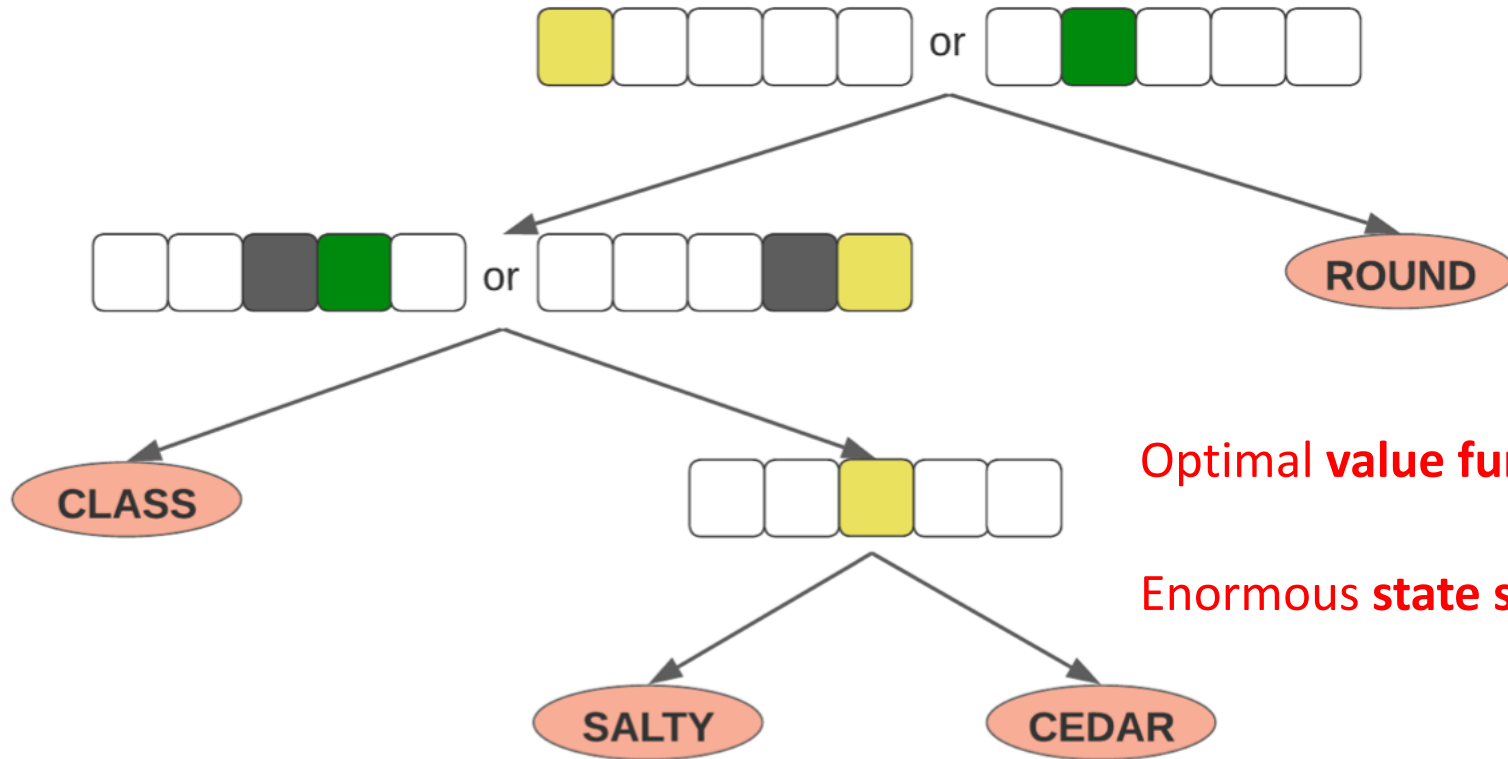


Figure adapted from *Bertsimas, D. and Paskov, A., 2022. An Exact and Interpretable Solution to Wordle.*  
Selby A., 2022. "The best strategies for Wordle".

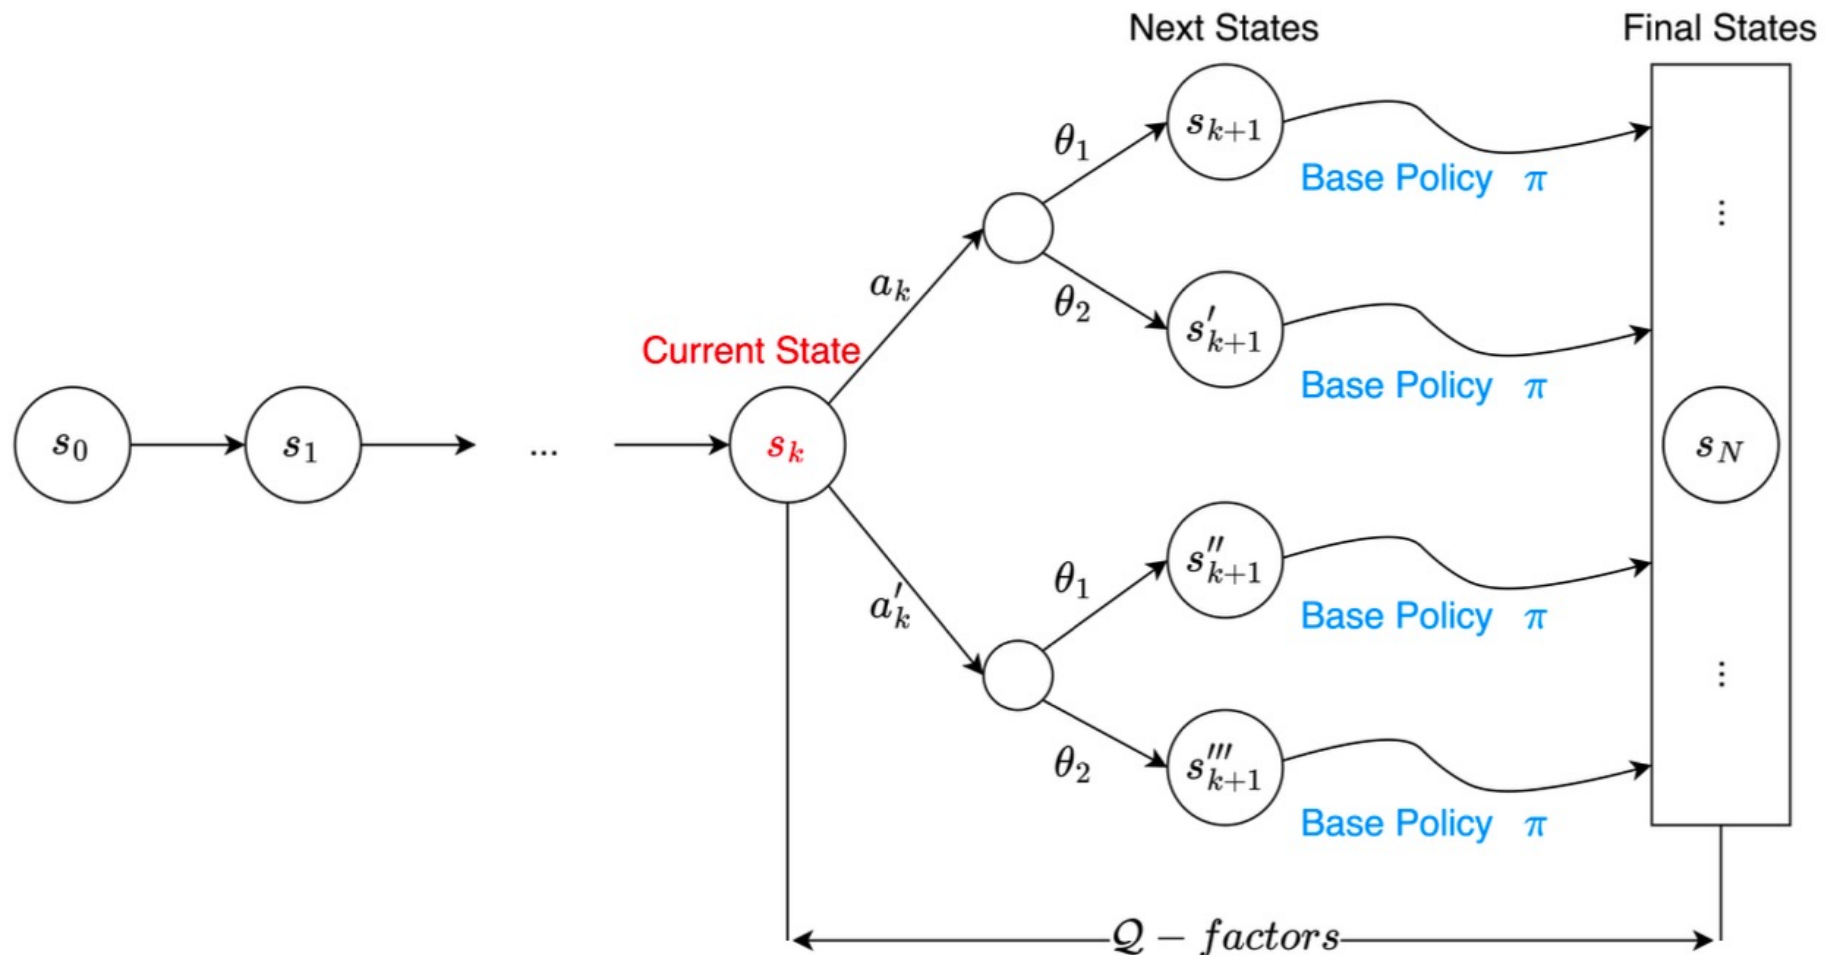
# Optimal Solution using Dynamic Programming



Optimal **value function** required to compute!

Enormous **state space**:  $2^{2315} \cong 10^{697}$

# Solution? - Approximate the Value Function!



Value Function --> Rollout Cost function: A Newton Step to solve the Bellman Eq.

# A Reinforcement Learning Approach Towards POMDP and Adaptive Control

Let us denote by  $\mathbf{x}_k$  and  $\mathbf{u}_k$  the state and control of the system at time  $k$ , respectively.

Let us also denote by  $\theta$  the unknown system parameter. We assume that  $\theta$  stays fixed over time, at one of  $m$  known values  $\theta_1, \dots, \theta_m$ :

$$\theta \in \{\theta^1, \dots, \theta^m\}.$$

The state  $\mathbf{x}_k$  is assumed to be perfectly observed by the controller at each time  $k$ , and evolves according to a system equation

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \theta, \mathbf{u}_k),$$

Finally, we assume that the following information vector,

$$I_k = \{\mathbf{x}_0, \dots, \mathbf{x}_k, \mathbf{u}_0, \dots, \mathbf{u}_{k-1}\},$$

is available at time  $k$ , and is used to compute the conditional probabilities

$$b_{k,i} = P\{\theta = \theta^i \mid I_k\}, \quad i = 1, \dots, m.$$

# A Reinforcement Learning Approach Towards POMDP and Adaptive Control

These probabilities form a vector

$$b_k = (b_{k,1}, \dots, b_{k,m}),$$

Following the choice of  $u_k$ , a **cost**  $g(x_k, \theta, u_k)$  is incurred, and we wish to choose controls to minimize the sum of the incurred costs over a given number of stages  $N$ .

Note,  $b_k$  can be updated according to an equation of the form

$$b_{k+1} = B_k(x_k, b_k, u_k, x_{k+1}),$$

where  $B_k$  is an appropriate function, which can be viewed as a **recursive estimator of  $\theta$** .

# A Reinforcement Learning Approach Towards Wordle

- We view the **mystery word  $\theta$**  as the unknown system parameter,
- and we view the list of the **mystery words as the set  $\{\theta_i \mid i = 1, \dots, m\}$**  of possible values for  $\theta$ .
- The **initial distribution of  $\theta$  is uniform over the list of the mystery words**, as is the case in the New York Times version of the puzzle. It can then be shown that the **belief distribution  $b_k$  at stage  $k$  continues to be uniform over the list of eligible mystery words** (those that have not been excluded by the preceding word guesses). This is an important simplification, which **obviates the need for the estimator**.
- An important consequence is that we may use as state  $x_k$  the list of eligible mystery words at stage  $k$ , which evolves according to a deterministic system **equation  $x_{k+1} = f(x_k, u_k)$ , with  $u_k$  being the guess word at stage  $k$** .

# The Exact DP Algorithm and its Approximation in Value Space

The algorithm operates in the space of information vectors  $I_k$ . In particular, we denote by  $J_k(I_k)$  the optimal cost starting at information vector  $I_k$  at time  $k$ . This vector evolves over a finite number of stages  $N$  according to the equation

$$I_{k+1} = (I_k, x_{k+1}, u_k) = (I_k, f(x_k, \theta, u_k), u_k), \quad k = 0, \dots, N - 1.$$

It admits a DP algorithm that takes the form

$$J_k^*(I_k) = \min_{u_k \in U(x_k)} E_{\theta} \left\{ g(x_k, \theta, u_k) + J_{k+1}^*(I_k, f(x_k, \theta, u_k), u_k) \mid I_k, u_k \right\},$$

for  $k = 0, \dots, N - 1$ , with

$$J_N^*(I_N) = g_N(x_N),$$

where we use  $E_{\theta}\{\cdot \mid I_k, u_k\}$  to denote expected value over  $\theta$ , conditioned on  $I_k$  and  $u_k$ .

# The Exact DP Algorithm, Approximation in Value Space & Rollout

We can rewrite this DP algorithm in terms of the conditional belief probabilities  $b_{k,i}$  as

$$J_k^*(I_k) = \min_{u_k \in U(x_k)} \sum_{i=1}^m b_{k,i} \left\{ g(x_k, \theta^i, u_k) + J_{k+1}^*(I_k, f(x_k, \theta^i, u_k), u_k) \right\}.$$

The control applied by the optimal policy is given by

$$u_k^* \in \arg \min_{u_k \in U(x_k)} \sum_{i=1}^m b_{k,i} \left\{ g(x_k, \theta^i, u_k) + J_{k+1}^*(I_k, f(x_k, \theta^i, u_k), u_k) \right\}.$$

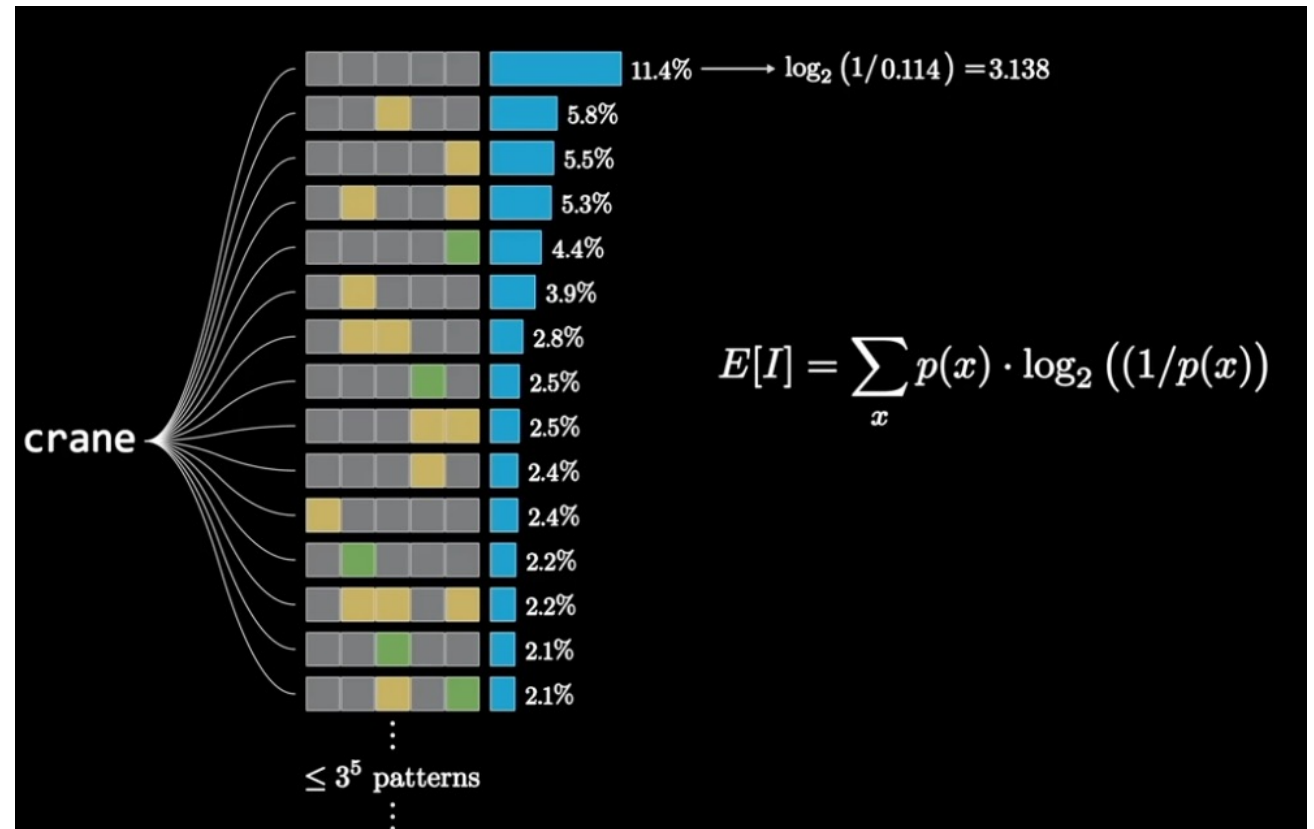
The corresponding approximation in value space scheme with one-step lookahead minimization is given by

$$\tilde{u}_k \in \arg \min_{u_k \in U(x_k)} \sum_{i=1}^m b_{k,i} \left\{ g(x_k, \theta^i, u_k) + \tilde{J}_{k+1}^i(f(x_k, \theta^i, u_k)) \right\};$$



# Base Heuristic for Wordle – Information Gain!

**Information gain –**  
calculating entropy of the distribution  
(roughly based on how much  
using a word reduces the uncertainty  
about the mystery word)



# Solving Wordle Using Rollout

Line 1: empty set to store the average Q-factors for each possible action at stage k.

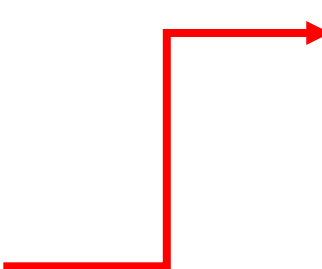
---

**Algorithm 1:** Rollout with One Step Look-ahead

---

**Data:** Current state  $s_k \in \mathcal{S}$ , Currently possible goal states  $\mathcal{G}_k \subseteq \mathcal{S}$ , Action space  $\mathcal{A}$ , Transition function  $\mathcal{T}$ , Cost function  $\mathcal{C}$ , Base Heuristic  $\mathcal{H}$ .

**Result:** Next state  $s_{k+1}$ .



```
1  $\hat{Q}_{factors} \leftarrow []$ ;
2 for  $a$  in  $\mathcal{A}$  do
3    $Cost\_total \leftarrow []$ ;
4   for  $g \in \mathcal{G}_k$  do
5      $s_{k+1} \leftarrow \mathcal{T}(s_k, a_k)$ ,  $cost \leftarrow 0$ ;
6     while  $s_{k+1} \neq g$  do
7        $s_{k+1} \leftarrow \arg \min_{a_k \in \mathcal{A}} \mathcal{H}(\mathcal{T}(s_k, a_k))$ ;
8        $cost \leftarrow cost + \mathcal{C}(s_k, a_k)$ 
9      $Cost\_total.append(cost)$ ;
10   $mean\_cost \leftarrow \frac{1}{|\mathcal{G}_k|} \sum(Cost\_total)$ ;
11   $\hat{Q}_{factors}.append(mean\_cost)$ ;
12  $\hat{a}_k \leftarrow \arg \min_{\forall a_k \in \mathcal{A}} \hat{Q}_{factors}$ ,  $s_{k+1} \leftarrow \mathcal{T}(s_k, \hat{a}_k)$ ;
13 return  $s_{k+1}$ 
```

---

# Solving Wordle Using Rollout

Line 4-8: for all possible  $g \in G_k$ , we perform the rollout by applying the next action as selected by our base heuristic cost function  $H$  and compute the Q-factor or cost until we reach the terminating state.

---

**Algorithm 1:** Rollout with One Step Look-ahead

---

**Data:** Current state  $s_k \in \mathcal{S}$ , Currently possible goal states  $\mathcal{G}_k \subseteq \mathcal{S}$ , Action space  $\mathcal{A}$ , Transition function  $\mathcal{T}$ , Cost function  $\mathcal{C}$ , Base Heuristic  $\mathcal{H}$ .

**Result:** Next state  $s_{k+1}$ .

```
1  $\hat{Q}_{factors} \leftarrow []$ ;  
2 for  $a$  in  $\mathcal{A}$  do  
3    $Cost\_total \leftarrow []$ ;  
4   for  $g \in \mathcal{G}_k$  do  
5      $s_{k+1} \leftarrow \mathcal{T}(s_k, a_k)$ ,  $cost \leftarrow 0$ ;  
6     while  $s_{k+1} \neq g$  do  
7        $s_{k+1} \leftarrow \arg \min_{a_k \in \mathcal{A}} \mathcal{H}(\mathcal{T}(s_k, a_k))$ ;  
8        $cost \leftarrow cost + \mathcal{C}(s_k, a_k)$   
9      $Cost\_total.append(cost)$ ;  
10     $mean\_cost \leftarrow \frac{1}{|\mathcal{G}_k|} \sum(Cost\_total)$ ;  
11     $\hat{Q}_{factors}.append(mean\_cost)$ ;  
12  $\hat{a}_k \leftarrow \arg \min_{\forall a_k \in \mathcal{A}} \hat{Q}_{factors}$ ,  $s_{k+1} \leftarrow \mathcal{T}(s_k, \hat{a}_k)$  ;  
13 return  $s_{k+1}$ 
```

---

# Solving Wordle Using Rollout

Line 10: find the average cost for solving the game for  $g$

---

**Algorithm 1:** Rollout with One Step Look-ahead

---

**Data:** Current state  $s_k \in \mathcal{S}$ , Currently possible goal states  $\mathcal{G}_k \subseteq \mathcal{S}$ , Action space  $\mathcal{A}$ , Transition function  $\mathcal{T}$ , Cost function  $\mathcal{C}$ , Base Heuristic  $\mathcal{H}$ .

**Result:** Next state  $s_{k+1}$ .

```
1  $\hat{Q\_factors} \leftarrow []$ ;  
2 for  $a$  in  $\mathcal{A}$  do  
3    $Cost\_total \leftarrow []$ ;  
4   for  $g \in \mathcal{G}_k$  do  
5      $s_{k+1} \leftarrow \mathcal{T}(s_k, a_k)$ ,  $cost \leftarrow 0$ ;  
6     while  $s_{k+1} \neq g$  do  
7        $s_{k+1} \leftarrow \arg \min_{a_k \in \mathcal{A}} \mathcal{H}(\mathcal{T}(s_k, a_k))$ ;  
8        $cost \leftarrow cost + \mathcal{C}(s_k, a_k)$   
9      $Cost\_total.append(cost)$ ;  
10     $mean\_cost \leftarrow \frac{1}{|\mathcal{G}_k|} \sum(Cost\_total)$ ;  
11     $\hat{Q\_factors}.append(mean\_cost)$ ;  
12  $\hat{a}_k \leftarrow \arg \min_{\forall a_k \in \mathcal{A}} \hat{Q\_factors}$ ,  $s_{k+1} \leftarrow \mathcal{T}(s_k, \hat{a}_k)$  ;  
13 return  $s_{k+1}$ 
```

---

# Solving Wordle Using Rollout

Line 12: we select the action  $\hat{a}_k$  that corresponds to the minimum average cost and apply it to state  $s_k$ .

---

**Algorithm 1:** Rollout with One Step Look-ahead

---

**Data:** Current state  $s_k \in \mathcal{S}$ , Currently possible goal states  $\mathcal{G}_k \subseteq \mathcal{S}$ , Action space  $\mathcal{A}$ , Transition function  $\mathcal{T}$ , Cost function  $\mathcal{C}$ , Base Heuristic  $\mathcal{H}$ .

**Result:** Next state  $s_{k+1}$ .

```
1  $\hat{Q\_factors} \leftarrow []$ ;  
2 for  $a$  in  $\mathcal{A}$  do  
3    $Cost\_total \leftarrow []$ ;  
4   for  $g \in \mathcal{G}_k$  do  
5      $s_{k+1} \leftarrow \mathcal{T}(s_k, a_k)$ ,  $cost \leftarrow 0$ ;  
6     while  $s_{k+1} \neq g$  do  
7        $s_{k+1} \leftarrow \arg \min_{a_k \in \mathcal{A}} \mathcal{H}(\mathcal{T}(s_k, a_k))$ ;  
8        $cost \leftarrow cost + \mathcal{C}(s_k, a_k)$   
9      $Cost\_total.append(cost)$ ;  
10     $mean\_cost \leftarrow \frac{1}{|\mathcal{G}_k|} \sum(Cost\_total)$ ;  
11     $\hat{Q\_factors}.append(mean\_cost)$ ;  
12  $\hat{a}_k \leftarrow \arg \min_{\forall a_k \in \mathcal{A}} \hat{Q\_factors}$ ,  $s_{k+1} \leftarrow \mathcal{T}(s_k, \hat{a}_k)$   
13 return  $s_{k+1}$ 
```

---


# Results for Rollout vs Optimal Scores

Opening Word	Easy Mode			Hard Mode		
	MIG as Base Heuristic	Rollout with MIG as Base Heuristic	Optimal Score	MIG as Base Heuristic	Rollout with MIG as Base Heuristic	Optimal Score
salet	3.6108	3.4345	3.4212	3.6078	3.5231	3.5084
reast	3.6	3.4462	3.4225	3.6181	3.53	3.5136
crate	3.6177	3.4414	3.4238	3.6289	3.5361	3.5175
trape	3.6319	3.4604	3.4454	3.6199	3.5356	3.5179
slane	3.6255	3.4444	3.4311	3.622	3.5378	3.5201
prate	3.6333	3.4535	3.4376	3.6173	3.5348	3.5210
crane	3.6091	3.4380	3.4255	3.6333	3.5374	3.5227
carle	3.6108	3.4419	3.4285	3.6384	3.5369	3.5261
train	3.6181	3.4622	3.4436	3.6216	3.5369	3.5248
clout	3.6955	3.5248	3.5097	3.7123	3.6125	3.5931


Table: Results using 'Maximum Information Gain' as base heuristic, and with rollout.



# Advantage of Rollout vs Only Base Heuristic



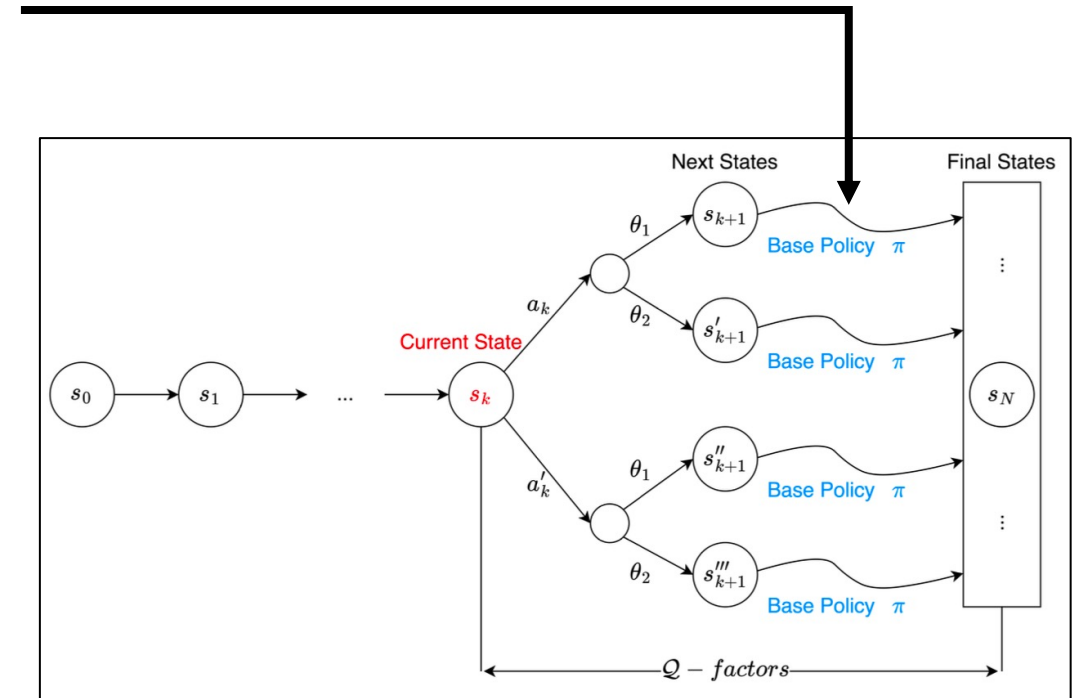
# Advantage of Rollout vs Only Base Heuristic





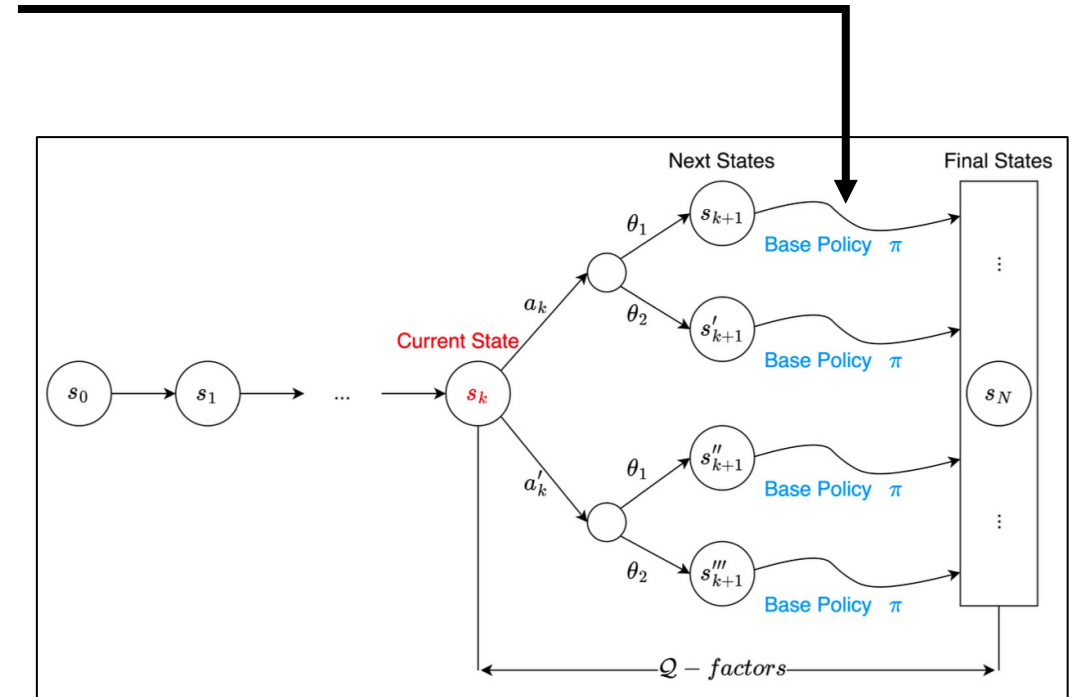
# Limitations of Rollout

- ❖ The need for a **reasonable base policy** – our experience with Wordle has been that the **rollout algorithm is relatively insensitive to the base policy** (e.g., the GEP heuristic in the paper).



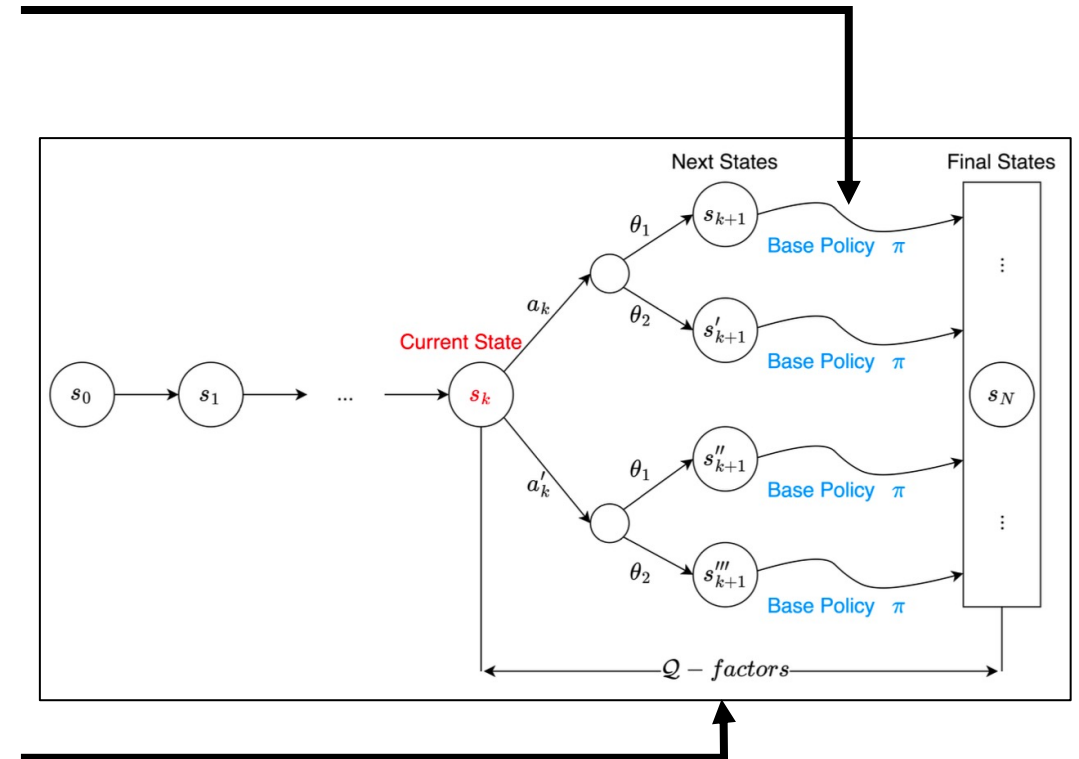
# Limitations of Rollout

- ❖ The need for a **reasonable base policy** – our experience with Wordle has been that the **rollout algorithm is relatively insensitive to the base policy** (e.g., the GEP heuristic).
- ❖ The need for a **posterior distribution estimator** - this is a limitation of most POMDP algorithms.



# Limitations of Rollout

- ❖ The need for a **reasonable base policy** – our experience with Wordle has been that the **rollout algorithm is relatively insensitive to the base policy** (e.g., the GEP heuristic).
- ❖ The need for a **posterior distribution estimator** - this is a limitation of most POMDP algorithms.
- ❖ The **number of Q-factors that need to be computed** by the algorithm online, particularly for a large action space - this difficulty may possibly be mitigated by **intelligently pruning the action space** or by **offline training using a neural network**.



# Summary

- ❖ We introduced a DP-based online rollout strategy as a computationally efficient solution to deterministic POMDPs with unknown parameters, whose exact solution is intractable.

# Summary

- ❖ We introduced a DP-based online rollout strategy as a computationally efficient solution to deterministic POMDPs with unknown parameters, whose exact solution is intractable.
- ❖ We demonstrated our approach using the challenging online puzzle Wordle, and empirically show that our approach provides near-optimal performance and impressive improvement over the heuristic approaches that have been used so far.

# Summary

- ❖ We introduced a DP-based online rollout strategy as a computationally efficient solution to deterministic POMDPs with unknown parameters, whose exact solution is intractable.
- ❖ We demonstrated our approach using the challenging online puzzle Wordle, and empirically show that our approach provides near-optimal performance and impressive improvement over the heuristic approaches that have been used so far.
- ❖ Through the Wordle computational demonstration, we identified the key obstacles in the way of solving other challenging POMDP problems that involve sequential estimation, possibly in conjunction with simultaneous adaptive control.

# Summary

- ❖ We introduced a **DP-based online rollout strategy** as a computationally efficient **solution to deterministic POMDPs with unknown parameters, whose exact solution is intractable.**
- ❖ We demonstrated our approach using the challenging **online puzzle Wordle**, and empirically show that our approach **provides near-optimal performance and impressive improvement over the heuristic approaches** that have been used so far.
- ❖ Through the Wordle computational demonstration, we **identified the key obstacles in the way of solving other challenging POMDP problems that involve sequential estimation**, possibly in conjunction with simultaneous adaptive control.

Bhambri, S., Bhattacharjee, A. and Bertsekas, D., 2023, August. Playing Wordle Using an Online Rollout Algorithm for Deterministic POMDPs. In *2023 IEEE Conference on Games (CoG)* (pp. 1-4). IEEE.



Access our paper:

<https://tinyurl.com/solving-wordle>



## Thank You!