

# Learning Fast Monomial Orders for Gröbner Basis Computations

R. Caleb Bunch<sup>1</sup>, Alperen A. Ergür<sup>2</sup>, Melika Golestani<sup>2</sup>, Jessie Tong<sup>3</sup>,  
Malia Walewski<sup>4</sup>, and Yunus E. Zeytuncu<sup>5</sup>

<sup>1</sup>Georgia Institute of Technology

<sup>2</sup>The University of Texas at San Antonio

<sup>3</sup>San Diego State University

<sup>4</sup>Emory University

<sup>5</sup>University of Michigan–Dearborn

• except A.E. and Y.E. all are  
undergraduate researchers!

# Plan for Today


1) RL recap → Half the talk almost.

2) Solving polynomial equations, why bother?

3) Gröbner basis definitions

- Monomial order
- Gröbner basis
- Gröbner fan
- F4, pair-selection


4) RL problem set-up and experiments



**Figure 3.1:** The agent–environment interaction in a Markov decision process.

## Markov Decision Process (MDP)

is a sequential decision making environment



# RL setting

◦ State-space → states agent travels

◦ Action-space → actions agent can take  
↓  
can be discrete or continuous | can be updated after every step or could be fixed

◦ Reward → the agent aims to maximize cumulative reward


$0 < \gamma < 1$  → discount factor

$$G_t \doteq R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1},$$

• Policy  $\rightarrow$  agent's strategy to pick the action in a given state

$$\pi: A \rightarrow S$$

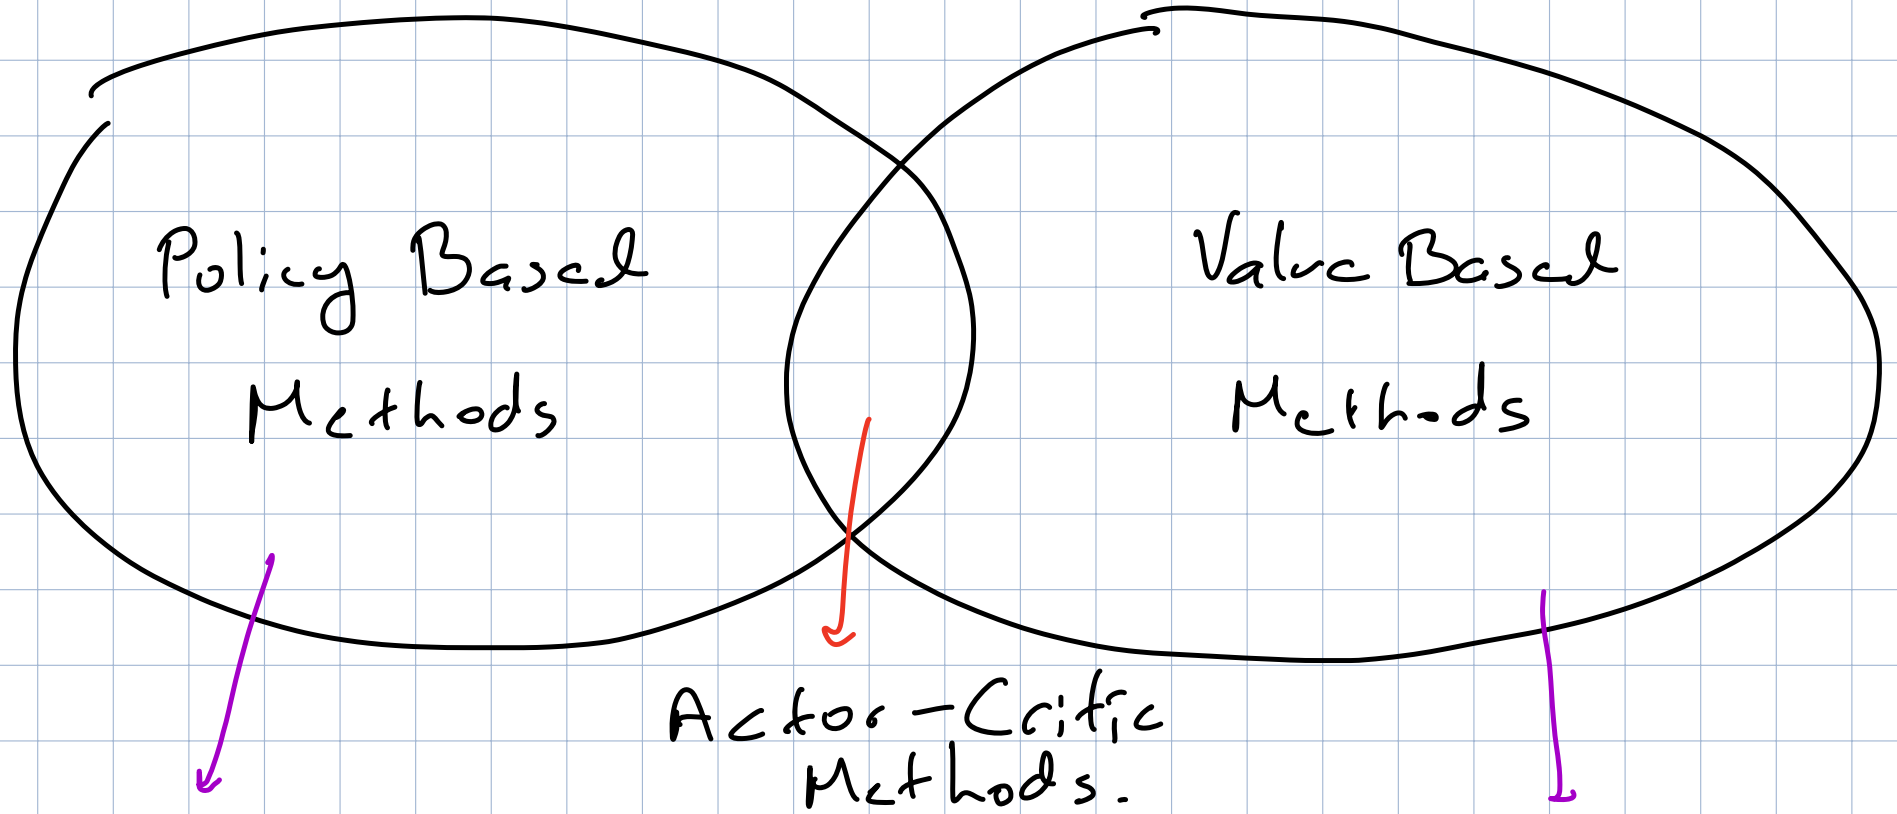
$$\sum_{a \in A} \pi(a, s) = 1$$



Backup diagram for  $v_\pi$

$\pi$  assigns value to states

$$q(s, a) = \mathbb{E} \left[ \sum_{k=0}^{\infty} \gamma^k R_{t+k} \mid S_t = s, A_t = a \right]$$



optimize the policy  
without learning the  
value function

Iterate over value  
function space, turn  
to a policy at the end.

# Policy Based optimization

- $\theta$  parameter of policy n-net
- Assume a probability measure on state-space
- Average return of policy  $\pi_\theta$  is

$$J(\theta) = \mathbb{E}_s [\text{cumulative reward of } s \mid \pi_\theta]$$

- Goal :

$$\arg\text{-max}_\theta J(\theta)$$

# Policy Gradient Theorem

- A trajectory  $z : s_0, a_0, R_0, s_1, a_1, R_1, \dots, R_T$

$\pi_{\theta}(s, a) \rightarrow$  probability of taking action  $a$   
at state  $s$  w.r.t  $\pi_{\theta}$

$p(s, a, s', r) \rightarrow$  when action  $a$  is taken the  
probability of going to  $s'$  and reward  $r$

$$IP(z) = \prod_{i=0}^T \pi_{\theta}(a_i) \cdot p(s_i, a_i, s_{i+1}, r_i)$$

Goal: Find  $\theta$  so that expected reward is maximized.

# Policy Gradient Theorem

$$\nabla_{\theta} \log P(z) = \frac{\nabla_{\theta} P(z)}{P(z)} = \nabla_{\theta} \sum_{t=1}^T \log \pi_{\theta}(a_t)$$

$R(z) :=$  discounted reward of path  $z$


$$\nabla_{\theta} J(\theta) = \sum_z \nabla_{\theta} \log P(z) \cdot P(z) \cdot R(z)$$

Empirical Version:

$$\frac{1}{n} \sum_{i=1}^n \nabla_{\theta} \log P(z_i) R(z_i)$$

## ■ Gradient tries to:

- Increase probability of paths with positive R
- Decrease probability of paths with negative R



Baseline Trick


$$\overline{J(\theta)} = \sum \pi_{\theta} p(z) (R(z) - b)$$

$$\nabla \overline{J(\theta)} = \nabla J(\theta)$$

anything that does not depend on  $\theta$ .

- a good baseline reduces the noise in training

# Actor-Critic Style




## Prioritized - Replay - Buffer Trick


- a data structure to create "artificial experiences" for the agent.

- **DDPG**: Policy gradients works for deterministic policies in continuous action space as well.

- **TD3**: DDPG is very noisy. Here is how we stabilize it



actor n.net



two critics

- at every step we update the critics with the reward collected, but delay the update of actor.

---

**Algorithm 1 TD3**

---

Initialize critic networks  $Q_{\theta_1}, Q_{\theta_2}$ , and actor network  $\pi_\phi$  with random parameters  $\theta_1, \theta_2, \phi$

Initialize target networks  $\theta'_1 \leftarrow \theta_1, \theta'_2 \leftarrow \theta_2, \phi' \leftarrow \phi$

Initialize replay buffer  $\mathcal{B}$

**for**  $t = 1$  **to**  $T$  **do**

    Select action with exploration noise  $a \sim \pi_\phi(s) + \epsilon$ ,

$\epsilon \sim \mathcal{N}(0, \sigma)$  and observe reward  $r$  and new state  $s'$

    Store transition tuple  $(s, a, r, s')$  in  $\mathcal{B}$

    Sample mini-batch of  $N$  transitions  $(s, a, r, s')$  from  $\mathcal{B}$

$\tilde{a} \leftarrow \pi_{\phi'}(s') + \epsilon, \quad \epsilon \sim \text{clip}(\mathcal{N}(0, \tilde{\sigma}), -c, c)$

$y \leftarrow r + \gamma \min_{i=1,2} Q_{\theta'_i}(s', \tilde{a})$

    Update critics  $\theta_i \leftarrow \text{argmin}_{\theta_i} N^{-1} \sum (y - Q_{\theta_i}(s, a))^2$

**if**  $t \bmod d$  **then**

        Update  $\phi$  by the deterministic policy gradient:

$\nabla_\phi J(\phi) = N^{-1} \sum \nabla_a Q_{\theta_1}(s, a)|_{a=\pi_\phi(s)} \nabla_\phi \pi_\phi(s)$

        Update target networks:

$\theta'_i \leftarrow \tau \theta_i + (1 - \tau) \theta'_i$

$\phi' \leftarrow \tau \phi + (1 - \tau) \phi'$

**end if**

**end for**

---

# Solving Polynomials

$$\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{Z}^n$$

$$x^\alpha := x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n}$$

Ex:  $\alpha = (2, 3)$        $x^\alpha = x_1^2 x_2^3$

Support set:  $A = \{a_1, a_2, \dots, a_t\} \subseteq \mathbb{Z}^n$

$$a_i = (a_{i1}, a_{i2}, \dots, a_{in}).$$

$$p_1(x) = \sum_{\alpha \in A} c_\alpha^{(1)} x^\alpha, \quad p_2(x) = \sum_{\alpha \in A} c_\alpha^{(2)} x^\alpha, \quad \dots$$

$$\dots, \quad p_k(x) = \sum_{\alpha \in A} c_\alpha^{(k)} x^\alpha$$

Goal: Find all  $u = (u_1, u_2, \dots, u_m) \in \mathbb{R}^m$  such that

$$\mathcal{Z}(p) = \left\{ u : p_1(u) = p_2(u) = \dots = p_k(u) = 0. \right\}$$



- Note that for any  $u \in \mathcal{Z}(p)$  and for any  $h$  s.t.

$$h = p_1(x) g_1(x) + p_2(x) g_2(x) + \dots + p_n(x) g_n(x)$$



we have  $h(u) = 0$ .

Why bother?



Optimization




Statistics



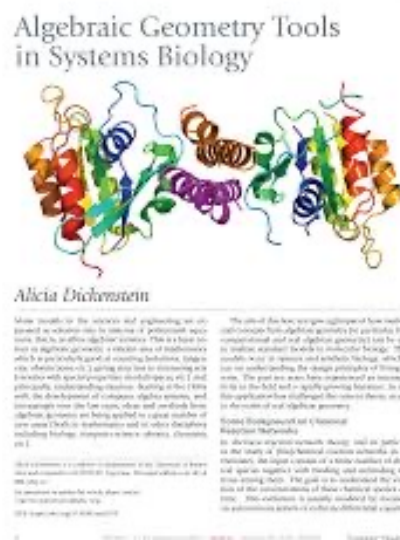

# Computer Vision



# Cryptography



# Systems Biology




# Kinematics

# Particle Physics

# Power-Flow networks

The Numerical Solution of Systems of Polynomials Arising in Engineering and Science



Andrew J. Sommese • Charles W. Wampler, II

# Gröbner Basis

• Monomial Order  $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$   $\beta = (\beta_1, \beta_2, \dots, \beta_n)$

How to compare  $\alpha$  and  $\beta$ ?

- the order should be a well ordering

no ties!

- If  $\alpha \succ \beta$  then  $\alpha + \gamma \succ \beta + \gamma$   
for any  $\gamma$ .

Dfn: 
$$p(x) = \sum_{\alpha \in A} c_{\alpha} x^{\alpha}$$

The  $\alpha \in A$  such that  $\alpha \succ \beta$  for all  $\beta \in A$

is called the leading term.  $LT(p)$ .

Dfn.:  $\langle P_1, P_2, \dots, P_k \rangle = \{ h \in \mathbb{R}[x] : h(x) = p_1(x)g_1(x) + \dots + p_k(x)g_k(x) \}$

Dfn.: A finite set of polynomials  $g_1, g_2, \dots, g_m$  is a Gröbner basis for  $P_1, P_2, \dots, P_k$  w.r.t  $\succ$  if

$$\langle LT(g_1), LT(g_2), \dots, LT(g_m) \rangle = \langle LT(P_1), \dots, LT(P_k) \rangle$$

### Buchberger's Algorithm

- Fix a monomial order, and a pair selection strategy.
- Start with  $P_1, \dots, P_k$ , order, pair-select, reduce until it stops.

## A simple way to encode monomial orders

- Pick  $v \in \mathbb{R}^n$   
compute  $\langle v, \alpha \rangle$  for all  $\alpha \in A$
- If  $\langle \alpha, v \rangle > \langle \alpha, \beta \rangle$   
then say  $\alpha \succ \beta$
- Create a tie-breaking rule in case  
 $\langle \alpha, v \rangle = \langle \alpha, \beta \rangle$

Note: Scaling  $v$  does not change the order.

Infinitely many orders actually create  
 finitely many different Gröbner basis.




Figure 1: A schematic two-dimensional slice ( $w_1 + w_2 + w_3 = 1, w \geq 0$ ) of the Gröbner fan for  $I = \langle x + y + z, x^3z + x + y^2 \rangle$ , illustrating how regions of constant  $\text{in}_w(I)$  arise. The diagram is not to scale and is intended to convey the geometric structure of cones rather than exact fan boundaries; see [23].

In practice ,

Instead of optimizing for fastest monomial order, community defaulted to a few monomial orders.

Lex, Graded Lex, Reverse Graded Lex

We object !


Make monomial order selection a RL - problem.

## Computational Setup

- We had to use Julia because Python packages only allow few fixed monomial orders
- We used `Groebner.jl`. It implements  $F_4$  algorithm for Gröbner basis computation
- $F_4$  does several reductions at the same time to speed-up computations.

Our reward function utilizes the Gröbner trace statistics of the F4 algorithm recorded by `Groebner.jl`, specifically: the number of iterations ( $t$ ), the number of columns in the elimination matrix ( $n_M(i)$ ) for  $1 \leq i \leq t$ , the number of pairs ( $n_P(i)$ ), and the degree of pairs ( $d_P(i)$ ). The reward function is defined as:

$$-\sum_{i=1}^t n_M(i)n_P(i) \ln(d_P(i))$$



## State - Space

The Groebner.jl package accepts only integer entries in weight vectors for the monomial ordering. However, the theoretical search space is the simplex:

$$\Delta_n := \{x = (x_1, x_2, \dots, x_n) : x_i \geq 0, \sum x_i = 1\}$$

Therefore, we define the action space as the simplex and define a monomial ordering by scaling the action vector by  $10^3$  and rounding to the nearest integer.

## Action - Space

Steps inside  $\Delta_n$  that has norm

at most  $\delta$ .

• We tested on four problems from  
Computer Vision and Systems Biology.

Polynomial System	GrevLex baseline			GrLex baseline		
	Win/Tie/Loss	Improvement	Degradation	Win/Tie/Loss	Improvement	Degradation
Relative Pose	98.4 / 0.0 / 1.6	19.43	-1.27	32.2 / 16.1 / 51.6	7.95	-7.04
Triangulation	27.0 / 66.1 / 7.0	0.86	-0.72	100 / 0 / 0	37.62	-
$n$ -Site Phos. ( $n=14$ )	100 / 0 / 0	70.77	-	100 / 0 / 0	70.77	-
Wnt Shuttle	91.6 / 0.0 / 8.4	54.64	-26.40	86.9 / 0.0 / 13.1	55.25	-20.83

Relative Pose > vision       $n$ -Site Phos > Bio  
 Triangulation                      Wnt Shuttle

Thank you!

Danke!

Gracias!

Teşekkürler!

ευχαριστώ!

متشکرم!

有り難う!