

# Evolution Strategies for Approximate Solution of Bayesian Games

Zun Li   Michael P. Wellman

Computer Science & Engineering  
University of Michigan, Ann Arbor

AAAI, Feb 2021



# Main Results Summary

- Goal: Compute Bayes-Nash equilibria (BNE) for Bayesian games with
  - Many ( $N > 2$ ) symmetric players
  - High-dimensional types, actions
  - Black-box access to
    - Type distribution
    - Payoff values

# Main Results Summary

- Goal: Compute Bayes-Nash equilibria (BNE) for Bayesian games with
  - Many ( $N > 2$ ) symmetric players
  - High-dimensional types, actions
  - Black-box access to
    - Type distribution
    - Payoff values
- Approach:
  - Represent strategy as a neural network
  - Exploit evolutionary computation and game symmetry

# Main Results Summary

- Goal: Compute Bayes-Nash equilibria (BNE) for Bayesian games with
  - Many ( $N > 2$ ) symmetric players
  - High-dimensional types, actions
  - Black-box access to
    - Type distribution
    - Payoff values
- Approach:
  - Represent strategy as a neural network
  - Exploit evolutionary computation and game symmetry
- Main contribution: Algorithms for
  - Computing pure BNE via minimax optimization
  - Computing mixed BNE via incremental strategy generation

# Symmetric Bayesian Games (SBGs)

- Strategy is a mapping from types to actions

# Symmetric Bayesian Games (SBGs)

- Strategy is a mapping from types to actions
- Types are agents' private knowledge of the game, which are i.i.d. sampled
  - Ex: in a multi-object auction, type is a vector defining one's valuation for sets of goods, and actions are bids for these goods

# SBG Payoff Functions

- Payoff for one player: a function of (joint action, own type)
- Typically, structured functional form
  - Ex: valuation(outcome bundle) – payment

# SBG Payoff Functions

- Payoff for one player: a function of (joint action, own type)
- Typically, structured functional form
  - Ex: valuation(outcome bundle) – payment
- Anonymity: Payoff is permutation-invariant wrt other-player actions
  - Ex: only the highest other-player bid matters
- Symmetry: Each player has same payoff function



# Strategy Space for SBGs

- A pure strategy  $s$  is a deterministic mapping from types to actions
- A mixed strategy  $\sigma$  is a probability measure over a set of pure strategies

# Strategy Space for SBGs

- A pure strategy  $s$  is a deterministic mapping from types to actions
- A mixed strategy  $\sigma$  is a probability measure over a set of pure strategies
- We consider cases where both types and actions represented by real vectors
- We represent a pure strategy as a neural network

# Bayes-Nash Equilibrium

- We seek symmetric equilibria
- Denote the  $u(\sigma', \sigma)$  the expected payoff received when one play  $\sigma'$  while all the rest  $N - 1$  adopt  $\sigma$ , marginalized over all joint type realization
- $Regret(\sigma) = \max_s u(s, \sigma) - u(\sigma, \sigma)$ , and we seek for  $\sigma$  that minimize its regret

# Natural Evolution Strategies

- Goal: optimize a black-box function  $F(\theta)$  with respect to network weights  $\theta$
- Approach: optimize a Gaussian smoothed function  $\mathbb{E}_{\varepsilon \sim \mathcal{N}(0, I)}[F(\theta + \nu \varepsilon)]$  by constructing finite difference approximation of the gradients

---

## Algorithm 1: Natural Evolution Strategies

---

**Input:** Black-box function  $F$ , hyperparameters  $J, \alpha, \nu$

**Output:** Approximate maximum  $\theta$  of  $F$

1 **Algorithm** NES( $F, J, \alpha, \nu$ )

2     Initialize  $\theta$ ;

3     **for**  $i = 1, 2, \dots$  **do**

4         Sample  $\varepsilon_1, \dots, \varepsilon_J \sim \mathcal{N}(0, I)$ ;

5          $\forall j, r_{j+} \leftarrow F(\theta + \nu \varepsilon_j), r_{j-} \leftarrow F(\theta - \nu \varepsilon_j); r_j \leftarrow r_{j+} - r_{j-}$ ;


6          $\theta \leftarrow \theta + \alpha \frac{1}{J\nu} \sum_j r_j \varepsilon_j$ ;

7     **end**


8     **return**  $\theta, F(\theta)$ ;

# Natural Evolution Strategies


Iteration 1, fitness 0.67




Iteration 2, fitness 0.94



Iteration 3, fitness 1.25




Iteration 4, fitness 1.32



# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES




Best response NES

# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES

Multiple calls on points around  $s_1$



Best response NES


# Computing PBNE via Minimax Optimization

- A minimax formulation:


$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Simulation  
payoffs



Regret minimization NES




Best response NES




# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$




Regret values of  
these points



# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES

Multiple calls on points around  $s_2$




Best response NES

# Computing PBNE via Minimax Optimization

- A minimax formulation:


$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES



Simulation  
payoffs




Best response NES

# Computing PBNE via Minimax Optimization


- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES

Regret values of  
these points




Best response NES

# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES

Multiple calls on points around  $s_3$



Best response NES


# Computing PBNE via Minimax Optimization

- A minimax formulation:


$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Simulation  
payoffs



Regret minimization NES




Best response NES

# Computing PBNE via Minimax Optimization


- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$



Regret minimization NES

Regret values of  
these points




Best response NES

# Computing PBNE via Minimax Optimization

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$

Approximate PBNE  $s^*$



Regret minimization NES



Best response NES



# Minimax-NES for Computing PBNE

- A minimax formulation:

$$\min_s \text{Regret}(s) = \min_s \max_{s'} u(s', s) - u(s, s)$$

---

**Algorithm 2:** Minimax-NES for PBNE

---

**Input:** Payoff Oracle  $\mathcal{O}$ , hyperparameters  $J_1, J_2, \alpha_1, \alpha_2, \nu_1, \nu_2$

**Output:** Approximate PBNE  $s_\theta$

1 **Function** MinusRegret( $\theta$ )

2      $V \leftarrow \mathcal{O}(s_\theta, s_\theta);$

3      $\theta', DEV \leftarrow \text{NES}(\mathcal{O}(\cdot, s_\theta), J_1, \alpha_1, \nu_1);$

4     **return**  $V - DEV;$


5 **Algorithm** MiniMax()

6     **return**  $\text{NES}(\text{MinusRegret}, J_2, \alpha_2, \nu_2)$

---

# Results for Games with Analytical Solutions

- Canonical solution for  $N$ -player single-item first-price auction  
 $FP[N]$  is  $s(t) = \frac{N-1}{N}t$ .
- Second-price  $SP[N] : s(t) = t$ .
- Third-price  $TP[N] : s(t) = \frac{N-1}{N-2}t$ .



# Incremental Strategy Generation for Computing MBNE

- ISG discretizes the strategy space as a finite strategy set  $S$ , and iteratively enlarges  $S$  via best responses
- two components: a meta-solver and a best response oracle

---

## Algorithm 3: Incremental Strategy Generation

---

**Input:** Payoff Oracle  $\mathcal{O}$ . Meta-solver  $MS$ . Hyperparameters  $J, \alpha, \nu$  ;

**Output:** A finite strategy set  $S$ , a mixed strategy  $\sigma$  over  $S$

```
1 Initial strategy set  $S = \{s_0\}$ , a singleton distribution  $\sigma$  with  $\sigma(s_0) = 1$ ;  
2 for  $i = 1, 2, \dots$  do  
3    $\sigma \leftarrow MS(\mathcal{O}, S, \sigma)$ ;  
4    $s', DEV \leftarrow NES(\mathcal{O}(\cdot, \sigma), J, \alpha, \nu)$ ;  
5    $S \leftarrow S \cup \{s'\}$ ;  
6 end
```

---

# Incremental Strategy Generation for Computing MBNE

- Given an iteration with restricted strategy set  $S$ , a meta-solver outputs a probability mixture over  $S$ , which could be:

---

## Algorithm 3: Incremental Strategy Generation

---

**Input:** Payoff Oracle  $\mathcal{O}$ . Meta-solver  $MS$ . Hyperparameters  $J, \alpha, \nu$  ;

**Output:** A finite strategy set  $S$ , a mixed strategy  $\sigma$  over  $S$

```
1 Initial strategy set  $S = \{s_0\}$ , a singleton distribution  $\sigma$  with  $\sigma(s_0) = 1$ ;  
2 for  $i = 1, 2, \dots$  do  
3    $\sigma \leftarrow MS(\mathcal{O}, S, \sigma)$ ;  
4    $s', DEV \leftarrow NES(\mathcal{O}(\cdot, \sigma), J, \alpha, \nu)$ ;  
5    $S \leftarrow S \cup \{s'\}$ ;  
6 end
```

---

# Incremental Strategy Generation for Computing MBNE

- Given an iteration with restricted strategy set  $S$ , a meta-solver outputs a probability mixture over  $S$ , which could be:
  - Self-play: all mass on the last pure strategy
  - Fictitious play: uniform mixture on  $S$
  - Replicator dynamics: a Nash equilibrium calculated by RD on the finite game defined by  $S$

---

## Algorithm 3: Incremental Strategy Generation

---

**Input:** Payoff Oracle  $\mathcal{O}$ . Meta-solver  $MS$ . Hyperparameters  $J, \alpha, \nu$  ;

**Output:** A finite strategy set  $S$ , a mixed strategy  $\sigma$  over  $S$

```
1 Initial strategy set  $S = \{s_0\}$ , a singleton distribution  $\sigma$  with  $\sigma(s_0) = 1$ ;  
2 for  $i = 1, 2, \dots$  do  
3    $\sigma \leftarrow MS(\mathcal{O}, S, \sigma)$ ;  
4    $s', DEV \leftarrow NES(\mathcal{O}(\cdot, \sigma), J, \alpha, \nu)$ ;  
5    $S \leftarrow S \cup \{s'\}$ ;  
6 end
```

---

# Computing MBNE via Incremental Strategy Generation

- And then NES generates a best response strategy against this mixture into  $S$

---

## Algorithm 3: Incremental Strategy Generation

---

**Input:** Payoff Oracle  $\mathcal{O}$ . Meta-solver  $MS$ . Hyperparameters  $J, \alpha, \nu$  ;

**Output:** A finite strategy set  $S$ , a mixed strategy  $\sigma$  over  $S$

```
1 Initial strategy set  $S = \{s_0\}$ , a singleton distribution  $\sigma$  with  $\sigma(s_0) = 1$ ;  
2 for  $i = 1, 2, \dots$  do  
3    $\sigma \leftarrow MS(\mathcal{O}, S, \sigma)$ ;  
4    $s', DEV \leftarrow NES(\mathcal{O}(\cdot, \sigma), J, \alpha, \nu)$ ;  
5    $S \leftarrow S \cup \{s'\}$ ;  
6 end
```

---

# Model Learning

- To calculate an NE on the game defined by  $S$ , one needs access to its  $u(s, \sigma)$ .
- Direct Monte-Carlo via sampling could be expensive and induce large variance for each update in RD

# Model Learning

- To calculate an NE on the game defined by  $S$ , one needs access to its  $u(s, \sigma)$ .
- Direct Monte-Carlo via sampling could be expensive and induce large variance for each update in RD
- Approach: learn a value network  $\hat{u} : \Delta(S) \rightarrow \mathbb{R}^M$  via simulation to estimate these deviation payoffs, where  $M = |S|$ .





# Model Learning

- To calculate an NE on the game defined by  $S$ , one needs access to its  $u(s, \sigma)$ .
- Direct Monte-Carlo via sampling could be expensive and induce large variance for each update in RD
- Approach: learn a value network  $\hat{u} : \Delta(S) \rightarrow \mathbb{R}^M$  via simulation to estimate these deviation payoffs, where  $M = |S|$ .
- then run RD on  $\hat{u}$  to get an approximate NE

# Evaluation Metrics



- Four methods: minimax-NES and self-play computing PBNE; fictitious play and replicator dynamics solving for MBNE
- Environments:  $N$ -player  $K$ -good market-based scheduling ( $MBS[N, K]$ ) & homogeneous-good auctions ( $HG[N, K]$ ). Both multidimensional types and actions possess no analytic solutions

# Results





# Results

- Self-play cycles





# Results

- Self-play cycles
- RD slightly outperforms FP



# Results

- Self-play cycles
- RD slightly outperforms FP
- Minimax-NES behaves less robust than mixed equilibria in MBS, but outperforms the others in HG



## Comparison to Hand-Crafted Strategies

- Compare against the state-of-the-art class of hand-crafted strategies called *self-confirming bidding strategies* (SC).

Instance	SP	MM	FP	RD	SC
<i>MBS</i> [5, 5]	19.1	3.51	3.47	2.51	5.30
<i>HG</i> [5, 5]	49.7	7.62	28.9	26.0	11.0

Table: Regret of SC compared with other methods within  $\bar{S}$

Instance	SP	MM	FP	RD
<i>MBS</i> [5, 5]	0.0	3.46	0.74	1.71
<i>HG</i> [5, 5]	0.45	3.05	0.0	0.0

Table: Regret of our methods with respect to SC

# Conclusion

- Our algorithms need not to know type distribution & utility functional form but only black-box samples
- Both PBNE and MBNE computation methods employ NES as a powerful optimization procedure, which can be replaced by other methods (e.g. genetic algorithms)
- Both methods conduct a many-to-two game reduction by exploiting player symmetry