Competitive Gradient Optimization

Abhijeet Vyas

Ph.D. student, Purdue

Dr. Brian Bullins

Purdue

Dr. Kamyar Azizzadenesheli

Nvidia

Talk Overview.

Min-max optimization

What is it?

Challenges?

Min-max optimization

Recent Work.

CGO

Motivation.

Algorithm.

CGO

Unconstrained Setting

CGO

Constrained setting

How to use in research?

Accessibility

The min-max optimization problem.

$$\min_{x \in \mathcal{X}} f(x, y), \quad \max_{y \in \mathcal{Y}} f(x, y)$$

What does it mean to solve the min-max problem?

- Obtain global Nash equilibrium
- Obtain local Nash equilibrium
- Obtain a stationary point
- Some other notions exist for sequential games.
  - Global minimax points
  - Local minimax points
  - Stackelberg equilibrium

Min-Max optimization problem is relevant in,

COMPETITIVEREINFORCEMENTLEARNING

ROBUSTLEARNING

RISKASSESEMENT

GENERATIVEADVERSARIALNETWORKS

REINFORCEMENTLEARNING

DEEPLEARNING

ADVERSERIALML

NASHEQUILIBRIUM

DISTRIBUTEDOPTIMIZATION

ECONOMICS

OPTIMIZATION

COMPETITIVEMARKOVDECISIONPROCESSES

BELLMANEQUATIONS

GAMETHEORY

MINIMAX-OPTIMIZATION

MULTILOSSOPTIMIZATION

GANS

MACHINELEARNING

SADDLEPOINTPROBLEMS

What are the challenges?

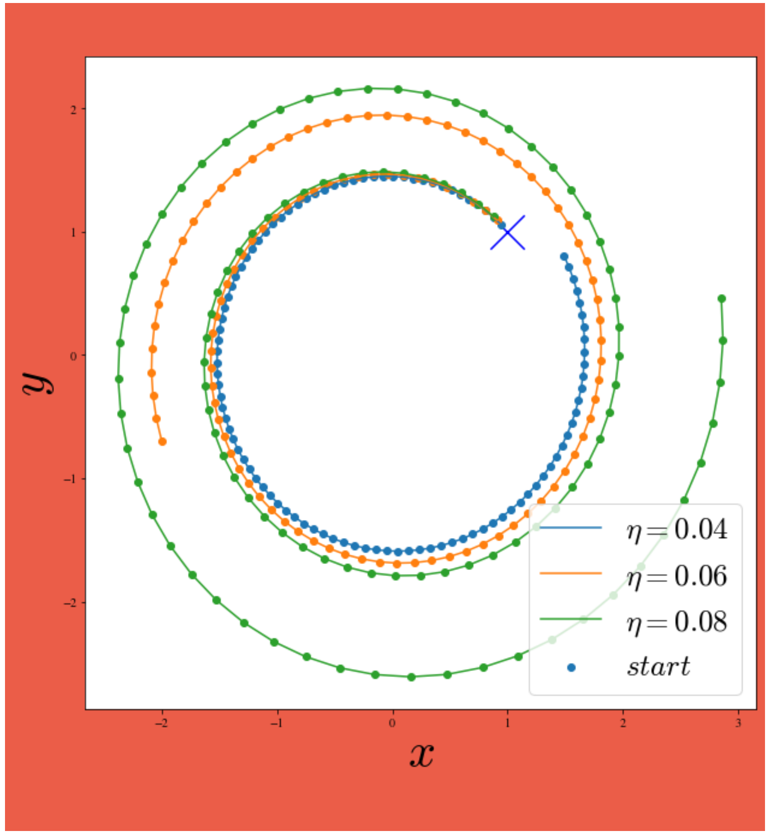

GDA diverges on simple functions.

GDA on f(x,y) = xy

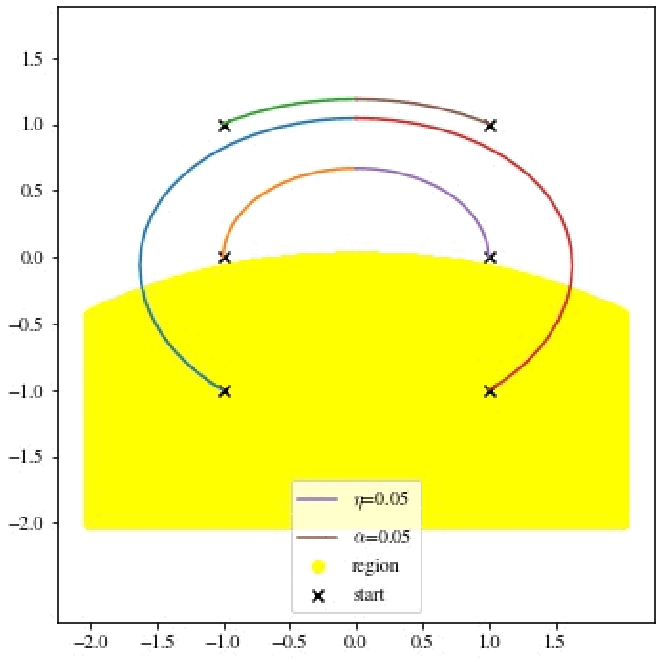

Non-convex non-concave Optimization

Competitive Gradient Descent on f(x,y) = x^2y

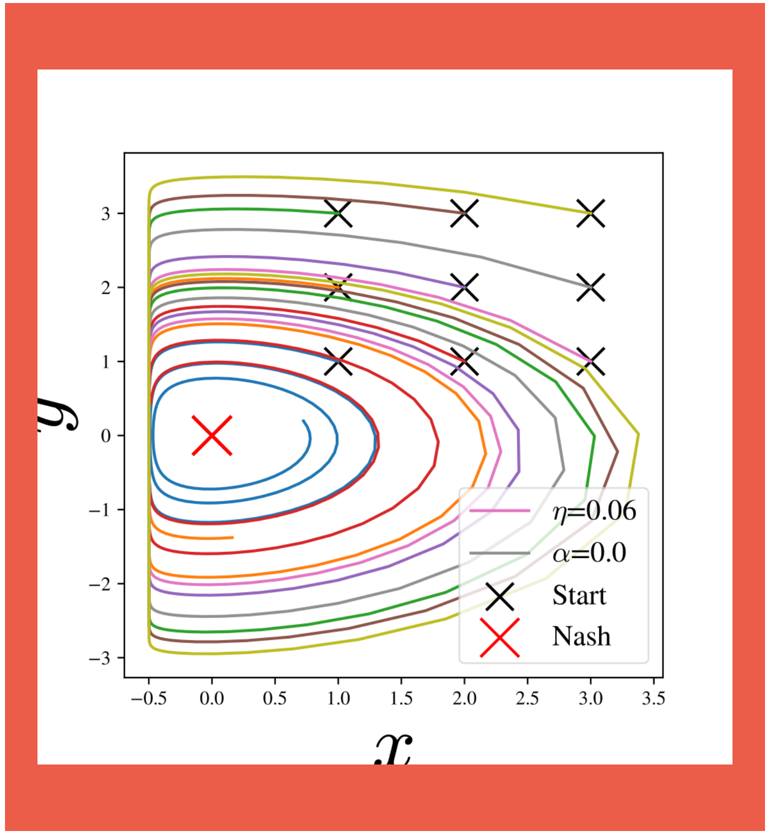

Algorithms cycle

Optimistic Mirror Descent on f(x,y) = x^2y+xy

Recent non-convex non-concave min-max optimization.

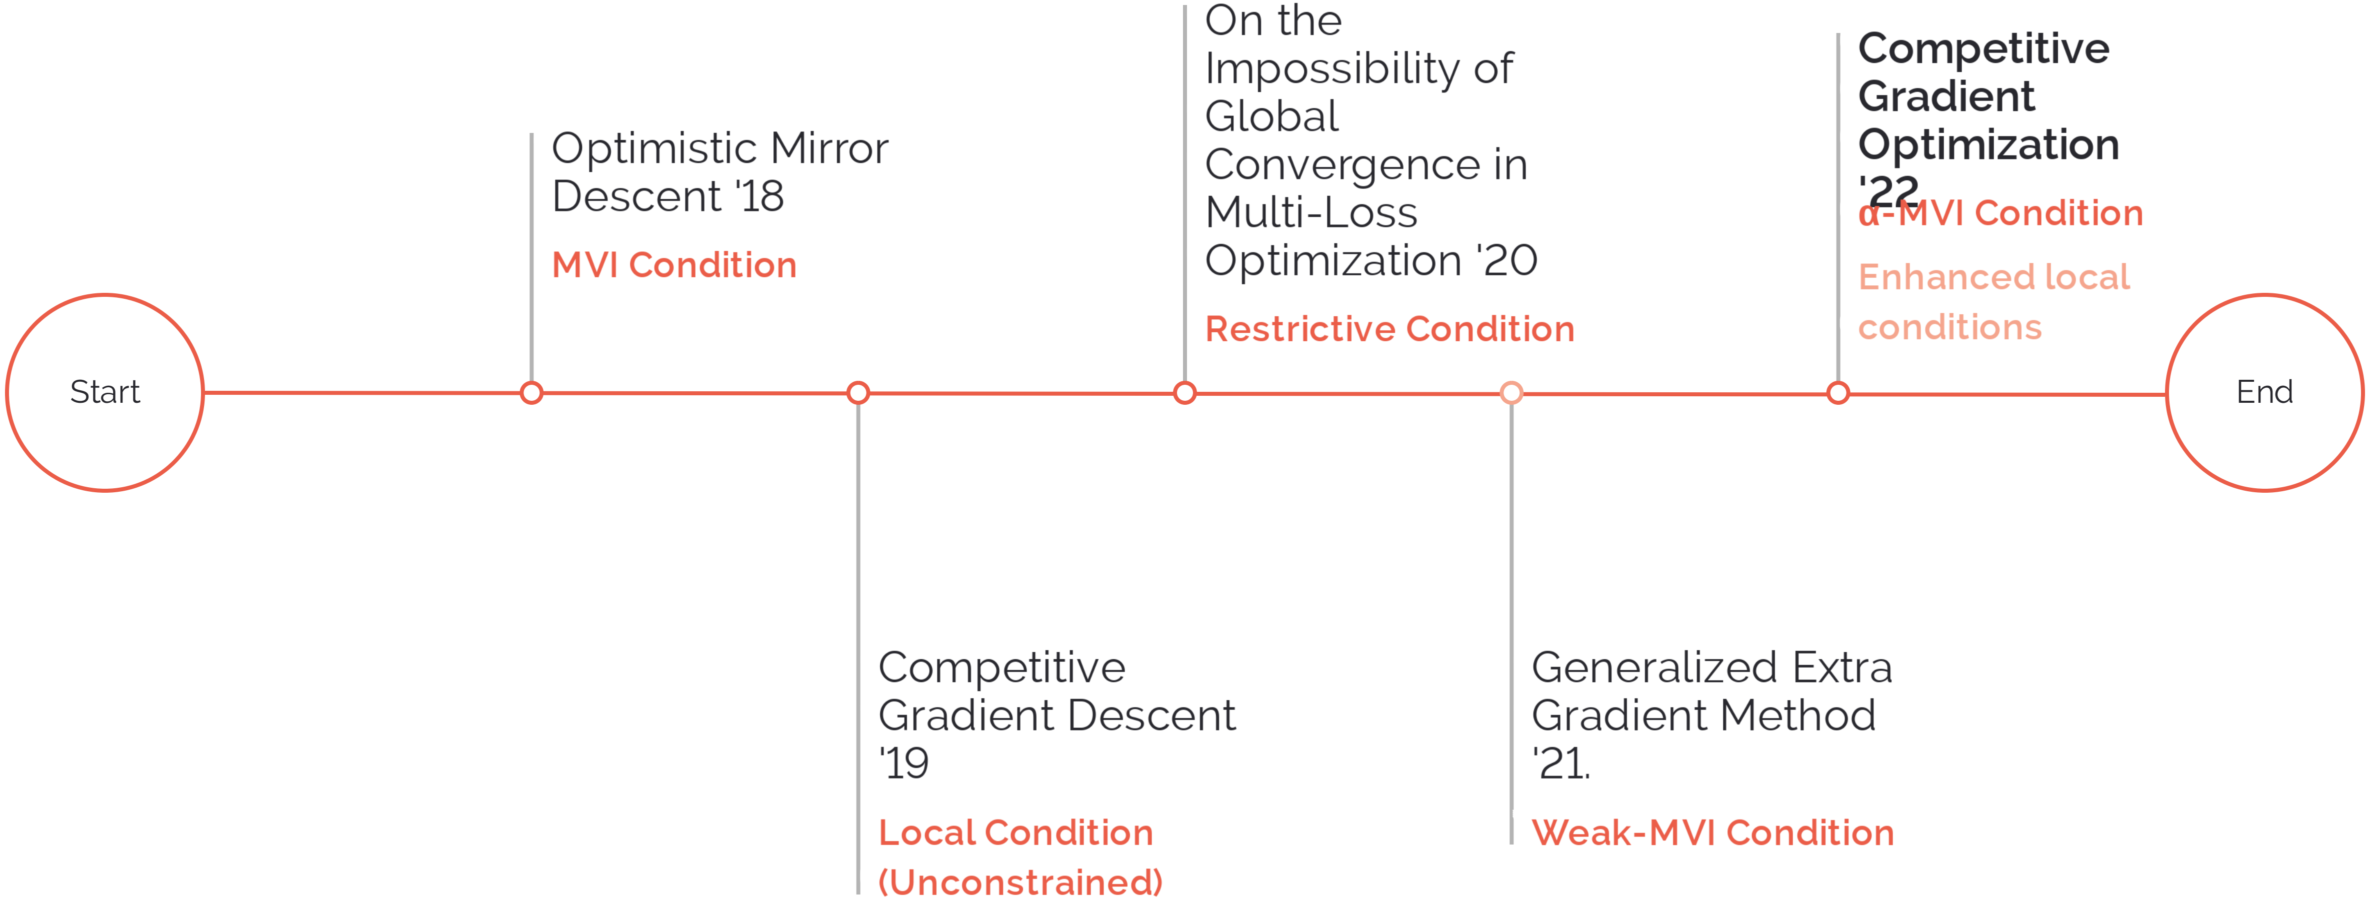

The local game.

At an iteration point CGO solves,

$$arg\,min_{\delta x \in \mathcal{X}} \;\; \delta x^\top \nabla_x f + \frac{\alpha}{\eta} \delta x^\top \nabla^2_{xy} f \delta y + \delta y^\top \nabla_y f + \frac{1}{2\eta} \delta x^\top \delta x$$

$$arg\,max_{\delta y \in \mathcal{Y}} \;\; \underbrace{\delta y^\top \nabla_y f + \frac{\alpha}{\eta} \delta y^\top \nabla^2_{yx} f \delta x + \delta x^\top \nabla_x f}_{Taylor\ expansion\ of\ f\ around\ (x,y),\ \delta x\ terms\ ignored} \; - \underbrace{\frac{1}{2\eta} \delta y^\top \delta y}_{Regularization}$$

*Set α=η to obtain exact Taylor expansion

CGO generalizes GDA and CGD

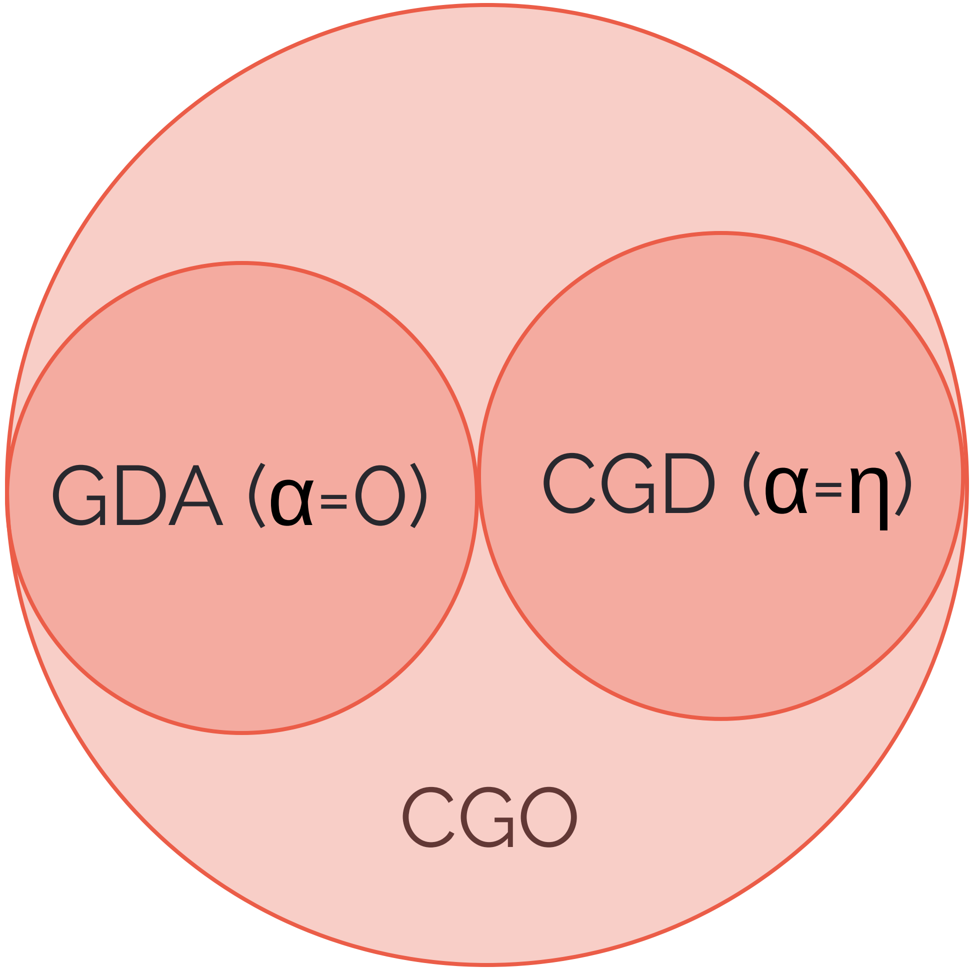

The global Nash of the local game is,

$$\delta x = -\eta \left(I + \alpha^2 \nabla^2_{xy} f \nabla^2_{yx} f\right)^{-1} \left(\nabla_x f + \alpha \nabla^2_{xy} f \nabla_y f\right)$$
$$\delta y = -\eta \left(I + \alpha^2 \nabla^2_{yx} f \nabla^2_{xy} f\right)^{-1} \left(-\nabla_y f + \alpha \nabla^2_{yx} f \nabla_x f\right)$$

The unconstrained CGO algorithm.

-Converges for arbitrary deviations from convex-concave condition for small learning rates.

-Enhances local conditions in discrete-time.

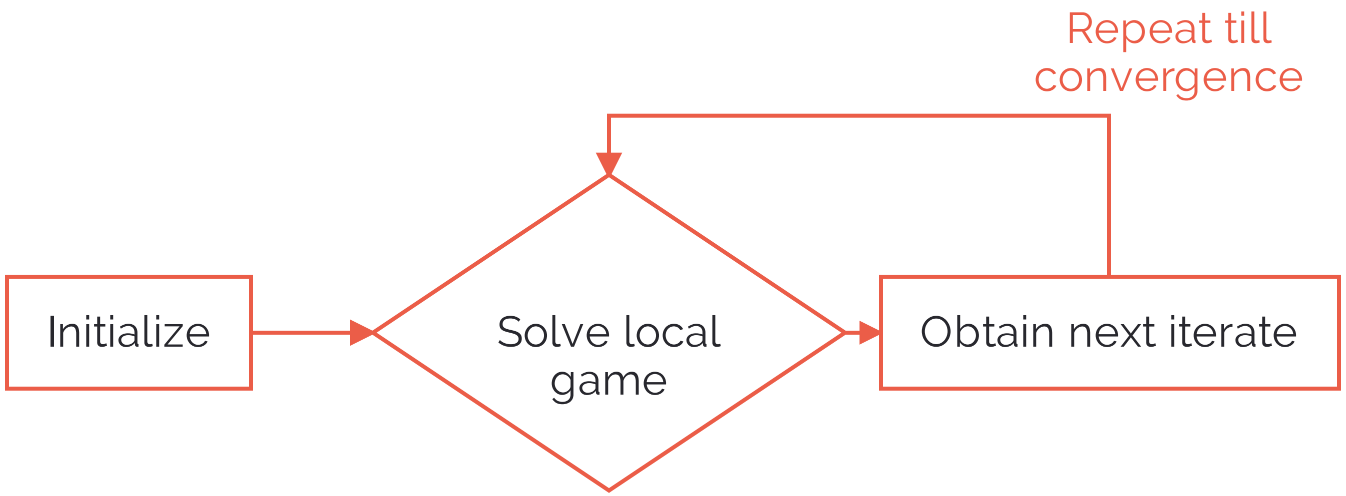

The continuous-time regime.

Description

$$\lim_{\eta \to 0}\left(\beta \frac{\delta x}{\eta}\right) = \dot{x} = -\beta\big(I + \alpha^2 \nabla_{xy}^2 f \nabla_{yx}^2 f\big)^{-1}\big(\nabla_x f + \alpha \nabla_{xy}^2 f \nabla_y f\big)$$

$$\lim_{\eta \to 0}\left(\beta \frac{\delta x}{\eta}\right) = \dot{y} = -\beta\big(I + \alpha^2 \nabla_{yx}^2 f \nabla_{xy}^2 f\big)^{-1}\big(-\nabla_y f + \alpha \nabla_{yx}^2 f \nabla_x f\big)$$

Comments

- Setting t=$^{\eta}/_{\beta}$ where **β** is the time scaling factor.
- For **α** = **η**, the update approaches GDA in the limit

Results

$$\textit{Continuous}-\textit{time CGO converges to a stationary point exponentially}$$

$$\textit{with rate } \lambda = \beta \min\left(2\underline{\lambda}_{xx} - 2\alpha \overline{\lambda}_{xx}^{2} + c\frac{\underline{\lambda}_{xy}}{1+\alpha^2 \underline{\lambda}_{xy}},\ -2\overline{\lambda}_{yy} - 2\alpha \overline{\lambda}_{yy}^{2} + c\frac{\underline{\lambda}_{yx}}{1+\alpha^2 \underline{\lambda}_{yx}}\right)$$

$$\textit{where } \overline{\lambda_1} = \max(\overline{\lambda}_{xx}, -\underline{\lambda}_{yy}),\ \overline{\lambda_2} = \max(\overline{\lambda}_{xx}, \overline{\lambda}_{yy}) \textit{ and } c = \beta(\alpha - 2\alpha^2 \overline{\lambda_1} - 2\alpha^3 \overline{\lambda_2}^{2}).$$

We can show

- That appropriately setting α allows λ to stays positive for deviations of magnitude min($\underline{\lambda}_{xy}$,$\underline{\lambda}_{yx}$)$^{0.5}$ (the square root of the singular values of the cross-terms of the Hessian) from the convex-concave condition.
- yy <= M, $\underline{\lambda}_{xy}$>=M, where M = min($\underline{\lambda}_{xy}$,$\underline{\lambda}_{yx}$)$^{0.5}$

The constrained CGO algorithm.

Converges for α-MVI functions

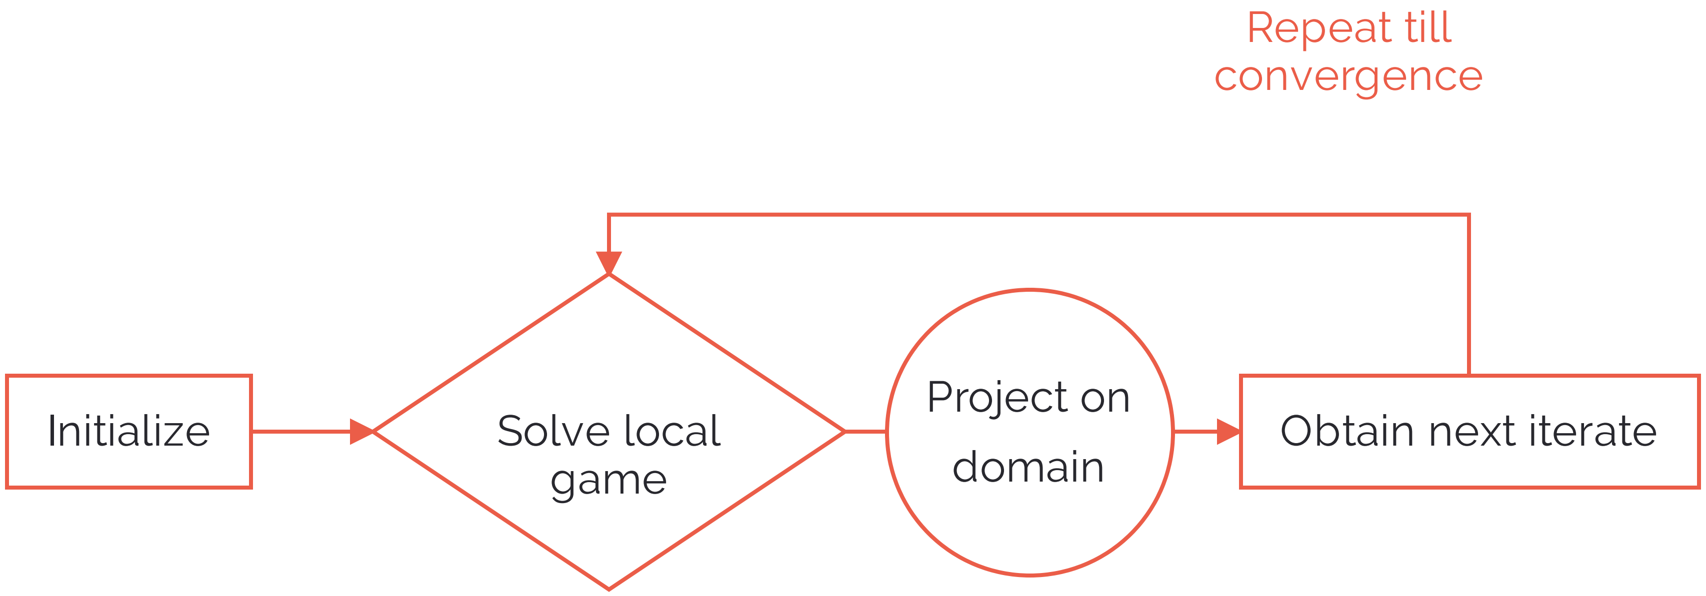

Examples

1. Strictly convex-concave : $f(x,y) = x^2 - y^2$

2. Convex-concave : $f(x,y) = x^T Ay$ (bilinear), any matrix A

3. MVI : $f(x,y) = (x^4y^2+x^2+1)(x^2y^4-x^2+1)$, Domain : $[-1,1]^2$

4. α-MVI (exclusive) : $f(x,y) = x^2y$, x^2y+xy

5. Weak-MVI (exclusive) : certain Neumann ratio games

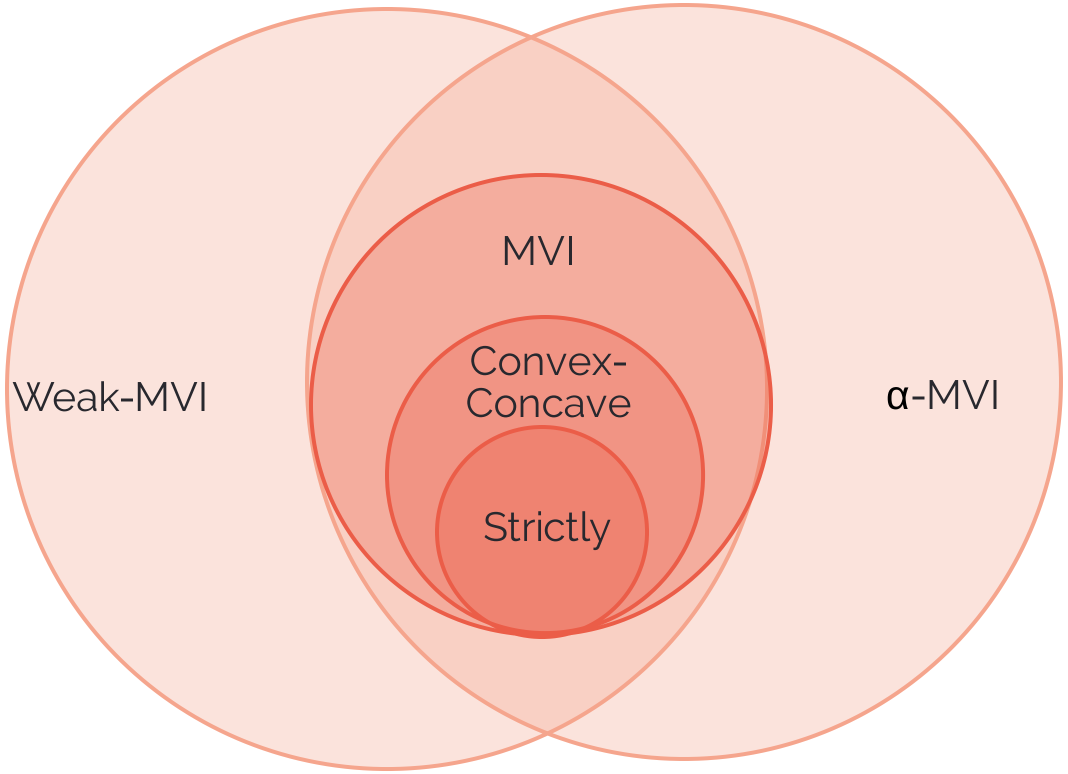

CGO in action,

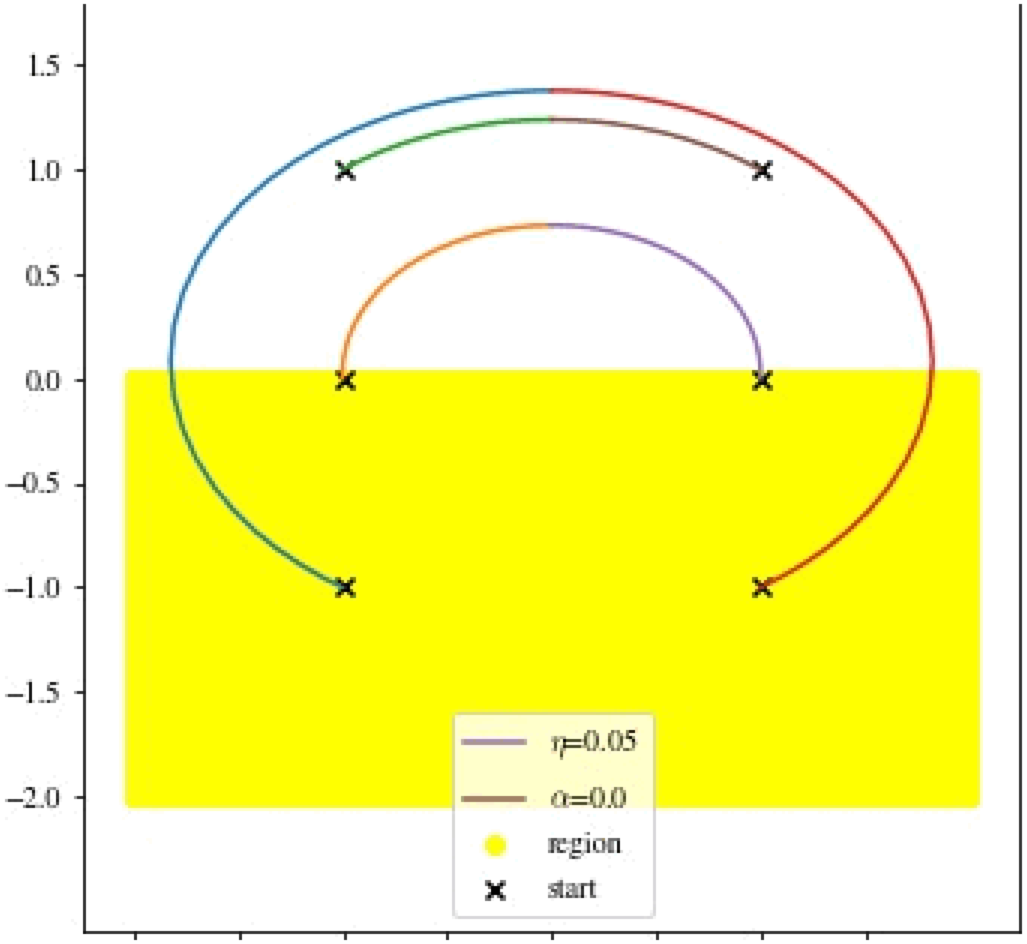

$f(x,y)=x^2y$, (α-MVI with α approaches infinity)

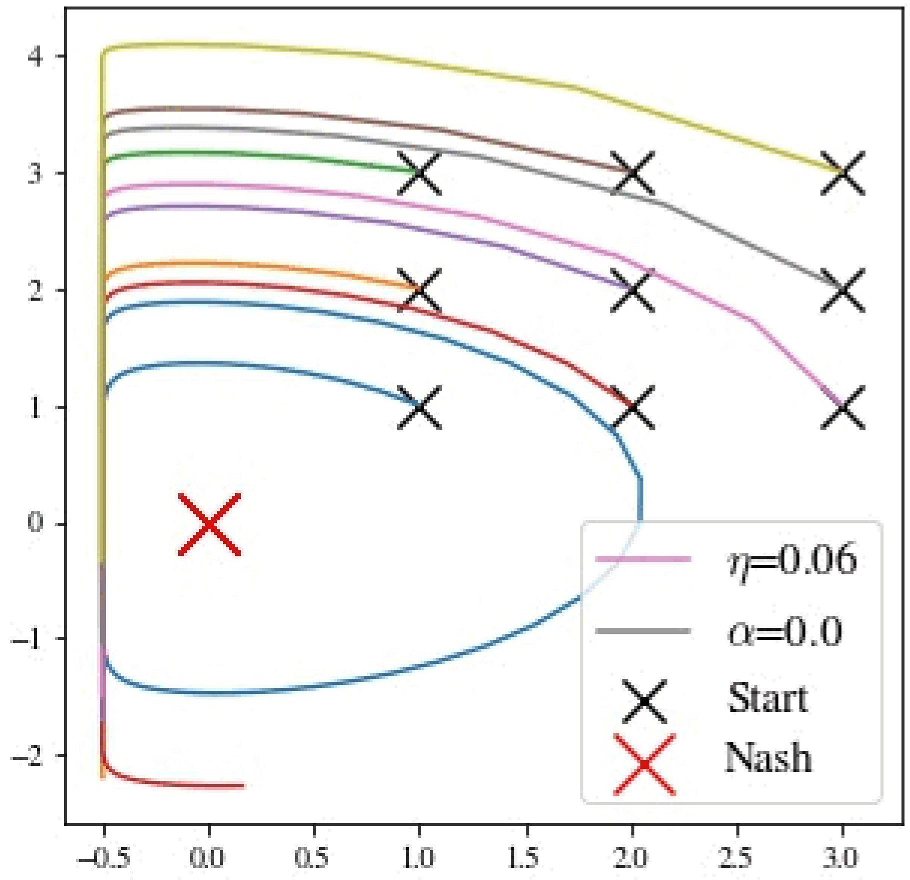

$f(x,y)=x^2y+xy$ (α-MVI with α>=2)

The CGO algorithm (and it's analysis),

- Solves a novel local game.
- Produces a distinct algorithm from GDA in continuous-time.
- Allows arbitrary deviation from convex-concave based on the cross-terms.
- Obtains enhanced local convergence guarantees to stationary points in unconstrained optimization
- Explains convergence of CGD on bilinear functions
- Obtains convergence guarantees on α=MVI class of functions for constrained optimization

How to use this in your research?

- Use to jointly train competitive reinforcement learning
- Use for adversarial learning by appropriately computing the cross terms
- Use as the optimization algorithm for Generative Adversarial
- Beyond the theoretical guarantees provided, setting different

Code available on github.

AbhijeetiitmVyas/CompetitiveGradientOptim

Pytorch-version coming soon!

Thank you!