Using Bifurcations for Diversity in Differentiable Games

Jonathan Lorraine[1,2,3], Jack Parker-Holder[1,4], Paul Vicol[2,3], Aldo Pacchiano[5], Luke Metz[6], Tal Kachman[7], Jakob Foerster[1]

FAIR[1], University of Toronto[2], Vector Institute[3], University of Oxford[4], UC Berkeley[5], Google Brain[6], Radboud University[7]

Motivation

- Finding different solution types has been useful in minimization - ex., shape vs. texture for CNNs.
- Differentiable games generalize minimization.

$$\boldsymbol{\theta}_A^* \in \arg\min_{\boldsymbol{\theta}_A} \mathcal{L}_A(\boldsymbol{\theta}_A, \boldsymbol{\theta}_B^*),$$
$$\boldsymbol{\theta}_B^* \in \arg\min_{\boldsymbol{\theta}_B} \mathcal{L}_B(\boldsymbol{\theta}_A^*, \boldsymbol{\theta}_B)$$

- Games are increasingly important in ML – ex., GANs, hyperparameter optimization, self-play, meta-learning, adversarial examples, many others.
- Goal: Find diverse solutions in differentiable games – ex., where players work together or battle each-other.

Background

- Ridge rider (RR) [2] finds diverse solution for a single objective by following negative EVals of the Hessian at saddles.
- The Hessian is symmetric with real EVals, so we have conservative dynamics.
- Bifurcations are where small changes cause solution differences.
- Saddles are a key bifurcation in conservative systems.

Our Method – Game Ridge Rider (GRR)

We generalize Ridge Rider to games with the following:

- Complex EVals may have EVecs with complex entries. We use an EVec selection for conjugate pairs that has all real entries, so we can follow it.
- We detect and allow for branching at new types of bifurcations – ex., Hopf where the negative real part of an EVal crosses the imaginary axis.
- We apply an arbitrary optimization algorithm after branching.

Generalizing the Hessian for Games

- The Hessian's generalization for games – i.e., the Game Hessian – may have complex EVals from a lack of symmetry. This gives non-conservative dynamics.

$$\overbrace{\hat{\mathcal{H}}}^{\text{Game Hessian}} = \begin{bmatrix} \text{Player } A \text{ Hessian} \nabla^2_{\boldsymbol{\theta}_A} \mathcal{L}_A & \nabla_{\boldsymbol{\theta}_A} \nabla_{\boldsymbol{\theta}_B} \mathcal{L}_A \\ \nabla_{\boldsymbol{\theta}_B} \nabla_{\boldsymbol{\theta}_A} \mathcal{L}_B^\top & \text{Player } B \text{ Hessian} \nabla^2_{\boldsymbol{\theta}_B} \mathcal{L}_B \end{bmatrix}$$

- In many games – ex., the Iterated Prisoner's Dilemma (IPD) – we are no longer be in a conservative gradient field, allowing more solution and bifurcation types.

New Toy Problems

Following the gradient in red, LOLA [1] in blue

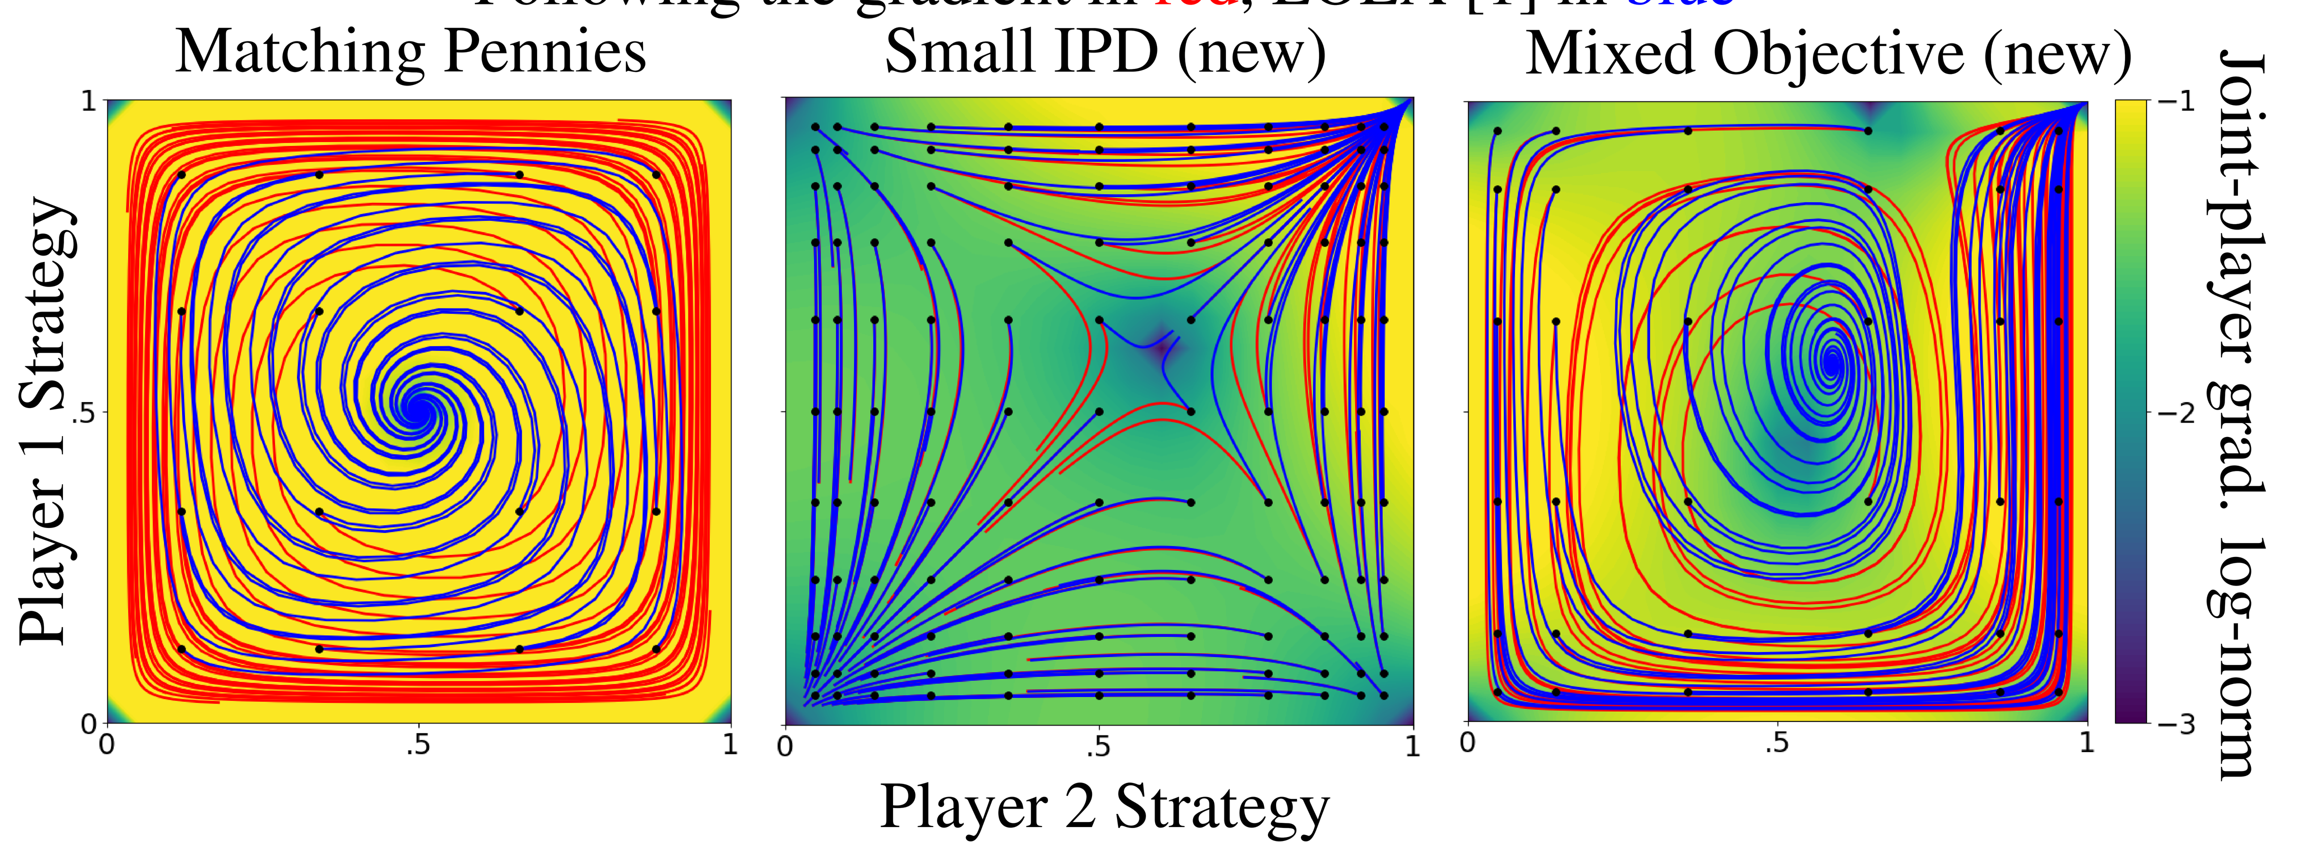

- Matching Pennies is a 2 param. game with imaginary EVals, but only 1 solution.
- Small IPD is a 2 param. IPD with TT and DD solutions, but only real EVals.
- Mixing these gives a 2 param. problem like the full IPD with multiple solutions, complex EVals, and a Hopf bifurcation.

Applying our Method on Toy Problems

- For both the small IPD (left) and mixed objective (right) our method – Game Ridge Rider (GRR) – finds all solutions.

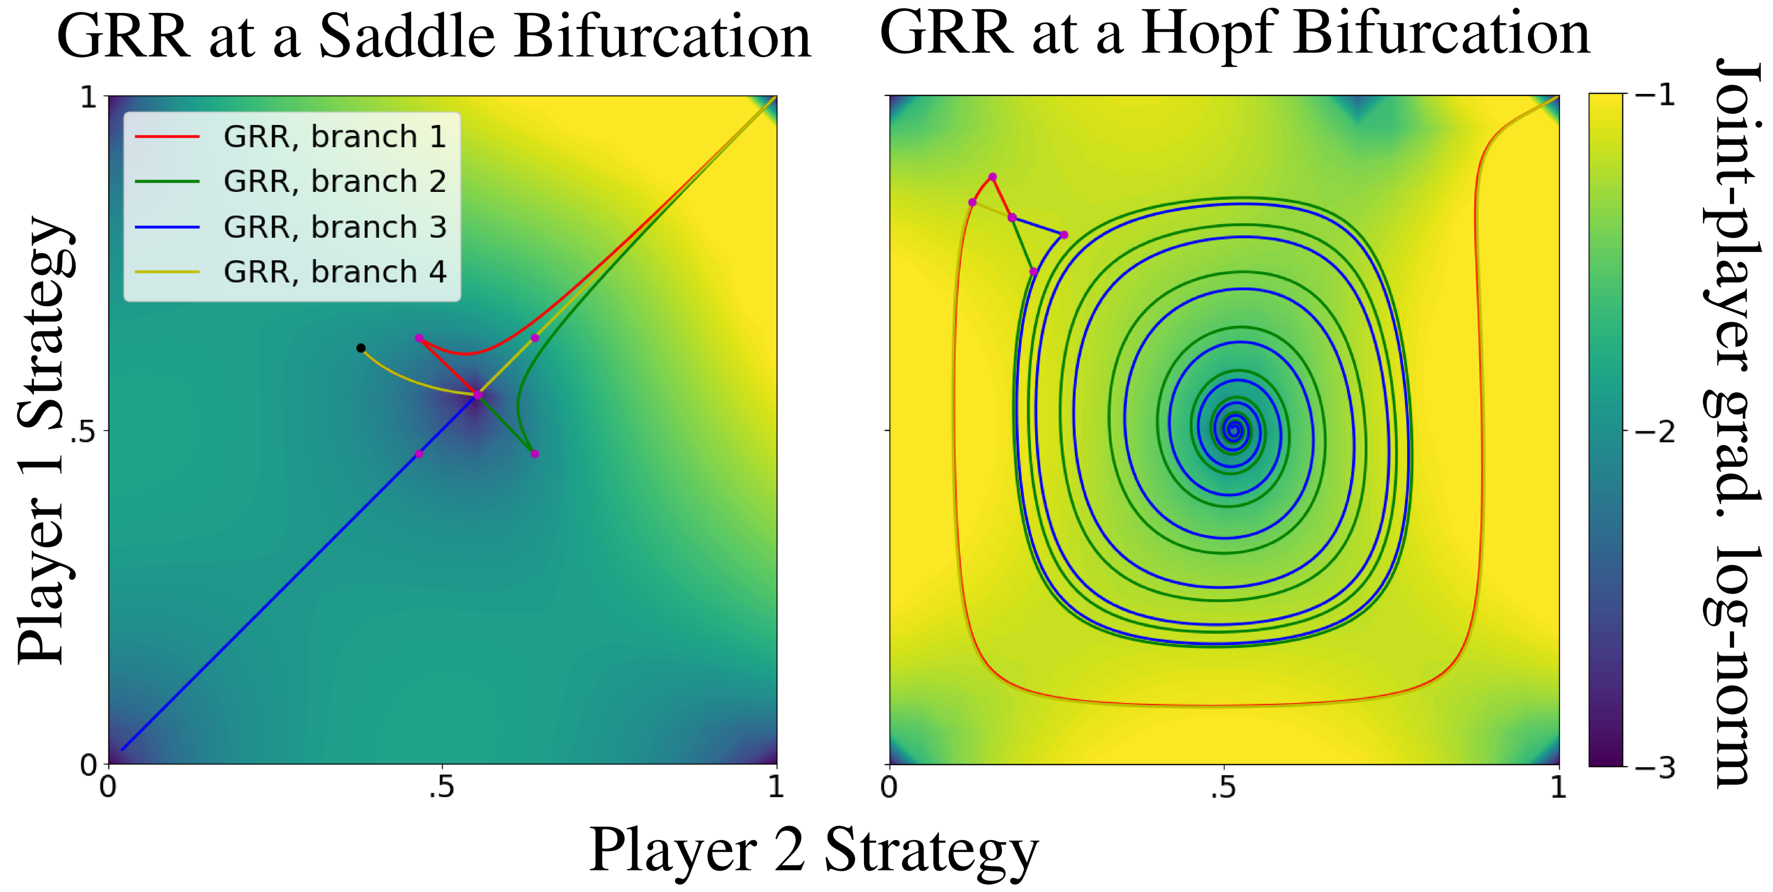

Finding Diverse Solutions in the Iterated Prisoners Dilemma (IPD)

Search Strategy	Coop.	Defect
×20 Rand init + LOLA [1]	✓	×
×20 Rand init + follow grad.	×	✓
Ours: Saddle + branch	✓	✓
Rand init + branch	×	✓

- Randomly initializing then applying a training method only finds 1 solution mode.
- Our method finds both solution modes.
- If we don't start at a saddle, then branching doesn't affect the solution.

Takeaways

- Differentiable games generalize minimization, but with non-conservative dynamics from complex EVals.
- We can view methods for diverse solutions in minimization – i.e., Ridge Rider (RR) – as finding bifurcations in conservative systems and branching.
- The viewpoint allows usage of tools from dynamical systems to generalize RR to non-conservative setups.
- Our method generalizes RR by branching at Hopf bifurcations and applying arbitrary optimizers after branching.

References

[1] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. In International Conference on Autonomous Agents and MultiAgent Systems, pages 122–130, 2018.

[2] Jack Parker-Holder, Luke Metz, Cinjon Resnick, Hengyuan Hu, Adam Lerer, Alistair Letcher, Alexander Peysakhovich, Aldo Pacchiano, and Jakob Foerster. Ridge rider: Finding diverse solutions by following eigenvectors of the hessian. In Advances in Neural Information Processing Systems, volume 33, pages 753–765, 2020.