DeepMind Google UCL

Discretization Drift in Two-Player Games

Mihaela Rosca, Yan Wu, Benoit Dherin, David Barrett

Quantify discretization drift in two-player games trained using gradient descent.

Quantify discretization drift in two-player games trained using gradient descent.

→

Use it as a framework to understand and improve two-player games (e.g. GANs).

Modified ODEs which better describe the discrete updates.

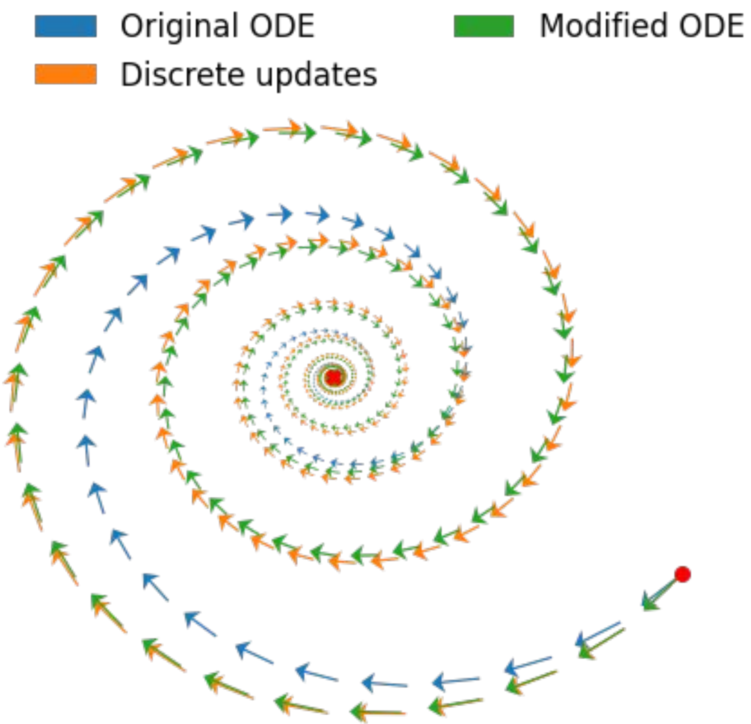

Modified ODEs which better describe the discrete updates.

- Stability analysis.
- Explicit regularization to cancel harmful forms of drift and improve GAN training.
- Understand the difference between simultaneous and alternating gradient descent.

Two-player games

$$\dot{\phi} = f(\phi, \theta)$$

$$\dot{\theta} = g(\phi, \theta)$$

Two-player games

$$\dot{\phi} = f(\phi, \theta)$$

$$\dot{\theta} = g(\phi, \theta)$$

discretization

discretization

Simultaneous Euler updates

Alternating Euler updates

Backward error analysis

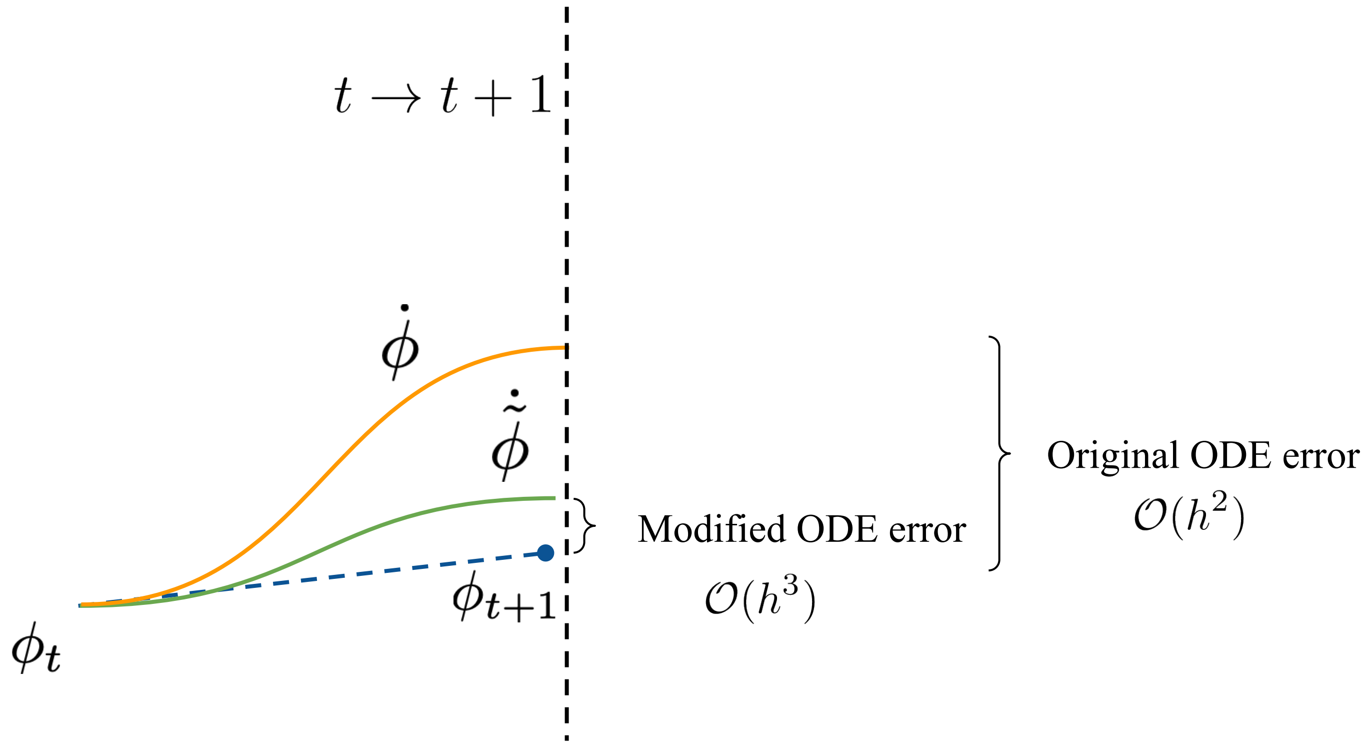

Approach introduced in supervised learning by
David Barrett and Benoit Dherin: *Implicit gradient regularization, ICLR 2021.*

Modified ODEs: Simultaneous Euler updates

Discrete dynamics **Backward error analysis** **Continuous dynamics**

$$\phi_t = \phi_{t-1} + \alpha h f(\phi_{t-1}, \theta_{t-1})$$
$$\theta_t = \theta_{t-1} + \lambda h g(\phi_{t-1}, \theta_{t-1})$$

$$\dot{\tilde{\phi}} = f(\phi, \theta) - \frac{\alpha h}{2}\left(f \nabla_\phi f + g \nabla_\theta f\right)$$
$$\dot{\tilde{\theta}} = g(\phi, \theta) - \frac{\lambda h}{2}\left(g \nabla_\theta g + f \nabla_\phi g\right)$$

Self and interaction terms.

Modified ODEs: Alternating Euler updates

Discrete dynamics **Backward error analysis** **Continuous dynamics**

$$\phi_t = \phi_{t-1} + \alpha h f(\phi_{t-1}, \theta_{t-1})$$
$$\theta_t = \theta_{t-1} + \lambda h g(\phi_t, \theta_{t-1})$$

$$\dot{\tilde{\phi}} = f(\phi, \theta) - \frac{\alpha h}{2}(f \nabla_\phi f + g \nabla_\theta f)$$
$$\dot{\tilde{\theta}} = g(\phi, \theta) - \frac{\lambda h}{2}\left(g \nabla_\theta g + (1 - \frac{2\alpha}{\lambda}) f \nabla_\phi g\right)$$

Self and interaction terms.

Stability analysis using the modified ODEs

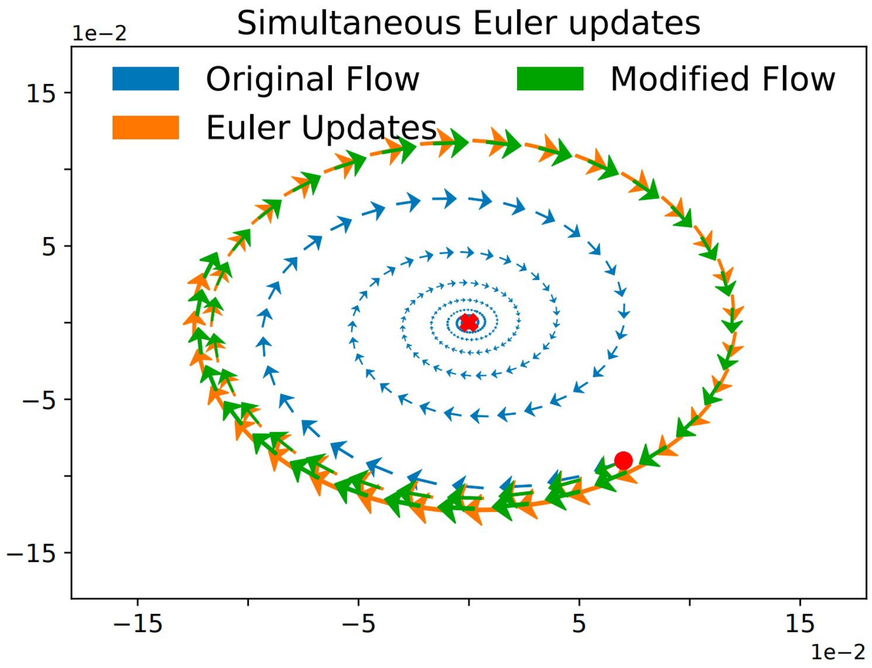

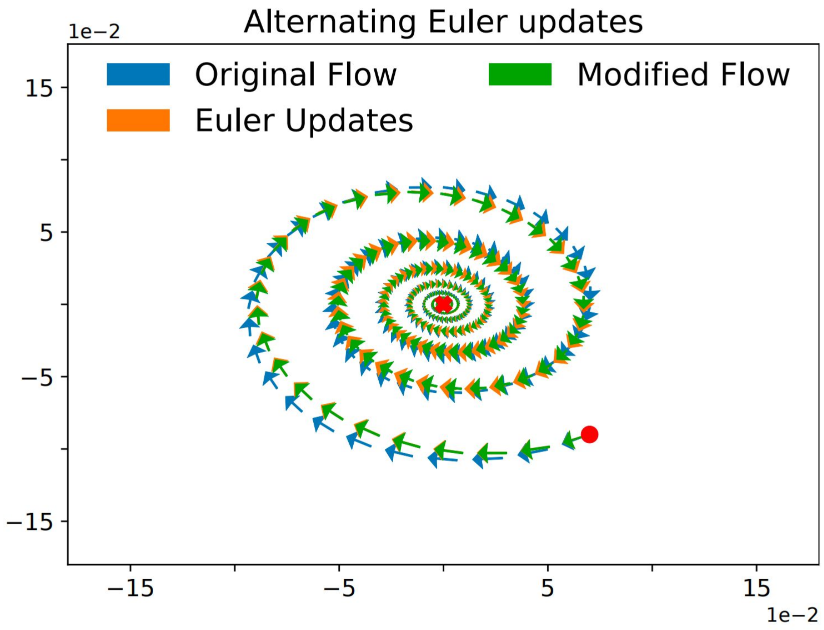

Stability analysis using the modified ODEs

- Does not ignore discretization drift
- Accounts for the different behaviour of simultaneous and alternating updates

Zero-sum games

$$L_f = -E$$

$$L_g = E$$

$$f = \nabla_\phi E$$

$$g = -\nabla_\theta E$$

Simultaneous gradient descent in zero-sum games

$$\tilde{L}_\phi = -E + \frac{\alpha h}{4} \left\|\nabla_\phi E\right\|^2 - \frac{\alpha h}{4} \left\|\nabla_\theta E\right\|^2$$

$$\tilde{L}_\theta = E + \frac{\lambda h}{4} \left\|\nabla_\theta E\right\|^2 - \frac{\lambda h}{4} \left\|\nabla_\phi E\right\|^2$$

The *interaction terms* maximize the gradient norm of the other player.

Simultaneous gradient descent in zero-sum games

$$\tilde{L}_\phi = -E + \frac{\alpha h}{4}\left\|\nabla_\phi E\right\|^2 -$$ [redacted] $$E\|^2$$

$$\tilde{L}_\theta = E + \frac{\lambda h}{4}\left\|\nabla_\theta E\right\|^2 -$$ [redacted] $$\|^2$$

GAN training
Higher is better

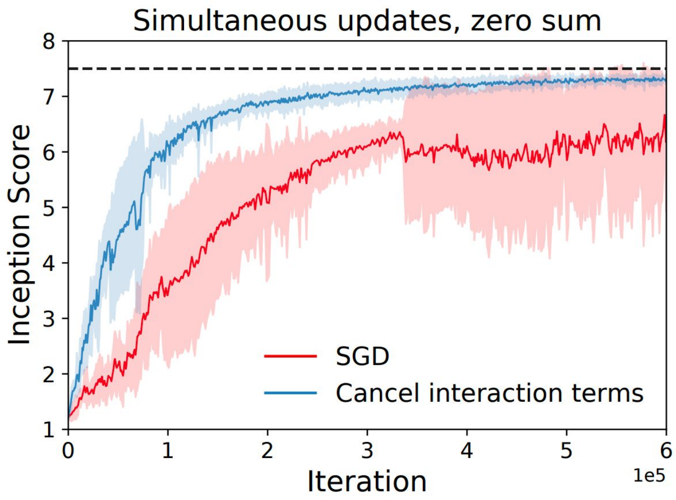

Explicit regularization cancelling the *interaction terms* improves performance

Alternating gradient descent in zero sum games

$$\tilde{L}_\phi = -E + \frac{\alpha h}{4}\left\|\nabla_\phi E\right\|^2 - \frac{\alpha h}{4}\left\|\nabla_\theta E\right\|^2$$

$$\tilde{L}_\theta = E + \frac{\lambda h}{4}\left\|\nabla_\theta E\right\|^2 - \frac{\lambda h}{4}\left(1 - \frac{2\alpha}{\lambda}\right)\left\|\nabla_\phi E\right\|^2$$

learning rate ratio

GAN training
Higher is better

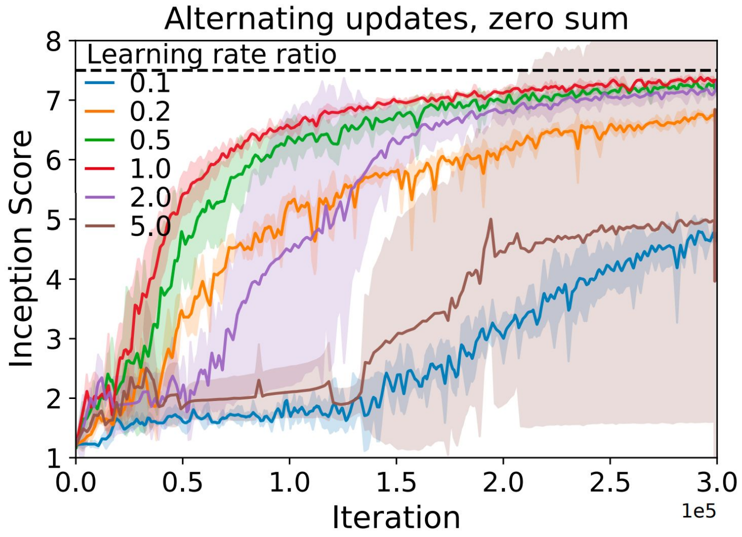

We can predict which learning rate ratios will perform best!

Quantifying discretization drift in two-player games via modified ODEs is a useful theoretical and empirical tool.