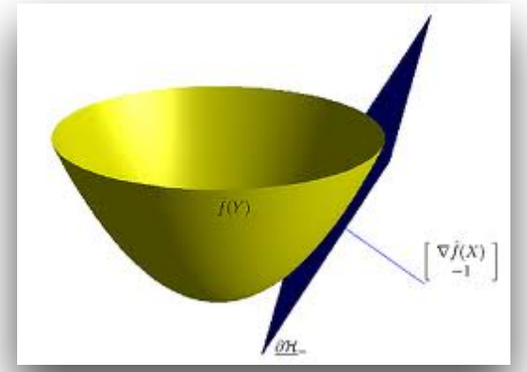

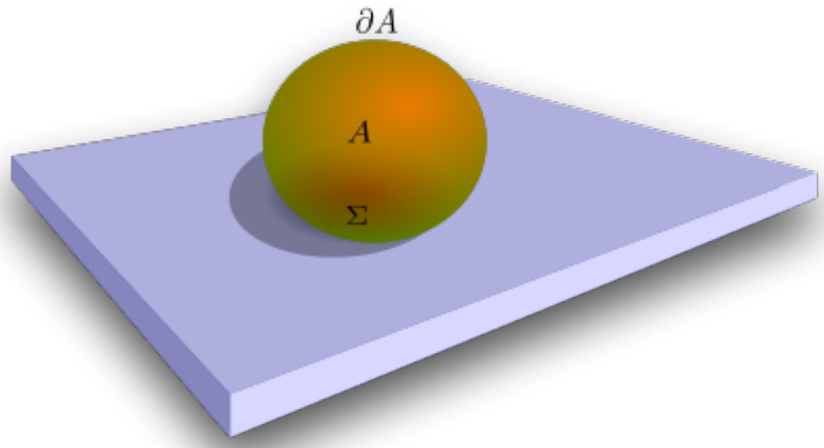

AAMAS 2020 Tutorial: Network Economics, Variational Inequalities and Deep Learning

Sridhar Mahadevan

Adobe Research

UMassAmherst
College *of* Information and Computer Sciences

Plan of the Tutorial

- **Part I: Motivation**
- Part II: Mathematical Background
- Part III: Algorithms
- Part IV: Applications and Challenges

Acknowledgments

- Anna Nagurney: Professor @ UMass Management School, has written many books and papers on VIs

- Boyd: Lecture notes on optimization

- UMass Amherst PhD students:

  - Ian Gemp (now @DeepMind): developed new methods for solving VIs, and carried out computational experiments

  - Other students: Bo Liu, Philip Thomas, Thomas Boucher, Clifton Carey, Nicholas Jacek, William Dabney

References

Downloaded 02/04/14 to 128.119.40.251. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

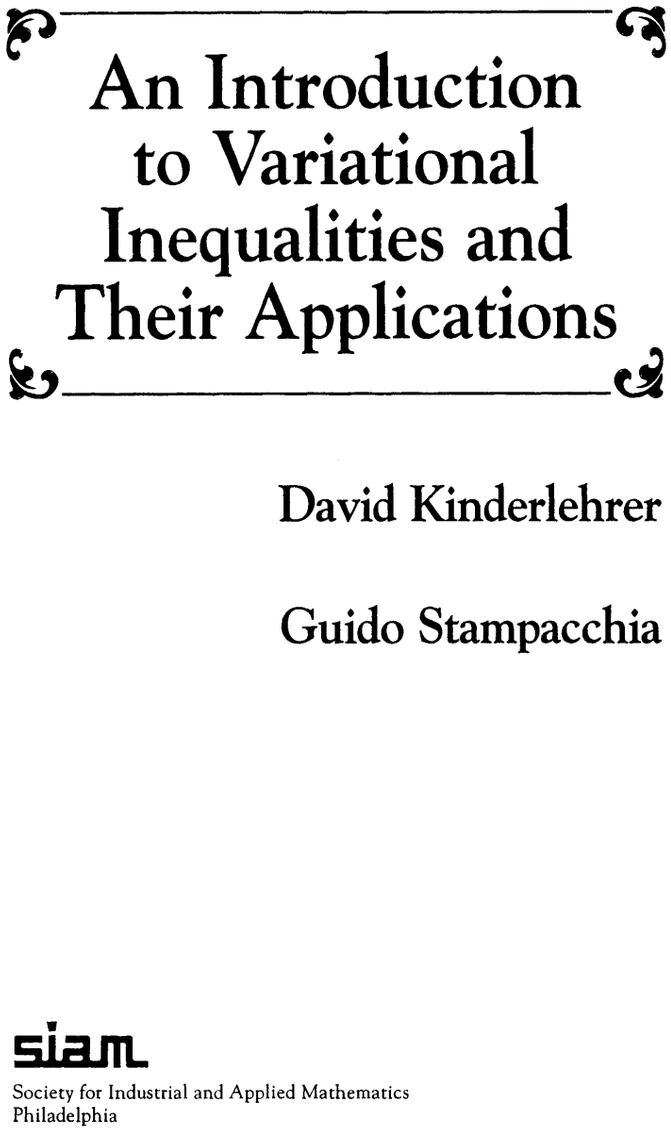

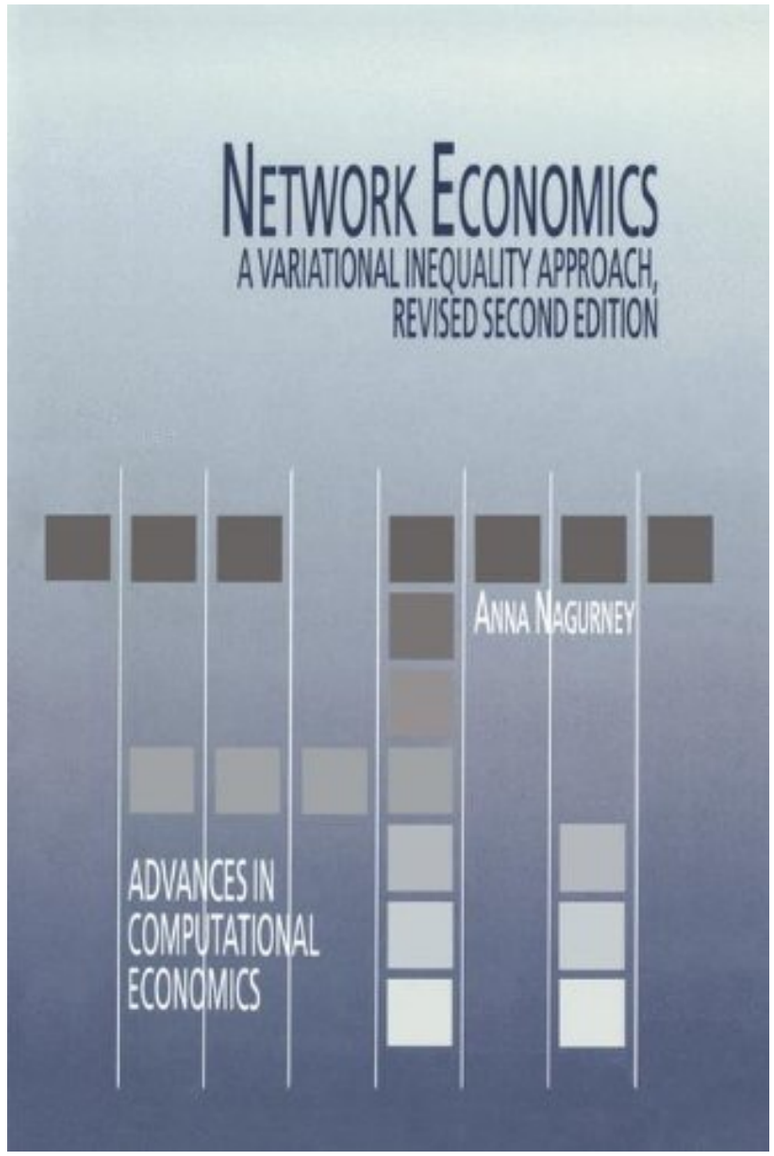

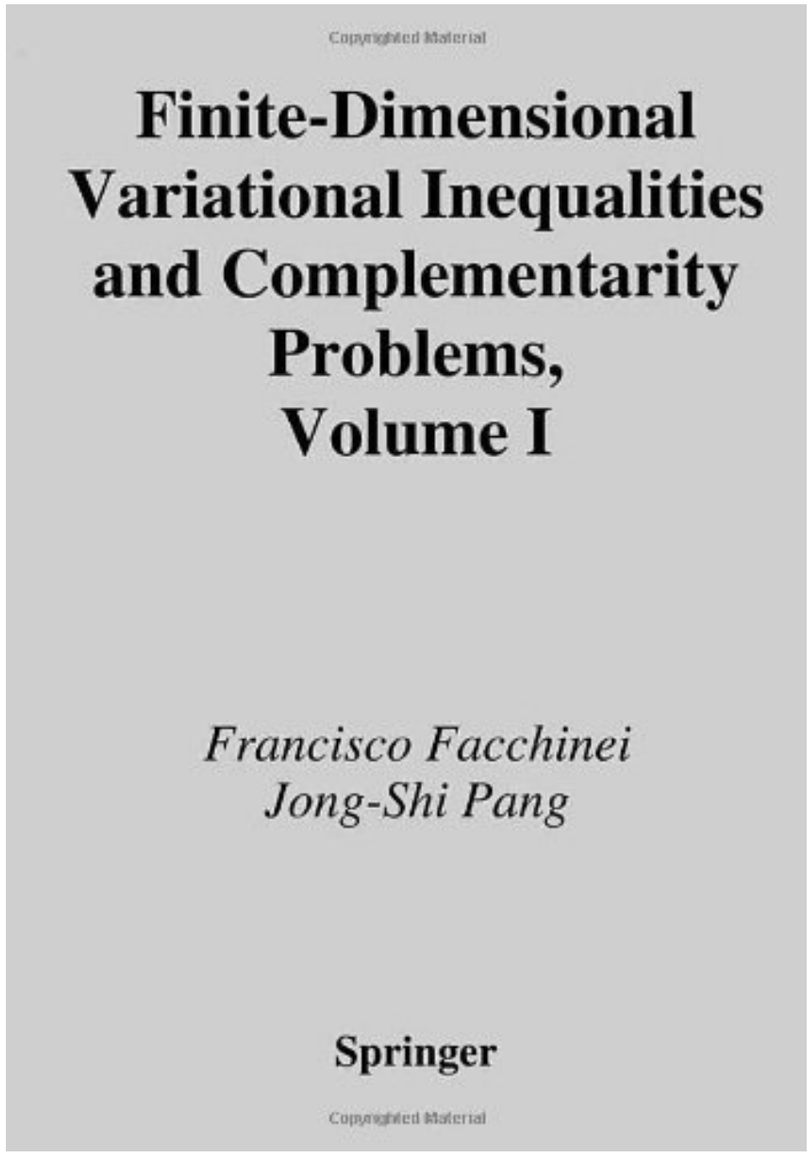

AI in the 21st Century

- AI is facing a “networked” future
  - Computation in the “cloud” by the “crowd”
  - Heterogeneous data storage in the “cloud”
- Competing objectives
  - Privacy, reliability, security
- Optimization may be too restrictive

Adobe Experience Cloud

Analytics Cloud

- Meaningful Measurement
- Precise Audience Creation
- Instant Insights In Action

Analytics

Audience Manager

Advertising Cloud

- Cross-channel
- Performance without compromise
- Independent

Media Optimizer DSP

Media Optimizer Search

Media Optimizer Creative

Marketing Cloud

- Manage Experiences At Scale
- Personalize Each Experience
- Orchestrate The Journey

Experience Manager

Campaign

Target

Primetime

© 2017 Adobe Systems Incorporated. All Rights Reserved. Adobe Confidential.

Adobe

Data Science driven Experiences

DELIVERY

> 50 Trillion

Activities tracked

DECISIONING & ALGORITHMS

> 5 Trillion

Dynamic campaign assets delivered

CONSUMER PROFILE

CRM

CONTENT

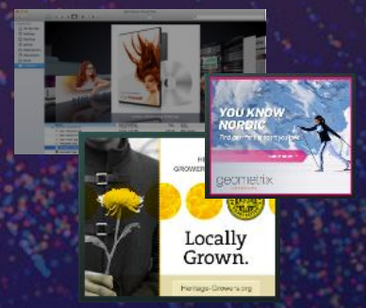

100 Petabytes

Data managed

© 2017 Adobe Systems Incorporated. All Rights Reserved. Adobe Confidential.

Adobe

Distributed Optimization via ADMM
(Boyd et al., ML FT 2010)

"ADMM was developed over a generation ago, with its roots stretching far in advance of the Internet, distributed and cloud computing systems, massive high-dimensional datasets, and associated large-scale applied statistical problems. Despite this, it appears well-suited to the modern regime."

Optimization to Equilibration

- Optimization forms a foundation of much of AI
  - Machine learning, NLP, planning, robotics vision
- Optimization has significant limitations
  - Limited types of utility functions, assumes symmetry of the Hessian etc.
- A new foundation for multiagent AI: **equilibration**
  - Generalizes optimization, covers many other problems

Optimization

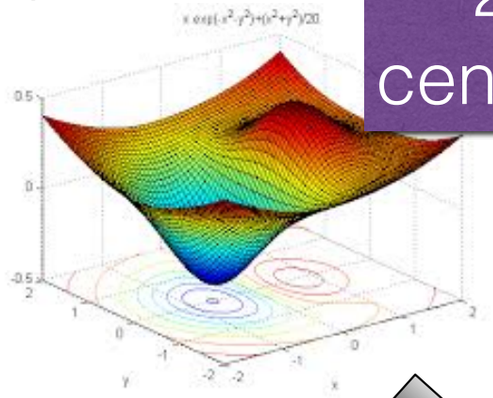

20th century AI

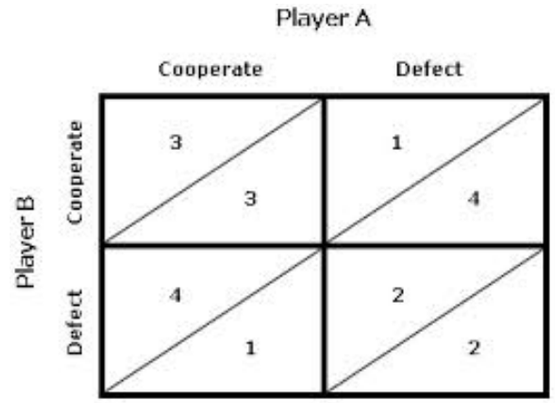

Game theory

21st century AI

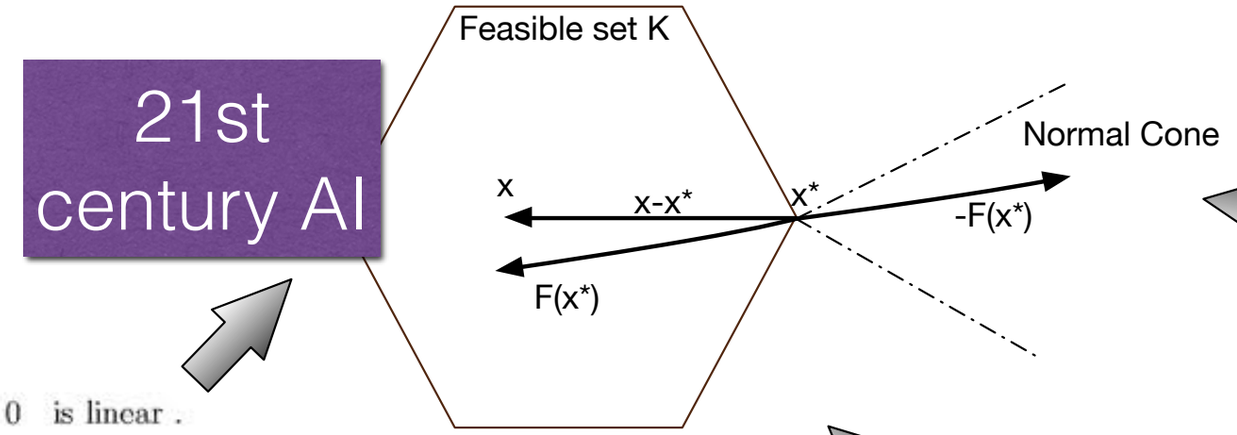

Variational Inequalities

Complementarity problems

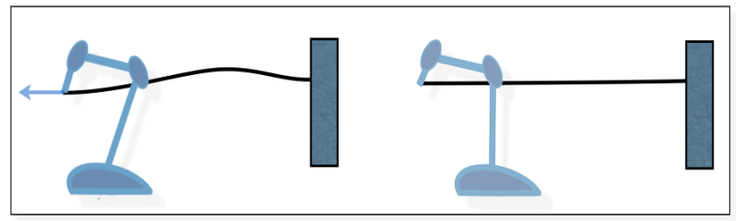

$$\frac{\partial u}{\partial x_1} + \frac{\partial u}{\partial x_2} = 0 \quad \text{is linear .}$$

$$\frac{\partial u}{\partial x_1} + \left(\frac{\partial u}{\partial x_2}\right)^2 = 0 \quad \text{is nonlinear .}$$

$$\frac{\partial u}{\partial x_1} + \frac{\partial u}{\partial x_2} + u^2 = 0 \quad \text{is nonlinear .}$$

$$\frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2} = x_1 \quad \text{is linear .}$$

$$\frac{\partial^2 u}{\partial x_1^2} + u\frac{\partial^2 u}{\partial x_2^2} = 0 \quad \text{is quasilinear .}$$

Nonlinear equation solving

Traffic equilibrium problem

Equilibria in Economics

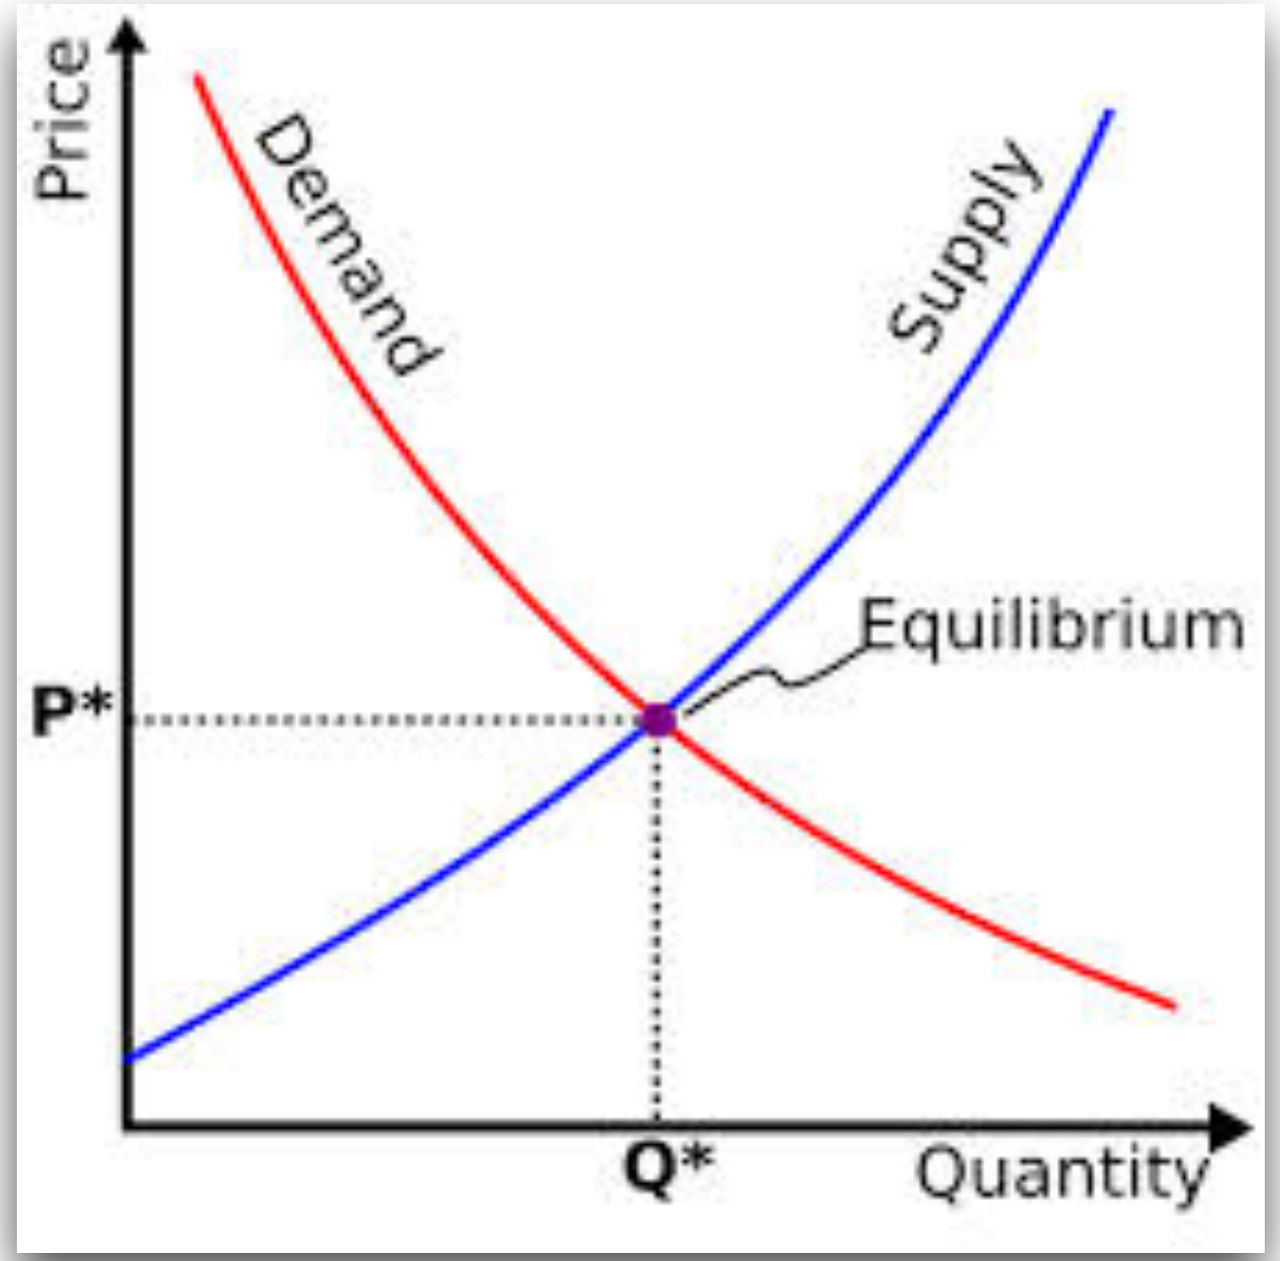

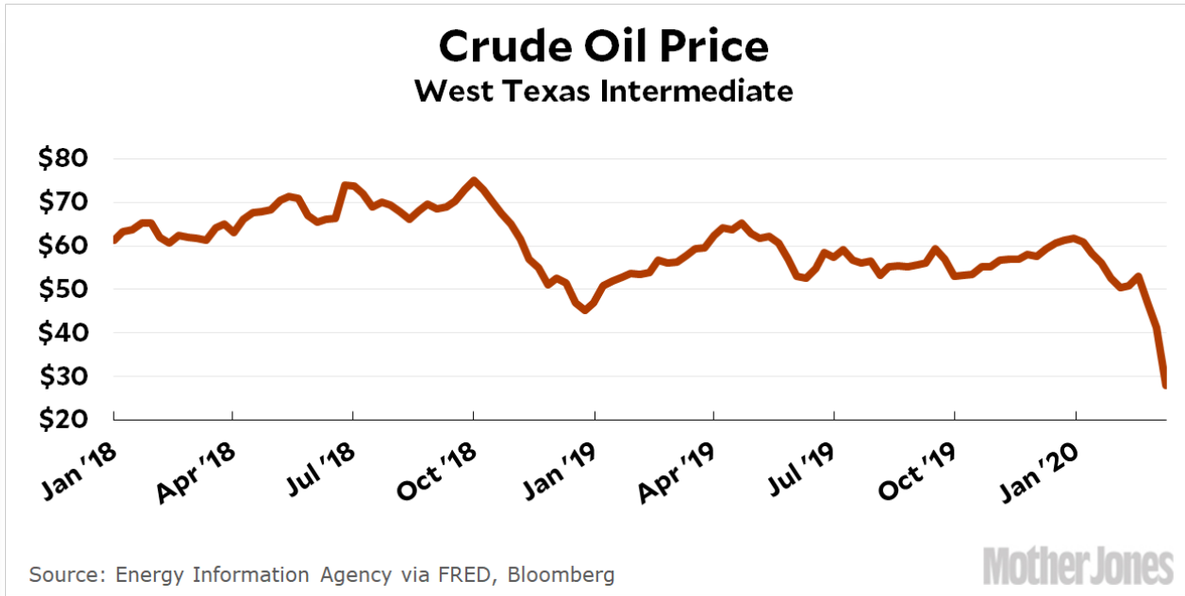

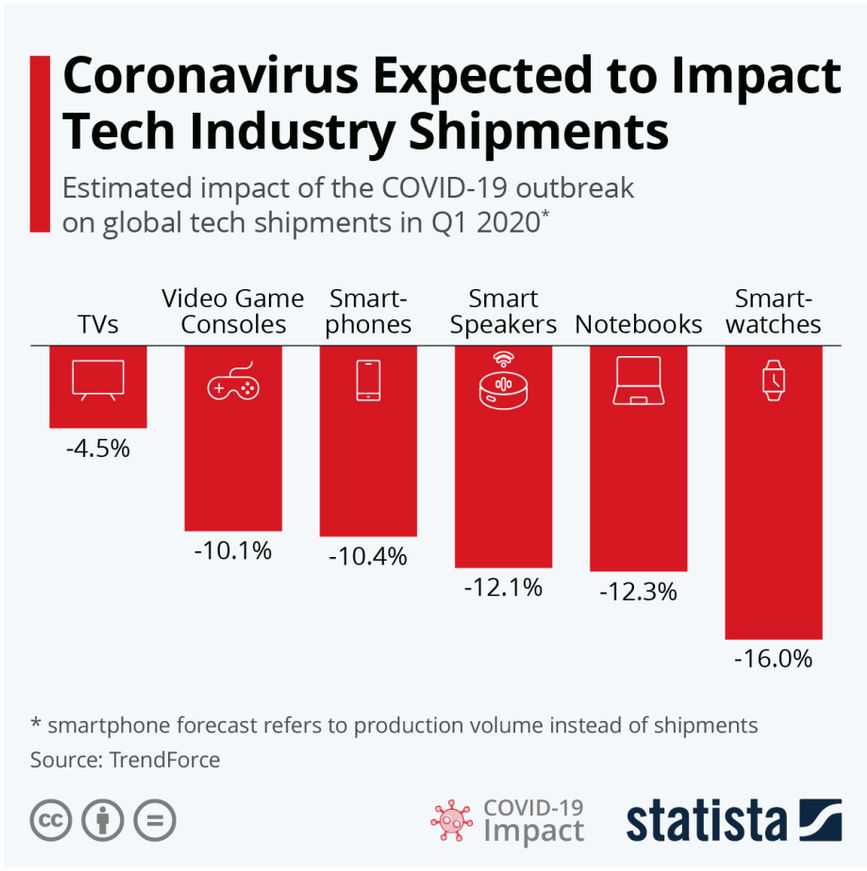

Some winners in this era

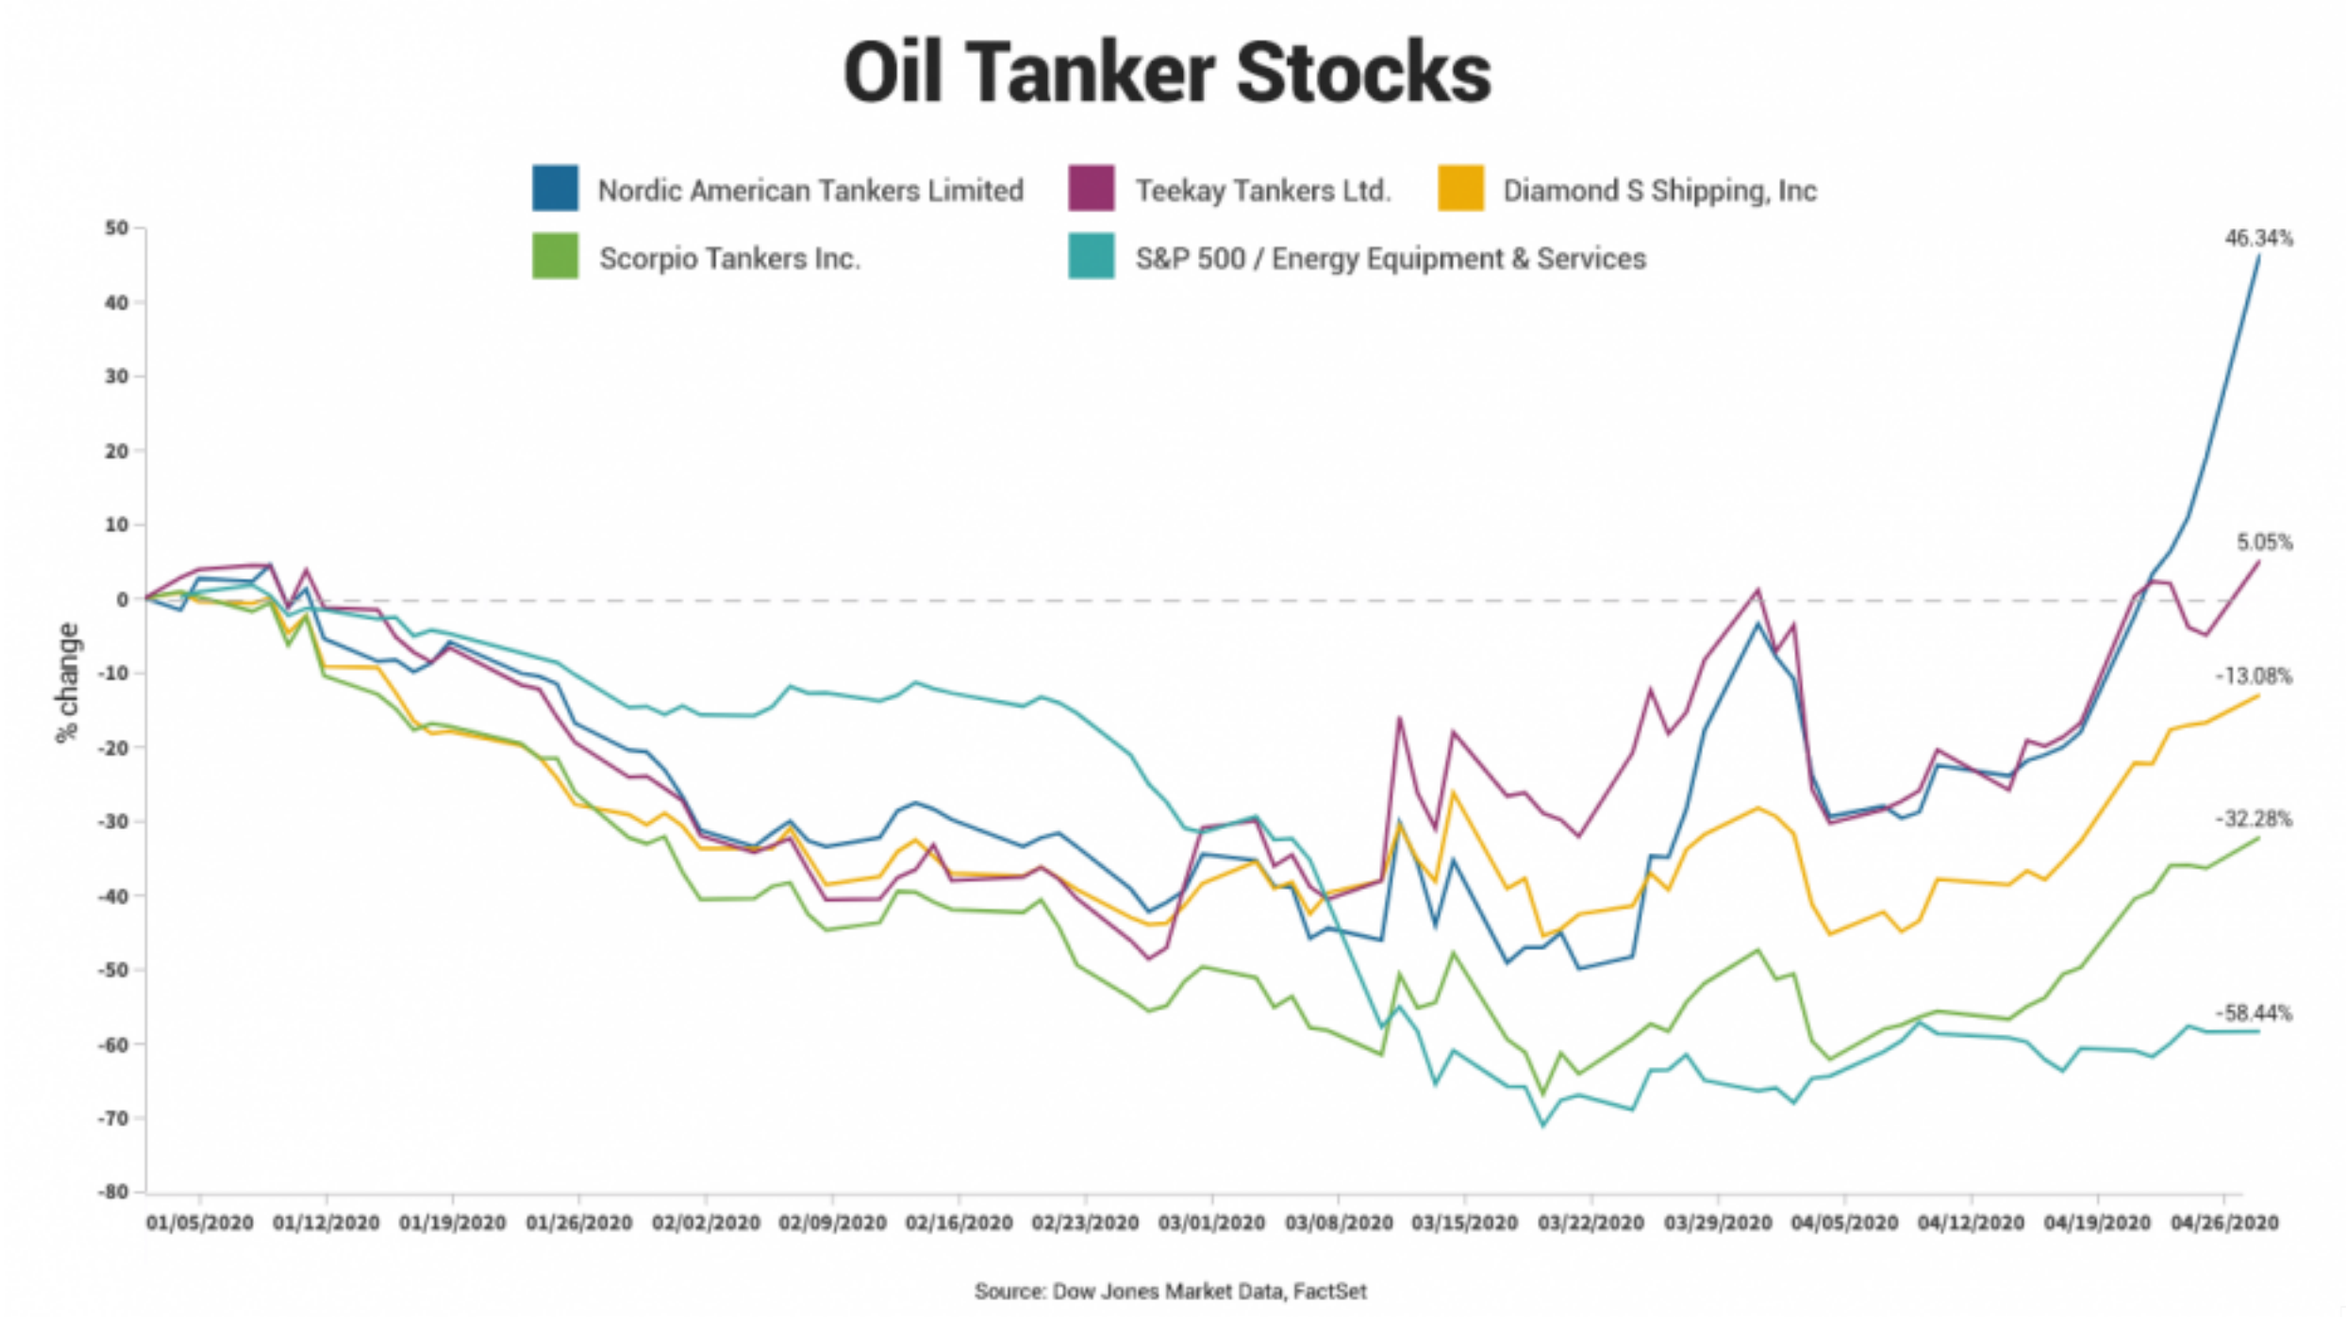

“The Internet is an equilibrium — we just have to identify the game (Scott Shenker)”

“The Internet was the first computational artifact that was not created by a single entity, but emerged from the strategic interaction of many (Christos Papadimitriou)”

The “Invisible Hand” of the Internet

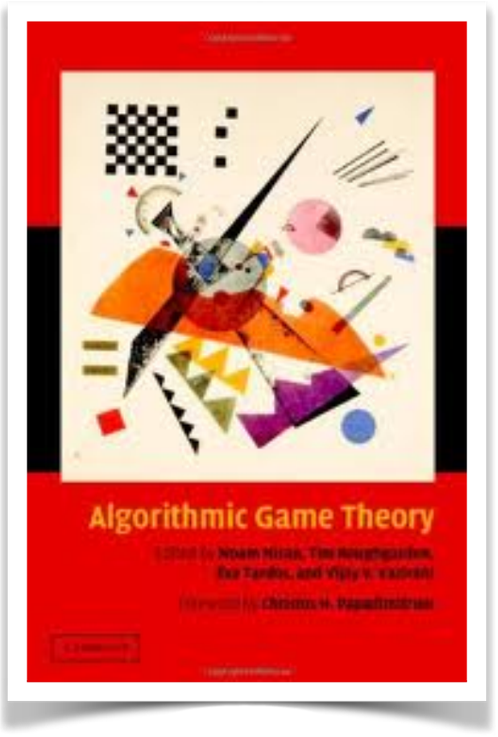

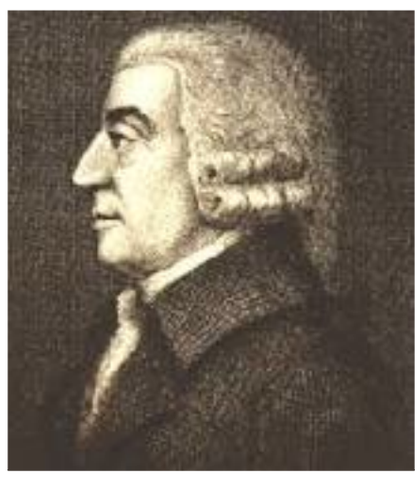

Adam Smith
The Wealth of Nations
1776

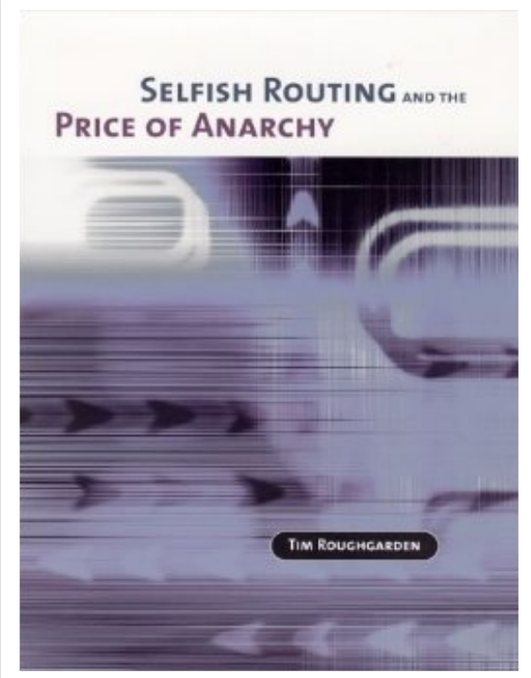

Competing Goals of the Internet:1992-2014

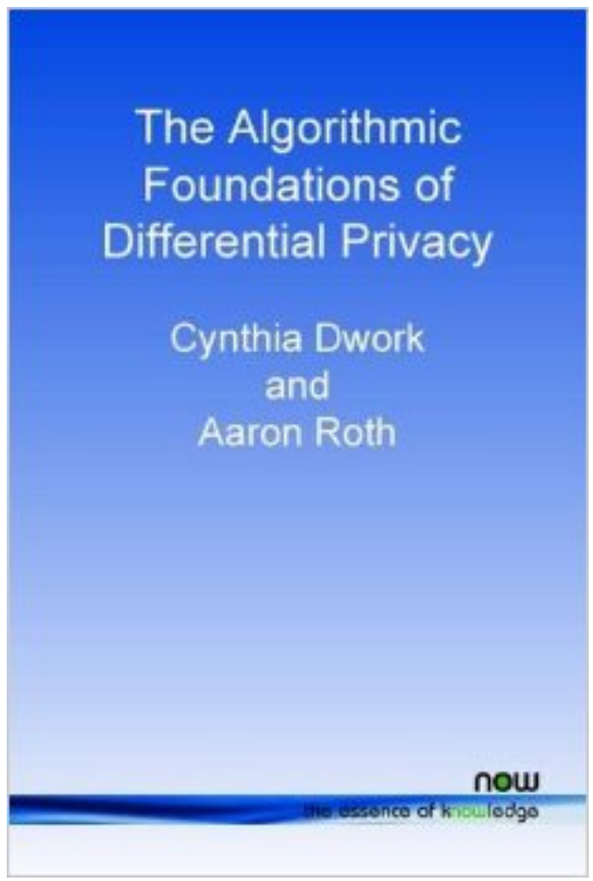

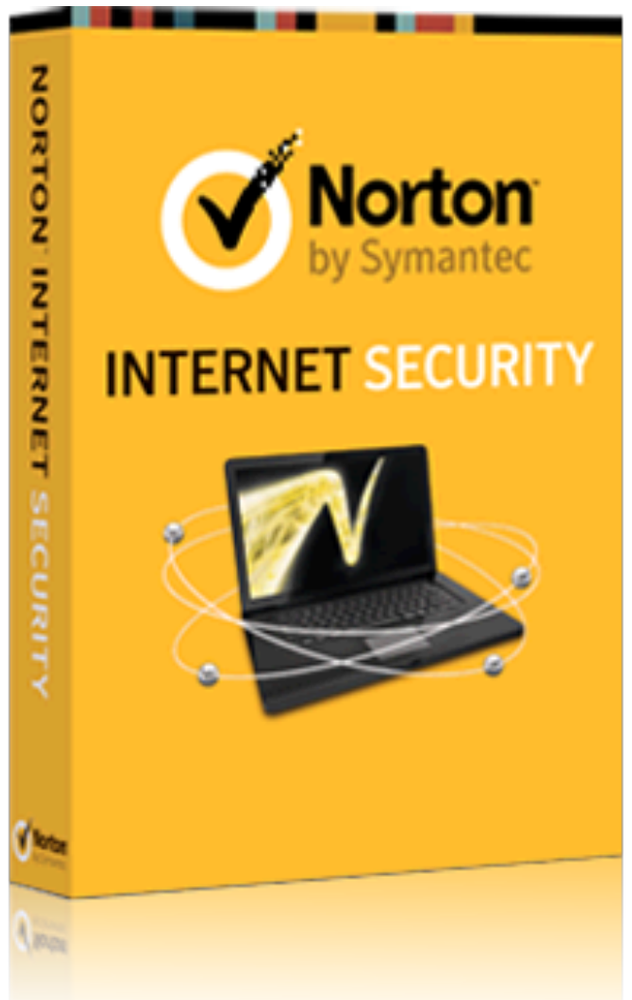

Blood Bank System

39,000 donations/day in the US

Often, only a 2 day supply

American Red Cross

In 2007, out of 1700 hospitals,
492 cancelled surgeries
due to blood shortages

The need is constant.
The gratification is instant.
Give blood.™

1-800-GIVE-LIFE | RedCrossBlood.org

Blood Distribution

- Multicriterion optimization problem
  - Determine optimal levels of blood processed
    - subject to minimization of total cost of blood collection, shipment, processing and testing, in addition to total discarding cost
    - as well as minimization of total supply risk
    - subject to uncertain demand being satisfied

Traffic Equilibrium Problem

Braess Paradox

Traffic Network Equilibrium

(Dafermos, Nagurney)

- Link travel cost functions
  - $C_a(F_a) = 10*F_a$
  - $C_b(F_b) = F_b + 50$
  - $C_c(F_c) = F_c + 50$
  - $C_d(F_d) = 10\ F_d$
- Travel demand $D_{14} = 6$
- Find equilibrium flows

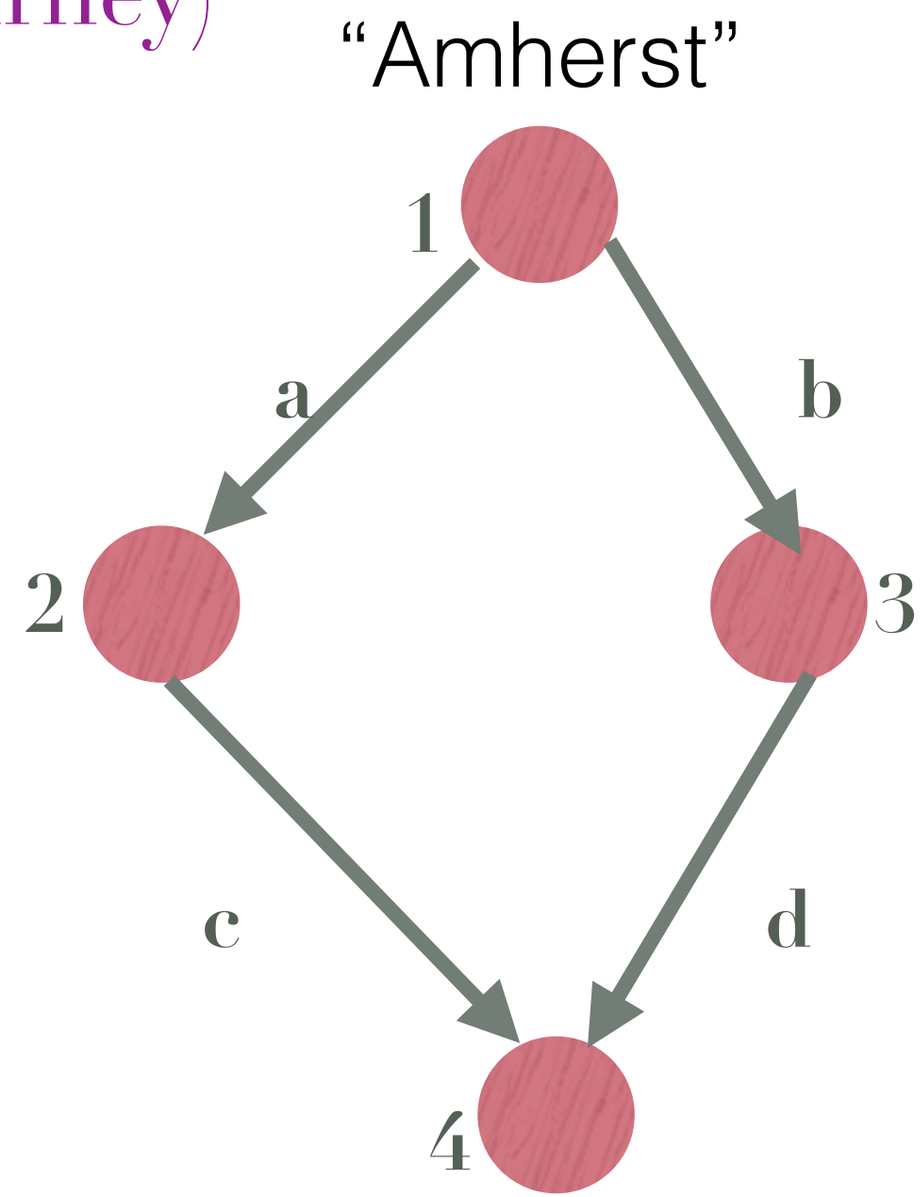

Traffic Network Equilibrium

(Dafermos, Nagurney)

- Flows at equilibrium
  - $F_a = F_b = 3$
  - $F_c = F_d = 3$
  - $C_a = 30, C_b = 53$
  - $C_c = 53, C_d = 30$
- Path costs = 83
- Nash equilibrium

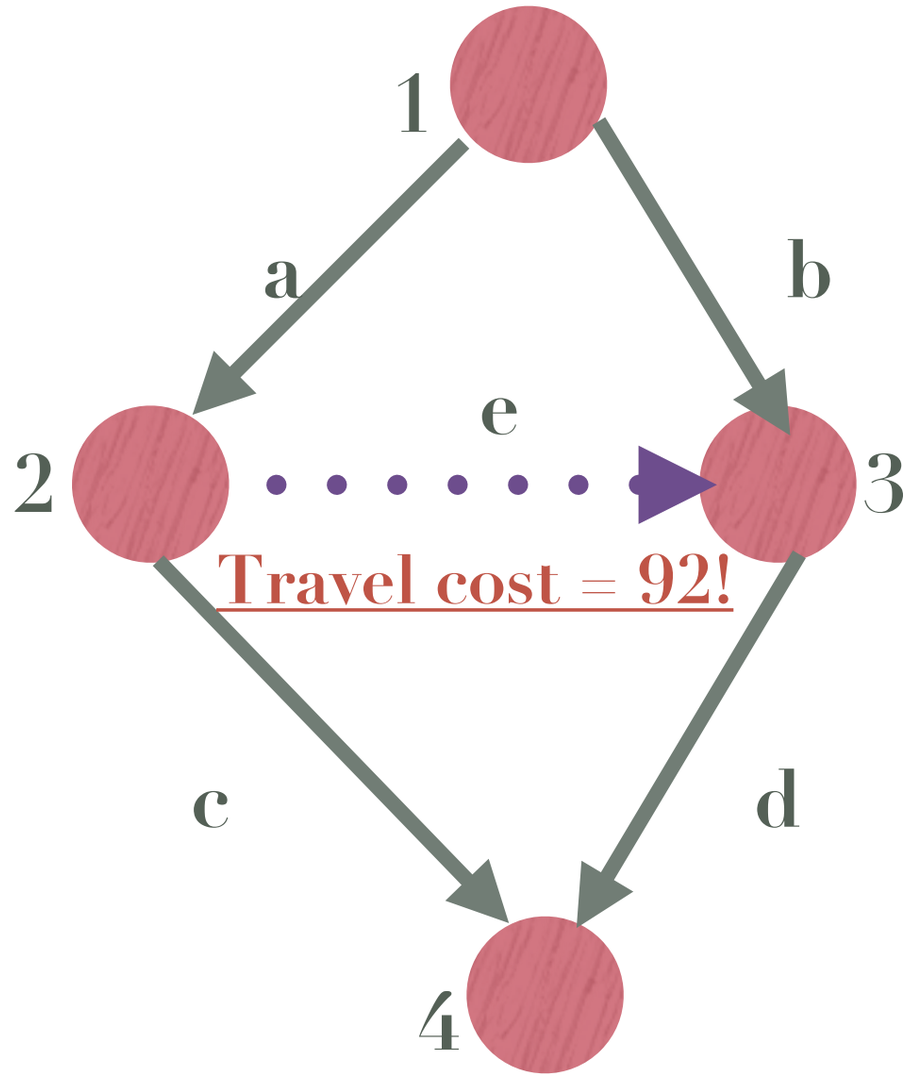

Braess' Paradox!

Guardian
March 2014

Seoul demolishes its urban expressways as city planners opt for greener schemes

The levelling of the Ahyeon overpass marks the latest attempt to improve the environment of South Korea's sprawling capital

Expressway in south-east Seoul. Photograph: Lo Mak/Corbis

Real-life examples of Braess Paradox

Seoul, South Korea in 2002, “they had over years built a six lane highway,” said Nagurney. Always jammed, “the highway carried over a hundred sixty thousand cars per day.” It came down. Result: Better traffic flow.

close to Hongik, a vibrant student neighbourhood in the city centre. In those days, the city council was led by Kim Hyun-ok. Aged only 39 when he got the job, he soon came to be known as Bulldozer Kim. He was often seen on building sites,

IMAGINATION IN ART

Cy Twombly
NY Times, April 11 2018

Jean Michel Basquiat
Sold in NY for $110M

Generative Adversarial Networks

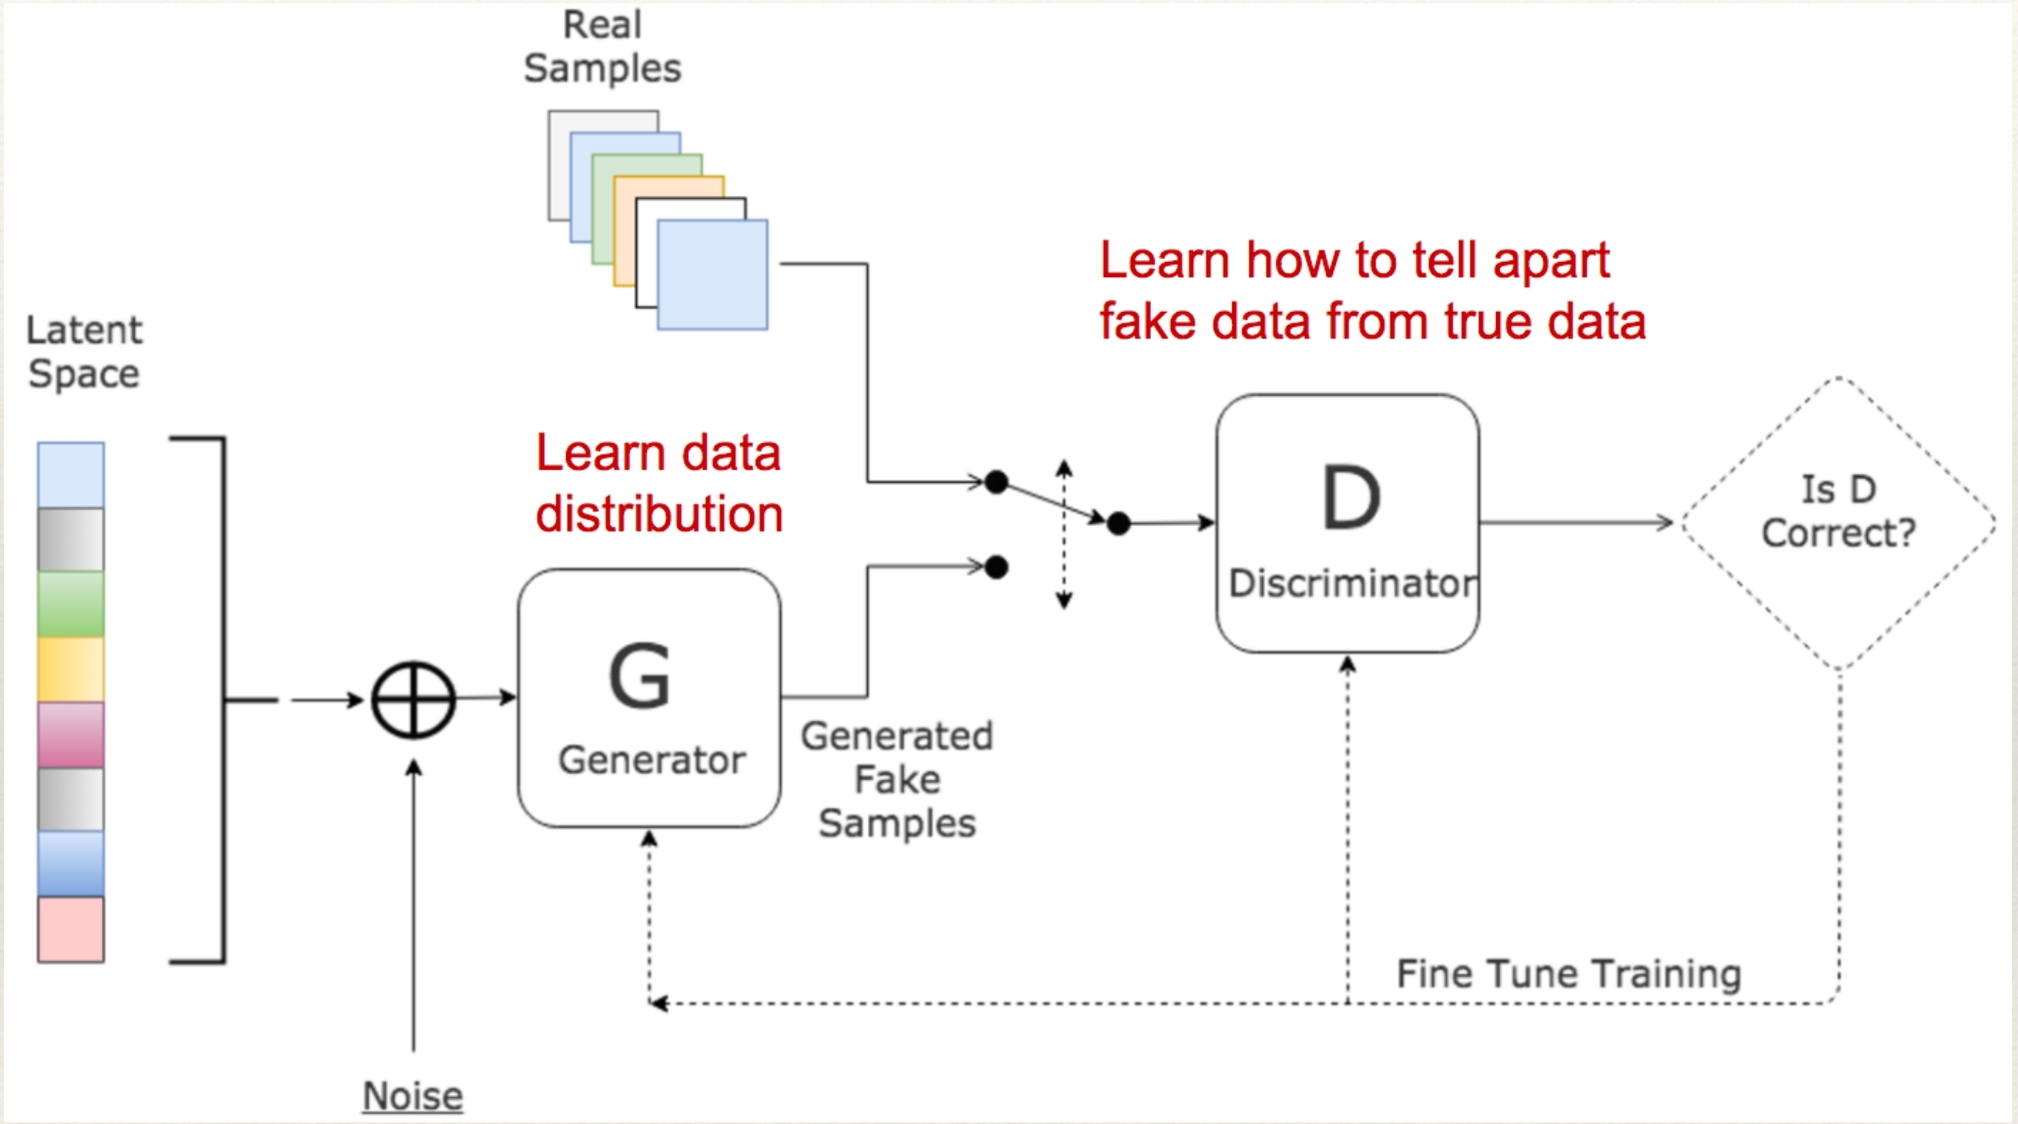

www.kdnuggets.com/2017/01/generative-…-learning.html

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}}[\log D(x)] +$$

$$\mathbb{E}_{z\sim p_z}[\log(1 - D(G(z)))]$$

Noise ~ N(0,1)

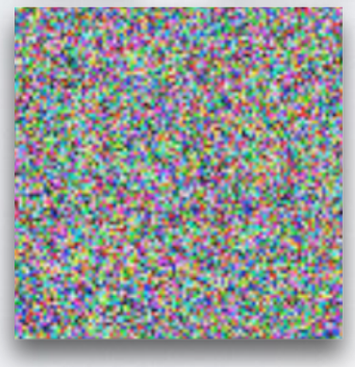

Generative Model

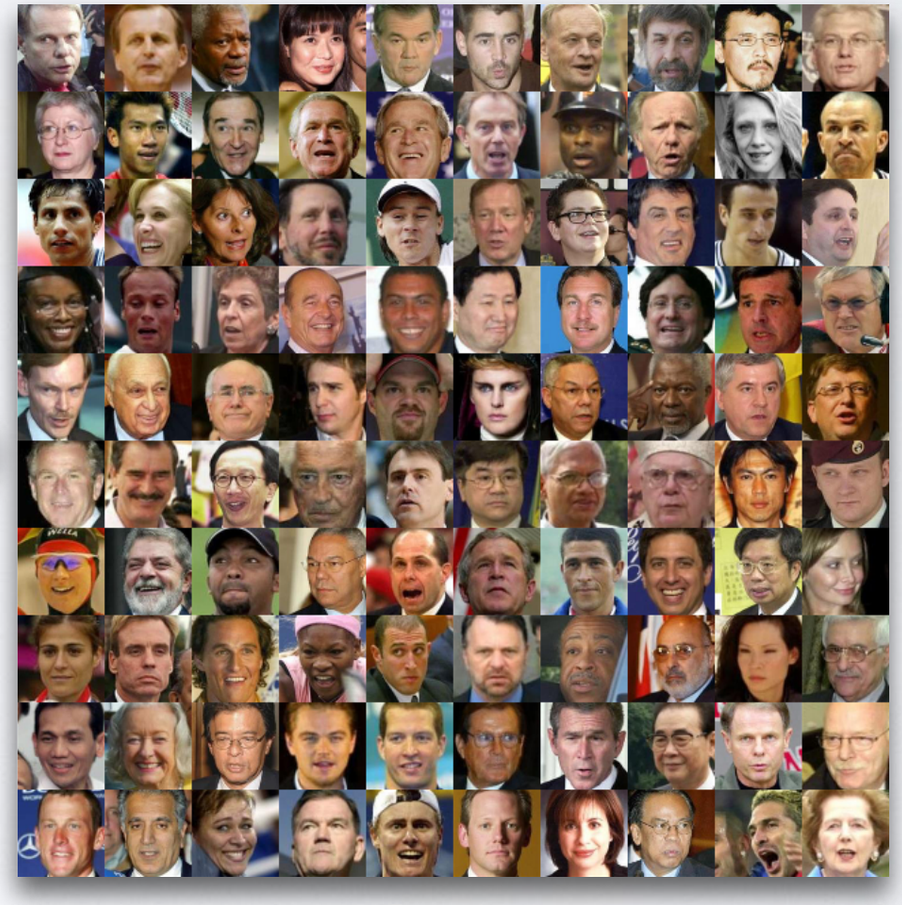

GENERATIVE MULTI-ADVERSARIAL NETWORKS (GMAN)

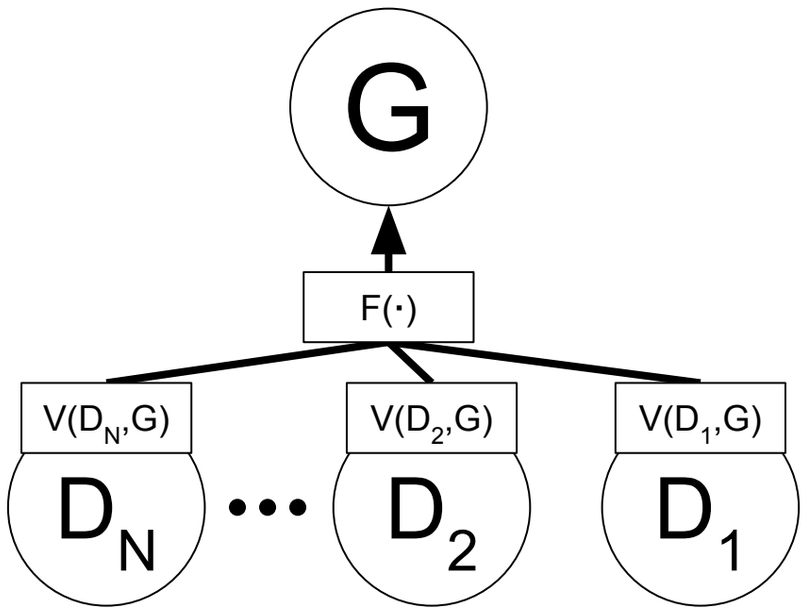

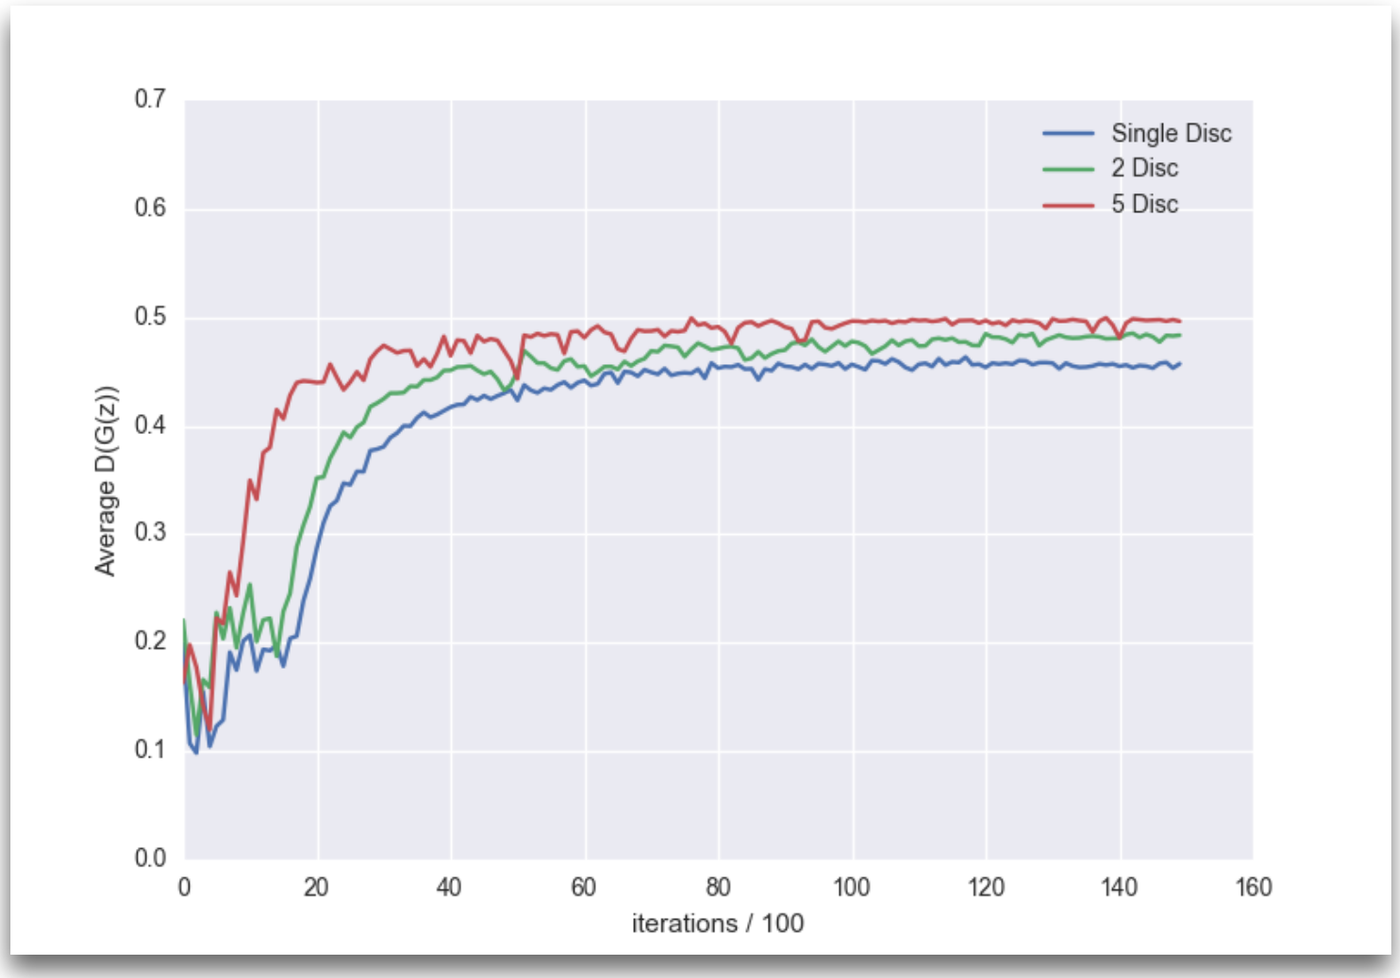

(Durugkar, Gemp, and Mahadevan, ICLR, 2017)

Plan of the Tutorial

- Part I: Motivation
- **Part II: Mathematical Background**
- Part III: Algorithms:
- Part IV: Applications and Challenges

Structure of VIs

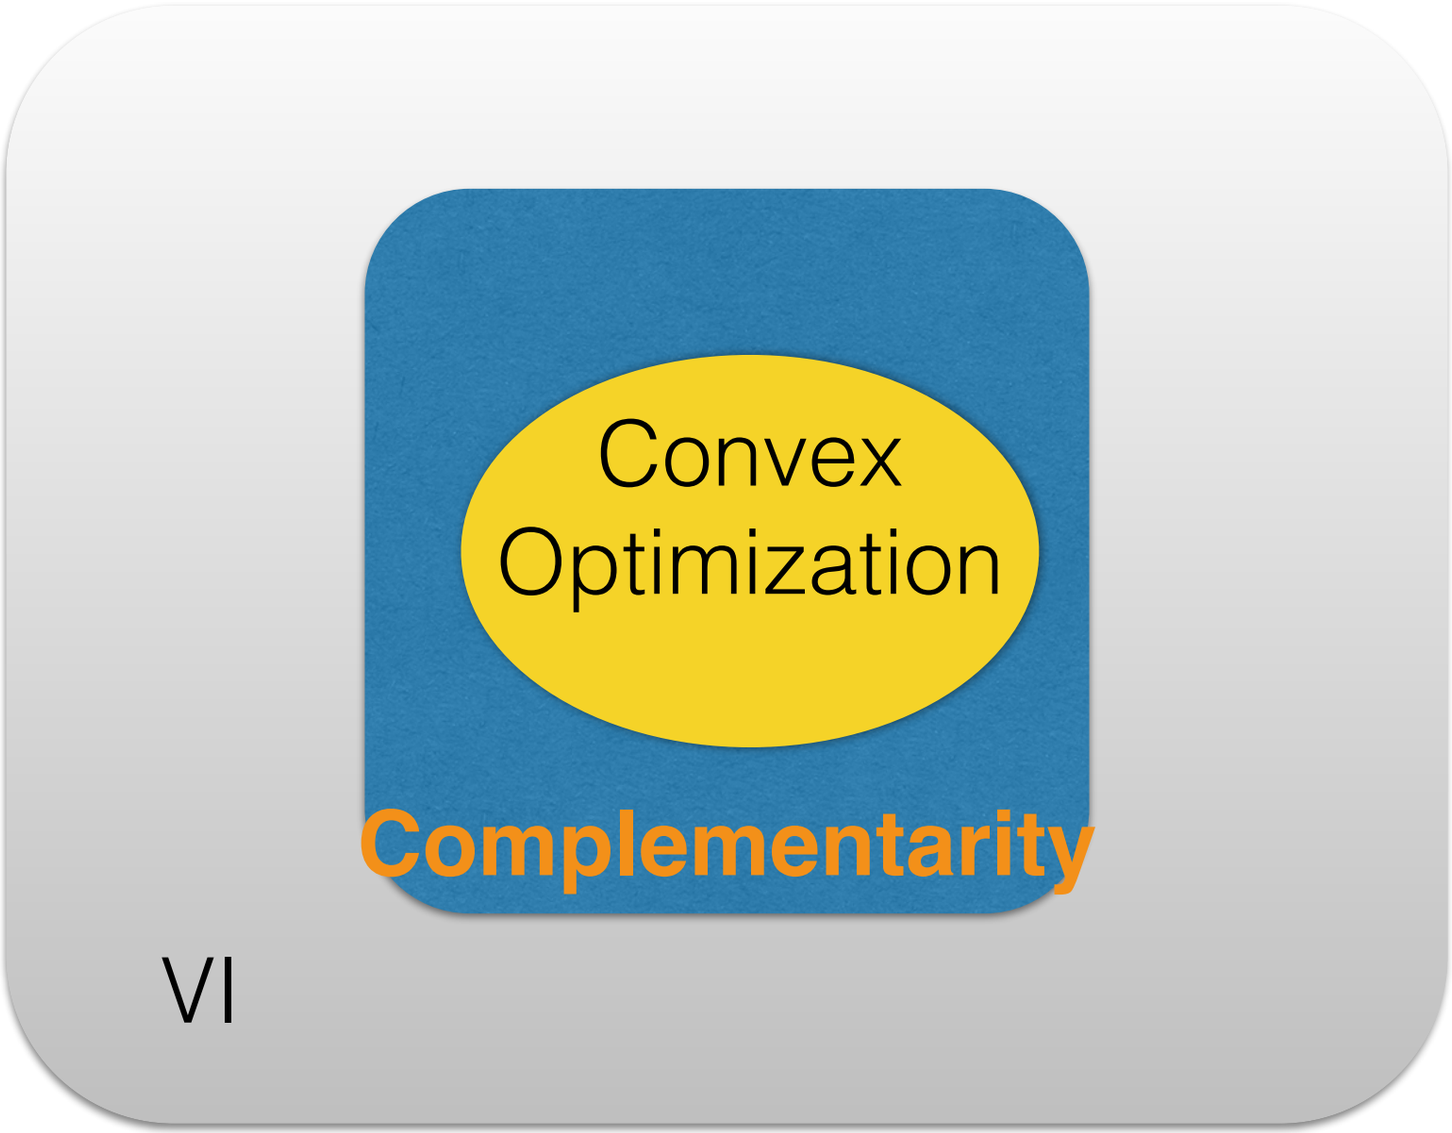

Optimization

Game Theory

Engineering

Linear Programming

$$\min_x c^T x$$

$$\text{s.t. } Ax = b, x \geq 0$$

AI

Business

Convex set

contains the line segment between any two points in the set

$$x_1, x_2 \in C, \quad 0 \leq \theta \leq 1 \quad \Longrightarrow \quad \theta x_1 + (1-\theta)x_2 \in C$$

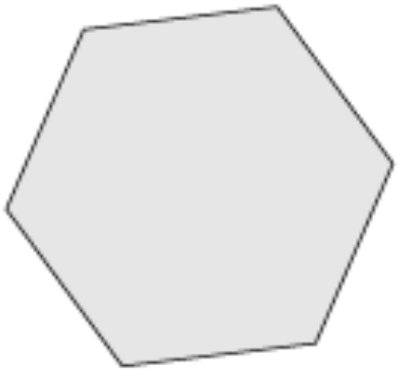

convex

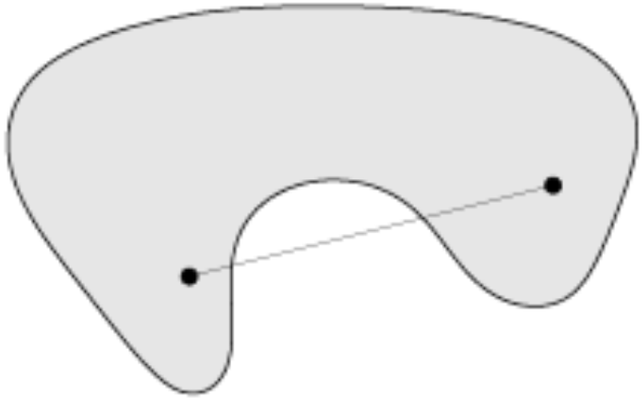

not convex

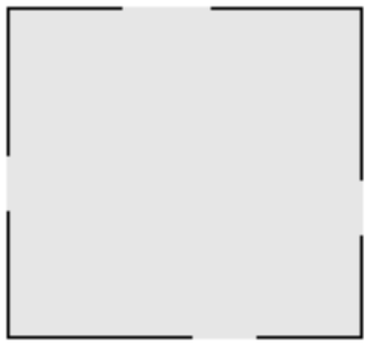

not convex

Projection onto Convex Sets

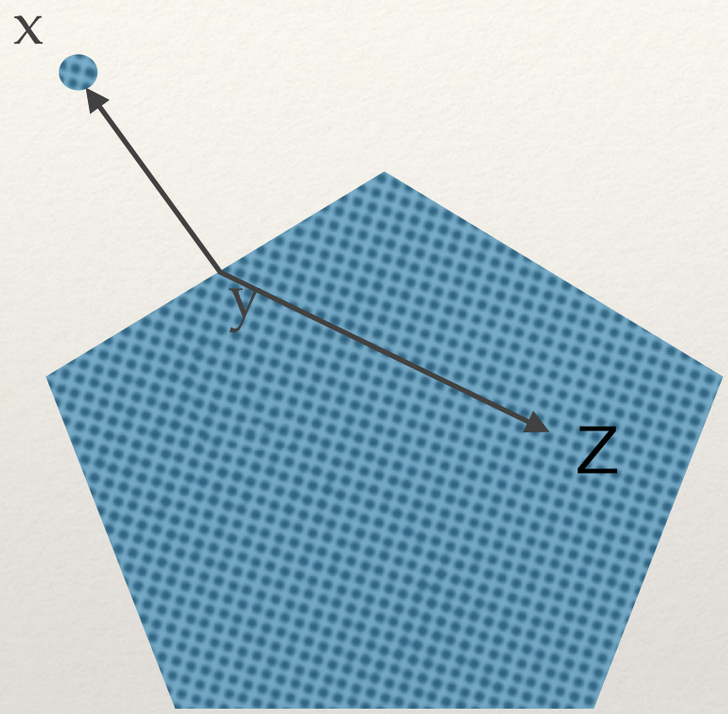

$$\text{if } y = \Pi_{\hat{S}}(x), \text{ then } \ (y-x)^T(z-y) \geq 0, z \in \hat{S}$$

Convex Feasibilty Revisited

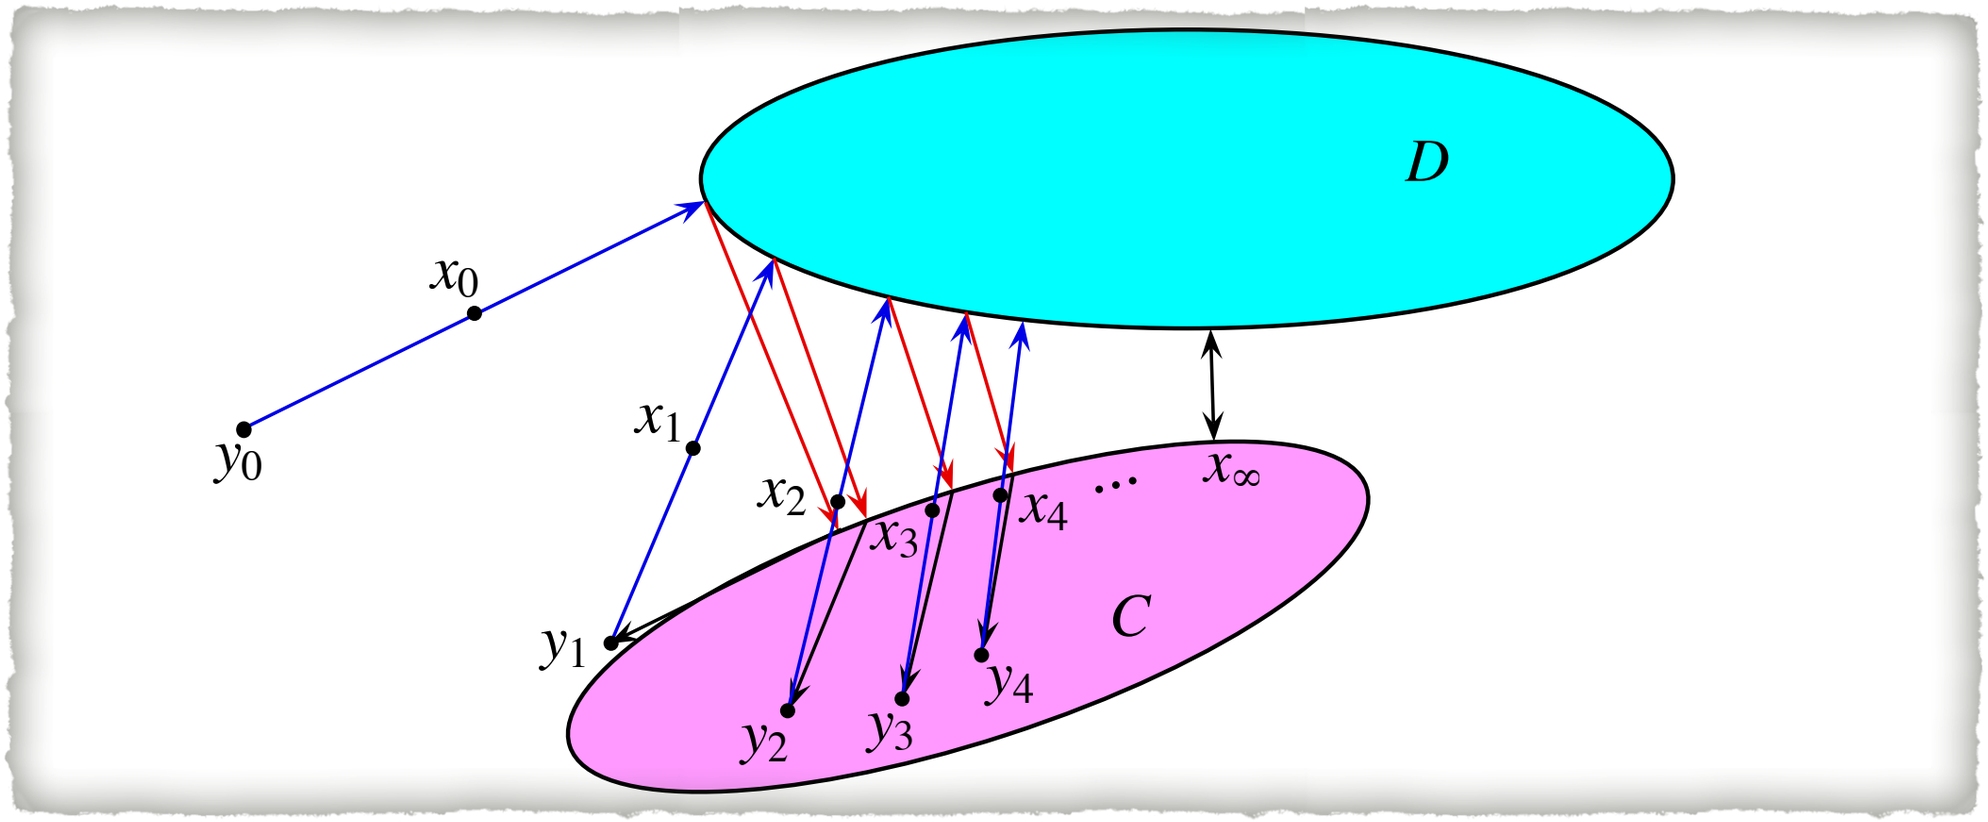

Proximal splitting methods in signal processing
Combetti and Pesquet

Convex function

convex function: domain is a convex set and Jensen's inequality holds:

$$f(\theta x + (1-\theta)y) \leq \theta f(x) + (1-\theta) f(y)$$

for all $x, y \in \mathbf{dom}\, f$, $0 \leq \theta \leq 1$

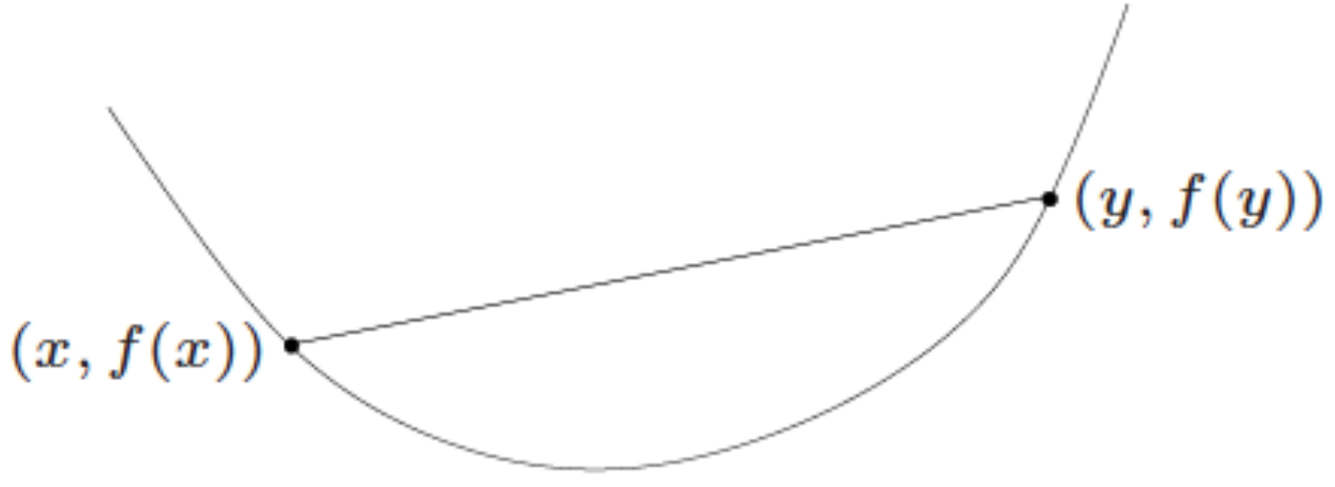

f is **strictly convex** if Jensen's inequality is strict for $0 < \theta < 1$, $x \neq y$

f is (strictly) **concave** if $-f$ is (strictly) convex

Indicator function

the indicator function of a set C is

$$I_C(x) = \begin{cases} 0 & x \in C \\ +\infty & \text{otherwise} \end{cases}$$

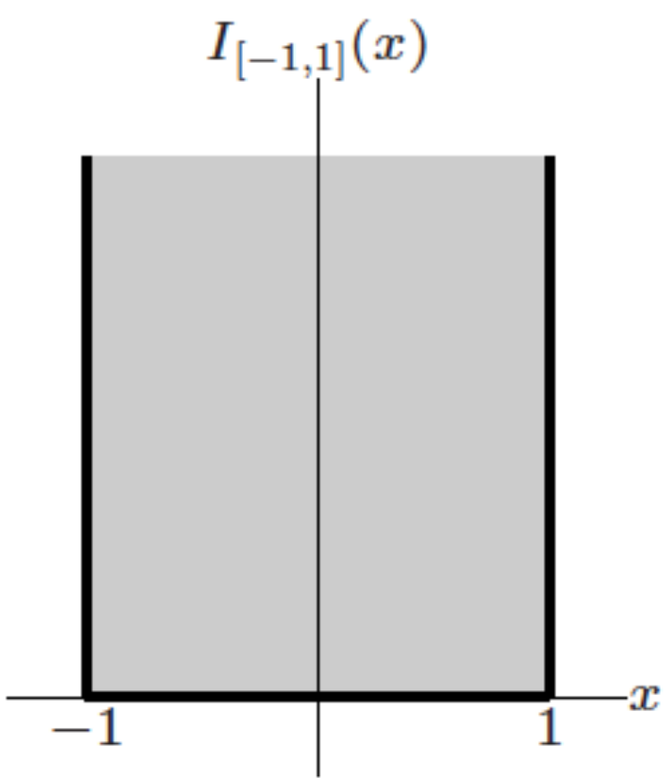

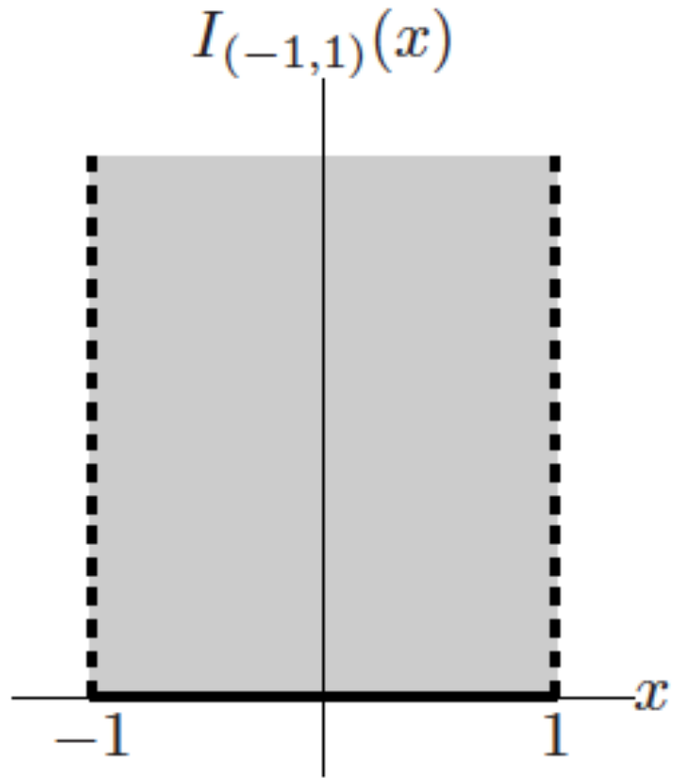

the indicator function of a convex set is a convex function

Bregman Divergence

$$D_\varphi(\mathbf{x}, \mathbf{y}) = \varphi(\mathbf{x}) - \varphi(\mathbf{y}) - (\mathbf{x} - \mathbf{y})^T \nabla\varphi(\mathbf{y})$$

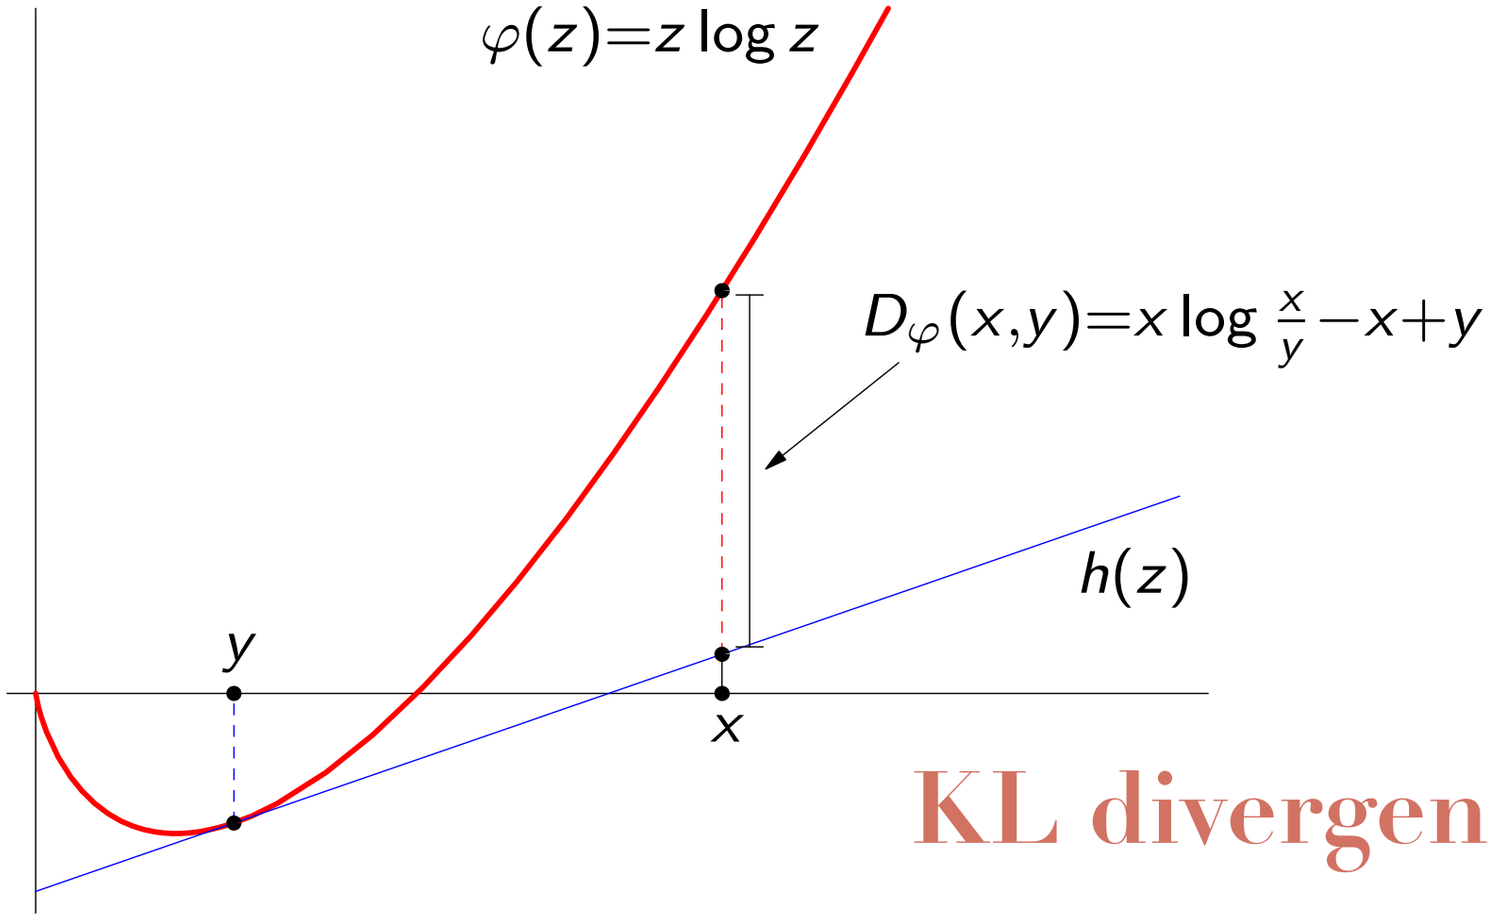

KL divergence

Proximal Mapping

Moreau, 1965:

$$\text{prox}_f(v) = \text{argmin}_x(f(x) + \frac{1}{2}\|x - v\|_2^2)$$

Proximal operator generalizes projection:

$$f(x) = I_{\mathcal{C}}(x) : \text{prox}_{I_{\mathcal{C}}}(v) = \Pi_{\mathcal{C}}(v) = \text{argmin}_{x \in \mathcal{C}}\|x - v\|_2$$

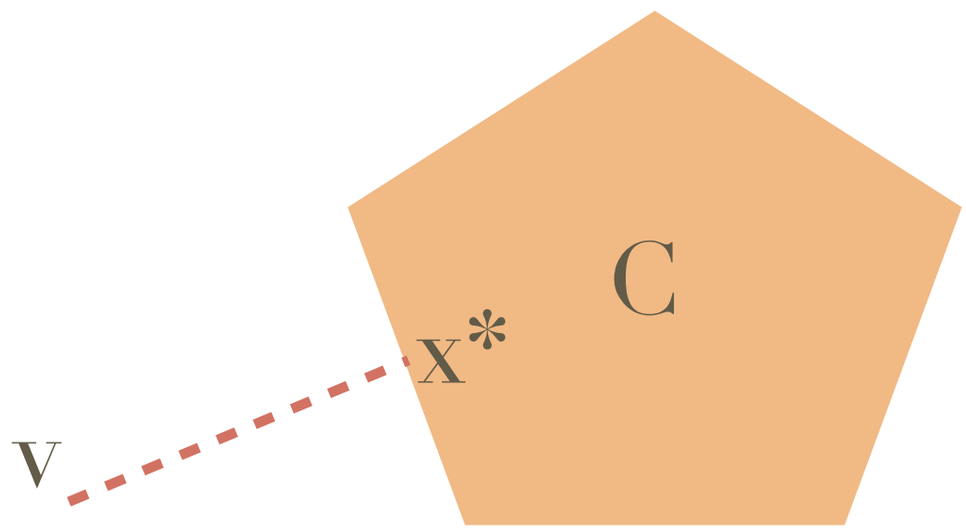

Normal Cone

Subdifferential of a convex function:

$$\partial f(x) = \{v \in \mathbb{R}^n : f(z) \geq f(x) + v^T(z - x), \forall z \in \mathrm{dom}(f)\}$$

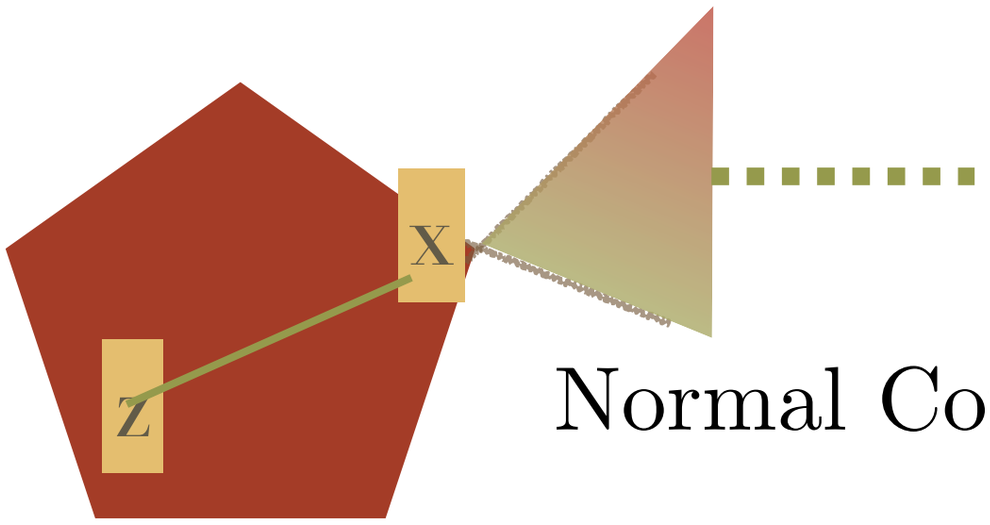

$I_{\mathcal{C}}(x)$

Normal Cone: $\partial I_{\mathcal{C}}(x) = N_{\mathcal{C}}(x)$

Resolvent

Given a set-valued mapping R

the resolvent of R is defined as $(I + \lambda R)^{-1}$

$$(x, y) \in R$$

$$(x, \lambda y) \in \lambda R$$

$$(x, x + \lambda y) \in I + \lambda R$$

$$(x + \lambda y, x) \in (I + \lambda R)^{-1}$$

Resolvent of the Subdifferential

$$\text{if } z \in (I + \lambda \partial f)^{-1}(x)$$

$$x \in (I + \lambda \partial f)(z)$$

$$0 \in \partial f(z) + \frac{1}{\lambda}(z - x)$$

$$0 \in \partial_z (f(z) + \frac{1}{2\lambda}\|z - x\|_2^2)$$

Proximal step is the resolvent of the subdifferential

Optimization

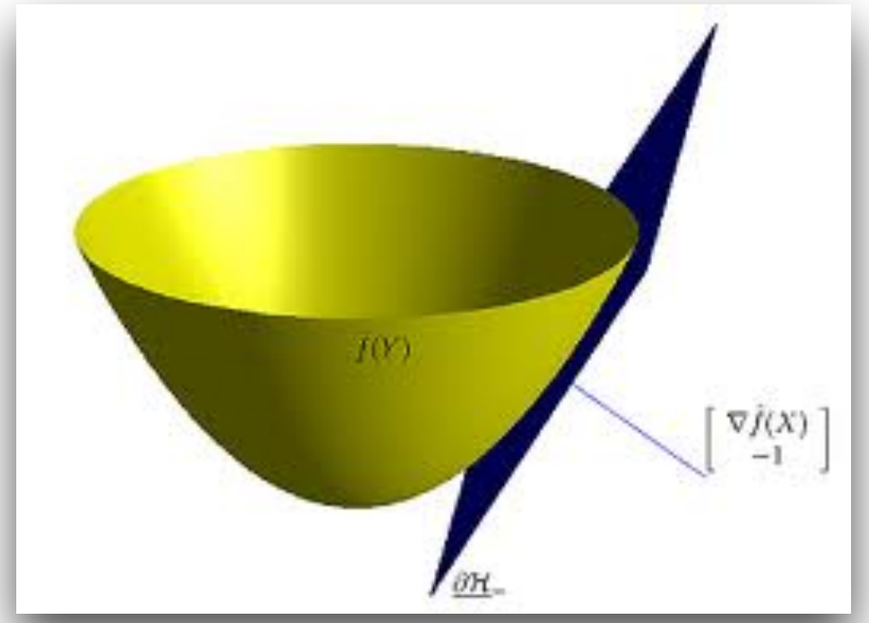

min f(x)
x in feasible set K

(Gauss, Newton, Shor, Hestenes 1800s-1950s)

Equilibration

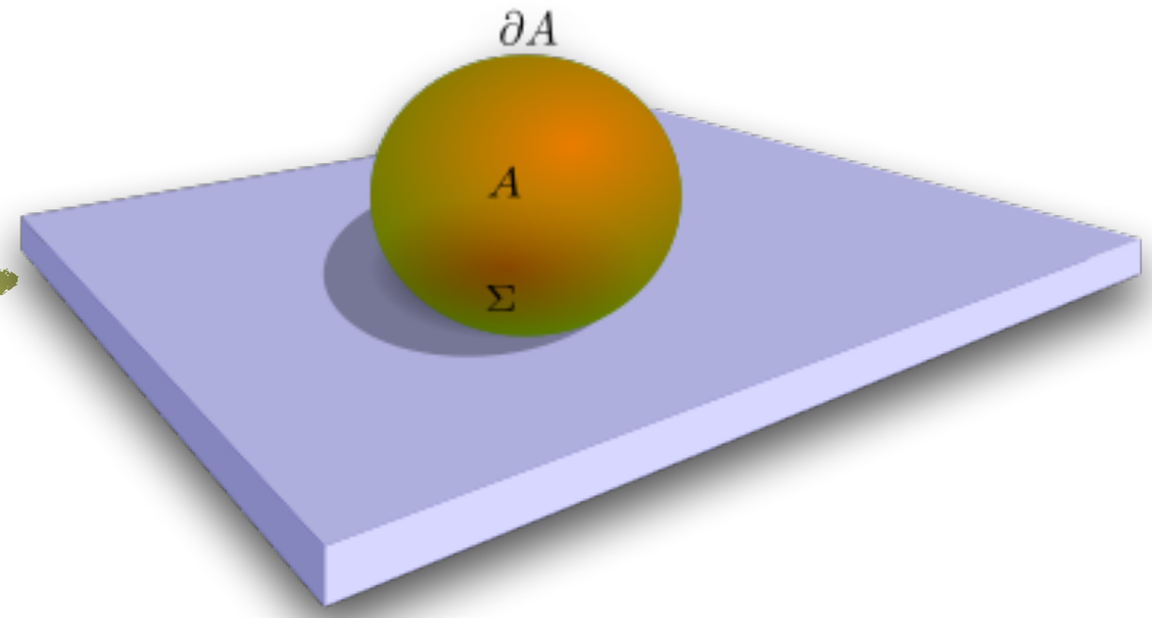

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

(Stampacchia, 1960s)

Variational Inequality

(Stampacchia, 1960s)

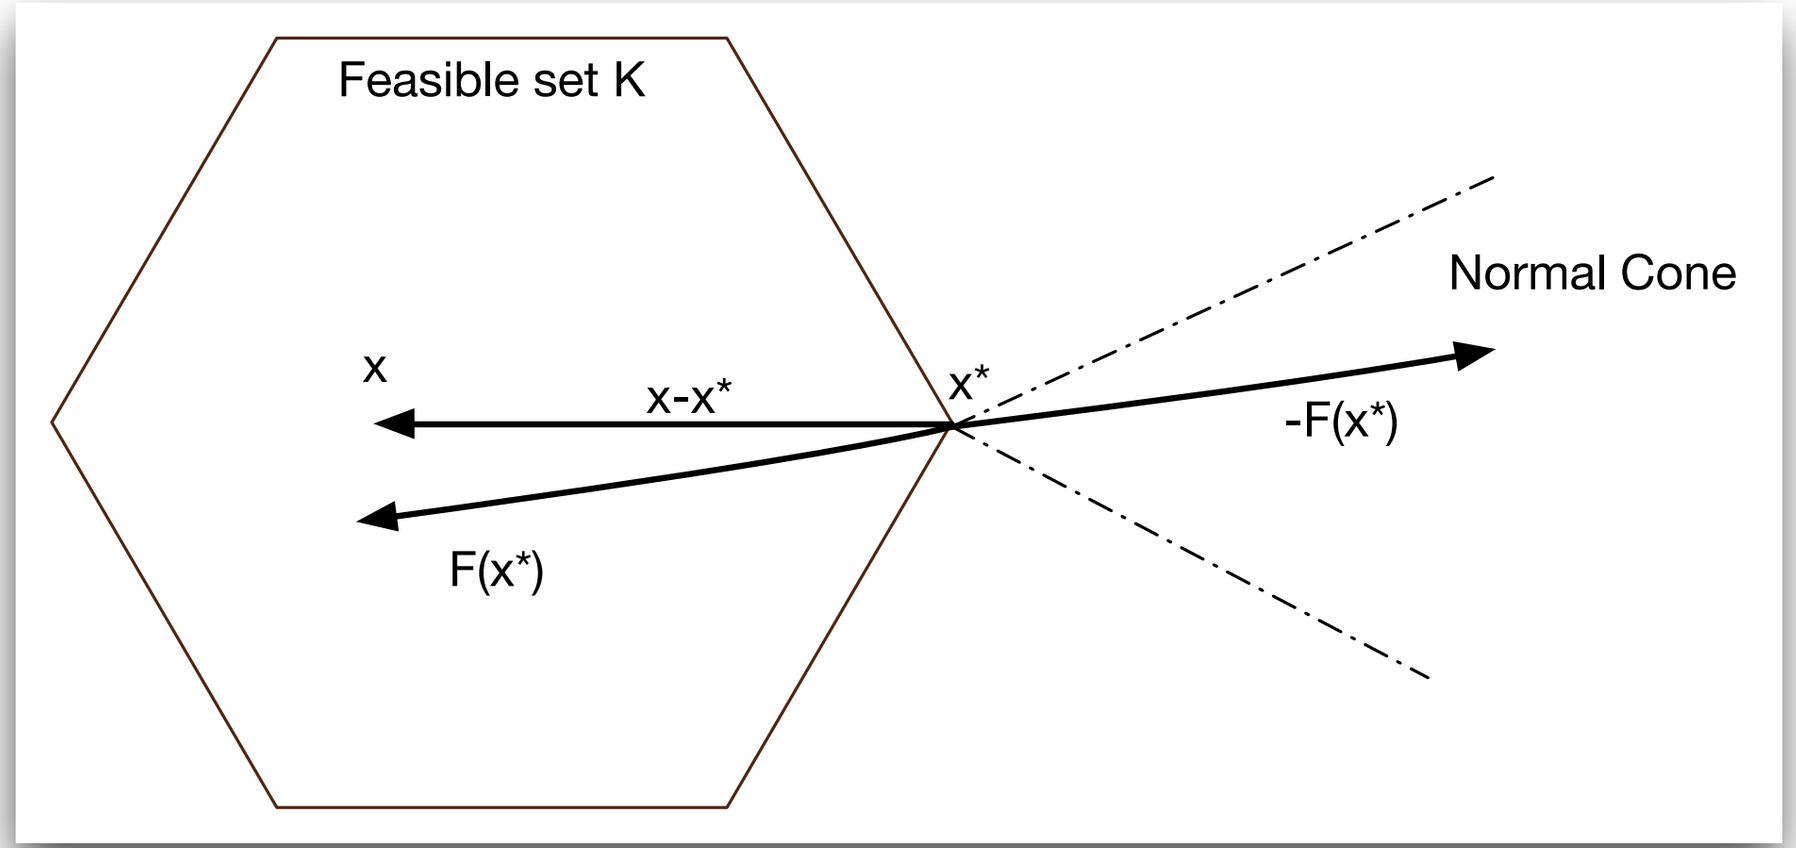

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

Vector field

Game theory => VI

- A CN game consists of m players, where player i chooses a strategy $x_i \in X_i$
- Let the joint payoffs for player i be $F_i(x_1,\ldots,x_m)$
- A set of strategies x* is in Nash equilibrium if

$$\langle (x_i - x_i^*), \nabla_i F_i(x_i^*) \rangle \geq 0$$

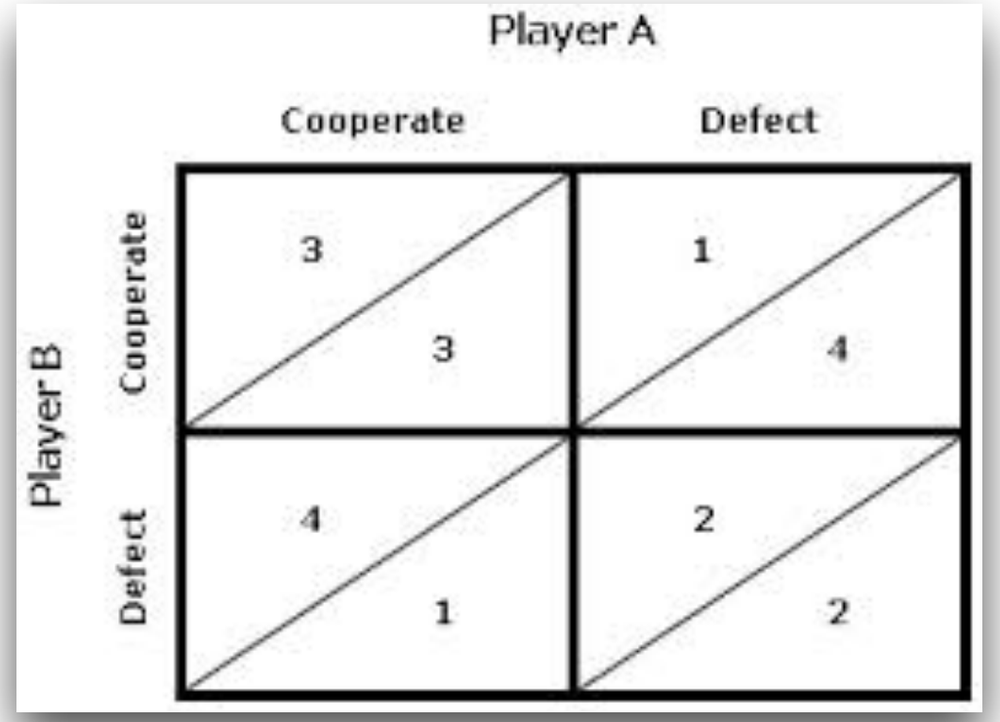

Generalized Nash Equilibrium Problems

- Arrow and Debreu wrote a famous paper in the 1950s, laying the foundations of mathematical economics.

- Central to their model is the GNEP

  - Each player's strategy sets are not independent, but may depend on what the other players do

  - A GNEP can be reduced to solving a VI

Equilibrium Problem

Let $K \subseteq R^n,\ f : K \times K \to R$ such that $f(x,x) = 0, \quad \forall x \in K$

Definition: An equilibrium problem is defined as that of finding $\bar{x} \in R^n$ such that

$$f(\bar{x}, y) \geq 0 \quad \forall y \in K$$

Eugen Blum and Werner Oettli, From optimization and variational inequalities to equilibrium problems, The Mathematics Student 63 (1994) 123–145.

VIs and Equilibrium Problems

Clearly, a VI is an equilibrium problem where

$$f(x, y) = \langle F(x), y - x \rangle$$

We will primarily focus on VIs in this tutorial

Complementarity: Optimization without objectives

Find a vector x^* in R^N such that

x^* is non-negative, $F(x^*)$ is non-negative,
x^* is perpendicular to $F(x^*)$

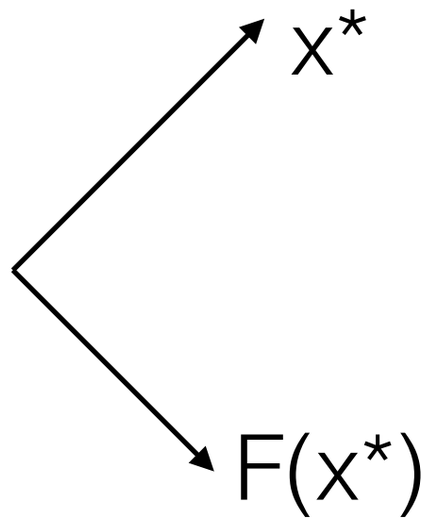

Gap Functions

- A function $\psi: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ which satisfies
  - $\psi(X) \geq 0$ for all $X \in C$
  - $\psi(X^*) = 0$, $X^* \in C$ iff X^* solves the VI
- A desired, but not required property is that the gap function diminishes on a sequence of points which converge to X^*

Example Gap Functions

- Auslender Gap Function

$$\theta_D(x) = \sup_{y \in C} \langle F(y), x - y \rangle.$$

- Fukushima & Auchmuty: Unbounded Domains (i.e. X ∈ $\mathbb{R}_+$)

$$g_\alpha(x) = \sup_{y \in C} \{ \langle F(x), x - y \rangle - \frac{\alpha}{2} ||y - x||^2 \}$$

- Note: sign of F(X) changes depending on ‘ascent’ or ‘descent’

Optimization vs VIs

Property	Optimization	VI
Mapping	(Strong) Convexity	(Strong) Monotonicity
Jacobian	Positive definite and symmetric	Asymmetric
Objective function	Single fixed	Multiple or none

Spatial Price Equilibria

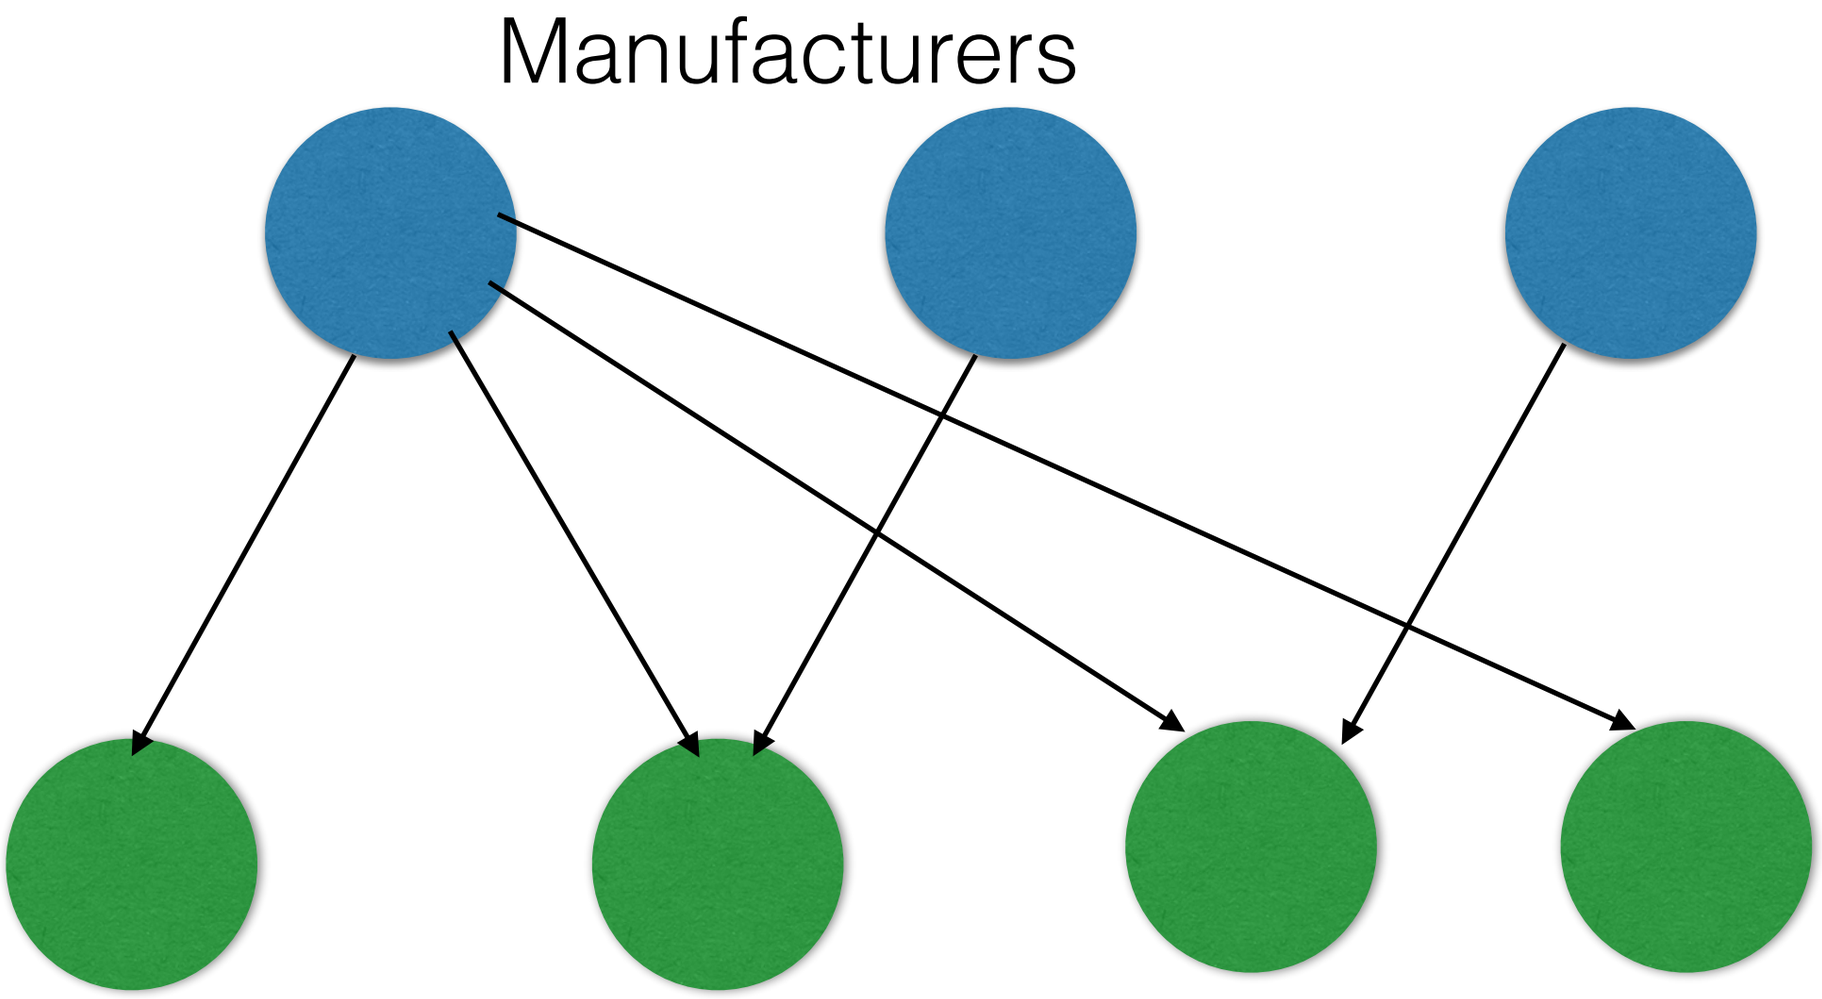

Spatial Price Equilibrium Model

$$\pi_i + c_{ij} = \begin{cases} \rho_j & \text{if } Q^*_{ij} > 0 \\ \geq \rho_j & \text{if } Q^*_{ij} = 0 \end{cases}$$

Supply level	s_i
Supply price	π_j
Quantity demanded	d_i
Demand price	ρ_j
Cost of transportation	Q_{ij}

Spatial Price Equilibrium Model as a VI

$$\begin{aligned} \langle \pi(s^*), s - s^* \rangle \; + \\ \langle c(Q^*), Q - Q^* \rangle \; - \\ \langle \rho(d^*), d - d* \rangle \geq 0 \quad \forall (s, Q, d) \in K \end{aligned}$$

if and only if

$$\pi_i + c_{ij} = \begin{cases} \rho_j & \text{if } Q^*_{ij} > 0 \\ \geq \rho_j & \text{if } Q^*_{ij} = 0 \end{cases}$$

Can SPE be reduced to an Optimization Problem?

$$\frac{\partial \pi_i}{\partial s_k} = \frac{\partial \pi_k}{\partial s_i}$$

Only when

$$\frac{\partial \rho_j}{\partial d_l} = \frac{\partial \rho_l}{\partial d_j}$$

This is obviously not true in general

Optimization => VI

Suppose $x^* = \text{argmin}_{x \in \mathcal{K}} f(x)$

where f is differentiable

Then x^* solves the VI

$\langle \nabla f(x^*), x - x^* \rangle \geq 0. \; \forall x \in \mathcal{K}$

Proof: Define $\phi(t) = f(x^* + t(x - x^*))$

Since $\phi(0)$ achieves the minimum

$\phi'(0) = \langle \nabla f(x^*), x - x^* \rangle \geq 0$

When VI => optimization?

$$\text{Given } VI(F,K), \text{ define } \nabla F(x) = \begin{bmatrix} \frac{\partial F_1}{\partial x_1} & \cdots & \frac{\partial F_1}{\partial x_n} \\ \vdots & \cdots & \vdots \\ \frac{\partial F_n}{\partial x_1} & \cdots & \frac{\partial F_n}{\partial x_n} \end{bmatrix}$$

When ∇F is symmetric and positive semi-definite VI(F,K) can be reduced to an optimization problem,

Monotone Operators

- relation F on $\mathbf{R}^n$ is *monotone* if

$$(u-v)^T(x-y) \geq 0 \quad \text{for all } (x,u),\ (y,v) \in F$$

- F is *maximal monotone* if there is no monotone operator that properly contains it

for f convex, $F(x) = \partial f(x)$ is monotone

- suppose $u \in \partial f(x)$ and $v \in \partial f(y)$
- then

$$f(y) \geq f(x) + u^T(y-x), \qquad f(x) \geq f(y) + v^T(x-y)$$

- add these and cancel $f(y) + f(x)$ to get

$$0 \leq (u-v)^T(x-y)$$

Monotone Inclusion Problem

Given a monotone operator F, find x $s.t.$ $0 \in F(x)$

This means that $(x, 0) \in F$

For convex f, if x^* minimizes f, $0 \in \partial f(x^*)$

VI as monotone inclusion

$$0 \in F(x^*) + N_K(x^*)$$

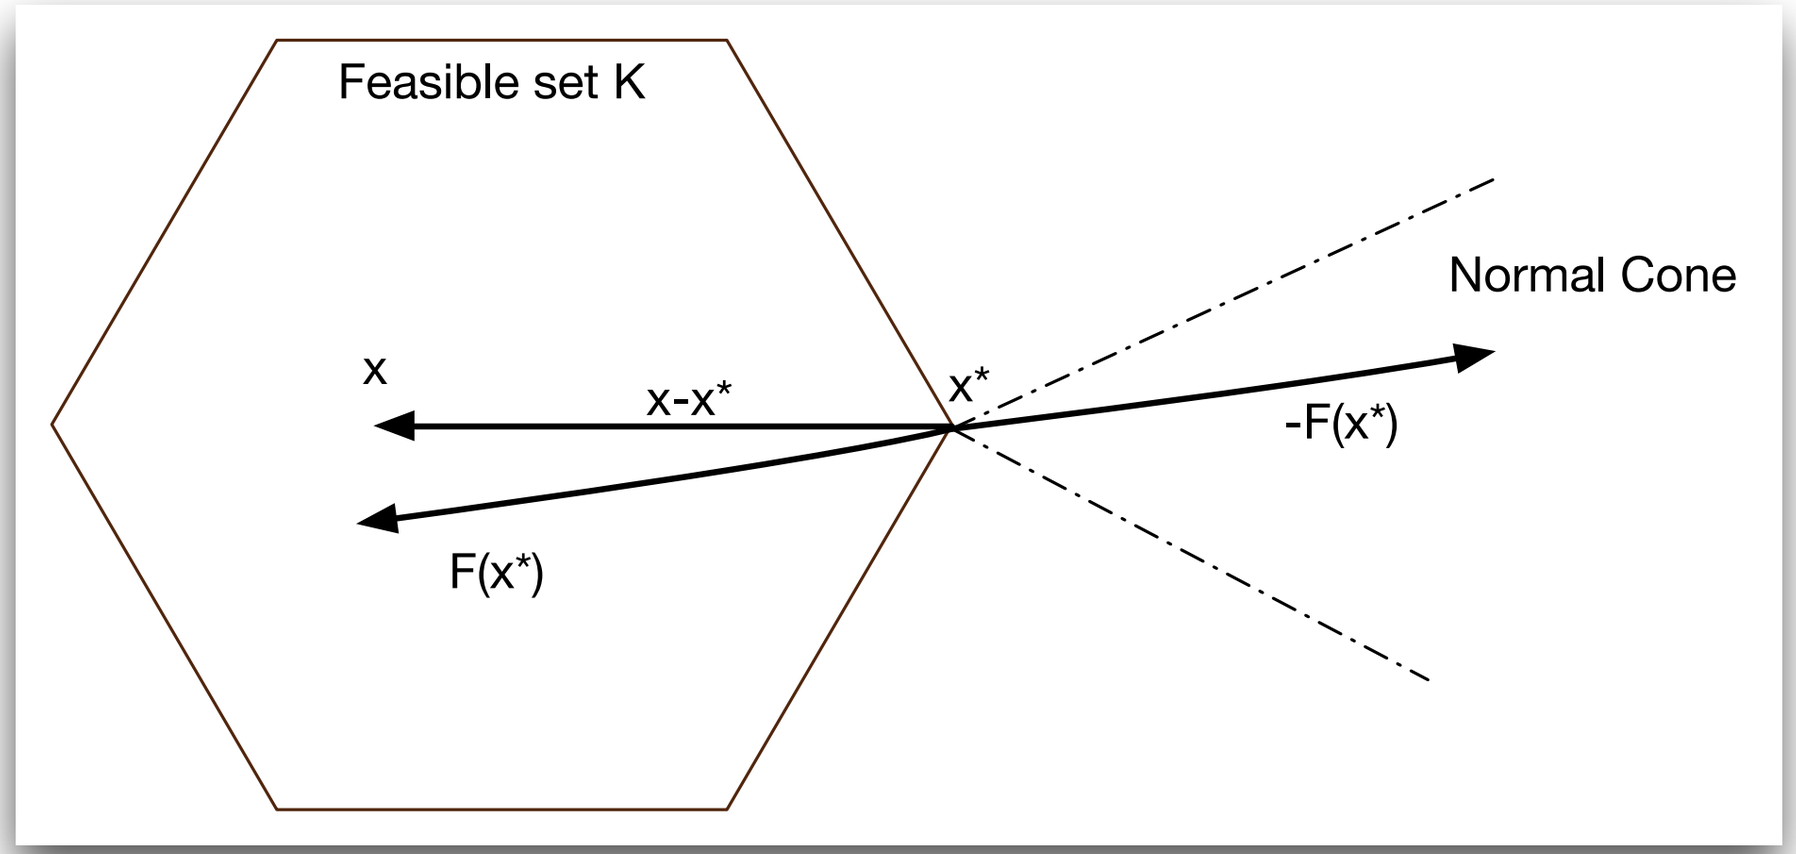

$0 \in T(x)$

Monotone Inclusion

$0 \in A(x) + B(x)$

Operator Splitting

$\text{prox}_f(v) = \text{argmin}_x(f(x) + \frac{1}{2}\|x - v\|_2^2)$

Proximal Mappings

$x_{k+1} \leftarrow \nabla\psi^*(\nabla\psi(x_k) - \alpha_k \partial f(x))$

Mirror Descent

ADMM (Dual Decomposition)

Plan of the Tutorial

- Part I: Motivation
- Part II: Mathematical Background
- **Part III: Algorithms**
- Part IV: Applications and Challenges

Fixed Point Formulation of VI

Theorem 5. *The vector x^* is the solution of VI(F,K) if and only if, for any $\gamma > 0$, x^* is also a fixed point of the map $x^* = \Pi_K(x^* - \gamma F(x^*))$, where Π_K is the projector onto convex set K.*

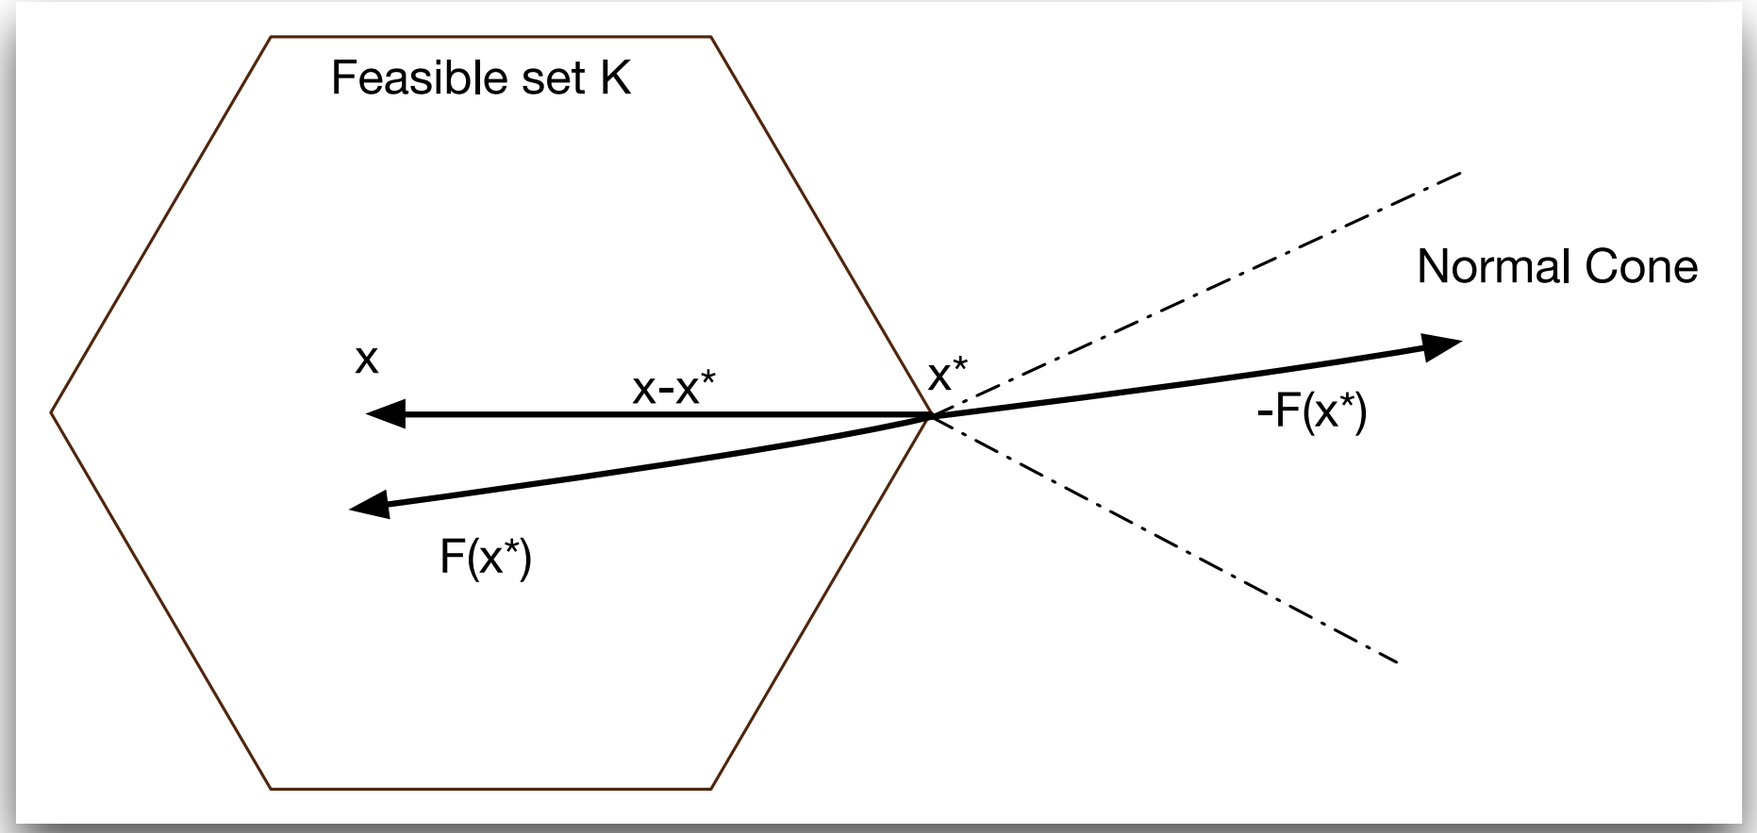

Projections and VIs

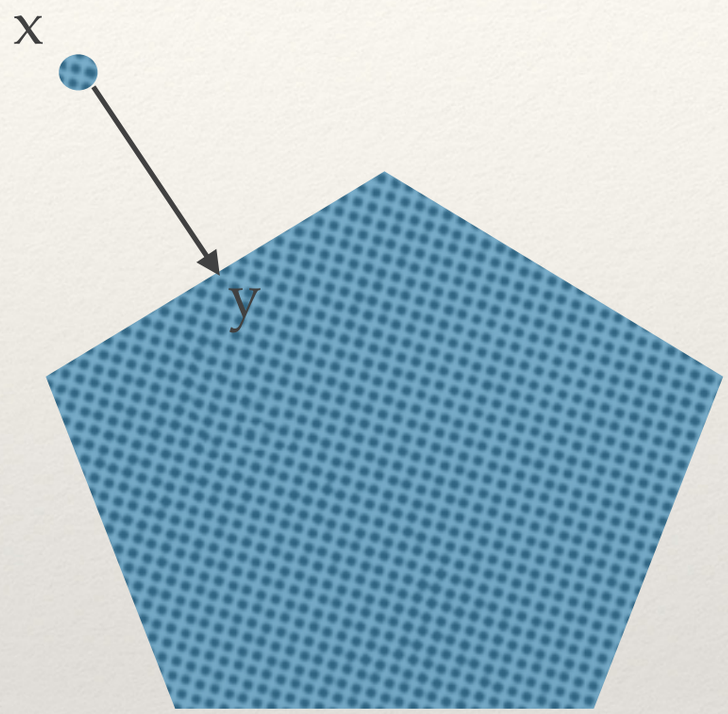

$$\text{if } y = \Pi_{\hat{S}}(x), \text{ then } \ (y-x)^T(z-y) \geq 0, z \in \hat{S}$$

Projection Algorithm

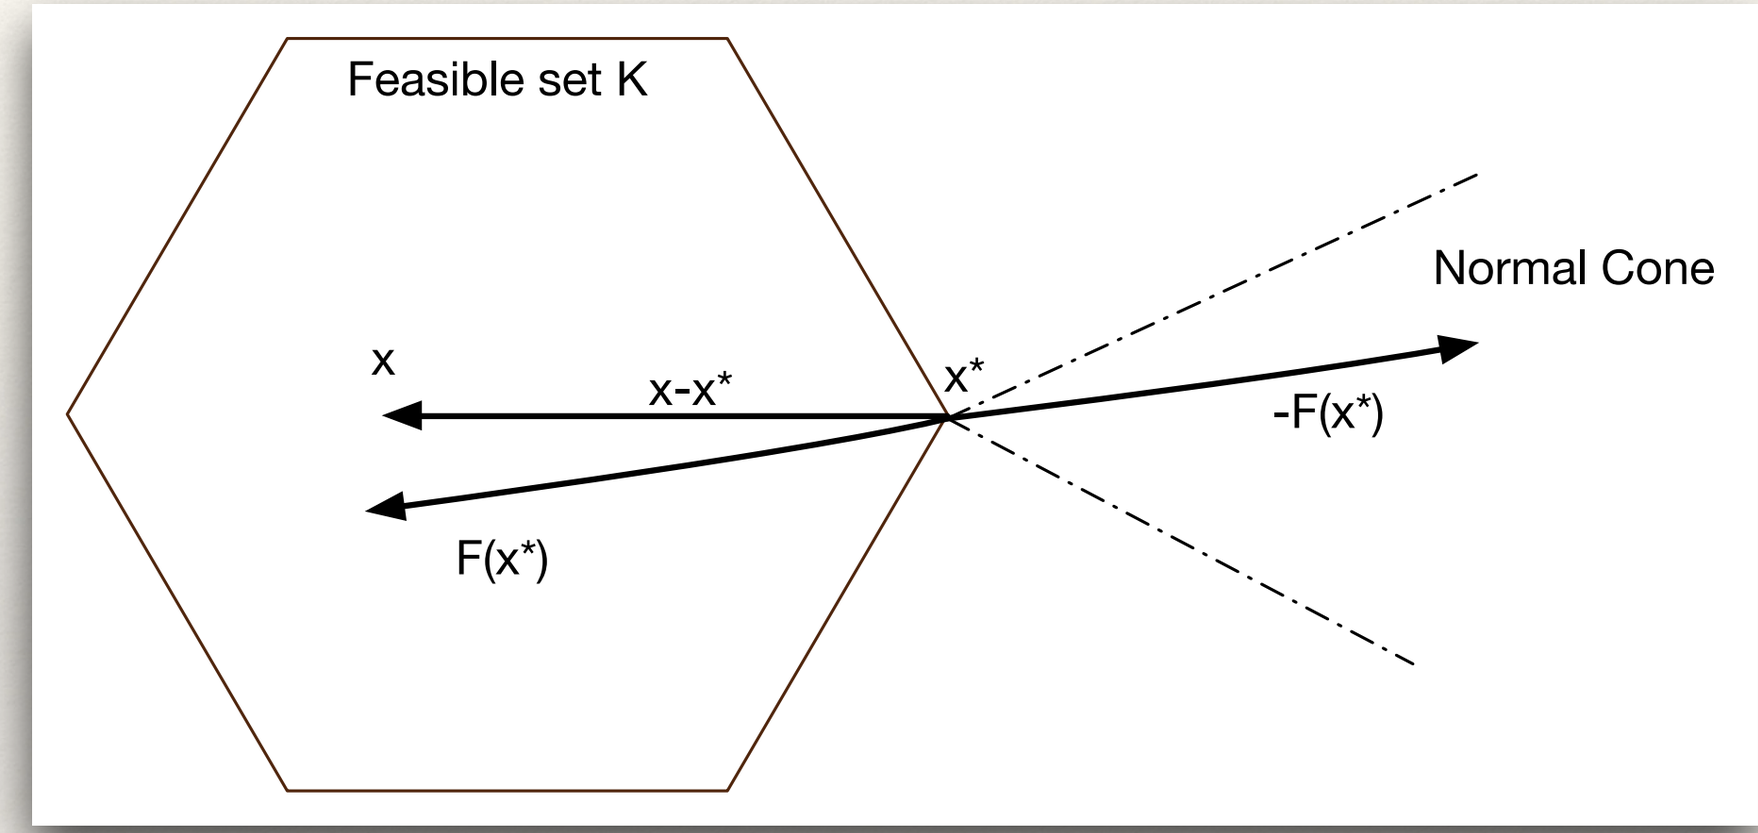

Projection Algorithm

Algorithm 1 The Basic Projection Algorithm for solving VIs.

INPUT: Given VI(F,K), and a symmetric positive definite matrix D.

1: Set $k = 0$ and $x_k \in K$.
2: **repeat**
3: Set $x_{k+1} \leftarrow \Pi_{K,D}(x_k - D^{-1}F(x_k))$.
4: Set $k \leftarrow k + 1$.
5: **until** $x_k = \Pi_{K,D}(x_k - D^{-1}F(x_k))$.
6: Return x_k

Gradients to Vector Fields

$$w_{t+1} \leftarrow w_t - \alpha_t \nabla f(w_t)$$

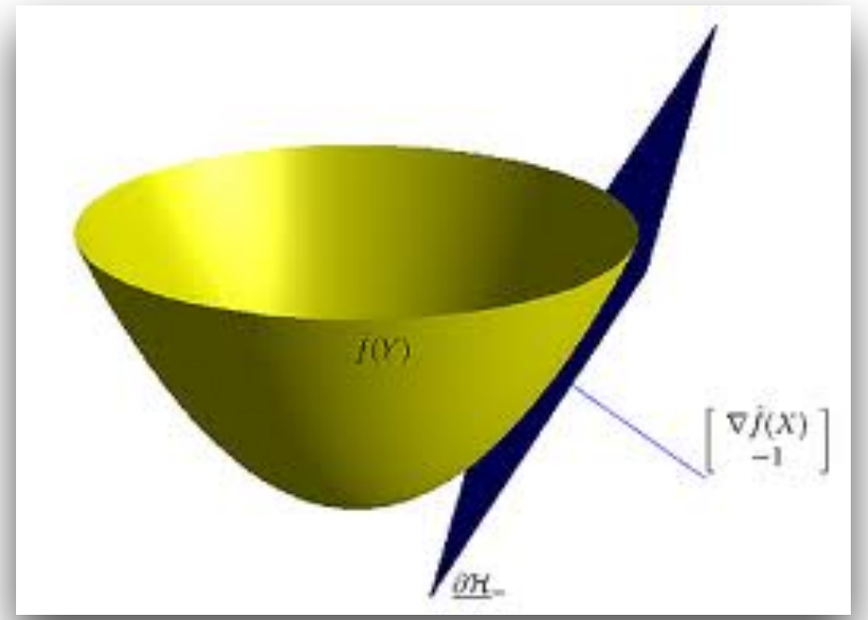

$$w_{t+1} \leftarrow w_t - \alpha_t F(w_t)$$

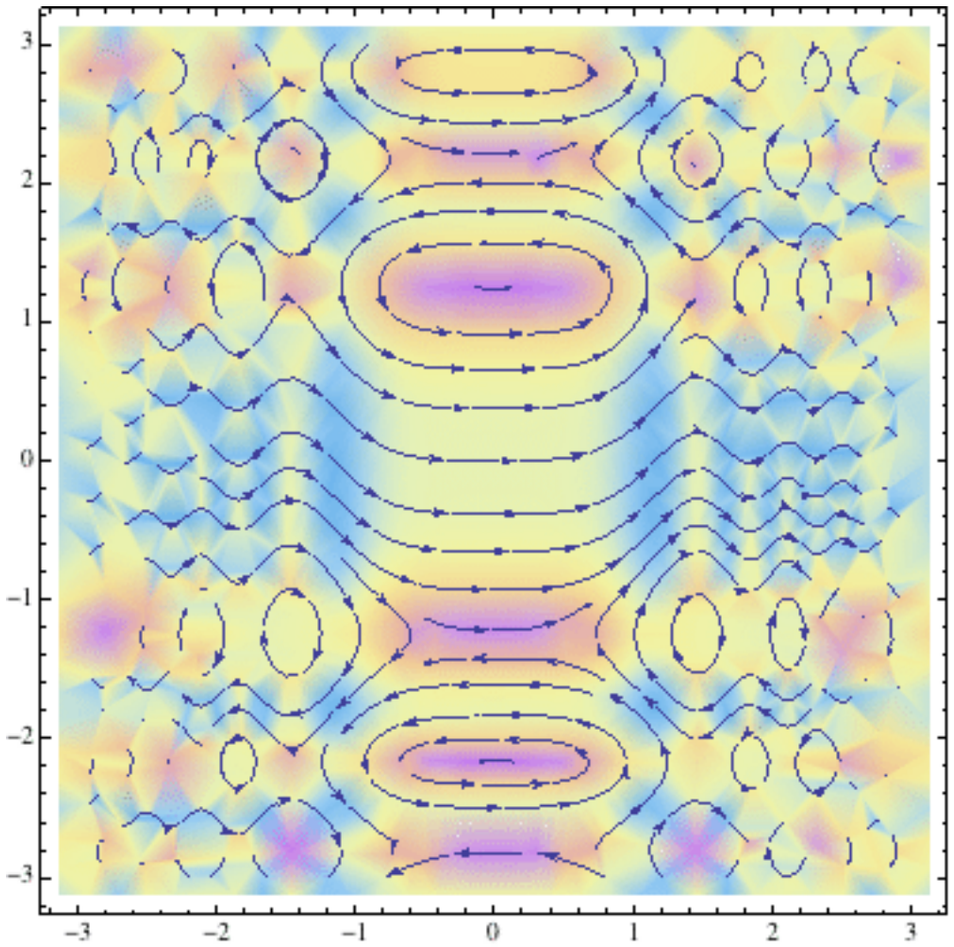

Projection Method Fails

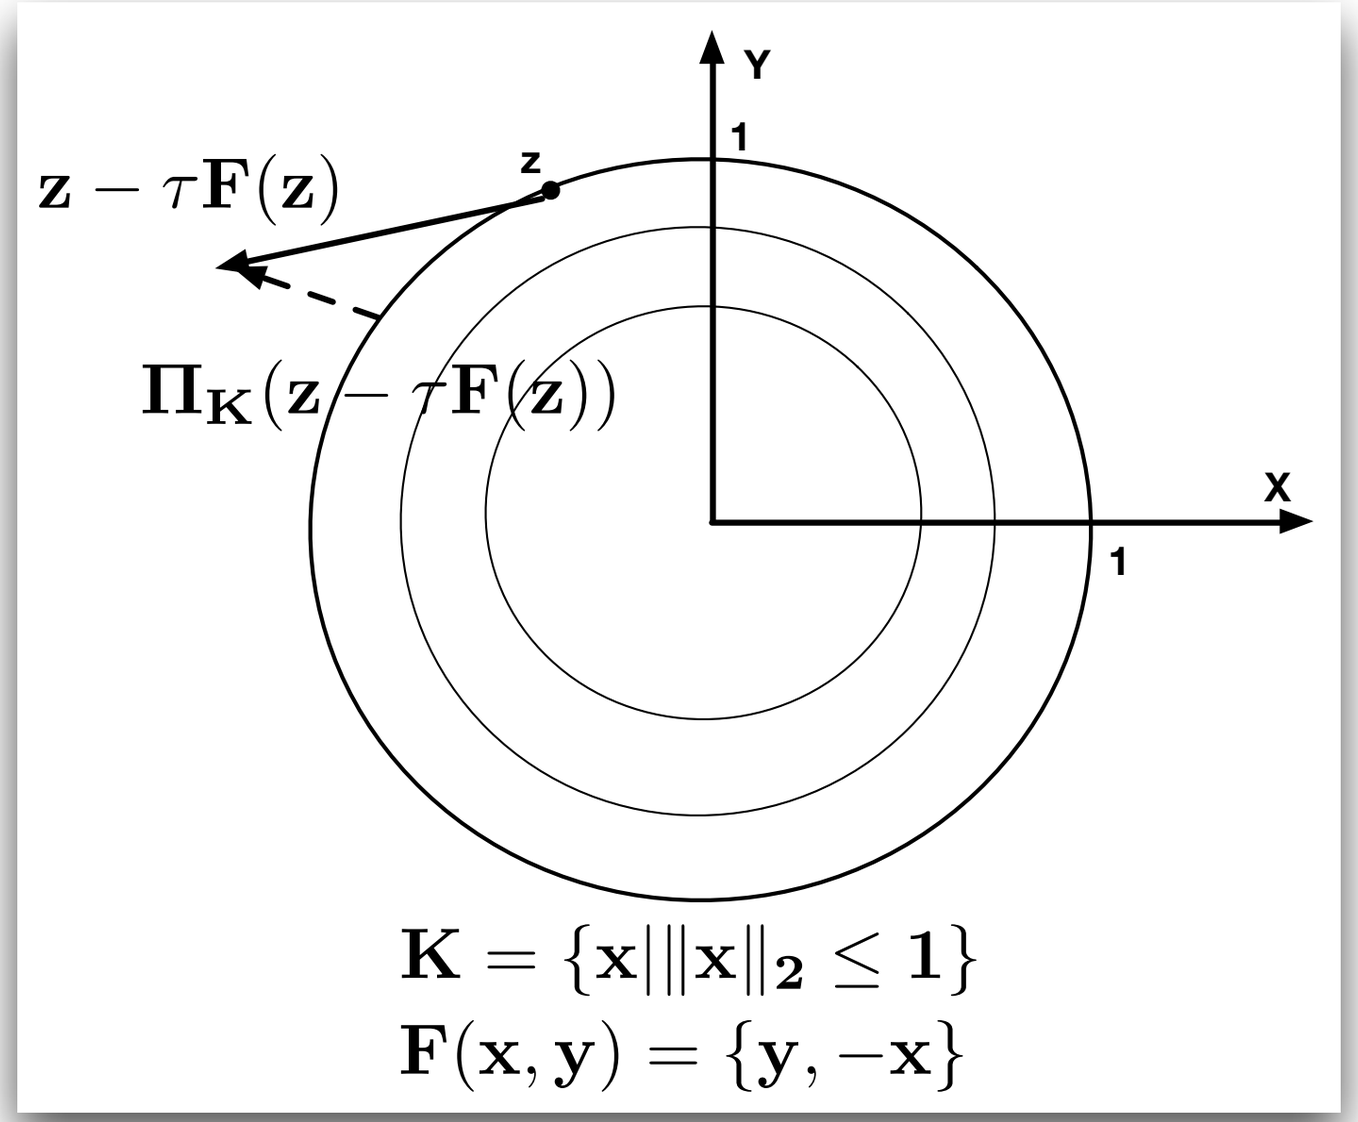

Bertsekas and Tsitsiklis, Parallel and Distributed Computation, Athena Scientific.

Generalizing Gradient Descent

$$w_{t+1} \leftarrow w_t - \alpha_t \nabla f(w_t)$$

$$w_{t+1} \leftarrow w_t - \alpha_t F(w_t)$$

Extragradient method for VIs

$$w_{t+1} \leftarrow \Pi_K(w_t - \alpha_t \Pi_K(w_t - \alpha_t F(w_t)))$$

Extragradient Method for solving VIs

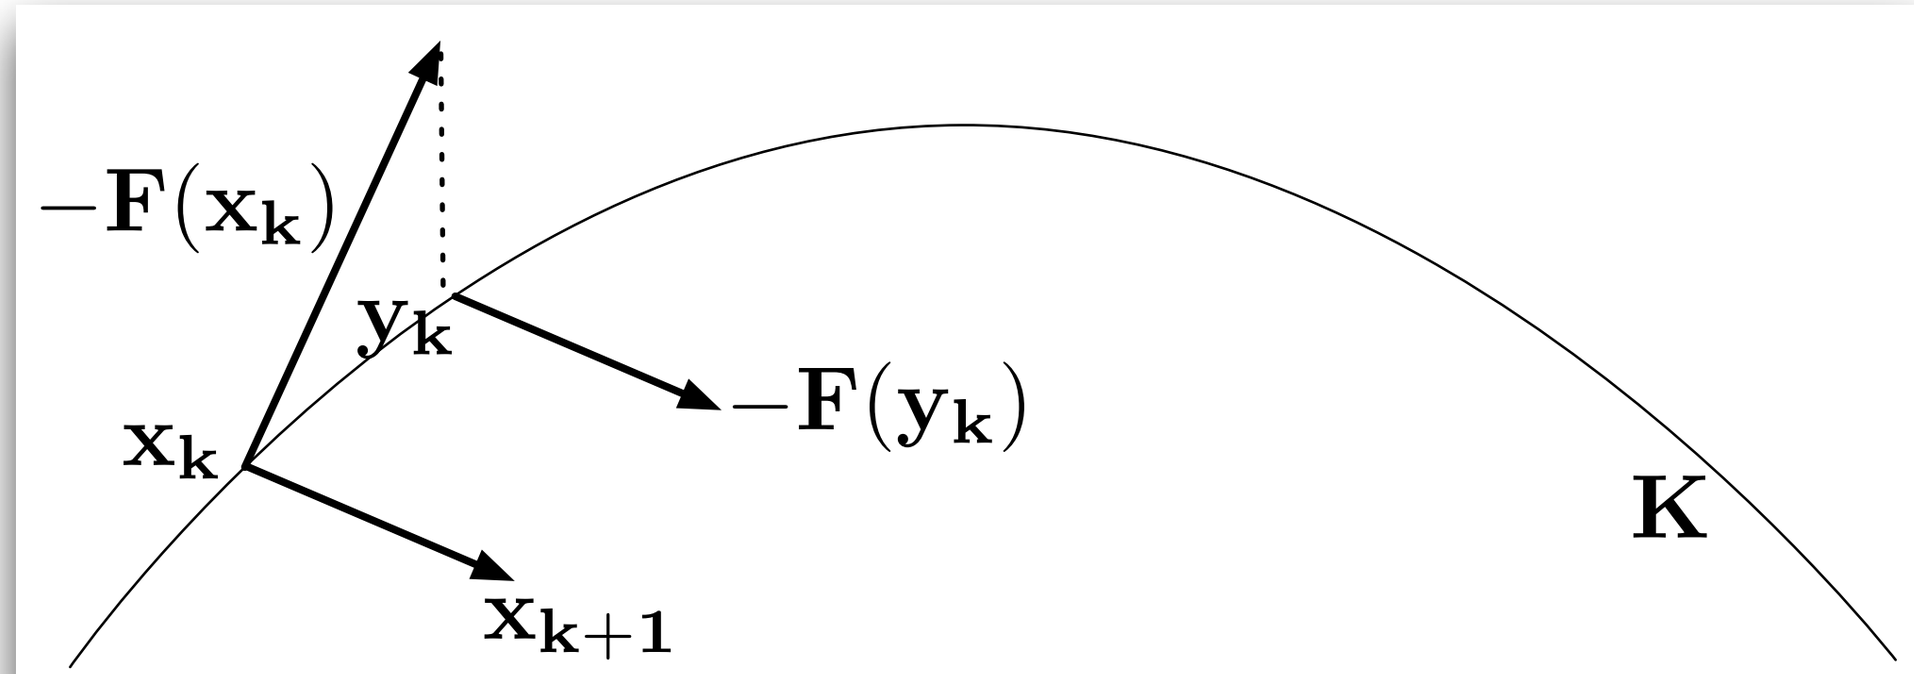

Korpolevich (1970s) developed the extragradient method

Extragradient Method

Algorithm 2 The Extragradient Algorithm for solving VIs.

INPUT: Given VI(F,K), and a scalar α.

1: Set $k = 0$ and $x_k \in K$.
2: **repeat**
3: Set $y_k \leftarrow \Pi_K(x_k - \alpha F(x_k))$.
4: Set $x_{k+1} \leftarrow \Pi_K(x_k - \alpha F(y_k))$.
5: Set $k \leftarrow k + 1$.
6: **until** $x_k = \Pi_K(x_k - \alpha F(x_k))$.
7: Return x_k

Khobotov developed a learning rate rule under which the extragradient method works for all pseudo-monotone mappings

$$\langle F(y), x - y\rangle \geq 0 \Rightarrow \langle F(x), x - y\rangle \geq 0, \ \forall x, y \in K$$

Mirror Descent = Proximal Algm + Bregman Divergence

- Mirror descent can be viewed as a proximal method using a Bregman divergence

$$x_{k+1} = \operatorname{argmin}_{x \in X} \left(\langle x, \partial f(x_k) \rangle + \frac{1}{t_k} D_\phi(x, x_k) \right)$$

- Mirror descent can outperform regular subgradient method by O(n/log(n)) (Beck and Teboulle, 2003)

Mirror Maps

(Nemirovski and Yudin, 1980s; Bubeck, 2014)

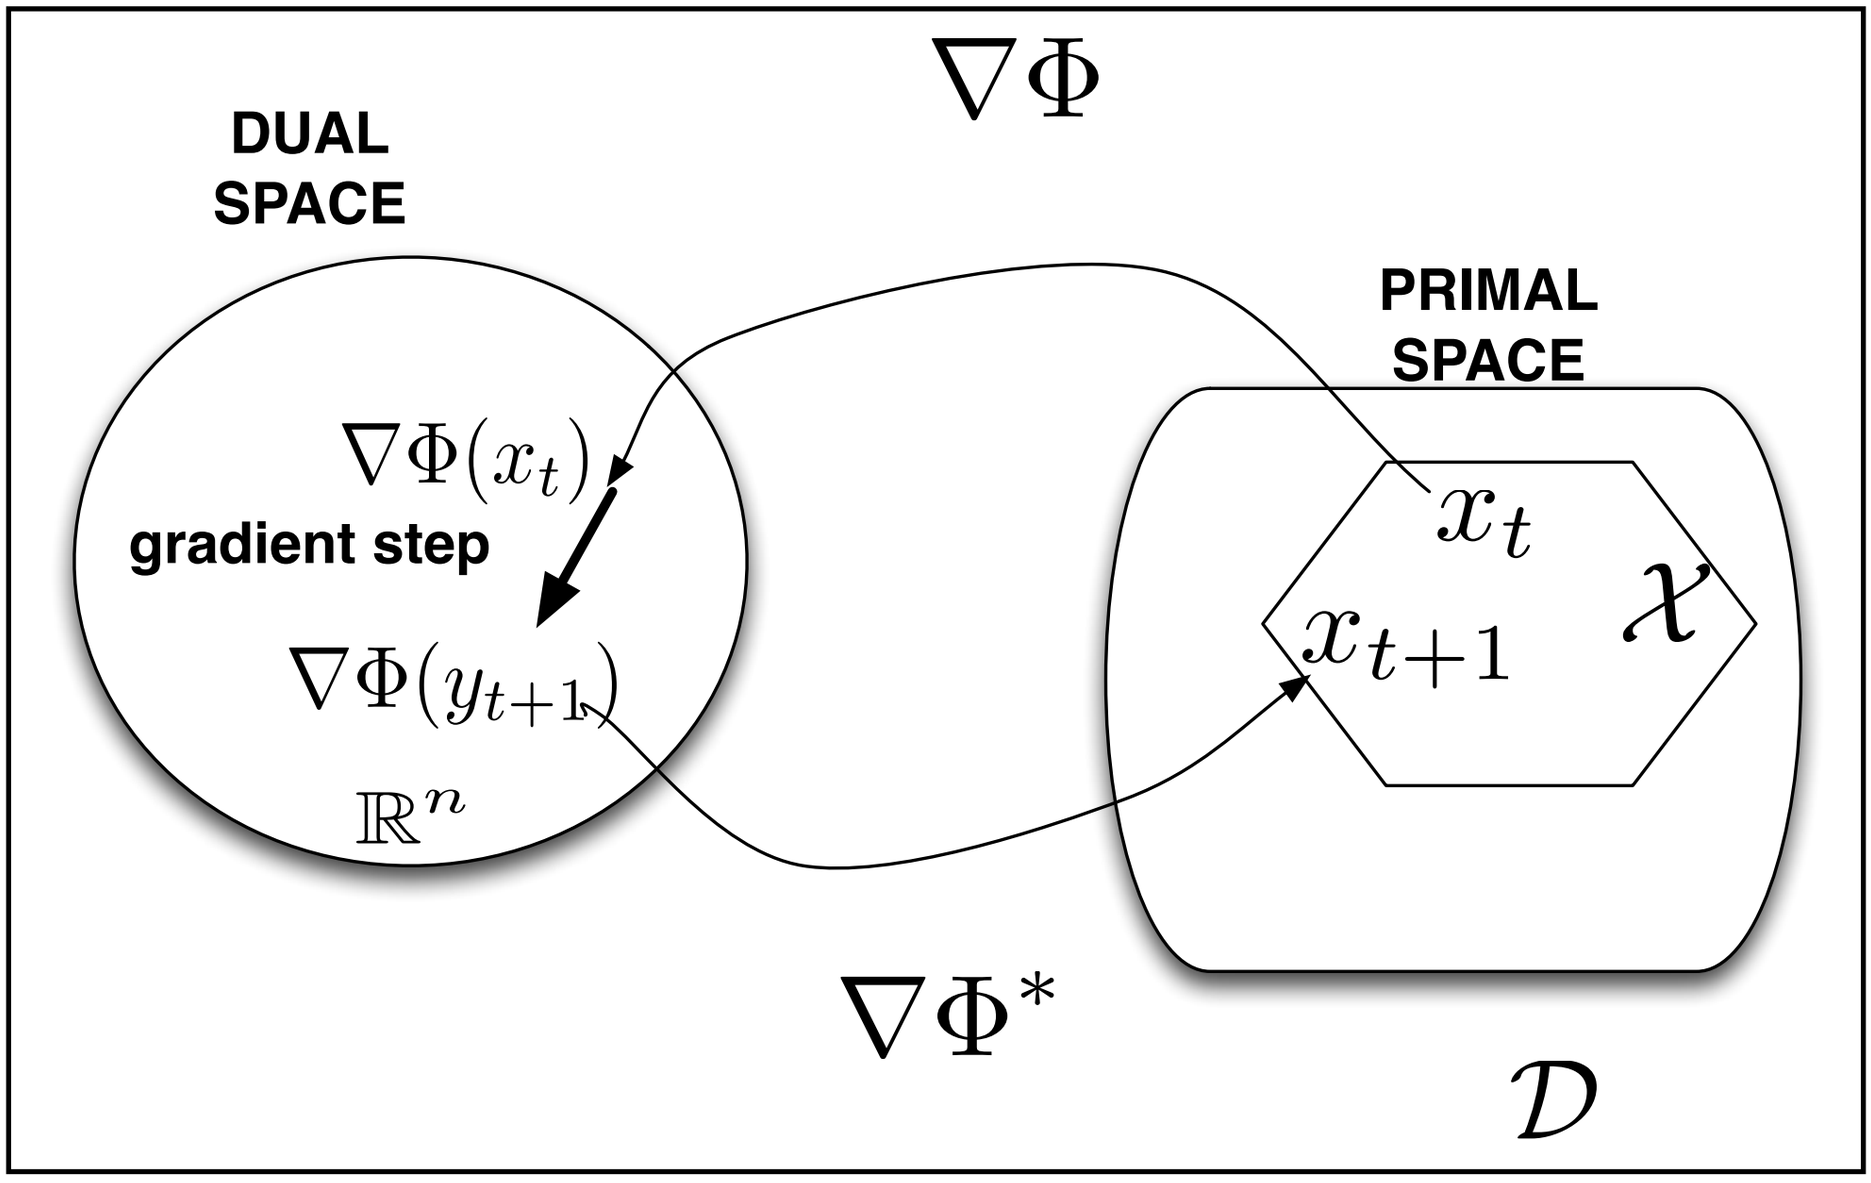

Mirror Descent = "Natural" Gradient

(Nemirovsky and Yudin; Amari, 1980s)

Mirror Descent

Mirror Map

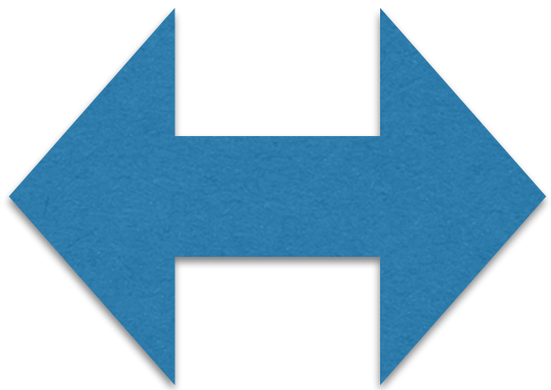

Natural gradient

Thomas, Dabney, Mahadevan, Giguere, NIPS 2013

Mirror-Prox: Non-Euclidean Extragradient

(Nemirovski, 2005)

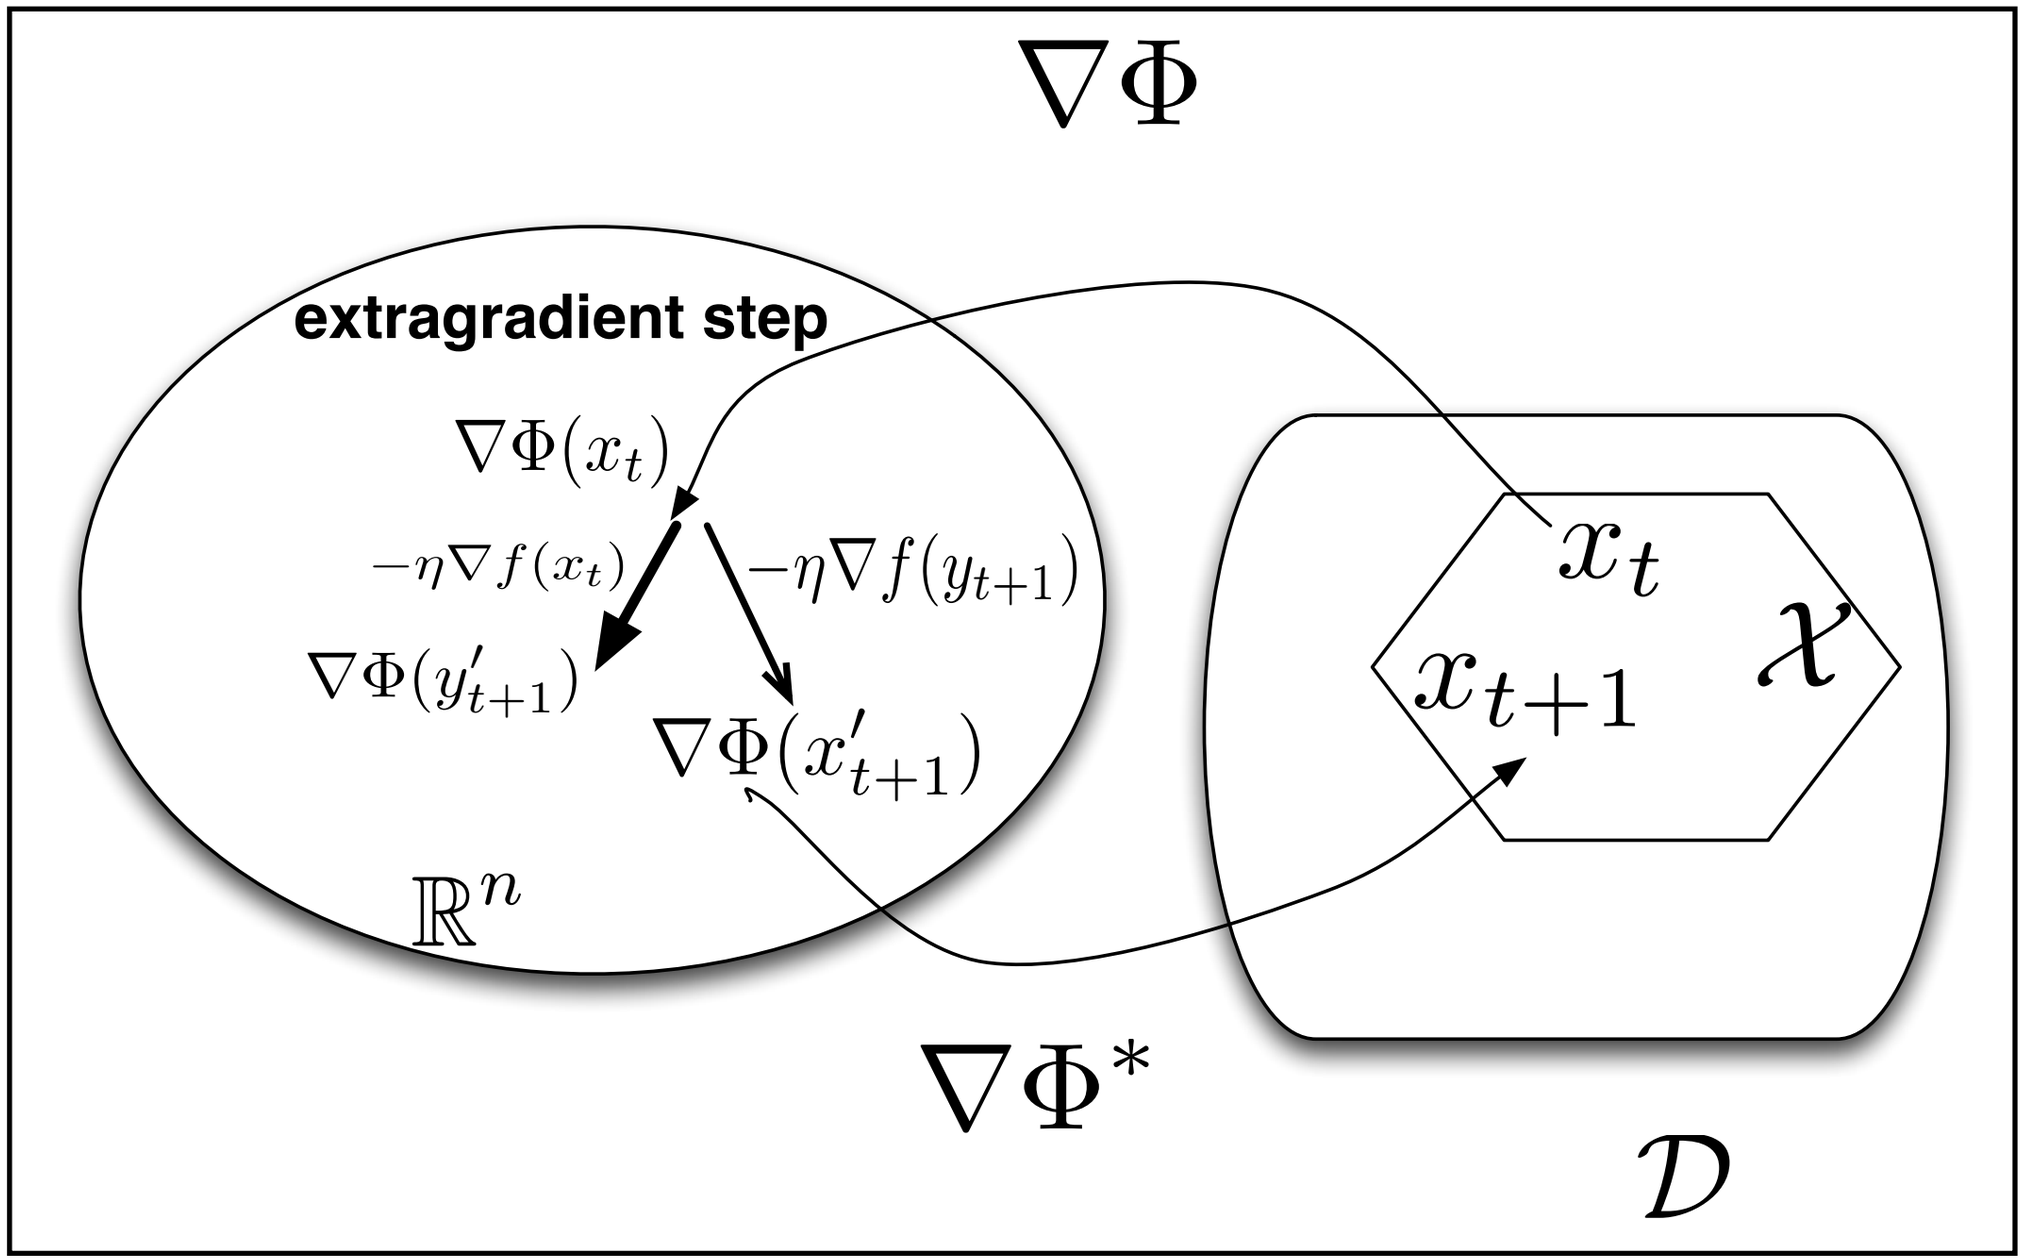

Our Latest VI Methods

(Gemp and Mahadevan, 2014)

General Runge-Kutta Mirror Descent (RKMDA)

$$
\begin{aligned}
k_1 &= \alpha_k F(x_k) \\
k_2 &= \alpha_k F(\nabla\psi_k^*(\nabla\psi_k(x_k) - a_{21}k_1)) \\
k_3 &= \alpha_k F(\nabla\psi_k^*(\nabla\psi_k(x_k) - a_{31}k_1 - a_{32}k_2)) \\
&\vdots \\
k_s &= \alpha_k F(\nabla\psi_k^*(\nabla\psi_k(x_k) - a_{s1}k_1 - a_{s2}k_2 - \ldots - a_{s,s-1}k_{s-1})) \\
\\
x_{k+1} &= \nabla\psi_k^*(\nabla\psi_k(x_k) - \Sigma_{i=1}^{s} b_i k_i)
\end{aligned}
$$

Runge-Kutta Method

ODE:

$$\frac{dy_i}{dx} = f_i(x, y_1, \ldots, y_n), \quad i = 1, \ldots, n$$

Euler's method:

$$y_{n+1} = y_n + hf(x_n, y_n)$$

RK(4) method:

$$\begin{aligned} k_1 &= hf(x_n, y_n), k_2 = hf(x_n + \frac{h}{2}, y_n + \frac{k_1}{2}) \\ k_3 &= hf(x_n + \frac{h}{2}, y_n + \frac{k_2}{2}), k_4 = hf(x_n + h, y_n + k_3) \\ y_{n+1} &= y_n + \frac{k_1}{6} + \frac{k_2}{3} + \frac{k_3}{3} + \frac{k_4}{6} \end{aligned}$$

Heun-Euler Adaptive VI Method

(Ian Gemp)

Heun-Euler Adaptive Non-Euclidean VI Method:

$k_1 = \alpha F(x_k)$

$k_2 = \alpha F(\nabla\psi_k^*(\nabla\psi_k(x_k) - k_1))$

$x^*_{k+1} = \nabla\psi_k^*(\nabla\psi_k(x_k) - k_1))$

$x_{k+1} = \nabla\psi_k^*(\nabla\psi_k(x_k) - \frac{1}{2}(k_1 + k_2))$

$\alpha = \alpha\left|\frac{\Delta_0}{\Delta_{k+1}}\right|^{1/2}$, where Δ_0 is the desired accuracy and

$\Delta_{k+1} = \max(x^*_{k+1} - x_{k+1})$.

Runge-Kutta Method for VIs

(Ian Gemp)

Runge Kutta (4) Gradient Descent

$$
\begin{aligned}
k_1 &= \alpha \nabla F(x_k) \\
k_2 &= \alpha \nabla F(x_k - \tfrac{1}{2} k_1) \\
k_3 &= \alpha \nabla F(x_k - \tfrac{1}{2} k_2) \\
k_4 &= \alpha \nabla F(x_k - k_3) \\
\\
x_{k+1} &= x_k - \tfrac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4)
\end{aligned}
$$

Runge Kutta (4) Non-Euclidean Extragradient

$$
\begin{aligned}
k_1 &= \alpha F(x_k) \\
k_2 &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - \tfrac{\alpha}{2} k_1)) \\
k_3 &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - \tfrac{\alpha}{2} k_2)) \\
k_4 &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - \alpha k_3)) \\
x_{k+1} &= \nabla \psi_k^*(\nabla \psi_k(x_k) - \tfrac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4))
\end{aligned}
$$

General Runge Kutta Gradient Descent

$$
\begin{aligned}
k_1 &= \alpha \nabla F(x_k) \\
k_2 &= \alpha \nabla F(x_k - a_{21} k_1) \\
k_3 &= \alpha \nabla F(x_k - a_{31} k_1 - a_{32} k_2) \\
&\vdots \\
k_s &= \alpha \nabla F(x_k - a_{s1} k_1 - a_{s2} k_2 - \ldots - a_{s,s-1} k_{s-1}) \\
\\
x_{k+1} &= x_k - \Sigma_{i=1}^{s} b_i k_i
\end{aligned}
$$

General RK Non-Euclidean Extragradient

$$
\begin{aligned}
k_1 &= \alpha F(x_k) \\
k_2 &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{21} k_1)) \\
k_3 &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{31} k_1 - a_{32} k_2)) \\
&\vdots \\
k_s &= \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{s1} k_1 - a_{s2} k_2 - \ldots - a_{s,s-1} k_{s-1})) \\
\\
x_{k+1} &= \nabla \psi_k^*(\nabla \psi_k(x_k) - \Sigma_{i=1}^{s} b_i k_i)
\end{aligned}
$$

Plan of the Tutorial

- Part I: Motivation
- Part II: Mathematical Background
- Part III: Algorithms
- **Part IV: Applications and Challenges**

Applications

- Network equilibrium Internet models
- Computational sustainability
- Imagination using GAN models
- Reinforcement learning and causal inference

NOVEL GAN ALGORITHMS BASED ON VI

- Joint work with my PhD student Ian Gemp

- In equilibrium problems, traversing the steepest descent direction is not sufficient for convergence

- **Orthogonal** directions turn out to be crucial for convergence

VECTOR FIELD GAN ALGORITHM

Example: $\begin{bmatrix} x^{(1)} \\ x^{(2)} \end{bmatrix} \leftarrow \begin{bmatrix} x^{(1)} - \rho \nabla_{x^{(1)}} f^{(1)}(x) \\ x^{(2)} - \rho \nabla_{x^{(2)}} f^{(2)}(x) \end{bmatrix}$ Simultaneous Gradient Descent

Player **1**'s loss is: $f^{(1)}(x) = 10x^{(1)}x^{(2)}$.

Player **2**'s loss is: $f^{(2)}(x) = -10x^{(1)}x^{(2)}$.

Let $x = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ and $\rho = 0.1$.

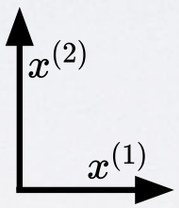

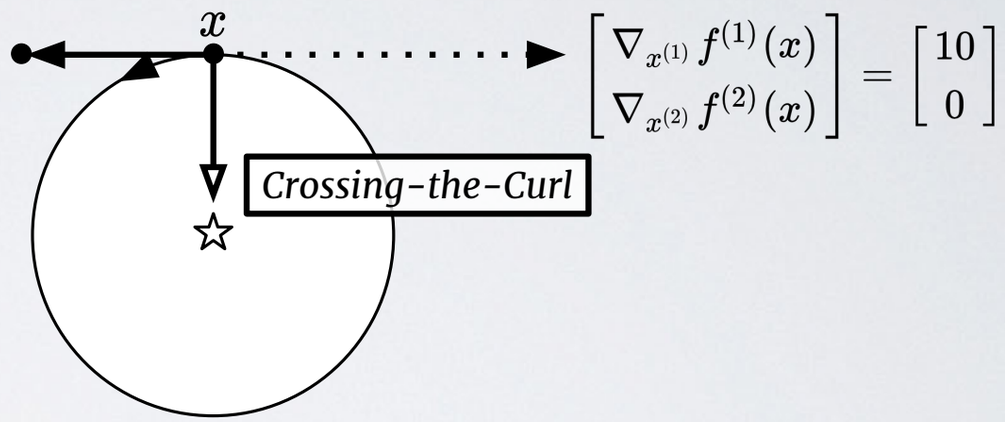

$\begin{bmatrix} \nabla_{x^{(1)}} f^{(1)}(x) \\ \nabla_{x^{(2)}} f^{(2)}(x) \end{bmatrix} = \begin{bmatrix} 10 \\ 0 \end{bmatrix}$

CAN: Creative Adversarial Networks Generating "Art" by Learning About Styles and Deviating from Style Norms*

Ahmed Elgammal[1†] Bingchen Liu[1] Mohamed Elhoseiny[2] Marian Mazzone[3]
The Art & AI Laboratory - Rutgers University
[1] Department of Computer Science, Rutgers University, NJ, USA
[2] Facebook AI Research, CA, USA
[3] Department of Art History, College of Charleston, SC, USA

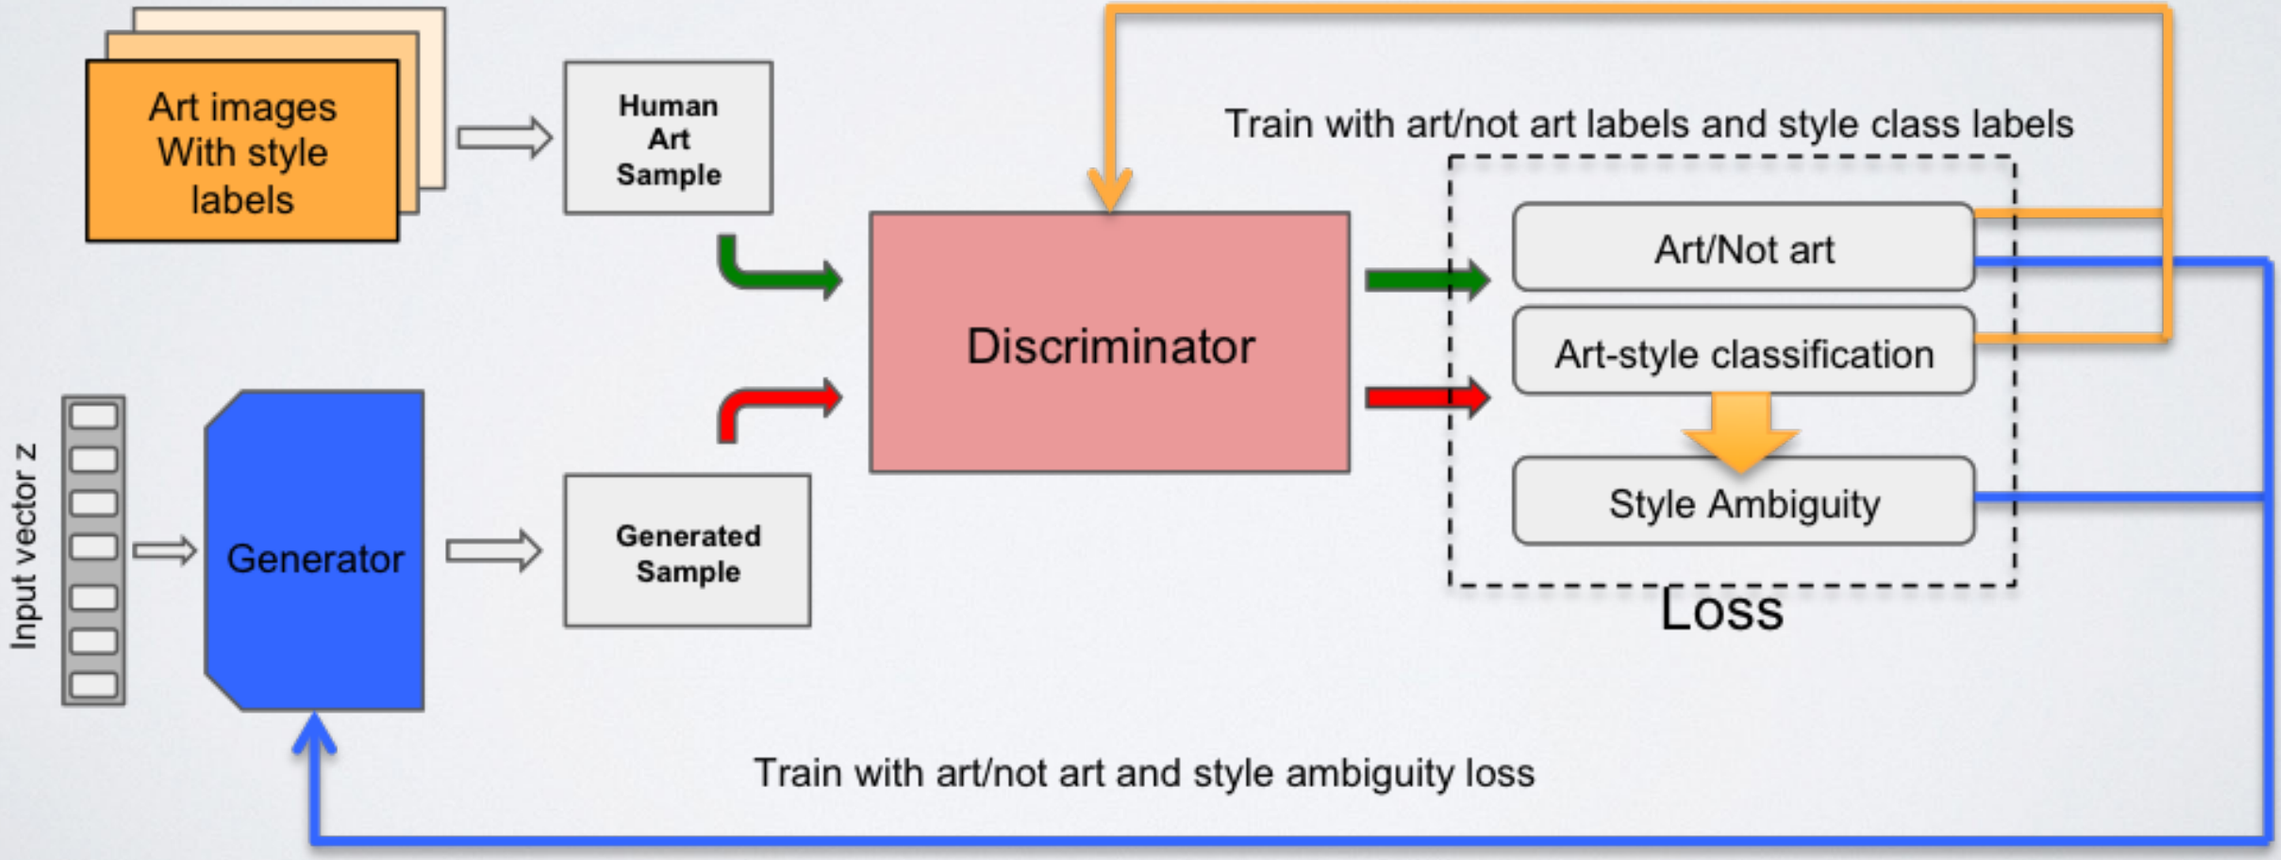

Sample art
created by CAN

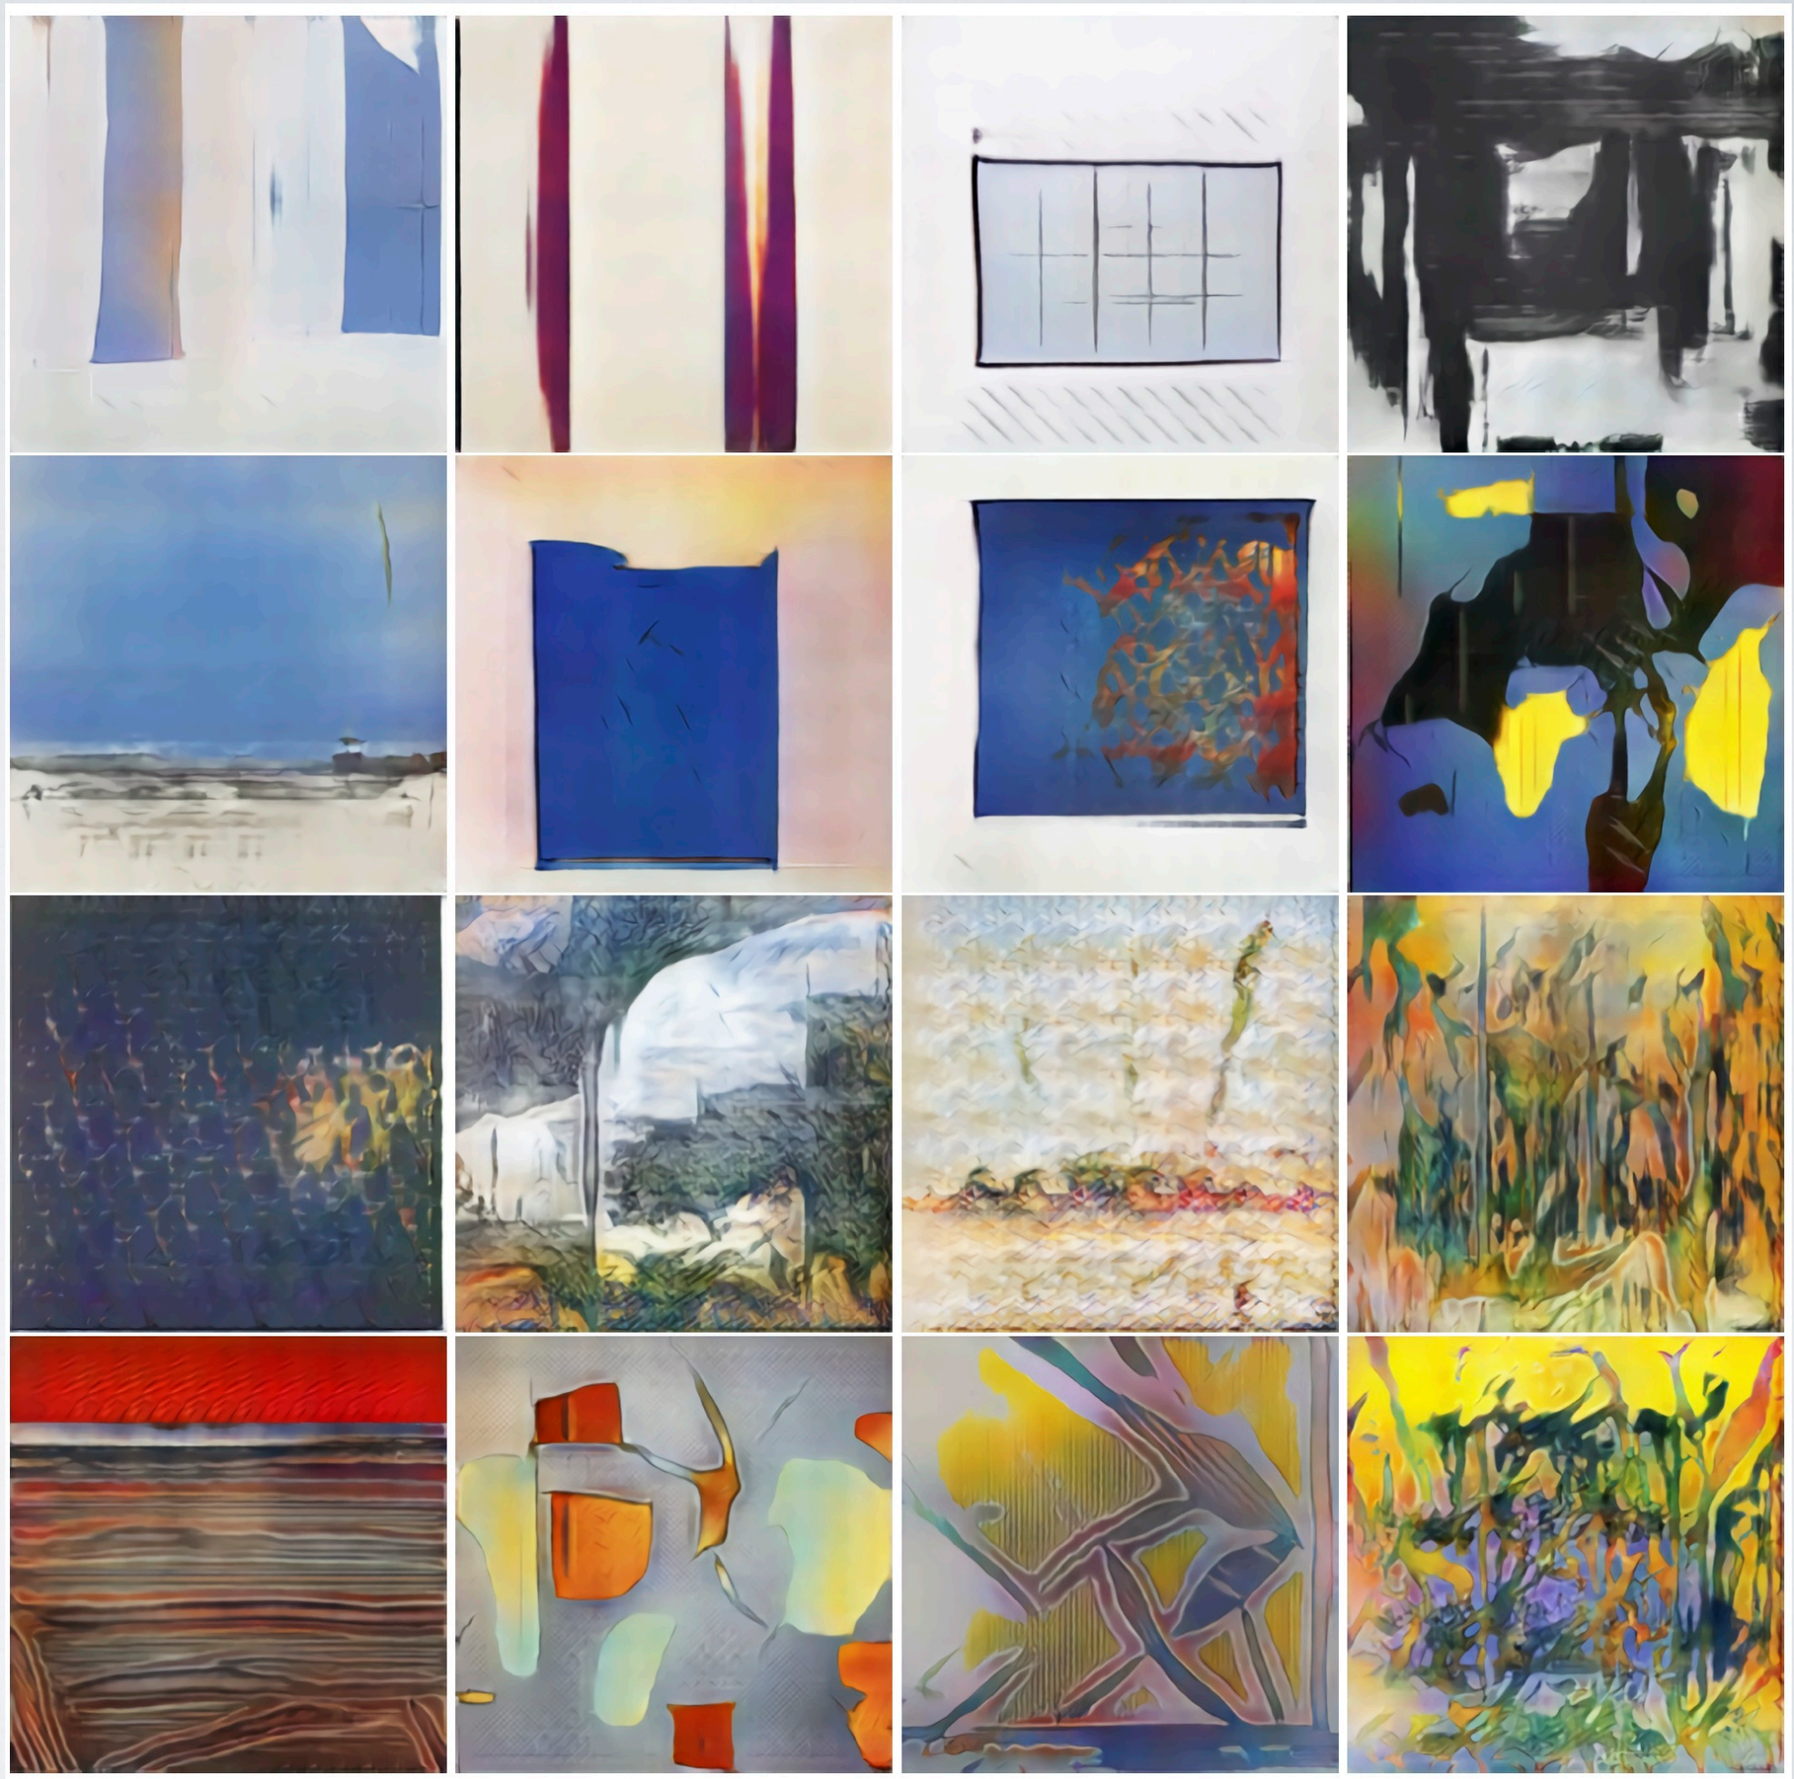

CAN loss function

$$\min_G \max_D V(D,G) =$$

$$\mathbb{E}_{x,\hat{c}\sim p_{data}}[\log D_r(x) + \log D_c(c=\hat{c}|x)] +$$

$$\mathbb{E}_{z\sim p_z}[\log(1 - D_r(G(z))) - \sum_{k=1}^{K}\big(\frac{1}{K} log(D_c(c_k|G(z))+$$

$$(1-\frac{1}{K})log(1 - D_c(c_k|G(z)))\big)],$$

MIX GAN MODELS

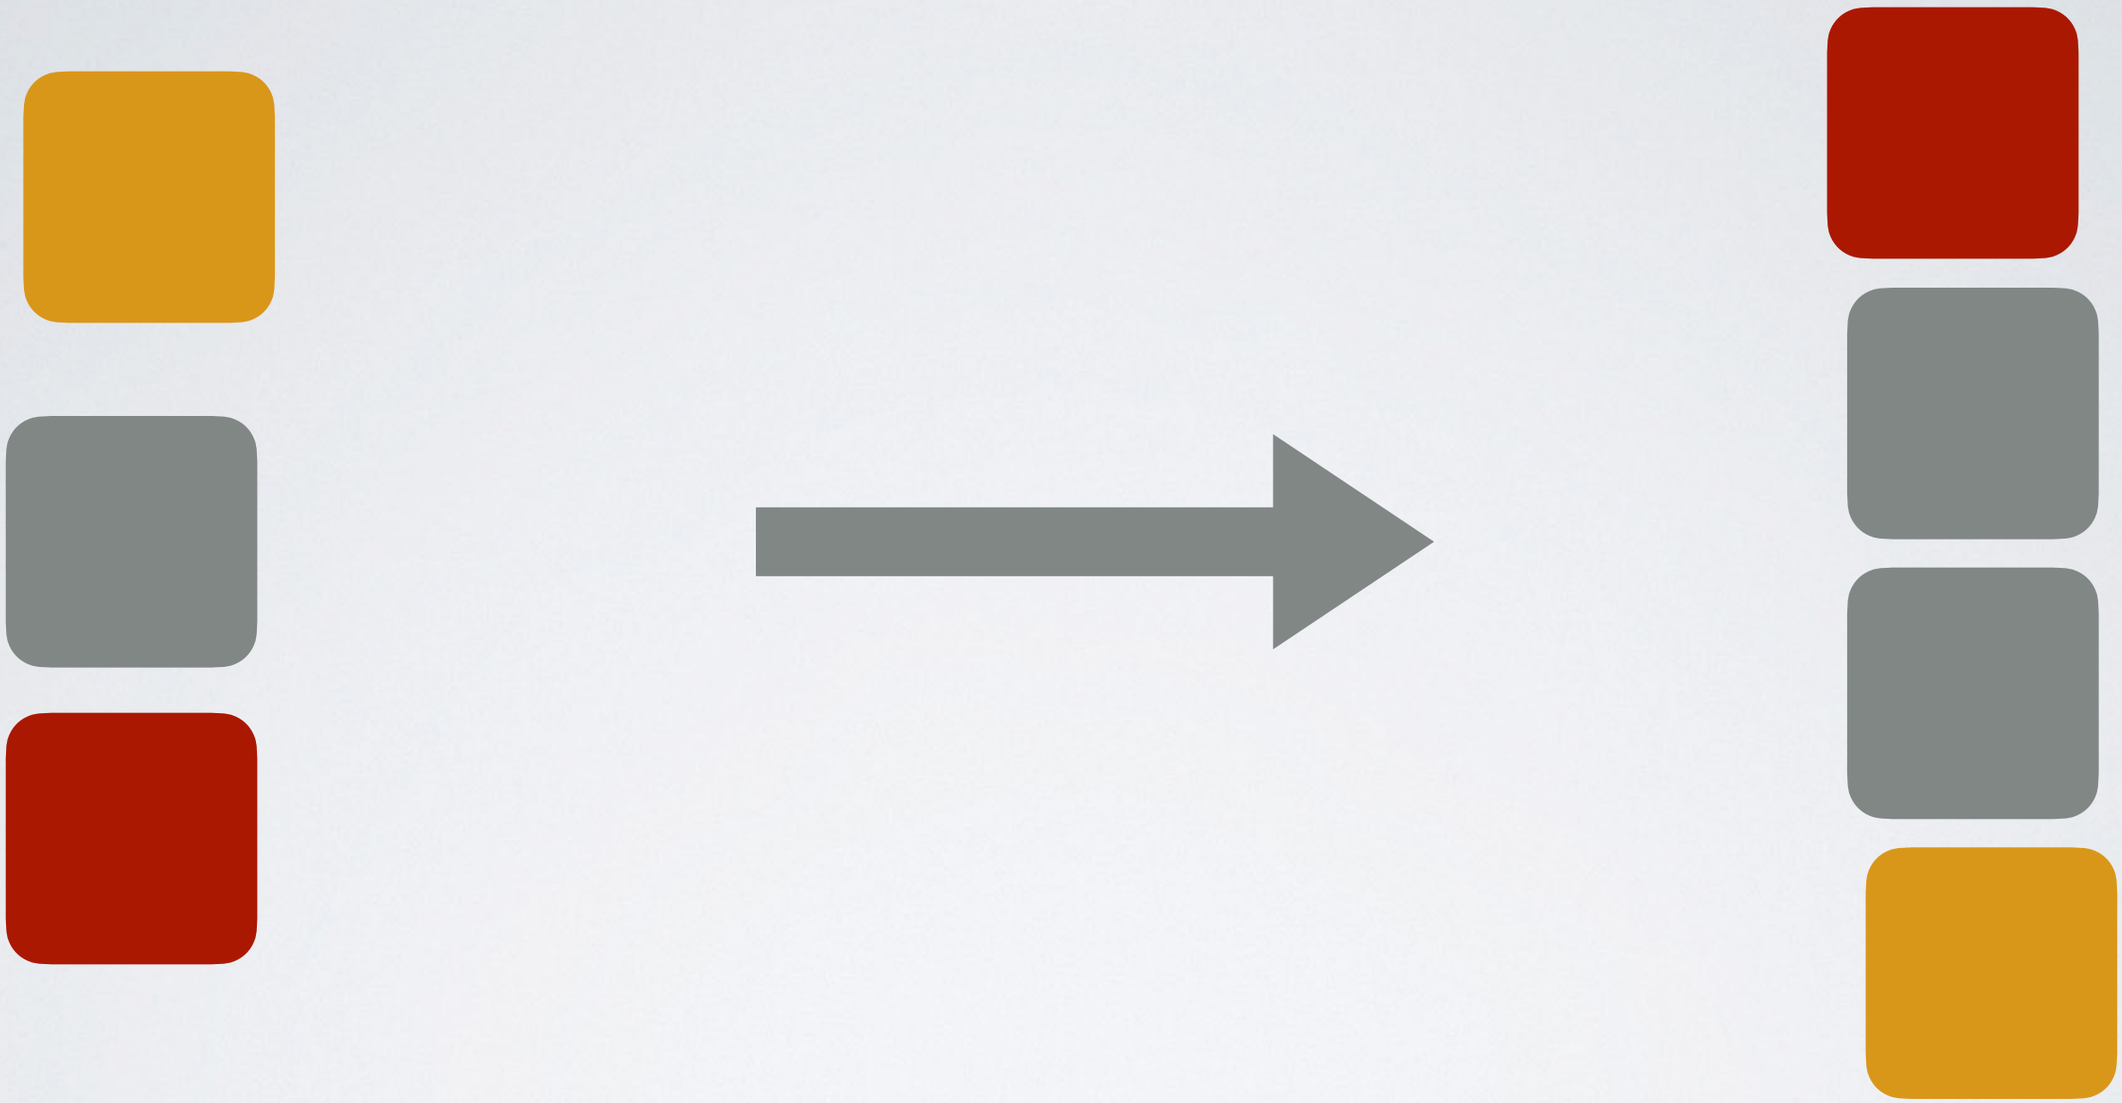

Generators

Discriminators

(Arora et al., ICML 2017)

Next Generation Internet Model [Nagurney et al., 2014]

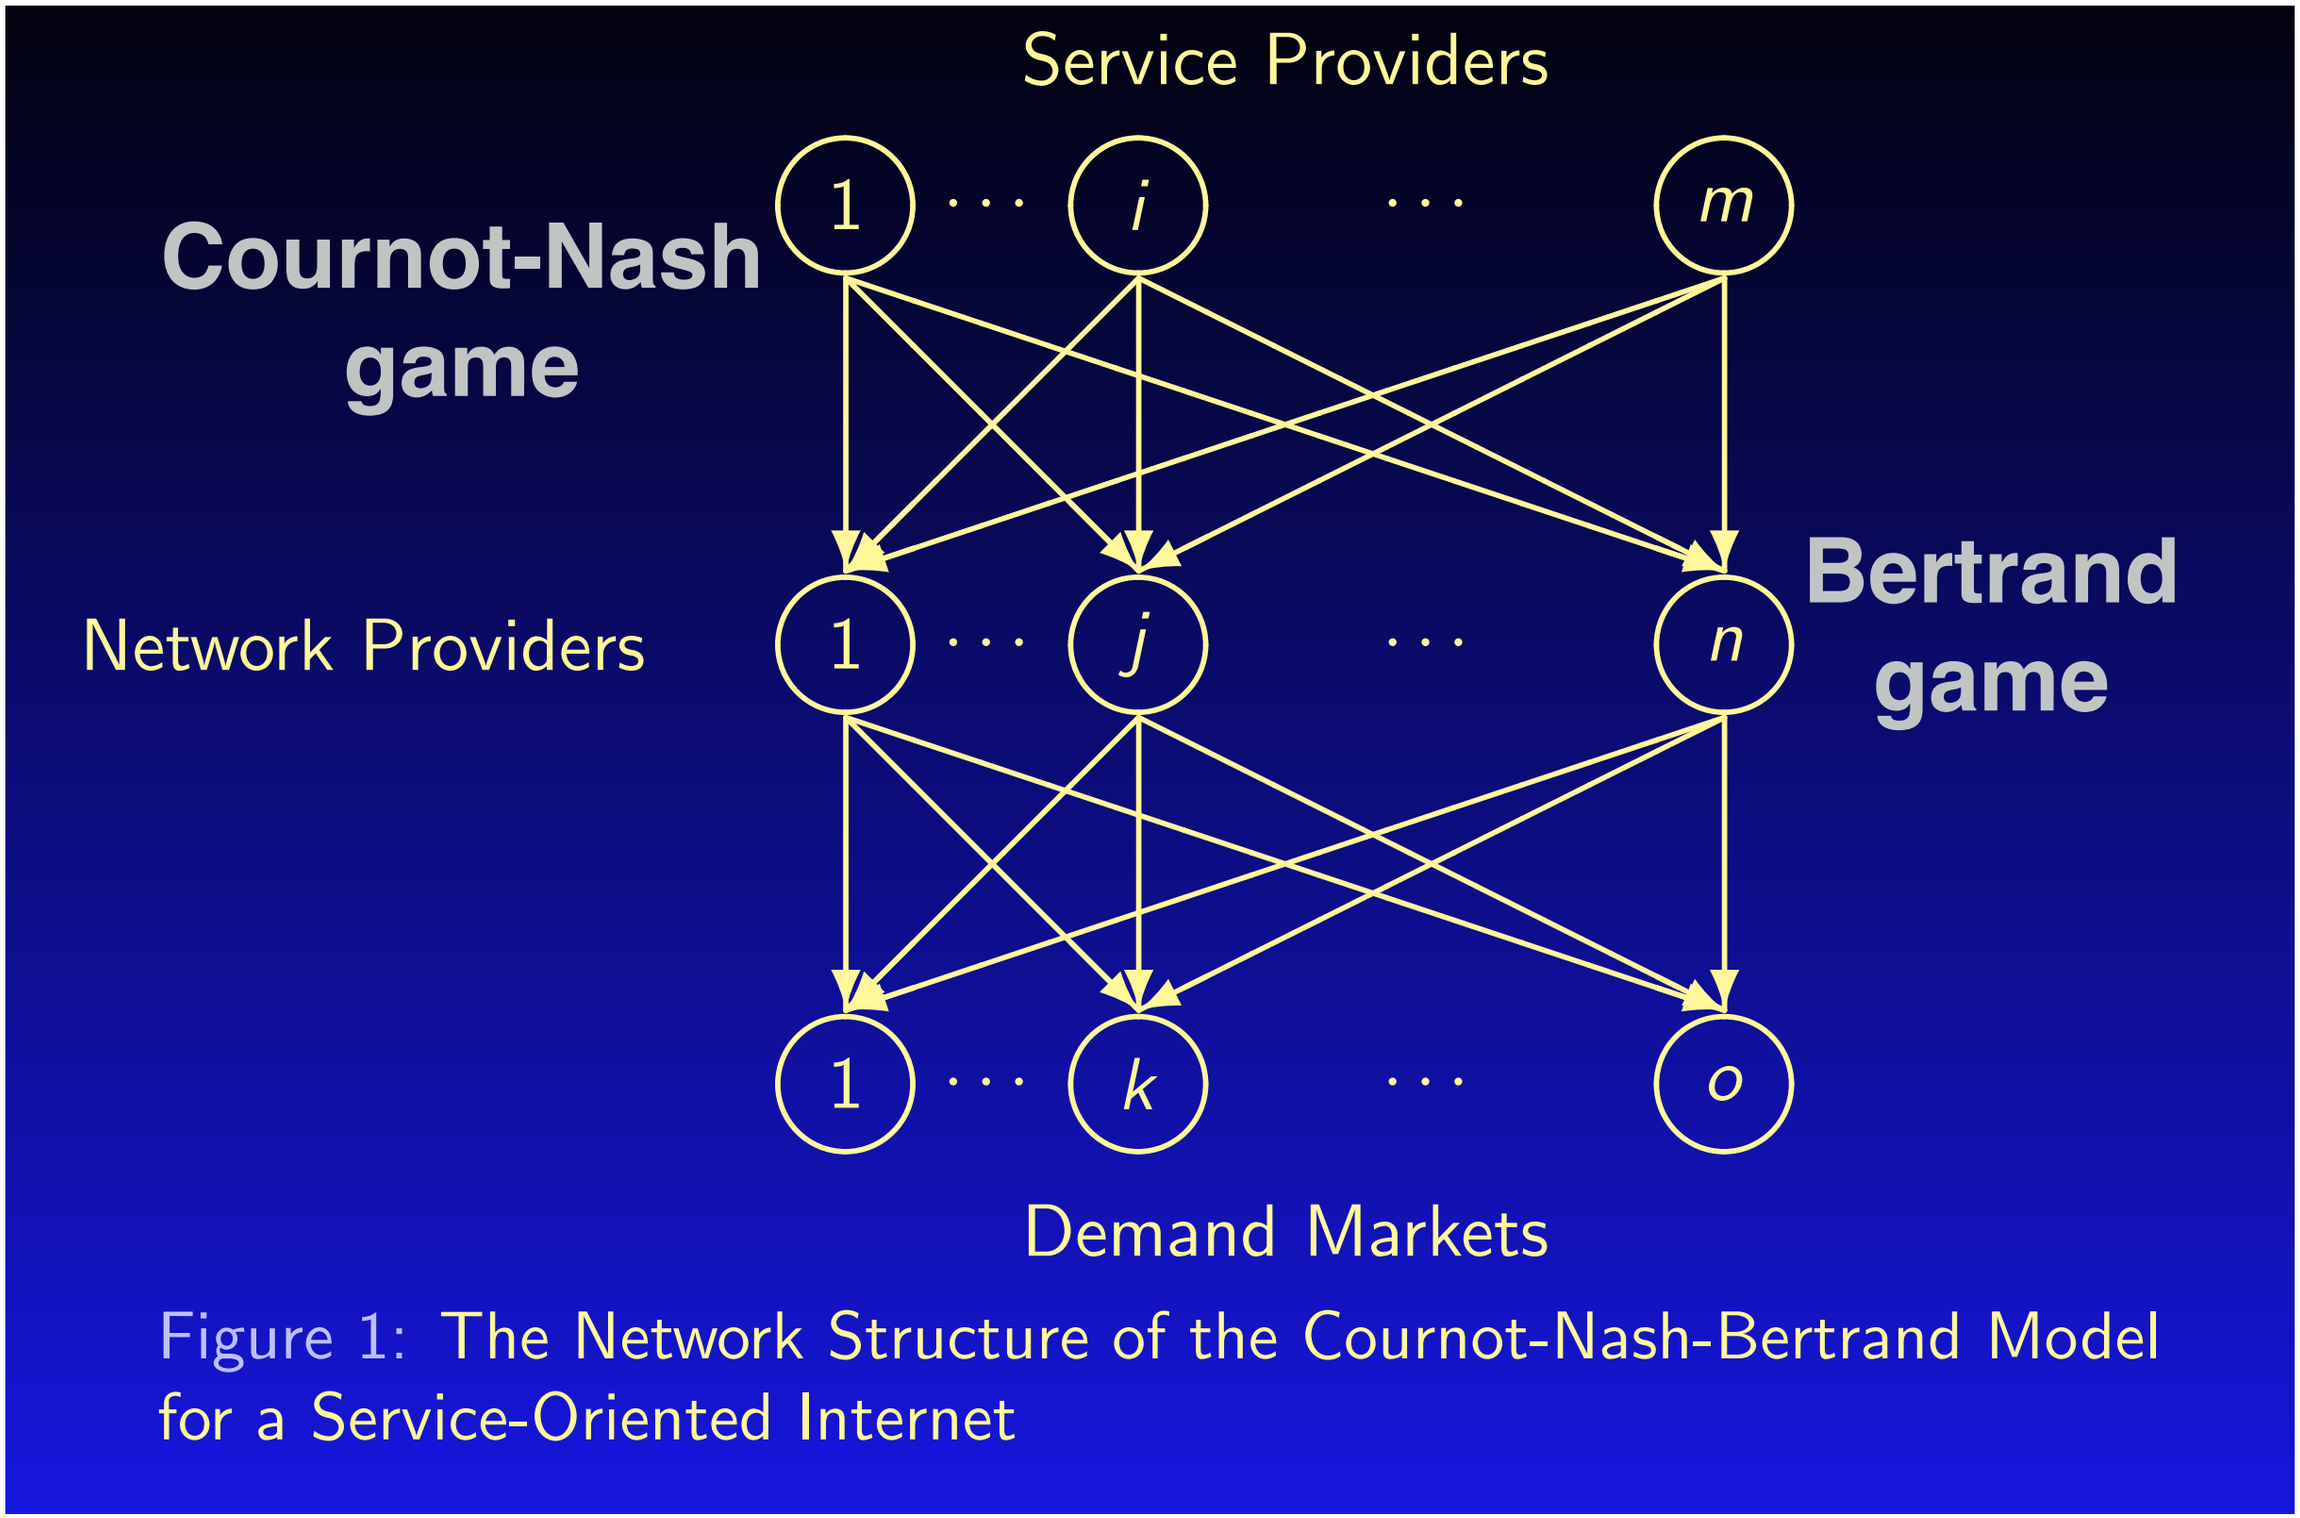

Figure 1: The Network Structure of the Cournot-Nash-Bertrand Model for a Service-Oriented Internet

Proximal Reinforcement Learning: A New Theory of Sequential Decision Making in **Primal-Dual** Spaces,

Arxiv, May 26, 2014 (126 pages)

Sridhar Mahadevan, Bo Liu, Philip Thomas, Will Dabney,
Stephen Giguere,
Nicholas Jacek, Ian Gemp,
Ji Liu

True Gradient TD-Learning: RL meets VI

Algorithm 2 GTD2-MP (2014)

$$
\begin{aligned}
(1)\quad & w_{t+\frac{1}{2}} = w_t + \beta_t(\delta_t - \phi_t^T w_t)\phi_t, \\
& \theta_{t+\frac{1}{2}} = \text{prox}_{\alpha_t h}\left(\theta_t + \alpha_t(\phi_t - \gamma\phi_t')(\phi_t^T w_t)\right) \\
(2)\quad & \delta_{t+\frac{1}{2}} = r_t + \gamma\phi_t'^T\theta_{t+\frac{1}{2}} - \phi_t^T\theta_{t+\frac{1}{2}} \\
(3)\quad & w_{t+1} = w_t + \beta_t(\delta_{t+\frac{1}{2}} - \phi_t^T w_{t+\frac{1}{2}})\phi_t\ , \\
& \theta_{t+1} = \text{prox}_{\alpha_t h}\left(\theta_t + \alpha_t(\phi_t - \gamma\phi_t')(\phi_t^T w_{t+\frac{1}{2}})\right)
\end{aligned}
$$

(Mahadevan et al., Arxiv 2014; Liu et al., JAIR 2019)

Baird counter example

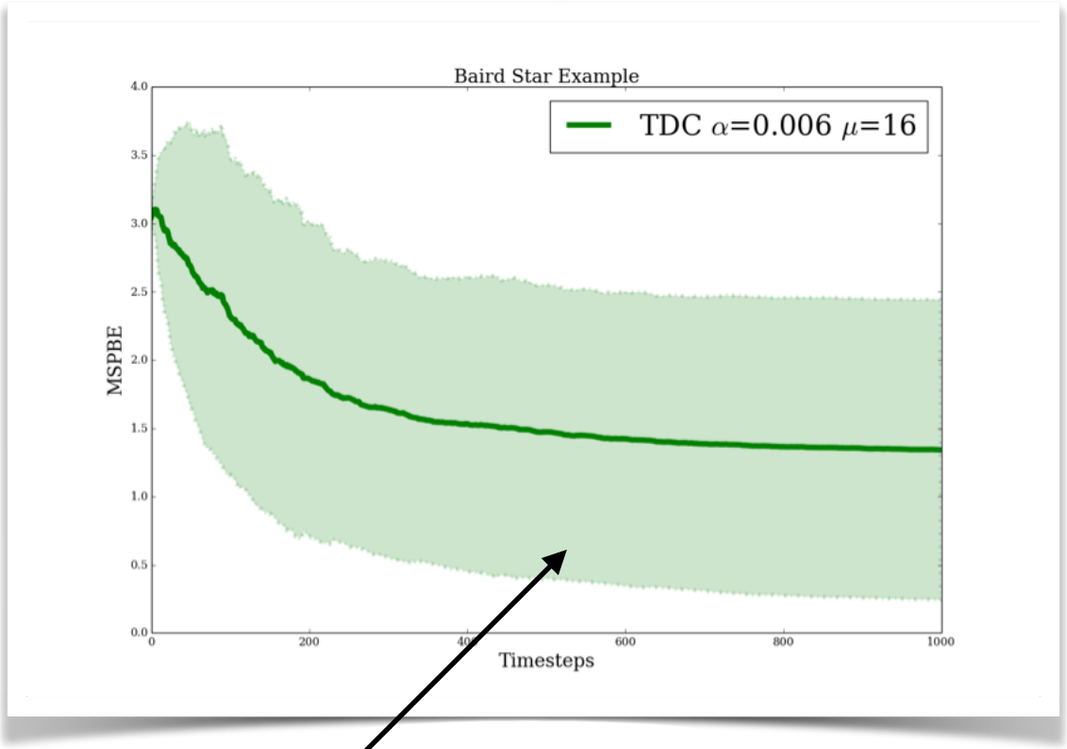

Sutton et al., 2009

$$\text{Variance of previous methods: } \propto O(\frac{1}{\sqrt{n}})$$

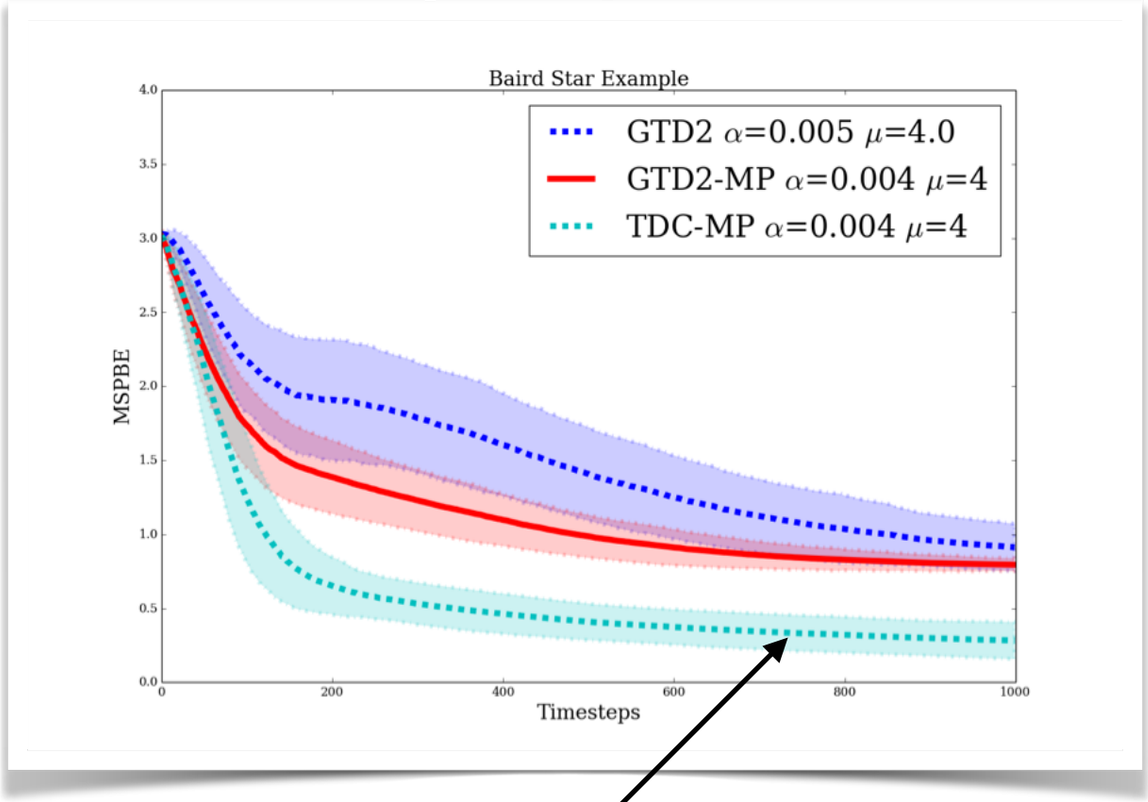

Our new methods 2014

$$\text{Variance of our methods: } \propto O(\frac{1}{n})$$

Problem Formulation

Table 1: Notation for the Game Theoretic Cournot-Nash-Bertrand Model

Notation	**Definition**
Q_{ijk}	the nonnegative service volume from i to k via j. We group the $\{Q_{ijk}\}$ elements for all j and k into the vector $Q_i \in R_+^{no}$ and then we group all the vectors Q_i for all i into the vector $Q \in R_+^{mno}$.
s_i	the service volume (output) produced by service provider i. We group the $\{s_i\}$ elements into the vector $s \in R_+^m$.
q_{ijk}	the nonnegative quality level of network provider j transporting service i to k. We group the q_{ijk} for all i and k into the vector $q_j \in R_+^{mo}$ and all the vectors q_j for all j into the vector $q \in R_+^{mno}$.
π_{ijk}	the price charged by network provider j for transporting a unit of service provided by i via j to k. We group the π_{ijk} for all i and k into the vector $\pi_j \in R_+^m$ and then we group all the vectors π_j for all j into the vector $\pi \in R_+^{mno}$.
$f_i(s)$	the total production cost of service provider i.
$\hat{\rho}_{ijk}(Q, q)$	the demand price at k associated with service i transported via j.
$c_{ijk}(Q, q)$	the total transportation cost associated with delivering service i via j to k.
$oc_{ijk}(\pi_{ijk})$	the opportunity cost associated with pricing by network provider j services transported from i to k.

VI Fomulation

$$\langle F(X^*), X - X^* \rangle \geq 0, \quad \forall X \in \mathcal{K},$$

$$X \equiv (Q, q, \pi)$$

- Production cost function *f*(Q) - cost of providing a certain volume of content

- Demand price function - user offer depends on content quality and market volume

$$F^1_{ijk}(X) = \frac{\partial \hat{f}_i(Q)}{\partial Q_{ijk}} + \pi_{ijk} - \hat{\rho}_{ijk}(Q, q) - \sum_{h=1}^{n} \sum_{l=1}^{o} \frac{\partial \hat{\rho}_{ihl}(Q, q)}{\partial Q_{ijk}} \times Q_{ihl},$$

$$F^2_{ijk}(X) = \sum_{h=1}^{m} \sum_{l=1}^{o} \frac{\partial c_{hjl}(Q, q)}{\partial q_{ijk}},$$

$$F^3_{ijk}(X) = -Q_{ijk} + \frac{\partial oc_{ijk}(\pi_{ijk})}{\partial \pi_{ijk}}.$$

Results on Internet VI Problem

Extragradient

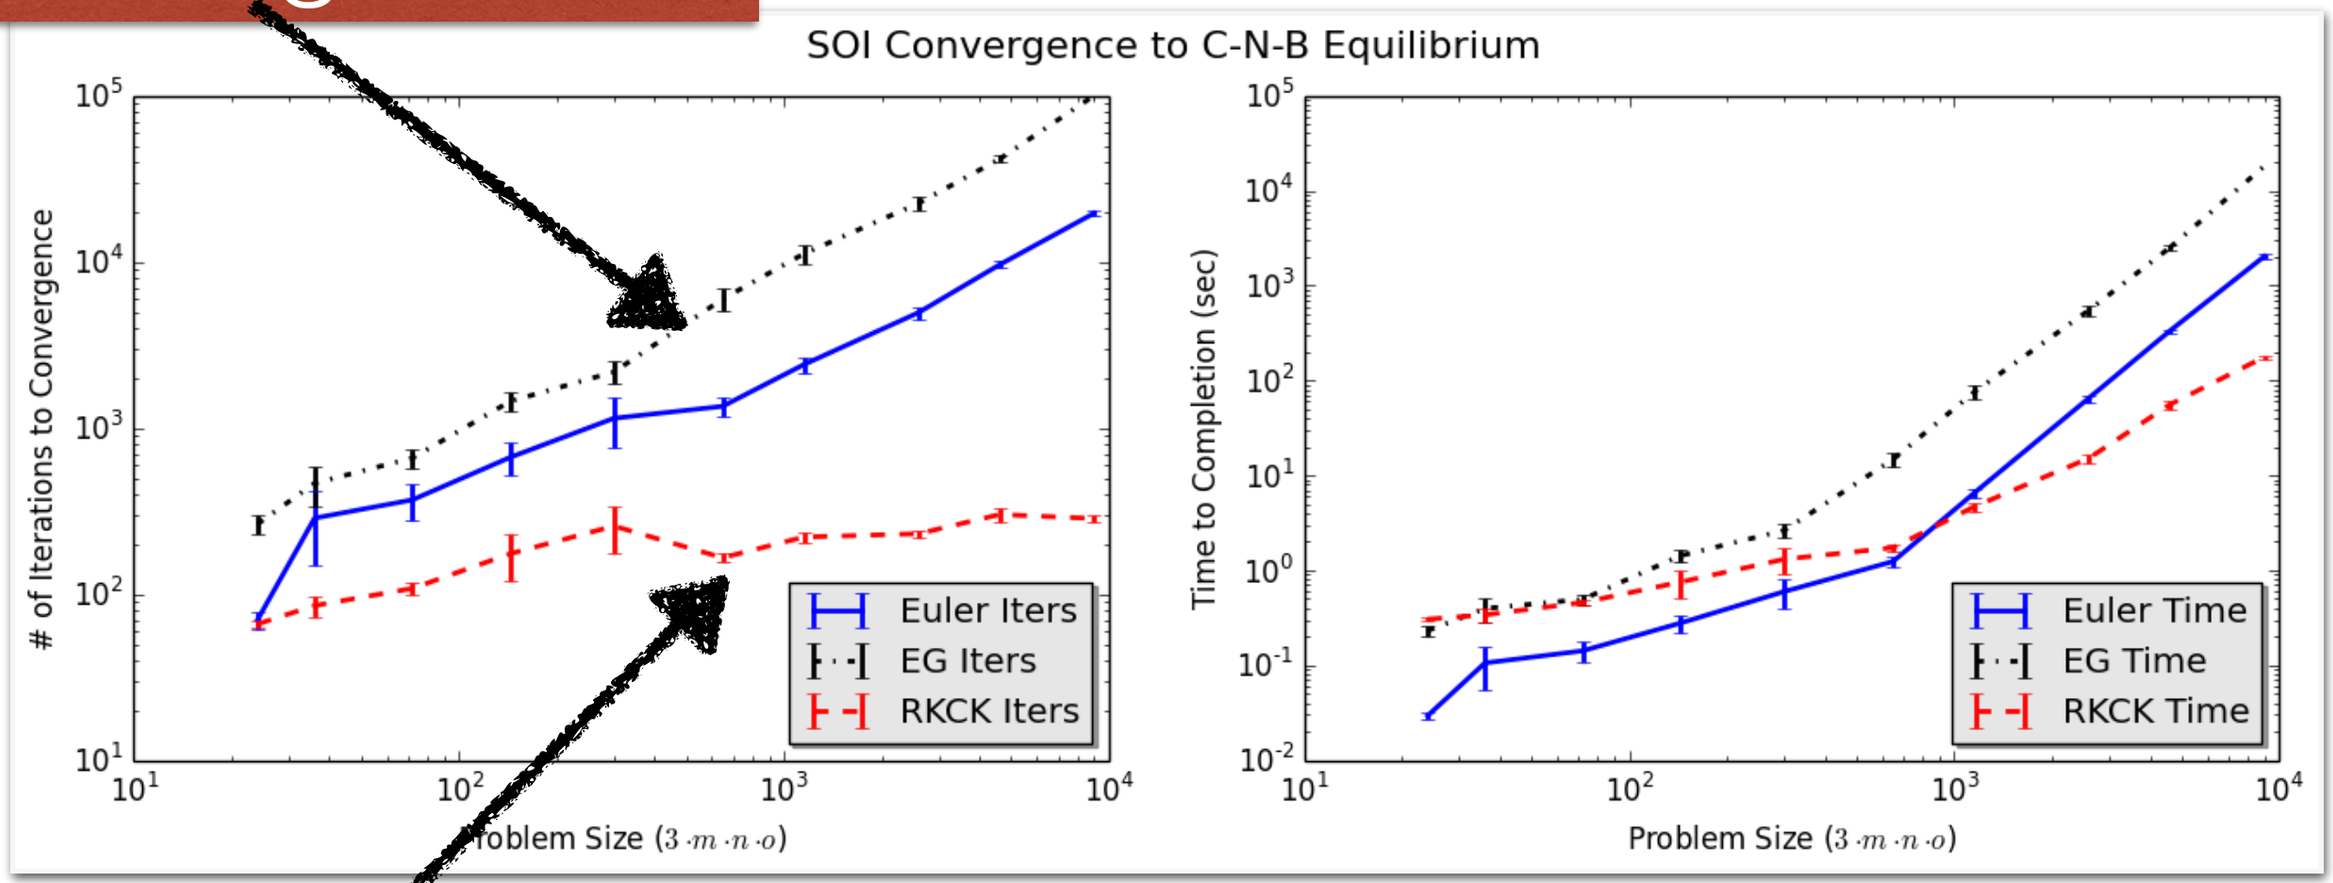

Our new algorithm

Substainable Blood Banking

(Nagurney and Masoumi, 2012)

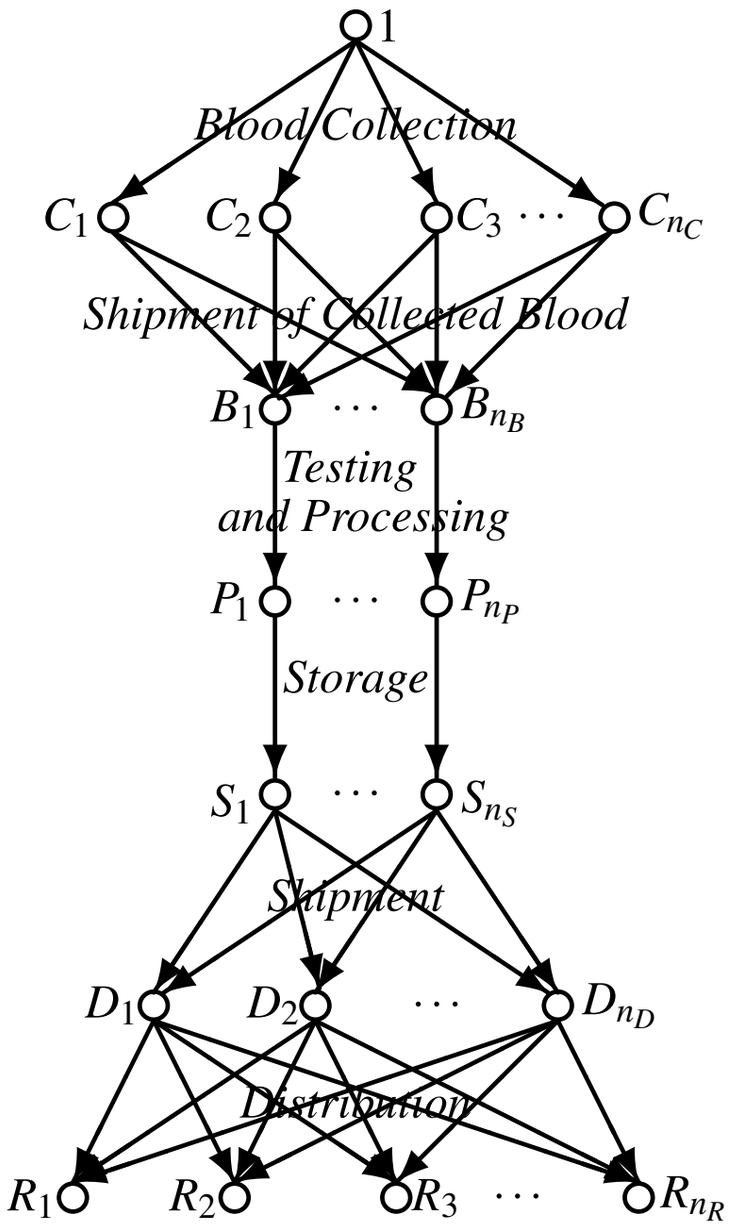

The Organization

The Blood Collection Sites

The Blood Centers

The Component Labs

The Storage Facilities

The Distribution Centers

The Demand Points

Sustainable Blood Banking

$$\text{Minimize} \quad \sum_{a\in L} \hat{c}_a(f_a) + \sum_{a\in L} \hat{z}_a(f_a) + \sum_{a\in L} \hat{\pi}_a(u_a)$$

$$+ \sum_{k=1}^{n_R} \left(\lambda_k^- E(\Delta_k^-) + \lambda_k^+ E(\Delta_k^+)\right) + \theta \sum_{a\in L_1} \hat{r}_a(f_a)$$

such that

$$x_p \geq 0, \qquad \forall p \in \mathscr{P}$$

$$v_k \equiv \sum_{p\in\mathscr{P}_{w_k}} x_p \mu_p, \qquad k = 1,\ldots,n_R.$$

$$f_a = \sum_{p\in\mathscr{P}} x_p \ \alpha_{ap}, \qquad \forall a \in L.$$

$$f_a \leq \bar{u}_a + u_a, \qquad \forall a \in L,$$

$$-\bar{u}_a \leq u_a, \qquad \forall a \in L.$$

Sustainable Blood Banking

$$\langle F(X^*), X - X^* \rangle \geq 0, \forall X \in \mathcal{K}, \text{ where } X = (x, \mu, \gamma) \in \mathbb{R}_+^{N_X}$$

$$F^1(X) = \frac{\partial(\sum_{q \in \mathcal{P}} \hat{C}_q(x))}{\partial x_p} + \frac{\partial(\sum_{q \in \mathcal{P}} \hat{Z}_q(x))}{\partial x_p}$$

$$+\lambda_k^+ \mu_p P_k\big(\sum_{p \in \mathcal{P}_{w_k}} x_p \mu_p\big) - \lambda_k^- \mu_p \big(1 - P_k\big(\sum_{p \in \mathcal{P}_{w_k}} x_p \mu_p\big)\big)$$

$$+\sum_{a \in L} \gamma_a \delta_{ap} + \theta \frac{\partial(\sum_{q \in \mathcal{P}} \hat{R}_q(x))}{\partial x_p}; p \in P_{w_k}; k = 1, \ldots, n_R$$

$$F^2(X) = \frac{\partial \hat{\pi}_a(u_a)}{\partial u_a} - \gamma_a; a \in L$$

$$F^3(X) = \bar{u}_a + u_a - \sum_{p \in \mathcal{P}} x_p \alpha_{ap}; a \in L$$

Experimental Results

(Gemp and Mahadevan, AAAI Workshop on Computational Sustainability, 2015)

Extragradient

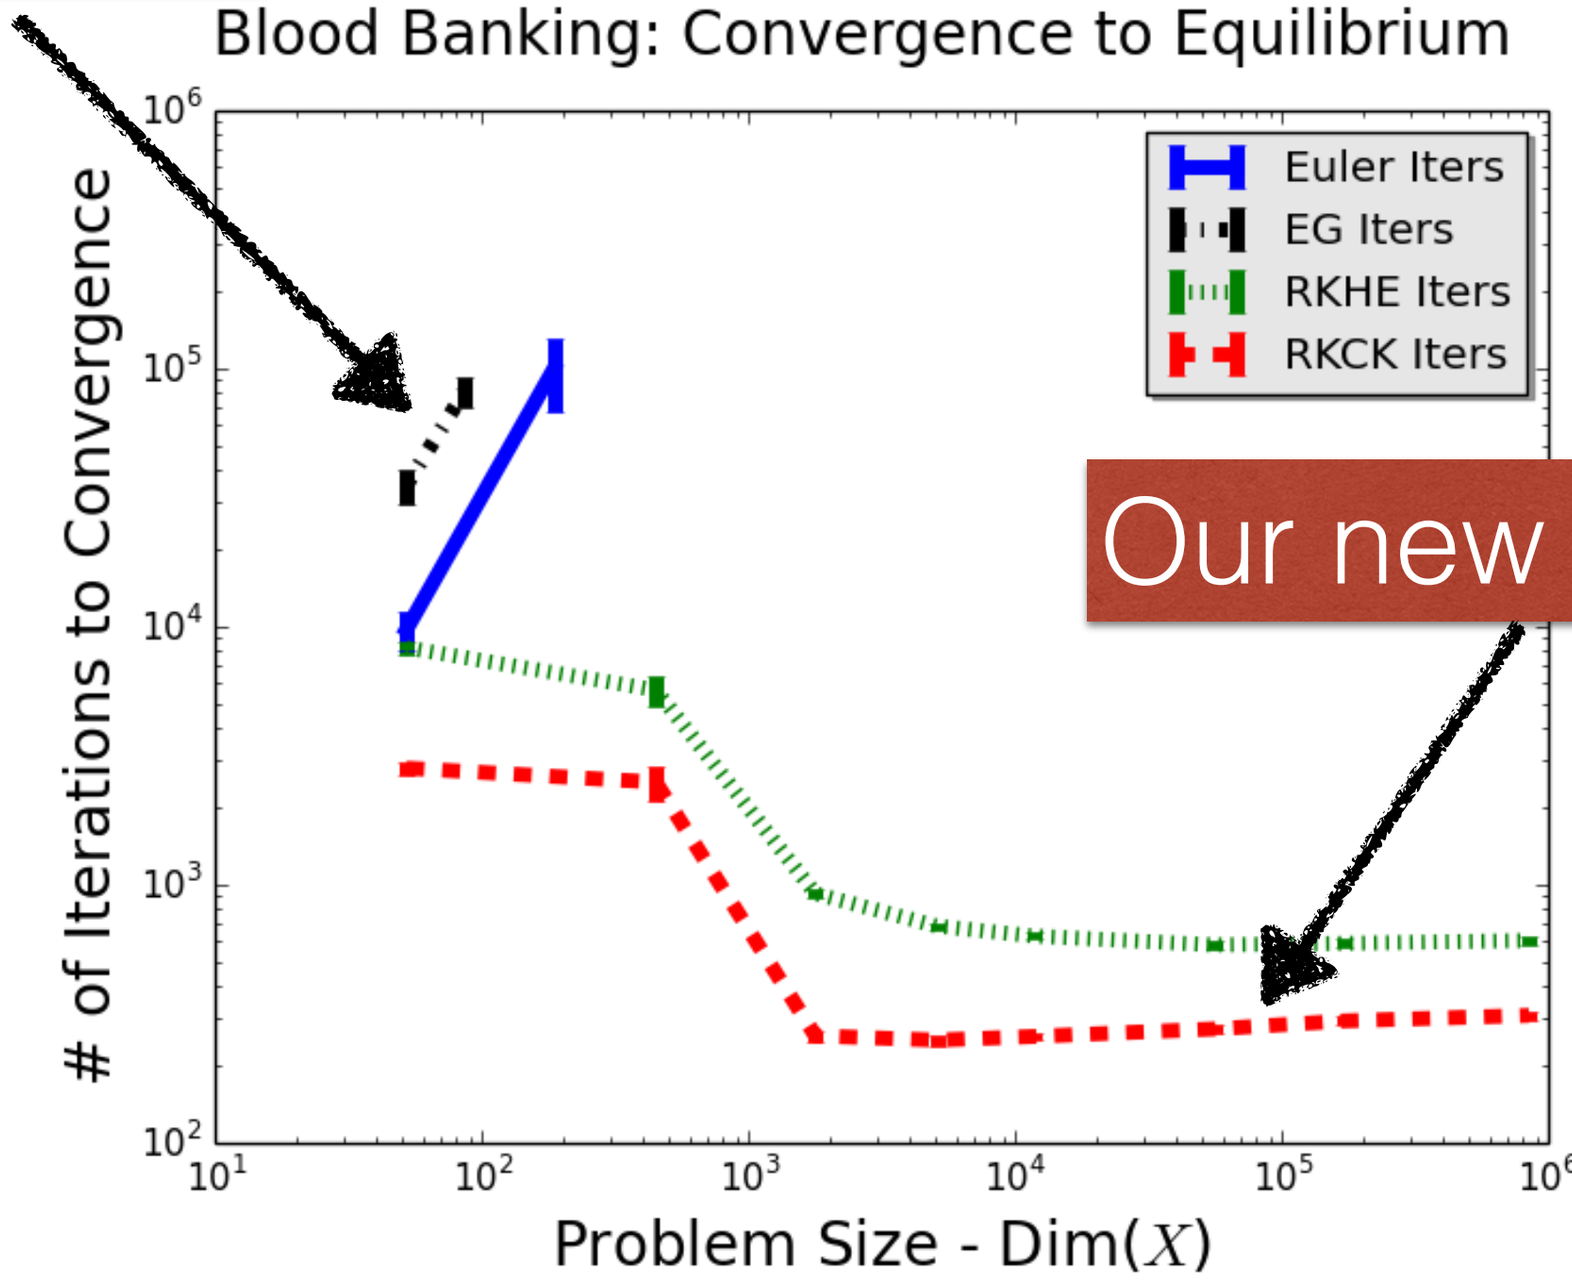

Our new algorithm

Experimental Results

(Gemp and Mahadevan, AAAI Workshop on Computational Substainability, 2015)

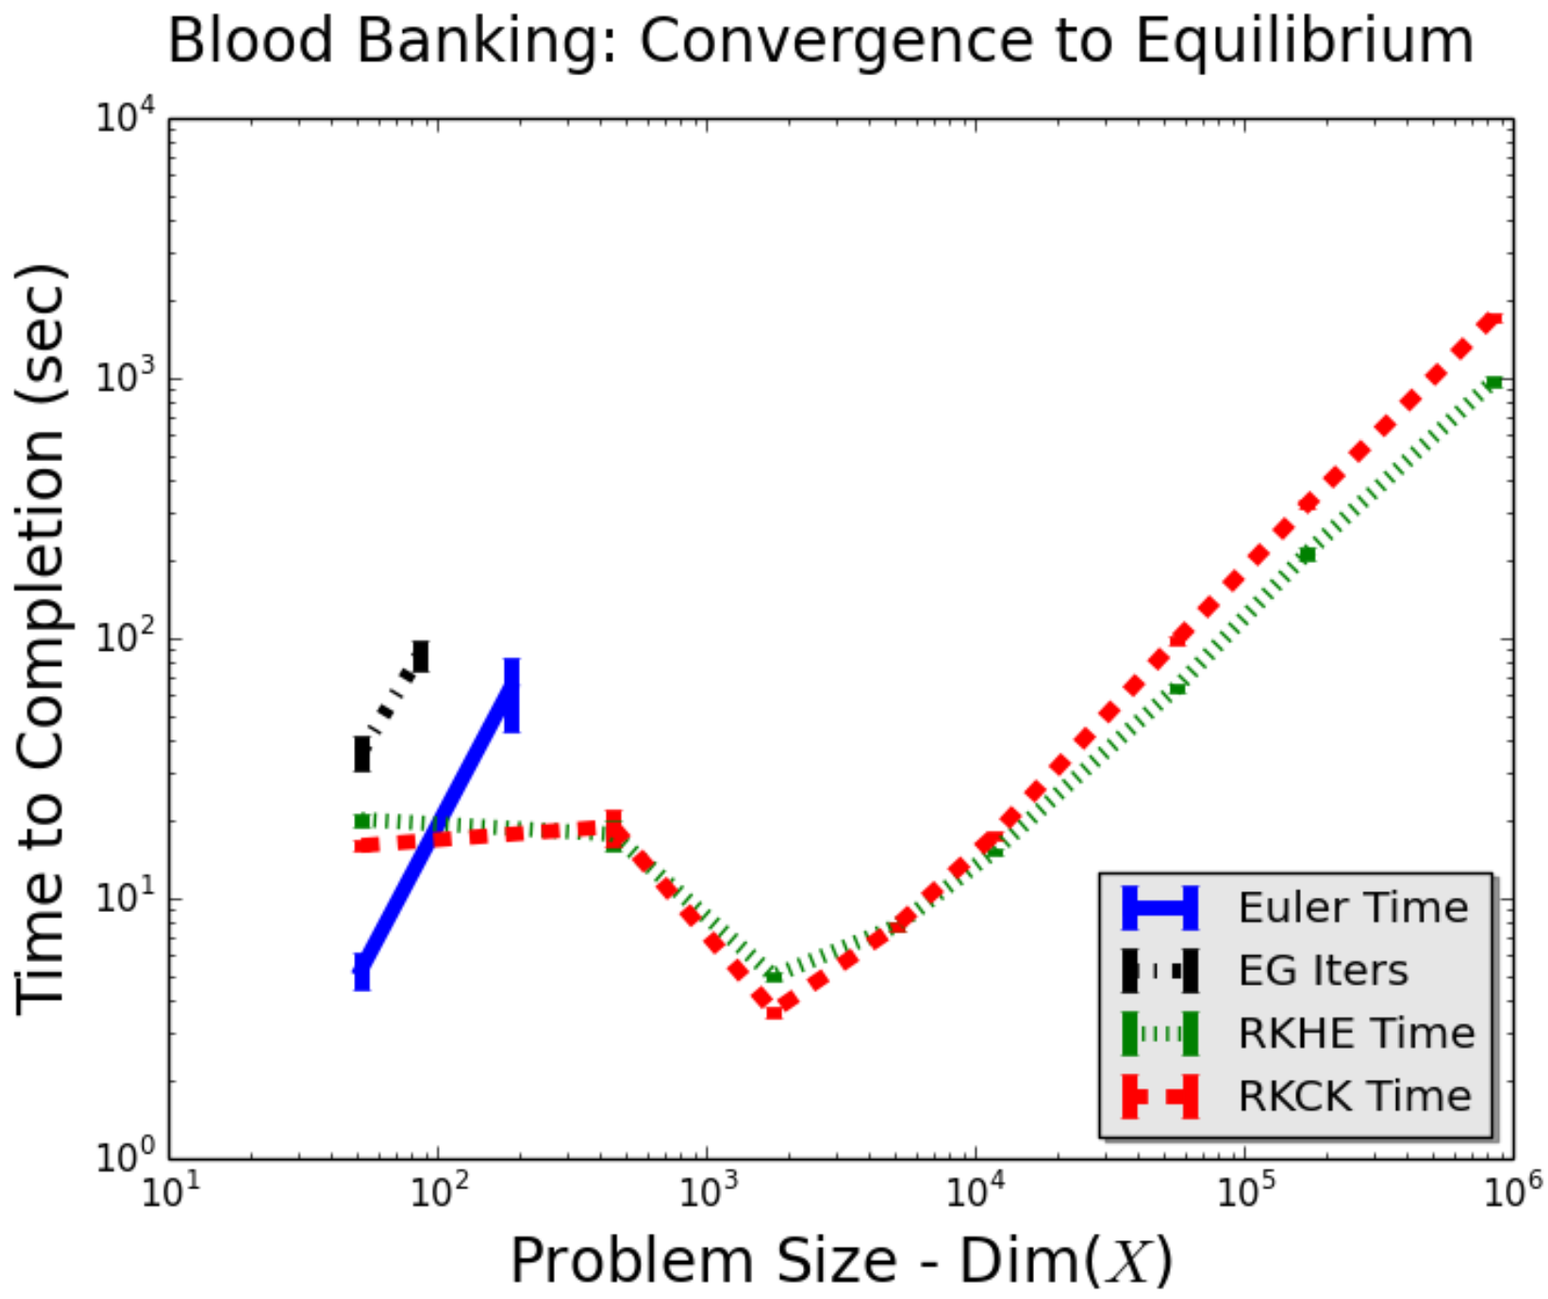

Sustainable Supply Chain

Nagurney et al.

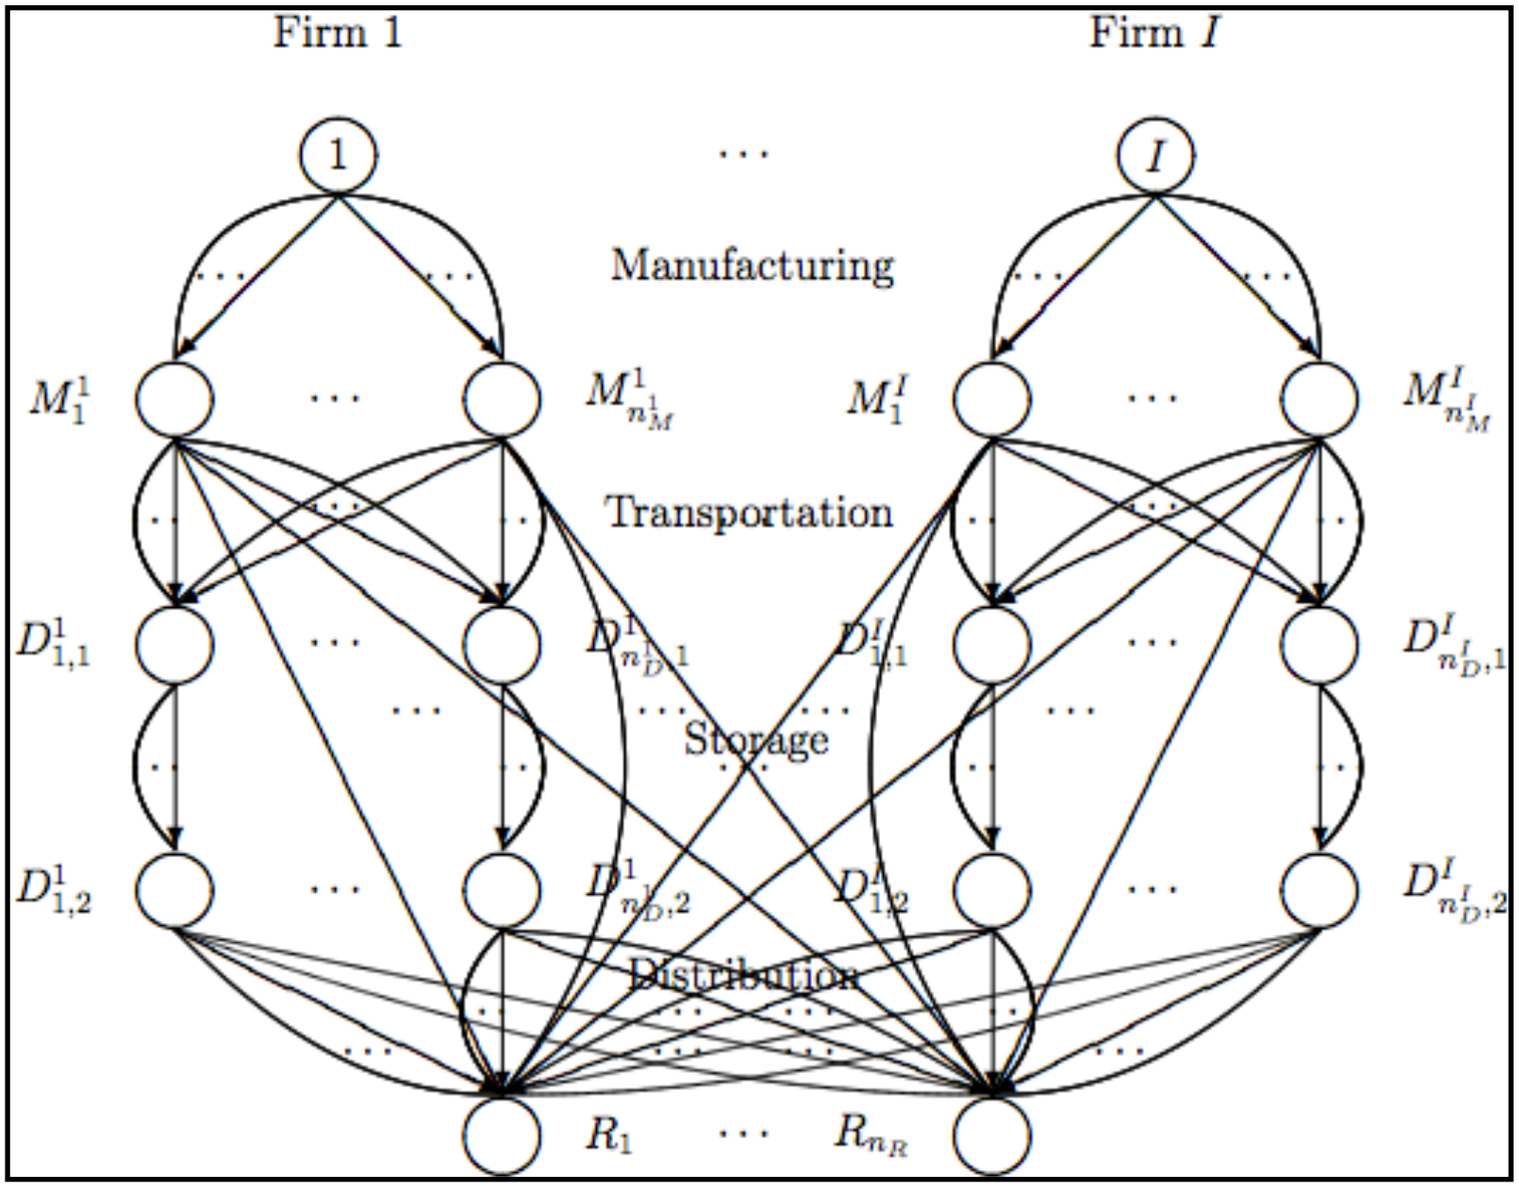

Results on Sustainable Supply Chain VI Problem

Extragradient

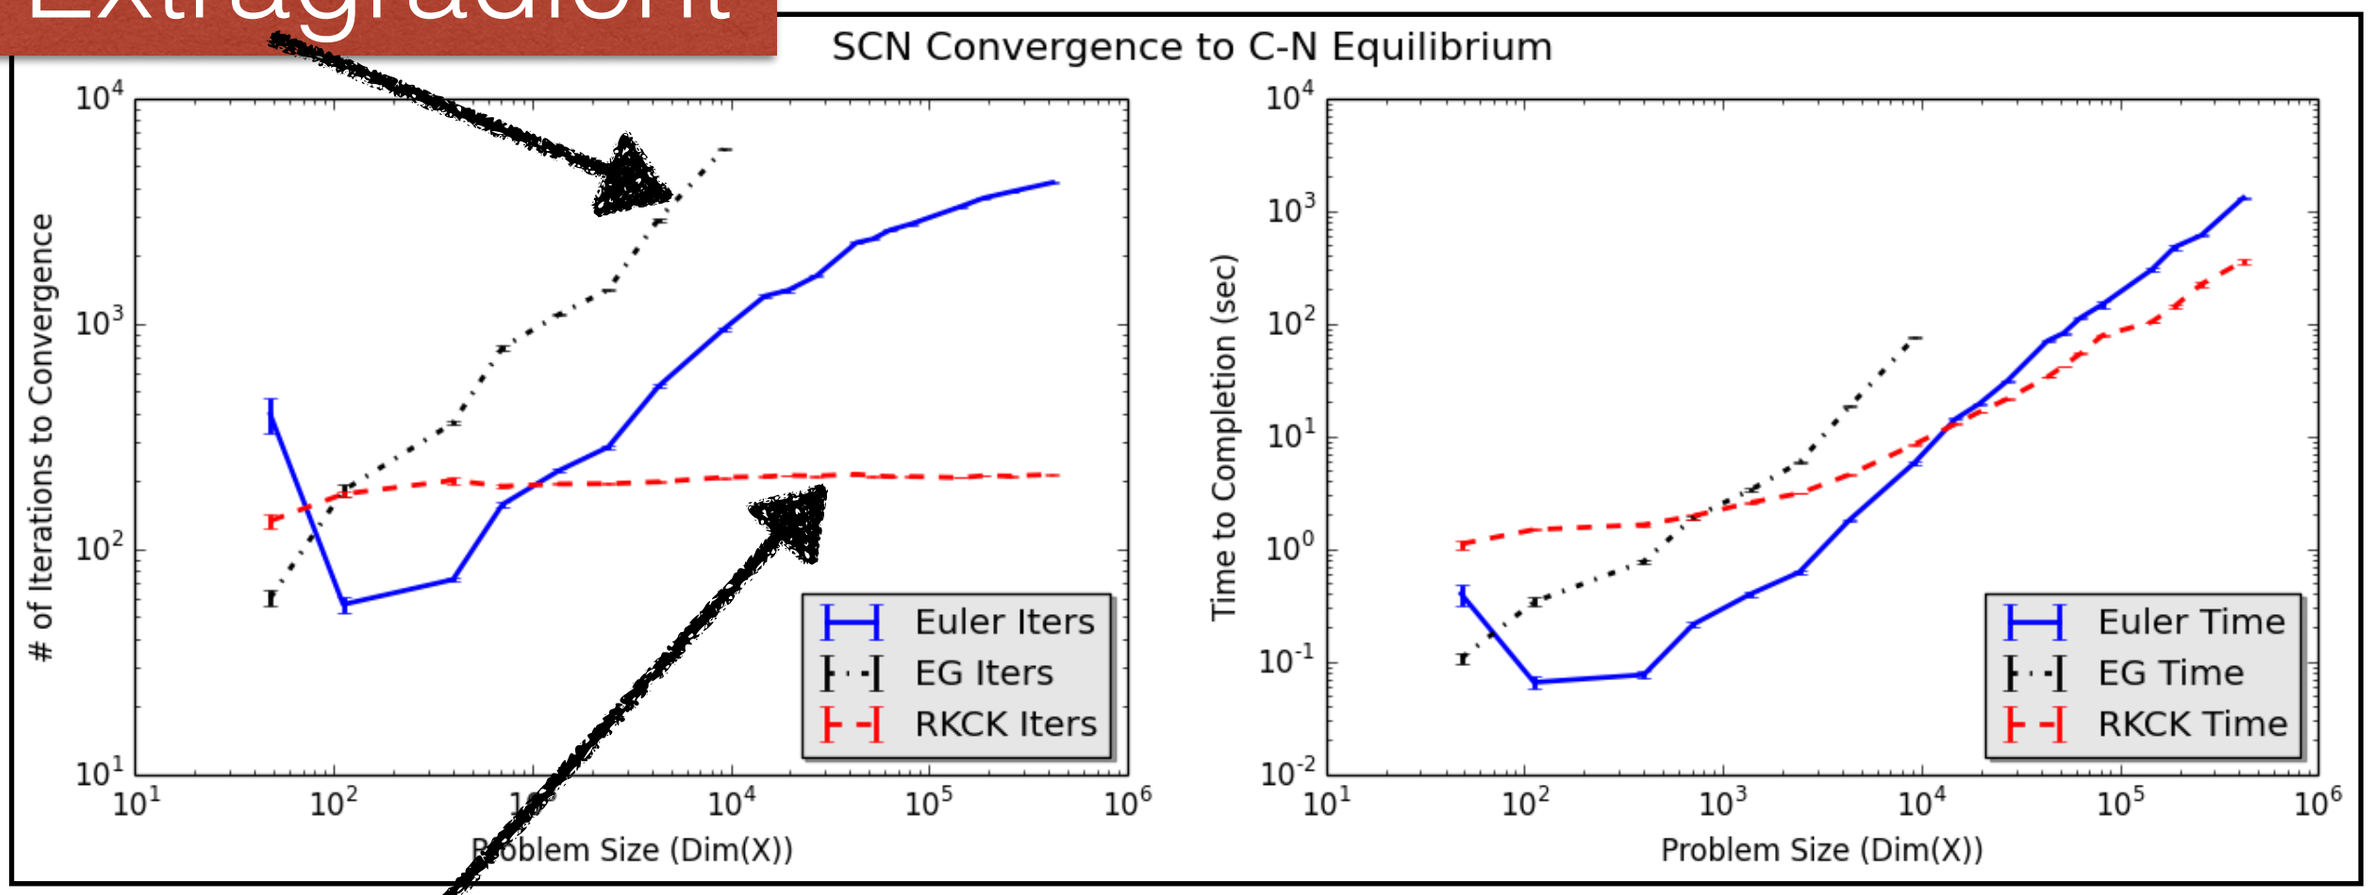

Our new algorithm

Proximal Reinforcement Learning in Primal-Dual Spaces

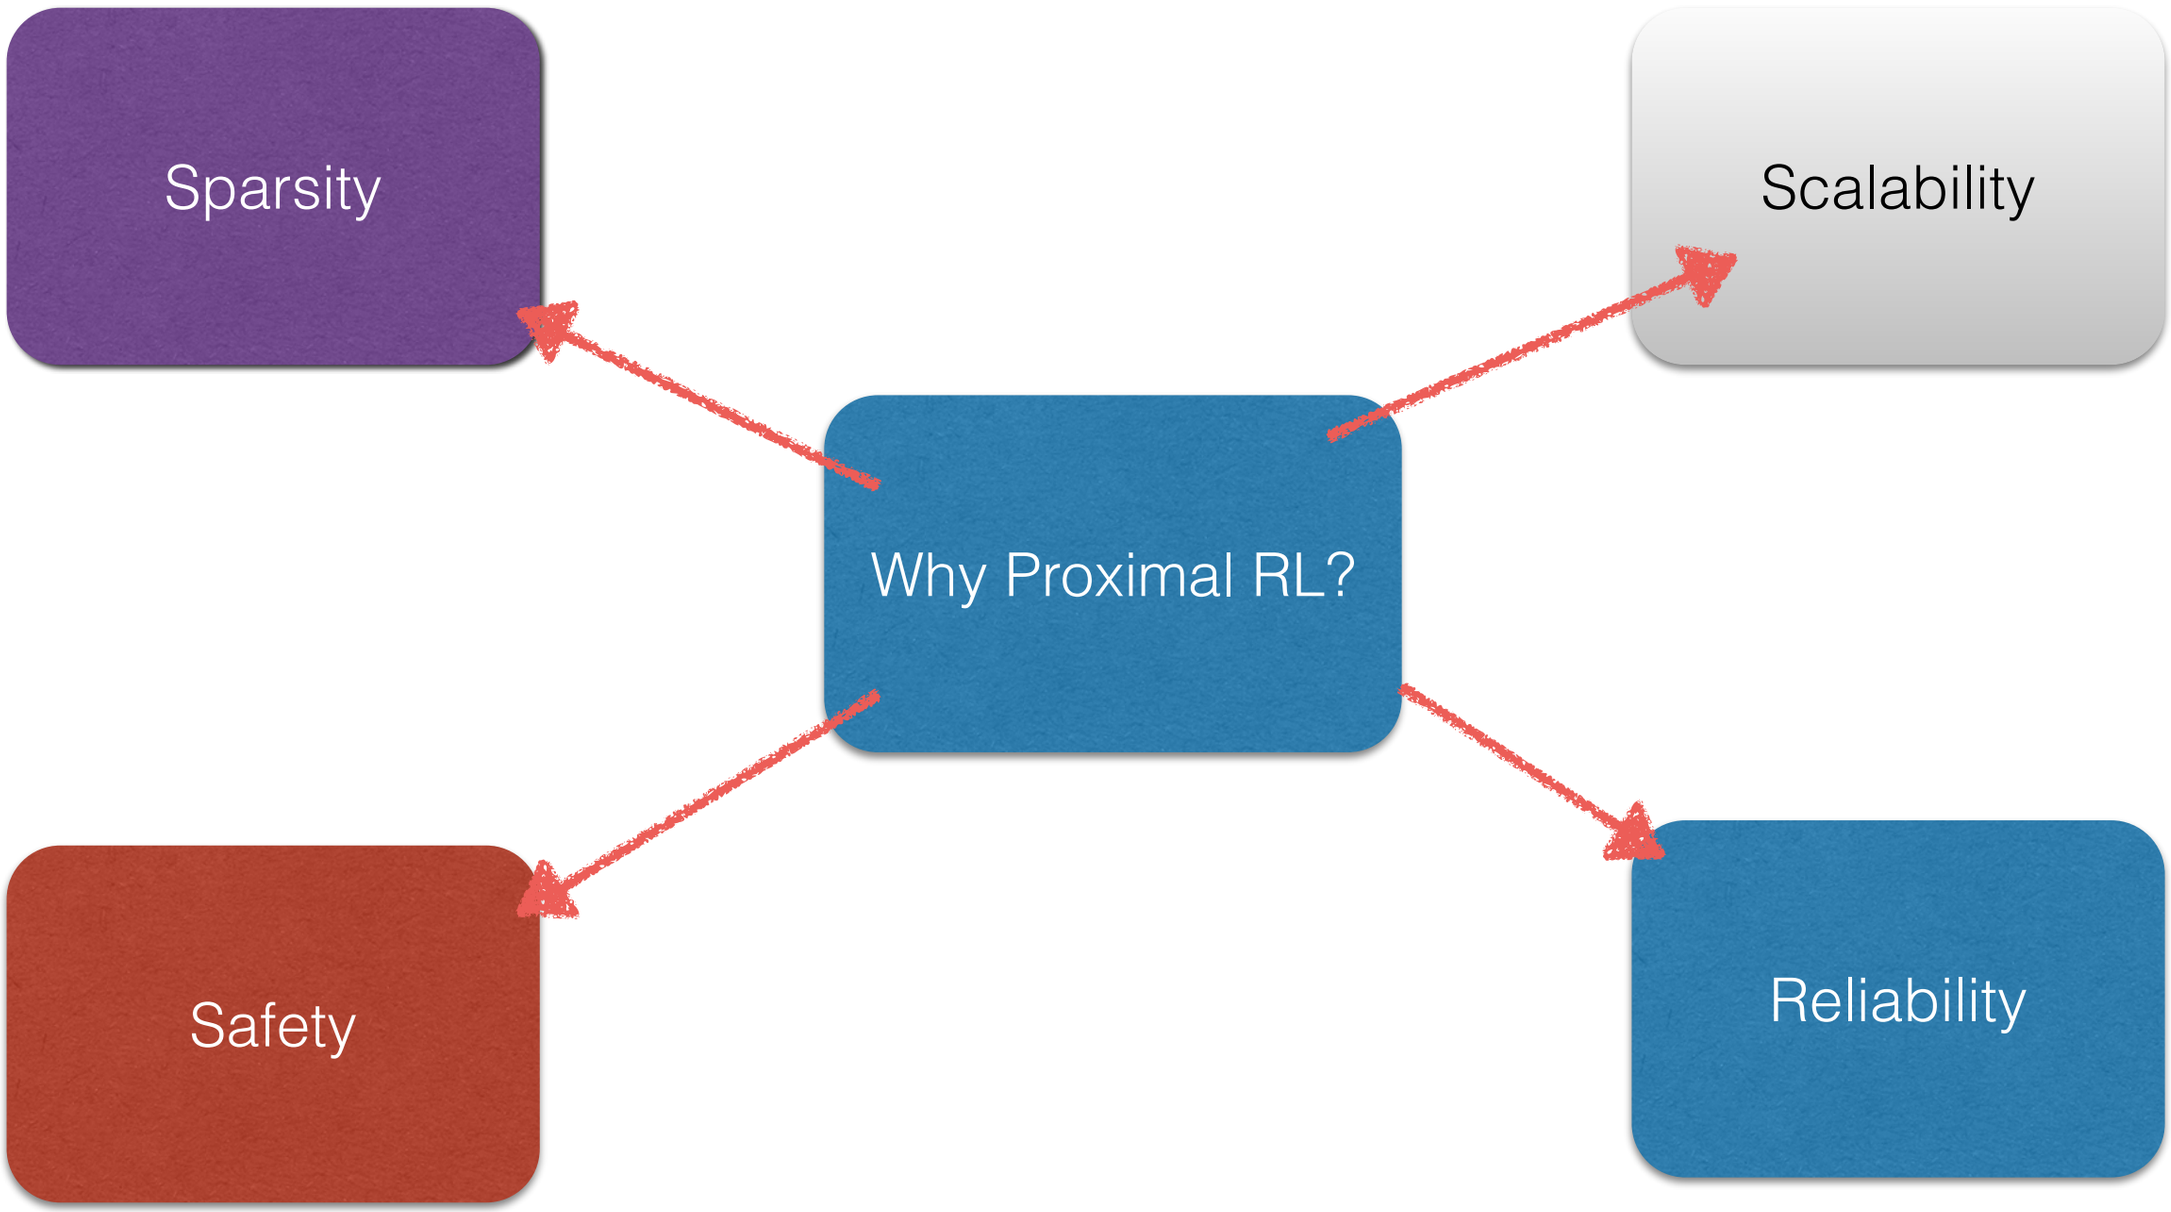

20-Dimensional Robot Arm

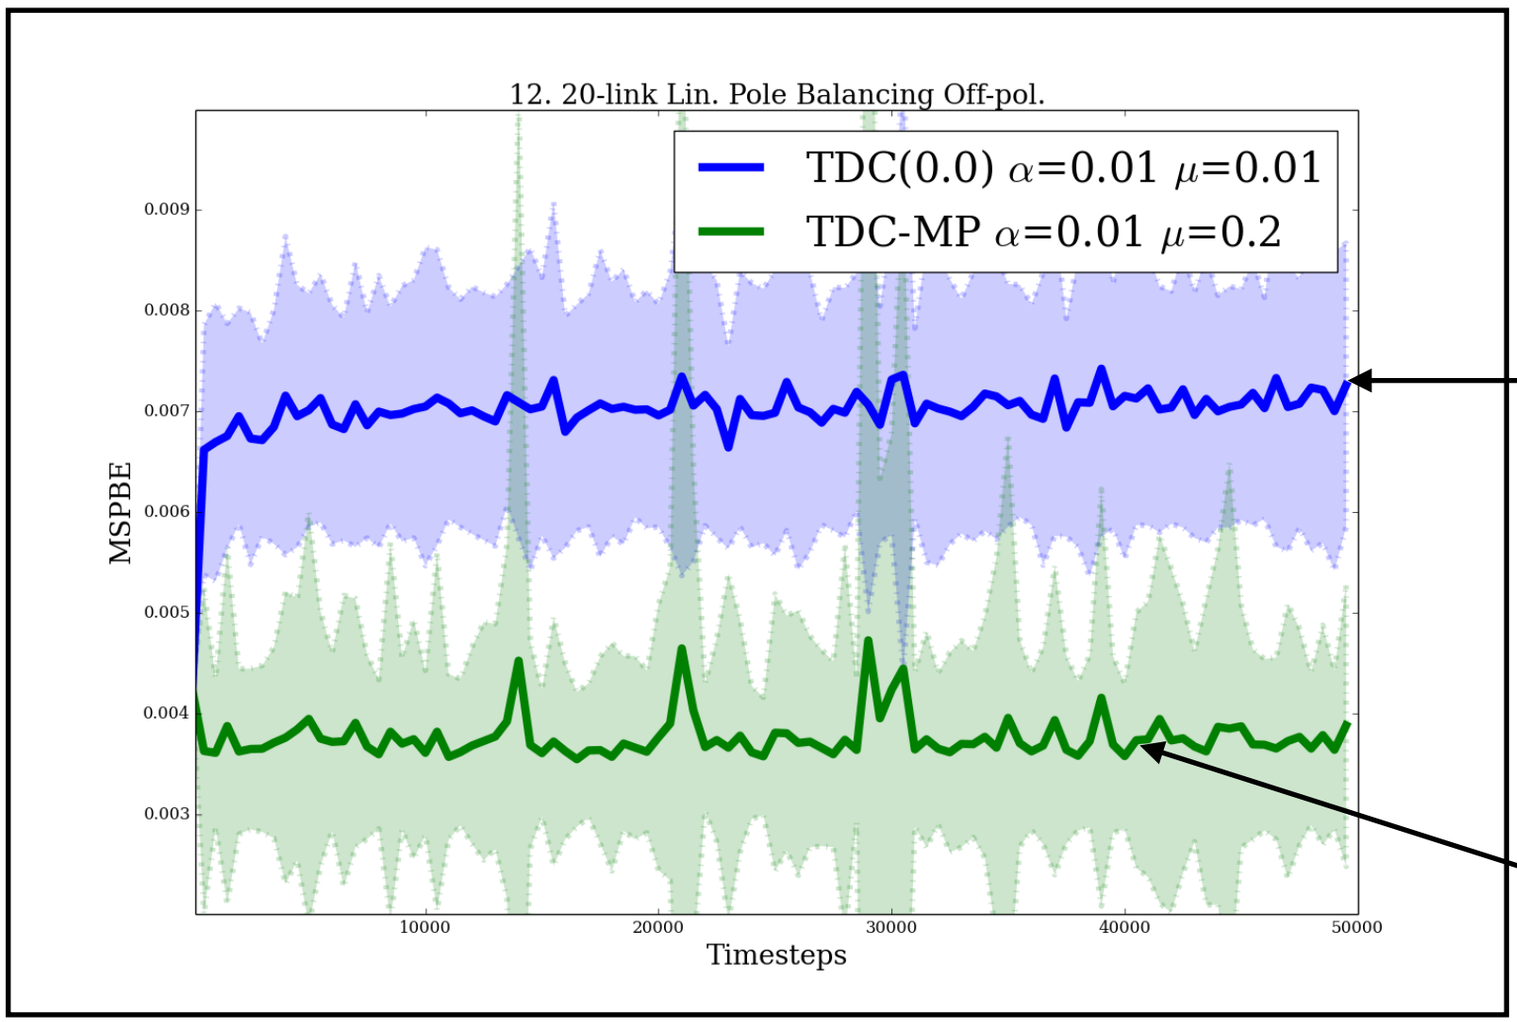

Sutton et al., 2009

Our new method

Safe Robot Learning

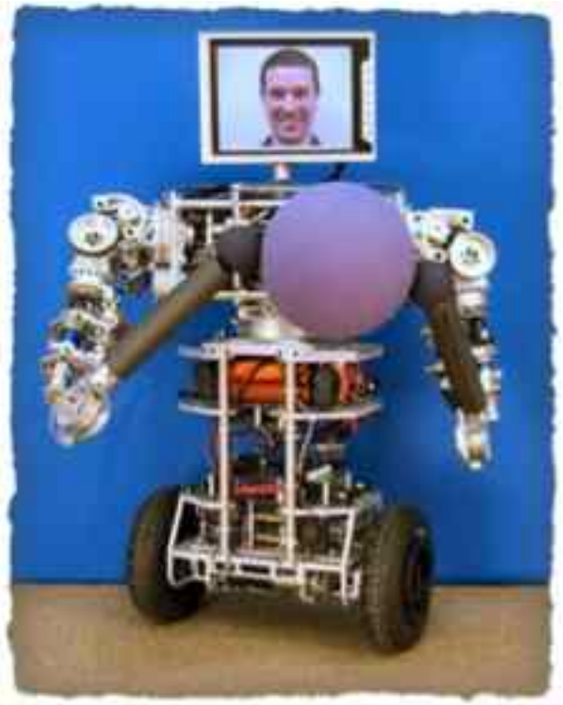

UBot, Laboratory of Perceptual Robotics

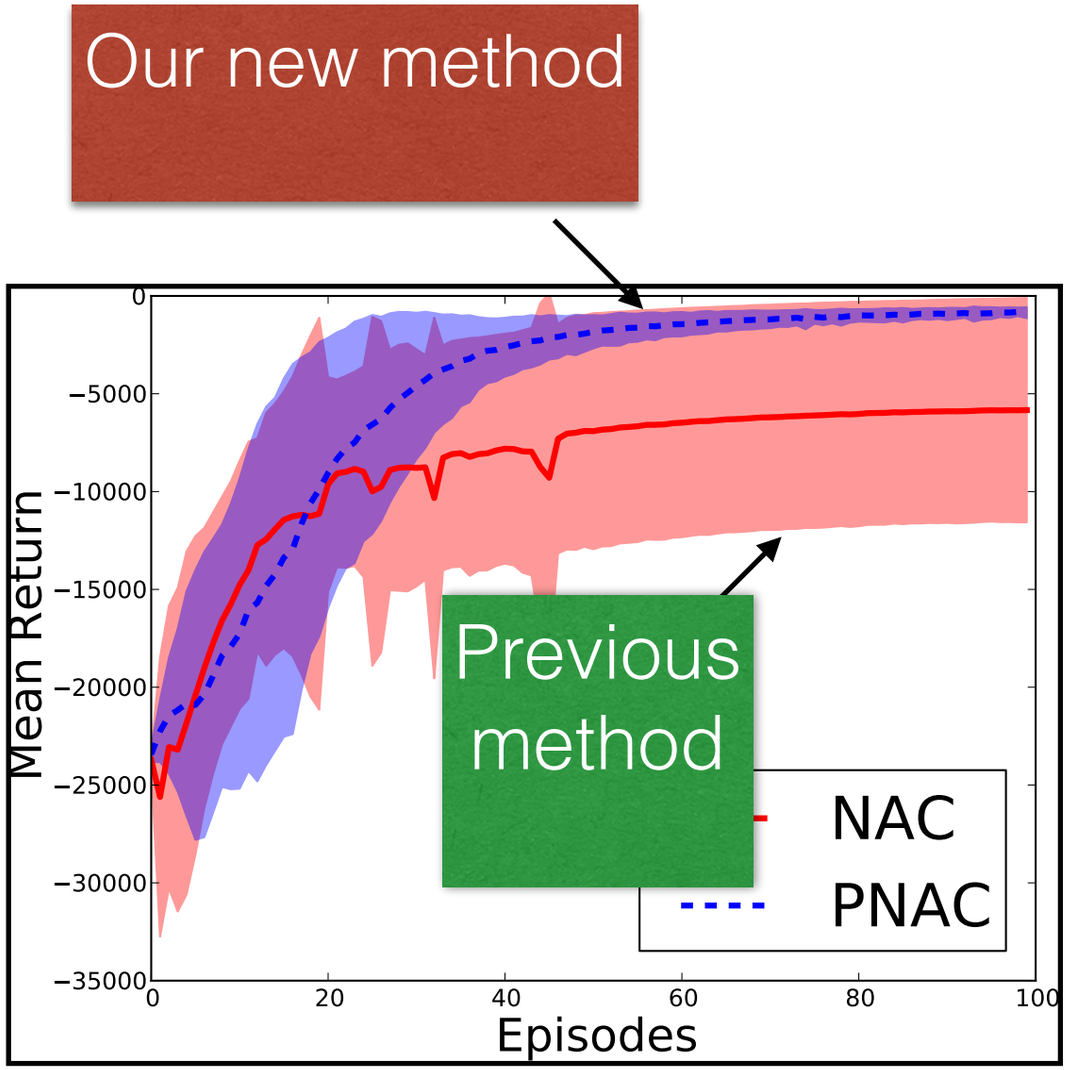

Thomas, Dabney, Mahadevan, Giguere, NIPS 2013

Summary and Challenges

- VIs provide a new foundation for multiagent AI
- Many applications of VI
  - Game theory, complementarity problems, (convex) optimization, equation solving, network equilibrium problems
- Challenges
  - Combine AI methods for modeling complex structured objects with VI theory
  - How to deal with non monotone vector fields (analog of non-convex optimization)?
  - Real-time incremental VI algorithms