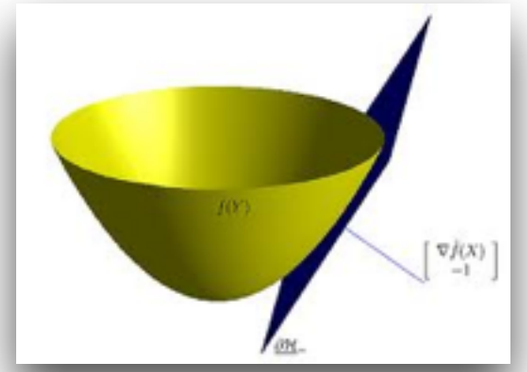



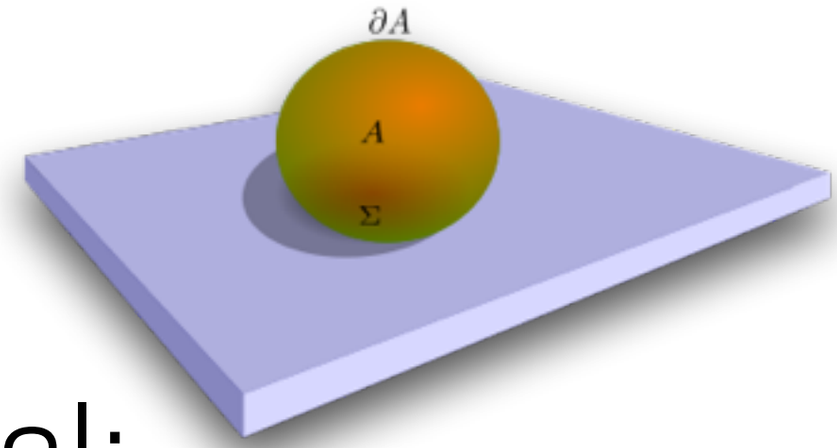

AAAI 2015 Tutorial: From Optimization to Equilibration: A New Foundation for AI in the 21st Century

Sridhar Mahadevan

UMASS CS 50 YEARS
SCHOOL OF COMPUTER SCIENCE

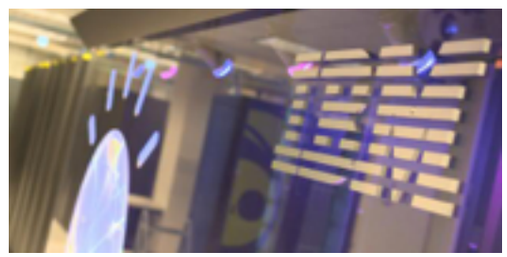

NSF National Science Foundation
WHERE DISCOVERIES BEGIN

NASA

AFRL THE AIR FORCE RESEARCH LABORATORY
LEAD | DISCOVER | DEVELOP | DELIVER
AFOSR

Plan of the Tutorial

- **Overview: 2:00-2:10 p.m.**
- Part I: Motivation: 2:10-2:30 p.m.
- Part II: Mathematical Background: 2:30-3:30 p.m.
- COFFEE BREAK: 3:30-4:00 p.m.
- Part III: Algorithms: 4:00-5:00 p.m.
- Part IV: Applications and Future Work: 5:00-6:00 p.m.

Acknowledgments

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

- Boyd: Lecture notes on optimization

- PhD students:

  - Ian Gemp: 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

Tutorial Web Page

- Slides are accessible online at

http://people.cs.umass.edu/~mahadeva/AAAI_2015_Tutorial/Welcome.html

- This page contains pointers to reference material
  - Books on VI
  - Selected papers

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

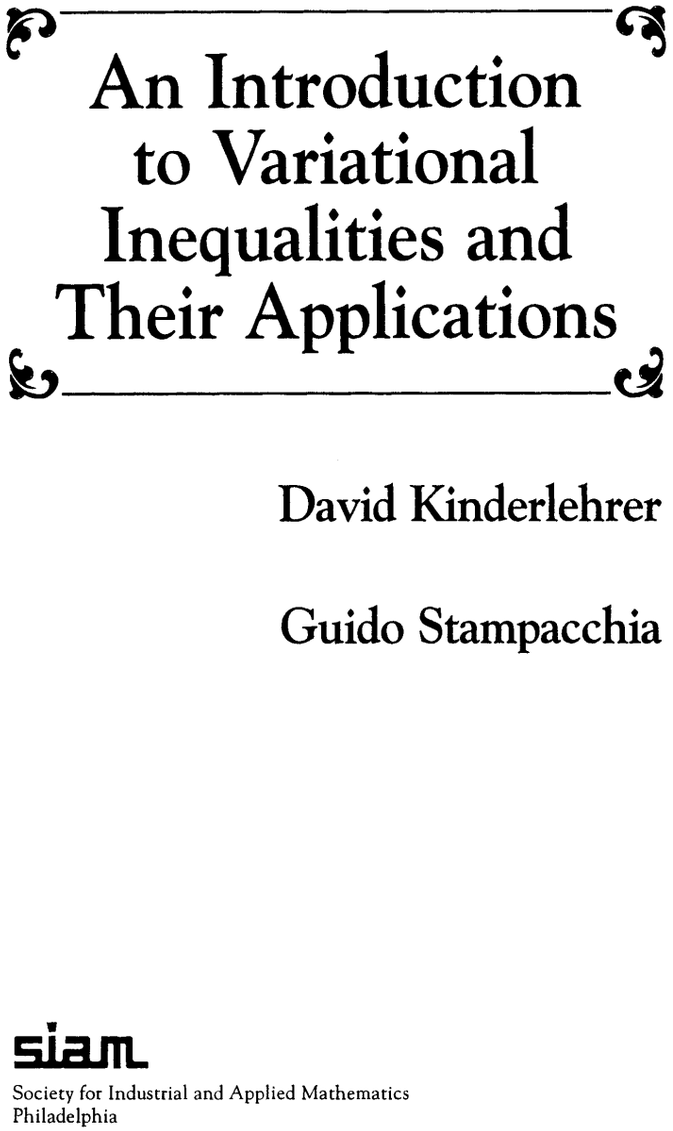

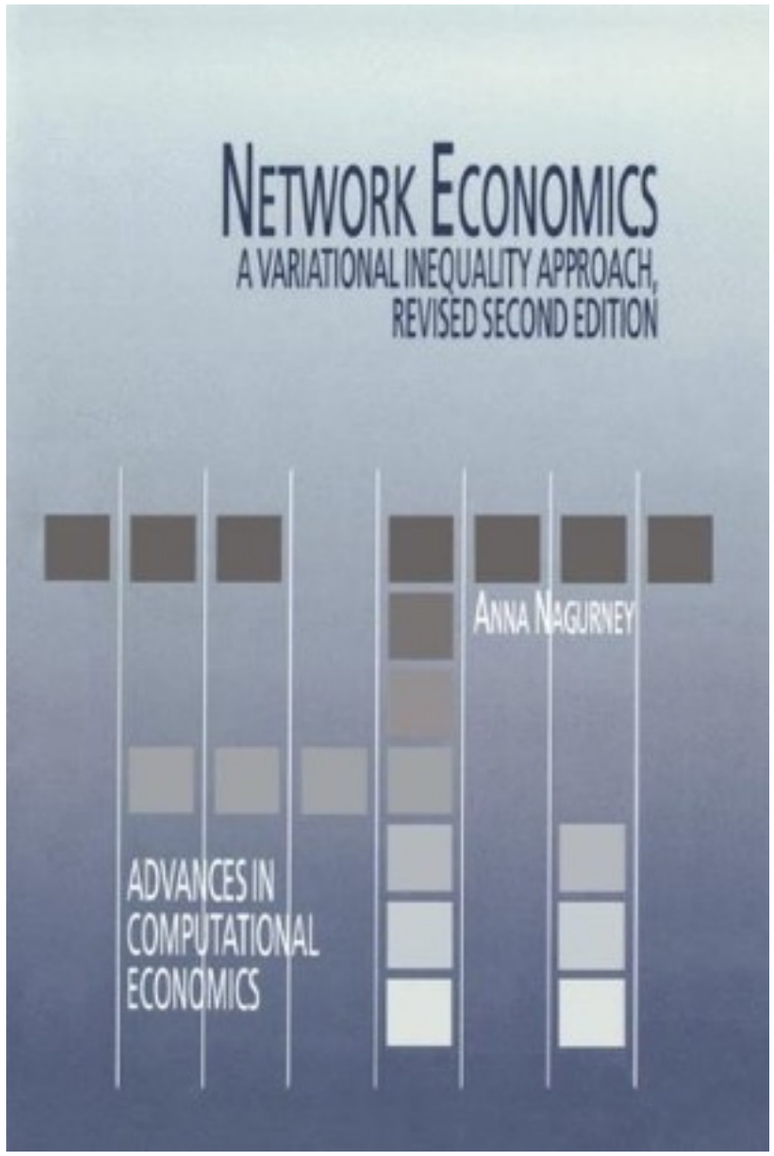

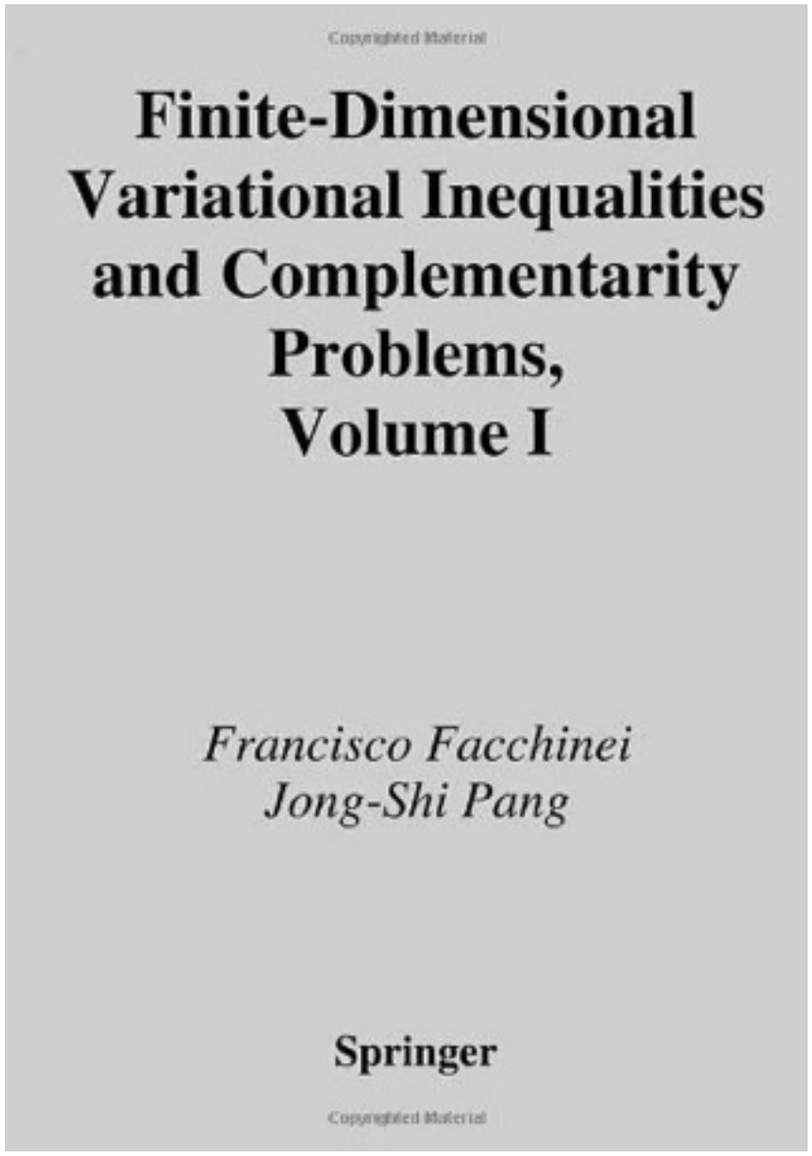

References

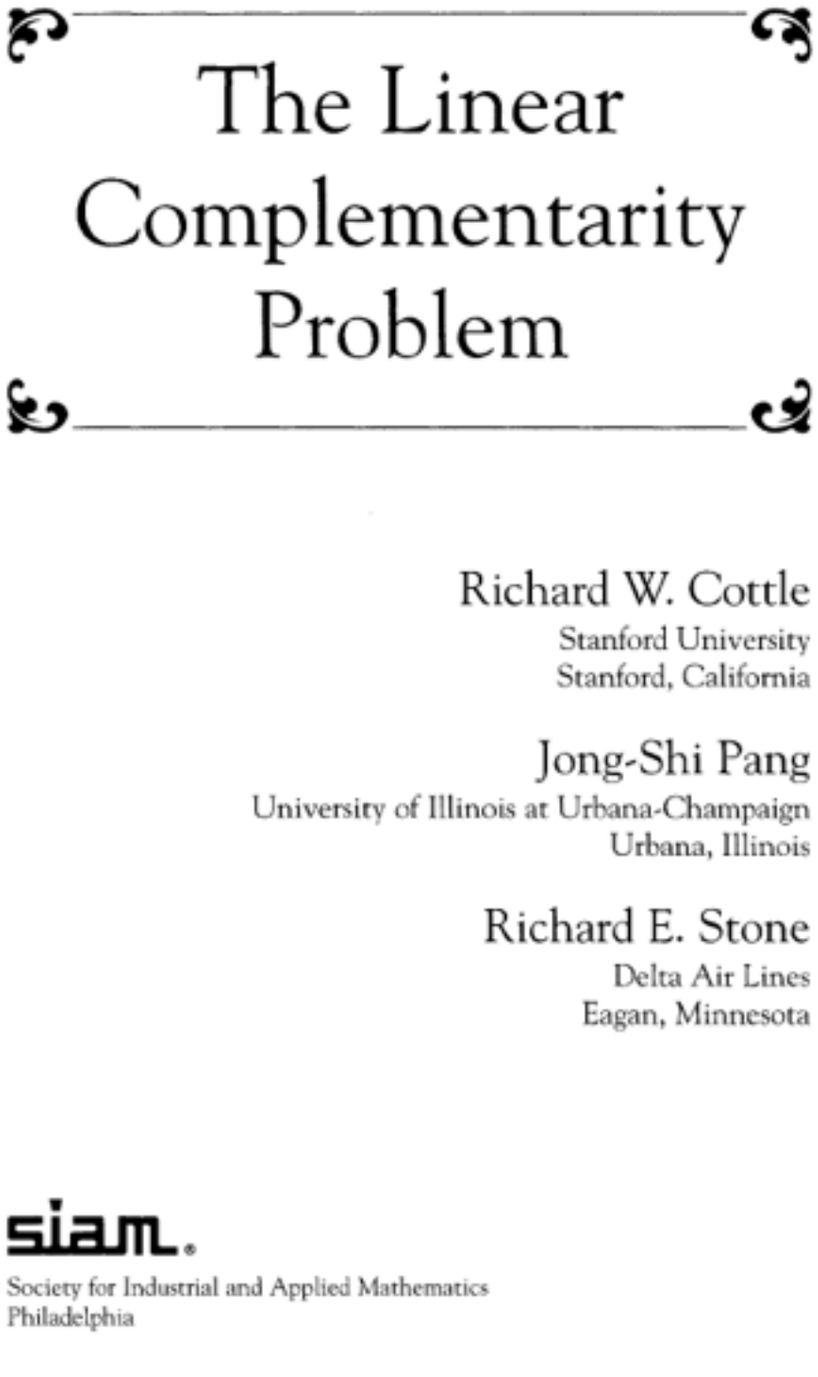
The Linear
Complementarity
Problem
Richard W. Cottle
Stanford University
Stanford, California
Jong-Shi Pang
University of Illinois at Urbana-Champaign
Urbana, Illinois
Richard E. Stone
Delta Air Lines
Eagan, Minnesota
siam.
Society for Industrial and Applied Mathematics
Philadelphia

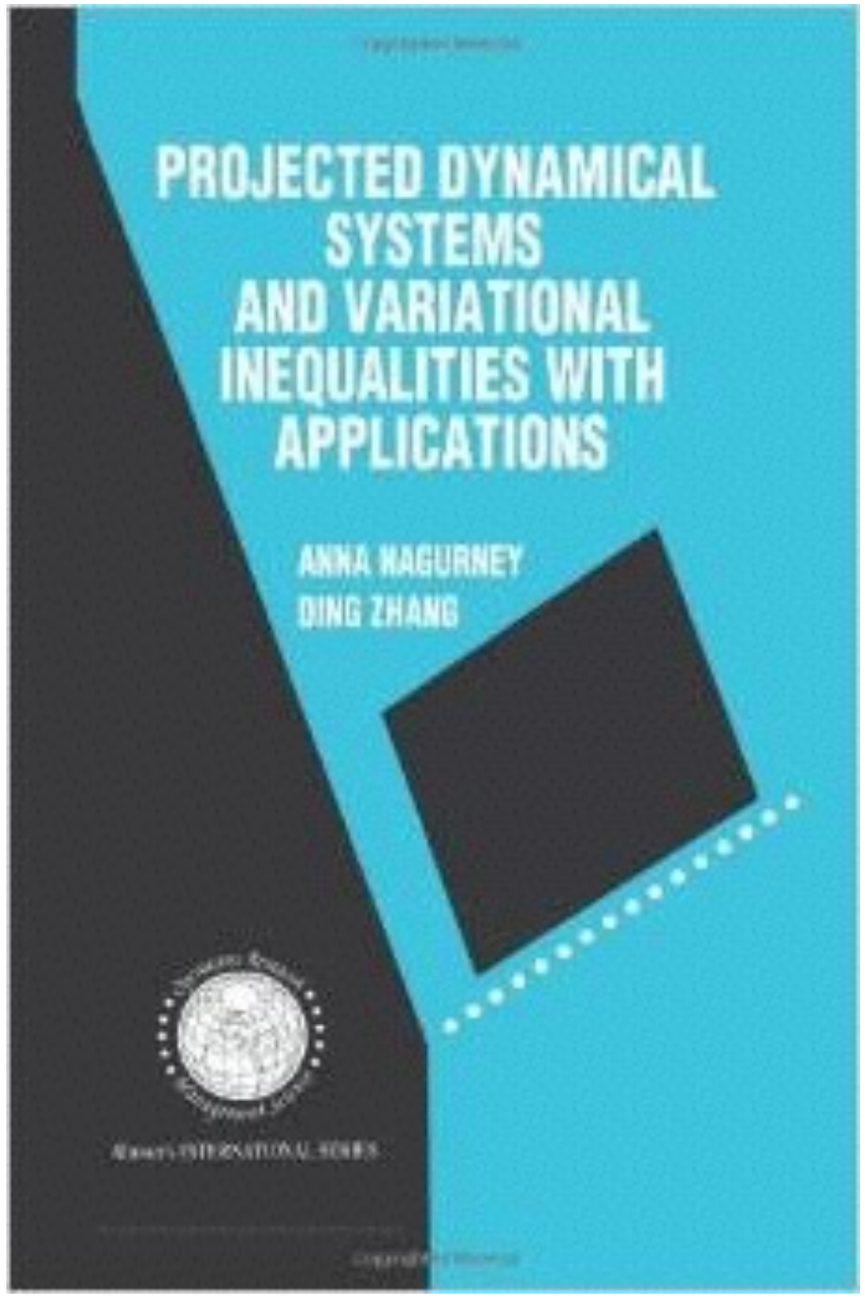
PROJECTED DYNAMICAL
SYSTEMS
AND VARIATIONAL
INEQUALITIES WITH
APPLICATIONS
ANNA NAGURNEY
DING ZHANG

Optimization

Game Theory

Engineering

Linear Programming

$$\min_x c^T x$$

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

AI

Business

Equilibria in Economics

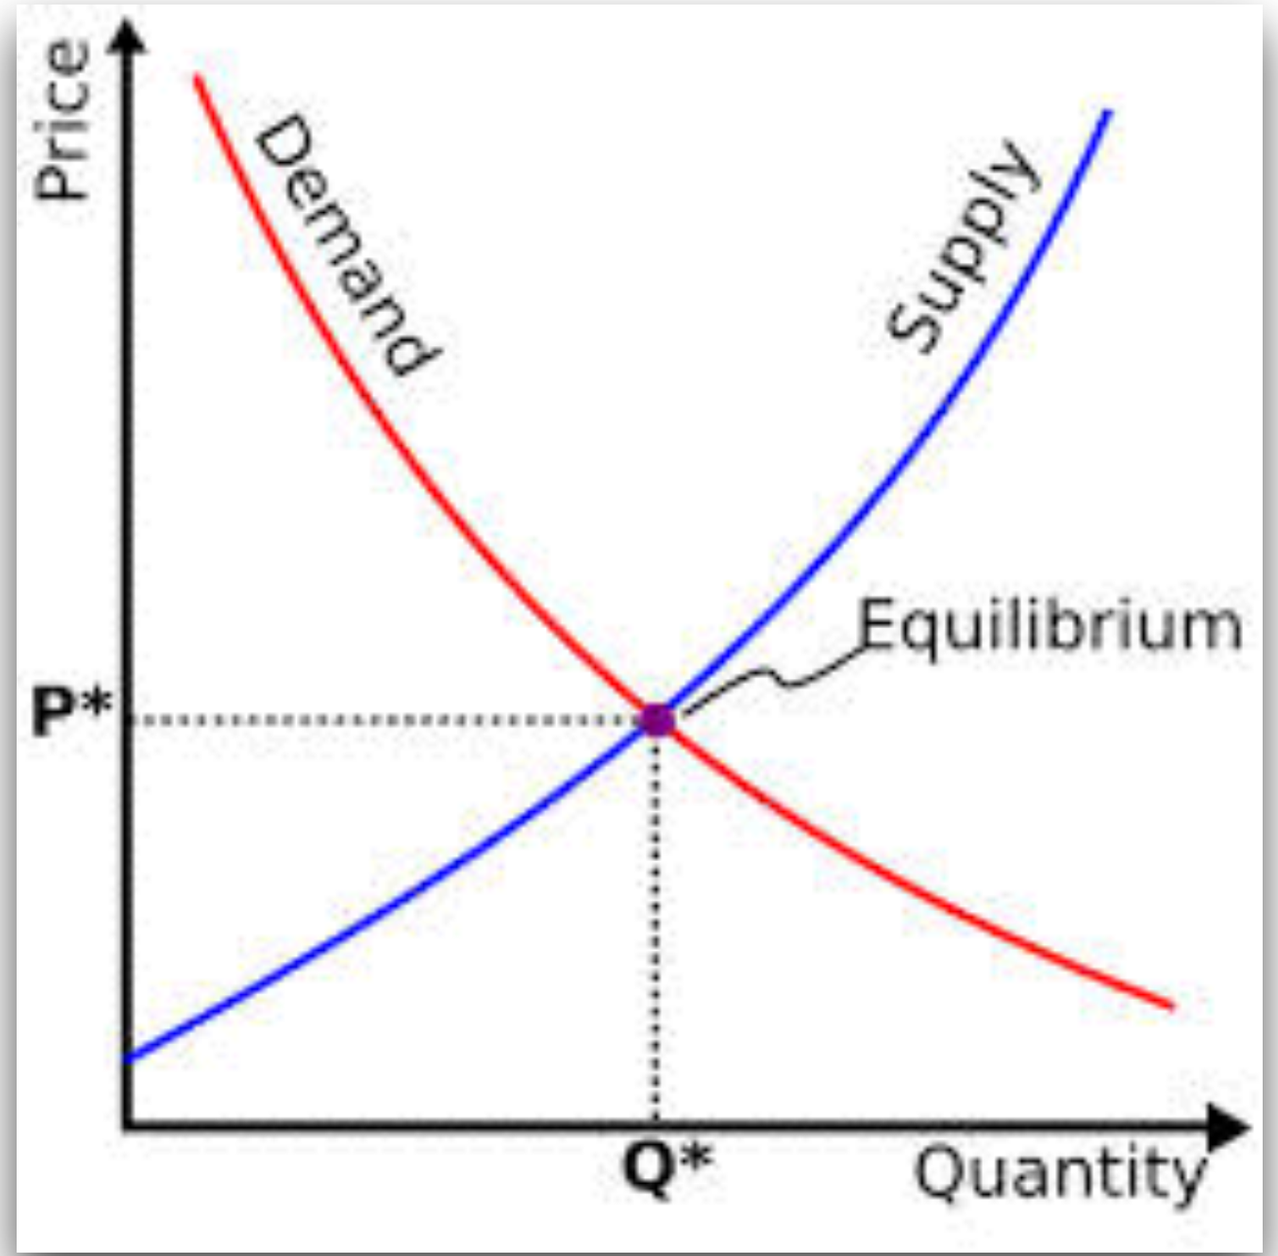

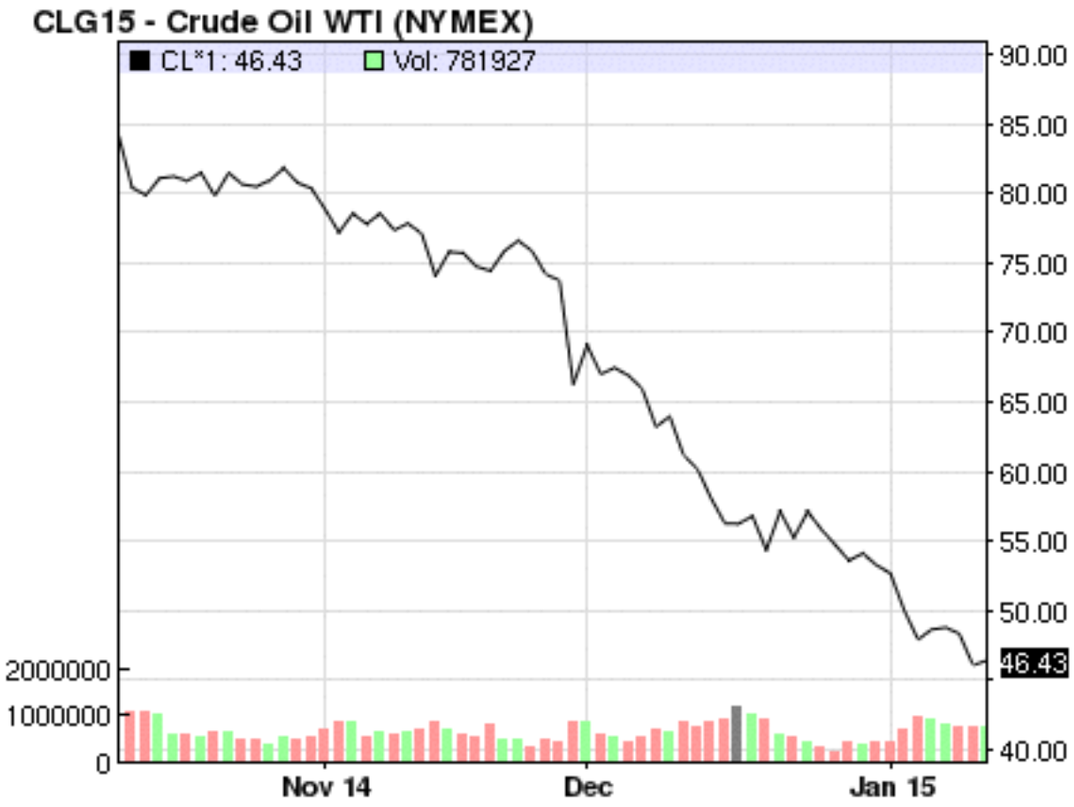

Price of Crude Oil

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.
- Tutorial presents a new foundation for AI: **equilibration**
  - Generalizes optimization, covers many other problems

Optimization

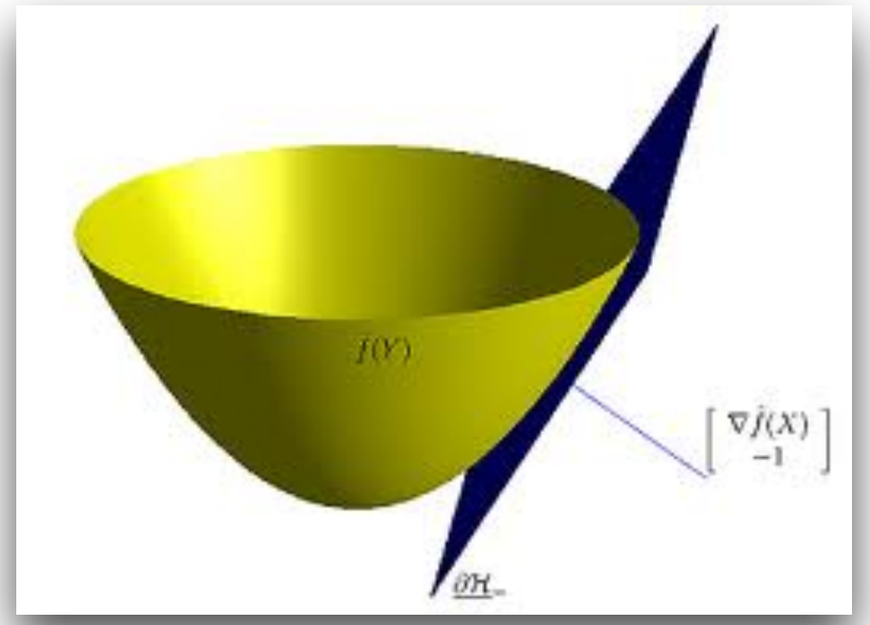

min f(x)
x in feasible set K

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

Equilibration

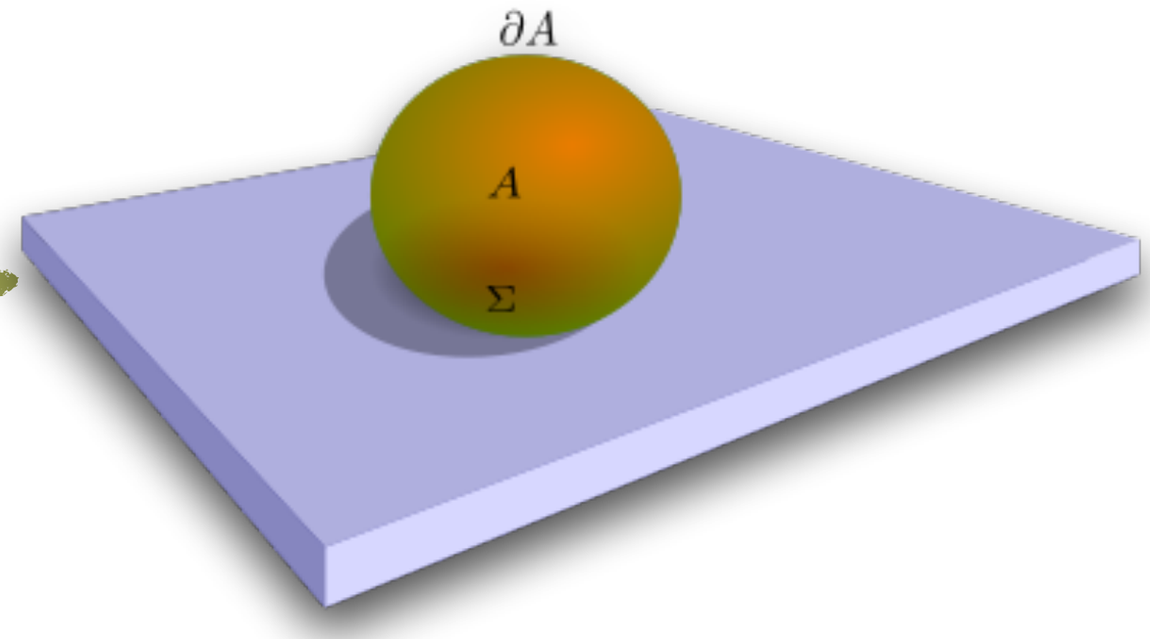

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

(Stampacchia, 1960s)

Variational Inequalities

- Most general problem formulation that admits a convex structure

- Applies to problems with unknown or multiple competing objective functions

- Generalizes optimization and includes fixed point problems, game theory, nonlinear equations, and network equilibrium problems

Variational Inequality

(Stampacchia, 1960s)

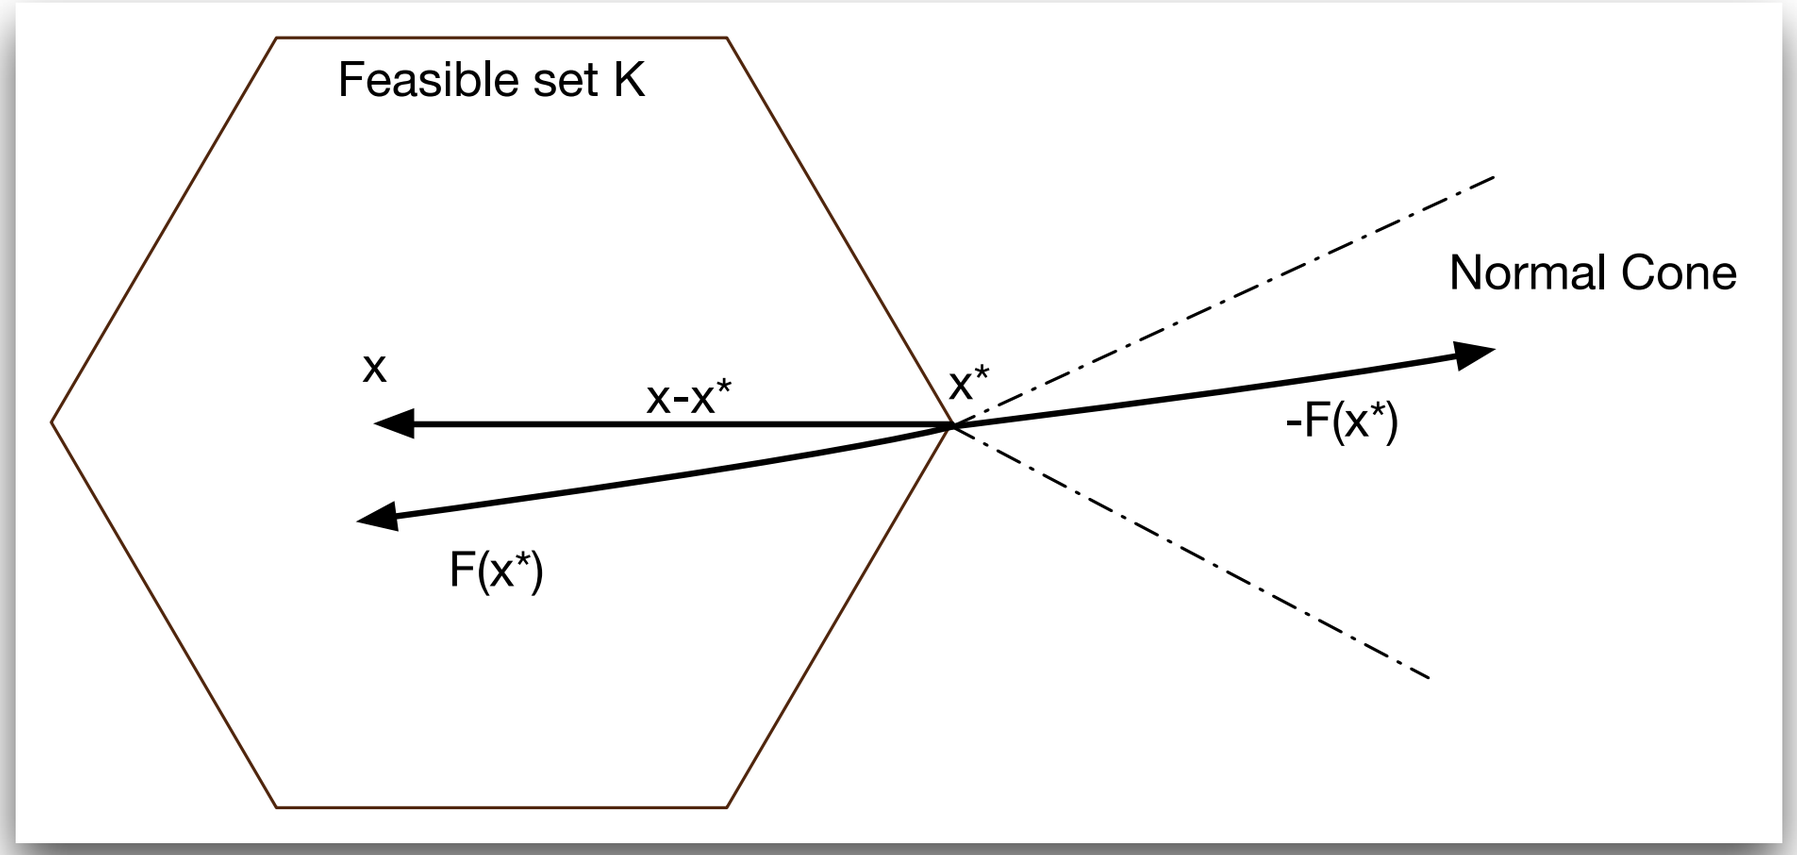

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

Vector field

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

Gradients to Vector Fields

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

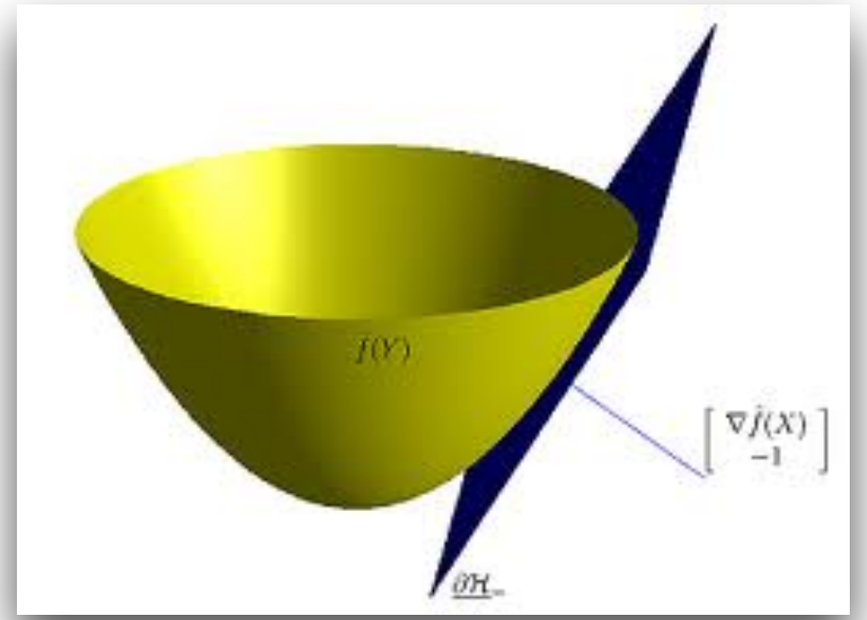

$$w_{t+1} \leftarrow w_t - \alpha_t F(w_t)$$

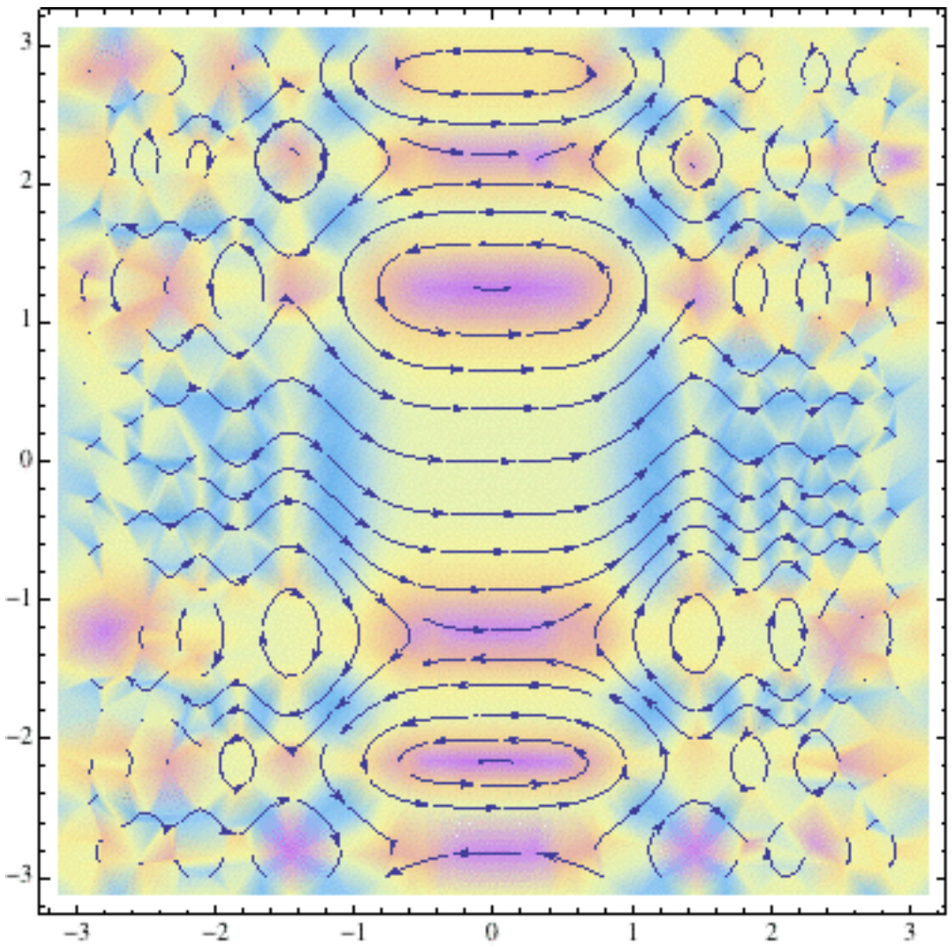

Generalizing Gradient Descent

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



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



Extragradient method for VIs

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

Spatial Price Equilibria

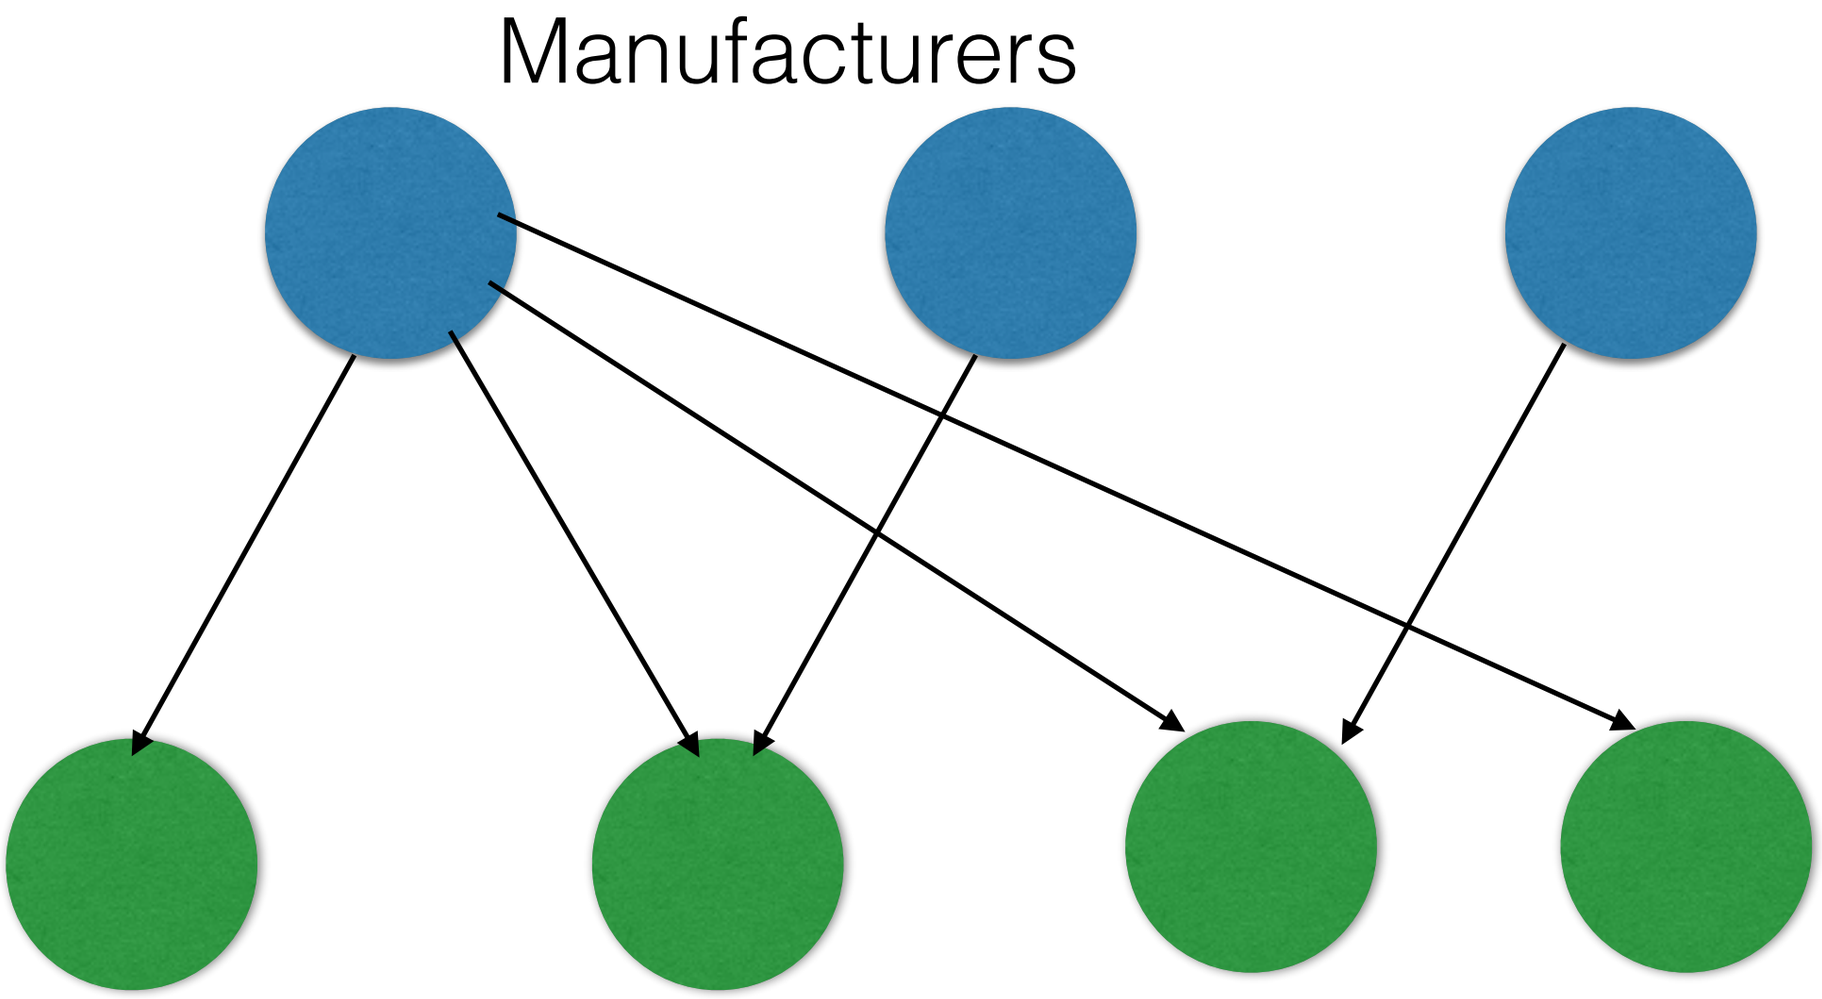

Spatial Price Equilibrium Model

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

Supply level	s_i
Supply price	π_j
Quantity demanded	d_i
Demand price	ρ_j
Cost of transportation	c_{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

Blood Bank System

39,000 donations/day in the US

Often, only a 2 day supply

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



Sustainable Blood Banking

(Nagurney and Masoumi, 2012)

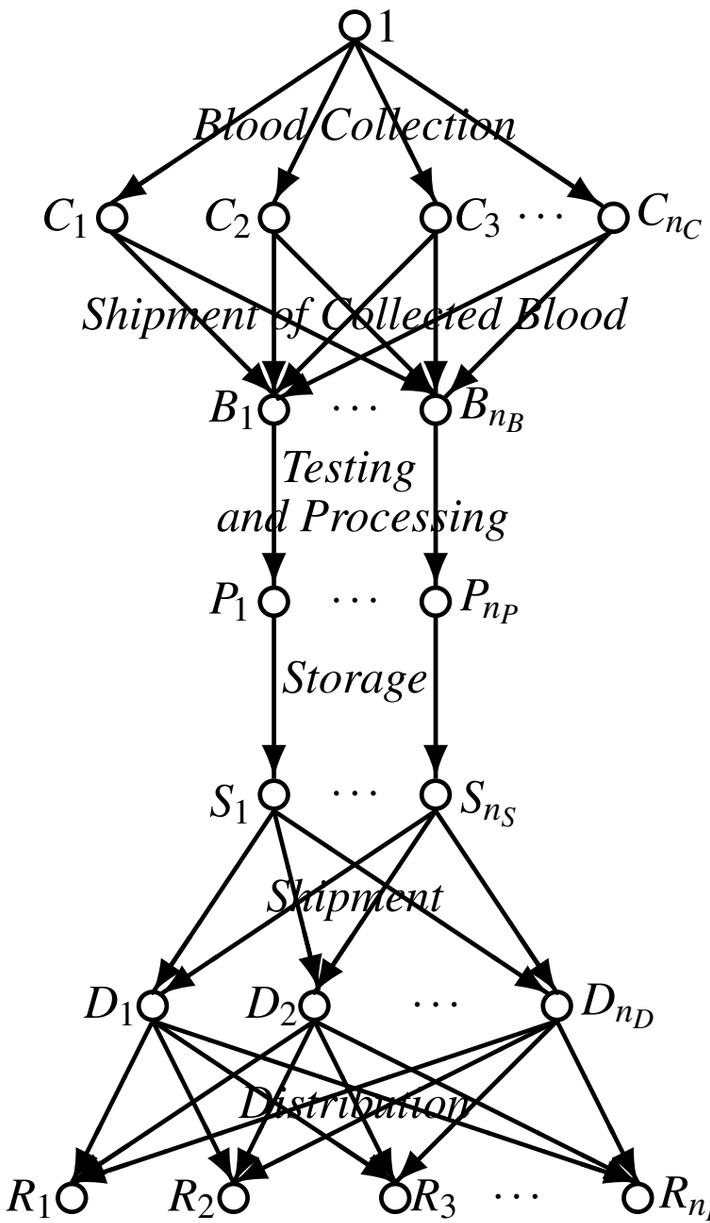

The Organization

The Blood Collection Sites

The Blood Centers

The Component Labs

The Storage Facilities

The Distribution Centers

The Demand Points

Blood Bank Demo

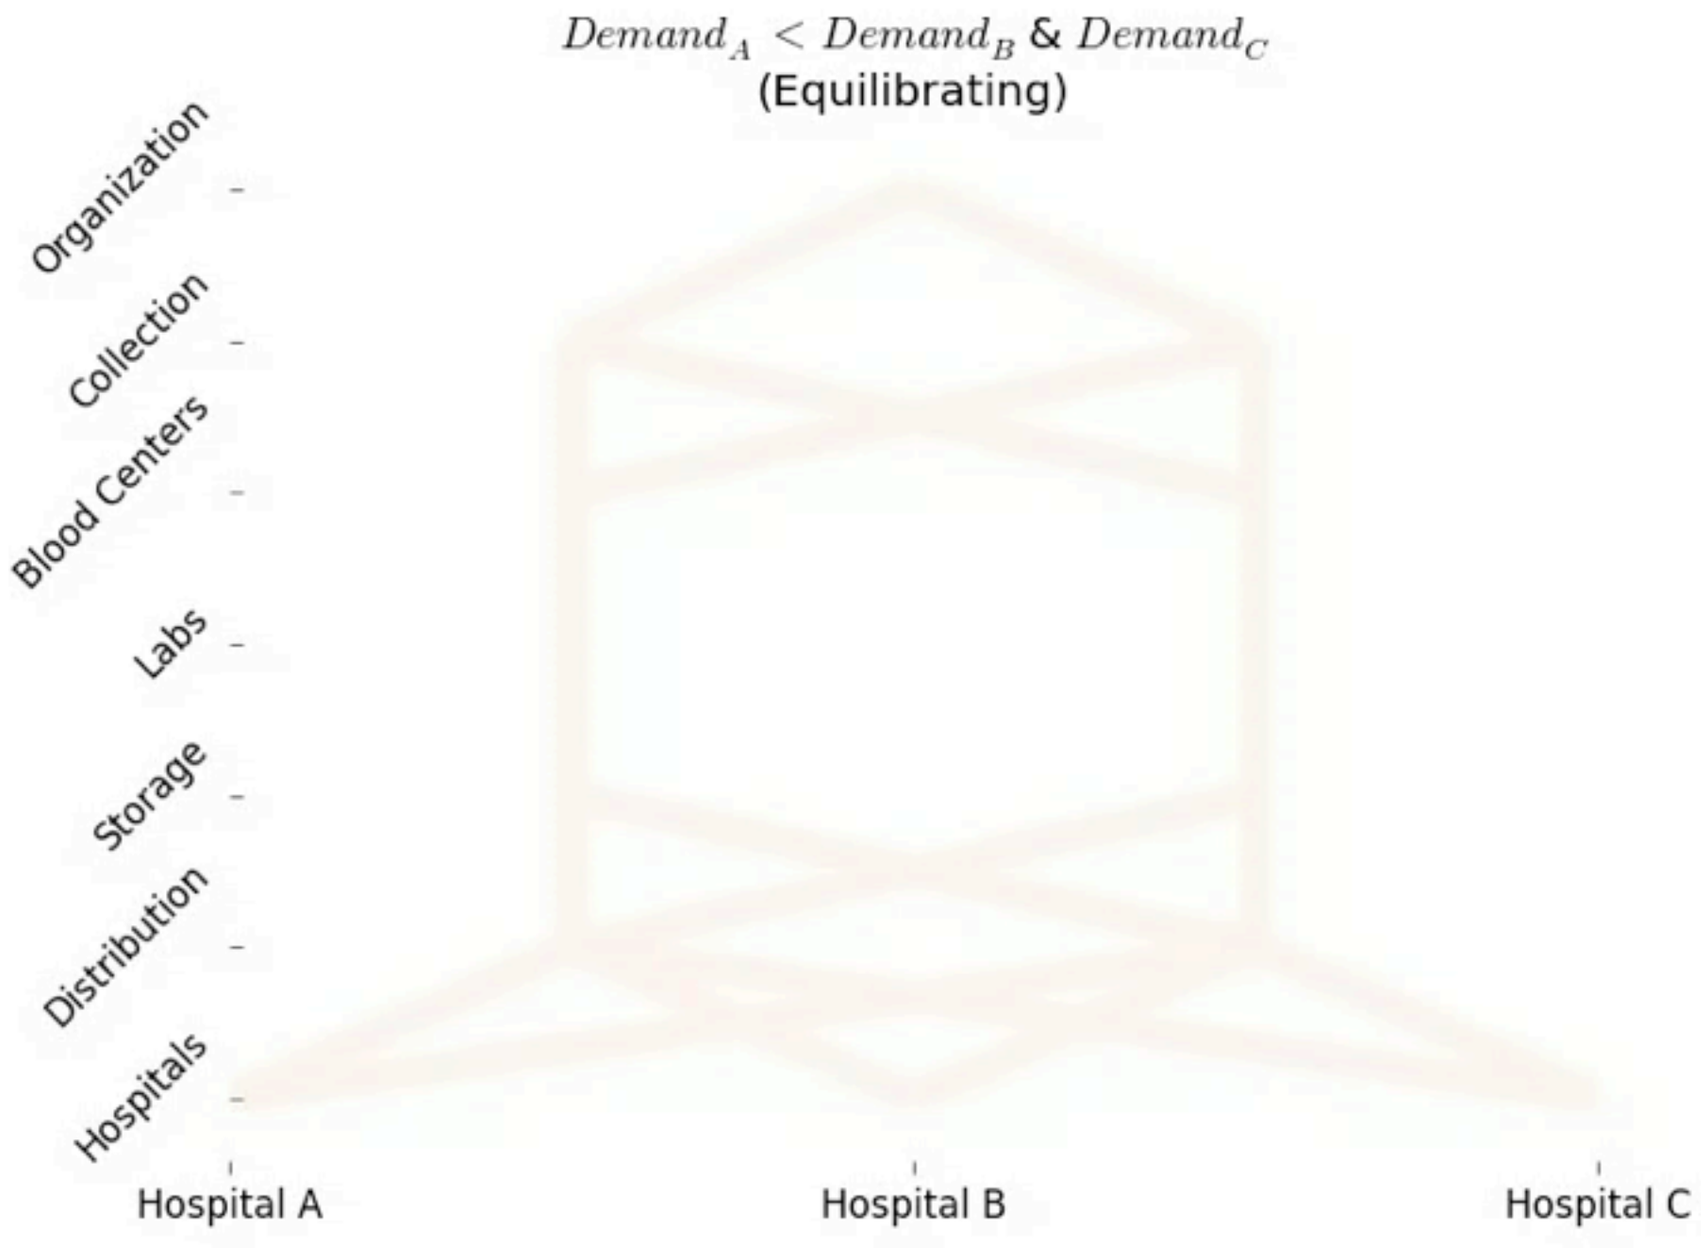

Video by Ian Gemp

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

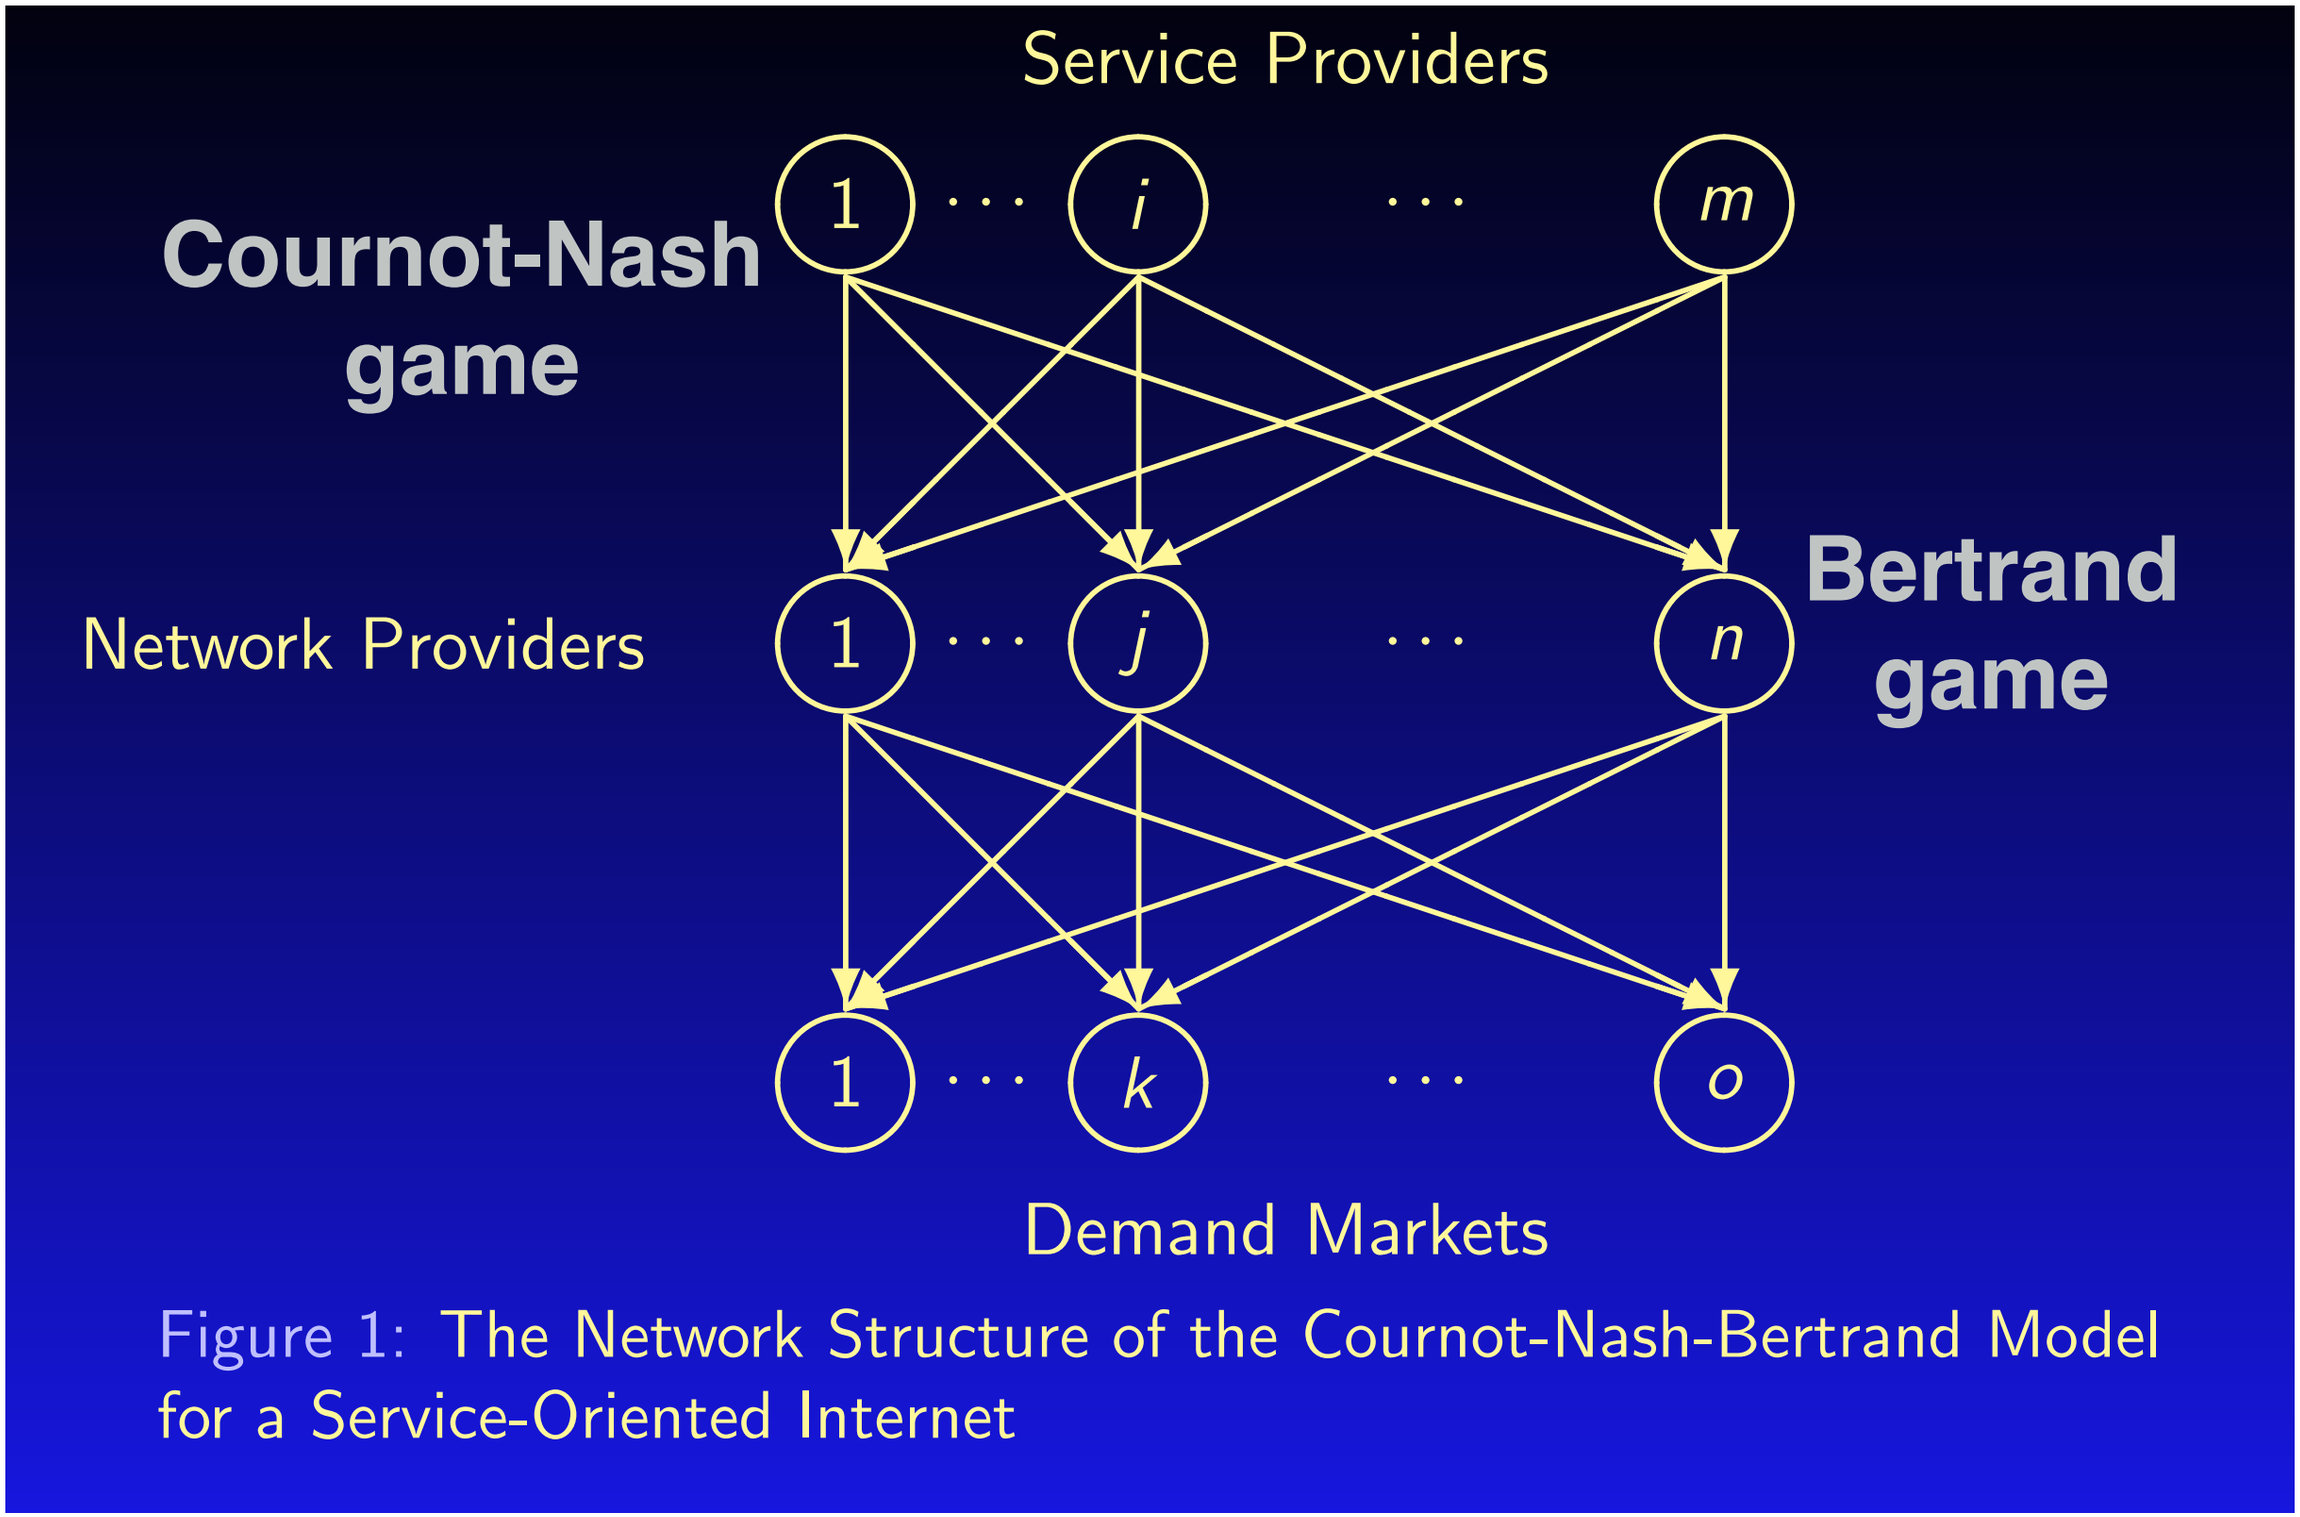

Figure 1: The Network Structure of the Cournot-Nash-Bertrand Model for a Service-Oriented Internet

Plan of the Tutorial

- Overview: 2:00-2:10 p.m.
- **Part I: Motivation: 2:10-2:30 p.m.**
- Part II: Mathematical Background: 2:30-3:30 p.m.
- COFFEE BREAK: 3:30-4:00 p.m.
- Part III: Algorithms: 4:00-5:00 p.m.
- Part IV: Applications and Future Work: 5:00-6:00 p.m.

Optimization

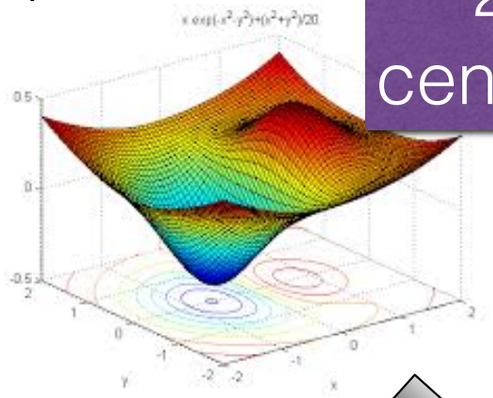

20th century AI

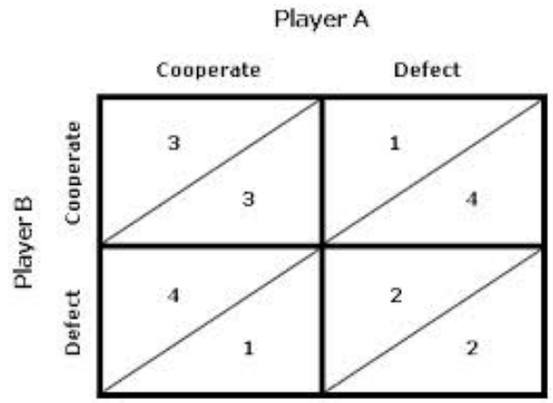

Game theory

21st century AI

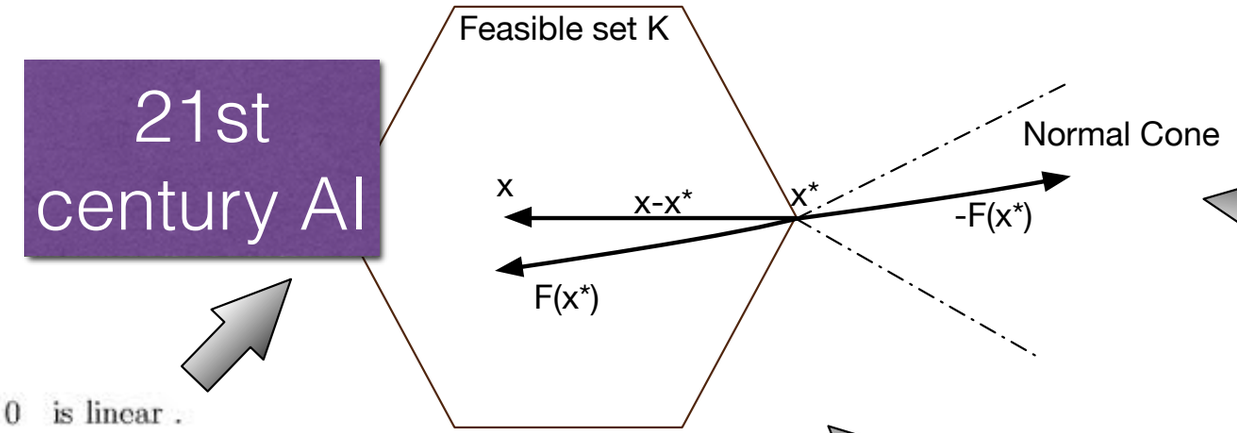

Variational Inequalities

Complementarity problems

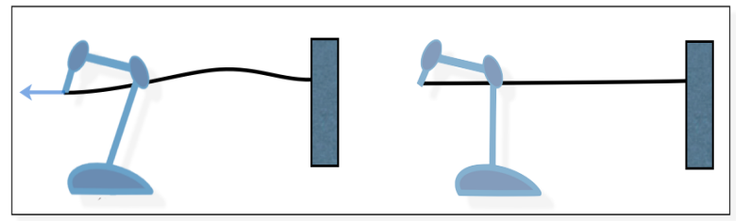

$$\frac{\partial u}{\partial x_1} + \frac{\partial u}{\partial x_2} = 0 \quad \text{is linear .}$$

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

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

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

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

Nonlinear equation solving



Traffic equilibrium problem

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

Objective Functions

- Typically, optimization involves minimization (or maximization) of some objective
    - KL-divergence, mean-squared error, likelihood
- Most problems require optimizing composite objectives
    - Model complexity vs fit to data
- Are objective functions even necessary?

Is AI dangerous?

http://chronicle.com/article/Is-Artificial-Intelligence-a/148763/

Imagine a machine programmed with the seemingly harmless, and ethically neutral, goal of getting as many paper clips as possible. First it collects them. Then. realizing that it could get more clips if it were smarter, it tries to improve its own algorithm to maximize computing power and collecting abilities. Unrestrained, its power grows by leaps and bounds, until it will do anything to reach its goal: collect paper clips, yes, but also buy paper clips, steal paper clips, perhaps transform all of earth into a paper-clip factory. "Harmless" goal, bad programming, end of the human race.

Nick Bostrom directs the Future of Humanity Institute at the University of Oxford.



Transfer Learning on Mars

(Dyar, Mount Holyoke; Boucher, Carey, Giguere, Mahadevan, UMass)

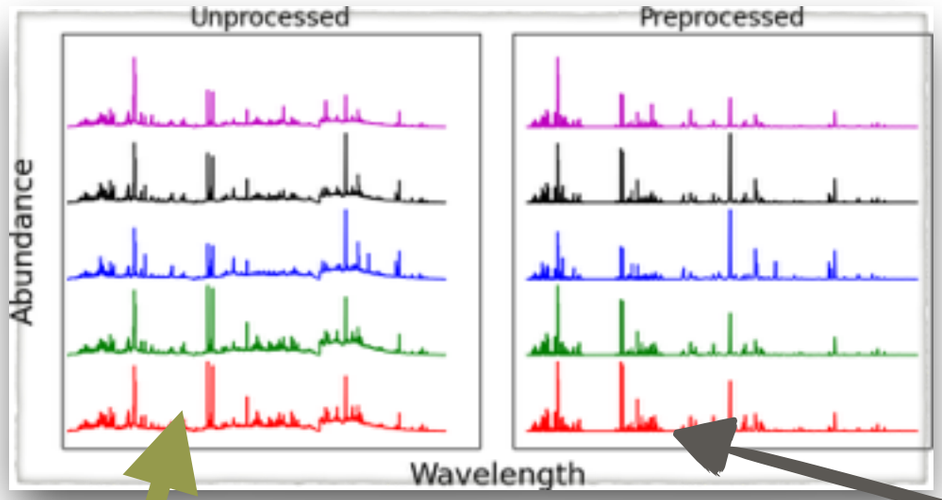

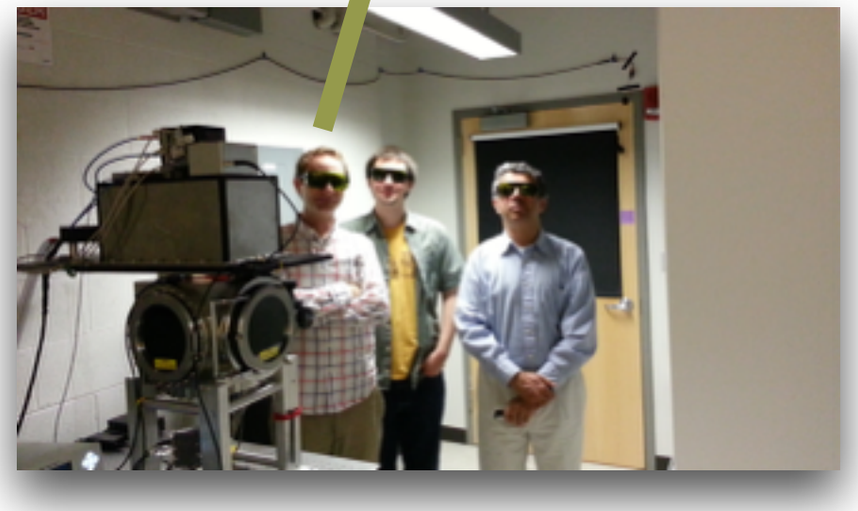

Same laser on Earth as on Mars

Curiosity zapping a rock with a laser



Transfer Learning from Mixture of Manifolds

(Boucher, Carey, Mahadevan, and Dyar, AAAI 2015)

Single manifold
(LLE, Laplacian Eigenmaps, Isomap)

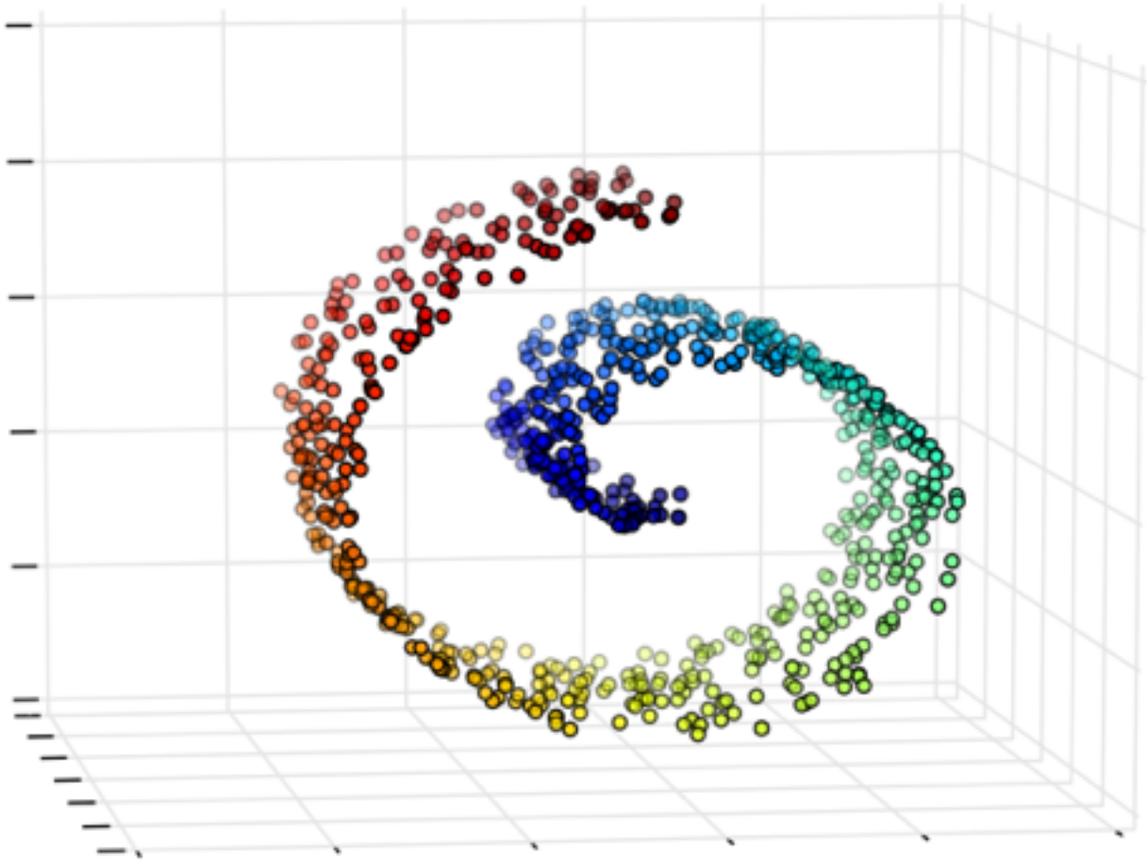

Low-rank Alignment (LRA)

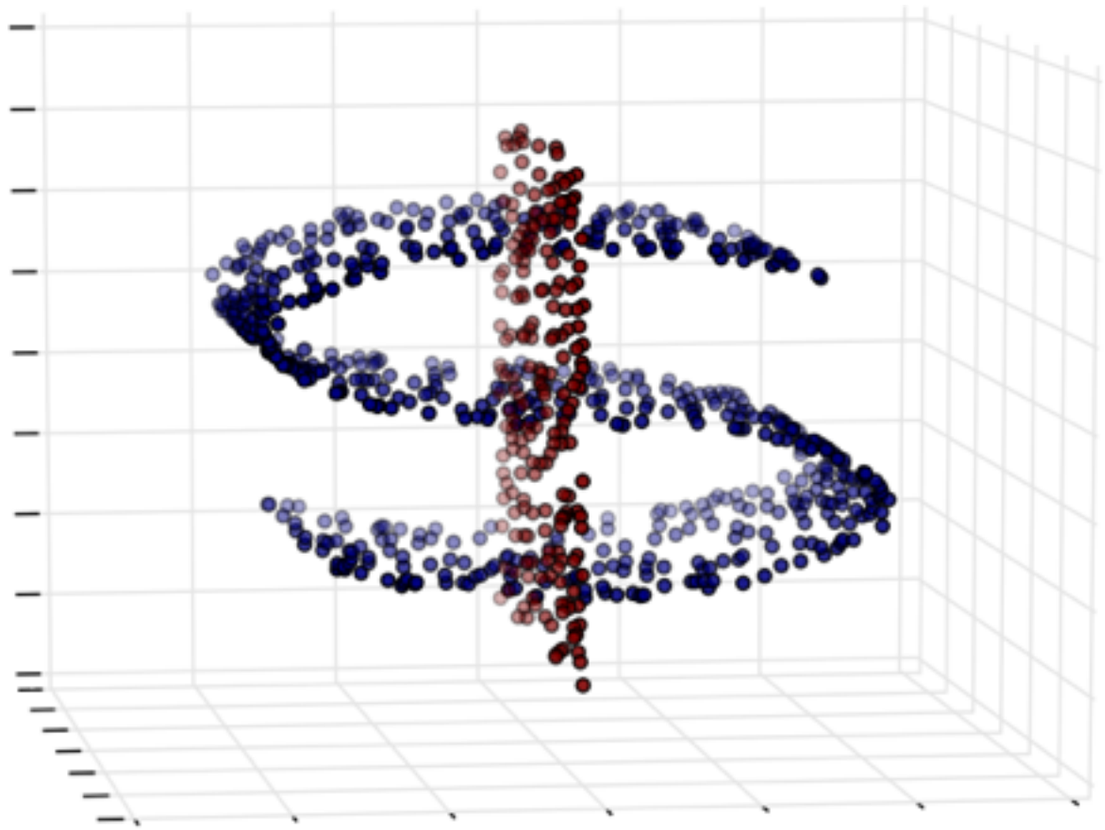

$$\min_R \frac{1}{2}||X - XR||_F^2 + \lambda||R||_*,$$

Multiple Objectives

$$\min_R \frac{1}{2}||X - XR||_F^2 + \lambda||R||_*,$$

Minimize reconstruction error

Minimize model complexity

How to handle non-differentiable objectives?

Composite Objective Functions in Machine Learning

"Sparse" Supervised learning

Lasso: $$\min_{x \in X} f(x) + g(x) : \min_{\beta \in \mathbb{R}^k} \|X\beta - y\|_2^2 + \lambda \|\beta\|_1$$

RO-TD: "Saddle Point" Reinforcement Learning

$$\min_x \|Ax - b\|_m + h(x) = \min_x \max_{\|y\|_n < 1} y^T (Ax - b) + h(x)$$

Unsupervised learning

Low-rank embedding: $$\min_R \frac{1}{2} ||X - XR||_F^2 + \lambda ||R||_*$$

$0 \in T(x)$

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)

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

Multiple data sources on Mars Curiosity Rover

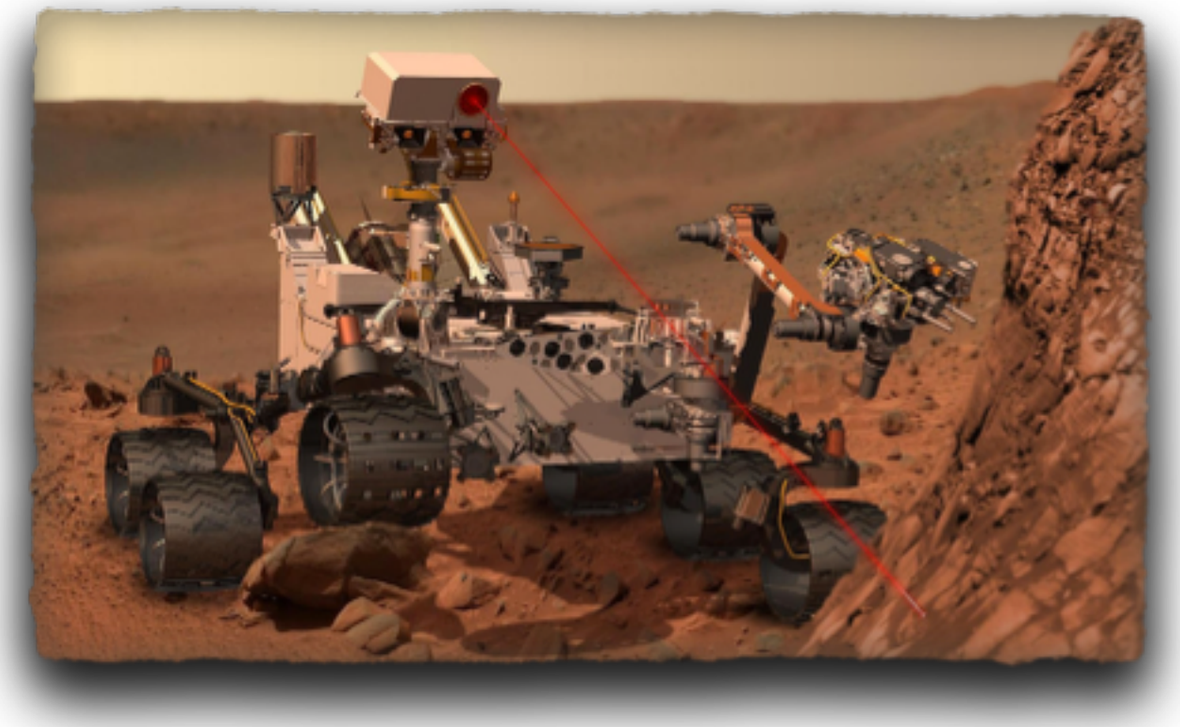

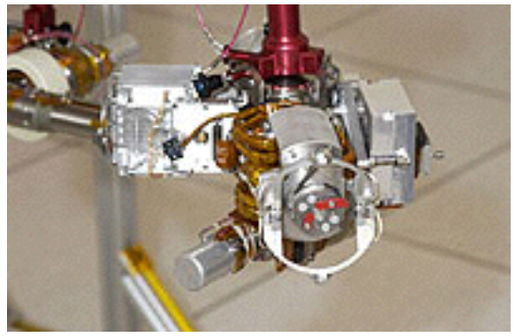

Rock Abrasion Tool

Miniature Thermal Emission Spectrometer

Moessbauer Spectrometer

Alpha Particle X-ray Spectrometer

Microscopic Imager

How to handle competition across across instruments and scientists?

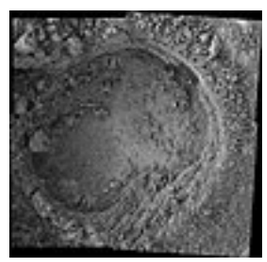

“The Internet is an equilibrium — we just have to identify the game (Scott Shenker)”

“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

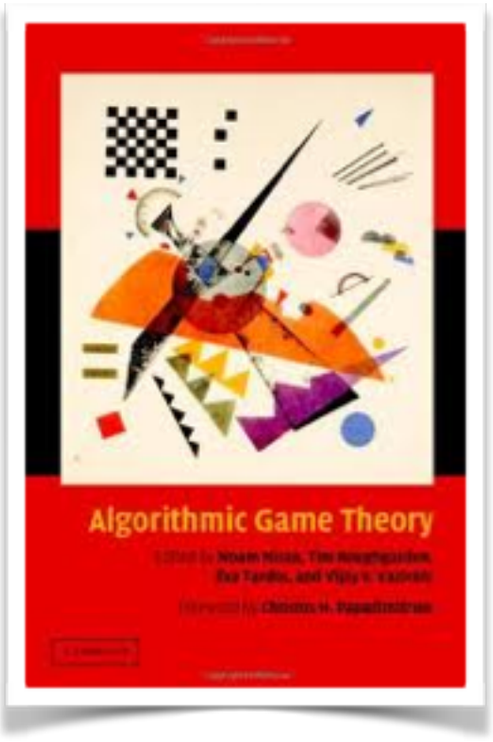

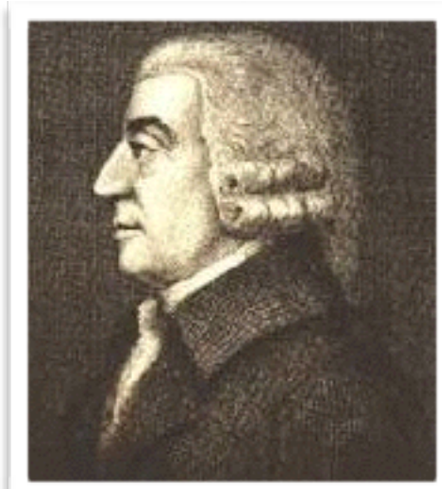

Adam Smith
The Wealth of Nations
1776

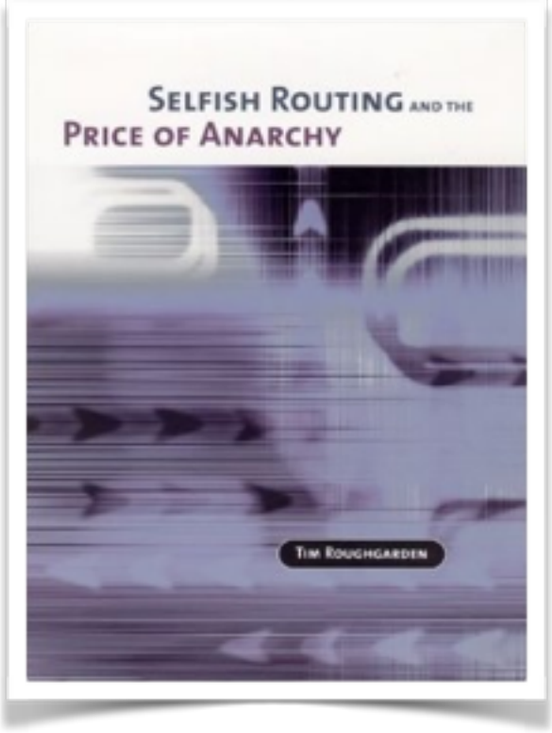

Competing Goals of the Internet: 1992-2014





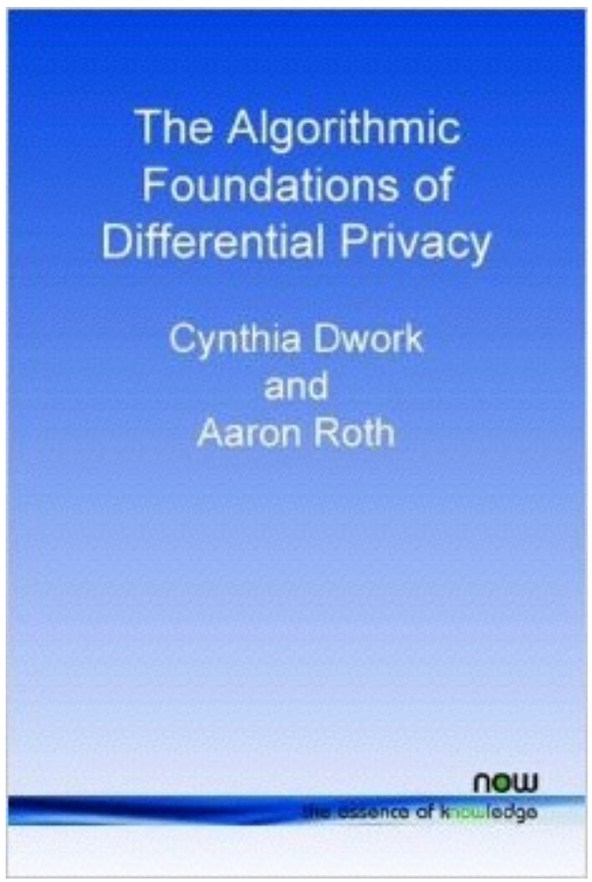

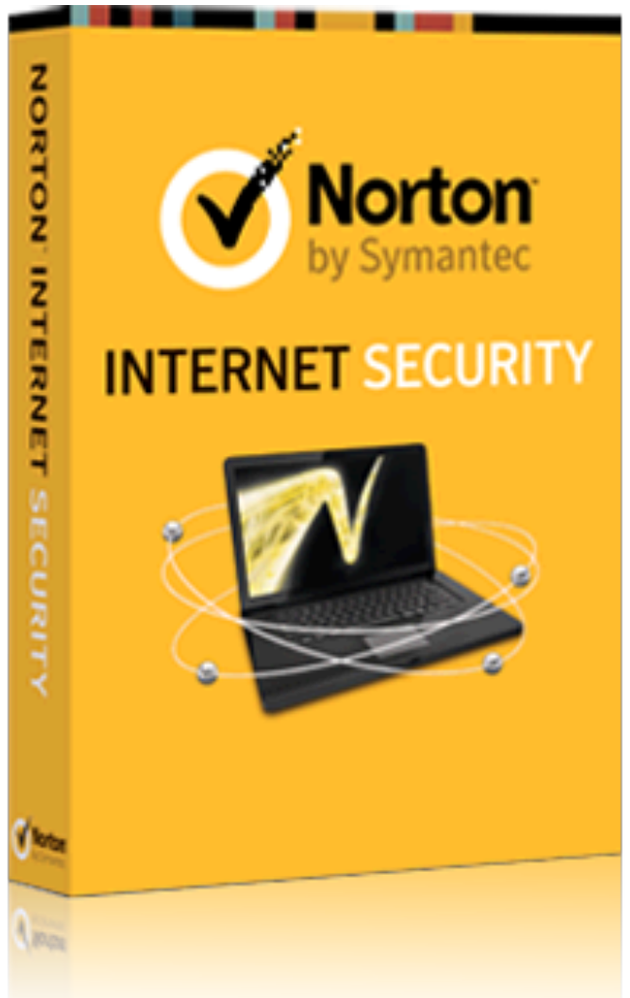

LA Times Story. June 05 2014

Verizon tells Netflix to stop blaming it for streaming issues

Verizon wants Netflix to stop blaming the Internet Service Provider for any streaming issues customers may be having when watching TV shows and movies.

"Netflix has been aware for some time that a few Internet middlemen have congestion issues with some IP Networks and nonetheless, Netflix has chosen to continue sending its traffic over those congested routes," said Verizon General Counsel Randal Milch. (Paul Sakuma / AP)

“Netflix” Cache Problem

(Dernbach, Kurose, Mahadevan, Technicolor)



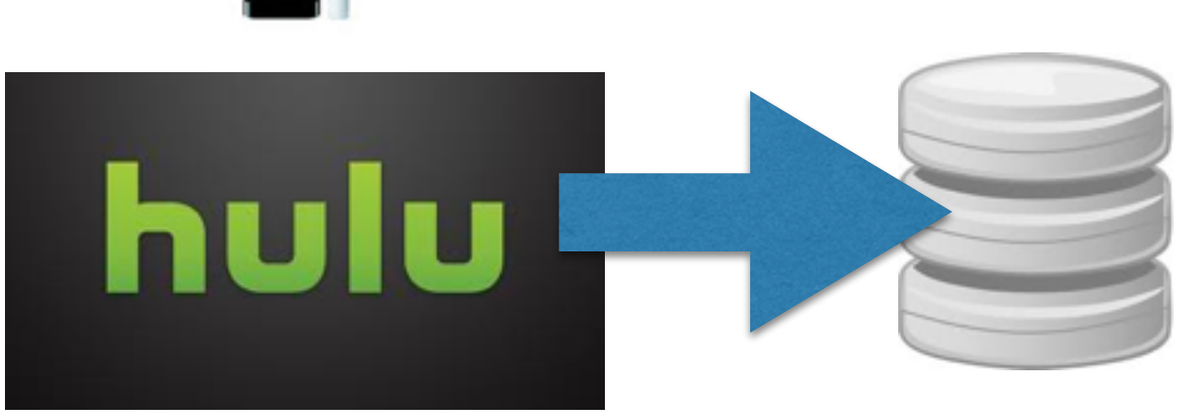





QA by Knowledge Base Completion

Input (Head and Label)	Predicted Tails
J. K. Rowling `influenced by`	*G. K. Chesterton*, J. R. R. Tolkien, *C. S. Lewis*, **Lloyd Alexander**, Terry Pratchett, Roald Dahl, Jorge Luis Borges, *Stephen King*, Ian Fleming
Anthony LaPaglia `performed in`	*Lantana*, *Summer of Sam*, *Happy Feet*, *The House of Mirth*, Unfaithful, **Legend of the Guardians**, Naked Lunch, X-Men, The Namesake
Camden County `adjoins`	**Burlington County**, *Atlantic County*, *Gloucester County*, Union County, Essex County, New Jersey, Passaic County, Ocean County, Bucks County
The 40-Year-Old Virgin `nominated for`	*MTV Movie Award for Best Comedic Performance*, *BFCA Critics' Choice Award for Best Comedy*, *MTV Movie Award for Best On-Screen Duo*, MTV Movie Award for Best Breakthrough Performance, **MTV Movie Award for Best Movie**, MTV Movie Award for Best Kiss, D. F. Zanuck Producer of the Year Award in Theatrical Motion Pictures, Screen Actors Guild Award for Best Actor - Motion Picture
David Foster `has the genre`	*Pop music*, *Pop rock*, Adult contemporary music, Dance music, **Contemporary R&B**, Soft rock, Rhythm and blues, Easy listening
Costa Rica football team `has position`	*Forward*, *Defender*, *Midfielder*, **Goalkeepers**, Pitchers, Infielder, Outfielder, Center, Defenseman
Lil Wayne `born in`	**New Orleans**, Atlanta, Austin, St. Louis, Toronto, New York City, Wellington, Dallas, Puerto Rico
WALL-E `has the genre`	Animations, Computer Animation, *Comedy film*, *Adventure film*, *Science Fiction*, **Fantasy**, Stop motion, *Satire*, Drama
Richard Crenna `has cause of death`	*Pancreatic cancer*, **Cardiovascular disease**, Meningitis, Cancer, Prostate cancers, Stroke, Liver tumour, Brain tumor, Multiple myeloma

Linguistic Relations

Google Dataset
Mikolov, 2013

a is to a*
as
b is to b*

Relation
capital-common-countries
capital-world
city-in-state
currency
family (gender inflections)
gram1-adjective-to-adverb
gram2-opposite
gram3-comparative
gram4-superlative
gram5-present-participle
gram6-nationality-adjective
gram7-past-tense
gram8-plural (nouns)
gram9-plural-verbs

Problem: Relational Embedding

$$\|\mathbf{h} + \mathbf{r} - \mathbf{t}\|_{\ell_{1/2}},\ \mathbf{r} \in \mathbb{R}^k$$

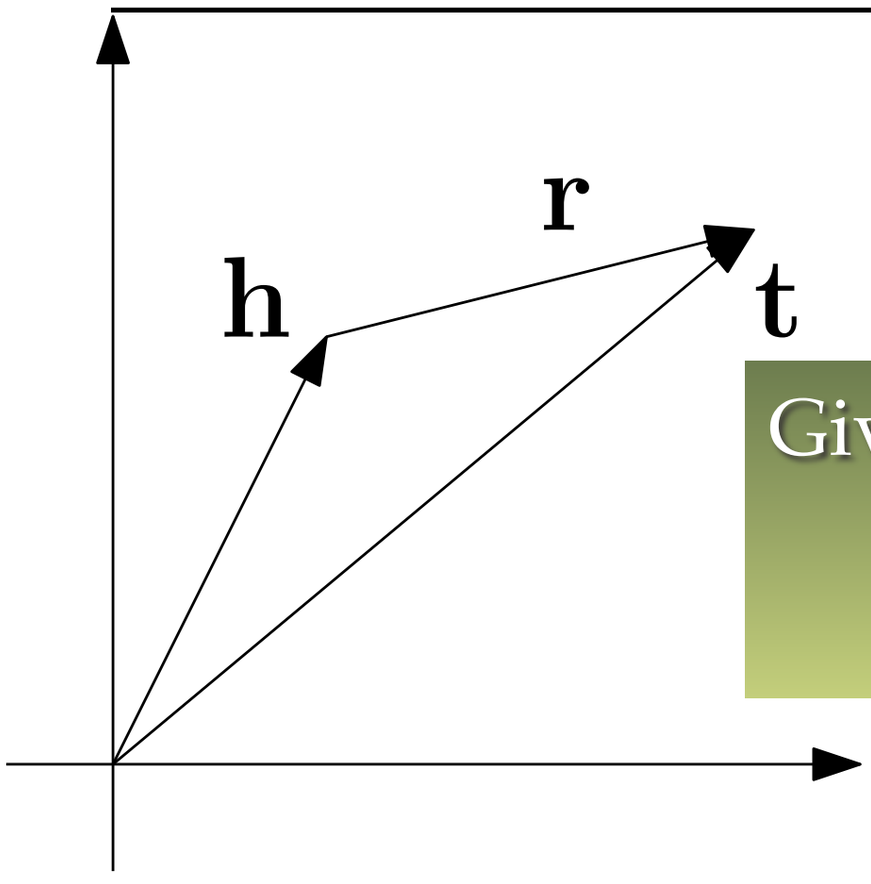

(a) TransE

Given a knowledge base of relational tuples,
construct a vector-space embedding
that is faithful

Linguistic Reasoning by Vector Arithmetic

$$queen \approx king - man + woman$$

$$\arg\max_{b^* \in V} \left(\cos\left(b^*, b - a + a^*\right)\right)$$

Mikolov et al., 2013

Additive vs. Multiplicative Objective Functions

(Levy and Goldberg, 2014)

$$\arg\max_{b^* \in V} \frac{\cos(b^*, b)\cos(b^*, a^*)}{\cos(b^*, a) + \varepsilon}$$

Active Research Area

Model	Score function $f_r(\mathbf{h}, \mathbf{t})$	# Parameters
TransE (Bordes et al. 2013b)	$\|\mathbf{h} + \mathbf{r} - \mathbf{t}\|_{\ell_{1/2}}$, $\mathbf{r} \in \mathbb{R}^k$	$O(n_e k + n_r k)$
Unstructured (Bordes et al. 2012)	$\|\mathbf{h} - \mathbf{t}\|_2^2$	$O(n_e k)$
Distant (Bordes et al. 2011)	$\|W_{rh}\mathbf{h} - W_{rt}\mathbf{t}\|_1$, $W_{rh}, W_{rt} \in \mathbb{R}^{k \times k}$	$O(n_e k + 2n_r k^2)$
Bilinear (Jenatton et al. 2012)	$\mathbf{h}^\top W_r \mathbf{t}, W_r \in \mathbb{R}^{k \times k}$	$O(n_e k + n_r k^2)$
Single Layer	$\mathbf{u}_r^\top f(W_{rh}\mathbf{h} + W_{rt}\mathbf{t} + \mathbf{b}_r)$ $\mathbf{u}_r, \mathbf{b}_r \in \mathbb{R}^s, W_{rh}, W_{rt} \in \mathbb{R}^{s \times k}$	$O(n_e k + n_r(sk + s))$
NTN (Socher et al. 2013)	$\mathbf{u}_r^\top f(\mathbf{h}^\top \mathbf{W}_r \mathbf{t} + W_{rh}\mathbf{h} + W_{rt}\mathbf{t} + \mathbf{b}_r)$ $\mathbf{u}_r, \mathbf{b}_r \in \mathbb{R}^s, \mathbf{W}_r \in \mathbb{R}^{k \times k \times s}, W_{rh}, W_{rt} \in \mathbb{R}^{s \times k}$	$O(n_e k + n_r(sk^2 + 2sk + 2s))$
TransH (this paper)	$\|(\mathbf{h} - \mathbf{w}_r^\top \mathbf{h}\mathbf{w}_r) + \mathbf{d}_r - (\mathbf{t} - \mathbf{w}_r^\top \mathbf{t}\mathbf{w}_r)\|_2^2$ $\mathbf{w}_r, \mathbf{d}_r \in \mathbb{R}^k$	$O(n_e k + 2n_r k)$

Plan of the Tutorial

- Overview: 2:00-2:10 p.m.
- Part I: Motivation: 2:10-2:30 p.m.
- **Part II: Mathematical Background: 2:30-3:30 p.m.**
- COFFEE BREAK: 3:30-4:00 p.m.
- Part III: Algorithms: 4:00-5:00 p.m.
- Part IV: Applications and Future Work: 5:00-6:00 p.m.

Problem Structure

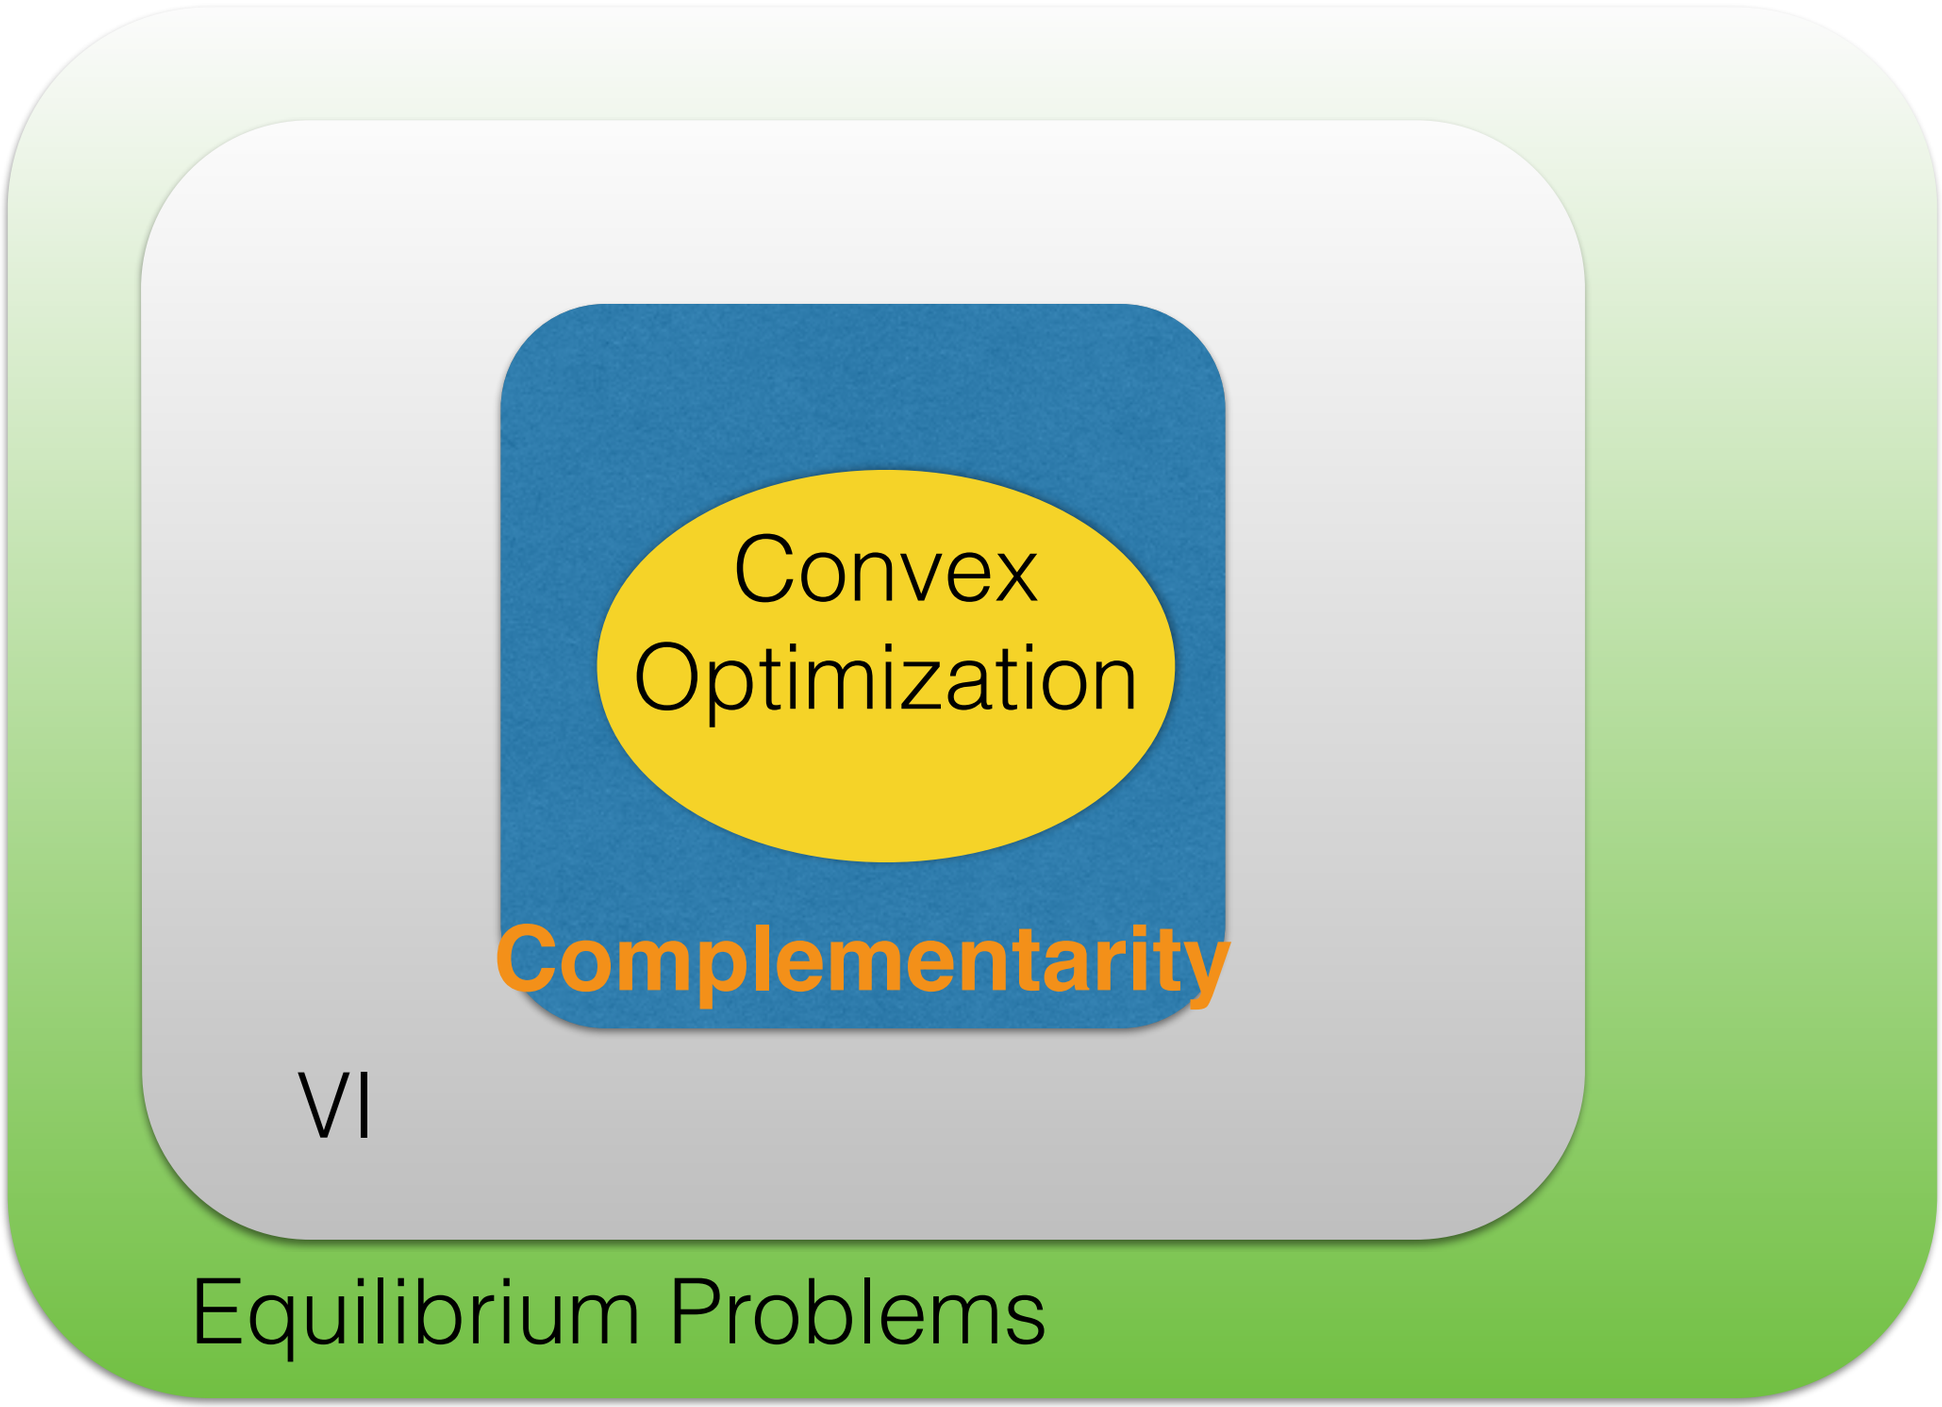

Complementarity: Optimization without objectives

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

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

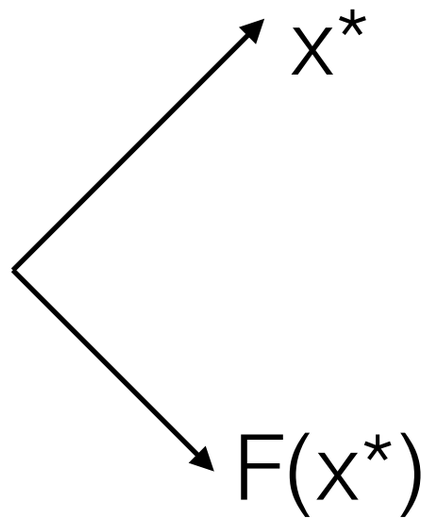

Arrow-Debreau Economic Equilibrium

L commodities

N manufacturers

M consumers

A: “technology” matrix of size L x N

p: price vector

z: production or activity vector

z is non-negative,
$-A^T p$ is non-negative
p is perpendicular to Az

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

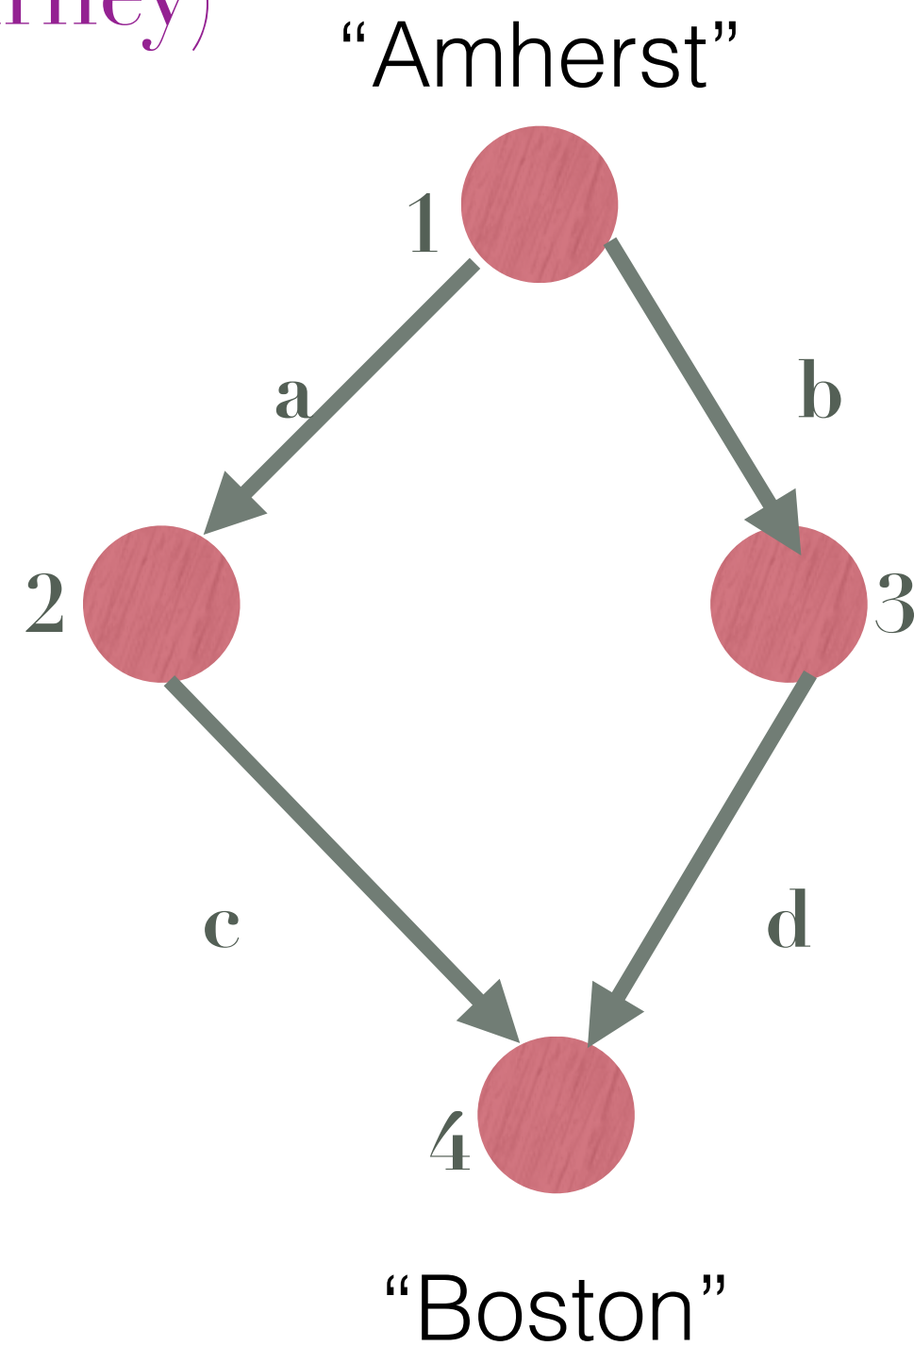

Traffic Network Equilibrium

(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

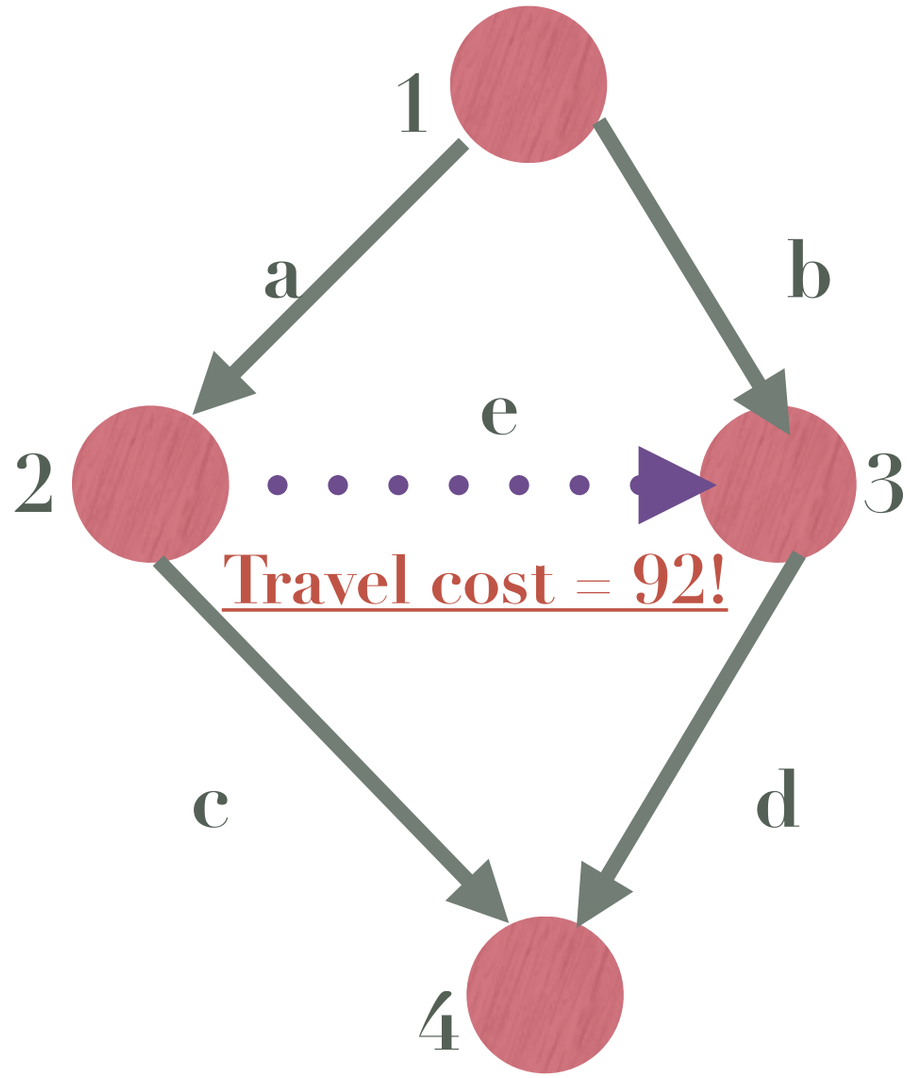

Braess' Paradox!

Guardian
March 2014

Real-life examples of Braess Paradox

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

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,

What if They Closed 42d Street and Nobody Noticed?

By GINA KOLATA
Published: December 25, 1990

New York Times

ON Earth Day this year, New York City's Transportation Commissioner decided to close 42d Street, which as every New Yorker knows is always congested. "Many predicted it would be doomsday," said the Commissioner, Lucius J. Riccio. "You didn't need to be a rocket scientist or have a sophisticated computer queuing model to see that this could have been a major problem."

But to everyone's surprise, Earth Day generated no historic traffic jam. Traffic flow actually improved when 42d Street was closed.

To mathematicians, this may be a real-world example of Braess's paradox, a statistical theorem that holds that when a network of streets is already jammed with vehicles, adding a new street can make traffic flow even more slowly. Seeking Out a New Street

The reason is that in crowded conditions, drivers will pile into a new street, clogging both it and the streets that provide access to it. By the same token, removing a major thoroughfare may actually ease congestion on the streets that normally provide access to it. And because other major streets are already overcrowded, diverting still more traffic to them may not make much difference.

FACEBOOK
TWITTER
GOOGLE+
EMAIL
SHARE
PRINT
REPRINTS



Convex set

contains the line segment between any two points in the set

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

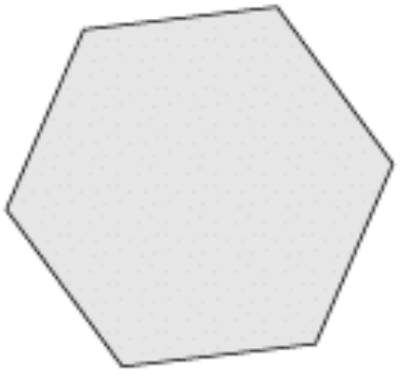

convex

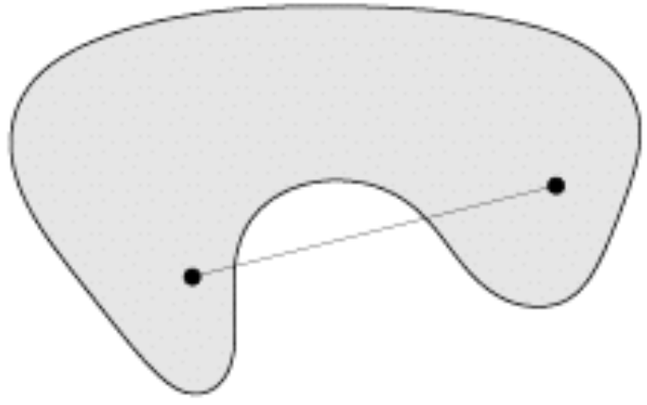

not convex

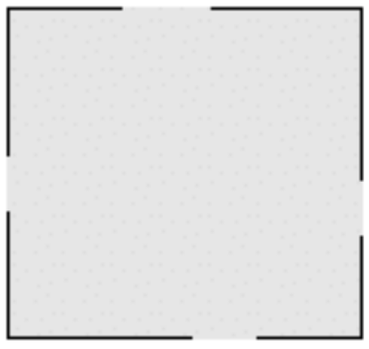

not convex

Three Types of Cones

- A cone K is a convex set such that if x ε K, then α x ε K for any nonnegative α
- Example: the set of non-negative matrices
- the dual cone K^* is defined as $K^* = \{y \in \mathbb{R}^n | \langle y, x \rangle \geq 0, \forall x \in K\}$
- Normal cone: Subdifferential of indicator function

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

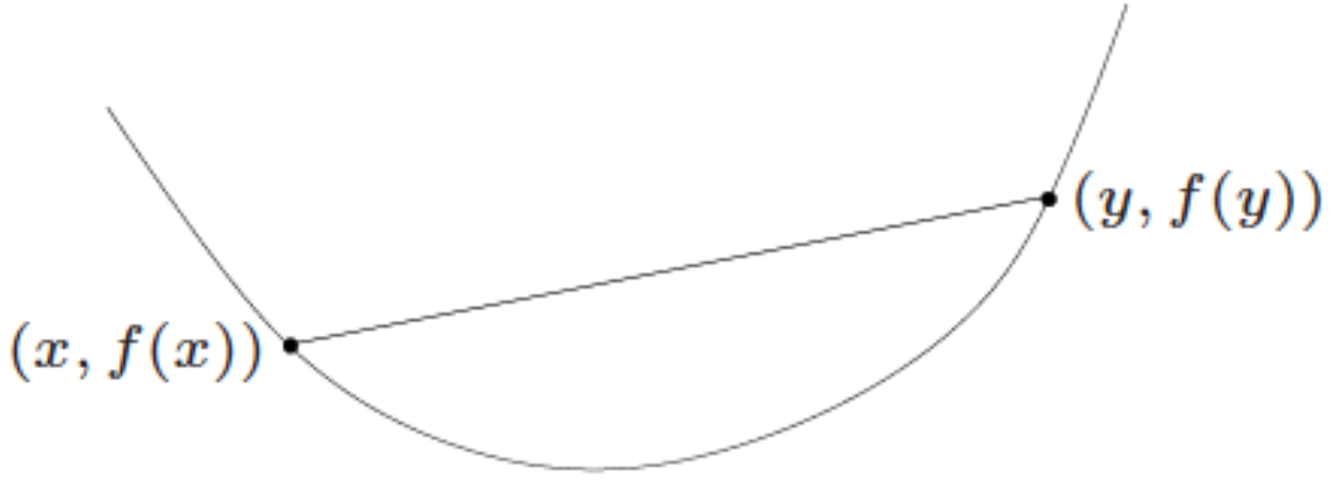

f is **strictly convex** if Jensen's inequality is strict for $0 < \theta < 1$, $x \neq y$

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

Indicator function

the indicator function of a set C is

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

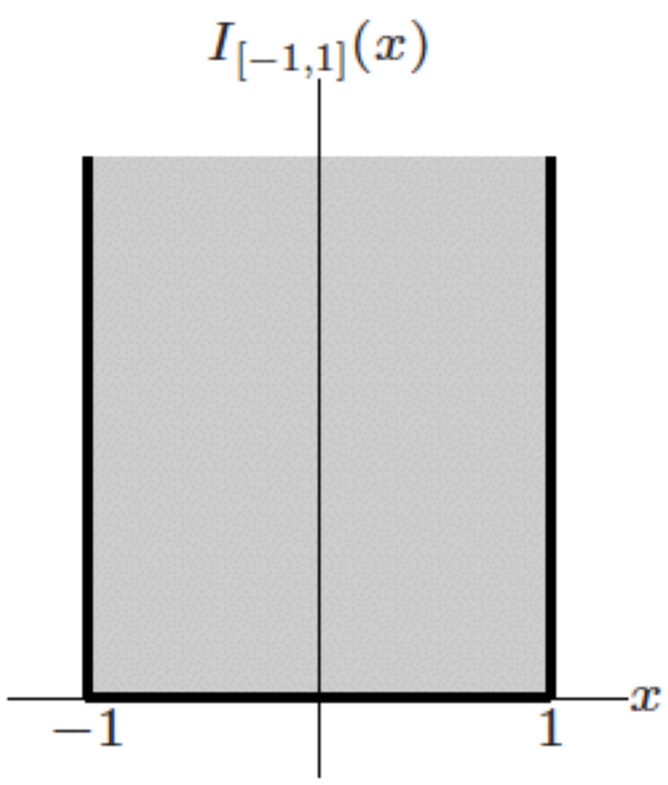

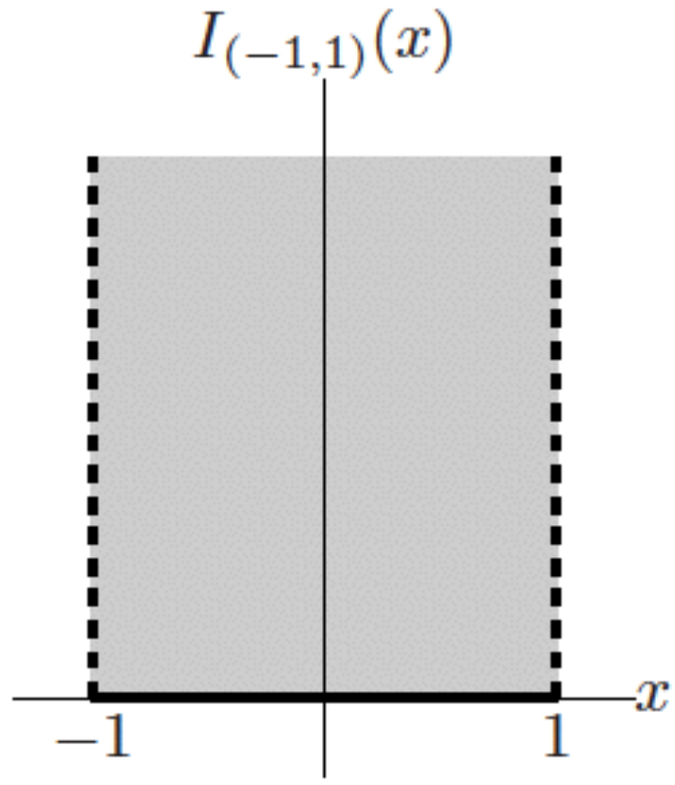

the indicator function of a convex set is a convex function

Subgradient of a function

s is a subgradient of a convex function f at $x \in \mathbf{dom}\, f$ if

$$f(y) \geq f(x) + s^T(y - x) \quad \text{for all } y \in \mathbf{dom}\, f$$

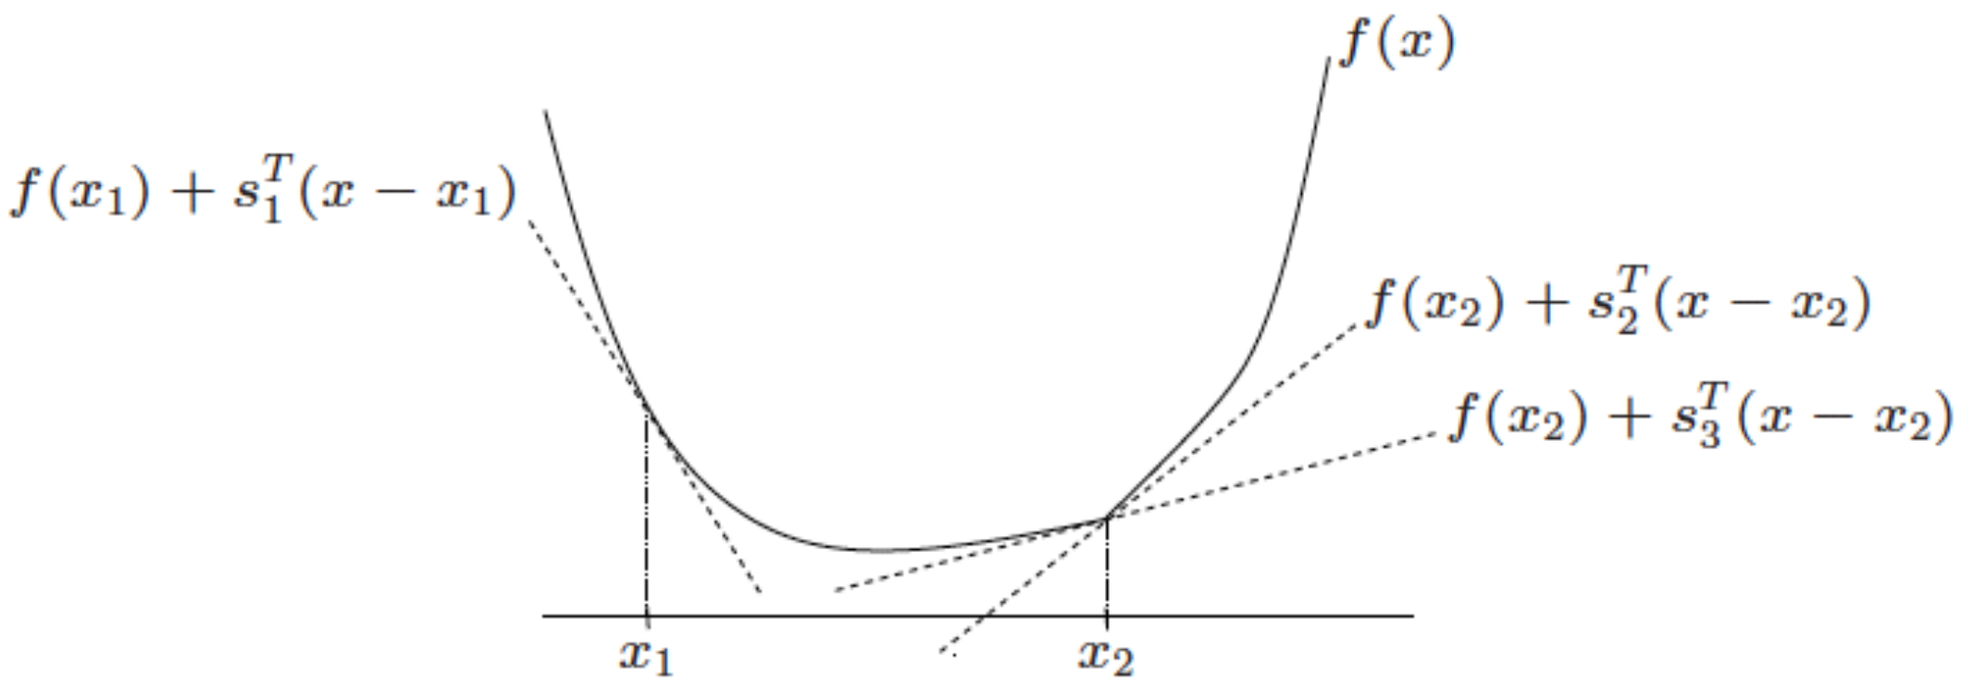

s_2, s_3 are subgradients at x_2; s_1 is a subgradient at x_1

Point-to-Set Analysis

- Point-to-set analysis studies mappings F:X —> Y where Y is a subset of X
- We can view such mappings as a relation
  - $F = \{(x,y) \mid y \in F(x)\}$
- Operations on such relations are defined as usual
  - $\lambda F = \{(x, \lambda y) \mid y \in F(x)\}$
  - $I + \lambda F = \{(x, x+ \lambda y) \mid y \in F(x)\}$
  - $(I + \lambda F)^{-1} = \{(x+ \lambda y, x) \mid y \in F(x)\}$

Normal Cone

Subdifferential of a convex function:

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

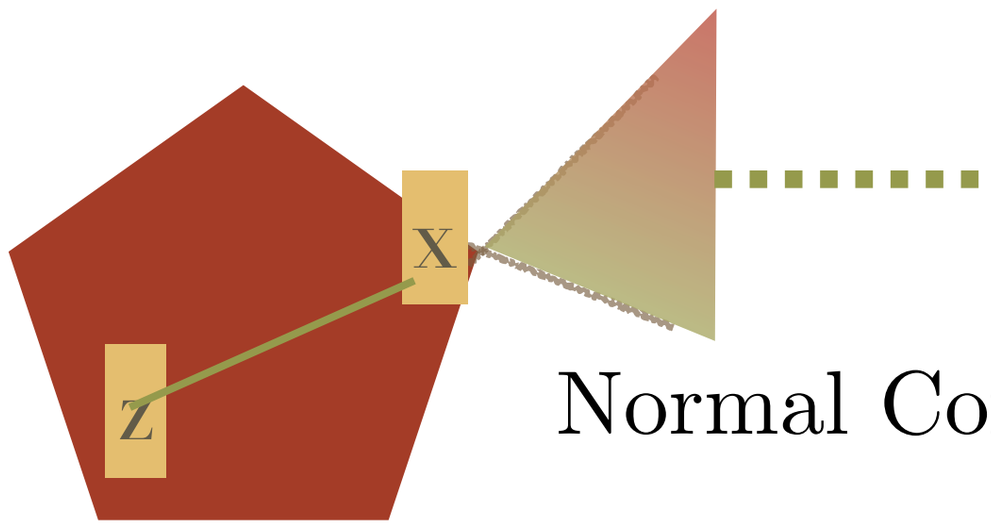

$I_{\mathcal{C}}(x)$

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

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, $0 \in \partial f(x^*)$

Variational Inequality

(Stampacchia, 1960s)

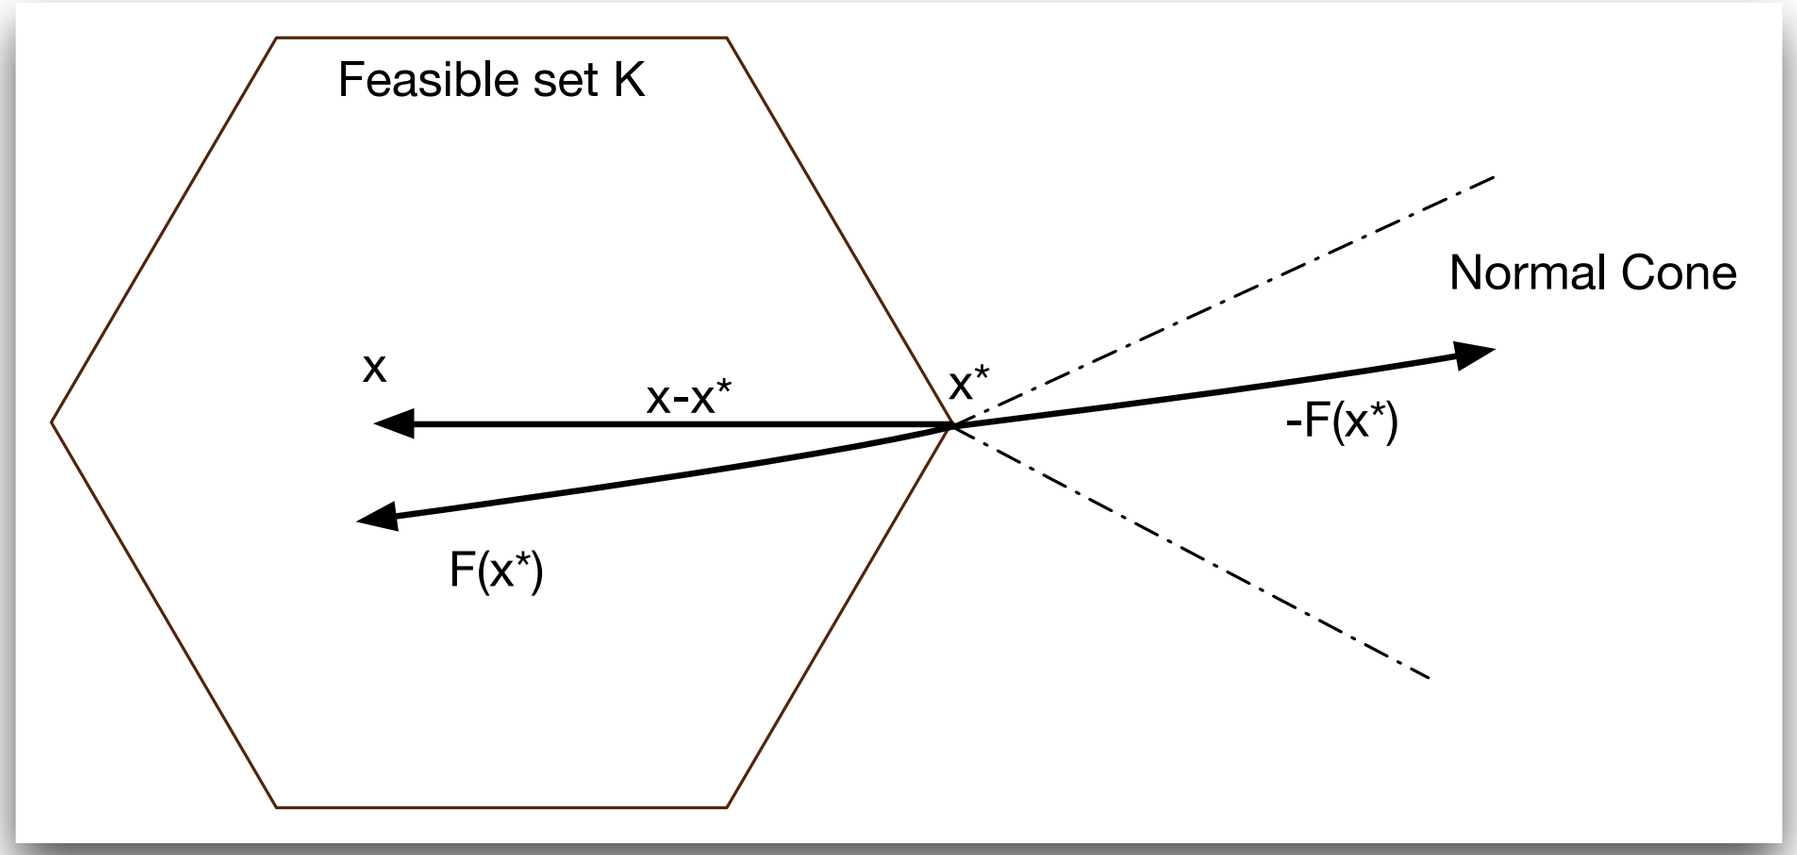

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

VI as monotone inclusion

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

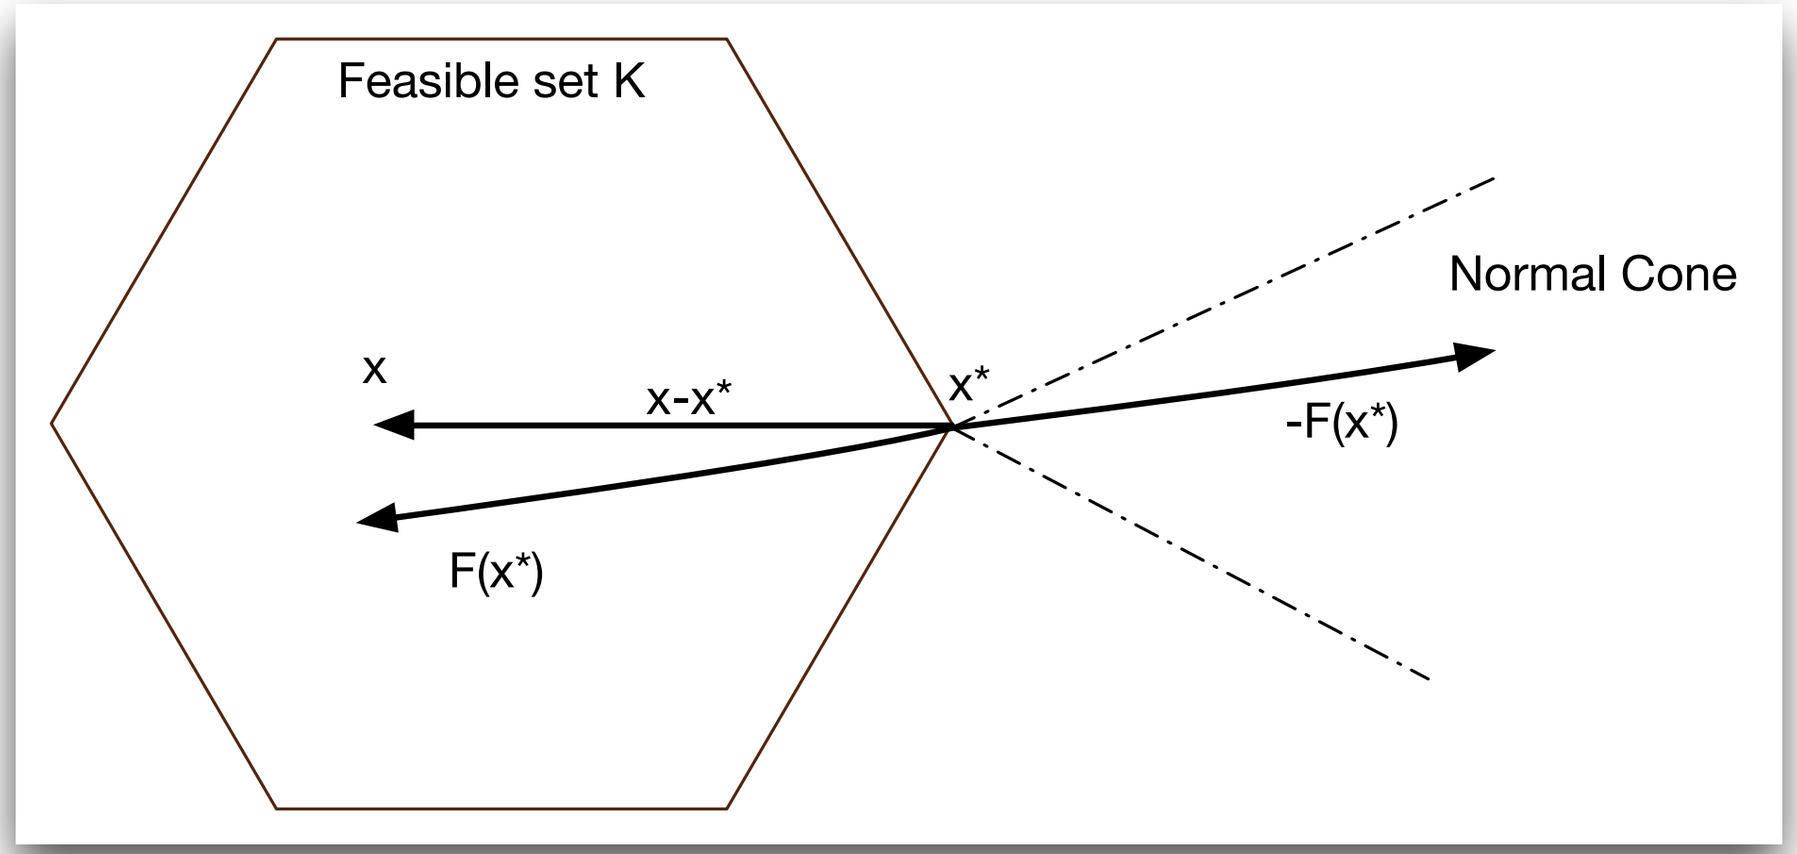

Fixed Point Formulation of VI

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.*

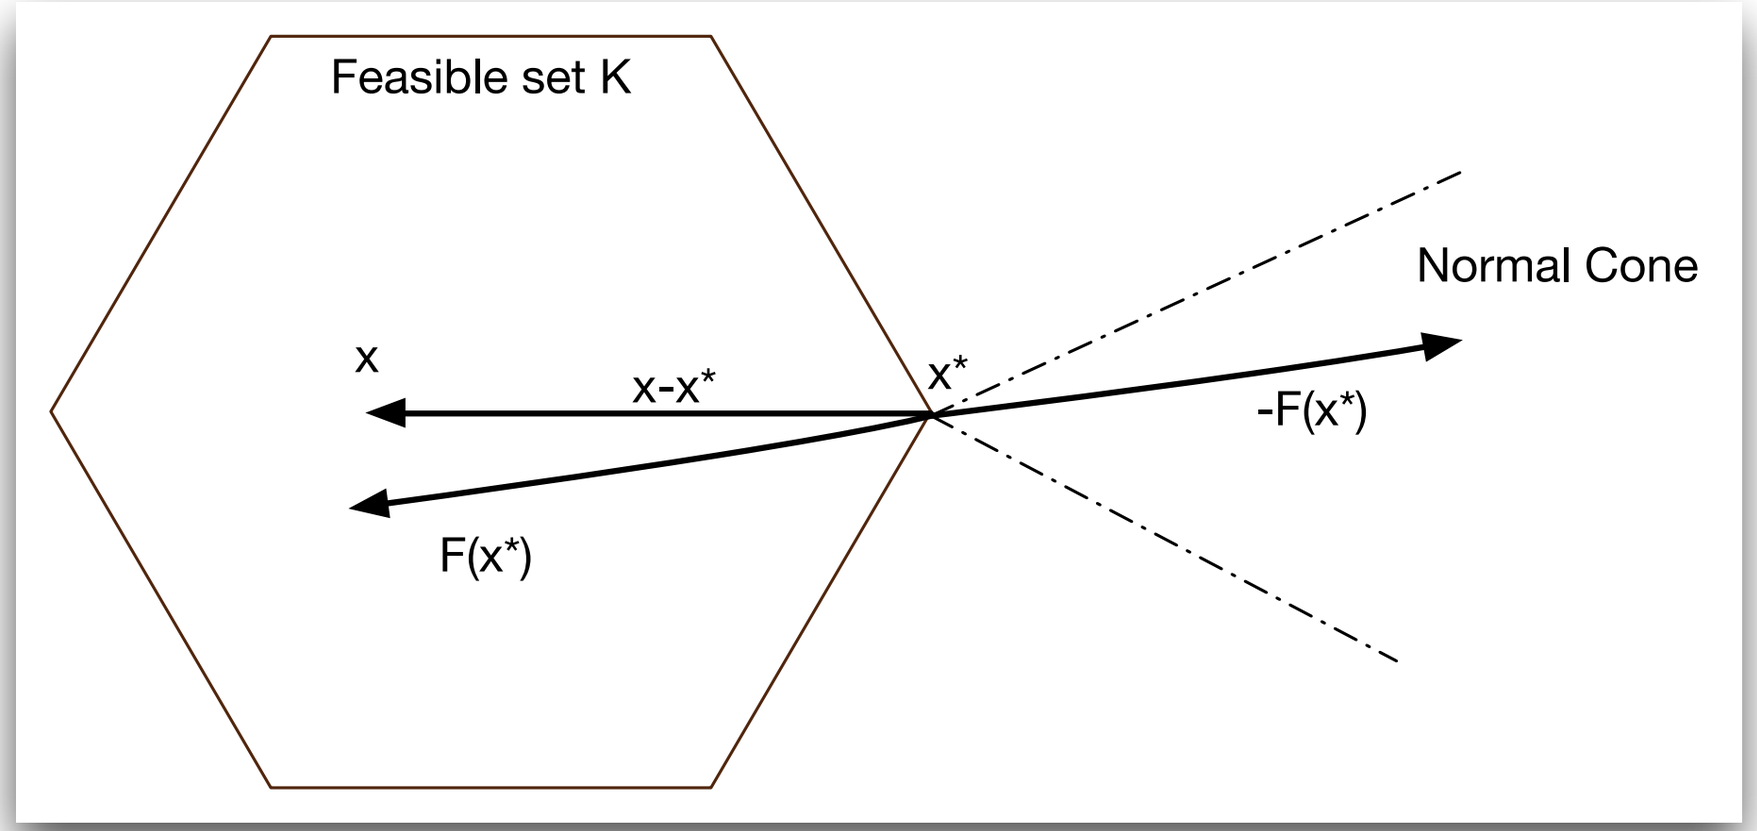

VI: Existence of a Solution

Definition 1. *The finite-dimensional variational inequality problem VI(F,K) involves finding a vector* $x^* \in K \subset \mathbb{R}^n$ *such that*

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

where $F : K \to \mathbb{R}^n$ *is a given continuous function and* K *is a given closed convex set, and* $\langle ., . \rangle$ *is the standard inner product in* $\mathbb{R}^n$*.*

Theorem 1. *Suppose* K *is compact, and that* $F : K \to \mathbb{R}^n$ *is continuous. Then, there exists a solution to VI*(F, K)*.*

Proof: Follows from the Banach fixed point theorem.

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

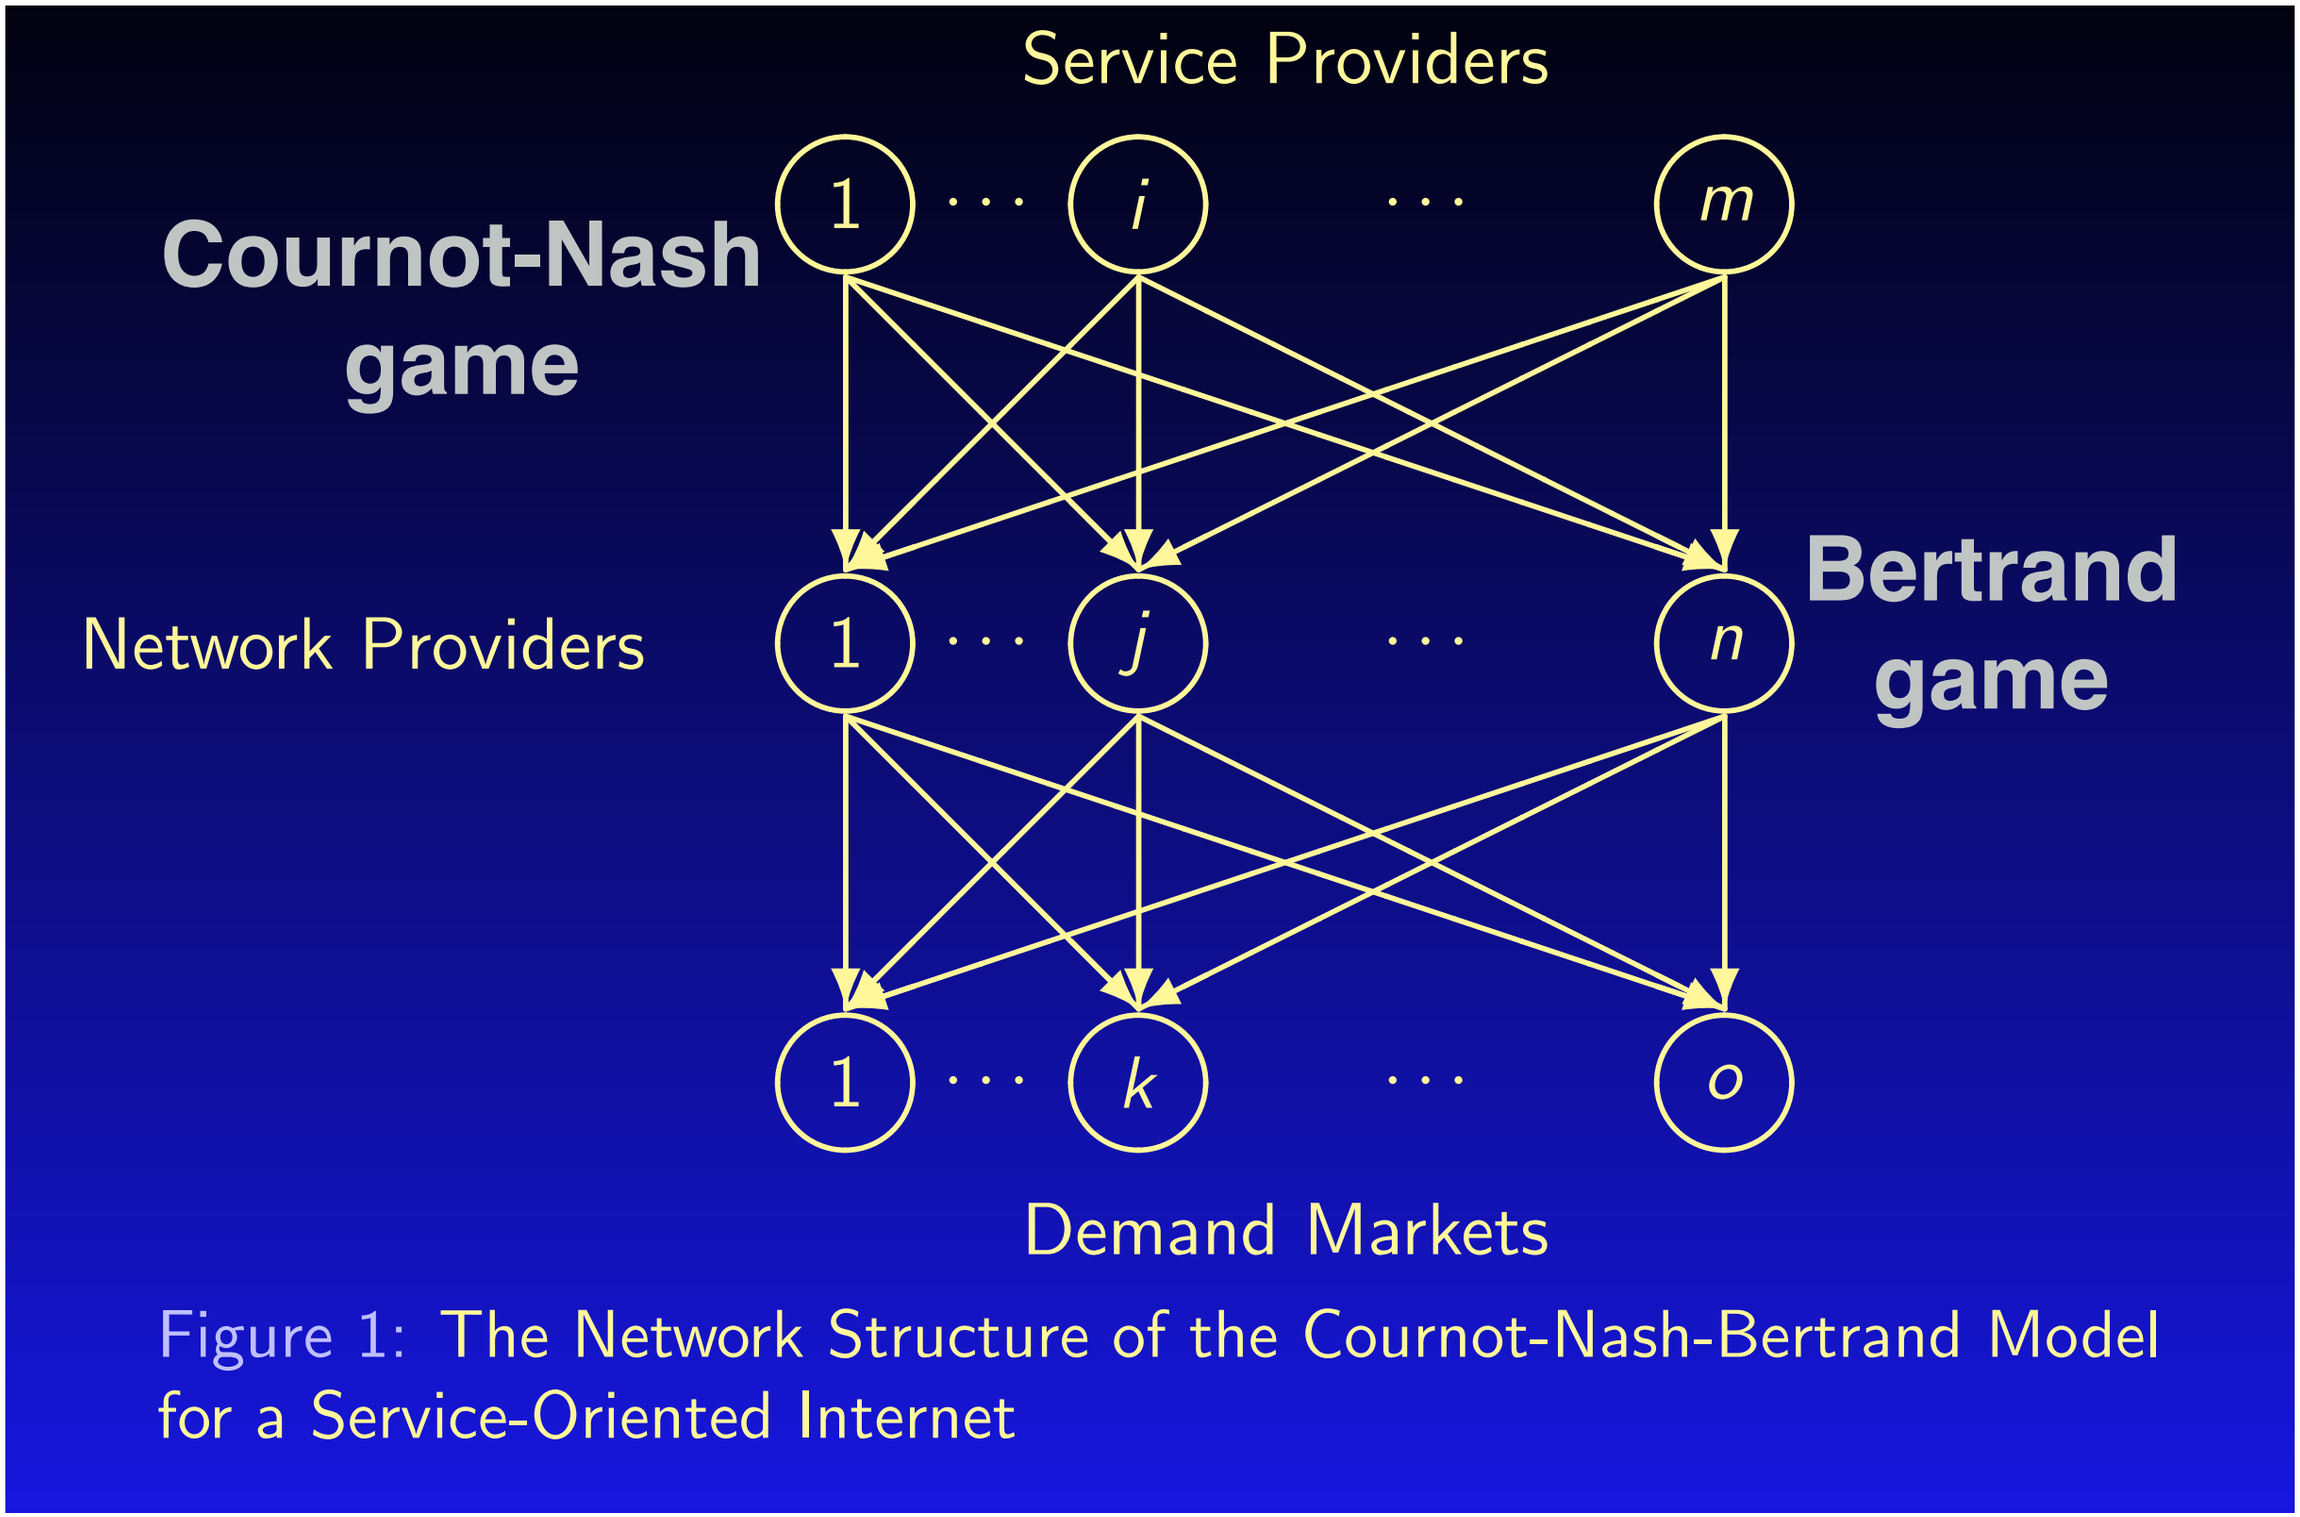

Figure 1: The Network Structure of the Cournot-Nash-Bertrand Model for a Service-Oriented Internet

Example: Internet Model

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.

Example

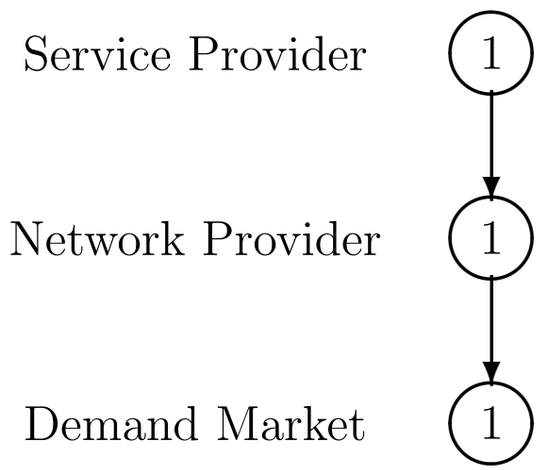

$$\hat{f}_1(Q_{111}) = Q_{111}^2 + Q_{111}.$$

$$\hat{\rho}_{111}(Q_{111}, q_{111}) = -2Q_{111} + q_{111} + 78$$

$$c_{111}(Q_{111}, q_{111}) = (q_{111} - 1)^2$$

$$F_{111}^1(X) = 2Q_{111} + 1 + \pi_{111} + 2Q_{111} - q_{111} - 78 + 2Q_{111}$$

$$J(Q_{111}, q_{111}, \pi_{111}) = \begin{pmatrix} 6 & -1 & 1 \\ 0 & 2 & 0 \\ -1 & 0 & 2 \end{pmatrix}$$

$$F_{111}^2(X) = 2q_{111} - 2$$

$$F_{111}^3(X) = -Q_{111} + 2\pi_{111}$$

Strictly diagonally dominant

Positive-definite matrix

Monotonicity

Definition 3.2 *Let $X \subseteq \mathbb{R}^n$ be a nonempty set and $F : X \to \mathbb{R}^n$ be a given function. Then F is said to be*

(i) monotone *(on X) if $(x-y)^T(F(x)-F(y)) \geq 0$ for all $x, y \in X$;*

(ii) strictly monotone *(on X) if $(x-y)^T(F(x)-F(y)) > 0$ for all $x, y \in X$ with $x \neq y$;*

(iii) strongly monotone *(on X) if there exists a constant $\mu > 0$ (called* modulus*) such that $(x-y)^T(F(x)-F(y)) \geq \mu \|x-y\|^2$ for all $x, y \in X$.*

Jacobian Properties

∇F	**F**	$F = \nabla f$
Positive semi-definite	Monotone	Convex
Positive-definite	Strictly monotone	Strictly convex
Strongly positive-definite	Strongly monotone	Strongly convex

Optimization vs. VIs

Theorem 2. *Let x^* be a solution to the optimization problem of minimizing a continuously differentiable function $f(x)$, subject to $x \in K$, where K is a closed and convex set. Then, x^* is a solution to $VI(\nabla f, K)$ since $\langle \nabla f(x^*), x - x^* \rangle \geq 0,\ \forall x \in K$.*

Theorem 3. *If $f(x)$ is a differentiable convex function, and x^* is the solution of $VI(\nabla f, K)$, then x^* minimizes f.*

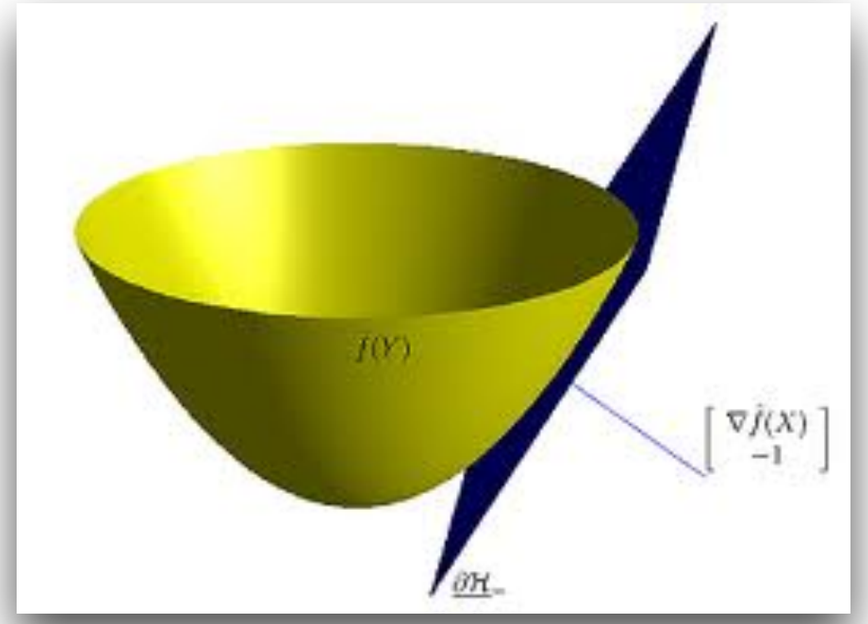

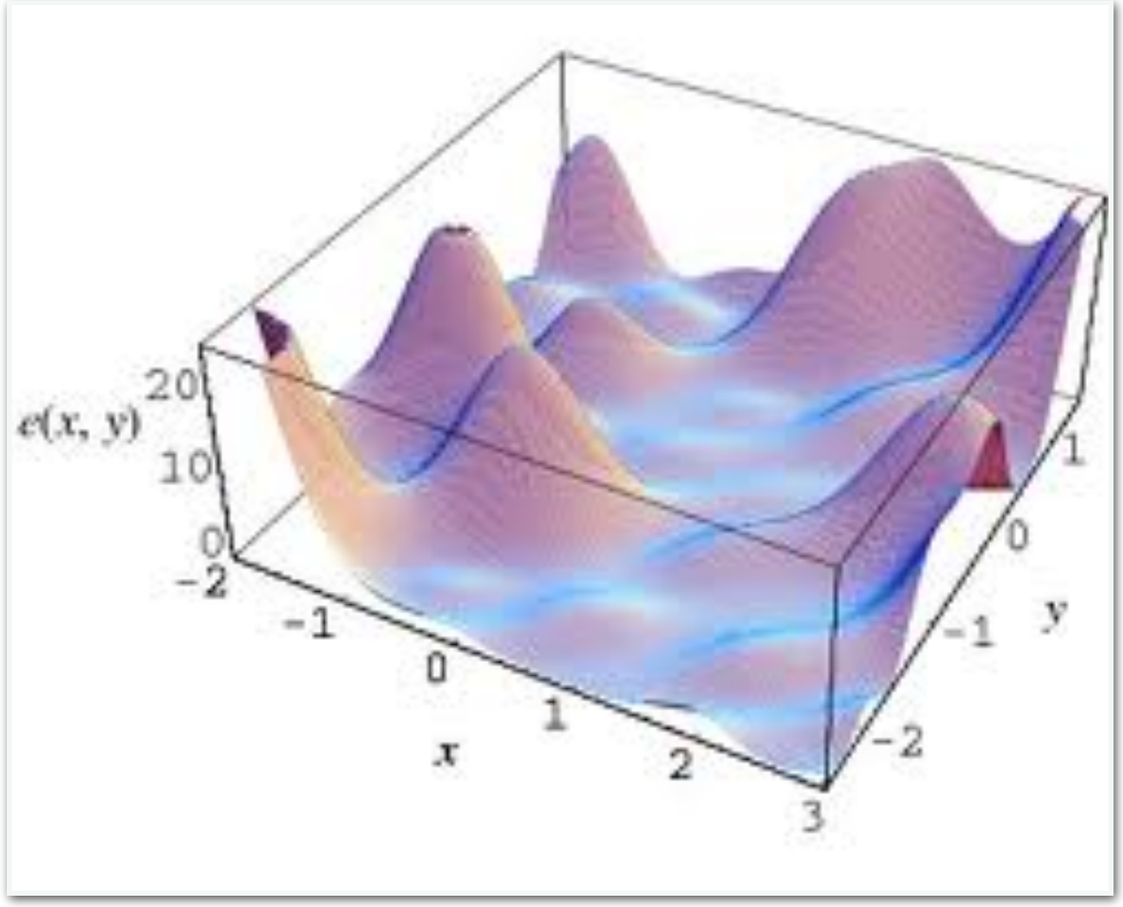

Optimization => VI

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

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

Convex Optimization => VI

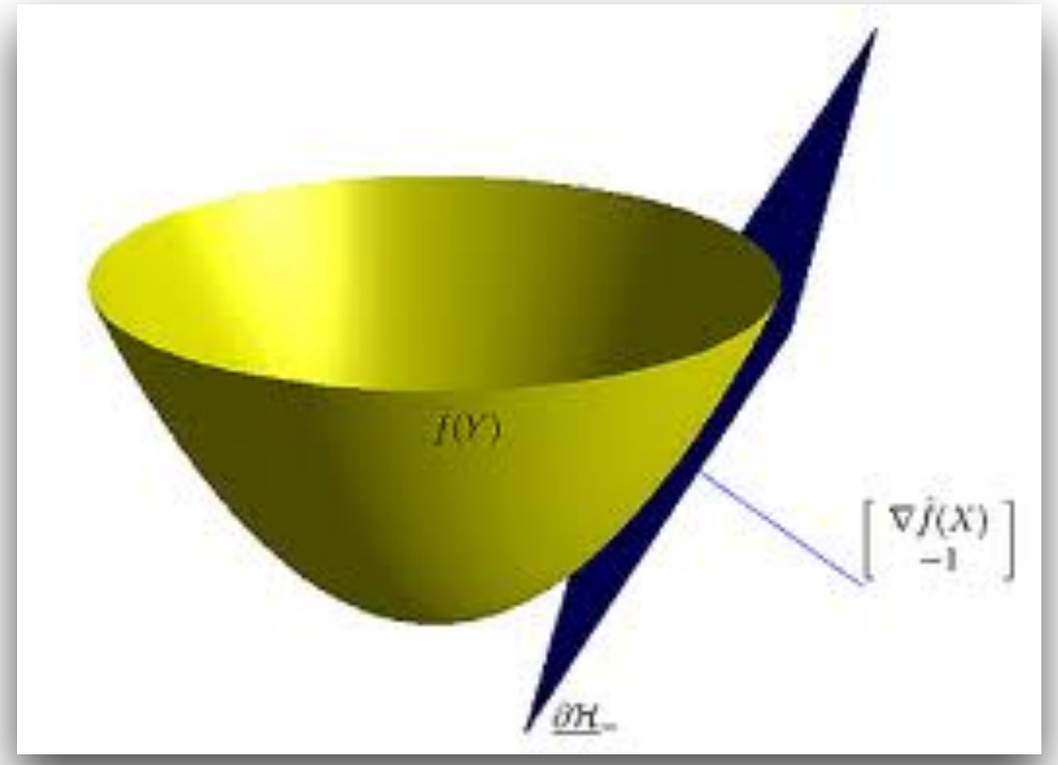

$$f(x) \geq f(x^*) + \langle \nabla f(x^*), x - x^* \rangle, \ \forall x \in K$$

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,

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

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

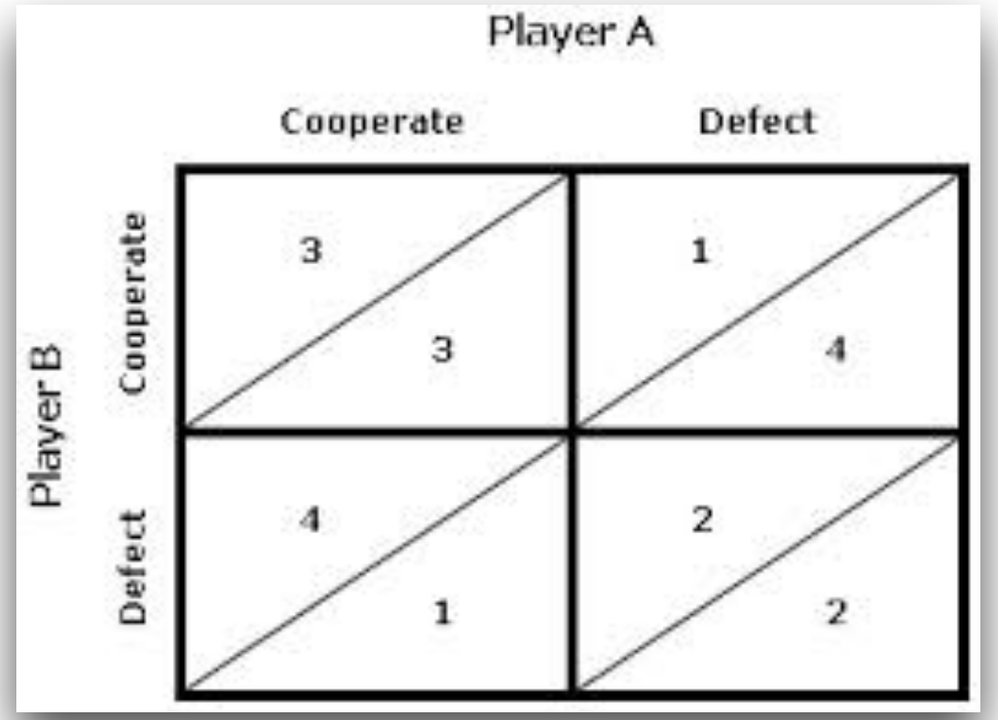

Generalized Nash Equilibrium Problems

- 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

GNEP

- Formally, a GNEP consists of N players, each player i controlling the vector of variables x^i
- Each player's strategy $x^i \in X_i(x^{-i})$
- Here, x^{-i} refers to the strategies of the other players
- Each player i wants to solve the following problem:

$$\text{Minimize}_{x^i} \ \theta_i(x^i, x^{-i}) \ s.t. \ x^i \in X_i(x^{-i})$$

Complementarity Problem

- A variational inequality VI(F,K) over a cone K is a **complementarity problem**
- $F(x) \in K^*$ if $x \in K$
- If x* solves VI(F,K), then F(x*) is perpendicular to x* (nonlinear CP)

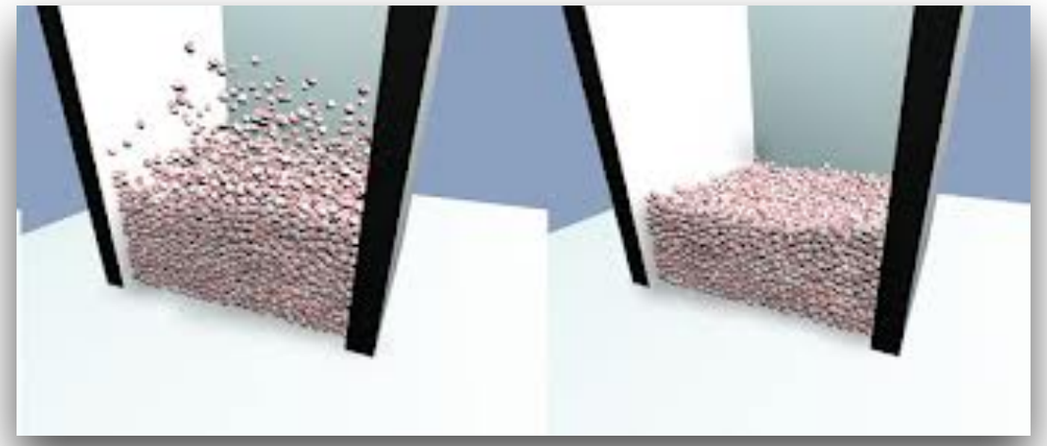

Nonlinear Complementarity Problem Existence

Theorem 3.3 *Let $F : X \to \mathbb{R}^n$ be a continuous function on $X = \mathbb{R}^n_+$. Then the following statements hold:*

(a) If F is monotone on X, then NCP has a convex (possibly empty) solution set.

(b) If F is strictly monotone on X, then NCP has at most one solution.

(c) If F is strongly monotone on X, then NCP has a unique solution.

VIs and Complementarity

the dual cone K^* is defined as $K^* = \{y \in \mathbb{R}^n | \langle y, x \rangle \geq 0, \forall x \in K\}$

Definition 2. *Given a cone $K \subset \mathbb{R}^n$, and a mapping $F : K \to \mathbb{R}^n$, the complementarity problem CP(F,K) is to find an $x \in K$ such that $F(x) \in K^*$, the dual cone to K, and $\langle x, F(x) \rangle \geq 0$.* [3]

Monotonicity Properties

Strongly monotone mapping:

$$\langle F(x) - F(y), x - y \rangle \geq \mu \|x - y\|_2^2, \mu > 0, \forall x, y \in K$$

Lipschitz mapping:

$$\|F(x) - F(y)\|_2 \leq L\|x - y\|_2, \forall x, y \in K$$

Example: if F is the gradient map of a function f, then strong monotonicity of F implies f is strongly convex

Optimization vs VIs

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

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’

Constructing a VI from Competing Objectives

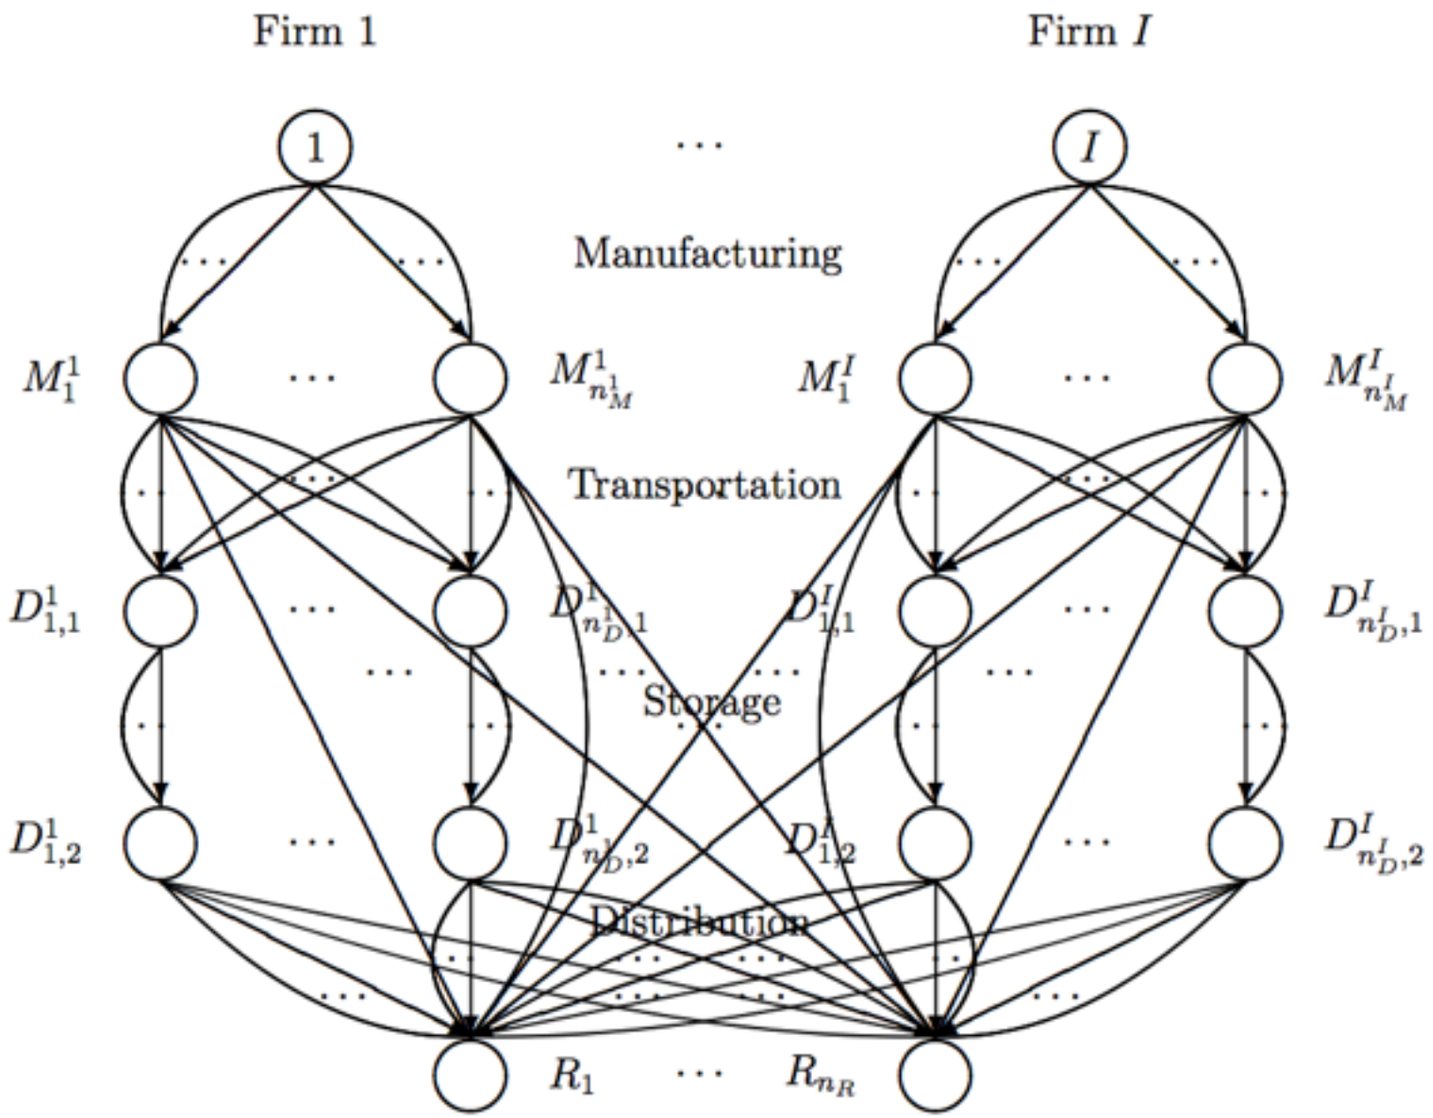

Nagurney, Yu, and Floden 2013

Identify Distinct Objectives

- Each Firm i Attempts to…

Revenue (price*quantity) − Total Operational Cost − Frequency of Operation Cost

- Maximize Profit

$$\text{Maximize} \quad \sum_{k=1}^{n_R} \hat{\rho}_{ik}(x) \sum_{p \in P_k^i} x_p - \sum_{a \in L^i} \hat{c}_a(f) - \sum_{a \in L^i} \hat{g}_a(\gamma_a).$$

- Minimize Emissions

$$\text{Minimize} \quad \sum_{a \in L^i} \hat{e}_a(f_a, \gamma_a).$$

- given the variables it controls - shipping routes (link flows f and path flows x) and frequencies of operation (γ)

Utility Function

Profit – factor*Emissions

$$U_i = \sum_{k=1}^{n_R} \hat{\rho}_{ik}(x) \sum_{p \in P_k^i} x_p - \sum_{a \in L^i} \hat{c}_a(f) - \sum_{a \in L^i} \hat{g}_a(\gamma_a) - \omega_i \sum_{a \in L^i} \hat{e}_a(f_a, \gamma_a).$$

- Relational Constraints such as

Link Flow Definition

$$f_a = \sum_{p \in P} x_p \delta_{ap}, \quad \forall a \in L,$$

and

Link Capacity Constraint

$$f_a \leq \bar{u}_a \gamma_a, \quad \forall a \in L,$$

- can be incorporated into the utility by use of Lagrange multipliers

Forming VI(F,K)

- Step 1: Define X_i as the variables controlled by firm i and concatenate them into a new vector X
  - $X = < X_1, X_2, \ldots, X_n >$
- Step 2: Reformulate firm i's utility for minimization (VI is analogous to minimization problem)
  - Let $D_i = -U_i$
- Step 3: Calculate gradient of D_i with respect to firm i's variables
  - $\nabla_i D_i = - \nabla_i U_i$
- Step 4: Concatenate the gradient of each firms "disutility" function into a new vector mapping, F
  - $F = < \nabla_1 D_1, \nabla_2 D_2, \ldots, \nabla_n D_n >$
- Step 5: Identify the feasible set K. In our case, all product flows, frequencies of operation, and Lagrange multipliers are required to be non-negative.
  - $K = R^+$

Jacobian of F

- Symmetric Jacobian implies Equivalent Optimization Problem
  - Is our Jacobian symmetric?
- Consider the case with only 2 firms, each with its own single variable that it controls (e.g. x_i) and disutility function, D_i

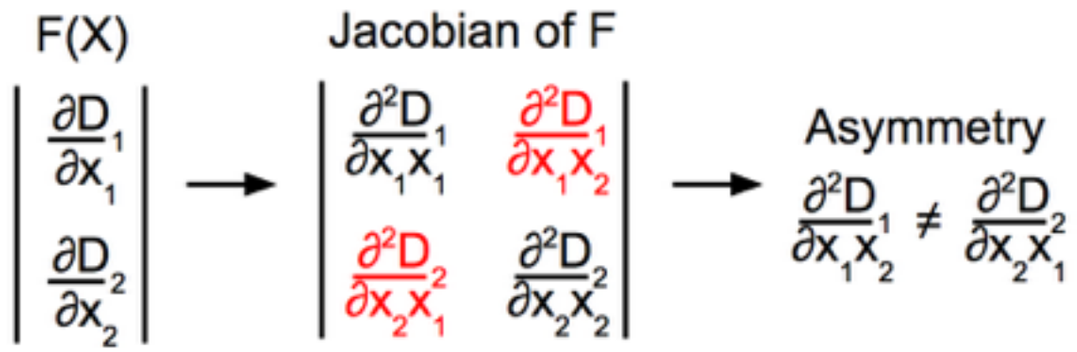

- Although it's possible, there's no reason for firm 1's disutility function to have any relationship to firm 2's disutility
  - e.g. $D_1 = \sin(x_1+x_2)$ and $D_2 = e^{(x1+x2)}$

Procedure

- Identify distinct objective functions with distinct variables
  - e.g. $\min_{x_1} f(x_1,x_2)$ and $\min_{x_2} g(x_1,x_2)$
- Concatenate gradients and variables of these objectives into mapping, F, and vector, X
  - e.g. $F = < \nabla_{x_1} f, \nabla_{x_2} g >$ and $X = < x_1,x_2 >$
- Identify feasible set, K, and update X with Projection Method (analogous to gradient descent)

Linear Complementarity Problem

$$w = Mz + q$$

M is a given NxN matrix
q is a given Nx1 vector
w is non-negative
z is non-negative
w is perpendicular to z

LCPs are a special case of VIs
LCPs can express any quadratic or linear program

Linear Complementarity Problem

- Given a complementarity problem VI(F,K)
- If $F(z) = Mz + q$, then the VI (F,K) is a **linear complementarity problem** (LCP)
- **Interestingly, LCPs have no objective function!**
- LCPs are powerful enough to represent any linear or quadratic programming problem
- LCP(M,q) means
  - Find a $z >= 0$ s.t. $Mz + q >= 0$ and
  - $(Mz + q)^T z = 0$

Sparse Regression as a VI

(Kim and Park, AI Stats 2010)

"Sparse" Supervised learning

$$\min_{x \in X} f(x) + g(x) : \min_{\beta \in \mathbb{R}^k} \|X\beta - y\|_2^2 + \lambda \|\beta\|_1$$

VI (LCP) Formulation

$$\begin{aligned} \beta &= (X^T X)^{-1} X^T y + (X^T X)^{-1})(-c) \\ &= q + Mz \\ \text{where } M &= (X^T X)^{-1} \\ q &= (X^T X)^{-1} X^T y \end{aligned}$$

Linear Complementarity for Regularized Policy Evaluation and Improvement

NIPS 2011



Jeff Johns **Christopher Painter-Wakefield** **Ronald Parr**
Department of Computer Science
Duke University
Durham, NC 27708
{johns, paint007, parr}@cs.duke.edu

Abstract

Recent work in reinforcement learning has emphasized the power of L_1 regularization to perform feature selection and prevent overfitting. We propose formulating the L_1 regularized linear fixed point problem as a linear complementarity problem (LCP). This formulation offers several advantages over the LARS-inspired formulation, LARS-TD. The LCP formulation allows the use of efficient off-the-shelf solvers, leads to a new uniqueness result, and can be initialized with starting points from similar problems (warm starts). We demonstrate that warm starts, as well as the efficiency of LCP solvers, can speed up policy iteration. Moreover, warm starts permit a form of modified policy iteration that can be used to approximate a "greedy" homotopy path, a generalization of the LARS-TD homotopy path that combines policy evaluation and optimization.

L1 Regularized RL

$$w = \arg\min_{x \in \mathbb{R}^k} \frac{1}{2} \|\Phi x - \max_{\pi}(R + \gamma \Phi'^{\pi} w)\|_2^2 + \beta \|x\|_1.$$

$A = \Phi^T (\Phi - \gamma \Phi'^{\pi})$ and $b = \Phi^T R$, the LCP is written:

$$\begin{bmatrix} x^+ \\ x^- \end{bmatrix} = \left(\beta - \begin{bmatrix} b \\ -b \end{bmatrix} \right) + \begin{bmatrix} A & -A \\ -A & A \end{bmatrix} \begin{bmatrix} w^+ \\ w^- \end{bmatrix}$$

$$\begin{bmatrix} x^+ \\ x^- \end{bmatrix}, \begin{bmatrix} w^+ \\ w^- \end{bmatrix} \geq \mathbf{0}$$

$$x_i^+ w_i^+ = 0 \quad \forall i$$

$$x_i^- w_i^- = 0 \quad \forall i$$

Projections and VIs

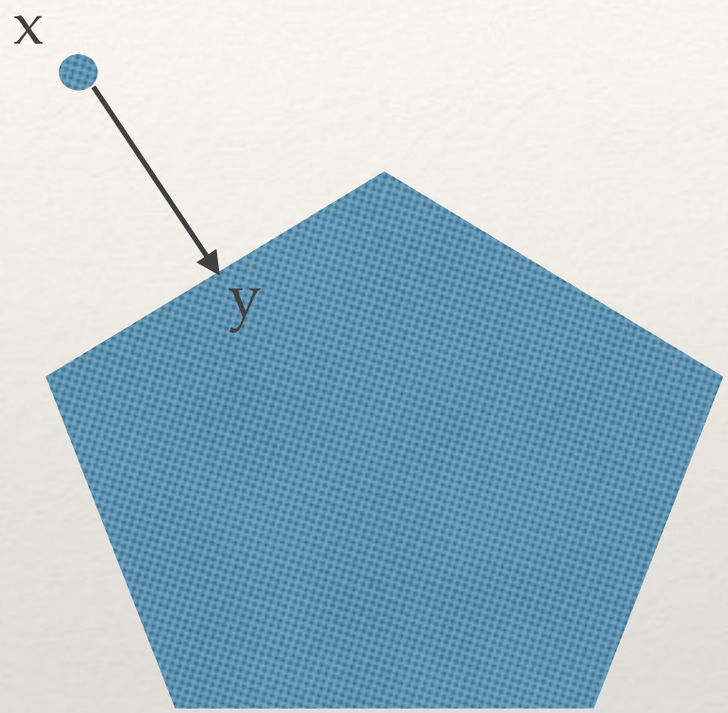

$$\text{if } y = \Pi_{\hat{S}}(x), \text{ then } \ (y-x)^T(z-y) \geq 0, z \in \hat{S}$$

Projected Dynamical Systems

Two Player Game

Row \ Column	Heads	Tails
Heads	(1, –1)	(–1, 1)
Tails	(–1, 1)	(1, –1)

Classical Dynamical System

$$\alpha_{t+1} = \alpha_t + \eta_t \frac{\partial V_r(\alpha_t, \beta_t)}{\partial \alpha}$$

$$\beta_{t+1} = \beta_t + \eta_t \frac{\partial V_c(\alpha_t, \beta_t)}{\partial \beta}$$

Projected Dynamical System

$$\dot{x} = \Pi_K(x, -F(x)), x(0) = x_0 \in K$$

Projected dynamical systems are a powerful framework for studying dynamics of equilibria in games

Contrast with ODE used in [Singh et al., UAI 2000]

PDS Formulation

ODE/IVP

$$\dot{X} = \Pi_{\mathcal{K}}(X, -F(X)), \quad X(0) = X^0,$$

Projection Operator

$$\Pi_{\mathcal{K}}(X, -F(X)) = \lim_{\delta \to 0} \frac{P_{\mathcal{K}}(X - \delta F(X)) - X}{\delta}$$

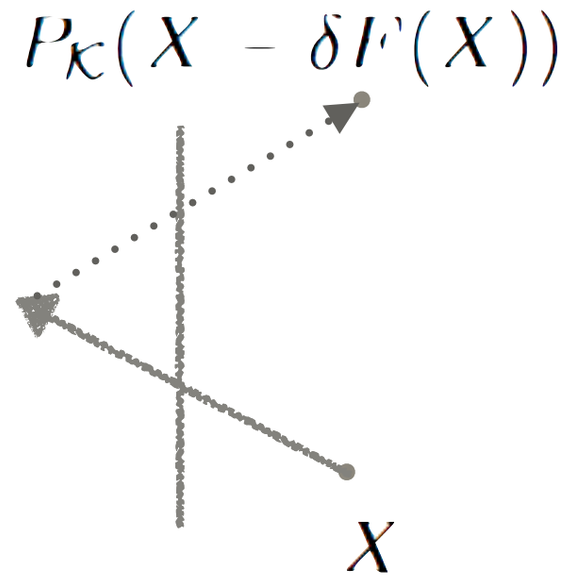

Fixed Point Problem

$$\dot{X} = 0 = \Pi_{\mathcal{K}}(X^*, -F(X^*))$$

- X^* solves the VI iff it is a stationary point of the projected ODE
- Lipschitz continuity of F(X) guarantees the existence of a unique solution
- Stability of equilibrium is given by the monotonicity of F(X) which can be determined from the positive-definiteness of the Jacobian of F(X)

Skorokhod Analysis

$$\dot{\phi}_x(t) = \Pi_K(\phi_x(t), -F(\phi_x(t))), \ \phi_x(0) = x$$

$$F_1(x_1, x_2) = -x_2, F_2(x_1, x_2) = 4x_1$$

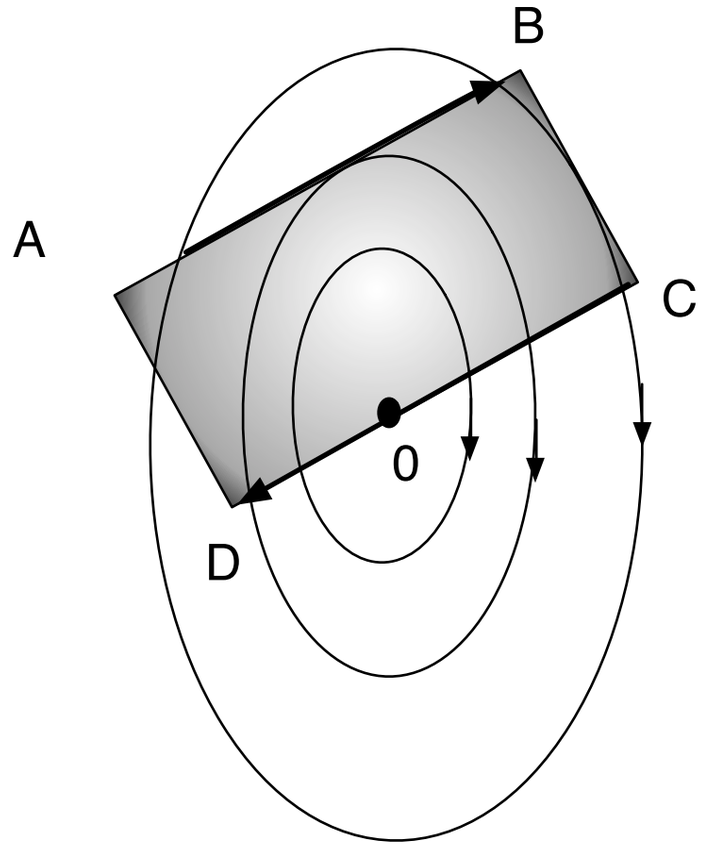

Questions?



Plan of the Tutorial

- Overview: 2:00-2:10 p.m.
- Part I: Motivation: 2:10-2:30 p.m.
- Part II: Mathematical Background: 2:30-3:30 p.m.
- COFFEE BREAK: 3:30-4:00 p.m.
- **Part III: Algorithms: 4:00-5:00 p.m.**
- Part IV: Applications and Future Work: 5:00-6:00 p.m.

Projection Algorithm

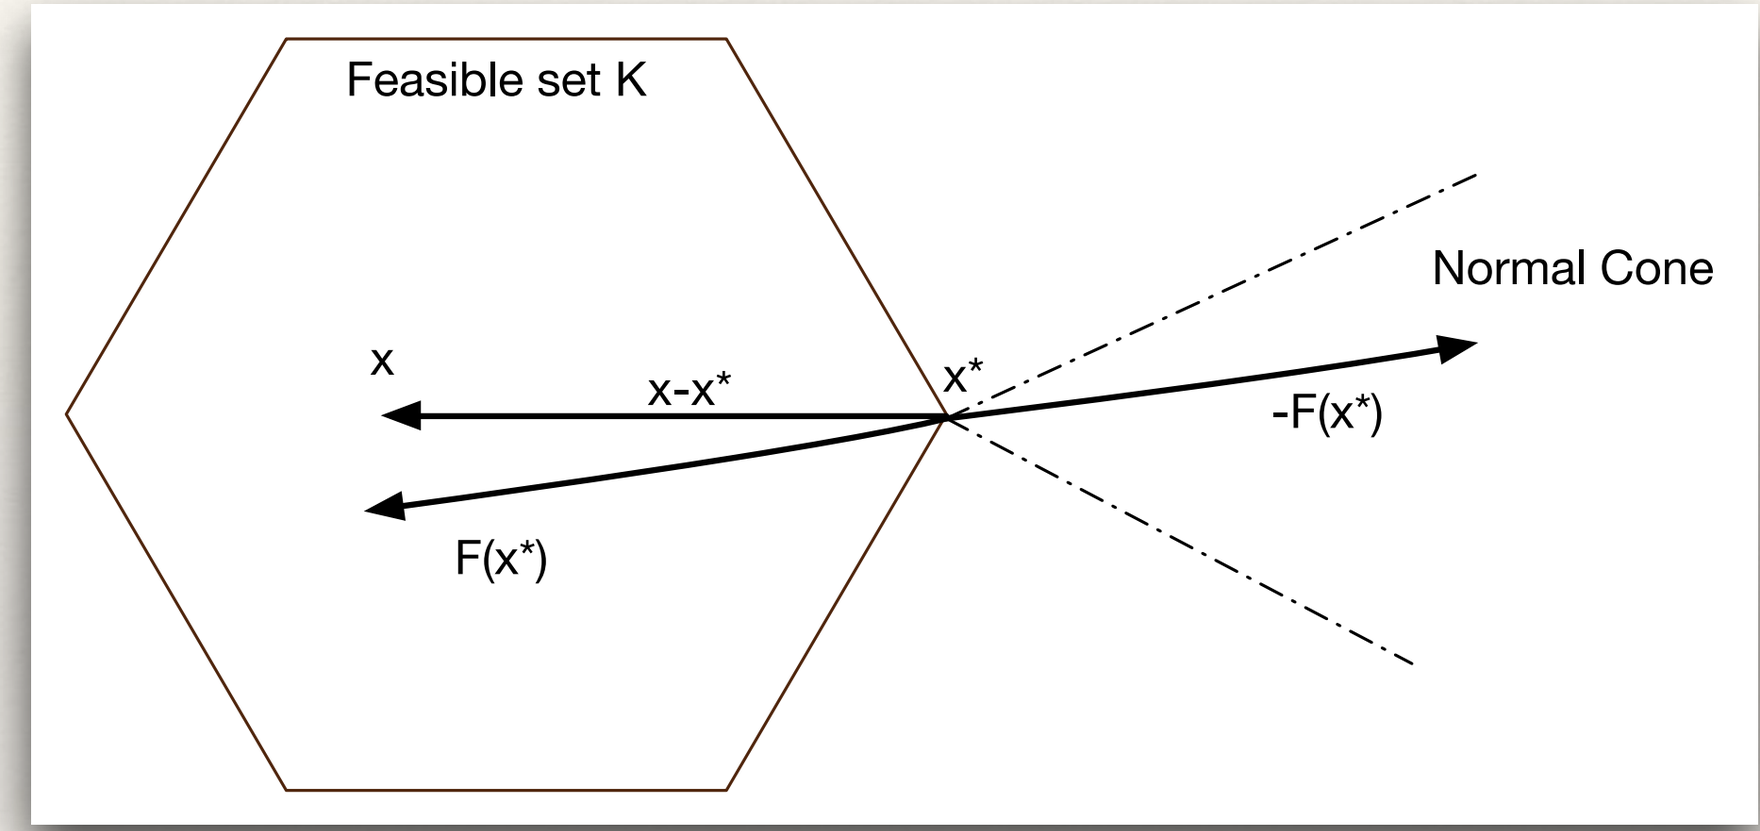

Projection Algorithm

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

Projection Method Fails

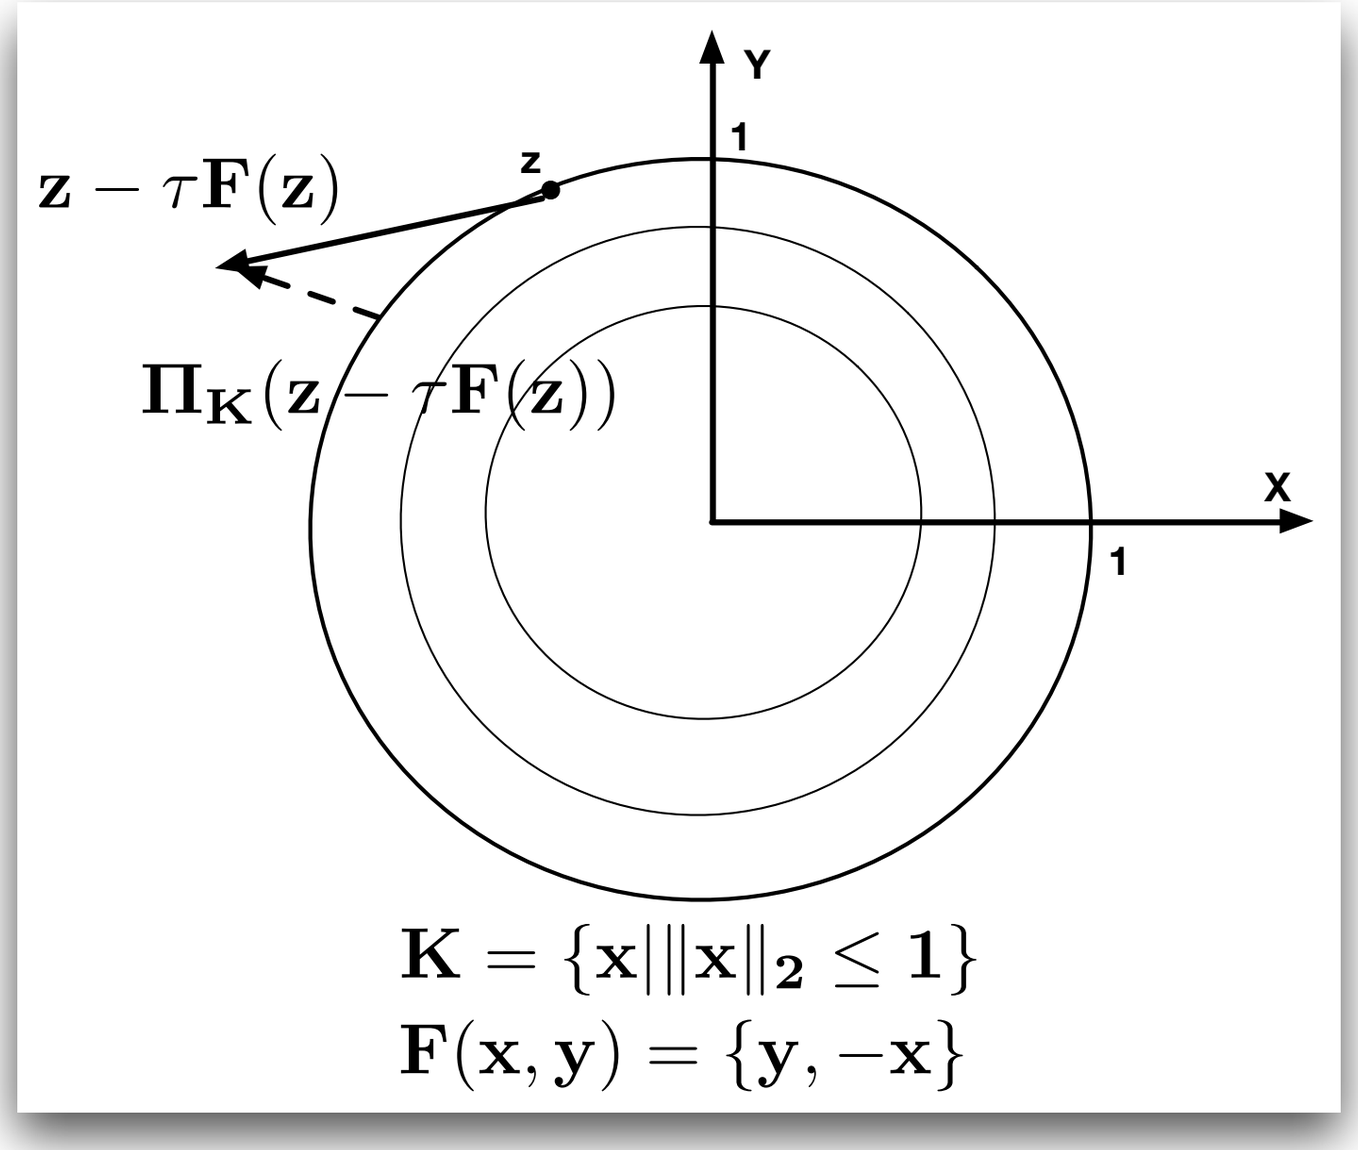

Bertsekas and Tsitsiklis, Parallel and Distributed Computation, Athena Scientific.

Extragradient Method for solving VIs

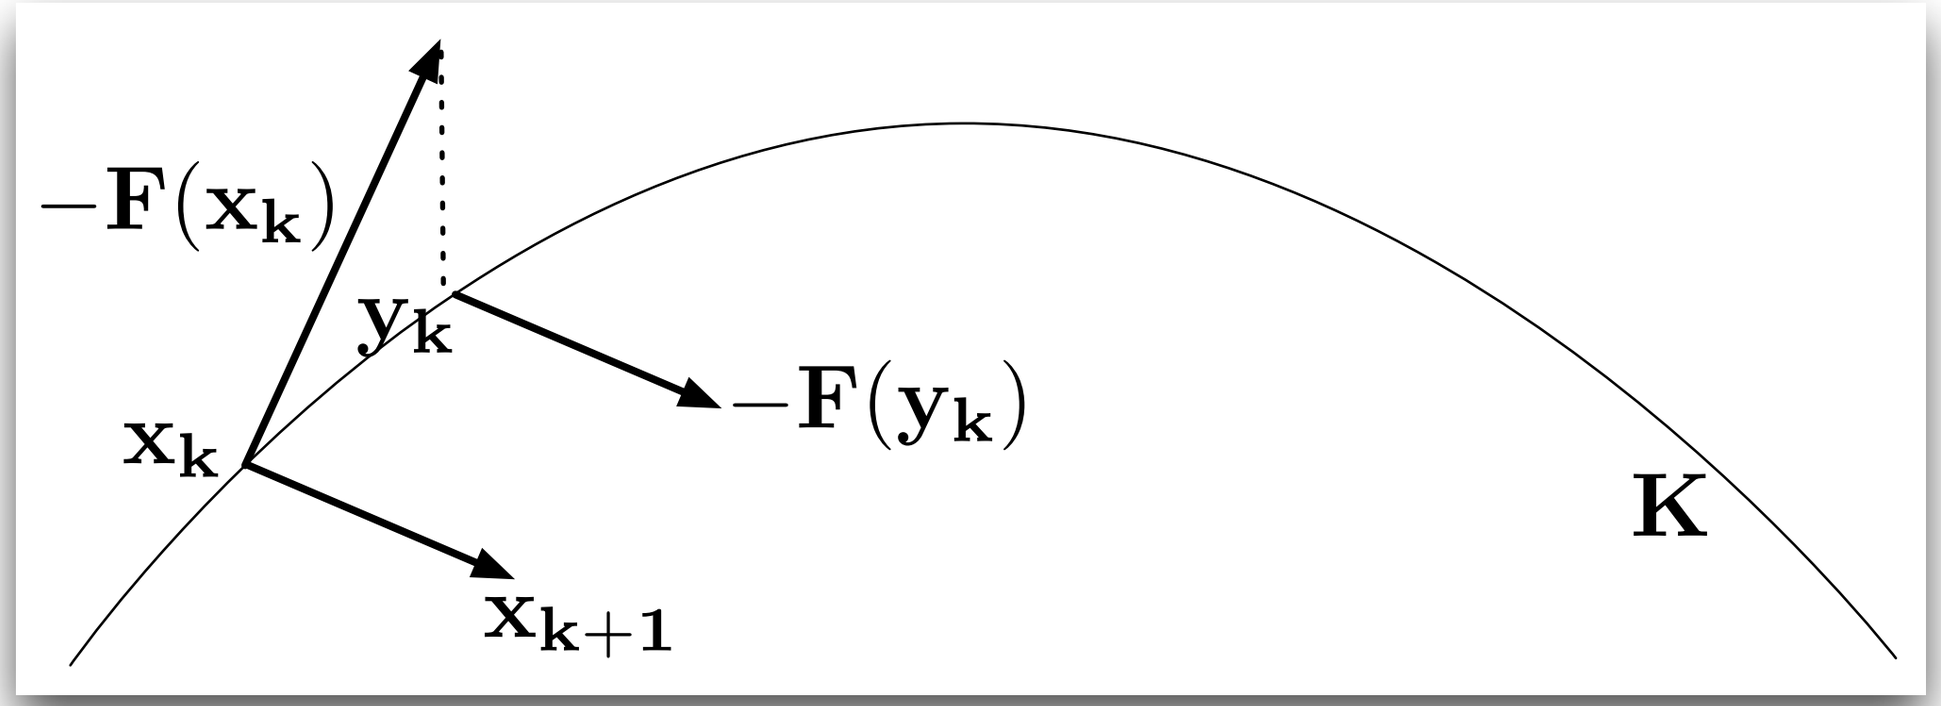

Korpolevich (1970s) developed the extragradient method

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

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

Designing New ML Algorithms using VIs

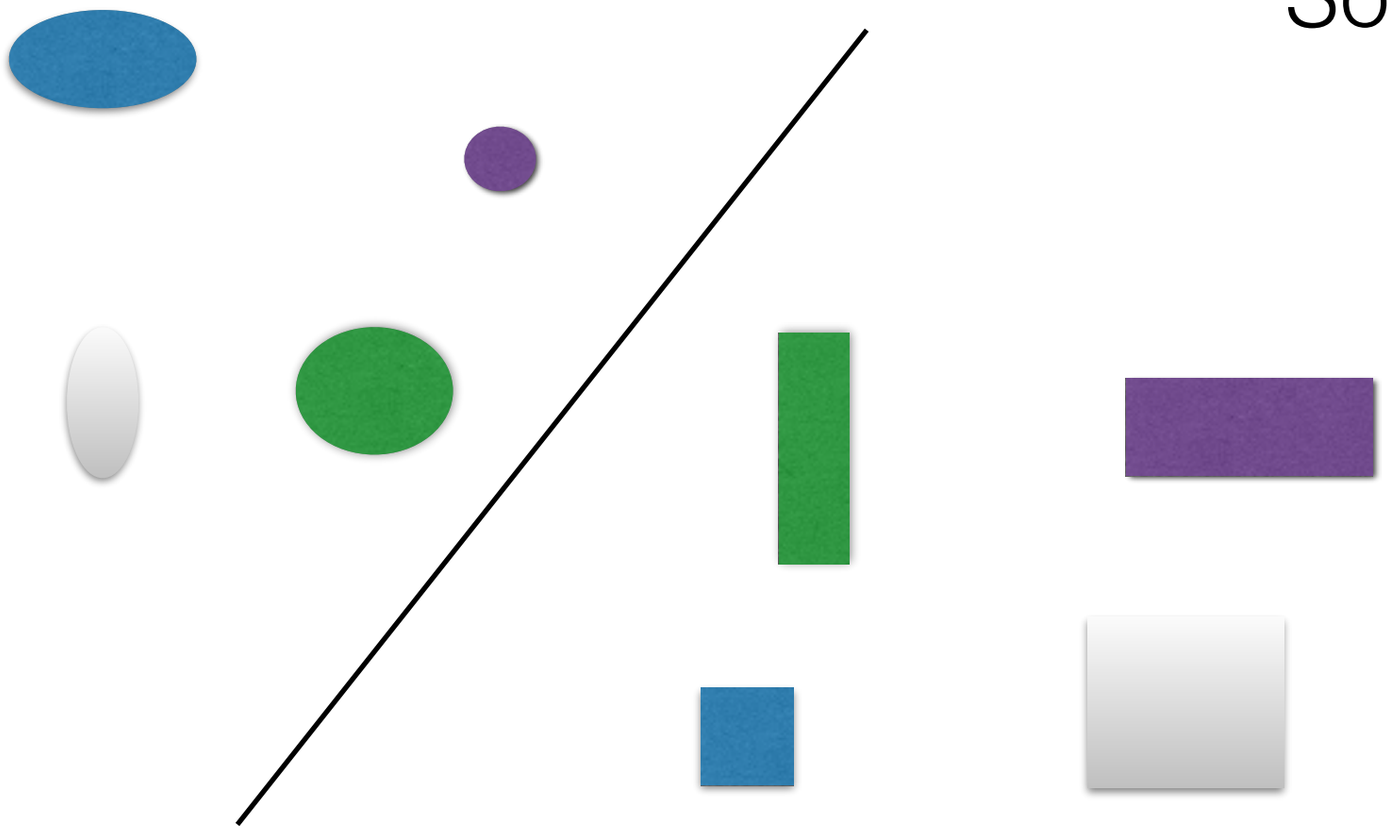

Soft-margin SVM

$$\begin{cases} \min \dfrac{1}{2}\|w\|^2 + C\sum_{i=1}^{N} \xi_i \\ \text{s.t. } y_i(w^T x_i + b) \geq 1 - \xi_i \\ \qquad \xi_i \geq 0,\ i = 1, 2, \cdots, N \end{cases}$$

Convert Convex Optimization into VI

Consider the convex constrained optimization problem:

$$\begin{cases} \min \sum_{i=1}^{m} f_i(x_i) \\ s.t. a_j^T X \le b_j, j = 1,2,\cdots,r, X_i \in K_i, i = 1,2,\cdots,m \end{cases}$$

Lemma 1 Optimization problem (9) is equal to the following variational inequality:

$$\sum_{i=1}^{m} \frac{\partial L(X^*)}{\partial X_i} \times (X_i - X_i^*) + \sum_{j=1}^{r} (b_j - a_j^T X^*) \times (u_j - u_j^*) \ge 0$$

where $L(X) = \sum_{i=1}^{m} f_i(X_i) + \sum_{j=1}^{r} u_j (a_j^T X - b)$ is a Lagrange function and $u_j (j = 1,2,\cdots,r)$ is a Lagrange multiplier.

Xie, Zhao, and Miao, IEEE CIIA, 2008

Stochastic EG Soft-Margin SVM

(Mahadevan, 2014)

Pick a random feature: $$w_i^{t+1} \leftarrow (1-\alpha_t) w_i^t + \alpha_t \sum_j \lambda_j y_j x_{ij}$$

Primal update

Pick a random example: $$\lambda_j^{t+1} \leftarrow \lambda_j^t - \alpha_t (\sum_i w_i^t x_{ij} + b^t) - 1 + \xi_j^t$$

Dual update

Repeat primal and dual update in extragradient step

Stochastic VIs

- Extending VIs to the stochastic setting
- Assume noisy samples of the true F mapping
  - $F(x) = E(F(x, \xi))$ where $\xi:\Omega \rightarrow R^d$, $F: (X, R^d) \rightarrow R^n$ and (Ω,F,P) is the associated probability space
- SVI: find x^* such that $(x - x^*)^T E(F(x^*,\xi(\omega))) >= 0, x \varepsilon X$

Stochastic Projection Algorithm

(Jiang and Xu, IEEE Trans. on Automatic Control, 2008)

- Jiang and Xu extended the projection algorithm for VIs to the stochastic VI setting
- $x_{k+1} \leftarrow\!\!- \Pi_X [x_k - \alpha_k (F(x_k) + w_k)]$
- where $w_k = F(x_x,\omega_k) - F(x_k)$
- Conditional mean $E(w_k \mid F_k) = 0$ for every k
- Conditional variance $E(|w_k{}^2| \mid F)$

Stochastic Extragradient Method

(Kannan and Shanbhag, Amer. Control Conf., 2014)

- This year, Kannan and Shanbhag developed a stochastic extragradient method ESA
  - $x_{k+1/2} \leftarrow\!\!- \Pi_X [x_k - \alpha_k (F(x_k) + w_k)]$
  - $x_{k+1} \leftarrow\!\!- \Pi_X [x_k - \alpha_k (F(x_{k+1/2}) + w_{k+1/2})]$
- Assuming standard stochastic approximation conditions (square summable learning rates), they showed convergence of ESA

Affine mappings

$$T(x) = Ax + b$$

$$T^{\pi}(V) = R + \gamma PV$$

$$\Phi^T F(\Phi x) = \Phi^T \Xi F(\Phi x) = Cr - d$$

$$\text{where } C = \Phi^T \Xi (I - A)\Phi, \quad d = \Phi^T \Xi b$$

Constructing Mappings from Samples

How to compute T if A is unknown?

$$T(x) = Ax + b$$

$$C_k = \frac{1}{k+1} \sum_{t=0}^{k} \phi(i_t)(\phi(i_t) - \frac{a_{i_t j_t}}{p_{i_t j_t}} \phi(j_t))^T$$

$$d_k = \frac{1}{k+1} \sum_{t=0}^{k} \phi(i_t) b_{i_t}$$

Extragradient Method for Affine Mappings

$$x_{k+1} = x_k - \gamma(C(x_k - \gamma(Cx_k - d)) - d))$$

Theorem 7. *Let C_k and d_k converge to C and d under the given sampling scheme. Then, the linear simulation-based extragradient iteration in Equation 6 converges to a solution of the VI $F(\Phi r^*)^T \Phi(r - r^*) \geq 0,\ \forall r \in \hat{R}$, assuming C_k and d_k are used in place of C and d, and assuming Φ is of full rank.*

Generalized Stochastic Extragradient TD Method

(Mahadevan, 2014)

$$\delta_{k+\frac{1}{2}} = \phi(i_k)^T x_k - \frac{a_{i_k,j_k}}{p_{i_k j_k}} \phi(j_k)^T x_k - b_{i_k} \quad (7)$$

$$x_{k+\frac{1}{2}} = x_k - \gamma_k \phi(i_k) \delta_{k+\frac{1}{2}} \quad (8)$$

$$\delta_k = \phi(i_k)^T x_{k+\frac{1}{2}} - \frac{a_{i_k,j_k}}{p_{i_k j_k}} \phi(j_k)^T x_{k+\frac{1}{2}} - b_{i_k} \quad (9)$$

$$x_{k+1} = x_k - \gamma_k \phi(i_k) \delta_k \quad (10)$$

Extragradient TD Method for RL

(Mahadevan, 2014)

Algorithm 2 The Extragradient TD(λ) algorithm.

1: Set $t = 0$ and the weights $w_t = 0$
2: **repeat**
3: In state s_t, do action $\pi(s_t)$, and observe next state s_{t+1} and receive reward r_t.
4: Update the eligibility trace $e_{t+1} \leftarrow \lambda\gamma e_t + \phi(s_t)$
5: Compute $\delta_{t+\frac{1}{2}} = r_t + \gamma\phi(s_{t+1})^T w_t - \phi(s_t)^T w_t$
6: Update the weights $w_{t+\frac{1}{2}} = w_t + \alpha_t \delta_{t+\frac{1}{2}} e_{t+1}$
7: Set $\delta_t = r_t + \phi(s_{t+1})^T w_{t+\frac{1}{2}} - \phi(s_t)^T w_{t+\frac{1}{2}}$.
8: Update weights: $w_{t+1} = w_t + \alpha_t \delta_t e_{t+1}$
9: **until convergence**
10: Return w_t

$0 \in T(x)$

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)

Proximal Mapping

Moreau, 1965:

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

Proximal operator generalizes projection:

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

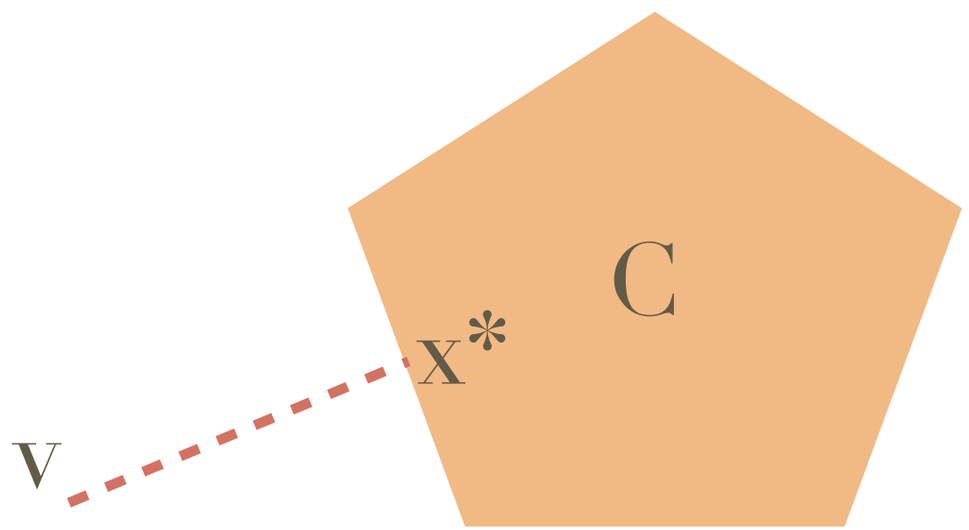

Proximal Abstraction

Proximal operators

Table 10.2 Proximity operator of $\phi \in \Gamma_0(\mathbb{R})$; $\alpha \in \mathbb{R}$, $\kappa > 0$, $\underline{\kappa} > 0$, $\overline{\kappa} > 0$, $\omega > 0$, $\underline{\omega} < \overline{\omega}$, $q > 1$, $\tau \geq 0$ [37,53,55].

	$\phi(x)$	$\operatorname{prox}_{\phi} x$
i	$\iota_{[\underline{\omega},\overline{\omega}]}(x)$	$P_{[\underline{\omega},\overline{\omega}]}\, x$
ii	$\sigma_{[\underline{\omega},\overline{\omega}]}(x) = \begin{cases} \underline{\omega} x & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ \overline{\omega} x & \text{otherwise} \end{cases}$	$\operatorname{soft}_{[\underline{\omega},\overline{\omega}]}(x) = \begin{cases} x - \underline{\omega} & \text{if } x < \underline{\omega} \\ 0 & \text{if } x \in [\underline{\omega},\overline{\omega}] \\ x - \overline{\omega} & \text{if } x > \overline{\omega} \end{cases}$
iii	$\psi(x) + \sigma_{[\underline{\omega},\overline{\omega}]}(x)$ $\psi \in \Gamma_0(\mathbb{R})$ differentiable at 0 $\psi'(0) = 0$	$\operatorname{prox}_{\psi}\big(\operatorname{soft}_{[\underline{\omega},\overline{\omega}]}(x)\big)$
iv	$\max\{\lvert x\rvert - \omega, 0\}$	$\begin{cases} x & \text{if } \lvert x\rvert < \omega \\ \operatorname{sign}(x)\omega & \text{if } \omega \leq \lvert x\rvert \leq 2\omega \\ \operatorname{sign}(x)(\lvert x\rvert - \omega) & \text{if } \lvert x\rvert > 2\omega \end{cases}$
v	$\kappa \lvert x\rvert^q$	$\operatorname{sign}(x)p$, where $p \geq 0$ and $p + q\kappa p^{q-1} = \lvert x\rvert$
vi	$\begin{cases} \kappa x^2 & \text{if } \lvert x\rvert \leq \omega/\sqrt{2\kappa} \\ \omega\sqrt{2\kappa}\lvert x\rvert - \omega^2/2 & \text{otherwise} \end{cases}$	$\begin{cases} x/(2\kappa+1) & \text{if } \lvert x\rvert \leq \omega(2\kappa+1)/\sqrt{2\kappa} \\ x - \omega\sqrt{2\kappa}\operatorname{sign}(x) & \text{otherwise} \end{cases}$
vii	$\omega\lvert x\rvert + \tau\lvert x\rvert^2 + \kappa\lvert x\rvert^q$	$\operatorname{sign}(x)\operatorname{prox}_{\kappa\lvert\cdot\rvert^q/(2\tau+1)}\dfrac{\max\{\lvert x\rvert - \omega, 0\}}{2\tau+1}$
viii	$\omega\lvert x\rvert - \ln(1+\omega\lvert x\rvert)$	$(2\omega)^{-1}\operatorname{sign}(x)\Big(\omega\lvert x\rvert - \omega^2 - 1 + \sqrt{\lvert\omega\lvert x\rvert - \omega^2 - 1\rvert^2 + 4\omega\lvert x\rvert}\Big)$
ix	$\begin{cases} \omega x & \text{if } x \geq 0 \\ +\infty & \text{otherwise} \end{cases}$	$\begin{cases} x - \omega & \text{if } x \geq \omega \\ 0 & \text{otherwise} \end{cases}$
x	$\begin{cases} -\omega x^{1/q} & \text{if } x \geq 0 \\ +\infty & \text{otherwise} \end{cases}$	$p^{1/q}$, where $p > 0$ and $p^{2q-1} - xp^{q-1} = q^{-1}\omega$
xi	$\begin{cases} \omega x^{-q} & \text{if } x > 0 \\ +\infty & \text{otherwise} \end{cases}$	$p > 0$ such that $p^{q+2} - xp^{q+1} = \omega q$
xii	$\begin{cases} x\ln(x) & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ +\infty & \text{otherwise} \end{cases}$	$W(e^{x-1})$, where W is the Lambert W-function
xiii	$\begin{cases} -\ln(x-\underline{\omega}) + \ln(-\underline{\omega}) & \text{if } x \in]\underline{\omega}, 0] \\ -\ln(\overline{\omega} - x) + \ln(\overline{\omega}) & \text{if } x \in]0, \overline{\omega}[\\ +\infty & \text{otherwise} \end{cases}$ $\underline{\omega} < 0 < \overline{\omega}$	$\begin{cases} \dfrac{1}{2}\Big(x + \underline{\omega} + \sqrt{\lvert x - \underline{\omega}\rvert^2 + 4}\Big) & \text{if } x < 1/\underline{\omega} \\ \dfrac{1}{2}\Big(x + \overline{\omega} - \sqrt{\lvert x - \overline{\omega}\rvert^2 + 4}\Big) & \text{if } x > 1/\overline{\omega} \\ 0 & \text{otherwise} \end{cases}$ (see Figure 10.1)
xiv	$\begin{cases} -\kappa\ln(x) + \tau x^2/2 + \alpha x & \text{if } x > 0 \\ +\infty & \text{otherwise} \end{cases}$	$\dfrac{1}{2(1+\tau)}\Big(x - \alpha + \sqrt{\lvert x - \alpha\rvert^2 + 4\kappa(1+\tau)}\Big)$
xv	$\begin{cases} -\kappa\ln(x) + \alpha x + \omega x^{-1} & \text{if } x > 0 \\ +\infty & \text{otherwise} \end{cases}$	$p > 0$ such that $p^3 + (\alpha - x)p^2 - \kappa p = \omega$
xvi	$\begin{cases} -\kappa\ln(x) + \omega x^q & \text{if } x > 0 \\ +\infty & \text{otherwise} \end{cases}$	$p > 0$ such that $q\omega p^q + p^2 - xp = \kappa$
xvii	$\begin{cases} -\underline{\kappa}\ln(x - \underline{\omega}) - \overline{\kappa}\ln(\overline{\omega} - x) & \text{if } x \in]\underline{\omega}, \overline{\omega}[\\ +\infty & \text{otherwise} \end{cases}$	$p \in]\underline{\omega}, \overline{\omega}[$ such that $p^3 - (\underline{\omega} + \overline{\omega} + x)p^2 +$ $\big(\underline{\omega}\overline{\omega} - \underline{\kappa} - \overline{\kappa} + (\underline{\omega} + \overline{\omega})x\big)p = \underline{\omega}\overline{\omega}x - \underline{\omega}\overline{\kappa} - \overline{\omega}\underline{\kappa}$

Low Rank Matrix Regularization via Schatten p-norm

Schatten p-norm of $A_{m\times n} = \|\sigma_s(A)\|_p = (\sum_{i=1}^{n} |\sigma_i|^p)^{\frac{1}{p}}$

Proximal Operator of nuclear norm (Schatten 1-norm):

$$\text{prox}_{\lambda F}(A) = \sum_{i=1}^{n} (\sigma_i - \lambda)_+ u_i v_i^T$$

Conjugate Functions



Moreau Decomposition:

$$v = \text{prox}_f(v) + \text{prox}_{f*}(v)$$

$$\text{where } f^*(y) = \sup_x (y^T x - f(x))$$

$$f(x) = \langle a, x \rangle - b,\ a \in \mathbb{R}^n, b \in \mathbb{R}$$

$$f^*(x^*) = \begin{cases} b, & x^* = a \\ +\infty, & x^* \neq a. \end{cases}$$

Proximal Abstractions

	Property	$f(x)$	$\operatorname{prox}_f x$
i	translation	$\varphi(x-z),\ z\in\mathbb{R}^N$	$z+\operatorname{prox}_\varphi(x-z)$
ii	scaling	$\varphi(x/\rho),\ \rho\in\mathbb{R}\smallsetminus\{0\}$	$\rho\operatorname{prox}_{\varphi/\rho^2}(x/\rho)$
iii	reflection	$\varphi(-x)$	$-\operatorname{prox}_\varphi(-x)$
iv	quadratic perturbation	$\varphi(x)+\alpha\|x\|^2/2+u^\top x+\gamma$ $u\in\mathbb{R}^N,\ \alpha\geq 0,\ \gamma\in\mathbb{R}$	$\operatorname{prox}_{\varphi/(\alpha+1)}\big((x-u)/(\alpha+1)\big)$
v	conjugation	$\varphi^*(x)$	$x-\operatorname{prox}_\varphi x$
vi	squared distance	$\frac{1}{2}d_C^2(x)$	$\frac{1}{2}(x+P_C x)$
vii	Moreau envelope	$\widetilde{\varphi}(x)=\inf_{y\in\mathbb{R}^N}\varphi(y)+\frac{1}{2}\|x-y\|^2$	$\frac{1}{2}\big(x+\operatorname{prox}_{2\varphi}x\big)$
viii	Moreau complement	$\frac{1}{2}\|\cdot\|^2-\widetilde{\varphi}(x)$	$x-\operatorname{prox}_{\varphi/2}(x/2)$
ix	decomposition in an orthonormal basis $(b_k)_{1\leq k\leq N}$	$\sum_{k=1}^N\phi_k(x^\top b_k)$ $\phi_k\in\Gamma_0(\mathbb{R})$	$\sum_{k=1}^N\operatorname{prox}_{\phi_k}(x^\top b_k)b_k$
x	semi-orthogonal linear transform	$\varphi(Lx)$ $L\in\mathbb{R}^{M\times N},\ LL^\top=\nu I,\ \nu>0$	$x+\nu^{-1}L^\top\big(\operatorname{prox}_{\nu\varphi}(Lx)-Lx\big)$
xi	quadratic function	$\gamma\|Lx-y\|^2/2$ $L\in\mathbb{R}^{M\times N},\ \gamma>0,\ y\in\mathbb{R}^M$	$(I+\gamma L^\top L)^{-1}(x+\gamma L^\top y)$
xii	indicator function	$\iota_C(x)=\begin{cases}0 & \text{if } x\in C\\ +\infty & \text{otherwise}\end{cases}$	$P_C x$
xiii	distance function	$\gamma d_C(x),\ \gamma>0$	$\begin{cases}x+\gamma(P_C x-x)/d_C(x) & \\ \qquad\text{if } d_C(x)>\gamma & \\ P_C x\quad\text{otherwise} & \end{cases}$
xv	function of distance	$\phi(d_C(x))$ $\phi\in\Gamma_0(\mathbb{R})$ even, differentiable at 0 with $\phi'(0)=0$	$\begin{cases}x+\left(1-\dfrac{\operatorname{prox}_\phi d_C(x)}{d_C(x)}\right)(P_C x-x) & \\ \qquad\text{if } x\notin C & \\ x\quad\text{otherwise} & \end{cases}$
xv	support function	$\sigma_C(x)$	$x-P_C x$
xvii	thresholding	$\sigma_C(x)+\phi(\|x\|)$ $\phi\in\Gamma_0(\mathbb{R})$ even and not constant	$\begin{cases}\dfrac{\operatorname{prox}_\phi d_C(x)}{d_C(x)}(x-P_C x) & \\ \qquad\text{if } d_C(x)>\max\operatorname{Argmin}\phi & \\ x-P_C x\quad\text{otherwise} & \end{cases}$

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

Proximal Point Algorithm

$$0 \in F(x) \;\leftrightarrow 0 \in \alpha F(x) \leftrightarrow x \in (I + \alpha F)(x) \leftrightarrow x = (I + \alpha F)^{-1}(x)$$

$$x_{k+1} \leftarrow (I + \alpha_k \partial f)^{-1}(x_k)$$

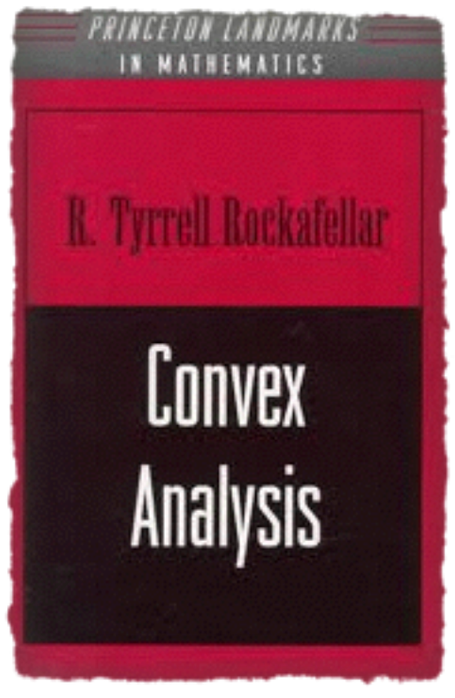

$$x_{k+1} \leftarrow \operatorname{prox}_{\alpha_k \partial f}(x_k)$$

Rockafellar

Operator Splitting

If we can write $F(x) = A(x) + B(x)$

how can we solve $0 \in \ A(x) + B(x)$?

Resolvent

$$R_A = (I + \lambda A)^{-1}$$

$$R_B = (I + \lambda B)^{-1}$$

Cayley Operator ("Reflection operator")

$$C_A = 2R_A - I$$

$$C_B = 2R_B - I$$

Douglas Rachford Splitting

(1956!)

$$z_{k+1} = \frac{1}{2}(I + C_A C_B)(z_k)$$

always converges when $0 \in A(x) + B(x)$ has a solution

$$x_{k+\frac{1}{2}} = R_B(z_k)$$

$$z_{k+\frac{1}{2}} = 2x_{k+\frac{1}{2}} - z_k$$

$$x_{k+1} = R_A(z_{k+\frac{1}{2}})$$

$$z_{k+1} = z_k + x_{k+1} - x_{k+\frac{1}{2}}$$

Alternating Direction Method of Multipliers

Minimize $f(x) + g(x)$

Solve $0 \in \partial f(x) + \partial g(x)$

Choose $A(x) = \partial g(x), B(x) = \partial f(x)$

ADMM is an instance of Douglas Rachford splitting

$$x_{k+\frac{1}{2}} = \text{argmin}_x(f(x) + \frac{1}{2\lambda}\|x - z_k\|_2^2)$$

$$z_{k+\frac{1}{2}} = 2x_{k+\frac{1}{2}} - z_k$$

$$x_{k+1} = \text{argmin}_x(g(x) + \frac{1}{2\lambda}\|x - x_{k+\frac{1}{2}}\|_2^2)$$

$$z_{k+1} = z_k + x_{k+1} - x_{k+\frac{1}{2}}$$

Constrained Optimization

$$\text{Minimize } f \ s.t. \ x \in \mathcal{C}$$

$$\text{Minimize } f + I_{\mathcal{C}}$$

$$\text{Find } x^* \ s.t. \ 0 \in \partial f(x^*) + N_{\mathcal{C}}(x^*)$$

Douglas Rachford splitting (ADMM)

$$x_{k+\frac{1}{2}} = \text{prox}_f(z_k)$$

$$z_{k+\frac{1}{2}} = 2x_{k+\frac{1}{2}} - z_k$$

$$x_{k+1} = \Pi_{\mathcal{C}}(z_{k+\frac{1}{2}})$$

$$x_{k+1} = z_k + x_{k+1} - x_{k+\frac{1}{2}}$$

Dual Decomposition

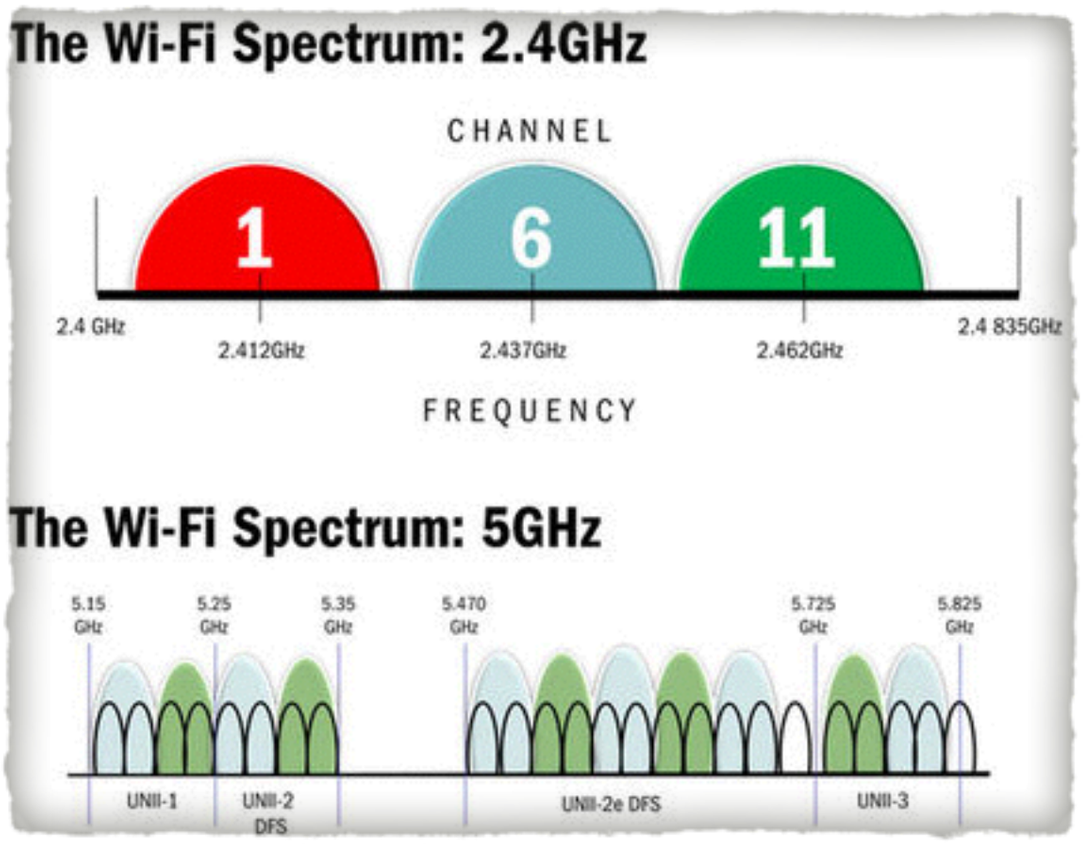



$$\max \sum_i u_i(x_i) \ s.t.$$

$$\sum_i x_i \leq c$$

$$x_i \geq 0, \ \forall i$$

$$q_i(\lambda) = \min_i(-u_i(x_i) + \lambda x_i)$$

$$q(\lambda) = \sum_i q_i(\lambda) - \lambda c$$

$$\lambda \leftarrow \lambda + \alpha_t(c - \sum_i x_i^*)$$

Convex Feasibilty Revisited

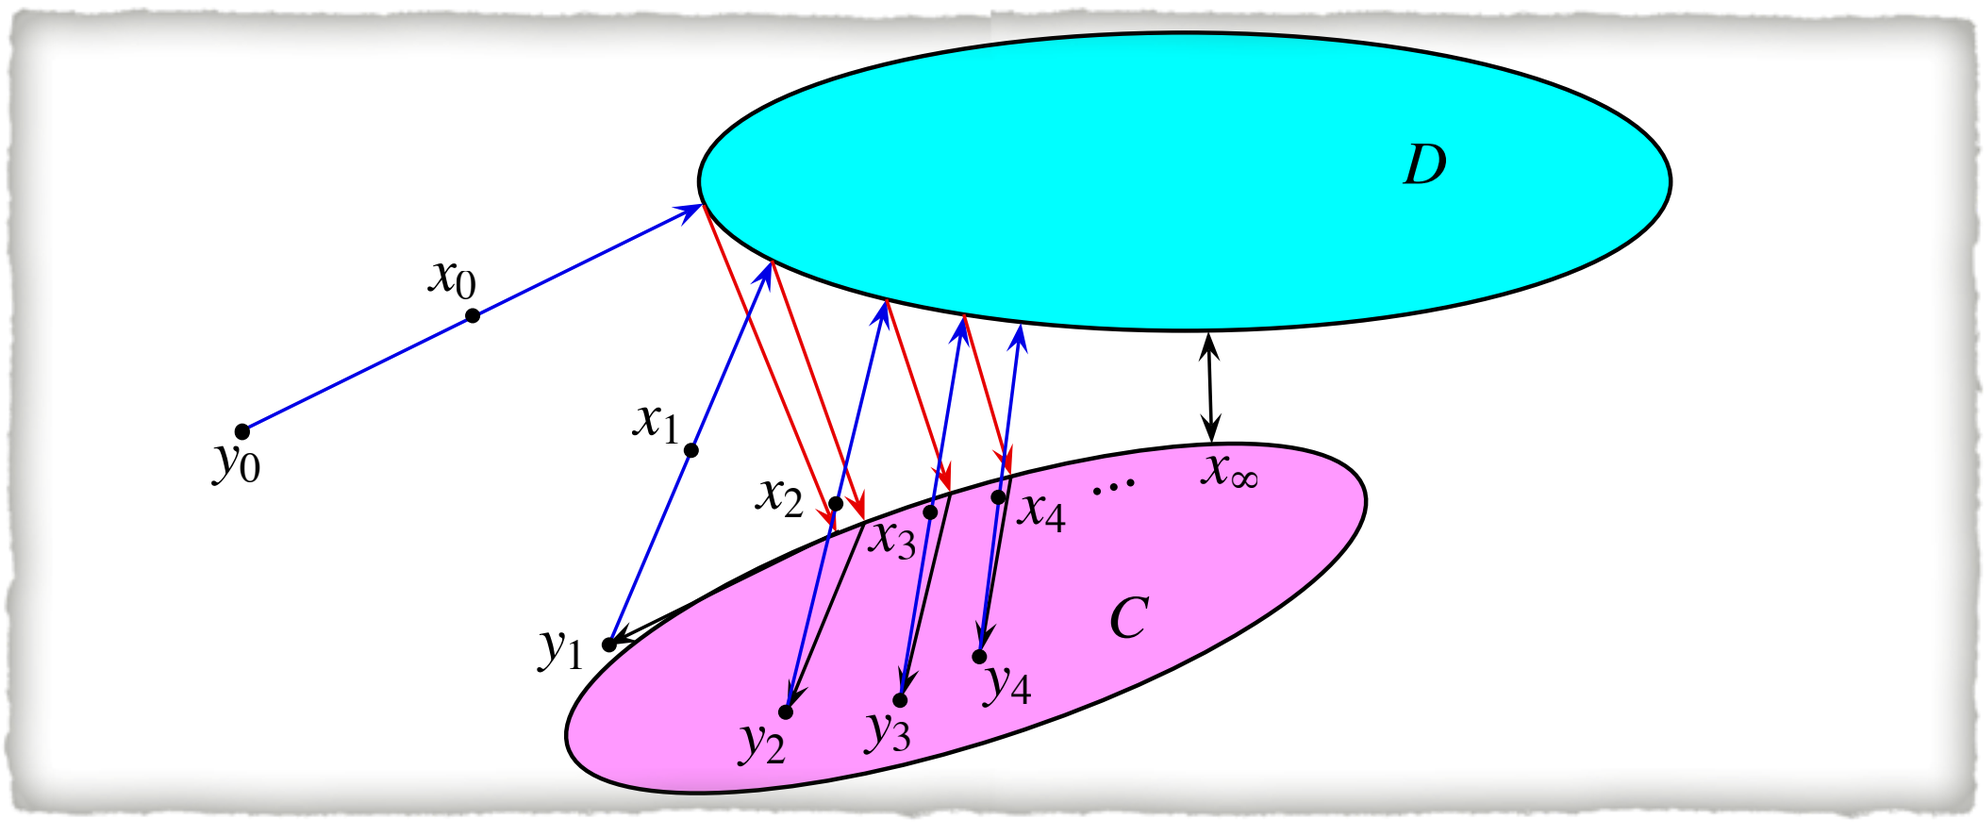

Proximal splitting methods in signal processing
Combetti and Pesquet

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."

Generalized ADMM Method for Separable VIs

(Tseng, 1988)

$$\langle x - x^*, R(x^*)\rangle + \langle z - z^*, S(z^*)\rangle \geq 0,$$
$$\forall (x,z) \in X \times Z \ s.t. \ Ax + Bz = b$$

↓

$$\text{Minimize } \langle R(x^*), x\rangle + \langle S(z^*), z\rangle$$
$$s.t. \ x \in X, z \in Z, Ax + Bz = b$$

↓

Let $N(.|X), N(.|Z)$ be subdifferentials of $\delta(.|X), \delta(.|Z)$

↓

Let p^* be the optimal Lagrange multiplier for $Ax + Bz = b$

Generalized ADMM for Separable VIs

Karush Kuhn Tucker conditions imply:

$$A^T p^* \in N(x^*|X) + R(x^*)$$

$$B^T p^* \in N(z^*|Z) + S(z^*)$$

$$Ax^* + Bz^* = b$$

Define maximal monotone operators

$$F(x) = R(x) + N(x|X)$$

$$G(z) = S(z) + N(z|Z)$$

Above equations can be rewritten as:

$$AF^{-1}(A^T p^*) + BG^{-1}(B^T p^*) = b$$

Splitting Algorithm for Separable VIs

Find x_t $s.t.$ $\langle x - x_t, R(x_t) - A^T p(t) \rangle \geq 0, \forall x \in X$

Compute z_t $s.t.$

$\langle z - z_t, S(z_t) - B^T(p(t) - c(t)(Ax_t + Bz_t - b) \rangle \geq 0,$

$\forall z \in Z$

Update $p(t+1) = p(t) + c(t)(b - Ax_t - Bz_t)$

Proximal Gradient Method

$$x^* = \text{argmin}_x(f(x) + g(x))$$

$$0 \in \nabla f(x^*) + \partial g(x^*)$$
$$0 \in \lambda \nabla f(x^*) + \lambda \partial g(x^*)$$
$$0 \in \lambda \nabla f(x^*) - x^* + x^* + \lambda \partial g(x^*)$$
$$x^* = (I + \lambda \partial g)^{-1}(I - \lambda \nabla f)(x^*)$$
$$x^* = \text{prox}_{\lambda g}(x^* - \lambda \nabla f(x^*))$$

Forward-Backwards method

Bregman Divergence

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

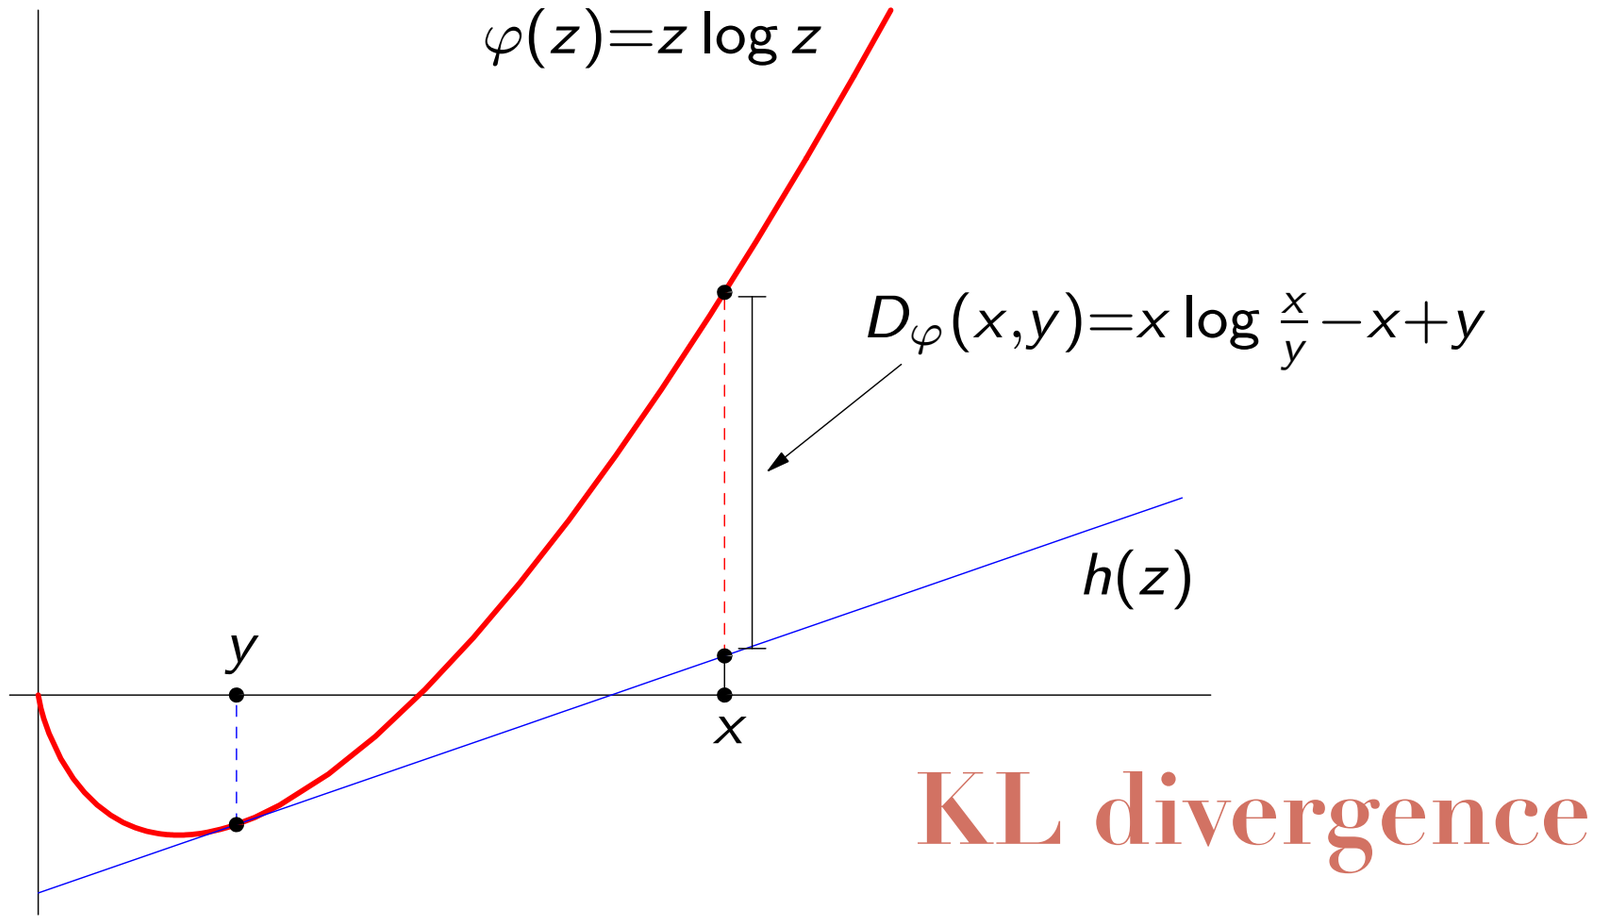

KL divergence

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)

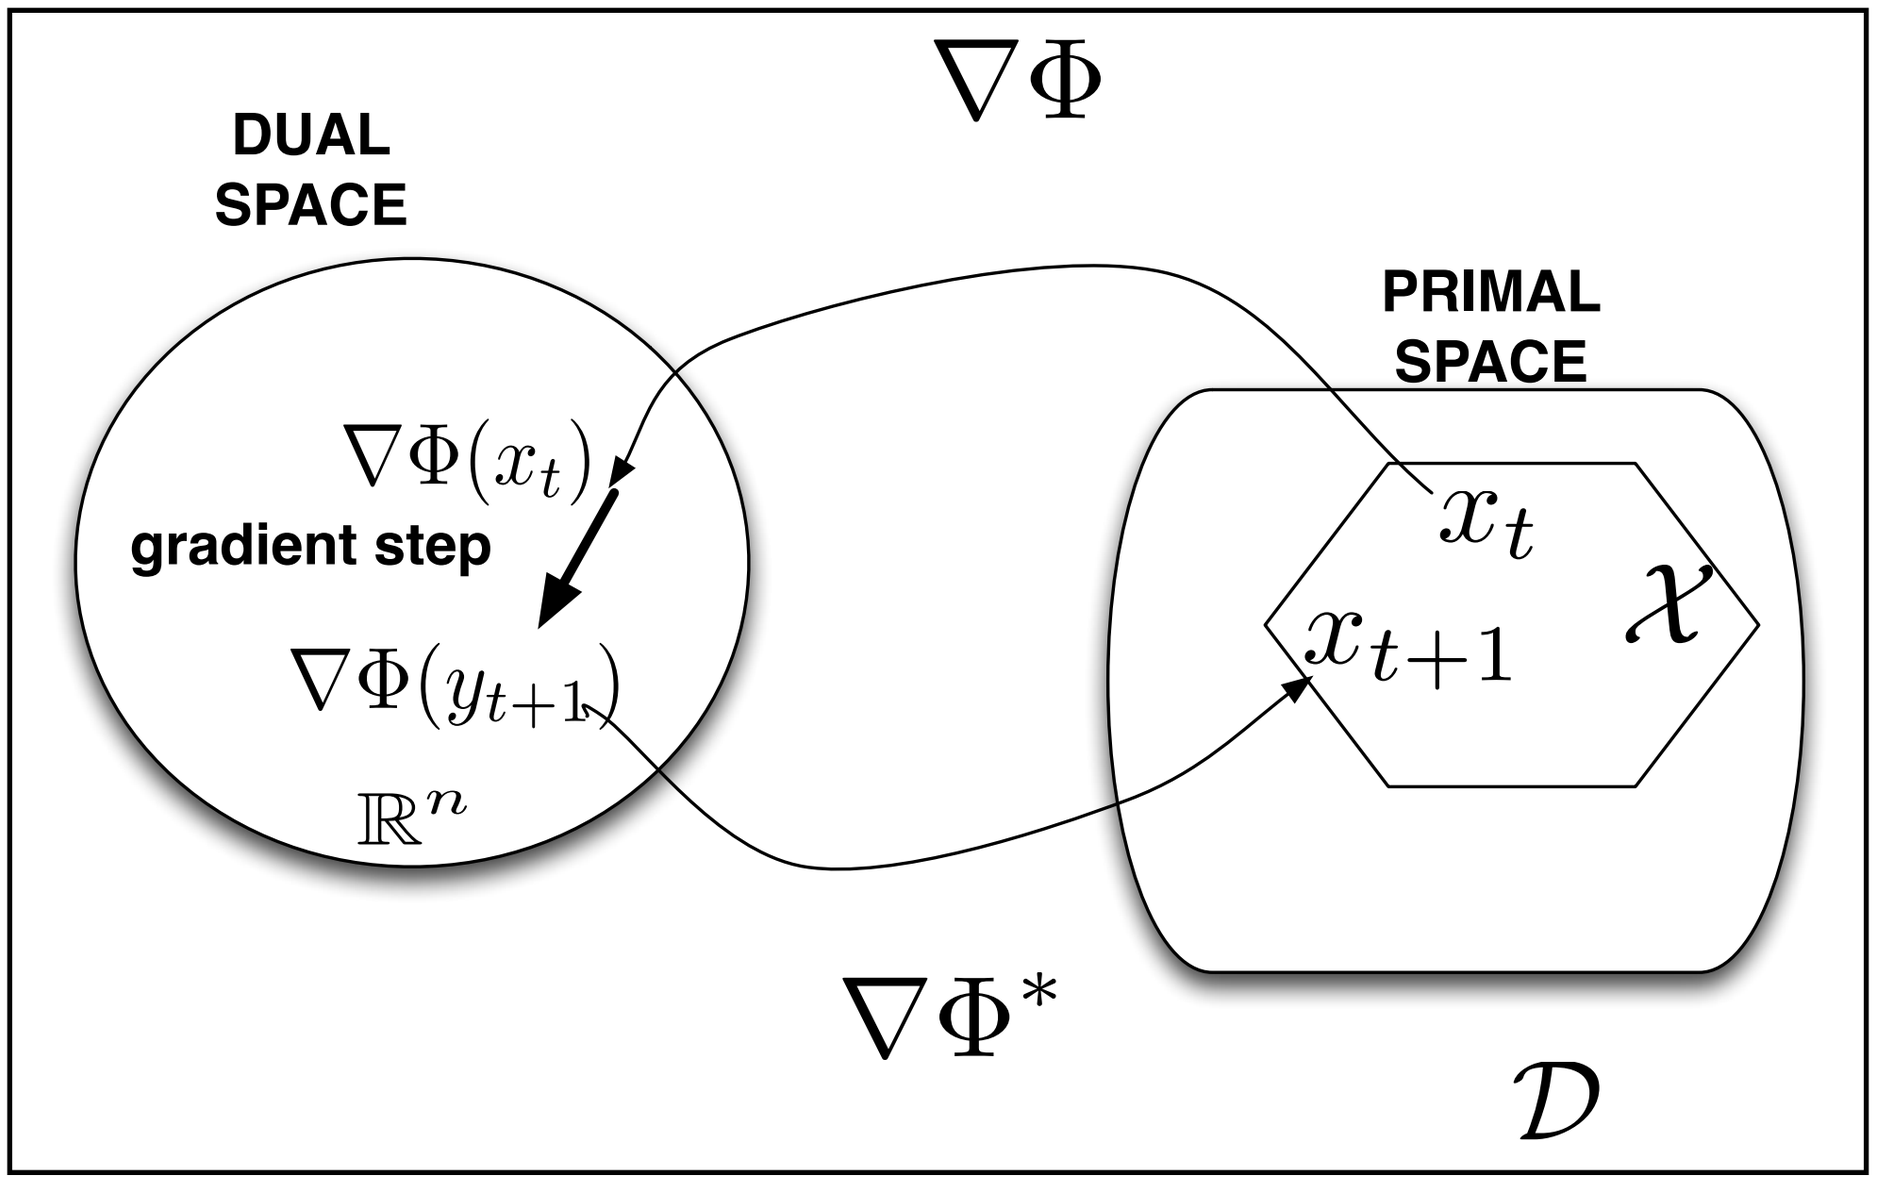

Mirror Descent = "Natural" Gradient

(Nemirovsky and Yudin; Amari, 1980s)

Mirror Descent

Mirror Map

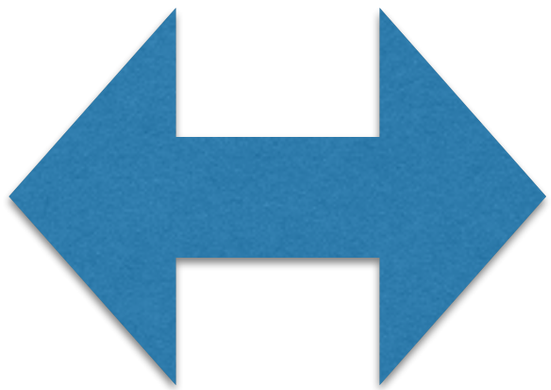

Natural gradient

Thomas, Dabney, Mahadevan, Giguere, NIPS 2013

Mirror-Prox: Non-Euclidean Extragradient

(Nemirovski, 2005)

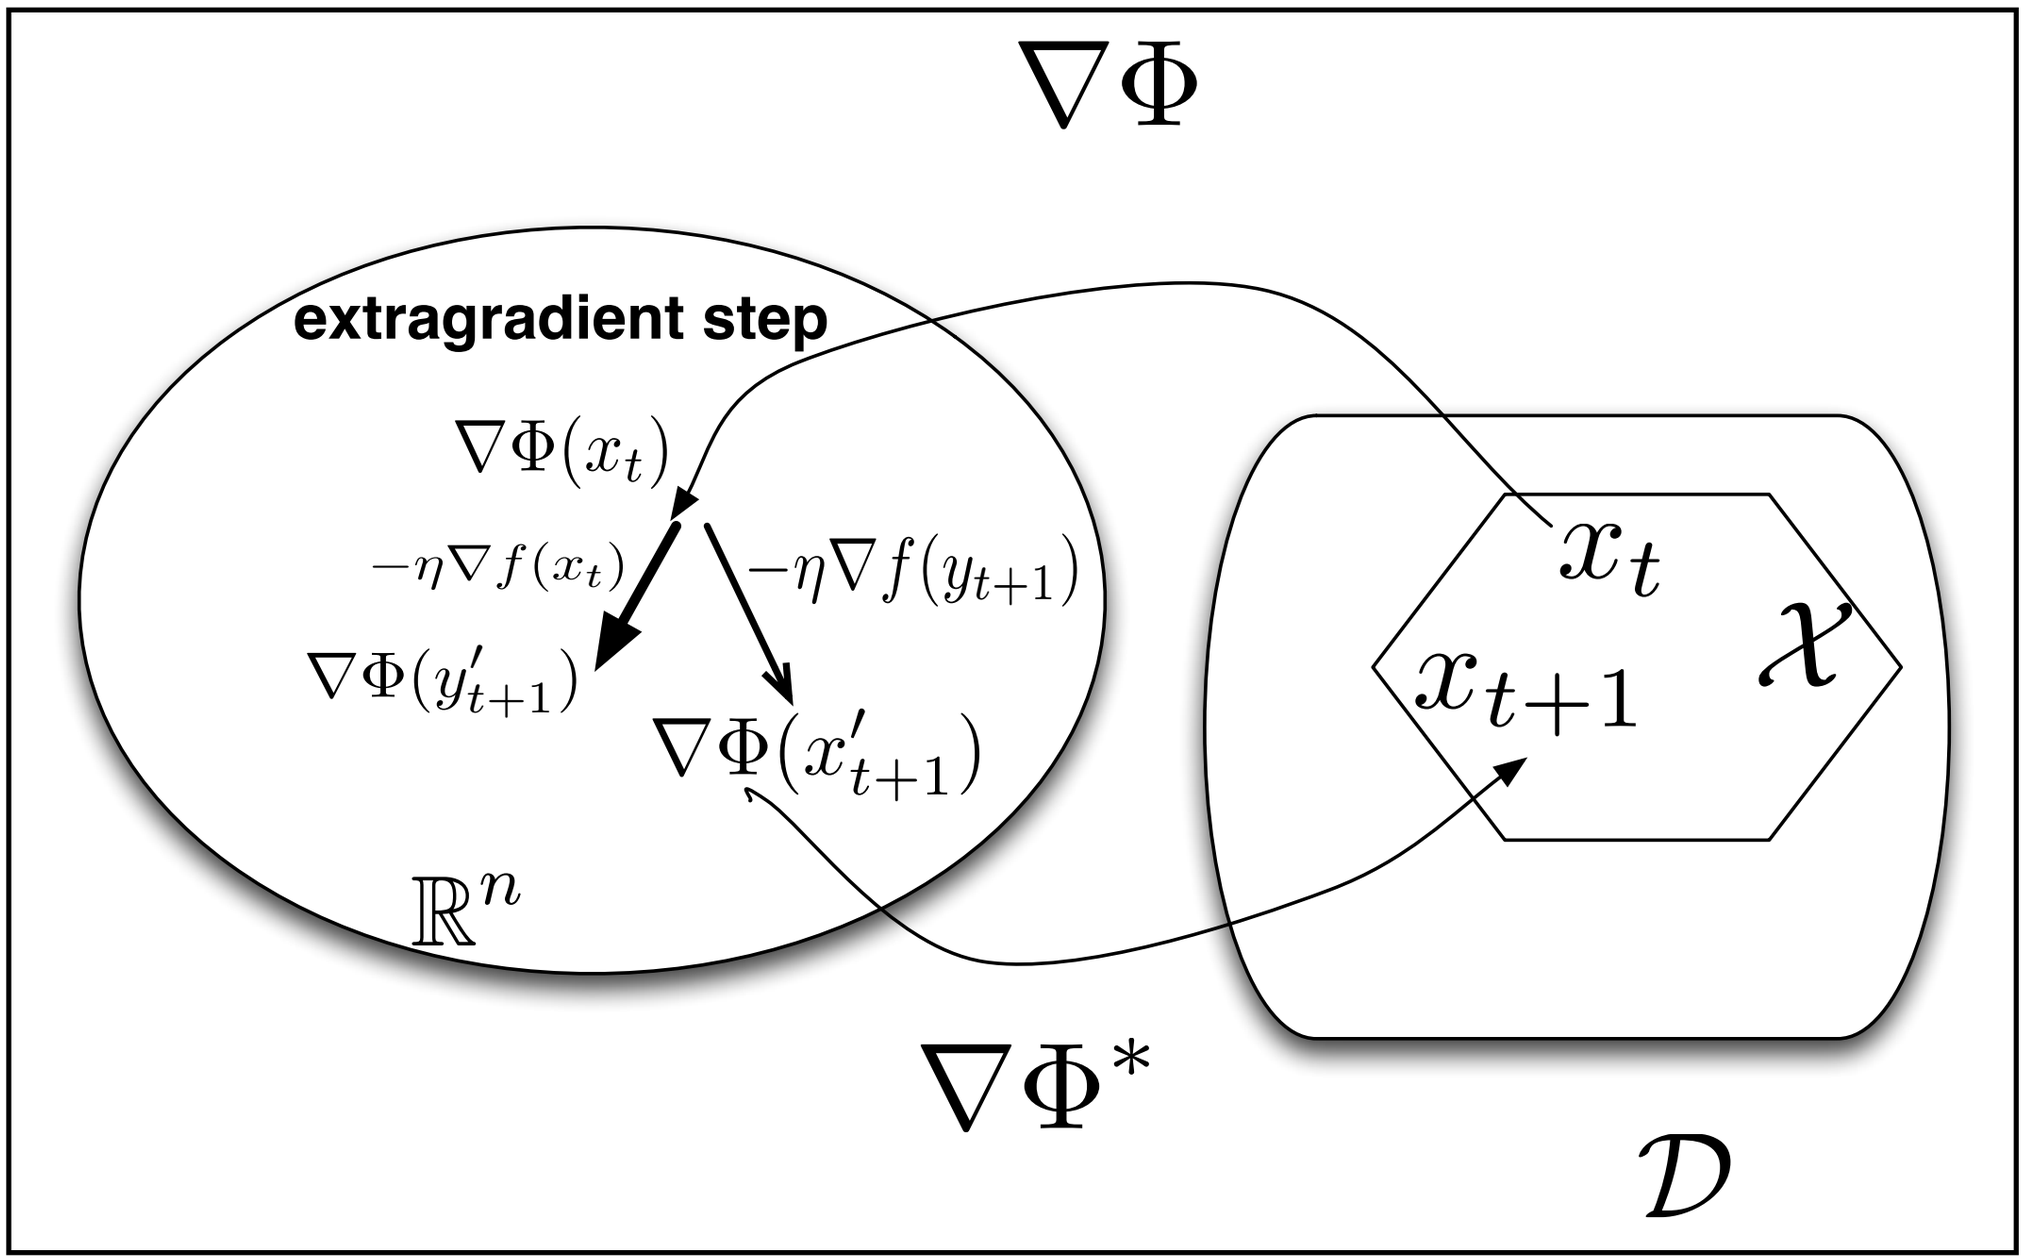

Dual Extragradient

(Nesterov, 2003; Dang and Lan, Taskar, JMLR)

Algorithm 3 Non-Euclidean Extragradient VI Method proposed by Dang and Lan [9]

INPUT: Given initial point $x_1 \in K$, initial step size $\gamma_0 \in (0,1)$, and $\lambda \in (0,1)$.

1: Set $k = 1$.

2: Compute $y_k = \nabla\psi^*(\nabla\psi(x_k) - \gamma_k F(x_k))$

3: Compute the residual $R_{\gamma_0}(x_k) = \frac{1}{\gamma_0}(x_k - \Pi_K(\gamma F(x_k))$. If $\|R_{\gamma_0}(x_k)\| = 0$, terminate the algorithm. Otherwise, do a line search to find the learning rate γ_k with the largest value from the list $\{\gamma_0, \gamma_0\lambda, \gamma_0\lambda^2, \ldots\}$ such that $\|F(x_k) - F(y_k)\|_*^2 \leq \frac{\alpha}{\gamma_k^2} D_\psi(x_k, y_k)$.

4: Set $x_{k+1} = \nabla\psi^*(\nabla\psi(x_k) - \gamma_k F(y_k))$

5: Set $k \leftarrow k + 1$ and go to step 2.

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

$$k_1 = \alpha \nabla F(x_k)$$
$$k_2 = \alpha \nabla F(x_k - \tfrac{1}{2}k_1)$$
$$k_3 = \alpha \nabla F(x_k - \tfrac{1}{2}k_2)$$
$$k_4 = \alpha \nabla F(x_k - k_3)$$

$$x_{k+1} = x_k - \tfrac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4)$$

Runge Kutta (4) Non-Euclidean Extragradient

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

General Runge Kutta Gradient Descent

$$k_1 = \alpha \nabla F(x_k)$$
$$k_2 = \alpha \nabla F(x_k - a_{21}k_1)$$
$$k_3 = \alpha \nabla F(x_k - a_{31}k_1 - a_{32}k_2)$$
$$\vdots$$
$$k_s = \alpha \nabla F(x_k - a_{s1}k_1 - a_{s2}k_2 - ... - a_{s,s-1}k_{s-1})$$

$$x_{k+1} = x_k - \Sigma_{i=1}^{s} b_i k_i$$

General RK Non-Euclidean Extragradient

$$k_1 = \alpha F(x_k)$$
$$k_2 = \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{21}k_1))$$
$$k_3 = \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{31}k_1 - a_{32}k_2))$$
$$\vdots$$
$$k_s = \alpha F(\nabla \psi_k^*(\nabla \psi_k(x_k) - a_{s1}k_1 - a_{s2}k_2 - \ldots - a_{s,s-1}k_{s-1}))$$

$$x_{k+1} = \nabla \psi_k^*(\nabla \psi_k(x_k) - \Sigma_{i=1}^{s} b_i k_i)$$

Benchmark VI Problem

Sun problem: This problem was proposed in [45]. The affine operator $F : \mathbb{R}^n \to \mathbb{R}^n$ is again given y $F(x) = Ax + b$, where

$$A = \begin{pmatrix} 1 & 2 & 2 & \dots & 2 \\ 0 & 1 & 2 & \dots & 2 \\ 0 & 0 & 1 & \dots & 2 \\ \vdots & \vdots & \vdots & \ddots & 2 \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}$$

and $b = (-1, \dots, -1)^T$. The problem instances ranged from $n = 8000$ to $n = 30,000$.

Results on Benchmark VI

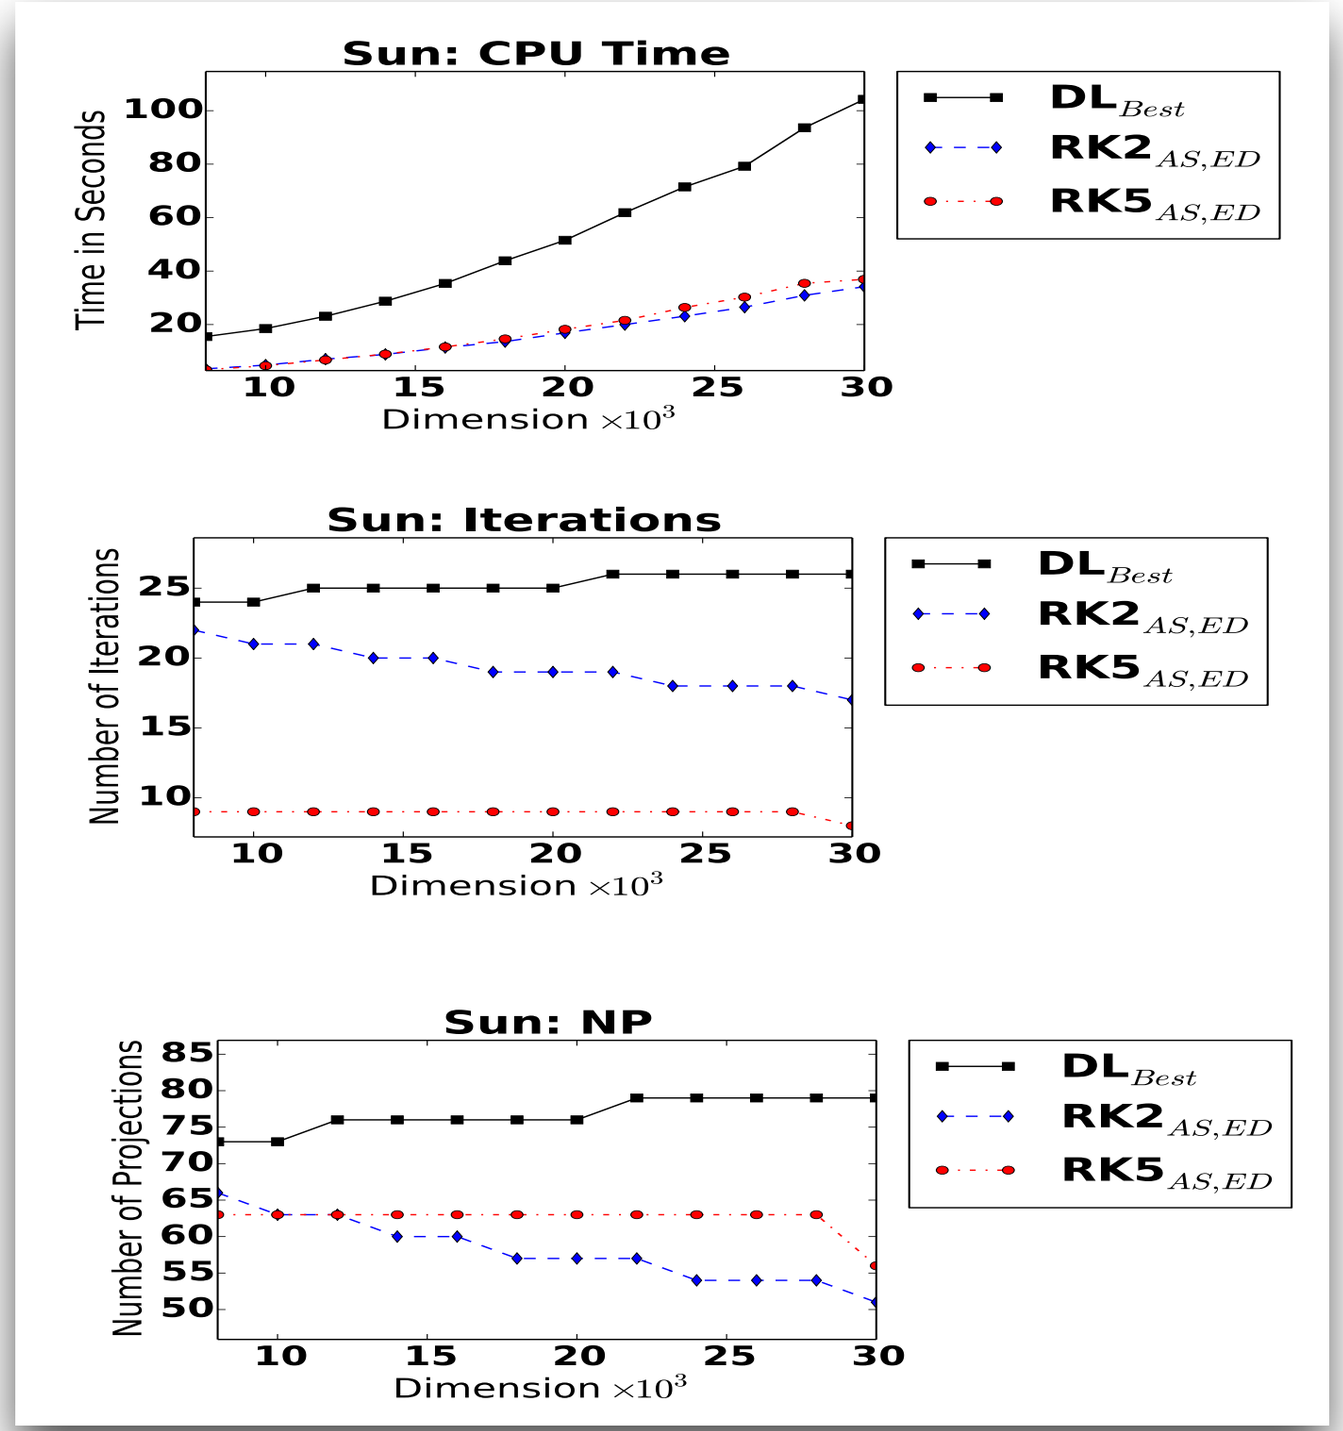

Plan of the Tutorial

- Overview: 2:00-2:10 p.m.
- Part I: Motivation: 2:10-2:30 p.m.
- Part II: Mathematical Background: 2:30-3:30 p.m.
- COFFEE BREAK: 3:30-4:00 p.m.
- Part III: Algorithms: 4:00-5:00 p.m.
- **Part IV: Applications & Future Work: 5:00-6:00 p.m.**

Applications

- Network equilibrium Internet models
- Computational sustainability
- Reinforcement learning
- Natural language processing
- Other applications

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

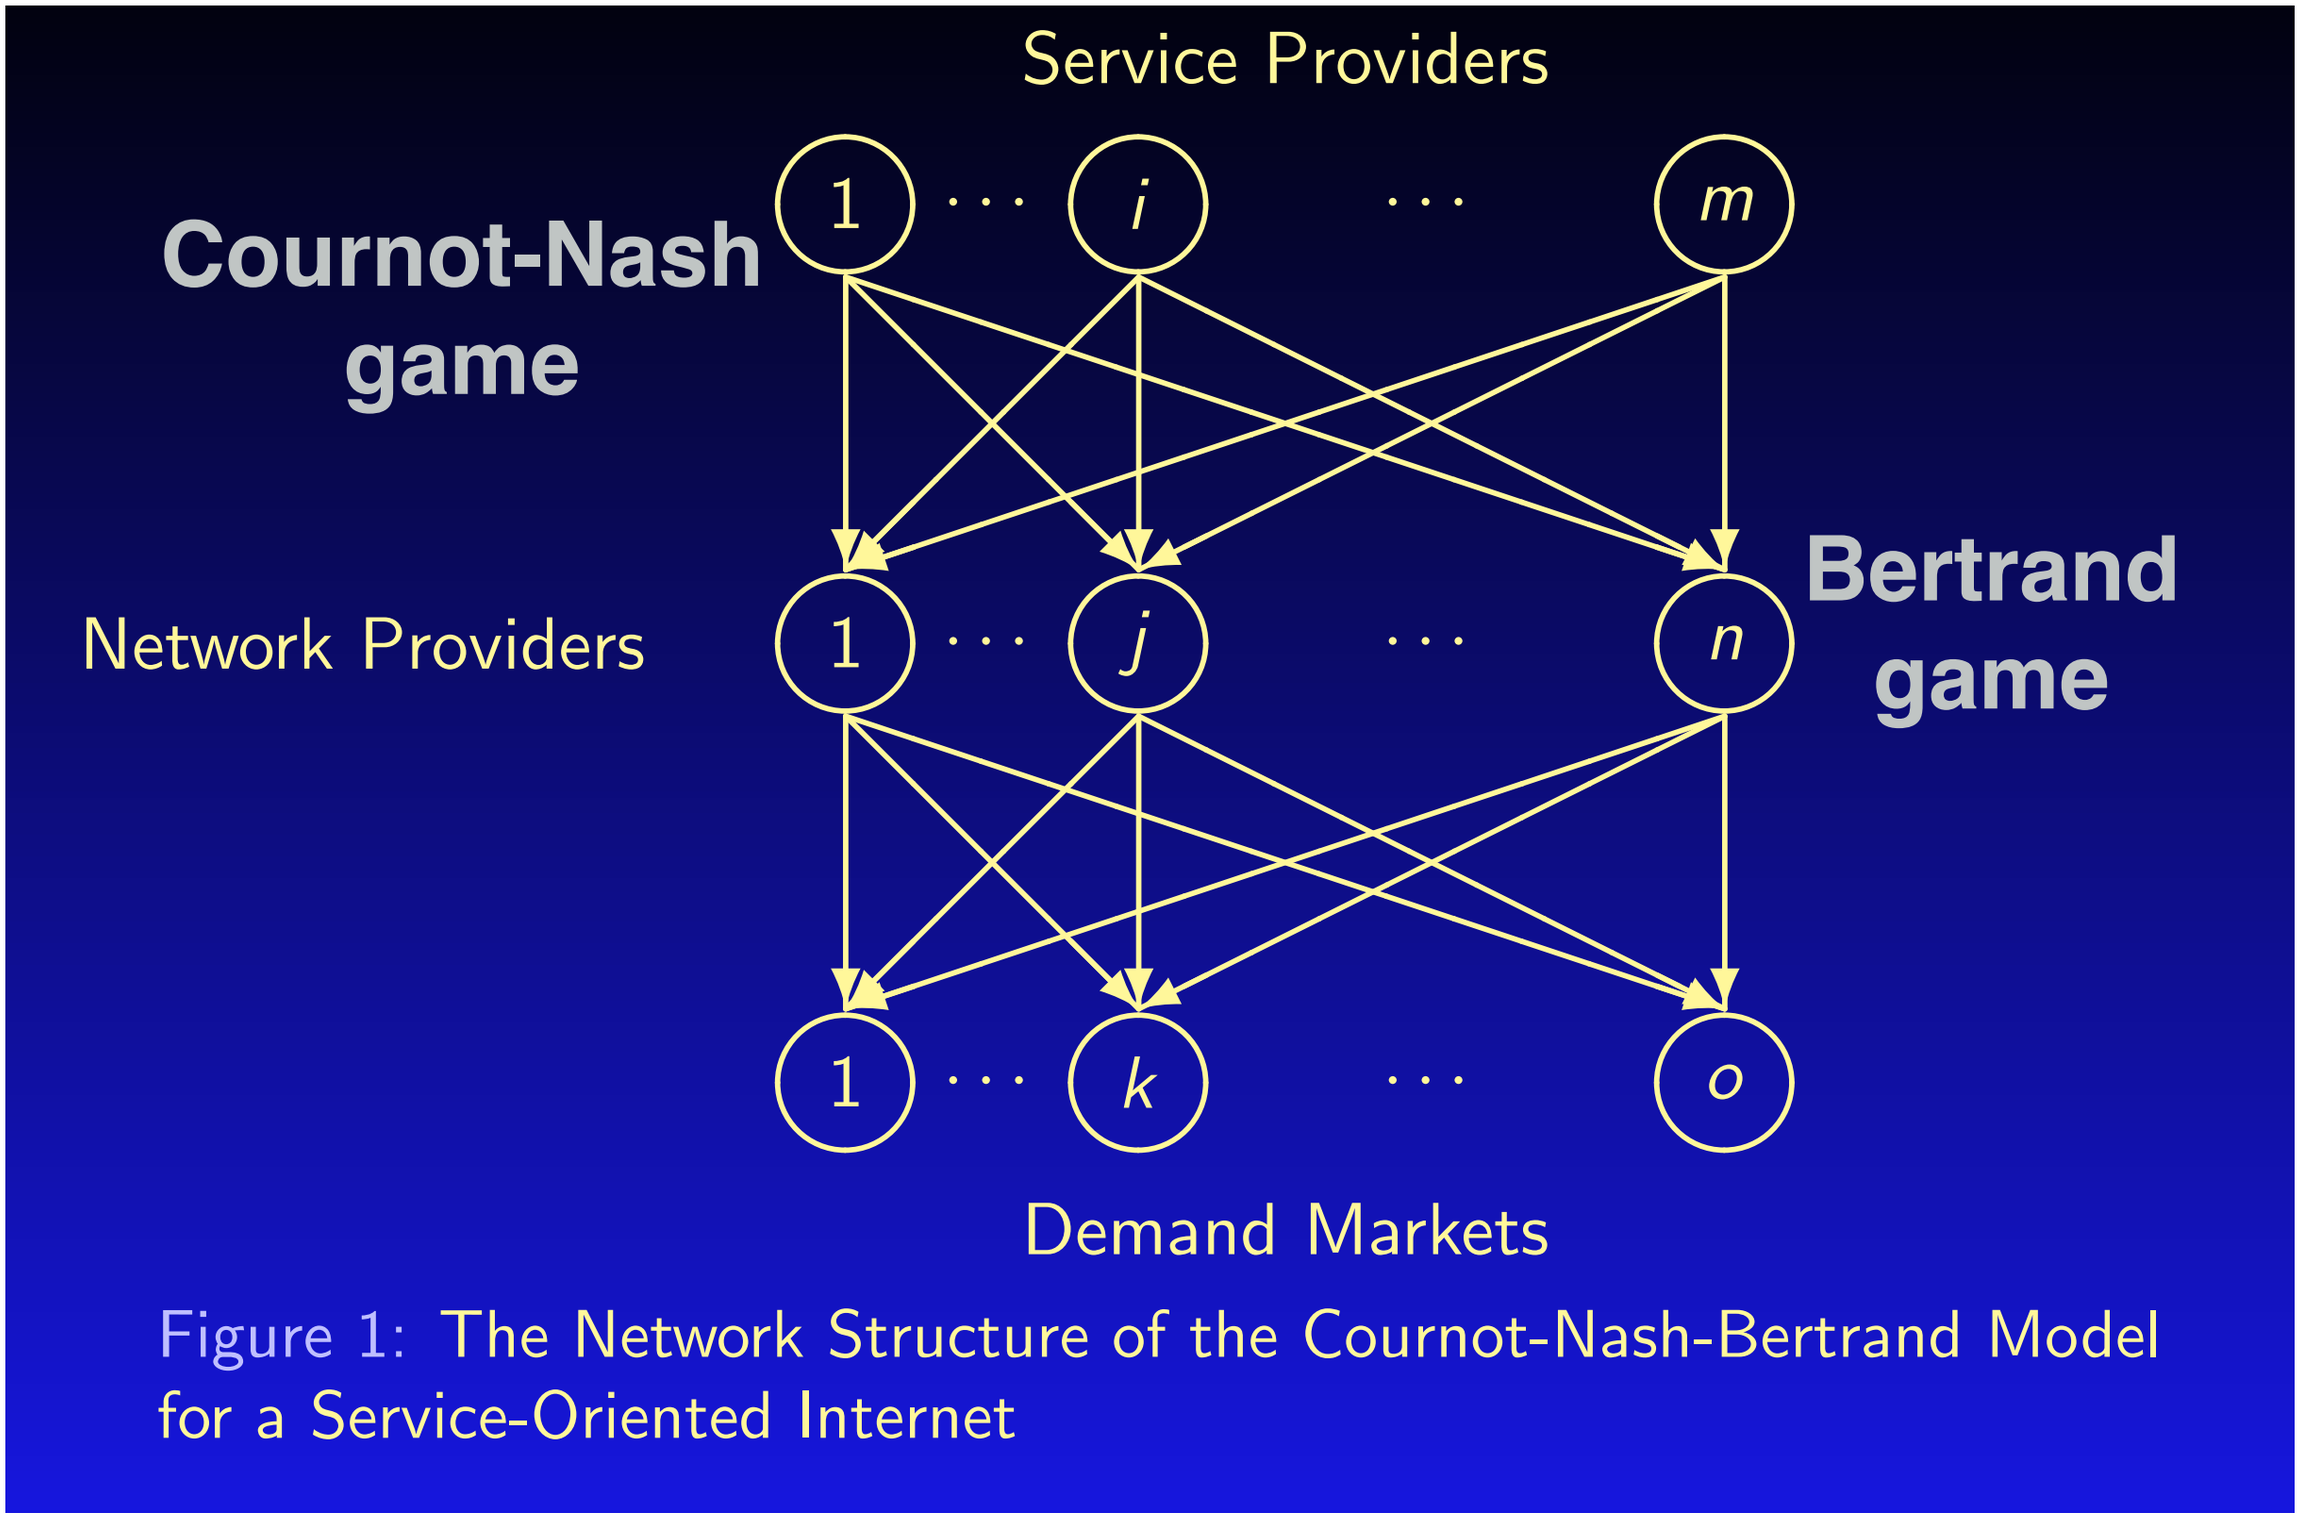

Figure 1: The Network Structure of the Cournot-Nash-Bertrand Model for a Service-Oriented Internet

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

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

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

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

Simple Example

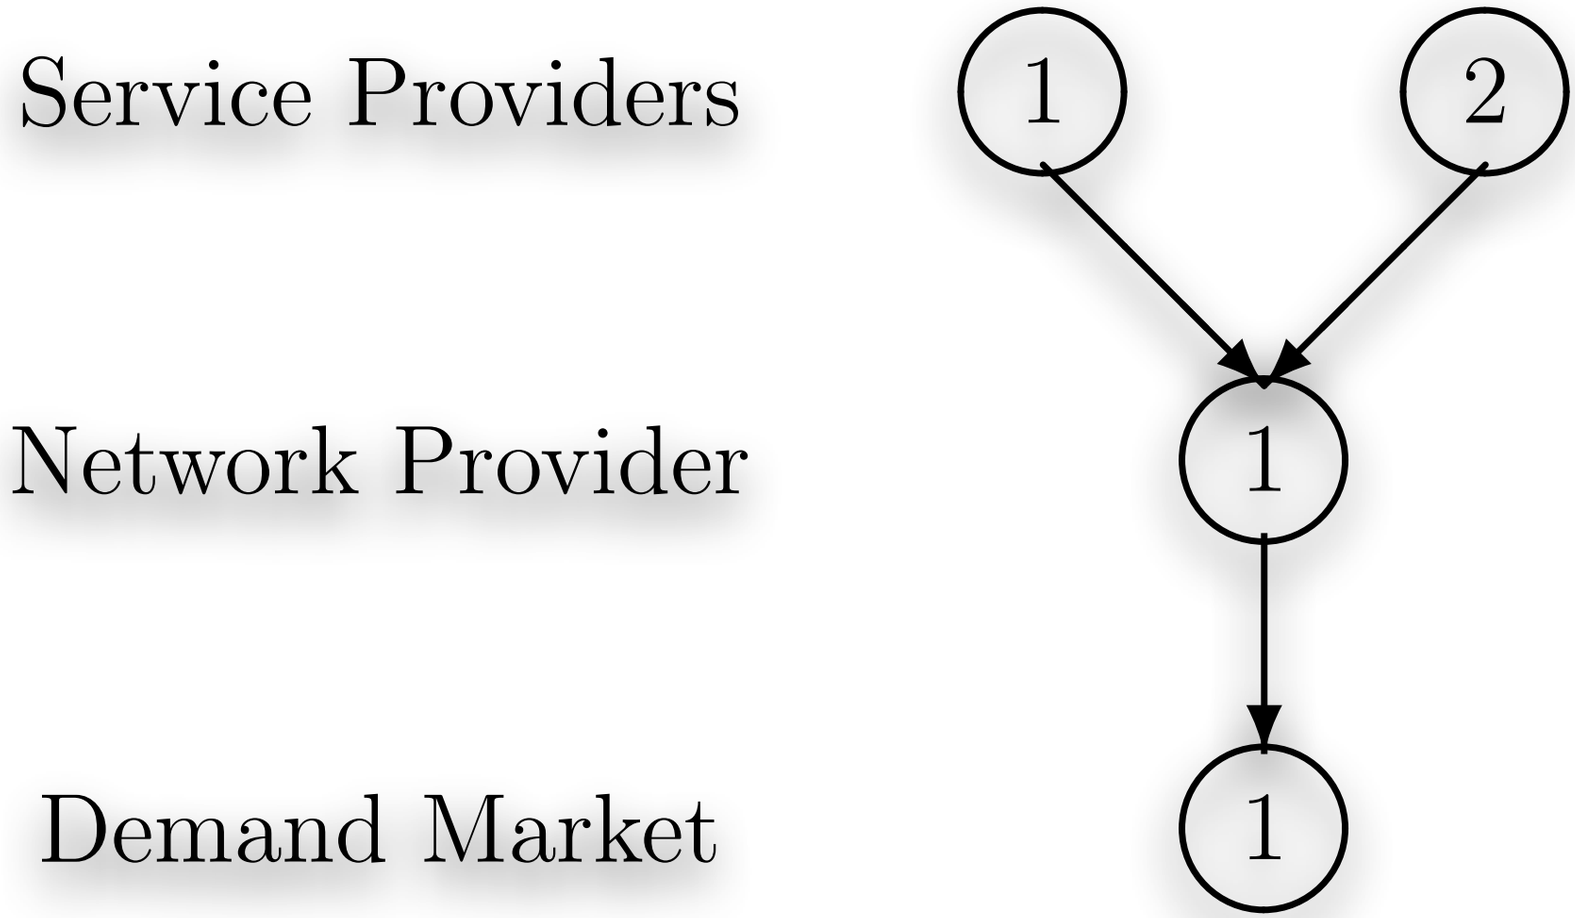

Example

The production cost functions are:

$$\hat{f}_1(Q) = Q_{111}^2 + Q_{111}, \quad \hat{f}_2(Q) = 2Q_{211}^2 + Q_{211}.$$

The demand price functions are:

$$\hat{\rho}_{111}(Q,q) = -Q_{111} - .5Q_{211} + .5q_{111} + 100, \quad \hat{\rho}_{211}(Q,q) = -Q_{211} - .5Q_{111} + .5q_{211} + 200.$$

The transportation cost functions are:

$$\hat{c}_{111}(Q,q) = .5(q_{111} - 20)^2, \quad \hat{c}_{211}(Q,q) = .5(q_{211} - 10)^2,$$

with the opportunity cost functions being:

$$oc_{111}(\pi_{111}) = \pi_{111}^2, \quad oc_{211}(\pi_{211}) = \pi_{211}^2.$$

Example Results

$$F^1_{111}(X) = 2Q_{111} + 1 + \pi_{111} + Q_{111} + .5Q_{211} - .5q_{111} - 100 + Q_{111},$$

$$F^1_{211}(X) = 4Q_{211} + 1 + \pi_{211} + Q_{211} + .5Q_{111} - .5q_{211} - 200 + Q_{211},$$

$$F^2_{111}(X) = q_{111} - 20, \quad F^2_{211}(X) = q_{211} - 10,$$

$$F^3_{111}(X) = -Q_{111} + 2\pi_{111}, \quad F^3_{211}(X) = -Q_{211} + 2\pi_{211}.$$

$$Q^*_{111} = 21.00, \quad Q^*_{211} = 30.00,$$

$$q^*_{111} = 20.00, \quad q^*_{211} = 10.00,$$

$$\pi^*_{111} = 10.50, \quad \pi^*_{211} = 15.00.$$

Results on Internet VI Problem

Extragradient

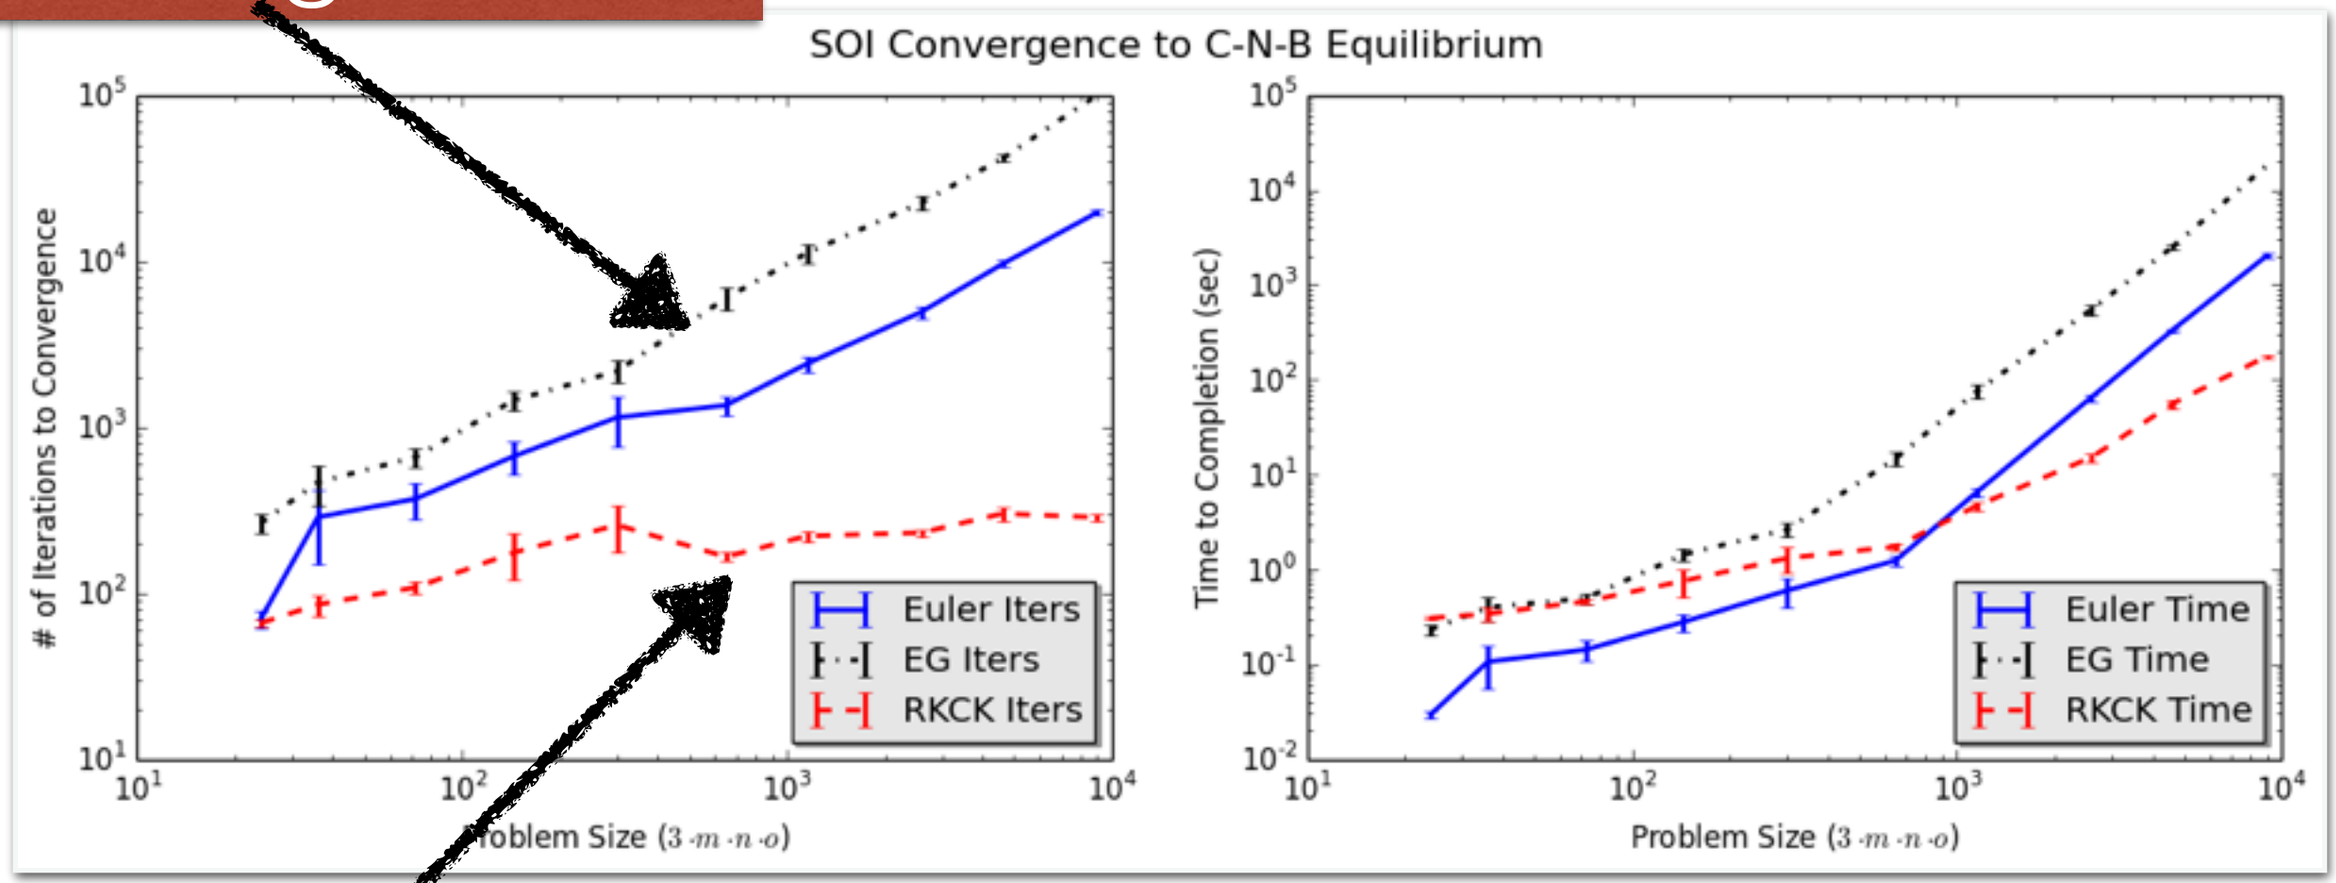

Our new algorithm

Substainable Blood Banking

(Nagurney and Masoumi, 2012)

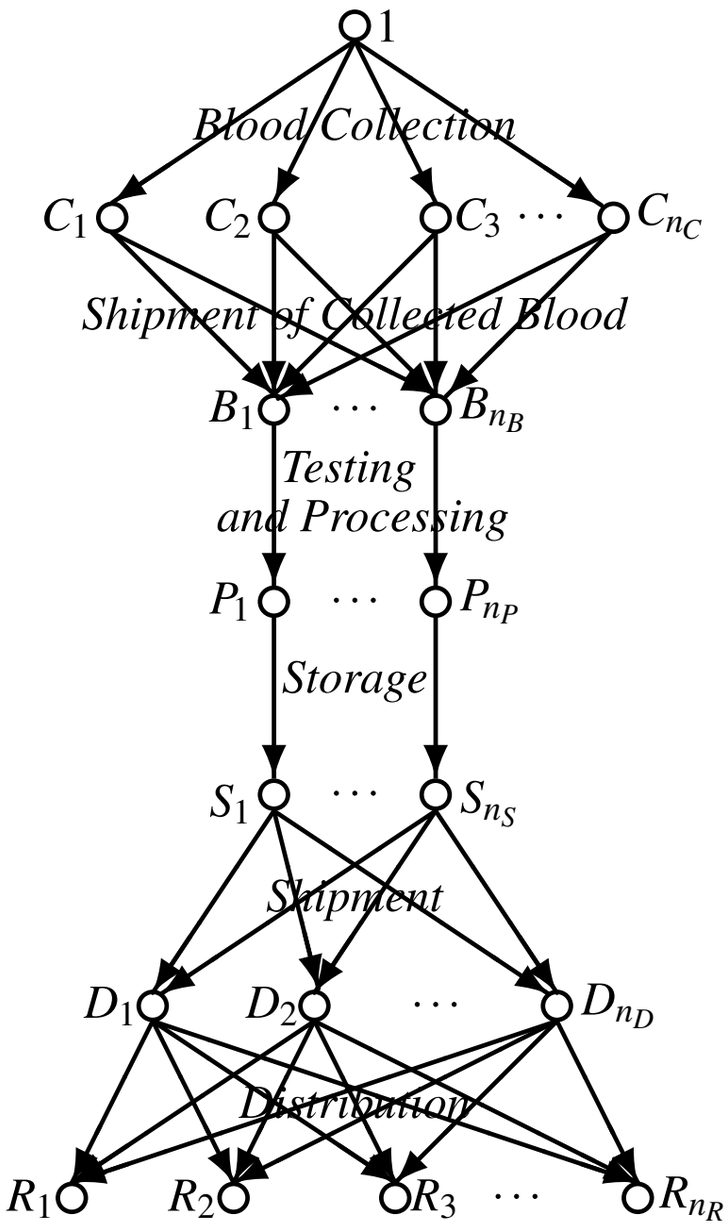

The Organization

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

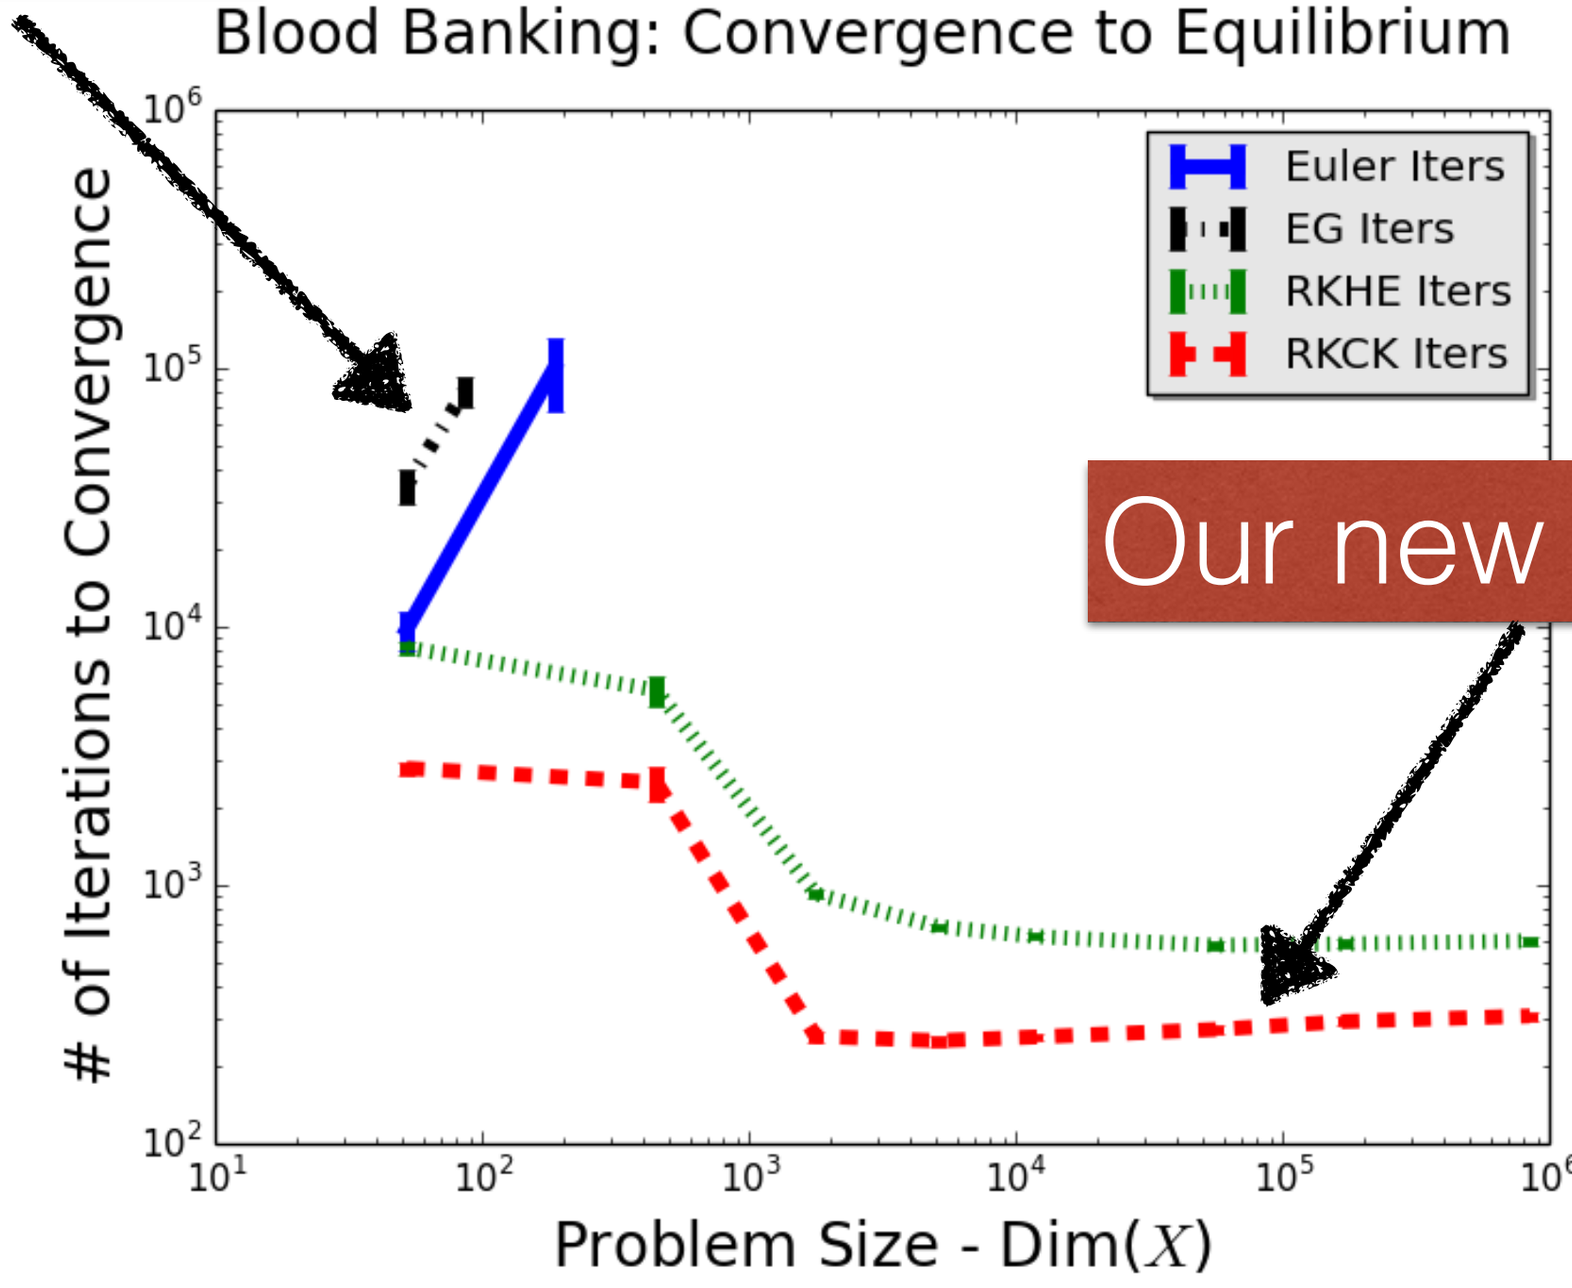

Our new algorithm

Experimental Results

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

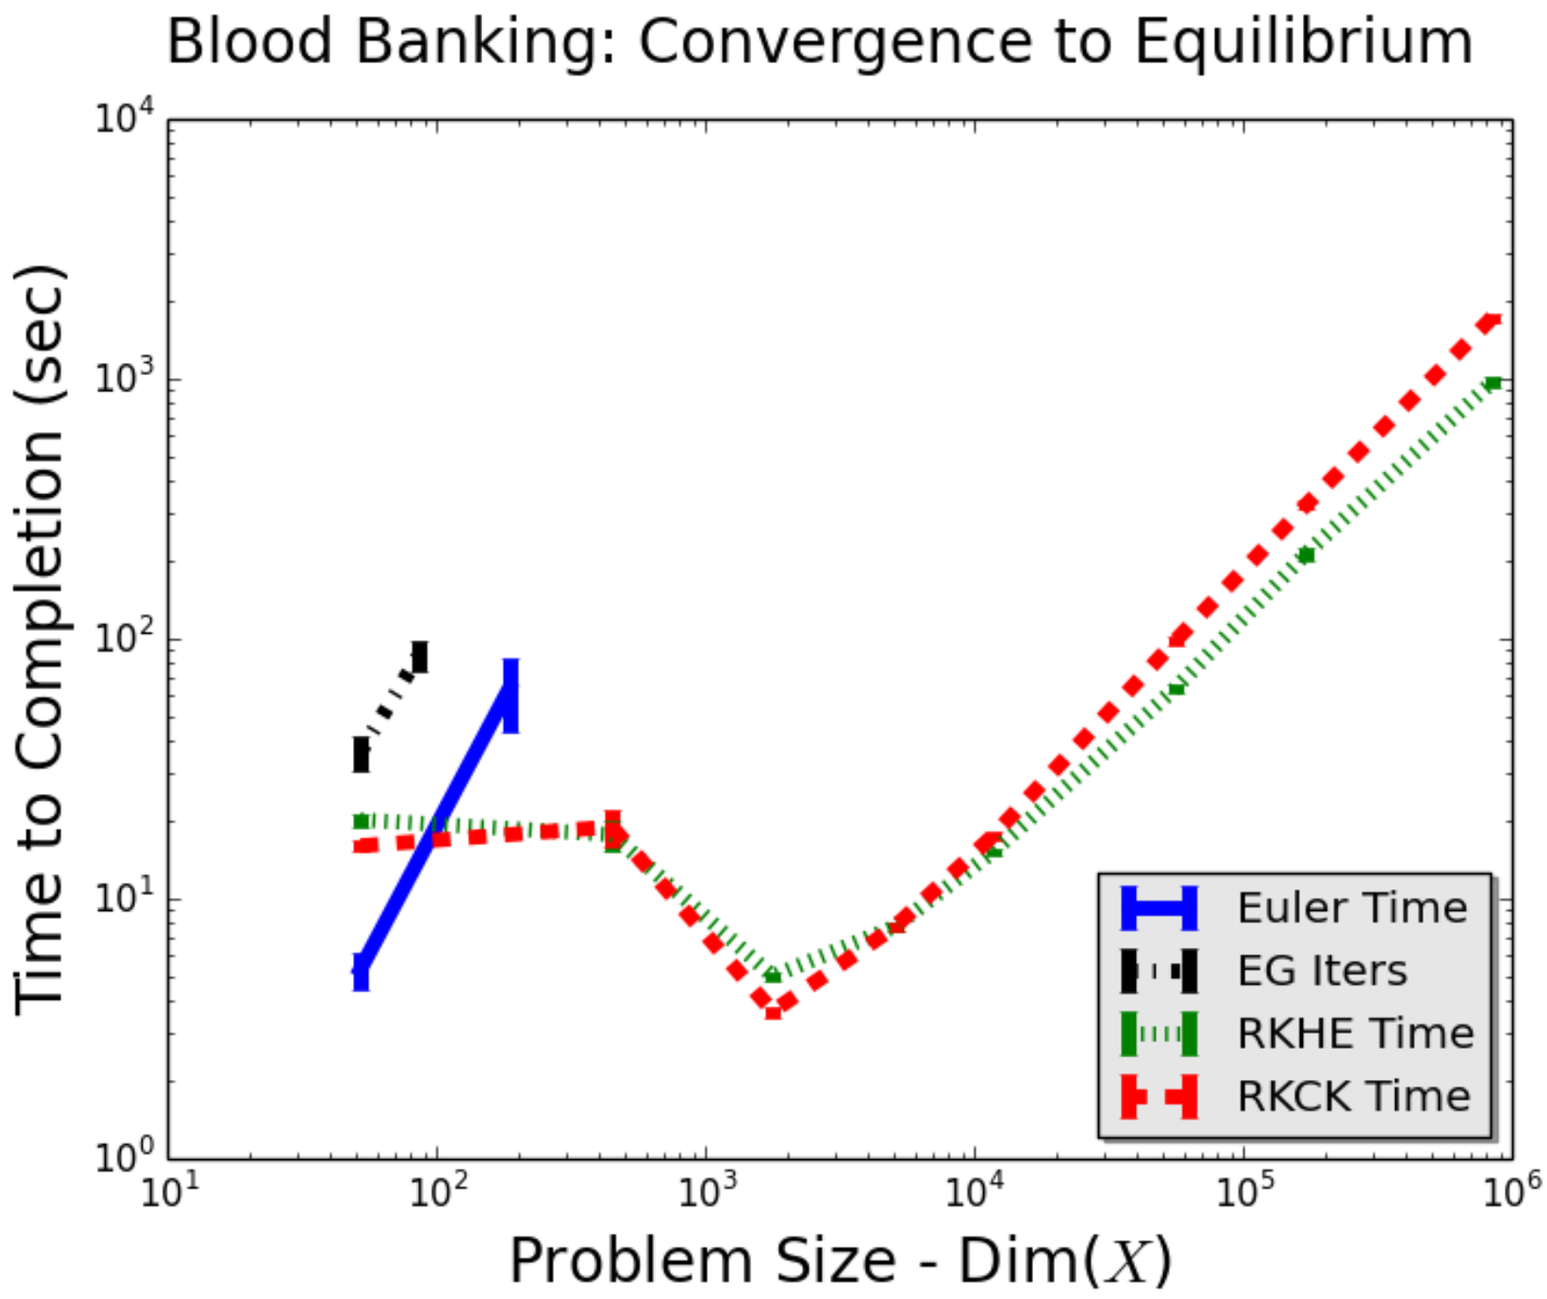

Sustainable Supply Chain

Nagurney et al.

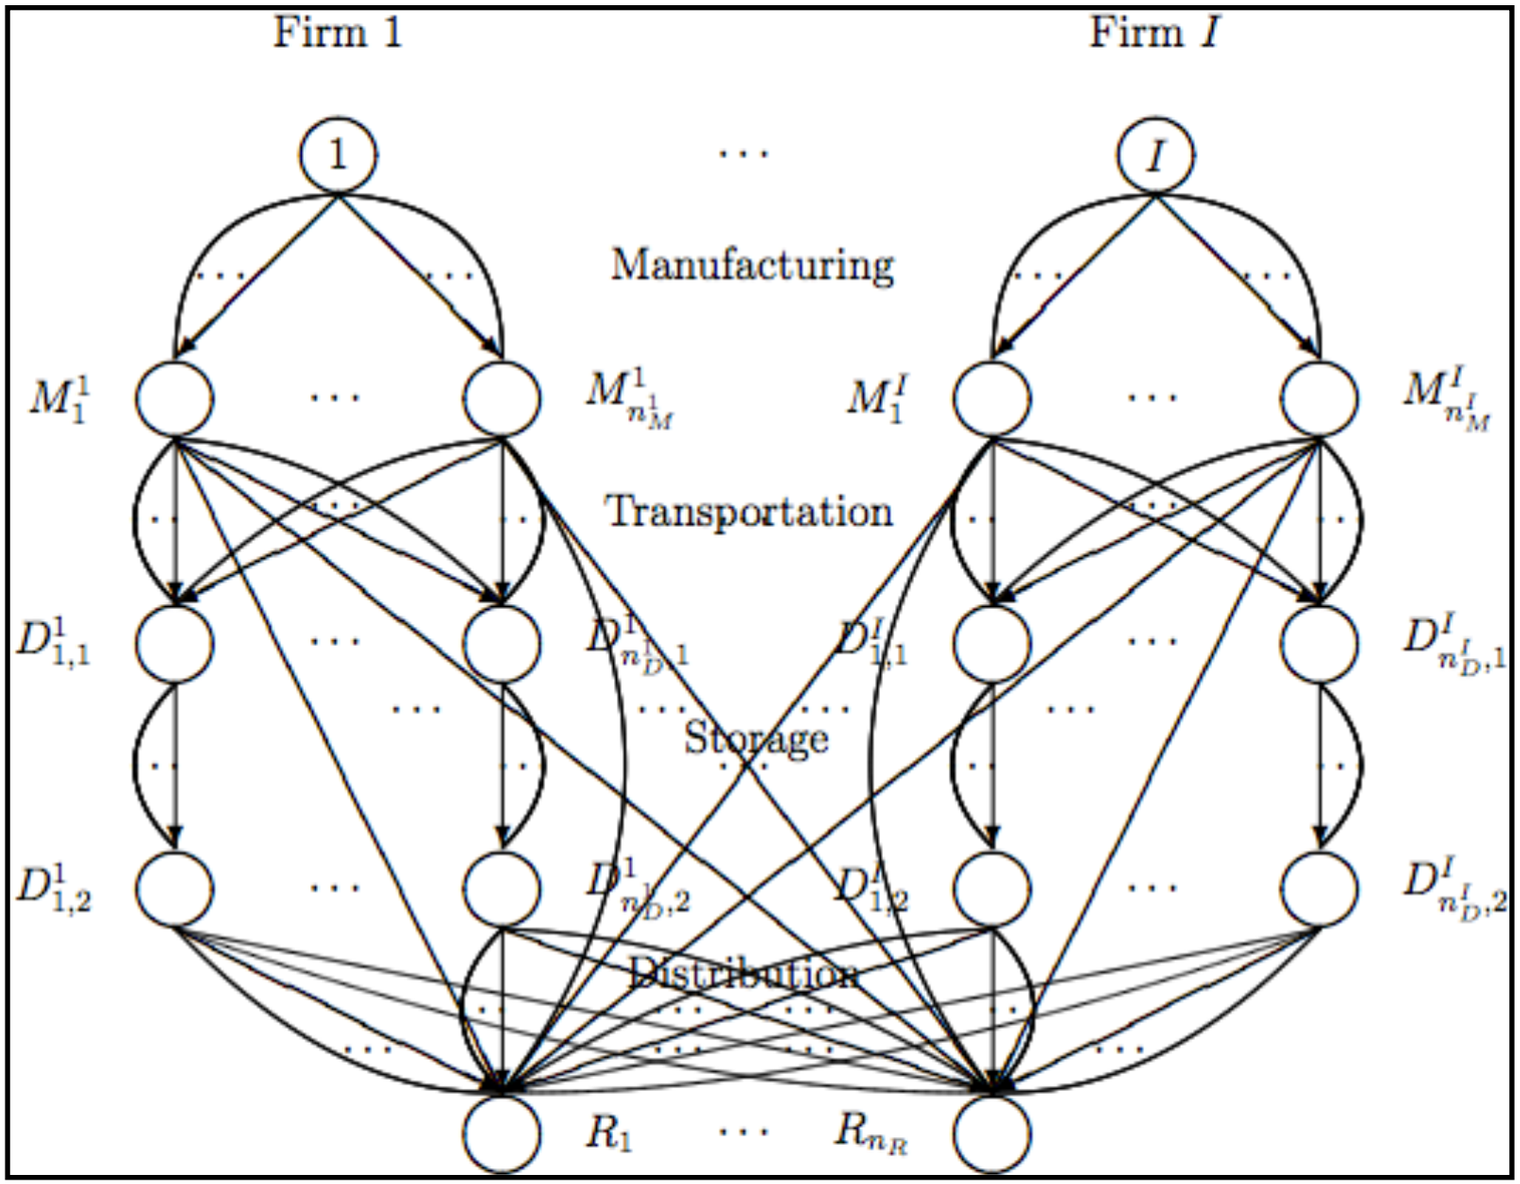

Results on Sustainable Supply Chain VI Problem

Extragradient

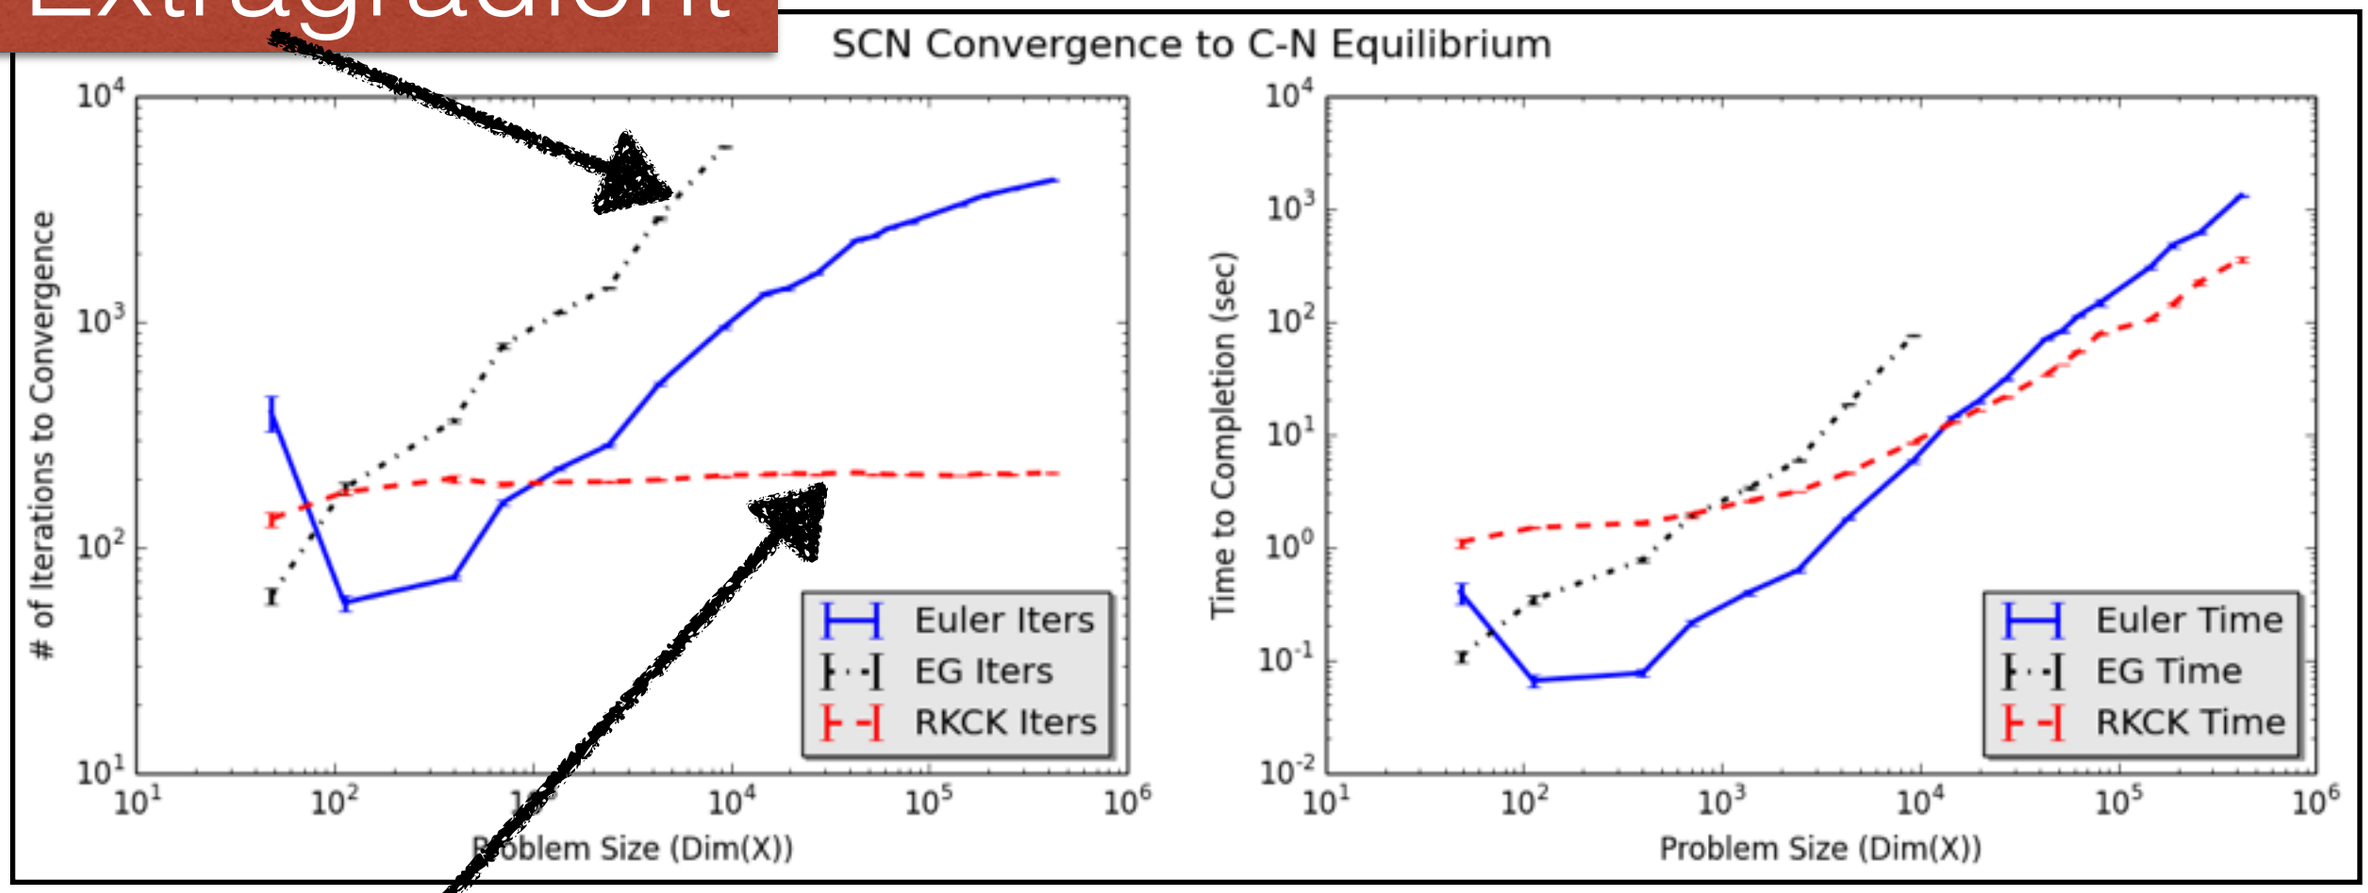

Our new algorithm

Proximal Reinforcement Learning: 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

Proximal Reinforcement Learning in Primal-Dual Spaces

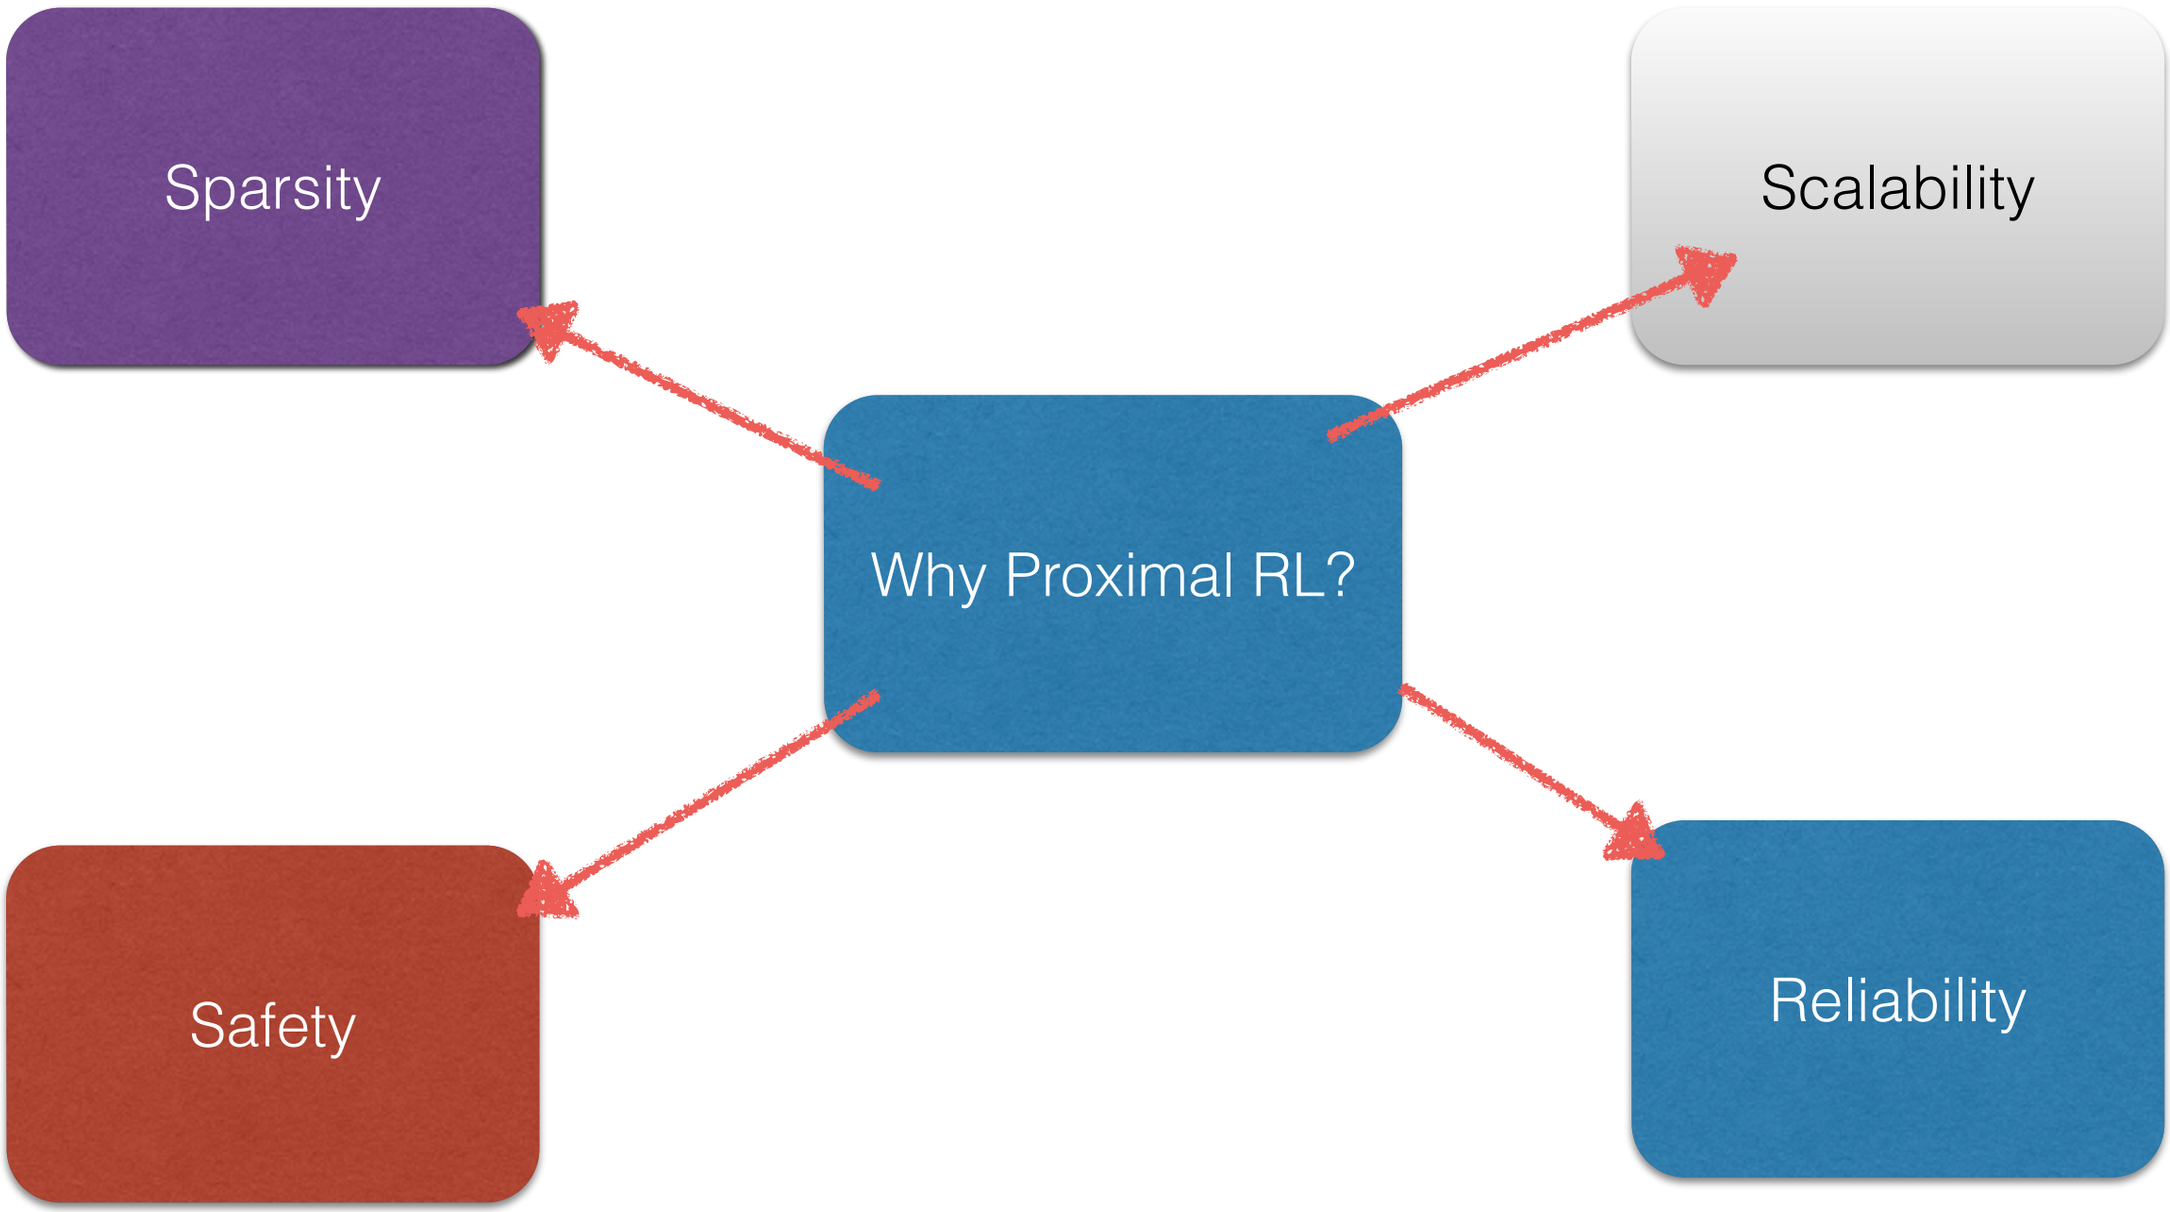

Reinforcement Learning

Bellman Equation

$$V^\pi = T^\pi(V^\pi) = R^\pi + \gamma P^\pi V^\pi$$

TD Algorithm

$$\delta_t = r_t + \gamma\phi(s_{t+1})^T w_t - \phi(s_t)^T w_t$$

$$w_{t+1} = w_t + \alpha_t \delta_t \phi(s_t)$$

TD Fixed Point

It is well known that TD solves for the fixed point

$$\Phi w^* = \Pi_\Phi (R + \gamma P \Phi w^*)$$

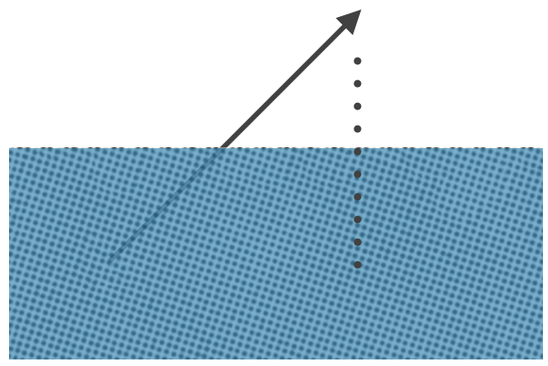

Projected Equations and Deterministic VIs

(Bertsekas, IEEE Trans. Automatic Control, 2011)

$$F(\Phi r^*)^T \Phi(r - r^*) \geq 0, \ \forall r \in \hat{R}$$

$$x^* = \Pi_{\hat{S}} T(x^*)$$

$$(x^* - T(x^*)^T \Xi(x - x^*) \geq 0, \ \forall x \in \hat{S}$$

$$\hat{S} = \{\Phi r | r \in \hat{R}\}, \quad \hat{R} = \{r | \Phi r \in \hat{S}\} \Rightarrow \Phi r^* = \Pi_{\hat{S}} T(\Phi r^*)$$

Mirror-Prox: Non-Euclidean Extragradient

(Nemirovski, 2005)

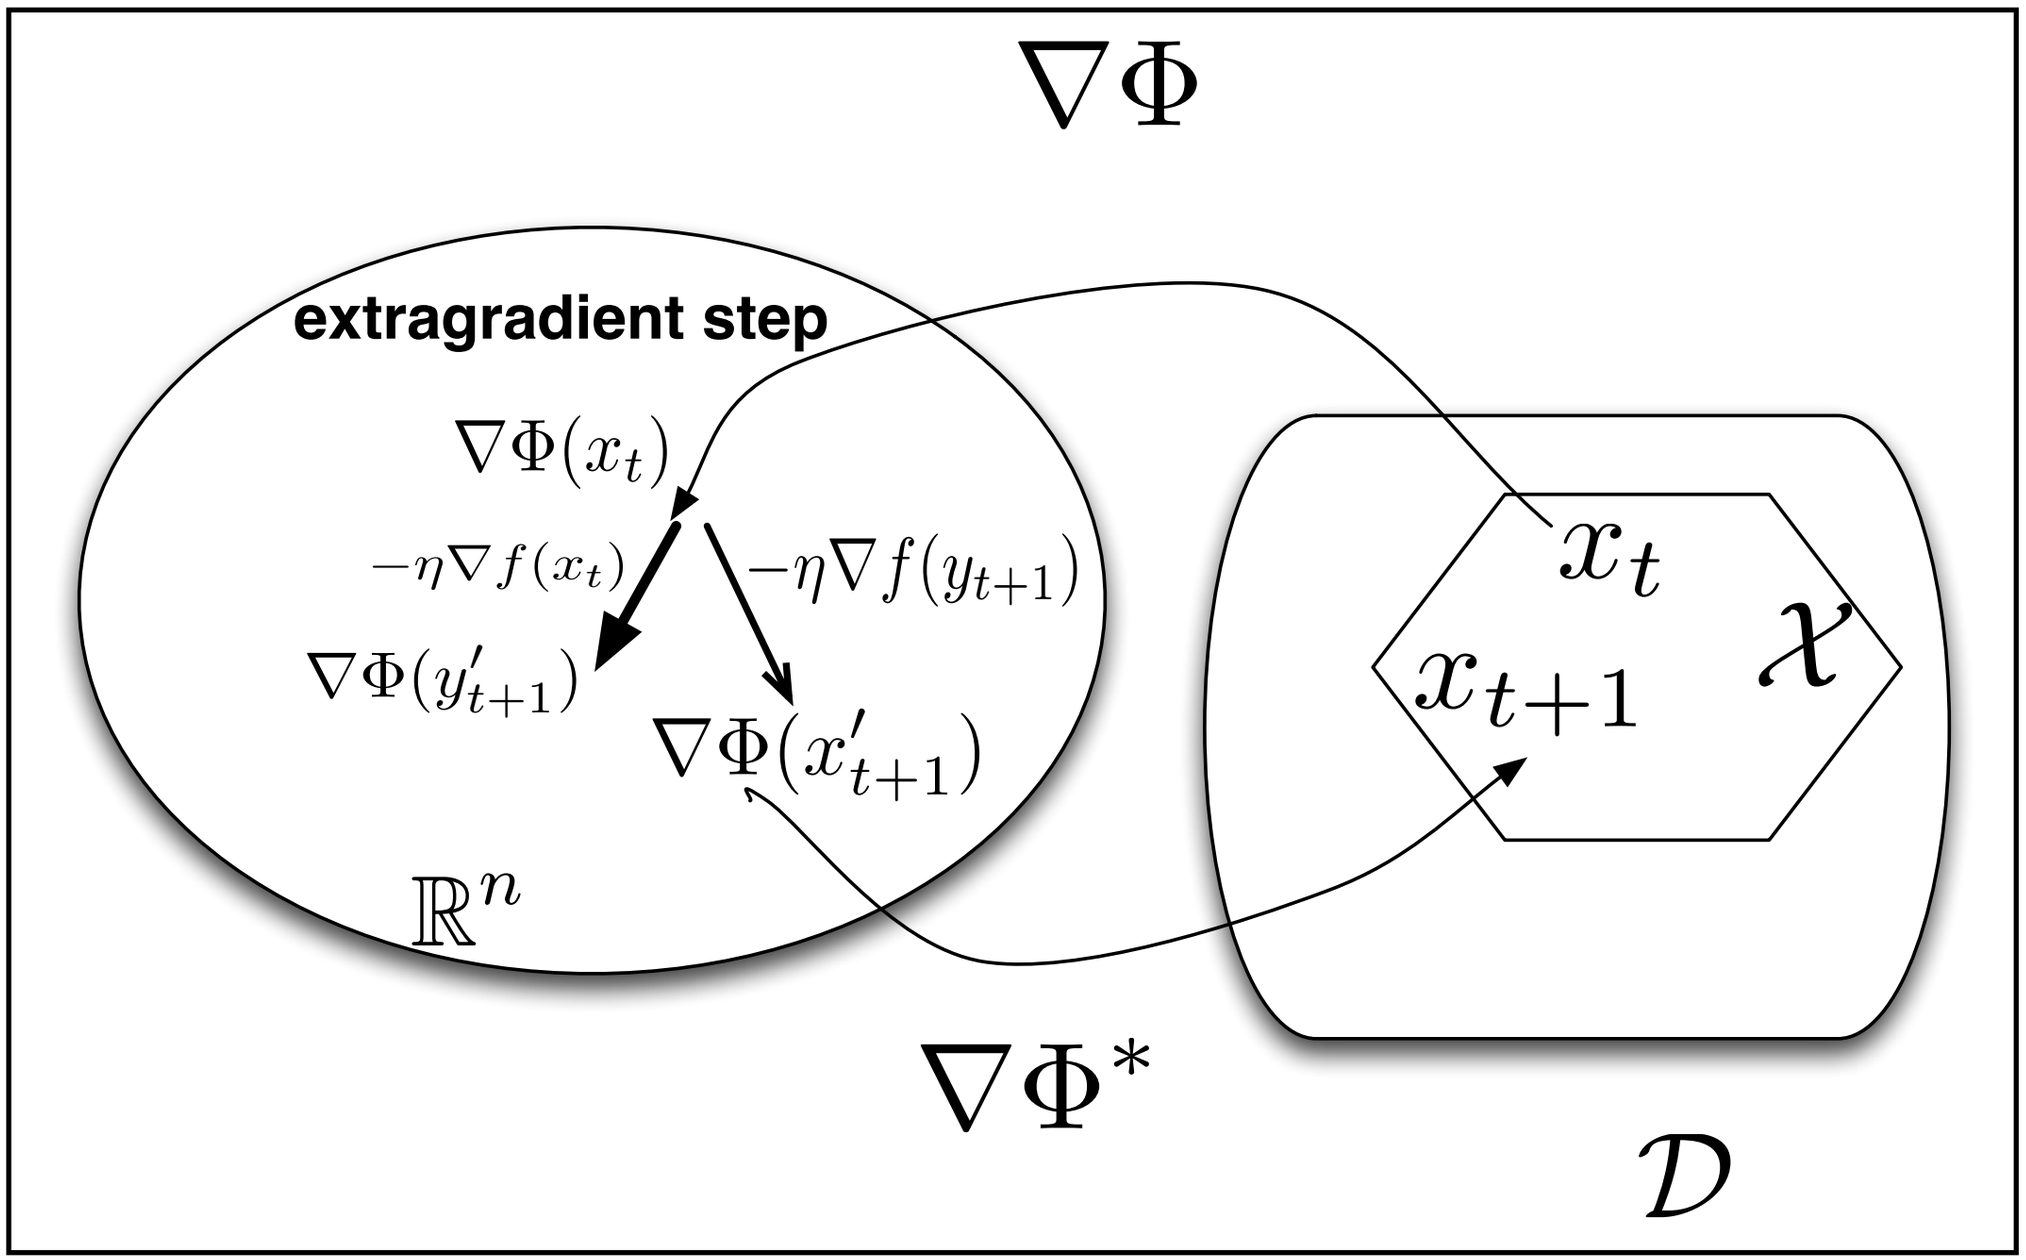

Mirror-Prox TD Method

(Mahadevan, 2014)

Algorithm 3 The Mirror-Prox TD(λ) algorithm.

repeat

Do action $\pi(s_t)$ and observe next state s_{t+1} and reward r_t.

Update the eligibility trace $e_{t+1} \leftarrow \lambda\gamma e_t + \phi(s_t)$

Compute $\delta_{t+\frac{1}{2}} = r_t + \gamma\phi(s_{t+1})^T w_t - \phi(s_t)^T w_t$

Update the dual weights θ_t:

$$\theta_{t+\frac{1}{2}} = \nabla\psi_t(w_t) + \alpha_t \delta_{t+\frac{1}{2}} e_{t+1}$$

where ψ is a distance generating function.

Update primal weights $w_{t+\frac{1}{2}} = \nabla\psi_t^*(\theta_{t+\frac{1}{2}})$.

$\delta_{t+1} = r_t + \gamma\phi(s_{t+1})^T w_{t+\frac{1}{2}} - \phi(s_t)^T w_{t+\frac{1}{2}}$

Update the dual weights θ_t:

$$\theta_{t+1} = \nabla\psi_t(w_{t+\frac{1}{2}}) + \alpha_t \delta_{t+\frac{1}{2}} e_{t+1}$$

where ψ is a distance generating function.

Update primal weights $w_{t+1} = \nabla\psi_t^*(\theta_{t+1})$.

Set $t \leftarrow t + 1$.

until convergence.

Return w_t

Experiment: Cart Pole

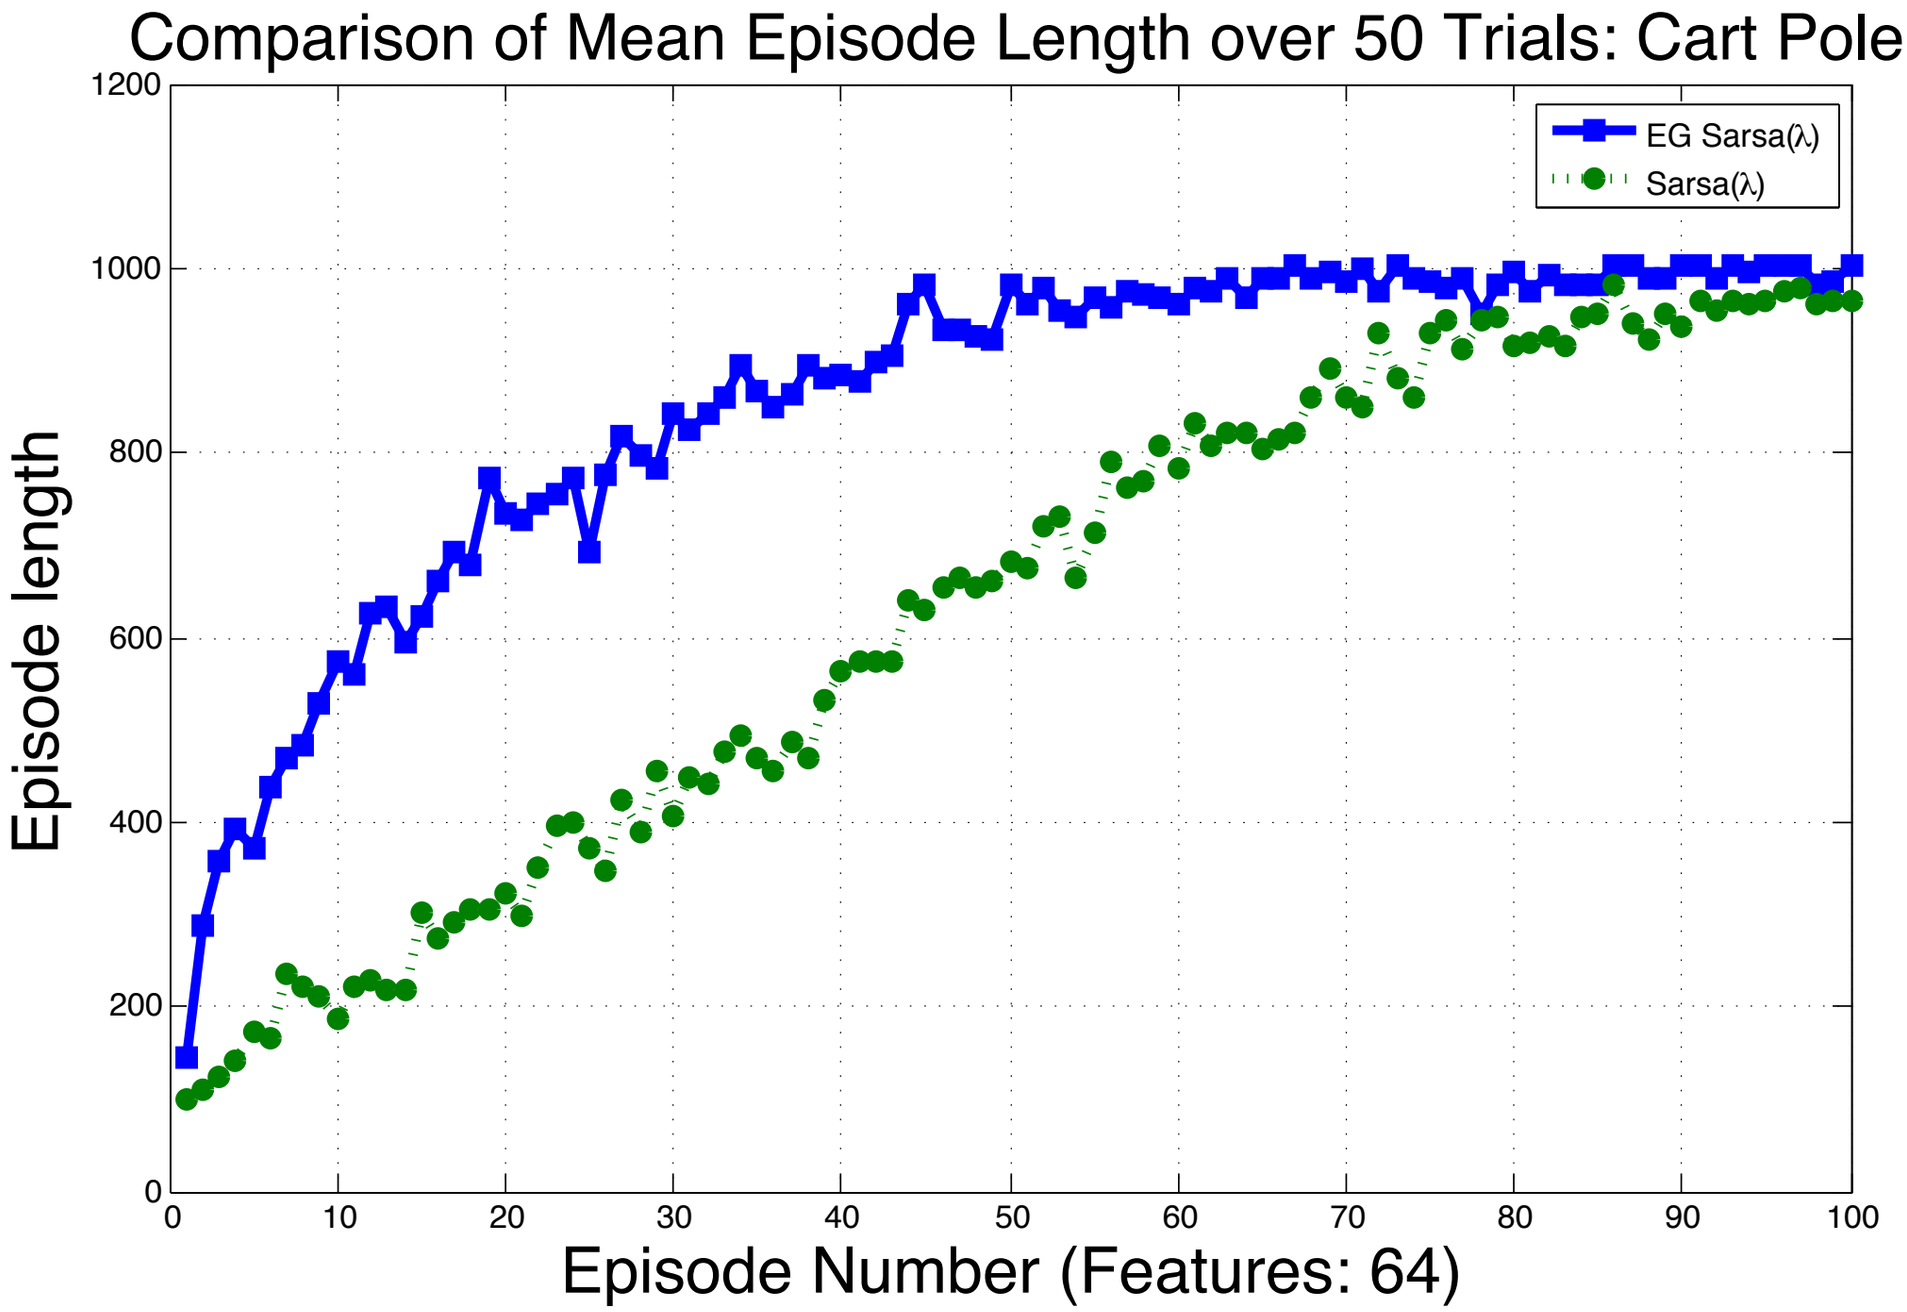

Experiment: Mountain Car

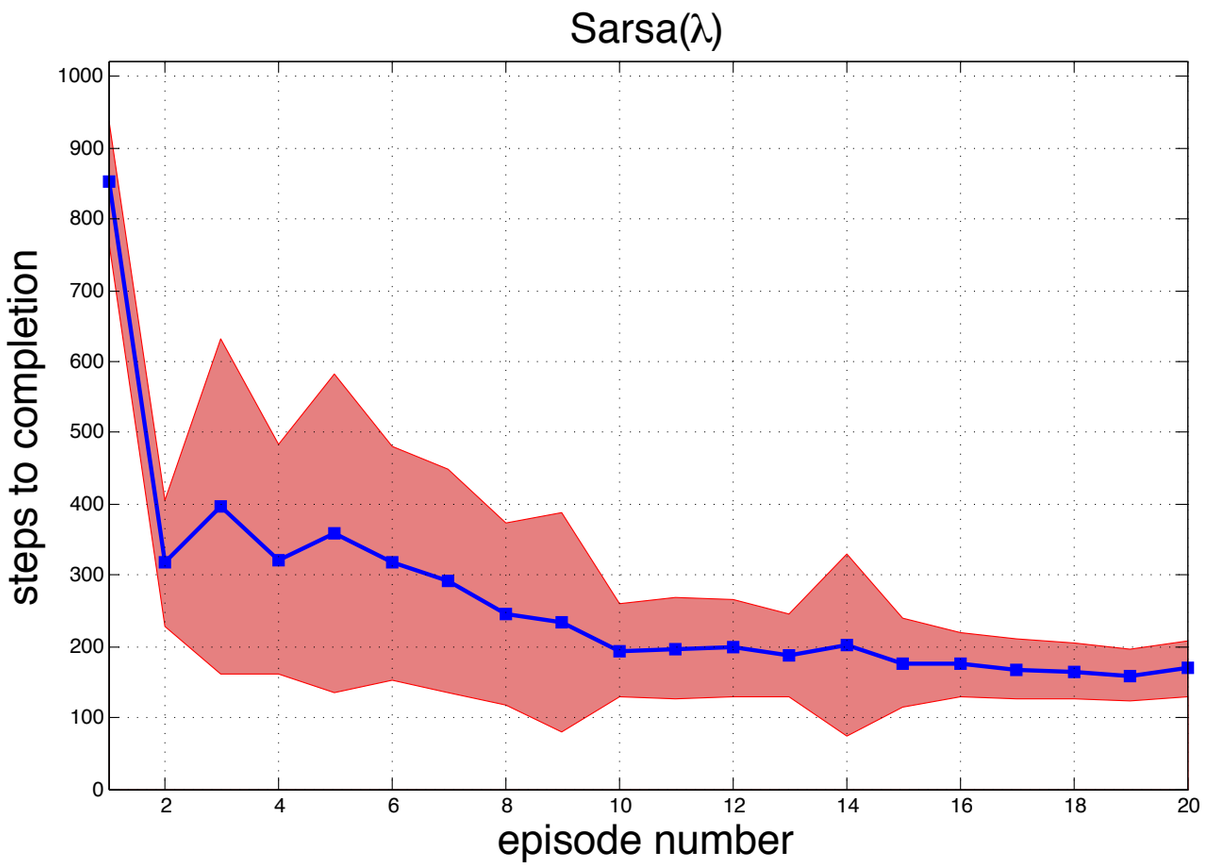

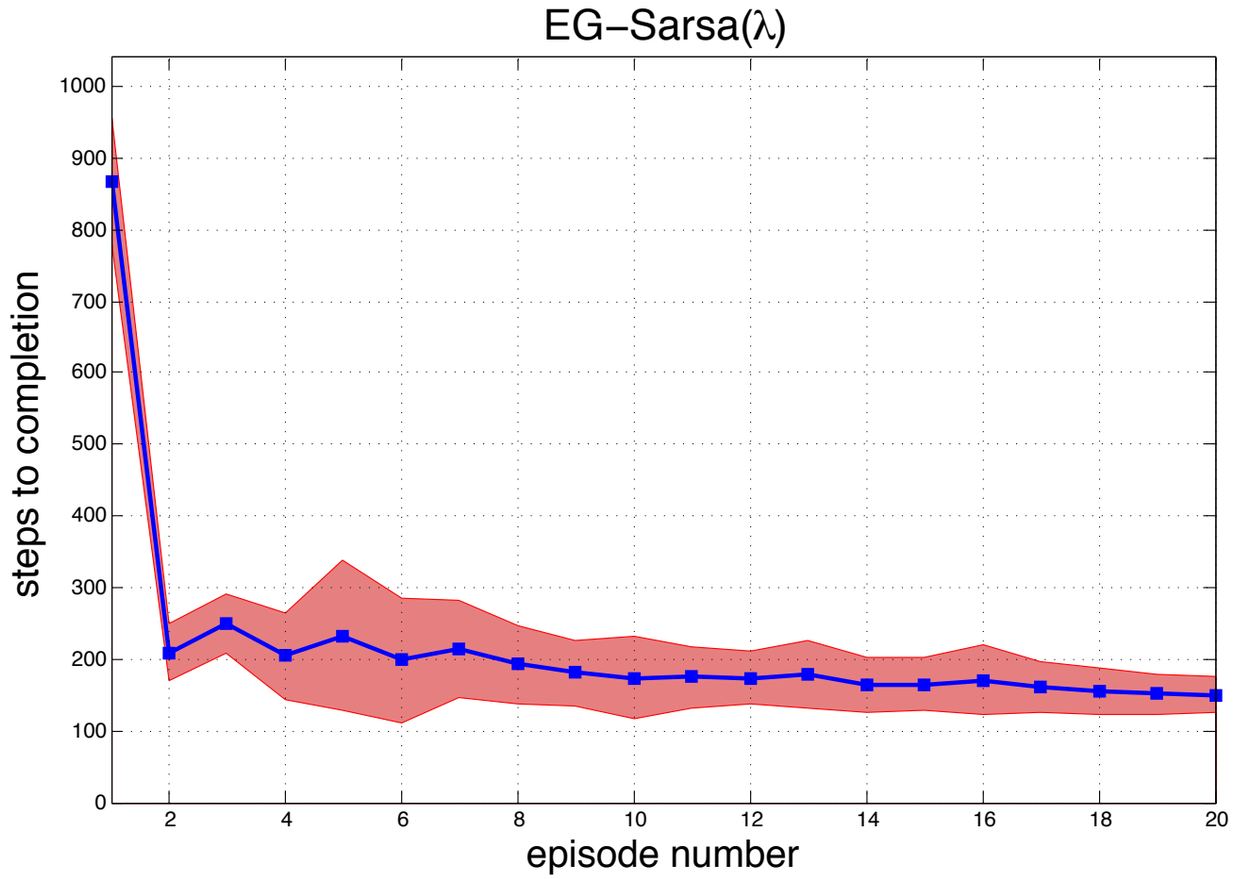

Mountain Car (16 features)

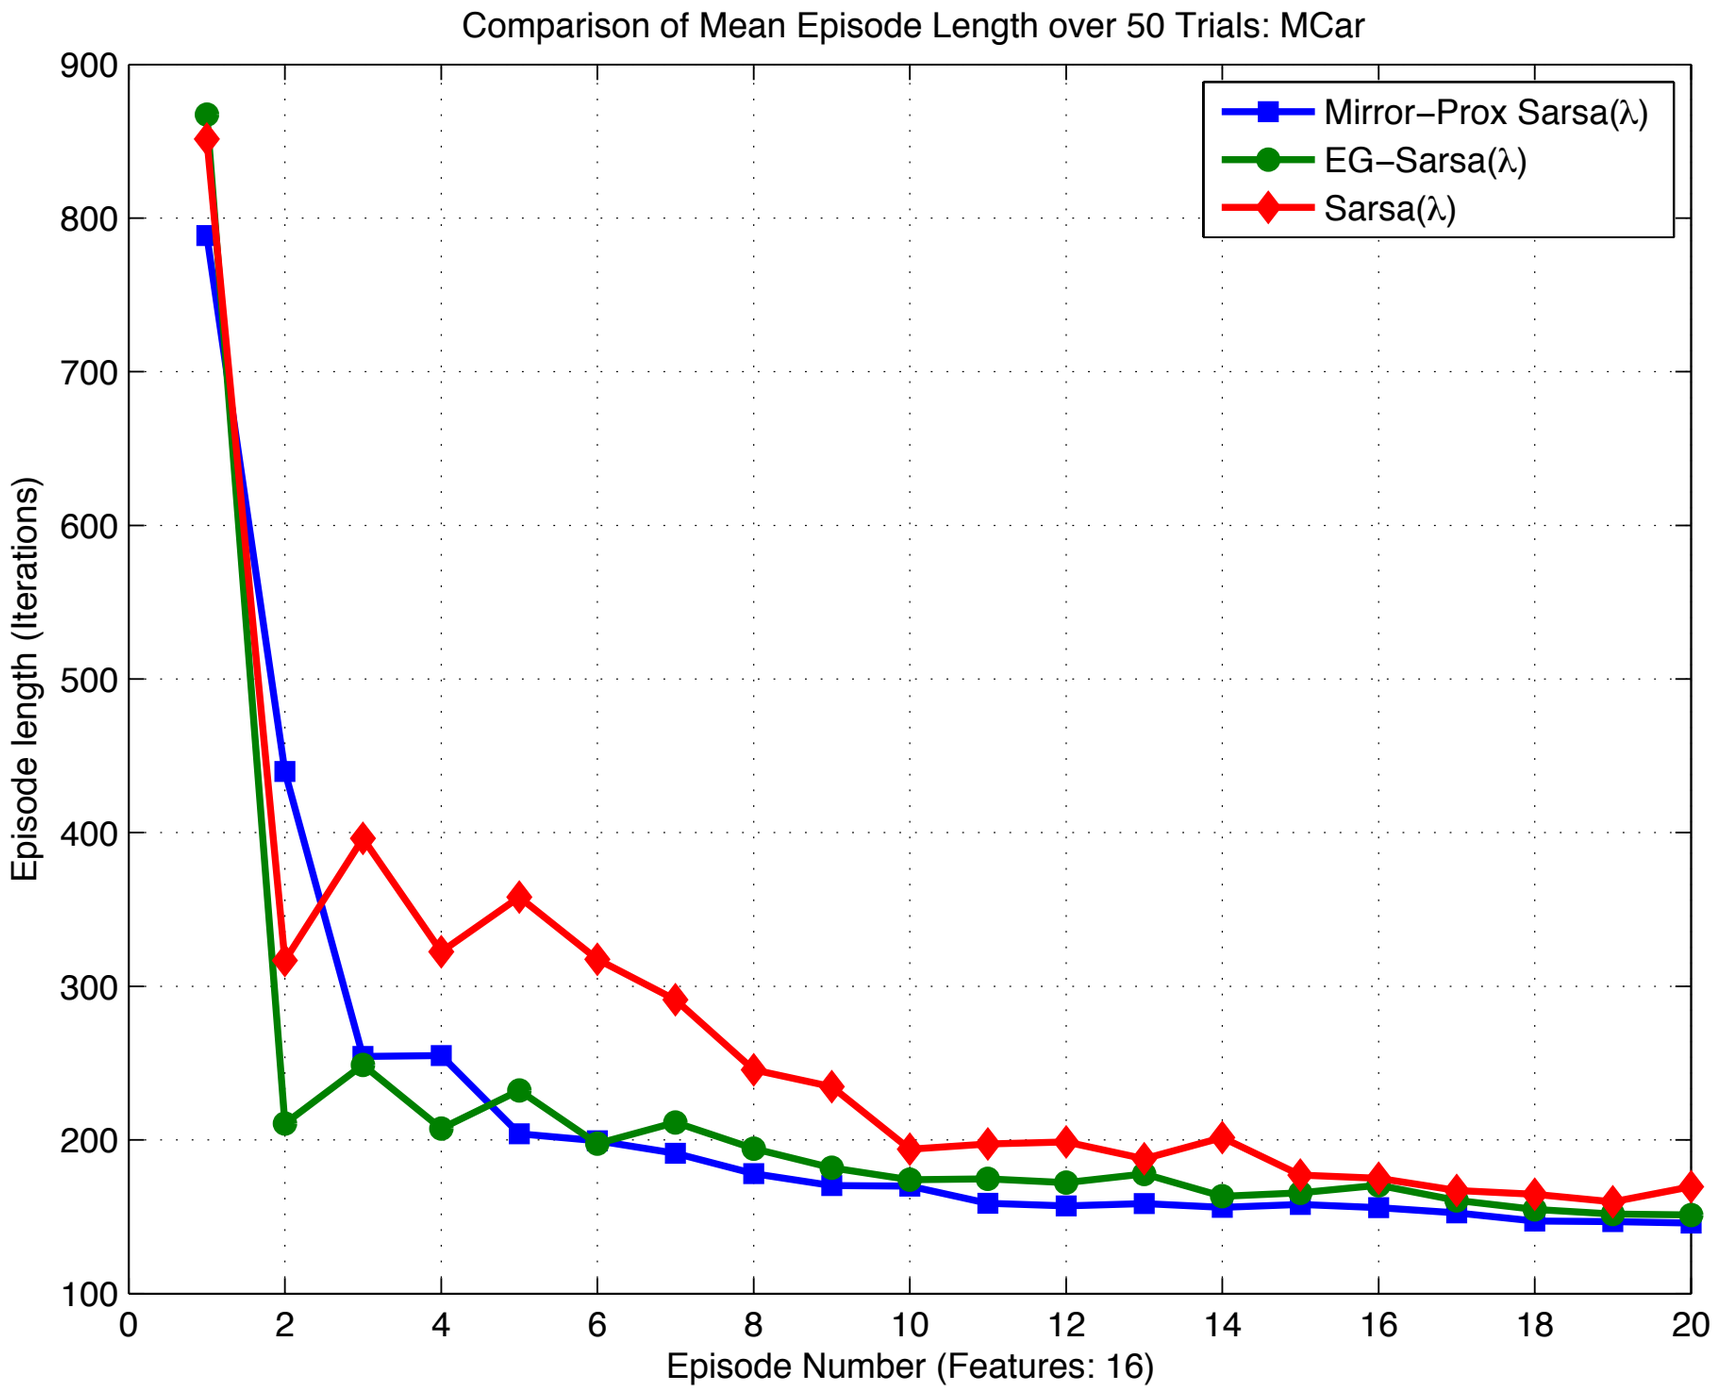

Experiment: Mountain Car

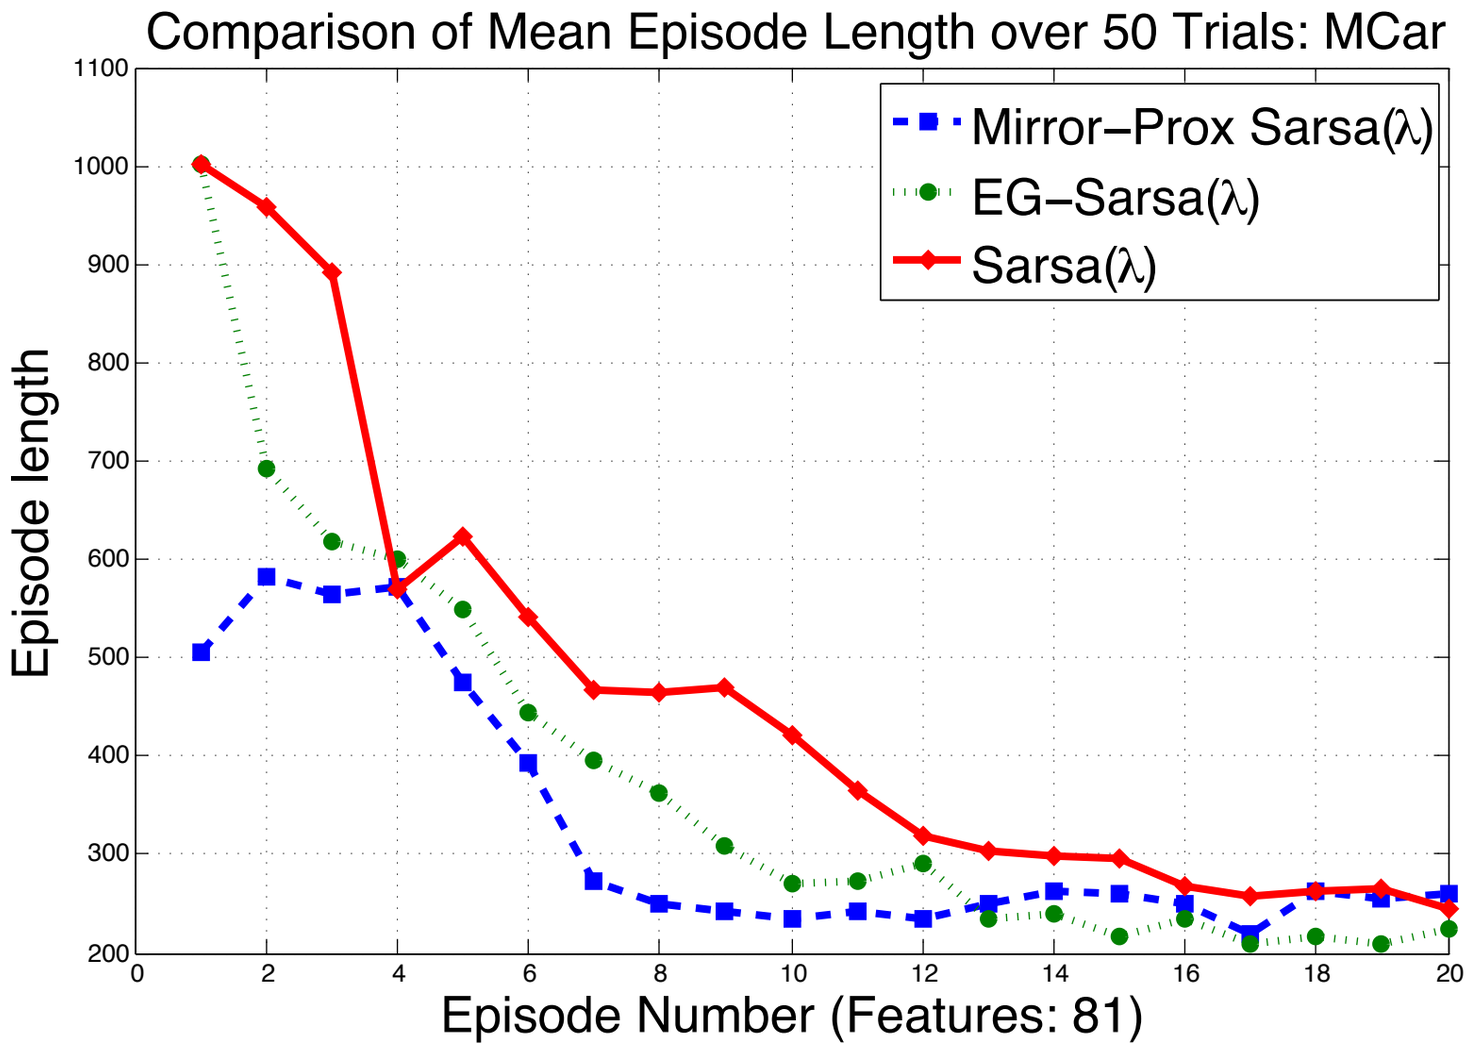

Experiment: Acrobot

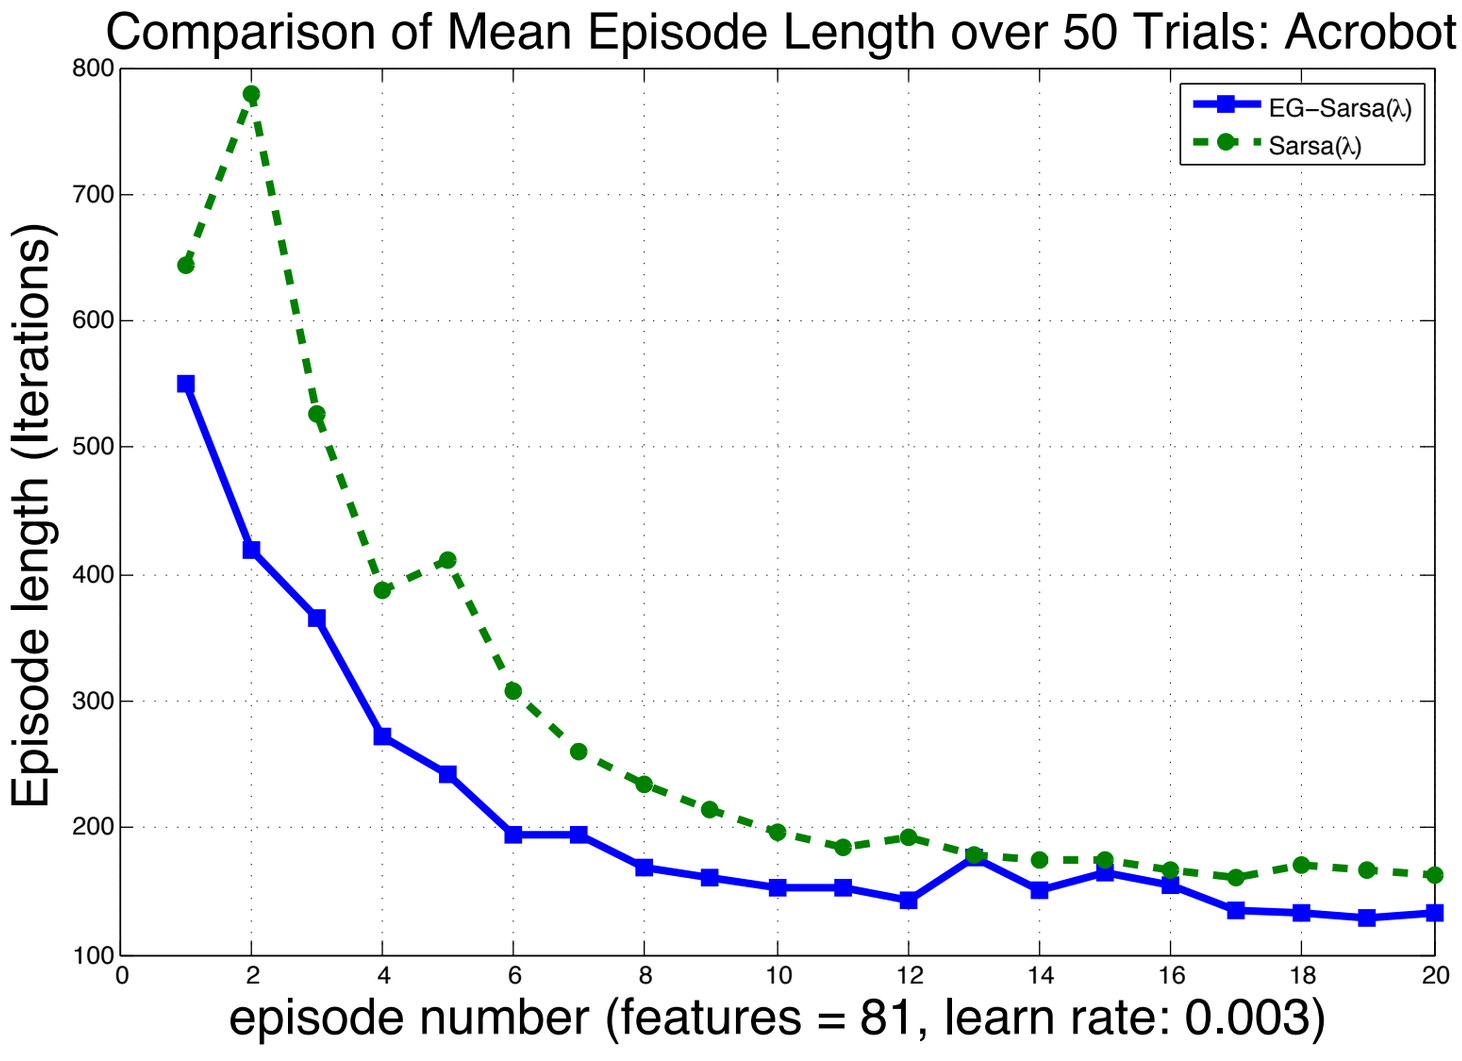

True Gradient TD-Learning: RL meets VI

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)

Baird
counter example

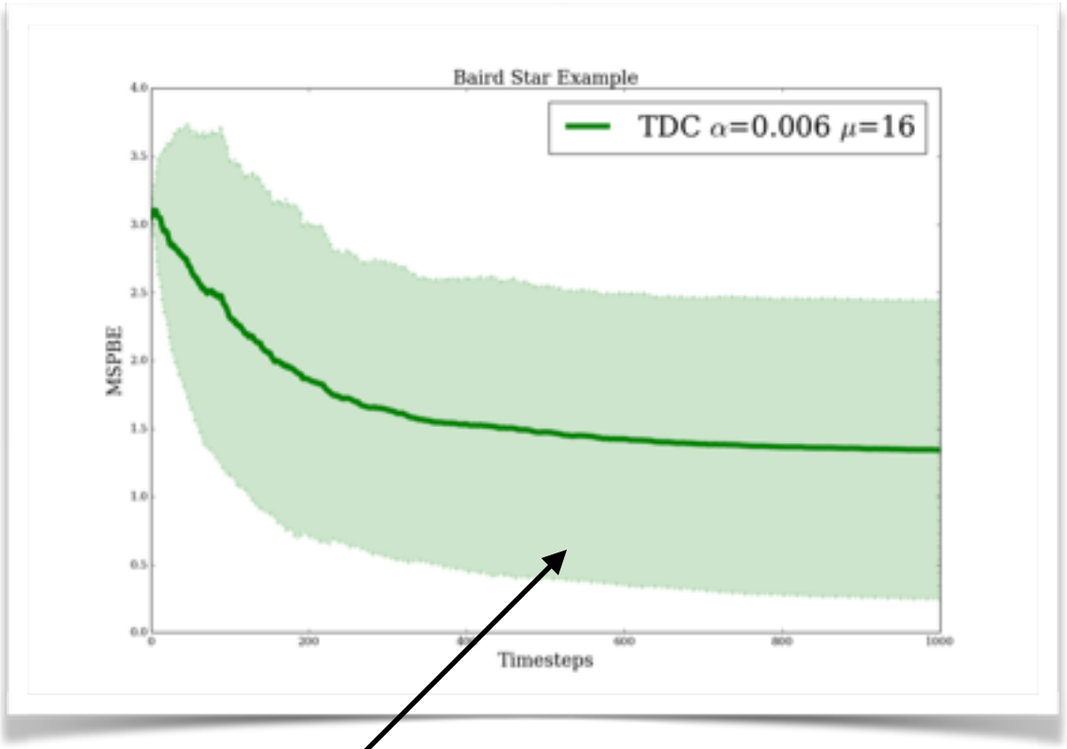

Sutton et al., 2009

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

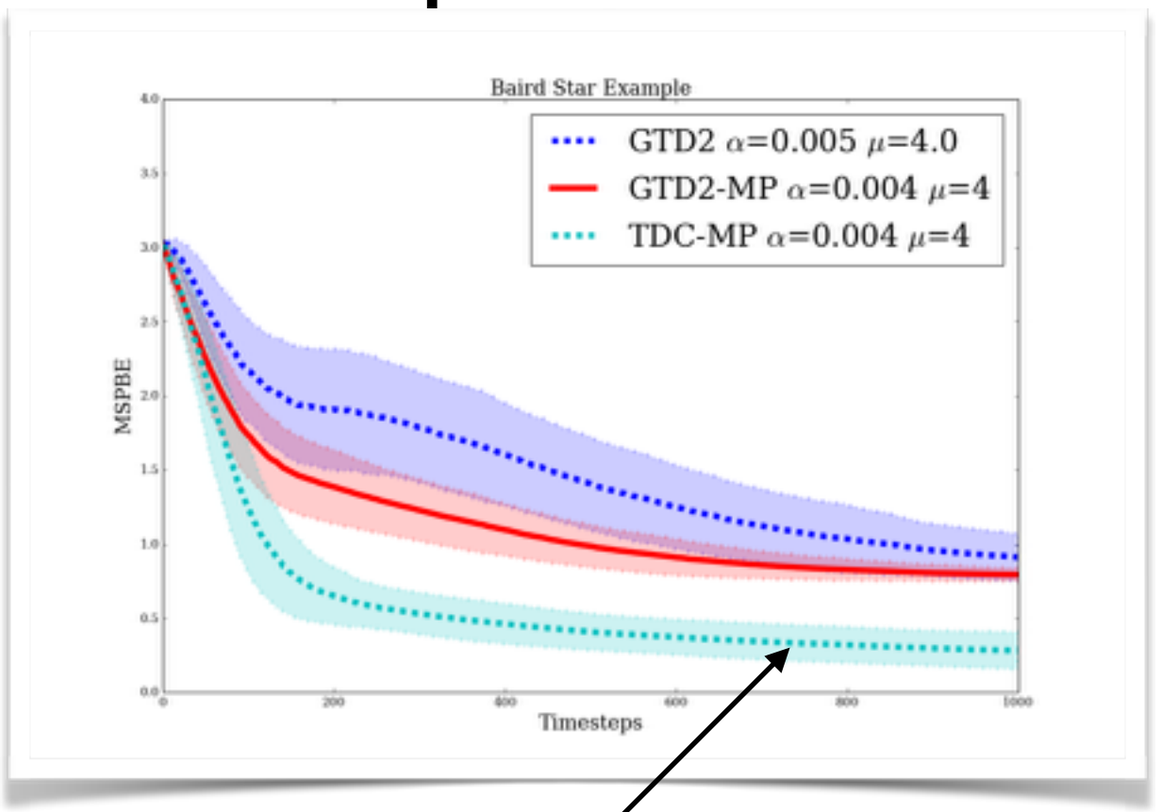

Our new methods
2014

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

20-Dimensional Robot Arm

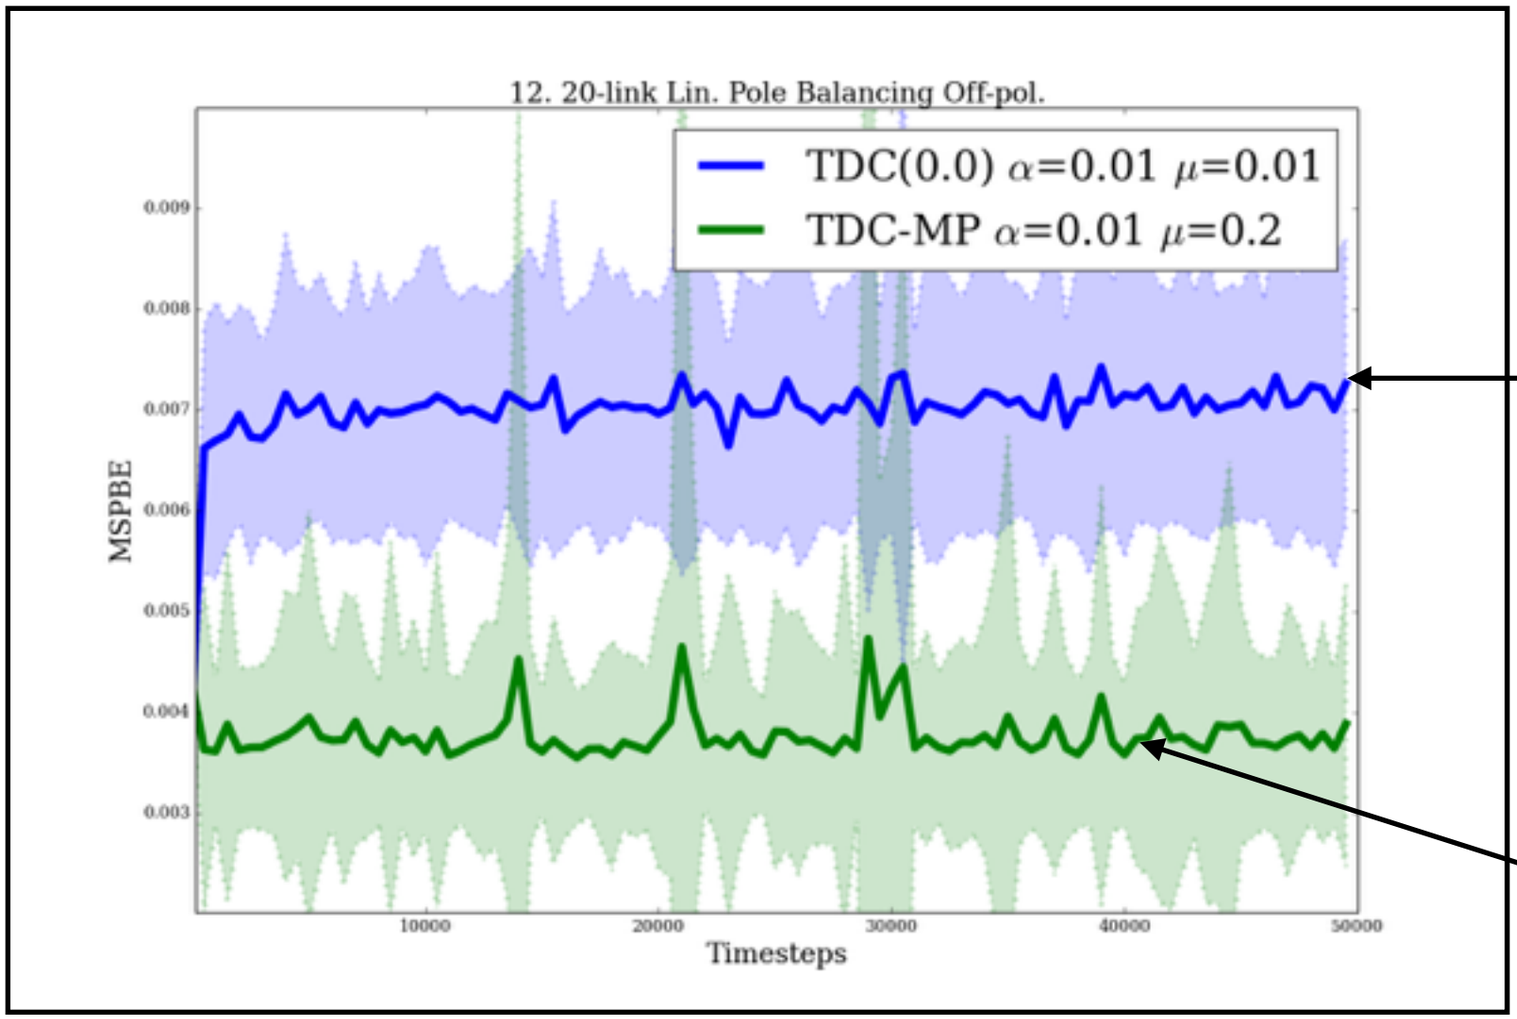

Sutton et al., 2009

Our new method

Safe Robot Learning

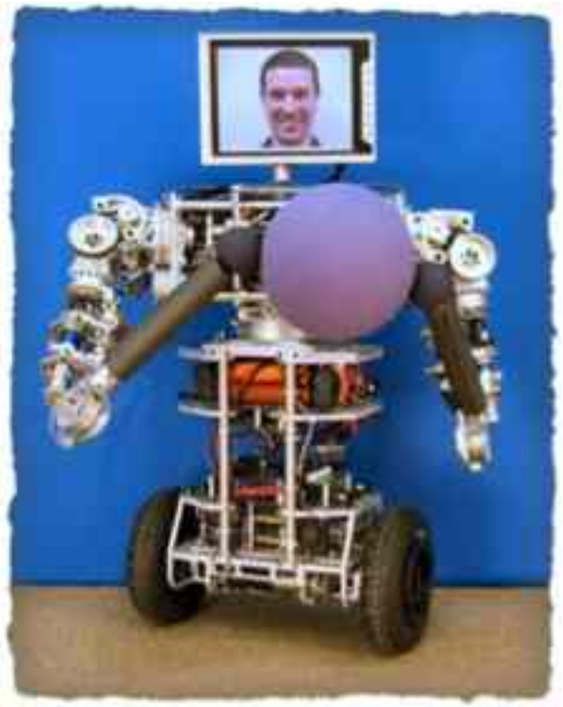

UBot, Laboratory of Perceptual Robotics

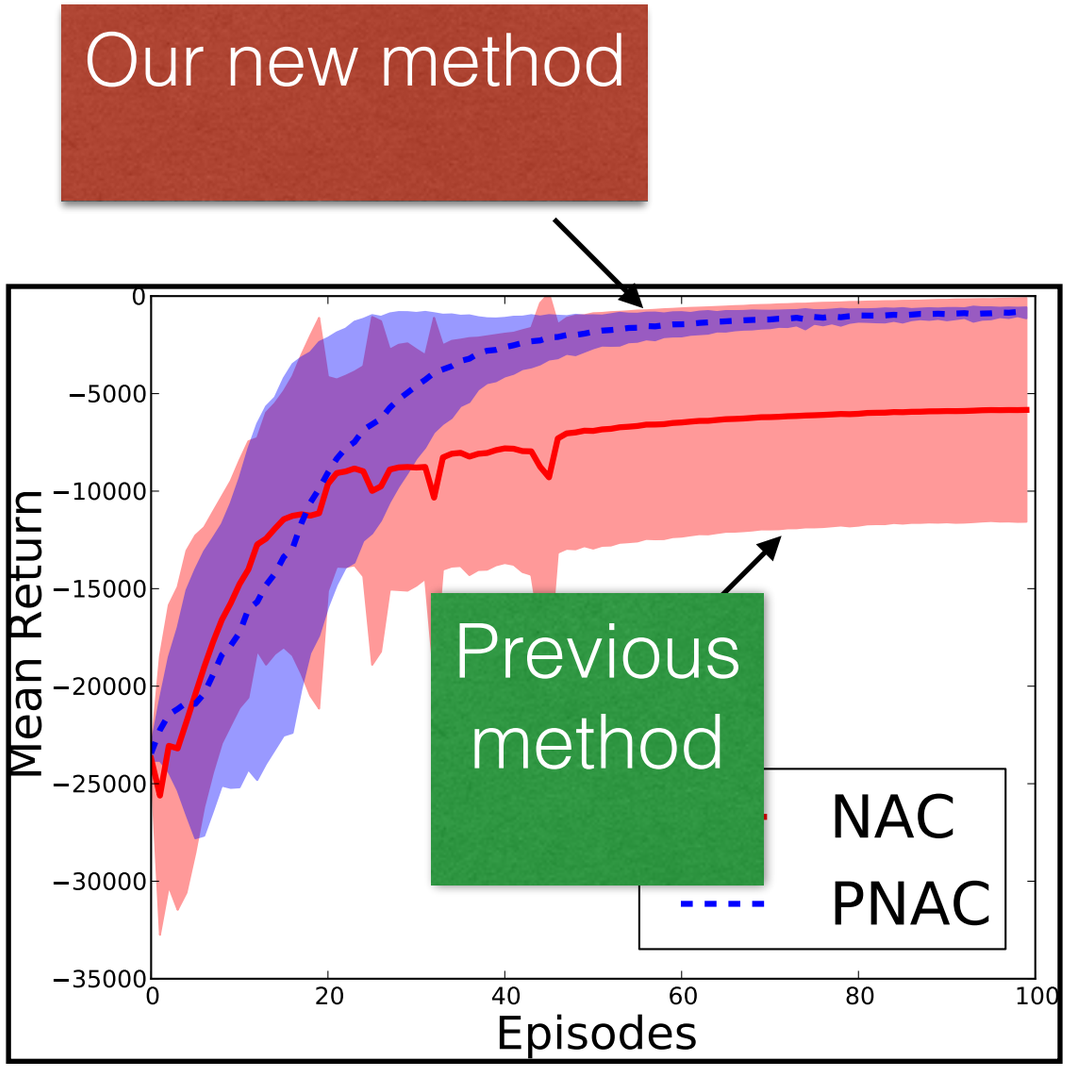

Thomas, Dabney, Mahadevan, Giguere, NIPS 2013

VIs applied to Natural Language Processing

Overview

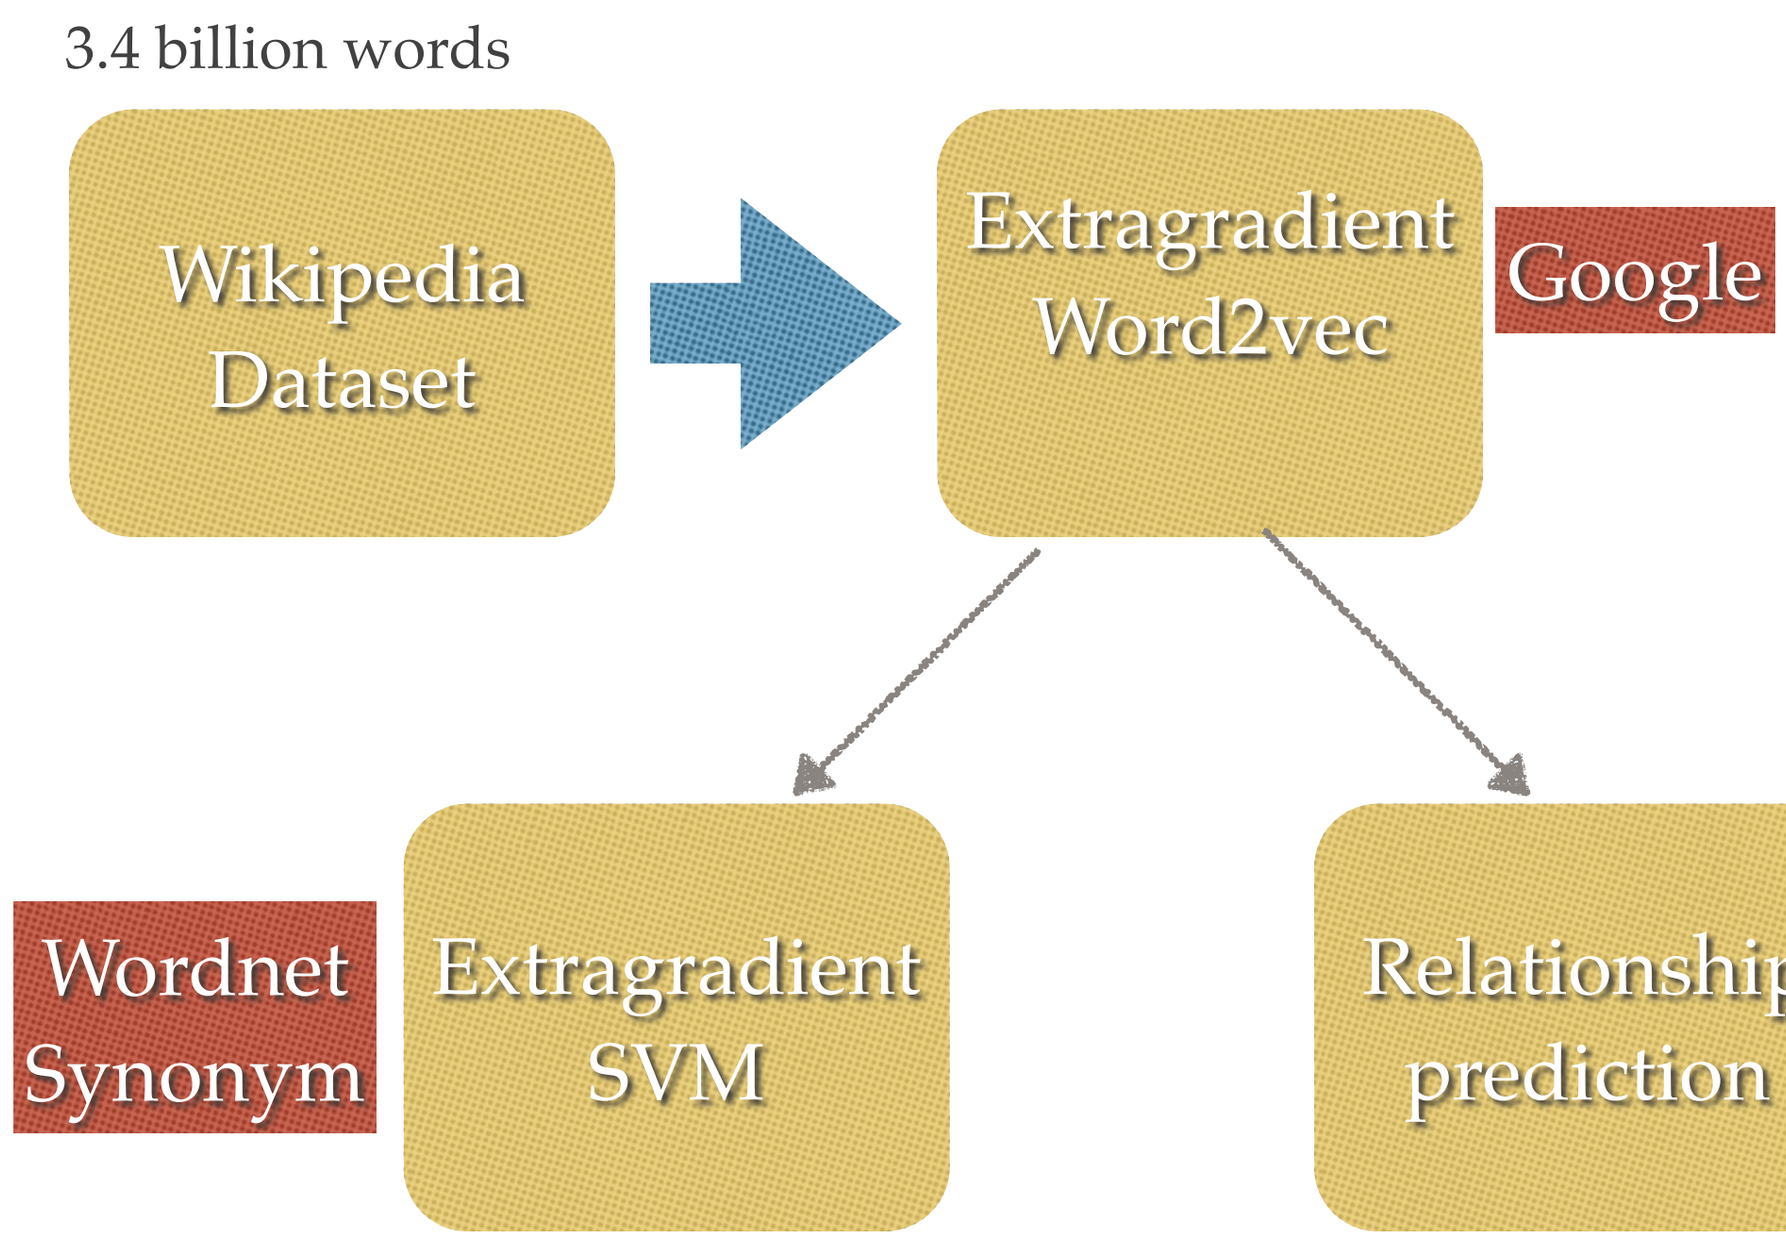
3.4 billion words
Wikipedia Dataset
Extragradient Word2vec
Google
Wordnet Synonym
Extragradient SVM
Relationship prediction

word2vec

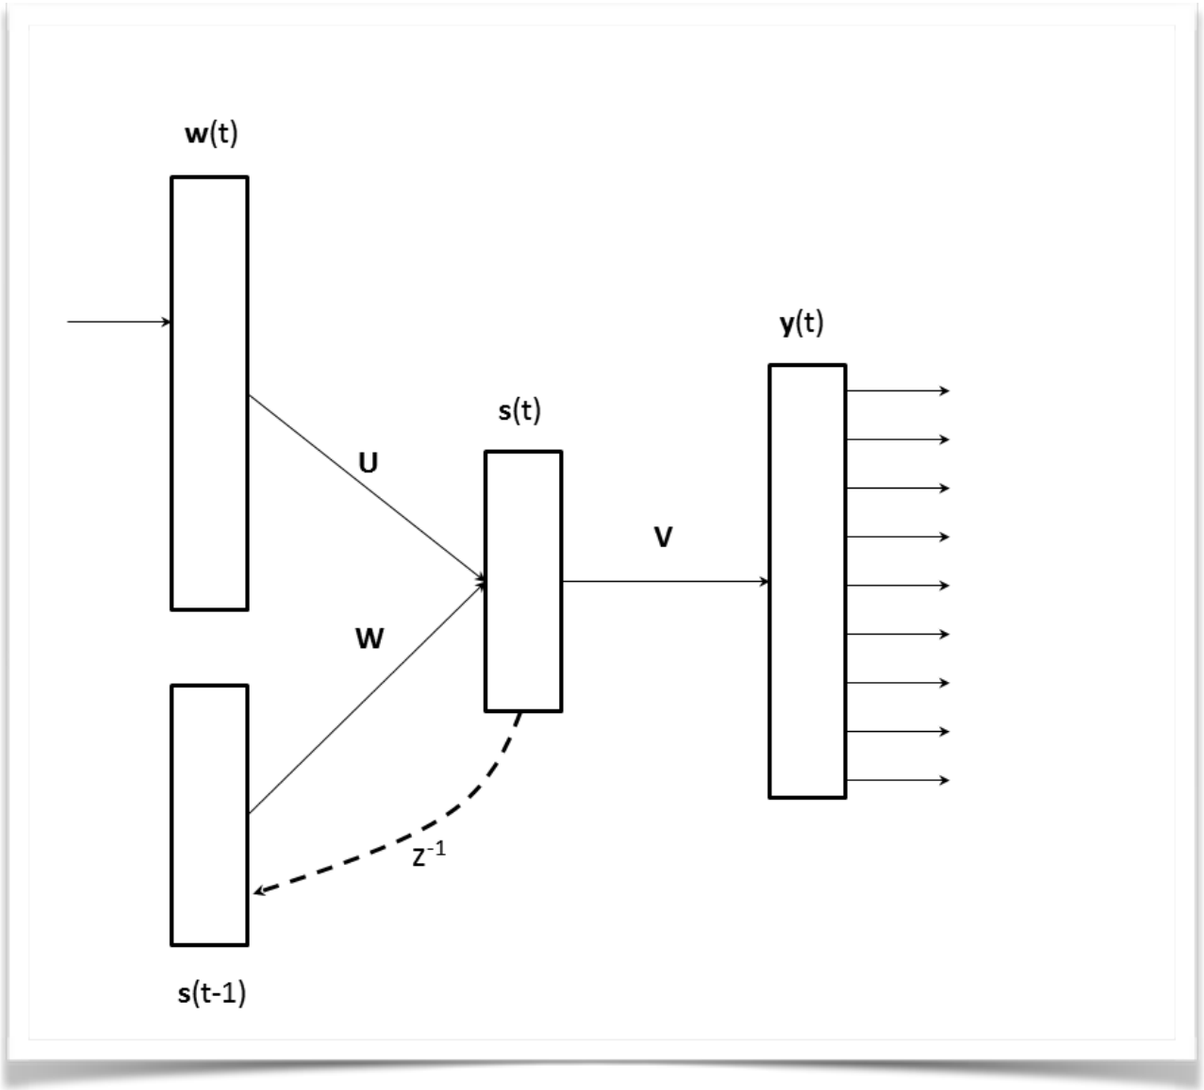

$$\mathbf{s}(t) = f\left(\mathbf{U}\mathbf{w}(t) + \mathbf{W}\mathbf{s}(t-1)\right) \qquad (1)$$

$$\mathbf{y}(t) = g\left(\mathbf{V}\mathbf{s}(t)\right), \qquad (2)$$

where

$$f(z) = \frac{1}{1+e^{-z}}, \quad g(z_m) = \frac{e^{z_m}}{\sum_k e^{z_k}}. \qquad (3)$$

Linguistic Regularities in Continuous Space Word Representations

Tomas Mikolov*, Wen-tau Yih, Geoffrey Zweig
Microsoft Research
Redmond, WA 98052

Synonym or Antonym?

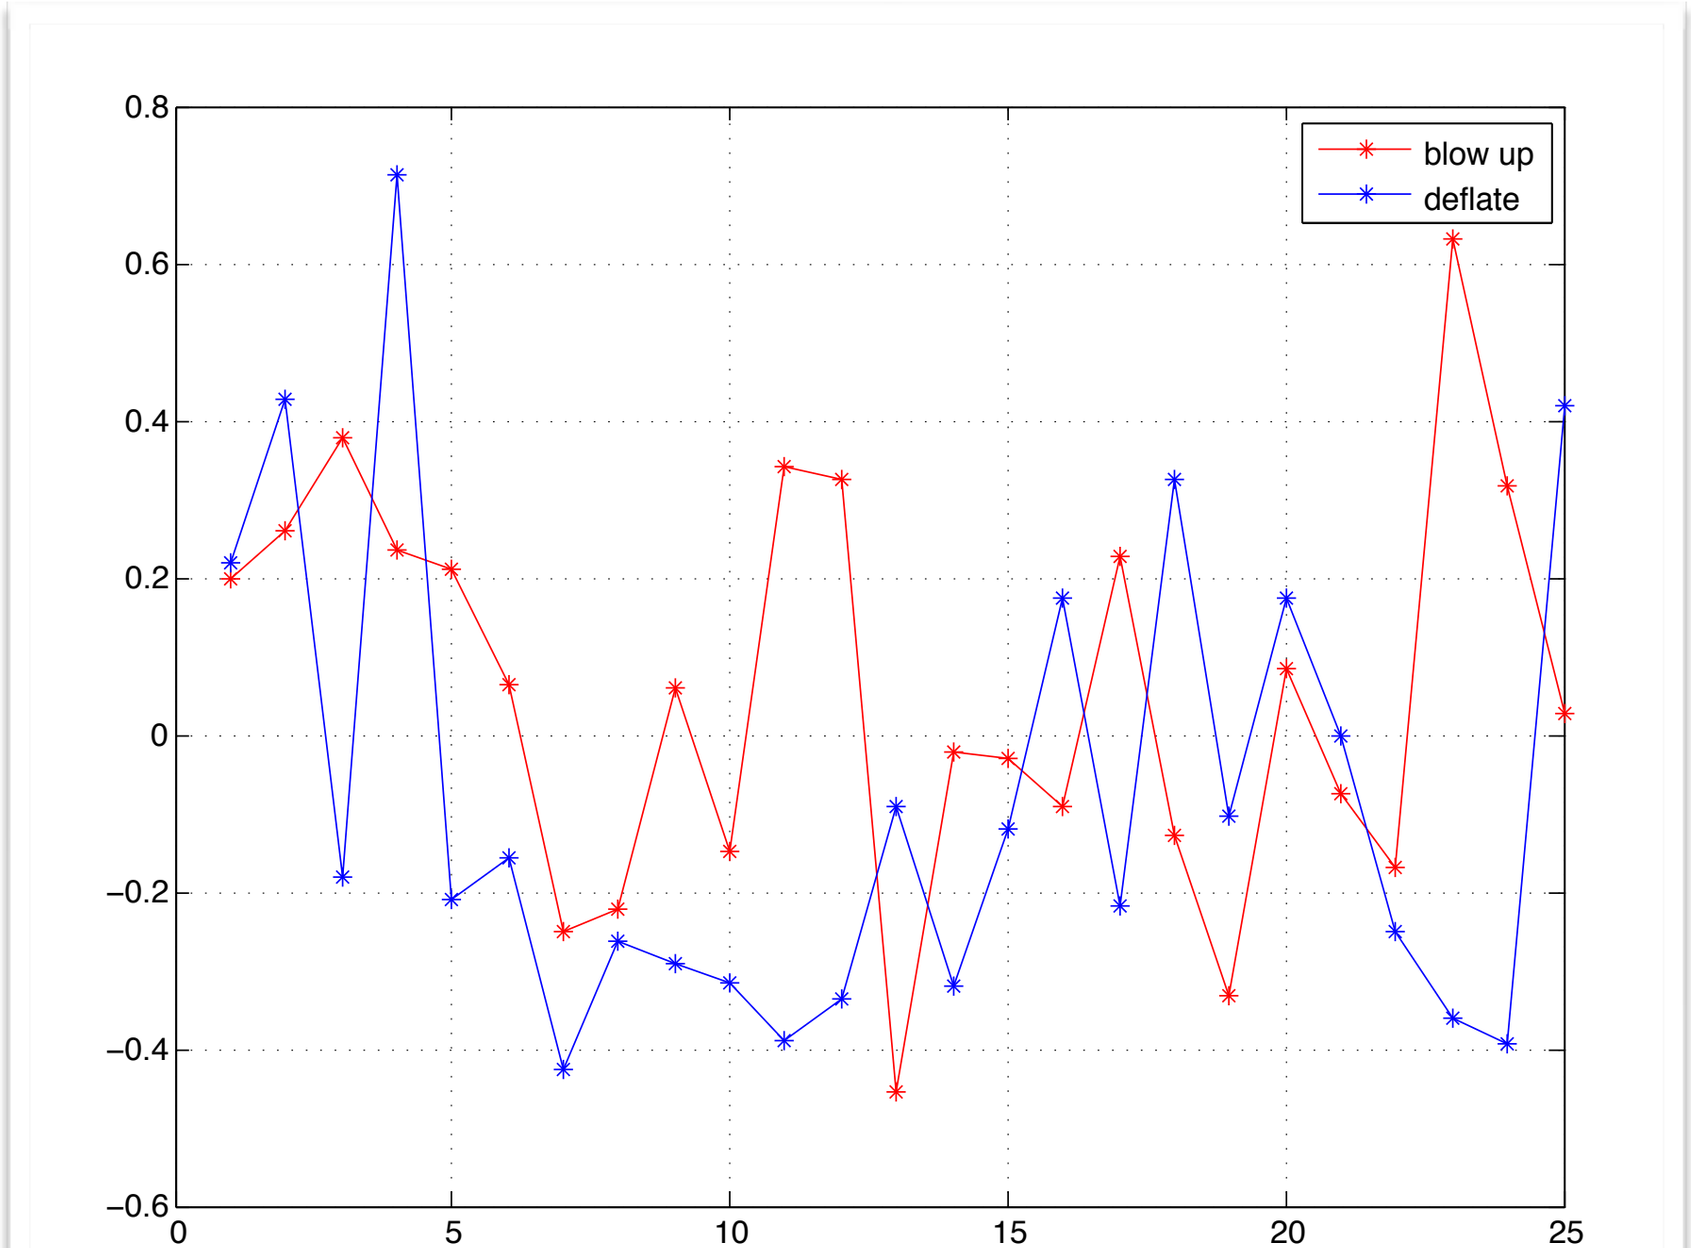

Linguistic Regularities

Google Dataset
Mikolov, 2013

a is to a*
as
b is to b*

Relation
capital-common-countries
capital-world
city-in-state
currency
family (gender inflections)
gram1-adjective-to-adverb
gram2-opposite
gram3-comparative
gram4-superlative
gram5-present-participle
gram6-nationality-adjective
gram7-past-tense
gram8-plural (nouns)
gram9-plural-verbs

Linguistic Reasoning by Vector Arithmetic

$$queen \approx king - man + woman$$

$$\arg\max_{b^* \in V} \left(\cos\left(b^*, b - a + a^*\right)\right)$$

Mikolov et al., 2013

Stochastic Extragradient SVM

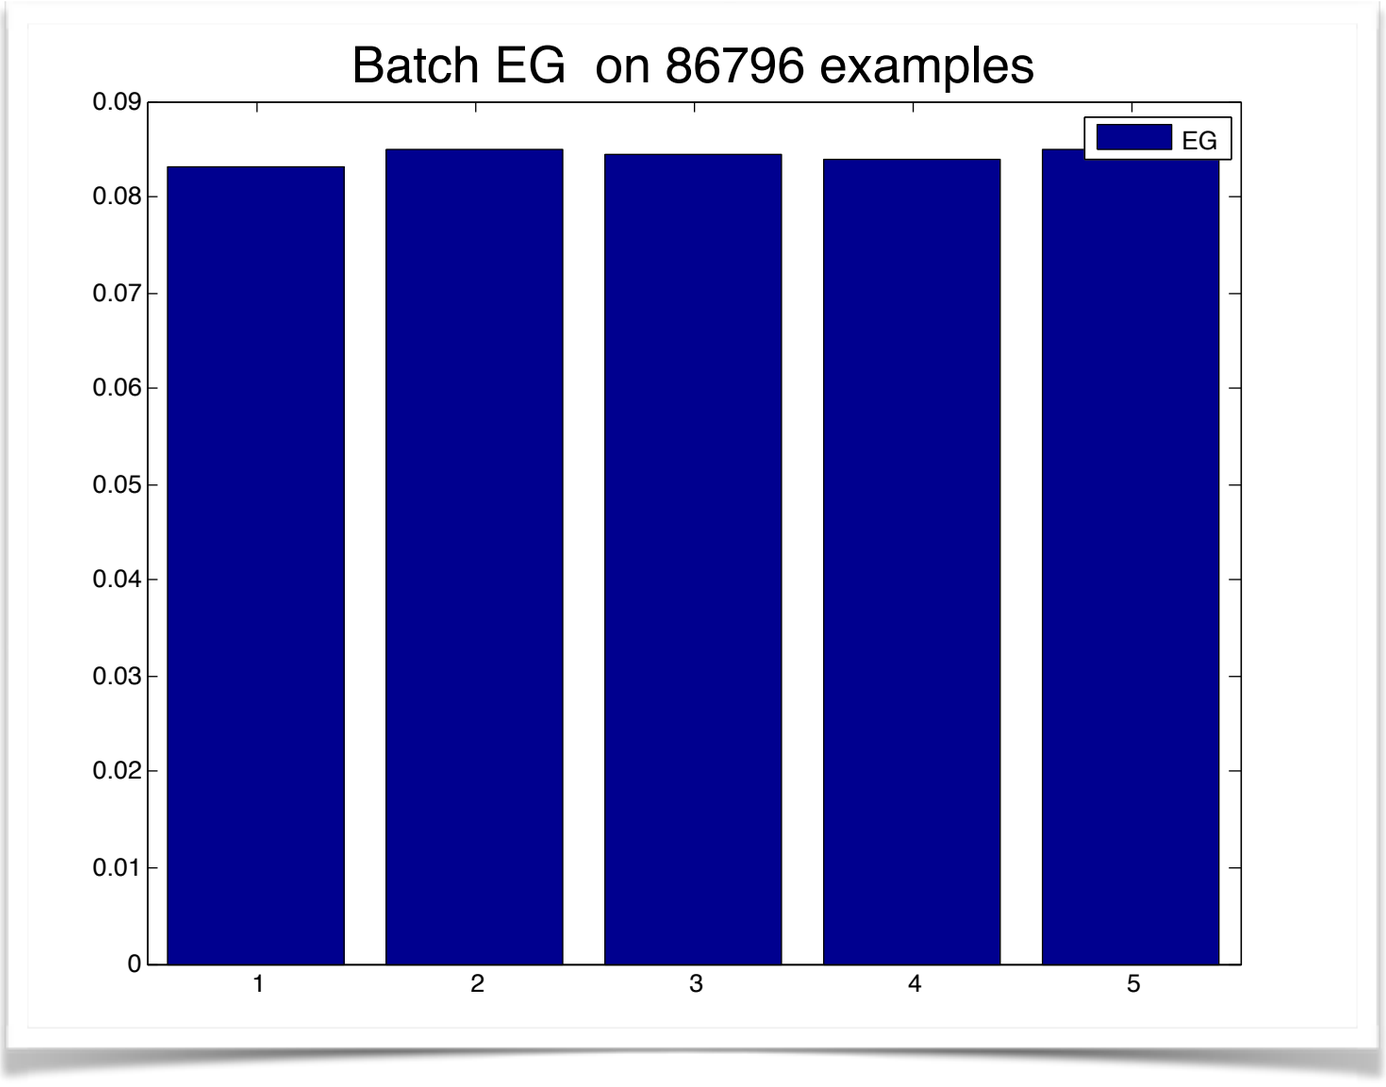

Accuracy: ~92%

Results using EG SVM

Summary

- VIs and PDS provide a new foundation for 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

Questions?

