

Approximation in Multiobjective Optimization

Ilias Diakonikolas
University of Edinburgh

*Horizons in TCS,
Princeton, August 2013*

Outline

- Introduction
- Approximate Pareto Sets
- Approximate *Convex* Pareto Sets
- Research Directions

Outline

- Introduction
- Approximate Pareto Sets
- Approximate *Convex* Pareto Sets
- Research Directions

Multiobjective Optimization

Decision making with multiple criteria (objectives).


Evaluate different solutions from a design space and pick the “best” one according to the criteria of interest.

Arises in many areas: economics, management, engineering, healthcare, biology, etc

Only minimization objectives for this talk.

Example (Bi-objective Shortest Path)

- Graph $G=(V, E)$. Each edge e has *length* $l(e)$ and *cost* $c(e)$. Find the “shortest, cheapest” $s - t$ path.




2^5 $s - t$ paths.

Many of them
incomparable.

Example (Bi-objective Shortest Path)

- Graph $G=(V, E)$. Each edge e has *length* $l(e)$ and *cost* $c(e)$. Find the “shortest, cheapest” $s-t$ path.




2^5 $s-t$ paths.
Many of them
incomparable.

- Decision space: \mathcal{Z} subset of $\{0,1\}^{|E|}$
(characteristic vectors of $s-t$ paths)

- Objectives: $l, c \in \mathbb{Q}_+^{|E|}$

- Objective space:

$$\mathcal{X} = \{(x, y) \in \mathbb{R}_+^2 \mid x = l \cdot z, y = c \cdot z, z \in \mathcal{Z}\}$$



Example (Bi-objective LP)

Minimize two linear functions subject to a set of linear constraints.

- Decision space: $\mathcal{Z} = \{ z \in \mathbb{R}^{n \times 1} \mid A \cdot z \geq b, z \geq \mathbf{0}_{n \times 1} \}$
where $A \in \mathbb{Q}^{m \times n}$ and $b \in \mathbb{Q}^{1 \times m}$

Example (Bi-objective LP)

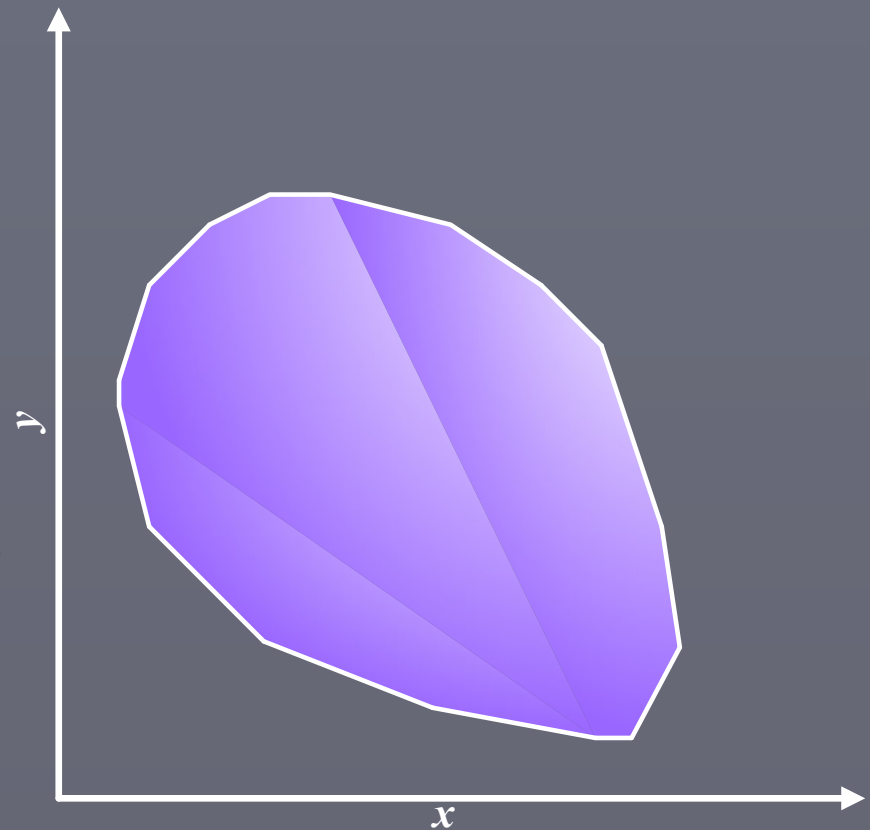
Minimize two linear functions subject to a set of linear constraints.

- Decision space: $\mathcal{Z} = \{ z \in \mathbb{R}^{n \times 1} \mid A \cdot z \geq b, z \geq \mathbf{0}_{n \times 1} \}$
where $A \in \mathbb{Q}^{m \times n}$ and $b \in \mathbb{Q}^{1 \times m}$

- Objectives: $c, d \in \mathbb{Q}^{1 \times n}$

- Objective Space:

$$\mathcal{X} = \{ (x, y) \in \mathbb{R}^2_+ \mid x = c \cdot z, y = d \cdot z, z \in \mathcal{Z} \}$$



Multiobjective Optimization

Active research area with many applications in various diverse disciplines (economics, management, engineering, healthcare, etc).

What does it mean to solve such a problem?

- One approach: Treat as single objective.
- Alternative approach: Pareto Set.

Multiobjective Optimization

Active research area with many applications in various diverse disciplines (economics, management, engineering, healthcare, etc).


What does it mean to solve such a problem?

- One approach: Treat as single objective.
- Alternative approach: **Pareto Set**.

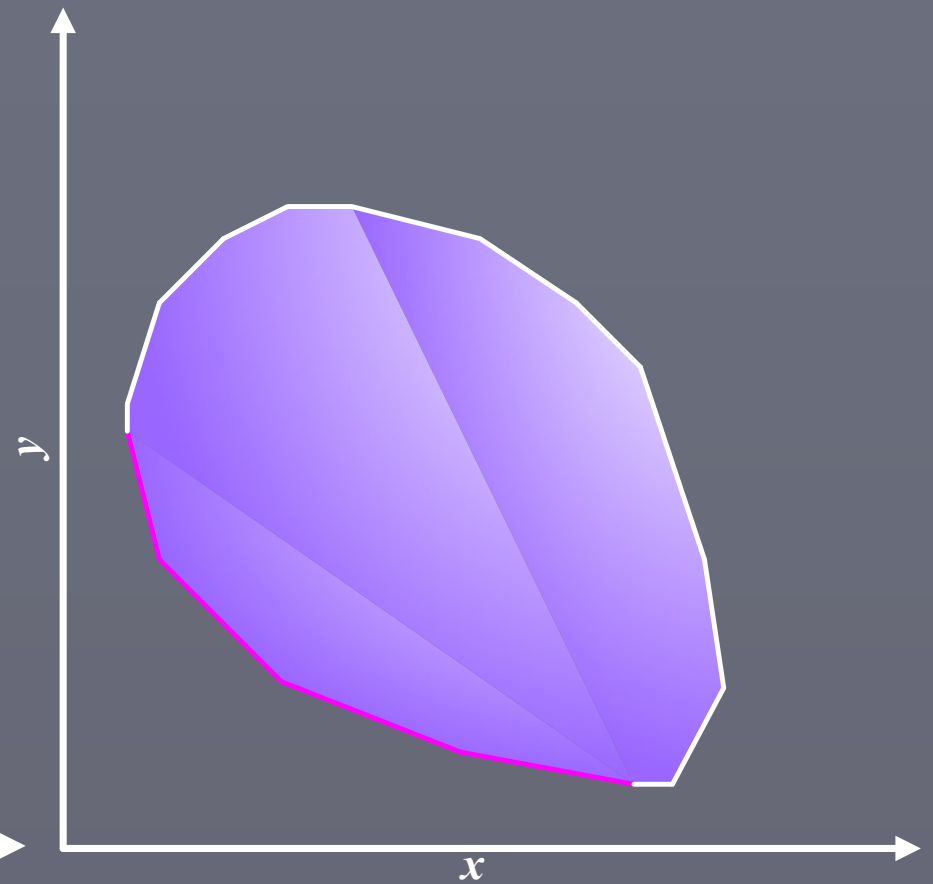
“The set of *undominated* solution points in the objective space.”

Pareto Set

Discrete Space



Convex Space



Pareto Set (I)

- Represents “reasonable optimal choices” in the objective/design space.
- Contains optimal solutions for all possible combining functions.
- Decision maker can choose based on preferences.

Pareto Set (II)

But...

- Exponentially large (or infinite) even for two objectives.
- NP-hard to decide whether a point is in the Pareto set.

“Representative” Approximation

In practice, some kind of approximation is computed.

Underlying goal:

Efficiently compute

a “*good*” approximation with “*few*” points.

Examples

- **Networking:** Network routing with multiple *QoS* criteria.
- **Databases:** “Skyline Query”
(google scholar: “skyline, databases” returns ~10300 results)
- **Healthcare:** [Craft *et al.*--Medical Physics '06] Approximating convex Pareto surfaces in multiobjective radiotherapy planning.
- **Computer Aided Design,**
- etc.

Outline

- Introduction
- **Approximate Pareto Sets**
- Approximate *Convex* Pareto Sets
- Research Directions

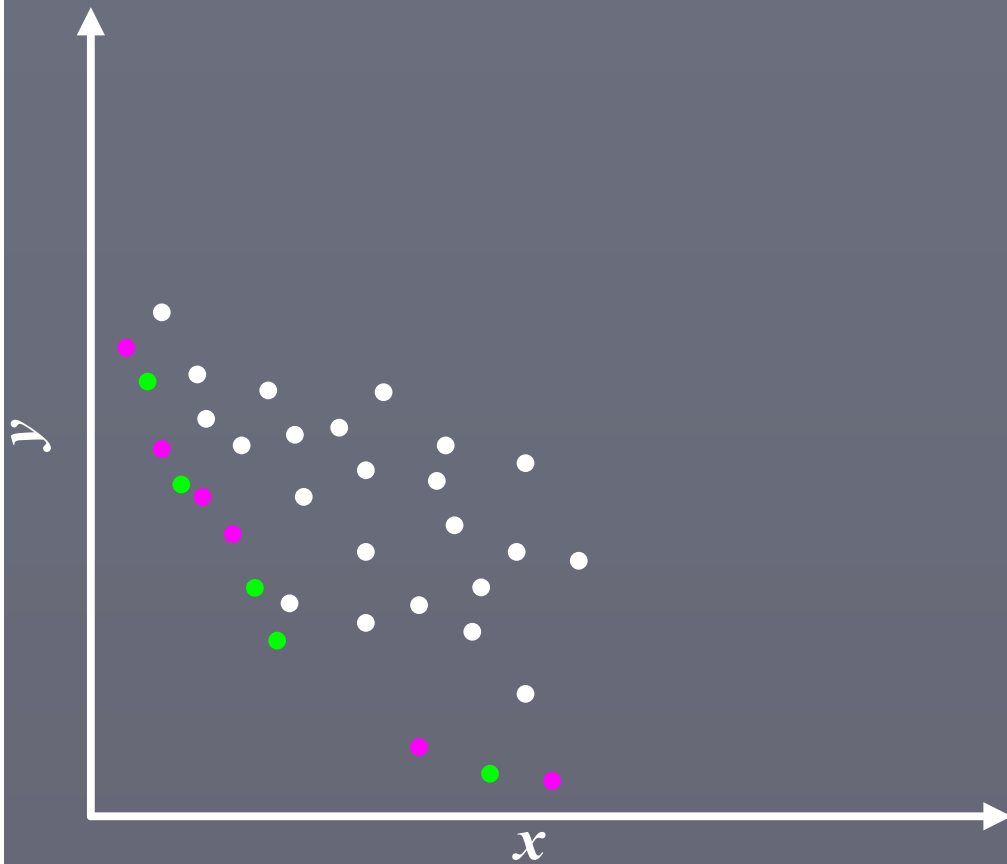
ϵ -Pareto Set

ϵ -Pareto set [Papadimitriou – Yannakakis '00]:


- A set of solutions P_ϵ that *approximately* dominates every other solution.
- For any solution point s , there exists a point in P_ϵ that is within a factor of $(1+\epsilon)$ in all the objectives.

ϵ -Pareto Set

Discrete Space



Convex Space



Approximate Pareto Set

ϵ -Pareto set [Papadimitriou – Yannakakis '00]:


- A set of solutions P_ϵ that *approximately* dominates every other solution.
- For any solution point s , there exists a point in P_ϵ that is within a factor of $(1+\epsilon)$ in all the objectives.

“Always” exists a polynomially succinct one: $\sim (m / \epsilon)^{d-1}$

#bits in objective functions

criteria

Polynomial size ϵ -Pareto sets



- Divide objective space geometrically with ratio $1+\epsilon$ into hyper-rectangles.
- Pick (at most) one point per rectangle.
- Number of points:

$$O \left((m / \epsilon)^{d-1} \right)$$

$m = \#$ bits in obj. values

$d = \#$ objectives


Efficient Constructibility – GAP primitive

Theorem [Papadimitriou – Yannakakis '00]

An ϵ -Pareto set can be computed in polynomial time (for every $\epsilon > 0$) iff the following GAP problem can be so solved (for every $\epsilon > 0$):

Given an instance and bounds b_1, b_2, \dots, b_d either:

- Find a solution point s with $s_i \leq b_i$ for all i , or
- Decide that there is no solution with $s_i \leq b_i / (1 + \epsilon)$ for all i



Discrete Linear Multiobjective Problems

Theorem [PY'00]: We can construct an ϵ -approximate Pareto set for a linear multiobjective combinatorial problem in polynomial time if the (single objective) exact version can be solved in pseudopolynomial time.

Exact Version: Given an instance and an integer B in unary, is there a solution with value exactly equal to B ?

Corollary: Shortest Path, Spanning Tree, Matching.

Other applications in several related contexts: Stochastic Optimization, Smoothed analysis, Mechanism Design, ...

Approximate Pareto Set

ϵ -Pareto set [Papadimitriou – Yannakakis '00]:

- A set of solutions P_ϵ that *approximately* dominates every other solution.
- For any solution point s , there exists a point in P_ϵ that is within a factor of $(1+\epsilon)$ in all the objectives.

Always exists a polynomially succinct one: $\sim (m / \epsilon)^{d-1}$

#bits in objective functions

criteria

The Succinctness Argument: *Obtain a “good” approximation of the Pareto set with as few points as possible.*

“Representative” Approximation

Two general goals:

- **Primal Problem:** Given an instance and $\epsilon > 0$, find an “ ϵ -approximation” to the Pareto set using as few points as possible. (“covering” problem)
- **Dual Problem:** Given an instance and k , find k points that approximate the Pareto set as well as possible, i.e. find the “best” k points. (“clustering” problem)

Succinct approximate Pareto sets

- Approximate Pareto sets are not unique. Want one of (approximately) minimum cardinality.
- **Problem 1 (Primal):** Given an instance and an $\epsilon > 0$, construct an ϵ -Pareto set of (approximately) minimum cardinality $k = \text{OPT}_\epsilon$.
- **Problem 2 (Dual):** Given an instance and a bound k , find k points that form an ϵ -Pareto set for the minimum possible ϵ .

Two Objectives (Primal Problem)

Theorem [Vassilvitskii-Yannakakis' 04] For any bi-objective problem with an efficient GAP routine, we can efficiently compute an ϵ -Pareto set with at most **3** OPT_ϵ many points. Moreover, the factor 3 is tight in general.

Theorem [D-Yannakakis' 07] For bi-objective shortest paths, spanning tree, matching (and other natural problems), we can efficiently compute an ϵ -Pareto set with at most **2** OPT_ϵ many points. Moreover, the factor 2 is tight for these problems, i.e., it is NP-hard to do better.

d Objectives

Theorem [DY' 07]: Let d be a constant. For any $\epsilon' > \epsilon$, we can find an ϵ' -Pareto set with $O(d \log \text{OPT}_\epsilon / \epsilon')$ points. For $d=3$, there is a constant factor approximation.

[VY' 04]: For the dual problem, can do $\log^* k$ approximation (for all d).

Theorem [DY' 07]: For $d = 3$, there is a factor 9 approximation for the dual problem. NP-hard to do better than $3/2$.

Main Lemma [DY' 07]: For any fixed d , an (a, b) - bicriterion approximation algorithm for the primal problem implies a c -approximation for the dual problem, where

$$c = \log^* b + a + 4.$$

Main Open Problem

- **Main Open Question:** Is there a “constant” factor “bi-criterion” approximation for 4 and more objectives?
- **Conjecture: YES**
“For any fixed d , we can efficiently compute an $(1+\epsilon)^d$ -cover whose size is at most $2^d OPT_\epsilon$ ”


Outline

- Introduction
- Approximate Pareto Sets
- *Approximate Convex Pareto Sets*
- Research Directions

Convex Space?


In many settings, ϵ -Pareto set not the right notion of approximation.
Convex objective space (e.g. Multiobjective LP, convex programs, etc.)

Pareto Set




Motivation (I)

In many settings, ϵ -Pareto set not the right notion of approximation.
Convex objective space (e.g. Multiobjective LP, convex programs, etc.)



Motivation (II)

In many settings, ϵ -Pareto set not the right notion of approximation.
Convex objective space (e.g. Multiobjective LP, convex programs, etc.)



Motivation (III)

In many settings, ϵ -Pareto set not the right notion of approximation.

- Convex objective space (e.g. Multiobjective LP, convex programs, etc.)
- Convexity can arise in various other ways.
 - In several applications, points dominated by convex combinations of other points considered inferior.
 - May actually want a representation of the “lower envelope”.
 - The decision is randomized (mixed) and the figures of merit are the expected values of the objective functions.


In these cases, we need an approximation of the **convex Pareto set**.

Convex Pareto Set

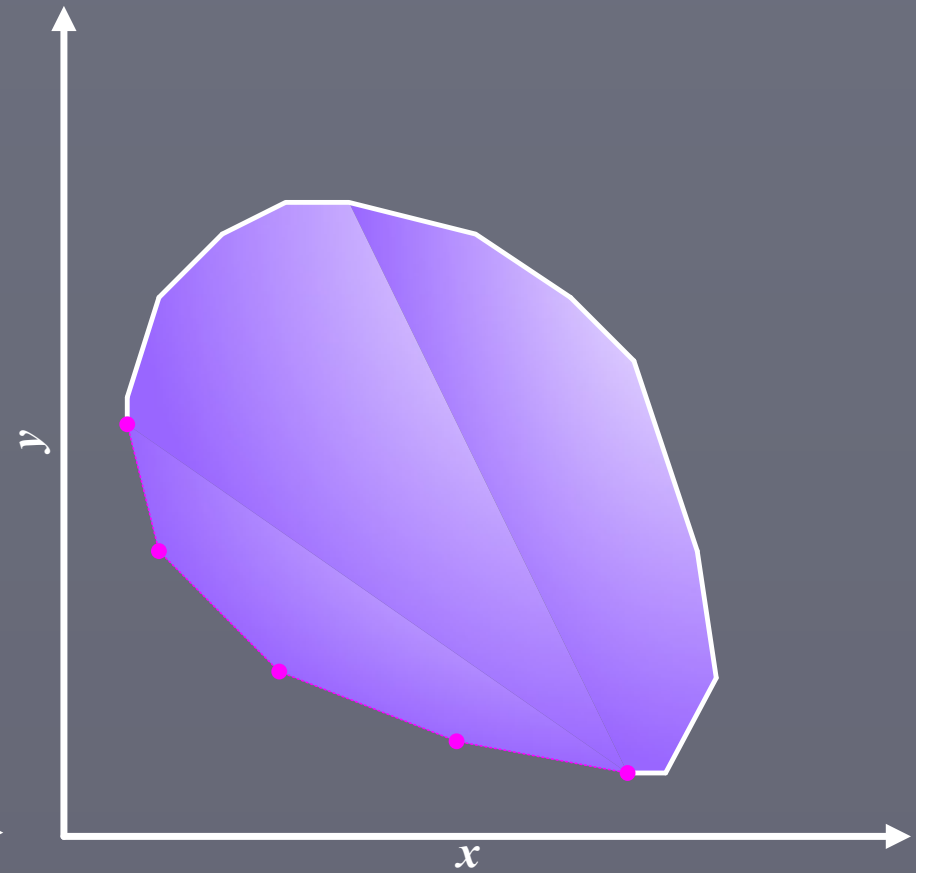
Convex Pareto Set: “Points not dominated by *convex combinations* of other points.”

Convex Pareto Set

Discrete Space



Convex Space



Convex Pareto Set

Convex Pareto Set: “Points not dominated by *convex combinations* of other points.”

Studied in “Parametric” Optimization.

Chandrasekaran ' 77; Megiddo ' 78; Gusfield ' 80; Carstensen ' 83;
Ruhe ' 88; ...


Approximate *Convex* Pareto Set

ϵ -convex Pareto set (ϵ -CP) [D-Yannakakis' 08]:


- A set of solutions CP_ϵ whose *convex combinations* approximately dominate every other solution.
- For any solution point s , there exists a *c.c.* of points in CP_ϵ that is within a factor of $(1+\epsilon)$ in all the objectives.

ϵ -convex Pareto Set

Discrete Space



Convex Space



Approximate *Convex* Pareto Set

ϵ -convex Pareto set (ϵ -CP) [D-Yannakakis' 08]:

- A set of solutions CP_ϵ whose *convex combinations* approximately dominate every other solution.
- For any solution point s , there exists a *c.c.* of points in CP_ϵ that is within a factor of $(1+\epsilon)$ in all the objectives.
- ϵ -Pareto also ϵ -CP ; not vice-versa.
- ϵ -CP can be *arbitrarily* smaller than optimal ϵ -Pareto.

Algorithmic Questions

- **Question 0:** Does a polynomial size ϵ -convex Pareto set always exist?
- **Question 1:** Under what condition is an ^{*} ϵ -convex Pareto set efficiently constructible?

^{}(any polynomial one; not necessarily the smallest one!)*
- **Question 2:** Assuming that the condition is satisfied, can we efficiently compute an ϵ -convex Pareto of approximately minimum size?

Existence (I)

- Question 0: Does a polynomial size ϵ -convex Pareto set always exist?

Existence (II)

- Question 0: Does a polynomial size ϵ -convex Pareto set always exist?

- Easy answer;

- In fact, upper bound is $\sim (m / \sqrt{\epsilon})^{d-1}$

[rescaling + adaptation of Dudley' 74]

Efficient Computability (I)

- Question 1: Under what condition is an^{*} ϵ -convex Pareto set efficiently constructible?

^{*}(*any* polynomial one; not necessarily the smallest one!)

Efficient Computability (II)

- **Question 1:** Under what condition is an ϵ -convex Pareto set efficiently constructible?


Theorem [DY' 08, PY' 00]: An ϵ -convex Pareto set can be computed in polynomial time *iff* the following “Comb problem” has a PTAS:

Given w in \mathbb{R}^d_+ , minimize the combined objective

$$v = w \cdot f = \sum_{i=1}^d w_i \cdot f_i.$$


Comb Oracle (I)

Illustration for $d=2$.




Comb Oracle (II)

Illustration for $d=2$.




Comb Oracle (III)

Illustration for $d=2$.




Comb Oracle (IV)

Illustration for $d=2$.



Comb Oracle (V)

Illustration for $d=2$.



Efficient Computability (III)

Theorem [DY' 08]: An ϵ -convex Pareto set can be computed in polynomial time *iff* the following “Comb problem” has a PTAS:

Given w in \mathbb{R}_{+}^d , minimize the combined objective

$$v = w \cdot f.$$

Proof:

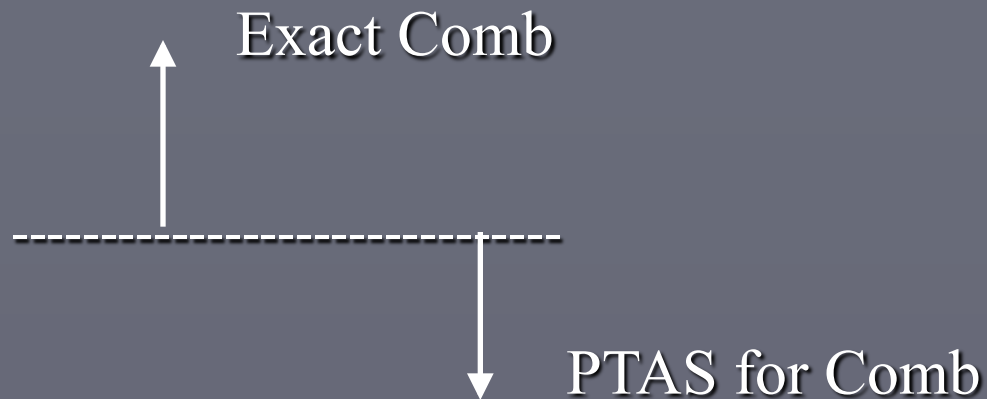
(\rightarrow) Best point of an ϵ -CP under $v = \sum_i w_i \cdot f_i$ is an ϵ -approximate optimum for v .

(\leftarrow) *Oblivious* algorithm ; uses $O_d((m/\epsilon)^{d-1})$ queries to Comb (w) and outputs an ϵ -CP.

Efficient Computability

Corollary: The following multi-objective problems have a PTAS for the construction of an ϵ -convex Pareto set:

- Shortest Path
- Spanning Tree
- Matching
- s - t Min-cut
- Linear Programs
- MDP' s
- ...



- Euclidean TSP
- Convex Programs
- ...

Succinct Approximation

- **Question 2:** Assuming that the condition is satisfied, can we efficiently compute an ϵ -convex Pareto of approximately minimum size?

Problem Statement: Given a problem with d objectives that has a PTAS for Comb, an instance I and error ϵ , compute an ϵ - CP with as few points (solutions) as possible.

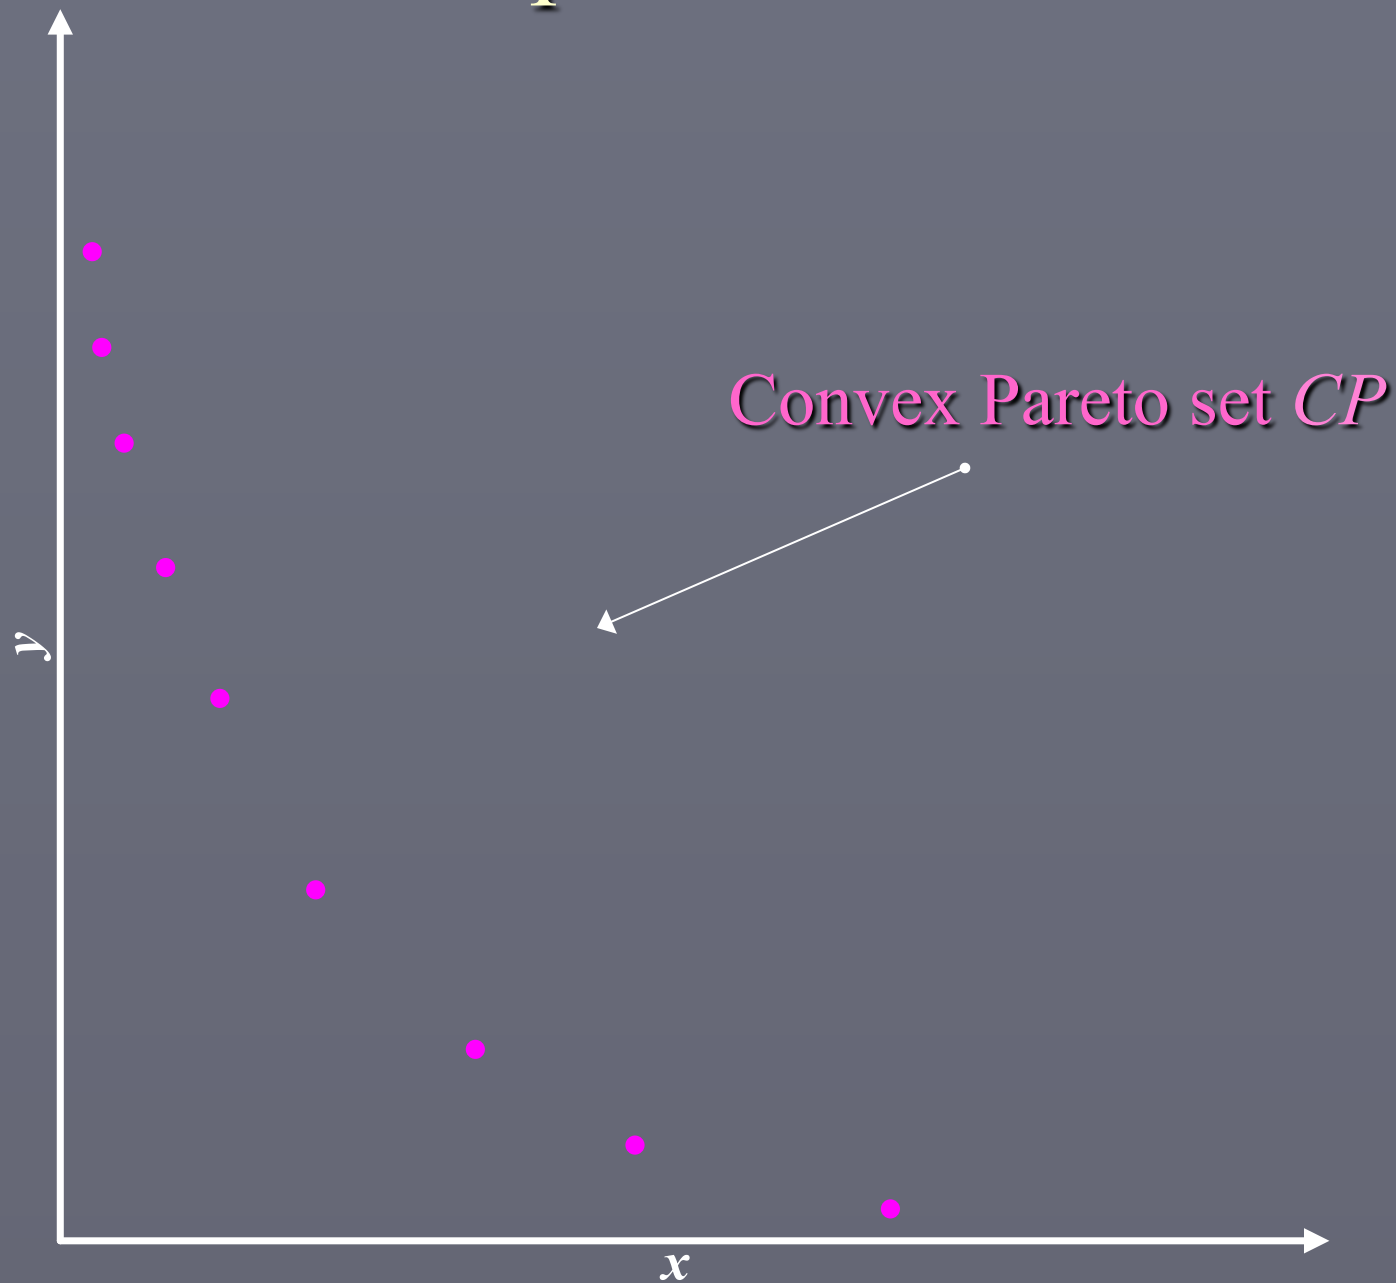
Results

| Primitive | Performance Ratio | | | | |
|------------------|-------------------|---------|----------------------------|--|--|
| | Space | $d = 2$ | $d = 3$ | $d > 3$ (fixed) | Unbounded d |
| Exact Comb | convex | 1^* | PTAS relaxed ϵ | $\log \text{OPT}_\epsilon$ relaxed ϵ | $\Omega(\log n)$ (even for explicit) |
| | discrete | 2 | | | |
| PTAS for Comb | convex | 3 | | | |
| | discrete | 6 | | | |

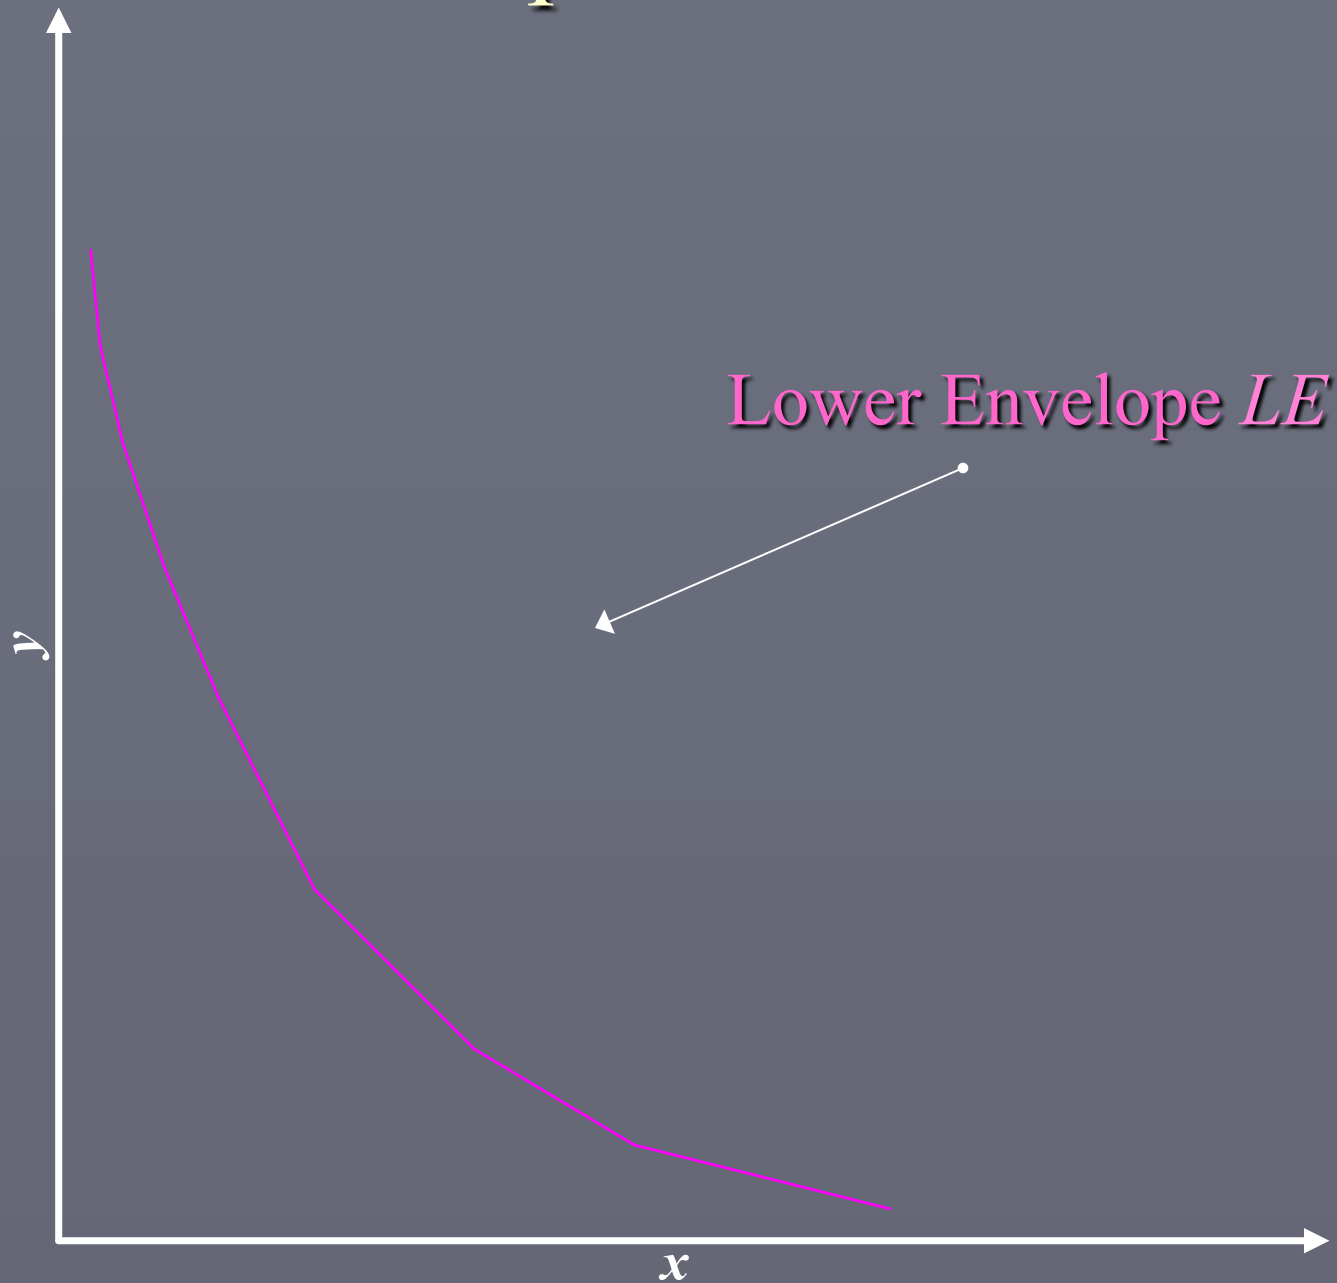
Results

| Primitive | Performance Ratio | | | | |
|------------------|-------------------|---------|----------------------------|--|---|
| | Space | $d = 2$ | $d = 3$ | $d > 3$ (fixed) | Unbounded d |
| Exact Comb | convex | 1^* | PTAS relaxed ϵ | $\log \text{OPT}_\epsilon$ relaxed ϵ | \mathbb{W} ($\log n$) (even for explicit) |
| | discrete | 2 | | | |
| PTAS for Comb | convex | 3 | | | |
| | discrete | 6 | | | |

$2d$ – Explicit – Convex



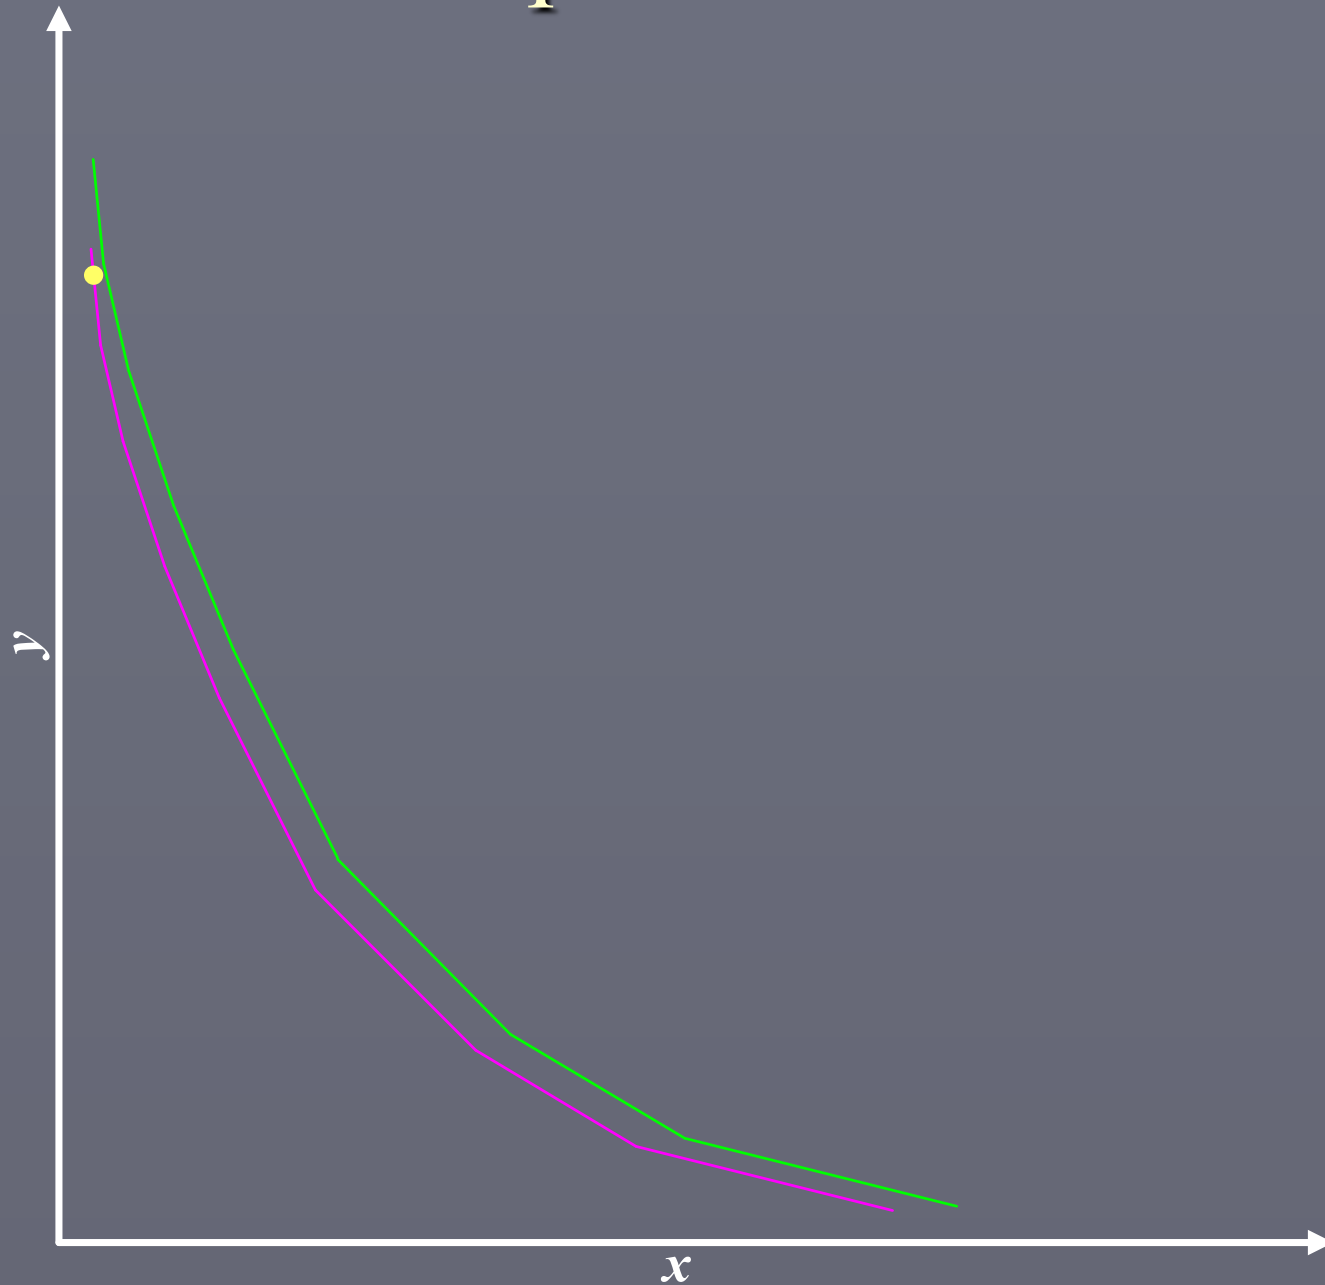
$2d$ – Explicit – Convex



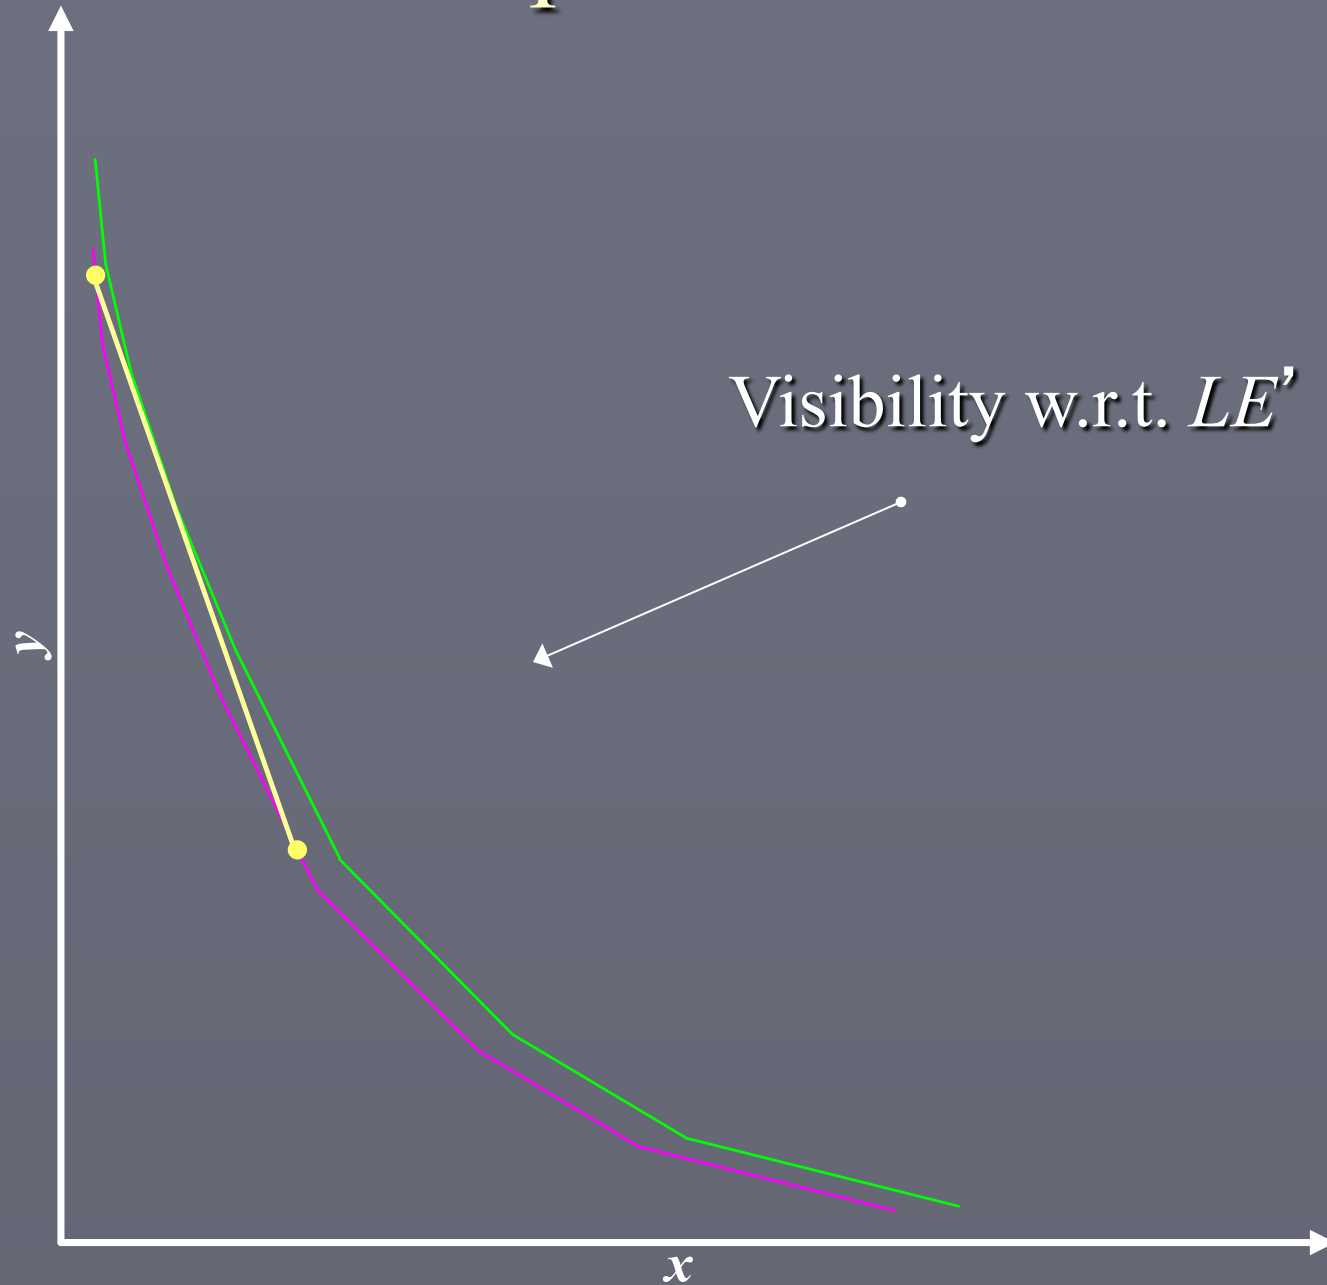
$2d$ – Explicit – Convex




$2d$ – Explicit – Convex



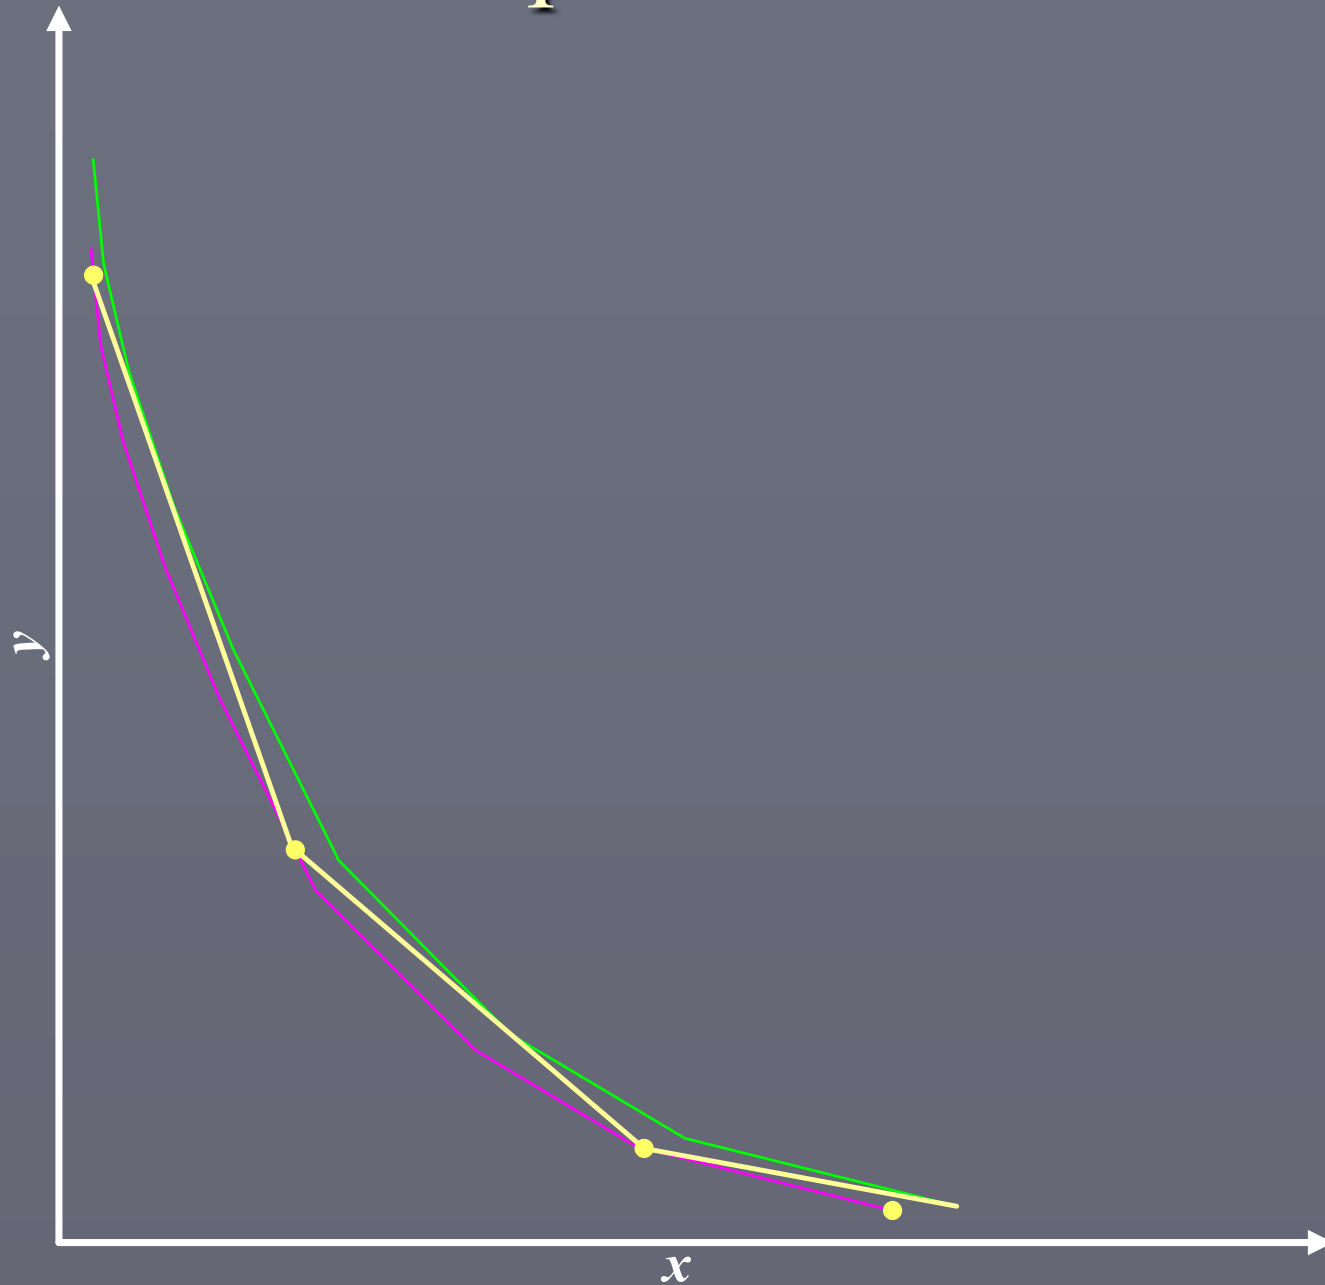
$2d$ – Explicit – Convex



$2d$ – Explicit – Convex

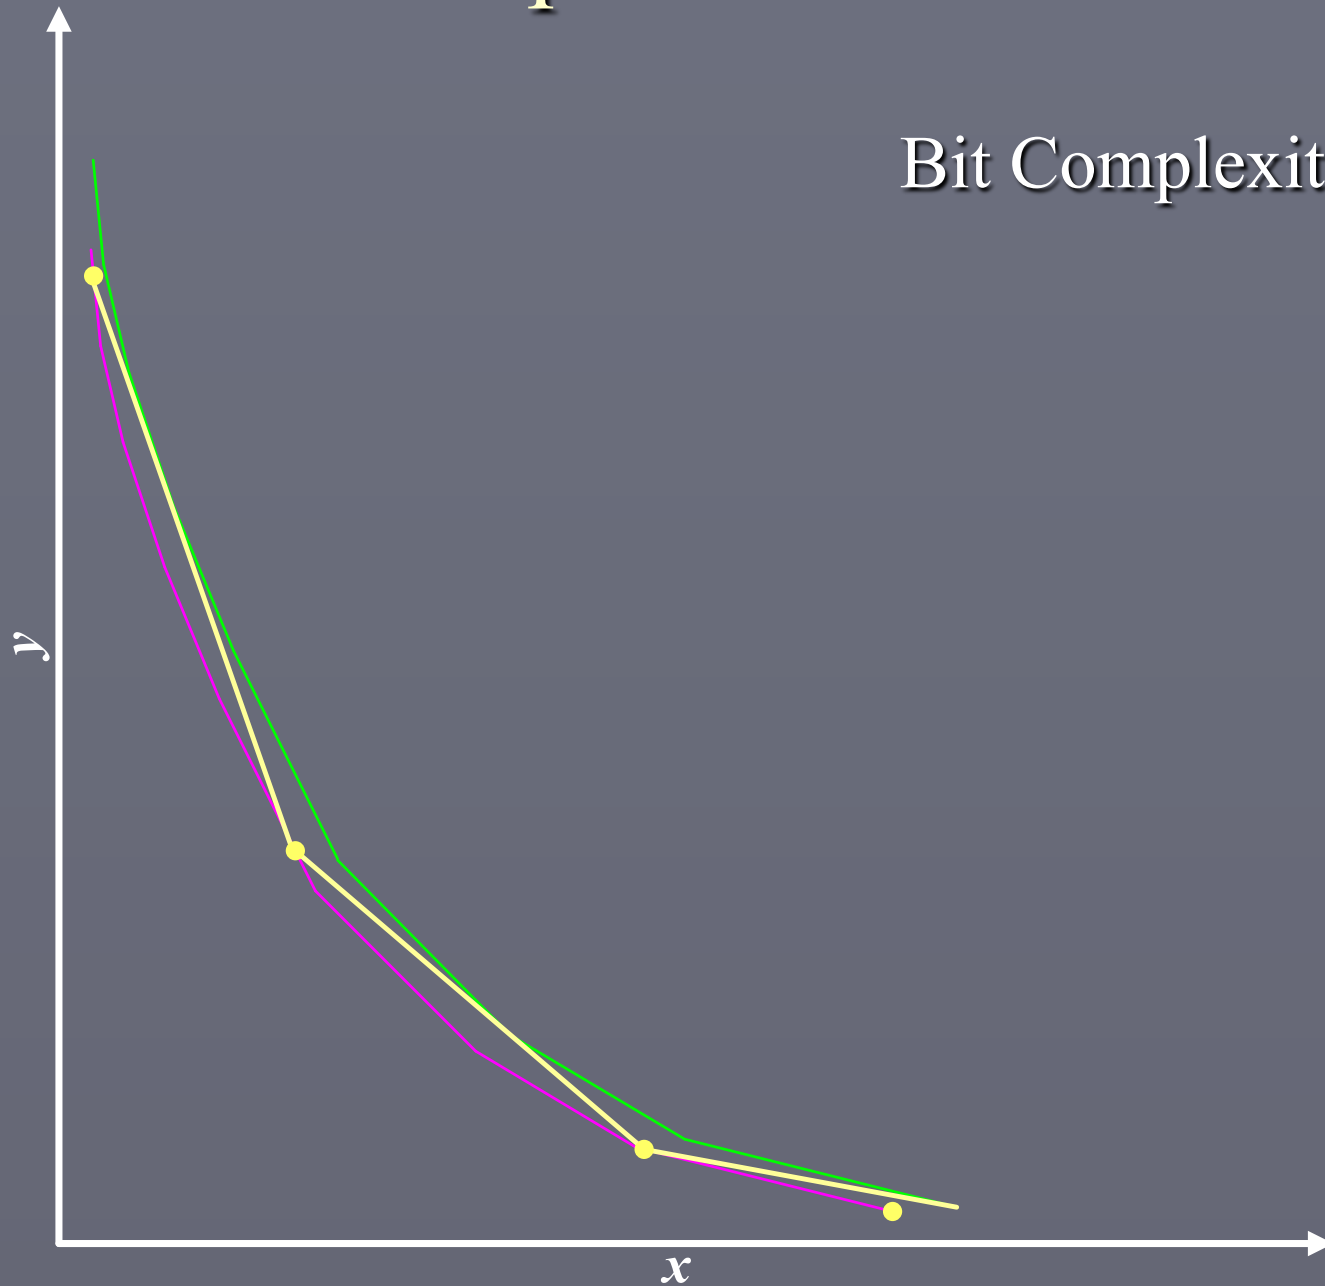


$2d$ – Explicit – Convex

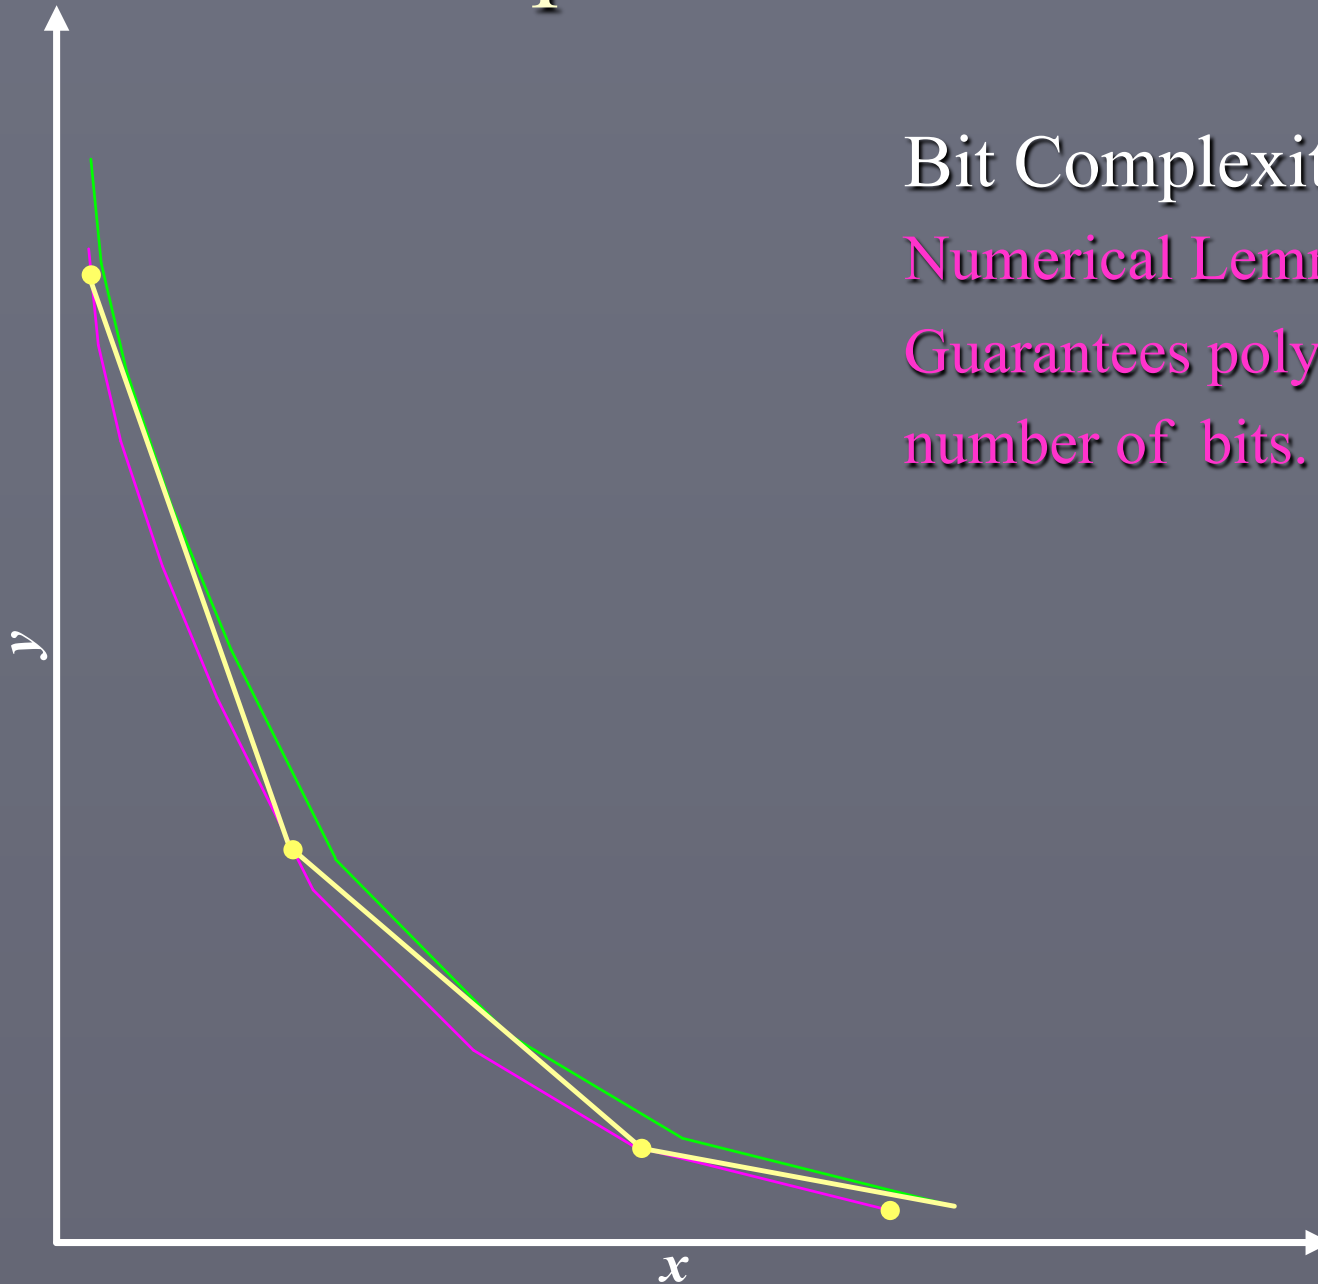


$2d$ – Explicit – Convex

Bit Complexity?



$2d$ – Explicit – Convex



Bit Complexity?

Numerical Lemma--

Guarantees polynomial
number of bits.

Bi-objective LP

Theorem [DY' 08]: Can compute an ϵ - CP of optimal size OPT_ϵ by solving $2 \cdot \text{OPT}_\epsilon$ Linear Programs.

Main idea:

- Simulate Explicit Points Algorithm.
- Exploit LP-duality.

General Case: Approximate Comb

- First compute a δ -CP for appropriate $\delta < \epsilon$.
(using generic oblivious algorithm)
- Then post-process using explicit algorithm.

Lemma [D-Y' 08]: Let A in \mathbb{R}^2_+ be a set of points and $\epsilon > 0$.

For all $\delta > 0$ satisfying $\delta < (1 + \epsilon)^{1/2} - 1 \approx \epsilon/2$ we have:

$$|CP^*(A, \delta)| \leq 3 \cdot |CP^*(A, \epsilon)|$$

- Factor 3 for convex objective space, factor 6 for discrete.


Minimize number of Queries

In the general case, number of Comb calls is $\Omega(m/\epsilon)$...

Problem Statement:


Find an ϵ -CP using *as few queries* to the Comb oracle as possible.
(online algorithm / competitive ratio).

What would Archimedes do ?



Archimedes Thoughtful (Fetti, 1620)

What would Archimedes do ?



Archimedes Thoughtful (Fetti, 1620)

[Daskalakis-D-Yannakakis' 10]

Analysis of a natural greedy heuristic for this problem.

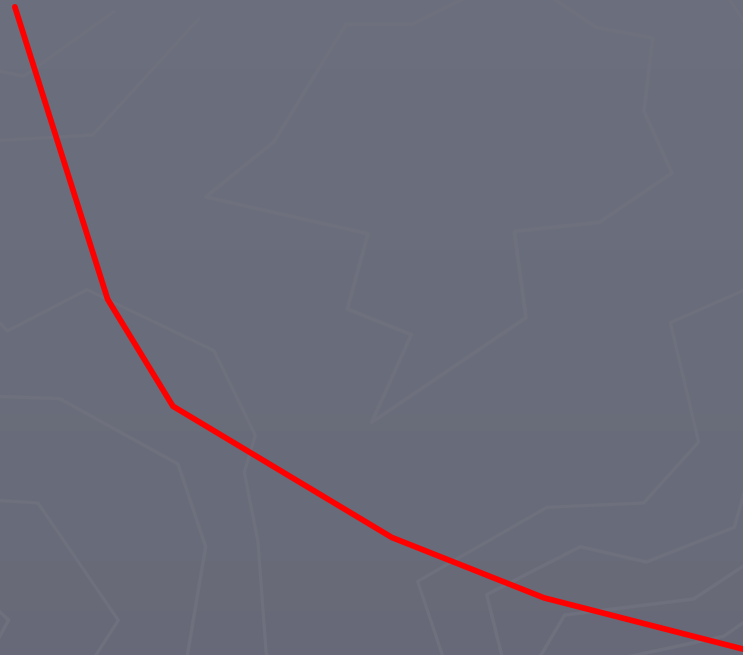
“Chord” Algorithm

Also very popular in other settings.

- Convex function approximation
[BHR' 91, Ro' 92, YG' 97, ...]
- Curve simplification
[Ramer' 72, Douglas-Peucker' 73]
- Parametric Optimization [ES' 79]

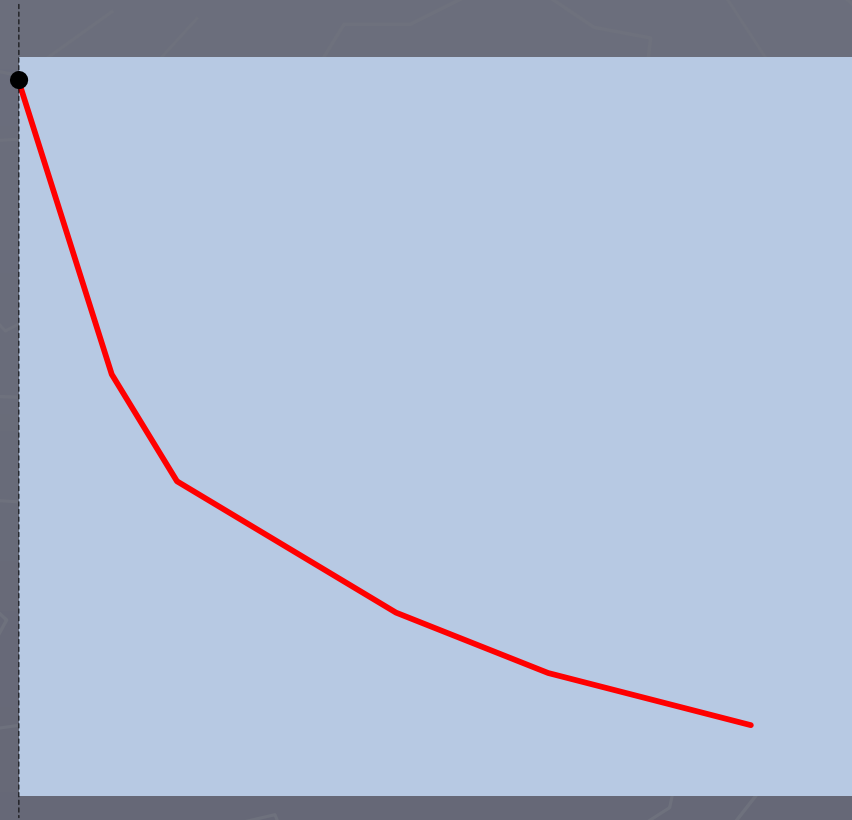
Illustration of Chord Algorithm

Convex
Curve C :




Chord Algorithm (I)

Find *leftmost*
point of C :



Chord Algorithm (II)

Find *rightmost*
point of C :




Chord Algorithm (III)

Initial

Information:

C in shaded triangle.




Chord Algorithm (IV)

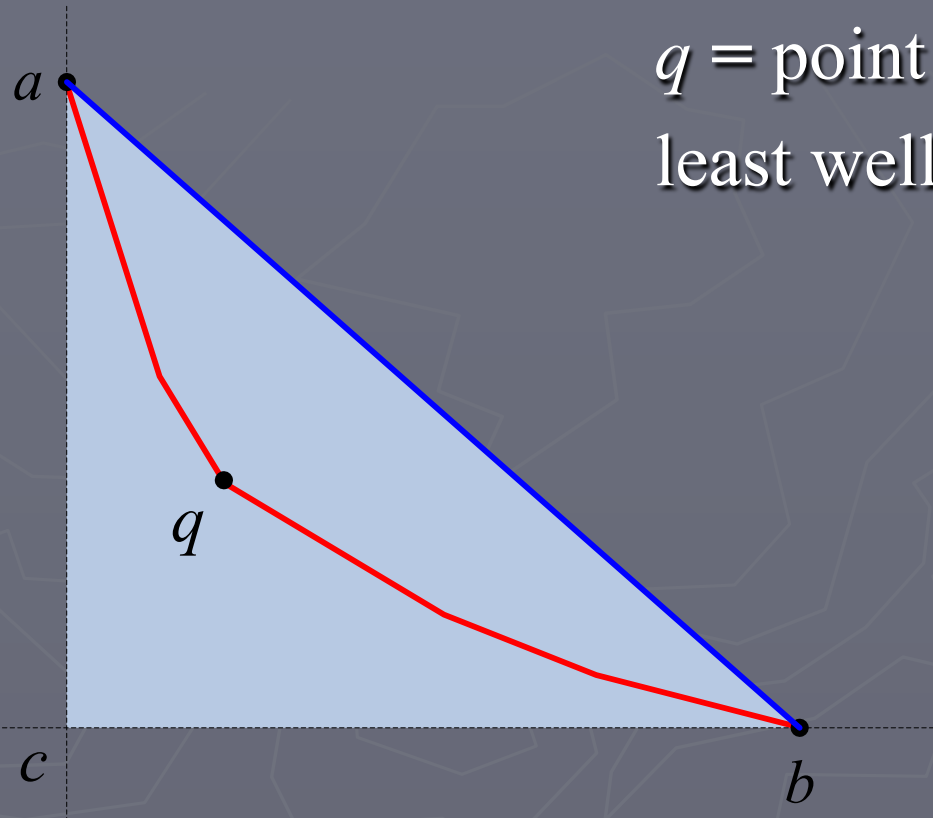
Initial
approx:
segment ab .

$$d(c, ab) \leq \epsilon$$

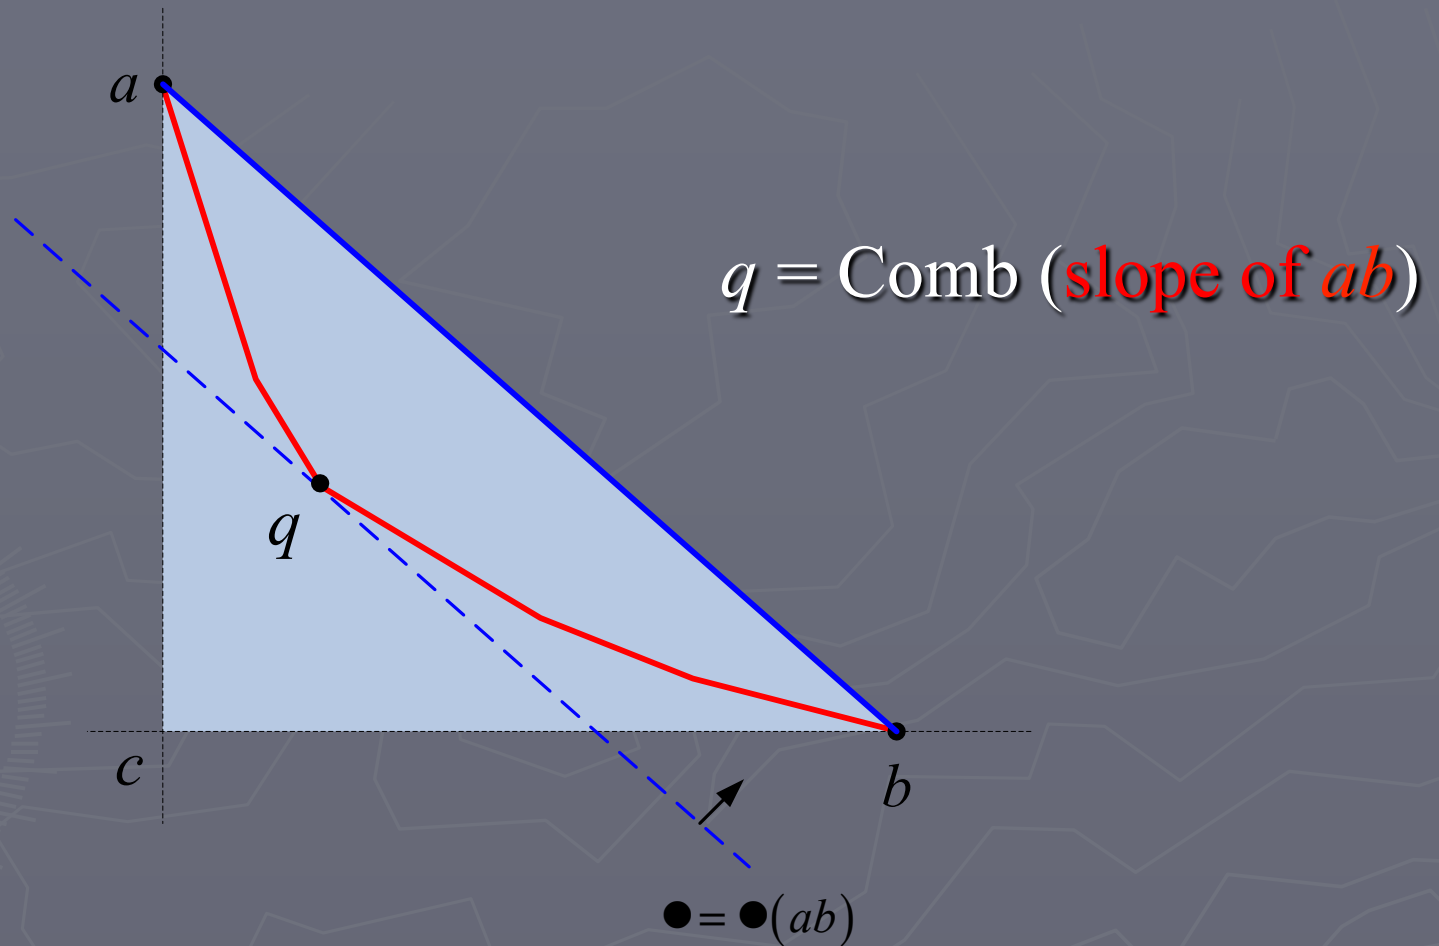
?



Chord Algorithm (V)

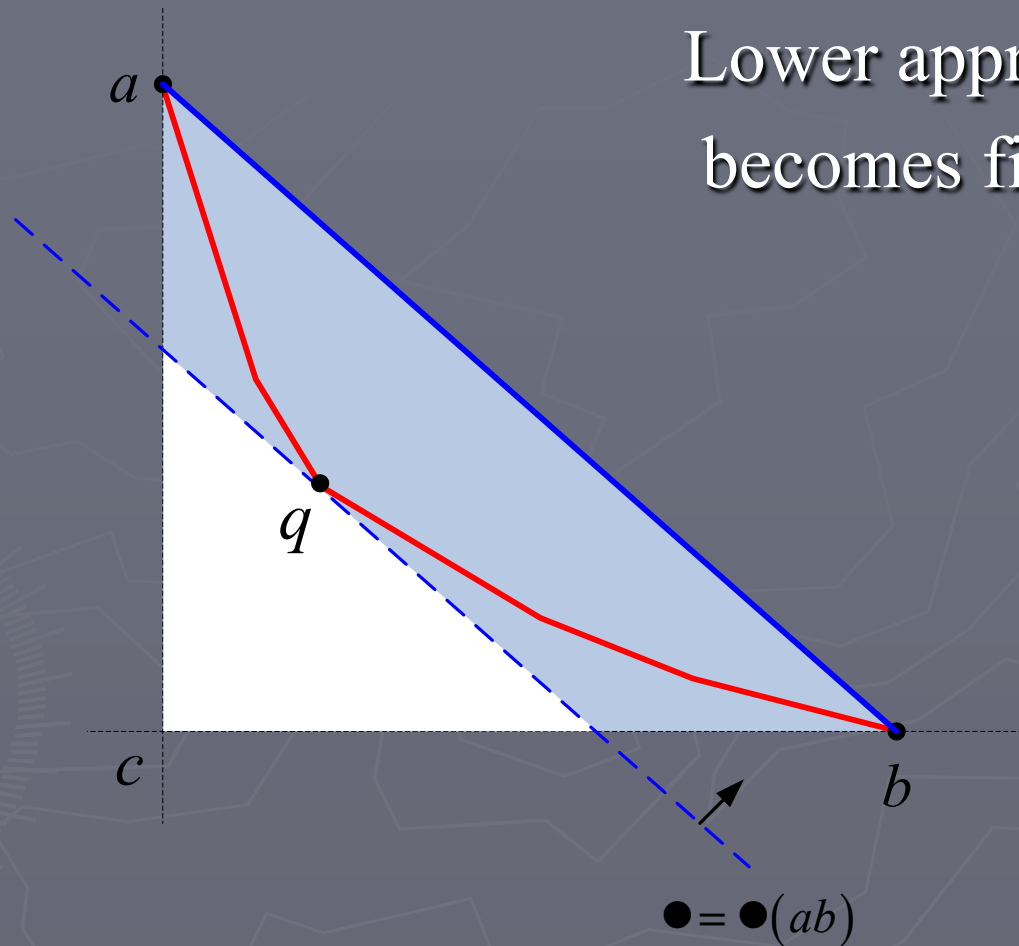


Chord Algorithm (VI)




Chord Algorithm (VII)

Lower approximation
becomes finer.



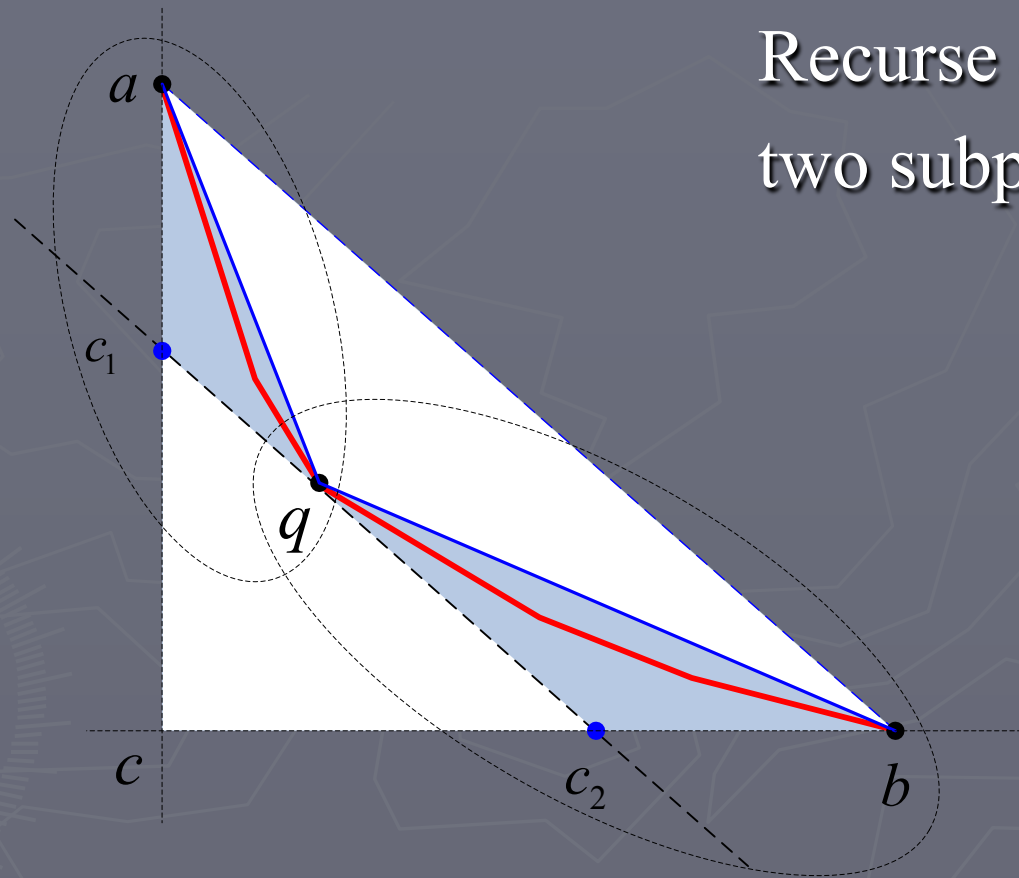
Chord Algorithm (VIII)

Same for upper approximation.




Chord Algorithm (IX)

Recurse on
two subproblems.



What would Archimedes do ?



Archimedes Thoughtful (Fetti, 1620)

[Rote' 92]

“When approximating a parabola, the sequence of upper approximations is just the sequence which Archimedes used to exhaust the area of a parabolic segment in his second proof of the area formula.”

Archimedes constructed his sequence of polygons according to the Chord algorithm.

Analysis of Chord Algorithm

Theorem 1 [DDY' 10]: The *worst-case* performance of the Chord algorithm is $\Theta(m/\log m + \log(1/\epsilon)/\log \log(1/\epsilon))$.

n points from “un-concentrated”
product distribution

Theorem 2 [DDY' 10]: The *average case* performance of the Chord algorithm is $\Theta(\log m + \log \log 1/\epsilon)$.

Optimal Algorithm ?

Lower Bound [DDY' 10]: No algorithm with access to a Comb oracle can have *worst-case* performance better than $\Omega(\log m + \log \log 1/\epsilon)$.

Is there is an algorithm with *worst-case* performance $O(\log m + \log \log 1/\epsilon)$? **Yes** [D-Yannakakis'12]

How about for more than 2 objectives?

Open Problems

- $d > 3$ objectives?
- Faster algorithms for important combinatorial problems.
- Online learning of multi-objective problems.
- Merging approximate Pareto sets.
- Connections to other areas ?

Any Questions?

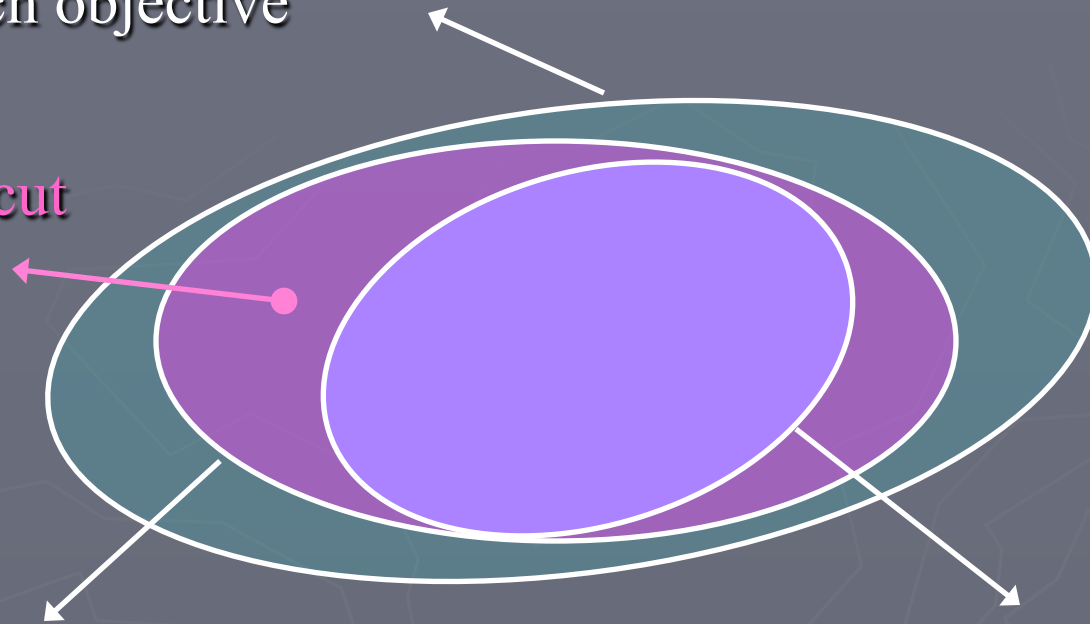


Thank you!

Efficient Computability — Comparison

PTAS for each objective

e.g. $s - t$ min-cut



PTAS for ϵ -convex Pareto
e.g. convex programs,
all linear problems with a
PTAS for single objective

PTAS for ϵ -Pareto
e.g. shortest path, spanning
tree, matching, etc.


Bi-objective LP

Theorem: Can compute an ϵ - CP of (optimal) size OPT_ϵ , by solving 2OPT_ϵ LPs.

Recall notation:

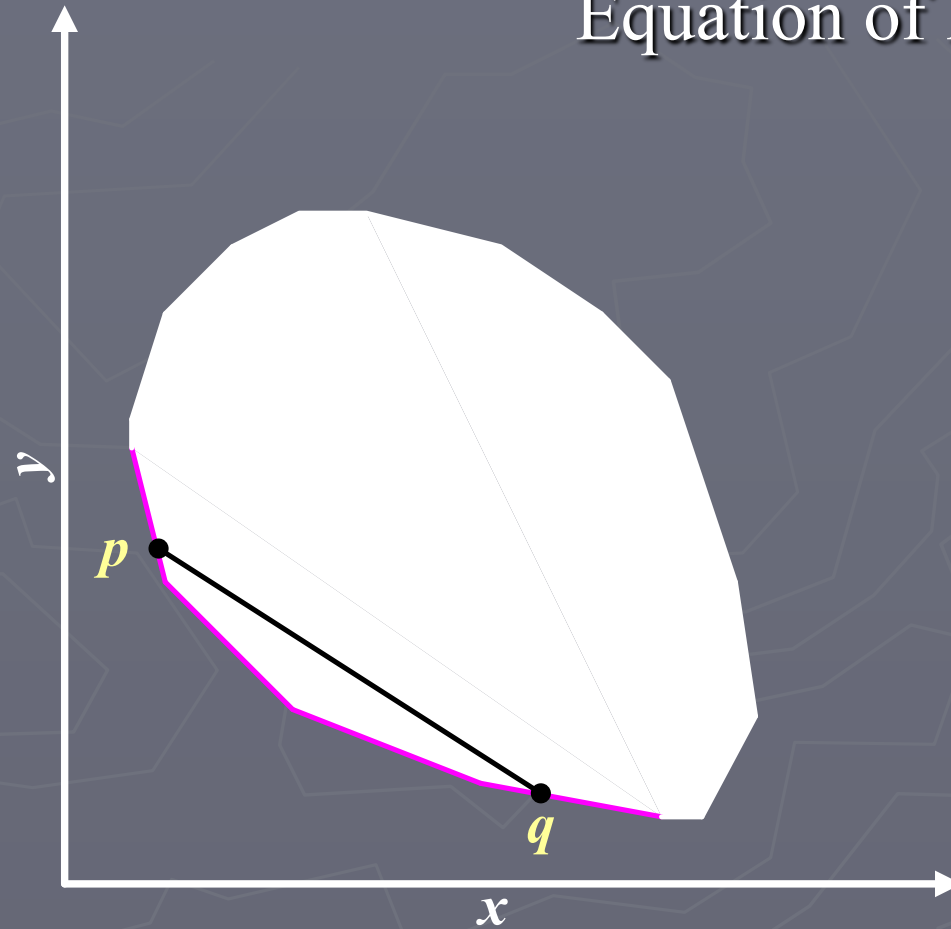
- Decision space: $\mathcal{Z} = \{ z \in \mathbb{R}^{n_{\mathbb{Q}}^1} \mid A \cdot z \geq b, z \geq \mathbf{0}_{n_{\mathbb{Q}}^1} \}$
- Minimization Objectives: $c, d \in \mathbb{Q}^{1_{\mathbb{Q}}^n}$
- Objective Space: $\mathcal{X} = \{ (x, y) \in \mathbb{R}_+^2 \mid x = c \cdot z, y = d \cdot z, z \in \mathcal{Z} \}$

Bi-objective LP



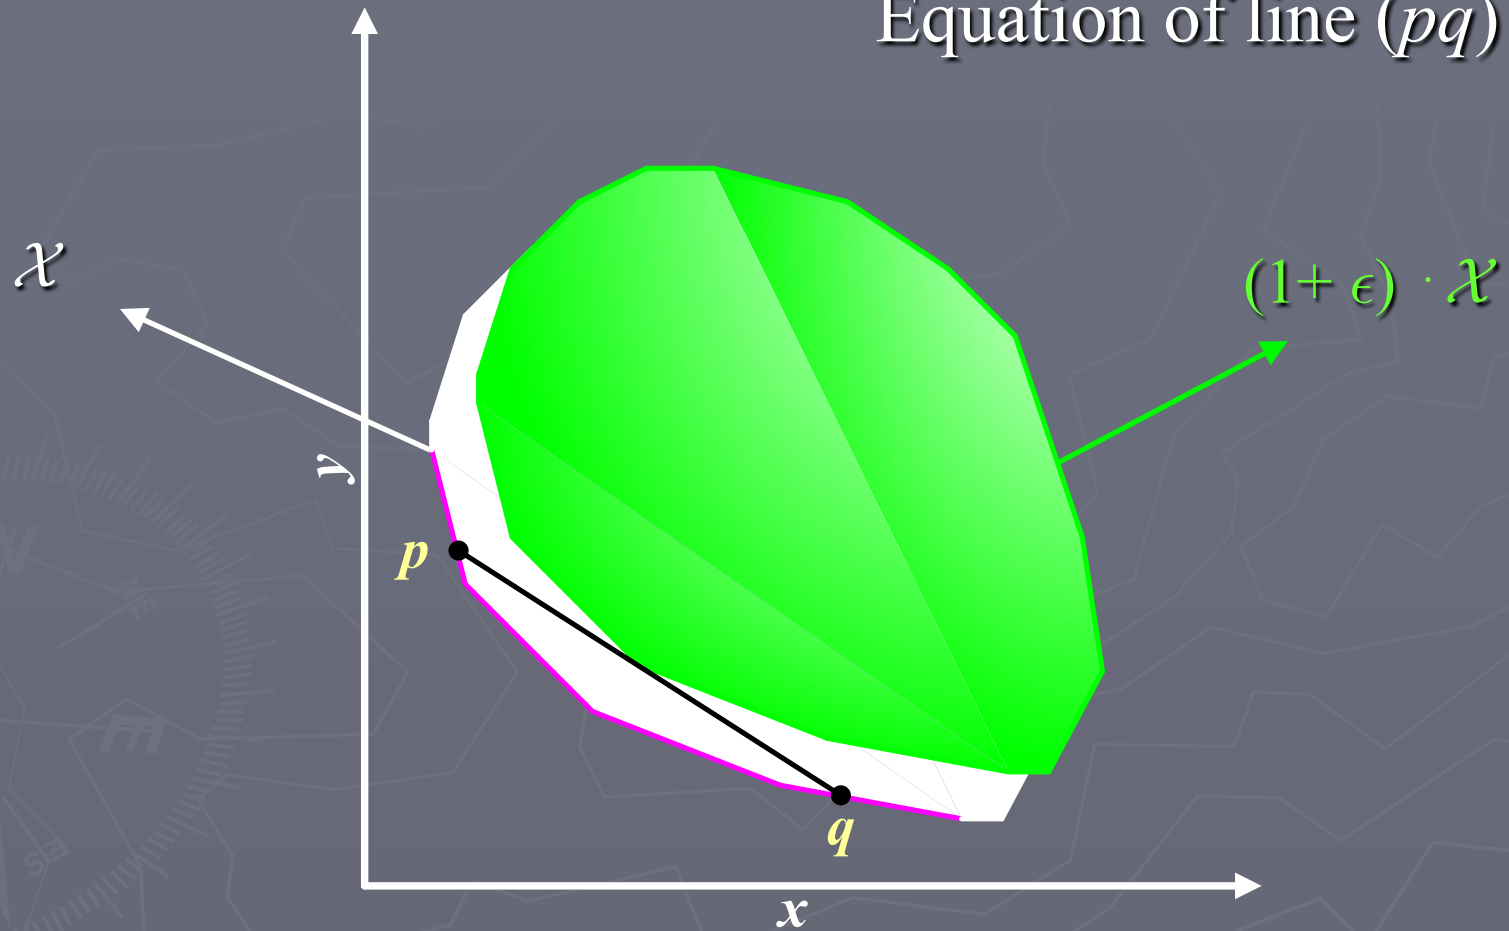
Bi-objective LP

Equation of line (pq) ?



Bi-objective LP

Equation of line (pq) ?




Bi-objective LP

Suppose (pq) : $r^* x + y = t^*$

Claim: $y(q) = \min y$

s.t.

$$r^* x + y = t^*$$




$$\begin{aligned} x &= c \cdot z \\ y &= d \cdot z \\ z &\in \mathcal{Z} \end{aligned}$$

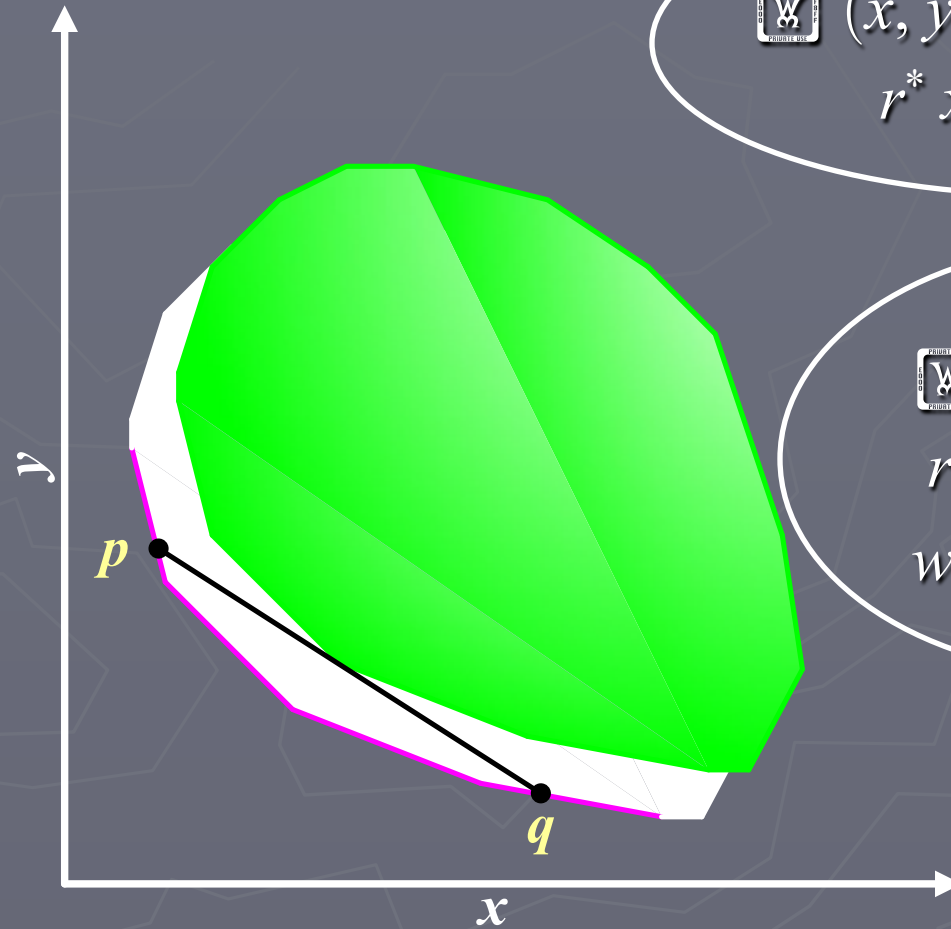
$$(x, y) \in \mathcal{X}$$

Bi-objective LP

$$\begin{aligned} \min (x, y) \in \mathcal{X}: \\ r^* x + y \leq t^* / (1 + \epsilon) \end{aligned}$$




Bi-objective LP



$$\begin{aligned} \min (x, y) \in \mathcal{X}: \\ r^* x + y \leq t^* / (1 + \epsilon) \end{aligned}$$

$$\begin{aligned} \min w \in \mathbb{R}^1 \times \mathbb{R}^m_+ : \\ r^* c + d \leq w \cdot A \\ w \cdot b \leq t^* / (1 + \epsilon) \end{aligned}$$

Bi-objective LP



$$r^* = \min r \text{ s.t.}$$

$$r x(p) + y(p) = t$$

$$r c + d \leq w \cdot A$$

$$w \cdot b \leq t / (1 + \epsilon)$$

$$w \geq 0$$

$$r \geq 0$$