Feature-based Dynamic Matching *in the Dark*



Akshit Kumar
University of Toronto

Digital Matching Markets

handy

Urban Company

kandua

SweepSouth
CLEAN HAPPY HOME

Digital Matching Markets

Home service platforms which provide on-demand services like cleaning, fixing, etc.

hANDY

UC Urban Company

kandua

SweepSouth
CLEAN HAPPY HOME

Digital Matching Markets

Home service platforms which provide on-demand services like cleaning, fixing, etc.

Salient features of these platforms:

- Have access to a pool of heterogeneous service providers which are differentiated by their features (eg. location, ratings, HOOs)

Digital Matching Markets

Home service platforms which provide on-demand services like cleaning, fixing, etc.

Salient features of these platforms:

- Have access to a pool of heterogeneous service providers which are differentiated by their features (eg. location, ratings, HOOs)
- Customers arrive online and specify their preferences (eg. location, price, time)

Digital Matching Markets

Home service platforms which provide on-demand services like cleaning, fixing, etc.

Salient features of these platforms:

- Have access to a pool of heterogeneous service providers which are differentiated by their features (eg. location, ratings, HOOs)
- Customers arrive online and specify their preferences (eg. location, price, time)
- Customers need to be matched in near real time to a service provider

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated?

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated?

The key trade-off is
making good matches now v/s preserving supply for later

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated?

The key trade-off is
making good matches now v/s preserving supply for later

Need to be forward looking !!!

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated *with a known demand distribution*?

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated *with a known demand distribution*?

A

Feature-Based Dynamic Matching

Yilun Chen, Yash Kanoria, Akshit Kumar, Wenxin Zhang

Published Online: 18 Nov 2025 | https://doi.org/10.1287/opre.2024.0730

Abstract

Motivated by matching platforms that match agents in a centralized manner, we study dynamic two-sided matching in a setting where both customers (demand) and service providers (supply) are heterogeneous and the pool of service providers is limited. We model heterogeneity on the two sides of the market by demand weight vectors drawn independently and identically distributed (i.i.d.) from some distribution and supply feature vectors drawn i.i.d. from a (possibly) different distribution. The matching of a demand-supply pair generates a matching value that depends on their weight and feature vectors. We adopt a notion of regret, specifically the additive loss relative to the value (per match) achievable in the limiting hindsight optimum as our performance metric for matching policies. Simple myopic policies suffer nonvanishing $\Omega(1)$ regret in large markets. We propose a forward-looking supply-aware policy dubbed simulate-optimize-assign-repeat (SOAR) that balances between producing high match value for the current match and preserving valuable supply for future customers. We prove that SOAR achieves the optimal regret scaling under different assumptions on the demand and supply distributions. En route to proving our guarantees, we develop a novel framework for analyzing the performance of our SOAR policy that may be of broader interest. As a corollary of our techniques, we also resolve an open problem posed previously.

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated *without a known demand distribution*?

Key Operational Question

Q How should centralized matching platforms match customers dynamically to maximize overall quality of matches generated *without a known demand distribution*?

A This talk

In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*



In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification



In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification

- Design a simple and near-optimal policy **ROAR**



In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification

- Design a simple and near-optimal policy **ROAR**

  - Textbook like proof for ROAR



In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification

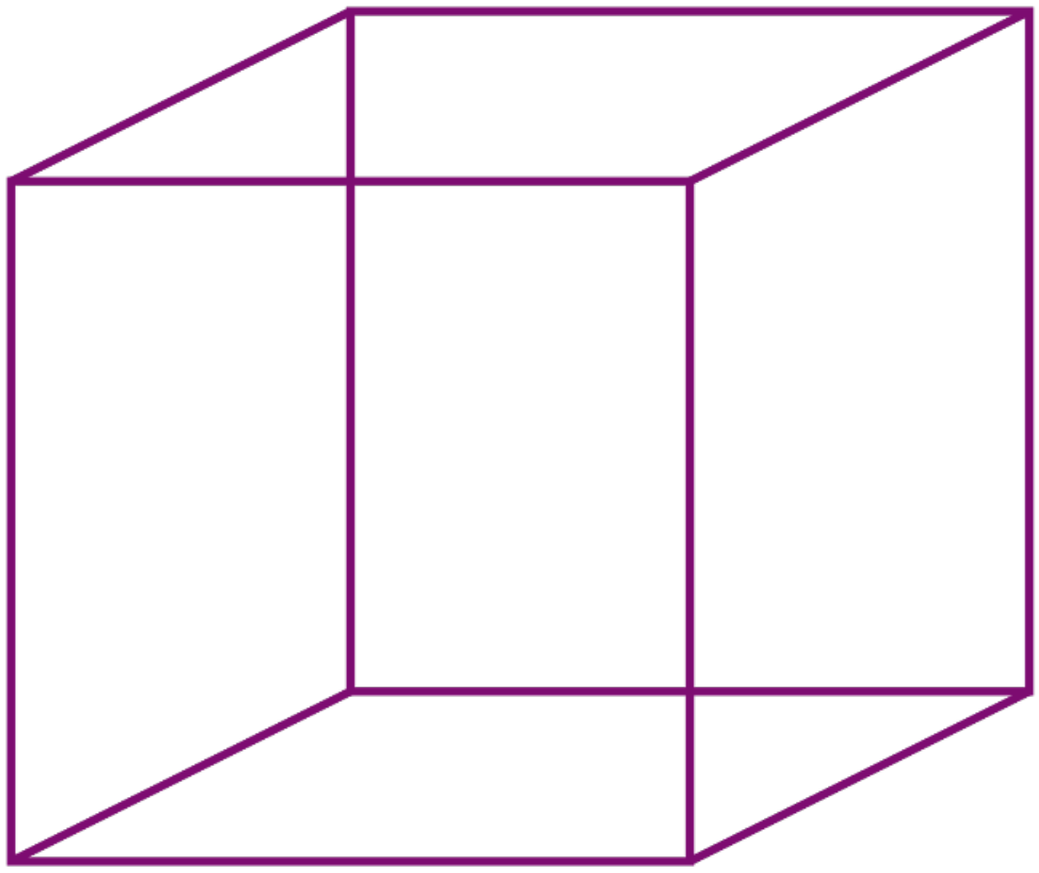
- Design a simple and near-optimal policy **ROAR**

  - Textbook like proof for ROAR

  - No separate learning terms in the regret analysis

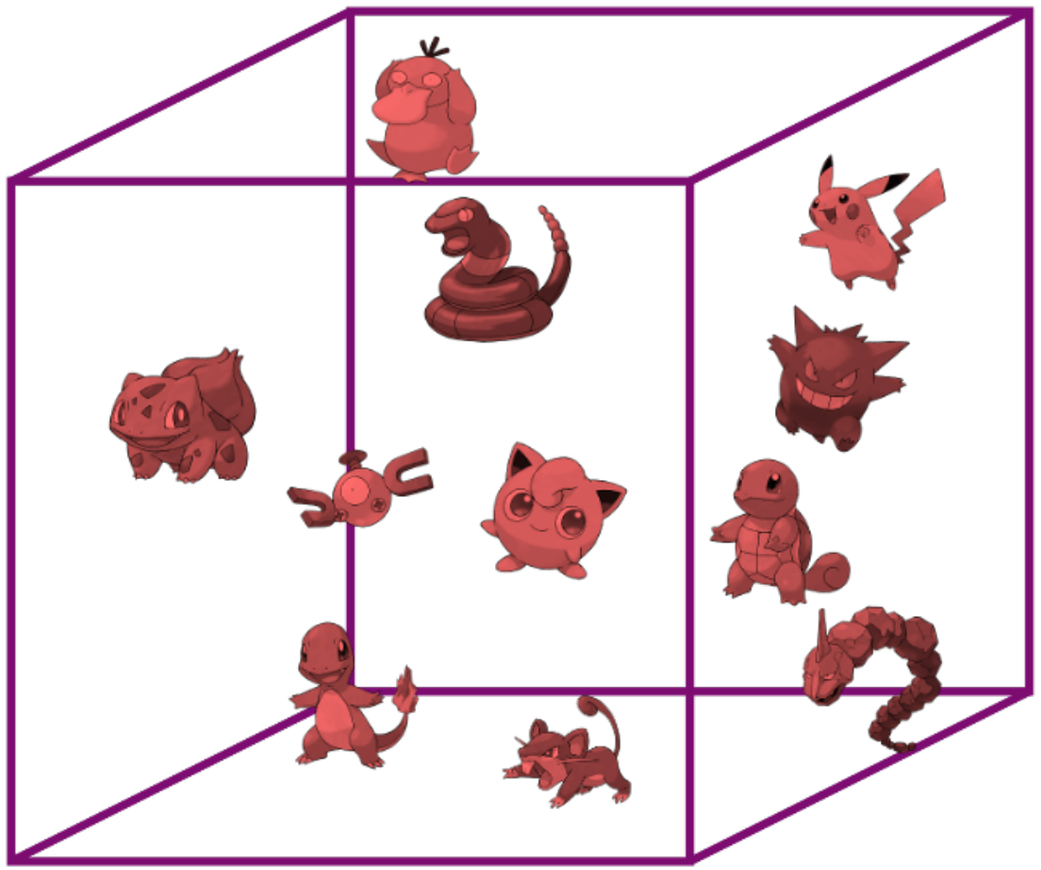


Model



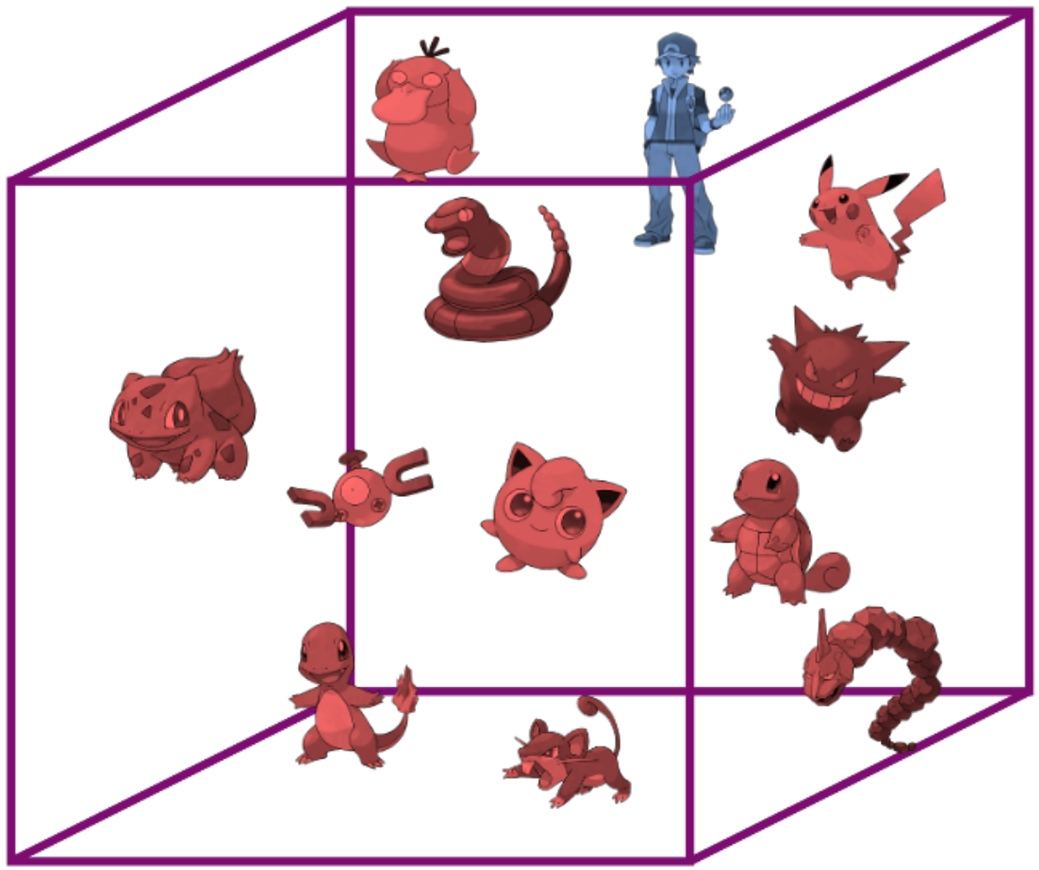
Supply is represented by a feature vector in a d-dimensional space

Model



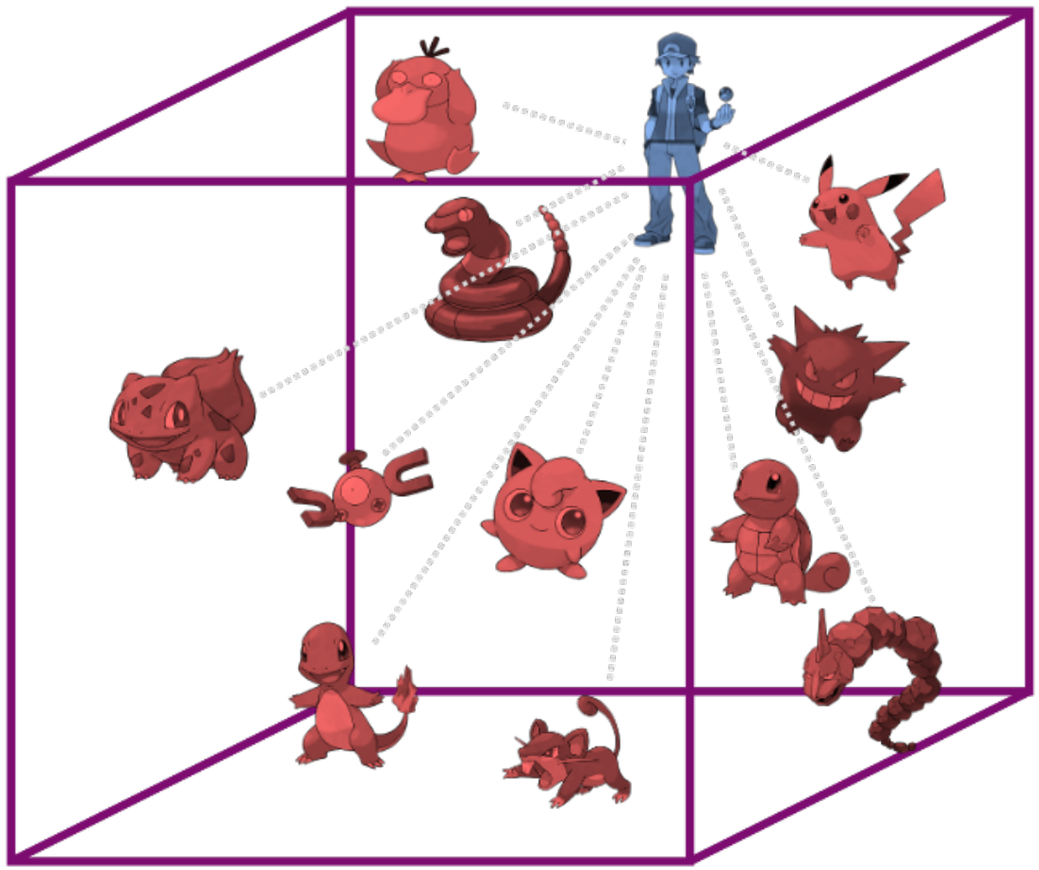
Platform has n supply units drawn i.i.d from distribution Q

Model



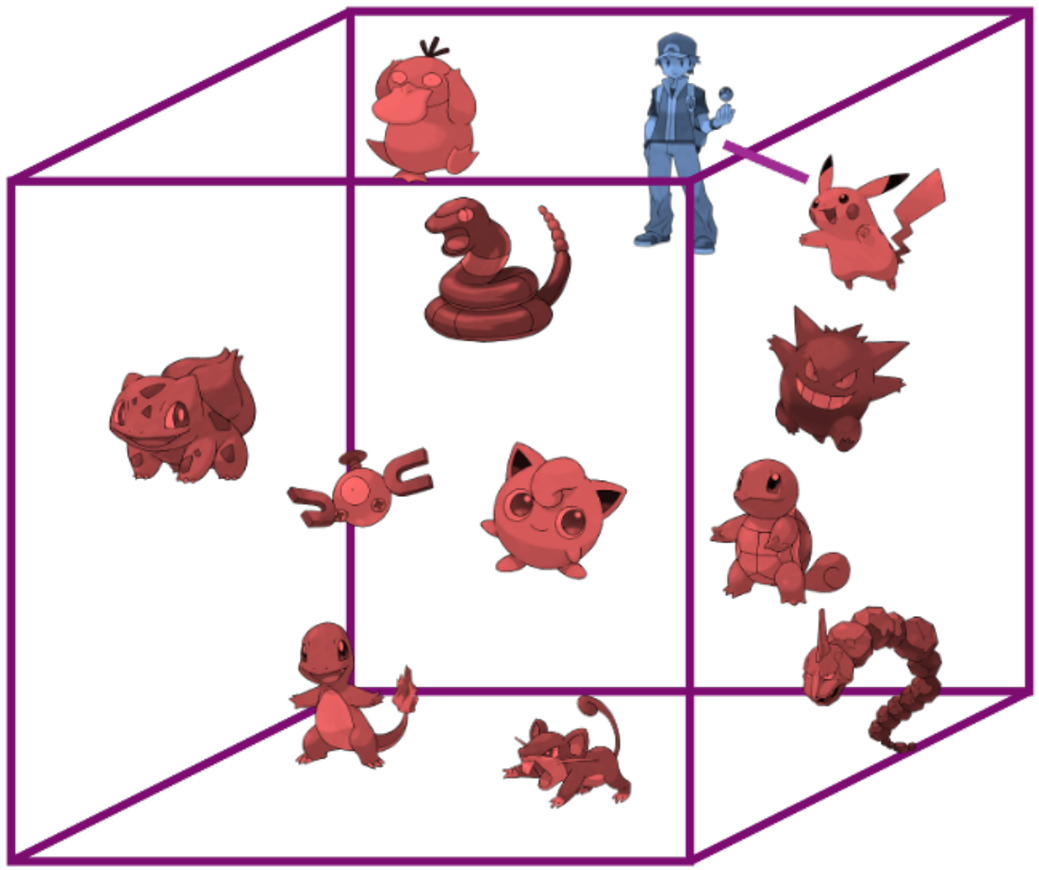
At time t, a demand unit (weight vector) is drawn i.i.d from **unknown** P

Model



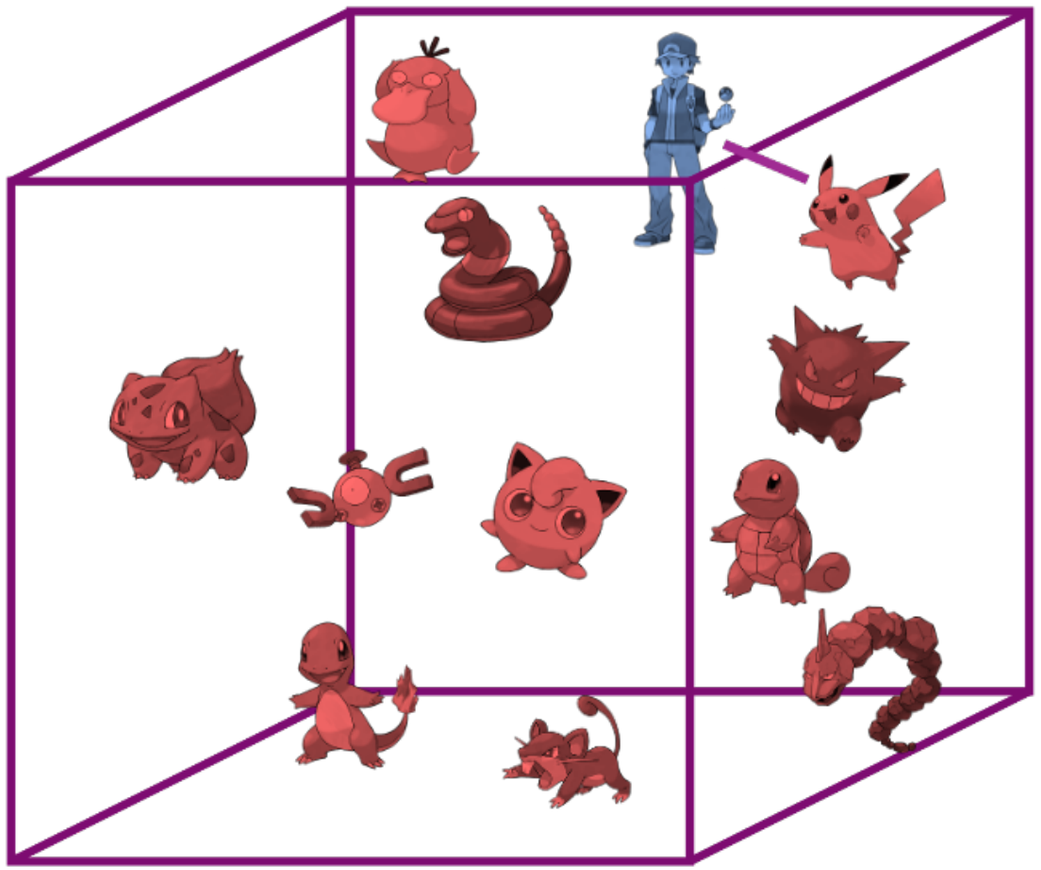
Platform must irrevocably match a demand unit to a supply unit

Model



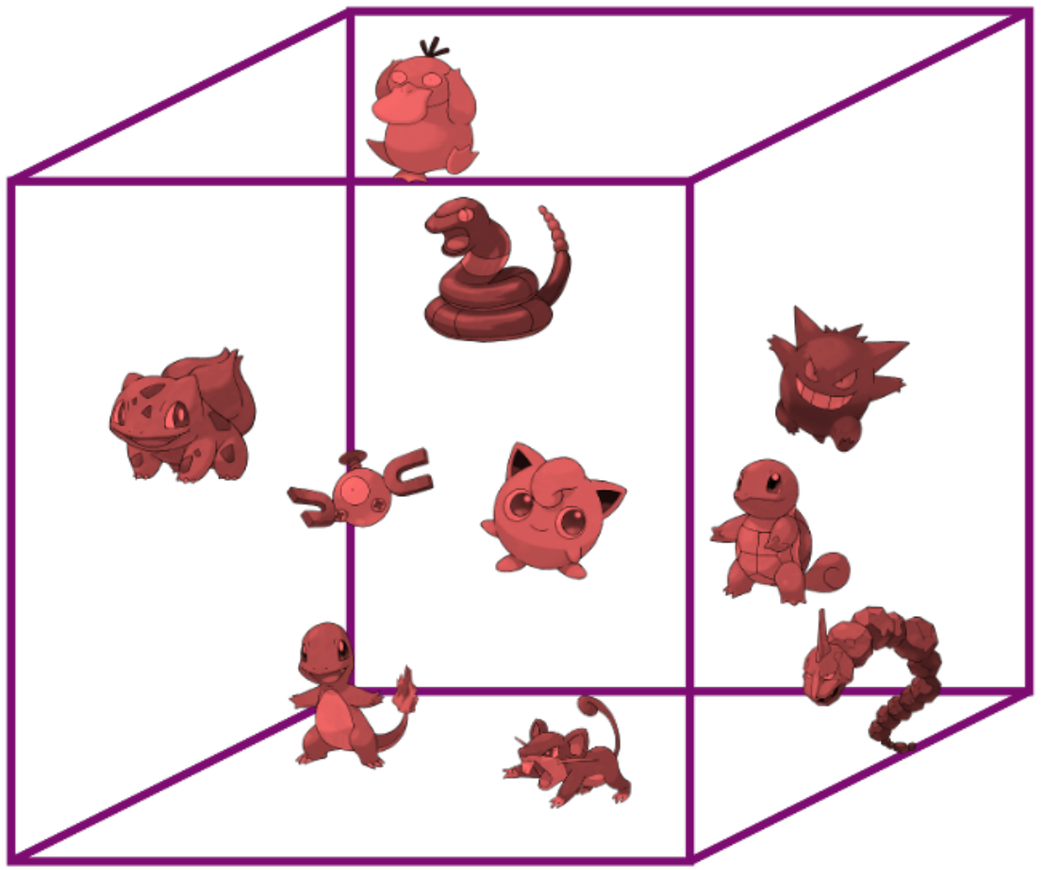
Platform must irrevocably match a demand unit to a supply unit

Model



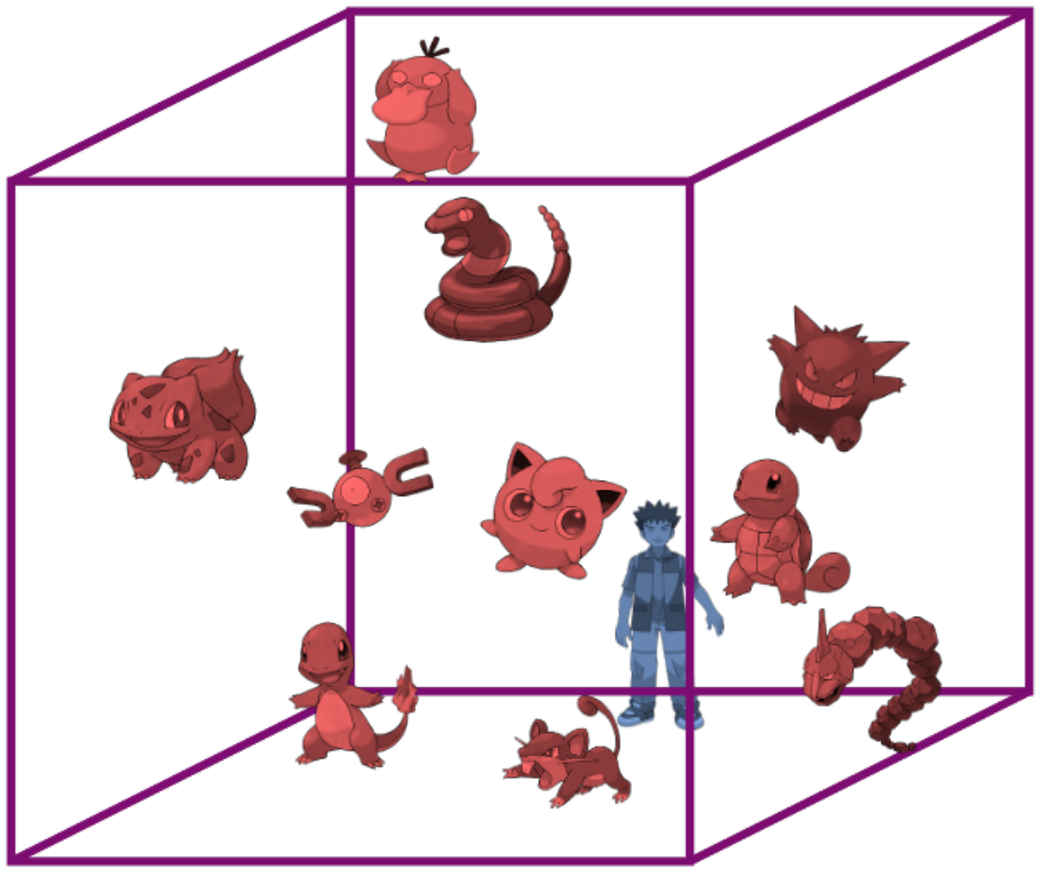
The matching quality is measured by the dot product ⟨ , ⟩

Model



The matched supply unit becomes unavailable for matching later

Model



The process repeats for a total of n time steps

Objective and Desiderata

Platform's Objective

maximize the expected average match utility **without knowing** P

$$\max_{\pi} \frac{1}{n} \mathbb{E}\left[\sum_{k=1}^{n} \langle D_k, S_{\pi(k)} \rangle\right]$$

equivalently, minimize the regret with respect to the **fluid benchmark**

fluid benchmark is the **optimal transport** between the demand and the supply distribution

Objective and Desiderata

Platform's Objective

maximize the expected average match utility **without knowing** P

$$\max_{\pi} \frac{1}{n} \mathbb{E}\left[\sum_{k=1}^{n} \langle D_k, S_{\pi(k)} \rangle\right]$$

equivalently, minimize the regret with respect to the **fluid benchmark**

fluid benchmark is the **optimal transport** between the demand and the supply distribution

Algorithmic Desiderata

`simple" dynamic matching algorithms with $o(1)$ **(vanishing) regret**

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

This Work

infinitely many types

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

Stochastic Online Matching

Feldman, Mehta, Mirrokni & Muthukrishnan (2009), Manshadi, Gharan & Saberi (2012), Jaillet & Lu (2014)

This Work

infinitely many types

spatial structure on type space

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

Stochastic Online Matching

Feldman, Mehta, Mirrokni & Muthukrishnan (2009), Manshadi, Gharan & Saberi (2012), Jaillet & Lu (2014)

Static Spatial Matching & Empirical Opt. Transport

Ajtai, Komlos & Tusnady (1984), Talagrand (1992,1994), Shor (1986, 1991), Ledoux (2019), Manole & Niles-Weed (2021)

This Work

infinitely many types

spatial structure on type space

dynamic demand arrivals

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

Stochastic Online Matching

Feldman, Mehta, Mirrokni & Muthukrishnan (2009), Manshadi, Gharan & Saberi (2012), Jaillet & Lu (2014)

Static Spatial Matching & Empirical Opt. Transport

Ajtai, Komlos & Tusnady (1984), Talagrand (1992,1994), Shor (1986, 1991), Ledoux (2019), Manole & Niles-Weed (2021)

Dynamic Spatial Matching with identical demand and supply distributions

Gupta, Guruganesh, Peng & Wajc (2019), Akbarpour, Alimohammad, Li & Saberi (2022), Kanoria (2022)

This Work

infinitely many types

spatial structure on type space

dynamic demand arrivals

different distributions

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

Stochastic Online Matching

Feldman, Mehta, Mirrokni & Muthukrishnan (2009), Manshadi, Gharan & Saberi (2012), Jaillet & Lu (2014)

Static Spatial Matching & Empirical Opt. Transport

Ajtai, Komlos & Tusnady (1984), Talagrand (1992,1994), Shor (1986, 1991), Ledoux (2019), Manole & Niles-Weed (2021)

Dynamic Spatial Matching with identical demand and supply distributions

Gupta, Guruganesh, Peng & Wajc (2019), Akbarpour, Alimohammad, Li & Saberi (2022), Kanoria (2022)

Stochastic Assignment Problem in 1D

Derman, Lieberman & Ross (1972), Su & Zenios (2005), Chen , Wang , Zeevi & Zhou (2021)

This Work

infinitely many types

spatial structure on type space

dynamic demand arrivals

different distributions

high dimensional features

Prior Work

Dynamic two-sided matching with few types

Talluri & van Ryzin (2004), Vera & Banerjee (2021), Bumpensanti & Wang (2021)

Stochastic Online Matching

Feldman, Mehta, Mirrokni & Muthukrishnan (2009), Manshadi, Gharan & Saberi (2012), Jaillet & Lu (2014)

Static Spatial Matching & Empirical Opt. Transport

Ajtai, Komlos & Tusnady (1984), Talagrand (1992,1994), Shor (1986, 1991), Ledoux (2019), Manole & Niles-Weed (2021)

Dynamic Spatial Matching with identical demand and supply distributions

Gupta, Guruganesh, Peng & Wajc (2019), Akbarpour, Alimohammad, Li & Saberi (2022), Kanoria (2022)

Stochastic Assignment Problem in 1D

Derman, Lieberman & Ross (1972), Su & Zenios (2005), Chen , Wang , Zeevi & Zhou (2021)

Feature-Based Dynamic Matching

Chen, Kanoria, Kumar & Zhang (2025), Yu and Yang (2025)

This Work

infinitely many types

spatial structure on type space

dynamic demand arrivals

different distributions

high dimensional features

unknown demand distribution

In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification

- Design a simple and near-optimal policy ROAR

  - Textbook like proof for ROAR

  - No separate learning terms in the regret analysis



Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

1. It simulates a forecast of demand units

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

3. Makes assignment decisions based on the solution of the optimization problem

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

3. Makes assignment decisions based on the solution of the optimization problem

Let the supply units be drawn uniformly from Unif($[0,1]$). Let the true demand distribution be Unif($[\epsilon, 1-\epsilon]$) for some $\epsilon > 0$. Let the misspecified demand distribution that SOAR uses be Unif($[0,1]$). Then the regret of SOAR scales as

$$\lim_{n\to\infty} \text{Reg}(\text{SOAR}; P, Q, P') = \Theta(\epsilon^2)$$

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

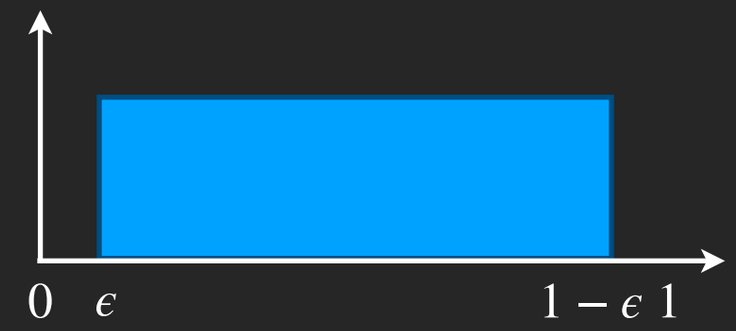
1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

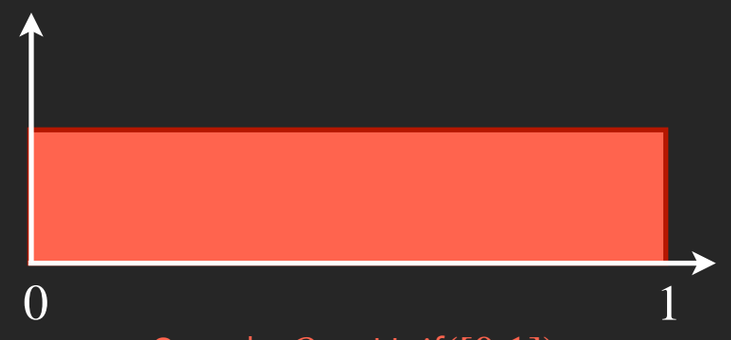
3. Makes assignment decisions based on the solution of the optimization problem

Let the supply units be drawn uniformly from Unif($[0,1]$). Let the true demand distribution be Unif($[\epsilon, 1-\epsilon]$) for some $\epsilon > 0$. Let the misspecified demand distribution that SOAR uses be Unif($[0,1]$). Then the regret of SOAR scales as

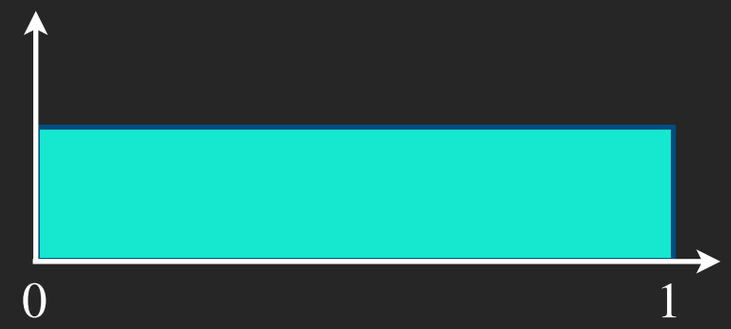
$$\lim_{n\to\infty} \text{Reg}(\text{SOAR}; P, Q, P') = \Theta(\epsilon^2)$$



True Demand $P \sim$ Unif($[\epsilon, 1-\epsilon]$)



Supply $Q \sim$ Unif($[0,1]$)



Misspecified Demand $P' \sim$ Unif($[0,1]$)

Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

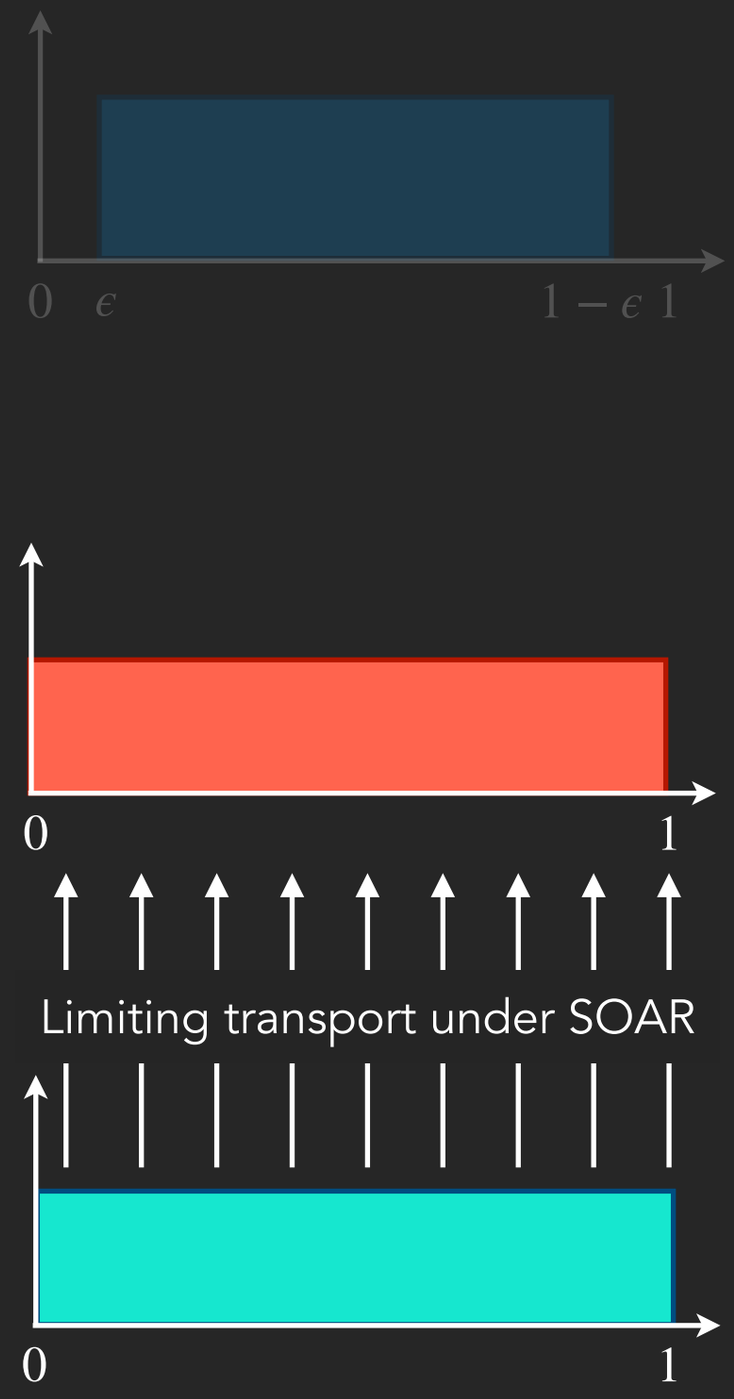
1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

3. Makes assignment decisions based on the solution of the optimization problem

Let the supply units be drawn uniformly from Unif([0,1]). Let the true demand distribution be Unif($[\epsilon, 1-\epsilon]$) for some $\epsilon > 0$. Let the misspecified demand distribution that SOAR uses be Unif([0,1]). Then the regret of SOAR scales as

$$\lim_{n\to\infty} \text{Reg}(\text{SOAR}; P, Q, P') = \Theta(\epsilon^2)$$



Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

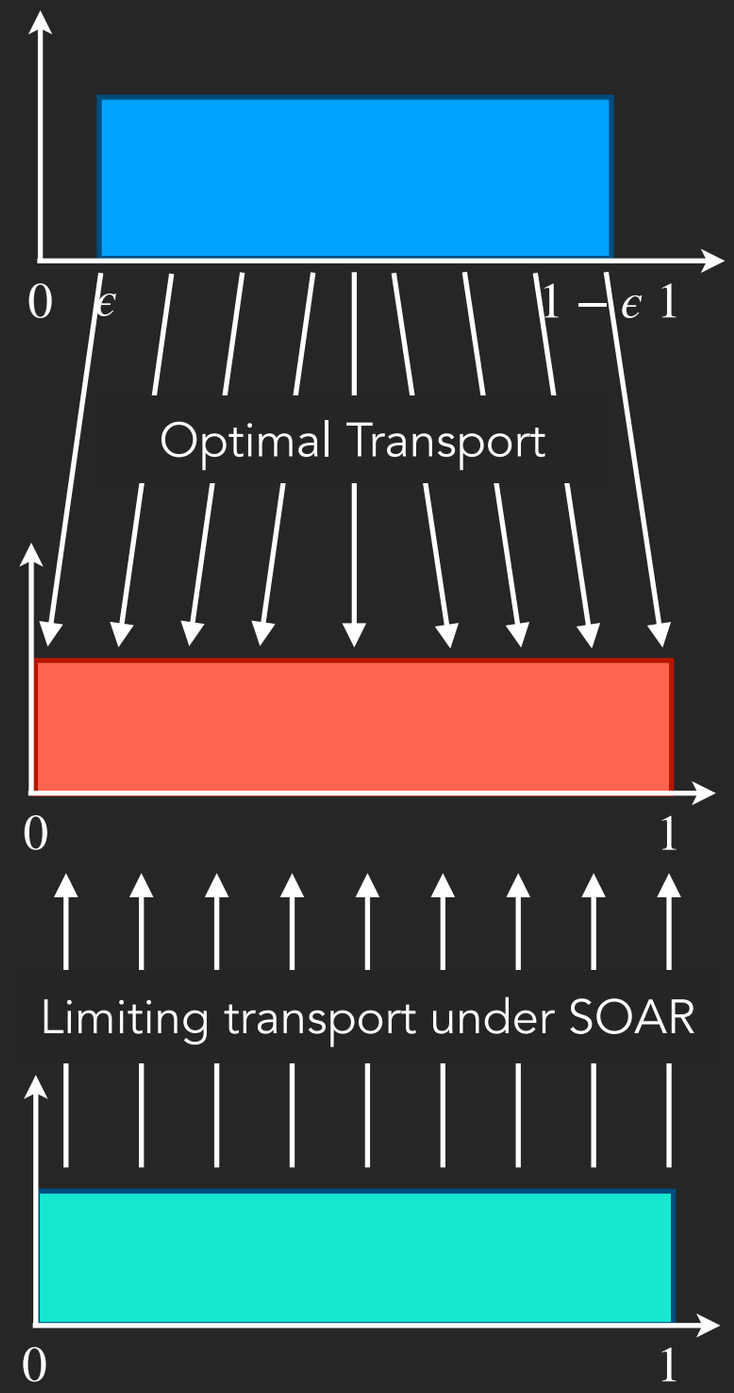
1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

3. Makes assignment decisions based on the solution of the optimization problem

Let the supply units be drawn uniformly from Unif($[0,1]$). Let the true demand distribution be Unif($[\epsilon, 1-\epsilon]$) for some $\epsilon > 0$. Let the misspecified demand distribution that SOAR uses be Unif($[0,1]$). Then the regret of SOAR scales as

$$\lim_{n\to\infty} \text{Reg}(\text{SOAR}; P, Q, P') = \Theta(\epsilon^2)$$



Fragility of SOAR to demand misspecification

SOAR is a near-optimal algorithm with known demand distribution.

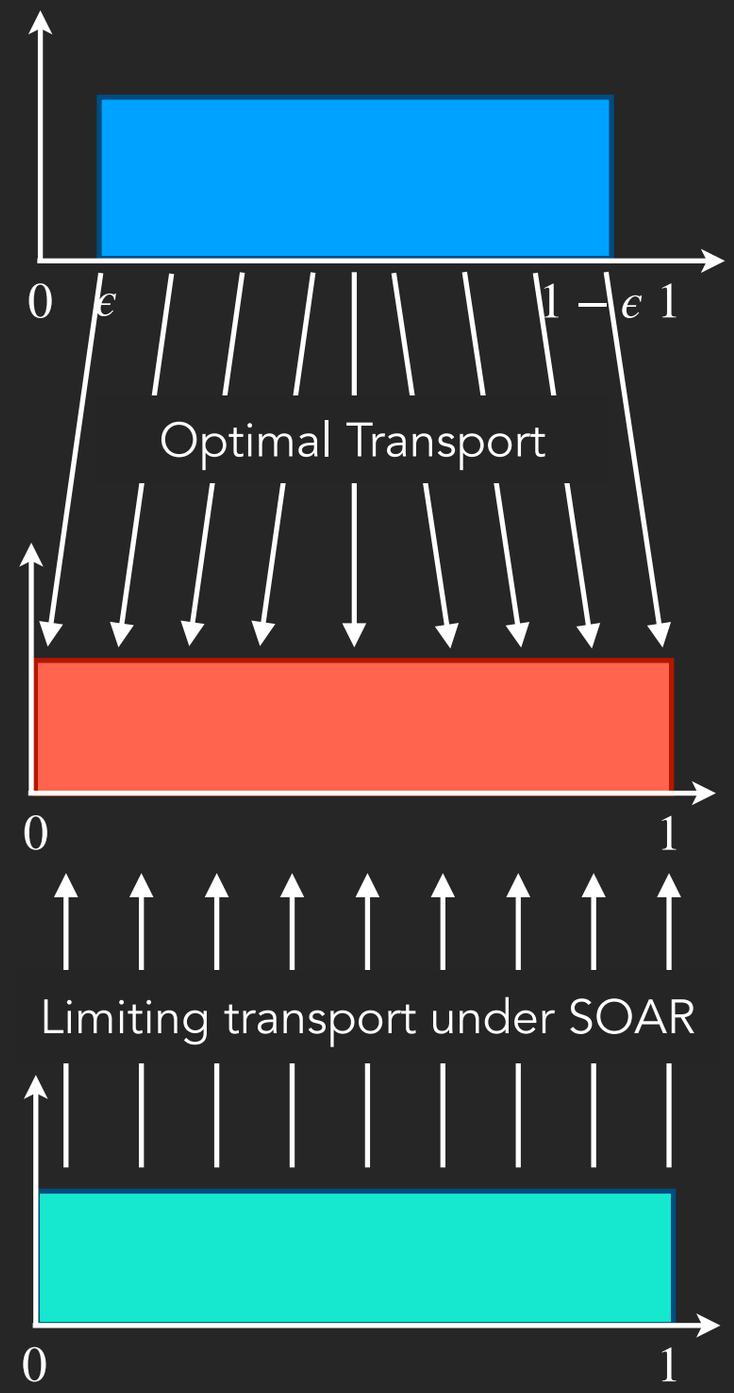
1. It simulates a forecast of demand units

2. Solves an optimization problem with simulated demand units

3. Makes assignment decisions based on the solution of the optimization problem

Let the supply units be drawn uniformly from Unif($[0,1]$). Let the true demand distribution be Unif($[\epsilon, 1-\epsilon]$) for some $\epsilon > 0$. Let the misspecified demand distribution that SOAR uses be Unif($[0,1]$). Then the regret of SOAR scales as

$$\lim_{n\to\infty} \text{Reg}(\text{SOAR}; P, Q, P') = \Theta(\epsilon^2)$$



In A Nutshell

- We study centralized dynamic matching with heterogeneous demand and supply *(unknown demand distribution)*

- Previous near-optimal algorithms (with known demand distribution) can be highly suboptimal under misspecification

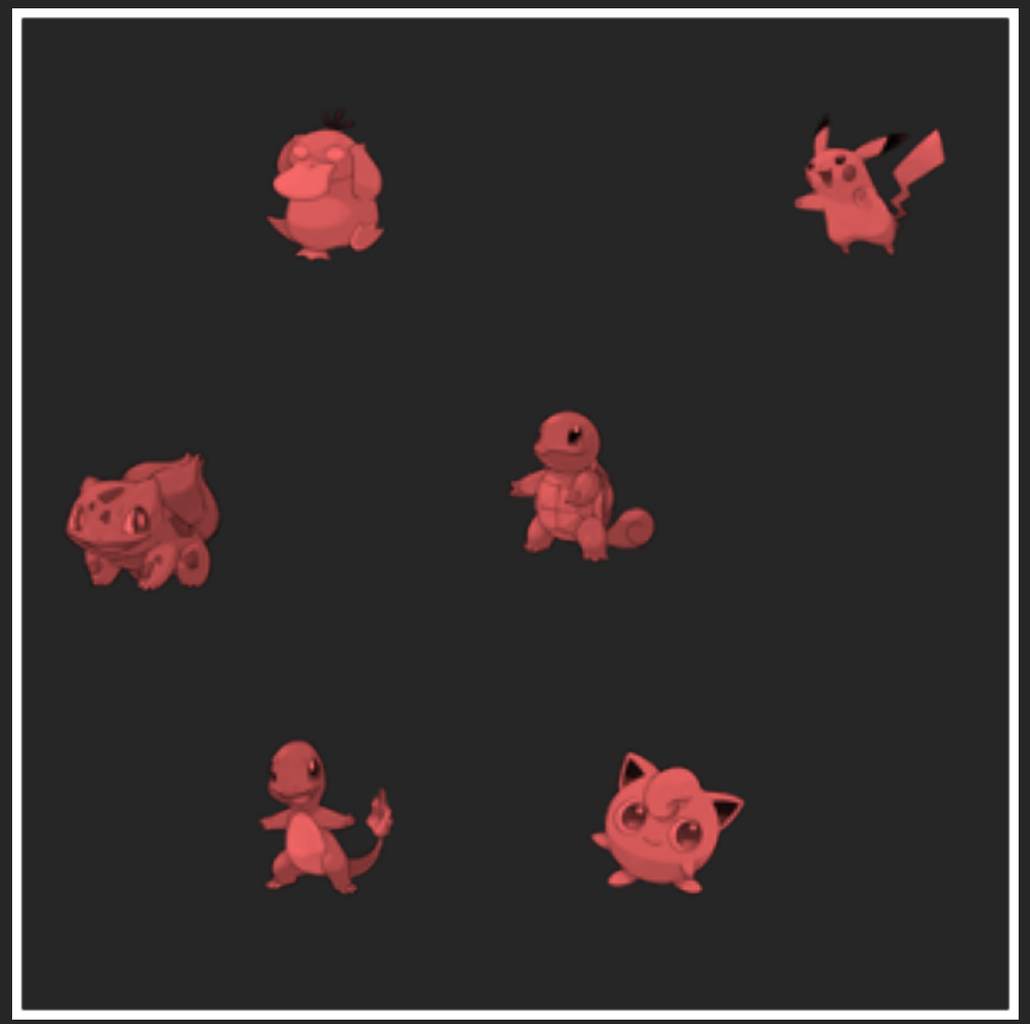
- Design a simple and near-optimal policy **ROAR**

  - Textbook like proof for ROAR

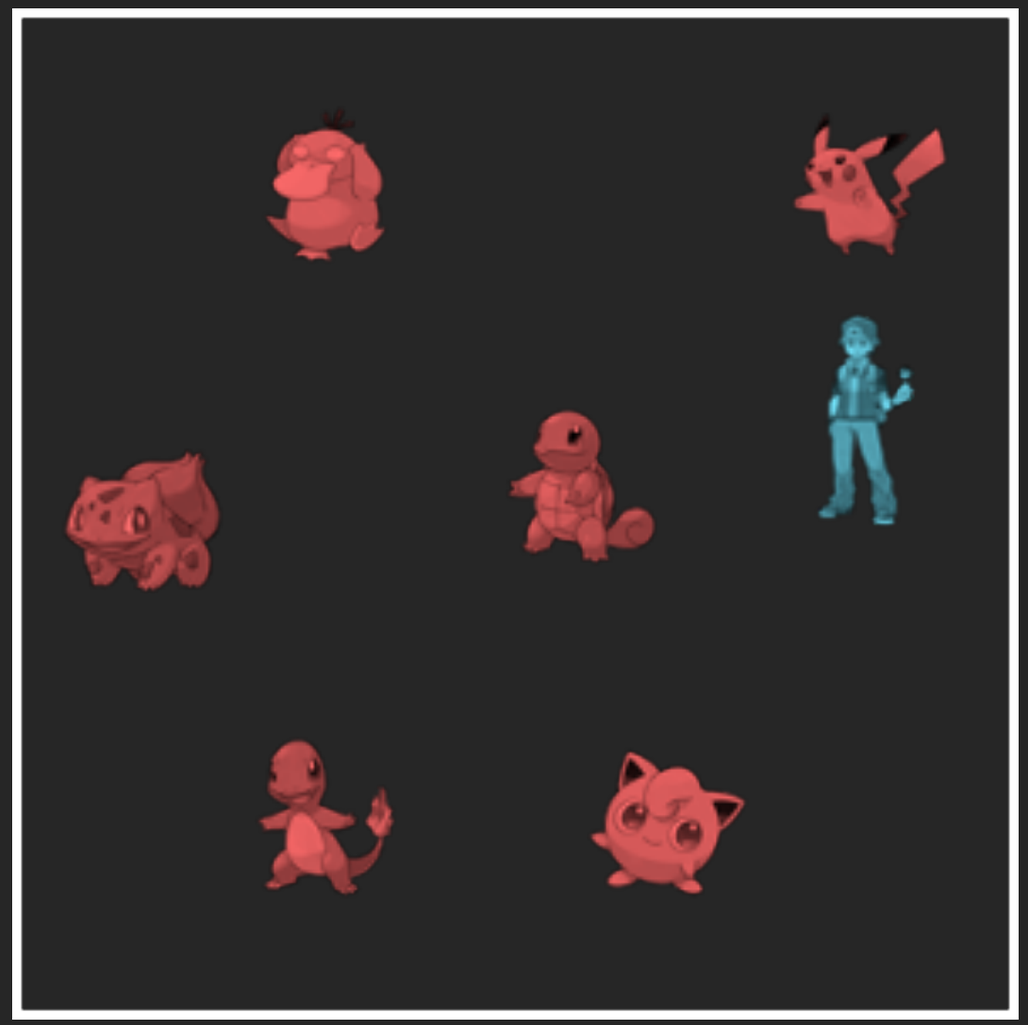
  - No separate learning terms in the regret analysis



ROAR



ROAR

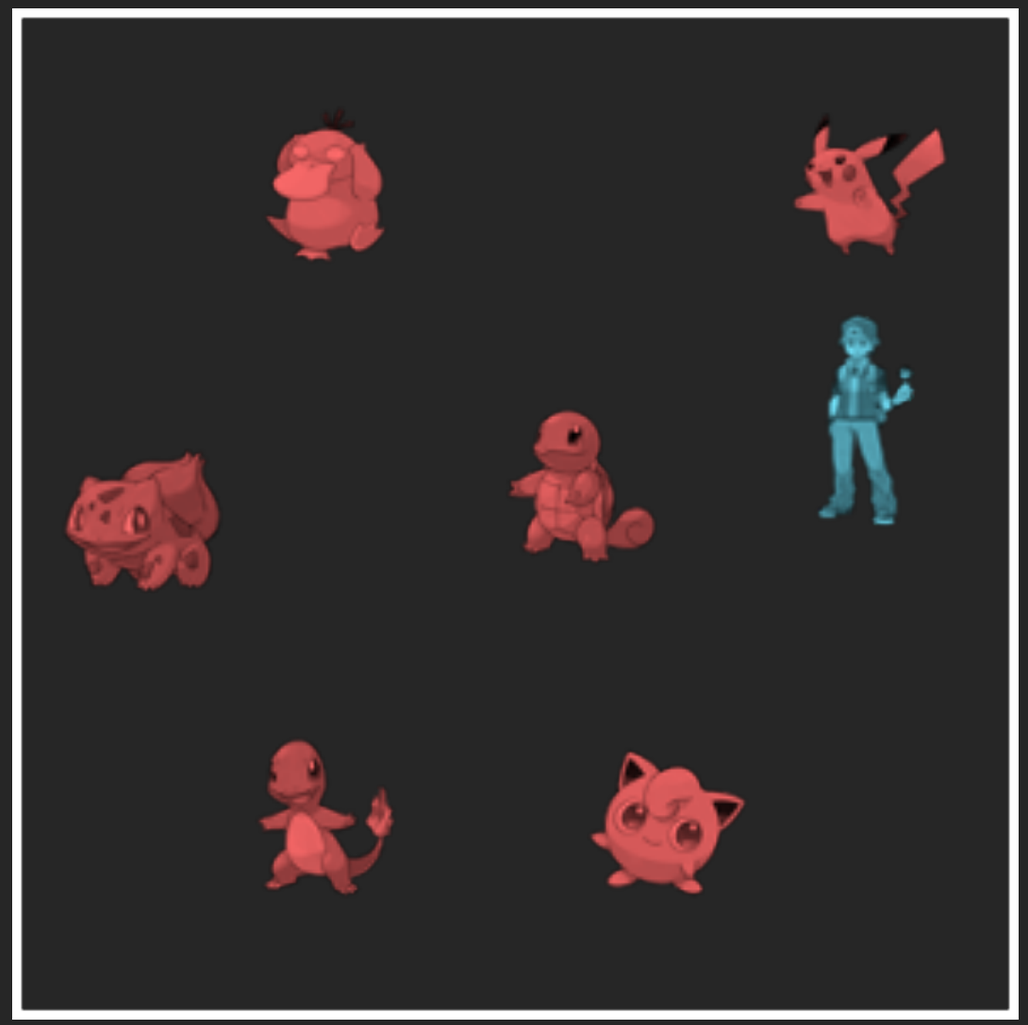


a demand unit arrives

ROAR

Resample

the supply units

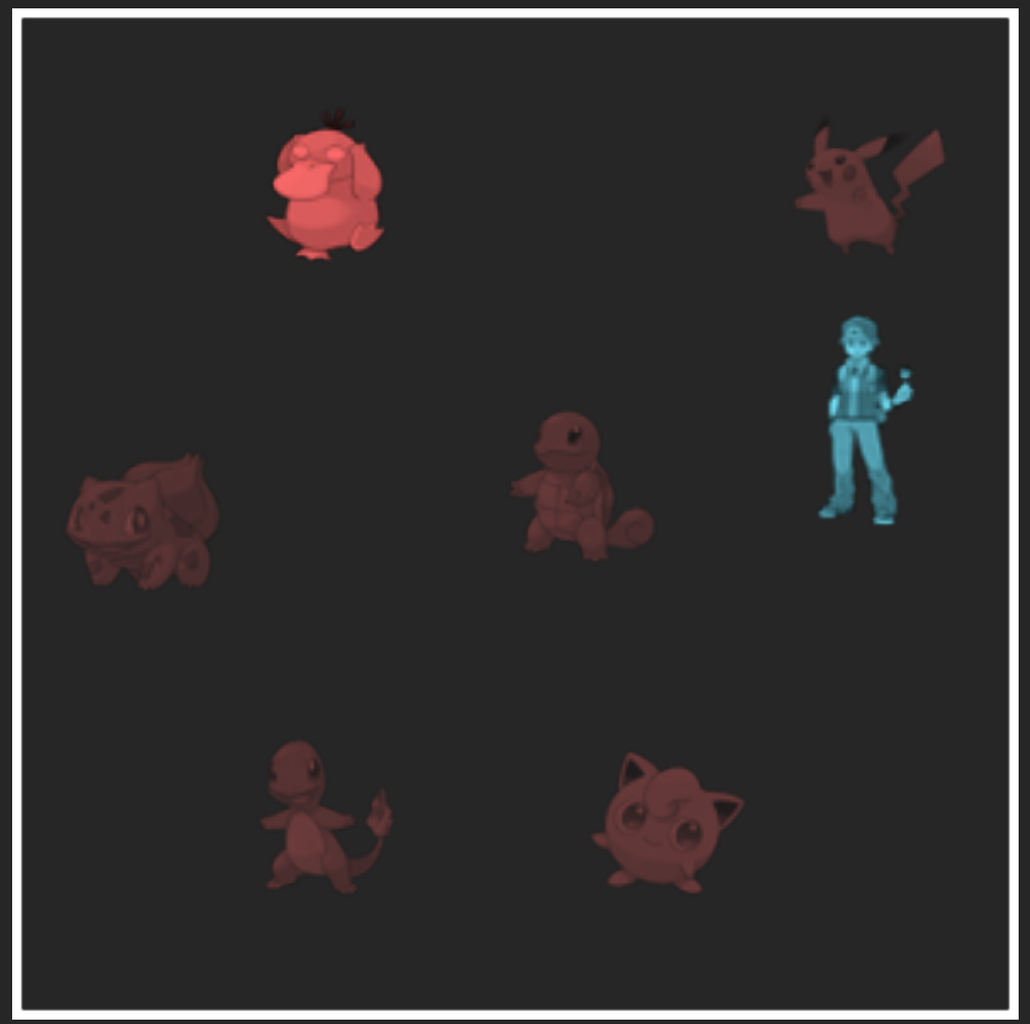


a demand unit arrives

ROAR

Resample

the supply units



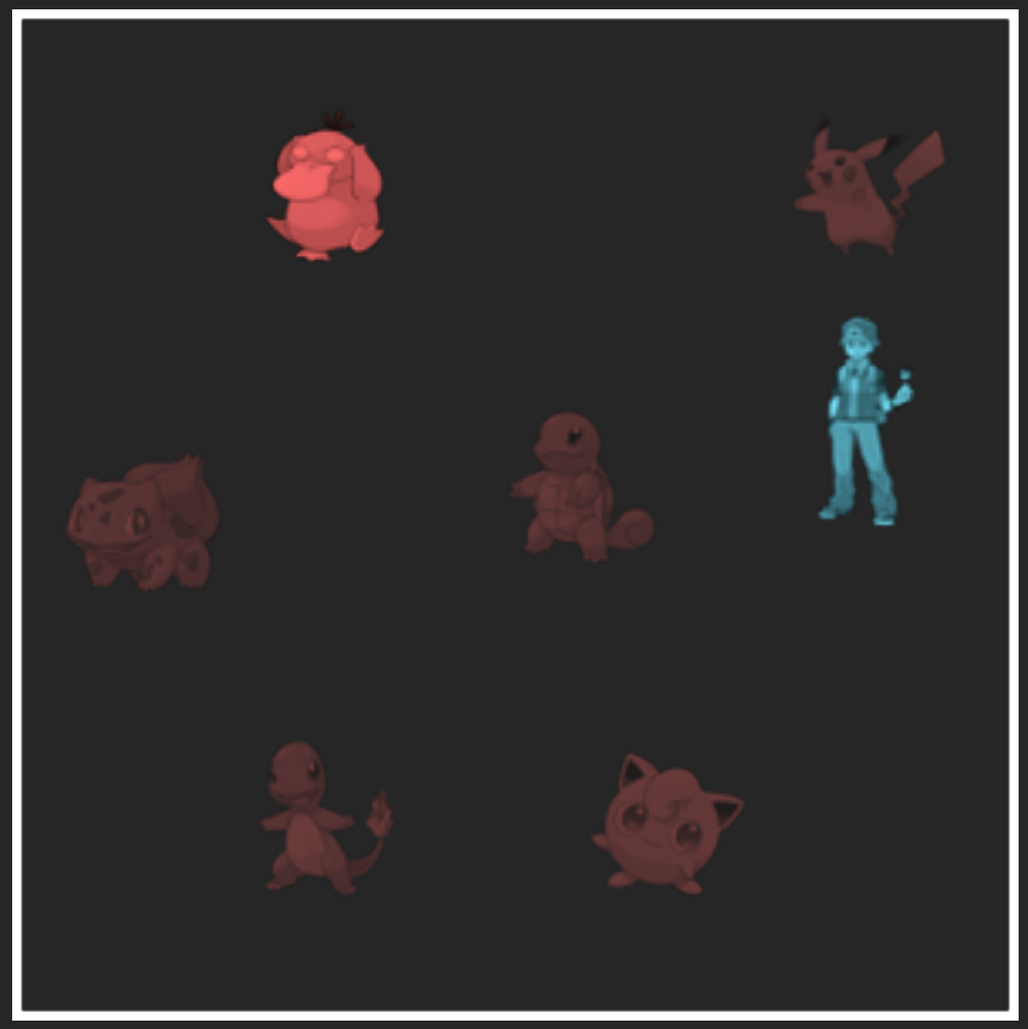
ROAR

Resample

the supply units

Optimize

for the scenario



compute the optimal
weighted bipartite matching

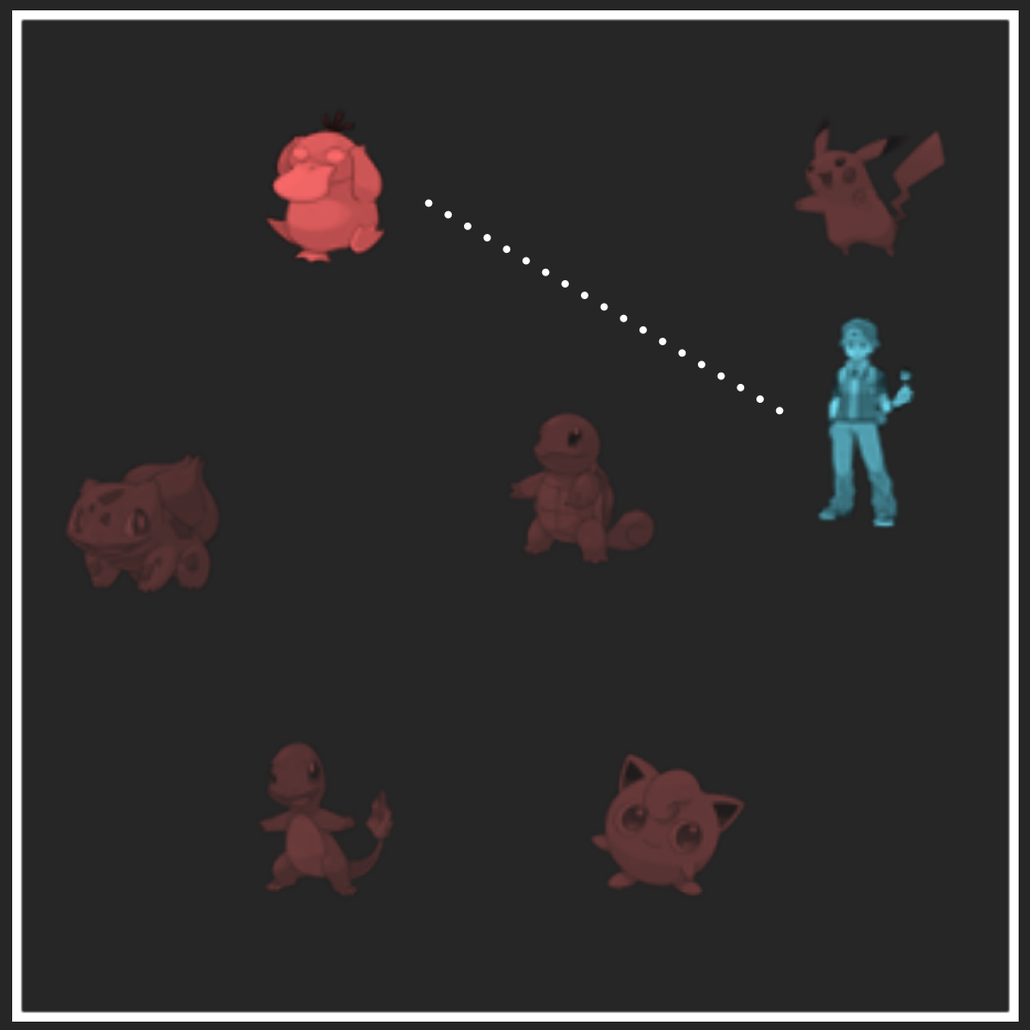
ROAR

Resample

the supply units

Optimize

for the scenario



compute the optimal
weighted bipartite matching

ROAR

Resample

the supply units

Optimize

for the scenario

Assign

based on optimal matching



ROAR

Resample

the supply units

Optimize

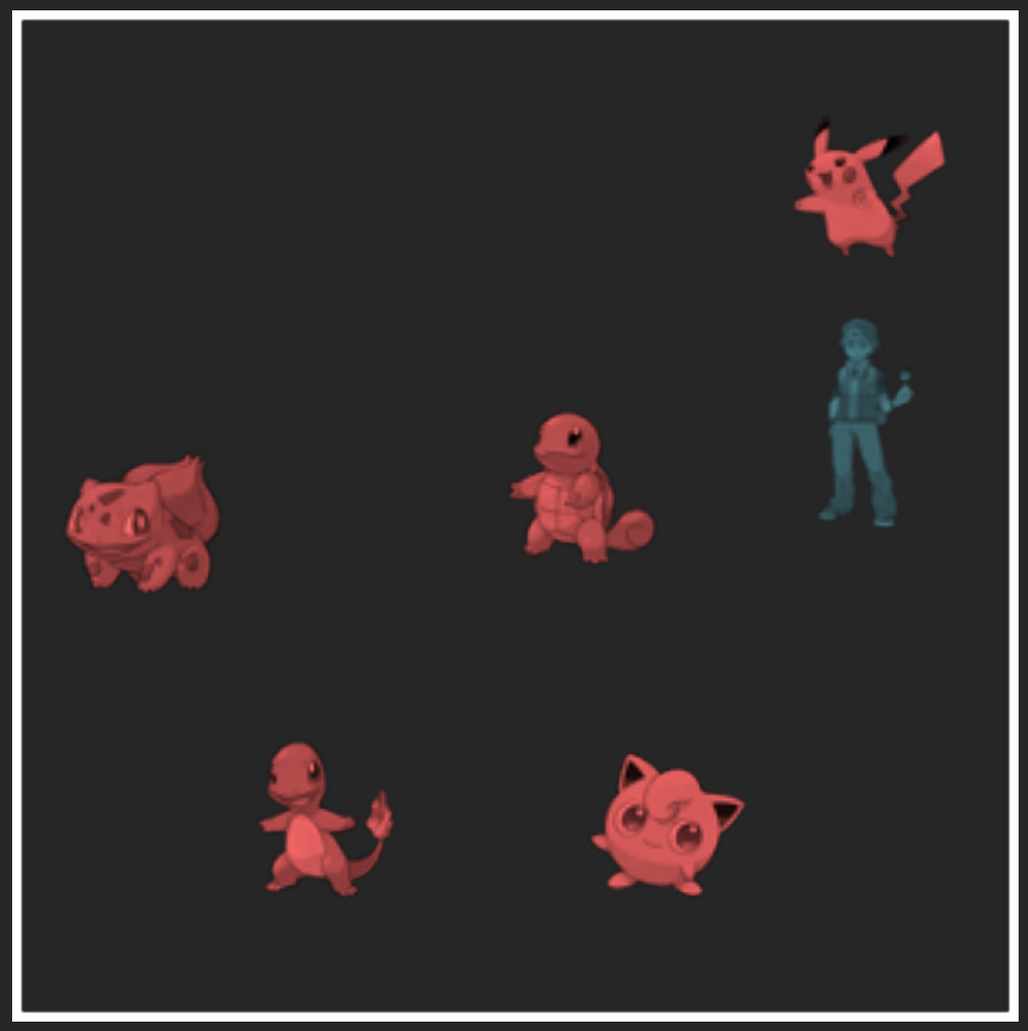
for the scenario

Assign

based on optimal matching

Repeat

with remaining supply



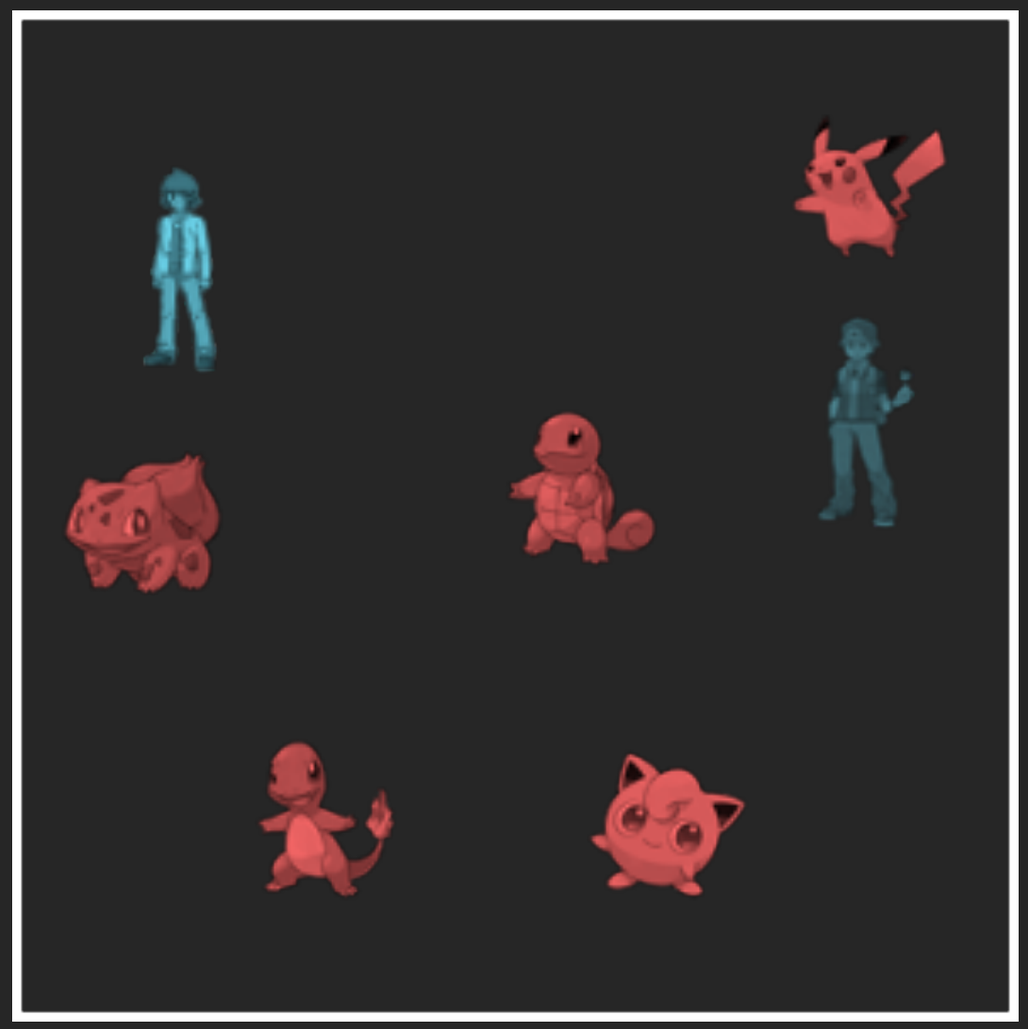
ROAR

Resample
the supply units

Optimize
for the scenario

Assign
based on optimal matching

Repeat
with remaining supply

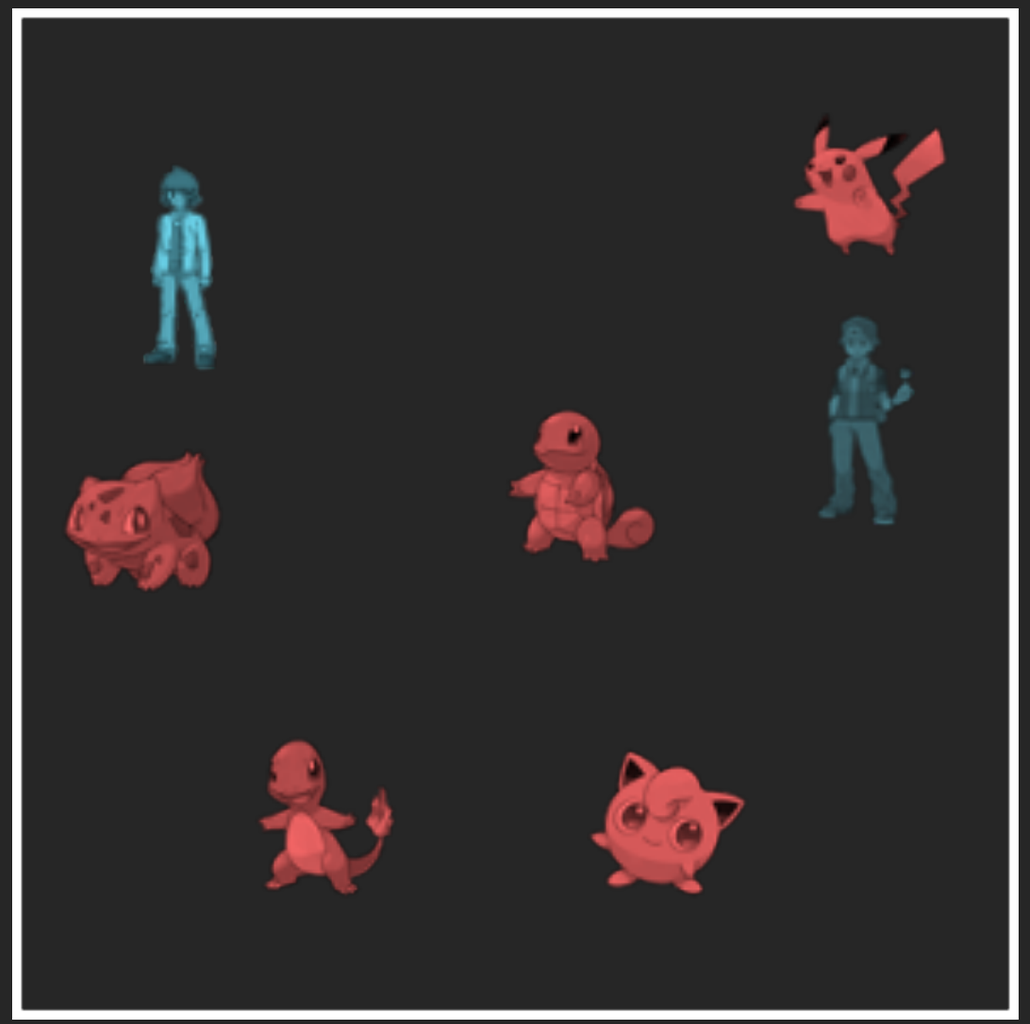


a demand unit arrives

ROAR

Resample

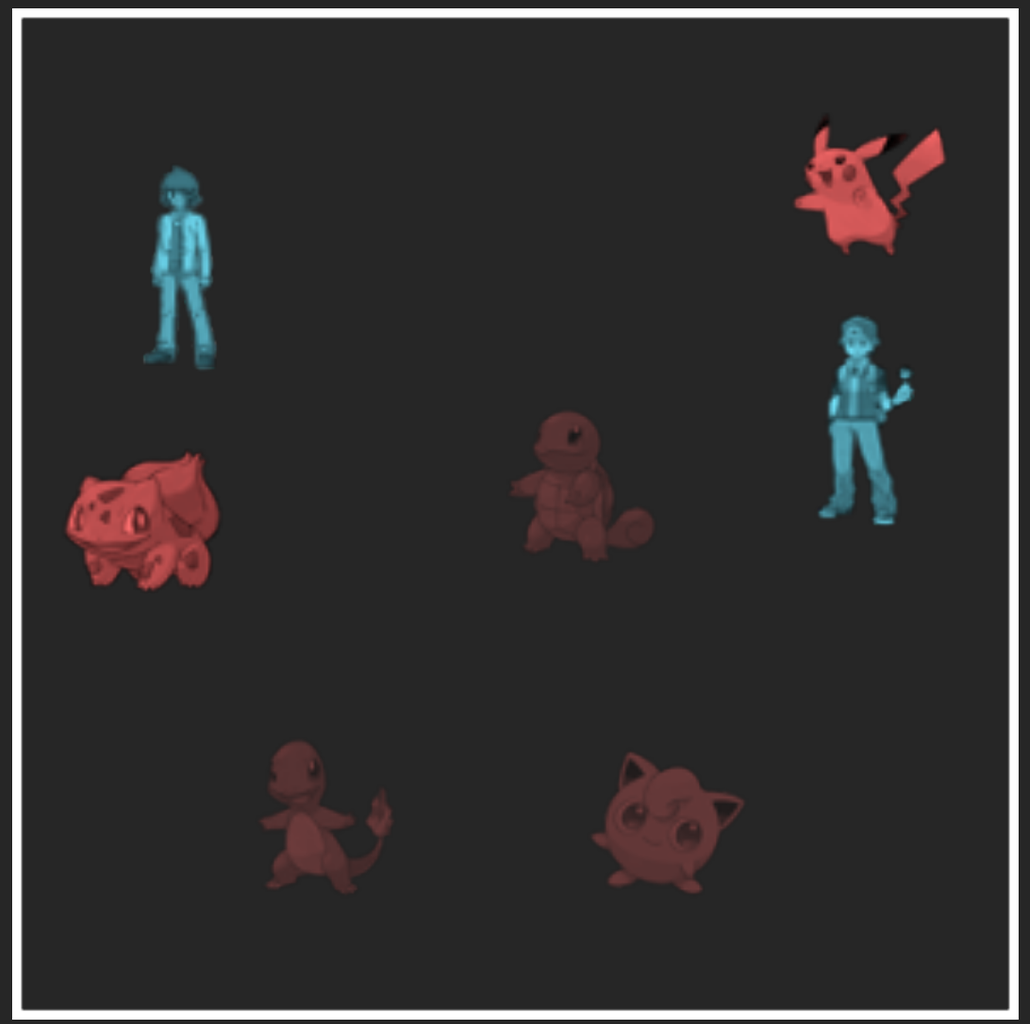
the supply units



ROAR

Resample

the supply units



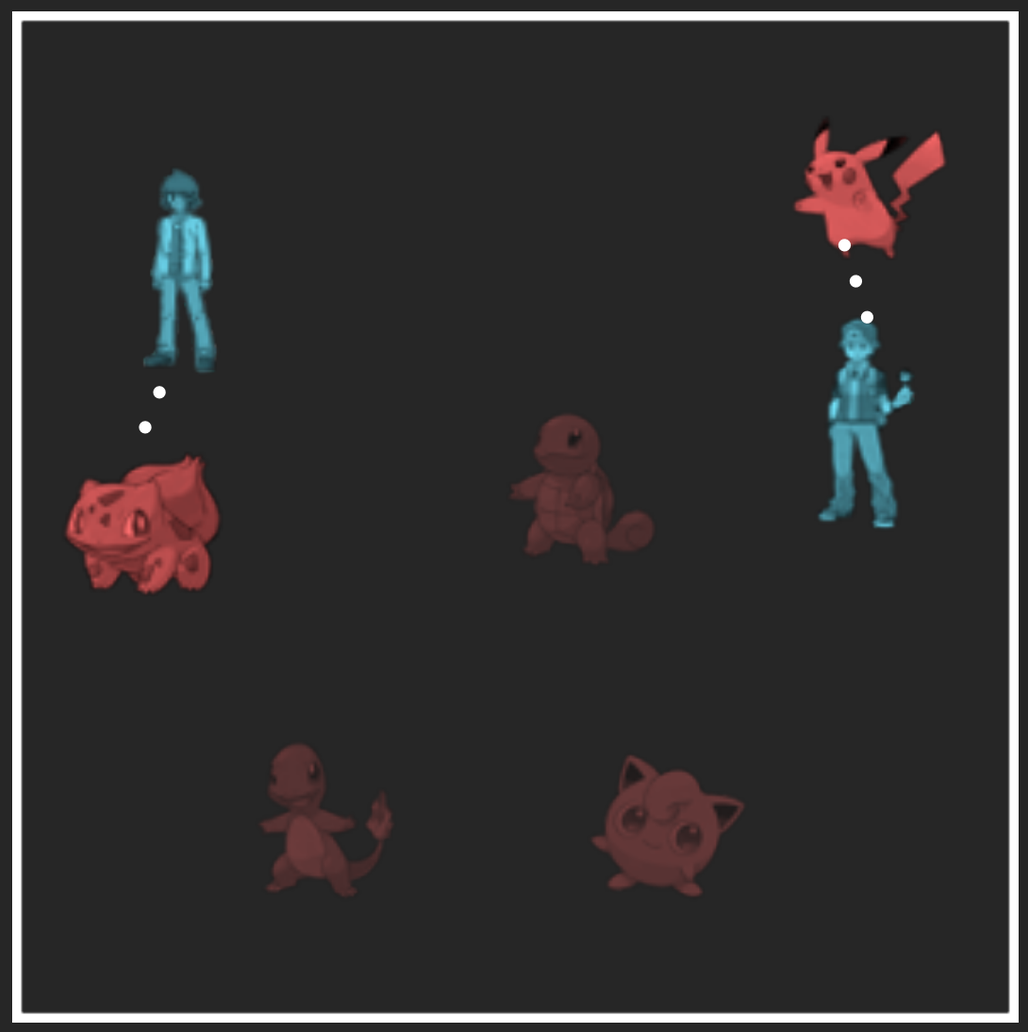
ROAR

Resample

the supply units

Optimize

for the scenario



compute the optimal
weighted bipartite matching

ROAR

Resample

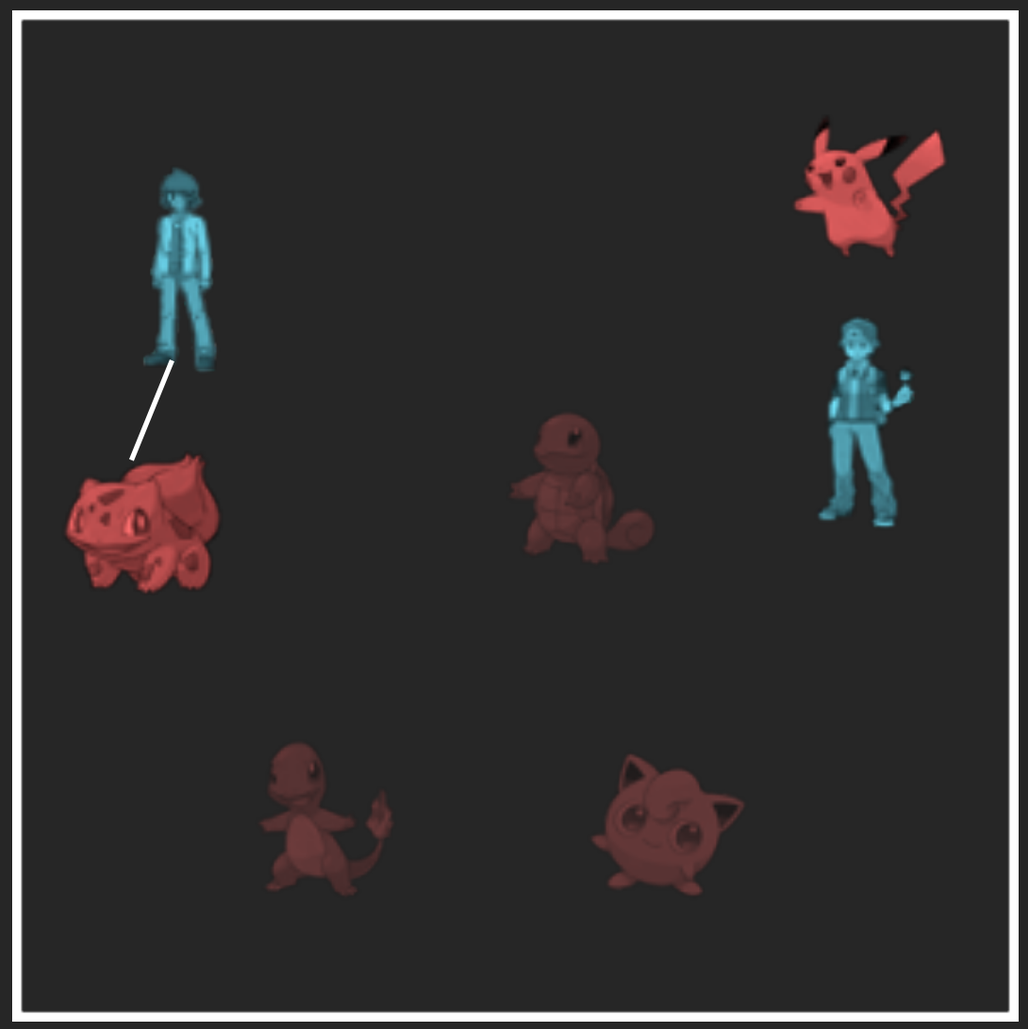
the supply units

Optimize

for the scenario

Assign

based on optimal matching



ROAR

Resample

the supply units

Optimize

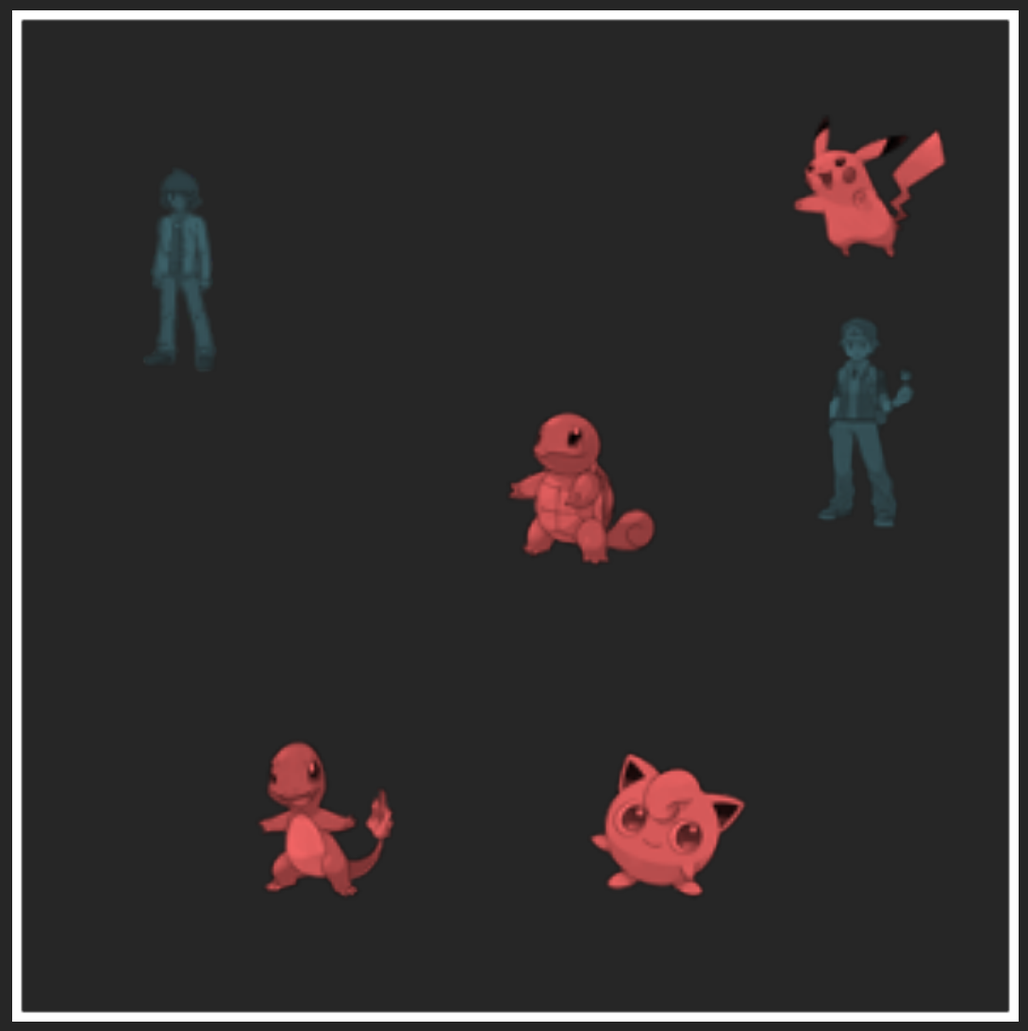
for the scenario

Assign

based on optimal matching

Repeat

with remaining supply



Main Result

Characterizing the average match quality of ROAR

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively.

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Then the expected average match quality under ROAR is

$$U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n} U^{\text{off}}_{\min\{k,n-k+1\}} = \frac{2}{n}\sum_{k=1}^{n/2} U_k^{\text{off}}$$

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Then the expected average match quality under ROAR is

$$U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n} U^{\text{off}}_{\min\{k,n-k+1\}} = \frac{2}{n}\sum_{k=1}^{n/2} U_k^{\text{off}}$$

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value.

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Then the expected average match quality under ROAR is

$$U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n} U^{\text{off}}_{\min\{k,n-k+1\}} = \frac{2}{n}\sum_{k=1}^{n/2} U_k^{\text{off}}$$

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value. Then the expected expected regret of ROAR is

$$\text{Reg}_n(\text{ROAR}) = U_\infty^{\text{off}} - U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n}\left(U_\infty^{\text{off}} - U^{\text{off}}_{\min\{k,n-k+1\}}\right) = \frac{2}{n}\sum_{k=1}^{n/2}\left(U_\infty^{\text{off}} - U_k^{\text{off}}\right)$$

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Then the expected average match quality under ROAR is

$$U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n} U^{\text{off}}_{\min\{k,n-k+1\}} = \frac{2}{n}\sum_{k=1}^{n/2} U_k^{\text{off}}$$

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value. Then the expected expected regret of ROAR is

$$\text{Reg}_n(\text{ROAR}) = U_\infty^{\text{off}} - U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n}\left(U_\infty^{\text{off}} - U^{\text{off}}_{\min\{k,n-k+1\}}\right) = \frac{2}{n}\sum_{k=1}^{n/2}\left(U_\infty^{\text{off}} - U_k^{\text{off}}\right)$$

- Regret scaling depends only on the convergence of the finite offline matching to the fluid benchmark $\Longrightarrow$ Can leverage convergence results in empirical OT literature.

Main Result

Characterizing the average match quality of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Then the expected average match quality under ROAR is

$$U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n} U^{\text{off}}_{\min\{k,n-k+1\}} = \frac{2}{n}\sum_{k=1}^{n/2} U_k^{\text{off}}$$

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value. Then the expected expected regret of ROAR is

$$\text{Reg}_n(\text{ROAR}) = U_\infty^{\text{off}} - U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n}\left(U_\infty^{\text{off}} - U^{\text{off}}_{\min\{k,n-k+1\}}\right) = \frac{2}{n}\sum_{k=1}^{n/2}\left(U_\infty^{\text{off}} - U_k^{\text{off}}\right)$$

- Regret scaling depends only on the convergence of the finite offline matching to the fluid benchmark $\Longrightarrow$ Can leverage convergence results in empirical OT literature.

- No separate learning term that needs to be analyzed/bounded!

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Proof of the Main Result

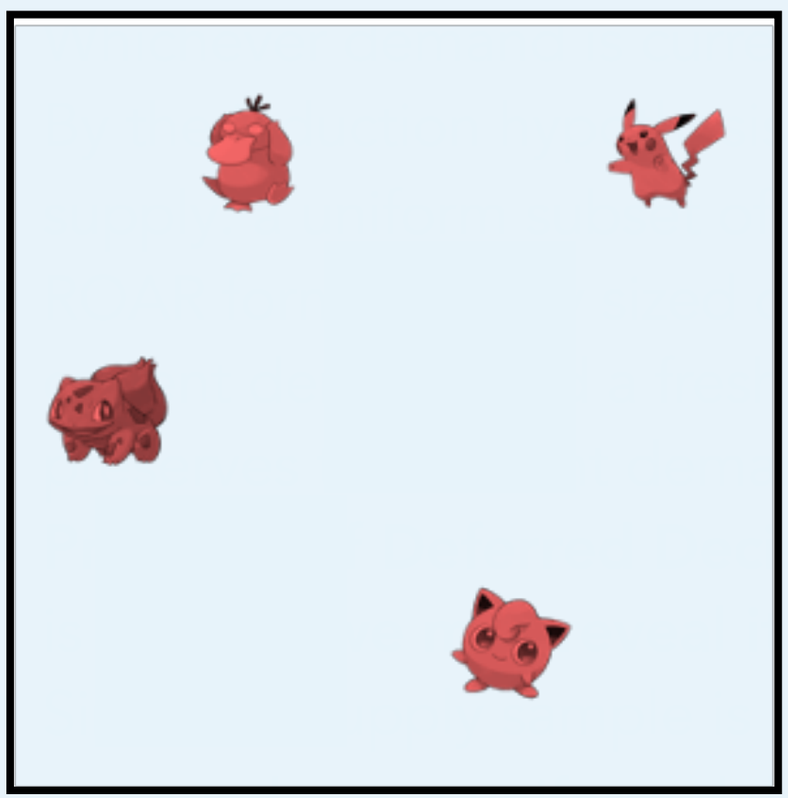
Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.

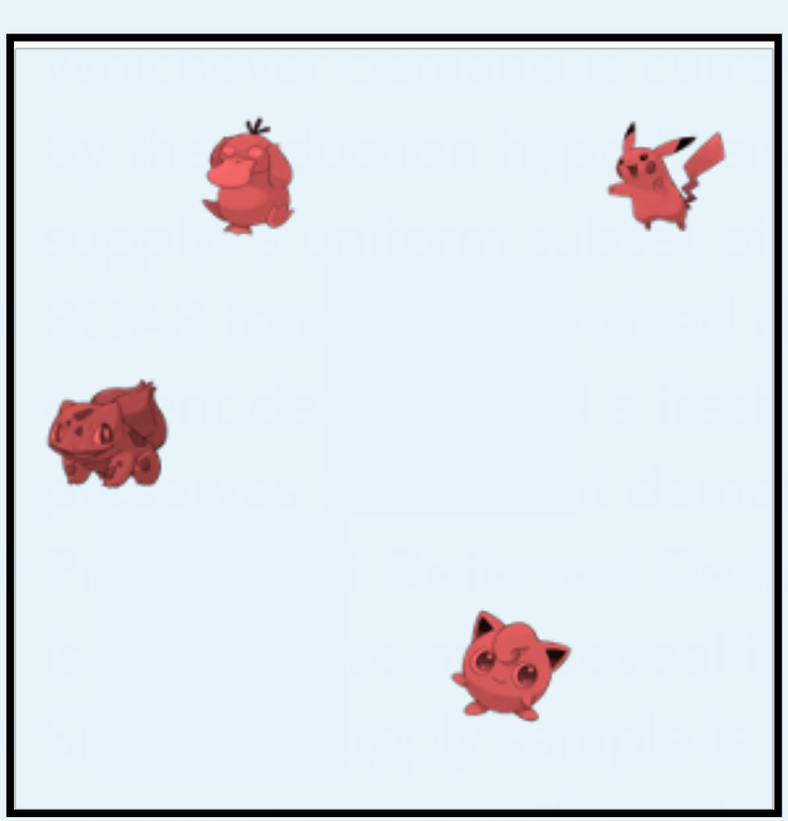


Initial set of supply units

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.



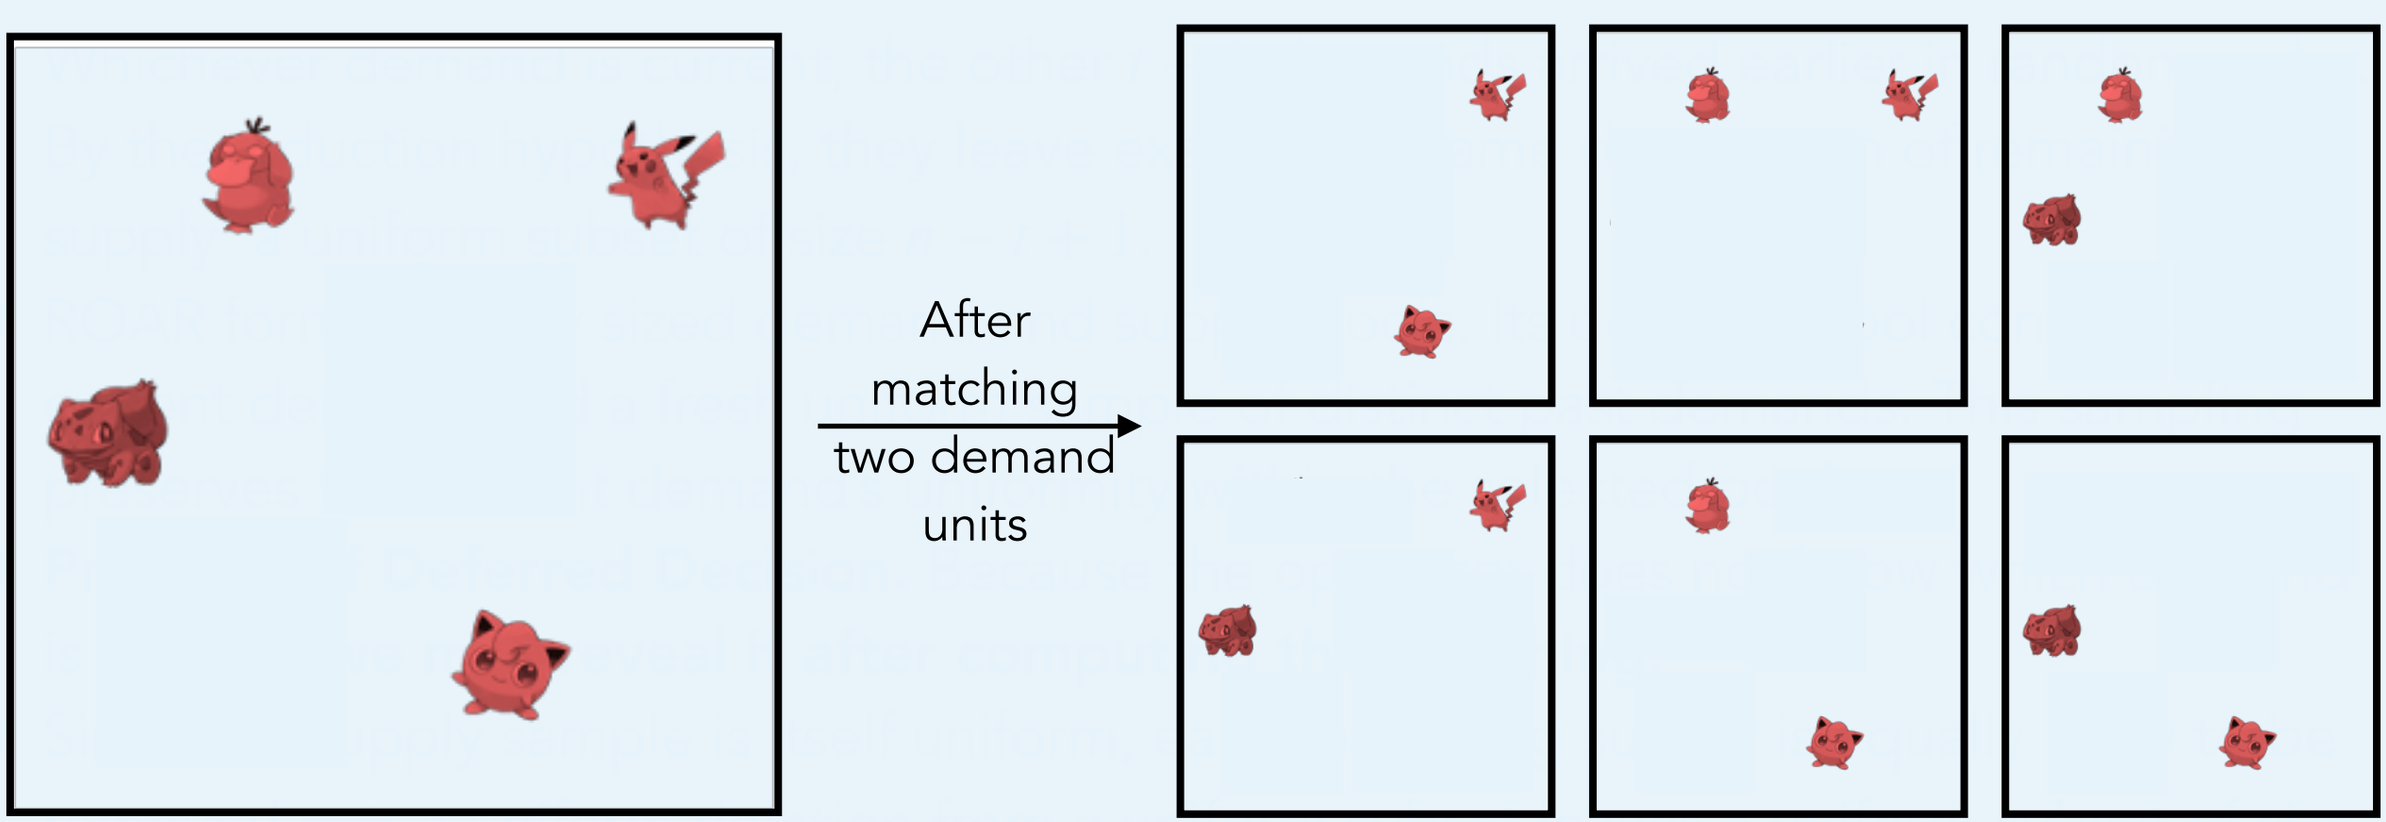
Initial set of supply units

After matching two demand units

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.



Initial set of supply units

Each of them is equally likely

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.

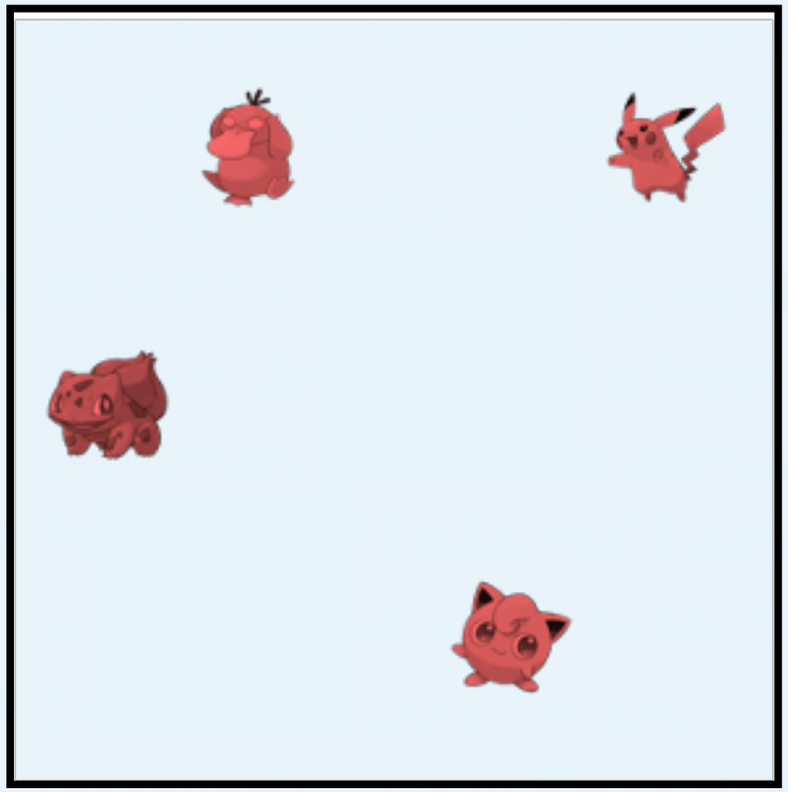
Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n - 1$ original supply units.

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.

Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n - 1$ original supply units.

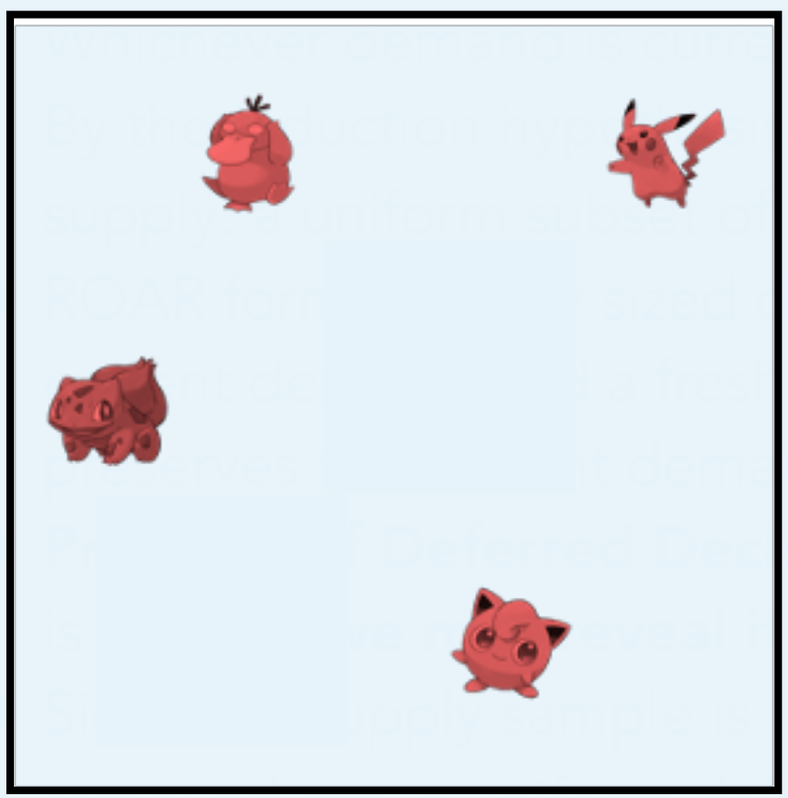


Proof of the Main Result

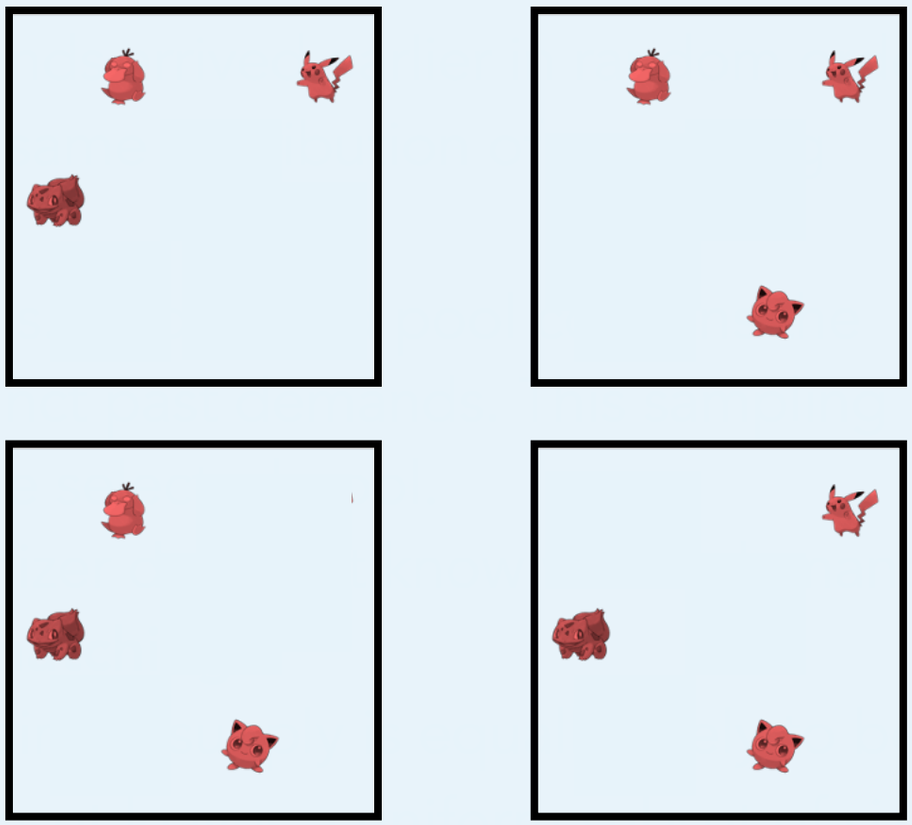
Fix the supply and demand set, and the demand arrives in uniformly random order.

Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.

Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n - 1$ original supply units.



Randomly sampled a supply unit to match



Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.
Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.
Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n - 1$ original supply units.

Induction Step.

- Fix the unordered collection of the first t demands. Whichever demand is current, the other $t - 1$ demands leave the same uniform distribution of supply units (by induction hypothesis). The remaining supply units therefore provides no information about which demand is current.

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.
Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.
Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n - 1$ original supply units.
Induction Step.

- Fix the unordered collection of the first t demands. Whichever demand is current, the other $t - 1$ demands leave the same uniform distribution of supply units (by induction hypothesis). The remaining supply units therefore provides no information about which demand is current.
- **Principle of Deferred Decisions.** Compute the matching without using which demand is current. Treat the entire matching as a black box, where for the sake of analysis, we do not really see who is being matched.

Proof of the Main Result

Fix the supply and demand set, and the demand arrives in uniformly random order.
Induction Claim. After processing a subset of demand units in random order, the remaining supply is a uniformly random subset of the original set, of the appropriate size.
Base Case. In the first round, ROAR matches the arriving demand to a uniformly chosen supply. Removing that supply leaves a uniform subset of $n-1$ original supply units.
Induction Step.

- Fix the unordered collection of the first t demands. Whichever demand is current, the other $t-1$ demands leave the same uniform distribution of supply units (by induction hypothesis). The remaining supply units therefore provides no information about which demand is current.
- **Principle of Deferred Decisions.** Compute the matching without using which demand is current. Treat the entire matching as a black box, where for the sake of analysis, we do not really see who is being matched.
- Each supply unit is equally likely to be matched. Uniform supply sampling makes the removal uniform among all remaining supply. Thus a uniform subset of size $n-t+1$ becomes a uniform subset of size $n-t$.

Proof of the Main Result

Expected reward. A uniformly chosen edge earns the average reward of the optimal matching. The selected demand pool is uniform, and the supply pool is uniform even after specifying the demand pool, so the two pools are independent. Averaging over independent i.i.d. initial populations gives

$$\mathbb{E}[\text{reward in round } t] = U^{\text{off}}_{\min\{t,n-t+1\}}.$$

Averaging over the rounds proves the identity.

Proof of the Main Result

Expected reward. A uniformly chosen edge earns the average reward of the optimal matching. The selected demand pool is uniform, and the supply pool is uniform even after specifying the demand pool, so the two pools are independent. Averaging over independent i.i.d. initial populations gives

$$\mathbb{E}[\text{reward in round } t] = U^{\text{off}}_{\min\{t,n-t+1\}}.$$

Averaging over the rounds proves the identity.

The right conditioning is the key to the proof. We should not condition on the full ordered history of demand. Because the remaining set of supply units need not be uniform conditional on the full ordered history of demand.

Proof of the Main Result

Expected reward. A uniformly chosen edge earns the average reward of the optimal matching. The selected demand pool is uniform, and the supply pool is uniform even after specifying the demand pool, so the two pools are independent. Averaging over independent i.i.d. initial populations gives

$$\mathbb{E}[\text{reward in round } t] = U^{\text{off}}_{\min\{t,n-t+1\}}.$$

Averaging over the rounds proves the identity.

The right conditioning is the key to the proof. We should not condition on the full ordered history of demand. Because the remaining set of supply units need not be uniform conditional on the full ordered history of demand.

This is not an issue when demand distribution is known because one can always simulate fresh demand scenarios and hence the remaining set of supply units are uniform.

Near-optimal performance of ROAR

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value. Then the expected expected regret of ROAR is

$$\text{Reg}_n(\text{ROAR}) = U_\infty^{\text{off}} - U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n}\left(U_\infty^{\text{off}} - U_{\min\{k,n-k+1\}}^{\text{off}}\right) = \frac{2}{n}\sum_{k=1}^{n/2}\left(U_\infty^{\text{off}} - U_k^{\text{off}}\right)$$

Near-optimal performance of ROAR

Characterizing the regret of ROAR

Let U_k^{off} denote the expected average match quality when matching k demand and supply units drawn i.i.d from P and Q respectively. Let U_∞^{off} be the fluid benchmark value. Then the expected expected regret of ROAR is

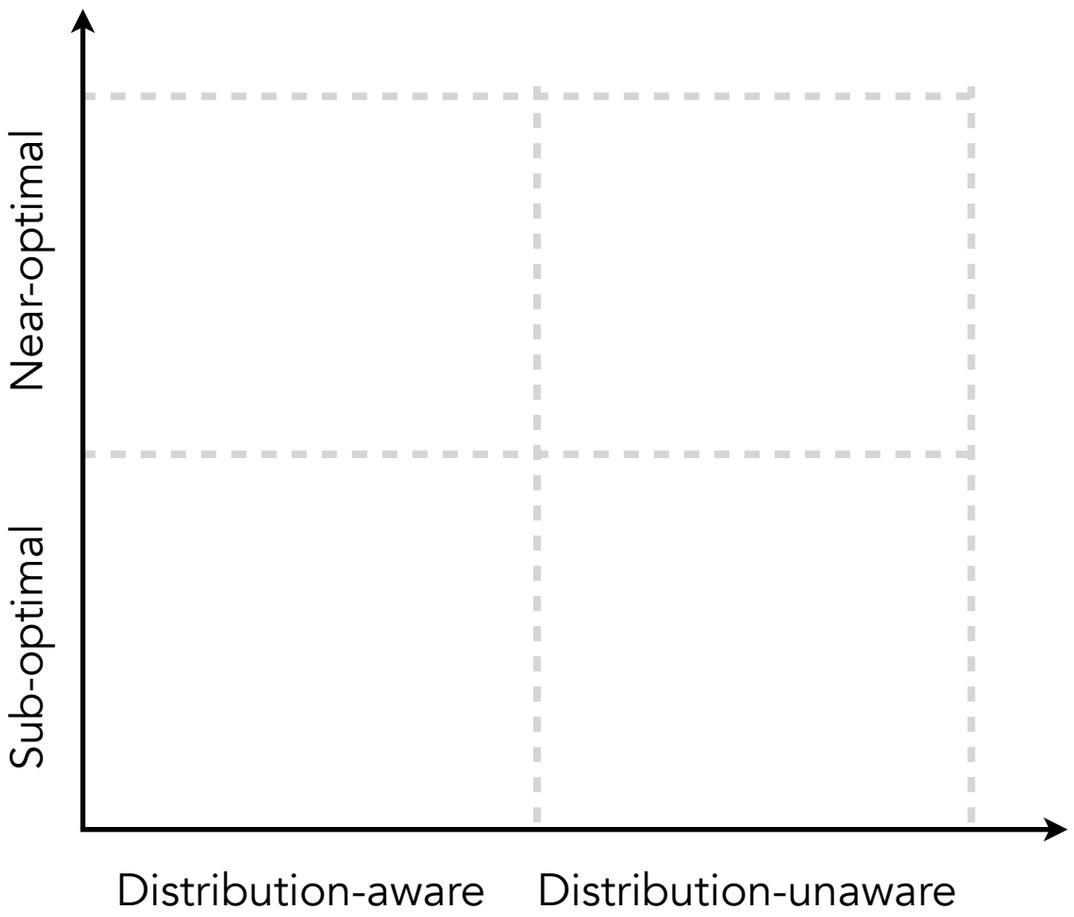
$$\text{Reg}_n(\text{ROAR}) = U_\infty^{\text{off}} - U_n^{\text{ROAR}} = \frac{1}{n}\sum_{k=1}^{n}\left(U_\infty^{\text{off}} - U_{\min\{k,n-k+1\}}^{\text{off}}\right) = \frac{2}{n}\sum_{k=1}^{n/2}\left(U_\infty^{\text{off}} - U_k^{\text{off}}\right)$$

Well studied object in empirical optimal transport literature

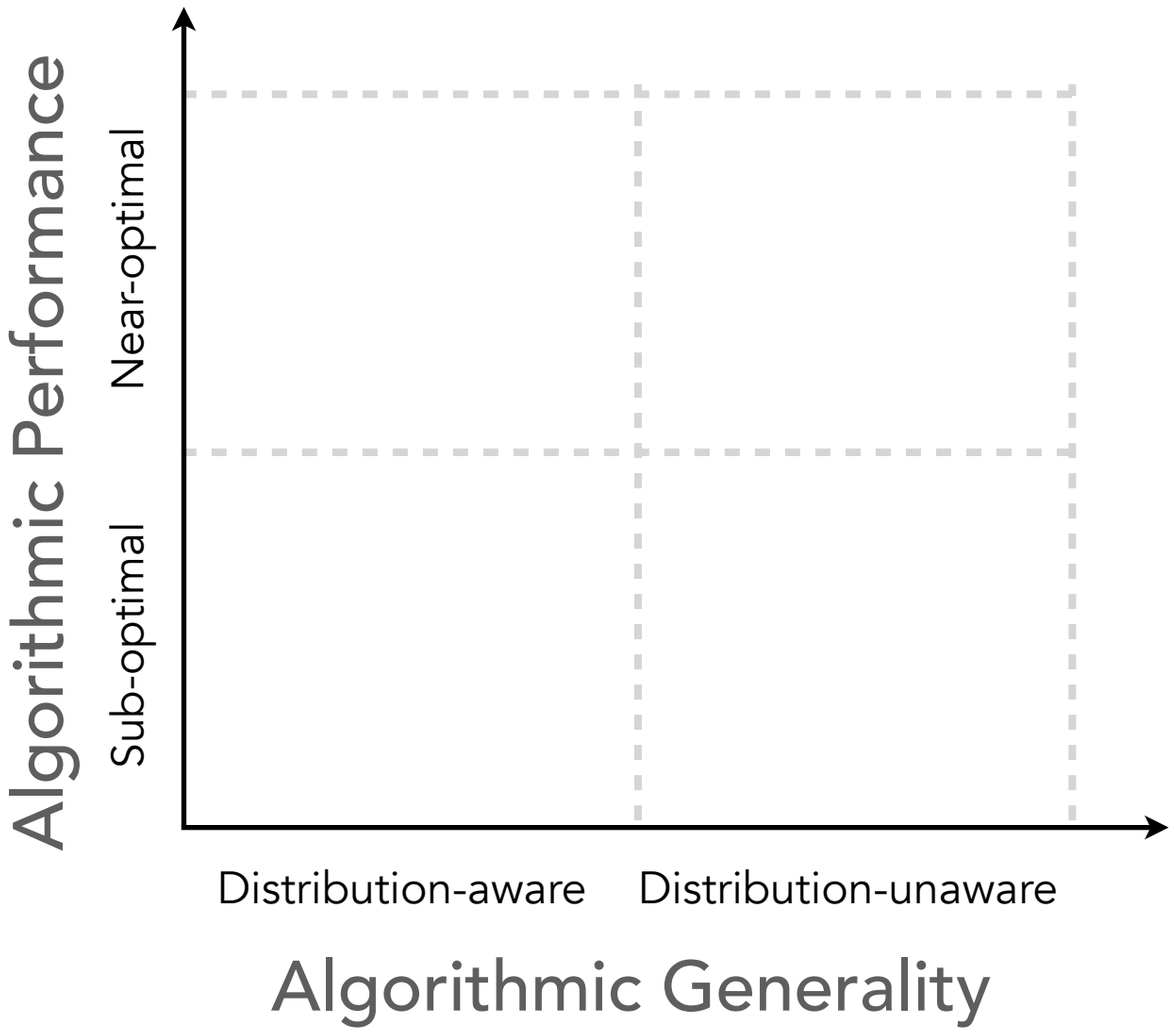
Let the match quality be measured as $\varphi(X, Y) = -\|X - Y\|^p$

	Demand and Supply distributions are smooth	Demand and Supply distributions are arbitrary
ROAR	$\tilde{\mathcal{O}}(n^{-(p/d\wedge 1)})$	$\tilde{\mathcal{O}}(n^{-(p/d\wedge 1/2)})$
Fundamental Limit	$\tilde{\Omega}(n^{-(p/d\wedge 1)})$	$\tilde{\Omega}(n^{-(p/d\wedge 1/2)})$

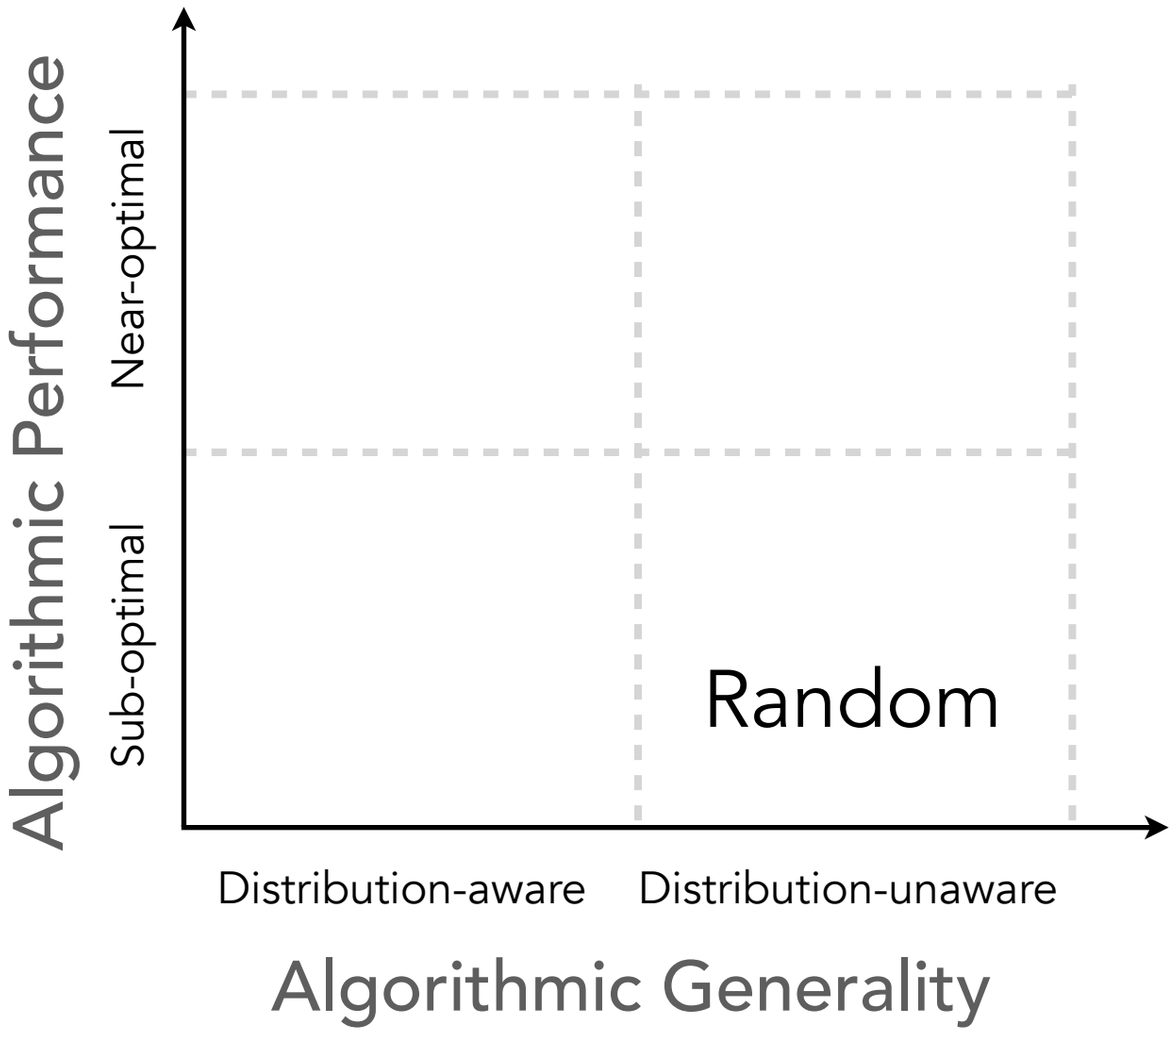
Take-Home Message



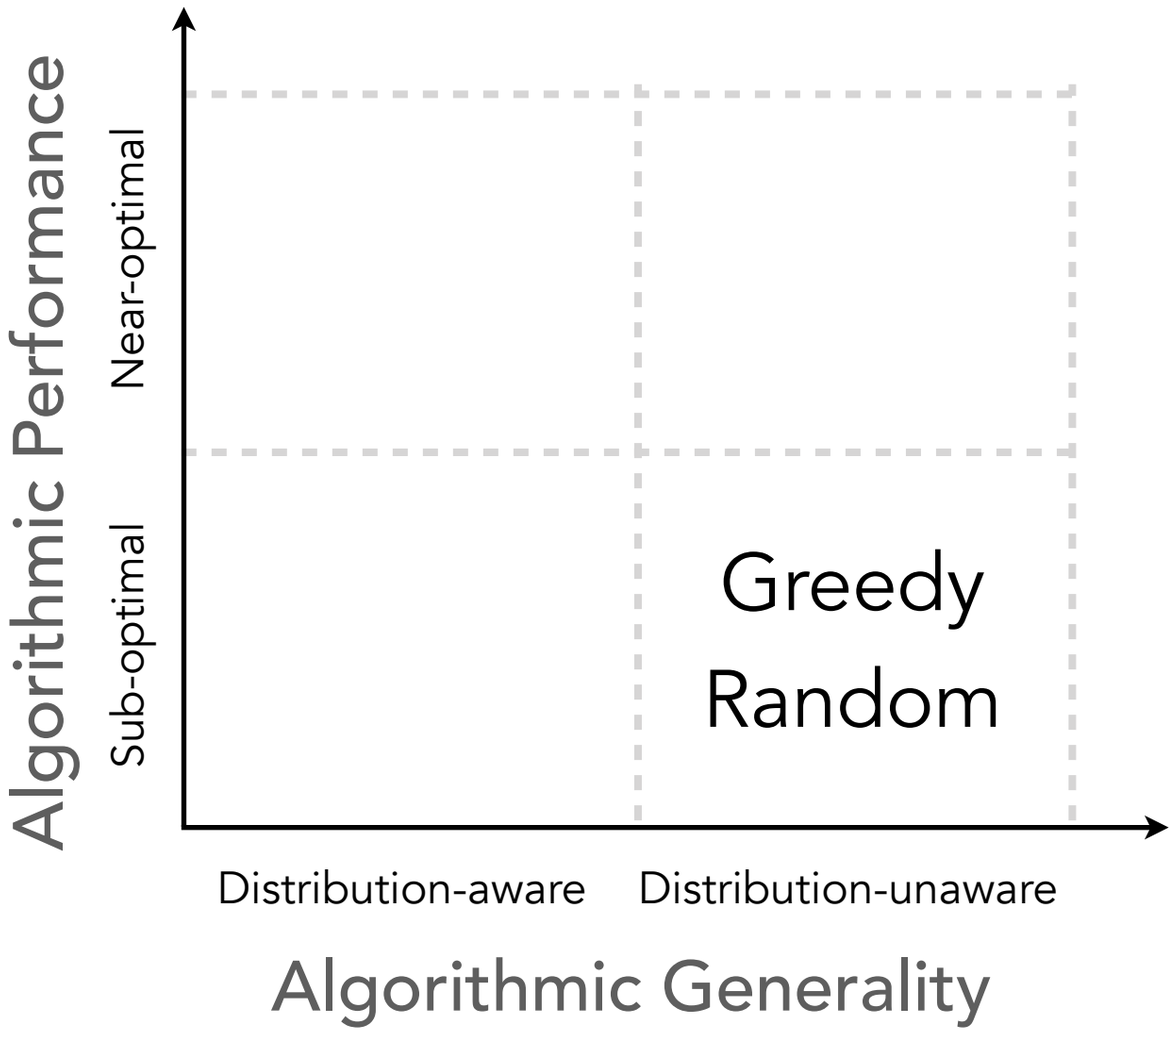
Take-Home Message



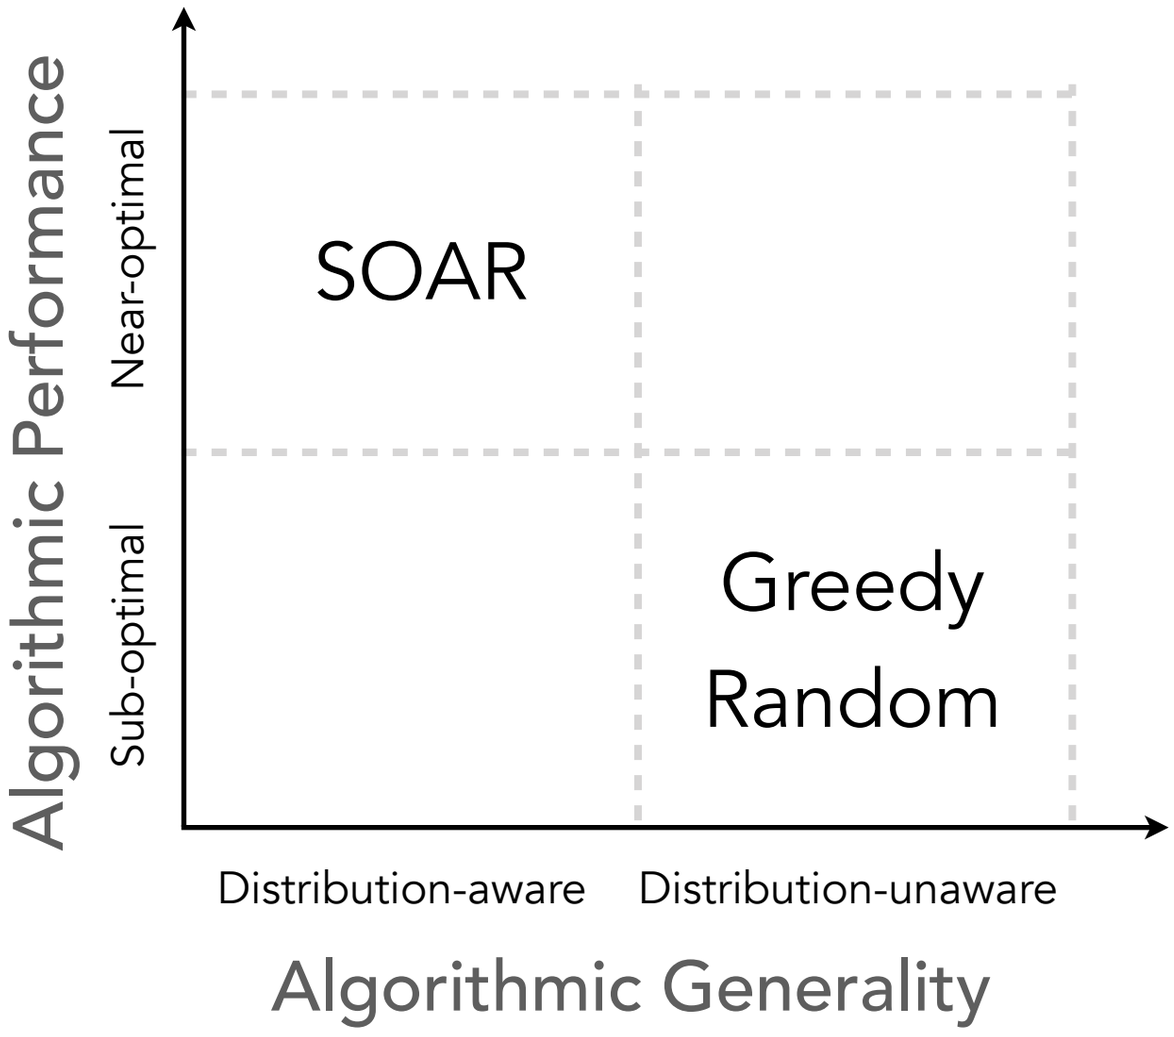
Take-Home Message



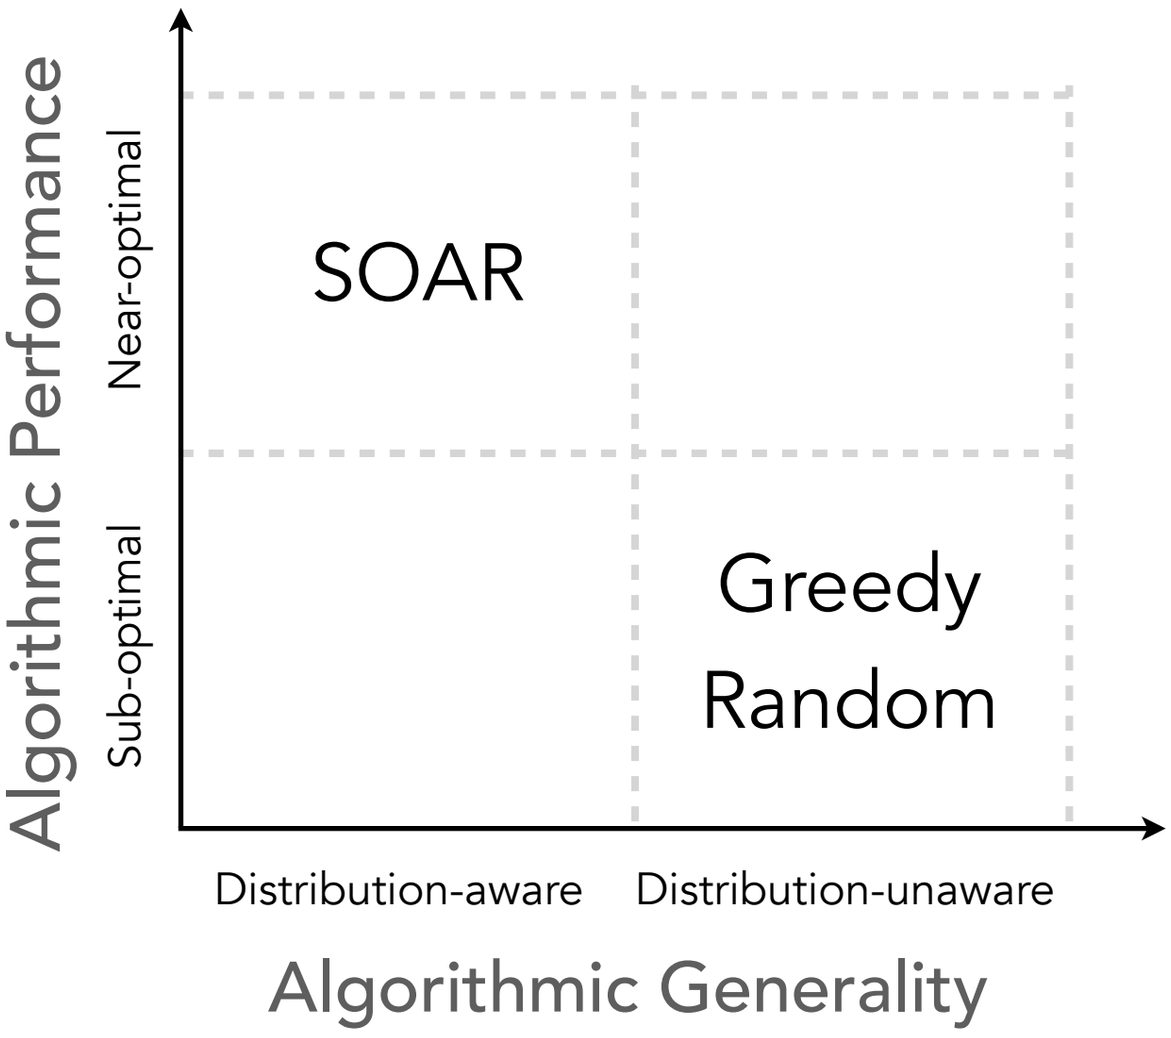
Take-Home Message

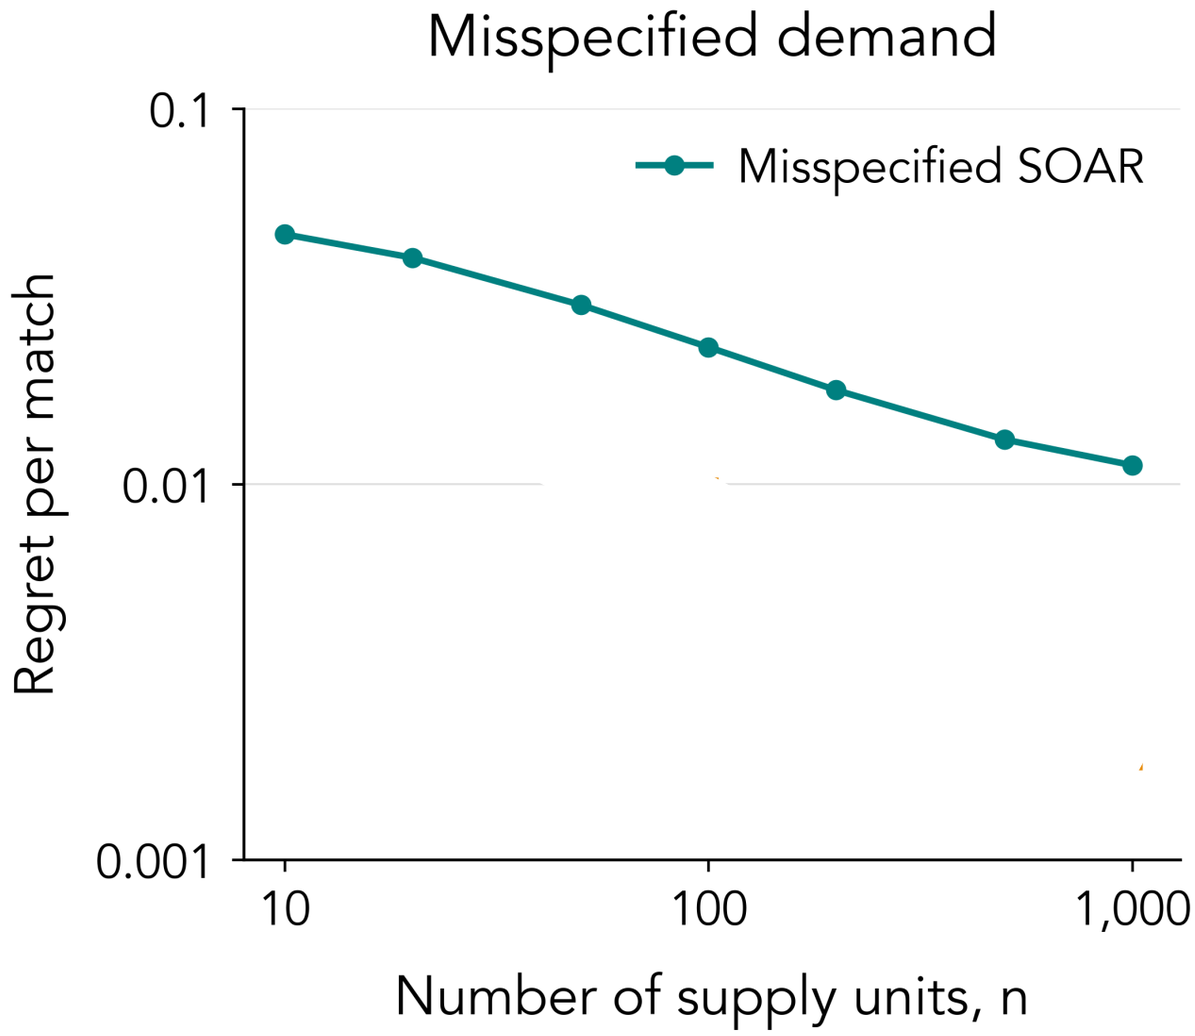


Take-Home Message

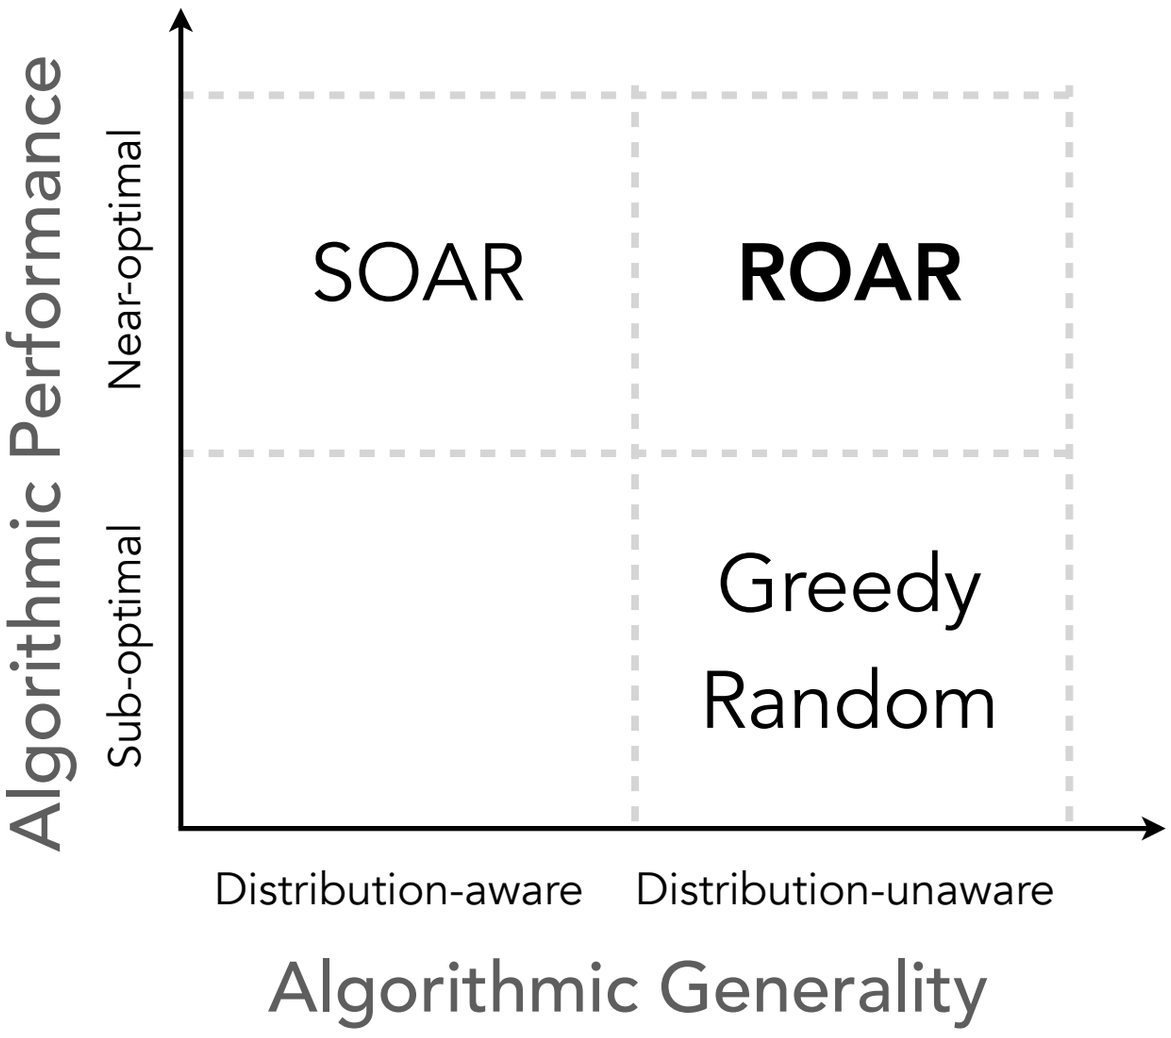


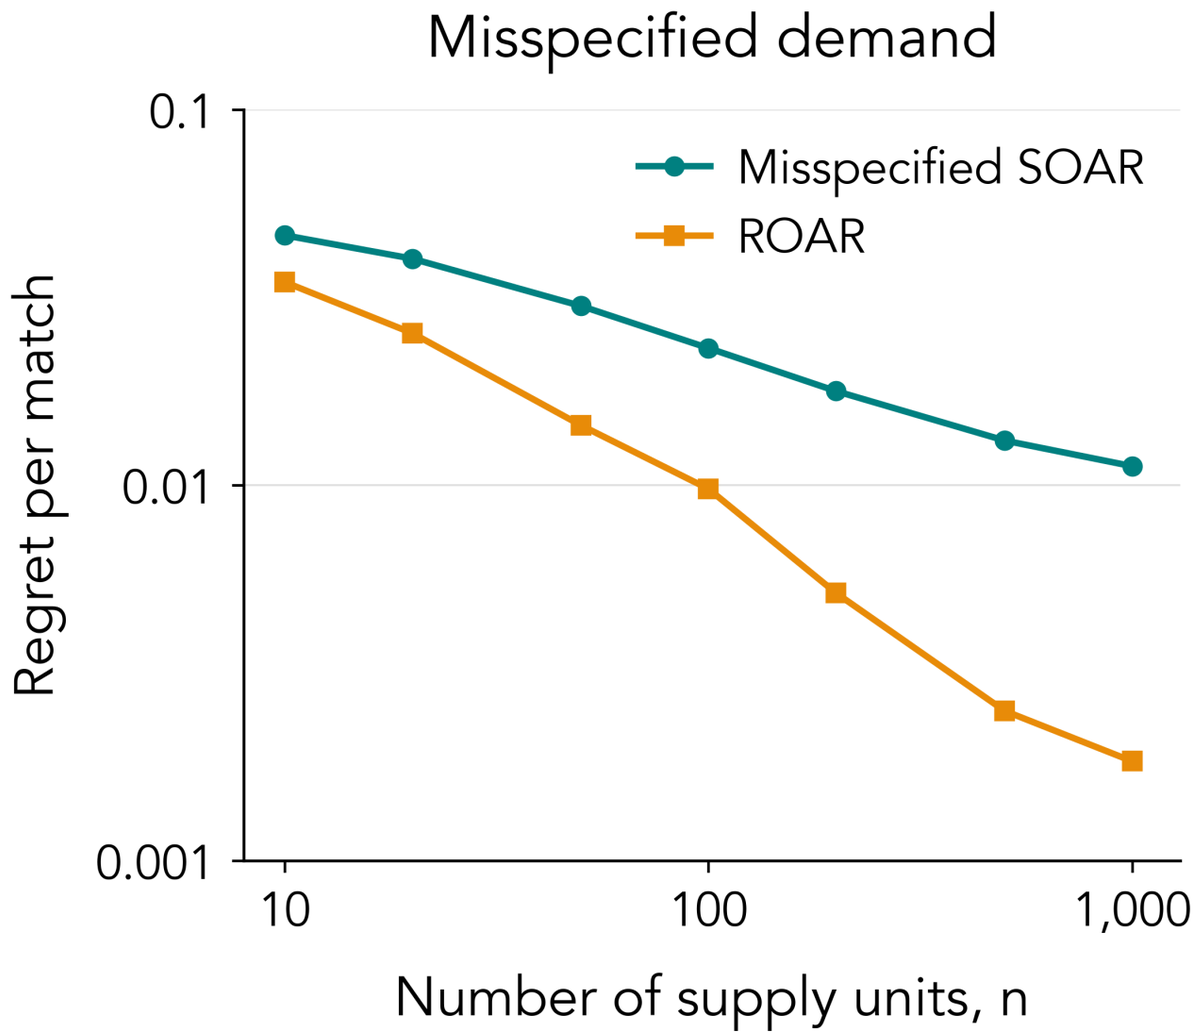
Take-Home Message



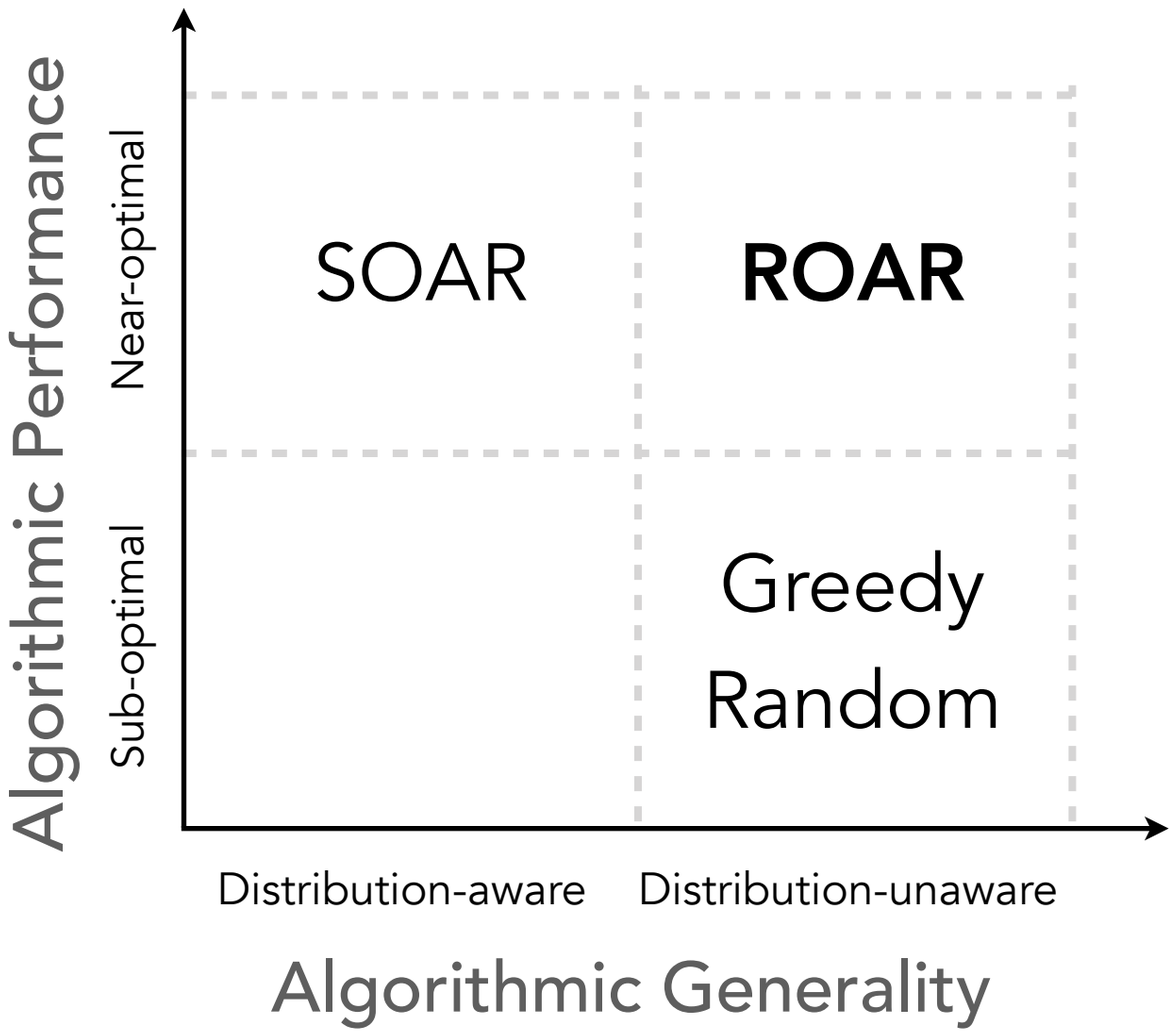


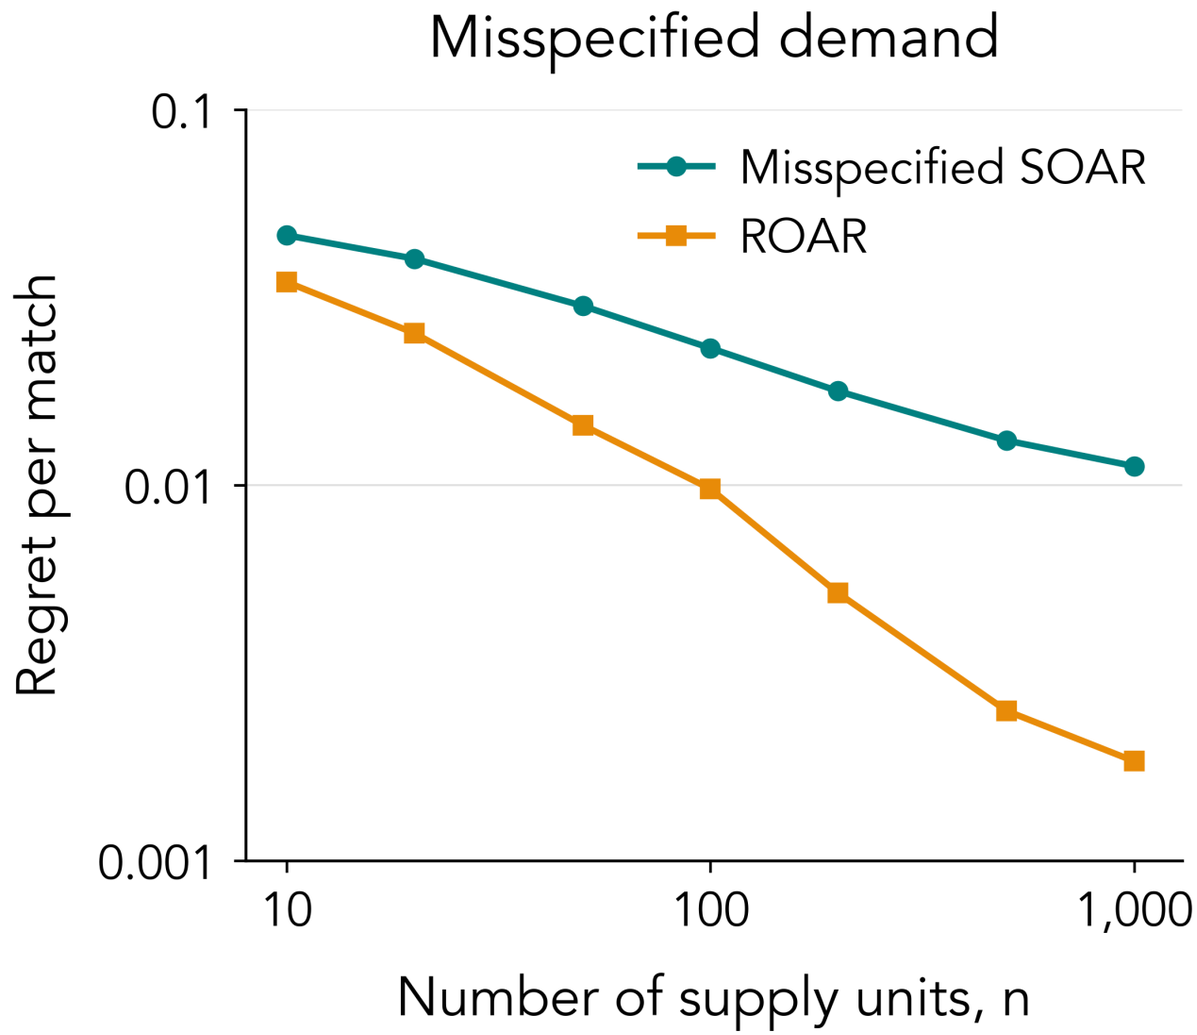
Take-Home Message





Thank You


Algorithmic Performance
Near-optimal
Sub-optimal
SOAR
ROAR
Greedy
Random
Distribution-aware
Distribution-unaware
Algorithmic Generality


Misspecified demand
Misspecified SOAR
ROAR
Regret per match
0.1
0.01
0.001
10
100
1,000
Number of supply units, n