Into the Unknown: Assigning Reviewers to Papers with Uncertain Affinities

Cyrus Cousins

Justin Payan

Yair Zick

UMassAmherst | College of Information & Computer Sciences

Peer Review

Papers must be reviewed by suitable reviewers

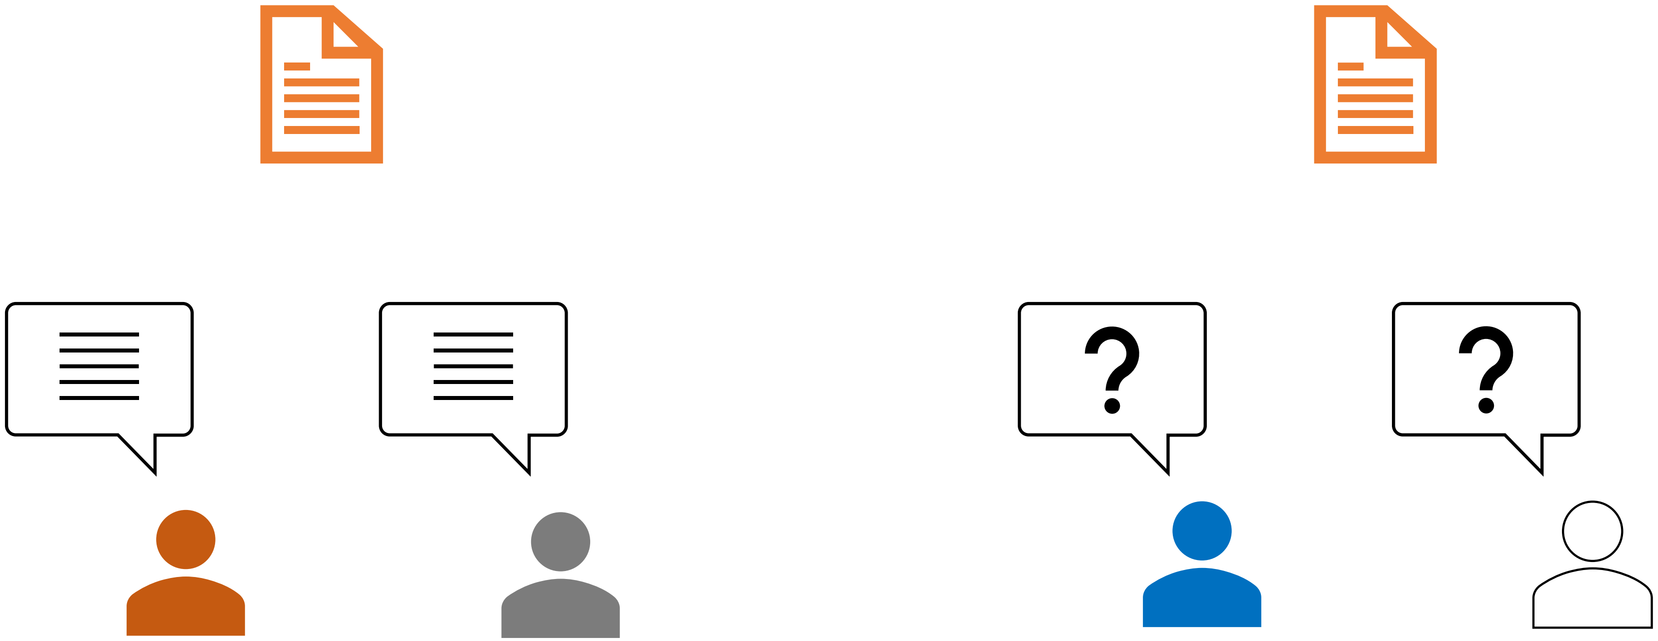

Reviewer Assignment Problem

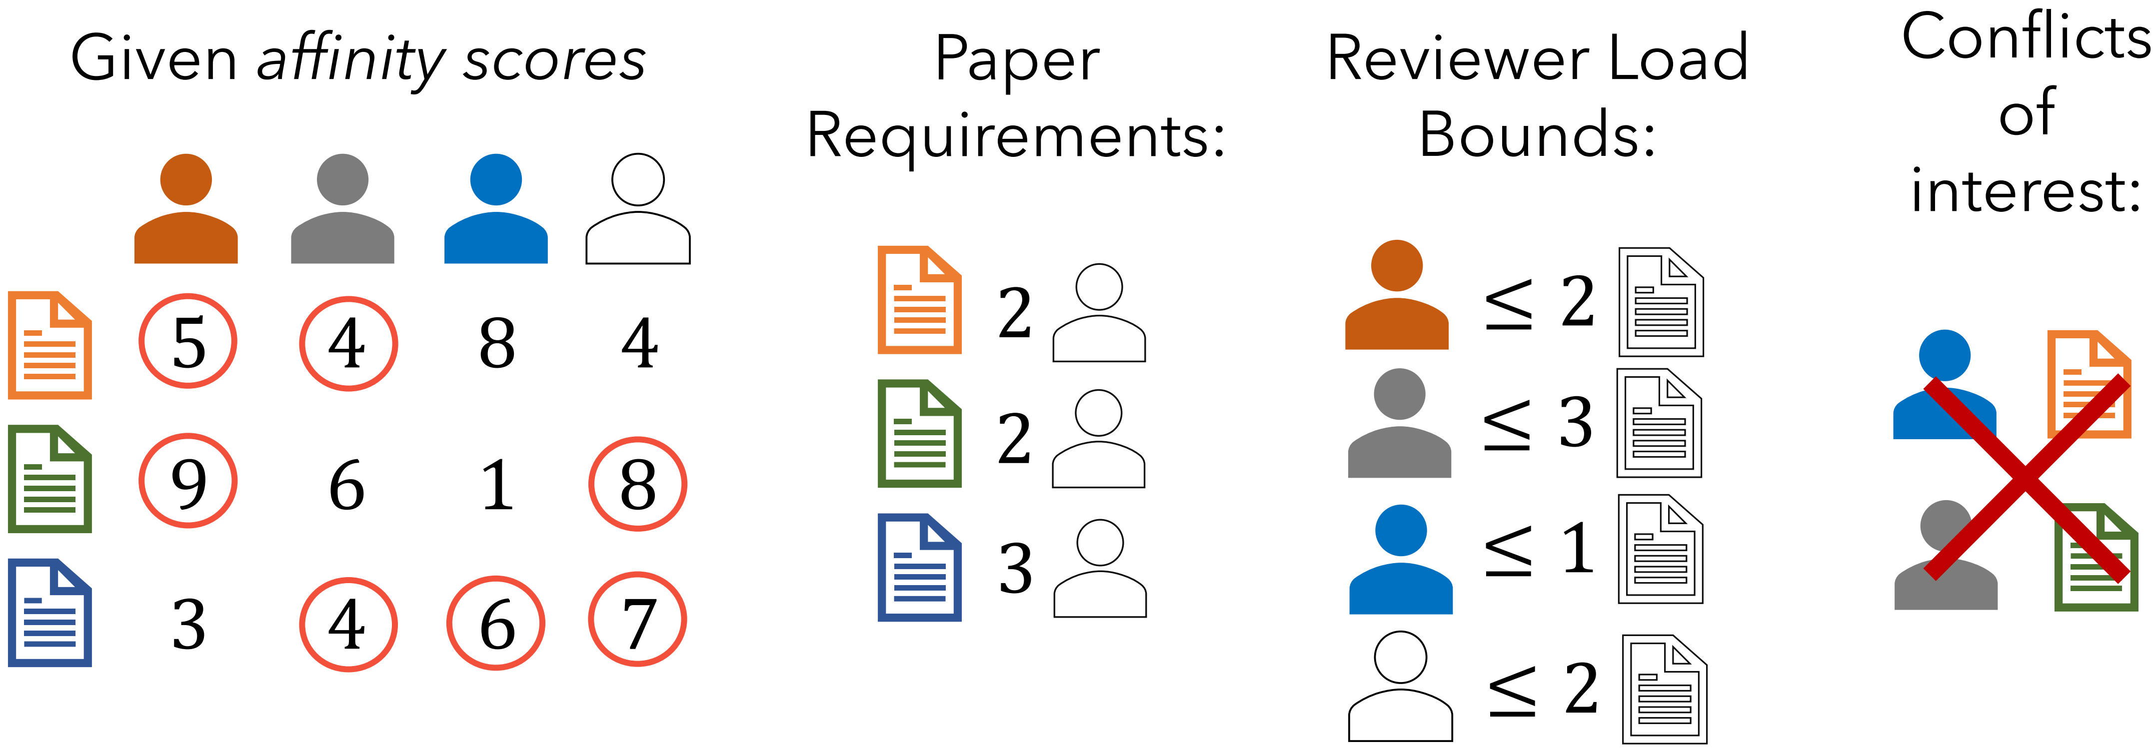

Goal:
Allocation of reviewers to papers with high affinity

What are affinity scores?

For given *reviewer-paper pair*
the affinity score measures:

Reviewer interest

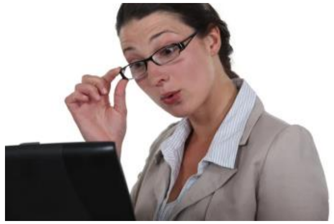

Reviewer expertise

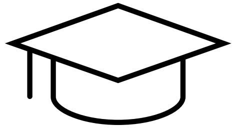

As proxies for:

Predicted review quality

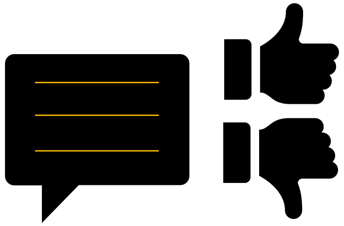

What are affinity scores?

For given *reviewer-paper pair* → the affinity score measures: 5

Reviewer interest

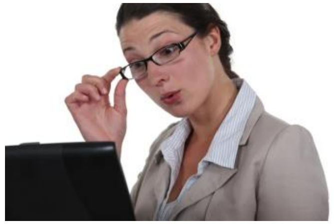

Reviewer expertise

As proxies for:

Predicted review quality

Hard to measure/predict

Reviewer Assignment Under Uncertainty (RAU)

Compute an *uncertainty set* $\mathcal{S}$

s.t. true, unknown affinity scores S are contained in $\mathcal{S}$ with prob. $1 - \delta$

5	4	8	4
9	6	1	8
3	4	6	7

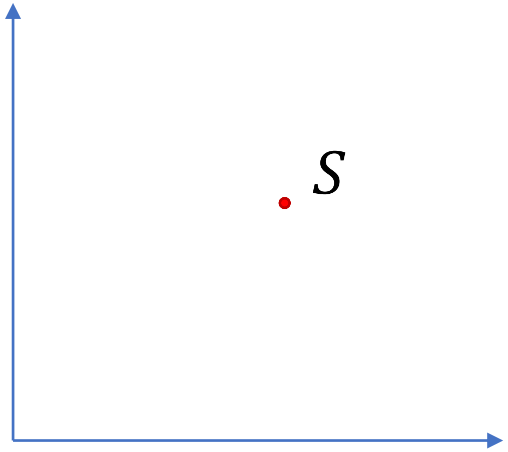

Reviewer Assignment Under Uncertainty (RAU)

Compute an *uncertainty set* $\mathcal{S}$
s.t. true, unknown affinity scores S are contained in $\mathcal{S}$ with prob. $1 - \delta$

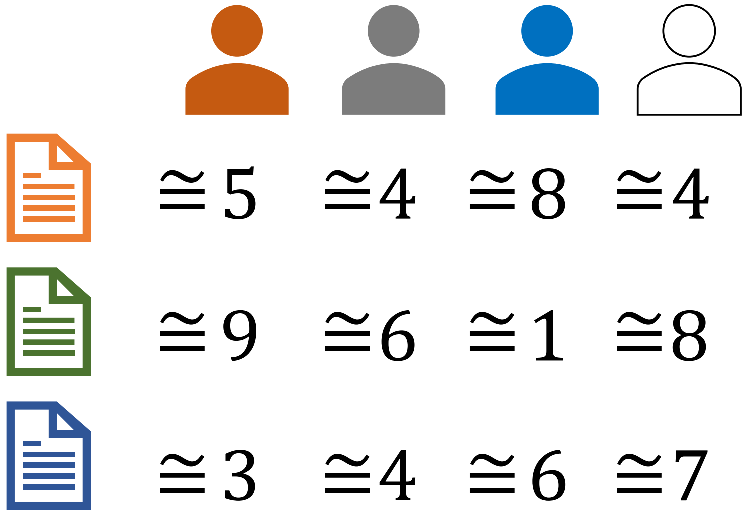

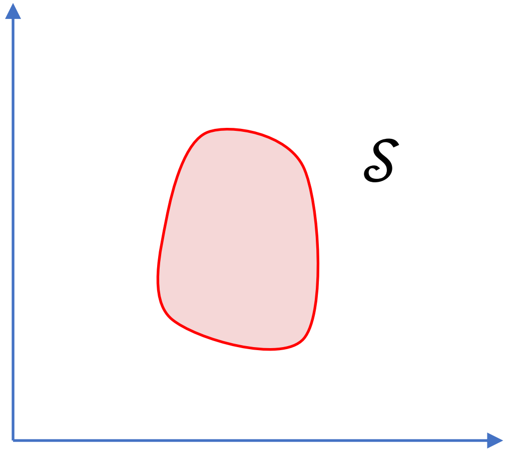

Reviewer Assignment Under Uncertainty (RAU)

Compute an *uncertainty set* $\mathcal{S}$
s.t. true, unknown affinity scores S are contained in $\mathcal{S}$ with prob. $1-\delta$

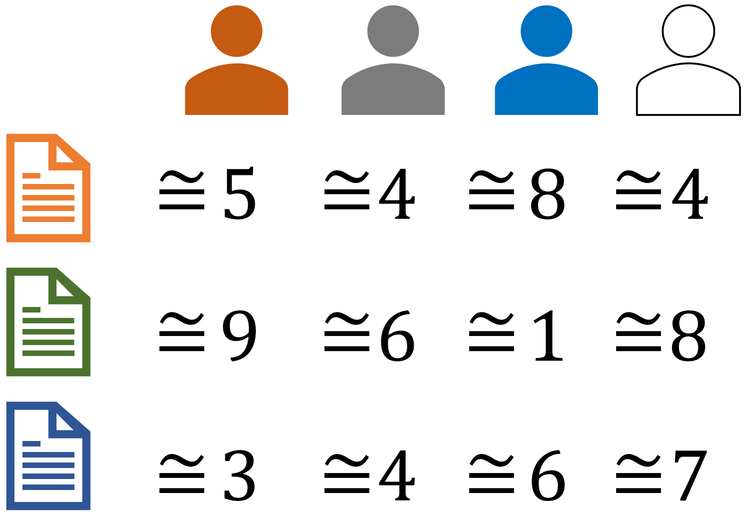

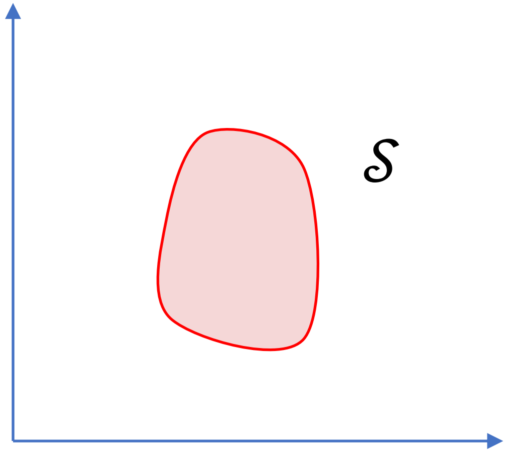

Maximize the worst-case welfare over the *uncertainty set*

Reviewer Assignment Under Uncertainty (RAU)

Papers P and reviewers R

Valid assignments $\mathcal{A} \subseteq \{0,1\}^{|P|\times|R|}$

Affinity score uncertainty set $\mathcal{S} \subseteq \mathbb{R}_+^{|P|\times|R|}$

$$USW(A,S) = \frac{1}{|P|}\sum_{p\in P}\sum_{r\in R} A_{pr}S_{pr}$$

Objective: $\max_{A\in\mathcal{A}} \min_{S\in\mathcal{S}} USW(A,S)$

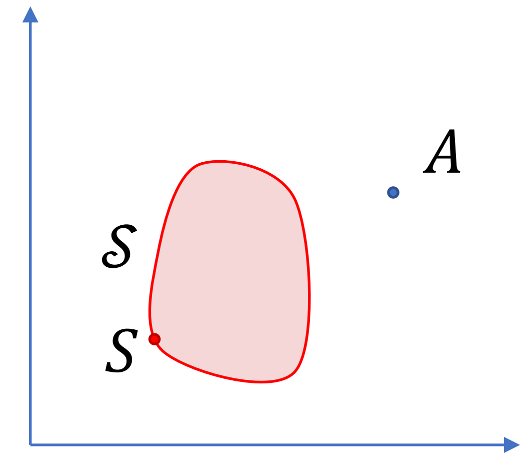

Reviewer Assignment Under Uncertainty (RAU)

Papers P and reviewers R
Valid assignments $\mathcal{A} \subseteq \{0,1\}^{|P|\times|R|}$
Affinity score uncertainty set $\mathcal{S} \subseteq \mathbb{R}_+^{|P|\times|R|}$

$$USW(A,S) = \frac{1}{|P|}\sum_{p\in P}\sum_{r\in R} A_{pr}S_{pr}$$

Objective: $\max_{A\in\mathcal{A}} \min_{S\in\mathcal{S}} USW(A,S)$

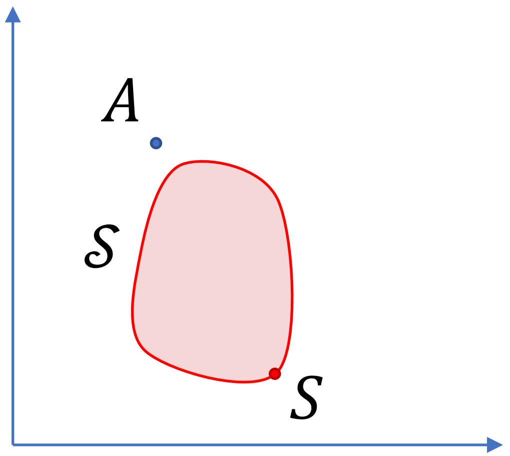

Why Maximin?

Why Maximin?

1. May not always know expected affinities

2. Theory and experiments show maximin corresponds to *true* welfare

Implementing RAU

Implementing RAU: Bid Prediction

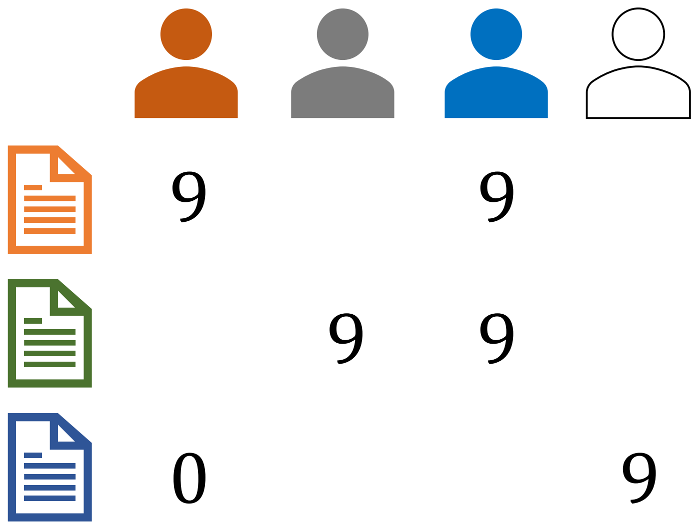

Solicit bids

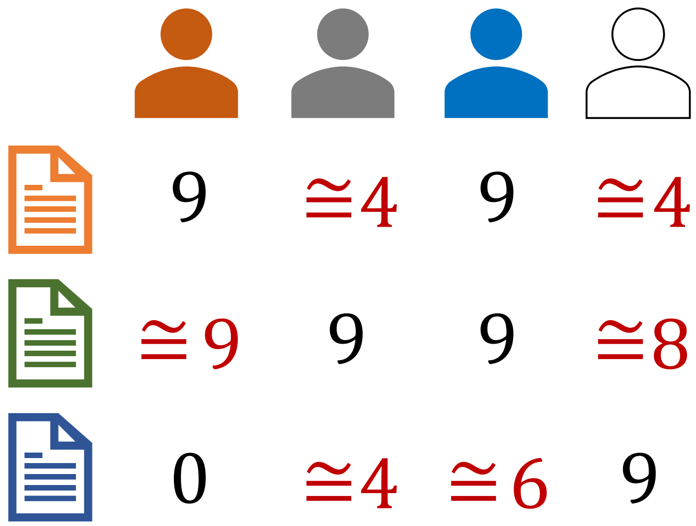

Collaborative Filtering with
Average Error Guarantees

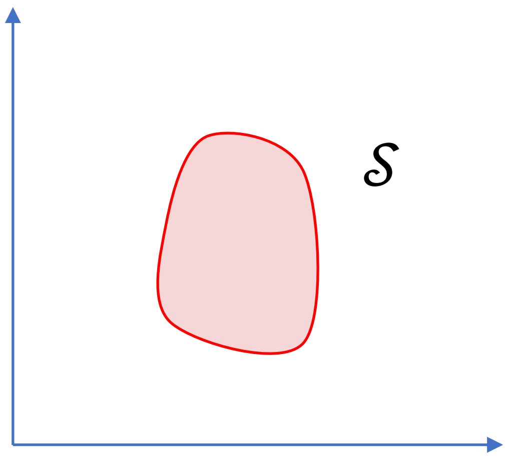

Implementing RAU: Review Quality Prediction

Label past assignments

Learn function from reviewer and paper features to affinities

Apply f with error bounds

BUILD-A Uncertainty Set WORKSHOP®

WHERE BEST FRIENDS ARE MADE®

Transductive Predictors (CF on Bids)

$\mathcal{S}$

$$\mathcal{S} = \{S: \left(S - \hat{S}\right)^T \Sigma^{-1} (S - \hat{S})/nm \leq \hat{\xi} + \eta_t\}$$

Inductive Predictors (Learned f from Data)

$\mathcal{S}$

$$\mathcal{S} = \{S: \left(S - \hat{S}\right)^T \Sigma^{-1} (S - \hat{S})/nm \leq \hat{\xi} + \eta_i\}$$

Transductive Predictors
(CF on Bids)

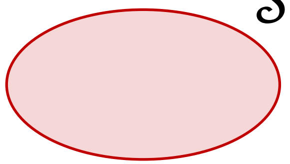

$$\mathcal{S} = \{S: (S - \hat{S})^T \Sigma^{-1} (S - \hat{S})/nm \le \hat{\xi} + \eta_t\}$$

Inductive Predictors
(Learned f from Data)

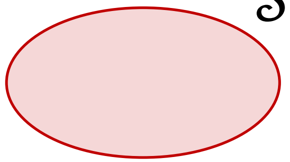

$$\mathcal{S} = \{S: (S - \hat{S})^T \Sigma^{-1} (S - \hat{S})/nm \le \hat{\xi} + \eta_i\}$$

Intersecting Uncertainty Sets

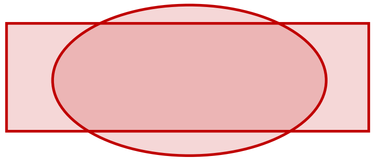

$$p(S \in \mathcal{S}) = 1 - (\delta_1 + \delta_2)$$

Expanding Uncertainty Sets

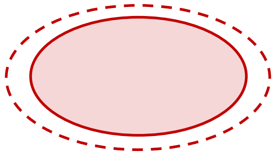

Solving RAU

RAU is NP-hard

Theorem 1:
RAU is NP-hard for convex uncertainty sets $\mathcal{S}$

RAU is NP-hard

Theorem 1:
RAU is NP-hard for convex uncertainty sets $\mathcal{S}$

Nonlinear objective

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{S}} USW(A, S)$$

Robust Reviewer Assignment (RRA)

Relax discrete allocations → continuous

$\mathcal{A} \subseteq \{0,1\}^{m \times n} \rightarrow \tilde{\mathcal{A}} \subseteq [0,1]^{m \times n}$

Projected subgradient-ascent optimization

Solve $\max_{A \in \tilde{\mathcal{A}}} \min_{S \in \mathcal{S}} USW(A,S)$
by stepping in $\partial_A \min_{S \in \mathcal{S}} USW(A,S)$
and projecting back to $\tilde{\mathcal{A}}$

Randomized rounding for discrete solution

Round $A \in \tilde{\mathcal{A}}$ to $A' \in \mathcal{A}$

RRA Guarantees

Define

$\tilde{A}$: the continuous maximin solution
A': the rounding of $\tilde{A}$
S^*: the true affinity score matrix
A^*: the true optimal allocation
L: $\mathcal{L}_1$-diameter of $\mathcal{S}$
δ: $\mathbb{P}(S^* \in \mathcal{S}) \geq 1 - \delta$

RRA Guarantees

Define

$\tilde{A}$: the continuous maximin solution
A': the rounding of $\tilde{A}$
S^*: the true affinity score matrix
A^*: the true optimal allocation
L: $\mathcal{L}_1$-diameter of $\mathcal{S}$
δ: $\mathbb{P}(S^* \in \mathcal{S}) \geq 1 - \delta$

Integrality Gap

$$\mathbb{E}_{A'} \left\| A' - \tilde{A} \right\|_1 = 2\left(\|A'\|_1 - \left\|\tilde{A}\right\|_1 \right)$$
$$= |P||R| - 2\left\| \frac{1}{2} - \tilde{A} \right\|_1$$

RRA Guarantees

Define

$\tilde{A}$: the continuous maximin solution
A': the rounding of $\tilde{A}$
S^*: the true affinity score matrix
A^*: the true optimal allocation
L: $\mathcal{L}_1$-diameter of $\mathcal{S}$
δ: $\mathbb{P}(S^* \in \mathcal{S}) \geq 1 - \delta$

Integrality Gap

$$\mathbb{E}_{A'}\left\|A' - \tilde{A}\right\|_1 = 2\left(\|A'\|_1 - \left\|\tilde{A}\right\|_1\right)$$
$$= |P||R| - 2\left\|\frac{1}{2} - \tilde{A}\right\|_1$$

True Welfare Gap

$$\mathbb{P}\left(USW(A^*, S^*) - \mathbb{E}_{A'}USW(A', S^*) > \frac{L}{|P|}\right) < \delta$$

RRA Guarantees

Integrality Gap

We may have to round heavily…

But the maximin objective still ensures high *true* welfare!

$< \delta$

ICLR Experiments

Year	# Revs	# Papers	100 * Adversarial USW: ILP	100 * Adversarial USW: RRA	100 * Average USW: ILP	100 * Average USW: RRA
2018	1657	546	**17**	16	**179**	160
2019	2620	851	22	**27**	**184**	161
2020	4123	1327	17	**23**	**187**	166
2021	4662	1557	23	**33**	**192**	174
2022	5023	1576	28	**38**	**191**	172
Average			21.4	27.4 (+28%)	187	167 (−11%)

USW numbers are divided by # Papers

ICLR Experiments

			100 * Adversarial USW		100 * Average USW	
Year	**# Revs**	**# Papers**	**ILP**	**RRA**	**ILP**	**RRA**
2018	1657	546	**17**	16	**179**	160
2019	2620	851	22	**27**	**184**	161
2020	4123	1327	17	**23**	**187**	166
2021	4662	1557	23	**33**	**192**	174
2022	5023	1576	28	**38**	**191**	172
Average			21.4	27.4 (+28%)	187	167 (−11%)

USW numbers are divided by # Papers

True Welfare under Noise

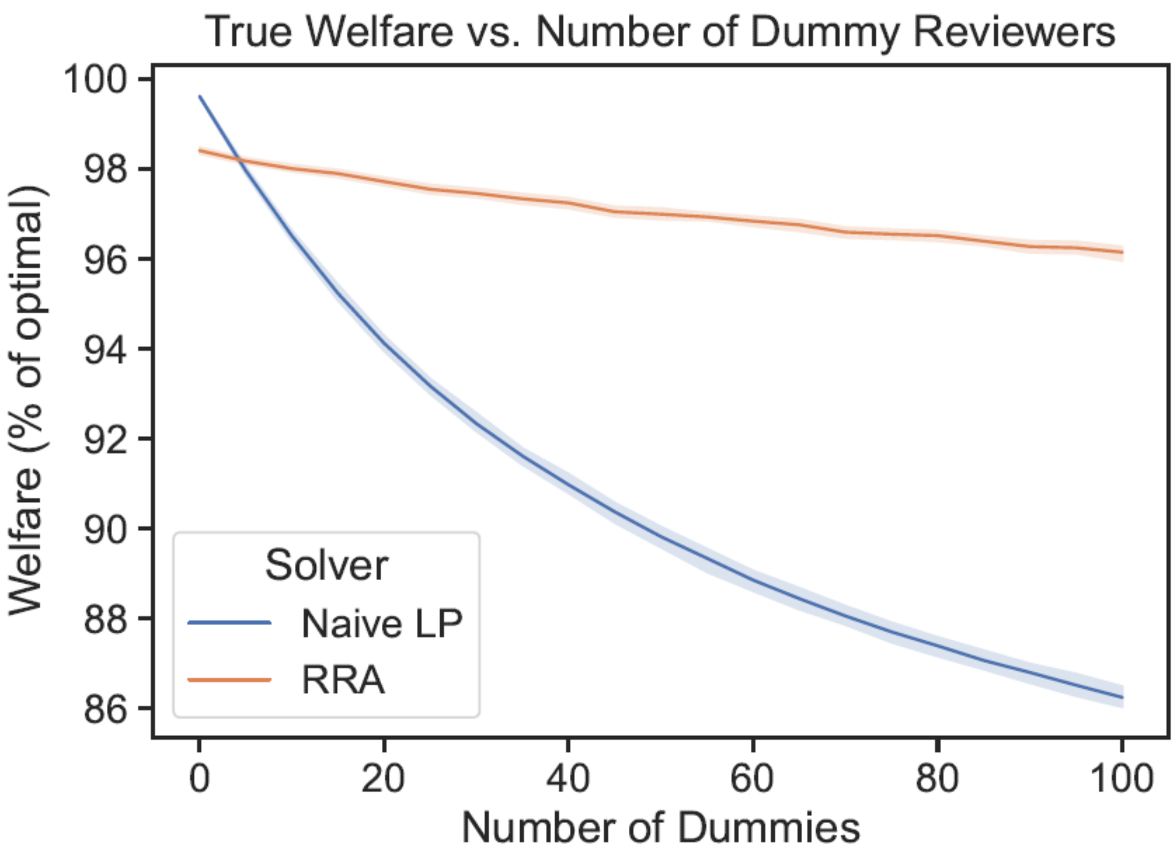

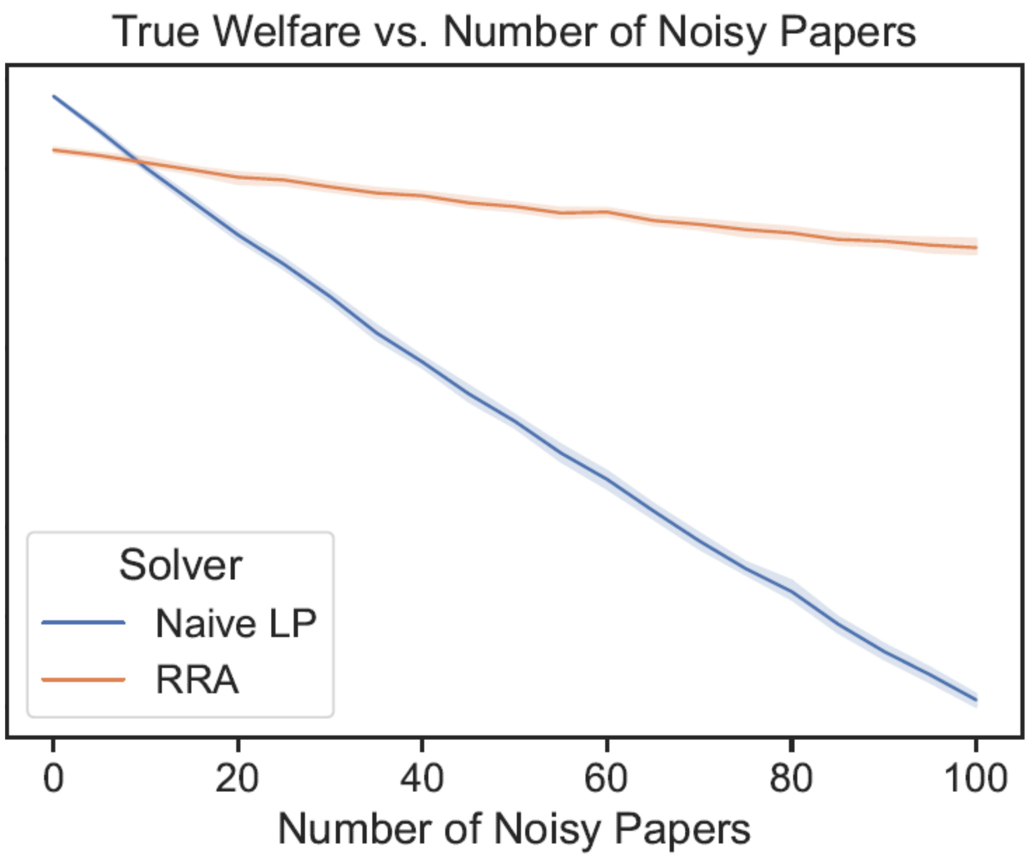

https://t.ly/g_sLd

Thanks!

Email:
jpayan@cs.umass.edu

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an "adversary"

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an "adversary"

↓

Gradient ascent for
non-differentiable functions

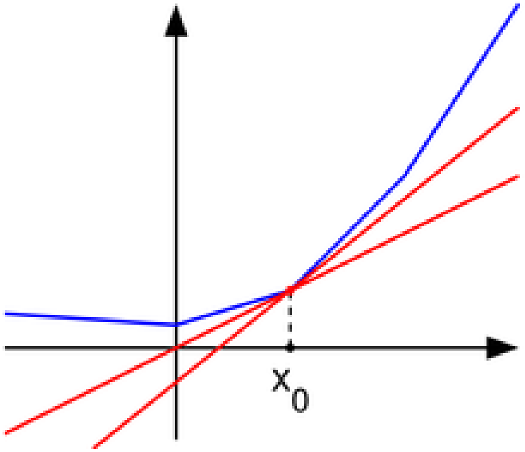

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an “adversary”

At each step, ensure constraints are satisfied

Gradient ascent for non-differentiable functions

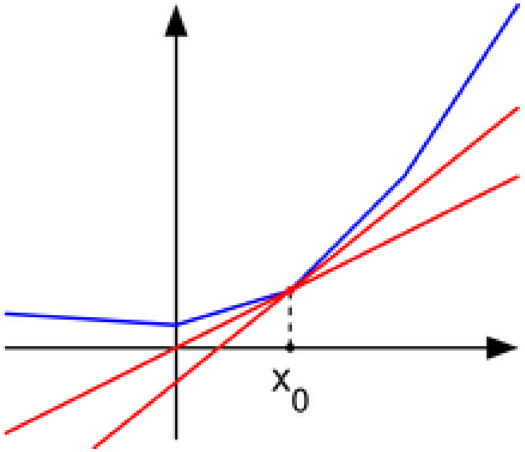

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an "adversary"

At each step, ensure constraints are satisfied

Gradient ascent for non-differentiable functions

It's a maximin problem

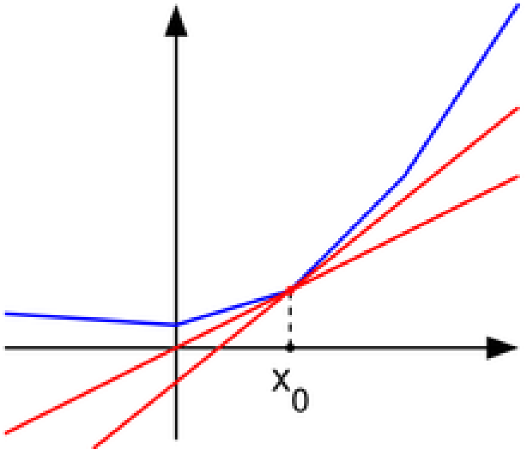

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an "adversary"

Converges in poly-time if projection and adversary are efficient

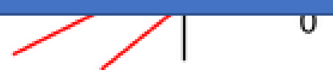

Robust Reviewer Assignment (RRA)

Randomized rounding for discrete solution

Round $A \in \tilde{\mathcal{A}}$ to $A' \in \mathcal{A}$

Robust Reviewer Assignment (RRA)

Randomized rounding for discrete solution

Round $A \in \tilde{\mathcal{A}}$ to $A' \in \mathcal{A}$

Generalization of Birkhoff von Neumann decomposition

Maintains constraints of $\tilde{\mathcal{A}}$ and $\mathcal{A}$

$$\mathbb{E}[A'] = A$$

$$\min_{S \in \mathcal{S}} USW(A', S) \leq \min_{S \in \mathcal{S}} USW(A, S)$$

Budish et al. Implementing random assignments: A generalization of the birkhoff-von neumann theorem. 2009.
Jecmen et al. Mitigating manipulation in peer review via randomized reviewer assignments. 2020.

ICLR Experiments

Robust Reviewer Assignment (RRA)

What about general convex uncertainty sets $\mathcal{S}$?

Relax, then apply projected subgradient ascent against an "adversary"

Simple Uncertainty Sets

Independent
Bounding Boxes

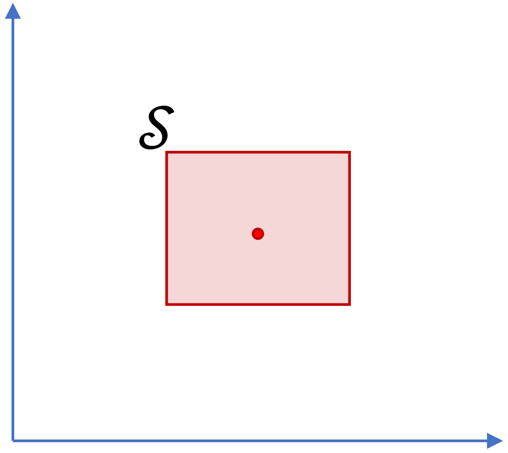

Simple Uncertainty Sets

Independent
Bounding Boxes

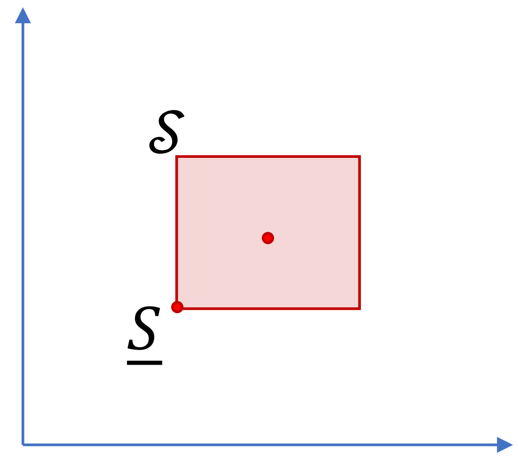

Optimize
assuming $\underline{S}$

Simple Uncertainty Sets

Independent Bounding Boxes

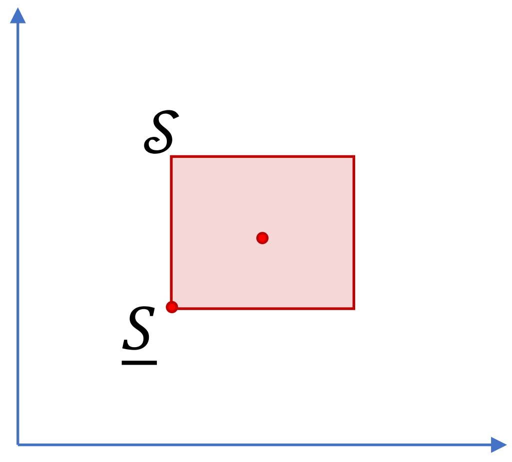

Optimize assuming $\underline{S}$

Sphere

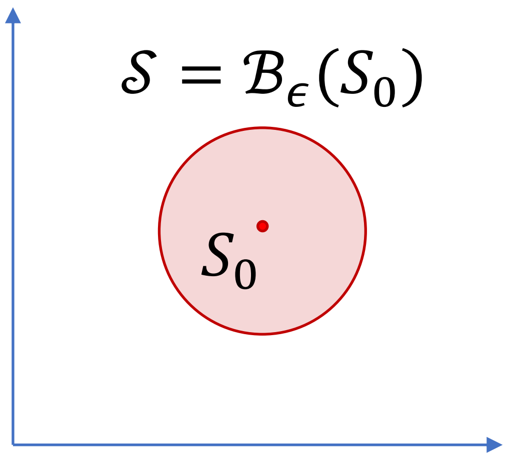

Simple Uncertainty Sets

Independent
Bounding Boxes

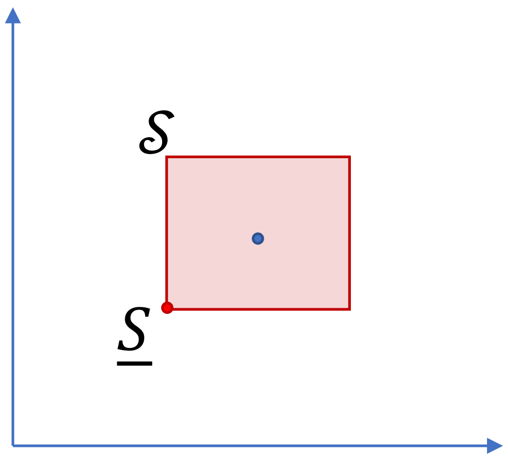

Optimize
assuming $\underline{S}$

Sphere

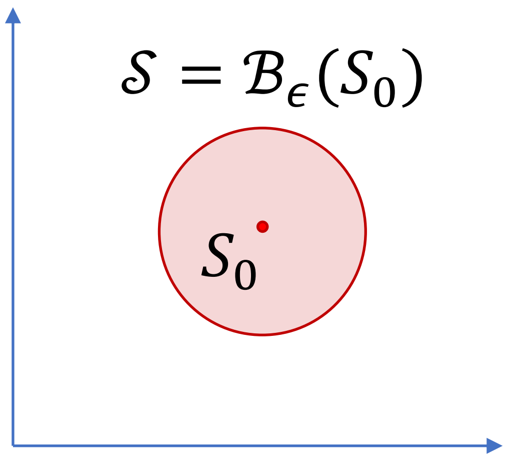

Optimize
assuming S_0

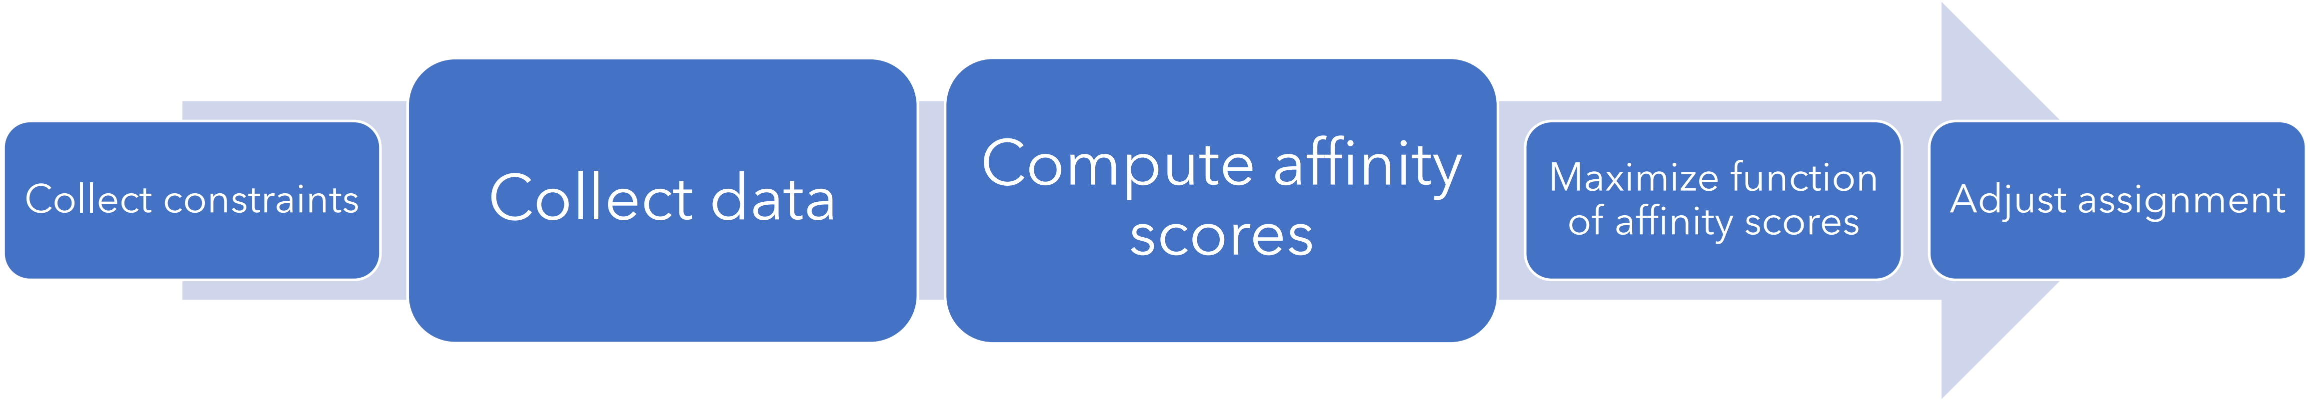
Collect constraints
Collect data
Compute affinity scores
Maximize function of affinity scores
Adjust assignment

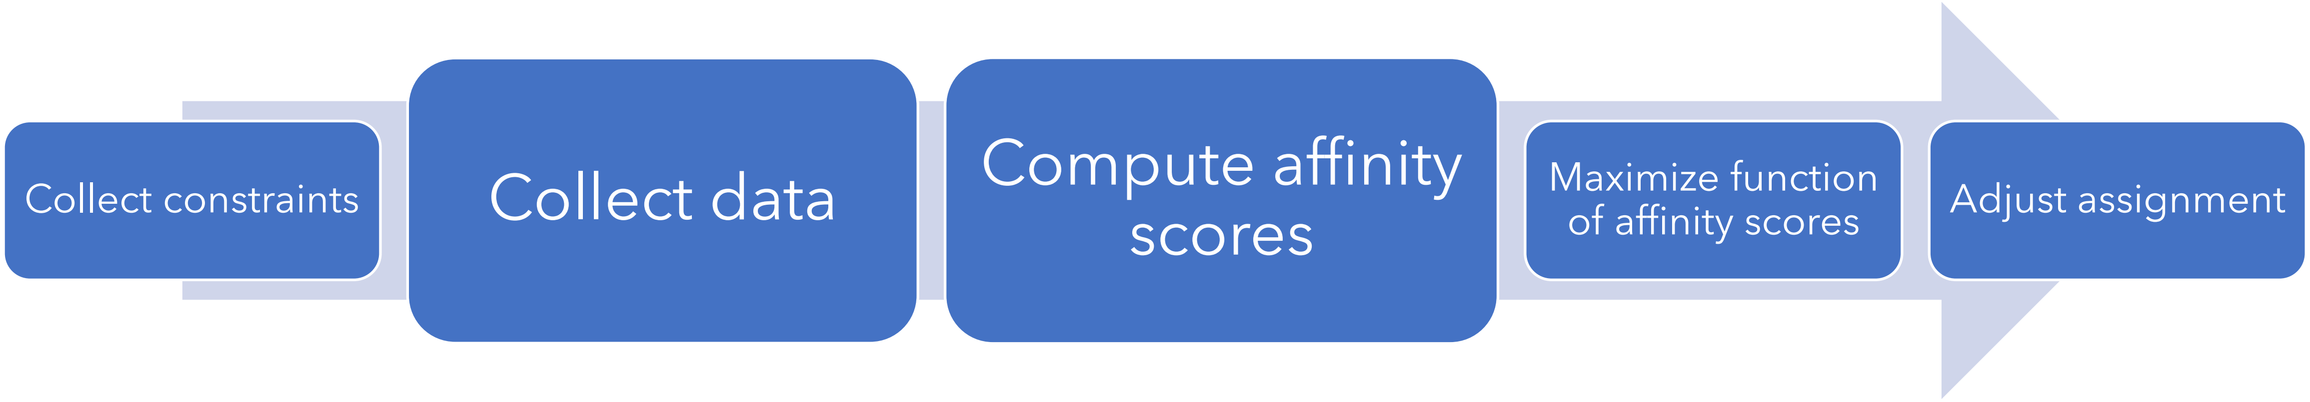

$$
\text{Base Aggregated Score} = \begin{cases} \frac{1}{4}\text{TPMS} + \frac{1}{4}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if all data is present} \\ \frac{1}{2}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if TPMS is missing} \\ \frac{1}{2}\text{TPMS} + \frac{1}{2}\text{SAM} & \text{if ACL is missing} \\ \text{SAM} & \text{if both ACL and TPMS are missing} \end{cases}
$$

Compute $\mathcal{S}$ as a τ% CI of $\mathcal{N}(\mu, \Sigma)$

μ is the base scores

Σ is diagonal, based on number of missing sources

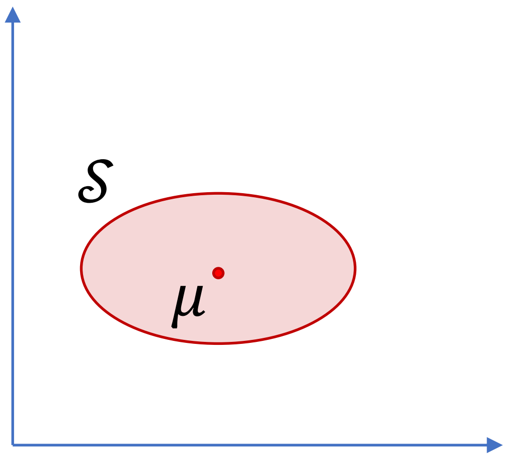

Extensions

(and Related Problems)

Learning Bids

Reviewer bid | Document similarity | Keyword matching

λ_1 1 + λ_2 7 + λ_3 4

Bids solicited according to distribution $\mathcal{D}$

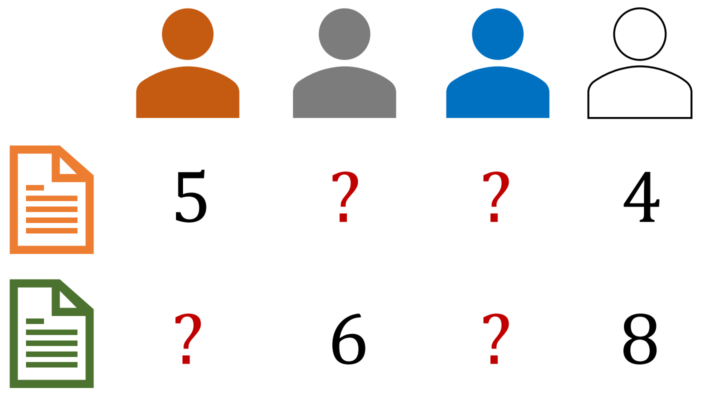

Fill in ? with L_2-error guarantees over $\mathcal{D}$

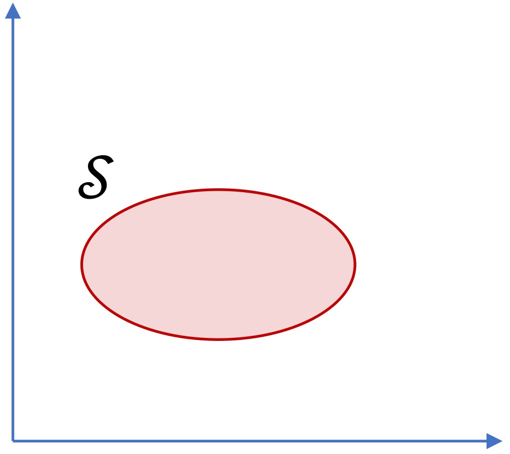

Assign for Predicted Quality

Predict important review aspects

Agreement with final decision (d)

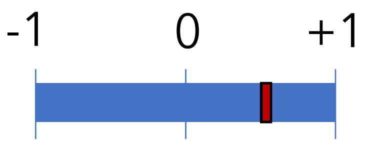

Probability of covering aspects ($p(a_i)$)

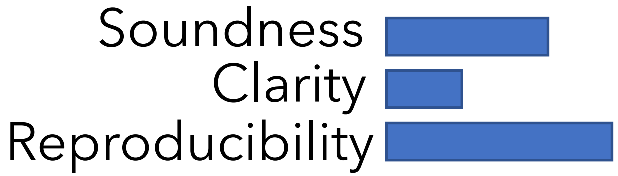

Value of assigning reviewer set C to paper p is

$$\lambda_1 \sum_{r \in C} d_{rp} + \lambda_2 \sum_{i} \max_{r \in C} p(a_i | r, p)$$

Constrained allocation of indivisible resources with submodular valuations

Online Reviewer Assignment

m reviewers available, must review n_t papers in month t

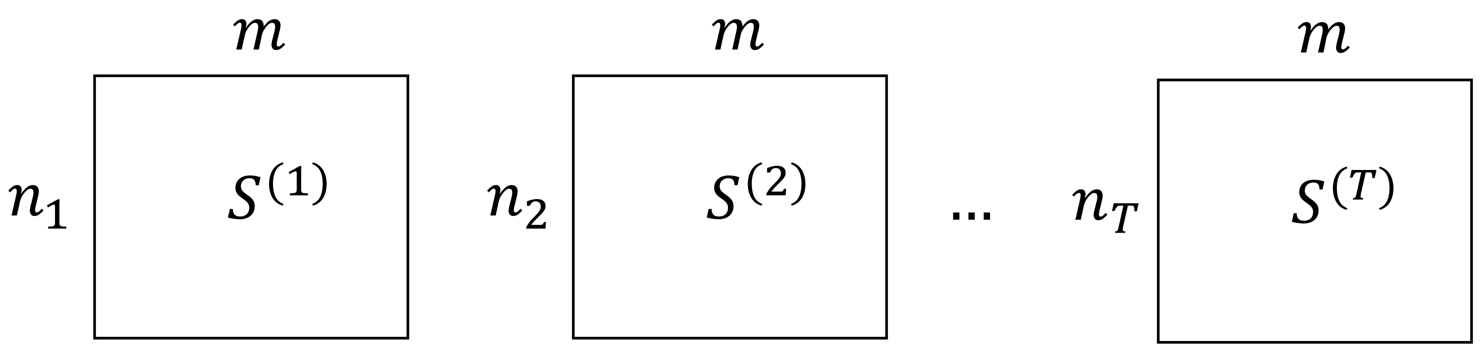

Goal: $$\max_{A^{(1)},A^{(2)},\ldots A^{(T)}} \sum_t USW(A^{(t)}, S^{(t)})$$ subj. to

$$\sum_t \sum_p A_{pr}^{(t)} \leq U_r \quad \text{for all } r$$

$$\sum_p A_{pr}^{(t)} \leq U_r^{(t)} \quad \text{for all } r, t$$

Improving RRA

Precise trade-offs in randomized rounding with constraints

Randomized rounding for discrete solution

Round $A \in \tilde{\mathcal{A}}$ to $A' \in \mathcal{A}$

Generalization of Birkhoff von Neumann decomposition

Maintains constraints of $\tilde{\mathcal{A}}$ and $\mathcal{A}$

$$\mathbb{E}[A'] = A$$

Can replace this with guarantees on USW?

Uncertainty: Nowhere to be Found

Cannot trust high affinity scores

Low scores are too pessimistic (esp. with missing data)

Outline

1

Step through a typical conference workflow

2

Introduce RRA, a framework that accounts for uncertainty

Also allows us to use new affinity score estimation methods

Robust Reviewer Assignment (RRA)

Compute an *uncertainty set* $\mathcal{S}$ containing true, unknown affinity scores S

Robust Reviewer Assignment (RRA)

Compute an *uncertainty set* $\mathcal{S}$ containing true, unknown affinity scores S

5	4	8	4
9	6	1	8
3	4	6	7

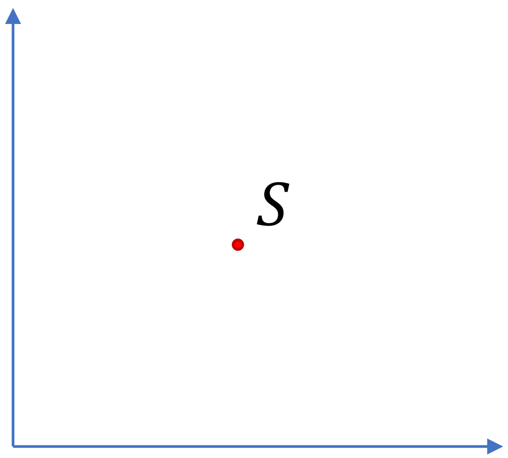

What is RAU, really?

What is RAU, really?

Maximin USW is a lower bound for *true* USW

Flexibly trades off between
certainty and high expected value

Approaches original problem as $\mathcal{S} \to S$

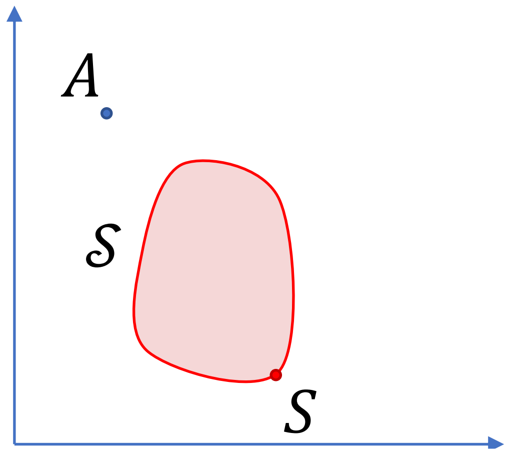

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{S}} USW(A, S)$$

Outline

1

Step through a typical conference workflow

2

Introduce RRA, a framework that accounts for uncertainty

Also allows us to use new affinity score estimation methods

Conference Workflow 1.0

1. Collect constraints
2. Collect data
3. Compute affinity scores
4. Maximize function of affinity scores
5. Adjust assignment

Conference Workflow 1.0

1. Collect constraints
2. Collect data
3. Compute affinity scores
4. Maximize function of affinity scores
5. Adjust assignment

Kobren et al. Paper matching with local fairness constraints. 2019.
Stelmakh et al. PeerReview4All: Fair and accurate reviewer assignment in peer review. 2019.
Leyton-Brown et al. Matching papers and reviewers at large conferences. 2022.
Charlin and Zemel. The Toronto paper matching system: an automated paper-reviewer assignment system. 2013


Collect constraints
Collect data
Compute affinity scores
Maximize function of affinity scores
Adjust assignment

Collect constraints → **Collect data** → Compute affinity scores → Maximize function of affinity scores → Adjust assignment

Affinity scores

	5	4	8	4
	9	6	1	8
	3	4	6	7

4 = Reviewer bid 1 + Document similarity 7 + Keyword matching 4

Collect constraints → **Collect data** → Compute affinity scores → Maximize function of affinity scores → Adjust assignment

Affinity scores

	5	4	8	4
	9	6	1	8
	3	4	6	7

4 = Reviewer bid [?] + Document similarity [7] + Keyword matching [4]

Bids often missing!


Collect constraints
Collect data
Compute affinity scores
Maximize function of affinity scores
Adjust assignment

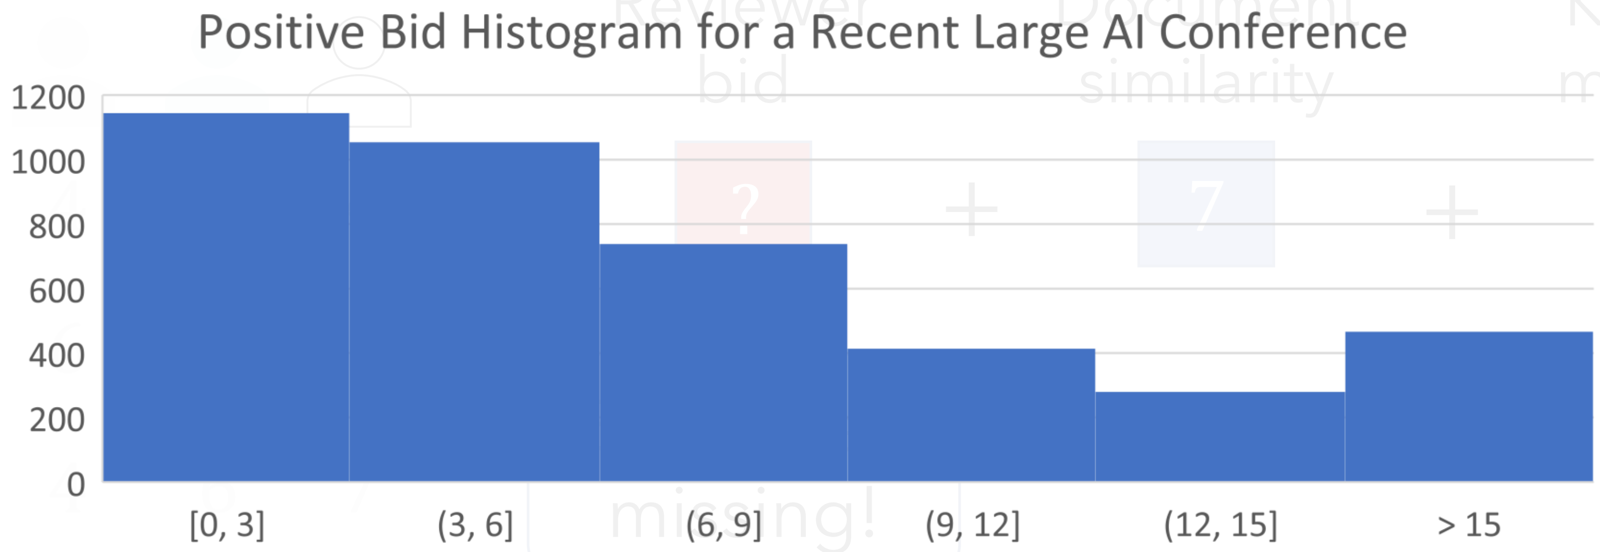
Positive Bid Histogram for a Recent Large AI Conference
1200
1000
800
600
400
200
0
[0, 3]
(3, 6]
(6, 9]
(9, 12]
(12, 15]
> 15

Collect constraints → **Collect data** → Compute affinity scores → Maximize function of affinity scores → Adjust assignment

Affinity scores

	5	4	8	4
	9	6	1	8
	3	4	6	7

4 = Reviewer bid ? + Document similarity 7 + Keyword matching 4

Bids often missing!

Affinity scores

	5	4	8	4
	9	6	1	8
	3	4	6	7

4 = Reviewer bid ? + Document similarity ? + Keyword matching 4

Reviewer bid ↑ Bids often missing!

Document similarity ↑ Error from ML models

Collect constraints → **Collect data** → Compute affinity scores → Maximize function of affinity scores → Adjust assignment

Affinity scores

5	4	8	4
9	6	1	8
3	4	6	7

4 = Reviewer bid ? + Document similarity ? + Keyword matching ?

- Reviewer bid ↑ Bids often missing!
- Document similarity ↑ Error from ML models
- Keyword matching ↑ Vocab mismatch

Collect constraints → **Collect data** → Compute affinity scores → Maximize function of affinity scores → Adjust assignment

Affinity scores

	Reviewer 1	Reviewer 2	Reviewer 3	Reviewer 4
Paper 1	5	4	8	4
Paper 2	9	6	1	8
Paper 3	3	4	6	7

4 = Reviewer bid [?] + Document similarity [?] + Keyword matching [?]

- Reviewer bid ← Bids often missing!
- Document similarity ← Error from ML models
- Keyword matching ← Vocab mismatch

All just proxies!

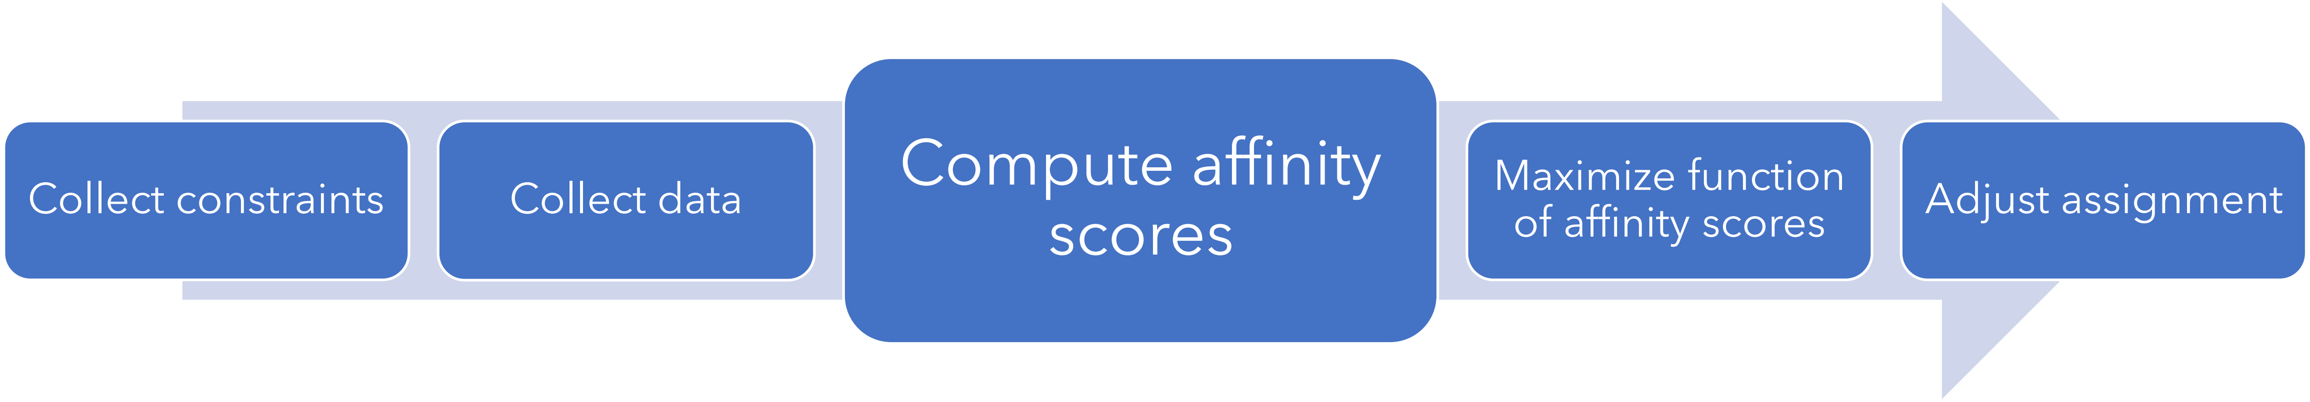
Collect constraints
Collect data
Compute affinity scores
Maximize function of affinity scores
Adjust assignment

Collect constraints → Collect data → **Compute affinity scores** → Maximize function of affinity scores → Adjust assignment

Affinity scores

	Reviewer 1	Reviewer 2	Reviewer 3	Reviewer 4
Paper 1	5	4	8	4
Paper 2	9	6	1	8
Paper 3	3	4	6	7

4 = Reviewer bid 1 + Document similarity 7 + Keyword matching 4

Collect constraints → Collect data → **Compute affinity scores** → Maximize function of affinity scores → Adjust assignment

	Reviewer bid		Document similarity		Keyword matching

$$4 = \lambda_1 \boxed{1} + \lambda_2 \boxed{7} + \lambda_3 \boxed{4}$$

Collect constraints → Collect data → **Compute affinity scores** → Maximize function of affinity scores → Adjust assignment

Open Review

CMT

	Reviewer bid		Document similarity		Keyword matching
$4 = \lambda_1$	1	$+ \lambda_2$	7	$+ \lambda_3$	4

Any unavailable data is set to 0

Underestimation means high affinity scores are high with certainty

- Collect constraints
- Collect data
- Compute affinity scores
- **Maximize function of affinity scores**
- Adjust assignment

Given *affinity scores*

5	4	8	4
9	6	1	8
3	4	6	7

Papers P and reviewers R

Valid assignments $\mathcal{A} \subseteq \{0, 1\}^{|P|\times|R|}$

Affinity score matrix $S \in \mathbb{R}_+^{|P|\times|R|}$

Maximize Utilitarian Social Welfare

$$\max_{A\in\mathcal{A}} \left[\sum_{p\in P, r\in R} A_{pr} S_{pr} = USW(A, S) \right]$$

Collect constraints → Collect data → Compute affinity scores → **Maximize function of affinity scores** → Adjust assignment

Given *affinity scores*

5	4	8	4
9	6	1	8
3	4	6	7

Papers P and reviewers R

Valid assignments $\mathcal{A} \subseteq \{0,1\}^{|P|\times|R|}$

Affinity score matrix $S \in \mathbb{R}_+^{|P|\times|R|}$

Maximize Utilitarian Social Welfare

$$\max_{A\in\mathcal{A}} \left[\sum_{p\in P, r\in R} A_{pr} S_{pr} = USW(A,S) \right]$$

ILP with totally unimodular constraints (poly-time solvable)

Taylor. On the optimal assignment of conference papers to reviewers. 2008.

Robust Reviewer Assignment (RRA)

Compute an *uncertainty set* $\mathcal{S}$ containing true, unknown affinity scores S

5	4	8	4
9	6	1	8
3	4	6	7

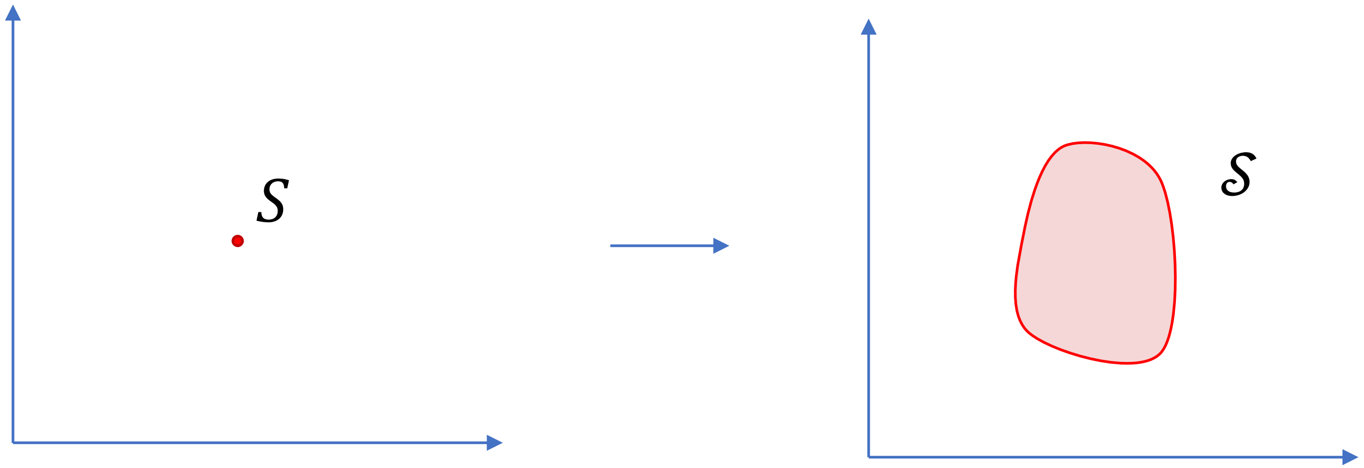

Maximize the worst-case welfare over the *uncertainty set*

Robust Reviewer Assignment (RRA)

Compute an *uncertainty set* $\mathcal{S}$ containing true, unknown affinity scores S

	5	4	8	4
	9	6	1	8
	3	4	6	7

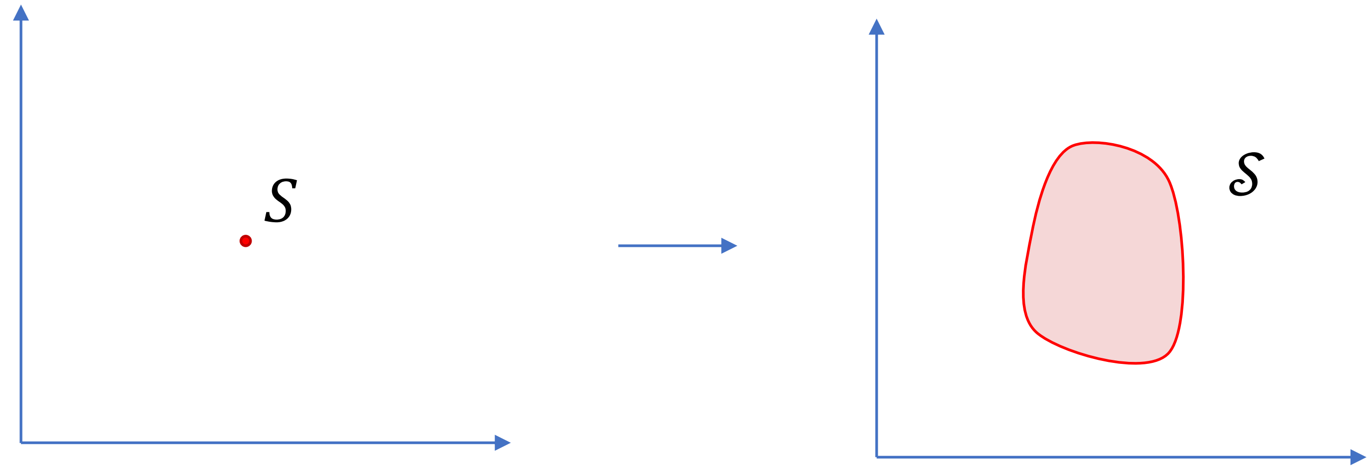

Maximize the worst-case welfare over the *uncertainty set*

ICLR Experiments

Got all submitted papers to last 5 years of ICLR

Used authors as “reviewers”

Keywords = paper topics and author expertise

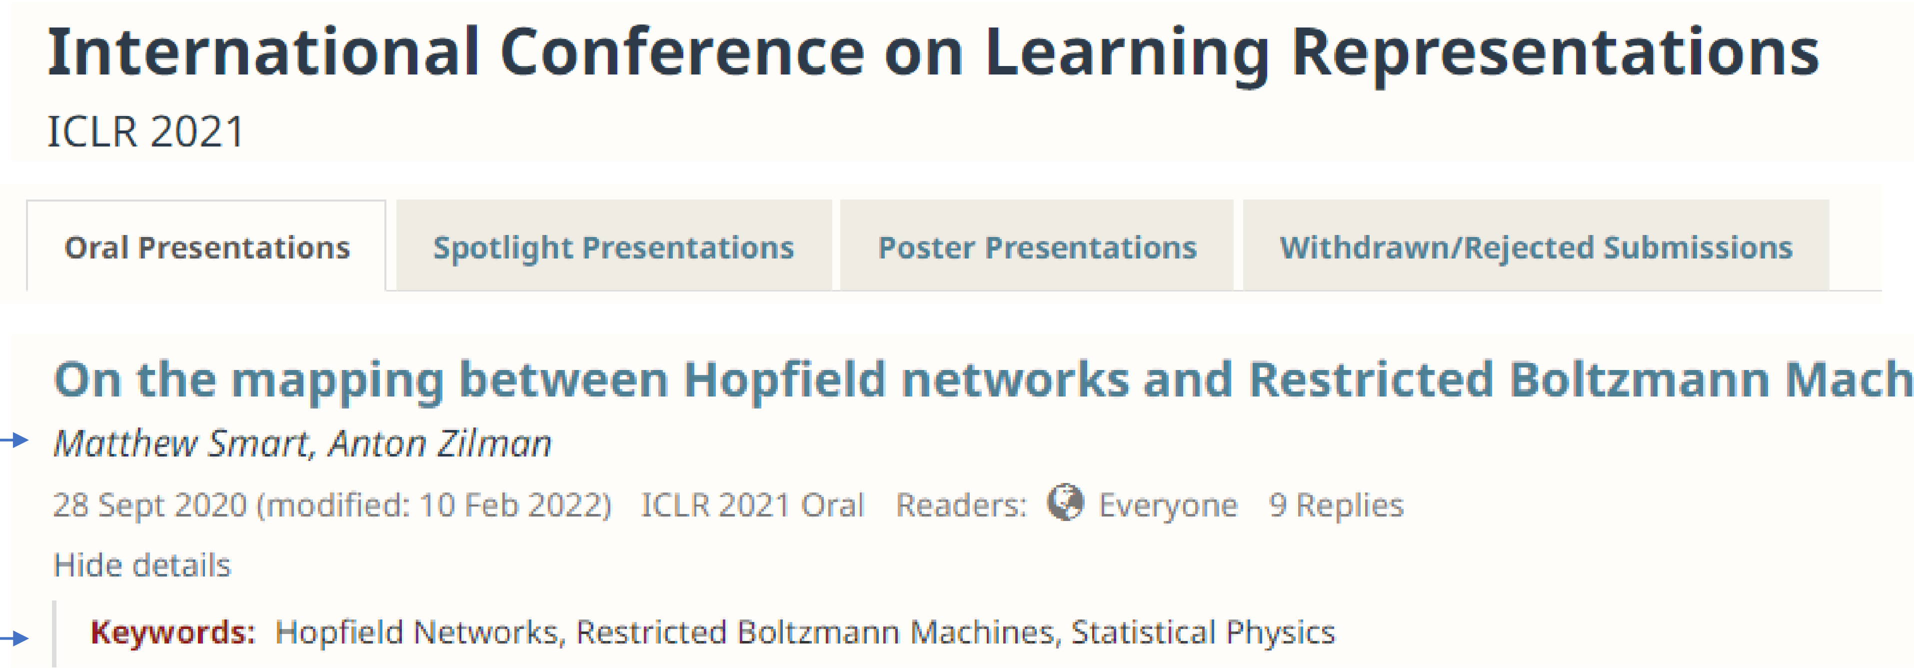

ICLR Experiments

Got all submitted papers to last 5 years of ICLR

Used authors as "reviewers"

Keywords = paper topics and author expertise

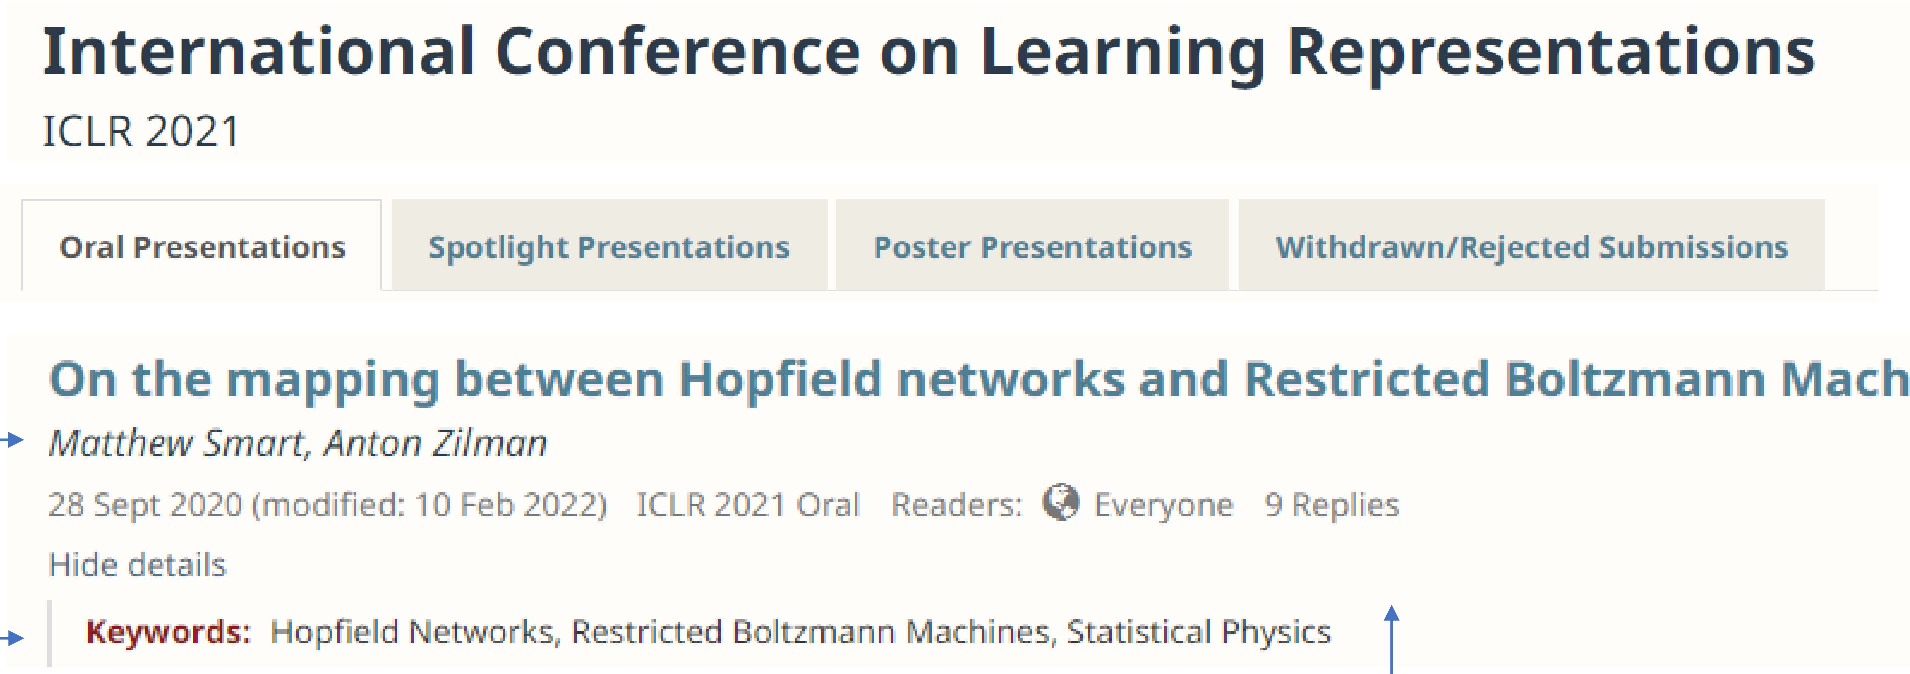

Compute $\mathcal{S}$ as a contour of $\mathcal{N}(\mu, \Sigma)$

μ based on keyword overlap

$\mathrm{diag}(\Sigma) \propto 1/(\#\ \text{keywords})$

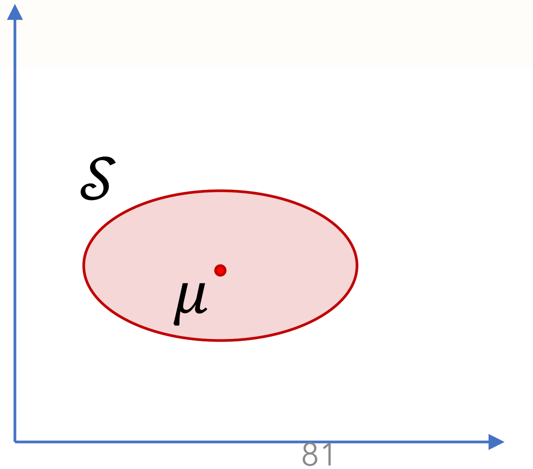

Robust Reviewer Assignment (RRA)

Relax discrete -> continuous

Subgradient-ascent optimization

Randomized rounding for discrete solution

Robust Reviewer Assignment (RRA)

Relax discrete -> continuous

$\mathcal{A} \subseteq \{0,1\}^{m\times n} \rightarrow \tilde{\mathcal{A}} \subseteq [0,1]^{m\times n}$

Subgradient-ascent optimization

Randomized rounding for discrete solution

Robust Reviewer Assignment (RRA)

Relax discrete -> continuous

$\mathcal{A} \subseteq \{0,1\}^{m\times n} \rightarrow \tilde{\mathcal{A}} \subseteq [0,1]^{m\times n}$

Subgradient-ascent optimization

Solve $\max_{A\in\tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

Randomized rounding for discrete solution

Robust Reviewer Assignment (RRA)

Subgradient-ascent optimization

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

Robust Reviewer Assignment (RRA)

Subgradient-ascent optimization

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

Robust Reviewer Assignment (RRA)

Subgradient-ascent optimization

T iterations:

- Compute $f(A)$ as the adversary
- Step in direction of $\partial_A f(A)$
- Project to feasible set $\tilde{\mathcal{A}}$

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

Robust Reviewer Assignment (RRA)

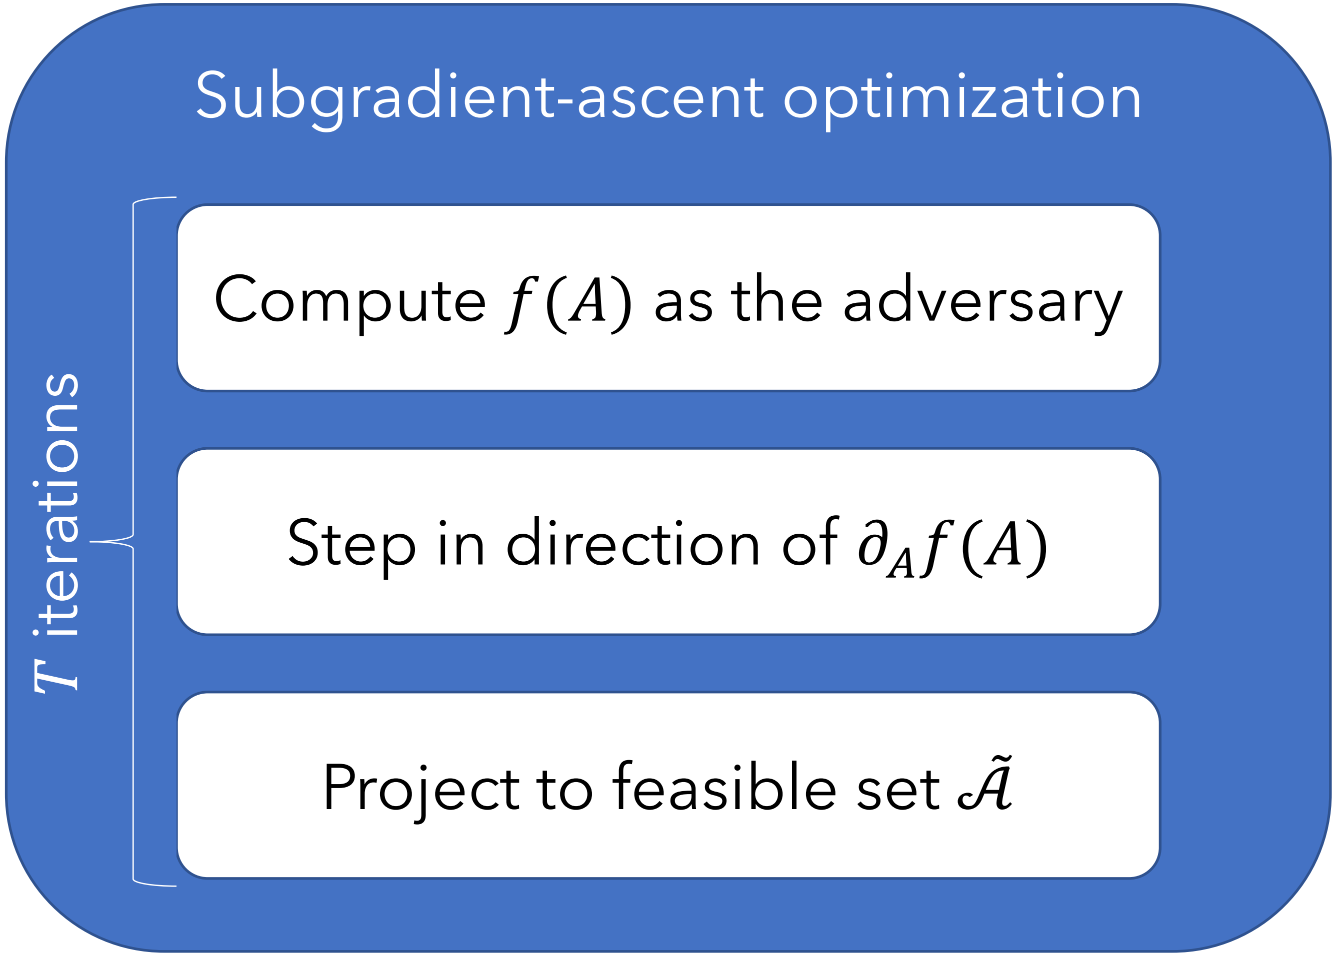

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

$$f(A) = \min_{S \in \mathcal{S}} USW(A, S)$$

Robust Reviewer Assignment (RRA)

Subgradient-ascent optimization

T iterations:

1. Compute $f(A)$ as the adversary
2. Step in direction of $\partial_A f(A)$
3. Project to feasible set $\tilde{\mathcal{A}}$

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

$$f(A) = \min_{S \in \mathcal{S}} USW(A, S)$$

$$\nabla_A USW(A, \operatorname{argmin}_{S \in \mathcal{S}} USW(A, S)) \in \partial_A f(A)$$

$$A' \leftarrow A + \nabla_A USW(A, \operatorname{argmin}_{S \in \mathcal{S}} USW(A, S))$$

Robust Reviewer Assignment (RRA)

Subgradient-ascent optimization

T iterations:

- Compute $f(A)$ as the adversary
- Step in direction of $\partial_A f(A)$
- Project to feasible set $\tilde{\mathcal{A}}$

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

$$f(A) = \min_{S \in \mathcal{S}} USW(A, S)$$

$$\nabla_A USW(A, \operatorname{argmin}_{S \in \mathcal{S}} USW(A, S)) \in \partial_A f(A)$$

$$A' \leftarrow A + \nabla_A USW(A, \operatorname{argmin}_{S \in \mathcal{S}} USW(A, S))$$

$$A'' \leftarrow \operatorname{argmin}_{A \in \tilde{\mathcal{A}}} \|A - A'\|_2$$

Robust Reviewer Assignment (RRA)

Relax discrete -> continuous

$\mathcal{A} \subseteq \{0,1\}^{m \times n} \rightarrow \tilde{\mathcal{A}} \subseteq [0,1]^{m \times n}$

Subgradient-ascent optimization

Solve $\max_{A \in \tilde{\mathcal{A}}} f(A)$ by steps in $\partial_A f(A)$

Randomized rounding for discrete solution

Round $A \in \tilde{\mathcal{A}}$ to $A' \in \mathcal{A}$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

Proof sketch:

$$\min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + X)$$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A\in\mathcal{A}} \min_{S\in\mathcal{B}_\epsilon(S_0)} USW(A,S) = \max_{A\in\mathcal{A}} USW(A,S_0)$$

Proof sketch:

$$\min_{S\in\mathcal{B}_\epsilon(S_0)} USW(A,S) = \min_{X\in\mathcal{B}_\epsilon(0)} USW(A,S_0+X)$$

$$= USW(A,S_0) + \min_{X\in\mathcal{B}_\epsilon(0)} USW(A,X)$$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

Proof sketch:

$$\min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + X)$$

$$= USW(A, S_0) + \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, X)$$

Minimized when X points as far in the direction opposite A as possible:

$$X = -A \frac{\epsilon}{\|A\|_2}$$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

Proof sketch:

$$\min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + X)$$

$$= USW(A, S_0) + \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, X)$$

$$= USW(A, S_0) + USW\left(A, -A\frac{\epsilon}{\|A\|_2}\right)$$

Minimized when X points as far in the direction opposite A as possible:

$$X = -A\frac{\epsilon}{\|A\|_2}$$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

Proof sketch:

$$\min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + X)$$

$$= USW(A, S_0) + \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, X)$$

$$= USW(A, S_0) + USW\left(A, -A\frac{\epsilon}{\|A\|_2}\right)$$

$$= USW(A, S_0) - \|A\|_2^2 \frac{\epsilon}{\|A\|_2}$$

Minimized when X points as far in the direction opposite A as possible:

$$X = -A\frac{\epsilon}{\|A\|_2}$$

Spherical Uncertainty Sets

Theorem 2:

$$\max_{A \in \mathcal{A}} \min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

Proof sketch:

$$\min_{S \in \mathcal{B}_\epsilon(S_0)} USW(A, S) = \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + X)$$

$$= USW(A, S_0) + \min_{X \in \mathcal{B}_\epsilon(0)} USW(A, X)$$

Minimized when X points as far in the direction opposite A as possible:

$$X = -A\frac{\epsilon}{\|A\|_2}$$

$$= USW(A, S_0) + USW\left(A, -A\frac{\epsilon}{\|A\|_2}\right)$$

$$= USW(A, S_0) - \|A\|_2^2 \frac{\epsilon}{\|A\|_2}$$

$$= USW(A, S_0) - \|A\|_2 \epsilon$$

Spherical Uncertainty Sets

Theorem 2:
$$\max_{A\in\mathcal{A}} \min_{S\in\mathcal{B}_\epsilon(S_0)} USW(A,S) = \max_{A\in\mathcal{A}} USW(A,S_0)$$

Proof sketch:

$$\min_{S\in\mathcal{B}_\epsilon(S_0)} USW(A,S) = \min_{X\in\mathcal{B}_\epsilon(0)} USW(A,S_0+X)$$

$$= USW(A,S_0) + \min_{X\in\mathcal{B}_\epsilon(0)} USW(A,X)$$

Minimized when X points as far in the direction opposite A as possible:

$$X = -A\frac{\epsilon}{\|A\|_2}$$

$$= USW(A,S_0) + USW\left(A, -A\frac{\epsilon}{\|A\|_2}\right)$$

$$= USW(A,S_0) - \|A\|_2^2\frac{\epsilon}{\|A\|_2}$$

$$= USW(A,S_0) - \|A\|_2\epsilon$$

Constant, due to constraints

Future Work

- Be sure to publicize FairSeq, our recently submitted paper, and the workshop.
- Further experiments
    - Take datasets of affinity scores, drop some and fill in remaining values
    - Build a predictive model of affinity scores from old conferences
    - Collect data from conference organizers
        - Discuss the collaborative filtering bid model?
- Combining fairness and robustness
    - Add an envy penalty to the objective function
    - Maximize expected egalitarian welfare
- Rounding tradeoffs?

Fair Division at IJCAI

- Anyone want to co-organize?
- Plan to submit papers!
- Mention proposed schedule, etc.

Collect constraints → Collect data → **Compute affinity scores** → Maximize function of affinity scores → Adjust assignment

Affinity scores

Open Review

Reviewer bid	Document similarity	Keyword matching

$$4 = \lambda_1 \boxed{1} + \lambda_2 \boxed{7} + \lambda_3 \boxed{4}$$

$$\text{Base Aggregated Score} = \begin{cases} \frac{1}{4}\text{TPMS} + \frac{1}{4}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if all data is present} \\ \frac{1}{2}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if TPMS is missing} \\ \frac{1}{2}\text{TPMS} + \frac{1}{2}\text{SAM} & \text{if ACL is missing} \\ \text{SAM} & \text{if both ACL and TPMS are missing} \end{cases}$$

$$\text{aggscore} = (\text{Base aggregated score})^{\text{bidscore}}$$

Collect constraints → Collect data → **Compute affinity scores** → Maximize function of affinity scores → Adjust assignment

Affinity scores

Reviewer bid

Document similarity

Keyword matching

$$4 \quad = \lambda_1$$

Does this necessarily correlate with review quality?

Base Aggregated Score

(SAM

if both ACL and TPMS are missing

$$\text{aggscore} = (\text{Base aggregated score})^{\text{bidscore}}$$

Allocation Deciderata

High affinity (welfare/USW)

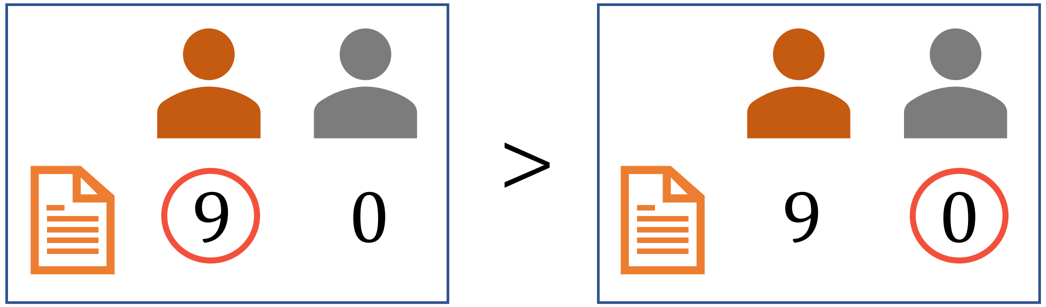

Fair to papers

Fast to compute

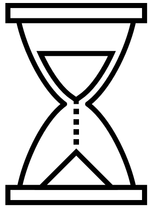

Adapt to constraints

Allocation Deciderata

High affinity (welfare/USW)

Fair to papers

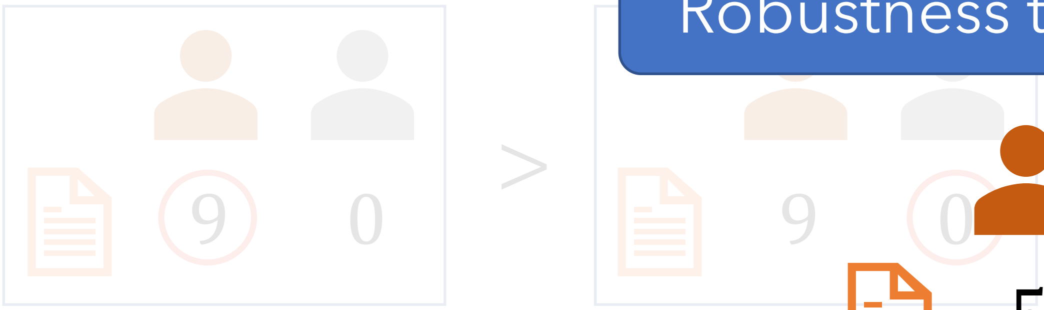

Robustness to noisy affinities

Fast to compute

Adapt to constraints

Fairness to Papers

Envy-freeness up to 1 Item (EF1)

Ensures pairwise balance in affinity scores

Large literature with simple/fast algorithms

Envy-freeness up to 1 Item (EF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A, then it is due to at most one reviewer

Envy-freeness up to 1 Item (EF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A, then it is due to at most one reviewer

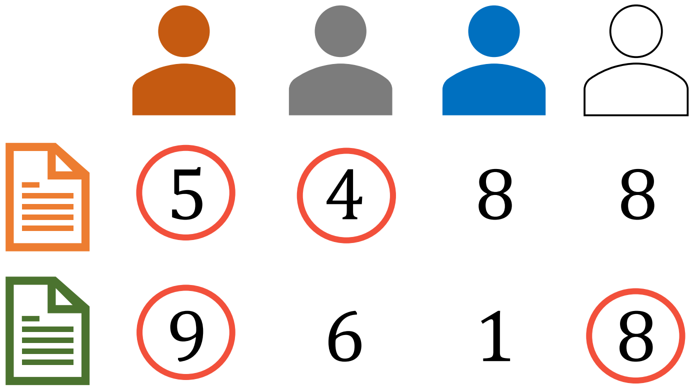

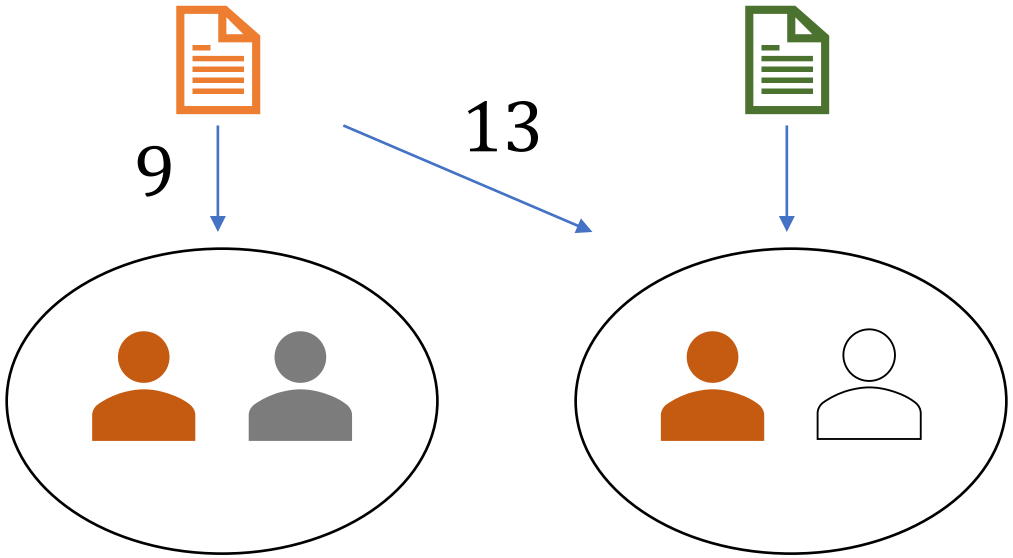

Envy-freeness up to 1 Item (EF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A, then it is due to at most one reviewer

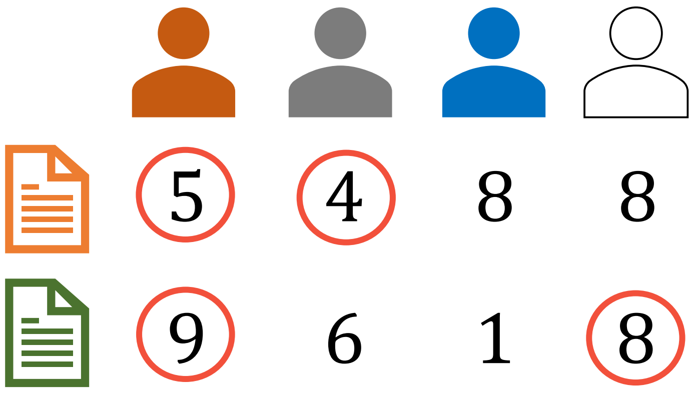

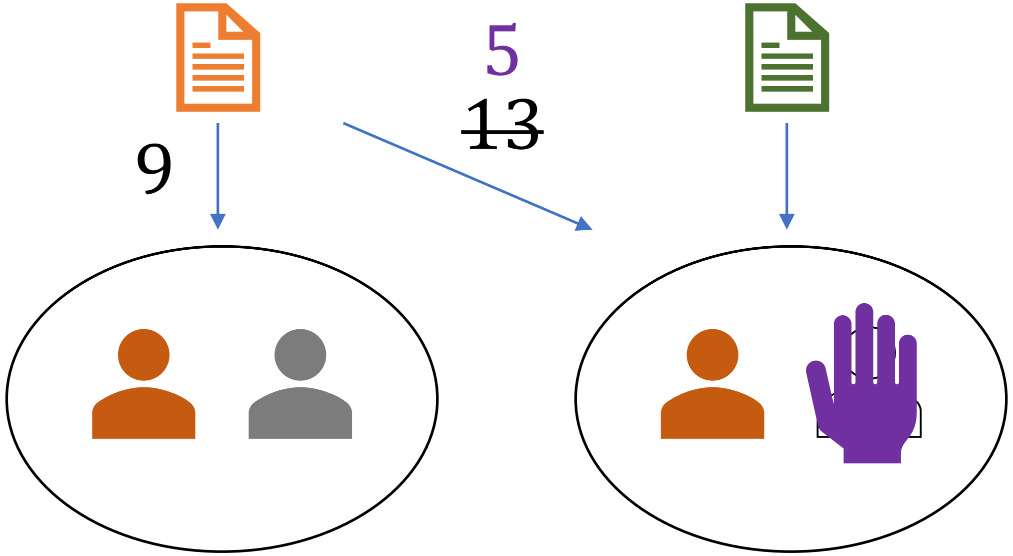

Envy-freeness up to 1 Item (EF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A, then it is due to at most one reviewer

Does not make sense when papers have variable demands

Weighted EF1 (WEF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A **after adjusting for paper demands**, then it is due to at most one reviewer

Weighted EF1 (WEF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A **after adjusting for paper demands**, then it is due to at most one reviewer

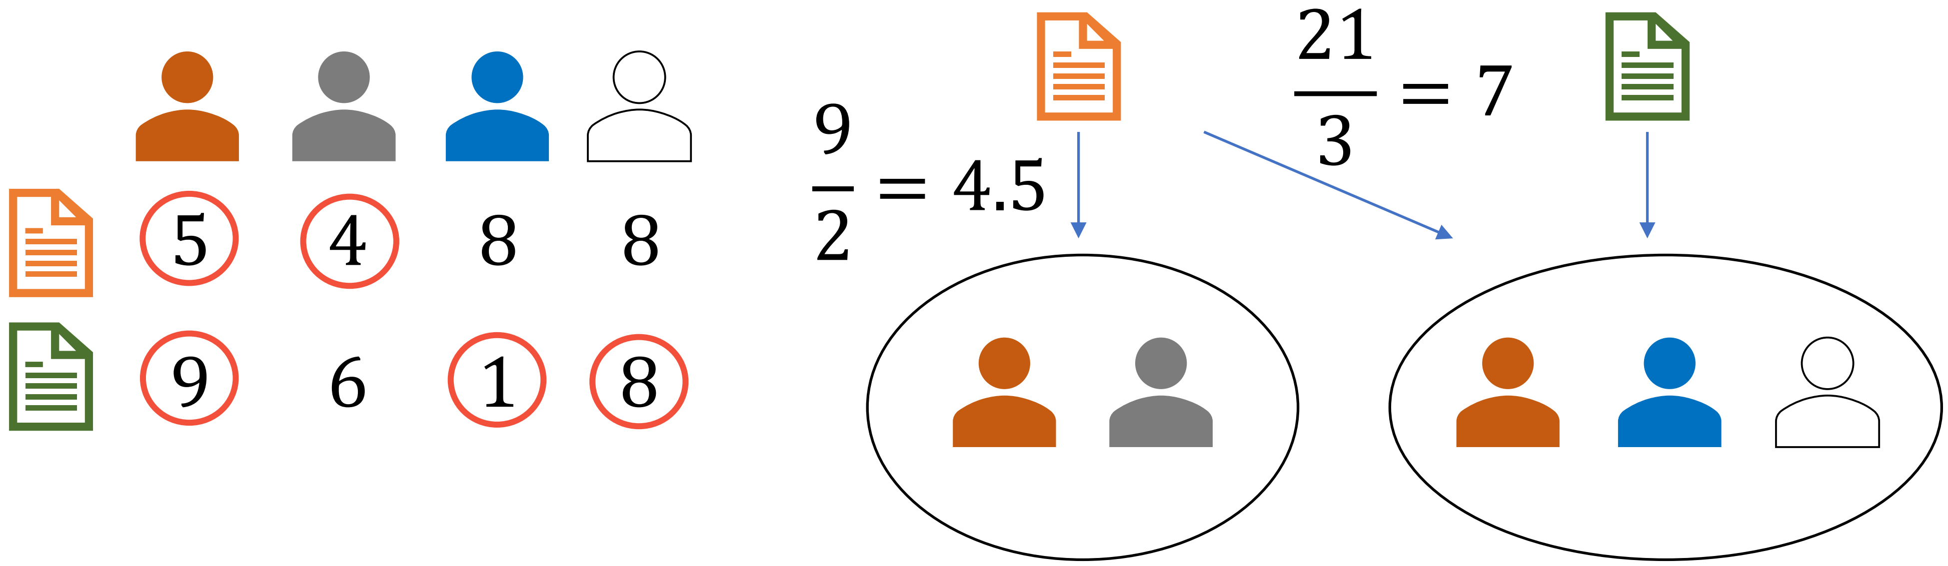

Weighted EF1 (WEF1)

For all pairs of papers A and B, if B receives a set of reviewers that is better suited for A **after adjusting for paper demands**, then it is due to at most one reviewer

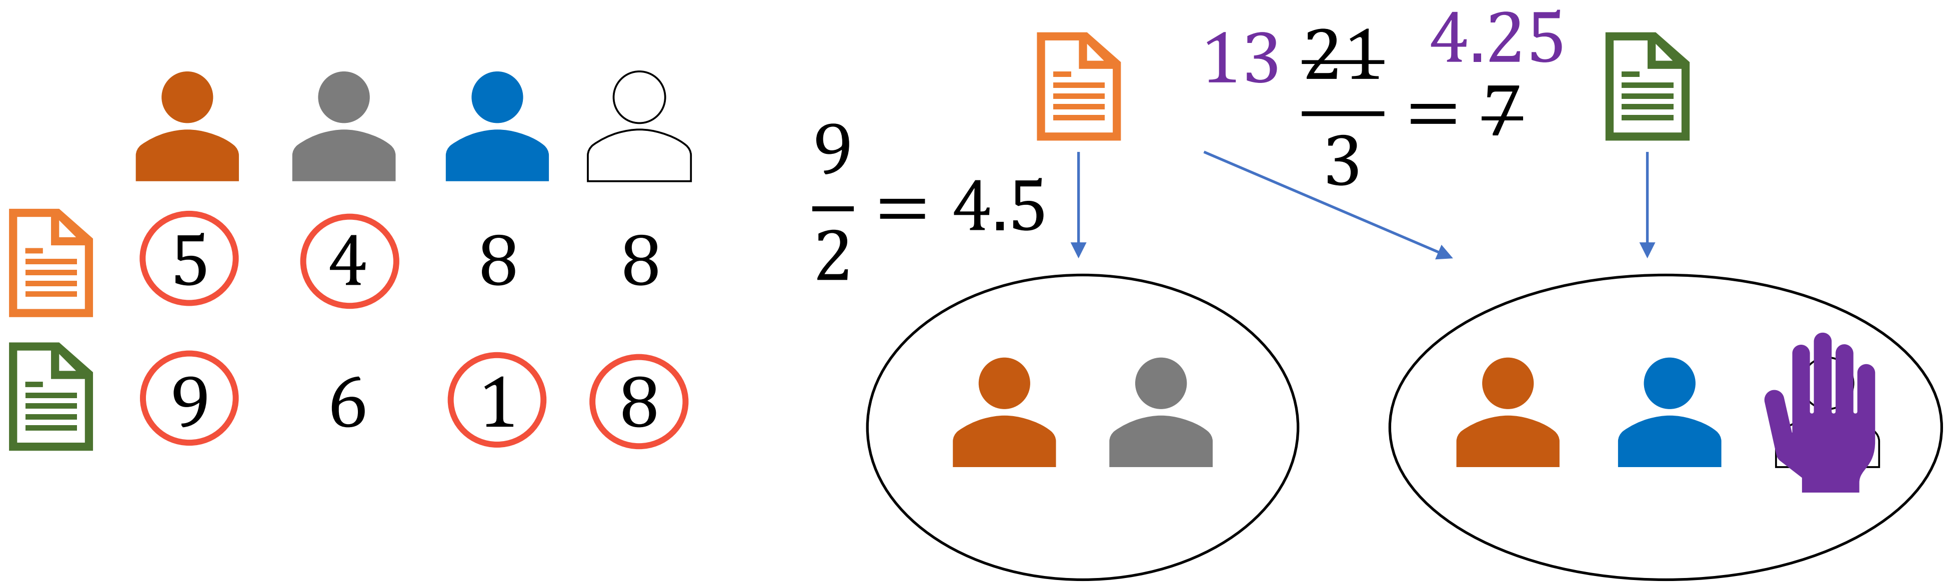

Picking Sequences for EF1 and Welfare

Put papers in some order,
and assign reviewers in that order

Choose the order for EF1 &
(approximate) max welfare

Picking Sequences for EF1 and Welfare

I Will Have Order! Optimizing Orders for Fair Reviewer Assignment.
Payan and Zick, IJCAI 2022.

Fixed order, repeated over rounds (*round-robin order*) is EF1

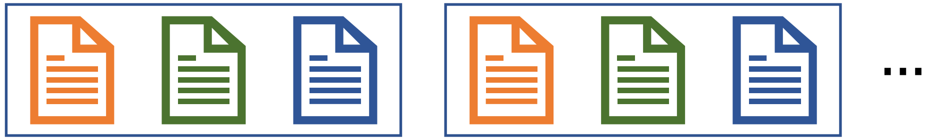

Does not work for variable paper demands

Finding a high-welfare round-robin order is slow

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

0/2	5	4	8	4
0/2	9	6	1	8
0/3	3	4	6	7

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

0/2	5	4	8	4
1/2	9	6	1	8
0/3	3	4	6	7

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

0/2		5	4	8	4
1/2		9	6	1	8
0/3		3	4	6	7

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

1/2		5	4	8	4
1/2		9	6	1	8
0/3		3	4	6	7

FairSequence

Picking sequence that assigns in order of fraction of demand satisfied

Ties broken to greedily maximize affinity

Very fast

Satisfies WEF1

High welfare in practice

FairSequence – Very Fast!

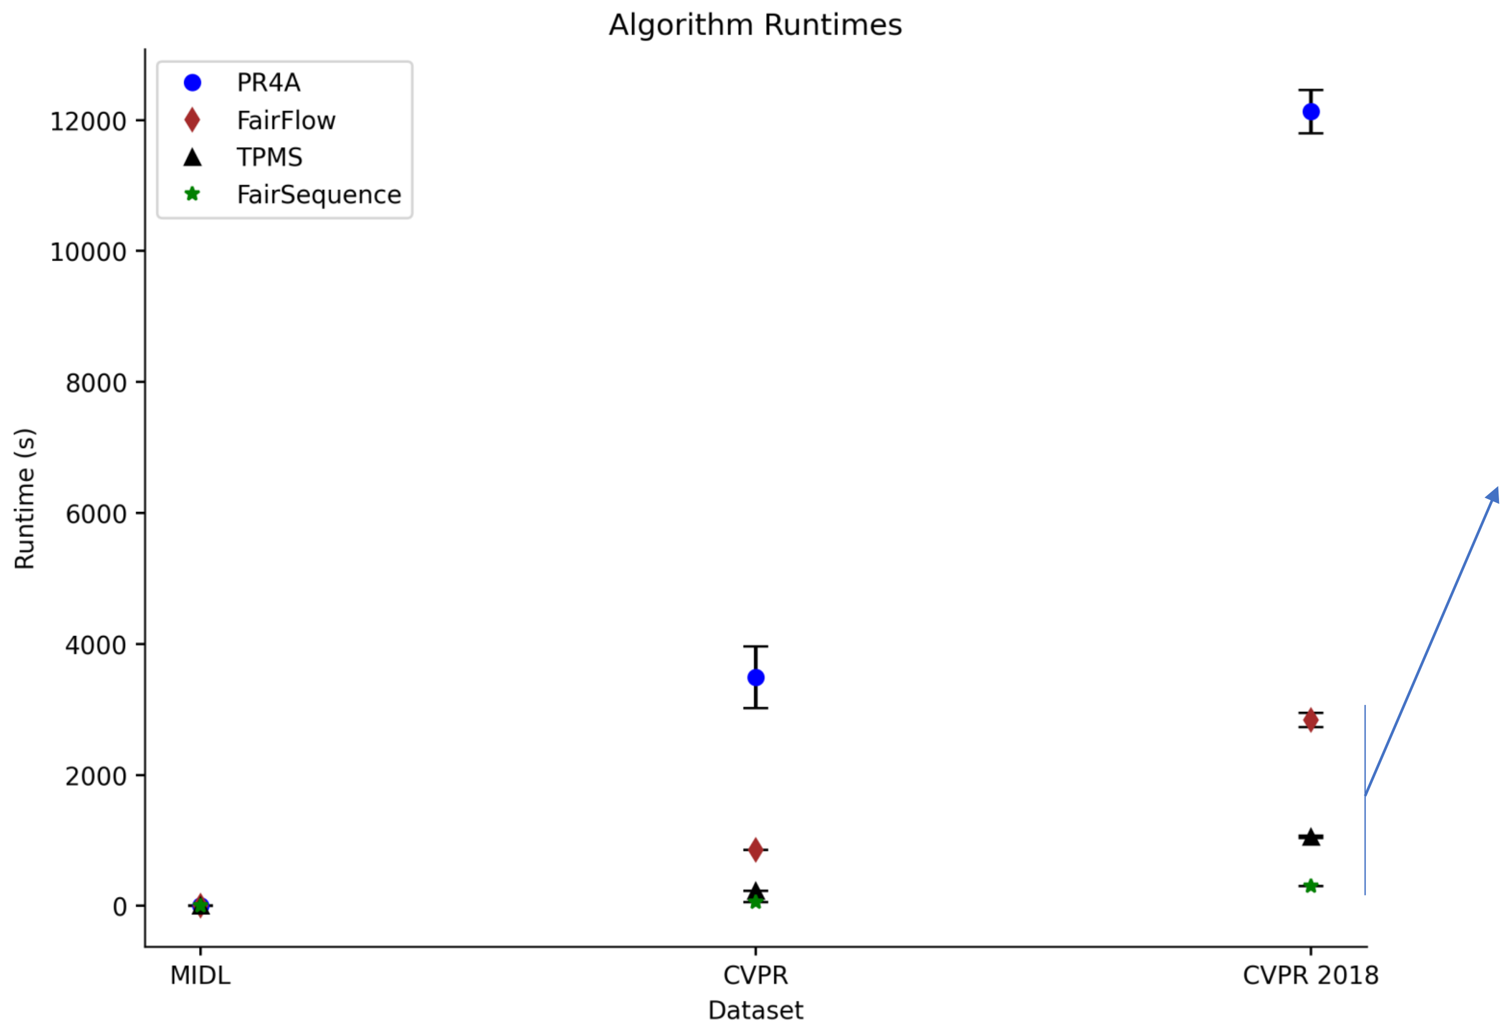

~10x faster than FairFlow!

~3x faster than TPMS!

Welfare and Fairness

		TPMS (OPT)	FairFlow	PR4A	GRRR (Our Approaches)	FairSeq (Our Approaches)
MIDL	USW (% of OPT)	100%	100%	98%	98%	99%
	# EF1 Viol.	0	0	0	0	0
CVPR	USW (% of OPT)	100%	96%	94%	88%	92%
	# EF1 Viol.	473545	23344	82	0	0
CVPR '18	USW (% of OPT)	100%	97%	97%	94%	96%
	# EF1 Viol.	134	25	2	0	0

FairSequence Fits the Criteria

Fairness	Ours are the only approaches that satisfy EF1
Welfare	High USW w.r.t. TPMS (OPT) and algorithms used in practice
Speed	~10x speedup compared to fair competitors
Flexibility	Simplicity ➔ flexibility

Look for FairSequence!

Available in Open Review

Can we trust affinity scores?

Given *affinity scores*

5	4	8	4
9	6	1	8
3	4	6	7

$= f\Big($

Reviewer bids

5	0	5	1
5	0	1	5
0	0	3	0

Document-based scores

1	7	6	7
2	2	1	4
3	3	9	8

Keyword matching

2	3	5	4
7	8	1	9
3	5	1	7

$\Big)$

Can we trust affinity scores?

Given *affinity scores*

5	4	8	4
9	6	1	8
3	4	6	7

$= f$

Reviewer bids

5	0	5	1
5	0	1	5
0	0	3	0

Document-based scores

1	7	6	7
2	2	1	4
3	3	9	8

Keyword matching

2	3	5	4
7	8	1	9
3	5	1	7

Can we trust affinity scores?

Reviewer bids

5	0	5	1
5	0	1	5
0	0	3	0

Document-based scores

1	7	6	7
?	2	1	4
3	3	9	8

Keyword matching

2	3	5	4
7	8	1	9
3	5	1	7

Can we trust affinity scores?

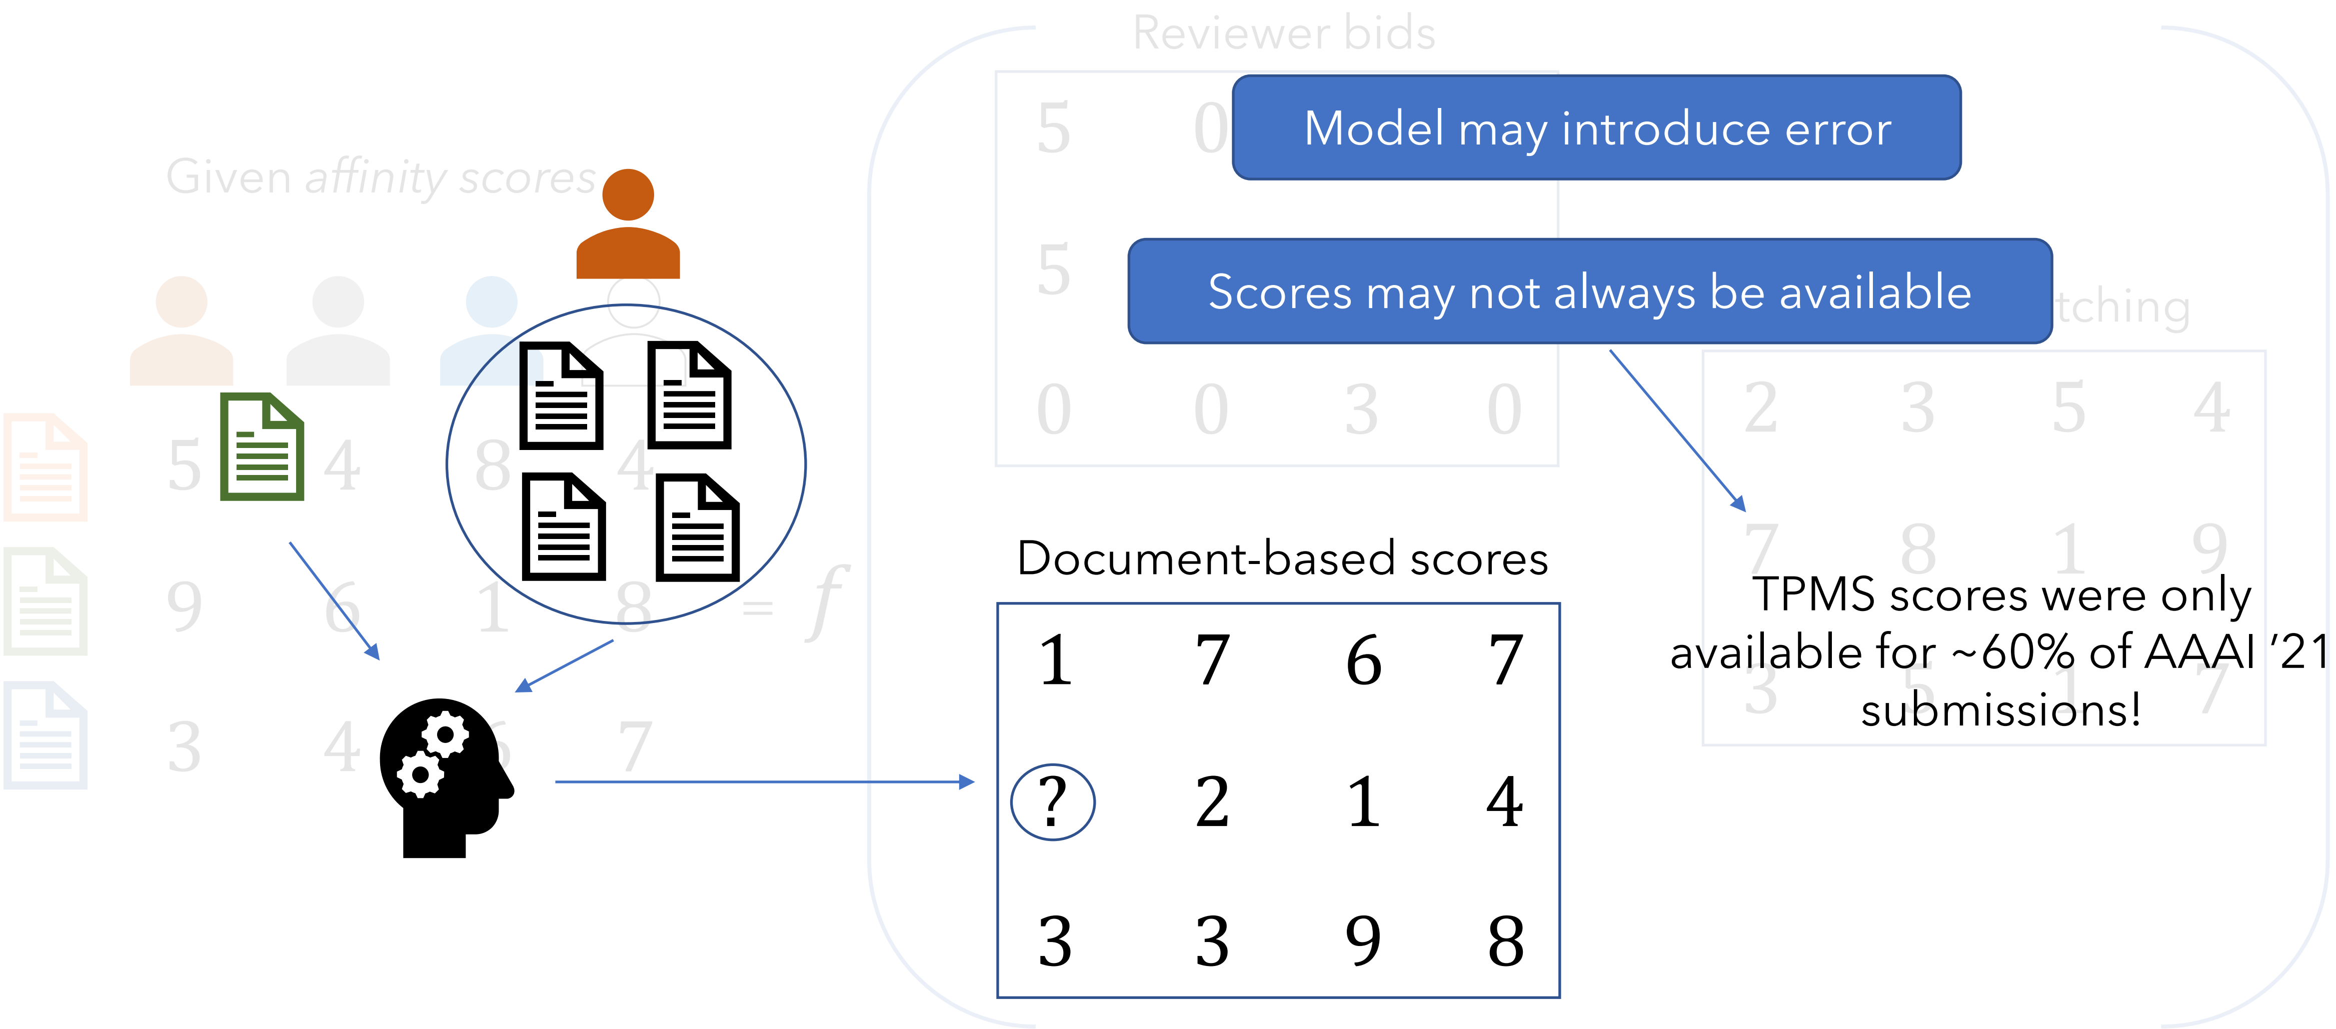

1	7	6	7
?	2	1	4
3	3	9	8

Leyton-Brown et al. 22, Matching Papers and Reviewers at Large Conferences

Robust Reviewer Assignment

Feasible region $\mathcal{S}$ for affinity scores

Solution space $\mathcal{A}$ of valid assignments

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{S}} USW(A, S_A)$$

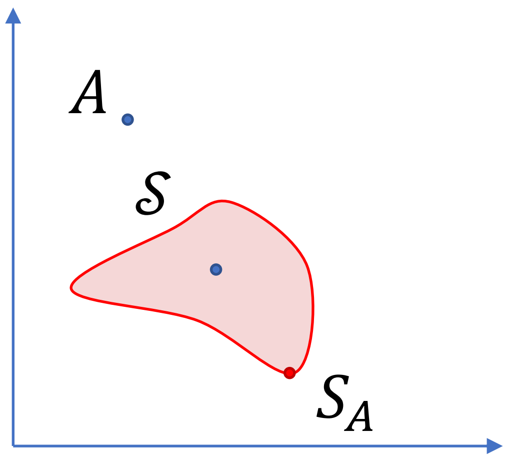

Goals:

1. Find plausible models for $\mathcal{S}$ with simple solutions
2. Build general purpose robust optimization algorithm
3. Enable alternative welfare functions

$\mathcal{S}$ as Box Constraints

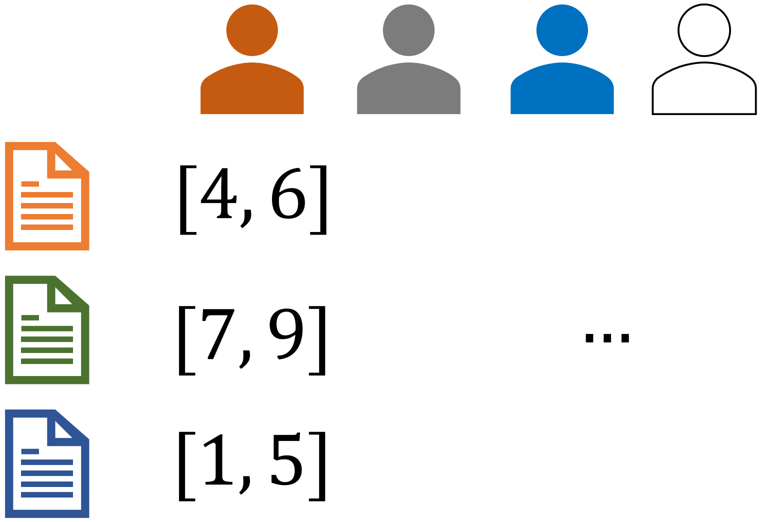

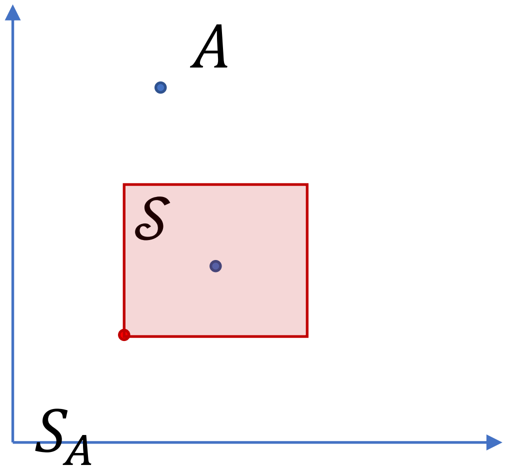

$\mathcal{S}$ as Box Constraints

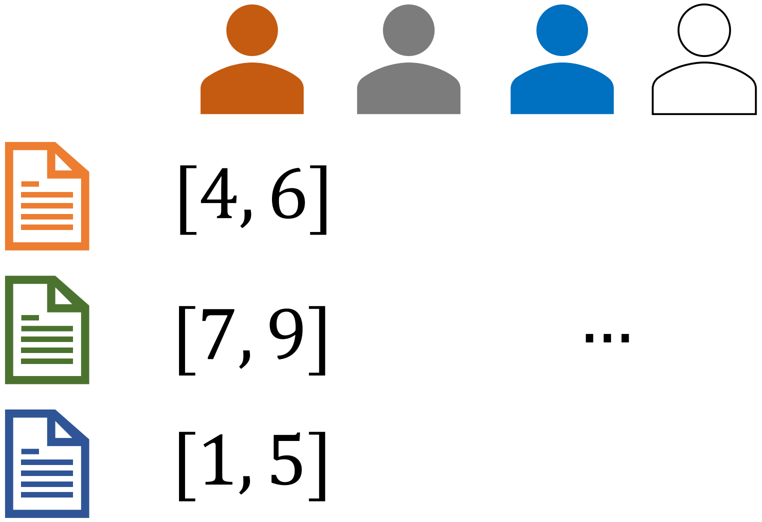

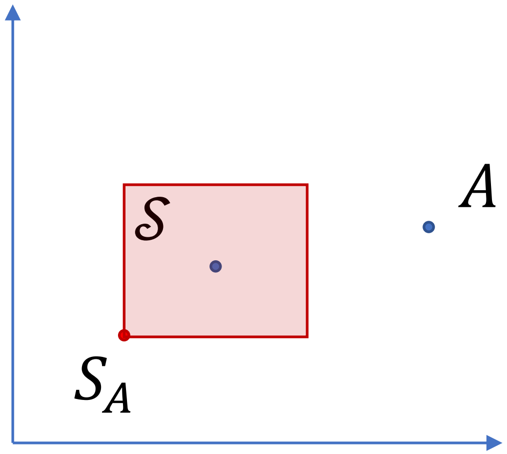

$\mathcal{S}$ as Box Constraints

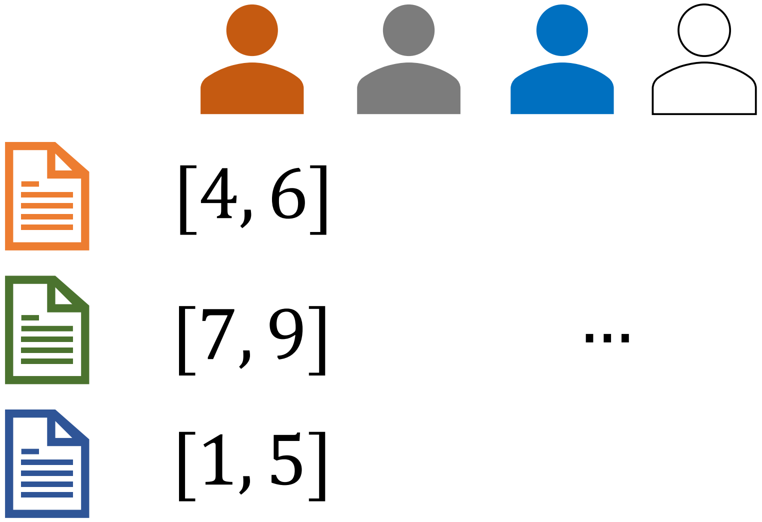

$\underline{S} =$

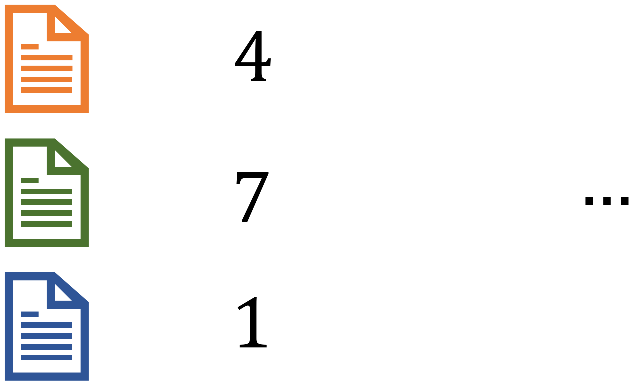

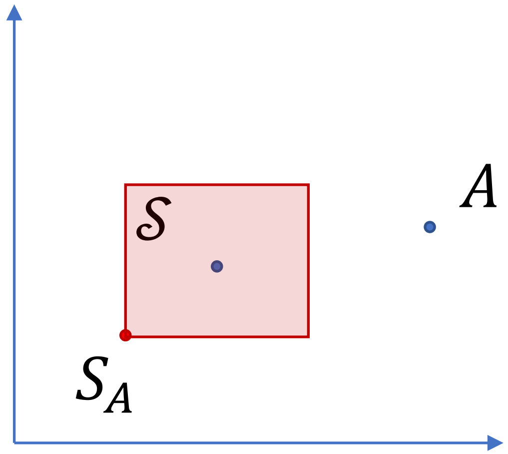

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{S}} USW(A, S_A)$$

$$= \max_{A \in \mathcal{A}} USW(A, \underline{S})$$

$\mathcal{S}$ as Spherical Noise

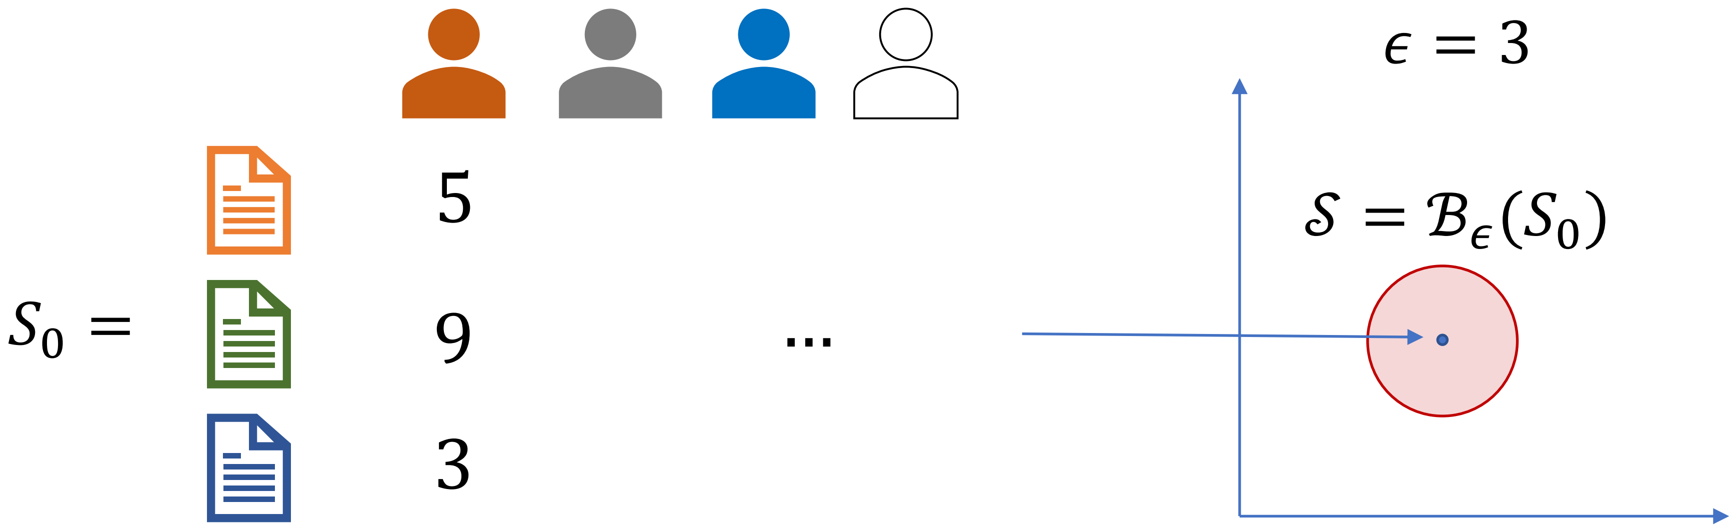

$\mathcal{S}$ as Spherical Noise

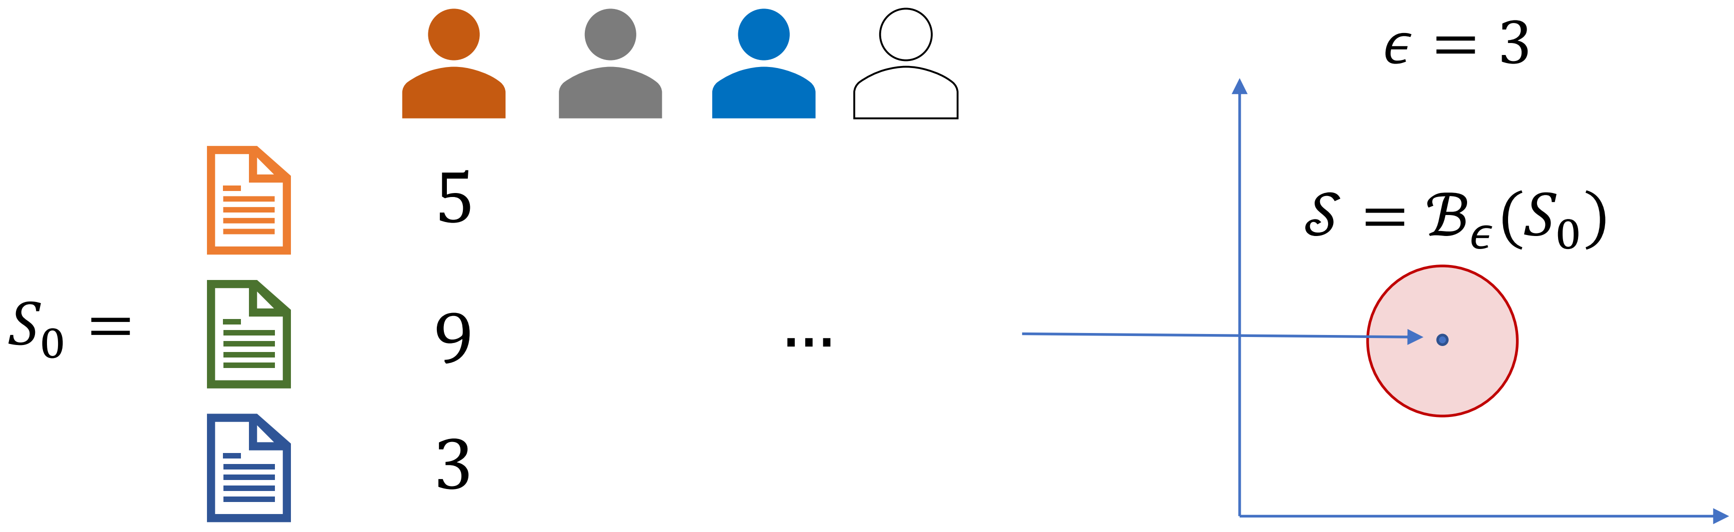

Theorem:

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

The welfare maximizer is robust to spherical noise.

$\mathcal{S}$ as Spherical Noise

Theorem:

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

$$\begin{aligned} \min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) &= \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + x) \\ &= USW(A, S_0) + \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, x) \end{aligned}$$

$\mathcal{S}$ as Spherical Noise

Theorem:

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

$$\min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) = \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + x)$$

$$= USW(A, S_0) + \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, x)$$

Minimized when x points as far in the direction opposite A as possible:

$$x = -A \frac{\epsilon}{||A||_2}$$

$\mathcal{S}$ as Spherical Noise

Theorem:

$$\max_{A \in \mathcal{A}} \min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) = \max_{A \in \mathcal{A}} USW(A, S_0)$$

$$\begin{aligned}\min_{S_A \in \mathcal{B}_\epsilon(S_0)} USW(A, S_A) &= \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, S_0 + x) \\ &= USW(A, S_0) + \min_{x \in \mathcal{B}_\epsilon(0)} USW(A, x) \\ &= USW(A, S_0) + USW\left(A, -A\frac{\epsilon}{||A||_2}\right) \\ &= USW(A, S_0) - ||A||_2^2 \frac{\epsilon}{||A||_2} \\ &= USW(A, S_0) - ||A||_2\, \epsilon\end{aligned}$$

Minimized when x points as far in the direction opposite A as possible:

$$x = -A\frac{\epsilon}{||A||_2}$$

Constant, due to constraints

AAAI's Affinity Scores

$$\text{Base Aggregated Score} = \begin{cases} \frac{1}{4}\text{TPMS} + \frac{1}{4}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if all data is present} \\ \frac{1}{2}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if TPMS is missing} \\ \frac{1}{2}\text{TPMS} + \frac{1}{2}\text{SAM} & \text{if ACL is missing} \\ \text{SAM} & \text{if both ACL and TPMS are missing} \end{cases}$$

Noise varies by score availability

Also modeled by elliptical noise

AAAI's Affinity Scores

$$\text{Base Aggregated Score} = \begin{cases} \frac{1}{4}\text{TPMS} + \frac{1}{4}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if all data is present} \\ \frac{1}{2}\text{ACL} + \frac{1}{2}\text{SAM} & \text{if TPMS is missing} \\ \frac{1}{2}\text{TPMS} + \frac{1}{2}\text{SAM} & \text{if ACL is missing} \\ \text{SAM} & \text{if both ACL and TPMS are missing} \end{cases}$$

Noise varies by score availability

Also modeled by elliptical noise

(Groupwise) Egalitarian Welfare

$$ESW(A,S) = \min_{i \in N} \sum_{j \in R} A_{ij} S_{ij}$$

If we have groups $G = \{G_1, G_2 \ldots G_g\}$
(e.g. subject areas/tracks),
we define groupwise ESW:

$$ESW(A,S,G) = \min_{G_k \in G} \sum_{i \in G_k} \sum_{j \in R} A_{ij} S_{ij}$$

(Groupwise) Egalitarian Welfare

Fair to papers

$$ESW(A,S) = \min_{i \in N} \sum_{j \in R} A_{ij} S_{ij}$$

If we have groups $G = \{G_1, G_2 \dots G_g\}$
(e.g. subject areas/tracks),
we define groupwise ESW:

Fair to subject areas

$$ESW(A,S,G) = \min_{G_k \in G} \sum_{i \in G_k} \sum_{j \in R} A_{ij} S_{ij}$$

Future Work

- Quantify integrality gaps
  - What is the impact of rounding?
- Closed-form solution for elliptical noise?
- More realistic experimental setup
  - Build document-similarity and bidding models with noise estimates
- Robust fairness
  - Envy term in the objective
  - Group-wise egalitarian welfare
- Other applications of robust allocation
  - Course allocation with inferred preferences, etc

Measuring Robustness

Identify allocation A_{S_0} produced by solving for optimal under S_0

Concentrate noise around pairs assigned by A_{S_0} with parameter $1 - \alpha$

Compare worst-case and random-case welfare of A_{S_0} with robust solution A_R

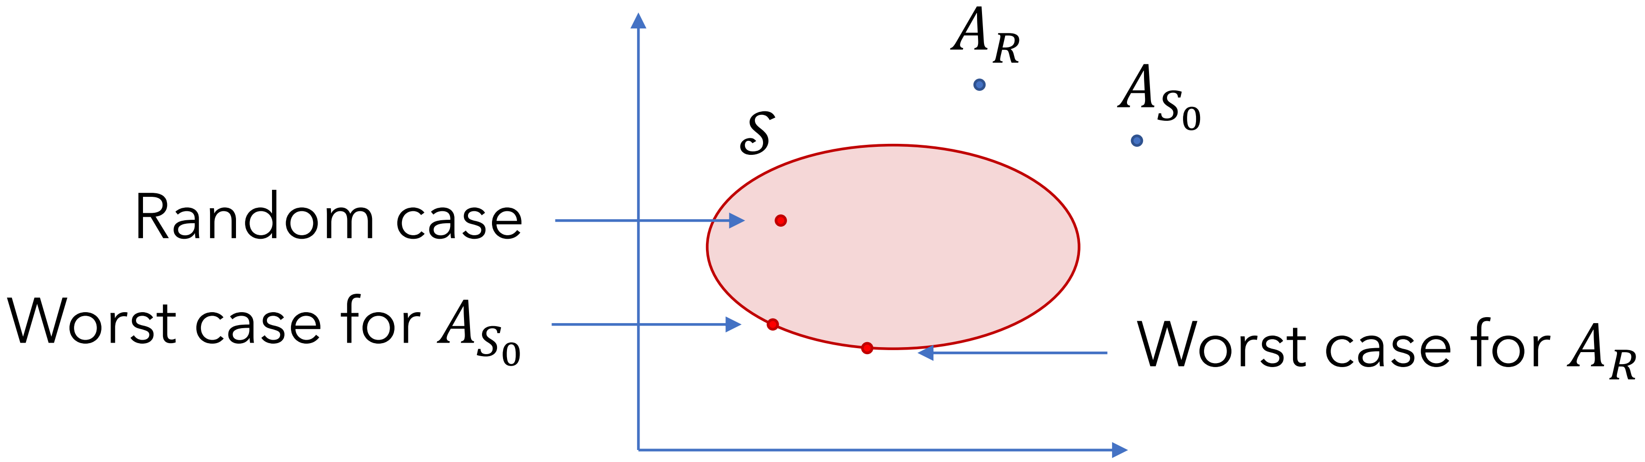

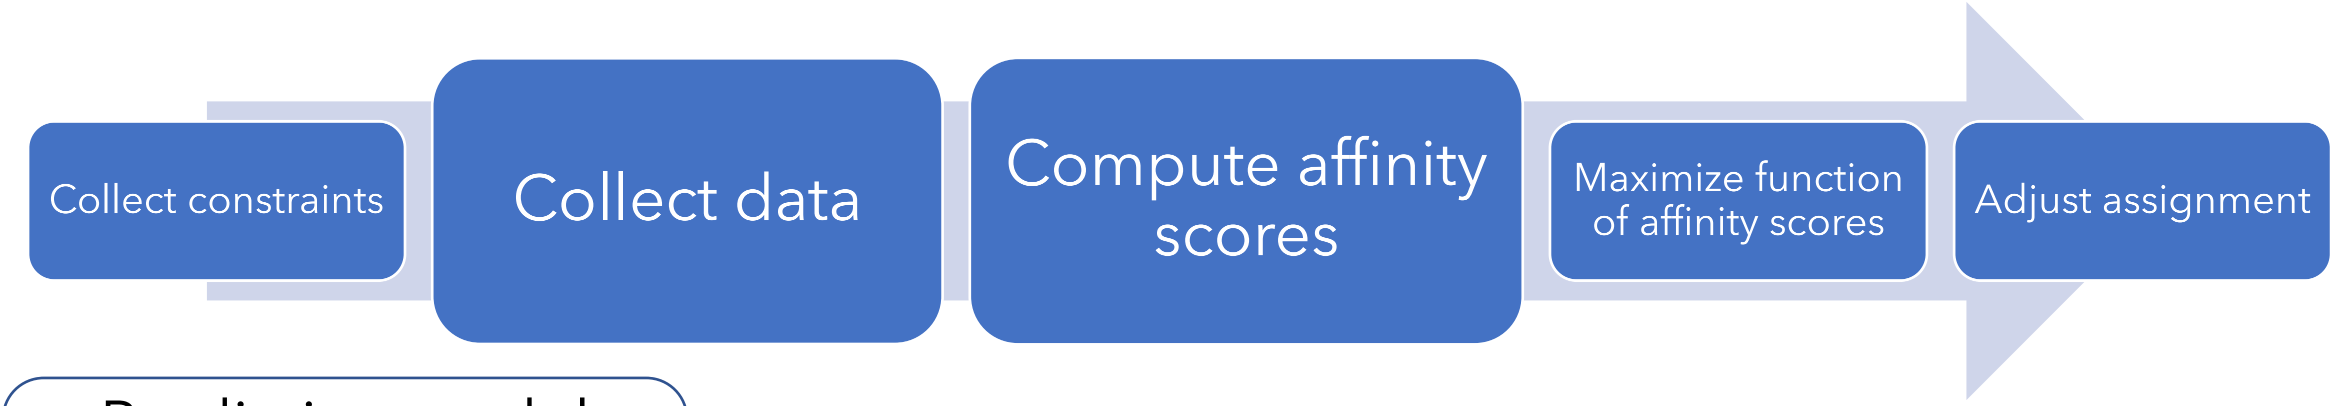

Predictive model based on historical performance

Measuring Robustness

Robust achieves higher
worst-case USW
and higher USW on randomly
sampled true scores

α controls concentration
of noise on pairs assigned
by point estimate method

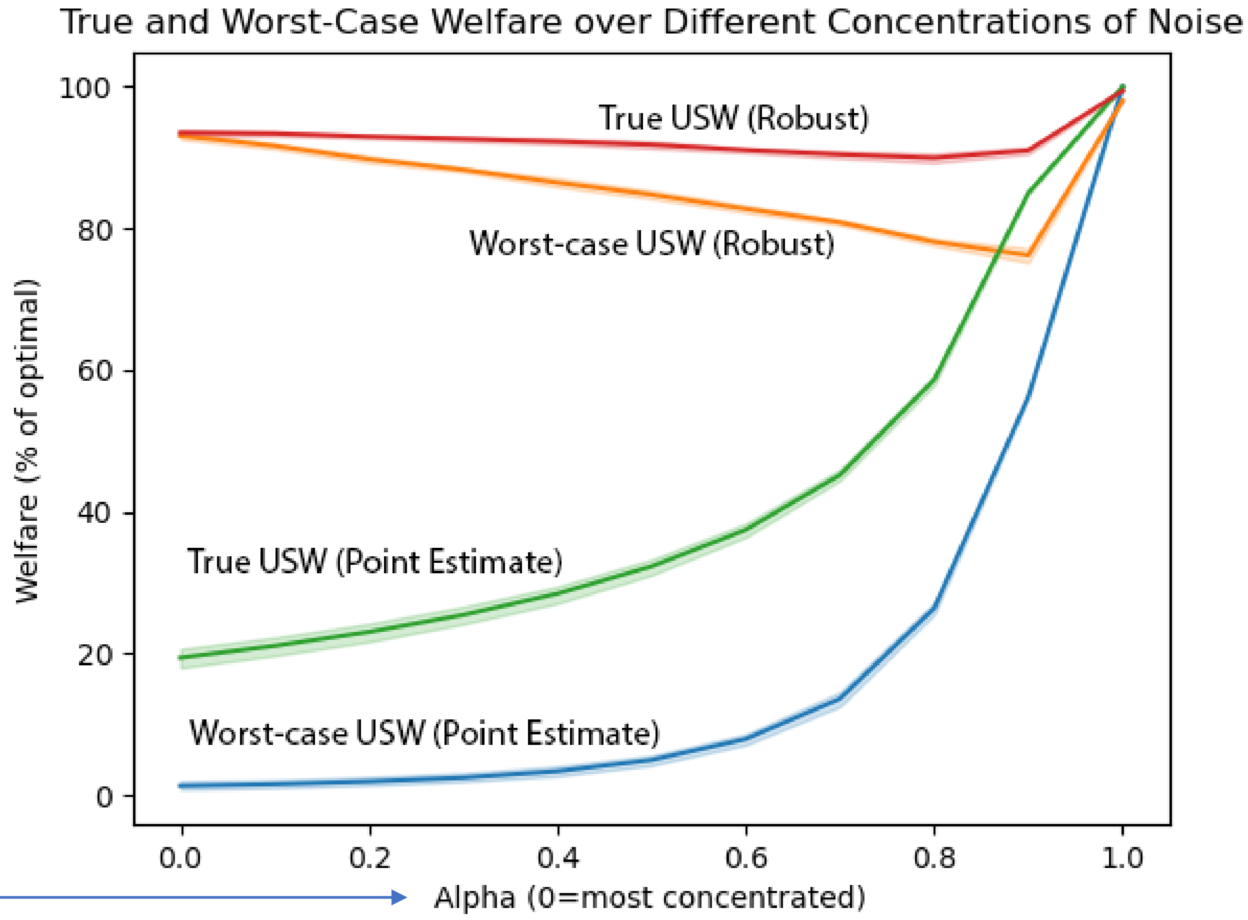

Computing the Subgradient

Initialize $A_1 = \max_{A \in \mathcal{A}} USW(A, S_0)$

For $t \in \{1, 2, \dots T\}$:

$S_t = \operatorname{argmin}_{S \in \mathcal{S}} USW(A_t, S)$

$A_{t+1} = A_t + \frac{1}{t+1} S_t$

$A_{t+1} = P_{\mathcal{A}}(A_{t+1})$

Return rounded A_T

Projection onto feasible set of allocations $\mathcal{A}$

Round to an integer, constraint-satisfying allocation[1]

1. Jecmen et al., NeurIPS 2020.