Fair Resource Allocation under Uncertainty


Cyrus Cousins


Elita Lobo


Justin Payan


Yair Zick

Fair & Explainable Decision-Making (FED) Lab

UMassAmherst | College of Information & Computer Sciences

Can we ensure fairness and efficiency when preferences are uncertain?

Resource Allocation under Uncertainty (RAU)

Given n agents N, partitioned into groups G, with m goods M
Additive, *unknown* utilities $U^* \in \mathcal{U}$ w.h.p.

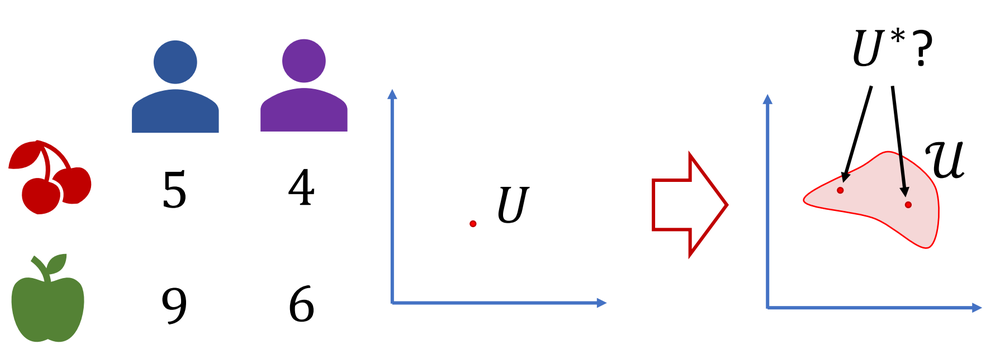

Select allocation $A \in \mathcal{A}$ ($\{0,1\}^{m \times n}$ + additional constraints)
such that welfare $W(A, U^*)$ is maximized

Objectives

Utilitarian Social Welfare (USW)

$$\max_{A \in \mathcal{A}} \sum_{i \in N, j \in M} A_{ij} U^*_{ij}$$

Group Egalitarian Social Welfare (G-ESW)

$$\max_{A \in \mathcal{A}} \min_{g \in G} \sum_{i \in g, j \in M} A_{ij} U^*_{ij}$$

Robust USW

$$\max_{A \in \mathcal{A}} \min_{U \in \mathcal{U}} \sum_{i \in N, j \in M} A_{ij} U_{ij}$$

Robust G-ESW

$$\max_{A \in \mathcal{A}} \min_{U \in \mathcal{U}} \min_{g \in G} \sum_{i \in g, j \in M} A_{ij} U_{ij}$$

Constructing Uncertainty Sets

Transductive Predictors
Construct $\hat{U}$ from samples of U^*

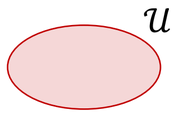

$$\mathcal{U} = \{U : (U - \hat{U})^T \Sigma^{-1} (U - \hat{U}) / nm \leq \hat{\xi} + \eta_t\}$$

Inductive Predictors
Construct $\hat{U}$ using past data

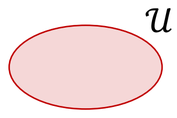

$$\mathcal{U} = \{U : (U - \hat{U})^T \Sigma^{-1} (U - \hat{U}) / nm \leq \hat{\xi} + \eta_i\}$$

Intersecting & Expanding Uncertainty Sets

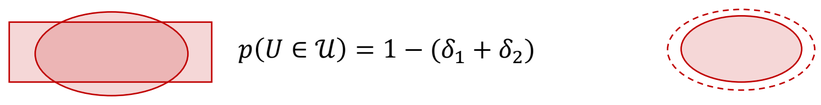

Solving RAU

RAU is NP-hard for USW under finite convex combinations of linear half-spaces (reduction from max egalitarian reviewer assignment)

Robust Resource Allocation (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 $\max_{\tilde{A} \in \tilde{\mathcal{A}}} \min_{U \in \mathcal{U}} W(\tilde{A}, U)$ by stepping in $\partial_{\tilde{A}} \min_{U \in \mathcal{U}} W(\tilde{A}, U)$ and projecting back to $\tilde{\mathcal{A}}$

Randomized rounding for discrete solution

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

Although RRA may have to round significantly, we obtain high probability lower bounds on $W(A, U^*)$

Case Study: Reviewer Assignment

RRA has strong worst-case USW on recent ICLR's

Year	m	n	Adversarial USW · 100					Average USW · 100				
			LP	FF	PR4A	FS	RRA	LP	FF	PR4A	FS	RRA
2018	1657	546	**17**±3	7±3	**17**±3	16±3	16±3	**179**±2	134±12	177±2	177±2	160±4
2019	2620	851	22±2	12±2	22±2	22±2	**27**±3	**184**±1	139±9	**184**±1	183±1	161±3
2020	4123	1327	17±2	11±2	18±2	17±2	**23**±2	**187**±1	158±8	**187**±1	186±1	166±5
2021	4662	1557	23±2	18±2	23±2	23±2	**33**±3	**192**±1	177±2	**192**±1	191±1	174±6
2022	5023	1576	28±2	23±2	28±2	27±2	**38**±2	**191**±1	177±1	190±1	190±1	172±3

RRA also improves "true" USW and G-ESW for ICLR

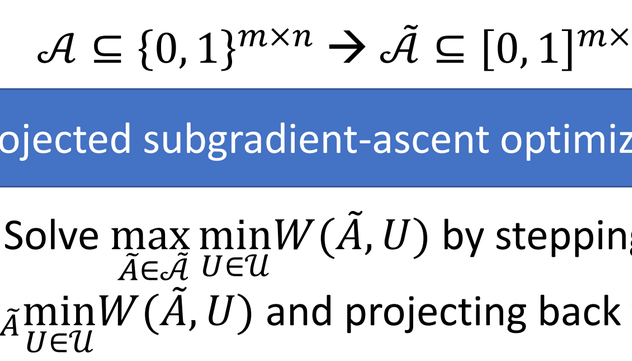

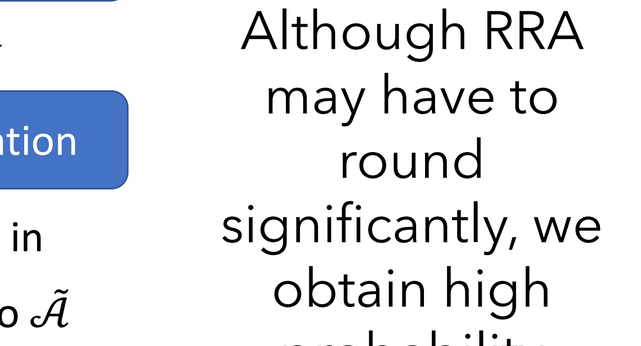

In presence of noisily-estimated papers/reviewers, RRA performs better on hidden "true" USW

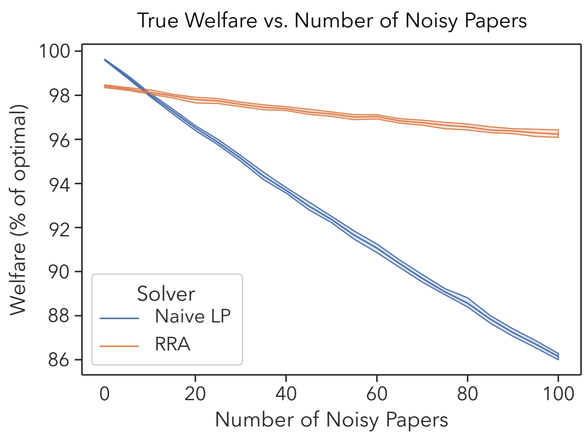

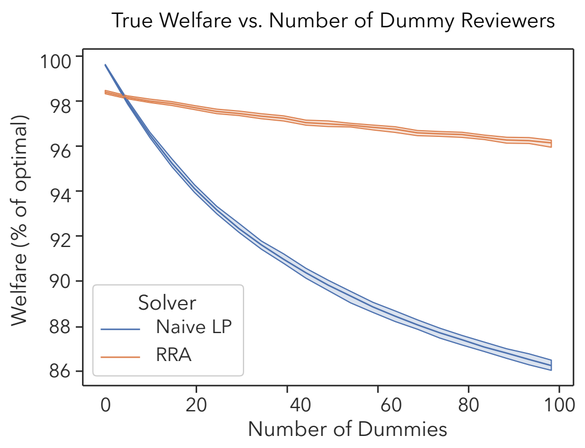

Read it on arxiv!