Fairness in Preference Queries:
Social Choice Theories Meet Data Management

Presenters:

Senjuti Basu Roy*, Baruch Schieber*, Nimrod Talmon^

*NJIT, ^Ben Gurion University

Slides available at: https://centers.njit.edu/bdal

Preference Queries

Query: Which food is the most preferred choice?

User	Preference
Alex	 >  >  >  > 
Bob	 >  >  >  > 
Jill	 >  >  >  > 
Jane	 >  >  >  > 
Lia	 >  >  >  > 

Answer 

Applications

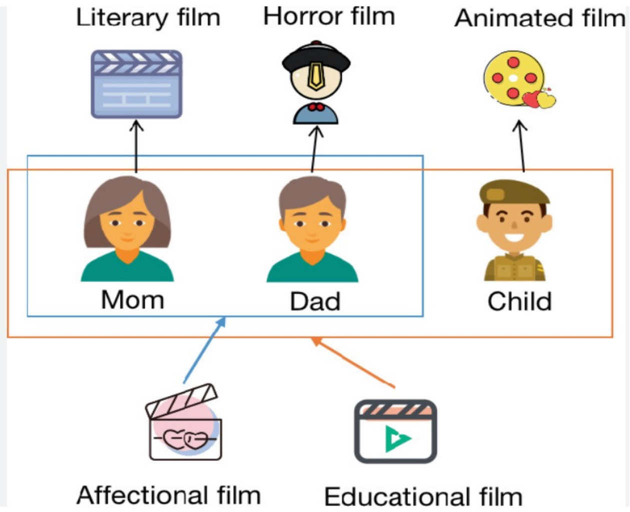

Group Recommendations

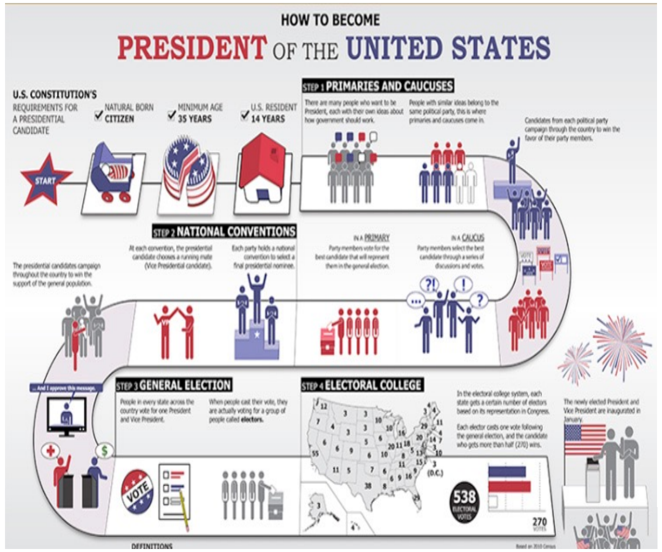

Electoral Systems

Hiring and recruitment

Admission

Some More Applications

Sports competition
Candidates shortlisting

Parliamentary elections
Capacitated facility location

Company portfolio
Movies selection
Facility location

Preference Queries in Database Literature

Application	Authors		
Querying	Lacroix and Lavency 1987; Kießling and G ¨untzer 1994; K¨ostler et al. 1995, Chomicki 2002; Kießling 2002; Kießling and Hafenrichter 2002; Kießling and K¨ostler 2002	Qualitative	the preferences between tuples in the answer to a query are specified directly, typically using binary preference relations
		Quantitative	preferences are specified indirectly using scoring functions that associate a numeric score with every tuple of the query answer.
Group Recommendation	Amer-yahia et. al., 2009; basu roy et. al. 2015, Yuan et. al. 2014,	Quantitative	Preferences are specified using scoring functions such as aggregated voting or least misery that associate a numeric score with every recommended item.

Notion of ranking is implicit!

Overview

- Part I – Preference aggregation methods (30 minutes)
- Part II - Fairness in answering preference queries (30 minutes)
- Part III - Future research directions (20 minutes)

Part I – Preference Aggregation Method
(30 min)

Input Formats – Preference Elicitation

VLDB2024
GUANGZHOU

Plurality

(declare most loved one)

a

Approval

(declare set of good ones)

[1,0,1]

Scoring

(score each option)

[3, 1, 5]

Ordinal

(order by preference)

a>b>c

Cumulative

(distribute a token)

[0.5, 0.2, 0.3]

Big Data Analytics Lab

Ordinal Voter Preferences

VLDB2024
GUANGZHOU

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

Big Data Analytics Lab

A *profile* is composed of a set of *candidates/alternatives* and a collection of *voters*

General Setting

Given inputs (preference elicitation), an aggregation method (*voting rule)* outputs a single or multiple winner(s)

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

Preference elicitation(Input)

Any ideas?

Aggregation method (Voting rule)

Winners

Output

Famous Aggregation Methods

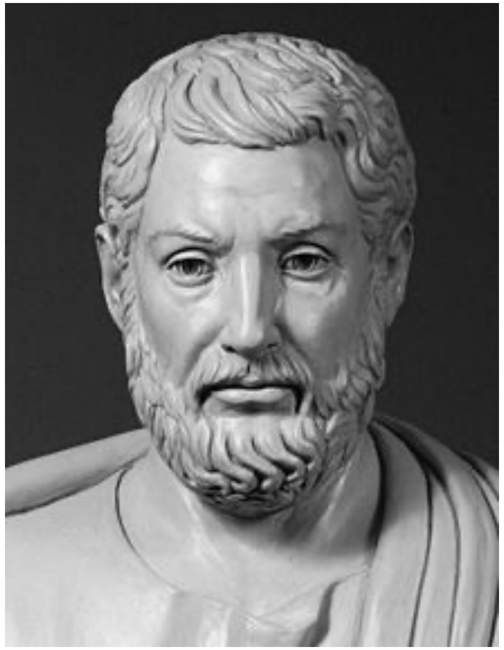

Plurality / Sortition (~-500)

Llull (1299)

Condorcet (1785)

Borda (1770)

Dogdson (1876)

Shulze (1997)

Computational Social Choice

Specific Aggregation Methods

Given an election, a voting rule outputs a single winner

Plurality: count 1^{st} positions

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

s(a) = 0
s(b) = 1
s(c) = 2
s(d) = 1
s(e) = 1

1 0 0 0 0

Scoring vector

Specific Aggregation Methods

Given an election, a voting rule outputs a single winner

Plurality: count 1^{st} positions
Borda: count (Borda) score

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

4 3 2 1 0

Scoring vector

s(a) = 10
s(b) = 14
s(c) = 11
s(d) = 8
s(e) = 7

Specific Aggregation Methods

Given an election, a voting rule outputs a single winner

Plurality: count 1st positions
Borda: count (Borda) score
Llull: count pairwise wins

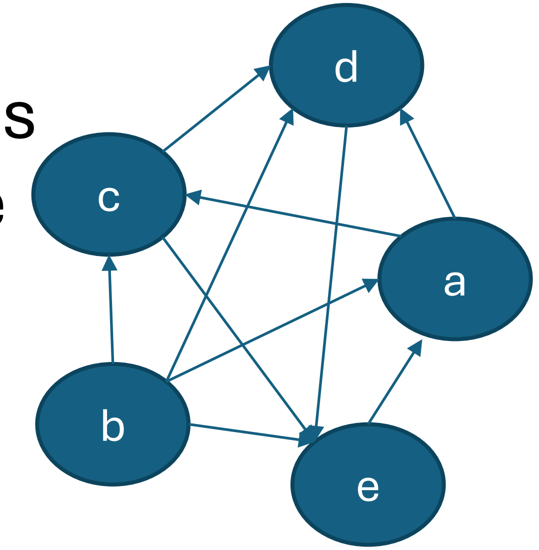

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

Specific Aggregation Methods

Given an election, a voting rule outputs a single winner

Plurality: count 1st positions
Borda: count (Borda) score
Llull: count pairwise wins
IRV: eliminate until majority

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

And Many More..

How to make a choice on how to make a choice?

How to evaluate different voting rules?

How to make a choice on how to make a choice?

Optimization

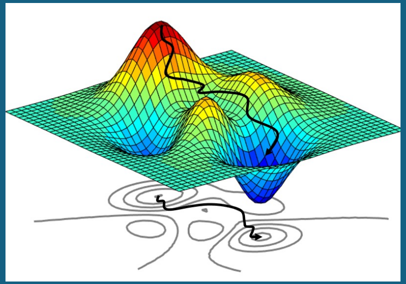

Defining optimization goals

How to make a choice on how to make a choice?

Optimization

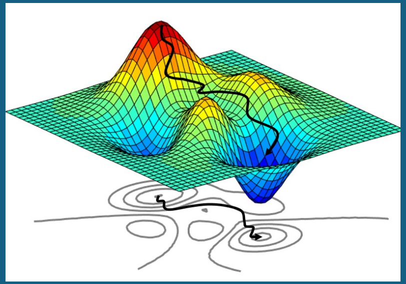

Defining optimization goals

Axioms

Formulating good properties

“My dream algorithm: proportional, discount monotone, etc…

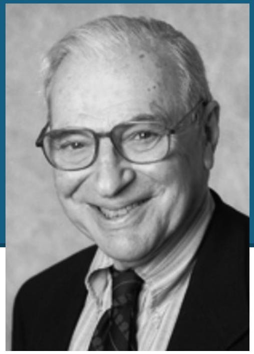

How to make a choice on how to make a choice?

Optimization

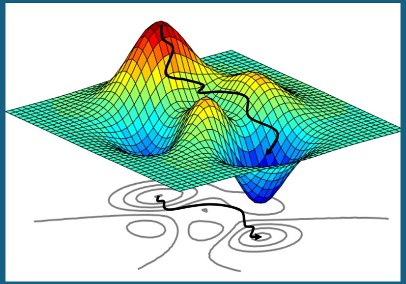

Defining optimization goals

Axioms

Formulating good properties

“My dream algorithm: proportional, discount monotone, etc…

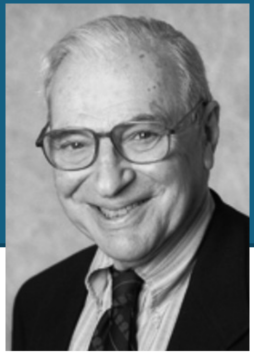

Simulations & Experiments

Does it look good?

How to make a choice on how to make a choice?

Optimization

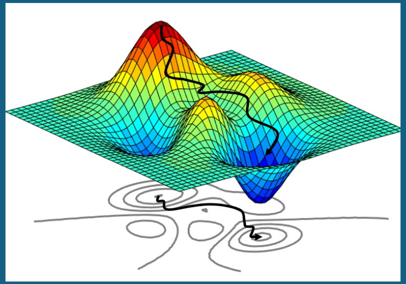

Defining optimization goals

Axioms

Formulating good properties

“My dream algorithm: proportional, discount monotone, etc…

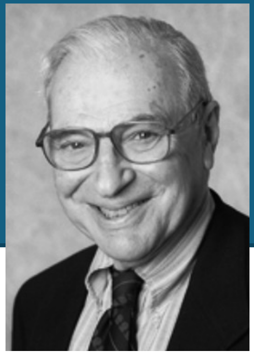

Simulations & Experiments

Does it look good?

Complexity and Xplainability

Ease of use..

The Axiomatic Way

Let’s formulate axioms expressing desirable (fairness) properties

The Axiomatic Way

Let's formulate axioms expressing desirable (fairness) properties

Anonymity

voter names do not matter

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

Voter	Ordinal Ballot
v_1	e > a > b > d > c
v_2	b > a > c > d > e
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

The Axiomatic Way

Let's formulate axioms expressing desirable (fairness) properties

Neutrality

candidate names do not matter

Voter	Ordinal Ballot
v_1	b > a > c > d > e
v_2	e > a > b > d > c
v_3	d > a > b > c > e
v_4	c > b > d > e > a
v_5	c > b > e > a > d

Voter	Ordinal Ballot
v_1	**a** > **b** > c > d > e
v_2	e > **b** > **a** > d > c
v_3	d > **b** > **a** > c > e
v_4	c > **a** > d > e > **b**
v_5	c > **a** > e > **b** > d

The Axiomatic Way

Let’s formulate axioms expressing desirable (fairness) properties

Positive Responsiveness (Monotonicity; campaigns don’t hurt)

raised (tied) winner is winner

b

Voter	Ordinal ballots
v_1	a > b > c
v_2	a > b > c
v_3	b > a > c
v_4	b > a > c

b

Voter	Ordinal ballots
v_1	b > a > c
v_2	b > a > c
v_3	b > a > c
v_4	b > a > c

~~IRV~~

Characterization Results

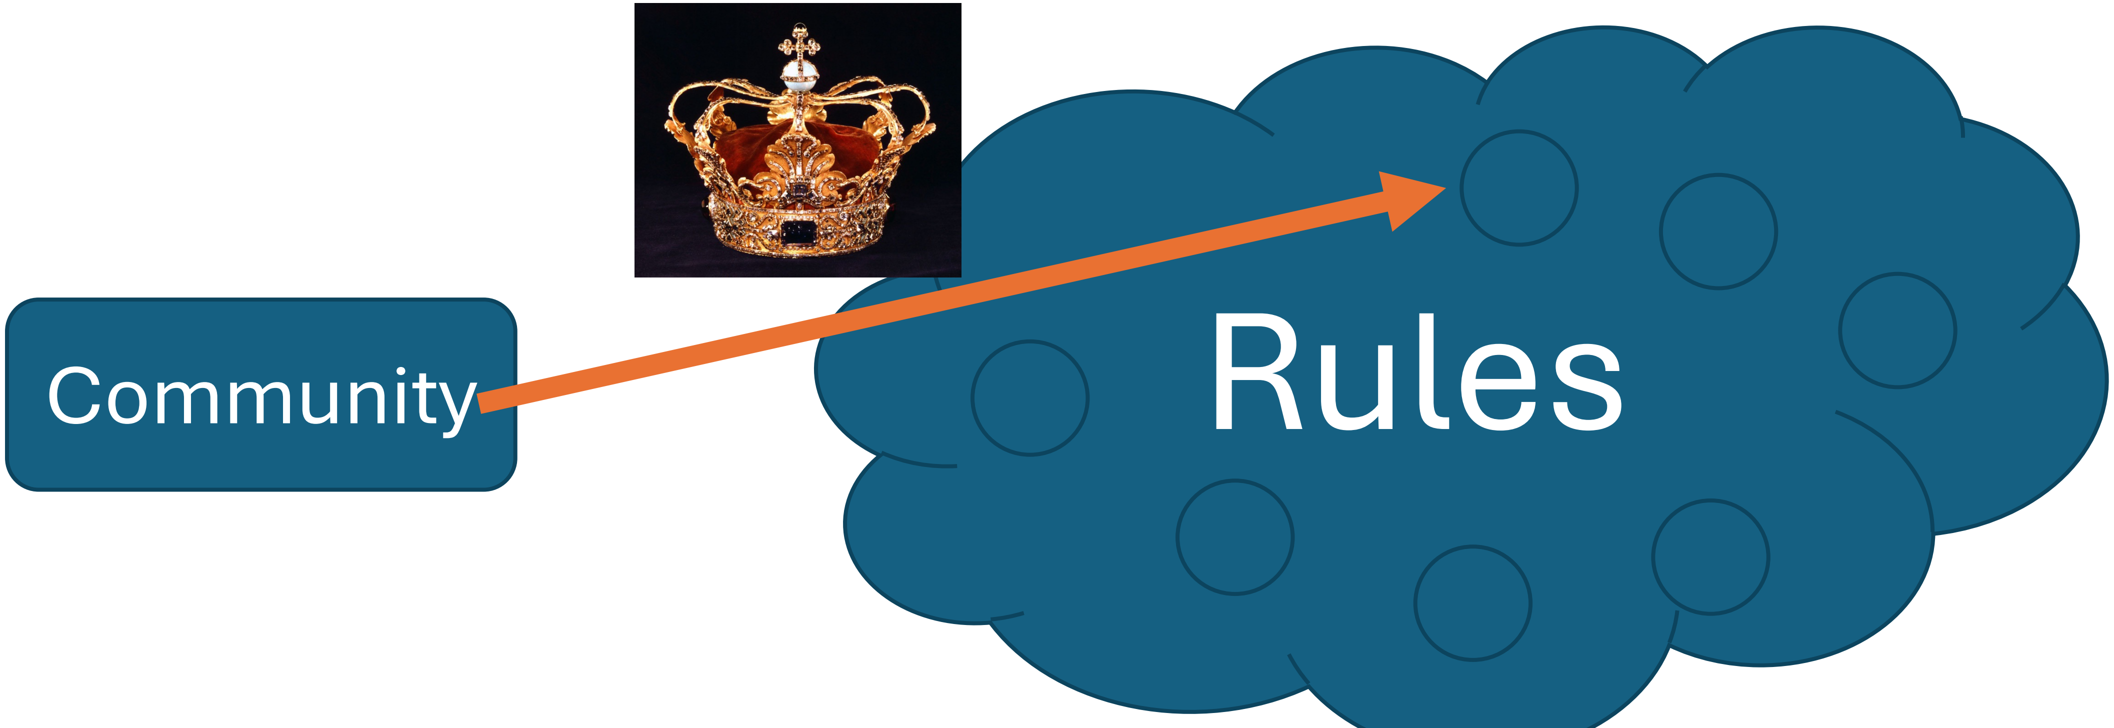

The ONLY rule that satisfies {a, b, c}

Mathematical Verdicts

May's Theorem (1952):

A voting rule R for two alternatives satisfies:

anonymity, neutrality, and positive responsiveness

if and only if R is ???

simple majority

Mathematical Verdicts

Arrow's Impossibility Theorem (1951): No rule satisfies Pareto, IIA, and non-dictatorship

(<u>Pareto</u> – **a** dominates **b** if all voters prefer **a** to **b**; rule must select a non-dominated winner)
(<u>IIA</u> – if **a** wins over **b** when **c** is considered then **a** wins over **b** also when **c** is not considered)
(<u>Non-dictatorship</u> – not one voter dictates the output)

Voter	Ordinal ballots
v_1	a > b > c
v_2	b > c > a
v_3	c > a > b

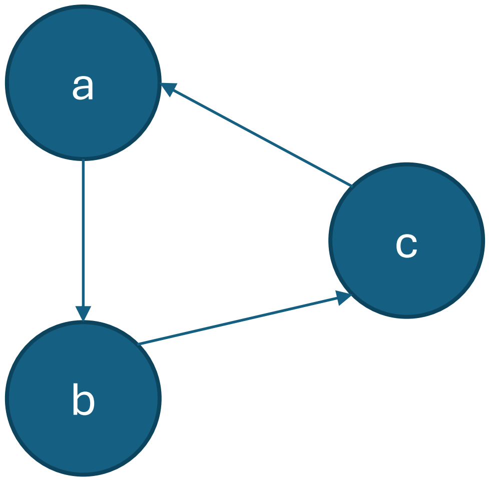

Mathematical Verdicts

GS Theorem (1973-5): No rule satisfies strategyproofness and non-dictatorship

(strategyproof – never beneficial to do strategic voting [i.e., to “lie”])

Voter	Ordinal ballots
v_1	a > b > c
v_2	b > c > a
v_3	c > a > b

Big Data Analytics Lab

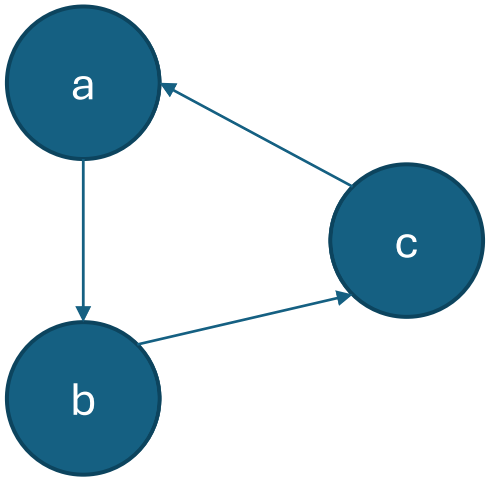

Mathematical Verdicts

Arrow's Theorem (1951): No sane rule?

GS Theorem (1973-5): No strategyproof rule?

Not all is lost

Quantitative axioms!

Randomization!

Domain restrictions!

Mathematical Verdicts

Black’s Theorem (1948): Median is <u>strategyproof</u> for singlepeaked domains

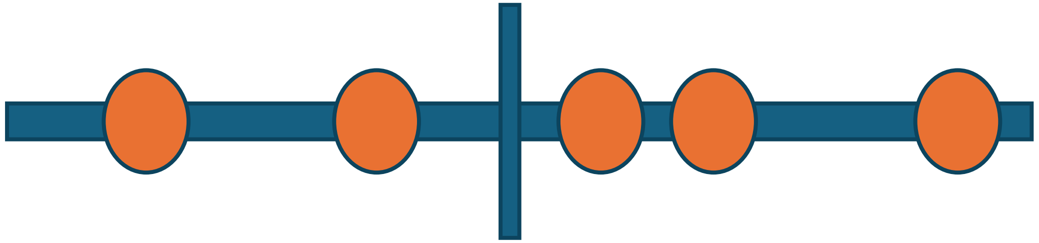

Center > Left > Right

~~Left > Right > Center~~

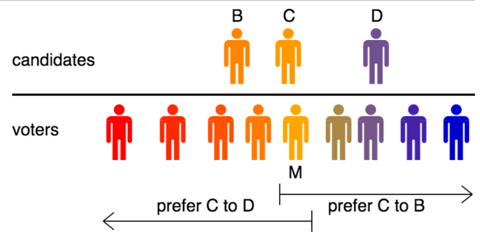

Output Formats – Type of Winner(s)

- Not just a single winner!

Single-winner → Multiwinner (also ranked) → Participatory budgeting (weighted Multiwinner)

General decisions → text, networks, allocations, etc…

Further Topics

Topic	Paper	Essence
Proportional ranking	https://arxiv.org/pdf/1612.01434	Aggregate rankings to a single ranking while respecting minority opinion at all prefixes
Movie search via voting	https://arxiv.org/pdf/2202.03385	Using proportional multiwinner rules for group recommendation
Music recommendation by vote delegation	https://arxiv.org/pdf/1503.08604	Using dampened delegations for group recommendation
Clustering via social choice	https://arxiv.org/pdf/2310.18162	Being fair to the clusters by voting
Applicable fair division	Spliddit.org	Battle-tested fair division algorithms
Applicable soritition	Panelot.org	Battle-tested sortition app

Summary:
Single-Winner Voting Aggregation Methods

Aggregation method	Input format	Definition	Properties
Plurality	1-approval	Top-ranked candidate	Simple
Borda	ranking	Linear points	Average-best candidate
Approval	approval	Most approved	Avoids Arrow
IRV	Ranking (usually weak)	Eliminating and transferring until majority	Avoids certain strategic behaviors
Llull	ranking	Count pairwise wins	Picking Condorcet winner (candidate not losing in any pairwise contest)

Summary: Other Aggregation Methods

Aggregation method	Setting	Essence	Axiomatic properties	Computational properties
Kemeny	Social welfare function	Swap-closest ranking (centroid)		NP-hard (but efficient)
SNTV	Multiwinner	Top-ranked candidates	Not so great	P
k-Borda	Multiwinner	Average-ranked candidates	Picking candidates that are average-good	P
STV	Multiwinner	Eliminating and transferring until k seats	Picking a representative committee	P
Chamberlin-Courant	Multiwinner	Diverse committee	Most voters have at least one candidate they like	NP-hard (but theoretically-efficient)
Proportional approval voting (PAV)	Multiwinner	Representative minority-respecting committee	Picking a representative committee	NP-hard (but efficient)

Part II – Existing Research on Fairness in Preference Queries (30 min)

Algorithmic Fairness

Algorithmic fairness interprets fairness as lack of discrimination asking that an algorithm should not discriminate entities based on attributes that are not relevant to the task at hand. Such attributes are called **protected or sensitive**, and often include among others gender, religion, age, sexual orientation and race.

COMPAS in Judicial Decision Making

WISCONSIN.GOV Agency Directory Online Services

State of Wisconsin
Department of Corrections

VICTIM SERVICES | OFFENDER INFORMATION | CAREERS | ABOUT DOC | NEWS | CONTACT US

Search Department of Corrections

DOC > COMPAS

COMPAS

DOC uses the Correctional Offender Management Profiling for Alternative Sanctions tool, commonly known as COMPAS, for criminogenic risk and needs assessments and unified case planning. This actuarial risk assessment system contains offender information specifically designed to determine their risk and needs and inform dynamic case plans that will guide the offender throughout his or her lifecycle in the criminal justice system.

The lifecycle is a framework for how an offender moves through the Wisconsin criminal justice system and the decision points informed by COMPAS along the way. From the time of arrest through eventual discharge, DOC will use the lifecycle as a framework for establishing meaningful practices and interventions across jurisdictions.

Evidence Based Practices clearly state that having a sound assessment that accurately identifies an offender's risk to reoffend is the cornerstone of effective supervision. Without a proper assessment, appropriate interventions and services cannot be delivered. Services are targeted for moderate to higher-risk offenders who are likely to reoffend if appropriate interventions are not available.

Research shows low risk offenders are less likely to commit new crimes and should be given minimal services and be excluded from intensive treatment and supervision.

COMPAS Case Study: Investigating Algorithmic Fairness of Predictive Policing

Mallika Chawla · Follow
7 min read · Feb 23, 2022

44

The defendant, Eric Loomis, had been assessed by COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) as a high risk individual and was consequently sentenced to eight years in prison — a ruling that he challenged as a violation of his due process rights.

Responsible Data Management Landscape ----

Task	Work	What it does
Score based ranking	Measuring fairness in ranked outputs [Yang. et. al.]	proportional representation fairness within the NDCG framework, imposing proportionality constraint over every prefix of the ranking and accounting for position bias with a logarithmic discount.
Score based ranking	On obtaining stable rankings [Asudeh. et. al]	Proposes a framework that can be used to assess the stability of a provided ranking and to obtain a stable ranking within an "acceptable" range of weight values (called "the region of interest")
Score based ranking	Designing Fair Ranking Schemes [Asudeh et. al.]	considers ranking functions that compute the score of each item as a weighted sum of (numeric) attribute values, and then sort items on their score. Each ranking function can be expressed as a point in a multidimensional space. For multiple fairness criteria, including proportionality, the work shows how to efficiently identify regions in this space that satisfy these criteria.
Range Queries	Fairness-Aware Range Queries for Selecting Unbiased Data [Shetiya et. al]	Query reformulation - finding the most similar fair range to a user-provided range query for the database to satisfy group fairness defined over a single binary protected attribute

Responsible Data Management Landscape ….

Task	Work	What it does
Link analysis	Fairness-Aware PageRank [Tsioutsiouliklis et. al.]	Revisits Pagerank through the lens of fairness definition and study variants that are fair. The work defines the utility loss of a fair algorithm as the difference between its output and the output of the Pagerank algorithm, and it considers the problem of achieving fairness while minimizing utility loss.
Group recommendation	Fairness in Package-to-Group Recommendations [serbos et. al.]	Studies fairness aware package-to-group recommendations, that of fairness with the goal that every group member is satisfied by a sufficient number of items in the package.
Database repair	Interventional Fairness : Causal Database Repair for Algorithmic Fairness [Salimi et. al]	describes an approach to removing discrimination by repairing the training data in order to remove the effect of any discriminatory causal relationship between the protected attribute and classifier predictions, without assuming adherence to an underlying causal models.
Data structure design	FairHash: A Fair and Memory/Time-efficient Hashmap [Shahbazi et. al]	a data- dependant hashmap that guarantees uniform distribution at the group-level across hash buckets, and hence, satisfies the statistical parity notion of group fairness.

But THIS tutorial studies fairness considering Different and Multiple preferences with the goal of producing a single output ....

Terminology

- Users(Voters) input preferences over items/candidates
- Input preferences are elicited differently
- Preference aggregation model takes inputs and generates output
  - Output is a
    - single item/candidate
    - a set of k items/candidates
    - a ranked order of items/candidates

A Possible Taxonomy of Fairness

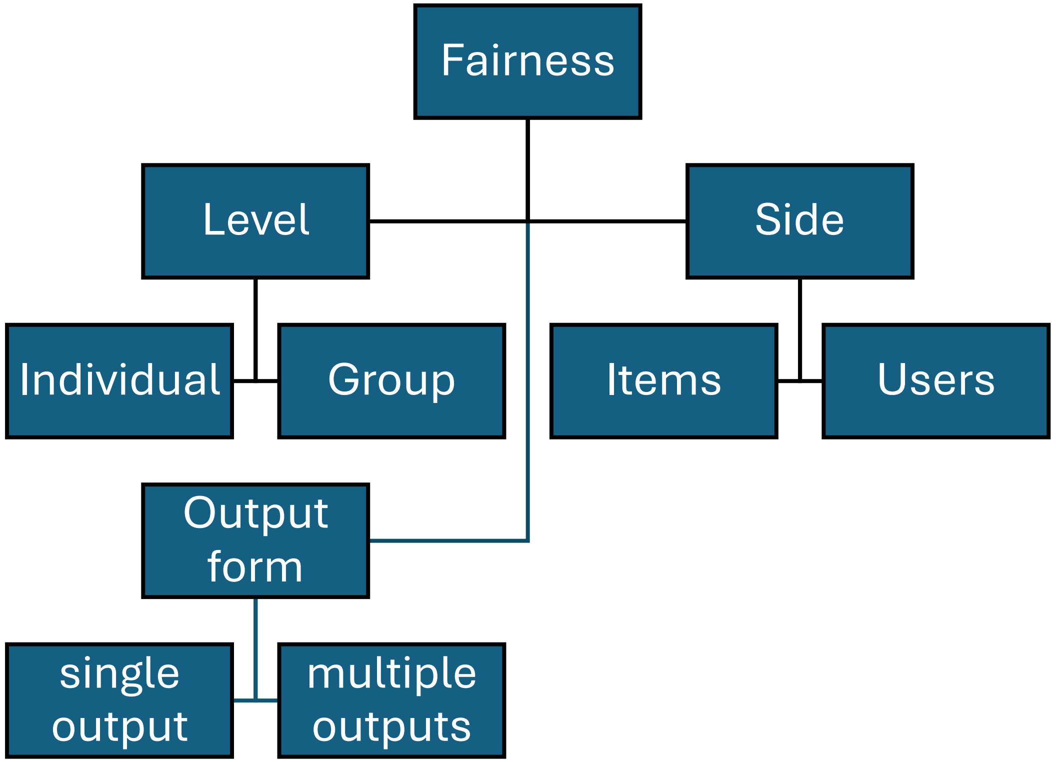

Type	Definition	
Level Based	Individual	Group
	Individuals with “similar” utility must have similar chance of getting favorable outcome	Members from different protected attribute groups must have equal chance of getting favorable outcome
Side Based	Items	User
	similar items receive similar favorable outcomes	Focuses on returning items to the users that express community interests and do not favor any specific user/ user group
Output forms	Single output	Multiple outputs
	One output (set/ordered) is produced to enable fairness	Multiple outputs are produced to achieve fairness

Fairness in Preference Queries

Fairness	Applicable to	Definition
Top-k Parity	Group Fairness	Proportionate representation of every protected attribute in top-k
p-fairness	Group Fairness	Proportionate representation of every protected attribute in **every prefix** of top-k
Probability-based fairness	Group fairness	defined by means of statistical significance tests that ask how likely it is that a given ranking was created by a fair process, such as by tossing a coin to decide whether to put a protected-group or a privileged-group item at position i
m-proportionality	Individual fairness	for what proportion of the users the output satisfies the following: there exists at least m items in the output that each of those users like.
m-envyfreeness	Individual fairness	for what proportion of the users the output satisfies the following: for every such user, there exists at least m items such that each item is in the top-δ% of their preferences, for an input parameter δ.
Dissatisfaction fairness	Individual Fairness	Minimize maximum dissatisfaction across all the users in the community
Uniform selection probability	Individual Fairness	Items with similar "utility" have "equivalent" chance of exposure
N level pareto optimal	Individual Fairness	An output with N items that are dominated by at most N − 1 other items. An item j dominates another item i if all users rank j higher than i.

Big Data Analytics Lab

Pitoura, Evaggelia, Kostas Stefanidis, and Georgia Koutrika. "Fairness in rankings and recommendations: an overview." *The VLDB Journal* (2022): 1-28.
Islam, Md Mouinul, Mahsa Asadi, and Senjuti Basu Roy. "Equitable top-k results for long tail data." *Proceedings of the ACM on Management of Data* 1.4 (2023): 1-24.
Wei, Dong, et al. "Rank aggregation with proportionate fairness." *Proceedings of the 2022 international conference on management of data*. 2022.

Computational Implications come from ……….

- Aggregation Method
  - Plurality voting
  - Rank aggregation
  - Multi round voting
- Group Structure
  - Cardinality of protected attributes.
    - Binary
    - Multi-valued
  - Number of protected attributes
    - 1
    - More than 1
- How to achieve fairness?
  - Change input
    - Constructive margin of victory
      - find minimal manipulation of inputs that guarantees fairness
  - Change output
    - May not reflect the community opinion well

Three interspersed dimensions

Blom, Michelle, Peter J. Stuckey, and Vanessa J. Teague. "Toward computing the margin of victory in single transferable vote elections." *INFORMS Journal on Computing* 31.4 (2019): 636-653.

How to Achieve Fairness

- Achieve fairness by
  - (Minimally) Changing input preferences
    - Margin of victory or MOV (preserves original input preferences in the most truthful way)
  - (Minimally) changing output
    - May not preserve original input preferences truthfully
  - Design (new) aggregation model
    - Negative results from Arrow's Theorem seem to indicate no aggregation functions satisfy even simple axioms related to fairness

Rank Aggregation with Proportionate Fairness (SIGMOD 2022)

AUTHORS:
Dong Wei, Md Mouinul Islam, Baruch Schieber, Senjuti Basu Roy

Fair Rank Aggregation (NeurIPS 2022)

AUTHORS:
D. Chakraborty, S. Das, A. Khan, A. Subramanian

Rank aggregation of search engine outputs

- Dwork, Kumar, Naor, Sivakumar, “Rank aggregation methods for the web”, WWW, 2001.
  - Q: How can search-engine bias be overcome?
  - A: By combining results from multiple search engines

Search: Waterloo

Google
1. Wikipedia: Battle of Waterloo
2. Wikipedia: Waterloo, ON
3. www.city.waterloo.on.ca (City of Waterloo website)
4. www.uwaterloo.ca (University of Waterloo)
5. www.waterlooindustries.com

Yahoo!
1. www.uwaterloo.ca
2. Wikipedia: Battle of Waterloo
3. www.city.waterloo.on.ca
4. Wikipedia: Waterloo, ON
5. www.waterloorecords.com

Big Data Analytics Lab

Minimize the number of “disagreements”

Google
1. Wikipedia: Battle of Waterloo
2. Wikipedia: Waterloo, ON
3. www.city.waterloo.on.ca
4. www.uwaterloo.ca

Aggregate ranking
1. www.uwaterloo.ca (4)
2. Wikipedia: Battle of Waterloo (1)
3. Wikipedia: Waterloo, ON (2)
4. www.city.waterloo.on.ca (3)

Yahoo!
1. www.uwaterloo.ca (4)
2. Wikipedia: Battle of Waterloo (1)
3. www.city.waterloo.on.ca (3)
4. Wikipedia: Waterloo, ON (2)

Rank Aggregation Model

Kemeny Distance. *For rankings $\rho_1, \rho_2, \ldots, \rho_m$, the Kemeny Distance of the ranking σ to these rankings is*

$$\kappa(\sigma, \rho_1, \rho_2, \ldots, \rho_m) = \sum_{i=1}^{m} K(\sigma, \rho_i)$$

$K(\sigma, \rho_i)$ = *Kendall Tau distance between σ, ρ_i*

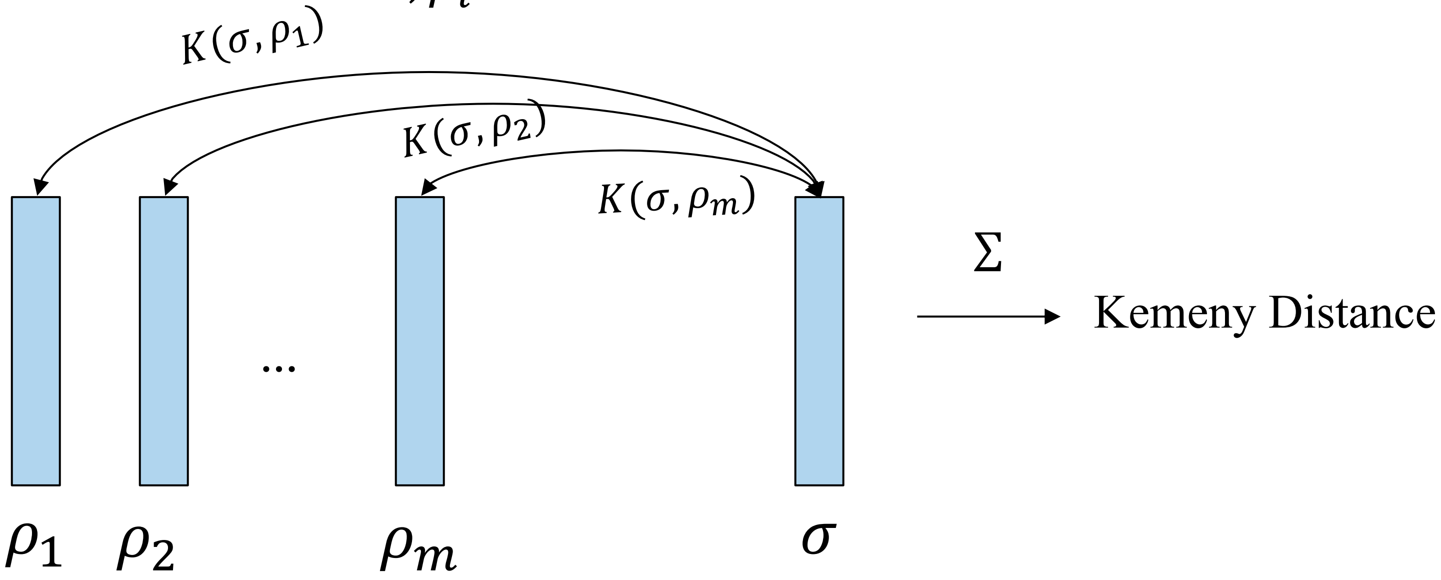

[4] Caitlin Kuhlman and Elke Rundensteiner. 2020. Rank aggregation algorithms for fair consensus. PVLDB 2020

Why Kemeny Distance

- Kemeny distance is a maximum likelihood estimator of the true preference order
- Simultaneously satisfies
  - Neutrality
    - $\kappa(\sigma_x, \rho_1, \rho_2, ..., \rho_m) = \kappa(\rho_1, \rho_2, ..., \rho_m, \sigma_x)$
  - Consistency
    - $\kappa(\sigma, \rho_1) + \kappa(\sigma, \rho_2) = \kappa(\sigma, (\rho_1, \rho_2))$
  - Condorcet winner property
    - If there is a candidate that wins majority of the preference, then that candidate wins
- Metric distance satisfies the triangle inequality
  - $K(\sigma, \eta) \leq K(\sigma, \eta') + K(\eta', \eta)$
- Has been studied extensively in rank aggregation literature [9,10]

[9] Dwork, Cynthia, et al. "Rank aggregation methods for the web." *Proceedings of the 10th international conference on World Wide Web*. 2001.
[10] Ailon, Nir, Moses Charikar, and Alantha Newman. "Aggregating inconsistent information: ranking and clustering." *Journal of the ACM (JACM)* 55.5 (2008): 1-27.

Fairness based on Group Protected Attributes

Rank	Protected Attribute
Amy	Female
Molly	Female
Abigail	Female
Kim	Male
Lee	Male
Park	Male
.	.
.	.

Binary Protected Attribute : Gender

Rank	Protected Attribute
Amy	Caucasian
Park	Asian
Molly	Caucasian
Kabir	Asian
Abigail	Native Hawaiian
Lee	Asian
.	.
.	.

Multi-valued Protected Attribute : Race

Proportionate Fair or p-fair [1,5] Rank

- p = *protected attribute value*
- $f(p)$ = *fraction of items with* p
- *for every* $k \in [1..n]$, *the number of items with protected attribute value* p *is* $\lfloor f(p) \cdot k \rfloor$ *or* $\lceil f(p) \cdot k \rceil$

F(female) = 50%

k	Min	Max
1	0	1
2	1	1
3	1	2

Aggregated Rank	Protected Attribute
Amy	Female
Park	Male
Molly	Female

K=1
Female = 1
K=2
Female = 1
K=3
Female = 2

P-fair Rank

Applications

- Worker selection, Faculty hiring, School and college admission, medical residency

[1]. Sanjoy K Baruah, et.al. 1996. Proportionate progress: A notion of fairness in resource allocation. Algorithmica.
[5]. Robert Tijdeman. 1980. The chairman assignment problem. Discrete Mathematics 32, 3 (1980), 323–330.

p-fair Ranking

Proportionate representation in *every position*

50% male, 50% female

Rank	Protected Attribute
Amy	Female
Molly	Female
Abigail	Female
Kim	Male
Lee	Male
Park	Male
.	.
.	.

- 1 male, 1 female in top 2 position ✗
- 2 male, 2 female in top 4 position ✗
- 3 male, 3 female in top 6 position ✓

Not p-fair

Rank	Protected Attribute
Amy	Female
Park	Male
Molly	Female
Kabir	Male
Abigail	Female
Lee	Male
.	.
.	.

- 1 male, 1 female in top 2 positions ✓
- 2 male, 2 female in top 4 positions ✓
- 3 male, 3 female in top 6 positions ✓

p-fair

Problem Definitions

- **Individual p-fair rank** (or IPF). *Given a ranking ρ find a p-fair ranking that is closest to ρ in Kendall-Tau distance*

- **Rank aggregation under p-fairness** (or RAPF). *Given m rankings $\rho_1, \rho_2, \ldots, \rho_m$ find a p-fair ranking that minimizes the Kemeny distance to these m rankings*

Studied Problems

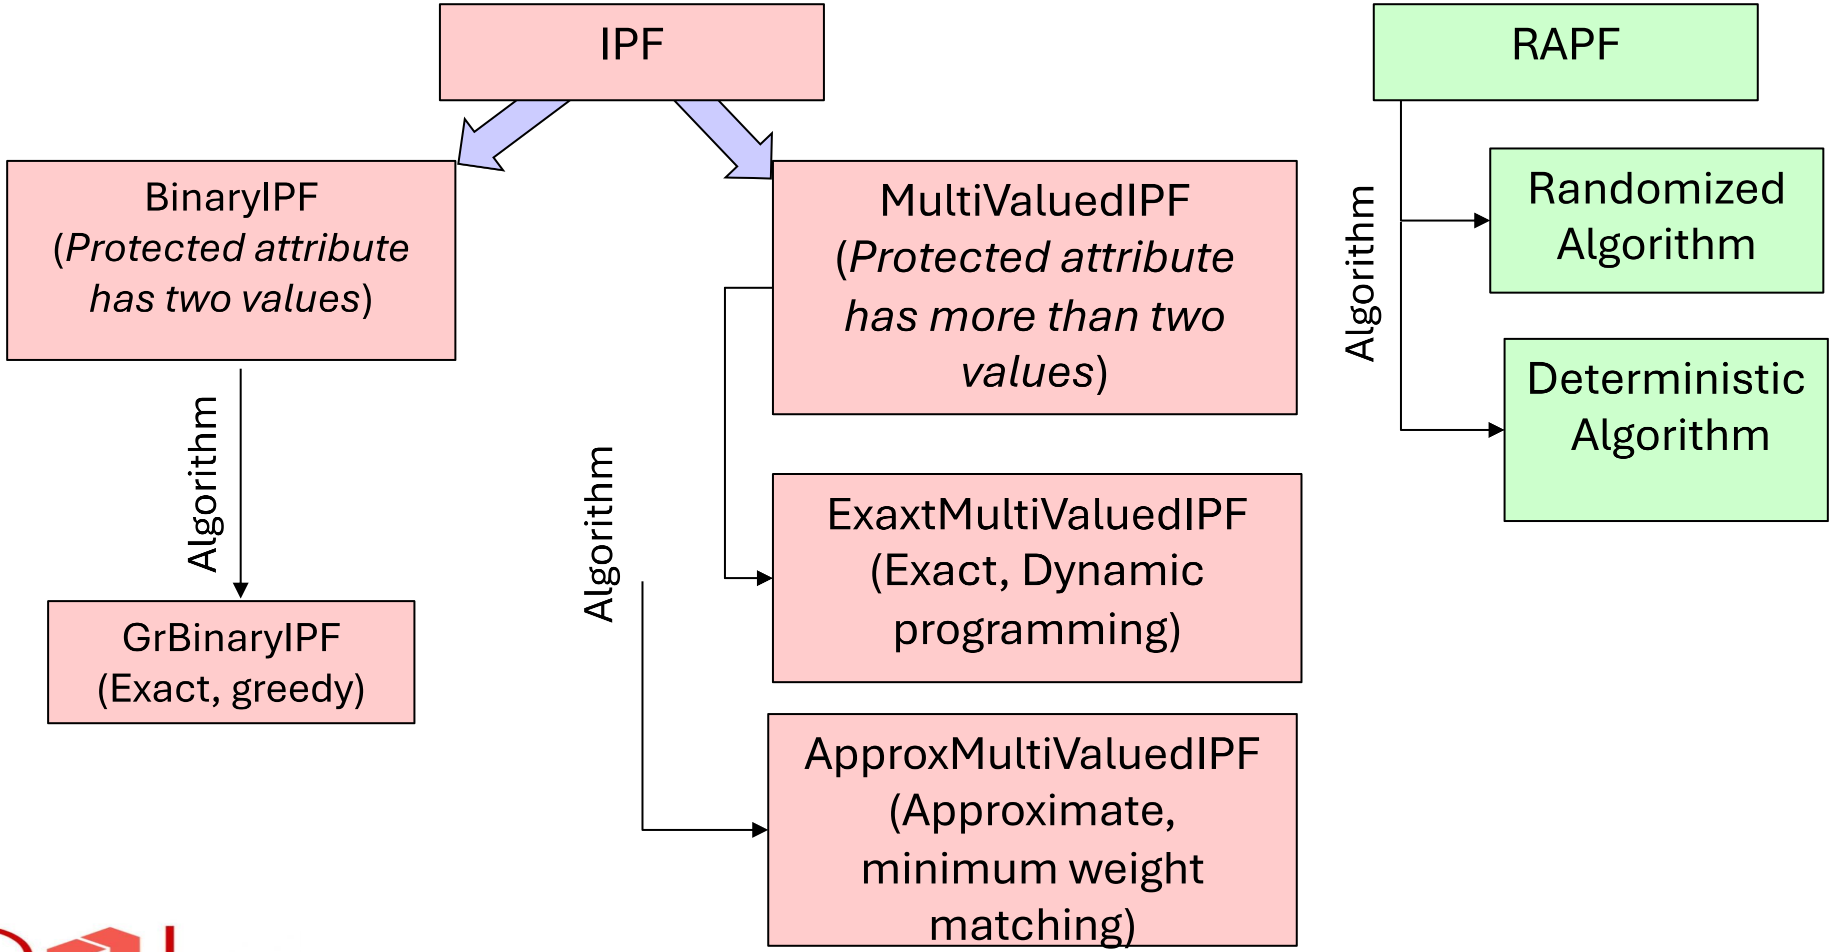

Comparison with Existing Work

Fairness in a single rank		Fairness in Rank Aggregation	
DetConstSort [3]	**IPF**	**FairILP[4]**	**RAPF**
Does not produce a p-fair ranking (only satisfies the lower bound of p-fairness)	Satisfies p-fairness.	Does not ensure p-fairness (Only satisfy pairwise statistical parity constraint)	Satisfy p-fairness.
Does not produce the closest ranking that satisfies the p-fairness lower bound	GrBinaryIPF, ExactMultiValuedIPF produces closest p-fair ranking.	Produces closest ranking that satisfy pairwise statistical parity constraint.	Does not produce the closest ranking that satisfies p-fairness.
Does not provide theoretical guarantees	**Exact solutions** GrBinaryIPF, ExactMultiValuedIPF **2-approximation factor** ApproxMultiValuedIPF	Exact algorithm	**2-approximation factor** RAPF(binary) **3-approximation factor** RandAlgRAPF+ ExactMultiValuedIPF **4-approximation factor** RandAlgRAPF+ ApproxMultiValuedIPF
Computationally scalable.	Computationally scalable.	Computationally not scalable.	Computationally scalable.

[3] Sahin Cem Geyik,et. al. Fairnessaware ranking in search & recommendation systems with application to linkedin talent search. SIGKDD 2019

[4] Caitlin Kuhlman and Elke Rundensteiner. 2020. Rank aggregation algorithms for fair consensus. Proceedings of the VLDB Endowment 13, 12 (2020), 2706–2719.

Satisfying Complex Top-k Fairness Constraints by Preference Substitutions

(VLDB 2023)

AUTHORS:
Md Mouinul Islam, Dong Wei, Baruch Schieber, and Senjuti Basu Roy

Top-k Selection via Plurality Voting

Voters and candidates. m *users (voters) provide* ***top-1*** *preference over* ***n*** *candidates*

Find k candidates with the k highest numbers of votes

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	Count
C1	1	1	0	1	0	0	1	0	0	0	0	0	4
C2	0	0	1	0	1	0	0	0	1	0	0	0	3
C3	0	0	0	0	0	1	0	0	0	0	0	1	2
C4	0	0	0	0	0	0	0	0	0	1	1	0	2
C5	0	0	0	0	0	0	0	1	0	0	0	0	1
C6	0	0	0	0	0	0	0	0	0	0	0	0	0

If k=4, C1, C2, C3, C4 are the winners.

Complex Fairness Constraints in top-k

VLDB2024
GUANGZHOU

Query: Find top-k (4) individuals with 2 males and 2 females (constraint on Gender) and 2 seniors and 2 juniors (constraint on Seniority Level) and 2 married, 1 single, and 1 divorced (constraint on Marital status)

Query comes with one or more fairness constraints defined over one or more protected attributes

Attribute	**Value**	**Constraint**
Gender **Binary attribute**	Male (M)	2
	Female (F)	2
Seniority Level **Binary attribute**	Senior (Sr)	2
	Junior (Jr)	2
Marital Status Multivalued Attribute	Married (ma)	2
	Single (si)	1
	Divorced (di)	1

Big Data Analytics Lab

Candidate	**Value**
C1	M, Sr, si
C2	M, Jr, si
C3	M, Jr, ma
C4	F, Jr, si
C5	F, Jr, ma
C6	F, Sr, di

The current top-4 candidates {C1, C2, C3, C4} DO NOT satisfy the query!!!

Preference (Ballot) Substitution to Satisfy Query Constraints

Single Ballot substitution. *Remove one vote from candidate i and assign it to candidate j;*

- Perform 3 single ballot substitutions

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	Count
C1 (M,Sr,si)	1	1	0	1	0	0	1	0	0	0	0	0	4
C2 (M,Jr,si)	0	0	1	0	1	0	0	0	1	0	0	0	3
C3 (M,Jr,ma)	0	0	0	0	0	1	0	0	0	0	0	1	2
C4 (F,Jr,si)	0	0	0	0	0	0	0	0	0	1	1	0	2
C5 (F,Jr,ma)	0	0	0	0	0	0	0	1	0	0	0	0	1
C6 (F,Sr,di)	0	0	0	0	0	0	0	0	0	0	0	0	0

- by removing 2 votes from C2 and 1 vote from C4
- assigning 2 votes to candidate C6 and 1 vote to C5.

- The resulting top-4 candidates (C1,C3,C5,C6) with votes (4, 2, 2, 2) satisfy fairness constraints.

Preference (Ballot) Substitution to Satisfy Query Constraints

VLDB2024 GUANGZHOU

Single Ballot substitution. *Remove one vote from candidate i and assign it to candidate j;*

- Perform 3 single ballot substitutions

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	Count
C1 (M,Sr,si)	1	1	0	1	0	0	1	0	0	0	0	0	4
C2 (M,Jr,si)	0	0	1	0	1	0	0	0	1	0	0	0	3
C3 (M,Jr,ma)	0	0	0	0	0	1	0	0	0	0	0	1	2
C4 (F,Jr,si)	0	0	0	0	0	0	0	0	0	1	1	0	2
C5 (F,Jr,ma)	0	0	0	0	0	0	0	1	0	0	0	0	1
C6 (F,Sr,di)	0	0	0	0	0	0	0	0	0	0	0	0	0

  - by removing 2 votes from C2 and 1 vote from C4
  - assigning 2 votes to candidate C6 and 1 vote to C5.
- The resulting top-4 candidates (C1,C3,C5,C6) with votes (4, 2, 2, 2) satisfy fairness constraints.

Big Data Analytics Lab

Problem definition

Single ballot substitution. Given two candidates i and j , a single ballot substitution is defined as removing one vote from candidate i and assigning it to candidate j

Margin. Margin is the minimum number of single ballot substitutions needed to guarantee that the top-k results satisfy the fairness constraints.

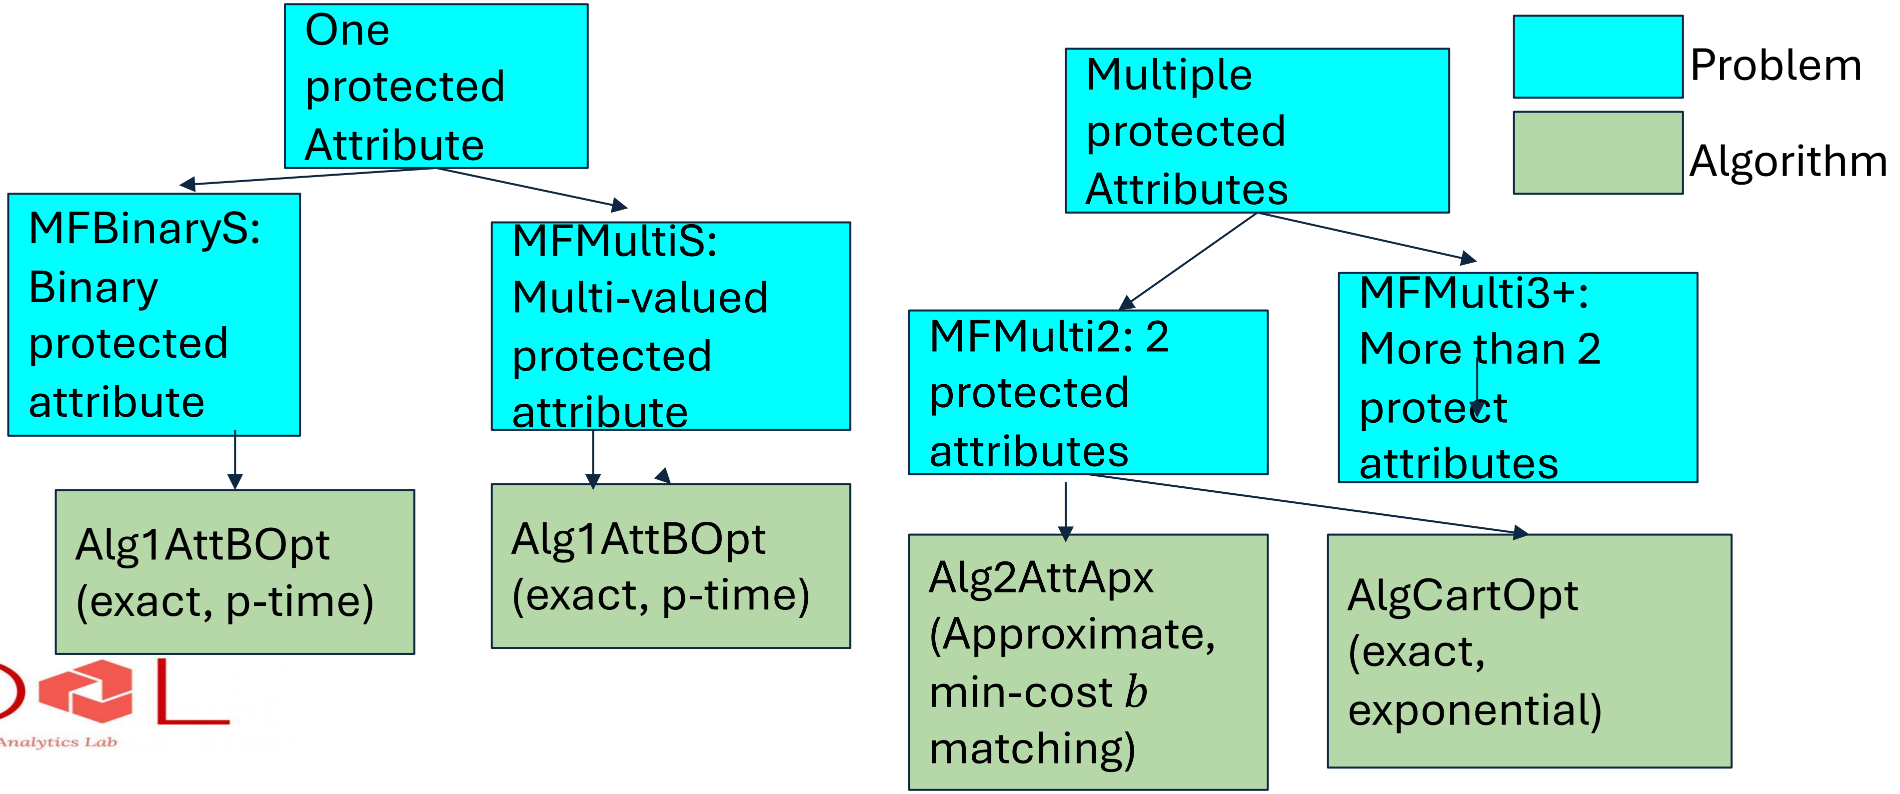

Related Work

Subgroup fairness [3,4]		Margin finding problem [5,6,7]	
Existing solutions	**Our Solution**	**Existing solutions**	**Our Solution**
Studied primarily in the context of classification models to ensure that a classifier is fair not only on each individual group but also it stays fair when more structured subgroups are defined over the protected attributes.	Fairness constraints are imposed over multiple protected attributes and could be imposed over subgroups with the goal of returning top-k results that satisfy constraints.	Margin of victory is studied in electoral voting systems to understand robustness of an underlying voting mechanism, specifically for the Single Transferable Vote (STV).	Margin finding problem is studied in the context of plurality voting to satisfy complex fairness constraints.
Audit a fixed classifier to see if it satisfies subgroup fairness (i.e., the false positive rates are equivalent across all subgroups) is deemed to be computationally hard [4]	Margin finding problem for 2 attributes is NP-hard. Deciding the feasibility of a fair outcome for 3 or more attributes is NP-Complete.	Orlin and Bartholdi proved margin finding is NP-hard even for a single candidate selection for STV [5]	P-time solution for single protected attribute, NP-hard for multiple protected attribute
Heuristic solutions	Exact or approximate solutions	Fairness constraints are not considered	Fairness constraints are imposed

Big Data Analytics Lab

[3] Subgroup fairness. Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2019. An Empirical Study of Rich Subgroup Fairness for Machine Learning. In Proceedings of the Conference on Fairness, Accountability, and Transparency. ACM. https://doi.org/10.1145/3287560.3287592

[4] Michael J. Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In Jennifer G. Dy and Andreas Krause, editors, Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, volume 80 of JMLR Workshop and Conference Proceedings, pages 2569– 2577. JMLR.org, 2018. URL http://proceedings.mlr.press/v80/kearns18a.html

[5] John J Bartholdi and James B Orlin. 1991. Single transferable vote resists strategic voting. Social Choice and Welfare 8, 4 (1991), 341–354

[6] Michelle Blom, Peter J Stuckey, and Vanessa J Teague. 2017. Towards computing victory margins in STV elections. arXiv preprint arXiv:1703.03511 (2017)

[7] David Cary. 2011. Estimating the Margin of Victory for Instant-Runoff Voting. In 2011 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 11).

Promoting Fairness and Priority in k-Winners Selection Using IRV

(KDD 2024)

AUTHORS:
Md Mouinul Islam, Soroush Vahidi, Baruch Schieber, and Senjuti Basu Roy

Instant Run off Voting (IRV)

- Instant Run-off Voting (IRV):
  - Tally the first choice votes
  - The candidate that has the lowest number of first choice votes is eliminated. Ties are broken arbitrarily
  - All the ranked orders that include the eliminated candidate are updated, and the candidates following this eliminated candidate are advanced one place up

1	2	3	4	5	6	7	8	9	10
		Sara	Sara	Sara	Sara	Sara	Sara	Sara	Sara

Molly Mira Gina Kim Laura Zoey

Motivating Example

- Faculty hiring committee ranking

1	2	3	4	5	6	7	8	9	10
Zoey	Laura	Zoey	Zoey	Mira	Sara	Gina	Sara	Kim	Laura
Mira	Gina	Molly	Molly	Molly	Gina	Sara	Gina	Gina	Kim
Laura	Molly	Kim	Sara	Sara	Kim	Kim	Kim	Sara	Gina
Gina	Kim	Gina	Gina	Kim	Zoey	Mira	Molly	Molly	Sara
	Zoey	Sara		Zoey		Zoey	Zoey	Zoey	

- Hiring priorities

1st	**DM**	Molly	Zoey	
2nd	**ML**	Gina	Laura	Kim
3rd	**AI**	Mira	Sara	

Instant Run-Off Voting

- IRV is gaining popularity around the world (Australia, Ireland, Cambridge USA, NYS USA)
- IRV properties
  - proportional representation of solid coalition
  - promotes anti-plurality
  - reduces conflict within the electorate
  - reduces strategic voting
  - amenable to incomplete ranked order
  - representativeness in query results [Behar & Cohen, SIGMOD 20222]

k-Winners Selection

- Users (Voters) cast ballots, which are ranked orders of candidates
- Each candidate belongs to one out of k groups
- Select k winners, one winner per group, that represent the voters' preferences
- Motivation: fairness and priority
- Approach: find a "minimal manipulation" of the ballots that guarantees the selection from the required groups

Constructive Margin of Victory (MOV)

- Number of ballot modifications required to guarantee winner from a given group

1st	DM	Molly	Zoey	

						1	2	3	4	5	6	7	8	9	10
						Molly	Molly	Molly	Molly	Molly	Molly	Molly	Molly	Molly	Molly

Mira Gina Kim Laura Zoey Sara

Computing Constructive MOV

- **Theorem:** Computing the constructive margin of victory is NP-Complete, even for ballot size 2
- **Theorem:** Computing the constructive margin of victory can be formulated as an integer linear program (IP)

$$\min \sum_{s \in S} a_s \quad \text{subject to}$$

$$\begin{array}{lll}
m_s + a_s - d_s = y_s & \forall\, s \in S & (1) \\
m \geq y_s \geq 0 & \forall\, s \in S & (2) \\
m_s \geq d_s \geq 0 & \forall\, s \in S & (3) \\
m - m_s \geq a_s \geq 0 & \forall\, s \in S & (4) \\
\sum_{s \in S} a_s = \sum_{s \in S} d_s & & (5) \\
u_{c_i,c_j} + u_{c_j,c_i} = 1 & \forall\, \{c_i, c_j\} \subseteq C & (6) \\
u_{c_i,c_j} + u_{c_j,c_r} + u_{c_r,c_i} \geq 1 & \forall\, \{c_i, c_j, c_r\} \subseteq C & (7) \\
v_{s,c_i,\tilde{c}} = u_{c_i,\tilde{c}} \cdot \Pi_{x=1}^{i-1} u_{\tilde{c},c_x} & \forall\, s \in S\ \forall\, \tilde{c} \in C & \\
 & \forall\, i \in \{1, \ldots, |s|\} & (8) \\
\sum_s (y_s \cdot v_{s,\hat{c},\tilde{c}}) \geq u_{\hat{c},\tilde{c}} \cdot \sum_s (y_s \cdot v_{s,\tilde{c},\hat{c}}) & \forall\, \{\hat{c}, \tilde{c}\} \subseteq C & (9)
\end{array}$$

A Branch and Bound Algorithm

- Based on the framework given in [Magrino et al. EVT/WOTE 2011] and [Blom et al. INFORMS 2019] for destructive margin of victory computation
- Enumerate over all the elimination orders that end with a winner from the desired set to find the one with the least number of modifications

A Branch and Bound Algorithm (cont.)

```
AlgExact(B,C,W)

ub = MqIRVUB(B,C,W)
initialize a priority queue PQ with (ω,0) ∀ω∈W
while PQ is not empty
   retrieve a minimum element (π,lb)
   for c∈C-π
      π = c+π  /* extend the elimination sequence π by c
      lb = DistToLB(B,C,π) /* lower bound on # of modifications
      if lb ≤ ub insert (π,lb) to PQ
      if |π| = |C|
         ub = min(ub,DistTo(B,C,π)) /* compute the exact # of
                                    /* modifications using IP
return ub
```

- A tighter lower bound and the use of an upper bound results in improved efficiency than previous work

Computing DistTo

- DistTo(B,C,π): the minimum number of ballot "manipulations" required to achieve an elimination sequence
- **Theorem:** Computing DistTo in case of ballot modifications is NP-Complete, even for ballot size 3.
- **Theorem:** DistTo in case of ballot additions can be computed in polynomial time

Summary: Computational Approach

Preference elicitation	Aggregation model	Output	Fairness	Approach
Score based	Average relevance, Least Misery	Top-k set	m-proportionality, m-envyfreeness, dissatisfaction score	Minimizing unfairness subject to quality constraints
Ranked order	Kemeny	Ranked k items	Top-k parity defined on a single binary or multi protected attribute	Find closest ranking that satisfies fairness
Ranked order	Kemeny, Spearman's footrule	Full rank	p-fairness defined on a single binary or multi protected attribute	Find closest ranking that satisfies fairness
Plurality	Plurality voting	Top-k set	Multiple protected attributes that may not be independent	Compute constructive MOV
Ranked order upto index l	Instant run off voting	Single winner	Fairness defined on a single protected attribute	Compute constructive MOV

Part III – Future Research Directions (30 min)

III-A : New Preference Aggregation Methods

- Producing structurally-involved outputs
- Dealing with uncertainties
- Going perpetual

III-A : New Preference Aggregation Methods

- **Producing structurally-involved outputs**
- Dealing with uncertainties
- Going perpetual

III-A : New Preference Aggregation Methods

- Beaten path
  - Single-winner > multiwinner > participatory budgeting
- Application-driven route (off the beaten path)
  - Resource allocation
  - Text aggregation

III-A : New Preference Aggregation Methods

Water distribution

Fair Division with Storage

Eyal Briman[1], Nimrod Talmon[1], Stephane Airiau[2], Umberto Grandi [3], Jerome Lang [2], Jerome Mengin [3] and Faria Nasiri Mofakham[4]

[1]Ben Gurion University of the Negev
[2]Paris Dauphine University
[3]Toulouse University Capitole
[4]University of Isfahan

Big Data Analytics Lab

III-A : New Preference Aggregation Methods

Water distribution

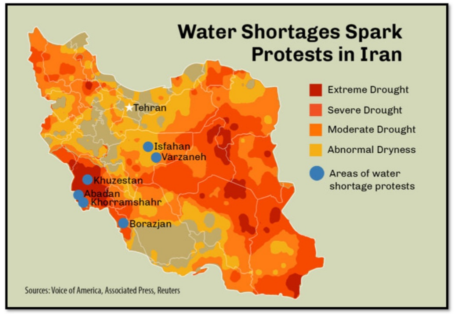

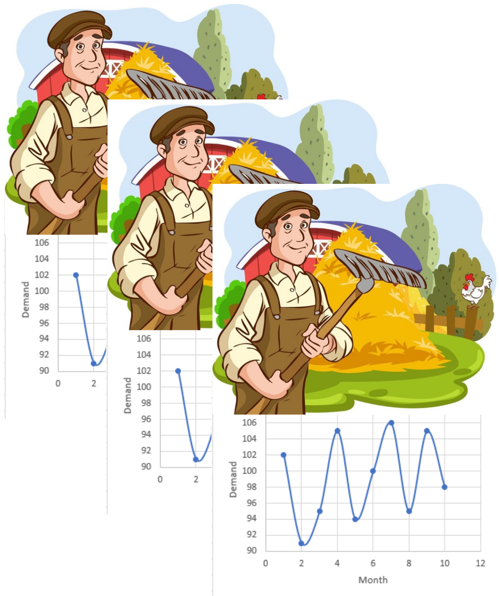

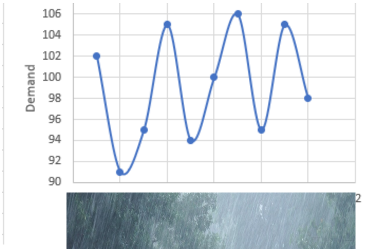

III-A : New Preference Aggregation Methods

Text aggregation

Aggregation over Metric Spaces: Proposing and Voting in Elections, Budgeting, and Legislation

Laurent Bulteau LAURENT.BULTEAU@U-PEM.FR
LIGM, CNRS, Univ Gustave Eiffel
5 Bd Descartes, 77454 Marne la Vallée, France

Gal Shahaf GAL.SHAHAF@WEIZMANN.AC.IL
Ehud Shapiro EHUD.SHAPIRO@WEIZMANN.AC.IL
Weizmann Institute of Science
Herzl St 234, Rehovot, Israel

Nimrod Talmon TALMONN@BGU.AC.IL
Ben-Gurion University
Ben-Gurion Boulevard, Be'er Sheva, Israel

Abstract

We present a unifying framework encompassing a plethora of social choice settings. Viewing each social choice setting as voting in a suitable metric space, we offer a gen-

AI-Generated Compromises for Coalition Formation

Anonymous

Collaborative Document Writing: an Iterative Thresholds Approach

Avital Finanser Nimrod Talmon

June 25, 2024

Big Data Analytics Lab

III-A : New Preference Aggregation Methods

Text aggregation

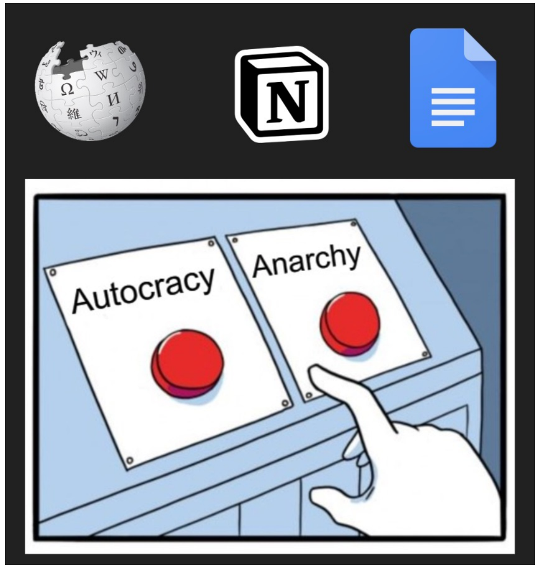

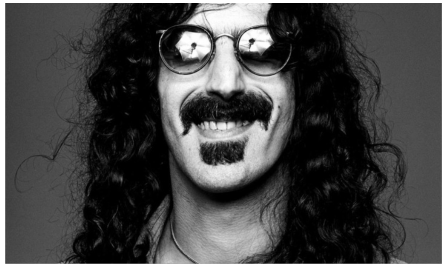

Put everything in a metric space!

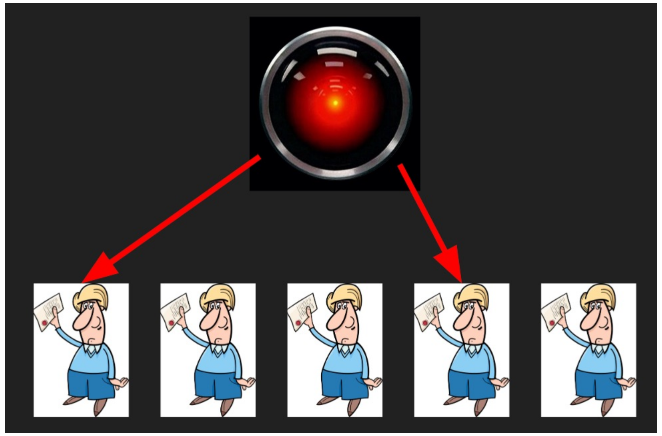

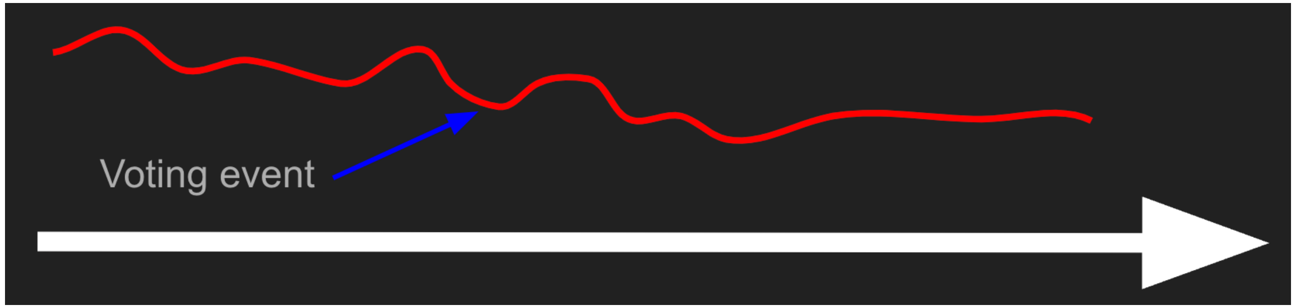

III-A : New Preference Aggregation Methods

- Aggregating more-involved preferences
- Outputing structurally-involved outputs
- **Dealing with uncertainties**
- Going perpetual

III-A : New Preference Aggregation Methods

- Dealing with uncertainties

A Recommendation System for Participatory Budgeting

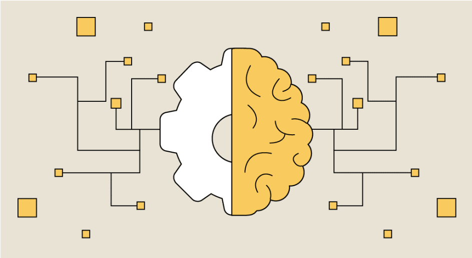

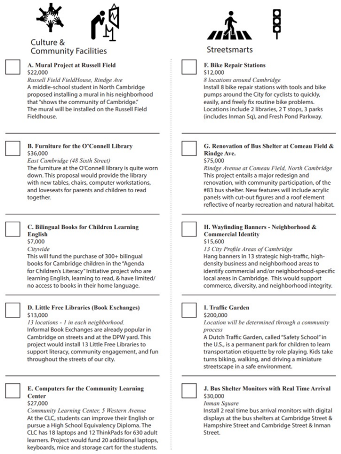

Culture & Community Facilities

☐ **A. Mural Project at Russell Field**
$22,000
Russell Field FieldHouse, Rindge Ave
A middle-school student in North Cambridge proposed installing a mural in his neighborhood that "shows the community of Cambridge." The mural will be installed on the Russell Field Fieldhouse.

☐ **B. Furniture for the O'Connell Library**
$36,000
East Cambridge (48 Sixth Street)
The furniture at the O'Connell library is quite worn down. This proposal would provide the library with new tables, chairs, computer workstations, and loveseats for parents and children to read together.

☐ **C. Bilingual Books for Children Learning English**
$7,000
Citywide
This will fund the purchase of 300+ bilingual books for Cambridge children in the "Agenda for Children's Literacy" Initiative project who are learning English, learning to read, & have limited/ no access to books in their home language.

☐ **D. Little Free Libraries (Book Exchanges)**
$13,000
13 locations - 1 in each neighborhood.
Informal Book Exchanges are already popular in Cambridge on streets and at the DPW yard. This project would install 13 Little Free Libraries to support literacy, community engagement, and fun throughout the streets of our city.

☐ **E. Computers for the Community Learning Center**
$27,000
Community Learning Center, 5 Western Avenue
At the CLC, students can improve their English or pursue a High School Equivalency Diploma. The CLC has 18 laptops and 12 ThinkPads for 630 adult learners. Project would fund 20 additional laptops, keyboards, mice and storage cart for the students.

Streetsmarts

☐ **F. Bike Repair Stations**
$12,000
8 locations around Cambridge
Install 8 bike repair stations with tools and bike pumps around the City for cyclists to quickly, easily, and freely fix routine bike problems. Locations include 2 libraries, 2 T stops, 3 parks (includes Inman Sq), and Fresh Pond Parkway.

☐ **G. Renovation of Bus Shelter at Comeau Field & Rindge Ave.**
$75,000
Rindge Avenue at Comeau Field, North Cambridge
This project entails a major redesign and renovation, with community participation, of the #83 bus shelter. New features will include acrylic panels with cut-out figures and a roof element reflective of nearby recreation and natural habitat.

☐ **H. Wayfinding Banners - Neighborhood & Commercial Identity**
$15,600
13 City Profile Areas of Cambridge
Hang banners in 13 strategic high-traffic, high-density business and neighborhood areas to identify commercial and/or neighborhood-specific local areas in Cambridge. This would support commerce, diversity, and neighborhood integrity.

☐ **I. Traffic Garden**
$200,000
Location will be determined through a community process
A Dutch Traffic Garden, called "Safety School" in the U.S., is a permanent park for children to learn transportation etiquette by role playing. Kids take turns biking, walking, and driving a miniature streetscape in a safe environment.

☐ **J. Bus Shelter Monitors with Real Time Arrival**
$30,000
Inman Square
Install 2 real time bus arrival monitors with digital displays at the bus shelters at Cambridge Street & Hampshire Street and Cambridge Street & Inman Street.

Big Data Analytics Lab

Using Liquid Democracy for Attention-Aware Social Choice

Shiri Alouf-Heffetz
Ben Gurion University
shirih@post.bgu.ac.il

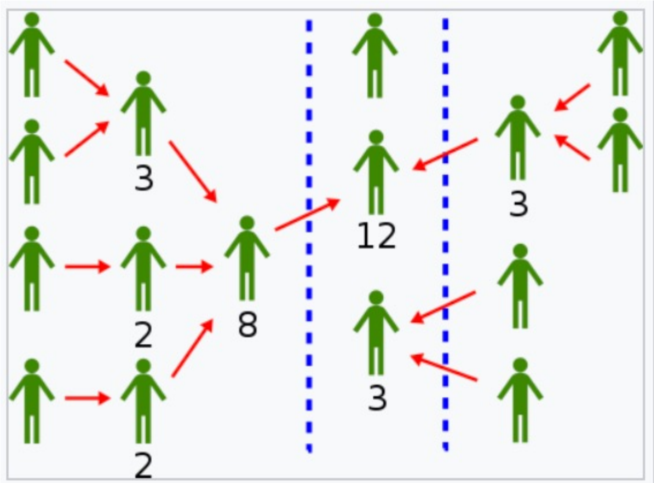

III-A : New Preference Aggregation Methods

- Aggregating more-involved preferences
- Outputing structurally-involved outputs
- Dealing with uncertainties
- **Going perpetual**

Big Data Analytics Lab

III-A : New Preference Aggregation Methods

- Perpetual voting
  - **Making repeated decisions**
  - While being fair to external properties!

	Monday	Tuesday	Wednesday	Thursday	Friday
Student 1	C	C	C	C	A,C
Student 2	A	A	A	A	B
Student 3	A,B	A	A,C	A	A
Student 4	A	A	B	C	A
Student 5	B,C	B	B	B,C	A

Justified Representation for Perpetual Voting

LAURENT BULTEAU[1], NOAM HAZON[2], RUTVIK PAGE[3], ARIEL ROSENFELD[4], AND NIMROD TALMON[5]

[1]Laboratoire d'Informatique Gaspard Monge, 77454 Marne-la-Vallée, France
[2]Department of Computer Science, Ariel University, Ariel 40700, Israel
[3]Department of Computer Science and Engineering, Indian Institute of Information Technology, Nagpur, Nagpur 440006, India
[4]Department of Information Science, Bar-Ilan University, Ramat Gan 5290002, Israel
[5]Ben Gurion University of the Negev, Be'er Sheva 8410501, Israel

Corresponding author: Noam Hazon (noam.hazon@g.ariel.ac.il)

The work of Noam Hazon was supported by the Ministry of Science, Technology and Space, Israel. The work of Nimrod Talmon was supported by the Israel Science Foundation (ISF) under Grant 630/19.

III-A : New Preference Aggregation Methods

- Perpetual voting
  - Making repeated decisions
  - **While being fair to external properties!**

III-B : Perpetual Axiomatic Properties related to Individual Fairness

- (**Could be treated as candidate's long term fairness**) Simple proportionality – In a sequence of n rounds, if each of the n voters prefers only one candidate and no voter changes their preference across rounds, the number of times a candidate gets selected is equal to the proportion of votes the candidate has received.

- (**Could be treated as voter's long term fairness**) Bounded dry spell - Given a sequence of n rounds, a voter v has a dry spell of length l if v's preferred candidate(s) get selected in at most every l rounds.

Lackner, Martin. "Perpetual voting: Fairness in long-term decision making." *Proceedings of the AAAI conference on artificial intelligence*. Vol. 34. No. 02. 2020.

Simple Proportionality

***Voters and candidates. m** users (voters) provide **top-1** preference over **n** candidates*

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	Count
C1	1	1	0	1	0	0	1	0	0	0	0	0	4
C2	0	0	1	0	1	0	0	0	1	0	0	0	3
C3	0	0	0	0	0	1	0	0	0	0	0	1	2
C4	0	0	0	0	0	0	0	0	0	1	1	0	2
C5	0	0	0	0	0	0	0	1	0	0	0	0	1
C6	0	0	0	0	0	0	0	0	0	0	0	0	0

After 12 rounds, C1 wins 4 times, C2 wins 3 times, C3 wins 2 times, C4 wins 2 times, C5 wins 1 time, and C6 wins 0 time

Bounded Dry Spell

VLDB2024 GUANGZHOU

***Voters and candidates. m** users (voters) provide **top-1** preference over **n** candidates*

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11	V12	Count
C1	1	1	0	1	0	0	1	0	0	0	0	0	4
C2	0	0	1	0	1	0	0	0	1	0	0	0	3
C3	0	0	0	0	0	1	0	0	0	0	0	1	2
C4	0	0	0	0	0	0	0	0	0	1	1	0	2
C5	0	0	0	0	0	0	0	1	0	0	0	0	1
C6	0	0	0	0	0	0	0	0	0	0	0	0	0

If the sequence of winning candidates in 12 rounds are C1, C2, C3, C4, C5, C1, C2,C3, C4, C5, C1, C2, every voter has a bounded dry spell of 4!

Big Data Analytics Lab

III-C : Technical Problems

- Is it possible to design a sequence of winners that satisfies
  - Perpetual axioms (e.g., simple proportionality/ bounded dry spell) **AND**
  - Group fairness (p-fairness or its relaxed version)?

  - Depends on input - Does not seem to be always possible!
- How to minimally change input preferences such that a sequence of designed candidates simultaneously satisfies simple proportionality and p-fairness?
- How to design a sequence of candidates that satisfies **simple proportionality** and **satisfies p-fairness as closely** as possible?

III-B : Alternative models for p-Fairness

- p-fairness may be too restrictive requirement
- Relaxed p-fairness 1: every prefix has “roughly” the proportionate number of occurrences from every protected attribute
  - The amount of deviation can be controlled
- Relaxed p-fairness 2: impose the proportionality constraints only at selected prefixes, e.g., every 10-percentile

III-C : Technical Problems

- Does there exist any relationship between bounded dry spell and p-fairness?
  - When a voter selects only 1 preferred candidate
  - When a voter selects a subset of candidates as their preference
- can we approximate the margin of victory (either destructive or constructive) of IRV and STV in polynomial time?
- can we improve the computational efficiency of margin computation

III-C : Technical Problems

- Handling more than one protected attribute in parallel. For example, guarantee fairness based on two protected attributes: gender and race.
- A simplistic approach is to assume that the attributes are independent and then consider the Cartesian product of the domains.
- Can this be done without assuming independence?

Team and Collaborators

Dong Wei
(Graduated Fall 2022,
now at Google NYC)

Md Mouinul Islam
(Graduated Summer 2023,
research scientist at Paypal)

Soroush Vahidi
PhD Student

Manish Kumar
Postdoc

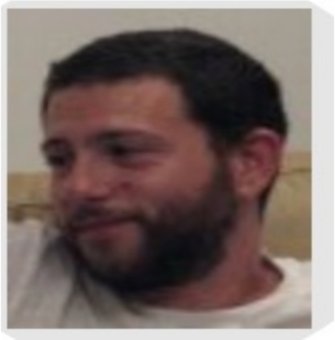

Baruch Scieber

Nomrod Talmon

Senjuti Basu Roy

Thank you – Questions?