Temporal Panel Selection in Ongoing Citizens' Assemblies

Yusuf Hakan Kalayci
University of Southern California
Los Angeles, United States
kalayci@usc.edu

Evi Micha
University of Southern California
Los Angeles, United States
evi.micha@usc.edu

ABSTRACT

Permanent citizens' assemblies are ongoing deliberative bodies composed of randomly selected citizens, organized into panels that rotate over time. Unlike one-off panels, which represent the population in a single snapshot, permanent assemblies enable shifting participation across multiple rounds. This structure offers a powerful framework for ensuring that different groups of individuals are represented over time across successive panels. In particular, it allows smaller groups of individuals that may not warrant representation in every individual panel to be represented across a sequence of them. We formalize this temporal sortition framework by requiring proportional representation both within each individual panel and across the sequence of panels.

Building on the work of Ebadian and Micha (2025), we consider a setting in which the population lies in a metric space, and the goal is to achieve both proportional representation, ensuring that every group of citizens receives adequate representation, and individual fairness, ensuring that each individual has an equal probability of being selected. We extend the notion of representation to a temporal setting by requiring that every initial segment of the panel sequence, viewed as a cumulative whole, proportionally reflects the structure of the population. We present algorithms that provide varying guarantees of proportional representation, both within individual panels and across any sequence of panels, while also maintaining individual fairness over time.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory**; **Facility location and clustering**.

KEYWORDS

Permanent Citizens' Assemblies, Sortition, Proportional Representation, Clustering, Computational Social Choice

ACM Reference Format:
Yusuf Hakan Kalayci and Evi Micha. 2026. Temporal Panel Selection in Ongoing Citizens' Assemblies. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/XBME2274

1 INTRODUCTION

In recent years, *citizens' assemblies* have emerged as a promising model for community-driven governance [32, 34]. The core idea is to randomly select a panel from the population, provide them with balanced information on an issue, allow time for deliberation, and then collect their recommendations, which are taken into account by the relevant authorities. This approach has been adopted in a variety of local contexts, with many city governments and local authorities implementing citizens' assemblies to address policy issues, as well as in several prominent recent examples at the national level—such as in Ireland [1]— and at the continental level through the European Union[2]. At the heart of citizens' assemblies is *sortition*: a randomized selection process creating a panel that reflects the general population's diverse perspectives, not just technical expertise. To support this, the process aims for representativeness across key features (e.g., geography, education level) while ensuring every individual has a meaningful selection chance. In practice, organizers often implement stratified sampling with quotas on individual features—for instance, reserving a fixed percentage of seats for a particular age group or requiring that at least a certain proportion of representatives have a specific occupation [17]. However, this method fails to account for more complex combinations of features—for example, individuals who are both from a specific region and work in a particular occupation. On the other hand, forcing representation across all such combinations is typically infeasible; with six features each taking just three values, there are already over 700 disjoint groups, far more than a typical panel in practice.

As a middle ground, recent works [13, 14] propose leveraging an underlying representation metric space for defining when a panel represents an underlying population in a more rigorous way. This space measures how well an individual represents another individual or group, with smaller distances indicating stronger representation and can be constructed based on features relevant to the application. Given such a metric space, Ebadian and Micha [14] adapt a notion of proportional representation previously used in multi-winner elections [29] and clustering [2, 9, 30] to define when a panel is proportionally representative of an underlying population. At a high level, the notion requires that, given a population of size n and a panel of size k, any subset of the population of size at least $q \cdot n/k$ is entitled to up to q representatives.

As citizens' assemblies gain increasing acceptance, the next step in this progression is the growing establishment of *permanent citizens' assemblies*, where members are periodically rotated through sortition. The Ostbelgien Model in the German-speaking community of Belgium is one of the most prominent examples [31]. Established in 2019, it provides citizens a permanent voice in decision making through rotating panels, and the European Union has recently experimented with a similar format [1]. This new innovation motivates the main question in this work:

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. https://doi.org/10.65109/XBME2274

[1] https://www.citizensinformation.ie/en/government-in-ireland/irish-constitution-1/citizens-assembly/
[2] https://citizens.ec.europa.eu/european-citizens-panels_en

> *In ongoing citizens' assemblies, can we ensure that smaller groups receive representation over time, even if they do not qualify in every individual panel, while still guaranteeing that every sufficiently large group is proportionally represented in each panel?*

Temporal Sortition. To illustrate the idea, consider a toy example in which the population consists of four disjoint groups: the first makes up half of the population, the second a quarter, and the remaining two one-eighth each. Suppose we are selecting a panel of size 4. Proportionality would assign 2 seats to the first group, 1 to the second, and the last to either the third or fourth group, but not both. However, if we construct two consecutive panels, we can alternate this final seat—assigning it to the third group in the first panel and to the fourth in the second. In this way, both smaller groups receive representation over time, even if not in every individual panel.

This simple example demonstrates how permanent citizens' assemblies can extend proportional representation beyond the limits of any single panel. To capture this formally, we introduce the *temporal sortition framework*. Given a population V of size n in a representative metric space, our objective is to construct a panel sequence $P_1, P_2, \dots, P_\ell$, where panel P_i has size k_i. Leveraging notions of proportional representation similar to those utilized by Ebadian and Micha [14] for static panels, and assuming the existence of such a metric space, our goal is to define a distribution over these ℓ consecutive panels that satisfies the following properties:

(1) *Representation per Panel:* Every subset of the population of size at least n/k_i (or $q \cdot n/k_i$) receives at least 1 (or q) representatives in panel P_i.
(2) *Representation Over Time:* Every subset of the population of size at least $n/\sum_{i=1}^{t} k_i$ (or $q \cdot n/\sum_{i=1}^{t} k_i$) receives at least 1 (or q) representatives among the first t panels, i.e., in $\cup_{i=1}^{t} P_i$.
(3) *Individual Fairness:* Each individual is included in the union of all panels $\cup_{i=1}^{\ell} P_i$ with equal probability, equal to $\sum_{i=1}^{\ell} k_i/n$.

To define representation rigorously, we draw on two notions from the literature: α-Proportionally Fair Clustering (α-PFC) [9], which guarantees that any group of at least n/k individuals has a nearby representative, and β-Proportionally Representative Fairness (β-PRF) [3], which additionally requires proportionally many representatives for larger groups (formally defined in Section 2).

Technical Challenge and Our Contributions. In Section 3, as a warm-up, we consider the simpler setting where the objective is to ensure proportional representation for each individual panel as well as for the *global panel* $\bigcup_{i\in[\ell]} P_i$. We investigate what happens when a smaller panel is constructed from a larger one: given a panel P_1 of size k_1 satisfying α-PRF, and a sub-panel P_2 of size k_2 satisfying β-PRF with respect to P_1, we show that P_2 represents the original population within a $(2\alpha\beta + \beta)$-approximation to PRF when k_1 is divisible by k_2. Surprisingly, when k_1 is *not* divisible by k_2, there exist instances where P_2 fails to satisfy *any* finite approximation. Using this, for ℓ equal-sized panels, combining existing algorithms yields constant-factor PRF for both individual panels and the global panel.

In Section 4, we strengthen the requirements to demand representation for each panel, the global panel, and every prefix $P_{\le t} = \bigcup_{i\in[t]} P_i$. This significantly increases the difficulty of the problem, since the goal is to construct a representative global panel, decompose it into representative individual panels, and order them so that every prefix $P_{\le t}$ is also representative, as if the task were to build a panel of size equal to the total size of that prefix, all while also ensuring individual fairness. Though this appears challenging, we design an algorithm for ℓ panels of size k that guarantees individual fairness and a PFC approximation, with respect to both each individual panel and each prefix $P_{\le t}$, that scales exponentially with the number of panels. The main property that the algorithm exploits is that under proportionality, any subset of individuals that deserves representation in smaller panels should also be represented in larger panels. The algorithm thus constructs groups for each prefix size, enforcing that groups for larger panels are nested within those for smaller panels. The algorithm then carefully leverages this nested structure to guarantee proportional representation for each panel and for every prefix. This hierarchical approach, however, comes at the cost of an exponential blow-up in the approximation factor with respect to the number of panels ℓ.

Finally, in Section 5 we consider a relaxed but still highly demanding requirement. Instead of enforcing proportional representation for each individual panel we require it only for every prefix $P_{\le t} = \bigcup_{i\in[t]} P_i$ with $t \le \ell$. While this abandons per-panel guarantees, it still captures a strong notion of representation by ensuring that cumulative representation is preserved over time. We show that this relaxation makes the problem more tractable and we present an algorithm that achieves a constant-factor approximation of PFC with respect to every prefix, while still guaranteeing that each individual is included in the global panel with equal probability. The algorithm is particularly intriguing and the key idea is to replace the hierarchical structure with a more flexible construction that carefully links together the groups that should be represented at different panel sizes, that is, across prefixes. This design ensures that representation achieved for earlier prefixes automatically extends to later ones, preventing error from accumulating and ultimately yielding a constant-factor bound.

Taken together, these results show that temporal sortition algorithms can achieve both per-panel and cumulative (over-time) representation guarantees, along with individual fairness. At the same time, our findings uncover subtle structural challenges that make the problem far from straightforward. Throughout the paper we highlight several intriguing open questions that we hope will inspire further research.

Related Work. The design of citizens' assemblies that are representative of the broader population has received significant attention in the computer science literature [4, 7, 13, 14, 17–19]. However, to the best of our knowledge, no prior work has addressed the problem of designing a sequence of panels that are representative of the population both individually and collectively over time. The only conceptually related work is the recent proposal by Halpern et al. [21], who introduce federated assemblies: a hierarchical model in which assemblies are connected through a directed acyclic graph, and members of higher-level assemblies are drawn from lower-level ones. Their algorithms ensure individual fairness (each person has an equal probability of selection), ex ante fairness (each subassembly is expected to receive representation proportional to its size), and ex post fairness (the realized allocation closely approximates the expected proportions). In contrast, our work focuses on a sequence of stand-alone panels that, taken together over time, provide comprehensive representation, shifting the focus from structural to longitudinal representational equity.

Metric proportional representation has emerged as a central theme across computational social choice, clustering, and data summarization. Early clustering research introduced proportionally fair clustering [9] and individual fairness [23] establishing the idea of ensuring representation for sufficiently large point sets. These concepts subsequently influenced committee selection, spawning frameworks for proportionally representative committees [24] and proportionally representative fairness [3].

In sortition, Ebadian et al. [13] first initiated the idea of utilizing a metric space for achieving representation but they focus on selecting a panel that maximize the social welfare. Ebadian and Micha [14] later introduced the idea of the proportional representation over a metric space and designed the Fair Greedy Capture algorithm that maintains individual fairness while ensuring constant-factor approximation to the ex post core. These single-shot formulations ($\ell = 1$) serve as foundation for our temporal framework extending to multi-panel settings. We refer readers to [26] for the landscape of metric proportional representation and their relationships.

Fairness across time horizons has gained momentum in perpetual voting [6, 27, 28], temporal committee selection [12, 15, 16], temporal clustering [10, 11], repeated matching [8, 20, 33], and sequential allocation [5, 22]. Motivated by these works and the need for time-robust citizens' assemblies, we extend single-shot metric sortition to multi-step settings where groups underrepresented in individual panels achieve fair representation across time.

2 PRELIMINARIES

For $m \in \mathbb{N}$, let $[m] = \{1, \ldots, m\}$. We denote the population by $V = [n]$. The individuals are embedded in an underlying *representation metric space* equipped with a distance function d, where the distance between any two individuals i and j is denoted by $d(i, j)$. We assume that d is a pseudo-metric [3]. An instance of our problem is fully specified by the set of individuals and their pairwise distances. For simplicity, we use d to refer to both the instance and the underlying distance function. For any individual $v \in V$ and radius $r \geq 0$, we define the *ball* $B(v, r) = \{u \in V \mid d(v, u) \leq r\}$, i.e. the set of all individuals within distance r in the metric space from a given individual v.

A *temporal selection algorithm* $\mathcal{A}$ takes as input the population, the metric, the number of panels ℓ and panel sizes $k_1, k_2, \ldots, k_\ell$ as input and outputs a probability distribution over a sequence of ℓ disjoint panels $P_1, \ldots P_\ell$ where P_i has size equal to k_i. A sequence of panels, $P_1, \ldots P_\ell$, is in the support of $\mathcal{A}$, if the algorithm returns this sequence with positive probability. Throughout this work, we assume that $\sum_{t=1}^{\ell} k_t \leq n$, so that the total number of panel seats does not exceed the population size. This is the most natural setting in practice, as citizens' assemblies typically involve a few hundred members per panel, and even over decades of operation, the cumulative number of participants remains orders of magnitude smaller than the relevant population.

Representation Axioms. We start by defining two notions of proportional representation that have been proposed in the literature for individual panels.

Definition 2.1 (α-Proportionally Fair Clustering (PFC)[9]). *A panel $P \subseteq V$ of size k is α-proportionally fair for $\alpha \geq 1$ if, for any subset $S \subseteq V$ of size at least n/k, there exists an individual $v \in S$ and a panel member $p \in P$ such that*

$$d(v, p) \leq \alpha \cdot \min_{y \in V} \max_{u \in S} d(u, y).$$

In words, proportionally fair clustering ensures that no group of individuals of size at least n/k can identify an alternative representative not in the current panel that all its members would strictly prefer over their closest representative in the selected panel.

Definition 2.2 (β-Proportionally Representative Fairness (PRF) [3]). *A panel $P \subseteq V$ of size k is β-proportionally representative if for any set of individuals $S \subseteq V$ of size at least $q \cdot n/k$ where the maximum pairwise distance within S is r, we have:*

$$\left| P \cap \bigcup_{v \in S} B(v, \beta \cdot r) \right| \geq q.$$

In words, proportionally representative fairness ensures that for every group of individuals of size at least $q \cdot n/k$ with maximum pairwise distance r, there exist at least q representatives in the panel, each within distance r of some individual in the group.

Intuitively, proportionally fair clustering ensures that coalitions of size at least n/k receive representation by at least one panel member, while proportionally representative fairness requires that larger coalitions (of size at least $q \cdot n/k$) receive proportionally many representatives (specifically, q representatives). Recent work has established theoretical relationships between these concepts: Aziz et al. [3] and Kellerhals and Peters [26] showed that 1-PRF implies $(1+\sqrt{2})$-PFC, and this can be generalized to demonstrate that β-PRF implies $(1 + \sqrt{2}) \cdot \beta$-PFC.

Axioms for Temporal Sortition. We evaluate the fairness and representation guarantees of a distribution across series of panels $\{P_1, \ldots, P_\ell\}$ using the following axioms :

(1) **Individual Fairness:** Each individual should have equal selection probability across the entire process. We say that a selection algorithm satisfies *individual fairness* if for all $v \in V$, $\mathbf{Pr}_{P_1,\ldots,P_\ell \sim \mathcal{A}}[v \in P] = \sum_{i=1}^{\ell} k_i / n$, where $P = \bigcup_{i=1}^{\ell} P_i$.
(2) **Individual Panel Representation:** For any representation axiom Π, we say that a selection algorithm $\mathcal{A}$ satisfies Π at the "panel level" if for every $\mathcal{P} = \{P_1, \ldots, P_\ell\}$ in the support of $\mathcal{A}$ every P_i satisfies Π for population V and panel size k_i, simultaneously. In this work, we focus on α-PFC and β-PRF definitions.
(3) **Global Panel Representation:** The union of all panels should satisfy representation axioms, providing long-term representation guarantees. Given a representation axiom Π, we say that a selection algorithm $\mathcal{A}$ satisfies Π *at the global level* if for every $\mathcal{P} = \{P_1, \ldots, P_\ell\}$ in the support of

[3] A *metric* is a non-negative function d on pairs satisfying: (i) $d(x, x) = 0$ for all x, (ii) *symmetry*: $d(x, y) = d(y, x)$ for all x, y, (iii) *triangle inequality*: $d(x, z) \leq d(x, y) + d(y, z)$ for all x, y, z, and (iv) *positivity*: $d(x, y) > 0$ whenever $x \neq y$. A *pseudo-metric* relaxes the positivity requirement, allowing distinct points to have zero distance.

$\mathcal{A}$, $\bigcup_{i=1}^{\ell} P_i$ satisfies axiom Π for population V and panel size $\sum_{i=1}^{\ell} k_i$.

(4) **Prefix Representation:** For more stringent requirements beyond global panel representation, we demand that for every $t \in [\ell]$, the cumulative selection panel of the first t panels maintains proportional representation, ensuring representation quality at every stage. Given a proportionality axiom Π, we say that that a selection algorithm $\mathcal{A}$ satisfies Π at the "prefix level" if for every $\mathcal{P} = \{P_1, \ldots, P_\ell\}$ in the support of $\mathcal{A}$, the cumulative panel $P_{\leq t}$ for each $t \in [\ell]$ satisfies representation axiom Π for population V with panel size $\sum_{i=1}^{t} k_i$.

3 WARM-UP: PROPORTIONAL REPRESENTATION PER PANEL AND GLOBAL PANEL

As a warm-up, we first consider the case where the goal is to achieve representation with respect to each individual panel and the global panel, without prefix representation concerns. This setup can be interpreted as a federated sortition problem where ℓ panels operate simultaneously, and we seek to ensure representation both within each individual panel and across the union of all panels [21].

To design algorithms for this case, we begin by examining how representation guarantees behave when rules satisfying PRF axioms are applied in sequence. In particular, consider a population V and suppose we apply a selection algorithm satisfying the α-PRF axiom to obtain a panel P_1 of size k_1. We then apply another selection algorithm satisfying the β-PRF axiom to the selected panel P_1, yielding a smaller panel $P_2 \subseteq P_1$ of size k_2. This raises a natural question: what proportionality guarantees does P_2 achieve with respect to the original population V and panel size k_2? In the next theorem, we show that this approach, provides an approximation guarantee that depends on α and β when k_1 is divisible by k_2. But surprisingly, when k_1 is not divisible by k_2, then P_2 may not provide any finite approximation with respect to V.

Theorem 3.1. *Given a population V and panel sizes k_1 and k_2, let P_1 be a panel of size k_1 that satisfies α-PRF for population V, and let P_2 be a smaller panel of size k_2 that satisfies β-PRF for population P_1. Then,*

- *when k_1 is divisible by k_2, P_2 satisfies $(2\alpha \cdot \beta + \beta)$-PRF for population V.*
- *there exist panel sizes k_1 and k_2 where k_1 is not divisible by k_2 and an instance in which P_2 does not satisfy α-PRF (or α-PFC) for population V, for any finite α.*

The proof of this theorem can be found in the full version [25].

Based on the above theorem, for the simple case where all panels have the same size k across all ℓ panels, we can utilize algorithms from the literature for achieving approximately PRF for each panel and the global panel and individual fairness properties. In particular, we can first apply the algorithm called, Fair Greedy Capture by Ebadian and Micha [14], for choosing a panel P with size $\ell \cdot k$ which satisfies 6-PRF and ensures that each individual is selected with the same probability [4]. We then apply the algorithm called Metric Expanding Approval by Aziz et al. [3] with parameters P and k for partitioning the $\ell \cdot k$ individuals in P into k groups of size ℓ each. At a high level, when panel size evenly divides population size, Expanding Approval Rule proceeds by simultaneously expanding balls around each representative in P. Once a ball captures ℓ individuals, these are grouped together, and the process continues with the remaining representatives until k such groups are formed. Aziz et al. [3] show that when Metric Expanding Approval is applied to an underlying population, selecting one representative from each resulting group yields a panel that satisfies 2-PRF with respect to the population. Building on this, we assign to each individual panel P_i one representative drawn uniformly at random from each of the k groups. This guarantees that every panel receives exactly one representative from each group, while each representative is assigned to a panel with probability $1/\ell$. Combined with the fact that each individual in the population is selected into P with probability $\ell \cdot k/n$, it follows that every individual is assigned into a panel P_i with probability k/n. By Theorem 3.1, we then obtain the following corollary.

Corollary 3.2. *Given population V, there exists a polynomial time algorithm that returns ℓ panels, $P_1, \ldots P_\ell$ of size k each, such that the global panel, i.e. $\cup_{i \in [\ell]} P_i$, is 6-PRF, each panel P_i is 26-PRF, and each $v \in V$ is included in panel P_i with probability k/n.*

However, despite significant effort, we were not able to generalize this result for the case where the panels have different sizes, and the following question remains open.

Open Question. Does there exist a distribution over a sequence of panels, $P_1, \ldots, P_\ell$, where P_i has size k_i such that each individual panel P_i satisfies $O(1)$-PRF and the global panel $\cup_{i \in [\ell]} P_i$ satisfies $O(1)$-PRF?

4 PREFIX AND PANEL LEVEL REPRESENTATION

In this section, we turn our attention to the setting where the goal is to ensure not only global representation but also *prefix representation.* That is, for every time step $t \in [\ell]$, we require that the cumulative panel $P_{\leq t} = \bigcup_{j=1}^{t} P_j$ provides appropriate representation. An ideal answer would be that there is a distribution over a sequence of panels, $P_1, \ldots, P_\ell$, where P_j has size k (or more generally k_j), such that each individual panel P_j and each prefix panel $P_{\leq t}$ satisfy a constant-factor approximation of PRF or at least PFC. Despite considerable effort, we are unable to prove whether such a distribution always exists, leaving a tantalizing question open.

Open Question. Does there exist a distribution over a sequence of panels, $P_1, \ldots, P_\ell$ such that each individual panel P_j and each prefix panel $P_{\leq t}$ satisfies $O(1)$-PRF or $O(1)$-PFC?

Instead, here we present a novel algorithm that provides weaker approximation guarantees with respect to individual and prefix representation for the case where each panel has size k. In particular, we show that there exists a polynomial-time algorithm that achieves

[4] Technically, Ebadian and Micha [14] establish this approximation under a slightly different notion of representation; however, by adapting their arguments, the guarantee carries over to PRF as well.

an approximation to PFC within each individual panel and each prefix panel, which grows exponentially with the number of panels ℓ, while also ensuring that each individual is selected to participate in one of the panels with equal probability.

Theorem 4.1. *There exists a polynomial time algorithm that returns a distribution over a sequence of ℓ panels $P_1, \ldots, P_\ell$ of size k each such that:*

- *For each $t \in [\ell]$, panel P_t satisfies $O(4^\ell)$-PFC;*
- *For each $t \in [\ell]$, prefix panel $P_{\leq t}$ satisfies $O(4^{\ell-t})$-PFC;*
- *Each individual is selected in $\cup_{t\in[\ell]} P_t$ with probability $\ell \cdot k/n$.*

At a high level, our main algorithm, NestedBasedRepresentation, operates in two phases. In the first phase, it constructs collections $\mathcal{G}^t$ of disjoint groups of individuals, each of size at least $n/t \cdot k$ for every $t \in [\ell]$. Each group in $\mathcal{G}^t$ represents a subset of the population that should be represented within the first t panels. The algorithm proceeds in reverse order, starting from $t = \ell$ and forming the collection of groups $\mathcal{G}^t$. As t decreases, it continues forming new groups while ensuring that previously constructed groups are preserved and hierarchically nested within the new groupings.

In the second phase, the algorithm traverses this hierarchy (visualized as a tree) starting from groups corresponding to smaller values of t (i.e., larger groups) and proceeding toward groups corresponding to larger values of t (i.e., smaller groups) that are nested within the previous ones. For each such path, the algorithm identifies the terminal group, i.e. a group corresponding to the larger value of t that does not have any other group nested and selects one individual from this terminal group as the path's representative. The algorithm then assigns these selected representatives to specific panels in sequential order, ensuring that the assignment respects both per-panel representation and prefix-level representation.

Below, we describe these two main phases separately in detail.

Tree Construction. First, we describe an algorithm called, ModifiedGreedyCapture which is a variation of Chen et al. [9]'s algorithm that takes population V, metric d, target panel size parameter K, and a partition $\mathcal{G}$ of the population into disjoint groups. The algorithm expands balls around each individual, initially marked incomplete. A ball captures a group in $\mathcal{G}$ only when it captures all its members; partial capture does not count. When an incomplete ball captures at least n/K individuals, it becomes *complete*, its captured groups are consolidated as a single group and added to $\mathcal{G}'$, then disregarded from further processing. The algorithm continues expanding all balls, disregarding any groups in $\mathcal{G}$ that are captured by complete balls. New groups are formed when incomplete balls capture total of at least n/K uncovered individuals. The process terminates when all groups are disregarded, returning $\mathcal{G}'$. By construction, each input group $G \in \mathcal{G}$ is either entirely contained within some output group $G' \in \mathcal{G}'$ or entirely excluded. For each group G created during the execution of ModifiedGreedyCapture, we denote by c_G its center, i.e., the point around which the ball was grown when the group was formed, and by r_G its radius, defined as the radius of the ball at the moment G was created. For complete algorithmic details, see full version of this paper [25].

NestedBasedRepresentation starts by constructing a tree with a root at $\mathcal{R}$, initially containing singleton groups $\{\{v\} : v \in V\}$ as children. Then, from $t = \ell$ down to 1, NestedBasedRepresentation calls ModifiedGreedyCapture, which at iteration t takes

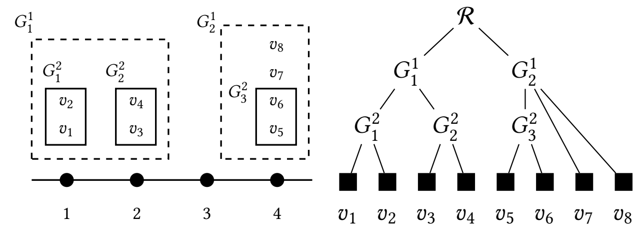

Figure 1: This figure illustrates the hierarchical group structure built in Phase 1 of NestedBasedRepresentation for eight individuals (v_1, v_2 at location 1; v_3, v_4 at 2; v_5–v_8 at 4) with $\ell = 2$ and $k = 2$. The left side shows two calls to ModifiedGreedyCapture: first with $t = 2$, which forms $\mathcal{G}^2 = G_1^2, G_2^2, G_3^2$ (leaving v_7 and v_8 ungrouped), and then with $t = 1$, applied to these groups and the remaining individuals to obtain $\mathcal{G}^1 = G_1^1, G_2^1$. The right side depicts the resulting tree, with groups as internal nodes and individuals as leaves.

$t \cdot k$ as parameter K and uses the current children of $\mathcal{R}$ as its group partition, outputting a collection $\mathcal{G}^t$ in which each group contains at least $n/t \cdot k$ individuals. After each execution, NestedBasedRepresentation updates the tree structure as follows. For each group $H \in \mathcal{G}^t$, all groups contained in H are removed from $\mathcal{R}$'s children, and H itself becomes a new child of $\mathcal{R}$. An example is illustrated at Figure 1. This tree structure ensures that if an individual v is selected, every group on the path from $\mathcal{R}$ to v is approximately represented. This is formalized in the following lemma, whose proof can be found in the full version [25].

Lemma 4.2. *Let P be a panel containing a representative from each group $G \in \mathcal{G}^j$. Then P satisfies $O(4^{\ell-j})$-PFC for population V and panel size $j \cdot k$.*

Panel Assignment. The algorithm ensures representation for both individual panels and prefix panels by utilizing the constructed tree structure $\mathcal{R}$ and the above lemma by enforcing the following two properties:

(a) Each individual panel contains one representative from each group $G \in \mathcal{G}^1$, ensuring approximate PFC representation for every panel.
(b) For each group $G \in \mathcal{G}^t$, there exists a representative from G among the first t panels, ensuring approximate PFC representation for every prefix up to time t.

To establish these properties, we first introduce the FindRepresentative algorithm, which the algorithm employs in the second phase. FindRepresentative takes as input the tree $\mathcal{R}$ and a target group G, corresponding to a node in the tree. Starting from G, it recursively selects a subgroup contained in G that belongs to some $\mathcal{G}^j$ with the smallest possible index j, continuing this process until it reaches a group consisting only of leaf nodes. Intuitively, FindRepresentative traces a path from the target group down to a leaf node which we call *trajectory*, always moving through subgroups that are required to be represented earlier. Upon reaching the terminal node, it arbitrarily selects $n/(\ell \cdot k)$ leaves (equivalently saying individuals) as the sampling group Q. Notice that by construction any formed group either contains at least $n/(\ell \cdot k)$ individuals or another group, implying that such a trajectory always exists.

Algorithm 1 FindRepresentative

Input: Target group G
Output: Modified group G' and sample group Q
if G contains a group $H \in \mathcal{G}^i$ for some $i \leq \ell$ **then**
  Let $H^* \in G$ be the group in $\mathcal{G}^i$ with minimum index i
  $H', Q \leftarrow$ FindRepresentative(H^*)
  Define modified target group $G' \leftarrow G \setminus \{H^*\} \cup H'$
  Return G', Q
else
  Let $Q \subseteq G$ be an arbitrarily selected subset of size $\frac{n}{\ell \cdot k}$
  Return $G \setminus Q, Q$
end if

Then, in two sub-phases, NestedBasedRepresentation enforces the two properties respectively, as follows.

Satisfying Property (a) - Panel Representation: This property requires that every group $G \in \mathcal{G}^1$ has a distinct representative in each of the ℓ panels. The algorithm addresses this requirement by processing groups individually through a sequential panel assignment procedure. For each group $G \in \mathcal{G}^1$, the algorithm iterates through panels $P_1, P_2, \ldots, P_\ell$ in order, seeking to assign one representative from G to each panel. This systematic approach is feasible because group G initially contains at least n/k individuals, while each execution of FindRepresentative consumes only $n/(\ell \cdot k)$ individuals through the sampling group. Consequently, the algorithm can perform exactly ℓ iterations to populate all required panels.

However, a complication arises when G contains subgroups that demand earlier representation. Specifically, when processing panel P_t, the algorithm first examines whether any subgroup $H \in G$ belongs to $\mathcal{G}^j$ for some $j < t$. Such a subgroup might violate the prefix representation guarantee when it is represented in panel P_t, since it requires representation in the first j panels. To resolve this conflict, the algorithm extracts this problematic subgroup from G and relocates it as direct children of the root $\mathcal{R}$, setting it aside for representation during the prefix sub-phase. Crucially, H has not been represented yet and we are simply deferring its assignment to ensure it gets placed in an appropriately early panel.

The key point is that even after removing H and its entire subtree, G still retains enough individuals to support ℓ executions of FindRepresentative, as we establish in our analysis in the full version [25]. After removing problematic subgroups, all remaining groups in G live in some $\mathcal{G}^i$ with $i \geq t$, allowing for the safe execution of FindRepresentative to obtain sampling group Q. The algorithm then samples an individual $v \in Q$, assigns v to panel P_t (where v represents all groups along execution trajectory of FindRepresentative), and updates the tree structure.

The tree update process permanently removes all sampled individuals in Q from future consideration, eliminates all intermediate groups along the trajectory except the root group G, and flattens the structure by having G directly contain the remaining children groups of these intermediate groups. After completing all ℓ panels for group G, the algorithm removes G from $\mathcal{R}$ and promotes any remaining children to become direct children of the root.

Satisfying Property (b) - Prefix Representation: Groups requiring prefix representation (i.e., those set aside in the previous step) plus any other children of $\mathcal{R}$, accumulate as children of $\mathcal{R}$. While the root contains groups ($\mathcal{R} \neq \emptyset$), the algorithm executes FindRepresentative starting from $\mathcal{R}$ to obtain sampling group Q. It samples an individual $v \in Q$ and assigns v to panel P_t where t is the minimum index such that $|P_t| < k$, and then it applies the same tree structure updates as before. As we show in the analysis, if a group on execution trajectory of FindRepresentativebelongs to $\mathcal{G}^t$ for some minimum value t, then there must be available capacity in the first t panels to accommodate the selected representative. The algorithm terminates when $\mathcal{R} = \emptyset$, at which point all groups have received appropriate representation in accordance with their prefix requirements.

Algorithm 2 NestedBasedRepresentation

Input: V, d, Groups $\mathcal{G}$, Parameter ℓ, Panel size k.
Output: Panels $P_1, \ldots, P_\ell$.
—Phase 1: Tree Formation—
Initialize tree $\mathcal{R}$ with children V, i.e., $\mathcal{R} \leftarrow V$
for $t = \ell$ **down to** 1 **do**
  Let $\mathcal{G}^t \leftarrow$ ModifiedGreedyCapture$(V, d, t \cdot k, \mathcal{R})$
  for $H \in \mathcal{G}^t$ **do**
    Remove groups in H from the root $\mathcal{R}$
    Add H as a child of the root $\mathcal{R}$
  end for
end for
—Phase 2: Panel Assignment—
—Phase 2.1: Panel Representation—
for $G \in \mathcal{G}^1$ **do**
  for $t = 1$ **to** ℓ **do**
    while there is a group $H \in G \cap \mathcal{G}^i$ for some $i < t$ **do**
      $G \leftarrow G \setminus \{H\}$ and $\mathcal{R} \leftarrow \mathcal{R} \cup \{H\}$.
    end while
    $G', Q \leftarrow$ FindRepresentative(G).
    Sample $v \in Q$ uniformly, assign to P_t
    Update $G \leftarrow G'$
  end for
  Remove G from $\mathcal{R}$, add its children to $\mathcal{R}$
end for
—Phase 2.2: Prefix Representation—
while $\mathcal{R} \neq \emptyset$ **do**
  $\mathcal{R}', Q \leftarrow$ FindRepresentative$(\mathcal{R})$
  Sample $v \in Q$, assign to P_t with min t s.t. $|P_t| < k$
  Update the root $\mathcal{R} \leftarrow \mathcal{R}'$
end while
return $P_1, \ldots, P_\ell$

These two properties together with Lemma 4.2 imply the representation guarantees stated in Theorem 4.1. For a comprehensive discussion and the full proofs of Theorem 4.1 and Theorem 4.2, please refer to the full version [25].

5 PREFIX LEVEL REPRESENTATION

In the previous section, we show that by nesting groups across panel sizes, we can ensure approximate PFC representation for both prefixes and individual panels. However, the approximation factor increases exponentially with the number of panels ℓ.

In this section, we adopt a different approach and show that a *constant*-factor approximation to PFC can be achieved with respect to every prefix panel $P_{\leq t}$, though at the cost of abandoning representation guarantees for individual panels. The key idea is to replace the nested-group construction with a distinct family of groups $\mathcal{G}^t$ for each prefix $t \in [\ell]$, capturing the groups that must be represented within the first t panels. Unlike the nested setting, this method does not require that for overlapping groups $G \in \mathcal{G}^t$ and $G' \in \mathcal{G}^{t'}$ with $t < t'$, the group from the larger prefix size (G') be fully contained in the group from the smaller prefix size (G), i.e., $G' \subseteq G$. Instead, the algorithm links groups across different prefix sizes into *chains*, i.e., sequences of the form $G^t \to G^{t+1} \to \cdots \to G^\ell$, with $G^j \in \mathcal{G}^j$, where each group overlaps with some of its predecessors and has radius no larger than those it overlaps with. Representing the final group in such a chain then suffices to approximate the representation of all groups in the sequence, ensuring that coverage for the last group automatically yields coverage for every earlier one.

Note that simply linking overlapping groups with radii no larger than their predecessors is not enough, as the approximation error can accumulate at each step, leading to a bound that grows linearly with ℓ. To prevent this, we build the chains in a more careful way.

Theorem 5.1. *There exists a polynomial time algorithm that returns a distribution over a sequence of ℓ panels $P_1, \ldots, P_\ell$, where panel P_t has size k_t, such that:*

- *For each $t \in [\ell]$, $P_{\leq t} = \cup_{j=1}^{t} P_j$ satisfies $O(1)$-PFC;*
- *Each individual is selected in $\cup_{t \in [\ell]} P_t$ with probability $\sum_{t=1}^{\ell} k_t / n$.*

Description of the Algorithm. Our algorithm, called ChainBasedRepresentation consists of three main phases.

In the first phase, for each $t \in [\ell]$, the algorithm runs ModifiedGreedyCapture on the population V with metric d, panel size $\sum_{j=1}^{t} k_j$, and with each individual initially assigned to their own singleton group (i.e., there is no pre-existing grouping that must be nested). This yields a collection of groups $\mathcal{G}^t$, representing the groups that should be represented among the first t panels.

In the subsequent phases, the algorithm aims to build panels $P_1, \ldots, P_\ell$ so that, for every prefix $P_{\leq t}$ and every group $G^t \in \mathcal{G}^t$, there exists an individual in $P_{\leq t}$ whose distance to c_{G^t} is at most αr_{G^t} for a universal constant $\alpha \geq 1$. We say that a group is *covered* once this condition holds, and the goal is to cover all groups in $\bigcup_{t \in [\ell]} \mathcal{G}^t$. To achieve this, the algorithm relies on the following key property of the groups generated by ModifiedGreedyCapture, which is formalized in the following lemma. The proof of this lemma can be found in the full version [25].

Lemma 5.2. *Let V be a population in a metric space with distance function d, and consider panel sizes $k_1 \leq k_2$. Let $\mathcal{G}$ and $\mathcal{H}$ be the outputs of ModifiedGreedyCapture with panel sizes k_1 and k_2, respectively. For every group $G \in \mathcal{G}$ with center c_G and radius r_G, there exists a group $H \in \mathcal{H}$ with center c_H and radius r_H such that $d(c_H, c_G) \leq 2 \cdot r_G$ and $r_H \leq r_G$.*

The above lemma indicates that for every group $G^t \in \mathcal{G}^t$ and every $j \geq t$, there exists a group $G^j \in \mathcal{G}^j$ such that (i) G^t and G^j overlap and (ii) the radius of G^j is at most equal to the radius of G^t. Essentially this means that by (approximately) representing group G^j, then G^t is also (approximately) represented.

By exploiting this lemma, the second phase proceeds as follows. Initially, all the groups are marked as uncovered. For each group G^t, the algorithm attempts to find a group in $\mathcal{G}^\ell$ such that once this group is represented, G^t is also approximately represented. More precisely, the algorithm starts from the set $\mathcal{G}^t$ with uncovered groups of smallest index and picks an arbitrary uncovered group G^t. It then calls the subroutine ConstructChain, which tries to construct a chain of groups $G^t \to \ldots \to G^\ell$, with each $G^j \in \mathcal{G}^j$, ensuring that selecting an individual from G^ℓ represents every group in the chain within a constant approximation. The subroutine begins by marking G^t as the *anchor* and uses it to guide the chain extension. For each $j \in \{t+1, \ldots, \ell\}$, it finds a group $G^j \in \mathcal{G}^j$ that satisfies the conditions of Lemma 5.2 relative to the anchor. If such a group is found and is uncovered, it is added to the chain; furthermore, if the radius of G^j is at most half that of the current anchor, G^j is promoted as the new anchor. This process continues until either a group in $\mathcal{G}^\ell$ is reached, at which point the chain construction succeeds, or a covered group is encountered, in which case ConstructChain returns an unsuccessful chain consisting only of G^t. ChainBasedRepresentation marks all groups in the returned chain as covered and repeats the process, always by calling ConstructChain on an uncovered group in $\mathcal{G}^j$ with the smallest index j. The phase ends once all groups are marked as covered.

Algorithm 3 ChainBasedRepresentation

Input: V, d, Panel sizes $k_1, \ldots, k_\ell$
Output: A distribution over a sequence of panels $P_1, \ldots, P_\ell$
— Phase 1: Construct Groups —
for level $t = 1$ to ℓ **do**
  $\mathcal{G}^t \leftarrow$ ModifiedGreedyCapture$(V, d, \sum_{j=1}^{t} k_j, \cup_{v \in V}\{v\})$
end for
—Phase 2: Build Chains and Assign Priorities —
for $t = 1$ to $\ell - 1$ **do**
  for each uncovered $G^t \in \mathcal{G}^t$ **do**
    status, $T \leftarrow$ ConstructChain $(G^t, \{\mathcal{G}^1, \ldots, \mathcal{G}^\ell\}, V, d)$
    Mark all groups in chain T as covered
    If status is succeed, assign to the last group in T priority t
  end for
end for
—Phase 3: Sampling and Assignment to Panels—
for each $G^\ell \in \mathcal{G}^\ell$ **do**
  Sample v from G^ℓ uniformly at random
  Assign v the priority label of its group G^ℓ
end for
Sample $\sum_{t=1}^{\ell} k_t - |\mathcal{G}^\ell|$ representatives uniformly at random from $V \setminus \bigcup_{G^\ell \in \mathcal{G}^\ell} G^\ell$, and assign each a priority label of ℓ.
Assign representatives to panels by ordering them in increasing order of priority labels, and placing each into the earliest panel P_t with available capacity (i.e., the smallest such index t)
return $P_1, \ldots, P_\ell$

In the third phase, the algorithm uses the previously constructed chains to generate a distribution over panels $P_1, \ldots, P_\ell$. To ensure individual fairness, it begins by sampling $\sum_{t=1}^{\ell} k_t$ individuals from the population as follows: one individual is selected uniformly at random from each group in $\mathcal{G}^\ell$, and the remaining $\sum_{t=1}^{\ell} k_t - |\mathcal{G}^\ell|$

individuals are sampled uniformly at random from the rest of the population[5] Next, to ensure proportional representation, the algorithm assigns each sampled individual a priority label based on the chains created in the previous phase. In particular, if an individual is sampled from a group G^ℓ that belongs to some chain $G^t \to \ldots \to G^\ell$, the individual is assigned the index t, corresponding to the head of the chain. If the individual is not sampled from any such group, they are assigned the priority label ℓ. Finally, the algorithm allocates individuals to panels sequentially, from P_1 to P_ℓ, in increasing order of their assigned priority labels. Intuitively, this gives higher priority to individuals representing groups that must be covered earlier, ensuring that such representatives appear sooner in the panel sequence.

Algorithm 4 ConstructChain

```
Input: Target group G^t ∈ 𝒢^t, {𝒢^1, ..., 𝒢^ℓ}, V, d
Output: status (succeed when a new chain was constructed
    and failed otherwise) and chain T
Initialize the chain and anchor: T ← G^t and H ← G^t
for j = t + 1 to ℓ do
    Let G^j ∈ 𝒢^j be such that d(c_H, c_{G^j}) ≤ 2 · r_H and r_{G^j} ≤ r_H
    if G^j is already covered then
        Remove from T any group except for G^t
        return fail, T
    end if
    Add G^j to the end of chain T
    if r_{G^j} < r_H/2 (radius shrinks significantly) then
        Update anchor: H ← G^j
    end if
end for
return succeed, T
```

Before proving Theorem 5.1, we state two more lemmas.

Lemma 5.3. *For every group $G^t \in \mathcal{G}^t$, there exists a representative v among the first t panels, such that $d(c_{G^t}, v) \leq 16 \cdot r_{G^t}$.*

Proof Sketch. The key idea is that whenever the anchor of a chain changes, the radius of the new anchor decreases by a factor of two. As a result, the radii decrease geometrically, ensuring that the distance between the first and final anchor centers is bounded by a constant multiple of the first anchor's radius. For any group $G^t \in \mathcal{G}^t$, we consider two cases:

Case 1: G^t appears in a chain. Suppose $G^\ell \in \mathcal{G}^\ell$ is the endpoint of the chain. If G^t serves as an anchor, the bound above applies directly, bounding the distance from c_{G^t} to any individual in G^ℓ by a constant factor of r_{G^t}. If G^t is not an anchor, its radius is no larger than that of the current anchor when G^t joins the chain, and the distance between c_{G^t} and a representative of G^ℓ is again bounded by a constant multiple of r_{G^t}.

Case 2: G^t appears in no chain. The algorithm attempts to start a chain from G^t, eventually reaching some $G^j \in \mathcal{G}^j$ which is already covered. Using our chain length bound, G^t lies within constant times its radius from the center of G^j and G^j's radius cannot be larger than two times the radius of G^t. Since G^j is covered, by Case 1 it will be represented by some representative in G^ℓ that is not farther away than a constant factor of r_{G^j} from G^j's centers. Combining these observations, the representative from G^ℓ also provides a constant representation to G^t.

The final component of the algorithm ensures representation of entitled groups in each $\mathcal{G}^t$ among the first t panels through a counting argument on group sizes. □

The complete version of this sketch and proof of the following lemma can be found in the full version [25].

Lemma 5.4. *Let V be a population and k be a panel size. Consider the collection of groups $\mathcal{G}$ formed by ModifiedGreedyCapture over V with panel size k and each individual initially assigned to their own group. If $P \subseteq V$ is a panel such that, for every $G \in \mathcal{G}$, there exists $p \in P$ with $d(p, c_G) \leq \alpha \cdot r_G$, then P satisfies $(\alpha + 3)$-PFC.*

Proof of Theorem 5.1. We establish both properties of the theorem. For the first property, we see that each prefix panel $P_{\leq t}$ satisfies 19-PFC by combining Theorem 5.3 and Theorem 5.4.

For the second property, we analyze the selection probability for each individual. If an individual appears in a group $G \in \mathcal{G}^\ell$, then it appears in some panel with probability $\sum_{t=1}^{\ell} k_t / n$ since each such group $\mathcal{G}^\ell$ contains exactly $n / \sum_{t=1}^{\ell} k_t$ individuals, and one of them is chosen uniformly at random. For individuals not covered by groups in $\mathcal{G}^\ell$, the number of such individuals equals $n / \sum_{t=1}^{\ell} k_t$ times the number of empty panel seats, and they are sampled with equal probability to fill the remaining positions. Hence, each individual appears in some panel with probability $\sum_{t=1}^{\ell} k_t / n$. □

It remains an open question whether one can design a distribution over a sequence of panels such that every prefix satisfies a constant-factor approximation of the proportionality notion PRF, which, in contrast to PFC, requires that each subset of individuals receive not just one representative, but a number of representatives proportional to its size.

Open Question. Does there exist a distribution over a sequence of panels, $P_1, \ldots, P_\ell$ where each prefix panel $P_{\leq t}$ satisfies $O(1)$-PRF?

6 CONCLUSION

Permanent citizens' assemblies represent a promising democratic innovation by ensuring that minority voices can persist over time, even when too small to warrant representation in any single panel. We formalize this challenge as a problem of *temporal sortition* in metric spaces and make progress on it by proposing three algorithms, each providing different guarantees for panel-level and prefix-level representation.

Beyond the questions highlighted above, several intriguing directions remain open. For instance, can we ensure representation not just for prefix panels, but for *any* consecutive subsequence of panels? Additionally, is it possible to achieve meaningful guarantees without knowing in advance the total number or sizes of panels?

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their detailed and constructive feedback, which helped improve the presentation of this paper.

[5]When n is divisible by $\sum_{t=1}^{\ell} k_t$, this can be achieved by applying standard dependent rounding techniques. Otherwise, we first create groups of size exactly $n / \sum_{t=1}^{\ell} k_t$ by fractional assignment, and then apply Birkhoff's decomposition, as in [14].

REFERENCES

[1] Alberto Alemanno. 2022. *Towards a Permanent Citizens' Participatory Mechanism in the EU.* Technical Report PE 735.927. European Parliament, Policy Department for Citizens' Rights and Constitutional Affairs. https://www.europarl.europa.eu/RegData/etudes/STUD/2022/735927/IPOL_STU(2022)735927_EN.pdf

[2] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[3] Haris Aziz, Barton E. Lee, Sean Morota Chu, and Jeremy Vollen. 2024. Proportionally Representative Clustering. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE).*

[4] Carmel Baharav and Bailey Flanigan. 2024. Fair, Manipulation-Robust, and Transparent Sortition. In *Proceedings of the 25th ACM Conference on Economics and Computation.* 756–775.

[5] Evripidis Bampis, Bruno Escoffier, and Sasa Mladenovic. 2018. Fair Resource Allocation Over Time. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* (Stockholm, Sweden). Richland, SC, 766–773.

[6] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. 2021. Justified Representation for Perpetual Voting. *IEEE Access* 9 (2021), 96598–96612. https://doi.org/10.1109/ACCESS.2021.3095087

[7] Ioannis Caragiannis, Evi Micha, and Jannik Peters. 2024. Can a Few Decide for Many? The Metric Distortion of Sortition. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, Vol. 235. 5660–5679.

[8] Ioannis Caragiannis and Shivika Narang. 2023. Repeatedly Matching Items to Agents Fairly and Efficiently. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory (SAGT)* (Egham, United Kingdom). Berlin, Heidelberg, 347–364. https://doi.org/10.1007/978-3-031-43254-5_20

[9] Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML).* 1032–1041.

[10] Tamal K. Dey, Alfred Rossi, and Anastasios Sidiropoulos. 2017. Temporal Clustering. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA)*, Vol. 87. Dagstuhl, Germany, 34:1–34:14. https://doi.org/10.4230/LIPIcs.ESA.2017.34

[11] Tamal K. Dey, Alfred Rossi, and Anastasios Sidiropoulos. 2017. Temporal Hierarchical Clustering. In *Proceedings of the 28th Annual International Symposium on Algorithms and Computation (ISAAC)*, Vol. 92. Dagstuhl, Germany, 28:1–28:12. https://doi.org/10.4230/LIPIcs.ISAAC.2017.28

[12] Virginie Do, Matthieu Hervouin, Jérôme Lang, and Piotr Skowron. 2022. Online Approval Committee Elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI).* 251–257. https://doi.org/10.24963/ijcai.2022/36 Main Track.

[13] S. Ebadian, G. Kehne, E. Micha, A. D. Procaccia, and N. Shah. 2022. Is Sortition Both Representative and Fair?. In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS).* 25720–25731.

[14] S. Ebadian and E. Micha. 2025. Boosting Sortition via Proportional Representation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).* 667–675.

[15] Edith Elkind, Svetlana Obraztsova, Jannik Peters, and Nicholas Teh. 2025. Verifying Proportionality in Temporal Voting. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 13 (Apr. 2025), 13805–13813. https://doi.org/10.1609/aaai.v39i13.33509

[16] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Temporal fairness in multiwinner voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI).* Article 2525, 8 pages. https://doi.org/10.1609/aaai.v38i20.30273

[17] B. Flanigan, P. Gölz, A. Gupta, B. Hennig, and A. D. Procaccia. 2021. Fair Algorithms for Selecting Citizens' Assemblies. *Nature* 596 (2021), 548–552.

[18] B. Flanigan, P. Gölz, A. Gupta, and A. D. Procaccia. 2020. Neutralizing Self-Selection Bias in Sampling for Sortition. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS).* 6528–6539.

[19] Bailey Flanigan, Jennifer Liang, Ariel D Procaccia, and Sven Wang. 2024. Manipulation-robust selection of citizens' assemblies. In *Proceedings of the aaai conference on artificial intelligence*, Vol. 38. 9696–9703.

[20] Sreenivas Gollapudi, Kostas Kollias, and Benjamin Plaut. 2020. Almost Envy-Free Repeated Matching in Two-Sided Markets. In *Web and Internet Economics*, Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta (Eds.). Cham, 3–16.

[21] Daniel Halpern, Ariel D Procaccia, Ehud Shapiro, and Nimrod Talmon. 2025. Federated assemblies. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 39. 13897–13904.

[22] Ayumi Igarashi, Martin Lackner, Oliviero Nardi, and Arianna Novaro. 2024. Repeated fair allocation of indivisible items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI).* Article 1089, 9 pages. https://doi.org/10.1609/aaai.v38i9.28837

[23] Christopher Jung, Sampath Kannan, and Neil Lutz. 2020. Service in Your Neighborhood: Fairness in Center Location. In *1st Symposium on Foundations of Responsible Computing (FORC 2020)*, Aaron Roth (Ed.), Vol. 156. Dagstuhl, Germany, 5:1–5:15. https://doi.org/10.4230/LIPIcs.FORC.2020.5

[24] Yusuf Kalayci, David Kempe, and Vikram Kher. 2024. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 38. 9815–9823.

[25] Yusuf Hakan Kalayci and Evi Micha. 2026. Temporal Panel Selection in Ongoing Citizens' Assemblies. arXiv:2602.16194 [cs.GT] https://arxiv.org/abs/2602.16194

[26] Leon Kellerhals and Jannik Peters. 2025. Proportional fairness in clustering: a social choice perspective. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS)* (Vancouver, BC, Canada) *(NIPS '24).* Red Hook, NY, USA, Article 3534, 19 pages.

[27] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)* 34, 02 (Apr. 2020), 2103–2110. https://doi.org/10.1609/aaai.v34i02.5584

[28] Martin Lackner and Jan Maly. 2023. Proportional Decisions in Perpetual Voting. *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)* 37, 5 (Jun. 2023), 5722–5729. https://doi.org/10.1609/aaai.v37i5.25710

[29] Martin Lackner and Piotr Skowron. 2022. Approval-based committee voting. In *Multi-Winner Voting with Approval Preferences*. Springer, 1–7.

[30] E. Micha and N. Shah. 2020. Proportionally fair clustering revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP).* 85:1–85:16.

[31] Christoph Niessen and Min Reuchamps. 2019. Designing a Permanent Deliberative Citizens' Assembly: The Ostbelgien Model in Belgium. (2019).

[32] Peter Stone. 2011. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press.

[33] Yohai Trabelsi, Abhijin Adiga, Sarit Kraus, S. S. Ravi, and Daniel J. Rosenkrantz. 2023. Resource sharing through multi-round matchings. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI).* Article 1311, 10 pages. https://doi.org/10.1609/aaai.v37i10.26380

[34] David Van Reybrouck. 2016. *Against Elections: The Case for Democracy.* The Bodley Head / Random House.