CSC2555 Summer'26
Additional Sample Project Ideas
WARNING: LLM-GENERATED. DETAILS MAY BE INACCURATE.

Note: These are just a few sample ideas. Try to find an idea similar to these but in your own research area; you'll have more fun working on a project in your own area. If you are having a really difficult time, then you can claim one of these ideas on this Google sheet (please do not overwrite another group's claim). The sample idea is presented at a high level; in your proposal, you'll have to detail out the exact goals.

Sample Idea 1: Connected Envy-Free Cake-Cutting for Interval Valuations

Assume each agent values a single interval $[l_i, u_i]$ uniformly, and require each agent to receive one connected piece. Connected envy-free allocations always exist under standard assumptions, but no finite Robertson–Webb protocol can find one for general valuations and $n \geq 3$ [22, 23]. For single-interval valuations, an exact polynomial-time algorithm is known under an ordering condition that excludes strict containment; without it, a polynomial-time algorithm achieving additive envy at most $1/4$ is known [5, 13].

Main question: Can an exact connected envy-free allocation be computed in polynomial time for unrestricted single-interval valuations?

Possible directions:

- design a greedy algorithm based on leftmost or rightmost cuts;
- exploit the favorite-piece graph induced by the cuts;
- study laminar intervals, bounded overlap or containment depth, or a bounded number of containments;
- study the case in which the left-to-right agent order is fixed.

A successful project could solve one nontrivial restriction, give an FPT algorithm, prove a useful structural invariant, or find counterexamples to natural algorithms.

Sample Idea 2: Fair and Efficient Allocation of Restricted Mixed Manna

For additive mixed manna, whether an EF1 and Pareto-optimal (PO) allocation always exists remains open, even for three agents [9]. Positive results are known for two agents [7], pure goods [10], pure chores [20], and structured domains such as identical or common ternary valuations [6].

Main question: Which simple restrictions beyond these cases guarantee EF1+PO, and permit efficient computation?

Candidate restrictions, individually or in combination:

- three agents;
- one or two goods, or one or two chores;

- all but one or two items forming a pure-goods or pure-chores instance;
- two agent types, or valuations differing on only one or two items;
- values in $\{-1, 0, 1\}$ or $\{-a_i, 0, b_i\}$;
- one or two items whose signs differ across agents;
- sparse agent–item incidence graphs, such as paths or trees;
- common rankings of goods and chores;
- $m \leq n+1$ or $m \leq n+2$.

Small-instance search can test candidate algorithms, exchange rules, and stronger structural conjectures. A successful project could prove EF1+PO for one nontrivial restriction, give an FPT algorithm, or explain why natural extensions of known algorithms fail.

Sample Idea 3: Proportional Participatory Budgeting with Pairwise Project Interactions

Let

$$u_i(W) = \sum_{g \in W} v_{i,g} + \sum_{\{g,h\} \subseteq W} w_{i,g,h},$$

where positive and negative pairwise weights encode complements and substitutes. Existing work studies richer interaction models, welfare maximization, and proportionality under substitution dependencies [17, 14, 12]. Standard PB proportionality axioms such as EJR and core stability are well understood for additive utilities [21].

Let G contain edge $\{g, h\}$ whenever some $w_{i,g,h} \neq 0$.

Main question: How should JR, PJR, EJR, FJR, or core-type guarantees extend to pairwise interactions, and what is achievable for restricted G?

Two possible frameworks are:

- retain standard approval-based cohesiveness, but measure representation using u_i;
- call S cohesive at level ℓ if some bundle T with $c(T) \leq |S|B/n$ gives every $i \in S$ utility at least ℓ, and require an EJR-style representative in S to receive utility at least ℓ.

Any extension should recover the standard axiom when all interaction weights vanish.

Possible questions:

- existence and implication relations among the resulting axioms;
- guarantees of PAV or the Method of Equal Shares;
- complexity of finding an outcome or a violating coalition;
- matchings, paths, trees, or bounded-treewidth interaction graphs;

- complements only, substitutes only, or voter-independent weights.

A natural starting point is unit-cost projects, binary base values, and a matching interaction graph.

Sample Idea 4: Proportional Representation under Uncertain Approvals

Approval uncertainty may be represented by probabilities $p_{i,c}$, incomplete ballots, or an uncertainty set of plausible profiles. Existing work studies possible and necessary proportionality under incomplete information and committees maximizing the probability of JR under several uncertainty models [16, 8]. Recent work also studies the sample complexity of obtaining JR-family guarantees from sampled voters [19].

Main question: What stronger robustness guarantees are possible for PJR, EJR, or related axioms under noisy or incomplete approvals?

Possible directions:

- maximize the probability of satisfying PJR or EJR;
- require an axiom for every profile in a structured uncertainty set;
- bound how many approval changes can destroy proportionality;
- derive sample-complexity or perturbation bounds;
- study independent flips, party-list profiles, or a constant committee size.

A successful project could give an algorithm or complexity result for one uncertainty model, derive a robustness bound for a standard rule, or compare rules empirically under controlled noise.

Sample Idea 5: Intersectional Fairness Auditing without Enumerating All Groups

A classifier may satisfy a fairness condition for each protected attribute separately while violating it on an intersectional subgroup. Auditing an exponential family of structured subgroups is closely connected to agnostic learning and is computationally hard in general [18].

Fix a classifier, dataset, fairness metric, minimum group size, and subgroup class—for example, conjunctions of at most d conditions or leaves of a small decision tree.

Main question: How can one efficiently find a subgroup with a statistically certified fairness violation?

Possible directions:

- derive uniform confidence bounds accounting for multiple testing;
- compare exact integer-programming and heuristic search;
- exploit bounded conjunction width, tree depth, or few attributes;
- handle statistical parity, false-positive-rate disparity, or calibration error;
- distinguish genuine violations from artifacts of adaptively searching many groups.

A successful project could provide a valid auditor for one subgroup class, a tractability result under a structural restriction, or an empirical comparison with finite-sample guarantees.

Sample Idea 6: Proportional Fairness When Agents Care About All Centers

Let the cost of agent i for a nonempty center set C be

$$d_{\text{avg}}(i, C) = \frac{1}{|C|} \sum_{c \in C} d(i, c).$$

Optimization for the corresponding sum-of-distances objective has been studied [15], while core-style proportionality has been studied for closest-center and q-th-closest-center costs [11].

Call a k-center set C an α-core outcome if there are no $\ell \in [k]$, group S with $|S| \geq \ell n/k$, and set C' with $1 \leq |C'| \leq \ell$ such that

$$\alpha \, d_{\text{avg}}(i, C') < d_{\text{avg}}(i, C) \qquad \text{for every } i \in S.$$

Averaging is essential: an unnormalized sum mechanically favors deviations using fewer centers.

Main question: What approximate-core guarantees are possible for average-distance costs?

Possible directions:

- existence of an exact core, even on a line;
- whether checking only single-center deviations suffices;
- the best universal approximation factor;
- algorithms on lines, trees, or for constant k;
- welfare loss from imposing proportional fairness.

A successful project could prove an existence or approximation result, give an algorithm for a restricted metric, or show that a natural clustering algorithm has no bounded guarantee.

Sample Idea 7: Proportional Representation in LLM-Selected Committees

LLMs are increasingly asked to select a small set of options for a group: news articles, policy proposals, speakers, courses, or recommendations. Such outputs can be viewed as committees selected from heterogeneous user preferences. Recent work has evaluated LLM allocations using fair-division notions [1]; this project applies the same methodology to multiwinner voting. Related work on generative social choice uses LLMs to generate representative statements, but does not systematically evaluate off-the-shelf LLM committee selection using standard multiwinner axioms [2].

Main question: How well do LLM-selected committees satisfy proportional representation?

Possible directions:

- generate approval profiles from party-list, spatial, or preference-mixture models;

- ask several LLMs to select a committee from the profile, with and without explicit fairness instructions;
- evaluate JR, PJR, EJR, proportionality degree, welfare, and coverage;
- compare against AV, PAV, Chamberlin–Courant, and the Method of Equal Shares;
- test sensitivity to voter and candidate order, natural-language framing, explanations, and profile size.

A successful project could identify systematic representation failures, determine which prompts or models perform best, or design a simple postprocessing or prompting method that improves proportionality with little loss of welfare.

Sample Idea 8: Selecting Diverse and Fair Classifier Ensembles

An ML pipeline often produces many candidate classifiers by varying the model family, training sample, hyperparameters, or fairness intervention. Rather than select one model, one may choose a committee of classifiers and combine their predictions by voting. Ensemble selection traditionally emphasizes accuracy and predictive diversity [3]. Multiwinner voting instead offers rules designed to represent heterogeneous voters [4].

Treat candidate classifiers as candidates and data points, demographic groups, or application environments as voters. A voter approves classifiers that meet a chosen performance threshold on that voter or group. The selected committee can then be evaluated using JR, PJR, EJR, coverage, accuracy, and standard group-fairness metrics.

Main question: Do proportional multiwinner rules select classifier ensembles that perform more evenly across heterogeneous populations than standard ensemble-selection methods?

Possible directions:

- compare top-k accuracy, greedy diversity, AV, PAV, Chamberlin–Courant, and the Method of Equal Shares;
- vary whether voters are individuals, demographic groups, data clusters, or deployment environments;
- compare approval thresholds based on correctness, calibrated loss, or performance relative to each voter's best classifier;
- evaluate ensemble accuracy, worst-group accuracy, equalized-odds gaps, coverage, and proportional-representation axioms;
- study robustness under distribution shift or underrepresented validation groups.

A successful project could establish when proportional selection improves worst-group performance, characterize its accuracy–representation tradeoff, or propose a hybrid ensemble rule that performs better than both standard ML and voting-based baselines.

References

[1] Benjamin Cookson, Soroush Ebadian, and Nisarg Shah. Fairness perceptions of large language models. In *AAAI*, 2026.

[2] Sara Fish, Paul Gölz, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. Generative social choice. *Journal of the ACM*, 2025.

[3] Nai Jin Sean Lim and Robert J. Durrant. A diversity-aware model for majority vote ensemble accuracy. In *AISTATS*, 2020.

[4] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2023.

[5] Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahmad S. Tajik. Envy-free mechanisms with minimum number of cuts. In *AAAI*, 2017.

[6] Martin Aleksandrov and Toby Walsh. Two algorithms for additive and fair division of mixed manna. arXiv:2007.04129, 2020.

[7] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In *IJCAI*, 2019.

[8] Haris Aziz, Venkateswara Rao Kagita, Baharak Rastegari, and Mashbat Suzuki. Approval-based committee voting under uncertainty. arXiv:2407.19391, 2024.

[9] Siddharth Barman, Vishwa Prakash HV, Aditi Sethia, and Mashbat Suzuki. Fair and efficient allocation of indivisible mixed manna. arXiv:2507.03946, 2025.

[10] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3), 2019.

[11] Soroush Ebadian and Evi Micha. Boosting sortition via proportional representation. arXiv:2406.00913, 2024.

[12] Roy Fairstein, Reshef Meir, and Kobi Gal. Proportional participatory budgeting with projects interaction. arXiv:2106.05360v2, 2024.

[13] Paul W. Goldberg, Alexandros Hollender, and Warut Suksompong. Contiguous cake cutting: Hardness results and approximation algorithms. *Journal of Artificial Intelligence Research*, 69:109–141, 2020.

[14] Mohak Goyal, Sahasrajit Sarmasarkar, and Ashish Goel. A mechanism for participatory budgeting with funding constraints and project interactions. arXiv:2305.11296, 2023.

[15] Yue Han and Elliot Anshelevich. Compatibility of max and sum objectives for committee selection and k-facility location. arXiv:2507.17063, 2025.

[16] Aviram Imber, Jonas Israel, Markus Brill, and Benny Kimelfeld. Approval-based committee voting under incomplete information. In *AAAI*, 2022.

[17] Pallavi Jain, Krzysztof Sornat, Nimrod Talmon, and Meirav Zehavi. Participatory budgeting with project interactions. In *IJCAI*, 2020.

[18] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *ICML*, 2018.

[19] Gregory Kehne. Proportionality from sampled approvals. arXiv:2606.10446, 2026.

[20] Ryoga Mahara. Existence of fair and efficient allocation of indivisible chores. arXiv:2507.09544, 2025.

[21] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *NeurIPS*, 2021.

[22] Walter Stromquist. How to cut a cake fairly. *American Mathematical Monthly*, 87(8):640–644, 1980.

[23] Walter Stromquist. Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics*, 15:R11, 2008.