CSC2421 Spring'24
Sample Project Ideas

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: Relaxed Fairness Notions in Cake-Cutting

As mentioned in class, envy-freeness is difficult to achieve in cake-cutting (the best known algorithm requires too many queries), but proportionality is very easy to achieve. Can we achieve fairness notions in between the two?

For example, given an ordering π of the n agents (where $\pi(i)$ is rank of agent i in the ordering), an allocation A is called *forward envy-free* for π if $v_i(A_i) \geq v_i(A_k)$ for all i, k where $\pi(i) < \pi(k)$. In other words, agents should not envy those who appear later in the ordering.

Clearly, this is a relaxed version of envy-freeness. Without additional constraints, it is too relaxed: we can just give the entire cake to the first agent in the ordering.

Research Question: How many queries do we need in the Robertson-Webb model to find an allocation that is proportional and forward envy-free with respect to a given ranking of the agents?

Instead of a ranking, you can consider other relations (e.g., a star graph where you just don't want a given agent to envy any other agents, etc).

Sample Idea 2: Fair Division with Probabilistic Preferences

We do not the exact values of agents for goods, but we know that the value of agent i for good g is drawn from a known distribution $D_{i,g}$, independently of all other values. Given these distributions, we want to find "the most fair" allocations, e.g., those that maximize the probability of being EF, or maximize the probability of being EF1, or maximize the probability of being EF+PO, etc.

The simplest case is where each distribution $D_{i,g}$ is just a Bernoulli distribution with value 1 with some probability $p_{i,g}$ and 0 otherwise.

Sample Idea 3: Approval-Based Committee Selection

Many of the rules we saw in class were quite sophisticated. In some domains, one may prefer to use simpler rules. Approval voting (where you select the k candidates with the highest total number of approvals) is one such simple rule. Each voter knows exactly what happens when they approve a candidate: the candidate's score goes up by one.

One criticism of this rule is that it gives more "voice" to voters who approve many candidates compared to those who approve few. One could define a more general family of rules, where whenever a voter approves r candidates, the score of each of those candidates would increase by $f(r)$ for some function f. Then, we still tally up the scores and choose the k candidates with the highest scores.

Approval voting would correspond to $f(r) = 1$ for all r, another extreme choice would be $f(r) = 1/r$ for all r (a voter can equally "spread" a total approval of 1 between as many candidates as she wishes), and one can also consider choices that are intermediate (e.g., $f(r) = 1/\sqrt{r}$) or even more extreme (e.g., $f(r) = 1/r^2$). Which functions f give better fairness guarantees, either theoretically or empirically?

Sample Idea 4: Fair and Efficient Course Allocation with Constraints

In class, we saw the open problem of fair and efficient course allocation. Say each student i has a value $v_{i,c}$ for each course c. At most q_c students can be assigned to each course c, and each student can be assigned to at most q^* courses (note that the student capacities are identical). Does there always exist a matching of students to courses that is both EF1 and PO?

This is a nice idealized problem, but it fails to capture a number of real-world considerations.

- *Flexible student capacities:* Each student may have a *lower* and upper bound on how many courses they can/must take, and these bounds may be different for different students.
- *Flexible course capacities:* Each course may have a lower and upper bound on how many students it can have.
- *Qualifications:* Some students may not be qualified to take some courses, e.g., due to not meeting some pre-requisites.
- *Scheduling constraints:* Some pairs of courses cannot be taken together due to having overlapping lecture timings.

Note that some of these modifications can make students "unequal", requiring a change in how we define notions such as envy-freeness. The goal of this project would be to study fairness and efficiency properties of various heuristic or optimization-based algorithms for course assignment in a highly realistic model that incorporates the abovementioned constraints (and more) in an empirical setup.