MS&E 336/CS 366: Computational Social Choice. Autumn 2021-2022
Course URL: http://www.stanford.edu/~ashishg/msande336.html.
Instructor: Ashish Goel, Stanford University.

Lecture 6, Wednesday, 6 October 2021. Scribed by Mac Broido.

7 Approval Voting and Participatory Budgeting

In today's class we briefly introduce approval voting and participatory budgeting, and overview some ideas for class projects motivated by open problems and practical applications.

7.1 Projects

1. The Hegselmann-Krause Dynamics Model

    - So far, we have been talking about voting, separate from any process of deliberation. In practice, opinion dynamics are often an important part of voting: it would be useful if we could model how one's opinion is changed by input from others.
    - Let $x_i^{(t)} \in \mathbb{R}$ denote the opinion of agent i at time t
    - We want to model an agent's opinion in the current time step as a function of the opinions of other agents in the previous time step:

    $$f(x_i^{(t+1)}) = f(x_1^{(t)}, x_2^{(t)}, \ldots, x_n^{(t)})$$

    - Hegselmann-Krause proposes that an agent's opinions evolves as the average of others' opinions close to theirs:

    $$f\left(x_i^{(t+1)}\right) = \text{avg}\{x_j^{(t)} \text{ s.t. } \mid x_j^{(t)} - x_i^{(t)} \mid \leq 1\}$$[1][2]

    - This converges to a steady state in $O(N^3)$ for N participants. The open question: is this tight? Proving this can be computed more efficiently (e.g. $O(N^2)$) would be a major result.
    - A nice property of this model is that we cannot have increased variance as we converge.
    - This model is not generative; its formulation is not necessarily from any realistic situation. Note also that no utility functions are considered.
    - An important consideration in analysis of opinion dynamics is *polarization*, which we will elaborate on at a later point.

2. Voting under strategic positioning

    - Each agent's utility decreases linearly (in all directions) from their ideal point.
    - However, they do not wish to be seen voting too far from the winner.

[1]This set always contains x_i^t, as $0 \leq 1$

[2]1 is an arbitrary distance; we could scale all of the x_i^t's and it would be a different number.

- So, they choose a region in which they will tolerate any result (such that there is more space in which they can "be seen choosing the winner")
- This is still truthful, because the region does contain their true point.
- Project idea: what does this look like in complex/multidimensional spaces?

3. Model of deliberation and argumentation theory, broadly
    - We attempt to accurately model what happens when people deliberate.
    - Are arguments convincing, and in what order are they presented?
    - We hypothesize that deliberation works because of empathy: you tell me what's important to you, and then I am able to understand your worldview and better understand your opinion.
    - We seek a model that is usable in empirical research, rather than one that is simply interesting as mathematical theory.
    - Perhaps the attention of the mathematically-inclined could produce a useful result.
4. Deliberation as negotiation, more theoretically
    - Method: agent X proposes an idea s. Agents Y and Z negotiate, with s as an outside alternative: if they cannot agree, then s will be chosen.
    - Recall our discussion of Nash bargaining, in which the Goel children are Y and Z and s is Chipotle
    - This would be more economic theory-oriented
5. Analysis of participatory budgeting data
    - Can we perform clustering or other unsupervised analysis on available data to inform us about how budgeting decisions are made?
6. Liquid democracy
    - Model situations in which agents can delegate their votes.
7. Gerrymandering
    - On October 13th, an ex-student will be giving a talk on this subject. We will not elaborate on this for now.
8. Market-based voting algorithms
    - Endow each voter with a budget, which they can allocate partially to candidates (issues) as they see fit.
    - The open problem: "what mechanism can you design that would be useful"
    - Assuming some utility model, such as a line, how would you implement delegative democracy?

7.2 Approval voting

- *Elicitation rule:* each agent gives a set of candidates they approve of. There are no scores or rankings, only the approval set.
- *Aggregation rule:* Choose the candidate that appears in the greatest number of approval sets. Ties are broken arbitrarily.
- Generally, this is not strategyproof, as agents could remove the winning candidate if they approve of it but it isn't their top choice, and changing their vote would result in a candidate they truly prefer winning.
- However, if every agent has *dichotomous preferences* (they have equal utility for all candidates they approve, and no utility for those they do not[3]), approval voting is incentive-compatible. For each agent, either the winner is in their true approval set, and they have utility 1, or it is not, and have no utility. Thus changing their approval set cannot increase their utility. The only exceptions are if a voter either approves of every candidate or no candidates. We conclude that approval voting can be a good mechanism if preferences are truly dichotomous, but less so if voters are forced to collapse non-dichotomous preferences into approval and non-approval sets.
- For dichotomous preferences, approval voting is Condorcet consistent and reinforcing, and is also very similar to a scoring rule.

7.3 Participatory Budgeting

- Participatory budgeting is an approach to situations in which a group of self-interested participants wish to provide their opinion on how a budget should be allocated to various projects. In this section we describe a few elicitation and aggregation methods, and analyze their theoretical and practical properties.
- Model setup and notation:
  - There are N voters and M expense items (e.g. city projects)
  - There is a total available budget $B > 0$
  - Each item x_j has cost $c_j > 0$
  - Feasible solutions are of the form $\mathbf{x} = \langle x_1, x_2, \dots, x_M \rangle$ s.t. $\mathbf{x} \geq 0, \mathbf{x} \leq \mathbf{c}$, and $\sum_{\text{projects } j} x_j \leq B$
  - We make the assumption that items can be chosen fractionally, as without this assumption the underlying problem is NP-hard.
- Prof. Goel has built a web platform as research for cities to implement participatory budgeting. This system was tested with over 100,000 people and has garnered significant interest. Soon, the platform may be sold to a survey platform provider such as SurveyMonkey, Google Forms, or Qualtrics. We will use examples from the platform to describe three participatory budgeting mechanisms.

[3]Specifically, there are two utility levels, so a voter equally less happy with all candidates not included in the approval set

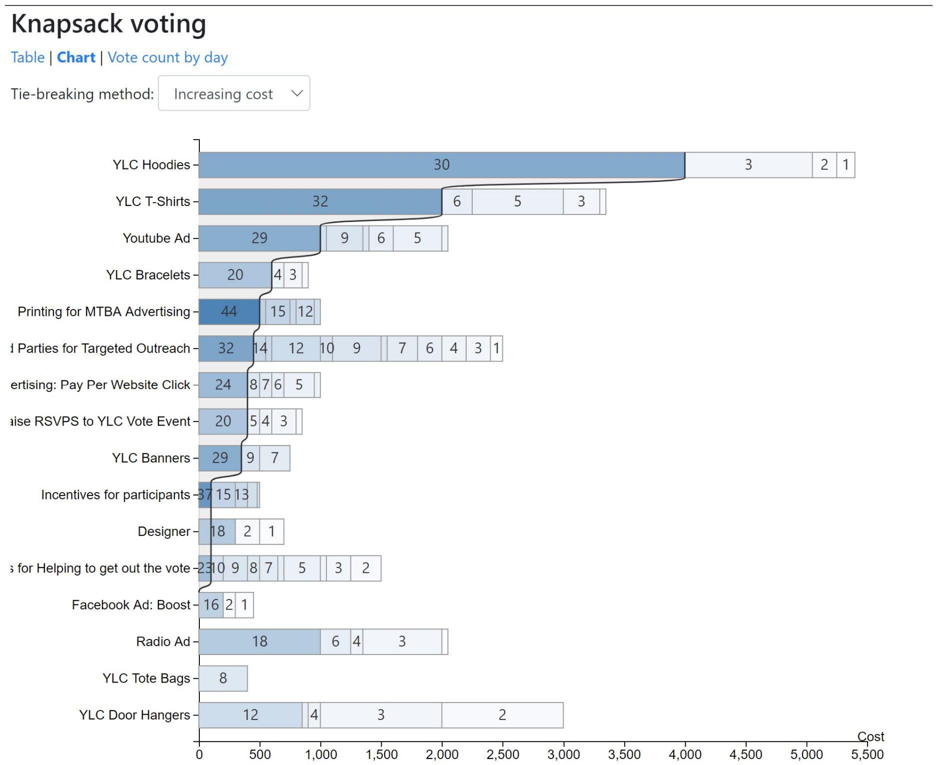

Figure 1: Water-filling aggregation method for knapsack voting with partial allocations

7.3.1 Approval Voting with Fixed Approval Set Size

- In the Boston experimental implementation of approval voting, each participant votes for exactly 4 projects.
- No constraints are placed on the costs of the chosen projects
- Coming up with an aggregation method for this mechanism is not straightforward: because of the unrestrained cost, it is possible that even if all voters agree on which 4 projects they want funded, those 4 projects cannot comprise the selected budget as the sum of their costs exceeds the total possible budget.
- Additionally, we still have the typical caveats of approval voting: we don't get any information on voters' preferences between items within and outside of the approval set.
- However, the elicitation method is user-friendly, as choosing your 4 favorite projects is intuitive and can be easily recorded on a paper ballot.

7.3.2 Knapsack Voting

- In *knapsack voting* as implemented in the (NYC District 8 experiment), each participant votes to approve or disapprove complete funding for each project, subject to the constraint that the sum of costs for the selected projects must not exceed the total budget.

- This elicitation mechanism conveys more information about voters' preferences: if they vote truthfully, it captures how they would allocate the budget if made dictator, subject only to the atomic project funding constraint.
- A natural aggregation method would be to select the projects included in the most approval sets ("knapsacks") until the sum of costs of these selected projects exceeds the total budget.

7.3.3 Knapsack Voting with Partial Allocations

- In *knapsack voting with partial allocations*, we modify knapsack voting as described in (7.3.2) by removing the constraint that a vote for a project must fund it entirely: rather, each voter can allocate any amount to each project, between 0 and the project's maximum budget.
- The aggregation method is given by the water-filling algorithm:
  1. For each project $x_j, j = 1, 2, \dots, M$, count how many voters approved at least \$1. Repeat this for the second, third $\dots, c_j$th dollar. We can see this in Figure 1, where for "YLC Hoodies", 30 voters approved dollars $1, 2, \dots, 4000$, and 3 voters approved up to \$5000. I will use *subproject* to refer to these dollar amounts within each project.
  2. For subproject that has not yet been selected for the result budget, identify the subproject included in the most approval sets (breaking ties arbitrarily). For example, the first subproject selected in this example is the first dollar of "Printing for MTBA Advertising", as no other subproject was included in at least 44 approval sets.
  3. Repeat step 2 for each dollar in the total budget, at which point we have the selected budget.
- While these knapsack elicitation methods can provide a granular description of voter preferences, they can be hard to implement in a city election as it is very difficult to use on a paper ballot. It is much more practical to have a computer interface in which voters can play around with different allocations until they finalize their submission, but for many voters using such an interface can be challenging to the point that expecting all citizens to truthfully describe their ideal allocation is infeasible.
- We should note that with all three mechanisms, no information is collected about how participants would vote on different total budgets. This is significant as often in policy-making, the allocation of a budget is discussed alongside its size.

7.4 Models of Voter Utility

- It would be useful to model voter utility or cost, and prevalent voting mechanisms are typically unable to measure this. We present two simple functions for voter utility given a known personal utility function and observed outcome.
  - *Linear utility:* voter i has utility $u_{i,j}$ for project j, and

$$U_i(\mathbf{x}) = \sum_{\text{project } j} u_{i,j} x_j$$

- *Overlap utility:* User i has ideal budget

$$\mathbf{z}_i = \langle z_{i,1}, z_{i,2}, \dots, z_{i,M} \rangle$$

  User i's utility for budget X is given by

$$U_i(x) = \sum_{\text{project } j} \min\{z_{ij}, x_j\} = \text{cost}_i(\mathbf{x}) \triangleq \|z_i - \mathbf{x}\|_1$$

  Intuitively, this quantifies how much the ideal budget deviates from the selected result, assuming dichotomous utility for each subproject.

- Discussion points
  - What are our natural notions of cost, and how do these connect to utilities?
  - How can we elicit ideal utilities or ideal budgets using voting mechanisms?
  - What are our natural notions of strategyproofness, and how do these interact with voting mechanisms both in theory and practice?
  - Where might we see effects of Bayesian persuasion, where agents strategically hide information to cause other voters to make decisions more preferable to the "persuasive agent"?