Effective Heuristics for Committee Scoring Rules

Piotr Faliszewski — AGH University
Martin Lackner — TU Wien
Dominik Peters — University of Oxford
Nimrod Talmon — Weizmann Institute of Science

Committee Elections

- Voters rank candidates
- Want to find a committee of k candidates
- Finding the best committee is NP-hard for many objective functions
- We study heuristics: do they produce committees that are 'similar' to the optimum?
- We focus on the t-Borda rules: Every voter is 'represented' by their t most-preferred committee members. and we maximise total utility using Borda scores.
- Special case: 1-Borda = Chamberlin—Courant

Heuristics

- **Greedy**: build up the committee step-by-step by adding the candidate who most improves objective.
- **Removal**: repeatedly kick out the candidate who least contributes to the objective, until k remain.
- **Banzhaf**: Add candidate to the committee who most improves objective in expectation, taken over a uniform distribution of all possible committees compatible with previous choices.
  - this heuristic (based on concepts from cooperative game theory) is new
  - we show that it can be evaluated in polynomial time for separable committee scoring rules
- **Simulated Annealing**: Start with a random committee, and then repeatedly randomly swap candidates to improve objective. Drawback: not deterministic!

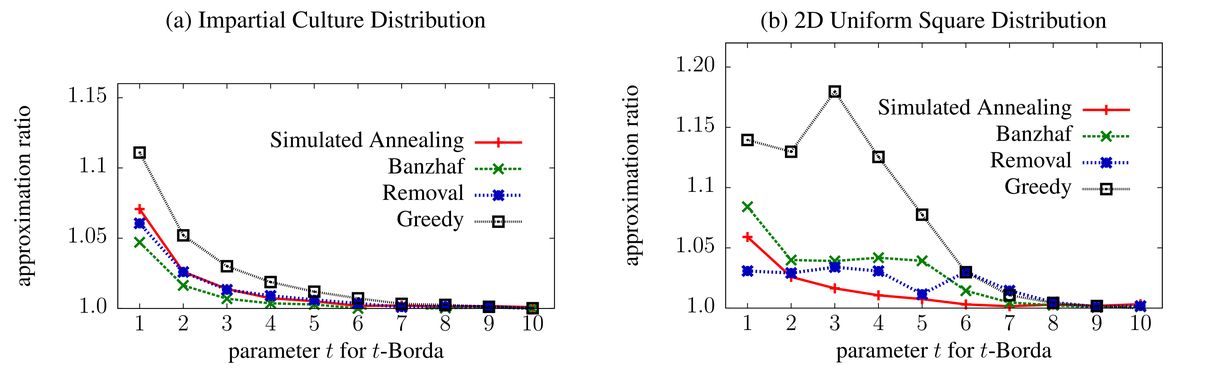

Figure 1: Ratios of average reverse t-Borda scores computed by our heuristics to those computed using the exact ILP-based algorithm for our two preference generation models ($m = 100$, $n = 100$, $k = 10$).

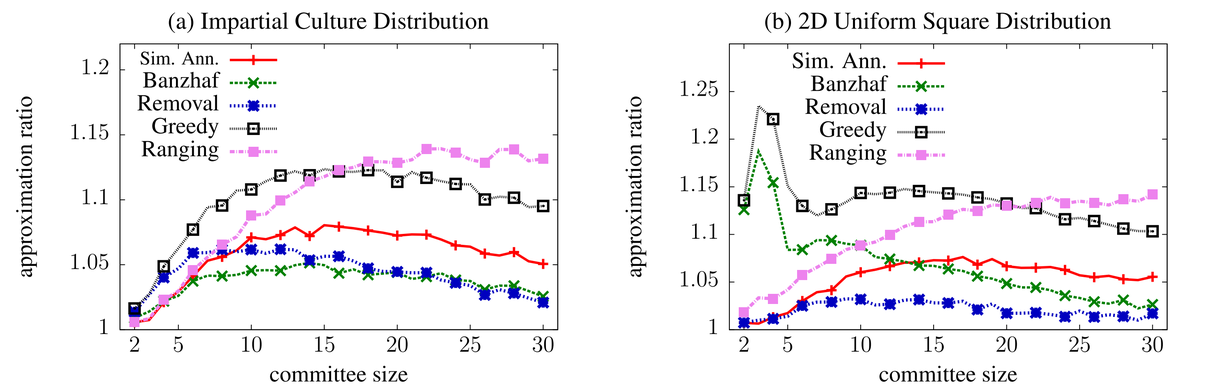

Figure 3: Ratios of average reverse CC scores (average positions of the voters' representatives) computed by our heuristics to those computed using the exact algorithm for our two preference generation models ($m = 100$, $n = 100$).

Experimental Evaluation

- **Impartial Culture**: for each voter, draw their preference ranking uniformly at random from all rankings
- **2D uniform square**: each candidate and voter is placed at a random point in the square; voters prefer candidates closer to them
- Determine approximation ratios of our heuristics
- Via visual representations, check whether heuristics generate similar committees to the exact rule
- ***Results***: Simulated Annealing performs very well in all settings. Greedy performs poorly, but Removal and Banzhaf often match or outperform Simulated Annealing.

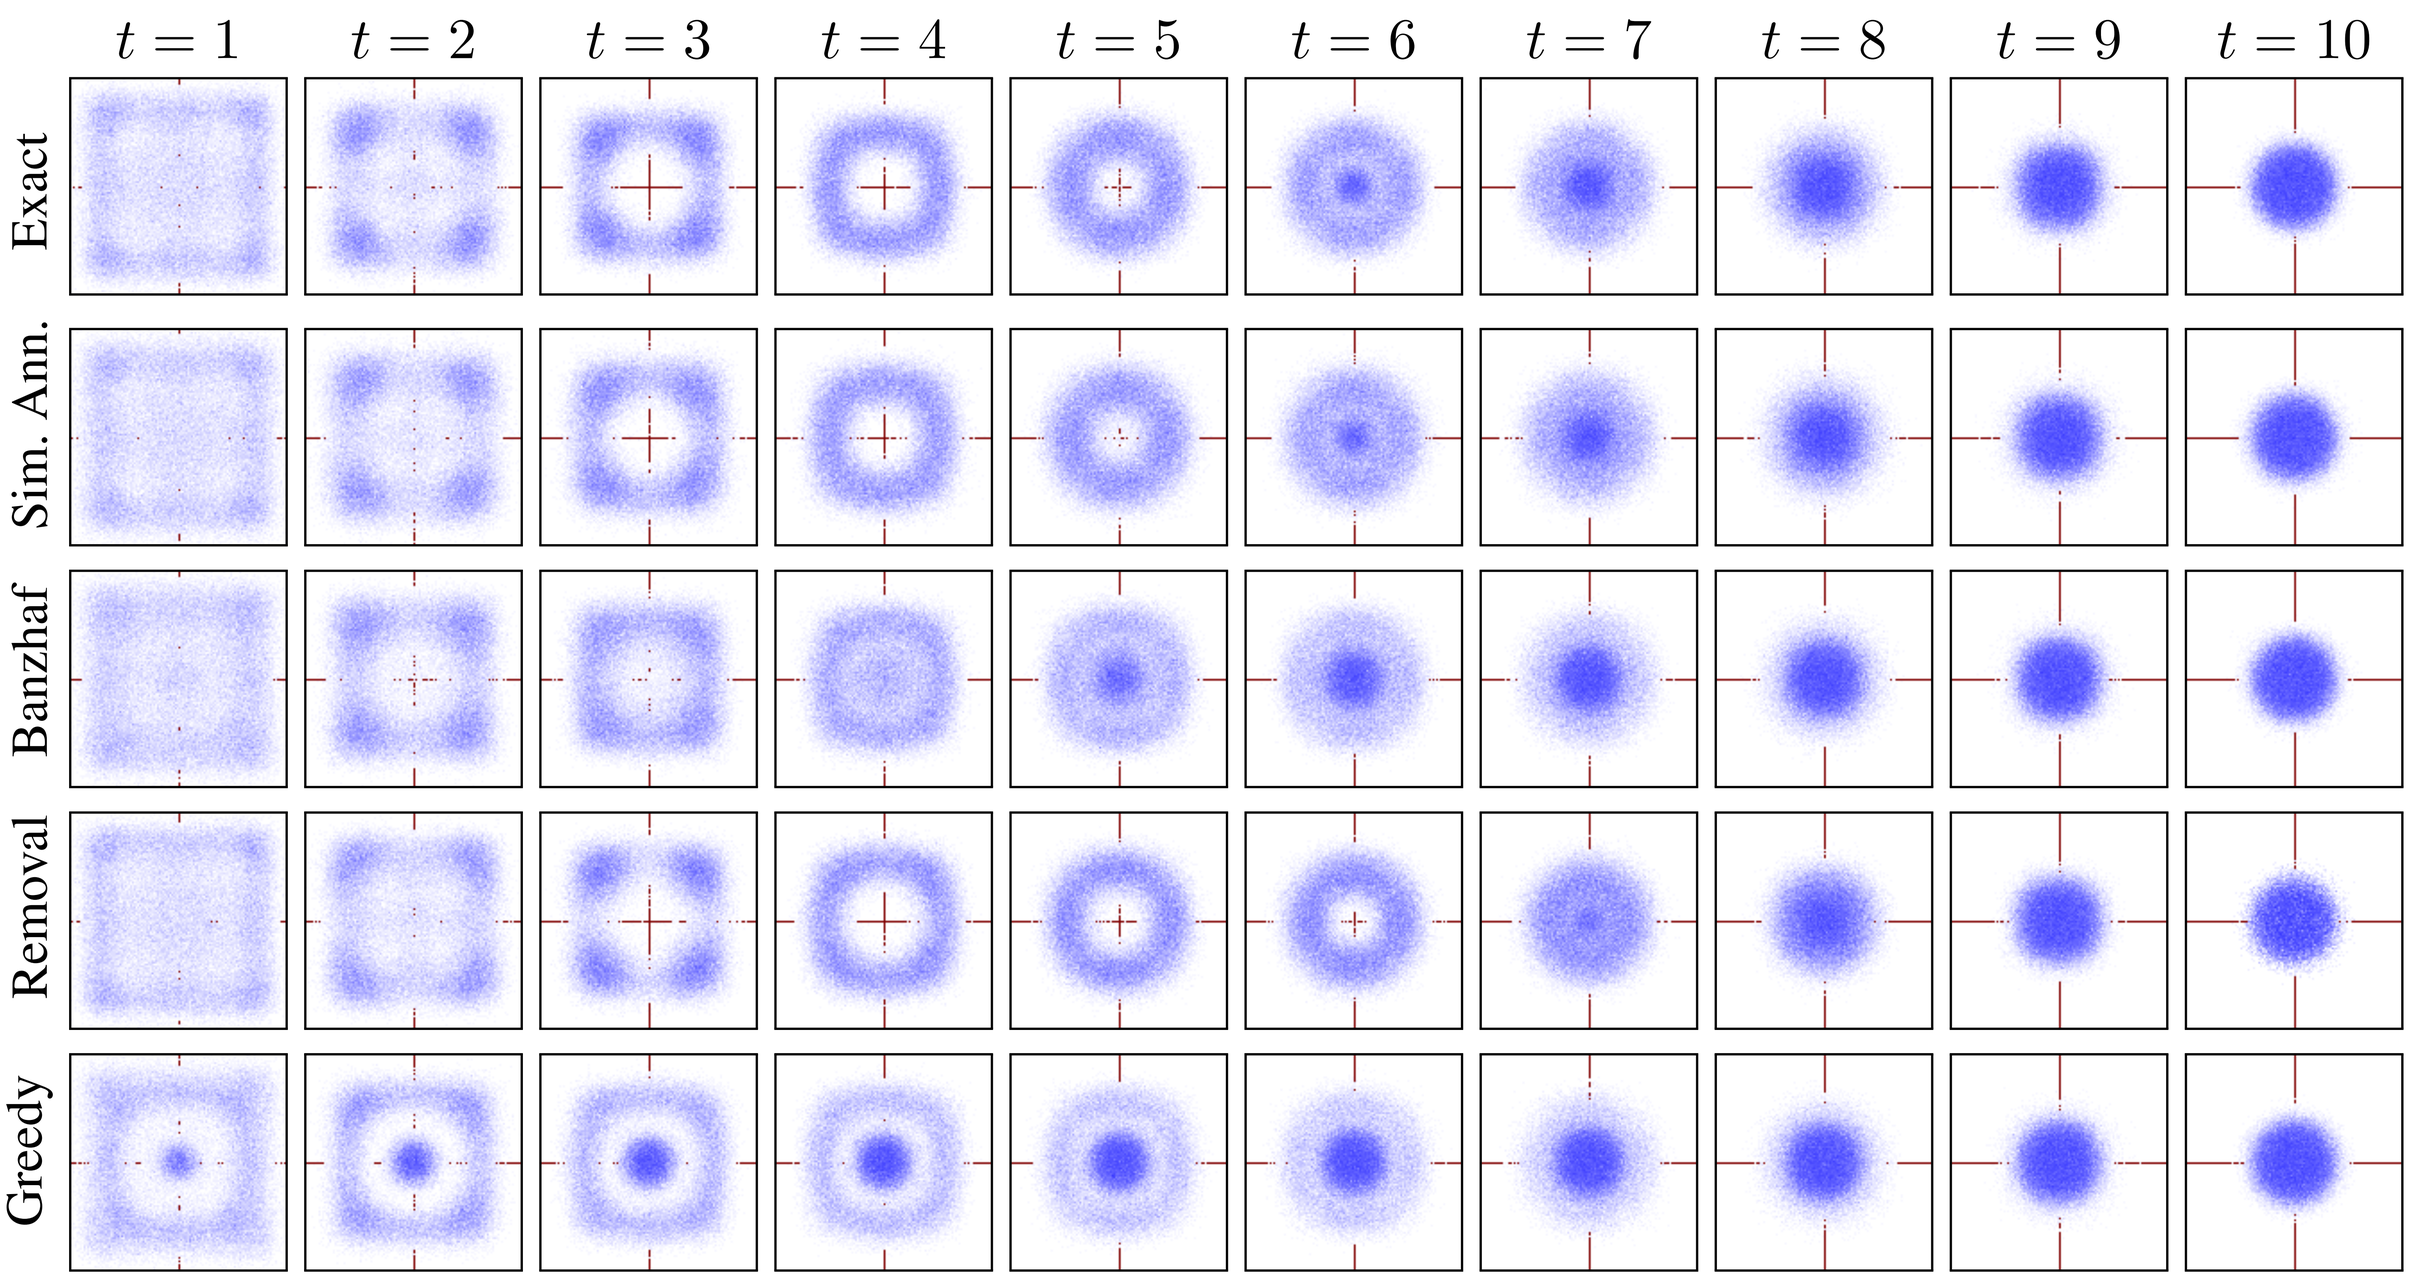

Figure 2: Histograms showing frequency of including candidates from given areas of the $[-3, 3] \times [-3, 3]$ square in the winning committees, depending on the algorithm and the value of t (2D uniform square distribution, $m = 100$, $n = 100$, $k = 10$).

Presented at AAAI 2018, New Orleans