Learning to Manipulate under Limited Information

Wesley H. Holliday[1], Alexander Kristoffersen[1], Eric Pacuit[2]

[1]University of California, Berkeley
[2]University of Maryland

New Directions in Social Choice at EC 2024

Learning to Manipulate under Limited Information

We use machine learning to gauge how resistant a preferential voting method is to manipulation under limited information about how other voters will vote.

Wesley Holliday, Alexander Kristoffersen, Eric Pacuit. *Learning to Manipulate under Limited Information*. arxiv.org/abs/2401.16412, 1st Workshop on Social Choice and Learning Algorithms (SCaLA 2024).

How to manipulate

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$\downarrow$

$$Borda(\mathbf{P}) = \{b\}$$

Winners

How to manipulate

Rankings

$R \in \{a\ b\ c\ d,\ \ldots,\ d\ c\ b\ a\}$

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$$Borda(\mathbf{P}) = \{b\}$$

Winners

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
a	b	b	b	b
d	d	d	a	a
b	a	a	c	c

$$Borda(R, \mathbf{P}_{-v_1}) = \{c, d\}$$

Winners

How to manipulate

Rankings

$R \in \{a\ b\ c\ d,\ \ldots,\ d\ c\ b\ a\}$

v_1	v_2	v_3	v_4	v_5
c	c	c		
b	b	b		
a	d	d		
d	a	a	c	c

Which ranking R *should* v_1 submit?

v_1	v_2	v_3	v_4	v_5
		c	d	d
		b	b	b
		d	a	a
b	a	a	c	c

$Borda(\mathbf{P}) = \{b\}$

Winners

$Borda(R, \mathbf{P}_{-v_1}) = \{c, d\}$

Winners

Profitable manipulations

Given a profile of utilities for each voters, we can define the profile of rankings submitted by each voter, where alternative a is ranked above alternative b when the utility of a is greater than the utility of b:

Voters	a	b	c	d
v_1	0.1	0.65	0.9	0.08
v_2	0.7	0.9	1.0	0.8
v_3	0.01	0.03	0.5	0.02
v_4	0.1	0.5	0	0.9
v_5	0.1	0.2	0.05	1.0

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

P

Profitable manipulations

Given a profile of utilities for each voters, we can define the profile of rankings submitted by each voter, where alternative a is ranked above alternative b when the utility of a is greater than the utility of b:

Voters	a	b	c	d
v_1	0.1	0.65	0.9	0.08
v_2	0.7	0.9	1.0	0.8
v_3	0.01	0.03	0.5	0.02
v_4	0.1	0.5	0	0.9
v_5	0.1	0.2	0.05	1.0

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

P

Profitable manipulations

A ranking R is a *profitable manipulation* for voter i in preference profile $\mathbf{P}$ generated from a utility profile $\mathbf{U}$ for voting method F provided that

$$\mathbf{EU}_i(F_\ell\ (\ R, \mathbf{P}_{-i}\ )) > \ \mathbf{EU}_i(F_\ell\ (\ \mathbf{P}\ ))$$

Profitable manipulations

A ranking R is a *profitable manipulation* for voter i in preference profile $\mathbf{P}$ generated from a utility profile $\mathbf{U}$ for voting method F provided that

$$\mathbf{EU}_i(F_\ell(R, \mathbf{P}_{-i})) > \mathbf{EU}_i(F_\ell(\boxed{\mathbf{P}}))$$

The profile where all voters submit their "true" ranking.

Profitable manipulations

A ranking R is a *profitable manipulation* for voter i in preference profile $\mathbf{P}$ generated from a utility profile $\mathbf{U}$ for voting method F provided that

$$\mathbf{EU}_i(F_\ell(\boxed{R, \mathbf{P}_{-i}})) > \mathbf{EU}_i(F_\ell(\boxed{\mathbf{P}}))$$

The profile where all voters except i submit their "true" ranking and i submits R.

The profile where all voters submit their "true" ranking.

Profitable manipulations

A ranking R is a *profitable manipulation* for voter i in preference profile $\mathbf{P}$ generated from a utility profile $\mathbf{U}$ for voting method F provided that

Expected utility for voter i of the winners according to the voting method F, using even-chance tiebreaking if needed

$$\mathbf{EU}_i(F_\ell(\boxed{R, \mathbf{P}_{-i}})) > \mathbf{EU}_i(F_\ell(\boxed{\mathbf{P}}))$$

The profile where all voters except i submit their "true" ranking and i submits R.

The profile where all voters submit their "true" ranking.

Profitable manipulations

The *profitability* of voter i's submitting ranking R given utility profile $\mathbf{U}$ that induces preference profile $\mathbf{P}$ is given by

$$\frac{\mathbf{EU}_i(F_\ell(R, \mathbf{P}_{-i})) - \mathbf{EU}_i(F_\ell(\mathbf{P}))}{\max(\{\mathbf{U}_i(x) \mid x \in X\}) - \min(\{\mathbf{U}_i(x) \mid x \in X\})},$$

adopting the normalization of Relative Utilitarianism (Dhillon and Mertons 1999).

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$

Winners

Ideal Voter

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$

Winners

Choose an R that maximizes profitability

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$
Winners

Plurality Scores

a	b	c	d
0	0	3	2

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$
Winners

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$
Winners

Plurality Scores

a	b	c	d
0	0	3	2

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$
Winners

v_1	v_2	v_3	v_4	v_5
c	d	c	c	d
b	a	b	b	a
a	b	d	a	b
d	c	a	d	c

,

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	a	d	a	a
a	b	b	b	b
d	d	a	c	c

, $\cdots$

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$

Winners

Plurality Ranking

$$c > d > (a \ b)$$

v_1	v_2	v_3	v_4	v_5
c	d	c	c	d
b	a	b	b	a
a	b	d	a	b
d	c	a	d	c

,

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	a	d	a	a
a	b	b	b	b
d	d	a	c	c

, $\cdots$

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$

Winners

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$

Winners

Sincere Winners

a	b	c	d
0	1	0	0

v_1	v_2	v_3	v_4	v_5
c	d	b	b	a
b	b	a	d	d
a	a	d	a	c
d	c	c	c	b

,

v_1	v_2	v_3	v_4	v_5
c	b	b	d	b
b	a	c	a	d
a	c	a	c	c
d	d	d	b	a

, $\cdots$

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$

Winners

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$

Winners

Majority Matrix

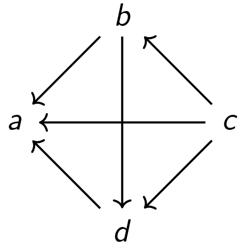

v_1	v_2	v_3	v_4	v_5
c	b	c	a	b
b	c	b	c	d
a	d	d	b	c
d	a	a	d	a

,

v_1	v_2	v_3	v_4	v_5
c	d	c	b	c
b	a	b	d	a
a	b	d	a	b
d	c	a	c	d

, $\cdots$

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$

Winners

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$
Winners

Margin Matrix

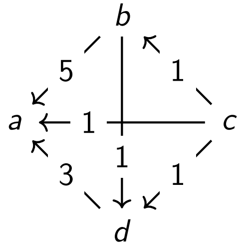

v_1	v_2	v_3	v_4	v_5
c	d	c	c	d
b	b	b	b	b
a	a	d	d	a
d	c	a	a	c

,

v_1	v_2	v_3	v_4	v_5
c	c	d	c	b
b	d	b	b	d
a	b	a	d	a
d	a	c	a	c

, $\cdots$

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$
Winners

Limited information

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$F(\mathbf{P})$

Winners

Qualitative Margin Matrix

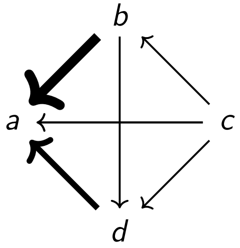

v_1	v_2	v_3	v_4	v_5
?	c	c	d	d
?	b	b	b	b
?	d	d	a	a
?	a	a	c	c

$F(R, \mathbf{P}_{-1})$

Winners

v_1	v_2	v_3	v_4	v_5
c	d	c	c	d
b	b	b	b	b
a	a	d	d	a
d	c	a	a	c

,

v_1	v_2	v_3	v_4	v_5
c	c	d	c	b
b	d	b	b	d
a	b	a	d	a
d	a	c	a	c

, $\cdots$

Learning to manipulate under limited information

Learning to manipulate under limited information

- We trained $\approx$ 85,000 multi-layered perceptrons (MLP) of 26 sizes to manipulate against 8 different voting methods, under 6 types of limited information, in profiles with 5-21 voters and 3-6 alternatives.

Learning to manipulate under limited information

- We trained $\approx$ 85,000 multi-layered perceptrons (MLP) of 26 sizes to manipulate against 8 different voting methods, under 6 types of limited information, in profiles with 5-21 voters and 3-6 alternatives.
- These networks act as function approximators for profitable manipulation policies for a given voting method and type of limited information.

Learning to manipulate under limited information

- We trained $\approx$ 85,000 multi-layered perceptrons (MLP) of 26 sizes to manipulate against 8 different voting methods, under 6 types of limited information, in profiles with 5-21 voters and 3-6 alternatives.
- These networks act as function approximators for profitable manipulation policies for a given voting method and type of limited information.
- We evaluate the manipulation resistance of a voting method by the size and complexity of the network required to learn a profitable manipulation policy, as well as the average profitability of learned policies.

Setup

1. **Generate Utility Profiles**: We generate utility profiles for the voters for training, validation, and evaluation according to some probability model:

Setup

1. **Generate Utility Profiles**: We generate utility profiles for the voters for training, validation, and evaluation according to some probability model:
    - **Random Utility Model**: for each voter, the utility of each alternative is drawn independently from the uniform distribution on the $[0, 1]$ interval.

Setup

1. **Generate Utility Profiles**: We generate utility profiles for the voters for training, validation, and evaluation according to some probability model:
    - **Random Utility Model**: for each voter, the utility of each alternative is drawn independently from the uniform distribution on the $[0, 1]$ interval.
    - **2D Spatial Model**: each alternative and each voter is independently placed in $\mathbb{R}^2$ according to the multivariate normal distribution with no correlation between the two dimensions; then the utility of a alternative for a voter is the square of the Euclidean distance between the alternative and the voter

Setup

1. **Generate Utility Profiles**: We generate utility profiles for the voters for training, validation, and evaluation according to some probability model:
   - **Random Utility Model**: for each voter, the utility of each alternative is drawn independently from the uniform distribution on the $[0, 1]$ interval.
   - **2D Spatial Model**: each alternative and each voter is independently placed in $\mathbb{R}^2$ according to the multivariate normal distribution with no correlation between the two dimensions; then the utility of a alternative for a voter is the square of the Euclidean distance between the alternative and the voter
   - **Mallows Model**: generate a linear profile with the Mallows model ($\phi = 0.8$); then for each ranking generate a random utility that represents the ranking

Setup

1. **Generate Utility Profiles**: We generate utility profiles for the voters for training, validation, and evaluation according to some probability model:
   - **Random Utility Model**: for each voter, the utility of each alternative is drawn independently from the uniform distribution on the $[0, 1]$ interval.
   - **2D Spatial Model**: each alternative and each voter is independently placed in $\mathbb{R}^2$ according to the multivariate normal distribution with no correlation between the two dimensions; then the utility of a alternative for a voter is the square of the Euclidean distance between the alternative and the voter
   - **Mallows Model**: generate a linear profile with the Mallows model ($\phi = 0.8$); then for each ranking generate a random utility that represents the ranking
2. **Labeling**: For a given training profile and voting method, compute the optimal rankings that the manipulator could possibly submit.

Setup

3. **Training**:
   - The input to an MLP is (1) the manipulator's own utility function plus (2) some limited information about the profile.

Setup

3. **Training**:
   - The input to an MLP is (1) the manipulator's own utility function plus (2) some limited information about the profile.
   - Applying a softmax to the output yields a probability distribution π over all rankings that the manipulator could submit, which we reduce to the probability of choosing an optimal ranking or not.

Setup

3. **Training**:
    - The input to an MLP is (1) the manipulator's own utility function plus (2) some limited information about the profile.
    - Applying a softmax to the output yields a probability distribution π over all rankings that the manipulator could submit, which we reduce to the probability of choosing an optimal ranking or not.
    - We compute the final loss as the mean-squared error between the reduced distribution and the distribution assigning probability 1 to choosing an optimal-labeled ranking and 0 to choosing a non-optimal-labeled ranking.

Setup

3. **Training**:
   - The input to an MLP is (1) the manipulator's own utility function plus (2) some limited information about the profile.
   - Applying a softmax to the output yields a probability distribution π over all rankings that the manipulator could submit, which we reduce to the probability of choosing an optimal ranking or not.
   - We compute the final loss as the mean-squared error between the reduced distribution and the distribution assigning probability 1 to choosing an optimal-labeled ranking and 0 to choosing a non-optimal-labeled ranking.

4. **Evaluation**: When evaluating the MLP, we take the most probable ranking R according to π to be submitted, and we compute the profitability of R.

Results: Random Utility Model, 6 alternatives

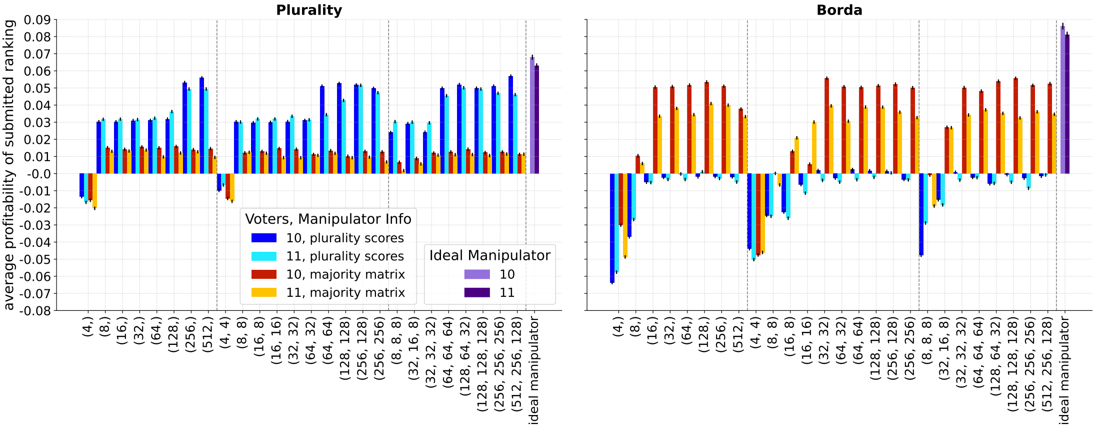

Results: 2D Spatial Model, 6 alternatives

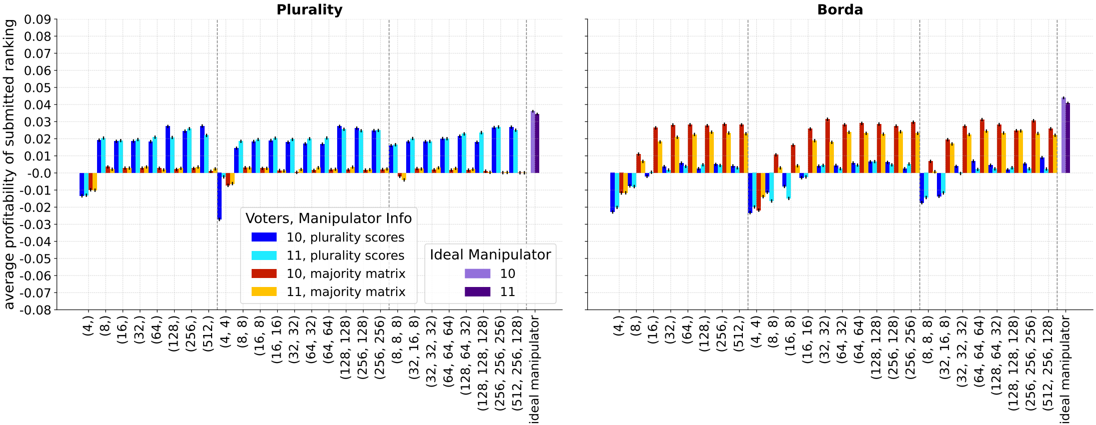

Results: Mallows Model, 6 alternatives

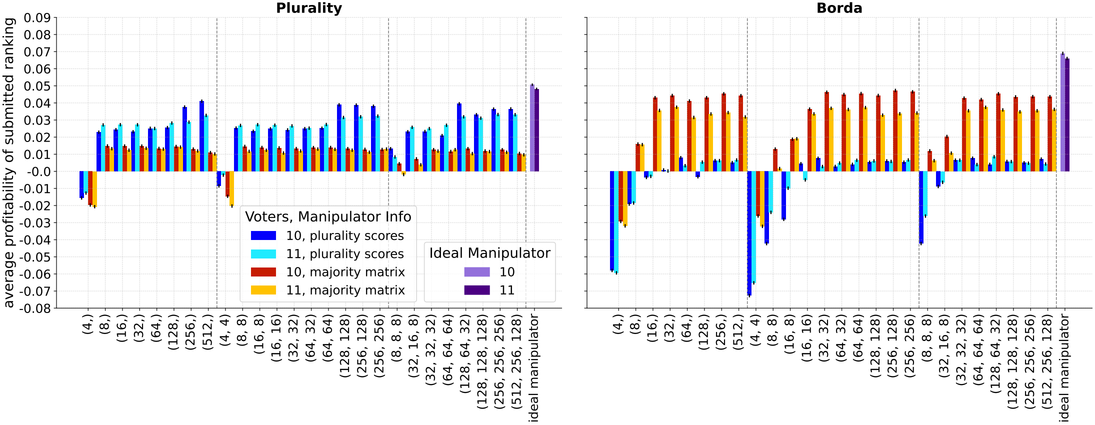

Results: Random Utility Model, 3-6 alternaitves

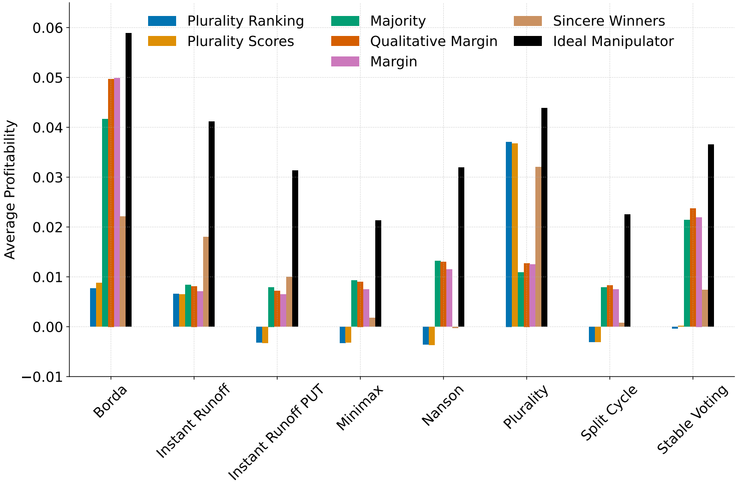

Average profitability of the best performing MLP with any hidden layer configuration for a given voting method and information type.

Results: 2D Spatial Model, 3-6 alternatives

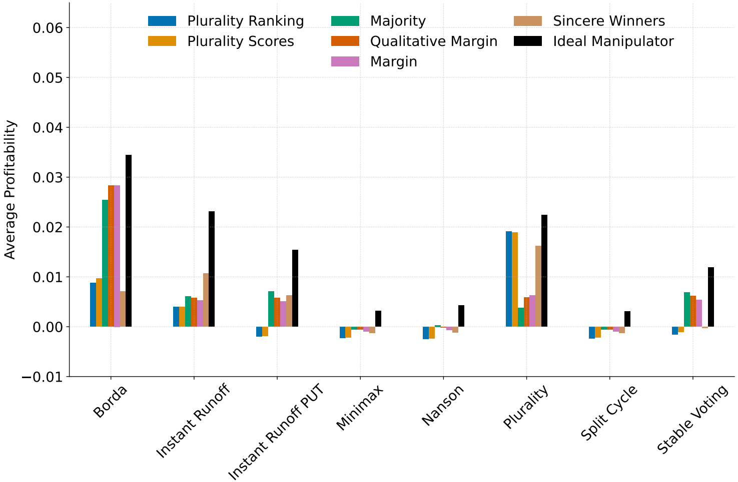

Average profitability of the best performing MLP with any hidden layer configuration for a given voting method and information type.

Results: Mallows Model, 3-6 alternatives

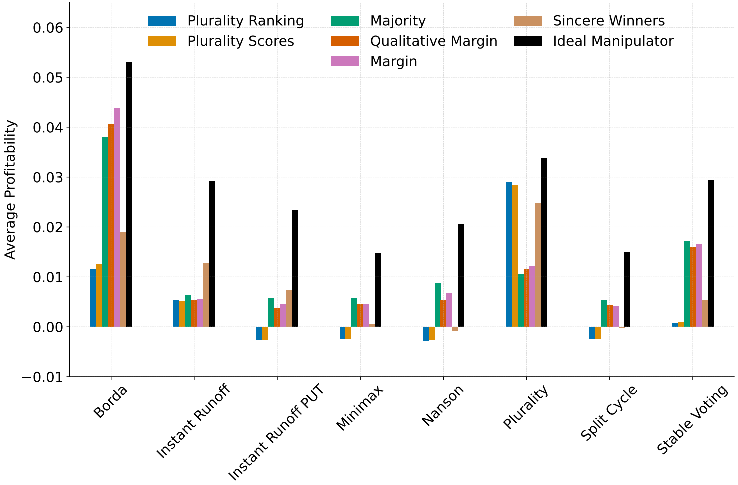

Average profitability of the best performing MLP with any hidden layer configuration for a given voting method and information type.

Results: Random Utility Model, 3-6 alternatives

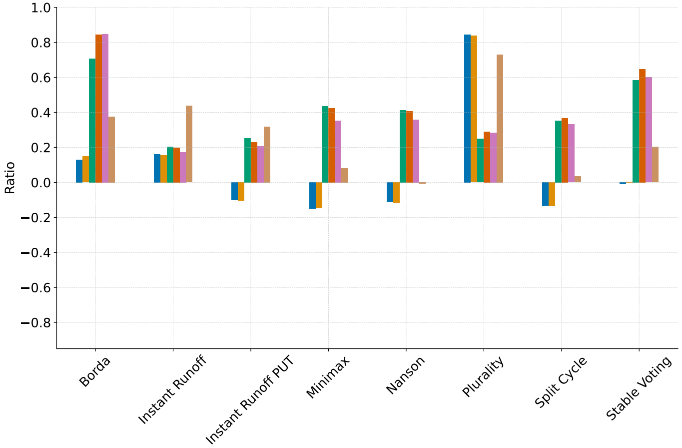

The ratio of the average profitability of the MLP's submitted ranking to that of the ideal manipulator's submitted ranking.

Results: 2D Spatial Model, 3-6 alternatives

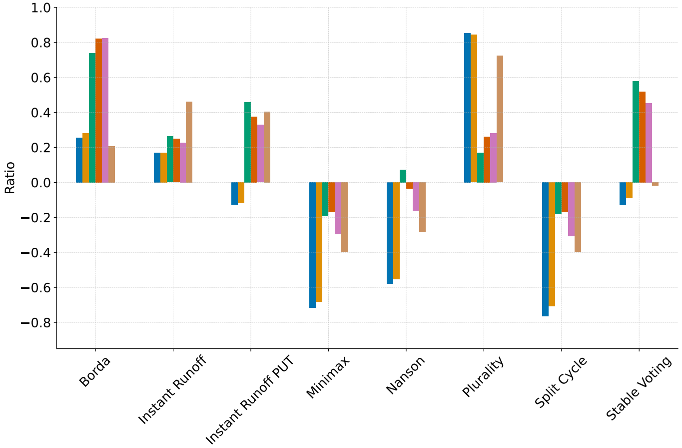

The ratio of the average profitability of the MLP's submitted ranking to that of the ideal manipulator's submitted ranking.

Results: Mallows Model, 3-6 alternatives

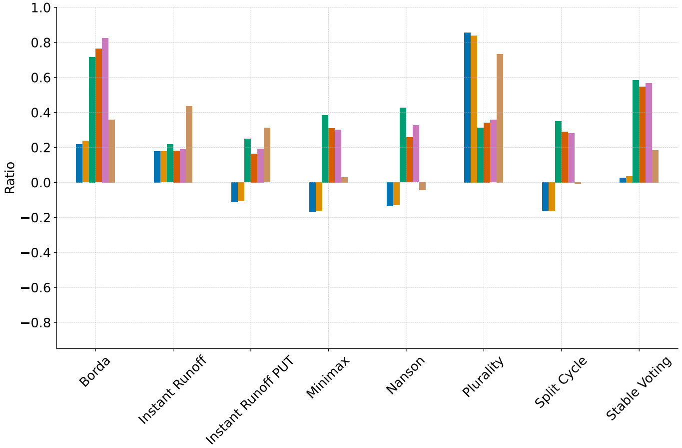

The ratio of the average profitability of the MLP's submitted ranking to that of the ideal manipulator's submitted ranking.

Conclusion

It is possible for MLPs to learn to manipulate preferential voting methods on the basis of limited information, though the profitability of doing so varies significantly between different voting methods and types of information.

Conclusion

It is possible for MLPs to learn to manipulate preferential voting methods on the basis of limited information, though the profitability of doing so varies significantly between different voting methods and types of information.

Roughly three types of methods:

- **Highly manipulable even under limited info**: e.g., Borda;
- **Significantly manipulable under full info but not under limited**: e.g., Instant Runoff (though somewhat manipulable with sincere winners info);
- **Highly resistant to manipulation, especially under limited info**: e.g., Minimax.

Conclusion

Additional research questions:

- What about manipulation by a *coalition* of voters?

Conclusion

Additional research questions:

- What about manipulation by a *coalition* of voters?
- What if all voters simultaneously strategize?

Conclusion

Additional research questions:

- What about manipulation by a *coalition* of voters?
- What if all voters simultaneously strategize?
- What is the social cost or benefit of the learned manipulations?

Conclusion

Additional research questions:

- What about manipulation by a *coalition* of voters?
- What if all voters simultaneously strategize?
- What is the social cost or benefit of the learned manipulations?

Cf. K. Dowding and M. van Hees (2008), "In Praise of Manipulation," *British Journal of Political Science*, 38(1), pp. 1-15.

Conclusion

Based on considerations of manipulability, William H. Riker's (1988) wrote:

> I conclude that the meaning of social choices is quite obscure. They may consist of the amalgamation of the true tastes of the majority. . . or they may consist simply of the tastes of some people (whether a majority or not) who are skillful or lucky manipulators. If we assume social choices are often the latter, they may consist of what the manipulators truly want, or they may be an accidental amalgamation of what the manipulators (perhaps unintentionally) happened to produce. Furthermore, since we can by observation know only expressed values (never true values), we can never be sure, for any particular choice, which of these possible interpretations are correct. (p. 167)

Conclusion

Based on considerations of manipulability, William H. Riker's (1988) wrote:

> I conclude that the meaning of social choices is quite obscure. They may consist of the amalgamation of the true tastes of the majority. . . or they may consist simply of the tastes of some people (whether a majority or not) who are skillful or lucky manipulators. If we assume social choices are often the latter, they may consist of what the manipulators truly want, or they may be an accidental amalgamation of what the manipulators (perhaps unintentionally) happened to produce. Furthermore, since we can by observation know only expressed values (never true values), we can never be sure, for any particular choice, which of these possible interpretations are correct. (p. 167)

Can we mitigate these worries to some extent by the use of more manipulation-resistant preferential voting methods?

Thank you!

Wesley Holliday, Alexander Kristoffersen, Eric Pacuit. *Learning to Manipulate under Limited Information*. arxiv.org/abs/2401.16412, 1st Workshop on Social Choice and Learning Algorithms (SCaLA 2024).

`https://github.com/epacuit/ltm`

`https://pref-voting.readthedocs.io/`