

Manipulating Elections by Changing Voter Perceptions

Junlin Wu, Andrew Estornell, Lecheng Kong and Yevgeniy Vorobeychik

Washington University in St. Louis

{junlin.wu, aestornell, jerry.kong, yvorobeychik}@wustl.edu

Abstract

The integrity of elections is central to democratic systems. However, a myriad of malicious actors aspire to influence election outcomes for financial or political benefit. A common means to such ends is by manipulating perceptions of the voting public about select candidates, for example, through misinformation. We present a formal model of the impact of perception manipulation on election outcomes in the framework of spatial voting theory, in which the preferences of voters over candidates are generated based on their relative distance in the space of issues. We show that controlling elections in this model is, in general, NP-hard, whether issues are binary or real-valued. However, we demonstrate that critical to intractability is the diversity of opinions on issues exhibited by the voting public. When voter views lack diversity, and we can instead group them into a small number of categories—for example, as a result of political polarization—the election control problem can be solved in polynomial time in the number of issues and candidates for arbitrary scoring rules.

1 Introduction

Elections are among the core functional elements of democratic systems. Consequently, there is broad consensus that their integrity is among the top democratic priorities. However, malicious actors may attempt to subvert elections for their own means, whether financial or political [Caldwell *et al.*, 2019; Harper *et al.*, 2019; Khetani-Shah and Deutsch, 2019]. A common approach for manipulating elections is by spreading false information about select candidates, an extreme example of which is the infamous “Pizzagate” campaign targeting Hillary Clinton in the 2016 U.S. presidential election [Robb, 2017]. Less extreme, but far more common, is the spread of misinformation about the positions of candidates on specific issues, such as taxation and debt.

The issue of election vulnerability to malicious manipulation has been studied in the computational social choice literature from a computational complexity perspective under the terms *election control* (when the election structure itself is manipulated) [Bartholdi *et al.*, 1992; Hemaspaandra *et al.*,

2007; Chen *et al.*, 2017] and *bribery* (when manipulation is through changing voter preferences over candidates) [Bredereck *et al.*, 2016; Faliszewski and Rothe, 2016].

The traditional study of election control takes voter preferences as given, while considerations of bribery investigate direct manipulations of preference rankings of individual voters. However, neither is a natural model of the impact of misinformation *about particular issues* on the perceptions of candidates by the voting public. To address this gap, we propose a new model of election control in the *spatial voting theory* framework. Spatial voting theory explicitly captures voter and candidate positions on issues, with voter preferences over candidates determined by their relative distance in issue space [Anshelevich and Postl, 2016; Anshelevich *et al.*, 2018; Enelow and Hinich, 1984]. In our model of election manipulation, a malicious party can change voter perceptions of a target candidate on issues, subject to a budget constraint (more precisely, we constrain the l_p norm of the manipulation to be below a specified bound).¹ We consider both constructive control, where the malicious goal is to cause the target candidate to win, and destructive control, in which the goal is to cause the target candidate to lose.

We show that when the issues are binary-valued, the problem is hard even with two candidates, for both forms of control and for any l_p norm with integer $1 \leq p < \infty$ used to measure distance in issue space. When issues are real-valued, however, the conclusions for constructive and destructive control differ slightly. For destructive control, the problem is hard even with two candidates. For constructive control, we show hardness for plurality elections when the number of candidates is arbitrary, for l_p norm with integer $1 < p \leq \infty$. However, if there are only two candidates and we measure distance using l_∞ norm, the control problem can be solved in polynomial time. Furthermore, we show that if we restrict either the number of issues or the number of voters to be bounded by a constant, all control problems become tractable, whether issues are binary (for arbitrary l_p norm) or real-valued (for l_2 and l_∞), for arbitrary scoring rules used to determine election outcomes. Moreover, we show that the

¹This model can also be viewed as an example of bribery, in the sense that the manipulation affects voter preference rankings over candidates. Our use of the term *election control* is general, referring to any setting in which a malicious party wishes to subvert an election, whatever means they use for doing so.

tractability generalizes even when the number of voters is arbitrary, but their opinions on issues are limited to only a constant (that is, small) number of options.

These seemingly highly technical results offer a broader insight: vulnerability of elections to malicious manipulation of voter perceptions hinges on the extent to which voters exhibit a high diversity of political views. When this is the case, elections are highly resistant to manipulation. However, *when voters are Balkanized into a small number of groups with effective uniformity of views within each, for example, due to political polarization, elections become easy to manipulate through misinformation.*

Our model of election control is related to several recent studies of election control in the spatial voting theory framework [Lu *et al.*, 2019; Estornell *et al.*, 2020]. However, the means of manipulation in this closely related work is changing the relative importance of issues to voters, whereas our focus is on changing voter perceptions of candidates.

Related Work. The study of election control was initiated by Bartholdi *et al.* [1992], who studied the impact of adversarially adding, deleting, or partitioning candidates or voters on election outcomes in the constructive control framework. Numerous follow-up efforts extended this analysis in a number of directions, such as destructive control [Hemaspaandra *et al.*, 2007], a variety of voting rules and settings [Menton, 2012; Erdélyi *et al.*, 2015; Chen *et al.*, 2017], as well as when voter preferences can be modified (commonly called the *bribery problem* [Bredereck *et al.*, 2016; Faliszewski and Rothe, 2016] or *optimal lobbying* [Christian *et al.*, 2007; Binkele-Raible *et al.*, 2014]).

In most election control settings voter preferences are specified directly as preference rankings over the candidates. An alternative approach based on spatial theory of voting, specifies voter and candidate positions on issues, with preference rankings then induced from relative distances between voter and candidate positions [Davis and Hinich, 1968; Enelow and Hinich, 1984; Anshelevich and Postl, 2016; Anshelevich *et al.*, 2018]. Lu *et al.* [2019] were the first to investigate election control within the spatial theory voting model, with the adversary’s ability restricted to selecting a subset of issues that become the focus of voting preferences. Estornell *et al.* [2020] study a variation in which an adversary can modify the relative importance of issues in determining voter preferences over candidates. Both are distinct from our model in which the adversary modifies not the importance of issues, but the *perceptions* of a particular candidate by the voters.

Several models of election control are also motivated by the spread of misinformation about candidates on social networks [Wilder and Vorobeychik, 2018; Castiglioni *et al.*, 2020]. However, these focus on stochastic spread of misinformation in the social influence modeling framework [Kempe *et al.*, 2003], but use the conventional model of elections in which voter preferences are directly specified, with misinformation having a direct impact on a target candidate’s relative ranking for a given voter, rather than an indirect impact stemming from the change in perceived positions on issues, as in our model.

2 Preliminaries

We consider an election with a set of n candidates $\mathcal{C} = \{c_1, \dots, c_n\}$ and m voters $\mathcal{V} = \{v_1, \dots, v_m\}$. Following spatial voting theory [Enelow and Hinich, 1984], we associate each candidate and voter with a d -dimensional vector corresponding to their positions (opinions) on issues, that is, $c_i, v_j \in \mathcal{I} \subseteq \mathbb{R}^d$. Each voter v_j ranks candidates in $\mathcal{C}$ according to their l_p distance from v_j , $\|v_j - c_i\|_p$, with $1 \leq p \leq \infty$ an integer; the closest candidate is ranked 1, and the farthest is ranked n in the list of v_j ’s preferences. If not mentioned, the parameters of the problem (e.g., $|\mathcal{V}|$, $|\mathcal{C}|$ and d) are arbitrary. We assume that there are no ties.

In our election control problem, the adversary has a target candidate whose voter perceptions they can manipulate. Without loss of generality, let c_1 be the target candidate. We assume that the adversary can change the perception of c_1 into $\tilde{c}_1$, subject to the constraint that $\|\tilde{c}_1 - c_1\|_p \leq \epsilon$ for $\epsilon > 0$. This “budget” constraint is natural: for example, if the means for changing perceptions is social media misinformation, the change to perception is likely gradual, and one cannot target arbitrary subsets of issues with an arbitrarily large stream of malicious content. We consider two types of control: constructive, in which the adversary’s goal is for c_1 to win the election, and destructive, where the goal is for c_1 to lose. While we assume no ties in the *actual* preference rankings, ties can arise due to adversarial activities; in that case, we always break ties in the adversary’s favor.

We consider election control problems for arbitrary *scoring rules*. In scoring rules, each candidate c_i ranked t_{ij} by voter v_j receives a score $f(t_{ij})$, where $f : [n] \rightarrow \mathbb{R}$ is a non-increasing function. c_i then receives a total score $s_i = \sum_j f(t_{ij})$ from all voters, and the candidate with the highest score s_i wins the election. Many common voting rules are positional, such as plurality ($f(1) = 1$, and $f(t) = 0$ for $t \neq 1$), veto ($f(n) = 0$ and $f(t) = 1$ for $t \neq n$), Borda ($f(t) = n - t$), and k -approval (for some $1 \leq k \leq n$, $f(t) = 1$ for all $t \leq k$, and $f(t) = 0$ for $t > k$). Note that plurality is a special case of k -approval with $k = 1$.

We study the problem both when issues are binary, i.e., $\mathcal{I} = \{0, 1\}^d$, salient if issues are framed in the form of yes-no questions, such as “do you support leaving the European Union?”, and when issues are real-valued ($\mathcal{I} = \mathbb{R}^d$).

3 Binary-Valued Issues

We begin by studying a special case of our problem in which the issues are binary, that is, $\mathcal{I} = \{0, 1\}^d$, a variant we call Binary Value Perception Manipulation (BVPM).

Definition 3.1 (BVPM). *Given a set of candidates $\mathcal{C}$, voters $\mathcal{V}$, and d issues, is there a $\tilde{c}_1 \in \mathcal{I} = \{0, 1\}^d$ where $\|\tilde{c}_1 - c_1\|_p \leq \epsilon$ for $\epsilon > 0$ such that $\tilde{c}_1$ wins the election?*

Note that for l_∞ the problem is trivial: either $\epsilon \geq 1$, in which case we can set c_1 to match any currently winning candidate (and c_1 wins by best-case tiebreaking), or $\epsilon < 1$ and we cannot change c_1 . Thus all the results in this section are for an arbitrary l_p norm for p integer, $1 \leq p < \infty$. Without loss of generality (since the label of 1 or 0 for each issue is arbitrary), we assume that c_1 takes a position labeled as 1 for each issue, i.e., $c_1 = [1, \dots, 1]$.

We begin by showing that even with 2 candidates and majority voting the BVPM problem is NP-complete. We reduce from Binary Issue Selection Control (BISC), shown by Lu *et al.* [2019] to be NP-Complete with best-case tie-breaking even when $|\mathcal{C}| = 2$.

Definition 3.2 (BISC [Lu *et al.*, 2019]). *Given a set of candidates $\mathcal{C}$, voters $\mathcal{V}$, and d issues, is there a nonempty subset of binary issues $S \subseteq [1 : d]$ such that a target candidate c_1 wins the plurality election?*

Theorem 1. *BVPM is NP-complete for constructive and destructive control even with 2 candidates and majority voting.*

Proof. It is easy to check if a given $\tilde{c}_1$ wins the election; thus, BVPM is in NP. We now show hardness for constructive control by reduction from BISC. Let $\epsilon = (d-1)^{1/p}$, target candidate $c_1 = [1, 1, \dots, 1]$ and rival candidate $c_2 = [0, 0, \dots, 0]$. Let the voter set of BVPM be the same as the one in BISC.

Suppose 2-candidate BISC has a solution $S \subseteq [1 : d]$, $S \neq \emptyset$ that will let c_1 win the election. For any voter $v_j \in \mathcal{V}$, $j \in [m]$ that votes for target candidate in BISC, by the problem definition of BISC, we have $\sum_{k \in S} |c_{1,k} - v_{j,k}|^p \leq \sum_{k \in S} |c_{2,k} - v_{j,k}|^p$. We set $\tilde{c}_{1,k} = 1$ if $k \in S$; $\tilde{c}_{1,k} = 0$ if $k \in S^c$. Since $S \neq \emptyset$, $|S^c| \leq d-1$, $\tilde{c}_1$ is within the budget constraint. For $k \in S$, we have $\sum_{k \in S} |\tilde{c}_{1,k} - v_{j,k}|^p = \sum_{k \in S} |c_{1,k} - v_{j,k}|^p \leq \sum_{k \in S} |c_{2,k} - v_{j,k}|^p$. For $k \in S^c$, since $\tilde{c}_{1,k} = c_{2,k} = 0$, we then have $|\tilde{c}_{1,k} - v_{j,k}|^p = |c_{2,k} - v_{j,k}|^p$. Thus we have $\sum_{k=1}^d |\tilde{c}_{1,k} - v_{j,k}|^p \leq \sum_{k=1}^d |c_{2,k} - v_{j,k}|^p$. This means voter v_j votes for $\tilde{c}_1$ in BVPM ($\|\tilde{c}_1 - v_j\|_p \leq \|c_2 - v_j\|_p$). $\tilde{c}_1$ wins the election and is a solution to two candidates BVPM.

If two candidates BVPM has a solution $\tilde{c}_1$ that wins the election, by the problem definition of BVPM, we have $\sum_{k=1}^d |\tilde{c}_{1,k} - v_{j,k}|^p \leq \sum_{k=1}^d |c_{2,k} - v_{j,k}|^p$. We let $S = \{k \in [d] \mid \tilde{c}_{1,k} = 1\}$. Since $\tilde{c}_{1,k} = c_{2,k} = 0$, $k \in S^c$, we have $\sum_{k \in S^c} |\tilde{c}_{1,k} - v_{j,k}|^p = \sum_{k \in S^c} |c_{2,k} - v_{j,k}|^p$; since $\tilde{c}_{1,k} = c_{1,k} = 1$, $k \in S$, we then have $\sum_{k \in S} |c_{1,k} - v_{j,k}|^p = \sum_{k \in S} |\tilde{c}_{1,k} - v_{j,k}|^p \leq \sum_{k \in S} |c_{2,k} - v_{j,k}|^p$. Due to the budget constraint $\epsilon = (d-1)^{1/p}$, we must have $S \neq \emptyset$, which satisfies the BISC solution requirement. S is a solution set to two candidates BISC.

For destructive control, the same argument applies after switching the positions of c_1 and c_2 . $\square$

While BVPM is hard in general, we next show that the problem is tractable for a constant number of voters.² While at first glance a constant number of voters seems an impractical restriction, we subsequently show that this result offers real insight even when the number of voters is arbitrary.

Theorem 2. *When the number of voters is constant, BVPM can be solved in polynomial time for arbitrary scoring rules, for both constructive and destructive control.*

Proof sketch. Given an issue j , let $\mathbf{v}^j$ be a vector corresponding to the position of each voter on issue j . The key idea is that when the number of voters m is a constant, there is also a constant number $l \leq 2^m$ of issue equivalence sets, where

²Note that this is trivial for a constant number of binary issues.

an equivalence set I is a set of issues with identical $\mathbf{v}^j$ (consequently, each issue in I has an identical and interchangeable impact on the election outcome). Since the number of issues of each equivalence set we are allowed to flip is at most $\lfloor \epsilon^p \rfloor \leq d$, it is direct that we can exhaustively enumerate all $O(d^{2^m})$ possibilities for arbitrary ϵ , which is polynomial since m is constant. $\square$

While there is a simple poly-time algorithm for solving BVPM, we can actually considerably improve on its time complexity by leveraging additional problem structure. We begin with the constructive control case for arbitrary scoring rules. A key feature of arbitrary scoring rules is that as long as c_1 receives (one of) the highest scores, c_1 wins the election. Since the distances between each voter v_j and all candidates c_i other than c_1 are fixed, the relative rankings of candidates c_i ($i \geq 2$) w.r.t. voter v_j , $j \in [m]$ are fixed as well. Given the rankings of candidates c_i ($i \geq 2$) w.r.t. to voter v_j from closest to furthest as $c_{i_1}, \dots, c_{i_{m-1}}$, we can enumerate all scenarios of insertion positions of c_1 in this sequence. We denote the final ranking of c_1 after insertion w.r.t. voter v_j as r_j , meaning c_1 will receive a score of $f(r_j)$ from v_j , $j \in [m]$. By going through all scenarios of c_1 getting a final ranking position of $(r_1, r_2, \dots, r_m)$, $r_j \in [n]$, $\forall j \in [m]$, which corresponds to c_1 getting a score of $(f(r_1), f(r_2), \dots, f(r_m))$, we cover all the scenarios of c_1 winning. As shown in Lemma 1, this is equivalent to enumerating all scenarios of c_1 getting a ranking position *higher* than $r_j \in [n]$ w.r.t. v_j for $j \in [m]$. The missing proofs of this and other results (referenced below) are provided in the Supplement (<https://arxiv.org/abs/2205.00102>).

Lemma 1. *For constructive control, if a ranking position $(r_1, r_2, \dots, r_m)$ is feasible for c_1 under the budget constraint and lets c_1 win the election, then for a ranking position $(r'_1, r'_2, \dots, r'_m)$ that is feasible with $r'_j \leq r_j$, $\forall j \in [m]$, it will also let c_1 win the election.*

This means by enumerating all scenarios of c_1 getting a ranking position *higher* than $(r_1, r_2, \dots, r_m)$, which corresponds to c_1 getting a score of *at least* $(f(r_1), f(r_2), \dots, f(r_m))$, we cover all the possible scenarios of c_1 winning. In fact, we can further simplify the calculation by only enumerating some ranking positions each corresponding to a *unique* f score value. Given an arbitrary scoring function f that has r unique values ($|f_{\text{uniq}}| = r$), since f is non-increasing, we can partition the domain of f by $0 = s_0 < s_1 < \dots < s_r = n$, so that $\{f(k)\}_{k=s_{i+1}}^{s_i+1}$ have the same value, $i \in \{0, 1, \dots, r-1\}$. This means s_{i+1} is the lowest ranking position that corresponds to score $f(s_{i+1})$ and $\{f(s_i)\}_{i=1}^r$ contains all the unique values of f .

Lemma 2. *For constructive control, by enumerating all scenarios of c_1 getting a ranking position higher than t_j w.r.t. voter v_j for $t_j \in \{s_1, \dots, s_r\}$, $\forall j \in [m]$, we cover all the possible scenarios of c_1 winning.*

Next we solve the problem for each $(t_1, \dots, t_m)$ scenario with $t_j \in \{s_1, \dots, s_r\}$, $\forall j \in [m]$. For each voter v_j , $j \in [m]$, we rank candidates c_i ($i \geq 2$) by their distances to voter v_j from closest to furthest, and use $d_j^{t_j}$ to denote the distance between v_j and the candidate ranked t_j -th closest to it. Since

the tie breaks in the adversary's favor, as long as c_1 's distance to v_j is no more than $d_j^{t_j}$, c_1 will receive a score of at least $f(t_j)$ from v_j . Notice that since the rankings of c_i ($i \geq 2$) do not include c_1 , only $d_j^{t_j}$ for $1 \leq t_j \leq n-1$ are properly defined. For $t_j = s_r = n$, we let $d_j^n = +\infty$, since c_1 is guaranteed to get at least the lowest score $f(n)$.

For each scenario, the problem can be represented as the following integer linear constraint problem:

$$x_i \leq \min\{b_i, \lfloor \epsilon^p \rfloor\}, \quad 1 \leq i \leq 2^m \quad (1a)$$

$$\sum_{i=1}^{2^m} z_{ij} x_i + (d_j^0)^p \leq (d_j^{t_j})^p, \quad 1 \leq j \leq m \quad (1b)$$

$$\sum_{i=1}^{2^m} x_i \leq \lfloor \epsilon^p \rfloor \quad x_i \in \mathbb{Z}^+, \quad (1c)$$

where x_i is the number of issues in an issue equivalence class i that we want to flip to 0, b_i is the size of the i -th issue equivalence class (where $\sum_i b_i = d$), d_j^0 is the original distance between target candidate c_1 and voter v_j . $z_{ij} \in \{-1, +1\}$ is the sign of impact of flipping the issue: $z_{ij} = +1$ if previously the j -th voter in any issue in the i -th equivalence class is 1 (since flipping the c_1 to 0 will increase the distance) and $z_{ij} = -1$ if previously it is 0, since flip the target candidate to 0 will decrease the distance. Since the size of the input of this integer feasibility problem is $O(\log(d))$ and the number of variables is constant, it can be solved in time $O(\log(d))$ [Lokshtanov, 2009, Theorem 2.8.1].

The total number of times we need to run the ILP and check whether $\tilde{c}_1$ wins the election is bounded by the number of unique $(t_1, \dots, t_m)$ scenarios, which is $|f_{\text{uniq}}|^m$. For an arbitrary scoring function f , the time complexity for calculating voter-candidate distances is $O(nd)$, ranking distances takes $O(n \log(n))$ time. The calculation of the issue equivalence sets takes $O(d)$ time. For each scenario, solving the ILP takes $O(\log(d))$ time; calculating the distance between $\tilde{c}_1$ and all voters takes $O(d)$ time; checking whether $\tilde{c}_1$ wins the election takes $O(n)$ time. The total time complexity of the algorithm is $O(n(d + \log(n)) + |f_{\text{uniq}}|^m(d + n))$.

If, in addition, $|f_{\text{uniq}}| = r$ is constant (e.g., for plurality), then the number of scenarios $|f_{\text{uniq}}|^m$ is constant. Moreover, we do not need to do a total sort for the distances. Through finding the t_j -th order statistics and Quicksort Partition, the total time complexity of the algorithm is $O(nd)$.

In Supplement B we present a similar analysis and algorithm as above for destructive control. In either case, the complexity is linear in the dimension d of the issue space.

As noted earlier, considering a constant number of voters may seem unrealistic. However, we note that these algorithms are straightforward to generalize to a setting with an *arbitrary number of voters*, but in which the positions of voters on issues, v_j , can only take on values from a small collection of possibilities Q (that is, $|Q|$ is a constant, and for each voter j , $v_j \in Q$). Specifically, the only change is to calculate the *weighted* final score in each case above, where the weight for each distinct voter position (opinion) type $q \in Q$ is the number of voters j with position $v_j = q$. This is expressed in the following corollary.

Corollary 1. *BVPM can be solved in polynomial time when the number of distinct voter opinions is constant for constructive and destructive control for arbitrary scoring rules.*

The key insight that our results offer is that *the complexity of manipulating elections by changing voter perceptions about candidates hinges on the diversity of opinions among voters*. In particular, when voters hold a broad diversity of views, manipulation is intractable; if, in contrast, voters are siloed into a relatively small collection of echo chambers, subverting elections becomes easy. Below, we show that this observation extends to real-valued issues.

4 Real-Valued Issues

Next, we turn to Real Value Perception Manipulation (RVPM), the problem identical to BVPM except that now the issue space $\mathcal{I}$ is real-valued.

4.1 Hardness Results

We begin by showing that nearly every variant of RVPM is, in general, computationally intractable. First, we show that the election control problem is hard under l_p norm with integer $1 < p \leq \infty$ under destructive control even with 2 candidates and majority voting, and constructive control even for plurality voting. Nonetheless, we show that constructive control with l_∞ norm and only two candidates is in P.

Our hardness result for destructive control uses a reduction from 3-SAT, the proofs are deferred to the supplement (A.3 and A.4).

Theorem 3. *The destructive control variant of RVPM is NP-complete under l_p norm for integer $1 < p \leq \infty$ even with two candidates and majority voting.*

The following theorem (proved in Supplements A.5 and A.6) shows that constructive control is also hard.

Theorem 4. *The constructive control variant of RVPM under l_p norm for integer $1 < p \leq \infty$ is NP-complete for plurality voting.*

Note that Theorem 4 is stated for an arbitrary number of candidates. For two candidates and l_∞ norm, however, constructive variant of RVPM is easy.

Theorem 5. *The constructive control l_∞ norm variant of RVPM with 2 candidates can be solved in time $O(md)$ for arbitrary scoring rules.*

The proof is provided in Supplement A.7. Note that our results do not resolve the question of constructive control with $p < \infty$; we leave it as an open question.

Next, we consider two restrictions of RVPM: 1) assuming a constant number of issues, and 2) assuming a constant number of distinct voter opinions. In all these restricted cases, we show how to solve RVPM in polynomial time for l_2 and l_∞ norm. We leave the problem open for arbitrary l_p norms.

4.2 Constant Number of Issues

When the number of issues is constant, we show that RVPM is tractable for l_2 and l_∞ norms (our focus on these two norms follows the precedent from prior literature [Crama et al., 1995; Crama and Ibaraki, 1997]). Note that unlike with

binary issues, tractability of RVPM with a constant number of issues is *non-trivial* since the issue space is continuous and cannot be exhaustively searched in finite time.

Constructive Control

We start by studying the problem of constructive control, with RVPM in that case closely related to the well-known *product positioning* and *ball intersection problems* [Crama et al., 1995]. The goal in product positioning is to find $x \in \mathbb{R}^m$ that maximizes the number of consumers for whom x is closer (in l_p) to their ideal product than any of the competitors. The ball intersection problem aims to maximize the weighted sum of l_p balls to which x belongs. Neither exactly captures our problem given the presence of the attacker budget constraint and different scoring scenarios, but both are useful tools in constructing the algorithms for our problem below.

Theorem 6. *RVPM can be solved in polynomial time under l_2 norm when the number of issues is constant for constructive control for arbitrary scoring rules.*

Proof. We first convert our problem into the ball intersection problem. Let $B_j^{t_j}$ be the ball that corresponds to voter v_j , with radius $d_j^{t_j}$ as defined in the constructive control variant of BVPM, that is,

$$B_j^{t_j} = \{\tilde{c}_1 \in \mathbb{R}^d \mid \|\tilde{c}_1 - v_j\|_2 \leq d_j^{t_j}\}.$$

Similarly, we define the candidate budget ball

$$B_c = \{\tilde{c}_1 \in \mathbb{R}^d \mid \|\tilde{c}_1 - c_1\|_2 \leq \epsilon\}.$$

Since the tie breaks in the adversary's favor, $\tilde{c}_1$ will receive a score of at least $f(t_j)$ from voter v_j iff $\tilde{c}_1$ falls within $B_j^{t_j}$ and B_c . According to Lemma 2 (which also applies when issues are real-valued), by finding a representative point $\tilde{c}_1$ within $\{B_c\} \cup (\bigcup_{j=1}^m \{B_j^{t_j}\})$ for all scenarios of $(t_1, \dots, t_m)$ with $t_j \in \{s_1, \dots, s_r\}$ (partition of the domain of f based on unique values as defined in the constructive control variant of BVPM), $\forall j \in [m]$, we cover all the scenarios of c_1 winning. We can now directly apply the ball intersection algorithm by Crama et al. [1995] for l_2 and constant d to our problem only once for $\{B_c\} \cup (\bigcup_{j=1}^m \bigcup_{l=1}^r \{B_j^{s_l}\})$, the resulting set P (representative points of intersections) contains a representative point for all the scenarios. We check for each point in P whether it is within B_c and wins the election. The time complexity of the algorithm is exponential only in d . $\square$

A similar result, based on a similar connection to box intersection, can be obtained for the l_∞ norm.

Theorem 7. *RVPM can be solved in polynomial time under l_∞ norm when the number of issues is constant for constructive control for arbitrary scoring rules.*

Destructive Control

Next we study the problem under destructive control. Define an open voter ball corresponding to voter v_j with radius $d_j^{t_j}$ under l_2 norm as $\hat{B}_j^{t_j} = \{\tilde{c}_1 \in \mathbb{R}^d \mid \|\tilde{c}_1 - v_j\|_2 < d_j^{t_j}\}$, and recall that $d_j^{t_j}$ is the distance between v_j and the candidate ranked t_j -th closest to it. The next lemma, proved in the Supplement, provides an important building block.

Lemma 3. *Given k open balls $\{\hat{B}_1, \dots, \hat{B}_k\}$ and a closed ball B_c , let P be the representative points of intersections (defined in Crama et al. [1995]) w.r.t. $\{B_c, B_1, \dots, B_k\}$, where B_j is the closed ball corresponds to $\hat{B}_j$, $\forall j \in [k]$. For any given family of open balls $\{\hat{B}_{i_1}, \dots, \hat{B}_{i_r}\} \subseteq \{\hat{B}_1, \dots, \hat{B}_k\}$, let $P' = \{x \mid x \notin \hat{B}_j, \forall j \in \{i_1, \dots, i_r\}, x \in B_c\}$. If $P' \neq \emptyset$, then $P \cap P' \neq \emptyset$.*

Next, we show that when the number of issues is constant, the destructive control variant of RVPM is tractable.

Theorem 8. *When the number of issues is constant, the destructive control variant of RVPM can be solved in polynomial time under l_2 norm for arbitrary scoring rules.*

Proof. Since the tie breaks in the adversary's favor, within B_c , $\tilde{c}_1$ will get a score of no more than $f(t_j)$ from voter v_j iff it falls outside of $\hat{B}_j^{t_j}$. Similar to Theorem 6, we cover all the scenarios of c_1 losing by finding the set that contains a representative point within B_c that falls outside of $\bigcup_{j=1}^m \{\hat{B}_j^{t_j}\}$ for all scenarios of $(t_1, \dots, t_m)$ with $t_j \in \{s_1, \dots, s_r\}$ (partition of the domain of f based on unique values as defined in the destructive control variant of BVPM), $\forall j \in [m]$. Lemma 3 shows us that the set P (representative points of intersections) for the family of balls $\{B_c\} \cup (\bigcup_{j=1}^m \bigcup_{l=1}^r \{\hat{B}_j^{s_l}\})$ contains a representative point for all the scenarios. The problem could be solved in polynomial time with minor modifications to the algorithm in Theorem 6. $\square$

The destructive control problem with l_∞ norm involves solving a non-convex feasibility problem. The next lemma shows that the problem of relevance can nevertheless be solved in polynomial time.

Lemma 4. *For the feasibility problem*

$$\|\tilde{y} - y\|_\infty \leq \epsilon \quad (2a)$$

$$\|\tilde{y} - a_i\|_\infty \geq b_i, \quad i \in [k] \quad (2b)$$

$$\epsilon > 0, b_i > 0 \quad \tilde{y}, y, a_i \in \mathbb{R}^d, \quad (2c)$$

if $\tilde{y} \in \mathbb{R}^d$ satisfy constraint (2a) and $\bigcup_{j=1}^d S_j(\tilde{y}_j) = [k]$, where $S_j(\tilde{y}_j) = \{i \in [k] \mid |\tilde{y}_j - a_{i,j}| \geq b_i\}$, $\forall j \in [d]$, then $\tilde{y}$ is a solution to the feasibility problem. Moreover, for all $j \in [d]$, let set $P_j = \bigcup_{i=1}^k (\{-b_i + a_{i,j}, b_i + a_{i,j}\} \cap [-\epsilon + y_j, \epsilon + y_j])$, or $P_j = \{y_j\}$ if the set is empty, then $P = \{p \in \mathbb{R}^d \mid p_j \in P_j, \forall j \in [d]\}$ contains a representative solution point $\tilde{y}$ to the problem.

We use Lemma 4 to show that the destructive control variant of RVPM is also tractable for the l_∞ norm.

Theorem 9. *RVPM can be solved in polynomial time under l_∞ norm when the number of issues is constant for destructive control for arbitrary scoring rules.*

Proof. We solve the problem by using Lemma 4 to find the set P which contains a representative solution point for all the scenarios of $(t_1, \dots, t_m)$ (defined as in Theorem 8). Each scenario represents finding a $\tilde{c}_1$ within the budget constraint that gets a score of no more than $f(t_j)$ from voter v_j , $j \in [m]$:

$$\|\tilde{c}_1 - c_1\|_\infty \leq \epsilon \quad (3a)$$

$$\|\tilde{c}_1 - v_j\|_\infty \geq d_j^{t_j}, \quad j \in [m] \quad (3b)$$

where $d_j^{t_j}$ is defined in the destructive control variant of BVPM. Notice that there are in total $m \cdot |f_{\text{uniq}}|$ open hypercubes involved. According to Lemma 4, the representative solution set P for all scenarios has at most $2m \cdot |f_{\text{uniq}}|$ choices for each dimension. The time complexity of the algorithm is exponential only in d . $\square$

4.3 Constant Number of Distinct Voters

Next, we turn to the case when the number of distinct voter opinion vectors v_j is bounded by a constant. As in Section 3, we simplify the discussion by assuming that the number of voters is constant; generalization to an arbitrary number of voters whose opinions can be grouped into a small set Q of possibilities is straightforward using the same idea as for BVPM. For l_2 norm, we use the ball intersection algorithm as in Section 4.2.

Theorem 10. *When the number of voters is constant, the constructive and destructive control variants of RVPM can be solved in polynomial time under the l_2 norm for arbitrary scoring rules.*

Proof. For constructive control, we solve the problem for each scenario of $(t_1, \dots, t_m)$ (defined as in Theorem 6) through finding a representative point within $\{B_c\} \cup (\bigcup_{j=1}^m \{B_j^{t_j}\})$, which takes $O(d^3)$ time using a variation of the ball intersection algorithm by Crama *et al.* [1995] as in Theorem 6. Since there are in total $O(|f_{\text{uniq}}|^m)$ scenarios, the problem can be solved in polynomial time. For destructive control, a similar argument holds and the problem can be solved by using a variation of the ball intersection algorithm as in Theorem 8. $\square$

For l_∞ norm, the constructive case can be solved by the application of linear programming.

Theorem 11. *When the number of voters is constant, the constructive control variant of RVPM can be solved in polynomial time under the l_∞ norm for arbitrary scoring rules.*

Proof. For each scenario of $(t_1, \dots, t_m)$ (defined as in Theorem 6), we solve the below linear programming:

$$\|\tilde{c}_1 - c_1\|_\infty \leq \epsilon \quad (4a)$$

$$\|\tilde{c}_1 - v_j\|_\infty \leq d_j^{t_j}, \quad j \in [m] \quad (4b)$$

Since each LP has $O(d)$ linear constraints, and there are $O(|f_{\text{uniq}}|^m)$ scenarios, the problem is polynomial time solvable. Alternatively, we can check the interval endpoints similar to Lemma 4 for each dimension. Since m is a constant, the algorithm returns $\tilde{c}_1$ in $O(d)$ time if a solution to the LP exists, or NO if not feasible. $\square$

The destructive case is somewhat more involved. We begin with a lemma that again shows that a non-convex feasibility problem we need to solve is tractable.

Lemma 5. *The feasibility problem in Lemma 4 can be solved in linear time if the number of constraints k is constant.*

Proof sketch. According to Lemma 4, for each dimension $j \in [d]$, we could compute set P_j and its corresponding set $S_j = \{S_j(p_j) \mid p_j \in P_j\}$. Finding a solution to the feasibility problem is equivalent to finding $S_j(p_j) \in S_j$ for all $j \in [d]$, with some $l \leq d$ that satisfy $\bigcup_{j=1}^l S_j(p_j) = [k]$, and $[p_1, p_2, \dots, p_l, y_{l+1}, y_{l+2}, \dots, y_d]$ is a solution. Since k is constant, the algorithm is linear and of complexity $O(d)$. We can also determine in linear time if none of the representative points satisfy the feasibility condition, and return NO in that case. The detailed algorithm and proof are provided in Supplement C. $\square$

Theorem 12. *When the number of voters is constant, the destructive control variant of RVPM can be solved in polynomial time under l_∞ norm for arbitrary scoring rules.*

Proof. Using Lemma 5, we can solve the linear feasibility problem in Equation (3) where m is constant for each scenario of $(t_1, \dots, t_m)$ (defined as in Theorem 8). Each feasibility problem can be solved in $O(d)$ time and there are $O(|f_{\text{uniq}}|^m)$ scenarios. The full time complexity of the algorithm is $O(n(d + \log(n)) + |f_{\text{uniq}}|^m(d + n))$, or $O(nd)$ if $|f_{\text{uniq}}|$ is constant. $\square$

We can extend the results to an arbitrary number of voters with a constant number of distinct opinions by the same argument as for BVPM.

5 Conclusion

We model the impact of political misinformation on elections as election control in the spatial model of voting in which an adversary manipulates perceptions of the positions of a target candidate by the voters. Our central observation, which obtains both when issues are real-valued and binary, and for different ways we can measure distance in generating preferences over candidates based on their relative positions to voters, is that what matters is the extent of opinion diversity in the voting population. Specifically, when voter positions on issues are highly diverse, the manipulation problem is intractable in most settings. In contrast, when voter views can be reduced by a small number of opinion groups, the control problem becomes linear in dimension when issues are binary, and polynomial with real-valued issues. Our characterization of the complexity landscape leaves several open questions, such as hardness of constructive control with two candidates in the setting with real-valued issues (we only show that it is tractable for l_∞ norm). Furthermore, our negative results for real-valued issues do not address the case of l_1 , while our positive results only apply to l_2 and l_∞ in this setting.

Our model has several important limitations that suggest further useful future directions. First, we assume that the same norm is used both by voters to rank candidates, and to limit the extent of perception manipulation; however, these distances may often be useful to measure in different ways. Second, we assume that perception manipulation has identical impact on all voters. A more sophisticated model would blend this with the election control approaches in which misinformation spreads through a social network, with only a subset of voters impacted, potentially in different ways.

Acknowledgements

This work was partially supported by the National Science Foundation (grants IIS-1905558, IIS-1903207, and IIS-1939677) and Amazon.

References

- [Anshelevich and Postl, 2016] Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. In *International Joint Conference on Artificial Intelligence*, page 46–52, 2016.
- [Anshelevich et al., 2018] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [Bartholdi et al., 1992] John J. Bartholdi, Craig A. Tovey, and Michael A. Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8):27–40, 1992.
- [Binkele-Raible et al., 2014] Daniel Binkele-Raible, Gábor Erdélyi, Henning Fernau, Judy Goldsmith, Nicholas Mattei, and Jörg Rothe. The complexity of probabilistic lobbying. *Discrete Optimization*, 11:1–21, 2014.
- [Bredereck et al., 2016] Robert Bredereck, Piotr Faliszewski, Rolf Niedermeier, and Nimrod Talmon. Large-scale election campaigns: Combinatorial shift bribery. *Journal of Artificial Intelligence Research*, 55:603–652, 2016.
- [Caldwell et al., 2019] L. A. Caldwell, H. Przybyla, and K. Stewart. Senate intelligence report finds ‘extensive’ Russian election interference. In *NBC News*. 2019.
- [Castiglioni et al., 2020] Matteo Castiglioni, Diodato Ferrioli, and Nicola Gatti. Election control in social networks via edge addition or removal. In *AAAI Conference on Artificial Intelligence*, pages 1878–1885, 2020.
- [Chen et al., 2017] Jiehua Chen, Piotr Faliszewski, Rolf Niedermeier, and Nimrod Talmon. Elections with few voters: Candidate control can be easy. *Journal of Artificial Intelligence Research*, 60(1):937–1002, 2017.
- [Christian et al., 2007] Robin Christian, Mike Fellows, Frances Rosamond, and Arkadii Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.
- [Crama and Ibaraki, 1997] Yves Crama and Toshihide Ibaraki. Hitting or avoiding balls in euclidean space. *Annals of Operations Research*, 69:47–64, 1997.
- [Crama et al., 1995] Yves Crama, Pierre Hansen, and Brigitte Jaumard. Complexity of product positioning and ball intersection problems. *Mathematics of Operations Research*, 20(4):885–894, 1995.
- [Davis and Hinich, 1968] Otto A. Davis and Melvin J. Hinich. On the power and importance of the mean preference in a mathematical model of democratic choice. *Public Choice*, 5(1):59–72, 1968.
- [Enelow and Hinich, 1984] James M. Enelow and Melvin J. Hinich. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press, 1984.
- [Erdélyi et al., 2015] Gábor Erdélyi, Michael R. Fellows, Jörg Rothe, and Lena Schend. Control complexity in bucklin and fallback voting: A theoretical analysis. *Journal of Computer and System Sciences*, 81(4):632–660, 2015.
- [Estornell et al., 2020] Andrew Estornell, Sanmay Das, Edith Elkind, and Yevgeniy Vorobeychik. Election control by manipulating issue significance. In *Conference on Uncertainty in Artificial Intelligence*, pages 340–349, 2020.
- [Faliszewski and Rothe, 2016] Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 146–168. Cambridge University Press, 2016.
- [Harper et al., 2019] T. Harper, C. Wheeler, and R. Kerbaj. Revealed: the Russia report. In *The Sunday Times*. 2019.
- [Hemaspaandra et al., 2007] Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171:255–285, 2007.
- [Kempe et al., 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.
- [Khetani-Shah and Deutsch, 2019] S. Khetani-Shah and J. Deutsch. Brexit timeline: From referendum to EU exit. In *Politico Pro*. 2019.
- [Lokshtanov, 2009] Daniel Lokshtanov. *New methods in parameterized algorithms and complexity*. PhD thesis, University of Bergen, Norway, 2009.
- [Lu et al., 2019] Jasper Lu, David Kai Zhang, Zinovi Rabinovich, Svetlana Obraztsova, and Yevgeniy Vorobeychik. Manipulating elections by selecting issues. In *International Conference on Autonomous Agents and Multiagent Systems*, page 529–537, 2019.
- [Menton, 2012] Curtis Menton. Normalized range voting broadly resists control. *Theory of Computing Systems*, 53(4):507–531, 2012.
- [Robb, 2017] Amanda Robb. Anatomy of a fake news scandal. In *Rolling Stone*. 2017.
- [Wilder and Vorobeychik, 2018] Bryan Wilder and Yevgeniy Vorobeychik. Controlling elections through social influence. In *International Conference on Autonomous Agents and Multiagent Systems*, page 265–273, 2018.