Detecting Approximate Clones under Approval Voting

Extended Abstract

Théo Delemazure
ILLC, University of Amsterdam
Amsterdam, Netherlands
theo.delemazure@uva.nl

Piotr Faliszewski
AGH University of Krakow
Krakow, Poland
faliszew@agh.edu.pl

Łukasz Janeczko
AGH University of Krakow
Krakow, Poland
ljaneczk@agh.edu.pl

Dušan Knop
Czech Technical University in Prague
Prague, Czech Republic
dusan.knop@fit.cvut.cz

Kristýna Pekárková
AGH University of Krakow
Krakow, Poland
pekarkova@agh.edu.pl

Jan Pokorný
Czech Technical University in Prague
Prague, Czech Republic
jan.pokorny@fit.cvut.cz

Šimon Schierreich
AGH University of Krakow
Krakow, Poland
schiesim@fit.cvut.cz

Ildikó Schlotter
ELTE Centre for Economic and Regional Studies
Budapest, Hungary
schlotter.ildiko@krtk.elte.hu

ABSTRACT

In approval elections, two candidates are called perfect clones if they are approved by exactly the same set of voters. We propose a general framework for studying *approximations* of this notion, and demonstrate its power using two natural approximation measures with various appealing axiomatic properties. For both of these measures, we consider two fundamental tasks: deciding whether a large approximate clone set exists in a given election, and computing a partition of the candidate set into approximate clone sets. We show that both tasks are, in general, computationally intractable. To have a better understanding of the boundary between tractable and intractable instances, we analyze the parameterized complexity of these problems with respect to several parameters, including the number of voters and candidates, the approximation threshold, the number and size of partition parts, and structural properties of the instances, such as the number of approvals per voter or per candidate. Finally, we explore how our approximation measures behave in real-world approval elections.

KEYWORDS

Candidate Clones, Approximate Clones, Approval Elections, Computational Complexity, Fixed-Parameter Tractability

ACM Reference Format:
Théo Delemazure, Piotr Faliszewski, Łukasz Janeczko, Dušan Knop, Kristýna Pekárková, Jan Pokorný, Šimon Schierreich, and Ildikó Schlotter. 2026. Detecting Approximate Clones under Approval Voting: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/MBRB4492

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/MBRB4492

1 INTRODUCTION

We study the complexity of identifying clones in approval elections, i.e., in elections where voters express preferences by indicating which candidates they find sufficiently appealing and which ones they don't, i.e., a voter's ballot contains the candidates they approve of. By a clone—or, more precisely, a *perfect clone set*—we mean a set of candidates that are approved by exactly the same voters [1, 8]. However, while identifying perfect clones in an election is easy—indeed, given some candidate we can test which other ones are approved by exactly the same voters—the notion is so demanding that we hardly expect to see such perfect clones in elections (except for those involving very few voters or being particular in some other way. Consequently, it is far more interesting to seek *approximate clones*, i.e., sets of candidates that are approved by *similar* sets of voters. Our goal is to formalize the notion of approximate clone sets and establish the complexity of (a) deciding if a given election has an approximate clone of a given size and quality (i.e., a given level of proximity to being a perfect clone), and (b) deciding if the candidates in a given election can be partitioned into a given number of approximate clones of up to a given size and a given quality. These two types of problems were introduced by Faliszewski et al. [5]. See the recent work of Procaccia et al. [7] for a discussion of approximate clones in the context of AI alignment, and for an axiomatic and experimental analysis of approximate clones in ordinal elections, see the paper of Delemazure [2].

Why Seek Approximate Clones? One of the main reasons why seeking approximate clones is interesting is that knowing them helps in understanding the structure of a given election (similar arguments were put forward by, e.g., Janeczko et al. [6] and Faliszewski et al. [5]). For example, in a participatory budgeting (PB) election—where voters indicate which projects they would like to see funded in their city—approximate clone sets may consist of projects of similar type (e.g., building bike paths, funding festivals, etc.), located in the same area (e.g., improving safety of a given neighborhood, planting flowers in this neighborhood, providing classes for seniors in this neighborhood), or involving the same demographic (e.g., building a playground, organizing an art festival for children). After a PB election, a city might want to evaluate what projects form approximate clone sets and, e.g., investigate if many projects forming an approximate clone set are funded (under commonly used PB voting rules, a well-organized group of voters can obtain disproportionally large funding for their projects, and a city might want to detect such cases) or if none of the projects from a given approximate set are funded (this might suggest an active group of citizens who, nonetheless, may need some support).

Another application of identifying partitions of candidates into approximate clone sets stems from the fact that in approval elections candidates and voters are "symmetric." In particular, given an approval election with candidate set C and voter set V, we can form a "transposed" election where voters become candidates and are approved by those candidates for which they originally voted. Now, an approximate clone set consists of those original voters that expressed similar preferences. By finding a partition of voters into approximate clone sets, we identify a hidden party-like structure of the electorate. This is useful in a number of settings, ranging from analyzing the Supreme Court of the US to considering self-government of people in various housing estates (who, e.g., run yes/no elections to decide what improvements to implement in their buildings; concatenation of these yes/no elections gives an approval election that we can analyze).

How to Define Approximate Clones? Given two candidates, a natural measure of their closeness to being perfect clones is the number of voters that rank them differently, that is, the symmetric difference of their sets of supporters. We consider two ways of extending this measure to clone sets including more candidates: Either we sum the above measure for each pair of candidates in a given set C, leading to the measure $\mu^H_\Sigma(C)$, or we take the maximum over all pairs of candidates within C, yielding the measure $\mu^H_{\max}(C)$. Then, the lower the value of our measures for a given set of candidates, the more this set resembles a set of perfect clones.

These measures are quite particular in that they treat approvals and disapprovals symmetrically. Such an approach is natural if we know that a voter's disapproval is a conscious decision against a given candidate (as would be the case in typical yes/no elections), but raises issues when lack of an approval may indicate that a voter simply did not pay much attention to a given candidate (as may be the case in PB elections, which can have over 100 candidates). In this latter case, our measures may suggest that two candidates are approximate clones even if no voter approves both of them at the same time (e.g., if very few voters approve either of these candidates, and an overwhelming majority does not approve either). Fortunately, there is an easy solution to this problem: Given an election where we want to identify approximate clones, we may consider only candidates who receive a number of approvals from a certain range (indeed, if some candidates are approximate clones, they should receive similar numbers of approvals).

Our Contribution. We study the (parameterized) complexity of the following two problems (see Tables 1 and 2 for our results):

(1) Decide if a given election contains an approximate clone set of a given size k—i.e., consisting of k candidates—and of a given quality q—i.e., where the measure that we consider has value at most q; we call this problem Approximate μ-Clone Set.
(2) Decide if the set of candidates can be partitioned into up to t approximate clone sets, each required to have a size not exceeding a given bound b and a given quality q—the Approximate μ-Clones Partition problem.

Table 1: The complexity picture of the Approximate μ-Clone Set problem. Each cell corresponds to a combination of a clones approximation measure and a parameter (m = number of candidates, n = number of voters, q = approximation threshold, k = minimum clone set size, A_c = maximum supporter set size, A_v = maximum ballot size).

	m	n	q	k	A_c	A_v
$\mu^H_{\max}$	FPT	FPT	XP	XP + W[1]-hard	XP	NP-hard
μ^H_Σ	FPT	FPT	XP	XP + W[1]-hard	?	NP-hard

Table 2: An overview of our basic complexity results for the Approximate μ-Clones Partition problem. We use b to denote the maximum clone set size and t the maximum number of clone sets. The rest of the notation is the same as in Table 1.

	$\mu^H_{\max}$	μ^H_Σ
$q = 0$	P	P
$q \geq 1$	NP-complete	NP-complete
$b \leq 2$	P	P
$b \geq 3$	NP-complete	NP-complete
$t = 1$	P	P
$t = 2$	P	?
$t \geq 3$	NP-complete	NP-complete
$A_v = 0$	P	P
$A_v = 1$	P	NP-complete
$A_v \geq 2$	NP-complete	NP-complete
$A_c = 0$	P	P
$A_c = 1$	P	?
$A_c \geq 2$	NP-complete	NP-complete
m	FPT	FPT
n	FPT	?

For both of these problems, and both of our measures of proximity to being a perfect clone set, we find strong hardness results. For example, while there is a polynomial-time algorithm for the latter of our problems for the case of perfect clones, allowing even a minimal level of imperfection leads to intractability. On the positive side, we do find a few XP and FPT results (e.g., for parameterizations either by the number of candidates or the number of voters).

We supplement our theoretical results with an analysis of a number of real-life elections. In particular, we seek approximate clones in the yes/no elections performed in the US Supreme Court [3] and in the PB elections from Pabulib [4].

ACKNOWLEDGMENTS

This project was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854), by the European Union under the project Robotics and Advanced Industrial Production (reg. no. CZ.02.01.01/00/22_008/0004590), by the European Research Council (ERC) Synergy Grant ADDI, ID: 10.3030/101166894, and by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID: 90254), project OPEN-34-18. JP acknowledges additional support from the Grant Agency of CTU in Prague, grant No. SGS23/205/OHK3/3T/18. IS is supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2021-2) and its János Bolyai Research Scholarship.

REFERENCES

[1] F. Brandl and D. Peters. 2022. Approval Voting Under Dichotomous Preferences: A Catalogue of Characterizations. *Journal of Economic Theory* 205, Article 105532 (2022), 26 pages. https://doi.org/10.1016/j.jet.2022.105532

[2] T. Delemazure. 2026. Independence of Approximate Clones With Ordinal Preferences. In *Proceedings of AAMAS-2026*. IFAAMAS, Richland, SC, USA. https://doi.org/10.65109/DGQX4543

[3] T. Delemazure, C. Dong, D. Peters, and M. Tydrichová. 2024. Comparing Ways of Obtaining Candidate Orderings From Approval Ballots. In *Proceedings of IJCAI-2024*. ijcai.org, 2757–2765. https://doi.org/10.24963/ijcai.2024/305

[4] P. Faliszewski, J. Flis, D. Peters, G. Pierczyński, P. Skowron, D. Stolicki, S. Szufa, and N. Talmon. 2023. Participatory Budgeting: Data, Tools and Analysis. In *Proceedings of IJCAI-2023*. ijcai.org, 2667–2674. https://doi.org/10.24963/ijcai.2023/297

[5] P. Faliszewski, Ł. Janeczko, G. Lisowski, K. Pekárková, and I. Schlotter. 2026. Identifying Imperfect Clones in Elections. In *Proceedings of AAAI-2026*. AAAI Press, Washington, DC, USA.

[6] Ł. Janeczko, J. Lang, G. Lisowski, and S. Szufa. 2024. Discovering Consistent Subelections. In *Proceedings of AAMAS-2024*. IFAAMAS, Richland, SC, USA, 935–943. https://dl.acm.org/doi/10.5555/3635637.3662948

[7] A. D. Procaccia, B. Schiffer, and S. Zhang. 2025. Clone-Robust AI Alignment. In *Proceedings of ICML-2025 (Proceedings of Machine Learning Research, Vol. 267)*. PMLR, 49903–49926. https://proceedings.mlr.press/v267/procaccia25a.html

[8] T. N. Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4 (1987), 185–206. https://doi.org/10.1007/BF00433944