



HAL
open science

Epistemic Approval Voting: Applications to Crowdsourcing Data Labeling

Tahar Allouche

► **To cite this version:**

Tahar Allouche. Epistemic Approval Voting: Applications to Crowdsourcing Data Labeling. Operations Research [math.OA]. Université Paris sciences et lettres, 2022. English. NNT : 2022UPSLD056 . tel-04164259

HAL Id: tel-04164259

<https://theses.hal.science/tel-04164259>

Submitted on 18 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à Université Paris Dauphine

Epistemic Approval Voting: Applications to Crowdsourcing Data Labeling

Soutenu par

Tahar Allouche

Le 14/12/2022

École doctorale n°543

Ecole doctorale SDOSE

Spécialité

Informatique

Composition du jury :

Dorothea Baumeister Professeure junior, Université de Düsseldorf	<i>Examinatrice</i>
Ioannis Caragiannis Professeur, Université d'Aarhus	<i>Rapporteur</i>
Yann Chevaleyre Professeur, Université Paris Dauphine	<i>Président</i>
Jérôme Lang Directeur de recherche CNRS, Université Paris Dauphine	<i>Directeur de thèse</i>
Emilie Morvant Maître de Conférences, Université Jean Monnet	<i>Examinatrice</i>
Matías Núñez Chargé de recherche CNRS, Ecole polytechnique	<i>Rapporteur</i>
Marcus Pivato Professeur, Université Cergy Paris	<i>Examineur</i>
Florian Yger Maître de Conférences, Université Paris Dauphine	<i>Co-encadrant</i>

Acknowledgment

First and foremost, I must say I am utterly grateful to both my supervisors, Jérôme and Florian, for being understanding and supportive when I had to make a tough career and life decision. Thank you for that! Your friendship also made my thesis journey a pleasant experience, even though it started when we were going through a rough period in the midst of the pandemic. You gave me all the freedom I needed, while always being there whenever I needed guidance. I knew straight from the beginning that I'm lucky to have you as supervisors, for your academic excellence obviously, and most importantly, for your human qualities.

Also, I want to thank all my family for believing in me and supporting me all along the way, especially, my mother Souad, my father Mohsen, my sister Mariem, my brothers Hamma and Hamza, my grandmother Fathia, my aunts Nahla and Melika, and my uncle Badii. I shall certainly not forget my grandfather Hamadi and my grandmother Baya, may they rest in peace, whom I hope I'm making proud.

Abstract

Epistemic social choice deals with the problem of unveiling a hidden *ground truth* state from a set of possible states, given the votes cast by some agents, that can be seen as *noisy* reports on the ground truth. The correlation between these reports and the ground truth is modelled by a noise model, sometimes tuned by a parameter reflecting the competence (expertise, reliability) of the voters. While some classical anonymous rules have been shown to output the maximum likelihood estimator of the ground truth under some assumptions, the aggregation rule may sometimes assign different weights to the voters according to their expertise.

We focus on a particular kind of input, consisting of *approval votes*, where each voter selects a subset of alternatives, and we propose and test novel methods to optimally weigh the voters with the aim of tracking a simple (single-winner) or composite (multi-winner) truth via *maximum likelihood estimation*, with a particular interest in applications to crowdsourcing.

First, we focus on single-winner epistemic approval voting, and argue for the *size matters* principle. It states that more reliable voters tend to select fewer alternatives, hence voting rules should assign more weight to smaller ballots. We characterize a large family of Mallows-like approval noise distributions based on this principle. Moreover, for the specific case of a Condorcet-like approval noise, we design a simple method to estimate voters' reliabilities by simply observing the size of their ballots. We test these aggregation methods on different real crowdsourced annotations datasets, and show that they outperform standard approval voting in single-question wisdom of the crowd situations.

Second, we deal with multi-winner epistemic approval voting, where the truth consists of a set of objective winners (sentiments in a text, objects in an image, best three papers at a conference..). In such contexts, there may be several interpretations of the size of the output. One of the contributions of the thesis is to distinguish and discuss two different interpretations, leading to different solutions.

In the first interpretation, we suppose that we have a prior knowledge on the ground truth, consist-

ing of a lower and an upper bound on its cardinality. After positing a parameterized noise model that incorporates this prior information, we propose an iterative procedure that intertwines the estimation of the parameters and the ground truth. We test the algorithm on an annotation dataset that we collected, and show that incorporating the prior knowledge into the estimation process significantly enhances its accuracy.

In the second interpretation, we focus on cases where some exogenous size constraints bear upon the number of winners in the output, even though the size of the ground truth may lie outside this interval (*e.g.*, accepting up to thirty student applying to a master program, even though there might be more than thirty good applicants). Whilst the usual aim of epistemic voting rules is to uncover the “whole” ground truth, here the decision consists in outputting the objectively best subset of alternatives. To this aim we introduce two solution concepts, which we test on synthetic data, and we prove that they outperform a baseline.

keywords: Epistemic approval voting, Maximum likelihood estimation, Crowdsourcing, Data Labeling.

Résumé

Nous nous concentrons sur un type particulier d'entrée, consistant en votes par approbation, où chaque votant sélectionne un sous-ensemble d'alternatives, et nous proposons et testons de nouvelles méthodes pour pondérer de manière optimale les votants dans le but de retrouver une vérité simple (vainqueur unique) ou composite (vainqueurs multiples) via l'estimation du maximum de vraisemblance, avec un intérêt particulier pour les applications au crowdsourcing.

Tout d'abord, nous nous concentrons sur le vote par approbation épistémique à un seul gagnant et défendons le principe "size matters". Il stipule que les électeurs les plus fiables ont tendance à sélectionner moins d'alternatives, et que les règles de vote devraient donc accorder plus de poids aux bulletins de vote plus petits. Nous caractérisons une grande famille de distributions de bruit d'approbation de type Mallows basée sur ce principe. De plus, pour le cas spécifique d'un bruit de type Condorcet, nous concevons une méthode simple pour estimer la fiabilité des électeurs en observant simplement la taille de leurs bulletins. Nous testons ces méthodes d'agrégation sur différents ensembles de données d'annotations réelles provenant du crowdsourcing, et nous montrons qu'elles sont plus performantes que le vote d'approbation standard dans les situations de sagesse de la foule à question unique.

Deuxièmement, nous traitons le vote par approbation épistémique à plusieurs vainqueurs, où la vérité consiste en un ensemble de gagnants objectifs (sentiments dans un texte, objets dans une image, les trois meilleurs articles d'une conférence...). Dans de tels contextes, il peut y avoir plusieurs interprétations de la taille de la sortie. Une des contributions de cette thèse est de distinguer et de discuter deux interprétations différentes, menant à des solutions différentes.

Dans la première interprétation, nous supposons que nous avons une connaissance préalable sur la vérité objective, consistant en une borne inférieure et supérieure sur sa cardinalité. Après avoir posé un modèle de bruit paramétré qui incorpore cette information a priori, nous proposons une procédure itérative qui entrelace l'estimation des paramètres et de la vérité. Nous testons l'algorithme sur un ensemble de données d'annotation que nous avons collecté, et nous mon-

trons que l'incorporation de la connaissance a priori dans le processus d'estimation améliore considérablement sa précision.

Dans la deuxième interprétation, nous nous concentrons sur les cas où certaines contraintes de taille exogènes pèsent sur le nombre de gagnants dans la sortie, même si la taille de la vérité peut se situer en dehors de cet intervalle (par exemple, accepter jusqu'à trente étudiants postulant à un programme de master, même s'il peut y avoir plus de trente bons candidats). Alors que l'objectif habituel des règles de vote épistémique est de découvrir la "totalité" de la vérité de base, la décision consiste ici à produire le meilleur sous-ensemble objectif d'alternatives. Nous introduisons deux concepts de solution, que nous testons sur des données synthétiques, et nous prouvons qu'ils sont plus performants qu'une base de référence.

mots-clefs: Vote d'approbation épistémique, Estimation du maximum de vraisemblance, Crowdsourcing, Labellisation des données.

Contents

Acknowledgement	1
Abstract	3
Résumé	5
List of Figures	13
List of Tables	15
Introduction	17
I Background	21
1 Voting	23
1.1 Introduction	23
1.2 Voting Rules	24
1.2.1 Ranking-based Aggregation Rules	25
1.2.1.1 The Ranking Framework	25
1.2.1.2 Single-winner Rules	26
1.2.1.3 Social Welfare Functions	28
1.2.1.4 Multi-winner Rules	30
1.2.2 Approval-based Aggregation Rules	33
1.2.2.1 The Approval Framework	33
1.2.2.2 Single-winner Rules	34
1.2.2.3 Multi-winner Rules	36

2	Epistemic Voting	39
2.1	Condorcet Jury Theorem	39
2.1.1	Heterogeneous Reliabilities	40
2.1.2	Correlated Votes	41
2.1.3	More than Two Alternatives	41
2.1.4	Approval Profiles	42
2.1.5	Abstract Voting Rules: General Results	42
2.1.6	Sampling Complexity	43
2.2	Optimal Estimation Rules	45
2.2.1	Brief Refresher on Maximum Likelihood Estimation	45
2.2.2	Voting on Binary Issues	46
2.2.3	Input Profiles are Full Rankings	49
2.2.4	Input Profiles are Partial Orders	49
2.2.5	Input Profiles are Pairwise Comparisons	50
2.2.6	Input Profiles are Truncated k-Approval Votes	52
2.2.7	Input Profiles are Approval Votes	52
2.2.8	Input Profiles are Preferences on Directed Hypercubes	54
2.2.9	Summary	54
2.3	Other Directions	54
2.3.1	The Converse Problem: Which Rules Correspond to MLEs?	54
2.3.2	Axiomatic Analysis for Collective Annotations	55
2.3.3	Incentives	56
2.3.4	The Surprisingly Popular Method	57
2.3.5	Judgment Aggregation	58
II	Contributions	59
3	Single-winner Truth-Tracking	63
3.1	Introduction	63
3.2	Framework	65
3.3	Anonymous Noise and Size-decreasing Approval Rules	67
3.4	Non-anonymous Separable Noise	71
3.4.1	The General Case	71
3.4.2	The Hamming Distance Case - Condorcet Noise Model	76
3.5	Experiments	77
3.5.1	Ballot Size and Reliability	78
3.5.2	Aggregation	81
3.5.3	Results	82
3.6	Conclusion	82
4	Multi-winner Epistemic Voting: Committee-size as Prior Knowledge	85
4.1	Introduction	85

4.2	The Model	87
4.3	Estimating the Ground Truth	88
4.3.1	Estimating the Ground Truth Given the Votes and the Parameters	89
4.3.2	Estimating the Parameters Given the Ground Truth	93
4.3.2.1	Estimating the prior parameters over alternatives	93
4.3.2.2	Estimating the voter parameters	95
4.3.3	Alternating Maximum Likelihood Estimation	96
4.4	Experiments	100
4.4.1	Experiment Design and Data Collection	100
4.4.2	Descriptive Analysis of Collected Answers	100
4.4.3	Anna Karenina’s Initialization	100
4.4.4	Time Complexity of AMLE	107
4.4.5	Results	107
4.4.5.1	Hamming and 0-1 Subset Accuracies	107
4.4.5.2	Harmonic Loss	110
4.4.5.3	Comparing the accuracies	110
4.5	The Case of Fixed Committee-size: Extension of Theorem 3.4 to Multi-winner Settings	111
4.5.1	Noise Model and Estimations	111
4.5.2	Experiments	112
4.6	Conclusion	113
5	Multi-winner Epistemic Voting: Committee-size as Constraint	115
5.1	Framework and Objective	115
5.1.1	Votes, Ground Truth and External Constraints	115
5.1.2	Utility	116
5.1.3	The Noise Model	116
5.1.4	The Solution Concepts	117
5.2	MLAUM: Most Likely Admissible Utility Maximizer	118
5.3	AEUM: Admissible Expected Utility Maximizer	125
5.4	Simulations	128
5.4.1	Heuristics	129
5.4.2	Generated Data and Parameters	129
5.4.3	Results and Observations	129
5.5	Conclusion	132
	Conclusion	133
	Résumé Long en Français	137
5.6	État de l’art sur le vote épistémique	138
5.7	Vote par approbation épistémique à un seul gagnant : la taille importe	139
5.7.1	Le modèle de bruit	140
5.7.2	Cas d’un bruit anonyme	141

5.7.3	Cas d'un bruit décomposable non-anonyme	142
5.7.4	Expériences	143
5.8	Vote par approbation épistémique à vainqueurs multiples: deux interprétations . .	145
5.9	Vote par approbation épistémique à vainqueurs multiples: les contraintes de taille comme information a priori	146
5.9.1	Le modèle de bruit	147
5.9.2	Estimer la vérité objective étant donné les paramètres	148
5.9.3	Estimer les paramètres étant donné la réalité objective	149
5.9.4	Estimation de maximum de vraisemblance alternée	150
5.9.5	Expériences	150
5.10	Vote par approbation épistémique à vainqueurs multiples: les contraintes de taille comme facteur exogène	152
5.10.1	Le modèle de bruit	154
5.10.2	Deux objectives	155
5.10.2.1	Le maximiseur le plus probable de l'utilité MPPU	155
5.10.2.2	Le maximiseur de l'utilité espérée MUE	156
5.11	Conclusion	156

III Appendix 157

6	Social Rankings Manipulability	159
6.1	Introduction	160
6.2	Preliminaries	162
6.3	Manipulability	163
6.3.1	Copeland-like and Kramer-Simpson-like rules	164
6.3.1.1	Copeland-like method	164
6.3.1.2	Kramer-Simpson-like method	165
6.3.2	Ordinal Banzhaf	166
6.3.3	Lexicographic excellence solution	167
6.4	Computational complexity of manipulation	168
6.5	Simulations	172
6.5.1	Integer linear Program	172
6.5.2	Ordinal Banzhaf solution	173
6.5.3	Copeland-like solution	174
6.5.4	Kramer-Simpsonsolution: linearization of $x = \max(a, b)$	174
6.5.4.1	Linearize $x = \max(a, b)$	174
6.5.5	Results	174
6.5.5.1	The proportion of manipulable cases:	174
6.5.5.2	Number of possible manipulators:	175
6.5.5.3	Manipulating to be the best ranked:	175
6.5.5.4	Cross-simulation	175
6.6	Conclusion	176

List of Figures

1.1	Image annotation task example [Shah et al., 2015]	24
1.2	Approval profile in a Doodle poll	34
3.1	Instance from the Animals task [Shah et al., 2015]	78
3.2	Instance from the Textures task [Shah et al., 2015]	79
3.3	Instance from the Languages task [Shah et al., 2015]	79
3.4	Size-normalized accuracy for three datasets	80
3.5	Accuracy of different rules on real-life datasets	83
4.1	The images to be labeled before editing	101
4.2	The images to be labeled after editing	102
4.3	Example of Annotation Task	103
4.4	Top quiz scores	103
4.5	Histogram of the ballots' sizes	104
4.6	Accuracies of different initializations	106
4.7	Time complexity of AMLE	108
4.8	Accuracies of different aggregation methods	109
4.9	Normalized Harmonic accuracy	111
4.10	0-1 subset and Hamming accuracy on Football quiz dataset	114
5.1	Hamming loss and MAE for different number of voters	130
5.2	Hamming loss and MAE for different constraints	131
5.3	Précision des différentes règles d'agrégation	144
5.4	Exemple de la tâche d'annotation d'images	152
5.5	Précisions de différentes méthodes d'agrégation	153
6.1	Proportion of manipulable cases	175
6.2	Proportion of possible manipulators.	176

6.3 Probability to be ranked first with manipulation. 177

List of Tables

1.1	All linear orders over 3 alternatives	25
1.2	Election with 4 voters and 3 alternatives	26
1.3	Approval-based election with 4 voters and 3 alternatives	33
2.1	Optimal rules according to the nature of their input-output	55
4.1	Approval Ballots of 3 Voters on 4 Instances	105
4.2	Hamming and 0/1 accuracy for entire dataset	112
5.1	Règles optimales en fonction de la nature de leurs entrées-sorties	139
6.1	Percentage of manipulable power relations.	177

Introduction

Computational social choice deals with the problem of collective decision making. It is an area of research at the intersection of computer science, political science, economics and mathematics, that focuses on aggregating individual inputs into a single collective outcome. Many subareas of social choice concern the design of methods for handling these preferences in different settings ranging from *matchings* and *stable coalition formation* to *fair division* and *voting*, the latter being the focal point of the thesis.

Social choice theorists distinguish between two interpretations of voting: *classical* and *epistemic*. In classical voting, we aggregate *preferences*, while in epistemic voting we aggregate *beliefs*. The typical instance of the classical view is a political election, where a number of candidates (persons or lists) run for candidacy, and an electorate gets to decide on to whom it will concede all or some of the political power at stake. Forming this decision, arising from an often heterogeneous population, and usually conflicting opinions and interests, while democratically guaranteeing equal chances and treatment to all the political stake-holders alike (anonymity of voters and neutrality with regards to candidates), is where the design of the voting procedure crucially intervenes. Besides such high-stake and low-frequency contexts, voting is often used to decide on matters at a smaller scale such as company boards or hiring committees. It is even becoming more and more incorporated into our routinal tasks including the scheduling of meetings via dedicated online platforms ¹.

Nonetheless, besides aggregating subjective preferences, partially determined by ideological stances and perceived interests, voting can also mediate different attempts towards unveiling an *objective truth*. These truth-tracking attempts, yielding beliefs about the truth, are called *epistemic votes*. They represent the second voting interpretation discussed since early in the 18th century [Condorcet, 1785] and followed by a large body of research, where this thesis is located. A classical example is a jury deliberation over a convict's trial: normally a judge does not prefer

¹Such as <https://doodle.com/> or <https://whale.imag.fr/>.

that the prosecuted be called guilty (or innocent) but she believes it should be the case given the evidence and according to the applicable law. Condorcet initially established the wisdom of the majority principle, which became known as *Condorcet jury theorem*, stipulating that the majority's opinion about a two-outcome issue tends to be correct when the population is large enough. While the jury is mostly guided in her truth-tracking by her moral, social and professional duties, other contexts directly align the voters' economic incentives to their truthfulness: it is the case in Blockchain block validations and in crowdsourcing data labeling platforms². In Blockchain block validation, a validator gets to vote on which final block is correct when the chain is temporarily forked into two or more ramifications, and would risk seeing her energy cost (or her staked money, depending on the protocol) lost for nothing if she approves the wrong block. On crowdsourcing platforms, a worker risks her reputation (and her payments) to be harmed if she deliberately picks the wrong labels when annotating some instances.

In the epistemic setting, a voter can express her beliefs in different ways. For instance, she can give her most plausible alternative, or else she can order the alternatives from the most to the least probably correct. This dissertation deals with an intermediate format called *approval voting*, where voters can select any number of alternatives that they seem possible. Although approval ballots are more informative than a single alternative vote and less cognitively burdening than ranking all the alternatives, approval voting has attracted less attention in the epistemic social choice literature.

Another advantage of approval voting is its compatibility with both a *simple* and a *composite* truth. As classical voting distinguishes single-winner voting (presidential elections, scheduling a meeting ..) and multi-winner voting (parliamentary elections, recruitment shortlisting..), we will similarly distinguish whether the truth consists of a single alternative (the city in which a photo is taken, the language in which a speech is spoken ..) or of multiple alternatives (the notes in transcribed chord, the objects in an image, the sentiments in a text, the objectively best applicants for a position..). In both cases, voters are able to express their beliefs by approving the alternatives they think might be correct and disapprove the others. The interpretation of the size of the output in the multi-winner context is not as straightforward as in classical voting: it might express a prior knowledge on the number of objectively true alternatives, or an exogenous constraint bearing upon the output of the rule (e.g. selecting applicants for a limited number of scholarships) regardless of the true size of the ground truth.

One of the questions we try to address in this thesis is whether aggregating epistemic approval votes should respect *anonymity* (treating all voters equally) and *neutrality* (treating all alternatives equally). There are obvious situations where treating all the voters equally is sub-optimal in terms of the efficient uncovering of the ground truth: it is justifiable to put more emphasis on the beliefs of expert voters than on those who are less informed on the matter. Moreover, an alternative might be privileged over the others if some prior knowledge or exogenous signs, other than the votes, suggest that it is more likely to be objectively true.

The main two challenges are: how to efficiently assess the voters' expertise without eliciting

²Such as <https://www.clickworker.com/clickworker/> or <https://www.mturk.com/>.

any further information about them but their votes, and how to incorporate any prior knowledge or external signs into the truth uncovering to correctly privilege some alternatives to others. In the upcoming chapters, we propose different solutions, with an emphasis on their truth-tracking performances on real-world annotation data, namely:

- We assess a voter’s reliability by observing the number of alternatives that she selects. We assume that the larger it is, the less likely is the voter well-informed about the question. This principle goes beyond the wisdom of the majority, and might, for instance, avoid succumbing to the judgement of an erroneous majority facing a tricky question.
- In case of a composite truth, we propose a more sophisticated iterative procedure which incorporates both weighing the voters and taking into account certain kinds of prior information about the ground truth. Typically, we focus on a specific prior knowledge, consisting of lower and upper bounds limiting the number of correct alternatives. These bounds might for instance arise from domain knowledge.
- Sometimes the decision we want to make based on the estimation of the ground truth is constrained by some exogenous factors. The typical situation is when we desire to keep the objectively eligible alternatives and discard the non-eligible ones, but the number of available places is limited (e.g. a master program’s committee selecting applicants based on the likelihood of their success). We show that the optimal aggregation rules in such situations depend on the utility (respectively, cost) of correctly including (respectively, discarding) a correct alternative and discarding (respectively, including) a wrong one.

Manuscript organization

The thesis is organised as follows:

- Chapter 1 introduces the main concepts of voting theory. The basic notions relating to elections with different natures of input (rankings or approvals) and different output formats (single alternative, multiple alternatives or rankings) are presented. For each case, we list the most important families of voting rules and illustrate them with examples.
- Chapter 2 reviews the epistemic social choice literature. We begin by presenting the seminal Condorcet Jury Theorem and its extensions and ramifications, along with works dealing with the sampling complexity problem. Then we proceed to reviewing the subsequent works which studies optimal rules for aggregating voters’ beliefs in order to uncover a ground truth of different natures.
- Chapter 3 deals with single-winner epistemic approval voting. Its driving idea is that the most reliable voters tend to select fewer alternatives. We consider a family of approval noise models and prove that under certain conditions the alternative maximizing the likelihood of being the ground truth is the winner by a weighted approval rule, where the weight of a ballot decreases with its cardinality. These rules are shown to outperform the standard approval rule on crowdsourced image annotation datasets.

- Chapter 4 deals with multi-winner epistemic approval voting. It considers contexts where the truth consists of a set of objective winners, knowing a lower and upper bound on its cardinality. Aggregation methods which incorporate this prior knowledge into the estimation procedure are then defined and tested on crowdsourced data.
- Chapter 5 also tackles the problem of multi-winner epistemic approval voting. Nonetheless, unlike Chapter 4, the inequality size constraints are not interpreted as a prior knowledge, but rather as exogenous limitations imposed on the output of the aggregation rule. Two solution concepts are introduced, and optimal estimation methods are defined for each of them.
- Appendix 6 (not formally part of the thesis work) addresses the *social ranking* problem, which was the topic I covered in my master's thesis [Allouche et al., 2020]. It considers situations where we have an ordinal record of the performances of different subsets of agents (*e.g.*, workers in a crowdsourcing platforms) and the goal is to deduce an order over individual performances (*e.g.*, to reward the agents accordingly).

Part I

Background

1.1 Introduction

Voting is a way to deal with the problem of collective decision making. It handles individual preferences or beliefs and aggregates them into a single collective outcome. It can be faced in a range of different situations, namely:

- *Political elections*: This was the initial motive for the first advances in Social Choice theory starting with the earlier works of Marquis de Condorcet [Condorcet, 1785] and Jean-Charles de Borda [de Borda, 1781] and the seminal works of Kenneth Arrow [Arrow, 1950] and Gibbard and Satterthwaite [Gibbard, 1973b, Satterthwaite, 1975a]. In a political election, voters are given a number of candidates over whom they are asked to submit their *preferences*. There are several possible input formats for electoral systems: each voter can for instance order the candidates from the most to the least preferred, give her approval for a subset of candidates or cast an evaluation (or grade) for each one of them (choosing between *good*, *average* and *bad* for example) as suggests, for instance, [Balinski and Laraki, 2007]. The nature of the outcome can also vary: in presidential elections the outcome would be a single candidate, whereas the outcome of a parliamentary election is a subset of more than one candidate.
- *Meeting scheduling platforms*: Some platforms like *Doodle*¹ or *Whale*² offer users the possibility of organizing events on the date that fits most of the concerned people. After creating a poll that offers a number of possible time slots, each user selects her availabilities or the slots that fit her most [Zou et al., 2015]. Usually the slot that suits the greatest number of participants is selected.
- *Crowdsourcing*: Supervised learning algorithms need labeled data to be trained on. Sometimes, labelling the dataset's instances is costly and companies might decide to outsource

¹<https://doodle.com/>

²<https://whale.imag.fr/>



Figure 1.1: Image annotation task example [Shah et al., 2015]

it. Some platforms like *Amazon Mechanical Turk*³ pay *workers* to perform some *Human Intelligence Tasks* (HITs) like image annotation, speech transcription and translation. In the image annotation task show in Figure 1.1, each voter (worker) has to select the alternatives (in this example, animals) that she believes might be the correct one. After collecting a certain number of answers, it remains to find an efficient way of aggregating them to recover the correct label which will be discussed more in details in later chapters.

- *Ensemble learning*: In order to improve the predictive power of supervised learning algorithms, ensemble learning techniques like *Bagging* [Breiman, 1996] are used. Roughly speaking, they consist in training numerous *base learners*, and aggregating their individual outputs at the moment of prediction by voting. Notice here that the agents (voters) are algorithms and not people.

Notice that in the examples above, voting scenarios may differ with regards to many aspects like:

- *The nature of the input*: it can be a single alternative, a complete or partial ranking of alternatives, a subset of alternatives, a numerical or qualitative evaluation of each alternative.
- *The nature of the output*: it can be a single alternative, a set of alternatives, a ranking over alternatives, a probability distribution over alternatives.
- *The interpretation of a vote*: in elections a vote is an expression of a *preference*. Whilst in the crowdsourcing and ensemble learning frameworks, a vote is the expression of a *belief* about an *objective truth*.

In Section 1.2 we will formalize the different natures of inputs in voting contexts and we will define voting rules for different output types.

1.2 Voting Rules

We usually consider a set N of n voters and a set \mathcal{X} of m alternatives, or candidates. To express their preferences (or beliefs) about the alternatives, voters may cast a vote in different forms.

³<https://www.mturk.com/>

In this section we will formally define the two voting formats that we'll use in the upcoming chapters: *rankings* and *approval ballots*, and list some of the common voting rules for each of them. The aggregation rules will be partitioned based on the nature on the input (ranking-based or approval-based) and the nature of the output (ranking, single-winner or multi-winner).

1.2.1 Ranking-based Aggregation Rules

1.2.1.1 The Ranking Framework

Voters orders the candidates from their most to the least preferred one. Rankings are also mathematical objects called *orders*.

Definition 1.1. A linear order \succ over the elements of \mathcal{X} is a binary relation over $\mathcal{X} \times \mathcal{X}$, which is:

- **transitive:** $\forall a, b, c \in \mathcal{X}, a \succ b \text{ and } b \succ c \implies a \succ c$.
- **complete:** $\forall a \neq b, a \succ b \text{ or } b \succ a$.
- **asymmetric:** $\forall a, b \in \mathcal{X}, a \succ b \implies \text{not } b \succ a$.

In Table 1.1 we explicit all the 6 possible linear orders over the set $\mathcal{X} = \{a, b, c\}$ of 3 alternatives.

Linear orders
$a \succ b \succ c$
$a \succ c \succ b$
$b \succ a \succ c$
$b \succ c \succ a$
$c \succ a \succ b$
$c \succ b \succ a$

Table 1.1: All linear orders over 3 alternatives

It is worth mentioning that in some settings, voters can cast incomplete preferences, which can be represented by *partial orders*.

Example 1.1. In a crowdsourcing scenario, imagine a voter is given three capital cities: Cairo, Jakarta and Tunis, and is asked to order them from most to least populated. Then the crowdsourcing platform can give the worker the possibility of expressing uncertainty by casting a partial order. Let's imagine that the worker knows that both Cairo and Jakarta are way more populated than Tunis, but is still unable to which one is bigger. In this case, her truthful vote would be the following partial order, containing only two pairwise comparisons, one between Tunis and Jakarta, and another between Tunis and Cairo:

$$R_i = \{\text{Cairo} \succ_i \text{Tunis}, \text{Jakarta} \succ_i \text{Tunis}\}$$

A subfamily of partial orders is *top- k truncated queries*, where each voter only provides the top- k ranked alternatives from her complete order for some $k \leq m$. For instance, consider a voter whose complete preference is $a \succ b \succ c \succ d \succ e$. Her top-3 truncated order would be $a \succ b \succ c$.

In a voting situation, every voter $i \in N$ casts her ranking $\succ_i \in L(\mathcal{X})$, where $L(\mathcal{X})$ is the set of linear orders, or $\succeq_i \in W(\mathcal{X})$, where $W(\mathcal{X})$ is the set of partial orders. The collection of the n rankings is called the profile of votes or preferences and is denoted $P = (\succ_1, \dots, \succ_n)$.

A ranking election is then a triplet (N, \mathcal{X}, P) , and can be presented as in Table 1.2.

$$\begin{aligned} v_1 : & a \succ_1 b \succ_1 c \\ v_2 : & a \succ_2 c \succ_2 b \\ v_3 : & a \succ_3 c \succ_3 b \\ v_4 : & b \succ_4 a \succ_4 c \end{aligned}$$

Table 1.2: Election with 4 voters and 3 alternatives

1.2.1.2 Single-winner Rules

Definition 1.2 (Social Choice Correspondence). *A social choice correspondence is a function f which, for each election E , assigns a non-empty set of co-winners $S \subseteq \mathcal{X}$:*

$$\begin{aligned} f : \quad \mathcal{P} = L(\mathcal{X})^n & \longrightarrow 2^{\mathcal{X}} \setminus \emptyset \\ P = (\succ_1, \dots, \succ_n) & \longmapsto S \subseteq \mathcal{X}, S \neq \emptyset \end{aligned}$$

Definition 1.3 (Social Choice Rule). *A social choice rule is a function f which, for each election E , assigns a single (winner) alternative $a \in \mathcal{X}$:*

$$\begin{aligned} f : \quad \mathcal{P} = L(\mathcal{X})^n & \longrightarrow \mathcal{P}(\mathcal{X}) \\ P = (\succ_1, \dots, \succ_n) & \longmapsto a \in \mathcal{X} \end{aligned}$$

Social choice correspondence are also called *irresolute* rules, whereas social choice functions are referred to as *resolute* rules. The latter are usually constructed from composing an irresolute rule with a tie-breaking mechanism either by using an exogenous predefined priority relation on candidates or a predefined priority relation on voters. It is worth mentioning that such tie-breaking methods violate either *anonymity* (all voters are a priori interchangeable) or *neutrality* (all candidates are a priori interchangeable).

We describe now two families of rules: *Positional scoring rules* and *Pairwise comparison rules*.

Definition 1.4 (Positional Scoring Rules). *Consider an m -length vector of real numbers $\vec{s} = (s_1, \dots, s_m)$, such that $s_1 \geq \dots \geq s_m$ and $s_1 > s_m$. The positional scoring function $f_{\vec{s}}$ maps each profile P to the (co-)winner(s):*

$$f_{\vec{s}} : P \mapsto \arg \max_{a \in \mathcal{X}} S(a)$$

where the score S of each alternative is:

$$S(a) = \sum_{i \in N} s_{r_i(a)}$$

with $r_i(a)$ being the rank of a in \succ_i :

$$r_i(a) = m - |x \in \mathcal{X}, a \succ_i x|$$

Some of the most common scoring functions are:

- *Plurality*: The scoring vector of the Plurality rule is $\vec{s}_{\text{Plurality}} = (1, 0, \dots, 0)$. The Plurality function simply selects the alternative ranked in first position by the largest number of voters.
- *Borda*: The scoring vector of the Borda rule is $\vec{s}_{\text{Borda}} = (m-1, m-2, \dots, 0)$. The Borda function is known to output the alternative with highest average position in the individual rankings.
- *Harmonic*: The scoring vector of the Harmonic rule is $\vec{s}_{\text{Harmonic}} = (1, 1/2, \dots, 1/m)$.

Example 1.2. Consider the following voting profile:

$$\begin{aligned} v_1 &: a \succ_1 b \succ_1 c \\ v_2 &: a \succ_2 b \succ_2 c \\ v_3 &: c \succ_3 b \succ_3 a \\ v_4 &: b \succ_4 c \succ_4 a \end{aligned}$$

Considering the positional scoring voting rules associated to the Plurality, Borda and Harmonic scores, the alternatives would have the following scores:

Candidate	a	b	c
Plurality	2	1	1
Borda	4	5	3
Harmonic	$8/3$	$5/2$	$13/6$

The three rules yield the following outcomes:

$$\begin{aligned} f_{\text{Plurality}}(P) &= a \\ f_{\text{Borda}}(P) &= b \\ f_{\text{Harmonic}}(P) &= a \end{aligned}$$

In order to introduce the next two rules we need to define the notion of *Majority graph*.

Definition 1.5 (Majority Graph). *For an election $E = (N, \mathcal{X}, P)$, and for any two alternatives $a, b \in \mathcal{X}$, let $N_P(a, b) = \#\{i \in N, a \succ_i b\}$ be the number of voters preferring a to b . The majority graph associated to P is a directed graph $M(P)$ whose vertices are \mathcal{X} and with a directed edge between a and b if and only if $N_P(a, b) > N_P(b, a)$.*

Definition 1.6 (Copeland Rule [Copeland, 1951]). *For an alternative $a \in \mathcal{X}$, let $S(a)$ be the outdegree of a in $M(P)$ plus half the number of alternatives b with no edge between a and b :*

$$S(a) = \#\{b \in \mathcal{X}, N_P(a, b) > N_P(b, a)\} + \frac{1}{2}\#\{b \neq a, N_P(a, b) = N_P(b, a)\}$$

The Copeland rule outputs the alternative with the highest score $S(a)$.

Definition 1.7 (Kramer-Simpson Rule [Kramer, 1977, Simpson, 1969]). *For an alternative $a \in \mathcal{X}$, let $S(a)$ be the minimum pairwise majority margin of a :*

$$S(a) = \min_{b \neq a} N_P(a, b)$$

The Kramer-Simpson (also known as Maximin) rule outputs the alternative with the highest score.

Example 1.3. *Recall the election presented in Example 1.2. The pairwise comparisons are as follows:*

$$\begin{cases} N_P(a, b) = 2 \\ N_P(a, c) = 2 \\ N_P(b, c) = 3 \end{cases}$$

The alternatives have thus the following Copeland scores:

$$S(a) = \frac{1}{2} + \frac{1}{2} = 1, S(b) = \frac{1}{2} + 1 = \frac{3}{2}, S(c) = \frac{1}{2}$$

The winner is b .

For the Kramer-Simpson rule, the scores are:

$$S(a) = \min\{2, 2\} = 2, S(b) = \min\{2, 3\} = 2, S(c) = \min\{2, 1\} = 1$$

Two winners a and b are irresolute.

1.2.1.3 Social Welfare Functions

Put simply, a *Social Welfare Function* (SWF) takes as input a preference profile P and aggregates it into a single linear order $\succ^* \in L(\mathcal{X})$.

Definition 1.8 (Social Welfare Function). *Consider a set N of n voters and a set \mathcal{X} of m alternatives. A social welfare function (SWF) f is function that maps every preference profile to a linear order:*

$$\begin{aligned} f : \quad \mathcal{P} = L(\mathcal{X})^n &\longrightarrow L(\mathcal{X}) \\ P = (\succ_1, \dots, \succ_n) &\longmapsto \succ^* \end{aligned}$$

We will list some usual SWFs.

Definition 1.9 (Modal SWF [Caragiannis et al., 2014]). *The modal social welfare function is the function that maps each profile to its most frequently cast order:*

$$f_{mode}(\succ_1, \dots, \succ_n) = \arg \max_{\succ \in L(\mathcal{X})} \#\{i \in N, \succ_i = \succ\}$$

Example 1.4. *Recall the preference profile P in Table 1.2. The ranking $a \succ b \succ c$ was cast by one voter, $a \succ c \succ b$ was cast by two voters, $b \succ a \succ c$ was cast by one voter, and all the remaining orders were cast by none of the voters. The most frequent ranking in the profile is thus $a \succ c \succ b$:*

$$f_{mode}(P) = a \succ c \succ b$$

Another family of SWFs is based on the positional scoring functions presented in Definition 1.4, but instead of outputting the alternative with the highest score, the associated SWF would rank the candidates according to their scores.

Definition 1.10 (Positional Scoring SWF). *Consider an m -length vector of real numbers $\vec{s} = (s_1, \dots, s_m)$, such that $s_1 \geq \dots \geq s_m$ and $s_1 > s_m$. The positional scoring function $f_{\vec{s}}$ maps each profile P to a ranking \succ such that:*

$$a \succeq b \iff S(a) \geq S(b), \text{ for all } a, b \in \mathcal{X}$$

Example 1.5. *Recall the example election in Example 1.2. Suppose that we fix an exogenous tie-breaking priority as follows:*

$$a \triangleright b \triangleright c$$

The outcome of the three rules are:

$$\begin{aligned} f_{Plurality}(P) &= a \succ_{Pl} b \succ_{Pl} c \\ f_{Borda}(P) &= b \succ_B a \succ_B c \\ f_{Harmonic}(P) &= a \succ_H b \succ_H c \end{aligned}$$

Definition 1.11 (Kemeny SWF). *The Kemeny SWF is the function f_K that maps every preference profile P to the linear order that minimizes the average Kendall-Tau distance to the voters' rankings:*

$$f_K(P) = \arg \min_{\succ \in L(\mathcal{X})} \frac{1}{n} \sum_{i \in N} d_{KT}(\succ_i, \succ)$$

where the Kendall-Tau distance is given by:

$$d_{KT}(\succ, \succ') = \frac{1}{2} \sum_{a, b \in \mathcal{X}} \mathbb{1}\{a \succ b\} \times \mathbb{1}\{b \succ' a\}$$

Example 1.6. Recall the profile in example 1.5. The Kemeny SWF yields the following linear order:

$$f_K(P) = a \succ_K b \succ_K c$$

Remark. Notice that by changing the distance Kendall-Tau into another distance on $L(\mathcal{X})$ like the swap distance, or by changing the ℓ_1 norm used in averaging the distance into any other norm like ℓ_∞ or ℓ_2 , we can navigate a whole family of Kemeny-like social welfare functions [Elkind et al., 2016].

1.2.1.4 Multi-winner Rules

There are some situations where a group of alternatives needs to be elected at once. The most classical example is the parliamentary elections, but other use-cases occur in different domains such that:

- Recruiting team shortlisting job applicants for interviews: the recruiter begins by ordering the applicants based on their submitted resumes. Some of them are then selected to proceed to the next round of the recruitment process, based on their excellence without violating some diversity constraints.
- Multi-label base-classifiers aggregating their outputs: given a multi-label classification task, a group of simple classifier algorithms are trained, their output is then aggregated to predict the correct labels.
- Crowdsourcing multi-label annotations tasks: for instance, participants are shown a video clip and are asked to rank the sentiments that each of them feels when watching it. The votes are aggregated to estimate the true sentiments expressed in the video.

First we will define elections and voting rules for rankings, then we will present the main families of multi-winner rules.

Definition 1.12 (Multi-winner Ranking-based Election). An election E is a quadruple $E = (\mathcal{X}, N, P, k)$ where:

- \mathcal{X} is the set of m candidates.
- N is the set of n voters.
- $P = (\succ_1, \dots, \succ_n)$ is the preferences profile consisting of n linear orders over \mathcal{X} .
- $k \leq m$ is the committee-size, representing the size of the desired outcome.

Definition 1.13 (Multi-Winner Voting Rule). *A voting rule is a mapping f that maps each election to a subset of \mathcal{X} with cardinality k representing the winning committee.*

Depending on the context of the election, [Elkind et al., 2017] and [Faliszewski et al., 2017] distinguish three main types of voting rules:

- *Excellence-based*: In this case, the rule is supposed to elect the “best” candidates, regardless of their interactions, according to the judgement of the voters.
- *Diversity-based*: The aim of such rules is to select a committee which is as diverse as possible.
- *Proportionality-based*: This is illustrated by the application of multi-winner voting to parliamentary elections, for which the winning committee should represent the views of the society proportionally. Hence alternatives cannot be considered independently.

Similarly as for single-winner elections, in case ties exist, we might either consider irresolute rules whose output is a set of tied winning committees, or use tie-breaking procedures which are applied after the voting rule.

Examples of some multi-winner voting rules will be given below with their main properties and illustrated with examples.

Definition 1.14 (Single Transferable Vote). *Given an election E the STV rule proceeds in different steps. First, a quota is fixed. Different quotas are studied (see [Lundell and Hill, 2007]), among the most used ones is the Droop quota $q = \lfloor \frac{n}{k+1} \rfloor + 1$. Then, several rounds are executed until k members are elected. Each round begins by checking if some candidate is ranked first by at least q voters. If so, the candidate is elected. Her surplus votes, that is, the votes she got above the quota, are then transferred to subsequent candidates. Different proportional surplus transferring methods exist [Wikipedia contributors, 2022]. The elected candidate is also removed from the preferences of the remaining voters. If there is no such candidate, then we eliminate the one who is ranked first by the smallest number of voters. We iterate the process until k candidates have been elected.*

Example 1.7. *Consider the election E with 4 candidates $\mathcal{X} = \{a, b, c, d\}$ and 12 voters whose preferences are presented in the table below:*

ranking	number of voters
$a \succ b \succ c \succ d$	6
$b \succ a \succ c \succ d$	3
$c \succ a \succ b \succ d$	2
$d \succ b \succ a \succ c$	1

We set the target committee size to $k = 2$. The STV rule proceeds as follows: First we compute the quota given by $q = \lfloor \frac{n}{k+1} \rfloor + 1 = 5$. Then several rounds are executed:

- **Round 1:** *The candidate a is ranked first by 6 voters, she is thus elected. 5 of her voters are removed, and she is also removed from the preference of the remaining one (who is*

henceforth ranking b first).

- **Round 2:** No candidate is ranked first by 5 voters. So we remove d from the preferences since she is ranked first by the smallest number.
- **Round 3:** The candidate b has 5 voters ranking her first so she gets elected.

The winning committee is the pair $\{a, b\}$.

Scoring rules are based on scoring functions that assign a value to each committee from each single voter. The rules then sum up all the individual scores and pick the committee with the highest one.

Definition 1.15 (Committee Position). *Let \succ be a preference order and S be a k -sized committee. The position of S w.r.t \succ is the vector $pos_{\succ}(S) = (i_1, \dots, i_k)$, the vector of positions of the members of S in the total order given by \succ sorted in an increasing manner.*

Definition 1.16 (Committee Scoring Functions, [Elkind et al., 2017]). *A committee scoring function $\gamma_{m,k} : [m]_k \mapsto \mathbb{R}$ (where $[m]_k$ the set of possible committee positions) is function that maps each committee position to a score such that*

$$\forall I, J \in [m]_k, I_l \leq J_l, \forall l \implies \gamma_{m,k}(I) \geq \gamma_{m,k}(J)$$

There exists numerous committee scoring functions (see [Faliszewski et al., 2017] for a survey). These are two simple examples:

- **k-Borda** (excellence-based): $\gamma_{m,k}(i_1, \dots, i_k) = \sum_{t=1}^k \beta_m(i_t)$, where $\beta_m(i) = m - i$ is the individual Borda score.
- **Chamberlin-Courant** (diversity-based): $\gamma_{m,k}(i_1, \dots, i_k) = \beta_m(i_1)$

Now we can proceed to defining the family of voting rules which are based on scoring in the multi-winner framework.

Definition 1.17 (Committee Scoring Rules). *A committee scoring rule f associated to the family of scoring functions $(\gamma_{m,k})$, is the rule that to each election E associates the winning committee defined as:*

$$W^* = \arg \max_{|W|=k} score_f(W) = \arg \max_{|W|=k} \sum_{i \in N} \gamma_{m,k}(pos_{\succ_i}(W))$$

Simply put, the scoring rule selects the committee that maximizes the sum of the scores it gets from all the voters. We will apply this type of rules to an example both with the k-Borda and the Chamberlin-Courant scores.

Example 1.8. *Recall that for the k-Borda rule, the score function is $\gamma_{m,k}(i_1, \dots, i_k) = \sum_{t=1}^k m - i_t$. Consider the election with $\mathcal{X} = \{a, b, c\}$, $N = \{v_1, \dots, v_5\}$, $k = 2$ and all the voters apart from v_5 share the same preference profile: $a >_v c >_v b$ whereas v_5 holds the preference $b >_{v_5}$*

$c >_{v_5} a$. The scores of the 2-sized committees are:

$$\begin{cases} \text{score}_{k\text{-borda}}(a, b) = 10 \\ \text{score}_{k\text{-borda}}(a, c) = 13 \\ \text{score}_{k\text{-borda}}(b, c) = 7 \end{cases}$$

Hence, the k -Borda rule selects the winning committee $W^* = \{a, c\}$.

Example 1.9. The score function for Chamberlin-Courant is $\gamma_{m,k}(i_1, \dots, i_k) = \beta_m(i_1)$. For the same previous example 1.8 the scores of the 2-sized committees are:

$$\begin{cases} \text{score}_{CC}(a, b) = 4 \times (3 - 1) + (3 - 1) = 10 \\ \text{score}_{CC}(a, c) = 9 \\ \text{score}_{CC}(b, c) = 6 \end{cases}$$

The winning committee is thus $W^* = \{a, b\}$.

Other families of rules, generalizing the *Condorcet consistency* property to multi-winner domains have also been defined (see for instance [Gehrlein, 1985], [Coelho, 2004] and [Elkind et al., 2015]).

1.2.2 Approval-based Aggregation Rules

1.2.2.1 The Approval Framework

A simpler way of expressing preferences (or beliefs) is selecting a subset of alternatives. Instead of ranking all the candidates, a voter may *approve* as many candidate as she wishes, which is less cognitively challenging. This mode of voting is called *approval voting* [Brams and Fishburn, 2007]. An approval ballots can be interpreted in multiple ways: it can reflect *dichotomous* preferences (voters only approve the time slots corresponding to their availabilities on a meeting scheduling platform), or in case a voter has an underlying ordinal preference (or cardinal utilities), the ballot might consist of the alternatives which are above some threshold (starting from a ranking over all the candidates, a voter might only approve her top-3 preferred alternatives).

Given the set of m alternatives \mathcal{X} , each voter $i \in N$ cast her approval ballot $A_i \subseteq \mathcal{X}$. The n -uple of such subsets $A = (A_1, \dots, A_n)$ is called an approval profile, and an approval election is simply a triplet (N, \mathcal{X}, A) . In Table 1.3 we present an approval-based election with 4 voters and 3 candidates.

$$\begin{aligned} v_1 &: \{a, b, c\} \\ v_2 &: \{b, c\} \\ v_3 &: \{b, c\} \\ v_4 &: \{c\} \end{aligned}$$

Table 1.3: Approval-based election with 4 voters and 3 alternatives

We give a concrete example of an approval profile in meeting scheduling context in Figure 1.2. We can see that each voter participating in the Doodle poll ticks the slots that he or she prefers.

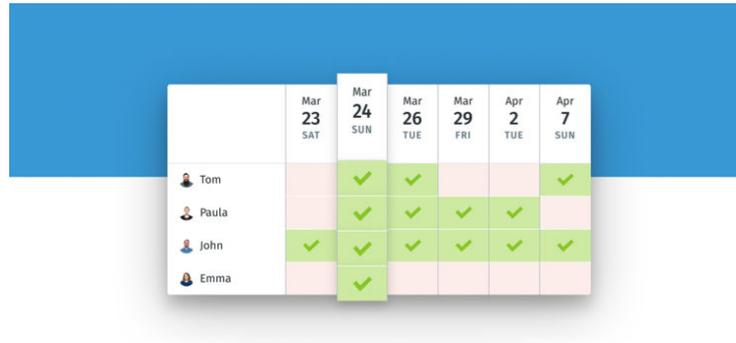


Figure 1.2: Approval profile in a Doodle poll

1.2.2.2 Single-winner Rules

Definition 1.18 (Single-winner Approval-based Rule). Consider a set N of n voters and a set \mathcal{X} of m alternatives. A single-winner approval rule f is function that maps every approval profile to a single alternative:

$$f : \begin{array}{l} \mathcal{P}(\mathcal{X})^n \longrightarrow \mathcal{X} \\ A = (A_1, \dots, A_n) \longmapsto a^* \end{array}$$

The most intuitive rule based on approvals is *Standard Approval Voting* (SAV), which outputs the alternative which is approved by the greatest number of voters.

Definition 1.19 (SAV). Given an approval profile $A = (A_1, \dots, A_n) \in \mathcal{P}(\mathcal{X})^n$, f_{SAV} outputs:

$$f_{SAV}(A) = \arg \max_{a \in \mathcal{X}} app(a)$$

where:

$$app(a) = \sum_{i=1}^n \mathbb{1} \{a \in A_i\} = |\{i \in N, a \in A_i\}|$$

We can generalize SAV to its weighted version, where *anonymity* among voters is violated, and the importance given to the vote of each of them differs.

Definition 1.20 (Weighted Approval Rule). Consider an n -length vector $\vec{w} = (w_1, \dots, w_n)$ of weights. Given an approval profile $A = (A_1, \dots, A_n) \in \mathcal{P}(\mathcal{X})^n$, $f_{\vec{w}}$ outputs the alternative:

$$f_{\vec{w}}(A) = \arg \max_{a \in \mathcal{X}} app_w(a)$$

where:

$$app_w(a) = \sum_{i=1}^n \mathbb{1} \{a \in A_i\} w_i = \sum_{i:a \in A_i} w_i$$

Remark. *Weighted variants of most of the rules that were presented in the previous sections can also be defined similarly. We chose to focus on approval voting since we will heavily use it in the next chapters.*

Example 1.10. *Consider the following the approval profile with $n = 5$ and $m = 3$:*

$$\begin{aligned} v_1 &: \{a, b, c\} \\ v_2 &: \{b, c\} \\ v_3 &: \{a, b\} \\ v_4, v_5 &: \{c\} \end{aligned}$$

We will compute the output of these election when SAV, weighted approval and size-decreasing approval are applied:

- *SAV: The approval scores of the alternatives are as follows:*

$$app(a) = 1, app(b) = 3, app(c) = 4$$

so the winner by standard approval voting is candidate c .

- *weighted AV: Imagine that the voters are weighted according to $\vec{w} = (1, 1, 2, 1, 0.5)$, so the approval of voter 3 counts twice as that of the remaining voters. The weighted approval scores are thus:*

$$\begin{aligned} app_w(a) &= 1 + 0 + 2 + 0 + 0 = 3 \\ app_w(b) &= 1 + 1 + 2 + 0 + 0 = 4 \\ app_w(c) &= 1 + 1 + 0 + 1 + 0.5 = 3.5 \end{aligned}$$

So the winning alternative is b .

Whilst weighted approval voting rules violate *anonymity* by assigning pre-defined weights to the voters, another family of rule called *size-decreasing approval rules* and introduced in [Alcalde-Unzu and Vorsatz, 2009], assign a weight to each approval ballot (and not to the voter) according to its cardinality: *smaller ballots are assigned bigger weights*. We will investigate, in later chapters, some situations where the use of size-decreasing approval rules is an efficient way of estimating a ground truth.

Definition 1.21 (Size-decreasing Approval Rule). *Consider a function*

$$\begin{aligned} f : \mathcal{P}(\mathcal{X})^n &\longrightarrow \mathcal{X} \\ (A_1, \dots, A_n) &\mapsto f(A_1, \dots, A_n) \end{aligned}$$

that, for each approval profile $A = (A_1, \dots, A_n)$, assigns a winning alternative $f(A_1, \dots, A_n)$ in \mathcal{X} . We say that f is a size-decreasing approval rule if there exists a vector $\vec{w} = (w_0, \dots, w_m)$ in \mathbb{R}^{m+1} verifying $w_j > w_{j+1}$ for all $1 \leq j \leq m - 2$, such that:

$$f(A_1, \dots, A_n) = \arg \max_{a \in \mathcal{X}} app_{\vec{w}}(a)$$

where $app_{\vec{w}}$ is the weighted approval score defined by:

$$app_{\vec{w}}(a) = \sum_{i:a \in A_i} w_{|A_i|}$$

Example 1.11. Recall the election of Example 1.10. Consider the satisfaction approval rule [Brams and Kilgour, 2015], which is a size-decreasing approval rule that, for a ballot A_i , assigns a weight $w_{A_i} = \frac{1}{|A_i|}$. For instance, $A_1 = \{a, b, c\}$ contains three alternatives and will be assigned a weight $\frac{1}{3}$. The alternatives' scores are then:

$$\begin{aligned} app_w(a) &= \frac{1}{3} + \frac{1}{2} \\ app_w(b) &= \frac{1}{3} + 2 \times \frac{1}{2} \\ app_w(c) &= \frac{1}{3} + \frac{1}{2} + 1 \end{aligned}$$

The winner is thus alternative c .

1.2.2.3 Multi-winner Rules

Definition 1.22 (Multi-winner Approval Election [Lackner and Skowron, 2020a]). An election E is a quadruplet $E = (\mathcal{X}, N, A, k)$ where:

- \mathcal{X} is the set of m candidates.
- N is the set of n voters.
- $A = (A_1, \dots, A_n)$ is the approval profile consisting of either n subsets of \mathcal{X} (approval-based voting), representing the preferences of each voter.
- $k \leq m$ is the committee-size.

Definition 1.23 (Multi-Winner Voting Rule). A voting rule is a mapping f that maps each election to a subset of A with cardinality k representing the winning committee.

The approval-based voting framework differs from the ranking-based one by its simpler input: each voter casts a ballot $A_i \subset A$ containing all the alternatives (or candidates) that the voter approves of.

We'll begin by introducing the simple *top- k* rule, which ignores any interaction among the candidates, and selects the k alternatives with the largest approval scores.

Definition 1.24 (Top- k AV Rules). Given an approval profile $A = (A_1, \dots, A_n) \in \mathcal{P}(\mathcal{X})^n$ and a committee-size k , and suppose that $app(a_1) > \dots > app(a_m)$ the top- k rule outputs:

$$f_{top-k}(A) = \{a_1, \dots, a_k\}$$

Notice that this definition can naturally be extended to a *weighted top- k* rule as in Definition 1.20.

We will present a family of approval voting rules, introduced for the first time more than a century ago.

Definition 1.25 (w-AV Rules [Thiele, 1895]). *Let $w = (w^1, \dots, w^j)$ be a family of vectors such that w^j consists of j real numbers. The w-AV rule associated to this family is the rule that for an election E selects the committee maximizing the w-AV score given by:*

$$Score_{w-AV}(S) = \sum_{i \in N} \sum_{j=1}^{|S \cap A_i|} w_j^k, \forall S \subset \mathcal{X}, |S| = k$$

Each voter assigns a score $\sum_{j=1}^{|S \cap A_i|} w_j^k$ to the committee S which only depends on the number of members of S whom the voter approves. The overall score is simply the sum of the voters' scores.

Example 1.12.

- *The Approval Voting (AV) method is a w-AV rule associated to the sequence of vectors $w^j = (1, \dots, 1)$. So the score a voter assigns to a committee is simply the number of its members he approves.*
- *The Proportional Approval Voting (PAV) method is a w-AV rule associated to the sequence of vectors $w^j = (1, \frac{1}{2}, \dots, \frac{1}{j})$.*

Another family of rules with a load-balancing approach was introduced by Lars Edvard Phragmén [Brill et al., 2017a]. For example, the *max-Phragmén* rule [Brill et al., 2017b] consists in optimizing a load distribution among voters, and has nice properties with regards to proportionality.

Although we solely defined elections with a pre-fixed committee-size k , some works have considered the case when no such constraint is imposed [Kilgour, 2016]. Some rules adequate for a *variable number of winners* have been introduced (for instance, the *net approval rule* [Kilgour, 2010] and the *net satisfaction approval rule* [Brams and Kilgour, 2015]).

Abstract

We review the existing literature related to epistemic social choice. We begin with the first development initiated by the *Condorcet Jury Theorem* which states that the majority vote of equally reliable voters asymptotically uncovers the ground truth almost certainly. Then we list subsequent works which extend it into broader settings: heterogeneous voters' competencies, correlated votes and approval-based votes. We also review works dealing with the related problem of *sampling complexity*, studying the number of required votes to uncover the truth with a given error tolerance using a fixed aggregation rule. Moreover, we present the approaches defining the *optimal estimation rules* in different settings going from binary issues to more complex forms of votes. We partition them according to the nature of the input: single alternative, rankings, partial orders, pairwise comparisons, truncated top- k approvals, approval ballots and CP-nets. The last section presents some other epistemic voting research directions that have been studied like designing proper incentive mechanisms or axiomatic analysis of collective annotations.

2.1 Condorcet Jury Theorem

Epistemic social choice witnessed its initial development with Marquis de Condorcet, dating back to the 18th century. It considers votes as noisy signals of an objective *ground truth*, and uses maximum likelihood estimation techniques to uncover the latter. The initial seminal result, known as *Condorcet Jury Theorem*, proved in Condorcet's original paper [Condorcet, 1785] by a combinatorial argument lead to the later ramifications and extensions. It considers two alternatives, with one candidate being *objectively better* than the other, but the two of them are *a priori* equally likely to be the best one. It also considers n independent voters that are equally reliable, i.e they are equally likely to select the objectively best alternative, and states that if every voter votes for the correct alternative with probability $p > \frac{1}{2}$, then the majority rule maximizes the likelihood of coinciding with the ground truth among all estimators and that it outputs the

correct decision with a probability that increases with the number of voters and tends to 1 when the latter grows to infinity. Formally, the theorem can be stated as follows:

Theorem 2.1 (Condorcet Jury Theorem). *Consider a set of two alternatives $\mathcal{X} = \{a, b\}$ and a set of n voters N where n is odd. We suppose that there is a true ranking over the alternatives (called the ground truth), denoted \succ^* , and we suppose that the two possibilities are a priori equally likely:*

$$P(a \succ^* b) = P(b \succ^* a) = \frac{1}{2}$$

Each voter casts a pairwise comparison \succ_i over \mathcal{X} independently of the other voters, with the following conditional probability:

$$P(a \succ_i b | a \succ^* b) = P(b \succ_i a | b \succ^* a) = p$$

where $p > \frac{1}{2}$. Let M_n be the outcome of the simple majority rule:

$$M_n = \begin{cases} a \succ b & \text{if } |i \in N, a \succ_i b| > \frac{n}{2} \\ b \succ a & \text{else.} \end{cases}$$

Then we have that

$$P(M_n = \succ^*) < P(M_{n+2} = \succ^*)$$

and that:

$$\lim_{n \rightarrow +\infty} P(M_n = \succ^*) = 1$$

The CJT has been since generalised in many directions, a survey can be found in [Elkind and Slinko, 2016].

2.1.1 Heterogeneous Reliabilities

[Owen et al., 1989] extend the CJT for a group of heterogeneous voters. Still in a dichotomous situation, with two a priori equally likely alternatives $\mathcal{X} = \{a_1, a_2\}$, and an unknown ground truth $a^* \in \mathcal{X}$, the n independent voters N (where n is odd) no longer have a common competence level $p \in (\frac{1}{2}, 1)$ but each voters $i \in N$ has her own reliability $p_i \in (0, 1)$:

$$P(A_i = a^* | a^*) = p_i$$

The main result states that regardless of the reliabilities' distribution, it suffices that the averaged competence across all voters exceeds $\frac{1}{2}$ to guarantee the asymptotic part of the CJT.

Theorem 2.2 (Distribution-free Generalization of CJT). *Let \bar{p} be the average reliability of the group of n voters N , and let M_n be the outcome of the majority vote, then:*

- If $\bar{p} > \frac{1}{2}$ then $\lim_{n \rightarrow +\infty} P(M_n = a^*) = 1$.
- If $\bar{p} < \frac{1}{2}$ then $\lim_{n \rightarrow +\infty} P(M_n = a^*) = 0$.
- If $\bar{p} = \frac{1}{2}$ then $1 - e^{-\frac{1}{2}} < \lim_{n \rightarrow +\infty} P(M_n = a^*) < e^{-\frac{1}{2}}$.

2.1.2 Correlated Votes

The independent votes assumption has also been relaxed. In [Shapley and Grofman, 1984], the authors present some situation of statistically dependent votes where majority is no longer optimal. Consider a dichotomous situation with three voters. Suppose that we know for sure that one of them is wrong (we do not know exactly which one), while the two others have a high probability p of being correct. We can check that the probability of the majority outcome to coincide with the ground truth is equal to p^2 . Now consider the rule that follows majority in a two-to-one situation but opposes it in case of unanimity. Its probability of uncovering the ground truth is equal to $p^2 + (1 - p)^2$, which makes the majority rule suboptimal.

[Ladha, 1995] studies different probabilistic distributions of correlated votes. For equally reliable voters in a binary issue settings, with reliability $p \in (0.5, 1)$, let ρ_{ij} be the correlation coefficient of voters $i, j \in N$, it is proved that as long as the average correlation is less than a given threshold, the majority outcome is more likely to uncover the ground truth than any single vote.

Proposition 2.3. *Let $\bar{\rho}$ be the average correlation coefficient. Then it suffices to have:*

$$\bar{\rho} < 1 - \frac{n}{n-1} \frac{p - 0.25}{p^2}$$

to guarantee that:

$$P(M_n = a^*) > p$$

Example 2.1. *Consider a set of correlated 20 voters N with a common reliability $p = 0.8$. To guarantee that the majority outcome is more likely to coincide with the ground truth than any single vote, it suffices that:*

$$\bar{\rho} < 1 - \frac{n}{n-1} \frac{p - 0.25}{p^2} = 0.095$$

If the 20 voters are more reliable $p = 0.95$ than it suffices that:

$$\bar{\rho} = 0.184$$

It is also shown that provided that this average correlation goes to 0 as the number of voters grows to infinity, the probability $P(M_n = a^*)$ will tend to 1. Similar results are given for unimodal symmetric distributions of votes, as well as the hypergeometric and Polya distributions.

2.1.3 More than Two Alternatives

[List and Goodin, 2001] extends the CJT to the case with arbitrarily many alternatives. Let $\mathcal{X} = \{a_1, \dots, a_m\}$ be the set of alternatives with $m \geq 2$, a priori, equally likely to be the ground truth a^* . Consider a set of n independent voters N . Suppose that there exist $p_1, \dots, p_m \in (0, 1)$ where p_j represents the probability that voter $i \in N$ selects the alternative a_j . The result states that if voters are more likely to select the ground truth than any other alternative, then the plurality outcome will most likely coincide with a^* than any other alternative, and this with a probability that grows to 1 as the number of voters grows to infinity.

Theorem 2.4. *Let j_* be the index of the ground truth alternative, and let Pl_n be the outcome of plurality. If $p_{j_*} > p_j, \forall j \neq j_*$ then:*

- $P(Pl_n = a_{j_*}) > P(Pl_n = a_j), \forall j \neq j_*$.
- $\lim_{n \rightarrow +\infty} P(Pl_n = a_{j_*}) = 1$

2.1.4 Approval Profiles

In [Everaere et al., 2010] the CJT is formulated in *belief merging* framework. It is proven, among other results, that for a set of n homogeneous voters and m alternatives \mathcal{X} , if the voters submit approval profiles, then the SAV rule would output the ground truth with probability approaching 1 as n grows to infinity, provided that the voters are competent enough: a voter is competent when the probability of including the ground truth alternative in her ballot is strictly higher than that of including any other alternative.

Theorem 2.5. *For any voter $i \in N$, suppose that:*

$$P(a^* \in A_i) > P(a \in A_i), \forall a \neq a^*$$

Then:

$$\lim_{n \rightarrow +\infty} P(SAV_n = a^*) = 1$$

where SAV_n is the winner by standard approval voting.

2.1.5 Abstract Voting Rules: General Results

By introducing a class of abstract voting rules, the *mean partition rules*, [Pivato, 2017] provides generalizations of CJT to settings unifying approval voting, rankings and quantitative assessments. The author considers the set of alternatives \mathcal{X} to be a metric space and the set of possible votes (inputs) \mathcal{V} to be a subset of an inner-product space. After defining the mean partition rules, which are functionally determined by the average vote, a general notion of noise model is introduced: each voter has a *behaviour model* $\rho^i : \mathcal{X} \rightarrow \Delta(\mathcal{V})$ which describes the probability distribution of the voter's input given each true state. These behaviour models are elements of some family of distributions referred to as *populace*. In the case of independent votes, it is proven that if the variance of all the behaviour models is bounded and if the expected vote for all the possible distributions is such that it is interpreted by the mean partition rule as the ground truth state, then the probability of this rule to output the ground truth grows to 1 as the number of voters goes to infinity. This results holds even if there are infinitely many possible states. An extension to the case of correlated votes is also provided. It is proven that mean partition rules will asymptotically output the true state if the profile is sampled from a distribution which guarantees that the average voter's input is correct, that its variance is bounded and also that the correlations among the voters' errors vanish as their number grows to infinity.

2.1.6 Sampling Complexity

Sampling complexity is the problem of estimating the number of votes needed to uncover the ground truth with high probability via a specific voting rule, under a given noise model.

[Caragiannis et al., 2016] suppose the existence of a true ranking \succ^* over a set of m alternatives \mathcal{X} . A set of n (which is not fixed) voters submit their opinions about the ground truth as a profile of linear orders $\pi = (\succ_1, \dots, \succ_n)$. For a given social welfare function (SWF) r the authors define its worst-case accuracy as the minimum probability that it succeeds in uncovering the ground truth by aggregating a given set of n votes:

$$Acc^r(n) = \min_{\succ^* \in L(\mathcal{X})} \sum_{\pi \in L(\mathcal{X})^n} P(\pi | \succ^*) P(r(\pi) = \succ^*)$$

For some $\epsilon > 0$, let $N^r(\epsilon) = \min\{n | Acc^r(n) \geq 1 - \epsilon\}$ be the minimum number of samples needed to uncover the ground truth by a probability greater than $1 - \epsilon$. The results concern votes generated from a Mallows noise model [Mallows, 1957] where the probability of casting an order \succ given a ground truth ranking \succ^* is:

$$P(\succ | \succ^*) = \frac{1}{\beta} \phi^{d_{KT}(\succ, \succ^*)}$$

with d_{KT} being the Kendall-Tau distance, ϕ a common reliability parameter, and β its associated normalization factor:

$$\beta = \sum_{\succ \in L(\mathcal{X})} \phi^{d_{KT}(\succ, \succ^*)}$$

It is first proven that given Mallows noise model and any tolerance $\epsilon > 0$, the Kemeny rule has the optimal sample complexity.

Theorem 2.6. *For any number of alternatives m and any tolerance $\epsilon > 0$, given noisy samples generated from a Mallows model, the Kemeny rule has the optimal sampling complexity:*

$$N^{Kemeny}(\epsilon) \leq N^r(\epsilon)$$

for any social welfare function r .

More generally, the paper defines a family of rules that contains the Kemeny rule (along with Copeland and other rules), called *Pairwise-Majority consistent* rules (PM-c): whenever there exist a ranking \succ over \mathcal{X} such that a majority of votes agree with \succ on every pair of alternatives $a, b \in \mathcal{X}$, a PM-c rule must output \succ . These rules only require a logarithmic number of samples in m and $\frac{1}{\epsilon}$ to uncover the truth under a Mallows noise model.

Theorem 2.7. *For any tolerance $\epsilon > 0$ and any number of alternatives m , a PM-c rule determines the true ranking with a probability at least $1 - \epsilon$ given $O(\log \frac{m}{\epsilon})$ votes from a Mallows distribution.*

Furthermore, it is shown that voting rules cannot do better: the logarithmic sampling complexity is proven to be a lower bound for all the voting rules for $\epsilon \leq \frac{1}{2}$.

Theorem 2.8. *For a tolerance $0 < \epsilon \leq \frac{1}{2}$, any voting rule requires $\Omega(\log \frac{m}{\epsilon})$ votes issued from a Mallows model to uncover the ground truth with probability higher than $1 - \epsilon$.*

The paper [Caragiannis and Micha, 2017] supposes a ground truth ranking, noisy rankings generated from Mallows distribution and truncated approval ballots as input, where each voter approves the top k alternatives in her underlying ordinal preference. The paper focuses on sampling complexity where the ground truth is a ranking and the rule ranks the alternatives according to their approval scores. A negative result is first given, stating that an exponential number of votes is needed to uncover the true ranking via k -approval voting (ranking the alternatives according to the number of approvals they get from the top- k truncated approval ballots) for any $k \in [m - 1]$. Formally, for any $\epsilon \in [0, 1/4]$ at least $\mathcal{O}(\phi^{-m/2})$ Mallows k -approval ballots are needed to uncover the true ranking with probability at least $1 - \epsilon$. Nevertheless, when each voter is asked to approve r alternatives, where r is each time drawn uniformly at random, this complexity becomes linear even for cases where voters' reliabilities are heterogeneous.

Theorem 2.9. *For any $\epsilon > 0$, when the average reliability is equal to p^* , it suffices to have n voters such that:*

$$n \geq \frac{2(m-1)^2}{(2p^*-1)^2} \ln\left(\frac{m}{\epsilon}\right)$$

to uncover the ground truth ranking with probability at least $1 - \epsilon$ via randomized approval voting.

Example 2.2. *Consider n voters with average reliability $p^* = 0.8$ who face $m = 5$ alternatives. It suffices to have:*

$$n \geq \frac{2(m-1)^2}{(2p^*-1)^2} \ln\left(\frac{m}{\epsilon}\right) = 347.73$$

votes to uncover the ground truth ranking with probability $1 - \epsilon = 0.9$ via randomized approval voting. If the average reliability is $p^ = 0.95$ and the number of alternatives is $m = 4$ than randomized approval voting uncovers the ground truth ranking with probability $1 - \epsilon = 0.9$ if the number of voters is greater than $n = 82$.*

These bounds are then assessed on simulated data.

When the ground truth is a single alternative $a^* \in \mathcal{X}$, the input are approval ballots submitted by n heterogeneous independent voters (A_1, \dots, A_n) and the estimation rule is the standard approval voting rule, [Karge and Rudolph, 2022] offer a tight lower-bound on the number of needed experts to guarantee a minimal probability P_{min} of uncovering the ground truth. This bound holds when the voters are supposed to be Δp -reliable, meaning that the average probability of approving the ground truth alternative a^* is at least higher by Δp than the average probability of approving any other alternative $b \neq a^*$:

$$\frac{1}{n} \sum_{i=1}^n P(a \in A_i | a^* = a) \geq \Delta p + \frac{1}{n} \sum_{i=1}^n P(b \in A_i | a^* = a)$$

Theorem 2.10. *In a $\Delta p > 0$ reliable group of n voters with m alternatives \mathcal{X} , the worst case approval vote success probability is at least P_{min} whenever:*

$$n \geq \min \left(\frac{2}{\Delta p^2} \log Q, 1 + \left(\frac{1}{\Delta p^2} - 1 \right) Q \right)$$

where:

$$Q = 2 \frac{m - 1}{1 - P_{min}}$$

Example 2.3. *If voters face 5 alternatives while being $\Delta p = 0.1$ reliable then it would suffice to have $n = 877$ voters to guarantee at least a probability $P_{min} = 0.9$ of uncovering the ground truth alternative via approval voting.*

2.2 Optimal Estimation Rules

Whereas the previous section focused on existing works dealing with CJT-like results showing under which conditions certain known rules are able to uncover the ground truth with a large enough number of voters, this section will address another direction: given some settings (voters, alternatives and input profiles) which rule is the most likely to output the ground truth?

Since Maximum Likelihood Estimation is the most widely used technique to answer this question, we begin this section by briefly introducing it.

2.2.1 Brief Refresher on Maximum Likelihood Estimation

When the studied phenomenon is assumed to follow a parametric probability distribution (the parameters being a priori unknown and used to approximately model the behaviour of the quantity at hand), we define the *likelihood* function as the probability of observing some realisations of the random phenomenon for different values of the parameter. Hence, once the observations are fixed, two different model parameters would yield two different data likelihoods. The probability distribution associated to the parameter with a greater likelihood is then more likely to be representative of the underlying model if we base our information solely on the considered sample.

Maximum likelihood estimation is thus a statistical method for estimating these parameters once we observe some data samples. The idea is to take the parameters' value that would make the observed data most likely given that it does follow the assumed probabilistic model.

Formally, suppose that a set of observations y_1, \dots, y_n is sampled from a joint parametric probability distribution f_θ where $\theta \in \Theta$ is unknown. The goal is to determine $\hat{\theta}$ that makes the joint probability of the observed samples maximal. We thus define the likelihood of the sample y_1, \dots, y_n given a parameter θ as:

$$\mathcal{L}(y_1, \dots, y_n; \theta) = f_\theta(y_1, \dots, y_n)$$

In most cases we will encounter in the next sections, the samples y_1, \dots, y_n are independent and identically distributed according to an unknown P_θ . The likelihood then reads:

$$\mathcal{L}(y_1, \dots, y_n; \theta) = \prod_{i=1}^n P_\theta(y_i)$$

The Maximum Likelihood Estimator (MLE) is defined as:

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}(y_1, \dots, y_n; \theta)$$

Example 2.4. *Suppose that we toss a biased coin three times independently. The probability of the unfair coin to yield a Head is unknown and our goal is to estimate it from the observed realisations y_1, y_2, y_3 . We thus introduce the parameter $\theta \in [0, 1]$ where $\theta = P_\theta(Y = \text{Head})$. We want to find $\hat{\theta}$ that maximizes the likelihood of the observed sample:*

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}(y_1, \dots, y_n; \theta) = \prod_{i=1}^3 P_\theta(y_i)$$

First suppose that all the tosses gave a Head: $y_1 = y_2 = y_3 = \text{Head}$:

$$\mathcal{L}(y_1, \dots, y_n; \theta) = \prod_{i=1}^3 P_\theta(y_i) = \theta^3$$

The likelihood is then maximal if and only if $\theta = 1$ so $\hat{\theta} = 1$. Similarly, if all the tosses gave a Tail, then the maximum likelihood estimator would be $\hat{\theta} = 0$.

Now suppose that $y_1, y_2 = \text{Head}$ and $y_3 = \text{Tail}$. We have:

$$\mathcal{L}(y_1, \dots, y_n; \theta) = \theta^2 \times (1 - \theta)$$

This function reaches a unique maximum in $\hat{\theta} = \frac{2}{3}$.

Since in epistemic social choice the votes are seen as noisy estimates of a correct state, the probability distribution which the votes follow depend on the *ground truth*. Each different possible value of the ground truth is associated to a different likelihood of the actual votes. The model is then parameterized (we will see that the parameter sometimes includes other variables than the ground truth such as the voters' reliabilities), and the maximum likelihood estimation method is often used to uncover these parameters, hence the ground truth.

2.2.2 Voting on Binary Issues

In [Nitzan and Paroush, 1982], each voter has a probability $p_i \in (0, 1)$ of selecting the correct alternative among the two possibilities. The maximum likelihood estimation rule is shown to be a weighted majority which assigns the weight $w_i = \log \frac{p_i}{1-p_i}$ to each voter $i \in N$.

[Ben-Yashar and Nitzan, 1997] generalizes this line of research on three axes:

- There are distinct payoffs (or costs) for the four different combinations of ground truth and chosen alternative: the objective is hence to maximize this payoff rather than simply maximizing the likelihood.
- The prior probabilities of alternatives differ.
- Voters' reliabilities (which are heterogeneous) depend on the ground truth.

Formally, the paper considers a set of n independent voters N . The set of alternatives is $\mathcal{X} = \{-1, 1\}$, where 1 represents a good project worthy of investment. Each voter casts a response $x_i \in \mathcal{X}$ and has a two parameter reliability:

$$p_i = P(x_i = 1|a^* = 1), 1 - q_i = P(x_i = -1|a^* = -1)$$

The prior probability that the project is good is denoted $P(a^* = 1) = \alpha$. The payoff/cost of selecting a project by the committee is as follows:

- Selecting a good project: $u(1, 1)$
- Rejecting a good project: $u(-1, 1)$
- Selecting a bad project: $u(1, -1)$
- Rejecting a bad project: $u(-1, -1)$

Theorem 2.11. *The aggregation rule $f : (x_1, \dots, x_n) \mapsto x \in \mathcal{X}$ that maximizes the expected payoff is given by:*

$$f(x_1, \dots, x_n) = \text{sign} \left(\sum_{i=1}^n \frac{\beta_i^1 + \beta_i^2}{2} x_i + \psi + \gamma + \delta \right)$$

where:

$$\begin{cases} \beta_i^1 &= \log \frac{p_i}{1-p_i} \\ \beta_i^2 &= \log \frac{1-q_i}{q_i} \\ \psi &= \frac{1}{2} \sum_{i=1}^n \log \frac{p_i(1-p_i)}{q_i(1-q_i)} \\ \gamma &= \log \frac{\alpha}{1-\alpha} \\ \delta &= \log \frac{u(1,1)-u(-1,1)}{u(-1,-1)-u(1,-1)} \end{cases}$$

Example 2.5. *Consider a set of $n = 3$ voters N , deciding whether to invest in a project which has a prior probability of being good equal to $\alpha = 0.6$. The voters have the following heterogeneous expertise parameters:*

$$(p_1, q_1) = (0.9, 0.4), (p_2, q_2) = (0.6, 0.1), (p_3, q_3) = (0.7, 0.3)$$

Let us suppose that the utility are:

$$u(1, 1) = 2, u(-1, 1) = -\frac{1}{2}, u(-1, -1) = 1, u(1, -1) = -1$$

To find the decision that maximizes the expected payoff we need to compute the parameters $\beta_i^1, \beta_i^2, \psi, \gamma$ and δ . We have:

$$(\beta_1^1, \beta_1^2) = (2.19, 0.4), (\beta_2^1, \beta_2^2) = (0.4, 2.19), (\beta_3^1, \beta_3^2) = (0.85, 0.85)$$

and:

$$\psi = 0, \gamma = 0.4, \delta = 0.22$$

If only the voter v_1 approves the project, that is, if:

$$x_1 = 1, x_2 = -1, x_3 = -1$$

the optimal decision is:

$$f(1, -1, -1) = \text{sign} \left(\sum_{i=1}^3 \frac{\beta_i^1 + \beta_i^2}{2} x_i + \psi + \gamma + \delta \right) = \text{sign}(-0.23) = -1$$

which means that the project should be rejected.

If the voter v_2 changes her mind and approves the project then the optimal decision is:

$$f(1, 1, -1) = \text{sign} \left(\sum_{i=1}^3 \frac{\beta_i^1 + \beta_i^2}{2} x_i + \psi + \gamma + \delta \right) = \text{sign}(2.36) = 1$$

which means that they should invest in the project.

More recent work [Halpern et al., 2021] relaxed the assumption that voters' reliabilities are totally known to the aggregation mechanism. The paper supposes that only a ranking over the voters according to their expertise is given. The set of alternatives is $\mathcal{X} = \{0, 1\}$, and the voters' reliabilities $p = (p_1, \dots, p_n)$ are unknown but are supposed to be greater than $\frac{1}{2}$. Since a ranking over the voters is given, let \mathcal{P}_n be the set of reliabilities compatible with it. Each voter casts an answer $x_i \in \{0, 1\}$. The selected alternative by the committee is denoted y . Three objectives are defined:

- Minimizing the distortion given by:

$$\sup_{p \in \mathcal{P}_n} \frac{P(x_1, \dots, x_n | a^* = 1 - y, p)}{P(x_1, \dots, x_n | a^* = y, p)}$$

It is the worst-case ratio, among all feasible reliabilities, of the likelihood of a wrong estimate to the likelihood of picking the correct one. The optimal rule for this criterion is given by:

$$f(x_1, \dots, x_n) = \arg \max_{y \in \{0, 1\}} s(y)$$

where :

$$s(y) = \max_{k \in \{0, \dots, n\}} \sum_{i=1}^k (\mathbb{1}[x_i = y] - \mathbb{1}[x_i = 1 - y])$$

- Maximizing the pessimistic likelihood:

$$\inf_{p \in \mathcal{P}_n} P(x_1, \dots, x_n | a^* = y, p)$$

It is the worst-case likelihood of selecting the correct alternative. It is proven that the optimal rule for this objective function is the one that outputs the majority decision. In case of ties, the opposite of the less expert voter's answer is output.

- Maximizing the optimistic likelihood:

$$\sup_{p \in \mathcal{P}_n} P(x_1, \dots, x_n | a^* = y, p)$$

A polynomial-time algorithm is presented to maximize this objective.

2.2.3 Input Profiles are Full Rankings

While the initial Condorcet article focused on the two alternative settings, it did nonetheless propose a noise model for generating rankings over more than two alternatives consisting of the following procedure:

- All the votes are independent.
- For any two alternatives, the voter compares them correctly with a probability $p > 1/2$.
- These pairwise comparisons are mutually independent.
- The votes are transitive.

However [Young, 1986, Young, 1988, Young, 1995] noticed the possible inconsistency within the previous assumptions (independent pairwise comparisons might yield intransitive relations over the alternatives). Nevertheless, when the ground truth is a ranking, Young has proven that the Kemeny rule is the maximum likelihood estimation rule for homogeneously reliable voters following the Mallows noise model. Recall that the probability of casting an order \succ given a ground truth ranking \succ^* for this distribution is:

$$P(\succ | \succ^*) = \frac{1}{\beta} \phi^{d_{KT}(\succ, \succ^*)}$$

with d_{KT} being the Kendall-Tau distance, ϕ a common reliability parameter, and β its associated normalization factor:

$$\beta = \sum_{\succ \in L(\mathcal{X})} \phi^{d_{KT}(\succ, \succ^*)}$$

2.2.4 Input Profiles are Partial Orders

[Xia and Conitzer, 2011] defines noise models that tolerate partial orders (transitive and antisymmetric binary relation) as input. For a set of alternatives $\mathcal{X} = \{a_1, \dots, a_m\}$ and a finite set of

outcomes (possible ground truths) \mathcal{O} (single alternatives, rankings or a subset of alternatives), the n voters cast partial orders by independently making pairwise comparisons with possible incomparability:

$$P(a_j \succeq_i a_k | o) + P(a_k \succeq_i a_j | o) + P(a_j \sim_i a_k | o) = 1, \forall o \in \mathcal{O}$$

The correspondent maximum likelihood estimation rule for such noise models, whenever they satisfy neutrality, is a neutral pairwise scoring rule r where for any profile of partial orders P :

$$r(P) = \arg \max_{o \in \mathcal{O}} \sum_{i \in N} \sum_{(a_j, a_k) \in \succeq_i} s(a_j, a_k, o)$$

where $s : \mathcal{X} \times \mathcal{X} \times \mathcal{O} \rightarrow \mathbb{R}$ is a pairwise scoring function satisfying $s(a, a, o) = 0$, for every alternative $a \in \mathcal{X}$ and outcome $o \in \mathcal{O}$.

In the paper's last section, the multi-winner selection problem is also tackled. The input consist of partial orders and the set of possible ground truth is $\mathcal{S}_k = \{S \subseteq \mathcal{X}, |S| = k\}$. A noise model is defined where, given the correct set of k alternatives, a voter is more likely to rank a winner alternative better than a loser one than the opposite. The optimal rule is then proven to be a pairwise scoring rule. The main result of the section is a complexity result, stating that the decision version of the problem L-EVALUATION is NP-complete. In L-EVALUATION we are given a committee-size k and a likelihood threshold t and a profile P and asked whether there exist a committee S with a likelihood greater than t .

2.2.5 Input Profiles are Pairwise Comparisons

[Procaccia et al., 2012] consider a ground truth ranking \succ^* , with a winning alternative a^* , different input forms and different objectives. A general noise model, called *noisy choice model* is introduced where the probability of a dataset (rankings, or collection pairwise comparisons) given the ground truth is as follows:

$$P(D | \succ^*) = \frac{\gamma^{d(\succ^*, D)}}{Z_\gamma}$$

where:

$$d(\succ^*, D) = \sum_{a, b \in \mathcal{X}: a \succ^* b} n_{ba}$$

and where n_{ba} is the number of voters preferring b to a in the input D . This model generalizes both the Condorcet pairwise model and the Mallows ranking model. The paper also considers three different objectives, namely:

- **Objective 1:** Finding a subset of size k that maximizes the probability of containing the winning alternative in the ground truth ranking a^* .
- **Objective 2:** Finding a subset of size k with highest likelihood of coinciding with the top- k alternatives in the ground truth ranking.

- **Objective 3:** Finding the ordered tuple of k alternatives with highest likelihood of coinciding with the k ordered top alternatives in the ground truth.

All three of the above objectives are proved to be NP-Hard. Nonetheless, it is also proved that in high noise settings, intuitive and tractable rules are able to retrieve the optimal sets. The authors introduce the *extended scoring rules* which, given a dataset D , assigns to each alternative $a \in \mathcal{X}$ the score:

$$s(a) = \sum_{b \neq a} n_{ab}$$

Theorem 2.12. *For all n and m , there exist a noise level $\gamma' < 1$ such that for all $\gamma \geq \gamma'$, the optimal solutions to Objective 1 and Objective 2 are within the sets of k alternatives with highest scores $s(a)$.*

For Objective 3, the *scored tuples* method is defined: for a k -tuple of alternatives (a_1, \dots, a_k) the score is:

$$s(a_1, \dots, a_k) = \sum_{j=1}^k \sum_{b \neq a_j} n_{a_j b} - \sum_{1 \leq j < l \leq k} n_{a_l a_j}$$

Theorem 2.13. *For all n and m , there exist a noise level $\gamma' < 1$ such that for all $\gamma \geq \gamma'$, the optimal solutions to Objective 3 are within:*

$$\arg \max_{(a_1, \dots, a_k) \in \mathcal{X}^k} s(a_1, \dots, a_k)$$

[Elkind and Shah, 2014] argues that both the ground truth and the votes might contain intransitive cycles (e.g. when alternatives are compared with regards to multiple criteria). The ground truth is thus a tournament T^* (a set of pairwise comparisons on \mathcal{X}). The n independent voters N also cast votes as tournaments. The noise model is a generalization of the Mallows distribution: given T^* and for each pair of alternatives $a, b \in \mathcal{X}$, a voter has a probability p of retaining the same comparison of a and b as in T^* , and a probability $1 - p$ of flipping it. The probability of a vote T_i is thus:

$$P(T_i | T^*) = p^{\binom{m}{2} - d(T_i, T^*)} (1 - p)^{d(T_i, T^*)}$$

The paper first proves that whenever a Condorcet winner exists in T^* , it can be estimated by maximum likelihood estimation in polynomial time. Two extreme cases are then considered:

- When voters are accurate ($p \rightarrow 1$): the MLE of the tournament winner is a refined Tideman winner, where the Tideman rule outputs the alternative with lowest score $s(a)$ given by the expression:

$$s(a) = \sum_{b \neq a} \max(0, n_{ba} - n_{ab})$$

The estimate can be computed in polynomial time.

- When voters are inaccurate ($p \rightarrow 0$): the MLE of the tournament winner is a refined Borda count winner and can also be computed in polynomial time.

2.2.6 Input Profiles are Truncated k-Approval Votes

[Procaccia and Shah, 2015] consider a k -approval framework where the ground truth is a full ranking over the set of alternatives \mathcal{X} and while each voter has also a full ranking \succ_i , they only cast a k -approval ballot consisting of their top- k alternatives. The paper considers variants of the Mallows noise model with homogeneous reliability ϕ :

$$P(\succ_i \mid \succ^*) = \frac{1}{\beta} \phi^{d(\succ_i, \succ^*)}$$

where d is a distance over permutations $L(\mathcal{X})$, namely, the Kendall Tau, Spearman, Maximum Displacement, Cayley and Hamming distances. The authors prove that for the specific case of the Mallows model with Kendall Tau distance, the maximum likelihood estimators of the winning alternative given k -approval ballots coincide exactly with the set of winning alternatives by standard approval voting. However, for the remaining distances, they construct profiles where it is not the case. Nonetheless, for any neutral noise model respecting mild assumptions about the voters' reliability, approval voting is optimal given plurality ($k = 1$) and veto ($k = m - 1$) votes. The paper concludes with experiments on synthetic and real data that shows that approval voting is close to optimal in practice.

2.2.7 Input Profiles are Approval Votes

[Ben-Yashar and Paroush, 2001] consider the problem of selecting k from m projects given some committee of experts' opinions. The set of projects is $T = \{t_1, \dots, t_m\}$, and the ground truth is a subset of objectively "good" projects. Each project has a hidden true ground truth $s_j \in \{-1, 1\}$, where 1 means that the project is good. Project t_j has a prior probability α_j of being good. Each voter $i \in N$ casts an opinion $x_i^j \in \{-1, 1\}$ on each project. The noise model has two parameters per voter and per project (p_i^j, q_i^j):

$$P(x_i^j = 1 \mid s_j = 1) = p_i^j, P(x_i^j = 0 \mid s_j = 0) = 1 - q_i^j$$

The utility/cost of selecting or rejecting a good/bad projects are defined exactly like in [Ben-Yashar and Nitzan, 1997]. The goal is to select the k projects that maximize the expected payoff. The article proves that the optimal rules select the k projects with the highest Δ_j where:

$$\Delta_j = \frac{\alpha_j(u(1, 1) - u(-1, 1))g_j(x, 1) - (1 - \alpha_j)(u(-1, -1) - u(1, -1))g_j(x, -1)}{\alpha_j g_j(x, 1) + (1 - \alpha_j)g_j(x, -1)}$$

where:

$$\begin{cases} g_j(x, 1) &= \prod_{i:x_i^j=1} p_i^j \prod_{i:x_i^j=0} (1 - p_i^j) \\ g_j(x, -1) &= \prod_{i:x_i^j=1} q_i^j \prod_{i:x_i^j=0} (1 - q_i^j) \end{cases}$$

Example 2.6. Consider $n = 2$ voters facing $m = 3$ projects $\{t_1, t_2, t_3\}$. They have to chose $k = 2$ projects to invest in. Suppose that voters' reliabilities are:

$$(p_1, q_1) = (0.9, 0.4), (p_2, q_2) = (0.7, 0.3)$$

The projects have the following prior probabilities of being good:

$$\alpha_1 = 0.9, \alpha_2 = 0.5, \alpha_3 = 0.4$$

The utilities are:

$$u(1, 1) = 2, u(-1, 1) = -\frac{1}{2}, u(-1, -1) = 1, u(1, -1) = -1$$

Suppose that voter v_1 approves $\{t_2, t_3\}$, whereas voter v_2 only approves $\{t_3\}$. We have:

$$\Delta_1 = 0.59, \Delta_2 = -11.5, \Delta_3 = 2.06$$

Thus the optimal decision is to invest in projects t_1 and t_3 .

Another work where the noise models produce random *approval votes* from a ground truth consisting of a *set of alternatives* is [Caragiannis et al., 2020]. The paper focuses on the robustness of the class of *Approval-based Counting Choice* (ABCC) rules introduced in [Lackner and Skowron, 2021] and described by a bi-variate function f such that $f(x, y)$ is the score that a committee S gets from a ballot of size y containing x elements of S .

A general family of noises is also considered, where the probability of an approval ballot given the ground truth committee is related to their *distance* according to some distance metric d as follows:

$$P(A|S^*) > P(B|S^*) \iff d(A, S^*) < d(B, S^*)$$

where A, B are two approval ballots and S^* is the ground truth subset of alternatives. Such noise is said to be *d-monotonic*. Examples of such distances are the Jaccard, Hamming or Dice distances [Choi et al., 2010]. The idea is that the closer a subset is to the ground truth, the more likely it is to be submitted by a voter. The main result of the work states that the modal rule (outputting the most often selected committee) is the only ABCC rule which is *monotone robust* against all *d-monotonic* noise models for all distance metrics d . This means that, for any such noise model, the probability of uncovering the ground truth via the the modal rule grows to 1 as the number of voters approaches infinity. A characterization of the distances to which SAV is robust is given, then it is proven that it is a necessary and sufficient condition for a rule that $f(k, k) > f(k - 1, k)$ to be monotone robust against all *natural* distances, where a natural distance simply verifies:

$$|A| = |B| \text{ and } |A \cap S| > |B \cap S| \implies d(A, S) \leq d(B, S)$$

Moreover, the paper states that any non-trivial ABCC rule is robust against any *similarity* distance, where a similarity distance is such that:

$$|A| = |B| \text{ and } |A \cap S| > |B \cap S| \implies d(A, S) < d(B, S)$$

2.2.8 Input Profiles are Preferences on Directed Hypercubes

[Xia et al., 2010] study the maximum likelihood approach in multi-issue domains. For a set $\mathcal{I} = \{x_1, \dots, x_p\}$ of p issues, each x_i taking values in a domain D_i , the set of alternatives is $\mathcal{X} = D_1 \times \dots \times D_p$. In the first sections, the paper studies whether issue-by-issue and sequential voting rules can be MLEs for some noise models.

Then it focuses on the special case of binary issues which is related to the problem of multi-winner epistemic approval voting given the common nature of their outputs. A voter's conditional preferences can be represented by a hypercube with directed edges in a p -dimensional space: each vertex represents an alternative, and two adjacent vertices represent two alternatives differing on only one issue. For any issue $i \leq p$, and for any $d_{-i} \in D_{-i}$, the edge is directed from $(0_i, d_{-i})$ to $(1_i, d_{-i})$ if and only if $(0_i, d_{-i})$ is preferred to $(1_i, d_{-i})$.

A distance-based noise model is defined on such preferences: there exist some probabilities q_0, \dots, q_{p-1} such that for any ground truth d^* and for any $a_{-i} \in D_{-i}$ where $|a_{-i} - d_{-i}^*| = k \leq p-1$, the probability of having the local conditional preference $a_{-i} : d_i \succ \bar{d}_i$ is:

$$\pi_{d^*}^{a_{-i}}(d_i \succ \bar{d}_i) = q_k$$

where $|a_{-i} - d_{-i}^*|$ is the Hamming distance separating a_{-i} and d_{-i}^* .

For a subfamily of distance-based noise, called *distance-based threshold noise models*, where there exist some $1 \leq k \leq p$ and $q > \frac{1}{2}$ such that for any $i \leq k$, $q_i = q$ and $q_i = \frac{1}{2}$ for $i > k$, it is shown that the maximum likelihood estimator of the ground truth given a profile P is the alternative with the highest the sum of the consistencies of degree k with all inputs in the profile.

In case of binary issues, each possible outcomes can be assimilated to a subset of alternatives.

2.2.9 Summary

We can partition the above cited papers according to the nature of the input (the votes to be aggregated) and the nature of the output (whether be it the full ranking ground truth, the winner alternative, a set of winners etc..). This partition is shown in Table 2.1.

2.3 Other Directions

2.3.1 The Converse Problem: Which Rules Correspond to MLEs?

In [Conitzer and Sandholm, 2005, Conitzer et al., 2009], the converse path is considered, as instead of fixing a probabilistic model and finding its associated maximum likelihood rule, the authors aim to answer the question of knowing *for which common voting rules there exist a noise model such that the rule is the MLE*. The paper considers a set of m alternatives \mathcal{X} with a ground truth ranking \succ^* and a ground truth winner a^* . A group of n homogeneously reliable voters

¹Only in the case of binary issues, where each outcome can be assimilated to a set of winning alternatives.

Input \ Output	winning alternative	ranking	set of winners
single alternative	[Nitzan and Paroush, 1982] [Ben-Yashar and Nitzan, 1997] [Halpern et al., 2021]		
ranking		[Young, 1986] [Young, 1988] [Young, 1995]	
partial order	[Xia and Conitzer, 2011]		
pairwise comparisons	[Elkind and Shah, 2014]		[Procaccia et al., 2012]
truncated approvals	[Procaccia and Shah, 2015]		
approval ballots			[Ben-Yashar and Paroush, 2001] [Caragiannis et al., 2020]
directed hypercubes			[Xia et al., 2010] ¹

Table 2.1: Optimal rules according to the nature of their input-output

cast orders $(\succ_1, \dots, \succ_n)$. First it is proven that any voting rule is a maximum likelihood rule for some noise model which tolerates inter-dependence among the votes. The paper then restricts the noise distributions to the models with independent voters. Some interesting results are then stated.

Theorem 2.14. *Any scoring rule (e.g. Borda or Plurality) is a maximum likelihood estimation of the winner under i.i.d votes (MLEWIV) and a maximum likelihood estimation rule of the ranking under i.i.d votes (MLERIV).*

Theorem 2.15. *The STV rule is a MLERIV but not MLEWIV rule.*

Theorem 2.16. *The Copeland rule is neither a MLEWIV nor a MLERIV rule.*

2.3.2 Axiomatic Analysis for Collective Annotations

The axiomatic analysis approach usually taken in the social choice literature have been adopted in [Kruger et al., 2014] and applied to the collective annotations framework (use-cases include voters detecting sentiments in a speech, or extracting keywords from a text etc..). The authors consider a group of voters annotating a number of items. Each voter can select a category to each item (or abstain), then an *aggregator* F is applied to try to uncover the correct category for each item. Axioms accounting for different aspects are introduced such that:

- *Decisiveness:* Exactly one category is assigned to each item by F . A weaker version is also presented for cases where no voter annotated the item, in which case the aggregator can simply discard it.
- *Groundedness:* F assigns to an element a category appearing in at least one voter's annotation. This axiom implies *Unanimity*.

- *Independence*: The category of each item is solely determined by its annotations.
- *Monotonicity*: If a category is accepted, than any additional support it gets wouldn't change it.

The plurality rule aggregator is then characterized with some of the previous axioms with additional neutrality and anonymity properties. It is also proven that an aggregator is nontrivial and grounded if and only if it is a *weighted plurality rule*. Some aspects of biased votes were also taken into account with axioms such that *over-use* and *under-use sensitivity*. The idea is that when a voter is seen to assign different categories across the items, it is evidence that she is less biased towards any of the categories and more confidence must be put upon her vote. The authors then define a class of *bias-correcting* aggregators which satisfies these axioms and give some examples. These rules are show to outperform the simple plurality rule on collected telephone calls annotations via Amazon Mechanical Turk. In [Qing et al., 2014], some of these bias-correcting rules are also compared, across different linguistic datasets, to another agreement-based rule whose idea consists in estimating the annotator's reliability via its agreement with the plurality output, and then applying weighted voting accordingly.

2.3.3 Incentives

[Shah et al., 2015] focuses on the incentives aspect of epistemic approval voting. They consider a set of instances (questions) T each of which has $m \geq 2$ possible options. Among the instances there is a subset G of gold standard questions whose ground truth is known to the mechanism designer. The idea is basically to uniformly mix these instances and evaluate the participants according to them. The goal of the paper is to define payment functions that incentivizes the voter to cast exactly the support of her beliefs about each instance (exactly the alternatives she believes might be the correct one), where a voter's belief on question $t \in T$ is a probability distribution over the m options $p_i^t = (p_{i1}^t, \dots, p_{im}^t)$. Its support is simply the set:

$$\{j \leq m, p_{ij}^t \neq 0\}$$

Although the authors begin by stating an impossibility result showing that no payment function can be incentive compatible in the previous sense, they nonetheless restrict the domain of the voters' beliefs and propose an adequate mechanism. Backed by insights from an extensive literature in psychology, the authors assume that the probability of any option for any question, according to the worker's belief, is either zero or greater than a positive threshold ρ : the voters are said to have *coarse beliefs*. Formally, the assumption reads $p_{ij}^t \in \{0\} \cup (\rho, 1]$ for some known and fixed $\rho \in (0, 1/m)$. The paper then defines a payment scheme that is optimal in the sense that there is no other incentive-compatible mechanism that expends as small an amount on a worker who does not attempt any question. The payment is given by the expression:

$$f(x_1, \dots, x_G) = \alpha(1 - \rho)^{\sum_{t=1}^G (x_t - 1)} \prod_{t=1}^G \mathbb{1}\{x_t \geq 1\}$$

where α is the payment due to a perfect answer, and $x_t \in \{-m + 1, \dots, -1, 1, \dots, m\}$ is the number of alternatives selected by a voter, given a negative sign if and only if her ballot does

not contain the correct answer. Notice though that whenever the participant fail to approve the correct alternative in one of the questions, her payment would be zero, regardless of the quality of her answers to the remaining questions.

Interestingly, it was shown that, even for a voter with arbitrary beliefs, the above payment function incentives her to approve the alternatives for which her relative beliefs are above some threshold that depends on ρ .

Theorem 2.17. *Under the payment function f , a voter with beliefs $1 \geq p_1 \geq \dots \geq p_B \geq 0$ is incentivized to approve the alternatives options $1, \dots, k$ where:*

$$k = \arg \max_{z \leq m} \left(\frac{p^z}{\sum_{l=1}^z p^l} > \rho \right)$$

Example 2.7. *Consider a crowdsourcing task with $G = 2$ accessible gold standards. Suppose that we want to incentivize voters to select options one by one in decreasing order of their beliefs as long as the selected option contributes a fraction more than 0.5 to the total belief of the selected options. If we are willing to pay the voter that gets all the answers correct 2 dollars then the optimal payment function is:*

$$f(x_1, x_2) = 2 \times 0.5^{x_1+x_2-2} \times \mathbb{1}\{x_1 \geq 1\} \times \mathbb{1}\{x_2 \geq 1\}$$

Now suppose that the set of alternatives is $\{a, b, c, d\}$, and that for both gold standards, the ground truth is the alternative a . Consider a voter who approves $A^1 = \{a\}$ in the first question, and $A^2 = \{a, b, c\}$ for the second one. Her reward is:

$$f(1, 3) = 1\$$$

2.3.4 The Surprisingly Popular Method

[Prelec et al., 2017] tackles the single-question wisdom of the crowd problem. The main challenge is to achieve high quality aggregated labels even when the majority of the participants are not experts, which is a case when most of the classical methods fail. The idea behind the SP (surprisingly popular) method is to elicit additional information, consisting of the voter's predictions about the remaining voters' answers, and to then select the alternative that actually gets reported more than collectively predicted. SP was tested on four real datasets of Yes/No questions related respectively to state capitals, general knowledge, lesion assessing by dermatologists and art judging. It was proven to significantly outperform other baselines such as confidence-weighted majority voting.

Recently [Hosseini et al., 2021] extend the *surprisingly popular* method to the rankings framework, with an adapted elicitation protocol. The method needs to not only gather each voter's ranking, but also her prediction on the remaining voters' answers. The paper doesn't require voters to submit a probability distribution over all rankings (which would require $m!$ entries) but only the ranking most frequently submitted by the remaining participants. This method is shown to outperform conventional aggregation rules such as Borda, Plurality and Copeland on crowdsourced datasets.

2.3.5 Judgment Aggregation

[Pigozzi and Hartmann, 2007] consider a set of n agents N making their judgments on a given set of finite set of logical propositions. They focus on the case $(P \wedge Q \iff R)$. They posit a probabilistic model with the following assumptions:

- The prior probabilities that P and Q are true are equal to some $q \in (0, 1)$.
- The voters have a common competence $p \in (0, 1)$ to assess the truth of P and Q .
- P and Q are logically and probabilistically independent.
- The voters' judgments are logically consistent.

Hence the set of alternatives is $\mathcal{X} = \{S_1, S_2, S_3, S_4\}$, where:

$$\begin{cases} S_1 = (1, 1, 1) = \{P, Q, R\} \\ S_2 = (1, 0, 0) = \{P, \neg Q, \neg R\} \\ S_3 = (0, 1, 0) = \{\neg P, Q, \neg R\} \\ S_4 = (0, 0, 0) = \{\neg P, \neg Q, \neg R\} \end{cases}$$

The paper then compares the truth tracking capabilities of three merging methods:

- *Premise-based procedure*: Each voter votes for P and Q (excluding her conclusion R). R is deduced for each vote and the majority decision is output.
- *Conclusion-based procedure*: Each voter privately decides on P and Q and then submits her judgment on the conclusion R .
- *Fusion procedure*: The majority fusion operator selects the model that minimizes the Hamming distance to the individual bases.

It was shown that the fusion procedure outperforms the two other operators for low and mid-dling competence values p , and that the premise-based method is the best truth-tracker for high competence voters.

[Terzopoulou and Endriss, 2019] considers the problem of uncovering the truth of a binary conclusion depending on two independent premises (with uniform prior) from noised judgments of some agents, either on one or both of the premises. The authors suppose that the precision of an agent decreases when multi-tasking. The optimal rule in such a context is a weighted majority rule where the weight of voter only depends on whether she answers one or two questions. The strategic and elicitation aspects are then considered.

Part II

Contributions

Introduction

The next three chapters will present our contributions to the epistemic voting literature.

In Chapter 3, we consider a single-winner approval voting setting, where a set N of n voters approve subsets of a set of $m \geq 2$ alternatives $\mathcal{X} = \{a_1 \dots, a_m\}$, one of which is a hidden ground truth a^* .

Our main goal is to propose a one-shot estimation method, where the mere observation of the voters' answers to one question (consisting in approval ballots) suffices to estimate their heterogeneous reliabilities which are not priorly known. The voters are then weighted accordingly in an optimal fashion and their votes are aggregated via the corresponding weighted approval voting rule.

In order to estimate a voter's expertise given her ballot, we make the assumption that *reliable voters tend to cast smaller ballots* (i.e they approve a smaller number of alternatives). We prove that this hypothesis holds theoretically under a large family of usual noise models. We also assess its efficiency in recovering the ground truth compared to a baseline aggregator for multiple crowdsourced image annotation datasets.

In Chapter 4 and Chapter 5, we move to a more general and less studied problem, namely, multi-winner epistemic approval voting. We begin by transferring the notion of *committee-size* from the classical to the epistemic voting framework. Though in classical voting, it is usually expressed as a integer k representing the number of alternatives to be elected, we define as an inequality constraint stating that the output's size must be greater than some given lower-bound l and less than a given upper-bound u ($l = u = k$ is a particular case).

Whilst in single-winner epistemic voting, the interpretation of the ground truth is clear: there is one correct alternative — the only one that is true in the real world — and the aim is to identify it. In multi-winner voting, things become more complicated.

Recall that the input of a multi-winner voting instance consists of a set of votes and a constraint on the committee size. Now, two different interpretations of this constraint on committee size coexist, which call for different solutions. The main difference between both interpretations is whether the constraints on the cardinality of the set of alternatives to be identified bear *on the ground truth itself*, or *on the output*.

Under the first interpretation, for which Chapter 4 is devoted, this constraint represents some *prior knowledge on the cardinality of the ground truth*, that is, on the number of winning alternatives. Here are two examples:

- Consider a guitar chord transcription task: participants hear a chord and are asked to select the set of notes that constitute it from a total of 8 alternatives. We know *a priori* that the true set of notes is made of at least 3 and at most 6 alternatives.
- Consider a crowdsourcing problem where participants are showed a picture of a football match, where they have to identify the team(s) appearing in it. Here we know that there must be either one or two teams in the ground truth.

Under the second interpretation, studied in Chapter 5, this constraint bears on the *number of winners in the output*. That is, whatever the ground truth is, we have to output a number of alternatives in a given interval, *even though the number of alternative in the ground truth may lie outside this interval*; the aim is to identify an *admissible* set of alternatives *closest to the ground truth*, in some sense to be defined later. Here are three examples:

- alternatives are students who apply to a master program. The ground truth consists of those students who have objectively a good enough level to graduate. The number of students to be accepted in the program is however constrained to be in an interval $[l, u]$: we need at least l for the program to open, and at most u because of the size of classrooms.
- alternatives are papers submitted to a conference. Again we have a minimal and maximal number of papers to be accepted.
- alternatives are papers to be given an award. The conference chair give exactly three awards. In that case the ground truth consists of the papers that truly deserve the award, and the output consists of the best three papers. A variant of the problem would allow the conference chairs to give *at most* three awards, so that they avoid giving an award to papers that do not deserve it.
- Alternatives are Covid-19 patients in urgent need of intensive care; there is a limited number of intensive care units.

Under this interpretation, committee size plays the role of an exogenous size constraint that specifies the minimum and maximum number of allowed winning alternatives. It remains to define precisely what we mean by being closest to the ground truth; we will propose two different solution concepts.

Although this distinction between two interpretations would already make sense in single-winner epistemic voting, we will show that in this special case they lead to the same solution: they are technically identical. This will however not be the case in the general case; therefore we shall develop solutions for each of these two interpretations separately.

Single-winner Truth-Tracking

Abstract

Epistemic social choice aims at unveiling a hidden ground truth given votes, which are interpreted as noisy signals about it. We consider here a simple setting where votes consist of approval ballots: each voter approves a set of alternatives which they believe can possibly be the ground truth. Based on the intuitive idea that more reliable votes contain fewer alternatives, we define several noise models that are approval voting variants of the Mallows model. The likelihood-maximizing alternative is then characterized as the winner of a weighted approval rule, where the weight of a ballot decreases with its cardinality. We have conducted an experiment on three image annotation datasets; they conclude that rules based on our noise model outperform standard approval voting; the best performance is obtained by a variant of the Condorcet noise model.

Most of the results and notions introduced in this chapter were published in [Allouche et al., 2022b].

3.1 Introduction

Epistemic social choice deals with the problem of unveiling a hidden *ground truth* state from a set of some possible states, given the reports of some voters. Votes are seen as *noisy* reports on the ground truth. The distribution of these reports is modelled by a noise model, sometimes tuned by some parameter reflecting the competence (expertise, reliability) of these voters.

The space of frameworks for epistemic social choice varies along two dimensions: the nature of the ground truth and the format of the reports (ballots expressed by voters). Depending on the framework chosen, the ground truth may be a single alternative, a set of alternative or a ranking over alternatives. *We assume the simplest form of ground truth: it is a single alternative* (the unique correct answer). Still depending on the framework, ballots may also contain a single alternative, a set of alternatives, or a ranking over alternatives. *We assume that they are subsets of*

alternatives, that is, approval ballots. Requiring voters to give only one answer (that is, a single alternative) is often too constraining because voters may be uncertain and believe that several alternatives may possibly be the ground truth. This is the path followed by [Procaccia and Shah, 2015, Shah et al., 2015, Caragiannis and Micha, 2017].

Beyond social choice, collective annotation has also been studied in the machine learning community. [Dawid and Skene, 1979] used an expectation-maximization (EM) approach for retrieving true binary labels. This approach has been improved along with other methods namely in [Raykar et al., 2010, Welinder et al., 2010, Bonald and Combes, 2017, Tao et al., 2018].

While some classical anonymous rules have been shown to be optimal under some assumptions, the aggregation rule may, when possible, assign different weights to the voters according to their expertise. Whilst this is doable if we have additional information about voter expertise or when we keep a record of their answers to different questions, estimating the individual expertise gets complicated when we have no prior information about voters and when the sole information we have are votes on a single issue. This leads to the *single-question wisdom of the crowd problem* for which the seminal work [Prelec et al., 2017] proposes a novel solution, namely selecting the alternative which is *surprisingly popular*. Although it proved to be an efficient way to get around the problem of estimating the voters' reliabilities, its major drawback is that it requires the elicitation of further information: each voter has to report her answer and her beliefs about the answers of the remaining voters.

Now we suggest that there is an alternative approach that does not require this surplus of information and that simply relies on the truthfulness of voters.

[Shah et al., 2015] have defined a proper mechanism to incentivize the participants to select an alternative if and only if they believe it can be the winning one. An intuitive idea might be to consider that *smaller ballots, i.e. answers that contain less alternatives, are more reliable*: a voter who knows the true answer (or, more generally, who believes to know it) will probably select only one alternative and a voter who selects all alternatives has probably no idea whatsoever of the correct answer. For instance, if voters hear a speech and are asked to detect the language in which it is spoken, we may give more weight to a voter approving *Arabic* and *Hebrew* than to one approving *Arabic, Hebrew, Persian* and *Turkish*.

Based on this intuition, more weight must be assigned to smaller ballots. Rules that work this way, which we call *size-decreasing approval rules*, are part of the family of *size approval rules* [Alcalde-Unzu and Vorsatz, 2009]. Our goal is to motivate the use of such rules from an epistemic social choice point of view. To this purpose, we will study a family of noise models which are approval voting variants of the Mallows model, and prove that in many cases the optimal rule is *size-decreasing*.

The chapter is structured as follows. In Section 3.2 we define the framework and the family of noise models we consider. Section 3.3 characterizes all anonymous noises whose associated optimal rule is size-decreasing. In Section 3.4, we consider a more general noise where voters have different noise parameters, prove that under some mild assumption, *the expected number of selected alternatives grows when the voter is less reliable*, and then give an explicit expression

for the expected size of the ballot as a function of the reliability parameter of a voter for a Condorcet-like noise model. Section 3.5 focuses on real datasets on which first we test the hypothesis that smaller ballots are more reliable then we apply different size-decreasing rules associated to various noise models to assess their performances. Section 3.6 concludes.

3.2 Framework

Consider a set N of n voters and a set of $m \geq 2$ alternatives $\mathcal{X} = \{a_1 \dots, a_m\}$.

The (hidden) ground truth consists of a single alternative a^* . Voters cast approval ballots $A_i \subseteq \mathcal{X}$ consisting of their noisy estimates of the ground truth. Voters who approve no alternative or all alternatives do not bring any meaningful information, therefore without loss of generality, we assume that for all i , $A_i \neq \emptyset$ and $A_i \neq \mathcal{X}$.

All along this paper, we will model the distribution of these approval ballots by approval voting variants of the *Mallows* noise model. The Mallows distribution was originally defined on *rankings*: we adapt it to subsets of alternatives, keeping the idea that the probability of a subset decreases as its distance from a central point increases, the dispersion being modelled by a parameter ϕ .

In general, we will call an approval Mallows noise model any model where voters' ballots are independent (we keep this hypothesis all along the paper) and there exist n parameters $\phi_i \in (0, +\infty)$ and a function $d : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \mapsto \mathbb{R}$ such that and for any voter $i \in N$:

$$P_{\phi_i, d}(A_i | a^* = a) = \frac{1}{Z_i} \phi_i^{d(a^*, A_i)}, \forall a \in \mathcal{X}$$

where Z_i is the corresponding normalization factor. If $\phi_i = \phi$ for all $i \in N$, we say the model is *anonymous*.

In the remaining of the paper we will only focus on *neutral* noise models. The neutrality of a noise is defined as its invariance by any permutation of the alternatives:

$$\forall \pi \in \sigma(\mathcal{X}), P_{\phi, d}(A | a^* = a) = P_{\phi, d}(\pi(A) | a^* = \pi(a))$$

We can immediately see that a Mallows noise is neutral if and only if its associated function d is neutral (invariant by a permutation of the alternatives).

A noise model is neutral if $d(a, A)$ depends only on $|a \cap A|$ (that is, 1 if $a \in A$ and 0 if $a \notin A$)¹ and $|A|$:

Proposition 3.1. *A noise model associated to a function d is neutral if and only if there exists a unique function²:*

$$\psi_d : \{0, 1\} \times \{0, \dots, m\} \setminus (1, 0) \mapsto \mathbb{R}$$

¹We omit the curly brackets and write $a \cap A$ for $\{a\} \cap A$.

² $(1, 0)$ is excluded from the domain of ψ_d simply because $(|A| = 0 \text{ and } |a \cap A| = 1)$ is impossible.

such that:

$$d(a, A) = \psi_d(|a \cap A|, |A|)$$

Proof. If $d(a, A) = \psi_d(|a \cap A|, |A|)$ then since ψ_d is neutral, d is neutral. Conversely, assume d is neutral. We claim that for any two pairs (a, A) and (b, B) :

$$(|a \cap A|, |A|) = (|b \cap B|, |B|) \implies d(a, A) = d(b, B)$$

Assume first $a \in A$ (and therefore, $b \in B$). Consider a permutation π such that $\pi(a) = b$ and $\pi(A) = B$. Then:

$$d(b, B) = d(\pi(a), \pi(A)) = d(a, A)$$

The argument for the case $a \notin A$ (and $b \notin B$) is similar. Thus, $d(a, A)$ depends only on $|a \cap A|$ and $|A|$, which means that there is a function:

$$\psi_d : \{0, 1\} \times \{0, \dots, m\} \setminus (1, 0) \mapsto \mathbb{R}$$

such that $d(a, A) = \psi_d(|a \cap A|, |A|)$.

Uniqueness is immediate. □

Example 3.1. For the Hamming distance we have that:

$$d_H(a, A) = |\bar{a} \cap A| + |a \cap \bar{A}| = 1 - 2|a \cap A| + |A|$$

so $d_H(a, A) = \psi_{d_H}(|a \cap A|, |A|)$, where:

$$\psi_{d_H}(t, j) = 1 - 2t + j$$

$d_H(a, A)$ takes its minimal value 0 for $A = \{a\}$ and its maximal value m for $A = \mathcal{X} \setminus \{a\}$.

Example 3.2. For the Jaccard distance [Jaccard, 1901] we have:

$$d_J(a, A) = \frac{|\bar{a} \cap A| + |a \cap \bar{A}|}{|\bar{a} \cap A| + |a \cap \bar{A}| + |a \cap A|}$$

so $d_J(a, A) = \psi_{d_J}(|a \cap A|, |A|)$ where:

$$\psi_{d_J}(t, j) = 1 - \frac{t}{1 - t + j}$$

$d_J(a, A)$ takes (again) its minimal value 0 for $A = \{a\}$ and its maximal value 1 for $A = \mathcal{X} \setminus \{a\}$.

3.3 Anonymous Noise and Size-decreasing Approval Rules

In this section, we suppose that voters share a common (unknown) noise parameter $\phi \in (0, 1)$ and that there exists some function $d : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \mapsto \mathbb{R}$ and its associated function ψ_d such that, for any $a \in \mathcal{X}$:

$$P_{\phi,d}(A_i | a^* = a) = \frac{1}{Z} \phi^{d(a^*, A_i)} = \frac{1}{Z} \phi^{\psi_d(|a^* \cap A_i|, |A_i|)}$$

After formally defining the notion of size-decreasing rules, we state the main result of this section which characterizes all the Mallows anonymous noises (that is, all functions d) whose associated maximum likelihood rule is size-decreasing. We will see that this is the case for some natural functions d , that we will test later on in the experiments.

Definition 3.1 (Size Approval Rule). *Consider a function*

$$\begin{aligned} v : \quad \mathcal{P}(\mathcal{X})^n &\longrightarrow \mathcal{X} \\ (A_1, \dots, A_n) &\mapsto v(A_1, \dots, A_n) \end{aligned}$$

that, for each approval profile $A = (A_1, \dots, A_n)$, assigns a winning alternative $v(A_1, \dots, A_n) \in \mathcal{X}$. We say that v is a size approval rule if there exists a vector $w = (w_0, \dots, w_m) \in \mathbb{R}^{m+1}$ such that:

$$v(A_1, \dots, A_n) = \arg \max_{a \in \mathcal{X}} \text{app}_w(a)$$

where app_w is the weighted approval score defined by:

$$\text{app}_w(a) = \sum_{i: a \in A_i} w_{|A_i|}$$

A size approval rule v is size-decreasing if its associated vector $w = (w_0, \dots, w_m) \in \mathbb{R}^{m+1}$ is such that $w_j > w_{j+1}$ for all $1 \leq j \leq m - 2$.

Example 3.3. *The size approval rule associated to the vector of weights given by $w_j = n^{m-j}$ is size-decreasing in the most extreme sense, as it is lexicographic: it outputs the alternative which appears most often in singleton ballots, in case of ties it considers ballots of size 2 and so on.*

Definition 3.2 (Maximum Likelihood Estimation Rule). *We define the function:*

$$\begin{aligned} \zeta_d : \quad \mathcal{P}(\mathcal{X})^n &\longrightarrow \mathcal{X} \\ (A_i)_{i \in N} &\mapsto \arg \max_{a \in \mathcal{X}} P_d(A_1, \dots, A_n | a^* = a) \end{aligned}$$

which given an approval profile outputs the maximum likelihood estimator of the ground truth alternative.

The next theorem aims to characterize the functions d for which the maximum likelihood estimation rule ζ_d is a size-decreasing approval rule.

Theorem 3.2. For $n \geq 3$, the maximum likelihood estimation rule ζ_d is a size-decreasing approval rule if and only if the function

$$\Delta\psi_d : j \mapsto \psi_d(0, j) - \psi_d(1, j)$$

is decreasing.

Proof. First, for any approval profile $A = (A_1, \dots, A_n)$,

$$\begin{aligned} \zeta_d(A) &= \arg \max_{a \in \mathcal{X}} \prod_{i=1}^n \frac{1}{Z} \phi^{d(a, A_i)} = \arg \max_{a \in \mathcal{X}} \phi^{\sum_{i=1}^n d(a, A_i)} \\ &= \arg \min_{a \in \mathcal{X}} \sum_{i=1}^n d(a, A_i) \\ &= \arg \min_{a \in \mathcal{X}} \sum_{i=1}^n \psi_d(|a \cap A_i|, |A_i|) \\ &= \arg \min_{a \in \mathcal{X}} \sum_{i: a \in A_i} \psi_d(1, |A_i|) + \sum_{i: a \notin A_i} \psi_d(0, |A_i|) \\ &= \arg \min_{a \in \mathcal{X}} \underbrace{\sum_{i=1}^n \psi_d(0, |A_i|)}_{\text{constant}} - \underbrace{\sum_{i: a \in A_i} \psi_d(0, |A_i|) - \psi_d(1, |A_i|)}_{\Delta\psi_d(|A_i|)} \\ &= \arg \max_{a \in \mathcal{X}} \sum_{i: a \in A_i} \psi_d(0, |A_i|) - \psi_d(1, |A_i|) \\ &= \arg \max_{a \in \mathcal{X}} \sum_{i: a \in A_i} \Delta\psi_d(|A_i|) \end{aligned}$$

\Leftarrow : If $\Delta\psi_d$ is decreasing then we can immediately prove that ζ_d is a size-decreasing approval rule with a weight vector w such that $w_j = \Delta\psi_d(j)$ for any $1 \leq j \leq m - 1$.

\Rightarrow : Suppose that ζ_d is a size-decreasing approval rule. Thus, there exists a weight vector $w = (w_0, \dots, w_m) \in \mathbb{R}^{m+1}$ such that $w_j > w_{j+1}$ for $1 \leq j \leq m - 2$ and for any approval profile $A = (A_1, \dots, A_n)$ we have:

$$\zeta_d(A) = \arg \max_{a \in \mathcal{X}} \sum_{i: a \in A_i} w_{|A_i|}$$

If $m = 3$: Let $\mathcal{X} = \{a, b, c\}$. To prove that $\Delta\psi_d(1) > \Delta\psi_d(2)$ consider the following approval profile:

$$\begin{cases} A_1 &= \{a\} \\ A_i &= \{a, b\} \quad , \forall i \in \{2, \dots, n-1\} \\ A_n &= \{b, c\} \end{cases}$$

Which yields the following weighted approval scores:

$$\begin{cases} app_w(a) &= w_1 + (n-2)w_2 \\ app_w(b) &= (n-1)w_2 \\ app_w(c) &= w_2 \end{cases}$$

Since $w_1 > w_2$ we have $\zeta_d(A) = a$, which implies that $\arg \max_{e \in \mathcal{X}} \sum_{i: e \in A_i} \Delta\psi_d(|A_i|) = a$. In particular:

$$\sum_{i: a \in A_i} \Delta\psi_d(|A_i|) > \sum_{i: b \in A_i} \Delta\psi_d(|A_i|)$$

So:

$$\Delta\psi_d(1) + (n-2)\Delta\psi_d(2) > (n-1)\Delta\psi_d(2)$$

which implies that $\Delta\psi_d(1) > \Delta\psi_d(2)$.

If $m > 3$: Let $\mathcal{X} = \{a, b, c, e_1, \dots, e_{m-3}\}$. To prove that $\Delta\psi_d(1) > \Delta\psi_d(2)$ we use the same approval profile as above. To prove that $\Delta\psi_d(j) > \Delta\psi_d(j+1)$ for $j \geq 2$, consider the following approval profile:

$$\begin{cases} A_1 = \{a, e_1, \dots, e_{j-1}\} \\ A_i = \{a, b\} \\ A_n = \{b, c, e_1, \dots, e_{j-1}\} \end{cases}, \forall i \in \{2, \dots, n-1\}$$

Which yields the following weighted approval scores:

$$\begin{cases} app_w(a) = w_j + (n-2)w_2 \\ app_w(b) = w_{j+1} + (n-2)w_2 \\ app_w(c) = w_{j+1} \\ app_w(e_l) = w_j + w_{j+1} \end{cases} \quad l \in \{1, \dots, j-1\}$$

Since $w_2 > w_j > w_{j+1}$ we have $\zeta_d(A) = a$, which implies that $\arg \max_{e \in \mathcal{X}} \sum_{i: e \in A_i} \Delta\psi_d(|A_i|) = a$. In particular, we have

$$\sum_{i: a \in A_i} \Delta\psi_d(|A_i|) > \sum_{i: b \in A_i} \Delta\psi_d(|A_i|)$$

So:

$$\Delta\psi_d(j) + (n-2)\Delta\psi_d(2) > \Delta\psi_d(j+1) + (n-2)\Delta\psi_d(2)$$

which implies that:

$$\Delta\psi_d(j) > \Delta\psi_d(j+1)$$

□

Interpretation: Consider an anonymous noise $P_{\phi, d}$, where d is such that $\Delta\psi_d$ is decreasing. Now consider any alternative $a \in \mathcal{X}$, and for any $k \in [1, m-2]$, let $A_k, A_{k+1}, B_k, B_{k+1}$ be four sets such that $a \in A_k \cap A_{k+1}$ and $a \notin B_k \cup B_{k+1}$ and $|A_k| = |B_k| = k$ and $|A_{k+1}| = |B_{k+1}| = k+1$. We can easily check that since $\phi \in (0, 1)$ and $d(a, B_k) - d(a, A_k) < d(a, B_{k+1}) - d(a, A_{k+1})$, we have

$$\frac{P_{\phi, d}(B_k | a^* = a)}{P_{\phi, d}(A_k | a^* = a)} < \frac{P_{\phi, d}(B_{k+1} | a^* = a)}{P_{\phi, d}(A_{k+1} | a^* = a)}$$

which implies the following:

- If it is more likely that a voter casts a k -sized ballot not containing the ground truth than a k -sized ballot that contains it, then it is even more likely that she casts a $(k+1)$ -sized ballot not containing the ground truth than a $(k+1)$ -sized ballot that contains it.

- If it is more likely that a voter casts a $(k + 1)$ -sized ballot containing the ground truth than a $(k + 1)$ -sized ballot that does not, then it is even more likely that she casts a k -sized ballot containing the ground truth than a k -sized ballot that does not.

We now give some examples with some usual functions d . We will see that the maximum likelihood estimation rule associated to the Jaccard distance is size-decreasing with weights $w_{|A|} = \frac{1}{|A|}$, and that the maximum likelihood estimation rule associated to the Hamming distance is not size-decreasing.

Example 3.4. Consider the Jaccard distance given by:

$$d_J(a, A) = \psi_{d_J}(|a \cap A|, |A|) = 1 - \frac{|a \cap A|}{|A| - |a \cap A| + 1}$$

which gives:

$$\Delta\psi_{d_J}(j) = \psi_{d_J}(0, j) - \psi_{d_J}(1, j) = 1/j$$

By Theorem 3.2, we conclude that the maximum likelihood estimation rule ζ_{d_J} is a size-decreasing approval rule with weights $w_j = 1/j$.

Example 3.5. Consider a strictly concave function $g : \mathbb{R}^+ \rightarrow \mathbb{R}$ and let:

$$d(a, A) = g(d_{Ham}(a, A))$$

be a concave transformation of the Hamming distance $d_{Ham}(a, A) = |a \cap \bar{A}| + |\bar{a} \cap A|$. Then ζ_d is a size approval rule with weights $w_j = g(j + 1) - g(j - 1)$.

By Theorem 3.2 it suffices to prove that $\Delta\psi_d$ is decreasing. Let $1 \leq x \leq m - 2$. We have $\psi_d(0, x) = g(x + 1)$ and $\psi_d(1, x) = g(x - 1)$. Now let $y > x$. We have

$$\begin{aligned} \frac{\Delta\psi_d(y) - \Delta\psi_d(x)}{y - x} &= \frac{g(y + 1) - g(x + 1)}{y - x} - \frac{g(y - 1) - g(x - 1)}{y - x} \\ &= \tau_{x+1}^g(y + 1) - \tau_{x-1}^g(y - 1) \\ &= \tau_{x+1}^g(y + 1) - \tau_{x+1}^g(y - 1) + \tau_{x+1}^g(y - 1) - \tau_{x-1}^g(y - 1) \\ &= \underbrace{\tau_{x+1}^g(y + 1) - \tau_{x+1}^g(y - 1)}_{<0} + \underbrace{\tau_{x+1}^g(y - 1) - \tau_{x-1}^g(y - 1)}_{<0} \end{aligned}$$

where:

$$\tau_x^g(y) = \frac{g(y) - g(x)}{y - x}$$

is the change rate of g at x , evaluated in y . For instance, consider the Euclid distance $d_{Euc}(a, A) = \sqrt{d_H(a, A)}$. We have that the associated maximum likelihood estimation rule ζ_d is a size approval rule with weights:

$$w_j = \sqrt{j + 1} - \sqrt{j - 1}$$

Example 3.6. Consider the Hamming distance given by:

$$d_H(a, A) = |a \cap \bar{A}| + |\bar{a} \cap A| = 1 + |A| - 2|a \cap A|$$

Which gives us that:

$$\Delta\psi_{d_H}(j) = \psi_{d_H}(0, j) - \psi_{d_H}(1, j) = 2$$

Therefore, the maximum likelihood estimation rule ζ_{d_H} is a size approval rule with constant weights: it is the standard approval rule (SAV), that selects the alternative with the maximum number of approvals. It can be seen immediately that SAV is not size-decreasing; however, it is, so to say, size-non-increasing, and thus can be seen as the limit of size-decreasing rules.

As a consequence of Theorem 3.2, we can easily prove that, for an anonymous noise, the maximum likelihood estimation rule associated to a function d defined as a linear combination of the quantities $|a \cap A|$ and $|A|$ is not a size-decreasing rule (this is the case for the Hamming distance). More generally, this applies to any function d such that ψ_d can be additively separated into two terms $\psi_d(|a \cap A|, |A|) = f(|a \cap A|) + g(|A|)$. In the next section, we will consider this particular family of separable functions with a non-anonymous noise, where each voter has her own noise parameter ϕ_i .

3.4 Non-anonymous Separable Noise

3.4.1 The General Case

Consider a set of m alternatives $\mathcal{X} = \{a_1, \dots, a_m\}$ and a ground truth answer a^* . Consider also a neutral function $d : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ with an associated function $\psi_d : \{0, 1\} \times \{0, \dots, m\} \setminus (1, 0) \mapsto \mathbb{R}$ which can be separated into two quantities:

$$\psi_d(|a \cap A|, |A|) = f(|a \cap A|) + g(|A|)$$

We define a non-anonymous Mallows noise model, where for each voter $i \in N$ there exists a parameter $\phi_i \in]0, +\infty[$ such that, for any $a \in \mathcal{X}$:³

$$P_{\phi_i, d}(A_i | a^* = a) = \frac{1}{Z_i} \phi_i^{f(|a^* \cap A_i|) + g(|A_i|)}$$

Notice that in this case, a bigger individual noise parameter ϕ_i models a less reliable voter (her distribution is less condensed around the ground truth). The aim of the next result is to motivate the use of size-decreasing approval rules to aggregate approvals generated from such distributions. More precisely, the goal is to find sufficient conditions on f and g that makes the expected size of the voter's ballot $\mathbb{E}_{\phi, d}[|A_i|]$ grow as the voter becomes less reliable (*i.e.* as her noise parameter ϕ_i grows).

We will denote $f(1) = f_1, f(0) = f_0$ and $\Delta f = f_0 - f_1$ which would naturally be positive $\Delta f > 0$. We will also denote by $\Delta g_{k,t} = g(k) - g(t)$.

³Recall that voters cannot cast empty or full approval ballots. Therefore we suppose that $P(\emptyset | a^* = a) = P(\mathcal{X} | a^* = a) = 0$.

Theorem 3.3. *If for every $1 \leq t < k \leq m - 1$ we have that:*

$$\Delta g_{k,t} \geq \frac{k-t}{2} \Delta f$$

Then:

$$\frac{\partial \mathbb{E}_{\phi,d}[|A_i|]}{\partial \phi} \geq 0$$

Proof. Consider a voter in N with a noise distribution:

$$P_{\phi,d}(A|a^* = a) = \frac{1}{Z} \phi_i^{f(|a^* \cap A|) + g(|A|)}$$

For any prior distribution on a^* , it suffices to prove that for any $a \in \mathcal{X}$ we have that

$$\frac{\partial \mathbb{E}_{\phi,d}[|A_i| | a^* = a]}{\partial \phi} \geq 0$$

since

$$\mathbb{E}_{\phi,d}[|A_i|] = \sum_{a \in \mathcal{X}} \mathbb{E}_{\phi,d}[|A_i| | a_* = a] P(a^* = a)$$

Let $a \in \mathcal{X}$. First let us compute Z :

$$\begin{aligned} Z &= \sum_{A \notin \{\emptyset, \mathcal{X}\}} \tilde{P}(A | a_* = a) = \sum_{t=1}^{m-1} \tilde{P}(|A| = t | a_* = a) \\ &= \sum_{t=1}^{m-1} \sum_{\substack{a_* \in A \\ |A|=t}} \tilde{P}(A | a_* = a) + \sum_{\substack{a_* \notin A \\ |A|=t}} \tilde{P}(A | a_* = a) \\ &= \sum_{t=1}^{m-1} \binom{m-1}{t-1} \phi^{f_1+g(t)} + \binom{m-1}{t} \phi^{f_0+g(t)} \\ &= \sum_{t=1}^{m-1} h_t(\phi) \end{aligned}$$

where

$$h_t(\phi) = \binom{m-1}{t-1} \phi^{f_1+g(t)} + \binom{m-1}{t} \phi^{f_0+g(t)}$$

We have that:

$$\mathbb{E}_{\phi,d}[|A_i| | a_* = a] = \sum_{t=1}^{m-1} t P(|A_i| = t | a_* = a) = \sum_{t=1}^{m-1} \left[t \times \frac{h_t(\phi)}{\sum_{k=1}^{m-1} h_k(\phi)} \right]$$

Thus we have that:

$$\begin{aligned}
 \frac{\partial \mathbb{E}_{\phi,d} [|A_i| | a_* = a]}{\partial \phi} &= \frac{\sum_{t=1}^{m-1} t h'_t(\phi) \sum_{k=1}^{m-1} h_k(\phi) - \sum_{t=1}^{m-1} t h_t(\phi) \sum_{k=1}^{m-1} h'_k(\phi)}{\underbrace{\left(\sum_{k=1}^{m-1} h_k(\phi) \right)^2}_{>0}} \\
 &\propto \sum_{t=1}^{m-1} \sum_{k=1}^{m-1} t h'_t(\phi) h_k(\phi) - \sum_{t=1}^{m-1} \sum_{k=1}^{m-1} t h_t(\phi) h'_k(\phi) \\
 &\propto \sum_{t=1}^{m-1} \sum_{k=1}^{m-1} t h'_t(\phi) h_k(\phi) - \sum_{t=1}^{m-1} \sum_{k=1}^{m-1} k h_k(\phi) h'_t(\phi) \\
 &\propto \sum_{t=1}^{m-1} \sum_{k=1}^{m-1} (t - k) h'_t(\phi) h_k(\phi) \\
 &\propto \sum_{t=1}^{m-2} \sum_{k=t+1}^{m-1} \underbrace{(k - t)}_{>0} \underbrace{[h'_k(\phi) h_t(\phi) - h'_t(\phi) h_k(\phi)]}_{\Delta h_{k,t}(\phi)}
 \end{aligned}$$

We can already notice that to guarantee that $\frac{\partial \mathbb{E}_{\phi,d} [|A_i| | a_* = a]}{\partial \phi} \geq 0$ it suffices that:

$$\Delta h_{k,t}(\phi) = [h'_k(\phi) h_t(\phi) - h'_t(\phi) h_k(\phi)] \geq 0, \forall 1 \leq t < k \leq m - 1$$

We have that:

$$\begin{aligned}
 h'_k(\phi) h_t(\phi) &= \left[(g(k) + f_1) \binom{m-1}{k-1} \phi^{g(k)+f_1-1} + (g(k) + f_0) \binom{m-1}{k} \phi^{g(k)+f_0-1} \right] \\
 &\quad \times \left[\binom{m-1}{t-1} \phi^{f_1+g(t)} + \binom{m-1}{t} \phi^{f_0+g(t)} \right] \\
 &= (g(k) + f_1) \binom{m-1}{k-1} \binom{m-1}{t-1} \phi^{g(k)+g(t)+2f_1-1} \\
 &\quad + (g(k) + f_1) \binom{m-1}{k-1} \binom{m-1}{t} \phi^{g(k)+g(t)+f_1+f_0-1} \\
 &\quad + (g(k) + f_0) \binom{m-1}{k} \binom{m-1}{t} \phi^{g(k)+g(t)+2f_0-1} \\
 &\quad + (g(k) + f_0) \binom{m-1}{k} \binom{m-1}{t-1} \phi^{g(k)+g(t)+f_1+f_0-1}
 \end{aligned}$$

So we have that:

$$\Delta h_{k,t}(\phi) = \binom{m-1}{k-1} \binom{m-1}{t-1} \phi^{g(k)+g(t)+2f_1-1} [g(k) - g(t)]$$

$$\begin{aligned}
 & + \binom{m-1}{k} \binom{m-1}{t} \phi^{g(k)+g(t)+2f_0-1} [g(k) - g(t)] \\
 & + \binom{m-1}{k} \binom{m-1}{t-1} \phi^{g(k)+g(t)+f_0+f_1-1} [g(k) - g(t) + f_0 - f_1] \\
 & + \binom{m-1}{k-1} \binom{m-1}{t} \phi^{g(k)+g(t)+f_0+f_1-1} [g(k) - g(t) - f_0 + f_1] \\
 & = \underbrace{\phi^{g(k)+g(t)+f_0+f_1-1} \binom{m-1}{k-1} \binom{m-1}{t-1}}_{>0} \times \left[\phi^{-\Delta f} \Delta g_{k,t} + \frac{m-k}{k} \frac{m-t}{t} \phi^{\Delta f} \Delta g_{k,t} \right. \\
 & \left. + \left(\frac{m-k}{k} + \frac{m-t}{t} \right) \Delta g_{k,t} + \left(\frac{m-k}{k} - \frac{m-t}{t} \right) \Delta f \right] \\
 & \propto \left[r(\phi) + \frac{m-k}{k} + \frac{m-t}{t} \right] \Delta g_{k,t} - \left(\frac{m-t}{t} - \frac{m-k}{k} \right) \Delta f
 \end{aligned}$$

where:

$$r(\phi) = \phi^{-\Delta f} + \frac{m-k}{k} \frac{m-t}{t} \phi^{\Delta f}$$

We show that:

$$r(\phi) \geq 2\sqrt{\frac{m-k}{k}} \sqrt{\frac{m-t}{t}}$$

To do so, we study the variations of the function r and its derivative. In fact, denote by $\alpha = \frac{m-k}{k} \frac{m-t}{t}$. We have that:

$$r'(\phi) = -\Delta f \phi^{-\Delta f-1} + \alpha \Delta f \phi^{\Delta f-1}$$

So:

$$r'(\phi) = 0 \iff -\phi^{-\Delta f-1} + \alpha \Delta f \phi^{\Delta f-1} = 0$$

Thus:

$$r'(\phi) = 0 \iff -1 + \alpha \phi^{2\Delta f} = 0 \iff \phi = \alpha^{\frac{1}{2\Delta f}}$$

Since r is continuous on $]0, +\infty[$ and $\lim_{\phi \rightarrow 0^+} r(\phi) = \lim_{\phi \rightarrow +\infty} r(\phi) = +\infty$ and r' vanishes in exactly one point $\phi_{min} = \alpha^{\frac{1}{2\Delta f}}$, then:

$$\min_{\phi \in]0, +\infty[} r(\phi) = r(\phi_{min}) = 2\sqrt{\alpha} = 2\sqrt{\frac{m-k}{k}} \sqrt{\frac{m-t}{t}}$$

So to guarantee that $\Delta h_{k,t}(\phi) \geq 0$ it suffices that:

$$\left[\underbrace{2\sqrt{\frac{m-k}{k}} \sqrt{\frac{m-t}{t}}}_{r(\phi_{min})} + \frac{m-k}{k} + \frac{m-t}{t} \right] \Delta g_{k,t} - \left(\frac{m-t}{t} - \frac{m-k}{k} \right) \Delta f \geq 0$$

Since we have that:

$$\begin{aligned} & \left[2\sqrt{\frac{m-k}{k}}\sqrt{\frac{m-t}{t}} + \frac{m-k}{k} + \frac{m-t}{t} \right] \Delta g_{k,t} - \left(\frac{m-t}{t} - \frac{m-k}{k} \right) \Delta f \\ &= \left[\sqrt{\frac{m-k}{k}} + \sqrt{\frac{m-t}{t}} \right]^2 \Delta g_{k,t} - \left(\frac{m-t}{t} - \frac{m-k}{k} \right) \Delta f \\ &\propto \left[\sqrt{k(m-t)} + \sqrt{t(m-k)} \right] \Delta g_{k,t} - \left[\sqrt{k(m-t)} - \sqrt{t(m-k)} \right] \Delta f \end{aligned}$$

if we prove that $\frac{\sqrt{k(m-t)} - \sqrt{t(m-k)}}{\sqrt{k(m-t)} + \sqrt{t(m-k)}} \leq \frac{k-t}{2}$ then we would have that:

$$\left[\sqrt{k(m-t)} + \sqrt{t(m-k)} \right] \Delta g_{k,t} - \left[\sqrt{k(m-t)} - \sqrt{t(m-k)} \right] \Delta f \geq \Delta g_{k,t} - \frac{k-t}{2} \Delta f$$

Thus, since we want to prove that having $\Delta g_{k,t} \geq \frac{k-t}{2} \Delta f$ is a sufficient condition to guarantee a positive derivative of the expected ballot's size it only remains to prove that: $\frac{\sqrt{k(m-t)} - \sqrt{t(m-k)}}{\sqrt{k(m-t)} + \sqrt{t(m-k)}} \leq \frac{k-t}{2}$.

For $k \geq t + 2$, we have that:

$$\frac{k-t}{2} \geq 1 \geq \frac{\sqrt{k(m-t)} - \sqrt{t(m-k)}}{\sqrt{k(m-t)} + \sqrt{t(m-k)}}$$

For $k = t + 1$, we have to prove that:

$$\frac{\sqrt{-t^2 + (m-1)t + m} - \sqrt{-t^2 + (m-1)t}}{\sqrt{-t^2 + (m-1)t + m} + \sqrt{-t^2 + (m-1)t}} \geq \frac{1}{2}$$

For this it suffices to prove that:

$$\sqrt{-t^2 + (m-1)t + m} \leq 3\sqrt{-t^2 + (m-1)t}$$

which is equivalent to proving that:

$$8t^2 - (m-1)t + m \leq 0$$

We can prove that this function decreases to reach a minimum in $t_0 = \frac{m-1}{2}$ then increases. We also have that it is negative in $t = 1$ (equals $-8m$) and for $t = m - 2$ (equals $-8(m-2)^2 + m$).

This gives us the sufficient condition to have $\frac{\partial \mathbb{E}_{\phi, a}[|A_i| | a_* = a]}{\partial \phi} \geq 0$. \square

Example 3.7. For $\alpha, \beta > 0$ define the distance:

$$\begin{aligned} d_{\alpha,\beta}(a, A) &= \alpha|\bar{a} \cap A| + \beta|a \cap \bar{A}| \\ &= -(\alpha + \beta)|a \cap A| + \beta + \alpha|A| \end{aligned}$$

which generalizes the Hamming distance in the same way the Tversky index [Tversky, 1977] generalizes Jaccard's. $d_{\alpha,\beta}$ is associated to the separable function:

$$\psi_{d_{\alpha,\beta}}(x, k) = -(\alpha + \beta)x + \beta + \alpha k = f(x) + g(k)$$

where $f(x) = -(\alpha + \beta)x + \beta$ and $g(k) = \alpha k$. We have:

$$\Delta g_{k,t} = \frac{\alpha}{\alpha + \beta}(k - t)\Delta f$$

So for every $d_{\alpha,\beta}$ such that $\alpha \geq \beta$ we have that:

$$\frac{\partial \mathbb{E}_{\phi, d_{\alpha,\beta}}[|A_i|]}{\partial \phi} \geq 0$$

3.4.2 The Hamming Distance Case - Condorcet Noise Model

The prototypical example of a separable noise is the noise associated to the Hamming distance, which is equivalent to the Condorcet-like noise model. We will prove that for this specific noise, we can express the expected size of a voter's ballot $\mathbb{E}[|A_i|]$ as a linear function of her reliability parameter. This enables us to estimate this parameter directly from the actual size of the ballot, without any prior belief about the ground truth.

Formally, consider the Condorcet noise model where for each voter $i \in N$ there exists a noise parameter $p_i \in (0, 1)$ such that:

$$P_{p_i}(a \in A_i | a = a^*) = P_{p_i}(a \notin A_i | a \neq a^*) = p_i, \forall a \in \mathcal{X}$$

and where the belonging or not of different alternatives to the voter's ballot are independent events. Notice that the model supposes equal error-rates for false positives and false negatives. In particular, voters who select many alternatives would *ipso facto* have a low p_i (since their ballots contain many false positives) which can even be below 0.5. Moreover, we can easily prove that in this case, the noise model is a non-anonymous Mallows noise model with the Hamming distance and with $\phi_i = \frac{1-p_i}{p_i}$ (We can have $\phi_i \geq 1$ since p_i can be below 0.5):

$$P_{p_i}(A_i | a = a^*) = p_i^m \left(\frac{1-p_i}{p_i} \right)^{d_H(a^*, A_i)}, \forall a \in \mathcal{X}$$

We will show that in this particular case, we can give an explicit formula of the expected size of a voter's approval ballot as a linear function of her precision parameter p_i .

Theorem 3.4. For $m \geq 2$, we have that:

$$\mathbb{E}_p[|A_i|] = (m - 1) - (m - 2)p$$

Proof. Let $a \in \mathcal{X}$:

$$\begin{aligned} \mathbb{E}_p[|A_i| | a^* = a] &= \mathbb{E}\left[\sum_{b \in \mathcal{X}} \mathbb{1}\{b \in A_i\} | a^* = a\right] \\ &= \sum_{b \in \mathcal{X}} P(b \in A_i | a^* = a) \\ &= P(a \in A_i | a^* = a) + \sum_{b \neq a} P(b \in A_i | a^* = a) \\ &= p + (m - 1)(1 - p) \\ &= (m - 1) - (m - 2)p \end{aligned}$$

Thus we have that:

$$\begin{aligned} \mathbb{E}_p[|A_i|] &= \sum_{a \in \mathcal{X}} \mathbb{E}_p[|A_i| | a^* = a] P(a^* = a) \\ &= \sum_{a \in \mathcal{X}} [(m - 1) - (m - 2)p] P(a^* = a) \\ &= (m - 1) - (m - 2)p \end{aligned}$$

□

Theorem 3.4 gives us a simple approach to estimate p_i by maximum likelihood estimations given some observations of A_i without a need to know the ground truth a^* . In a single-question scenario, we would only access a single observation of the voter's ballot A_i , and the estimation would not be statistically significant. Nonetheless, we show in the next section that it yields accurate ground truth estimations on real annotation datasets.

3.5 Experiments

We took the three image annotation datasets, originally collected in [Shah et al., 2015] for incentive-design purposes⁴, and used them to test our hypothesis and to assess the accuracy of different aggregation rules of interest.

Each dataset consists of a set of approval profiles of a number of voters (participants) who had to select *all the alternatives that they thought were correct* in a number of instances (images), namely:

- Animal task: 16 images/questions and 6 alternatives, see Figure 3.1.

⁴Accessible on the author's webpage: https://cs.cmu.edu/~nihars/data/data_approval.zip



Figure 3.1: Instance from the Animals task [Shah et al., 2015]

- Texture task: 16 images/questions and 6 alternatives, see Figure 3.2.
- Language task: 25 images/questions and 8 alternatives, see Figure 3.3.

From now on, a dataset denotes the set N of n voters, $\mathcal{X} = \{a_1, \dots, a_m\}$ the set of alternatives, L approval profiles $A^z = (A_1^z, \dots, A_n^z)$ each associated to an image z with ground truth alternative a_z^* .

3.5.1 Ballot Size and Reliability

To test the hypothesis that smaller ballots are more reliable, we introduce the *size-normalized accuracy* which is defined for each dataset and each $k \in [1, m - 1]$ as:

$$\text{SNA}(k) = \frac{1}{k} \frac{|\{A_i^z, |A_i^z| = k, a_z^* \in A_i^z\}|}{|\{A_i^z, |A_i^z| = k\}|}$$

It can be interpreted as the probability of recovering the ground truth after drawing randomly (uniformly) an alternative from a ballot of size k . Notice that if smaller ballots were not more reliable, one would expect that, for instance, ballots of size 2 are twice more probable to contain the ground truth than ballots of size 1, so the chance of finding the ground truth after randomly picking an alternative from a 2-sized ballot is equal to the chance that a singleton ballot selects the ground truth. So we would expect that SNA is almost constant for all k .

However when we compute the SNA for the three datasets (see Figure 3.4) we can clearly see that it decreases for the bigger ballots, which confirms that the alternatives selected in smaller approval ballots are more likely to coincide with the ground truth.

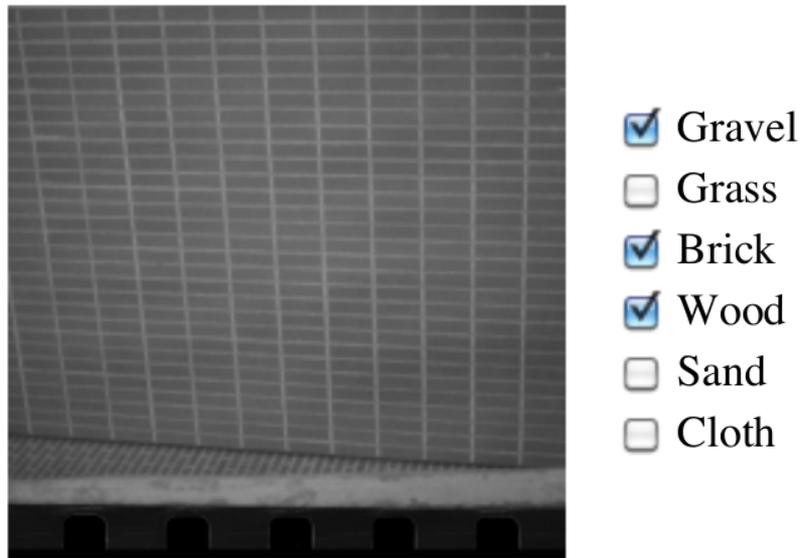
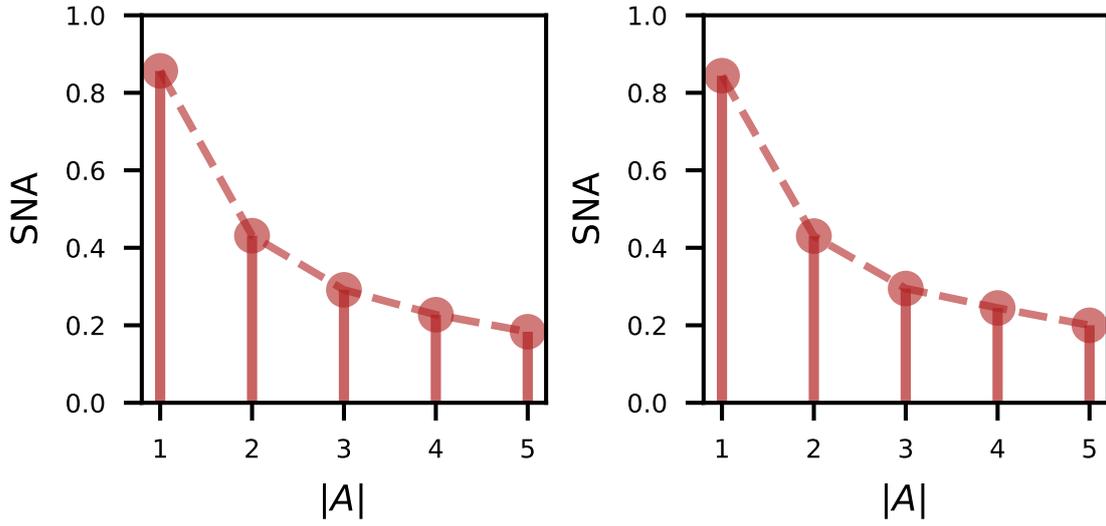


Figure 3.2: Instance from the Textures task [Shah et al., 2015]

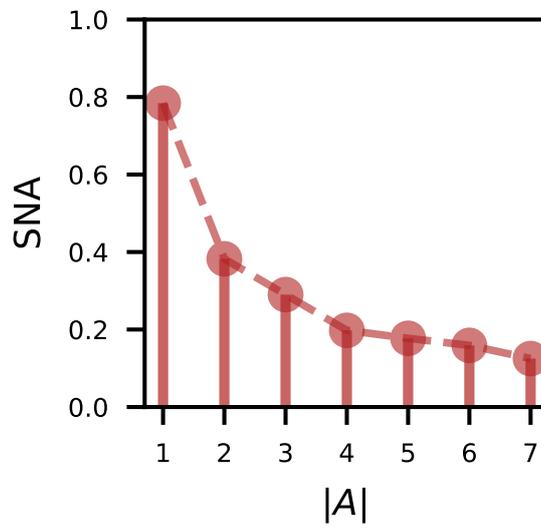


Figure 3.3: Instance from the Languages task [Shah et al., 2015]



(a) Animals

(b) Textures



(c) Languages

Figure 3.4: Size-normalized accuracy for three datasets

3.5.2 Aggregation

Since we are mostly interested in the single-question wisdom of the crowd problem, we will only consider aggregation rules that operate question-wise (voters' answers on different questions do not affect the output of the rule for a given question). We will use the following aggregation methods (we include more methods in the Appendix)⁵:

Condorcet: For each instance with approval profile $A^z = (A_1^z, \dots, A_n^z)$ and ground truth a_z^* , we suppose that each voter has a precision parameter p_i^z such that:

$$P_{p_i^z}(a \in A_i^z | a = a_z^*) = P_{p_i^z}(a \notin A_i^z | a \neq a_z^*) = p_i^z, \forall a \in \mathcal{X}$$

and where the belonging or not of different alternatives to the voter's ballot are independent events. We know that if these parameters were known, then the maximum likelihood estimation rule returns the alternative \hat{a}_z such that:

$$\hat{a}_z = \arg \max_{a \in \mathcal{X}} \sum_{i: a \in A_i^z} \ln \frac{p_i^z}{1 - p_i^z}$$

To estimate the parameters p_i^z , we will make use of the expression in Theorem 3.4 that states that:

$$\mathbb{E}_{p_i^z}[|A_i^z|] = (m - 1) - (m - 2)p_i^z$$

and set:

$$\hat{p}_i^z = \text{proj}_{[\varepsilon, 1-\varepsilon]} \frac{m - 1 - |A_i^z|}{m - 2}$$

where proj is the projection operator on the corresponding interval. The projected quantity is simply the maximum likelihood estimation of p_i^z with a single sample (the actual observation of the voter's ballot). We project it into a closed interval to avoid having $\hat{p}_i^z = 1$ (which yields an infinite weight to the voter) whenever $|A_i^z| = 1$, and $\hat{p}_i^z = 0$ whenever $|A_i^z| = m - 1$. So the aggregation rule finally outputs:

$$\hat{a}_z = \arg \max_{a \in \mathcal{X}} \sum_{i: a \in A_i^z} \ln \frac{\hat{p}_i^z}{1 - \hat{p}_i^z}$$

which is size-decreasing.

Jaccard: Here we suppose that for each instance $A^z = (A_1^z, \dots, A_n^z)$, the noise model is as follows:

$$P_{\phi, d_J}(A_i^z | a_z^* = a) = \frac{1}{Z} \phi^{d_J(a_z^*, A_i^z)}, \forall a \in \mathcal{X}$$

where $d_J(a, A) = 1 - \frac{|a \cap A|}{|A| + 1 - |a \cap A|}$.

We saw in Example 3.4 that the maximum likelihood estimation rule ζ_{d_J} is a size approval rule with weights $w_j = 1/j$.

⁵The code can be found at https://github.com/taharallouche/Truth_Tracking-via-AV

Simple approval: We will compare all these rules to the benchmark SAV rule where for each instance $A^z = (A_1^z, \dots, A_n^z)$: $\hat{a}_z = \arg \max_{a \in \mathcal{X}} |\{A_i^z, a \in A_i^z\}|$

3.5.3 Results

For each task, we took 25 batches for each different number of voters, and applied the aforementioned rules. We measure the accuracy of each rule, outputting the estimates \hat{a}^z for each instance, defined as

$$\frac{1}{L} \sum_{z=1}^L \mathbb{1}\{a_z^* = \hat{a}_z\}$$

The results are shown in Figures 3.5a, 3.5b and 3.5c respectively for the Animals, Textures and Languages datasets.

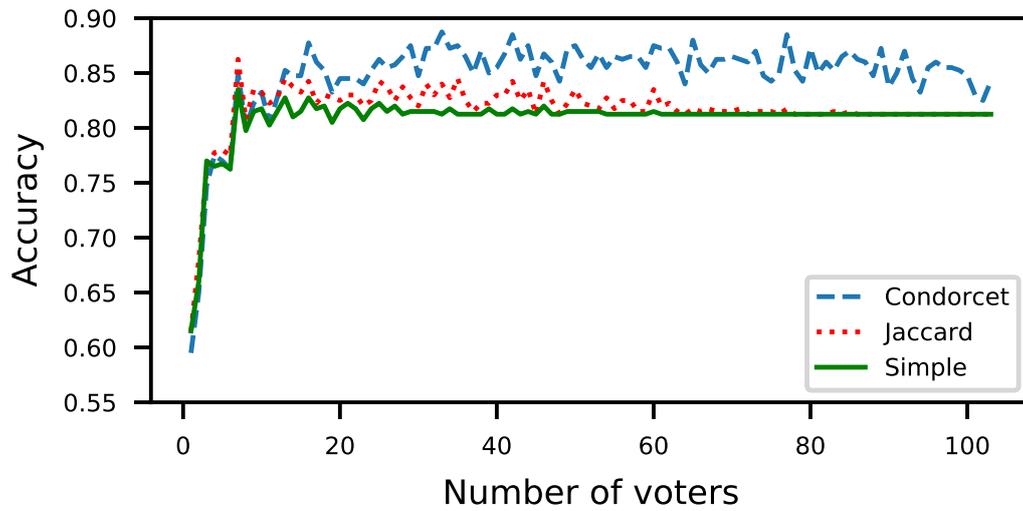
Observations: First we notice that for all the three datasets, the aggregation rules associated to Jaccard anonymous noise show slightly better accuracy than the simple approval rule especially for small number of voters.

We can also see that the aggregation rule associated to the non-anonymous Condorcet noise show significant improvement in the accuracy compared to this rule for Animals and Languages (specially for relatively big numbers of voters). However it fails to outperform it for the Textures dataset, where it only shows similar accuracies to the standard rule as the number of voters grows. This can be the result of the poor estimation quality which uses only one sample.

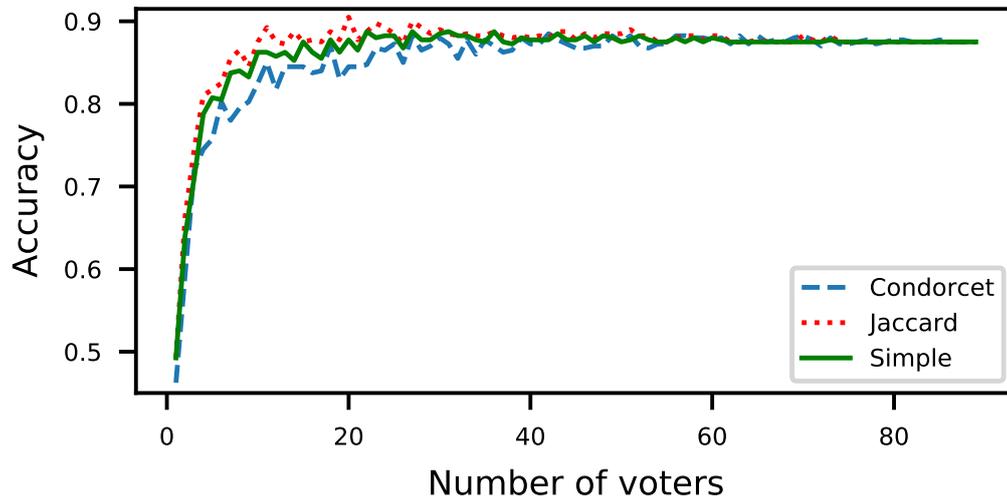
3.6 Conclusion

We propose a novel approach for epistemic approval voting based on the intuition that more reliable votes contain fewer alternatives. First, we show that for different anonymous variants of Mallows-like noise models, the maximum likelihood rule is size-decreasing, *i.e* it assigns more weight to smaller ballots. Then we consider non-anonymous noises and give a sufficient condition to have an expected size of the ballot which increases as a voter gets less reliable. In particular, we prove that for a Condorcet-like noise, the expected number of approved alternatives decreases linearly with the voter's precision. Finally, we conduct experiments to test our hypothesis on real data and to assess the performances of different aggregation rules.

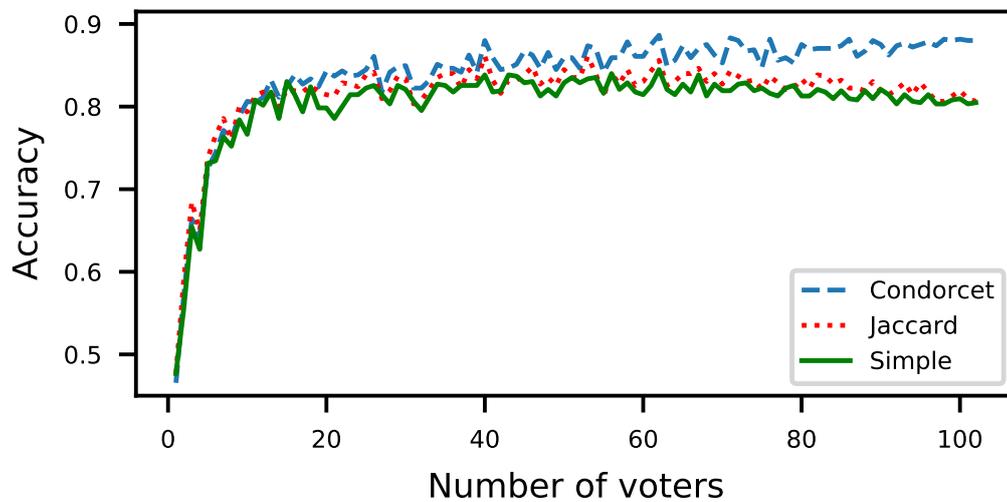
These methods may fail in two possible scenarios. First, if the voters do not respond truthfully. In this case, a voter can select a single alternative even though she doesn't know the correct answer at all. Second, if a large enough group of non-expert voters are mistakenly over-self-confident, whereas the experts are uncertain about their responses.



(a) Animals



(b) Textures



(c) Languages

Figure 3.5: Accuracy of different rules on real-life datasets

Multi-winner Epistemic Voting: Committee-size as Prior Knowledge

Abstract

Epistemic voting interprets votes as noisy signals about a ground truth. We consider contexts where the truth consists of a set of objective winners, knowing a lower and upper bound on its cardinality. A prototypical problem for this setting is the aggregation of multi-label annotations with prior knowledge on the size of the ground truth. We posit noise models, for which we define rules that output an optimal set of winners. We report on experiments on multi-label annotations (which we collected).

The results of this chapter were published in [Allouche et al., 2022a].

4.1 Introduction

The epistemic view of voting assumes the existence of a ground truth which, usually, is either an alternative or a ranking over alternatives. Votes reflect opinions or beliefs about this ground truth; the goal is to aggregate these votes so as to identify it. Usual methods define a noise model specifying the probability of each voting profile given the ground truth, and output the alternative that is the most likely state of the world, or the ranking that is most likely the true ranking.

Now, there are contexts where the ground truth does not consist of a single alternative nor a ranking, but of a *set of alternatives*. Typical examples are multi-label crowdsourcing (find the items in a set that satisfy some property, *e.g.* the sport teams appearing on a picture) or finding the objectively k best candidates (best papers at a conference, best performance in artistic sports, k patients with highest probabilities of survival if being assigned a scarce medical resource). Beyond social choice, collective multi-label annotation was first addressed in [Nowak and R uger, 2010], which studies the agreement between experts and non-experts in some multi-labelling

tasks, and in [Deng et al., 2014], where a scalable aggregation method is presented to solve the multi-label estimation problem.

The alternatives that are truly in the ground truth are called ‘winning’ alternatives. Depending on the context, the number of winning alternatives can be fixed, unconstrained, or more generally, constrained to be in a given interval. This constraint expresses some *prior knowledge on the cardinality of the ground truth*. Here are some examples:

- *Picture annotation via crowdsourcing*: participants are shown a picture taken from a soccer match and have to identify the team(s) appearing in it. The ground truth is known to contain one or two teams.
- *Guitar chord transcription*: voters are base classifier algorithms [Nguyen et al., 2020] which, for a given chord, select the set of notes constitute it. The true set of notes can contain three to six alternatives.
- *Jury*: participants are members of a jury which has to give an award to three papers presented at a conference: the number of objective winners is fixed to three. (In a variant, the number of awards would be *at most* three.)
- *Resource allocation*: participants are doctors and alternatives are Covid-19 patients in urgent need of intensive care; there is a limited number k of intensive care units. The ground truth consists of those patients who most deserve to be cured (for example those with the k highest probabilities of survival if cured).

We assume that voters provide a simple form of information: *approval ballots*, indicating which alternatives they consider plausible winners. These approval ballots are not subject to any cardinality constraint: *a voter may approve a number of alternatives, even if it does not lie in the interval bearing on the output*. This is typically the case for totally ignorant voters, who are expected to approve all alternatives.

Sometimes, the aggregating mechanism has some prior information about the likelihood of alternatives and the reliability of voters. We first study a simple case where this information is specified in the input: in the noise model, each voter has a probability p_i (resp. q_i) of approving a winning (resp. non-winning) alternative, and each alternative has a prior probability to be winning. This departs from classical voting, where voters are usually treated equally (*anonymity*), and similarly for alternatives (*neutrality*).

This simple case serves as a building component for the more complex case where these parameters are not known beforehand but *estimated from the votes*: votes allow to infer information about plausibly winning alternatives, from which we infer information about voter reliabilities, which leads to revise information about winning alternatives, and so on until the process converges. Here we move back to an anonymous and neutral setting, since all alternatives (resp. voters) are treated equally before votes are known.

The outline of the chapter is as follows. We introduce the model (Section 4.2) and give an estimation algorithm (Section 4.3), first in the case where the parameters are known, and then in

the case where they are estimated from the votes. In Section 3.5 we present a data gathering task and analyse the results of the experiments. Section 4.6 concludes.

4.2 The Model

Let $N = \{1, \dots, n\}$ be a set of voters, and $\mathcal{X} = \{a_1, \dots, a_m\}$ a set of alternatives (possible objects in images, notes in chords, papers, patients...). Consider a set of L instances: an instance z consists of an approval profile $A^z = (A_1^z, \dots, A_n^z)$ where $A_i^z \subseteq \mathcal{X}$ is an approval ballot for every $i \in N$. For example, in a crowdsourcing context, a task usually contains multiple questions, and an instance comprises the voters' answers to one of these questions.

For each instance $z \in L$, there exists an *unknown* ground truth S_z^* belonging to $\mathcal{S} = 2^{\mathcal{X}}$, which is the set of objectively correct alternatives in instance z . It is common knowledge that the number of alternatives in each of them lies in the interval $[l, u]$: $S_z^* \in \mathcal{S}_{l,u} = \{S \in \mathcal{S}, l \leq |S| \leq u\}$, for given bounds $0 \leq l \leq u \leq m$.

Our goal is to unveil the ground truth for each of these instance using the votes and the prior knowledge on the number of winning alternatives. We define a noise model consisting of two parametric distributions, namely, a conditional distribution of the approval ballots given the ground truth, and a prior distribution on the ground truth. Here we depart from classical noise models in epistemic social choice, as we suppose that the parameters of these distributions may be unknown and thus need to be estimated.

For each voter $i \in N$, we suppose that there exist two unknown parameters (p_i, q_i) in $(0, 1)$ such that the approval ballot A_i^z on an instance $z \in L$ is drawn according to the following distribution: for each $a \in \mathcal{A}$,

$$P(a \in A_i^z | S_z^* = S) = \begin{cases} p_i & \text{if } a \in S \\ q_i & \text{if } a \notin S \end{cases}$$

where p_i (resp. q_i) is the (unknown) probability that voter i approves a correct (resp. incorrect) alternative. Then we make the following assumptions:

- (1) A voter's approvals of alternatives are mutually independent given the ground truth and parameters $(p_i, q_i)_{i \in N}$.
- (2) Voters' ballots are mutually independent given the ground truth.
- (3) Instances are independent given the parameters $(p_i, q_i)_{i \in N}$ and the ground truths.

To model the prior probability of any set S to be the ground truth S^* , we define parameters $t_j = P(a_j \in S^*)$. t_j can be understood as the prior probability of a_j to be in the ground truth set S^* before the cardinality constraints are taken into account. Let $t = (t_1, \dots, t_m)$ be the vector of all m prior parameters. These, together with an independence assumption on the events $\{a_j \in S^*\}$, gives:

$$P(S = S^*) = \prod_{a_j \in S} t_j \prod_{a_j \notin S} 1 - t_j$$

Note that the choice of the parameters t_j is not crucial when running the algorithm for estimating the ground truth: we will see in Section 4.3.3 that it converges whatever their values. The distribution conditional to the prior knowledge on the size of the ground truth can be seen as a projection on the constraints followed by a normalization:

$$\tilde{P}(S) = P(S^* = S | l \leq |S^*| \leq u) = \frac{P(S^* = S \cap |S^*| \in [l, u])}{P(|S^*| \in [l, u])}$$

It follows:

$$\tilde{P}(S) = \begin{cases} \frac{1}{\beta(l, u, t)} \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j) & \text{if } S \in \mathcal{S}_{l, u} \\ 0 & \text{if } S \notin \mathcal{S}_{l, u} \end{cases}$$

where:

$$\beta(l, u, t) = \sum_{S \in \mathcal{S}_{l, u}} \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j)$$

The ground truths associated with different instances are assumed to be mutually independent given the parameters.

Two particular cases are worth discussing.

- When $(l, u) = (0, m)$, the problem is *unconstrained* and we have $\beta(0, m, t) = P(|S^*| \in [0, m]) = 1$, so $\tilde{P}(S) = P(S = S^*)$. In this case the problem degenerates into a series of independent binary label-wise estimations (see Subsection 4.3.1).
- In the single-winner case $(l, u) = (1, 1)$, we have:

$$\tilde{P}(\{a_j\}) = P(\{a_j\} = S^* | |S^*| = 1) = \frac{t_j \prod_{h \neq j} (1 - t_h)}{\beta(1, 1, t)}$$

therefore, for any approval profile A :

$$P(S^* = \{a_j\} | A, |S^*| = 1) \propto \frac{t_j}{1 - t_j} P(A | S^* = \{a_j\})$$

We recover the same estimation problem if we simply introduce $\alpha_j = P(S^* = \{a_j\})$ with $\sum \alpha_j = 1$ as in [Ben-Yashar and Paroush, 2001], in which case we have:

$$P(S^* = \{a_j\} | A, |S^*| = 1) \propto \alpha_j P(A | S^* = \{a_j\})$$

4.3 Estimating the Ground Truth

Our aim is the intertwined estimation of the ground truth and the parameters via maximizing the total likelihood of the instances:

$$\mathcal{L}(A, S, p, q, t) = \prod_{z=1}^L \tilde{P}(S_z) \prod_{i=1}^n P(A_i^z | S_z)$$

where:

$$P(A_i^z | S_z) = p_i^{|A_i^z \cap S_z|} q_i^{|A_i^z \cap \overline{S_z}|} (1 - p_i)^{|\overline{A_i^z} \cap S_z|} (1 - q_i)^{|\overline{A_i^z} \cap \overline{S_z}|}$$

To this aim, we will introduce an iterative algorithm whose main two steps will be presented in sequence, in the next subsections, before the main algorithm is formally defined and its convergence shown. These two steps are:

- Estimating the ground truths given the parameters.
- Estimating the parameters given the ground truths.

Simply put, the algorithm consists in iterating these two steps until it converges to a fixed point.

4.3.1 Estimating the Ground Truth Given the Votes and the Parameters

Since instances are independent given the parameters, we focus here on one instance with ground truth S^* and profile $A = (A_1, \dots, A_n)$. Before diving into maximum likelihood estimation (MLE), we introduce some notions and prove some lemmas. In this subsection, we suppose that the parameters $(p_i, q_i)_{i \in N}$ and $(t_j)_{j \in \mathcal{X}}$ are known (later on, these parameters will be replaced by their estimations at each iteration of the algorithm). Thus, all in all, input and output are as follows:

- Input: approval profile A ; parameters $(p_i, q_i)_{i \in N}$ and $(t_j)_{j \in \mathcal{X}}$.
- Output: MLE of the ground truth S^* .

Definition 4.1 (weighted approval score). *Given an approval profile (A_1, \dots, A_n) , noise parameters $(p_i, q_i)_{1 \leq i \leq n}$ and prior parameters $(t_j)_{1 \leq j \leq m}$, define:*

$$app_w(a_j) = \ln \left(\frac{t_j}{1 - t_j} \right) + \sum_{i: a_j \in A_i} \ln \left(\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right)$$

The scores $app_w(a_j)$ can be interpreted as weighted approval scores for a $(n + m)$ -voter profile where:

- for each voter $1 \leq i \leq n$: i has a weight $w_i = \ln \left(\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right)$ and casts approval ballot A_i .
- for each $1 \leq j \leq m$: there is a virtual voter with weight $w_j = \ln \left(\frac{t_j}{1 - t_j} \right)$ who casts approval ballot $A_j = \{a_j\}$.

While the weight of each voter $i \in N$ depends on her reliability, each prior information on an alternative plays the role of a virtual voter who only selects the concerned alternative, with a weight that increases as the prior parameter increases.

From now on, we suppose without loss of generality that the alternatives are ranked according to their score:

$$app_w(a_1) \geq app_w(a_2) \geq \dots \geq app_w(a_m)$$

Definition 4.2 (threshold and partition). *Define the threshold:*

$$\tau_n = \sum_{i=1}^n \ln \left(\frac{1 - q_i}{1 - p_i} \right)$$

and the partition of the set of alternatives in three sets:

$$\begin{cases} S_{max}^{\tau_n} &= \{a \in A, app_w(a) > \tau_n\} \\ S_{tie}^{\tau_n} &= \{a \in A, app_w(a) = \tau_n\} \\ S_{min}^{\tau_n} &= \mathcal{X} \setminus (S_{max}^{\tau_n} \cup S_{tie}^{\tau_n}) \end{cases}$$

and let $k_{max}^{\tau_n} = |S_{max}^{\tau_n}|$, $k_{tie}^{\tau_n} = |S_{tie}^{\tau_n}|$, $k_{min}^{\tau_n} = |S_{min}^{\tau_n}|$.

The next result characterizes the sets in \mathcal{S} that are MLEs of the ground truth given the parameters.

Theorem 4.1. $\tilde{S} \in \arg \max_{S \in \mathcal{S}} \mathcal{L}(A, S, p, q, t)$ *if and only if there exists $k \in [l, u]$ such that \tilde{S} is the set of k alternatives with the highest k values of app_w and:*

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau_n}| &= \min(u, k_{max}^{\tau_n}) \\ |\tilde{S} \cap S_{min}^{\tau_n}| &= \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases} \quad (4.1)$$

So the estimator \tilde{S} is made of some top- k alternatives, where the possible values of k are determined by Eq. (4.1). The first equation imposes that \tilde{S} includes as many elements as possible from $S_{max}^{\tau_n}$ (without exceeding the upper-bound u), whereas the second one imposes that \tilde{S} includes as few elements as possible from $S_{min}^{\tau_n}$ (without getting below the lower-bound l).

Proof. Since $\tilde{P}(S) > 0 \iff S \in \mathcal{S}_{l,u}$, we have that $\arg \max_{S \in \mathcal{S}} L(S) = \arg \max_{S \in \mathcal{S}_{l,u}} L(S)$. Using the assumption on the ballots' independence, we can write:

$$L(S) \propto \tilde{P}(S) \prod_{i=1}^n P(A_i | S)$$

which, by using the independence assumption on alternatives, becomes, for any $S \in \mathcal{S}_{l,u}$: Moreover, we have that for any $S \in \mathcal{S}_{l,u}$:

$$\begin{aligned} L(S) &= \tilde{P}(S) \prod_{i=1}^n p_i^{|A_i \cap S|} q_i^{|A_i \cap \bar{S}|} (1 - p_i)^{|\bar{A}_i \cap S|} (1 - q_i)^{|\bar{A}_i \cap \bar{S}|} \\ &= \tilde{P}(S) \prod_{i=1}^n p_i^{|A_i \cap S|} q_i^{|A_i| - |A_i \cap S|} (1 - p_i)^{|S| - |A_i \cap S|} (1 - q_i)^{|\bar{A}_i| - |S| + |A_i \cap S|} \\ &\propto \tilde{P}(S) \prod_{i=1}^n \left[\frac{1 - p_i}{1 - q_i} \right]^{|S|} \left[\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right]^{|A_i \cap S|} \\ &\propto \underbrace{\frac{1}{\beta(l, u, t)}}_{>0} \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j) \prod_{i=1}^n \left[\frac{1 - p_i}{1 - q_i} \right]^{|S|} \left[\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right]^{|A_i \cap S|} \end{aligned}$$

$$\begin{aligned} &\propto \underbrace{\prod_{a_j \in \mathcal{X}} (1 - t_j)}_{>0} \prod_{a_j \in S} \frac{t_j}{1 - t_j} \prod_{i=1}^n \left[\frac{1 - p_i}{1 - q_i} \right]^{|S|} \left[\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right]^{|A_i \cap S|} \\ &\propto \prod_{a_j \in S} \frac{t_j}{1 - t_j} \prod_{i=1}^n \left[\frac{1 - p_i}{1 - q_i} \right]^{|S|} \left[\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right]^{|A_i \cap S|} \end{aligned}$$

Thus, with C a constant stemming from previous simplifications, the log-likelihood reads:

$$\begin{aligned} l(S) &= \sum_{a_j \in S} \ln \frac{t_j}{1 - t_j} + \sum_{i=1}^n |S| \ln \frac{1 - p_i}{1 - q_i} + |A_i \cap S| \ln \frac{p_i(1 - q_i)}{q_i(1 - p_i)} + C \\ &= \sum_{a_j \in S} \left[\underbrace{\ln \frac{t_j}{1 - t_j} + \sum_{i: a_j \in A_i} \ln \frac{p_i(1 - q_i)}{q_i(1 - p_i)}}_{app_w(a_j)} - \underbrace{\sum_{i=1}^n \ln \frac{1 - q_i}{1 - p_i}}_{\tau_n} \right] + C \\ &= \sum_{a_j \in S} [app_w(a) - \tau_n] + C \\ &= \sum_{a_j \in S} l(a) + C \end{aligned}$$

This means that $a \in S_{max}^{\tau_n}$ if and only if $l(a) > 0$, $a \in S_{min}^{\tau_n}$ if and only if $l(a) < 0$ and $a \in S_{tie}^{\tau_n}$ if and only if $l(a) = 0$.

Now, let S_M be a maximizer of the likelihood. Since $l(a_j) \geq l(a_h) \iff app_w(a_j) \geq app_w(a_h)$ we have that S_M , which maximizes $\sum_{a_j \in S} l(a_j)$, is made of top- k alternatives for some $k \in [l \dots u]$.

Furthermore, $|S_M \cap S_{min}^{\tau_n}| = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n})$.

In fact, start by noticing that $|S_M \cap S_{min}^{\tau_n}| \geq \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n})$, since:

$$|S_M \cap S_{min}^{\tau_n}| \geq l - |S_M \cap S_{max}^{\tau_n}| - |S_M \cap S_{tie}^{\tau_n}| \geq l - k_{max}^{\tau_n} - k_{tie}^{\tau_n}$$

Now suppose that $|S_M \cap S_{min}^{\tau_n}| > \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n})$. In this case we have that $|S_M| > l$ because otherwise, if $|S_M| = l$, then:

$$|S_M \cap S_{max}^{\tau_n}| + |S_M \cap S_{tie}^{\tau_n}| = l - |S_M \cap S_{min}^{\tau_n}| < k_{max}^{\tau_n} + k_{tie}^{\tau_n}$$

which would mean that there are elements in $S_{tie}^{\tau_n}$ and $S_{max}^{\tau_n}$ which are not in S_M , which is a contradiction since $|S_M \cap S_{min}^{\tau_n}| > 0$ and S_M is a top- k set.

Now consider $a \in S_M \cap S_{min}^{\tau_n}$, we have that $|S_M \setminus \{a\}| \geq l$ and $l(S_M) = l(S_M \setminus \{a\}) + l(a) < l(S_M \setminus \{a\})$ which is a contradiction.

With the same idea we can prove that $|S_M \cap S_{max}^{\tau_n}| = \min(u, k_{max}^{\tau_n})$. In fact, it is obvious that $|S_M \cap S_{max}^{\tau_n}| \leq \min(u, k_{max}^{\tau_n})$ since:

$$|S_M \cap S_{max}^{\tau_n}| \leq |S_{max}^{\tau_n}| = k_{max}^{\tau_n} \text{ and } |S_M \cap S_{max}^{\tau_n}| \leq |S_M| \leq u$$

Now suppose that $|S_M \cap S_{max}^{\tau_n}| < \min(u, k_{max}^{\tau_n})$. First, notice that $|S_M| < u$ since if $|S_M| = u$ then $|S_M \cap S_{max}^{\tau_n}| < |S_M|$ and $|S_M \cap S_{max}^{\tau_n}| < k_{max}^{\tau_n}$ which would mean that S_M cannot be made of top- k alternatives since it would have had to contain alternatives from $S_{tie}^{\tau_n}$ and $S_{min}^{\tau_n}$ while omitting some alternatives from $S_{max}^{\tau_n}$.

Now consider $a \in S_{max}^{\tau_n} \setminus S_M$, we have that $|S_M \cup \{a\}| \leq u$ and $l(S_M \cup \{a\}) = l(S_M) + l(a) > l(S_M)$ which is a contradiction.

Conversely, consider a set $S_M = \{a_1, \dots, a_k\}$ such that:

$$\begin{cases} k & \in [l, u] \\ |S_M \cap S_{max}^{\tau_n}| & = \min(u, k_{max}^{\tau_n}) \\ |S_M \cap S_{min}^{\tau_n}| & = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases}$$

and let us prove that it maximizes the likelihood. To do so, consider $S' \in \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S|A)$, so by the first part of the proof there exists some k' such that $S' = \{a_1, \dots, a_{k'}\}$ such that:

$$\begin{cases} k' & \in [l, u] \\ |S' \cap S_{max}^{\tau_n}| & = \min(u, k_{max}^{\tau_n}) \\ |S' \cap S_{min}^{\tau_n}| & = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases}$$

Since $|S' \cap S_{max}^{\tau_n}| = |S_M \cap S_{max}^{\tau_n}|$ and $|S' \cap S_{min}^{\tau_n}| = |S_M \cap S_{min}^{\tau_n}|$ and given the structure of $S' = \{a_1, \dots, a_{k'}\}$ and $S_M = \{a_1, \dots, a_k\}$, we have that $S' \cap S_{max}^{\tau_n} = S_M \cap S_{max}^{\tau_n}$ and $S' \cap S_{min}^{\tau_n} = S_M \cap S_{min}^{\tau_n}$.

Given the expression of the likelihood of a set of alternatives, and given that $l(a) = 0$ for any $a \in S_{tie}^{\tau_n}$ we have that $l(S_M) = l(S')$, which implies that $S_M \in \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S|A)$. \square

Notice that when $(l, u) = (0, m)$, the problem degenerates into a collection of label-wise problems, one for each alternative: a_j is selected if $a_j \in S_{max}^{\tau_n}$, rejected if $a_j \in S_{min}^{\tau_n}$, and those that are on the fence can be arbitrarily selected or not.

Example 4.1. Consider the following situation:

- 5 alternatives $\mathcal{X} = \{a, b, c, d, e\}$.
- 10 voters N .
- All voters share the same parameters $(p, q) = (0.7, 0.4)$

We thus have that all voters share the same weight:

$$w = \ln \left(\frac{p(1-q)}{q(1-p)} \right) = 1.25$$

and the threshold is

$$\tau_n = \sum_{i=1}^n \ln \left(\frac{1-q}{1-p} \right) = 6.93$$

We consider the constraints $(l, u) = (1, 4)$:

First, suppose that $t_a = 0.6$ and that $t_j = 0.5$ for all the remaining candidates. Consider also the approval counts (and weighted approval scores) in the table below.

Candidate	a	b	c	d	e
Approval count	9	8	7	5	5
app_w	11.25	10	8.75	6.65	6.25

We can easily check, by Theorem 4.1 that $\tilde{S} = \arg \max_{S \in \mathcal{S}} P(S = S^* | A) = \{a, b, c\}$. We have that $S_{max}^{\tau_n} = \{a, b, c\}$, $S_{tie}^{\tau_n} = \emptyset$ and $S_{min}^{\tau_n} = \{d, e\}$. We know that there exists some $k \in [1, 4]$ such that \tilde{S} would consist of the top k alternatives. We also have that:

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau_n}| = \min(u, k_{max}^{\tau_n}) = 3 \implies \{a, b, c\} \subseteq \tilde{S} \\ |\tilde{S} \cap S_{min}^{\tau_n}| = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) = 0 \implies d, e \notin \tilde{S} \end{cases}$$

So the only possibility is $\tilde{S} = \{a, b, c\}$.

Now suppose that $t_j = 0.5$ for all the alternatives $a_j \in \mathcal{X}$. We also change the approval counts:

Candidate	a	b	c	d	e
Approval count	5	4	3	2	2
app_w	6.25	5	3.75	2.5	2.5

We have that $S_{max}^{\tau_n} = S_{tie}^{\tau_n} = \emptyset$ and $S_{min}^{\tau_n} = \{a, b, c, d, e\}$, so the MLE must verify that $|\tilde{S} \cap S_{min}^{\tau_n}| = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) = 1$, which means that \tilde{S} consists of a single alternative belonging to $S_{min}^{\tau_n}$. Since the MLE is a top- k set (in this case it is top-1), we get that $\tilde{S} = \{a\}$.

Again, if we change the approval count in the following way:

Candidate	a	b	c	d	e
Approval count	9	8	7	7	6
app_w	11.25	10	8.75	8.75	7.5

We have that $S_{max}^{\tau_n} = \{a, b, c, d, e\}$ and $S_{tie}^{\tau_n} = S_{min}^{\tau_n} = \emptyset$, which implies that $|\tilde{S} \cap S_{max}^{\tau_n}| = \min(u, k_{max}^{\tau_n}) = 4$. The MLE is that made of the top-4 alternatives so $\tilde{S} = \{a, b, c, d\}$.

4.3.2 Estimating the Parameters Given the Ground Truth

4.3.2.1 Estimating the prior parameters over alternatives

Once the ground truths are estimated at one iteration of the algorithm, the next step consists in estimating the prior parameters $(t_j)_{j \in \mathcal{X}}$, with the ground truths being given (in Subsection

4.3.3 the ground truth will be replaced by its estimation at each iteration). The next proposition explicits the closed-form expression of the MLE of the prior parameter of each alternative given the ground truth of each instance S_z^* once the prior parameters of all other alternatives are fixed.

- Input: Approval profile (A_1, \dots, A_n) , ground truths S_z^* for all $z \in L$, and all but one prior parameters $(t_h)_{h \neq j}$.
- Output: MLE of t_j .

Proposition 4.2. For every $a_j \in \mathcal{X}$:

$$\arg \max_{t_j \in (0,1)} \mathcal{L}(A, S, p, q, t_j, t_{-j}) = \frac{occ(j)\bar{\alpha}_j}{(L - occ(j))\underline{\alpha}_j + occ(j)\bar{\alpha}_j}$$

$$\text{where: } \begin{cases} \bar{\alpha}_j &= \sum_{\substack{S \in \mathcal{S}_{l,u} \\ a_j \in S}} \prod_{\substack{a_h \in S \\ h \neq j}} t_h \prod_{a_h \notin S} (1 - t_h) \\ \underline{\alpha}_j &= \sum_{\substack{S \in \mathcal{S}_{l,u} \\ a_j \notin S}} \prod_{a_h \in S} t_h \prod_{\substack{a_h \notin S \\ h \neq j}} (1 - t_h) \\ occ(j) &= |\{z \in \{1, \dots, L\}, a_j \in S_z\}| \end{cases}$$

Notice that:

$$\bar{\alpha}_j = P(l \leq |S^*| \leq u | a_j \in S^*)$$

and that:

$$\underline{\alpha}_j = P(l \leq |S^*| \leq u | a_j \notin S^*)$$

so $\beta(l, u, t) = \bar{\alpha}_j t_j + \underline{\alpha}_j (1 - t_j)$.

Notice also that $occ(j)$ is the number of instances whose ground truth contains a_j .

Proof. Fix all sets $S_z \in \mathcal{S}_{l,u}$ and all the noise parameters $(p_i, q_i)_i$ and all the prior parameters t_h but for one t_j for some $j \leq m$, and let $t_j \in (0, 1)$:

$$\begin{aligned} \mathcal{L}(t_j, t_{-j}) &\propto \prod_{z=1}^L \frac{1}{\beta(l, u, t)} \prod_{a_h \in S_z} t_h \prod_{a_h \notin S_z} (1 - t_h) \\ &\quad \prod_{i \in N_z} p_i^{|A_i^l \cap S_z|} q_i^{|A_i^z \cap \bar{S}_z|} (1 - p_z)^{|\bar{A}_i^z \cap S_z|} (1 - q_i)^{|\bar{A}_i^z \cap \bar{S}_z|} \\ &\propto \prod_{z=1}^L \frac{1}{\beta(l, u, t_j, t_{-j})} \prod_{a_h \in S_z} t_h \prod_{a_h \notin S_z} (1 - t_h) \\ &\propto \left(\frac{1}{\beta(l, u, t_j, t_{-j})} \right)^L \underbrace{\prod_{z: a_j \in S_z} t_j}_{t_j^{occ(j)}} \underbrace{\prod_{z: a_j \notin S_z} (1 - t_j)}_{(1 - t_j)^{L - occ(j)}} \\ &\propto \left(\frac{1}{\beta(l, u, t_j, t_{-j})} \right)^L t_j^{occ(j)} (1 - t_j)^{L - occ(j)} \end{aligned}$$

Taking the log we can write the function as:

$$\ell(t_j) = -L \log \beta + \text{occ}(j) \log t_j + (L - \text{occ}(j)) \log(1 - t_j)$$

Its derivative reads:

$$\frac{\partial \ell}{\partial t_j} = -L \frac{\alpha_j - \bar{\alpha}_j}{\alpha_j t_j + \bar{\alpha}_j (1 - t_j)} + \text{occ}(j) \frac{1}{t_j} + (\text{occ}(j) - L) \frac{1}{1 - t_j}$$

Canceling it, we obtain:

$$t_j = \frac{\text{occ}(j) \bar{\alpha}_j}{(L - \text{occ}(j)) \alpha_j + \text{occ}(j) \bar{\alpha}_j}$$

The derivative vanishes in a single point in $(0, 1)$ and $\lim_{t_j \rightarrow 0} \ell(t_j) = \lim_{t_j \rightarrow 1} \ell(t_j) = -\infty$ thus ℓ reaches a unique maximum. \square

We will see later that the algorithm applies Proposition 4.2 sequentially to estimate the alternatives' parameters one by one (see Example 4.2).

4.3.2.2 Estimating the voter parameters

Once the ground truths are known (or estimated), we can estimate the voters' parameters (p, q) .

- Input: Instances (A^1, \dots, A^L) , ground truths (S_1^*, \dots, S_L^*) .
- Output: MLE of voter reliabilities (p, q) .

The next result simply states that the maximum likelihood estimator of p_i of some voter is the fraction of alternatives that the voter approves and that actually belong to the ground truth; the estimation of q_i is similar. See Example 4.2.

Proposition 4.3. *Fix sets $S_z \in \mathcal{S}_{l,u}$ and prior parameters t . Then:*

$$\arg \max_{(p,q) \in (0,1)^{2 \times n}} \mathcal{L}(A, S, p, q, t) = (\hat{p}, \hat{q})$$

where:

$$\hat{p}_i = \frac{\sum_{z \in L} |A_i^z \cap S_z|}{\sum_{z \in L} |S_z|}, \hat{q}_i = \frac{\sum_{z \in L} |A_i^z \cap \bar{S}_z|}{\sum_{z \in L} |\bar{S}_z|}, \forall i \in N$$

Proof. The independence assumptions in the noise model made the likelihood expression separable voter-wise such that for any voter $i \in \mathcal{N}$:

$$\arg \max_p \mathcal{L}(A, S, p, q, t) = \arg \max_p \prod_{z \in L} p^{|A_i^z \cap S_z|} (1 - p)^{|\bar{A}_i^z \cap S_z|}$$

so, applying the log to the expression, we see that it suffices to maximize:

$$h(p) = \sum_{z=1}^L |A_i^z \cap S_z| \log p + |\bar{A}_i^z \cap S_z| \log(1 - p)$$

whose derivative reads for all $p \in (0, 1)$:

$$h'(p) = \sum_{z=1}^L |A_i^z \cap S_z| \frac{1}{p} - |\overline{A_i^z} \cap S_z| \frac{1}{1-p}$$

We have that:

$$\begin{aligned} h'(p) = 0 &\iff \sum_{z=1}^L |A_i^z \cap S_z| (1-p) - |\overline{A_i^z} \cap S_z| p = 0 \\ &\iff \left(\sum_{z=1}^L |A_i^z \cap S_z| \right) (1-p) = \left(|\overline{A_i^z} \cap S_z| \right) p \\ &\iff p = \frac{\sum_{z=1}^L |A_i^z \cap S_z|}{\sum_{z=1}^L |A_i^z \cap S_z| + |\overline{A_i^z} \cap S_z|} \\ &\iff p = \frac{\sum_{z=1}^L |A_i^z \cap S_z|}{\sum_{z \in L} |S_z|} \\ &\iff p = \hat{p}_i \end{aligned}$$

The derivative vanishes in a single point in $(0, 1)$ and $\lim_{p \rightarrow 0} h(p) = \lim_{p \rightarrow 1} h(p) = -\infty$ thus h reaches a unique maximum.

We proceed exactly in the same way for the estimation of q_i . □

4.3.3 Alternating Maximum Likelihood Estimation

Now the estimation of the ground truths and that of the parameters are intertwined to maximize the overall likelihood $\mathcal{L}(A, S, p, q, t)$ by the *Alternating Maximum Likelihood Estimation algorithm*. AMLE is an iterative procedure similar to the *Expectation-Maximization* procedure introduced in [Baharad et al., 2011] but with two maximization steps per iteration, whose aim is to intertwinedly estimate the voter reliabilities, the alternatives' prior parameters and the instances' ground truths. The idea behind this estimation consists in alternating a MLE of the ground truths given the current estimate of the parameters, and an updating of these parameters via a MLE based on the current estimate of the ground truths.¹ Each of these steps have been discussed in the previous subsections and are now incorporated into Algo. 1.

The algorithm continues to run until a convergence criterion is met in the form of a bound on the norm of the change in the parameters' estimations. In practice we chose ℓ_∞ , but any other norm could be used in Algorithm 1 as in finite dimensions, all norms are equivalent (if a sequence converges according to one norm then it does so for any norm).

We define the vector of parameters $\hat{\theta}^{(v)} = (\hat{p}^{(v)}, \hat{q}^{(v)}, \hat{t}^{(v)})$ containing the voters' estimated noise parameters as well as the prior information estimated parameters at iteration v . In particular $\hat{\theta}^{(0)}$

¹In case of ties between subsets when estimating the ground truth, a tie-breaking priority over subsets is used. No ties occurred in our experiments.

Algorithm 1 AMLE procedure

Input: Approval ballots $(A_i^z)_{1 \leq z \leq L, i \in N}$
 Initial parameters $\hat{\theta}^{(0)}$, Bounds (l, u) , Tolerance ε

Output: Estimations $(\hat{S}_z), (\hat{p}_i, \hat{q}_i), (\hat{t}_j)$

repeat

for $z = 1 \dots L$ **do**

 Compute $\hat{S}_z^{(v+1)} = \{a_1, \dots, a_k\}$ with $k \in [l, u]$ and:

$$\begin{cases} |\hat{S}_z^{(v+1)} \cap S_{max,z}^{(v)}| &= \min(u, k_{max,z}^{(v)}) \\ |\hat{S}_z^{(v+1)} \cap S_{min,z}^{(v)}| &= \max(0, l - k_{tie,z}^{(v)} - k_{max,z}^{(v)}) \end{cases}$$

end for

for $i = 1 \dots N$ **do**

 Update the parameters (p_i, q_i) given $\hat{S}^{(v+1)}$:

$$\hat{p}_i^{(v+1)} = \frac{\sum_{z \in L} |A_i^z \cap \hat{S}_z^{(v+1)}|}{\sum_{z \in L} |\hat{S}_z^{(v+1)}|}, \hat{q}_i^{(v+1)} = \frac{\sum_{z \in L} |A_i^z \cap \overline{\hat{S}_z^{(v+1)}}|}{\sum_{z \in L} |\overline{\hat{S}_z^{(v+1)}}|}$$

end for

for $j = 1 \dots m$ **do**

 Update $\hat{t}_j^{(v+1)}$ by:

$$\hat{t}_j^{(v+1)} = \frac{occ^{(v+1)}(j) \overline{\alpha}_j^{(v+1)}}{occ^{(v+1)}(j) \overline{\alpha}_j^{(v+1)} + (L - occ^{(v+1)}(j)) \underline{\alpha}_j^{(v+1)}}$$

where :

$$\begin{cases} occ^{(v+1)}(j) &= \sum_{z=1}^L \mathbb{1}\{a_j \in \hat{S}_z^{(v+1)}\} \\ \overline{\alpha}_j^{(v+1)} &= \beta((l-1)^+, u-1, \hat{t}_{<j}^{(v+1)}, \hat{t}_{>j}^{(v)}) \\ \underline{\alpha}_j^{(v+1)} &= \beta(l, u, \hat{t}_{<j}^{(v+1)}, \hat{t}_{>j}^{(v)}) \end{cases}$$

end for

until $\|\hat{\theta}^{(v+1)} - \hat{\theta}^{(v)}\| \leq \varepsilon$

is the input initial values. The choice of the exact initial values depends on the application at hand.

Note that at convergence, only local optimality is guaranteed.

Theorem 4.4. *For any initial values $\hat{\theta}^{(0)}$, AMLE converges to a fixed point after a finite number of iterations.*

Proof. First we have by Theorem 4.1 that $\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)}) = \max_{S \in \mathcal{S}} \mathcal{L}(A, S, \hat{\theta}^{(v)})$, and we have in particular that:

$$\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)}) \geq \mathcal{L}(A, \hat{S}^{(v)}, \hat{\theta}^{(v)})$$

To prove that $\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v+1)}) \geq \mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)})$ we use the fact that we update (p, q, t) by their MLE. By Proposition 4.3 we have that:

$$(\hat{p}^{(v+1)}, \hat{q}^{(v+1)}) = \arg \max_{(p, q)} \mathcal{L}(A, \hat{S}^{(v+1)}, p, q, \hat{t}^{(v)})$$

Also by Proposition 4.2, and since we apply it sequentially to update t_j we have:

$$\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v+1)}) \geq \mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)})$$

To prove convergence, it suffices to show that $\hat{S}^{(v)} = \hat{S}^{(v+1)}$ for some v (which guarantees the estimators staying unchanged hereafter). Notice that the ground truth has a finite number of possible values (exactly 2^{mL}), leading the algorithm to cycle at some iteration. For the sake of simplicity, suppose that this cycle is of length 2, in other words, suppose that $\hat{S}^{(v+2)} = \hat{S}^{(v)}$ for some v ; this also implies that $\hat{\theta}^{(v+2)} = \hat{\theta}^{(v)}$. So:

$$\mathcal{L}(A, \hat{S}^{(v)}, \hat{\theta}^{(v)}) = \mathcal{L}(A, \hat{S}^{(v+2)}, \hat{\theta}^{(v+2)}) \geq \mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)})$$

By optimality of $\hat{S}^{(v+1)}$, we have also that:

$$\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)}) \geq \mathcal{L}(A, \hat{S}^{(v)}, \hat{\theta}^{(v)})$$

Hence, we get that:

$$\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)}) = \mathcal{L}(A, \hat{S}^{(v)}, \hat{\theta}^{(v)})$$

and thus, $\hat{S}^{(v+1)} = \hat{S}^{(v)} = \arg \max_{S \in \mathcal{S}_{l,u}} \mathcal{L}(A, S, \hat{\theta}^{(v)})$ and the estimators will remain the same after any number of iterations following v . \square

Because $\mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v+1)}) \geq \mathcal{L}(A, \hat{S}^{(v+1)}, \hat{\theta}^{(v)}) \geq \mathcal{L}(A, \hat{S}^{(v)}, \hat{\theta}^{(v)})$, the likelihood increases at each step of the algorithm. This guarantees that whenever the execution stops, the likelihood is closer to the maximum than it initially was. Therefore the algorithm can not only be run until convergence, but it can also be run as an anytime algorithm.

Example 4.2. *Take $n = 3$, $m = 5$, $l = 1$, $u = 2$, $L = 4$, and the following profile and initial parameters:*

$$\begin{cases} \hat{p}_1^{(0)} = 0.5 & \hat{p}_2^{(0)} = 0.5 & \hat{p}_3^{(0)} = 0.5 \\ \hat{q}_1^{(0)} = 0.44 & \hat{q}_2^{(0)} = 0.41 & \hat{q}_3^{(0)} = 0.32 \\ \hat{t}_1^{(0)} = \dots = \hat{t}_5^{(0)} = 0.5 \end{cases}$$

	A^1	A^2	A^3	A^4
Voter 1	$\{a_1, a_4\}$	$\{a_1\}$	$\{a_3\}$	$\{a_1\}$
Voter 2	$\{a_2\}$	$\{a_5\}$	$\{a_4\}$	$\{a_1\}$
Voter 3	$\{a_2, a_3, a_4\}$	$\{a_2, a_3, a_5\}$	$\{a_2, a_3\}$	$\{a_3\}$

Estimating the ground truth: *The first step is the application of Theorem 4.1 to estimate the ground truth of the instances given the initial parameters, yielding*

$$\hat{S}_1^{(1)} = \{a_2, a_4\}, \hat{S}_2^{(1)} = \{a_2, a_5\}, \hat{S}_3^{(1)} = \{a_2, a_3\}, \hat{S}_4^{(1)} = \{a_1, a_3\}$$

Estimating the voter reliabilities: *In the next step we use these estimates of the ground truths to compute the MLEs of the voter reliabilities. For instance:*

- voter 1 has 2 false positive labels from a total of 12 negative labels so $\hat{q}_1^{(1)} = \frac{2}{12} = 0.17$.
- voter 1 also has 3 true positive labels out of 8 positive ones so $\hat{p}_1^{(1)} = \frac{3}{8} = 0.38$.

In the end, we get:

$$\begin{cases} \hat{p}_1^{(1)} = 0.38 & \hat{p}_2^{(1)} = 0.38 & \hat{p}_3^{(1)} = 0.88 \\ \hat{q}_1^{(1)} = 0.17 & \hat{q}_2^{(1)} = 0.08 & \hat{q}_3^{(1)} = 0.17 \end{cases}$$

Estimating the prior parameters: *The final step of this iteration consists in updating the estimations of the prior parameters by applying Proposition 4.2 sequentially. First we estimate $\hat{t}_1^{(1)}$ given $\hat{S}^{(1)}$ and $\hat{t}_2^{(0)}, \dots, \hat{t}_5^{(0)}$ by maximum likelihood estimation. We first compute $\bar{\alpha}_1, \underline{\alpha}_1$ and $\text{occ}(a_1)$:*

$$\begin{cases} \bar{\alpha}_1 & = \beta(0, 1, t_2, \dots, t_5) = 0.3125 \\ \underline{\alpha}_1 & = \beta(1, 2, t_2, \dots, t_5) = 1 \\ \text{occ}(a_1) & = 1 \end{cases}$$

Then the maximum likelihood estimation of t_1 is:

$$\hat{t}_1 = \frac{\text{occ}(a_1)\bar{\alpha}_1}{(L - \text{occ}(a_1))\underline{\alpha}_1 + \text{occ}(a_1)\bar{\alpha}_1} = 0.09$$

The next steps are to estimate $\hat{t}_2^{(1)}$ given $\hat{t}_1^{(1)}, \hat{t}_3^{(0)}, \hat{t}_4^{(0)}, \hat{t}_5^{(0)}$ and so on. Finally, we get:

$$\hat{t}_1^{(1)} = 0.09, \hat{t}_2^{(1)} = 0.56, \hat{t}_3^{(1)} = 0.28, \hat{t}_4^{(1)} = 0.14, \hat{t}_5^{(1)} = 0.20$$

Fix $\varepsilon = 10^{-5}$. We repeat all steps until convergence (according to ℓ_∞), after 5 full iterations. In the fixed point, the estimations of the ground truths are:

$$\hat{S}_1 = \{a_2, a_3\}, \hat{S}_2 = \{a_2, a_3\}, \hat{S}_3 = \{a_2, a_3\}, \hat{S}_4 = \{a_3\}$$

4.4 Experiments

4.4.1 Experiment Design and Data Collection

We designed an image annotation task as a football quiz.² We selected 15 pictures taken during different matches between two of the following teams: Real Madrid, Inter Milan, Bayern Munich, Barcelona, Paris Saint-Germain. In each picture, it may be the case that players from both teams appear, or players from only one team, therefore $l = 1$ and $u = 2$. Each participant is shown the instances one by one, and is each time asked to select all the teams she can spot (see Figure 4.3).

The images' saturation and hue were edited in order to make the task less obvious (see the before and after versions in Figure 4.1 and Figure 4.2 respectively ³).

We designed a simple incentive for participants, consisting in ranking them according to the following principle:

- The participants get one point whenever their answer contains all correct alternatives for a picture. They are then ranked according to their cumulated points.
- To break ties, the participant who selected a smaller number of alternatives overall is ranked first.

We gathered the answers of 76 participants (only two of them spammed by simply selecting all the alternatives). We show how the top participants are ranked in Figure 4.4

4.4.2 Descriptive Analysis of Collected Answers

To see how the participants behave given the ranking incentives that we defined in the football quiz, we plotted the histogram of the sizes of the answers (see Figure 4.5). It appears that although the platform enables to select every alternative, only two voters did so for all the questions. Moreover, figures 4.5b and 4.5a show that the majority of the voters tend to select exactly the number of teams that appear in an image.

4.4.3 Anna Karenina's Initialization

Inspired by the *Anna Karenina Principle* in [Meir et al., 2019], we devised an initialisation strategy for the voters' reliabilities. In his book, Leo Tolstoy stated that "Happy families are all alike; every unhappy family is unhappy in its own way". In the same spirit, it seems reasonable to make the hypothesis that accurate users tend to make similar answers, whereas inaccurate users have each their own way of being inaccurate.

²The annotations dataset and the code are available at:
<https://github.com/taharallouche/Football-Quiz-Crowdsourcing>.

³Note that the difference is only visible on a coloured printout

CHAPTER 4. MULTI-WINNER EPISTEMIC VOTING: COMMITTEE-SIZE AS PRIOR KNOWLEDGE

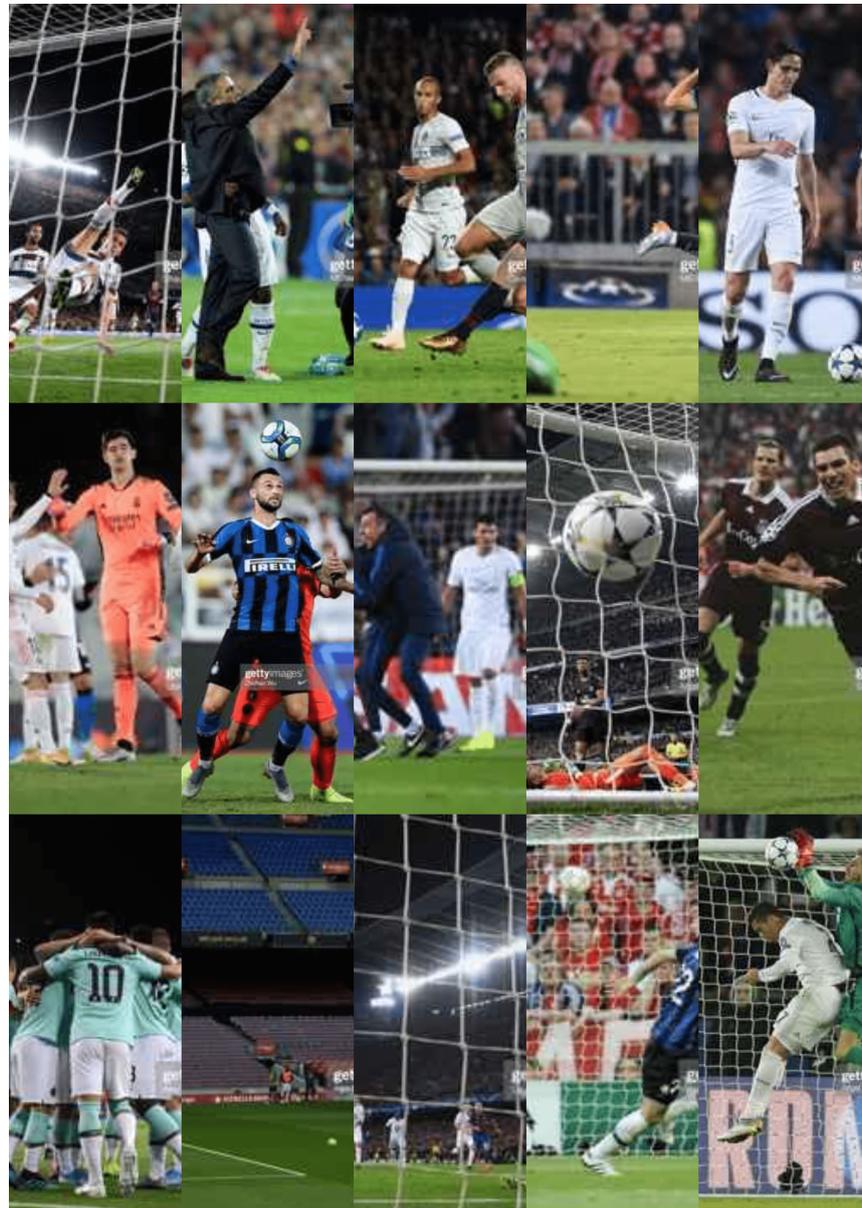


Figure 4.1: The images to be labeled before editing



Figure 4.2: The images to be labeled after editing



Figure 4.3: Example of Annotation Task

~ Ranking ~


 UEFA
 CHAMPIONS
 LEAGUE

Rank	Participant	Score	Nb. of selected teams
1	Nedjou	14	29
2	Nitze_jr	14	62
3	yann	14	74
4	Yhafsi	13	29
5	Momo	13	29
6	Skander	12	29
7	Florian	12	53
8	ch7	11	33

Figure 4.4: Top quiz scores

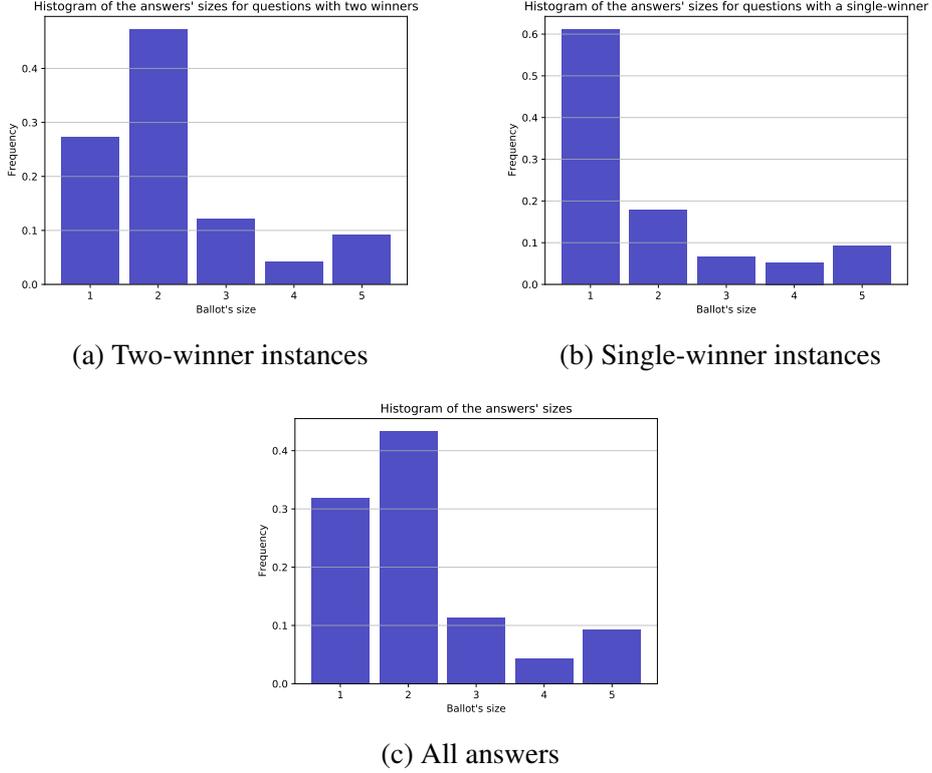


Figure 4.5: Histogram of the ballots' sizes

We use the following heuristic (see Algorithm 2) for the initialization. We used the Jaccard distance given by:

$$d_{Jacc}(A, B) = \frac{|\bar{A} \cap B| + |A \cap \bar{B}|}{|A \cup B|}$$

Algorithm 2 Initializing $(p_i, q_i)_i$

Input: Approval ballots $(A_i^z)_{z,i}$

Output: Initialization $(\hat{p}_i^{(0)}, \hat{q}_i^{(0)})$

- Compute $w_{max} = \frac{n}{1+n}, w_{min} = \frac{1}{1+n}$
 - Compute $d_i = \sum_{j \neq i} d_{Jacc}(A_i, A_j)$
 - Compute $d_{max} = \max d_i, d_{min} = \min d_i$
 - Compute $w_i = (w_{max} - w_{min}) \left(\frac{\frac{1}{d_i} - \frac{1}{d_{max}}}{\frac{1}{d_{min}} - \frac{1}{d_{max}}} \right) + w_{min}$
 - Fix $\hat{p}_i^{(0)} = \frac{1}{2}$ and $\hat{q}_i^{(0)} = \frac{1 - e^{w_i - 1}}{e^{w_i - 1} + 1}$
-

Remark. The formulas in Algorithm 2 guarantee that a voter's parameters $(\hat{p}_i^{(0)}, \hat{q}_i^{(0)})$ are such that her initial weight is equal to w_i , and that $\frac{w_{max}}{w_{min}} = n$ which means that initially, the voter

closest in average to the other voters counts n times the voter with biggest average distance.

Example 4.3. Consider following the approval profile (Table 4.1) for 3 voters, 5 alternatives and 4 Instances. Here we have that:

	A^1	A^2	A^3	A^4
Voter 1	$\{a_1, a_4\}$	$\{a_1\}$	$\{a_3\}$	$\{a_1\}$
Voter 2	$\{a_2\}$	$\{a_5\}$	$\{a_4\}$	$\{a_1\}$
Voter 3	$\{a_2, a_3, a_4\}$	$\{a_2, a_3, a_5\}$	$\{a_2, a_3\}$	$\{a_3\}$

Table 4.1: Approval Ballots of 3 Voters on 4 Instances

$$w_{max} = \frac{n}{n+1} = 0.75, w_{min} = \frac{1}{n+1} = 0.25$$

First, compute the mean Jaccard distance of all voters: $d_1 = 1.71, d_2 = 1.69, d_3 = 1.65$. So $d_{max} = d_1 = 1.71$ and $d_{min} = d_3 = 1.65$, which means that voter 3 (the closest in average to all the voters) will get the biggest weight $w_3 = w_{max} = 0.75$ and voter 1 gets the smallest weight $w_1 = w_{min}$. Next, compute the weight that will be assigned to each voter, for instance:

$$w_2 = (w_{max} - w_{min}) \frac{\frac{1}{d_2} - \frac{1}{d_{max}}}{\frac{1}{d_{min}} - \frac{1}{d_{max}}} + w_{min} = 0.38$$

Now we can set the initial values for the reliability parameters accordingly:

$$\hat{p}_2^{(0)} = \frac{1}{2}, \hat{q}_2^{(0)} = \frac{1 - \frac{e^{w_2} - 1}{e^{w_2} + 1}}{2}$$

We can check that these parameters are such that:

$$\ln \left[\frac{p_2(1 - q_2)}{q_2(1 - p_2)} \right] = w_2$$

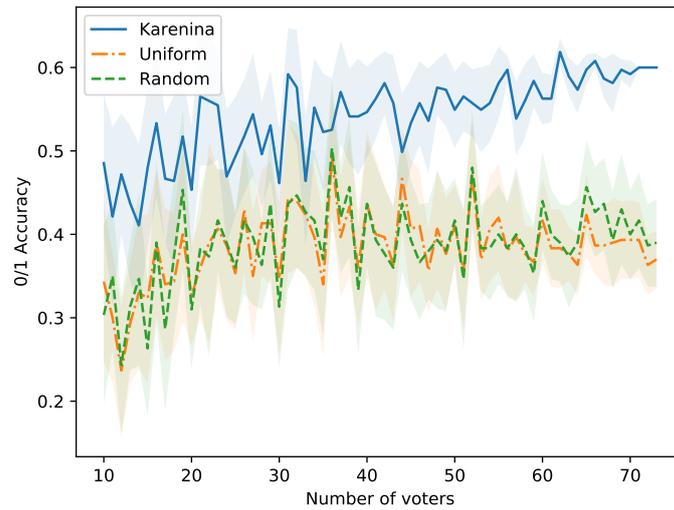
After proceeding in the same fashion with all the voters, we get the initial parameters:

$$\begin{cases} \hat{p}_1^{(0)} = 0.5 & \hat{p}_2^{(0)} = 0.5 & \hat{p}_3^{(0)} = 0.5 \\ \hat{q}_1^{(0)} = 0.44 & \hat{q}_2^{(0)} = 0.41 & \hat{q}_3^{(0)} = 0.32 \end{cases}$$

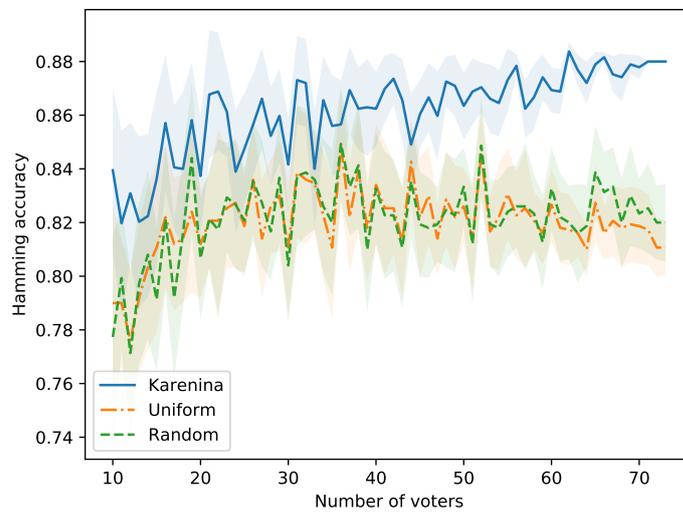
Since the AMLE only guarantees convergence to a local maximum, which makes the result depending on the initial point, we compared the results of this initialization (Anna Karenina) to other procedures to motivate its choice, see Figure 4.6, namely we tested:

- Uniform weights: Initially all the voters in the batch are given the same weight.
- Random weights: Initially, for each voter in the batch, p_i is randomly picked from $(0.5, 1)$ and q_i is randomly picked from $(0, 0.5)$.

We can notice that these two baseline procedures show very similar performances, and that they are both outperformed by the Anna Karenina initialization.



(a) 0-1 accuracy



(b) Hamming accuracy

Figure 4.6: Accuracies of different initializations

4.4.4 Time Complexity of AMLE

We assessed the execution time of the AMLE algorithm with and without constraints (referred to as AMLE and $AMLE_f$), run on Intel Core i7-10610U CPU @1.80Ghz 4 cores, 8 threads and 32Gb RAM. Results are shown in Figure 4.7. We can see that whereas the number of iterations does not seem to grow as the number of voters increases, the execution time of AMLE does, especially around 40 voters.

4.4.5 Results

4.4.5.1 Hamming and 0-1 Subset Accuracies

To assess the importance of prior information on the size of the ground truth, we tested the AMLE algorithm with free bounds $(l, u) = (0, m)$ (will be referred to as $AMLE_f$) and the $AMLE_c$ algorithm with $(l, u) = (1, 2)$.

We also apply the modal rule [Caragiannis et al., 2020] which outputs the subset of alternatives that most frequently appears as an approval ballot

$$\arg \max_{S \in \mathcal{S}} |i \in N, S = A_i|$$

and a variant of label-wise majority rule which outputs the subset of alternatives S such that:

$$a \in S \iff |i \in N, a \in A_i| > \frac{n}{2}$$

If this subset is empty it is replaced by the alternative with highest approval count, and if it has more than two alternatives then we only keep the top-2 alternatives.

We took 20 batches of $n = 10$ to $n = 74$ randomly drawn voters and applied the four methods to all of them (see Figure 4.8a, 4.8b).

As classically done in the literature [Nguyen et al., 2020], we use the Hamming accuracy:

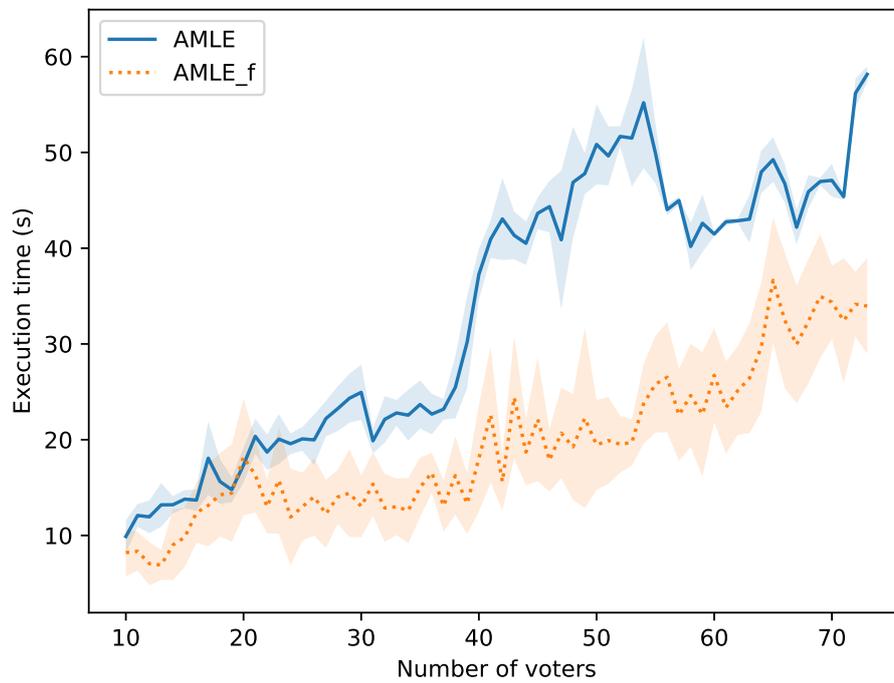
$$\frac{1}{mL} \sum_{z=1}^L \left[|S_z^* \cap \hat{S}^z| + |\overline{S}_z^* \cap \overline{\hat{S}}^z| \right]$$

and the 0/1 accuracy:

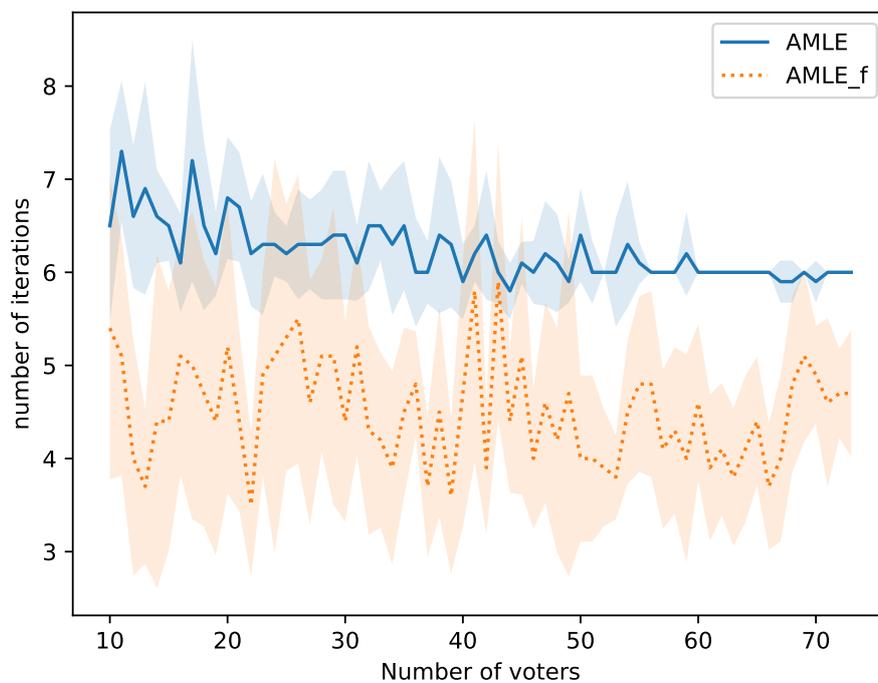
$$\frac{1}{L} \sum_{z=1}^L \mathbb{1}\{S_z^* = \hat{S}^z\}$$

as metrics and report their 0.95 confidence intervals.

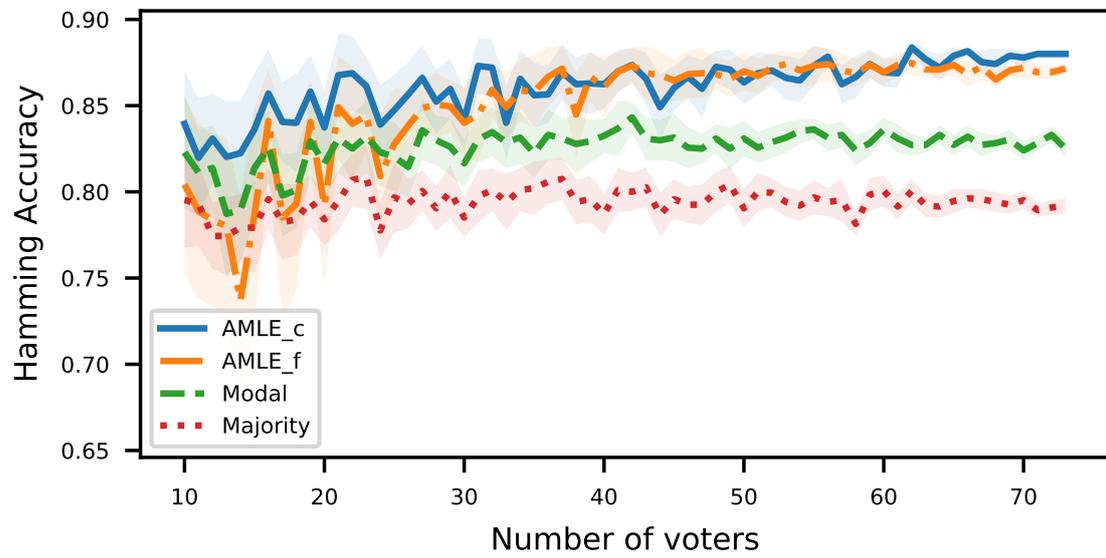
We notice that the majority and the modal rule are outperformed by AMLE, which can be explained by the fact that they do not take into account the voters' reliabilities. Comparing the performances of $AMLE_c$ and $AMLE_f$ emphasizes the importance of the prior knowledge on the committee size to improve the quality of the estimation.



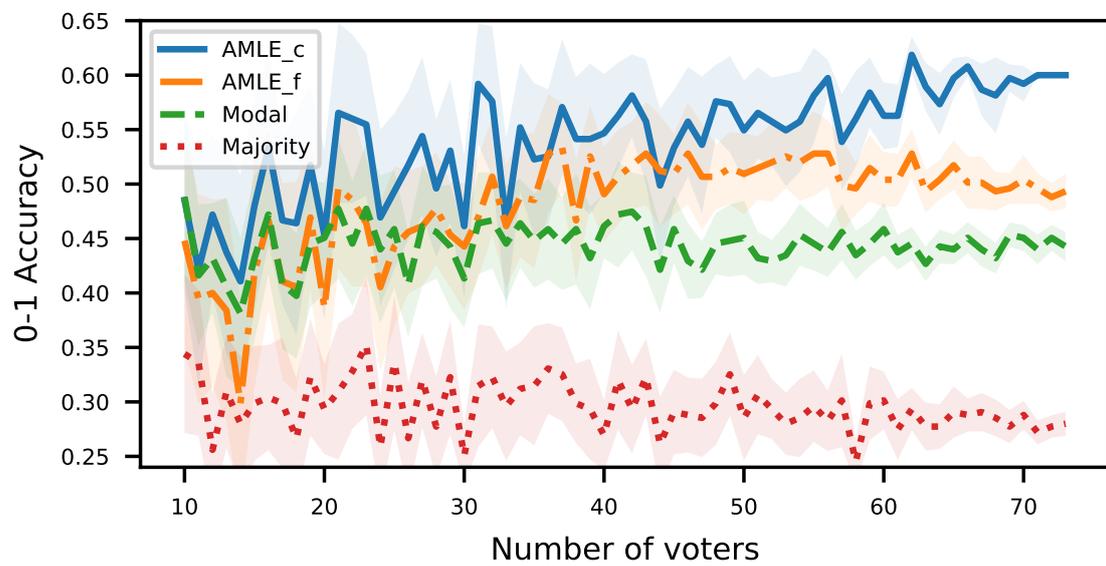
(a) Execution time of AMLE



(b) Number of iterations of AMLE



(a) Hamming accuracy



(b) 0/1 accuracy

Figure 4.8: Accuracies of different aggregation methods

4.4.5.2 Harmonic Loss

In addition to the Hamming and 0-1 subset accuracies, we introduced a new metric which can be considered as an intermediate one. The Hamming metric considers each label independently and the 0-1 subset loss considers them jointly in a strict fashion, whereas the harmonic accuracies that we introduced considers all the instance's labels jointly but with different convex weights depending on the number of correctly predicted ones:

$$T(S, S^*) = \sum_{k=1}^{|S \cap S^*|} \frac{1}{6-k}$$

So out of the 5 labels:

- if 0 labels are correct then $T = 0$.
- if 1 labels is correct then $T = \frac{1}{5}$.
- if 2 labels are correct then $T = \frac{1}{5} + \frac{1}{4}$.
- if 3 labels are correct then $T = \frac{1}{5} + \frac{1}{4} + \frac{1}{3}$.
- if 4 labels are correct then $T = \frac{1}{5} + \frac{1}{4} + \frac{1}{3} + \frac{1}{2}$.
- if 5 labels are correct then $T = \frac{1}{5} + \frac{1}{4} + \frac{1}{3} + \frac{1}{2} + 1$.

Defined as such, this accuracy favours the estimators that are able to correctly estimate most of the instance's labels without being as rigid as the 0-1 subset accuracy.

This metric is reminiscent of the Proportional Approval Voting rule for multi-winner elections, which defines the score of a subset of candidates W for a voter as $1 + \frac{1}{2} + \dots + \frac{1}{j}$, where j is the number of candidates in W approved by the voter. We could consider more generally a class of metrics defined by a vector \vec{w} , such that $T(S, S^*) = w_{|S \cap S^*|}$. This class generalizes Hamming, 0-1 and Harmonic and is reminiscent of the class of *Thiele* rules (see for instance [Lackner and Skowron, 2020b] for an extended presentation of multi-winner approval-based committee rules).

In Figure 4.9 we show the evolution of the Harmonic accuracies when the number of randomly picked voters in each batch increase.

4.4.5.3 Comparing the accuracies

We show in Table 4.2 the accuracies of the considered methods when applied to the entire annotation dataset. The results show that AMLE outperforms the baselines according to the three metrics. Moreover, when AMLE incorporates the prior knowledge, its accuracy increases.

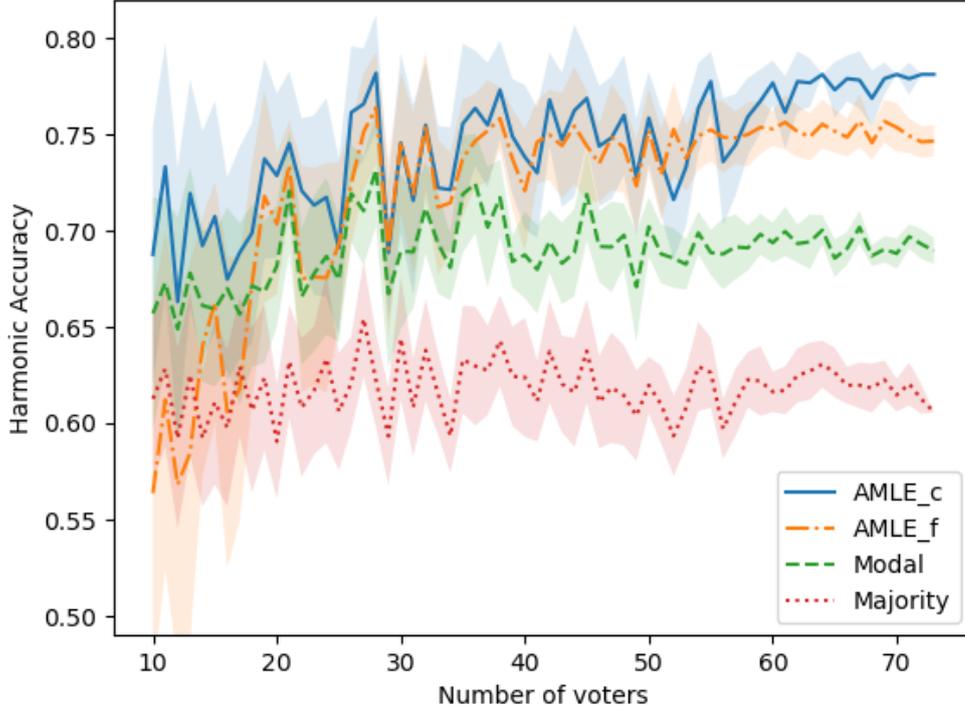


Figure 4.9: Normalized Harmonic accuracy

4.5 The Case of Fixed Committee-size: Extension of Theorem 3.4 to Multi-winner Settings

In this section, we extend the result presented in Theorem 3.4 to the classical multi-winner case with fixed committee-size k .

4.5.1 Noise Model and Estimations

We suppose that the ground truth is a set of k alternatives $S^* \subseteq \mathcal{S}_k = \{S \subseteq \mathcal{X}, |S| = k\}$. The input still consists of approval ballots $A = (A_1, \dots, A_n)$. The approvals are generated from the following noise model:

$$P(a \in A_i | S^* = S) = \begin{cases} p_i & \text{if } a \in S \\ 1 - p_i & \text{if } a \notin S \end{cases}$$

Theorem 4.5. For $m \geq 2$ alternatives and for a committee-size $k \geq 1$, we have that:

$$\mathbb{E}_p[|A_i|] = (m - k) - (m - 2k)p$$

	AMLE _c	AMLE _f	Modal	Majority
Hamming	0.88	0.86	0.84	0.80
Harmonic	0.78	0.74	0.69	0.61
0/1	0.60	0.53	0.46	0.26

Table 4.2: Hamming and 0/1 accuracy for entire dataset

Just like in the single-winner case, this formula expressing the expected size of the ballot as a linear function of the reliability p would enable us to have an estimation of this latter by only observing the ballot, not needing any further information about the ground truth.

Proof. Let $S \in \mathcal{S}_k$:

$$\begin{aligned}
 \mathbb{E}_p[|A_i| | S^* = S] &= \mathbb{E}\left[\sum_{b \in \mathcal{X}} \mathbb{1}\{b \in A_i\} | S^* = S\right] \\
 &= \sum_{b \in \mathcal{X}} P(b \in A_i | S^* = S) \\
 &= \sum_{b \in S} P(b \in A_i | S^* = S) + \sum_{b \notin S} P(b \in A_i | S^* = S) \\
 &= k \times p + (m - k)(1 - p) \\
 &= (m - k) - (m - 2k)p
 \end{aligned}$$

Thus we have that:

$$\begin{aligned}
 \mathbb{E}_p[|A_i|] &= \sum_{S \in \mathcal{S}_k} \mathbb{E}_p[|A_i| | S^* = S] P(S^* = S) \\
 &= \sum_{S \in \mathcal{S}_k} [(m - k) - (m - 2k)p] P(S^* = S) \\
 &= (m - k) - (m - 2k)p
 \end{aligned}$$

□

Remark. We recover the same expression as in Theorem 3.4 in the single-winner case $k = 1$.

4.5.2 Experiments

In the same fashion as in Section 3.5, we test this size-related weighting technique on real image annotations. We used a subset of the football quiz dataset. We only kept instances where the photo displays players from exactly two teams. Thus we have that $k = 2$.

We estimated the reliability parameter p_i^z of each voter for each single instance using Theorem 4.5 as follows:

$$\hat{p}_i^z = \text{proj}_{[\varepsilon, 1-\varepsilon]} \frac{m - k - |A_i^z|}{m - 2k}$$

The projected quantity is simply the maximum likelihood estimation of p_i^z with a single sample (the actual observation of the voter’s ballot). We project it into a closed interval to avoid having $\hat{p}_i^z = 1$ or $\hat{p}_i^z = 0$. The voters are then weighed accordingly, that is:

$$w_i^z = \ln \frac{\hat{p}_i^z}{1 - \hat{p}_i^z}$$

and we output the k alternatives with highest weighted approval scores.

We compared this method to two baseline multi-winner voting rules namely:

- Top-2: This rule simply outputs the two alternatives having the two greatest number of approvals.
- Modal: This rule selects the ballot that was casted by the greatest number of voters.

The results are shown in Figure 4.10.

We observe that, just like the other two baseline rules, the size-decreasing rule has a tendency to be more accurate, for both accuracy metrics, as the number of considered voters increases. It is also clear that it significantly outperforms the baseline.

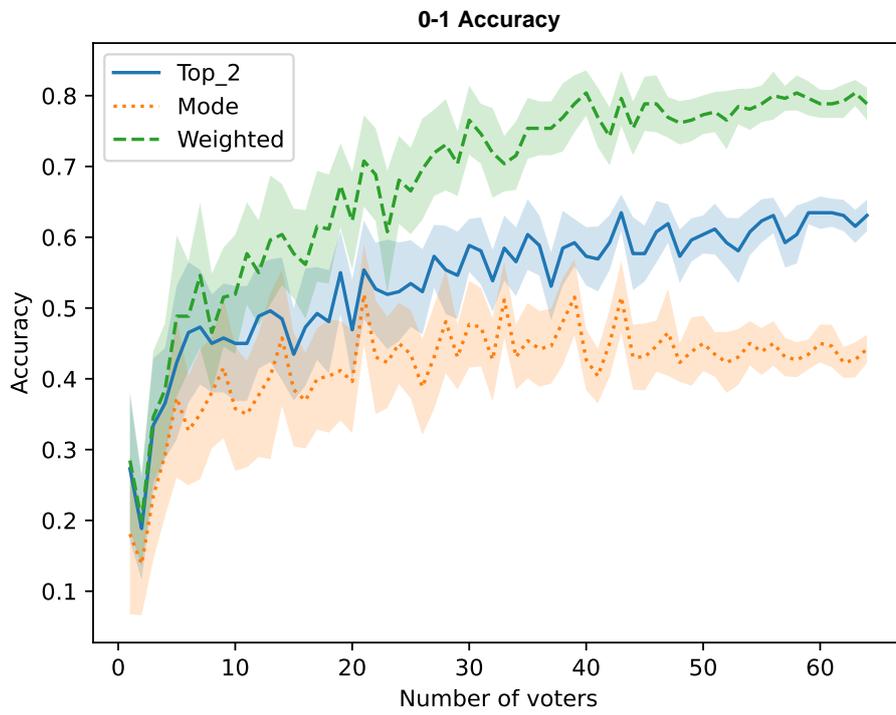
4.6 Conclusion

We study multi-winner approval voting from an epistemic point of view. We propose a noise model that incorporates the prior belief about the size of the ground truth. Then we derive an iterative algorithm to intertwinedly estimate the ground truth labels, the voter noise parameters and the prior belief parameters and we prove its convergence. Our algorithm is based on a simplification of Expectation-Maximization (EM), and its simple steps are more easily explainable to voters than EM and other similar statistical learning approaches.

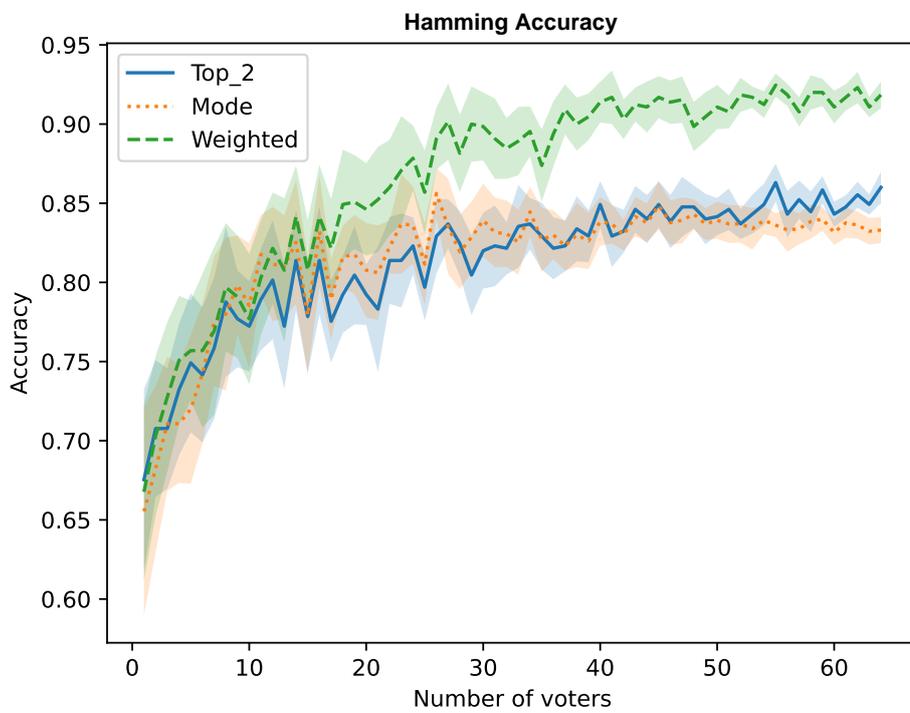
We also generalized the *size matters* principle to multi-winner situations with a fixed ground truth size. We showed how to estimate the voters’ reliabilities by merely observing the size of their ballots, and tested it on our image annotation dataset to prove that it outperforms the baselines.

Although we mainly considered a general multi-instance task that fits the collective annotation framework, where each voter answers several questions on the same set of alternatives, we can nonetheless apply the same algorithm to single-instance problems (such as the allocation of scarce medical resources) where only one question is answered. In this case, the prior parameters cannot be updated and it suffices to fix them once and for all and alternate between the estimation of the ground truth and the voter parameters.

In some contexts (*e.g.*, patients in a hospital), alternatives and votes are not observed at once but streamed. To cope with this online setup we consider extending our AMLE algorithm in the spirit of [Cappé and Moulines, 2009].



(a) 0/1 accuracy



(b) hamming accuracy

Multi-winner Epistemic Voting: Committee-size as Constraint

Abstract

Epistemic social choice considers votes as noisy estimates of a ground truth. In some contexts the truth consists of a set of objective winners. Whilst the usual aim of epistemic voting rules is the uncovering of the whole ground truth, we focus here on cases where some exogenous size constraints bears upon the number of winners in the output even though the size of the ground truth may lie outside this interval. The prototypical illustration of such cases is a master program's selection committee deciding on accepting or rejecting a number of applicants, based on their academic excellence, with a limited number of seats. We introduce two solution concepts: the *most likely admissible utility maximizer* and the *admissible expected utility maximizer*, that we test on synthetic data.

5.1 Framework and Objective

5.1.1 Votes, Ground Truth and External Constraints

Consider a set of n voters $N = \{1, \dots, n\}$ and a set of m alternatives $\mathcal{X} = \{a_1, \dots, a_m\}$. Among these alternatives, there exists an *unknown* subset of *objectively good/eligible* alternatives (patients really in need of medical resources, students admissible to a program, non-defective products that can already be put on sale, promising job applicants in a shortlisting phase [..]). Voters are asked to approve all the alternatives that they think are eligible, which yields the approval profile $A = (A_1, \dots, A_n)$.

We suppose in this chapter that due to some external requirements (limited resources, client order satisfaction..), the number of alternatives to be selected in the output, once the approval ballots of the voters have been aggregated, ought to be bounded by a given lower bound l and a given

upper bound u . We say that a subset of alternatives is *admissible* if and only if satisfies these size constraints. The annex problem of approval-based shortlisting of candidates has been studied in [Lackner and Maly, 2021] with a classical (non-epistemic) voting approach.

The goal is hence to aggregate the votes in order to select the best *admissible* subset of alternatives. To formally define what we mean by “best”, we define the following notion of utility.

5.1.2 Utility

Let $\alpha^+, \alpha^-, \beta^+, \beta^- \in \mathbb{R}$ such that $\alpha^+ > \alpha^-$ and $\beta^+ < \beta^-$. For a ground truth subset S^* , we define the utility of any subset of alternatives S as:

$$u(S|S^*) = \sum_{a \in S} u^+(a|S^*) + \sum_{a \notin S} u^-(a|S^*)$$

where:

$$u^+(a|S^*) = \begin{cases} \alpha^+ & \text{if } a \in S^* \quad \text{True Positive} \\ \beta^+ & \text{if } a \notin S^* \quad \text{False Positive} \end{cases}, \quad u^-(a|S^*) = \begin{cases} \alpha^- & \text{if } a \in S^* \quad \text{False Negative} \\ \beta^- & \text{if } a \notin S^* \quad \text{True Negative} \end{cases}$$

So α^+, α^- can be interpreted as respectively the utility of selecting an alternative in S^* and the cost of not selecting it. From now on, we will denote $u(S)$ instead of $u(S|S^*)$ for ease of reading. Notice that:

$$u(S) = \sum_{a \in \mathcal{X}} u^-(a) + \sum_{a \in S} [u^+(a) - u^-(a)]$$

We define $f(a) = u^+(a) - u^-(a)$ which is equal to $\alpha = \alpha^+ - \alpha^- > 0$ if $a \in S^*$ and is equal to $-\beta = \beta^+ - \beta^- < 0$ if $a \notin S^*$. We extend the domain of f to the set of all possible subsets of alternatives $\mathcal{S} = 2^{\mathcal{X}}$ by defining $f(S) = \sum_{a \in S} f(a)$. We can notice that maximizing $u(S)$ is equivalent to maximizing $f(S)$.

So all in all, we want to find an admissible subset in $\mathcal{S}_{l,u} = \{S \subseteq \mathcal{X}, l \leq |S| \leq u\}$ which maximizes f . However, the challenge is that the utility cannot be directly computed since it depends on the ground truth S^* which is unknown. To solve this issue, we will define a noise model consisting of a prior distribution on S^* and conditional distribution on the approval ballots given the ground truth, which will be used later on to propose two solution concepts.

Remark. *This problem can be seen as a generalization of a knapsack problem to situations where the utility of each item is not a priori known, but is only accessible through the noisy reports of voters.*

5.1.3 The Noise Model

For all $1 \leq j \leq m$, let $t_j = P(f(a_j) = \alpha) = P(a_j \in S^*)$ the prior probability of a_j being eligible and suppose that the events $\{f(a_j) = \alpha\}$ are mutually independent. This assumption

is realistic because we do not have any prior on the size of S^* : so for instance the probability of student s being objectively above the bar does not change once we learn about other students being below or above the bar. Also let $A = (A_1, \dots, A_n)$ be the approval profile drawn according to the following distribution:

$$P(a \in A_i | S^* = S) = \begin{cases} p_i & \text{if } a \in S \\ q_i & \text{if } a \notin S \end{cases}$$

where p_i (resp. q_i) is the probability that a voter approves an objectively good (resp. bad) alternative. We also keep the following independence assumptions as in Chapter 4:

- (1) Voters' ballots are independent given the ground truth.
- (2) A voter's decisions to approve or not different alternatives are independent.

We suppose in this part that all the parameters $(p_i, q_i)_{i \in N}$ and $(t_j)_{j \leq m}$ are known. In practice this parameters can be:

- Fixed by a central decision maker: For instance when a jury decides on the admission of students to a master program, a central decision maker can choose to make the vote anonymous and assign common reliabilities (p, q) to all the voters, and to make the vote neutral she can assign equal prior t to all the candidates. This can also not be the case if she prefers to put more weight on the approvals of some voters or to favor some candidates.
- Estimated from past observations: For instance in applications where the vote is repeated through time. Consider a technical support center platform, where each time a client submits a question to a chatbot, different algorithms analyse the message and pick a set of suitable solutions, then their votes are aggregated into a certain number of final suggestions to be given to the client.

5.1.4 The Solution Concepts

Recall that we focus on finding solutions to estimate the admissible subset of alternatives maximizing f . In the following paragraphs we will introduce, study and compare two different solution concepts:

- The most likely admissible utility maximizer: among all the possible sets of winners that respect the size constraints, which one has the highest probability of maximizing the utility?
- The expected utility maximizer: which admissible set of winners maximizes the expected utility among all the possible sets that verify the size constraints?

Even though maximizing an expected utility is more common in decision making theory, we chose to scout another method consisting of taking the decision that is the most likely optimal among all possible actions. We will see that this method has some advantages like robustness to inexact utility values.

We will also see that there is a strong connection between this interpretation (committee-size as constraint) and the former one (committee-size as prior knowledge). We will show that in the case of equality constraints $l = u = k$, the two interpretations yield the same output, consisting of the top- k alternatives according to their weighted approval scores.

Recall the definitions of section 4.2, namely the weighted approval score:

$$app_w(a_j) = \ln\left(\frac{t_j}{1-t_j}\right) + \sum_{i:a_j \in A_i} \ln\left(\frac{p_i(1-q_i)}{q_i(1-p_i)}\right)$$

and the threshold

$$\tau_n = \sum_{i=1}^n \ln\left(\frac{1-q_i}{1-p_i}\right)$$

and its associated partition of \mathcal{X} into:

$$\begin{cases} S_{max}^{\tau_n} &= \{a \in A, app_w(a) > \tau_n\} \\ S_{tie}^{\tau_n} &= \{a \in A, app_w(a) = \tau_n\} \\ S_{min}^{\tau_n} &= \mathcal{X} \setminus (S_{max}^{\tau_n} \cup S_{tie}^{\tau_n}) \end{cases}$$

and let:

$$k_{max}^{\tau_n} = |S_{max}^{\tau_n}|, k_{tie}^{\tau_n} = |S_{tie}^{\tau_n}|, k_{min}^{\tau_n} = |S_{min}^{\tau_n}|$$

We will suppose, without loss of generality, that:

$$app_w(a_1) \geq \dots \geq app_w(a_m)$$

5.2 MLAUM: Most Likely Admissible Utility Maximizer

Let $M_{l,u} = \max_{S \in \mathcal{S}_{l,u}} f(S)$ be the unknown maximum value of f for the admissible sets of alternatives. For a given admissible set of alternatives S we define $\psi(S) = P(f(S) = M_{l,u} | A)$, which is the probability of S maximizing the function f . Our goal is to estimate $S \in \mathcal{S}_{l,u}$ with highest probability of $f(S) = M_{l,u}$, that is, we seek:

$$\arg \max_{S \in \mathcal{S}_{l,u}} P(f(S) = M_{l,u} | A) = \arg \max_{S \in \mathcal{S}_{l,u}} \psi(S)$$

We will now enumerate some results that will finally lead us into formulating the full algorithm of estimation of the MLAUM. The first result is a reformulation of Theorem 4.1 with a slight adaptation to this interpretation. It characterizes the admissible sets which are most likely to coincide with the ground truth (but as we will see later, they do not always coincide with the MLAUM, but are potential candidates).

Proposition 5.1. $\tilde{S} \in \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S | A)$ if and only if there exists $k \in [l, u]$ such that $\tilde{S} = \{a_1, \dots, a_k\}$ and:

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau_n}| &= \min(u, k_{max}^{\tau_n}) \\ |\tilde{S} \cap S_{min}^{\tau_n}| &= \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases}$$

Proof. The likelihood of an admissible subset $S \in \mathcal{S}_{l,u}$:

$$L(S) \propto \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j) \prod_{i=1}^n \left[\frac{1 - p_i}{1 - q_i} \right]^{|S|} \left[\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right]^{|A_i \cap S|}$$

Maximizing this likelihood coincides exactly with the optimization problem we solved in proof of Theorem 4.1. \square

See Example 4.1 for an illustration.

Now we will state two lemmas that will lead to the final result on the estimation of the maximizer of ψ . The first one expresses the conditional probability of some alternative being in the ground truth, given the votes, as an increasing function (precisely, a logistic function) of its weighted approval score. The second lemma explicits the probability of a given admissible set of alternatives S being a maximizer of the utility f , namely $\psi(S)$, as a function of the aforementioned alternatives' probabilities.

Lemma 5.2. *For any alternative $a_j \in \mathcal{X}$ denote by:*

$$A^j = (A_1^j, \dots, A_n^j) = (\mathbb{1}\{a_j \in A_1\}, \dots, \mathbb{1}\{a_j \in A_n\})$$

the answers of the voters concerning alternative a_j . Also denote:

$$P^j = P(a_j \in S^* | A^j) = \frac{P(A^j | a_j \in S^*) P(a_j \in S^*)}{P(A^j)}$$

the posterior probability of a_j being in the ground truth S^ . Then:*

$$P^j = \frac{\exp(\text{app}_w(a_j) - \tau_n)}{1 + \exp(\text{app}_w(a_j) - \tau_n)}$$

In particular, P^j increases with $\text{app}_w(a_j)$.

Proof. Let $a_j \in \mathcal{X}$. We have that:

$$\begin{aligned} P^j &= \frac{P(A^j | a_j \in S^*) P(a_j \in S^*)}{P(A^j)} \\ &= \frac{P(A^j | a_j \in S^*) P(a_j \in S^*)}{P(A^j | a_j \in S^*) P(a_j \in S^*) + P(A^j | a_j \notin S^*) P(a_j \notin S^*)} \\ &= \frac{t_j \prod_{i: a_j \in A_i} p_i \prod_{i: a_j \notin A_i} (1 - p_i)}{t_j \prod_{i: a_j \in A_i} p_i \prod_{i: a_j \notin A_i} (1 - p_i) + (1 - t_j) \prod_{i: a_j \in A_i} q_i \prod_{i: a_j \notin A_i} (1 - q_i)} \\ &= \frac{\frac{t_j}{1 - t_j} \prod_{i: a_j \in A_i} \frac{p_i}{q_i} \prod_{i: a_j \notin A_i} \frac{1 - p_i}{1 - q_i}}{1 + \frac{t_j}{1 - t_j} \prod_{i: a_j \in A_i} \frac{p_i}{q_i} \prod_{i: a_j \notin A_i} \frac{1 - p_i}{1 - q_i}} \end{aligned}$$

But we have that:

$$\prod_{i:a_j \in A_i} \frac{p_i}{q_i} \prod_{i:a_j \notin A_i} \frac{1-p_i}{1-q_i} = \prod_{i:a_j \in A_i} \frac{p_i(1-q_i)}{q_i(1-p_i)} \prod_{i=1}^n \frac{1-p_i}{1-q_i}$$

So:

$$\frac{t_j}{1-t_j} \prod_{i:j \in A_i} \frac{p_i}{q_i} \prod_{i:j \notin A_i} \frac{1-p_i}{1-q_i} = \exp(\text{app}_w(a_j) - \tau_n)$$

and we conclude that:

$$P^j = \frac{\exp(\text{app}_w(a_j) - \tau_n)}{1 + \exp(\text{app}_w(a_j) - \tau_n)}$$

□

Example 5.1. Consider two voters and two alternatives $\mathcal{X} = \{a, b, c\}$. We suppose that $t_a = t_b = t_c = 0.6$ and that the voters' reliabilities are $(p_1, q_1) = (0.9, 0.3)$ and $(p_2, q_2) = (0.6, 0.4)$. First, we can compute the weight of each voter:

$$w_1 = \ln \left(\frac{p_1(1-q_1)}{q_1(1-p_1)} \right) = 3.04, w_2 = \ln \left(\frac{p_2(1-q_2)}{q_2(1-p_2)} \right) = 0.81$$

and the threshold:

$$\tau = \sum_{i=1}^2 \ln \left(\frac{1-q_i}{1-p_i} \right) = 2.35$$

Suppose that $A_1 = \{a, b, c\}$ and $A_2 = \{a\}$. The weighted approval score of the alternatives are:

$$\text{app}_w(a) = 4.25, \text{app}_w(b) = \text{app}_w(c) = 3.44$$

Now, given the formula:

$$P^j = \frac{\exp(\text{app}_w(a_j) - \tau_n)}{1 + \exp(\text{app}_w(a_j) - \tau_n)}$$

we can compute the posterior probabilities that each alternative is in the ground truth:

$$P^1 = 0.87, P^2 = P^3 = 0.75$$

Lemma 5.3. Suppose $l < u$. Then for any $S \in \mathcal{S}_{l,u}$:

$$\psi(S) = \begin{cases} \prod_{a_j \notin S} (1 - P^j) & , \text{ if } |S| = l \\ \prod_{a_j \in S} P^j & , \text{ if } |S| = u \\ \prod_{a_j \in S} P^j \prod_{a_j \notin S} (1 - P^j) & , \text{ if } l < |S| < u \end{cases}$$

where $\psi(S) = P(f(S) = M_{l,u}|A)$ is the probability that the admissible subset S maximizes the utility f .

Proof. Suppose $l < u$ and let $S \in \mathcal{S}_{l,u}$:

- If $|S| = l$ then $f(S) = M_{l,u} \iff S^* \subseteq S$: suppose that $\exists b \in S^* \setminus S$ such that $f(b) = \alpha > 0$, then $f(S \cup b) > f(S)$. The converse is immediate. So:

$$\begin{aligned}
 \psi(S) &= \frac{P(A \cap (S^* \subseteq S))}{P(A)} = \sum_{S' \subseteq S} \frac{P(A \cap (S^* = S'))}{P(A)} \\
 &= \sum_{S' \subseteq S} \frac{P(A \cap \bigcap_{a_j \in S'} (a_j \in S^*) \cap \bigcap_{a_j \notin S'} (a_j \notin S^*))}{P(A)} \\
 &= \sum_{S' \subseteq S} \frac{P(\bigcap_{j=1}^m A^j \cap \bigcap_{a_j \in S'} (a_j \in S^*) \cap \bigcap_{a_j \notin S'} (a_j \notin S^*))}{P(\bigcap_{j=1}^m A^j)} \\
 &= \sum_{S' \subseteq S} \frac{P(\bigcap_{a_j \in S'} ((a_j \in S^*) \cap A^j) \cap \bigcap_{a_j \notin S'} ((a_j \notin S^*) \cap A^j))}{P(\bigcap_{j=1}^m A^j)} \\
 &= \sum_{S' \subseteq S} \prod_{a_j \in S'} \frac{P((a_j \in S^*) \cap A^j)}{P(A^j)} \prod_{a_j \notin S'} \frac{P((a_j \notin S^*) \cap A^j)}{P(A^j)} \\
 &= \sum_{S' \subseteq S} \prod_{a_j \in S'} P(a_j \in S^* | A^j) \prod_{a_j \notin S'} P(a_j \notin S^* | A^j) \\
 &= \sum_{S' \subseteq S} \prod_{a_j \in S'} P^j \prod_{a_j \notin S'} (1 - P^j) = \prod_{a_j \in S} (1 - P^j)
 \end{aligned}$$

- If $|S| = u$ then $f(S) = M_{l,u} \iff S \subseteq S^*$ (by similar argument). So:

$$\begin{aligned}
 \psi(S) &= \frac{P(A \cap (S \subseteq S^*))}{P(A)} = \sum_{S \subseteq S'} \frac{P(A \cap (S^* = S'))}{P(A)} \\
 &= \sum_{S \subseteq S'} \prod_{a_j \in S'} P^j \prod_{a_j \notin S'} (1 - P^j) = \prod_{a_j \in S} P^j
 \end{aligned}$$

- If $l < |S| < u$ then $f(S) = M_{l,u} \iff S = S^*$: suppose that $\exists b \notin S$ such that $f(b) = \alpha > 0$ then $f(S \cup b) > f(S)$, and if $\exists a \in S$ such that $f(a) = -\beta < 0$ then $f(S \setminus a) > f(S)$. The converse is immediate. So

$$\psi(S) = \frac{P(A \cap S = S^*)}{P(A)} = \prod_{a_j \in S} P^j \prod_{a_j \notin S} (1 - P^j)$$

□

Example 5.2. Recall the settings of Example 5.1, and suppose that the bounds are $(l, u) = (1, 3)$. Since these bounds only bear on the size of the output, they do not affect the probabilities of the

alternatives to be in the ground truth that we computed in Example 5.1. We will now compute the probability that each admissible subset maximizes the utility. The empty set is the only non-admissible subset. For every singleton $\{a\}$, $\{b\}$ or $\{c\}$, the probability of maximizing the utility f is :

$$\begin{cases} \psi(\{a\}) &= (1 - P^b)(1 - P^c) = 0.0625 \\ \psi(\{b\}) &= (1 - P^a)(1 - P^c) = 0.0325 \\ \psi(\{c\}) &= (1 - P^a)(1 - P^b) = 0.0325 \end{cases}$$

The probabilities of the pairs of alternatives maximizing the utility are:

$$\begin{cases} \psi(\{a, b\}) &= P^a P^b (1 - P^c) = 0.156 \\ \psi(\{a, c\}) &= P^a P^c (1 - P^b) = 0.156 \\ \psi(\{b, c\}) &= P^b P^c (1 - P^a) = 0.073 \end{cases}$$

Finally, the probability that the set of all alternatives $\{a, b, c\}$ maximizes f is:

$$\psi(\{a, b, c\}) = P^a P^b P^c = 0.467$$

Now we can state the main result. It claims that the MLAUM coincides either with the admissible set that is most likely to be the ground truth (which were characterized in Proposition 5.1) or with one of the sets of top- l or top- u alternatives according to their weighted approval score.

Theorem 5.4. *Suppose $l < u$ and define $\tilde{S}_l = \{a_1, \dots, a_l\}$ and $\tilde{S}_u = \{a_1, \dots, a_u\}$ and $\tilde{S} = \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S|A)$. We have that:*

$$\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S) = \arg \max_{S \in \tilde{S} \cup \{\tilde{S}_l, \tilde{S}_u\}} \psi(S)$$

Proof. Let $S \in \mathcal{S}_{l,u}$:

If $|S| = l$ then by Lemma 5.3 we have that $\psi(S) = \prod_{a_j \notin S} (1 - P^j)$. We also know by Lemma 5.2 that P^j increases with $app_w(a_j)$ which implies that $\psi(S) \leq \psi(\tilde{S}_l)$.

Similarly, if $|S| = u$ then $\psi(S) = \prod_{a_j \in S} P^j \leq \psi(\tilde{S}_u)$.

Now, suppose that $l < |S| < u$, we have:

$$\psi(S) = \prod_{a_j \in S} P^j \prod_{a_j \notin S} (1 - P^j) = P(S^* = S|A)$$

So for any $\tilde{S} \in \tilde{S} = \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S|A)$:

$$\psi(S) \leq P(S^* = \tilde{S}|A) \leq \prod_{a_j \in \tilde{S}} P^j \prod_{a_j \notin \tilde{S}} (1 - P^j) \leq \psi(\tilde{S})$$

□

Remark. If $l = u = k$, for any $S \in \mathcal{S}_k$ we have that:

$$\begin{aligned} f(S) = M_k &\iff S \subseteq S^* \text{ or } (S \not\subseteq S^* \text{ and } \bar{S} \subseteq \bar{S}^*) \\ &\iff S \subseteq S^* \text{ or } S^* \subsetneq S \end{aligned}$$

So we can write:

$$\psi(S) = P(S \subseteq S^* | A) + \sum_{S' \subsetneq S} P(S^* = S' | A)$$

and we have that $\arg \max_{S \in \mathcal{S}_k} \psi(S) = \tilde{S}_k$

These results enable us to define the following algorithm (see algorithm 3) for the estimation of $\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S)$. It boils down to computing the most likely admissible set to coincide with the ground truth, and then compare it to the top- l and top- u sets. The one with the highest probability of maximizing the utility is output.

Algorithm 3 Estimation of $\arg \max_{S \in \mathcal{S}_{l,u}} P(f(S) = M_{l,u} | A)$ in case $l < u$

Input: Approval ballots $(A_i)_{1 \leq i \leq N}$ and bounds $l < u$

Output: Estimation \hat{S}

- Compute $app_w(a)$ for each alternative $a \in \mathcal{X}$, and threshold τ_n
 - Rank alternatives according to $app_w(a)$: $app_w(a_1) \geq \dots \geq app_w(a_m)$
 - Apply Theorem 4.1 to compute $\tilde{S} = \arg \max_{S \in \mathcal{S}_{l,u}} P(S = S^* | A)$
 - Apply Lemma 5.3 to compute $\psi(\tilde{S}_l)$, $\psi(\tilde{S}_u)$ and $\psi(\tilde{S})$ for any $\tilde{S} \in \tilde{\mathcal{S}}$
 - return** $\hat{S} = \arg \max_{S \in \tilde{\mathcal{S}} \cup \{\tilde{S}_l, \tilde{S}_u\}} \psi(S)$
-

Example 5.3. Consider 5 students $\mathcal{X} = \{a, b, c, d, e\}$ applying to a scholarship and a selection committee composed of 10 voters N all sharing the same parameters $(p, q) = (0.7, 0.4)$. The weight of each voter is $w = \ln \frac{p(1-q)}{q(1-p)} = 1.25$ and $\tau_n = 6.93$. We consider the constraints $(l, u) = (1, 4)$, where at least one scholarship must be accorded, and for budget limitations, the department can only afford a maximum of four scholarships.

First, suppose that the candidate d is favored for some outstanding past achievement, so $t_d = 0.6$ and $t_j = 0.5$ for all the remaining candidates a, b, c and e . Consider also the following approval counts (and probabilities $P^j = P(a_j \in S^* | A^j) = \frac{\exp(app_w(a_j) - \tau_n)}{1 + \exp(app_w(a_j) - \tau_n)}$):

Candidate	a	b	c	d	e
Approval count	9	8	7	5	5
P^j	0.98	0.95	0.86	0.43	0.34

Noticing that in general $a_j \in S_{max}^{\tau_n}$ if and only if $P^j > 0.5$, we can easily check, by Proposition 5.1 that $\tilde{S} = \arg \max_{S \in \mathcal{S}_{l,u}} P(S = S^* | A) = \{a, b, c\}$: we have that $S_{max}^{\tau_n} = \{a, b, c\}$, $S_{tie}^{\tau_n} = \emptyset$

and $S_{min}^{\tau_n} = \{d, e\}$, and we know that \tilde{S} consists of the top- k alternatives, such that:

$$\begin{cases} k \in [1, 4] \\ |\tilde{S} \cap S_{max}^{\tau_n}| = \min(u, k_{max}^{\tau_n}) = 3 \implies \{a, b, c\} \subseteq \tilde{S} \\ |\tilde{S} \cap S_{min}^{\tau_n}| = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) = 0 \implies d, e \notin \tilde{S} \end{cases}$$

So the only possibility is $\tilde{S} = \{a, b, c\}$. We also have $\tilde{S}_l = \{a\}$ and $\tilde{S}_u = \{a, b, c, d\}$. We want to find:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S)$$

By Theorem 5.4, we only need to compare ψ of the three mentioned sets:

$$\begin{cases} \psi(\tilde{S}_l) = (1 - P^b)(1 - P^c)(1 - P^d)(1 - P^e) = 0.002 \\ \psi(\tilde{S}_u) = P^a \times P^b \times P^c \times P^d = 0.34 \\ \psi(\tilde{S}) = P^a \times P^b \times P^c(1 - P^d)(1 - P^e) = 0.3 \end{cases}$$

Since $\psi(\tilde{S}_u) > \psi(\tilde{S}) > \psi(\tilde{S}_l)$, we have:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S) = \tilde{S}_u = \{a, b, c, d\}$$

Now suppose rather that $t_j = 0.5$ for all the candidates. And that candidate e loses an approval:

Candidate	a	b	c	d	e
Approval count	9	8	7	5	4
P^j	0.98	0.95	0.86	0.34	0.12

We now have that:

$$\begin{cases} \psi(\tilde{S}_l) = (1 - P^b)(1 - P^c)(1 - P^d)(1 - P^e) = 0.004 \\ \psi(\tilde{S}_u) = P^a \times P^b \times P^c \times P^d = 0.27 \\ \psi(\tilde{S}) = P^a \times P^b \times P^c(1 - P^d)(1 - P^e) = 0.46 \end{cases}$$

Since $\psi(\tilde{S}) > \psi(\tilde{S}_u) > \psi(\tilde{S}_l)$, we have:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S) = \tilde{S} = \{a, b, c\}$$

Surprisingly, estimating the most likely admissible utility maximizer does not depend on the exact values of the utility. It only requires that α and β are positive, which means that the net utility of a good alternative and the net cost of a bad alternative are positive. This makes the MLAUM a suitable solution in situations where assessing the exact utility values is difficult. We will see that this is not the case for the admissible expected utility maximizer.

5.3 AEUM: Admissible Expected Utility Maximizer

Another natural approach, instead of searching the admissible set of alternatives $S \in \mathcal{S}_{l,u}$ with maximum probability of maximizing f , would be to seek the admissible set maximizing the expected value of f , that is:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E} [f(S)|A]$$

Definition 5.1. *In light of Definition 4.2, we introduce a new threshold:*

$$\tau'_n = \tau_n + \log \frac{\beta}{\alpha}$$

where:

$$\tau_n = \sum_{i=1}^n \ln \left(\frac{1 - q_i}{1 - p_i} \right)$$

We define also the associated partition:

$$\begin{cases} S_{max}^{\tau'_n} &= \{a \in \mathcal{X}, app_w(a) > \tau'_n\} \\ S_{tie}^{\tau'_n} &= \{a \in \mathcal{X}, app_w(a) = \tau'_n\} \\ S_{min}^{\tau'_n} &= \mathcal{X} \setminus S_{max}^{\tau'_n} \cup S_{tie}^{\tau'_n} \end{cases}$$

We also denote $k_{max}^{\tau'_n} = |S_{max}^{\tau'_n}|$, $k_{tie}^{\tau'_n} = |S_{tie}^{\tau'_n}|$, $k_{min}^{\tau'_n} = |S_{min}^{\tau'_n}|$.

In contrast with the threshold τ_n introduced in Definition 4.2, which is agnostic to the utility levels α and β , the new threshold τ'_n contains a new element $\log \frac{\beta}{\alpha}$ which calibrates its selectivity according to the ratio of the (dis)utility of the alternatives:

- When $\beta > \alpha$ the cost of an objectively bad alternative is higher than the utility of an objectively good alternative. In this case, we have that $\log \frac{\beta}{\alpha} > 0$, and the threshold τ'_n is set higher than τ_n . This implies that it requires an alternative to be backed by more approval votes or prior evidence to be considered as a potentially good one.
- When $\alpha > \beta$ the utility of an objectively good alternative exceeds the cost of an objectively bad alternative. We have that $\log \frac{\beta}{\alpha} < 0$, and the threshold τ'_n is less risk-averse than τ_n since the alternatives need less weighted approval scores to get above it.

Theorem 5.5. $\tilde{S} \in \arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E} [f(S)|A]$ if and only if there exists $k \in [l, u]$ such that $\tilde{S} = \{a_1, \dots, a_k\}$ and:

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau'_n}| &= \min(u, k_{max}^{\tau'_n}) \\ |\tilde{S} \cap S_{min}^{\tau'_n}| &= \max(0, l - k_{tie}^{\tau'_n} - k_{max}^{\tau'_n}) \end{cases}$$

Remark. *In particular, if $\alpha = \beta$ then the set of admissible alternatives maximizing the expected value of f coincides with the (admissible) set of alternatives most likely to be the ground truth (characterized in Theorem 4.1):*

$$\arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E} [f(S)|A] = \arg \max_{S \in \mathcal{S}_{l,u}} P(S^* = S|A)$$

Proof. First, notice that for a set S , the possible values that $f(S)$ can take is the set of quantities $\alpha|S| - (\alpha + \beta)v$ where $v \in [0, |S|]$ is the number of "non-eligible" alternatives in S . So the expectation can be written as:

$$\begin{aligned} \mathbb{E}[f(S)|A] &= \sum_{v=0}^{|S|} (\alpha|S| - (\alpha + \beta)v) P(|S \cap S^*| = |S| - v|A) \\ &= \sum_{v=0}^{|S|} (\alpha|S| - (\alpha + \beta)v) \times \sum_{\substack{x \in \{0,1\}^{|S|} \\ \sum x_j = |S| - v}} \prod_{a_j \in S} (P^j)^{x_j} (1 - P^j)^{1-x_j} \end{aligned}$$

Since the alternatives are ranked according to $app_w(a_j)$ (and hence according to P^j) it easy to prove that for any k and S such that $|S| = k$ we have that $\mathbb{E}[f(S)|A] \leq \mathbb{E}[f(S_k)|A]$ where $S_k = \{a_1, \dots, a_k\}$. Moreover, if $l < u$ we have that for any $k \in [l, u - 1]$:

$$\begin{aligned} \mathbb{E}[f(S_{k+1})|A] &= \sum_{v=0}^{k+1} (\alpha k + \alpha - (\alpha + \beta)v) P(|S_{k+1} \cap S^*| = k + 1 - v|A) \\ &= \sum_{v=0}^{k+1} (\alpha k + \alpha - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^{k+1} \\ \sum x_j = k+1-v}} \prod_{a_j \in S_{k+1}} (P^j)^{x_j} (1 - P^j)^{1-x_j} \\ &= \sum_{v=0}^{k+1} (\alpha k + \alpha - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} P^{k+1} \\ &+ \sum_{v=0}^{k+1} (\alpha k + \alpha - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v+1}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} (1 - P^{k+1}) \\ &= P^{k+1} \sum_{v=0}^k (\alpha k - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\ &+ \alpha P^{k+1} \sum_{v=0}^k \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\ &+ (1 - P^{k+1}) \sum_{v=1}^{k+1} (\alpha k + \alpha - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v+1}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\ &= P^{k+1} \mathbb{E}[f(S_k)|A] + \alpha P^{k+1} \sum_{v=0}^k \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \end{aligned}$$

$$\begin{aligned}
& + (1 - P^{k+1}) \sum_{v=0}^k (\alpha k - \beta - (\alpha + \beta)v) \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\
& = P^{k+1} \mathbb{E}[f(S_k)|A] + \alpha P^{k+1} \sum_{v=0}^k \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\
& \quad + (1 - P^{k+1}) \mathbb{E}[f(S_k)|A] - \beta(1 - P^{k+1}) \times \sum_{v=0}^k \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j} \\
& = \mathbb{E}[f(S_k)|A] + \underbrace{((\alpha + \beta)P^{k+1} - \beta) \sum_{v=0}^k \sum_{\substack{x \in \{0,1\}^k \\ \sum x_j = k-v}} \prod_{j=1}^k (P^j)^{x_j} (1 - P^j)^{1-x_j}}_{>0}
\end{aligned}$$

So we have that:

$$\begin{aligned}
\mathbb{E}[f(S_{k+1})|A] > \mathbb{E}[f(S_k)|A] & \iff P^{k+1} > \frac{\beta}{\alpha + \beta} \\
& \iff \exp(\text{app}_w(a_{k+1}) - \tau'_n) \geq \frac{\beta}{\alpha} \\
& \iff a_{k+1} \in S'_{max}
\end{aligned}$$

and:

$$\mathbb{E}[f(S_{k+1})|A] = \mathbb{E}[f(S_k)|A] \iff a_{k+1} \in S'_{tie}$$

Thus we can proceed exactly like in proof of Theorem 4.1 to characterize $\arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E}[f(S)|A]$. \square

Example 5.4. Recall the settings of Example 5.3 with 5 candidates and 10 voters N all sharing the same parameters $(p, q) = (0.7, 0.4)$. All the voters share the same weight $w = 1.25$, and the threshold is equal to $\tau_n = 6.93$. We keep the same constraints $(l, u) = (1, 4)$, the same prior information, namely $t_d = 0.6$ and $t_a = t_b = t_c = t_e = 0.5$, and the same approval counts:

Candidate	a	b	c	d	e
Approval count	9	8	7	5	5
Prior	0.5	0.5	0.5	0.6	0.5
app_w	11.25	10	8.75	6.65	6.25

If we set α and β such that $\alpha = \beta$ we have that $\tau'_n = \tau_n$ and thus:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E}[f(S)|A] = \arg \max_{S \in \mathcal{S}_{l,u}} P(S = S^*|A) = \{a, b, c\}$$

Now suppose that $\beta = 10\alpha$, then $\tau'_n = \tau_n + \ln(10) = 9.23$, thus we have that $S_{max}^{\tau'_n} = \{a, b\}$ and $S_{min}^{\tau'_n} = \{c, d, e\}$. We can easily check that in this case:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \mathbb{E} [f(S)|A] = \{a, b\}$$

Recalling that the most likely admissible utility maximizer in this case is the set $\{a, b, c, d\}$, we can see how the admissible expected utility maximizer $\{a, b\}$ adapts to a disutility β which is greater than the utility α by selecting fewer alternatives.

5.4 Simulations

In order to assess the efficiency of the two approaches for recovering an admissible set maximizing f , we generate artificial data, namely L instances (ground truths and approval profiles). Assessing the performances of these methods on real data is challenging because, usually, the ground truth is not accessible, and the decision that we make affects it. For instance, in the master's program selection problem, once an application is rejected, it is not possible to know whether the candidate was objectively bad or wrongly discarded. Also, if we are deciding on projects to allocate investments to, the decision of not supporting a project affects its performance (it is more likely to fail for lack of available liquidity) regardless of whether it was objectively promising or not.

We test and compare both methods with a simpler heuristic that we will introduce further, according to two metrics:

- Hamming loss: the proportion of instances where each method succeeds in recovering a set maximizing f :

$$\frac{1}{L} \sum_{z=1}^L \mathbb{1}\{f(\hat{S}^z) = M_{l,u}^z\}$$

- Mean absolute error MAE: the mean absolute distance between the value of f for the estimated set and the real maximum of f on the admissible sets:

$$\frac{1}{L} \sum_{z=1}^L |f(\hat{S}^z) - M_{l,u}^z|$$

We follow the evolution of these metrics when:

- The number of voters grows.
- The upper bound grows with a fixed lower bound (the constraints loosen).

5.4.1 Heuristics

A simpler heuristic that we will test consists in estimating the set of all eligible alternatives S^* , then extracting an admissible set with a maximum number of eligible alternatives (according to the estimation) and a minimum number of non-eligible alternatives. More precisely:

- If the size of this estimated set lies between the bounds l and u then it suffices to take all its elements.
- If its size is smaller than the lower bound l , then we take all its elements and randomly pick the remaining alternatives ($l - |\hat{S}|$ alternatives) to satisfy the constraint.
- If its size is bigger than the upper bound u than we randomly pick u elements from the estimated set.

5.4.2 Generated Data and Parameters

We consider a set of 30 alternatives and we generate 1000 ground truths (binary vectors) with uniform priors. After that, in Figure 5.1, we fix common values $p = 0.6$ and $q = 0.3$ for all the voters and generate up to 100 votes on each instance and estimate the desired sets for $l = 5$ and $u = 10$, whereas for Figure 5.2, we generate 20 votes and estimate the desired sets for a fixed $l = 5$ and varying upper bound u going from l to $m = 30$.

5.4.3 Results and Observations

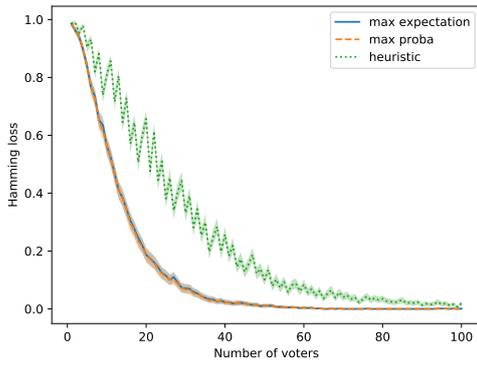
In coherence with the core spirit of epistemic social choice, as the number of voters grows, the output of the three methods gets closer to the ground truth admissible maximizer of f .

For $\alpha = \beta$ (see Figures 5.2a,5.2b,5.1a,5.1b) the performances of the two solution concepts (most likely maximizer of f or expectation maximizer) are nearly identical and are always better than the results given by the heuristic.

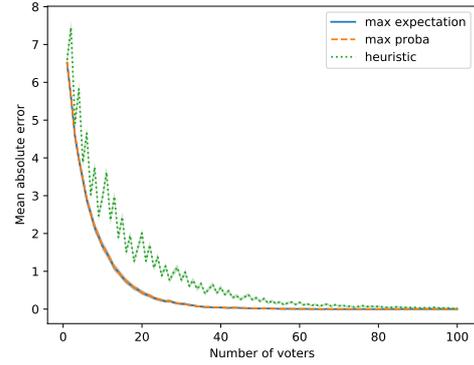
For $\alpha \neq \beta$ (see Figures 5.2c,5.2d,5.1c,5.1d) we can see that computing the most likely maximizer performs better than computing the expectation maximizer w.r.t Hamming loss (see Figures 5.2c, 5.1c) whereas the opposite is true w.r.t MAE (see Figures 5.2d, 5.1d).

As the constraints get looser (see Figure 4.1) the heuristic's performance gets closer to that of computing the most likely maximizer.

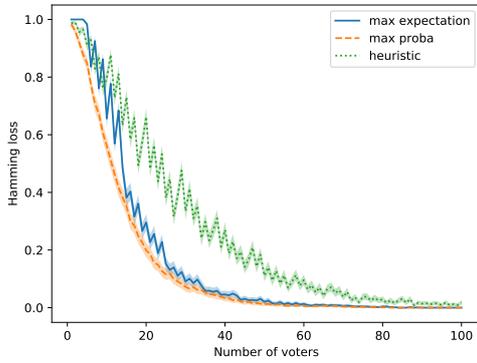
So the experiments show that despite the fact that MLAUM is able to return the exact maximizer of the utility more often than AEUM, on average, AEUM outputs subsets of alternatives whose utilities are closer to the maximum. This means that AEUM is always close to the maximum whereas MLAUM makes fewer, but more grosser, mistakes.



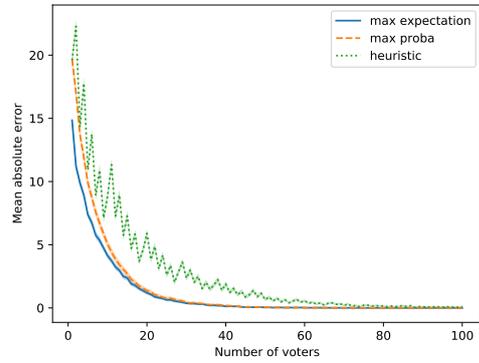
(a) Hamm $\alpha = \beta = 1$



(b) MAE $\alpha = \beta = 1$

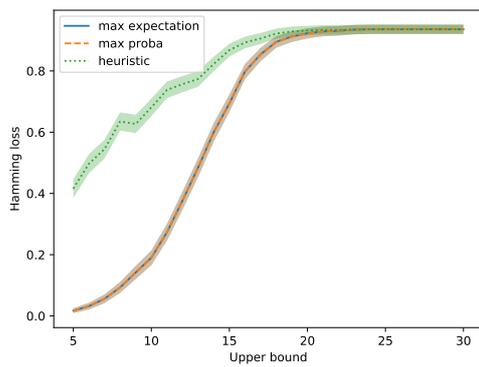


(c) Hamm $\alpha = 5, \beta = 1$

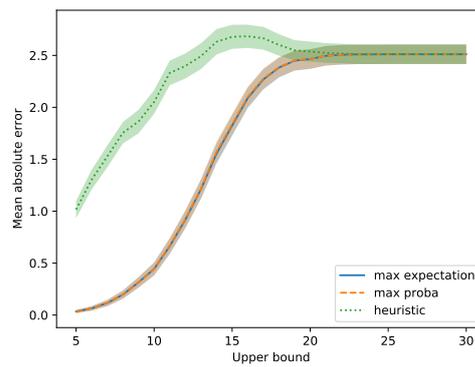


(d) MAE $\alpha = 5, \beta = 1$

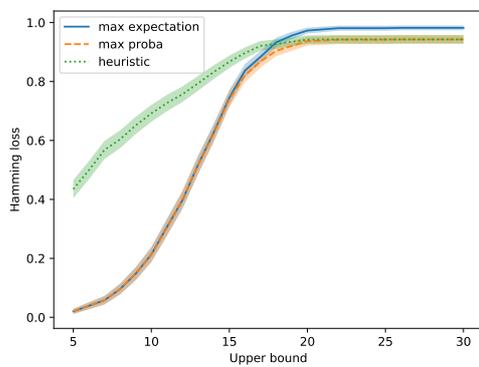
Figure 5.1: Hamming loss and MAE for different number of voters



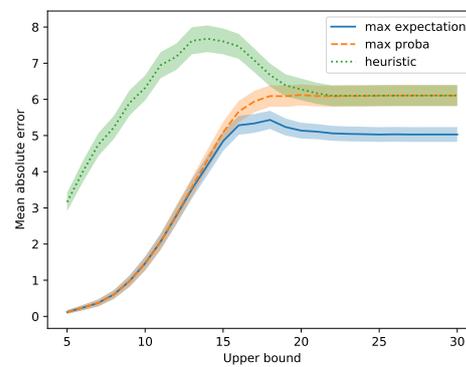
(a) Hamm $\alpha = \beta = 1$



(b) MAE $\alpha = \beta = 1$



(c) Hamm $\alpha = 5, \beta = 1$



(d) MAE $\alpha = 5, \beta = 1$

Figure 5.2: Hamming loss and MAE for different constraints

5.5 Conclusion

We tackled the problem of multi-winner epistemic approval voting with committee-size interpreted as exogenous size constraints. Under this interpretation, we consider the problem of, given the approval votes, selecting a set of alternatives maximizing an a priori unknown underlying utility under size constraints (selecting students to a master program, selecting patients to medical resources, selecting technical solutions to propose to a client..). To solve this problem we suggest and characterize two solution concepts namely: *most likely admissible utility maximizer* and *admissible expected utility maximizer*. We tested these two solutions on synthetic data and showed that they perform differently according to different metrics: although MLAUM returns the exact utility maximizer more often than AEUM, on average, AEUM outputs subsets of alternatives whose utilities are closer to the maximum. The next step would be to generalize this study to cases with more sophisticated utilities than our two-level model, and to situations with more flexible constraints.

Conclusion

This thesis focuses on the epistemic view of approval voting with a particular interest in its application to crowdsourcing collective annotations. After reviewing the existing literature on the subject, we considered approval voting in two contexts, offering adequate novel solutions to each of them, namely: the single-winner and the multi-winner cases.

First, we supposed that the ground truth consists of a single alternative (*e.g.*, detecting the language of a speech, guessing the city in a photo ...) and that each voter approves all the alternatives that she believes might be the correct one. Driven by the intuition that *smaller ballots are more accurate*, we studied aggregation rules which *assign more weight to voters who select fewer alternatives*. Such rules are called *size-decreasing* rules. We began by positing an approval variant of the Mallows noise model, where the probability of casting some approval ballot depends on its (set-)distance to the singleton ground truth, and we observed that when such noises are *neutral*, this probability only depends on the size of the ballot and the belonging or not of the correct answer to the ballot. When all the participants share the same level of expertise, we characterized all the noises whose associated *maximum likelihood estimation* rule is size-decreasing, showing that a necessary and sufficient condition is to have a distance where the importance of the belonging or not of the ground truth to the ballots decreases for the bigger ones. As an immediate consequence, we proved that the rules associated to some usual distances such as *Jaccard* and *Dice* are in fact size-decreasing with respective optimal weights $w_i = \frac{1}{|A_i|}$ and $w_i = \frac{2}{1+|A_i|}$. Nonetheless, all the additively-separable distances, including the well-known *Hamming* metric, do not satisfy the condition. We thus focused on this family of noises and relaxed the assumption of homogeneously reliable voters. In this context, we gave a sufficient condition to guarantee that, for any voter, the expected number of approved alternatives increases as the voter gets less reliable, which motivates the use of size-decreasing aggregations. Although the *Hamming* distance satisfies it, we proved that in its specific case, we can give an explicit closed form formula of the expected size of the ballot as an affine decreasing function of the voter's competence $\mathbb{E}_p[|A_i|] = (m - 1) - (m - 2)p$. This formula gives us a new approach to directly estimate the

reliability of a voter by a mere observation of the size of her answer, with no information on the ground truth and no iterative procedures, and then the voters are weighted accordingly with the optimal weights $w_i = \log \frac{p_i}{1-p_i}$. We compared these novel aggregation rules with *standard approval voting* on three image annotation datasets. The experiments shown that they outperform the baseline most of the times, with a significant performance improvement for the rule associated to the Condorcet noise. However, we also noticed that these rules might fail in situation where many incompetent participants are over-confident, that is, when they are mistakenly sure they know the correct answer.

Second, we turned to the more general and less studied context of epistemic multi-winner approval voting. In classical social choice, the input in such situations is an approval profile and a constraint on the *committee-size* usually defined as an integer k representing the number of candidates to get elected. Our first departure from this is that we relax it into inequality size constraints represented by a lower-bound $l \geq 1$ and an upper-bound $u \leq m$. This lead us to distinguish two possible interpretations in the epistemic settings, depending on whether the constraints on the cardinality of the set of alternatives to be identified bear *on the ground truth itself*, or *on the output*. We studied each interpretation and its optimal rules.

When the constraint plays the role of a *prior knowledge* on the number of correct alternatives, we proposed methods to incorporate into the estimation procedure of the ground truths of multiple related instances. In fact, we defined a noise model which capture the two types of possible errors in the approval votes, along with a prior distribution which reflects the certainty that the ground truth subset of alternatives has a cardinality that satisfies the constraints. Our first step was to characterize the maximum likelihood estimator of the ground truth when the noise parameters are known: we proved that there is a threshold τ on the weighted approval score of the alternatives which enables us to partition them into three sets; the "good", the "borderline", and the "bad" ones, with the maximum likelihood estimator being a set of top- k candidates containing as many as possible of "good" alternatives without violating the upper-bound u , and as few as possible of "bad" alternatives as to satisfy the lower-bound l . The second step was to compute the maximum likelihood estimator of the noise parameters given the ground truth. These two steps were then incorporated into a single iterative procedure that we called *alternating maximum likelihood estimation* (AMLE), that alternates the two estimations, increasing the overall likelihood at each step, until convergence to a fixed point after a finite number of iterations. To assess the performance of AMLE, we collected a new dataset of image annotations. To this end, we designed a football quiz asking voters to select the teams they think might appear in the photos they are shown. We managed to gather the answers of 76 participants. The experiments showed that AMLE outperforms the baselines (namely, the modal rule and an enhanced variant of label-wise majority taking into account the size prior knowledge). They also demonstrated that the incorporation of the size constraints in the model significantly increase the estimation quality.

When the constraints bear on the size of the output of the voting rule we need to aggregate the votes in a way that yields the best admissible subset of alternatives which maximizes some utility. The challenge in such context is that the ground truth worth of each alternative is not known, but

needs to be estimated given the voters' input. For this aim, we propose two solution concepts. The first possibility is to compute the subset of alternatives which is most likely to maximize the utility. We proved that this set is either the top- l set or the top- u set or the admissible maximum likelihood estimator of the ground truth set of all winning alternatives, and we showed how we can compute this latter based on a strong connection with the results of the first size constraints' interpretation. The second solution concepts consists in finding the subset which maximizes the expected utility, and we demonstrated that such maximum can be computed in a similar fashion, involving a threshold which, contrary to the previous solution concept, depends on the ratio between the utility of an objectively good alternative and the cost (or dis-utility) of an objectively bad one. We compared the two solution concepts, along with a heuristic approach on artificially generated data and observed that whilst they both outperform it, their performances differ according to the metric we use.

We showed all along the thesis how the use of statistical learning inspired techniques, namely maximum likelihood estimation and expectation-maximization-like procedures can enhance our understanding of the voters' behaviours through probabilistic modeling, and thus ameliorate the design of proper aggregation rules in the computational social choice field.

We also emphasised the importance of putting our novel voting rules to the test, confronting them with real world annotation data. When data was not available we designed our own crowdsourcing experiment to collect new annotations on image labeling tasks that we created. This further supports the usability of the voting rules, backed by the theoretical guarantees that we provide. It also reveals the limitations of the underlying noise models and might give insights into ways to fit them more to reality.

The works in this dissertation paves the way to other future research direction such as:

- Whilst we studied methods that are fully agnostic to the participants attributes (age, gender, ethnicity, socio-professional category ...) and only uses their mere votes, guaranteeing an equal treatment *ex ante*, we can also study if eliciting some non-trespassing attributes enhances the quality of the expertise estimation and thus the accuracy of the aggregation rule.
- In light of the *surprisingly popular*, the *Anna Karenina*, and our *size matters* principles, we can continue scouting methods for distinguishing the experts that go beyond the majority principle.
- Applying epistemic voting to blockchain protocols to design efficient and strategy proof block validation mechanisms [Caragiannis and Schwartzbach, 2022].
- While some works in the epistemic social choice literature study the case of correlated voters, modelling interaction among alternatives, beyond the size constraints, is a promising direction for enhancing the estimation accuracy. In fact, in many applications, the independence assumption is way too simplifying. For instance, in the chord transcription task, the structure of the musical intervals within chords (triads, extensions ..) is an essential feature and a source of strong interdependence between the notes.

Résumé Long en Français

La théorie du choix social traite du problème de la prise de décision collective. Il s'agit d'un domaine de recherche à l'intersection de l'informatique, des sciences politiques, de l'économie et des mathématiques, qui se concentre sur le traitement des préférences ou des croyances individuelles et leur agrégation en un seul résultat collectif. De nombreux sous-domaines du choix social concernent la conception de méthodes permettant de traiter ces préférences dans différents contextes, allant des *appariements* (matchings) et des *formations de coalitions stables* aux *divisions équitables* et au *vote*. La thèse considère uniquement des questions liés au vote. Ils peuvent être rencontrés dans plusieurs situations telles que:

- Les élections politiques (parlementaires, présidentielles..): selon le système électoral, les votants ont la possibilité de choisir leurs candidates préférés ou encore d'ordonner les candidates selon leurs préférences.
- Les plateformes de planification des réunions: les participants votent sur la date d'une réunion qui convient le plus.
- Les plateformes de crowdsourcing: les participants doivent annotés des instances (images, vidéos, discours ..) en choisissant les labels correctes.

Il est à noter que dans les exemples ci-dessus, les scénarios de vote peuvent différer sur de nombreux aspects, à savoir la nature du vote exprimé (approbation d'alternatives, classement des alternatives ..), la nature du résultat (un unique candidat vainqueur, un ensemble de candidats, un classement d'alternatives ..) et l'interprétation même d'un vote, en effet, dans les élections, un vote est l'expression d'une *préférence* alors que dans le cadre du crowdsourcing et de l'apprentissage d'ensemble, un vote est l'expression d'une *croyance* sur une *vérité* objective. La majeure partie de la thèse traite de cette seconde interprétation appelée *vote épistémique*.

5.6 État de l'art sur le vote épistémique

Le choix social épistémique a connu son développement initial avec Marquis de Condorcet, remontant au XVIII^{ème} siècle. Il considère les votes comme des signaux bruités d'une *vérité de base* objective, et utilise des techniques d'estimation du maximum de vraisemblance pour découvrir cette dernière.

Le résultat initial, connu sous le nom de *théorème du jury de Condorcet*, démontré dans l'article original de Condorcet [Condorcet, 1785] par un argument combinatoire a conduit aux ramifications et extensions ultérieures. Il considère 2 alternatives, avec un candidat étant *objectivement meilleur* que l'autre, mais les deux sont *a priori* également susceptibles d'être le meilleur. Il considère également n électeurs indépendants qui sont également fiables, c'est-à-dire qu'ils sont également susceptibles de choisir la meilleure alternative, et affirme que si chaque électeur vote pour la bonne alternative avec une probabilité $p > \frac{1}{2}$, alors la règle de la majorité maximise la probabilité de coïncider avec la vérité de base parmi tous les estimateurs et qu'elle produit la bonne décision avec une probabilité qui augmente avec le nombre d'électeurs et tend vers 1 lorsque ce dernier croît à l'infini. Formellement, le théorème peut être énoncé comme suit :

Theorem 5.6 (Jury de Condorcet). *Soit un ensemble de deux alternatives $\mathcal{X} = \{a, b\}$ et un ensemble de n électeurs N où n est impair. Supposons qu'il existe un classement réel des alternatives (appelé vérité de base), noté \succ^* , et supposons que les deux possibilités sont a priori également probables :*

$$P(a \succ^* b) = P(b \succ^* a) = \frac{1}{2}$$

Chaque électeur effectue une comparaison \succ_i sur \mathcal{X} indépendamment des autres électeurs, avec la probabilité conditionnelle suivante :

$$P(a \succ_i b | a \succ^* b) = P(b \succ_i a | b \succ^* a) = p$$

où $p > \frac{1}{2}$. Soit M_n le résultat de la règle de la majorité simple :

$$M_n = \begin{cases} a \succ b & \text{si } |i \in N, a \succ_i b| > \frac{n}{2} \\ b \succ a & \text{sinon.} \end{cases}$$

On a alors que:

$$P(M_n = \succ^*) < P(M_{n+2} = \succ^*)$$

et que :

$$\lim_{n \rightarrow +\infty} P(M_n = \succ^*) = 1$$

Le TJC a depuis été généralisée dans de nombreuses directions à des contextes avec des votants avec de compétences différentes [Owen et al., 1989], des votes corrélés [Shapley and Grofman, 1984, Ladha, 1995], plus que deux alternatives [List and Goodin, 2001], le vote par approbation [Everaere et al., 2010] et à un cadre abstrait unifiant plusieurs résultats [Pivato, 2017]. Le problème annexe de savoir comment faut-il avoir de votants pour estimer la vérité avec une précision

donné a aussi été étudié dans [Caragiannis et al., 2016] pour une vérité sous forme de classement des alternatives et des vote sous forme d'ordre linéaires aussi, dans [Caragiannis and Micha, 2017] pour un classement comme vérité de base mais des votes sous forme d'approbations tronqués, et dans [Karge and Rudolph, 2022] pour des bulletins d'approbations et un unique candidat vainqueur.

Ces travaux traitent des résultats de type TJC montrant sous quelles conditions certaines règles connues sont capables de récupérer la vérité de base avec un nombre suffisant de votants, le tableau suivant résume grossièrement les papiers qui essaient de répondre à la problématique suivante: *étant donné certains paramètres (votants, alternatives et entrées), quelle règle est la plus susceptible de produire la vérité de base?*

Résultat Entrées	winner alternative	ranking	set of winners
vainqueur-unique	[Nitzan and Paroush, 1982] [Ben-Yashar and Nitzan, 1997] [Halpern et al., 2021]		
classement		[Young, 1986] [Young, 1988] [Young, 1995]	
ordre partiel	[Xia and Conitzer, 2011]		
comparaison par paires	[Elkind and Shah, 2014]		[Procaccia et al., 2012]
approbation tronquée	[Procaccia and Shah, 2015]		
approbation			[Ben-Yashar and Paroush, 2001] [Caragiannis et al., 2020]
CP-nets			[Xia et al., 2010]

Table 5.1: Règles optimales en fonction de la nature de leurs entrées-sorties

Dans les sections suivantes on présentera nos contributions sur ce sujet.

5.7 Vote par approbation épistémique à un seul gagnant : la taille importe

On considère un ensemble de n votants N et un ensemble de m alternatives $\mathcal{X} = \{a_1, \dots, a_m\}$. On suppose qu'il existe une unique alternative correcte $a^* \in \mathcal{X}$ qui n'est pas a priori connu, et que les votants soumettent un profil de votes par approbation (A_1, \dots, A_n) , où $A_i \subseteq \mathcal{X}$ est l'ensemble d'alternatives qui peuvent coïncider avec la réalité objective selon le votant $i \in N$.

Pour agréger ces votes en vue d'estimer la bonne réponse, une idée intuitive pourrait être de considérer que les bulletins plus petits, c'est-à-dire les réponses qui contiennent moins d'alternatives, sont plus fiables : un électeur qui connaît la vraie réponse (ou, plus généralement, qui croit la connaître) ne sélectionnera probablement qu'une seule alternative et un électeur qui sélectionne

toutes les alternatives n'a probablement aucune idée de la bonne réponse. En se basant sur cette intuition, un poids plus important doit être attribué aux bulletins plus petits. Les règles qui fonctionnent de cette manière, que nous appelons *size-decreasing approval rules*, font partie de la famille des *size approval rules*. et ont été étudiées axiomatiquement dans [Alcalde-Unzu and Vorsatz, 2009]. Notre objectif est de motiver l'utilisation de telles règles du point de vue du choix social épistémique. À cette fin, nous étudierons une famille de modèles de bruit qui sont des variantes de vote d'approbation du modèle de Mallows, et nous prouverons que dans de nombreux cas, la règle optimale est *size-decreasing*.

5.7.1 Le modèle de bruit

Nous modéliserons la distribution de ces bulletins d'approbation par des variantes de vote d'approbation du modèle de bruit *Mallows*. La distribution de Mallows a été définie à l'origine sur des *rankings* : nous l'adaptions à des sous-ensembles d'alternatives, en conservant l'idée que la probabilité d'un sous-ensemble diminue à mesure que sa distance par rapport à un point central augmente, la dispersion étant modélisée par un paramètre ϕ .

En général, nous appellerons modèle de bruit de Mallows d'approbation tout modèle où les bulletins des électeurs sont indépendants (nous conservons cette hypothèse tout au long de cette section) et où il existe n paramètres $\phi_i \in (0, +\infty)$ et une fonction $d : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \mapsto \mathbb{R}$ telle que pour tout électeur $i \in N$:

$$P_{\phi_i, d}(A_i | a^* = a) = \frac{1}{\beta_i} \phi_i^{d(a^*, A_i)}, \forall a \in \mathcal{X}$$

où β_i est le facteur de normalisation correspondant. Si $\phi_i = \phi$ pour tout $i \in N$, on dit que le modèle est *anonyme*.

Dans la suite, nous nous concentrerons uniquement sur les modèles de bruit *neutre*. La neutralité d'un bruit est, par définition, son invariance par toute permutation des alternatives:

$$\forall \pi \in \sigma(\mathcal{X}), P_{\phi, d}(A | a^* = a) = P_{\phi, d}(\pi(A) | a^* = \pi(a))$$

On voit immédiatement qu'un bruit de Mallows est neutre si et seulement si sa fonction associée d est neutre (invariante par une permutation des alternatives).

Un modèle de bruit est neutre si $d(a, A)$ ne dépend que de $|a \cap A|$ (c'est-à-dire 1 si $a \in A$ et 0 si $a \notin A$)¹ et $|A|$:

Proposition 5.7. *Un modèle de bruit associé à une fonction d est neutre si et seulement s'il existe une fonction unique:*

$$\psi_d : \{0, 1\} \times \{0, \dots, m\} \setminus (1, 0) \mapsto \mathbb{R}$$

tel que:

$$d(a, A) = \psi_d(|a \cap A|, |A|)$$

¹Nous omettons les accolades et écrivons $a \in A$ pour $\{a\} \cap A$.

5.7.2 Cas d'un bruit anonyme

Dans cette sous-section, nous supposons que les électeurs partagent un paramètre de bruit commun (inconnu) $\phi \in (0, 1)$ et qu'il existe une fonction $d : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \mapsto \mathbb{R}$ et sa fonction associée ψ_d telle que, pour tout $a \in \mathcal{X}$:

$$P_{\phi,d}(A_i | a^* = a) = \frac{1}{\beta} \phi^{d(a^*, A_i)} = \frac{1}{\beta} \phi^{\psi_d(|a^* \cap A_i|, |A_i|)}$$

Après définir formellement la notion de règles décroissantes, nous énonçons le résultat principal de cette sous-section qui caractérise tous les bruits anonymes de Mallows (c'est-à-dire toutes les fonctions d) dont la règle de maximum de vraisemblance associée est *size-decreasing*. Nous verrons que c'est le cas pour certaines distances usuelles d , que nous testerons plus tard dans les expériences.

Definition 5.2 (Size Approval Rule). *Considérons une fonction*

$$\begin{aligned} v : \quad \mathcal{P}(\mathcal{X})^n &\longrightarrow \mathcal{X} \\ (A_1, \dots, A_n) &\mapsto v(A_1, \dots, A_n) \end{aligned}$$

qui, pour chaque profil d'approbation $A = (A_1, \dots, A_n)$, attribue une alternative gagnante $v(A_1, \dots, A_n)$ dans \mathcal{X} . Nous disons que v est une *size approval rule* s'il existe un vecteur $w = (w_0, \dots, w_m) \in \mathbb{R}^{m+1}$ tel que :

$$v(A_1, \dots, A_n) = \arg \max_{a \in \mathcal{X}} \text{app}_w(a)$$

où app_w est le score d'approbation pondéré défini par :

$$\text{app}_w(a) = \sum_{i: a \in A_i} w_{|A_i|}$$

La règle v est dite *size-decreasing* si son vecteur associé $w = (w_0, \dots, w_m) \in \mathbb{R}^{m+1}$ est tel que $w_j > w_{j+1}$ pour tout $1 \leq j \leq m-2$.

Theorem 5.8. *Pour $n \geq 3$, la règle d'estimation du maximum de vraisemblance ζ_d est size-decreasing si et seulement si la fonction*

$$\Delta\psi_d : j \mapsto \psi_d(0, j) - \psi_d(1, j)$$

est strictement décroissante.

Nous donnons maintenant quelques exemples avec des fonctions usuelles d :

- La distance de Jaccard est définie comme suit:

$$d_J(a, A) = 1 - \frac{|a \cap A|}{|A| - |a \cap A| + 1}$$

La règle d'estimation du maximum de vraisemblance associée est *size-decreasing* avec des poids $w_{|A|} = \frac{1}{|A|}$.

- La distance de Dice est définie comme suit:

$$d_{Dice}(a, A) = 1 - \frac{2|a \cap A|}{|A| + 1}$$

La règle d'estimation du maximum de vraisemblance associée est size-decreasing avec des poids $w_{|A|} = \frac{2}{|A|+1}$.

Cependant la règle associée à la distance de Hamming:

$$d_H(a, A) = |a \cap \bar{A}| + |\bar{a} \cap A| = 1 + |A| - 2|a \cap A|$$

n'est pas size-decreasing (elle est la règle Standard Approval Voting). Plus généralement, ceci est vrai pour toute fonction d telle que ψ_d peut être décomposée additivement en deux termes $\psi_d(|a \cap A|, |A|) = f(|a \cap A|) + g(|A|)$. Dans la prochaine sous-section, nous allons considérer cette famille particulière de fonctions décomposables avec un bruit où chaque électeur a son propre paramètre de bruit ϕ_i .

5.7.3 Cas d'un bruit décomposable non-anonyme

Nous définissons un modèle de bruit de Mallows non anonyme, où pour chaque électeur $i \in N$, il existe un paramètre $\phi_i \in]0, +\infty[$ tel que, pour tout $a \in \mathcal{X}$:²

$$P_{\phi_i, d}(A_i | a^* = a) = \frac{1}{\beta_i} \phi_i^{f(|a^* \cap A_i|) + g(|A_i|)}$$

Il est à noter que dans ce cas, un paramètre de bruit individuel plus grand ϕ_i modélise un électeur moins fiable (sa distribution est moins concentrée autour de la vérité de base). Le but du résultat suivant est de motiver l'utilisation de règles d'approbation size-decreasing pour agréger les approbations générées à partir de telles distributions. Plus précisément, le but est de trouver des conditions suffisantes sur f et g qui font que la taille espérée du bulletin de l'électeur $\mathbb{E}_{\phi, d}[|A_i|]$ croît à mesure que l'électeur devient moins fiable (*i.e.* à mesure que son paramètre de bruit ϕ_i croît).

Theorem 5.9. *Si pour chaque $1 \leq t < k \leq m - 1$ on a:*

$$g(k) - g(t) \geq \frac{k-t}{2} [f(0) - f(1)]$$

Alors:

$$\frac{\partial \mathbb{E}_{\phi, d}[|A_i|]}{\partial \phi} \geq 0$$

²Rappelez-vous que les électeurs ne peuvent pas déposer des bulletins d'approbation vides ou pleins. Nous supposons donc que $P(\emptyset | a^* = a) = P(\mathcal{X} | a^* = a) = 0$.

On peut montrer que la distance de Hamming vérifie bien cette propriété. En plus, on remarque que le bruit de Mallows avec une distance de Hamming est équivalent au bruit de Condorcet où on a n paramètres de compétences $p_i \in]0, 1[$ tel que:

$$P_{p_i}(a \in A_i | a = a^*) = P_{p_i}(a \notin A_i | a \neq a^*) = p_i, \forall a \in \mathcal{X}$$

L'appartenance ou pas des différentes alternatives au bulletin d'un votant étant supposé des événements indépendants. Cette forme équivalente nous permet de donner une formule explicite de l'espérance du nombre d'alternatives qu'un votant sélectionnera $\mathbb{E}_{\phi, d}[|A_i|]$ en fonction de son paramètre d'expertise p_i .

Theorem 5.10. *Pour $m \geq 2$, on a:*

$$\mathbb{E}_p[|A_i|] = (m - 1) - (m - 2)p$$

Ce théorème nous donne une approche simple pour estimer p_i par des estimations de maximum de vraisemblance étant donné certaines observations de A_i sans avoir besoin de connaître la vérité de base a^* .

5.7.4 Expériences

Nous avons utilisé les trois jeu de données d'annotation d'images, initialement collectés par [Shah et al., 2015] à des fins de conception de mécanismes de paiements optimaux, afin de tester notre hypothèse et d'évaluer la précision de différentes règles d'agrégation d'intérêt. Chaque ensemble de données consiste en un ensemble de bulletins d'approbation d'un certain nombre de votants (participants) qui devaient sélectionner *toutes les alternatives qu'ils pensaient être correctes* dans un certain nombre d'instances (images), à savoir: des images d'animaux, de textures ou de langues. Puisque nous nous intéressons principalement au problème de la sagesse de la foule pour une seule question, nous ne considérerons que les règles d'agrégation qui fonctionnent par question (les réponses des votants à différentes questions n'affectent pas le résultat de la règle pour une question donnée). Nous utiliserons les méthodes d'agrégation suivantes:

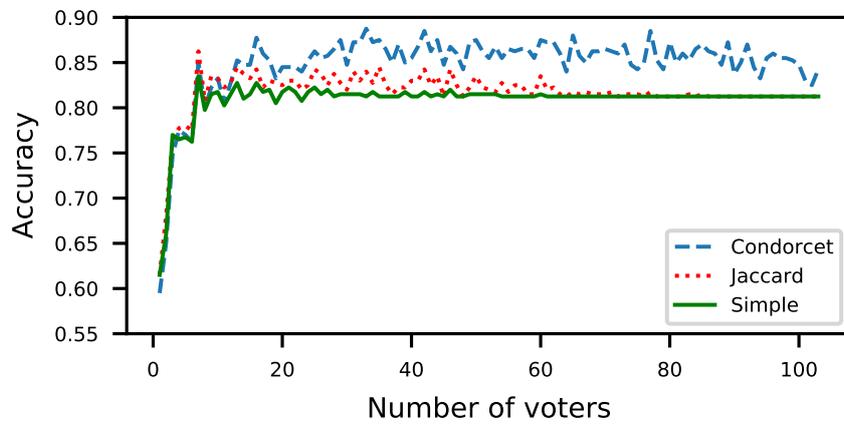
- *Condorcet*: Le paramètre de fiabilité p_i de chaque votant est estimé en observant la taille de son bulletin par la formule:

$$\hat{p}_i = \frac{m - 1 - |A_i|}{m - 2}$$

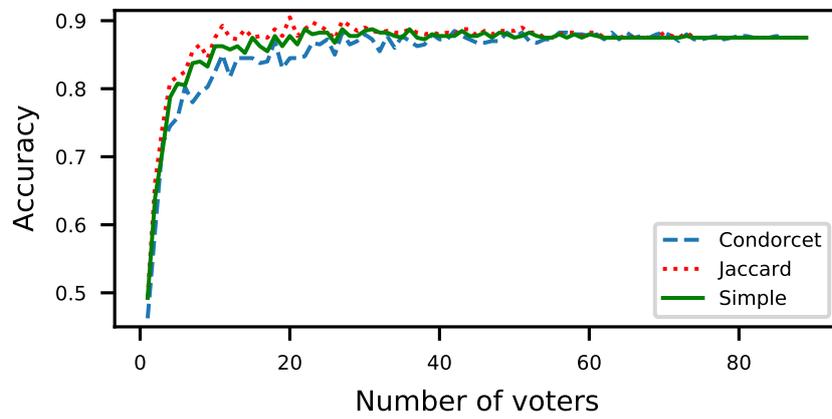
le poids assigné à chaque votant est $\log \frac{\hat{p}_i}{1 - \hat{p}_i}$.

- *Jaccard*: Le poids de chaque est votant $i \in N$ est égale à $\frac{1}{|A_i|}$.
- *SAV*: Tous les votant sont pondérés de façon égale.

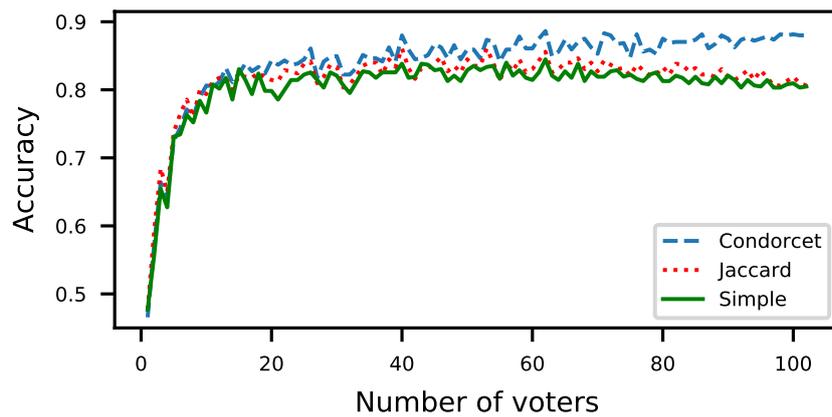
Pour chaque jeu de données, nous avons pris des lots de 25 pour chaque nombre différent de votants, et appliqué les règles susmentionnées. Nous mesurons la précision de chaque règle comme étant la fraction des instances dont la réponse juste à été correctement estimée. Les



(a) Animaux



(b) Textures



(c) Langues

Figure 5.3: Précision des différentes règles d'agrégation

résultats apparaissent sur la figure 5.3. Nous remarquons tout d'abord que pour les trois jeux de données, les règles d'agrégation associées au bruit anonyme de Jaccard montrent une précision légèrement meilleure que la règle d'approbation simple, en particulier pour un petit nombre de votants.

Nous pouvons également constater que la règle d'agrégation associée au bruit de Condorcet non-anonyme présente une amélioration significative de la précision par rapport à cette règle pour Animaux et Langues (en particulier pour des nombres relativement élevés de votants). Cependant, elle ne parvient pas à la surpasser pour le jeu de données de Textures, où elle ne montre que des précisions similaires à la règle standard lorsque le nombre de votants augmente. Cela peut être le résultat de la mauvaise qualité de l'estimation qui n'utilise qu'un seul échantillon ou à des votant trop confiants mais qui se trompent.

5.8 Vote par approbation épistémique à vainqueurs multiples: deux interprétations

Alors que dans le vote épistémique à vainqueur unique, l'interprétation de la vérité de base est claire : il existe une seule alternative correcte — la seule qui soit vraie dans le monde réel — et le but est de l'identifier. Dans le cas d'un vote à plusieurs gagnants, les choses se compliquent.

Rappelons que l'entrée d'une instance de vote multi-vainqueurs consiste en un ensemble de votes et une contrainte sur la taille du comité. Or, deux interprétations différentes de cette contrainte sur la taille du comité coexistent, qui appellent des solutions différentes. La principale différence entre les deux interprétations est de savoir si les contraintes sur la cardinalité de l'ensemble des alternatives à identifier portent *sur la vérité terrain elle-même*, ou *sur la sortie* de la règle d'agrégation. Dans la première interprétation, cette contrainte représente une certaine *de connaissance préalable sur la cardinalité de la vérité de base*, c'est-à-dire sur le nombre d'alternatives gagnantes. Voici deux exemples :

- Considérons une tâche de transcription d'un accord de guitare : les participants entendent un accord et doivent sélectionner l'ensemble des notes qui le constituent parmi un total de 8 alternatives. Nous savons *a priori* que le véritable ensemble de notes est composé d'au moins 3 et d'au plus 6 alternatives.
- Considérons un problème de crowdsourcing dans lequel on montre aux participants l'image d'un match de football, où ils doivent identifier la ou les équipes qui y figurent. Ici, nous savons qu'il doit y avoir une ou deux équipes dans la vérité terrain.

Selon la deuxième interprétation, cette contrainte porte sur le *nombre de gagnants dans la sortie*. En d'autres termes, quelle que soit la vérité de base, nous devons produire un nombre d'alternatives dans un intervalle donné, même si le nombre d'alternatives dans la vérité de base peut se situer en dehors de cet intervalle ; le but est d'identifier un ensemble d'alternatives admissibles les plus proches de la vérité de base, dans un sens qui sera défini plus tard. Voici deux exemples :

- Les alternatives sont des étudiants qui postulent à un programme de master. La vérité de base est constituée des étudiants qui ont objectivement un niveau suffisamment bon pour être diplômés. Le nombre d'étudiants à accepter dans le programme est cependant contraint d'être dans un intervalle $[l, u]$: il faut au moins l pour que le programme s'ouvre, et au plus u à cause de la taille des salles de classe.
- Les alternatives sont des articles soumis à une conférence. Là encore, nous avons un nombre minimal et maximal d'articles à accepter.
- Les alternatives sont les articles qui recevront un prix. Le président de la conférence décerne exactement trois prix. Dans ce cas, la vérité de base est constituée des articles qui méritent vraiment le prix, et le résultat est constitué des trois meilleurs articles. Une variante du problème permettrait aux présidents de conférence de décerner au maximum trois prix, afin d'éviter de décerner un prix à des articles qui ne le méritent pas.

Dans cette interprétation, la taille du comité joue le rôle d'une contrainte de taille exogène qui spécifie le nombre minimum et maximum d'alternatives gagnantes autorisées. Il reste à définir précisément ce que nous entendons par être le plus proche de la vérité du terrain ; nous proposerons deux concepts de solution différents.

Bien que cette distinction entre deux interprétations ait déjà un sens dans le cas du vote épistémique à un seul gagnant, nous montrerons que dans ce cas particulier, elles conduisent à la même solution : elles sont techniquement identiques. Ce ne sera cependant pas le cas dans le cas général ; nous développerons donc des solutions pour chacune de ces deux interprétations séparément.

5.9 Vote par approbation épistémique à vainqueurs multiples: les contraintes de taille comme information a priori

Soit $N = \{1, \dots, n\}$ un ensemble de votants, et $\mathcal{X} = \{a_1, \dots, a_m\}$ un ensemble d'alternatives (objets possibles dans les images, notes dans les accords...). Considérons un ensemble de L instances : une instance z consiste en un profil d'approbation $A^z = (A_1^z, \dots, A_n^z)$ où $A_i^z \subseteq \mathcal{X}$ est le bulletin d'approbation de $i \in N$. Par exemple, dans un contexte de crowdsourcing, une tâche contient généralement plusieurs questions, et une instance comprend les réponses des votants à l'une de ces questions.

Pour chaque instance $z \in L$, il existe une *inconnue* vérité fondamentale S_z^* appartenant à $\mathcal{S} = 2^{\mathcal{X}}$, qui est l'ensemble des alternatives objectivement correctes dans l'instance z . Bien que ces ensembles de vérités fondamentales soient a priori inconnus, nous supposons dans cette section qu'il est communément admis que le nombre d'alternatives dans chacun d'eux se situe dans l'intervalle $[l, u]$: $S_z^* \in \mathcal{S}_{l,u} = \{S \in \mathcal{S}, l \leq |S| \leq u\}$, pour des bornes données $0 \leq l \leq u \leq m$. Notre objectif est de dévoiler la vérité de base pour chacune de ces instances en utilisant les votes et la connaissance préalable du nombre d'alternatives gagnantes.

5.9.1 Le modèle de bruit

Nous définissons un modèle de bruit constitué de deux distributions paramétriques, à savoir une distribution conditionnelle des bulletins d'approbation conditionnellement à la vérité de base, et une distribution a priori sur la vérité de base. Nous nous écartons ici des modèles de bruit classiques dans le choix social épistémique, car nous supposons que les paramètres de ces distributions peuvent être inconnus et doivent donc être estimés.

Formellement, nous supposons que pour chaque électeur $i \in N$, il existe deux paramètres inconnus (p_i, q_i) dans $(0, 1)$ de sorte que le vote d'approbation A_i^z sur une instance $z \in L$ soit tiré selon la distribution suivante, pour chaque $a \in \mathcal{X}$,

$$P(a \in A_i^z | S_z^* = S) = \begin{cases} p_i & \text{if } a \in S \\ q_i & \text{if } a \notin S \end{cases}$$

où p_i (resp. q_i) est la probabilité (inconnue) que le votant i approuve une alternative correcte (resp. incorrecte). Ensuite, nous faisons les hypothèses suivantes :

- (1) Les approbations des alternatives par un électeur sont mutuellement indépendantes étant donné la vérité de base et les paramètres $(p_i, q_i)_{i \in N}$.
- (2) Les bulletins de vote des électeurs sont mutuellement indépendants étant donné la vérité de base.
- (3) Les instances sont indépendantes étant donné les paramètres $(p_i, q_i)_{i \in N}$ et les vérités de base.

Pour modéliser la probabilité a priori de tout ensemble S à être la vérité de base S^* , nous introduisons des paramètres $t_j = P(a_j \in S^*)$. t_j peut être compris comme la probabilité antérieure de a_j d'être dans l'ensemble de vérité fondamentale S^* avant la prise en compte des contraintes de cardinalité. Celles-ci, avec une hypothèse d'indépendance sur les événements $\{a_j \in S^*\}$, donnent :

$$P(S = S^*) = \prod_{a_j \in S} t_j \prod_{a_j \notin S} 1 - t_j$$

. La distribution conditionnelle à la connaissance préalable sur la taille de la vérité terrain peut alors être vue comme une projection sur les contraintes suivie d'une normalisation :

$$\tilde{P}(S) = P(S^* = S | l \leq |S^*| \leq u) = \frac{P(S^* = S \cap |S^*| \in [l, u])}{P(|S^*| \in [l, u])}$$

. Il s'ensuit :

$$\tilde{P}(S) = \begin{cases} \frac{1}{\beta(l, u, t)} \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j) & \text{if } S \in \mathcal{S}_{l, u} \\ 0 & \text{if } S \notin \mathcal{S}_{l, u} \end{cases}$$

où :

$$\beta(l, u, t) = \sum_{S \in \mathcal{S}_{l, u}} \prod_{a_j \in S} t_j \prod_{a_j \notin S} (1 - t_j)$$

Notre objectif est l'estimation conjointe de la vérité objective et des paramètres en maximisant la vraisemblance totale des instances :

$$\mathcal{L}(A, S, p, q, t) = \prod_{z=1}^L \tilde{P}(S_z) \prod_{i=1}^n P(A_i^z | S_z)$$

où :

$$P(A_i^z | S_z) = p_i^{|A_i^z \cap S_z|} q_i^{|A_i^z \cap \bar{S}_z|} (1 - p_i)^{|\bar{A}_i^z \cap S_z|} (1 - q_i)^{|\bar{A}_i^z \cap \bar{S}_z|}$$

Dans ce but, nous allons introduire un algorithme itératif dont les deux principales étapes seront présentées successivement, dans les prochaines sous-sections, avant que l'algorithme principal ne soit formellement défini et sa convergence démontrée. Ces deux étapes sont :

- Estimation des vérités de base en fonction des paramètres.
- Estimation des paramètres en fonction des vérités de base.

En termes simples, l'algorithme consiste à itérer ces deux étapes jusqu'à ce qu'il converge vers un point fixe.

5.9.2 Estimer la vérité objective étant donné les paramètres

Comme les instances sont indépendantes compte tenu des paramètres, nous nous concentrons ici sur une instance avec une vérité de base S^* et un profil d'approbations $A = (A_1, \dots, A_n)$. On va introduire quelques notions avant d'énoncer le théorème caractérisant l'estimateur de maximum de vraisemblance de la réalité.

Definition 5.3 (Score d'approbation pondéré). *Étant donné un profil d'approbation (A_1, \dots, A_n) , des paramètres de bruit $(p_i, q_i)_{1 \leq i \leq n}$ et des paramètres a priori $(t_j)_{1 \leq j \leq m}$, on définit:*

$$app_w(a_j) = \ln \left(\frac{t_j}{1 - t_j} \right) + \sum_{i: a_j \in A_i} \ln \left(\frac{p_i(1 - q_i)}{q_i(1 - p_i)} \right)$$

.

Alors que le poids de chaque votant $i \in N$ dépend de sa fiabilité, chaque information préalable sur une alternative joue le rôle d'un votant virtuel qui ne sélectionne que l'alternative concernée, avec un poids qui augmente à mesure que le paramètre a priori augmente.

A partir de maintenant, nous supposons sans perte de généralité que les alternatives sont classées en fonction de leur score :

$$app_w(a_1) \geq app_w(a_2) \geq \dots \geq app_w(a_m)$$

Definition 5.4 (seuil et partition). *Définissez le seuil :*

$$\tau_n = \sum_{i=1}^n \ln \left(\frac{1 - q_i}{1 - p_i} \right)$$

et la partition de l'ensemble des alternatives en trois ensembles :

$$\begin{cases} S_{max}^{\tau_n} &= \{a \in A, app_w(a) > \tau_n\} \\ S_{tie}^{\tau_n} &= \{a \in A, app_w(a) = \tau_n\} \\ S_{min}^{\tau_n} &= \mathcal{X} \setminus (S_{max}^{\tau_n} \cup S_{tie}^{\tau_n}) \end{cases}$$

and let $k_{max}^{\tau_n} = |S_{max}^{\tau_n}|$, $k_{tie}^{\tau_n} = |S_{tie}^{\tau_n}|$, $k_{min}^{\tau_n} = |S_{min}^{\tau_n}|$.

Le résultat suivant caractérise les ensembles dans \mathcal{S} qui sont des EMV de la vérité de base étant donné les paramètres.

Theorem 5.11. $\tilde{S} \in \arg \max_{S \in \mathcal{S}} \mathcal{L}(A, S, p, q, t)$ si et seulement si il existe $k \in [l, u]$ tel que \tilde{S} est l'ensemble des k alternatives ayant les k valeurs les plus élevées de app_w et :

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau_n}| &= \min(u, k_{max}^{\tau_n}) \\ |\tilde{S} \cap S_{min}^{\tau_n}| &= \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases} \quad (5.1)$$

Ainsi, l'estimateur \tilde{S} est des top- k alternatives, où les valeurs possibles de k sont déterminées par l'Eq. (5.1). La première équation impose que \tilde{S} inclut autant d'éléments que possible parmi $S_{max}^{\tau_n}$ (sans dépasser la borne supérieure u), tandis que la seconde impose que \tilde{S} inclut aussi peu d'éléments que possible parmi $S_{min}^{\tau_n}$ (sans descendre en dessous de la borne inférieure l).

5.9.3 Estimer les paramètres étant donné la réalité objective

Une fois les vérités fondamentales estimées à une itération de l'algorithme, l'étape suivante consiste à estimer les paramètres $(t_j)_{j \in \mathcal{X}}$, les vérités fondamentales étant données. La proposition suivante explicite l'expression de l'estimateur de maximum de vraisemblance du paramètre t_j de chaque alternative étant donné S_z^* une fois que les paramètres t_{-j} de toutes les autres alternatives sont fixés.

Proposition 5.12. Pour tout $a_j \in \mathcal{X}$:

$$\arg \max_{t \in (0,1)} \mathcal{L}(A, S, p, q, t, t_{-j}) = \frac{occ(j)\bar{\alpha}_j}{(L - occ(j))\underline{\alpha}_j + occ(j)\bar{\alpha}_j}$$

$$\text{où: } \begin{cases} \bar{\alpha}_j &= \sum_{\substack{S \in \mathcal{S}_{l,u} \\ a_j \in S}} \prod_{\substack{a_h \in S \\ h \neq j}} t_h \prod_{a_h \notin S} (1 - t_h) \\ \underline{\alpha}_j &= \sum_{\substack{S \in \mathcal{S}_{l,u} \\ a_j \notin S}} \prod_{a_h \in S} t_h \prod_{\substack{a_h \notin S \\ h \neq j}} (1 - t_h) \\ occ(j) &= |z \in \{1, \dots, L\}, a_j \in S_z| \end{cases}$$

Nous pouvons aussi estimer les paramètres des votants (p, q) par maximum de vraisemblance.

Proposition 5.13. Pour $S_z \in \mathcal{S}_{l,u}$ et $t_j \in (0, 1)$ on a :

$$\arg \max_{(p,q) \in (0,1)^{2n}} \mathcal{L}(A, S, p, q, t) = (\hat{p}, \hat{q})$$

où:

$$\hat{p}_i = \frac{\sum_{z \in L} |A_i^z \cap S_z|}{\sum_{z \in L} |S_z|}, \hat{q}_i = \frac{\sum_{z \in L} |A_i^z \cap \overline{S}_z|}{\sum_{z \in L} |\overline{S}_z|}$$

5.9.4 Estimation de maximum de vraisemblance alternée

Maintenant, l'estimation des vérités fondamentales et celle des paramètres sont entrelacées pour maximiser la vraisemblance globale $\mathcal{L}(A, S, p, q, t)$ par l'algorithme de l'*Estimation de maximum de vraisemblance alternée* (EMVA). EMVA est une procédure itérative similaire à la procédure *Espérance-Maximisation* introduite dans [Baharad et al., 2011] mais avec une itération de type coordinate-steepest-ascent, dont l'objectif est d'estimer de manière entrelacée la fiabilité des électeurs, les paramètres a priori des alternatives et les vérités fondamentales des instances. L'idée derrière cette estimation est assez intuitive, et consiste à alterner une EMV des vérités de base étant donné l'estimation actuelle des paramètres, et une mise à jour de ces paramètres via un EMV basé sur l'estimation actuelle des vérités de base. L'algorithme poursuit son exécution jusqu'à ce qu'un critère de convergence soit satisfait sous la forme d'une borne sur la norme de la variation des estimations des paramètres $\|\hat{\theta}^{(v+1)} - \hat{\theta}^{(v)}\|$. En pratique, nous avons choisi ℓ_∞ , mais toute autre norme pourrait être utilisée dans l'algorithme 4 car en dimension finie, toutes les normes sont équivalentes (si une séquence converge selon une norme, alors elle le fait pour toute norme). Nous avons prouvé que l'algorithme converge en un nombre fini d'itérations, mais il est à noter que l'optimalité globale n'est pas garantie au point de convergence.

Theorem 5.14. *Pour toute valeur initiale $\hat{\theta}^{(0)}$, EMVA converge vers un point fixe après un nombre fini d'itérations.*

5.9.5 Expériences

Nous avons conçu une tâche d'annotation d'images sous la forme d'un quiz de football. Nous avons sélectionné 15 images prises lors de différents matchs entre deux des équipes suivantes : Real Madrid, Inter Milan, Bayern Munich, Barcelone, Paris Saint-Germain (PSG). Sur chaque photo, il se peut que des joueurs des deux équipes apparaissent, ou des joueurs d'une seule équipe, donc $l = 1$ et $u = 2$. Chaque participant regarde les instances une par une et doit à chaque fois sélectionner toutes les équipes qu'il peut repérer (voir figure 5.4). La saturation et la teinte des images ont été modifiées afin de rendre la tâche moins évidente. Nous avons recueilli les réponses de 76 participants (seuls deux d'entre eux ont fait du spamming en sélectionnant simplement toutes les alternatives).

Inspirés par le *Principe d'Anna Karénine* dans [Meir et al., 2019], nous avons conçu une heuristique d'initialisation pour la fiabilité des électeurs. Dans son livre, Léon Tolstoi affirme que "les familles heureuses sont toutes semblables ; chaque famille malheureuse est malheureuse à sa manière". Dans le même esprit, il semble raisonnable de faire l'hypothèse que les utilisateurs compétents ont tendance à donner des réponses similaires, alors que les utilisateurs imprécis ont chacun leur propre façon d'être imprécis.

Algorithm 4 EMVA

Input: Bulletins d'approbations $(A_i^z)_{1 \leq z \leq L, i \in N}$
 Paramètres initiaux $\hat{\theta}^{(0)}$, Bounds (l, u) , Tolérance ε

Output: Estimations $(\hat{S}_z), (\hat{p}_i, \hat{q}_i), (\hat{t}_j)$

1: **repeat**

2: **for** $z = 1 \dots L$ **do**

3: Calculer $\hat{S}_z^{(v+1)} = \{a_1, \dots, a_k\}$ with $k \in [l, u]$ and:

$$\begin{cases} |\hat{S}_z^{(v+1)} \cap S_{max,z}^{(v)}| &= \min(u, k_{max,z}^{(v)}) \\ |\hat{S}_z^{(v+1)} \cap S_{min,z}^{(v)}| &= \max(0, l - k_{tie,z}^{(v)} - k_{max,z}^{(v)}) \end{cases}$$

4: **end for**

5: **for** $i = 1 \dots N$ **do**

6: Mettre à jour les paramètres (p_i, q_i) étant donné $\hat{S}^{(v+1)}$:

$$\hat{p}_i^{(v+1)} = \frac{\sum_{z \in L} |A_i^z \cap \hat{S}_z^{(v+1)}|}{\sum_{z \in L} |\hat{S}_z^{(v+1)}|}, \hat{q}_i^{(v+1)} = \frac{\sum_{z \in L} |A_i^z \cap \overline{\hat{S}_z^{(v+1)}}|}{\sum_{z \in L} |\overline{\hat{S}_z^{(v+1)}}|}$$

7: **end for**

8: **for** $j = 1 \dots m$ **do**

9: Mettre à jour $\hat{t}_j^{(v+1)}$:

$$\hat{t}_j^{(v+1)} = \frac{occ^{(v+1)}(j) \bar{\alpha}_j^{(v+1)}}{occ^{(v+1)}(j) \bar{\alpha}_j^{(v+1)} + (L - occ^{(v+1)}(j)) \alpha_j^{(v+1)}}$$

où :

$$\begin{cases} occ^{(v+1)}(j) &= \sum_{z=1}^L \mathbf{1}\{a_j \in \hat{S}_z^{(v+1)}\} \\ \bar{\alpha}_j^{(v+1)} &= \beta((l-1)^+, u-1, \hat{t}_{<j}^{(v+1)}, \hat{t}_{>j}^{(v)}) \\ \alpha_j^{(v+1)} &= \beta(l, u, \hat{t}_{<j}^{(v+1)}, \hat{t}_{>j}^{(v)}) \end{cases}$$

10: **end for**

11: **until** $\|\hat{\theta}^{(v+1)} - \hat{\theta}^{(v)}\| \leq \varepsilon$



Figure 5.4: Exemple de la tâche d'annotation d'images

Pour évaluer l'importance de l'information a priori sur la taille de la vérité de base, nous avons testé l'algorithme EMVA avec des contraintes libres $(l, u) = (0, m)$ (on l'appellera $AMLE_f$) et l'algorithme $AMLE_c$ avec $(l, u) = (1, 2)$. Nous appliquons également la règle modale qui retourne le bulletin le plus fréquent, et une variante de la règle de la majorité par label qui, compte tenu d'un profil d'approbation (A_1, \dots, A_n) , produit le sous-ensemble d'alternatives S tel que :

$$a \in S \iff |i \in N, a \in A_i| > \frac{n}{2}$$

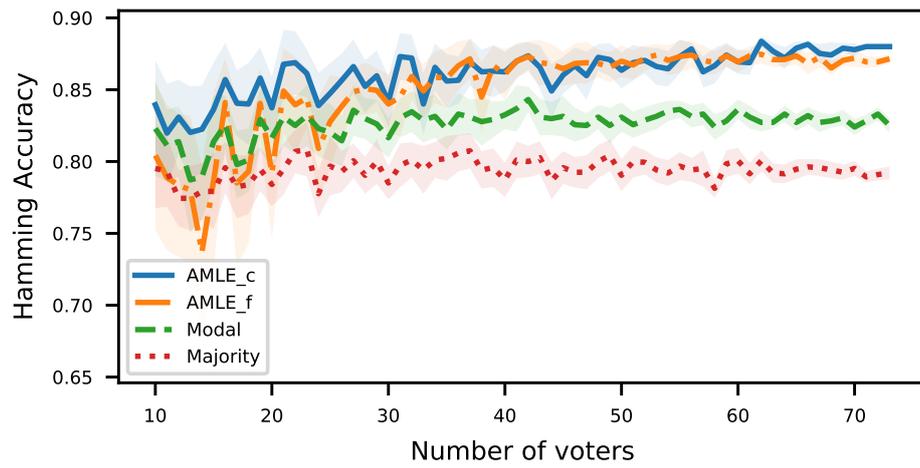
. Si ce sous-ensemble est vide, il est remplacé par l'alternative ayant le plus grand nombre d'approbations, et s'il a plus de deux alternatives, nous ne gardons que les deux premières.

Nous avons pris 20 lots de $n = 10$ à $n = 74$ électeurs tirés au hasard et nous avons appliqué les quatre méthodes à chacun d'entre eux (voir figure 5.5a, 5.5b).

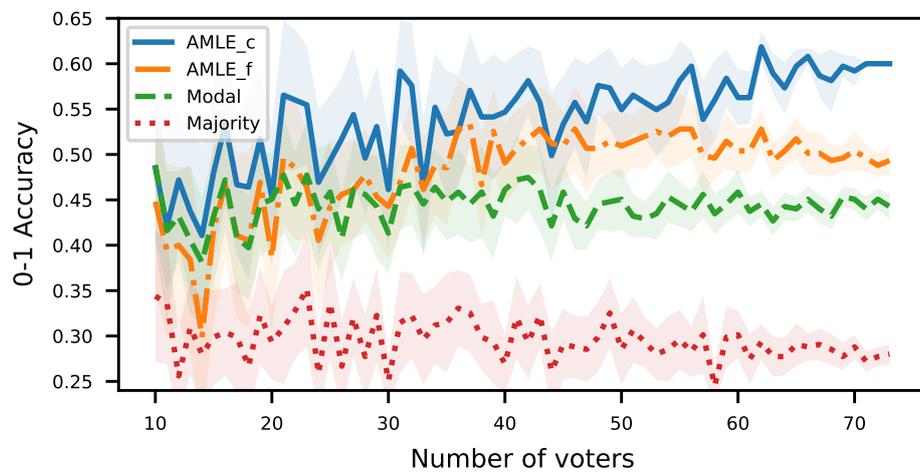
On remarque que la règle de la majorité et la règle modale sont dépassées par EMVA, ce qui peut s'expliquer par le fait qu'elles ne tiennent pas compte de la fiabilité des électeurs. La comparaison des performances de $AMLE_c$ et de $AMLE_f$ souligne l'importance de la connaissance préalable de la taille du comité pour améliorer la qualité de l'estimation.

5.10 Vote par approbation épistémique à vainqueurs multiples: les contraintes de taille comme facteur exogène

Considérons un ensemble de n électeurs $N = \{1, \dots, n\}$ et un ensemble de m alternatives $\mathcal{X} = \{a_1, \dots, a_m\}$. Parmi ces alternatives, il existe un sous-ensemble *inconnu* d'alternatives objectivement "bonnes"/"éligibles" (patients ayant réellement besoin de ressources médicales, étudiants admissibles à un programme, produits non défectueux pouvant déjà être mis en vente...). On demande aux votants d'approuver toutes les alternatives qu'ils pensent être éligibles, ce qui donne le profil d'approbation $A = (A_1, \dots, A_n)$.



(a) Précision Hamming



(b) Précision 0/1

Figure 5.5: Précisions de différentes méthodes d'agrégation

Nous supposons dans cette section qu'en raison de certaines contraintes externes (ressources limitées, satisfaction de la commande du client...), le nombre d'alternatives à sélectionner et à produire, une fois que les bulletins d'approbation des votants ont été agrégés, doit être limité par une borne inférieure donnée l et une borne supérieure donnée u .

L'objectif est donc d'agrèger les votes afin de sélectionner le "meilleur" sous-ensemble d'alternatives *admissible* (vérifiant la contrainte de taille). Pour définir formellement ce que nous entendons par "meilleur", nous définissons la notion d'utilité suivante.

Soit $\alpha^+, \alpha^-, \beta^+, \beta^- \in \mathbb{R}$ tel que $\alpha^+ > \alpha^-$ et $\beta^+ < \beta^-$. On définit l'utilité d'un sous-ensemble S comme suit:

$$u(S) = \sum_{a \in S} u^+(a) + \sum_{a \notin S} u^-(a)$$

où:

$$u^+(a) = \begin{cases} \alpha^+ & \text{if } a \in S^* \\ \beta^+ & \text{if } a \notin S^* \end{cases}, u^-(a) = \begin{cases} \alpha^- & \text{if } a \in S^* \\ \beta^- & \text{if } a \notin S^* \end{cases}$$

α^+, α^- peuvent donc être interprétés comme, respectivement, l'utilité de prendre une alternative dans S^* et le coût de la laisser. Il est à noter que:

$$u(S) = \sum_{a \in \mathcal{X}} u^-(a) + \sum_{a \in S} [u^+(a) - u^-(a)]$$

On définit $f(a) = u^+(a) - u^-(a)$ qui est égal à $\alpha = \alpha^+ - \alpha^- > 0$ si $a \in S^*$ et $-\beta = \beta^+ - \beta^- < 0$ si $a \notin S^*$. On étend f à l'ensemble des tous les sous-ensembles d'alternatives $\mathcal{S} = 2^{\mathcal{X}}$ en définissant $f(S) = \sum_{a \in S} f(a)$. On peut remarquer que maximiser $u(S)$ revient à maximiser $f(S)$.

Donc, en somme, nous voulons trouver un sous-ensemble admissible dans $\mathcal{S}_{l,u} = \{S \subseteq \mathcal{X}, l \leq |S| \leq u\}$ qui, dans un certain sens, maximise f . Cependant, le problème est que l'utilité ne peut pas être calculée directement puisqu'elle dépend de la vérité S_* qui est inconnue. Pour résoudre ce problème, nous définirons un modèle de bruit composé d'une distribution préalable sur S_* et d'une distribution conditionnelle sur les bulletins d'approbation étant donné la vérité de base, qui sera utilisée plus tard pour proposer deux concepts de solution.

5.10.1 Le modèle de bruit

Pour modéliser la fiabilité des votants, on utilise la même distribution de la section précédente à savoir:

$$P(a \in A_i | S^* = S) = \begin{cases} p_i & \text{if } a \in S \\ q_i & \text{if } a \notin S \end{cases}$$

En outre, soit $t_j = P(f(a_j) = \alpha) = P(a_j \in S_*)$ la probabilité a priori que a_j soit admissible et supposons que les événements $\{f(a_j) = \alpha\}$ sont mutuellement indépendants. Cette hypothèse est réaliste car nous n'avons pas de prior sur la taille de S^* : ainsi, par exemple, la probabilité que l'étudiant s soit objectivement au-dessus de la barre ne change pas une fois que nous connaissons le résultat (au-dessus ou au-dessous de la barre) d'un certain nombre d'autres étudiants. On

suppose aussi l'indépendance des votants. Cependant, contrairement à la section précédente, on va supposer que ces paramètres sont connus. Ils peuvent être fixés par une entité centrale ou estimés à partir d'observations antérieures.

5.10.2 Deux objectives

Rappelons que nous nous concentrons sur la recherche de solutions pour estimer le sous-ensemble admissible d'alternatives maximisant f . Dans les paragraphes suivants, nous allons présenter, étudier et comparer deux concepts de solutions différents :

- Le maximiseur le plus probable de l'utilité.
- Le maximiseur de l'utilité espérée.

5.10.2.1 Le maximiseur le plus probable de l'utilité MPPU

Soit $M_{l,u} = \max_{S \in \mathcal{S}_{l,u}} f(S)$ la valeur maximale inconnue de f pour les ensembles d'alternatives admissibles. Pour un ensemble admissible donné d'alternatives S , nous définissons $\psi(S) = P(f(S) = M_{l,u} | A)$, c'est-à-dire la probabilité que S maximise la fonction f . Notre objectif est d'estimer $S \in \mathcal{S}_{l,u}$ avec la plus grande probabilité de $f(S) = M_{l,u}$, c'est-à-dire que nous cherchons :

$$\arg \max_{S \in \mathcal{S}_{l,u}} P(f(S) = M_{l,u} | A) = \arg \max_{S \in \mathcal{S}_{l,u}} \psi(S)$$

Le premier résultat est simplement emprunté au théorème 5.11 avec une légère adaptation à cette interprétation. Il caractérise les ensembles admissibles qui ont le plus de chances de coïncider avec la vérité objectives (mais comme nous le verrons plus tard, ils ne coïncident pas toujours avec la MPPU, mais sont des candidats potentiels).

Proposition 5.15. $\tilde{S} \in \arg \max_{S \in \mathcal{S}_{l,u}} P(S_* = S | A)$ si et seulement si il existe un $k \in [l, u]$ tel que $\tilde{S} = \{a_1, \dots, a_k\}$ and:

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau_n}| = \min(u, k_{max}^{\tau_n}) \\ |\tilde{S} \cap S_{min}^{\tau_n}| = \max(0, l - k_{tie}^{\tau_n} - k_{max}^{\tau_n}) \end{cases}$$

Nous pouvons maintenant énoncer le résultat principal. Il affirme que la MPPU coïncide soit avec l'ensemble admissible le plus susceptible d'être la vérité de base (qui ont été caractérisés dans la Proposition 5.15), soit avec l'un des ensembles de top- l ou top- u alternatives en fonction de leur score d'approbation pondéré.

Theorem 5.16. Supposons que $l < u$ et soit $\tilde{S}_l = \{a_1, \dots, a_l\}$ and $\tilde{S}_u = \{a_1, \dots, a_u\}$ and $\tilde{S} = \arg \max_{S \in \mathcal{S}_{l,u}} P(S_* = S | A)$. On a:

$$\arg \max_{S \in \mathcal{S}_{l,u}} \psi(S) = \arg \max_{S \in \tilde{S} \cup \{\tilde{S}_l, \tilde{S}_u\}} \psi(S)$$

5.10.2.2 Le maximiseur de l'utilité espérée MUE

Une autre approche naturelle, au lieu de rechercher l'ensemble admissible d'alternatives $S \in \mathcal{S}_{l,u}$ avec la probabilité maximale de maximiser f , serait de rechercher l'ensemble admissible maximisant la valeur espérée de f , c'est-à-dire:

$$\operatorname{argmax}_{S \in \mathcal{S}_{l,u}} \mathbb{E}[f(S)|A]$$

A l'instar de la définition 5.4, on introduit un nouveau seuil:

$$\tau'_n = \tau_n + \log \frac{\beta}{\alpha}$$

et la partition associée:

$$\begin{cases} S_{max}^{\tau'_n} &= \{a \in \mathcal{X}, \operatorname{app}_w(a) > \tau'_n\} \\ S_{tie}^{\tau'_n} &= \{a \in \mathcal{X}, \operatorname{app}_w(a) = \tau'_n\} \\ S_{min}^{\tau'_n} &= \mathcal{X} \setminus S_{max}^{\tau'_n} \cup S_{tie}^{\tau'_n} \end{cases}$$

et on dénote $k_{max}^{\tau'_n} = |S_{max}^{\tau'_n}|$, $k_{tie}^{\tau'_n} = |S_{tie}^{\tau'_n}|$, $k_{min}^{\tau'_n} = |S_{min}^{\tau'_n}|$.

Theorem 5.17. $\tilde{S} \in \operatorname{arg max}_{S \in \mathcal{S}_{l,u}} \mathbb{E}[f(S)|A]$ si et seulement s'il existe $k \in [l, u]$ tel que $\tilde{S} = \{a_1, \dots, a_k\}$ et:

$$\begin{cases} |\tilde{S} \cap S_{max}^{\tau'_n}| &= \min(u, k_{max}^{\tau'_n}) \\ |\tilde{S} \cap S_{min}^{\tau'_n}| &= \max(0, l - k_{tie}^{\tau'_n} - k_{max}^{\tau'_n}) \end{cases}$$

En particulier, si $\alpha = \beta$, alors l'ensemble des alternatives admissibles maximisant la valeur espérée de f coïncide avec l'ensemble (admissible) des alternatives les plus susceptibles d'être la vérité de base (caractérisée dans le théorème 4.1) :

$$\operatorname{arg max}_{S \in \mathcal{S}_{l,u}} \mathbb{E}[f(S)|A] = \operatorname{arg max}_{S \in \mathcal{S}_{l,u}} P(S_* = S|A)$$

5.11 Conclusion

Après avoir passé en revue la littérature sur le choix social épistémique, nous avons proposé différentes nouvelles méthodes pour traiter à la fois le cas à un seul gagnant et le cas à plusieurs gagnants pour le vote par approbation.

Tout d'abord, nous avons introduit la méthode *size-matters* pour le cas à vainqueur unique, où plus de poids est attribué aux votants qui approuvent moins d'alternatives, et nous avons prouvé qu'elle pouvait surpasser les règles de base connues sur plusieurs jeu de données d'annotations d'images collectives.

Ensuite, nous nous sommes intéressés au problème du vote à vainqueurs multiples et nous avons distingué deux interprétations possibles de la notion de taille du comité dans le cadre épistémique. Lorsque les contraintes de taille sont considérées comme des connaissances préalables, nous avons montré comment nous pouvons les intégrer dans le modèle et les règles d'agrégation. De même, lorsque les contraintes sont exogènes, nous avons défini deux concepts de solution différents et proposé des règles optimales associées à chacun d'eux.

Part III

Appendix

Social Rankings Manipulability

This chapter presents the work I did during my master thesis, on the computational social choice problem of *social rankings*, where the goal is to rank individual agents given a ranking of the groups they form. For instance, given the records of collective performances of different crowdsourcing workers groups, social ranking rules rank the individuals from best to worst. The platform may then reward them accordingly.

More precisely, we study the robustness of the social ranking solutions to strategic behaviours. We focus on a particular kind of manipulation where an agent reaches a higher position in the output by undermining the performance of the groups to which she belongs. Robustness to this kind of malicious behaviour is crucial in a crowdsourcing setting [Kittur et al., 2008, Zhu and Carterette, 2010, Le et al., 2010, Downs et al., 2010], where spammers, aiming to maximize the payment they receive with minimum effort (*e.g.*, by randomly approving alternatives in a labelling task), can make the cost of acquiring labels very expensive and can potentially degrade the quality of the collected annotations.

Hence, even though the topic of *social ranking* is not directly related to epistemic approval voting, we include this work in the appendix because it has applications to the study of reward mechanisms in crowdsourcing and data labelling.

The results of this chapter are published in [Allouche et al., 2020].

Abstract

We investigate the issue of manipulability for social ranking rules, where the goal is to rank individuals given the ranking of coalitions formed by them and each individual prefers to reach the highest positions in the social ranking. This problem lies at the intersection of computational social choice and the algorithmic theory of power indices. Different social ranking rules have been recently proposed and studied from an axiomatic point of view. In this article, we focus on rules representing three classical approaches in social choice theory: the marginal contribution

approach, the lexicographic approach and the (*ceteris paribus*) majority one. We first consider some particular members of these families analysing their resistance to a malicious behaviour of individuals. Then, we analyze the computational complexity of manipulation, and complete our theoretical results with simulations in order to analyse the manipulation frequencies and to assess the effects of manipulations.

6.1 Introduction

In decision making and social choice theory, a number of studies are devoted to ranking individuals based on the performance of the coalitions formed by them. For instance, given values on coalitions of individuals, power indices map these values of coalitions on values of individuals. The seminal works of [Shapley, 1953] and [Banzhaf III, 1964] paved the way of a whole research domain and a related literature with many issues, including axiomatization [Laruelle and Valenciano, 2001, Holler and Packel, 1983], applications [Bilbao et al., 2002, Moretti and Patrone, 2008], algorithmic analysis [Matsui and Matsui, 2000] and computational complexity [Deng and Papadimitriou, 1994, Bachrach and Rosenschein, 2009, Faliszewski and Hemaspaandra, 2009]. The non-manipulability (or *strategy proofness*) is another fundamental issue. In social choice, since the seminal theorems of Gibbard and Satterthwaite ([Gibbard, 1973a] and [Satterthwaite, 1975b]), we know that every interesting social choice function is manipulable by misrepresentation of preferences. The manipulability is also analysed for power indices. We quote in particular the literature on the paradoxical behaviour of power indices under the modification of some elements of the game, like the number of players or the size of coalitions [Felsenthal et al., 1998, Laruelle and Valenciano, 2005], or the study of manipulation in weighted voting games [Aziz et al., 2011, Zuckerman et al., 2012]. In these models, players are analyzed from a strategic perspective to establish under which conditions they can increase their power adopting malicious behaviors like, for example, splitting or merging.

Power indices (and other indices of individual productivity based on the evaluation of revenues generated by teams [Flores-Szwagrzak and Treibich, 2020]) require a numerical evaluation of coalitions of individuals. Following classical situations in social choice where ordinal data are provided (for instance, voting theory), several articles address the question of defining ordinal notions of power indices when we only have ordinal information over coalitions. This has been formalized as the *social ranking from coalitions* (SRC) problem, where the objective is to evaluate the “influence” of individuals involved in a collective decision process like an electoral system, a parliament, a governing council, a management board, etc. ([Moretti, 2015, Moretti and Öztürk, 2017]). Basically, an SRC problem consists of a finite set N of individuals and a binary relation \succeq over some subsets (hereafter called *coalitions*) of N ; the binary relation \succeq is called *power relation* and represents the relative power of coalitions in a decision process. A *solution* or *rule* for an SRC problem is a “suitable” method aimed to convert the information contained in a power relation \succeq into a ranking over the single elements of N representing their overall individual power.

Several solutions for SRC problems have been proposed in the literature. For instance, in the

work by [Haret et al., 2018] (and the one by [Fayard and Öztürk, 2018]) two individuals are compared using information from subsets ranking them under a *ceteris paribus* interpretation. [Bernardi et al., 2019] axiomatically characterize a solution based on the idea that the most influential individuals are those appearing more frequently in the highest positions of the power relation. A rule based on the idea of ordinal marginal contribution and analogous to the Banzhaf index has been recently introduced in the paper by [Khani et al., 2019].

Following these lines of research, in this chapter we are interested in the analysis of the *strategic manipulation* of SRC rules, in the sense that an individual may be interested in behaving maliciously within one or more teams (weakening their group's effectiveness) in order to obtain a better position in the individual ranking. The notion of manipulability for SRC considered in this paper assumes that an individual can only weaken the ranking of teams to which she belongs. In other words, an individual i cannot affect the performances of teams not containing i and, in addition, cannot improve the position of a coalition already containing i .

Example 6.1. Consider a manager who must decide how to allocate two bonuses over three employees (denoted by 1, 2 and 3). Suppose that she can only compare the performance of teams in an ordinal way: $\{1, 2, 3\} \succ \{1, 3\} \succ \{1\} \succ \{1, 2\} \succ \{2, 3\} \succ \{2\} \succ \{3\} \sim \emptyset^1$. Suppose the manager wants to keep into account the attitude of employees to cooperate. So, an option is to count the number of (ordinal) positive and negative marginal contributions provided by each employee to all possible coalitions, i.e. 1 contributes positively to four teams (i.e., $\{1, 2, 3\} \succ \{2, 3\}$, $\{1, 3\} \succ \{3\}$, $\{1, 2\} \succ \{2\}$ and $\{1\} \succ \emptyset$), 3 contributes positively to three coalitions while 2 also contributes positively to three coalitions, but negatively to coalition $\{1, 2\}$. Therefore the manager would end up to award players 1 and 3. Such a rule could push individual 2 to behave strategically and to undermine the cooperation within coalition $\{2, 3\}$. So the new ranking being $\{1, 2, 3\} \succ' \{1, 3\} \succ' \{1\} \succ' \{1, 2\} \succ' \{2\} \succ' \{2, 3\} \succ' \{3\} \sim' \emptyset$, individuals 1 and 2 should now get the bonus.

To our knowledge, this chapter is the first one which investigates the manipulation for SRC rules. As in social choice, the manipulability is an important issue in many real-world coalitional frameworks. For instance, within a parliament, small political parties are often able to blackmail a majority coalition by threatening to withdraw from the coalition or to run a candidate of their own [Ferrara, 2004]. In international organizations, the capacity to block the proposals of a group ensures a great visibility of the blocking state's positions. In any organizational context, like a business company or an academic institution, any perceived disparity about the merits of a team's members (e.g., over-remuneration of a leader) may engender internal competition, or even the sabotage of the team by some of it's members. As a sensible example of this kind of behaviors, consider the situation following the Italian general election of March 2018, where no political party got an absolute majority. After 3 months of negotiation, a coalition government was finally formed by the two parties casting the highest number of votes, precisely, the Five Star Movement and the League. However, such a government ended one year after, in August 2019, when the League withdrew its support to the coalition government. As a consequence, the League substantially increased its position in the opinion polls done immediately after the

¹ \succ : strict preference, \sim : indifference.

government crisis [Chiaramonte et al., 2020].

The chapter is organized as follows. We present basic notions and notations in Section 6.2. In Section 6.3, we introduce a formal definition of manipulability for an SRC rule and provide theoretical results on four social ranking rules: Copeland-like, Kramer-Simpson-like, Lexicographic Excellence [Bernardi et al., 2019] and Ordinal Banzhaf [Khani et al., 2019]. These social ranking rules display a wide variety of characteristics. The two first ones are based on *ceteris paribus* comparisons and use the majority principle in a different way. Copeland-like solution is a kind of flow analysis of majority graphs, whereas Kramer-Simpson is a minmax score. Ordinal Banzhaf rule is based on a marginal contribution principle. Lexicographic Excellence (lexcel) considers only information from the best ranked coalitions. We show that only lexcel is not manipulable. In Section 6.4, we analyse the computational complexity of manipulating each of the three manipulable social ranking rules, and prove that for each of them determining whether an individual can manipulate or not is an NP-hard problem. In Section 6.5, we present some simulation results on manipulable social ranking rules showing the manipulation frequencies and their vulnerability against the manipulation. Section 6.6 concludes the chapter.

6.2 Preliminaries

Let $N = \{1, \dots, n\}$ be a finite set of elements called *individuals* and let $R \subseteq N \times N$ be a *binary relation* on N .

A *preorder* is a reflexive and transitive binary relation. A preorder that is total is called *total preorder*. An antisymmetric² total preorder is called *linear order*. We denote by $\mathcal{T}(N)$ the set of all total preorders on N and by 2^N the powerset of N , i.e. the set of all subsets (also called, *coalitions*) of N . Let $\mathcal{P} \subseteq 2^N$ be a non-empty collection of subsets of N . A *power relation* on \mathcal{P} is a total preorder $\succeq \subseteq \mathcal{P} \times \mathcal{P}$. We denote by $\mathcal{T}(\mathcal{P})$ the family of all power relations on every non-empty collection $\mathcal{P} \subseteq 2^N$. Given a *power relation* $\succeq \in \mathcal{T}(\mathcal{P})$ on $\mathcal{P} \subseteq 2^N$, we denote by \sim its symmetric part (i.e. $S \sim T$ if $S \succeq T$ and $T \succeq S$) and by \succ its asymmetric part (i.e. $S \succeq T$ and not $T \succeq S$). So, for each pair of subsets $S, T \in \mathcal{P}$, $S \succ T$ means that S is *strictly stronger* than T , whereas $S \sim T$ means that S and T are *indifferent*.

Let $\succeq \in \mathcal{T}(\mathcal{P})$ be of the form $S_1 \succeq S_2 \succeq S_3 \succeq \dots \succeq S_{|\mathcal{P}|}$. The *quotient order* of \succeq is denoted as $\Sigma_1 \succ \Sigma_2 \succ \Sigma_3 \succ \dots \succ \Sigma_m$ in which the subsets S_j are grouped in the *equivalence classes* Σ_k generated by the symmetric part of \succeq . This means that all the sets in Σ_1 are indifferent to S_1 and are strictly better than the sets in Σ_2 and so on. So, $\Sigma_i = S_i$ for any $i = 1, \dots, |\mathcal{P}|$ if and only if \succeq is a linear order.

A *social ranking solution* or *solution on N* , is a function $R : \mathcal{T}(\mathcal{P}) \rightarrow \mathcal{T}(N)$ associating to each power relation $\succeq \in \mathcal{T}(\mathcal{P})$ a total preorder $R(\succeq)$ (or R^\succeq) over the elements of N . By this definition, the notion $iR^\succeq j$ means that applying the social ranking solution to the power relation \succeq gives the result that i is ranked higher than or equal to j . We denote by I^\succeq the symmetric part

² $\forall i, j \in N, iRj \text{ and } jRi \Rightarrow i = j.$

of R^\succeq , and by P^\succeq its asymmetric part. The *social score* $p_i(R^\succeq)$ of individual $i \in N$ in R^\succeq is defined as the number of individuals in $N \setminus \{i\}$ that are ranked lower than i minus the number of individuals in $N \setminus \{i\}$ that are ranked higher than i , that is

$$p_i(R^\succeq) = |\{j \in N \setminus \{i\} : iR^\succeq j\}| - |\{j \in N \setminus \{i\} : jR^\succeq i\}|.$$

6.3 Manipulability

In this chapter we focus on a particular notion of manipulation, intended as the “unlimited” capacity of individuals to undermine the position of coalitions to which they belong in a power relation \succeq on $\mathcal{P} \subseteq 2^N$.

Definition 6.1. Let $\succeq \in \mathcal{T}(\mathcal{P})$ on \mathcal{P} be a power relation with the associated quotient order \succ :

$$\Sigma_1 \succ \Sigma_2 \succ \dots \succ \Sigma_j \succ \dots \succ \Sigma_m. \quad (6.1)$$

Let i be an individual, and $\mathcal{C} \subseteq \mathcal{P}$ be a collection of coalitions in \mathcal{P} all containing i . For all $S \in \mathcal{C}$, let $j(S) \in \{1, \dots, m\}$ be such that $S \in \Sigma_{j(S)}$ ³.

A *manipulation* of \succeq by individual i via collection \mathcal{C} is another power relation \succeq^c on \mathcal{P} with $\succeq^c \neq \succeq$ and with the associated quotient order \succ^c such that the following two conditions hold:

- i) $\Sigma_1 \setminus \mathcal{C} \succ^c \Sigma_2 \setminus \mathcal{C} \succ^c \dots \succ^c \Sigma_j \setminus \mathcal{C} \succ^c \dots \succ^c \Sigma_m \setminus \mathcal{C}$;
- ii) $T \succ^c S$ for all $S \in \mathcal{C}$ and $T \in \bigcup_{i=1, \dots, j(S)} \Sigma_i \setminus \mathcal{C}$.

A social ranking \mathcal{R} is *manipulable by i* on a power relation \succeq on \mathcal{P} if there exists a collection of coalitions $\mathcal{C} \subseteq \mathcal{P}$ containing i , a manipulation \succeq^c of \succeq by i via \mathcal{C} such that

$$p_i(R^{\succeq^c}) > p_i(R^\succeq).$$

[recall that $p_i(R^\succeq) = |\{j, iR^\succeq j\}| - |\{j, jR^\succeq i\}|$.]

A social ranking solution \mathcal{R} is *manipulable* on a power relation \succeq if it is manipulable by some individual i .

Condition (ii) says that \succeq^c is obtained from \succeq moving each coalition $S \in \mathcal{C}$ from the equivalence class to which it belongs in \succeq , to a strictly lower equivalence class (that can also be a new singleton equivalence class containing only S in \succeq^c), while the relation among all the other coalitions not in \mathcal{C} is maintained as in \succeq (condition (i)). The family of all manipulations of \succeq via collection \mathcal{C} is denoted by $\mathcal{M}_{\mathcal{C}}(\succeq)$.

Example 6.2. Consider the power relation \succeq such that⁴ $23 \succ (123 \sim 12) \succ 13$ (hence, $\Sigma_1 = \{23\}$, $\Sigma_2 = \{123, 12\}$, $\Sigma_3 = \{13\}$). Imagine that individual 1 wants to manipulate by

³ $j(S)$ represents the rank of the equivalence class to which S belongs in the initial power relation \succeq .

⁴To avoid cumbersome notations later, sets will be written for short without commas and parentheses, e.g., 123 instead of $\{1, 2, 3\}$, and $S \cup i$ instead of $S \cup \{i\}$.

deteriorating the positions of coalitions in $\mathcal{C} = \{12, 13\}$. Condition i) of Definition 6.1 imposes to maintain $23 \succ^c 123$, condition ii) imposes $23 \succ^c 12$, $23 \succ^c 13$, $123 \succ^c 12$ and $123 \succ^c 13$. Hence, the family of all possible manipulations of \succeq by 1 via collection \mathcal{C} is $\mathcal{M}_{\mathcal{C}}(\succeq) = \{\succeq_a^c, \succeq_b^c, \succeq_c^c\}$, with $23 \succ_a^c 123 \succ_a^c 12 \succ_a^c 13$, $23 \succ_b^c 123 \succ_b^c 13 \succ_b^c 12$ and $23 \succ_c^c 123 \succ_c^c 12 \sim_c^c 13$.

We will now analyse the manipulability of different social ranking rules.

6.3.1 Copeland-like and Kramer-Simpson-like rules

Copeland-like and Kramer-Simpson (KS)-like rules are both based on *Ceteris Paribus*-majority relation, where individuals i and j are ranked according to their relative success over comparisons of coalitions of the type $S \cup i$ vs. $S \cup j$ (CP-comparisons), more precisely:

Definition 6.2 (CP-Majority [Haret et al., 2018]). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *Ceteris Paribus* (CP-) majority relation is the binary relation $R_{CP}^{\succeq} \subseteq N \times N$ such that for all $i, j \in N$:

$$iR_{CP}^{\succeq}j \Leftrightarrow d_{ij}(\succeq) \geq d_{ji}(\succeq),$$

where $d_{ij}(\succeq)$ represents the cardinality of the set $D_{ij}(\succeq)$, the set of all coalitions $S \in 2^{N-\{i,j\}}$ for which $S \cup i \succeq S \cup j$ ⁵.

Example 6.3. Consider: $(123 \sim 12 \sim 3 \sim 1) \succ (2 \sim 23) \succ 13$. Then, we obtain:

$$1I_{CP}^{\succeq}2, 1P_{CP}^{\succeq}3, 2I_{CP}^{\succeq}3$$

For instance, $1I_{CP}^{\succeq}2 : D_{12}(\succeq) = \{\emptyset\}$ ($1 \succ 2$); $D_{21}(\succeq) = \{3\}$ ($23 \succ 13$).

The CP-Majority relation has a major drawback: it can generate cycles within the individual ranking (except under some particular domain restrictions, as suggested in [Haret et al., 2018]).

For this reason, we investigate the manipulability of two transitive solutions derived from the CP-Majority, which are inspired, respectively, by the Copeland [Copeland, 1951] and Kramer-Simpson [Simpson, 1969] [Kramer, 1977] voting schemes. These two rules are known to be Condorcet coherent, meaning that when a Condorcet winner (a candidate beating all the other candidates by the majority rule) exists, it is chosen by them. Interestingly, while it can be easily proved that CP-majority relation is not manipulable, Copeland like and KS like solutions are manipulable.

6.3.1.1 Copeland-like method

Strongly inspired by the Copeland score of social choice theory, we define Copeland-like solution based on the net flow of CP-majority graph. According to the Copeland solution, individuals are ordered according to the number of pairwise winning comparisons, minus the one of pairwise losing comparisons, over the set of all CP-comparisons.

⁵Note that S can be \emptyset .

Definition 6.3 (Copeland-like solution). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *Copeland-like* relation is the binary relation $R_{Cop}^{\succeq} \subseteq N \times N$ such that for all $i, j \in N$:

$$iR_{Cop}^{\succeq}j \Leftrightarrow Score_{cop}^{\succeq}(i) \geq Score_{cop}^{\succeq}(j).$$

where $Score_{cop}^{\succeq}(i) = p_i(R_{CP}^{\succeq}) =$

$$|\{j \in N \setminus \{i\} : iR_{CP}^{\succeq}j\}| - |\{j \in N \setminus \{i\} : jR_{CP}^{\succeq}i\}|.$$

Theorem 6.1. *The Copeland-like solution is manipulable.*

Proof. See Example 6.4 for an instance of manipulation. □

Example 6.4. *Consider \succeq of Example 6.3. Then,*

$$Score_{cop}^{\succeq}(1) = 1, Score_{cop}^{\succeq}(2) = 0, Score_{cop}^{\succeq}(3) = -1.$$

Hence, the Copeland-like relation is: $1P_{Cop}^{\succeq}2P_{Cop}^{\succeq}3$.

Now imagine that 3 deteriorates the performance of 23 ($C = \{23\}$)

$$(123 \sim^c 12 \sim^c 3 \sim^c 1) \succ^c 2 \succ^c 13 \succ^c 23.$$

Now we have: $Score_{cop}^{\succeq^c}(1) = 2, Score_{cop}^{\succeq^c}(2) = -1, Score_{cop}^{\succeq^c}(3) = -1$.

So, now, 3 shares the second position with 2.

6.3.1.2 Kramer-Simpson-like method

Strongly inspired by the Kramer-Simpson method of social choice theory (*Minmax*), individuals are ranked inversely to their greatest pairwise defeat over all possible CP-comparisons.

Definition 6.4 (Kramer-Simpson-like solution). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *KS-like* relation (KS relation) is the binary relation $R_{KS}^{\succeq} \subseteq N \times N$ such that for all $i, j \in N$:

$$iR_{KS}^{\succeq}j \Leftrightarrow Score_{KS}^{\succeq}(i) \leq Score_{KS}^{\succeq}(j),$$

where $Score_{KS}^{\succeq}(i) = \max_{j \in N} (d_{ji}(\succeq))$

Theorem 6.2. *The Kramer-Simpson (KS)-like solution is manipulable.*

Proof. See Example 6.5 for an instance of manipulation. □

Example 6.5. Consider $\succeq : 2 \succ (1 \sim 3) \succ 12 \succ (13 \sim 23) \succ \emptyset \succ 123$. Then, $Score_{KS}^{\succeq}(1) = 1$, $Score_{KS}^{\succeq}(2) = 0$, $Score_{KS}^{\succeq}(3) = 2$. Hence, KS-like solution is: $2P_{KS}^{\succeq}1P_{KS}^{\succeq}3$. Now consider the following manipulation operated by 1 on $\mathcal{C} = \{12\}$:

$$2 \succ^c (1 \sim^c 3) \succ^c (13 \sim^c 23) \succ^c 12 \succ^c \emptyset \succ^c 123.$$

The new scores are: $Score_{KS}^{\succeq^c}(1) = 1$, $Score_{KS}^{\succeq^c}(2) = 1$, $Score_{KS}^{\succeq^c}(3) = 1$. So, individual 1 now gets the first position.

Remark. For $n = 2$, the Copeland-like solution and the KS-like solution coincide with the CP-majority relation, hence these solutions are not manipulable for $n = 2$.

6.3.2 Ordinal Banzhaf

In the same spirit of the Banzhaf index [Banzhaf III, 1964], the ordinal Banzhaf solution is based on counting the number of positive and negative ordinal marginal contributions.

Definition 6.5 (Ordinal marginal contribution [Khani et al., 2019]). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *ordinal marginal contribution* $m_i^S(\succeq)$ of player i w.r.t. coalition S , $i \notin S$, in power relation \succeq is defined as:

$$m_i^S(\succeq) = \begin{cases} 1 & \text{if } S \cup \{i\} \succ S, \\ -1 & \text{if } S \succ S \cup \{i\}, \\ 0 & \text{otherwise.} \end{cases} \quad (6.2)$$

Definition 6.6 (Ordinal Banzhaf relation). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *ordinal Banzhaf relation* is the binary relation R_{Banz}^{\succeq} such that for all $i, j \in N$:

$$iR_{Banz}^{\succeq}j \Leftrightarrow Score_{Banz}^{\succeq}(i) \geq Score_{Banz}^{\succeq}(j),$$

where $Score_{Banz}^{\succeq}(i) = u_i^{+,\succeq} - u_i^{-,\succeq}$ and $u_i^{+,\succeq}$ ($u_i^{-,\succeq}$) is defined as the number of coalitions S with $i \notin S$ such that $m_i^S(\succeq) = 1$ ($m_i^S(\succeq) = -1$).

Theorem 6.3. The Ordinal-Banzhaf solution is manipulable.

Proof. See Example 6.6 for an instance of manipulation. □

Example 6.6. Consider $\succeq : 13 \succ 1 \succ 12 \succ 23 \succ 2 \succ 3 \succ 123 \succ \emptyset$.

Then, $1I_{Banz}^{\succeq}3P_{Banz}^{\succeq}2$ since $Score_{Banz}^{\succeq}(1) = 2$, $Score_{Banz}^{\succeq}(2) = 0$, $Score_{Banz}^{\succeq}(3) = 2$. However, if 2 undermines the cooperation with 1 and 3 ($\mathcal{C} = \{12, 23\}$): $13 \succ^c 1 \succ^c 2 \succ^c 23 \succ^c 3 \succ^c 12 \succ^c 123 \succ^c \emptyset$. Then the three individuals would have a null Banzhaf score and would be ranked equally.

6.3.3 Lexicographic excellence solution

The idea of lexicographic excellence is based on the lexicographic comparison of the frequency of individuals within equivalence classes, and taking care to reward individuals within the most excellent ones. Given the power relation \succeq and its associated quotient ranking $\Sigma_1 \succ \Sigma_2 \succ \Sigma_3 \succ \dots \succ \Sigma_m$, we denote by i_k the number of sets in Σ_k containing i :

$$i_k = |\{S \in \Sigma_k : i \in S\}|$$

for $k = 1, \dots, l$. Now, let $\theta^\succeq(i)$ be the l -dimensional vector $\theta^\succeq(i) = (i_1, \dots, i_l)$ associated to \succeq . Consider the lexicographic order \geq_L among vectors \mathbf{i} and \mathbf{j} : $\mathbf{i} \geq_L \mathbf{j}$ if either $\mathbf{i} = \mathbf{j}$ or there exists $t : i_r = j_r, r = 1, \dots, t-1$, and $i_t > j_t$.

Definition 6.7 (Lexicographic-excellence solution [Bernardi et al., 2019]). Let $\succeq \in \mathcal{T}(\mathcal{P})$. The *lexicographic excellence (lexcel) relation* is the binary relation R_{lexcel}^\succeq such that for all $i, j \in N$:

$$iR_{lexcel}^\succeq j \Leftrightarrow \theta^\succeq(i) \geq_L \theta^\succeq(j).$$

Example 6.7. Consider the power relation of Example 6.2. We have $\theta^\succeq(1) = (0, 2, 1)$ (since 1 is twice in Σ_2 and once in Σ_3), $\theta^\succeq(2) = (1, 2, 0)$, $\theta^\succeq(3) = (1, 1, 1)$, which yields the following lexcel ranking: $2P_{lexcel}^\succeq 3P_{lexcel}^\succeq 1$.

Theorem 6.4. The lexcel solution is not manipulable.

Proof. Let $\succeq \in \mathcal{T}(\mathcal{P})$ be a power relation on $\mathcal{P} \subseteq 2^N$ with the associated quotient order \succ :

$$\Sigma_1 \succ \Sigma_2 \succ \dots \succ \Sigma_j \succ \dots \succ \Sigma_k \dots \succ \Sigma_m, \quad (6.3)$$

and let $\mathcal{C} = \{S_1, S_2, \dots, S_l\} \subseteq \mathcal{P}$ and $i \in N$ be such that $i \in \bigcap_{S \in \mathcal{C}} S$ (wlog, assume $S_1 \succeq S_2 \succeq \dots \succeq S_l$).

Suppose there exists a manipulation \succeq^c of \succeq by i such that $p_i(R^{\succeq^c}) \geq p_i(R^\succeq)$. Then there must be some $k \in N \setminus \{i\}$ such that

$$kR_{lexcel}^\succeq i \quad \text{and} \quad iR_{lexcel}^{\succeq^c} k. \quad (6.4)$$

First, notice that there exists some coalition $S \in \mathcal{C}$ such that $k \notin S$ (otherwise, if $i, k \in \bigcap_{S \in \mathcal{C}} S$, the manipulation would have no impact on the relative comparison of i and k , since in this case $\theta^\succeq(i) \geq_L \theta^\succeq(k) \Leftrightarrow \theta^{\succeq^c}(i) \geq_L \theta^{\succeq^c}(k)$).

Now let $S^* \in \mathcal{C}$ be a coalition not containing k with the smallest index in \mathcal{C} , and let $j(S^*) \in \{1, \dots, m\}$ be such that $S^* \in \Sigma_{j(S^*)}$. Since $kR_{lexcel}^\succeq i$, we distinguish two cases:

- i) $kI_{lexcel}^\succeq i$ (k and i are indifferent in \succeq according to the lexcel relation). Then, by Definition 6.7, $\theta_v^{\succeq^c}(k) = \theta_v^\succeq(k) = \theta_v^\succeq(i) = \theta_v^{\succeq^c}(i)$ for all $v < j(S^*)$, while $\theta_{j(S^*)}^{\succeq^c}(k) > \theta_{j(S^*)}^{\succeq^c}(i)$. So, $kP_{lexcel}^{\succeq^c} i$;

- ii) $kP_{lexcel}^{\succeq}i$ (k is strictly stronger than i in \succeq according to the lexcel relation). So, let t be the smallest index such that $\theta_t^{\succeq}(k) > \theta_t^{\succeq}(i)$. Moreover, let $q = \min(t, j(S^*))$. By Definition 6.7, $\theta_v^{\succeq^c}(k) = \theta_v^{\succeq}(k) = \theta_v^{\succeq}(i) = \theta_v^{\succeq^c}(i)$ for all $v < q$ and $\theta_q^{\succeq^c}(k) > \theta_q^{\succeq^c}(i)$. So, again, $kP_{lexcel}^{\succeq^c}i$.

In both cases we get a contradiction with the fact that $iR_{lexcel}^{\succeq^c}k$ by relation (6.4). Since \mathcal{C} is an arbitrary collection of coalitions in \mathcal{P} , we have proved that the lexcel solution is not manipulable. \square

A drawback of the lexcel solution is that it makes use of a limited amount of information, giving a quasi-dictatorial power to the best ranked coalitions. On the other hand, Copeland-like and KS-like rules are based on the intuitive idea of *ceteris paribus* comparisons and have the advantage of choosing the Condorcet winner (with respect to the CP-majority relation) when it exists. The Ordinal Banzhaf rule, based on the intuitive notion of ordinal marginal contribution of individuals, also offers another interesting interpretation in terms of fairness. In order to further study the effective impact of manipulation on these three appealing solutions, in Section 6.4 we analyse the computational complexity of manipulation and in Section 6.5 we introduce some simulations on several numerical instances.

6.4 Computational complexity of manipulation

As we have seen in the previous section, the Copeland-like solution, the KS-like solution, and the Ordinal Banzhaf solution are manipulable. We strengthen these results in this section by showing that, for each of these social ranking solutions, determining whether an individual can manipulate or not is an NP-hard problem. Let us state the problem precisely. As an instance, we have a set $N = \{1, \dots, n\}$ of individuals with a manipulator $t \in N$, a set $\mathcal{P} \subseteq 2^N$, and a power relation \succeq on \mathcal{P} . The question is to determine whether a given solution is manipulable by t on \mathcal{P} , as defined in Definition 6.1.

We prove in the next three theorems that the manipulation is NP-hard for the three solutions.

Theorem 6.5. *For the Ordinal Banzhaf solution, the Manipulation problem is NP-hard.*

Proof. We build the following instance of the manipulation problem under the Ordinal Banzhaf solution. First, let us consider the following individuals:

- we associate to each edge $e_i \in E$ an individual that we call e_i as well (for convenience), and to each vertex v an individual that we call v as well (for convenience);
- two other individuals: t (the manipulator), and α .

For each vertex v , let us call \mathcal{P}_v the set containing the subsets of individuals $\{v, \alpha, t\}$, $\{v, t\}$ and all subsets $\{v, e_i, \alpha, t\}$ for each edge e_i incident to v . \mathcal{P}_v is ordered as follows in \succeq : $\{v, \alpha, t\}$ is

the first one (strictly preferred to any other sets in \mathcal{P}_v), $\{v, t\}$ is the last one, and all $\{v, e_i, \alpha, t\}$ are equivalent, ranked between $\{v, \alpha, t\}$ and $\{v, t\}$.

Note that each set in \mathcal{P}_v contains v , so for the scores the relative positions of 2 sets in \mathcal{P}_v and $\mathcal{P}_{v'}$ do not matter; we do not specify it.

The contribution of these sets \mathcal{P}_v to the scores are: $+|V|$ for α (due to $\{v, \alpha, t\} \succ \{v, t\}$ for each vertex/individual v), -2 for each object e_i (due to $\{v, \alpha, t\} \succ \{v, e_i, \alpha, t\}$, for each of the two extremities of edge e_i), and 0 for t (each set contains t).

The idea of the reduction is that, in order to manipulate, t has to become first (defeating α). To do this she shall put $\{v, \alpha, t\}$ below $\{v, t\}$ in some \mathcal{P}_v . But doing this, the score of the edges incident to v increases. t cannot do this for the two extremities of an edge, otherwise e_i defeats him.

To make this true, we need to add dummy individuals to adjust the initial scores of α and t . For the score of α , we add $\lambda = 2k - |V|$ individuals b_1, \dots, b_λ . For each b_i , we order $\{b_i, \alpha\} \succ \{b_i\}$. This gives an extra score of $\lambda = 2k - |V|$ to α , while b_i has score 0 . Finally, we add an object γ , and order $\{\gamma, t\} \succ \{\gamma\}$, giving an extra score of 1 to the manipulator t .

Note that as previously we do not need to further specify \succeq , since the relative positions of sets containing different b_i , and/or γ , and/or in different \mathcal{P}_v , does not matter with respect to the scores (there is no other set inclusion).

To sum up, we have $|E| + |V| + 2 + (2k - |V|) + 1$ individuals: each individual e_i has score -2 , individual t has score 1 , and α has score $|V| + 2k - |V| = 2k$. All dummy objects and objects v have score 0 . Note that the size of \mathcal{P} is polynomial in the size of G .

We claim that t can manipulate if and only if there is an independent set of size k in G .

Suppose that there is an independent set S of size k in G . Then consider the manipulation where, for each v in S , t puts $\{v, e_i^j, \alpha, t\}$ down to the last position in \mathcal{P}_v . Then the score of α decreases by $2k$ and becomes 0 . The score of e_i is modified in at most one \mathcal{P}_v , since S is an independent set, so it is at most 0 . The score of t is still 1 , and t manipulated the election.

Conversely, suppose that t can manipulate. Note that t cannot increase her own score, so she must make the score of α at most 1 . This means that she has to put $\{v, \alpha, t\}$ in the last position in at least k sets \mathcal{P}_v . Let S be the corresponding set of vertices. If S contains both extremities of one edge e_i , then the score of e_i becomes $+2$, and t is not better off. So, in order to manipulate, the set S must be an independent set, and it is of size at least k . \square

Theorem 6.6. *For the Copeland-like solution, the Manipulation problem is NP-hard.*

Proof. We make a reduction from the NP-complete vertex cover problem where, given a graph $G = (V, E)$ and an integer k , the goal is to determine whether there exists a set C of k vertices such that each edge has at least one endpoint in C (vertex cover).

Given a graph G and an integer k , we build an instance of the manipulation problem under Copeland-like solution as follows. We first consider two individuals v, \bar{v} for each vertex v , one

individual e_i for each edge e_i , and three individuals t , α and j . t is the manipulator.

For each vertex v , we define \mathcal{P}_v as the set containing the subsets of individuals $\{v, \bar{v}, \alpha, t\}$, $\{v, \bar{v}, \alpha, j\}$ and all subsets $\{v, \bar{v}, e_i, t\}$ for each edge e_i incident to v . \mathcal{P}_v is ordered as follows in \succeq : $\{v, \bar{v}, \alpha, t\}$ is the first one, $\{v, \bar{v}, \alpha, j\}$ is the second one, and all $\{v, \bar{v}, e_i, t\}$ (e_i incident to v) are equivalent, in third position. Note that, since all sets in \mathcal{P}_v contain v and \bar{v} , the relative positions of sets in \mathcal{P}_v and $\mathcal{P}_{v'}$ do not matter, and we do not specify them.

Up to now, we have: $d_{tj} = -d_{jt} = |V|$ (1 for each v), $d_{\alpha e_i} = -d_{e_i \alpha} = 2$ (1 for each extremity of e_i). We also have non zero values for some $d_{e_i e_j}$ but this will not matter.

Now, we add dummy individuals to modify d_{jt} : we add $|V| - k - 1$ pairs of individuals $\beta_i, \bar{\beta}_i$ with $\{\beta_i, \bar{\beta}_i, j\} \succ \{\beta_i, \bar{\beta}_i, t\}$. Then $d_{tj} = |V| + k + 1 - |V| = k + 1$. In a similar manner, we add dummy individuals to modify $d_{\alpha e_i}$, in order to make $d_{\alpha e_i} = d_{e_i \alpha} = 0$ for all edges e_i .

Finally, we add dummy individuals in order to increase the (social) Copeland-scores of j , t and α . In particular, the Copeland-scores of other individuals (including e_i) will not matter, as they are much lower. We add $\{\delta_i, \bar{\delta}_i, \alpha\} \succ \{\delta_i, \bar{\delta}_i, \phi_i\}$ for $i = 1, \dots, 3|E|$. This adds $3|E|$ to the Copeland-score of α . Similarly, we add $2|E|$ to the Copeland-score of t , and $2|E| + 1$ to the Copeland-score of j . Then:

- All individuals up to j , α and t have Copeland-score at most $|E| + 1$.
- t has Copeland-Score $1 + 2|E|$ (he defeats j), α has Copeland-Score $3|E|$ (he is tied with all edges e_i), and j has Copeland-Score $-1 + 2|E| + 1 = 2|E|$ (t defeats him).

Hence, t is ranked in second position (with no tie) according to Copeland-scores.

We claim that t can manipulate if and only if there is a vertex cover of size k in G .

If there is such a vertex cover C , let us consider the manipulation where t puts $\{v, \bar{v}, \alpha, t\}$ in the last position in \mathcal{P}_v for each $v \in C$. Then: d_{tj} decreases by k , so $d_{tj} = -d_{jt} = 1$ and t still defeats j . Since C is a vertex cover, for each edge e_i $d_{\alpha e_i}$ decreased by at least one, and becomes negative. So now each e_i defeats α , and α has a Copeland-score of $3|E| - |E| = 2|E|$, tied with t . t is now ranked first, tied only with α , and is better of.

Suppose now that t can manipulate. Note that when manipulating, t cannot increase his Copeland-score (the d_{ik} cannot increase). So, he must reduce the Copeland-Score of α by at least $|E|$. This decreasing can only be obtained by putting some $\{v, \bar{v}, \alpha, t\}$ below $\{v, \bar{v}, e_i, t\}$ for *each* edge e_i (so that the global decreasing is $|E|$). Let us call C the set of vertices for which t makes such a modification. Since it must affect each edge e_i , C is a vertex cover of G . Suppose that it has size $|C| > k$. Then d_{tj} decreases by at least $k + 1$, so now $d_{tj} = d_{jt} \leq 0$. t has Copeland-Score (at most) $2|E|$, while j has Copeland-Score (at least) $2|E| + 1$, and t is not better of in the global ranking, a contradiction. So C is a vertex cover of G of size (at most) k . \square

Theorem 6.7. *For Kramer-Simpson-like solution, the Manipulation problem is NP-hard.*

Proof. We reduce the independent set problem where, given a graph $G = (V, E)$ and an integer k , the goal is to determine whether there exists a set S of k pairwise non adjacent vertices (independent set).

Given a graph G and an integer k , we build an instance of the manipulation problem under Kramer-Simpson-like solution as follows. We first consider two individuals v, \bar{v} for each vertex v , two individuals e_i, f_i for each edge e_i , and three individuals t (the manipulator), α and β .

For each vertex v , we define \mathcal{P}_v as the set containing the subsets of individuals $\{v, \bar{v}, \alpha, t\}$, $\{v, \bar{v}, \beta, t\}$ and all subsets $\{v, \bar{v}, e_i, t\}$ and $\{v, \bar{v}, f_i, t\}$ for each edge e_i incident to v . \mathcal{P}_v is ordered as follows in \succeq : $\{v, \bar{v}, \alpha, t\}$ is the first one, $\{v, \bar{v}, \beta, t\}$ is the last one, and all $\{v, \bar{v}, e_i, t\}$, $\{v, \bar{v}, f_i, t\}$ (e_i incident to v) are equivalent, between $\{v, \bar{v}, \alpha, t\}$ and $\{v, \bar{v}, \beta, t\}$. Note that, since all sets in \mathcal{P}_v contain v and \bar{v} , the relative positions of sets in \mathcal{P}_v and $\mathcal{P}_{v'}$ do not matter, and we do not specify them.

Up to now, we have: $d_{\alpha\beta} = |V|$ (1 for each v), $d_{\alpha e_i} = 2$ (1 for each extremity of e_i). We also have non zero values for some other pairs of individuals but this will not matter.

Now, we add dummy individuals to modify the (social) KS-scores of some individuals. Let M sufficiently large (say greater than $|V| + 1$). We add $\{\delta_i, \bar{\delta}_i, \mu\} \succ \{\delta_i, \bar{\delta}_i, t\}$ for $i = 1, \dots, M$. This makes the KS-score of t equal to $d_{\mu, t} = M$.

For each e_i we add $\{\gamma_j, \bar{\gamma}_j, \alpha\} \succ \{\gamma_j, \bar{\gamma}_j, e_i\}$ for $j = 1, \dots, M$. This makes $Score_{KS}(e_i) = d_{\alpha e_i} = M + 2$. We do similarly with individuals f_i , to get $Score_{KS}(f_i) = d_{\alpha f_i} = M + 2$.

Finally, we artificially make (with new individuals) $Score_{KS}(\alpha) = d_{\beta\alpha} = M - k$.

Note that t can modify the scores only for α, β and e_i, f_i . He cannot increase the score of β (which is $d_{\alpha\beta} = |V|$), so since M is larger, t is and will ever be worse than β . So, with respect to manipulation, only the scores of α, t, e_i and f_i matter. Initially, their respective KS-scores are $M - k, M, M + 2$ and $M + 2$.

We claim that there is an independent set of size k if and only if t can manipulate.

If there is an independent set S of size k , consider the manipulation where t puts $\{v, \alpha, t\}$ in last position in \mathcal{P}_v for each v in S . By doing this, $d_{\beta\alpha}$ increases by k , and $Score_{KS}(\alpha) = M = Score_{KS}(t)$. Since S is an independent set, $d_{\alpha e_i}$ decreases by at most 1 (idem for $d_{\alpha f_i}$), so $Score_{KS}(e_i) \geq M + 1$, $Score_{KS}(f_i) \geq M + 1$, and t is better of.

Conversely, in order to manipulate t must increase the KS-score of α by at least k . To achieve this, the unique possibility is to put $\{v, \bar{v}, \alpha, t\}$ in last position in at least k sets \mathcal{P}_v . Let S be the corresponding set of vertices (of size at least k). Suppose that S is not an independent set. Then for some edge e_i , $\{v, \bar{v}, \alpha, t\}$ is ranked last for both extremities. This means that for this edge $d_{\alpha e_i} = d_{\alpha f_i} = M$. Then $Score_{KS}(t) = M = Score_{KS}(e_i) = Score_{KS}(f_i)$, so t cannot be better of (he may have won 2 by defeating α , but is no tied with e_i and f_i , loosing at least 2). \square

6.5 Simulations

Inspired by previous works in voting theory [Chamberlin, 1985], we study to what extent the three rules are manipulable. In other terms, based on computer simulations of various power relations, we estimate the probability that a manipulation occurs and we analyse the vulnerability of each solution to manipulation. We only consider situations with a single manipulator and power relations on the whole set 2^N . In order to perform our simulations we need to find a manipulation strategy for each rule under the assumption that the manipulator has a complete knowledge about the power relation.

To find a manipulation, we set up an integer linear programming (ILP) formulation of the problem, the variables of which represent the ranking after manipulation. This ILP is efficient enough for small values of n , and we perform our simulation on total power relations over 2^N up to $n = 5$, i.e., power relations on up to $2^5 = 32$ coalitions. The data generation is done using Monte Carlo methods (see Algorithm 5), following uniform (*impartial culture*) model, which assumes that all power relations over coalitions are equally likely to occur. For each number of individuals n , we generated 1000 random total power relations.

Algorithm 5 Generate a random Power Relation PR over $\mathcal{P} \subseteq 2^N$

```

 $L = [S_1, \dots, S_{|P|}], PR = [], K = \emptyset, i = 1$ 
Shuffle L to get a random order over  $\mathcal{P}$ 
while  $K \neq \mathcal{P}$  do
    Randomly pick  $l \leq |L| - |K|$  the size of the indifference class
     $PR \leftarrow [PR, L[i : i + l - 1]]$ 
     $i \leftarrow i + l - 1$ 
     $K \leftarrow K \cup L[i : i + l - 1]$ 
end while
return PR

```

6.5.1 Integer linear Program

For all the three social ranking solutions, the inputs are the same:

Inputs: We are given a power relation \succeq , the potential manipulator i and the set of individuals N . From which we derive a matrix $M \in \{0, 1\}^{2^N \times 2^N}$ that is defined as follows:

- $M(S, S') = 1 \iff S \succeq S'$

The main variables of the program are also common to the three social ranking solutions:

Variables: A matrix $M' \in \{0, 1\}^{2^N \times 2^N}$ describing the manipulation which is defined as follows:

- $M'(S, S') = 1 \iff S \sqsupseteq S'$

From this matrix, we derive other variables:

- A vector of scores (Marginal contribution, Copeland (p_i) or Kramer-Simpson score): $s^{\sqsupset} \in \mathbb{Z}^N$
- A matrix of the social ranking $R \in \{0, 1\}^{N \times N}$ such that: $R(i, j) = 1 \iff iR^{\sqsupset}j \iff s_i^{\sqsupset} \geq s_j^{\sqsupset}$

The three solutions also share these common constraints:

Constraints: First we need to respect the definition of manipulation:

$$\begin{cases} M'(S, S') = M(S, S'), & \forall S, S' \notin S^i & \text{condition (i)} \\ M'(S', S) \geq M(S', S) \wedge M'(S, S') \leq M(S, S'), & \forall S \in S^i, S' \notin S^i & \text{condition (ii)} \end{cases}$$

To maintain the reflexivity, transitivity and completeness of the power relation:

$$\begin{cases} M'(S, S) = 1, & \forall S & \text{reflexiveness} \\ M'(S, S'') \geq M'(S, S') + M'(S', S'') - 1, & \forall S, S', S'' \in 2^N & \text{transitivity} \\ M'(S, S') + M'(S', S) \geq 1, & \forall S, S' \in 2^N & \text{total preorder} \end{cases}$$

Moreover, to deduce the matrix of the individual social ranking we impose the following constraints:

$$\begin{cases} R(j, k) \geq \epsilon(s_j^{\sqsupset} - s_k^{\sqsupset} + 1), & \forall j, k \in N \\ R(j, k) \leq 1 - \epsilon(s_k^{\sqsupset} - s_j^{\sqsupset}), & \forall j, k \in N \end{cases}$$

where ϵ is a small constant chosen in such a way that $|\epsilon(s_j^{\sqsupset} - s_k^{\sqsupset} + 1)| \in (0, 1)$.

Objective function Since the manipulator aims towards maximizing his social score, the objective function is:

$$\max \sum_{j \in N \setminus \{i\}} R(i, j) - R(j, i)$$

In addition to that, in order to compute the vector of scores s^{\sqsupset} each solution has its specific constraints and intermediate variables.

6.5.2 Ordinal Banzhaf solution

We impose the following constraints to compute the marginal contributions:

$$\forall k \in N, s_k^{\sqsupset} = \sum_{S \in \mathcal{S}_k} M'(S, S \setminus \{k\}) - M'(S \setminus \{k\}, S)$$

6.5.3 Copeland-like solution

First, we need to compute the matrix $D \in \mathbb{N}^{N \times N}$ such that $D_{ij} = |\{S \in 2^{N \setminus \{x,y\}}, S \cup \{x\} \succeq S \cup \{y\}\}|$ is the number of coalitions preferring i to j :

$$\forall j, k \in N, D_{jk} = \sum_{S \in 2^{N \setminus \{j,k\}}} M'(S \cup \{j\}, S \cup \{k\})$$

Next, we compute the matrix $CP \in \{0, 1\}^{N \times N}$ such that $CP_{ij} = 1 \iff iR_{CP}^{\sqsupseteq} j$, via constraints exactly similar to those that calculate the matrix R from the vector of scores s^{\sqsupseteq} . The Copeland score is, by definition, computed as follows:

$$\forall j \in N, s_j^{\sqsupseteq} = \sum_{k \in N} CP_{jk} - CP_{kj}$$

6.5.4 Kramer-Simpsonsolution: linearization of $x = \max(a, b)$

We consider the case of 3 individual $N = \{1, 2, 3\}$. The matrix D is computed exactly like for the Copeland-like solution. However, the score of $j \in N$ is $s_j^{\sqsupseteq} = \max_{k \neq j} D_{kj}$. For example: $s_1^{\sqsupseteq} = \max(D_{21}, D_{31})$. Hence, we have to linearize constraints like $x = \max(a, b)$.⁶

6.5.4.1 Linearize $x = \max(a, b)$

We introduce the variables:

- $(a - b)^+, (a - b)^- \in \mathbb{N}$, the positive and negative parts of $(a - b)$
- $y \in \{0, 1\}$ such that $y = 1 \iff (a - b) \geq 0$

and the constraints:

$$\left\{ \begin{array}{l} (a - b) = (a - b)^+ - (a - b)^- \\ y \geq \frac{a-b}{M}, |a - b| < M \\ y \leq 1 + \frac{a-b}{M} \\ (a - b)^+ \leq Ay, |a - b| < A \\ (a - b)^- \leq A(1 - y) \\ x - \frac{a+b}{2} = \frac{(a-b)^+ + (a-b)^-}{2} \end{array} \right.$$

Since $\max(a, b, c) = \max(\max(a, b), c)$ we can use this type of constraints to linearize the constraints for 4 or more individuals.

Remark. *The problem is always feasible since in the worst case for i , she can maintain the same \succsim .*

6.5.5 Results

6.5.5.1 The proportion of manipulable cases:

A power relation is manipulable if there exists at least one individual who can manipulate it. The results are shown in Figure 6.1. The probability of having a manipulable power relation increases

⁶For 4 or more agents we have to linearize constraints $x = \max(a, b, c)$

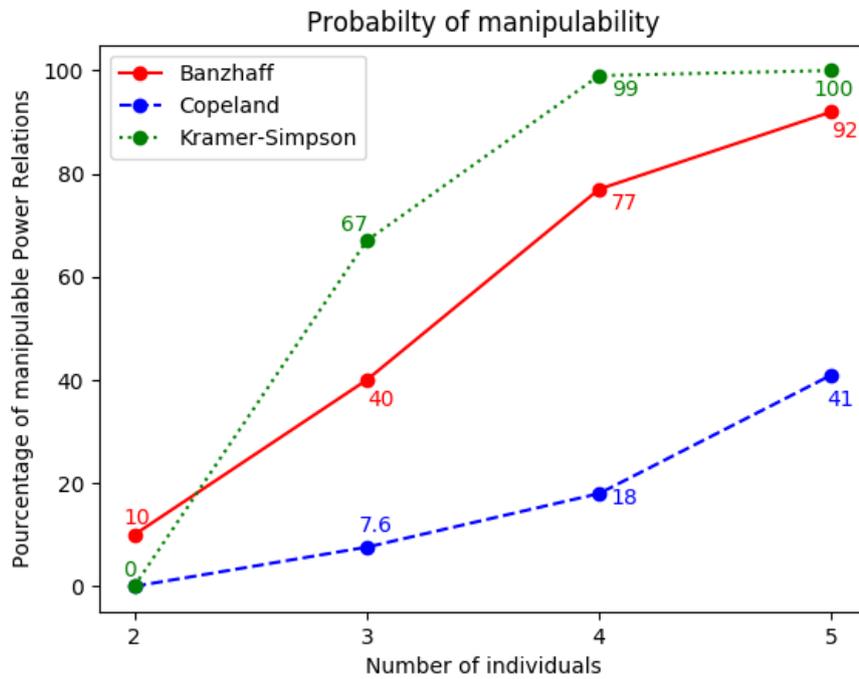


Figure 6.1: Proportion of manipulable cases

rapidly with the number of individuals for the three solutions, especially for the KS relation for which it reaches 99.6% for $n = 4$ and 100% for $n = 5$. However, for Copeland it reaches 41% and 92% for Banzhaf.

6.5.5.2 Number of possible manipulators:

We look at the number of possible manipulators for each **manipulable** case. The results are shown in Figure 6.2. The proportion of manipulators grows with the number of individuals. We note also that for Copeland solution, there are on average less possible manipulators for each power relation, and thus has a lower probability of being actually manipulated by one of them.

6.5.5.3 Manipulating to be the best ranked:

We analyse in the following the probability of becoming the best ranked one (ties are possible) thanks to a manipulation. See Figure 6.3.

6.5.5.4 Cross-simulation

We end our analysis by a cross simulation where for a given power relation we analyze the manipulability with respect to each social ranking rule. Table 6.1 shows our results:

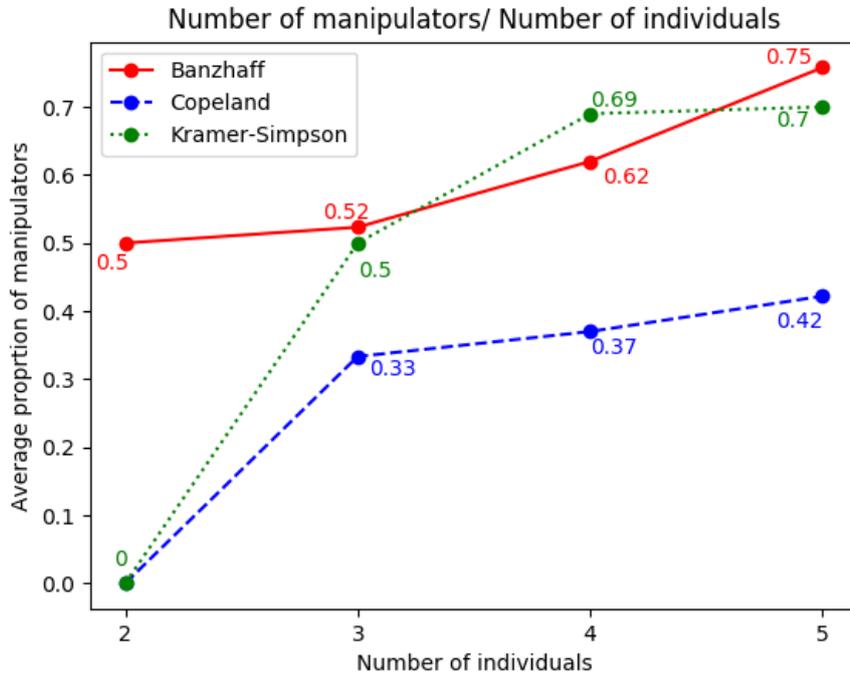


Figure 6.2: Proportion of possible manipulators.

Table 6.1 is coherent with our previous results (see Figure 6.1). Based on the results of simulations, it seems that if a power relation is manipulable by the Copeland-like solution, it is also by the KS-like solution. The most common case is to be manipulable by ordinal the Ordinal Banzhaf solution and the KS-like solution. For these reasons, our conjecture (suggested by the experimental results) is that the Copeland-like solution is not manipulable alone.

6.6 Conclusion

We have studied the problem of manipulating social ranking solutions. We have shown that lexcel is not manipulable and the manipulation of three other rules is NP-hard. Using simulation, we have remarked that Copeland-like is more resistant to manipulation than the Ordinal Banzhaf solution and the KS-like solution.

Our study opens the way for many future works. We quote some of them : An axiomatic characterization of SRC rules taking into account strategy-proofness (like the one of Gibbard and Satterthwaite ([Gibbard, 1973a] and [Satterthwaite, 1975b])), analysis of the impact on the manipulability of some domain restrictions, study of coalitional manipulation or of simultaneous manipulation (game theoretical issues).

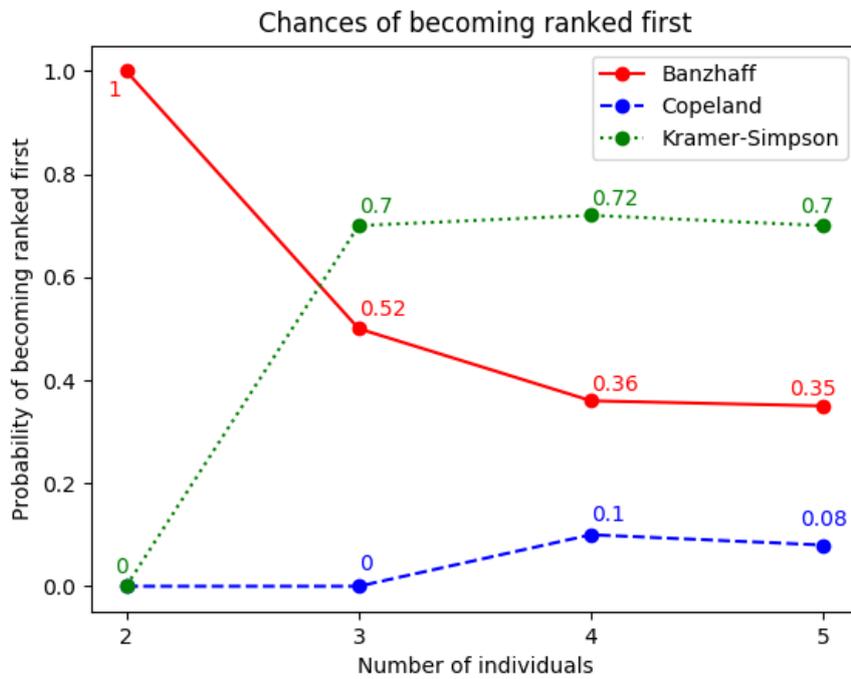


Figure 6.3: Probability to be ranked first with manipulation.

manipulable by	$N = 3$	$N = 4$	$N = 5$
Banz. +Cop. +KS	7.4	16.4	33.1
Banz. +Cop.	0	0	0
Banz. +KS	29.6	61.2	56.2
Banz.	1	0	0
Cop. +KS	0.2	0.4	4.7
Cop.	0	0	0
KS	28.7	22.0	6.0
none	33.1	0	0

Table 6.1: Percentage of manipulable power relations.

Bibliography

- [Alcalde-Unzu and Vorsatz, 2009] Alcalde-Unzu, J. and Vorsatz, M. (2009). Size approval voting. *J. Econ. Theory*, 144(3):1187–1210.
- [Allouche et al., 2020] Allouche, T., Escoffier, B., Moretti, S., and Öztürk, M. (2020). Social ranking manipulability for the cp-majority, banzhaf and lexicographic excellence solutions. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 17–23. ijcai.org.
- [Allouche et al., 2022a] Allouche, T., Lang, J., and Yger, F. (2022a). Multi-winner approval voting goes epistemic. In *UAI*.
- [Allouche et al., 2022b] Allouche, T., Lang, J., and Yger, F. (2022b). Truth-tracking via approval voting: Size matters. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 4768–4775. AAAI Press.
- [Arrow, 1950] Arrow, K. J. (1950). A Difficulty in the Concept of Social Welfare. *Journal of Political Economy*, 58:328–328.
- [Aziz et al., 2011] Aziz, H., Bachrach, Y., Elkind, E., and Paterson, M. (2011). False-name manipulations in weighted voting games. *J. Artif. Intell. Res.*, 40:57–93.
- [Bachrach and Rosenschein, 2009] Bachrach, Y. and Rosenschein, J. S. (2009). Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems*, 18(1):106.
- [Baharad et al., 2011] Baharad, E., Goldberger, J., Koppel, M., and Nitzan, S. (2011). Distilling the wisdom of crowds: weighted aggregation of decisions on multiple issues. *Autonomous Agents and Multi-Agent Systems*, 22(1):31–42.

- [Balinski and Laraki, 2007] Balinski, M. and Laraki, R. (2007). A theory of measuring, electing, and ranking. *Proceedings of the National Academy of Sciences*, 104(21):8720–8725.
- [Banzhaf III, 1964] Banzhaf III, J. F. (1964). Weighted voting doesn't work: A mathematical analysis. *Rutgers L. Rev.*, 19:317.
- [Ben-Yashar and Paroush, 2001] Ben-Yashar, R. and Paroush, J. (2001). Optimal decision rules for fixed-size committees in polychotomous choice situations. *Soc. Choice Welf.*, 18(4):737–746.
- [Ben-Yashar and Nitzan, 1997] Ben-Yashar, R. C. and Nitzan, S. I. (1997). The optimal decision rule for fixed-size committees in dichotomous choice situations: The general result. *International Economic Review*, 38(1):175–186.
- [Bernardi et al., 2019] Bernardi, G., Lucchetti, R., and Moretti, S. (2019). Ranking objects from a preference relation over their subsets. *Social Choice and Welfare*, 52(4):589–606.
- [Bilbao et al., 2002] Bilbao, J. M., Fernandez, J. R., Jiménez, N., and Lopez, J. J. (2002). Voting power in the european union enlargement. *European Journal of Operational Research*, 143(1):181–196.
- [Bonald and Combes, 2017] Bonald, T. and Combes, R. (2017). A minimax optimal algorithm for crowdsourcing. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4352–4360.
- [Brams and Kilgour, 2015] Brams, S. and Kilgour, D. (2015). *Satisfaction Approval Voting*, pages 275–298.
- [Brams and Fishburn, 2007] Brams, S. J. and Fishburn, P. C. (2007). *Approval voting, 2nd edition*. Springer.
- [Breiman, 1996] Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2):123–140.
- [Brill et al., 2017a] Brill, M., Freeman, R., Janson, S., and Lackner, M. (2017a). Phragmén's voting methods and justified representation. In Singh, S. and Markovitch, S., editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 406–413. AAAI Press.
- [Brill et al., 2017b] Brill, M., Freeman, R., Janson, S., and Lackner, M. (2017b). Phragmén's voting methods and justified representation. In Singh, S. P. and Markovitch, S., editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 406–413. AAAI Press.
- [Cappé and Moulines, 2009] Cappé, O. and Moulines, E. (2009). On-line expectation–maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):593–613.

BIBLIOGRAPHY

- [Caragiannis et al., 2020] Caragiannis, I., Kaklamanis, C., Karanikolas, N., and Krimpas, G. A. (2020). Evaluating approval-based multiwinner voting in terms of robustness to noise. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 74–80. ijcai.org.
- [Caragiannis and Micha, 2017] Caragiannis, I. and Micha, E. (2017). Learning a ground truth ranking using noisy approval votes. In Sierra, C., editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 149–155. ijcai.org.
- [Caragiannis et al., 2014] Caragiannis, I., Procaccia, A. D., and Shah, N. (2014). Modal ranking: A uniquely robust voting rule. In Brodley, C. E. and Stone, P., editors, *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 616–622. AAAI Press.
- [Caragiannis et al., 2016] Caragiannis, I., Procaccia, A. D., and Shah, N. (2016). When do noisy votes reveal the truth? *ACM Trans. Economics and Comput.*, 4(3):15:1–15:30.
- [Caragiannis and Schwartzbach, 2022] Caragiannis, I. and Schwartzbach, N. I. (2022). Adjudication with rational jurors. *CoRR*, abs/2201.06597.
- [Chamberlin, 1985] Chamberlin, J. R. (1985). An investigation into the relative manipulability of four voting systems. *Behavioral Science*, 30(4):195–203.
- [Chiaromonte et al., 2020] Chiaromonte, A., De Sio, L., and Emanuele, V. (2020). Salvini’s success and the collapse of the five-star movement: The european elections of 2019. *Contemporary Italian Politics*, pages 1–15.
- [Choi et al., 2010] Choi, S., Cha, S., and Tappert, C. C. (2010). A survey of binary similarity and distance measures. *Journal of Systemics, Cybernetics and Informatics*, 8(1):43–48.
- [Coelho, 2004] Coelho, D. (2004). *Understanding, Evaluating and Selecting Voting Rules Through Games and Axioms*. PhD thesis, Universitat Autònoma de Barcelona.
- [Condorcet, 1785] Condorcet (1785). *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie royale.
- [Conitzer et al., 2009] Conitzer, V., Rognlie, M., and Xia, L. (2009). Preference functions that score rankings and maximum likelihood estimation. In Boutilier, C., editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, pages 109–115.
- [Conitzer and Sandholm, 2005] Conitzer, V. and Sandholm, T. (2005). Common voting rules as maximum likelihood estimators. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence, UAI’05*, page 145–152, Arlington, Virginia, USA. AUAI Press.
- [Copeland, 1951] Copeland, A. (1951). A reasonable social welfare function. Technical report, mimeo, 1951. University of Michigan.

- [Dawid and Skene, 1979] Dawid, A. P. and Skene, A. M. (1979). Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society*, 28(1):20–28.
- [de Borda, 1781] de Borda, J.-C. (1781). Mémoire sur les élections au scrutin. *Histoire de l'Académie Royale des Sciences*.
- [Deng et al., 2014] Deng, J., Russakovsky, O., Krause, J., Bernstein, M. S., Berg, A. C., and Li, F. (2014). Scalable multi-label annotation. In *CHI Conference on Human Factors in Computing Systems*, pages 3099–3102. ACM.
- [Deng and Papadimitriou, 1994] Deng, X. and Papadimitriou, C. H. (1994). On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266.
- [Downs et al., 2010] Downs, J. S., Holbrook, M. B., Sheng, S., and Cranor, L. F. (2010). Are your participants gaming the system? screening mechanical turk workers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, page 2399–2402, New York, NY, USA. Association for Computing Machinery.
- [Elkind et al., 2017] Elkind, E., Faliszewski, P., Skowron, P., and Slinko, A. (2017). Properties of multiwinner voting rules. *Soc. Choice Welf.*, 48(3):599–632.
- [Elkind et al., 2015] Elkind, E., Lang, J., and Saffidine, A. (2015). Condorcet winning sets. *Soc. Choice Welf.*, 44(3):493–517.
- [Elkind and Shah, 2014] Elkind, E. and Shah, N. (2014). Electing the most probable without eliminating the irrational: Voting over intransitive domains. In Zhang, N. L. and Tian, J., editors, *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI 2014, Quebec City, Quebec, Canada, July 23-27, 2014*, pages 182–191. AUAI Press.
- [Elkind and Slinko, 2016] Elkind, E. and Slinko, A. (2016). Rationalizations of voting rules. In *Handbook of Computational Social Choice*, pages 169–196.
- [Elkind et al., 2016] Elkind, E., Slinko, A., and Moulin, H. (2016). *Rationalizations of Voting Rules*, page 169–196. Cambridge University Press.
- [Everaere et al., 2010] Everaere, P., Konieczny, S., and Marquis, P. (2010). The epistemic view of belief merging: Can we track the truth? In Coelho, H., Studer, R., and Wooldridge, M. J., editors, *ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 621–626. IOS Press.
- [Faliszewski and Hemaspaandra, 2009] Faliszewski, P. and Hemaspaandra, L. (2009). The complexity of power-index comparison. *Theoretical Computer Science*, 410(1):101–107.
- [Faliszewski et al., 2017] Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. (2017). Multiwinner voting: A new challenge for social choice theory. In *Trends in Computational Social Choice*, pages 27–47. AI Access.

- [Fayard and Öztürk, 2018] Fayard, N. and Öztürk, M. (2018). Ordinal Social ranking: simulation for CP-majority rule. In *Multiple Criteria Decision Aid to Preference Learning, DA2PL 2018*, Poznan, Poland.
- [Felsenthal et al., 1998] Felsenthal, D. S., Machover, M., et al. (1998). The measurement of voting power. *Books*.
- [Ferrara, 2004] Ferrara, F. (2004). Electoral coordination and the strategic desertion of strong parties in compensatory mixed systems with negative vote transfers. *Electoral studies*, 23(3):391–413.
- [Flores-Szwagrzak and Treibich, 2020] Flores-Szwagrzak, K. and Treibich, R. (2020). Teamwork and individual productivity. *Management Science*.
- [Gehrlein, 1985] Gehrlein, W. V. (1985). The condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199 – 209.
- [Gibbard, 1973a] Gibbard, A. (1973a). Manipulation of voting schemes. *Econometrica*, 41:587–601.
- [Gibbard, 1973b] Gibbard, A. (1973b). Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601.
- [Halpern et al., 2021] Halpern, D., Kehne, G., Peters, D., Procaccia, A. D., Shah, N., and Skowron, P. (2021). Aggregating binary judgments ranked by accuracy. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 5456–5463. AAAI Press.
- [Haret et al., 2018] Haret, A., Khani, H., Moretti, S., and Öztürk, M. (2018). Ceteris paribus majority for social ranking. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 303–309.
- [Holler and Packel, 1983] Holler, M. J. and Packel, E. W. (1983). Power, luck and the right index. *Zeitschrift für Nationalökonomie*, 43(1):21–29.
- [Hosseini et al., 2021] Hosseini, H., Mandal, D., Shah, N., and Shi, K. (2021). Surprisingly popular voting recovers rankings, surprisingly! In Zhou, Z., editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 245–251.
- [Jaccard, 1901] Jaccard, P. (1901). Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin del la Société Vaudoise des Sciences Naturelles*, 37:547–579.
- [Karge and Rudolph, 2022] Karge, J. and Rudolph, S. (2022). The more the worst-case-merrier: A generalized condorcet jury theorem for belief fusion. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*.

- [Khani et al., 2019] Khani, H., Moretti, S., and Öztürk, M. (2019). An ordinal banzhaf index for social ranking. In Kraus, S., editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 378–384. ijcai.org.
- [Kilgour, 2010] Kilgour, D. M. (2010). Approval balloting for multi-winner elections. In *Handbook on approval voting*, pages 105–124. Springer.
- [Kilgour, 2016] Kilgour, D. M. (2016). Approval elections with a variable number of winners. *Theory and Decision*, 81(2):199–211.
- [Kittur et al., 2008] Kittur, A., Chi, E. H., and Suh, B. (2008). Crowdsourcing user studies with mechanical turk. In Czerwinski, M., Lund, A. M., and Tan, D. S., editors, *Proceedings of the 2008 Conference on Human Factors in Computing Systems, CHI 2008, 2008, Florence, Italy, April 5-10, 2008*, pages 453–456. ACM.
- [Kramer, 1977] Kramer, G. H. (1977). A dynamical model of political equilibrium. *Journal of Economic Theory*, 16(2):310–334.
- [Kruger et al., 2014] Kruger, J., Endriss, U., Fernández, R., and Qing, C. (2014). Axiomatic analysis of aggregation methods for collective annotation. In Bazzan, A. L. C., Huhns, M. N., Lomuscio, A., and Scerri, P., editors, *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14, Paris, France, May 5-9, 2014*, pages 1185–1192. IFAA-MAS/ACM.
- [Lackner and Maly, 2021] Lackner, M. and Maly, J. (2021). Approval-based shortlisting. In Dignum, F., Lomuscio, A., Endriss, U., and Nowé, A., editors, *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 737–745. ACM.
- [Lackner and Skowron, 2020a] Lackner, M. and Skowron, P. (2020a). Approval-based committee voting: Axioms, algorithms, and applications. *CoRR*, abs/2007.01795.
- [Lackner and Skowron, 2020b] Lackner, M. and Skowron, P. (2020b). Approval-based committee voting: Axioms, algorithms, and applications. *CoRR*, abs/2007.01795.
- [Lackner and Skowron, 2021] Lackner, M. and Skowron, P. (2021). Consistent approval-based multi-winner rules. *Journal of Economic Theory*, 192:105173.
- [Ladha, 1995] Ladha, K. K. (1995). Information pooling through majority-rule voting: Condorcet’s jury theorem with correlated votes. *Journal of Economic Behavior & Organization*, 26(3):353–372.
- [Laruelle and Valenciano, 2001] Laruelle, A. and Valenciano, F. (2001). Shapley-shubik and banzhaf indices revisited. *Mathematics of operations research*, 26(1):89–104.
- [Laruelle and Valenciano, 2005] Laruelle, A. and Valenciano, F. (2005). A critical reappraisal of some voting power paradoxes. *Public Choice*, 125(1-2):17–41.

- [Le et al., 2010] Le, J., Edmonds, A., Hester, V., and Biewald, L. (2010). Ensuring quality in crowdsourced search relevance evaluation: The effects of training question distribution. In *SIGIR 2010 workshop on crowdsourcing for search evaluation*, volume 2126, pages 22–32.
- [List and Goodin, 2001] List, C. and Goodin, R. E. (2001). Epistemic democracy: Generalizing the condorcet jury theorem. *Journal of Political Philosophy*, 9(3):277–306.
- [Lundell and Hill, 2007] Lundell, J. and Hill, I. (2007). Notes on the droop quota. *Voting matters*, 24:3–6.
- [Mallows, 1957] Mallows, C. L. (1957). Non-null ranking models. *Biometrika*, 44(1-2):114–130.
- [Matsui and Matsui, 2000] Matsui, T. and Matsui, Y. (2000). A survey of algorithms for calculating power indices of weighted majority games. *Journal of the Operations Research Society of Japan*, 43(1):71–86.
- [Meir et al., 2019] Meir, R., Amir, O., Cohensius, G., Ben-Porat, O., and Xia, L. (2019). Truth discovery via proxy voting. *arXiv:1905.00629*, abs/1905.00629.
- [Moretti, 2015] Moretti, S. (2015). An axiomatic approach to social ranking under coalitional power relations. *Homo Oeconomicus*, 32(2):183–208.
- [Moretti and Öztürk, 2017] Moretti, S. and Öztürk, M. (2017). Some axiomatic and algorithmic perspectives on the social ranking problem. In *Algorithmic Decision Theory - 5th International Conference, ADT 2017, Luxembourg, Luxembourg, October 25-27, 2017, Proceedings*, pages 166–181.
- [Moretti and Patrone, 2008] Moretti, S. and Patrone, F. (2008). Transversality of the shapley value. *Top*, 16(1):1.
- [Nguyen et al., 2020] Nguyen, V., Hüllermeier, E., Rapp, M., Mencía, E. L., and Fürnkranz, J. (2020). On aggregation in ensembles of multilabel classifiers. In *Discovery Science*, volume 12323, pages 533–547.
- [Nitzan and Paroush, 1982] Nitzan, S. and Paroush, J. (1982). Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review*, 23(2):289–297.
- [Nowak and Rüger, 2010] Nowak, S. and Rüger, S. M. (2010). How reliable are annotations via crowdsourcing: a study about inter-annotator agreement for multi-label image annotation. In Wang, J. Z., Boujemaa, N., Ramirez, N. O., and Natsev, A., editors, *Proceedings of the 11th ACM SIGMM International Conference on Multimedia Information Retrieval, MIR 2010, Philadelphia, Pennsylvania, USA, March 29-31, 2010*, pages 557–566. ACM.
- [Owen et al., 1989] Owen, G., Grofman, B., and Feld, S. L. (1989). Proving a distribution-free generalization of the Condorcet Jury Theorem. *Mathematical Social Sciences*, 17(1):1–16.
- [Pigozzi and Hartmann, 2007] Pigozzi, G. and Hartmann, S. (2007). Aggregation in multi-agent systems and the problem of truth-tracking. In *Proceedings of the 6th International Joint*

- Conference on Autonomous Agents and Multiagent Systems (AAMAS 07), 14–18 May 2007, Honolulu, Hawaii, USA*, pages 674–676.
- [Pivato, 2017] Pivato, M. (2017). Epistemic democracy with correlated voters. *Journal of Mathematical Economics*, 72:51 – 69.
- [Prelec et al., 2017] Prelec, D., Seung, H. S., and McCoy, J. (2017). A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532–535.
- [Procaccia et al., 2012] Procaccia, A. D., Reddi, S. J., and Shah, N. (2012). A maximum likelihood approach for selecting sets of alternatives. In de Freitas, N. and Murphy, K. P., editors, *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, August 14-18, 2012*, pages 695–704. AUAI Press.
- [Procaccia and Shah, 2015] Procaccia, A. D. and Shah, N. (2015). Is approval voting optimal given approval votes? In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1801–1809.
- [Qing et al., 2014] Qing, C., Endriss, U., Fernández, R., and Kruger, J. (2014). Empirical analysis of aggregation methods for collective annotation. In *COLING*, pages 1533–1542. ACL.
- [Raykar et al., 2010] Raykar, V. C., Yu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., and Moy, L. (2010). Learning from crowds. *Journal of Machine Learning Research*, 11(4).
- [Satterthwaite, 1975a] Satterthwaite, M. A. (1975a). Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217.
- [Satterthwaite, 1975b] Satterthwaite, M. A. (1975b). Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217.
- [Shah et al., 2015] Shah, N. B., Zhou, D., and Peres, Y. (2015). Approval voting and incentives in crowdsourcing. In Bach, F. R. and Blei, D. M., editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 10–19. JMLR.org.
- [Shapley and Grofman, 1984] Shapley, L. and Grofman, B. (1984). Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice*, 43(3):329–343.
- [Shapley, 1953] Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317.
- [Simpson, 1969] Simpson, P. B. (1969). On Defining Areas of Voter Choice: Professor Tullock on Stable Voting. *The Quarterly Journal of Economics*, 83(3):478–490.

- [Tao et al., 2018] Tao, D., Cheng, J., Yu, Z., Yue, K., and Wang, L. (2018). Domain-weighted majority voting for crowdsourcing. *IEEE transactions on neural networks and learning systems*, 30(1):163–174.
- [Terzopoulou and Endriss, 2019] Terzopoulou, Z. and Endriss, U. (2019). Optimal truth-tracking rules for the aggregation of incomplete judgments. In *SAGT*, volume 11801, pages 298–311. Springer.
- [Thiele, 1895] Thiele, T. (1895). *Om flerfoldsvalg*, volume 1814-1836. København :A.F. Høst.
- [Tversky, 1977] Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4):327–352.
- [Welinder et al., 2010] Welinder, P., Branson, S., Belongie, S. J., and Perona, P. (2010). The multidimensional wisdom of crowds. In *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada*, pages 2424–2432. Curran Associates, Inc.
- [Wikipedia contributors, 2022] Wikipedia contributors (2022). Single transferable vote — Wikipedia, the free encyclopedia. https://en.wikipedia.org/w/index.php?title=Single_transferable_vote&oldid=1116556843. [Online; accessed 17-October-2022].
- [Xia and Conitzer, 2011] Xia, L. and Conitzer, V. (2011). A maximum likelihood approach towards aggregating partial orders. In Walsh, T., editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 446–451. IJCAI/AAAI.
- [Xia et al., 2010] Xia, L., Conitzer, V., and Lang, J. (2010). Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In van der Hoek, W., Kaminka, G. A., Lespérance, Y., Luck, M., and Sen, S., editors, *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 10-14, 2010, Volume 1-3*, pages 399–408. IFAAMAS.
- [Young, 1986] Young, H. P. (1986). *Information pooling and group decision making*. Jai.
- [Young, 1988] Young, H. P. (1988). Condorcet’s theory of voting. *American Political science review*, 82(4):1231–1244.
- [Young, 1995] Young, P. (1995). Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64.
- [Zhu and Carterette, 2010] Zhu, D. and Carterette, B. (2010). An analysis of assessor behavior in crowdsourced preference judgments. In *SIGIR 2010 workshop on crowdsourcing for search evaluation*, pages 17–20.
- [Zou et al., 2015] Zou, J. Y., Meir, R., and Parkes, D. C. (2015). Strategic voting behavior in doodle polls. In Cosley, D., Forte, A., Ciolfi, L., and McDonald, D., editors, *Proceedings of*

the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW 2015, Vancouver, BC, Canada, March 14 - 18, 2015, pages 464–472. ACM.

[Zuckerman et al., 2012] Zuckerman, M., Faliszewski, P., Bachrach, Y., and Elkind, E. (2012). Manipulating the quota in weighted voting games. *Artif. Intell.*, 180-181:1–19.

RÉSUMÉ

Dans le vote par approbation épistémique, les agents sont confrontés à une question comportant plusieurs alternatives et une vérité objective cachée. Les votants sélectionnent les alternatives qui, suivant leurs croyances, peuvent correspondre à la vérité. Les croyances des votants sont agrégées pour estimer cette vérité. D'abord, nous considérons le cas où une seule alternative est correcte. Nous recommandons des règles d'agrégation qui attribuent plus de poids aux votants ayant sélectionné peu d'alternatives, car ils ont tendance à être plus précis. Il en découle de nouvelles méthodes fondées sur des résultats théoriques et validées par des expériences numériques. Ensuite, nous considérons le cas où la vérité est constituée de plusieurs alternatives (les notes d'un accord, les meilleurs candidats..). La taille de la sortie dans de telles situations peut être connue au préalable ou limitée par une contrainte exogène. Nous proposons des solutions adaptées à chacun de ces deux cas.

MOTS CLÉS

Vote d'approbation épistémique, Estimation du maximum de vraisemblance, Crowdsourcing, Labellisation des données

ABSTRACT

In epistemic approval voting, there is a hidden ground truth, and voters select the alternatives which, according to their beliefs, can correspond to the ground truth. These votes are then aggregated to estimate it. We first focus on tracking a *simple* truth, where exactly one alternative is correct. We advocate using aggregation rules that assign more weight to voters who select fewer alternatives, as they tend to be more accurate. This yields novel methods backed by theoretical results and experiments on image annotation datasets. Second, we consider cases where the ground truth contains *multiple* alternatives (*e.g.*, notes in a chord, objectively best applicants). The size of the output can be either a prior knowledge on the number of true alternatives, or an exogenous constraint bearing on the output of the rule regardless of the true size of the ground truth. We propose suitable solution concepts for each of these two interpretations.

KEYWORDS

Epistemic approval voting, Maximum likelihood estimation, Crowdsourcing, Data Labeling.