



## UvA-DARE (Digital Academic Repository)

### Collective Information

Endriss, U.

**DOI**

[10.1609/aaai.v34i09.7074](https://doi.org/10.1609/aaai.v34i09.7074)

**Publication date**

2020

**Document Version**

Final published version

**Published in**

AAAI-20, IAAI-20, EAAI-20 proceedings

**License**

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/in-the-netherlands/you-share-we-take-care>)

[Link to publication](#)

**Citation for published version (APA):**

Endriss, U. (2020). Collective Information. In *AAAI-20, IAAI-20, EAAI-20 proceedings: Thirty-Fourth AAAI Conference on Artificial Intelligence, Thirty-Second Conference on Innovative Applications of Artificial Intelligence, The Tenth Symposium on Educational Advances in Artificial Intelligence : February 7–12th, 2020, New York Hilton Midtown, New York, New York, USA* (Vol. 9, pp. 13520-13524). (Proceedings of the AAAI Conference on Artificial Intelligence; Vol. 34). AAAI Press. <https://doi.org/10.1609/aaai.v34i09.7074>

**General rights**

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

**Disclaimer/Complaints regulations**

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

# Collective Information

Ulle Endriss

Institute for Logic, Language and Computation  
University of Amsterdam  
The Netherlands

## Abstract

Many challenging problems of scientific, technological, and societal significance require us to aggregate information supplied by multiple agents into a single piece of information of the same type—the *collective information* representing the stance of the group as a whole. Examples include expressive forms of voting and democratic decision making (where citizens supply information regarding their preferences), peer evaluation (where participants supply information in the form of assessments of their peers), and crowdsourcing (where volunteers supply information by annotating data). In this position paper, I outline the challenge of modelling, handling, and analysing all of these diverse instances of collective information using a common methodology. Addressing this challenge will facilitate a transfer of knowledge between different application domains, thereby enabling progress in all of them.

## Introduction

When attempting to summarise the views of a committee of experts, when trying to recommend a restaurant based on online reviews, or when hoping to compute a meaningful grade for a student in a massively open online course (MOOC) that uses peer assessment, we always face the same daunting task: to aggregate multiple pieces of information of a specific type, each contributed by a different agent, into a single piece of information of the same type that accurately reflects the position taken by the group as a whole. Indeed, a wide range of challenging practical problems, originating in the most diverse corners of science, technology, and society, all fit this general pattern. Other examples include participatory budgeting (Cabannes 2004), collective argumentation (Bodanza, Tohmé, and Auday 2017), and even ontology merging (Flouris et al. 2008). While all of these application scenarios have been (and continue to be) addressed by dedicated research communities, to date there has been only relatively little effort directed towards understanding the *general pattern* they all have in common. In this position paper, I argue that closing this gap has the potential to facilitate a much-needed transfer of knowledge between research areas, thereby enabling progress on all of these applications.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

So what is this general pattern? We always start out with a set of *agents*, such as the students taking part in a MOOC. In fact, these agents need not be people. Rather, we may think of them simply as distinct *sources of information*. These agents are *autonomous* entities, so we have an obligation to treat them *fairly*. Every agent provides us with a *piece of information*—expressed in some *formal language*—and we need to *aggregate* these individual pieces of information into a single piece of *collective information*, expressed in the same language. The formal language in question is determined by the domain of application. It could be a language to encode expert judgments, restaurant appraisals, student evaluations, budget allocations, or ontologies.

This abstract perspective raises a number of questions: What are the best methods to implement such a process of aggregation? What *general principles* should guide our choice of method? But also: How should the *particular features* of the type of information at hand (i.e., of the formal language used) affect our choice of method? Is it possible to formulate general principles that are *parametric* in the type of information at hand? Answering the latter question, in particular, would allow us to generalise beyond specific solutions for specific applications and transfer insights about designing good aggregation methods across applications.

A specific domain where principles of aggregation have been investigated in depth is that of preference aggregation, which is the main object of study in *social choice theory* (Arrow, Sen, and Suzumura 2002). The methodology of classical social choice theory encompasses conceptual analysis informed by Philosophy, Political Science, and Economics (to determine what makes for a good method of aggregation), a variety of mathematical tools (to formalise desiderata and explore their logical consequences), and empirical studies (to better understand how different methods fare in practice). Modern *computational social choice*, which has advanced the idea of thinking of methods of aggregation *as algorithms*, extends this toolbox with a variety of techniques from Computer Science and AI (Brandt et al. 2016). This makes computational social choice a natural starting point for investigating the phenomenon of collective information.

The remainder of this position paper is organised as follows. I first discuss the phenomenon of collective informa-

tion and how it may be formalised in a little more detail (*what?*). I then review a number of application domains that would directly benefit from a deeper understanding of this phenomenon (*why?*) as well as a number of disciplines that are likely to be relevant in developing such an understanding (*how?*). Finally, I outline what I consider to be the main challenge in this area: to understand how specific features of the type of information to be aggregated should impact on our choice of aggregation mechanism (*whither?*).

### What? — Collective Information

Collective information is what we obtain when we aggregate several individual pieces of information. Let  $\mathcal{L}$  be a (finite) *formal language* for describing pieces of information relevant to a specific application. For example, in the context of the work of a hiring committee that has to rank three job candidates—Alice, Bob, and Carol—the language  $\mathcal{L}$  might be the set of the  $3! = 6$  possible rankings of the candidates:

$$\mathcal{L} = \left\{ \begin{array}{|c|} \hline \text{Alice} \\ \hline \text{Bob} \\ \hline \text{Carol} \\ \hline \end{array}, \begin{array}{|c|} \hline \text{Alice} \\ \hline \text{Carol} \\ \hline \text{Bob} \\ \hline \end{array}, \begin{array}{|c|} \hline \text{Bob} \\ \hline \text{Alice} \\ \hline \text{Carol} \\ \hline \end{array}, \begin{array}{|c|} \hline \text{Bob} \\ \hline \text{Carol} \\ \hline \text{Alice} \\ \hline \end{array}, \begin{array}{|c|} \hline \text{Carol} \\ \hline \text{Alice} \\ \hline \text{Bob} \\ \hline \end{array}, \begin{array}{|c|} \hline \text{Carol} \\ \hline \text{Bob} \\ \hline \text{Alice} \\ \hline \end{array} \right\}$$

For other applications, we may need a very different language. For example, suppose we want to annotate a large corpus of linguistic data with semantic information that can be used by researchers in natural language processing (NLP) as a gold standard when testing machine learning techniques to recognise such semantic features automatically. We could try to do this by crowdsourcing many (likely low-quality) annotations using a tool such as AMAZON’S MECHANICAL TURK and then aggregate this information into a (hopefully) high-quality gold standard annotation (Snow et al. 2008). In this context,  $\mathcal{L}$  might be the set of all conceivable (partial) annotations. Imagine, for instance, we want to annotate the 800 sentence-pairs collected by Dagan, Glickman, and Magnini (2006) for the purpose of training NLP systems to automatically recognise *textual entailment*. Then a partial annotation might be modelled by the list of the 800 identifiers used in this dataset, with each of them being labelled by  $\checkmark$  (if the first sentence of this pair entails the second),  $\times$  (if this is not the case), or a question mark (if this pair has not been annotated at all). Thus,  $\mathcal{L}$  now has  $3^{800}$  elements:

$$\mathcal{L} = \left\{ \begin{array}{|c|} \hline 001 : \checkmark \\ \hline \vdots \\ \hline 799 : \checkmark \\ \hline 800 : \checkmark \\ \hline \end{array}, \begin{array}{|c|} \hline 001 : \checkmark \\ \hline \vdots \\ \hline 799 : \checkmark \\ \hline 800 : \times \\ \hline \end{array}, \begin{array}{|c|} \hline 001 : \checkmark \\ \hline \vdots \\ \hline 799 : \checkmark \\ \hline 800 : ? \\ \hline \end{array}, \begin{array}{|c|} \hline 001 : \checkmark \\ \hline \vdots \\ \hline 799 : \times \\ \hline 800 : \checkmark \\ \hline \end{array}, \dots \right\}$$

Every agent reports an expression belonging to  $\mathcal{L}$ . The elements of  $\mathcal{L}$  are also the pieces of information an aggregation rule might return. Thus, an *aggregation rule* for language  $\mathcal{L}$  and  $n$  agents (with  $n \in \mathbb{N}$ ) is a function  $F$  mapping each profile of individual pieces of information, one for each agent, to a single collective piece of information:

$$F : \mathcal{L}^n \rightarrow \mathcal{L}$$

In fact, often a small refinement of this very simple model will be useful. For the input to  $F$  we may want to restrict attention to elements of  $\mathcal{L}$  that meet certain *input constraints*

and for the output of  $F$  we may want to impose certain *output constraints* (Endriss 2018).

For example, in the context of voting in our hiring committee, we may think of  $\mathcal{L}$  as the set of all *weak orders* on  $\{\text{Alice}, \text{Bob}, \text{Carol}\}$ , thereby allowing for ties in a ranking. We could then impose *antisymmetry* as an input constraint (forcing each agent to provide a strict ranking) and *dichotomousness* as an output constraint (forcing at most two levels, thereby distinguishing only candidates accepted and candidates rejected for the job). As a second example, consider again the case of crowdsourcing annotations of a linguistic corpus. Maybe we want to allow each individual worker to annotate any number of sentence-pairs (so no input constraint is required), but in the output we might require all pairs to be annotated (so the output constraint should exclude annotations with question marks).

### Why? — Application Scenarios

Being able to model, handle, and reason about collective information is important, because collective information is a core component of a wide range of application domains of scientific, technological, and societal relevance. Let us review some of them here.

The first three examples concern applications with a clear societal component. They belong, broadly speaking, to the domain of politics and the design of democratic institutions. They particularly relate to collective decision making at the local level, thereby promoting citizen involvement.

- **Expressive voting.** Running an election amounts to aggregating information on the preferences of individual voters into collective information regarding the preferences of the electorate as a whole. Specifically, if  $\mathcal{L}$  is the set of weak orders on a set of candidates, then  $F$  is a so-called *social welfare function*. By adding suitable input and output constraints, we can model voting rules for returning single election winners or sets of winners of a given size. However, many real-world settings call for more expressive models that allow, for instance, for the possibility to not rank all candidates (Terzopoulou and Endriss 2019) or to delegate your vote (Brill 2018).
- **Community-driven policy design.** Several innovative ideas fit under this heading. One that already is widely used in practice is *participatory budgeting* (Cabannes 2004), which is about allowing citizens to directly participate in decisions about the allocation of budgets to different public projects. A few selected contributions already employ the terminology of social choice theory to discuss suitable mechanisms (Goel et al. 2016; Benadè et al. 2017; de Haan 2018). Another example is the recent proposal by Conitzer, Brill, and Freeman (2015) to aggregate the opinions of individual citizens to determine socially acceptable numerical *tradeoffs* between different kinds of bads (say, burning  $x$  litres of petrol *vs.* clearing  $y$  hectares of forest) to inform governmental policy.
- **Collective argumentation.** Collective decision making does not occur in a vacuum. Before taking a decision, people want to research relevant facts, exchange points

of view, and debate. This can be facilitated using an online debating platform such as KIALO ([www.kialo.com](http://www.kialo.com)). Some of the research on *computational models of argument* in the AI community can provide adequate formal underpinnings for such platforms (Bodanza, Tohmé, and Auday 2017). Indeed, there already are several proposals for tackling specific problems in the domain of collective argumentation using methods developed in social choice theory (see, e.g., Chen and Endriss 2019).

The next three examples concern different kinds of technological uses of collective information.

- **Crowd recommendation.** Online recommendation services, such as TRIPADVISOR ([tripadvisor.com](http://tripadvisor.com)), aggregate individual information provided in the form of reviews, e.g., of hotels or restaurants, into collective evaluations. Here,  $\mathcal{L}$  might be the set of all possible (incomplete) star-ratings of the restaurants in a given city. Research on this topic often focuses on probabilistic models and the use of maximum-likelihood methods to estimate an underlying ground truth, as well as on interpreting the textual feedback given by users (Ghose, Ipeirotis, and Li 2012). Issues typically studied in social choice theory, such as fairness or strategic behaviour, have received less attention to date (Pennock, Horvitz, and Giles 2000).
- **Rank aggregation for information retrieval.** Reputation systems used to determine the relative importance of different websites can be modelled as aggregation problems for information inherent in the link structure between such sites (Tennenholtz and Zohar 2016). The most famous example is the PAGERANK algorithm used by the GOOGLE search engine. Integrating the search results returned by several different Internet search engines is yet another example for an aggregation problem (Dwork et al. 2001).
- **Peer evaluation.** The advent of MOOCs raises the significant challenge of how to grade the work of vast numbers of students. Various peer grading models have been proposed, under which we have to grade a student based on the grading information each student provides on (a small subset of) the work of her peers (Caragiannis 2017). Selecting project proposals by scientists for funding or papers for publication in the proceedings of a conference, based on the reviews produced by members of the same population of scientists, is a closely related challenge. Certainly when selecting papers for large conferences, we increasingly rely on the rankings suggested by the software tools used to manage the reviews received. Thus, this is an important domain where a more principled use of technology may result in significant improvements.

The remaining examples are about helping scientists who use tools that involve various forms of collective information. This includes Cognitive Science and the Social Sciences more generally, as well as AI and Computer Science.

- **Crowdsourced annotation of data.** Scientists from a broad range of disciplines require large corpora of annotated data for their work. Such corpora can serve as benchmarks when testing new theories. They also are required

for research that makes use of supervised machine learning (as training data). For example, linguists use crowdsourcing tools such as AMAZON'S MECHANICAL TURK to collect annotations of linguistic corpora (Snow et al. 2008). This information, consisting of individual annotations obtained from multiple workers, then needs to be turned into a *collective annotation* that accurately reflects the “wisdom of the crowd” (Endriss and Fernández 2013).

- **Consensus clustering.** Scientists from various disciplines need to cluster data collected in experiments, to better understand its inherent structure. Numerous algorithms for computing such a clustering have been proposed. When trying to combine the results returned by alternative algorithms, we speak of *consensus clustering* (Goder and Filkov 2008). This is yet another form of collective information, with  $\mathcal{L}$  representing the range of possible clusterings, e.g., in the form of equivalence relations.
- **Ontology merging.** Ontologies are formal descriptions of a narrowly defined domain of knowledge (e.g., medical knowledge) in a logical language. Such ontologies are useful for the development of expert systems and intelligent question answering systems. When trying to combine information from several ontologies, we once again perform aggregation. Ontology merging is a very demanding task and current research usually focuses on integrating information from just two (rather than  $n$ ) ontologies (Noy 2004), but in the future, where we can imagine larger numbers of (possibly small) ontologies being discovered dynamically on the Semantic Web, the agent-centric perspective advocated here is clearly attractive and needs to be developed further (Porello et al. 2018).

### How? — Contributing Disciplines

Research on all of these challenges will have to integrate domain-specific and domain-independent approaches. Regarding the former, for instance, making progress on the challenge of ontology merging requires deep insights into description logics and other Semantic Web technologies. Similarly, advancing the state of the art in expressive voting requires insights from Political Science to better understand what kind of technological innovations would in fact benefit users in practice. But at the same time there also is much that these diverse applications have in common, and a full understanding of the nature of aggregation will require us to isolate this common pattern from domain-specific aspects.

As mentioned in the introduction already, computational social choice (Brandt et al. 2016) appears to be a promising starting point for developing a methodology for analysing the domain-independent aspects of collective information. Computational social choice is itself an interdisciplinary undertaking, bringing together ideas and techniques from Philosophy, Political Science, Economics, Mathematics, Logic, Computer Science, and AI. To obtain a better understanding of collective information and to develop theoretical and practical tools for working with this concept, I expect that at least the following approaches will play a prominent role:

- **Axiomatic method.** When employing the axiomatic

method (Thomson 2001), as understood in economic theory at large and social choice theory in particular, we formulate normative desiderata of mechanisms in mathematically rigorous terms to then explore the logical consequences of these definitions. The desiderata themselves may be informed by a range of disciplines, including not only Economics, but also Political Science, Philosophy, and particularly Ethics. They often formalise certain intuitive conceptions of the notion of fairness.

- **Game theory.** Game theory is the study of mathematical models to analyse strategic interactions between rational agents (Leyton-Brown and Shoham 2008). Whenever we think of a source of information as an agent that has a certain degree of autonomy, we should account for the possibility that this agent may act in a strategic manner to influence the outcome of an aggregation process. Game theory is the framework of choice to analyse such effects and, where possible and desirable, to contain them.
- **Probability theory and statistics.** The predominant methodology in several of the application areas discussed earlier is grounded in probability theory and statistics. The basic idea is that the pieces of information supplied by individual agents can be understood as representing noisy copies of some underlying *ground truth*. Given an appropriate probabilistic model, tailored to the application at hand, we can then try to recover this ground truth by means of maximum-likelihood estimation (Raykar et al. 2010; Elkind and Slinko 2016; Xia 2019).
- **Logic and automated reasoning.** Logic has the potential to play a significant role in the study of collective information in more than one way. For one, logic is a natural framework in which to encode richly structured information. Indeed, arguably the most general form of aggregation that has been studied systematically in recent years is logic-based *judgment aggregation* (Grossi and Pigozzi 2014). Logic also sits at the core of an emerging research trend within computational social choice that is concerned with the use of automated reasoning tools—notably SAT solvers—to derive formal results regarding social choice mechanisms automatically (Geist and Peters 2017).
- **Algorithms and complexity theory.** Every mechanism for turning individual pieces into collective information ultimately is an algorithm. Thus, any such mechanism is subject to the kind of analysis of algorithms routinely carried out in Computer Science. This includes, in particular the use of complexity theory, to understand the consumption of resources required by a mechanism (Betzler et al. 2012), but also the analysis of its communication requirements (Boutilier and Rosenschein 2016).

### Whither? — Research Challenges

I have sketched a definition for a notion of *collective information* and argued that several important practical problems currently studied in very different fields may be classified as instances of this notion. I also have pointed out that many of the techniques currently employed to study preference aggregation in the field of computational social choice are

natural starting points for a systematic investigation of collective information. There are many interesting and likely fruitful avenues of research to pursue here, focusing either on specific approaches or specific application scenarios. But the research directions I believe to have the most potential for impact are those aimed at building new bridges *between* different application scenarios by exploring the general pattern of collective information they have in common.

The ultimate grand challenge here, it appears, is to try and understand how the specific features of the type of information at play in a given application scenario do and should influence the design of good aggregation mechanisms. For example, what changes when we move from the *structurally simple* kind of information used to represent preferences in an election or corpus annotations in a crowdsourcing experiment to the *structurally rich* kind of information required to model an ontology? Of course, very many things change, but what is it that specifically changes about the desiderata we have for our aggregation mechanism? Similarly, what changes when we move from a scenario in which every agent provides *complete* information (e.g., where everyone ranks the same candidates in an election or answers the same questions in a questionnaire) to a scenario with very *incomplete* information, such as the crowd recommendation or peer evaluation scenarios, where every agent provides information on only a tiny fraction of all of the issues at stake? I believe that developing this kind of understanding will enable us to transfer insights—and ultimately also techniques and solutions—from one application to the next, thereby advancing the state of the art in all of them. Doing so is challenging, but not impossible. For example, for *graph aggregation* there already are results of this kind, illustrating how the logical richness of the space of inputs individual agents can provide affects the possibility of designing appealing aggregation mechanisms (Endriss and Grandi 2017).

On the path towards tackling this grand challenge, there are a multitude of important and interesting research questions concerning somewhat more concrete instances of knowledge transfer between pairs of application scenarios. The basic idea is to take one technique that has been successfully applied in one scenario and explore what happens when we try to port it to another scenario. There are several examples in the literature where this approach did bear fruit. To conclude, let me just mention a handful of them: Altman and Tennenholtz (2005) have used the axiomatic method, as classically used for the analysis of voting rules, to obtain an axiomatic characterisation of PAGERANK. Kleinberg (2002) has applied the same methodology to cluster analysis. Mao, Procaccia, and Chen (2013) have experimented with voting rules originally designed for political elections to design better crowdsourcing mechanisms. Caragiannis, Krimpas, and Voudouris (2015) have used methods originally developed to understand the informational requirements of voting rules to design better peer grading mechanisms. Finally, de Haan (2018) has used insights gained through a complexity-theoretic analysis of judgment aggregation to shed new light on participatory budgeting.

I hope that these examples will inspire others to contribute to our understanding of collective information as well.

## References

- Altman, A., and Tennenholtz, M. 2005. Ranking systems: The PageRank axioms. In *Proc. 6th ACM Conference on Electronic Commerce (EC-2005)*. ACM.
- Arrow, K. J.; Sen, A. K.; and Suzumura, K., eds. 2002. *Handbook of Social Choice and Welfare*, volume 1. North-Holland.
- Benadè, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Preference elicitation for participatory budgeting. In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*.
- Betzler, N.; Brederick, R.; Chen, J.; and Niedermeier, R. 2012. Studies in computational aspects of voting. In *The Multivariate Algorithmic Revolution and Beyond*. Springer. 318–363.
- Bodanza, G. M.; Tohmé, F. A.; and Auday, M. R. 2017. Collective argumentation: A survey of aggregation issues around argumentation frameworks. *Argument & Computation* 8(1):1–34.
- Boutilier, C., and Rosenschein, J. 2016. Incomplete information and communication in voting. In *Handbook of Computational Social Choice*. Cambridge University Press. 223–257.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Brill, M. 2018. Interactive democracy. In *Proc. 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*. Blue Sky Ideas Track.
- Cabannes, Y. 2004. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization* 16(1):27–46.
- Caragiannis, I.; Krimpas, G. A.; and Voudouris, A. A. 2015. Aggregating partial rankings with applications to peer grading in massive online open courses. In *Proc. 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*. IFAA-MAS.
- Caragiannis, I. 2017. Recent advances in large-scale peer grading. In *Trends in Computational Social Choice*. AI Access.
- Chen, W., and Endriss, U. 2019. Preservation of semantic properties in collective argumentation: The case of aggregating abstract argumentation frameworks. *Artificial Intelligence* 269:27–48.
- Conitzer, V.; Brill, M.; and Freeman, R. 2015. Crowdsourcing societal tradeoffs. In *Proc. 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*. Blue Sky Ideas Track.
- Dagan, I.; Glickman, O.; and Magnini, B. 2006. The PASCAL recognising textual entailment challenge. In *Machine Learning Challenges: Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*. Springer.
- Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proc. 10th International Conference on World Wide Web (WWW-2001)*.
- Elkind, E., and Slinko, A. 2016. Rationalizations of voting rules. In *Handbook of Computational Social Choice*. Cambridge University Press. 169–196.
- Endriss, U., and Fernández, R. 2013. Collective annotation of linguistic resources: Basic principles and a formal model. In *Proc. 51st Annual Meeting of the Association for Computational Linguistics (ACL-2013)*.
- Endriss, U., and Grandi, U. 2017. Graph aggregation. *Artificial Intelligence* 245:86–114.
- Endriss, U. 2018. Judgment aggregation with rationality and feasibility constraints. In *Proc. 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*.
- Flouris, G.; Manakanatas, D.; Kondylakis, H.; Plexousakis, D.; and Antoniou, G. 2008. Ontology change: Classification and survey. *Knowledge Engineering Review* 23(2):117–152.
- Geist, C., and Peters, D. 2017. Computer-aided methods for social choice theory. In *Trends in Computational Social Choice*. AI Access.
- Ghose, A.; Ipeirotis, P. G.; and Li, B. 2012. Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science* 31(3):493–520.
- Goder, A., and Filkov, V. 2008. Consensus clustering algorithms: Comparison and refinement. In *Proc. 10th Workshop on Algorithm Engineering and Experiments (ALENEX-2008)*. SIAM.
- Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2016. Knapsack voting: Voting mechanisms for participatory budgeting. Working Paper, Stanford University.
- Grossi, D., and Pigozzi, G. 2014. *Judgment Aggregation: A Primer*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- de Haan, R. 2018. Hunting for tractable languages for judgment aggregation. In *Proc. 16th International Conference on Principles of Knowledge Representation and Reasoning (KR-2018)*.
- Kleinberg, J. M. 2002. An impossibility theorem for clustering. In *Proc. 15th Annual Conference on Neural Information Processing Systems (NIPS-2002)*.
- Leyton-Brown, K., and Shoham, Y. 2008. *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Mao, A.; Procaccia, A. D.; and Chen, Y. 2013. Better human computation through principled voting. In *Proc. 27th AAAI Conference on Artificial Intelligence (AAAI-2013)*.
- Noy, N. F. 2004. Tools for mapping and merging ontologies. In *Handbook on Ontologies*. Springer-Verlag.
- Pennock, D. M.; Horvitz, E.; and Giles, C. L. 2000. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proc. 17th National Conference on Artificial Intelligence (AAAI-2000)*.
- Porello, D.; Troquard, N.; Peñaloza, R.; Confalonieri, R.; Galliani, P.; and Kutz, O. 2018. Two approaches to ontology aggregation based on axiom weakening. In *Proc. 27th International Joint Conference on Artificial Intelligence (IJCAI-2018)*.
- 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* 11:1297–1322.
- Snow, R.; O’Connor, B.; Jurafsky, D.; and Ng, A. Y. 2008. Cheap and fast—but is it good? Evaluating non-expert annotations for natural language tasks. In *Proc. Conference on Empirical Methods in Natural Language Processing (EMNLP-2008)*.
- Tennenholtz, M., and Zohar, A. 2016. The axiomatic approach and the Internet. In *Handbook of Computational Social Choice*. Cambridge University Press. 427–452.
- Terzopoulou, Z., and Endriss, U. 2019. Aggregating incomplete pairwise preferences by weight. In *Proc. 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*.
- Thomson, W. 2001. On the axiomatic method and its recent applications to game theory and resource allocation. *Social Choice and Welfare* 18(2):327–386.
- Xia, L. 2019. *Learning and Decision-Making from Rank Data*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.