Artificial Intelligence 245 (2017) 86–114

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Graph aggregation ☆

Ulle Endriss [a], Umberto Grandi [b]

[a] ILLC, University of Amsterdam, The Netherlands
[b] IRIT, University of Toulouse, France

ARTICLE INFO

Article history:
Received 20 June 2016
Received in revised form 6 January 2017
Accepted 8 January 2017

Keywords:
Social choice theory
Collective rationality
Impossibility theorems
Graph theory
Modal logic
Preference aggregation
Belief merging
Consensus clustering
Argumentation theory

ABSTRACT

Graph aggregation is the process of computing a single output graph that constitutes a good compromise between several input graphs, each provided by a different source. One needs to perform graph aggregation in a wide variety of situations, e.g., when applying a voting rule (graphs as preference orders), when consolidating conflicting views regarding the relationships between arguments in a debate (graphs as abstract argumentation frameworks), or when computing a consensus between several alternative clusterings of a given dataset (graphs as equivalence relations). In this paper, we introduce a formal framework for graph aggregation grounded in social choice theory. Our focus is on understanding which properties shared by the individual input graphs will transfer to the output graph returned by a given aggregation rule. We consider both common properties of graphs, such as transitivity and reflexivity, and arbitrary properties expressible in certain fragments of modal logic. Our results establish several connections between the types of properties preserved under aggregation and the choice-theoretic axioms satisfied by the rules used. The most important of these results is a powerful impossibility theorem that generalises Arrow's seminal result for the aggregation of preference orders to a large collection of different types of graphs.

1. Introduction

Suppose each of the members of a group of autonomous agents provides us with a different directed graph that is defined on a common set of vertices. Graph aggregation is the task of computing a single graph over the same set of vertices that, in some sense, represents a good compromise between the various individual views expressed by the agents. Graphs are ubiquitous in computer science and artificial intelligence (AI). For example, in the context of decision support systems, an edge from vertex x to vertex y might indicate that alternative x is *preferred* to alternative y. In the context of modelling interactions taking place on an online debating platform, an edge from x to y might indicate that argument x

☆ This work refines and extends papers presented at COMSOC-2012 [1] and ECAI-2014 [2]. We are grateful for the extensive feedback received from Davide Grossi, Sylvie Doutre, Weiwei Chen, several anonymous reviewers, and the audiences at the SSEAC Workshop on Social Choice and Social Software held in Kiel in 2012, the Dagstuhl Seminar on Computation and Incentives in Social Choice in 2012, the KNAW Academy Colloquium on Dependence Logic held at the Royal Netherlands Academy of Arts and Sciences in Amsterdam in 2014, a course on logical frameworks for multiagent aggregation given at the 26th European Summer School in Logic, Language and Information (ESSLLI-2014) in Tübingen in 2014, the Lorentz Center Workshop on Clusters, Games and Axioms held in Leiden in 2015, the SEGA Workshop on Shared Evidence and Group Attitudes held in Prague in 2016, and lectures delivered at Sun Yat-Sen University in Guangzhou in 2014 as well as École Normale Supérieure and Pierre & Marie Curie University in Paris in 2016. This work was partly supported by COST Action IC1205 on Computational Social Choice. It was completed while the first author was hosted at the University of Toulouse in 2015 as well as Paris-Dauphine University, Pierre & Marie Curie University, and the London School of Economics in 2016.

E-mail addresses: ulle.endriss@uva.nl (U. Endriss), umberto.grandi@irit.fr (U. Grandi).

http://dx.doi.org/10.1016/j.artint.2017.01.001
0004-3702/

undercuts or otherwise *attacks* argument y. And in the context of social network analysis, an edge from x to y might express that person x is *influenced* by person y. How to best perform graph aggregation is a relevant question in these three domains, as well as in any other domain where graphs are used as a modelling tool and where particular graphs may be supplied by different agents or originate from different sources. For example, in an election, i.e., in a group decision making context, we have to aggregate the preferences of several voters. In a debate, we sometimes have to aggregate the views of the individual participants in the debate. And when trying to understand the dynamics within a community, we sometimes have to aggregate information coming from several different social networks.

In this paper, we introduce a formal framework for studying graph aggregation in general abstract terms and we discuss in detail how this general framework can be instantiated to specific application scenarios. We introduce a number of concrete methods for performing aggregation, but more importantly, our framework provides tools for evaluating what constitutes a "good" method of aggregation and it allows us to ask questions regarding the existence of methods that meet a certain set of requirements. Our approach is inspired by work in social choice theory [3], which offers a rich framework for the study of aggregation rules for preferences—a very specific class of graphs. In particular, we adopt the *axiomatic method* used in social choice theory, as well as other parts of economic theory, to identify intuitively desirable properties of aggregation methods, to define them in mathematically precise terms, and to systematically explore their logical consequences.

An aggregation rule maps any given *profile* of graphs, one for each agent, into a single graph, which we are often going to refer to as the *collective graph*. The central concept we focus on in this paper is the *collective rationality* of aggregation rules with respect to certain properties of graphs. Suppose we consider an agent rational only if the graph she provides has certain properties, such as being reflexive or transitive. Then we say that a given aggregation rule F is collectively rational with respect to that property of interest if and only if F can guarantee that that property is preserved during aggregation. For example, if we aggregate individual graphs by computing their *union* (i.e., if we include an edge from x to y in our collective graph if at least one of the individual graphs includes that edge), then it is easy to see that the property of *reflexivity* will always transfer. On the other hand, the property of *transitivity* will not always transfer. For example, if we aggregate two graphs over the set of vertices $V = \{x, y, z\}$, one consisting only of the edge (x, y) and one consisting only of the edge (y, z), then although each of these two graphs is (vacuously) transitive, their union is not, as it is missing the edge (x, z). Thus, the union rule is collectively rational with respect to reflexivity, but not with respect to transitivity.

We study collective rationality with respect to some such well-known and widely used properties of graphs, but also with respect to large families of graph properties that satisfy certain *meta-properties*. We explore both a semantic and a syntactic approach to defining such meta-properties. In our semantic approach, we identify three high-level features of graph properties that determine the kind of aggregation rules that are collectively rational with respect to them. For example, transitivity is what we call a "contagious" property: under certain circumstances, namely in the presence of edge (y, z), inclusion of (x, y) *spreads* to (x, z). Transitivity also satisfies a second meta-property, which we call "implicativeness": the inclusion of two specific edges, namely (x, y) and (y, z), *implies* the inclusion of a third edge, namely (x, z). The third meta-property we introduce, "disjunctiveness", expresses that, under certain circumstances, at least one of two specific edges has to be accepted. This is satisfied, for instance, by the property of *completeness*: every two vertices x and y need to be connected in at least one of the two possible directions. In our syntactic approach, we consider graph properties that can be expressed in particular syntactic fragments of a logical language. To this end, we make use of the language of modal logic [4]. This allows us to establish links between the syntactic properties of the language used to express the integrity constraints we would like to see preserved during aggregation and the axiomatic properties of the rules used.

We prove both *possibility* and *impossibility results*. A possibility result establishes that every aggregation rule belonging to a certain class of rules (typically defined in terms of certain axioms) is collectively rational with respect to all graph properties that satisfy a certain meta-property. An impossibility result, on the other hand, establishes that it is impossible to define an aggregation rule belonging to a certain class that would be collectively rational with respect to any graph property that meets a certain meta-property—or that the only such aggregation rules would be clearly very unattractive for other reasons. Our main result is such an impossibility theorem. It is a generalisation of Arrow's seminal result for preference aggregation [5], which we shall recall in Section 3.1. Our approach of working with meta-properties has two advantages. First, it permits us to give conceptually simple proofs for powerful results with a high degree of generality. Second, it makes it easy to instantiate our general results to obtain specific results for specific application scenarios. For example, Arrow's Theorem follows immediately from our more general result by checking that the properties of graphs that represent preference orders (namely transitivity and completeness) satisfy the meta-properties featuring in our theorem, yet our proof of the general theorem is arguably simpler than a direct proof of Arrow's Theorem. This is so, because the meta-properties we use very explicitly exhibit specific features required for the proof, while those features are somewhat hidden in the specific properties of transitivity and completeness. Similarly, we show how alternative instantiations of our general result easily generate both known and new results in other domains, such as the aggregation of plausibility orders (which has applications in nonmonotonic reasoning and belief merging) and the aggregation of equivalence relations (which has applications in clustering analysis).

Related work. Our work builds on and is related to contributions in the field of social choice theory, starting with the seminal contribution of Arrow [5]. This concerns, in particular, contributions to the theory of voting and preference aggregation [6–10,3], but also judgment aggregation [11–17]. In fact, in terms of levels of generality, graph aggregation may be regarded as occupying the middle ground between preference aggregation (most specific) and judgment aggregation (most