Artificial Intelligence 215 (2014) 79–119

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach

Mohammad T. Irfan [a,1], Luis E. Ortiz [b,*]

[a] Department of Computer Science, Bowdoin College, Brunswick, ME 04011, United States
[b] Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, United States

ARTICLE INFO

Article history:
Received 13 September 2012
Received in revised form 31 May 2014
Accepted 17 June 2014
Available online 25 June 2014

Keywords:
Computational game theory
Social network analysis
Influence in social networks
Nash equilibrium
Computational complexity

ABSTRACT

We introduce a new approach to the study of influence in strategic settings where the action of an individual depends on that of others in a network-structured way. We propose *network influence games* as a *game-theoretic* model of the behavior of a large but finite networked population. In particular, we study an instance we call *linear-influence games* that allows *both* positive and negative *influence factors*, permitting reversals in behavioral choices. We embrace *pure-strategy Nash equilibrium*, an important solution concept in non-cooperative game theory, to formally define the *stable outcomes* of a network influence game and to predict potential outcomes without explicitly considering intricate dynamics. We address an important problem in network influence, the identification of the *most influential individuals*, and approach it algorithmically using pure-strategy Nash-equilibria computation. *Computationally*, we provide (a) complexity characterizations of various problems on linear-influence games; (b) efficient algorithms for several special cases and heuristics for hard cases; and (c) approximation algorithms, with provable guarantees, for the problem of identifying the most influential individuals. *Experimentally*, we evaluate our approach using both synthetic network influence games and real-world settings of general interest, each corresponding to a separate branch of the U.S. Government. *Mathematically*, we connect linear-influence games to important models in game theory: *potential and polymatrix games*.

1. Introduction

The influence of an entity on its peers is a commonly noted phenomenon in both online and real-life social networks. In fact, there is growing scientific evidence that suggests that influence can induce behavioral changes among the entities in a network. For example, recent work in medical social sciences posits the intriguing hypothesis that much of our behavior such as smoking [16], obesity [15], and even happiness [24] is contagious within a social network.

Regardless of the specific problem addressed, the underlying system studied by Christakis and Fowler exhibits several core features. First, it is often very *large and complex*, with the entities exhibiting different behaviors and interactions. Second, the *network structure of complex interactions* is central to the emergence of collective (global) behavior from individual (local) behavior. For example, in their work on obesity, individuals locally interact with their friends and relatives within their

* Corresponding author. Tel.: +1 631 632 1805; fax: +1 631 632 8334.
E-mail addresses: mirfan@bowdoin.edu (M.T. Irfan), leortiz@cs.stonybrook.edu (L.E. Ortiz).
[1] Parts of this work were done while the author was a PhD student at Stony Brook University.

http://dx.doi.org/10.1016/j.artint.2014.06.004
0004-3702/© 2014 Published by Elsevier B.V.

social network. These local interactions appear to give rise to a global phenomenon, namely, the clustering of medically obese individuals [15]. Third, the *directions and strengths of local influences* are highlighted as very relevant to the global behavior of the system as a whole. Fourth, given that one's behavioral choice depends on others, the individuals potentially act in a *strategic* way.

The prevalence of systems and problems like the ones just described, combined with the obvious issue of often-limited control over individuals, raises immediate, broad, difficult, and longstanding policy questions: e.g., *Can we achieve a desired goal, such as reducing the level of smoking or controlling obesity via targeted, minimal interventions in a system? How do we optimally allocate our often limited resources to achieve the largest impact in such systems?*

Clearly, these issues are not exclusive to obesity, smoking or happiness; similar issues arise in a large variety of settings: drug use, vaccination, crime networks, security, marketing, markets, the economy, public policy-making and regulations, and even congressional voting![2] The work reported in this paper is in large part motivated by such questions/settings and their broader implication.

We begin by providing a brief and informal description of our approach to influence in networks. In the next section, we place our approach within the context of the existing literature.

1.1. Overview of our model of influence

Consider a social network where each *individual* has a *binary* choice of *action or behavior*, denoted by -1 and 1. Let us represent this network as a *directed* graph, where each node represents an individual. Each node of this graph has a *threshold level*, which can be positive, negative, or zero; and the threshold levels of all the nodes are not required to be the same. Each arc of this graph is weighted by an *influence factor*, which signifies the level of *direct* influence the tail node of that arc has on the head node. Again, the influence factors can be positive, negative, or zero and are not required to be the same (i.e., symmetric) between two nodes.

Given such a network, our model specifies the *best response* of a node (i.e., what action it should choose) with respect to the actions chosen by the other nodes. The best response of a node is to adopt the action 1 if the *total influence* on it exceeds its threshold and -1 if the opposite happens. In case of a tie, the node is *indifferent* between choosing 1 and -1; i.e., either would be its best response. Here, we calculate the total influence on a node as follows. First, sum up the incoming influence factors on the node from the ones who have adopted the action 1. Second, sum up those influence factors that are coming in from the ones who have adopted -1. Finally, subtract the second sum from the first to get the total influence on that node.

Clearly, in a network with n nodes, there are 2^n possible *joint actions*, resulting from the action choice of each individual node. Among all these joint actions, we call the ones where every node has chosen its best response to everyone else a *pure-strategy Nash equilibria (PSNE)*. We use PSNE to mathematically model the *stable outcomes* that such a networked system could support.

1.2. Overview of the most-influential-nodes problem

We formulate the most-influential-nodes problem *with respect to a goal of interest*. The goal of interest indirectly determines what we call the *desired stable outcome(s)*. Unlike the mainstream literature on the most-influential-nodes problem [49], maximizing the *spread* of a particular behavior is *not* our objective. Rather, *the desired stable outcome(s) resulting from the goal of interest is what determines our computational objective*. In addition, our solution concept abstracts away the dynamics and does not rely on the "diffusion" process by which such a "spread of behavior" happens.

Roughly speaking, in our approach, we consider a set of individuals S in a network to be a *most influential* set, with respect to a particular goal of interest, if S is the *most preferred* subset among all those that satisfy the following condition: were the individuals in S to choose the behavior $\mathbf{x}_S$ prescribed to them by a desired stable outcome $\mathbf{x} \equiv (\mathbf{x}_S, \mathbf{x}_{-S})$ which achieves the goal of interest, then the *only* stable outcome of the system that remains consistent with their choices $\mathbf{x}_S$ is $\mathbf{x}$ itself.

Said more intuitively, once the nodes in the most influential set S follow the behavior $\mathbf{x}_S$ prescribed to them by a desired stable outcome $\mathbf{x}$ achieving the goal of interest, they become collectively "so influential" that their behavior "forces" every other individual to a unique choice of behavior! Our proposed concept of the most influential individuals is illustrated in Fig. 1 with a very simple example.

Now, there could be many different sets S that satisfy the above condition. For example, S could consist of all the individuals, which might not be what we want. To account for this, we also specify a preference function over all subsets of individuals. While this preference function could in principle be arbitrary, a natural example would be the one that prefers a set S of minimum cardinality.

[2] The headline-grabbing U.S. "debt-ceiling crisis" in 2011, especially the last-minute deal to increase the debt ceiling, is evidence of influence among senators in a strategic setting. We can also view the bipartisan "gang-of-six" senators, specifically chosen to work out a solution, as an *intervention* as such a group would not naturally arise otherwise.