Towards a formal semantics of social influence

Adam Saulwick *, Kerry Trentelman

Defence Science and Technology Organisation, Edinburgh, Australia

ARTICLE INFO

ABSTRACT

Our high-level goal is to answer questions concerned with social influence such as: "Who influences whom?", "Who can be influenced?", "Why is an individual attracted to a particular group?", and "Who is the most influential individual in a particular social network?". To ask these questions we need to define social influence. In this paper we provide a formal definition appropriate to our world of Big Data and automated reasoning. Despite the pervasiveness of influence throughout society and given the vast and disparate literature on the topic, we observe a dearth of work on formalising its semantics. To remedy this, based on the literature, we have categorised and formalised five essential types. To our knowledge this is the first attempt to implement a nuanced representation, and it provides us with a conceptual basis for automated reasoning about social interactions.

1. Introduction

Social influence is pervasive throughout society. It is a fundamental aspect of human social relationships. Our high-level goal is to answer questions concerned with social influence such as: "Who influences whom?", "Who can be influenced?", "Why is an individual attracted to a particular group?", and "Who is the most influential individual in a particular social network?". We believe that answers to these questions help us identify the intent of individuals and groups. To ask these questions we need to define social influence. There are many descriptions of social influence in the literature yet to our knowledge a comprehensive formally implemented definition of the phrase is lacking. In this paper we identify and categorise five essential types.

In the context of Big Data we aim to represent and reason about social influence with our automated high-level information fusion system. We aim to bring together large volumes of relevant, unstructured, heterogeneous information in this system for analysis.[1] In our digital information era we can draw on vast volumes of information about individuals and groups, but our capacity to interpret the relevance of this information in a timely fashion is limited.

We believe we need a system which can automatically process large volumes of information from a variety of sources, *not* in order to supersede human reasoning, but to offer analysts a tool which augments their natural cognitive capabilities and minimises the burden of information processing drudge. In terms of Baranyi and Csapo's model of Cognitive Infocommunications [1], our system can be classified as offering an inter-cognitive mode, and sensor-bridging type, of communication.

We claim that simply storing and retrieving source information in its original form is inadequate because this burdens the human user with its interpretation and discovering the 'needle in the haystack', i.e. identifying relevant information for a given question. Our solution is to utilise the computational capabilities of high speed search, automated reasoning, and natural language question and answering. However in order to represent heterogeneous information consistently and utilise automated reasoning, we need to capture the information's semantics formally. This necessitates establishing a formal theory of the world (including, for current purposes, of social influence) and implementing this in machine processable code. In order for non-specialist users to interact with the automated system, we need an interface which translates between human and machine language. We have developed such a high-level information fusion system—called Consensus—which utilises a formal theory of the world (described in Section 3). Consensus has a multi-modal Controlled Natural Language interface and automated reasoning capabilities, see [2]. Consensus is able to process large volumes of heterogeneous information in the form of text and track data. A virtual adviser can answer questions put to it and present situation reports and alerts about tracked objects in spoken Controlled Natural Language. The virtual adviser is

---

* Corresponding author.
E-mail addresses: adam.saulwick@dsto.defence.gov.au (A. Saulwick), kerry.trentelman@dsto.defence.gov.au (K. Trentelman).

[1] We define *information* as any kind of data stream with semiotic content. This is in contrast to *data* which we define as a set of values or attributes possibly with some structure but not necessarily semiotic content. We leave open precisely who determines whether some piece of matter is considered to have semiotic content as this is beyond the focus of this document.

coordinated with a 3-D geospatial display [3–5]. Consensus can also generate reports written in Controlled Natural Language.

To reach our high-level goal of answering our questions involving social influence we pursue the following:

1. Posit four fundamental drivers underpinning social influence (i.e. kinship, emotional states, agreements, and resource access) and give them a formal semantics.[2] Provide mappings of these to Controlled Natural Language terms.
2. Formally define social influence types and map these to Controlled Natural Language terms.
3. Specify axioms which infer likely influence relationships from combinations of the drivers.
4. Automatically infer likely and existing influence relationships by processing Controlled Natural Language texts containing those terms which have mappings to the drivers and influence types.
5. Interrogate the processed data with Controlled Natural Language queries.

To date our research has focussed on (1) and (2) which we consider essential preliminary steps. (2) is the focus of this paper. The next stage of our research will focus on (3) and (4) as this will provide the mechanism to undertake (5) in order to achieve our goal. A technical study can then be conducted to evaluate this approach.

The remainder of the paper is structured as follows. Section 2 reviews the literature on social influence and informally introduces our distinct types of social influence. In Section 3 we provide formal definitions of key influence terms. Section 4 summarises our findings and posits further research questions.

2. Social influence

Our understanding of social influence is informed by the literature and we refer to relevant theories below.

2.1. Influence in the literature

There is a voluminous literature on social influence specialised in many domains, including social science, culture, leadership, psychology, anthropology, linguistics, behavioural science, management and marketing studies. Different disciplines use different terminology. We now discuss research germane to our concerns.

Social psychologists French and Raven [6–8] identified five 'bases' of power: *reward power*, *coercive power*, *the power of legitimate authority*, *referent power*, and *expert power*. With minor adjustments [9,10] these five bases reoccur as the received wisdom on social power relations. For French and Raven, the process of influencing draws on a power base stemming from an individual's position, access to material possessions, skills, expertise and charisma. This places power squarely inside each individual and downplays the social context. Of relevance to us is their distinction between positive and negative forms of referent and expert power [11]. For example, a positive referent power involves the recipient liking or identifying with the influencer. Negative referent power has the opposite outcome.

For Deutsch and Gerard [12], social influence is *normative*, pressuring individuals to conform to the expectations of others. They suggest people avoid behaving in ways which will lead to social punishment or disapproval. Social influence is also a product of an individual's desire to be informed about a complex, ambiguous world. When a person is uncertain they seek solace in the interpretations of others. This locates social influence in the realm of social groups, peer pressure and opinion. Our model picks up on the group dynamic aspect of this theory.

Kelman [13] separates the influence process into a 1–1 mapping between antecedents and consequents. The antecedents, *means-control*, *attractiveness* and *credibility*, represent sources from which an individual exerts power.[3] The consequents represent three types of accepting influence: *compliance*, *identification* and *internalisation*. Compliance occurs in order to gain a reward or avoid punishment. Identification occurs in order to establish or maintain positive interpersonal relationships, whereas internalisation occurs in order to maintain a congruent value system [14]. Kelman's identification of influence types strikes a chord with our model which highlights influence as a process of interpersonal relationships within a social group.

Turner's approach has been influential for our theory because it focuses on an individual's self-categorisation with respect to their group membership [15]. Turner posits a three-process theory of influence involving persuasion, authority and coercion. For Turner, influence stems from psychological group formation [16]. He suggests persuasion is a function of shared social identity and consensus which flows from changes in an individual's relative 'prototypicality'. Prototypicality reflects how central or peripheral a given group member is to the group with respect to its values. Of relevance for us is the notion of collective social identity stemming from the group's value system. We also adopt the idea that groups admit or expel members based on their prototypicality or lack thereof [17].

Finally, Cialdini posits seven persuasion tactics, popularised in [18]. These are: *reciprocity*, *commitment*, *consistency*, *social proof*, *authority*, *liking* and *scarcity*. We see these as tactics which exploit people's feelings of obligation, potential embarrassment, desire to conform, comply and be liked, and fear of loss of opportunity. Cialdini suggests that individuals have automatic responses and default modes of behaviour, labelled as 'click-whirr' responses. Though claimed as necessary and important to achieve unimpeded functional social coherence, they leave an individual susceptible to manipulation.

Some other agent-based work [19–23] models social processes without necessarily dissecting the semantics of influence, whereas our emphasis is to capture the differences formally.

2.2. Our theory of social influence

Having briefly surveyed the literature, we now describe our notions of cognitive individual and social group because these are crucial to any occurrence of social influence. We give these concepts a formal treatment in Section 3.2.

For us, a cognitive individual is an entity which has dispositions towards propositional claims about the world. Such an individual can be thought of as the sum of their attitudes which have been and continue to be internalised. These attitudes are based on a myriad of material interactions with their environment and are influenced continually by other cognitive individuals.

For us, social groups are collections of cognitive individuals. They are bound by conventionalised behaviour which stems from two forces: an imperative for social interactions, and the need for some interactions to function with click-whirr ease. Social groups are dynamic, open systems reacting to new stimuli. Repeated group behaviour becomes conventionalised over time and this reinforces social cohesion.[4] Behaviour cannot become conventionalised

---

[2] We believe that combinations of these drivers mediate the likelihood of particular influence relations. Space constraints prohibit further explanation here.

[3] These correspond to the coercive-reward, referent and expert power bases of the French and Raven system.

[4] Conventionalised behaviour can be codified in scripture, doctrine, or law; or it may be implicitly adopted, as in natural language, or other social mores, norms and tropes. Clearly some conventions are detrimental to the success of a society, however we deliberately refrain from taking a value-based stance on a particular convention, as this is outside the scope of this paper.