Decision Support Systems 66 (2014) 102–113

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Filtering trust opinions through reinforcement learning

Han Yu [a,*], Zhiqi Shen [a], Chunyan Miao [a], Bo An [a], Cyril Leung [b]

[a] School of Computer Engineering, Nanyang Technological University 637659, Singapore
[b] Department of Electrical and Computer Engineering, the University of British Columbia, Vancouver, BC V6T1Z4, Canada

A R T I C L E I N F O

Article history:
Received 23 September 2013
Received in revised form 4 June 2014
Accepted 10 June 2014
Available online 26 June 2014

Keywords:
Trust
Reputation
Credibility
Collusion

A B S T R A C T

In open online communities such as e-commerce, participants need to rely on services provided by others in order to thrive. Accurately estimating the trustworthiness of a potential interaction partner is vital to a participant's well-being. It is generally recognized in the research community that third-party testimony sharing is an effective way for participants to gain knowledge about the trustworthiness of potential interaction partners without having to incur the risk of actually interacting with them. However, the presence of biased testimonies adversely affects a participant's long term well-being. Existing trust computational models often require complicated manual tuning of key parameters to combat biased testimonies. Such an approach heavily involves subjective judgments and adapts poorly to changes in an environment. In this study, we propose the Actor–Critic Trust (ACT) model, which is an adaptive trust evidence aggregation model based on the principles of reinforcement learning. The proposed method dynamically adjusts the selection of credible witnesses as well as the key parameters associated with the direct and indirect trust evidence sources based on the observed benefits received by the trusting entity. Extensive simulations have shown that the ACT approach significantly outperforms existing approaches in terms of mitigating the adverse effect of biased testimonies. Such a performance is due to the proposed accountability mechanism that enables ACT to attribute the outcome of an interaction to individual witnesses and sources of trust evidence, and adjust future evidence aggregation decisions without the need for human intervention. The advantage of the proposed model is particularly significant when service providers and witnesses strategically collude to improve their chances of being selected for interaction by service consumers.

1. Introduction

In open online communities where users are from diverse backgrounds and may have conflicting interest, trust-based interaction decision support is needed to sustain long term interactions among them. Nowadays, such systems are quite common (e.g., service oriented computing systems [1], e-commerce systems [2], wireless communication networks [3], and multi-agent systems [4] etc.). In such environments in which services and devices usually have limited capabilities, users often have to interact with each other in order to complete complex tasks. These interactions usually involve an exchange of services, information, or goods with value. Selfish users may renege on their commitments, thereby breaching the trust placed in them by others. Therefore, trust and reputation management mechanisms are often used to minimize the negative impact of selfish users.

1.1. Background

Generally, users in an open online community that can be modeled as multi-agent systems (MASs) may play two types of roles [1]:

- *service providers* (SPs), who provide services, goods or information requested by others and do not need to rely on others to perform these services; and
- *service consumers* (SCs), who need to rely on service providers to accomplish certain tasks.

The main objective of evidence-based trust models is to estimate the trustworthiness of a potential interaction partner which represents its true behavior pattern. Evidences about a service provider from the perspective of a service consumer are usually from two sources:

- *direct trust evidence*: which consists of a service consumer's direct interaction experience with the service provider; and
- *indirect trust evidence*: which consists of third-party testimonies about the service provider from other service providers in the system.

In practical systems, it is not possible to definitively know the trustworthiness of a service provider. Therefore, it is often estimated using trust evidences. The estimation of a service provider's trustworthiness derived from the direct trust evidence of a service consumer alone is called *direct trust*, while that derived from the indirect trust evidence is called *indirect trust*. According to [4], an estimation derived from both sources of trust evidence is commonly known as the *reputation* of a service provider. In the eyes of a service consumer, other service consumers who provide it with indirect trust evidence (i.e. testimonies)

* Corresponding author.
E-mail address: han.yu@ntu.edu.sg (H. Yu).

http://dx.doi.org/10.1016/j.dss.2014.06.006
0167-9236/

about a service provider are regarded as *witnesses*. A witness's reliability in terms of providing useful testimonies is referred to as its *credibility*.

Since such systems tend to be very large in practice, service consumers often have to interact with service providers with whom they may not be very familiar (i.e. have little or no prior interaction experience with) [5]. Thus, it is both necessary and advantageous to allow service consumers to act as witnesses to provide their own first-hand interaction experience as testimonies to other service consumers who lack such information. However, such an approach is not without its perils.

Third-party testimonies may be biased and, thus, degrade the accuracy of trust decisions [1]. Therefore, testimonies from witnesses need to be filtered before being used to evaluate a service consumer's reputation.

To this end, a number of evidence-based trust and reputation management (TRM) models have been proposed over the years. The general flow for a service consumer to decide which service provider to select for interaction is illustrated in Fig. 1. Each service consumer continuously records its direct interaction experience with service providers over time. When a service provider's trustworthiness needs to be evaluated, the service consumer may request third-party testimonies from witnesses, depending on the service consumer's confidence on its own direct trust evidence. These testimonies are preprocessed in an attempt to filter out unfair ratings. The resulting direct and indirect trust evidences are then aggregated to form a trustworthiness evaluation for that particular service provider. At the end of this process, the service consumer decides which service provider to interact with based on their trustworthiness evaluations.

1.2. Research objectives

Existing approaches for third-party testimony filtering and aggregation commonly involve a crucial step in which the weight assigned to each third-party testimony and the weight assigned to the direct and the indirect sources of trust evidence need to be determined [6–9].

However, existing approaches often require manual tuning of key parameters in their models which heavily involves subjective judgments and adapts poorly to changes in the environment.

In this paper, we address this limitation by proposing the Actor–Critic Trust (ACT) model based on the principles of the Actor-Critic Learning Method [10]. The ACT approach automates the adjustment of key threshold based parameters to eliminate human subjectivity and enhance the effectiveness of existing reputation evaluation models. Specifically, it enables existing evidence-based trust models to dynamically make two important decisions when presented with third-party testimonies for a service provider: 1) how much weight to assign to its own personal direct trust evidence and the collective opinions from witnesses, and 2) how much weight to assign to the testimonies from each witness. Experimental results, presented in Section 4, show that the ACT approach outperforms state-of-the-art approaches by around 20% in terms of improving the accuracy of finding trustworthy service providers in the presence of biased testimonies, especially when witnesses collude with malicious service providers.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 presents the basic notations used in this paper and the details of the proposed ACT approach. Section 4 describes the simulation test-bed and analyzes the results. The implications of the proposed approach for practical decision support in online product review systems are discussed in Section 5. Finally, Section 6 presents a summary of our contributions and possible future work.

2. Related work

It is widely recognized within the research community that the importance of incorporating mechanisms to mitigate the adverse effects of biased testimonies. In this section, we discuss some recent research work on aggregating trust evidence from different sources and filtering out biased testimonies. For a more comprehensive review of this field, readers may refer to [1–3].

2.1. Trust evidence aggregation approaches

Evidence-based trust models often make use of two distinct sources of information to evaluate the trustworthiness of a service provider: *direct trust evidence* and *indirect trust evidence*. The majority of existing trust models adopt a weighted average approach when aggregating these two sources of trust evidence [3]. Direct trust evidence is often assigned a weight of $0 \leq \gamma \leq 1$, and indirect evidence is assigned a corresponding weight of $1 - \gamma$. Existing approaches for aggregating direct and indirect trust evidence can be divided into two broad categories: 1) *static approaches*, where the value of γ is pre-defined; and 2) *dynamic approaches*, in which the value of γ is continually adjusted by the service consumer.

In many papers, static γ values for trust evidence aggregation. The majority of them tend to take a balanced approach by assigning a value of 0.5 to γ [6,9,7,11,12]. In some studies, the authors assign the value 0 [13,14] or 1 [15] to γ to exclusively use only one source of trust information. Barber and Kim [16] have empirically shown, without considering the presence of biased testimonies, that direct trust evidence is the most useful to a service consumer over the long term

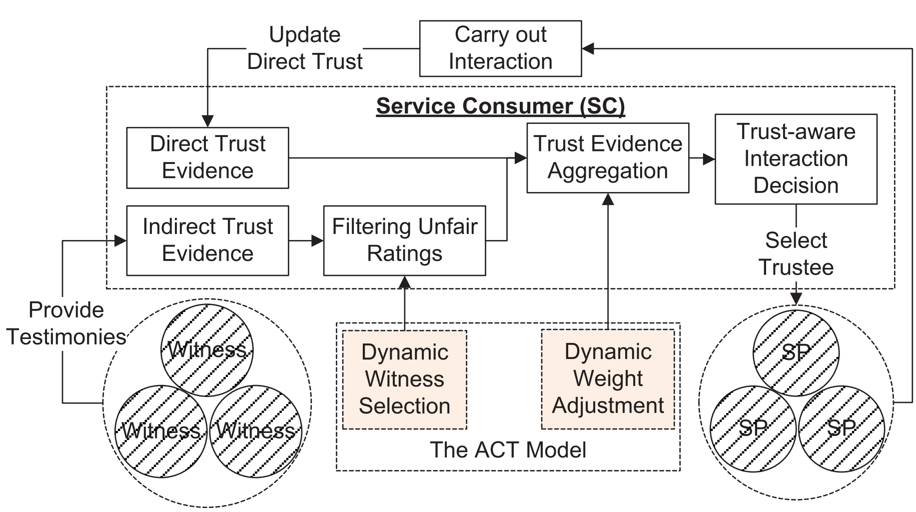

Fig. 1. The general flow of trust–aware interaction decision making for evidence-based trust and reputation management models, and the contributions by the proposed ACT approach.