Automatica 83 (2017) 351–360

Contents lists available at ScienceDirect

Automatica

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

Peer-Assisted Individual Assessment in a multi-agent system ☆

Wenjie Li [a], Francesca Bassi [b,a], Laura Galluccio [c], Michel Kieffer [a,d]

[a] L2S, CNRS, Supelec, Univ Paris-Sud, 3 rue Joliot-Curie, 91192 Gif-sur-Yvette, France
[b] ESME-Sudria, 38 rue Molière, 94200 Ivry-sur-Seine, France
[c] DIEEI, Università degli studi di Catania, viale Andrea Doria 6, Catania, Italy
[d] Institut Universitaire de France, 103 bld Saint-Michel, 75005 Paris, France

ARTICLE INFO

Article history:
Received 25 July 2016
Received in revised form 16 February 2017
Accepted 19 May 2017
Available online 12 July 2017

Keywords:
Classification
Delay tolerant networks
Distributed estimation
Equilibrium
Multi-agent systems
Sensing

ABSTRACT

Consider a multi-agent system where agents perform a given task with different levels of ability. Agents are initially not aware of how well they perform in comparison with their peers, and are willing to self-assess. This scenario is relevant, *e.g.*, in wireless sensor networks, or in crowdsensing applications, where devices with embedded sensing capabilities collaboratively collect data to characterize the environment: the global performance is very sensitive to the measurement accuracy, and agents providing outliers should restrain to participate.

This paper presents a distributed algorithm enabling each agent to self-assess its own ability. The algorithm tracks the outcomes of a local comparison test performed by pairs of agents when they randomly meet, and able to gauge their relative level of ability. The dynamics of the proportions of agents with similar assessments are described using continuous-time state equations. The existence of an equilibrium is shown. Closed-form expressions for the various proportions of agents with similar assessments are provided at equilibrium. In simulations, a community of agents equipped with sensors, and trying to determine the performance of their equipment is considered. Simulation results show a good fitting with theoretical predictions.

1. Introduction

Consider a community of agents collaborating to execute some task, *e.g.*, sensing, detection, classification (see Ang, Gopalkrishnan, Ng, & Hoi, 2009; Luo, Xiong, Lü, & Shi, 2007; Shah, Balakrishnan, & Wainwright, 2016). Agents are expected to have different levels of ability (LoAs) in carrying out atomic operations. We assume that LoAs are represented by positive integers or real numbers[1]. This paper considers a Peer-Assisted Individual Assessment (PAIA) problem, in which each agent of a community aims at learning its own LoA from pairwise, sporadic interactions, as in delay tolerant networks (Khabbaz, Assi, & Fawaz, 2012), or in networks where exchanges are performed via gossiping (Dimakis, Kar, Moura, Rabbat, & Scaglione, 2010). In this paper a distributed PAIA algorithm is proposed. The latter is able to address the PAIA problem in the absence of any central ranking authority.

The PAIA problem is of interest in several scenarios. In a wireless sensor network (WSN) (Yick, Mukherjee, & Ghosal, 2008), for instance, devices with embedded sensors collaboratively collect data to characterize the surrounding environment. Agents who have only incomplete knowledge of the characteristics of the sensing noise (*e.g.*, either biased or unbiased, as in Chiuso, Fagnani, Schenato, & Zampieri, 2011), may use the PAIA algorithm to estimate it. Similarly, in crowdsensing applications (Guo, Wang, Yu, Wang, Yen, & Zhou, 2015) data generated by personal mobile devices are collected in order to estimate some process. Since the reliability of the service depends on the accuracy of the measurements, the server prefers to pull data from the devices with the most accurate sensors. The problem of device selection is usually addressed by centralized reputation-based mechanisms, (see, *e.g.*, Kantarci, Mouftah, & August, 2014; Ren, Zhang, Zhang, & Shen, 2015; Yu & Schaar, 2012), where the devices apply via an auction system and the server selects on the basis of their established reputation level. Using the PAIA algorithm to assess their own accuracy, agents aware to be temporarily producing outliers may decide to restrain from an auction, to preserve their reputation at the central authority; agents, aware that their accuracy is above the average, may negotiate a better reward.

☆ This work has been partly supported by the Newcom# NoE, and by Paris-Saclay ICODE 2 LABEX. The material in this paper was partially presented at the 55th IEEE Conference on Decision and Control, December 12–14, 2016, Las Vegas, NV, USA. This paper was recommended for publication in revised form by Associate Editor Shreyas Sundaram under the direction of Editor Christos G. Cassandras.

E-mail addresses: weli@l2s.centralesupelec.fr (W. Li), bassi@l2s.centralesupelec.fr (F. Bassi), laura.galluccio@dieei.unict.it (L. Galluccio), kieffer@l2s.centralesupelec.fr (M. Kieffer).

1 Here, one considers discrete-valued LoAs, which may be obtained, *e.g.*, by the quantization of real-valued LoAs.

http://dx.doi.org/10.1016/j.automatica.2017.06.031
0005-1098/

The PAIA problem can be viewed as a generalization of distributed faulty node detection (DFD). In DFD part of the nodes of a network are equipped with defective sensors producing measurement outliers (Mahapatro & Khilar, 2013; Zhang, Meratnia, & Havinga, 2010). Each node is willing to estimate the status of its own sensor (good or defective) (Chen, Kher, & Somani, 2006; Lee & Choi, 2008; Li, Bassi, Dardari, Kieffer, & Pasolini, 2016a). DFD can thus be seen as a PAIA problem for two LoAs. This paper extends to PAIA the ideas introduced for DFD in Li, Galluccio, Bassi, and Kieffer (2016b) by considering more than two possible LoAs for each agent. In the proposed PAIA algorithm a pairwise interaction results in a local comparison test (LCT) able to gauge the relative strength of the participants. Each agent observes only the outcomes of the LCTs it has been involved in. Based on the proportion of interactions during which it has been deemed better, the agent is able to iteratively determine its own LoA. The algorithm parameters depend on the proportions of agents with the same LoA and on the probabilities of error of the LCT.

Ranking or classification by pairwise comparisons has been of interest for a long time. In this work we presuppose, as in classical models (Bradley & Terry, 1952; Luce, 1959; Thurstone, 1927), that an inherent partition of the agents according to their LoA exists, and that the outcomes of the comparisons are probabilistic. Unlike in parametric models, however, we limit the assumptions made on the matrix of the outcome probabilities. Besides classical *win or lose* pairwise comparisons, we account also for weaker LCTs, where the outcome only indicates whether the participants have comparable strength. Recent years are seeing renewed interest in ranking by pairwise comparisons (see, *e.g.* Heckel, Shah, Ramchandran, & Wainwright, 2016; Jamieson & Nowak, 2011; Negahban, Oh, & Shah, 2012; Wauthier, Jordan, & Jojic, 2013, and references therein). In these works a central authority observes the whole collection of outcomes, and usually directs the measurement process. In the PAIA problem, on the contrary, no agent centralizes all the data, and the agents cannot select the peers they interact with. Since distributed classification usually refers to agents cooperating to rate a set of objects, as in Ang et al. (2009) and Luo et al. (2007), whereas PAIA refers to agents rating *themselves*, PAIA can be defined as a distributed self-classification problem from pairwise comparisons.

The paper is organized as follows. Section 2 reviews additional related work. Section 3 introduces the system model and the LCT. Section 4 describes and analyzes the proposed PAIA algorithm. Its effectiveness is measured by the proportions of agents who assess their LoAs correctly. The analysis is performed by assuming a well-mixed population of agents, with intercontact delay following an exponential distribution (Galluccio, Lorenzo, & Glisic, 2016; Hernandez-Orallo, Serra. Olmos, Cano, Calafate, & Manzoni, 2015; Zhu, Fu, Xue, Zhu, Li, & Ni, 2010). This communication model allows one to derive continuous-time state equations approximating the evolution in time of the proportions of agents with similar self-assessments. The existence of an equilibrium is shown in Section 5, and closed-form expressions for the proportions of agents with similar assessments at equilibrium are provided. The dependence of the correct decision rate and of the false decision rate on the characteristics of the LCT provides insights on the way the PAIA algorithm should be tuned to trade-off between them. Section 6 reports simulation results for a population of agents aiming to determine the LoA of their embedded sensors. The numerical results show and excellent match with the theoretical predictions. Finally, conclusions are drawn in Section 7.

2. Related work

In Chiuso et al. (2011), Fagnani, Fosson, and Ravazzi (2014a, 2014b) each node of a WSN estimates from noisy measurements the value of some constant parameter, jointly with the bias (Chiuso et al., 2011; Fagnani et al., 2014a) or the variance (Fagnani et al., 2014b) of the noise. The bias or the variance determines the LoA of the agent. These works involve at least partially instances of the PAIA problem. In Chiuso et al. (2011) the nodes belong to one of two classes, defined by the absence or presence of the bias. The algorithm involves a gossip consensus, robust against node mobility, and a distributed ranking of the agents (Fagnani & Zampieri, 2008) according to their measured value. The signal model is extended to vector measurements and to multiple bias in Fagnani et al. (2014a). Each node uses consensus (Huang & Manton, 2009) to cooperate with the neighbors for the common estimation problem, while iteratively estimating its own local bias. Consensus algorithms are used also in Fagnani et al. (2014b), where the two classes depend on the possible values of the noise variance. Notice that in Chiuso et al. (2011) and Fagnani et al. (2014a, 2014b) the PAIA problem is solved by estimation and is thus bound to the considered signal model. The PAIA algorithm does not make assumptions on the nature of the measurements, but only presupposes a generic LCT, characterized by its probabilities of error. For example, the LCT may compare noisy measurements of some constant parameter as in Chiuso et al. (2011), or results of the same supervised image classification performed by two agents, or be a match when agents have to assess their level in a game.

When the number of possible LoAs equals the number of agents, the PAIA problem is equivalent to the distributed self-ranking problem addressed in a centralized way in Heckel et al. (2016) and with a distributed approach as in Fagnani and Zampieri (2008).

3. System model and local comparison test

Consider a set $\mathcal{A}$ of N_{A} moving agents. Let $\theta_i \in \Theta = \{1 \dots K\}$ be the LoA of Agent i. $\mathcal{A}$ is partitioned into K groups denoted $\mathcal{A}_1 \dots \mathcal{A}_K$, with $\mathcal{A}_\theta = \{i \in \mathcal{A} : \theta_i = \theta\}$. Denote p_θ the proportion of the agents belonging to $\mathcal{A}_\theta$. Without loss of generality, we assume that the groups are sorted in decreasing LoA: thus, the agents in $\mathcal{A}_1$ are the best-performing and those in $\mathcal{A}_K$ are the worst-performing. The following assumption is made:

- (A1) $\theta_i(t) = \theta_i$, *i.e.*, the LoA of Agent i does not change during the experiment.

Agent i is not aware of the actual value of θ_i but is willing to estimate it as fast as possible. To accomplish this it exploits data obtained interacting with other agents. As in Li et al. (2016b), consider the following assumptions:

- (A2) only pairwise meetings are considered;
- (A3) the agents form a well-mixed population, *i.e.* the probability that the next meeting of Agent i will be with an agent belonging to $\mathcal{A}_\theta$ is proportional to $|\mathcal{A}_\theta|$;
- (A4) the time interval between two successive meetings of Agent i with any other agent follows an exponential distribution with inter-contact rate λ (Galluccio et al., 2016; Zhu et al., 2010).

During a meeting two agents may engage in an interaction. Interactions take different forms depending on the application scenario, for example, the exchange of noisy measurements m_i and m_j of the same physical quantity when the agents are nodes of a WSN, or a blitz-game between humans, willing, *e.g.*, to assess their playing level. A meeting does not necessarily entail interaction. Define $\alpha(\widehat{\theta}_i, \widehat{\theta}_j)$ as the probability of interaction of Agents i and j meeting at instant t. It is a function of the current estimates $\widehat{\theta}_i(t)$ and $\widehat{\theta}_j(t)$ of the agents. When Agents i and j meet,