

# Applying inferential processes to partner selection in large agents communities

Pasquale De Meo<sup>a</sup>, Rino Falcone<sup>b</sup> and Alessandro Sapienza<sup>b</sup>

<sup>a</sup>Department of Ancient and Modern Civilizations (University of Messina), Messina, 98122, Italy

<sup>b</sup>Institute of cognitive sciences and technologies (ISTC-CNR), Rome, 00185, Italy

## Abstract

The current literature clearly highlighted the need to define a fast and efficient tool for trust assessment, even in lack of direct information, as much as possessing mechanisms allowing a matching between a selected task and a reliable agent able to carry it out. Direct experience plays a big part, yet it requires a long time to offer a stable and accurate performance and this characteristics may represents a strong drawback especially within huge agents' communities. We support the idea that category-based evaluations and inferential processes represent a useful resource for trust assessment. Within this work, we exploit simulations to investigate how efficient this inferential strategy is, with respect to direct experience, focusing on when and to what extent the first prevails on the latter. Our results suggest that in some situations categories represent a valuable asset, providing even better results.

## Keywords

trust, inference, multi-agent systems

## 1. Introduction and Background

A great deal of effort has been made to assess trust within agents' societies [1][2][3][4]. The great majority of the approaches makes extensive use of direct experience as the main source of information, considering recommendation/reputation and inferential processes just later, as a secondary mechanism to refine trust assessment. Unfortunately, direct experience might not always represent the best solution to assess trustworthiness. For instance, using direct experience within huge networks may become unfeasible, because of their very nature: it is hard and extremely costly to possess enough experience to judge a sufficient number of agents. It is therefore fundamental to find an effective approach for trust assessment even in lack of direct experience [5]. In particular, we argue that category-based evaluations and inferential processes may represent an effective solution. Let us then introduce the formal framework within which we will work.

We consider a population of agents  $\mathcal{A} = \{a_1, \dots, a_n\}$  which can collaborate by reciprocally delegating the execution of some tasks.

We suppose that any agent  $a_i \in \mathcal{A}$  needs to achieve a goal  $g$ , which identifies a state of the environment (in short, the *world*) which the agent  $a_i$  plans to achieve. The agent  $a_i$  can reach

---

WOA 2020: Workshop "From Objects to Agents", September 14–16, 2020, Bologna, Italy

✉ pdemeo@unime.it (P. D. Meo); rino.falcone@istc.cnr.it (R. Falcone); alessandro.sapienza@istc.cnr.it (A. Sapienza)

🆔 0000-0001-7421-216X (P. D. Meo); 0000-0002-0300-458X (R. Falcone); 0000-0001-9360-779X (A. Sapienza)



© 2020 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

the goal  $g$  by executing a *task*  $\tau$ , which affects the world. The most interesting case occurs if the agent  $a_i$  want/must delegate the execution of  $\tau$ .

At an abstract level, each agent possesses some skills and resources, defined here as *features* which determine its ability in carrying out the tasks it has to face. Nevertheless, not all the features associated with an agent are crucial for the execution of  $\tau$  and the majority of these will not even be necessary. If I were to ask someone to cook for me, it would be interesting to know how fast she/he is, how good is her/his culinary imagination or if she/he knows how to cook specific dishes; however, knowing that she/he loves reading astrophysics books would not be of any help.

It is therefore a fundamental precondition that an agent  $a_i$  identifies which features are necessary to carry out  $\tau$ . Then,  $a_i$  needs a mental representation of any other  $a_j$ , which comprises, at least, the subset of the features which are relevant to execute  $\tau$ . It is also important to underline that just the possession of these features is not enough, it is also very relevant  $a_j$ 's willingness (following its motivations) to actually realize  $\tau$ . Of course, two different tasks, say  $\tau_1$  and  $\tau_2$ , require different features to be efficiently addressed.

Thanks to its mental model,  $a_i$  is able to estimate the likelihood  $\phi(i, j)$  that  $a_j$  will positively bring to completion that specific agreed task, for each agent  $a_j \in \mathcal{A}$ . The function  $\phi(i, j)$  measures the *degree of trust* [6] that  $a_i$  (hereafter, the *trustor*) puts in  $a_j$  (hereafter the *trustee*), i.e., quantifies to what extent  $a_i$  is confident that  $a_j$  is capable of successfully executing  $\tau$ .

It is crucial to point out that the assessment of trust is not only *task-dependent* but also *context-dependent*, because external causes may amplify or downsize the trust between the trustor and trustee. For instance, assume that  $a_i$  wants to get to the airport one hour before the departure of her/his flight and suppose that  $a_i$  is confident that  $a_j$  is able to safely drive and she/he is knowledgeable of obstacles to traffic flow (e.g., limited access roads), and thus,  $a_i$  puts a high degree of trust in  $a_j$ . However, unforeseen circumstances (e.g.,  $a_j$  stuck in a traffic jam) may prevent  $a_i$  from being at the airport at the scheduled time: such an event negatively influence the trust from  $a_i$  to  $a_j$ , even if, of course, the liability of  $a_j$  is limited.

The procedure to select the agent to which the task  $\tau$  has to be delegated is thus entirely driven from the calculation of the function  $\phi(i, j)$ : the trustor should select the agent  $a_j^*$  for which  $\phi(i, j)$  achieves its largest value, i.e.,  $j^* = \arg \max_j \phi(i, j)$ . Such a protocol is, unfortunately, infeasible in real-life applications: in fact,  $a_i$  is capable of estimating the trust of those agents – in short  $E_i$  – with which it interacted in the past and of which it knows features. In real-life applications, we expect that the size of  $\mathcal{A}$  is much larger than that of  $E_i$  and, thus, the search of a successful partner is likely to end up in a failure.

An elegant solution to the problem of selecting partners in large agent communities is described in [7, 8] and it relies on the concept of *agent category* or, in short, *category*.

Broadly speaking, a category is a subset of agents in  $\mathcal{A}$  such that each category member possesses homogeneous features. Their unique nature makes categories very interesting and particularly useful. Since the members of a category possess similar features, even their performance concerning the same task will be similar. For sure, we have to consider a certain degree of uncertainty, due to the specific peculiarities of the individuals.

The specific categories to take into consideration change with the context and with the task of interest. For instance, suppose that  $\mathcal{A}$  correspond to a community of people working in food service with different roles; chefs, waiters, and sommeliers are possible examples of categories

in this context.

Because of the existence of categories, the set of agents that the trustor can evaluate significantly expands in size and it consists of the following type of agents:

1. The aforementioned set  $E_i$ , which consists of the agents with which  $a_i$  has had a direct experience.
2. The set  $C_i$  of agents, such that each agent  $a_j \in C_i$  belongs to at least one of the categories  $CS = \{C_1, C_2, \dots, C_p\}$ ; here we suppose that  $a_i$  has had a direct experience with at least one agent in each of the categories in  $CS$ .
3. The set of agents  $R_i$  with which  $a_i$  had no direct experience but which have been recommended to  $a_i$  by other agents in  $\mathcal{A}$  (for instance, on the basis of their reputation).
4. The set of agents  $RC_i$ , such that each agent in  $RC_i$  belongs to a category which contain at least one agent in  $R_i$ .

Advantages arising from the introduction of categories have been extensively studied in past literature [9, 8, 7]: the trustor, in fact, could be able to estimate the performance of any other agent  $a_j$ , even if it has never met this agent (and, as observed in [7] without even suspecting its existence), through an *inferential mechanism*.

As the authors of [9] say, it is possible to take advantage of categories just if a few conditions are met. First of all,  $\mathcal{A}$  must be partitioned into the categories  $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ , classifying the agents according to their features. We assume that this classification is given and accepted by all the agents in  $\mathcal{A}$ . It must be possible to clearly and unequivocally link  $a_j$  to a category  $c_l$ . Finally, we must somehow identify the average performance of the category  $C_l$  with respect to the task  $\tau$ : we will discuss in detail in Section 3 a procedure to estimate the performance – called true quality –  $\theta_l(\tau)$  of the category  $C_l$  for task  $\tau$ .

When all three of these conditions are met, then the category  $C_l$ 's evaluation can be used for the agent  $a_j$ , concerning the task  $\tau$  since, by definition of category, all agents in  $C_l$  will share the same features of  $a_j$  and, thus, if the other agents in  $C_l$  are able to successfully execute the task  $\tau$  (or not), we can reasonably assume that even  $a_j$  can do it (or not).

Of course, only some of the categories  $C_1, \dots, C_m$  possess the qualities to successfully execute the task  $\tau$  while others do not. As a consequence, the first step to perform is to *match* the task  $\tau$  with a set of categories capable of executing  $\tau$ .

At a basic level, such a matching could be implemented through a function  $\psi(C_l, \tau)$  which takes a category  $C_l$  and a task  $\tau$  and returns `True` if agents in  $C_l$  are capable of executing  $\tau$ , `False` vice versa. The computation of the function  $\psi$  requires an analytical and explicit specification of: (a) the chain of actions to perform to execute  $\tau$  and (b) for each action mentioned in (a), the features an agent should possess to perform such an action.

The protocol above easily generalizes to the case in which the trustor has a limited experience (or, in the worst case it has no previous experience): in this case, in fact, the trustor  $a_i$  could leverage the sets of agents  $R_i$  and  $RC_i$ .

## 2. Related Work

The growing need to deal with bigger and bigger agents' networks makes it difficult to find reliable partners to delegate tasks. It becomes clear that, in such situations, direct experience

[10] is not enough to allow us facing this problem. Going beyond this dimension becomes essential, on the light of the knowledge we already have, identifying models and methodologies able to evaluate our interlocutors and possibly to select adequate partners for the collaborative goals we want to pursue.

Several authors proposed trust propagation as a solution to this topic. Trust propagation [11][12] starts from the assumption that if  $a_i$  trusts  $a_j$  and  $a_j$  trusts  $a_k$ , then it is reasonable to assume that  $a_i$  can trust  $a_k$  to some extent. Exploiting this and other assumptions, this technique allows propagating a trust value from an agent to another one, without requiring a direct interaction. The confusion in the reasoning process here is due to the consideration of a generic trust value for an individual, leaving aside the reason why we trust it: the task we want to delegate to it.

Many articles have discussed the use of categories/stereotypes in trust evaluations [13][14]. This is a very useful instrument, allowing to generalize an individual's evaluation, concerning a specific task to other agents owning similar characteristics. It represents a useful proxy for individuating knowledge about specific trustees [15], elicited in particular in all those circumstances precluding the development of person-based trust [16]. Here the intuition is that, given a specific task  $\tau$ , the performance of the agent we are evaluating are related to the values of the features it needs to carry out the task itself. Along these lines, it is natural to assume that other individual owning similar values, i.e. belonging to the same category, have the same potential to solve  $\tau$ .

Pursuant to these considerations, our contribution within this work concerns the investigation of how efficient this inferential strategy is, with respect to direct experience, focusing on when and to what extent the first prevails on the latter.

### 3. Inferring the quality of categories

In this section we illustrate our procedure to estimate the performance (in short called the *true quality*)  $\theta_l(\tau)$  of agents in the category  $C_l$  to successfully execute a particular task  $\tau$ .

Because of the assumptions of our model (illustrated in Section 1), agents belonging to the same category share the same features and, thus, their performances in executing  $\tau$  are roughly similar; this implies that if an agent  $a_j \in C_l$  is able (resp., not able) to execute  $\tau$ , then we expect that any other agent  $a_q \in C_l$  is also able (resp., not able) to execute  $\tau$ .

In the following, we suppose that agents in  $C_l$  are able to execute  $\tau$ , i.e. in compliance with notation introduced in Section 1, we assume that  $\psi(C_l, \tau) = \text{True}$ . In contrast, if  $\psi(C_l, \tau) = \text{False}$ , it does not make sense to estimate  $\theta_l(\tau)$ .

The next step of our protocol consists of selecting one of the agents, i.e., the *trustee*, in  $C_l$  to which delegate  $\tau$ ; to this purpose, we could select, uniformly at random, one of the agents in  $C_l$ , as illustrated in [7]. However, agents are individual entities and, thus, slight differences in their features exist. Because of these differences, an agent (say  $a_j$ ) may have better (resp., worse) performance than another agent (say,  $a_q$ ) in executing  $\tau$ .

In the light of the reasoning above, the true quality  $\theta_l(\tau)$  quantifies the *expected performance of an arbitrary trustee in  $C_l$  in the execution of  $\tau$* .

We assume that  $\theta_l(\tau)$  ranges in  $(-\infty, +\infty)$ : positive (resp., negative) values of  $\theta_l(\tau)$  are an

indicator of good (resp., bad) performances.

The first step to compute  $\theta_l(\tau)$  consists of modelling the performance  $f_j(\tau)$  of an arbitrary agent  $a_j \in C_l$  in executing  $\tau$ . To capture uncertainty in the performance of  $a_j$ , we represent  $f_j(\tau)$  as a Gaussian random variable with mean  $\mu_j$  and variance  $\sigma_j^2$ .

The assumption that all of the agents in the same category should reach the same performance implies that  $\mu_j = \theta_l(\tau)$  for each agent  $a_j \in C_l$ .

The variance  $\sigma_j^2$  controls the amount of variability in the performances of the agent  $a_j$ : large (resp., small) values in  $\sigma_j^2$  generate significant (resp., irrelevant) deviations from  $\theta_l(\tau)$ . In this paper we considered two options for  $\sigma_j^2$ , namely:

1. *Fixed Variance Model*: we suppose that  $\sigma_j^2 = \sigma^2$  for each  $a_j \in C_l$ .
2. *Random Variance Model*: we suppose that  $\sigma_j^2$  is a uniform random variable in the interval  $[\alpha, \beta]$ .

Based on these premises, the procedure to estimate  $\theta_l(\tau)$  is *iterative* and, at the  $k$ -th iteration it works as follows:

- a) We select, uniformly at random, an agent, say  $a_j$  from  $C_l$
- b) We sample the performance  $\hat{f}_j(k) \sim f_j(\tau)$  of  $a_j$

Steps a) and b) are repeated  $N$  times, being  $N$  the number of agents we need to sample before making a decision. In addition, in Step a), agents are sampled with replacement, i.e., an agent could be selected more than once. The algorithm outputs the average value of sampled performances, i.e.:

$$\hat{\theta}_l(\tau) = \frac{1}{N} \sum_{k=1}^N \hat{f}_j(k) \quad (1)$$

Our algorithm actually converges to the true value  $\theta_l(\tau)$  as stated in the following theorem:

Let  $N$  be the number of agents queried by our algorithm and let  $\hat{\theta}_l(\tau)$  be the estimation of the true quality  $\theta_l(\tau)$  the algorithm returns after  $N$  rounds. We have that in both the fixed variance and random variance models  $\hat{\theta}_l(\tau)$  converges to  $\theta_l(\tau)$  at a rate of convergence of  $\frac{1}{\sqrt{N}}$ .

Let us first analyze the individual agent performances  $f_j(\tau)$  and we are interested in computing the mean and variance of  $f_j(\tau)$ . If we opt for the Fixed Variance Model, then  $f_j(\tau)$  is a Gaussian random variable with mean  $\theta_l(\tau)$  and the variance is equal to a constant value  $\sigma^2$ . In contrast, if we are in the Random Variance Model, then the estimation of the mean and the variance of  $f_j(\tau)$  can be obtained by law of total mean and the law of total variance [17], which state that for two arbitrary random variables  $X$  and  $Y$ , the following identities hold true:

$$E(X) = E(E(X | Y)) \quad (2)$$

$$\text{Var}(Y) = E(\text{Var}(Y | X)) + \text{Var}(E(Y | X)) \quad (3)$$

We apply Equations 2 and 3 to  $X = f_j(\tau)$  and  $Y = \sigma$ ; if we condition on  $\sigma = \bar{\sigma}$ , then  $f_j(\tau)$  is a Gaussian random variable with mean equal to  $\theta_l(\tau)$  and variance equal to  $\bar{\sigma}$  and therefore:

$$\mathbb{E}(f_j(\tau)) = \mathbb{E}(\mathbb{E}(f_j(\tau) \mid \sigma = \bar{\sigma})) = \mathbb{E}(\theta_l(\tau)) = \theta_l(\tau)$$

In addition,

$$\mathbb{E}(\text{Var}(f_j(\tau) \mid \sigma = \bar{\sigma})) = \mathbb{E}(\sigma) = \frac{\alpha + \beta}{2}$$

and

$$\text{Var}(\mathbb{E}(f_j(\tau) \mid \sigma = \bar{\sigma})) = \text{Var}(\mathbb{E}(\theta_l(\tau))) = \text{Var}(\theta_l(\tau)) = 0$$

which jointly imply

$$\text{Var}(f_j(\tau)) = \frac{\alpha + \beta}{2}$$

As a consequence, independently of the agent  $a_j$ , we have that the agent performances  $f_j(\tau)$  have the same distribution which we denote as  $f(\tau)$ . Therefore, in both the Fixed Variance Model and Random Variance Model, the algorithm selects a random sample of agents  $Z_1, Z_2, \dots, Z_N$  of size  $N$  in which, for each  $k$  such that  $1 \leq k \leq N$ ,  $Z_k$  is the average performance of the agent selected at the  $k$ -th iteration and it is distributed as  $f(\tau)$ . The algorithm calculates:

$$S_N = \frac{Z_1 + Z_2 + \dots + Z_n}{N} \quad (4)$$

Because of the Central Limit Theorem [17], the distribution of  $S_N$  gets closer and closer to a Gaussian distribution with mean  $\theta_l(\tau)$  as  $N \rightarrow +\infty$  with a rate of convergence in the order of  $\frac{1}{\sqrt{N}}$  and this ends the proof.

## 4. Experimental Analysis

We designed our experiments to answer two main research questions, namely:

- RQ<sub>1</sub>** What are the benefits arising from the introduction of categories in the selection of a trustee against, for instance, a pure random search or a direct-experience based strategy?  
**RQ<sub>2</sub>** How quickly our algorithm to estimate  $\theta_l(\tau)$  converges?

In what follows, we first describe a reference scenario in which our task consists of recruiting a chef from a database of applicants (see Section 4.1). Then, in Sections 4.2 and 4.3, we provide an answer to **RQ<sub>1</sub>** and **RQ<sub>2</sub>**.

### 4.1. The reference scenario

We assume that features associated with our task are as follows: (i) *Culinary Education*, measured as the (overall) number of hours spent in training courses with qualified chef trainers, (ii) *Expertise*, i.e., the number of years of professional experience, (iii) *Language Skills*, defined as the number of foreign languages in which the applicant is proficient, (iv) *Culture*, measured on

**Table 1**

Some tasks associated with the recruitment of a professional chef and their requirements

Task	Culinary Education	Expertise	Language Skills	Culture	Creativity
$\tau_1$	150	5	2	6	7
$\tau_2$	200	4	2	6	6
$\tau_3$	300	4	2	7	6

**Table 2**Some features associated with categories  $C_1$ - $C_5$ 


Category ID.	Culinary Education	Expertise	Language Skills	Culture	Creativity
$C_1$	250	5	2	8	7
$C_2$	250	3	2	6	5
$C_3$	300	3	1	4	5
$C_4$	100	6	1	4	5
$C_5$	400	3	1	4	5

a scale from 0 (worse) to 10 (best) and which is understood as the ability of preparing different kind dishes (e.g. fish, meat, vegetarian and so on) in different styles (e.g. Indian, Thai or Italian) and ( $v$ ) *Creativity*, measured on a scale from 0 (worse) to 10 (best). The list of features is, of course, non-exhaustive. We suppose that each feature is associated with a plausible range: for instance, in Table 1, we consider three potential tasks and the corresponding requirements.

In the following, due to space limitations, we concentrate only on the task  $\tau_1$  and we suppose that five categories exist, namely: *Professional Chefs* -  $C_1$ , who are trained to master culinary art. Members in  $C_1$  are able to provide creative innovation in menu, preparation and presentation, *Vegan Chefs* -  $C_2$ , specialized in the preparation of plant-based dishes, *Pastry Chefs* -  $C_3$ , who are capable of creating chocolates and pastries, *Roast Chefs* -  $C_4$ , who have expertise in preparing roasted/braised meats and *Fish Chefs* -  $C_5$ , who are mainly specialized in the preparation of dish fishes. Each category consists of 100 agents and, thus, the overall number of agents involved in our experiments is 500. Features associated with categories  $C_1$ - $C_5$  are reported in Table 2.

In our scenario, *only agents in  $C_1$  are able to fulfill  $\tau_1$* ; agents in other categories are, for different reasons, unable to execute  $\tau_1$ : for instance, the expertise of agents in categories  $C_2$ ,  $C_3$  and  $C_5$  is not sufficient while agents forming categories  $C_3$ - $C_5$  correspond to applicants with a high level of specialization in the preparation of some specific kind of dishes (e.g., fish-based dishes) but they are not sufficiently skilled in the preparation of other type of foods and, thus, agents in these categories showcase an insufficient level of culture.

To simplify discussion we suppose that, through a proper normalization, the performance  $f(\tau_1)$  (see Section 3) of an individual agent as well as the true quality  $\theta_l(\tau_1)$  of a category  $C_l$  (for  $l = 1 \dots 5$ ) range from 0 to 1. Here, the best performance of an agent can provide (resp., the highest true quality of a category) is 1.




**Figure 1:** Feedback  $f(\tau)$  provided by the trustee. Fixed Variance Model with  $\sigma = 0.05$

#### 4.2. A comparison of category-based search with random-based search and direct-experience search

In our first experiment, we compare three strategies to search for a trustee, namely: *a) Random-Based Search:* here, the trustor selects, uniformly at random, a trustee to execute  $\tau_1$ . The trustor measures the performance  $f(\tau_1)$  provided by the trustee in the execution of  $\tau_1$ . *b) Category-Based Search:* here, the trustor considers only agents in the most appropriate category (which in our reference scenario coincides with  $C_1$ ); as suggested in [7], the trustor selects, uniformly at random, one of the agents in  $C_1$  to act as trustee. Once again, the trustor measures the performance  $f(\tau_1)$  provided by the trustee in the execution of  $\tau_1$ . *c) Direct-Experience Search:* we suppose that the trustor consults up to  $B$  agents in the community, being  $B$  a fixed integer called *budget*. The trustor records the performance of each consulted agent but it *does not memorize* its category (it may be that the trustor is unable to perceive/understand the trustee’s category). At the end of this procedure, the trustor selects as trustee the agent providing the highest performance  $f(\tau_1)$  among all consulted agents. The Direct-Experience Search strategy can be regarded as an evolution of the Random-Based Search strategy in which the trustor learns from its past interactions it uses its knowledge to spot the trustee. Here, the budget  $B$  regulates the duration of the learning activity the trustor pursues.

In our experimental setting, we considered two values of  $B$ , namely  $B = 10$  and  $B = 30$  and we discuss only results in the Fixed Variance Model with  $\sigma = 0.05$  and  $\sigma = 0.15$ . We applied the Random-Based, the Category-Based and the Direct-Experience Search strategies 20 times; a sketch of the probability density function (pdf) of  $f(\tau_1)$  for each strategy is shown in Figure 1 and 2.

As expected, Category-Based Search performs consistently better than the Random-Search one. In addition, the standard deviation of  $f(\tau_1)$  in Category-Based Search is much smaller than that observed in the Random-Based strategy and such a behaviour depends on the different degree of matching of categories  $C_1$ - $C_5$  with the task  $\tau_1$ : in other words, if the trustee is in  $C_1$ , the performances it provides are constantly very good; in contrast, in Random-Based Search strategy, the measured performance may significantly fluctuate on the basis of the category to



**Figure 2:** Feedback  $f(\tau)$  provided by the trustee. Fixed Variance Model with  $\sigma = 0.15$


which the trustee belongs to and this explains oscillations in  $f(\tau_1)$ .

The analysis of the Direct-Experience Search strategy offers many interesting insights which are valid for both  $\sigma = 0.05$  and  $\sigma = 0.15$ . Firstly, notice that if  $B = 10$ , then the Direct-Experience Search strategy achieves significantly better performances than the Random-Based strategy, which indicates that an even short learning phase yields tangible benefits. If  $B$  increases, the trustor is able to see a larger number of agents before making its decision and, in particular, if  $B$  is sufficiently large, then the trustor might encounter the best performing agent  $a^*$  in the whole community. In this case, the Direct-Experience strategy would outperform the Category-Based Search strategy: in fact, in the Category-Based strategy, the trustor chooses, uniformly at random, one of the agents in  $C_1$  which provides a performance worse than (or equal to)  $a^*$ . In short, for large values of  $B$ , the Direct-Experience strategy achieves performances which are comparable and, in some cases, even better than those we would obtain in the Category-Based strategy, as shown in Figure 1 and 2. However, the budget  $B$  has the meaning of a *cost*, i.e., it is associated with the time the trustor has to wait before it chooses the trustee and, thus, in many practical scenarios, the trustor has to make its decisions as quick as possible.


It is also instructive to consider a further strategy, called *Mixed-Based Search*, which combines the Random-Based Search strategy with the Category-Based Search strategy.

In Mixed-Based Search, we assume the existence of a *warm-up* phase in which the trustor selects the trustee by means of the Random-Based Search strategy; unlike the Direct-Experience Search strategy, the trustor collects not only  $f(\tau_1)$  but it also records the category of the trustee. In this way, the trustor is able to identify (after, hopefully, a small number of steps) the category with the highest true quality, i.e.,  $C_1$ . From that point onward, the trustor switches to a Category-Based strategy and it selects only agents from  $C_1$ . From a practical standpoint, we suppose that a performance  $f(\tau_1) \geq 0.6$  is classified as an indicator of good performance (in short, *positive signal*); as soon as the trustor has collected 2 positive signals, it makes a decision on the best performing category and it switches to the Category-Based Search strategy.

We are interested at estimating, through simulations, the length  $\ell$  of the warm-up phase, i.e., the number of agents that the trustor has to contact before switching to the Category-Based search strategy. In Figure 3 and 4 we plot the pdf of  $\ell$ .



**Figure 3:** Probability Density Function of the Warm-up length ( $\ell$ ) in the Mixed-Search strategy. Fixed Variance Model with  $\sigma = 0.05$



**Figure 4:** Probability Density Function of the Warm-up length ( $\ell$ ) in the Mixed-Search strategy. Fixed Variance Model with  $\sigma = 0.15$

Here, the variance  $\sigma$  has a minor impact and we notice that, the pdf achieves its largest value at  $\ell \simeq 10$ , i.e., 10 iterations are generally sufficient to identify the best performing category.

### 4.3. The rate of convergence of our algorithm


We conclude our study by investigating how the Fixed Variance Model and the Random Variance model influence the rate at which our algorithm estimates the true quality  $\theta_l(\tau)$  of a category.

To make exposition of experimental outcomes simple, we suppose that  $\theta_l = 1$  (which models a scenario in which agents in  $C_1$  showcase an exceptionally high ability in executing  $\tau_1$ ).

We considered the Fixed Variance Model with  $\sigma \in \{0.05, 0.1, 0.15\}$  and the Random Variance Model in which  $\sigma$  is uniformly distributed in  $[0.01, 0.3]$ .

We investigated how  $\hat{\theta}_l(\tau)$  varied as function of the number  $N$  of queried agents; obtained results are reported in Figure 5.

The main conclusions we can draw from our experiment are as follows:



**Figure 5:** Variation of  $\hat{\theta}_l(\tau)$  as function of  $N$ .

1. Individual agent variability (modelled through the parameter  $\sigma$ ) greatly affects the rate at which  $\hat{\theta}_l(\tau)$  converges to  $\theta_l(\tau)$ . Specifically, Figure 5 suggests that less than 5 iterations are enough to guarantee that  $|\hat{\theta}_l(\tau) - \theta_l(\tau)| < 10^{-2}$  if  $\sigma = 0.05$ . In addition, as  $\sigma$  gets larger and larger, we highlight more and more fluctuations in  $\hat{\theta}_l(\tau)$ : as an example, if  $\sigma = 0.15$  (green line), we highlight the largest fluctuation in  $\hat{\theta}_l(\tau)$  and, at a visual inspection, at least  $N = 30$  queries are needed to achieve a significant reduction in  $|\hat{\theta}_l(\tau) - \theta_l(\tau)|$ .
2. An interesting case occurs in the Random Variance Model: in some iterations of the algorithm, agents with a small variability are selected (i.e., we would sample agents with  $\sigma \simeq 0.01$ ) while in other cases agent with a larger variability are selected (here  $\sigma \simeq 0.3$ ). Overall, agents with small variability fully balance agents with high variability and, thus, the algorithm converges to  $\theta_l(\tau)$  (red line) generally faster than the case  $\sigma = 0.1$  (orange line) and  $\sigma = 0.15$  (green line).

## 5. Conclusions

Although highly populated networks are a particularly useful environment for agents' collaboration, the very nature of these networks may represent a drawback for trust formation, given the lack of data for evaluating the huge number of possible partners. Many contributions in the literature [8, 18] showed that category-based evaluations and inferential processes represent a remarkable solution for trust assessment, since they allow agents to generalize from trust in individuals to trust in their category and vice versa, basing on their observable features. On that note, we cared about stressing the tight relationship between trust and the specific task, target of the trust itself. With the purpose of investigating the role of agents' categories, we considered a simulated scenario, testing in particular the performance of a category-based evaluation, with respect to a random-based search - which it is easily outperformed - and a direct-experience one, showing that, in case of little direct experience, categories grant a better result. Moreover, we proved that, if not available, it is possible to estimate the category's true quality  $\theta_l(\tau)$  in a

reasonably short amount of time. Future research will attempt to test these findings on a real data set.

## References

- [1] J. Yan, D. Wu, S. Sanyal, R. Wang, Trust-oriented partner selection in d2d cooperative communications, *IEEE Access* 5 (2017) 3444–3453.
- [2] F. Messina, G. Pappalardo, C. Santoro, D. Rosaci, G. M. Sarné, A multi-agent protocol for service level agreement negotiation in cloud federations, *International Journal of Grid and Utility Computing* 7 (2016) 101–112.
- [3] F. Messina, G. Pappalardo, D. Rosaci, G. M. Sarné, A trust-based, multi-agent architecture supporting inter-cloud vm migration in iaas federations, in: *International Conference on Internet and Distributed Computing Systems*, Springer, 2014, pp. 74–83.
- [4] A. Sapienza, R. Falcone, Evaluating agents’ trustworthiness within virtual societies in case of no direct experience, *Cognitive Systems Research* 64 (2020) 164–173.
- [5] S. Abar, G. K. Theodoropoulos, P. Lemarinier, G. M. O’Hare, Agent based modelling and simulation tools: A review of the state-of-art software, *Computer Science Review* 24 (2017) 13–33.
- [6] C. Castelfranchi, R. Falcone, *Trust theory: A socio-cognitive and computational model*, volume 18, John Wiley & Sons, 2010.
- [7] R. Falcone, A. Sapienza, Selecting trustworthy partners by the means of untrustworthy recommenders in digitally empowered societies, in: *Proc. of the International Conference on Practical Applications of Agents and Multi-Agent Systems*, Springer, Avila, Spain, 2019, pp. 55–65.
- [8] R. Falcone, A. Sapienza, C. Castelfranchi, The relevance of categories for trusting information sources, *ACM Transactions on Internet Technology (TOIT)* 15 (2015) 13.
- [9] R. Falcone, M. Piunti, M. Venanzi, C. Castelfranchi, From manifesta to krypta: The relevance of categories for trusting others, *ACM Transactions on Intelligent Systems and Technology (TIST)* 4 (2013) 27.
- [10] C.-W. Hang, Y. Wang, M. P. Singh, Operators for propagating trust and their evaluation in social networks, Technical Report, North Carolina State University. Dept. of Computer Science, 2008.
- [11] R. Guha, R. Kumar, P. Raghavan, A. Tomkins, Propagation of trust and distrust, in: *Proceedings of the 13th international conference on World Wide Web*, 2004, pp. 403–412.
- [12] M. Jamali, M. Ester, A matrix factorization technique with trust propagation for recommendation in social networks, in: *Proceedings of the fourth ACM conference on Recommender systems*, 2010, pp. 135–142.
- [13] W. L. Teacy, M. Luck, A. Rogers, N. R. Jennings, An efficient and versatile approach to trust and reputation using hierarchical bayesian modelling, *Artificial Intelligence* 193 (2012) 149–185.
- [14] H. Fang, J. Zhang, M. Sensoy, N. M. Thalmann, A generalized stereotypical trust model, in: *2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications*, IEEE, 2012, pp. 698–705.

- [15] R. M. Kramer, Collective trust within organizations: Conceptual foundations and empirical insights, *Corporate Reputation Review* 13 (2010) 82–97.
- [16] B. D. Adams, R. D. Webb, Trust in small military teams, in: 7th international command and control technology symposium, 2002, pp. 1–20.
- [17] D. Bertsekas, J. Tsitsiklis, *Introduction to probability*, Athena Scientific, 2008.
- [18] R. Falcone, A. Sapienza, C. Castelfranchi, Trusting information sources through their categories, in: *International Conference on Practical Applications of Agents and Multi-Agent Systems*, Springer, 2015, pp. 80–92.