

# Exploring the Dynamics of Learned, Pre-existing, and Partial Knowledge in Dependence Networks within Multi-Agent Systems

Alessandro Sapienza<sup>1,\*</sup>, Rino Falcone<sup>1</sup>

<sup>1</sup>*Institute of Cognitive Sciences and Technologies, National Research Council of Italy (ISTC-CNR), 00185 Rome, Italy*

## Abstract

Dependence networks play a pivotal role in shaping interactions among agents in multi-agent systems, influencing collaboration, resource management, and overall system performance. In this study, we explore the effects of limited knowledge on the dynamics of dependence networks within such systems. Through a series of experiments, we investigate how agents with restricted awareness of their environment and potential partners navigate within dependence networks and the implications of their limitations on system outcomes. Our findings reveal that agents with limited knowledge face significant challenges, including reduced collaboration opportunities, suboptimal resource utilization, and limited goal achievement. Moreover, we demonstrate the tangible costs associated with acquiring trustworthiness knowledge and the critical role it plays in optimizing agent interactions. Overall, our study sheds light on the intricate interplay between knowledge, dependence, and trust in multi-agent systems, offering insights into strategies for enhancing system efficiency in real-world applications.

## Keywords

dependence networks, trust, multi-agent systems, social-simulation

## 1. Introduction

Social dynamics play a particularly important role in social sciences, since they shed light on how relationships develop, evolve, and shape human behavior, providing insights crucial for a deeper understanding of society [1, 2, 3]. This is why the social sciences have devoted so much interest to the study of these dynamics, in both a theoretical and practical perspective.

Moreover, the dynamics of social bonds play a pivotal role in individual and community decision-making processes. People are often influenced and conditioned by the beliefs and actions of their social circles [4]. Understanding how these dynamics affect behavior is vital for making more accurate predictions and planning effective social and economic interventions.

This is precisely the context in which our research fits. Specifically, we are interested in investigating dependence networks within hybrid societies comprising both human and artificial agents, focusing on the social dynamics they give rise to, also in relation to the concept of trust.

Specifically, this work builds upon the results obtained in previous studies [5, 6], to which we refer for a more comprehensive formalization of the role of dependence. In particular, as

---

WOA 2024: 25th Workshop "From Objects to Agents", July 8–10, 2024, Forte di Bard (AO), Italy

\*Corresponding author.

✉ alessandro.sapienza@istc.cnr.it (A. Sapienza); rino.falcone@istc.cnr.it (R. Falcone)

🆔 0000-0001-9360-779X (A. Sapienza); 0000-0002-0300-458X (R. Falcone)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

preliminary works, some limitations were highlighted. Firstly, preliminary studies investigated dependence relationships within a limited network of agents: in a world where agents are closely interrelated for the accomplishment of the tasks they need, having few partners available is certainly a disadvantage. This point was addressed by significantly increasing the number of available partners. In the scenarios considered here, we introduce a total of 30 agents, notably more than the 6 in the previous works.

Secondly, there was no effective reinforcement or punitive mechanism for partner selection. In other words, there was no penalty for agents interacting with untrustworthy partners. In such a case, adopting a restrictive partner selection strategy became counterproductive because, without incurring any cost for incorrect interactions, the best strategy was to rely even on unreliable partners.

This mechanism was addressed by introducing a monetary capital that agents must use to manage their interactions.

Specifically, within this article, we aim to investigate the role of knowledge in relation to dependence, addressing the following research questions:

1. How does the necessity to acquire knowledge on others' trustworthiness impact? Since trust is closely related to dependence, it is crucial to accurately evaluate partners to decide whom to rely on. Without this knowledge, there is a "cost" to acquiring it.
2. What happens if agents possess partial and limited knowledge of the world? How does this aspect affect their performance and that of other agents in the network?


Indeed, the importance of investigating agents operating with restricted knowledge of their environment and potential partners is evidenced by the great interest in the literature on this subject[7, 8, 9, 10, 11, 12, 13].

We address the two research questions by introducing a multi-agent simulation and modeling a network of interconnected agents that need to carry out a series of tasks.

## 2. Practical Formulation of the Model

In this section, we focus on describing the framework we have developed. We chose to implement the problem of dependence networks within the block world framework [14, 15, 16]. This domain was chosen due to its characteristics, which make the block world an interesting study context. Indeed, despite its conceptual simplicity, the block world offers sufficient complexity to model a wide range of scenarios and behaviors. Precisely because of this nature, it allows for the effective exploration of various aspects such as planning, learning, and social cooperation. The specific world we consider comprises a table and several blocks with different characteristics (shape, color, weight). Within the simulation context, agents must move specific blocks on the table to achieve their goals. During the execution of their tasks, they may need to collaborate with each other. As we will see, from this perspective, different strategies enable agents to achieve more or less valid performances.

Figure 1 shows how the simulated world appears at the beginning of the simulation.



**Figure 1:** Example of the simulated world

## 2.1. The blocks

The blocks are characterized by different shapes (cylinders, cones, cubes, spheres, stars), colors (red, blue, green, yellow, white), and weights (light, medium-light, medium, medium-heavy, heavy). Overall, 125 blocks are present in the world. Moreover, blocks can have an owner, i.e. a single agent who is authorized to change the status of the block in terms of position or ownership. Initially, all blocks are out of the table. Some of them are assigned to an agent from the beginning, while others are free and can be claimed by the agents.

Indeed, blocks are a limited resource. They are abundant at the beginning but get consumed for the agents' goals. One type of block will be more valuable the rarer it is in the world. We modeled this phenomenon by introducing a mechanism that takes into account the rarity of blocks in agents' transactions: a block will have a minimum value of 1 unit of capital and will reach a maximum value, set in the simulation, when it remains the only available block of that type.

## 2.2. The agents

Agents act in the world in order to realize their goals, i.e. to place specific blocks on the table. Specifically, each agent is defined in terms of:

- A goal: the type of block the agent needs to move on the table.
- A given competence, which defines how capable an agent is in performing certain tasks.
- A category of membership: we considered two categories: human or artificial agents. The category influences the characteristics of the agent. Specifically, we assume that humans

can move cylinders and cones, while robots can move cubes and spheres. Both types of agents are able to move the stars.

- Resources (blocks): initially, each agent possesses a given number of blocks. As the simulation progresses, this number changes because the agent can yield some blocks or acquire others.
- Monetary capital: it is the initial capital that citizens have available to manage interactions with other agents. This capital decreases when the agent delegates a task, while it increases when they perform tasks for other agents.
- Beliefs: The entire perception and reasoning of the agents are based on beliefs, whether they are about themselves, the world, or others. Therefore, there is a strong influence of their personal interpretation of reality. Certainly, these beliefs can vary in accuracy or even be absent.
- A  $\sigma$  threshold, which determines how trustworthy its potential partners must be to consider the dependence with them usable. Such a threshold value, specific to each agent, has the purpose of verifying that the partner is capable of performing certain actions. Of course, there remains a certain probability of error.

Agents must collaborate to achieve their goals, considering their subjective dependence networks, i.e. namely their personal perception of dependencies on others and of others on themselves. Agents are acquainted with all other agents and blocks in the world.

Introducing categories in this framework allows us, on the one hand, to differentiate the characteristics of the agents, such as their manipulation and action capabilities in the world. On the other hand, this enables us to introduce and model processes of inferential reasoning [17, 18, 19, 20].

### 2.3. Agents' trust and trustworthiness

In this study, we refer to the concept of trust as modeled in [21]. Trust is considered by the agents for the selection of dependencies, serving as a mechanism to decide whether to interact with one partner rather than another.

Within the simulation context, we assume the absence of malicious agents; hence, we choose not to consider the influence of motivational aspects on the determination of an agent's trustworthiness. For the sake of completeness, it is worth underlining that an agent might have conflicting motivations regarding a task: for instance, it may not want to give up a block of its interest as it will be needed for the completion of its sub-goals. However, this does not imply malicious intent. In such a case, the agent will simply decline the proposed task.

Therefore, we characterize agent trustworthiness in terms of *competence*, i.e., how effectively they can accomplish tasks in the world. Competence is defined as a real value within the range [0,1], where 0 implies a total inability to act, while 1 signifies a guaranteed success.

In the simulated world, we consider three types of tasks:

1. Acquisition of a block.
2. Repositioning of a block.

Since we have no interest in differentiating the values of competence for these tasks, for computational simplicity, we assume that an agent's trustworthiness is the same for each of

them. We would like to point out that this is not necessarily true in reality. Indeed, skills on different tasks usually tend to differ. Nevertheless, considering such a difference would have no practical impact within our scenario.

An agent is considered capable of achieving a task if it has a probability greater than a given threshold  $\sigma$  of accomplishing it. Such a probability is assessed through its trustworthiness evaluation. As mentioned earlier, agents possess a trustworthiness. This is an intrinsic characteristic of the agent that determines its task execution capability. As such, it cannot be accessed directly, not even by the agent itself, but it can only be estimated. To estimate the trustworthiness of agents, we consider a computational model based on the Beta distribution. The Beta distribution is commonly employed in the analysis of agent trustworthiness [22, 23, 24, 25], especially when it comes to modeling and estimating success or failure probabilities in complex situations. The Beta distribution is defined by two parameters, denoted as  $\alpha$  and  $\beta$ . As described in Equations 1 and 2, they depend on the estimation of the number of observed successes  $n\_successes_{a_x}$  and failures  $n\_failures_{a_x}$  of the agent  $a_x$ .

$$\alpha_{a_x} = n\_successes_{a_x} + 1 \quad (1)$$

$$\beta_{a_x} = n\_failures_{a_x} + 1 \quad (2)$$

In this context, the expected value of the distribution, representing the estimation of the average trustworthiness  $Trustworthiness_{a_x}$  of an agent  $a_x$ , is given by Equation 3:

$$Trustworthiness_{a_x} = \frac{\alpha_{a_x}}{\alpha_{a_x} + \beta_{a_x}} \quad (3)$$

When an agent  $a_x$  identifies its lack of power to perform a task  $\tau$  and the consequent presence of such power in another agent  $a_y$ , a dependence is explicitly manifested. In this case, if agent  $a_x$  sufficiently trusts agent  $a_y$ , i.e., the assessment that  $a_x$  has made of the trustworthiness of  $a_y$  exceeds the trust threshold  $\sigma$  of  $a_x$ , then  $a_x$  will contact  $a_y$  for the execution of  $\tau$ . In our framework, in the presence of multiple potential trustees, the choice is made randomly among those selected as sufficiently trustworthy.

Note that lack of power may either mean a total impossibility to perform the task (for example, the agent needs to move a certain type of block but does not have the physical characteristics to do so), or it could also mean that the agent evaluates itself as not sufficiently competent, based on its internal threshold  $\sigma$ , to execute  $\tau$ . In this case, the trustor agent will first attempt to delegate the task to someone else. In the absence of potential partners, it will still try to execute the task, as there is no cost associated with doing so.

## 2.4. Beliefs

Beliefs represent the perceptions and knowledge that agents possess about the state of the environment, about other agents, and about the relationships between them. These beliefs influence the decisions and actions of the agents and, consequently, guide the overall evolution of the simulation. Indeed, their fundamental role becomes even more crucial if beliefs on dependence networks are also considered.

In our framework, agents possess beliefs about:

- their own goals;
- their own abilities;
- the blocks that exist in the world;
- who the owners of the blocks are;
- the other agents that exist in the world;
- the goals of the other agents;
- the abilities of the other agents;
- the plans of the other agents (they know which ones they possess, not how these plans are articulated);
- dependencies on actions;
- dependencies on resources (blocks)
- dependencies on plans.

In this contribution, beliefs play a particularly significant role as we investigate what happens in the presence of limited knowledge. Indeed, in a network of agents, especially in relation to dependence networks, limited knowledge can significantly impact the beliefs that agents hold. Remarkably, agents with limited knowledge may fail to recognize all relevant dependencies within the network. This can result in missed opportunities for collaboration or reliance on less optimal partners. A similar problem concerns the use of resources. One must believe that a resource exists in the world and is available in order to utilize it.

### **3. Simulations**

Once the practical model is introduced, we proceed to examine its simulation implementations, aiming to investigate the effectiveness of using dependence networks. In particular, we also intend to explore the role of trust in this context. One of the limitations of previous experiments lies in the small number of agents present in the network, specifically six. Thus, in this case, we considered a total of 30 agents in the world. The goal is randomly assigned for all agents, thus avoiding any agent being advantaged or disadvantaged by the complexity of the goal assigned by the system. In each goal, two out of three components among weight, shape, and color are specified, leaving the third dimension free. When an agent completes a task, it is assigned another one, always randomly.

In the experiments, we need a strategy that allows us to identify which agents perform better and which perform worse. Specifically, we will assign a point each time an agent successfully completes a goal. Additionally, another metric of fundamental importance is the capital of the agents. On one hand, this dimension is crucial to account for how well the implemented delegation strategy optimizes the success/failure ratio of delegated tasks. Furthermore, capital is a dimension that also models how reliable an agent is considered within the network.

### **4. Results**

We have considered 3 simulation scenarios in which to assess the effectiveness of using dependence networks, specifically comparing their effectiveness with trust. The experiments

were conducted by using agent-based simulation, implementing what was described in the previous sections on the 3D version of the NetLogo platform [26]. The experiments aim at investigating the relationship between trust[27] and dependence networks[28, 29]. Specifically, we are interested in analyzing:

1. The importance of identifying an effective strategy in trusting. This aspect has been implemented through the trust threshold. We are interested in understanding whether, in relation to dependencies, it is better to use more restrictive or more open evaluation strategies.
2. The importance of being trustworthy within the network of agents.

The results we report pertain to a window of 30 interactions among the agents, which is sufficient to stabilize the interactions. Moreover, what we considered are the results averaged over 1000 simulations, in such a way as to eliminate the variability introduced by the random effects on the individual runs.

#### 4.1. First simulation

In this first experiment, we evaluate how the process of acquiring knowledge about partner trustworthiness impacts on agent performance. To investigate this aspect, we compare two scenarios: one where agents have to estimate trust from scratch and another where agents, through a training phase, learn about this dimension beforehand.

Moreover, we will divide agents into two groups of equal size, thus considering the 15 most trustworthy and the 15 least trustworthy, and compare the performance of these groups.

The experiments were conducted using the following settings:

- duration: 30 time units;
- training phase: 30 time units;
- number of agents: 30;
- number of blocks: 125;
- maximum value of a block: 5;
- percentage of human agents: 50%;
- percentage of artificial agents: 50
- initial capital: 10 unit;
- 2 blocks per agent;
- agents' *competence* randomly assigned in the range [0,1];
- $\sigma$  threshold randomly assigned between 0.25, 0.5 and 0.75.

As Table 1 shows, undoubtedly, the most trustworthy agents achieve better results. This characteristic impacts both the final capital, as these agents receive on average more delegation requests and have higher incomes, and the average score because, as seen in Section 2.3, agents who consider themselves less trustworthy tend to delegate more than those who are more trustworthy.

We then proceed to analyze the trust strategies of the agents, with reference to the partner selection thresholds: 0.25, 0.5, and 0.75. From Table 2, it emerges that agents with a lower

**Table 1**

Effect of the training phase on the performance of the most and least trustworthy agents

Training	Type	Average score	Final Capital	Income	Expenses	Total score
no	least trustworthy	2.50	7.50	5.58	-8.09	9.99
yes	least trustworthy	2.50	6.56	4.19	-7.63	9.06
no	most trustworthy	3.48	12.56	10.59	-8.03	16.03
yes	most trustworthy	3.35	13.51	10.43	-6.92	16.86

**Table 2**

Effect of the training phase on the performance as the trust threshold varies

Training	Type	Average score	Final Capital	Income	Expenses	Total score
no	$\sigma = 0.25$	3.41	7.63	8.67	-11.04	11.04
yes	$\sigma = 0.25$	3.25	7.76	7.69	-9.94	11.01
no	$\sigma = 0.5$	3.22	9.20	7.97	-8.77	12.42
yes	$\sigma = 0.5$	3.03	9.22	7.05	-7.84	12.25
no	$\sigma = 0.75$	2.31	13.17	7.54	-4.37	15.48
yes	$\sigma = 0.75$	2.49	13.03	8.09	-4.07	15.51

threshold generally manage to achieve better performance in terms of average score. However, this happens strongly at the expense of the final capital. Despite having a higher score of 32.3% compared to agents with a threshold of 0.75, these end up with a lower capital of 72.6%. On the other hand, using a threshold of 0.75 allows for maintaining a higher capital. This occurs because in this way agents delegate more cautiously: they manage to delegate less, but when they do, they have a higher likelihood of success.

Regarding the effect of training, we can observe that, in general, the average scores are homogenized. Furthermore, examining the ratio of how many tasks are delegated by the most trustworthy agents compared to the least trustworthy ones, this value stands at 0.82 without training and 0.69 when training is introduced. Additionally, the ratio of how many tasks are delegated to the most trustworthy agents compared to the least trustworthy ones stands at 2.20 without training and 3.03 after. Thus, in general, less trustworthy agents delegate more, and more trustworthy agents receive more delegation requests. The introduction of training further polarizes these effects.

In practice, the presence of a training phase implies that there is already knowledge available within the network, and this can be effectively utilized. For example, agents who consider themselves less trustworthy will directly turn to others for the execution of their tasks. This, in turn, results in better performance of these agents, which leads to fewer resources available in the world for other agents.

In general, we observe that by introducing training, expenses are minimized. This is interesting because, in fact, the difference in terms of expenses represents the cost that agents should pay to learn about the trustworthiness of their potential partners.



## 4.2. Second simulation

In this second experiment, we are interested in investigating the effects of limited knowledge on the utilization of dependence networks. To explore the effects of this lack of knowledge, we introduced agents who have a limited view of the world. These agents are able to observe only a limited part of the world, which means knowing a restricted number of agents and having access to a limited amount of resources. In this sense, the knowledge about dependencies that can be exploited is reduced.

The experiments were conducted using the following settings:

- duration: 30 time units;
- number of agents: 30;
- percentage of agents with limited knowledge about the world: 16.7%, 33.3%, 50%;
- percentage of known world: 25%, 50%, 75%;
- number of blocks: 125;
- maximum value of a block: 5;
- percentage of human agents: 50%;
- percentage of artificial agents: 50
- initial capital: 10 unit;
- 2 blocks per agent;
- agents' *competence* randomly assigned in the range  $[0,1]$ ;
- $\sigma$  threshold randomly assigned between 0.25, 0.5 and 0.75.

Within Table 3 and Table 4, the Setting refers to the combination of how many agents with limited knowledge there are in the world and how limited their knowledge is.

From Table 3, it indeed emerges that knowledge limitations have a significant and strong impact on the agents' performance: generally, starting from very low average scores (between 0.85 and 0.89) when agents can observe only 25% of the world, their performance improves significantly as their knowledge increases, reaching a maximum (between 2.35 and 2.39) when agents have access to 75% of the world.

Conversely, Table 4 indicates the opposite trend for the remaining agents, who start from a situation of substantial advantage and gradually achieve fewer goals. This dynamic occurs because, in the presence of agents with severe knowledge limitations, the remaining agents can exploit their knowledge more effectively, utilizing more resources than others. By reducing the knowledge gap between these two categories of agents, the difference in terms of scores decreases considerably.

Thus, indeed, having limited knowledge severely restricts the effective utilization of dependencies. While this represents a significant disadvantage for these agents, the remaining agents can exploit this situation to their advantage.

## 5. Discussion and Conclusions

This work builds upon the findings of previous studies on dependence networks, focusing on the role of knowledge. In comparison to prior works, this article contributes by investigating the

**Table 3**

Performance of agents with limited knowledge about the world

Setting	Average score	Final Capital	Income	Expenses	Total score
16.7% - 25%	0.89	12.46	4.65	-2.19	13.35
16.7% - 50%	1.81	11.11	5.81	-4.71	12.92
16.7% - 75%	2.39	10.39	6.68	-6.29	12.79
33.3% - 25%	0.86	12.25	4.27	-2.02	13.12
33.3% - 50%	1.81	11.01	5.48	-4.47	12.82
33.3% - 75%	2.37	10.46	6.48	-6.02	12.83
50% - 25%	0.85	11.79	3.72	-1.93	12.65
50% - 50%	1.81	10.87	5.10	-4.23	12.68
50% - 75%	2.35	10.33	6.17	-5.84	12.68

**Table 4**

Performance of agents with complete knowledge about the world, when in presence of agents with limited knowledge

Setting	Average score	Final Capital	Income	Expenses	Total score
16.7% - 25%	3.12	9.51	7.16	-7.65	12.63
16.7% - 50%	3.01	9.78	7.41	-7.63	12.79
16.7% - 75%	2.98	9.92	7.71	-7.78	12.90
33.3% - 25%	3.27	8.87	6.22	-7.35	12.14
33.3% - 50%	3.05	9.50	6.84	-7.34	12.55
33.3% - 75%	2.98	9.77	7.31	-7.54	12.75
50% - 25%	3.52	8.21	5.32	-7.12	11.73
50% - 50%	3.14	9.13	6.30	-7.17	12.27
50% - 75%	3.00	9.67	7.02	-7.35	12.68

difference between needing to learn and already knowing information, and between possessing it entirely or partially, within dependence networks in multi-agent systems. These problems are also explored in relation to trust, particularly focusing on the trust strategy and the relevance of being trustworthy.

Remarkably, in both experiments, albeit on different aspects, the value of knowledge clearly emerges. In its absence, there are significant limitations on the utilization of dependencies, i.e. agents cannot optimally leverage the network for collaborations, which negatively affects their performance.

In the first experiment, we demonstrate how the lack of knowledge about agents' trustworthiness results in an actual cost to acquire it. This cost is paid through the direct experience that agents have with the world and their partners. Initially, agents enter the environment with no pre-established understanding of who among their peers is trustworthy or not. As a result, they must engage in interactions and collaborations based on failure and error, which involves both successful and failed attempts at task delegation. Each failed attempt represents a tangible cost: it could mean wasted resources, missed opportunities, and time lost in reattempting the same task or seeking out new partners. Moreover, the process of learning from these interactions

incurs additional cognitive and computational costs as agents need to continually update their beliefs and strategies based on the outcomes of their engagements. This iterative process of refining trust assessments underscores a significant investment in terms of effort and efficiency. The cumulative effect of these costs highlights the critical importance of initial knowledge about trustworthiness in optimizing agent performance within the network.

Conversely, this situation represents an advantage for the remaining agents, who can effectively exploit resources in the world.

In the second experiment, we analyze what happens to agents with limited knowledge of the world. These agents face significant interaction limitations, as their restricted understanding of the environment hampers their ability to identify and connect with potential collaborators. As a consequence, they often miss out on crucial opportunities for cooperation and resource sharing, leading to suboptimal performance. Their lack of comprehensive knowledge means they cannot fully leverage the dependencies within the network, which is vital for achieving complex tasks that require coordinated efforts. This limited knowledge also results in inefficient use of the capital they possess, as they are unable to identify the most strategic ways to invest or utilize their resources. Conversely, this situation represents a distinct advantage for the remaining agents, who possess a broader and more accurate understanding of the world. These agents can capitalize on the limitations of their counterparts, effectively exploiting available resources and maximizing their performance within the network.

Another finding is that more relaxed delegation strategies regarding trust lead to higher performance, but at a significantly higher cost. It is noteworthy that this scenario considered a world with strictly finite resources. Indeed, in the absence of such a condition, this strategy would be further penalized.

This analysis has provided valuable insights into the role of knowledge within dependence networks. These aspects certainly deserve further investigation. Additionally, it would be interesting to explore the effects of erroneous beliefs on dependencies. Erroneous beliefs can significantly influence agents' behavior, since they can lead agents to take suboptimal or even harmful actions, as they are based on incorrect information about dependencies among themselves and other agents. For example, an agent may mistakenly believe to be dependent on another agent when in fact it is not, or vice versa. This could lead to erroneous decisions in resource management or collaborative task planning. Therefore, understanding how erroneous beliefs influence the dynamics of dependence networks is crucial for developing more effective strategies and improving the overall performance of the multi-agent system. We intend to investigate this aspect in future work.

## **Acknowledgments**

This work has been partially supported by the project FAIR - Future Artificial Intelligence Research (MIUR-PNRR).

## References

- [1] X. Jin, Y. Wang, Research on social network structure and public opinions dissemination of micro-blog based on complex network analysis, *Journal of Networks* 8 (2013) 1543.
- [2] S. Wasserman, K. Faust, *Social network analysis: Methods and applications* (1994).
- [3] D. J. Watts, *Six degrees: The science of a connected age*, WW Norton & Company, 2004.
- [4] R. B. Cialdini, N. J. Goldstein, Social influence: Compliance and conformity, *Annu. Rev. Psychol.* 55 (2004) 591–621.
- [5] R. Falcone, C. Castelfranchi, Grounding human machine interdependence through dependence and trust networks: Basic elements for extended sociality, *Frontiers in Physics* 10 (2022) 946095.
- [6] R. Falcone, A. Sapienza, The role of trust in dependence networks: A case study, *Information* 14 (2023) 652.
- [7] P. De Meo, F. Messina, D. Rosaci, G. M. Sarné, Recommending users in social networks by integrating local and global reputation, in: *Internet and Distributed Computing Systems: 7th International Conference, IDCS 2014, Calabria, Italy, September 22-24, 2014. Proceedings* 7, Springer, 2014, pp. 437–446.
- [8] R. Falcone, A. Sapienza, C. Castelfranchi, Recommendation of categories in an agents world: The role of (not) local communicative environments, in: *2015 13th Annual Conference on Privacy, Security and Trust (PST), IEEE, 2015*, pp. 7–13.
- [9] A. Herzig, A. Y. Ginel, Multi-agent abstract argumentation frameworks with incomplete knowledge of attacks, in: *Thirtieth International Joint Conference on Artificial Intelligence (IJCAI 2021), 2021*, pp. 1922–1928.
- [10] D. Rosaci, Cilos: Connectionist inductive learning and inter-ontology similarities for recommending information agents, *Information systems* 32 (2007) 793–825.
- [11] D. Rosaci, Web recommender agents with inductive learning capabilities, *Emergent Web Intelligence: Advanced Information Retrieval* (2010) 233–267.
- [12] A. Sapienza, R. Falcone, A bayesian computational model for trust on information sources., in: *WOA, 2016*, pp. 50–55.
- [13] Y. Zhang, M. M. Zavlanos, Cooperative multi-agent reinforcement learning with partial observations, *IEEE Transactions on Automatic Control* (2023).
- [14] S. V. Chenoweth, On the np-hardness of blocks world., in: *AAAI, 1991*, pp. 623–628.
- [15] J. Slaney, S. Thiébaux, Blocks world revisited, *Artificial Intelligence* 125 (2001) 119–153.
- [16] T. Winograd, *Five lectures on artificial intelligence*, Stanford University, Standord Artificial Intelligence Laboratory Stanford . . . , 1974.
- [17] C. Burnett, T. J. Norman, K. Sycara, Stereotypical trust and bias in dynamic multiagent systems, *ACM Transactions on Intelligent Systems and Technology (TIST)* 4 (2013) 1–22.
- [18] R. Falcone, A. Sapienza, C. Castelfranchi, Trusting information sources through their categories, in: *Advances in Practical Applications of Agents, Multi-Agent Systems, and Sustainability: The PAAMS Collection: 13th International Conference, PAAMS 2015, Salamanca, Spain, June 3-4, 2015, Proceedings* 13, Springer, 2015, pp. 80–92.
- [19] R. Falcone, A. Sapienza, F. Cantucci, C. Castelfranchi, To be trustworthy and to trust: The new frontier of intelligent systems, *Handbook of Human-Machine Systems* (2023) 213–223.

- [20] X. Liu, A. Datta, K. Rzadca, E.-P. Lim, Stereotrust: a group based personalized trust model, in: Proceedings of the 18th ACM conference on Information and knowledge management, 2009, pp. 7–16.
- [21] C. Castelfranchi, R. Falcone, Trust theory: A socio-cognitive and computational model, John Wiley & Sons, 2010.
- [22] D. Choi, S. Jin, Y. Lee, Y. Park, Personalized eigentrust with the beta distribution, ETRI journal 32 (2010) 348–350.
- [23] W. Fang, C. Zhang, Z. Shi, Q. Zhao, L. Shan, Btres: Beta-based trust and reputation evaluation system for wireless sensor networks, Journal of Network and Computer Applications 59 (2016) 88–94.
- [24] W. Fang, W. Zhang, Y. Yang, Y. Liu, W. Chen, A resilient trust management scheme for defending against reputation time-varying attacks based on beta distribution, Science China Information Sciences 60 (2017) 1–11.
- [25] V. Kanchana Devi, R. Ganesan, Trust-based selfish node detection mechanism using beta distribution in wireless sensor network, J. ICT Res. Appl 13 (2019) 79–92.
- [26] U. Wilensky, Netlogo. evanston, il: Center for connected learning and computer-based modeling, northwestern university, 1999.
- [27] A. Sapienza, F. Cantucci, R. Falcone, Modeling interaction in human–machine systems: A trust and trustworthiness approach, Automation 3 (2022) 242–257.
- [28] R. Conte, J. S. Sichman, Depnet: How to benefit from social dependence, Journal of Mathematical Sociology 20 (1995) 161–177.
- [29] S. Za, F. Marzo, M. De Marco, M. Cavallari, Agent based simulation of trust dynamics in dependence networks, in: Exploring Services Science: 6th International Conference, IESS 2015, Porto, Portugal, February 4-6, 2015, Proceedings 6, Springer, 2015, pp. 243–252.