

# Dependence Networks and Trust in Agents Societies: Insights and Practical Implications

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

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

## Abstract

In human, artificial, and hybrid societies, agents continuously interact to carry out complex tasks. In this sense, knowing and analyzing the dependencies that an agent has towards others (and vice versa) becomes a tool of fundamental importance. Making use of a structural theory, within this study, we investigate the role of dependence networks in the interaction between agents, with particular interest in the relationship between dependence and trust. The results of this research provide valuable insights into the use of dependence networks and how they influence collaboration and resource management. In particular, it is interesting to note that agents that exploit dependence, even when they have to interact with untrustworthy partners, obtain better performance in situations where resources are limited. On the contrary, in contexts where the use of dependency is limited, trust takes on a more relevant role. These conclusions emphasize the importance of dependence networks and their practical applications in areas such as robotics, resource management, and collaboration between humans and artificial agents.

## Keywords

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

## 1. Introduction

The study of social bonds is particularly relevant in the field of social sciences, both from a theoretical and empirical point of view [1, 2, 3, 4, 5]. This broad body of research offers a comprehensive insight into the different aspects and uses of social networks. Research in this area has highlighted the critical importance of the network structure in influencing various phenomena, such as the spread of ideas [6, 7, 8] and the propagation of consumption trends [9, 10]. The analysis of these networks provides us with a distinctive viewpoint on collective phenomena and social behaviors. Remarkably, social network studies offer a rich and diversified perspective on the intricacies of human relationships and the ways in which these relationships influence and guide a wide array of phenomena.

The primitives of these relational structures are of significant importance in cooperative, neutral, or conflictual interactions. These components constitute the basis of what we could define "extended sociality" [11], a concept that extends to both artificial and human agents. In order to achieve this type of interaction, even artificial agents need to possess some of the human social skills, such as the ability to recognize and understand the intentions, beliefs, and

---

WOA 2023: 24th Workshop From Objects to Agents, November 6–8, Rome, Italy

\*\* Corresponding author.

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

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



© 2023 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)

motivations of other individuals[12, 13], namely the "theory of mind"[14, 15, 16]. This skill not only pertains to the interpretation of objective data from reality but also involves anticipating the cognitive processes of other actors at play. In other words, the capacity to acquire knowledge about the convictions and desires of other agents is fundamental, as these pieces of information play a crucial role in social interactions. This concept is intriguing as it sheds light on the intricate nuances of both human and artificial interactions, emphasizing how an understanding of cognitive dynamics is pivotal in fostering effective and productive relationships within an increasingly interconnected reality.

We are interested in investigating the fundamentals of collaboration within a world inhabited by both artificial and human agents. More in detail, we dedicate our investigation to dependence networks[17, 18, 19] enriched by the beliefs that agents hold about the trustworthiness[20] of their counterparts. By employing a structural theory including agents' beliefs, we can tackle critical questions about the influence of each specific agent within a network, but we can also acquire a deeper understanding of the dynamic aspects of relational capital.

Starting from the theoretical framework introduced in [11], the aim of this contribution is to introduce a simulation-based implementation of dependence networks, In order to analyze its usage and impact. Specifically, our focus lies in conducting a comparative analysis between the concepts of dependence and trust, examining the roles they play in shaping interactions among agents. In our analysis, we refer to the concepts of agent and multi-agent systems, considering in particular the BDI model of the rational agent[21, 22, 23].

## 2. Agents and Social Dependence

We begin our discussion by introducing dependence networks, as considered in [11]. Within a shared "common world", agents move and act with the purpose of realizing their goals, endowed with limited power and control over the world and its components. Indeed, in order to pursue their goals, agents can collaborate with their peers, so that the latter carry out actions useful for their purposes. The action of an agent can support (positive interference) or hinder/compromise the goals of another agent (negative interference). Therefore, it is essential that agents also possess social powers, which give them the possibility of exploiting the actions of other agents to pursue their own objectives.

Indeed, it is important that the agents are able to correctly identify who they depend on for the realization of their tasks and evaluate when and with whom it is appropriate to interact. Not only. knowing which other agents each agent depends (or believes to depend) on is also critical.

In general, we will say an agent  $a_i$  possesses an *objective dependence* relationship with another agent  $a_j$  in regard to a task,  $\tau_k$ , if achieving  $\tau_k$  requires actions, plans, and/or resources owned by  $a_j$  and not available or less suitable for use by  $a_i$ . Such dependence occurs regardless of  $a_i$ 's and  $a_j$ 's awareness. If, in addition to this dependence relation, it also occurs that  $a_j$  holds an objective dependence relationship concerning  $a_i$ , we define this relationship as a *mutual dependence* relationship.

Objective dependence plays a key role in social interactions, the rationale for society. It fosters cooperation across various forms. However, the knowledge of objective dependence relationships is not sufficient to predict the arising (or absence) of relationships among agents.

To this end, it will also be necessary to consider the dependence relationships that agents know/believe or assume they have. In that case, we talk about *subjective dependence*.

It's important to highlight that when we introduce this concept related to the subjective perception of an agent about dependence relationships, we are considering what the agent believes and represents about its dependence on others. This introduces an additional dimension, that of beliefs. Instead of dealing with objective facts, we are delving into the personal mental representation of the agents and what they consider true in the world. Additionally, it is crucial to take into account what an agent believes about the dependencies of other agents within the network, in other words, how it perceives the dependencies of others. The set of objective dependencies, subjective dependencies, and the believed dependencies of others constitutes the foundation for initiating negotiation processes. It is important to note that while the objective level is based on factual information and is always present, regardless of whether agents are aware of it, the same cannot be said for subjective and believed dependencies. In these cases, we are dealing with beliefs, which may or may not be present in the minds of the agents. It is evident that what an agent subjectively believes may not necessarily align with the objective reality of the world. In this regard, different agents may hold varying perspectives on the same dependence relationships.

Dependence networks are highly dynamic and can change in a completely unpredictable manner as the context taken into consideration changes. Clearly, these change based on the goal that the agent intends to pursue. At the same time, they evolve according to the resources in the world and the individual skills of the agents. In addition, the variation of the agents that make up the society (open world) has a particularly important effect in determining the dependence relations. Simply becoming aware of possessing (or not) a certain capability (or what the agent/others believe about this capability) can alter possible scenarios. In fact, given that, to decide to pursue a goal, a cognitive agent must believe (at least with a certain degree of certainty) in having that power, then they do not truly possess that power if they are unaware of it. Conversely, the acquisition of power, and therefore autonomy and power over other agents, can also be simply due to the awareness of that power and not necessarily the acquisition of external resources or skills and competencies.

In essence, a complex and multifaceted framework is being outlined, wherein concepts like theory of mind come into play and prove essential for understanding these tools. Possessing the ability to analyze dependence networks is crucial for comprehending, predicting, and optimizing interactions with other agents. This is a fundamental tool that ensures those who employ it (appropriately) gain a distinct advantage over other agents.

### **3. Practical Formulation of the Model**

In this contribution, in order to investigate the complex dynamics of dependence networks, we realized an implementation of the block world[24, 25, 26]. We are therefore going to introduce the practical formulation of our model, which will then be implemented in the simulations. The world we consider is composed of a table and a number of blocks that have different characteristics (shape, color, weight). Within the context of the simulation, the agents aim to create one or more combinations/sequences of blocks on the table.

### 3.1. The blocks

As mentioned, the blocks are characterized by different shapes (cylinders, cone, cubes, spheres), colors (red, blue, green, yellow) and weights (light, heavy). Overall, 32 blocks are present in the world. These blocks can be on the table or out of the table. Also, 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.

The blocks are subject to physical constraints in the world, which are known by all agents:

- It is possible to have stacks of up to 3 elements.
- A light element can have at most 1 light element on top.
- A heavy element can have up to 2 elements of any kind on top. Clearly, due to the previous constraint, a combination of heavy-light-heavy blocks cannot be realized.
- Cones and spheres cannot have other blocks on top of them.

The blocks represent the basic structure with which it is possible to compose stacks, which are structures formed by stacking multiple blocks, arranged in a certain order. The idea behind this structuring is that the configurations of blocks abstractly represent the basic elements for the realization of goals. As we will show later, we intend to account for the fact that in reality there are simpler tasks, the correct execution of which involves the direct satisfaction of a sub-goal, and others which instead require a more complex construction, the outcome of which depends on several phases and on knowing the specific methods of achieving the goal.

### 3.2. The agents

Agents, human or artificial, act in the world to accomplish their goals. Specifically, each agent is defined in terms of:

- A goal: the specific combination of blocks the agent wants to achieve in the world. This configuration is in turn constituted by a series of more or less articulated sub-goals.
- A set of plans to achieve its goal/sub-goals (ranging from 0 to  $n$ ): if there are no plans, a dependence is created (towards someone) for retrieving a plan.
- A given competence, defining 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.
- Resources (blocks): initially, each agent possesses 5 blocks.
- Beliefs about themselves, the world, and others. The agents' entire perception of the world, elaboration, and planning are based on beliefs, therefore on their personal interpretation of reality. Certainly, these beliefs can be more or less correct or even missing.
- A  $\sigma$  threshold, that 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 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[27, 28, 29, 30]: knowing that an agent belongs to a certain category allows us to deduce specific characteristics, and thus - in our case - whether it is capable of achieving certain states of the world or not. This possibility is interesting in relation to the concept of dependence, since knowing by an agent A2 the category to which a given agent A1 belongs allows A2 to deduce whether it depends on A1 (knowing its own objective/plan) or if A1 depends on A2 (knowing the goal/plan of A1).

### 3.3. Goals and plans

Each agent has the goal of placing certain blocks on the table. These blocks can be stacked in a specific order or simply positioned on the table. In this regard, the agent's goal is divided into sub-goals, which can be:

- *Atomic*: for example, moving a single block onto the table. This kind of task is useful for modeling the presence in the world of simpler tasks that do not require complex planning skills and do not need to be performed in multiple steps.
- *Complex*: creating a stack, which is an ordered sequence of blocks. The stack introduces the need to perform a series of actions in a specific sequence (in fact a plan) to achieve a single sub-goal. Fulfilling only part of it is insufficient; all actions must be executed.

The stacks of blocks in the blocks world are meant to represent complex and challenging tasks for AI systems. The agent must be able to do complex planning on blocks to achieve specific goals, which cannot be done with a single action, but require more complex planning. Since these are complex tasks, in our framework we assume that in the absence of a specific plan that contains the implementation instructions, the agents are not able to achieve these sub-goals.

We establish that each agent requires 3 to 5 blocks to complete its goal. The agent's goal is considered fully satisfied when all sub-goals have been achieved. Conversely, it is considered partially satisfied if only some of the sub-goals have been realized. The goals are not shared, in the sense that the presence of a block or a sequence on the table satisfies the goal/sub-goal only of its specific possessor, but not of the other agents.

A plan is considered feasible for Agent A1 if:

1. there exists a set of unused blocks that, properly used, satisfies it;
2. there is someone (Agent A1 or an Agent A2 dependent on A1) who can potentially move these blocks;
3. the plan is physically achievable, i.e. the plan respects the block composition rules defined in Section 3.1.

### 3.4. Agents' trust and trustworthiness

In this study, we refer to the concept of trust as modeled in [31]. 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. Obtaining a plan.
2. Acquisition of a block.
3. 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 [32, 33, 34, 35], 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:

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

### 3.5. 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;
- their own plans;
- 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.

### 3.6. The Blackboard

Of course, only an observer outside the world and who has control over the simulation system is capable of knowing what the objective dependencies are, since these by their very nature are not directly knowable. Agents possess limited and partial knowledge of their dependencies. We suppose that agents are capable of autonomously determining from whom they depend, solely through the observation of the world, which is then processed through their beliefs. This assumption is reasonable, because they are able to determine both who is able to do what through the use of inferential processes on categories, and because the possession of resources in the world is public. However, this represents limited and partial knowledge which, in the event of an error, could also prove to be incorrect. Moreover, it is not guaranteed that agents can independently determine who depends on them. In order to establish a protocol of agent interaction based on these principles, we introduced in the world the presence of a blackboard. The blackboard represents the system that agents use for communication and for verifying dependencies. At the beginning of each simulation, each agent will declare the goal it intends to pursue. Then, whenever it needs to perform a task to continue its plan, it will check the blackboard.

Firstly, it will verify the existence of a mutual dependence. The agent will determine if among the agents with a trust rating higher than its internal threshold  $\sigma$ , there are requests that it can fulfill, and if there is someone who can satisfy its task. In the case of a negative outcome, the agent will simply post its request on the blackboard, awaiting another agent to select it for a

mutual dependence in the future. Conversely, if mutual dependencies are identified, the agent will proceed to initiate a negotiation phase. If both agents agree, each will proceed to fulfill the other's request.

### 3.7. Workflow

The agent starts each cycle by updating its beliefs. This is crucial because, in the previous cycle, the state of the world may have been altered by the agents' actions. For example, blocks' ownership may have changed, or blocks may have been moved.

In each cycle, the agents are limited to performing only one action: retrieving a plan from other agents, obtaining a block, moving a block.

1. At first, each agent evaluates whether it possesses at least one feasible plan to achieve its goals. If this condition does not hold, in this cycle it focuses on obtaining a plan. If the agent does not require external resources to achieve its goal, then it simply declares its goal on the blackboard, starting to execute the first task necessary to achieve it. Otherwise:
  - a) The agent establishes how to proceed in obtaining the required elements, following internal priority criteria.
  - b) It checks previous requests, updating its subjective view on the dependence network, and verifies in the blackboard if any of the agents having active dependence on it can provide the needed resource. Where this is the case, a mutual dependence is explicit, and the agent attempts a negotiation to formalize the exchange. If a partner is found, both requests are executed. If no partner is found, the agent declares on the blackboard its request and the goal it is pursuing. Then, it waits for a future mutual dependence.

It is worth emphasizing that within the complex system defined, dependence is not just a necessity but also a resource in itself. The fact that someone depends on me provides me with the opportunity to access the resources the other has to offer. Therefore, while being independent of everyone could be considered an advantage, the simulation world has been designed to make this possibility unlikely. Conversely, the fact that no one depends on us represents a significant competitive disadvantage.

Regarding the ranking criteria, the agent ranks the sub-goals it can take care of in the specific cycle, according to the following principles:


- *Abstraction level* of the sub-goal: it will prioritize less abstract sub-goals since, as the need for the sub-goal becomes more specific, the availability of blocks in the world that can satisfy this request becomes more restricted.
- *Reasoning about others' goals*: starting from knowledge about other agents' goals, an agent estimates which blocks it needs that are most likely to be used by other agents. It might even find it better to take possession of the final block of a stack, even if the base has not been constructed yet (typical market problem: offer/demand dynamics).



## 4. 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. As this is an initial experiment, we begin by considering the dynamics that emerge within a relatively small network of agents and in a controlled setting.

In the simulations, we considered a total of 6 agents in the world. We introduced a unique goal for all agents, thus avoiding any agent being advantaged or disadvantaged by the complexity of the goal assigned by the system. The goal involves creating a stack of two blue blocks and placing two lightweight blocks on the table, as in Figure 1.



**Figure 1:** In the figure, agent 5 completely realized its goal

In general, agents are capable of observing the world accurately, i.e., their subjective perception coincides with what objectively occurs. There is only one exception to this. In fact, it is possible for an agent to fail in performing an action to fulfill a sub-goal. However, this action could still satisfy another sub-goal of the same agent. In this sense, the action will be perceived as a failure by the executing agent, but other agents in the world will interpret it as correct (since, when agents start interacting in the world, every agent has perfect knowledge of the goals of every other agent). In more detail, agents reveal their potential requests when pursuing their own goals and, as a result, they become active on the blackboard.

## 4.1. Comparison metric

We need to define a metric to evaluate and compare the performance of agents. The metric should be designed to reward agents that are able to perform more complex actions. For instance, creating a stack of two blocks is more complex than simply placing two blocks on the table, both in terms of planning and because resources may run out in the meantime, given the stricter constraints on block composition. Therefore, we decided that:

1. Successfully placing a block of interest on the table is worth one point;
2. Successfully completing a stack earns an additional point;
3. Successfully completing all goals grants an additional point;

Overall, given the set goals for the agents, the maximum achievable score for an agent is 6. Naturally, considering that not all agents have correct plans available to achieve their goals, that agents make mistakes, and that resources in the world become engaged at some point, we expect the average score to decrease significantly.

## 5. 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 [36]. The experiments are designed in such a way as to increasingly disadvantage the utility of dependence: in the first simulation, it plays a significant role, while in the others, it is progressively limited. In the simulations, we consider a total of 6 agents, comprising 3 humans and 3 artificial agents. We decided to allocate 5 blocks to each agent right from the start. Indeed, we had initially considered the possibility of assigning a lower number of blocks. However, from the initial results, it became evident that such a setup significantly disadvantaged trust in favor of dependence.

As far as it concerns trust, we considered agents with three different  $\sigma$  threshold levels: 0.25, 0.5, 0.75. The first threshold identifies agents willing to interact with almost all available partners. The second one, on the other hand, pertains to agents who are willing to interact only with partners with above-average performance, thus, on average, interacting with only half of the available agents. The last group will make a strict selection of their partners, which, however, will significantly reduce the availability of agents to interact with. 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.

### 5.1. First simulation

In this first experiment, agents will identify mutual dependencies by means of the blackboard. Then, they will select their potential partners by filtering them based on their personal  $\sigma$  threshold. Finally, they will rank the remaining partners according to their trust evaluation and

**Table 1**  
First experiment results

$\sigma$	Average score	Completed tasks	Percentage of delegated task	Success rate of delegated tasks
0.25	1.39	13.31	0.32	0.53
0.5	1.35	13.0	0.18	0.68
0.75	1.27	12.64	0.06	0.76

contact partners based on the established order until a partner is found or until all available partners have refused.

The experiments were conducted using the following settings:

- number of human agents: 3; number of artificial agents: 3;
- 5 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 this experiment, we aim to investigate the value of trustworthiness within this type of network, comparing it with the effect of dependence.

From the results in Table 1, it emerges paradoxically that agents operating with a lower threshold of 0.25 achieve superior outcomes. The 0.25 agents exhibit a higher score by 9.3% compared to those with a threshold of 0.75. Even 0.5 agents achieve a 6.3% higher score, compared to 0.75 agents.

This phenomenon occurs due to the presence of a context in which errors do not lead to substantial or effective losses. In the event of task failure, agents do not incur any economic loss, nor lose the possibility of accomplishing the task in the future. Under these conditions, attempting reliance until the selected trustee succeeds proves to be the winning strategy, ensuring better results. Furthermore, in this experiment, it is noteworthy that all agents always approach their trustees in order of trust (requesting assistance from the most trusted to the least trusted). Therefore, having a high threshold introduces a disadvantage, as those with a low threshold can always attempt with other agents in case of rejection, whereas having a high threshold means forfeiting task execution if a partner with sufficient trustworthiness is not found.

In fact, upon observing other metrics, it is noted that although the success rate of delegated tasks is significantly lower (0.53 for the 0.25 threshold compared to 0.76 for the 0.75 threshold), the 0.25 agents manage to delegate tasks 5 times more than the 0.75 agents. This, in turn, results in these agents being able to complete an average of 0.67 more tasks.

## 5.2. Second Simulation

Given the significant weight of dependence in this context, in this second experiment, we attempted to identify certain conditions that can mitigate its effect. Specifically, agents are constrained to use dependence only once. Furthermore, we are examining what occurs when partner selection is randomized among those chosen via the threshold  $\sigma$ .

**Table 2**

Second experiment results

$\sigma$	Average score	Completed tasks	Percentage of delegated task	Success rate of delegated tasks
0.25	1.21	12.2	0.13	0.52
0.5	1.22	12.38	0.09	0.65
0.75	1.25	12.5	0.04	0.68

**Table 3**

Third experiment results

$\sigma$	Average score	Completed tasks	Percentage of delegated task	Success rate of delegated tasks
0.25	1.55	8.45	0.22	0.63
0.75	1.57	8.8	0.07	0.78

In this case, the trend identified in the previous scenario is reversed. Remarkably, as we can see in Table 2, the influence of trust becomes more significant compared to dependence. Agents with  $\sigma = 0.75$  achieve a superior average performance by 3.15% and are able to complete an average of 0.3 more tasks, compared to agents with  $\sigma = 0.25$ . It is worth noting that, although the effects are relatively small, they pertain to a single delegated task, thus the actual difference in performance remains limited.

### 5.3. Third Simulation

Even in this last case, as we did in the previous experiment, we consider experimental conditions that limit the use of the dependence to only once. Additionally, here too, partner selection occurs randomly among those surpassing the threshold  $\sigma$ . Furthermore, we attempt to further reinforce the utilization of trust by fixing the agents' performances at 0.4 (3 randomly chosen agents) and 0.9 (3 randomly chosen agents), so as to make a clear division between trustworthy and untrustworthy agents. Due to this simplification, we will consider only two threshold values, 0.25 and 0.75, as the threshold of 0.5 would yield results similar to that of 0.75.

Compared to the previous scenario, there are more agents considered trustworthy. This has an impact on several dimensions of the simulation. Firstly, as emerges from Table 3, more delegation is achieved. This is precisely a direct consequence of the fact that, with the same potential use of the dependence, in a network with more reliable agents it is easier to exploit the dependence itself. Additionally, the average number of completed tasks decreases in favor of higher scores, as fewer tasks are needed to achieve the same results. Once again, we can observe a slight increase in terms of both score and average performance. The difference is more pronounced when observing completed tasks. With a threshold of 0.75, an average of 0.35 more tasks are completed. The performance of delegated tasks also differs significantly: 0.63 versus 0.78.

## 6. Discussion and Conclusions

This research starts from the consideration that dependence networks are a fundamental resource for cognitive agents[11]. Indeed, in most real-world situations, agents need the collaboration of others for the success of their tasks [37, 38, 39]. This phenomenon occurs since agents have limited abilities and control in the world [40, 41]. Furthermore, for an agent to expand its scope of attainable tasks, it must possess the capability to accurately represent its own capabilities, those of others, and the resulting dependencies that emerge.

Indeed, the knowledge of dependencies enables agents to cooperate more efficiently[42]. When agents understand the impact of their actions on others and vice versa, they can decide better where to focus their efforts, optimizing efficiency and obtaining better results[11].

Awareness of dependence on other agents and, vice versa, of the dependencies that other agents have on them, allows cognitive agents to predict in advance possible reactions or responses from others. This can help them make better decisions and manage potential conflicts or problems.

At the same time, the ability to recognize dependencies allows agents to identify potential weaknesses or risks within the chains of actions necessary to achieve a goal. This awareness is also useful when it is necessary to develop risk mitigation strategies and emergency plans in the event of problems or failures. As an example of this, we can consider the case of supply chain management in disaster scenarios. This knowledge is of fundamental importance to implement resilience capabilities, allowing better planning of communication, coordination, and cooperation processes[43].

In light of these considerations, in this study, we decided to investigate the role of dependence within a society populated by both human and artificial agents. Regarding this aspect, the block world represented for us a particularly interesting scenario for our analysis. This is an analysis context commonly used in the field of artificial intelligence and cognitive science. The comprehension of dependence networks in this context can have important repercussions on sectors such as robotics[44], resource management[45], and collaboration[46, 47] between artificial agents.

Indeed, the results of our simulations provide practical confirmation of the importance of dependence networks. Most notably, it emerges that the effect of dependence is very significant in the interactions between agents.

Remarkably, results suggest that agents making greater use of dependence, even at the expense of the trustworthiness of their partners, manage to achieve better performances than those who prefer a more restrictive partner selection based on their estimated trustworthiness. This experiment allows us to verify that, in the specific situation at hand – a closely interconnected world with limited resources – in the presence of dependence on one or more agents, it is preferable to rely on one of them, rather than give up due to their lack of trustworthiness. This phenomenon happens because of the potential unavailability of an alternative way to carry out the task.

Subsequently, we introduced two further experiments, aimed at reducing agents' reliance on dependence. While this increased the influence of trust, as seen in both experiments where agents with a higher trust threshold performed better, this distinction did not turn out to have a substantial effect.

Although we expected a strong impact of dependency networks on network performance, we did not predict that this effect would conflict with the use of trust. It appears that using trust as a filter, which implies restricting interactions with potential partners, limits the effectiveness of dependence networks. This leads us to an additional conclusion: in a scenario defined by the environmental characteristics under consideration, it is more advantageous for agents to possess the necessary skills, abilities, and resources required by other agents, rather than focusing on being perceived as trustworthy. This is especially relevant when, as in our scenario, the risk of failure does not permanently compromise the achievement of one's objectives. In this specific context, there are no severe penalties for agents' erroneous choices. While opting for unreliable partners may lead to a higher rate of failed tasks, and time can be a valuable resource in a resource-constrained environment, the limited number of available agents significantly influences the outcomes. Because agents can only select partners from a restricted pool, filtering these partners based on their trustworthiness can potentially undermine success.

These first results have allowed us to obtain useful insights and interesting preliminary considerations. Starting from this, we reserve for the future the possibility of investigating simulated worlds composed of more extensive networks of agents, evaluating the effectiveness of dependency networks in these contexts.

## Acknowledgments

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

## References

- [1] Guare, J. (1990) *Six Degrees of Separation* (Vintage, New York).
- [2] Wasserman, S. & Faust, K. (1994) *Social Network Analysis* (Cambridge Univ.Press, Cambridge)
- [3] Watts, D. J. (1999). *Small Worlds* Princeton University Press. New Jersey.
- [4] Jin, X., & Wang, Y. (2013). Research on social network structure and public opinions dissemination of micro-blog based on complex network analysis. *Journal of Networks*, 8(7), 1543.
- [5] Safarnejad, L., Xu, Q., Ge, Y., Krishnan, S., Bagarvathi, A., & Chen, S. (2020). Contrasting misinformation and real-information dissemination network structures on social media during a health emergency. *American journal of public health*, 110(S3), S340-S347.
- [6] Acemoglu, D., & Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1, 3-49.
- [7] Das, A., Gollapudi, S., & Munagala, K. (2014, February). Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM international conference on Web search and data mining* (pp. 403-412).
- [8] Xia, H., Wang, H., & Xuan, Z. (2011). Opinion dynamics: A multidisciplinary review and perspective on future research. *International Journal of Knowledge and Systems Science (IJKSS)*, 2(4), 72-91.

- [9] Drescher, L. S., Grebitus, C., & Roosen, J. (2023). Exploring Food Consumption Trends on Twitter with Social Media Analytics: The Example of Veganuary. *EuroChoices*, 22(2), 45-52.
- [10] Yu, L., Zhao, Y., Tang, L., & Yang, Z. (2019). Online big data-driven oil consumption forecasting with Google trends. *International Journal of Forecasting*, 35(1), 213-223.
- [11] Falcone, R., & Castelfranchi, C. (2022). Grounding Human Machine Interdependence Through Dependence and Trust Networks: Basic Elements for Extended Sociality. *Frontiers in Physics*, 10, 946095.
- [12] Lapierre, M. A. (2015). Development and persuasion understanding: Predicting knowledge of persuasion/selling intent from children's theory of mind. *Journal of Communication*, 65(3), 423-442.
- [13] Sylwester, K., Lyons, M., Buchanan, C., Nettle, D., & Roberts, G. (2012). The role of theory of mind in assessing cooperative intentions. *Personality and Individual Differences*, 52(2), 113-117.
- [14] Carlson, S. M., Koenig, M. A., & Harms, M. B. (2013). Theory of mind. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(4), 391-402.
- [15] Leslie, A. M., Friedman, O., & German, T. P. (2004). Core mechanisms in 'theory of mind'. *Trends in cognitive sciences*, 8(12), 528-533.
- [16] Nichols, S., & Stich, S. (2003). How to read your own mind: A cognitive theory of self-consciousness. *Consciousness: New philosophical essays*, 157-200.
- [17] Castelfranchi, C. (2000, August). Founding agents' "autonomy" on dependence theory. In *ECAI* (Vol. 1, pp. 353-357).
- [18] Sichman, J. S., Conte, R., Demazeau, Y., & Castelfranchi, C. (1998). A social reasoning mechanism based on dependence networks. In *Proceedings of 11th European Conference on Artificial Intelligence* (pp. 416-420).
- [19] Sichman, J. S., & Conte, R. (2002, July). Multi-agent dependence by dependence graphs. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1* (pp. 483-490).
- [20] Falcone, R., Sapienza, A., Cantucci, F., & Castelfranchi, C. (2023). To Be Trustworthy and To Trust: The New Frontier of Intelligent Systems. *Handbook of Human-Machine Systems*, 213-223.
- [21] Michael Bratman. *Intentions, Plans and Practical Reason*. Harvard U. Press, 1987
- [22] Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial intelligence*, 42(2-3), 213-261.
- [23] Wooldridge, M. (2009). *An introduction to multiagent systems*. John wiley & sons.
- [24] Chenoweth, S. V. (1991, July). On the NP-Hardness of Blocks World. In *AAAI* (pp. 623-628).
- [25] Slaney, John, and Sylvie Thiébaux. "Blocks world revisited." *Artificial Intelligence* 125.1-2 (2001): 119-153.
- [26] Winograd, T. (1974). *Five lectures on artificial intelligence* (p. 0100). Stanford University, Standord Artificial Intelligence Laboratory.
- [27] Burnett, C., Norman, T. J., & Sycara, K. (2013). Stereotypical trust and bias in dynamic multiagent systems. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 4(2), 1-22.
- [28] Falcone, R., Sapienza, A., & Castelfranchi, C. (2015). 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* (pp. 80-92). Springer International Publishing.

- [29] Liu, X., Datta, A., Rządca, K., & Lim, E. P. (2009, November). Stereotrust: a group based personalized trust model. In *Proceedings of the 18th ACM conference on Information and knowledge management* (pp. 7-16).
- [30] Sapienza, A., & Falcone, R. (2020). Evaluating agents' trustworthiness within virtual societies in case of no direct experience. *Cognitive Systems Research*, 64, 164-173.
- [31] Castelfranchi, C., & Falcone, R. (2010). *Trust theory: A socio-cognitive and computational model*. John Wiley & Sons.
- [32] Choi, D., Jin, S., Lee, Y., & Park, Y. (2010). Personalized eigentrust with the beta distribution. *ETRI journal*, 32(2), 348-350.
- [33] Fang, W., Zhang, C., Shi, Z., Zhao, Q., & Shan, L. (2016). BTRES: Beta-based trust and reputation evaluation system for wireless sensor networks. *Journal of Network and Computer Applications*, 59, 88-94.
- [34] Fang, W., Zhang, W., Yang, Y., Liu, Y., & Chen, W. (2017). A resilient trust management scheme for defending against reputation time-varying attacks based on BETA distribution. *Science China Information Sciences*, 60, 1-11.
- [35] Kanchana Devi, V., & Ganesan, R. (2019). Trust-based selfish node detection mechanism using beta distribution in wireless sensor network. *Journal of ICT Research and Applications*, 13 (1), 79-92.
- [36] Wilensky, U. (1999). *NetLogo*. Evanston, IL: Center for connected learning and computer-based modeling, Northwestern University.
- [37] Antonelli, G. (2013). Interconnected dynamic systems: An overview on distributed control. *IEEE Control Systems Magazine*, 33(1), 76-88.
- [38] Ji, Z., Wang, Z., Lin, H., & Wang, Z. (2009). Interconnection topologies for multi-agent coordination under leader-follower framework. *Automatica*, 45(12), 2857-2863.
- [39] Sapienza, A., Cantucci, F., & Falcone, R. (2022). Modeling interaction in human-machine systems: A trust and trustworthiness approach. *Automation*, 3(2), 242-257.
- [40] Conte, R., & Castelfranchi, C. (1996). Simulating multi-agent interdependencies. A two-way approach to the micro-macro link. In *Social science microsimulation* (pp. 394-415). Springer Berlin Heidelberg.
- [41] Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2), 115-152.
- [42] Gaur, V., & Soni, A. (2012). A novel approach to explore inter agent dependencies from user requirements. *Procedia Technology*, 1, 412-419.
- [43] Scholten, K., Sharkey Scott, P., & Fynes, B. (2014). Mitigation processes-antecedents for building supply chain resilience. *Supply Chain Management: An International Journal*, 19(2), 211-228.
- [44] Motes, J., Sandström, R., Lee, H., Thomas, S., & Amato, N. M. (2020). Multi-robot task and motion planning with subtask dependencies. *IEEE Robotics and Automation Letters*, 5(2), 3338-3345.
- [45] Ravnborg, H. M., & Westermann, O. (2002). Understanding interdependencies: stakeholder



identification and negotiation for collective natural resource management. *Agricultural Systems*, 73(1), 41-56.

- [46] Raposo, A. B., & Fuks, H. (2002). Defining Task Interdependencies and Coordination Mechanism for Collaborative Systems. In *COOP* (pp. 88-103).
- [47] Sapienza, A., & Falcone, R. (2022, September). An Autonomy-Based Collaborative Trust Model for User-IoT Systems Interaction. In *International Symposium on Intelligent and Distributed Computing* (pp. 178-187). Cham: Springer International Publishing.