Epistemic Diversity and Explanatory Adequacy in Distributed Information Processing

Asimina Mertzani[1], Jeremy Pitt[1,2], Andrzej Nowak[2,3], and Tomasz Michalak[4]

[1] Department of Electrical & Electronic Engineering, Imperial College London, UK
[2] Robert B. Zajonc Institute for Social Studies, University of Warsaw, Poland
[3] Department of Psychology, Florida Atlantic University, FL 33431, USA
[4] Institute of Informatics, University of Warsaw, Poland

Abstract A common problem facing an organisation of autonomous agents is to track the dynamic value of a signal, by aggregating their individual (and possibly inaccurate or biased) observations (sensor readings) into a commonly agreed result. A meta-problem is to *explain* the observation of the value: to say what rules produced the signal value that has been observed. In this paper, we use the Regulatory Theory of Social Influence and self-organising multi-agent systems to simulate a Distributed Information Processing unit (DIP) trying to solve such a meta-problem. Specifically, we examine what configuration of initial conditions on the DIP produce what type of epistemic condition for the collective, and determine the *explanatory adequacy* of this condition, i.e. to what extent does the DIP's explanation of the rules match the actual rules. The results offer some further insight into the need for epistemic diversity for self-improvement in dynamic self-organising systems.

Keywords: distributed information processing · explanatory adequacy · knowledge processing · social influence · multi-agent systems.

1 Introduction

A commonly recurring problem confronting an organisation, composed of autonomous agents connected by a (social) network but lacking a central authority, is to map a set of individual measurements, judgements, votes, opinions or preferences into a single collective output. This problem is typically encountered in social systems (e.g. jury trials, deliberative assemblies, etc.) as well as cyber-physical systems (e.g. cybernetic systems, sensor networks, etc.)

An instance of this general problem is truth tracking, when the task of an organisation of autonomous agents is to track the dynamic value of a signal, by aggregating their individual (and possibly inaccurate or biased) observations (sensor readings) into a commonly agreed result. In this sense, the organisation can be seen as a *Distributed Information Processing* (DIP) unit. However, such a DIP can also face a meta-problem: to *explain* the observation of the value – i.e. to say what rules produced the signal value that has been observed. In this case, the DIP is not trying to pool its diverse opinions to order to produce a social choice, but to pool its diverse knowledge to produce a 'plausible' explanation.

This paper investigates a solution to this problem using the Regulatory Theory of Social Influence (RTSI) [12]. RTSI is chosen because it has two unique propositions: firstly, that social influence is bilateral, i.e. that as well as sources seeking targets to influence, targets are seeking sources by whom to be influenced; and secondly that in addition to exchanging opinions, people also exchange information processing rules. Both of these propositions are essential for addressing the problem: the first because we want experts or 'specialists' to emerge, because they know more and are better at solving the problem; and the second because we want their knowledge (of the rules) to flow over the social network.

Therefore, we implement an algorithmic model of RTSI within a self-organising multi-agent systems' to simulate a DIP trying to solve such a meta-problem by proposing (collectively) a set of rules to explain the observed value that may (or may not) match the actual rules that produce the value. Specifically, we experimentally investigate what configuration of initial conditions on the DIP produce what type of epistemic condition of the DIP.

We then want to evaluate the *explanatory adequacy* of the DIP's solution. The term 'explanatory adequacy' is used in linguistics to describe an analysis which provides a 'reasonable' account of a linguistic phenomenon [18]. We want to know if the DIP can produce a 'reasonable' or 'plausible' explanation, based on the extent to which its collective explanation matches the actual cause (i.e. the *ground truth*). We measure the difference using a suitable metric (cosine similarity) and use that as an indicator of explanatory adequacy.

Accordingly, this paper is structured as follows. Section 2 establishes the background of DIP and RTSI, and gives a formal specification of the problem. Section 3 describes the experimental design, Section 4 defines the multi-agent simulation, and Section 5 presents a set of experimental results. After a consideration of related and further work in Section 6, Section 7 concludes that these results offer some further insight into the need for epistemic diversity for self-improvement in dynamic self-organising systems.

2 Background: DIP, RTSI and Plato's Cave

In this section we review the background to this work: organisations as distributed information processing units (DIP), a theory of social influence in such units, the Regulatory Theory of Social Influence (RTSI), and a specification of the problem we are trying to solve, which has similarities, at an abstract level, to the problem posed in the allegory of Plato's Cave (see `https://tinyurl.com/yckmzkyf`).

2.1 Distributed Information Processing Units (DIPs)

Many organisations, in the form of complex cyber-physical, socio-technical or social systems, often have to function as *Distributed Information Processing* units (DIPs), i.e., although composed of many different autonomous components, the components have to act as a collective to transform a set of data inputs into a single output. Although, depending on the context, the precise definition differs (cf. [22] vs. [12]), in this paper the term DIP refers to an organisation of autonomous, (socially) networked agents encountering a requirement to self-manage their diverse, dispersed, and potentially incomplete and inconsistent knowledge.

In general, successful knowledge management enables a DIP to make correct decisions, identify expertise, maintain collective memory, provide education, spark innovation and even accumulate "wisdom". A more mundane function, perhaps, is to converge on a *ground truth* from a set of observations that may have been influenced by environmental or community bias (cf. [20]). Here, though, rather than converging on the truth, we want to study how a DIP can produce *explanatory adequacy*: can the DIP converge on the rules that produced that truth, rather than the truth itself. In this situation, we need a theory which considers social influence not just in terms of the exchange of opinions, but also in the exchange of processing rules. The theory we use is the Regulatory Theory of Social Influence.

2.2 Regulatory Theory of Social Influence (RTSI)

RTSI is a psychological theory proposed by Nowak [12] that focuses on the target's perspective of social influence, and specifically, examines how the targets look for sources by whom to be influenced. This theory emphasises a target's intentions and strategies, and posits that targets actively monitor others' opinions and behaviours, and and are fully engaged in the controlling the influence process.

In this way, a target tries to optimise its decision-making and conserve its own resources by delegating the tasks of information gathering and/or information processing to individuals that they credit for the such tasks. This enables targets to leverage others' network, processing capacity or knowledge, maximising their access to information and information processing rules. Therefore, social influence becomes an instrument of targets to maximise their individual cognitive efficiency and quality of their outcomes, which are reflected by improvements in individual and collective performance over time.

2.3 Problem Specification

In this study, the situation to be addressed by a DIP, using RTSI, is illustrated in Figure 1. The DIP is embedded in an environment, in which there is a process $\mathcal{P}$ that converts some set of inputs into an output. The process $\mathcal{P}$ is parameterised by a set of n processing rules, each with an associated weight in [0..1]. This set of rules, denoted by $\mathcal{K}$, is the *ground truth knowledge* given by:

$$\mathcal{K} = \{(r_i, w_i) \mid i \in [1..n] \wedge \Sigma_{i=1}^{n} w_i = 1.0\}$$

We denote by $\mathcal{K}_r$ the set of rules in $\mathcal{K}$ (without the weights).

Each agent a in the DIP 'knows' $\mathcal{K}^a$, which is some subset of m rules of $\mathcal{K}_r$, $m \leq n$. Each agent associates a random weight with each of its rules, with the weights normalised to sum to 1.0, so that the knowledge of agent a is:

$$\mathcal{K}^a = \{(r_1, w_1), \ldots, (r_m, w_m)\} \quad \text{such that} \quad \forall i, 0 \leq i \leq m. r_i \in \mathcal{K}_r$$

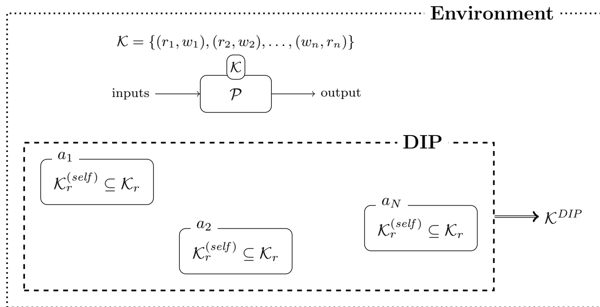

Figure 1: The Problem: Is $\mathcal{K}^{DIP}$ an 'adequate explanation' for output of $\mathcal{P}(\mathcal{K})$?

Note, that if $i = 0$, then the agent knows nothing.

The problem for the $\mathcal{N}$ agents comprising the DIP is to use their partial and distributed knowledge to 'explain' the solution to process $\mathcal{P}$ as parameterised by $\mathcal{K}$. This is done by each agent offering its own explanation $\mathcal{K}^a$ for parameters to process $\mathcal{P}$, and these are 'aggregated' into a collective explanation $\mathcal{K}^{DIP}$. In addressing this task, the agents have three 'tools' at their disposal:

– *sharing*: using RTSI, an agent can ask one of the neighbouring agents in its social network, for a processing rule (or rules) that it (the neighbouring agent) used in its 'explanation' of $\mathcal{K}$.
– *feedback*: each agent receives feedback from the environment on the quality of the collective knowledge and their own contribution, which is used to update 'attitudes' to itself and a neighbouring agent (if it asked one); and
– *'discovery'*: new agents joining the system may bring new knowledge to the system, which may then be shared as above, using RTSI.

Given this context, we investigate:

– what different initial conditions of the DIP, including population variation (e.g. static, dynamic), rate of social learning, and rate of 'discovery', ...
– ...produce what different epistemic condition on the individual knowledge bases, i.e. the similarity of $\{\mathcal{K}^a \mid a \in \mathcal{N}\}$, which we identify as either diversity, incongruence, or stagnation, and...
– ...evaluate explanatory adequacy of $\mathcal{K}^{DIP}$, i.e. the (dis)similarity of $\mathcal{K}^{DIP}$ to $\mathcal{K}$.

In passing, we note that this problem can be seen, at its most abstract, as a form of Plato's Cave, wherein a group of people in a cave try to derive the true nature of an object from the shadow it casts on the cave wall. Note, though, there are three perspectives on knowledge: $\mathcal{K}$, the ground truth knowledge, $\mathcal{K}^{DIP}$, the aggregated knowledge of the DIP, and $\mathcal{K}^{DIP}$, and the "knowledge potential" $\mathcal{K}^{\cup}$, which is an epistemological limit on what it is possible for an agent to know, because this knowledge exists somewhere in the DIP.

However, as in Plato's Cave, this is not a once-off, one-shot problem. The overall situation is as illustrated in Figure 2, where it can see seen that the DIP composition is dynamic (agents may leave and join), and the knowledge made available ("discovered", or introduced along with new agents) also varies.

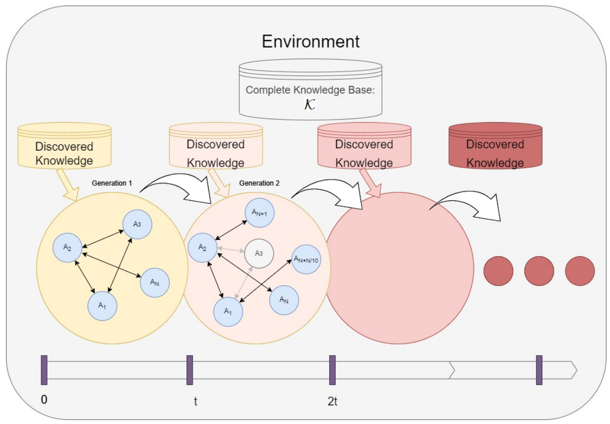

Figure 2: The DIP Unit and Knowledge changing over time

Therefore, "expertise" in the group is also temporary, and knowledgeable agents who may be good at the task may also be lost to the group. As such, the there are two perspectives on the collective: one being a functional perspective as a DIP, where the collective pool knowledge and identify expertise in order to accomplish a common goal (cf. [1]), and the other being a societal perspective, where the group is using social influence as a way to persuade and change attitudes about a common problem (cf. [13]). Accordingly, we will use the terms 'experts' or 'specialists' in the DIP and the sources of influence in RTSI inter-changeably, and equate knowledge with the processing rules; likewise the terms DIP, community and collective are all inter-changeable.

3 Experimental Design

To address the problem defined in the previous section, this section details the experimental design, firstly specifying the initial conditions on the DIP, (i.e. the independent variables), and secondly specifying a metric for computing the DIP's epistemic condition and explanatory adequacy (i.e. the dependent variables).

3.1 Initial Conditions for the DIP (Independent Variables)

For specifying the initial conditions on the DIP, we define two independent experimental variables, $\mathcal{F}$ and $\mathcal{R}$. The former determines the rate of change of the population and rate of change of knowledge. The latter defines a constraint on the RTSI algorithm which affects how the agents communicate the processing rules and how they influence one another.

The DIP will operate in a succession of T *epochs*, and every $t < T$ epochs (except in the *static* condition) some new agents are added and some are removed. In each epoch the DIP will produce and evaluate $\mathcal{K}^{DIP}$ against $\mathcal{K}$, so $\mathcal{F}$ can have one of eight values:

- *static:* The population consists of $\mathcal{N}$ agents, they remain active throughout all T epochs. Each agent *a* is initialised with knowledge $\mathcal{K}^a$ being any subset of $\mathcal{K}$.
- *dynamic:* The population consists of $\mathcal{N}$ agents, and every t epochs a new generation of $\frac{\mathcal{N}}{10}$ agents joins the network and $\frac{\mathcal{N}}{10}$ of the existing agents leave. The agents are initialised having any subset of the eight processing rules.
- *restart:* The population consists of $\mathcal{N}$ agents, and every t epochs $\frac{\mathcal{N}}{10}$ new agents join the network and $\frac{\mathcal{N}}{10}$ leave. The 1^{st} generation of $\mathcal{N}$ agents is initialised so that each agent's *a* knowledge is a subset of $\{(r_1, w_1), (r_2, w_2)\}$. The next generation, which consists of $\frac{\mathcal{N}}{10}$ agents, is initialised knowing either a new rule or no rule, so each new agent's knowledge $\mathcal{K}^a$ is either $\{(r_3, 1)\}$ or $\{\}$, and so on till the generation the generation that knows $\{\}$, or $\{(r_8, 1)\}$. After that generation, the upcoming generations are initialised with knowledge $\mathcal{K}^a$ being any subset of $\mathcal{K}$.
- *iterate:* The population consists of $\mathcal{N}$ agents, and every t epochs $\frac{\mathcal{N}}{10}$ new agents join the network and $\frac{\mathcal{N}}{10}$ leave. The 1^{st} generation is initialised knowing $\{(r_1, 1)\}$, the next knows a new rule, so their knowledge is $\{(r_2, 1)\}$, and so on, so every new generation knows only a new processing rule and the only way to access past knowledge is to interact with others.
- *add rapid:* The population consists of $\mathcal{N}$ agents, and every t epochs $\frac{\mathcal{N}}{10}$ new agents join and $\frac{\mathcal{N}}{10}$ leave. The 1^{st} generation of $\mathcal{N}$ agents is initialised so that each agent's *a* knowledge is a subset of $\{(r_1, w_1), (r_2, w_2)\}$. The next generation knows what their ancestors knew and a new processing rule, so the knowledge of each new agent *a* is a subset of the rules $\{(r_1, w_1), (r_2, w_2), (r_3, w_3)\}$, and so on. So, new knowledge is progressively added to the population through the new generations.
- *add slow:* The is scenario is similar with *add rapid*, but, in this setting, the new generations are added every $t * 1.6$ epochs instead of t.
- *add rapid/slow long:* New generations are added every *t/t* $* \ 1.6$ epochs as per *add rapid/slow*, but the simulator runs for $T * 2$ epochs.

Additionally, for each of the different values of $\mathcal{F}$ we specify two ways of communicating the processing rules $\mathcal{R}$:

- *max*: The sources can only share only one processing rule, therefore they select the rule with greatest weight, which corresponds to the rule that they perceive as the most important piece of knowledge.
- *all*: The sources share their knowledge base, so the target gains access to all the rules that the source knows.

3.2 Epistemic Condition and Explanatory Adequacy

To 'measure' the epistemic condition and the explanatory adequacy of the DIP under different initial conditions, we requires a metric to measure diversity in two dimensions:

- the epistemic diversity, i.e. how different the agents' knowledge bases are from each other, given by $\sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{diff}(i, j)$; and

– the explanatory adequacy, i.e. the divergence of the DIP's knowledge from the ground truth knowledge, given by $\mathbf{diff}(\mathcal{K}^{DIP}, \mathcal{K})$.

For the $\mathbf{diff}$ function, there are many metrics to measure diversity, such as Euclidean Distance, Manhattan Distance, KL divergence etc. We use cosine similarity, because we wanted to to identify the variations between the vectors of weights on processing rules – which represent agents' knowledge and ground truth of the environment – and therefore need a metric that focuses on the orientation rather than the magnitude.

Cosine similarity is a metric used for the comparison of the similarity between two non zero vectors $\mathbf{A}$ and $\mathbf{B}$ in $\mathbb{R}^n$. Specifically, it measures the cosine of the angle between the two vectors, and its value is given by 1:

$$cos\_sim(\mathbf{A}, \mathbf{B}) = \cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\|\|\mathbf{B}\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} {A_i}^2}\sqrt{\sum_{i=1}^{n} {B_i}^2}} \tag{1}$$

After defining the metric for evaluating the performance of the collective and the individuals, we need then to define the two groups of the population that we are going to observe. The first group is the *participants* which refers to all the agents that have been randomly selected to participate to the next epoch. The other group, is the 'specialists', which is a subset of the *participants* and refers to the sources of processing rules. In particular, in this context, if agent *i* asks for processing rules agent *j*, and *j* asks agent *k*, then specialist is considered the agent *k* which constitute the actual source of influence. These agents don't have any notion of expertise, but they are the ones credited by their network.

In the beginning of the experiments, all the agents give equal credence for processing rules to all the agents of their network, they are initialised with different knowledge ,and consequently they give different processing rules to the agents that ask them, and the credence that others give to them is adjusted overtime based on the utility of the information that they offer.

Aiming to identify the capability of the agents to adequately explain the environment (explanatory adequacy), we computed the cosine similarity of the knowledge bases $\mathcal{K}^a$ of agents with the ground truth $\mathcal{K}$, which we denote with $\mathcal{CE}_1$ as well as the cosine similarity of the knowledge bases $\mathcal{K}^a$ of specialists with the ground truth $\mathcal{K}$, which we denote with $\mathcal{CS}_1$. Moreover, to observe knowledge distribution and diversity through the exchange of processing rules (epistemic diversity), we measured the ensemble average cosine similarity of the knowledge bases $\mathcal{K}^a$ of the agents, which we denote by $\mathcal{CE}_2$, the ensemble average cosine similarity of the knowledge bases $\mathcal{K}^a$ of the specialists, which we denote by $\mathcal{CS}_2$. Moreover,The calculation of $\mathcal{CE}_1$ and $\mathcal{CE}_2$ is given by Equations 2 and 3 respectively, and the calculation of $\mathcal{CS}_1$ and $\mathcal{CS}_2$ can be computed by substituting *participants* with *specialists* on those equations.

$$\mathcal{CE}_1 = \frac{\sum_{i=1}^{participants} cos\_sim(\mathcal{K}^i, \mathcal{K})}{\sum_{i=1}^{participants} i} \tag{2}$$

$$\mathcal{CE}_2 = \frac{\sum_{i=1}^{participants,} \sum_{j=1}^{participants, j \neq i} cos\_sim(\mathcal{K}^i, \mathcal{K}^j)}{(\sum_{i=1}^{participants} i)^2 - \sum_{i=1}^{participants} i} \tag{3}$$

4 Formal Specification

This section provides the formal specification of the multi-agent model. This section defines the agents of the system, the environment in which they exist as well as the RTSI algorithm for knowledge processing based on which the agents act in this environment.

4.1 The Environment

The environment $\mathcal{E}$ consists of a network of agents which try to identify the complete knowledge base $\mathcal{K}$ corresponding to the ground truth. The agents are connected through a network $\mathcal{G}(\mathcal{N}$,m,$\mu)$ which is a Klemm-Eguiluz network [8] with $\mathcal{N}$ nodes (where each node is a agent), m the number of fully connected agents used for the generation of the network and characterised as "active", and μ the probability of a new agent to be attached to one of the "active" agents (otherwise the agents attaches to an inactive agent and becomes active, substituting a randomly selected agent from the active agents) as described in [17]. This network type was selected because it combines all three properties of many "real world" irregular network, that is high clustering coefficient, short average path length, and scale-free degree distribution.

4.2 Agent Specification

The autonomous networked units of the population are described by the term "agents". The specification of the agents is based on the specification in [15], and is given by the 6-tuple defined in 4:

$$i = \langle SN, \mathcal{K}^i, sc_i, TN_i, a, b \rangle \tag{4}$$

where SN_i is its social network (connected neighbours), $\mathcal{K}^i$ its knowledge, which is a subset of the knowledge in the environment $\mathcal{K}$ (possibly with different weights), sc_i is a measure of self-confidence of its knowledge (relative to its neighbours), in whom it also gives credence $\tau_{i,j}$ for each agent $j \in SN_i$ (cf. [3]). These values are stored in an ordered list of credence to neighbours TN_i, and are updated each time agent i asks a neighbour j for knowledge (i.e. for a processing rule or rules) depending on how well (similarly) this neighbour approximates the complete knowledge of the environment $\mathcal{K}$. Each agent orders its neighbours in descending order of credence. Each agent has also two reinforcement coefficients a, b which define the rate of change of self-confidence and credence to the network after each epoch.

4.3 Algorithm

The algorithm is an iterative process of T epochs, and in every epoch each participating agent goes through the steps described in 1. Therefore, in every epoch, a subset of agents $\mathcal{A}$ is randomly selected to participate in the next epoch $participants \subset \mathcal{A}$. The aim of the agents is to manage to produce a good approximation of the complete knowledge base of the environment $\mathcal{K}$, while they are only given only a subset of this knowledge.

Throughout the epochs each agent look for sources in the DIP that can provide the processing rules that produce the best approximation of the complete knowledge $\mathcal{K}$.

The knowledge of the DIP can be accessed by asking a neighbouring agent. Therefore, in each epoch, each agent iterates over its social network SN_i according to the order of credence TN_i, to find the source to ask. The neighbour selected *j* is questioned how similar is its knowledge with the complete knowledge base $\mathcal{S}^{\cup}$ and also the agent asking computes the similarity of its own knowledge with the ground truth $\mathcal{S}^{(self)}$. If the neighbour asked can offer a better approximation of the ground truth than the agent asking has, then the agent proceeds in asking the neighbour for processing rules. Depending on the value of the independent variable $\mathcal{R}$, the agent that asked for processing rules (target) either receives as a reply a processing rule with a weight (which is the processing rule of the neighbour *j* that has the greatest weight for *j*), if $\mathcal{R} = max$, otherwise it receives all the processing rules and their weights. Then, *i* is integrating this knowledge $\mathcal{K}^{\cup}$ to its knowledge $\mathcal{K}^i$. After that follows the process of reflection, in which each agent updates its credence to the neighbour selected τ_{ij} and its self-confidence sc_i depending on whether it can more adequately explain the environmental knowledge than its neighbour.

$$w_{r_i avg} = \frac{\sum_{j=1}^{participants} w_{r_i j}}{\sum_{j=1}^{participants} j} \quad (5) \qquad\qquad w_{r_i c} = \frac{w_{r_i avg}}{\sum_{i=1}^{n} w_{r_i avg}} \quad (6)$$

In this way, the collective forms a knowledge $\mathcal{K}^{DIP}$ which is the outcome of the aggregation of the participating agents' knowledge and normalising the weights, as shown in 5 and 6. The collective/DIP knowledge is defined as in 7.

$$\mathcal{K}^{DIP} = \{(r_1, w_{r_1 c}), ..., (r_r, w_{r_r c})\} \quad (7)$$

Algorithm 1: RTSI for knowledge seeking: for each agent i

```
j = selected neighbour from network;
S^(self) = cos_sim(K^(self), K);
S^∪ = cos_sim(K^∪, K);
if S^(self) < S^∪ then
    if R = max then
        K^∪ = {(r_x, w_{r_x}) | (r_x, w_{r_x}) ∈ K^j ∧ ¬∃(r_y, w_{r_y}) ∈ K^j . w_{r_y} > w_{r_x}} ;
    else
        K^∪ = K^j;
    end
end
if S_i^(self) > S^∪ then
    sc_i = sc_i + a * (1 − sc_i);
    τ_{i,j} = τ_{i,j} − b * τ_{i,j};
end
if S^(self) < S^∪ then
    sc_i = sc_i − b * sc_i;
    τ_{i,j} = τ_{i,j} + a * (1 − τ_{i,j});
end
```

According to this formal specification, a multi-agent simulator has been implemented in Python3, which is an extension of the system presented in [14] to include the exchange of processing rules. This simulator was used to run a series of experiments, the results of which are present in section 5.

Table 1 presents the simulator parameters for the agents and the RTSI algorithm. This specifies either a fixed representative value used in the experiments (e.g. the number or agents, reinforcement coefficients, etc.) or a range of values for those that are randomly assigned (e.g. the individual agent knowledge bases). Other experiments could examine different combinations of initialisation of these variables, e.g. to look for effects of scale, but this is left for further work.

Symbol	Description: Factor of Agent i	Initial Condition/Range
$\mathcal{N}$	network of agents	100
m	total number of edges	$\frac{\mathcal{N}}{10}$
μ	number of edges to 'active' agents	0.75
$participants$	agents participating in the next epoch	$\frac{\mathcal{N}}{2}$
$\mathcal{K}^i$	individual knowledge base	$\{(r_1, w_1), \ldots, (r_m, w_m)\}, \forall i, 0 \leq i \leq m.r_i \in \mathcal{K}_r$
r_k	processing rule k	k = {1,2,...,r}
w_{r_k}	weight of rule r_k	$0 \leq w_{r_k} \leq 1$
sc_i	self-confidence for similarity of knowledge	0.5
a, b	self-confidence & credence reinforcement coefficients	0.1, 0.1
SN_i	social network	1 to $\mathcal{N}$ agents
TN_i	ordered list of credence to social network	list length from 1 to $\mathcal{N}$
τ_{ij}	credence to agent j	$0 \leq \tau_{ij} \leq 1$
$\mathcal{S}^{(self)}$	cos_sim between self and environmental knowledge	$0 \leq \mathcal{S}^{(self)} \leq 1$
$\mathcal{S}^{\cup}$	cos_sim between knowledge of agent (neighbour) asked and ground truth	$0 \leq \mathcal{S}^{\cup} \leq 1$
$r_{\cup}$	rule proposal of (neighbour) agent asked	$r_k. \in \mathcal{K}_r$
$w_{\cup}$	weight proposal of (neighbour) agent asked	$0 \leq w_{\cup} \leq 1$
$\mathcal{K}^{\cup}$	knowledge proposal of (neighbour) agent asked	$\{(r_1, w_1), \ldots, (r_m, w_m)\}, \forall i, 0 \leq i \leq m.r_i \in \mathcal{K}_r$
$\mathcal{K}^{DIP}$	collective knowledge	$\{(r_1, w_1), \ldots, (r_m, w_m)\}, \forall i, 0 \leq i \leq m.r_i \in \mathcal{K}_r$
$w_{r_i avg}$	average weight of rule r_i	$0 \leq w_{r_i avg} \leq \frac{\mathcal{N}}{2}$
$w_{r_i c}$	normalised average weight of rule r_i	$0 \leq w_{r_i c} \leq 1$

Table 1: Simulator Parameters and Variables

5 Experimental Results

This section describes three experiments which investigate what initial conditions on the DIP produce what type of epistemic condition, and how 'adequately' does that epistemic condition explain the ground truth knowledge. The experiments range over the variables $\mathcal{F}$ and $\mathcal{R}$ of Section 3.1 under the initial conditions specified in Table 1, with $T = 5000, t = 300$:

- Experiment 1: Static population of agents, with complete fixed knowledge, and dynamic population with complete fixed knowledge (all the knowledge is available from the first epoch of the simulation).
- Experiment 2: Dynamic population with progressive addition of new knowledge but non-persistence of 'discovered' knowledge.
- Experiment 3: Dynamic population with progressive addition of new knowledge and with persistence of already 'discovered' knowledge.

The following subsections describe the results of each experiment in turn, before discussing some over-arching results in Section 5.4.

5.1 Experiment 1: Static and dynamic populations

In the first set of experiments, we explore the dynamics of the system for $\mathcal{F}$ being set to *static* and *dynamic*, and agents are initialised with any combination of the available processing rules.

Figure 3 illustrates the evolution of common and specialists knowledge for the different settings. Specifically, the 1^{st} column illustrates the results for *static* and the 2^{nd} for *dynamic* for $\mathcal{R}$ being *max* and *all*. The black line is calculated according to $\mathcal{CE}_1$, the red based on $\mathcal{CS}_1$, the green according to $\mathcal{CE}_2$, and, finally, the blue line based on $\mathcal{CS}_2$.

Therefore, the black line indicates how 'adequately' the DIP identifies the ground truth $\mathcal{K}$, while the red line indicates whether how 'adequately' the specialists identify the ground truth $\mathcal{K}$. The green and blue lines demonstrate the diversity of knowledge within the collective and within the specialists, showing the (dis)similarity between the knowledge of each group.

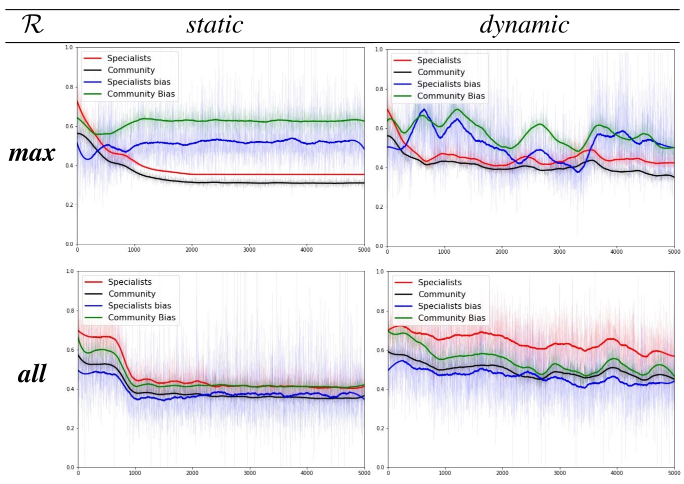

Figure 3: Exp. 1: Knowledge dynamics for static and dynamic population.

Starting with the *static* condition, when $\mathcal{R} = max$, the similarity between the processing rules of the agents is high, since the group as a whole is influenced by the specialists to promote a single rule. By contrast, when the sources share all their knowledge ($\mathcal{R}$ =*all*), the community and the specialists similarity is decreased. However, the lines corresponding to how well do specialists and community track the environmental knowledge (red and black) remain low in both cases. This is because the population is static, therefore the community is prone to ask the sources credited during the first epochs, regardless of whether they maintained their knowledge. Static populations are stable but also stagnant and agent don't increase significantly their processing capacity although they could (since all the knowledge is discovered).

With the *dynamic* condition, for both *max* and *all*, the agents can better explain the environmental knowledge. In the former case, the knowledge of the specialists and the community is continuously modified as illustrated by the fluctuating green and blue lines. This demonstrates that different *epistemes* are generated in this condition, and the system could be characterised as quasi-stable and moving from one temporary equilibrium to another with different values for its control variables (cf. [16]). In the latter case, the specialists are well-identified and have significantly higher similarity with the environmental knowledge than the community; however it seems that the other agents cannot assimilate this knowledge and the DIP knowledge seems stagnant.

5.2 Dynamic population, progressive addition, non-persistence

In this experiment, we observe how the system works with a dynamic population (in which the specialist sources are not so easily identified), there is progressive addition of new knowledge brought by joining agents, but knowledge is non-persistent (i.e. new agents only bring new 'discovered' knowledge).

Figure 4 illustrates the results for *restart* and *iterate* in the first and the second column respectively.

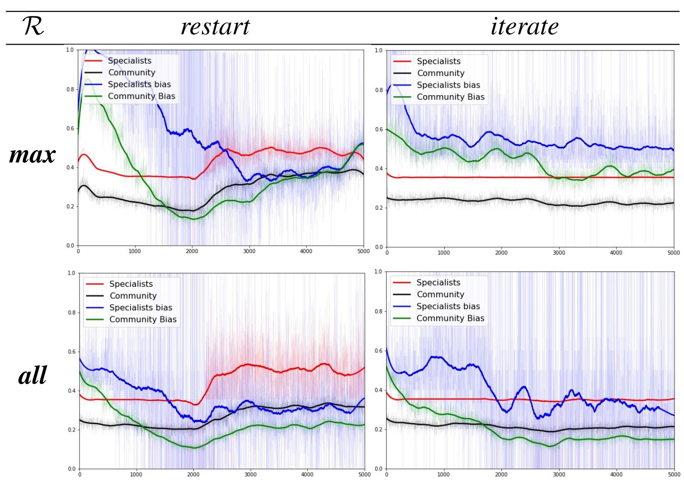

Figure 4: Exp. 2: Knowledge dynamics for DIP with non-persistent knowledge

In the *restart* condition, the first generation is initialised with two rules available, the next joins with the third rule, and so on until the 8^{th} generation that has all the rules available (as defined in section 3.1). For $\mathcal{R} = max$, the similarity of processing rules within the community is high since agents are given only one rule from the sources. This phenomenon is less striking for $\mathcal{R} = all$, where agents quickly assimilate new knowledge. Note that when new generations possess only one processing rule (only the new piece of knowledge), agents consider that they cannot learn from others, and their knowledge remains narrow (low similarity with the environmental knowledge). However, after epoch 2100, when all the processing rules are made available for the new generations, there is a significant increase in the community and specialists' knowledge because agents have different levels of knowledge (i.e. different similarity with environmental knowledge) and they seek sources to provide them with missing bits of knowledge.

The phenomenon of agents not asking for processing rules because they perceive others as having similar knowledge is even clearer under the *iterate* condition. Particularly, in both *max* and *all*, most agents seem to have equal knowledge (i.e. equal similarity of own processing rules and environmental processing rules), due to the fact that they all have either zero or one processing rule, and consequently only agents having an empty knowledge base ask for knowledge. This variation between the empty knowledge base and the knowledge base containing one processing rule generates the difference in the similarity of the knowledge of the sources and the community (red and green lines), with the specialists. Additionally, the fluctuation of the intrinsic similarity of the collective as well as the group of specialists is caused by the randomness in the selection of agents to leave and join the network.

5.3 Dynamic population, progressive addition, persistence

In this experiment, we observe the behaviour of the system under progressive addition of knowledge, but new agents may bring any discovered knowledge. Figure 5 demonstrates $\mathcal{CE}_1$, $\mathcal{CS}_1$, $\mathcal{CE}_2$, and $\mathcal{CS}_2$ for the *add rapid/slow (long)* scenarios, for $\mathcal{R}$ *max* in the first row and *all* in the second.

add rapid	*add slow*	*add rapid long*	*add slow long*

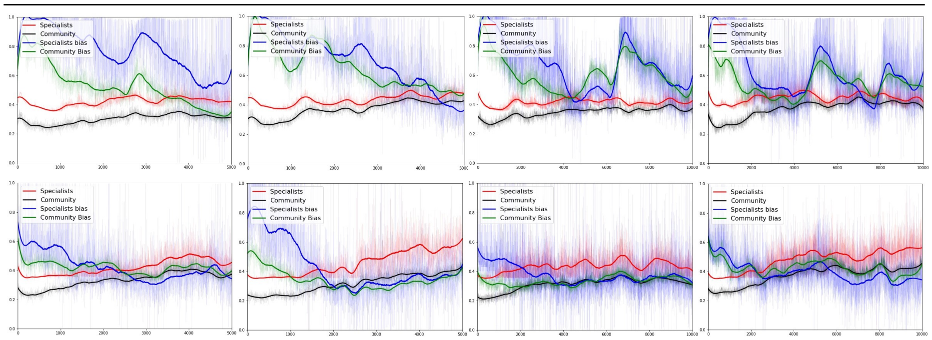

Figure 5: Exp. 3: Knowledge dynamics for progressively added knowledge.

The rapid progressive addition of knowledge allows minor improvement both in short-term and long-term (*add rapid* and *add rapid long*). Particularly, in both *max* and *all*, the specialists and community knowledge remains low (red and blue lines) because new rules cannot be assimilated. By contrast, the slow addition fosters epistemic improvement (*add slow* and *add slow long*). Moreover, in all these scenarios, when $\mathcal{R} = max$ different *epistemes* are produced because sources share parts of their knowledge and both the community and the sources develop different beliefs over time.

5.4 Summary of experiments

To conclude this section, Table 2 summarises what configuration of initial conditions for the DIP produces what type of epistemic condition of the DIP, and assesses the capability of the DIP to explain adequately the environment.

$\mathcal{F}$	$\mathcal{R}$	**Epistemic Condition**	**Explanatory Adequacy**
static	max	epistemic stagnation	$\mathcal{K}^{DIP} \ncong \mathcal{K}$
static	all	epistemic stagnation	$\mathcal{K}^{DIP} \ncong \mathcal{K}$
dynamic	max	epistemic incongruence	conditionally $\mathcal{K}^{DIP} \cong \mathcal{K}$
dynamic	all	epistemic incongruence	conditionally $\mathcal{K}^{DIP} \cong \mathcal{K}$
restart	max	epistemic diversity	$\mathcal{K}^{DIP} \cong \mathcal{K}$
restart	all	epistemic diversity	$\mathcal{K}^{DIP} \cong \mathcal{K}$
iterate	max	epistemic stagnation	$\mathcal{K}^{DIP} \ncong \mathcal{K}$
iterate	all	epistemic stagnation	$\mathcal{K}^{DIP} \ncong \mathcal{K}$
add rapid	max	epistemic incongruence	conditionally $\mathcal{K}^{DIP} \cong \mathcal{K}$
add rapid	all	epistemic incongruence	conditionally $\mathcal{K}^{DIP} \cong \mathcal{K}$
add slow	max	epistemic diversity	$\mathcal{K}^{DIP} \cong \mathcal{K}$
add slow	all	epistemic diversity	$\mathcal{K}^{DIP} \cong \mathcal{K}$

Table 2: Summary of experimental results

Starting from *static*, we observe that with a static population the DIP has a high similarity of knowledge, and therefore they seem to be congruent, but knowledge does not seem to be exchanged over the social network. This does not allow further improvement in the system and potential adaptation to a dynamic environment. Moreover, the collective has a low similarity of knowledge with the environment, which means that they are not adequately explaining the knowledge $\mathcal{K}$. In contrast, dynamic populations that have all the knowledge available from the first epochs (*dynamic*) seem to be more diverse, and they transition from a status of higher to lower congruence and vice-versa. Although for certain periods of time they manage to accurately explain the environment, there are other periods that they do not succeed in identifying the ground truth.

Moreover, when agents perceive their knowledge to be similar to others knowledge, they do not ask for processing rules and the collective knowledge stagnates. Specifically, in the *restart* condition, during the first epochs where they are given only one processing rule, agents do not communicate their knowledge. This is also the case for the *iterate* condition, where knowledge remains stagnant while the collective is fragmented. Therefore, we argue that knowledge remains static and the agents do not manage to model the phenomenon which they observe in the environment, when they consider that others are incapable of helping them (perceiving their knowledge similar with their own knowledge), although they might have different knowledge that is useful for them.

This would suggest that systemic evolution and epistemic diversity require both knowledge differentiation and the capability of agents to perceive this differentiation. However, in *restart*, when all the knowledge becomes available (after 2500 epochs), the agents quickly increase the utility of the collective knowledge with respect to explanatory adequacy and they produce a collective knowledge that is a better approximation of the knowledge situated in the environment. It is worth noting that the sources seem to 'emerge', i.e. to increase the utility of their knowledge, significantly more than the community, which shows that the ones identified by the collective as specialists are also more likely to assimilate new knowledge.

Furthermore, the rapid addition of new processing rules fosters diversity, but the agents do not have enough time to adapt and assimilate new knowledge; therefore they can be congruent in the short term but incongruent in the long-term. This cannot guarantee that the DIP will manage to produce an adequate explanation of the environmental knowledge for an extended period. In this case, we observe different *epistemes* being generated, which could be considered a demonstration of Foucault's Theory of Knowledge & Power in cyber-physical systems. However, when agents share all their knowledge (*all*) the population becomes more congruent with the environment and has the potential to evolve since it can provide an adequate explanation of the environment.

6 Related and Further Research

Related research has extensively studied issues of consensus formation in complex systems [11]. More specifically, previous work has focused on the conditions which lead to the alignment of the network [4], as well as the division of it into multiple opinions [7]. Much effort has also been made to identify the probability of forming a majority depending on the network topology [5].

A baseline for using RTSI as a model of distributed information processing, proposing the exchange of subjective opinions for the formation of a collective decision and the self-organisation have been established in [15]. This work extends this model of RTSI in a different direction, and specifically proposes the communication of processing rules not for forming a collective opinion but for developing a collective knowledge and social explanations. Lopez-Sanchez and Müller [10] suggest that social influence in the form of hate speech can propagate through the whole virtual community and propose countermeasures such as education, deferring hateful content and cyber activism as mechanisms for altering it. In this research, we argue that social influence can be also used as an instrument for spreading knowledge and providing explanations instead of propagating hate and negative opinions.

Additionally, there is a substantial body of literature in topics of information sharing and norm emergence. Villatoro et al. [21] proposed the use of social instruments to facilitate norm convergence and proved that the subconventions delay global convergence and jeopardise stability. Incremental social instruments and creating ties between agents has also provided a mechanism for dissolving self-reinforcing structures and facilitating global norm emergence [9]. Norm or convention emergence can be also achieved though social learning [19], and under various topologies [2]. Although these works offer deep insight into the emergence of a collective property (socially-constructed behaviours) from local interactions, our approach differs by proposing RTSI as a mechanism for producing 'adequate' collective explanations for external properties from local interactions.

Further research could establish a set of evaluation criteria and metrics for multi-agent populations that face problems of producing social and environmental explanations. Additionally, further work on different conditions in the population such as agents having personalities or intentionally sharing only that part of their knowledge they want to, in order to direct opinions and thoughts, or more advanced methods for developing self-confidence and credence to the network, such as models of costly signaling or block-based approaches. Moreover, future research could extend the communication of the network and allow not only the exchange of processing rules but also the exchange of the reasons for selecting these processing rules.

Furthermore, the present setting could be modified so that not only can the community adapt its knowledge but also the environmental knowledge can change, towards or away from to the knowledge of the collective. Moving towards the might cause a loss of expertise that becomes critical when the environmental knowledge moves away from the DIP knowledge. Finally, a really ambitious step is to move from explanation to innovation, how knowledge of the rules can be used to shape the environment for purposes of self-improvement.

7 Summary & Conclusion

In summary, the contributions of this paper are:

- We have specified a problem of explanatory adequacy for self-organising multi-agent systems, as disparate agents use their social network to aggregate their possibly incomplete and inconsistent knowledge bases to 'explain' some observed phenomenon;
- We have implemented an algorithm based on the Regulatory Theory of Social Influence (RTSI), which includes bilateral influence between targets and sources and the exchange of information processing rules, and implemented it in a simulator for a Distributed Information Processing unit (DIP); and
- We have run three experiments to explore what initial conditions of the DIP and the RTSI algorithm lead to what type of epistemic condition for the collective, and use a similarity metric to determine how well these conditions do indeed provide explanatory adequacy.

In conclusion, these experiments point to the following postulates that will be explored in further work, but we regard as crucial for developing DIP for socio-technical and cyber-physical systems embedded in dynamic environments. These postulates are that systemic self-improvement through epistemic evolution requires diversity, a willingness to learn, and having good intentions.

Primarily, we argue that systemic self-improvement epistemic evolution requires *diversity*. We observed that DIP composed of almost identical agents, in terms of having the same knowledge, could not improve their explanatory adequacy. It is also important that knowledge should be preserved somewhere in the network, because this knowledge might yet be relevant and useful at a later time. Moreover, not only should the knowledge of the agents be diverse, but *the agents should be capable of understanding the diversity of knowledge sources*, and be able to identify from where whom or where they can reliably acquire or consult expertise.

Secondly, systemic evolution requires each individual to be *willing to learn*. Epistemic evolution requires agents who are, in the first place, willing to make the effort to ask and to answer, but are also willing to make the effort to assimilate the answer. Both of these are assumptions made by the RTSI algorithm, and factoring in obdurate agents (who will not ask) or disruptive agents, who block or break communication chains requires further experimentation. Another line of investigation would be if the DIP included 'denialist' cliques, a group of self-supporting agents whose inaccurate knowledge is altered by neither evidence nor argument, and what impact such cliques might have on effective performance.

Finally, another important requirement for RTSI to enable a DIP to solve the explanatory adequacy problem is that both sources and targets must have *good intentions*. From one side, a knowledge source, who is responsible for transferring knowledge, should not have any intention to manipulate either the target or worse, in fact, perturb the value of $\mathcal{K}^{DIP}$ for its own interests rather than the collective (public) interest. From the other side, the knowledge seeker, or target, should intend to evolve positively, but critically and not naively.

To highlight the importance of good intentions from the source's side, we note that all the experiments implicitly share a common characteristic: the strong relationship between knowledge and power. In particular, under all conditions, the DIP does manage to identify the 'specialist' individuals (who are best, or least bad, at the task) and credits them for sharing their knowledge. Consequently, the most knowledgeable agents are also the ones who could, in effect control and manipulate common knowledge and public opinion. This way, these agents can not only occupy the prosocial role of knowledge gatekeeper, but could also become an antisocial 'knowledge dictator'. This dynamic is clearly illustrated in Foucault's [6] observation that power is based on and reproduces knowledge, while knowledge in turn begets power. Therefore, if the sources have other motives for sharing their knowledge, the expertise of the network can degenerate into an oligarchy (a 'knowligarchy').

Acknowledgements

We are particularly grateful to the three anonymous reviewers whose many insightful comments helped to revise and improve the presentation of this work.

References

1. Afshar, M., Asadpour, M.: Opinion Formation by Informed Agents. *Journal of Artificial Societies and Social Simulation* **13**(4) (2010)
2. Airiau, St., Sen, S., Villatoro, D.: Emergence of conventions through social learning. *Autonomous Agents and Multi-Agent* , **28**(5), 779–804 (2014)
3. Asch, S.E.: Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological monographs: General & Applied* **70**(9), pp.1–70 (1956)

4. Chacoma, A., Zanette, DH.: Opinion Formation by Social Influence: From Experiments to Modeling. *PLoS One* **10**(10), e0140406 (2015)
5. Fadda, E., He J., Tessone, Cl., Barucca, P.: Consensus formation on heterogeneous networks. *CoRR* **abs/2111.11949** (2021)
6. Foucault, M.: *Power/Knowledge: Selected Interviews & Other Writings*. Colin Gordon (1980)
7. Hegselmann, R., Krause, U.: Opinion Formation by Social Influence: Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation* **5**(3), (2002)
8. Klemm, K., Eguiluz, V.M.: Growing scale-free networks with small-world behavior. *Phys. Rev. E* **65**(5), 057102 (2002)
9. Liu, Y., Liu, J., Wan, K., Qin, Z., Zhang, Z., Khoussainov, B., Zhu, L.: From Local to Global Norm Emergence: Dissolving Self-reinforcing Substructures with Incremental Social Instruments. In: *International Conference on Machine Learning* (PMLR), pp. 6871–6881. (2021).
10. Lopez-Sanchez M., Müller A.: On Simulating the Propagation and Countermeasures of Hate Speech in Social Networks. *Applied Sciences* **11**(24) (2021).
11. Medo, M., Mariani, M.S. and Lü, L.: The fragility of opinion formation in a complex world. *Commun Phys* **4**(75) (2021).
12. Nowak, A.K., Vallacher, R.R., Rychwalska, A., Roszczyńska-Kurasińska, M., Ziembowicz, K., Biesaga, M., Kacprzyk-Murawska, M.: Target in control: *Social influence as distributed information processing*. Springer International Publishing (2019)
13. Cacioppo, J.T., Petty, R.E.: The elaboration likelihood model of persuasion. In T. Kinnear (ed.), *NA - Advances in Consumer Research* **11**:673-675, Association for Consumer Research (1984)
14. Pitt, J., Nowak, A., Michalak, T., Borkowski, W., Vallacher, R.: Knowing What the Bits Know: Social Influence as the Source of Collective Knowledge. *Second International Workshop on Agent-Based Modelling of Human Behaviour* (ABMHuB) `http://abmhub.cs.ucl.ac.uk/2020/papers/Pitt.pdf` (2020)
15. Pitt, J.: Interactional Justice and Self-Governance of Open Self-Organising Systems. In: *11th IEEE International Conference on Self-Adaptive and Self-Organizing Systems* (SASO), pp. 31-40 (2017)
16. Pitt, J. and Ober, J.: Democracy By Design: Basic democracy and the self-organisation of collective governance. In: *12th IEEE International Conference on Self-Adaptive and Self-Organizing Systems* (SASO), pp. 20-29 (2018)
17. Prettejohn, B., Berryman, M., Mcdonnell, M.: Methods for Generating Complex Networks with Selected Structural Properties for Simulations: A Review and Tutorial for Neuroscientists. *Frontiers in computational neuroscience* **5** (2011)
18. Rizzi, L.: The concept of explanatory adequacy. In: I.Roberts (ed.), *The Oxford Handbook of Universal Grammar*, Oxford University Press (2016).
19. Sen, S., Airiau, St.: Emergence of norms through social learning. In: *20th International Joint Conference on Artificial Intelligence* (IJCAI), vol. 1507, pp.1507-1512. (2007)
20. Sîrbu, A., Pedreschi, D., Giannotti, F., and Kertész, J.: Algorithmic bias amplifies opinion fragmentation and polarization: A bounded confidence model. *PloS one* **14**(3), e0213246 (2019).
21. Villatoro, D., Sabater-Mir, J., Sen, S.: Social instruments for robust convention emergence. *22nd International Joint Conference on Artificial Intelligence* (IJCAI), pp.420-425 (2011)
22. Wiggins, R.E.: Distributed information processing: trends and implications. *Aslib Proceedings* **37**(2), 73–90 (1985)