Knowing What the Bits Know: Social Influence as the Source of Collective Knowledge

Jeremy Pitt[1,2], Andrzej Nowak[2,3], Tomasz Michalak[4], Wojciech Borkowski[2] and Robin Vallacher[3]

[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
j.pitt@imperial.ac.uk, andrzejn232@gmail.com, tpm@mimuw.edu.pl

Abstract

Every organisation seeks to narrow the gap between the potential and the actual value of its collective knowledge for 'correct' or 'optimal' decision-making capability. This problem has deep historical roots in social choice theory, organisational theory and political science, but recent and rapid developments in social networks, big data and artificial intelligence have only amplified and exacerbated the scale and complexity of knowledge management. This paper addresses the magnified problem from a psychological perspective, using a model called RTSI (Regulatory Theory of Social Influence). This focuses on how targets (of information) seek sources to be influenced by, rather than the alternative traditional model, of sources influencing targets. Using agent-based modelling, a series of experiments shows how a group of agents using RTSI can pool their collective knowledge to becomes a *multi-dimensional distributed information processing unit* that can track a 'true' signal or a 'community' signal.

Introduction

Every organisation seeks to narrow the gap between the potential and the actual value of its collective knowledge for 'correct' or 'optimal' decision-making capability. The gap is perhaps most widely known through the corporate lament "If only HP [Hewlett-Packard] knew what HP knows", recognising the difference between the organisation's actual performance with fragmented knowledge and the organisation's potential performance if that knowledge could be synthesised (Sieloff, 1999).

This *knowledge gap* problem has deep historical roots in political science (Ober, 2008), social choice theory (List and Goodin, 2001), and organisational theory (Davenport and Prusak, 1998). However, recent and rapid developments in social networks, big data and artificial intelligence have only amplified and exacerbated the scale and complexity of knowledge management. For example, misinformation and confirmation bias have caused sub-optimal decision-making on a global scale, adversely affecting public health decisions (Hussain et al., 2019); participatory sensing applications have resulted in an asymmetry of benefit between the data generators and the data aggregators (Macbeth and Pitt, 2015); while biased datasets have been the basis for training machine learning algorithms that simply reproduce that bias, for example in employment and policing (Asaro, 2019).

This paper addresses the magnified knowledge gap problem from a psychological perspective. Since influence is so important in human behaviour in social interactions, we examine a new model called RTSI (Regulatory Theory of Social Influence) (Nowak et al., 2020). This model focuses on how targets (of information) seek sources to be influenced by, rather than the alternative traditional model, of sources influencing targets. A target's decision-making is based firstly by a competence-based decision (do-it-yourself or delegate to others), and, if delegating, a trust-based decision. Experiments with a multi-agent simulator shows that a group of agents using this model of social influence the collective can become a *multi-dimensional distributed information processing unit* that can accurately track either the 'true' signal, or the 'community' signal.

The argument is developed as follows. The next section expands on the multi-disciplinary background to this work, including a more detailed description of RTSI. We introduce a test environment and objectives, and specify algorithms for software agents implementing a circumscribed version RTSI. A series of experiments explores the behaviour of RTSI in the test environment. We summarise and conclude that agent-based modelling can be used to represent a theory of human psychology and shed some interesting light on knowledge management in organisations.

Social Influence

Influence plays a key role in human behaviour in social interactions, and such it has attracted considerable attention in the literature. In his classic work, Allport (1937) defined social influence almost as broadly as the field of social psychology, as a change in thoughts, feelings and behaviour resulting from real or imagined presence of others. Research in social psychology has concentrated on social influence described from the perspective of the agent of influence (i.e. the source) whose overarching concern is how to alter the opinions, decisions, or courses of action of the intended target of the influence. In this vein influence can take many forms such as conformity (Asch, 1956), obedience (Milgram, 1963), persuasion (Petty and Cacioppo, 1986), compliance (Cialdini, 2016) and control. The rules of influence specify the principles by which the source can influence thinking and decisions of the target and overcame the passivity or resistance of the target. Implicit in this perspective is the assumption that the source's agenda is not shared by the target and is beneficial to the source. In this perspective social influence is close to power and manipulation.

Social influence, however, may be beneficial to the target in such processes as, for example, seeking information (Baldwin and Hunt, 2002) and advice (Dyer and Ross, 2008). The observation that influence may serve the interests of the target of influence, led to the development of the Regulatory Theory of Social Influence (RTSI) (Nowak et al., 2020) that assumes that individuals often desire to be influenced and actively search for influence. In particular, from the target's perspective, social influence is tantamount to the delegation of information gathering and processing to others. The target chooses the topic of influence, the sources of influence and the form of influence. The goal of seeking influence is the formation of a judgment or reaching a decision on an issue. Delegating information processing to others is functional; it saves the processing resource and processing time and may improve the quality of the decision or judgment. It is, however, risky as one may be misled, exploited, or receive information and advice of poor quality. The demand for efficiency pushes individuals toward delegating, risk avoidance induces individuals to gather and process the information themselves.

The rules of RTSI describe how individuals can solve the dilemma. In short individuals tend to devote their own processing resources to matters that are associated with the highest risk while delegating to others judgments and decisions associated with a lower risk. At intermediate levels of risk they use a mixed strategy, delegating only part of decision-*informing* information processing, while reserving final decision-*making* for themselves. If they are unable to reach a decision by themselves, individuals can even delegate decision-making, while checking the advice.

In this process *trust* is essential. Trust toward the other is the most important variable controlling influence seeking. Individuals tend seek information and advice from the trusted others, if this is not possible they gather and process the information themselves. Information from the trusted others is also weighted the most. Trust is a dynamical variable which reacts to the quality of information received from others. Individuals increase the trust toward those who give them true and accurate information and they decrease trust toward those who give them inaccurate information. Individuals also take into account their own expertise: the lower it is, the more likely they are to delegate information processing to others. Importance of the decision, own knowledge and coherence of information are the other variables in social influence. The joint influence of these variables decides when and how individuals delegate information gathering and processing to others while minimizing the associated risk. RTSI proposes that social influence initiated by the target turns social groups into optimizing distributed information processing systems.

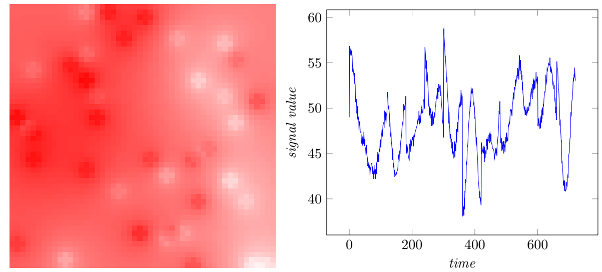

Figure 1: Variation of ambient signal with location and time.

Test Environment and Objectives

This section defines a test environment which exhibits features of uncertainty, incompleteness and variation, in order to test the essential properties of RTSI: social networking, self-organisation and distributed information processing.

Test Environment

To explore a computational model of RTSI, we set up a test environment which consists of a grid, of variable size $[1..maxX, 1..maxY]$, with $maxX$ and $maxY$ being specified as experimental parameters.

There is a signal, ρ with a value that varies with time from a starting value (called the base signal). The ambient value of the signal is given by combining a random number of sine functions (called *scramblers*) with a randomly generated amplitude, frequency and phase. Because adding two sine waves produces another sine wave, a random amount of noise is also added, with the scramblers' parameters reset periodically. (The intention being that, while not impossible, it should be relatively complex (computationally) to learn the function: the sparsity of data relative to the size of search space means that the rate of change of an unknown number of environmental parameters is 'sooner' than a "learning algorithm of corresponding complexity" could converge on the number and values of those parameters.)

Moreover, at a random number of locations are added signal noise generators: these may amplify the signal or suppress the signal. The amplification/suppression of the signal at a location changes according to an inverse square law with the distance of the location from each noise generator. The result is that the ambient signal not only varies over time, but varies with location. This is illustrated in Figure 1, which shows a heatmap of a grid at $t = 0$, and the change of ambient signal over time from $t = 0$ to $t = 720$, with recalculation of the signal variation when $t \bmod 60 = 0$.

Value	Range
$maxX, maxY$	$\mathbb{N}^+$
base signal	$\mathbb{R}$
ambient signal	$\mathbb{R}$
signal amplifiers	$\{((x, y), S, I)\}$
signal suppressors	$\{((x, y), S, I)\}$
randomisation	$\{(A, f, \varphi)\}$

Table 1: Environmental specification.

The environment ϵ is defined by the data shown in Table 1. The base signal is the value which, if not distorted by signal amplifiers and suppressors, or the randomisation function, would be the ground truth for every location. The ambient signal is the base signal at a location (x, y) modified by the inverse-square of the distance to each of the signal amplifiers or suppressors according to its surface strength S and intensity I. The variation of the ambient signal is computed by the set of sine functions, each with a randomly-generated amplitude A, frequency f and phase φ.

Test Objectives

Such an environment is populated with a number of agents at random locations. Each agent has an imperfect sensor that can read the ambient signal at its location, and can communicate with the other agents in its social network. Their aim is to determine a common value, either by averaging individual values (the *wisdom of crowds*) or by exchanging signal readings and agreeing a value. The group therefore faces two possible situations:

- *informative*: collectively decide on the base signal given feedback from the environment about the accuracy of their ambient signals;
- *conformative*: collectively decide on the 'community' signal given only the previous collective decision(s).

Experimental Results

We have implemented a simulator for agent-based modelling of RTSI in Qu-Prolog[1], called QuRTSI (pronounced "curtsey"). Details are omitted here but essentially agents have to decide whether to sense the signal (do the work themselves) or ask the most trusted agent in their social network (i.e. delegate the task). Asked agents can also ask, up to a maximum hop-count of 6 (i.e. 6 degrees of separation). Agents' sensors are inherently unreliable, and there are three types of agent: those that consistently over-estimate, those that consistently under-estimate, and those that can do either.

Agent have 'trust' in other agents and 'self-confidence' in themselves, represented by a value in $[0..1]$. If an agent?s sensed signal is better than the signal of the one it asked, then it increases its self-confidence and decreases its trust in the other agent. It does the opposite if the asked agent's signal was better than its own.

Four experiments have been run to investigate:

- The effect of conformative- vs. informative-RTSI on signal-tracking. Hypothesis: that informative-RTSI will track the ambient signal, while conformative-RTSI will track the wisdom-of-crowds (WoC) estimate.
- The effect of different social networks on signal-tracking. Hypothesis: that informative-RTSI will track the ambient signal independent of network properties.
- The emergence of expertise. Hypothesis: that low error and high degree will be correlated with high trust.
- The trade-off between asking and sensing. Hypothesis: that agents will ask more often.

Experiment 1: Signal Tracking

In this experiment, we investigated the difference between informative- and conformative-RTSI with respect to the actual ambient signal, and the averaged individual estimates, i.e. the wisdom of crowds (WoC). The simulation was run for 720 clock 'ticks' with scramblers reset every 60 ticks, $\mid A \mid = 40$ (approximate size of a platoon), and maximum delegation hop-count $= 6$ (future experiments will investigate the effect of hop-count as an independent variable).

Figure 2 show the output for a typical run with a random network, $p = 0.25$. The blue line records the base signal. The black line records the aggregated and averaged sensed signal when each agent is sampling the signal on its own without the RTSI option of delegating to another agent. Individually, the agents are able to track the basic shape of the ambient signal, but the inaccuracies in their sensing, the distortions introduced by the signal amplifiers and suppressors, and population bias (in this case to over-estimate) means that the estimated signal is diverges from the ambient signal. Hence in this case the estimated signal values are higher than ambient, but can equally be lower.

The green line is the agents' collective signal estimate using informative-RTSI. The agents get feedback about the actual ambient signal, and update their self-confidence in themselves, or in the agent that they asked, according to the relative accuracy. As RTSI predicts, the estimate converges relatively quickly on the base signal and tracks that closely. An expert group emerges, with expertise on the base signal.

The red line is the agents' collective signal estimate using conformative-RTSI. Each agent gets feedback about the estimated signal from the previous round and compares their estimate, and the asked agent's estimate, relative to that, and update self-confidence and trust accordingly. As RTSI predicts, the collective estimate converges on the WoC signal and tracks that – in other words, an expert group emerges, but they are experts in expressing the population's opinion.

[1] http://staff.itee.uq.edu.au/pjr/HomePages/QuPrologHome.html

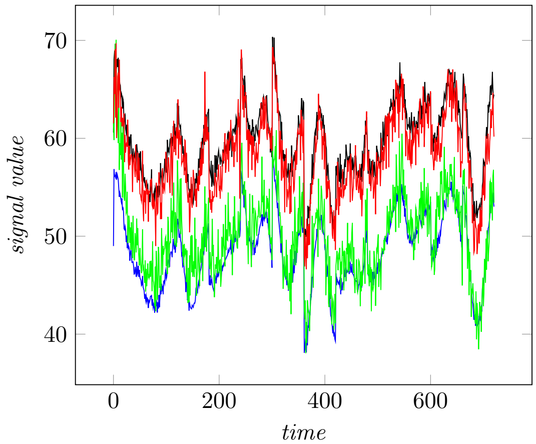

Figure 2: Signal tracking (single example) – Blue line: ambient signal; green line: informative-RTSI estimate; black line: WoC estimate; red line: conformative-RTSI estimate.

To verify the tracking of informative-RTSI against the ambient signal, and of conformative-RTSI against the wisdom-of-crowds, we run the simulator 20 times with same configuration, and measure the overall average percentage error. This is done for four cases: (1) Informative-RTSI vs. WoC; (2) Informative-RTSI vs. ambient signal; (3) Conformative-RTSI vs. ambient signal; and (4) Conformative-RTSI vs. WoC. The results are illustrated in Figure 3.

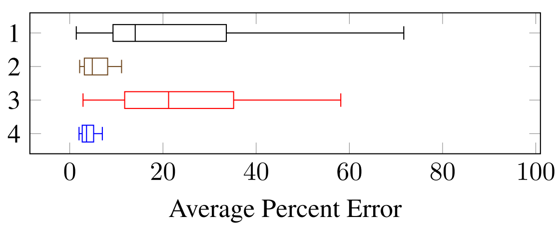

Figure 3: Percentage difference (20 runs). (1) Informative-RTSI vs. WoC; (2) Informative-RTSI vs. ambient signal; (3) Conformative-RTSI vs. ambient signal; (4) Conformative-RTSI vs. WoC.

This shows that informative-RTSI diverges from the WoC estimate, but tracks the base signal; while conformative-RTSI diverges from the base signal, but tracks the WoC estimate instead, supporting the experimental hypothesis.

Experiment 2: Network Variation

The aim of the second experiment was to examine the effect of different network types on informative-RTSI. We used six different types of network: ring, all-to-all (A2A), random network with $p = 0.33$ (ER.33) and $p = 0.66$ (ER.66), small-world (WS) network with $k_N = 3, p_w = 0.25$,[2] and scale-free (BA) network with $m_0 = 3, m = 2$.[3] Each of these were analysed with three different sets of agents, with cardinality of $\mid A \mid$ 20, 40, and 60. Each permutation (network vs. size) was run ten times and the percentage error of the informative-RTSI vs. the ambient signal was measured. Otherwise the same configuration as in Experiment 1 was used (maximum hop-count = 6, clock ticks = 720).

Note that using a pseudo-random number generator and setting the seed explicitly, the environment and the agent profiles were the same on each nth run in all of the eighteen combinations: only the social network changed. The results are shown in Figure 4.

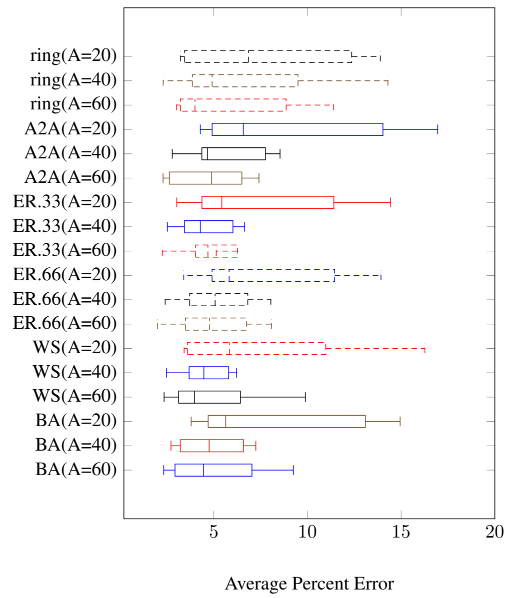

Figure 4: Informative-RTSI and Network Independence

Although the informative-RTSI estimate outperformed the WoC estimate in all cases, with a population size of 20, there is wider variation in accuracy of tracking. This can be attributed to the sparsity of coverage of, and fortune of location within the environment. Interestingly, the median performance of a ring network was on a par with other networks, but evidently there were some random distributions that even with a maximum hop-count of 6, made it harder for accurate information to percolate across the network. Even then, accuracy improved with larger group sizes.

However, with population sizes of 40 and 60, the "graph is essentially flat", the median and quartile ranges are very similar. This is evidence to support the hypothesis that tracking with informative-RTSI is independent of network type.

[2]The algorithm for generating a small-world network links each node to k_N nearest neighbours, and then rewires each link with probability p_w. See Prettejohn et al. (2011).

[3]The algorithm for generating a scale-free network starts with an all-2-all connected network of m_0 nodes and successively links new nodes to m existing nodes with a probability proportional to the existing node's number of links. See Prettejohn et al. (2011).

Experiment 3: The Emergence of Expertise

In this experiment, we investigated the relationship between network degree (the number of links to other agents) and self-confidence, using a scale-free (Barabási-Albert) network with 40 agents.

Recall that network generation algorithm starts with two parameters, m_0 and m, where $m0$ is the number of agents in the base network which are all connected to each other, and each new agent thereafter is connected to m other already-networked agents with a probability proportional to the existing agent's current link proportion. The result is that the network has an exponential degree distribution.

Here, we have used $m_0 = 3$ and $m = 2$. A first result for three separate runs over 720 cycles is shown in Figure 5, which plots self-confidence against the percentage individual error (deviation of sensed signal from ambient signal). The black triangles are the agents in m_0, the red triangles are the rest. This suggests that self-confidence is *not* correlated with network centrality.

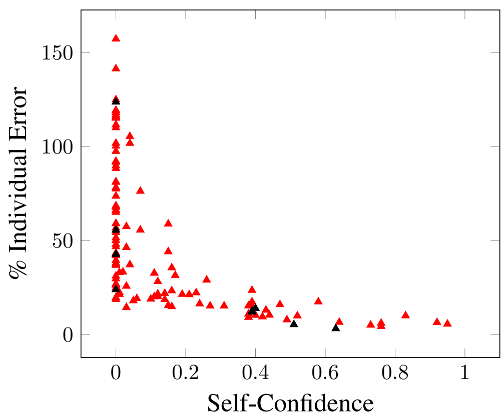

Figure 5: Informative-RTSI and Self-Confidence. Scale-free (Barabási-Albert) network, 3 runs, $| A |= 40, m_0 = 3, m = 2$. Black triangles: agents in m_0, red triangles: the rest.

A second result for the same three runs is shown in Figure 6. This plots (for the same agents) the 'trustworthiness' of an agent against its degree, where the trustworthiness of the agent is the average trust that the agents in one agent's social network have in it. While not correlated with high degree, generally agents with a high degree do have a high trustworthiness. However, as the second graph shows, the hub agents generally have the highest proportion of total trust in the system: it is evident where the 'social capital' in the system resides.

Trustworthiness appears to be correlated with degree only, pointing to the old adage "it is not what you know, but who you know". Furthermore, it brings us, in a sense, full circle: we started with knowledge aggregation in classical Athens and the importance of mechanisms for maintaining heterogeneity. These results suggest that even if expertise is not

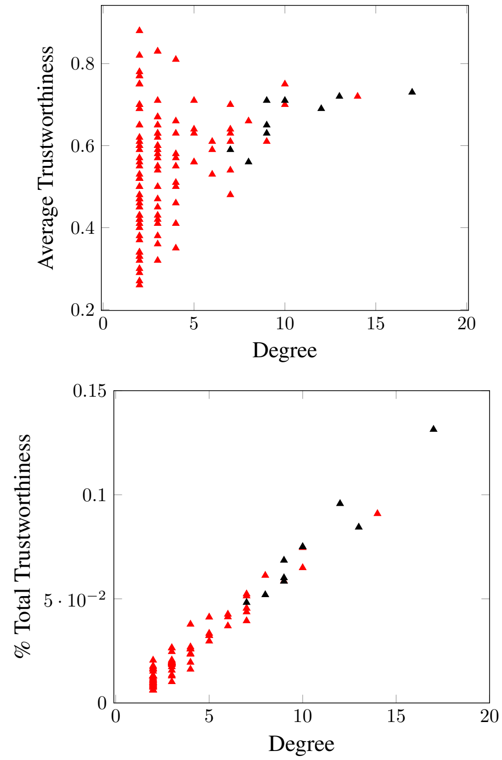

Figure 6: Informative-RTSI and Trustworthiness. Same experimental configuration as Figure 5.

well-connected, given adequate connectivity and an appropriate model of social influence, it is possible to access expertise for 'the whole' to know what 'the parts' know.

Experiment 4: The Cost/Saving of Expertise

In this experiment, we investigated the the ratio of sensing to asking, i.e. in any round, what was the ratio of the number of agents that sensed the environment, to the number of agents that asked another agent for its measurement of the environment. The rationale for this is that, in general, it is expected that the 'cost' of sensing and reasoning is greater than that of communicating.

For this experiment, we used a Erdos-Renyi (random) network with $p = 0.33$, a Watts-Strogatz (small-world) network with $k = 3$ and $q_k = 0.25$, a Barabási-Albert (scale-free) network with $m_0 = 3$ and $m = 2$, each with network sizes of 36, 64, and 100 agents. We were again testing informative-RTSI only.

One result is shown in Figure 7. This shows at each time point the *rolling average* for this and the previous 19 timepoints, of the number of agents that used their sensed value as opposed to asking a neighbour. This is a graph for a Barabási-Albert network, but bith Erdos-Renyi and Watts-Strogatz networks show the same profile. The again suggests that the outcome of informative-RTSI is largely independent of network topology. Moreover, the shape of the graph demonstrates that the agents are engaging in significantly more 'asking' than 'sensing' (an even split would have been a single horizontal line; more sensing than asking would be the mirror image of the graphs).

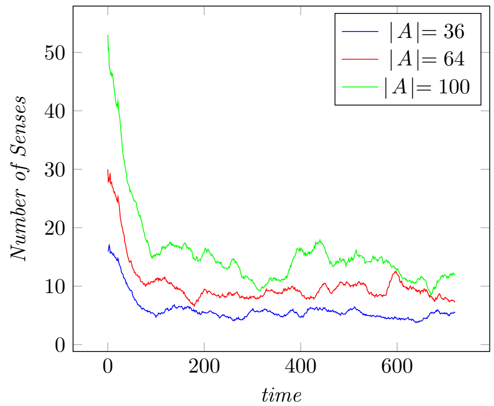

Figure 7: Sensing vs. Asking: Scale-free (Barabási-Albert) network, $m_0 = 3, m = 2$, rolling average of 20 time-points.

However, the particularly interesting feature is the convergence value, which suggests that only a few agents are doing most of the work. This observation needs to be pursued in further research, as well as looking at the 'amount' of asking that happens in a network, i.e. how many delegations occur before an agent uses its sensed value. This is of course linked to the value specified for the hop-count.

Summary and Conclusions

This inter-disciplinary work has brought together two lines of work: one from multi-agent systems and models of trust, opinion formation and self-organisation, the other from social psychology and models of social influence, dynamic networks and distributed information processing. We have used techniques from the former to specify and implement the first-pass of a new theory (RTSI) from the latter.

We have also implemented a multi-threaded Qu-Prolog simulator and used it to conduct some experimental investigations into what we distinguished as informative-RTSI, where the agents try to track a 'correct' signal, and conformative-RTSI, where the agents try to track a 'community preference' signal. Our analysis demonstrates that a collective of agents using RTSI in 'informative mode', the collective can identify expertise within it that enables it to become an effective 'multi-dimensional distributed information processing unit' that can outperform the 'wisdom of crowds' aggregate in tracking the 'true' signal (i.e. the base signal). However, in 'conformative' mode, the collective can identify expertise that enables it to more effectively track this same 'wisdom of crowds' aggregate, i.e. the 'community' signal.

There are several limitations in the current work. Firstly, for reasons of space, we have not situated this work fully in the literature from several disciplines (including complex systems, ALife, computational social choice, etc.) that has investigated information-seeking behaviour in social networks. This includes work from marketing Roch (2005); Van den Bulte and Joshi (2007), diffusion of innovations Nan et al. (2014), and social networks (Bonchi, 2011), as well as work aiming to quantify the power of influencers (Cha et al., 2010; Bakshy et al., 2011) and to understand their role in forming public opinion (Watts and Dodds, 2007; Jalili, 2013). The wisdom of crowd phenomenon has been the focus of several works, for example (Mendes et al., 2010; Yampolskiy and El-Barkouky, 2011). A particular issue is whether an asymmetry in crowd members' skills may lead to the situation when more skilled sub-crowds beat the wisdom of the whole crowd (Goldstein et al., 2014).

Although a distinction between these previous works and RTSI can be drawn, because the analyses are predominantly made from the perspective of the source of influence, we also think this first-pass implementation of RTSI has added to this literature: notably in the 'informative' and' 'conformative' modes of signal tracking, the independence of network topology, the correlation of network centrality and trust, and the distribution of effort across the network. However, the second limitation acknowledges that this is a first-pass: RTSI is a much richer theory than has been implemented here, and more of its features need to be implemented and explored in QuRTSI.

A third limitation is also an opportunity. We used sine waves to try to generate a signal that could not easily be learnt. However, an alternative would be to use an n-bit binary signal, whose proximity to a scrambler would increase the likelihood of a bit switching. It might then be possible to analyse the properties of transmission and processing using information theory, and develop an information-theoretic basis of RTSI. This might also enable an alternative approach to experimentation: here we have set up a model based on RTSI and explore how it behaves, but with this approach experiments could be designed to (attempt to) falsify predictions made by the theory.

However, beyond these theoretical and practical results and limitations, the primary contribution of this paper is twofold. The first is to have demonstrated how a theory of human psychology and human behaviour can be represented in computational form; while the second is to have shown that agent-based modelling can be used to animate that theory of human psychology and shed some interesting light on knowledge management in organisations.

Acknowledgements

This research was supported by Narodowe Centrum Nauki NCN DEC-2011/02/A/HS6/00231. Thanks to the anonymous reviewers for many encouraging and helpful remarks.

References

Allport, G. W. (1937). *Personality: A psychological interpretation*. Henry Holt & Co.

Asaro, P. (2019). Ai ethics and predictive policing: From models of threat to an ethics of care. *IEEE Technol. Soc. Mag.*, 38(2):40–53.

Asch, S. E. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9):1.

Bakshy, E., Hofman, J. M., Mason, W. A., and Watts, D. J. (2011). Everyone's an influencer: quantifying influence on twitter. In *Proc. 4th ACM Int. Conf. on Web Search and Data Mining*, pages 65–74.

Baldwin, J. R. and Hunt, S. K. (2002). Information-seeking behavior in intercultural and intergroup communication. *Human Communication Research*, 28(2):272–286.

Bonchi, F. (2011). Influence propagation in social networks: A data mining perspective. *IEEE Intelligent Informatics Bulletin*, 12(1):8–16.

Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, K. P. (2010). Measuring user influence in twitter: The million follower fallacy. In *Proc. 4th AAAI Conf. on Weblogs and Social Media*.

Cialdini, R. (2016). *Pre-Suasion: A Revolutionary Way to Influence and Persuade*. Simon & Schuster.

Davenport, T. and Prusak, L. (1998). *Working knowledge [electronic resource] : how organizations manage what they know*. Harvard Business School Press.

Dyer, L. M. and Ross, C. A. (2008). Seeking advice in a dynamic and complex business environment: Impact on the success of small firms. *Journal of Developmental Entrepreneurship*, 13(02):133–149.

Goldstein, D. G., McAfee, R. P., and Suri, S. (2014). The wisdom of smaller, smarter crowds. In *Proc. 15th ACM conference on Economics and Computation*, pages 471–488.

Hussain, A., Ali, S., Ahmed, M., and Hussain, S. (2019). The anti-vaccination movement: A regression in modern medicine. *Cureus*, 10(7):e2919.

Jalili, M. (2013). Effects of leaders and social power on opinion formation in complex networks. *Simulation*, 89(5):578–588.

List, C. and Goodin, R. (2001). Epistemic democracy: Generalizing the condorcet jury theorem. *Journal of Political Philosophy*, 9(3):277–306.

Macbeth, S. and Pitt, J. (2015). Self-organising management of user-generated data and knowledge. *Knowledge Engineering Review*, 30(3):237–264.

Mendes, P. N., Passant, A., and Kapanipathi, P. (2010). Twarql: tapping into the wisdom of the crowd. In *Proc. 6th International Conference on Semantic Systems*, pages 1–3.

Milgram, S. (1963). Behavioral study of obedience. *The Journal of abnormal and social psychology*, 67(4):371.

Nan, N., Zmud, R., and Yetgin, E. (2014). A complex adaptive systems perspective of innovation diffusion: an integrated theory and validated virtual laboratory. *Computational and Mathematical Organization Theory*, 20(1):52–88.

Nowak, A., Vallacher, R., Rychwalska, A., Roszczyńska-Kurasińska, M., Ziembowicz, K., Biesaga, M., and Kacprzyk-Murawska, M. (2020). *Target in Control: Social Influence as Distributed Information Processing*. Springer.

Ober, J. (2008). *Democracy and Knowledge*. Princeton Univ. Press.

Petty, R. and Cacioppo, J. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19:123–205.

Prettejohn, B., Berryman, M., and McDonnell, M. (2011). Methods for generating complex networks with selected structural properties for simulations: a review and tutorial for neuroscientists. *Front. Comput. Neurosci.*, 5(art.11):1–18.

Roch, C. H. (2005). The dual roots of opinion leadership. *The Journal of Politics*, 67(1):110–131.

Sieloff, C. (1999). "if only hp knew what hp knows": the roots of knowledge management at hewlettpackard. *Journal of Knowledge Management*, 3(1):47–53.

Van den Bulte, C. and Joshi, Y. V. (2007). New product diffusion with influentials and imitators. *Marketing science*, 26(3):400–421.

Watts, D. J. and Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of consumer research*, 34(4):441–458.

Yampolskiy, R. V. and El-Barkouky, A. (2011). Wisdom of artificial crowds algorithm for solving np-hard problems. *International Journal of Bio-inspired computation*, 3(6):358–369.