Geoinformatik

TRUST & REPUTATION MODELS FOR HUMAN SENSOR OBSERVATIONS

Inaugural-Dissertation zur Erlangung des Doktorgrades der
Naturwissenschaften im Fachbereich Geowissenschaften der
Mathematisch-Naturwissenschaftlichen Fakultät der Westfälischen
Wilhelms-Universität Münster

vorgelegt von
Mohamed Bishr
aus Kairo, Ägypten
im März 2011

TRUST & REPUTATION MODELS FOR HUMAN SENSOR OBSERVATIONS

MOHAMED BISHR

Doctoral Degree
March 2011

`Datum der mündlichen Prüfung: 13.05.2011`

Mohamed Bishr: *Trust & Reputation Models for Human Sensor Observations,* Doctoral Degree, © March 2011

DEKAN:
Prof. Dr. Hans Krep
SUPERVISORS:
Prof. Dr. Werner Kuhn
Prof. Dr. Matt Dukham

LOCATION:
Münster, Germany

TIME FRAME:
March 2011

To the fine people who put their life on the line to free our country -
Egypt!

ABSTRACT

The emergence of Web 2.0 and its hallmark phenomenon - user generated content - extended to the field of GIScience. User generated content is exemplified in GIScience through Volunteered Geographic Information (VGI). Users who were traditionally consumers of geospatial information are now volunteering geospatial information on a large scale by acting as sensors to their environment. However, this comes at a price. Inherent in VGI is a lack of knowledge about the traditional GI quality criteria. Also, volunteers are not necessarily experts, have different backgrounds, and varying motivations to contribute data. Therefore, not all provided information is valuable, and may even be fraudulent or misleading. This research proposes a trust and reputation based approach to develop quality assessment methods for human sensor observations. The approach presented in this thesis has two aspects.

First, we provide an ontological account of trust and reputation as measures for quality assessment of human sensor observations. The ontological account is based on earlier work on a functional ontology for observation and measurement for the sensor web that is grounded in the DOLCE foundational ontology. We extend the ontology with trust and reputation and implement a use case scenario of the human sensor web for water monitoring. Our ontological account shows how to integrate trust and reputation models as quality measure in the human sensor web. The ontological account is developed as an algebraic specification in the Haskell functional language which acts also as a simulation for the developed use case.

Second, we develop the notion of *informational trust* as being mediated by interpersonal trust to enable trusting information. We then propose the novel paradigm of spatially and temporally sensitive informational trust and reputation models. We develop two different computational models for triage, filtering and quality assessment of VGI and human sensor observations based on our spatial-temporal informational trust and reputation. The first model is temporally sensitive where trust in human sensor observations decays over time. The second model is spatially sensitive where the spatial contexts of the volunteers and the observations are explicitly integrated in the model's computations. We develop agent based simulation models to simulate the human sensor web for water monitoring use case and apply our computational models. Our analysis shows that the developed models are effective tools for triage, filtering and quality assessment of human sensor observations. Furthermore, the integration of space and time in our trust and reputation models has a positive impact on the models' performance.

ZUSAMMENFASSUNG

Die Entstehung des Web 2.0 und sein Hauptmerkmal - Nutzer-generierte Inhalte - reichen bis in die Geoinformationswissenschaften. Nutzergenerierte Inhalte werden in den Geoinformationswissenschaften durch so genannte Volunteered Geographic Information (VGI), d.h. freiwillig zur Verfügung gestellten Geoinformationen, verkörpert. Nutzer, die traditionell nur als Konsumenten von Geoinformationen auftreten, stellen nun selbst Geoinformationen zu Verfügung und betätigen sich in großem Umfang als Sensoren ihrer Umwelt. Dieses Phünomen hat jedoch seinen Preis, denn Qualitätskriterien herkömmlicher Geoinformationen fehlen den VGI. Weiterhin sind freiwillig Beitragende oftmals keine Experten auf dem Gebiet, haben verschiedene Hintergründe und Erfahrungen so wie unterschiedliche Beweggründe Daten bereitzustellen. Deshalb ist nicht jeder Beitrag von VGI wertvoll, und kann sogar irreführend oder falsch sein. Die hier vorgestellte Forschungsarbeit schlägt einen Ansatz basierend auf Vertrauen (Trust) und Reputation (Reputation) der Beitragenden vor um die Qualität von Beobachtungen von menschlichen Sensoren zu beurteilen. Der in dieser Arbeit vorgestellte Ansatz hat zwei Aspekte.

Zuerst werden die Begriffe Vertrauen und Reputation als Maße für die Bewertung der Qualität von Beobachtungen menschlicher Sensoren ontologisch beschrieben. Die ontologische Beschreibung basiert auf früheren Forschungsarbeiten zu einer in der Basisontologie DOLCE verankerten funktionalen Ontologie von Beobachtungen und Messungen (functional ontology for observations and measurements) für Sensornetzwerke. Diese Ontologie wird um Vertrauen und Reputation erweitert und ein Anwendungsfall eines menschlichen Sensornetzwerks zur Wasserüberwachung umgesetzt. Die ontologische Beschreibung verdeutlicht wie Modelle von Vertrauen und Reputation als Maße für die Qualität in menschlichen Sensornetzwerken eingesetzt werden können. Die ontologische Beschreibung wurde in Form einer algebraischen Spezifikation in Haskell entwickelt und dient gleichzeitig als Simulation des Anwendungsfalls.

Als zweites wird der Begriff des *informational Trust* entwickelt, welcher vom zwischenmenschlichen Vertrauen in den Informationsanbieter abstrahiert. Daraufhin wird ein neues Paradigma für raum- und zeitsensitive Modelle von Vertrauen und von Reputation vorgestellt. Basierend auf raum- und zeitsensitiven *informational Trust* und Reputation werden zwei Computermodelle zur Sichtung, Filterung und Qualitatsbewertung von VGI und menschlichen Beobachtungen entwickelt. Im zeitsensitiven Modell nimmt das Vertrauen mit der Zeit ab. Im raumsensitiven Modell werden der räumliche Kontext des Beitragenden und der Beobachtung explizit im Computermodell berücksichtigt. Eine Analyse der Computermodelle zeigt die Effizienz dieser Modelle als Werkzeuge zur Sichtung, Filterung und Qualitätsbewertung von menschlichen Beobachtungen. Des weitern hat die Berücksichtigung von Raum und Zeit in den vorgestellten Modellen für Vertrauen und Reputation einen positiven Einfluss auf die Performanz der Modelle.

> *Twenty years from now you will be more disappointed,*
> *by the things that you didn't do than by the ones you did.*
> *So throw off the bowlines. Sail away from the safe harbour.*
> *Catch the trade winds in your sails.*
> *Explore. Dream. Discover.*
>
> — *Mark Twain*

ACKNOWLEDGMENTS

Phew! thank god this thing is finished, seriously! I mean it seemed like impossible until very recently, but then again I surprise me!! There are so many people to thank, and I also would like to thank me ... you know! I was at the center of it but for some reason people think those who thank themselves are narcissistic, which they probably are.

So many people tried to stop me from finishing this thesis. My first supervisor, Werner Kuhn did an excellent job of making the work environment enjoyable. From projects to lavish retreats and interesting research insights, it all meant I didn't want to finish. He also helped me on a personal level many many times. Finally, when all his attempts to derail the thesis failed he helped me find an excellent job and this was the last nail in the coffin of the thesis but I still won and I finished it! Thanks Werner!

I also would like to thank Matt Duckham for accepting to be my second supervisor and for his insight at an early stage of the thesis which prompted me to find a proper use case. I think this was his attempt to derail the thesis, but I still found the use case and I still won!

I also would like to thank my brother and mentor Yaser Bishr for his exceptional scientific and moral support at the most critical phase of this thesis. His attempts to push me to the limit almost ruined me and could have fatally derailed the thesis, but once more, I won! Thanks Yaser!

Several other friends at the Institute for Geoinformatics have also hopelessly taken their shot at my thesis. I want to personally mention Carsten Kessler who made reviews of the last chapters and Krzysztof Janowicz who gave personal and professional advise when it was absolutely not needed. Imaduddin Humayun found countless mistakes and errors through painstaking review of my thesis, and he thought this would frustrate me but it didn't. Jens Ortmann has also dragged me into endless partying and we all know what this meant. Alas, they have all failed to derail the thesis and I'm deeply indebted to everyone of them. I would also like to thank Caroline Wahle for her unconditional friendship and support. We have had so many expensive dinners together. I'm almost bankrupt! Not to mention the time spent with her which I could have spent working. I could have finished years ago without her support.

I also thank the folks at the netlogo support groups and also Tamás Máhr for his excellent analytical skills and expertise. Without them all, life could have been placid.

I would like to thank my Mom and Dad for their support. They have provided endless moral, spiritual and financial support so I can live comfortably and worse still drive an exotic convertible and spend as

much time out of office as I could. Their efforts to derail my thesis were genuine and inventive. I know they are proud I have reached this stage despite their attempts otherwise, and for this I'm deeply grateful.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LISTINGS

1

INTRODUCTION

A paradigm shift in how people cope with the information overload is underway. People rely more on online social network relations and their artefacts such as trust for information discovery, and tasks traditionally reserved for search engines like Google or Yahoo. Noticeable also is how Google as a search engine is tapping into the social web by enabling user feedback on search results. On the social web where users are information producers as much as they are information consumers [10] issues of reputation and trust become essential in helping people make judgements about the quality and relevance of information produced by other users. This thesis focuses on a specific aspect of user generated content, namely on Volunteered Geographic Information (VGI) and humans as sensors [46]. Particularly we address the problem of the quality of human sensor observations characterized by proneness to errors or fraud and the lack of traditional Geospatial Information (GI) quality criteria [10]. We present a novel approach using trust and reputation as proxy measures of observation quality, particularly for human sensor observations.

The concept of social networking is evolving and morphing. It's now about making the entire Web social instead of just creating a ghetto of destination sites where people have to go to socialize. [81]

1.1 MOTIVATION AND PROBLEM

In this section we present our research motivation through a brief discussion of VGI applications leading to the identification of our research problem. We argue that the information overload resulting from the massive flow of VGI, leads to an information quality problem that is at the heart of reaping the full benefits of VGI.

1.1.1 *Research Motivation*

The next generation social web will be a large-scale social platform and location will be ubiquitous. It is estimated that 46% of mobile web users use the mapping functionality of their mobile device [81]. Current emerging applications increasingly allow users to provide real time location information, such as their own location, meet friends in proximity, filter comments by geography or discover and create information with strong spatial and temporal affinities. For example FourSquare[1] shown in figure 1 is a location based social network application that is gaining over a 100,000 new service subscribers per week[2].

Furthermore, the massive community effort following the Haitian earthquake produced invaluable VGI that helped shape the relief effort. As one example, the Ushahidi[3] project. The platform allows users to submit reports through SMS, MMS, email and an online web interface. Figure 2 shows a geotagged report during the Haitian earthquake crisis. The ability of Ushahidi to receive alerts from volunteers on the ground using the cheap and ubiquitous SMS technology meant that a massive influx of reports flooded the system during the crisis. Furthermore,

---

[1] http://www.foursquare.com/
[2] http://techcrunch.com/2010/06/22/foursquare-growth/
[3] http://www.ushahidi.com/

Figure 1: *FourSquare smart phone application locating a user. FourSquare is a location based social networking service. The phone's GPS sensor is used to log users' locations. For example, users can tag photos, leave notes and connect to friends in the location based social network.*

rescue messages were also sent by volunteers on the ground through Twitter[4], the micro blogging service, where other volunteers in various parts of the world translated the messages from Creole, geotagged them and placed them on the Ushahidi platform. In addition, the U.S. state department intervened to assist in the geolocation and verification of the massive flow of messages to assist in delivering trustworthy information to the Red Cross and the U.S. Coast Guard [58]. The Ushahidi platform and the entire VGI effort had a significant impact on the effectiveness of the relief effort. Ushahidi was hailed by army officials involved in the relief effort as having saved lives every day [103]. Ushahidi is centred around real-time geotagged reports that contain actionable messages and has been widely implemented for different purposes beyond disaster response[5].

The success of VGI, and the potential of humans as sensors resulted in a massive information overload. How can we effectively manage, curate and verify VGI is the motivation of this thesis.

1.1.2 *Problem*

The success of VGI applications comes with its own challenges. The massive effort to manage, curate and verify the flow of VGI in the Ushahidi platform succeeded in producing positive impact on the ground. Yet, the VGI generated remains massively under utilized which prompted a spin off project to Ushahidi called SwiftRiver[6]. SwiftRiver aims to develop a system to filter and verify real-time human sensor observations by authority and accuracy across different channels akin

---

[4] http://www.twitter.com/

[5] e.g. a new deployment is planned by activists for election monitoring in Egypt. http://crisismapper.wordpress.com/2010/11/20/ushahidiegyptwhenopendata-isnotsoopenorwhenpeoplejusdon'tgetit/

[6] http://swift.ushahidi.com/

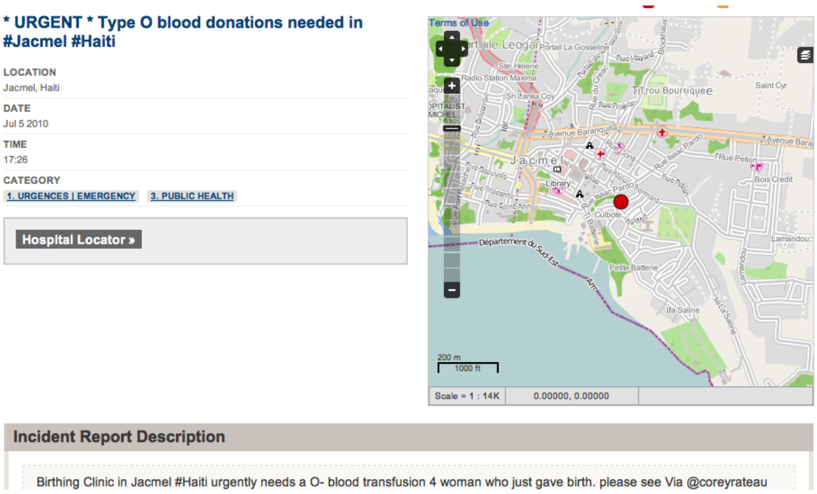

*** URGENT * Type O blood donations needed in #Jacmel #Haiti**

LOCATION
Jacmel, Haiti

DATE
Jul 5 2010

TIME
17:26

CATEGORY
1. URGENCES | EMERGENCY 3. PUBLIC HEALTH

Incident Report Description

Birthing Clinic in Jacmel #Haiti urgently needs a O- blood transfusion 4 woman who just gave birth. please see Via @coreyrateau

Figure 2: *A report instance from Haiti on the Ushahidi platform. The report is geo-tagged with the location of the incident where help is required along with a short description.*

to the Ushahidi channels (i.e. SMS, Twitter, etc.). Another challenge during the crisis was the integration of human sensor observation from different sources[103]. This problem resembles the semantic problems associated with integration of sensor observation in the sensor web, but it also adds the challenge of identifying ways to assess and control the quality of VGI. Verifying the authority and accuracy of VGI can be viewed through the perspective of information quality, where quality is viewed as fitness for purpose.

As it stands today, attempts to judge human observations in general are based on user feedback by people who have consumed these observations (e.g., in digg.com or outalot.com). People discover observations posted by other users and might elect to experience those recommended observations. They would then give a rating quantifying their experience as positive or negative. We notice here that there is no mechanism by which the user can judge an observation prior to consenting to use it. With an unprecedented flow of human sensor observations, inaccurate, misleading, outdated or even fraudulent observations become a problem.

A volunteer may report an observation about a traffic jam, or a roadblock due to construction or a restaurant location and review. First, we notice that this observation varies with respect to sensitivity to the temporal dimension, i.e. relevance and value of a traffic jam observation probably decays faster than a roadblock observation, which in turn decays faster than a restaurant review. Second, we notice that the location of the person giving this information should affect how accurate or useful it is. This recalls the notion of activity or familiarity spaces [52]. If observations are within those users' spatial activity space, should it not be more trusted? Such that a local or a frequent visitor might be more trusted to give useful and otherwise difficult to get observations, than infrequent visitors. Social network research shows that network structure and dynamics affect trust relationships and user reputations on the network. Also, social networks online and offline have strong geographic affinities, and geography affects network structure and dynamics (see e.g. [71]). It follows then that

"Trust is the glue that holds everything together, the bond that creates healthy communities" [85]

geography should have an effect on how people form and maintain trust relationships in social networks. Some research tried to tackle the issue of the temporal nature of the phenomenon of trust (see e.g. [68]), where trust is formalized to develop or decay overtime. Research has also suggested indirect relations between space and trust, for example [78] suggested that companies in closer proximity are always preferred partners implying a higher degree of trust. In [17], distance was used as a surrogate for social network density, and the research asserts that the higher the network density, the higher the overall trust on a network. However, we are unaware of research that tried to investigate and utilize the effects of both space and time on trusting relationships in trust and reputation models. This may directly affect how much an observation provided by some volunteer should be trusted by a potential user of this observation.

Hence there is potential in tapping into the space and time dimensions to build effective trust and reputation models to solve the quality problems inherent in VGI. The notion that trust can be used as a proxy measure of observation quality [10], is novel and understudied. VGI observations are spatial and temporal in nature - by extending the notion of trust with spatial and temporal dimensions we can develop trust models that can be used for filtering and triage of VGI and provide higher quality information to consumers. We propose a novel paradigm of quality assessment of human sensor observations through integrating space and time in trust and reputation models for human sensor observations.

1.2 RESEARCH QUESTIONS AND HYPOTHESIS

We make a trifurcate observation; firstly, it is that VGI human sensor observations have strong spatial and temporal affinities. Secondly, current trust models are static in nature - they rely on user ratings and do not take into account the space and time dimensions. Finally, to our knowledge, there are no trust models applied on or developed for human sensor observations. We postulate that trust and reputation models that are spatially and temporally sensitive will be adequate quality measures for VGI in particular and the social web in general. Specifically we hypothesise that:

Spatially and Temporally Sensitive Trust Models are Suitable Proxy Measures for VGI Quality

We pose the following research questions stemming from the hypothesis:

What is trust?

As we discuss in this thesis, one of the problems of studying trust is the loaded nature of the term. Researchers argue that there are many, sometimes contradicting, understandings of trust because each one of us sees trust differently based on our own experiences, but maybe most importantly because there are indeed that many different types of trust. Before attempting to use trust in any research endeavour, it is prudent to clarify what one means with the term. Sub-questions we raise here are:

A. What is the definition of trust this thesis subscribes to?

B. What types of trust matter to our proposal of proxy measures of observation quality?

C. Is there evidence for space and time effects on trust and how to employ them in our proxy measures proposal?

How can we integrate trust as a proxy measure of quality of human sensor observations into the sensor web?
Our vision is that human sensors and their observation data will seamlessly integrate with other sensory assets of the sensor web as technologies converge in the future. Attempts to introduce human sensors into sensor web observation and measurement ontologies is presented in [60, 79]. Sub-questions we raise here are:

A. What is the ontological nature of both trust and reputation when used as a proxy measure of human sensor observation quality?

B. How can volunteers of VGI be integrated as sensors in sensor web ontologies?

C. How can trust and reputation models be integrated in the sensor web ontologies?

How can we develop spatio-temporal computational models of trust and reputation for quality assessment of human sensor observations?
Sub-questions we raise here are:

A. Are computational models of trust and reputation useful for quality assessment of human sensor observation?

B. Does accounting for the spatial and temporal nature of VGI and volunteers impact performance?

1.3 RESEARCH APPROACH

In the first phase of research we establish a common understanding of trust that is used throughout this research. As we mentioned earlier, one of the problems of researching trust is that everyone experiences trust. Thus each one has his own first hand experience with trust and their own definition of what trust is [45]; this could be one reason why there are different and sometimes contradicting definitions of trust. In [26, 87] another reason is provided which is that there are many definitions of trust simply because there are that many types of trust. Both reasons are true to a large extent, with intuitive understanding: trusting you to drive my car, does not mean trusting you to repair my car; trusting you to manage my financial portfolio does not entail trusting your advice for my career progress. It is also different between different disciplines, where philosophical or sociological meanings of trust do not necessarily explain behaviour of trust game participants in economics. Thus we aim to define what trust means in the context of being a proxy measure of information quality. We also mine earlier work for understanding of the spatial and temporal aspects of trust, on which there is scarce literature. This phase leads us to establish our understanding of trust as dealt with in this research, followed by identification of key characteristics of our understanding of trust that will guide the formalization of our approach for proxy measures of quality.

A precise definition of trust in the context of our research as a proxy measure of observation quality is essential before any attempt at further discussion

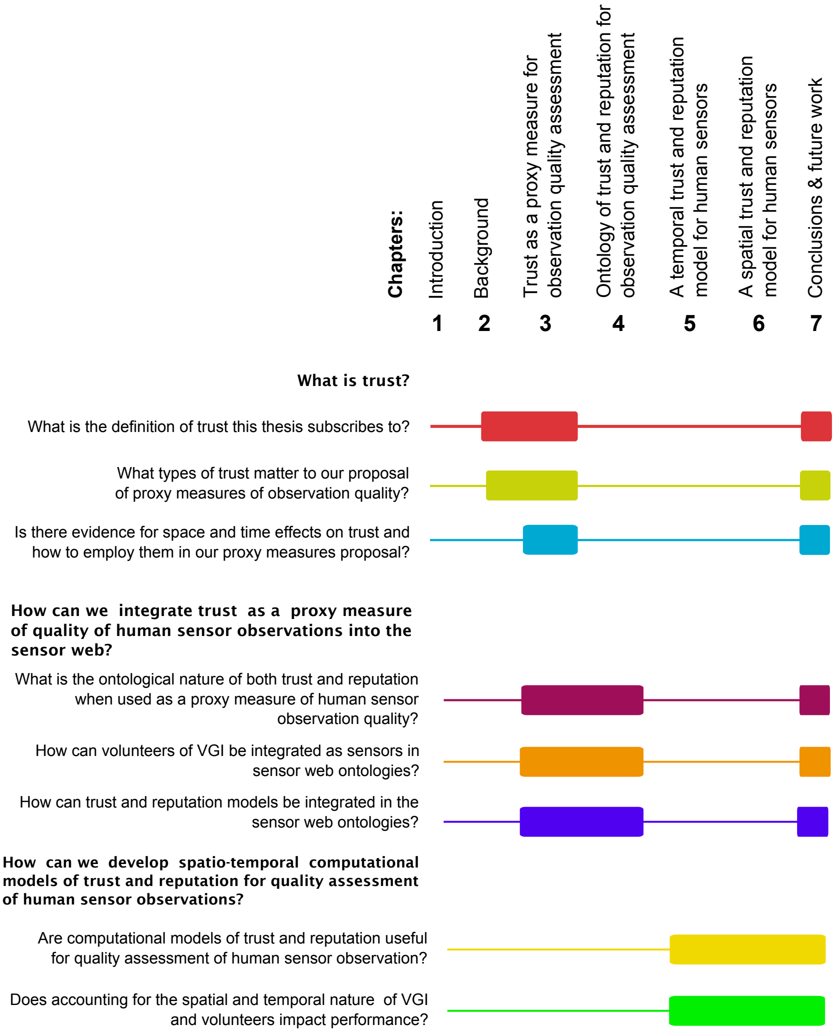

Figure 3: *Research questions across thesis chapters.*

In the second phase we amalgamate our findings on the phenomenon of trust and provide an ontological account of trust and reputation as proxy measures for human sensor observation. We extend the work on the sensor web observation and measurement ontology [60, 79] by formalizing interpersonal trust and our proposed notion of informational trust. We also introduce the ontological account of trust and reputation models within the same ontology. This work represents the general framework of how the computational models developed later in the thesis can be integrated along with human sensors in the sensor web. The final phase of this research is the development and evaluation of novel spatial and temporal computational trust and reputation models for quality control and triage of human sensor observations. We will use agent based simulations to subject the models to testing and evaluation. These two phases of the thesis use the human sensor web use case of the H2.0 project, presented later, as our use case scenario.

As a guiding principle in this thesis, we subscribe to the view that it is common when embarking on research in a new area to oversimplify in favour of reaching useful insights. In this research we are developing a novel paradigm of trust and reputation models that are spatially and temporally sensitive. In order to achieve this, we address implications of space and time on trust for our computational models using the simplest set of assumptions wherever feasible. Any simplifying assumptions we make are clarified where appropriate in the thesis.

"When working in a very new area, it is entirely forgiveable to make outrageous simplifications in pursuit of insights, with the faith that the model can be brought closer to facts on later passes" [59, p.39]

1.4 CONTRIBUTION AND RELEVANCE

The main contribution of this research is developing a novel approach for the curation, triage and quality assessment of human sensor observations. The proposed approach stems from our proposal of trust as a proxy measure for observation quality. The three contributions in this thesis as exemplified by our approach are:

- *Ontological Perspective:* We develop an ontological account of trust and reputation as quality assessment measures for human sensor observations grounded in the DOLCE foundational ontology. We use the developed ontology to implement a use case of the human sensor web and show, first how human sensors are integrated with other sensory assets. Second, show how trust and reputation are modelled as relational qualities. Third, we show how computational models of trust and reputation are integrated within the sensor web as sensors.

- *Computational Models Perspective:* We develop, implement and evaluate computational models of trust and reputation. Our analysis shows the resilience of our models to malicious behaviour and their ability to consistently deliver accurate observation results consistent with the true value of the observed phenomenon.

- *Spatial and Temporal aspects of Trust:* Our computational trust and reputation models are novel contributions to the field of VGI in their own regard. However, a major contribution is the proof that the spatial and temporal sensitivity of trust and reputation models are detrimental in improving model performance. We believe this finding will have far reaching effects on the future of trust and reputation models for VGI in particular but also in general.

The three contributions of our approach constitute a novel strand of research on the quality of VGI. As discussed in the motivation section 1.1.1, the potential of VGI and humans as sensors is just beginning to be realized. The challenge of managing the information overload resulting from the influx of human sensor observations hinders the full potential of VGI and our research directly addresses this problem.

1.5 THESIS OUTLINE

This thesis is divided into 7 chapters. In chapter 2 we present an interdisciplinary literature review covering different aspects of trust to lay the foundations for the assumptions made in this research concerning trust. In chapter 3 we consolidate the findings of our literature review and present a coherent picture of our formal proposal for spatial and temporal trust and reputation models as proxy measures of the quality of human sensor observations.

In chapter 4, we build on the line of thought developed in chapters 2 and 3 with a rigorous ontological account of how to integrate trust, reputation and their computational models into the sensor web, for human sensors.

In chapter 5 and chapter 6 we present two different temporal and spatial computational trust and reputation models. The viability of the models developed is evaluated using agent based simulations followed by analysis to show how such models can be effective measures of observation quality. In these chapters we also provide evidence for the impact of spatial and temporal sensitivity of trust and reputation models on improving model performance.

At the end, chapter 7 presents our main research conclusions, summarizes our findings and discusses future research directions. This chapter consolidates the contribution of this research and lays the foundations of a new paradigm of spatial and temporal trust and reputation models for the social web in general.

2

TRUST: BACKGROUND AND CHARACTERIZATION

In this chapter we present a background investigation of trust and other topics surrounding our research such as the sensor web and human sensors. We establish the understanding of what we mean by trust in this research, and ground our work in its theoretical foundations.

2.1 TRUST: A MULTIDISCIPLINARY OVERVIEW

Trust is a widely studied phenomenon in various disciplines like sociology, political science, economics, philosophy and computer science [17, 34, 36, 39, 92, 10]. Computer science has taken a multidisciplinary approach to the problem and adopted various trust definitions for models of trust. This includes models of social trust in online social networks, recommender systems and artificial intelligence. Several authorities on the topic of trust provide various reasons as to why we study trust. Trust provides a mechanism to understand or at least accept the complexity of our society [65], underlies the ability to cooperate [26] and improves performance of task accomplishments [75]. Societies rely heavily on trust to mediate transactions on a wide scale [34] and online communities are no exception [42].

One of the problems occurring when studying a notion like trust is that everyone experiences trust. Each one has a first-hand experience with trust, and hence a personal view of what trust actually is [45]. This is the first intuitive explanation of why trust has multiple and varying definitions. A second explanation is fact that there are multiple definitions of trust simply because there are that many types of trust [26, 87]. It is as such essential to establish a common understanding of what one means by trust, or of which trust we are talking before embarking on any research endeavour on the topic. In the following sections, we discuss different definitions of trust to establish the common understanding necessary for this work.

Trust is a difficult notion to study. For one, each person has a personal view on what trust is, and also because there are so many types of trust

2.1.1 *Trustworthiness and Credibility*

Credibility of information has been widely studied in different contexts [20, 32, 72, 73]. Different fields have discussed credibility in different ways. Sociology, communication science, and information science [4, 83, 90] have all addressed the issue of credibility, often leading to contradicting views on credibility and its effects. This contradiction was primarily due to the fact that each field arrived at the topic with its own presuppositions, assumptions and areas of interest [82].

In information science credibility research has focused on how information consumers assess information believability and to some degree the perception of information quality [63, 84]. Here credibility has been addressed as a measure of relevance of judgement on information *believability*, where believability is the quality of being believable as defined in common English dictionaries [4, 5]. Within our research this view on credibility bears resemblance to how we view trust, hence causing a confusion that calls for a discussion on the subject.

Expertise is also viewed in the same context as an integral component of credibility. For one to say a source is credible, this source must be viewed as an expert on the subject matter [82]. However, this view on expertise is again problematic. Expertise is difficult to define itself, and asserting that trusting someone implies understanding that he is an expert on the subject matter is a dubious generalization. For example, trusting some restaurant review, does not necessarily entail that the reviewer is an expert in reviewing restaurants. It is true he must have experienced the restaurant, but his review remains a subjective opinion, and this is understood by the reader of the review. The point we want to make is that the authority of the reviewer is much less related to his expertise than it is to his reputation as we later discuss.

Credibility of information entails trusting the person, this decision to trust is more affected by reputation than by expertise

Finally, credibility is a loaded term and much of the discussions on credibility or its components rendered above, do not formally define what is credible, is it people? or information? and use both interchangeably. In essence, where credibility inheres and how the term should be used is not formally established. The lack of a formal description of credibility, its components, the relations between them and information or people makes building any models utilizing credibility infeasible. Some researchers have proposed that credibility is composed of trust and expertise [82]. Although the notions of trust and credibility have been sometimes used interchangeably, they remain distinct in our view. Concepts like accepting advice and information believability are more a reference to credibility than they are to trust [32]. One can argue to make a distinction, namely that credibility is a property of the information, while trust is a property of the people behind this information, which can then be transferred to the information making it trustworthy or untrustworthy. This understanding leads us to conclude that considering trust a component of credibility is a reasonable simplification, yet it entangles the two concepts in a way that makes both hard to study given the distinction we make. This lack of clear ontological commitments with respect to credibility is of crucial importance when studying another loaded term such as trust in this thesis. Our work in chapters 3 and 4 lays the formal foundations of what we mean with trust and its ontological nature for this thesis.

Credibility is a property of information, while trust is a property of the people

2.1.2 *Trust, a sociological, psychological and philosophical perspective*

Two strands of sociology characterize the field. The first dominated the scene in sociology until the mid of the 20th century and focused on societal whole, complex structures and social systems. The second strand of sociology started gaining momentum by the mid of the 20th century, and focused on societal members and individual actions. This made apparent the importance of trust as an element emerging from individual interactions and based on individual actions [92]. Humans in this second sociology need to rely on those involved in *representative activities* [24] or those who act on our behalf in matters of economy, politics, government and science. Such dependence implies high degrees of trust on part of the individuals. We can naturally extend that view to the online world, where people who are involved in representative activities are so many. In fact, they might not seem as large a percentage, but they are a considerable number of people. For example, it is estimated that 1% of the users of Wikipedia are the actual contributors who create over 90% of the content [7]. Yet we choose to trust them as they act on our behalf, and we consequently trust the information they provide. Similar percentages prevail in most user-contributed content on the web and can be generalized to the wider web. It is then imperative that individuals willingly choose to trust a small group of representatives on the web to form beliefs and make commitments to using information online.

A widely cited definition of trust is that of Deutsch [26]. Deutsch states that trust occurs when ambiguity about a path a person has to take arises and the outcome of this path could be good or bad, the occurrence of either result is contingent upon the actions of another person, and the bad outcome is more harmful than the good/desirable outcome. Deutsch states the following in his definition of trust:

1. The individual is confronted with an ambiguous path, a path that can lead to an event perceived to be beneficial (Va^+) or to an event perceived to be harmful (Va^-);
2. He perceives that the occurrence of (Va^+) or (Va^-) is contingent on the behaviour of another person;
3. He perceives the strength of (Va^-) to be greater than the strength of (Va^+)

If the person chooses to take an ambiguous path with such properties, I shall say he makes a trusting choice; if he chooses not to take the path, he makes a distrustful choice [26, p.303]

This definition implies a personal view on trust. Trust here is dependent on the trustor's perception of the situation, and this implies a person centric view of the world. Thus two different individuals will perceive the situation differently and take different actions accordingly based on their evaluations of Va^+ and Va^- [67].

In Sztompka [92, p.27] trust is defined as *a bet about the future contingent actions of others*. This definition has two components, belief and commitment. A belief that a certain person will act in a favourable way and a commitment from my side to a certain action based on that belief. Trust here occurs in situations of uncertainty, where one is uncertain about the outcome of a transaction and trust acts as a mediating factor. Both definitions from Deutsch and Sztompka take a game theoretic approach to defining trust where the trust game represents a social dilemma where choices by a trustor involves different outcomes for both the trustor and the trustee, this approach is further explained in the next section.

Deutsch addresses nine different types of trust (see also [45]), but he focuses on *trust as confidence* where he asserts that a person trusts because he is confident of the positive outcome of his trust. This view supports our notion of trust between people and information. This is while taking into consideration the view expressed by Sztompka, where trusting objects stems from trusting the people behind these objects. Thus, when a user decides to use another user's observation on the web, then he is confident it will meet his expectations. Trust here can be viewed as a function of users' confidence in the quality of this observation (it's fitness for their purpose), thus we can say that:

$$T = f(i) \tag{2.1}$$

Where i is the confidence in the quality of information. It is intuitive to assume that i is positively affected by Va^{+} based on Deutsch's view that the higher the outcome Va^{+} is, the more likely the person to place trust (i.e. he has more confidence to place trust).

In Luhmann [65] a different outlook on trust is presented. Luhmann views trust as means of reducing complexity in society. He states that *the only problem that does arise is the relation of the world as a whole to individual identities within it, and this problem expresses itself as that of the increase in complexity in space and time, manifested as the unimaginable superabundance of its realities and its possibilities [65, p.5]*. Luhmann goes further to explain that *further increases in complexity call for new mechanisms for the reduction of complexity [65, p.7]* and he suggests that trust is a more effective mechanism for this purpose than alternatives such as utility theory [6] for example. This view is indeed interesting, given the fact that this research views trust as a tool for overcoming information overload in today's world. If trust is used by people to reduce the increasing complexity of life in space and time, we posit that online social trust can be viewed as the means people use to make sense of the complexity of the online world. In other words people rely on trust to organize their online world in space and time to reduce the complexities arising from information overload.

people rely on trust to organize their online world in space and time to reduce the complexities arising from information overload.

Luhmann also views trust as emerging from interactions of individuals in a society, and not strictly as a collective societal property. This is to say that trust in communities is driven by individual actions, which is reflected in online communities where users assert trust in others to be able to receive personalized views from those individuals about different types of content. This is a similar view to that of Deutsch, who addresses trust as a function of individual personal variables, yet it is different in that it places a much larger weight on societal systems. Trust thus can be viewed according to Deutsch and less according to Luhmann as a phenomenon that emerges bottom-up in the community. However, one cannot separate individuals' actions from societal norms [68] and by changing individual actions, the societal systems change and vice versa. Thus, we can identify two types of trust in that context. One depends on societal systems and the other on individual variables. Each is affected by a different set of independent variables. In Buskens [17] trust is studied as a collective property of social networks and that social network properties will affect the overall trust behaviour of the community. Nevertheless, Buskens cannot avoid the individual actions' effects on trust since he explains his work in game theoretic terms (see 2.1.3). The question of what these independent variables would be is a difficult question to answer, again because there are many types and definitions of trust. However, our central hypothesis in this research is about space and time as specific variables that have an effect as we elucidate later.

Gambetta [36] was among the first to use values for trust. Although trust is a very subjective measure, such that a value for one person could mean another thing for another person, his approach is a very important step in characterizing trust. He viewed trust as a probability (value between 0 and 1). According to Gambetta, *Trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both before he can monitor such action (or independently of his capacity ever to be able to monitor it) and in a context in which it affects his own action [36,*

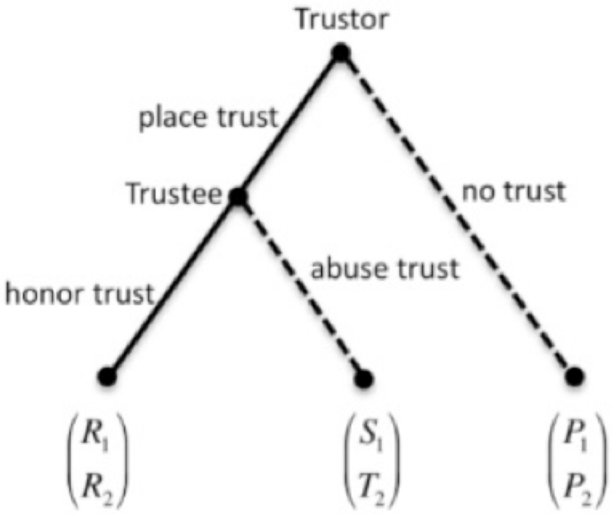

Figure 4: *a trust game with the different scenarios of play by the trustor and the trustee [17]*

p.217]. In this research however, we are not studying distrust, because Gambetta considers lower values than a preset threshold (say $p = .5$) to be distrust. While on the web, a lower calculated value of trust could simply be equal to the inability to determine whether or not a particular person is trustworthy. Such a conclusion is certainly not equivalent to distrust. Nevertheless, Gambetta's approach is interesting from our perspective because his approach places values on trust decisions by individuals. Whether we consider trust a probability value or simply given by a certain person for another person on a continuous or binary scale (see e.g. [39]) makes no difference to the subjective nature of such ratings. It is important though for any trust rating system to clearly define the semantics of this rating. We also note here that since we are not studying distrust, and because of the nature of trust as we study it in this work, we favour the view that lower trust ratings for all practical purposes imply just low trust rather than distrust.

We favour the view that lower trust ratings for all practical purposes imply an unknown status rather than distrust.

2.1.3 *Trust, an economics perspective*

In economics a myriad of researchers have studied trust, mostly in game theoretic settings. Game theory allows for predictions about behaviour of actors in situations of uncertainty about choices [33]. In figure 4 the trustor has two initial options at the beginning of the game, either to place trust, or to not place trust, in which case the game ends and both the trustor and the trustee receive a return P_1 or P_2 respectively. If the trustor decides to place trust the trustee has two options, either to honour trust and both players receive R_1 and R_2 respectively.

The other option is that the trustee decides to abuse trust and both players receive a return of S_1 and T_2 respectively. In this game $R_1 > P_1, R_2 > P_2, P_1 > S_1$ and $T_2 > R_2$. The game is a social dilemma since both players receive higher returns than when trust is not placed, and the trustee receives a maximal return if he chooses to abuse trust while the trustor achieves maximal loss in this case.

It is clear in this example, that trust here involves actions by one person whose outcome depends on the actions of another person. This definition does not strictly fit our view on trust, where someone decides to place trust in someone else's information, without direct dependence on the concurrent actions of that person, the outcome is not dependant on the trustee taking any further or subsequent action. Rather the trustee has already taken his action in the past by contributing the information and the decision to trust comes later, unlike the trust game in which the order of events is reversed. Also, we notice that the person choosing to use the information has some slight expectation about the quality of future information provided by another person.

distance is used as a proxy measure for social network density of buyers and suppliers networks. Proximity are always preferred partners implying a higher degree of trust.

Other relevant work in economics studies the relation between trust and proximity as a variable. In (Buskens, 2002) distance is used as a proxy measure for social network density of buyers and suppliers networks. The reason is that partners in proximity are always preferred partners (Nohria & Eccles, 1992), implying a higher degree of trust. Such conclusion is derived based on *Safe Guards* or precautionary measures taken when making contracts. The less the safe guards the higher the implied trust between partners [17]. Also important to mention here is the network embeddedness, where there are higher probabilities that firms located closer have more common ties with third parties who might also have more contacts in common among each other.

Potentially, geographic proximity has a positive effect on trust

Buskens also asserts that the cost to the buyer as a result of abusing trust by the trustee increases with distance. This means that the larger the distance the more difficult it is to resolve problems, leading to lower trust. A similar conclusion is made in [66], where he finds that distance has a positive effect on the probability of a subcontractor and a customer governing their relation by a formal contract, such that the probability increases with distance, implying less trust. A related explanation is provided in [64] where the subjects indicated that personal contact is important for establishing trust, and as such geographic distance was necessary in easing this personal contact, increasing trust with distance proximity. It is unknown as of yet without further research if such properties apply to online social networks. However indications of similar effects of geography on how online friendships evolve is available in [62].

Research implies that the closer the geographic distance between two subjects in a social network the higher the probability that trust relations will form

This implies a relation between geographic proximity and trust. The closer the geographic distance between two subjects in a social network the higher the probability that trust relations will form and that trust will increase with geographic proximity. Studies demonstrating this effect of geographic distance on trust do not discuss the psychological factors involved. We do not know if there are underlying psychological factors that govern our tendency to trust those in proximity. Research on evolutionary Psychology however, indicates that kinship or in-group can increase reciprocal cooperation [28], and hence the probability to place trust. We propose here that proxies of kinship such as geographic proximity, and for example, belonging to the same city, region or country increases one's sense of reciprocation and evokes emotions that can promote trust. Furthermore, studies of how similarity between people breeds connection, also known as homophily, suggest that geographic proximity is a basic source of homophily.

Other researchers in economic behaviour looked into geographic proximity and trust. Some have made qualitative statements suggesting a strong relation between geographic proximity and trust, specifically that geographic proximity between business partners foster more trust in inter-firm relationships. For example, Dyer and Singh [29] suggest that there is more face to face communication between suppliers and automakers in Japan than there is in the US or Korea, which may positively affect trust. They argue that this is facilitated by geographic proximity. Similarly in [61] the same conclusion is reached about supplier relations in both Germany and Britain, suggesting that geographic proximity fosters inter-firm trust. Similarly, a more qualitative study of the effects of geographic proximity concluded its positive effects on fostering inter-firm trust [11]. The same study in [11] refers to *trust as prudence* which implies trusting a partner because he has no incentive to abuse trust (the question here is if this is really an act of trust). This is compared to *trust as hope*, which implies trusting a partner who has the potential to abuse trust in hope it will not be abused. The study showed that proximity would have a positive effect on both types of trust.

Trust is the glue that holds everything together, the bond that creates healthy communities and successful businesses.[85]

Such studies on inter-firm trust on initial thought do not say much about trust between individuals as much as they say about trust between organizations. To make the connection we refer to section 2.1.2 where we relied on Sztompka's assertion that trust is always between people or individuals. Given this observation, we notice that [11] also states that it is inappropriate to state that organizations trust each other, but it is more appropriate to state that individuals within one organization build a trusting relationship with one another. The context of this statement though remains that of organizations and inter-firm relations, while we have no basis to generalize this statement to either real world social networks or online social networks. We do not know what are the effects of geography on trust between individuals offline or online when the inherent conditions and limitation of inter-firm relations such as contractual agreements and formal business practices are not involved.

2.1.4 *Trust in computer science*

In [41] three realms for the study of trust online are identified. Namely, trust in content, trust in services and trust in people. Such distinctions, particularly between trust in content and trust in people are in our opinion not clear cut. In our view trusting content stems from trusting people behind this content which we later call observations. Trust, reputation, user observations and information quality are entangled together such that making arbitrary distinctions leads to inappropriate conclusions. In chapter 4 we give an ontological account of these elements elucidating the relationships between them, however, for now it is not necessary to explain these relationships, but we simply avoid making such distinctions at this stage.

When discussing trust, the decision of whether or not to trust the information one encounters online is currently an unconscious decision based on different criteria given the lack of a web of trust that makes trust in a social sense explicit on the web. Studies on credibility of information on the web have dealt with web site appearance and design, as well as other signs of authority such as the reputation of the entity behind the website content be it a person or a company (see e.g. [32, 50, 63, 72, 82]. In [22] three aspects are determinants of how much some websites are trusted: perception of credibility, ease of use and risk.

Such criteria might apply for websites, but when we speak of trust as a proxy for information quality in general terms it is more difficult to relate to this criteria. For example, ease of use as a criterion is applicable when speaking about Amazon.com in general, but when assessing the decision to buy a book, the user ratings of the book as independent information entities are not subject to ease of use, but rather to ease of interpretation, while perception of credibility and risk remain relevant. To avoid unnecessary complexities, a simplifying assumption in our research is to look at trust assessment of some information entity a user encounters online as a single unit of decision or a black-box, without looking deeper into the underlying cognitive aspects of the decision. We can safely say that a user either decides to trust or not trust an information entity and the outcome of this decision is what concerns us in this research, not the process by which the decision is derived.

We assume the process leading to a decision to trust to be a black-box, our concern in this research is with the outcome, either to trust or not to trust

If the content we address is spatial and temporal in nature, it is reasonable to expect that the spatial and temporal dimensions should have some effect on the trustworthiness of this information [10, 8]. After all, I might trust you about Berlin simply because you live there, or have been there several times, while I would trust you less about London because you are based considerably faraway and have never been there or that you have been there but way back in the past. Such an intuitive assumption implies that trust mechanisms for information with spatial and temporal affinities needs to take into account the spatial and temporal properties of the information, the people creating/using the information and the interplay between the people and their geographic and temporal spaces. Such information could be part of the information pedigree/provenance, which we discuss later.

When speaking of trust in people on the web we refer to Web Based Social Networks (WBSN). WBSNs have grown tremendously over the past few years. The number of WBSNs online doubled between 2004 and 2006 from 125 to 223 [41]. Over the same period the total number of members among these sites grew from 115 million to 490 million [40] , and is expected to be much higher in the future. FOAF (Friend of a Friend) [27] is a method to resolve the conundrum that people maintain profiles on multiple social networking sites, FOAF allows for sharing different social networking data among sites via each user having his own FOAF profiles using its framework of representing information about people and their social connections. FOAF is being extended by various groups to allow for more complex description of social relations among people, e.g. there is a FOAF extension that allows users to describe trust relation on a scale of 1-10 between themselves and their social connections [44]. Also, OpenSocial that aims to provide a common API for all social applications on the web promises to make mashups from different social applications accessible to everyone on the web. Adding to that the universal authentication of web users across sites, like OpenID, is becoming more common to access all their accounts with a single logon paving the way for sharing trust information about those users.

In WBSNs the problem of computing trust in social networks can be summarized in figure 5. If A trusts B and B trusts D how much would A be willing to trust D based on his trust of B? This question requires assuming some properties summarized in [39]:

- *Composability and transitivity:* composability pertains to the ability to combine different trust values, which in turn entails transitivity

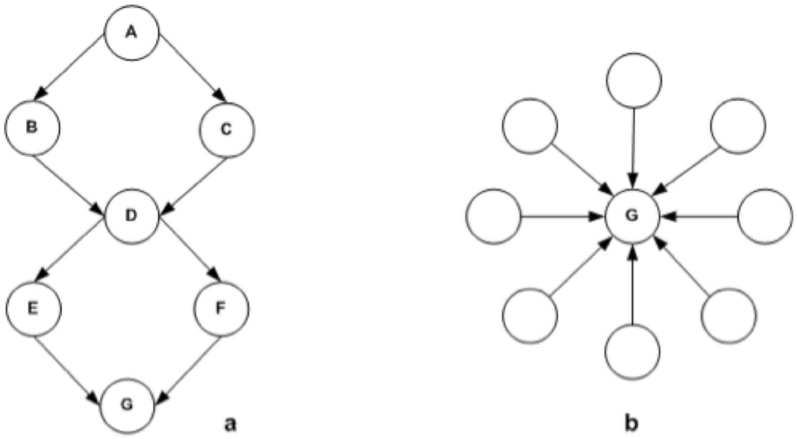

Figure 5: *(a) the problem is computing trust along the paths from source A to sink G. (b) the problem of computing trust to the sink at one degree of separation [39]*

as well. Trust is not entirely transitive in a mathematical sense. To say that Alice highly trusts Bob and Bob highly trusts Randy does not necessarily entail that Alice trusts Randy also highly. There is however a limited sense of transitivity in trust. People tend to share opinions about other people, conveying their trust in those people to others. Computationally this allows us to make trust transitive along chains of connections which has been widely implemented (see e.g. [39, 47, 49]).

- *Personalization and asymmetry:* this pertains to two intuitive facts about trust which were also earlier addressed in our previous analysis. Trust is a personal opinion, so that two people will mostly have different trust values on a similar issue, making the amalgamation of different people's opinions about some information entity difficult. This also means that trust is also asymmetrical. If Bob trusts Alice with some value x it can happen that Alice trust Bob with a lower value y. Both such issues we will have to deal with in chapter 4.

- *Contextualization:* trust is highly contextualized. Trusting you to review my paper does not entail trusting you to fix my car. Generally trust algorithms impose some context on trust when doing computations [41], however a trust model might incorporate context explicitly in the formalism to differentiate between the results of the model in different contexts, an issue worth investigating when developing the ontological account and trust models.

2.1.5 *Ontologies of trust*

Within the last years, several ontologies of trust have been proposed [93, 51, 44, 95, 94]. While they differ in terms of application areas, they all focus on computational aspects of trust such as trust values, rankings, belief reasoning, contextual influence on trust-based recommendations, and so forth. In contrast, our research will lay the ontological foundations of trust and reputation in more general terms, regardless of any given models, while focusing on observations in for VGI and user contributed content, and the human sensor web is used as a use case. As such the few works on ontologies of trust and reputation are of little relevance to our work.

2.2 PROVENANCE

One cannot discuss trust in content and in people without discussing provenance, which is the history of something (some times referred to as information lineage, or pedigree [16]). Provenance describes the origins of information and the process by which it came into being, and helps to verify the authenticity of the data. This in turn makes it possible to judge the trustworthiness of the information. Also, provenance might describe the people who created the information, and if we know how trustworthy these people are then we can decide on how trustworthy is the information they produce. It is important to notice that the value of provenance to any application depends on the degree of granularity at which provenance information is collected [88]. Several approaches exist for modelling provenance of information on the web, and of these approaches several rely on semantic web technologies such as illustrated in [43, 35, 56, 100]. Such approaches however, are concerned with provenance of information on websites and of ontologies on the web to judge if certain inferences or assertions can be relied upon.

Provenance although important, is not within the scope of this research.

Considering our earlier discussion on trust from an economics perspective, we can say here that trust is based on past behaviour (provenance) and the shadow of the future, such that:

$$T = f(p, s) \tag{2.2}$$

Where T is trust, and (p, s) are the provenance (previous behaviour of the information contributor) and the shadow of the future (aka. expectation to provide quality contributions) respectively. However, a person deciding to contribute or recommend information could naturally be trustworthy or not. A user deciding to use this information can view trust in this person's recommendations in light of his previous behaviour as a good recommender/contributor, hence he commits to trust in hope of a good outcome in the future. It is also evident that the longer the duration of interaction between a trustor and a trustee the higher is the trust given that the future is important for the trustee. This in fact is the case with most Internet applications, where users tend to care for their long-term reputation with their desire of remaining trustworthy to their trustees [41]. In that regard we can assume that human sensors are not different and will always try to maintain individual reputations as rational economic actors. In this view, we take the view that provenance is a separate issue and choose not to deal with it explicitly within the scope of this research. However, the models presented later in this thesis implicitly rely on some form of provenance information.

2.3 THE SENSOR WEB AND VOLUNTEERED GEOGRAPHIC INFORMATION

The term *Sensor Web* was perhaps first coined by Kevin Delin of the NASA Jet Propulsion Laboratory in 1997 to define the idea of *"developmental collections of sensor pods that could be scattered over land or* *water areas or other regions of interest to gather data on spatial and temporal patterns of relatively slowly changing physical, chemcial, or biological phenomena in those regions"* [25]. Sensor Webs have since taken a life of their own and are now a widely studied research field with several standards being developed through standardization bodies like the Open Geospatial Consortium (OGC). According to OGC, *a Sensor Web refers to a web of accessible sensor networks and archived sensor data that can be discovered and accessed using standard protocols and application program interfaces* [13]. The OGC Sensor Web Enablement (SWE) [13] framework was conceived as a collection of well-defined standard model encodings and standard web service interface. The SWE framework provides these standard encodings:

- *SWE Common*: models and XML schema for defining SWE data types used throughout the SWE specifications.
- *SensorML*: models and XML schema for defining sensor systems and the processes surrounding observations.
- *TransducerML:* XML schema for providing observation values and packaging.

The following web service interfaces are also provided by the SWE:

- *Sensor Observation Service (SOS)*: common interface for retrieving observations from sensors, processes, or models.
- *Sensor Alert Service (SAS)*: common interface for publishing and subscribing to asynchronous alerts from sensor systems, processes, or models.
- *Sensor Planning Service (SPS):* common interface for tasking a sensor system or model
- *Sensor and Observation Registry:* common interface for discovery of sensor systems, processes, models and observations

In addition several other standards have been proposed to complement the SWE. The Sensor Event Service (SES) aims to filter incoming notifications depending on a criteria defined by the subscribers [30], while the Web Notification Service (WNS) [89] allows a client to conduct asynchronous message exchanges with one or more other services. In addition, within the SWE, models and simulations are considered as sensors themselves, they are essentially another source of observations. The encodings mentioned earlier have originated in the OGC Observation and Measurement (O&M) [23] and an observation and measurement ontology for the sensor web is presented in [60, 79] and is addressed in detail in chapter 4.

It is sobering to be reminded that one of the basic instincts of human nature-mutual cooperation for no cost-is thriving on a global scale. [55]

Collaboratively generated content for the geospatial domain and the associated quality problem [10] and the subsequent vision of humans as sensors [46] have gained momentum and continue to grow. In [46] VGI is enabled by various trends and technologies, one of which is humans as sensors. Humans in this vision are synthesizers and interpreters of local information. Enabled by Web 2.0 technologies humans use their own senses to report information about the world around them, and in other cases they are empowered by extra sensory devices that have become mundane everyday devices like smart phones the modern ones of which contain 10 or more different sensors. The potential for using humans as sensors is just starting to emerge and project like the Human Sensor Web[1] which we introduce in chapter 4 are beginning to integrate humans as sensory assets within the larger sensor web framework. This integration poses many challenges, and in this thesis our research challenge is the quality assessment of human sensor observations. Our proposal for spatio-temporal trust and reputation models as quality assessment measures of human sensor observation has been discussed in [9, 10] and is presented in the next chapter.

2.4 DISCUSSION & CONCLUSIONS

It is essential that we study the different facets of trust in order to come to a consensus as to what we mean by trust in our research. Although the different fields that studied trust seem diverse, it is clear that they are all tied by the notion of society [68], such that trust is essential for the existence of society, or that societies lead to the occurrence of trust, which is a somehow circular argument depending on which research you refer to. We cannot dismiss this argument since it seems ubiquitous in the study of trust, however a question remains, how do people trust information? For example, trusting your car to start in the morning implies trusting your manufacturer, or your repair shop, so is trusting your airline to take you to your destination implies trusting the people managing and working at this airline [92]. This underlies the definition we adopt in this work. The explicit claim we make is that trusting information is not different.

> DEFINITION 1: to trust information is to trust the agents behind this information, information here pertains to observations made by human sensors

This latter assertion is implicit in various computer science trust studies that assume people making claims about the trustworthiness of each other, then later assessing the trustworthiness of the observations provided by those people within the community [39, 74, 101].

Thus, in our work we are concerned with the notion of social trust. Social trust reflects opinion similarity making recommendations from trusted users relevant for information requesters [102]. Particularly, research has shown that when a user's opinion is different from the average opinion of the population, social trust can be an effective tool in giving this user relevant information from his viewpoint. Social trust then has been used successfully in applications for rating, sorting and filtering information, especially user-generated content that is proliferating on the web today [41]. We also argue that social trust is generally useful for filtering all types of information that can be rated and assessed by users in online social systems.

The understanding of the external and internal factors affecting online social trust is still rudimentary. This is due to the fact mentioned earlier, namely that trust in general is a difficult concept to define let alone formalize. Thus, it is a difficult task to identify such internal and external factors. Research in network science and modern social network dynamics [3, 98] hold promise in helping us understand how complex behaviours in social networks arise from simple individual actions, which can be reflected on our understanding of trust.

---

[1] http://www.h2oinitiative.org/article/17001/Human_ Sensor_ Web

VGI and human sensor observations, as we earlier discussed are spatial and temporal in nature. Within this wide scope of internal and external factors that are potentially affecting trust, we position our hypothesis on spatial and temporal aspects of the problem. It is intuitive to assert that trust decays and develops over time, or that trust develops slowly, but can easily be tarnished if abused, showing a temporal affinity to the concept of trust. Some researchers tried to formalize these temporal effects on trust; see for example [68]. It might be less intuitive though to think of the effects of space on trust, but some are still apparent in our lives although being subtler as we have discussed earlier. Consider trusting a person who lives in Berlin in providing information about Berlin, versus the same person providing information about London compared to a Londoner. Of course, one can postulate one moving from Berlin to another place while still remaining knowledgeable about Berlin in some sense or the other, but then again the temporal currency of his experience about Berlin declines nonetheless. Also, time comes again into play to define what kind of information is he trusted about given the increased distance from Berlin (e.g. can he still be trusted about traffic information). The question we raise is how can we employ some forms of space and time in trust and reputation models? and would they affect the performance of such models? Clearly, there are effects of space and time on trust that can be utilized to build more effective trust and reputation models, particularly for VGI and human sensor observations.

Quality assessment of VGI and human sensor observations is the focus of this research. In the next chapter we distil our findings from this chapter leading to the introduction of our proposal for spatial and temporal trust and reputation models as a novel approach to observation quality assessment.

3

TRUST AS A PROXY MEASURE FOR OBSERVATION QUALITY ASSESSMENT

The aim of this chapter is to introduce our proposal for trust as a proxy measure of observation quality for human sensors. Using the information gained from the previous chapter, we start by discussing the contextualization of trust and our assumptions about it in this research. We then discuss our view on the spatial and temporal aspects of trust and how they relate to our proposal. We also discuss reputation and trustworthiness and their role in our proposal. Finally we address trust as a social phenomenon between people and our proposal to use trust in information as a proxy measure of quality by introducing the notion of *informational trust* as a new form of trust mediated through *interpersonal trust*. We conclude by introducing our vision for spatial-temporal trust and reputation models for quality assessment of human sensor observations through effective information triage and filtering.

3.1 CONCLUSIONS ON TRUST

In this section we discuss three aspects of trust derived from our background research. Namely we address the contextualization of trust, the spatial aspects of trust and the temporal aspects of trust. The aim is to clarify our basic assumptions about the three aspects and their relevance to our research.

3.1.1 *Trust is Contextualized*

When using the sentence *I trust this person*, we humans rarely define the sentence in any clearer terminology, and from the individual sociological perspective [92, 86] we are rarely clear about what exactly we are willing to trust this particular person about. One can postulate that trusting a person about keeping the keys to your residence does not in anyway imply trusting the same person about the keys to your office, yet in our minds the term carries a vague conceptual meaning that resembles a general mood of trust that is not made explicit, but is inferred from the context in which the sentence is made. This context parameter is important and will determine the truth value of the sentence *I trust this person*. We can then say that:

$$T = f(i, c) \tag{3.1}$$

Where c is a context parameter that defines the context where a trust assertion is made. We note that this context parameter makes trust values from different contexts incomparable (i.e. truth value of the same statement is not portable across contexts). One might be tempted to make a context comparability assumption to simplify any formalism of the problem, especially if we are talking about information in a general sense in online environments, especially VGI. In our research we resort to neutralizing the context problem such that there is always one context whenever we are making trust statements in this research.

We simplify our assumptions such that there is always one context shared by all parties whenever we are making trust statements in models developed in this research.

This context is that of trusting a specific information entity provided by person A and acted upon by person B about a specific VGI observation of a single phenomenon. As such we do not try to formally model the context of trust in our research either in the ontological work presented in the next chapter or the computational models presented in chapters 5 and 6.

3.1.2 *Spatial Aspects of Trust*

According to our observations from Buskens [17] discussed in chapter 2 we recognize that social network density has a positive effect on trust relationships in social networks. In this research we are more concerned with online communities, however we note that online communities and real world communities could both, to a limited degree, be used as proxies for one another. As such, from our knowledge that network density is affected by geography, we can conclude that geographic proximity positively affects trust relationships in communities.

The probability of befriending a certain person is proportional to the number of closer people.

In [62] friend-formation patterns in a large scale spatially situated social network harvested from online sources is studied. In this work the probability of befriending a certain person is proportional to the number of people in geographic proximity. This does not of course say much about trust in particular, but it does imply some relation since we postulate that friendship implies an inherent trust component at some level. The question here is to what level does geographic proximity impact trust between people, and consequently impacts each other's recommendations about certain observations.

Our proposal here is that if the spatial dimension is explicitly represented in trust and reputation models for spatial information like VGI then we can build effective models for quality assessment of the VGI.

3.1.3 *Temporal Aspects of Trust*

As we earlier discussed In Marsh [68] an attempt to formalize the notion of trust in distributed artificial intelligence (DAI) is introduced. Marsh attempts to produce a theoretical model of trust in DAI and incorporates in his model sociological and psychological aspects of trust. This includes the temporal nature of trust. A problem with Marsh's model is that in the open environment of the web much of the information needed for his model is not available [39]. Users tend to give each other trust ratings in a limited context that mostly has to do with trusting the user recommendations about some information like product recommendations. However, we find Marsh's model highly relevant in that it formalizes how humans build and maintain trust, particularly interesting is the time aspect dealt with in the model. Also, an implication of the work done in [19] is that trust is developed over time as a result of continuous interactions and does not arise spontaneously. Hence, we assert that trust is built slowly over time, but as a matter of common life practice, it is also tarnished with immediacy when abused. In fact, some trust models have tried to incorporate this intuitive assumption (see [96]).

Thus, trust relationships develop and decay with the time dimension. Such effect of time on trust relations will influence how we develop trust models for VGI that are sensitive to the temporal dimension.

3.2 REPUTATION AND TRUSTWORTHINESS

Trustworthiness of a certain person is not an intrinsic quality of that person. This is to say that trusting someone does not actually mean the person is trustworthy, it simply means the trustor decided to place trust in that person, whether or not the trustee will honour this trust or defect is another question.

When trustworthiness is viewed this way, reputation is then said to be the perception of trustworthiness of a person by the community [74]. Reputation of a person is not an act of that person, but a quality bestowed upon that person by the community, and it depends on many factors including previous behaviour, community perception of the person, the capacity of the community to sanction bad behaviour and propagation through word of mouth.

In this research we propose that reputation is the collective trust vested in a person by the community, so that if Bob trusts Alice with a value x, then the collective of x values by all other members of the community is said to be Alice's reputation. It goes for reputation as well that it is a contextualized problem, although slightly different from trust. One can see that acquiring a bad reputation as a student could imply a bad reputation in other contexts such as work, or family. However, for our work we consider that reputation in one context has no influence on other contexts. Thus reputation of person x in context C can be represented as:

Reputation is the collective trust vested in a person as perceived by the community.

$$R_x = (\overbrace{\mathrm{T}_x}, C) \tag{3.2}$$

In other words reputation of person x is a function of the collective trust in person x by the community, this collective trust is denoted $\overbrace{\mathrm{T}_x}$ in context C. Following traditional economic thought, rational agents would naturally like to maintain good reputations, as such they have an incentive to make quality observation contributions in order to maintain their reputation. It could be said that quality of information contributed by a person is related to their reputation and people would always try to maintain a good reputation by providing high quality information.

The quality of observations contributed by a person is directly proportional to her reputation.

3.3 INFORMATIONAL TRUST

Several researchers argue that trust holds only between people which make our everyday use of the terms in statements like *I trust this information* to be essentially flawed. Trusting a company like Lufthansa to take you to your destination is, in fact, trusting the people behind the company or in a way personifying the company as an institutional entity. How can we then argue for using trust in information as a measure of information quality?

In the previous chapter we discussed the trust definition of Sztompka [92] which defines trust as *a bet about the future contingent actions of others*. This is also the definition we adopt here for *interpersonal trust* as a social tie between a trustor and a trustee [70]. We reiterate here the two components of this definition as discussed earlier, namely belief and commitment (the belief that a particular person will act in a favourable way and my commitment to a certain action based on that belief). Therefore, one could argue that trust in entities is based on trust in the persons responsible for these entities. Following this argumentation, we propose the notion of *people-object transitivity of trust* which differs from the trust transitivity commonly applied to Web-based social networks [39]. In our view, *interpersonal trust* implies the transition of trust from the trustee to information entities conveyed by the trustee. The trustor can then assert trust directly in the information conveyed by the trustee. We call this *informational trust*; where a trusting tie between a trustor and an information entity such as VGI is mediated by interpersonal trust between the VGI originator and the VGI consumer.

Given our earlier discussion on the spatial and temporal aspects of trust, we further propose to extend *informational trust* by spatial and temporal characteristics of the trust phenomenon. It is intuitive to assert that trust decays and develops over time, or that trust develops slowly, but can easily be tarnished if abused, showing a temporal affinity to the concept of trust. It is also intuitive that people's location with respect to the observed phenomena will impact how others trust their observations.

3.4 SPATIAL-TEMPORAL TRUST MODELS FOR ASSESSMENT OF OBSERVATION QUALITY

The observations provided by human sensors lacks traditional GI quality criteria (completeness, consistency, lineage and accuracy). In addition, VGI systems are subject to fraud by malicious users or to contamination by low quality observations from inexperienced users. This deficiency in the unconventional process of VGI production calls for unconventional solutions to the quality problem.

Informational-Trust through spatial and temporal trust models is our proxy measure of the quality of observations made by human sensors

We have addressed in chapter 2 how trust through different trust models is used to filter then provide relevant recommendations which is in our view a process of information triage and filtering. We have also discussed spatial and temporal aspects of trust. We propose to use the introduced notion of informational trust to develop spatial-temporal trust and reputation models that can leverage the spatial and temporal nature of VGI observations to act as proxy measures of information quality. Our definition of quality in that context is fitness for purpose, as such quality here is a subjective measure, and it always is to some extent [10]. When observations are fit for the purpose of an observation consumer at a specific location at a specific time, these observations are then said to be of high quality. By integrating the spatial-temporal trust and reputation models into the human sensor web we can mitigate the risks of fraudulent or low quality observations. This is done through filtering of trusted observations by reputable users while diffusing the negative effects of less reputable users.

We call our proposal *trust as a proxy measure of observation quality*. This proposal requires that:

1. Integration of trust and reputation as quality measures of human sensor observations into the sensor web observation and measurement standard.
2. Development of computational trust and reputation models for human sensor observation filtering and triage

3.5 CONCLUSIONS

In this chapter we discussed the contextualization of trust and our simplifying assumption. We took a pragmatic approach of treating context as out of scope of this research. We assume a constant context when making trust statements in our work and do not attempt to formally model the context of trust. We also discussed the spatial and temporal aspects of trust leading to extending our proposed notion of informational trust with spatial and temporal dimensions. We then discussed reputation and trustworthiness. We commit to the view that *the quality of observations contributed by a person is directly proportional to her reputation.*.

We concluded our discussion by introducing our proposal for *trust as a proxy measure of observation quality* and the two research objectives. One is the ontological approach to introduce trust and reputation as quality measures to sensor web ontologies for human sensors. The second is computational models of trust and reputation for human sensor observations. In the following chapter we present an ontological account of observations and measurement taking into account humans as sensors and trust and reputation as measures for quality assessment. The following two chapters 5 and 6 present two computational spatial-temporal trust models and their evaluation.

4

ONTOLOGY OF TRUST & REPUTATION FOR OBSERVATION QUALITY ASSESSMENT

Our proposal to build spatio-temporal trust and reputation models for quality assessment of human sensor observations requires as a first step integrating trust and reputation as quality measures into the sensor web, particularly for human sensors. The ontology presented in this chapter extends the functional ontology for observation and measurement (FOOM) presented in [60, 79] to accommodate quality assessment measures in the form of trust and reputation models for humans as sensors. We present an ontological account of trust and reputation using the water well use case scenario of the H2.0 project [54]. We show through the scenario implementation in FOOM how trust and reputation models can be integrated within the sensor web. The ontology is grounded in the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE). Thus, it is important to highlight the distinction between *observation quality* in the information quality sense and that of the notion of an *observable quality* in the DOLCE foundational ontology as the basic entities that can be perceived and measured which inheres in other entities, like the color of an apple [69]. We show how our approach formalizes the notions of quality measures for observations by humans as sensors [46], the computational models of trust and reputation also as sensors and trust and reputation as qualities.

4.1 MOTIVATING SCENARIO: A HUMAN SENSOR WEB FOR WATER AVAILABILITY MONITORING

This section introduces guiding scenario for the work in the following chapters. This scenario is adopted from the H2.0 [1] initiative and is based on the Human Sensor Web [2] for water availability monitoring. The system developed for this scenario has been implemented in Africa [3] and is described in [54, 53]. The project is co-funded by google.org and UN-Habitat. The projects develops Water Supply Monitoring System (WSMS) based on the vision of humans as sensors [46] . The project utilizes sensor web technology to implement a VGI system by which the local population can report on the quality of water wells in the area as well as inquire about the available water wells that have drinkable water within the same area.

Figure 6 illustrates the WSMS system showing the major user groups and interactions of the users with the system. The users of the system can be viewed as the general public, the subscribers and the reporters. Reporters act as sensors and make observations and report them back to the WSMS. The subscribers to the WSMS receive notifications upon request on the water quality of the water wells in the area. The general public are users who are neither subscribers nor reporters and can use a web portal to access the information from the WSMS.

---

[1] http://www.h2oinitiative.org/
[2] http://www.h2oinitiative.org/article/17001/Human_ Sensor_ Web
[3] http://geonetwork.itc.nl/zanzibar/

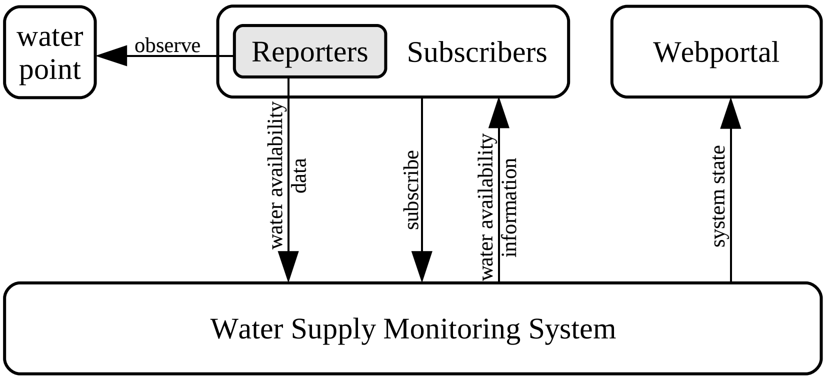

Figure 6: *Overview of the Water Supply Monitoring System [54].*

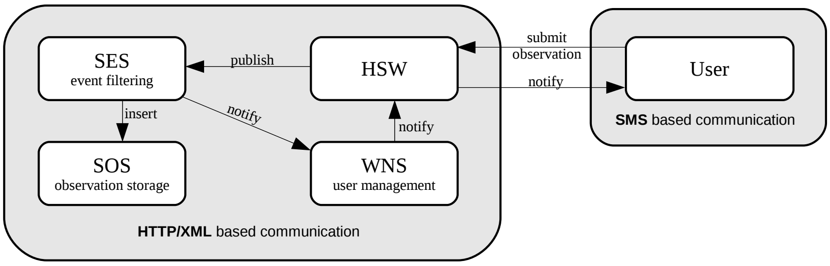

Figure 7: *Implemented Architecture of the WSMS [54].*

The system architecture for the WSMS depicted in figure 7 is based on the Sensor Web Enablement (SWE) of the OGC [13]. The system uses GSM messages from the reporters to report their observations which are collected and processed through the Sensor Event Service (SES) interactions as shown in figure 8. After the processing done through the SES, the observations are sent to the Sensor Observation Service (SOS) where they are stored. The SES also notifies the Web Notification Services (WNS) which handles the system subscribers to notify them about water well status in their area. The information from the WNS is first passed to the Human Sensor Web component developed specifically for the WSMS system which forwards this information to the SMS gateway in response to queries by the system subscribers. A detailed description of the system is provided in [54].

The WSMS system in general assumes an ideal environment where all users act altruistically for the common good of all users. While this assumption is reasonable in many cases, it falls short of its promise in this scenario where the motivation for the abuse of the system by some users is high. One can postulate various scenarios where the system is misused or abused, among which:

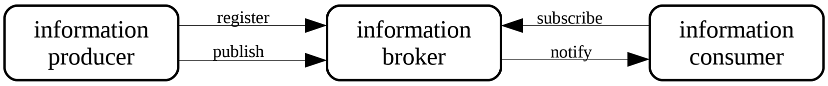

Figure 8: *Interactions of the SES [54].*

- inexperienced users make faulty reporting. This results in low quality information in the system despite the good intention of the users
- malicious users from competing tribes attempt to manipulate the system to keep the high quality water wells for themselves by misguiding others
- much like many web 2.0 applications the flow of VGI becomes so large so as to make effective information triage impossible

The result of these factors is low quality information or an information overload that makes the WSMS an unreliable source of information for all its users. To mitigate such risks the SES implements a simplified version of the trust model presented in chapter 5. In the next chapters we introduce our trust and reputation models for VGI using the WSMS scenario to illustrate the working of these models and deploy agent based simulations to test the effectiveness of such models.

The presented WSMS scenario is the guiding scenario for the remaining parts of this thesis. In the next sections we introduce an ontological account of trust and reputation as proxy measures for observation quality for humans as sensors to enable the integration of trust and reputation models in the sensor web platform. We use the WSMS scenario to implement our ontology to show how the spatially and temporally sensitive trust models, presented in chapters 5 and 6, can be integrated in the sensor web platform for quality assessment of the observation of human sensors.

4.2 THE FOUNDATIONAL ONTOLOGY DOLCE

Ontologies as formal theories are based on some formal language. Much like all languages, one can only define a new term with the help of a set of well-known terms. A set of terms with a well fixed meaning is essential and foundational ontologies in a nutshell provide such set of terms. As [76] points out, with foundational ontologies general errors in ontology engineering can be averted since basic philosophical distinctions are clarified and fixed on a higher level compared to the domain level. A foundational ontology like any other ontology assumes a shared vocabulary, yet this set of terms for which a common understanding is assumed are relatively few. Most importantly, this set of terms represent the most basic ontological distinctions and are domain independent. Formal ontologies are an important tool when recognizing disagreements in the use of vocabulary is important and when precise distinctions between terms are required. Figure 9 shows the case where foundational ontologies are of increasing value. Because effective communication, sharing and discovery of sensor observations rests on our ability to specify the intended meaning of observations rigorously across domains with minimal ambiguity. The Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) [4] is the foundational ontology adopted by the observation ontology FOOM presented in section 4.3 and consequently also adopted in our extension of the FOOM ontology presented in this chapter. In this section we briefly introduce the basic aspects of DOLCE which are required to understand the work presented in this chapter.

---

[4] http://www.loa-cnr.it/DOLCE.html

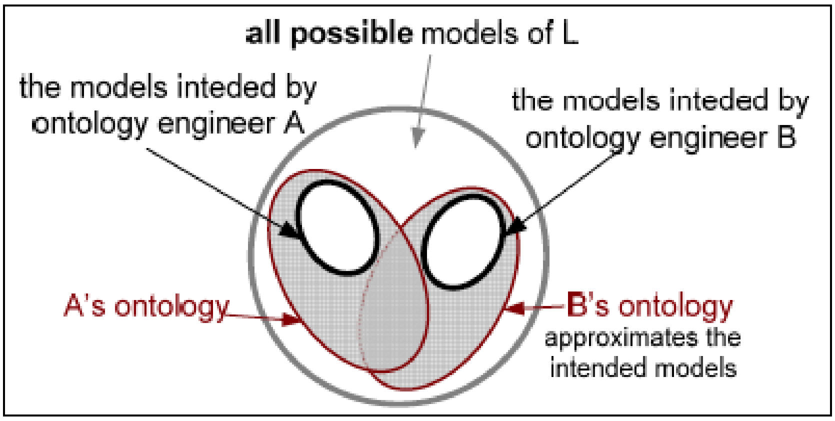

Figure 9: *The intended model of person A does not overlap with the intended model of B although their ontology axiomatizations overlap. based on [80] as modified from [48]. Foundational ontologies are meant to improve such a situation when compared with shared vocabularies.*

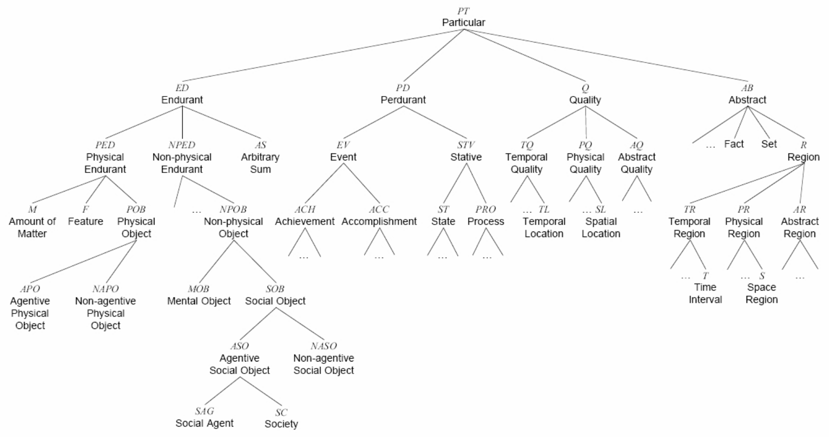

Figure 10: *The basic categories in DOLCE [69].*

In DOLCE categories are cognitive artefacts that depend on human perception, cultural imprints and social conventions. Figure 10 presented the basic DOLCE categories. The four basic categories of DOLCE are Endurant, Perdurant, Quality and Abstract. The difference between endurants and perdurants is apparent with respect to their behaviour in time. Endurants are wholly present at any moment in time and their parts move with them in time, e.g. a an amount of air is an endurant. Perdurants extend in time by accumulating different temporal parts. At any time perdurants are present, they are only partially present in the sense that their proper part are not present any more or are partially present e.g. a storm is a perdurant. Endurants as such can participate in perdurants or are said to live in time by participating in a perdurant e.g. an amount of air participates in a storm.

Another DOLCE top level category that we are concerned with is quality. Qualities are basic entities we can perceive or measure, for example the volume of an amount of air is a quality inhering in an amount of air. As such qualities inhere in entities, so in DOLCE physical qualities can only inhere in physical endurants. Non-physical endurants can only have abstract qualities while perdurants can only have temporal qualities. DOLCE makes a strict distinction between a quality like the volume of an amount of air and its value (called a quale in DOLCE) which can be approximated with a volume measure to, e.g., three cubic meters.

Formal ontology provides the general means to state what can be observed and measured, while measurement theory provides the general means to state how to measure. The specification of precise semantics for observation and measurements requires a framework that combines the specifications of what is measured with those of how it is measured [80]. The FOOM ontology presented in section 4.3 aims to brings such precise specification to the sensor web.

4.3 FUNCTIONAL ONTOLOGY OF OBSERVATION AND MEASUREMENT

In this section we introduce the Functional Ontology for Observation and Measurement (FOOM) and discuss the basic assumptions underlying the ontology. Observations are the link between information and reality. Basic ontological questions like *what can be observed?* or *how do observations relate to entities in reality?* can not yet be answered properly [77]. The observation ontology presented in [60] specifies observation as processes and defines their semantics by the four top level DOLCE categories of qualities, endurants, perdurants and abstracts involved in the observation processes. The ontology generalizes from technical sensors to human sensors to observing agents in general. It also uses an approach of generalization of observation and sensing behaviour where people, devices, sensor systems and sensor networks can all observe. Such an approach has been shown to improve the architecture of the sensor web [91, 15].

The FOOM ontology is encoded in the functional programming language Haskell which is a powerful algebraic specification language and also acts as a testing platform for the ontology through simulation. We will refer to Haskell code using the following conventions adopted in [60]. Capital true type fonts stand for Haskell type classes like `OBSERVATIONS`, while normal true type fonts like `Drinkability` stand for data types, and lower case true type fonts like `perceive` stand for operations and individuals. We begin by showing upper-level type classes from the FOOM ontology. We then show examples of sensor implementations in the ontology and show its capabilities. In FOOM endurants are implemented as data types of kind $*$ and have identities. Perdurants are implemented as types of kind $*- > *$ which translates into being functions over endurants. Qualities are by definition dependant entities and are modelled as constructor functions applied to their host, i.e. the entities in which they inhere. Haskell offers the higher level notion of type classes which amalgamate types sharing some behaviour. The type class `QUALITIES` is defined for all functions returning a single quality value.

In FOOM, the observation process first invokes a quale, defined as an analog signal in an observer's mind or technical sensor. This invoked quale is followed by symbolization to complete the observation process. The core concepts of FOOM are quality, stimulus, agent, quale and value. A quality physical or temporal can be observed such as temperature or humidity of an amount of air (physical quality) or the duration of a storm (temporal quality). Abstract qualities can also be

Listing 1: The Haskell definition of `AmountOfAir` and the definition of the top-level DOLCE classes.

```haskell
class PHYSICAL_ENDURANTS physicalEndurant
class PHYSICAL_ENDURANTS amountOfMatter => AMOUNTS_OF_MATTER
    amountOfMatter
class PHYSICAL_ENDURANTS physicalObject => PHYSICAL_OBJECTS
    physicalObject
class PHYSICAL_OBJECTS agent => AGENTS agent -- agentive
    physical objects

data AmountOfAir = AmountOfAir {heat:: Float, moisture:: Float}
    deriving (Eq, Show)
muensterAir = AmountOfAir {heat = 20.0, moisture = 70.0}

instance PHYSICAL_ENDURANTS AmountOfAir
instance AMOUNTS_OF_MATTER AmountOfAir
```

admitted but are currently not covered in the ontology. A stimulus is required which can be defined as a *detectable change in the internal or external environment* following the definition adopted in [60, 79]. An example of a stimuli is the heat flowing between an amount of air and a thermometer. An agent is defined in FOOM in line with artificial intelligence literature and allows for the inclusion of both humans and technical sensors or any other entities that can observe. According to FOOM the conceptualization of the observation process is inline with OGC's observation and measurement standard [23]. It has two steps performed by agents:

1. first, agents generate a quale for the observed quality through stimuli.
2. second, they express the quale by symbols representing the observation value.

The following illustrative examples clarify the basic FOOM concepts:

- A *Thermometer* measures the temperature of an amount of air using heat as the stimulus. The heat expands the amount of gas by some amount. This is the signal that gets converted to a number of degrees on, for e.g. the Celsius scale.
- A *Person* perceives temperature of the surrounding air and a quale is invoked in her mind corresponding to the perceived temperature. The person expresses the quale as an observation value of either "high" or "low" temperature.

Listing 1 shows the `AmountOfAir` that have heat and moisture. It is of type `AMOUNT_ OF_ MATTER` type class which is a `PHYSICAL_ ENDURANTS`, this is an example of how FOOM defines the universal `AmountOfAir` as quality carrying endurants according to DOLCE. We can then proceed to define individual values such as the air in the city of Muenster `muensterAir`.

Listing 2 shows the quality `Temperature` which inheres in an endurant, in this case `AmountOfAir` that is earlier defined in listing 1. Such a quality is perceived by observing `AGENTS`, in the case of our examples above these are a `Thermometer` and a `Person` as defined in listing 3. Both `Thermometer` and `Person` show how the FOOM generalizes to observing

Listing 2: The Haskell definition of the `Temperature` quality.

```
class QUALITIES quality entity

data PHYSICAL_ENDURANTS physicalEndurant => Temperature
     physicalEndurant = Temperature physicalEndurant deriving
     Show
instance QUALITIES Temperature AmountOfAir
```

Listing 3: The Haskell definition of both the `Person` as sensor and the `Thermometer` as sensor.

```
data Thermometer = Thermometer {tid :: Id, tQuale :: Float, tValue
     :: Float} deriving Show
fmoThermometer = Thermometer {tid = 1, tQuale = 0.0, tValue =
     0.0}

instance Eq Thermometer where Thermometer id1 _ _ == Thermometer
      id2 _ _ = id1 == id2
instance PHYSICAL_ENDURANTS Thermometer
instance PHYSICAL_OBJECTS Thermometer
instance AGENTS Thermometer

data Person = Person {pid :: Id, legLength :: Float, loc :: Id,
     pQuale :: Float, pValue :: String} deriving Show
ann = Person {pid = 1, legLength = 0.8, loc = 0, pQuale = 0.0,
     pValue = ""}

instance Eq Person where Person id1 _ _ _ _ == Person id2 _ _ _
     _ = id1 == id2
instance PHYSICAL_ENDURANTS Person
instance PHYSICAL_OBJECTS Person
instance AGENTS Person
```

agents in general where a human sensor and a technical sensor are interchangeable as observing agents.

To this end the agents perceive stimuli then produce an observation value. Listing 4 shows the `STIMULI` type class that provides the `perceive` behaviour that sets the quale in the mind of the person or in the thermometer.

The final step is the `observe` behaviour provided by the `OBSERVATIONS` type class so that the agent can express the value of the observation result as shown in listing 5. In this example the person `ann` makes an observation on the temperature of the air in the city of Muenster

Listing 4: The Haskell definition of the `perceive` behaviour of the `STIMULI` type class.

```
class (QUALITIES quality entity, AGENTS agent) => STIMULI
     quality entity agent where
          perceive :: quality entity -> agent -> agent

instance STIMULI Temperature AmountOfAir Thermometer where
          perceive (Temperature amountOfAir) thermometer =
               thermometer {tQuale = heat amountOfAir}
ptt = perceive (Temperature muensterAir) fmoThermometer

instance STIMULI Temperature AmountOfAir Person where
          perceive (Temperature amountOfAir) person = person {
               pQuale = heat amountOfAir}
ptp = perceive (Temperature muensterAir) ann
```

Listing 5: The Haskell definition of the `observe` behaviour of the `OBSERVATIONS` type class.

```
class STIMULI quality entity agent => OBSERVATIONS quality
    entity agent where
        observe :: quality entity -> agent -> agent

instance OBSERVATIONS Temperature AmountOfAir Thermometer where
        observe (Temperature amountOfAir) thermometer =
            thermometer {tValue = tQuale (perceive (Temperature
            amountOfAir) thermometer)}
ott = observe (Temperature muensterAir) fmoThermometer

instance OBSERVATIONS Temperature AmountOfAir Person where
        observe (Temperature amountOfAir) person = person {
            pValue =
                if (pQuale (perceive (Temperature amountOfAir)
                    person)) > 15 then  "warm" else "cold"}
otp = observe (Temperature muensterAir) ann
```

as being *warm* since the individual `muensterAir` mentioned earlier in listing 1 had a heat value of 20 degrees.

In the next sections we extend the FOOM ontology with trust and reputation as measures and show how the ontology can be extended to implement the WSMS use case presented earlier.

4.4 trust and reputation models for vgi: an ontological perspective

In contrast to the earlier discussed ontologies of trust in chapter 2, this work presents an observation centric approach to trust to enable integrating trust models into the sensor web with a particular focus on human sensor observations. In the following sections we address the foundational aspects of our ontological approach.

4.4.1 *Formalizing Trust and Reputation in the FOOM Ontology*

To be able to implement our trust and reputation models in the context of the human sensor observations it is essential to introduce the notions of trust and reputation to the FOOM ontology. This entails identifying the ontological nature of trust and reputation in the FOOM ontology so that we can integrate our trust and reputation models in the human sensor web. In this section we summarize our ontological commitments derived from our earlier study of trust and reputation in the previous chapters and show how can such notions be introduced to the FOOM ontology. We address our ontological commitments with respect to interpersonal trust then to informational trust as introduced in chapter 3 followed by our ontological commitment concerning reputation.

Interpersonal trust is a relational quality. What we mean by trust when we use the term generally in our work is interpersonal trust as the source of all other types of trust. As such interpersonal trust is a relational quality that can be represented as $T(a, b)$ so that we say volunteer a or the trustor has some level of trust in the volunteer b or the trustee. We model this in Haskell as a quality of the meriological sum of the Haskell type `Volunteer`. Listing 6 shows the meriological sum of two volunteers defined as the class `PersonalTie` which is a DOLCE `ARBITRARY_ SUM`. We then define interpersonal trust quality

Listing 6: The Haskell implementation of the `InterTrustQuality`

```
class QUALITIES quality entity

data PersonalTie = PersonalTie (Volunteer, Volunteer)
instance ARBITRARY_SUM PersonalTie
instance PERSONAL_TIE PersonalTie

data ARBITRARY_SUM personalTie => InterTrustQuality personalTie
    = InterTrustQuality personalTie deriving Show
instance QUALITIES InterTrustQuality PersonalTie
```

Listing 7: The Haskell implementation of the `InfoTrustQuality`

```
class QUALITIES quality entity

data InstitutionalTie = InstitutionalTie (Volunteer,
    WaterQualityAggregator)
instance ARBITRARY_SUM InstitutionalTie
instance INSTITUTIONAL_TIE InstitutionalTie

data ARBITRARY_SUM institutionalTie => InfoTrustQuality
    institutionalTie = InfoTrustQuality institutionalTie
    deriving Show
instance QUALITIES InfoTrustQuality InstitutionalTie
```

`InterTrustQuality` to inhere in the meriological sum i.e. `PersonalTie`. It is important to note here that the `PersonalTie` is a directed relation so that (a, b) does not imply (b, a).

Informational Trust is a relational quality. If there is $I(b, i)$ to say that b is the originator of information entity (i) then, *informational trust* is defined as $IT(a, i) \longleftrightarrow \exists b.T(a, b) \wedge I(b, I)$. This is to say informational-trust is implied by interpersonal-trust where the trustee is acting as a mediator between the trustor and the information entity generated by the trustee. In our case a volunteer creates a report and another volunteer can query the system, in which case the `WaterQualityAggregator` which is the trust model and is discussed later in this chapter is the link between the two volunteers and the trust relationship is between a `Volunteer` and a trust model, i.e. Haskell data type `WaterQualityAggregator`. Informational-Trust as a relational quality is defined on the sum of the volunteer and the water quality aggregator at a specific moment when querying the system. In a similar manner to interpersonal trust we take a short cut and create `InstitutionalTie` to model the sum of a `Volunteer` and a `WaterQualityAggregator` as a DOLCE `ARBITRARY_SUM` which is also a directed relationship. Listing 7 shows the implementation of the informational trust as a relational quality over the `InstitutionalTie`.

Both interpersonal trust and informational trust are contextualized. With respect to our discussion of the contextualization of trust in chapter 3, trust statements are valid only within the context of the WSMS so that volunteer a trusts volunteer b about water quality of a well.

Reputation is a relational quality. Reputation is the amalgamation of the perception of trustworthiness of a person by the community. It is in that sense a relational quality between a community and a volunteer. It is then said that the reputation of a person belongs to the sum of the information community in question and a volunteer. We introduce to the ontology the DOLCE class `SOCIETY` to represent an information community. We then introduce the `CommunityTie` as the

Listing 8: The Haskell implementation of the `Reputation` quality

```
class QUALITIES quality entity

data Society = Society
instance SOCIETY Society

data CommunityTie = CommunityTie (Society, Volunteer)
instance ARBITRARY_SUM CommunityTie
instance COMMUNITY_TIE CommunityTie

instance QUALITIES Reputation CommunityTie
```

Listing 9: The Haskell data `Volunteer`.

```
data Volunteer = Volunteer {vid:: Id, vloc:: Id, odorQuale::
    Bool, clarityQuale:: Bool, fullnessQuale:: Bool, vQuale::
    Bool, vReputation:: Float, honesty:: Bool, vReport:: Report}
     deriving Show
instance Eq Volunteer where Volunteer id1 _ _ _ _ _ _ _ _ ==
    Volunteer id2 _ _ _ _ _ _ _ _ = id1 == id2
instance PHYSICAL_ENDURANTS Volunteer
instance PHYSICAL_OBJECTS Volunteer
instance AGENTS Volunteer
```

DOLCE `ARBITRARY_SUM` of a `Society` and a `Volunteer` as a directed relationship as shown in listing 8; the `Reputation` quality inhering in a `CommunityTie`.

In this section we formalized trust and reputation in the FOOM ontology. In the next section we proceed to show our extensions of the FOOM ontology to accommodate the WSMS scenario.

4.4.2 *Implementation of the WSMS in the FOOM Ontology*

In this section we extend the FOOM ontology to implement the WSMS use case. Our aim is to show how can human sensors be integrated within the observation and measurement ontology with the associated quality control measures in the form of trust and reputation models. The implementation shows humans as sensors in the ontology in the form of the class `Volunteer` making observations and how the trust model itself is modelled as a sensor in the class `WaterQualityAggregator`. This vision enables the integration of any trust and reputation model like the models presented in chapters 5 and 6 in a sensor web platform for quality assessment of human sensor observations.

The first element of our ontology is the `Volunteer` as a type of `AGENTS` that can observe the environment. The `Volunteer` is modelled in the FOOM ontology as a sensor. The Haskell implementation is defined in listing 9.

A `Volunteer` can observe certain qualities of a water well through stimuli defined as detectable changes in the environment. As discussed earlier the type class `STIMULI` provides a `perceive` behaviour which is defined by its signature as a DOLCE perdurant. The `perceive` behaviour relates a quality inhering in some entity to an agent by generating a quale in the agent. The `Volunteer` as a sensor class has three qualia subject to the `perceive` behaviour, namely, odorQual, clarityQuale and the fullnessQuale which are set by observing the qualties `Odor`, `Clarity` and `Fullness` respectively. The implementation of the type

Listing 10: The Haskell Type Class `STIMULI` and the associated stimuli of the use case

```
class (QUALITIES quality entity, AGENTS agent) => STIMULI
    quality entity agent where
        perceive :: quality entity -> agent -> agent

instance STIMULI Odor WaterWell Volunteer where
        perceive (Odor waterWell) volunteer = volunteer {
            odorQuale = odor waterWell}

instance STIMULI Clarity WaterWell Volunteer where
        perceive (Clarity waterWell) volunteer = volunteer {
            clarityQuale = clarity waterWell}

instance STIMULI Fullness WaterWell Volunteer where
        perceive (Fullness waterWell) volunteer = volunteer {
            fullnessQuale = fullness waterWell}

instance STIMULI Drinkability WaterWell Volunteer where
    perceive (Drinkability waterWell) volunteer = volunteer {
        vQuale = drinkability waterWell}
```

Listing 11: The Haskell Type Class `QUALITIES` and the definition of water well qualities

```
class QUALITIES quality entity

data PHYSICAL_OBJECTS waterWell => Odor waterWell = Odor
    waterWell deriving Show
instance QUALITIES Odor WaterWell


data PHYSICAL_OBJECTS waterWell => Clarity waterWell = Clarity
    waterWell deriving Show
instance QUALITIES Clarity WaterWell


data PHYSICAL_OBJECTS waterWell => Fullness waterWell = Fullness
     waterWell deriving Show
instance QUALITIES Fullness WaterWell

data PHYSICAL_OBJECTS waterWell => Drinkability waterWell =
    Drinkability waterWell deriving Show
instance QUALITIES Drinkability WaterWell
```

class `STIMULI` is shown along with the perceive behaviour defined for all the qualities of our water wells in listing 10.

The constructor of the `QUALITIES` type class requires the qualities of `Odor`, `Clarity` and `Fullness` to be defined with their host, in this case the class defined by the Haskell data type `WaterWell`. The defintion of the type class `QUALITIES` and the respective qualities inhering in a water well is shown in listing 11.

In listing 11 we notice a `Drinkability` quality in addition to the qualities of a water well mentioned earlier. Ultimately the volunteer in our use case needs to make an observation about the drinkability of a water well through perceiving the three basic qualities of `Odor`, `Clarity` and `Fullness`. The relation between `Drinkability` as a quality of a water well and the `Odor`, `Clarity` and `Fullness` qualities of the water well requires us to address the notion of *Quality Spaces*, and particularly *basic vs. composed quality spaces* as discussed in [80].

Listing 12: The Haskell Type Class `OBSERVATONS` and the implemented drinkability observation

```haskell
class STIMULI quality entity agent => OBSERVATIONS quality
    entity agent where
        observe :: quality entity -> agent -> agent

instance OBSERVATIONS Drinkability WaterWell Volunteer where
    observe (Drinkability waterWell) volunteer = volunteer {
        vReport =
      if (odorQuale (perceive (Odor waterWell) volunteer) &&
          clarityQuale (perceive (Clarity waterWell) volunteer)
           && fullnessQuale (perceive (Fullness waterWell)
          volunteer)) then (if (honesty volunteer) then "
          drinkable" else "undrinkable") else (if (honesty
          volunteer) then "undrinkable" else "drinkable")}
```

In the foundational ontology DOLCE, the notion of quality space is introduced in analogy to the notion of conceptual spaces as discussed in [37], with the consideration that the DOLCE [69] ontologically define a quality space as the mereological sum of all quality regions at which qualities of a certain type are located. Taking that into consideration, according to [80] a basic quality space consists of a single quality dimension, i.e. the quality space orders a fundamental magnitude type that cannot be further decomposed such as a the quality space for temperature or mass. The same analogy applies to our qualities of `Odor`, `Clarity` and `Fullness` of the water well as basic quality spaces of a single quality dimension. Following this ontological account, the `Drinkability` quality is by contrast a composed quality space having `Odor`, `Clarity` and `Fullness` as its dimensions. From the perspective of the volunteer the `Drinkability` quality is a cognitively composed quality space. By observing the different quality dimensions, the volunteer makes an observation on the drinkability of the water well. This is done by expressing the drinkability quale (i.e. vQuale) in the intrinsic semantic reference system of the volunteer, distinguishing between "drinkable" and "undrinkable".

Given our definition of the `Drinkability` quality as a composed quality space, the perception of the stimuli is followed by observations which will illustrate the implementation of the composed quality of `Drinkability` as shown in listing 12. Here the volunteer can be honest or dishonest and by perceiving the three basic qualities of a water well invoking a quale corresponding to each quality. The earlier presented perception of a stimuli through the `perceive` behaviour in listing 10 does not yet result in a symbol or observation value, thus the composed quality space of the `Drinkability` quality is set accordingly producing the value of "drinkable" or "undrinkable" as an observation result by the volunteer through the `observe` behaviour defined in listing 12.

To this end we have modelled volunteers that can perceive stimuli and observe qualities of water wells. The perception followed by observation, both result in a symbol or observation value about a water well.

The reports of the volunteers are aggregated by the system to form final conclusions about the water well as drinkable or undrinkable. For example, some volunteers might report water well W_1 as drinkable while another group of volunteers might report the same water well W_1 as undrinkable. The question we have now is which group of volunteers should we believe to decide whether or not W_1 is drinkable

Listing 13: The Haskell data type `WaterQualityAggregator`

```
data WaterQualityAggregator = WaterQualityAggregator {aid:: Id,
    aQuale:: Bool, aReputation:: Float, aReport:: [Report]}
wqa = WaterQualityAggregator {aid = 1, aQuale = True,
    aReputation = 0.0, aReport = [report]}

instance PHYSICAL_ENDURANTS WaterQualityAggregator
instance PHYSICAL_OBJECTS WaterQualityAggregator
instance AGENTS WaterQualityAggregator
```

Listing 14: The Haskell implementation of the `WaterQualityAggregator` perceive behaviour

```
class (QUALITIES quality entity, AGENTS agent) => STIMULI
    quality entity agent where
        perceive :: quality entity -> agent -> agent

instance STIMULI Drinkability WaterWell WaterQualityAggregator
    where
    perceive (Drinkability waterWell) waterQualityAggregator =
        waterQualityAggregator {aQuale = drinkability waterWell}
```

or undrinkable? Here the system utilizes a computational model of trust and reputation as the model presented in chapter 5.

The FOOM ontology in general terms defines human and technical sensors as agentive physical objects in DOLCE terminology that act on stimuli. This is because the FOOM ontology takes a stimulus centric view of sensors [91] as opposed to the user-centric view taken in OGC's O & M [23] and the sensor-centric view in the OGC SensorML specification [12]. In analogy we define models to be sensors, in essence this is not different from the view of sensors in the mentioned specifications and as such seamlessly integrates trust and reputation models as sensors in the FOOM ontology. In our ontology we introduce the class `WaterQualityAggregator` as a sensor in listing 13 which is in our case a computational trust and reputation model.

Listing 13 shows how to encode a computational trust and reputation model in the FOOM ontology. The model is in essence a sensor and hence is an agentive physical object like all sensors in the FOOM ontology. It perceives drinkability reports from the volunteers as stimuli and utilizes an `observe` behaviour to produce a final conclusion on a water point based on all the drinkability reports of the volunteers to indicate if a water well is drinkable or undrinkable. Listing 14 shows the `perceive` behaviour of the trust and reputation model sensor, i.e. `WaterQualityAggregator`.

The perception of stimuli by the `WaterQualityAggregator` is not enough to produce observation results. The `observe` behaviour requires the `OBSERVATIONS` type class as presented in listing 15. Unlike listing 12 where the `observe` behaviour is implemented fully, it is unnecessary for us at this point to implement how the `WaterQualityAggregator` sensor produces its values. This task is done in detail in chapter 5 where the computational trust and reputation model is presented in details and implemented in a simulation environment. The `behaviour` in listing 15 is a short cut and serves to complete our ontology without cluttering it with the implementation details addressed in the next chapter.

Listing 15: The Haskell implementation of the `WaterQualityAggregator` `observe` behaviour

```
class STIMULI quality entity agent => OBSERVATIONS quality
    entity agent where
        observe :: quality entity -> agent -> agent

instance OBSERVATIONS Drinkability WaterWell
    WaterQualityAggregator where
    observe (Drinkability waterWell) waterQualityAggregator =
        waterQualityAggregator {aReport = [report]}
```

4.5 CONCLUSIONS

To be able to discuss our spatio-temporal trust and reputation models for the quality assessment of human sensor observation it is necessary as a first step to integrate such models into the larger vision of the sensor web and particularly of humans as sensors. The FOOM ontology seamlessly defines technical sensors and human sensors as observing agents. In this chapter we have presented an extension of the FOOM ontology [60, 79] that presents trust and reputation as quality assessment measures for observations of human sensors. In our ontology interpersonal trust, informational trust and reputation are presented as relational qualities. Trust and reputation models are presented in our ontology as sensors. The approach enables seamless integration of trust and reputation models into sensor web ontologies. We have implemented the WSMS use case and the ontology in the Haskell functional programming language which provides means for a rigorous algebraic specification as well as a simulation environment that enables testing the functionality of the ontology in a sensor web context, but this is however not a proof of the feasibility of spatio-temporal trust and reputation models themselves which we deal with in the coming chapters.

An important consequence of our ontological approach is the observability of trust and reputation. Since both are qualities, this implies that they are observable and can be treated as such in our ontology in later extensions. We envision a `Volunteer` having a quale for trust or reputation resulting from observing actions based on having trust in some other volunteer or an information entity. The ontological consequences of this approach to observing trust and reputation needs further study.

The `WaterQualityAggregator` trust and reputation model presented as a sensor in this chapter is implemented in the next chapters. We show two different computational models applied to the WSMS use case in an agent based simulation environment to assess the viability of our proposal of spatio-temporal trust models as proxy measure for quality of human sensor observations.

5

A TEMPORAL TRUST & REPUTATION MODEL FOR HUMAN SENSORS

In this chapter we present a temporally sensitive computational trust and reputation model for VGI, particularly for human sensor observations. The computational model is applied to the WSMS scenario introduced in the previous chapter. This model is the implementation of the type class `WaterQualityAggregator` introduced in the FOOM ontology integrating trust and reputation models as sensors in the sensor web technology stack. The computational model aims to discourage fraudulent, malicious or faulty reporting of volunteers by sanctioning abusers and using the computational components to dilute such negative effects as they occur. The effectiveness of the model in achieving its aim of predicting the true status of the water wells is tested and evaluated using an agent based simulation model (ABM) of the WSMS use case. Throughout this chapter, *model* and *computational model* refer to the computational trust and reputation model introduced in section 3.1.1 while *ABM* and *simulation model* are used interchangeably to refer to the simulation of the WSMS scenario introduced in section 5.2 which is used for evaluation of the computational model.

In section 3.1.1 we introduce the formal computational model and show how reputation, informational trust and the *temporal* dimension of informational trust is integrated in the computational model. We then use dynamics of the WSMS use case to show how the model parameters are initialized in this use case. In section 5.2 we discuss ABMs in general and our developped ABM in particular to introduce the types of agents in the model, the parameter space and the logical flow of the simulation as a system.

The rest of the chapter is dedicated to the evaluation of the computational model using the developed ABM. Before embarking on a deeper analysis we start with the experimental design and preliminary sampling of the model in section 5.3. This section leads to the development of formal performance measures of the computational model introduced in section 5.4. Using the results from the previous sections and the developed performance measures we conduct a full spectrum analysis of the simulation model in section 5.5. This analysis reveals a wide range of desirable properties of our temporal computational model of trust and reputation in quality control of human sensor observations. The effectiveness of the integration of the temporal dimension is investigated and proven to improve our performance measures. Finally, section 5.6 ends with conclusions and discussion.

5.1 MODEL DESCRIPTION

This section provides a formal description of the temporal trust and reputation model for human sensor observations. It is divided into two subsections, the first describes the artefacts of the computational model. The second, describes dynamic aspects of the model leading to initialization of the model parameters. The aim is to disentangle the computational artefacts from the dynamics of the WSMS use case.

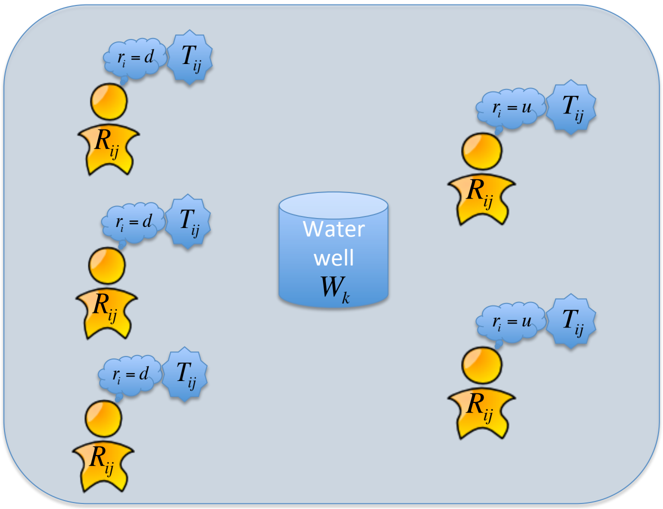

Figure 11: *An Illustration of the main elements of the computational model interactions.* w_k *is a water well.* R_{ij} *is the reputation of the user* i *in the system at history moment* j*, and* T_{ij} *is the informational trust value of a report* r_i *by person* i *at a history moment* j*.*

This shows that our model can fit other use cases as well, yet the dynamics are best illustrated through a concrete use case. The second section through the WSMS dynamics elucidates how to initialize the computational model.

5.1.1 *Computational Model Description*

Despite that fact that the the computational model design was dictated by the dynamics of the WSMS use case, we focus here on the computational artefacts of the model and congenitally touch on the dynamic aspects without much details opting to save this for the next section.

5.1.1.1 *Water Well Status Prediction in Initialization Phase*

Our model has an initialization phase and a post-initialization phase. In figure 11 we illustrate the main artefacts of the computational model at both phases. Water wells in our model are denoted w_z where the set of all wells is $W = \{w_1, ..., w_z\}$. The person shapes in the same figure 11 represent $i \in \mathtt{volunteer}$ where each volunteer makes a report r_i about a water well. Thus the set of all drinkability reports for a specific water well w_z is $D_z = \{r_1, ..., r_i\}$. This is to say person i makes a drinkability report r_i about water well where $r_i \in \{d, u\}$. Here $d = \mathtt{drinkable}$ and $u = \mathtt{undrinkable}$.

During the initialization phase we have no information other than new volunteers i and their reports r_i where $r_i \in \{d, u\}$. To initialize the system we use the majority rule by taking the toll of drinkable $r_i = d$ versus undrinkable $r_i = u$ reports for any water well w_z and make a prediction on the water well status from the volunteer reports. This rule is illustrated in equation 5.1. In other words if the number of all drinkable reports $r_i = d$ for a water well w_z is higher than that of the undrinkable reports $r_i = u$ for the same water well then the water well status will be predicted by our model to drinkable and vice versa.

$$w_z = \begin{cases} \mathtt{drinkable}, \text{ if } \#\{r_i \in D_z | r_i = d\} \geqslant \#\{r_i \in D_z | r_i = u\} \\ \mathtt{undrinkable}, \text{ if } \#\{r_i \in D_z | r_i = d\} < \#\{r_i \in D_z | r_i = u\} \end{cases} \tag{5.1}$$

In the following subsections we introduce the different artefacts of the computational model.

5.1.1.2 *Introducing Volunteer Reputations*

This majority rule is followed only during the initialization phase of the model and is used to initialize the next component of the computational model. To illustrate the next component of the computational model we briefly address one of the dynamic aspects discussed in the next section. We gather information about the correctness of a volunteer's reporting (i.e. a volunteer reports a water well status matching the reality of the current water well status) during the initialization phase. This information is used to gradually build user reputations using an approach similar to earlier work in [96]. Equation 5.2 is used to reward users who do honest reporting with increasing reputation while equation 5.3 is used to punish those who make fraudulent of false reporting. We use two equations so that we can adjust equation parameters to control the rate of increase or decrease in reputation to cause a slow build up in user reputations and a fast tarnish when abuse occurs. Both are desirable effects from our earlier study of reputation.

$$\widetilde{R}_{ij} = \alpha R_{ij} + (1 - \alpha) P_{pos} \tag{5.2}$$

where $P_{pos} \geqslant 0$

$$\widetilde{R}_{ij} = (1 - \alpha) R_{ij} + \alpha P_{neg} \tag{5.3}$$

where $P_{neg} < 0$

In equations 5.2 and 5.3 R_{ij} is the current reputation of a volunteer i at a specific moment j in the model's history where $R_{ij} \in [-1, 1]$. And $\widetilde{R}_{ij}$ is the newly computed value of the reputation for the same person i at the current history moment j based on older reputation value and the other model parameters. The parameters $P_pos \in [0 - 1]$ and $P_neg \in [-1, < 0]$ control the reward rate and punishment rate for increasing or tarnishing reputations of volunteers and α is a control parameter to be determined experimentally to optimize the reward and tarnish rates.

5.1.1.3 *Introducing Informational Trust & Temporal Dimension*

Since users in the model now have reputation values R_{ij} we introduce trust values for each individual report from each volunteer. Based on the method explained in the next section each report r_i created by a volunteer i at a specific moment in the model's history j is assigned an *informational trust* value T_{ij} corresponding to the notion of *informational trust* we discussed in section 3.3 in chapter 3.

Having introduced *informational trust* to the model we proceed to integrate the temporal dimension into the model in compliance with our proposal for extending *informational trust* with a temporal dimension. The water wells are by their nature not stable with respect to water quality such that a well that is drinkable today might be undrinkable as time passes. The *informational trust* of any report is subject to decay in its truth value given the change in water well status even if the volunteer making the report was honest in the first place. Such a decay rate for the *informational trust* of the reports has to be determined for each model implementation depending on the phenomena being observed. We introduce an exponential time decay function 5.4. In this equation informational trust T_{ij} assigned by a volunteer i at a specific moment in the model's history j of a given observation report decays in time t_j where t_j is the time passing from the moment j when the report is made. The parameter k controls the decay rate of *informational trust*. The expected effect will be that the conclusion on the water well will change with time as the truth value is impacted by the gradually decaying trustworthiness of volunteer reports.

$$\widetilde{T}_{ij} = T_{ij} e^{-k \times t_j} \tag{5.4}$$

Equation 5.4 integrates the temporal dimension of *informational trust* in the model, and as a result when studying the effects of time on the model performance we are studying the effect of the decay rate *informational trust*, namely the parameter k, on the performance of our computational model.

5.1.1.4 *Final Prediction of Water Well Status*

In the post-initialization phase we have volunteer reputations R_{ij} and informational trust values T_{ij} for each report $r_i \in \{d, u\}$ for all reports in a water well w_k. In this stage we need to determine the final prediction of the status of w_k. Let us denote the indices of the drinkable reports of all volunteers in all reporting histories j as $I^d := \{(i, j) \in \mathtt{volunteer} \times \mathtt{history} | r_{ij} = d\}$ and the indices of the undrinkable reports of all volunteers in all reporting histories to be $I^u := \{(i, j) \in \mathtt{volunteer} \times \mathtt{history} | r_{ij} = u\}$. In both cases the reports are sub-indexed with j to denote the moment in the model history when the report is made. The final prediction of the water well status is then determined based on equation 5.5.

$$w_z = \begin{cases} \mathtt{drinkable}, \text{ if } \sum\limits_{(i,j) \in I^d} R_{ij} \cdot T_{ij} \geqslant \sum\limits_{(i,j) \in I^u} R_{ij} \cdot T_{ij} \\ \mathtt{undrinkable}, \text{ if } \sum\limits_{(i,j) \in I^d} R_{ij} \cdot T_{ij} < \sum\limits_{(i,j) \in I^u} R_{ij} \cdot T_{ij} \end{cases} \tag{5.5}$$

Equation 5.5 is slightly similar to equation 5.1 in that it produces binary values of drinkable or undrinkable. However, instead of using majority rules and absolute count of reports, it uses sums of reputations and informational trust values to arrive at a final prediction of the status of a water well.

5.1.2 *WSMS Dynamics: Initialization & Model Operation*

In this subsection we discuss some dynamic aspects of the computational model to elucidate how the reputation values R_{ij} and the informational trust values T_{ij} are initialized in the model. These initialization processes for the values are later used in our ABM during model evaluation.

The reputation values can be initialized based on two approaches. The first approach is as follows:

1. through equation 5.1 we have reached a tentative prediction of the water well status as drinkable or undrinkable in a certain water well. The volunteers who lead to this prediction are then grouped into two factions. The first is the group that voted with the majority. The second is the group who voted against the majority.

2. volunteers who voted with the majority are then subject to equation 5.2 causing an increase in their reputation from an initial set value of $R_{ij} = 0$.

3. on the other hand those volunteers who voted against the majority are subject to equation 5.3, causing a drop in their reputation from an initial set value of also $R_{ij} = 0$.

4. after the initialization the process continues recursively during the model's life cycle after the initialization phase, causing a slow build up of reputations for honest contributors and a faster decrease in the reputations of the fraudulent or weakly skilled reporters. We remind here that the speed of reputation increase for honest volunteers and the speed of the reputation decrease for the dishonest or inaccurate volunteers is controlled through the P_pos and P_neg parameters of the model.

The weakness of this approach is that it could lead to unpredictable results if during the initialization phase groups of volunteers lobby to skew the majority results intentionally, but it is beneficial when the second approach using initial field validation is not feasible. The second approach is to use field validation during the intialization phase then fall back to approach one once a critical mass of validated volunteers of high reputation exists. This approach works as follows:

1. through equation 5.1 we have reached a tentative prediction of the water well status as drinkable or undrinkable in a certain water well. A team of validators checks the actual status of the water well and compares it to our model's prediction.

2. volunteers who reported honestly are then subject to equation 5.2 causing an increase in their reputation from an initial set value of $R_{ij} = 0$.

3. on the other hand those volunteers who reported incorrectly or dishonestly are subject to equation 5.3, causing a drop in their reputation from an initial set value of also $R_{ij} = 0$.

4. after the initialization the process during which we have ensured a group of honest volunteers has gained high enough reputations the model is left to work recursively as in the first approach leading to a continued increase of the reputations of honest volunteers and punishment of the dishonest ones.

The second approach above is the one followed in our ABM during the evaluation phase. An ABM allows us to assume knowledge about the water well status since we know the true status of any water well in the simulation at any given moment in time.

Having passed the initialization phase and having established volunteer reputations throughout the life cycle of the model, we elucidate how the information trust values T_{ij} for each report are initialized, where $T_{ij} \in [0, 1]$ such that:

- during the initialization phase of the system all reports are given $T_{ij} = 0$, which is unproblematic since the tentative conclusion on a specific water well does not in anyway depend on T_{ij}. It depends only on absolute counts of reports.
- reports of volunteers who have reputations $R_{ij} \geqslant 0.5$ are initialized at $T_{ij} = 1$
- reports of volunteers who have reputations below $R_{ij} < 0.5$ are initialized at $T_{ij} = 0.5$

Following establishing the values of informational trust T_{ij} the temporal dimension is integrated and enters the model via equations 5.4 and predictions of water wells statuses' are then computed based on equation 5.5.

5.2 OVERVIEW: WSMS SIMULATION MODEL

We have introduced the computational model in the previous sections. This computational model needs to be evaluated. The goal of the evaluation is to verify the success of our computational model in the quality control of human sensor observations in the WSMS use case introduced in chapter 4. This success is defined as the ability of our computational model to predict water well statuses, based on volunteer reports, that correspond to the true status of the water well. We use agent simulations to simulate the WSMS use case. This means that water well statuses are known at all times during simulation runs. We then implement our computational model to predict water well statuses and observe the correctness of our model's predictions and performance compared to the known statuses of the water well. Later in the chapter we develop rigorous performance measures of the computational model to assess its success. We start with a brief introduction to agent based modelling, then proceed to give an overview of the developed simulation model of the WSMS use case.

5.2.1 *Agent Based Modelling*

In our simulation we use the paradigm of agent-based modelling (ABM), which is a powerful simulation technique in which many agents interact according to rules resulting in the emergence of complex aggregate-level behaviour. ABM has been described as a third way of doing science, in contrast to induction and deduction [1]. Like deduction, simulations start with assumptions, but does not produce theorems, instead they produce data that can be analysed inductively. Unlike induction, the data of ABMs is the result of a strict set of rules and not of direct measurements of the real world. A crucial step in the modelling process is an analysis of how the system's behaviour is affected by the various model parameters. However, the number of controlling parameters and range of parameter values in an ABM is often large, the computation required to run a model is often significant, and agent-based models are typically stochastic in nature, meaning that multiple trials must be performed when making conclusions about model results. Our ABM is an agent based social simulation, where social simulation is defined as *the study of artificial societies of autonomous agents* [38, p.3]. We use the NetLogo[1] agent based modelling environment. Several agent based modelling environments exist, our choice of NetLogo is based on our experiments with other environments concluding that NetLogo is a good fit for our model.

It is generally understood that although ABMs are simulation models, most of them are not meant to provide an accurate representation or model of a particular empirical application. ABMs aim at enriching our understanding about the fundamentals of processes, and thus, the KISS principal (Keep it Simple Stupid) is vital in simulation design [1]. In a nutshell, while the topic investigated might be complex, the underlying assumptions of the ABM must be kept simple. Complexity of ABMs is in the simulated results not in the assumptions of the model [1]. In the following section we discuss our ABM with more details to highlight the main types of agents and the basic model interactions and its parameter space.

5.2.2 *Agent Based Modelling*

Figure 12 shows the interface designed for our ABM. The black ticks screen shows human shaped figures in three colours representing the three types of volunteers in this simulation run. The hut shaped figures are the water wells where green is a water well whose status drinkable and the red wells have the undrinkable status. The colours of the water wells change dynamically to reflect the status change. On the left hand side of the screen are the sliders that control the different parameters of our model's parameter space. The first group of sliders changes the number of wells for each run along with the requested number of the different types of volunteers as we later discuss. The second group of sliders control the computational model parameters. The parameters α, P_{pos}, $P_k neg$ are the parameters of the reputation equations 5.2 and 5.3 while the parameter k slider controls the decay rate of informational trust over time based on equation 5.4. With respect to the decay rate of trust over time, we have to define the nature of time in our simulation. How simulation time maps to real world time in the WSMS scenario is not the focus of our investigation. We assume each 24 simulation ticks to represent 1 day in simulation time. Thus, in our terminology when speaking about the simulation from this point onward we will speak of days as in *simulation days*.

There are two main agent classes in the simulation model, namely the volunteers and the water wells. There is only one type of water well in the model, which can have a status of *drinkable* or *undrinkable*. Moreover, the water well status switches over time between drinkable and undrinkable based on a probability distribution.

The volunteers in the model are of three different types of agents, namely:

---

[1] http://ccl.northwestern.edu/netlogo/

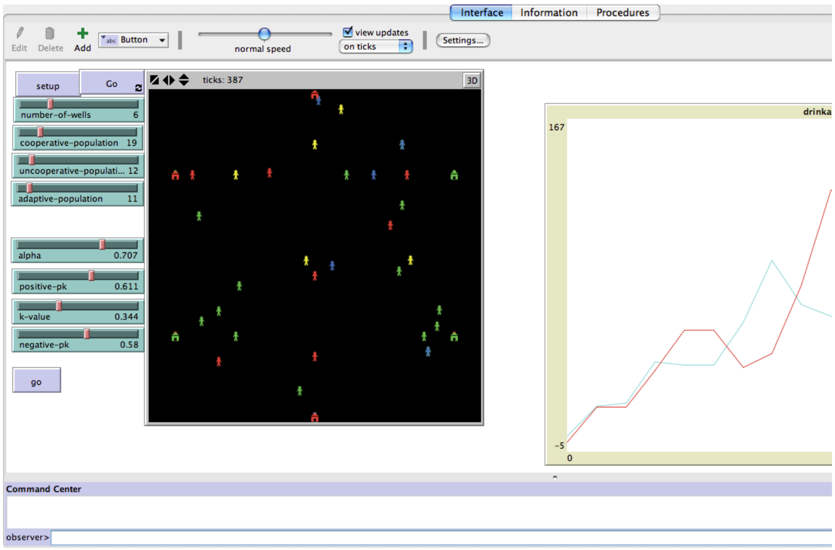

number-of-wells 6

cooperative-population 19

uncooperative-populati... 12

adaptive-population 11

alpha 0.707

positive-pk 0.611

k-value 0.344

negative-pk 0.58

Figure 12: *The WSMS Simulation Model Interface.*

- *Cooperative Volunteer:* this volunteer is consistently honest. Upon observing a water well, this volunteer reports the true status of the water well.
- *Non-cooperative Volunteer:* this volunteer is consistently dishonest. Upon observing a water well, this volunteer reports the opposite of the true status of the water well.
- *Adaptive Volunteer:* this volunteer has an inconsistent behaviour. According to a probability distribution, the volunteer autonomously chooses to act as honest or as dishonest.

We use a Bernoulli discreet probability distribution in the behavioural function of both the water wells (drinkable or undrinkable) and the adaptive volunteers (honest or dishonest) to produce dynamic behaviour change throughout the length of a simulation run. The standard Bernoulli function takes a value of 1 (drinkable for the water well or honest for the volunteer) with success probability p and value 0 (undrinkable, or dishonest) with a failure probability $q = 1 - p$.

To elucidate the functional design of the ABM, figure 13 shows its logical flow. A typical simulation run involves the following:

1. the population size of each volunteer type is chosen.
2. the starting behaviour of the adaptive volunteers is randomly initialized to either honest or dishonest. According to the random probability distribution discussed later, the behaviour alternates throughout the length of the run.
3. each volunteer randomly selects a target water well to move to. The volunteer maintains a memory of which water wells have been visited in a given day.
4. volunteers of all types are only permitted to report each well once in a single day.
5. based on a chosen probability distribution, the water well status changes through the life of each simulation run.

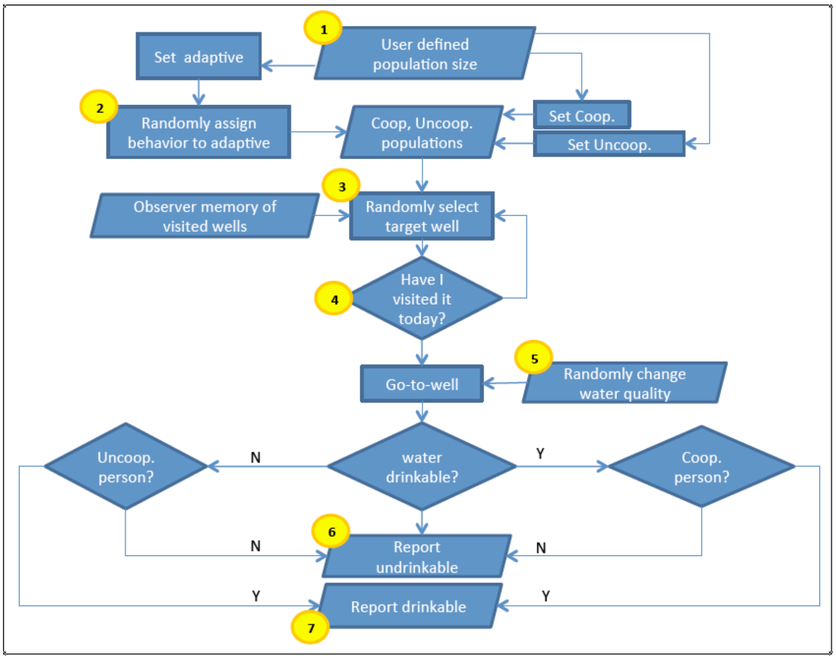

Figure 13: *A flowchart showing the logical flow of the WSMS agent simulation model.*

6. based on the true status of the water well and the volunteer type a drinkable or an undrinkable report is made.
7. the process repeats with the volunteer selecting a new target water well and moving towards it.

Finally, ABM validation is a topic of increasing importance as models and programming environments become more complex. In our ABM we are concerned with the internal validity of the model, defined as the correctness of the programming code in implementing the computational model, since the agent interactions are clearly simple to observe [2]. We implemented the computational model in Microsoft Excel spreadsheets and produced different sample results of the different model parameters and the corresponding values produced by the equations of the computational model. The values corresponded with the values produced by our ABM for the same parameter values providing evidence for the internal validity of the programming code of the computational model as well as proof of the soundness of agent interactions.

5.3 EXPERIMENTAL DESIGN AND BASE RUNS

Simulations of any real system like the WSMS scenario are subject to three conditions. First, there could be no data available from the real system, only output data available or both input and output data available. Our ABM falls in the first category, and as such, experiments with the simulation model need to be designed to obtain simulated data. Generally, these experiments should be guided by principals of Design of Experiments (DoE) Kleijnen and for Economic Research at Amsterdam [57]. However, before we propose a complete experimental approach we an initial exploration and sampling of the model using simple initial assumptions is essential. This is needed to understand what behaviours we are interested in for model evaluation before embarking on the deeper analysis that follow.

We resort to a novel approach for using Genetic Algorithm (GA) search. While DoE approaches traditionally attempt to do efficient sampling of the parameter space of an ABM (i.e. parameter space refers to parameters k, P_pos, P_neg, α, number of wells and populations types), this is a separate direction from the genetic algorithm search-oriented approach that we are pursuing here to assist in reaching an optimal experiment design. In particular, we are interested in the use of GAs to search our ABM parameter space for behaviours of interest. For a GA to yield meaningful results we have to define an objective function that describes the behaviour of interest that the GA should be searching for. The dilemma is that we do not know anything at this stage about the behaviour of the model, except that by definition the value of $\sum\limits_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ when higher than the value of $\sum\limits_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ for any given water well this means the opinion of the cooperative/honest volunteers wins over the opposing uncooperative/dishonest volunteers. With this knowledge we design a rudimentary objective function in equation 5.6 to sample the simulation behaviour. This means that the honest/cooperative volunteers dominate by the largest possible margin, resulting in the expected behaviour that our computational trust and reputation model effectively identifies the correct status of the water wells as it is in reality (by reality we mean the actual status of the water well in the ABM at any given moment in the simulation run).

$$Q = |\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij} - \sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}| \tag{5.6}$$

For this initial sampling of the model with our rudimentary objective function we use the BehaviorSearch[2] tool. We interface the tool with our ABM and use a standard GA to search the parameter space of the ABM. The GA uses an adaptive evolutionary approach to finding the best parameters that satisfy our predetermined objective function using a low mutation rate of 0.03 as recommended for this type of GA. The result of the GA search of the model's parameter space is depicted in table 1. The minimum and maximum range searched for each parameter is listed along with the optimal value found. The combination of these values satisfies the objective function set prior to running the GA.

Having established optimal values fulfilling our rudimentary objective function, we perform two different initial baseline runs. The first baseline run uses the exact optimal values found for all model parameters using the BehaviorSearch GA. Figure 14 shows the behavior of the first baseline run for one water well in a run. The horizontal axis is the number of ticks summarized in days (each simulation day is 24 simulation ticks) elapsed since the start of the simulation and the vertical axis are the values of $\sum\limits_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ and $\sum\limits_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ for drinkable and undrinkable observation reports respectively. The arrows in the graph indicate the time when the status of the water well changes. The model successfully reflects the true status of the water well for long period of

---

[2] http://www.BehaviorSearch.org/

Parameter Name	Min.	Step	Max.	Optimum
Number of wells	1	1	3	2
Cooperative Population	0	1	100	95
Uncooperative Population	0	1	100	35
Adaptive Population	0	0	0	0
α	0	0.01	1	0.35
$k-value$	0	0.01	1	0.19
Positive P	0	0.01	1	0.8
Negative P	0	0.01	1	0.61

Table 1: Optimal values of model parameters that satisfy the objective function of the GAs after searching the model's parameter space.

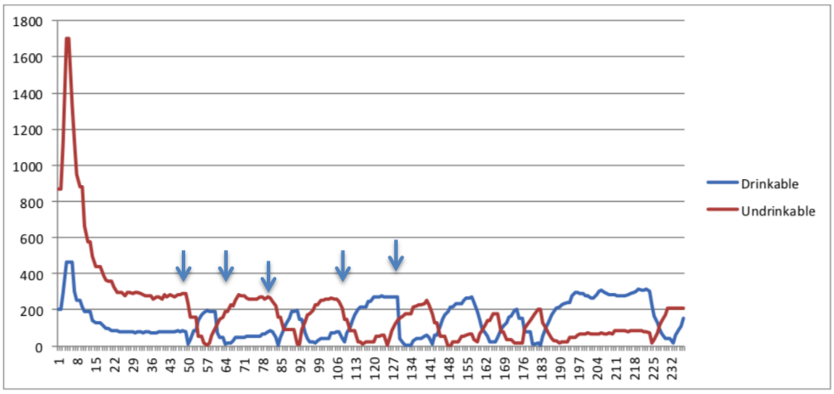

Figure 14: *The behaviour of the trust and reputation computational model in the base run. Days are shown on the* $x-axis$ *and values of* $\sum_{(i,j)\in I^d} R_{ij}\cdot T_{ij}$ *and* $\sum_{(i,j)\in I^u} R_{ij}\cdot T_{ij}$ *for a single well on the* $y-axis$*. The vertical arrows are the moment in time when the water well status switched between drinkable or undrinkable. Our computational model adjusts to reflects the true status of the water well as it changes.*

time using the optimal values of the rudimentary objective function, i.e. the honest strategy wins and misleading reporting is avoided. One important observation is an observed model latency where the water well status changes and the model takes a few days of reporting from volunteers until it reflects the true status of the water wells, which is to be expected in a real life situation. For example, in figure 14 the true status of the water well changed from drinkable to undrinkable at around day 45. The model continued to report an inaccurate water quality status with a latency of 3 days but then it accurately reflected the status of the water on the third day.

The second base model run uses the same optimal mathematical variables of the model found by the GA search of the parameter space in table 1. However, we set all types of volunteers to a maximum population of a 100 each (100 cooperative/honest, 100 uncooperative/dishonest, 100 adaptive volunteers). Figure 15 shows the experiment behaviour space engine while running the experiment, the graph showing model performance. Figure 16 shows a typical maximum population run.

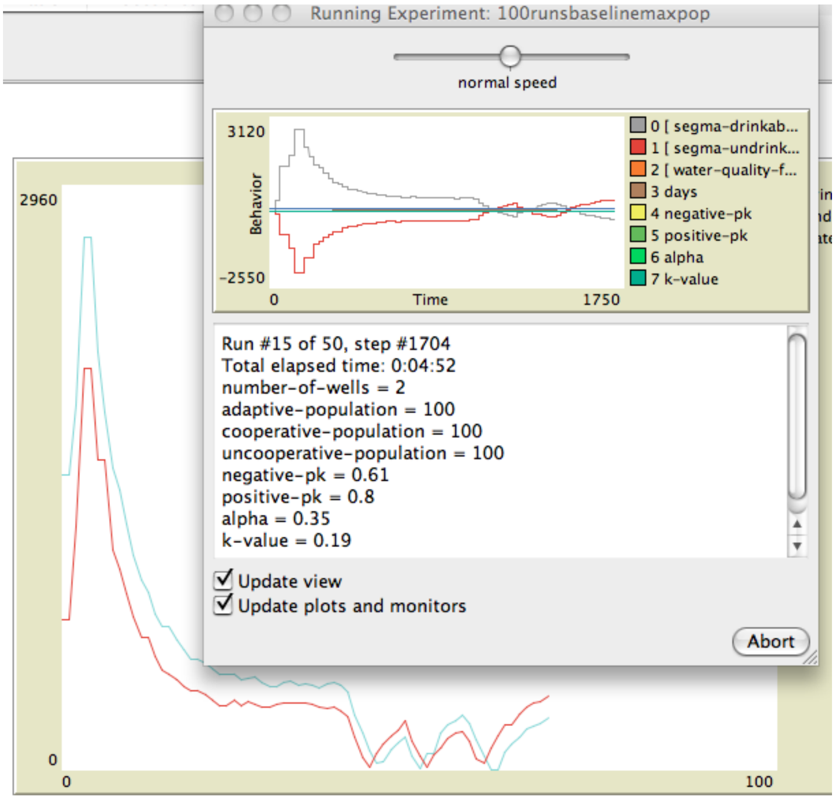

Run #15 of 50, step #1704
Total elapsed time: 0:04:52
number-of-wells = 2
adaptive-population = 100
cooperative-population = 100
uncooperative-population = 100
negative-pk = 0.61
positive-pk = 0.8
alpha = 0.35
k-value = 0.19

☑ Update view
☑ Update plots and monitors

Figure 15: *The behaviour space engine during the maximum population experiment. The top graph and the bottom graph snippet show the same behaviour for the drinkable and undrinkable* $\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ *and* $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ *values. The values shift to reflect our model adjusting to consistently report the correct current true status of the water well as it changes.*

Under this configuration, and based on averaging of a 100 runs the model is consistently resilient to strong attacks by both the dishonest and the adaptive population. Especially telling of the stability of our computational trust and reputation model is the resilience to the adaptive volunteers which represent a standard on-off attack, known in the trust models literature for technical sensor networks [18]. We also notice that the absolute values of $\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ and $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ are higher in the baseline run than in the maximum population run. This behaviour is to be expected since the standard deviation of both and $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ values is consistently lower for higher populations as a result of an increased number of observations.

Based on our initial investigation of the model using a rudimentary objective function and two baseline runs we observe the following:

1. the objective function used is not optimal. It is a maximization of the difference between the values $\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ and $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ for drinkable observations and undrinkable observations. This is an extreme objective given that our model would still report

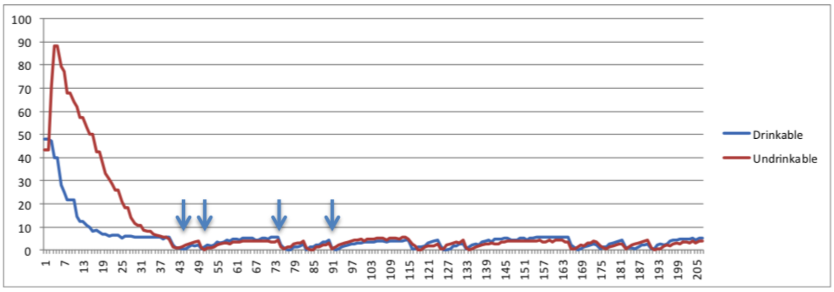

Figure 16: *The behaviour of the trust and reputation computational model in the maximum population run. Days are shown on the* $x - axis$ *and value of* $\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ *and* $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$ *for a single well on the* $y - axis$. *The vertical arrows show example moments in time when the water well status switched between drinkable or undrinkable. Our computational model still correctly reflects the true status of the water well.*

accurately as long as the value $\sum_{(i,j)\in I^d} R_{ij} \cdot T_{ij}$ is even marginally higher than that for and $\sum_{(i,j)\in I^u} R_{ij} \cdot T_{ij}$.

2. there is a clearly observable model latency. Which is defined as the number of simulation days between the change of status of the water wells and our computational trust and reputation model receiving observation from volunteers to reflect the current status of the well.

3. through experimenting with the simulation under different populations there is an adaptation phase. The computational model adapts to the true status of the water well and stays there until another change in well status. Under different model parameters the adaptation might or might not happen leading to model failure in predicting the true status of the water wells.

We are interested to see the effect the different numbers of volunteers have on the performance of the model, and also interested in understanding the effects of the time decay of trust $'k'$ has on this performance. In light of our observations, any further analysis of the simulation output to evaluate the performance of our computational trust and reputation model require as a first step a clear definition of what we mean by model performance. In the next section we introduce our formal performance measures that will guide the evaluation process..

5.4 MODEL PERFORMANCE MEASURES

Our ABM aims to assist us in understanding the properties and behaviour of our computational trust and reputation model. There are two different properties of the model behaviour we are interested in. One is how well the model predicts the actual status of the wells, which we refer to as model precision. The other is how long it takes for the model to adapt to a change in water well status, which we will refer to as the latency of the model.

5.4.1 *Model precision*

There are two different ways we calculate the precision of the model. The first tells us how often the status of the water well predicted by our model matched the actual status of the water well in the ABM at any moment in time. Let s_t^i denote the state of well i at tick t, and p_t^i denote the prediction of the model for well i at tick t. For every time-point T, the average *hit rate* $\bar{p}(T)$ of the model is:

$$\bar{p}(T) = \frac{1}{T|wells|} \sum_{i \in wells} \sum_{t=1..T} I(s_t^i = p_t^i), \tag{5.7}$$

Where $I(x)$ equals 1, if x is true, and zero otherwise. And a hit is a single instance of our model correctly predicting the current status of the water well. This measure is basically an average hit rate of the predictions taken over all wells, and time.

The other precision measure reflects the correct operation of the model in terms of adapting to a new well state. Whenever the state of a well changes, the model should adapt to the new situation and stay in the correct state until the real state changes again. This is because when the model is unstable due to e.g. attacks by uncooperative volunteers the hit rate can be reasonable, but the model is not stable in making predictions over time. So, between two real-state changes of the water wells, the model is called *adapted* if after an initial incorrect prediction (latency), it switches to the correct prediction and keeps it until the real state changes again. We denote the number of times the model adapted to the state of well i until time t as a_t^i. Furthermore, let c_t^i denote the number of times well i changed state until time t. Then the *adaption rate* at time T $\bar{a}(T)$ is:

$$\bar{a}(T) = \frac{1}{|wells|} \sum_{i \in wells} \frac{a_T^i}{c_T^i} \tag{5.8}$$

This is an average of the adaption rate computed over all the water wells.

5.4.2 *Model latency*

The latency measure quantifies the phenomenon that after each well-state change, it takes some time for the model until it receives enough volunteer reports to reflect the change. On each occasion, we count the number of ticks it takes for the model to adapt. It is important to note, that we count the required number of ticks only if the model adapted to the new state. That means that when the prediction changes again before the new state change, or if the model gives an incorrect prediction the latency is not counted. Let l_t^i denote the number of ticks the model required to adapt to the change of well i that happened in time t. Then the latency of the model in time T, $l(T)$ is:

$$l(T) = \frac{1}{\sum_{t=1..T} \sum_{i \in wells} a_T^i} \sum_{t=1..T} \sum_{i \in wells} l_t^i \tag{5.9}$$

This is simply an average of the latencies for all the adaptions where the model succeeded.

5.5 MODEL EVALUATION ANALYSIS

In the following sections we present the overall setup of the experiments and the different experiments conducted to study the behaviours of the model using our ABM.

5.5.1 *Experimental Setup*

We use different analysis packages to explore our ABM for behaviours of the computational model with respect to the earlier defined performance measures. The term *simulation run* in all our experiments refers to a maximum of 6000 ticks in the simulation. This number is determined based on the model reaching a stable state during the previous baseline runs. This number of ticks in each run enables the ABM to converge, where a converged ABM means our computational model followed the well state change or as we earlier mentioned it *adapted* at least six times in a row. Control conditions to halt a running simulation were defined as follows:

1. if the model changed its prediction falsely six times in a row, the model stopped. This means that if the model predicted opposite status to the current water well six times in a row during well status changes it was forced to stop.
2. if the simulation reached the 6000 tick count, it was forced to stop.

Generally, the goal of a simulation experiments design is to minimize the number of simulation runs required, while maintaining the ability to draw conclusions from the generated data. In our work we opt for a multi-phase approach. In the first phase, a 2-level factorial experiment design is run to determine the effect of the input parameters of our computational model on the three model performance measures. The conclusions of this experiment are utilized in the next phase, when several response-surface designs are applied to understand the detailed properties of the model. The number of runs for each experiment is discussed in the relevant section for each experiment along with other experimental design issues concerning each experiment.

5.5.2 *Model Parameter Interactions*

One of the characteristics of our computational model is the existence of many variables like the p_k and α. then there is the time decay factor of trust in the model k. Furthermore, through ABM we need to understand the performance of the model with respect to such variables and additionally with respect to variables like population sizes.

To discover main and interaction effects of model parameters on the performance measures of the model, a factorial experiment design can be applied [14]. A 2-level factorial experiment design constrains the levels of each input parameter that is tested to 2 (high and low level) e.g. 10 cooperative volunteers as a low and a 100 cooperative volunteers as a high for this particular parameter. In a full factorial experiment all possible combinations of the high and low level values of the input parameters are tested. In our case with 8 input parameters, a full 2-level factorial would require $2^8 = 256$ input combinations to test. Since the ABM uses random numbers, a sufficient number of replicates should also be computed in order to discover and discount the model's dependency on the random seed. With a modest replication of 50, this would imply $256 \times 50 = 12800$ simulation runs at 6000 ticks each. According to our initial tests, one simulation takes approximately 30 seconds to run, such an experiment would require approximately 106 hours to finish on a single computer. To lower the time needed to explore the model, i) we use a fractional factorial design instead of a full factorial design, and ii) we run the simulation on multiple computers.

To drastically decrease the time required to run the 2-level full factorial design, we employed a 2^{8-2} fractional factorial design [14]. In statistics a fractional factorial as opposed to a full factorial, standard designs have been developed to enable a selection of a subset (fraction) of the possible combinations of the variables in a full factorial while maintaining the ability to detect the interesting interactions between the different experimental parameters.

A 2^{8-2} fractional factorial design reduces the full factorial design to its $1/4$th, and requires only 2^6 input combinations to test (without replication). An important property of a fractional design is its resolution or ability to separate main effects and low-order interactions from one another. The higher the resolution, the more low-order interactions are unconfounded. The most important fractional designs are those of resolution III, IV, and V: Resolutions below III are not useful and resolutions above V are wasteful in that they can estimate very high-order interactions which rarely occur in practice. The 2^{8-2} design used here is a resolution V design, which can:

- estimate main effects unconfounded by three-factor (or less) interactions,
- estimate two-factor interaction effects unconfounded by two-factor interactions, and
- estimate three-factor interaction effects, but these may be confounded with other two-factor interactions.

By applying this design, the total number of simulations to run is reduced to 3200 full runs.

The fractional factorial design is implemented using the MEME[3] toolset. The main results are depicted on Figure 17, 18, and 19. In all three figures, we depict three interaction statistics charts. The top one presents the predicted effects, when we sampled the parameters at their limits. The middle figure shows the predicted effects when the parameters were sampled just within the limits. Finally, the bottom figures show the predicted effects in the 10% neighbourhood of the parameter combination that was found to be optimal using our first GA experiment results in table 1.

To illustrate how to read these charts we use figure 19 and a single parameter k as an example. The first chart is the predicted interactions of the parameters at range limits. This means that for example the time decay of trust k (denoted $k - value$ in all charts) is sampled at its two edges, these being 0 and 1. The second chart would show the interaction of the same parameter at two different values that are within the range limits $0 - 1$ range limit, here being 0.2 and 0.8. The last chart in the figure the $k - value$ is sampled at the 10% neighbourhood of the value of k in table 1. In the first chart, within parameter limits, the time

---

[3] http://www.aitia.hu/meme

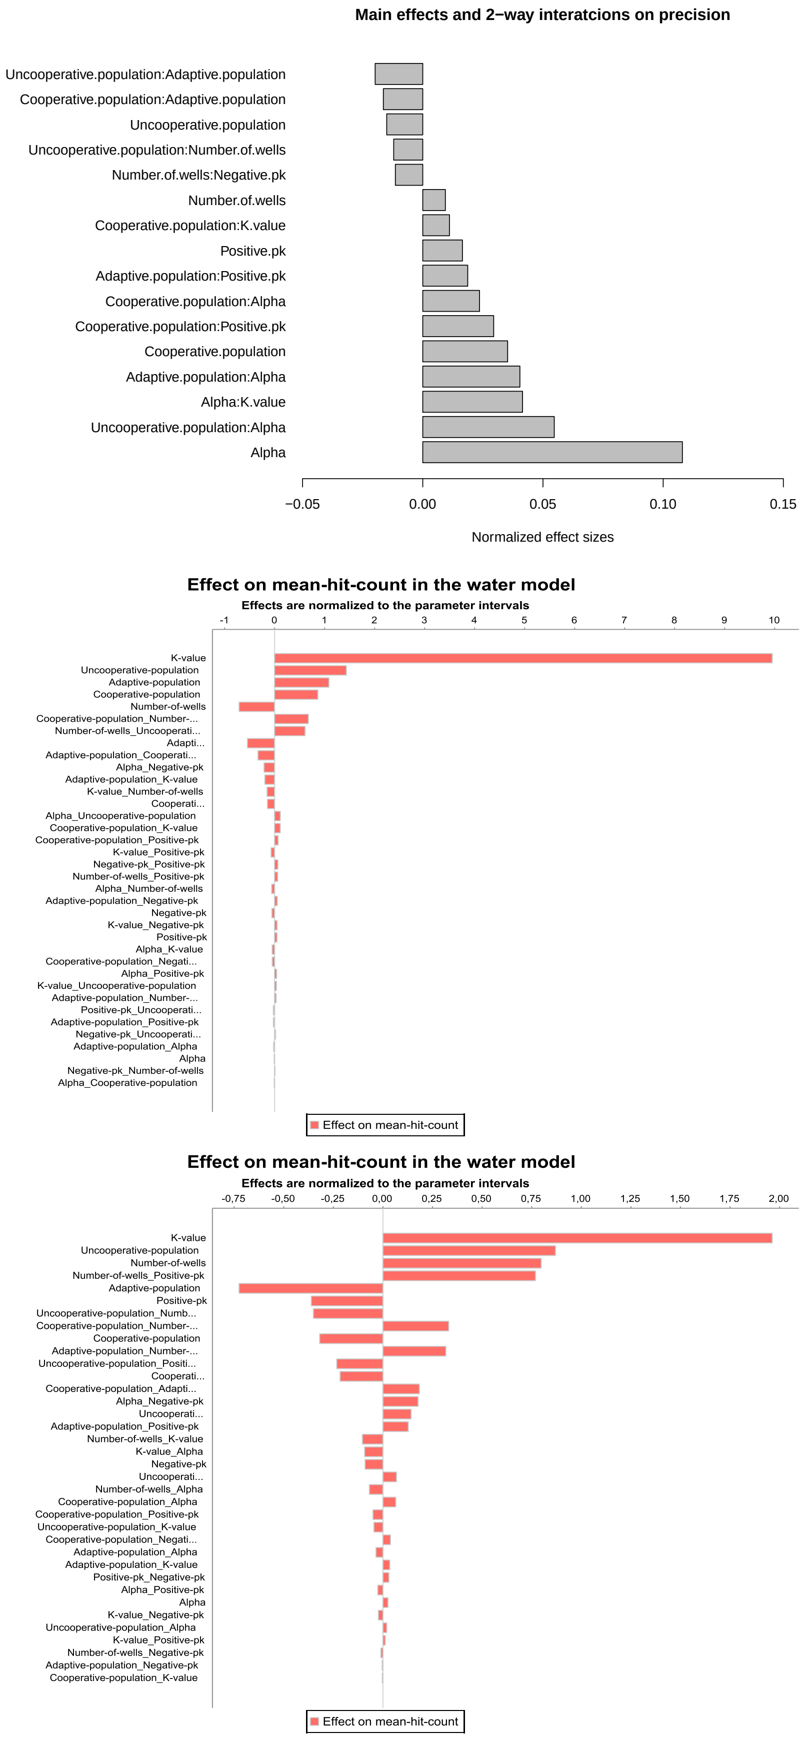

Figure 17: The effect of all parameters on hit rate (equation 5.7) measured at the range limits, within the range limits, and around the optimal parameter values in table 1

decay of trust k has a weak positive interaction value with the latency measure, which means it has very modest effect on increasing model latency, which is an *undesirable effect*. The second chart shows that the same parameter k has a very strong negative interaction, which means it has a very strong effect on reducing model latency which is a *desirable effect*. This is confirmed by the third chart in the figure as well. This is a surprising finding in its own regard. Namely that at range limits of the parameters the time decay of trust k has an undesirable effect on the model's performance but it has a very strong desirable effect on model performance within range limits.

From the current analysis we can conclude that 5 parameters have major effects on precision and adoption, and latency as our three model performance measures. The three population size parameters, the number of wells, and the time decay of trust k. What is remarkable in figures 17, 18, and 19 is that the predicted effects are drastically different when the model is sampled at the parameter range limits, than when it is sampled within the limits. While the middle and bottom figures differ only slightly, the top figure is very different. This suggests that the model has transient behaviour at the end of the ranges, particularly at the lower end of the range as we later discover and illustrate.

In the next sections, we find parameter combinations, that cause the model to perform optimally in terms of adaptation and latency as the two most important model performance measures. We do not do the same for the hit rate measure because adaptation as a measure encompasses hit rates that reflect a more stable model. We also proceed to study both populations and time decay of trust k to shed more light on the nature of their effects on the model.

5.5.3 *Optimization of Model Adaptation Measure*

In this experiment, we use the GA of the MEME toolset. As with the earlier GA experiment we must implement an objective fitness function for the GA to search for a specific behaviour in the model. The objective of the GA here was to maximize adaptation of the model, namely equation 5.8. The input parameters as genes were constrained as shown in table 2 along with the minimum and maximum range for the parameter values.

Concerning the GA settings, the population size was set at 50, and 40 generations were computed. The initial generation had randomly set genes. The selection operator selected the best 40% in terms of adaptation rate, then an averaging cross-over regenerated the lost 60% of the population selecting parents randomly from the survivors of the previous generation. After the cross-over operation, mutation with 5% chance were applied.

The collected data contained the input and output parameter values of not only the last, but, in fact, all generations. The best solution found is presented in table 2. The GA searches the parameter space of the model with the given settings to achieve a maximum value of adaptations of the model as an important performance measure. The last column (optimal) contains the best found settings of the model parameters. The last three rows show the values of the output variables of all performance measures of the model.

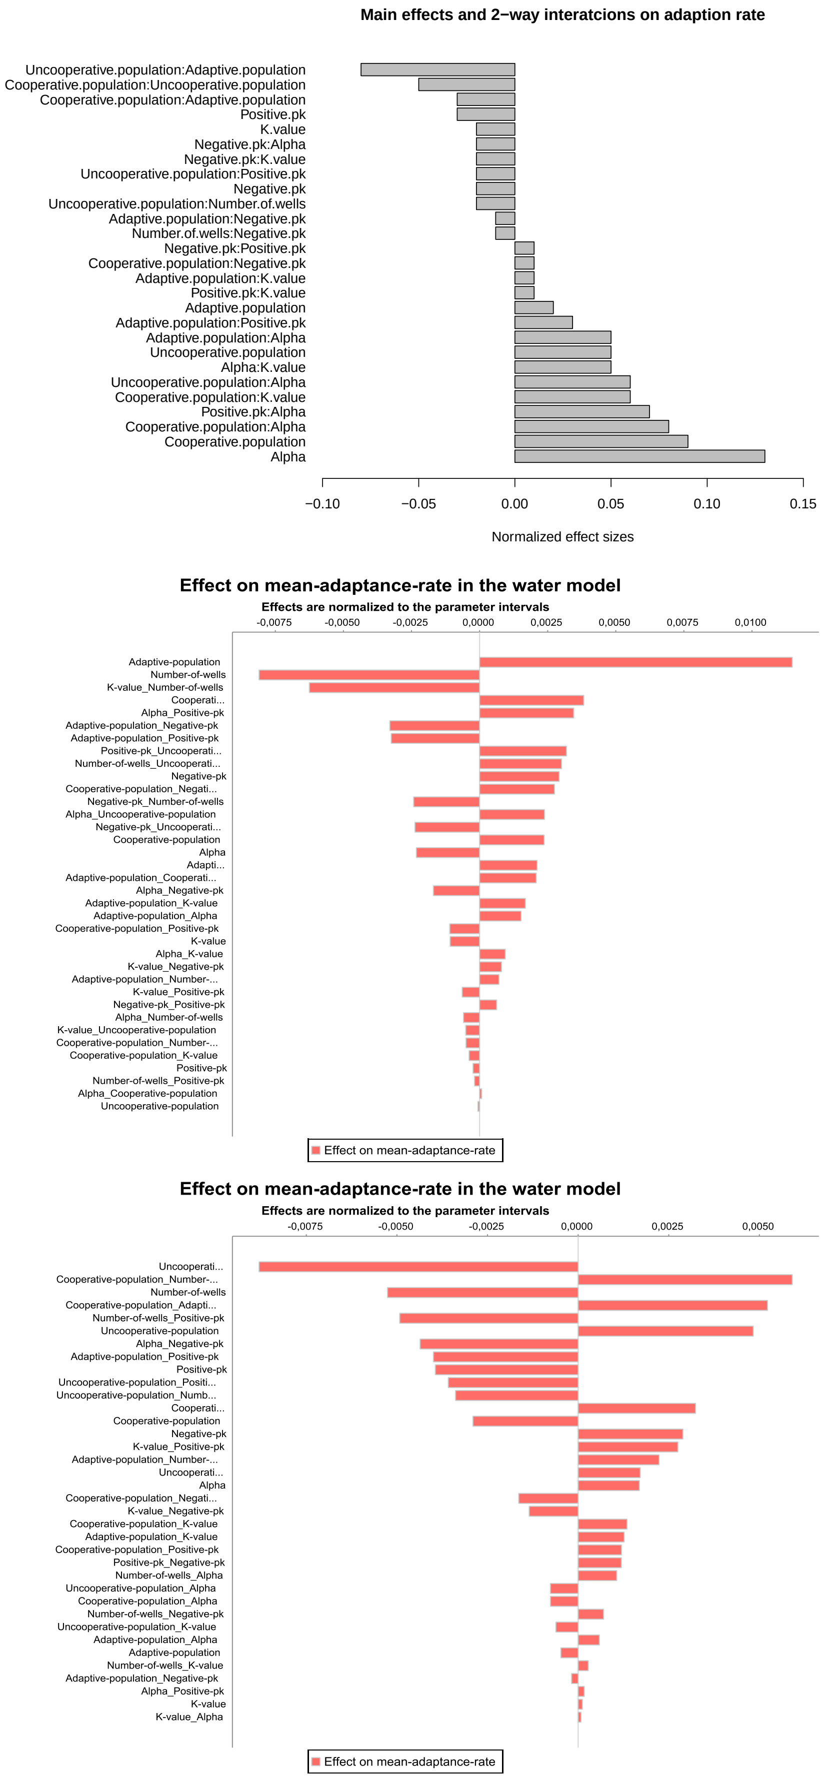

Figure 18: The effect of all parameters on adaptation (equation 5.8) measured at the range limits, within the range limits, and around the optimal parameter values in table 1

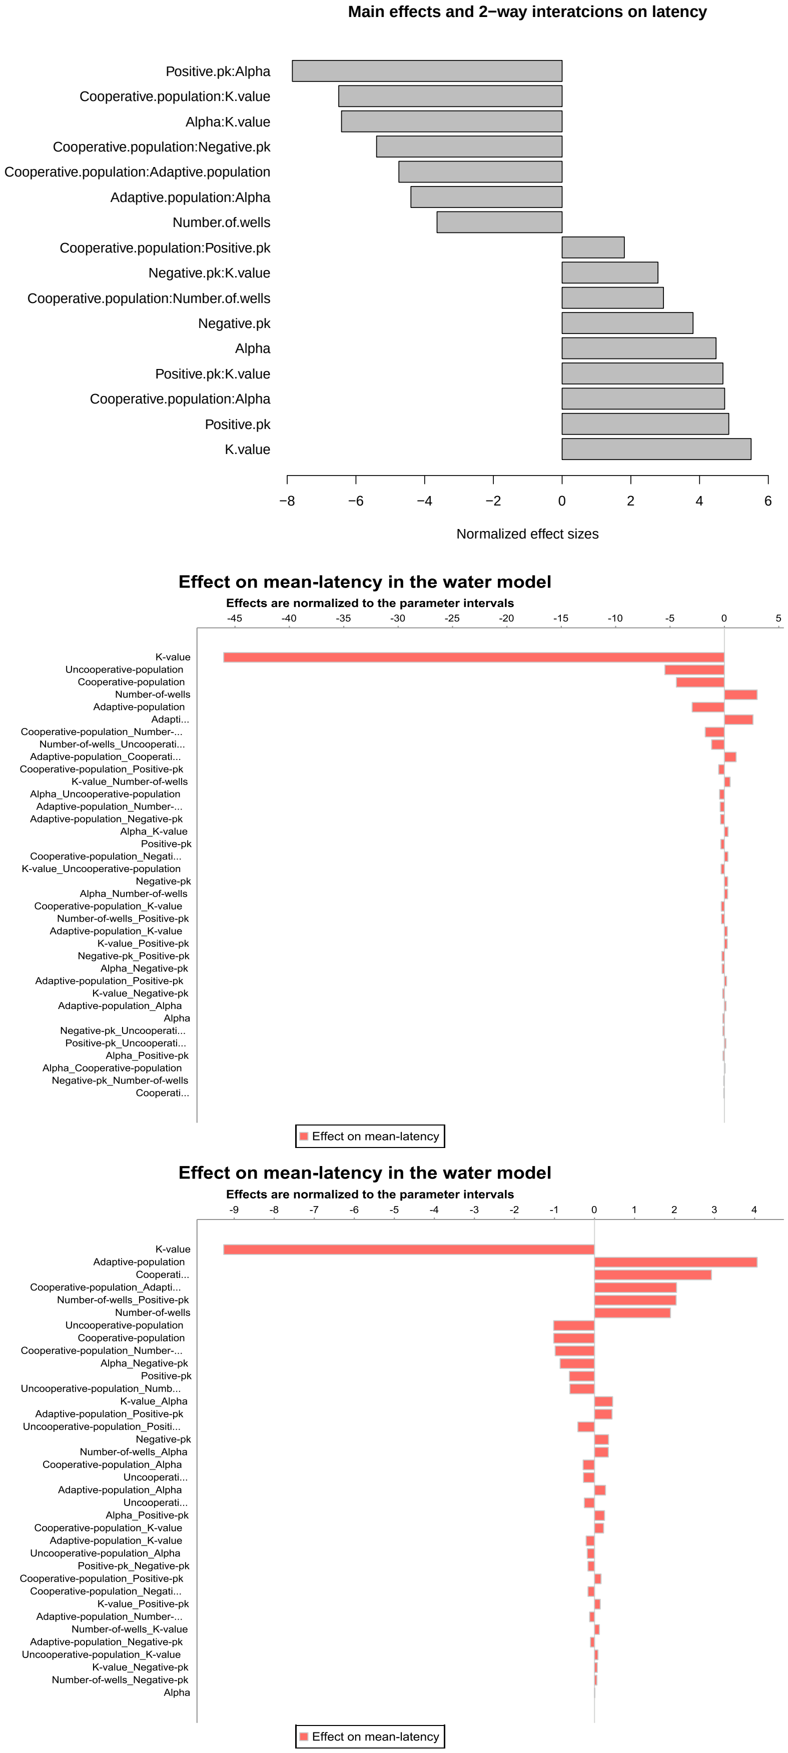

Figure 19: The effect of all parameters on latency equation (5.9) measured at the range limits, within the range limits, and around the optimal parameter values in table 1

gene	min	max	optimal
number of wells	2	4	2
cooperative population	0	100	55
uncooperative population	0	100	78
adaptive population	0	100	15
alpha	0	1	0.35
k	0	1	0.92
positive-P	0	1	0.7
negative-P	0	1	0.4

adaptation	1
latency	31.18
hit rate	0.92

Table 2: The gene constraints and optimal values found by the GA with the objective fitness function of maximum adaptation.

gene	min	max	opt
number of wells	2	4	2 (2)
cooperative population	0	100	64 (55)
uncooperative population	0	100	72 (78)
adaptive population	0	100	16 (15)
alpha	0	1	0.42 (0.35)
k	0	1	0.86 (0.92)
positive-P	0	1	0.56 (0.7)
negative-P	0	1	0.46 (0.4)

adaption rate	1 (1)
latency	22.125 (31.18)
hit rate	0.95 (0.92)

Table 3: The gene constraints and optimal values found by the GA with the objective fitness function of minimum latency.

5.5.4 *Optimization of Model Latency Measure*

This experiment is similar to the last GA experiment in section 5.5.3. The objective fitness function of the genetic algorithm is to minimize the latency of the model, namely minimize equation 5.9. The input parameters as genes were constrained as shown in table 3 along with the minimum and maximum range for the parameter values.

One caveat in analysing the results of this GA experiment is that optimizing for a minimal latency can be misleading. If the model does not converge on the real states of the wells, that is when no adaptations occur, the latency is zero. These cases should be omitted when searching for the best specimen in the recorded parameter values across all GA searches.

The population size was set at 50, and 40 generations were computed. The initial generation had randomly set genes. The selection operator selected the best 40% in terms of latency, then an averaging cross-over re-generated the lost 60% of the population selecting parents randomly from the survivors of the previous generation. After the cross-over operation, mutation with 5% chance were applied.

The collected data contained the input and output parameter values of not only the last, but, in fact, all generations. The best solution found is presented in table 3. The last columns contain the best found setting of the parameters. The last three lines show the values of the output variables. In parenthesis are the output values for the input that was found to be best by the previous GA experiment of maximum adaptation in section 5.5.3 for comparison.

Comparing the two GA runs, we can see that both found a solution where adaptation rate of the model equals 1, and that this GA experiment found a better solution than the previous experiment both in terms of latency and hit rate. The minimal latency achieved through our GA search is around 22 ticks or less than 1 simulation day. Meaning it takes less than 1 day for our model to reflect in its prediction the current status of the water well after it has changed.

In addition, the results of the GA experiments highlight the effect of populations of cooperative/honest and uncooperative/dishonest as well as adaptive behaviour populations on model performance. Further we need to shed light on the relation between adaptation and latency since the GA experiments indicate that there is a compromise between minimizing latency and maximizing precision as the two most desirable properties of our model. Finally, we need to understand the effect of the time decay of trust k on model performance, particularly on latency since our earlier fractional factorial experiments indicated strong interaction between k and model latency in particular, and at this point it is unclear what is the nature of the relation between time decay of trust k and latency.

5.5.5 *Effects of the Time Decay of Trust on Model Performance*

Our factorial experiment indicated a very strong interaction between the time decay of trust k denoted $k - value$ as illustrated in figure 19 in the second and third charts. While the first chart in the same figure (experiment at range limits) showed interaction in the opposite direction suggesting that our model might have transient behaviour at range limits of the k parameter. Thus, we focus on the latency as a performance measure in this experiment and its relation with the time decay of trust k to try and understand the nature of time decay of trust on model performance.

To discover how 'k' influences latency, we use the Iterative Uniform Interpolation (IUI) design in the MEME tool set. This design takes one parameter, and divides the whole range into predefined sections. It samples the model at the section boundaries, and then tests the output for gradient. If there is a significant difference between the two ends of a section, then in the next iteration this section will be further subdivided into more sections. This iterative sub-division of sections is continued until there is no more significant change found between the ends of any sub-section, or a user defined iteration count is reached.

We ran two experiments with the following common settings. The sections of parameter k which has a range of $0 - 1$ were divided into 10 subsections. A maximum iteration count of 3 was defined, and 10%

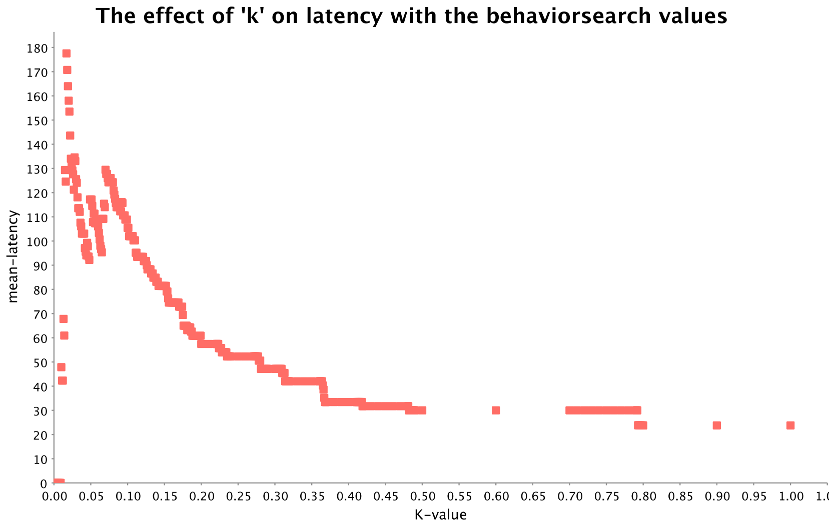

Figure 20: The effect of k on latency, with the rest of the parameters set to the BehaviorSearch-optimal

difference between section ends were taken as a significant difference. The IUI manipulates only one parameter, and keeps the rest fixed. In one of the experiments we used the optimal parameter values from the GA in table 1 tool found to be best for this model, the results of this experiment are summarized in figure 20. While in the second experiment, we used the values of the GA in table 3 which had the objective fitness function to minimize latency and the result of this experiment are summarized in figure 21. Both of these figures represent the k-latency relation in some optimal setting of the model parameters. Furthermore, both settings feature a low number of adaptive volunteers and higher levels of cooperative and uncooperative volunteers.

Both figures show the same trends. At the lower end of the range of the time decay of trust k, latency has a transient behaviour. It starts out low, and quickly rises to 150-180 ticks. After a peak around 0.05-0.1, latency of the model starts to drop. The drop is steeper in figure 20, but in the end they both finish around 30 ticks or about 1 simulation day as a latency value. This experiment proves two essential findings about the effects of time decay of trust k on model performance:

- the transient behaviour of the model is explained mainly by the time decay of trust.
- For low values of k the model performance is negatively and slightly erratically impacted. However, after this initial threshold the time decay of trust k has a strong impact on improving model performance through reducing the model latency to its minimum.

To see whether this shape of the k-latency curve is maintained in less optimal points of the parameter space, we ran a third IUI experiment. The parameter settings were basically taken from the optimal values for minimizing latency in table 3, except that we set the adaptive population parameter to 80, and the cooperative and uncooperative population to 10-10. As shown in figure 22, at this point of the parameter space, the model behaves very similar to the previous two parameter settings, especially to the one produced by the BehaviorSearch tool of NetLogo.

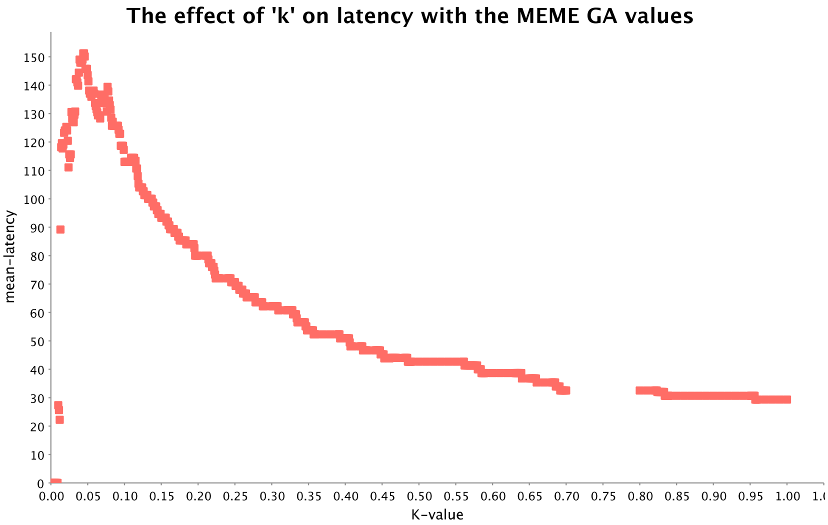

Figure 21: The effect of k on latency, with the rest of the parameters set to the MEME latency-optimal

Although there are differences in the exact numbers, the main features of the curves are the same. This suggests that the behaviour that higher k values lowers latency values leading to improved model performance is consistent.

This relation between k and latency as a model performance measure raises another question. What is the relation between the latency as a performance measure of the model and the precision measures, particularly the hit rate of the model? In the next section we conduct an analysis of this relation to shed light on the model performance goal of having low latency and high hit rates.

5.5.6 *The Model Precision-Latency Relation*

The findings so far in the previous sections raise a further interesting question. What is the relationship between model latency and model precision? Ideally we would like our model to take little time to reflect the change of a water well status (low latency) while also correctly reflecting the status of the water well most of the time with high precision (a high hit rate). Intuitively, we expect a negative relation between the two, since higher values of latency imply that it took a larger number of ticks for the model to adapt, which should result in a lower precision.

To discover this relation, we analysed data from extensive experiments of the effects of volunteer populations on the model (population experiments are presented in the next section). The dots in the left chart in figure 23 represents simulation results with the output that corresponds to the given coordinate. The results can be clustered into two clouds. The lower cloud (much fewer in number) are results where the model did not converge as indicated by the low level of precision (hit rate). These are results obtained from parameter combinations where k was very low, they correspond to the transient region in figure 21 reflecting the transient model behaviour. As such this cloud is essentially noise in the data. These dots are removed from the right-hand-side picture in figure 23 for the sake of clarity.

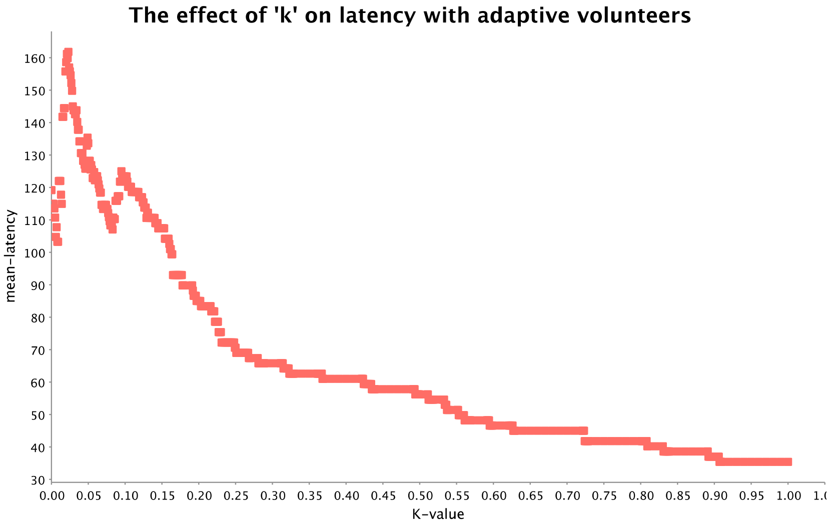

Figure 22: The effect of k on latency, with 80% adaptive volunteers and 10-10% (un)cooperative volunteers

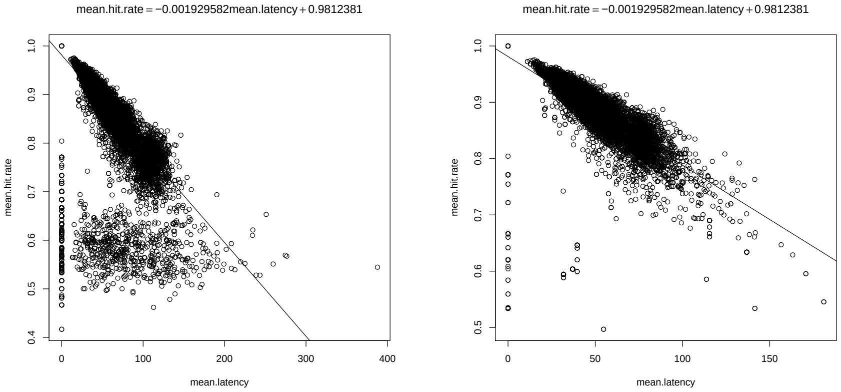

Figure 23: The relation of the two output, precision (mean.hit.rate in the Figure) and latency with (on the left) and without (on the right) noise

The filter expression used selected only results that were obtained from parameter combinations with $k > 0.1$. The remaining cloud, that contained the majority of the measurements, was used to fit a linear model on the data leading to a linear equation which can predict the relation between latency and hit rate for our model in all conditions:

$$\bar{p}(T) = -0.0019295821(T) + 0.9812381 \tag{5.10}$$

Where $\bar{p}(T)$ is the hit rate from equation 5.7 and $l(T)$ is the latency from equation 5.9. The straight line running through the cloud represents this fitted linear model. The fitted linear model is counter intuitive to our earlier assumption, namely our model can have a high hit rate while maintaining low latency achieving optimal performance through higher values for the time decay of trust variable k.

5.5.7 *The Effect of Volunteer Populations on Model Precision*

In this section we examine the effects of populations on the model performance with focus on precision of the model. In addition, we examine the relation between the time decay of trust and population changes while focusing on their effects on model precision. There are multiple questions one can ask about the effect of the population sizes on precision. How does:

1. the total number of volunteers
2. the ratio of different types of volunteers
3. the total number of volunteers and time decay of trust k

influence the precision of the model? To answer these questions, we set up another factorial experiment with 11 levels (0-100) of cooperative, uncooperative, and adaptive population counts and 11 levels (0-1) of k while focusing on model precision measures. The rest of the parameters were set according to the best parameter combination found to minimize model latency in table 3.

5.5.7.1 *The total number of different volunteer types*

To see how the total number of volunteers influence precision, we plotted results from parameter combinations where all three types of volunteers were represented in equal numbers. The result of analysing the first set of data is depicted in figure 24. The box plot showing a slight increase of model precision, particularly the hit rate as the population size grows (the total number of volunteers is 3 times the number on the x-axis, e.g 10 honest and 10 dishonest and 10 adaptive volunteers are represented as 10 on the x-axis). The increase is significant in the lower regions of $0 - 30$ volunteers, and insignificant at the higher end of the population scale.

This means that for our computational model to reach best precision a minimum of 30 volunteers is needed. This shows the resilience of our model to attacks from dishonest volunteers. Rarely in real life situations would the number of dishonest volunteers be equal to that of honest volunteers, and even then our model shows resilience as the number of dishonest volunteers increases.

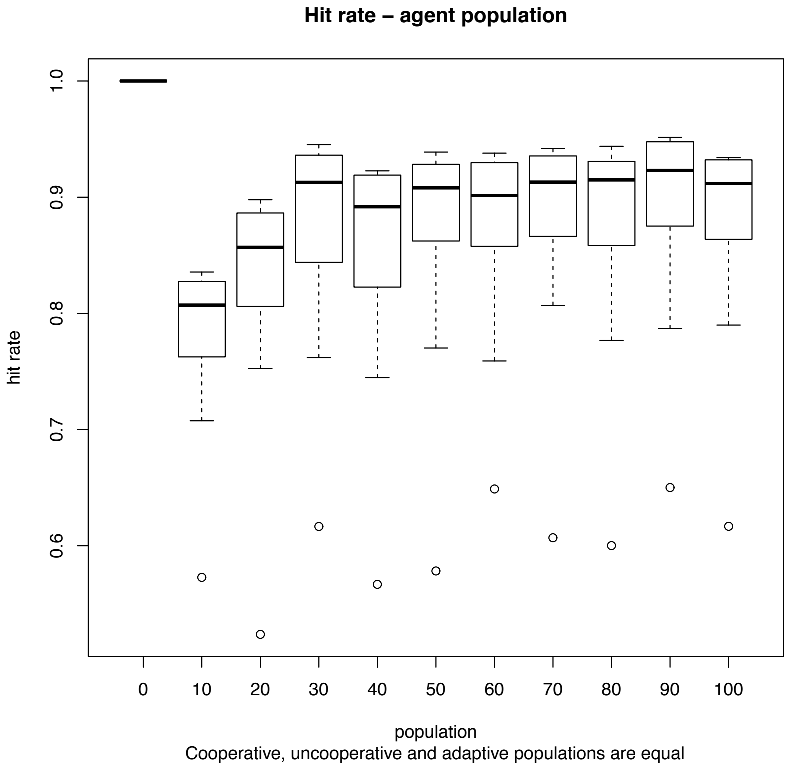

Figure 24: Precision as a function of population with equal numbers of all three types of volunteers.

5.5.7.2 *The total number of volunteers and time decay of trust* k

After we have seen that everything else fixed an increasing number of agents increase the precision of the model up to a point, it is interesting to look into how the size of the total population and k interact with the model's precision. The result of the analysis on these effects from our experiment is summarized in figure 25. The figure shows the interaction plot of the population size and k. In each sub-plot, the y-axis represents the precision of the model, particularly hit rate, the x-axis represents the input of the corresponding column, and the colours represent the corresponding row. E.g. in the top-left corner, the x-axis is the population size, the y-axis is the precision, and the colours are the different k values. In these plots, again, we used results from parameter combinations with equal number of adaptive, cooperative, and uncooperative volunteers, and the population sizes on the x-axis reflect the number of one type only, therefore they have to be multiplied by three to get the total number of volunteers. The plots in the diagonal show the main effects of the given input: in the lower-left corner the population size, in the top-right corner the k value. These charts correspond to results seen previously and confirm our findings about the model shedding more light on its functionality. That is an increasing number of population increases precision at lower numbers of volunteers, but when more than 30 volunteers per type are present further increase in the numbers do not increase the precision any further. Furthermore, as k increases, precision converges to 1, which is to be expected from our earlier analysis of the hit rate relation with latency in figure 23. Because k improves latency, and because latency is correlated with hit rate, it

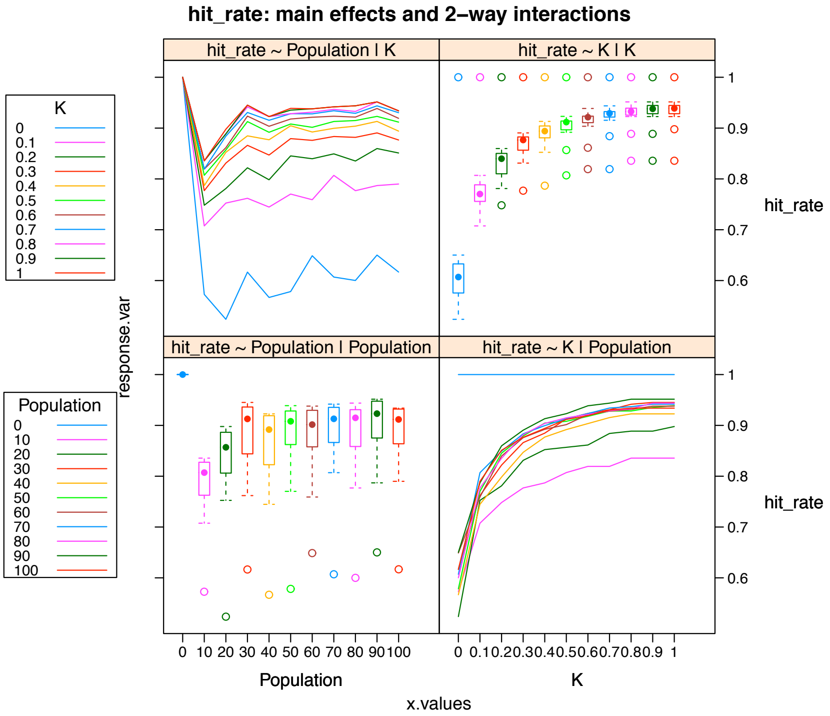

Figure 25: The combined effect of the population size and the time decay of trust k.

is to be expected that the time decay of trust k when increased also improves the model's hit rate.

The other two plots show the interaction between the population size and k. Both the top-left and bottom-right plot shows the same interaction, from different viewpoints. On the top-right the blue line representing the $k = 0$ case stands out from the rest. Here precision is hardly better than 50%, and an increase of population size (the x-axis) does not help much. Additionally, one can observe an instability in precision when the population size changes represented by the zig-zag feature of the blue line. For larger number of ks this irregularity disappears from the output, and a larger k clearly results in better precision. The same conclusions can be drawn from the bottom-right plot. For $k = 0$, the precision is low, and the different coloured lines are very close. As k increases the precision increases and all higher population levels converge to 1. This is the true reason behind the model's transient behaviour, namely for low k values the model is impacted negatively and as the value increases the effect is reversed and the time decay of trust improves model performance dramatically. This transient nature is clearly and emergent property of our computational trust and reputation model.

To summarize the interaction of the population size and k, we can say that at a low value of k ($k < 0.1$) the population size does not matter much, the precision is equally low at any number of volunteers. For increasing k, the effect of population on model precision is positive and the effect of k along on the model precision is strongly positive.

5.5.7.3 *The ratio of different types of volunteers*

After studying the model's behaviour for populations with equal numbers of cooperative, uncooperative, and adaptive agents, we focus now on the effect of the ratio of the different volunteer types. The analysis of the experiments' results are summarized in figure 26, the interaction of the cooperative-uncooperative and the adaptive population is depicted. The CP dimension encodes the number of cooperative and uncooperative volunteers such that the total number of non-adaptive agents is always 100, so a CP value of 10 for example means a setting with 10 cooperative and 90 uncooperative agents.

The diagonals show the main effects of the number of non-adaptive and adaptive volunteers. In the bottom-left plot, we can observe that the hit rate is generally stable for all cooperative-uncooperative combinations, except at the extremes, where it drops by a small amount. According to the top-right plot, the size of the adaptive population does not seem to influence precision at all. From the interaction plots we can conclude that the number of adaptive volunteers matter only when there are only cooperative, or only uncooperative agents. In these cases, any number of adaptive volunteers improves the precision of the model. In the range, when both cooperative and uncooperative volunteers are present in the simulations, the number of adaptive volunteers does not matter. This is a surprising find since the adaptive volunteers constitute a standard on and off attach on our model. From our current analysis it is unclear if this effect is due to our computational model or is an artefact of the design of the ABM.

As a conclusion we can say that population size has limited effect on the precision of the model. The number of adaptive volunteers matters only when there are either no cooperative or no uncooperative volunteers. Otherwise given more than approximately 30 volunteers per type, the model's performance is not further improved by the population sizes significantly meaning that a low number of volunteers is generally enough to ensure proper functioning of our model and that the model is resistant to attacks by adaptive volunteers.

5.6 CONCLUSIONS

In this chapter we have presented and evaluated our temporally sensitive computational trust and reputation model for human sensor observations with the WSMS use case scenario. We developed an ABM simulation of the WSMS scenario and applied our computational model to study model behaviour and performance. To achieve this we developed a rigorous set of performance measures, namely hit rate, adaptation and latency to measure different performance criteria for our computational model. Our analysis elucidates that our computational model shows remarkable resilience to different types of malicious behaviours and that it is an effective tool for triage of human sensor observations and VGI. Key to the success of our model is the temporal dimension integrated within informational trust whose effects are evaluated in our analysis as the effects of the k model parameter. Through our analysis of the time decay of trust we can confidently conclude that the sensitivity to the temporal dimension in our trust and reputation model directly and positively impacts model performance. We have

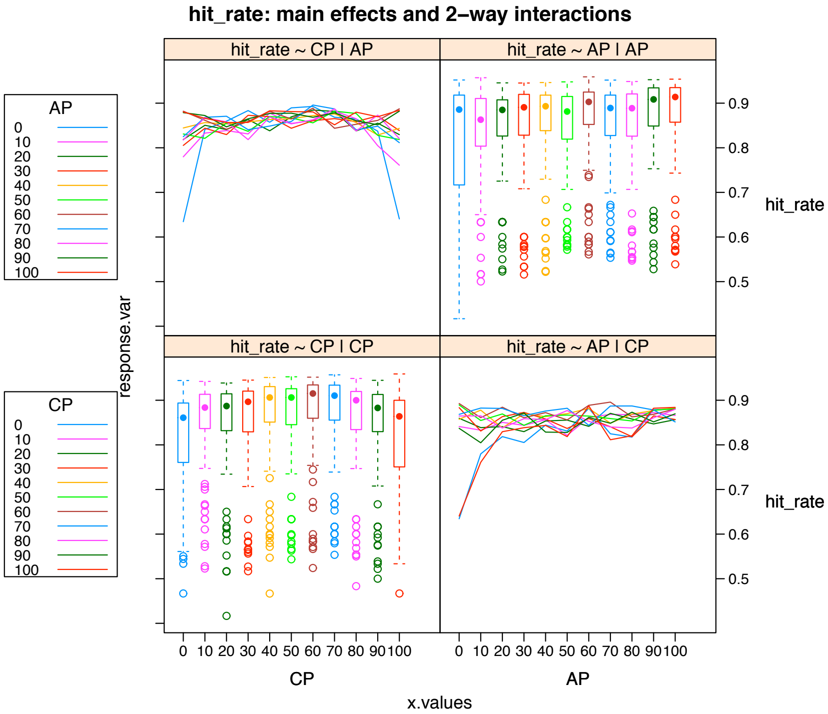

Figure 26: The combined effects of different ratios of volunteers on the model precision.

also proven the effectiveness of our computational model as a tool for quality control of human sensor observations.

6

A SPATIAL-TEMPORAL TRUST & REPUTATION MODEL FOR HUMAN SENSORS

In this chapter we present a computational spatial-temporal trust and reputation model for human sensor observations and VGI. Much like the model in the previous chapter, this model is another implementation of the type class `WaterQualityAggregator` introduced in the FOOM ontology integrating trust and reputation models as sensors in the sensor web technology stack. The computational model is applied to an extended WSMS scenario that adds spatial information and social networks to the WSMS scenario.

While the previous chapter introduced a temporally sensitive model, this chapter introduces a spatially and temporally sensitive model. In our evaluation we focus only on the spatial aspects of the model since the temporal aspects were extensively studied in the previous chapter. We isolate the effects of space and report on the impact of the spatial sensitivity of our model on the model performance measures of hit rate and adaptation developed and applied in the previous chapter. The terms *model* and *computational model* refer to the computational trust and reputation model introduced in section 6.2 while *ABM* and *simulation model* are used interchangeably to refer to the simulation of the WSMS scenario developed in this chapter.

We start by describing the extended WSMS scenario followed by description of our spatially and temporally sensitive trust and reputation model. We then present our developed ABM followed by experimental design and analysis of the experimental results. Our computational model encompasses a social network and sensitivity to the distance between the location of the volunteer at the time of reporting and that of the water well being reported. For evaluation of the computational model, the developed ABM reflects the extended WSMS use case in section 6.1. Our experiment focuses on the effect of the spatial sensitivity of the model on the model's performance and we show that the reporting distance has a positive effect on the model's performance under certain conditions as presented in section 6.4.

6.1 MOTIVATING SCENARIO: EXTENSIONS OF THE WSMS

In this section we present the extension of the WSMS use case from chapter 4. The H2.0 project team predicts that in the future the system will use smart phones to identify the location of volunteers at the time of making an observation. In the extended WSMS use case we address two main functionalities, namely:

1. *Volunteer location identification:* current smart phones in the market are increasingly equipped with Global Positioning Systems (GPS) sensors. The WSMS system depends on GSM networks to enable volunteers to report water status through their mobile phones. With GPS enabled smart phones the WSMS smart phone application can capture the location of a volunteer at the time of making an observation.

2. *Social Network capability:* in the WSMS use case presented in chapter 4 the volunteers are not associated with each other through social relations of any sort. Currently social network applications like FaceBook[1] are ubiquitous and as we postulated in chapter 1 the future of the web will lead to merging of the social silos into a unified social graph. The WSMS application will tap into the social network and enables volunteers to build social relations with other volunteers with associated trust ratings about how each volunteer finds another one to be trustworthy akin to the work done in [39] as discussed in chapter 2.

In this scenario a volunteer reports a water well much like in the use case presented in chapter 4 and the smart phone application logs the location of the volunteer and our computational model uses this information along with triage of the social graph to generate informational trust values and produce a prediction of the water wells' statuses.

6.2 MODEL DESCRIPTION

In this section we present our spatial-temporal computational model for trust and reputation. We follow a similar path to section 5.1 in chapter 5. We differ in the first section 6.2.1 where we first present the conceptual model behind the computational model. This helps in clarifying the functionality and computational artefacts of the model, which are then presented in the following section 6.2.2. We then end by describing the model dynamic initialization in section 6.2.3.

6.2.1 *Conceptual Model Description*

In the extended WSMS system, when a volunteer makes a water well observation report the resulting data model can be viewed as a bipartite graph. The resulting bipartite graph of volunteers and observations is a two mood non-dyadic affiliation network, see figure 27.

In this figure on the left hand side we observe $D_z = r_1, ..., r_{ij}$ is the set of all reports r_{ij} made by volunteers V at each moment in history j of the model as it runs. And $r_{ij} \in \{d, u\}$ represents drinkable or undrinkable value or each report. In the same figure for each well z there is D_z as the set of all reports where $D_z^d \subset D_z$ is the subset of drinkable reports and $D_z^u \subset D_z$ is the subset of undrinkable reports. This results in two sets of reports for a single well z as shown in figure 27. Let $\mathbb{V}$ the set of all volunteers in our model reporting on all wells, then $V \subset \mathbb{V}$ on the right hand side of figure 27 is the set of volunteers reporting on a single well. For this set V there is a sink volunteer (e.g. v_2) for which we will later calculate interpersonal trust values from the social network based on ratings on the social ties shown on the right hand side of the figure (e.g. Jack, Alice are friends).

In the same figure the arrows in the middle from the set V nodes to the D_z nodes are virtual links between volunteers are report sets that carry a weight of the geographic distance between a volunteer on the right side and a well z (e.g. link between v_4 and D_k^d). This distance weighting is later used to integrate the spatial dimension in the computational model as we elucidate in subsection 6.2.2.2.

---

[1] http://www.facebook.com/

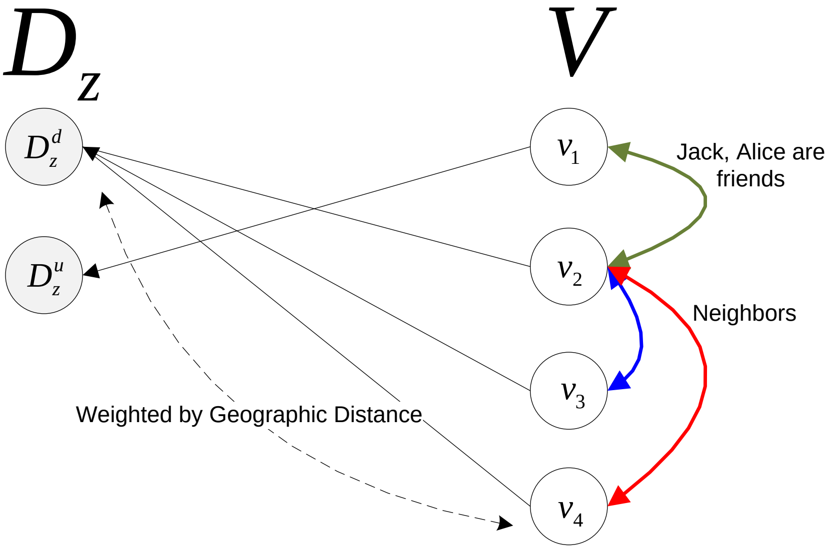

Figure 27: *An affiliation network of volunteers (right hand side) and the observation reports they contribute (left hand side).* D_z *is the set of all reports for a single well* z *where* $D_z^d \subset D_z$ *is the subset of drinkable reports and* $D_z^u \subset D_z$ *is the subset of undrinkable reports.* $V \subset \mathbb{V}$ *is the set of volunteers reporting on a single well. The volunteers are connected with a social network with trust ratings shown as flexible arrows on the right hand side.*

In the next section we proceed to describe the computational model which is built on this conceptual model in detail.

6.2.2 *Computational Model Description*

In this section we describe the spatial and temporal computational model. This model, in some parts, builds on the model introduced in the previous chapter. We start with the prediction of the water well statuses in the initialization phase of the model following similar approach to chapter 5 by not discussing the dynamic aspects in detail for clarity and saving it for section 6.2.3. We then introduce the spatial dimension to the computational model followed by the social network and interpersonal trust. We then introduce informational trust and the integration of both the spatial and temporal dimensions in our unified model and end with the prediction of the final status of water wells.

6.2.2.1 *Water Well Status Prediction in the Initialization Phase*

During the initialization phase we assume limited knowledge during the initialization phase like chapter 5. We avoid repeating the same discussion and refer the reader to section 5.1.1.1 which is the same approach used to initialize this model. The prediction of water well statuses is then done using equation 5.1 as done in the previous chapter.

It is important to notice that the result of this initialization phase is the emergence of reputation values R_{ij} for volunteers which are then used to introduce the next computational artefact of our model in the following section. We also build each computational artefact separately for clarity and end by combining them in the last section.

6.2.2.2 *Introducing Spatial Dimension & Reputation*

This component is inspired in part by the notion of nodal degree in graph theory denoted $'d(m_h)$ (see [97]). The nodal degree is simply the count of links connected to a single graph node. For example, in figure 27 the node representing the set of drinkable reports D_z^d has a nodal degree $'d(m_h) = 3$ based on the three links it is connected to.

We take this a step further to integrate by not exactly counting the links but by taking absolute sum of the reputation values R_{ij} of the volunteers V creating those links as weights to the links and using the distance between the volunteer and the well as an inverse weight. As in the previous chapter, we denote the indices of the drinkable reports of all volunteers using subscript i in all reporting histories using subscript j as $I^d := \{(i, j) \in volunteer \times history | r_{ij} = d\}$ and the indices of the undrinkable reports of all volunteers in all reporting histories to be $I^u := \{(i, j) \in volunteer \times history | r_{ij} = u\}$.

$$O_z^d = \sum_{(i,j)\in I^d} \frac{R_{ij}}{\log(c_{ij})} \text{ where } c_{ij} > 1 \tag{6.1}$$

$$O_z^u = \sum_{(i,j)\in I^u} \frac{R_{ij}}{\log(c_{ij})} \text{ where } c_{ij} > 1 \tag{6.2}$$

Then the first component of our computational model is a virtual intermediate component for each D_z^d and D_z^u nodes denoted O_z^d and O_z^u respectively. This component is an intermediate computational artefact to integrate reputation and geographic distance into our computational model as we later illustrate. Equations 6.1 and 6.2 show how to compute both components for drinkable and undrinkable reports which are identical in each case. In both equations c_{ij} is the geographic distance between a volunteer and the water well he is reporting on at the time of reporting using a log function to normalize large distance variations. R_{ij} is the reputation of the volunteer at the time of reporting. The computation is done once for drinkable and once for undrinkable reports.

The virtual computational artefacts O_z^d and O_z^u are then used later in the computation of the informational trust values for drinkable and undrinkable reports when making final predictions about water well statuses. In the next section we integrate the social network and the notion of interpersonal trust also discussed in chapter 3 into the model.

6.2.2.3 *Introducing the Social Network & Interpersonal Trust*

We then proceed to integrate interpersonal trust through the social network as another model component which is the combined vested trust (CVT) S_i for each volunteer 6.3. It represents a trust calculation at the social network level between the volunteers resulting in a CVT value for each volunteer using the subscript i and denoted S_i.

$$S_i = \frac{\sum_{i\in V} s_{il} \cdot y_{il}}{|V|} \tag{6.3}$$

In figure 27 let social tie links be denoted l, then for example $s_i l$ is the trust in volunteer v_2 as assigned by volunteer v_1 through a link l. And $y_i l$ is a weight that can be assigned to each social link by the volunteers denoting the importance of this relationship to them. The value $|V|$ is the cardinality of the set V. Then the CVT is computed for each volunteer through all social tie links leading to a volunteer node using equation 6.3. At the end of this recursive process across the social network of all volunteers in the subset V of volunteers who reported on a well z each volunteer has a CVT value S_i.

6.2.2.4 *Introducing Informational Trust & Combining Temporal & Spatial Dimensions*

In this step we compute informational trust values for the two possible conclusions on a water well (drinkable and undrinkable). This is akin to two computational trust values for the two nodes denoted D_z^d and D_z^u in figure 27 . In this step we also integrate the temporal dimension of informational trust. The results are informational trust value for node D_z^d denoted T_z^d in equation 6.4 and for node D_z^u denoted T_z^u in equation 6.5.

$$T_z^d = \sum_{(i,j) \in I^d} S_i \cdot O_z^d \cdot e^{-k \cdot t_j} \tag{6.4}$$

$$T_z^u = \sum_{(i,j) \in I^u} S_i \cdot O_z^u \cdot e^{-k \cdot t_j} \tag{6.5}$$

In both equations 6.4 and 6.5 we combined the previously presented artefacts to produce two informational trust values on the drinkable then on the undrinkable conclusions. In the process the artefact $e^{-k \cdot t_j}$ integrates the temporal dimension using the same method as in equation 5.4 in chapter 5.

The next section uses the computed informational trust value to make predictions of water well statuses.

6.2.2.5 *Final Prediction of Water Well Status*

The determination of the status of a water well is then done by comparing the informational trust values of the two nodes representing a specific well w_z as shown in equation 6.6

$$w_z = \begin{cases} \mathtt{drinkable}, \text{ if } |T_z^d| \geqslant |T_z^u| \\ \mathtt{undrinkable}, \text{ if } |T_z^d| < |T_z^u|| \end{cases} \tag{6.6}$$

In the next section we discuss some initialization and operational aspects of the model to further clarify the model's description.

6.2.3 *Extended WSMS Dynamics: Initialization & Model Operation*

For the computational model to function some parameters need to be initialized, or assumed to exist. We highlight here the following parameters and our assumptions about them building on the work done in chapter 5. These parameters and our assumptions about them are:

1. The reputation values R_{ij} are initialized based on the same process as explained in chapter 5 in section 5.1.2. Once users in the WSMS have reputation values having passed an initialization phase similar to the previous chapter, then our model is functional using the computational artefacts presented in this chapter.

2. The distances are logged for each observation using the GPS sensor of the smart phones, hence this is a known parameter

3. Trust values on the social network are assumed to exist in our use case in the fabric of social graph in the WSMS system. Each volunteer rates a person with a rating $s_{ij} \in \{.1 - 1\}$ and for computational reasons values lower than .1 are not allowed on the social network. Low values indicate low trust, since we do not deal with the notion of distrust as discussed in chapter 3.

4. The weight of the social tie links y_{ij} is part of the model, but is not initialized or used during our experiments resulting in the use of s_{ij} in the computation of equation 6.3.

In the next section we present the ABM developed to simulate the extended WSMS scenario discussed earlier and proceed with evaluation experiments of the model to verify specifically the role of the spatial dimension on the performance of the computational model.

6.3 OVERVIEW: EXTENDED WSMS SIMULATION MODEL

In this section we illustrate the developed ABM to evaluate the use case presented in this chapter by extending the ABM developed in the previous chapter. In our explanation we focus only on the extensions beyond the ABM developed in chapter 5. Figure 28 shows the interface of the extended ABM. The current model focuses on distance as opposed to the model presented in chapter 5 focusing on time. This model also has a social network at the heart of the model's design. Hence, the two important aspects with which the ABM was extended are the social network and space. Essentially our design implemented two simulations embedded into a single structure. The first is for initialization (referred to as ABM 'initializer') and generates the social network, its trust ratings and the distances of the agents to test the effect of distance on model precision. The second runs our WSMS use case and implements the model presented in this chapter to compute drinkability of water wells based on the social network and distance structure generated from the first simulation. Of importance is the generated social network which we illustrate in this section.

A social network is essentially represented as a graph of nodes and vertices. A question when designing the social network in our simulation is what are the properties of this social network? Perhaps the most widely known property of real world social networks are the so called *small-world effect* which denotes the fact that the average path length in a network (through how many people do you have to connect to reach another person) is small relative to the system's size. This has been widely popularized as the six degrees of separation (for a good explanation see [98]). The first theoretical model to produce this property is credited to Paul Erdos and Alfred Renyi [31]. In this model a network is said to have the small-world property if the average path length scales logarithmically with the network size for a fixed mean

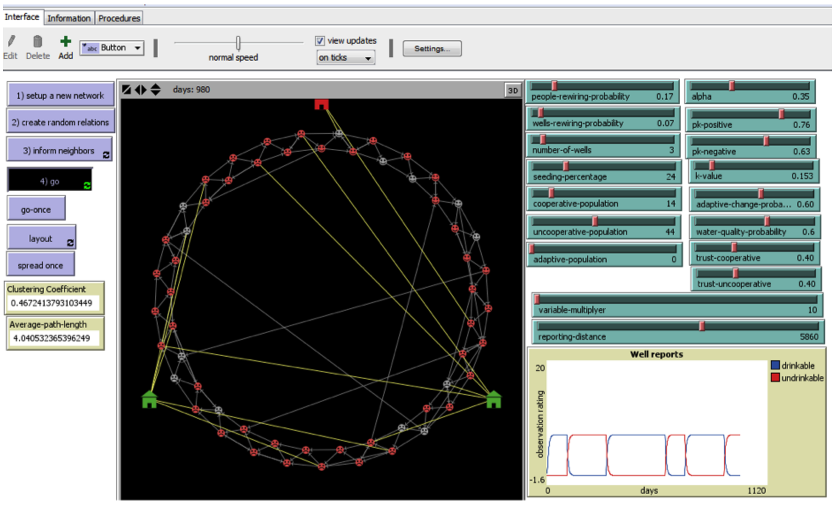

Clustering Coefficient 0.4672413793103449
Average-path-length 4.040532365396249
people-rewiring-probability 0.17
wells-rewiring-probability 0.07
number-of-wells 3
seeding-percentage 24
cooperative-population 14
uncooperative-population 44
adaptive-population 0
alpha 0.35
pk-positive 0.76
pk-negative 0.63
k-value 0.153
adaptive-change-proba... 0.60
water-quality-probability 0.6
trust-cooperative 0.40
trust-uncooperative 0.40
variable-multiplyer 10
reporting-distance 5860

Figure 28: *The extended WSMS simulation interface.*

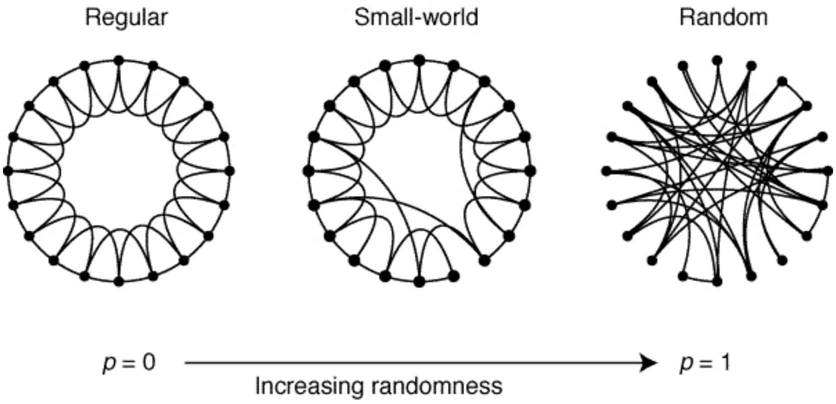

Figure 29: *The Watts-Strogatz Model [99]. The rewiring probability* p *increases randomness by reconnecting an edge with a vertex chosen uniformaly at random over the entire ring, with duplicate edges forbidden.*

network degree. The problem with this model however is that it results in very small clustering coefficients, and in the late 90's Watts [99] observed that an important property of complex networks including social networks is a high level of clustering. Clustering here means that nodes form groups that are densely connected internally but have fewer links to outer groups. The Watts-Strogatz model in figure 29 shows the effect of the rewiring probability p in a uniformly connected lattice resulting in a small world effect.

Our candidate for a social network in this simulation has to have a low average path length and a high clustering coefficient. To produce this effect we implemented the Watts-Strogatz model [99]. Figure 30 shows our ABM interface with the implementation of the Watts-Strogatz model to produce a social network with small world property and a high clustering coefficient for our simulations of the computational model. The original Watts-Strogatz model produced an undirected graph while our implementation produces a directed graph, because our graph is a trust network and trust networks are directed due to the asymmetry of trust we discussed in earlier chapters.

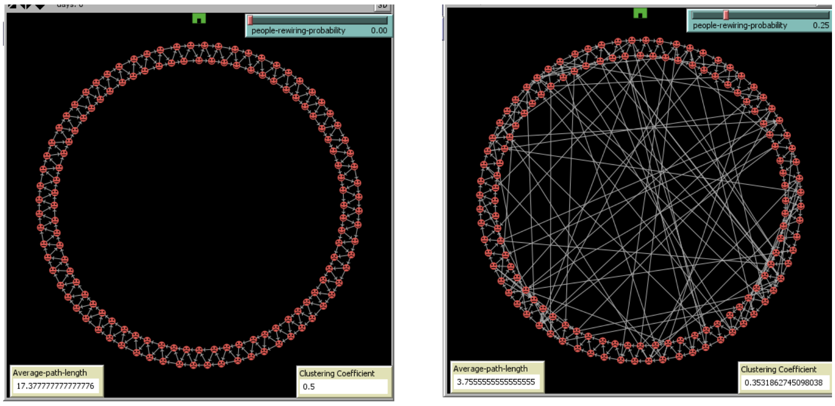

people-rewiring-probability 0.00
Average-path-length 17.377777777777776
Clustering Coefficient 0.5

people-rewiring-probability 0.25
Average-path-length 3.7555555555555555
Clustering Coefficient 0.3531862745098038

Figure 30: *The ABM interface showing our implementation of the Watts-Strogatz model to produce the social network for our simulation. On the left the network with rewiring probability* $p = 0$ *and on the right rewiring probability* $p > 0$.

Furthermore we use an algorithm to generate trust ratings on the social network. The sliders in the simulation interface labelled 'trust-cooperative' and 'trust-uncooperative' allow for the control of the trust value which is assigned according to the relationship between the volunteers. To calculate this value, the 'seeding-percentage' slider is used during initialization of the ABM to establish the percentage of the population that has *prior knowledge* about the water wells status. This knowledge is used during initialization to determine the trust between individuals as information is passed from one volunteer to the other. For example, in a directed graph, the outgoing link from a node that passes accurate information to its neighbour about the water quality of a well takes a trust value equal to set by the 'trust-cooperative' slider and the same happens for uncooperative volunteers who receive the value of the 'trust-uncooperative' slider. The adaptive volunteers are initialized as cooperative or uncooperative and receive the same treatment corresponding to their initialized type. Finally, the cooperative volunteers are given a value of distance c as 10 distance units while the uncooperative and adaptive volunteers are controlled through the slider in the ABM titled 'reporting-distance' to make the uncooperative and adaptive volunteers report from distance of 10 to 5000 distance units. Our ABM experiments in the following sections aim at detecting the effect of distance on the precision of the computational model. It is important to notice that all the rules explained so far apply to the ABM initializer only. Once the initialization is done our simulation runs on the generated network structure using similar steps as explained in chapter 5 in figure 13.

In the next section we directly study the effect of space on model precision by conducting experiments and discussing the results.

6.4 EFFECTS OF SPACE ON THE MODEL PERFORMANCE

Much like our analysis proved in chapter 5 the model successfully captures the status of the changing water wells in the presence of different types of volunteers in the system. Our aim at this stage is to understand the effect of adding space sensitivity to our model without analysing other behaviours which we have tested in the previous chapter. Thus all model parameters, including the time decay of trust k which we studied in the previous chapter, are constant except the four parameters addressed later in this section namely the three population types and the reporting-distance.

Initial sampling of the model revealed that, for example, if we have 10 honest volunteers (by default reporting from 10 distance units) and 20 dishonest and adaptive volunteers, with the dishonest and adaptive volunteers reporting also from 10 distance units, the model fails to predict the true status of the water wells. While if the 'reporting-distance' slider is moved to 400 distance units the model correctly predicts the true status of the water wells. It is also to be noticed that during sampling of our model and when distance is applied, the latency measure in equation 5.9 has its values always kept at a stable minimum, hence there is not need to further investigate this measure. As such for this model our evaluations are based on the two performance measures of hit rate and adaptation in equations 5.7 and 5.8. The aim of our experiments is then to reveal the interactions between different volumes of volunteer types and the combined effect of the reporting distance on the hit rate and adaptation of the model.

To achieve this we design and implement a manual experiment in the MEME toolset. The manual experimental design builds a parameter tree for the parameters to be tested and their values and creates all possible combinations of these parameters for testing the interactions between the selected parameters. Initially our experimental design selects all three types of populations as three parameters (cooperative, uncooperative and adaptive volunteers) in addition to the ABM slider 'reporting-distance' which controls the distance from which the dishonest and adaptive volunteers are reporting. Thus, our experiment studies interactions of four parameters. The initial experiment was designed to sample populations between their limits of $0 - 100$ for each volunteer type at a sampling interval of 10. The distance is sampled between its limits of $10 - 10000$ distance units with a sampling interval of 100. And we have earlier mentioned that we are only interested in the two model performance measures of hit rate and adaptation. This design resulted, with random seed replications, in 3240 runs at 6000 ticks each (same tick count in the previous chapter). Due to the computationally intensive nature of this ABM with the social network and the bipartite graph, this design required extensive resources. It was estimated that on multiple computing cores (approximately 8), this experiment would run for over a week.

Upon further sampling of the ABM we realized that, unlike the previous ABM from chapter 5 it reached a stable state after about 1000 ticks. We also realized that there is a significant performance improvement if populations are kept to a maximum of 50 rather than a 100. Also, given the limited population size, we noticed that distances bigger than 2000 units did not have much impact on model performance. As a pragmatic decision we also chose to sample the distance between $10 - 5000$ with a 1000 as interval. The result of this reduced design was 1080 simulation runs, with random seed replications, at a 1000 ticks each, that was completed in about approximately 1 day.

The results of our experiments are summarized in the interaction statistics diagrams in figure 31 for combined effects on hit rate and in figure 32 for combined effects on adaptation.

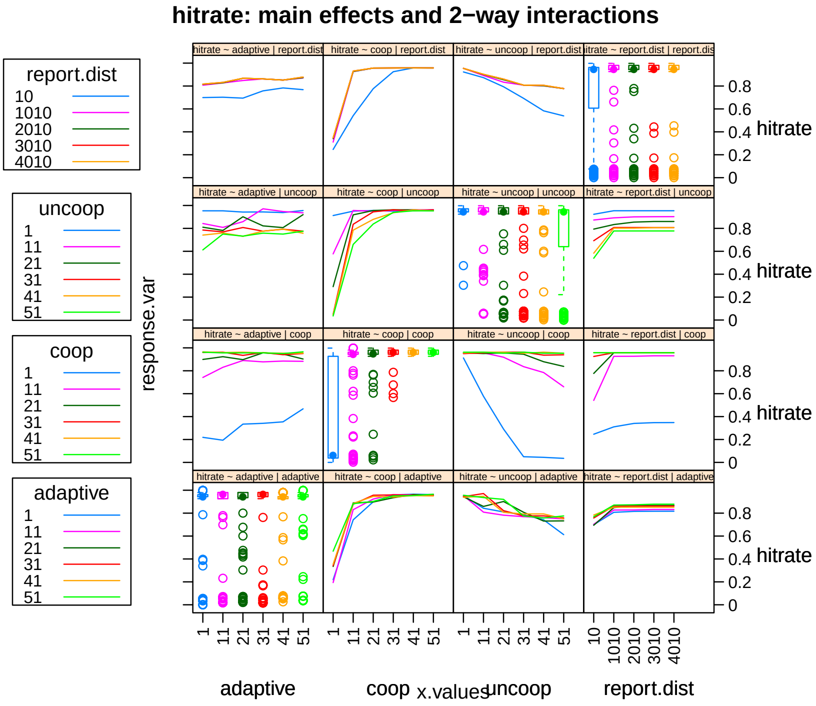

Figure 31: *The 2-way interactions of the different volunteer populations and reporting distance and the combined effects on the model's hit rate performance measure.*

Looking at the main effects on hit rate (diagonal charts from top left most chart to bottom right most in figure 31), we observe that regardless of what reporting distance is used, the number of adaptive volunteers has very limited effect on the hit rate. The only effect of adaptive volunteers is that for higher reporting distances the overall hit rate increases but is stable for all variations of adaptive populations. For the cooperative volunteers, several things are observed. Once we have at least 11 cooperative volunteers the effect on the hit rate stabilizes for our model at all reporting distances. Essentially, reporting distance has a higher importance in improving the model's performance when the number of cooperative volunteers is lower. In the third chart from the left at the top we observe that the increasing number of uncooperative volunteers reduces the overall hit rate of the model, yet most importantly this effect is offset by the increase in reporting distance. In other words reporting distance has a direct positive impact on the model's hit rate when populations of malicious behaviour increase in the ABM.

In addition, more interesting findings are revealed when examining the interaction charts closely. Looking at the charts left to right, from top to bottom, the number of adaptive agents and the reporting distance does not really have any interactions, except when the reporting distance is as low as 10. In other words when the adaptively dishonest agents report from the same distance as the honest agents by going to the water well's location they have the ability to disrupt the model, yet not to a large degree. Generally, this scenario is not likely to happen because it requires the physical effort of all malicious volunteers to move to the water wells location and cause a flood of reporting from this location for an extended period of time. Similarly, there is no interaction effect between the number of adaptive and cooperative or uncooperative volunteers (as indicated by the horizontal lines in the charts), but we observe that a lower number of uncooperative, or a higher number of cooperative volunteers directly improves the model's performance. Combining this with the improving effect of the reporting distance the model is solid in the face of high number of uncooperative volunteers given increasing reporting distances.

Finally, in the second column of charts, we clearly observe that the higher the number of cooperative agents, the higher the hit rate will be on all three interaction charts. The top one shows that if the reporting distance is 10, than hit rate grows slower with the number of cooperative volunteers, than if the reporting distance is higher. This means that higher reporting distances improve the prediction in favour of the cooperative/honest volunteers. As a general case, if there are more than 40 cooperative volunteers in the system and a smaller number of uncooperative and adaptive volunteers, the reporting distance improves the hit rate but is not detrimental to the success of our model as long as it is more than 10 (i.e. uncooperative volunteers are further away from the water well than cooperatives, even by a small margin). Generally, from the second chart of the second column we observe that the number of uncooperative volunteers only matters if there are less than 30 cooperative volunteers. In this case the lower the number of uncooperative volunteers, the higher the hit rate. In the third column, the top chart shows that the number of uncooperative volunteers has, in general, a negative effect on the hit rate, and the level of reporting distance, unless it is 10, has a weak effect when the hit rate is less than otherwise. The chart also confirms our earlier assertions that the reporting distance really matters more strongly with higher levels of uncooperative volunteers where in all cases the reporting distance improves the model's hit rate.

The same analysis conducted for the adaptation performance measure reveals similar results as shown in figure 32. The analysis yields similar conclusions to the hit rate analysis discussed earlier. One result of note is the third row of charts in the figure. For very low numbers of cooperative volunteers and in the existence of higher number of uncooperative and adaptive volunteers, the reporting distance does not improve the model's adaptation rate, while the effect is very strong for cooperative populations above 11, in which case the reporting distance has a consistent effect resulting in higher adaptation rate.

6.5 CONCLUSIONS

In this chapter we extended the WSMS use case with a scenario that depends on our computational trust and reputation model knowing the distance between the location of the volunteer at the time of reporting and that of the water well being reported on. We extended the ABM from the previous chapter to accommodate the requirements of the extended WSMS use case and computational model presented in this chapter. Our ABM depends on ratings in a social network as well as distance to compute informational trust ratings for water wells to predict the status of the water wells as they change throughout the simulation time.

We have conducted analysis of the effect of the reporting distance on the model's performance measures with respect to different volunteer

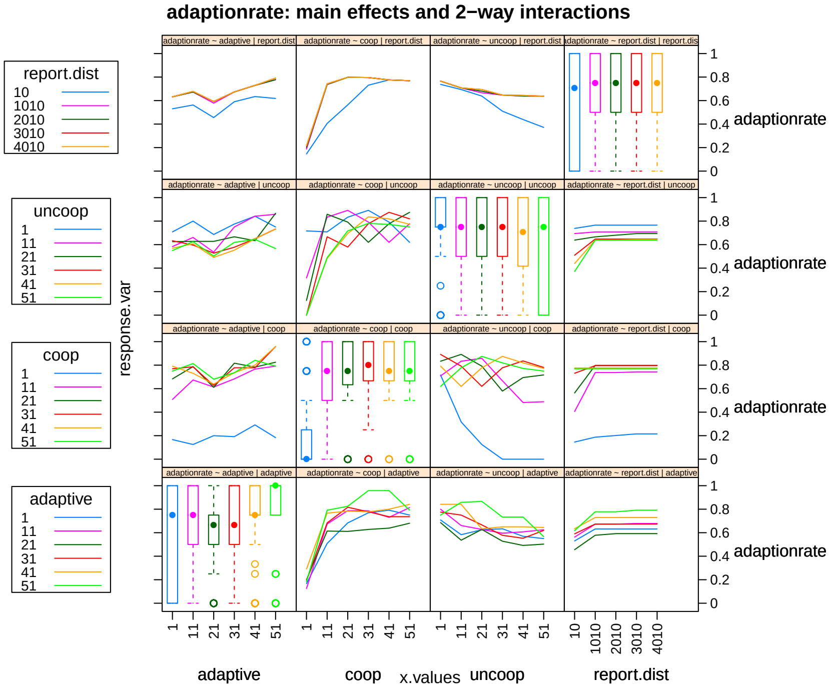

Figure 32: *The 2-way interactions of the different volunteer populations and reporting distance and the combined effects on the model's adaptation performance measure.*

populations. We have shown that the reporting distance parameter does impact the performance of the model positively with lower distances and we have shown through our analysis the different conditions where this effect is strongest.

7

CONCLUSIONS & FUTURE WORK

In this chapter, we summarize the research results of this thesis and discuss future research directions. In section 7.1, a summary of our research and a conclusion on our research hypothesis is presented followed by a closer look at our research questions and the obtained research results. In section 7.2 we discuss future research directions stemming from the work done in this thesis.

7.1 SUMMARY & DISCUSSION

In this thesis we address the problem of the quality of user generated content. Particularly we address the problem of the quality of VGI and human sensor observations characterized by proneness to errors or fraud and the lack of traditional Geospatial Information (GI) quality criteria [10]. We posed this hypothesis for our research:

Spatially and Temporally Sensitive Trust Models are Suitable Proxy Measures for VGI Quality

Our research build on a trifurcate approach to quality assessment of human sensor observations:

- *Integration of trust and reputation as quality measures in the sensor web* particularly for human sensors. We propose an ontological account that extends the functional ontology for observation and measurement (FOOM) presented in [60, 79]. This work integrates trust and reputation models for humans as sensors into sensor web ontologies. The ontology is grounded in the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE). We show how our approach formalizes the notions of trust and reputation as ontological qualities and the computational models of trust and reputation as sensors. We make clear distinctions between trust as a social relation between people and trusting information through the introduction of our notion of *informational trust* mediated through interpersonal trust.

- *Computational models of trust and reputation* are presented in this thesis. Both models presented were evaluated using agent based simulation models. Our analysis showed the resilience of our computational models to malicious behaviour and their ability to consistently predict accurate observation results for the observed phenomenon. This is the proof of the first part of our hypothesis, namely that trust and reputation models are a suitable proxy measure of observation quality.

- *Spatial and temporal sensitivity* of computational trust and reputation models have been shown to improve model performance measures, namely hit rate, adaptation and latency which have been defined to reflect the effectiveness of our computational models in reflecting accurate observations. We have isolated the effects of time and then of space. In each case we have shown

> how the temporal and spatial sensitivity of our computational models positively impacted model performance measures. These positive effects of space and time on model performance measures constitute the proof of the second part of our hypothesis, namely that the temporal dimension of informational trust improved model performance measures as shown in chapter 5 and the same applied to the spatial dimension as shown in chapter 6.

The remainder of this section discusses the research questions individually and sheds light on how this research tackled each question.

What is trust? as the first research question included the three subquestions a) what is the definition of trust this thesis subscribes to?, what types of trust matter to our proposal of proxy measures of observation quality? and is there evidence for space and time effects on trust and how to employ them in our proxy measures proposal? In answering these questions we grounded our work in earlier research by studying trust from various disciplines and adopted the definition of Sztompka [92, p.27] where trust is defined as *a bet about the future contingent actions of others*. We also discussed the two components of this definition, belief and commitment. A belief that a certain person will act in a favourable way and a commitment by the believer to a certain action based on that belief. Furthermore, in chapter 3 we used the knowledge gained from chapter 2 to build an understanding of trust as a proxy measure of quality of human sensor observations. We introduced properties of trust like contextualization. However, our research did not attempt to model the context of trust explicitly, but we make an assumption in which the context of trust is implicitly understood and fixed in our use cases in later chapters. We have also discussed evidence of spatial and temporal aspects of trust. Concerning the spatial dimension, we focused on to what level does geographic proximity impact trust between people, and consequently impact each other's recommendations about certain observations. Our proposal is to explicitly integrate the spatial dimension in trust and reputation models for human sensor observations. A similar approach was followed for the temporal dimension where trust relationships develop and decay with with time. Such effect of time and space on trust influenced how we developed spatial-temporal trust and reputation models for VGI. Finally, to close our first research questions we introduced the notion of *informational trust*. We posited that *interpersonal trust* implies the transition of trust from the trustee to information entities conveyed by the trustee. The trustor then asserts trust in the information conveyed by the trustee. This type of trust we termed *informational trust*. A trusting tie between a trustor and an information entity such as VGI is mediated by interpersonal trust between the VGI originator and the VGI consumer. We then proceeded to extend the notion of *informational trust* with spatial and temporal dimensions asserting that trust decays and develops over time and that a volunteer's location with respect to the observed phenomena will impact how others trust her observations.

How can we integrate trust as a proxy measure of quality of human sensor observations into the sensor web? Sub-questions raised here were a) what is the ontological nature of both trust and reputation when used as a proxy measure of human sensor observation quality? b) how can volunteers of VGI be integrated as sensors in sensor web ontologies? and how can trust and reputation models be integrated in the sensor web ontologies? Initially, we presented our proposal of *trust as a proxy measure of observation quality* in chapter 3 . This proposal required:

1. the integration of trust and reputation as quality measures of human sensor observations into the sensor web observation and measurement standard.

2. the development of computational trust and reputation models for human sensor observation filtering and triage

In chapter 4 we addressed the first requirement through the developed Functional Ontology of Observation and Measurement (FOOM). The ontology builds on work done in [60, 79] and generalizes technical sensors and human sensors as observing agents. Our work integrates trust and reputation as quality assessment measures for observations of human sensors. In our ontology interpersonal trust, informational trust and reputation are presented as relational qualities while trust and reputation models are presented in our ontology as sensors. The approach integrates trust and reputation as proxy measures of quality, and computational trust and reputation models into sensor web ontologies. The ontology is implemented in the Haskell functional programming language and applied to the water well scenario use case of the H2.0 project. The Haskell functional programming language provided means for a rigorous algebraic specification as well as a simulation environment that enabled testing the functionality of the ontology.

How can we develop spatio-temporal computational models of trust and reputation for quality assessment of human sensor observations? This question included sub-questions a) are computational models of trust and reputation useful for quality assessment of human sensor observations? b) does accounting for the spatial and temporal nature of VGI and volunteers impact model performance? We have used a computational model approach to develop both a temporal trust and reputation model in chapter 5, and a spatial-temporal trust and reputation model in chapter 6. In chapter 5 we have presented and evaluated our temporally sensitive computational model in a Water Supply Monitoring System (WSMS) scenario. We developed an ABM simulation of the WSMS scenario and applied our computational model to study model behaviour and performance. To achieve this we developed performance measures, namely hit rate, adaptation and latency to measure different performance criteria for our computational model centred on its ability to predict the statuses of water wells as drinkabe or undrinkable based on volunteer reports. The model shows a remarkable degree of effectiveness in predicting water well statuses under challenging conditions of malicious or flawed reporting. Also, simulation experiments show that essential to the performance of our model is the temporal dimension of informational trust evaluated in our analysis as the effects of the k model parameter. Through our analysis we concluded that the sensitivity to the temporal dimension in our trust and reputation model directly and positively impacts model performance. Thus, we have proven the effectiveness of our first computational model as a tool for quality control of human sensor observations. In chapter 6 we presented a spatial-temporal computational trust and reputation model and studied the effects of the spatial dimension separately. The scenario included a social network and uses knowledge of the distance between the location of the volunteer at the time of reporting and that of the water well being reported on. Our analysis of the effect of the reporting distance on the model's performance measures developed in the previous chapter with respect to different volunteer populations proved our model's effectiveness. We have shown that the reporting distance parameter does impact the performance of the model positively with lower distances and we have shown through our analysis the conditions under which this is true.

In this discussion concerning our research questions from chapter 1 we have shown that our research answered all the research questions raised. Since our models showed remarkable ability to predict the status of water wells that corresponded to reality we also conclude that spatial and temporal informational trust and reputation models allow for a novel approach for quality control of human sensor observations and are suitable proxy measures for quality of human sensor observations on particular and of VGI in general.

7.2 BENEFITS & FUTURE WORK

In this section we discuss some open research directions stemming from the research done in this thesis. The core benefit of our research is in it being, as far as we know, the first attempt to methodically solve the problem of quality assessment of VGI and human sensor observations. Our ontological approach is the first attempt to integrate trust and reputation as novel measures for quality of human sensor observations. And the proposed notion of informational trust provides a conceptual framework to achieve this vision. Furthermore, the computational models developed in this thesis are novel in their approach to using trust and reputation for information triage in real time and also for their integration of the spatial and temporal dimensions of informational trust.

In this research we have focused on a specific application area as well as a specific use case. Namely we have focused on the sensor web as an application area and a water supply management system (WSMS) of the H2.0 project as our use case. The chosen application area is wide and our focus on the emerging field of humans as sensors leads to many potential benefits of our research to numerous use cases.

For example, *disaster response* use cases can considerably benefit from our research. Our models show the ability to filter a large flow of human sensor observations in real time. This has important implications for disaster response applications like the Haitian earth quake which we discussed in chapter 1. The ability to perform information triage on large scale real time information can provide immediate aid to disaster response teams trying to check the veracity of volunteer reports to plan and prioritise rescue and relief efforts. We believe this research will motivate a larger research agenda on quality assessment of human sensor observations in particular and VGI in general.

Some future research topics emerging from this research are discussed in the remainder of this section.

Interactions of the spatial and temporal dimensions of informational trust: in this research we developed two trust and reputation models. The first integrated the temporal dimension and the second integrated the spatial and temporal dimensions. Yet, we isolated the effect of temporal and spatial dimensions of informational trust and studied them separately. Further research is necessary to understand the interaction of both dimensions and their combined effects on the performance of our models.

VGI about non-stationary objects: the two models developed in this thesis assumed that the volunteers are mobile agents reporting on stationary phenomena, the water wells in this case. Reporting information about non-stationary objects by mobile volunteers constitutes a different type of adhoc networks that can occur in different situations, especially emergency situations where rescue teams are mobile and the volunteers reporting on them are mobile. Such a highly dynamic system requires further research into another class of models that captures its dynamics to use spatial and temporal informational trust for quality assessment. We postulate that such models could also be tested using agent based simulations which are effective tools for studying emergent behaviours and properties of such dynamic models.

Implementing the Haskell ontology in OWL: while the Haskell language provided the necessary expressiveness and functional testing capability that enabled developing the FOOM ontology, the Web Ontology Language (OWL) can be used to encode our developed trust and reputation ontology to enable a seamless integration with the emerging semantic sensor web standards.

Studying the social network effects: in chapter 6 we introduced a model capable of functioning on a social network and built a static small world social network in our simulations. Research on a dynamic social network where social ties are built and lost is essential to understand the dynamics of the social network. It is also prudent to study the dynamics of trust on the social network by enabling changing interpersonal trust ratings on the social network in the simulation. This will provide essential insights on the behaviour of interpersonal trusting agents in social networks with respect to VGI trust and reputation models.

Sustaining Volunteer Participation: the success of VGI highly depends on the participation from volunteers and this was proven true through our simulations. Some mass of volunteers is required before our models could reach optimum performance. Sustained cooperative behaviour in social networks is referred to as *cooperative regime* [21]. Including the evolving and emerging social structures in cooperative social networks has been proposed by several researchers to develop cooperative regimes. There is no research however that considers this approach to ensure a high level of participation from volunteers in VGI systems. This is a large area of potential future research to understand the dynamics of volunteer participation.

Human Sensors in VGI and Spatial Data Infrastructure (SDI) Research: the SDI has been conceived as a top-down research agenda. In contrast VGI and human sensing are terms coined for a phenomenon that emerged fully from the bottom up. Reconciling both fields and creating a unified research agenda is no easy task and we argue that this should not in fact be done in order to preserve the dynamic bottom up nature of human sensors and VGI as a research field. What we might hope to do is outline, in broad terms, the potential research areas without imposing the constituting details of each area. Figure 33 elucidates a proposal for this outline. VGI and human sensor research will have to consider, software, hardware and institutional and business models. The purpose of VGI is to generate geospatial information, which is the common perspective from which VGI is viewed. Another important perspective on VGI is to organize information around geography.

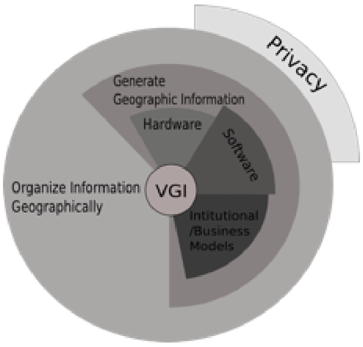

Figure 33: *Outline of potential research areas surrounding integration of human sensors and VGI into the SDI.*

From the software perspective, software must evolve around the VGI principles of community work and collaboration while embracing the current SDI standards. Some questions regarding software for VGI are a) can we develop systems that integrate state of the art in web research on credibility, trust and reputation models to assess the quality of VGI and reward credible and productive citizens? such a question has been addressed in our research and further work is needed to build APIs to integrate our ontological account and computational models within the SDI and the sensor web. Another question is b) how can traditional GIS integrate with VGI and emergent VGI web applications? Recent developments of ArcGIS (starting at version 9.2) publish all services and data on ArcGIS server as KML streams and HTML to the wider web. This enables mashing up traditional GIS resources with mapping APIs of Google, Yahoo or Google earth.

From a hardware perspective GPS devices, navigation devices and mobile phones, which now increasingly come bundled with GPS, cameras, motion sensors and potentially other forms of sensors, provide the technical hardware platform for VGI and human sensors. Some questions we raise here are a) how can we develop non-intrusive technologies that integrate seamlessly with people's lifestyles to leverage the VGI potential? b) how can different devices integrate to provide VGI functionality to users? An example here is new navigation devices with TomTom mapshare technology. Such devices allow users to report information (e.g. road block) on the device, which can then be shared online with the mapshare community. Users can then download the updates on their devices, and there are controls to specify which types of updates are allowed and what degree of verification the updates go through. Integration can allow users' phones to send pictures with location tags to the mapshare device, which can then be shared with the community, giving navigation data users real life images of the locations they visit.

In addition, the importance of institutional frameworks and business models becomes apparent when one considers the recent developments mentioned earlier, which promise to bridge the gap between VGI applications and traditional GIS. This comes with the traditional problems of intellectual property rights and information sharing which have hampered the GIS world for many years. When information from legacy GIS systems is made available on the web for mainstream use in VGI applications, issues of copyright and data ownership become crucial.

Finally, the collection of VGI through human sensors and the organization of information geographically are directly impacted by our research. Two perspectives exist when looking at VGI and human sensors. VGI is not just about collaborative creation and maintenance of GI, but it is as much about organizing different types of information around geography. In many of the examples like Wikimapia or Openstreetmaps one can see geographic information in the making, but many other examples like Flickr or mapping of Wikipedia illustrate organization of information geographically. Volunteers are not just sensors, but are also indirectly organizing their information geographically. Much of our information has geographic components, what Web 2.0 and VGI tell us is that there has always been an untapped need to search for information based on its geographic properties. Once tools that made this possible were available, we were able to observe a plethora of applications satisfying this need. Examples of questions that arise here are a) how can we place VGI at the center of information archival and retrieval systems? b) what are the implications of VGI on information discovery? Naturally, quality of human sensor observations is essential to be able to ensure effective organization of information geographically. This aspect is one where the main body of our research contribution serves its role.

Privacy Research: finally, is the issue of privacy. In work by Microsoft Research[1] about 5% of the participants in a study involving collecting GPS tracks where successfully identified and their true identities revealed by analysing the GPS tracks. The percentage, although not very high, is alarming because a person compromised is one person too many, and the percentage can also be improved if more sophisticated techniques are used. Even the simplest applications of human sensing in VGI like tagging photos, can reveal a lot of private information about the individuals. Research on the implications of VGI on privacy has to deal with collecting and maintaining VGI while preserving privacy. The computational models developed in this thesis require knowledge about volunteers such as unique identifiers. The models also maintain histories of volunteer interactions with the WSMS. The implications of collecting and processing this information on privacy are unclear and could prove very challenging to VGI research. Some immediate research questions on privacy are a) what are the implications of VGI on user privacy, both technically and legally? and b) how can we develop technical means to preserve the anonymity of VGI human sensors, while preserving the quality and integrity of observations?

---

[1] http://where.blip.tv/#570532

BIBLIOGRAPHY

[1] R. Axelrod. Advancing the art of simulation in the social sciences. *Complexity*, 3(2):16–22, 1997. ISSN 1076-2787. (Cited on pages 48 and 49.)

[2] R Axelrod and WD Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981. (Cited on page 51.)

[3] A. L. Barabási and R. E. Crandall. Linked: The new science of networks, 2003. (Cited on page 20.)

[4] CL Barry. User-defined relevance criteria: An exploratory study. *Journal of the American Society for Information Science*, (45):149–159, 1994. (Cited on page 9.)

[5] J Bateman. Modeling the importance of end-user relevance criteria. In *Proceedings of the 62nd Annual Meeting of the American Society for Information Science*, volume 36, pages 396–406, 1999. Proceedings of the ASIS Annual Meeting. (Cited on page 9.)

[6] D. E. Bell, H. Raiffa, and A. Tversky. *Decision making: Descriptive, normative, and prescriptive interactions*. Cambridge University Press, New York, 1988. (Cited on page 12.)

[7] Y. Benkler. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, 2006. (Cited on page 10.)

[8] M Bishr and L Mantelas. A trust and reputation model for filtering and classifying knowledge about urban growth. *GeoJournal*, 72(3):229–237, 2008. (Cited on page 16.)

[9] Mohamed Bishr and Krazystof Janowicz. Can we trust information? - the case of volunteered geographic information. In Anusuriya Devaraju, Aljandro Llaves, Patrick Maue, and Carsten Kessler, editors, *Towards Digital Earth: Search, Discover and Share Geospatial Data*. Workshop at Future Internet Symposium, Berlin, Germany, September 20th 2010. (Cited on page 20.)

[10] Mohamed Bishr and Werner Kuhn. Geospatial information bottom-up:a matter of trust and semantics. In *AGILE*, Aalborg, Denmark, 2007. Springer-Verlag. (Cited on pages 1, 4, 9, 16, 19, 20, 26, and 87.)

[11] Werner Bönte. Inter-firm trust in buyer–supplier relations: Are knowledge spillovers and geographical proximity relevant? *Journal of Economic Behaviour and Organization*, 67(3-4):855–870, 2008. (Cited on page 15.)

[12] M. Botts. Opengis sensor model language (sensorml) implementation specification. Technical report, Open Geospatial Consortium, 2007. (Cited on page 41.)

[13] M. Botts and G. Percivall. Ogc sensor web enablement: Overview and high level architecture. White paper, Open Geospatial Consortium, December 2007. (Cited on pages 19 and 30.)

[14] G.E.P. Box, W.G. Hunter, and J.S. Hunter. Statistics for experimenters: an introduction to design, data analysis, and model building. 1978. (Cited on pages 57 and 58.)

[15] A. Bröring, K. Janowicz, C. Stasch, and W. Kuhn. Semantic challenges for sensor plug and play. *Web and Wireless Geographical Information Systems*, pages 72–86, 2009. (Cited on page 33.)

[16] P Buneman, S Khanna, and WC Tan. Why and where: A characterization of data provenance. *Lecture Notes in Mathimatics - Springer Verlag*, pages 316–330, 2000. (Cited on page 18.)

[17] V. W. Buskens. *Social networks and trust*. Springer, 2002. (Cited on pages xiii, 4, 9, 12, 13, 14, and 24.)

[18] Haiguang Chen, Huafeng Wu, Xi Zhou, and Chuanshan Gao. Agent-based trust model in wireless sensor networks. *Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, ACIS International Conference on*, 3:119–124, 2007. doi: http://doi.ieeecomputersociety.org/10.1109/SNPD.2007.122. (Cited on page 54.)

[19] S Cheskin. ecommerece trust: Building trust in digital environments. *Archetype/Sapient*, 1999. (Cited on page 24.)

[20] T Chesney. An empirical examination of wikipedia's credibility. *First Monday*, 11(11), 2006. (Cited on page 9.)

[21] M.D. Cohen, R.L. Riolo, and R. Axelrod. The role of social structure in the maintenance of cooperative regimes. *Rationality and Society*, 13(1):5, 2001. ISSN 1043-4631. (Cited on page 91.)

[22] CL Corritore, B Kracher, and S Wiedenbeck. On-line trust: concepts, evolving themes, a model. *International Journal of Human-Computer Studies*, 58(6):737–758, 2003. (Cited on page 15.)

[23] Simon Cox. Observation and measurement. implementation specification 07-022rl, Open Geospatial Consortium, 2007. (Cited on pages 19, 34, and 41.)

[24] R Dahrendorf. *Reflections on the Revolution in Europe*. Transaction Publishers, 2005. (Cited on page 10.)

[25] Kevin Delin. Sensor webs. (Cited on page 19.)

[26] M Deutsch. Cooperation and trust: Some theoretical notes. volume 10, pages 275–319. Lincoln: University of Nebraska Press, 1962. Nebraska Symposium on Motivation. (Cited on pages 5, 9, and 11.)

[27] E Dumbill. Xml watch: Finding friends with xml and rdf. *IBM Developer Works, http://www-106. ibm. com/developerworks/xml/library/xfoaf. html, June*, 2002. (Cited on page 16.)

[28] R Dunbar and L Barrett. *Oxford Handbook of Evolutionary Psychology*. Oxford University Press, 2007. (Cited on page 14.)

[29] J.H. Dyer and H. Singh. The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *The Academy of Management Review*, 23(4):660–679, 1998. (Cited on page 15.)

[30] J Echterhoff and T Everding. Opengis sensor event service interface specification. Discussion paper, Open Geospatial Consortium, October 2008. (Cited on page 19.)

[31] P. Erdos and A. Renyi. On random graphs I. *Publ. Math. Debrecen*, 6(290-297):156, 1959. (Cited on page 80.)

[32] BJ Fogg and H Tseng. The elements of computer credibility. pages 80–87. ACM Press New York, NY, USA, 1999. Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit. (Cited on pages 9, 10, and 15.)

[33] D Fudenberg and J Tirole. *Game Theory*. Mit Press, 1991. (Cited on page 13.)

[34] Francis Fukuyama. Trust: The social virtues and the creation of prosperity, 1996. (Cited on page 9.)

[35] J Futrelle and J Myers. Tracking provenance semantics in heterogeneous execution systems. *Concurrency and Computation*, 20(5): 555, 2008. (Cited on page 18.)

[36] Diego Gambetta. *Trust: Making and breaking cooperative relations*, chapter Can We Trust Trust? Basil Blackwell New York, 1990. (Cited on pages 9 and 12.)

[37] P. Gärdenfors. *Conceptual spaces: The geometry of thought*. The MIT Press, 2004. (Cited on page 40.)

[38] N. Gilbert, J. S. Sichman, and R. Conte. Multi-agent systems and agent-based simulation. In *Multi-agent based Simulation (MABS)*, La Villette, Paris, France, 1998. Springer-Verlag. (Cited on page 49.)

[39] Jennifer Golbeck. *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, 2005. (Cited on pages xiii, 9, 13, 16, 17, 20, 24, 26, and 76.)

[40] Jennifer Golbeck. The dynamics of web-based social networks: Membership, relationships, and change. *First Monday*, 12(11), 2007. (Cited on page 16.)

[41] Jennifer Golbeck. Trust on the world wide web: A survey. Technical report, 2008. (Cited on pages 15, 16, 17, 18, and 20.)

[42] Jennifer Golbeck. Weaving a web of trust. *Science*, pages 1640–1641, 2008. (Cited on page 9.)

[43] Jennifer Golbeck and J Hendler. A semantic web approach to tracking provenance in scientific workflows. *Concurrency and Computation: Practice and Experience*, 2007. (Cited on page 18.)

[44] Jennifer Golbeck, B Parsia, and J Hendler. Trust networks on the semantic web. *Lecture Notes in Computer Science*, pages 238–249, 2003. (Cited on pages 16 and 17.)

[45] RT Golembiewski and M McConkie. The centrality of interpersonal trust in group processes. *Theories of group processes*, pages 131–185, 1975. (Cited on pages 5, 9, and 11.)

[46] M.F. Goodchild. Citizens as voluntary sensors: spatial data infrastructure in the world of web 2.0. *International Journal*, 2. (Cited on pages 1, 19, and 29.)

[47] E Gray, JM Seigneur, Y Chen, and C Jensen. Trust propagation in small worlds. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 239–254, 2003. (Cited on page 17.)

[48] N. Guarino. Formal ontology in information systems: Proceedings of the 1st international conference june 6-8, 1998, trento, italy. 1998. (Cited on pages xiii and 32.)

[49] R Guha, R Kumar, P Raghavan, and A Tomkins. Propagation of trust and distrust. pages 403–412. ACM Press New York, NY, USA, 2004. Proceedings of the 13th international conference on World Wide Web. (Cited on page 17.)

[50] D T Hawkins. What is credible information? *Online(Weston, CT)*, 23(5):86–89, 1999. (Cited on page 15.)

[51] J. Huang and M.S. Fox. An ontology of trust: formal semantics and transitivity. In *Proceedings of the 8th international conference on Electronic commerce: The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet*, pages 259–270. ACM New York, NY, USA, 2006. (Cited on page 17.)

[52] RJ Johnston. *The Dictionary of Human Geography*. Blackwell Publishers, 2000. (Cited on page 3.)

[53] Eike H. Jürrens, Arne Bröring, and Simon Jirka. A human sensor web for water availability monitoring. In *2nd International Workshop on Blending Physical and Digital Spaces on the Internet*, September 2009. (Cited on page 29.)

[54] Eike Hinderk Jürrens. Swe-connection to a gsm-interface to establish water availability monitoring in africa. Master's thesis, Institute for Geoinformatics, University of Muenster, Muenster, Germany, 2009. (Cited on pages xiii, 29, and 30.)

[55] Victor Keegan. Meet the wikipedia of the mapping world, November 2010. URL `http://www.guardian.co.uk/technology/2010/feb/04/mapping-open-source-victor-keegan`. (Cited on page 19.)

[56] J Kim, E Deelman, Y Gil, G Mehta, and V Ratnakar. Provenance trails in the wings/pegasus system. *CONCURRENCY AND COMPUTATION*, 20(5):587, 2008. (Cited on page 18.)

[57] J.P.C. Kleijnen and Center for Economic Research at Amsterdam. *Validation of simulation, with and without real data*. Center for Economic Research, Tilburg University, 1998. (Cited on page 51.)

[58] David Knowles. Twitter message sparks big rescue effort. URL `http://www.aolnews.com/tech/article/twitter-message-sparks-big-rescue-effort-in-haiti/19337648`. (Cited on page 2.)

[59] P. Krugman. *Development, geography, and economic theory*. MIT press, 2002. (Cited on page 7.)

[60] Werner Kuhn. A functional ontology of observation and measurement. In Krazystof Janowicz, Martin Raubal, and S. Levashkin, editors, *GeoSpatial Semantics Â· Third International Conference (GeoS)*, volume 5892 of *Lecture Notes in Computer Science*, pages 26–43. Springer-Verlag, 2009. (Cited on pages 5, 7, 19, 29, 33, 34, 42, 87, and 89.)

[61] C. Lane and R. Bachmann. The social constitution of trust: supplier relations in Britain and Germany. *Global Supply Chain Management*, 17(3):161, 2006. (Cited on page 15.)

[62] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *Proceedings of the National Academy of Sciences*, 102(33):11623, 2005. (Cited on pages 14 and 24.)

[63] Z Liu. Perceptions of credibility of scholarly information on the web. *Information Processing and Management*, 40(6):1027–1038, 2004. (Cited on pages 9 and 15.)

[64] EH Lorenz. *Trust: Making and Breaking Cooperative Relations*, chapter Neither Friends nor Strangers: Informal Networks of Subcontracting in French Industry. Basil Blackwell, New York, 1988. (Cited on page 14.)

[65] N Luhmann. *Trust and Power*. JOHN WILEY AND SONS, 1979. (Cited on pages 9 and 12.)

[66] BR Lyons. Contracts and specific investment: An empirical test of transaction cost theory. *Journal of Economics and Management Strategy*, 3(2):257–278, 1994. (Cited on page 14.)

[67] S. Marsh and M. R. Dibben. The role of trust in information science and technology. *Annual Review of Information Science and Technology*, 37(1):465–498, 2003. (Cited on page 11.)

[68] Stephen Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, 1994. (Cited on pages 4, 12, 20, 21, and 24.)

[69] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. Wonderweb deliverable d18 ontology library (final). Technical report, ICT Project, 33052, 2003. (Cited on pages xiii, 29, 32, and 40.)

[70] R.C. Mayer, J.H. Davis, and F.D. Schoorman. An integrative model of organizational trust. *Academy of management review*, 20 (3):709–734, 1995. ISSN 0363-7425. (Cited on page 25.)

[71] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather:homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001. (Cited on page 3.)

[72] MJ Metzger. Making sense of credibility on the web: Models for evaluating online information and recommendations for future research. *JOURNAL-AMERICAN SOCIETY FOR INFORMATION SCIENCE AND TECHNOLOGY*, 58(13):2078, 2007. (Cited on pages 9 and 15.)

[73] MJ Metzger, AJ Flanagin, K Eyal, DR Lemus, and RM McCann. Credibility for the 21st century: Integrating perspectives on source, message, and media credibility in the contemporary media environment. *COMMUNICATION YEARBOOK*, 27:293–336, 2003. (Cited on page 9.)

[74] N Mezzetti. A socially inspired reputation model. *Lecture Notes in Computer Science*, pages 191–204, 2004. (Cited on pages 20 and 25.)

[75] Argyle Michael. *Cooperation, the basis of sociability*. Routledge, London, 1991. (Cited on page 9.)

[76] S.K. Milton and B. SMITH. Top-level ontology: The problem with naturalism. In *Formal Ontology In Information Systems: Proceedings of the Third International Conference (FOIS-2004)*, page 85. IOS Press, 2004. ISBN 1586034685. (Cited on page 31.)

[77] H. Neuhaus and M. Compton. The Semantic Sensor Network Ontology. In *AGILE Workshop on Challenges in Geospatial Data Harmonisation, Hannover, Germany*, 2009. (Cited on page 33.)

[78] N. Nohria and R. G. Eccles. *Networks and organizations: structure, form, and action*. Harvard Business School Press, 1992. (Cited on page 4.)

[79] Jens Ortmann and Werner Kuhn. Affordances as qualities. volume 6th International Conference on Formal Ontology in Information Systems (FOIS), Toronto, Canada, 2010. (Cited on pages 5, 7, 19, 29, 34, 42, 87, and 89.)

[80] Florian Probst. *Semantic Reference Systems for Observations and Measurements*. PhD thesis, University of Muenster, 2007. (Cited on pages xiii, 32, 33, 39, and 40.)

[81] RazorFish. Razorfish consumer survey, 2008. (Cited on page 1.)

[82] S. Y. Rieh and D. Danielson. Credibility: A multidisciplinary framework. *Annual Review of Information Science and Technology*, 41:307–364, 2007. (Cited on pages 9, 10, and 15.)

[83] SY Rieh. Judgment of information quality and cognitive authority in the web. *Journal of the American Society for Information Science and Technology*, 53(2):145–161, 2002. (Cited on page 9.)

[84] SY Rieh and NJ Belkin. Understanding judgment of information quality and cognitive authority in the www. *Journal of the American Society for Information Science*, 35:279–289, 1998. (Cited on page 9.)

[85] Klaus Schwab and Terry Malleret. The trust deficit. *Wall Street Journal Europe*, 28(A10), January 2003. (Cited on pages 3 and 15.)

[86] A. B. Seligman. *The Problem of Trust*. Princeton University Press, 1997. (Cited on page 23.)

[87] S. P. Shapiro. The social control of impersonal trust. *The American Journal of Sociology*, 93(3):623, 1987. (Cited on pages 5 and 9.)

[88] YL Simmhan, B Plale, and D Gannon. A survey of data provenance in e-science. *SIGMOD RECORD*, 34(3):31, 2005. (Cited on page 18.)

[89] I Simonis and J Echterhoff. Opengis web notification service implementation. Best practice paper, Open Geospatial Consortium, November 2006. (Cited on page 19.)

[90] MD Slater and D Rouner. How message evaluation and source attributes may influence credibility assessment and belief change. *Journalism and Mass Communication Quarterly*, 73(4):974–91, 1996. (Cited on page 9.)

[91] C. Stasch, K. Janowicz, A. Bröring, I. Reis, and W. Kuhn. A stimulus-centric algebraic approach to sensors and observations. *GeoSensor Networks*, pages 169–179, 2009. (Cited on pages 33 and 41.)

[92] P. Sztompka. *Trust: A Sociological Theory*. Cambridge University Press, 1999. (Cited on pages 9, 10, 11, 20, 23, 25, and 88.)

[93] M. Taherian, R. Jalili, and M. Amini. PTO: A Trust Ontology for Pervasive Environments. In *Advanced Information Networking and Applications-Workshops, 2008. AINAW 2008. 22nd International Conference on*, pages 301–306. IEEE, 2008. (Cited on page 17.)

[94] S. Toivonen and G. Denker. The impact of context on the trustworthiness of communication: An ontological approach. In *Proceedings of the Trust, Security, and Reputation on the Semantic Web workshop, held in conjunction with the 3rd International Semantic Web Conference (ISWC 2004), Hiroshima, Japan*, volume 127. Citeseer, 2004. (Cited on page 17.)

[95] L. Viljanen. Towards an ontology of trust. *Lecture notes in computer science*, 3592:175, 2005. (Cited on page 17.)

[96] Frank Walter, Stefano Battiston, and Frank Schweitzer. A model of a trust-based recommendation system on a social network. *Autonomous Agents and Multi-Agent Systems*, 16(1):57–74, 2007. (Cited on pages 24 and 45.)

[97] S. Wasserman and K. Faust. *Social Network Analysis: methods and applications*. Cambridge University Press, 1994. (Cited on page 78.)

[98] D. J. Watts. *Six Degrees: The Science of a Connected Age*. W. W. Norton & Company, 2004. (Cited on pages 20 and 80.)

[99] D. J. Watts and S. H. Strogatz. Collective dynamics of'small-world'networks, 1998. (Cited on pages xiv and 81.)

[100] J Zhao, C Goble, R Stevens, and D Turi. Mining tavernas semantic web of provenance. *CONCURRENCY AND COMPUTATION*, 20 (5):463, 2008. (Cited on page 18.)

[101] C. N. Ziegler and G. Lausen. Spreading activation models for trust propagation. In *The IEEE International Conference on e-Technology, e- Commerce, and e-Service*, , Taipei, Taiwan, 2004. (Cited on page 20.)

[102] CN Ziegler and Jennifer Golbeck. Investigating correlations of trust and interest similarity. *Decision Support Services*, 2(43), 2006. (Cited on page 20.)

[103] M. Zook, M. Graham, T. Shelton, and S. Gorman. Volunteered geographic information and crowdsourcing disaster relief: a case study of the Haitian earthquake. *World Medical & Health Policy*, 2 (2):2, 2010. (Cited on pages 2 and 3.)

DECLARATION

Hiermit versichere ich, dass ich bisher noch keinen Promotionsversuch unternommen habe.

Münster, Germany, March 2011

Mohamed Bishr

Hiermit versichere ich, dass ich die vorgelegte Dissertation selbst und ohne unerlaubte Hilfe angefertigt, alle in Anspruch genommenen Quellen und Hilfsmittel in der Dissertation angegeben habe und die Dissertation nicht bereits anderweitig als Prüfungsarbeit vorgelegen habe.

Münster, Germany, March 2011

Mohamed Bishr

Hiermit erkläre ich, nicht wegen einer Straftat rechtskräftig verurteilt worden zu sein, zu der ich meine wissenschaftliche Qualifikation missbraucht habe.

Münster, Germany, March 2011

Mohamed Bishr