Testing Citizen Science: Social Interaction In An Agent-Based Environment

Roosa Tammela

5th Year Project Report
Master of Informatics with Honours Informatics
School of Informatics
University of Edinburgh

2017

Abstract

Recent years have seen a surge in the popularity of virtual citizen science projects - a class of social machines dedicated to the pursuit of scientific goals by harnessing human aptitude for problem-solving and pattern recognition. The most prevalent example is the vast collection of projects hosted on Zooniverse, with subjects ranging from classifying bat calls to identifying meteors in radio data. Studies have been increasingly focusing on the various forms of motivation inspiring the public to take part in such projects, both in an attempt to explain the participant behaviours observed and to identify factors that might increase the success of citizen science projects.

This paper surveys the landscape of research into motivational factors involved in citizen science participation, particularly focusing on the role of social presence of participants. An agent-based model, previously built as a tool for testing the properties of citizen science social machines, is extended to simulate the network of social interactions produced by participants. Empirical data from existing studies is used as the basis for the simulated behavioural patterns, while concepts from social network analysis are used for quantitative processing of the resulting networks. The observations produced hope to shed light on the characteristics of community structure in citizen science as well as the factors involved in how these communities are shaped.

Table of Contents

Chapter 1

Introduction

Section 1.1 introduces the context of this report in terms of citizen science as the focus of study and on the previous work completed for Part 1 of this project. Section 1.2 presents the motivations for this project and the goals identified in Part 1. Section 1.3 presents the main contributions of this project.

1.1 Testing Citizen Science: An Agent-Based Approach

This report is a continuation of the work set out in Part 1 of my project on the study of social machines, entitled Testing Citizen Science: An Agent-Based Approach (2016).

1.1.1 Original Project Goals

The two primary goals outlined in the original project proposal were firstly, defining properties that characterise the intended behaviour of social machines, and secondly, developing tools and processes to test these properties. The topic of social machines is broad, covering systems such as Wikipedia, Youtube, and a multitude of other collaborative and crowdsourcing-based applications. The project goals involve modelling aspects of social machines, both in low-level user responses who participate in a system, and in the resultant high-level behaviour observed in the system overall. A completed model representative of a social machine can then be used to further investigate how a varying range of social responses affects the behaviour of the system.

The preliminary stages of the project narrowed down Zooniverse as an ideal system for study and modelling. Zooniverse projects allow its users to interact with the system in a clearly defined way by making contributions through small, discrete tasks in a structured environment. This format lends itself well to modelling, as there is an explicit set of ways in which a user can interact with the system. While approaches from various disciplines may be taken to produce a model, agent-based modelling was chosen for its capability of modelling generative behaviour, as that observed in social machines where the sum of many individual behaviours leads to a complex system, as well as for the clear translation of Zooniverse participants in a structured platform into agents in a simulated environment.

1.1.2 Implementation and Testing

Repast Simphony, an open-source tool utilising the agent-based modelling language Relogo, was chosen as the programming tool to model Zooniverse-style projects. Empirical data and behavioural observations from existing studies on Zooniverse projects was used as the basis for the model, both in

guiding the functionality required of the model as well as in determining the patterns of behaviour taken by agents in the model.

Overall, the model built for Part 1 was successful and yielded significant results. The generative agent-based modelling approach succeeded in capturing more complex system-wide behaviour. Testing focused primarily on the frequency of users' contributions, the rate at which they join and leave the project, and the makeup of the user population. The results were processed through a combination of numerical metrics and visualisations revealing patterns in user behaviour within each test setting. Analysis of these results revealed how changes in each property of user or system behaviour will affect overall behaviour of the system. The results produced by baseline settings for the various parameters were consistent with characteristics observed in Zooniverse projects, indicating that testing and analysis of the model's properties is predictive of behaviour in real systems. The model produced thus has the capability to provide insight into how performance of social machines can be improved if participant behaviour corresponding to the model's parameters are influenced in certain ways.

The conclusions drawn from my analysis reaffirm predictions laid out in previous studies on social machine participants, in that long-term persistent users are critical for the success of citizen science projects. Increasing the activity levels of these long-term users has the most significant impact on overall project contribution levels. Meanwhile, incentivising users of all types to increase their contribution rates also has a significant positive effect on citizen science project success.

However, the model so far is still relatively simplistic as it has no consideration for social elements, communities, or interactions between users, which have been shown to play a potentially significant role in social machines, including in citizen science specifically. The aforementioned topics remain an area for exploration, and were stated as goals for extension for the Second part of this project.

1.2 Project Motivations

Part 2 of this project will continue to pursue the goal identified during Part 1, namely characterising and testing properties governing the behaviour of social machines. As the primary goal for Part 2, this project will focus on the importance of social interactions in the context of social machines: interactions between volunteers, the communities they form, and the effects of interaction within these communities on the success of projects.

An additional goal for extension laid out in Part 1 was robustness testing, namely whether citizen science projects can still achieve their objectives when some proportions of users behave irrationally by making incorrect contributions to tasks. On further consideration, this goal is not a priority for two reasons. Firstly, citizen science projects already account for some degree of user error and irrationality by collecting a consensus from participants on each item to be classified, effectively eliminating erroneous contributions. Secondly, the subject of social interaction provides far greater room for exploration and interesting findings. As such, this project will focus solely on studying the social interaction involved in citizen science projects.

The model built in Part 1 will be used as an existing platform for Part 2, enabling social interactions between users to be modelled in conjunction with individual users' interactions with the system. Implementing this model will involve a series of steps to ensure a successful outcome. Firstly, background research on the importance of communities within social machines, particularly in citizen science and Zooniverse projects, will provide an overview of behavioural patterns and empirical data to act as guides during implementation. Further research into the technical aspects of analysing and studying the formation of communities, networks, and social interactions will inform the data collection and analysis process. Aspects from graph theory and social network analysis are promising subjects of study in this regard.

The implementation stage will follow a similar procedure to that taken in Part 1; parallel with implementing extensions to the model to simulate social interactions, empirical data will be drawn from to establish baseline values for parameters while acting as a form of validation to ensure that the model produces results comparable with real-life phenomenon.

The motivations for this project are further driven by the expanding research into how social machines operate and the communities formed within, as will be further discussed in Chapter 2.

1.3 Contributions

The contributions of this project are outlined below:

- Extension of an agent-based model built in Repast Simphony to simulate the social interactions of citizen science participants.
- Establishing a strong theoretical basis for the design and implementation choices, in addition to affirming the validity of the model based on empirical comparison.
- Systematic testing of the model produced to explore the significance of user behaviour in terms of social network formation.
- Analysis of resultant data in light of real-world observations.

1.4 Report Outline

The remaining sections of this project report are summarised below.

- **Chapter 2** explores recent research into motivation factors driving citizen science participation and social engagement, while addressing the challenges and opportunities in each.
- **Chapter 3** describes the process of implementing extensions to the model, and the additional elements introduced.
- **Chapter 4** outlines a process for using social network analysis to collect, process, and analyse data produced by the model.
- **Chapter 5** presents the results of systematic testing, as guided by the process outlined in Chapter 4.
- **Chapter 6** discusses the significant findings of this project and their relevance.
- **Chapter 7** presents the conclusions and wider contributions of this project.

Chapter 2

Background

Section 2.1 examines the field of research on motivational factors in citizen science. Section 2.2 provides an overview of social participation in citizen science and related concerns. Section 2.3 examines related work in agent-based modelling.

2.1 Motivation in Citizen Science

Motivation is a key area for study due to the very nature of citizen science; participants are volunteers who receive no monetary compensation for their efforts, and rely on other factors to drive them where each citizen science project may show great variance in the strength of these factors (Katmada et al., 2016). Indeed, as Roy et al. (2012) found in a review of various studies on motivation in citizen science, the motivations experienced by participants differ wildly among individuals and projects alike, with personal enjoyment and enthusiasm for a project's goals standing out as notable motivations.

The degree to which participants' motivations must be fulfilled also varies among projects according to the size of individual tasks involved, with larger tasks requiring greater motivation to approach. Nov et al. (2011) describe contributory projects, as well as collaborative projects such as the Zooniverse collection, as having large task granularity where the minimum investment in time and effort is relatively large. This is in contrast to passive computing projects of low granularity where virtually no effort is required from a participant. Sustained contribution by individuals is crucial for the success of citizen science, and so it is important to understand the motivations driving contributions.

Ensuring that volunteers are motivated to participate in virtual citizen science presents a great challenge for a multitude of reasons. Wiggins and Crowston (2014) observe that the vast majority of citizen science projects offer no monetary and few other extrinsic rewards for participation. In addition, Katmada et al. (2016) found that participants' motivations may each be fulfilled from one or more different incentives. Finally, the 'long-tail' phenomenon is widely observed in citizen science (Jackson et al., 2016; Kamar et al., 2016; Tinati et al., 2014; Luczack-Roesch et al., 2014) where roughly 70% of volunteers in virtual projects stay for only a single session. The above issues indicate volunteer motivation as a vital field of study to promote the success of citizen science.

2.1.1 Virtual vs Non-Virtual

Before examining the key motivational factors in citizen science, it is important to recognise the distinction between virtual and non-virtual citizen science, and how a participant's experience is

shaped by a project's design. Zooniverse projects fall into the virtual category, and share the scientific concerns common to the virtual category of replicating tasks for validation, in turn requiring volunteer motivation to be maintained to achieve a critical mass of contributors (Wiggins and Crowston, 2011).

A key difference between virtual and non-virtual citizen science is that non-virtual forms more easily allow for contributory projects where participants collect observations from their environment according to the project's guidelines. Contributory projects are organised by a variety of authorities and take place on any geographic or indeed global scale. For instance, The Monarch Larva Monitoring Program (Oberhauser and LeBuhn, 2012), organised by staff and students at the University of Minnesota, aims to study the local distribution of breeding monarch butterflies to inform conservation efforts, with the secondary goals of enhancing undergraduate education through data collection, research opportunities, and class projects. The iSPEX project is an example of a more geographically distributed citizen science initiative. Funded by the European Commission, iSPEX enables volunteers to use a smartphone-attached tool for measuring aerosol levels in European cities (Land-Zandstra et al., 2016).

The two examples above illustrate how the nature of communication and coordination within a citizen science project is determined by geographical distribution; in the Monarch Larva Monitoring Program, although a website and social media are readily available, volunteers have the opportunity to contribute and interact in person due to geographic proximity. The limited possibilities for face-to-face interaction between participants in virtual citizen science on the other hand poses an additional challenge in the design of virtual projects (Wiggins and Crowston, 2014).

Further complicating the understanding of participant motivation is that motivations may change over time. Rotman et al. (2012), in their survey and interviews of citizen science participants, found that the initial choice to take part in a project hinges on intrinsic motivations, mainly egoism and personal interest in the project topic. The choice to continue participation, on the other hand, was found to rely in part on extrinsic motivations of attribution and a sense of involvement with the scientific community. Likewise, the motivational factors driving each user may change dynamically over the duration of their participation.

2.1.2 Intrinsic vs Extrinsic Motivations

As alluded to above, motivational factors can be divided into intrinsic factors, such as personal interest, and extrinsic factors, such attribution and acknowledgment by others. Each of these groups of motivation have been found to be relevant in citizen science at various stages of participation.

Extrinsic Motivations

Citizen science projects offer relatively few extrinsic motivations due to lack of monetary compensation for participants. Katmada et al. (2016) found that in crowdsourcing platforms in general, common extrinsic motivations include self-marketing and social motives, which are driven by a need for respect by organisers and peers. However, Wiggins and Crowston (2012) observed that among a variety of citizen science projects assessed, many projects lack clear data ownership policies, making it difficult to provide volunteers with due credit for their contributions.

Intrinsic Motivations

Volunteers' curiosity about the natural world and inherent interest in a project's subject have been found to be consistent motivators in all forms of citizen science. Nov et al. (2011) found that intrinsic motivations are the most significant in increasing participation, in addition to collective motivations of feeling part of the scientific community. Devictor et al. (2010), in a review of how citizen science is used to address ecological conservation issues, found that contributory projects rely on participants' curiosity in observing features in their own surroundings and in discovering a connection to their environment. This finding is supported by Price and Lee (2013), who reaffirm individual curiosity as a driving source of interest, particularly in contributory projects.

Personal enjoyment and the desire to contribute to scientific projects have also been identified as key motivations (Reed et al., 2013). In a survey of EyeWire volunteers, a gamified project to map neurons within the brain, gamified elements were a significant source of motivation for students participants, as it provided a form of procrastination which offers the satisfaction of being 'productive' by contributing to a scientific goal (Tinati et al., 2016). The desire to contribute and interest in the subject were primary motivations for 20% and 19% of participants respectively. Although the desire to learn was not found to be a prevalent motivation, EyeWire inspired 85% of surveyed participants to learn more about the human brain.

Similar findings were observed in a survey of participants of the iSPEX project, where the most important motivations were contribution to a scientific cause (27%), potential improvement of local environmental quality (11%), and the novelty factor involved in the project (10%) (Land-Zandstra et al., 2016). However, the authors note that iSPEX participants were more likely to come from backgrounds closely related to science, which may affect the motivations experienced by participants.

Katmada et al. (2016) likewise found similar motivations for crowdsourcing platforms in general. These include learning through access to the knowledge and feedback of experts, a sense of altruism through contribution to a good cause, and enjoyment and intellectual curiosity derived from new experiences.

Novelty Factor

The 'novelty factor' experienced from being the first to see a stunning image of a faraway galaxy or witnessing a rare natural or astronomical phenomenon is a significant motivator across the breadth of citizen science, tying in closely with curiosity and personal enjoyment. The potential for utilising the novelty factor has been shown in Planet Hunters, where Jackson et al. (2014) found that making discoveries lead to contributors becoming sustained and more active.

This phenomenon has been verified by other studies. In a paper aiming to examine the effect of the novelty factor on participation rates of Higgs Finders volunteers, Jackson et al. (2016) conducted a controlled experiment showing messages to existing users whenever a never-before-viewed image was shown to them. The average number of annotations made by the experimental group increased to

64.7, compared to 34.4 made by the control group. As a method for retaining existing participants, drawing on volunteer's enjoyment of novel experiences is extremely effective. Jackson et al.'s paper corroborates experimental findings from Part 1 of this project, which suggested that increasing the contribution of persistent volunteer leads to more significant results than recruiting new, and potentially short-staying volunteers. However, as a criticism to overly relying on the novelty factor as a motivator in citizen science, any project will at some point run out of novel items to show users unless project managers constantly update the dataset.

2.2 Social Participation in Citizen Science

A strong link has been found between high levels of social participation and the fulfilment of motivations linked to a sense of community and participation in the scientific process. However, most studies agree that participants of the vast majority of citizen science projects do not engage in high levels of social participation for a variety of reasons, and suggest methods to increase participation.

2.2.1 Communication & Intervention

The importance of communication by project organisers has repeatedly been stressed as a driving force for participation. Clearly relaying a project's mission and goals increase participants' sense of commitment, and highlight the importance of contributions (Nov et al., 2011). E-mail communication from scientists and organisers regarding upcoming projects and papers published as a result of volunteer contributions are strong motivators (Eveleigh et al., 2014). Participants have even stated their desire to receive such communication; 42% of iSPEX volunteers reported that they want to receive e-mails reminding to take aerosol level measurements (Land-Zandstra et al, 2016).

The role of well-timed intervention messages to retain participants in existing projects has also been highlighted. Kamar et al. (2016) conducted an experiment with Galaxy Zoo participants, where machine learning was used to send intervention messages to users predicted to leave. Messages emphasising the value of volunteer contributions were found to be effective only when shown according to disengagement predictions; displaying messages about helpfulness of contribution at predicted times made users generate 19.6% more contributions than if the messages were displayed at random times. Authors hypothesise that this increase draws on motivations of contributing to a scientific cause, as messages aiming to ease anxiety over contribution quality, or informing of the platform's opportunities for social interaction showed no statistically significant increase. Rotman et al. (2012) corroborates this finding, suggesting that motivation be provided during points of decline in participation.

Social Media

Social media has increasingly been explored as an avenue for both attracting participants and as a platform upon which to conduct citizen science projects in order to boost participation (Nov et al., 2011). Social media, such as Twitter, also acts as a tool for distributing feedback as a means of increasing engagement (Roy et al., 2012). However, the uptake of social media by participants

remains lacking. Land-Zandstra et al. (2016) found that 69% of iSPEX (an aerosol monitoring project) participants were not active on social media related to the iSPEX project, although no explanation for low social media participation was found, nor a link between contribution levels and activity on social media.

In contrast to other researchers, Stafford et al. (2010) warn against over-reliance on social media as a platform for citizen science. Contributory ecological projects relying solely on social networks such as Facebook and Flickr for data submission face shortcomings such as missing tags on photos and misclassification due to lack of a framework supporting expert opinions. However, Stafford et al. note that several location-specific projects have successfully relied solely on Flickr as a platform. Indeed, for short-term projects, utilising public social network groups and existing communities can result in attracting significant numbers of contributors.

Community Engagement

In a similar vein to Stafford et al.'s suggestion, Crain et al. (2014) state that designing citizen science projects around existing communities may improve participation as contributors will have a common inherent interest in participating. Despite the low rates of community engagement observed, there is active discussion on whether engaging in social participation alongside regular contribution tasks can lead to higher rates of contribution by volunteers.

Roy et al. (2012) found a sense of involvement with the scientific community to be an important motivation, whether this is through direct communication with scientists on project forums, or through blogs posted by scientists. Despite a recent push for projects to be more involved in social media, Roy et al. found that only 9% of citizen science participants became involved through social media. Case studies examined by the authors further highlight the need for constant communication to maintain contributor interest, while some projects were found to have an explicit focus on building online communities as a way to enable communication. However, Roy et al. found that a small minority of users carry out all interaction on forums with a silent majority simply observing. Eveleigh et al. (2014) agrees with this finding, stating that 'dabblers' in Zooniverse projects describe their experience as solitary, and do not actively take part in platforms for social interaction. However, Østerlund et al. (2014) suggest that based on the high number of forum views, the silent majority find some benefit from simple observation.

A solution to increasing forum participation in Zooniverse is suggested by Greenhill et al. (2016), who propose that gamification and social interaction for the purposes of 'fun', as opposed to discussions strictly about science or projects, can encourage motivation. This might also make the forums more welcoming by placing less focus on scientific expertise, which newcomers may find daunting.

Even if community participation is increased, it is not likely to be a primary source of motivation for contributing to a project. Among EyeWire users, being part of a community was the least important reason for taking part in the project. However, a sense of community was the biggest reason to interact on chat despite the relatively low rate of social interaction observed, with only 12% of participants active on the live chat (Tinati et al., 2016).

The rate of social participation in Zooniverse projects has been observed to be much higher, with 40.5% of all users having contributed to both classification tasks and Talk discussions (Luczak-Roesch et al., 2014). However, social participation does not necessarily lead to active social interaction; over 90% of Talk posts observed were on task items, where threads are unlikely to form conversations between participants. Luczak-Roesch et al. note that poor interface design may lead to low levels of social interaction, a new users may not be aware of Talk features. However, among those who contribute socially, the power-law characteristics common to citizen science contribution rates are apparent: less than 1% of users were found to contribute 70% of Talk posts.

Feedback and Learning

Much of the reluctance by volunteers to participate in forums appears to stem from insecurity of their own subject expertise, or even that of others'. Roy et al. (2012) Participants recognise importance of citizen science, but misguidedly believe the quality of citizen science contributions are low (particularly for contributory projects). This issue is exacerbated by the prevalent project design of not allowing participants to see how others have completed a task, in an effort to prevent bias in contributions which might lead to inaccuracies. Mugar et al. (2015) describe the resulting lack of feedback as the primary hurdle discouraging new contributors from engaging both in contributions and socially. On the other hand, social participation may provide the solution for this dilemma, as it has been shown to act as a learning tool for participants and a way of improving confidence in their expertise of the subject.

Price and Lee (2013) suggest that without a shared physical space for participants to interact in, community formation can prove difficult, and the potential for learning is missed. Communities centered around citizen science projects promote improvement in scientific literacy as an additional benefit beyond increased levels of participation (Tinati et al., 2014). In Price and Lee's survey of Citizen Sky participants, social interaction and working with others were found to make projects feel more interesting by allowing participants to be involved beyond taking a role as anonymous data collectors.

Jackson et al. (2016) hypothesise that, as newcomers read Talk posts by experienced contributors, their quality of contributions increases. However, newcomers from various backgrounds with a shared interest in astronomy were reluctant to participate on forums due to lack of technical vocabulary and perceived lack of experience or subject knowledge. In contrast, 70% of users describe Talk pages as helpful for newcomer learning (Østerlund et al., 2014). Those newcomers who stayed beyond a single session were found to engage in varying levels and durations of contribution and social interaction. Jackson et al. found a clear divide in how newcomers and expert users participate in forums; Talk pages related to classification items were frequently visited by newcomers while forums about scientific discussions were visited mostly by experienced users.

Community Roles

There is a stark contrast apparent in the roles that participants can take in citizen science and other forms of social machines. Jackson et al. (2016) discuss the difference in Wikipedia contributors, who may adopt clear paths to become registered users or technical administrators with a well-developed

support framework in place. However, citizen science offers the single role of 'participant', with no explicit introduction or a showcase of work by other participants. Jackson et al. suggest that without formal roles for newcomers to fill, and the requirement to self-organise activity, it is important to display existing work of the community despite this clashing with design goals of preventing bias among participants.

On the other hand, many citizen science projects with opportunities for social engagement have seen small groups of users take on roles of core members of their respective communities with extremely high rates of social participation. For example, in EyeWire, 10.9% of all users made 95.6% of all chat messages (Tinati et al., 2016). Likewise in Zooniverse projects, 1% of all users were responsible for 72% of Talk posts (Tinati et al., 2014).

2.3 Related Work

This project continues to use agent-based modelling as a tool for investigating social machines. Existing applications support the applicability of agent-based modeling for this purpose, as there is a close interaction between social agents in modelling and observations of behaviour in real societies (Moss and Edmonds, 2005). If expert knowledge and empirical findings about social relationships are incorporated into an agent-based model, the model can then serve as a generator for data of the type observed in real life.

Agent-based simulation has been successfully applied to the study of various social phenomena. An example is PsychSim (Marsella et al., 2004), an agent-based simulator to study social influence in groups. PsychSim incorporates each agent with a decision-theoretic model of the world, and so allows agents to be represented with a psychologically motivated belief system which they use in attempt to exert their influence on other agents. While PsychSim was largely built as a theoretical tool, it demonstrates the expressiveness of agent-based modelling environments and illustrates how complex behaviours and interactions between agents may be implemented.

Tools have also been built specifically for studying social networks in the context of agent-based systems. Driven by an increasing need for toolkits incorporating modelling and network analysis, Holzhauer (2010) developed ReSoNetA, a social network analysis and visualisation module for Repast J. The benefit of utilising agent-based modelling for network analysis lies in the easy representation of graphs, relationships, and communication present in object-oriented languages. ReSoNetA is available as a library for Repast J and supports integration for the model built as part of this project. However, Gephi is a more suitable tool due to its maturity in terms of the features offered, with a wider range of statistics and filters available in addition to intuitive visualisation tools (Gephi, 2016).

Chapter 3

Implementation

Section 3.1 provides an overview of the model developed for Part 1 of the project to be used as the basis for extension. Section 3.2 identifies the functionality that the modelled social interaction must fulfil. Section 3.3 outlines the components and their interactions that provide this functionality. Section 3.4 describes the parameters governing the social interaction functionality and identifies other relevant parameters for testing.

3.1 Base Model

Part 1 of this project focused on building an agent-based model to simulate and test properties of citizen science social machines. This model was built in Repast Simphony based on the defining characteristics of social machines identified in a review of related literature. A modular approach allowed separate classes to be responsible for representing the components and interactions involved in social machines.

3.1.1 Components

`UserObserver`

A default class for Repast Simphony projects written in Relogo, a language designed for agent-based modelling. This class controls the simulation environment, parameters, set-up functions including the creation and addition of Users, and the running of the simulation.

`Platform`

This class represents a citizen science project, which holds a collection of Items to be classified, keeps track of the contributions made to Items, and sends incentive messages to users.

`Item`

A representation of an individual classification task, which a User may interact with to perform a Contribution.

`Contribution`

Uses the Repast Simphony built-in class Link to form a connection between a User and an Item which they have classified. This allows the number of contributions to each item to be tracked as well as enforcing the requirement that a User cannot make multiple contributions to the same Item.

User

This class represents the participants of a citizen science social machine. During the simulation of the model, Users engage in behaviours in a probabilistic manner governed by several parameters, determined by which group the User is part of. A review of research during Part 1 of the project found an agreement that citizen science users can be classified into two major groups: 'transient' or 'dabbler' users, who contribute to projects only briefly and leave immediately after; and 'active' or 'stable' users, who form a much smaller minority but are sustained participants.

Each of the 'transient' and 'active' users have distinct parameters for the following behaviours:

- Leave rate: the likelihood that the User stops contributing to the project
- Enthusiasm: the response rate to incentive messages sent by the Platform

3.1.2 Interactions

The base model simulates the behaviour of a citizen science social machine through a series of interactions between its components during each time-step of the simulation. On initialisation, the UserObserver sets up the Platform with a specified number of Items, and creates initial Users and sets the parameters of each Transient and Active User. During each time-step, a User chooses to find a new Item to contribute to, or leaves the project based on their Leave rate. If the number of Users active within the project falls below a set threshold, the Platform sends an incentive message to Users who have left the platform, who may then rejoin based on their Enthusiasm parameter. When enough Contributions have been made to each Item, the Platform marks that Item as classified. After all Items have been classified, the project is considered completed and the simulation ends.

3.2 Modelling Social Interaction

The existing functionality of the model as described in the sections above is extended to allow the modelling of social interactions between users. The model will produce a graph in which nodes represent Users and edges represent social interactions between users. A single type of directed edge is sufficient to represent these interactions.

Based on existing research, participants of citizen science projects can be split into 'active' and 'transient' users with different behaviours. This is modelled by assigning different activity rates to each group of users. All Users will perform social interactions through Talks and Replies in a homogenous manner. Jackson et al. (2016) has shown that Zooniverse participants have wildly varying reasons for interacting on Talk, based on complex personal motivations, and there is no clear pattern as to how newcomers and active participants adopt roles within the community. As a simplification, all Users in the model will have their social behaviour approximated by parameters common to the Active and Transient groups respectively.

We can generalise that users will typically come across random Talk pages, with individual users' propensity to talk representing their overall rate of social participation across all Talk pages.

In order to account for the users who have been seen to take particularly active roles within citizen science communities, a 'feedback' parameter will be implemented, which strengthens the likelihood that a User makes social interactions, based on the social interactions that others have made with them. This parameter reflects the findings that a small minority of participants in citizen science forums take on the roles of core community members, becoming highly active social participants.

3.3 Components and Interactions

3.3.1 User

The existing features of the User class are retained, with additional extensions for dealing with social interaction. As before, the Users are partitioned into Transient users and Active users. The additional functions of Users implemented are:

- Create a Talk post, based on the user's Talk rate and the value of the Feedback parameter
- Create a Reply to some other user's Talk post, based on the Reply rate and the value of the Feedback parameter
- Reach to Replies from other users to update the Feedback mechanism

3.3.2 Talk

The Talk class represents a Talk post made by a User, and allows for a means to draw edges between socially active users.

3.3.3 Reply

The Reply interaction serves as the edge connecting Users who have interacted with each other. This is a directed relationship, with a single User at one end and a single Talk item at the other end.

3.4 Parameters

The additional parameters implemented are listed in Table 3.1 below.

Parameter	Type	Description
activeTalkRate	Double	Sets the chance that a contributing Active user creates a Talk item during each tick. Range: [0.0, 1.0]
activeReplyRate	Double	Sets the chance that a contributing Active user creates a Reply to an existing Talk item during each tick. Range: [0.0, 1.0]
transientTalkRate	Double	Sets the chance that each contributing Transient user creates a Talk item during each tick. Range: [0.0, 1.0]
transientReplyRate	Double	Sets the chance that each contributing Transient user creates a Reply to an existing Talk item during each tick. Range: [0.0, 1.0]
feedBack	Double	The strength of the feedback mechanism: on each Reply from another user, the Owner's feedback parameter is increased by this amount. The feedback acts as multiplier for the above parameters, increasing the likelihood of Talking and Replying. Range: [0.0, 1.0]

Table 3.1: Additional parameters implemented.

Chapter 4

Methodology

This section describes the procedure for using the model in testing. Section 4.1 outlines ways in which social network analysis is applied to the study of online communities. Sections 4.2 and 4.3 give the data collection and processing methods, respectively. Section 4.4 discusses the graph visualisations used for further analysis.

4.1 Social Network Analysis

The primary purpose of the extended model is to examine the network structure produced by social interactions of User agents, thereby providing a representation of and comparison for real-world networks. Given a graph of nodes representing users, and directed edges representing social interactions, some method for extracting relevant network characteristics must be applied. For this, we draw on the subject of social network analysis.

A variety of methods have been suggested for understanding the social structure underlying citizen science projects. An example is Murray-Rust et al's (2015) concept of a meshwork, where threads, representing users' pathway-like experience of the online interface, can entangle through social opportunities to strengthen the social fabric. This approach takes into account the user's perspective, and recognises that online interfaces are typically not engaged with in a structured and rigid manner. However, while a meshwork may lead to a more natural and user-oriented representation of social networks, it lacks the explanatory strength of more traditional methods.

Social network analysis has been widely applied in studying community structure, and offers a set of tools for meaningful analysis. Clauset et al. (2004) discuss existing algorithms used in social network analysis in finding communities in large networks, noting that the computational complexity involves quickly becomes overwhelming when used on large datasets, to the order of $O(m^3n)$, with m edges and n vertices. A proposed alternative uses the concept of modularity, a measure of the fraction of edges that fall within communities in an observed network compared to that expected from a fully random network. In Clauset et al.'s algorithm, a greedy optimisation algorithm finds sets of vertices whose joining leads to the greatest increase in modularity, resulting in a hierarchical structure of the communities discovered this way.

Other concepts from social network analysis have been used for analysing the behaviours of online community members. Angeletou et al. (2011) used a variety of numerical metrics to characterise forum users in order to analyse the emergent behaviour and evolution of community activity. They found that complex quantitative and qualitative parameters assigned to forum users can act as a predictor for community health, visible in metrics such as the number and frequency of post, population of long-term users, and reciprocity in user interaction.

Similarly, Hamill and Gilbert's (2010) performed an analysis of several models of social networks by simulating their characteristics using an agent-based model, and using metrics from social network

analysis to compare existing models with their own. Despite no clear conclusion being drawn on how a perfectly accurate social network might be modelled or constructed, Hamill and Gilbert's paper demonstrate the usefulness of social network analysis as a tool for characterising social networks. Handcock and Gile (2010) provide additional insight into how social networks might be studied by identifying two ways in which a social network might be viewed: firstly, as a stochastic process which allows for modelling, and secondly, as a non-stochastic and fixed structure from which observations can be made.

The research reviewed above supports social network analysis as an appropriate tool, particularly as a means of characterising the social network produced from the model in this project. Community detection has been stated as a complex problem, however, the networks produced from the model represent interactions already within some citizen science community. The processes by which agents interact to form a social network are generalised through a set of parameters, aligning with Handcock and Gile's view that social networks can be approached as a stochastic process.

4.1.1 Metrics

The following metrics from social network analysis will be used in evaluating the results produced by the model:

Graph Density

Graph density is the ratio of edges in a graph and the total possible edges. This metric gives an indication of what proportion of other users any single user is connected to, and is typically low in social networks (Hamill and Gilbert, 2010).

Range: [0,1], with 1 representing no connections between any nodes.

Clustering Coefficient

Clustering coefficient represents the extent to which nodes with an edge to some other node also have an edge between each other. This gives a measure of how tightly connected clusters of nodes are, and is likely to be high in social networks (Hamill and Gilbert, 2010). However, given that the social networks in citizen science will include a high number of lone nodes, the clustering coefficient of graphs produced by the model is likely to be lower than in typical social networks.

Range: [0,1].

Path Length

Path length is the minimum number of edges between two nodes. Gephi provides the average path length within the graph as a whole. Typically, a short average path length represents a network structure in which there is a high degree of connection between nodes overall (akin to Milgram's famous experiment on 6 degrees of separation). However, in citizen science social networks, the average path length metric is likely to be skewed downwards due to the high number of lone nodes.

Average Degree

The average degree of all nodes in the graph is the average of both in- and out-edges. This metric is representative of how many connections each agent makes, i.e. how many social interactions a citizen science user makes on average.

Modularity

While modularity is primarily used for finding communities in networks (Clauset et al., 2004), there is likely to be a high number of disconnected nodes in the graphs produced by the model. These lone nodes are representative of the majority of citizen science users who do not take part in social interaction (see the discussion in Section 2.2.3). In addition, modularity is typically used as a measure for how well sets of nodes can be joined to produce communities, while there is only a single community of interest in the graphs under analysis. The modularity algorithm used by Gephi (Blondel et al., 2008) produces a measure of how densely connected the nodes within a community are, thus the modularity metric will provide insight into the cohesiveness of the community formed.

Range: [-0.5, 1], with 1 representing perfect connectedness between nodes in the community.

Percent of Users who have Talked

As a further metric, the percentage of the agent population to have created a Talk item will be kept track of. This is accomplished by having each User create an edge to itself whenever they create a Talk item; using Gephi's filter tools, this metric can then be extracted. The edges created this way are filtered out from the calculations of all other metrics. The percentage of Talkers acts as a link to literature reviewed in Chapter 2, where many studies provide concrete numbers for proportions of Zooniverse volunteers who have participated in creating Talk threads.

4.2 Data Collection

The network structure produced by the model is collected by exporting a GraphML file from the Repast Simphony simulation. Despite Repast Simphony's powerful capabilities for making batch runs, the GraphML file must be exported manually for each simulation.

The resulting GraphML file contains node and edge data for all objects in the simulation, including Items (classification tasks) and Contributions (contributions to these tasks). To extract only the relevant nodes and edges (Users and Replies between them), a Python script is run on each GraphML file. A GraphML file with Users and directed edges between them is thus produced.

4.3 Data Processing

Gephi is used to perform quantitative as well as qualitative data processing on the GraphML files collected. Each of the metrics listed in Section 4.1.1 (with the exception of Percent of Users who have Talked) can be calculated from within Gephi's user interface, with the values collated in a results table in Chapter 5.

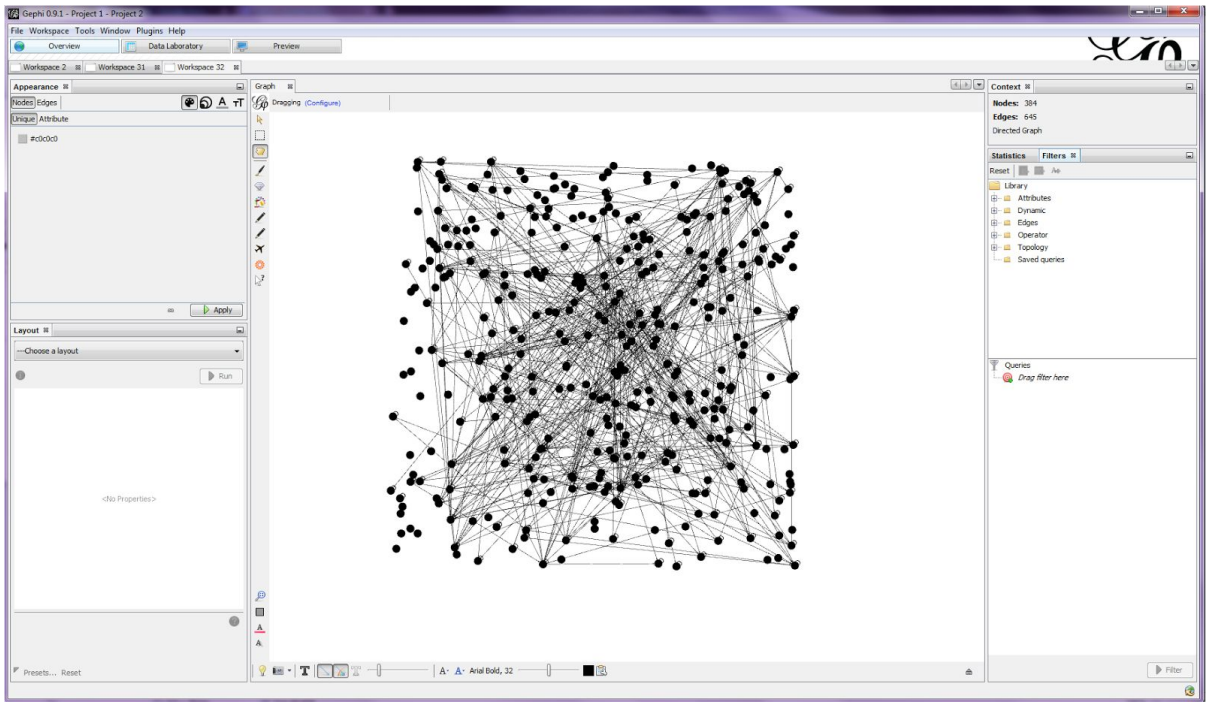

Figure 4.1: Gephi's user interface.

4.4 Graph Visualisation

The visualisation tools provided by Gephi will be used to provide an easily digestible overview of the network structures produced, as a form of qualitative analysis. The Yifan-Hu (Hu, 2005) Proportional graph layout, as implemented in Gephi, will be used for visualisation of graphs produced by the model. The Yifan-Hu algorithm is both efficient and produces the best results for large, densely connected graphs. A key benefit of the algorithm is its graph coarsening technique, which places highly connected nodes more accurately at the center of a dense network. For greater clarity and emphasis on the structure of connected portions of the graph, lone nodes will be left out from graph visualisations.

The aim of graph visualisations is not to provide any definitive measure of a network's characteristics, but to give some intuitive sense of how parameters alter the shape of a graph. The visualisations will act as a complementary tool alongside numerical metrics. The limitations of relying on graphs of extremely dense networks as a form of analysis should be recognised. For example, the visualisation below reveals only that the centre is dense and tightly connected, with a multitude of nodes with single connections along the edges, but provides little information about numerical metrics such as average degree.

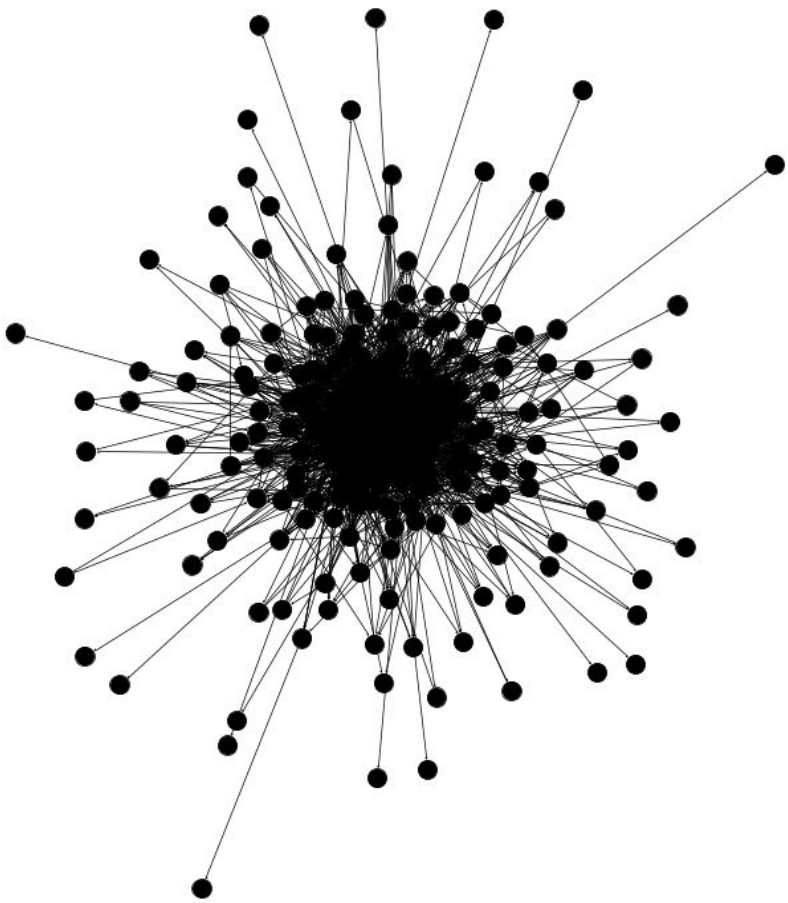

Figure 4.2: Example of a dense graph visualised in Gephi using Yifan-Hu Proportional

A close-up view of the centre of the above graph reveals that a high number of edges are concentrated within the central cluster, but simply expanding the visualisation to display each edge does not allow any more information to be discerned from it. The numerical metrics listed in Section 4.1.1 will thus remain the primary tool for analysis of results.

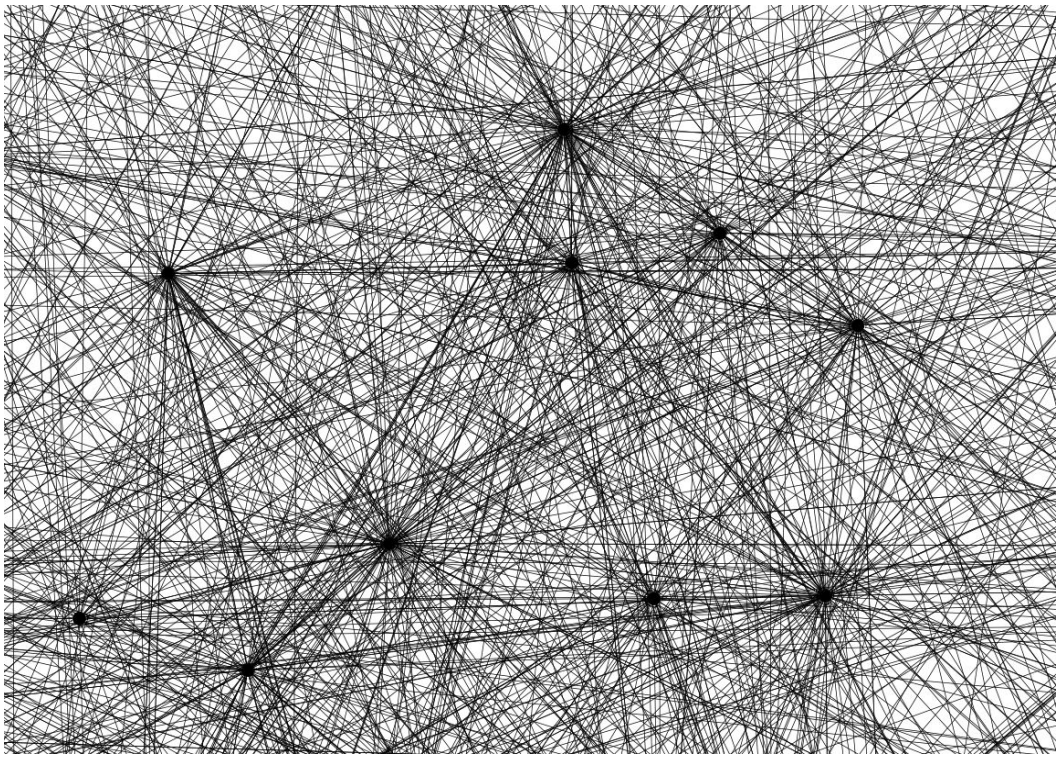

Figure 4.3: Close-up view of a densely connected graph.

4.4.1 Degree Distribution

In addition to a visualisation of the social network's graph, degree distribution provides an overview of the pattern of interaction between users. This is expected to follow a long-tail distribution when baseline parameters are used, as has been observed both in contribution rates and social interaction within citizen science.

Chapter 5

Results

5.1 Baseline for testing

Prior to experimentation, a baseline set of parameters must be established to form a valid point of comparison. Running the model with these baseline parameters should produce results comparable to empirical findings from a variety of literature reviewed in Chapter 2. Due to the limited amount of research focusing on the structures of social networks in citizen science communities, there is little empirical data available for establishing an accurate baseline; the data that is available is in the form of interaction rates over aggregate populations, whereas no studies have performed social network analysis to derive numerical metrics on citizen science social networks.

5.1.1 Existing Parameters

Part 1 of this project implemented functionality for simulating the contribution behaviour of citizen science participants. The baseline parameters for the variables involved were established and tested through an iterative validation process using empirical data more readily available for the related variables. The results achieved using these values were found to be in line with behaviour observed in real-life citizen science projects; the baseline values for the following parameters will therefore be retained for Part 2 of this project:

	User Parameters			
	User %	Leave rate	Enthusiasm	Rational %
Transient	0.7329	0.67	0.2	0.9
Active	0.2671	0.07	0.3	0.95
	Global Parameters			
	Initial users	Join rate	Join amount	
	100	0.2	3	
	Project Parameters			

Project items	Item validation	Automatic incentives	Incentive user amount	Incentive trigger (users)	Classification majority
1000	5	on	100	2	0.5

Table 5.1: Baseline parameters established in Part 1

The above variables continue to control the basic functioning of the model. Both the User parameters and Global parameters affect the amount and duration of contribution of Users, which directly influences the formation of social networks by affecting how long users interact within the system, and with how many other users. A combination of User, Global, and Project parameters influences the total ticks that the model's simulation runs for, also affecting the size and characteristics of the network as observed at the end of the simulation. Using a set of parameters already established as an accurate baseline is therefore the logical choice, given the goal of simulating and analysing the structure of real-life social networks.

5.1.2 Talk and Reply Rates

The additional variables introduced in Part 2 of this project control the behaviour of agents in making social interactions. These variables are Talk and Reply rates, in addition to a Feedback mechanism. As discussed, the only empirical data available to form a baseline with consists of aggregate data over all Zooniverse participants, with no distinction made between different groups of users. A few studies have focused specifically on observing newcomer behaviour, however, a newcomer may become either an active or transient participant. It is assumed that newcomers may fall into these two groups in proportion to the overall distribution of active and transient users observed in Zooniverse projects; empirical data from studies on newcomer behaviour can thus be treated as aggregate data across both user groups.

Both Luczack-Roesch et al. (2014) and Tinati et al. (2014) found that 40% of Zooniverse participants took part in both project contributions and on Talk forums. Contrastingly, Jackson et al. (2016) found that 16% of newcomers to Zooniverse posted in Talk forums, however, no specific timeline for the duration of observation was given.

The patterns of social behaviour have been found to be similar across social machines (Tinati et al., 2014). The findings by Land-Zandstra et al. (2016) on the iSPEX aerosol monitoring project, where 31% of participants were active on the project's social media, are therefore also indicative of what values the experimental results should reflect.

No data is available on the frequency at which Zooniverse participants reply to other users' Talk posts; it is assumed that Reply behaviour is comparable to Talk behaviour, and the same parameter values will be used for both variables. In addition, studies have not made a distinction between the social behaviour of different groups of users, although users have been found to take on distinct social roles within their communities (Tinati et al., 2014). The baseline Talk and Reply values will be equivalent for Active and Transient users, with the Feedback mechanism acting as a way to differentiate a small set of users into roles of active community participants.

A significant hurdle on the data available on how frequently citizen science participants engage in social participation is that data in studies was collected over long time periods of up to several months. These values cannot be directly mapped to the Talk parameter in the model, as the parameter controls the likelihood that a User creates a Talk post during each tick of the simulation. An appropriate parameter value must be derived through experimentation, and observing the metrics calculated at the end of a simulation's run.

5.1.3 Feedback mechanism

Small groups of participants within social machines have been observed to take active roles within communities over time. Jackson et al. (2016) has observed the user role of an active community participant to occur across the breadth of social machine systems from Wikipedia to Zooniverse, where a small minority of participants contributes a significant portion of a system's social interaction. Tinati et al. (2014) found that about 1% of Zooniverse users contributed 72% of Talk posts. Tinati et al. (2016) found a similar phenomenon in EyeWire platform, where 10.9% of users had made 95.6% of all chat messages.

The feedback parameter allows the the shift in some users' behaviour towards higher social participation to be represented. However, the same issue arises of a lack of directly comparable empirical data. An appropriate baseline value for the feedback variable will be assigned based on analysis of the degree distributions obtained from experimentation with a range of values.

5.2 Initial Model

In searching for a suitable set of values to use for the Talk, Reply, and Feedback variables, the following goals were set:

Firstly, approximately 40% of all Users should have created a Talk post during a simulation.
Secondly, the degree distribution observed should follow a long-tail distribution.
Thirdly, a small set of Users with disproportionately high levels of social interaction should be observed.

Without a clear starting point for parameter values drawn from literature, the following set of values was estimated as a starting point:

Talk Rate	**Reply Rate**	**Feedback Strength**
0.16	0.15	0.05

Table 5.2: Initial parameter values.

These initial values produced a network in which 52% of Users had created Talk posts, suggesting that the Talk rate (and correspondingly Reply rate) were set much too high. In addition, the degree distribution did not clearly reflect a long-tail distribution typical in social machines.

After several rounds of iterative tweaking of variable values, the following set of values was settled on, with the corresponding set of metrics below:

Talk Rate	**Reply Rate**	**Feedback Strength**
0.09	0.09	0.10

Table 5.3: Baseline parameter values.

Graph density	**Clustering coefficient**	**Path length**	**Average degree**	**Modularity**	**% of Users to Talk**
0.005	0.041	3.14	3.352	0.274	44.44%

Table 5.4: Metrics produced from baseline parameters.

The percentage of Users to create a Talk post, at 44.44%, is in line with the findings by Luczack-Roesch et al. (2014) and Tinati et al. (2014) that about 40% of Zooniverse users create Talk posts. Note that there is some variation in the metrics produced by a Repast Simphony simulation run with the same set of parameter values, due to the random seed for each run affecting the probabilistic calculations inherent in the model. This variation also has an effect on the number of ticks (time-steps) that the simulated citizen science project is completed in, altering the length of time that Users are active and able to participate in social interaction.

The degree distribution observed with the above set of parameters shows a clear long-tail distribution, with the majority of Users showing a degree of 1, and a small subset exhibiting high degrees, up to 34. This result is representative of findings in research, where small groups of users are observed to have high rates of social participation. Luczak-Roesch et al. (2014), for instance, found that about 1% of participants were responsible for 70% of all Talk posts; this is reflected in the degree distribution graph, where a few individual agents exhibit disproportionately high degrees.

There is no source of direct comparison for the social network analysis-related metrics, however, Hamill and Gilbert (2010) note that real-life social networks typically exhibit low graph density and a high clustering coefficient. Without some form of comparison, it is difficult to determine whether the metrics seen in Table 5.4 show the graph produced by the model with these parameter values as similar to a real-life social network. The graph density metric is obviously quite low, at 0.005, but little can be said of the other metrics. The three goals set out for baseline value testing have been met; the set of values in Table 5.3 will thus be adopted as the baseline values for experimentation.

The degree distribution and a visualisation of the graph are provided in Figures 5.5 and 5.6 respectively.

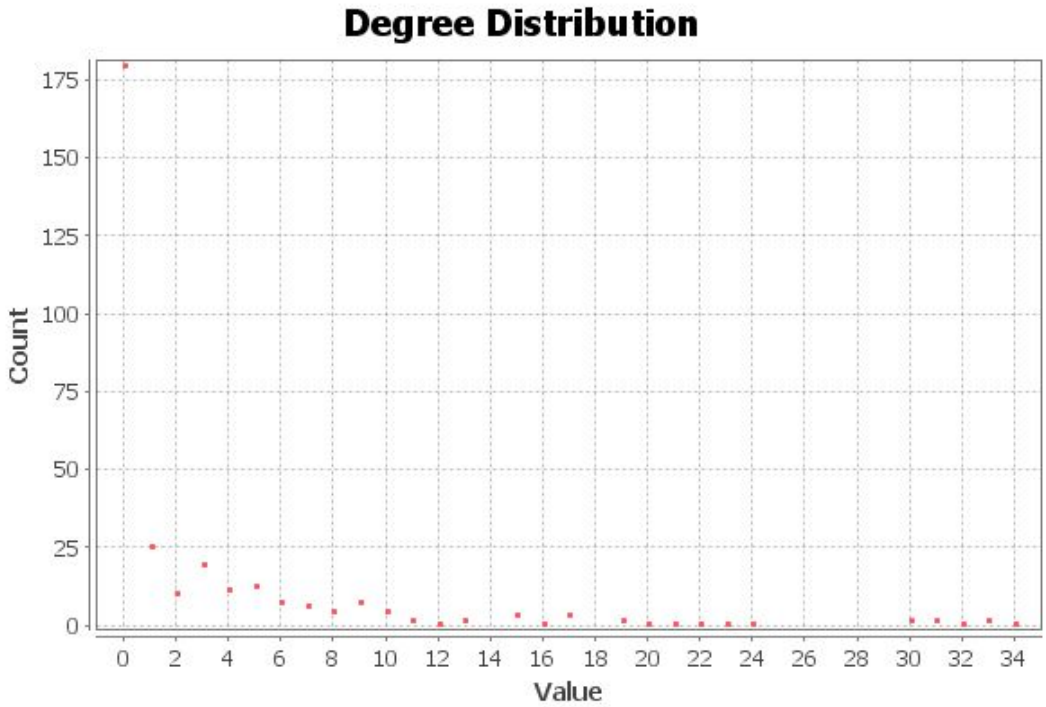

Figure 5.5: Baseline degree distribution.

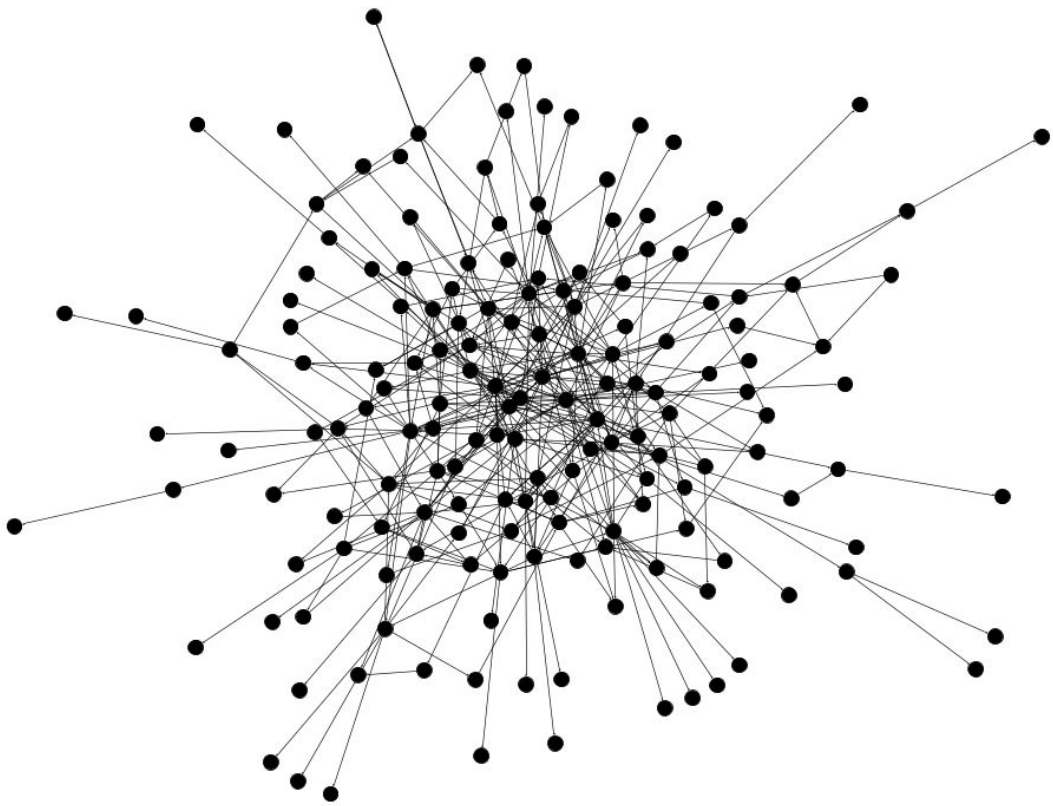

Figure 5.6: Baseline graph visualisation.

5.3 Test Results

The baseline values for parameters as established in the previous section will act as the default values, while one variable at a time is altered during testing. The data collection and processing methods described in Chapter 4 are then used to derive metrics and visualisations for each set of results. Analysis of each parameter individually will then reveal how change in that parameter affects the structure of the social network produced by the model.

In addition to the new parameters introduced in Part 2 of the project, the effect of user balance between Transient and Active users will also be tested. Altering the proportions of these user groups during testing is relevant in showing how the community structure may change if a larger number of users remains active upon joining a project.

Systematic testing of each parameter will use the test cases laid out in Table 5.7 below, with non-altered parameters maintained at their baseline values.

Test	Parameter altered	Baseline value	Minimum value	Maximum value	Value interval
1	Transient Talk Rate	0.09	0.00	0.50	0.05
2	Active Talk Rate	0.09	0.00	0.50	0.05
3	Transient Reply Rate	0.09	0.00	0.50	0.05
4	Active Reply Rate	0.09	0.00	0.50	0.05
5	Feedback Strength	0.10	0.00	1.00	0.10
6	Transient User %	0.73	0.00	1.00	0.10

Table 5.7: Test case specifications used for the testing of the model.

In addition to the test cases above, the evolution of the network structure over time will be tested by extracting metrics and visualisations from the model running under baseline parameters, every 100 ticks.

5.2.1 Transient User Properties

Transient or ‘dabbler’ users have been found to represent the majority of citizen science participants at about 70% (see Section 2.1), whose total contributions are nevertheless significant. Given that the Transient user group represents the majority of users, it is worthwhile testing how different levels of social engagement by Transients may affect the structure of social networks produced. The range of values for both Talk rate and Reply rate are tested from 0.0 to 0.5, where a value of 0.5 means that Transient users will on average Talk or Reply correspondingly every other tick.

Transient talk rate	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk
0.00	0.004	0.035	3.08	2.644	0.341	32.5%
0.05	0.004	0.044	3.26	2.706	0.338	40.11%
0.10 (baseline: 0.09)	0.004	0.022	3.89	2.652	0.348	46.26%
0.15	0.003	0.018	3.802	2.604	0.361	53.21%
0.20	0.004	0.029	3.648	2.714	0.368	51.89%
0.25	0.004	0.026	3.740	2.812	0.374	58.81%
0.30	0.003	0.013	4.375	2.500	0.448	64.8%
0.35	0.003	0.007	4.746	2.527	0.455	70.97%
0.40	0.004	0.012	4.477	2.635	0.445	71.07%
0.45	0.003	0.011	5.077	2.335	0.501	73.30%
0.50	0.003	0.016	4.232	2.497	0.469	74.05%

Table 5.8: Test results from altering Transient Talk Rate.

As seen above, no significant change is observed in the graph density. This is to be expected, as graph density is a proportion of edges with the total possible number of edges. Altering the Talk rate changes only the number of Talk posts made, not the replies formed between them. The average degree metric reflects this, as all Users still make replies at the baseline rate.

Increasing the Transient Talk rate from 0.0 towards 0.5 shows a decrease in the clustering coefficient from 0.035 to 0.016, and increase in path length by about 1. Comparing the baseline graph (Figure 5.6) with the graph for Transient Talk rate of 0.5 (Figure 5.9) shows that a wider and more sparsely connected graph is produced. This is explained by the rate of Replies remaining constant; with a higher Transient Talk rate, more replies are distributed to Talks by Transient users, who themselves are likely to have made only a few replies during their duration of activity. This results in the joining of sparsely clustered nodes to the central community while significantly increasing the graph diameter. The modularity, acting as a measure of cohesiveness, is also seen to increase with a higher Transient Talk rate.

From these results, we can interpret that Transient users creating more Talks significantly increases the users connected to the community, but does not increase the cohesiveness or clustering, which are indicators of tight connectedness between community members.

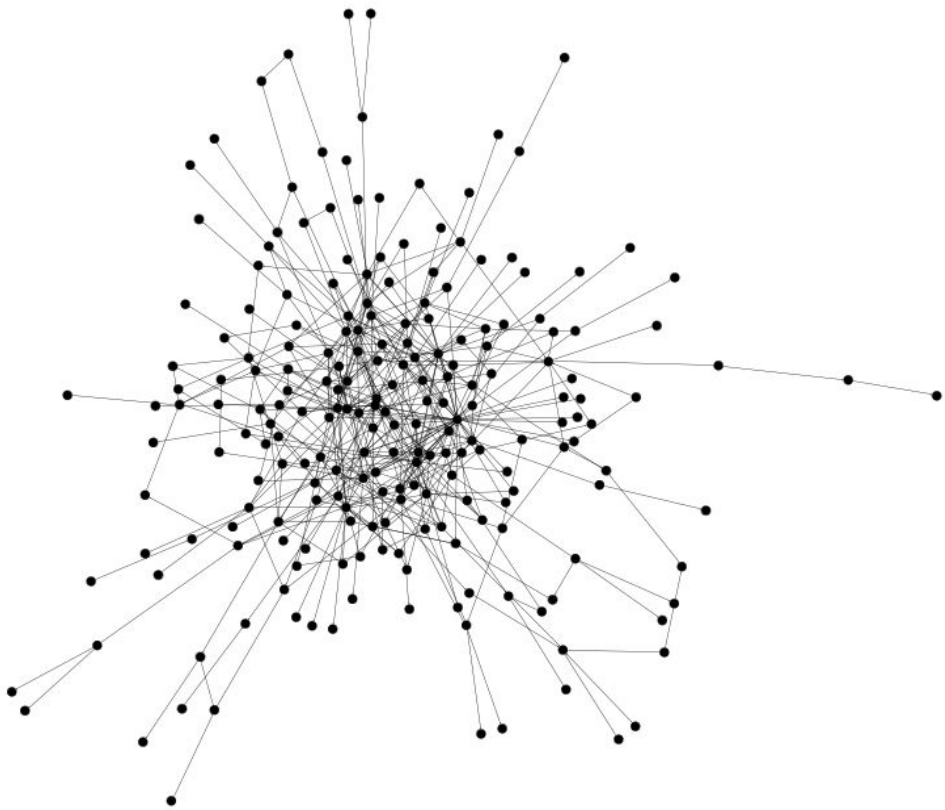

Figure 5.9: Graph visualisation with Transient Talk rate 0.5.

Transient reply rate	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk
0.00	0.004	0.044	3.179	2.630	0.303	45.06%
0.05	0.003	0.019	3.958	2.342	0.378	46.79%
0.10 (baseline: 0.09)	0.003	0.021	3.757	2.279	0.376	42.29%
0.15	0.003	0.024	4.066	2.503	0.402	42.75%
0.20	0.005	0.039	3.538	3.212	0.367	45.59%
0.25	0.004	0.029	3.915	2.995	0.385	43.01%
0.30	0.004	0.018	4.348	3.088	0.446	49.48%
0.35	0.004	0.021	3.903	3.328	0.389	43.49%
0.40	0.005	0.034	3.609	3.538	0.403	45.16%
0.45	0.006	0.033	3.647	4.050	0.383	49.72%
0.50	0.006	0.065	3.520	4.104	0.382	47.80%

Table 5.10: Test results from altering Transient Reply rate.

Altering the Transient Reply rate shows some increase in graph density (from about 0.003 to 0.006) and average degree (from 2.63 to 4.1) as the rate is increased from 0.0 to 0.5. There is a great degree of volatility in the other metrics, explained by the random elements incorporated into the model. In addition, the defining characteristic of Transient users is that they stay active within the system for a very short time. Even with a high reply rate, Transient users are likely to stop contributing before making a significant number of social interactions.

However, the observed increase in average degree is clearly reflected in Figure 5.11 illustrating the degree distribution for a Transient Reply rate of 0.5. Compared with the baseline degree distribution (Figure 5.5), there is a clear shift upwards in the degree distribution, with a more even spread of replies among users, particularly for lower values of replies.

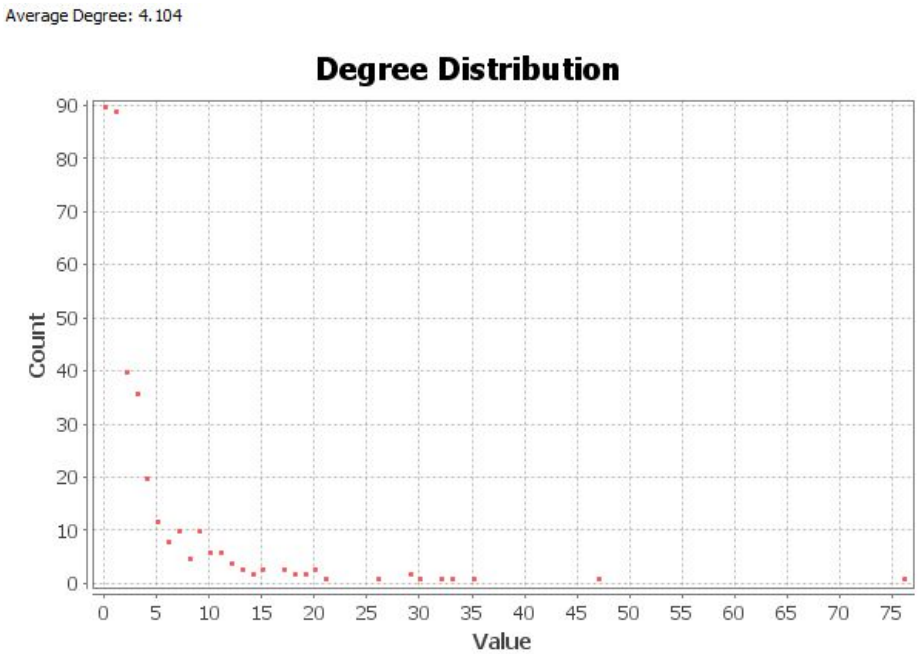

Figure 5.11: Degree distribution with Transient Reply rate 0.5.

The graph of Figure 5.12 illustrates that the visual structure of the network produced, with a Transient Reply rate of 0.5, is not significantly different from the baseline. An increased number of nodes with single edges appear to be connected around the community's periphery, corresponding to additional Transient users who have made a Reply before leaving the system. However, the structure near the center of the community appears to remain largely the same, as is also apparent in the clustering coefficient and modularity metrics of Table 5.10.

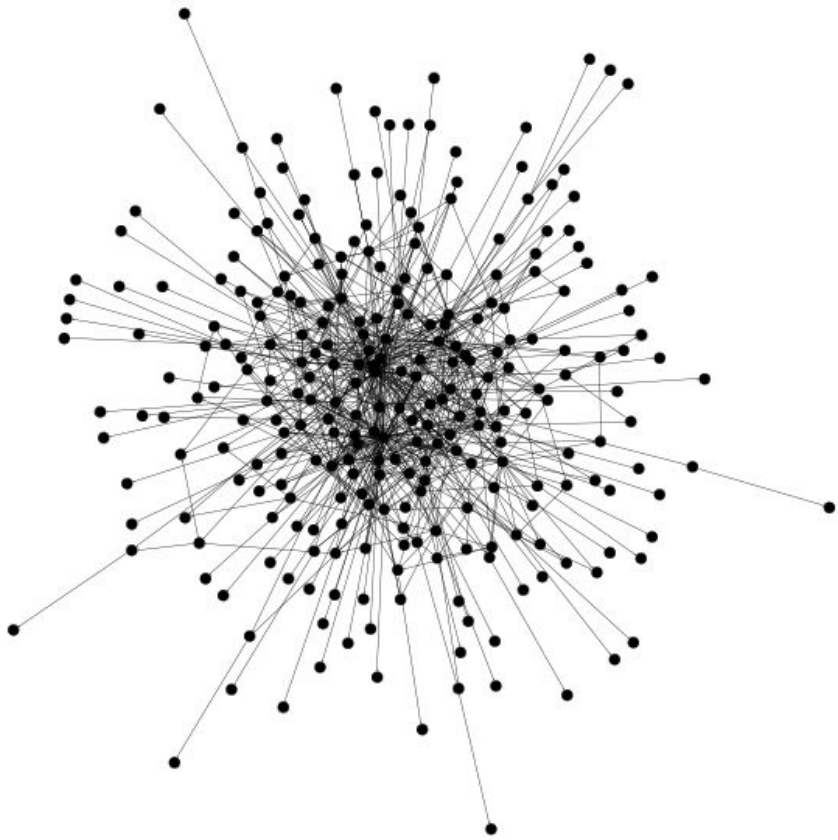

Figure 5.12: Graph visualisation with Transient Reply rate 0.5.

5.2.2 Active User Properties

Active users in the model are characterised by the propensity to stay active in the system for long periods of time. Any change in the pattern of their social interaction should result in a clear change in the structure of the network produced.

Active talk rate	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk
0.00	0.003	0.018	2.245	2.378	0.356	11.08%
0.05	0.004	0.029	3.890	2.988	0.334	42.26%
0.10 (baseline 0.09)	0.003	0.014	4.886	2.401	0.414	47.92%
0.15	0.003	0.033	3.607	2.606	0.352	50.80%
0.20	0.004	0.025	3.636	2.724	0.340	52.70%
0.25	0.003	0.017	3.660	2.401	0.387	52.28%
0.30	0.003	0.026	3.624	2.611	0.361	52.85%

0.35	0.003	0.029	3.701	2.561	0.382	54.76%
0.40	0.005	0.036	3.408	3.059	0.330	56.21%
0.45	0.004	0.027	3.900	2.864	0.375	57.10%
0.50	0.004	0.031	3.621	2.824	0.336	54.55%

Table 5.13: Test results from altering Active Talk rate.

As seen from Table 5.13 above, increasing the Talk rate of Active users only leads to a minor increase in the total percentage of Users who have created Talk posts. The same effect was observed when increasing the Talk rate of Transient users; increasing Talk rates does not change the interaction of users, as the rate of Replies is not affected. However, as Active users are very likely to create Talk threads during the total duration of their activity (with an exception for a Talk rate of 0), there is no observable change in the metrics as the Talk rate is altered. The small amount of variation in the metrics can be attributed to randomness in the model.

This result is not unexpected; during the implementation stage, a distinction was made between simply creating Talk threads, as a form of participation in the social aspects of citizen science, and between Replying to other users' threads, which constitutes social interaction. The latter behaviour is representative of community formation and a better indicator for community health, as measured through the clustering coefficient and modularity metrics (Hamill and Gilbert, 2010).

Active reply rate	**Graph density**	**Clustering coefficient**	**Path length**	**Average degree**	**Modularity**	**% of Users to Talk**
0.00	0.000	0.00	1.162	0.302	0.924	46.35%
0.05	0.002	0.007	4.738	1.479	0.483	45.57%
0.10 (baseline 0.09)	0.004	0.043	3.672	3.022	0.331	46.93%
0.15	0.007	0.056	3.114	4.483	0.258	48.26%
0.20	0.010	0.076	2.915	6.568	0.228	55.33%
0.25	0.010	0.083	2.832	7.142	0.191	45.36%
0.30	0.012	0.128	2.792	8.557	0.163	46.31%
0.35	0.010	0.099	2.705	8.227	0.166	41.16%
0.40	0.013	0.126	2.736	9.755	0.151	45.31%
0.45	0.016	0.149	2.767	11.067	0.154	47.49%
0.50	0.014	0.127	2.607	10.995	0.148	43.50%

Table 5.14: Test results from altering Active Reply rate.

The results of altering Active Reply rate can be seen in Table 5.14 above. The change in network structure with higher Reply rate is significant; there is a many-fold increase in graph density, clustering coefficient, and average degree with high values of Reply rate. Even a small increase in the reply rate by 0.05 leads to large changes in the aforementioned metrics; for instance, increasing the reply rate from 0.10 (which is close to the baseline value of 0.09) to 0.15 increases average degree by over 1.4. The overall decrease in path length with high Reply rate can be explained through the increased interconnectedness observed in the graph, which leads to a shorter distance between any two nodes. Particularly the graph density and clustering coefficient metrics show that Active users' levels of social interaction (i.e. Replies) are crucial in determining the structure and health of the community.

Interestingly, the value for modularity decreases as the Reply rate is increased. This can be explained by the increased number of edges and connectedness between Users making it more difficult to divide the community into separate sub-communities.

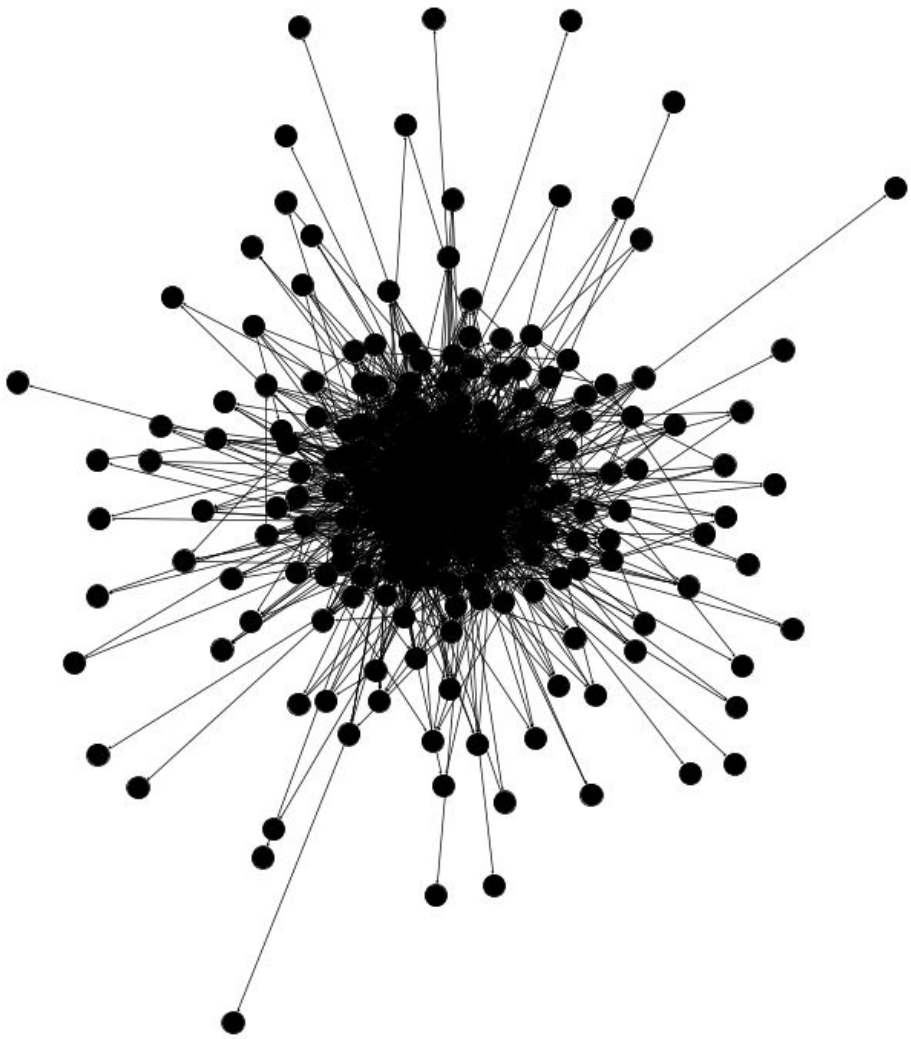

Figure 5.15: Graph visualisation with Active Reply rate 0.50.

The increasing connectedness with high values of Active Reply rate is seen in Figure 5.15 above, where the central nodes of the community show an extremely density of edges and nodes, compared with the baseline, although the visual structure becomes difficult to discern.

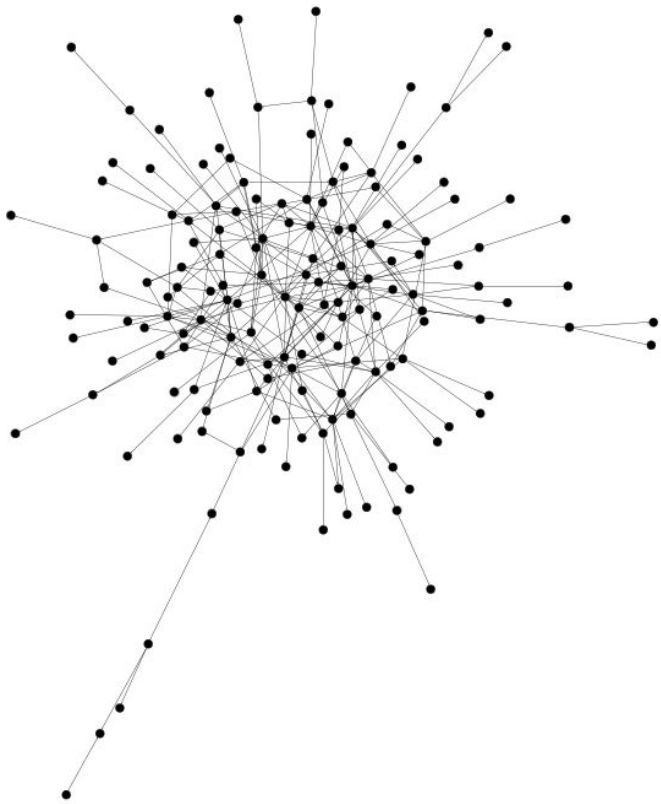

Figure 5.16: Graph visualisation with Active Reply rate 0.0.

The lower extreme end of Active Reply rate is illustrated in Figure 5.16 above. In this graph, where Active users make no replies at all, the structure appears extremely loose with most nodes having few edges.

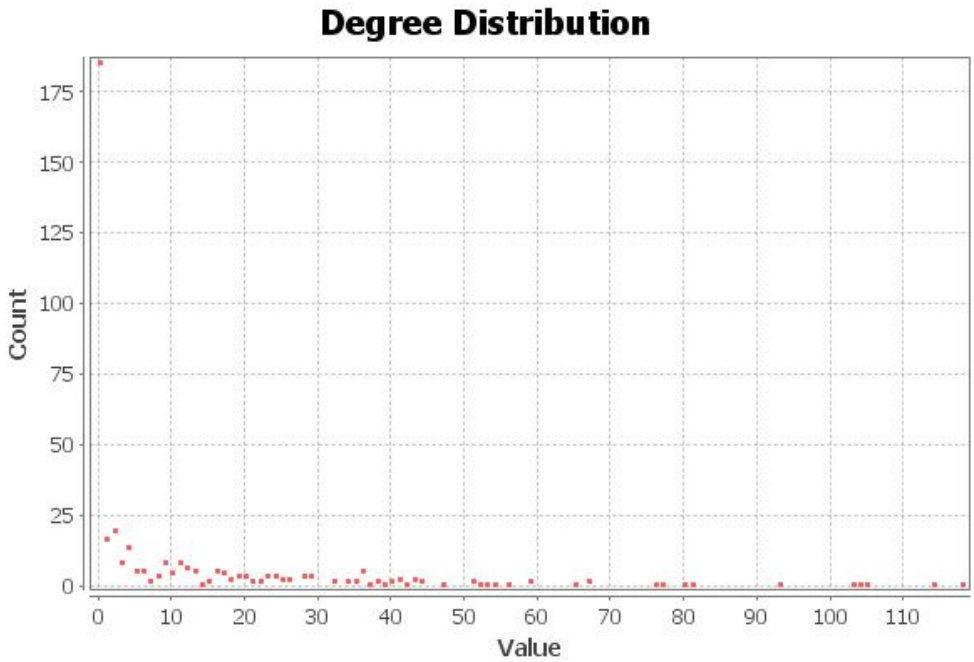

Figure 5.17: Degree distribution with Active Reply rate 0.50.

The effect of a high Active Reply rate on the degree distribution is seen in Figure 5.17 above; the interactions between users appear to become less evenly distributed, with a few individual users having far more edges than the majority. This can be attributed to the Feedback mechanism of the model, which increases a User's personal likelihood to both Talk and Reply whenever they receive a Reply from another User. However, this result is also consistent with the findings from research that a small proportion of users take on highly active roles within citizen science communities.

5.2.3 User Population Balance

The importance of the Active user population in forming social structure was reaffirmed in Section 5.2.2, where the Reply rates of Active users have a profound impact on community-related metrics. Investigating the effect of population balance on the network structure will likewise shed light on the relationship between the overall activity levels of users and community health.

Transient %	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk	Project age (ticks)
00%	0.007	0.026	3.542	3.717	0.345	61.96%	464
10%	0.008	0.027	3.734	3.836	0.368	69.26%	382
20%	0.004	0.013	4.316	2.601	0.408	55.36%	478
30%	0.005	0.020	4.509	3.053	0.388	59.67%	519
40%	0.006	0.030	3.913	3.312	0.363	60.51%	472
50%	0.005	0.019	3.861	2.880	0.370	51.27%	554

60%	0.004	0.020	4.031	2.629	0.354	44.62%	670
70%	0.003	0.030	3.305	2.522	0.358	44.78%	666
80%	0.003	0.016	4.302	2.268	0.407	44.74%	726
90%	0.003	0.027	3.612	2.475	0.339	43.43%	708
100%	0.004	0.015	4.669	2.709	0.442	63.08%	558

Table 5.18: Test results from altering the Transient/Active User balance.

From the results above, the only metric to significantly change is graph density, with a lower proportion of Transient users (i.e. higher proportion of Active users) corresponding to high graph density. The project age of each test was also recorded, and, as was observed in Part 1 of the project, a higher proportion of Active users leads to a shorter project duration. This can explain the lack of change in other metrics; as the project lasts for a less time, there is a shorter duration during which the community develops.

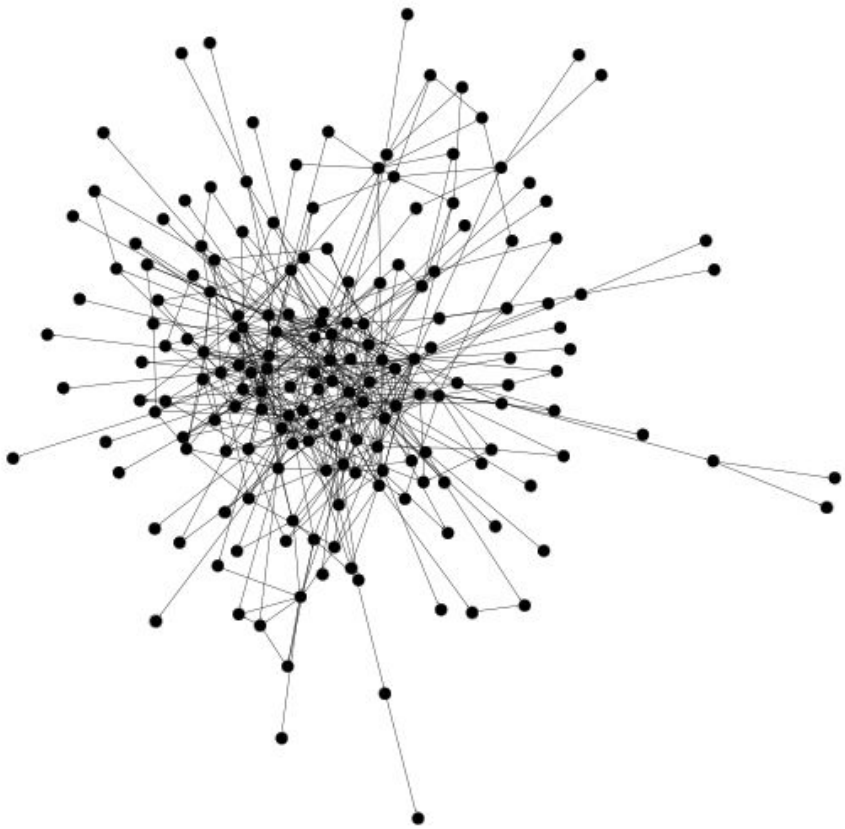

Figure 5.19: Graph visualisation with 30% Transient users.

The lack of change in the network structure is seen in Figure 5.19 above, where 30% of the user population are Transients. This graph appears largely similar to the baseline graph, particularly in terms of the edge density of the central portion.

A user population consisting of 100% Transients is an outlier, and shows a limitation of the model when dealing with extreme values. The incentive system implemented in Part 1 sends incentive messages to users whenever the active user count drops below a certain threshold, leading to a spike in user activity directly after. When the population consists only of Transient users who are likely to leave the system, the incentive system is triggered much more often, in turn leading to unexpectedly high levels of activity.

5.2.4 Feedback Strength

The Feedback parameter is a representation of users who take on a more active role within the citizen science community over time. As there is a lack of empirical evidence on the rate at which users take on this role and increase their social contribution, it is worthwhile experimenting on how the feedback parameter simulates this phenomenon with a variety of values.

Feedback strength	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk
0.00	0.003	0.017	4.634	2.254	0.418	46.49%
0.10 (baseline)	0.003	0.025	4.646	2.321	0.377	43.78%
0.20	0.005	0.047	3.264	3.121	0.307	45.76%
0.30	0.004	0.046	3.291	2.920	0.285	41.76%
0.40	0.005	0.086	3.043	3.371	0.271	44.35%
0.50	0.004	0.057	3.151	3.343	0.284	44.70%
0.60	0.005	0.088	3.133	3.478	0.274	48.35%
0.70	0.007	0.154	2.698	4.491	0.221	45.21%
0.80	0.008	0.120	2.822	5.095	0.216	49.68%
0.90	0.007	0.114	2.763	4.731	0.214	44.15%
1.00	0.007	0.118	2.674	4.876	0.197	42.37%

Table 5.20: Test results from altering the Feedback strength.

The results in Table 5.20 suggest that the effect of increasing the Feedback parameter is roughly equivalent to that of increasing the Active user Reply rate. Both have a positive impact on graph density, clustering coefficient, and average degree with a higher parameter value, and a negative impact on path length and modularity. The reasons for these effects are the same; a higher Feedback strength indirectly leads to an increase in Reply rates. The similarity between the effects of the two

parameters is further observed in Figure 5.22, where a graph visualisation of the network with Feedback strength 0.80 appears similar to the network with Active Reply rate 0.50.

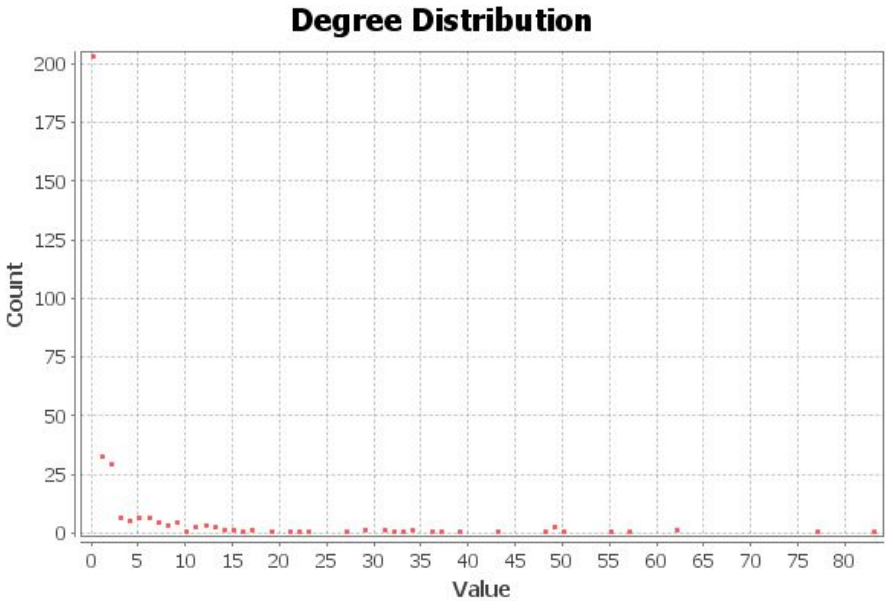

Figure 5.21: Degree distribution with Feedback strength 1.00.

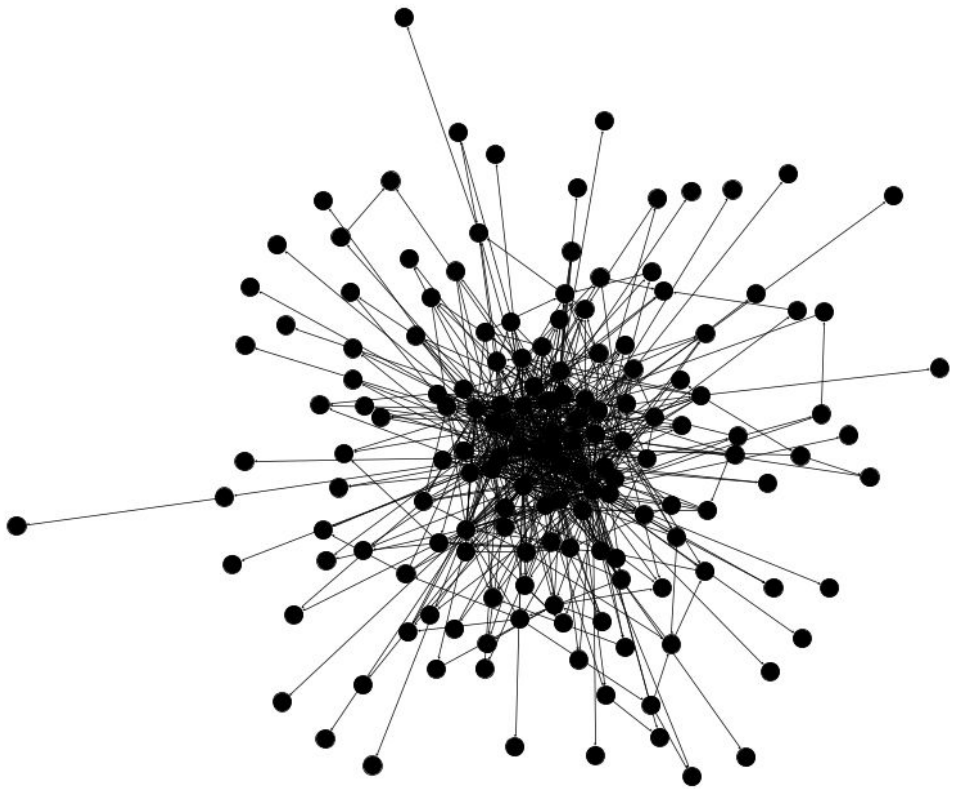

Figure 5.22: Graph visualisation with Feedback strength 0.80.

5.2.5 Evolution of Network

As a final experiment, the evolution of the network structure with baseline parameters over a single run of the model is observed. This will allow the increase in connectedness of the graph over time to be analysed. In addition, with the exception of studies on newcomer behaviour, existing research on

citizen science focuses on aggregate user behaviour collected over long time periods. Testing the change in network structure over small time-steps thus provides for a novel, if only theoretical, experiment.

Tick	Graph density	Clustering coefficient	Path length	Average degree	Modularity	% of Users to Talk
100	0.003	0.0010	2.945	0.845	0.437	26.06%
200	0.004	0.023	3.494	1.322	0.389	31.67%
300	0.003	0.018	3.570	1.513	0.381	33.48%
400	0.003	0.022	3.709	1.882	0.360	36.40%
500	0.003	0.019	3.738	2.006	0.348	37.82%
600	0.003	0.025	3.800	2.243	0.365	41.24%
624	0.003	0.023	3.814	2.212	0.377	41.03%

Table 5.23: Evolution of the baseline network structure over time.

The simulation of the baseline model for this test ran for a total of 624 ticks, with a snapshot of the network structure taken every 100 ticks, and once more at the end. From the results, it is apparent that the graph density remained perfectly constant over time. From previous experiments, only Active user Reply rate and Feedback strength were found to significantly change graph density, with higher values leading to higher graph density. This suggests that the graph density is determined by the aforementioned parameters, and is otherwise constant throughout the life of the modelled community.

Some variation was seen in the clustering coefficient, but no clear trend emerged after Tick 200. This also suggests that after the community has begun forming, the clustering coefficient will remain roughly constant; likewise for this metric, Active user Reply rate and Feedback strength were the only parameters found to affect the clustering coefficient.

The path length showed constant increase with Ticks; this can be attributed to lone nodes becoming joined to the periphery of the community at a faster rate than shorter paths are formed between already connected nodes. The average degree of nodes is nevertheless seen to show a steady increase over time, showing that despite new, mostly Transient, users constantly joining the system, the graph becomes more connected over time. This can be explained by the feedback mechanism, which provides a constant increase to the Reply rates of users during their social interactions.

The metric for percentage of users to have created at least one Talk post also shows an increase over time, seeming to stabilise when reaching ~40%. This observation highlights a concern with the findings of existing studies, in that aggregate statistics on citizen science participation over long time periods may not reflect the structure of communities still in the early stages of a project. On the other hand, the Zooniverse projects acting as subjects for such studies, such as that on Higgs Finders by Jackson et al. (2016), have continued for long periods of time as new data has been continuously added to the project. The communities of these long-term projects will likely have stabilised similarly

to what Table 5.23's experimental results show, so any aggregate findings on long-term studies on such projects are relevant and valuable.

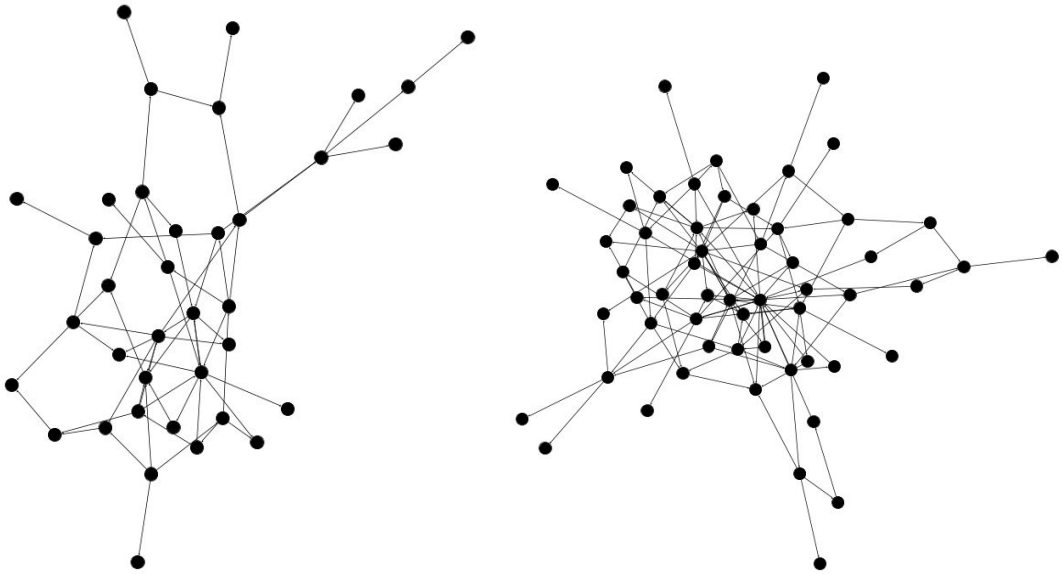

Figure 5.24: Graph visualisations of baseline model at 100 and 200 Ticks.

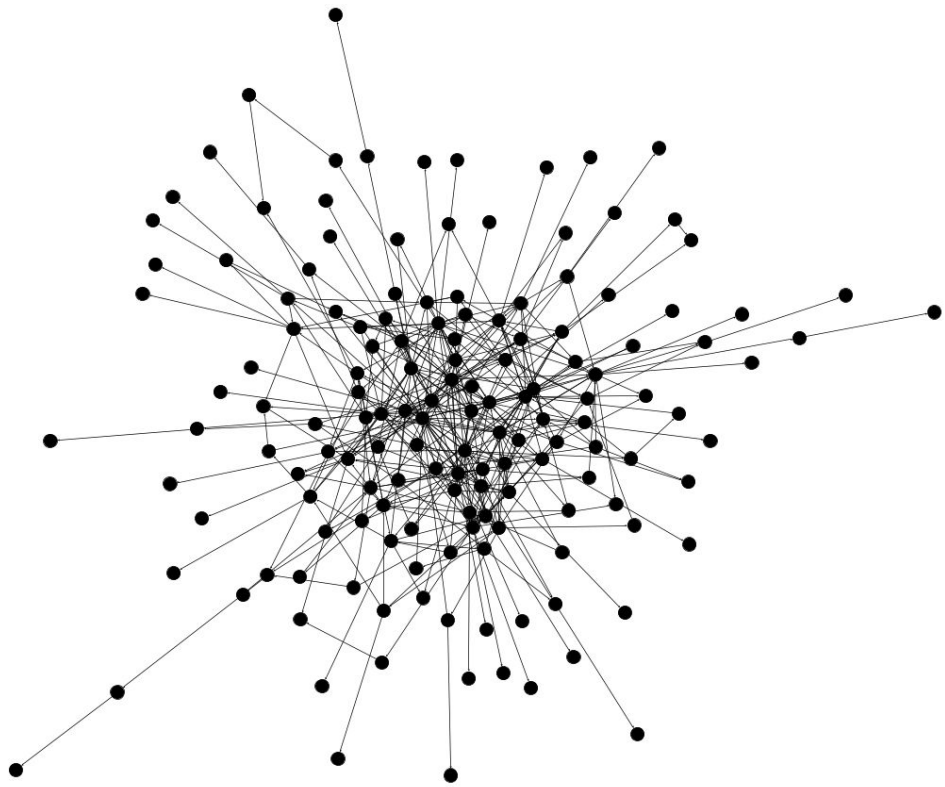

Figure 5.25: Graph visualisation of baseline model at 624 Ticks.

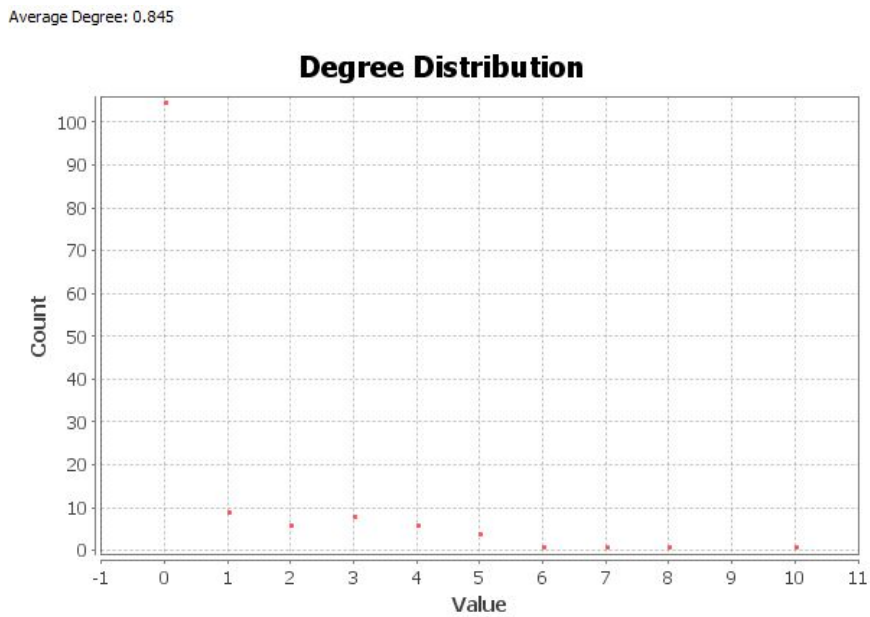

Figure 5.26: Degree distribution of baseline model at 100 Ticks.

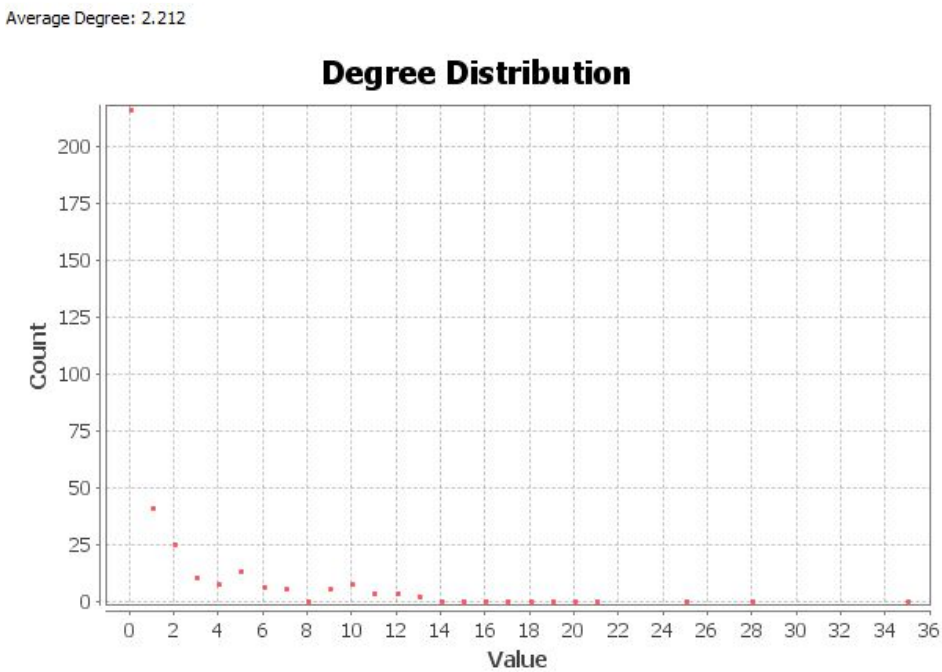

Figure 5.27: Degree distribution of baseline model at 624 Ticks.

Figures 5.24 and 5.25 illustrate the change in network structure over time, displaying both an increase in connectivity within the central cluster as well as the joining of nodes to the periphery, increasing path length. Figures 5.26 and 5.27 display the long-tail distribution of degrees over users, and shows the movement of a small group of users to becoming more active social participants with higher degree.

5.3 Overall Findings

Testing on the extended model produced results consistent with real-life observations where available, while establishing the significance of each parameter on affecting network structure. The only anomalous result was obtained when testing the User balance, where a 100% Transient population leads to the incentive mechanism of the base model triggering at abnormally high rates.

Although there was a lack of empirical data from which to draw baseline values for the parameters added as this part of the project, iterative testing found baseline parameters that produce a ~40% rate of Talking among agents, in line with the rate observed in various studies on Zooniverse projects. With a solid basis on theory, the model can therefore produce results that are valid in studying real-life systems.

The most significant finding is that the social structure is largely fixed around the rate of social interaction between users, as seen in the comparison of Section 5.2.5 on Network Evolution with that of 5.2.2 on Active user Reply rates and 5.2.4 on Feedback strength. The graph density, giving a measure of the overall connectedness, clustering coefficient, indicating how tightly nodes are clustered into groups, and average degree, describing the graph's overall interconnectedness, are the primary metrics illustrating the network structure. All of these metrics were observed to only significantly change with the Active user Reply rate, and Feedback strength, although increasing Transient user Reply rate resulted in a slight increase as well.

The above findings further highlight the importance of an active core set of participants for citizen science communities, in contrast to the significant role that Transient users play in terms of overall contribution levels, as found in Part 1 of this project. Furthermore, observations on the simulated social network's evolution over time conclude that while the graph grows over time, the fundamental structure of the graph is constant, and this growth appears to stabilise as a project reaches maturity.

Chapter 6

Evaluation

The evaluation chapter focuses on both the validity and quality of findings produced by the extended model, as well the implications of these findings. In addition, the success of this project in fulfilling the goals laid out in Part 1 of this project is assessed.

6.1 Validity

A major challenge in implementing any largely theoretical model is ensuring that it produces results that are accurate, valid, and significant in some regard. This challenge has been tackled wherever possible, starting from the theoretical background laid out in Sections 2.1 and 2.2 exploring the nature and patterns of social participation in citizen science. The base model itself, implemented in Part 1, was found to hold empirical validity in exploring the micro- and macro-level behaviour of citizen science social machines in terms of individual user characteristics and emergent system properties. Implementation of the extensions for studying social interaction utilised empirical data from existing research both in driving design decisions and in establishing baseline parameters for testing.

A gap in existing research regarding social network analysis as applied to citizen science communities presented both an opportunity for a novel project, and a challenge in lack of points of comparison. Despite a lack of solid empirical basis for establishing baseline parameters, the experimentation performed on the extended model revealed results consistent with observations from related research and theory. The simulated social networks were successful in replicating commonly observed characteristics of social networks, and in displaying interactions between model parameters and resultant social network analysis metrics in a manner explainable by theory.

6.1.1 Generalisations in the Model

As discussed in Chapters 3 and 4, some generalisations had to be made in implementing the extensions to the model. These generalisations include the way in which User agents find other Users to form social connections with; it was assumed that agents do so in a perfectly random manner. However, this generalisation is justifiable as the subject of the model is a virtual citizen science platform. In Section 2.2, the challenges of participant interaction and community formation in virtual settings were discussed, and it is reasonable to assume that in both a virtual setting and in the model produced, socially active users are willing to interact with all other users present on the platform.

Despite this generalisation, the results produced by the model were consistent, reaffirming its success. This also leaves the potential for future development of the model in incorporating more sophisticated methods for social interaction to involve the building of relationships between agents.

6.2 Completion of Goals

The first goal for extension stated in Part 1 of the project was to perform robustness testing on the model and thus analyse whether citizen science social machines can fulfil their function despite irrationality or error by users. This goal was not completed; the problem of robustness is already addressed in the design inherent to citizen science projects, where a consensus is collected from the contributions of several participants to account for error and bias. Furthermore, while there may be some value in performing robustness analysis on citizen science to guide future design decisions, there is little related research on the subject to aid in building a successful model for robustness testing.

The second goal for extension was to integrate participant motivations and their interaction with incentives into the model. The goal was originally identified due to the great importance of motivation in citizen science participation (as seen in Section 2.1). However, during the process of reviewing existing research, it became apparent that individual motivations are incredibly complex, shaped by personal experiences and backgrounds, and dynamic over time. A strong link was found to social interaction as both a form of motivation, and the result of motivation to participate in a project. The goal of Part 2 thus shifted towards studying the significance of social interactions within citizen science, and its links to contribution rates.

This change in project goal is justified by allowing for a feasible modelling task while addressing an issue relevant to citizen science, towards which more research has been called for. In fulfilling this goal, the project was successful; the model was extended to allow for the thorough analysis of social network formation in citizen science communities.

Chapter 7

Conclusions

7.1 Contributions

The primary contribution of this project is the extension of an empirically founded agent-based model to study the characteristics and formation of social networks in citizen science. A wide range of recent research in Chapter 2 was used as the backbone upon which the extended model's functions were built. Drawing from concepts of social network analysis identified in Section 4.1, data processing and analysis methods relevant for testing social networks were selected. The combination of agent-based modelling with social network analysis was shown to form a valid platform for studying the significance of social interactions within citizen science.

The completed model provides functionality both in characterising individual participation and social participation, and thus provides a powerful tool for analysing the behaviour of citizen science social machines.

7.2 Wider Considerations

This project extended on the concept of using agent-based modelling techniques to test characteristics of social machines, and came to a successful completion. A novel and theoretically-based approach was taken in studying social interactions in citizen science. The findings produced through testing shed light on the formation of social networks. Firstly, the role of persistent active users in ensuring a healthy community was once again reaffirmed, in this instance in the light of the social community. Secondly, the emergence of the socially active core users in virtual communities was successfully modelled. Thirdly, the findings suggest that the core structure of social networks in citizen science is fixed based on the activity of active users, and that the levels of activity observed stabilise over the community's evolution.

In conclusion, this project produced findings that produced provide explanatory power in how participants in citizen science interact, form social networks, and evolve in their role in the social landscape. The relevance of agent-based modelling and social network analysis in studying citizen science were shown, while driving forward the emerging research on the role of social participation in citizen science. There is yet room for future research into both agent-based modelling as a tool, incorporating a complex variety of motivations in studying user behaviour, and in further exploring the social behaviour of citizen science participants.

Bibliography

Angeletou, S., Rowe, M., & Alani, H. (2011, October). Modelling and analysis of user behaviour in online communities. In *International Semantic Web Conference* (pp. 35-50). Springer Berlin Heidelberg.

Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, *2008*(10), P10008.

Clauset, A., Newman, M. E., & Moore, C. (2004). Finding community structure in very large networks. *Physical review E*, *70*(6), 066111.

Crain, R., Cooper, C., & Dickinson, J. L. (2014). Citizen science: a tool for integrating studies of human and natural systems. *Annual Review of Environment and Resources*, *39*, 641-665.

Devictor, V., Whittaker, R. J., & Beltrame, C. (2010). Beyond scarcity: citizen science programmes as useful tools for conservation biogeography. *Diversity and distributions*, *16*(3), 354-362.

Eveleigh, A., Jennett, C., Blandford, A., Brohan, P., & Cox, A. L. (2014, April). Designing for dabblers and deterring drop-outs in citizen science. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (pp. 2985-2994). ACM.

Gephi (Version 0.9.1) [Computer software]. (2016). Retrieved from https://gephi.org.

Greenhill, A., Holmes, K., Woodcock, J., Lintott, C., Simmons, B. D., Graham, G., ... & Masters, K. (2016). Playing with science: Exploring how game activity motivates users participation on an online citizen science platform. *Aslib Journal of Information Management*, *68*(3), 306-325.

Hamill, L., & Gilbert, N. (2010). Simulating large social networks in agent-based models: A social circle model. *Emergence: Complexity and Organization*, *12*(4), 78.

Handcock, M. S., & Gile, K. J. (2010). Modeling social networks from sampled data. *The Annals of Applied Statistics*, *4*(1), 5.

Hu, Y. (2005). Efficient, high-quality force-directed graph drawing. *Mathematica Journal*, *10*(1), 37-71.

Jackson, C. B., Crowston, K., Mugar, G., & Østerlund, C. (2016, November). Guess what! You're the First to See this Event: Increasing Contribution to Online Production Communities. In *Proceedings of the 19th International Conference on Supporting Group Work* (pp. 171-179). ACM.

Jackson, C., Østerlund, C., Crowston, K., Mugar, G., & Hassman, K. D. (2014). Motivations for sustained participation in citizen science: Case studies on the role of talk. In *17th ACM Conference on Computer Supported Cooperative Work & Social Computing*.

Jackson, C., Østerlund, C., Maidel, V., Crowston, K., & Mugar, G. (2016, February). Which Way Did They Go?: Newcomer Movement through the Zooniverse. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (pp. 624-635). ACM.

Kamar, E., Horvitz, E., Bowyer, A., & Miller, G. (2016, April). Intervention strategies for increasing engagement in crowdsourcing: Platform, predictions, and experiments.

Katmada, A., Satsiou, A., & Kompatsiaris, I. (2016, September). Incentive mechanisms for crowdsourcing platforms. In *International Conference on Internet Science* (pp. 3-18). Springer International Publishing.

Land-Zandstra, A. M., Devilee, J. L., Snik, F., Buurmeijer, F., & van den Broek, J. M. (2016). Citizen science on a smartphone: Participants' motivations and learning. *Public Understanding of Science*, *25*(1), 45-60.

Luczak-Roesch, M., Tinati, R., Simperl, E., Van Kleek, M., Shadbolt, N., & Simpson, R. J. (2014, June). Why Won't Aliens Talk to Us? Content and Community Dynamics in Online Citizen Science. In *ICWSM*.

Marsella, S. C., Pynadath, D. V., & Read, S. J. (2004). PsychSim: Agent-based modeling of social interactions and influence. In *Proceedings of the international conference on cognitive modeling* (Vol. 36, pp. 243-248).

Moss, S., & Edmonds, B. (2005). Sociology and Simulation: Statistical and Qualitative Cross-Validation 1. *American journal of sociology*, *110*(4), 1095-1131.

Mugar, G., Østerlund, C., Jackson, C. B., & Crowston, K. (2015, June). Being present in online communities: learning in citizen science. In *Proceedings of the 7th International Conference on Communities and Technologies* (pp. 129-138). ACM.

Murray-Rust, D., Tarte, S., Hartswood, M., & Green, O. (2015, May). On wayfaring in social machines. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 1143-1148). ACM.

Nov, O., Arazy, O., & Anderson, D. (2011, February). Dusting for science: motivation and participation of digital citizen science volunteers. In *Proceedings of the 2011 iConference* (pp. 68-74). ACM.

Nov, O., Arazy, O., & Anderson, D. (2011, July). Technology-Mediated Citizen Science Participation: A Motivational Model. In *ICWSM*.

Nov, O., Arazy, O., & Anderson, D. (2014). Scientists@ Home: what drives the quantity and quality of online citizen science participation?. *PloS one*, *9*(4), e90375.

Oberhauser, K., & LeBuhn, G. (2012). Insects and plants: engaging undergraduates in authentic research through citizen science. *Frontiers in Ecology and the Environment*, *10*(6), 318-320.

Price, C. A., & Lee, H. S. (2013). Changes in participants' scientific attitudes and epistemological beliefs during an astronomical citizen science project. *Journal of Research in Science Teaching*, *50*(7), 773-801.

Reed, J., Raddick, M. J., Lardner, A., & Carney, K. (2013, January). An exploratory factor analysis of motivations for participating in Zooniverse, a collection of virtual citizen science projects. In *System Sciences (HICSS), 2013 46th Hawaii International Conference on* (pp. 610-619). IEEE.

Rotman, D., Preece, J., Hammock, J., Procita, K., Hansen, D., Parr, C., ... & Jacobs, D. (2012, February). Dynamic changes in motivation in collaborative citizen-science projects. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* (pp. 217-226). ACM.

Roy, H.E., Pocock, M.J.O., Preston, C.D., Roy, D.B., Savage, J., Tweddle, J.C. & Robinson, L.D. (2012) Understanding Citizen Science & Environmental Monitoring. Final Report on behalf of UK-EOF.NERC Centre for Ecology & Hydrology and Natural History Museum.

Stafford, R., Hart, A. G., Collins, L., Kirkhope, C. L., Williams, R. L., Rees, S. G., ... & Goodenough, A. E. (2010). Eu-social science: the role of internet social networks in the collection of bee biodiversity data. *PloS one*, *5*(12), e14381.

Tammela, R. (2016). Testing Citizen Science: An Agent-Based Approach. Undergraduate dissertation.

Tinati, R., Luczak-Roesch, M., Simperl, E., & Hall, W. (2016, May). Because science is awesome: studying participation in a citizen science game. In *Proceedings of the 8th ACM Conference on Web Science* (pp. 45-54). ACM.

Tinati, R., Simperl, E., Luczak-Roesch, M., Van Kleek, M., & Shadbolt, N. (2014). Collective Intelligence in Citizen Science--A Study of Performers and Talkers. *arXiv preprint arXiv:1406.7551*.

Wiggins, A., & Crowston, K. (2011, January). From conservation to crowdsourcing: A typology of citizen science. In *System Sciences (HICSS), 2011 44th Hawaii international conference on* (pp. 1-10). IEEE.

Wiggins, A., & Crowston, K. (2012, January). Goals and tasks: Two typologies of citizen science projects. In *System Science (HICSS), 2012 45th Hawaii International Conference on* (pp. 3426-3435). IEEE.

Wiggins, A., & Crowston, K. (2014). Surveying the citizen science landscape. *First Monday*, *20*(1).

Østerlund, C., Mugar, G., Jackson, C. B., Hassman, K. D., & Crowston, K. (2014, January). Socializing the Crowd: Learning to talk in citizen science. In *Academy of Management Annual Meeting, OCIS Division, Philadelphia, PA*.