Collaborative Research: Focusing Attention to Improve the Performance of Citizen Science Systems - Beautiful Images and Perceptive Observers

Dealing with the flood of digital data that confronts researchers is the fundamental challenge of 21st century research (Hey et al., 2009). New techniques, tools and strategies for dealing with massive data sets, whether they consist of vast numbers of base-pair sequences or terabytes of data from all-sky astronomical surveys, present an opportunity to establish a 'fourth paradigm' of scientific discovery, but the task is not easy. In many areas of research, the relentless growth of data sets has led to the adoption of increasingly automated and unsupervised methods of classification. In many cases, this has led to degradation in classification quality, with machine learning and computer vision unable to replicate the successes of human pattern recognition. The growth of citizen science on the web has provided a temporary solution to this problem; in particular, the highly successful Galaxy Zoo (Lintott et al. 2008, 2011) and the Zooniverse projects (Smith et al. 2011, Fischer et al. 2011, Davis et al. 2011), which have grown from it and which this proposal takes as its starting point, have demonstrated that it is possible to recruit hundreds of thousands of volunteers to make an authentic contribution to results, boosting human analysis through the collective wisdom of a crowd of classifiers. However, human classifiers alone will not be able to cope with expected flood of data from future scientific instruments.

The goal of the proposed project is to develop a next-generation socio-computational citizen science platform that combines the efforts of human classifiers with those of computational systems to maximize the efficiency with which human attention can be used. We recognize that to do so requires a thorough understanding of human motivation and learning in this context, and knowledge of how the proposed system will affect these. The proposed research will be carried out by a partnership between computer and social scientists, addressing research problems both in automated data analysis and social science through systems implementation, alongside field research and experiments with project participants.

Expected intellectual merit. The intellectual merit of this proposal lies in its contribution to advancing knowledge and understanding in multiple domains of science. First, the work will contribute to developing new methods of computational data analysis, initially with analysis of astronomical images, which also contributes to astronomy, and later extending to additional fields. Second, the project includes social science research to test and apply theories of human motivation and learning in an online context, which can then be applied to a broad range of social-computational problems. The research team brings a track record of success in creating, sustaining and studying citizen science projects in multiple fields. By mixing human and computational elements, the proposed system has the potential to transform the application of citizen science and its approach to data analysis.

Expected broader impacts. This project will advance science while promoting teaching, training and learning. One of the most significant broader impacts for all Zooniverse activities is enabling a community of hundreds of thousands of volunteers to participate in research, a powerful and rapidly developing form of informal science education. In addition, Zooniverse projects are widely used in classroom environments and the development of our existing tool for sharing formal teaching resources, ‘Zoo Teach’, will promote the use of these projects in the classroom. The inclusion of McNair fellows will help broaden participation of underrepresented groups. By choosing the relatively generic topic of image classification, we expect that the techniques developed under this grant will be of significant value to future investigations in similar research areas (as already demonstrated by the Zooniverse team), thus enhancing the infrastructure for research and education. In addition to scholarly publications, our projects receive extensive press coverage, providing opportunities for project teams to communicate their research to a wide audience, addressing dissemination to enhance scientific and technological understanding. Finally, the findings from the various projects supported by these tools may provide a benefit to society more generally; for example, citizen science was explicitly mentioned as a benefit to society in the National Academy’s recent review of astronomy and astrophysics.

Collaborative Research: Focusing Attention to Improve the Performance of Citizen Science Systems: Beautiful Images and Perceptive Observers

Dealing with the flood of digital data that confronts researchers is the fundamental challenge of 21st century research (Hey et al., 2009). New techniques, tools and strategies for dealing with massive data sets, whether they consist of vast numbers of base-pair sequences or terabytes of data gathered by astronomical surveys, present an opportunity to establish a 'fourth paradigm' of scientific discovery, but the task is not easy. In many areas of research, the relentless growth of data sets has led to the adoption of increasingly automated and unsupervised methods of classification. In many cases, this has led to degradation in classification quality, with machine learning and computer vision unable to replicate the successes of human pattern recognition, particularly when boosted by the collective wisdom of a crowd of classifiers. The growth of citizen science on the web has provided a temporary solution to this problem; in particular, the highly successful Galaxy Zoo (Lintott et al. 2008,2011) and Zooniverse (Fortson et al., 2011; Smith et al. 2011, Fischer et al. 2011, Davis et al. 2011) projects, which this proposal takes as its starting point, have demonstrated that it is possible to recruit hundreds of thousands of volunteers to make an authentic contribution to results.

At present, the Zooniverse projects in general and Galaxy Zoo in particular can be seen as an implementation of a common simple socio-computational design pattern, in which the computational system passes tasks to one of many randomly selected distributed human participants and aggregates the results which are then utilized by researchers. This pattern can be seen in systems such as the reCaptcha digitization program (van Ahn et al.), in researchers' use of Amazon Web Service's Mechanical Turk service, and even in more complicated citizen science games such as Fold.it (Cooper et al. 2010).

These successful programs have managed to reach the scale of analysis demanded by current data sets, with Galaxy Zoo, for example, recruiting more than 300,000 members of the public as "perceptive observers" to provide more than 180 million classifications of galaxies. Together, they achieve an accuracy comparable to that of professional astronomers, outperforming the best machine-learning solutions and making citizen science a powerful tool for modern research. The Zooniverse team have been successful in expanding this initial success to projects in fields as diverse as astrophysics, ecology and papyrology, with the resulting data has been used in more than 25 peer-reviewed publications (http://www.galaxyzoo.org/published_papers).

This solution is, however, clearly inadequate when faced with the next generation of surveys. While Galaxy Zoo today deals with a decade's worth of data from the Sloan Digital Sky Survey, amounting to approximately ~60TB, the next generation of survey telescope, the Large Synoptic Survey Telescope (LSST) will produce 30 TB of data *every night,* and these numbers themselves pale in comparison to projected data rates from future radio telescopes whose data volumes might exceed petabytes per day. The data challenges faced by astronomers today will confront scientists in many other fields tomorrow, and so a new generation of intelligent citizen science projects which are smarter about allocating tasks and more sophisticated in combining human and machine classification are necessary. Rather than expanding existing projects, we must redesign citizen science to maximize the use of human attention, and to take full advantage of collaboration with machine classifiers in a truly socio-computational system.

The development of such a system is the goal of this proposal, with the design and execution informed by our team's cross-disciplinary experience in the study of volunteer motivation, learning and behavior and in building, launching and developing the world's largest and most successful suite of citizen science projects. Without these developments, the gains made by the citizen science projects which are currently capable of providing essentially complete classifications will prove to have been a short-lived and chimerical solution to the data flood. The resulting system will be able to take advantage of the differing abilities of human and machine classifiers. While automated routines excel at tasks where parameter definitions are well understood (such as color and size), it is difficult if not practically impossible to build routines capable of classifying diverse and disparate data sets. In particular, human classifiers are capable of making serendipitous discoveries even while carrying out routine classifications'.

This means that objects such as Hanny's Voorwerp—shown left (Lintott et al. 2009) were missed by automated tools but easily identified as unusual by a member of the Galaxy Zoo community. In addition to improving the overall classifications, a combined socio-computational approach thus leaves open a space for serendipitous discovery in even the largest data sets.

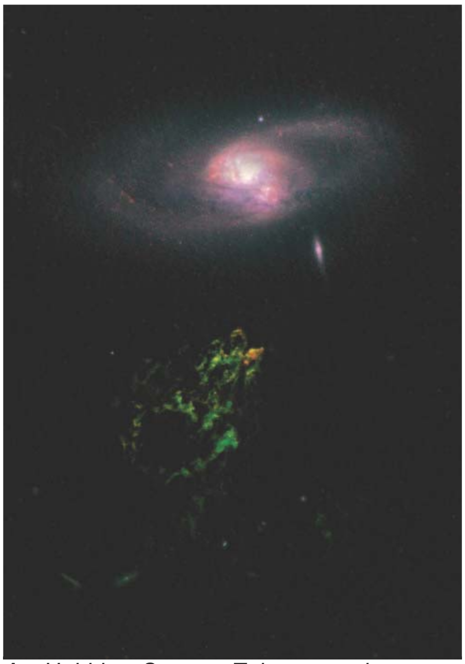

A Hubble Space Telescope image of Hanny's Voorwerp (the green cloud at the bottom of the image), an object discovered (& named) by citizen scientists in Galaxy Zoo and now the subject of detailed study with the world's largest telescopes.

Our approach builds on, but is fundamentally more advanced than existing collaborations between machine and citizen scientist, in which human and automated classifiers are siloed, working independently of each other or combined only in series. The Galaxy Zoo: Supernova project (Smith et al. 2010), for example, allows human classifiers to refine a list of candidate transients identified by an automatic image subtraction and analysis routine. Citizen science projects, with their inherently rich data sets containing not only classification but also estimates of difficulty, are also a source of large training sets for machine learning approaches to problems (Banerji et al. 2009). We aim to make use of both of these modes of collaboration, implemented in a live classification environment and informed by a better understanding of how to influence human behavior.

Real time social computing allows for a richer interaction between human and machine classifiers, as data flowing from an experiment will be examined and passed to either a machine or human classifier for analysis, depending on the properties of the image, their relative capabilities and, in the case of humans, our understanding of their motivation and ability to learn and improve. Results from both types of classifier will inform what is sent to the other, leading to a flexible and continuously improving classification system. This real time mode of operation is explicitly designed for future experiments, where classification will be necessary even before the decision is made to store or discard interesting data. It also allows the distribution of effort allocated to human classifiers to respond to changes in demand, using increased machine effort to produce consistent overall classifications when a citizen science project is popular (at 6pm, for example) and when it isn't.

In this proposal we outline a strategy for developing this hybrid system, a next-generation socio-computational citizen science platform; one that develops a better understanding of individual human classifier performance and combines the efforts of human classifiers with those of computational systems to maximize the efficiency of deployment of human attention. This proposal represents a collaboration between the world-leaders in the development and deployment of massively collaborative projects for academic research (the Zooniverse team), and a research group with the background and expertise to understand the participants. In combining efforts, we believe we can build a citizen science system capable of dealing with the challenges of the next ten years and beyond.

1.1 Background: Citizen science for image classification

Citizen science has been used to solve a variety of problems produced by large data sets, but we concentrate in this proposal on the generic problem of image identification and classification. There are four primary reasons for choosing this type of problem:

1. Image classification is a task that humans excel at; even the most sophisticated morphological classification computer algorithms in astronomy and the biosciences are inferior to a human classification. As the methods used are also relatively generic (Banerji et al. 2009), any approaches developed are easily shared between research areas.
2. For many image classification problems, we have a way to determine the correctness of a classification, when received, that allows us to judge the performance of an individual classifier. We will design projects that have an inherently measurable concept of a 'good' classification.
3. Thirdly, images often have meta-data or easily computable properties that makes it possible to sort them into subcategories for which the level of difficulty can be estimated, as discussed below.
4. Finally, generic image classification is intrinsically difficult, which suggests that applying both human and computational methods together will produce a better result than either on their own. A typical dataset will find that generic computational methods may account for the majority of the required classifications (once provided with sufficient training sets), but human input is almost always necessary to refine the classification scheme, to resolve difficult cases, and to identify the truly unusual and to distinguish these images from artifacts of the process. In other words, the two methods are can be profitably deployed together to exploit adaptive approaches in which machine algorithm performance and human expertise is improved symbiotically thus producing a highly efficient image classification tool, a truly socio-computational system.

These qualities apply across any image classification project, drawn from any field of research. The Zooniverse currently runs a competition for academics to submit proposals three times a year, and this has been successful in drawing a wide range of proposals, with the last round producing more than twenty proposals that passed scientific and technical peer review. The project team will thus be able to select projects that are especially suited for use in this research, whether because they present interesting and effective ways of studying motivation or behavior, or because their large data sets require the advanced tools we will develop. Projects already under development which would be a good fit include a search for small scale events seen in Solar Dynamic Observatory images; this satellite generates 8 megabytes a second of imaging data, most of which is easily available for only a short period of time. The kind of real-time analysis described in this proposal will thus be of vital interest to solar physicists, and the SDO citizen science project will provide the material for development and study in the unlikely event that no other suitable proposals are submitted.

1.2 Problem statement: Next-generation tools for citizen science

The overall objective of this project is to develop an environment where the efforts of human classifiers are combined with automated methods in the most effective manner possible. To do so requires not only a rethinking of the contribution of humans and computational agents but also a new generation of software to support an environment. Furthermore, it is impossible to successfully implement such a strategy without a thorough understanding of human motivation for participation and their potential for learning.

A key aspect of the proposed system will be continual modeling of the performance of individual human volunteer classifiers to inform task assignment (such modeling is feasible as every volunteer logs in with a unique ID to use the system). Volunteer classifiers with different abilities may be better suited to different tasks, with more experienced classifiers, for example, being given the most challenging tasks while their less experienced colleagues encounter a more representative sample. (It should be noted that experts are not necessarily better at all aspects of the task: exploratory work with an eye-tracking system suggests that novice users are most likely to pay attention to the entire image presented to them, making them more likely to make spontaneous discoveries of background features.) As described below, it is possible to set up metrics for performance that, whether based on responses to data with known answers or through measures of consistency, can measure the quality of a classifier, whether machine or human. For example, synthetic planets are periodically injected into the Planet Hunters (http://www.planethunters.org/) tasks in order to measure classifier sensitivity.

A naive approach to optimizing human attention would be to assume that the quality of an individual classifier is inherent to them, and thus constant, just like that of a static computational routine. However, human perceptions and abilities change: the very act of engaging with a citizen science project means that their behavior evolves. Volunteers learn, both from experience and more directly from training materials deliberately incorporated into their workflow, which might, for example, correct them when they make an error or which might present progressively easier or harder tasks.

Such techniques can be used to optimize the work assigned to each classifier, and the degree to which they switch from classifying to learning and back. However, without an understanding of motivation, an approach which optimizes classification strategy for the most accurate and fastest outcome may have unintended and damaging consequences. To give an example, one might build a system which assigns to the most accurate volunteer online at the time the hardest images available. However, in many cases, these images will be those distorted by artifacts, or will be otherwise unrepresentative of the whole sample, leading to a change in the volunteer experience While some degree of challenge likely increases motivation, a continual stream of images which are boring or that take a long time to classify may drive away volunteers; in this case, this will result in a continual decay in the quality of classifiers participating.

The problem is thus not only to build a technical infrastructure to support machine and volunteer contribution, but one which is capable of deploying an understanding of motivation and learning of classifiers in order to retain their interest and improve their capabilities. We plan to meet these challenges by addressing with this proposal the following four specific and interconnected goals:

Goal 1: Develop a new platform for humans and machines to work in concert. Work has already shown that data from the Galaxy Zoo community (Banerji et al. 2009) can be used to significantly improve the performance of automated routines. However, this is an ‘offline’ approach to classification, i.e., the results collected from the community can improve automated methods once a full classification of the dataset has been provided through human effort. Our aim is to combine the efforts of human and machine classifier agents in *real time* to create a system capable of processing the majority of the data by automated agents asking for help from a human classifier only when necessary. However, a hybrid system that can combine humans and automated classifier agents requires a complete redesign of our citizen science platform; one that can scale to support large number of simultaneous classifiers, support complex models of user behavior and provide programmatic interfaces that allow algorithms to ‘plug in’ to the architecture of citizen science.

Goal 2: Be more efficient with the human attention that we already receive. Currently all Zooniverse volunteer classifiers are treated equally when using the citizen science classification interface. Each are equally likely to be assigned a particular task, but yet later, in the analysis of the results, models are developed to assess the performance of the individual volunteer. A more efficient approach would be to develop a model of a volunteer’s performance in real time and to use this guide a selection of what tasks are most suitable. To achieve this goal will require systems development effort, mindful of a definition of success that must be defined not only in efficient classification but in extended and deepening engagement of volunteers.

Goal 3: Improve human performance. All Zooniverse projects include a tutorial that is designed to help volunteers understand the main analysis task. However, currently no effort is made at a later point to measure or improve the accuracy of the analyses provided through intervention and continual training. By asking volunteers to occasionally classify an object drawn from a training set, with the correct answer known, the performance of the individual can be continually measured and if necessary extra training introduced. The project plan again includes experiments with different approaches to training to identify effective approaches to improving performance.

Goal 4: Increase the number of contributors and their engagement. To be successful, systems such as Galaxy Zoo must provide a satisfying experience for volunteers to maintain their involvement. This problem is one faced by all projects that rely on volunteers, both online (e.g., free/open source software development (FOSS) and Wikipedia) and off (e.g., the Boy Scouts and other voluntary organizations, cf. Pearce, 1993). However, as noted above, changes to the kinds of tasks given to volunteers seem likely to interact with motivations for participation, making it critical to understand motivations in more detail. Furthermore, a significant fraction of the classifications collected come from one-time visitors who classify 10 objects or fewer and then never return. Turning visitors into repeat volunteers and increasing the number of classifications that each individual provides necessarily requires a better understanding of what makes for a satisfying user experience, specifically what factors are significant in turning an initial visitor into a regular contributor. The project plan includes interviews and experiments to explore these motivations.

2 Conceptual development: A framework for next-generation citizen science systems

To achieve the goals articulated above will require both technical and social science advances. This section sets out the overall framework for the research, based on the team's extensive expertise in studying this emerging field, deferring to the following section a detailed discussion of the work to be carried out. As seen in the research plan below, the aim of this proposal is to further develop this understanding so that it can be directly used as part of the design of next generation citizen science systems, creating a more sophisticated platform in which machines and humans can collaborate, while retaining the unique features that ensure that these projects reach the widest possible audience.

2.1 Technical principles for next-generation citizen science

Goals 1 and 2 involve building software to support a new approach to integrating human and machine classifiers in systems like Galaxy Zoo and to make best use of human contributions. This work will build on a strong record of success in building citizen science systems for the Zooniverse. The design goal when developing the Zooniverse software tools has always been to minimize the overhead of building new citizen science projects. A number of core principles have been followed that have helped the team achieve that goal and develop 12 very different projects in under 3 years; these principles thus underlie our future development.

1. *A flexible data model*: the launch of Galaxy Zoo 2 introduced a flexible data model that allowed a number of projects to be served by the same application thus reducing development overhead for new projects and reducing complexity in the Zooniverse web environment.
2. *Code reuse and automated test suites*: code reuse - the extraction of common functionality into shared libraries is promoted where possible within our development teams and leads to a more compact, familiar codebase. Significant use is made of automated test suites that make assertions about code behavior and test sequences such as the classification of objects within an interface or user signup at Zooniverse Home. The tests execute continually as software is being developed thus ensuring that as software is developed, existing functionality is not broken.

3. *A loosely-coupled service-based architecture*. Applications running in the Zooniverse web environment are loosely coupled, that is, there are very few dependencies *between* applications. This means that should a service such as user account registration become unavailable, community members are still able to log in to Zooniverse sites and classify data. This loose coupling of applications is a core principle of modern web-application development.
4. *A scalable hosting environment*: The Zooniverse exploits the flexibility and scalability of cloud computing (Amazon Web Services) to handle both web and database servers. With a web stack designed to auto-scale up and down on demand. In a matter of minutes the Zooniverse infrastructure is able to grow in capacity to serve hundreds of thousands of visitors per hour and then later decrease back down to minimal levels when the load is low. This 'utility computing' model, where cloud-based virtual servers may only be running for a few minutes offers a cost effective approach to operating large-scale web applications.

2.2 Theories for enhancing participant performance and learning

To meet goal 2, we will model users and select the tasks to be assigned to each user based upon certain criteria. Galaxy Zoo already maintains considerable information about each volunteer. To supplement those features for optimizing human attention, we can draw on information about each volunteer's tasks, accuracy on each task, persistence over time, along with a record of the feedback interventions (Kluger & DeNisi 1996, 1998) applied, and what changes in behavior have occurred (either as a result of the interventions, or spontaneously). We conceptualize these system capabilities as similar to the functionality of learning-support systems in tracking user behavior and system interactions with users in order to guide presentation of materials (though not in storing instructional content), and will draw on this literature.

For judging the accuracy of classification we will use a number of approaches: 1) comparison with a small expert training set, 2) comparison with a set of synthetic images or 3) a consistency check with other community classifiers. For example, for galaxy classification we can estimate classifier performance as a function of galaxy size, distance, color etc, or as a function of the properties of the image such as noisiness, background brightness and so on. In the absence of a large expert set, sometimes consistency (approach 3) is all that is available; for example, in automated transcription projects, consistency between two independent machine classifiers is usually taken to be the 'right' answer; for human performance, agreement with the consensus is used. This collaborative weighting (Land et al. 2008, Bamford et al. 2009) allows performance of even those classifiers who do not see any of the expert training data to be assessed.

To meet goal 3—improve human performance—we will experiment with providing training at various points. The first question is about the appropriate kind of training to provide. Since the tasks are all image classification, our current plan is to periodically present prototypical images (Ashby & Maddox 2005) to assist with learning the classification category. A second question is about appropriate points at which to introduce the training. We plan to experiment with different algorithmic evaluations of user models, to determine when best to introduce the training images.

And since feedback is an important supplement to instruction (Hattie & Timperley 2007), we will explore the effects of periodically presenting feedback at various junctures in the volunteer's interaction with the system. Drawing on knowledge about the effects of feedback interventions on performance (e.g., Kluger & DeNisi 1996, 1998), we will investigate the possibilities of providing task-specific feedback in a relatively unobtrusive form—not even always identified or labeled as such—while presenting overall performance feedback at a time and in a space more remote from the task situation.

2.3 Theories for explaining and enhancing contributory participation

Goal 4 is to develop a better understanding of why volunteers contribute to citizen science projects such as Galaxy Zoo and other Zooniverse projects (understanding these motivations is also important for accomplishing goal 2). Indeed, for Galaxy Zoo, contributors must decide for each galaxy whether to continue participation. For this project, we will apply a model of helping behaviors to structure our analysis of contributory participation in general (Crowston & Fagnot, 2008). In work on voluntary organizations, Wilson (2000) describes volunteering as "part of a cluster of helping behaviors, entailing more commitment than spontaneous assistance but narrower in scope than the care provided to family and friends" (p. 215). Drawing on the literature on helping behaviors, we suggest that such behavior results from the satisfaction of three precursor conditions (Schwartz & Howard, 1982):

1. First, an individual must recognize a need in others. This condition, called *attention*, focuses on recognizing situational cues that suggest the need for a helping response. These situational cues vary in salience and seriousness.
2. Next, an individual must have *an impetus* to respond, often arising from a combination of feelings of social obligation and/or responsibility together with a self-perceived capability to respond. The capability to respond arises from the volunteer's resources (Uslaner, 2003) and skills and knowledge relevant to the voluntary role (Wilson, 2000, p. 221).
3. Individuals weigh the obligation and capability of helping against the social and tangible costs of doing so in a phase called *evaluation* (Schwartz & Howard, 1982). Of course, helping may also have benefits to the volunteer. Unlike much of the literature on helping behaviors that has examined crises situations requiring quick decisions, evaluation of volunteering can be done deliberately over time.

Second, we distinguish motives at different individual stages of contribution to projects. Distinguishing different stages of individual contribution acknowledges the common observation that the distribution of contributions to projects is quite skewed, with a few people doing the majority of work, and the majority doing little. In light of this skewed distribution, we distinguish three stages of contributors, which we label initial, sustained and meta-contribution (i.e., contributions that structure and enable further contributions (Bryant, Forte, & Bruckman, 2005)). Preece and Schneiderman (2009) similarly noted a possible progression of participation in online groups from "reader to leader" characterized by different activities and motives at each stage. In the remainder of this section, we use the helping behavior model as an overall framework for synthesizing diverse motives for contribution, but then differentiate motives that are relevant at the different stages.

Non-participant Becomes Initial Contributor

The first and largest group of participants we consider is *initial contributors*. Contributors begin their involvement with a project with an initial contribution. It is important to consider the factors - such as interventions - that affect initial contributors and their contributions, because all contributors must pass through this stage. And because projects need visitors to become contributors to sustain and grow the collaboration. We therefore start by considering the steps in a visitor becoming a contributor, to lay the foundation for the more detailed explanation of our research design that follows.

Attention. According to the literature reviewed above, the first stage in volunteering is becoming aware of the project's need for help. The most basic factor for an initial contribution is simply having heard of the project at all and knowing that contribution is possible. The Galaxy Zoo project presents a clear illustration of this effect: as mentioned above, new visitors to the site increase significantly after every mention of the project in the media. However, potential contributors may also become aware of the project's needs by visiting the project's site before contributing (parallel to reading on Wikipedia or lurking in a discussion group). Indeed, such activities may be viewed a form of legitimate peripheral participation (Lave and Wenger, 1991) that allows prospective participants to learn about the project's needs and how their participation can help before taking the first step. We will therefore consider how to make such learning opportunities available to non-participants and measure the conversion to contributors.

Positive evaluation of contributing. The helping model suggests that potential contributors make an evaluation of the costs and benefits of contributing. For participation in an online project, costs include at least a computer and Internet access. Furthermore, a major cost to participation is the opportunity cost of the time spent doing so. In the context of FOSS, Dahlander and McKelvey (2005) found that the reason most commonly cited by non-contributors was a lack of time. Hertel et al. (2003) found that people more willing to tolerate the time cost of contributing made greater contributions. This observation supports the common-sense suggestion that the interface be designed to minimize the time required for contribution.

We next consider possible benefits to participation. In the case of citizen science projects, outcomes rarely include direct monetary or material benefit, but prior research has suggested a number of non-monetary benefits. We review these below when we discuss sustained contribution, but note that few of these seem likely to apply to an initial contributor who is not familiar with the project or with other contributors. Bryant et al. (2005) note that initial users of Wikipedia were often curious about claim that they could just edit a page, so we suggest instead that the benefit to initial participation is simply satisfaction of curiosity about the project. Curiosity has been identified as an important part of intrinsic motivation to use computer system (Malone, 1980). Malone (1980) separated curiosity into two components: sensory curiosity (the attention-attracting value of changes) and cognitive curiosity (the prospect of modifying higher-level cognitive structures) and argued that cognitive curiosity can be incited by indicating discrepancies or paradoxes in a learner's knowledge, which motivate the learner to learn more. Indeed, for Galaxy Zoo, a likely motivation for many is the attraction of the images and curiosity about their nature. Preliminary studies suggest that initial contributors shown a particularly attractive image of a galaxy early in their interactions are more likely to continue their participation; the work included in this proposal will continue to explore this important area, in particular studying whether varying the initial message presented to volunteers alters their sign-up and subsequent behavior. For example, we will study whether adding more information about the project's scientific goals leads more people to contribute.

Initial Contributor Continues Contributing

We next consider factors that might influence an initial contributor to sustain their contribution, thus becoming *sustained contributors*, a second stage in our model of contribution. It is striking that the majority of contributors to do not participate past an initial trial. For example, in the case of Galaxy Zoo, a long tail of contributors classify at most a few galaxies before ceasing participation. To improve the performance of the project, it is important to increase the number of sustained contributors, those in the second stage of our model, who account for the bulk of contributions, despite being a small proportion of the total number of contributors. This portion, therefore, is where the bulk of our intervention experiments will take place, as elaborated below under "Plan of work."

Attention. Again, the first stage in our model is attention. We can assume that sustained contributors are aware of the project from their initial encounters. However, we suggest that to continue contributing, a second important factor is whether the contributor perceives that further contributions are needed by the project. Dahlander and McKelvey (2005) note that the second most cited reason for not contributing to a FOSS project is that there did not seem to be a need, e.g., the software worked well enough or was so specialized that there did not seem to be an opportunity to contribute. Therefore, we suggest that projects should make the need for user contributions visible in order to motivate continued participation.

Impetus to respond. The second stage in the model is the impetus to respond, based on a perceived capacity to respond and a feeling of obligation. Considering the first, perceived capacity, we suggest that feelings of domain knowledge and self-efficacy are important. Enhancing knowledge was addressed in the previous, but we note that it may also improve motivation as well as performance. Turning to feelings of obligation, in contrast to initial contributions, we argue that feelings of social obligation are likely key in decisions to be a sustained contributor (Schroer & Hertel, 2009). To explore motivations for these feelings of obligation, we draw on the literature on social movements, defined by Marshall (1998) as an organized effort by a group of people to effect societal change (Eyerman & Jamison, 1991, p. 4). To the extent that a project can develop the characteristics of a social movement, we suggest that they will be better able to retain and motivate participants. Klandermans' (1997) model of motivations (as augmented by Simon et al., 1998) suggests four distinct areas of motivation for participation in a social movement: collective motives, identification with the group or a subgroup, reward motives and social motives. We consider the first two of these motives in this section and the others in the final two sections.

Collective motivations come from the individual's evaluation of the group's goals or ideology, relevant because many or most social movements coalesce around a shared ideology. For example, Xu et al. (2009) found that stated agreement with a project's values, norms and beliefs was a strong predictor of FOSS developer involvement. In Wikipedia, sustained contributors express feelings of agreement with the project's goals, contributing to the greater good (Bryant et al., 2005). Previous work has found that the most common motivation for Zooniverse volunteers is the desire to make a contribution to science (Raddick et al.). Therefore, we suggest that having a clear project ideology may make it easier for a participant to form an identification with the project, making it more likely s/he is to become a sustained contributor.

Group or community identification means that individuals join a movement because of their feelings of being part of or wanting to contribute to a valued group. The feeling of community identity is essential to the transformation of interests (group or individual) into collective action (Gotham, 1999, p. 19). Group identification differs from social motives in that the latter arise directly from interactions with other people—whether group members or not—while the former is a preferred state of mind based on a sense of belonging. This sense is part of the explanation for the feeling of obligation to the group that provides a motivation for sustained contribution. Wiesenfeld et al. (1999) described organizational identification as "the psychological tie that binds scattered employees together into an organization, rather than a collection of incidentally related individuals" (p. 786). Ellemers, DeGilder and Haslam (2004) stated that when individuals self-identify as part of a collective that they are more inclined to work towards improving the collective and its identity and Johnson et al. (2010) suggested that commitment to the group is an important motivation for work more generally. Such identification has been found across a range of collaborative projects. Bryant et al. (2005) noted that active Wikipedia contributors develop an identity in the project, e.g., by having a Wikipedia home page and use a talk page to interact with others. Ren et al. (2010) claimed that online community members who identify with the group and/or to particular members of the group have a higher commitment and hence continued participation. Therefore, we propose that actions to enable group identification will help bring more sustained contribution to the project.

Positive evaluation of contributing. The third condition in the model is the comparison of costs and benefits of contributing. We considered costs of contributing above, and those propositions hold for sustained contributors as well, though the time cost of contribution may be higher for a sustained contributor, as additional work is expected to meet the standards of the project, and these standards are likely higher for larger or more organized projects. Turning to benefits, we start by considering individual rewards for contribution, which aligns with reward motivations identified by Klandermans' (1997) model of motivations for social movements. For a citizen science project, we do not expect that many participants expect to benefit directly from their participation, and indeed, we find that while 'making a discovery' is a strong secondary motivation for Galaxy Zoo volunteers, the more communal desire to 'make a contribution' dominates. Instead, we consider the sometimes-overlooked motivation of personal satisfaction of contribution. To better understand why making contributions could be intrinsically rewarding for sustained contributors, we leverage Hackman and Oldham's (1980) model of work motivations. Hackman and Oldham identify five job dimensions—skill variety, task identity, task significance, autonomy and feedback—that they suggest create positive psychological states about the work and thus lead to work motivation. Hackman and Oldham note that the first three dimensions lead to experienced meaningfulness of work, while autonomy leads to perceived responsibility for work and feedback to knowledge of work outcomes. Together, those three factors lead to making the work attractive to potential volunteers. We therefore suggest that projects consider how to design tasks that have more skill variety, task identity, task significance and autonomy to increase the likelihood of a participant becoming a sustained contributor to the project. Furthermore, we suggest that if contributors are given positive feedback about their contributions, directly or from other contributors, then they are more likely to become a sustained contributor to the project. The project can also provide feedback about the level of performance and changes over time.

Sustained Contributor Becomes Meta-contributor

Finally, we turn to consideration of factors supporting participants' decision to become meta-contributors, a third stage of the model. We note that in successful projects, a very few contributors, perhaps only 1% of sustained contributors, shift their focus from substantive contributions to what we label "*meta-contributions*," those contributions that structure and enable further contributions (Bryant et al., 2005). For example, within the Zooniverse communities, contributors may curate content within discussion environments, review object collections created by other volunteers in the Zooniverse 'Talk' tool (a custom discussion tool designed to facilitate the discovery of interesting objects) and generally perform higher-level functions within the community environment. We suggest that projects explicitly recognize and reward those who perform meta-contributor roles in such venues. Individuals might be appointed to take on such a position, but the model suggests that making the role one that sustained contributors with suitable qualifications can aspire to may help increase motivation for those individuals.

Attention. As with sustained contributors, we note that becoming a meta-contributor starts with awareness of the project's need for this kind of work. The distinguishing characteristic of meta-contributors is that they are concerned with structure of the whole project, not just a few pieces, and with the state of the community, not just its output (Bryant et al., 2005). We propose that projects make visible their needs for meta-contributions, increase the likelihood that a contributor will become a meta-contributor.

Impetus to respond. Regarding the second stage in the helping process, we suggest that meta-contributors go through much the same evaluation as sustained contributors in determining their capacity to respond. However, rather than domain knowledge and media self-efficacy, meta-contributors must have a good knowledge of the community and its norms and rules. As with sustained contributors, we believe that meta-contributors feel a social obligation to respond based on their adoption of the project's shared ideology, though in their role of meta-contributor, they also help shape this ideology.

Positive evaluation of contributing. The third dimension of the model is the comparison of costs and benefits of contributing. We expect the same low evaluation of costs for meta-contributors as for sustained contributors. Considering benefits, above we hypothesized a set of individual benefits that motivate sustained contribution. While meta-contribution may still be intrinsically motivating, we suggest that individuals rely primarily on social rewards.

3 Plan of work

In this section, we present an initial plan for achieving the four goals articulated above, as informed by the theories and perspectives developed in the previous section. The project is divided into a several phases. The initial phase will address building necessary software infrastructure and carrying out fieldwork to refine the theoretical models presented above. Later phases will implement experiments with specific system features to improve learning, performance and motivation. We anticipate being able to implement 4–6 experiments per year, depending on the scope of the intervention required. Initially these experiments will be performed with Galaxy Zoo; in the final year of the project, additional projects will be added. Necessary criteria for projects include a focus on image analysis, a large volume of data to be classified, a need for live analysis, a way to assess classification performance (e.g., a large set of known images) and a large community of volunteers; these criteria describe many of the Zooniverse projects. Specific projects will be chosen to generalize and contrast with findings from the initial experiments.

3.1 Software development

This project requires the development of a highly sophisticated software framework that fuses human-based citizen science classification interfaces together with automated classification routines. First iterations will be in a 'dual' mode, that is, all data being classified in parallel by both humans and machines. Later software iterations will allow machine learning software to interface directly with the citizen science project with real-time decisions being made about what to pass to human classifiers.

By the end of this project, the software should be able to do the following:

- Filter (and classify) the majority of incoming data.
- Upon inspection of a new data object should choose which classifier should perform the analysis.
- Use some of the human-based classifications to refine automated tools in realtime, through the use of recent classifications as targeted training sets, selected according to the properties of the data being processed.

To achieve these goals this work will be carried out in a number of distinct phases:

Developing programmatic interfaces for automated methods: Combining the classifications of community volunteers with machine learning algorithms will require the development of defined programmatic interfaces (APIs) so that the framework can communicate with automated methods. Approaches will need to be developed for supporting the automatic ingestion of incoming data requiring analysis by citizen scientist. The team already has relevant experience in this space, however further development is likely to include larger volumes of data.

Machine learning software encapsulation: The development of bespoke machine learning approaches for image processing is not within the scope of this project; rather, we will integrate software developed in other projects. By using higher-level languages (such as Python and Ruby) for our web environment we are in a strong position to encapsulate machine learning routines written by academics using a range of other programming languages such as MATLAB or C++.

Live decision algorithm: Decisions will need to be made on the fly as how best to allocate data for analysis. Machine learning methods incorporating 'active learning' will be used to choose not only how data are allocated to classifiers (choosing between people and machines and also between difference algorithms) but also to continually refine the training sets used by automated methods.

Classifier performance and behavior modeling: Approaches will be developed to model the performance and behavior of the human classifiers as well as characteristics of the automated analysis packages. Exactly what factors are modeled during this phase of software development will rely heavily upon the theories developed in early phases of the project.

Design of tools to support controlled experiments: Members of the community will need to be selected accorded to a stratified sampling plan and their behavior tracked when changes are made to the user experience to ensure that they are treated consistently. Volunteers will be required at points to take surveys from within the main citizen science experience. Design of the experimentation environment will be a collaborative effort between both groups, but implementation will be the responsibility of the Smith team.

3.2 User motivation and learning studies

The studies of user motivation and learning will be conducted in two phases, each with multiple subphases.

Develop models of user motivation and learning: During the first year of the project, prior studies of the Zooniverse community will be reviewed and interviews conducted with selected community members to identify key motivations and interests of volunteers and opportunities for enhancing participant learning. These studies will extend and ground the conceptual models developed above, and suggest ways to manipulate these factors through the design of the citizen science project interaction. A main outcome of the first year will be identification of a set of factors that seem useful for improving participant learning and motivation for participation that can be systematically tested in the stage of the project.

Design and execute experiments to test hypotheses regarding motivation and learning: During the second and third years of the project, the tools developed during year 1 to support controlled experiments will be used to investigate motivations of community members and how learning and engagement can best be supported. The system will enable us to assign a randomly-selected set of participants to different conditions (e.g., providing task performance feedback) and to collect performance data to compare outcome (e.g., whether participants who receive the intervention perform better or stay with the project longer). The website design will enable us to hold constant other aspects of the interaction. The large number of project participants will enable us to draw samples large enough to ensure sufficient power for the experiments.

We plan to start some longer running experiments at the start of the year, while adding additional shorter experiments throughout the year (the limiting factor will be how quickly the necessary interface changes can be implemented). Data from the planned surveys will support additional tests of the theoretical models. We plan to carry out experiments and surveys with one Zooniverse project initially, then expand to additional projects in the final year of the project to explore the implications of various task types. Successful interventions can quickly be implemented for all participants.

4. Project management plan

In this section we present our plans for bringing the proposed project to a successful completion. Just as the proposed system is a partnership between human participants and a system, the project will be carried out by a partnership between social and computer scientists at two institutions. Based on a preliminary assessment of the effort required, we are requesting funding for an interdisciplinary team comprising two PIs (partial support), one in computer science and one in the domain of information systems and organizational behavior, partial support for a post-doctoral fellow at Syracuse and one full-time postdoctoral fellow/developer at Adler. In addition, we plan to recruit two undergraduate students in Chicago from the computational data analysis program, specifically for analysis of astronomical images, and two at Syracuse from the Ronald E. McNair Post-Baccalaureate Achievement Program[1]. As well, the Syracuse University School of Information Studies has a program that employs masters students to work on research projects; these students can be used to supplement the team.

---

[1] The McNair program is a federally-funded TRIO program, designed to prepare students for graduate education leading to doctoral studies. Undergraduate students eligible for the program must either be a potential first generation graduate student or be a member of a group underrepresented in graduate education.

The PIs, Arfon Smith and Kevin Crowston, will be funded during the summer on project management and research design. The PI will devote effort during the academic year to project management and oversight. Both PIs will share in project selection, overall project design and report writing. Each PI has specific responsibilities for the execution of the project:

PI Smith and his postdoctoral researcher/developer will lead the software development effort required for this work (roughly goals 1 & 2, though supporting and drawing on research on goals 3 & 4). Initially this work will require solutions to be designed and implemented for real-time modeling of the individual and to incorporate machine learning within the project environment.

Co-PI Crowston and his postdoctoral researcher will lead the social science research. They will be responsible for exploring existing data collected by the Zooniverse and for idea and theory development that will inform which factors are manipulated in the citizen science platform (roughly goals 3 & 4, though also informing work done on the other goals). Later his group will be responsible for designing experiments to test theories developed around user behavior and will work in close collaboration with the Zooniverse developer to implement their ideas. The post-doctoral associate at Syracuse will support co-PI Crowston in sample selection, definition of constructs and variables, and data collection and analysis. Co-PI Crowston and the post-doc will also act as faculty research mentors for two McNair undergraduate students. Their tasks will depend on their interests and capabilities, but we anticipate having them work on data collection and analysis. While student stipends are provided by the McNair program, travel funds are requested for these students to attend one conference per year, with the goal of presenting a research poster.

This timeline for the proposed project. For management and tracking, we have divided the work to be done into quarters. An initial project activity will be the development of a more detailed timeline against which progress will be measured. The budget includes support during summer to support these activities.

We will employ two main *project management techniques*. First, we will have regular meetings of the project members by conference call to share findings and to plan the work. Initially, these will be every other week, but the frequency of meetings will be adjusted depending on our experience and the pace of the work being carried out at the time. These formal meetings of all project participants will augment the regular interaction of the teams of PIs and students working on the implement and data analysis and expected frequent interactions of the students as they analyze data from the same projects. Second, a more formal review meeting will be held at least quarterly to assess progress and to make plans for the next quarter.

We plan to *disseminate our findings* in various ways for multiple purposes. First, we will share our findings with the cooperating citizen science projects and communities; more than 100 academic staff are currently engaged in developing and utilizing Zooniverse proposals and projects, producing a large and diverse audience for our research. Such findings may provide insight on project management and recruiting strategies, as well as training and task design. Second, we will present our findings at relevant national and international conferences for academics and practitioners, again influencing both research and practice. Third, we will send our theoretical findings to academic journals; thus the findings will contribute to the literature with long-lasting impact in related areas of studies. Finally, we will use blogs and live chats on project websites to share results with the citizen scientists contributing to the project.

The project team will be advised throughout by the existing Citizen Science Alliance advisory board, which consists of scientists, citizen science practicioners, educational researchers and developers. Consultation with this group will be primarily through email reports and discussion, with a presentation given by one team member at their annual meeting. The CSA itself is an inherently cross-disciplinary group of museums and universities, whose members will support the work within this proposal.

5. Conclusions

We conclude by reiterating the intellectual merit and expected broader impacts of the proposed project, and by briefly describing results from prior NSF funding.

The *intellectual merit* of this proposal lies in its contribution to advancing knowledge and understanding in multiple domains of science. First, the work will contribute to developing new methods of computational data analysis, initially with analysis of astronomical images, which also contributes to astronomy, and later extending to additional fields. Second, the project includes social science research to test and apply theories of human motivation and learning in an online context, which can then be applied to a broad range of social-computational problems. The research team brings a track record of success in creating, sustaining and studying citizen science projects in multiple fields. By mixing human and computational elements, the proposed system has the potential to transform the application of citizen science and its approach to data analysis. The project includes a work plan and a detailed project management plan to help bring the project to successful completion and to ensure its impact.

We expect the work carried out under this grant to have extensive *broader impacts*. The project will advance science while promoting teaching, training and learning. One of the most significant broader impacts for all Zooniverse activities is enabling a community of hundreds of thousands of volunteers to participate in research, a powerful and rapidly developing form of informal science education. In addition, Zooniverse projects are widely used in classroom environments and the development of a new tool for sharing formal teaching resources ‘Zoo Teach’ will promote the use of these projects in the classroom. The inclusion of McNair fellows will help broaden participation of underrepresented groups. By choosing the relatively generic topic of image classification, we expect that the techniques developed under this grant will be of significant value to future investigations in similar research areas, thus enhancing the infrastructure for research and education. In addition to scholarly publications, our projects receive extensive press coverage, providing opportunities for project teams to communicate their research to a wide audience, addressing dissemination to enhance scientific and technological understanding. Finally, the findings from the various projects supported by these tools may provide a benefit to society more generally; for example, citizen science was explicitly mentioned as a benefit to society in the National Academy's recent review of astronomy and astrophysics.

This proposal makes use of the infrastructure of the Zooniverse network of citizen science projects, which has been supported by three separate and ongoing NSF grants. A CDI grant at the Adler Planetarium (DRL-0941610; $1,889,993; 01/01/2010–12/31/2013; "CDI Type II: Zooniverse - Conquering the Data Flood with a Transformative Partnership between Citizen Scientists and Machines") supports the development of advanced tools for citizen scientists to lead their own research. Basic versions of such tools have been deployed on Galaxy Zoo and Ancient Lives, the latter of which also benefits from an in-post NSF supported papyrologist. The talk discussion tool deployed on Planet Hunters and other Zooniverse projects was also developed under this grant, and has been released as an open source resource for other developers. A more general framework to support advanced tools is now being developed in collaboration with the University of Washington, though the use of these tools is beyond the scope of the current grant. The CDI grant also supports the development of machine learning tools using citizen science data, through the establishment of a series of challenges. Results from this research effort will inform the choice of machine learning tools used in the projects to be developed here.

The second grant, under the ISE program, (DRL-0917608; $926,327; 09/01/2009–08/31/2012; "Investigating Audience Engagement with Citizen Science") funds education research into citizen science motivation. The team, which is divided between Adler and Johns Hopkins have carried out surveys and interviews with planetarium visitors exposed to the website to support the model of volunteer behavior described in the main text, and to inform design decisions. The third Zooniverse grant, under the CI-TEAM program, (OCI-1041419; $249,560; 10/01/2010 - 09/30/2012; "Citizen CI-Team Demonstration Project - Expanding the Zooniverse: Refining Resources to Support Educator Use of Citizen Science in the Classroom") is aimed at developing the use of citizen science in schools. Pilot programs with Chicago Public School teachers are underway, with resources being developed to use Zooniverse projects in a diverse variety of lessons, ranging from science to maths, geography and even for career advice.

Co-PI Crowston has been funded by several NSF grants in the past years. Three are directly relevant to the current proposal. The first, (09–43049; $150,000; 09/01/2009 – 08/31/2011; "VOSS: Theory and design of virtual organizations for citizen science"), was a two-phase theory-based study of virtual organizations for citizen science. The project was directed at advancing the understanding of what constitutes effective citizen science virtual organizations and under what conditions citizen science virtual organizations can enable and enhance scientific and education production and innovation. The study was theoretically grounded in small group theory and rooted empirically in a survey of and case studies in citizen science projects. The second grant (09–68470; $478,858; 09/10/2010 – 09/09/2013; "Socially intelligent computing to support citizen science") and (11–11107, $747,831), started with an analysis of systems that support citizen science projects, including case studies of eBird, Galaxy Zoo and the Great Sunflower project, as well as a workshop with citizen science project leaders to development system requirements. The current work for this grant is development of a set of tools and games to support classification of the species of organisms in pictures submitted to citizen science projects. These will be deployed to test ideas about motivation and satisfaction in citizen science projects. The project supports two doctoral students, 6 masters students and 3 undergraduate students (one through an REU supplement). The current proposal will extend this work to consider factors beyond motivation, while focusing specifically on task assignment. The final project, 09/01/2011–08/31/2014, for "Socially intelligent computing for coding of qualitative data" (with Nancy McCracken), just started, is developing a semi-automated tool that combines human coding, natural language processing and machine learning to support qualitative data analysis. The grant supports two doctoral students; we are currently seeking two McNair Fellows to join.

References

Ahn, L.; Blum, M.; Hopper, N. J.; Langford, J. (2003). CAPTCHA: Using Hard AI Problems for Security *Advances in Cryptology — EUROCRYPT 2003*. Lecture Notes in Computer Science. 2656. pp. 294–311.

Ashby. F. Gregory & W. Todd Maddox (2005) “Human category learning.” In Annual Review of Psychology. Vol. 56, pp. 149-78.

Bamford S., et al. (2009) Galaxy Zoo : The dependence of morphology and colour on environment, MNRAS, 393, 1324

Banerji M., et al., (2010). Galaxy Zoo: Reproducing galaxy morphologies via machine learning, *MNRAS*, 406, 342.

Bryant, S. L., Forte, A., & Bruckman, A. S. (2005). Becoming Wikipedian: Transformation of participation in a collaborative online encyclopedia*.* In *Proceedings of GROUP Conference*, (Sanibel Island, FL, 2005).

Cooper, S.; Khatib, F.; Treuille, A.; Barbero, J.; Lee, J.; Beenen, M.; Leaver-Fay, A.; Baker, D. et al. (2010). Predicting protein structures with a multiplayer online game. Nature 466 (7307): 756–760.

Crowston, K., & Fagnot, I. (2008, 1–2 July). The motivational arc of massive virtual collaboration*.* In *Proceedings of IFIP WG 9.5 Working Conference on Virtuality and Society: Massive Virtual Communities*, Lüneberg, Germany.

Dahlander, L., & McKelvey, M. (2005) Who is not developing open source software? Non-users, users, and developers. *Economics of Innovation and New Technology*.

Ellemers, N., DeGilder, D., & Haslam, S. A. (2004) Motivating individuals and groups at work: A social identity perspective on leadership and group performance. *Academy of Management Review* 29(3): 459–478.

Eyerman, R., & Jamison, A. (1991) *Social movements: A cognitive approach*, University Park, PA: Pennsylvania State University.

Fischer et al., (2011), *Planet Hunters : The first two planet candidates identified by the public using the* Kepler *public archive data,* Accepted by *Monthly Notices of the Royal Astronomical Society*, arXiv/1109.4621

Fortson et al. (2011) *Galaxy Zoo : Morphological Classification and Citizen Science*, to be published in *Advances in Machine Learning and Data Mining for Astronomy*, Ed. M.J. Way

Gotham, K. F. (1999) Political opportunity, community identity and the emergence of a local anti-expressway movement. *Social Problems* 46(3): 332–354.

Hackman, J. R., & Oldham, G. R. (1980) *Work Redesign*, Reading, MA: Addison-Wesley.

Hattie, John & Helen Timperley (2007) “The power of feedback.” Review of Educational Research. Vol. 77, No. 1, pp. 81-112.

Hertel, G., Niedner, S., & Herrmann, S. (2003) Motivation of software developers in Open Source projects: An Internet-based survey of contributors to the Linux kernel. *Research Policy* 32: 1159–1177.

Hey, T., Tansley, S., Tolle, K. (Ed.) (2009) The Fourth Paradigm: Data-Intensive Scientific Discovery, *Microsoft Research*

Howison, J., Inoue, K., & Crowston, K. (2006, June). Social dynamics of free and open source team communications. In *Proceedings of International Conference on Open Source Software*, (Lake Como, Italy, 2006). 319–330.

Johnson, R. E., Chang, C.-H. D., & Yang, L.-Q. (2010) Commitment and motivation at work: The relevance of employee identity and regulatory focus. *Academy of Management Review* 35(2): 226–245.

Klandermans, B. (1997) *The Social Psychology of Protest*, Oxford: Basil Blackwell.

Kluger, Avraham N. & DeNisi, Angelo. (1996) "The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory." Psychological Bulletin. Vol. 119, No. 2, pp. 254-284.

Kluger, Avraham N. & DeNisi, Angelo. (1998) "Feedback interventions: Toward the understanding of a double-edged sword" Current Directions in Psychological Science. Vol. 7, No. 3, pp. 67-72.

Land et al., (2008) Galaxy Zoo : The large-scale spin statistics of spiral galaxies in the Sloan Digitial Sky Survey, MNRAS, 388, 1686

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.

Lintott C.J., et al., (2008), Galaxy Zoo: Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey, *MNRAS*, 389, 1179.

Lintott C.J., et al., (2009), Galaxy Zoo: 'Hanny's Voorwerp', a quasar light echo?, *MNRAS*, 399, 129

Malone, T. W. (1980). *What makes things fun to learn? Heuristics for designing instructional computer games*. Paper presented at the Proceedings of the 3rd ACM SIGSMALL Symposium and the First SIGPC Symposium on Small Systems, New York, NY, USA.

Marshall, G. (Ed.). (1998). *A Dictionary of Sociology*. Oxford: Oxford.

Pearce, J. L. (1993) *Volunteers: The organizational behavior of unpaid workers*: Routledge.

Preece, J., & Shneiderman, B. (2009) The Reader-to-Leader Framework: Motivating Technology-Mediated Social Participation. *AIS Transactions on Human-Computer Interaction* 1(1): 13–32.

Raddick, K. et al., (2009) *Astronomy Education Review*, 9, 1

Ren, Y., Harper, F. M., Drenner, S., Terveen, L., Kiesler, S., Riedl, J., & Kraut, R. E. (2010). Increasing Commitment to Online Communities: Designing from Theory. Minneapolis, MN: University of Minnesota.

Schroer, J., & Hertel, G. (2009) Voluntary engagement in an open web-based encyclopedia: Wikipedians and why they do it. *Media Psychology* 12(1): 96--120, doi: 10.1080/15213260802669466.

Schwartz, S. H., & Howard, J. A. (1982) Helping and cooperation: A self-based motivational model. In V. J. Derlega & J. Grzelak, (eds.), *Cooperation and Helping Behavior: Theories and Research* (pp. 327–353), New York: Academic,pp.327–353.

Simon, B., Loewy, M., Stürmer, S., Weber, U., Freytag, P., Habig, C., . . . Spahlinger, P. (1998) Collective identification and social movement participation. *Journal of Personality and Social Psychology* 74: 646–658.

Smith A. M, et al., (2011), Galaxy Zoo Supernovae, *MNRAS*, 412, 1309.

Uslaner, E. M. (2003). Civic engagement in America: Why people participate in political and social life. College Park, MD: Department of Government and Politics, University of Maryland.

Wiesenfeld, B. M., Raghuram, S., & Garud, R. (1999) Communication patterns as determinants of organizational identification in a virtual organization. *Organization Science* 10(6): 777–790. Available from: http://www.jstor.org/stable/2640241

Wilson, J. (2000) Volunteering. *Annual Review of Sociology* 26: 215–240.

Xu, B., Jones, D. R., & Shao, B. (2009) Volunteers' involvement in online community based software development. *Information & Management* 46(3): 151 – 158, doi: 10.1016/j.im.2008.12.005.