2011 Seventh IEEE International Conference on e-Science Workshops

Reputation-aware Filtering Services for Citizen Science Data

Charles Brooking
School of ITEE
The University of Queensland
Brisbane, Australia
brooking@itee.uq.edu.au

Jane Hunter
School of ITEE
The University of Queensland
Brisbane, Australia
j.hunter@uq.edu.au

***Abstract* - The Internet, Web 2.0 and Social Networking technologies are enabling citizens to actively participate in "citizen science" projects by contributing data to scientific programs. However, the limited expertise of contributors can lead to poor quality or misleading data being submitted. Subsequently, the scientific community often perceive citizen science data as not worthy of being used in serious scientific research. In this paper, we describe how online reputation models can be adapted for citizen science projects to provide a simple and effective mechanism for assessing the reliability of community-generated data. We also describe the reputation-aware querying, filtering and visualization services that we have developed that enable users to distinguish between datasets based on the reputation of the source/contributor. The resulting services are evaluated in the context of the Coral Watch project which uses volunteers to collect data on coral reef bleaching.**

Keywords - citizen science; reputation models, reputation-aware filtering

I. Introduction

The term *citizen scientist* refers to a volunteer who collects and/or processes data to contribute to scientific research. The number and variety of *citizen science* projects has grown dramatically in recent years. Such projects typically combine web-based social networks with community-based information systems to harness collective intelligence and apply it to a particular scientific problem. Online communities of volunteers are now contributing data to projects that range from astronomy [1] to bird watching [2] and air quality [3]. In particular, the issues of climate change and associated environmental impacts are mobilizing individuals who want to contribute to the monitoring, management and maintenance of ecosystem health by capturing observational data. Citizen science data is often essential for assessing and validating predictive models that involve large-scale spatial and temporal extents. The necessary longitudinal datasets are often incomplete and outdated due to resource limitations with regard to funding and personnel availability. Citizen science can also play an important role in reducing costs associated with research projects and the development of more comprehensive data collection. Furthermore, citizen science programs often lead to increased public awareness of environmental and scientific challenges, civic involvement, fulfillment of academic requirements (in the case of students), and improvement in decision making skills [4].

However, there are some inherent weaknesses to citizen science and crowd sourcing projects. The limited training, knowledge and expertise of contributors and their relative anonymity can lead to poor quality, misleading or even malicious data being submitted [5]. The absence of formal "scientific methods" [6] and the use of non-standardized and poorly designed methods of data collection [7] often lead to incomplete or inaccurate data. Also, the lack of commitment from volunteers in collecting field data [4, 5] can lead to gaps in the data across time and space. Subsequently, these issues have caused many in the scientific community to perceive citizen science data as low quality and not worthy of being considered in serious scientific research [8].

In this paper we propose the application of online reputation models to citizen science projects to provide an indication of the reliability of the citizen science data based on its source – thus enabling unreliable data to be excluded via reputation-aware querying, filtering, visualization and reporting services that take into account the reliability of the data, which is explicitly displayed to users.

A. Hypothesis

A large amount of research has been undertaken into approaches to improve data quality. For example, fairly simple techniques can be applied to validate data input (e.g., syntax, format and values) by checking compliance against schemas. More complex data quality assessment may require comparison with data sets from alternative sources or comparison with historical trends. However these approaches are limited if there are no other sources of comparable data or there is no longitudinal data for trend analysis. An alternative and complementary approach to data quality enhancement services is to exploit social network analysis tools to infer the reliability of the data based on the contributor's reputation. A number of reputation models have been developed by researchers in the context of Web 2.0 [12,19] – but to date, none have been applied to citizen science data.

Our hypothesis is that trust and reputation metrics (such as those developed to provide recommender services in online social networks (e.g., eBay, Netflix) can usefully be applied to citizen science data. We believe that reputation

978-0-7695-4598-1/11 $26.00 © 2011 IEEE
DOI 10.1109/eScienceW.2011.25

models can provide a simple and effective mechanism for filtering unreliable data. Moreover, by combining online social trust/reputation metrics with data validation services, the quality and reliability of the community-generated data can be significantly improved – thus enabling its confident re-use by the scientific community.

II. Objectives

The objective of the work described here is to analyse and compare different reputation models and to develop the optimum reputation model for assessing the reliability of citizen science data. More specifically we aim to:

- Survey alternative reputation models and algorithms for measuring reputation and identify those approaches most applicable to citizen science projects;
- Identify a set of criteria or attributes for computing the reputation of a user and their citizen science data e.g.:
  - The contributor's age;
  - The contributor's role and qualifications (primary student, secondary student, PhD student, volunteer, council worker, scientist);
  - The quality and amount of past data that they have contributed;
  - The extent of training programs completed;
  - Frequency and period of contributing;
  - The contributor's ranking from other members (direct, inferred or calculated using social trust algorithms).
- Develop tools for capturing the reputation related attributes above, and for calculating a reputation score within citizen science projects (e.g., the optimum weightings that should be applied to the criteria listed above to determine the most accurate measure of the data's trustworthiness);
- Evaluate, analyze, refine and optimize these reputation algorithms, tools and services – by measuring the improvements in data quality that result from using reputation metrics to filter or remove untrustworthy data or untrusted contributors;
- To investigate and identify the optimum mechanisms for displaying and communicating the reliability of the data and contributors, to other members of the community, especially scientists who are considering re-using the community-generated data.

III. Classes of Reputation Systems

Existing reputation systems can be categorized into:

- content-driven systems that rely on automated content analysis to derive user and content reputations by comparing contributed content with "ground truth". Examples include Wikipedia's WikiTrust and Crowdsensus [14]. User reputations are computed according to the quality and quantity of contributions.
- user-driven reputation systems (such as eBay, Amazon) that rely on explicit user feedback and ratings as well as inferred ratings across networks.

These two different categories of reputation systems present different challenges. Content-driven reputation systems work even if the contributors are anonymous, because it is the data that is being assessed, not the user. Because feedback is derived from all of the data uniformly, there is less bias and they are difficult to manipulate by users. But the "black-box' nature of the automatic algorithms employed, often lead to users not trusting the scores because they don't understand how they were calculated. In addition, these systems rely on the availability of ground truth data.

User-driven reputation systems rely on a constant community of dedicated users who provide high-quality feedback. They also require user registration and user profiles and are open to manipulation by malicious users. Moreover, the feedback may be biased towards the "squeaky wheels" – happy or unhappy contributors who are more likely to provide ratings. The advantages are that user-driven approaches enable personalized views of the data reliability to be generated and they can be used even when no ground truth exists.

In the context of citizen science projects, the choice of reputation model depends on the responses to a number of questions: is ground truth data available?; is there a constant community of dedicated users who have been involved long-term?; are users registered?; are user profiles available?; what QA/QC is performed on the data?; does the system monitor the number or type of corrections performed on contributed data?; do the staff and users know each other? are there ranking tools available for data or users?

In the next section we illustrate via a particular citizen science case study, the process that we have used to identify and optimize a reputation model to filter unreliable data.

IV. Case Study

CoralWatch is a citizen science project managed by the University of Queensland that aims to "*improve the extent of information on coral bleaching events and coral bleaching trends*" [15]. Currently the CoralWatch project has over 1300 members from 80 countries and its members have collected over 29,000 surveys. CoralWatch provides simple color charts (Fig. 1) that can be used by anyone (scientists, tourists, divers, school students) to provide useful monitoring data on coral bleaching on a relatively large scale via an inexpensive, 'user friendly' and non-invasive devices. Data collected through the CoralWatch program includes coral species, coral color, latitude and longitude of the location, reef name, water temperature, data and time and the method by which the data is collected e.g., snorkeling, reef walking or fishing. As well as collecting monitoring data, the project aims to educate the public about the causes and impact of bleaching on coral reefs.

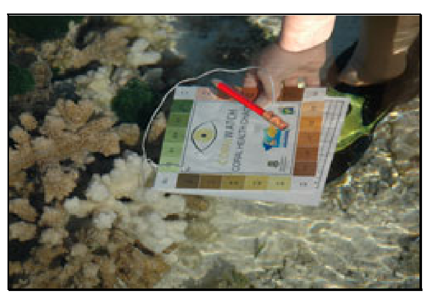

Figure 1. Use of Coral Health Chart in the field

New members register through the CoralWatch website[1]. Once registered, the member can request a DIY Coral Health Monitoring Kit through the website. The kit provides a field guide for recording observations. Each observation includes coral types and color intensity of the coral. The user records the color intensity of the coral species observed by comparing it with a chart. "*The colour charts are based on the actual colours of bleached and healthy corals. Each colour square corresponds to a concentration of symbionts contained in the coral tissue. The concentration of symbionts is directly linked to the health of the coral*" [13]. The user generates an online survey by recording observations (of species, color, lat, long, etc) along transects and inputting the data to the CoralWatch database via an online form.

A. Existing Data Quality Issues

A detailed analysis of the CoralWatch data (approx. 18560 records, collected between July 2003 and September 2009), was carried out in order to determine the quality. A significant number of errors were identified. Fig. 2 illustrates the distribution of error types and the extent of errors in the data. Fig. 2 shows that significant errors occurred in the GPS data (~64% of records) and temperature data. Such errors make the observations close to useless from a scientific perspective – although the reef name does provide a coarse positioning. There were also a significant number of errors in the volunteers' contact details – making it difficult to attribute errors to individuals, to hold individuals responsible for the data or to contact volunteers to clarify, confirm or correct outlying data. The causes of the majority of the errors were due to:

- Lack of validation and consistency checking;
- Lack of automated metadata/data extraction;
- Lack of user authentication and automatic attribution of data to individuals;
- Absence of a data model;
- Lack of data quality assessment measures;
- Lack of feedback to volunteers on their data;
- Lack of graph, trend analysis and visualization tools.

By our estimation, over 70% of the errors could be prevented by focusing on new services available via the CoralWatch Portal that focused on the gaps described above.

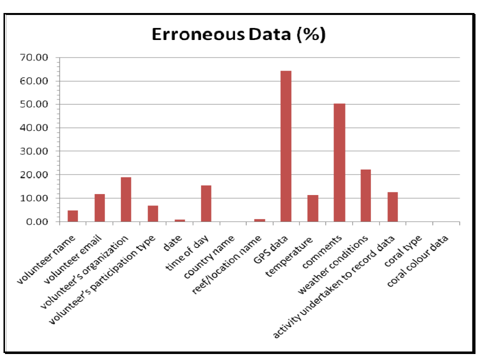

Figure 2. Errors in the legacy CoralWatch data

[1] www.coralwatch.org

To reduce syntactic errors, we implemented a metadata and data suggestion/validation process that employs XML Schemas and controlled vocabularies to restrict input to permitted values/ranges and to validate the data. Registered, authenticated members submit their data through a user friendly Web page that performs form validation and checking before the data is saved. For example, country lists and reef names are validated against the GPS data provided by the member. Input data is run through the data quality cycle on submission and the data is assigned a rating value based on the outcome of the quality measure process. If the data does not pass the data quality assessment, it will be marked "unvalidated". Unvalidated data is checked by a mediator – the number corrections to the data are monitored and used as a measure of the quality of past data from that contributor.

V. Developing the Reputation Model

The first step in developing the reputation model was to analyse the Coral Watch project by answering the key questions listed at the end of Section III. Coral Watch does have a large community of users (over 1300) – of whom about 70% have been involved for over 3 years. There is a turnover of approx. 30% of volunteers annually. User profiles are available, but many are incomplete. The data undergoes both automatic and manual QA/QC as described above – so it is possible to monitor the number of errors in and corrections to past data. In addition, a core set of staff and users have definite views on who the best volunteers are and know them personally. Ground truth data is available via related datasets including the ReefCheck data, NOAA satellite data (which provides sea surface temperature data), and AIMS bleaching events data. These data sets provide an imperfect benchmark against which we may be able to identify outliers or inconsistencies.

A. Calculating Reputation

The methods by which these reputation systems calculate and represent a reputation value, varies significantly [19]. For example, online marketplace sites such as eBay and Amazon consider reputation as a single value (represented as number, star or bar rating) that is independent of the context. The information used to calculate the value of reputation is derived from other agents that have interacted previously with the target agent [19].

Due to the nature of the Coral Watch project, we are in a position to implement a hybrid content-driven and user-driven reputation model. In order to calculate the reputation for entities (both users and data) we apply weightings to and aggregate the attributes listed below:

- a1 = the contributor's age;
- a2 = the contributor's role and qualifications (primary student, secondary student, PhD student, volunteer, council worker, scientist);
- a3= quality of past data that has been contributed (determined by monitoring the no. of corrections);
- a4= The amount of past data contributed;

- a5 = The frequency and period of contributing;
- a6 = Direct rating from other members;
- a7 = Inferred rating across the social network using a social trust algorithm [11];
- a8= Direct rating of survey data (if it exists).

Fig. 4 illustrates our model for calculating a unified reputation value for an object *rateeObj* based on the different criteria listed above. Each time an object (data or person) is created in the system, the *reputationCalculatorService* is called to create a *reputationProfile* for that object. Periodically, the *currentReputationValue* for each object is recalculated by executing an algorithm that calculates the different attribute values (a1-a8) for the *rateeObj*, applies a weighting (w_i) and aggregates them to generate an overall score (1-5 stars). For example:

- if (age <18) then (a1 = 1)….
- if (role is scientist) then (a2 = 5)
- a4 = (userContributions/totalContributions)*20.

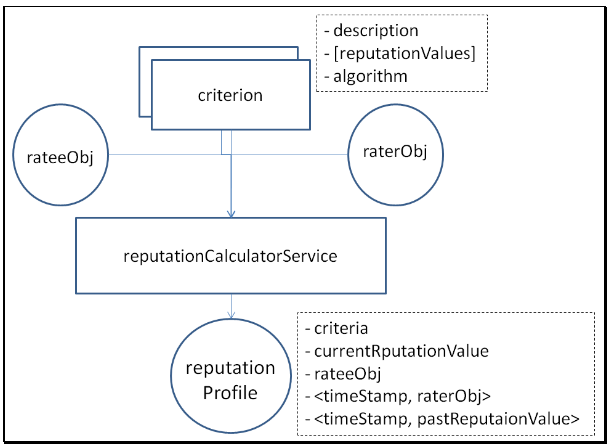

Figure 3. Unified Reputation Calculator Model

An automatic rating agent is a process that detects newly submitted data and uses the *reputationCalculatorService* to recalculate the *currentReputationValue* for a user based on the latest submission. It also keeps track of *pastReputationValues* recorded in the *reputationProfile*.

The process for calculating the inferred reputation value (a7) for a particular user is based on the reputation inference algorithm described in [17] but uses a 1-5 rating system and works as follows. The *sink* (*s*) is the contributor for whom a rating is desired, and the *source* (*i*) is the contributor for whom the rating will be made. The source polls each of his/her neighbors for which he/she has given a positive reputation rating. Neighbors with negative ratings are ignored, since their reputation means that they give unreliable information. Each of the source's trusted neighbours return their rating for the sink. The source will then average these ratings and round the final value to a value from 1-5. This rounded value is the inferred reputation rating from source to sink. If there is no direct arc between the neighbours and the sink, the value must be inferred. The inferred rating from the source to the sink is calculated via a weighted average of the reputation ratings returned for the sink by each of its neighbours.

Fig. 4 illustrates the application of this social trust inferencing algorithm to calculate "trustworthiness" between members of the Coral Watch network – some who do and do not directly know each other. For example, Macdonald wants to know if he can trust Mishra (and hence Mishra's data). There are direct reputation ratings from Macdonald to Administrator, Administrator to Charlie, and Charlie to Mishra. Each of these ratings is aggregated and divided by the number of arcs to infer a direct rating between the *source*, Macdonald and the *sink*, Mishra.

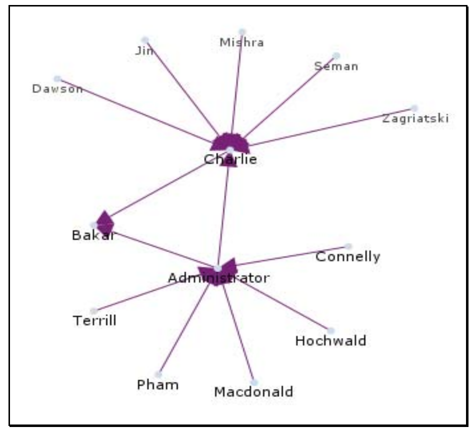

Figure 4. Visualization of CoralWatch Trust Network

B. System Architecture

The diagram in Fig. 5 provides an overview of the system architecture of the revised CoralWatch Web interface and database that we developed in collaboration with the CoralWatch project manager. The system utilizes the PostgreSQL object-relational database management system for storing and processing CoralWatch data and R statistical functions (e.g. statistical analysis of coral watch data to determine whether a bleaching event has occurred). PostGIS [16] also provides geospatial operations such as high speed spatial queries.

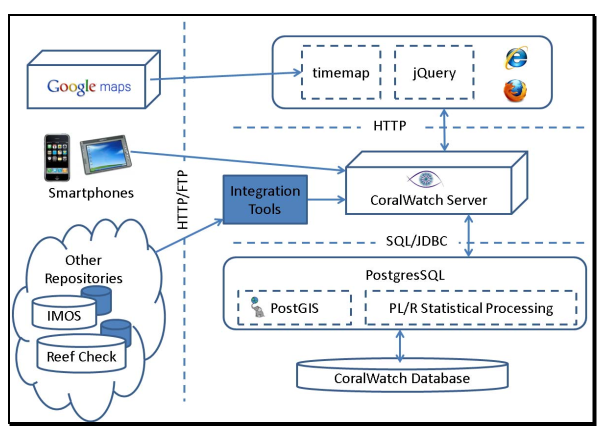

Figure 5. CoralWatch system architecture.

The Coral Watch portal provides the interface by which users can upload their data, view surveys and reports, download data and interact with other users. A Smartphone interface is currently under development to enable data to be uploaded directly from the field and automatic extraction of GPS data, and date and time. The system utilizes Google Maps to provide the geospatial interface to coral bleaching survey data. A SIMILE timeline enables users to specify the temporal period of interest. This provides a tracking mechanism of bleaching events and the speed at which they are happening.

The integration tools are a set of scripts that harvest data, images and files from other related repositories (e.g. IMOS satellite imagery data, AIMS coral bleaching event registry), and map it to a common observational data model. These datasets are used as a benchmark or "ground truth" to measure the quality of the volunteers' CoralWatch data. The benchmark data sets can also be easily displayed as layers via GoogleMaps – on which the CoralWatch surveys can be overlaid.

C. *User Interface for Assigning and Displaying Trust*

Users first need to register via the Coral Watch Web site. Registration requires users to enter contact details, current role (secondary student, undergrad, postgrad, post-doc, research fellow, professor, teacher, volunteer), expertise and professional qualifications. Once registered, a user profile is stored and they are assigned a user id and password and an initial base reputation score (1-2 stars, depending on their age/role/qualifications).

Authenticated users create a new survey (set of observations that document the current state of a particular reef) by first entering the metadata for the survey. This includes the location (lat and long), date/time, weather conditions and water temperature. A validation process then checks compliance of the input data against an XML schema. Once the user has created a new survey, they can enter the set of observations of coral species and color (Fig. 6). Every time the user submits an observation, the data is instantly analysed on the server side. The charts generated from the data analysis show the color distribution across the observed coral reef transect. Users can determine whether a bleaching event has occurred on a particular reef by analyzing the change in color over time.

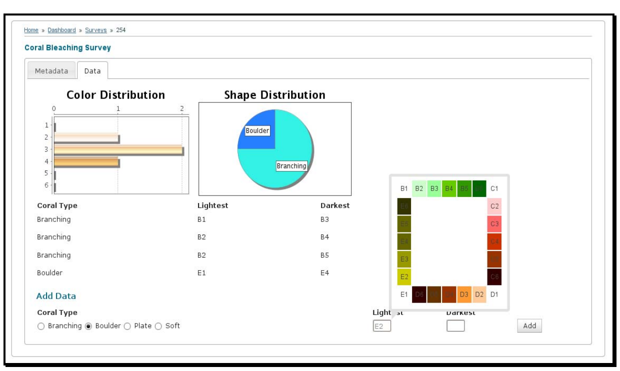

Figure 6. Submitting data via the CoralWatch Web Entry Form.

Once the survey data has been entered, the next step is to calculate a reputation score for it. To date, we have developed simple tagging tools whereby members of the network can assign trust rankings to other members. The aggregate community trust value on a member is calculated by weighting and aggregating both direct and inferred trust values plus additional attributes (e.g., role, expertise, quality of past data, frequency and extent of past contributions) as described in section V.A. The calculated aggregate score is a 1-5 star rating in the user's profile (Fig. 7) – this information is visible only to the system administrator. The reputation scores is then associated with the data uploaded by that user.

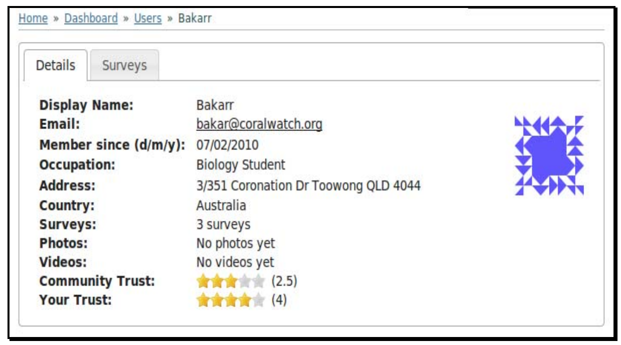

Figure 7. User profile showing trust as 5 star rating.

D. *Reputation-aware Querying, Filtering and Visualization*

The screenshot in Fig. 8 shows the mapping and timeline interface to the CoralWatch data. Through this interface users are able to perform spatio-temporal and keyword-based queries and analysis of the citizen science data. Users are also able to specify the reputation/trust level required. For example: "Show me all coral watch observations for Masthead Reef between 2007 and 2009 with a ranking of 3 or more stars".

The coral bleaching surveys (represented by coloured markers on the map) are layered simultaneously on both the map and the timeline above the map. When the timeline is dragged horizontally to a specific date, the surveys that were conducted around that date are displayed on the map. The user can click on the survey markers on either the timeline or the map to display a balloon that contains the survey metadata and observational data. The observations in Fig. 8 are colored according to their trust metric – red = 1 star, purple=2 star, yellow=3 star, white = 4 star, green = 5 star.

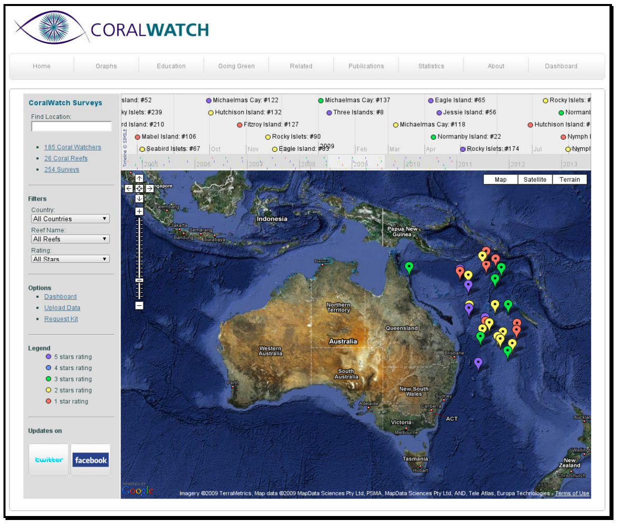

Figure 8. User Interface of CoralWatch

VI. DISCUSSION

Following the implementation of the revised Coral Watch system described above, evaluations were carried out on different aspects of the system to determine improvements to data quality, the user interface, performance and data reliability:

- A survey of feedback from the users and administrators of the CoralWatch data indicated that the trust metrics associated with individual users should be hidden – so as not to deter volunteers – but that trust metrics associated with specific datasets should be explicit. Poor trust metrics associated with individual volunteers could be used to target online training modules. Good trust metrics could be used to reward, encourage and retain volunteers.
- The response from users to the ranking/tagging tools and the improved filtering, search and browse interfaces – was that these tools were easy to use and greatly improved users' ability to understand the temporal, seasonal and spatial trends in coral bleaching events.
- Deliberate submission of consistently poor data by dummy users was picked up eventually by other members of the network who assigned low rankings to these contributors. But there was a delay period during which the data and the user was unknown and assigned an "average" ranking – which was not sufficient to filter it out.

Further effort is required to refine the algorithms for calculating the reputation score. This research will focus on:

- Optimizing the weightings applied to the different attributes (a1-a8) to improve the accuracy of trust/reputation score
- Measuring the performance, accuracy, efficiency and scalability of the reputation assessment tools as the size of the community and the database grows;.
- Monitoring changes in the number and frequency of contributing volunteers, the retention of existing volunteers and the attraction of new volunteers.

VII. CONCLUSION

Citizen science is democratizing science in that it enables public citizens and the scientific community to work together in monitoring, managing, maintaining and understanding the environment around us. A literature review has revealed that there is a critical need for a framework to improve the quality and trust of citizen science data – and that there exists a range of technologies from the data quality, social trust and online reputation fields, that can potentially be combined to maximize the quality and re-use of citizen science data.

Using the CoralWatch project as a case study, we have implemented a system that demonstrates that it is possible to significantly improve the quality of community generated observational data by calculating a measure of the trustworthiness of citizen science data using a weighted aggregation of both direct and inferred attributes. By explicitly enabling this metric to be displayed to users, and considered within querying and reporting services, we have enhanced the potential re-use of citizen science data by scientists.

REFERENCES

[1] Lintott, C.J., et al., *Galaxy Zoo : Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey.* Monthly Notices of the Royal Astronomical Society, 2008. **389**(3).

[2] Cooper, C.B., et al., *Citizen science as a tool for conservation in residential ecosystems.* Ecology and Society, 2007. **12**(2).

[3] (MESSAGE), M.E.S.S.A.G.E. *Project Overview.* 2010 [cited 2010 20/02/2010]; Available from: http://bioinf.ncl.ac.uk/message/?q=node/5.

[4] Galloway, A, Tudor, M. and Haegen, W., *The Reliability of Citizen Science: A Case Study of Oregon White Oak Stand Surveys.* Wildlife Society Bulletin, 2006. **34**(5): p. 1425-1429.

[5] Foster-Smith, J. and Evans, S., *The value of marine ecological data collected by volunteers.* Biological Conservation, 2003. **113**(2): p. 199-213.

[6] Paulos, E., *Designing for Doubt Citizen Science and the Challenge of Change*, in *Engaging Data: First International Forum on the Application and Management of Personal Electronic Information*. 2009: MIT - Cambridge, MA, USA.

[7] Silvertown, J., *A new dawn for citizen science.* Trends in Ecology and Evolution, 2009. **24**(9): p. 2832-2842.

[8] Delaney, D.G., et al., *Marine invasive species: validation of citizen science and implications for national monitoring networks* Biological Invasions, 2007. **10**(1): p. 117-128.

[9] Nardin, L., et al., *SOARI : A Service Oriented Architecture to Support Agent Reputation Models Interoperability*, in *AAMAS-TRUST*, 2008. p. 292-307.

[10] Josang A., Ismail, R., Boyd, C. *A survey of trust and reputation systems for online service provision*, *Decis. Support Syst.* 43, 2 (March 2007), pp 618-644.

[11] Levien, R., *Attack Resistant Trust Metrics*, in *Computing with Social Trust*, J. Golbeck, Editor. 2009, Springer: London. p. 121-132

[12] Sabater, J. and C. Sierra, *Review on Computational Trust and Reputation Models.* Artificial Intelligence Review, 2005. **24**(1): p. 33-60.

[13] Golbeck, J., *Trust and nuanced profile similarity in online social networks.* ACM Trans. Web, 2009. **3**(4).
[14] De Alfaro, L., Kulshreshtha, A., Pye, I., and Adler, B.T.. 2011. *Reputation systems for open collaboration*. Commun. ACM 54, 8 (August 2011), 81-87.
[15] Reid, C., et al., *Coral Reefs and Climate Change: The guide for education and awareness*. 2009, Brisbane: CoralWatch.
[16] PostGIS. *"What is PostGIS?* 2010 11/02/2010]; Available from: http://postgis.refractions.net/.
[17] Golbeck J., Hendler, J., *Accuracy of Metrics for Inferring Trust and Reputation in Semantic Web-based Social Networks*, LNCS Volume 3257/2004, 116-131, DOI: 10.1007/978-3-540-30202-5_8