

CollabMap: Augmenting Maps using the Wisdom of Crowds

(Extended Abstract)

Ruben Stranders, Sarvapali D. Ramchurn, Bing Shi, Nicholas R. Jennings

School of Electronics and Computer Science
University of Southampton
Southampton, UK
{rs2,sdr,bs07r,nrj}@ecs.soton.ac.uk

Introduction

The creation of high fidelity scenarios for disaster simulation is a major challenge for a number of reasons. First, the maps supplied by existing map providers (e.g., Ordnance Survey,¹ TeleAtlas) tend to provide only road or building shapes and do not accurately model open spaces which people use to evacuate buildings, homes, or industrial facilities (e.g. the space around a stadium or a commercial centre both constitute evacuation routes of different shapes and sizes). Secondly, even if some of the data about evacuation routes is available, the real-world connection points between these spaces and roads and buildings is usually not well defined unless data from buildings' owners can be obtained (e.g. building entrances, borders, and fences). Finally, in order to augment current maps with accurate spatial data, it would require either a good set of training data (which is not available to our knowledge) for a computer vision algorithm to define evacuation routes using pictures (working on aerial maps) or a significant amount of manpower to directly survey a vast area.

Against this background, we develop a novel model of geospatial data creation, called CollabMap, that relies on *human computation*. CollabMap is a crowdsourcing tool to get users contracted via Amazon Mechanical Turk or a similar service to perform micro-tasks that involve augmenting existing maps (e.g. GoogleMaps or Ordnance Survey) by *drawing evacuation routes*, using satellite imagery from GoogleMaps and panoramic views from GoogleStreetView. In a similar vein to (Von Ahn, Liu, and Blum 2006; Heipke 2010), we use human computation to complete tasks that are hard for a computer vision algorithm to perform or to generate training data that could be used by a computer vision algorithm to automatically define evacuation routes. In so doing, we advance the state of the art in the following ways. First, we propose the first crowdsourced mapping system that relies on large numbers of non-experts to generate maps that define and connect open spaces (occupied by pedestrians or vehicles) to buildings and roads.² Second, we extend the Find-Fix-Verify pattern by Bern-

stein et al. (2010) to include measures of trust and reputation of the task performers. Third, we propose a number of evaluation mechanisms for determining the quality of the maps generated using such free-form, creative drawing micro-tasks. Finally, we aim to use the maps generated by CollabMap to provide more accurate maps to emergency responders and to build simulations using an evacuation simulator such as RoboCupRescue³ to help design emergency response plans for the emergency planning department at Hampshire County Council in the UK.

CollabMap

CollabMap crowdsources the task of identifying building evacuation routes to a large number of users, by offering them freely available data, such as satellite imagery (e.g. Google Maps), panoramic views (e.g. Google Streetview) and building shapes (e.g. Ordnance Survey and OpenStreetMap⁴) to carry out this task. By so doing, even users not familiar with the location of a building or area can potentially contribute evacuation routes (though local inhabitants are expected to provide more accurate data and tasks could be targeted at them if their location is known). The scope of a task is a single building, and follows a workflow based on the Find-Fix-Verify pattern proposed by Bernstein et al. (Bernstein et al. 2010). Using this workflow, tasks are subdivided into multiple smaller activities, each carried out by a different user, and fall into three categories: task identification, task execution and task verification:

Task Identification The IDENTIFIER actor draws on a map the outline of a building for which no evacuation routes currently exist. This activity is executed once per task.⁵

Task Execution The EXECUTOR actor first checks the building outline drawn by the IDENTIFIER. If it is correct, she draws a single evacuation route on a map. This activity can be executed multiple times, depending on how many evacuation routes are identified.

CollabMap but focus on defining where areas, buildings, or roads are rather than connecting the spaces around them.

³<http://www.sf.net/roborescue>

⁴Building shapes are available for a few countries, including Germany and The Netherlands.

⁵This activity is only required if no data is available about building shapes, or if the data provider disallows sharing data with unlicensed users, as is the case with Ordnance Survey.

¹Ordnance Survey is the UK's national mapping service.

²Google's MapMaker (<http://www.google.com/mapmaker>), OpenStreetMap (<http://www.openstreetmap.org>) and Geo-Wiki (<http://www.geo-wiki.org>) are crowdsourcing initiatives similar to

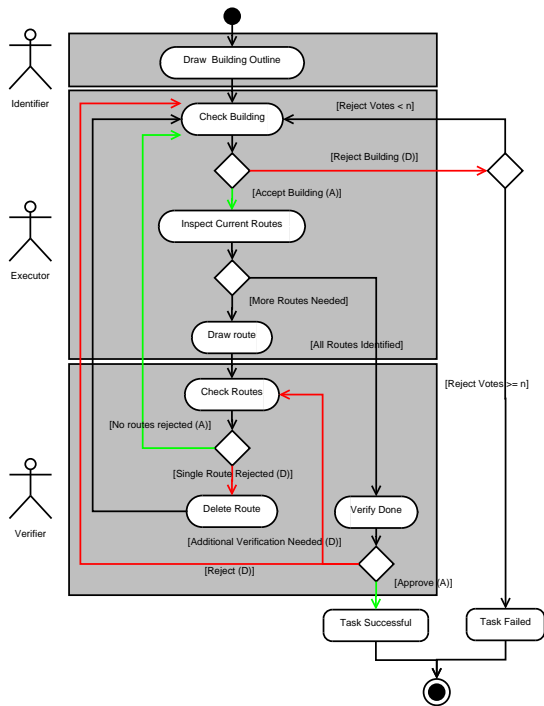


Figure 1: CollabMap’s task workflow represented as a UML activity diagram.

Task Verification The VERIFIER is shown all routes drawn by EXECUTORS. She either rejects a single route that she deems incorrect, or accepts all current routes. This activity is executed after every Task Execution activity.

Figure 1 shows the detailed interactions between these three activities represented as a UML activity diagram. This workflow is designed to ensure that each contribution by a user is independently verified by another. Moreover, users have only partial control of the end result: VERIFIERS assess the validity of *all* routes, but can only reject a single route, and IDENTIFIERS can only draw a single route. In so doing, CollabMap aims to achieve a higher level of trustworthiness and reliability of the end result compared to a system that collects evacuation routes without independent verification.

Figure 2 shows the user interface for EXECUTORS.⁶ Note that the questions above the drawing area correspond to the decisions points in the workflow for EXECUTORS.

Metrics, Reputation and Incentives

In order to further increase the reliability of the outputs, CollabMap keeps track of the reputation of users. CollabMap infers a user’s reputation from task execution traces using the *consensus ratio*: the ratio between agreements and disagreements (indicated by the green (A) and red (D) transitions in Figure 1 respectively) that can be attributed to her. For example, a user who rejects a building that is later accepted by the majority of users, or identifies a route that is rejected by another user are both counted as disagreements for which that user is responsible. Similarly, the accuracy of a task can be expressed in terms of the consensus ratio of

⁶The user interfaces for IDENTIFIERS and VERIFIERS are not shown due to space limitations.

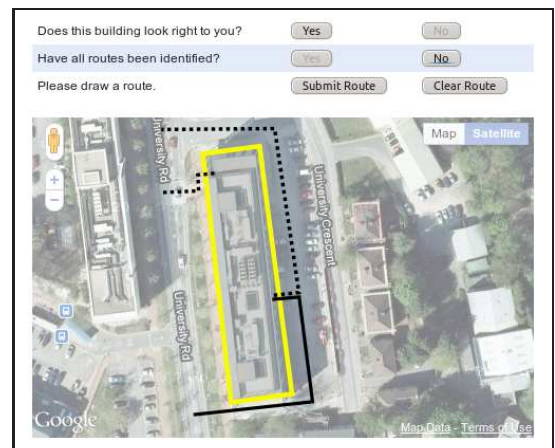


Figure 2: The user interface for the Task Execution activity. The solid black line is the route drawn by the EXECUTOR for the building drawn in yellow by the IDENTIFIER. The dotted lines are the routes drawn previously by other EXECUTORS. Panoramic views not shown due to space restrictions.

the (dis)agreements in its execution trace. To incentivise accuracy and reliability, CollabMap awards a monetary bonus to users with a high level of consensus and bars users whose consensus level falls below a threshold.

Future Work

CollabMap leverages the wisdom of crowds to augment existing maps with evacuation routes and spaces that are essential for the preparation and simulation of large-scale evacuation procedures. CollabMap is still at an initial stage⁷ and we are currently working on a production version that will be deployed on Amazon Mechanical Turk. Future work will investigate the use of GPS traces to more accurately assess the output of CollabMap and the reputation of users. In addition we intend to further improve reputation metrics based on the familiarity of a user with the area (inferred from her physical location). Finally, we wish to maintain task execution traces as the *provenance* of evacuation routes using the Open Provenance Model (Moreau et al. 2010), in order to make the processes that led to their creation more transparent and auditable.

Acknowledgements

This work was carried out as part of the ORCHID project funded by EPSRC (EP/I011587/1).

References

Bernstein, M. S.; Little, G.; Miller, R. C.; Hartmann, B.; Ackerman, M. S.; Karger, D. R.; Crowell, D.; and Panovich, K. 2010. Soylent: a word processor with a crowd inside. In *Proceedings of UIST 2010*, 313–322.

Heipke, C. 2010. Crowdsourcing geospatial data. *ISPRS Journal of Photogrammetry and Remote Sensing* 65(6):550–557.

Moreau, L.; Clifford, B.; Freire, J.; Futrelle, J.; Gil, Y.; Groth, P.; Kwasnikowska, N.; Miles, S.; Missier, P.; Myers, J.; Plale, B.; Simmhan, Y.; Stephan, E.; and den Bussche, J. V. 2010. The open provenance model core specification (v1.1). *Future Generation Computer Systems*.

⁷See a working video at: <http://www.vimeo.com/22624839>.

Von Ahn, L.; Liu, R.; and Blum, M. 2006. Peekaboom: a game for locating objects in images. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 55–64. ACM.