

Verifying Truth from the Ground: Leveraging Human Strengths in the Image Geolocation Process



Sukrit Venkatagiri, Jacob Thebault-Spieker, Kurt Luther

@thesukrit

sukrit@vt.edu

sukritv.com



Introduction

Photo and video verification is a key step in investigations in journalism, human rights advocacy, and disaster response. We focus on a subset of image verification, geolocation, in which expert investigators work to identify the exact geographic location where a photo/video was made [1]. This manual process does not scale easily and experts are often overwhelmed by the amount of information they must process [2]. Furthermore, computer vision-based geolocation techniques are ineffective because they are not generalizable.



Completed Work

Crowd Performance in Image Geolocation

Motivation

Geolocation is a largely manual process and experts often spend hours or days searching for an image in satellite imagery. They are overwhelmed by the amount of images they must geolocate and are unable to scale.

Research Questions

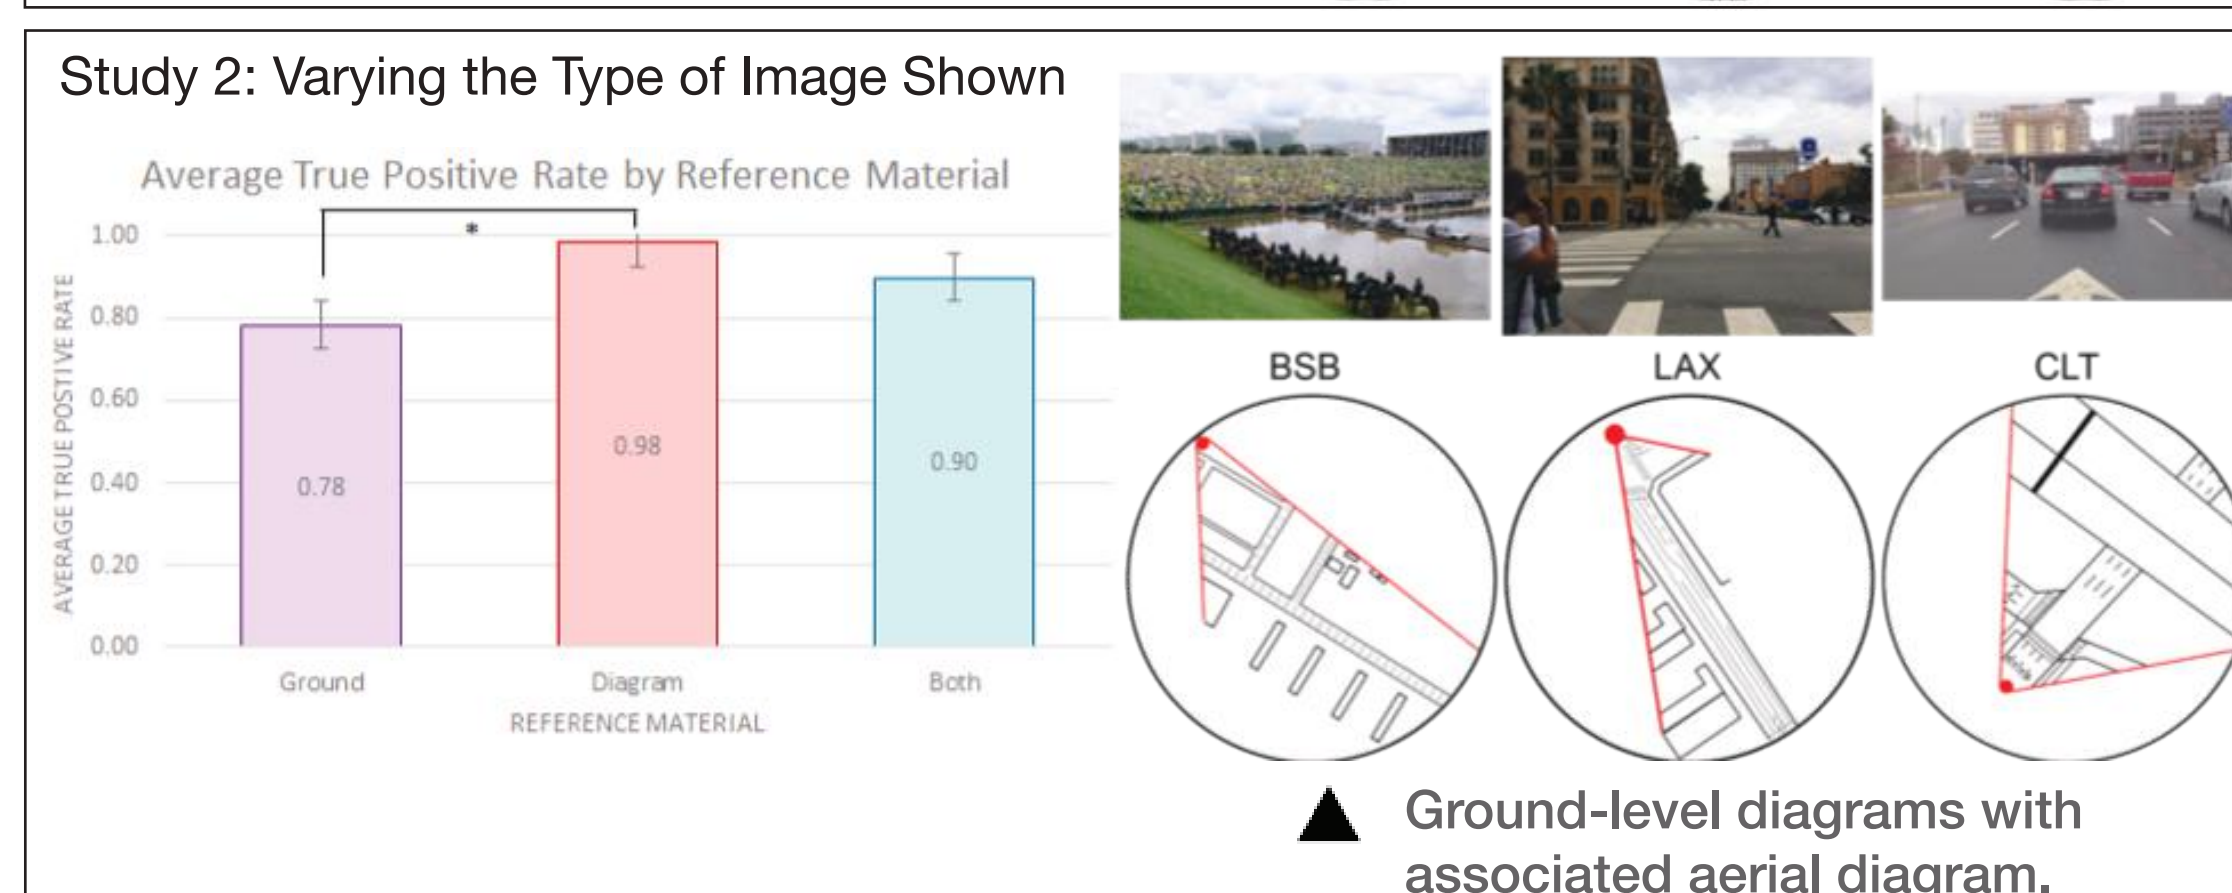
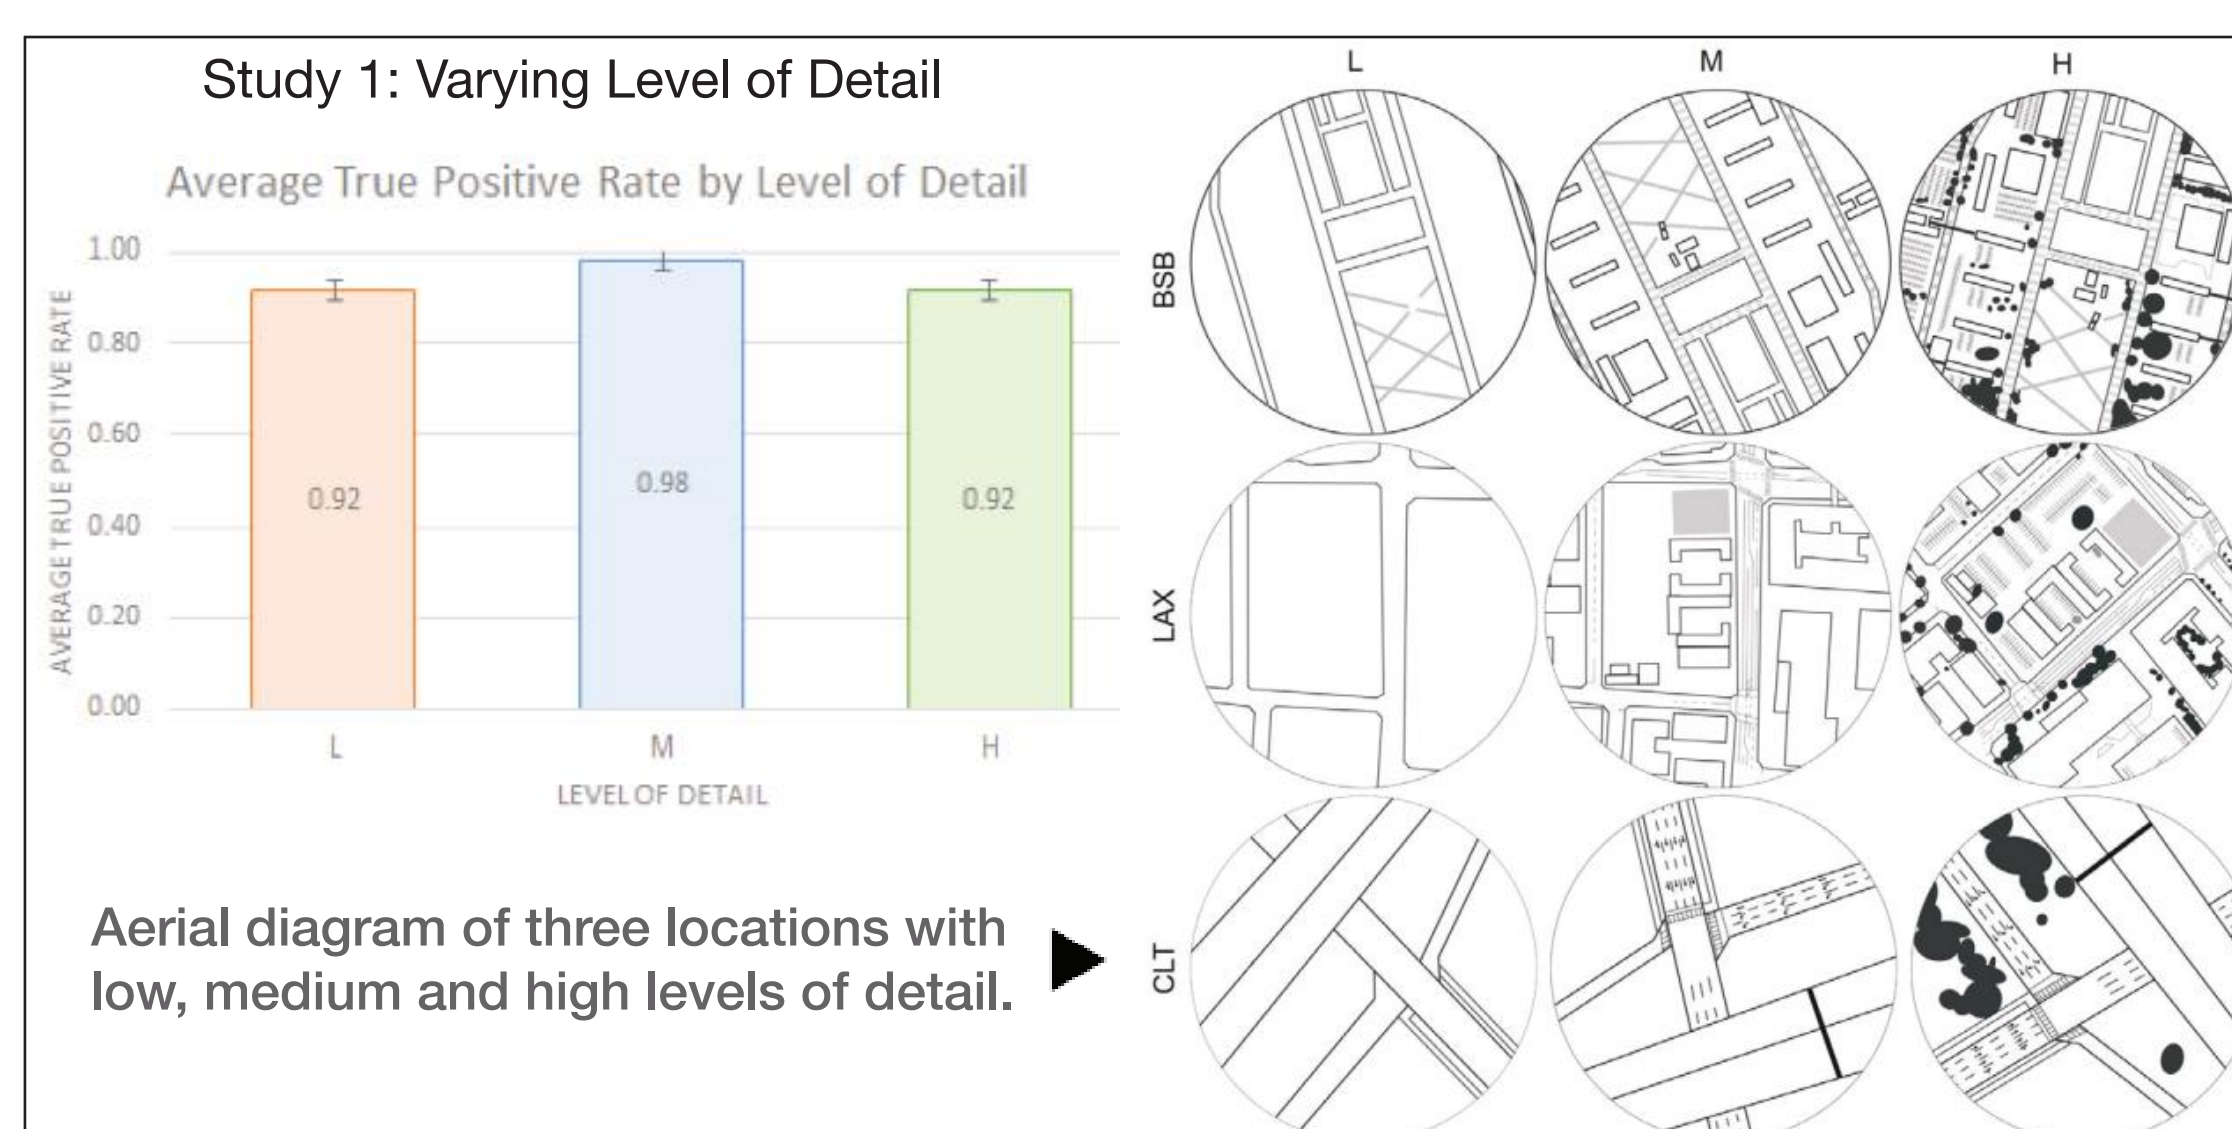
- RQ1: With what level of detail does the crowd perform best?
- RQ2: Does the crowd perform better with a ground-level photo, an aerial diagram, or both?
- RQ3: How fast can the crowd reduce the search space?

Method

We adapted an expert diagramming technique to overcome spatial reasoning limitations of novice crowds, allowing them to support an expert's search. We considered how this diagramming technique can be adapted for crowds who lack an expert geolocator's spatial thinking skills [1].

Results

- The crowd performed best with a medium level of detail.
- The crowd performed significantly better with aerial diagram than with ground-level image.
- The crowd reduced the search area by 50% in 10 minutes, without expert intervention.



Our System: GroundTruth

GroundTruth is a web-based system that enables experts to designate a search area for crowdworkers to investigate in parallel and then review the aggregated results for a potential location match.



In Progress

Study 1: User Study with Expert Geolocators

Motivation

Our prior work has shown that crowds are effective at image geolocation. However, it is not known if experts will listen to and trust crowd feedback or if they prefer their current practices over one that involves crowds.

Research Questions

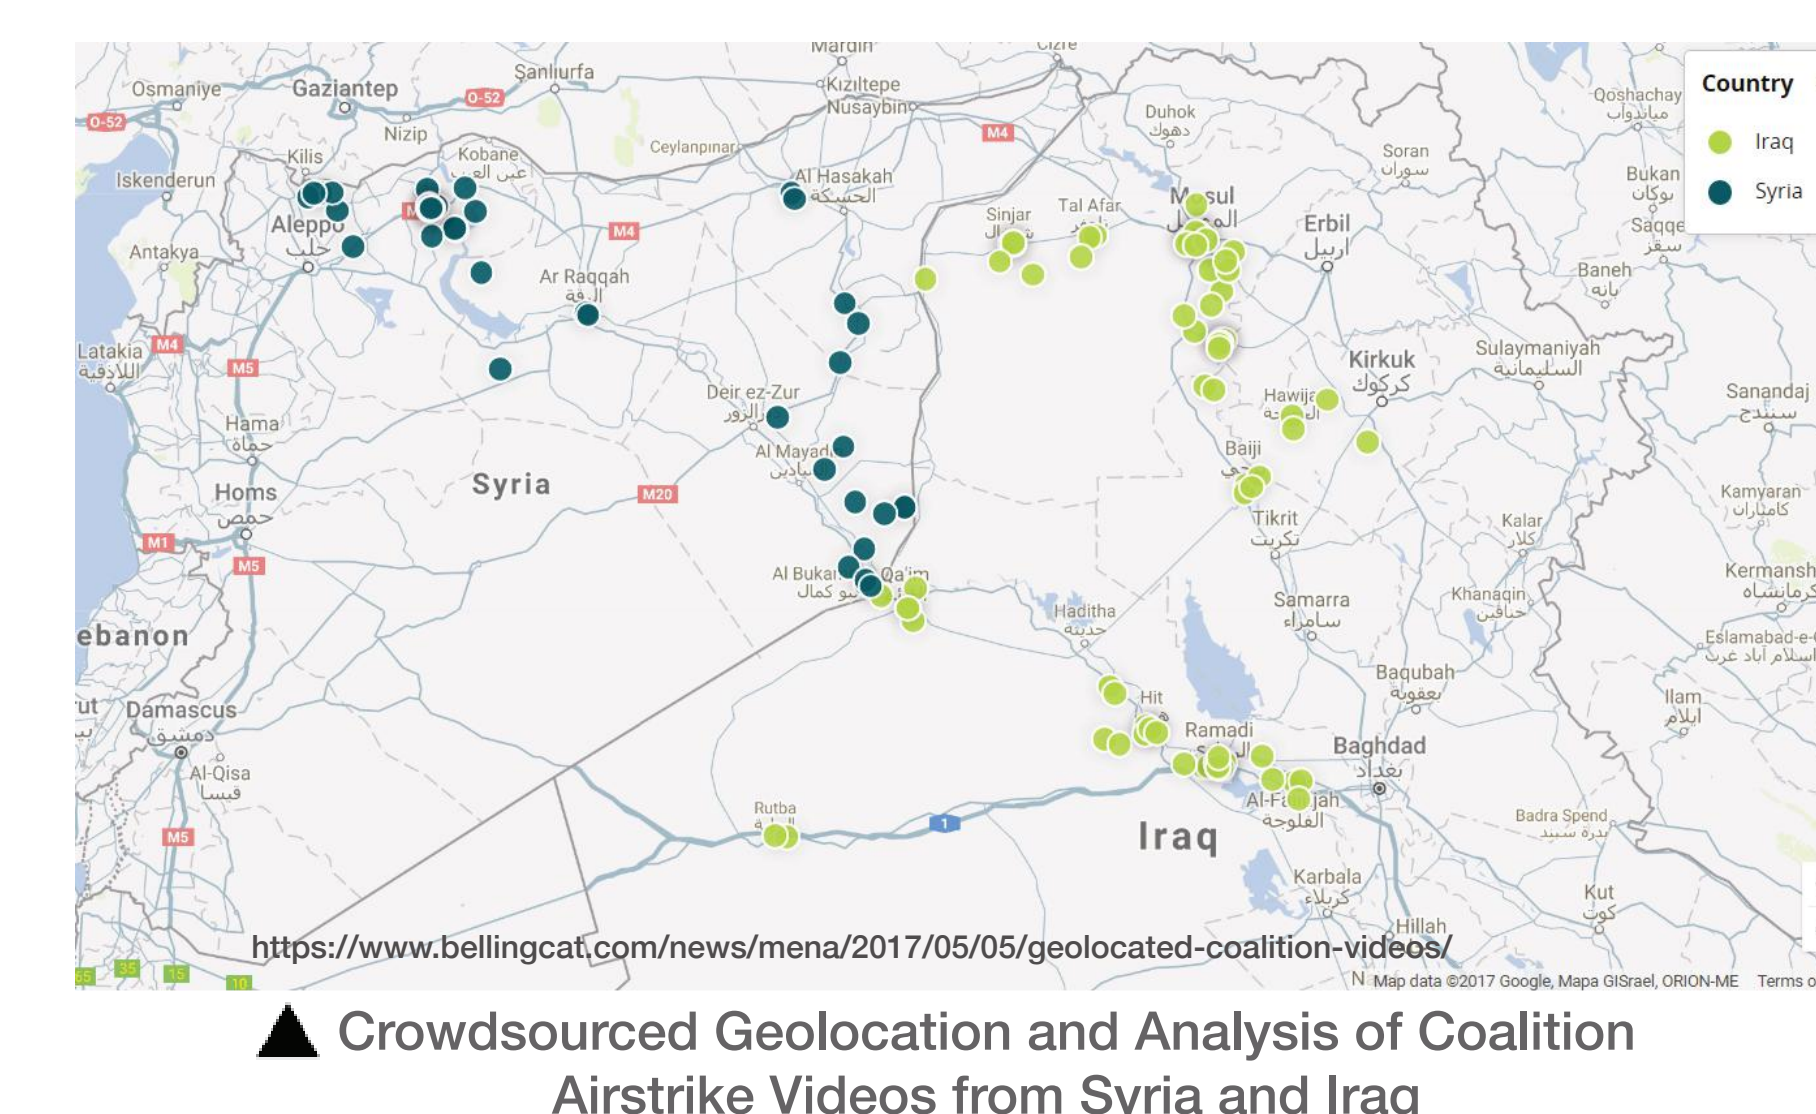
- RQ1: How do expert geolocators collaborate with the crowd in real-time?
- RQ2: How does GroundTruth fit in to current expert practices?

Method

Using a think-aloud protocol, we will observe expert geolocators as they use GroundTruth to geolocate images from a computer vision-based image geolocation dataset. This will allow us to gain insights into how they collaborate with the crowd. We will follow this up with a semi-structured interview.

Expected Results

- Successful geolocation outcomes and a positive overall experience, resulting in a strong preference for GroundTruth.
- Insights into how experts collaborate with crowds on geolocation tasks, and generalized design implications.
- Facilitating a higher rate of image geolocation by scaling expertise with the help of the crowd, thereby combating the spread of misinformation.



Study 2: Exploring Geographic Biases in Crowd Work

Motivation

Geographic crowd work is crucial to Humanitarian OpenStreetMaps, Uber, crowdsourced geolocation, and other domains.

However, the geographic nature of this work may be influenced by differences in human geographic knowledge and familiarity.

Research Questions

- RQ1: Does distance-based localness affect an individual's success at geographic crowd work?
- RQ2: Does density-based localness affect an individual's success at geographic crowd work?

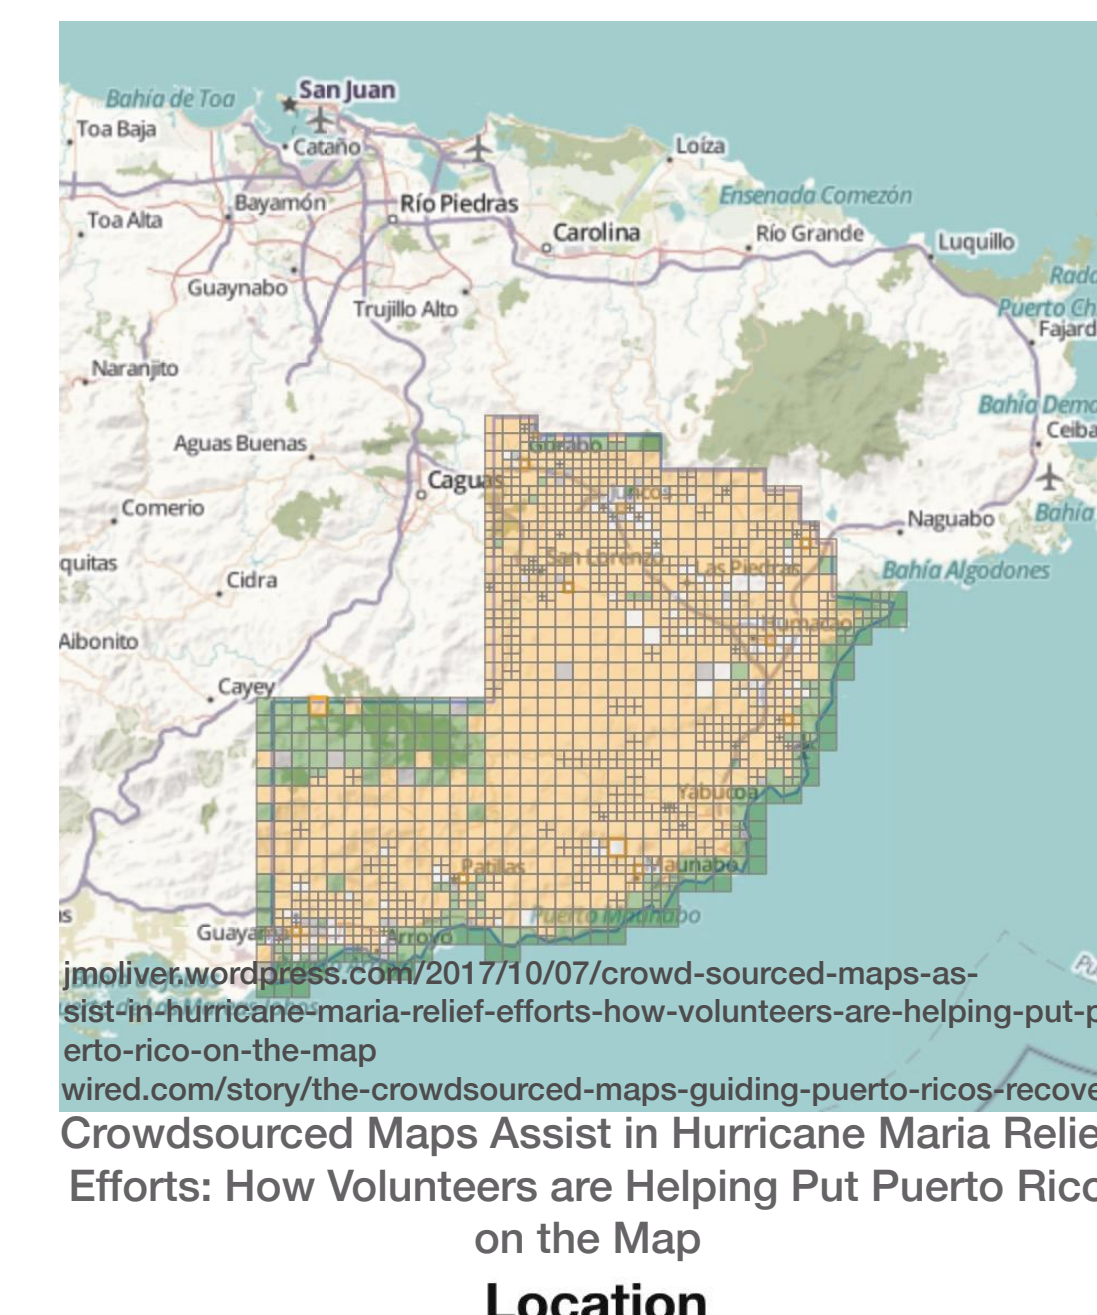
Method

In two high-stakes domains, we focus on understanding the correlation between distance-based localness and success, and density-based localness and success. In both cases, we seek to predict or experimentally verify success based on how far away a contributor is from a given area, and the demographic information about that area.

- Observational Study:** mapping in Humanitarian OpenStreetMap (HOT)
- Experimental Study:** geolocation in GroundTruth

Expected Results

- Distance-based localness affects success: the closer someone is, physically, the more likely they are to succeed at geolocation and mapping.
- Density-based localness affects success: In both mapping and geolocation, people from areas of similar density tend to be more successful.
- By taking localness into consideration, crowd-based geographic systems can greatly increase success factors.



	Location	
Urban	Minnesota	Urban Virginia
Rural	Rural Minnesota	Rural Virginia

Acknowledgments

We would like to thank Rachel Kohler, Anne Hoang, and Rifat Sabbir Mansur. We would also like to thank the crowdworkers on Amazon Mechanical Turk for their countless hours of effort. This work was funded by NSF awards 1527453 and 1651969.

References

- [1] R. Kohler, J. Purviance, and K. Luther. Supporting Image Geolocation with Diagramming and Crowdsourcing. AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2017)
- [2] R. Kohler and K. Luther. Crowdsourced Image Geolocation as Collective Intelligence. Collective Intelligence 2017. (CI 2017)
- [3] Mitchell Gordon, Jeffrey P. Bigham, and Walter S. Lasecki. 2015. LegionTools: A Toolkit + UI for Recruiting and Routing Crowds to Synchronous Real-Time Tasks. In Adjunct Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15 Adjunct)