Designing for Coordinated Action: Challenges and Opportunities of Working with Experts

Sukrit Venkatagiri
sukrit@vt.edu
Department of Computer Science and
Center for Human-Computer Interaction
Virginia Tech

Jacob Thebault-Spieker
jthebaultspieker@vt.edu
Department of Computer Science and
Center for Human-Computer Interaction
Virginia Tech

Kurt Luther
kluther@vt.edu
Department of Computer Science and
Center for Human-Computer Interaction
Virginia Tech

ABSTRACT

Expert investigators bring advanced skills and deep experience to analyze visual evidence, but they face limits on their time and attention. In contrast, crowds of novices can be highly scalable and parallelizable, but lack expertise. Here, we discuss a real-world deployment of a system we built, GroundTruth, that enables experts to direct crowds in determining the exact geographic location where an image was taken, a key step in verifying visual evidence. We discuss the challenges working with these experts and the methods we have employed to study our software deployment.

ACM Reference Format:
Sukrit Venkatagiri, Jacob Thebault-Spieker, and Kurt Luther. 2019. Designing for Coordinated Action: Challenges and Opportunities of Working with Experts. CSCW 2019 Workshop on Mapping the "How" of Collaborative Action. Austin, TX, USA, 5 pages.

CSCW 2019 Workshop on Mapping the "How" of Collaborative Action, 2019

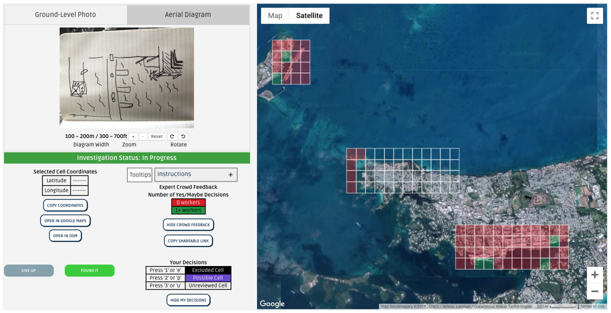

Figure 1: Expert interface where expert specifies locations for the crowd to search, view crowd feedback, and conduct their own parallel search.

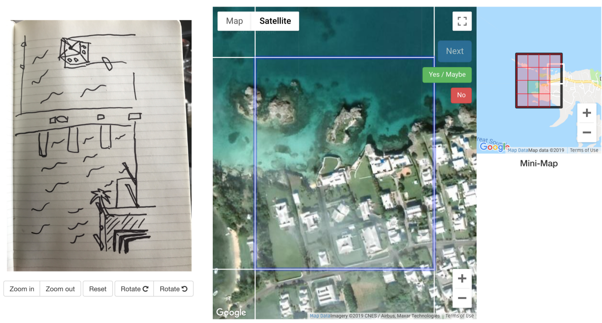

Figure 2: Crowd worker interface where a crowd worker can provide feedback on whether the aerial diagram matches the cell of satellite imagery or not.

INTRODUCTION

Expert investigators, such as journalists and human rights advocates, increasingly leverage social media in their investigations as sources of visual evidence as well as volunteers. For example, Bellingcat [2], a community of citizen journalists, uses imagery from social media to investigate claims made by and about governmental or terrorist activity. They often ask their followers on social media for help.

A key step in investigations is image geolocation [1, 4], which involves finding the exact location where photo or video imagery was taken. Initially, an expert inspects the image for clues, and then uses inference and experience to conduct a manual, brute-force search through large swathes of satellite imagery, looking for the location in the image. However, this process is time-consuming and does not scale easily. Although experts leverage online crowds in their investigations, these crowds lack expertise and their efforts may be uncoordinated, leading to wasted effort. Undirected, crowdsourced investigations can also lead to vigilantism [7] and misidentification [5].

PRIOR WORK: SYSTEM DEVELOPMENT AND EVALUATION

Over 18 months, we employed a user-centered design process to build GroundTruth [8], a software tool that allows experts to coordinate the efforts of novice crowds to geolocate images. We first developed a deep understanding of expert investigative work practice by studying existing literature and interviewing 12 experts, followed by sketching, prototyping, and empirical evaluations.

GroundTruth enables experts to efficiently coordinate the efforts of crowds in three ways. First, it leverages the existing expert practice of creating a diagram to represent the image under investigation, highlighting for the crowd what to search for. This improves crowd performance by augmenting their spatial reasoning skills. Second, the expert uses deep domain knowledge to narrow down a set of possible locations for the crowd to search through. This better directs crowd worker effort. The system then generates microtasks for individual crowdworkers to search satellite imagery in parallel and at scale, reducing the time an expert must spend. Finally, the system aggregates and displays crowd feedback as a heatmap for the expert to review. It both visualizes the crowd's progress and allows an expert to keep track of where they have searched.

In a lab-based evaluation, we provided 11 experts with a diverse set of images to choose from and geolocate with crowds in real-time. We displayed images in a gallery to cater to experts' varied backgrounds and preferences. We found that GroundTruth effectively combined the deep knowledge of experts with the speed and scale of crowds, narrowing a search area by up to 67%. Experts also said that GroundTruth facilitated the coordination of crowd feedback, allowing experts to better allocate their own time and effort. We also observed new collaboration dynamics where experts would choose to review crowd feedback, instead of conducting an independent search, or divert their search to locations where the crowd had not yet provided feedback.

Name	Dakota	Nick	Christiaan
Occupation	Journalist	Open-Source Investigator	Journalist
Employer	Storyful	Bellingcat	New York Times
Experience (years)	5	3	5
Avg. number of images investigated per week (# images)	25	30-50	5
Use image geolocation (% images / week)	75–100%	75–100%	75–100%
Geolocation success rate (% images / week)	50–75%	50–75%	75–100%
Ask on social media for help (% images / week)	0–25%	0–25%	25–50%
Ask colleagues for help (% images / week)	25-50%	0–25%	0–25%

Table 1: Expert participants.

Although we have learned much from our lab study, the controlled setting presented some limitations. The images we provided were not representative of what they typically encounter. We also recruited paid crowds on Amazon Mechanical Turk, but some experts said that they would prefer to use GroundTruth with colleagues or social media followers acting as the crowd, or on their own. Experts also faced time constraints, meaning that we could not observe how they would use GroundTruth within a larger investigation. These limitations led us to where we are now: a real-world deployment.

CURRENT WORK: REAL-WORLD DEPLOYMENT

Our motivations for conducting a real-world deployment are three-fold and involve developing a deep understanding of in-situ processes and contexts of use [6]. First, we wish to understand how experts with diverse skills and expertise will use GroundTruth to geolocate real-world images as part of their work practice. Second, we seek to understand how experts will work with volunteer crowds. Third, we see GroundTruth as a technology probe [3] to help experts to consider other parts of their investigative process that could benefit from crowdsourced support, eliciting needs and desires for the design of future tools. To date, we have deployed GroundTruth with three experts in three different organizations—Christiaan, Dakota, and Nick—who are currently integrating it into their workflows and using it for real-world investigations (Table 1).

Challenges

As our expert participants continue to use GroundTruth, we have faced a number of challenges related to working patterns, organizational contexts, skill sets, tool usage, and processes.

One key challenge is the diverse set of skills and contexts of use among the three experts. Although all experts perform image geolocation, they rarely seek crowdsourced assistance on social media. Even when they do so, they may not have a sufficient number of followers to quickly gain a critical mass. For example, Dakota relies on social media for 0–25% of her investigations, and asks individuals for information, but rarely engages in open calls for help on social media. However, she and her colleagues often ask each other for help when they are stuck or need the expertise of someone who is familiar with a certain geographic region. On the other hand, Nick relies heavily on open-source information posted on social media and leverages his large Twitter following. Yet, for more sensitive cases, he prefers to rely on his colleagues when he needs analytic help. However, he and his colleagues often compete against each other to see who can geolocate an image first, instead of collaborating together.

Second, all three experts deal with different time frames for verifying imagery. For example, Dakota works at the breaking news desk at Storyful and must verify imagery and claims within a short period of time, sometimes within just 15–30 minutes. For Dakota, who conducts 25+ investigations per week, it may not be feasible to rely on social media for help in such a short period of time. On the other hand, Christiaan, who works on the Visual Investigations team at The New York Times, conducts fewer investigations per week (5), but they are more in-depth and have a longer time frame. In his situation, he may not need to rely on social media followers because there is less pressure to verify the images quickly. He may also enjoy conducting these investigations himself. For all three experts, the adversarial crowd and getting "scooped" by other organizations may also be a concern.

Our Approach

To overcome some of these challenges, we have employed a mixed-methods approach. First, we initially held a workshop session with each expert to provide a demonstration of GroundTruth as well as form a common understanding of each others' objectives. We also discussed ways for experts to effectively use GroundTruth, such as how to draw good diagrams, how they plan to recruit crowd workers, and for what types of imagery they planned to use GroundTruth.

Second, we wished to elicit highly contextual feedback on how experts were using GroundTruth. However, due to experts' time constraints, it would have been infeasible for them to record their entire session and provide us detailed descriptions of how they conducted each investigation, or interact with us frequently. Instead, we provided them prompts to think about their most recent session and asked them to record a five-minute recording each time they used GroundTruth. These video "diary entries" allow experts to reflect on their session in a minimally intrusive manner, while still eliciting valuable insights. They also provide information on implicit and explicit feature requests that we are able to quickly implement, test, and deploy to all of our users.

Third, to understand how experts' experiences and processes have changed over time, we plan to conduct in-depth, semi-structured interviews halfway through and at the end of the deployment. This will allow experts to provide a summative description of their experiences across multiple uses.

Fourth, we are carefully studying recordings of sessions collected through MouseFlow, a commercial screen-recording service, and inspecting system database logs. These logs will tell us how often experts are using GroundTruth, how large of an area they are searching through, how much time they have spent looking at a particular cell, as well as how well crowd workers are performing.

CONCLUSION AND ACKNOWLEDGEMENTS

Although working with expert investigators has proven difficult in terms of access and visibility, we believe our methods will allow us to form a holistic understanding of how GroundTruth is used, appropriated, and integrated in the wild. In addition, we believe our methods will allow us to elicit new insights on how to to deploy systems with expert investigators, and what tools and technological support they desire.

We would like to thank the members of the Crowd Intelligence Lab for their contributions to this paper. We would also like to thank our expert participants for their time and engagement. This work was supported by NSF awards IIS-1651969 and IIS-1527453.

Author Biographical Sketches

Sukrit Venkatagiri is a Ph.D. student at Virginia Tech, where he studies expert investigative and collaborative work practice in fields such as journalism, social cybersecurity, and human rights advocacy. He builds and empirically validates collaboration tools and crowdsourcing systems to augment and scale up their expertise. He employs a variety of mixed-methods in his work—from large-scale experiments, log analysis, and surveys, to interviewing, diary studies, and participant observation.

Jacob Thebault-Spieker is a Postdoctoral Associate at Virginia Tech. The north star guiding his research is understanding and mitigating social biases in crowd work platforms such as in content moderation, ridesharing, and volunteered geographic information systems. In his postdoctoral work, he has begun to study the role of expertise and context in how biases manifest and designing intervention strategies.

Kurt Luther is an assistant professor in the Department of Computer Science and (by courtesy) the Department of History at Virginia Tech. His research group, the Crowd Intelligence Lab, builds and studies crowdsourcing systems that support creativity and discovery, with applications in domains such as national security, journalism, history, biology, and design. His research and teaching interests include crowdsourcing, social computing, human computation, human-AI collaboration, and creativity support tools.

REFERENCES

[1] Petter Bae Brandtzaeg, Marika Lüders, Jochen Spangenberg, Linda Rath-Wiggins, and Asbjørn Følstad. 2016. Emerging Journalistic Verification Practices Concerning Social Media. *Journalism Practice* 10, 3 (2016), 323–342. https://doi.org/10.1080/17512786.2015.1020331 arXiv:https://doi.org/10.1080/17512786.2015.1020331

[2] Eliot Higgins. 2014. A Beginner’s Guide to Geolocation. https://www.bellingcat.com/resources/how-tos/2014/07/09/a-beginners-guide-to-geolocation/

[3] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. 2003. Technology Probes: Inspiring Design for and with Families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’03)*. ACM, New York, NY, USA, 17–24. https://doi.org/10.1145/642611.642616

[4] Rachel Kohler, John Purviance, and Kurt Luther. 2017. Supporting Image Geolocation with Diagramming and Crowdsourcing. In *Fifth AAAI Conference on Human Computation and Crowdsourcing*.

[5] Johnny Nhan, Laura Huey, and Ryan Broll. 2015. Digilantism: An Analysis of Crowdsourcing and the Boston Marathon Bombings. *The British Journal of Criminology* 57, 2 (12 2015), 341–361. https://doi.org/10.1093/bjc/azv118 arXiv:http://oup.prod.sis.lan/bjc/article-pdf/57/2/341/10461339/azv118.pdf

[6] Yvonne Rogers and Paul Marshall. 2017. Research in the Wild. *Synthesis Lectures on Human-Centered Informatics* 10, 3 (2017), i–97. https://doi.org/10.2200/S00764ED1V01Y201703HCI037 arXiv:https://doi.org/10.2200/S00764ED1V01Y201703HCI037

[7] Daniel Trottier. 2017. Digital Vigilantism as Weaponisation of Visibility. *Philosophy & Technology* 30, 1 (01 Mar 2017), 55–72. https://doi.org/10.1007/s13347-016-0216-4

[8] Sukrit Venkatagiri, Jacob Thebault-Spieker, Rachel Kohler, John Purviance, Rifat Sabbir Mansur, and Kurt Luther. 2019. GroundTruth: Augmenting Expert Image Geolocation with Crowdsourcing and Shared Representations. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (November 2019), 107:1–107:30.