Crowdsourcing Images for Global Diversity

Pei-Yu (Peggy) Chi, Matthew Long, Akshay Gaur, Abhimanyu Deora, Anurag Batra, Daphne Luong
Google Research
Mountain View, CA, USA
crowdsource-help@google.com

ABSTRACT

Crowdsourcing enables human workers to perform designated tasks unbounded by time and location. As mobile devices and embedded cameras have become widely available, we deployed an image capture task globally for more geographically diverse images. Via our micro-crowdsourcing mobile application, users capture images of surrounding subjects, tag with keywords, and can choose to open source their work. We open-sourced 478,000 images collected from worldwide users as a dataset "Open Images Extended" that aims to add global diversity to imagery training data. We describe our approach and workers' feedback through survey responses from 171 global contributors to this task.

CCS CONCEPTS

• **Human-centered computing**; • **Information systems** → *Mobile information processing systems*;

KEYWORDS

Mobile crowdsourcing; open source; diversity; global community, images; mobile applications

MobileHCI '19, October 1–4, 2019, Taipei, Taiwan

ACM ISBN 978-1-4503-6825-4/19/10.
https://doi.org/10.1145/3338286.3347546

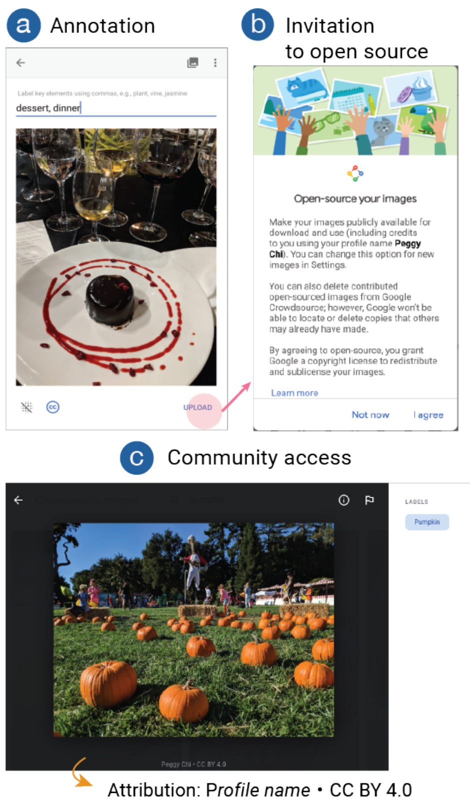

Figure 1: Users annotate a captured image via the task UI in our Crowdsource app (a). Along with the first contribution, a dialog invites users to open source the work under their profile name (b). Open-sourced images are publicly accessible with author attribution (c).

ACM Reference Format:
Pei-Yu (Peggy) Chi, Matthew Long, Akshay Gaur, Abhimanyu Deora, Anurag Batra, Daphne Luong. 2019. Crowdsourcing Images for Global Diversity. In *21st International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '19), October 1–4, 2019, Taipei, Taiwan.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3338286.3347546

INTRODUCTION

Crowdsourcing is a scalable approach for diverse data collection and content creation from human workers by removing the time and location constraints [18]. Tasks can be deployed and completed efficiently [29, 30] nearly in real-time [4, 22]. Recent research has demonstrated how mobile crowdsourcing further engages a wide community of diverse contributors by taking advantage of the global mobile device market [13, 14, 16, 19]. A system can allocate suitable tasks based on device location [1] and utilize short attention time [35].

Crowdsourced data has enabled new technologies that reflect the community viewpoints, including training Machine Learning models in Computer Vision [28], translations [5], topic mining [6], and trend prediction [11]. Data quality from crowd workers is often a critical issue, especially for model training [31, 37] and content publishing [33]. To ensure data quality, crowdsourcing platforms rely on user profiles that receive ratings through task completion. Profiles are used for workforce control without giving attribution to contributors. Such lack of worker visibility, however, may lead to concerns on work quality and ethics [15, 17, 26]. Recent research attempts to provide public access to an open, reliable community, which has been applied to design critics [12, 23] and open-ended research [34]. By introducing crowd visibility, workers produce quality data and sustain for a longer term. However, it is rarely seen that the integrated outcome from a community is officially attributed to the individuals. This practice is distinct from open-source communities of knowledge sharing (e.g., Wikipedia [20]) and software development (e.g., GitHub and Q&A sites [9, 24, 36]), where creators receive proper credit for their work, including articles, code, or answers in public distributions.

In this paper, we describe how we collect real-world images at the global scale via an existing mobile application for micro-crowdsourcing. Our goal is to provide a geo-diverse image dataset for model training. To encourage high-quality work, we enable users to open source their work at the task level beyond a general profile that conventional crowdsourcing systems provide. Users annotate and upload images captured by their device cameras given task instructions. They can open source the work with their profile name. The system attaches the author attribution to open-source images in every distribution (see Figure 1). We have received over 1 million images from worldwide users since April 2018 and have open-sourced 478,000 images as a dataset "Open Images Extended" [3].

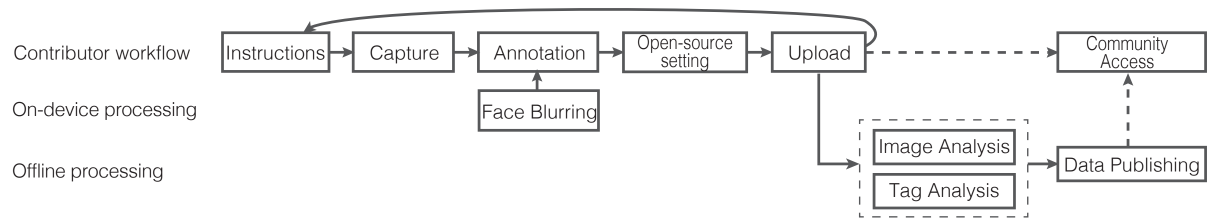

Figure 2: In our micro-crowdsourcing mobile platform, users can choose to open source their individual contribution (an annotated image) under their profile name. They can review and manage the submitted work, which will later be open-sourced for public access with attribution to contributors.

CROWDSOURCING IMAGES WITH ATTRIBUTION

We introduce a crowdsource image capture task where users can choose to open source their contributions. To enable a scalable field deployment that reaches a global community, we build the system upon an existing micro-tasking mobile platform that offers non-monetary rewards. We design a new, independent task that guides users through the workflow with the open-source option.

Background: Mobile Crowdsourcing Platform

We utilize an existing Android microtasking application, Google Crowdsource[1], designed for mobile phone experiences [7]. The app provides microtasks on *data labeling or verification*, such as handwriting transcription and verification of image labels or language translations. Each task requires short completion time and different knowledge or skills. Users can flexibly choose between tasks without a commitment to the number of tasks or work time. To engage users in task completion, the system leverages game mechanics suggested by prior work [38]. In our design, logged-in users receive *badges* for answering questions of different categories and get to *level up* by contributing more answers.

[1] Google Crowdsource. https://play.google.com/store/apps/details?id=com.google.android.apps.village.boond

Image Capture Task

Based on this mobile platform, we design a new task that requires *content generation* for creating a geo-diverse dataset. Public image datasets including ImageNet [10] and OpenImages [21] have been widely used for training Machine Learning models. The source images were collected from photo hosting services (e.g., Flickr) where creators release photograph usage under copyright licenses, such as Creative Commons [8]. However, these services are designed for social sharing purposes and may be limited to specific geo regions [2], which resulted in biased content [31]. To build a diverse image dataset, we distribute an image collection task named "Image Capture" to our global crowdsourcing community. Specifically, we scope to capture subjects around contributors.

Table 1: Our task instructions contain an ordered list of topics. Users review and capture a photo of subjects not limited to this list.

Title	*Content*
Food	Tag with name & ingredients
Plants	Tag with type & name
Professions	Capture people at work. Tag with their profession.
Apparel	Anything you wear: tag with brand & other details
Animals	Include type & breed tags if possible
Products	Anything you can buy: tag with name & place
Landmarks	Natural & manmade: tag with name & place

Contributor Workflow

Figure 2 presents the end-to-end task flow. Users launch our application and see a new task "Image Capture" on the home screen. By choosing this task, an instruction card is shown that describes the task ("*Capture & tag the world around you*") and the potential image use. Then, the device camera is triggered for users to take a photo, which proceeds to the Annotation UI (see Figure 1a). The image that users captured is shown, along with a text field showing a prompt "*Label key elements using commas, e.g., plant, vine, jasmine*" for free-form text labeling. Text length or word count is not limited. The "SUBMIT" button is enabled only when text is entered in the input field. To protect subject's privacy, face detection and blurring is performed on the capture devices before work submission.

Open Source Option. After users complete the annotation and tap to submit their first image, a dialog titled "*Open-source your images*" is immediately shown (see Figure 1b), followed by "*Make your images publicly available for download and use (including credits to you using your profile name* ***User's profile name****). You can change this option for new images in Settings.*" The attribution is retrieved from the user's profile name associated with the logged-in account. Users can follow the "Learn more" link[2] to a web page for detailed statements about the open source licensing and implication.

[2]Crowdsource Help. Open-source images. https://support.google.com/crowdsource/answer/7680293

Contribution Review. The annotated image is uploaded and counts toward in-app rewards (badges and levels) regardless the open source setting. Users then see a new instruction screen showing a list of subjects, such as food and plants (see Table 1). They can upload more photos with the same flow and change the open source setting of each contribution if desired. Users can review and manage their contributions (download or delete) at any time in the Review UI. The appendix shows our detailed screenflow diagram. Finally, once our offline pipeline processes the content, images will be accessible to the public[3] (see Figure 1c). The system ensures that all open-source images receive the author attribution in every distribution.

[3]Crowdsource Images. https://crowdsource.google.com/images

CROWDSOURCED IMAGE DATA

We released this new task in April 2018 to our global app users. Instructions and all the UI messages are translated into 69 world languages and are shown based on the user's device language setting. We ran a series public campaigns on both a social media platform and local communities to promote this task. Top contributors were highlighted with their profile names and pictures shown on our social media page. No monetary incentive was offered for the contributions.

Data Overview

In December 2018, we open-sourced the "Open Images Extended" dataset with 478,000 images and 1.27M labels under the CC BY 4.0 license [3], which permits content sharing and adapting while

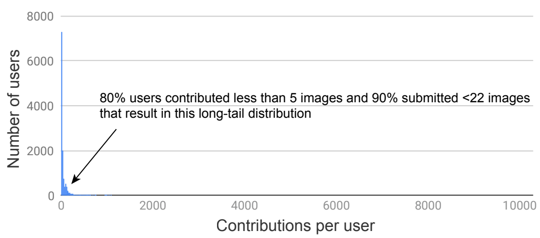

Figure 3: Histogram of contributions per user to our Image Capture task shows a long-tail distribution.

attribution is required. All the content (images and annotations) were reviewed by our human validators to ensure the data quality and exclude any personally identifying information.

To investigate the worker behaviors, we analyzed the raw data we received from April to August 2018. Within the first 5-month period, 587,486 images from users of over 120 countries or regions were received. A total of 14,301 users contributed at least 1 image to this task. Users submitted 41.86 images on average (median is 1). Similar to existing crowdsourcing services [15, 32], there exists a long-tail distribution as "Participation Inequality" suggested [27] (see Figure 3): 66.61% users uploaded 1 or 2 images. 78.59% made 5 or fewer contributions, while the remaining 21.41% contributed 96.96% of the dataset (average 189.59 images). More than 90% uploaded 22 or less images. Finally, the top 1% users contributed 58.29% of images, and the top .09% users each shared over 5,000 images, making it a total of 15.51% of the dataset. As high as 95.57% images were opted in for open source. These images are contributed by 87.5% of the total users to this task. We do not observe any pattern where users selectively open source only a subset of work.

International Survey

To understand the motivation of our users performing the image capture task, we conducted an online survey study to global users with Likert-scale and open questions. A link to the online survey hosted by Google Forms was released to a social group with 6,116 followers that our team maintains. No incentive was offered for survey completion. We received 171 responses in three weeks, of which 130 self-reported their residency that includes 17 countries covered by Bangladesh, Finland, India, Mexico, Russia, Spain, USA, and others. From a screening question, 97.08% (n=166) confirmed that they have uploaded at least one image in this task; the remaining respondents were directed to the end of the survey with no other data recorded.

Why users contribute. 33.7% of respondents indicated that the most important reason they uploaded images is to help improve technology developed based on their contributions, followed by 19.3% who worked because they found this task entertaining (confirmed by the Likert-scale question with the median as 5). We see a wide spectrum of reasons from their open-ended comments on motivations. Some have strong beliefs to make impact: "*I know its gonna help everyone*", "*Mainly to make the internet better*", "*I love (to) be a part of this community*", and "*I mainly take pictures of my hometown, and it's the way to show the whole world my hometown which I still love.*" Feedback is aligned with their top motivations: "*it's interesting and fun to upload images*", "*just done it for fun*", and "*This is one of my hobbies.*" A few indicated that they want to be recognized: "*For recognition of my labor in uploading creative images for the world.*"

In what context users perform tasks. From the multi-select question "*I captured images mostly in:*", 62.9% performed the tasks outdoor, 61.7% at home, and 41.3% in work area.

Figure 4: Global users organized photowalks to share their world (image captured by the community-produced video: *Photowalk by Google Crowdsource Community Malaysia* https://www.youtube.com/watch?v=6Mo93gdYATU)

Perception on open source. While 78% indicated that they understood the concept of "open-source image," 10.4% did not and 11.6% were unsure about the answer. Although the majority of survey respondents open sourced their contributions, only 3% thought it is the most important motivation to have images publicly available under their name. Meanwhile, 12% thought that it is the least important reason of task completion.

DISCUSSION

We aggregate the contribution analysis and international survey and summarize the below findings. As the first content creation task in our micro-crowdsourcing platform with the open source option, users contributed due to the new task design to show their culture and environments. Some of them express strong beliefs and become avid contributors, which explain the long-tail work distribution. One commented, "*i am a hardcore photographer so I take this as a challenge to complete 2000 photos.*"

The majority of users (87.5%) chose to open source their work. Overall, the data quality from the crowd work is positive: As low as 1.70% of the contributions violate our content policy; the detected subjects from images match the topic list from the instructions. Although global users do not show a consistent view on the attribution concept, many want to be recognized and build teamwork, similar to communities such as Google Local Guides and Translates, suggested by survey respondents. From several user-shared videos on social media that documented their self-organized photowalks for this task (see Figure 4), it demonstrates a different aspect of contributor visibility from open source licensing. This shares commonality with prior research on building forums or communities for information sharing among crowd workers [25].

Finally, users are eager to receive feedback of their work: "*We all want to see how my contributions are helping other people that keeps us motivate to do more.*" We recognize the importance of presenting the impact of worker contribution back to the community. Therefore, in addition to open-sourcing the images as a dataset for public download, a public web viewer is also critical to engage users in seeing this community effort interactively.

CONCLUSION

We introduce a micro-crowdsourcing platform where users can open source their work. We deployed an image capture task to an existing mobile phone application and open-sourced 478,000 image with 1.27M labels as the "Open Images Extended" dataset. Based on the aggregated results with survey responses from 171 global contributors to this task, we suggest that users provided high-quality data and followed instructions in our crowdsourcing model. Users are motivated to organize community events (e.g., photowalks) and see their contributions open-sourced with author attribution. We are continuing this crowdsourcing effort and hope to expand our work to document more scenes all over the world.

ACKNOWLEDGMENTS

This work has been possible thanks to the hard work of a lot of people including, but not limited to the following (in alphabetical order of last name): Surabhi Anand, James Atwood, Pallavi Baljekar, Parker Barnes, Tulsee Doshi, Tom Duerig, Vittorio Ferrari, Victor Gomes, Yoni Halpern, George Hodulik, Gursheesh Kaur, Xiaoli Liu, Chandra sekhar Ponnaganti, Mahima Pushkarna, Jigyasa Saxena, D. Sculley, Richa Singh, Steven Thomas Sullivan, Rachelle Summers, and Tanya Wendling.

REFERENCES

[1] Florian Alt, Alireza Sahami Shirazi, Albrecht Schmidt, Urs Kramer, and Zahid Nawaz. 2010. Location-based Crowdsourcing: Extending Crowdsourcing to the Real World. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (NordiCHI '10)*. ACM, New York, NY, USA, 13–22. https://doi.org/10.1145/1868914.1868921

[2] Morgan Ames and Mor Naaman. 2007. Why We Tag: Motivations for Annotation in Mobile and Online Media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 971–980. https://doi.org/10.1145/1240624.1240772

[3] Google Crowdsource app users. 2018. Open Images Extended - Crowdsourced. Retrieved December, 2018 from https://ai.google/tools/datasets/open-images-extended-crowdsourced/

[4] Michael S. Bernstein, Joel Brandt, Robert C. Miller, and David R. Karger. 2011. Crowds in Two Seconds: Enabling Realtime Crowd-powered Interfaces. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 33–42. https://doi.org/10.1145/2047196.2047201

[5] Chris Callison-Burch. 2009. Fast, Cheap, and Creative: Evaluating Translation Quality Using Amazon's Mechanical Turk. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1 (EMNLP '09)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 286–295. http://dl.acm.org/citation.cfm?id=1699510.1699548

[6] Shuo Chang, Peng Dai, Jilin Chen, and Ed H. Chi. 2015. Got Many Labels?: Deriving Topic Labels from Multiple Sources for Social Media Posts Using Crowdsourcing and Ensemble Learning. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion)*. ACM, New York, NY, USA, 397–406. https://doi.org/10.1145/2740908.2745401

[7] Pei-Yu (Peggy) Chi, Anurag Batra, and Maxwell Hsu. 2018. Mobile Crowdsourcing in the Wild: Challenges from a Global Community. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '18)*. ACM, New York, NY, USA, 410–415. https://doi.org/10.1145/3236112.3236176

[8] Creative Commons. 2018. About The Licenses. Retrieved September 1, 2018 from https://creativecommons.org/licenses/

[9] Laura Dabbish, Colleen Stuart, Jason Tsay, and Jim Herbsleb. 2012. Social Coding in GitHub: Transparency and Collaboration in an Open Software Repository. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)*. ACM, New York, NY, USA, 1277–1286. https://doi.org/10.1145/2145204.2145396

[10] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255. https://doi.org/10.1109/CVPR.2009.5206848

[11] Clifton Forlines, Sarah Miller, Leslie Guelcher, and Robert Bruzzi. 2014. Crowdsourcing the Future: Predictions Made with a Social Network. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3655–3664. https://doi.org/10.1145/2556288.2556967

[12] Andrew Garbett, Rob Comber, Edward Jenkins, and Patrick Olivier. 2016. App Movement: A Platform for Community Commissioning of Mobile Applications. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 26–37. https://doi.org/10.1145/2858036.2858094

[13] Jorge Goncalves, Simo Hosio, Niels Van Berkel, Furqan Ahmed, and Vassilis Kostakos. 2017. CrowdPickUp: Crowdsourcing Task Pickup in the Wild. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 51 (Sept. 2017), 22 pages. https://doi.org/10.1145/3130916

[14] Aakar Gupta, William Thies, Edward Cutrell, and Ravin Balakrishnan. 2012. mClerk: Enabling Mobile Crowdsourcing in Developing Regions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1843–1852. https://doi.org/10.1145/2207676.2208320

[15] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. 2018. A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 449, 14 pages. https://doi.org/10.1145/3173574.3174023

[16] Kazushi Ikeda and Keiichiro Hoashi. 2017. Crowdsourcing GO: Effect of Worker Situation on Mobile Crowdsourcing Performance. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 1142–1153. https://doi.org/10.1145/3025453.3025917

[17] Lilly C. Irani and M. Six Silberman. 2013. Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 611–620. https://doi.org/10.1145/2470654.2470742

[18] Durga M. Kandasamy, Kristal Curtis, Armando Fox, and David Patterson. 2012. Diversity Within the Crowd. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work Companion (CSCW '12)*. ACM, New York, NY, USA, 115–118. https://doi.org/10.1145/2141512.2141556

[19] Yongsung Kim, Darren Gergle, and Haoqi Zhang. 2018. Hit-or-Wait: Coordinating Opportunistic Low-effort Contributions to Achieve Global Outcomes in On-the-go Crowdsourcing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 96, 12 pages. https://doi.org/10.1145/3173574.3173670

[20] Aniket Kittur, Ed H. Chi, and Bongwon Suh. 2009. What's in Wikipedia?: Mapping Topics and Conflict Using Socially Annotated Category Structure. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1509–1512. https://doi.org/10.1145/1518701.1518930

[21] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Shahab Kamali, Matteo Malloci, Jordi Pont-Tuset, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. 2017. OpenImages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from https://storage.googleapis.com/openimages/web/index.html* (2017).

[22] Walter S. Lasecki, Christopher D. Miller, and Jeffrey P. Bigham. 2013. Warping Time for More Effective Real-time Crowdsourcing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2033–2036. https://doi.org/10.1145/2470654.2466269

[23] Narges Mahyar, Michael R. James, Michelle M. Ng, Reginald A. Wu, and Steven P. Dow. 2018. CommunityCrit: Inviting the Public to Improve and Evaluate Urban Design Ideas Through Micro-Activities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 195, 14 pages. https://doi.org/10.1145/3173574.3173769

[24] Lena Mamykina, Bella Manoim, Manas Mittal, George Hripcsak, and Björn Hartmann. 2011. Design Lessons from the Fastest Q&a Site in the West. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2857–2866. https://doi.org/10.1145/1978942.1979366

[25] David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a Turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 224–235. https://doi.org/10.1145/2531602.2531663

[26] Brian McInnis, Dan Cosley, Chaebong Nam, and Gilly Leshed. 2016. Taking a HIT: Designing Around Rejection, Mistrust, Risk, and Workers' Experiences in Amazon Mechanical Turk. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2271–2282. https://doi.org/10.1145/2858036.2858539

[27] Jakob Nielsen. 2006. The 90-9-1 Rule for Participation Inequality in Social Media and Online Communities. Retrieved September 1, 2018 from https://www.nngroup.com/articles/participation-inequality/

[28] Dim P. Papadopoulos, Jasper R. R. Uijlings, Frank Keller, and Vittorio Ferrari. 2016. We don't need no bounding-boxes: Training object class detectors using only human verification. *CoRR* abs/1602.08405 (2016). arXiv:1602.08405 http://arxiv.org/abs/1602.08405

[29] Dim P. Papadopoulos, Jasper R. R. Uijlings, Frank Keller, and Vittorio Ferrari. 2017. Extreme clicking for efficient object annotation. *CoRR* abs/1708.02750 (2017). arXiv:1708.02750 http://arxiv.org/abs/1708.02750

[30] Bryan C. Russell, Antonio Torralba, Kevin P. Murphy, and William T. Freeman. 2008. LabelMe: A Database and Web-Based Tool for Image Annotation. *Int. J. Comput. Vision* 77, 1-3 (May 2008), 157–173. https://doi.org/10.1007/s11263-007-0090-8

[31] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. 2017. No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World. In *NIPS 2017 workshop: Machine Learning for the Developing World*.

[32] Osamuyimen Stewart, David Lubensky, and Juan M. Huerta. 2010. Crowdsourcing Participation Inequality: A SCOUT Model for the Enterprise Domain. In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP '10)*. ACM, New York, NY, USA, 30–33. https://doi.org/10.1145/1837885.1837895

[33] Heli Väätäjä, Teija Vainio, Esa Sirkkunen, and Kari Salo. 2011. Crowdsourced News Reporting: Supporting News Content Creation with Mobile Phones. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 435–444. https://doi.org/10.1145/2037373.2037438

[34] Rajan Vaish, Snehalkumar (Neil) S. Gaikwad, Geza Kovacs, Andreas Veit, Ranjay Krishna, Imanol Arrieta Ibarra, Camelia Simoiu, Michael Wilber, Serge Belongie, Sharad Goel, James Davis, and Michael S. Bernstein. 2017. Crowd Research: Open and Scalable University Laboratories. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*. ACM, New York, NY, USA, 829–843. https://doi.org/10.1145/3126594.3126648

[35] Rajan Vaish, Keith Wyngarden, Jingshu Chen, Brandon Cheung, and Michael S. Bernstein. 2014. Twitch Crowdsourcing: Crowd Contributions in Short Bursts of Time. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3645–3654. https://doi.org/10.1145/2556288.2556996

[36] Bogdan Vasilescu, Alexander Serebrenik, Prem Devanbu, and Vladimir Filkov. 2014. How Social Q&A Sites Are Changing Knowledge Sharing in Open Source Software Communities. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 342–354. https://doi.org/10.1145/2531602.2531659

[37] A. Vempaty, L. R. Varshney, and P. K. Varshney. 2014. Reliable Crowdsourcing for Multi-Class Labeling Using Coding Theory. *IEEE Journal of Selected Topics in Signal Processing* 8, 4 (Aug 2014), 667–679. https://doi.org/10.1109/JSTSP.2014.2316116

[38] Luis von Ahn and Laura Dabbish. 2004. Labeling Images with a Computer Game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 319–326. https://doi.org/10.1145/985692.985733

SCREENFLOW DIAGRAM OF THE IMAGE CAPTURE TASK

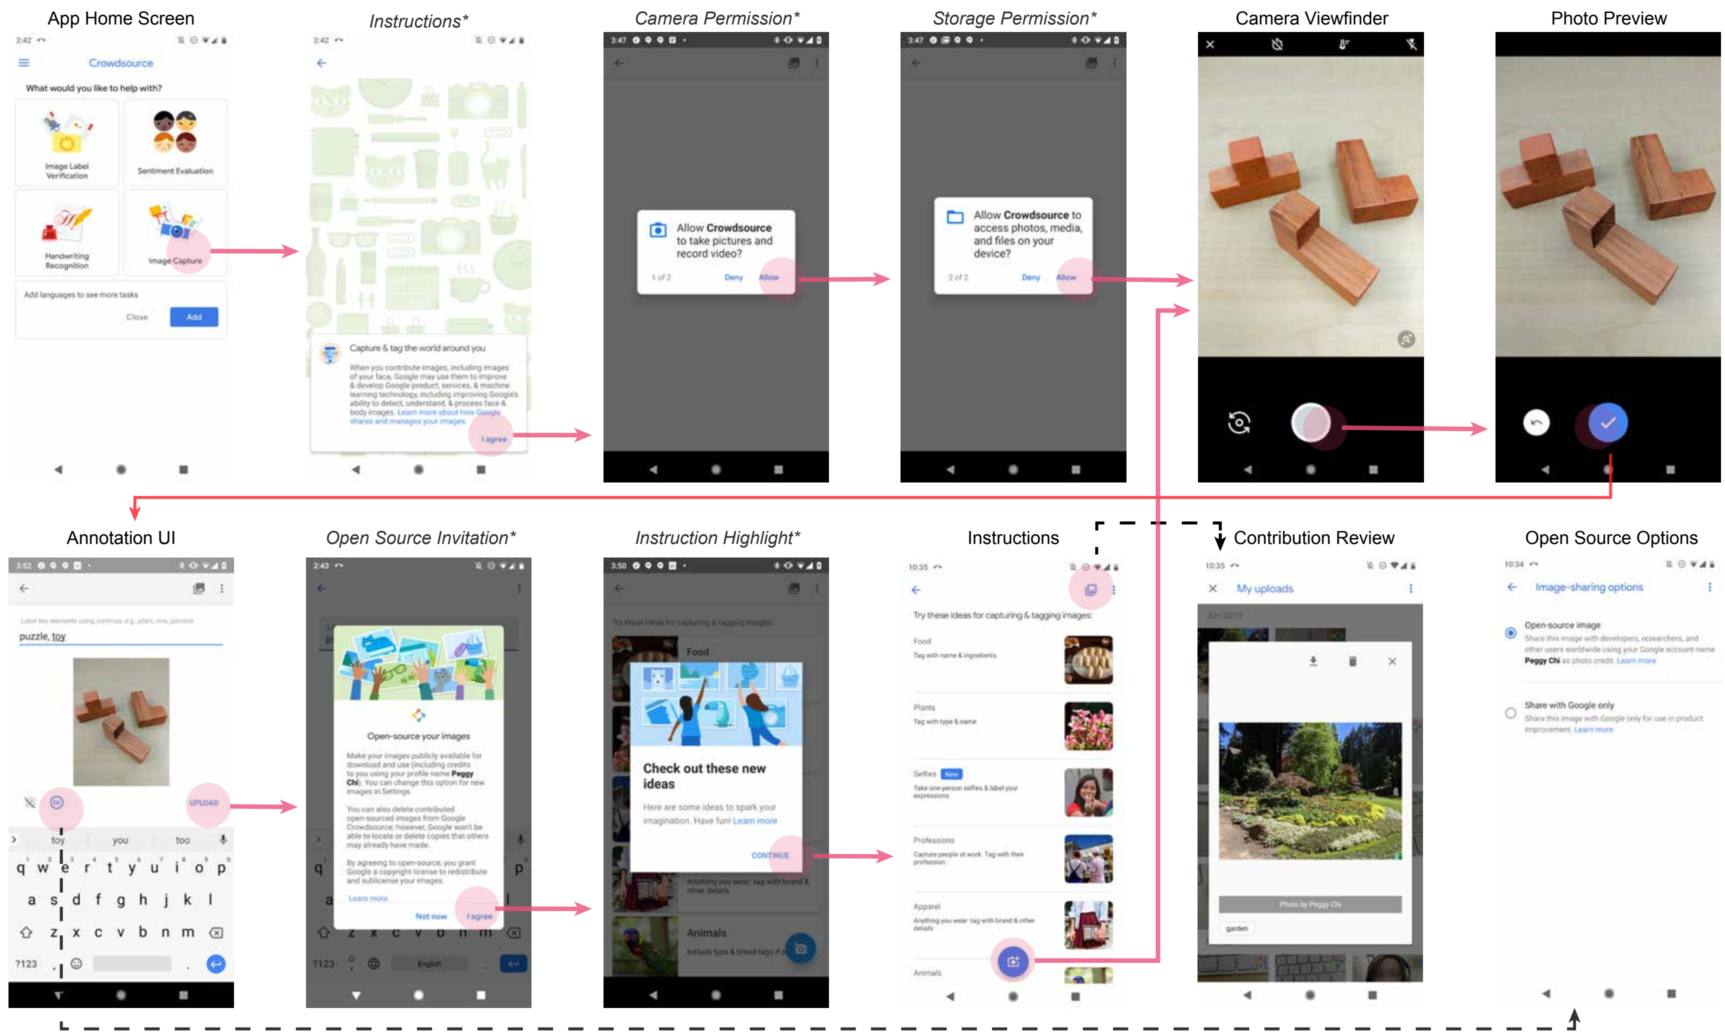

* Screens are only shown once in the workflow during the first task completion.