Contributed Geographic Information: Gray Zones in Collection and Usage

Thomas Kavanagh, Alex Weston

Open Cities AI Competition - Responsible AI Submission - March, 2020

When historic sea level rises encroach on coastal communities, how will we determine the scale of the subsequent refugee crisis? For settlements along seismic fault lines, how do we calculate the number of buildings potentially devastated by earthquakes? As climate crisis and natural disaster threaten the built environment, questions about **infrastructural disaster exposure** become existentially potent. Mitigation of these risks requires application of disaster risk management (DRM) techniques, which in large part depend upon accurate and up to date maps of the built environment.

Terrain maps can serve DRM efforts in a number of ways. As one example, proper flood mitigation relies on a strong understanding of how storm runoff and drainage overflow might move through urban environments. Scientists can use building coverage to calculate the water absorption capacity of the surrounding soil. Knowledge of construction material and building layout allows DRM practitioners to pinpoint structures most at risk in the event of a flood.

Maintaining the accuracy of maps as populations and cities change is costly and time consuming. As populations globally flock to cities, they change the urban environment, making old maps obsolete. The rapid pace of economic development outstrips the efforts of mappers, many of whom are volunteers.

Computer vision and object detection tools hold the potential to dramatically reduce map annotation time and cost. By using human annotations of aerial images of cities, a machine learning model can be trained to detect and segment building areas on a pixel by pixel basis. That model can then be used to create new maps of the most recent aerial imagery on demand. In Guatemala, the World Bank used machine learning techniques to identify homes at high risk for collapse in the event of an earthquake, leading to a dramatic reduction in cost and with performance competitive with human engineers.[1]

The Global Facility for Disaster Reduction and Recovery (GFDRR) has engaged the data science community through DrivenData to create such a tool, in order to “advance the creation of global public goods for improving disaster resilience.” Using labeled building maps of various African cities pulled from OpenStreetMap, the *Open Cities AI Challenge* asks competitors to build computer vision models capable of returning building footprints given high resolution drone imagery of cities. Developing a Building Segmentation Model tailored to African terrains and urban development patterns is preferred over one trained on European urban environments (as was this case with previous building segmentation models).

Alongside these gains, the implementation of such a tool (and the dataset it demands) introduces new ethical considerations. We wish to weigh the evident gains offered by such models against the risks posed to both consenting and non-consenting stakeholders. First, we will consider the potential privacy and security drawbacks of the data collection required to make such models possible. Second, we will outline the potential injustices that could befall disaster relief deployment if the users of this model are not considerate of its inherent biases. As we discuss each issue, we will introduce relevant guidelines for projects involving Contributed Geographic Information (CGI). We hope that GFDRR, in its use of data from OpenStreetMap and other sources, can discuss issues with project planners, data engineers, and surveyors in the outset of its responsible data collection and usage.

An Ethical Framework For Use of Contributed Geographic Information

The following are ethical framework checkpoints, presented in the order they should be addressed.

1. **Projects must be properly scoped**. That is, projects must specify the exact data they require, have clear intentions for the use of the data.
2. Project planners must **investigate avenues for harm posed by the data collection or model creation process**, and address potential security hazards in advance.
3. Any data collection project must **obtain informed consent** from residents or businesses whose privacy is impacted as part of the survey.
4. If informed consent from 1st party stakeholders is not possible, then the onus is on the data collector to **prove that the data project is in the public interest.**
5. It is the responsibility of the data collectors to **collect and/or store the minimum data possible**, thereby limiting a database's potential for abuse.

Data should be granted, not taken.

We build our framework from three foundational principles:

1) The data collected *belongs* to the participants.
2) Project planners must protect participants from abuse.
3) Project planners are responsible for the communication of risks posed by the data collection or modeling process.

It follows from principle (1) that participants have the right to elect their level of participation in a data project. Data collectors must therefore obtain consent from participants to use their data. In order for that consent to be ethically obtained, participants need to understand exactly what they are agreeing to.

"Informed consent" in this context means that a participant is...

- aware of their rights with respect to the data collected,
- aware of how the data collected will be used in the project, and
- aware of the potential risks involved in the disclosure of their data.

Obtaining informed consent for data release is quite an involved process. and is more like an ongoing social ambition than a one-time bureaucratic exercise.

A clear project scope prevents excessive surveillance.

The need for informed consent means that **it is necessary for data collectors to properly scope the project *before* collecting consent from participants**. Otherwise they can't explain to participants how their data will be used.

Surveyors must engage community members before any collection begins, to first and foremost inform residents of (1) the scope and ambitions of the data collection process, (2) the implication of the data's existence on individual or communal privacy, and (3) the mechanisms by which they can exert influence over their digital representations. Surveyors should **consider a mechanism for individual community members to opt out** from the data collection process. The opt-out mechanism would ideally be accessible to community members who may or may not have internet access.

There are already promising community engagement efforts from local volunteers, an example of which is the first person account of field mapping for Open Cities Accra from Chris Eshun, student and YouthMappers training coordinator at the University of Mines and Technology.[4] While collecting data as part of a team of Ghanaian volunteers, Eshun describes how mappers engaged community members who expressed resistance to surveying operations, particularly in response to then-recent demolition projects in the area. Accounts such as this highlight a tangible need to communicate the objectives and downstream community benefits of open source mapping efforts, while raising awareness of the associated risks to privacy and security.

Data stewards must imagine all potential abuses.

Informed consent also means that, in order to articulate the risks for harm associated with the data project to participants, **data collectors need to reckon with potential risks associated with the data prior to collection**. For example, Community members can only give informed consent to geographic surveillance when they are made aware that the same dataset used for DRM can potentially be used by predatory enterprises who could exploit up to date local surveys of residential areas for profit or some other socio-political agenda. Participants should be aware of what third parties, if any, have access to their data. In the context of an open data project (in which case the potential third parties are innumerable), they should understand how the data might be abused by those third parties.

Data collectors also need to assess potential shortcomings of a predictive model, stemming either from the data quality or the modeling process itself. **Modelers must investigate the potential consequences of (potentially unseen) model bias.** Even a high-scoring model trained on data from local data may lead to faulty conclusions if inherent bias isn't accounted for.

Suppose that the downstream use-case for the survey data is to calculate the number of inhabitants in a region. Without local knowledge of population densities / habitation patterns, total area covered by buildings may be a weak predictor of residential population. As it is presented to competitors, the Building Segmentation Model requested in the *Open Cities AI*

Challenge has limited access to contextual information. Such a model might have no knowledge of the housing density, building use, or building height of structures it labels. The responsibility for introducing key contextual information relies on the modeler.

DRM strategists must leverage domain knowledge and neighborhood-level insight to counteract model bias and address the following questions, among others:

- For residential zones, can single-family homes be adequately distinguished from mixed-use or multi-residence dwellings? Is the average population per residence sufficient to generalize for an entire region?
- Are informal, or smaller settlements, adequately represented in the labeled data?
- Are the homeless – and continuous resettlement – accounted for in regional predicted population variance?
- Are communities who lack resources to fight eviction and demolition actions against them underrepresented in assessments of residential building coverage?
- Are active residences distinguished from abandoned buildings?
- Are commercial buildings distinguished from residential ones?

These questions need only be answered in service of more truthful insight about the affected region. In this case, a comprehensive picture of population density might require additional on-the-ground data collection.

When consent is not viable, defer to trusted authorities.

It is often not feasible to collect informed consent from every participant. In fact, this may be true for the majority of open data projects. In addition to consent from participants, surveyors may elect to obtain additional approval from community leaders and/or governmental bodies.

In any case, if informed consent is not possible, surveyors must find some sort of authority that represents the interest of the local population. It is therefore increasingly important for the data collectors to make a strong case for the project to such authorities. If consent of the individual is to be waived, it must be for a reason directly benefiting the community.

More data, more problems. Collect the minimum.

As drones with high quality cameras enter consumer availability, aerial photography can be gathered with higher spatial resolution than was previously possible. In many cases, drone imagery can exceed the spatial resolution of publicly accessible commercial maps (e.g. Google or Apple Maps). Faces, or details about personal property, may be exposed in these images.

This privacy consequence might not be immediately obvious to participants, and they should be made aware. Even if participants are already tolerant of existing commercial satellite imagery databases, but may not realize that drone photography captures neighborhoods in much greater detail.

In general, the more information is collected from participants, the easier it is to trace that information back to individuals, and therefore the greater potential for abuse. Data collectors must weigh the need for predictive power in their model with the risks to participants, recognizing that it is their responsibility to prevent harm from coming to participants as a result of their activities. While it is difficult to know in advance what features will be most useful in a predictive model, data collectors can leverage domain experts' knowledge to start with the most probable features first, expanding as necessary. In an open data project, collectors might only publish those features that they have deemed most predictive after conducting their initial analysis.

Unless it is absolutely necessary for prediction, downstream tasks such as the Building Segmentation Model should avoid usage of high fidelity geodata. Even if up-to-date surveys through drone imagery is necessary for certain models, images should be downsampled before release on public databases. This maintains the privacy status quo without intruding further. Again, a clear project scope will help answer the question, "Is this level of detail and/or prediction actually necessary?"

These considerations are only meaningful if addressed before the data collection process has begun and made open source. But models that build into their pipeline a requirement for high fidelity data in turn creates plausible justification for further collection of this data. As graphics processors become cheaper to run at scale, there is no longer a computational barrier to the storage and processing of high fidelity data. So researchers and archivists must develop models with explicit objectives to minimize reliance on more intrusive data collection practices.

Nobody reads the Terms and Conditions. Data scientists should.

Data scientists have the power to do immense good by leveraging machine learning technologies. Yet those same technologies can inadvertently do harm. In the case of Contributed Geographic Data, unique privacy and ethical considerations arise. But similar principles apply to a diversity of projects. Data scientists should continue to revise guidelines on privacy and ethics as new tools are developed. Our framework constitutes a first attempt at defining a common code for ethical data collection, an attempt we hope will spark further discussion and introspection.

Citations

1. GFDRR. 2018. Machine Learning for Disaster Risk Management. Washington, DC: GFDRR. License: Creative Commons Attribution CC BY 3.0
2. “#VizRisk: Understanding and Communicating Disaster Risk.” *Understanding Risk*, Global Facility for Disaster Reduction and Recovery, 2019, understandrisk.org/vizrisk/what-is-risk/
3. Goodchild, M. F. (2007) Citizens as sensors: the world of volunteered geography. GeoJournal, 69(4), 211–221.)
4. Harvey, F. (2013). To volunteer or to contribute locational information? Towards truth in labeling for crowdsourced geographic information. In Crowdsourcing Geographic Knowledge: Volunteered Geographic Information (VGI) in Theory and Practice (Vol. 9789400745872, pp. 31-42). Springer Netherlands. https://doi.org/10.1007/978-94-007-4587-2_3
5. Weiser, Paul and Abdalla, Amin. (2013). Surveillance in the Context of Security and Profit – The Case of “Non-Volunteered” Geographic Information.
6. Eshun, Chris. “My Experience as a Student Mapper for Open Cities Accra.” *OpenDRI*, Open Data for Resilience Initiative, 19 Nov. 2018, opendri.org/my-experience-as-a-student-mapper-open-cities-accra/.