

Becoming Good at AI for Good

Meghana Kshirsagar*, Caleb Robinson*, Siyu Yang*, Shahrzad Gholami, Ivan Klyuzhin, Sumit Mukherjee, Md Nasir, Anthony Ortiz, Felipe Oviedo, Darren Tanner, Anusua Trivedi, Yixi Xu, Ming Zhong, Bistra Dilkina, Rahul Dodhia, Juan M. Lavista Ferres

meghana.kshirsagar@microsoft.com  Microsoft  USC

The promise of AI to improve lives and protect vulnerable people and ecosystems has not yet reached its potential

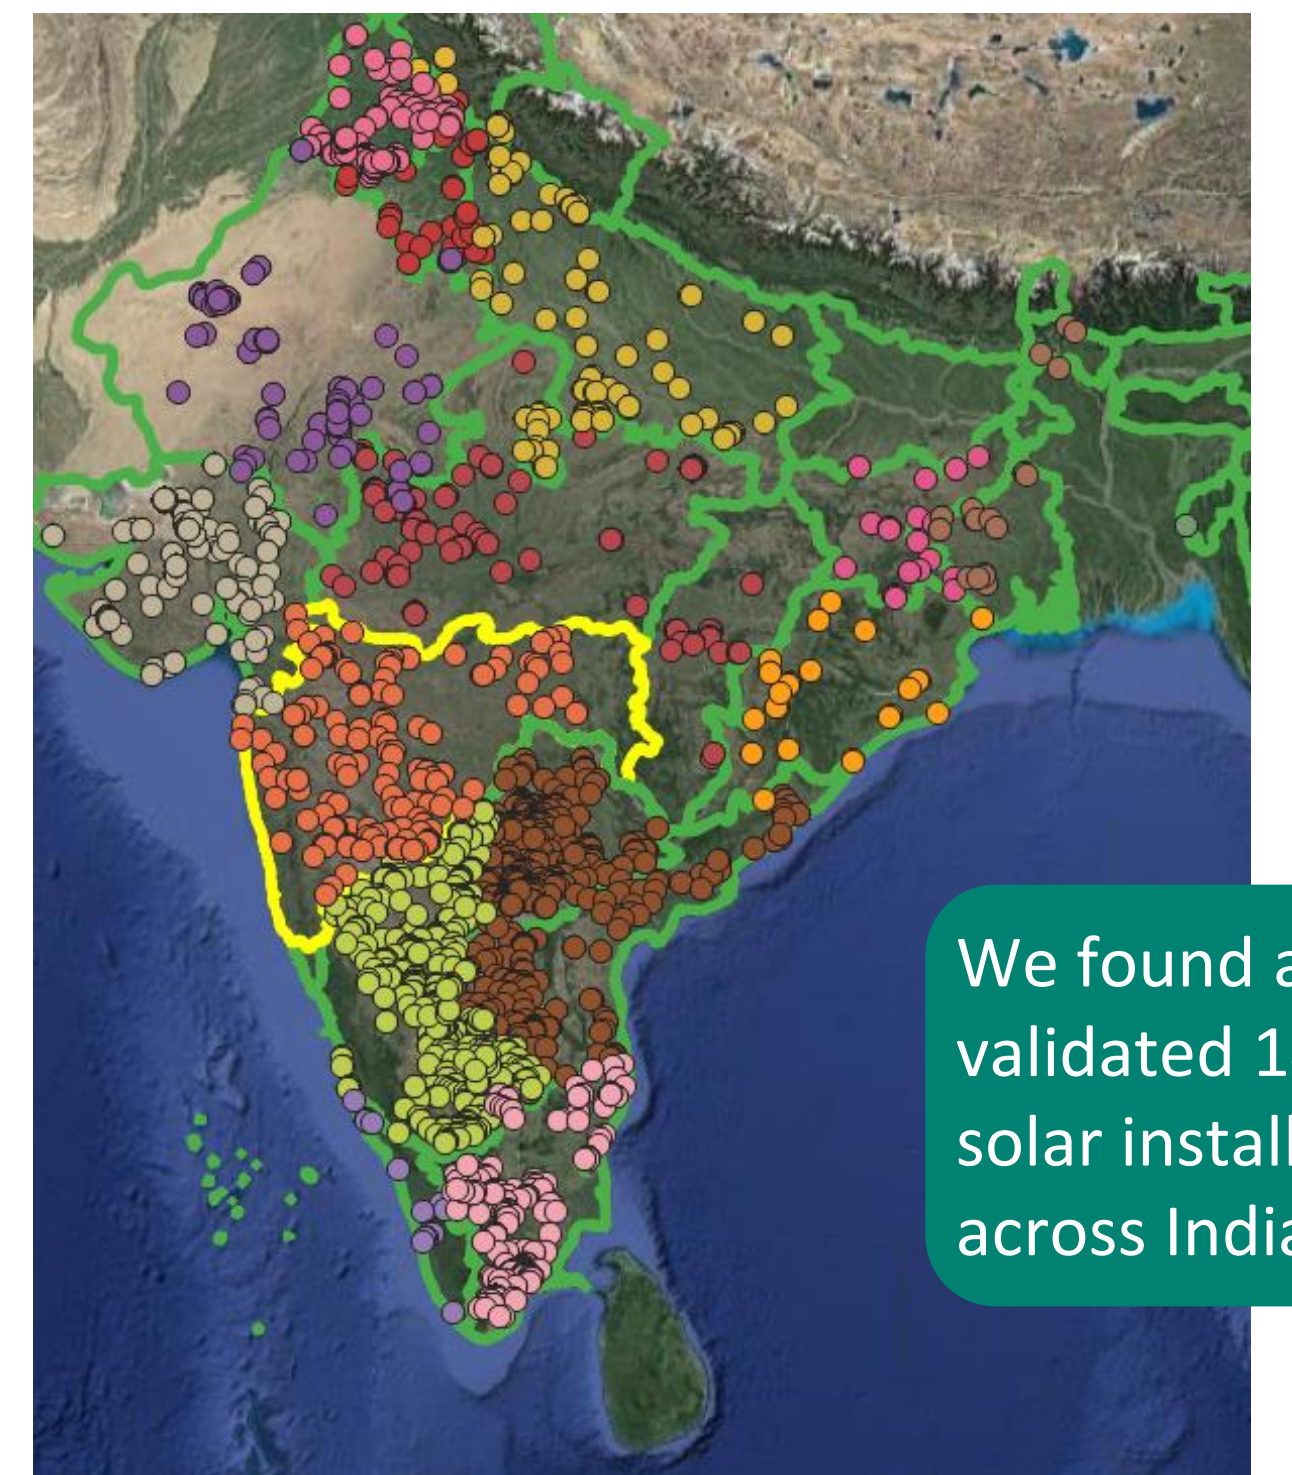
AI for Good (AI4G) is a movement within the field of AI

- **Goal:** further progress towards “good” outcomes
- Loosely guided by the UN Sustainable Development Goals (SDGs) and priorities within local communities

Our approach to AI4G

- We collaborate with *partner organizations (POs)* on AI4G projects
- POs define “good”

Case Study - Finding Solar Farms in India from Satellite Imagery



We found and validated 1422 solar installations across India

We distill our experience with AI4G projects in 11 takeaways across four areas:

Communication

- 1. Setting realistic expectations from AI**
Educating POs about AI's limits and opportunities is a core part of an AI4G project. Potentially unrealistic expectations for AI can often be reframed
- 2. Project scoping**
To ensure we develop solutions that are practically useful, project scoping needs to be an ongoing dialogue with the PO

Data

- 3. Adapting to previously collected datasets**
It is important to understand the associated metadata and collection process when using previously collected data
- 4. Dealing with subjective data annotation**
In several socially important domains, labels suffer from subjective annotation
- 5. Creating training and test sets with the application scenario in mind**
So that the model's ability to generalize to unseen instances of input is measured

Modeling

- 6. Incorporating domain expertise**
Endeavor to incorporate the PO's domain expertise in model dev
- 7. Model development with resource constraints**
Consider a project's constraints during deployment in advance
- 8. Evaluation and metrics**
Incorporate domain-specific metrics in training or determine which ML metrics are relevant
- 9. Humans in the loop**
AI4G projects require humans in the loop to some extent. Active learning enables POs to engage with modeling

Impact

- 10. Uphill path to deployment and adoption**
Maintaining deployed models requires long-term commitments. Focusing on time saved instead of pure ML metrics helps with adoption
- 11. Measuring impact**
Domain experts within POs should define mission-related impacts; work with the PO to quantify immediate (workflow or analysis enhancement) and farther-removed (mission-related) impacts

Problem

- NGO wants to select locations for solar farm construction that minimize ecological harm
- Needs to know where existing solar farms are
- Only 72 point-labels of locations of solar farms in two states were available

Takeaway 3 – working with previously collected data

Approach

- We used an interactive web app to fine-tune a model that clustered similar pixels to create weak labels
- We incorporated OpenStreetMap data to remove areas of roads, snow and water

Takeaway 6 – incorporating PO domain expertise

Paper: 


 AAAI / ACM conference on ARTIFICIAL INTELLIGENCE, ETHICS, AND SOCIETY

Takeaway 1 – Setting realistic expectations from AI

- Example:** Detecting fish species from underwater cameras with a high accuracy on rare species
- This goal likely not possible to achieve on very rare species
 - Instead, initiate data collection and labeling for rare species

Takeaway 5 – Creating training and test splits correctly

Generalize temporally



Generalize spatially

