Auditing and Designing for Equity in Government Service Allocation

Nikhil Garg, Cornell Tech (ngarg@cornell.edu)

Special thanks to Urban Tech Hub at Cornell Tech & NYC Department of Parks and Recreation

Talk goals + overview

Tell you about a cool (ecology-adjacent) application

Shared methodological challenges

- Estimation under missing data (missing species, occupancy modeling)
- Allocation of scarce on-the-ground resources

Shared “working with practitioners/having impact” challenges

- What’s the benefit of AI approaches?
- Decision-support tools vs independent computational social science

Government service allocation

Local government manages many services

- ~8k miles of streets in NYC
- ~700k trees lining streets in NYC
- Housing, sanitation, transportation, etc.

Operational tasks

- [Learning] What problems are there?
- [Allocation] Which ones to address?
- [Auditing] Did we do a good job?

Desiderata: Efficiency & Equity

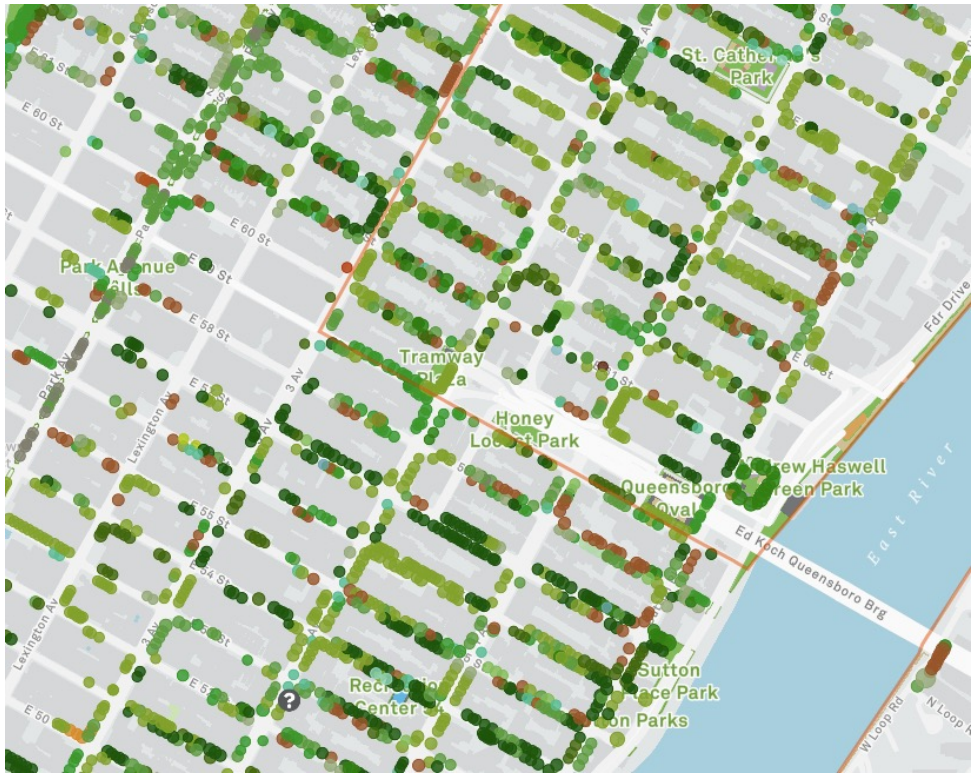

Street trees on Upper East Side in NYC

311 (crowdsourcing) systems

Cities have a phone number & app to complain to the local government

NYC's 311 system received about 2.7 million requests 2021

These are the primary way the government learns about problems

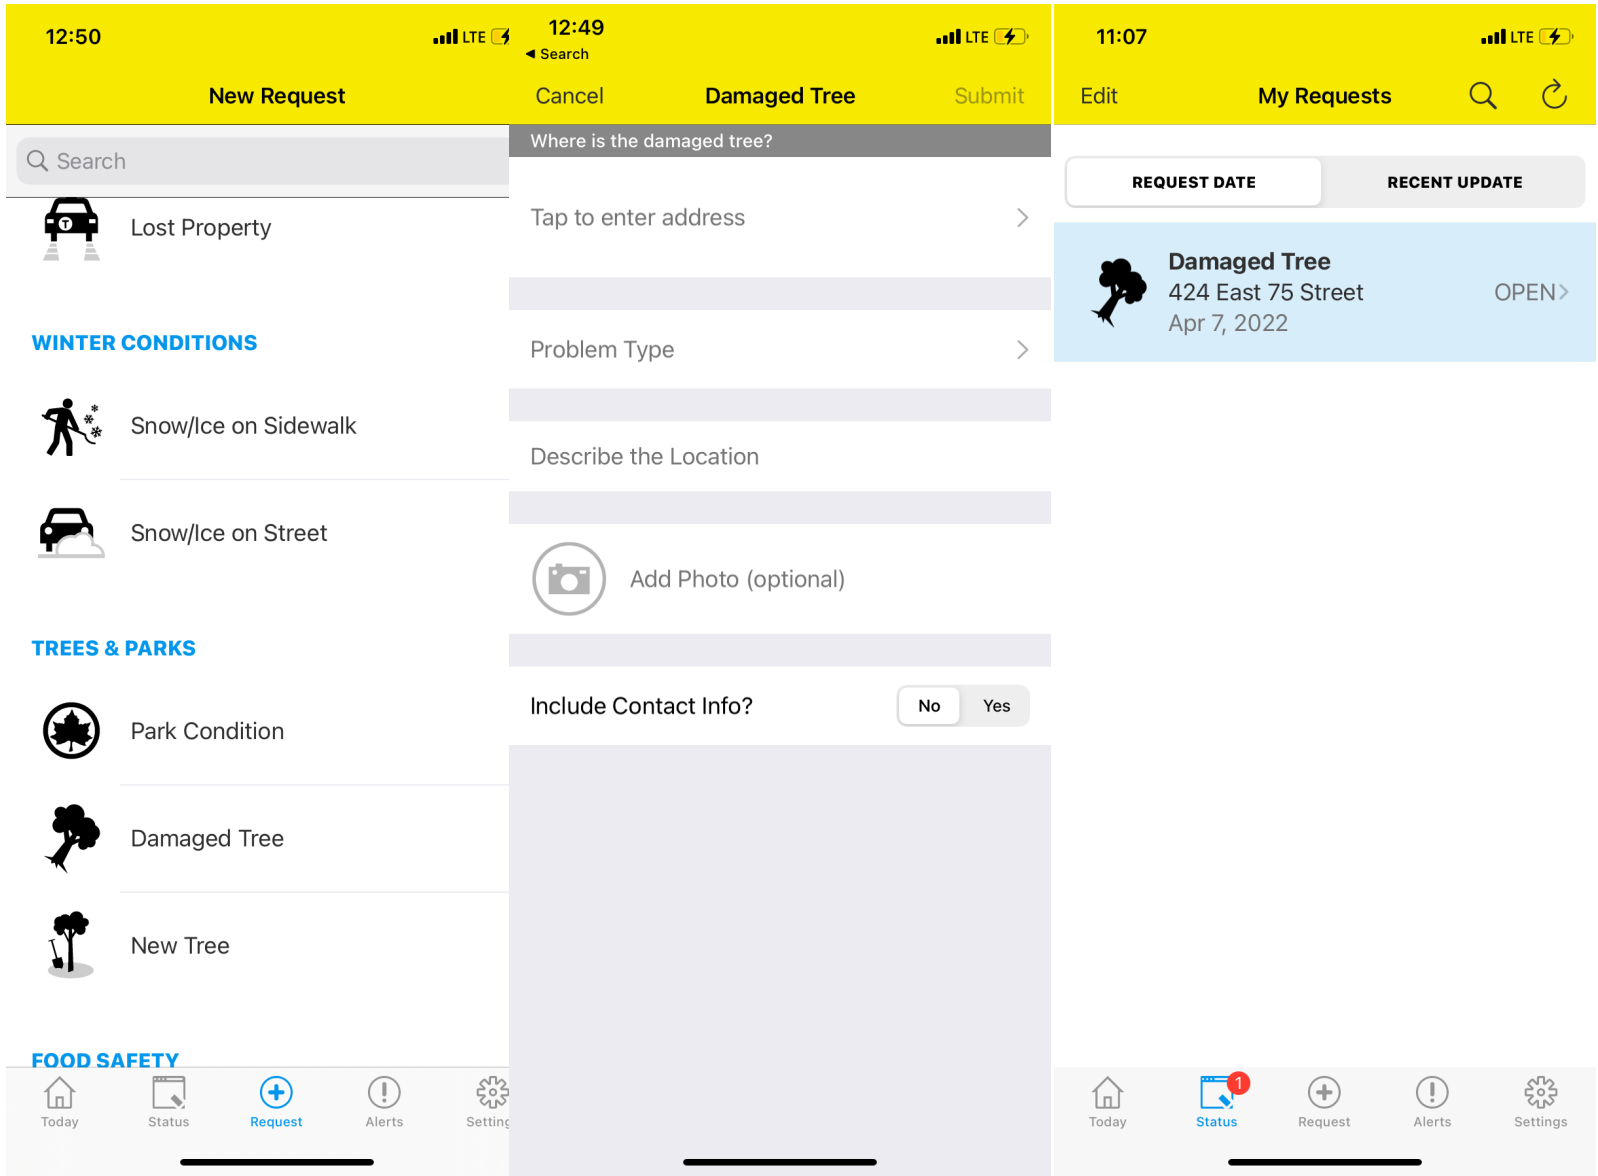

Pipeline: from incident to work orders

Incident

311 report

70-100k/year to forestry unit of NYC DPR

Inspection

~65% of reports

Work order

~50% of inspections

Why is this hard? Uncertainty, heterogeneous + strategic behavior, distribution shifts over time, capacity constraints, pipelined decisions

Research agenda: Audit and improve process's efficiency and equity

Existing collaboration: NYC Department of Parks and Recreation

Understanding reporting behavior

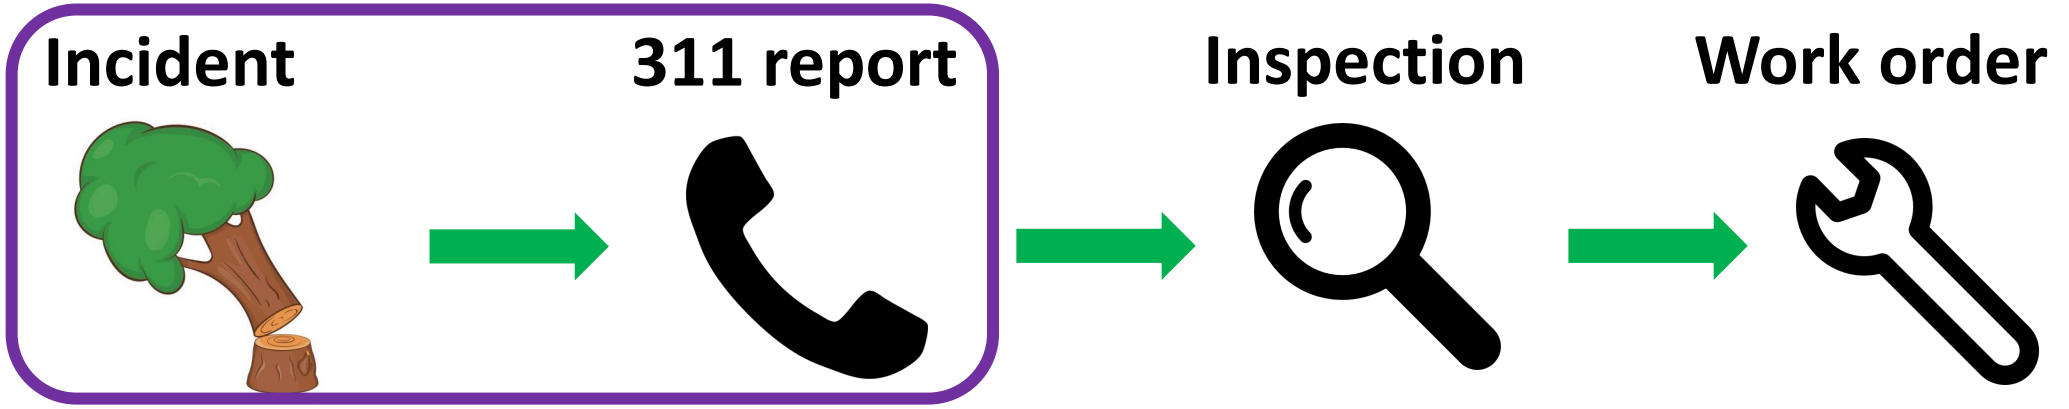

Why? If there are disparities in who reports problems, there will be disparities in what work gets done

"Equity in Resident Crowdsourcing:
Measuring Under-reporting without Ground Truth Data"
w/ **Zhi Liu** (ACM EC 2022)

Statistical challenge

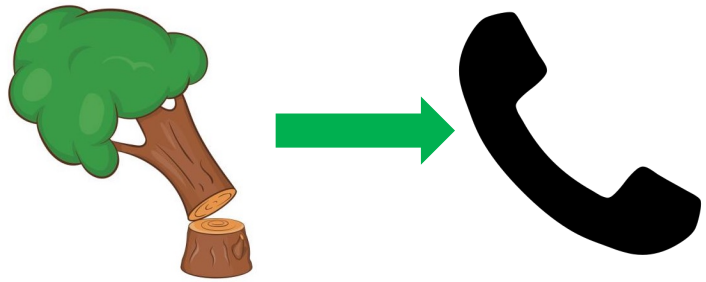

How do we distinguish between under-reporting, and some neighborhoods truly having fewer problems?

- By definition, we don't observe data on missing reports

If a tree falls in a forest, and no one reports it...
(how) does the city know about it?

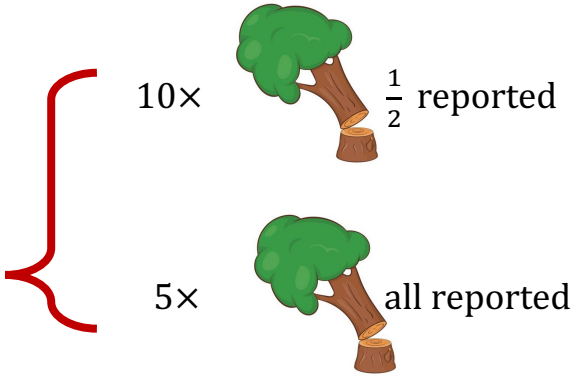

This "Benchmark" problem is a fundamental challenge across contexts

- **Policing**: crimes committed vs inequitable policing
- **Healthcare**: under-testing vs better health
- **Ecology**: recording effort vs species population

How to measuring under-reporting?

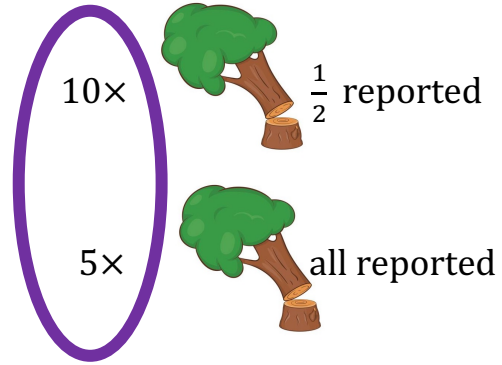

"Standard" approach: Use ground-truth data on incident rate: "how many incidents of each type (hazards, root issues, tree pruning requests…) do we expect to see in each neighborhood?"

- Go out and walk the streets and get a snapshot, uncensored view
- Construct proxy measures (number of trees, their size, species, etc)

Our question: Can we measure under-reporting, *without ground truth?*

Key idea ("Missing species"): Leverage the rate of duplicate reports about the same incident to identify the reporting rate, *given* that an incident has occurred

Temporal Setting*:* Incidents happen and are fixed; & have reporting *delays*

This work: We develop the statistical method, and then apply it to audit reporting behavior of street tree incidents over 3 years

Model + Method summary

How long does it take for an incident of type θ to be reported?

Step 1: Model under which reporting delay can be transformed into rate estimation

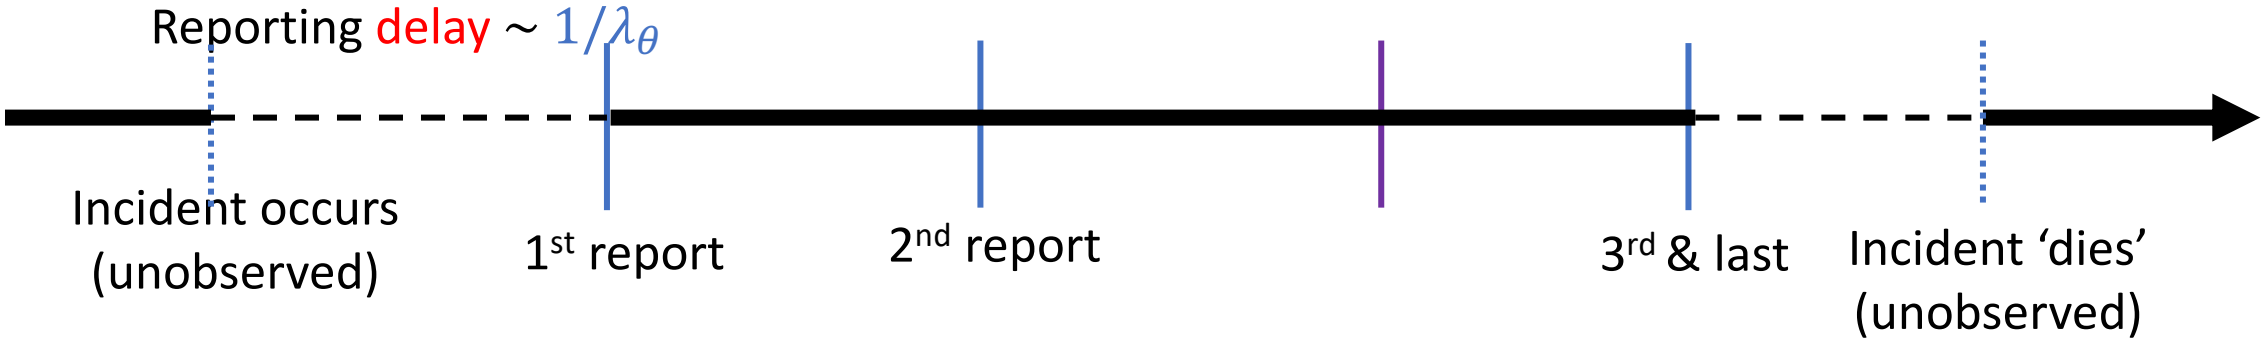

Step 2: Computationally + statistically tractable estimation

$$\#\ \text{reports(i)} \sim \text{Poisson}\left(\lambda_\theta \times (b_i - a_i)\right)$$

Spatial smoothing: ICAR Model [Morris et al. 2019]
Type θ contains an indicator for census tract (2000+ in NYC)
Then, α_k for each tract is drawn with mean of α_j of neighboring tracts

Results

Efficiency: Reporting rates higher for more urgent incidents

Equity: Reporting rates vary substantially by neighborhood

	Manhattan	Queens
High risk hazards	2.5 days	4.7 days
Medium risk tree damage	15 days	28 days
Low risk minor issue	112 days	209 days

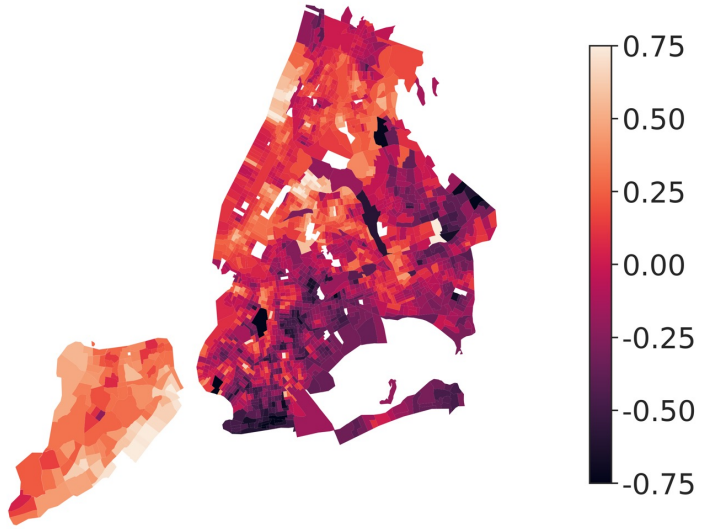

Tree agenda summary

Incident → **311 report** → **Inspection** → **Work order**

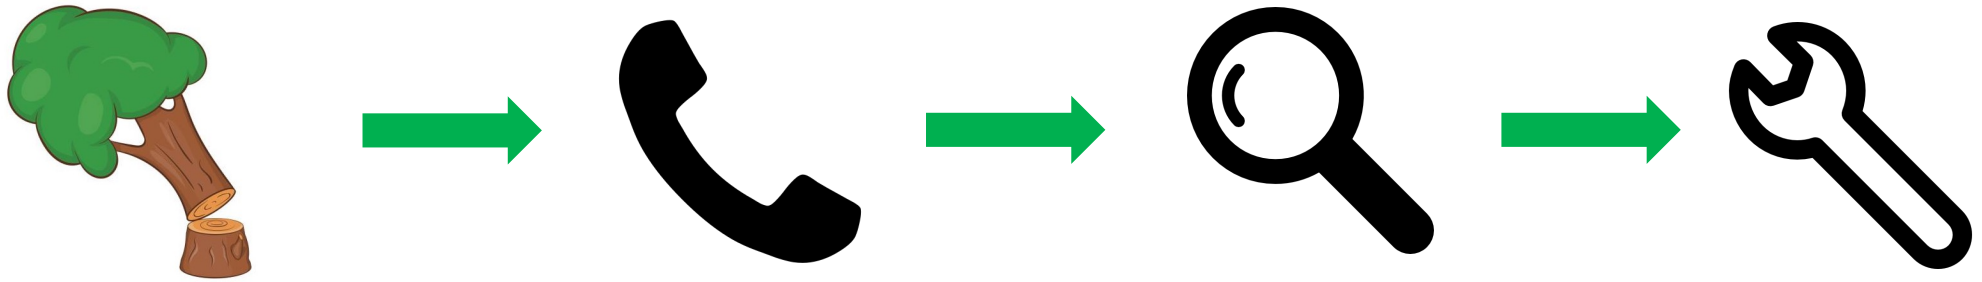

- Method to measure heterogeneous reporting behavior
  - [Today] for “public” incidents using Missing Species ideas
  - [Ongoing] for ”spatially correlated” incidents using Occupancy Modeling ideas
- [Ongoing] Method to understand agency responses
- [Ongoing] *Make* optimal allocation decisions
- [Ongoing] Software tool for the agency to understand/improve decisions
- Many more interesting theory + empirical questions!

Questions?

ngarg@cornell.edu