Time-varying, people-centric criticality of electricity distribution feeders

Hongrong Yang[1] Markus Schläpfer[1]

[1]Department of Civil Engineering and Engineering Mechanics
Columbia University, New York, USA

Abstract

Conventional assessments of electricity distribution networks rank components by engineering measures of failure, such as the energy not supplied or the number of customers interrupted. These measures are tied to meter locations and are effectively static: they do not account for the fact that the population depending on a given asset changes continuously as people move through the city. Here we ask whether the criticality of distribution feeders, measured by the number of people actually present and affected, is time-varying, and whether it departs from the conventional engineering ranking. Using large-scale mobility data to estimate the time-varying presence of people within feeder service areas, we evaluate the impact of a four-hour single-feeder outage across all possible outage start times for a representative commercial and residential feeder pair in San Francisco, drawn from a synthetic distribution-network model. The two measures disagree fundamentally. By energy not supplied, the commercial feeder is more critical at every hour of the day and the ranking never reverses. By person-hours affected, the ranking does reverse: the residential feeder affects more people overnight and into the evening, whereas the commercial feeder dominates during the working day. Each measure also departs substantially from its own time-averaged value, so that no static ranking—engineering or time-averaged human—identifies the most human-critical feeder throughout the day. These results indicate that human-consequence criticality is a distinct property of distribution assets, not reducible to load-based measures, and motivate a systematic, network-wide and predictive treatment.

USE-Lab Working Paper #2026–07–003
`https://urbansystems.civil.columbia.edu`

Time-varying, people-centric criticality of electricity distribution feeders

Hongrong Yang* Markus Schläpfer*

Abstract

Conventional assessments of electricity distribution networks rank components by engineering measures of failure, such as the energy not supplied or the number of customers interrupted. These measures are tied to meter locations and are effectively static: they do not account for the fact that the population depending on a given asset changes continuously as people move through the city. Here we ask whether the criticality of distribution feeders, measured by the number of people actually present and affected, is time-varying, and whether it departs from the conventional engineering ranking. Using large-scale mobility data to estimate the time-varying presence of people within feeder service areas, we evaluate the impact of a four-hour single-feeder outage across all possible outage start times for a representative commercial and residential feeder pair in San Francisco, drawn from a synthetic distribution-network model. The two measures disagree fundamentally. By energy not supplied, the commercial feeder is more critical at every hour of the day and the ranking never reverses. By person-hours affected, the ranking does reverse: the residential feeder affects more people overnight and into the evening, whereas the commercial feeder dominates during the working day. Each measure also departs substantially from its own time-averaged value, so that no static ranking—engineering or time-averaged human—identifies the most human-critical feeder throughout the day. These results indicate that human-consequence criticality is a distinct property of distribution assets, not reducible to load-based measures, and motivate a systematic, network-wide and predictive treatment.

1 Introduction

Urban power-grid resilience planning assumes people are where they live. But in modern cities, people spend much of their time away from home, and outages affect them wherever they are. A feeder serving an office district may support tens of thousands of workers at midday and almost nobody at night; a residential feeder follows the opposite rhythm. Conventional criticality measures are insensitive to this variation: energy not supplied (ENS), customers interrupted, and topological centrality are all defined with respect to connection points rather than to the people present at those points at the moment of failure.

The consequence is a potential structural blind spot. If asset criticality is assessed by measures that hold the population geography fixed, then prioritisation of hardening, restoration sequencing, and distributed-energy-resource placement may systematically protect the wrong assets at the wrong times. Recent work has begun to argue that infrastructure criticality should reflect human dependence rather than engineering measures alone[1], but existing approaches compute a static, time-averaged criticality for fixed service facilities from observed mobility, and therefore cannot express how the ranking of network assets changes through the day.

This paper tests the first and most basic question in that programme: is human-consequence criticality time-varying, and does it reorder assets relative to the conventional engineering measure? We address this for a representative commercial and residential feeder pair, using mobility-derived human presence to compute the person-hours affected by a fixed-duration outage as a function of when the outage begins. The single-pair setting is deliberate: it isolates the mechanism before scaling to full-network rankings, and it permits a direct, like-for-like comparison against ENS computed on the same network model.

*Department of Civil Engineering and Engineering Mechanics, Fu Foundation School of Engineering and Applied Science, Columbia University, New York, NY, USA.

We find that the two measures not only differ in magnitude but disagree in ordering, and that the disagreement is systematic rather than incidental: the commercial feeder is more critical by ENS at every hour, while by human consequence the two feeders exchange rank repeatedly over the course of a day.

2 Data and methods

2.1 Distribution network model

Feeder-level analysis requires a realistic representation of distribution topology and service areas. We use the SMART-DS synthetic distribution-network dataset for San Francisco[2], which provides validated OpenDSS models with realistic feeder topologies, load allocations, and capacity ratings.

2.2 Mobility-derived human presence

Time-varying human presence is estimated from large-scale, privacy-preserving GPS mobility data covering March 2023, provided by CityData.ai.

2.3 Outage impact measures

For each feeder f and outage start time t_0, we evaluate a single-feeder outage of fixed duration $T = 4$ h under two measures.

The engineering measure is the *energy not supplied*,

$$\mathrm{ENS}_f(t_0) = \int_{t_0}^{t_0+T} P_f(t)\,\mathrm{d}t, \tag{1}$$

where $P_f(t)$ is the load served by feeder f at time t, given by the SMART-DS allocated load shapes.

The human-consequence measure is the *person-hours affected*,

$$\mathrm{PH}_f(t_0) = \int_{t_0}^{t_0+T} N_f(t)\,\mathrm{d}t, \tag{2}$$

where $N_f(t)$ is the estimated number of people present within the service area of feeder f at time t. Both quantities are evaluated for all start times t_0 across the day, yielding an impact profile rather than a single number, and both are compared against their own time-averaged values.

A fixed outage duration is used so that the two measures differ only in what they count, not in the event they describe.

3 Results

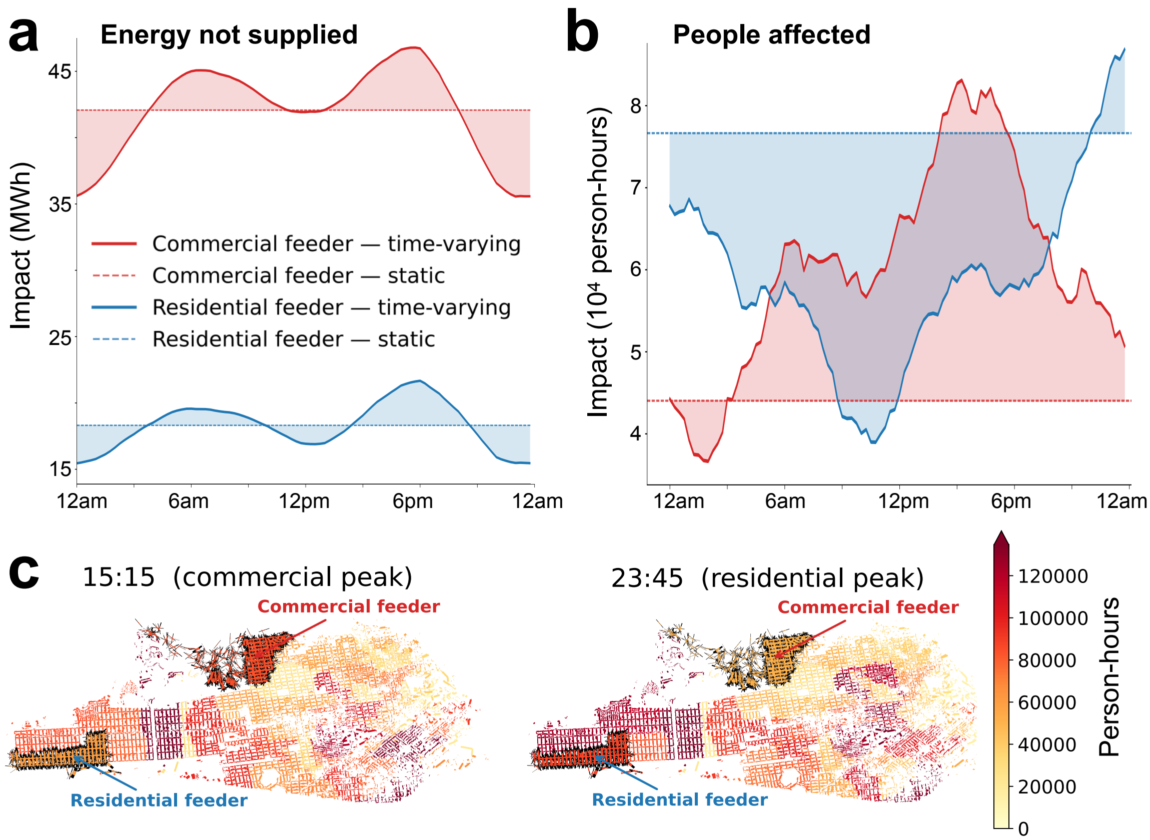

Figure 1: **Human-consequence criticality is time-varying and diverges from the engineering measure.** Impact of a four-hour outage of a commercial and a residential feeder in San Francisco (SMART-DS), by outage start time. **a**, By energy not supplied, the commercial feeder is more critical at every hour. **b**, By person-hours affected, derived from mobility-based human presence, the ranking reverses: the residential feeder dominates overnight and in the evening, the commercial feeder during the working day. Each curve crosses the other and its own time-average (dashed). **c**, Spatial distribution of person-hours at the commercial (15:15) and residential (23:45) peaks.

3.1 The engineering measure gives a fixed ordering

Evaluated by energy not supplied, the commercial feeder is more critical than the residential feeder at every outage start time (Fig. 1a). Both profiles vary through the day, following the diurnal shape of their respective load curves, but the two curves do not intersect: the ordering established by peak load is preserved across all start times. An assessment based on ENS would therefore prioritise the commercial feeder unconditionally, and would never flag the residential feeder as the more critical of the pair.

3.2 The human-consequence measure reverses the ordering

Evaluated by person-hours affected, the ordering is not fixed (Fig. 1b). The residential feeder affects more person-hours overnight and into the evening, when residents are at home, whereas the commercial feeder dominates during the working day, when the daytime population concentrates in the commercial district. The two curves cross repeatedly over the course of the day.

Each curve also departs substantially from its own time-average (Fig. 1b, dashed lines). A time-averaged human-consequence measure—the natural static summary of the mobility-based quantity—therefore misstates the impact at most hours of the day, understating it during each feeder's peak period and overstating it otherwise. The consequence is that no static ranking, whether engineering or time-averaged human, identifies the most human-critical feeder throughout the day.

3.3 Spatial structure of the affected population

The spatial distribution of affected person-hours at the two peak times (Fig. 1c) shows that the reversal reflects a genuine redistribution of people rather than a change in a single aggregate. At the commercial peak (15:15), person-hours concentrate within the commercial service area; at the residential peak (23:45), the concentration shifts to the residential area.

4 Discussion

The two measures do not merely differ in scale: they disagree in ordering, and they do so systematically. Because the number of people served per unit of load varies with land use and time of day, a feeder can rank far higher by human consequence than by energy not supplied. This makes human-consequence criticality a distinct property of a distribution asset rather than a rescaling of a load-based one, and it identifies a class of assets—those serving populations that are present at times when their load is comparatively low—whose societal importance is systematically understated by conventional assessment.

Three implications follow. First, criticality should be reported as a function of time and operating condition rather than as a single score, since the identity of the most critical asset changes through the day. Second, the divergence between the two measures is itself a quantity of interest: characterising which feeders, at which times, are most severely misjudged provides a direct target for resilience investment. Third, because human presence is estimated from mobility rather than from meter locations, the same framework extends naturally to conditions in which the population geography changes—during extreme events, or following structural interventions in the built environment that reorganise where people go.

A natural extension concerns not only how many people are present, but where they are relative to home. Presence counts treat every affected person identically, yet the consequence of an outage plausibly depends on a person's displacement: someone caught in a commercial district several kilometres from home during an evening outage faces disrupted transit, impaired wayfinding, and a curtailed ability to return, whereas someone at home faces a materially different situation. Mobility data resolve this directly, since the distance between a person's current location and their inferred home location is observable at the same temporal resolution as presence itself. Weighting affected person-hours by displacement would yield a second dimension of human consequence—one that captures exposure to disrupted mobility rather than exposure to the outage alone, and that would be expected to reinforce the daytime criticality of commercial and transit-served feeders while attenuating the night-time criticality of residential ones. More broadly, it suggests viewing the city as pulsating between a dispersed daytime state, in which a large fraction of the population is displaced from home and therefore mobility-dependent, and a concentrated night-time state in which it is not. How that pulsation interacts with the criticality reversal reported here is left to future work.

More generally, this proof of concept establishes the central mechanism underlying mobility-informed infrastructure resilience: infrastructure criticality depends not only on engineering characteristics but also on the dynamic distribution of people. Extending this framework from individual feeders to network-wide assessments, and from observed to predicted mobility under future urban change, is a natural next step toward people-centric infrastructure planning.

5 Limitations and ongoing work

This study establishes the mechanism for a single feeder pair in one city, and several extensions are required before the result can be treated as general.

The analysis compares two feeders selected to represent contrasting land-use types; it does not yet produce a network-wide ranking, and the magnitude of the reordering across a full distribution system remains to be quantified. The presence estimates inherit the representativeness limitations of passively collected mobility data, and the extent to which they misstate the population of particular service areas has not been fully characterised. The outage scenario is deterministic, of fixed duration, and considers single-feeder failures only; realistic assessment requires probabilistic failure models, particularly under stress conditions such as extreme heat, where elevated failure likelihood and concentrated human presence may coincide.

Finally, person-hours count people but do not distinguish among them. Incorporating demographic and social-vulnerability indicators would allow the *severity* of the consequence to be assessed alongside its magnitude, which is likely to matter most where dependent populations can least cope—for example, heat-vulnerable residents at night.

These extensions form part of a broader research programme on mobility-informed infrastructure resilience, in which time-varying human presence—predicted rather than only observed—becomes the basis for evaluating both the adequacy of infrastructure and the human consequence of its failure, before structural changes to the built environment reshape urban mobility.

References

[1] Junwei Ma, Bo Li, Xiangpeng Li, Chenyue Liu, and Ali Mostafavi. Quantifying functional criticality of lifelines through mobility-derived population-facility dependence for human-centered resilience. *Sustainable Cities and Society*, 147:107547, 2026.

[2] Bryan Palmintier, Tarek Elgindy, Carlos Mateo, Fernando Postigo, Tomás Gómez, Fernando de Cuadra, and Pablo Duenas Martinez. Experiences developing large-scale synthetic u.s.-style distribution test systems. *Electric Power Systems Research*, 190:106665, 2021.