

Institutions, Democracy and Peace

New ways to measure with remote sensing

Hannes Mueller

IAE (CSIC), BSE, CEPR

ReCIPE Virtual Course

Key Concepts for Economies in Conflict and Fragile Settings – Part 2

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

Why we should care about institutions


Two parallel global trends since the end of the Cold War:

- Democracy in decline – autocratization, weakening checks and balances (Lührmann et al. 2022; V-Dem 2023)
- Armed conflict on the rise again after a relatively peaceful interlude (ACLED 2024; Rustad 2024)

Robust institutions have historically been the main safeguard against violence: they shape incentives, allocate resources, and determine power dynamics.


Recent overview article Mayoral and Mueller (2026, Economic Policy) <https://doi.org/10.1093/epolic/eiaf004>.

Democracy and armed conflict: the basic correlation



Source: Binned scatter, Share of months in armed conflict (UCDP, > 0.5 fatalities/million inhab.) vs. V-Dem Egalitarian Democracy Index, 1989–2023.

A framework for the channels



The empirical problem

The framework gives us many candidate channels. The hard part is *measuring* them:

- **State capacity:** fiscal and legal reach – often invisible from administrative data in the very places we care about (fragile states).
- **Public services:** who gets services, who doesn't – official numbers are often biased or absent.
- **Violence and its impact:** conflict event datasets (ACLED, UCDP) rely on news reports and eyewitnesses – biased, lagged, incomplete.

This lecture: remote sensing + machine learning gives us a new instrument to look at these things directly. I'll go through three of my own projects.

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

The starting point: nightlights

Henderson, Storeygard and Weil (AER 2012), *Measuring Economic Growth from Outer Space*:

- Take DMSP-OLS satellite imagery of night-time lights
- Aggregate to country, sub-national, or grid-cell level
- Show it correlates strongly with GDP growth
- For countries with poor national income accounts: optimal estimate puts roughly equal weight on lights and conventional GDP

This opened a huge agenda: nightlights as a proxy for economic activity at any spatial scale.

From aggregation to extraction

Modern machine learning lets us get *much more* out of an image than a brightness count because it allows us to understand features that are a constellation of many pixels:

- **Object detection:** buildings, roads, vehicles, crops
- **Texture/pattern classification:** road quality, agricultural type, land use, informal vs. formal settlement

With these tools, satellite imagery can speak to the very institutional channels in the framework – state capacity, political exclusion, public service provision – where standard data fails us most.

Three illustrations from my own work follow.

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

Why building destruction?

Mueller, Groeger, Hersh, Matranga, Serrat (PNAS 2021).

Building destruction during war:

- Particularly harmful form of violence (urbicide – knock-on effects on water, power, hospitals)
- Used strategically to displace populations
- Existing data: UNOSAT manual annotation, eyewitness reports – scarce, lagged, biased

Goal: an automated, satellite-based classifier that produces destruction data with unprecedented scope, resolution, and frequency.


Setting: Syrian civil war, six major cities, 2011–2017.

The data setup

- **Imagery:** 64 high-resolution (50cm) Google Earth / Maxar images across 6 Syrian cities, 2011–2017
- **Labels:** UNITAR / UNOSAT building destruction annotations (point coordinates)
- Cut images into 64×64 pixel “patches” ($\approx 1,024 \text{ m}^2$) – this is the unit of observation
- Aleppo: 22 images $\times \approx 95,000$ patches $\approx 2.1\text{M}$ patch-period observations


The central difficulty: *extreme class imbalance*. Even in heavily affected Aleppo, only about 2.8% of populated patches were classified as destroyed in September 2016. We are looking for needles in a haystack.

Aleppo: the haystack and the needles



Red dots: UNOSAT destruction annotations. Yellow: populated areas under analysis.
Magenta: non-civilian no-analysis zones.

Aleppo: zoomed in



Red dots: UNOSAT destruction annotations.

Why naive classifiers fail under class imbalance

Big problem with remote sensing is often *class imbalance*.

A worked example to fix ideas: imagine 100,000 image patches with 1,000 destroyed.

Classifier with 15% false-positive and 90% true-positive rate produces:

- True positives: $0.90 \times 1,000 = 900$
- False positives: $0.15 \times 99,000 = 14,850$
- Precision = $\frac{900}{900+14,850} \approx 5.7\%$

Conditional on predicting “destroyed”, model is wrong 94% of the time! The *same* classifier achieves 86% precision on a balanced sample.

Measurement error is huge when underlying data is imbalanced.

Balanced-sample metrics from the ML literature are misleading.

Report performance on the unbalanced sample!

Don't trust articles that do not.

Approach: a convolutional neural network (intuition)

A CNN looks at an image patch and learns to spot visual features of destruction – rubble, collapsed roofs, bomb craters – through layers of learned filters.


- **Input:** a $64 \times 64 \times 3$ (RGB) patch + the same area pre-war as reference (*change detection*)
- **Hidden layers:** convolutions that detect edges, textures, then larger destruction-specific patterns
- **Output:** a continuous score in $[0, 1]$ – our predicted probability of destruction

We experimented with several architectures and ended up with a relatively *flat* CNN – deeper networks overfit given the limited number of labels.

Three methodological contributions

1. **Label augmentation.** Destruction *persists* – once a building is destroyed, it stays destroyed (during the war). We use this to propagate labels forward and backward in time. Boosts our usable labels from $\sim 50,000$ to $\sim 2.3\text{M}$ (51k destroyed).
2. **Two-stage classifier.** Feed the raw CNN scores into a *random forest* that uses spatial leads/lags (neighboring patches) and temporal leads/lags (same patch at other dates) as features. Exploits the fact that destruction is spatially and temporally clustered.
3. **Honest evaluation on imbalanced samples.** We report precision-recall on the empirical (unbalanced) distribution, not on artificial 1:1 splits.

Cross-validation: avoiding overfitting



Naive cross-validation would put the *same patch at different dates* into both training and test – the model would memorize locations, not generalize.

Our fix: split by patch ID – 70% of patches in training, 30% reserved for testing, *across all time periods*. All performance numbers come from the held-out test patches.

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

From measurement to economics

Groeger and Mueller (CEPR DP 21059), *War Destruction, Commercial Relocation and Urban Recovery*.

Now that we can *measure* destruction at fine spatial resolution, what can we learn about post-war recovery?

Aleppo is a particularly informative case:

- Syria's pre-war commercial capital ($\sim 2.5\text{M}$ people)
- Concentrated, neighbourhood-level destruction (2012–2016)
- Final phase of the *Battle of Aleppo* (late 2016) generated substantial additional damage *after* the last UNOSAT manual assessment – a gap our ML method fills
- Eight years of post-conflict data through 2025

The within-city question

The WWII city-size reversion literature (Davis-Weinstein, Brakman et al.) finds that bombed cities tend to recover in aggregate. *Within* cities, though, the picture is more nuanced.

We ask:


1. Where does economic activity go after destruction? Within neighbourhoods or across them?
2. Does it return to previously active locations (*place-based persistence*)?
3. Or does the city settle into a new spatial configuration (*re-optimization*, coordination failures)?
4. Are these dynamics different across sectors – commerce vs. schools vs. places of worship?

The data

A high-resolution spatial panel of Aleppo 2013–2025:

- **Destruction:** 22 dated satellite images, combined UNOSAT labels + our PNAS-style ML predictions, harmonized to a $\sim 30\text{m}$ grid. $\approx 95,746$ cells.
- **Economic activity:** OpenStreetMap full-history POI data (Raifer et al. 2024) – end-of-year snapshots of commercial and amenity tags, annually, 2013–2025.
- **Independent proxies, free of OSM contributor bias:** VIIRS nighttime lights (500m), Global Human Settlement Layer built-up area (100m).
- **Pre-war political alignment:** neighbourhood-level regime control from AAAS (2013) satellite imagery analysis – 37% regime-aligned, 63% non-regime.

The destruction of Aleppo



Source: Groeger and Mueller (2025). Figure 1. Predicted destruction in February 2026

Empirical strategy: destruction effects

For each cell i in year t , with outcome Y_{it} (POI count, NTL, built-up area):

$$Y_{it} = \beta_1 d_{it} + \beta_2 sd_{n(i)t} + \alpha_i + \delta_t + \varepsilon_{it}$$

- d_{it} : time-varying cell destruction indicator
- $sd_{n(i)t}$: share of cells in cell i 's neighbourhood destroyed at time t – captures **spillovers**
- α_i, δ_t : cell and year fixed effects

Identification. Conditional on neighbourhood, the granular geography of where weapons impacted is plausibly orthogonal to cell-specific economic shocks. Robustness with 300m cluster $\times$ year FE absorbs essentially all variation outside small geographic areas.

Main destruction effects

VARIABLES	(1) Total commercial	(2) Total amenities	(3) Nighttime light	(4) Total commercial	(5) Total amenities	(6) Nighttime light
Cell destroyed	-0.530 (0.120)	-0.704 (0.075)	-0.327 (0.028)	-0.353 (0.131)	-0.350 (0.083)	-0.046 (0.022)
Neighborhood share destroyed				-3.526 (1.276)	-7.034 (1.633)	-5.581 (0.479)
Observations	670,222	670,222	670,222	670,222	670,222	670,222
R-squared	0.789	0.764	0.662	0.789	0.765	0.667
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Unit of observation are 30m cells, 2013–2019. Dependent variables divided by sample mean. SEs clustered at 300m cluster cells.

- **Cell destruction:** –50% of mean for commercial tags, –70% for amenities, –30% for nighttime lights.
- **Neighbourhood share destroyed:** large negative coefficients – strong spillovers. For NTL, neighbourhood effect dominates (–85% of the cell coefficient absorbed).
- In a destroyed cell of a neighbourhood with 20% destruction $\Rightarrow$ effects close to –100% of mean.

Recovery dynamics: event study

For each cell, $D_i = 1$ if ever destroyed during the conflict, 0 otherwise:


$$Y_{it} = \sum_{\tau \neq 2014} \beta_{\tau} (D_i \times 1\{t = \tau\}) + \alpha_i + \delta_t + \varepsilon_{it}$$

Normalized to 2014 (last pre-conflict year of OSM coverage).
Standard errors clustered at 300m cell clusters.


This traces out the full recovery dynamics through 2025 – *eight years after combat ended* in Aleppo.

What we're looking for: do destroyed cells converge back to non-destroyed cells? Or does the gap persist?

Event study: amenities and commercial activity



(a) Total amenities



(b) Total commercial

Coefficients on $(\text{ever-destroyed}) \times \text{year}$, normalized to 2014. Cell and year FE; clustered at 300m cell cluster.

2013 coefficients are near zero – no apparent pre-trend. Sharp, *persistent* decline starting 2014–2015, no convergence back even in 2025.

Sectoral divergence: who comes back?


The aggregate decline masks striking sectoral heterogeneity. Decomposing into within- and between-neighbourhood components (Kitagawa-Das Gupta), share of activity in destroyed cells, 2014 → 2025:

	ΔS	Within	Between
Restaurants	−79%	−47%	−33%
Supermarkets	−92%	−46%	−45%
Schools	−20%	+21%	−41%
Places of worship	−15%	+20%	−35%


Interpretation: commerce coordinates on *new* clusters (agglomeration externalities lock in new nodes). Community amenities track *population* back to where it used to live.

This is the headline finding – and one that aggregate proxies (nightlights, total city counts) entirely miss.

Heterogeneity by pre-war political alignment



(c) Total amenities



(d) Total commercial

Event study estimated separately for Assad regime aligned (40 neighbourhoods) and non-regime (69 neighbourhoods) cells.

From 2022 onwards, regime-aligned cells partially recover towards the pre-2014 baseline, especially for commercial activity; non-regime cells continue to decline. CIs overlap and the classification is coarse – we read this as *suggestive*, not conclusive.

Take-away: the political economy of urban recovery

What the fine-grained data lets us do that aggregates can't:

- Distinguish *within*- and *between*-neighbourhood reallocation
- Separate sectors with different spatial logics (agglomeration vs. catchment-area)
- Cross-validate OSM with NTL and built-up area – catching mapping-intensity biases
- Compare across political-economy categories (regime-aligned vs. not)

In Aleppo specifically the regime-alignment evidence is *mild* – the war ended with one side firmly in control of the entire city, so the contrast is muted.

Where this is really exciting: take this methodology to a

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

Road Investment and Violence in DRC: Perishable Peace Dividends

Lebrand, Mueller, Schouten, Steinbuks (BSE WP).

The Democratic Republic of Congo:

- One of the most protracted humanitarian crises in the world since 1997
- Road network “152,400 km long, but for a large part only exists on maps” (WB 2017)
- Sustained donor investment in road rehabilitation since 2003 – explicit peacebuilding objective

Question: did these road investments actually reduce violence?
And if so, for how long?

The theory pulls in opposite directions

Roads $\Rightarrow$ less violence

- State penetration of “ungoverned” zones (Herbst 2000; Müller-Crepon et al. 2021)
- Stationary banditry – armed groups settle and tax rather than raid (Sánchez de la Sierra 2020)
- Opportunity costs: trade and jobs raise the cost of fighting

Roads $\Rightarrow$ more violence

- Lower transport costs of violence (Zhukov 2012)
- Remoteness as sanctuary (Ali et al. 2015, DRC; Nunn-Puga 2012)
- “Dark side of proximity” (Mueller-Rohner-Schönholzer 2022)

And: roads *decay*. Standard datasets ignore this. A road today is not the same road five years on.


A new database of road projects in DRC

A key contribution: assembling the first comprehensive geospatial panel of large road investments in DRC, 2003–2022.

1. Collect from multiple sources: academic mapping (William & Mary, Boston U, JHU AidData), government (Plateforme de Gestion de l'Aide), donors (WB, AfDB).
2. Geolocate each project – often from project documents, news, and triangulation.
3. Build a panel at the PRIO 55km cell-year level.

Starting from 184 projects $\Rightarrow$ 52 intercity / large rural projects with full information and geolocation, generating **245 treated PRIO cells** (avg. project crosses ~ 4 cells).

Unique road dataset



World Bank collaboration gave us detailed road beginning and completion dates.

Why we need road *quality*, not just completion

A road project has a start date and an end date in our database. But what we really want to know is the actual *state of the road* in any given year.

Why this matters:

- Tropical conditions + lack of maintenance $\Rightarrow$ unpaved roads degrade fast. Bachmann (2018), Schouten (2022): rehabilitated roads in DRC routinely deteriorate within years
- If road quality is what matters for violence (security, trade, state reach), then a binary “completed” dummy is the wrong variable

Goal: use remote sensing to recover a time-varying road *quality* measure, and estimate the rate of decay.

Measuring road quality from satellite images


Building on Muhebwa et al. (2023), a deep-learning classifier trained on road quality labels from Kenya and Liberia (International Roughness Index, IRI):

- Input: 256×256 pixel satellite tiles at 60cm/pixel
- Output: road quality class per tile (“good” / “bad”, plus 5-class option)
- For DRC: scanned three large rehabilitation projects in the East, repeatedly over time

Aggregation. Project roads are broken into 10km segments; tiles within a segment are aggregated, then segments are binned by year. This removes most measurement noise (we find no significant seasonal variation).

One project, decomposed into segments




Project 04 Segments



Project 04 Segments (Zoomed)



Visual: the same road over time



Estimating the decay factor

For each segment, fit an exponential decay from the post-rehabilitation peak:

$$y_t = a_0 \cdot \delta^t$$

Aggregating across segments (weighted by number of images):

$$\hat{\delta} = 0.735 \pm 0.05$$

What this means: after 3 years, only $\approx 40\%$ of the rehabilitated road quality remains. 60% is “unusable”.

This gives us a continuous treatment variable. Instead of a 0/1 “project completed” dummy, we can compute a *road stock*.

Empirical strategy: event study at the PRIO cell level


PRIO grid is $0.5^\circ \times 0.5^\circ$ ($\approx 55\text{km}$) cells – standard in conflict economics.

Event-study via Callaway and Sant'Anna (2021):


- Treatment: completion year of a road project in a given cell
- Control: never-yet-treated cells
- Outcomes: log UCDP events, log UCDP fatalities
- Standard errors clustered at the PRIO cell

Identification: the timing of *completion* in any one of the (avg.) 4 cells a project crosses is plausibly orthogonal to local violence shocks in that specific cell. Robustness: matched event study within province-year.

Event study: violence around road completion



(h) UCDP events



(i) UCDP fatalities

- Statistically significant drop in violence at completion: -6% for events, -10% for fatalities
- **But:** violence *rebounds* to pre-treatment levels within 3 years – the peace dividend is perishable
- Mild positive pre-trend visible – works themselves may attract or generate violence

Main regression: completion vs. stock

	(1)	(2)	(3)	(4)	(5)	(6)
	events	events	events	fatalities	fatalities	fatalities
having an ongoing project	-0.025 (0.016)			-0.019 (0.032)		
having a completed project		-0.046** (0.022)			-0.059 (0.045)	
road project stock			-0.031** (0.015)			-0.060** (0.027)
Observations	18722	18722	18722	18722	18722	18722
R2	.402	.402	.403	.348	.348	.348

Source: Lebrand, Mueller, Schouten, Steinbuks (2024), Table 4. PRIO cell-year panel; cell and year FE. SEs clustered at cell.

Three treatment specifications:

- **Ongoing:** dummy if project in progress. Weak / insignificant.
- **Year of completion:** ~ 5 pp reduction in events.
- **Road stock (decay-adjusted):** significant for *both* events and fatalities. The decay-adjusted measure is the only one with robust effects across both outcomes.

Discussion: which channels?

What's the mechanism? Most likely several at once.

1. **State capacity at a distance.** Roads let the central state deploy security and project authority. This is the deterrence channel of Section 3 of the survey. Practitioners we spoke with at the WB identify this as primary in DRC.
2. **Trade-based pacification.** Roads let communities interact through markets, not guns. Consistent with Rohner-Thoenig-Zilibotti.
3. **Rapacity vs. stationary banditry.** Roads also lower the cost of predation. We find roads + gold → *more* violence; roads + cobalt → less. Sign depends on local economic structure.

But the most robust finding is about *durability*, not channels: even where roads reduce violence, the effect dies as the road dies. State capacity that doesn't maintain its infrastructure isn't, in practice, present.

Policy take-away: the perishable peace dividend

- Road rehabilitation produces a real, measurable peace dividend in DRC – 5–10 pp reduction in violence around completion.
- But the dividend *perishes*: rebounds to baseline within ~ 3 years.
- Implication for donors: investing in initial rehabilitation without maintenance is investing in a 3-year peace.
- Two options for durable impact: **maintenance**, or upfront investment in more **durable** road types (paved roads have ~ 20 -year lifespans without maintenance).

This is exactly the kind of insight that becomes visible only with the satellite-derived road-quality measure. The standard “project completed” dummy literature would have missed it entirely.

Setting the stage: institutions and conflict

Remote sensing: from nightlights to deep learning

Application 1: Monitoring destruction from space

Application 2: Destruction and urban recovery – Aleppo

Application 3: Road investment and violence – DRC

Wrap-up

What remote sensing buys us

- **Coverage.** Places where conflict-event data is unreliable (rebel-held areas, remote provinces) become observable.
- **Resolution.** Patch- or pixel-level outcomes, not just country-year aggregates.
- **Frequency.** Closer to real-time monitoring – weeks, not years.
- **New mechanisms.** Destruction, infrastructure decay, neighbourhood-level recovery – proxies for state capacity and political inclusion that we just didn't have before.

But: precision is hard with class imbalance; labels are scarce; domain shift across cities and countries is real. ML choices matter and have to be validated.

Back to the survey

We started with: institutions → conflict, with channels that are easy to write down and hard to measure.

Three concrete steps closer:

- **PNAS 2021:** a transferable method for measuring war destruction.
- **Aleppo paper (CEPR DP 21059):** destruction → sectoral divergence in urban recovery; commerce relocates, community amenities return; political alignment shapes the trajectory.
- **DRC roads (BSE WP):** infrastructure → state capacity → violence, but the dividend is *perishable*.

Remote sensing might allow to check public service delivery.

Open research directions

- Transfer of destruction classifiers across conflicts (Ukraine, Gaza, Sudan) – domain shift is the binding constraint (MONWADES)
- Combining remote sensing with newspaper text and survey data – different biases, different blind spots.
- **Cross-city political exclusion** (à la Burgess et al 2015) using destruction and recovery data – a natural extension of the Aleppo work.
- Real-time monitoring as input to early warning – ViEWS, conflictforecast.org.
- Ethical limits: when does monitoring *reduce* violence and when does it provide targeting information to perpetrators?

Thank you

Questions and comments very welcome.

`h.mueller.uni@gmail.com`

`www.hannesfelixmueller.com`

`econai.org`

Appendix: PNAS performance by city

City	1st-stage prec.	+ spatial	+ spatio-temporal	AUC
Aleppo	16.1	16.9	35.7	91.5
Daraa	4.2	4.6	11.7	89.0
Deir-Ez-Zor	11.0	12.1	21.7	80.0
Hama	54.5	65.2	68.0	91.0
Homs	25.8	34.9	55.2	85.7
Raqqa	12.8	17.4	32.1	87.6
All	24.5	28.7	42.5	90.7

Source: Mueller et al. (PNAS 2021), Table 2.