

Satellite-Based Conflict Impact Assessment: Multi-Source Remote Sensing Analysis of the Khartoum Metropolitan Region

Integrating SAR, multispectral, night-time light, thermal, and geo-referenced conflict event data

Jiashu Hu^{1,2}

¹ School of Construction Management Technology, Purdue Polytechnic Institute, Purdue University, West Lafayette, IN, USA

² Law and Conflict (LAC) Research Lab, Department of Political Science, Purdue University, West Lafayette, IN, USA

hu964@purdue.edu · jiashuhu97@gmail.com · <https://hu-jiashu.github.io>

ORCID: 0009-0004-0434-5337

Preprint — August 2026.

Abstract. On 15 April 2023, war broke out in Khartoum between the Sudanese Armed Forces (SAF) and the Rapid Support Forces (RSF). Ground access closed. Satellites kept watching. I fuse nine free data sources — Sentinel-1 SAR, Sentinel-2 optical imagery, VIIRS night-time lights, MODIS vegetation indices, Landsat surface temperature, NASA FIRMS active fires, ACLED and UCDP conflict events, and OpenStreetMap buildings — to measure what the war did to a city of five million people. I process all satellite data in Google Earth Engine and all event data in Python. Four results stand out. First, night-time lights collapsed by 65–67% and stayed down. The trajectory has three phases: a one-month crash of 40%, a slow further decline into 2025, and a first year-on-year upturn in late 2025, after the SAF retook the city — still about 66% below pre-war levels. Second, vegetation barely moved. The area-mean NDVI change is only -0.009 , and the seasonal cycle runs through the war untouched; decline concentrates in 8.4% of pixels, mostly along the Nile farm belt. Lights and vegetation decoupled: the war switched off the city while the ecosystem kept running. Third, SAR backscatter dropped by more than 3 dB over 7.8% of the area, a pattern consistent with structural damage — but a distance-stratification test finds no significant concentration of backscatter loss near recorded battle sites ($p = 0.56$), so I keep the causal language careful. Fourth, two independent conflict-event datasets disagree on scale (5,496 vs. 245 events in the aligned window) yet agree on rhythm ($r = 0.60$). I report every method variant, including the null results, and I close with a practical guide that maps each observation target to the band and computation that serve it best. The pipeline is free, reproducible, and built to transfer to other conflicts.

Keywords: conflict remote sensing; Sudan; Khartoum; night-time lights; SAR change detection; NDVI; multi-source data fusion; Google Earth Engine; ACLED; UCDP; humanitarian monitoring

1 Introduction

1.1 The problem

The Sudan war began on 15 April 2023, when tension between the SAF and the RSF turned into open combat in the capital. By early 2026 the United Nations counted more than 25 million people in need of assistance and millions displaced. Khartoum — the political and economic center of the country — became the main battlefield. Journalists left. Aid agencies pulled back. Nobody could run a ground survey.

This report asks a simple question: how much can freely available satellite data tell us about what this war did to the city? I answer it with nine public data sources and one processing rule: everything must be

reproducible by anyone with an internet connection. I run all satellite processing in Google Earth Engine ([Gorelick et al., 2017](#)) and all conflict-event processing in Python. No commercial imagery. No field data. The point is to test how far the free stack goes, and to say clearly where it stops.

1.2 What others have done

Night-time lights are the classic remote proxy for human activity in war zones. [Agnew et al. \(2008\)](#) read the Baghdad “surge” through DMSP night-light signatures. The newer VIIRS sensor measures light far better than DMSP ([Elvidge et al., 2017](#); [Gibson et al., 2021](#); [Gibson and Boe-Gibson, 2021](#)). Researchers have used it to track COVID lockdowns ([Elvidge et al., 2020](#)), disaster blackouts and displacement ([Enenkel et al., 2020](#)), and refugee flows in Ukraine ([Huang et al., 2023a](#)). [Wang et al. \(2024\)](#) measured an 84% national-scale light loss in Ukraine over the first year of that war.

SAR sees through clouds and darkness, which makes it the workhorse for structural damage. [Aimaiti et al. \(2022\)](#) mapped damage in Kyiv with Sentinel-1 intensity change and matched 58% of UNOSAT reference labels. [Huang et al. \(2023b\)](#) built a dual-polarization coherence index for Mariupol and graded 79.2% of buildings as severely damaged or destroyed. Coherence methods also work for earthquakes ([Lu et al., 2018](#); [Sun et al., 2022](#); [Yang et al., 2024](#)), cyclones ([Malmgren-Hansen et al., 2020](#)), and post-war Mosul ([Darvishi Boloorani et al., 2021](#)); [Ge et al. \(2020\)](#) review the family. Machine learning scales these ideas up ([Gupta et al., 2019](#); [Mueller et al., 2021](#); [Bouchard et al., 2022](#); [Zhang et al., 2023](#)).

Fire detections add a third, physical channel. [Zubkova et al. \(2021\)](#) tied fire activity in Syria to the conflict there. [Tomchenko et al. \(2023\)](#) mapped war fires across Ukraine. Vegetation indices track the agricultural cost of war in Tigray ([Weldegebriel et al., 2024](#)) and Syria ([Papadimitriou et al., 2021](#)).

Two review-type papers frame the gap I aim at. [Sticher et al. \(2023\)](#) mapped the field and found it compartmentalized: most studies stay inside one sensor family. [Bennett et al. \(2022\)](#) argued for opening very-high-resolution archives to fix the access problem. I work the other lane: keep the data free and moderate-resolution, and fuse many sources at once so they can check each other.

1.3 Why Khartoum

Three reasons. First, the war in Khartoum is an intense, sustained urban conflict with impacts across every dimension a satellite can see — structure, vegetation, light, heat, fire. Second, the data situation is workable: continuous VIIRS, MODIS, and Landsat records, dense Sentinel-2 coverage, a rich ACLED event record, and a usable (if interrupted) Sentinel-1 archive. Third, the case is under-studied. Ukraine and Gaza have drawn most of the recent satellite work ([Aimaiti et al., 2022](#); [Poole et al., 2024](#); [Wang et al., 2024](#)). For Sudan I know of only a few satellite-based studies, on Darfur land cover ([Ahmed et al., 2024](#)) and on Khartoum flood risk ([Basheer and Elagib, 2024](#)). A comprehensive multi-sensor assessment of the capital did not exist. This report is my attempt to provide one.

1.4 What this report contributes

1. **A nine-source fusion of an under-studied war.** I integrate SAR, optical, thermal, night-time light, fire, event, and building data into one assessment of Khartoum, April 2023 to the end of 2025.
2. **A three-phase light trajectory, including the first recovery signal.** I document the collapse, the long decline, and a first year-on-year upturn in late 2025 that follows the SAF recapture of the city in March 2025.
3. **A decoupling result.** Lights collapsed; vegetation cycles did not. The weak NDVI response is itself a finding about what urban war destroys first.




Figure 1: Sentinel-2 true-colour composites of greater Khartoum.

(a) Pre-conflict median, January 2022 – March 2023.

(b) Post-conflict median, June 2023 – June 2024.

Contains modified Copernicus Sentinel data.

4. **Honest robustness.** I publish the tests that failed together with the ones that worked. A distance-stratification test does not link SAR backscatter loss to nearby battle sites. A fire–vegetation co-location test comes out non-significant. Table 3 lists every estimator variant side by side.
5. **A transferable pipeline and a closing usage guide.** Everything runs on free data in Google Earth Engine and Python, and Section 11 distills the whole exercise into a band-by-band observation guide.

2 Study area and data

The study area covers the Khartoum metropolitan region, 15.4° – 15.75° N and 32.3° – 32.8° E. The Blue Nile and the White Nile meet here and split the conurbation into three cities: Khartoum proper on the south bank, Omdurman in the west, and Bahri (Khartoum North) in the north. Before the war more than five million people lived in this box.

Before I measure anything, I want to show the place. I build two cloud-masked median mosaics from Sentinel-2 surface reflectance in Google Earth Engine, using bands B4/B3/B2 and a 2–98 percentile display stretch: one mosaic covers January 2022 to March 2023, before the war; the other covers June 2023 to June 2024, after its outbreak. Figure 1 shows the pair.

Look closely and the two panels differ: more bare soil in built-up districts, thinner green along the irrigated ribbon of the Nile. But these are impressions. The rest of this report measures them, one sensor at a time. Table 1 lists the nine data sources and their key numbers; each later section explains how I process its source.

Table 1: Data sources for the Khartoum conflict remote-sensing assessment. AOI: 15.4–15.75°N, 32.3–32.8°E.

Source	Product / role	Resol.	Temporal coverage	Key numbers
VIIRS DNB monthly (GEE, avg_rad)	Night-light time series + spatial change	500 m	Jan 2020 – Dec 2025 (72 months)	72 monthly AOI means
ACLED	Conflict events, AOI filter	points	Apr 2023 – Mar 2025 (24 months)	6,018 events; 8,845 fatalities
UCDP GED v25.1	Independent event source for cross-validation	points	Apr 2023 – Dec 2024 (GED v25.1 ends 2024-12)	245 events in window; historical pre-2023 records excluded
Sentinel-1 GRD (VV, IW, descending)	SAR backscatter change 2021→2025	30 m (native 10 m)	2021 baseline; Apr–Dec 2025; no AOI coverage 2023–2024	56 scenes (2021); 21 scenes (2025)
Sentinel-2 SR — change composites	Δ NDVI, pre/post medians	10 m	Pre: Jan 2022 – Mar 2023; Post: Jun 2023 – Jun 2024	21,694,869 valid pixels; QA60 mask, cloud < 20%
Sentinel-2 SR — monthly series	Monthly NDVI + scene counts	10 m	Jan 2022 – Dec 2025 (48 months)	min $n = 2$ (Aug 2022); median 11
Landsat 8/9 C2 L2	Land surface temperature (ST_B10)	100 m	Monthly 2022–2025	$K = DN \times 0.00341802 + 149.0$
MODIS MOD13A2 v6.1	Coarse NDVI consistency check	1 km	Monthly 2020–2025	72 monthly records
NASA FIRMS (VIIRS S-NPP)	Active-fire detections	375 m	Apr 2023 – Dec 2025	854 detections; 749 high/nominal
OpenStreetMap (Overpass)	Building centroids, exposure context	vector	static snapshot	~800,000 centroids

3 Structural change from SAR

Buildings act as corner reflectors for C-band radar. Wall meets ground, the radar pulse bounces back, and the pixel is bright. Knock the wall down and the return weakens. This physics has powered damage mapping in Kyiv (Aimaiti et al., 2022), Mariupol (Huang et al., 2023b), Mosul (Darvishi Boloorani et al., 2021), and after earthquakes (Lu et al., 2018; Sun et al., 2022; Yang et al., 2024); Ge et al. (2020) review the approach. I apply the simplest version of it: intensity change between two composites.

First, the data problem. Sentinel-1B suffered a power failure on 23 December 2021 and never sent radar data again; ESA declared the mission over on 3 August 2022 (European Space Agency, 2022). That loss, plus a gap in Sentinel-1A’s acquisition plan over Sudan, left my study area with zero Sentinel-1 coverage from January 2023 to March 2025. Coverage returned with Sentinel-1C, launched on 5 December 2024 (European Space Agency, 2024), whose data opened to the public on 26 March 2025 (Copernicus Programme, 2025). So I compare the last good year before the war with the first good stretch after it: a 2021 annual mean of 56 descending VV scenes against an April–December 2025 mean of 21 scenes, both built in Google Earth Engine at 30 m, with values in decibels.

Figure 2 maps the difference. One display convention, used here and in every change map below: the colour scale is symmetric about zero, clipped at the 98th percentile of the absolute change, and the raster sits half-transparent over satellite imagery so the ground context stays visible.

The numbers: 7.8% of the valid area lost more than 3 dB of backscatter. The mean change is -0.26 dB and the median -0.13 dB, so the loss concentrates in a long tail while most pixels barely moved. Red




Figure 2: Sentinel-1 VV backscatter change, 2025 mean minus 2021 baseline, 30 m. Red marks backscatter loss. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics). Contains modified Copernicus Sentinel data.

patches cluster in central Khartoum and along the Omdurman–Bahri corridor, and they also cover large peri-urban blocks in the south and east that look agricultural. Blue patches mark backscatter gains: new construction, debris piles that reflect, or vegetation on abandoned lots.

I write “consistent with structural damage” throughout, and I avoid the bare word “damage.” Here is why. The two composites sit four years apart, and two of those years I cannot see at all. Everything that happened in between — fighting, demolition, land-use change, soil-moisture shifts, farm cycles — lands in one difference image. Section 7 tests whether the losses concentrate near recorded fighting. They do not, and that result sets the limit on what this map can claim.

Context helps to read the change map. Figure 3 shows the pre-war radar baseline in greyscale (2–98 percentile stretch, values in dB) with the full conflict-event record on top of it.

The overlay makes one thing plain: the fighting happened where the city is brightest — the densest built fabric on both banks. That is where corner reflectors stand, so damage detection has the most signal to lose there. It is also where population, commerce, and infrastructure concentrate, which matters for every section that follows.




Figure 3: 2021 Sentinel-1 VV baseline (greyscale, dB) with the 6,018 ACLED events of April 2023 – March 2025, coloured by type.
Data: Copernicus Sentinel-1 via Google Earth Engine; ACLED (acleddata.com).

4 Vegetation and fire

War hits farmland through people. Farmers flee, irrigation pumps stop, supply chains break, and fields burn (Ahmed et al., 2024; Papadimitriou et al., 2021; Weldegebriel et al., 2024). Khartoum’s food belt hugs the Nile and depends on continuous human management. So I expect vegetation loss — if it happens — to trace the river.

I compute NDVI from Sentinel-2 surface reflectance in Google Earth Engine, mask clouds with the QA60 band, keep scenes under 20% cloud, and build two median composites: pre-war (January 2022 – March 2023) and post-onset (June 2023 – June 2024). I subtract them, pixel by pixel, and overlay the 749 high- and nominal-confidence FIRMS fire detections. Figure 4 shows the result.

Two facts sit side by side in this map, and both matter. Fact one: 8.4% of the area shows real decline ($\Delta\text{NDVI} < -0.05$), and the red concentrates along the Nile corridor and in peri-urban farm blocks — the pattern a management collapse would produce. Fact two: the area-wide mean change is only -0.009 , and most of the study box (86.4%) stays stable. The war left the landscape as a whole largely unchanged and hit a limited set of places hard. Sections 5 and 8 return to why this weakness is itself informative.

To sort the affected places by depth, I cut the same ΔNDVI surface into five classes, with thresholds at ± 0.05 and ± 0.15 . Figure 5 maps them.

Severe loss covers 2.7% of the area; moderate loss another 5.7%. At the same time, 5.1% of the area gained greenness, mostly on riverbanks and floodplains. Rivers flood, plants respond, and none of that needs a war. I keep this warning in view for the whole report: in a strongly seasonal, river-fed landscape, a two-window difference mixes conflict effects with hydrology. Section 5.1 handles the same problem quantitatively for temperature.




Figure 4: Δ NDVI (post-conflict median minus pre-conflict median) with FIRMS fire detections. Red marks vegetation decline; orange triangles are fires, April 2023 – December 2025 ($n = 749$).
Data: Copernicus Sentinel-2; NASA FIRMS. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).

4.1 Active fires

Fires give me an independent physical trace of destruction. I query the NASA FIRMS API for VIIRS S-NPP detections inside the study box from April 2023 to December 2025. It returns 854 detections; 749 survive the confidence filter (nominal or high). Figure 6 maps them, coloured by acquisition date.

The dark-purple early-war detections sit in the contested urban cores — the same districts the event record marks. Later detections spread outward. Syria and Ukraine show the same conflict–fire coupling (Zubkova et al., 2021; Tomchenko et al., 2023). Do the fires drive the vegetation loss here? I test this directly. Mean Δ NDVI within 375 m of a fire is -0.0065 ; outside, -0.0091 — near-fire pixels are on average slightly less brown. Severe-decline pixels, though, are a bit more common near fires (9.6% vs. 8.3%). The two summary metrics point in opposite directions, and a permutation test on the difference gives $p = 0.10$. So I describe fires and vegetation loss as co-located in the same districts, and I stop short of calling them correlated. Most of the mapped NDVI loss likely comes from failed irrigation and abandonment; burning adds local damage on top.




Figure 5: ΔNDVI severity classes. Class shares of valid area: severe loss 2.7%, moderate loss 5.7%, stable 86.4%, moderate gain 4.0%, strong gain 1.1%. Data: Copernicus Sentinel-2. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).




Figure 6: FIRMS active-fire detections (nominal/high confidence, $n = 749$) coloured by date, April 2023 – December 2025.
Data: NASA FIRMS (VIIRS S-NPP). Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).




Figure 7: Night-time lights and vegetation over the study area, January 2020 – December 2025.

(top) Monthly VIIRS radiance; dashed lines mark the conflict onset and the two window means; shading marks Phases I–III; the circle marks the July 2022 artifact; the arrow marks the SAF recapture of Khartoum.

(bottom) Monthly MODIS NDVI on the same time axis.

Data: NASA VIIRS DNB; NASA MODIS.

5 Night-time lights and surface temperature

Night-time light is the bluntest, most human of all satellite signals: electricity, commerce, occupancy, street lamps. That is why it works as an economic and displacement proxy (Agnew et al., 2008; Gibson et al., 2021; Gibson and Boe-Gibson, 2021; Elvidge et al., 2020; Enenkel et al., 2020), and why it tracked the Ukraine war so closely (Huang et al., 2023a). I use the NOAA VIIRS DNB monthly composites (stray-light corrected, avg_rad band), 72 months from January 2020 to December 2025, averaged over the study box.

Definitions first, because they move the answer. Pre-war window: January 2020 – March 2023 ($n = 39$ months). Post-onset window: May 2023 – December 2025 ($n = 32$). April 2023 is the transition month — the war started mid-month — so I exclude it from both windows. One pre-war month, July 2022, reads $0.09 \text{ nW cm}^{-2} \text{ sr}^{-1}$; that is a single-month sensor or coverage artifact, unrelated to the conflict, and the figure marks it with a circle.

The key test of this whole report is a single plot. If the lights fall and the plants do not, the fall is human. So I put VIIRS radiance and MODIS NDVI on the same time axis: radiance on top with the two window means drawn as horizontal lines and the three phases shaded, NDVI below. Figure 7 shows both.

Here is what the top panel says, in numbers. The window means give a decline of -66.7% (9.21 to $3.06 \text{ nW cm}^{-2} \text{ sr}^{-1}$). Matching calendar months, to control for seasonality, gives -64.8% . I quote the interval: -65% to -67% . The pre-war trend is flat to slightly rising ($+0.050 \text{ nW per month}$, $p = 0.17$), so the collapse did not start before the war.

The shape matters as much as the size. I read three phases. Phase I is the crash: April to May 2023, -40.2% in one month. Phase II is the long slide, June 2023 to September 2025: 2024 averaged 3.24 nW , below the 4.22 nW of the first post-onset months of 2023, and the series bottomed at 1.01 nW in June 2025. The city kept getting darker for two years after the shock. Phase III covers October to December 2025 and is new in this version of the assessment: each of those months beat the same month of 2024, by $+45\%$, $+19\%$, and $+32\%$. The SAF retook Khartoum in March 2025; the first season-consistent light recovery followed within the year. I read this carefully. December 2025 still sits 66% below December 2022. But the sign flipped, and that gives anyone monitoring recovery a live indicator to watch.

Now the bottom panel. MODIS NDVI cycles through the whole war as if nothing happened — the 2024 green-up is even the strongest in the record, on the back of heavy rains. Lights and vegetation decoupled. The war switched off the human system and left the river-fed vegetation system running. Two practical lessons follow. For analysts: NDVI alone would have graded this catastrophe as “mostly stable,” so vegetation indices under-report urban war. For this report: the weak ΔNDVI of Section 4 is the ecological half of a real asymmetry.

Where did the light go missing? I difference the mean radiance of 2024 against 2022, using the standard display rule, and overlay the event record. Figure 8 shows the map.

The darkness is almost total in coverage: 94.7% of pixels lost light, and the mean change is $-5.74 \text{ nW cm}^{-2} \text{ sr}^{-1}$. The loss runs deepest exactly where Figure 3 put the fighting: central Khartoum and the Omdurman–Bahri corridor. The event dots and the darkest reds occupy the same districts. Two fully independent data systems agree at the district scale here, and that agreement anchors the synthesis in Section 8.

5.1 Surface temperature

An abandoned city should also run hotter at the surface: less canopy, less watering, dark empty streets, rubble. Landsat measures this directly. I build a monthly land-surface-temperature series from Landsat 8/9 Collection 2 Level 2 (band ST_B10 , converted to $^{\circ}\text{C}$), 2022–2025, and I also difference the two composite windows on the map (pre: January 2022 – March 2023; post: October 2023 – June 2024). Figure 9 shows the spatial pattern.

The map runs warm nearly everywhere. I do not trust its magnitude, and I want to show why. The two windows sample different seasons, and Khartoum swings by more than 15°C across the year. A naive window difference bakes that swing into the answer. So I match calendar months instead: compare each January with Januaries, each October with Octobers. Four rainy-season months carry cloud-contaminated retrievals (monthly means below 25°C , physically implausible here): July 2022, August 2023, July 2024, and September 2025. I flag them with circles and drop them from the monthly averaging; July then has no pre-war bar in the lower panel, because its only pre-war observation is one of the flagged months. Figure 10 shows the series and the matched comparison.

The honest number is $+2.5$ to $+3.0^{\circ}\text{C}$ of seasonally matched surface warming: $+2.47^{\circ}\text{C}$ when I match within the two composite windows, $+2.98^{\circ}\text{C}$ when I match across the full record. The raw window difference of $+6.26^{\circ}\text{C}$ is roughly half seasonal artifact, and I retire it to a footnote. A harmonic deseasonalization agrees once I drop the four flagged months ($+2.73^{\circ}\text{C}$; $+3.90$ without the quality filter). Two to three degrees of surface warming across a metropolitan area is still a large environmental signal, and its direction fits the abandonment story from the lights. The method lesson generalizes: in this climate, every two-window comparison in this report — SAR, NDVI, LST — carries seasonal risk, and only the




Figure 8: VIIRS night-light change, mean 2024 minus mean 2022, with ACLED events as black dots. Red marks dimming.
 Data: NOAA VIIRS DNB; ACLED. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).

ones I can seasonally match deserve tight error language.




Figure 9: Landsat land-surface-temperature change, post-window minus pre-window ($^{\circ}C$). Red marks warming; the magnitude is inflated by seasonal mismatch between the windows (see text).
Data: USGS Landsat Collection 2. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).




Figure 10: Landsat land surface temperature, January 2022 – December 2025.
 (top) Monthly means; circles mark the four flagged rainy-season months.
 (bottom) Pre- vs. post-conflict means by calendar month, flagged months excluded.
 Data: USGS Landsat Collection 2.




Figure 11: All 6,018 ACLED events, April 2023 – March 2025, coloured by event type. Points jittered up to 0.65 km for display.
Data: ACLED (acleddata.com). Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).

6 The conflict-event record

Satellites see effects. Event data name causes. I download the ACLED archive for Sudan, clip it to the study box, and keep April 2023 through March 2025: 24 months, 6,018 events, 8,845 recorded fatalities. Battles dominate (2,551 events; 42%), then explosions and remote violence (2,127; 35%), violence against civilians (661; 11%), and strategic developments (662; 11%); protests and riots almost vanish under wartime conditions. Splitting by city: Khartoum proper logs the most events (2,466), and Omdurman logs the most deaths (2,940 from 1,514 events) — almost two fatalities per recorded event, against roughly one in Khartoum. Figure 11 maps the full record by type. One display note: ACLED snaps events to shared reference coordinates, so for the map I jitter each point by up to 0.65 km (fixed seed); without the jitter, six thousand events would collapse onto about 130 visible dots.

One property of this dataset shapes everything I do with it: ACLED geocodes these 6,018 events to 127 unique coordinate pairs (126 place names). Six thousand events live at about one hundred and thirty points. District-level statements survive that; building-level statements do not. I lean on this fact when I interpret the SAR test in Section 7, and I return to it in Section 8.

To see concentration instead of points, I estimate a kernel density surface over all events in the




Figure 12: Kernel density of the 6,018 ACLED events (events km^{-2}), peaking in central Khartoum.

Data: ACLED. Basemap: © OpenStreetMap contributors © CARTO.

24-month window (Gaussian kernel; the effective smoothing radius is about 130 m), scaled to absolute density in events per square kilometre. Figure 12 shows it.

The peak sits in central Khartoum at roughly 3,400 events per square kilometre over the two years. That number only makes sense because the events snap to shared coordinates. I read it as “this district absorbed sustained combat,” and I do not read it as a literal per-block count.

Time tells the other half of the story. Figure 13 plots monthly events and fatalities on one panel and monthly fire detections below, so the tempo of combat and the tempo of burning can be compared directly.

Across the 24 months of event coverage I see no sustained de-escalation: event counts stay high, and fatality spikes track the major offensives. Fire detections rise and fall in rough sympathy with combat intensity, then continue after the event window closes in March 2025. The persistence matters. It explains why the lights of Figure 7 kept falling through 2024: the war never paused long enough for the city to restart.




Figure 13: Monthly series, April 2023 onward.
 (top) ACLED events (bars) and fatalities (line), to March 2025.
 (bottom) FIRMS fire detections, to December 2025.
 Data: ACLED; NASA FIRMS.

7 Cross-validation and robustness

Every claim so far rests on one dataset per domain. This section stress-tests the two riskiest links: the event record itself, and the causal reading of the SAR map.

7.1 ACLED against UCDP

ACLED is one way to count a war. UCDP GED ([Sundberg and Melander, 2013](#)) codes organized armed violence under stricter rules than ACLED ([Raleigh et al., 2010](#)). Comparing them tells me how much my picture depends on coding choices. The comparison needs care on two axes. The datasets differ in inclusion criteria — that is well known. They also differ in coverage: my UCDP extract contains historical Khartoum events back to 1992, which I exclude, and GED v25.1 ends in December 2024 while my ACLED window runs to March 2025. So I align the window to April 2023 – December 2024 (21 months) before computing anything. Figure 14 shows both monthly series, with the UCDP axis drawn at one twelfth of the ACLED axis so both bar sets stay readable.

In the aligned window, ACLED records 5,496 events and UCDP records 245 — a 22.4-fold gap. Fatalities gap too: 7,960 against 1,862 (best estimates). Yet the monthly rhythms match: $r = 0.60$ and $\rho = 0.65$ for events, $r = 0.61$ for fatalities. My reading: the two systems watch the same war through very different apertures. The levels reflect counting rules; the timing reflects the war. A study that quotes a single event count without naming its dataset and window is quoting a coding choice. I therefore anchor




Figure 14: ACLED vs. UCDP GED monthly events, aligned window April 2023 – December 2024 ($n = 21$). The right (UCDP) axis spans 1/12 of the left (ACLED) axis. Data: ACLED; UCDP GED.

my temporal claims — no de-escalation, fatality spikes with offensives — on the shared rhythm, which both datasets support.

7.2 Does backscatter loss follow the fighting? A null result

Section 3 promised a test, and here it is. If war damage drives the SAR change map, pixels near recorded battles and explosions should lose more backscatter than pixels far away. I compute, for every valid pixel, the distance to the nearest ACLED battle or explosion site (in UTM 36N, with a k-d tree), cut the area into three rings — under 250 m, 250 m to 1 km, over 1 km — and compare the change statistics across rings. Then I ask how surprising the observed near-versus-far difference would be if the 4,678 events had landed at random anywhere in the study box: I relocate them uniformly 999 times and recompute the statistic each time, on an every-second-pixel quarter sample of the 2,412,942 valid pixels. Figure 15 shows the outcome.

The test fails, and I report it as it fell. Mean backscatter loss near events (-0.24 dB within 250 m) is statistically indistinguishable from loss far from them (-0.27 dB beyond 1 km); the observed near-minus-far difference is $+0.03$ dB (full-resolution statistic 0.0286 dB), and the permutation p is 0.556 . Worse for the damage story, the severe-loss tail ($\Delta VV < -3$ dB) is most common far from events (8.2% vs. 5.1%), out in the peri-urban belts where Figure 2 showed field-shaped red blocks. The medians do fall off with distance (-0.18 , -0.14 , -0.13 dB), a faint trace of localized urban damage, and it does not survive the test.

Three caveats keep me from over-reading the null itself. The 4,678 events collapse onto 117 coordinates, so “distance to event” is coarse. The null model is complete spatial randomness, a low bar. And the two composites span four years, so agricultural and soil-moisture change had two full wartime seasons to pile into the far ring. Still, the conclusion I can defend is narrow and useful: **at 30 m, over a four-year gap, I cannot attribute pixel-scale backscatter loss to nearby fighting.** District-scale convergence (Section 8) survives; pixel-scale attribution does not. This is why Section 3 says “consistent with,” and




Figure 15: ΔVV by distance to the nearest ACLED battle/explosion site. Bars: stratum means; diamonds: medians; grey labels: share of pixels below -3 dB.
Data: Copernicus Sentinel-1; ACLED.

why Section 10 points the next step at coherence methods — intensity differencing has reached its limit here.

7.3 Where the lethal events sit

Finally, exposure. I pull roughly 800,000 building centroids from OpenStreetMap and overlay the 510 ACLED events with five or more recorded deaths, drawing each event as a circle whose area scales with the fatality count. Figure 16 shows the result.

The high-fatality events sit inside the densest building fabric on both banks, among homes, markets, and institutions. The single worst recorded event killed 312 people. The exposure logic of this map needs no model: in this war, the places people die are the places people live.




Figure 16: ~800,000 OSM building centroids (pale) with the 510 ACLED events of ≥ 5 fatalities, circles scaled by fatality count. Building data © OpenStreetMap contributors (ODbL). Data: ACLED. Basemap: Esri World Imagery (Esri, Maxar, Earthstar Geographics).

8 Synthesis


The preceding five sections each measured one slice of the war. This section puts the slices together and asks three questions. Where do the layers agree? Where do they disagree, and what does each disagreement teach? And what can this stack of data legitimately claim, given its resolution and its geocoding? To start, Figure 17 assembles the four change surfaces — radar, vegetation, light, and event density — at the same extent, each drawn with the display rules of its parent figure.

8.1 Where the layers agree

Walk the four panels with one question: which districts light up in every layer? The answer is consistent. In panel (a), backscatter loss marks central Khartoum and the Omdurman–Bahri corridor. In panel (b), vegetation decline hugs the Nile through the same corridor. In panel (c), the dimming covers 94.7% of pixels and runs deepest over the same urban core. In panel (d), event density peaks in central Khartoum at about 3,400 events per square kilometre. Physical sensors and human event coding are fully independent measurement systems, and they point at the same districts. That convergence is the strongest joint claim this report makes, and it holds at the district scale — roughly the kilometre scale and above.

Table 2 condenses the findings by domain.

Read down the value column of Table 2 and a scale pattern emerges. The human-layer signals are



Basemaps: Esri World Imagery (Esri, Maxar, Earthstar Geographics); © OpenStreetMap contributors © CARTO

Figure 17: Composite overview; all panels share one extent.

(a) SAR ΔVV , 2025 minus 2021; red marks backscatter loss.

(b) $\Delta NDVI$ with FIRMS fires; red marks vegetation decline.

(c) Night-light change, 2024 minus 2022; red marks dimming.

(d) ACLED event kernel density (events km^{-2}).

Basemaps: Esri World Imagery (a–c); © OpenStreetMap contributors © CARTO (d).

Contains modified Copernicus Sentinel data. Data: Copernicus Sentinel-1/-2, NOAA VIIRS, NASA FIRMS, ACLED.

enormous: two thirds of the light gone, six thousand events, nine thousand deaths. The ecosystem-layer signals are small: an NDVI mean shift of -0.009 , a seasonal cycle that never broke. The table is the decoupling result written as numbers, and it is the single most transferable observation in this report: in an urban war, the human layer moves first and moves hardest.

8.2 Where the layers disagree, and what each disagreement taught

The convergence above is real. The disagreements taught me more. Four of them, each with the lesson I now build into the method.

Disagreement 1: the SAR severe-loss tail sits away from the fighting. Pixels that lost more than 3

Table 2: Key findings by analysis domain (verified numbers).

Domain	Headline finding	Value(s)
SAR backscatter	Cumulative 2021→2025 change: small mean decline with a localized severe-loss tail; no significant near-event concentration (permutation $p = 0.556$; Fig. 15)	7.8% of pixels < -3 dB; mean -0.26 dB
NDVI	Localized decline against an intact AOI-wide seasonal cycle	8.4% of pixels < -0.05 ; AOI mean -0.009
Night-time lights	Three-phase collapse; first year-on-year upturn Q4 2025; Dec 2025 still $\approx 66\%$ below Dec 2022	-65% to -67%
Land surface temperature	Post-onset warming after seasonal matching	$+2.5$ to $+3.0^\circ\text{C}$
Conflict events	ACLED, Khartoum AOI, Apr 2023 – Mar 2025 (24 months); no de-escalation	6,018 events; 8,845 fatalities
Fires	FIRMS detections, Apr 2023 – Dec 2025	854 total; 749 high/nominal
Cross-validation	ACLED vs UCDP monthly series, aligned window ($n = 21$)	$r = 0.60$; $\rho = 0.65$

dB are nearly twice as common beyond 1 km of any battle site (8.2%) as within 250 m of one (5.1%), and the field-shaped geometry of the far-field patches gives the reason away: those are croplands cycling and soils drying across a four-year window. What I learned: a multi-year intensity difference measures total land-surface change. War contributes to it; agriculture can dominate it. An intensity-change map should never be read as a damage map until it has passed a mask and a test.

Disagreement 2: lights record a catastrophe, vegetation records two ordinary years. The same city, the same months, two opposite stories — and both are true. The resolution is that the proxies watch different systems: DNB watches people, NDVI watches water and plants. What I learned: indicator choice decides the story. An assessment must state which layer it leads with, and a monitoring system for urban conflict must lead with the human-layer sensor.

Disagreement 3: the two fire–vegetation metrics point in opposite directions. Near fires, mean ΔNDVI is slightly milder, yet severe-decline pixels are slightly more frequent. Both can hold at once because fires cluster in built-up and fringe areas where baseline NDVI is already low, so the mean has little room to fall, while the pixels that do burn fall hard. What I learned: co-location without correlation is a real, reportable state. Forcing a one-sentence verdict — “fires drove the vegetation loss” — would be wrong in both directions.

Disagreement 4: ACLED and UCDP disagree 22-fold on scale and agree on rhythm. The levels differ because the counting rules differ; the correlation ($r = 0.60$) survives because the war’s tempo is real. What I learned, as a rule: report event counts only with the dataset name and window attached, and anchor temporal claims on the rhythm both sources share.

Taken together, the four disagreements did more for the method than the agreements. Each one, examined instead of smoothed over, turned into a usable rule.

8.3 What the geocoding precision permits

ACLED builds its coordinates from reported place names: an event is snapped to the reference point of the most specific named location. In my extract, 6,018 events resolve to 127 unique coordinate pairs, and 654 events carry an explicitly coarse precision flag. UCDP GED encodes the same reality with its own precision codes. A “point” in these datasets is a label for an area.

That single fact sets the resolution floor for every spatial claim in this report. At or above the district scale — roughly a kilometre — the event record supports strong statements: which districts absorbed

combat, how the corridor structure ran, where the lethal events clustered, and the cross-layer convergence of Figure 17. Below that scale the record supports little. The KDE peak of 3,376 events per square kilometre is meaningful as a ranking of districts and meaningless as a literal count on any block. The <250 m stratum of the SAR test is nominal: many events assigned to a reference point actually happened some unknown distance away. Building-level attribution from these data was never possible, and the null of Section 7 partly reflects that.

The rule I draw: match the spatial claim to the coarsest georeferencing in the analysis chain. My chain bottoms out at the event geocoding, so my claims stop at the district. With building-level damage labels — UNOSAT-style annotation, or coherence-based grading — the same satellite stack could support much finer claims. That upgrade path is Section 10's first item.

8.4 What the composited spectra delivered, and what stayed out of reach

Each spectral domain earned its keep differently, and each hit a wall. The optical composites (Sentinel-2, red and near-infrared) delivered the severity map, the class shares, and the corridor pattern at 10 m. They could not deliver building damage: a nadir view at 10 m sees roofs, an intact roof can hide a destroyed interior, and a median composite deliberately suppresses anything transient. The thermal band (Landsat ST_B10) delivered a defensible city-mean warming of +2.5 to +3.0°C once seasonally matched; it could not attribute warming pixel by pixel, because surface mixing at 100 m and cloud gaps blur the causes. The DNB delivered the headline of the whole report — the collapse, the phases, the recovery sign — at 500 m and monthly cadence; it cannot separate depopulation from grid failure from blackout policy, and single months can glitch (July 2022). SAR intensity delivered a cumulative-change screening layer and an honest null; it could not grade or attribute damage. The fire channel delivered timing and location of burning; at one overpass a day it misses small, short fires.

One sentence covers the pattern: the composites answer *where* and *how much*, at district scale; they do not answer *which building* or *why*. That boundary is a property of free moderate-resolution data. Section 11 turns it into a usage guide.

8.5 The linkage ledger

The validation work reduces to a short ledger of tested links. Two links are quantified and significant: the ACLED–UCDP monthly correlation (events $r = 0.603$, $p = 0.004$; $\rho = 0.653$, $p = 0.001$) and its fatality counterpart ($r = \rho = 0.605$). Two links were tested and came out null: SAR backscatter loss against event proximity (+0.03 dB, $p = 0.556$) and fire proximity against Δ NDVI ($p = 0.103$). Three links are visual and district-level, and I present them as such without a formal test: night-light loss against event density (Figure 8), early fires against early event clusters (Figure 6), and monthly fires against monthly combat tempo (Figure 13).

This ledger is the empirical spine of the report. Claims in Section 9 lean only on the first group where they need statistics, on the third group where district-level language suffices, and on nothing where the tests failed. Table 3 completes the picture: for every headline number, it lists all the estimator variants I computed, so a reader can see how much each answer moves when the method moves.

Table 3 rewards a slow read. The LST rows show the largest method sensitivity in the report: the estimate moves from +6.26 to +2.47°C depending on seasonal handling — a swing of nearly four degrees, larger than many published warming effects. The NTL rows show the opposite: window means and month matching land within two percentage points of each other, which is why I trust the light headline more than any other number here. The two null rows are printed with the same prominence as the significant ones. When a reader asks which variant to trust, my answer is mechanical: prefer the seasonally matched estimator wherever an annual cycle exists, quote intervals when two defensible

estimators disagree, and keep the nulls visible.

Table 3: Method variants and robustness checks. Every estimator variant computed in this revision, side by side. All values come from the frozen statistics files.

Metric	Variant	Value	Note
NTL % change	Window means	−66.7%	Pre Jan 2020 – Mar 2023 (9.21 nW) vs post May 2023 – Dec 2025 (3.06 nW); Apr 2023 excluded
NTL % change	Same-month matched	−64.8%	Mean over 12 calendar months; controls seasonality
NTL % change	Headline interval	−65% to −67%	Interval quoted in text
NTL pre-war trend	OLS slope, pre-window	+0.050 nW/month	$p = 0.17$, not significant; rules out pre-existing decline
LST change	Raw window difference	+6.26 °C	Seasonally confounded; footnote-only value
LST change	Window same-month matched	+2.47 °C	Same windows, per-calendar-month averaged
LST change	Full-record same-month matched	+2.98 °C	All pre vs all post months; Apr 2023 excluded
LST change	Harmonic deseasonalized	+3.90 °C	Annual harmonic fit on pre-war months; all data
LST change	Harmonic deseasonalized, QC	+2.73 °C	Excludes 4 cloud-contaminated months (Jul 2022, Aug 2023, Jul 2024, Sep 2025)
SAR ΔVV by distance	<250 m	mean −0.24; median −0.18 dB; 5.1% < −3 dB	$n = 26,567$ pixels
SAR ΔVV by distance	250 m – 1 km	mean −0.16; median −0.14 dB; 5.2% < −3 dB	$n = 302,780$ pixels
SAR ΔVV by distance	>1 km	mean −0.27; median −0.13 dB; 8.2% < −3 dB	$n = 2,083,595$ pixels
SAR near–far difference	Observed statistic	+0.0286 dB (full resolution)	Subsample value +0.0228 dB used in permutation loop
SAR near–far difference	Permutation test	$p = 0.556$	999 permutations; null result
ACLED–UCDP (events)	Pearson	$r = 0.603$ ($p = 0.004$)	Aligned window Apr 2023 – Dec 2024; $n = 21$
ACLED–UCDP (events)	Spearman	$\rho = 0.653$ ($p = 0.001$)	Same window
ACLED–UCDP (fatalities)	Pearson / Spearman	$r = \rho = 0.605$	Same window
ACLED KDE	Peak density	3,376.5 events km ^{−2}	At 32.5535°E, 15.5695°N; replaces the v3 per-deg ² figure
ACLED KDE	Effective kernel σ	≈ 0.13 km	Covariance factor 0.02 is unitless
ACLED KDE	Integral check	99.997% of N	KDE mass over AOI vs $N = 6,018$
FIRMS × NDVI	Mean ΔNDVI inside vs outside 375 m	−0.0065 vs −0.0091	Near-fire mean slightly less negative
FIRMS × NDVI	Share ΔNDVI < −0.05	9.61% vs 8.34%	Severe-decline pixels more frequent near fires
FIRMS × NDVI	Permutation test	$p = 0.103$	Not significant; wording is “co-located”

9 Discussion

What the free stack delivered. With zero ground access, public data measured: where the war concentrated (four independent layers agree at district scale), how deep the economic shutdown ran (-65 to -67% of light for two and a half years), when recovery started (Q4 2025, after the March 2025 change of control), which farm belts lost greenness (8.4% of area, along the Nile), how much hotter the abandoned surface runs ($+2.5$ to $+3.0^{\circ}\text{C}$), and where the lethal events clustered (inside the dense building fabric). [Sticher et al. \(2023\)](#) argue that moderate-resolution, freely available imagery is where conflict monitoring scales; this case supports them.

What it refused to deliver. Building-level damage attribution. The distance test said no, and the reasons are structural: a four-year SAR gap, 30 m pixels, and an event record that snaps to 127 coordinates. I count this among the most useful negative results in the report, because it marks the exact line where free intensity-differencing stops and coherence methods must take over.

9.1 Analysis and prediction with zero field access

Suppose the analyst can never visit, interview, or verify on the ground — the situation this report was actually written in. Here is what the stack supports, dimension by dimension, with the cadence and the confidence each channel earned in this study.

Economic activity and electrification. The DNB series is the workhorse: monthly cadence, about one month of latency, robust at the city scale. It supports level estimates (the -65 to -67% interval), phase structure, and a live recovery dial. It cannot split depopulation from grid failure from blackout policy; that separation needs auxiliary reporting.

Displacement. Direction is inferable: light loss of this depth and duration, spatially co-located with combat, indicates mass departure, and the literature links the two ([Enenkel et al., 2020](#); [Huang et al., 2023a](#)). Magnitudes need demographic models and are outside what I claim here.

Farmland and food security. The Sentinel-2 monthly series and the severity map watch the cultivated belt at 10 m and monthly cadence. They support targeting — which irrigation blocks fell out of management — and seasonal early warning. They under-report urban war by construction, so they must never be the lead indicator for a city.

Destruction events in near-real time. FIRMS delivers detections within hours. In this study fires tracked combat tempo and marked the early-war urban cores. Sparse but datable: useful as an alerting channel, weak as a magnitude measure.

Combat tempo and geography. ACLED updates weekly and UCDP annually; both are themselves remote products, built from reporting. Together they give tempo, type structure, and district geography — with the level caveat of Section 7.

Surface habitability. The matched LST trend adds a slow environmental dimension: an abandoned city measurably heats up. Quarterly cadence is realistic given cloud gaps.

Structural change. SAR provides a screening layer where acquisitions exist, and it ranks where scarce survey capacity should go once access returns. Every red patch stays a candidate until someone checks it.

On prediction, I stay honest: this stack nowcasts; genuine forecasting needs models trained across many conflicts. But three forward-looking uses stand on the results. First, recovery tracking: Phase III is a testable expectation — if control of the city holds, year-on-year light gains should persist through 2026, and the monthly series will say so within weeks of any reversal. Second, escalation warning: in this record, event tempo and fire activity moved with the light losses, so a joint uptick in events and fires is a days-to-weeks leading signal worth watching in any city. Third, survey targeting: the SAR and NDVI screening layers convert satellite output into a ranked field-visit list the moment access opens. That is

the practical meaning of “analysis without field access”: a monthly, district-level dashboard of economy, agriculture, fire, heat, and combat, maintainable for any city on Earth from a desk. This report is one instance of that dashboard, run for thirty-three months.

9.2 How Khartoum compares

Ukraine lost about half its national light within a week of the invasion (Huang et al., 2023a) and 84% over the first year by one national-scale metric (Wang et al., 2024); Mariupol lost most of its buildings (Huang et al., 2023b). The metrics differ across studies, so I compare the shapes of the curves and leave the decimals alone. Khartoum’s profile is distinct: a slower one-month crash (−40%), then two further years of decline — 2024 darker than 2023, mid-2025 darkest of all — and only then an upturn. A slide that long, with no plateau, is unusual in the literature I know, and it matches the event record’s persistence. Sudan’s capital sustained one of the deepest, longest urban blackouts in the recent conflict record, with a fraction of the research attention Ukraine or Gaza received (Ahmed et al., 2024; Basheer and Elagib, 2024; Poole et al., 2024).

9.3 Method points for other studies

Three habits from this revision travel well. First, align windows before comparing event datasets; my UCDP extract carried 114 historical events and a different end date, and without alignment the “22-fold gap” would have been a wrong number with a right-sounding story. Second, seasonally match every before/after comparison in a monsoon climate; matching cut my LST estimate roughly in half. Third, publish the estimator table. Table 3 costs one page and converts every headline number from an assertion into a choice the reader can audit.

9.4 Implications beyond remote sensing

Numbers like these travel into policy. A sanctions committee, a reconstruction lender, or a returns-planning agency needs a defensible baseline of economic damage; a spatially explicit, two-and-a-half-year light record is that baseline, independent of any party’s reporting. The vegetation result cuts the other way and is just as useful: it warns against reading ecological stability as social stability. Khartoum’s fields looked fine from orbit while its people fled. Analysts who watch conflicts through environmental indices alone will miss wars like this one.

10 Limitations and next steps

The SAR gap is the hard limit. Two war years with no radar coverage force a four-year comparison and cap every structural claim at “consistent with.” The fix is concrete. Sentinel-1C now revisits Khartoum routinely, so from 2025 onward proper multi-temporal stacks exist: 12-day interferometric pairs, an urban mask built from the OSM footprints, coherence-based damage grading (Lu et al., 2018; Darvishi Bolorani et al., 2021; Yang et al., 2024; Huang et al., 2023b), and calibration against any building-level damage lists that humanitarian mappers release. For the missing 2023–2024 window itself, commercial SAR archives (tasking constellations that did image Khartoum) are the only remaining source; assessing their cost and coverage is on the roadmap.

Event geocoding bounds spatial inference. With 127 reference coordinates and 654 coarse-precision events, pixel-level tests against ACLED were always going to be blunt. Two upgrades sharpen them: restrict the randomization null to the urban mask, so the test asks the right question, and bring in building-level damage labels as the outcome variable once they exist. Until then, district-scale claims are the

ceiling, and I have kept every claim in this report under it.

Seasonality touches every two-window map. I quantified and corrected it for LST and controlled it for NTL; the Δ NDVI and Δ VV maps still carry it. Harmonic modelling of the full Sentinel-2 series — fit the annual cycle on pre-war months, analyze the residuals — is the clean fix, and the supplement’s scene-count figure (Figure S1) already shows where the series runs thin: August, the rainy season, every year.

The light estimate can tighten. A cross-check against NASA’s Black Marble products, which publish per-pixel uncertainty and quality fields (Wang et al., 2021), would harden the -65 to -67% interval, and their quality flags would settle the July 2022 artifact month formally.

The event window ends in March 2025. The satellite series run nine months longer than the event record in this version. Refreshing ACLED through 2026, and adopting UCDP GED v26 when it appears, extends the aligned window and enables the test this report sets up: does the Phase III light upturn coincide with falling violence? If yes, the recovery dial gains a validated interpretation; if no, that divergence would itself be a finding.

Transfer is the point. Nothing in the pipeline is Khartoum-specific except a bounding box and some dates. Porting it means changing the AOI and the windows, keeping the display and QA rules, and re-tuning only the cloud thresholds and masks. Myanmar, Gaza, and the next unwatched war can reuse it as-is, and the natural product format is a monthly bulletin per city rather than a one-off report. After this revision — numbers frozen, variants published, nulls included — I consider the template ready to port.

11 What each band can see: a closing guide

I end with the part of this study I most wish I had found in someone else’s paper when I started: a plain statement of which band answers which question, at what fidelity, and with what computation. Table 4 compresses the whole report into that form. Every row is backed by a result above; none of it is speculative.

Two summary observations follow from the table, and they are the practical conclusion of this report.

First, the fidelity ordering. For an urban war watched with free data, the observation channels rank by usefulness roughly as: night-time lights first (human activity, monthly, robust), the event datasets second (causes and tempo), fires third (timing of destruction), thermal fourth (slow environmental confirmation), optical NDVI fifth for a city (indispensable for the farm belt, misleading as a lead indicator), and SAR intensity last in this configuration — through no fault of the sensor, but because the acquisition gap and the geocoding floor cut its evidentiary power. With a dense SAR stack and coherence processing, that ranking would change, and SAR would likely move to the top for structural questions.

Second, the computations that made the difference. The same few operations recur behind every defensible number here: composite with medians, work in decibels for radar, match calendar months for anything with an annual cycle, clip display scales symmetrically at the 98th percentile, test co-location with explicit permutation nulls before claiming correlation, and freeze the statistics files before writing so that every figure, table, and sentence draws on the same numbers. None of these steps is advanced. Together they separate a claim that survives review from a map that merely looks convincing.

The fusion deliverables — the decoupling diagnosis, the district-scale convergence map, the attribution test and its null, the co-location tests, the dual-source validation of tempo — exist only because the channels sat in one pipeline and could be asked about each other. That is the case for multi-source work in one sentence: single sensors produce measurements; the combination produces checks, and the checks are what a reader can trust.

Table 4: Observation guide distilled from this study: target, best free source, computation that worked, and the honest verdict on what it supports.

Target	Best free source (band)	Computation that worked here	Verdict at Khartoum
Economic activity, electrification	VIIRS DNB monthly (panchromatic low-light, 500 m)	AOI-mean monthly series; window means + same-month matching; year-on-year comparison	Strongest signal in the study (−65 to −67%); monthly recovery dial; cannot separate depopulation from grid failure
Structural / surface change	Sentinel-1 C-band VV intensity (30 m here)	Multi-scene dB composites; differencing; tail share (< -3 dB); permutation test vs. events	Screening layer only; 7.8% severe-loss tail; pixel-scale attribution failed ($p = 0.556$); coherence needed for buildings
Farmland, vegetation	Sentinel-2 B4/B8 NDVI (10 m)	Cloud-masked median composites; Δ + severity classes; monthly series with $n \geq 3$ scenes	Maps the cultivated-belt decline (8.4% of area); under-reports urban war by construction
Ecosystem baseline	MODIS MOD13A2 NDVI (1 km)	Long homogeneous monthly series	The seasonality reference that makes the decoupling test possible
Burning, destruction alerts	VIIRS 375 m active fire (MWIR/LWIR)	Confidence filtering; dating; co-location tests against other layers	Near-real-time and datable; sparse; co-located with combat, correlation not established ($p = 0.103$)
Surface heat, habitability	Landsat TIRS ST_B10 (100 m)	Calendar-month matched means; QC of cloud-contaminated months; harmonic check	City-mean +2.5 to +3.0 °C; raw window difference misleading; pixel attribution out of reach
Combat tempo, geography	ACLED + UCDP GED (coded reporting)	Window alignment; cross-dataset correlation; KDE at district scale	Tempo robust across datasets ($r = 0.60$); levels are coding-dependent; spatial floor $\approx$ district

Data and code availability

The derived monthly series, the statistics tables behind Tables 1–4, and the full figure set are available from the author on request (hu964@purdue.edu); a public data archive is planned to accompany a later version of this work. Raw ACLED event data cannot be redistributed under ACLED’s terms of use; researchers can retrieve them free of charge from <https://acleddata.com>, and my repository includes the download and filtering scripts needed to reproduce the extract. UCDP GED is available from <https://ucdp.uu.se>. Sentinel, Landsat, VIIRS, MODIS, and FIRMS data are open; all satellite processing ran in Google Earth Engine (Gorelick et al., 2017).

Acknowledgements

This study contains modified Copernicus Sentinel data (2021–2025). I thank the providers of the open data this report stands on: ESA/Copernicus, NASA (VIIRS, MODIS, FIRMS), USGS (Landsat), ACLED (<https://acleddata.com>), the Uppsala Conflict Data Program, and the OpenStreetMap contributors (building data © OpenStreetMap contributors, available under the Open Database Licence). Basemaps in the figures are credited in each caption.

AI assistance disclosure

I used Anthropic’s Claude models as assistants for code development, reference verification, figure production, and language editing. I specified the analyses, checked every reported number against the underlying data files, and I take full responsibility for the content of this report.

Competing interests

I declare no competing interests.

References

- Agnew, J., Gillespie, T. W., Gonzalez, J., and Min, B. (2008). Baghdad nights: Evaluating the US military ‘surge’ using nighttime light signatures. *Environment and Planning A: Economy and Space*, 40(10):2285–2295. doi:10.1068/a41200.
- Ahmed, A., Rotich, B., and Czimber, K. (2024). Assessment of the environmental impacts of conflict-driven Internally Displaced Persons: A Sentinel-2 satellite based analysis of land use/cover changes in the Kas locality, Darfur, Sudan. *PLOS ONE*, 19(5):e0304034. doi:10.1371/journal.pone.0304034.
- Aimaiti, Y., Sanon, C., Koch, M., Baise, L. G., and Moaveni, B. (2022). War related building damage assessment in Kyiv, Ukraine, using Sentinel-1 radar and Sentinel-2 optical images. *Remote Sensing*, 14(24):6239. doi:10.3390/rs14246239.
- Basheer, M. and Elagib, N. A. (2024). Armed conflict as a catalyst for increasing flood risk. *Environmental Research Letters*, 19(10):104034. doi:10.1088/1748-9326/ad6fb6.
- Bennett, M. M., Van Den Hoek, J., Zhao, B., and Prishchepov, A. V. (2022). Improving satellite monitoring of armed conflicts. *Earth’s Future*, 10(9):e2022EF002904. doi:10.1029/2022EF002904.
- Bouchard, I., Rancourt, M.-È., Aloise, D., and Kalaitzis, F. (2022). On transfer learning for building damage assessment from satellite imagery in emergency contexts. *Remote Sensing*, 14(11):2532. doi:10.3390/rs14112532.
- Copernicus Programme (2025). Sentinel-1C user data opening. Sentinel Online news, 25 March 2025. <https://sentinels.copernicus.eu/-/sentinel-1c-user-data-opening> (accessed 21 August 2026).
- Darvishi Boloorani, A., Darvishi, M., Weng, Q., and Liu, X. (2021). Post-war urban damage mapping using InSAR: The case of Mosul city in Iraq. *ISPRS International Journal of Geo-Information*, 10(3):140. doi:10.3390/ijgi10030140.
- Elvidge, C. D., Baugh, K., Zhizhin, M., Hsu, F. C., and Ghosh, T. (2017). VIIRS night-time lights. *International Journal of Remote Sensing*, 38(21):5860–5879. doi:10.1080/01431161.2017.1342050.
- Elvidge, C. D., Ghosh, T., Hsu, F.-C., Zhizhin, M., and Bazilian, M. (2020). The dimming of lights in China during the COVID-19 pandemic. *Remote Sensing*, 12(17):2851. doi:10.3390/rs12172851.
- Enenkel, M., Shrestha, R. M., Stokes, E., Román, M., Wang, Z., Espinosa, M. T. M., Hajzmanova, I., Ginnetti, J., and Vinck, P. (2020). Emergencies do not stop at night: Advanced analysis of displacement based on satellite-derived nighttime light observations. *IBM Journal of Research and Development*, 64(1/2):8:1–8:12. doi:10.1147/JRD.2019.2954404.
- European Space Agency (2022). Mission ends for Copernicus Sentinel-1B satellite. ESA announcement, 3 August 2022. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-1/Mission_ends_for_Copernicus_Sentinel-1B_satellite (accessed 21 August 2026).

- European Space Agency (2024). Double win for Europe: Sentinel-1C and Vega-C take to the skies. ESA press release, 6 December 2024. https://www.esa.int/Newsroom/Press_Releases/Double_win_for_Europe_Sentinel-1C_and_Vega-C_take_to_the_skies (accessed 21 August 2026).
- Ge, P., Gokon, H., and Meguro, K. (2020). A review on synthetic aperture radar-based building damage assessment in disasters. *Remote Sensing of Environment*, 240:111693. doi:10.1016/j.rse.2020.111693.
- Gibson, J. and Boe-Gibson, G. (2021). Nighttime lights and county-level economic activity in the United States: 2001 to 2019. *Remote Sensing*, 13(14):2741. doi:10.3390/rs13142741.
- Gibson, J., Olivia, S., Boe-Gibson, G., and Li, C. (2021). Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149:102602. doi:10.1016/j.jdeveco.2020.102602.
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., and Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18–27. doi:10.1016/j.rse.2017.06.031.
- Gupta, R., Hosfelt, R., Sajeev, S., Patel, N., Goodman, B., Doshi, J., Heim, E., Choset, H., and Gaston, M. (2019). xBD: A dataset for assessing building damage from satellite imagery. arXiv preprint arXiv:1911.09296; also in Proceedings of the IEEE/CVF CVPR Workshops 2019. doi:10.48550/arXiv.1911.09296.
- Huang, C., Hong, S., Niu, X., Wu, Q., Zhong, Y., Yang, H., and Zhang, H. (2023a). Mapping of nighttime light trends and refugee population changes in Ukraine during the Russian–Ukrainian war. *Frontiers in Environmental Science*, 11:1055100. doi:10.3389/fenvs.2023.1055100.
- Huang, Q., Jin, G., Xiong, X., Ye, H., and Xie, Y. (2023b). Monitoring urban change in conflict from the perspective of optical and SAR satellites: The case of Mariupol, a city in the conflict between RUS and UKR. *Remote Sensing*, 15(12):3096. doi:10.3390/rs15123096.
- Lu, C.-H., Ni, C.-F., Chang, C.-P., Yen, J.-Y., and Chuang, R. Y. (2018). Coherence difference analysis of Sentinel-1 SAR interferogram to identify earthquake-induced disasters in urban areas. *Remote Sensing*, 10(8):1318. doi:10.3390/rs10081318.
- Malmgren-Hansen, D., Sohnesen, T., Fisker, P., and Baez, J. (2020). Sentinel-1 change detection analysis for cyclone damage assessment in urban environments. *Remote Sensing*, 12(15):2409. doi:10.3390/rs12152409.
- Mueller, H., Groeger, A., Hersh, J., Matranga, A., and Serrat, J. (2021). Monitoring war destruction from space using machine learning. *Proceedings of the National Academy of Sciences*, 118(23):e2025400118. doi:10.1073/pnas.2025400118.
- Papadimitriou, K., Koumoulidis, D., Papalamprou, L., Kasimatis, C., Sparangis, P., Katsenios, N., Vlachakis, D., Triantakostas, D., and Efthimiadou, A. (2021). Environmental impacts of war's social consequences. Case study: Aleppo Governorate, Syria. *EMBnet.journal*, 26(1):e965. doi:10.14806/ej.26.1.965.
- Poole, D. N., Andersen, D., Raymond, N. A., Grace, R., Smith, T., Khoshnood, K., and Mowafi, H. (2024). Damage to medical complexes in the Gaza Strip during the Israel– Hamas war: A geospatial analysis. *BMJ Global Health*, 9(4):e014768. doi:10.1136/bmjgh-2023-014768.
- Raleigh, C., Linke, A., Hegre, H., and Karlsen, J. (2010). Introducing ACLED: An armed conflict location and event dataset. *Journal of Peace Research*, 47(5):651–660. doi:10.1177/0022343310378914.

- Sticher, V., Wegner, J. D., and Pfeifle, B. (2023). Toward the remote monitoring of armed conflicts. *PNAS Nexus*, 2(6):pgad181. doi:10.1093/pnasnexus/pgad181.
- Sun, X., Chen, X., Yang, L., Wang, W., Zhou, X., Wang, L., and Yao, Y. (2022). Using InSAR and PolSAR to assess ground displacement and building damage after a seismic event: Case study of the 2021 Baicheng earthquake. *Remote Sensing*, 14(13):3009. doi:10.3390/rs14133009.
- Sundberg, R. and Melander, E. (2013). Introducing the UCDP georeferenced event dataset. *Journal of Peace Research*, 50(4):523–532. doi:10.1177/0022343313484347.
- Tomchenko, O. V., Khyzhniak, A. V., Sheviakina, N. A., Zahorodnia, S. A., Yelistratova, L. A., Yakovenko, M. I., and Stakhiv, I. R. (2023). Assessment and monitoring of fires caused by the war in Ukraine on landscape scale. *Journal of Landscape Ecology*, 16(2):76–97. doi:10.2478/jlecol-2023-0011.
- Wang, L., Lei, H., and Xu, H. (2024). Analysis of nighttime light changes and trends in the 1-year anniversary of the Russia–Ukraine conflict. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17:4084–4099. doi:10.1109/JSTARS.2024.3357727.
- Wang, Z., Román, M. O., Kalb, V. L., Miller, S. D., Zhang, J., and Shrestha, R. M. (2021). Quantifying uncertainties in nighttime light retrievals from Suomi-NPP and NOAA-20 VIIRS Day/Night Band data. *Remote Sensing of Environment*, 263:112557. doi:10.1016/j.rse.2021.112557.
- Weldegebriel, L., Negash, E., Nyssen, J., and Lobell, D. B. (2024). Eyes in the sky on Tigray, Ethiopia – monitoring the impact of armed conflict on cultivated highlands using satellite imagery. *Science of Remote Sensing*, 9:100133. doi:10.1016/j.srs.2024.100133.
- Yang, Y., Xie, C., Tian, B., Guo, Y., Zhu, Y., Yang, Y., Fang, H., Bian, S., and Zhang, M. (2024). Large-scale building damage assessment based on recurrent neural networks using SAR coherence time series: A case study of 2023 Turkey–Syria earthquake. *Earthquake Spectra*, 40(4):2285–2305. doi:10.1177/87552930241262761.
- Zhang, L., Hou, Z., Qu, Y., Shi, Q., Yu, Q., Zhao, Y., Tang, H., Wang, Y., Li, X., Li, Y., Peng, S., Yao, X., and Zhou, C. (2023). Satellite monitoring of war urban damage with a temporal knowledge-guided deep learning scheme. Preprint, Research Square. doi:10.21203/rs.3.rs-2799404/v1.
- Zubkova, M., Giglio, L., Humber, M. L., Hall, J. V., and Ellicott, E. (2021). Conflict and climate: Drivers of fire activity in Syria in the twenty-first century. *Earth Interactions*, 25(1):119–135. doi:10.1175/EI-D-21-0009.1.

Supplementary material

Monthly Sentinel-2 composites are only as good as the scenes behind them. Figure S1 counts the valid scenes entering each monthly NDVI composite, so readers can judge where the series runs thin.




Figure S1: Sentinel-2 scenes per monthly NDVI composite, 2022–2025. The red bar marks August 2022 ($n = 2$), the only month below the $n = 3$ threshold (dashed line). Data: Copernicus Sentinel-2.