UDC 004.9:004.4 DOI: https://doi.org/10.64076/iedc250905.16

Real-time radiation anomaly mapping using Kp/GOES X-ray correction and H3 spatial aggregation

Maksym Ilin
National Technical University of Ukraine "Igor Sikorsky Kyiv polytechnic institute", Kyiv
https://orcid.org/0009-0001-0803-3726
Liubov Oleshchenko
National Technical University of Ukraine "Igor Sikorsky Kyiv polytechnic institute", Kyiv
https://orcid.org/0000-0001-9908-7422

Abstract. *Solar disturbances often raise background gamma readings across large regions. Alerts that ignore space weather produce false positives and erode trust. This paper describes developed real-time software pipeline that ingests public dose-rate data, standardizes station signals, and corrects anomaly scores using the planetary K-index and GOES X-ray flux. Outputs aggregate to H3 tiles in DuckDB and render as daily maps for operations.*
Keywords: *real-time software pipeline, radiation analytics, Kp index, GOES X-ray, H3, DuckDB.*

Emergency analysts need alerts that trigger on local events, not on global swings. Geomagnetic storms and solar flares change the upper atmosphere and can shift background gamma measurements over broad areas. Without a space-weather correction, station-level thresholds fire together and mimic a widespread hazard. That wastes attention when focus should be on local outliers.

The software system reported here ingests public dose-rate feeds, produces robust station signals, and applies a space-weather correction tied to two operational indices: the planetary K-index (Kp) and GOES X-ray flux. Kp summarizes geomagnetic disturbance at a 3-hour cadence; GOES X-ray flux tracks flare activity at 1-minute cadence and marks ionospheric impacts relevant for short-term spikes.

After correction, anomalies reflect station-local deviations rather than global uplift. The stack uses H3 to tile results and DuckDB for fast, in-process analytics, with a thin web UI for daily review [1–4, 5–7, 8–9].

Radiation inputs combine official and community sources. EURDEP publishes hourly gamma dose-rate averages from thousands of fixed stations across Europe; the interface exposes near-real-time maps and recent histories suitable for micro-batch pull. SaveEcoBot provides station and volunteer counter data across Ukraine with CPM or μSv/h values and device metadata when available. Safecast contributes large volumes of community measurements with open access for aggregation. All records carry station identifiers, coordinates, units, and provenance fields that remain attached through aggregation [1–2, 5]. Space-weather drivers arrive from NOAA SWPC. Kp is a composite planetary index derived from ground magnetometers and published with near-real-time updates. GOES X-ray flux reports 1-minute averages in two passbands

and highlights flare-related surges that correlate with global background shifts. These feeds are pulled on their native cadence and joined to station observations by timestamp. The combination gives a coarse (Kp) and a fine (X-ray) handle on solar influence [3–4]. Data land in an in-process DuckDB database. Raw events append to events_live; transformed features and daily products materialize as views and tables. H3 provides hexagonal tiling and neighborhood joins; tiles at resolution 7 balance spatial detail with daily stability. The UI reads the materialized daily surface and renders tile layers and station pop-ups via Streamlit and PyDeck/deck.gl [6–9].

Station streams vary in noise and calibration. Hourly medians reduce spikes from short outages, message bursts, or counters stepping. A rolling 14-day window builds a baseline for each station, producing a location-specific mean and spread that update slowly. The z-score normalizes units and variance and lets downstream logic use one threshold family for different sensors and networks. EURDEP ADER values (e.g., nSv/h) mix cleanly with SaveEcoBot CPM/μSv/h once standardized at the score level; raw units and device metadata remain available for audit [1–2].

Solar activity moves many stations together. To discount this common component, the anomaly score subtracts an online linear response to Kp and GOES X-ray flux. Coefficients update by a small stochastic step at each timestamp, which adapts to regime shifts without chasing noise. Kp supplies broad disturbance context at 3-hour resolution; X-ray flux supplies minute-scale flare information that often matches short spikes. Corrected scores then emphasize local deviations, not the global background uplift during storms [3–4]. Implementation used in the service:

Table 1. Station anomaly computation and solar correction

```
# DataFrame df: ['station','ts','lat','lon','value','kp','xray'] in UTC
import numpy as np, pandas as pd
df['ts'] = pd.to_datetime(df['ts'])
df = df.sort_values(['station','ts'])
def _roll_stats(g):
 s = g.set_index('ts').sort_index()
 mu = s['value'].rolling('14D').median()
 sg = s['value'].rolling('14D').std().clip(lower=1e-6)
 return s.assign(mu=mu, sigma=sg).reset_index()
df = df.groupby('station', group_keys=False).apply(_roll_stats)
df['z'] = (df['value'] - df['mu']) / df['sigma']
# Online correction on Kp and X-ray
b1 = b2 = 0.0; lr = 0.01
for i in range(len(df)):
 pred = b1*df.at[i,'kp'] + b2*df.at[i,'xray']
 err = df.at[i,'z'] - pred
 b1 += lr * err * df.at[i,'kp']
 b2 += lr * err * df.at[i,'xray']
df['z_corr'] = df['z'] - (b1*df['kp'] + b2*df['xray'])
```

Corrected observations write back to events_live and roll up into a daily H3 surface (gamma_anom_h3) containing the mean corrected score per tile.

The aggregation reduces sensitivity to single outliers and yields a stable, interpretable map layer for operations. Analysts view station points over the tile surface and use consistent thresholds for daily review and alerts. H3 indexing and DuckDB SQL make the daily build fast enough to run inside the service without external warehousing [6–7]. DuckDB statement used in production:

Table 2. Daily H3 aggregation of corrected anomalies

```
-- events_live(ts, lat, lon, z_corr) -> gamma_anom_h3(d, h3_r7, z_corr_mean)
LOAD h3;
CREATE OR REPLACE TABLE gamma_anom_h3 AS
SELECT h3_geo_to_h3(lat, lon, 7) AS h3_r7,
 date_trunc('day', ts) AS d,
 avg(z_corr) AS z_corr_mean
FROM events_live
GROUP BY 1,2;
```

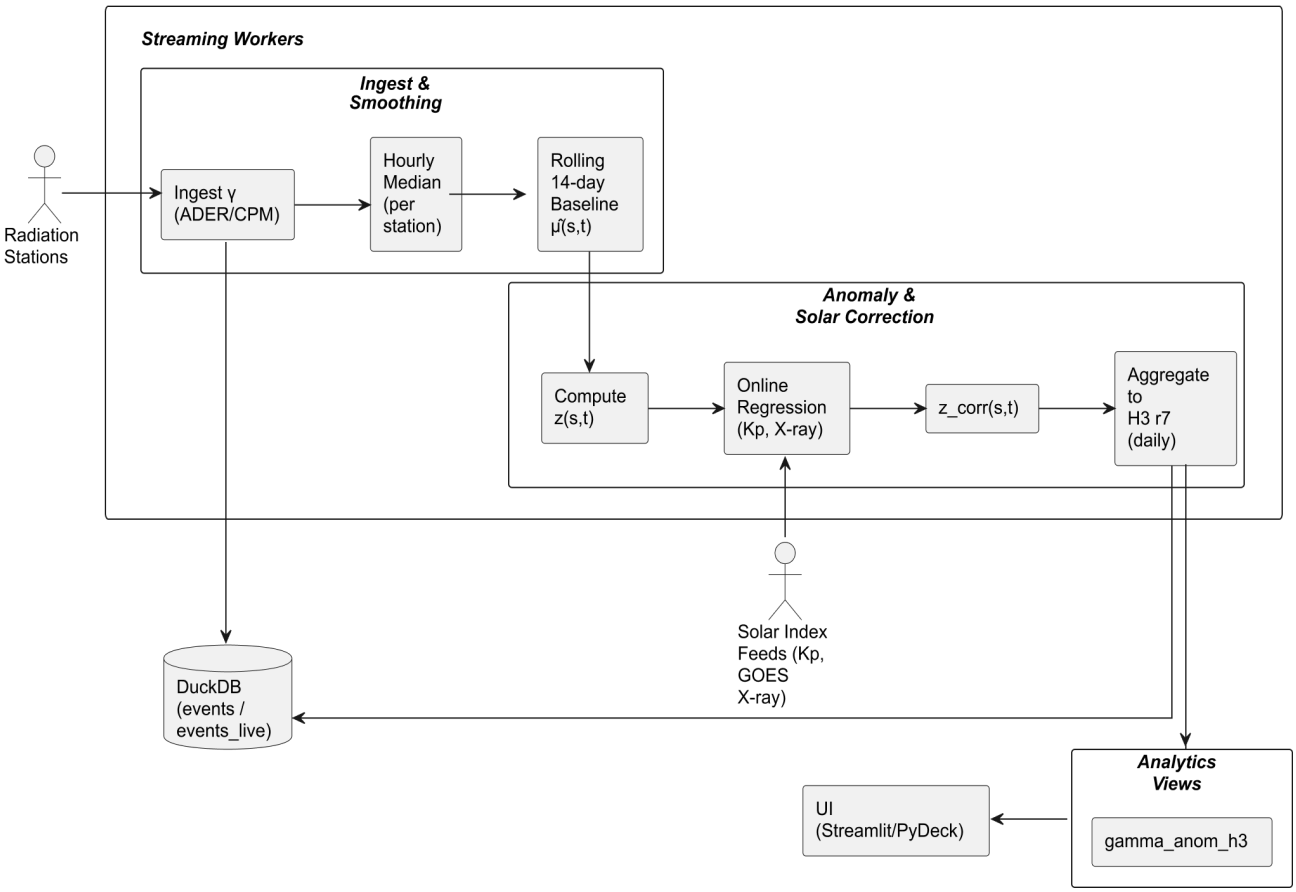

Fig. 1. Real-time pipeline from station ingest to solar-aware anomaly tiles and UI

Micro-batches pull EURDEP dose-rate values on their published cadence and SaveEcoBot streams as updates arrive. Records pass a light validation step, then land in DuckDB (Fig. 1). Hourly medians and the 14-day rolling baseline run inside

DuckDB and pandas for clarity, and the solar correction applies on the joined Kp and GOES tables. Daily H3 tiles materialize as a single table.

The UI layer loads gamma_anom_h3 and draws (a) the tile surface for the daily mean corrected score and (b) station points with pop-ups showing source, unit, sensor type, and a short sparkline of recent values. Streamlit's st.pydeck_chart embeds deck.gl layers directly, so the front end stays minimal while offering smooth interaction at city to national scales [8–9].

During geomagnetic disturbances, raw z-scores climb together across many stations. After correction, the daily tile surface returns to a muted background while stations with genuine local excursions continue to stand out. This separation improves triage: alerts reflect station-specific deviations rather than global uplift. Analysts track the corrected surface day over day, check outliers against source pop-ups, and export tiles as Parquet/GeoJSON for reports. The process relies only on public feeds and runs entirely in process, keeping deployment simple [1–4, 6–9].

Sensor networks mix counters, calibrations, and maintenance habits. Hourly medians reduce, but do not erase, device artifacts. Very strong solar events can leave residual influence even after Kp/X-ray correction. Extending the regression with additional space-weather terms is straightforward; neutron monitor counts, for example, offer another proxy for cosmic radiation variability. Time-of-day modeling may help separate human activity from environmental baselines in urban clusters. H3 resolution can be tuned where station density is high, and daily products can be complemented with intra-day surfaces for faster review cycles [5].

A solar-aware anomaly pipeline reduces false radiation alarms during space-weather events and highlights local deviations that matter operationally. Public inputs (EURDEP, SaveEcoBot, GOES X-ray, Kp) and an H3/DuckDB core keep the build reproducible and fast. Daily corrected tiles, station pop-ups, and simple thresholds give analysts a clear view without heavy infrastructure.

References

1. European Commission JRC. Radiological maps (EURDEP/REMon). https://remap.jrc.ec.europa.eu.

2. SaveEcoBot. Radiation maps and station data (Ukraine). https://www.saveecobot.com/en/radiation-maps.

3. NOAA Space Weather Prediction Center. Planetary K-index. https://www.swpc.noaa.gov/products/planetary-k-index.

4. NOAA Space Weather Prediction Center. GOES X-ray flux. https://www.swpc.noaa.gov/products/goes-x-ray-flux.

5. Safecast. Data portal. https://safecast.org/data.

6. H3. Introduction and docs. https://h3geo.org/docs.

7. DuckDB Foundation. DuckDB documentation. https://duckdb.org/docs/stable/

8. Streamlit. st.pydeck_chart. https://docs.streamlit.io/develop/api-reference/charts/st.pydeck_chart.

9. deck.gl. Documentation. https://deckgl.readthedocs.io.