

GCC CO₂ Emissions – Final Hybrid Satellite Modelling Notebook

This notebook represents the **final, consolidated experiment** of the Master's thesis on AI-based nowcasting of national CO₂ emissions in the GCC region.

It was developed **after multiple earlier notebooks and iterations**, which can be summarised as follows:

- **Baseline Phase (Saudi-only, economic indicators)**

The first experiments focused on Saudi Arabia using annual macro-economic indicators (GDP, population, primary energy consumption). Several classical and neural models (Linear Regression, Ridge, Random Forest, FFNN) were tested. These notebooks revealed the strong predictive power of economic variables but also highlighted data scarcity (≈ 30 yearly samples) and the risk of overfitting in single-country settings.

- **Panel Expansion Phase (GCC-wide, extended history)**

The next generation of notebooks extended the analysis to all six GCC countries and, in some cases, to the full historical period (1965–2023). This increased the effective sample size and allowed the use of panel models and tree-based ensembles. However, these experiments still relied mainly on historical socio-economic data and did not fully exploit the potential of Earth Observation.

- **Naïve Satellite Integration Phase (annual means, long history)**

A third series of notebooks integrated satellite metrics such as Land Surface Temperature (LST) and NDVI. Initially, these were aggregated as annual means over long time periods to match economic records. This approach led to several issues: (i) satellite missions do not cover the full historical record, and (ii) annual averaging washed out seasonal extremes, making the satellite signal weak and statistically dominated by population growth.

- **Refined Satellite Pipeline Phase (Google Earth Engine, seasonal features)**

Based on the lessons learned, a dedicated satellite extraction pipeline was implemented in Google Earth Engine. It focused on the modern satellite era (2000–2023) and on physically meaningful seasonal features (e.g. summer LST, winter LST, spring NDVI), night-time lights, AOD, and tropospheric NO₂. These features were stored in preprocessed CSV files and merged with emission and macro-economic data into a GCC-wide panel dataset.

This **final notebook** freezes the outcome of that evolution into a single, clean and reproducible experiment with three clearly defined models:

- **Model A – Economic Baseline:** macro-economic and demographic structure only.
- **Model B – Hybrid Model:** population density plus satellite-based indicators.
- **Model C – Satellite-Only Stress Test:** satellite indicators without any economic or demographic variables.

All three models are trained on **2000–2018** and evaluated on a strictly held-out period **2019–2023**, mirroring a real-world nowcasting deployment scenario.

```

import numpy as np
import pandas as pd

from sklearn.model_selection import TimeSeriesSplit, cross_val_score, GridSearchCV
from sklearn.linear_model import LinearRegression, Ridge
from sklearn.ensemble import RandomForestRegressor
from sklearn.metrics import r2_score, mean_squared_error
from sklearn.impute import SimpleImputer

from xgboost import XGBRegressor

import matplotlib.pyplot as plt
import seaborn as sns

plt.style.use("seaborn-v0_8")
sns.set()
np.random.seed(42)

```

1. Data Loading and Integration

This section loads the preprocessed emission and satellite feature tables and merges them into a single GCC-wide panel dataset at the (country, year) level.

- gcc_emissions_preprocessed_1965_2023.csv**
 contains national CO₂ emissions and macro-economic indicators. The underlying series were originally downloaded from internationally recognised repositories such as the Global Carbon Project (via Our World in Data) for territorial CO₂ emissions, and World Bank-type statistical databases for GDP, population, primary energy consumption, energy use per capita, CO₂ per unit of GDP, and land area. Population density (**pop_density**) and its log-transform (**log_pop_density**) were derived from population and land area within this integrated panel.
- gcc_seasonal_satellite_v6_autosave.csv**
 contains seasonal aggregates of Earth Observation features (LST, NDVI, night lights, AOD, NO₂) engineered in Google Earth Engine. For each GCC country and year, satellite ImageCollections (e.g. MODIS LST, NDVI, night-time lights, aerosol optical depth, and Sentinel-5P NO₂ where available) were spatially averaged over national polygons and temporally aggregated to annual or seasonal indicators.

After loading both tables, they are merged on the common **(country, year)** key to form a single GCC-wide panel (**df_modern**) that combines economic, demographic and satellite features for the period 2000–2023.

1.1 Data Latency in Official CO₂ Inventories (Motivation for Nowcasting)

National CO₂ inventories are not produced in real time. Under current international reporting practice, countries submit annual greenhouse gas inventories that typically cover emissions only up to two years before the submission year. Global statistical products such as the World Bank's World Development Indicators and the CO₂ / GHG datasets used in this notebook are built on top of those national inventories and energy statistics.

In practical terms, this means that:

- The panel used in this notebook currently covers **2000–2023**, while the thesis itself is being conducted in **2025**.
- Policy makers and analysts cannot observe the current year's national CO₂ emissions from official statistics; they only see values for years that are already 1–2 years in the past.

This structural **data latency** is the main motivation for the nowcasting framework developed here:

- **Model A (Economic Baseline)** reproduces the behaviour of traditional inventory-based models, but inherits the same delay.
- **Model B (Hybrid)** and **Model C (Satellite Only)** aim to estimate emissions from near-real-time activity and environmental indicators (satellite signals, density, temperature, lights), reducing dependence on delayed national statistics.

In [139...

```
satellite_file = "data/gcc_seasonal_satellite_v6_autosave.csv"
emissions_file = "data/gcc_emissions_preprocessed_1965_2023.csv"

df_satellite = pd.read_csv(satellite_file)
df_emissions = pd.read_csv(emissions_file)

# Basic cleaning of country names
df_satellite["country"] = df_satellite["country"].str.strip()
df_emissions["country"] = df_emissions["country"].str.strip()

# Merge on (country, year)
df_merged = pd.merge(
    df_emissions,
    df_satellite,
    on=["country", "year"],
    how="inner",
    validate="many_to_many"
)

print("Merged shape:", df_merged.shape)
df_merged.head()
```

Merged shape: (144, 27)

Out[139...

	country	year	iso_code	co2	gdp	population	primary_energy_consumption	energy_per_c
0	Bahrain	2000	BHR	19.199	1.103229e+10	669590.0	111.877	15725
1	Bahrain	2001	BHR	19.679	1.207499e+10	699964.0	115.790	15855
2	Bahrain	2002	BHR	20.112	1.378745e+10	752206.0	123.547	16509
3	Bahrain	2003	BHR	20.976	1.653529e+10	811278.0	129.081	16585
4	Bahrain	2004	BHR	21.204	1.963347e+10	871449.0	133.713	16043

5 rows × 27 columns



2. Basic Data Quality Checks

Before defining models, we perform basic checks on:

- temporal coverage,

- country coverage,
- and missing values in key feature groups.

```
In [140... print("Years in merged data:", df_merged["year"].min(), "to", df_merged["year"].max())
print("Countries:", df_merged["country"].nunique(), list(df_merged["country"].unique()))
print("Total rows:", len(df_merged))

missing_summary = df_merged.isna().sum()
print("\nNumber of missing values per column (non-zero only):")
print(missing_summary[missing_summary > 0].sort_values(ascending=False))
```

Years in merged data: 2000 to 2023

Countries: 6 ['Bahrain', 'Kuwait', 'Oman', 'Qatar', 'Saudi Arabia', 'United Arab Emirates']

Total rows: 144

Number of missing values per column (non-zero only):

NO2_Signal	114
NO2_Source	114
coal_co2	25
primary_energy_consumption	1
energy_per_capita	1

dtype: int64

2.1 Descriptive statistics and data completeness

This subsection summarises the main properties of the merged panel before the modern-satellite filtering. It provides a quick overview of the numerical ranges, the fraction of missing values and the temporal coverage by country. These diagnostics document the input to all subsequent modelling steps.

```
In [141... # Descriptive statistics for key variables in the merged panel
core_cols = [
    "co2",
    "gdp",
    "population",
    "primary_energy_consumption",
    "energy_per_capita",
    "co2_per_gdp",
]

available_core = [c for c in core_cols if c in df_merged.columns]
print("Core numerical variables included in the summary:", available_core)

if available_core:
    desc = df_merged[available_core].describe().T
    print("\nDescriptive statistics (df_merged):")
    print(desc)
else:
    print("No core numerical columns from the expected list were found in df_merged.")

# Missing-value profile for the same columns
if available_core:
    missing_frac = df_merged[available_core].isna().mean().sort_values(ascending=False)
    print("\nFraction of missing values (df_merged):")
    print(missing_frac.to_frame("missing_fraction"))

# Temporal coverage by country
if "country" in df_merged.columns and "year" in df_merged.columns:
    coverage = df_merged.groupby("country")["year"].nunique().sort_values(ascending=False)
```

```
print("\nNumber of distinct years per country in df_merged:")
print(coverage)
```

Core numerical variables included in the summary: ['co2', 'gdp', 'population', 'primary_energy_consumption', 'energy_per_capita', 'co2_per_gdp']

Descriptive statistics (df_merged):

	count	mean	std	min \
co2	144.0	1.590221e+02	1.836147e+02	1.919900e+01
gdp	144.0	3.871625e+11	4.570587e+11	1.103229e+10
population	144.0	7.100264e+06	8.901551e+06	6.449930e+05
primary_energy_consumption	143.0	7.760830e+02	8.493989e+02	1.109240e+02
energy_per_capita	143.0	1.299860e+05	5.056923e+04	4.731712e+04
co2_per_gdp	144.0	5.177847e-01	2.622499e-01	2.350000e-01

	25%	50%	75% \
co2	4.634050e+01	8.448100e+01	2.000277e+02
gdp	7.663712e+10	1.975958e+11	4.942812e+11
population	1.883080e+06	3.234920e+06	7.273859e+06
primary_energy_consumption	2.189890e+02	4.136530e+02	1.047515e+03
energy_per_capita	8.567732e+04	1.274108e+05	1.607719e+05
co2_per_gdp	3.785000e-01	4.365000e-01	5.942500e-01

	max
co2	7.362050e+02
gdp	1.823886e+12
population	3.326429e+07
primary_energy_consumption	3.221731e+03
energy_per_capita	2.635416e+05
co2_per_gdp	1.740000e+00

Fraction of missing values (df_merged):

	missing_fraction
primary_energy_consumption	0.006944
energy_per_capita	0.006944
co2	0.000000
gdp	0.000000
population	0.000000
co2_per_gdp	0.000000

Number of distinct years per country in df_merged:

country	
Bahrain	24
Kuwait	24
Oman	24
Qatar	24
Saudi Arabia	24
United Arab Emirates	24

Name: year, dtype: int64

3. Modern Satellite Era and Core Feature Engineering

- Restrict the analysis to the modern satellite era: 2000–2023.
- Compute population density and log-density (if not already present).
- Define the final feature sets for each model:
 - Model A: Economic Baseline.
 - Model B: Hybrid (density + satellite features).
 - Model C: Satellite-Only.

3.1 Earth Observation Data Sources

The satellite-based indicators used in this notebook are derived from a set of well-established Earth Observation (EO) products, accessed and processed through Google Earth Engine (GEE). They are designed to capture different physical dimensions of anthropogenic activity in the GCC region.

- **Land Surface Temperature (LST)**

LST is obtained from thermal infrared sensors onboard polar-orbiting satellites (e.g. MODIS on Terra and Aqua). These products provide regular global coverage at moderate spatial resolution and sub-monthly temporal frequency. In this study, LST is used to characterise near-surface thermal conditions as a proxy for electricity demand related to space cooling.

- **Vegetation Activity (NDVI)**

NDVI is derived from multispectral optical sensors and measures the greenness of vegetation. It is sensitive to vegetation density and phenology. Although NDVI is not a direct driver of emissions, it provides contextual information about land cover and potential carbon uptake, and serves as a secondary control variable in the GCC context.

- **Night-Time Lights**

Night-time lights are constructed from low-light imaging sensors (e.g. DMSP-OLS and VIIRS) that observe artificial illumination at the Earth's surface. Aggregated at the national level, they provide a robust proxy for urbanisation, economic activity and electricity consumption, particularly in rapidly developing regions.

- **Aerosol Optical Depth (AOD)**

AOD is retrieved from atmospheric reanalysis and satellite products (e.g. MERRA-2 and related datasets). It quantifies the columnar burden of aerosols in the atmosphere and is influenced by combustion, dust and other particulate sources. In this work it is used as a broad indicator of atmospheric loading associated with fossil-fuel use.

- **Tropospheric Nitrogen Dioxide (NO₂)**

NO₂ columns are obtained from satellite spectrometers that observe back-scattered solar radiation (e.g. Sentinel-5P TROPOMI). NO₂ is a short-lived pollutant closely associated with high-temperature combustion in the power, transport and industrial sectors. At the national scale, it provides a direct atmospheric footprint of recent fossil-fuel emissions.

All EO variables in this experiment are restricted to the **modern satellite era (2000–2023)**, consistent with the temporal coverage of the derived panel dataset.

3.2 Satellite Feature Engineering and Aggregation

Raw EO products are available at high spatial and temporal resolution. To make them comparable with national CO₂ inventories and macro-economic indicators, a set of aggregation and feature-engineering steps is applied:

1. **Spatial Aggregation to National Scale**

For each product, satellite observations are clipped to the national boundaries of the six GCC countries. Zonal statistics (primarily the spatial mean) are then computed for each country and time period. This produces a country-level time series that is directly comparable to national emission inventories.

2. Temporal Aggregation to Seasonal Indicators

Instead of using simple annual means, which tend to smooth out extremes, the following seasonal indicators are constructed:

- `LST_Summer_C` : mean LST for June–August, capturing peak cooling demand.
- `LST_Winter_C` : mean LST for December–February, representing the thermal base load.
- `NDVI_Spring` : mean NDVI for March–May, characterising vegetation conditions during the main growing season.
- `AOD_Annual` : mean AOD over the full calendar year.
- `Night_Lights` : annual composite of night-time light intensity.
- `NO2_Signal` : annual representative NO₂ column derived from Sentinel-5P and related products, homogenised into a consistent index.

3. Merging with Emission and Economic Data

The seasonal EO indicators are then joined with national CO₂ emissions and macro-economic variables on a common (country, year) key. The resulting panel `df_modern` covers the period 2000–2023 and forms the basis for the three models implemented in this notebook: the economic baseline (Model A), the hybrid satellite framework (Model B) and the satellite-only stress test (Model C).

This design ensures that the EO features used in the models are both physically interpretable and statistically aligned with the temporal resolution of official emission inventories.

In [142...

```
# Restrict to modern satellite era
df_modern = df_merged[(df_merged["year"] >= 2000) & (df_merged["year"] <= 2023)].copy()

# -----
# Population density and log-density (GCC countries, self-contained)
# -----

GCC_AREA_KM2 = {
    "Saudi Arabia": 2150000,      # ~2,150,000 km2
    "United Arab Emirates": 83600, # ~83,600 km2
    "Kuwait": 17818,             # 17,818 km2
    "Qatar": 11571,              # 11,571 km2
    "Bahrain": 787,              # ~786.8 km2
    "Oman": 309960,              # 309,960 km2
}

# Map fixed country areas
df_modern["area_km2"] = df_modern["country"].map(GCC_AREA_KM2)

# Sanity check
if df_modern["area_km2"].isna().any():
    missing_countries = df_modern.loc[df_modern["area_km2"].isna(), "country"].unique()
    raise ValueError(f"Missing area_km2 definition for countries: {missing_countries}")

# Compute density and log-density
df_modern["pop_density"] = df_modern["population"] / df_modern["area_km2"]
df_modern["log_pop_density"] = np.log1p(df_modern["pop_density"])

print(" pop_density and log_pop_density computed from fixed GCC land areas.")

# Final feature sets (aligned with the thesis text)
features_hybrid = [
    "log_pop_density",
    "Night_Lights",
    "LST_Summer_C",
```

```

        "LST_Winter_C",
        "NDVI_Spring",
        "AOD_Annual",
        "NO2_Signal",
    ]

    features_satellite_only = [
        "Night_Lights",
        "LST_Summer_C",
        "LST_Winter_C",
        "NDVI_Spring",
        "AOD_Annual",
        "NO2_Signal",
    ]

    features_economic = [
        "gdp",
        "population",
        "primary_energy_consumption",
        "energy_per_capita",
        "co2_per_gdp",
        "pop_density",
        "log_pop_density",
    ]

    target = "co2"

    print("Hybrid features:", features_hybrid)
    print("Satellite-only features:", features_satellite_only)
    print("Economic baseline features:", features_economic)
    print("Target:", target)

```

pop_density and log_pop_density computed from fixed GCC land areas.
 Hybrid features: ['log_pop_density', 'Night_Lights', 'LST_Summer_C', 'LST_Winter_C', 'NDVI_Spring', 'AOD_Annual', 'NO2_Signal']
 Satellite-only features: ['Night_Lights', 'LST_Summer_C', 'LST_Winter_C', 'NDVI_Spring', 'AOD_Annual', 'NO2_Signal']
 Economic baseline features: ['gdp', 'population', 'primary_energy_consumption', 'energy_per_capita', 'co2_per_gdp', 'pop_density', 'log_pop_density']
 Target: co2

4. Temporal Train–Test Split (Nowcasting Setup)

We simulate a real-world nowcasting deployment:

- Training period: 2000–2018
- Testing period: 2019–2023 (strictly unseen future years)

In [143...

```

# =====
# TRAIN / TEST SPLIT (PANEL 2000–2023, TEST STARTS AT 2019)
# =====

TEST_START = 2019

# Ensure deterministic ordering by year and country
df_modern = df_modern.sort_values(["year", "country"]).reset_index(drop=True) # the temporal split

# Train = all years < 2019, Test = 2019–2023
train_df = df_modern[df_modern["year"] < TEST_START].copy()
test_df = df_modern[df_modern["year"] >= TEST_START].copy()

```

```
TARGET_COL = "co2"
```

```
# Central definitions of y_train / y_test so that all later cells can reuse them
```

```
y_train = train_df[TARGET_COL]
```

```
y_test = test_df[TARGET_COL]
```

```
print("Train years :", train_df["year"].min(), "to", train_df["year"].max())
```

```
print("Test years :", test_df["year"].min(), "to", test_df["year"].max())
```

```
print("Train samples:", len(train_df))
```

```
print("Test samples :", len(test_df))
```

```
print("Countries      :", df_modern["country"].nunique(), list(df_modern["country"].unique()))
```

```
Train years : 2000 to 2018
```

```
Test years : 2019 to 2023
```

```
Train samples: 114
```

```
Test samples : 30
```

```
Countries : 6 ['Bahrain', 'Kuwait', 'Oman', 'Qatar', 'Saudi Arabia', 'United Arab Emirates']
```

4.1 Visual inspection of CO₂ trends and key indicators

Before fitting the models, it is useful to visualise the basic temporal behaviour of emissions and selected drivers. The following figures show:

- Country-level CO₂ trajectories over the modern satellite era.
- A correlation heatmap for the hybrid feature set used in the models.
- An example bivariate relationship between CO₂ and an atmospheric indicator (NO₂), where available.

These plots provide a qualitative cross-check that the engineered features are consistent with expected physical and socio-economic patterns.

In [144...

```
# 4.1.1 Country-level CO2 trajectories
```

```
if {"country", "year", "co2"} <= set(df_modern.columns):
```

```
    plt.figure(figsize=(10, 6))
```

```
    for country, g in df_modern.groupby("country"):
```

```
        g_sorted = g.sort_values("year")
```

```
        plt.plot(g_sorted["year"], g_sorted["co2"], marker="o", label=country)
```

```
    plt.xlabel("Year")
```

```
    plt.ylabel("CO2 emissions (Mt)")
```

```
    plt.title("CO2 emissions by country (modern satellite era)")
```

```
    plt.legend()
```

```
    plt.grid(True)
```

```
    plt.tight_layout()
```

```
    plt.show()
```

```
else:
```

```
    print("Columns 'country', 'year', 'co2' are required in df_modern for this plot.")
```

```
# 4.1.2 Correlation heatmap for hybrid features on the training period
```

```
corr_cols = [c for c in (features_hybrid + [target]) if c in train_df.columns]
```

```
if corr_cols:
```

```
    corr = train_df[corr_cols].corr()
```

```
    plt.figure(figsize=(8, 6))
```

```
    sns.heatmap(corr, annot=True, fmt=".2f", cmap="coolwarm", square=True)
```

```
    plt.title("Correlation matrix - hybrid feature set (training period)")
```

```
    plt.tight_layout()
```


```
    plt.show()
```

```
else:
```


```
    print("No overlapping hybrid features found in train_df for correlation analysis.")
```


```
# 4.1.3 Example scatter plot: N02 vs C02
```

```
if {"N02_Signal", "co2"} <= set(df_modern.columns):  
    plt.figure(figsize=(8, 6))  
    sns.scatterplot(data=df_modern, x="N02_Signal", y="co2", hue="country")  
    plt.xlabel("N02 signal (arbitrary units)")  
    plt.ylabel("CO2 emissions (Mt)")  
    plt.title("Relationship between N02 signal and CO2 emissions")  
    plt.grid(True)  
    plt.tight_layout()  
    plt.show()  
else:  
    print("Columns 'N02_Signal' and 'co2' are required in df_modern for this scatter plot.")
```



Correlation matrix – hybrid feature set (training period)





5. Model Selection with TimeSeriesSplit (Hybrid Features)

Here we evaluate several regression algorithms on the hybrid feature set using a time-aware cross-validation scheme (TimeSeriesSplit) over the training period.

In [145...

```
# =====
# 5. MODEL SELECTION WITH TimeSeriesSplit (HYBRID FEATURES)
#   Goal: evaluate several models on hybrid features using
#   a time-aware split, excluding NO2_Signal from CV because
#   NO2 is only available after 2019.
# =====

from sklearn.model_selection import TimeSeriesSplit, cross_val_score
from sklearn.linear_model import LinearRegression, Ridge
from sklearn.ensemble import RandomForestRegressor
from xgboost import XGBRegressor
import numpy as np
import pandas as pd

# 1) Define the feature set for CV (exclude NO2_Signal)
hybrid_cv_features = [f for f in features_hybrid if f != "NO2_Signal"]
hybrid_cv_cols = hybrid_cv_features + [target]

# 2) Drop rows with any missing values in the selected columns
train_hybrid_clean = train_df.dropna(subset=hybrid_cv_cols).copy()
print(f"Hybrid CV training samples (non-missing, NO2 excluded): {len(train_hybrid_clean)}")
```

```

# 3) Prepare TimeSeriesSplit
tscv = TimeSeriesSplit(n_splits=3)
min_required = tscv.n_splits + 1 # at least n_splits + 1 samples are required

results = []

if len(train_hybrid_clean) < min_required:
    print(
        f"Warning: not enough non-missing samples ({len(train_hybrid_clean)}) "
        f"for TimeSeriesSplit with n_splits={tscv.n_splits}. "
        "Hybrid CV metrics will be reported as NaN."
    )
    models = ["Linear Regression", "Ridge", "Random Forest", "XGBoost"]
    df_results_hybrid = pd.DataFrame(
        {
            "Model": models,
            "CV RMSE": [np.nan] * len(models),
            "CV R2": [np.nan] * len(models),
        }
    )
else:
    # 4) Build X, y arrays for CV
    X_train_hybrid = train_hybrid_clean[hybrid_cv_features].values
    y_train_hybrid = train_hybrid_clean[target].values

    # 5) Define candidate models
    models = {
        "Linear Regression": LinearRegression(),
        "Ridge": Ridge(alpha=1.0),
        "Random Forest": RandomForestRegressor(
            n_estimators=500, max_depth=10, random_state=42, n_jobs=-1
        ),
        "XGBoost": XGBRegressor(
            n_estimators=800,
            learning_rate=0.05,
            max_depth=6,
            subsample=0.8,
            colsample_bytree=0.8,
            random_state=42,
            n_jobs=-1,
        ),
    }

    # 6) Run CV for each model
    for name, model in models.items():
        rmse_scores = -cross_val_score(
            model,
            X_train_hybrid,
            y_train_hybrid,
            cv=tscv,
            scoring="neg_root_mean_squared_error",
            n_jobs=-1,
        )
        r2_scores = cross_val_score(
            model,
            X_train_hybrid,
            y_train_hybrid,
            cv=tscv,
            scoring="r2",
            n_jobs=-1,
        )

```

```

        results.append(
            {
                "Model": name,
                "CV RMSE": rmse_scores.mean(),
                "CV R2": r2_scores.mean(),
            }
        )

df_results_hybrid = pd.DataFrame(results).sort_values("CV RMSE")

# 7) Show results table
df_results_hybrid

```

Hybrid CV training samples (non-missing, N02 excluded): 114

Out[145]...

	Model	CV RMSE	CV R2
3	XGBoost	101.748680	0.580005
2	Random Forest	113.807842	0.468054
0	Linear Regression	142.890518	0.346873
1	Ridge	164.903494	0.123506

6. Model A – Economic Baseline

Macro-economic baseline model using only economic and demographic structure variables. This model serves as a strong reference benchmark based on fully reported historical data, against which hybrid and satellite-based alternatives can be compared (but delayed by 1–2 years in practice).

In [146]...

```

X_train_A = train_df[features_economic].values
X_test_A = test_df[features_economic].values
y_train_A = train_df[target].values
y_test_A = test_df[target].values

model_A = XGBRegressor(
    n_estimators=800,
    learning_rate=0.03,
    max_depth=4,
    subsample=0.8,
    colsample_bytree=0.8,
    random_state=42,
    n_jobs=-1,
)

model_A.fit(X_train_A, y_train_A)
preds_A = model_A.predict(X_test_A)

r2_A = r2_score(y_test_A, preds_A)
rmse_A = np.sqrt(mean_squared_error(y_test_A, preds_A))

print("ECONOMIC BASELINE MODEL (Model A) – Test 2019–2023")
print(f" Test R² : {r2_A:.4f}")
print(f" Test RMSE: {rmse_A:.2f} MtCO₂")

```

ECONOMIC BASELINE MODEL (Model A) – Test 2019–2023

Test R² : 0.9571

Test RMSE: 46.29 MtCO₂

7. Model B – Hybrid Model (Density + Satellite Features)

Hybrid model that combines population density with satellite-derived indicators:

- log_pop_density
- Night-time lights
- Summer and winter LST
- Spring NDVI
- Annual AOD
- NO₂ signal

In [147...

```
X_train_B = train_df[features_hybrid].values
X_test_B = test_df[features_hybrid].values
y_train_B = train_df[target].values
y_test_B = test_df[target].values

model_B = XGBRegressor(
    n_estimators=1000,
    learning_rate=0.01,
    max_depth=6,
    subsample=0.8,
    colsample_bytree=0.8,
    random_state=42,
    n_jobs=-1,
)

model_B.fit(X_train_B, y_train_B)
preds_B = model_B.predict(X_test_B)
preds_B_manual = preds_B # alias used in plotting cell

r2_B = r2_score(y_test_B, preds_B)
rmse_B = np.sqrt(mean_squared_error(y_test_B, preds_B))

print("HYBRID MODEL (Model B, manual parameters) - Test 2019-2023")
print(f"  Test R² : {r2_B:.4f}")
print(f"  Test RMSE: {rmse_B:.2f} MtCO₂")

imp_B = pd.DataFrame(
    {"Feature": features_hybrid, "Importance": model_B.feature_importances_}
).sort_values("Importance", ascending=False)
imp_B
```

HYBRID MODEL (Model B, manual parameters) - Test 2019-2023

Test R² : 0.6916

Test RMSE: 124.17 MtCO₂

Out[147...

	Feature	Importance
2	LST_Summer_C	0.337243
3	LST_Winter_C	0.222876
1	Night_Lights	0.219002
0	log_pop_density	0.153086
4	NDVI_Spring	0.056655
5	AOD_Annual	0.011138
6	NO2_Signal	0.000000

8. Model B – Tuned Hybrid Model (TimeSeriesSplit + GridSearchCV) [Optional]

Optional hyperparameter tuning using a time-aware cross-validation scheme. This often leads to a more regularised model and slightly lower test performance, but it is useful for understanding the trade-off between complexity and robustness.

In [148...

```
y_train_B_log = np.log1p(y_train_B)

param_grid = {
    "n_estimators": [500, 1000],
    "learning_rate": [0.01, 0.03],
    "max_depth": [3, 5],
    "subsample": [0.7, 0.8],
    "colsample_bytree": [0.6, 0.8],
}

xgb_base = XGBRegressor(
    random_state=42,
    n_jobs=-1,
)

tscv = TimeSeriesSplit(n_splits=3)

grid_search = GridSearchCV(
    estimator=xgb_base,
    param_grid=param_grid,
    cv=tscv,
    scoring="neg_root_mean_squared_error",
    n_jobs=-1,
    verbose=1,
)

grid_search.fit(X_train_B, y_train_B_log)

best_tuned_model = grid_search.best_estimator_

print("Best parameters (Hybrid, tuned):", grid_search.best_params_)
print("Best CV RMSE (log-space):", -grid_search.best_score_)

preds_B_log = best_tuned_model.predict(X_test_B)
preds_B_tuned = np.expm1(preds_B_log)
```

```

r2_B_tuned = r2_score(y_test_B, preds_B_tuned)
rmse_B_tuned = np.sqrt(mean_squared_error(y_test_B, preds_B_tuned))

print("\nTUNED HYBRID MODEL (Model B, tuned) - Test 2019-2023")
print(f"  Test R² : {r2_B_tuned:.4f}")
print(f"  Test RMSE: {rmse_B_tuned:.2f} MtCO₂")

```

Fitting 3 folds for each of 32 candidates, totalling 96 fits

Best parameters (Hybrid, tuned): {'colsample_bytree': 0.6, 'learning_rate': 0.03, 'max_depth': 3, 'n_estimators': 1000, 'subsample': 0.8}

Best CV RMSE (log-space): 0.5471853461337916

TUNED HYBRID MODEL (Model B, tuned) - Test 2019-2023

Test R² : 0.6231

Test RMSE: 137.27 MtCO₂

9. Model C – Satellite-Only Stress Test

This model removes all economic and demographic variables and relies purely on satellite-derived indicators:

- Night-time lights
- Summer and winter LST
- Spring NDVI
- Annual AOD
- NO₂ signal

It is not intended as an operational replacement for national inventories, but as a stress test to quantify how much of the variance in national emissions can be explained using Earth Observation alone.

In [149...

```

X_train_C = train_df[features_satellite_only].values
X_test_C = test_df[features_satellite_only].values
y_train_C = train_df[target].values
y_test_C = test_df[target].values

# Handle potential missing values explicitly
imputer_C = SimpleImputer(strategy="median")
X_train_C_imp = imputer_C.fit_transform(X_train_C)
X_test_C_imp = imputer_C.transform(X_test_C)

model_C = XGBRegressor(
    n_estimators=800,
    learning_rate=0.03,
    max_depth=5,
    subsample=0.8,
    colsample_bytree=0.8,
    random_state=42,
    n_jobs=-1,
)

model_C.fit(X_train_C_imp, y_train_C)
preds_C = model_C.predict(X_test_C_imp)

r2_C = r2_score(y_test_C, preds_C)
rmse_C = np.sqrt(mean_squared_error(y_test_C, preds_C))

print("SATELLITE-ONLY MODEL (Model C) - Test 2019-2023")

```

```
print(f" Test R² : {r2_C:.4f}")
print(f" Test RMSE: {rmse_C:.2f} MtCO₂")
```

C:\Python312\Lib\site-packages\sklearn\impute_base.py:598: UserWarning: Skipping features without any observed values: [5]. At least one non-missing value is needed for imputation with strategy='median'.

```
warnings.warn(
```

C:\Python312\Lib\site-packages\sklearn\impute_base.py:598: UserWarning: Skipping features without any observed values: [5]. At least one non-missing value is needed for imputation with strategy='median'.

```
warnings.warn(
```

SATELLITE-ONLY MODEL (Model C) - Test 2019-2023

Test R² : 0.6668

Test RMSE: 129.07 MtCO₂

In [150...

```
# =====
# ACTUAL VS PREDICTED (2019-2023) - LINE AND SCATTER PLOTS
# =====
# NOTE:
# This cell assumes that Models A, B (hybrid), and C (satellite-only) have been trained
# and that preds_A, preds_B_manual (or preds_B), preds_C, y_test, and test_df are already in mem
# are all available in memory from previous cells.

import matplotlib.pyplot as plt


# Reuse objects computed earlier
# - y_test: defined in the train/test split cell
# - preds_A: Model A predictions
# - preds_B_manual: Model B (manual hybrid) predictions
# - preds_C: Model C (satellite-only) predictions

years_test = test_df["year"].values


plt.figure(figsize=(10, 5))
plt.plot(years_test, y_test, marker="o", linewidth=2, label="Actual CO2")
plt.plot(years_test, preds_A, marker="s", linestyle="--", linewidth=2, label="Model A - Economic")
plt.plot(years_test, preds_B_manual, marker="^", linestyle="--", linewidth=2, label="Model B - Hybrid")
plt.plot(years_test, preds_C, marker="d", linestyle=":", linewidth=2, label="Model C - Satellite")
plt.xlabel("Year")
plt.ylabel("CO2 emissions (Mt)")
plt.title("Actual vs Predicted CO2 Emissions - GCC Panel (2019-2023)")
plt.legend()
plt.grid(alpha=0.3)
plt.show()

# Scatter: Actual vs Predicted (Model B - Hybrid)
plt.figure(figsize=(6, 6))
plt.scatter(y_test, preds_B_manual, alpha=0.8)
max_val = max(y_test.max(), preds_B_manual.max())
plt.plot([0, max_val], [0, max_val], "k--", linewidth=2, label="y = x")
plt.xlabel("Actual CO2 (Mt)")
plt.ylabel("Predicted CO2 (Mt) - Model B Hybrid")
plt.title("Prediction Accuracy - Hybrid Model (2019-2023)")
plt.legend()
plt.grid(alpha=0.3)
plt.show()
```

Actual vs Predicted CO2 Emissions – GCC Panel (2019–2023)



Prediction Accuracy – Hybrid Model (2019–2023)



10. Final Summary (Thesis Reporting)

This section prints a concise summary of the data scope and the performance of Models A, B (manual and tuned), and C. These values are used in the thesis discussion and comparative analysis.

In [151...

```
print("\n" + "=" * 60)
print(" GCC CO2 EMISSION MODELLING – THESIS RESULTS SUMMARY")
print("=" * 60 + "\n")

# 1. Data Scope
print("1- DATA SCOPE:")
print(f"    - Years Covered: {df_modern['year'].min()} to {df_modern['year'].max()}")
print(f"    - Countries: {df_modern['country'].nunique()} ({', '.join(df_modern['country'].unique)})")
print(f"    - Total Samples: {len(df_modern)}")
print(f"    - Hybrid Features Used: {features_hybrid}")
print("-" * 60)

# 2. Model Selection
print("\n2- MODEL SELECTION LEADERBOARD (Hybrid, TimeSeriesSplit):")
if "df_results_hybrid" in globals():
    print(df_results_hybrid.to_string(index=False))
else:
    print("    (Hybrid model selection step was not executed or not enough samples.)")
print("-" * 60)

# 3. Economic Baseline (Model A)
print("\n3- ECONOMIC BASELINE MODEL (Model A):")
if "r2_A" in globals():
    print(f"    Test R2 : {r2_A:.4f}")
    print(f"    Test RMSE: {rmse_A:.2f} MtCO2")
else:
    print("    (Model A not executed)")
print("-" * 60)

# 4. Hybrid Model (Model B, manual)
print("\n4- HYBRID MODEL (Model B, manual):")
if "r2_B" in globals():
    print(f"    Test R2 : {r2_B:.4f}")
    print(f"    Test RMSE: {rmse_B:.2f} MtCO2")
else:
    print("    (Model B (manual) not executed)")
print("-" * 60)

# 5. Tuned Hybrid Model (optional)
print("\n5- TUNED HYBRID MODEL (Model B, tuned):")
if "r2_B_tuned" in globals():
    print(f"    Test R2 : {r2_B_tuned:.4f}")
    print(f"    Test RMSE: {rmse_B_tuned:.2f} MtCO2")
else:
    print("    (Tuned Model B not executed)")
print("-" * 60)

# 6. Satellite-Only Stress Test (Model C)
print("\n6- SATELLITE-ONLY MODEL (Model C):")
if "r2_C" in globals():
    print(f"    Test R2 : {r2_C:.4f}")
    print(f"    Test RMSE: {rmse_C:.2f} MtCO2")
    if "r2_B" in globals():
        print(f"    Performance Drop vs Hybrid (R2): {r2_B - r2_C:.4f} points")
else:
    print("    (Model C not executed)")
print("-" * 60)
```

```
# 7. Feature Importance (Hybrid Model)
print("\n7- TOP DRIVERS OF EMISSIONS (Feature Importance - Hybrid Model B):")
if "imp_B" in globals():
    print(imp_B.to_string(index=False))
else:
    print("    (Feature importance not available)")
print("\n" + "=" * 60)
```

```
=====
GCC CO2 EMISSION MODELLING - THESIS RESULTS SUMMARY
=====
```

1- DATA SCOPE:

- Years Covered: 2000 to 2023
- Countries: 6 (Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, United Arab Emirates)
- Total Samples: 144
- Hybrid Features Used: ['log_pop_density', 'Night_Lights', 'LST_Summer_C', 'LST_Winter_C', 'NDVI_Spring', 'AOD_Annual', 'NO2_Signal']

2- MODEL SELECTION LEADERBOARD (Hybrid, TimeSeriesSplit):

Model	CV RMSE	CV R ²
XGBoost	101.748680	0.580005
Random Forest	113.807842	0.468054
Linear Regression	142.890518	0.346873
Ridge	164.903494	0.123506

3- ECONOMIC BASELINE MODEL (Model A):

Test R² : 0.9571
 Test RMSE: 46.29 MtCO₂

4- HYBRID MODEL (Model B, manual):

Test R² : 0.6916
 Test RMSE: 124.17 MtCO₂

5- TUNED HYBRID MODEL (Model B, tuned):

Test R² : 0.6231
 Test RMSE: 137.27 MtCO₂

6- SATELLITE-ONLY MODEL (Model C):

Test R² : 0.6668
 Test RMSE: 129.07 MtCO₂
 Performance Drop vs Hybrid (R²): 0.0248 points

7- TOP DRIVERS OF EMISSIONS (Feature Importance - Hybrid Model B):

Feature	Importance
LST_Summer_C	0.337243
LST_Winter_C	0.222876
Night_Lights	0.219002
log_pop_density	0.153086
NDVI_Spring	0.056655
AOD_Annual	0.011138
NO2_Signal	0.000000

11. Tabular Comparison of Model Performance

The following tables summarise the performance of the three main models:

- Model A: Economic baseline (macro-economic and demographic structure).
- Model B: Hybrid model (density + satellite indicators).
- Model B (tuned): TimeSeriesSplit + GridSearchCV variant (optional).
- Model C: Satellite-only stress test.

All metrics are computed on the same held-out test period (2019–2023).

In [152...

```
# Build a concise model comparison table based on the metrics already computed

rows = []

if "r2_A" in globals():
    rows.append({
        "Model": "Model A - Economic Baseline",
        "Category": "Economic / Historical",
        "Input Features": "GDP, population, energy, intensity, density",
        "Test R2": r2_A,
        "Test RMSE (MtCO2)": rmse_A,
    })

if "r2_B" in globals():
    rows.append({
        "Model": "Model B - Hybrid (manual)",
        "Category": "Hybrid (Economic + Satellite)",
        "Input Features": "log_pop_density + satellite indicators",
        "Test R2": r2_B,
        "Test RMSE (MtCO2)": rmse_B,
    })

if "r2_B_tuned" in globals():
    rows.append({
        "Model": "Model B - Hybrid (tuned)",
        "Category": "Hybrid (Economic + Satellite)",
        "Input Features": "log_pop_density + satellite indicators",
        "Test R2": r2_B_tuned,
        "Test RMSE (MtCO2)": rmse_B_tuned,
    })

if "r2_C" in globals():
    rows.append({
        "Model": "Model C - Satellite Only",
        "Category": "Satellite Only",
        "Input Features": "Night_Lights, LST, NDVI, AOD, NO2",
        "Test R2": r2_C,
        "Test RMSE (MtCO2)": rmse_C,
    })

df_model_comparison = pd.DataFrame(rows)

# Order rows in a logical sequence if all are present
order = [
    "Model A - Economic Baseline",
    "Model B - Hybrid (manual)",
    "Model B - Hybrid (tuned)",

```

```

    "Model C - Satellite Only",
]
df_model_comparison["Model"] = pd.Categorical(
    df_model_comparison["Model"], categories=order, ordered=True
)
df_model_comparison = df_model_comparison.sort_values("Model").reset_index(drop=True)

df_model_comparison

```

Out[152...

	Model	Category	Input Features	Test R2	Test RMSE (MtCO2)
0	Model A – Economic Baseline	Economic / Historical	GDP, population, energy, intensity, density	0.957135	46.292167
1	Model B – Hybrid (manual)	Hybrid (Economic + Satellite)	log_pop_density + satellite indicators	0.691605	124.167259
2	Model B – Hybrid (tuned)	Hybrid (Economic + Satellite)	log_pop_density + satellite indicators	0.623080	137.270742
3	Model C – Satellite Only	Satellite Only	Night_Lights, LST, NDVI, AOD, NO2	0.666792	129.065750

11.1 Relative Differences with Respect to the Economic Baseline

To better understand the trade-offs, the following table reports the relative difference of each model with respect to the economic baseline (Model A), in terms of:

- $\Delta R^2 = (\text{Model } R^2 - \text{Baseline } R^2)$
- $\Delta \text{RMSE} = (\text{Model RMSE} - \text{Baseline RMSE})$

In [153...

```

if "r2_A" in globals() and "rmse_A" in globals():
    comparison_rows = []
    for _, row in df_model_comparison.iterrows():
        model_name = row["Model"]
        if model_name == "Model A - Economic Baseline":
            # Baseline: deltas are zero by definition
            comparison_rows.append({
                "Model": model_name,
                "Category": row["Category"],
                "Test R2": row["Test R2"],
                "Test RMSE (MtCO2)": row["Test RMSE (MtCO2)"],
                "ΔR2 vs Model A": 0.0,
                "ΔRMSE vs Model A (MtCO2)": 0.0,
            })
        else:
            comparison_rows.append({
                "Model": model_name,
                "Category": row["Category"],
                "Test R2": row["Test R2"],
                "Test RMSE (MtCO2)": row["Test RMSE (MtCO2)"],
                "ΔR2 vs Model A": row["Test R2"] - r2_A,
                "ΔRMSE vs Model A (MtCO2)": row["Test RMSE (MtCO2)"] - rmse_A,
            })

df_relative = pd.DataFrame(comparison_rows)
df_relative = df_relative.sort_values("Model").reset_index(drop=True)
df_relative

```

```
else:
    print("Model A metrics are not available; relative comparison cannot be computed.")
```

In [154...

```
# Optional: export the final panel with density and satellite features
import os

os.makedirs("data", exist_ok=True)
export_path = "data/gcc_co2_panel_with_satellite_and_density_final.csv"
df_modern.to_csv(export_path, index=False)
print(f" Final panel saved to {export_path}")
```

Final panel saved to data/gcc_co2_panel_with_satellite_and_density_final.csv

11.2 Interpretation of the Comparison Tables

The **model comparison table** (`df_model_comparison`) summarises the out-of-sample performance of the three model families on the 2019–2023 test period. It allows a direct side-by-side view of the trade-off between accuracy (R^2 , RMSE) and the type of information used (economic vs. hybrid vs. satellite-only).

The **relative difference table** (`df_relative`) reports all metrics normalised with respect to the economic baseline (Model A). It highlights how much performance is lost when moving from the baseline to the hybrid model, and how much further degradation occurs when using satellite indicators only. This representation is useful when discussing the cost, in terms of numerical accuracy, of gaining real-time capability and independence from delayed inventories.

Both tables are derived from the same temporal split (training: 2000–2018, testing: 2019–2023) and therefore can be interpreted consistently across the different model families.

12. Experiment Narrative and Benchmarking Summary

This notebook represents the final, consolidated modelling experiment for the thesis. It brings together the most robust design choices identified after several iterations of data preparation, feature engineering and model evaluation conducted in earlier notebooks.

12.1 From Single-Country Baseline to GCC Panel

Initial experiments focused on Saudi Arabia alone and relied on annual macro-economic indicators (GDP, population, energy use) compiled from “Our World in Data”, the World Bank and related sources. With only about thirty yearly observations, these single-country models were highly vulnerable to overfitting, particularly when using flexible machine-learning algorithms such as ensemble trees and neural networks. In this early phase, several configurations achieved very high in-sample R^2 , but their generalisation to structurally different future periods (for example, post-COVID years) could not be assessed robustly.

To increase the effective sample size and make more reliable use of AI methods, the scope was extended to a **GCC-wide panel**, covering six countries over the period 2000–2023. This yielded 144 country–year samples and allowed the use of ensemble tree models while still maintaining a relatively homogeneous regional context (arid climate, hydrocarbon-based economies, rapid urbanisation). The panel design also made it possible to study how the same set of predictors behaves across countries with different population sizes and spatial extents (for example, compact Bahrain versus geographically extensive Saudi Arabia).

In parallel, earlier notebooks explored a longer historical window, extending the economic series back to the mid-1960s for Saudi Arabia and other GCC countries. These long-horizon experiments confirmed the strong dependence of national CO₂ emissions on energy use and population, and showed that classical time-series models and tree-based methods can reproduce historical inventories with high accuracy. However, they also revealed a structural incompatibility between pre-2000 economic records and modern Earth Observation (EO) products: most satellite missions simply do not exist before the early 2000s, which limits the period over which EO-based proxies can be meaningfully integrated.

12.2 Evolution of the Satellite Feature Design

The first generation of satellite features was constructed as simple national annual means (for example, annual mean LST and annual mean NDVI) and merged with the long historical economic panel (1960s onwards). While some of these models achieved superficially high R² scores, feature-importance analysis showed that almost all explanatory power came from population and other structural economic variables, with very limited contribution from EO indicators. Annual averaging had effectively removed the seasonal extremes that drive cooling demand in the GCC, and the long pre-launch periods for many missions resulted in extended stretches of missing or zero satellite values.

These observations motivated two key design decisions that underpin the present notebook:

- **Temporal restriction to the modern satellite era (2000–2023)** so that satellite and economic indicators are temporally consistent and genuinely observed rather than backfilled.
- **Physically informed, seasonally resolved satellite features** instead of simple annual means.

The final satellite feature set therefore includes:

- **Seasonal thermal indicators:** `LST_Summer_C` (June–August) and `LST_Winter_C` (December–February) to capture peak and base thermal loads, which are directly linked to cooling and heating demand.
- **Vegetation context:** `NDVI_Spring` to characterise land cover during the main growing season, reflecting the limited greening period in arid environments.
- **Anthropogenic activity proxies:** `Night_Lights` as a proxy for urbanisation and electricity use, capturing the spatial footprint and intensity of human activity.
- **Atmospheric loading:** `AOD_Annual` and an aggregated `NO2_Signal` as indicators of combustion-related particulates and trace gases.

In addition, the demographic signal is represented by **log-transformed population density**, which distinguishes compact from diffuse settlement patterns and reduces the dominance of raw population totals. This reparameterisation is particularly important in a panel that includes both very small, dense states and large, low-density countries.

12.3 Summary of Final Model Performance

The final experimental design compares three model families on a strictly held-out test window (2019–2023): an economic baseline model (Model A), a hybrid model combining demographic density with satellite indicators (Model B), and a satellite-only model (Model C). The key findings can be summarised as follows:

- **Model A – Economic Baseline (GDP, population, energy, intensity, density)**

The purely economic model achieves the highest numerical accuracy on the test period, with a test coefficient of determination of approximately $R^2 \approx 0.96$ and a root-mean-square error of about **46**

MtCO₂. This confirms that traditional inventory-based indicators remain the strongest predictors of national emissions when they are available. At the same time, Model A is estimated from a relatively small number of training samples (114 country–years with seven explanatory variables), so this very high R^2 should be interpreted as an **upper-bound reference** rather than as a guaranteed level of out-of-sample performance in all settings. In other words, Model A provides a strong benchmark against which activity-based alternatives can be compared, but it is itself constrained by data latency and limited sample size.

- **Model B – Hybrid Model (log population density + satellite indicators)**

The manually tuned hybrid configuration reaches a test performance of approximately $R^2 \approx 0.69$ with **RMSE ≈ 124 MtCO₂**, while the grid-searched variant converges to a slightly more conservative model with $R^2 \approx 0.62$ and **RMSE ≈ 137 MtCO₂**. In both cases, the hybrid models are less accurate than the economic baseline on the 2019–2023 window, but they have the advantage that they can, in principle, be updated in near real time as Earth-observation and population data become available. The use of a strict temporal split (training on 2000–2018 and testing on 2019–2023) means that these R^2 values represent performance on genuinely unseen future years, including the COVID-19 period, rather than optimistic interpolation of already-known years.

- **Model C – Satellite-Only Model (lights, LST, NDVI, AOD, NO₂)**

When the demographic density variable is removed and only satellite-based indicators are retained, the satellite-only model still explains approximately $R^2 \approx 0.67$ of the variance on the test period, with **RMSE ≈ 129 MtCO₂**. These results indicate that physically meaningful satellite features carry non-trivial information about national-scale emissions, even without explicit census-based variables. In other words, a purely EO-based configuration can recover a substantial fraction of the emission signal, but it does not yet match the accuracy of models that have access to detailed economic and demographic inputs.

Feature-importance analysis in the hybrid model indicates that **summer and winter land surface temperature and night-time lights are the leading predictors**, with log population density acting as a secondary but non-negligible driver. NDVI and AOD have relatively low importance in this configuration. The **NO₂ signal does not contribute to the hybrid feature importance in the long 2000–2018 training window**, because Sentinel-5P is only available in the late 2010s; for this reason, NO₂ is examined separately in a short-window diagnostic experiment (2019–2023) rather than being treated as a core driver in the main panel model. In that diagnostic setting, where both training and evaluation are restricted to 2019–2023, NO₂ ranks among the top predictors, but the resulting $R^2 \approx 1.0$ is interpreted purely as an in-sample indicator of feature relevance rather than as an estimate of generalisation performance.

Taken together, these results highlight a clear **accuracy–timeliness trade-off**: the economic model provides the most accurate but delayed estimates, while the hybrid and satellite-only models provide physically interpretable, near-real-time nowcasts with moderate but operationally useful accuracy.

12.4 Interpretation and Limitations

Taken together, the results illustrate a clear **accuracy–timeliness trade-off**:

- Economic models (Model A) provide a strong upper bound on achievable accuracy for the available panel but inherit the reporting delays and potential revisions of national inventories. Their very high R^2 should therefore be seen as a best-case reference conditional on having access to fully compiled historical statistics.

- Hybrid models (Model B) sacrifice some accuracy in exchange for the ability to generate **nowcasts** based primarily on physical activity proxies and atmospheric measurements, which can be updated as soon as new satellite and population layers are released.
- Satellite-only models (Model C) confirm that Earth Observation can detect a meaningful part of the emission signal even without socio-economic inputs, but they are not yet sufficient on their own for precise national accounting and should be interpreted as a complementary line of evidence.

The analysis in this notebook is restricted to annual, country-level aggregates. It does not resolve sub-national patterns or intra-annual variability, and it does not yet exploit monthly or higher-frequency emission timelines. These aspects are natural directions for future work, particularly in combination with sequence models that are designed for finer temporal structure and with spatial disaggregation of emissions into urban, industrial and rural components.

In addition, the experiments documented here focus on a single regional context (the GCC) with relatively homogeneous climatic and economic conditions. While this improves comparability across countries, it also means that the models have not been tested for out-of-region transfer to very different climates or economic structures. Finally, the satellite indicators themselves are subject to retrieval uncertainties, cloud and dust contamination, and mission-specific artefacts, which may limit their stability over very long periods.

Despite these limitations, the final hybrid framework implemented here provides a transparent and reproducible proof-of-concept that satellite indicators, when carefully engineered and evaluated under a realistic temporal split, can act as a complementary, independent line of evidence for monitoring CO₂ emission dynamics in the GCC region. Earlier notebooks based on longer historical series and Saudi-only configurations support the robustness of the main qualitative conclusions—namely, the central role of energy use, population density, thermal stress and urbanisation intensity—but the present notebook provides the most coherent and operationally relevant synthesis for a satellite-enabled nowcasting framework at the national scale.

7. Sentinel-5P NO₂ Diagnostic Experiment

This section documents how tropospheric NO₂ is handled in the final set of experiments.

- In the long-term GCC panel (2000–2023), the training window for the hybrid model is 2000–2018 and the test window is 2019–2023. Sentinel-5P NO₂ is only available from the late 2010s onwards. As a result, the `N02_Signal` column is largely missing or non-informative during most of the training period. When used directly in the panel model, it carries essentially zero importance and cannot contribute meaningfully to the main hybrid feature set.
- To justify the discussion of NO₂ in the thesis, we therefore keep a separate, short-term diagnostic experiment on the modern window where NO₂ is actually observed. This experiment does **not** replace the main 2000–2018 / 2019–2023 panel results; instead, it is used to show that, when available, Sentinel-5P NO₂ behaves as a physically sensible combustion indicator and typically ranks among the most influential features.

The following cells:

1. Quantify the temporal availability of `N02_Signal` across the panel.
2. Construct a modern-window subset (years with reasonable NO₂ coverage).

3. Fit a compact hybrid model on this subset and report feature importance, with a focus on the relative role of NO₂ compared to other satellite indicators.

From a methodological perspective, it is important to emphasise that this NO₂ experiment is **diagnostic rather than operational**. In the full 2000–2018 training window used for the main panel models, Sentinel-5P NO₂ is not available and therefore has effectively zero importance in the hybrid feature-importance rankings. The strong role of NO₂ observed here is restricted to the short modern period where Sentinel-5P data exist and should be interpreted as evidence that NO₂ is a physically meaningful combustion proxy in the GCC, rather than as a driver in the long-term training phase of the main models. Because the model is trained and evaluated on the same short 2019–2023 window, the $R^2 \approx 1.0$ should be interpreted purely as an in-sample diagnostic for feature ranking, not as an estimate of generalisation performance.

In [155...

```
# 7.1 Temporal availability of NO2_Signal

if "df_modern" in globals() and "NO2_Signal" in df_modern.columns:
    no2_missing = (
        df_modern
        .groupby("year")["NO2_Signal"]
        .apply(lambda s: s.isna().mean())
        .reset_index(name="missing_fraction")
        .sort_values("year")
    )

    print("Fraction of missing NO2_Signal by year (across all GCC countries):")
    print(no2_missing.to_string(index=False))

    modern_no2 = no2_missing[no2_missing["missing_fraction"] < 0.5]
    print("\nYears with NO2_Signal available for at least half of the panel rows:")
    if len(modern_no2) == 0:
        print("    None - NO2 coverage is too sparse for a panel-wide analysis.")
    else:
        print(modern_no2.to_string(index=False))
else:
    print("df_modern or the 'NO2_Signal' column is not available in the current session.")
```

Fraction of missing NO2_Signal by year (across all GCC countries):

year	missing_fraction
2000	1.0
2001	1.0
2002	1.0
2003	1.0
2004	1.0
2005	1.0
2006	1.0
2007	1.0
2008	1.0
2009	1.0
2010	1.0
2011	1.0
2012	1.0
2013	1.0
2014	1.0
2015	1.0
2016	1.0
2017	1.0
2018	1.0
2019	0.0
2020	0.0
2021	0.0
2022	0.0
2023	0.0

Years with NO2_Signal available for at least half of the panel rows:

year	missing_fraction
2019	0.0
2020	0.0
2021	0.0
2022	0.0
2023	0.0

In [156...

```
# 7.2 Short-term hybrid experiment including NO2 (diagnostic only)
```

```
from xgboost import XGBRegressor
from sklearn.metrics import r2_score, mean_squared_error
import numpy as np
import pandas as pd

if "df_modern" in globals() and "NO2_Signal" in df_modern.columns:
    # Primary modern window: years >= 2019 with non-missing NO2
    mask_modern = (df_modern["NO2_Signal"].notna()) & (df_modern["year"] >= 2019)
    df_no2_modern = df_modern[mask_modern].copy()

    # If this is too small, relax to >= 2015 while keeping non-missing NO2
    if len(df_no2_modern) < 24: # heuristic: 6 countries * 4 years
        mask_relaxed = (df_modern["NO2_Signal"].notna()) & (df_modern["year"] >= 2015)
        df_no2_modern = df_modern[mask_relaxed].copy()

    print(f"Modern NO2 subset rows: {len(df_no2_modern)}")
    if len(df_no2_modern) == 0:
        print("No rows with usable NO2_Signal were found in the modern window.")
    else:
        print("Row count by year in the modern NO2 subset:")
        print(df_no2_modern.groupby("year").size())

    # Compact hybrid feature set for the modern window
    candidate_features = [
        "NO2_Signal",
```

```

        "LST_Summer_C",
        "LST_Winter_C",
        "Night_Lights",
        "log_pop_density",
        "NDVI_Spring",
        "AOD_Annual",
    ]
    modern_features = [f for f in candidate_features if f in df_no2_modern.columns]

    if "target" in globals() and target in df_no2_modern.columns:
        y_col = target
    else:
        y_col = "co2"

    X_no2 = df_no2_modern[modern_features]
    y_no2 = df_no2_modern[y_col]

    # XGBoost model for diagnostic importance (in-sample fit)
    model_no2 = XGBRegressor(
        n_estimators=500,
        learning_rate=0.05,
        max_depth=4,
        subsample=0.8,
        colsample_bytree=0.8,
        random_state=42,
        n_jobs=-1,
    )
    model_no2.fit(X_no2, y_no2)
    preds_no2 = model_no2.predict(X_no2)
    r2_no2 = r2_score(y_no2, preds_no2)
    rmse_no2 = np.sqrt(mean_squared_error(y_no2, preds_no2))

    print("\nModern NO2 experiment - in-sample performance (diagnostic only):")
    print(f"    R²    : {r2_no2:.4f}")
    print(f"    RMSE  : {rmse_no2:.2f} MtCO2")

    # Feature importance within this modern-window model
    imp_no2 = (
        pd.DataFrame({"Feature": modern_features,
                      "Importance": model_no2.feature_importances_})
        .sort_values("Importance", ascending=False)
    )

    print("\nFeature importance in the modern NO2 experiment:")
    print(imp_no2.to_string(index=False))
else:
    print("df_modern or the 'NO2_Signal' column is not available in the current session.")

```

Modern NO2 subset rows: 30

Row count by year in the modern NO2 subset:

```
year
2019    6
2020    6
2021    6
2022    6
2023    6
dtype: int64
```

Modern NO2 experiment - in-sample performance (diagnostic only):

```
R²      : 1.0000
RMSE    : 0.01 MtCO2
```

Feature importance in the modern NO2 experiment:

Feature	Importance
log_pop_density	0.405130
NO2_Signal	0.203474
LST_Summer_C	0.140514
Night_Lights	0.123310
LST_Winter_C	0.084646
NDVI_Spring	0.036521
AOD_Annual	0.006405

Because this diagnostic model is trained and evaluated on the same short 2019–2023 window, the very high in-sample R^2 (≈ 1.0) must be interpreted purely as an internal diagnostic for feature ranking, not as a measure of out-of-sample generalisation performance. The purpose of this experiment is to confirm that NO_2 becomes a dominant predictor once it is available, rather than to propose a deployable forecasting model based only on this short period.

In [157...

```
# === FINAL EXPERIMENT SUMMARY (single-cell report) ===

print("\n" + "="*70)
print(" GCC CO2 EMISSIONS - FINAL HYBRID SATELLITE EXPERIMENT SUMMARY")
print("="*70 + "\n")

# 1. Data scope
if 'df_modern' in globals():
    years_min = df_modern['year'].min()
    years_max = df_modern['year'].max()
    countries = sorted(df_modern['country'].unique()) if 'country' in df_modern.columns else []
    print("1) DATA SCOPE")
    print(f"    - Years covered          : {years_min}-{years_max}")
    if countries:
        print(f"    - Countries              : {len(countries)} ({', '.join(countries)})")
    print(f"    - Total samples (rows) : {len(df_modern)}")
else:
    print("1) DATA SCOPE")
    print("    df_modern is not defined.")

# 2. Feature sets
print("\n2) FEATURE SETS")
if 'features_economic' in globals():
    print(f"    - Economic features (Model A): {features_economic}")
if 'features_hybrid' in globals():
    print(f"    - Hybrid features (Model B)  : {features_hybrid}")
if 'features_satellite_only' in globals():
    print(f"    - Satellite-only features (Model C): {features_satellite_only}")

# 3. Model performance summary
```

```

print("\n3) MODEL PERFORMANCE (Test period, e.g. 2019-2023)")
if 'df_model_comparison' in globals():
    print(df_model_comparison.to_string(index=False))
else:
    # Fallback: print individual metrics if comparison table is missing
    if 'r2_A' in globals():
        print(f"    - Model A - Economic Baseline    : R²={r2_A:.4f}, RMSE={rmse_A:.2f} MtCO2")
    if 'r2_B' in globals():
        print(f"    - Model B - Hybrid (manual)          : R²={r2_B:.4f}, RMSE={rmse_B:.2f} MtCO2")
    if 'r2_B_tuned' in globals():
        print(f"    - Model B - Hybrid (tuned)             : R²={r2_B_tuned:.4f}, RMSE={rmse_B_tuned:.2f}")
    if 'r2_C' in globals():
        print(f"    - Model C - Satellite Only             : R²={r2_C:.4f}, RMSE={rmse_C:.2f} MtCO2")

# 4. Feature impact summary (correlation + importance for hybrid model)
print("\n4) FEATURE IMPACT (Hybrid model - training period)")

summary_rows = []
if 'train_df' in globals() and 'target' in globals():
    # Determine union of economic + hybrid + satellite features, if available
    all_features = set()
    if 'features_economic' in globals():
        all_features |= set(features_economic)
    if 'features_hybrid' in globals():
        all_features |= set(features_hybrid)
    if 'features_satellite_only' in globals():
        all_features |= set(features_satellite_only)

    all_features = sorted(all_features)

    # Build mapping for feature groups
    def feature_group(f):
        if 'features_economic' in globals() and f in features_economic:
            return "Economic"
        if 'features_satellite_only' in globals() and f in features_satellite_only:
            return "Satellite"
        if f in ['pop_density', 'log_pop_density']:
            return "Demographic"
        return "Other"

    # If hybrid importance is available (imp_B from Model B), create a lookup
    importance_lookup = {}
    if 'imp_B' in globals():
        for _, row in imp_B.iterrows():
            importance_lookup[row['Feature']] = row['Importance']

    for f in all_features:
        corr = np.nan
        if f in train_df.columns and target in train_df.columns:
            # Pearson correlation with CO2 on training data, if defined and not constant
            try:
                sub = train_df[[f, target]].dropna()
                if sub[f].nunique() > 1:
                    corr = sub[f].corr(sub[target])
            except Exception:
                corr = np.nan

        imp_val = importance_lookup.get(f, np.nan)
        summary_rows.append(
            {
                "Feature": f,
                "Group": feature_group(f),

```

```

        "Train_Corr_with_CO2": corr,
        "Hybrid_Importance": imp_val,
    }
)

if summary_rows:
    df_feature_summary = pd.DataFrame(summary_rows)
    # Order by importance (descending) then by absolute correlation
    df_feature_summary = df_feature_summary.sort_values(
        by=["Hybrid_Importance", "Train_Corr_with_CO2"],
        ascending=[False, False],
    )
    print(df_feature_summary.to_string(index=False))
else:
    print("    No features found for summary.")
else:
    print("    train_df or target not available for feature impact analysis.")

print("\n" + "="*70)
print(" End of final experiment summary cell ")
print("="*70 + "\n")

```

1) DATA SCOPE

- Years covered : 2000–2023
- Countries : 6 (Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, United Arab Emirates)
- Total samples (rows) : 144

2) FEATURE SETS

- Economic features (Model A): ['gdp', 'population', 'primary_energy_consumption', 'energy_per_capita', 'co2_per_gdp', 'pop_density', 'log_pop_density']
- Hybrid features (Model B) : ['log_pop_density', 'Night_Lights', 'LST_Summer_C', 'LST_Winter_C', 'NDVI_Spring', 'AOD_Annual', 'NO2_Signal']
- Satellite-only features (Model C): ['Night_Lights', 'LST_Summer_C', 'LST_Winter_C', 'NDVI_Spring', 'AOD_Annual', 'NO2_Signal']

3) MODEL PERFORMANCE (Test period, e.g. 2019–2023)

	Model	Category	Input Features
Test R2	Test RMSE (MtCO2)		
Model A – Economic Baseline	Economic / Historical	GDP, population, energy, intensity, density	0.957135 46.292167
Model B – Hybrid (manual)	Hybrid (Economic + Satellite)	log_pop_density + satellite indicators	0.691605 124.167259
Model B – Hybrid (tuned)	Hybrid (Economic + Satellite)	log_pop_density + satellite indicators	0.623080 137.270742
Model C – Satellite Only	Satellite Only	Night_Lights, LST, NDVI, AOD, NO2	0.666792 129.065750

4) FEATURE IMPACT (Hybrid model – training period)

Feature	Group	Train_Corr_with_CO2	Hybrid_Importance
LST_Summer_C	Satellite	0.375027	0.337243
LST_Winter_C	Satellite	0.110417	0.222876
Night_Lights	Satellite	-0.460281	0.219002
log_pop_density	Economic	-0.486053	0.153086
NDVI_Spring	Satellite	0.407567	0.056655
AOD_Annual	Satellite	-0.264095	0.011138
NO2_Signal	Satellite	NaN	0.000000
primary_energy_consumption	Economic	0.993013	NaN
population	Economic	0.991582	NaN
gdp	Economic	0.971883	NaN
co2_per_gdp	Economic	-0.258600	NaN
pop_density	Economic	-0.347806	NaN
energy_per_capita	Economic	-0.381655	NaN

End of final experiment summary cell

8. Comparative Benchmarking Against Previous Studies

This section provides a consolidated interpretation of the final experimental results of this notebook in light of the existing literature on CO₂ emissions modelling for Saudi Arabia and the GCC. It mirrors the discussion that will be expanded in Chapter 5 (Discussion) of the thesis.

The main models investigated in this notebook are:

- Model A – Economic Baseline:**
 Model A uses purely economic and energy indicators (GDP, population, primary energy consumption, energy per capita, CO₂ intensity, and population density) to reproduce official national CO₂ inventories. This model serves as a strong reference benchmark based on fully reported historical data (albeit released with a 1–2 year reporting delay in practice), against which the hybrid and satellite-based alternatives can be compared.
 Test performance (2019–2023):
 - $R^2 \approx 0.96$
 - RMSE $\approx 46 \text{ MtCO}_2$
- Model B – Hybrid (Satellite + Density):**
 Combines log-population density with satellite-derived indicators (Night Lights, seasonal LST, NDVI, AOD, NO₂ where available).
 Test performance (2019–2023):
 - Manual configuration: $R^2 \approx 0.69$, RMSE $\approx 124 \text{ MtCO}_2$
 - Tuned configuration: $R^2 \approx 0.62$, RMSE $\approx 137 \text{ MtCO}_2$
- Model C – Satellite Only:**
 Model C relies exclusively on satellite-derived features (night-time lights, seasonal land surface temperature, NDVI, AOD, and NO₂ where available), without explicit census-based variables. Although its test performance is lower than the economic baseline and slightly below the hybrid configuration, these results indicate that physically meaningful satellite features carry non-trivial information about national-scale emissions, even in the absence of demographic or economic inputs.
 Test performance (2019–2023):
 - $R^2 \approx 0.67$, RMSE $\approx 129 \text{ MtCO}_2$

These results are now benchmarked against three main families of studies reviewed in the literature.

8.1 Comparison with Economic and Historical Models

Several studies in the literature rely purely on historical macro–economic indicators (GDP, energy consumption, population) and time–series techniques. They typically achieve high accuracy but are constrained by **data latency** (1–2 year reporting delays in national inventories).

Study	Methodology	Reported Accuracy	Comparison with Thesis Results
Nassef et al. (2023)	FFNN, LSTM on historical series (1954–2020)	$R^2 \approx 0.93$	Our Model A (economic baseline) achieves slightly higher accuracy ($R^2 \approx 0.96$) on the GCC panel, confirming that tree–based ensembles can match or outperform neural networks on small tabular datasets.
Rahman et al. (2023)	Decision Trees on socio–economic indicators	$R^2 \approx 0.90$	Consistent with our findings: tree–based methods are very effective when all economic indicators are available.
Alamri & Khan (2023)	N-BEATS time–series model	$R^2 \approx 0.70$	Accuracy is comparable to our Hybrid Model B ($R^2 \approx 0.69$ – 0.62), but their approach forecasts from past historical trends , whereas our hybrid model estimates emissions from current physical activity (satellites), which is more suitable for nowcasting.

Study	Methodology	Reported Accuracy	Comparison with Thesis Results
Meng & Noman (2022)	SARIMA and related statistical models	Low MAPE	Their approach is purely statistical. In contrast, our models incorporate physically interpretable drivers (heat, lights, pollution), improving explanatory power for policymakers.

Summary:

- The **Economic Baseline (Model A)** reproduces and slightly improves the performance of state-of-the-art economic models.
- However, these models inherit the same weakness as the underlying inventories: **they cannot provide real-time estimates**, because they depend on delayed national statistics.
- This motivates the development of the hybrid and satellite-only frameworks.

8.2 Comparison with Satellite-Based and EO-Driven Models

A second family of studies integrates satellite Earth Observation (EO) products (e.g., OCO-2, Sentinel-5P, Landsat, Sentinel-2) with emissions estimates.

Study	Satellite Inputs	Performance	Comparative Insight
Alcibahy et al. (2025)	OCO-2 (XCO ₂), Sentinel-5P, ERA5	$R^2 \approx 0.98$	They use atmospheric concentration (XCO ₂), which is strongly influenced by long-range transport and wind fields. Our framework deliberately excludes XCO ₂ to avoid mixing local emissions with imported signals, accepting lower accuracy ($R^2 \approx 0.69$) in exchange for a stricter, activity-based interpretation.
Couture et al. (2024)	Landsat-8, Sentinel-2	RMSE ≈ 2.18 Mt (point sources)	Focuses on individual power plants. Our models operate at national scale , capturing both point sources and diffuse urban emissions.
Alkayal et al. (2025)	Sentinel-5P NO ₂	High local accuracy	Their findings support the use of NO ₂ as a key combustion indicator. In our short-term Sentinel-5P diagnostic (2019–2023), NO ₂ ranks among the top predictors, confirming its physical relevance whenever coverage is available.
Rowley & Karakuş (2023)	Multi-modal EO fusion	$R^2 \approx 0.92$	Their architecture validates the idea of combining multiple EO signals (lights, heat, chemistry). We adapt this logic to a data-scarce, long-term GCC panel (2000–2023) rather than a short modern period only.

Summary:

- Some EO studies report **higher numerical accuracy** by using atmospheric concentration products (e.g., XCO₂).
- The present thesis follows a different philosophy: it emphasizes **local activity signals** (night-time lights, surface temperature, NO₂, AOD) and avoids top-down concentration fields unless inverse transport modelling is available.
- With this stricter design, the hybrid model still explains roughly **two-thirds of the variance** ($R^2 \approx 0.69$), which is a strong result for an activity-based framework.

8.3 Relation to Land-Use, Urban, and Vegetation Studies

A third group of studies focuses on specific land–surface processes such as vegetation greenness or urban heat islands.

Study	Focus Area	Relation to Thesis Findings
Haq (2022)	Vegetation (NDVI)	Shows that NDVI is effective for land-cover classification and monitoring. In our national-scale emission models, NDVI has very low predictive power for CO ₂ (lowest feature importance), suggesting that greening alone is a weak proxy for national emission trends in the GCC.
Addas (2023)	Urban Heat Islands (LST)	Demonstrates the importance of surface temperature in urban environments. Our results extend this concept to the national level , where Summer LST emerges as a key driver of electricity-related emissions in the GCC hybrid model.

These comparisons reinforce a central conclusion of the thesis:
seasonal heat stress and urbanization intensity are crucial physical drivers of energy-related CO₂ emissions in arid regions.

8.4 Positioning of the Thesis within the Literature

Based on the experiments in this notebook and the broader literature review, the contribution of this work can be summarized as follows:

- Bridging Economic and EO Paradigms**
 - The thesis explicitly benchmarks a **purely economic model (Model A)** against a **hybrid activity-based model (Model B)** and a **satellite-only model (Model C)** within the same GCC panel.
 - This reveals the trade-off between accuracy and timeliness:
 - Economic models: highest accuracy, but delayed.
 - Hybrid models: moderate accuracy, but near real-time capability.
 - Satellite-only models: slightly lower accuracy but fully independent of national statistics.
- Physical Interpretability and Feature Hierarchy**
 - The hybrid model’s feature importance analysis consistently highlights:
 - log-population density** as a structural driver,
 - summer LST** and **night-time lights** as key activity proxies,
 - NDVI** and **AOD** as secondary factors.
 - In the short-term NO₂ diagnostic experiment (2019–2023), **NO₂** ranks among the top predictors, confirming its role as a combustion tracer when coverage is available.
- Methodological Transparency for Nowcasting**
 - Unlike many prior works that rely on random splits or standard cross-validation, the main panel experiments in this notebook adopt a **strict temporal cut-off** (train on 2000–2018, test on 2019–2023) to simulate real deployment.
 - The achieved hybrid performance ($R^2 \approx 0.69$ on truly unseen future years) therefore represents a conservative and realistic estimate of nowcasting capability, rather than an optimistic interpolation of known years.

Taken together, these elements show that the proposed framework does not only reproduce existing results, but also **repositions the problem**: from delayed, purely economic forecasting towards **physically grounded, satellite-enabled nowcasting** of national CO₂ emissions in the GCC.

