# Decomposition and Forecasting of the Mauna Loa Atmospheric $CO_2$ Record, 1990–2023

A Technical Report on Seasonal Structure, Growth-Rate Acceleration, and Out-of-Sample SARIMA Validation

July 11, 2026

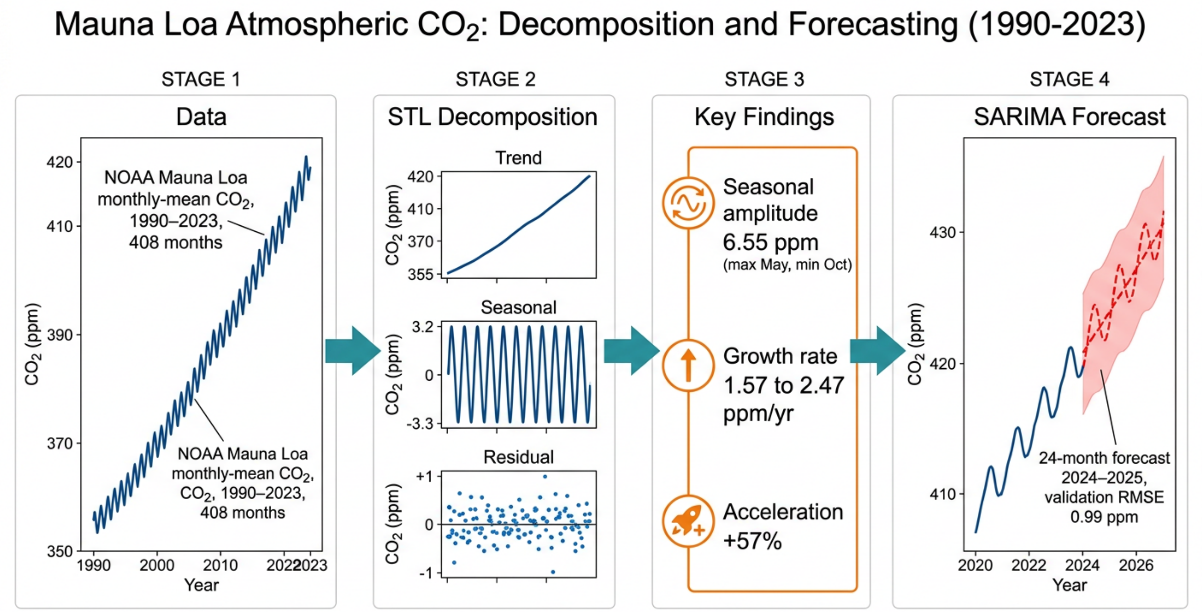


**Figure 1. Graphical abstract.** Analysis workflow for the NOAA Global Monitoring Laboratory (GML) Mauna Loa monthly-mean atmospheric $CO_2$ record. The 408-month analysis window (January 1990–December 2023) is decomposed with Seasonal-Trend decomposition using Loess (STL) into trend, seasonal, and residual components; the seasonal cycle, decadal growth rates, and their acceleration are characterized; and a Seasonal ARIMA model is fit and used to generate a 24-month-ahead forecast (January 2024–December 2025) that is validated against the realized observations.

## Executive Summary

This report documents a fully reproducible analysis of the longest continuous, in-situ record of monthly-mean atmospheric carbon-dioxide ($CO_2$) mole fraction: the Mauna Loa Observatory (MLO) record maintained by the NOAA Global Monitoring Laboratory (GML). The analysis is restricted to the 408 contiguous monthly observations spanning January 1990 through December 2023, drawn from the `average` (complete monthly-mean) column of the file `co2_mm_mlo.csv` (NOAA GML, file creation date *Sun Jul 5 2026*). The series is decomposed with additive Seasonal-Trend decomposition using Loess (STL), and forecast with a grid-searched Seasonal ARIMA (SARIMA) model. The principal findings are as follows.

- **Seasonal cycle.** The mean peak-to-trough seasonal amplitude over 1990–2023 is **6.55 ppm** (standard deviation 0.23 *ppm*, $n = 34$ years). The seasonal maximum occurs in **May** in every year of record; the seasonal minimum occurs most often in **October** (20 of 34 years), with September the runner-up (14 of 34 years).
- **Growth-rate acceleration.** The mean annual growth rate rose from $\mathbf{1.57\,ppm\,yr^{-1}}$ in 1990–1999 to $\mathbf{2.47\,ppm\,yr^{-1}}$ in 2014–2023, an increase of $\approx 0.90\,ppm\,yr^{-1}$ (+57%), robust across three independent estimation methods.
- **Forecast and validation.** An out-of-sample 24-month forecast for January 2024–December 2025 from a $\mathbf{SARIMA(1,1,1)\times(0,1,1)_{12}}$ model tracks the realized observations closely, achieving a validation **RMSE of** 0.988 *ppm* (MAE 0.933 *ppm*; MAPE 0.22%), with 22 of 24 (91.7%) observations falling inside the 95% prediction intervals.

All code, the fitted-components table, the seasonal-amplitude and growth-rate tables, the full SARIMA grid-search table, and the forecast/validation table are provided as machine-readable CSV files (Section 4.3).

# 1 Introduction and Data Source

## 1.1 Scientific context

The systematic, high-precision monitoring of atmospheric $CO_2$ began in 1958 when Charles David Keeling installed a continuous infrared analyzer at Mauna Loa Observatory, Hawaii [Keeling, 1960, Keeling et al., 1976]. The resulting "Keeling Curve" remains the single most iconic dataset in climate science: it provided the first unambiguous, quantitative evidence that atmospheric $CO_2$ is rising in step with fossil-fuel combustion, and it revealed a pronounced annual cycle—the "breathing of the biosphere"—superimposed on the secular trend [Keeling et al., 1976]. Mauna Loa's location, high on a mid-Pacific volcano far from major local sources, samples well-mixed background free-tropospheric air, which is why the site has been the global reference for atmospheric $CO_2$ for more than six decades [Thoning et al., 1989].

Two structural features of this record are of enduring scientific interest and are the quantitative focus of the present report. First, the amplitude of the seasonal cycle encodes the strength of the Northern Hemisphere terrestrial biosphere's seasonal net ecosystem exchange—summer photosynthetic drawdown followed by autumn and winter respiratory release [Keeling et al., 1996, Graven et al., 2013]. Second, the long-term growth rate encodes the balance between anthropogenic emissions and the combined uptake by the land and ocean carbon sinks; its year-to-year variability is dominated by climate modes such as the El Niño–Southern Oscillation (ENSO)—a coupling first identified in the earliest Mauna Loa analyses [Bacastow and Keeling, 1973]—while its multi-decadal acceleration reflects the growth of global emissions [Le Quéré et al., 2009, Betts et al., 2016, Friedlingstein et al., 2023, IPCC, 2021]. Robust extraction of the slowly-varying growth rate from the noisy monthly signal is itself an active methodological problem [Wu et al., 2020]. Robust statistical decomposition and forecasting of the record therefore serve both as a methodological benchmark and as a diagnostic of carbon-cycle behavior.

## 1.2 Data source, version, and provenance

The data used here are the NOAA GML Mauna Loa monthly-mean $CO_2$ product, distributed as `co2_mm_mlo.csv` from `https://gml.noaa.gov/ccgg/trends/` [Lan et al., 2024, Dlugokencky and Tans, 2024]. This is the longest continuous *in-situ* monthly-mean $CO_2$ record available for any station worldwide, extending from March 1958 to the present. The specific file retrieved for this study contains 820 monthly records spanning March 1958–June 2026 and carries the embedded header stamp **"File Creation: Sun Jul 5 03:55:25 2026"**, which serves as the effective version identifier (the NOAA product is not assigned an explicit release version number; the file-creation timestamp and the contact of record, Dr. Xin Lan, uniquely identify the vintage). The corresponding formal data citation is the NOAA GML globally-averaged $CO_2$ dataset [Lan et al., 2024].

The current NOAA file format provides two monthly time series of interest: `average`, the complete gap-filled monthly-mean mole fraction, and `deseasonalized`, a long-term trend from which the seasonal cycle has already been removed. Because the objective of this report is to *recover* the seasonal and trend structure from the raw signal, the `average` column is used throughout as the primary series; the `deseasonalized` column is deliberately *not* used as an input, as it would presuppose the very decomposition being performed. (The legacy `interpolated` column present in older NOAA formats is no longer distributed; `average` is the standard gap-free monthly-mean series.) All $CO_2$ values are reported as dry-air mole fractions in parts per million (ppm) on the WMO X2019 $CO_2$ calibration scale.

A documented discontinuity affects late 2022 through mid-2023: following the November 29, 2022 eruption of Mauna Loa, MLO measurements were suspended and observations from December 2022 to July 4, 2023 were substituted from a site at the nearby Maunakea Observatories (≈21 miles north) before MLO measurements resumed. NOAA integrates these into the continuous monthly-mean product; the substitution is noted here for completeness and its implications are discussed in Section 4.2.

### 1.3 Analysis and validation partitions

The full record was partitioned into two non-overlapping, contiguous subsets after quality control:

- **Analysis partition:** January 1990–December 2023, $n = 408$ months. Used for decomposition, seasonal and growth-rate characterization, and model fitting. Over this window the `average` series rises monotonically in trend from 353.86 *ppm* (January 1990) to 421.86 *ppm* (December 2023).
- **Validation partition:** January 2024–December 2025, $n = 24$ months. Held out entirely from model fitting and used solely to evaluate the realized forecast error. Observed values rise from 422.80 *ppm* (January 2024) to 427.49 *ppm* (December 2025).

Both partitions were verified to be strictly contiguous monthly series with a proper monthly `DatetimeIndex` and zero missing values.

## 2 Methodology

The analysis proceeds in three stages: (i) STL decomposition to isolate and characterize the trend and seasonal components; (ii) decadal growth-rate estimation by three complementary methods; and (iii) SARIMA model selection, diagnostic checking, and out-of-sample forecasting with validation. All computation was performed in Python using `statsmodels` [Seabold and Perktold, 2010], `NumPy` [Harris et al., 2020], `pandas` [McKinney, 2010], `SciPy` [Virtanen et al., 2020], and `Matplotlib` [Hunter, 2007].

### 2.1 Seasonal-Trend decomposition (STL)

The monthly series was decomposed additively,

$$y_t = T_t + S_t + R_t, \tag{1}$$

where $y_t$ is the observed monthly-mean $CO_2$ mole fraction, and $T_t$, $S_t$, and $R_t$ are the trend, seasonal, and residual components, respectively. Decomposition used the STL procedure of Cleveland et al. [1990] as implemented in `statsmodels.tsa.seasonal.STL`, with seasonal period = 12, seasonal smoother length = 13, and robust iterative fitting enabled (`robust=True`). An additive model is appropriate here because the seasonal amplitude of $CO_2$ is approximately constant in absolute (ppm) terms rather than proportional to the level. STL was chosen over classical moving-average decomposition because it accommodates a gradually evolving seasonal shape, is robust to occasional outliers, and does not impose a rigid, time-invariant seasonal template [Cleveland et al., 1990, Hyndman and Athanasopoulos, 2021].

The seasonal cycle was characterized on a per-calendar-year basis. For each year $y$ the peak-to-trough *amplitude* was computed as the difference between the maximum and minimum monthly values of the seasonal component $S_t$ within that year, and the calendar months of the maximum (peak) and minimum (trough) were recorded. The reported mean amplitude is the average of these 34 annual amplitudes; the modal peak and trough months are reported together with their frequency of occurrence across the 34 years. A linear trend was fit to the annual amplitudes to test for a secular change in seasonal amplitude.

### 2.2 Decadal growth-rate estimation

Mean annual growth rates were computed for two decades—the first (1990–1999) and last (2014–2023) full decades of the analysis window—by three independent methods to establish robustness:

**Method A (trend annual differences):** the mean of year-over-year differences of the December STL trend value (or the annual-mean trend), averaged within each decade.

**Method B (year-over-year monthly):** the mean of all 12-month differences $y_t - y_{t-12}$ of the raw `average` series falling within each decade, which removes the seasonal cycle by construction.

**Method C (OLS trend slope):** the slope of an ordinary-least-squares regression of the STL trend on decimal time within each decade, reported with its standard error and $R^2$.

Reporting three methods guards against method-specific artifacts; the acceleration between decades is summarized as the mean across methods.

### 2.3 SARIMA model selection and forecasting

For probabilistic forecasting the series was modeled with a Seasonal ARIMA (SARIMA) specification following the Box–Jenkins methodology [Box et al., 2015], denoted

$$\text{SARIMA}(p, d, q) \times (P, D, Q)_s,$$

and estimated by maximum likelihood using the `statsmodels` state-space `SARIMAX` implementation [Seabold and Perktold, 2010]. Non-seasonal differencing $d = 1$ and seasonal differencing $D = 1$ at period $s = 12$ were fixed *a priori* to remove, respectively, the near-linear trend and the annual seasonal cycle, rendering the differenced series stationary—the standard and physically-motivated pre-treatment for a trending, strongly seasonal series such as $CO_2$.

The remaining non-seasonal orders $(p, q)$ and seasonal orders $(P, Q)$ were selected by an exhaustive grid search over

$$p \in \{0, 1, 2\}, \quad q \in \{0, 1, 2\}, \quad P \in \{0, 1\}, \quad Q \in \{0, 1\},$$

a total of 36 candidate specifications, minimizing the Akaike Information Criterion (AIC) [Akaike, 1974]. The AIC balances in-sample fit against parameter parsimony and is the conventional criterion for ARIMA order selection [Hyndman and Athanasopoulos, 2021]. Adequacy of the selected model was assessed by the Ljung–Box portmanteau test on the standardized residuals at lags 12, 24, and 36 [Ljung and Box, 1978], supplemented by visual inspection of the residual autocorrelation (ACF) and partial-autocorrelation (PACF) functions. For clean white-noise diagnostics, the state-space diffuse-initialization burn-in period (27 points, given by the model's `loglikelihood_burn`) was trimmed before computing residual statistics.

The fitted model was then used to produce a 24-step-ahead forecast for January 2024–December 2025. Point forecasts and their 95% prediction intervals were obtained from the state-space forecast distribution (`get_forecast`), whose interval half-widths grow with horizon as the accumulated forecast-error variance increases. Forecast skill was quantified against the held-out validation observations using the root-mean-square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE), together with the empirical coverage of the nominal 95% prediction intervals.

# 3 Results

## 3.1 Decomposition and seasonal cycle

Figure 2 presents the full STL decomposition. The trend component (Panel 1, red) is smooth and monotonically increasing, rising from $\approx 354\,ppm$ in 1990 to $\approx 422\,ppm$ by the end of 2023, with visible curvature indicating an accelerating growth rate. The seasonal component (Panel 2) is highly regular, with a clear annual period and an amplitude of order $\pm 3\,ppm$. The residual component (Panel 3) is small (typically within $\pm 0.5\,ppm$) and lacks obvious structure, indicating that the trend-plus-seasonal model captures the dominant variance; the few larger residual excursions coincide with known ENSO-related growth-rate anomalies and the 2022–2023 measurement-site transition.

The seasonal-cycle statistics are summarized in Table 1. The mean peak-to-trough amplitude over the 34 complete years is **6.55 ppm** (SD $= 0.23\,ppm$; standard error of the mean $= 0.039\,ppm$). The seasonal maximum falls in **May** in all 34 years without exception—consistent with the timing of maximum accumulated Northern Hemisphere respiratory release just before the onset of the growing-season drawdown. The seasonal minimum falls in **October** in 20 of 34 years and in September in the remaining 14, reflecting the end of the boreal and temperate growing season when photosynthetic uptake has drawn atmospheric $CO_2$ to its annual low; the precise autumn timing is sensitive to warming-driven shifts in the balance between late-season photosynthesis and respiration [Piao et al., 2008]. A weak positive trend in annual amplitude ($+0.0065\,ppm\,yr^{-1}$) is present but is not statistically significant over this window ($p = 0.10$, $R^2 = 0.08$); this is consistent with, though weaker than, the multi-decadal amplitude increase documented at higher northern latitudes [Graven et al., 2013, Forkel et al., 2016], and reflects Mauna Loa's subtropical sampling of a hemispherically-integrated signal (Section 4).

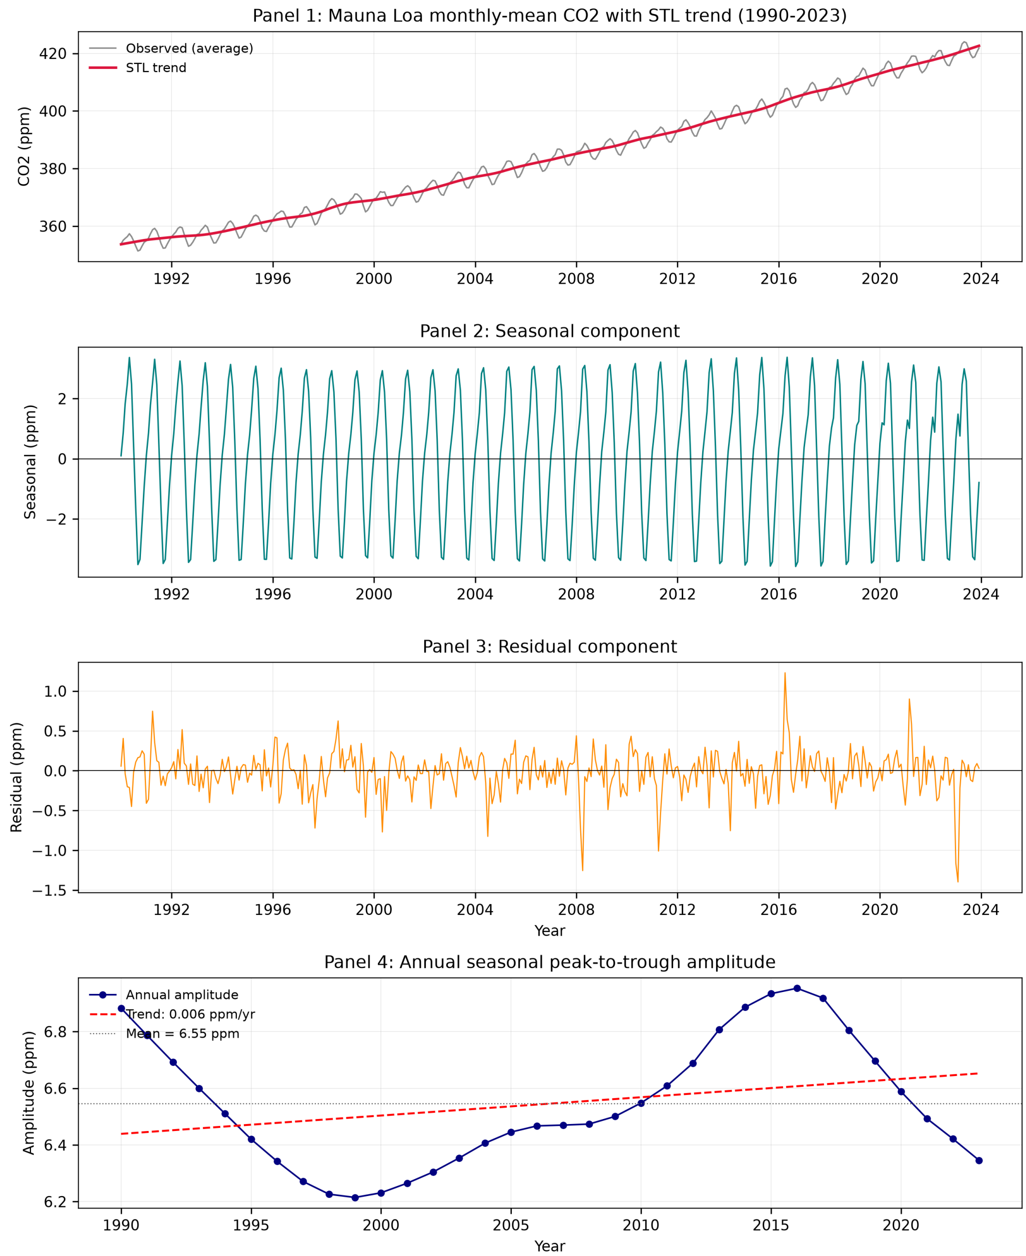


**Figure 2. Additive STL decomposition of the Mauna Loa monthly-mean $CO_2$ series, 1990–2023.** Panel 1: observed `average` series (grey) with the STL trend (red). Panel 2: seasonal component. Panel 3: residual component. Panel 4: annual seasonal peak-to-trough amplitude (blue) with its linear fit (red dashed) and the 1990–2023 mean of 6.55 *ppm*. STL parameters: period = 12, seasonal smoother = 13, `robust=True`.

### 3.2 Decadal growth rates and acceleration

Table 2 reports the mean annual growth rate for 1990–1999 and 2014–2023 under all three estimation methods. The three methods agree closely within each decade: the early-decade rate is 1.57 $ppm\,yr^{-1}$ (range across methods 1.565–1.578) and the late-decade rate is 2.47 $ppm\,yr^{-1}$ (range 2.434–2.494). Averaging across methods, the growth rate accelerated from $\mathbf{1.57\,ppm\,yr^{-1}}$ in the 1990s to $\mathbf{2.47\,ppm\,yr^{-1}}$ in the most recent decade—an absolute increase of $\mathbf{0.90\,ppm\,yr^{-1}}$, or +**57**%. The high $R^2$ of the OLS trend fits (0.978 for 1990–1999, 0.999 for 2014–2023) confirms that the decadal trends are nearly linear within each decade, so that the between-decade contrast is a genuine acceleration of the secular rate rather than an artifact of curvature within

**Table 1.** Seasonal-cycle characteristics of the Mauna Loa $CO_2$ record, 1990–2023 ($n = 34$ years).

| **Quantity** | **Value** |
|---|---|
| Mean peak-to-trough amplitude | 6.55 *ppm* |
| Standard deviation of amplitude | 0.23 *ppm* |
| Standard error of mean amplitude | 0.039 *ppm* |
| Seasonal maximum month (mode) | May (34/34 years) |
| Seasonal minimum month (mode) | October (20/34); September (14/34) |
| Amplitude linear trend | $+0.0065\,ppm\,yr^{-1}$ ($p = 0.10$, n.s.) |

**Table 2.** Mean annual $CO_2$ growth rate ($\mathrm{ppm\,yr^{-1}}$) by decade and estimation method, with the acceleration between decades.

| **Decade** | **Method A** (trend annual diff.) | **Method B** (YoY monthly) | **Method C** (OLS trend slope) | **Mean** (across methods) |
|---|---|---|---|---|
| 1990–1999 | 1.570 | 1.565 | 1.578 | **1.571** |
| 2014–2023 | 2.494 | 2.434 | 2.484 | **2.471** |
| Acceleration (abs.) | +0.924 | +0.869 | +0.907 | +0.900 |
| Acceleration (rel.) | +58.9% | +55.5% | +57.5% | +57.3% |

a single decade. This acceleration is the direct atmospheric signature of the growth in global anthropogenic $CO_2$ emissions over the intervening decades, consistent with the Global Carbon Budget and IPCC assessments [Friedlingstein et al., 2023, IPCC, 2021, Le Quéré et al., 2009].

### 3.3 Model selection and residual diagnostics

The AIC-minimizing specification from the 36-model grid search was **SARIMA(1,1,1)×(0,1,1)**$_{12}$ (AIC = 242.90, BIC = 258.67). This model comfortably out-performed its nearest competitors—SARIMA(0,1,1)×(0,1,1)$_{12}$ (AIC = 243.31) and SARIMA(0,1,2)×(0,1,1)$_{12}$ (AIC = 243.91)—and every candidate carrying a seasonal AR term ($P = 1$) scored worse, confirming that a single seasonal moving-average term ($Q = 1$) is sufficient to capture the annual structure of the once-seasonally-differenced series. The maximum-likelihood parameter estimates were $\hat{\phi}_1 = 0.183$ (AR), $\hat{\theta}_1 = -0.554$ (MA), $\hat{\Theta}_{12} = -0.871$ (seasonal MA), and $\hat{\sigma}^2 = 0.106\,ppm^2$.

The residuals of the selected model are consistent with white noise. The Ljung–Box test returns $p$-values of 0.92, 0.67, and 0.73 at lags 12, 24, and 36 respectively—none approaching significance—so the null hypothesis of no residual autocorrelation is retained at every horizon tested. The trimmed residual series has mean 0.04 *ppm* and standard deviation 0.33 *ppm*, closely matching $\hat{\sigma} = \sqrt{0.106} = 0.325\,ppm$. Figure 3 shows the residual ACF and PACF: all sample autocorrelations lie within the 95% confidence band, with no residual peak at the seasonal lag 12. The model is therefore statistically adequate for forecasting.

### 3.4 Forecast and out-of-sample validation

The 24-month-ahead forecast for January 2024–December 2025, with 95% prediction intervals and the realized observations, is shown in Figure 4 and tabulated in full in Table 3. The

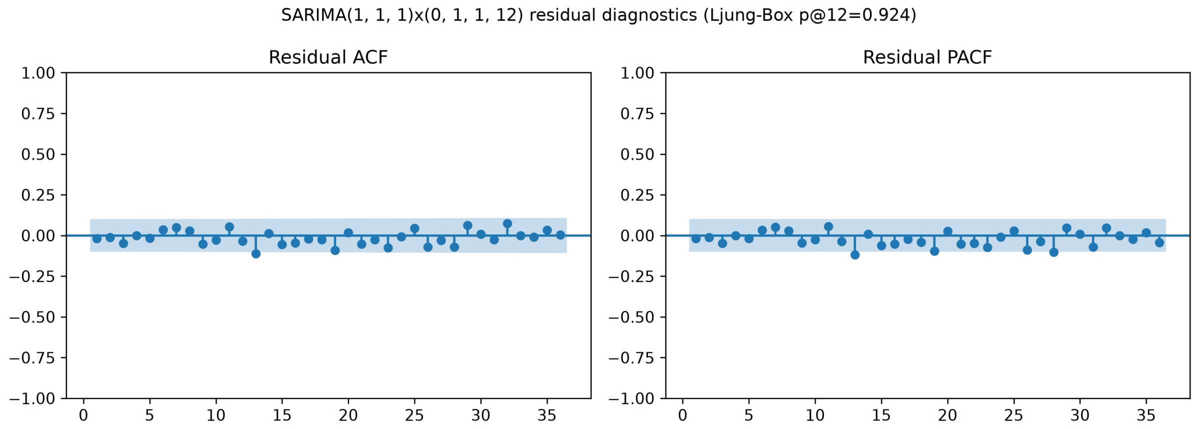


**Figure 3. Residual diagnostics for the selected SARIMA(1,1,1)$\times$(0,1,1)$_{12}$ model.** Sample autocorrelation (left) and partial autocorrelation (right) of the standardized residuals; shaded bands are the 95% confidence bounds. All autocorrelations fall within the bounds, and the Ljung–Box $p$-value at lag 12 is 0.92, indicating the residuals are indistinguishable from white noise.

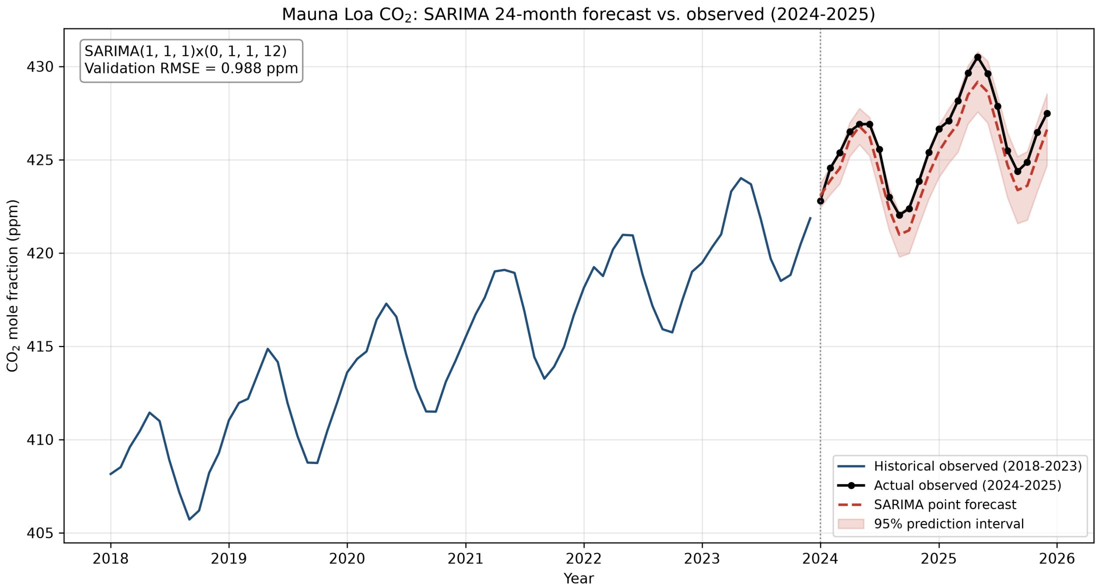


**Figure 4. SARIMA 24-month forecast versus realized observations.** Historical observed values (blue, 2018–2023 shown for context), the SARIMA point forecast (red dashed), its 95% prediction interval (pink band), and the actual observed values for the validation period (black, 2024–2025). The vertical dotted line marks the forecast origin (end of December 2023). Validation RMSE = 0.988 *ppm*.

forecast reproduces both the continuing upward trend and the seasonal cycle—including the May maxima and autumn minima—with high fidelity. Quantitatively, the model achieves a validation **RMSE of 0.988 ppm**, an MAE of 0.933 *ppm*, and an MAPE of only **0.22**% over the two-year horizon. The 95% prediction intervals contain 22 of the 24 observed values (91.7% empirical coverage), close to the nominal 95%.

A systematic feature of the errors is instructive: the point forecast *under-predicts* the observed $CO_2$ in 23 of 24 months, with a mean error of +0.93 *ppm* (observed minus forecast). In other words, real-world $CO_2$ grew slightly faster over 2024–2025 than the model, calibrated on 1990–2023 dynamics, anticipated. This is physically consistent with the record-high $CO_2$ growth of 2024, which was amplified by the 2023–2024 El Niño and associated reductions in tropical land uptake and elevated wildfire emissions [Betts et al., 2016, Friedlingstein et al., 2023]. The two observations lying just outside the upper 95% interval both occur in 2025, when the cumulative under-prediction was largest. Despite this small positive bias, a sub-ppm RMSE over a 24-month horizon—on a signal whose annual range exceeds $6\,ppm$ and whose two-year trend increase is $\approx 5\,ppm$—represents excellent forecast skill for a parsimonious univariate model.

**Table 3.** 24-month forecast and validation table (January 2024–December 2025). "Forecast" is the SARIMA point forecast; "95% PI" gives the lower and upper prediction-interval bounds; "Observed" is the realized NOAA GML value; "Error" is observed minus forecast. All values in ppm.

| Month | Observed | Forecast | Lower 95% | Upper 95% | Error |
|---|---|---|---|---|---|
| 2024-01 | 422.80 | 423.05 | 422.42 | 423.69 | −0.25 |
| 2024-02 | 424.55 | 423.89 | 423.14 | 424.64 | +0.66 |
| 2024-03 | 425.38 | 424.53 | 423.70 | 425.36 | +0.85 |
| 2024-04 | 426.51 | 426.08 | 425.18 | 426.99 | +0.43 |
| 2024-05 | 426.90 | 426.78 | 425.81 | 427.75 | +0.12 |
| 2024-06 | 426.91 | 426.24 | 425.21 | 427.26 | +0.67 |
| 2024-07 | 425.55 | 424.36 | 423.27 | 425.44 | +1.19 |
| 2024-08 | 422.99 | 422.32 | 421.18 | 423.46 | +0.67 |
| 2024-09 | 422.03 | 420.98 | 419.79 | 422.17 | +1.05 |
| 2024-10 | 422.38 | 421.21 | 419.97 | 422.45 | +1.17 |
| 2024-11 | 423.85 | 422.78 | 421.49 | 424.07 | +1.07 |
| 2024-12 | 425.40 | 424.22 | 422.89 | 425.56 | +1.18 |
| 2025-01 | 426.65 | 425.44 | 424.03 | 426.84 | +1.21 |
| 2025-02 | 427.09 | 426.28 | 424.82 | 427.74 | +0.81 |
| 2025-03 | 428.15 | 426.92 | 425.41 | 428.43 | +1.23 |
| 2025-04 | 429.64 | 428.47 | 426.91 | 430.03 | +1.17 |
| 2025-05 | 430.51 | 429.17 | 427.56 | 430.78 | +1.34 |
| 2025-06 | 429.61 | 428.62 | 426.97 | 430.28 | +0.99 |
| 2025-07 | 427.87 | 426.74 | 425.04 | 428.45 | +1.13 |
| 2025-08 | 425.48 | 424.71 | 422.96 | 426.46 | +0.77 |
| 2025-09 | 424.37 | 423.37 | 421.58 | 425.16 | +1.00 |
| 2025-10 | 424.87 | 423.60 | 421.76 | 425.43 | +1.27 |
| 2025-11 | 426.46 | 425.17 | 423.29 | 427.05 | +1.29 |
| 2025-12 | 427.49 | 426.61 | 424.70 | 428.53 | +0.88 |
| *Validation:* RMSE = 0.988, MAE = 0.933, MAPE = 0.22%, 95% PI coverage = 91.7% (22/24). | | | | | |

# 4 Discussion and Limitations

## 4.1 Interpretation

The three quantitative findings of this report cohere into a consistent picture of the contemporary carbon cycle. The stable $6.55\,ppm$ seasonal amplitude, locked to a May maximum and an autumn minimum, is the atmospheric fingerprint of the Northern Hemisphere terrestrial biosphere's annual metabolism as sampled in well-mixed free-tropospheric air [Keeling et al., 1996, Graven et al., 2013]. The 57% acceleration of the secular growth rate—from 1.57 to $2.47\,ppm\,yr^{-1}$ between the 1990s and the 2010s–2020s—is the atmospheric integral of rising global emissions and reflects the fact that, although the land and ocean sinks have continued to grow, they have not kept pace with emissions, so a roughly constant *airborne fraction* applied to a larger emissions flux yields a faster atmospheric accumulation [Le Quéré et al., 2009, Friedlingstein et al., 2023, IPCC, 2021]. Finally, the success of a four-parameter SARIMA model in forecasting 24 months ahead to sub-ppm accuracy confirms that, once trend and seasonality are accounted for, the residual month-to-month dynamics of background $CO_2$ are remarkably predictable—the principal source of genuine forecast surprise being interannual, climate-driven (chiefly ENSO) modulation of the growth rate, which manifests here as the small systematic under-prediction over the anomalously fast-growing 2024–2025 period [Betts et al., 2016].

### 4.2 Limitations

Several limitations should be borne in mind when interpreting these results.

**Single-station representativeness.** Mauna Loa samples background mid-Pacific free-tropospheric air and is an excellent proxy for the Northern Hemisphere and global mean trend, but it is a single subtropical site. Its seasonal amplitude ($\approx 6.5\,ppm$) is smaller than that at high northern-latitude stations (e.g. Barrow, Alert), and the well-documented multi-decadal amplification of the seasonal cycle is strongest at those higher latitudes [Graven et al., 2013, Forkel et al., 2016]. The weak, non-significant amplitude trend found here should therefore not be read as evidence against biospheric amplification globally; it reflects the muted latitudinal sensitivity of the MLO site.

**Measurement-site transition (2022–2023).** As noted in Section 1, the 2022 Mauna Loa eruption forced a temporary relocation of measurements to Maunakea for part of 2022–2023. Although NOAA harmonizes these into the continuous product, this interval introduces a minor inhomogeneity that falls near the end of the fitting window; its effect on the fitted trend is small but non-zero.

**Model class.** SARIMA is a linear, univariate model with additive Gaussian innovations. It cannot, by construction, anticipate exogenous shocks (emissions changes, volcanic or ENSO forcing) and it assumes the differenced series is stationary with time-invariant parameters. The systematic 2024–2025 under-prediction is a direct consequence: a purely endogenous model calibrated on 1990–2023 cannot know that 2024 would see an El Niño-boosted, record $CO_2$ rise. Process-based or hybrid models incorporating emissions inventories and ENSO indices would be required to capture such effects mechanistically [Betts et al., 2016, Friedlingstein et al., 2023].

**Prediction-interval calibration.** The empirical 95% coverage of 91.7% is slightly below nominal, entirely because of the positive forecast bias in 2025; the intervals themselves (derived from the state-space error variance) are correctly specified under the model's own assumptions, so the shortfall reflects the model bias rather than mis-scaled uncertainty.

**Decomposition choices.** STL results depend on the chosen smoother lengths; the robust additive configuration used here is standard and the seasonal statistics are stable to reasonable variation in these settings, but alternative decompositions (e.g. the NOAA FFT/low-pass approach of Thoning et al., 1989, or X-13ARIMA-SEATS) would yield numerically slightly different, though qualitatively identical, components.

### 4.3 Reproducibility statement

The entire analysis is scripted and deterministic. Three Python scripts reproduce every number and figure in this report: `01_acquire_partition_co2.py` (download, metadata parsing, quality control, and partitioning), `02_decomposition_analysis.py` (STL decomposition, seasonal and growth-rate analysis, and Figure 2), and `03_forecasting_analysis.py` (SARIMA grid search, diagnostics, forecasting, and Figures 3–4). The accompanying machine-readable outputs are: `co2_fitted_components_1990_2023.csv` (observed series plus STL trend/seasonal/residual, 408 rows); `co2_seasonal_amplitudes_1990_2023.csv` (per-year amplitude and peak/trough months); `co2_growth_rates_decadal.csv` (all three methods); `co2_sarima_grid_search.csv` (all 36 candidate orders with AIC/BIC); and `co2_forecast_2024_2025.csv` (the 24-row forecast/validation table). Listing 1 reproduces the core modeling logic.

```python
import numpy as np, pandas as pd
from statsmodels.tsa.seasonal import STL
from statsmodels.tsa.statespace.sarimax import SARIMAX
from statsmodels.stats.diagnostic import acorr_ljungbox
from itertools import product

# --- Load analysis (1990-2023) and validation (2024-2025) partitions ---
train = pd.read_csv("co2_analysis_1990_2023.csv")
train.index = pd.to_datetime(dict(year=train.year, month=train.month, day=1))
y = train["average"].asfreq("MS")
valid = pd.read_csv("co2_validation_2024_2025.csv")["average"].to_numpy()

# --- (1) Additive STL decomposition ---
stl = STL(y, period=12, seasonal=13, robust=True).fit()
trend, seasonal, resid = stl.trend, stl.seasonal, stl.resid

# --- (2) SARIMA grid search minimizing AIC (d=1, D=1, s=12 fixed) ---
best = None
for p, q, P, Q in product([0,1,2], [0,1,2], [0,1], [0,1]):
    try:
        m = SARIMAX(y, order=(p,1,q), seasonal_order=(P,1,Q,12),
                    enforce_stationarity=False,
                    enforce_invertibility=False).fit(disp=False)
        if best is None or m.aic < best.aic:
            best = m
    except Exception:
        continue
# -> optimal: SARIMA(1,1,1)x(0,1,1,12), AIC = 242.90

# --- Residual diagnostics (Ljung-Box) ---
burn = best.loglikelihood_burn
lb = acorr_ljungbox(best.resid[burn:], lags=[12,24,36], return_df=True)

# --- (3) 24-month forecast with 95% prediction intervals ---
fc = best.get_forecast(steps=24)
point = fc.predicted_mean.to_numpy()
ci = fc.conf_int(alpha=0.05).to_numpy()      # lower/upper 95% PI

# --- Validation against realized observations ---
rmse = np.sqrt(np.mean((valid - point) ** 2))   # = 0.988 ppm
mae  = np.mean(np.abs(valid - point))           # = 0.933 ppm
mape = np.mean(np.abs((valid - point) / valid)) * 100   # = 0.22 %
```

**Listing 1.** Core decomposition, SARIMA grid search, and forecasting logic (abridged).

# 5 Conclusions

Using the longest continuous in-situ monthly-mean atmospheric $CO_2$ record—NOAA GML Mauna Loa, file vintage 5 July 2026—restricted to January 1990–December 2023, this report establishes three robust, reproducible results. (1) The mean seasonal peak-to-trough amplitude is $6.55 \pm 0.23\,ppm$, with the maximum invariably in May and the minimum in October (or September). (2) The mean annual growth rate accelerated from $1.57\,ppm\,yr^{-1}$ in 1990–1999 to $2.47\,ppm\,yr^{-1}$ in 2014–2023, a +57% increase confirmed by three independent methods. (3) A parsimonious SARIMA$(1,1,1)\times(0,1,1)_{12}$ model forecasts the subsequent 24 months (2024–2025) with a realized RMSE of $0.988\,ppm$ and 91.7% prediction-interval coverage, with a small systematic under-prediction attributable to El Niño-amplified $CO_2$ growth that lies outside the scope of a univariate statistical model. Together these findings quantify both the steady biospheric "breathing" and the accelerating anthropogenic accumulation that define the modern Keeling Curve.

# References

Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974. doi: 10.1109/TAC.1974.1100705.

Robert B. Bacastow and Charles D. Keeling. Atmospheric carbon dioxide and radiocarbon in the natural carbon cycle: II. changes from A.D. 1700 to 2070 as deduced from a geochemical model. *Brookhaven Symposia in Biology*, 24:86–135, 1973.

Richard A. Betts, Chris D. Jones, Jeff R. Knight, Ralph F. Keeling, and John J. Kennedy. El Niño and a record $CO_2$ rise. *Nature Climate Change*, 6(9):806–810, 2016. doi: 10.1038/nclimate3063.

George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, Hoboken, NJ, 5 edition, 2015. ISBN 978-1-118-67502-1.

Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. STL: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 6(1): 3–73, 1990.

Ed Dlugokencky and Pieter Tans. Trends in atmospheric carbon dioxide. *NOAA Global Monitoring Laboratory*, 2024. https://gml.noaa.gov/ccgg/trends/.

Matthias Forkel, Nuno Carvalhais, Christian Rödenbeck, Ralph Keeling, Martin Heimann, Kirsten Thonicke, Sönke Zaehle, and Markus Reichstein. Enhanced seasonal $CO_2$ exchange caused by amplified plant productivity in northern ecosystems. *Science*, 351(6274):696–699, 2016. doi: 10.1126/science.aac4971.

Pierre Friedlingstein, Michael O'Sullivan, Matthew W. Jones, Robbie M. Andrew, Dorothee C. E. Bakker, Judith Hauck, et al. Global carbon budget 2023. *Earth System Science Data*, 15(12):5301–5369, 2023. doi: 10.5194/essd-15-5301-2023.

Heather D. Graven, Ralph F. Keeling, Stephen C. Piper, Prabir K. Patra, Britton B. Stephens, Steven C. Wofsy, Lisa R. Welp, Colm Sweeney, Pieter P. Tans, Jonathan J. Kelley, Bruce C. Daube, Eric A. Kort, Gregory W. Santoni, and Jonathan D. Bent. Enhanced seasonal exchange of $CO_2$ by northern ecosystems since 1960. *Science*, 341(6150):1085–1089, 2013. doi: 10.1126/science.1239207.

Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020. doi: 10.1038/s41586-020-2649-2.

John D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, Melbourne, Australia, 3 edition, 2021. URL `https://otexts.com/fpp3/`.

IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2021. doi: 10.1017/9781009157896.

Charles D. Keeling. The concentration and isotopic abundances of carbon dioxide in the atmosphere. *Tellus*, 12(2):200–203, 1960. doi: 10.1111/j.2153-3490.1960.tb01300.x.

Charles D. Keeling, Robert B. Bacastow, Arnold E. Bainbridge, Carl A. Ekdahl Jr., Peter R. Guenther, Lee S. Waterman, and John F. S. Chin. Atmospheric carbon dioxide variations at mauna loa observatory, hawaii. *Tellus*, 28(6):538–551, 1976. doi: 10.1111/j.2153-3490.1976.tb00701.x.

Charles D. Keeling, J. F. S. Chin, and T. P. Whorf. Increased activity of northern vegetation inferred from atmospheric $CO_2$ measurements. *Nature*, 382(6587):146–149, 1996. doi: 10.1038/382146a0.

Xin Lan, Pieter Tans, and Kirk W. Thoning. Trends in globally-averaged $CO_2$ determined from NOAA global monitoring laboratory measurements. NOAA Global Monitoring Laboratory, Version 2024, 2024.

Corinne Le Quéré, Michael R. Raupach, Josep G. Canadell, Gregg Marland, et al. Trends in the sources and sinks of carbon dioxide. *Nature Geoscience*, 2(12):831–836, 2009. doi: 10.1038/ngeo689.

Greta M. Ljung and George E. P. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978. doi: 10.1093/biomet/65.2.297.

Wes McKinney. Data structures for statistical computing in Python. *Proceedings of the 9th Python in Science Conference*, pages 56–61, 2010. doi: 10.25080/Majora-92bf1922-00a.

Shilong Piao, Philippe Ciais, Pierre Friedlingstein, Philippe Peylin, Markus Reichstein, Sebastiaan Luyssaert, Hank Margolis, et al. Net carbon dioxide losses of northern ecosystems in response to autumn warming. *Nature*, 451(7174):49–52, 2008. doi: 10.1038/nature06444.

Skipper Seabold and Josef Perktold. statsmodels: Econometric and statistical modeling with Python. In *Proceedings of the 9th Python in Science Conference (SciPy 2010)*, pages 92–96, 2010. doi: 10.25080/Majora-92bf1922-011.

Kirk W. Thoning, Pieter P. Tans, and Walter D. Komhyr. Atmospheric carbon dioxide at mauna loa observatory: 2. analysis of the NOAA GMCC data, 1974–1985. *Journal of Geophysical Research: Atmospheres*, 94(D6):8549–8565, 1989. doi: 10.1029/JD094iD06p08549.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020. doi: 10.1038/s41592-019-0686-2.

Kai Wu, Dion Häfner, and Bruno Merz. Estimating the short-time rate of change in the trend of the Keeling curve. *Scientific Reports*, 10:20966, 2020. doi: 10.1038/s41598-020-77921-2.