ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

Comparative Time-Series Modeling and Forecasting: Nonseasonal ARIMA for Global Internet Penetration and Seasonal ARIMA for Yellowstone Visitation

Yunting Wang

McGill University, Montréal, Canada, H3A 0G4

Abstract: *This paper presents two complementary applications of the Box-Jenkins approach to time-series modeling and forecasting. The first study analyzes annual global Internet penetration from 2005 to 2025 using a nonseasonal ARIMA framework. Successive differencing, formal stationarity testing, parameter screening, residual diagnostics, and chronological validation support the third-differenced moving-average model as the preferred specification relative to a fourth-differenced alternative. The second study analyzes monthly recreation visitation to Yellowstone National Park over the same period using a multiplicative SARIMA framework. Strong annual recurrence requires seasonal differencing and explicit seasonal dependence, and the evidence favors a model with one nonseasonal autoregressive term, one seasonal difference, and one seasonal moving-average term among the candidate specifications. Together, the studies demonstrate that model architecture must reflect sampling frequency and dependence structure: annual Internet penetration requires treatment of persistent nonseasonal trend, whereas monthly park visitation requires simultaneous representation of short-run and seasonal dynamics.*

Keywords: *ARIMA, SARIMA, Box-Jenkins Methodology, Internet Penetration, National Park Visitation, Seasonality, Forecast Validation*

1. Introduction

Digital connectivity and protected-area visitation represent distinct planning problems with a common statistical challenge: both evolve over time and require forecasts consistent with their dependence structures. Global Internet penetration informs assessments of digital inclusion and the future scale of the digital economy. Yellowstone National Park visitation informs transportation management, congestion control, staffing, maintenance, and environmental planning.

The two datasets contrast nonseasonal and seasonal modeling. The Internet series contains 21 annual observations from 2005 to 2025, so no within-year pattern can be identified and a nonseasonal ARIMA model is appropriate. The Yellowstone series contains 252 monthly observations from January 2005 through December 2025, so its annual cycle requires a multiplicative SARIMA model representing both ordinary and seasonal dynamics.

The combined analysis has three objectives. First, it demonstrates how stationarity assessment and correlogram-based identification differ between annual and monthly data. Second, it compares how model estimation, diagnostic checking, information criteria, and chronological validation contribute to model selection in each setting. Third, it draws broader lessons about the importance of aligning model structure with substantive context, sampling frequency, and forecast horizon.

2. Unified Methodological Framework

2.1. Nonseasonal and Seasonal ARIMA Models

An ARIMA(p, d, q) process is written as follows [1, 3, 6]:

$$\varphi(B)(1-B)^d x_t = \theta(B)\varepsilon_t$$

Here B is the backshift operator, d is the ordinary differencing order, φ(B) is the autoregressive polynomial, θ(B) is the moving-average polynomial, and ε_t denotes white noise. A SARIMA(p, d, q)(P,

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

D, Q)$_S$ process with seasonal period S satisfies:

$$\Phi(B^S)\varphi(B)(1-B)^d(1-B^S)^D x_t = \Theta(B^S)\theta(B)\varepsilon_t$$

Here P, D, and Q are the seasonal autoregressive, differencing, and moving-average orders, respectively, and Φ and Θ denote the seasonal polynomials. For monthly visitation, S = 12, and the first seasonal difference is:

$$(1-B^{12})x_t = x_t - x_{t-12}$$

2.2. Identification, Estimation, Diagnostics, and Validation

Both studies follow the Box-Jenkins sequence: identification, maximum-likelihood estimation, diagnostic checking, model comparison, chronological holdout validation, and final forecasting. Candidate models are evaluated using parameter significance, residual plots, residual autocorrelation, Ljung-Box tests, AIC, AICc, BIC, parsimony, forecast uncertainty, and out-of-sample accuracy [1, 3, 6]. ACF and PACF patterns are used only as preliminary order-identification evidence; final retention depends on the joint behavior of statistical support, residual adequacy, information criteria, and validation performance. Table 1 summarizes the nonseasonal ACF and PACF signatures, while Table 2 summarizes their seasonal counterparts. Residual serial dependence is assessed using the Ljung–Box portmanteau test [5].

Table 1: Indicative ACF and PACF signatures for non-seasonal models.

	AR(p)	**MA(q)**	**ARMA(p, q)**
ACF	Tails off	Cuts off after lag q	Tails off
PACF	Cuts off after lag p	Tails off	Tails off

Table 2: Indicative ACF and PACF signatures for seasonal AR, MA, and ARMA processes.

Model	**ACF**	**PACF**
SAR(P)	Tails off	Cuts off after lag PS
SMA(Q)	Cuts off after lag QS	Tails off
SARMA(P, Q)	Tails off	Tails off

3. Study 1: ARIMA Modeling of Global Internet Usage

3.1. Methodology and Data

The first empirical analysis uses the World Bank indicator Individuals using the Internet (% of population) [7]. The dataset contains 21 annual global observations from 2005 to 2025 and represents a population-weighted measure of worldwide Internet penetration. Because the series has only one observation per year, no seasonal pattern can be estimated. The original series rises from 15.6% in 2005 to 73.6% in 2025, whereas the PACF displays a dominant first-lag spike. These features indicate strong persistence and suggest that the level series does not satisfy the stationarity assumption required for ARIMA modelling.

3.2. Stationarity and Differencing

Successive ordinary differences were examined with time-series plots, ACF and PACF diagnostics, and ADF tests. Figure 1 presents the time-series plots and correlograms for the successive transformations, and Table 3 reports the corresponding ADF test results.

The original, first-differenced, and second-differenced series failed to reject the unit-root null. The third-differenced series rejected the null at the 5% level, making d = 3 the minimum differencing order adopted in this study. The fourth difference also rejected the null more strongly but was retained only as an alternative stationary specification because unnecessary differencing can remove meaningful serial dependence.

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

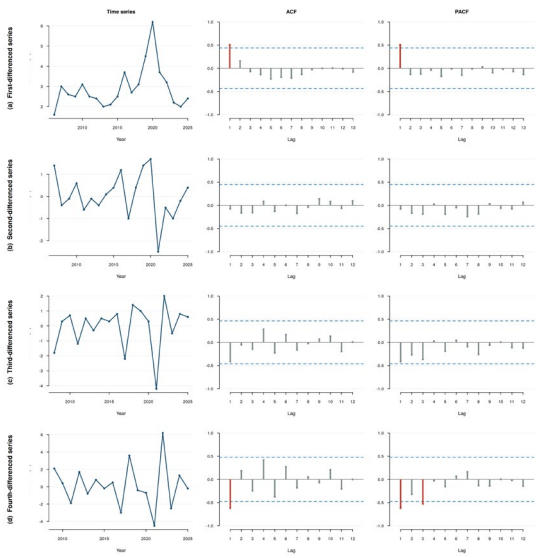


Figure 1. Time-series plots, sample ACFs, and sample PACFs of the progressively differenced global Internet penetration series. Dashed lines denote approximate 95% white-noise confidence limits.

Note. Bars highlighted in red indicate autocorrelations whose absolute values exceed the approximate 95% confidence limits.

Table 3: ADF stationarity tests for global Internet penetration.

Series	ADF	p-value	Decision
Original	-2.1648	0.5096	Fail to reject
First difference	-2.0013	0.5719	Fail to reject
Second difference	-2.8121	0.2630	Fail to reject
Third difference	-3.6799	0.0443	Reject
Fourth difference	-4.8164	0.0100	Reject

3.3. Candidate Identification and Selection

After third differencing, most ACF and PACF coefficients lie within approximate 95% confidence limits, but a pronounced negative lag-1 ACF supports a first-order moving-average component. ARIMA(0, 3, 1) was therefore identified as the primary candidate, while ARIMA(1, 3, 1) was retained to test whether an additional autoregressive term improved fit. For the fourth-differenced series, the ACF exhibits a pronounced negative spike at lag 1, while weaker dependence extends beyond the first lag without displaying a clear higher-order cut-off pattern. This residual short-run correlation provides a basis for considering an MA specification beyond first order. In contrast, the PACF does not exhibit a clear autoregressive cut-off pattern, providing little evidence for the inclusion of AR terms. Accordingly, ARIMA(0, 4, 2) was retained as an alternative candidate for subsequent parameter estimation and diagnostic evaluation.

Parameter estimation eliminated ARIMA(1, 3, 1) because its AR(1) coefficient was not significant (estimate = -0.0245, p = 0.9232). ARIMA(0, 3, 1) had a significant MA(1) estimate at the invertibility boundary, while ARIMA(0, 4, 2) had significant MA(1) and MA(2) terms. Table 4 reports the parameter estimates and significance tests used in this screening.

Table 4: Parameter estimates and statistical screening of candidate ARIMA models.

Model	Parameter	Estimate	p-value	Interpretation
ARIMA(0, 3, 1)	MA(1)	-1.0000	0.0001	Significant; boundary estimate
ARIMA(1, 3, 1)	AR(1)	-0.0245	0.9232	Not significant
	MA(1)	-1.0000	0.0002	Significant; boundary estimate
ARIMA(0, 4, 2)	MA(1)	-1.9003	0.0001	Significant
	MA(2)	0.9202	0.0233	Significant

Residual diagnostics indicated ARIMA(0, 3, 1) and ARIMA(0, 4, 2) both have approximately zero residual means, approximately constant variance, no significant residual ACF spikes among seven displayed lags, and Ljung-Box p-values above 0.05 for both retained models. Normality was approximate, with a lower-tail deviation around 2021, and Gaussian prediction intervals should be

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

interpreted cautiously in the presence of exceptional shocks. Figure 2 presents the residual diagnostics for both retained ARIMA models.

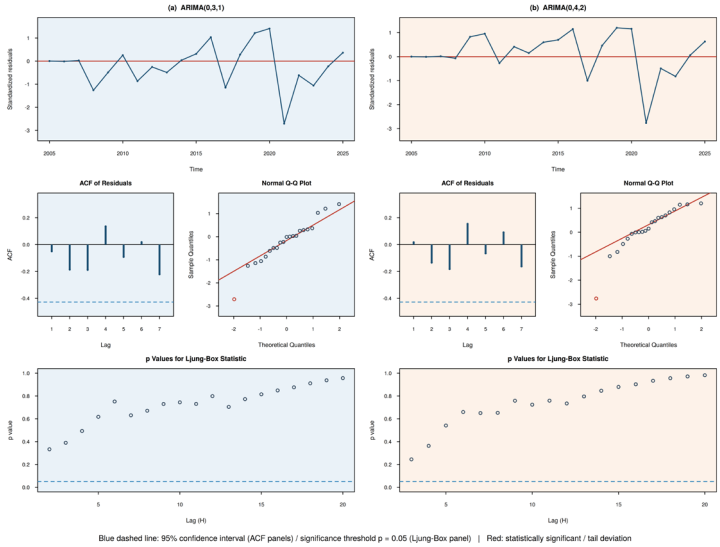


Figure 2. Residual diagnostics for ARIMA(0, 3, 1) and ARIMA(0, 4, 2), including standardized residuals, residual ACFs, normal Q–Q plots, and Ljung–Box test p-values.

Both models removed the observable residual autocorrelation sufficiently well, so the next stage of comparison focused on whether the additional differencing order and extra moving-average parameter in ARIMA(0, 4, 2) produced enough improvement in likelihood-based fit to justify the loss of parsimony. The Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) both penalize unnecessary complexity while rewarding improved fit. Following the information-theoretic model-selection framework [2], lower values therefore indicate a more favorable balance between explanatory adequacy and model economy. In this comparison, both criteria favored ARIMA(0, 3, 1), as shown in Table 5.

Table 5: Information-criterion comparison of the diagnostically adequate ARIMA models.

Model	AIC	BIC
ARIMA(0, 3, 1)	3.241	3.339
ARIMA(0, 4, 2)	4.112	4.259

The information-criterion evidence therefore establishes a first preference for ARIMA(0, 3, 1).

Information criteria evaluate in-sample likelihood and parsimony, but they do not directly measure forecast accuracy. A chronological holdout validation was therefore conducted before final model selection. The sample was divided according to its temporal order: observations from 2005 to 2019 were used for model estimation, and observations from 2020 to 2025 were reserved for out-of-sample evaluation. This design preserves the forecasting chronology and avoids information leakage from future observations.

Forecast accuracy was measured using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). Lower values on all three measures indicate better out-of-sample forecasting accuracy. Table 6 reports the holdout accuracy measures for the two retained specifications.

Table 6: Forecast accuracy under chronological train-test validation.

Model	MAE	RMSE	MAPE (%)	Rank
ARIMA(0, 3, 1)	4.818	6.406	6.804	1
ARIMA(0, 4, 2)	14.826	19.856	20.842	2

ARIMA(0, 3, 1) outperforms ARIMA(0, 4, 2) on every forecasting measure. The fourth-differenced model produces forecast errors approximately three times as large as those of the third-differenced specification. This contrast shows that stronger rejection of the unit-root null after fourth differencing does not translate into superior forecast performance. The chronological validation therefore provides direct predictive evidence in favor of ARIMA(0, 3, 1).

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

Final model selection was based on the full sequence of Box-Jenkins evidence rather than on any single diagnostic. Taken together, the evidence supports ARIMA(0, 3, 1) as the preferred specification. It is the most parsimonious stationary model retained after the ADF screen, its dynamic parameter is statistically supported, its residual behavior is adequate, its AIC and BIC values are lower, and its chronological holdout errors are substantially smaller. ARIMA(0, 4, 2) is not selected because it is associated with materially weaker out-of-sample performance.

3.4. Forecasts and Interpretation

ARIMA(0, 3, 1) was re-estimated using the complete 2005-2025 sample and used to forecast 2026-2035. The point forecasts increase from 76.04% in 2026 to 99.92% in 2035, broadly consistent with diffusion toward saturation. However, conventional ARIMA forecasts are unconstrained and do not explicitly respect the 0%-100% bound. The 95% prediction-interval width increases from 4.06 percentage points at h = 1 to 92.22 percentage points at h = 10. This substantial expansion indicates that forecast uncertainty accumulates rapidly as the prediction horizon extends, so later forecasts should be interpreted as indicative projections rather than precise estimates. Table 7 lists the annual point forecasts, and Figure 3 displays the historical series, forecast path, and prediction intervals.

Table 7: Ten-year point forecasts from the selected ARIMA(0, 3, 1) model.

Year	Forecast (%)
2026	76.04
2027	78.53
2028	81.05
2029	83.62
2030	86.23
2031	88.88
2032	91.58
2033	94.32
2034	97.09
2035	99.92

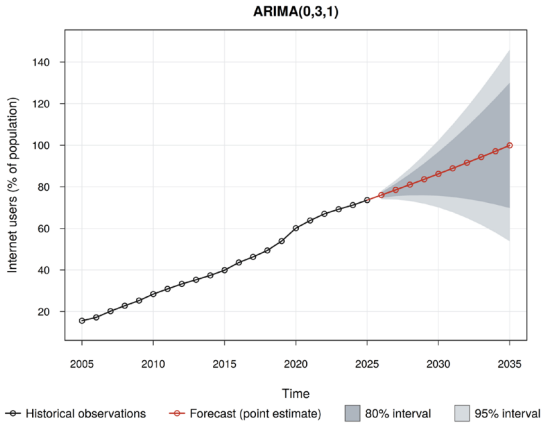


Figure 3. Historical observations and ten-year forecasts from the selected ARIMA(0, 3, 1) model, including point forecasts with 80% and 95% prediction intervals.

4. Study 2: SARIMA Modeling of Yellowstone National Park Visitation

4.1. Data and Seasonal Structure

The second empirical analysis uses monthly recreation visitation to Yellowstone National Park from the National Park Service Visitor Use Statistics [8]. The sample covers January 2005 through December 2025 and contains 252 monthly observations, or 21 complete annual cycles. The raw series exhibits a regular summer peak, low winter visitation, gradual long-run change, and an exceptional temporary disturbance around 2020. The mean monthly visitation is approximately 313,393 visits, with a sample standard deviation of 343,768 visits. The minimum is 0 visits in April 2020, and the

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

maximum is 1,080,767 visits in July 2021.

Additive decomposition confirms a stable and pronounced annual pattern. Figure 4 displays the observed, trend, seasonal, and irregular components.

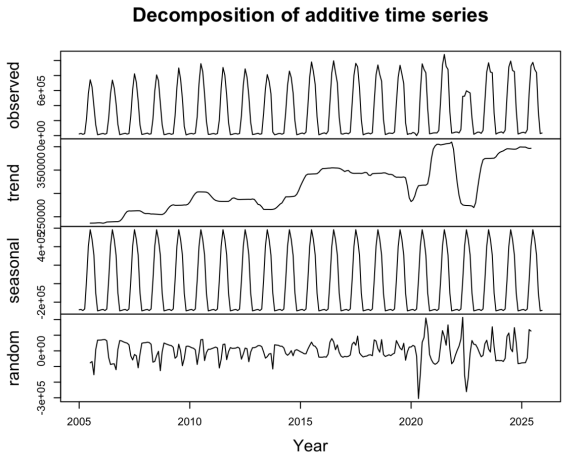


Figure 4. Additive decomposition of monthly Yellowstone recreation visitation into observed, trend, seasonal, and irregular components.

The seasonal component reaches its maximum in July at approximately 583,341 visits above the trend-cycle level and its minimum in November at approximately 294,211 visits below that level, giving a peak-to-trough seasonal amplitude of about 877,553 visits. This magnitude and regularity justify a seasonal time-series model rather than a purely nonseasonal ARIMA specification.

4.2. Stationarity, Differencing, and Identification

ADF tests rejected the unit-root null for the original series and for ordinary, seasonal, and combined differenced transformations. Because the ADF test primarily evaluates nonseasonal unit roots, this result alone did not determine SARIMA differencing orders. Figure 5 compares the original series and alternative differenced transformations alongside their ACFs and PACFs.

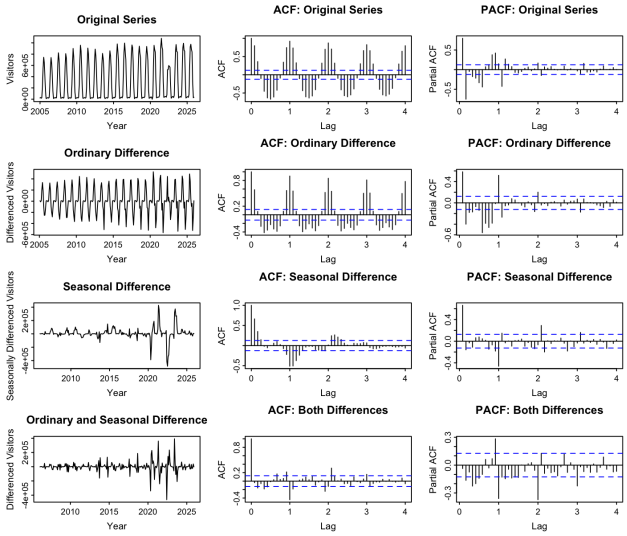


Figure 5. Original monthly visitation series and alternative differenced transformations with their ACF and PACF functions. Dashed lines denote approximate 95% confidence limits for zero autocorrelation. For monthly data with seasonal frequency S = 12, the displayed lag scale corresponds to annual cycles, where lags 1, 2, 3, and 4 represent 12, 24, 36, and 48 months, respectively.

The sample ACF displays large positive correlations at seasonal lags (12, 24, 36 and 48 months). The ndiffs() procedure recommended no ordinary difference (d = 0), while nsdiffs() recommended one seasonal difference (D = 1), following the automated differencing framework implemented in the

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

forecast package for R [4]. Subsequent identification therefore used the seasonally differenced series $(1 - B^{12})x_t$.

In the seasonally differenced series, the short-lag ACF decays gradually and the PACF has a dominant first-lag spike, supporting a nonseasonal AR(1) component. At the seasonal frequency, both ACF and PACF show pronounced negative spikes at the first seasonal lag, with coefficients about -0.511 and -0.461. Because the two correlograms do not clearly distinguish between these seasonal structures, three parsimonious candidate specifications were retained: SARIMA(1, 0, 0)(1, 1, 0)$_{12}$, SARIMA(1, 0, 0)(0, 1, 1)$_{12}$, and SARIMA(1, 0, 0)(1, 1, 1)$_{12}$.

4.3. Estimation, Diagnostics, and Selection

Maximum-likelihood estimation showed that SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ had significant AR(1) and SMA(1) parameters, and SARIMA(1, 0, 0)(1, 1, 0)$_{12}$ had significant AR(1) and SAR(1) parameters. SARIMA(1, 0, 0)(1, 1, 1)$_{12}$ was excluded because its added seasonal AR coefficient was not significant (estimate = -0.0598, p = 0.488). Table 8 summarizes the estimated coefficients, standard errors, and significance tests.

Table 8: Maximum-likelihood parameter estimates for the candidate SARIMA models.

Model	Parameter	Estimate	SE	p-value	Significant
SARIMA(1, 0, 0)(0, 1, 1)$_{12}$	AR(1)	0.6603	0.0481	<0.001	Yes
	SMA(1)	-0.6391	0.0454	<0.001	Yes
SARIMA(1, 0, 0)(1, 1, 1)$_{12}$	AR(1)	0.6589	0.0483	<0.001	Yes
	SAR(1)	-0.0598	0.0862	0.488	No
	SMA(1)	-0.6357	0.0462	<0.001	Yes
SARIMA(1, 0, 0)(1, 1, 0)$_{12}$	AR(1)	0.6402	0.0493	<0.001	Yes
	SAR(1)	-0.4343	0.0559	<0.001	Yes

Residual diagnostics indicated that both retained models captured the major seasonal dependence but did not produce fully independent Gaussian residuals. Figure 6 presents the residual diagnostics for the retained SARIMA specifications.

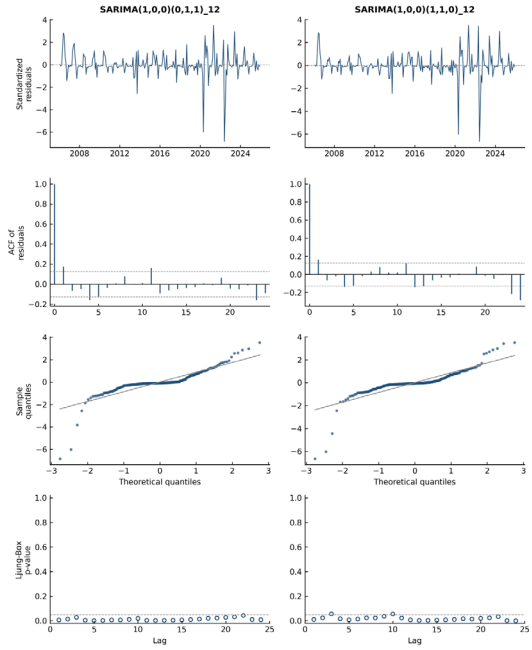


Figure 6. Residual diagnostics for the retained SARIMA specifications, including standardized residuals, residual ACFs, normal Q–Q plots, and Ljung–Box p-values.

In both cases, the residual mean is approximately zero. However, residual variability increases after 2020, indicating that the assumption of constant variance is questionable. This interpretation is reinforced by the largest absolute standardized residuals, which are approximately 6.95 for SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ and 6.85 for SARIMA(1, 0, 0)(1, 1, 0)$_{12}$.

The seasonal moving-average model leaves several significant residual ACF spikes, especially near

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

the second seasonal cycle, whereas the seasonal autoregressive model produces more pronounced and persistent residual spikes. At lag 24, Ljung-Box tests rejected the white-noise null for both models, with Q = 45.730 and p = 0.0021 for the seasonal MA model and Q = 84.278 and p = 3.16×10^{-9} for the seasonal AR model. The normal Q-Q diagnostics further indicate strong lower-tail deviations and additional upper-tail departures, so the Gaussian error assumption is not well supported.

The residual diagnostics therefore show the seasonal moving-average specification is comparatively preferable because it leaves fewer and weaker residual correlations, but this relative advantage should not be interpreted as full residual adequacy.

The information-criterion results first provided clear in-sample evidence in favor of the seasonal moving-average specification. SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ achieved lower AIC, AICc, and BIC values. The resulting Delta AICc of 35.133 indicates that the seasonal autoregressive alternative is much less supported once model fit and complexity are evaluated jointly. Table 9 reports the information criteria and the relative AICc differences [2].

Table 9: Information-criterion comparison of statistically admissible SARIMA models.

Candidate model	AIC	AICc	BIC	ΔAICc
SARIMA(1, 0, 0)(0, 1, 1)$_{12}$	5927.296	5927.398	5937.738	0.000
SARIMA(1, 0, 0)(1, 1, 0)$_{12}$	5962.429	5962.531	5972.871	35.133

Using January 2005-December 2022 as the training period and January 2023-December 2025 as the 36-month test period, the seasonal moving-average model again performed better in chronological holdout evaluation. It reduced RMSE and MAE by approximately 26.5% and 20.2%, respectively, relative to the seasonal autoregressive benchmark. Table 10 compares their out-of-sample RMSE and MAE.

Table 10: Out-of-sample forecast accuracy over the January 2023-December 2025 validation period.

Model	Seasonal structure	RMSE	MAE
SARIMA(1, 0, 0)(0, 1, 1)$_{12}$	Seasonal MA	91,791.53	62,513.59
SARIMA(1, 0, 0)(1, 1, 0)$_{12}$	Seasonal AR	124,869.51	78,375.21

Considering parameter support, information criteria, residual adequacy, parsimony, and predictive accuracy jointly, SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ was selected as the final forecasting model.

4.4. Forecasts and Interpretation

The selected SARIMA model was re-estimated using all observations from January 2005 through December 2025 and used to forecast January-October 2026. Figure 7 displays the historical series and forecast intervals, and Table 11 reports the monthly point forecasts and 95% prediction bounds.

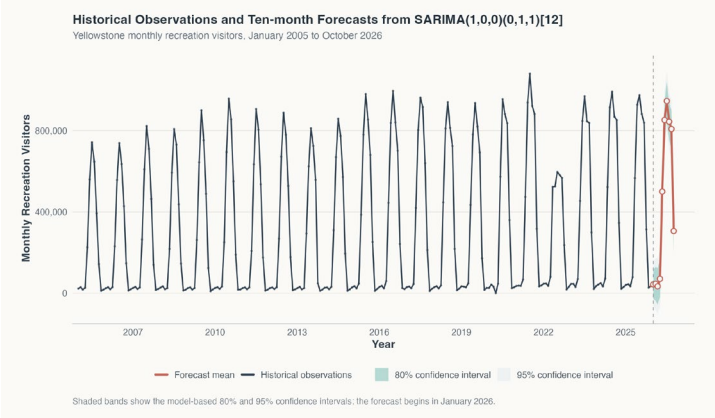


Figure 7. Historical observations and ten-month forecasts from the selected SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ model, including point forecasts and 80% and 95% prediction intervals.

The forecast path reproduces Yellowstone's established annual cycle: low visitation during winter and early spring, rapid growth in late spring, a July peak, and autumn decline. The first four lower 95% prediction bounds are negative because the model uses symmetric Gaussian intervals on a non-negative count-like scale; these lower bounds are not feasible visitation outcomes.

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

Table 11: Yellowstone forecasts and 95% prediction intervals for January-October 2026.

Month	Forecast	Lower 95%	Upper 95%
Jan 2026	42,924.85	-64,002.72	149,852.40
Feb 2026	45,994.54	-82,142.27	174,131.40
Mar 2026	33,405.21	-102,950.82	169,761.20
Apr 2026	70,374.88	-69,413.88	210,163.60
May 2026	500,318.96	359,059.49	641,578.40
Jun 2026	852,661.18	710,765.18	994,557.20
Jul 2026	946,109.19	803,936.53	1,088,281.90
Aug 2026	844,352.93	702,059.80	986,646.10
Sep 2026	808,207.71	665,862.08	950,553.30
Oct 2026	305,914.55	163,546.03	448,283.10

5. Discussion

The two studies show why ARIMA and SARIMA should be viewed as related but structurally distinct members of the Box-Jenkins family. Annual global Internet penetration has no estimable within-year recurrence and is dominated by persistent diffusion trend, while monthly Yellowstone visitation is accompanied by annual seasonality and meaningful month-to-month persistence. The Internet analysis therefore centers on ordinary differencing and the risk of over-differencing; the Yellowstone analysis centers on seasonal differencing and the choice between seasonal AR and seasonal MA dependence.

Model comparison also shows that final selection cannot rely on a single diagnostic. In the Internet study, ARIMA(0, 4, 2) satisfied a stronger differencing criterion and produced adequate residuals, but it had higher information criteria and much worse holdout accuracy than ARIMA(0, 3, 1). In the Yellowstone study, both retained seasonal specifications had significant parameters and residual limitations, but the seasonal MA model achieved lower information criteria and superior chronological validation accuracy. Across both applications, selected models are preferred because they provide the most favorable combination of parameter support, residual behavior, information criteria, forecast performance, parsimony, and interpretability within their candidate sets.

5.1. Limitations

The Internet analysis is limited by the short annual sample, marginal ADF rejection at d = 3, a boundary MA(1) estimate, sensitivity to the exceptional 2021 observation, and the use of an unconstrained linear model for a bounded 0%-100% outcome. The Yellowstone analysis is limited by only 21 seasonal replications, remaining residual autocorrelation, heavy-tailed residual behavior around the 2020 disruption, negative lower prediction bounds in low-visitation months, and reliance on a single chronological holdout period. Both models are univariate and do not explicitly include policy, technological, socioeconomic, weather, road-access, holiday, or disruption variables.

5.2. Directions for Future Research

Future Internet penetration research could compare ARIMA with bounded logistic or Gompertz diffusion models, intervention models, and multivariate time-series frameworks incorporating explanatory covariates. Future Yellowstone research could use variance-stabilizing transformations such as logarithmic or Box–Cox to address heteroscedasticity and eliminate infeasible negative prediction bounds, SARIMAX or dynamic regression with external drivers, explicit intervention variables for shocks such as COVID-19, and rolling-origin time-series cross-validation to reduce dependence on specific historical periods. These extensions would address the scale constraints, external shocks, and validation uncertainty identified in the present analysis.

6. Conclusion

The combined evidence demonstrates that effective time-series modeling depends on matching the model architecture to the temporal structure of the data. Annual global Internet penetration is appropriately treated as a nonseasonal process, and the selected ARIMA(0, 3, 1) model provides the best balance of stationarity, parsimony, information criteria, and chronological holdout accuracy among

ISSN 2616-5775 Vol. 9, Issue 8: 118-127, DOI: 10.25236/AJCIS.2026.090815

the evaluated specifications. Monthly Yellowstone visitation exhibits strong annual seasonality, and the selected SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ model provides the most competitive short-term forecasting baseline among the tested seasonal alternatives.

The selected models should be interpreted as context-specific forecasting tools rather than universally optimal solutions. ARIMA(0, 3, 1) produces an interpretable upward diffusion trajectory but long-horizon uncertainty and bounded-outcome issues remain important. SARIMA(1, 0, 0)(0, 1, 1)$_{12}$ reproduces Yellowstone's seasonal profile and offers operationally useful short-horizon forecasts, but residual dependence and external disruptions limit claims of complete adequacy. In both cases, robust model selection requires balancing statistical adequacy with forecast performance, interpretability, scale constraints, and the decision context in which forecasts will be used.

References

[1] Box, G.E.P., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2015) Time Series Analysis: Forecasting and Control. 5th Edition, Wiley.
[2] Burnham, K.P. and Anderson, D.R. (2002) Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach. 2nd Edition, Springer.
[3] Hyndman, R.J. and Athanasopoulos, G. (2021) Forecasting: Principles and Practice. 3rd Edition, OTexts.
[4] Hyndman, R.J. and Khandakar, Y. (2008) Automatic Time Series Forecasting: The forecast Package for R. Journal of Statistical Software, 27(3), 1-22.
[5] Ljung, G.M. and Box, G.E.P. (1978) On a Measure of Lack of Fit in Time Series Models. Biometrika, 65(2), 297-303.
[6] Shumway, R.H. and Stoffer, D.S. (2025) Time Series Analysis and Its Applications: With R Examples. 5th Edition, Springer.
[7] World Bank. (2026) World Development Indicators: Individuals Using the Internet (% of Population). World Bank DataBank. https://data.worldbank.org/indicator/IT.NET.USER.ZS
[8] National Park Service. (2026) Visitor Use Statistics. U.S. Department of the Interior. https://irma.nps.gov/Stats/