

Conformal Prediction for Enterprise Database Capacity Planning

Sriram Anjanadri sanjanadri@dx.com

Enterprise Database and Infrastructure Capacity Planning Hyderabad, India

Abstract - This paper looks at enterprise database capacity forecasting from a practical business-planning point of view. In large organisations, storage growth is not just a technical concern; it can affect procurement timing, budget approvals, service continuity, and operational risk. The proposed framework combines SARIMAX forecasting with selected business drivers and residual-based split conformal prediction. Instead of treating the forecast as one fixed future path, the approach turns it into a calibrated interval that reflects uncertainty from changing business volumes, retention policies, regulatory needs, and infrastructure lead times. By using historical forecast errors for calibration, the method presents capacity growth as a planning window rather than a single predicted date. For the evaluated 6.2-unit threshold, the framework gives a planning window from March 16, 2026, to May 3, 2026, with the central forecast crossing on April 9, 2026. This result shows how uncertainty-aware forecasting can support clearer and more defensible infrastructure investment decisions.

Key Words: Enterprise database, capacity planning, conformal prediction, SARIMAX, uncertainty quantification, agentic AI, forecasting models.

1. INTRODUCTION

Enterprise databases are central to daily operations in sectors such as banking, insurance, healthcare, and telecommunications. As digital services expand, database growth is increasingly shaped by customer activity, transactions, documents, applications, and regulatory records.

For infrastructure teams, forecasting this growth is not only a technical task. It affects budget planning, procurement lead time, storage allocation, and service-risk management [5], [6].

Traditional forecasts usually give a single expected value or date, but that is often too narrow for real planning because business demand, retention rules, and operational workloads can change over time.

This paper uses conformal prediction to express that uncertainty as a planning window. Rather than asking planners to act on one predicted threshold-crossing date, the method provides upper and lower bounds calibrated from past forecast errors [2], [3]. This makes the forecast easier to discuss across technical and business teams.

The study combines SARIMAX forecasting with selected business drivers and applies residual-based split conformal prediction to produce calibrated capacity intervals [1], [16]. It also outlines how future AI-assisted and agentic workflows could help planners test business-growth scenarios more easily.

The main contributions are:

- A database capacity forecasting framework that uses business drivers as external inputs.
- Use of conformal prediction to convert point forecasts into calibrated planning intervals.
- Evaluation of the approach using enterprise operational data anonymised for security.
- A roadmap for scenario testing, ensemble forecasting, and AI-assisted capacity planning.

2. LITERATURE REVIEW

Earlier work on forecasting shows that simple trend-based estimates are useful for monitoring, but they are not enough when planning decisions depend on changing business activity and long infrastructure lead times. In many enterprise environments, capacity estimates still begin with historical database-size trends or system growth reports [5], [6].

These methods are easy to explain, but they become less reliable when customer activity, retention rules, application usage, or operational workloads change.

Time-series methods provide a stronger foundation for this problem. ARIMA and SARIMA are useful when the main pattern is found in the historical series, while SARIMAX allows external business variables to be added [1], [16]. This is important for database forecasting because storage growth is rarely driven by time alone. Measures such as customers, policies, documents, transactions, and applications can explain why growth changes, and they make the forecast easier to connect with business assumptions.

Machine learning models such as Random Forests, XGBoost, LSTM networks, and Temporal Fusion Transformers can capture more complex relationships when sufficient data is available [7], [8]. However, better prediction accuracy does not fully solve the planning problem. Capacity planners also need to understand the uncertainty around the forecast, because acting too late can create service risk, while acting too early can increase cost.

Conformal prediction is relevant here because it can add calibrated prediction intervals around the output of an existing forecasting model [2], [3]. Instead of replacing the model, it uses past forecast errors to estimate a reasonable range for future outcomes. For capacity planning, this changes the forecast from a single expected value into a planning window that can support operational discussion.

AI governance and risk-management guidance also point toward the need for transparent and reviewable model outputs. Frameworks such as the NIST AI Risk Management Framework, OECD AI Principles, the European Union AI Act, and ISO/IEC 23894 emphasise reliability, accountability, and risk-aware use of AI systems [4], [9], [10], [11]. In this context, a forecast is more useful when its uncertainty can be shown clearly and reviewed by both technical and business teams.

Overall, the literature points to three useful elements for enterprise database capacity forecasting: historical growth behaviour, business drivers that help explain that growth, and a clear measure of forecast uncertainty. The approach used in this study follows that path by combining SARIMAX with residual-based split conformal prediction.

This progression is summarised in Figure 1, showing how capacity forecasting moves from simple trend monitoring toward business-aware and uncertainty-aware planning.

Capacity Planning Forecasting Pipeline with Conformal Prediction




Fig -1: Forecasting with business context and uncertainty.

3. BUSINESS CONTEXT AND PROBLEM DEFINITION

3.1 Enterprise Capacity Planning

Storage capacity is a major component of infrastructure cost, and growth planning is key to managing operational costs.

If future growth is not estimated accurately, organisations can face storage shortages, unexpected expansion,

performance issues, and operational risks [5], [6]. For this reason, having a reasonable estimate of future database growth is important for both technical and business planning. This shift is illustrated in Figure 2, where a single expected threshold date is reframed as a more practical planning window.

Enterprise Capacity Planning Forecast




Fig-2: From a fixed date to a planning window.

3.2 Capacity Planning as a Continuous Process

Business drivers change; customer, product, transaction, and document volumes grow; and data retention policies or regulatory requirements evolve over time. These factors directly impact database size and infrastructure consumption. Therefore, capacity planning needs to be continuous, and a six-month-to-one-year capacity view is essential for any enterprise to anticipate future infrastructure requirements and support business growth [5], [6].

Enterprise database storage typically resides on high-performance infrastructure designed to support business-critical workloads. While maximising utilisation is desirable, sufficient free capacity must also be maintained to support operational activities such as backup, recovery, maintenance, and other administrative processes.

Just as financial planners need a buffer for cost uncertainty, capacity planners need a time interval because hardware procurement involves lead time, budget approvals, storage allocation, and change activities that are often tied to business priorities and justification. A planning window is therefore more useful than a single threshold-crossing date, because it lets teams judge decisions within a reasonable range of operational risk.

Because business processes and growth patterns evolve, capacity forecasting should be viewed as an ongoing lifecycle activity rather than a one-time exercise.

3.3 Business Drivers and Database Growth

Business stakeholders typically communicate future growth in terms of customers, contracts, applications, documents, policies, or transactions rather than database objects.

In large enterprise environments containing hundreds or thousands of database tables, obtaining reliable growth estimates for individual database objects is often impractical. Capacity forecasting therefore requires identifying business drivers that exhibit a meaningful relationship with observed database growth.

This approach focuses on business-level metrics that reflect operational activity. These metrics are subsequently evaluated using correlation analysis to determine their relationship with database growth. Drivers demonstrating meaningful explanatory power can then be incorporated into forecasting models as exogenous variables [1], [16].

3.4 Challenges in Capacity Forecasting

Not all business drivers grow at the same rate, so planning and projecting future capacity under varying business-parameter changes is essential for any capacity-planning exercise. This helps align forecasting with business planning processes and enables future growth scenarios to be evaluated using business assumptions rather than relying solely on database-specific estimates.

Capacity forecasting must therefore operate under conditions of uncertainty, in which future growth depends not only on historical behaviour but also on future business decisions and operational events. This creates a need for forecasting approaches that communicate both expected outcomes and the uncertainty associated with those outcomes [2], [3], [8].

4. FROM POINT FORECASTING TO UNCERTAINTY-AWARE CAPACITY PLANNING

This section builds on the progression outlined in Figure 1, where enterprise capacity planning moves from trend monitoring toward business-aware forecasting and finally to uncertainty-aware planning.

4.1 Limitations of Single-Point Predictions

In most enterprises, capacity decisions need enough lead time for procurement, approvals, and implementation planning. Deterministic single-point predictions have limitations because they do not provide a statistical band or interval that can help planners account for uncertainty and variability in future growth [12], [13].

In the real world, the size of a database can never be predicted with complete accuracy, nor can the exact date on which its capacity will become full. This is mainly because business-driver metrics fluctuate over time and influence database growth. As a result, planners often question the forecasted outcome and apply a simple +10% extrapolation or buffer to account for uncertainty.

4.2 Conformal Prediction

Figure 3 shows how conformal prediction adds a calibration layer around the forecast, helping turn a single model output into a usable prediction band.




Fig -3: Calibrating a forecast with conformal prediction.

The errors of the forecasting models are interpreted as a statistical measure of uncertainty. This uncertainty is then represented as a prediction band or prediction interval around the forecasted value [2], [3]. The statistically quantified

uncertainty can be directly transferred into risk metrics and used by planners to support capacity planning and procurement decisions [4], [5], [11].

4.3 Conformal Prediction Variants and Selected Approach

Before choosing the final approach, the different statistical flavours of conformal prediction were reviewed, as shown in Figure 4.




Fig -4: Conformal prediction methods.

Several conformal prediction variants have been proposed in the literature, including split conformal prediction, Jackknife+, cross-conformal prediction, conformalized quantile regression (CQR), and time-series conformal prediction [3], [14]. These approaches differ primarily in how calibration samples are used and how prediction intervals are generated.

Residual-based split conformal prediction was chosen because it is simple to explain, efficient to run, and suitable for enterprise forecasting environments [2], [3], [14]. Unlike Jackknife+ and cross-conformal approaches, the method does not require repeated model retraining. Similarly, it avoids the additional complexity associated with quantile-specific modelling techniques.

The implementation follows a residual calibration workflow. The forecasting model is first trained using historical observations. Predictions generated on a calibration dataset are then compared with actual observations to calculate residual errors. A selected residual quantile corresponding to the desired confidence level is subsequently used as an uncertainty buffer around future forecasts [2], [3].

Prediction interval = Forecast $\pm$ Residual quantile

This approach enables statistically calibrated prediction intervals to be generated without modifying the underlying forecasting model. The implementation presented in this study utilizes the MAPIE conformal prediction framework [15].

The workflow in Figure 5 shows how past forecast errors are reused to create a practical prediction range around the forecast.

Residual-Based Split Conformal Prediction — Method and Workflow

Why we pick the residual-based split variant for enterprise forecasting — and how it plugs into the existing model



Fig -5: Building a prediction range from past errors.

5. EXPERIMENTAL ENVIRONMENT AND ENGINEERING CHALLENGES

5.1 Experimental Environment

The study utilized historical enterprise database capacity measurements collected from operational database environments together with business activity metrics obtained from both application and administrative databases. Business activity was captured through measures such as customers, contracts, policies, applications, documents, and related transactions. For each candidate business driver, both daily activity measures and cumulative growth measures were derived to capture short-term operational changes as well as long-term business expansion trends.

Correlation analysis was performed on the available business metrics, including quotes, clients, documents, applications, and policies. Both daily and cumulative measures were evaluated to identify the business drivers most closely associated with database growth and storage consumption. The selected drivers were subsequently considered as candidate exogenous variables for forecasting.

Figure 6 gives a quick view of how the business drivers relate to database growth and helps narrow down the variables used in the forecast.

Business Driver Correlation Analysis — Selecting Inputs for Capacity Planning

Pearson correlation across candidate drivers — cumulative totals, daily deltas, and size variants




Fig -6: Business-driver correlation view.

The correlation analysis identified four selected business drivers that were used as exogenous variables in the SARIMAX model. Changes in these underlying business metrics determine how the size of the database grows over

time. Business drivers with weak correlation and unwanted noise were excluded to improve the prediction of future database growth.

5.2 Forecasting Framework Selection

A key requirement of the study was the ability to evaluate future database growth under alternative business-growth scenarios. Since business expectations are typically expressed through customers, contracts, applications, and documents, the forecasting framework required support for exogenous variables.

Holt-Winters Exponential Smoothing (HWES) [1] from the Statsmodels library was initially used for forecasting. However, as the database size was already in the terabyte range, the daily growth measured in megabytes and the day-to-day changes in business drivers were relatively small. As a result, the model tended to converge and produced nearly constant forecasts. This made HWES unsuitable for the study, as the rate of change was minimal and the results remained unsatisfactory even after feature scaling.

In addition, the study required the incorporation of business-driver variables to evaluate alternative growth scenarios. The investigation was therefore extended to SARIMAX models, which support exogenous variables and business-driven forecasting [1], [16].

AutoARIMA was subsequently used to identify suitable model parameters, while the final forecasting implementation was performed using SARIMAX to retain access to residual diagnostics required for conformal calibration [17].

5.3 Conformal Calibration Integration

The forecasting framework required statistically calibrated prediction intervals in addition to point forecasts, consistent with the broader need for reliable uncertainty quantification in predictive modelling [8], [12], [13].

Residual errors generated by the SARIMAX model were reused as the calibration dataset for conformal prediction. This enabled uncertainty quantification to be introduced without modification of the underlying forecasting model [2], [3], [15].

```
residuals = abs(y_actual - sarimax_pred)  
q_hat = np.quantile(residuals, 0.80)  
lower_bound = sarimax_forecast - q_hat  
upper_bound = sarimax_forecast + q_hat
```

5.4 Experimental Assumptions and Limitations

This study is intended as a practical capacity-planning framework rather than a complete comparison of all forecasting and conformal prediction methods. The SARIMAX model requires future values of the selected business drivers for the forecast horizon. In this work, those future driver values are treated as planning assumptions derived from recent linear growth trends or business expectations, rather than as independently optimised forecasts. The resulting database-size forecast should therefore be interpreted as conditional on the assumed future paths of these business drivers. If the assumptions change, the forecast and planning window should be recalculated.

The conformal interval used in this paper is based on residual-based split conformal prediction. This approach is simple, transparent, and suitable for an applied enterprise setting, but it does not fully address every challenge of time-series dependence. Residual autocorrelation and changing business patterns may affect interval behaviour over longer forecast horizons. Time-series conformal methods such as adaptive conformal inference or ensemble-based prediction intervals can be evaluated in future work.

The 80% conformal interval should be interpreted as an interval around forecasted capacity values at each forecast horizon. The threshold-crossing window derived from the upper and lower bounds is a useful decision-support range for procurement and planning, but it should not be interpreted as a formal 80% confidence interval for the exact

crossing date.

The validation period available in this study is relatively short. The observed 15-day overlap provides an initial sanity check that recent actual values fall within the conformal band, but it is not sufficient to fully validate long-horizon interval coverage. A stronger evaluation would use multiple rolling-origin backtests and compare empirical coverage against nominal coverage across different forecast horizons.

The study uses an anonymised enterprise dataset and a selected set of correlated business drivers. While this protects sensitive operational information, it also means that the results should be interpreted as evidence for the proposed framework rather than as a universal benchmark. Additional datasets, longer history, and more varied business-growth scenarios would help test the generalisability of the approach.

5.5 Code Availability

The source code and supporting materials used in this study, including model development, conformal prediction calibration, and result generation, are publicly available in the project repository. The dataset provided has been anonymised to protect sensitive information.

GitHub Repository:

https://github.com/sridba503/Capacity-Planning_Conformal-Prediction

6. RESULTS OVERVIEW

The validation results show that the selected SARIMAX model provides a reasonable basis for short- to medium-term database capacity forecasting [1], [16], [17]. More importantly, the conformal calibration step helped translate the forecast into a planning range that is easier to use in business and infrastructure discussions [2], [3].

The final forecast output is shown in Figure 7, where the threshold line, uncertainty band, and planning window can be read together.

Table 1. Summary of model configuration and forecast results

Result item	Observed value / configuration	Interpretation for capacity planning
Data frequency	Daily observations	This keeps the forecast at a day-by-day level, which is useful for tracking capacity changes and planning ahead.
Training period	1 June 2024 to 13 September 2025	This period provided the historical capacity pattern and related business activity used to train the model.
Validation period	From 15 September 2025 onward	This held-out period was used to check how closely the forecast followed the actual database growth.
Target variable	Database size	This is the main capacity measure being forecast, so planners can estimate future storage demand.
Selected exogenous variables	Cumulative business-driver variables: cd_f1, cd_f2, cd_f4, and cd_f5	These business drivers help the forecast reflect real operational growth instead of relying only on a time trend.
Feature scaling	StandardScaler applied to target and exogenous variables	Scaling was used so that large and small variables could be handled more consistently by the model.

Final forecasting model	SARIMAX with order (1,1,1) and seasonal order (0,1,1,7)	The weekly seasonal setting helps capture repeated patterns that appear in daily capacity data.
Model-selection support	AutoARIMA used with weekly seasonality, AIC criterion, and exogenous variables	AutoARIMA helped test model settings, while SARIMAX was kept for the final forecast and error calibration.
MAPE	2.8%	On average, the validation forecast was about 2.8% away from the observed database size, which suggests the point forecast is reasonably close for this planning use case.
Conformal confidence level	80% interval, using $\alpha = 0.2$	Instead of showing only one forecast line, the result gives a calibrated band that reflects expected uncertainty.
Threshold evaluated	6.2 storage units	This is the capacity level at which teams should start treating expansion planning as necessary.
Earliest threshold date	16 March 2026	This gives the earliest date planners may need to be ready, based on the upper conformal bound.
Central forecast threshold date	9 April 2026	This is the date suggested by the central forecast, sitting roughly within the planning window.
Latest threshold date	3 May 2026	This gives the later date in the planning range, based on the lower conformal bound.
Planning window	16 March 2026 to 3 May 2026	This turns the forecast into a usable procurement and capacity-planning window, rather than one fixed date.
Observed overlap check	Recent observed values remained within the conformal interval in the forecast-overlap view	This gives a practical check that the interval is reasonable for planning, although more validation can be added when available.

The validation MAPE reported in Table 1 was generated as part of the SARIMAX validation workflow, and the corresponding code has been made available in the project repository to support reproducibility.

MAPE gives a simple percentage view of forecast error, but it should not be the only measure of model quality. Future work can add RMSE, MASE, and interval-based metrics such as Winkler score to assess both forecast accuracy and uncertainty-band performance more completely [18].




Fig -7: Forecast result snapshot.

Reading the forecast output

In this chart, the capacity forecast generated by the SARIMAX model is shown together with the conformal prediction interval. The blue line represents the central forecast, and the horizontal red line marks 6.2 storage units as the capacity threshold used in this study. The chart focuses on the period around the threshold crossing, which makes the planning window easier to read.

Interpreting the uncertainty bands

Yellow band, conformal 80% interval: the yellow area represents the calibrated uncertainty range around the forecast. It does not claim exact certainty for every future point. Rather, it provides a statistically guided range based on the model's past forecast errors. This is useful for planners because it shows how much variation should be considered when discussing future capacity needs.

Blue vertical band, projected threshold-crossing window: the blue band converts the vertical forecast uncertainty into a practical time range. The left side indicates the earliest date on which the upper bound reaches the 6.2-storage-unit threshold, while the right side indicates the later date on which the lower bound reaches the same threshold. For business use, this date range is more helpful than a single projected crossing date because it reflects procurement lead time, budget approval cycles, and operational planning risk.

Validation against observed values

The green segment shows a 15-day overlap where actual values remain within the conformal band. This is a useful sanity check, not a full validation of long-horizon 80% coverage.

Capacity planning window

Based on the calibrated interval, the upper bound reaches the 6.2-unit storage threshold on March 16, 2026, while the lower bound reaches the threshold on May 3, 2026. This produces a planning window from **March 16, 2026, to May 3, 2026**. From a business perspective, this window gives infrastructure and finance teams a clearer basis for deciding when to initiate procurement, migration planning, or budget approval activities.

Point forecast in relation to the planning window

The red dotted vertical line marks **April 9, 2026**, the date on which the central forecast reaches the threshold. Based on the point forecast alone, planning actions would likely be scheduled around this date.

However, forecasts are inherently uncertain. The conformal prediction interval shows that the threshold could be reached slightly earlier or later. Factoring in this range allows planning to start in advance, reducing the risk of last-minute capacity issues while avoiding unnecessary early investment and the stress that comes from operating too close to capacity limits.

7. FUTURE WORK

Figure 8 shows how the same idea could be extended into a simple agentic AI workflow for testing future business-growth scenarios.




Fig -8: Agentic AI scenario-planning workflow.

This future workflow extends the current forecast result into a more interactive planning process. The earlier forecast result gives a capacity window for the 6.2-storage-unit threshold. In practice, this is useful, but planners may also want to ask follow-up questions before making a decision. For example, they may want to know what happens if customer growth is higher than expected or if the policy count grows faster than the current trend.

This is where the next stage becomes more practical. Business teams often know their expected growth in simple terms, not in database terms. They may say that customer count could grow by 10%, policy count by 20%, or document volume may reduce because of a cleanup activity. A useful next step is to allow these assumptions to be entered directly and then see how the capacity window changes.

The modelling side also needs further testing. SARIMAX worked well enough for this study because it can use business drivers as external inputs [1], [16]. But one model may not suit every situation. A later study can compare SARIMAX with other models such as Random Forest, XGBoost, LSTM, or Temporal Fusion Transformer, especially when growth is uneven or when business behaviour changes suddenly [7], [8].

The uncertainty layer should also be checked with more methods. This paper used residual-based split conformal prediction because it is simple and easier to explain [2], [3]. Future testing can compare it with Jackknife+, cross-conformal prediction, conformalized quantile regression, and time-series conformal methods [14]. The main question is simple: which method gives a stable and useful planning window when the business changes?

This can be read as a possible working design. The planner does not need to operate every model manually. The planner can ask a question, and the system can decide which step is needed: update the input assumptions, run the forecast, check the uncertainty band, or monitor the threshold date. This keeps the interaction simple for the user, while the model work happens in the background.

A small proof of concept can start with one baseline and one changed scenario. The baseline can use the current business-driver trend. The changed scenario can use 10% customer growth and 20% policy growth, while keeping the other drivers unchanged. The result can then be compared with the capacity window from Figure 7. This would show whether the changed business assumptions move the threshold date earlier or later.

The next paper can build this as a simple chatbot-based prototype. A planner could ask, “What happens if customer growth increases by 10%?” The agentic AI flow shown in Figure 8 can then take that input, update the scenario, run the forecast, apply the conformal interval, and return the new capacity window. This would make the approach easier for business users to try, while still keeping the forecasting process clear enough to test and validate.

Although the focus here is on RDBMS capacity, it is worth noting that similar planning challenges are starting to emerge in AI-driven systems, such as LLM repositories and vector databases. These environments introduce different growth patterns, and extending this kind of uncertainty-aware view to them would be an interesting direction to explore.

8. CONCLUSIONS

The central idea of this paper is simple: database planners need more than one expected capacity date. Using SARIMAX with selected business drivers and residual-based split conformal prediction, the study showed how a 6.2-unit threshold can be converted into a practical window from March 16, 2026, to May 3, 2026, with the central forecast falling on April 9, 2026.

This range is the main outcome of the work. It gives infrastructure and business teams a better way to discuss procurement timing, budget approval, and capacity risk because it accepts that database growth cannot be predicted as one exact date.

As discussed in Section 7, the next step is being explored as a small MVP or proof of concept, where planners can use a chatbot-style interface to change business assumptions and see how the conformal capacity window moves. This can form the basis for a continuing paper on agentic AI-supported capacity planning.

REFERENCES

- [1] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed. Melbourne, Australia: OTexts, 2021. [Online]. Available: <https://otexts.com/fpp3/>. Accessed: Jun. 11, 2026.
- [2] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. New York, NY, USA: Springer, 2005.
- [3] A. N. Angelopoulos and S. Bates, “A gentle introduction to conformal prediction and distribution-free uncertainty

quantification,” arXiv:2107.07511, 2021. [Online]. Available: <https://arxiv.org/abs/2107.07511>. Accessed: Jun. 11, 2026.

[4] E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1. Gaithersburg, MD, USA: National Institute of Standards and Technology, Jan. 2023. [Online]. Available: <https://doi.org/10.6028/NIST.AI.100-1>. Accessed: Jun. 11, 2026.

[5] IBM Corporation, “Capacity management and capacity planning concepts,” *IBM Documentation*. [Online]. Available: <https://www.ibm.com/docs/>. Accessed: Jun. 11, 2026.

[6] Amazon Web Services, *AWS Well-Architected Framework*. Seattle, WA, USA: Amazon Web Services, Inc. [Online]. Available: <https://docs.aws.amazon.com/wellarchitected/>. Accessed: Jun. 11, 2026.

[7] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, “Temporal Fusion Transformers for interpretable multi-horizon time series forecasting,” *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, Oct.–Dec. 2021.

[8] J. Gawlikowski *et al.*, “A survey of uncertainty in deep neural networks,” *Artificial Intelligence Review*, vol. 56, pp. 1513–1589, 2023.

[9] European Parliament and Council of the European Union, *Regulation on Artificial Intelligence: European Union Artificial Intelligence Act*. Brussels, Belgium: European Union, 2024. [Online]. Available: <https://artificialintelligenceact.eu/>. Accessed: Jun. 11, 2026.

[10] Organisation for Economic Co-operation and Development, *OECD AI Principles*. Paris, France: OECD. [Online]. Available: <https://oecd.ai/en/ai-principles>. Accessed: Jun. 11, 2026.

[11] ISO/IEC, *ISO/IEC 23894:2023 Artificial Intelligence — Guidance on Risk Management*. Geneva, Switzerland: International Organization for Standardization, 2023. [Online]. Available: <https://www.iso.org/standard/77304.html>. Accessed: Jun. 11, 2026.

[12] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. New York, NY, USA: Springer, 2009.

[13] G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, *An Introduction to Statistical Learning: With Applications in Python*. Cham, Switzerland: Springer, 2023.

[14] A. N. Angelopoulos and S. Bates, “Conformal prediction tutorial.” [Online]. Available: <https://people.eecs.berkeley.edu/~angelopoulos/blog/posts/gentle-intro/>. Accessed: Jun. 11, 2026.

[15] MAPIE Developers, “MAPIE documentation.” [Online]. Available: <https://mapie.readthedocs.io/>. Accessed: Jun. 11, 2026.

[16] Statsmodels Developers, “statsmodels.tsa.statespace.sarimax,” *Statsmodels Documentation*. [Online]. Available: <https://www.statsmodels.org/stable/generated/statsmodels.tsa.statespace.sarimax.SARIMAX.html>. Accessed: Jun. 11, 2026.

[17] pmdarima Developers, “pmdarima.arima.auto_arima,” *pmdarima Documentation*. [Online]. Available: https://alkaline-ml.com/pmdarima/modules/generated/pmdarima.arima.auto_arima.html. Accessed: Jun. 11, 2026.

[18] R. J. Hyndman and A. B. Koehler, “Another look at measures of forecast accuracy,” *International Journal of Forecasting*, vol. 22, no. 4, pp. 679–688, 2006, doi: 10.1016/j.ijforecast.2006.03.001.