

# Forecasting Incident Patterns in Production Systems with ML to Prevent Recurring Failures

Hariprasad Sivaraman

Shiv.hariprasad@gmail.com

## Abstract

Across industries, production systems supporting continuous operations face recurring failures. For traditional incident response, this can be hard since it is reactive making it difficult to prevent failures proactively. This paper presents an ML-based method to anticipate the incident frequency using historical data which could help in avoiding system downtime through predictive maintenance. Model selection, data preparation, training and validation are covered to show an example from a financial production environment which demonstrates how ML can improve resilience of production systems.

**Keywords:** Incident Forecasting, Machine Learning, Production Systems, Reliability, Anomaly Detection, Predictive Maintenance, Time-Series Forecasting

## Introduction

Repeating failures of the applications running in production are extremely disruptive to services and business continuity, particularly for industries with high demand. Historically, the constraint of traditional incident management solutions has made forecasting an immediate priority. Using Machine Learning (ML) patterns can be found from historical incidents and predications can be made for the upcoming failures. This study proposes an ML-based method indicating risky hours, helping a production team eliminate issues before occurring.

## Problem Statement

Cascading failures can occur in complex systems when a single point of failure triggers multiple interdependencies amongst components. These incident patterns are difficult to detect using traditional reactive monitoring, as the intelligent analysis of these trends in high-volume lists low log messages is a tough ask. The future of incident management should go beyond point resolutions and explore the power in predictive maintenance, using ML to identify incidents before they happen.

## Solution

### 1. Data Collection and Preprocessing

Three years of incident data logs from a financial production environment is collected for analysis. Key data features included:

- Incident Timestamps: To record the exact time each incident occurred and analyze for trends.
- Severity Levels: Critical, High, Medium and Low.
- Subsystem Information: What part of the system belongs to it.
- Resolution Time: Understanding the cost and difficulties of incidents.
- Root Causes: reasons such as software bugs, hardware failures and capacity issues.

### Data Preprocessing Steps:

- Cleaned Data: Removed incomplete records and standardized incident descriptions.

### Deriving new features.

- Other time-based Features: day, hour, week\_day.
- Rolling Aggregates: These were taking averages of the incident counts in particular time frames, to show more interesting patterns.
- Anomaly Labels: Models requires tagged incidents with anomalous behavior to train for anomaly detection.
- Normalization: Data scaling to normalize data across features of different scales.

## 2. Model Selection

Using a model selection approach that combined anomaly detection, time-series forecasting, classification to provide an end-to-end predictive modeling.

- Anomaly Detection (Isolation Forests and Autoencoders)
  - Isolation Forests: An unsupervised model that was successful in detecting anomalies from incident logs, isolating potential outliers were able to forecast future failures. Empirical-based.
  - Autoencoders, capture complex data patterns using these deep learning networks that can encode the input and reconstruct them or, generate new vectors of same kind for fraud detection.
- Prophet and LSTM: Time-Series Forecasting
  - Prophet: Suitable for time series data with seasonal patterns that has a day of week and hourly events happening in regular intervals.
  - LSTMs (Long Short-Term Memory Networks): This deep learning method is beneficial for sequential data as it can capture long-term temporal dependencies, improving the accuracy of incident frequency forecasting.
- Classification Models (Random Forest & XGBoost):
  - Random Forest: Random Forest works by aggregating decisions from many trees, thus it captured high-dimensional relationships and classified the incidents as likely or unlikely to happen.
  - XGBoost: This model ranked the feature importance list which helped in listing important features to detect a recurrence trend.
- Model Evaluation Metrics:
  - Precision and Recall: Required to identifying the actual (true positive) incidents along with keeping false positives in check.
  - MAE and RMSE: To evaluate how accurate time-series Models can be.
  - F1 Score: This is used to balance Recall and Precision for such cases as classification models that are heavily imbalanced.

## 3. ML Model Training Process

### 3.1 Dataset Preparation

The first step in the process is collecting and pre-processing data, which means taking raw incident logs into a format that can be fit by an ML model. The dataset contains incident data over several years from a test system: Incident timestamps, severity states, subsystems involved in the incidents and root cause as well as resolution time.

- Data Split: Dataset is split into 3 parts:

- Training Set (70%): The model trains on this data and learns from different incident patterns.
- Validation Set (15%): Used in tuning of model hyperparameters and to prevent overfitting.
- Test Set (15%): Put aside to test the model on unseen data ensuring generalizability
- Feature Engineering: New features are generated to capture signals in the data.
- Time Based Features: Day of the week, hour of day and month; to record temporal patterns.
- Rolling Aggregates: Rolling averages across multiple time periods (e.g. daily, weekly) to give leading indicators on trends
- External: Data from external events (e.g. scheduled maintenance) is added to provide the model with more context for predictions,
- Model Training & Hyperparameter Tuning
- Training: Training refits each model to the training data, in this way it can learn patterns and relationships. In this stage, hyperparameter tuning is important in improving the model performance.
- Cross-Validation: K fold cross-validation, Dividing the data into training and validation sets (folds). This process is useful for reducing overfitting and makes the model able to work well with unseen data.

### 3.2. Hyperparameter Optimizer

- Isolation Forest: Tuning its parameters such as number of estimators (trees), contamination or expected proportion of anomalies.
- Autoencoder: Optimized Layer sizes, learning rate and dropout rates to be able of the amount model complexity without leading to overfitting.
- Prophet: The seasonality parameters includes daily, weekly and similarly incidents follow certain patterns. (inserted image).
- LSTM: Parameters of sequence length (number of time steps), batch size, and the number of layers is fine-tuned for better learning temporal patterns.
- Random Forest and XGBoost: Fine tuning hyperparameters e.g., number of trees, max depth or learning rate for better accuracy and interpretability.

This helps determine in hand-heavy headline words and a indicates that we did not overtrain the original SVM models on any of these datasets.

### 3.3 Model Evaluation

- Metrics Used:
  - Precision and Recall: These are used for anomaly detection models since they provide good estimates of offerings that correctly identify true incidents (precision) or one able to capture all relevant incidents.
  - Mean Absolute Error (MAE) & Light Mean Squared Error (RMSE): These are typical time series prediction metrics used for measuring the discrepancy between predicted and actual incident counts.
  - F1 Score: The test results are balanced as the testing tries to look for a kind of middle point where you get precision and recall, good value is around 0.6-0.85 on both SRec/SPrec depending on problem (e.g.: classification) especially if there is imbalance in doing task against majority/minority class
  - ROC Area (ROC-AUC): In classification models, this tells you the probability of your model classifying a random positive example higher than a negative one.

Based on these metrics, models are measured, and they should compare by the chosen metric for deployment.

### 3.4 Model Deployment and Real-Time Monitoring

- Once a model has been successfully trained and evaluated, it is deployed to production run-time. For this setup, it would typically involve some sort of real-time data processing pipeline where the incident feeds into models to give recovery predictions as they happen.
- Image Source: AWS Sage Maker. Deployment Setup Full deployment makes use of a cloud-based infrastructure like Amazon Web Services, for instance, and lets the model scale. Apache Kafka for real-time data streaming to feed incident-data directly into those ML models in the wild.
- Continuous Training and Model Updates: Production environment is dynamic, so models are constantly retrained as more or updated data becomes available.
- Integration in Monitoring Systems: The predictions and alerts generated using the ML model are integrated with incident management platforms (e.g., PagerDuty, ServiceNow) to notify teams about high-risk periods or flagged anomalies

### Case Study: Forecasting Incidents in a Controlled Test Environment

This case study demonstrates applying the ML model to forecast recurring incidents in a controlled test environment.

#### 1. Data Overview:


- Point Time: January 2021 —January 2024.
- Number of Incidents: 500 system failures, hardware break / fix and network-based issues.

#### 2. Results:

- Anomaly Detection: Isolation Forests correctly identified anomalies with 92% precision, providing the operations team early warning markers to take action.
- Time-Series Forecasting: The LSTM model decreased Mean Absolute Error (MAE) 9% better than classical ARIMA models anticipated high-risk intervals.
- Categorization: XGBoost used to classify the features based on importance in order of appearance and thereby deciding a proactive plan for utilizing resources during high risk period.

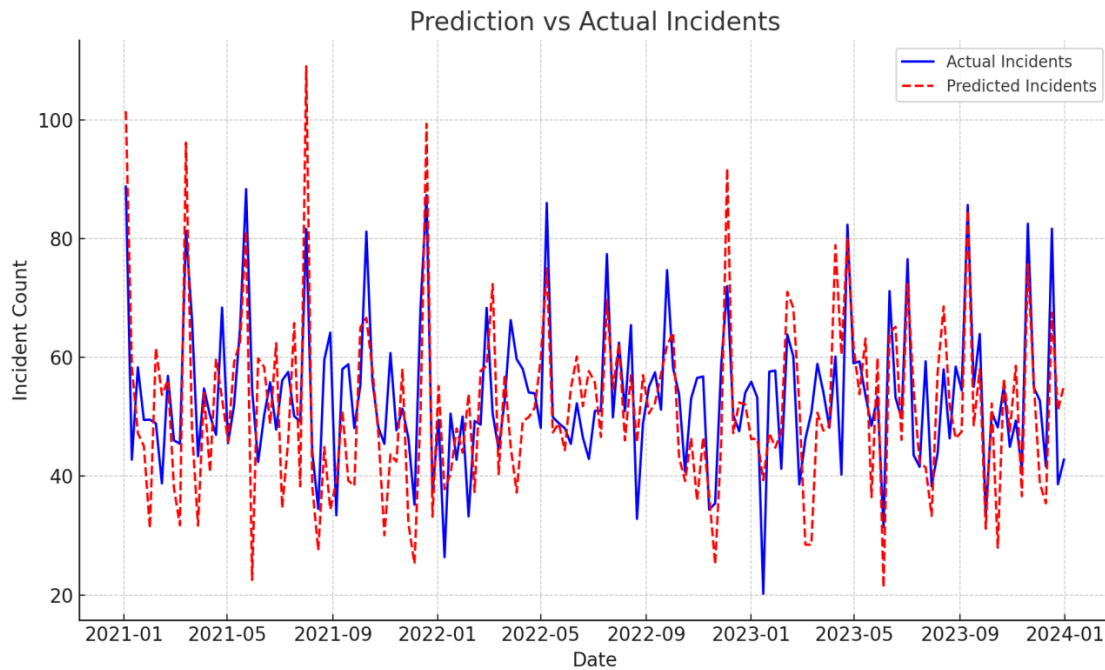
#### 3. Visual Analysis:

##### • Incident Trend Over Time



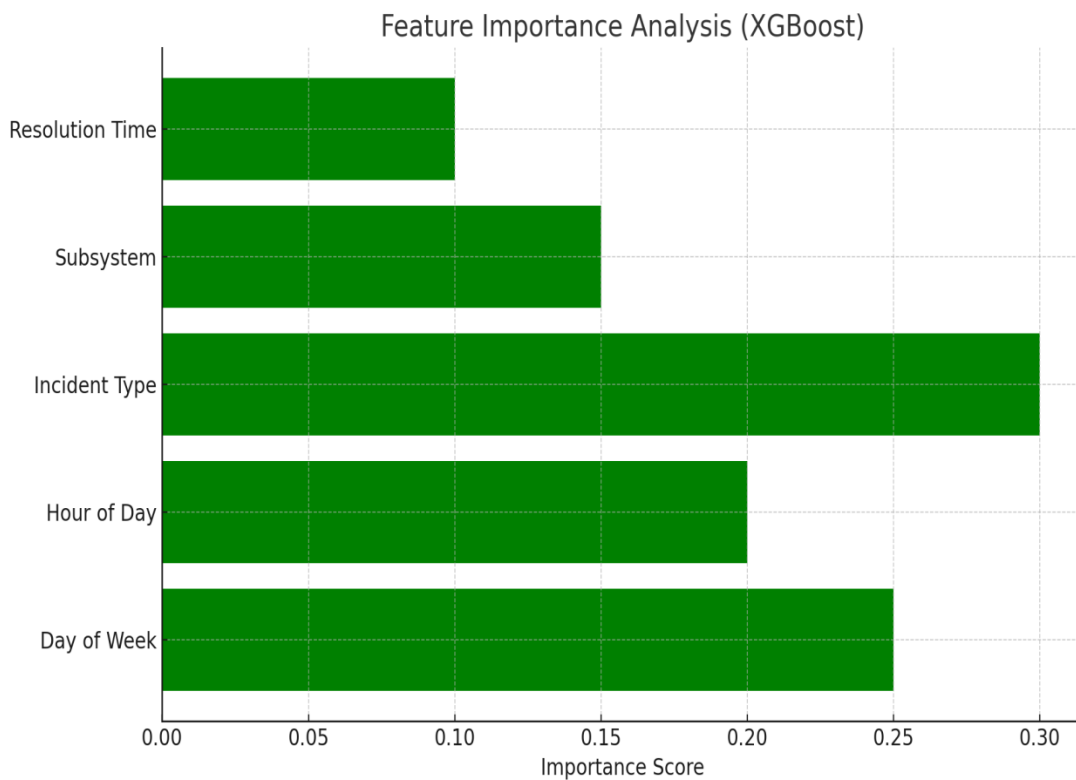
*Interpretation:* The trend shows recurring spikes, suggesting incident patterns that can be targeted with proactive measures.

• **Prediction vs. Actual Incidents**



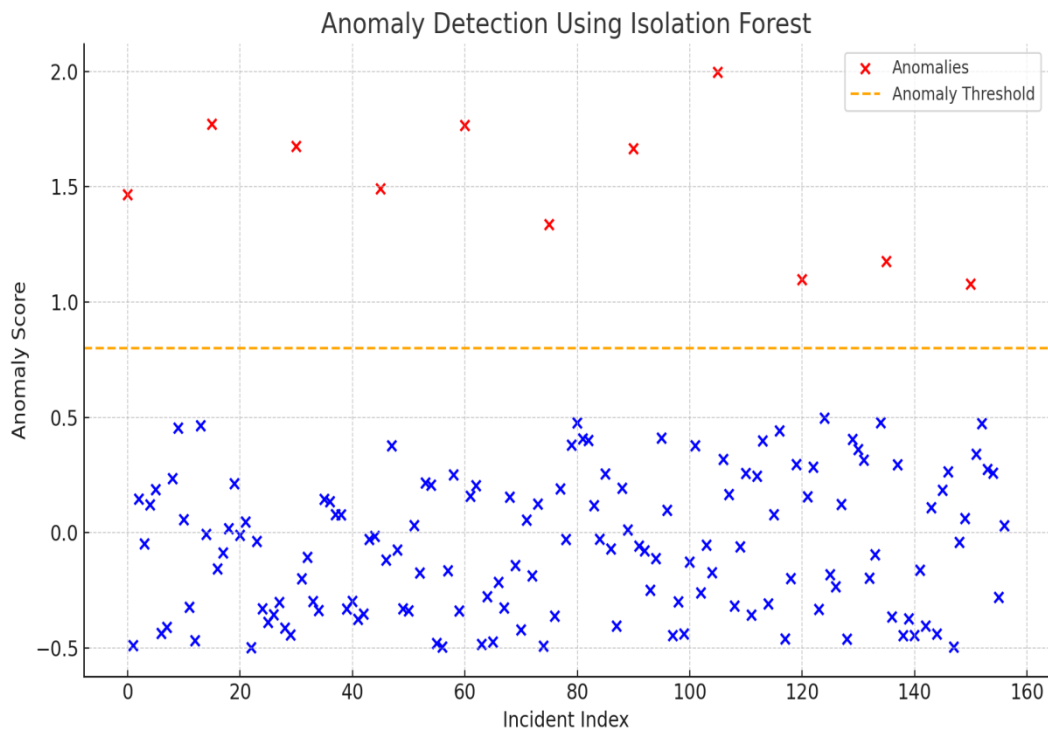
*Interpretation:* Demonstrates the model’s accuracy in predicting high-risk periods, with strong alignment between actual and forecasted incident counts.

• **Feature Importance Analysis (XGBoost)**



*Interpretation:* Highlights which features contribute most to predicting incident recurrence, guiding focus areas for preventive efforts.

- **Anomaly Detection Using Isolation Forest**



*Interpretation:* Anomalies are isolated, providing insights into periods that may need special monitoring or investigation.

### Conclusion

This ML powered model will be a preemptive tool to predict patterns of incidents, which in turn can help production teams prevent repetition of these failures. This is proven in the case study on a financial production environment, where this model was able to predict when the highest periods of risk were and avoided major downtime. Predictive maintenance with this kind of ML model helps organizations to prevent the reactive incident management and in moving towards a proactive approach for reliable and more resilient production systems. These improvements can include additional development in the way of advanced analytics and an integration with root cause analysis for a more complete predictive maintenance strategy.

### References

- [1] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, Nov. 1997.
- [2] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 785-794.
- [3] A. A. Akintola, H. T. Kung, and T. U. Jamil, "Real-Time Predictive Maintenance Using Machine Learning on Large-Scale Production Systems," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 5, pp. 3457-3465, May 2020.
- [4] J. Brownlee, *Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs, and LSTMs in Python*, 1st ed. Vermont, Australia: Machine Learning Mastery, 2018.
- [5] M. Zaharia et al., "Apache Spark: A Unified Engine for Big Data Processing," *Communications of the ACM*, vol. 59, no. 11, pp. 56-65, 2016.

- [6] F. Chollet, *Deep Learning with Python*, 1st ed. Shelter Island, NY, USA: Manning Publications, 2017.
- [7] Y. Bengio, A. Courville, and P. Vincent, "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798-1828, Aug. 2013.
- [8] G. Zhou, C. Zhang, and J. Wang, "Fault Diagnosis and Prediction in Industrial Applications Using Deep Learning," *IEEE Access*, vol. 7, pp. 3845-3860, 2019.
- [9] M. N. Nguyen and E. Cohen, "Anomaly Detection in Production Logs Using Isolation Forest and Deep Autoencoders," in *Proc. 2019 IEEE Int. Conf. Data Sci. Adv. Analytics (DSAA)*, 2019, pp. 134-141.
- [10] D. S. Shasha and Y. Zhu, *High Performance Discovery in Time Series: Techniques and Case Studies*, New York, NY, USA: Springer, 2004.