



# Build AI That Works in Production

Generate realistic synthetic operational and time-series data to train, test, and validate AI before deployment.

Most teams building AI on operational and time-series data have plenty of examples of normal behavior — but not enough of the failures, degradations, anomalies, and edge cases that matter most in production. Rockfish generates realistic synthetic data from your existing data or schema, helping teams train, test, and validate AI systems with better coverage before deployment.

## HOW IT WORKS

|  |   |  |
|--|---|--|
| <p><b>01</b></p> <p><b>Connect to your data or schema</b></p> <p>Start with existing datasets or simply a schema definition. When data is available, Rockfish learns its structure, patterns, timing, and correlations. When it is not, Rockfish generates directly from the underlying system design.</p> | <p><b>02</b></p> <p><b>Generate realistic synthetic scenarios</b></p> <p>Create synthetic datasets that reflect the behavior of your environment — including rare failures, anomalies, and edge cases that historical data often misses. Output is labeled and ready for training, testing, and evaluation.</p> | <p><b>03</b></p> <p><b>Train, test, and validate AI faster</b></p> <p>Use generated data to train ML models, evaluate AI agents, test pipelines, and validate systems safely — without exposing production environments or waiting on hard-to-access data.</p> |
|--|---|--|

## WHERE TEAMS GET STUCK

|   |  |   |
|---|--|---|
| <p><b>Validating AI before production</b></p> <p>Rare outages, degradations, and edge cases are often missing from historical data, making it hard to test models before they are trusted in live decision loops.</p> | <p><b>Evaluating AI agents on real operational scenarios</b></p> <p>Generic benchmarks do not reflect production. Agents that perform well on standard tests often fail on stateful, incident-driven tasks in enterprise environments.</p> | <p><b>Accelerating AI development without increasing risk</b></p> <p>Production data is hard to access, sensitive to share, and slow to approve, which delays testing, iteration, and deployment.</p> |
|---|--|---|

## BUSINESS IMPACT

|   |  |  |
|---|--|--|
| <p><b>Shorten development cycles</b></p> <p>Give teams faster access to realistic test data without waiting on production access or sensitive data approvals.</p> | <p><b>Reduce production risk</b></p> <p>Catch failures, anomalies, and edge cases in testing before they affect live systems or end users.</p> | <p><b>Increase confidence in production readiness</b></p> <p>Validate models and agents on broader, more realistic scenarios than historical data alone can provide.</p> |
|---|--|--|

## WHY ROCKFISH

|  |  |  |
|--|--|--|
| <p><b>Built for your environment</b></p> <p>Generate synthetic data from your existing data or schema, so testing reflects your systems, patterns, and operating conditions.</p> | <p><b>Expands coverage beyond historical data</b></p> <p>Simulate rare failures, cascading events, cold starts, and other edge cases that real-world data often lacks.</p> | <p><b>Privacy-safe by design</b></p> <p>Support secure deployment across SaaS, on-prem, hybrid, and GovCloud environments without requiring sensitive production data to leave your control.</p> |
|--|--|--|

### About Rockfish

Founded in 2022 and built on research from Carnegie Mellon University, Rockfish helps organizations build safer, more reliable AI systems on operational data.

**Customers:** Conviva, Ford, MantisGrid, US Army **Partners:** Snowflake, Databricks **Backed by:** Emergent Ventures, Foster Ventures, TEN13

Let's schedule a conversation to discuss how Rockfish could support your AI testing and validation efforts.

rockfish Data · rockfish.ai · (650) 762 - 5001 · sales@rockfish.ai