From Crash Data to Action: An Agentic AI Framework for Segment-Level Crash Frequency Prediction

Md Tufajjal Hossain, Ph.D. Candidate, Joyoung Lee, Ph.D., Dejan Besenski, Ph.D., Lazar Spasovic, Ph.D.

John A. Reif, Jr. Department of Civil & Environmental Engineering, New Jersey Institute of Technology

mh684@njit.edu, jo.y.lee@njit.edu, besenski@njit.edu, spasovic@njit.edu

Predict → explain → verify → act: an agentic AI that forecasts monthly crashes on every NJ Interstate segment, tells the engineer why, checks its own text, and learns from feedback.

1 Why X-SAFE?

- **Problem.** Finding and fixing high-crash road segments still takes many separate manual steps.
- **Gap 1: no reasons.** Accurate machine-learning models do not say why a segment is risky.
- **Gap 2: no proof.** AI chatbots write clearly, but can state numbers the data do not support.
- **Gap 3: no link** from a crash forecast to a checked diagnosis and a proven countermeasure.

What X-SAFE adds

1. **One pipeline** from crash records to a countermeasure
2. **The reasons** behind every forecast, variable by variable
3. **A Verifier** that checks every number the AI writes
4. **Engineer feedback** reused as lessons, no retraining
5. **Ready for new data:** AI suggests, analyst confirms

2 Data

862	137	169k	47%
Interstate segments (1 mile, by direction)	months, 2010 – May 2021	crashes matched to segments	segment-months with no crash

- **3.29 million** NJ police crash reports, 2010–2022.
- **Strict matching:** a crash is kept only if its location, route and travel direction all match a segment.
- **61 factors** in 8 groups: road design, traffic, crash history, crash types, season, driver behavior, weather, facility.

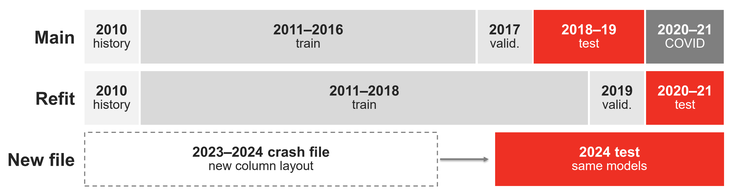


Fig. 1. Three tests on months the models never saw.

3 Models compared

HA	SPF-NB	LightGBM	CGA-NB
12-month crash average simple baseline	**Highway Safety Manual** road + traffic only	**boosted decision trees** accurate, black box	**proposed model** accurate and self-explaining

$$\ln \mu = \text{offset} + \beta_0 + \Sigma_g f_g(x_g), \quad y \sim NB(\mu, \alpha)$$

- **CGA-NB:** one small neural network per factor group; their effects are added, so each group's share of the forecast is exact.
- **Why CGA-NB:** within 1% of the most accurate model (LightGBM) on validation data, and it explains itself.
- **Reasons as factors:** ×1.45 means 45% more expected crashes than a typical segment-month (SHAP).

4 X-SAFE agentic framework

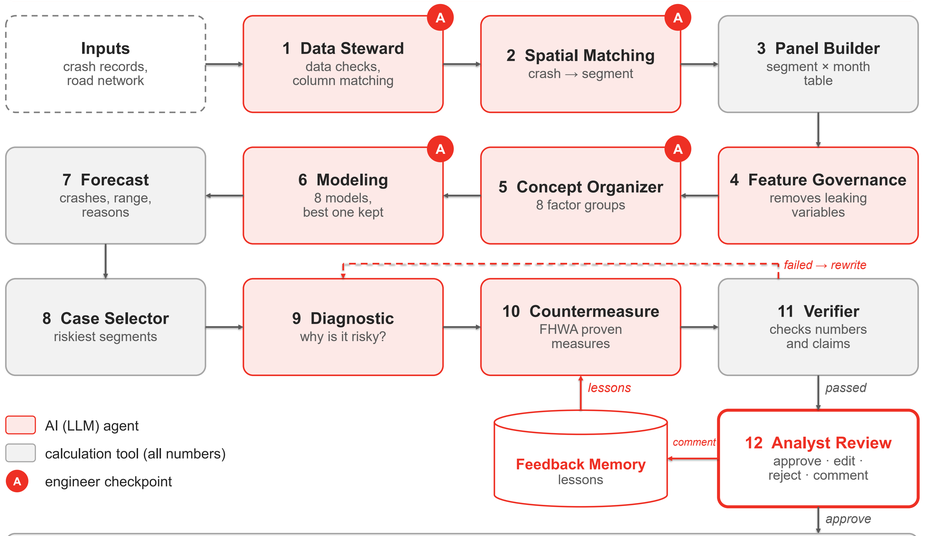


Fig. 2. The AI writes; calculation tools compute and check every number.

5 Example: I-280 WB, MP 14.1–15.1, Dec 2019

13.2	3–27	Top 1%	Very high
predicted crashes	90% range	of all segments	risk class

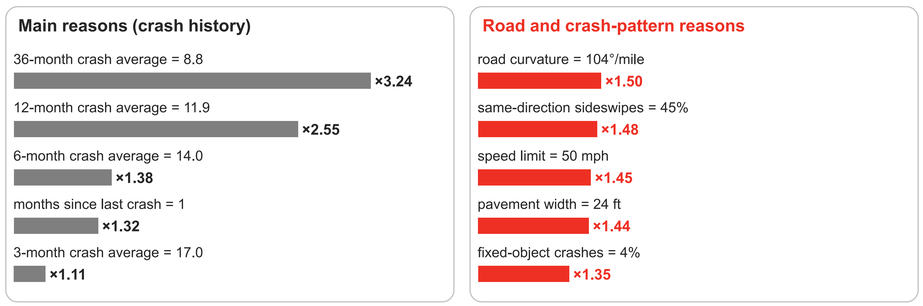


Fig. 3. How much each factor raises the expected crashes at this segment (SHAP).



Learning without retraining: the lesson goes into the AI's instructions; the forecast and its reasons do not change.

6 Results

0.96	0.98	93%	98%
average error, crashes per segment-month	ranking agreement (Spearman ρ)	of the top-5% hotspots found	of outcomes inside the 90% range

CGA-NB, test years 2018–2019, 862 segments

As accurate as LightGBM on 2018–2019 and the most accurate model in every later test, while explaining each forecast.

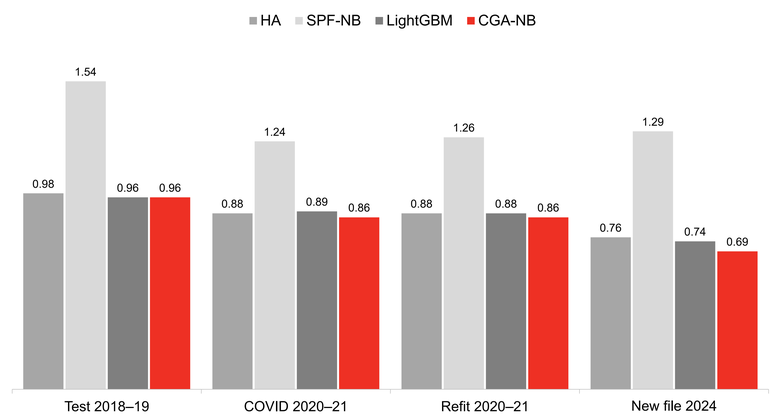


Fig. 4. Average error (MAE) by model and test; lower is better.

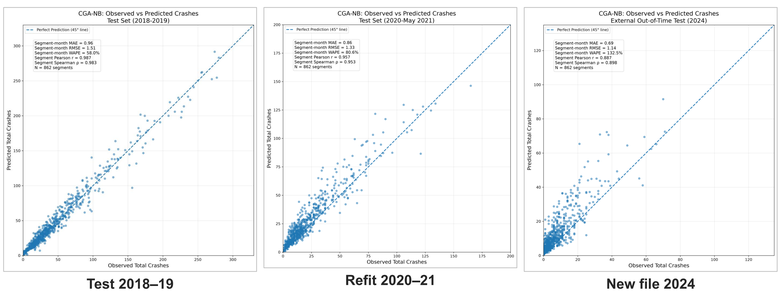


Fig. 5. CGA-NB observed vs. predicted crashes per segment.

- **COVID years:** ranking holds, but hotspots found drop from 93% to 81%; re-estimate after big changes.
- **New 2024 file:** ranking holds (ρ ≈ 0.90), but all models over-predict, as late-2024 reports are still incomplete; recalibrate first.

7 Take-home messages

- **Accurate and explained:** no trade-off between accuracy and reasons.
- **Trustworthy AI text:** every number is checked before an engineer sees it.
- **Next:** recover more crashes, add traffic and weather data, expert review, NJ state highways.