ZeroTrain.Ai

Knowledge-Based Inference for Real-Time Decisioning

The Rose That Became a Tulip

An Engineering Observation on Decision-State Preservation

Prepared For:

Enterprise Clients, Architects, Regulators, Engineering Teams

By Leonard Gambrell

Abstract

While validating deterministic model distillation for ZeroTrain.ai, a deliberately simple machine learning model was created solely as a verification artifact. The objective was to create a five-flower classifier using four attributes and determine whether every learned decision state could be distilled into an equivalent deterministic representation. During testing, the attributes describing a Rose unexpectedly produced a Tulip. Further analysis suggested that the expected five decision states had effectively collapsed into three distinguishable states despite multiple training attempts and configuration changes. This paper documents the observation and discusses why decision-state preservation deserves consideration independently from predictive accuracy.

1. Why This Experiment Was Conducted

ZeroTrain.ai already supported deterministic decision models. During development, an additional capability was explored: converting a probabilistic machine learning model into a deterministic ONNX representation. To validate the process, an intentionally simple flower classifier was constructed because every expected state could be verified manually.

2. Experimental Design

Dataset:

- 5 flower classes
- 4 input attributes

One expected output for every flower

The purpose was not to benchmark machine learning performance. The purpose was to verify that every known decision state survived training and could be reproduced deterministically.

3. Observation

Testing revealed an unexpected result. Entering the attributes representing a Rose produced a Tulip prediction. Training epochs were increased, reduced, and alternative configurations were evaluated. The behavior persisted. When the trained model was distilled for analysis, the expected five explicit states appeared to have become only three distinguishable states.

4. Engineering Discussion

This observation does not imply that every machine learning model loses decision states. It does illustrate that optimizing statistical loss and preserving explicit decision states are different engineering objectives. The experiment raised an important question: if engineers cannot independently verify complete state preservation in a five-state model, how can they confidently verify models containing thousands of business rules?

5. ZeroTrain.ai Perspective

The purpose of ZeroTrain.ai is different from traditional machine learning. Rather than learning statistical decision boundaries, ZeroTrain.ai represents explicit knowledge directly.

Key characteristics include:

- No training required
- Deterministic execution
- Explicit decision states
- Explainable outcomes
- ONNX export
- Rule-to-model generation
- Decision traceability

Because the decision space is explicitly represented, every defined state can be inspected, validated, certified, and reproduced without relying on statistical approximation.

Questions Raised

- Can every predefined state be independently verified after training?
- Can statistically similar states become merged?
- Can omitted states be detected without exhaustive testing?
- How many decision states could disappear before anyone notices?

Conclusion

The flower classifier was never intended to evaluate machine learning. It existed only to validate deterministic distillation. Instead, it exposed an unexpected engineering characteristic: predictive capability and decision-state preservation are not necessarily the same objective. For applications where every decision matters, preserving explicit decision states may be just as important as achieving high predictive accuracy.