



Systems and Methods for Replay-Verifiable Control of Stochastic Machine-Learning Outputs Using Frozen Evidence Manifests and Bounded Signal Transformation

Principal Investigator

Prof H Yan Chair Professor of Management Sciences and Director of MSc in Accounting and Finance with AI and Fintech Applications, City University of Hong Kong and Beijing National Accounting Institute
Director, AIFT

Dr D Wu Head of Research, AIFT

Registered Patent Region	Patent/ Filing No.	Description	Issued/ Filing Date	Status
US	64/072,132	Systems and Methods for Replay-Verifiable Control of Stochastic Machine-Learning Outputs Using Frozen Evidence Manifests and Bounded Signal Transformation	22 May 2026	Pending

Stage of Technology Transfer: Commercially Viable Technology/ Product Ready for Sale
Research Area: Risk-Bounded AI Control Systems and Replay-Verifiable Machine Learning Integration

Abstract

A computer-implemented system controls use of stochastic machine-learning language-model output in a downstream automated execution engine by converting unstructured digital-document evidence into a replayable, schema-validated, bounded control signal. Document records retrieved from dynamic network-accessible sources are canonicalized into a frozen evidence manifest with an immutable run identifier and manifest content hash before any model invocation. A schema-constrained forecast request is assembled from the manifest and a stored baseline prior distribution, and a machine-learning language model returns a probability vector over predetermined outcome bins. The response is validated against machine-readable schema constraints. Distribution-distance metrics are computed between the vector and the baseline prior distribution; when the metrics indicate an immaterial shift, the baseline prior replaces the vector. A scalar value is computed from the accepted distribution and transformed into a bounded control adjustment according to stored policy parameters. A replay manifest stores artifact identifiers and content hashes. Replay verification reloads artifacts, verifies hashes, recomputes the adjustment, and blocks or flags downstream automated use when recomputed values differ from stored values beyond tolerance. In one embodiment, the bounded control signal is joined to an expected-value block of a structural limit engine.