

EXECUTIVE BIO 4/2026

FOR AI RELIABILITY/DECISION CONSULTING / SHOWING SENIOR POSITIONING IN SUBJECT

Ronald Prescott Loui, Ph.D. *CS/CogSci at U of Rochester*, born 1961 a U.S. citizen, was trained in *decision and control at Harvard*. What began by auditing a class, taking it, and assisting teaching it, led to a *summa cum laude* senior thesis. This won the ACM Award for best undergraduate paper, published in the flagship journal of computing. The title was "Optimal Paths in Graphs with Stochastic or Multidimensional Weights" and it has a long history of citation. The fields involved are decision science, operations research, optimization, management science, and automation of data-driven decision-making: this is the space currently occupied by AI/ML neural net classifiers in applied data science.

This work was immediately relevant to network packet routing, robot path planning, and later was the basis for a patent and doctoral thesis at *MIT*. It models the decision problem more realistically than traditional path planning; it accepts that times are random variables and that many factors intrude on optimality.

His interests continued to revolve around automated rational decision for decades in academia and consulting, after a *Stanford postdoc* under Amos Tversky's famous umbrella. Internationally recognized academic advances in inductive inference and decision resulted: robustness and meta-appraisal are prime concerns; process and confidence are taken more seriously than by prior generations.

Loui's next work was on interval-valued decisions. This built on the criteria of Nobel Prizewinning Economist Hurwicz, adding confidence and process (following on the ideas of Nobel and Turing Award winner Simon). Like the first work, this is about reliability and robustness: careful appraisal, not expedience. Works appeared in the journal *THEORY AND DECISION* and in successive conferences on *UNCERTAINTY IN AI*, *MACHINE INTELLIGENCE AND PATTERN RECOGNITION*, and *AAAI SPRING SYMPOSIUM* (decision automation track at Stanford).

The doctoral thesis, "Theory and Computation of Uncertain Inference and Decision" was completed as an AI/Cognitive Science work under the massively respected scholar of probability: Henry Kyburg, an engineering mind in philosophy of science. Kyburg and Loui were motivated by actuaries solving real insurance problems. This was "direct inference" from data, years before "knowledge discovery and data mining" became AI+Data Science. Thesis work was sponsored by the *U.S. Army Signals Warfare Laboratory* as "data fusion" and "Uncertain inference in AI". The use of interval-valued probabilities is intended to provide robustness, reliability, and resilience. These same themes pervade his course on *Operating Systems* which he has taught for the past 13 semesters at *Case Institute of Tech.* (*CWRU Cleveland*, where he also teaches *Linux Scripting*, *User Experience*, and *AI Seminar/Current Issues*; also, he is on the campus *AI Taskforce* and advises the *Linux Club*.)

His most recent research at *Case* continues to put robustness, reliability, and resilience at the forefront, asking those who automate decision to exercise epistemic humility:

One work is about avoiding snapshots in cyberphysical systems that trigger emergency automated responses: e.g., preventing *Boeing MCAS* kinds of overreactions when sensors fail.

Another work, being unveiled to the public this month (*Weatherhead Management School* conference on AI for corporate partners) is titled "An Innovative Tool For Deciding When AI Predictions Can Be Trusted, Or Not".

This is a fully implemented system for plotting the part of the training data that a neural network actually uses to make a specific prediction. Insights come from showing things missed by classical statisticians, Bayesians, and Machine Learning approaches alike: subset relevance, subset stability, and subsample sizes in the training data. We examine these *after* a particular query is in focus (no more claiming 99.99% accuracy on cat photos while flipping a coin on whether a target is worth bombing, based on lots of cats and not much else). The paper in progress is titled "Inspecting Data-Driven Predictions: Visualizing Reference Classes and Auditing AI".

Many in AI will find Loui's claims of superiority to the Bayes Net bold; but the comparison is well-informed. The Bayes Net was first developed when Judea Pearl at UCLA was a close mentor and friend of Loui. (Pearl's famous book's bibliography cites Loui's graduate work and personal communication more than any other.)

Recently Loui published his "third thesis" titled "Against Narrow Optimization and Short Horizons: An Argument-based, Path Planning, and Variable Multiattribute Model for Decision and Risk" in a UK Journal. This modernizes the typical business school approach to risk analysis. It finds middle ground between discounting and mitigating. The idea is to secure and protect a desirable long path into the future: increasing its probability and meeting milestones within standards of due diligence.

His AI work can be found in MINDS AND MACHINES, AI JOURNAL, AI AND LAW JOURNAL, ACM COMPUTING SURVEYS Strategic Directions on AI, ICAIL, COMPUTATIONAL INTELLIGENCE, COGNITIVE SCIENCE, IEEE COMPUTER, ARGUMENTATION, AI MAGAZINE, SIGPLAN, SIGART, ANALYSIS, JOURNAL OF PHILOSOPHY, BEHAVIORAL AND BRAIN SCIENCE, LAW PROBABILITY AND RISK, many IEEE conferences, and many invited book chapters, including THE MIT ENCYCLOPEDIA OF COGNITIVE SCIENCE. Some of those papers were the best cited of the journal in their particular publication year.

His most recent chapters are in LAW OF AI, and LAW OF ALGORITHMS, the latest being "Adding to the EU AI Liability Directive: Degree of Autonomy, Chain of Confidence, Inherent Flaws of Indecent Induction, and Mandatory Insurance". This work identifies where neural net predictors have design flaws. It proposes tests that might certify ML-predictors (as decent extrapolators, not just very good interpolators).

After 9/11 he worked on a famous project for US intelligence. Over 40 years, his work was for the US Army/USASC, US Navy/USMC, USAF/AFRL, DARPA/IAO/DTO, NSA/SAIC, FBI Infraguard, DHS through ORNL, and NSF. He also consulted on a large AI project with the Cleveland Clinic.

Professor Loui has lectured on AI in a third of US states and in Sweden, Spain, Portugal, Argentina, Germany, and Canada, and at law schools in Tokyo, Melbourne, Buenos Aires, Chicago, DC, Boston, NYC, Amsterdam, Maastricht, Brussels, Paris, Nice, and Pisa. He has worked on a friend's national campaign, advised the Dean of Harvard College on drinking policy and leadership selection, stood with the Chancellor of WashUStL while the twin towers were being attacked, met with the COO of Illinois and Governor of Hawaii to discuss funding, befriended future billionaires when they were just programmers, shared an awards ceremony with the creator of Unix, and was co-captain of his state championship high school math team.