

[PDF] Path-specific harm decomposition: A partial identification framework

R Yan, D Frauen, M Schröder, S Feuerriegel - arXiv preprint arXiv:2609.29938, 2026

A central goal when designing treatment policies is often to "do no harm", that is, to avoid interventions that improve average outcomes while worsening outcomes for some individuals. A widely used notion for harm is the fraction of negatively affected ...

[PDF] Formulating Cross-World Mediation Estimands Through Single-World Mixtures

R Nabi, D Benkeser - arXiv preprint arXiv:2609.27075, 2026

In causal mediation analysis, natural direct and indirect effects are defined through nested counterfactuals that combine the outcome under one exposure level with the mediator value under another and are therefore inherently cross-world. Their ...

[PDF] Fairness-constrained explainable machine learning with differentially private federated training for employee attrition prediction in civil engineering and law

R Kumari, S Pratap, P Anand, P Anand, MB Yesuf... - Scientific Reports, 2026

Employee attrition in knowledge-intensive sectors such as civil engineering and law presents compounding risks to project continuity, regulatory compliance, and organisational profitability. Despite growing adoption of machine learning (ML) in ...

Governing Fairness in Insurance Pricing: The Trilemma of Actuarial Adequacy, Solidarity, and Causal Legitimacy

A Charpentier, O Côté, MP Côté - Journal of Business Ethics, 2026

Insurance is at once a bilateral transfer of risk and a collective device for sharing it. These two functions pull insurance pricing in opposite directions, and because every transaction is mediated by money, the fairness of an insurance contract is mostly ...

Stage-conditioned comparison of intervention scenarios in domino accidents: An evidence-anchored ABRC-DBN framework

X Wang, X Xue, X Sun - Journal of Loss Prevention in the Process Industries, 2026

Domino accidents in hazardous-chemical tank farms evolve through distinct control windows, but assessments often emphasize final consequences rather than stage-conditioned mitigation. This study develops an evidence-anchored accident-barrier ...

[HTML] Residual governance for responsible embodied AI: authority, contestability, and cognitive sovereignty

CW Sun - AI and Ethics, 2026

Artificial intelligence (AI) ethics increasingly concerns systems that do not merely recommend but intervene: clinical decision support, human-machine teaming, robotics and automated handover protocols can shift timing, attention and authority ...

[PDF] xWhyL: Causal Interactive Learning

N Tagliapietra, FP Busch, M Willig, M Zečević, L Halilaj... - arXiv preprint arXiv ..., 2026

Explanations are central to causal reasoning, and cognitive science has long established that the human drive to explain is itself a mechanism for learning about causality. Despite this, learning from those abductive signals is largely ignored in ...

[PDF] The Gold in Bias: Maturing the AI Design Process through Verification

S Maghool, P Ceravolo - arXiv preprint arXiv:2609.29730, 2026

Bias in AI systems is typically framed as a flaw to be minimized, yet it also serves as a critical indicator of underlying weaknesses in data, modeling assumptions, and system design. Existing approaches often treat bias as an isolated problem rather ...

[PDF] Formulating Cross-World Mediation Estimands Through Single-World Mixtures

R Nabi, D Benkeser - arXiv preprint arXiv:2609.27075, 2026

In causal mediation analysis, natural direct and indirect effects are defined through nested counterfactuals that combine the outcome under one exposure level with the mediator value under another and are therefore inherently cross-world. Their ...

[PDF] Functional Causal Discovery via Conditional Covariance Ordering

K Li, R Tan - arXiv preprint arXiv:2609.27256, 2026

We study causal discovery where each node is a random function. Previous studies on this topic rely on structural assumptions, eg, linearity or non-linearity, and distributional assumptions, eg, Gaussianity or non-Gaussianity. In contrast, we make ...

Offline Policy Evaluation and Learning with Harm Constraints

J Chaoqe, Q Yang, J Li, P Wu - Entropy, 2026

Offline policy learning aims to optimize individualized decisions using historical data. However, conventional methods primarily focus on maximizing expected rewards while neglecting individual-level counterfactual harm—cases where the assigned ...

Band-Specific Analysis of Magnetoencephalography Data: Application to Clinical and Auditory Experimentation

V Commuri - 2026

Neural activity in the brain can be divided into several frequency bands: delta (1-4 Hz), theta (4-8 Hz), alpha (8-12 Hz), beta (12-25 Hz), and gamma (25+ Hz). Historically, the delineation of these bands derives from measurements within each ...

[PDF] Supplementary material from: Density dependence and territorial competitors can modulate parrotfish social foraging and herbivory on coral reefs

H Thareja, A Paul, T Alcoverro, R Karkarey, W Pinto... - 2026

Figure S1 Map of study region (left) and atolls sampled with sampling locations (right). At the locations marked in green, we conducted both—behavioural observations of *Chlorurus sordidus*, and community sampling; at the locations ...

Strongly (Dis) agree: Distorting Effects of Extreme Response Style in Survey-Based Extremism Research

M Erdemandi - 2026

This dissertation investigates whether a common but often overlooked source of bias, extreme response style (ERS; a tendency for survey responders to consistently select endpoint options on Likert-type scales), distorts findings in the survey-based ...

Sequential Learning and Inference of Causal Graphical Models

Y Zhao - 2026

Many problems involve systems in which causal relationships among variables are represented by causal graphs. Because experimental data are often costly, learning procedures must combine statistical inference with sequential decision making. This ...

[PDF] Causal Inference for Health Policy Support Systems: A Directed Acyclic Graphs Based Approach

S Diaz, G Galo, W Urrutia, A Diaz, O Gradis, OL Garcia... - LACCEI, 2026

The accurate estimation of causal effects is fundamental for evidence-based health policy decisions. Traditional correlational approaches often produce biased estimates due to uncontrolled confounding, leading to potentially misleading policy ...

Machine Learning for Ozone Pollution Research: A Chemistry-Informed Framework for Model Selection and Interpretation

Y Sun, S Ding, C Wang, D Liu - ACS Earth and Space Chemistry

Ozone pollution remains a major challenge for air quality management, ecology and public health because of its nonlinear formation chemistry, strong meteorological modulation, multiscale transport and toxicity. Although chemical transport models ...

[PDF] Proximal Balancing for Causal Effect Estimation under Unmeasured Confounding

Y Jung

Estimating treatment effects from observational data, a central task in medicine and public policy, usually assumes that every confounder (every common cause of the treatment and the outcome) is measured. This assumption, known as no ...

[PDF] Take, or Pay

W Oosterhout, C Merz - Journal of Peer Rejected Studies• PRJ

Industrial gas is sold, overwhelmingly, on take-or-pay contracts: a supplier builds an air-separation plant at the customer's fence line and, for fifteen to twenty years, the customer pays for a committed volume of oxygen or nitrogen whether or not it is ever ...

[HTML] The role of DNA methylation-based biological aging, inflammation and immunity in post-stroke cognition: A unified causal framework

R Waziry, CH Miles, O Williams, H Tiemeier - PLOS Aging and Health, 2026

In the present study, we assessed the direct and indirect effects of stroke on short-term and intermediate-term cognition with respect to biological aging, inflammation and immunity. The study included individuals with stroke information at wave 12, who ...

[PDF] Interpretability and Safety in Language Models: Methods, Mechanisms and Evaluations

WJ Yeo - 2026

While the rapid scaling of Large Language Models has enabled significant automation of complex tasks, a widening gap persists between the growing capabilities of these models and our understanding of their internal decision-making ...

[\[PDF\] Layer-wise structural and mechanistic interpretation of OpenVLA using controlled perturbations](#)

T Yıldız - 2026

Limitations of conventional policies in robot learning have driven the adoption of large-scale pretrained models, known as foundation models, leading to the emergence of Vision-Language-Action (VLA) models. Although VLA models have ...

[Psychometric and conceptual challenges in dynamic risk factors for recidivism](#)

SC Lee, RK Hanson - *Journal of Sexual Aggression*, 2026

Dynamic risk factors play a central role in correctional practice because they are theoretically malleable and empirically linked to recidivism. Using them to guide supervision, treatment, and discharge decisions requires addressing several ...

[\[HTML\] Explainable Machine Learning in Mineral Prospectivity Mapping: A Critical Review of Methods, Geological Knowledge Embedding, Validation, and Future Directions](#)

M Lu, L Zhong, W He, Y Zhao, D Sun, J Ma, J Hu... - *Minerals*, 2026

Mineral prospectivity mapping (MPM) increasingly integrates geological, geochemical, geophysical, remote-sensing, and structural evidence through machine-learning workflows. Predictive accuracy alone, however, does not establish ...

[\[PDF\] RPCShield: Defending Microservices Against Cascading Failures](#)

M Chabbi, S Mahajan, I Beschastnikh, R Just, Y Wang... - *Proceedings of the ACM* ..., 2026

Microservice-based systems are tightly interdependent. A failure in one service can propagate to a dependent service, potentially bringing down critical, user-facing functionality. We have developed and deployed RPCShield, a suite of program ...

[Beyond regulation: adaptive responses of nonregulated SMEs to mandatory CSR reporting](#)

AM Huq, F Hartwig, W Bai, N Rudholm - *Journal of Accounting & Organizational* ..., 2026

Purpose This study aims to examine whether mandatory sustainability/corporate social responsibility (CSR) disclosure generates spillover effects beyond firms directly subject to regulation. While prior research focuses mainly on direct ...

[Causal analysis of an industrial iron-ore flotation circuit: Which setpoint effects can be identified from plant data](#)

J Caers, A Eskanolu - *Minerals Engineering*, 2027

Abstract Machine-learning soft sensors predict concentrate grade in froth flotation with good accuracy, and SHAP or partial dependence are then read as guidance for changing a setpoint. Those readings answer an associational question. A setpoint ...

[Event-time regression discontinuity for biomedical data: LATT estimation and covariate distribution tests](#)

Y Choe, Y Kwon, M Kho, S Lee, S Lee - *American Journal of Epidemiology*, 2026

Estimating causal effects of drugs on clinical biomarkers is often infeasible through randomized trials, while observational data require designs that approximate experimental conditions. Using linked UK Biobank records collected during 2006 ...

[HTML] [Explicit Mechanistic Causal Analyses or Interventional Trials Are Required for Objective, Clinical, Voice-Based Parkinson Disease Characterization](#)

M Little - *Journal of Medical Internet Research*, 2026

For most with Parkinson disease, a movement disorder, voice and speech are impaired at some point, which presents an opportunity to use digital recordings and sophisticated machine learning to assist objective clinical characterization of the ...

[PDF] [Demistifying Data and Simulator Assumptions in Supervised Causal Discovery](#)

P Ma, R Ding, B Huang, S Wang - *arXiv preprint arXiv:2609.37446*, 2026

Supervised causal discovery learns to infer causal structure for a new dataset from training datasets paired with structural labels. These training pairs are typically simulated, making the simulator both a source of supervision and a carrier of ...

[PDF] [Counterfactual Self-Evolving Agents for Evidence-Grounded Reasoning](#)

X Han, Y Wang, C Chen, W Dai, GK Gudur, S Li... - *arXiv preprint arXiv ...*, 2026

Self-play proposer--solver methods improve reasoning by generating tasks and learning from verified solutions. However, for evidence-identifiable tasks, where case-specific evidence and domain knowledge determine a checkable answer, self-play ...

[PDF] [Policy coordination of public data openness and supercomputing center for agricultural chain resilience](#)

Z Chen, J Zheng, C Chi - *Frontiers in Sustainable Food Systems*, 2026

Strengthening agricultural supply chain resilience is critical for safeguarding food security and promoting high-quality agricultural development in the face of escalating external shocks, including climate extremes, market volatility, and geopolitical ...

[PDF] [From Attack Success to Attack Severity: Counterfactual Memory Attacks on LLM Agents](#)

M Zou, L Liang, Z Wang, Y Zhao, L Qu, Z Xu - *arXiv preprint arXiv:2609.34132*, 2026

As LLM agents increasingly rely on persistent memory for long-horizon and personalized behavior, they can retain and reuse information across interactions, but this also creates a lasting channel through which malicious memory writes can ...

[PDF] [A Mathematical Theory of Cognitive Filiation](#)

P Walecki - 2026

This paper proposes a formal theory of cognitive filiation: a historically directed relation in which an organized cognitive process transmits causally operative structure to a distinguishable successor. Its mathematical primitives are time-indexed ...

[PDF] [FOCUS: Fixed-Confidence Online Causal Learning Using Sequential Adaptive Interventions](#)

H Xu, C Zhang - *arXiv preprint arXiv:2609.32165*, 2026

We study fully online fixed-confidence causal discovery without any historical observational data. Starting from zero samples, the learner sequentially selects interventions to recover both the causal DAG and its edge weights under a linear ...

[PDF] Energy Vision--Language--Action: A Controlled Multimodal Benchmark for Intent-Conditioned Residential Energy Management

LS Saoud, O Doukhi, E Reihani, S Sepasi, DJ Lee... - arXiv preprint arXiv ..., 2026

Vision-Language-Action (VLA) models are studied mainly in robotics, where visual observations and language instructions are mapped to physical actions. This paper introduces Energy Vision-Language-Action (EVLA), a controlled multimodal ...

[PDF] seq2cause: One Autoregressive Backbone, Four Causal Discovery Tasks in Event Sequences

H Math - arXiv preprint arXiv:2609.31801, 2026

Complex systems such as vehicles, patients, or genomes emit discrete event sequences whose operative question is causal, not predictive: which events cause which other events, and which cause higher-level outcomes such as failures or ...

[PDF] Causal Retention in Interactive Agents: Interface Factorization and Selective Adaptation

S Zhang, T Liu, D Xie, Y Dong, X Wang, C Zeng - arXiv preprint arXiv:2609.30650, 2026

Task performance need not determine which intervention mechanism an agent retains. We study causal retention: whether a frozen learned state answers a mechanism-probe map fixed independently of training, including action, context ...

[PDF] I Act Therefore I Am: When Is JEPA's Action-Conditioning Enough to Learn Causal Mechanisms?

Y Liu, Z Huang, JQ Shi - arXiv preprint arXiv:2609.31161, 2026

Recent empirical and theoretical advances suggest that joint-embedding predictive architectures (JEPAs) may learn meaningful representations for action-conditioned prediction of future outcomes, thus becoming one of the foundational structures for ...

[PDF] CCRV-Bench: Constraint-Based Evaluation of Causal Reasoning in Vision-Language Models

L Gao, Y Wu, Y Chang - arXiv preprint arXiv:2609.30979, 2026

Vision-language models (VLMs) have demonstrated excellent performance in visual tasks, but their visual causal reasoning capabilities still lack reliable evaluation. Existing evaluations struggle to distinguish whether a model is performing causal ...

[PDF] Beyond Accuracy: Counterfactual Fragility and Demographic Bias in Clinical Evaluation of LLMs

CD Purkayastha, BK Bolla, VSR Nandi - arXiv preprint arXiv:2609.32807, 2026

Clinical LLM evaluation often emphasizes answer accuracy; however, accuracy alone does not test counterfactual consistency or demographic robustness. We evaluated six LLMs on 150 MedQA USMLE questions using two automated ...

[HTML] From Algorithmic Performance to Hospital Value: A Narrative Review of Clinical Artificial Intelligence

C Morfaki - Hospitals, 2026

Clinical artificial intelligence (AI) has demonstrated performance across diagnostic, predictive, monitoring, and decision-support tasks, yet hospital adoption depends on whether that capability can produce sustainable value under local conditions ...

[\[HTML\] Admissibility defines structural limits on optimisation in artificial intelligence](#)

[D Moore](#) - [Discover Artificial Intelligence](#), 2026

Many persistent failures of contemporary artificial intelligence are treated as failures of optimisation, data coverage, or statistical robustness. This paper argues that a recurring class of these failures arises earlier: the observation regime under which a ...

[Physics-guided bidirectional modeling of strength evolution and constrained inverse optimization for alkali-activated materials](#)

[H Ma](#), [F Ma](#), [J Feng](#) - [Journal of Building Engineering](#), 2026

The mechanical property evolution of alkali-activated materials (AAMs) is governed by nonlinear coupling between precursor composition and reaction kinetics. Conventional empirical trial-and-error design suffers from low efficiency, poor ...

[\[HTML\] Individual and group fairness assessments via counterfactual explanations](#)

[F Sabbatini](#), [R Calegari](#) - [AI and Ethics](#), 2026

This study explores the potential of counterfactual explanations to assess artificial intelligence (AI) fairness, especially in critical decision-making systems. Predictive models may amplify biases inherent in data sets or algorithms, and given the ...

[\[PDF\] CADeT: Causal-Aware Deformation Transmission for Indirect Robotic Manipulation of Soft Tissue](#)

[J Hu](#), [D Jones](#), [P Valdastrì](#) - [arXiv preprint arXiv:2609.38483](#), 2026

Indirect manipulation of deep-seated deformable anatomy inaccessible to the robot is challenging in robot-assisted minimally invasive surgery (RAMIS) because intervening tissues spatially filter deformation transmission. Passive observations ...

[Criminal Legal System Involvement During Adolescence as a Social Determinant of Adulthood Body Mass Index?](#)

[PDIA Silver](#) - [Journal of Adolescent Health](#), 2026

Purpose A variety of environmental and genetic factors contribute to body mass index (BMI) over the life course, including experiences and events that occur during adolescence. One such experience that could shape early adulthood and adulthood ...

[Teleological Inference](#)

[B Claessens](#) - [Philosophy and Phenomenological Research](#), 2026

How do we acquire teleological beliefs? A teleological belief is a belief that some event or process occurs for the sake of a purpose. Such beliefs include beliefs about actions, natural functions, and the meaning of life. This article will advance a theory ...

[\[PDF\] Inference-Layer Security: Defending Against Adversarial Inference and Infrastructure Abuse](#)

[K Lee](#) - [arXiv preprint arXiv:2609.38239](#), 2026

A Technical Report: Operating a large language model (LLM) as a service requires more than inference infrastructure: the provider must also defend against adversarial interactions that seek to exploit the service, including jailbreaking for harmful use ...

From Credit Scoring to Governed Decision Systems: Integrating Predictive Performance, Uncertainty, and Governance

N Khalili, AS Sedghpour, M Jahanbakht, Z Hosseini - *Machine Learning and ...*, 2026

Credit-scoring systems must translate predictive risk into calibrated, uncertainty-aware, and governable decisions rather than ranking applicants alone. This study aims to develop and evaluate a governed credit-scoring architecture integrating ...

[PDF] СТРУКТУРНІ ПРИЧИННІ МОДЕЛІ ДЛЯ ЗАДАЧ ПРИЙНЯТТЯ РІШЕНЬ В ІНТЕРАКТИВНИХ СЕРЕДОВИЩАХ

В Ленартович - Редакційна колегія, 2026

691 часто виявляються «чорними скриньками», що обмежує їх застосування в критичних інтерактивних системах. Перехід до структурного причиннонаслідкового моделювання дозволяє не лише покращити якість ...

[PDF] Towards Deep-Learning Architectures That Respect Causal Principles

O Mian, M Kamp - *AXIOM: Foundations of Efficient Deep Learning*

Deep learning and causal inference have largely developed as separate fields. Deep learning provides scalability and flexibility, but it mainly learns statistical associations between variables. Causal learning, on the other hand, aims to ...

[PDF] SLICE-MIL: Counterfactual Semantic Anchoring of Disentangled Evidence for Child-Pugh Grading

Y Deng, H Huang, Y Peng, L He, W Cui, Y Leng...

Child-Pugh (CP) grading, a widely used clinical standard for therapeutic decision-making in hepatocellular carcinoma, integrates biochemical and clinical markers to assess hepatic reserve. Estimating CP grade from routinely acquired contrast ...

[PDF] Counterfactual Stress-Testing for Fixing Lung Cancer Segmentation Models

A Santhirasekaram, C Johnson, M Chen

Segmentation models for lung cancer on CT can achieve strong performance yet remain brittle to clinically relevant morphological and textural variations, with failures more prevalent in certain phenotypes. We introduce a counterfactual stress-testing ...

[PDF] A Human-like Qualitative Control Algorithm for Autonomous Pushing

J Lemeire, R Spolmink, A Zerri, D Nkansah - *The 38th Benelux Conference on Artificial ...*

A novel generic control algorithm based on qualitative causal models is proposed. These models describe the causal relationships between action and state changes in terms of sign relations, ie, the direction of change in the state variables. Qualitative ...

[PDF] Causality-Inspired Representation Learning for Student Performance Prediction

B Liu, Z Li, J Yang

Predicting student performance is crucial for timely educational interventions, yet most existing approaches primarily focus on statistical pattern learning that may not fully reflect the underlying factors associated with academic success. This limitation ...

[\[PDF\] Double Machine Learning for Downstream Causal Measurement: A Hybrid Econometric-ML Framework for Feature Valuation](#)

[DP Uddandarao](#), [A Mahajan](#), [P Krishnamsetty](#)

Digital platforms require defensible evidence on which product features causally improve downstream customer value, yet feature-prioritization systems often rely on predictive importance, usage frequency, or descriptive engagement indicators ...

[\[PDF\] Right Answer, Wrong Evidence: Revealing Visual Confabulation in Multi-Image Medical VLMs with the Evidence Attribution Score](#)

[C Kuppani](#), [B Halavar](#)

Vision-language models achieve increasing accuracy in medical imaging tasks, but accuracy alone cannot reveal whether models rely on the correct visual evidence. In clinical settings, decisions must be justified by evidence, not merely correct ...

[\[TXT\] IMPROVING OUTCOMES IN MACHINE LEARNING AND DATA-DRIVEN LEARNING SYSTEMS USING STRUCTURAL CAUSAL MODELS A DISSERTATION ...](#)

[HM MBOGU](#)

UNIVERSITY OF OKLAHOMA GRADUATE COLLEGE IMPROVING OUTCOMES IN MACHINE LEARNING AND DATA-DRIVEN LEARNING SYSTEMS USING STRUCTURAL CAUSAL MODELS A DISSERTATION SUBMITTED TO THE ...

[\[TXT\] IMPROVING OUTCOMES IN MACHINE LEARNING AND DATA-DRIVEN LEARNING SYSTEMS USING STRUCTURAL CAUSAL MODELS A DISSERTATION ...](#)

[HM MBOGU](#)

UNIVERSITY OF OKLAHOMA GRADUATE COLLEGE IMPROVING OUTCOMES IN MACHINE LEARNING AND DATA-DRIVEN LEARNING SYSTEMS USING STRUCTURAL CAUSAL MODELS A DISSERTATION SUBMITTED TO THE ...

[Potential Outcomes Effects in Mediation Designs with Three Experimental Conditions](#)

[D Alvarez-Bartolo](#), [HL Smyth](#), [AM Brincks...](#) - [Prevention Science](#), 2026

Causal mediation analysis, based on the potential outcomes framework, is an important new tool for prevention research. Causal mediation analysis is a non-parametric, general technique for estimating causal effects in mediation analysis with ...

[\[HTML\] Covariate Selection in Physical Activity Epidemiology: Too Little, Too Much, or Just Right?](#)

[JM Blodgett](#), [JJ Mitchell](#), [EM Steele](#), [K Milton](#), [M Hamer...](#) - [Journal of Physical Activity ...](#), 2026

Many observational studies are interested in the total effect of physical activity on any given health outcome. These studies rely heavily on covariate adjustment, yet decisions about which variables to include vary widely. Such decisions have ...

[An adaptive test for the conditional natural indirect effect in mediation analysis](#)

F Liang, X Zhu, L Zhu - *Science China Mathematics*, 2026

In this paper, we introduce a model for the conditional natural indirect effect (CNIE), along with relevant estimation and a test statistic to examine the presence of mediation effects within a nonlinear framework that includes outcome and mediator ...

[\[HTML\] Beyond correlation and causation: epistemic caution, anomalous findings, and COVID-19 vaccine safety](#)

AS Siebner, A Rubio-Casillas, M Raszek, D Cowley... - *Philosophy, Ethics, and ...*, 2026

The COVID-19 pandemic exposed science's adaptive strengths and epistemic limits. In fact, in the pandemic, mRNA vaccines are attributed to have saved millions of lives. The maxim "correlation does not imply causation" is an essential safeguard ...

[\[PDF\] CID: Clean-Seed-Free Backdoor Defense via Counterfactual Invariance for Neural Code Models](#)

J Ye, Z Li, X Tang, S Li, S Liang, D Zou, H Jin - *Proceedings of the ACM on Software ...*, 2026

Neural code models automate core software engineering tasks such as code classification and generation, but remain vulnerable to backdoor attacks. Existing defenses struggle with both injection-based and Semantically-Equivalent ...

[Learning Causal Structures from LLM-AHP Comparisons for Interpretable Forecasting](#)

L Yang, X Hu, W Pedrycz, Z Li - *Knowledge-Based Systems*, 2026

Large language models (LLMs) can generate pairwise comparisons via the analytic hierarchy process (AHP), yet existing methods treat these judgments independently, discard the causal structure embedded across scenarios, and provide no uncertainty ...

[\[PDF\] Cognitive Nutrition: Architecture of a Provenance-Embedded Data Refinery](#)

T MacPherson - *Metavolve Labs*, 2026

We present the architecture, operational metrics, and theoretical foundations of a production metadata factory that transforms raw media artifacts into provenance-verified, semantically dense training assets—what we term Cognitive Nutrition (CN) ...

[\[HTML\] Covariate Selection in Physical Activity Epidemiology: Too Little, Too Much, or Just Right?](#)

JM Blodgett, JJ Mitchell, EM Steele, K Milton, M Hamer... - *Journal of Physical Activity ...*, 2026

Many observational studies are interested in the total effect of physical activity on any given health outcome. These studies rely heavily on covariate adjustment, yet decisions about which variables to include vary widely. Such decisions have ...

[Attention-Guided Causal Structure Learning With Gated Feature Fusion](#)

Z Song, J Lin, Y Wei, L Guo - *Concurrency and Computation: Practice and ...*, 2026

Causal discovery, also referred to as causal structure learning, aims to identify causal relationships among variables from observational data and represent them as a directed acyclic graph (DAG). Existing differentiable causal discovery methods ...

[\[PDF\] Causal Memory Policy: Making Memory Utility Identifiable by Intervening on Retrieval](#)

A Behnam, B Wang - [arXiv preprint arXiv:2610.02070](#), 2026

Memory-augmented large language models must decide which memories to retain, and recent systems do so by estimating each memory's effect on task performance. However, these estimates rely entirely on retrieved memories. When a memory is ...

[\[PDF\] Admissible Portfolio Optimization: Information Constraints, Conditional Efficient Frontiers, and the Price of Causal Identification](#)

AR Dominguez - [arXiv preprint arXiv:2610.00147](#), 2026

Mean--variance portfolio choice takes the conditioning information as given and optimizes over weights, so two errors about that information pass into the portfolio unseen: using variables unavailable at the decision time and treating common ...

[\[PDF\] Every Ablation Is a Dose: Counterweights and the Semblance of Self-Repair](#)

A Ahmad, P Seth, VK Sankarapu - [arXiv preprint arXiv:2610.02173](#), 2026

Ablate a component of a language model, and other components often appear to adjust and compensate. This phenomenon, termed self-repair, has been observed repeatedly, but its mechanism remains unclear. The most systematic study to date ...

[\[PDF\] Stable and Counterfactually Robust Physical World Models from Imposed Structure and Learned Physics](#)

Y Wang, P Priye, L Wei, H Ling - [arXiv preprint arXiv:2610.00280](#), 2026

A world model learns to forecast how a physical system evolves from recorded trajectories, yet the systems it imitates obey physical laws that are neither fully supplied nor reliably respected. The model may create energy, drift or diverge over ...

[\[PDF\] Causal Trade-Offs in Superluminal Extensions of Relativity. A no-go theorem](#)

A Sen, M Salzger, V Vilasini - [arXiv preprint arXiv:2610.00494](#), 2026

Extensions of relativity that include superluminal reference frames have recently been proposed by Dragan and Ekert as a way to reconcile relativistic symmetry with quantum uncertainty. In this work, we develop a causal modelling framework to ...

[\[PDF\] Causal Price-of-Risk Mandates under Overlapping Information](#)

AR Dominguez - [arXiv preprint arXiv:2610.00158](#), 2026

We study whether causal risk mandates constructed from overlapping information blocks can be implemented by one self-financing portfolio that is optimal under pooled information. In an incomplete continuous Brownian market, signal-responsive ...

[CausalGA: An Evolutionary Framework for Causal Structure Learning](#)

LA Duarte, MC Teixeira, AIR Carvalho, E Birchal... - [Brazilian Conference on ...](#), 2026

Discovering causal relationships from observational data is a fundamental challenge in several domains, and existing causal discovery algorithms frequently converge to divergent structures depending on their specific theoretical assumptions, acting as ...

[\[PDF\] What Conceptual Change Cannot Recover: Interface Closure, Epistemic Recoupling and Representational Sealing](#)

P Kahl - 2026

Accounts of conceptual change tend to run together two operations: representational innovation—new concepts, mathematics and standards of correctness—and epistemic recoupling, a change in how a system is causally coupled to what it ...

[\[PDF\] Adaptive Machine Learning](#)

MH Peroni - 2026

Modern machine learning systems are increasingly deployed in settings where predictive performance must be balanced against real-world constraints, including limited data availability, heterogeneous acquisition costs, and the need for ...

[\[PDF\] Electricity Market Prices in Transition: Determinants, Regimes, and Causal Structure](#)

S Madadkhani - 2026

In the fight against climate change, rapid decarbonization of the electricity sector is imperative. Electricity and heat generation are the largest source of anthropogenic greenhouse gas (GHG) emissions worldwide (Our World in Data, 2024), and clean ...

[\[PDF\] Mathematics and AI Oracles](#)

ST Piantadosi - 2026

The best definition of mathematics I have seen is “the study of mental objects with reproducible properties”[DHM12], meaning entities that can be perfectly replicated from one human mind to another. Most concepts we have are shared only ...

[\[PDF\] Counterfactual Regret Evolutionary Agents \(CREA\)](#)

AK Anand

Financial decision-making is increasingly challenged by non-stationarity, tail risk, ambiguous information, and causal complexity. Traditional quantitative models often optimize policies against historical returns or point forecasts, but such optimization ...

[\[PDF\] Mitigating Algorithmic Bias in Information Systems: A Comprehensive Framework for Fair Artificial Intelligence](#)

M Qasim, L Tabish, F Ali

As information systems (IS) transition from deterministic rule-based engines to probabilistic machine learning architectures, the risk of systematic bias and disparate impact increases significantly. While advanced predictive models optimize ...

[Engineering Data Analysis Beyond Statistics vs. Machine Learning: A Review of Explanation, Prediction, Hybrid Modeling, and Responsible Decision Support](#)

HY Chen, C Chen - Eng, 2026

This structured narrative review explores how engineers select and evaluate traditional statistical, machine learning, mechanistic, and hybrid methods for explanation, prediction, and responsible decision support. We identified ...

Future Trends and Challenges in Explainable AI

V Ahmad, R Kumar - *Explainable Artificial Intelligence: Bridging Concepts ...*, 2026

The interpretability and explainability of AI have been the focus throughout the entire research agenda for artificial intelligence (AI), and explainable AI (XAI) is a key focal point. This chapter reviews the current state of this fast-evolving field, emerging ...

Machine Learning for Screening Bioactive Compounds in Traditional Chinese Medicine: Computational Prediction, Multi-Omics Integration, and Clinical Translation ...

ZM Qi, JZ Wu, P Li, TY Xu, P Yu - *Pharmaceuticals*, 2026

Identifying bioactive compounds in traditional Chinese medicine (TCM) requires computational prioritization to be connected with chemical and pharmacological evidence. This structured narrative and critical review examines machine learning ...

Explainable AI (XAI) Bridging the Gap Between Black Box Models and Human Trust: A Comprehensive Framework for Transparency, Interpretability, and Responsible ...

MH Devare, PKR Vishwakarma, AM Devare... - *Rethinking Internal Audit ...*, 2027

Artificial Intelligence has achieved significant progress through deep learning, enabling applications in healthcare, drug discovery, autonomous systems, and financial security. However, the complexity of these models often leads to the “black ...

Multimodal Subjective Intention Modeling in Social Interactions

A Mercier - *Proceedings of the 28th International Conference On ...*, 2026

Understanding human intentions in complex social interactions is essential for developing socially intelligent systems, yet remains challenging due to their subjective, context-dependent, and multi-interpretable nature. Existing approaches ...

Directed Acyclic Graphs for Causal Modeling of Musical Interventions in High-Tropical Grazing Dairy Cattle: A Causal Framework Review

P Betancourth-Chaves, JJ Montoya-Zuluaga... - *Veterinary Sciences*, 2026

Musical exposure in dairy cattle has been associated with potential effects on milk production, behavior, and physiological regulation; however, available evidence remains heterogeneous and potentially affected by biological and environmental ...

A Causal Inference Protocol for Root-cause Diagnosis in Non-stationary Industrial Time Series

C Yalim - 2026

Industrial fault diagnosis increasingly relies on data-driven methods, yet most approaches assume process stationarity and treat association-based rankings as sufficient evidence for root-cause identification. Neither assumption holds when ...

Linking projections from species distribution models and management in Indigenous communities

R Gamlen-Greene, W Godsoe, J Prevost, P Arcese - 2026

Species distribution models (SDMs) are widely used to forecast invasions, but translating these predictions into management remains challenging. We address this using the introduced northern red-legged frog (*Rana aurora*) on Haida Gwaii (British ...

[PDF] Exploring cause-effect relationships for trustworthy language-based maintenance chatbots

W Prokop - 2026

This thesis examines the effectiveness of enhancing Large Language Model (LLM)-based maintenance chatbots through Retrieval-Augmented Generation (RAG), incorporating both Vector Similarity Search (VSS) for semantic information retrieval ...

Causal Machine Learning and Explainable AI for Understanding the Impact of Technology Use on Obesity in Healthcare

ANA Tüylü - Global Digital Health Transformation: AI, Policy, and ..., 2027

Obesity is one of the most significant public health problems of our time due to its multifaceted causes and increasing prevalence. The integration of digital technologies into daily life has increased interest in the effects of technology use on ...

[PDF] Gender data bias e disuguaglianze in una prospettiva sociologica

S Tusini, M Dentale - SOCIOLOGIA, 2026

Il contributo mostra come il gender data bias non sia un semplice errore tecnico, ma il prodotto di processi sociali situati, nei quali i dati riflettono priorità istituzionali, asimmetrie di accesso, lacune informative e rappresentazioni culturalmente ...

[PDF] Causal Machine Learning for Personalized Recommendation of Bioactive Peptides in Aging-Related Disease Prevention

F Salgari, J Devi, K Wahner, J Utama - Computer Science and Engineering ..., 2026

Aging-related diseases represent a convergent burden shaped by biological aging, chronic inflammation, metabolic dysregulation, microbiome shifts, and cumulative environmental exposures. Bioactive peptides derived from food and animal sources ...

[PDF] Causal Perspectives on Federated Unlearning for Agentic Generative Systems

V Kaustubha

Federated learning is increasingly used to train large generative models that power agentic systems capable of planning, tool use, and multi-step decision making. However, existing machine unlearning methods in federated settings are typically ...

[PDF] Noir Hunting: A Falsifiable Cross-Domain Framework for Hidden-Structure Discovery

PT Trung

Open-ended scientific search is difficult when the target structure is not known in advance, candidate spaces are large, and apparently strong signals can be generated by confounding, leakage, representation choices, or distribution shift. This ...

[PDF] The Causal Survival of Persons: An Interventionist Analysis of the Causal Conditions of Personal Survival

S Sasaki

Parfit argued that what matters in personal survival is Relation R with the right kind of cause. After distinguishing the normal cause, a reliable cause, and any cause, he endorsed the last option. His treatment of this causal requirement, however, has ...

[\[PDF\] Pseudoupdate: A Naturalized Framework for the Experimental Study of Recursive Non-Actualization](#)

[Z Jezyk](#)

This paper develops a naturalized and empirically testable framework for a specific problem in cognitive science: the possibility that a system can undergo substantial internal modification while failing to acquire correspondingly new information about ...