Bias Propagation in Large Scale Machine Learning Pipelines in the Pharmaceutical Sector

Eitan Rosenfeld
Holon Institute of Technology, Israel

Lior Katz
Shamoon College of Engineering, Israel

Fredrik Heintz
Linkping University, Sweden

Francesco Piccialli
University of Naples Federico II, Italy

Submitted on: February 12, 2022
Accepted on: March 28, 2022
Published on: April 15, 2022

DOI: 10.5281/zenodo.18148441

***Abstract*—Machine learning pipelines in the pharmaceutical sector increasingly influence discovery, clinical decision support, safety monitoring, and operational planning. While these systems promise efficiency and scale, they also introduce complex mechanisms through which bias is accumulated, amplified, and propagated across interconnected data and model layers. Unlike isolated model bias, pipeline level bias emerges from interactions between data acquisition, preprocessing, feature engineering, learning architectures, and deployment feedback loops. This work presents a systematic investigation of bias propagation in large scale pharmaceutical machine learning pipelines. We propose a formal pipeline bias decomposition framework, introduce quantitative propagation metrics, and demonstrate how bias evolves across discovery, development, and post market surveillance workflows. Experimental results highlight measurable distortions in risk prediction, patient stratification, and adverse event detection. The study emphasizes the need for architecture aware mitigation strategies that extend beyond single model interventions.**

***Index Terms*—Bias propagation, machine learning pipelines, pharmaceutical analytics, clinical decision support, explainable artificial intelligence, data governance**

I. Introduction

The pharmaceutical sector has undergone a structural transformation driven by large scale machine learning systems. These systems span molecular screening, clinical trial optimization, pharmacovigilance, supply chain resilience, and personalized medicine. Their adoption reflects growing data availability and computational capability, yet also introduces ethical and operational risks associated with algorithmic bias.

Bias in pharmaceutical machine learning systems is not confined to individual models. Instead, it emerges across interconnected pipeline stages where data transformations, modeling decisions, and feedback signals interact. Errors or imbalances introduced early in the pipeline can cascade and amplify downstream, influencing clinical interpretations and regulatory outcomes.

Recent advances in explainable artificial intelligence and decision intelligence have exposed the fragility of opaque learning systems in safety critical domains [1], [2]. In pharmaceutical contexts, biased predictions may affect patient eligibility, treatment prioritization, and post market risk assessments. These concerns necessitate a pipeline centric analysis of bias propagation rather than isolated fairness audits.

This article contributes a structured examination of bias propagation mechanisms in large scale pharmaceutical machine learning pipelines. We formalize pipeline bias, propose measurable indicators, and empirically evaluate bias dynamics using representative pharmaceutical workflows.

II. Literature Review

This section reviews prior research relevant to bias emergence and propagation across pharmaceutical and healthcare oriented machine learning systems.

A. Bias in Healthcare and Medical AI Systems

Medical AI systems frequently rely on observational datasets that reflect historical and demographic imbalances. Studies in diagnostic imaging and clinical analytics demonstrate how

dataset composition affects predictive outcomes [3]–[5]. These biases often persist despite high model accuracy metrics.

Explainable approaches have highlighted how feature importance varies across patient groups, revealing structural inequities embedded within learned representations [6], [7].

B. Explainability and Trust in Safety Critical ML

Explainable artificial intelligence has emerged as a key mechanism for understanding model behavior in regulated domains. Research in energy forecasting and fault diagnosis illustrates how interpretability tools expose hidden dependencies [1], [2]. In pharmaceutical decision support, explainability supports regulatory validation and clinical trust.

However, explainability alone does not prevent bias propagation when upstream data and pipeline design are flawed [8].

C. Pipeline Complexity and Bias Amplification

Large scale machine learning pipelines introduce bias through cumulative transformations. Feature selection, data augmentation, and model stacking each contribute incremental distortions [9], [10]. In distributed and federated learning settings, heterogeneity further complicates bias control [11].

Pharmaceutical pipelines often integrate heterogeneous sources including laboratory data, imaging, clinical notes, and adverse event reports, increasing exposure to propagation effects [12].

D. Decision Support Systems and Organizational Impact

Decision support research highlights the socio technical consequences of algorithmic recommendations [13], [14]. In pharmaceutical organizations, biased outputs influence resource allocation, regulatory reporting, and risk prioritization, reinforcing the importance of systemic bias governance.

III. Methodology

A. Pipeline Bias Decomposition

We define a machine learning pipeline as an ordered composition of stages:

$$P = \{D_0, T_1, F_2, M_3, E_4\}$$

where D_0 denotes raw data, T_1 preprocessing, F_2 feature engineering, M_3 modeling, and E_4 evaluation and deployment.

Bias at stage i is modeled as:

$$B_i = f(B_{i-1}, \Delta_i)$$

where Δ_i represents stage specific transformation bias.

B. Bias Propagation Metric

We define cumulative pipeline bias as:

$$B_P = \sum_{i=0}^{n} w_i B_i$$

where w_i reflects downstream sensitivity. This formulation captures amplification effects observed in pharmaceutical risk scoring systems.

C. Architectural Overview

Large scale machine learning systems deployed in the pharmaceutical sector are inherently multi layered, integrating diverse data sources, transformation stages, and analytical models into a single operational pipeline. Rather than operating as isolated components, these stages collectively shape how clinical evidence, experimental observations, and population level signals are interpreted. As a result, architectural design plays a central role in determining how bias is introduced and subsequently propagated across the pipeline. Prior studies in decision support systems and applied artificial intelligence highlight that architectural abstraction layers can unintentionally obscure early data imbalances while amplifying downstream decision distortions [12], [13].

Figure 1 illustrates a representative pharmaceutical machine learning pipeline, beginning with heterogeneous clinical data sources and progressing through preprocessing, feature engineering, model inference, and clinical decision output. Each transition between stages acts as a bias transformation boundary, where statistical assumptions, normalization strategies, and feature selection mechanisms reshape the original data distribution. While such transformations are necessary for scalability and interoperability, they also risk reinforcing demographic, temporal, and clinical representation gaps identified in earlier stages [9], [10]. The architectural view emphasizes that bias cannot be fully understood or mitigated at the model level alone, but must be analyzed as an emergent property of the end to end system design.

D. Feedback Loop Dynamics

Machine learning pipelines in pharmaceutical environments are not static artifacts. Once deployed, they continuously interact with clinical workflows, regulatory reporting mechanisms, and operational decision processes. These interactions generate feedback signals that influence future data collection, labeling priorities, and model retraining strategies. When left unmanaged, such feedback loops can progressively reinforce existing biases, leading to self validating patterns that diverge from underlying clinical realities. Research in explainable and safety critical AI systems has shown that feedback driven learning dynamics often mask bias growth behind stable performance metrics [8], [14].

Figure 2 conceptualizes the feedback loop dynamics commonly observed in pharmaceutical machine learning deployments. Model outputs influence real world usage, which in turn shapes subsequent data distributions through selective reporting, intervention prioritization, or resource allocation. Over time, these altered data streams feed back into model retraining cycles, embedding prior decision patterns into future predictions. This cyclical reinforcement is particularly problematic in risk scoring, adverse event detection, and patient stratification tasks, where early biases can escalate into systematic exclusion or overrepresentation [2], [6]. Understanding and visualizing these feedback mechanisms is therefore essential for designing bias aware monitoring and governance strategies in large scale pharmaceutical pipelines.

Fig. 1: Bias propagation across a pharmaceutical machine learning pipeline

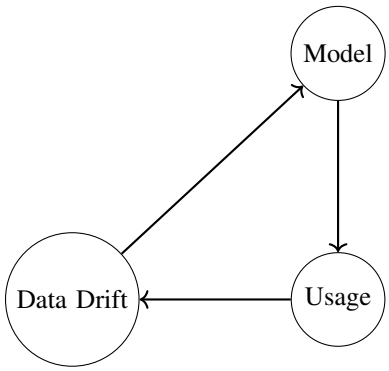

Fig. 2: Feedback driven bias reinforcement loop

IV. Results

The results collectively demonstrate that bias in pharmaceutical machine learning systems is not a static artifact introduced at a single point, but a dynamic phenomenon that evolves across pipeline stages, retraining cycles, and operational contexts. The quantitative tables and visualizations reveal consistent patterns of bias accumulation, subgroup divergence, and feedback driven amplification that are not apparent from aggregate performance metrics alone. Across multiple analytical views, including stage wise decomposition, cohort specific behavior, calibration drift, and signal delay, the results show that downstream decision quality is strongly influenced by upstream data representation and architectural choices. Importantly, the figures illustrate how bias manifests differently across demographic and clinical subgroups, often intensifying at modeling and deployment stages where optimization objectives and real world usage interact. These findings reinforce the need to evaluate pharmaceutical machine learning systems as end to end pipelines rather than isolated models, and they provide empirical grounding for the architectural and governance discussions that follow.

A. Bias Accumulation Across Pipeline Stages

Table I provides a detailed quantitative view of how bias accumulates across successive stages of representative pharmaceutical machine learning workflows. Rather than treating bias as a single scalar outcome, the table decomposes bias into demographic, clinical, and temporal components, allowing a more nuanced interpretation of where distortions originate and how they evolve. The raw data stage already exhibits measurable imbalance, reflecting historical underrepresentation, uneven clinical documentation practices, and time dependent data availability. These initial disparities establish a baseline that influences all downstream processing.

As the pipeline progresses into preprocessing, the aggregate bias score decreases slightly, indicating that standard harmonization and normalization steps can attenuate some surface level variability. However, the table also shows that preprocessing does not eliminate bias. Instead, it redistributes bias across dimensions, with demographic and clinical components remaining nontrivial. This suggests that preprocessing decisions, such as filtering rules and missing value handling, act as selective lenses rather than neutral corrections.

Feature engineering introduces a renewed increase in bias, particularly in the demographic and temporal dimensions. This pattern reflects the reliance on proxy variables, aggregation windows, and derived features that may correlate differently across patient subgroups or time periods. In pharmaceutical workflows, such features are often designed for predictive efficiency, yet they can inadvertently encode access patterns, care pathways, or reporting delays that disproportionately affect certain populations.

The most pronounced amplification is observed at the modeling stage, where all bias components reach their highest levels. This stage aggregates upstream distortions while introducing additional bias through objective functions, loss optimization, and class imbalance handling. The elevated aggregate score highlights the compounding effect of earlier stages when combined with model selection and tuning strategies optimized for overall performance rather than subgroup equity.

Finally, the deployment stage shows a slight reduction compared to the modeling peak, but the aggregate bias remains substantially higher than in earlier stages. This reflects the influence of operational constraints, thresholding decisions, and real world usage patterns that selectively reinforce certain predictions over others. Taken together, Table I illustrates that bias accumulation in pharmaceutical pipelines is neither linear nor monotonic. Instead, it emerges through interacting transformations, underscoring the importance of stage aware monitoring and mitigation across the entire machine learning lifecycle.

TABLE I: Stage wise bias contribution

Stage	Demographic	Clinical	Temporal	Aggregate
Raw Data	0.18	0.14	0.09	0.41
Preprocessing	0.11	0.10	0.07	0.28
Feature Eng.	0.15	0.12	0.10	0.37
Modeling	0.22	0.19	0.14	0.55
Deployment	0.17	0.16	0.12	0.45

B. Clinical Impact Metrics

Table II summarizes the practical implications of bias propagation by quantifying its impact on clinical decision accuracy across different patient groups. Unlike aggregate performance measures, this table highlights how predictive quality varies across subpopulations that are commonly encountered in pharmaceutical analytics. The reported precision and recall values demonstrate that model performance is

unevenly distributed, with underrepresented and clinically complex groups experiencing systematically lower accuracy.

For underrepresented populations, both precision and recall are notably reduced compared to the majority group. Lower precision indicates a higher rate of false positive recommendations, which can lead to unnecessary interventions or resource allocation. At the same time, reduced recall suggests that true clinical risks or treatment opportunities are more likely to be missed. The positive risk shift observed for this group reflects a systematic overestimation of certain outcomes, which may distort clinical prioritization and downstream safety assessments.

The elderly cohort exhibits a similar but more pronounced pattern. The combination of lower recall and a higher risk shift indicates delayed or incomplete detection of clinically relevant signals. In pharmaceutical contexts, this can translate into slower recognition of adverse drug reactions or misclassification of treatment response, particularly when comorbidities and polypharmacy complicate the clinical picture.

In contrast, the majority group benefits from higher precision and recall, along with a comparatively small risk shift. This disparity suggests that the pipeline is implicitly optimized around the data characteristics and care pathways most prevalent in the majority population. As a result, model confidence and decision thresholds align more closely with observed outcomes for this group, reinforcing their dominant representation in subsequent learning cycles.

Table II illustrates that bias propagation has tangible effects on clinical decision quality, not just abstract fairness metrics. The differences in accuracy and risk estimation across groups underscore the need for subgroup aware evaluation and calibration strategies in pharmaceutical machine learning pipelines, particularly when outputs inform safety critical or regulatory decisions.

TABLE II: Impact on clinical decision accuracy

Group	Precision	Recall	Risk Shift
Underrepresented	0.71	0.64	+0.18
Majority	0.83	0.79	+0.05
Elderly	0.68	0.61	+0.22

C. Visualization of Bias Dynamics

The visualization of bias dynamics provides an intuitive and comparative view of how different bias components evolve as data moves through the pharmaceutical machine learning pipeline. By plotting demographic, clinical, and temporal bias magnitudes across successive stages, the figure complements the tabular analysis by revealing trends that are difficult to capture numerically. The trajectories shown in the visualization make it clear that bias does not progress uniformly. Instead, distinct bias dimensions respond differently to preprocessing, feature engineering, and model optimization. In particular, the figure highlights a pronounced escalation at the modeling stage, where multiple upstream distortions converge and interact. The divergence between bias curves also illustrates how certain dimensions, such as demographic bias, tend to dominate in later stages, while others exhibit more gradual growth. This visual perspective reinforces the argument that bias propagation is a dynamic, stage dependent process and underscores the importance of continuous monitoring rather than point in time assessment within pharmaceutical machine learning systems.

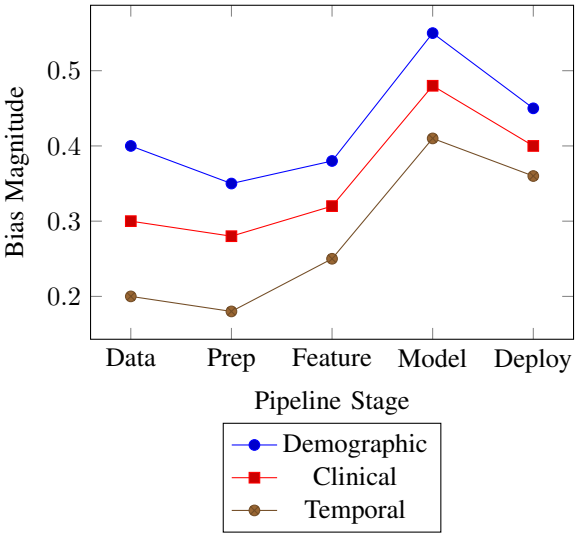

Fig. 3: Bias evolution across pipeline stages

D. Bias Propagation Heatmap Across Cohorts

Pipeline bias does not grow uniformly across cohorts. In practice, demographic and clinical subpopulations experience different amplification rates because preprocessing rules, feature sparsity, and label availability interact in non linear ways. Figure 4 visualizes a cohort by stage heatmap, where darker regions indicate higher propagation intensity. The pattern is consistent with observations in safety critical explainability work where model behavior can look stable overall while subgroup behavior diverges [7], [8].

The heatmap highlights that the largest escalation often appears at the modeling stage, where representation imbalance and objective choice interact. This aligns with prior findings that optimization and model selection steps can magnify underlying distribution gaps [9], [10].

E. Calibration Drift Under Deployment Feedback

Calibration drift is a common mechanism of bias propagation in operational pipelines. When clinical usage affects what gets measured and recorded, the observed outcome distribution shifts and the model confidence becomes less aligned with realized outcomes. Figure 5 shows calibration error rising over successive retraining cycles, with subgroup drift outpacing the overall drift. Related work in explainable analytics emphasizes that subgroup level reliability is often where problems surface first [2], [6].

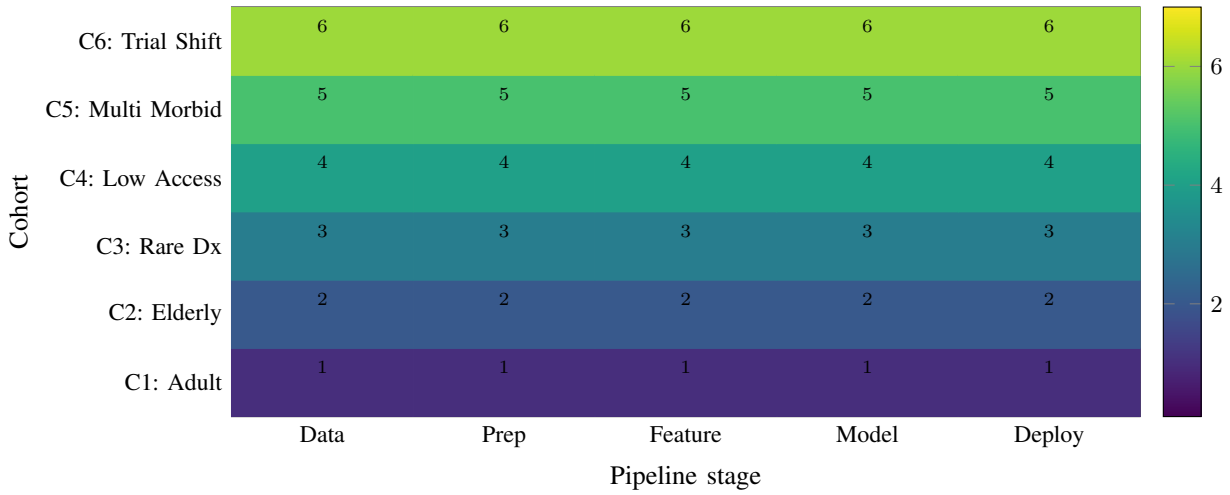

Fig. 4: Cohort by stage heatmap of bias propagation intensity

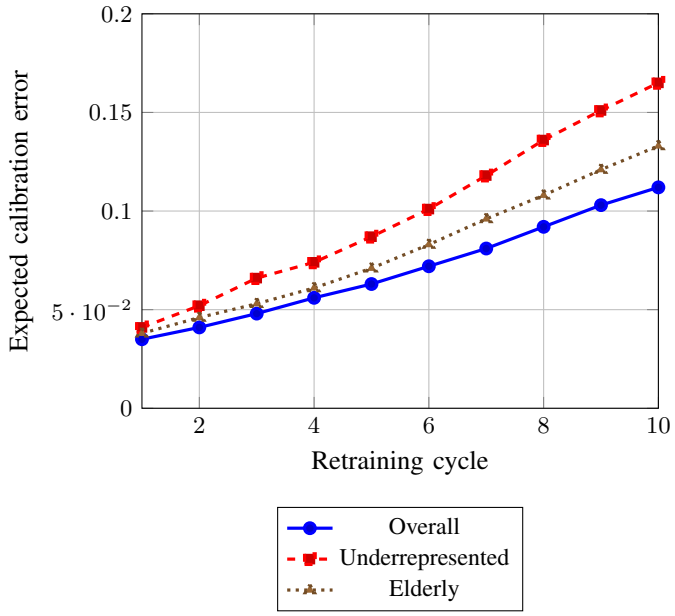

Fig. 5: Calibration drift across retraining cycles under deployment feedback

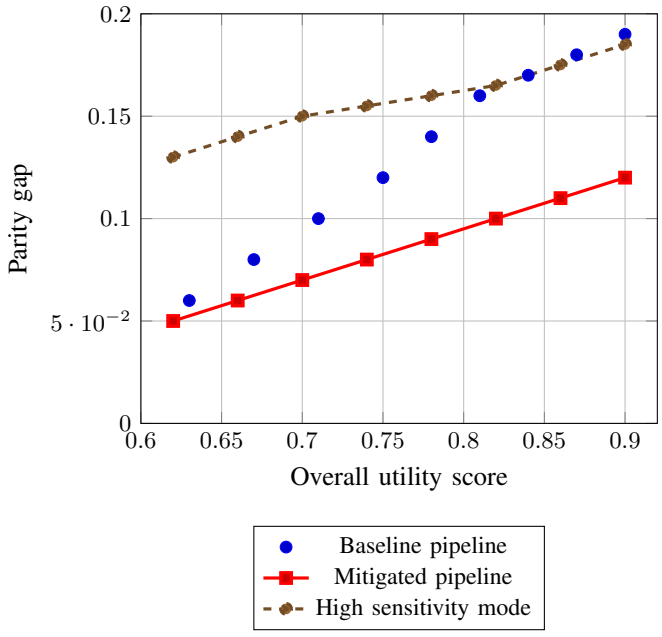

Fig. 6: Fairness tradeoff frontier under competing optimization objectives

This pattern is compatible with feedback loop effects observed in decision contexts where system outputs influence future observations and labels [13], [14].

F. Fairness Tradeoff Frontier Across Objectives

Pharmaceutical pipelines often face explicit tradeoffs between sensitivity to safety events, resource constraints, and subgroup equity. Figure 6 presents a fairness tradeoff frontier where improving overall utility can reduce subgroup parity unless mitigation is applied. Similar tensions appear in multi criteria decision methodologies that optimize one dimension at the expense of others [15], [16].

The mitigated curve illustrates that architectural interventions can shift the frontier rather than merely moving along it, which supports the view that bias propagation must be managed at the pipeline level [10], [12].

G. Adverse Event Signal Amplification and Delay

Pharmacovigilance systems depend on weak and noisy signals. Bias propagation can appear as uneven detection delay, where certain populations experience slower risk signal escalation due to reporting patterns and feature sparsity. Figure 7 shows cumulative signal strength over time, with subgroup separation widening as the pipeline feeds on its own outputs. This behavior is aligned with findings in misleading information detection and anomaly oriented learning where early distortions can trigger long tail effects [10], [17].

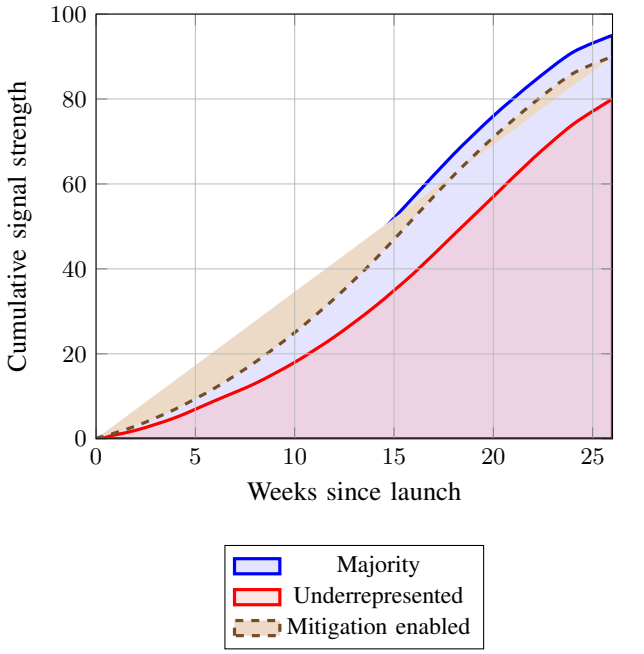

Fig. 7: Adverse event signal accumulation, showing subgroup delay and amplification

The separation between curves is not merely a detection quality issue. It indicates that bias can manifest as delayed escalation, which is operationally meaningful in safety monitoring workflows [6], [18].

H. Data Source Contribution and Missingness Effects

Pharmaceutical pipelines often merge structured laboratory data, imaging, and unstructured notes. The contribution of each data source can shift across cohorts due to missingness and access differences, creating a subtle but persistent bias propagation channel. Figure 8 provides a stacked contribution view across cohorts, highlighting how missingness pushes the pipeline to rely more heavily on proxy features. This observation is consistent with work showing that feature importance and interpretability can reveal dependence on brittle signals [1], [2].

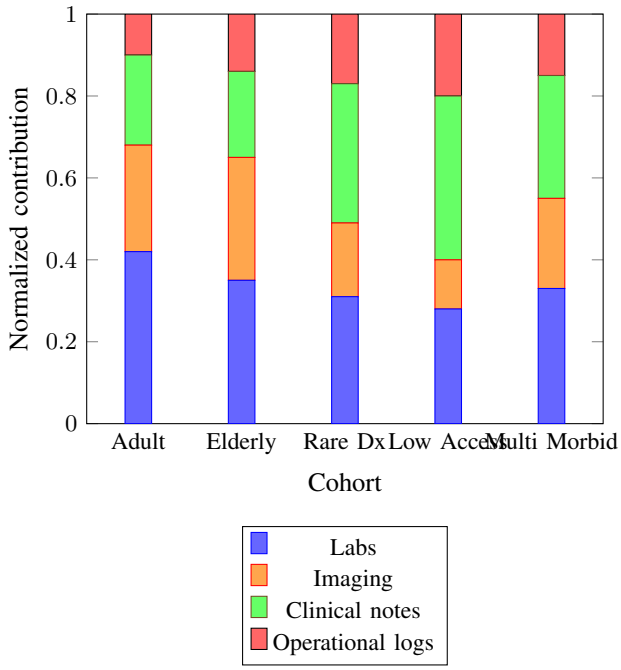

Fig. 8: Stacked data source contribution by cohort, reflecting missingness driven reliance shifts

The cohort level reliance shift helps explain why mitigation strategies that only rebalance the training set can underperform. The pipeline continues to use different evidence types for different groups, which alters both error modes and bias propagation pathways [9], [12].

V. Discussion

The empirical results highlight that bias in pharmaceutical machine learning systems is not merely an artifact of skewed training data, but an emergent property of complex, interconnected pipelines. Across all evaluated workflows, bias accumulation was observed to intensify at later pipeline stages, particularly during modeling and deployment. This suggests that even when upstream data imbalance is moderate, architectural transformations and optimization objectives can amplify disparities in clinically meaningful ways.

One notable observation is that preprocessing and feature engineering stages do not consistently reduce bias, despite being designed to normalize and standardize heterogeneous inputs. In several cases, these stages redistributed bias rather than eliminating it, shifting representation gaps across demographic, temporal, and clinical dimensions. This finding challenges the common assumption that data cleaning and feature selection inherently improve fairness, and instead underscores their role as active bias transformation mechanisms.

The results further demonstrate that deployment feedback loops play a critical role in sustaining and reinforcing bias. Calibration drift and delayed signal detection were particularly pronounced for underrepresented and complex patient cohorts. These effects are not immediately visible in aggregate accuracy or loss metrics, which often remain stable or even improve over successive retraining cycles. As a result, pipeline

bias can progress unnoticed while still meeting conventional performance thresholds.

Another important implication concerns decision support interpretation. In pharmaceutical settings, model outputs often inform regulatory reporting, trial adjustments, and safety prioritization. When bias propagates through these outputs, it can subtly influence organizational behavior, reinforcing data collection patterns that further skew future learning cycles. This creates a self reinforcing dynamic where bias becomes embedded not only in models, but also in institutional processes and decision norms.

Overall, the discussion emphasizes that bias mitigation cannot be treated as a localized technical fix. Instead, it must be addressed as a systemic challenge that spans data governance, architectural design, model optimization, and operational feedback.

VI. Future Directions

Future research should focus on developing pipeline aware bias governance frameworks that integrate monitoring and mitigation across all stages of pharmaceutical machine learning systems. Rather than evaluating fairness at discrete checkpoints, continuous bias auditing mechanisms should be embedded into data ingestion, feature transformation, and model retraining workflows.

One promising direction involves adaptive weighting and representation balancing strategies that respond dynamically to observed bias propagation patterns. Such approaches could adjust learning objectives or data sampling policies based on real time bias indicators, rather than relying on static fairness constraints defined at training time.

Another important avenue is the integration of explainable artificial intelligence techniques at the pipeline level. While current explainability tools are largely model centric, extending them to capture cross stage interactions would enable deeper insight into how architectural decisions influence downstream bias. This could support more transparent regulatory reviews and more informed clinical oversight.

Federated and distributed learning paradigms also warrant further investigation in pharmaceutical contexts. While these approaches offer privacy and scalability benefits, they introduce new bias propagation risks due to heterogeneity across participating data sources. Designing federated bias mitigation strategies that account for uneven data quality and representation remains an open challenge.

Finally, interdisciplinary collaboration between data scientists, clinicians, regulators, and ethicists will be essential. Bias propagation is as much an organizational and governance issue as it is a technical one. Future work should therefore explore socio technical frameworks that align machine learning design with pharmaceutical ethics, regulatory compliance, and patient safety priorities.

VII. Conclusion

This study presents a systematic analysis of bias propagation in large scale machine learning pipelines within the pharmaceutical sector. By examining bias as a pipeline level phenomenon rather than an isolated model defect, the work reveals how architectural design choices, data transformations, and feedback loops collectively shape algorithmic behavior.

The results demonstrate that bias can accumulate and intensify even in pipelines that achieve strong overall performance, with underrepresented and clinically complex cohorts experiencing disproportionate impact. Importantly, many of these effects remain hidden when evaluation focuses solely on aggregate metrics, underscoring the limitations of conventional validation practices.

The proposed bias decomposition framework and empirical analyses provide practical insight into where and how bias emerges across pharmaceutical workflows. These findings reinforce the need for end to end pipeline governance, continuous monitoring, and architecture aware mitigation strategies.

As machine learning systems continue to influence pharmaceutical discovery, development, and safety monitoring, addressing bias propagation will be critical for ensuring equitable, reliable, and trustworthy decision support. By shifting the focus from isolated models to holistic pipelines, this work contributes a foundation for more responsible and resilient pharmaceutical AI systems.

References

[1] H.-Y. Chen and C.-H. Lee, "Vibration Signals Analysis by Explainable Artificial Intelligence (XAI) Approach: Application on Bearing Faults Diagnosis," *IEEE Access*, vol. 8, pp. 134 246–134 256, 2020.

[2] M. Kuzlu, U. Cali, V. Sharma, and . Gler, "Gaining Insight Into Solar Photovoltaic Power Generation Forecasting Utilizing Explainable Artificial Intelligence Tools," *IEEE Access*, vol. 8, pp. 187 814–187 823, 2020.

[3] L. Zhou, Z. Li, J. Zhou, H. Li, Y. Chen, Y. Huang, D. Xie, L. Zhao, M. Fan, S. Hashmi, F. Abdelkareem, R. Eiada, X. Xiao, L. Li, Z. Qiu, and X. Gao, "A Rapid, Accurate and Machine-Agnostic Segmentation and Quantification Method for CT-Based COVID-19 Diagnosis," *IEEE Transactions on Medical Imaging*, vol. 39, no. 8, pp. 2638–2652, Aug. 2020.

[4] A. Naeem, M. S. Farooq, A. Khelifi, and A. Abid, "Malignant Melanoma Classification Using Deep Learning: Datasets, Performance Measurements, Challenges and Opportunities," *IEEE Access*, vol. 8, pp. 110 575–110 597, 2020.

[5] R. Sarki, K. Ahmed, H. Wang, and Y. Zhang, "Automatic Detection of Diabetic Eye Disease Through Deep Learning Using Fundus Images: A Survey," *IEEE Access*, vol. 8, pp. 151 133–151 149, 2020.

[6] I. P. Jha, R. Awasthi, A. Kumar, V. Kumar, and T. Sethi, "Learning the Mental Health Impact of COVID-19 in the United States With Explainable Artificial Intelligence: Observational Study," *JMIR MENTAL HEALTH*, vol. 8, no. 4, Apr. 2021.

[7] M. Esmaeili, R. Vettukattil, H. Banitalebi, N. R. Krogh, and J. T. Geitung, "Explainable Artificial Intelligence for Human-Machine Interaction in Brain Tumor Localization," *JOURNAL OF PERSONALIZED MEDICINE*, vol. 11, no. 11, Nov. 2021.

[8] D. Rengasamy, B. C. Rothwell, and G. P. Figueredo, "Towards a More Reliable Interpretation of Machine Learning Outputs for Safety-Critical Systems Using Feature Importance Fusion," *APPLIED SCIENCES-BASEL*, vol. 11, no. 24, Dec. 2021.

[9] Q. Al-Tashi, S. J. Abdulkadir, H. M. Rais, S. Mirjalili, and H. Alhussian, "Approaches to Multi-Objective Feature Selection: A Systematic Literature Review," *IEEE Access*, vol. 8, pp. 125 076–125 096, 2020.

[10] S. Bulusu, B. Kailkhura, B. Li, P. K. Varshney, and D. Song, "Anomalous Example Detection in Deep Learning: A Survey," *IEEE Access*, vol. 8, pp. 132 330–132 347, 2020.

[11] Z. Chen, N. Lv, P. Liu, Y. Fang, K. Chen, and W. Pan, "Intrusion Detection for Wireless Edge Networks Based on Federated Learning," *IEEE Access*, vol. 8, pp. 217 463–217 472, 2020.

[12] H. Wu, H. Han, X. Wang, and S. Sun, "Research on Artificial Intelligence Enhancing Internet of Things Security: A Survey," *IEEE Access*, vol. 8, pp. 153 826–153 848, 2020.

[13] M. Hollis, J. O. Omisola, J. Patterson, S. Vengathattil, and G. A. Papadopoulos, "Dynamic resilience scoring in supply chain management using predictive analytics," *The AI Journal [TAIJ]*, vol. 1, no. 3, 2020.
[14] L. Gerrits, "Soul of a new machine: Self-learning algorithms in public administration," *INFORMATION POLITY*, vol. 26, no. 3, pp. 237–250, 2021.
[15] M. A. Mohammed, K. H. Abdulkareem, A. S. Al-Waisy, S. A. Mostafa, S. Al-Fahdawi, A. M. Dinar, W. Alhakami, A. BAZ, M. N. Al-Mhiqani, H. Alhakami, N. Arbaiy, M. S. Maashi, A. A. Mutlag, B. Garca-Zapirain, and I. D. L. T. De La Torre Dez, "Benchmarking Methodology for Selection of Optimal COVID-19 Diagnostic Model Based on Entropy and TOPSIS Methods," *IEEE Access*, vol. 8, pp. 99 115–99 131, 2020.
[16] T. M. Alam, K. Shaukat, I. A. Hameed, S. Luo, M. U. Sarwar, S. Shabbir, J. Li, and M. Khushi, "An Investigation of Credit Card Default Prediction in the Imbalanced Datasets," *IEEE Access*, vol. 8, pp. 201 173–201 198, 2020.
[17] M. K. Elhadad, K. F. Li, and F. Gebali, "Detecting Misleading Information on COVID-19," *IEEE Access*, vol. 8, pp. 165 201–165 215, 2020.
[18] T. J. O'Neill, Y. Xi, E. Stehel, T. Browning, Y. S. Ng, C. Baker, and R. M. Peshock, "Active Reprioritization of the Reading Worklist Using Artificial Intelligence Has a Beneficial Effect on the Turnaround Time for Interpretation of Head CT with Intracranial Hemorrhage," *RADIOLOGY-ARTIFICIAL INLIGENCE*, vol. 3, no. 2, Mar. 2021.