Math. Appl. 15(2s)2026, 94–115

Algorithmic Bias and Data Privacy in Business Analytics
A Framework for Ethical Compliance in Automated Decision Systems

Mahesh Reddy Challa[1]

Abstract: *Businesses now use automated decision systems to approve credit, screen applicants, set prices and target offers, and these systems learn from data that carry the history of earlier human decisions and the personal information of the people concerned. Two ethical and legal problems follow: algorithms can reproduce or amplify discrimination, and the data they use can expose individuals. The two problems are usually handled by different teams, with different tools, and sometimes in tension, since measuring bias needs sensitive attributes that privacy rules restrict. This paper proposes a six-stage compliance framework that links scoping and legal mapping, data governance, model development, pre-deployment validation, monitoring and accountability, with decision gates and evidence at each stage. We test several of its controls on a simulated lending process with 30,000 applicants per run and five repetitions, in which historical decisions penalise one group although true repayment is the same in both groups. A model trained on those decisions and given the group attribute amplifies the disparity (disparate impact ratio 0.60, against 0.66 in the historical decisions). Removing the attribute raises the ratio only to 0.73 because a residential proxy predicts group membership with an AUC of 0.84; removing the proxy as well gives 0.79, reweighing 0.81, and group-specific thresholds about 0.98, with unchanged discrimination ability (AUC 0.77) and a modest rise in the default rate among approved applicants for the equal-opportunity variant (17.8% to 19.0%). Privacy tests show that overfitted tree ensembles are vulnerable to membership inference (attack AUC up to 0.86), while regularised models are close to the no-leakage value of 0.5 and generalise better. A differentially private audit, which releases only noisy group counts, estimates the disparate impact ratio almost as accurately as a non-private audit when at least 2,000 decisions are audited, even at a strong privacy setting. The data are synthetic and the legal discussion is general, so the results illustrate the framework and are not legal advice or evidence about any real system.*

Introduction

Business analytics has moved from describing what happened to deciding what happens next. Credit limits, insurance premiums, job shortlists, fraud holds, delivery priorities, dynamic prices and marketing offers are increasingly determined by models that score individuals without a human reviewing each case. The efficiency gains are real, but so are the risks. When a decision is made at scale by a statistical model, any systematic error is repeated thousands of times, often invisibly, and it falls on identifiable people.

[1]3151 Solana circle Denton Texas zip code 76203.

Keywords: algorithmic bias, data privacy, fairness, automated decision-making, business analytics, compliance, differential privacy, membership inference, algorithmic auditing.

Two kinds of risk dominate public and regulatory attention. The first is algorithmic bias: systematically worse outcomes for people in particular groups, whether defined by sex, race, age, disability, religion or another characteristic. Bias can arise even when no one intends it, because models learn from historical data that reflect past discrimination, unequal access to opportunity and uneven measurement [1], [2]. A well-known case in health care showed that an algorithm used to allocate care management gave lower risk scores to Black patients than to equally sick White patients because it used past health spending as a proxy for need [3]. The second is data privacy: the collection, inference and exposure of personal information. Automated systems require large amounts of personal data, can infer sensitive traits that individuals never disclosed, and can leak information about the people in their training data [4], [5].

These risks are usually managed separately. Fairness work is carried out by data scientists who measure disparities in model outputs. Privacy work is led by legal and security teams who ask about consent, retention and access. Yet the two are linked. Fairness assessment needs data on protected attributes, which privacy law treats as especially sensitive. Removing sensitive attributes to protect privacy may leave proxies that still encode them, so that the model remains biased while the organisation can no longer measure it. Techniques that strengthen privacy, such as noise addition, can reduce the accuracy of subgroup estimates. A compliance approach that treats the two in isolation will either miss these interactions or resolve them by accident.

Regulation is also converging on automated decisions. European data protection law sets principles for personal data processing, gives individuals rights regarding decisions based solely on automated processing, requires data protection by design and calls for impact assessments for high-risk processing [6]. The European Union's regulation on artificial intelligence takes a risk-based approach and imposes requirements on high-risk systems, a category that includes some uses in credit scoring and employment [7]. In the United States, anti-discrimination law in credit and employment applies to algorithmic decisions, and agencies have long used rules of thumb such as the four-fifths guideline to flag adverse impact. Voluntary frameworks, such as the risk management framework from the United States National Institute of Standards and Technology, provide structure for managing AI risk throughout the lifecycle [8]. Organisations therefore face a patchwork of obligations and a gap between legal principles and the concrete engineering choices that satisfy them.

This paper offers a bridge. It proposes a lifecycle compliance framework for automated decision systems in business analytics that treats bias and privacy together, specifies controls and evidence at each stage, and includes decision gates. It then tests several of the framework's technical controls in a controlled simulation of a lending process, where we can know the truth about both fairness and privacy.

The contributions are as follows.

1. A six-stage compliance framework that integrates fairness and privacy controls, with evidence artefacts and decision gates, mapped to major regulatory anchors.

2. An empirical comparison of bias mitigation options (feature removal, proxy removal, reweighing and group-specific thresholds) on a simulated lending process in which the cause of the disparity is known.
3. Privacy experiments on data minimisation, membership inference against different model types, and differentially private auditing of fairness metrics.
4. Practical guidance on the tensions between fairness and privacy, and on what organisations should document and monitor.

The paper does not give legal advice. Legal requirements vary by jurisdiction and sector and change over time, and the references to regulation illustrate how a framework can be anchored in them.

LITERATURE REVIEW

A. Sources of algorithmic bias

Bias in automated decisions can enter at several points. Surveys of the field distinguish historical bias, in which the world that generated the data was itself unfair; representation bias, in which some groups are under-represented; measurement bias, in which features or labels are poor or unequal proxies for what matters; aggregation bias, in which one model is used for groups with different relationships; and evaluation bias, in which test data or metrics do not represent the population served [1]. Legal scholarship stresses that data mining can produce discriminatory outcomes even without intent, through the choice of target variables, labelled examples, features and proxies, and that existing anti-discrimination doctrine struggles to address such cases [2]. The most common practical mechanism is the proxy: when a feature such as postal area is correlated with a protected attribute, a model can reproduce discrimination without ever seeing the attribute.

B. Defining and measuring fairness

Many formal definitions of fairness exist, and they capture different ideas. Group fairness criteria compare outcomes across groups. Demographic parity asks that approval rates be equal, and its practical ratio form, the disparate impact ratio, compares the lowest group rate with the highest and is the basis of the four-fifths guideline. Equality of opportunity asks that people who truly deserve a positive outcome, such as applicants who would repay, are equally likely to receive it in each group [9]. Predictive parity asks that a positive decision mean the same thing for each group. Individual fairness asks that similar individuals be treated similarly [10]. Counterfactual approaches ask whether a decision would change had the person belonged to another group, given a causal model [11].

A central result is that these criteria cannot generally be satisfied together. When base rates differ between groups, calibrated scores cannot also equalise error rates, except in trivial cases [12], [13]. Choosing a fairness criterion is therefore a normative and context-dependent decision and not a technical one, and the choice should be documented and justified. It also means that no single number can certify a system as fair.

C. Bias mitigation

Mitigation methods are commonly grouped by where they act. Pre-

processing methods change the training data, for example by reweighing examples so that group membership and outcome are statistically independent in the weighted data [14]. In-processing methods change the learning algorithm, for example by adding fairness constraints, as in reductions-based approaches [15]. Post-processing methods adjust the model's outputs, for example by choosing different decision thresholds for groups to equalise a chosen error rate [9]. Each has advantages and drawbacks. Pre-processing is model-agnostic but may be insufficient. In-processing can be effective but requires control over training. Post-processing is simple but uses the group attribute at decision time, which may be legally restricted in some domains and contexts. "Fairness through unawareness", simply omitting protected attributes, is easy to implement and widely recognised as inadequate because of proxies [10].

D. Privacy risks of analytics and machine learning

Privacy risk in analytics takes several forms. Direct identification arises when records contain names or identifiers. Re-identification arises when apparently anonymous records can be linked to other data; early work showed that a few attributes can be enough to single out a person in a large sparse dataset [4], and k-anonymity was proposed as a property that limits such linkage [16]. Inference arises when models predict sensitive traits from innocuous data. Finally, models themselves can leak: membership inference attacks test whether a particular record was used to train a model by exploiting the fact that models tend to behave differently on their training data than on new data [5]. Overfitted models are the most vulnerable, which links privacy to generalisation.

E. Differential privacy

Differential privacy provides a mathematical guarantee that the output of an analysis changes only slightly whether or not any single individual's data are included, with the strength of the guarantee controlled by a parameter, usually called epsilon, where smaller values mean stronger privacy [17], [18]. The basic mechanism for counting queries adds random noise calibrated to the maximum effect one person can have on the count. Differential privacy has been extended to training machine learning models, for example by clipping and noising gradient updates [19]. It offers a principled way to release statistics, including fairness audit statistics, without revealing individuals, but at a cost in accuracy that depends on the size of the data and the strength of protection.

F. Governance, documentation and auditing

Structured documentation supports accountability. Datasheets for datasets record the motivation, composition, collection and recommended uses of data [20], and model cards record intended use, performance across groups and limitations of a model [21]. Internal algorithmic auditing frameworks propose that audits span the whole lifecycle, from scoping to post-deployment monitoring, with documentation at each stage [22]. A review of AI ethics guidelines found convergence on principles such as transparency, fairness and privacy, but wide divergence in implementation [23], and the sociotechnical perspective adds that abstracting a social problem into a prediction task may itself omit what matters [24].

RELATED WORK

Fairness toolkits and metrics. The research literature provides many fairness definitions, metrics and mitigation algorithms [1], [9], [14], [15]. These are essential building blocks, but they are mostly presented as technical methods, with limited guidance on how to embed them in organisational processes with legal accountability. Our framework positions them as controls inside a lifecycle with defined evidence and decision gates.

Privacy-preserving machine learning. Work on differential privacy, membership inference and de-identification provides tools to quantify and limit privacy risk [4], [5], [16], [17], [18], [19]. Most studies address privacy and fairness separately. A line of research has noted that privacy mechanisms can have unequal effects on groups, but we do not pursue that here. Our experiments instead emphasise two practical interactions: the dependence of fairness measurement on sensitive data, and the possibility of auditing with privacy protection.

Auditing and documentation. End-to-end internal audit frameworks and documentation practices [20], [21], [22] are the closest in spirit to our framework. We contribute an explicit integration of privacy controls with fairness controls, an evidence-centred gate structure, and empirical illustrations of selected controls.

Regulatory frameworks. Data protection law, AI-specific regulation, sector anti-discrimination rules and voluntary risk management frameworks [6], [7], [8] define obligations and good practice. We do not interpret these in detail. We use them as anchors that indicate what a documented control should be able to demonstrate.

Gap. We are not aware of work that offers, in one place, a lifecycle framework that integrates fairness and privacy, tests representative controls on a setting with known ground truth for both, and quantifies the cost of privacy-preserving fairness auditing. This paper aims to contribute such a combination, with the limits of synthetic data stated openly.

METHODOLOGY

A. The compliance framework

The framework in Fig. 1 has six stages. Each stage has key questions, controls and evidence artefacts, and ends with a gate at which a named accountable person signs off or sends the work back. Table 1 summarises the stages.

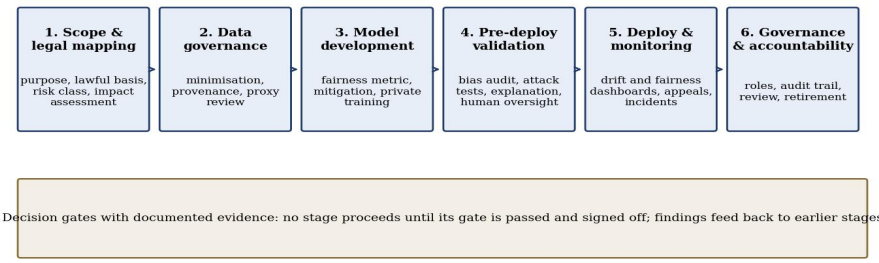


Fig. 1. Six-stage framework for ethical compliance in automated decision systems.

Table 1. Stages, controls, evidence and illustrative regulatory anchors.

Stage	Key questions	Controls	Evidence	Illustrative anchors
1. Scope and legal mapping	What decision is automated, who is affected, what is the lawful basis and the risk class?	Purpose statement; risk classification; impact assessment; human oversight design	Impact assessment; decision inventory entry	Data protection impact assessment; risk class under AI regulation; sector credit and employment rules
2. Data governance	Is each attribute needed, how was it collected, and could it act as a proxy?	Data minimisation; provenance record; proxy review; retention limits; access control	Dataset documentation; feature justification table	Purpose limitation and minimisation; data protection by design; special-category data limits
3. Model development	Which fairness criterion fits the context, and how are privacy and accuracy balanced?	Fairness metric choice with rationale; mitigation; regularisation or private training; reproducible pipeline	Model documentation; experiment log	Risk management and documentation duties for high-risk systems
4. Pre-deployment validation	Does the model meet the thresholds for fairness, privacy and explanation?	Subgroup bias audit; privacy attack test; reason codes; challenge testing	Validation report; sign-off	Non-discrimination; adverse-action explanations; rights regarding automated decisions
5. Deployment and monitoring	Do fairness and privacy hold in production as data drift?	Fairness and drift dashboards; private auditing; appeals channel; incident response	Monitoring reports; appeal logs	Post-market monitoring; accountability for outcomes
6. Governance and accountability	Who is responsible, and how are lessons fed back?	Roles and independent review; audit trail; periodic re-approval; retirement criteria	Governance charter; audit trail	Organisational accountability; voluntary risk frameworks

The framework has four design principles: bias and privacy are assessed together; evidence is produced as a by-product of work; gates are explicit, so nothing is deployed without documented review; and depth of control scales with the risk class of the decision.

B. Technical controls tested in this study

We test technical controls from stages 2, 3, 4 and 5 in a controlled setting. *Feature removal and proxy review (stage 2).* We compare a model that includes the group attribute, one that omits it ("unawareness"), and one that also omits a feature identified as a proxy. We also test how well the group can be predicted from the remaining features, which measures proxy leakage.

Bias mitigation (stage 3). We compare reweighing, which weights training examples so that group and historical outcome are independent, and two post-processing methods that set group-specific thresholds: one that equalises approval rates and one that equalises the approval rate among applicants who would repay.
Data minimisation (stage 2). We remove features progressively, starting with those that carry no signal, and measure effects on predictive ability and fairness.
Privacy attack testing (stage 4). We run a membership inference attack that labels a record as a training member if the model's loss on it is low, for four model types, and report the attack's area under the ROC curve, in which 0.5 means no leakage and 1.0 means perfect identification.
Private fairness auditing (stage 5). We simulate an audit in which a decision system's approvals are counted by group, the counts are released with noise that satisfies differential privacy, and the disparate impact ratio is estimated from the noisy counts. We compare the estimate with the non-private one for several privacy strengths and audit sizes.

C. Metrics

We report the following measures.

- *Approval rate* by group, and the *disparate impact ratio*, the lower group approval rate divided by the higher. A value of one indicates equal rates and the four-fifths guideline treats values below 0.8 as a warning sign.
- *Equal opportunity gap*: the difference between the approval rates of the two groups among applicants who would in fact repay. A value of zero indicates equal treatment of equally creditworthy people.
- *Default rate among approved applicants*: the share of approved applicants who would not repay, as a measure of business cost.
- *Discrimination ability*: the area under the ROC curve for predicting actual repayment.
- *Membership inference attack AUC*, as above.
- *Audit error*: the mean absolute difference between the estimated and true disparate impact ratio.

EXPERIMENTAL SETUP

A. Why a simulation

Testing fairness controls requires knowing what a fair decision would have been, which is never the case in real lending data, where outcomes are observed only for approved applicants and the true creditworthiness of the rejected is unknown. A simulation lets us build a world in which the cause of disparity is known and the truth is available. Privacy experiments likewise benefit from knowing exactly which records are members. The price is realism, which we discuss under limitations. No personal data were used.

B. Simulated lending process

Each run generates 30,000 applicants, of whom 30% belong to group 1 and 70% to group 0. Each applicant has a latent creditworthiness drawn from the same distribution in both groups, and true repayment depends only on

this latent quantity. Repayment is therefore equally likely in the two groups (66.9% and 66.5% in the first run). Applicants have ten features: income, credit history, debt ratio, employment, education, a residential region indicator and four pure-noise features. Because of structural disadvantage, the measured income of group 1 is lower by about a third of a standard deviation for the same creditworthiness, credit history and education are slightly lower, and the debt ratio slightly higher. The region indicator matches the group for 80% of applicants, so it is a proxy.

Historical decisions were taken by a process that scored applicants on latent creditworthiness with noise and subtracted a penalty for group 1, and approved the top 55%. These historical decisions are the training labels, which is the common situation in which models learn from past human decisions. They approve 61.4% of group 0 and 40.6% of group 1, a disparate impact ratio of 0.66 (the means over five runs).

C. Modelling and evaluation

The main model is a logistic regression trained on 60% of the data to predict the historical decision, with 20% used as a validation set to set thresholds and 20% for testing. Performance on fairness and discrimination is evaluated against true repayment, not against historical decisions. Results are averaged over five random repetitions of the simulation, and we report standard deviations for the main fairness measures. The equal-opportunity thresholds are set using validation outcomes, which assumes that repayment outcomes are observed for a validation sample, an assumption that does not hold when outcomes are observed only for approved applicants.

For privacy experiments, models are trained on 500, 2,000 or 8,000 records with an equal number of non-members, using unpruned random forests, deep gradient boosting, regularised random forests and heavily regularised logistic regression. Attack AUC is averaged over three repetitions. For the private audit, a sample of 500, 2,000 or 10,000 decisions is audited, the number of approvals per group is released with Laplace noise (the standard mechanism for counts, with noise scaled to the effect of one person), group sizes are treated as public, and the estimate is repeated 400 times for each setting.

D. Software

All experiments used Python with NumPy, pandas and scikit-learn, with fixed random seeds [25]; the complete experiment runs in about two minutes on one processor.

RESULTS

A. Bias propagation and mitigation

Table 2 reports fairness and business measures for the historical decisions and for each model, averaged over five runs, and Fig. 2 plots the disparate impact ratio.

Table 2. Fairness and business measures (mean over 5 runs; standard deviation of the disparate impact ratio in parentheses).

Method	Approval, group 0	Approval, group 1	Disparate impact ratio	Equal opportunity gap	Default rate among approved	AUC for repayment
Historical decisions	0.614	0.406	0.662 (0.017)	0.225	0.177	–
Model with group attribute	0.652	0.391	0.599 (0.013)	0.275	0.184	0.759
Unawareness (group removed)	0.624	0.453	0.726 (0.018)	0.181	0.178	0.766
Unawareness, proxy removed	0.610	0.484	0.794 (0.021)	0.132	0.178	0.769
Reweighing	0.607	0.490	0.807 (0.019)	0.123	0.178	0.769
Group thresholds, equal approval rates	0.572	0.565	0.976 (0.007)	0.014	0.176	0.766
Group thresholds, equal opportunity	0.624	0.617	0.980 (0.012)	0.016	0.190	0.766

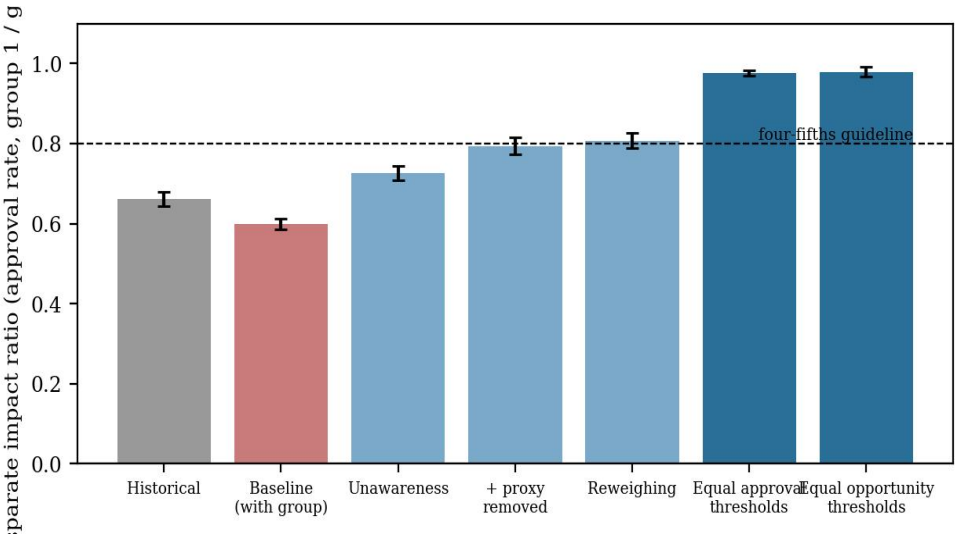


Fig. 2. Disparate impact ratio by method (error bars: standard deviation over 5 runs).

Three results stand out. First, the model that learned from biased decisions and was allowed to use the group attribute does not merely reproduce the disparity but amplifies it: the disparate impact ratio falls from 0.66 in the historical decisions to 0.60, and the equal opportunity gap rises from 0.225 to 0.275. The model faithfully learns the penalty that historical decision makers applied. Its discrimination ability for actual repayment (AUC 0.759) is lower than that of the unaware models, which shows that the learned penalty is noise with respect to the truth.

Second, removing the group attribute helps only partly. The ratio improves to 0.73, which is still below the four-fifths guideline, because the other features carry the group's imprint. The residential region indicator has a negative weight in the unaware model, and a separate model can predict group membership from the unaware features with an AUC of 0.84. When the region indicator is also removed, this falls to 0.60 and the ratio rises to

0.79. The remaining gap arises because income, credit history and education are lower for group 1 for the same creditworthiness. Those features are legitimate predictors of repayment in the data but they are also carriers of structural disadvantage, and a model that relies on them will still approve fewer members of group 1.

Third, active mitigation can close the gap. Reweighing, which needs the group attribute only at training time, brings the ratio to 0.81, just above the guideline, and the equal opportunity gap to 0.12. Group-specific thresholds bring the ratio to about 0.98 and the gap to about 0.015. The cost in this simulation is small. The discrimination ability is unchanged at about 0.77, the default rate among approved applicants stays at 17.6% for equal approval rates and rises from 17.8% to 19.0% for equal opportunity thresholds, in which the approval rate for both groups is higher than under equal approval rates. The accuracy of approve-versus-repay decisions, measured against true repayment, is slightly higher for the equal opportunity thresholds (0.714) than for the unaware model (0.697), because the biased historical labels had led the model to wrongly reject creditworthy members of group 1.

B. Interpreting the mitigation results

These results should be read carefully, since the simulation is constructed so that the two groups are equally creditworthy. In such a world, correcting disparities moves decisions closer to the truth, and there is little conflict between fairness and accuracy. In real settings, base rates may differ between groups for reasons that are themselves the result of historical disadvantage, and the impossibility results mentioned earlier apply [12], [13]. The choice among criteria then involves value judgements and trade-offs that the data cannot settle.

A second caution concerns legality. Group-specific thresholds use the protected attribute at decision time. In some domains, such as consumer credit in the United States, explicit use of protected characteristics in decisions is restricted, and in the European Union the use of special categories of data is limited. A method that is statistically effective may therefore be unlawful or require a specific legal basis. Reweighing, which uses the attribute only during training, and proxy review, which removes problematic inputs, are more likely to be acceptable, but each case needs legal review. This is a reason why the framework starts with legal mapping at the first stage and requires validation of the chosen mitigation at the fourth.

C. Data minimisation

Table 3 shows the effect of progressively removing features from the unaware model.

Table 3. Effect of data minimisation (single run).

Features retained	AUC for repayment	Disparate impact ratio	Equal opportunity gap
10 (all)	0.775	0.727	0.172
6 (noise features removed)	0.775	0.728	0.174
5	0.773	0.721	0.177
4	0.772	0.729	0.175
3	0.768	0.734	0.170
2 (income and region only)	0.747	0.722	0.178

Removing the four noise features changes nothing: predictive ability and fairness are identical, and the organisation collects and stores less personal data. Removal of more informative features reduces predictive ability slowly, from 0.775 to 0.768 with three features, and to 0.747 with two. The fairness measures barely move, because minimisation as such is not aimed at bias. This illustrates both the value and the limit of data minimisation: it is a privacy control that is often free of cost when features carry no signal, and it should be justified feature by feature, but it is not a substitute for bias mitigation.

D. Membership inference

Table 4 and Fig. 3 report the attack AUC against four model types.

Table 4. Membership inference attack AUC by model type and training size, and discrimination ability (training size 2,000).

Model	500 records	2,000 records	8,000 records	Train AUC	Test AUC
Unpruned random forest	0.817	0.813	0.818	1.000	0.748
Deep gradient boosting	0.860	0.782	0.672	1.000	0.726
Regularised random forest	0.520	0.524	0.514	0.820	0.762
Regularised logistic regression	0.505	0.504	0.502	0.777	0.768

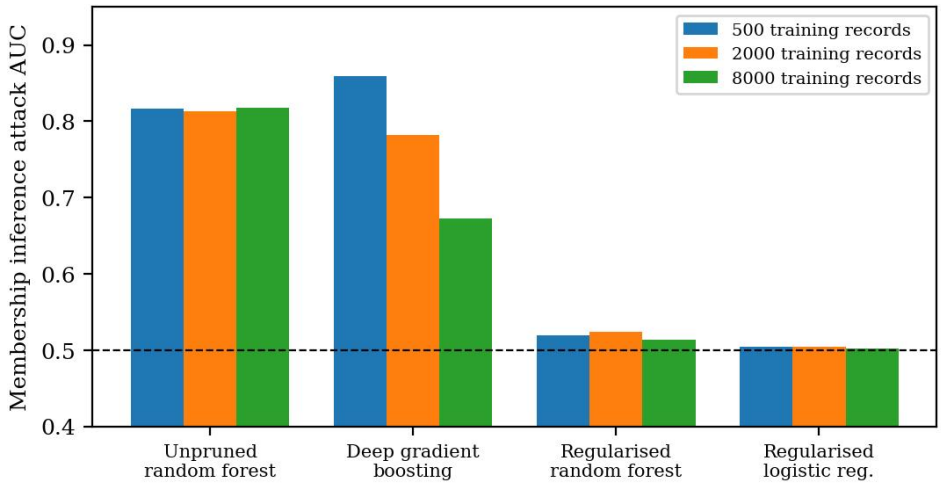


Fig. 3. Membership inference attack AUC by model type and training size (dashed line: no leakage).

The overfitted models leak. A simple attack that flags low-loss records as members reaches an AUC of 0.82 against the unpruned random forest and 0.86 against deep gradient boosting trained on 500 records. For the boosted model, leakage declines as the training set grows (0.78 at 2,000 and 0.67 at 8,000), since with more data the model overfits each record less. The regularised random forest and the regularised logistic regression are close to the no-leakage value of 0.5. These regularised models also generalise better: their test AUC (0.762 and 0.768) exceeds that of the overfitted models (0.748 and 0.726), so that in this setting privacy and accuracy point in the same direction. This need not hold for all data, but it

suggests that regularisation and capacity control are low-cost privacy controls that should be part of model development. The attack used is among the simplest, so a low attack AUC is not a guarantee of privacy; formal guarantees such as differential privacy are stronger.

E. Private fairness auditing

Table 5 and Fig. 4 show the accuracy of the disparate impact estimate when audit counts are released under differential privacy. The true ratio for the audited model is 0.60.

Table 5. Private audit of the disparate impact ratio: mean absolute error (and share of audits that flag a violation of the four-fifths guideline).

Audited decisions	No privacy	Epsilon 5	Epsilon 1	Epsilon 0.5	Epsilon 0.1
500	0.056 (99.5%)	0.050 (99.8%)	0.056 (99.8%)	0.058 (99.5%)	0.120 (90.5%)
2,000	0.026 (100%)	0.026 (100%)	0.026 (100%)	0.027 (100%)	0.039 (100%)
10,000	0.011 (100%)	0.012 (100%)	0.011 (100%)	0.012 (100%)	0.014 (100%)

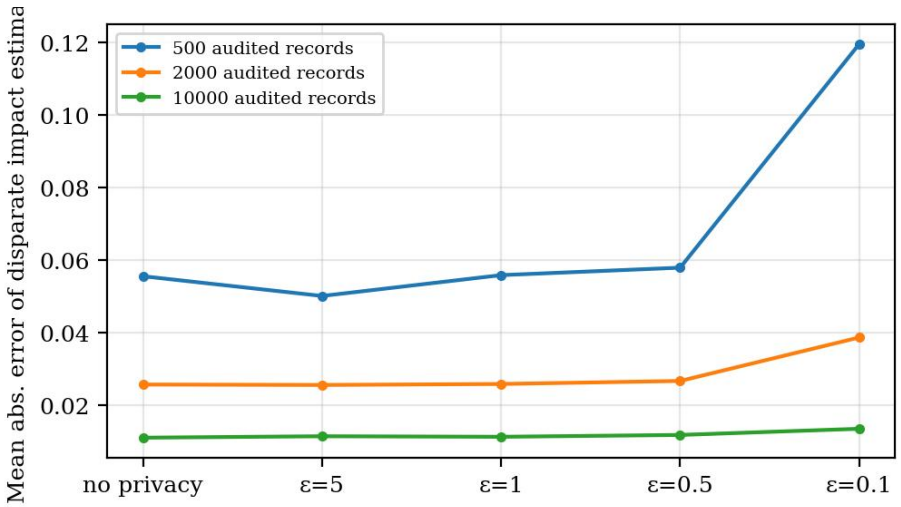


Fig. 4. Error of the disparate impact estimate under differential privacy.

The cost of privacy-preserving auditing is small unless the audit sample is small and protection is very strong. With 2,000 or more audited decisions, the added error is at most 0.013 even at the strongest setting tested, and the audit detects the violation every time. With 500 decisions and a strong setting, the error more than doubles (0.12 against 0.056) and one in ten audits fails to flag the violation. Because random noise in counts is of fixed size, its relative effect shrinks as the audited sample grows. Two limitations apply. The violation here is far from the threshold, with a true ratio of 0.60 against a guideline of 0.8, and a ratio close to the threshold would require larger samples to classify reliably. And the experiment protects only the released counts; it does not address who holds the raw sensitive attributes, which must still be governed.

The result matters for the apparent conflict between fairness and privacy. It suggests that organisations can monitor group disparities in production while releasing only noisy aggregate statistics to the monitoring team, with

the raw sensitive data kept under restricted access, at a modest cost in precision for audits of reasonable size.

DISCUSSION

A. Interpretation

Four messages emerge from the framework and the experiments. First, models trained on historical decisions inherit and may amplify their bias, and simple measures such as removing the protected attribute are insufficient because of proxies. Second, bias can be reduced substantially by documented, testable controls (proxy review, reweighing, threshold adjustment), but the choice among them involves legal constraints and value judgements as well as numbers. Third, several privacy controls are cheap: removing features that carry no signal costs nothing, and regularised models resist simple membership inference while generalising better. Fourth, the tension between fairness measurement and privacy can be managed technically, since private auditing yields accurate fairness statistics at reasonable audit sizes.

B. Practical implications

Start with the decision, not the model. Stage 1 forces the organisation to state what is decided, who is affected, on what lawful basis and in which risk class.

Treat proxies as a first-class risk. Review each feature for its relationship with protected attributes, test whether the group can be predicted from the retained features, and document why each retained feature is justified. Proxy testing needs sensitive data, which should be handled in a restricted environment.

Pick the fairness criterion deliberately. Record why a given criterion was chosen, what trade-offs were accepted, and who approved them. Because criteria conflict when base rates differ, this is a governance decision and should not be left to the modelling team alone.

Check legality before mitigation. Mitigation methods differ in whether they use protected attributes at decision time. Obtain legal advice on which are permissible for the decision and jurisdiction.

Test for leakage before release. A membership inference test is inexpensive and should be a routine validation step. Use regularisation, early stopping and capacity limits as default controls, and consider formal differential privacy for models trained on especially sensitive data.

Monitor in production, privately. Distributions drift, and fairness measured at launch may not hold six months later. Use dashboards that track outcomes by group, using noisy aggregate counts where appropriate, and tie alerts to a defined escalation process.

Preserve contestability. Individuals should be able to obtain meaningful information about a decision, have it reviewed by a person and contest it, which requires reason codes, appeal channels and logging.

C. Organisational and governance considerations

Each gate has a named owner and sign-off, and those who build a system should be separate from those who approve its release. Documentation should be readable by auditors and regulators and should record datasets,

intended uses, group-level performance, limitations and monitoring plans [20], [21], [22]. Training is needed because many decisions in the framework, such as choosing a criterion or accepting a trade-off, are not purely technical, and the social context of use shapes a system's effects [24].

D. Threats to validity

Synthetic data and constructed truth. The simulation is simple. The cause of disparity (a penalty in historical decisions and structural shifts in features) is known by construction, and the groups have equal true repayment. Real data have more complex causes, unobserved confounders and base-rate differences. The sizes of the effects reported here should not be applied to real systems.
Selective labels. In real lending, outcomes are observed only for approved applicants. The equal opportunity thresholds assumed outcomes for a validation sample. Methods that work under selective labels are needed in practice.
Two groups and one model family. We considered two groups and mainly logistic regression. Real contexts involve several protected attributes, intersections of groups and more complex models.
Simple attack and simple mechanism. The membership inference attack is basic, and stronger attacks exist. The private audit uses the simplest mechanism and treats group sizes as public.
Legal generality. The references to law and regulation are illustrative and not legal advice. Requirements vary and change.
No user study. We did not test whether the framework is usable in real organisations.

E. Future work

Priorities include case-study validation with organisations, extension to intersectional groups and selective-label settings, evaluation of formally private model training, and metrics for contestability and explanation quality, with legal scholars helping to refine the mapping from controls to obligations.

CONCLUSION

This paper proposed a six-stage framework for ethical compliance in automated decision systems that treats algorithmic bias and data privacy together, with controls, evidence and decision gates at each stage, and tested selected controls in a simulated lending process. A model trained on biased historical decisions amplified the disparity (disparate impact ratio 0.60, against 0.66 historically). Removing the protected attribute was insufficient because a residential proxy predicted group membership with an AUC of 0.84. Proxy removal, reweighing and group-specific thresholds raised the ratio to 0.79, 0.81 and about 0.98, with discrimination ability unchanged. Data minimisation was free for features without signal but did not address bias. Overfitted tree ensembles were vulnerable to membership inference, with attack AUC up to 0.86, whereas regularised models were close to 0.5 and generalised better. A differentially private audit estimated the disparate impact ratio nearly as accurately as a non-private one at audit sizes of 2,000 decisions or more.
The practical lesson is that bias and privacy should be managed as parts of

one lifecycle, with evidence produced at each gate. The tension between measuring bias and protecting sensitive data is real but can be managed with restricted access and private auditing. The legal acceptability of each mitigation must be checked, because the most effective statistical fix may not be permitted. As the data are synthetic, the numbers illustrate the framework and do not describe any real system. The next step is to apply the framework and tests to real decision processes under proper governance.

References

[1] Pawan Kumar Singh, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5713

[2] Pawan Kumar Singh, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5712

[3] Rangavinay Teja Royal Jagga, 2026, DOI: 10.1109/GCWCN66157.2025.11448495, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026 ,https://ieeexplore.ieee.org/document/11448495

[4] Rangavinay Teja Royal Jagga , 2026, Subtle Data Contamination in Visual AI: Covert Model Backdoors via Resampling Exploits, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026 ,https://ieeexplore.ieee.org/document/11448514

[5] Rangavinay Teja Royal Jagga, 2026, Robotic Tactile-Vision Integration for Interactive Object Disentanglement, https://ieeexplore.ieee.org/document/11448337

[6] Rangavinay Teja Royal Jagga, 2026, A Discounted Future Reward Framework for Intelligent Loyalty Optimization using CLTV-Aware Segmentation and Churn Prediction, https://ieeexplore.ieee.org/document/11497560

[7] Anjaneyulu Gudla; Manish Kumar; Vamshi Jeksani; Kiran Kumar Reddy; Gayathri Band, Harnessing Artificial Intelligence Cybersecurity: Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026 https://doi.org/10.1109/aiei69164.2026.11497833

[8] G. Mahender, Department of CSE(Data Science), Vardhaman College of Engineering ; Vamshi Jeksani; Manish Kumar; Koppuravuri Sri Lakshmi Sruthi; Kiran Kumar Reddy, SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable Cybersecurity 2026 https://doi.org/10.1109/aiei69164.2026.11497435

[9] Kiran Kumar Reddy; Kamalakar Gunupati; Manish Kumar; Pell Reddy Rajender Reddy; Ramesh Julakanti; Ram Reddy Jonnalagadda, SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025 https://doi.org/10.1109/computingcon64838.2025.11376613

[10] Amilineni, M.V. 2025. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. DOI: https://doi.org/10.29284/7fjzk883

[11] Pawan Kumar Singh, ACCELERATING LARGE LANGUAGE MODEL TRAINING USING HYBRID QUANTUM-CLASSICAL VARIATIONAL AUTOENCODERS WITH INTEGRATED PHOTONIC QRAM, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/501, DOI: https://doi.org/10.46121/pspc.54.3.45

[12] Pawan Kumar Singh , RESILIENCE ASSESSMENT OF MODERN POWER GRID INFRASTRUCTURE AGAINST EXTREME WEATHER USING PHYSICS-INFORMED NEURAL NETWORKS, Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3415,https://pspac.info/index.php/dlbh/article/view/502, DOI: https://doi.org/10.46121/pspc.54.1.64

[13] Pawan Kumar Singh, MITIGATING THE EXABYTE STORAGE CRISIS: ENHANCING READ/WRITE EFFICIENCY IN 5D OPTICAL GLASS STORAGE USING SILVER NANOPARTICLE DOPING, Vol. 38 No. 11s (2025), International Journal of Applied Mathematics (IJAM) , pISSN 1311-1728, eISSN 1314-8060 ,https://ijamjournal.org/ijam/publication/index.php/ijam/article/view/2153

[14] Mr. Pawan Kumar Singh (Author), Miss. Pari Nidhi Singh (Author) ,Intelligent Bot, ASIN : B0DBGCQY8X,Publisher : Namya Press, Publication date : November 25, 2022,Language : English, Print length : 105 pages,ISBN-10 : 9390124867,ISBN-13 : 978-9390124862,Item Weight : 7.5 ounces, Dimensions : 6.24 x 0.43 x 9.24 inches,https://www.amazon.com/Intelligent-Bot-Pawan-Kumar-Singh/dp/B0DBGCQY8X/ref=sr_1_5?crid=2VAXN5BRN5SMS&dib=eyJ2IjoiMSJ9.v3OCQvHBpy20RIjfMWhoqbOl-RLOlUJIC8Xa0Pjeqilvo2sbZ_suN3LbSCw_OA_ZD_4lgjRLMGlBwZQQ1ehZEfDpt_MJiOUfVuqqzMav0I5_mb6pN5DQ0g0U7SA826n1DQQ3QAtWwopHy4lC5IfhZXttKj9faiyATwhdtDATSSV1-r7FmzM2VQyIA0jxcfq6E6hp9ECfH1fby0li9kUfGpLxMq4GktlzUIJO5ggrhsQ.G1AjQiVqU4HFsnSczvz-X6NXCNs3EFSn-RI8mcWWGOc&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book&qid=1787815741&sprefix=pawan+kumar+singh%2C+bo%2Caps%2C450&sr=8-5

[15] Mr. Pawan Kumar Singh (Author), STORAGE CRISES DUE TO RISE OF AI: A Comprehensive Guide to Managing the Storage Challenges in the Age of Artificial Intelligence, ASIN : B0HFK12ZCT
Publisher : EB1 Profile, Accessibility : Learn more, Publication date : August 17, 2026,Edition : 1st
Language : English, File size : 1.7 MB, Screen Reader : Supported, Enhanced typesetting : Enabled
X-Ray : Not Enabled, Word Wise : Not Enabled ,Print length : 261 pages ,ISBN-13 : 979-8951479112,Page Flip : Enabled, https://www.amazon.com/STORAGE-CRISES-DUE-RISE-Comprehensive-ebook/dp/B0HFK12ZCT/ref=sr_1_1?crid=9MF6C4WDVJR5&dib=eyJ2IjoiMSJ9.DQ8fesBch7ELmRnilDt-YA.NBSCHLmetvnbhkxwIpOoj3s0IVsCV8y6-ZCfCP6Ul1c&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book%2C+eb1+Profile&qid=1787816040&sprefix=pawan+kumar+singh%2C+book%2C+eb1+profile%2Caps%2C420&sr=8-1

[16] Parag Bharatkumar Parikh, FRAMEWORKS FOR MITIGATING SUPPLY CHAIN DISRUPTIONS AND MATERIAL PRICE VOLATILITY IN SUSTAINABLE BUILDING CONSTRUCTION, Vol. 53 No. 2 (2025): April-June 2025 , , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/447, DOI: https://doi.org/10.46121/pspc.53.2.34

[17] Parag Bharatkumar Parikh, IMPLEMENTING DIGITAL INSPECTION PROTOCOLS FOR ENHANCED QUALITY CONTROL MANAGEMENT IN HIGH-RISE CONSTRUCTION, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/439, DOI: https://doi.org/10.46121/pspc.54.3.12
[18] Parag Bharatkumar Parikh , EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/issue/view/30, DOI: https://doi.org/10.46121/pspc.54.1.61
[19] Parag Bharatkumar Parikh, EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Vol. 54 No. 1 (2026): January-March 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/440, DOI: https://doi.org/10.46121/pspc.54.1.61
[20] Ronakkumar shah, REDEFINING BUSINESS PROCESS MANAGEMENT: THE SYNERGY OF AI AND INTELLIGENT AUTOMATION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/448, DOI: https://doi.org/10.46121/pspc.54.3.14
[21] Ronakkumar shah, OPERATIONAL ACCOUNTABILITY IN ACTION: METRICS AND MECHANISMS FOR SYSTEMIC SUCCESS, Vol. 53 No. 4 (2025): October-December 2025, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/452, DOI: https://doi.org/10.46121/pspc.53.4.40
[22] Ronakkumar shah, BEYOND THE BLUEPRINT: AUTOMATED WORKFLOW ENFORCEMENT AND POLICY GOVERNANCE, Vol. 53 No. 2 (2025): April-June 2025 , System Protection and Control, ISSN-1674-3415 ,https://pspac.info/index.php/dlbh/article/view/451, DOI: https://doi.org/10.46121/pspc.53.2.35
[23] Tushita Agarwal; Swaminathan Sethuraman, Mapping Customer Behavior: Visual Analysis of Online Retail Interactions, Date Added to IEEE Xplore: 17 February 2026, DOI: https://doi.org/10.1109/IC2NC67409.2025.11376261
[24] Tushita Agarwal, Efficient Data Summarization and Comparison via Grid-Based Statistical Proxies for Large-Scale Datasets, Date Added to IEEE Xplore: 03 December 2025, DOI: https://doi.org/10.1109/ICBATS66542.2025.11258564
[25] Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021,DOI: https://doi.org/10.30574/ijsra.2021.2.1.0047
[26] Dip Bharatbhai Patel, University of North America, Virginia, USA, Home / Archives / Vol. 5 No. 02 (2025): Volume 05 Issue 02 / Articles, https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/5958
[27] Sai Akshara Vemuri , KINEMATIC OPTIMISATION OF SOFT ROBOTIC MANIPULATORS FOR MEDICAL INTERVENTION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/450
[28] Sai Akshara Vemuri, https://www.periodicos.ulbra.org/index.php/acta/article/view/528
[29] Aziz, A., Chandusha, K., Naga Kiran, B., Prasad, M. L. M., Golla, K., Rajeswari, S., & Katta, S. K. R. (2026). Exploring Machine Learning Applications for Genomic Data Analysis in Personalized Medicine. International Journal of

Drug Delivery Technology, 16(38s), 1166-1173. https://doi.org/10.25258/ijddt.16.38s.127,https://www.periodicos.ulbra.org/index.php/acta/article/view/687

[30] Mohammad Naffizuddin , Chaganti Ashok ,Sowjanya Gunukula , Rvbs Sarma , Bhavana Sujanamulk , Bharani Krishna Takkella , Chukka Ram Sunil ,Varri Sujana , Ashwini Kumar, Etiological Factors of Orofacial Clefts in a Hospital-Based Population: A Descriptive Analytical Study, Published: December 10, 2025,DOI: 10.7759/cureus.98868,DOI : https://doi.org/10.7759/cureus.98868

[31] Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/970

[32] Manoj Kumar Kagitha, Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/969

[33] Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS),Publication year 2023, Volume 2023, Issue 12,https://computerfraudsecurity.com/index.php/journal/article/view/971

[34] Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, https://www.eudoxuspress.com/index.php/pub/article/view/5054

[35] Lakshmi Rahul Reddy Mareddy. (2026). Security and Performance Assessment of Encryption Techniques for Cloud Platforms. International Journal of Communication Networks and Information Security (IJCNIS), 18(1), 68–80. Retrieved from https://ijcnis.org/index.php/ijcnis/article/view/8735

[36] Lakshmi RahulReddy Mareddy, Context-Aware Learning Approaches for Improving Prediction Accuracy in Dynamic Systems, International Journal of Computational Mathematical Ideas (IJCMI), Vol.17 No.3(2025): IJCMI Vol17 NO3 2025 https://www.ijcmi.com/index.php/journal/article/view/15

[37] Lakshmi Rahul Reddy , Mareddy. (2026). Revisiting Autonomy in Artificial Intelligence through Agentic System Design. Journal of Computational Analysis and Applications (JoCAAA), 35(1), 707–713. https://eudoxuspress.com/index.php/pub/article/view/4780

[38] Mahendrakumar Kalal, Lean-Driven SAP Production Planning Optimization Using Machine Learning for Inventory and Throughput Efficiency, Date of Conference: 19-20 March 2026, Date Added to IEEE Xplore: 08 April 2026,ISBN Information:,ISSN Information: DOI:10.1109/ISDFS69419.2026.11459029, https://ieeexplore.ieee.org/abstract/document/11459029

[39] Mahendra kumar Kalal, SUPPLY CHAIN RESILIENCE THROUGH INTEGRATED SAP PRODUCTION PLANNING – PROCUREMENT SYNCHRONIZATION, Archives / Vol. 54 No. 1 (2026): January-March 2026 / Articles, https://pspac.info/index.php/dlbh/article/view/259

[40] Dr Manoj Goel (December 2018). THE ATTITUDE OF DIFFERENT STATES STUDENTS TOWARDS PHYSICAL EDUCATION; STUDYING IN PHYSICAL EDUCATION COLLEGES OF MAHARSHTRA STATE, https://ijeponline.com/index.php/journal/article/view/965/847

[41] Ilakiya Ulaganathan, Tagore Engg. College, Anna University, Chennai, India, ; FNU Antara; Santhosh Chitraju Gopal Varma; Zokir Mamadiyarov; Ziyadulla Turayev, Published in: 2025 International Conference on Sustainability, Innovation & Technology (ICSIT), Date of Conference: 22-23 August 2025,

Date Added to IEEE Xplore: 19 December 2025, ISBN Information: https://ieeexplore.ieee.org/document/11294748/references#references
[42] Shreyas Kasture, Sudhanshu Maurya, Amit Kumar Sharma, Santhosh Chitraju Gopal Varma, Kashish Mirza, Firdous Sadaf Mohammad Ismail, DOI: http://doi.org/10.11591/ijai.v15.i2.pp1746-1761, IAES International Journal of Artificial Intelligence (IJ-AI), ISSN/e-ISSN 2089-4872/2252-8938 https://ijai.iaescore.com/index.php/IJAI/article/view/26932
[43] Shiv Shankar Dwivedi, Mr. Santosh Kumar Sharma, Future-Proofing Data Security with Quantum-Resistant Cryptography, Vol. 11 No. 3 (2025): IJCESEN, DOI: https://doi.org/10.22399/ijcesen.3447, https://ijcesen.com/index.php/ijcesen/article/view/3447
[44] Vijender Kalmotia, Enhancing statistical accuracy in gene perturbation studies, Volume 266, August 2026, 105819, https://doi.org/10.1016/j.biosystems.2026.105819, https://www.sciencedirect.com/science/article/pii/S0303264726001292
[45] Naveen Undrathi, Design and implementation of generative AI-driven data transformation pipelines using spring boot for scalable financial applications, eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 || CrossRef DOI, September 2026 (Volume 31, Issue 3), DOI url: https://doi.org/10.30574/wjarr.2026.30.3.1629,
[46] Naveen Undrathi, Agentic generative AI-enabled microservices framework for intelligent financial data processing in java-based enterprise systems, September 2026 (Volume 31, Issue 3), eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 || https://doi.org/10.30574/wjarr.2026.30.3.1628
[47] Akshay Jain, Comparative Analysis of Linkage Methods in Hierarchical Clustering for Customer Segmentation, Conference paper, First Online: 07 October 2025, https://link.springer.com/chapter/10.1007/978-981-96-7289-9_44
[48] Ismail Hossain, Department of Computer Science, George Mason University, Fairfax, VA, USA, Mridul Banik, Department of Computer Science, Colorado State University, Fort Collins, CO, USA, https://ieeexplore.ieee.org/abstract/document/11446534 ,,DOI: 10.1109/FMLDS67896.2025.00106
[49] Ismail Hossain, Mridul Banik , Fahmid Al Farid , Jia Uddin , Hezerul bin Abdul Karim , Advancing Large Language Models for Low-Resource Languages: A Systematic Review of Pretraining, Adaptation, and Ethical Challenges, 10.32604/cmes.2026.075507, Date: 2026 https://www.sciencedirect.com/org/science/article/pii/S152614922600322X
[50] Ismail Hossain; Mridul Banik, Bridging Visual and Linguistic Intelligence for Chest X-rays: A Comprehensive Review of ViTs and LLM Synergies, Published in: 2025 IEEE International Conference on Future Machine Learning and Data Science (FMLDS) Date of Conference: 02-05 November 2025, Date Added to IEEE Xplore: 25 March 2026, ISBN Information: https://ieeexplore.ieee.org/abstract/document/11446534
[51] Rifat Sarker Aoyon1, Ismail Hossain2, M. Abdullah-Al-Wadud3, Jia Uddin4, A Secured and Continuously Developing Methodology for Breast Cancer Image Segmentation via U-Net Based Architecture and Distributed Data Training, https://www.sciopen.com/article/10.32604/cmes.2025.060917
[52] M. B. Vibha, M. Chandrika, Samitha Khaiyum & P. Rakshitha Kiran, Predicting sedentary behavior in adults using stacked LSTM modeling, published: 27 November 2024
Volume 16, pages 346–355 (2025), https://link.springer.com/article/10.1007/s13198-024-02622-2
[53] Bellary Chiterki Anil, Jayasimha Sondekoppa Rajkumar, T. L. Divya, Samitha khaiyum
Rakshitha Kiran P., Balakrishnan Ramadoss, A Radiomics-based Framework

for Liver Cancer Analysis using Explainable Artificial Intelligence (XAI) Methods, Vol. 15 No. 3 (2025): June, 2025, https://etasr.com/index.php/ETASR/article/view/10377

[54] Jayminkumar Patel, Sanjaykumar patel, Birju Patel, Rajan shah, Development and Validation of a Stability-Indicating Analytical Method for Quantitative Determination of Selected Anticancer Drugs in Pharmaceutical Dosage Form, Archives / Vol. 35 No. S2 (2026): Apr-jun 2026, https://welltestingjournal.com/index.php/WT/article/view/288

[55] Jayminkumar Patel, Sanjaykumar patel, Birju Patel, Rajan shah, Artificial Intelligence-Guided Optimization of Chromatographic Methods for Pharmaceutical Analytical Development, Archives / Vol. 35 No. S2 (2026): Apr-jun 2026, https://welltestingjournal.com/index.php/WT/article/view/294

[56] Sharoon Sharif, Inikpi Egbunu, Yujian Fu, and Zhigang Xiao. 2026. Hybrid Multi-Modal Learning for Extreme Long-Tail Fungi Classification. In 2026 ACM Southeast Conference (ACMSE 2026), April 23–25, 2026, Troy, AL, USA. ACM, New York, NY, USA, 38–45. https://doi.org/10.1145/3746467.3801502

[57] A. Chandrakumaran, "Innovative Strategies for Real-Time Analytics in Modern Data Warehousing Using Dockerized Ecosystems," 2025 International Conference on Advances in Next-Gen Computer Science (ICANCS), Bangalore, India, 2025, pp. 1-6, doi: https://doi.org/10.1109/ICANCS65819.2025.11377462

[58] A. Chandrakumaran, "Enhancing SAP Predictive Maintenance and Industrial Analytics with LLM-Driven Automation," 2025 IEEE Madhya Pradesh Section Conference (MPCON), Jabalpur, India, 2025, pp. 238-244, doi: https://doi.org/10.1109/MPCON66082.2025.11256577.

[59] A. Chandrakumaran, "Enhancing SAP Cloud Optimization with Explainable AI: Cost Control, Workload Forecasting, and Performance Enhancement," 2025 3rd International Conference on Inventive Computing and Informatics (ICICI), Bangalore, India, 2025, pp. 456-463, doi: https://doi.org/10.1109/ICICI65870.2025.11069968.

[60] A. Chandrakumaran, "Innovative Strategies for Real-Time Analytics in Modern Data Warehousing Using Dockerized Ecosystems," 2025 International Conference on Advances in Next-Gen Computer Science (ICANCS), Bangalore, India, 2025, pp. 1-6, doi: https://doi.org/10.1109/ICANCS65819.2025.11377462.

[61] Ashwini Chandrakumara, year :2025, Leveraging Autonomous AI for Agile Growth: Transforming Project Management in Fast-Scaling Startups, Proceedings of the 2nd International Conference on Sustainable Business Practices and Innovative Models (ICSBPIM-2025), pages :730-743, issn :2352-5428, isbn :978-94-6463-872-1, Atlantis Press https://doi.org/10.2991/978-94-6463-872-1_46

[62] Anasuya Patil 1, Gurinderdeep Singh 2, Rajendra Dnyandeo Dighe 3, Dhruv Dev 4, Bhaveshkumar A Patel 5, Samatha Rudrangi 6, Gaurav Tiwari 7, Preparation, optimization, and evaluation of ligand-tethered atovaquone-proguanil-loaded nanoparticles for malaria treatment, 2025 Apr;36(6):711-742. doi: 10.1080/09205063.2024.2422704. Epub 2024 Nov 10., https://pubmed.ncbi.nlm.nih.gov/39522102/

[63] Kancharla Vijayalakshmi, Bethapudi Samuel Anand Andrews, Estimation of Cyproheptadine Hydrochloride and Tricholine Citrate Simultaneously in Syrup Dose by Applying Stability Indicating HPLC Methodology, December 31, 2020, Vol. 11, Issue 4, Pages 1813–1819, DOI: 10.26452/ijrps.v11iSPL4.4383, https://ijrps.com/v11si4/3683

[64] Vijayalakshmi Kancharla ,Kumar Abbaraju V D N, Andrews B S A, Nagaraju Dasari, Nagu Korupolu, Melphan as Anticancer and New Insights of Drug as Parenteral Dosage form against SARS-CoV-2 by Using Development and

Validation , November 18, 2020, Vol. 11, Issue 1, Pages 1437–1446
DOI: 10.26452/ijrps.v11iSPL1.3686, https://ijrps.com/v11si1/2911
[65] Ramesh, D., A. Singh, and L. H. Channaiah. 2026. " Evaluating the Effectiveness of Chewy Chocolate Chip Granola Bar Manufacturing Process on Salmonella Survival." Food Science & Nutrition 14, no. 4: e71698. https://doi.org/10.1002/fsn3.71698.
[66] Inventor: Akhilesh RaiDilip Srinivas Krovi Sarthak Saptarshi Bandopadhyay Ravinder Singh Choudhary Richard LAWTON, 2025-02-28, Priority to US19/067,141,2025-11-14,Assigned to CITIBANK, N.A.,2025-12-30,Application granted,2025-12-30, Publication of US12511187B, https://patents.google.com/patent/US12511187B1/en?oq=US12511187
[67] Vivekanandan Govindan Ekambaram , REAL-TIME AUTOMATED DEFECT DETECTION IN INDUSTRIAL MANUFACTURING USING DEEP LEARNING AND TEXTURE ANALYSIS, Archives ,Vol. 54 No. 4 (2026): October-December 2026, Power System Protection and Control, ISSN-1674-3415,DOI: https://doi.org/10.46121/pspc.54.4.06, https://pspac.info/index.php/dlbh/article/view/545
[68] Vivekanandan Govindan Ekambaram, Yamuna Natarajan, AN OPTIMIZED HYBRID FRAMEWORK FOR REAL-TIME EDGE DETECTION AND FEATURE EXTRACTION IN AUTONOMOUS VEHICLE VISION SYSTEMS, Vol. 54 No. 4 (2026): October-December 2026 , Power System Protection and Control, ISSN-1674-3415, DOI: https://doi.org/10.46121/pspc.54.4.05, https://pspac.info/index.php/dlbh/article/view/544
[69] John Smart Mgbakor, Behavioral Health Equity In High-Disparity Urban Communities: Public Health Impact, National Scalability, And Clinical Practice Implications, Pages1166 – 1177Cover Date, August 2026, Volume34, Issue, Special Issue 7- Fundamental Studies, Print ISSN-0965-7452, Electronic- ISSN-2396-8893, DOI: 10.1922/ejprd.v34i7s.1767, https://www.ejprd.org/article.php?article_id=1767&journal_id=171
[70] Nikhil Kumar Singh, Binay K. Singh, Piyush Kumar, Avinash Pandey, Sudhir Kumar, Sujit Kumar Bishi, A. Pattanayak, V. P. Bhadana, Sujay Rakshit & Kishor U. Tribhuvan, Integrative genomic and machine learning approaches reveal evolutionary signatures in the winged bean mitochondrial genome, Published: 28 November 2025,Volume 26, article number 3 (2026), https://link.springer.com/article/10.1186/s12870-025-07738-6
[71] GAURAV KUMAR SINHA, Year 2026 https://mathsapplication.com/article-view/?id=223
[72] GAURAV KUMAR SINHA, Year 2026 , https://mathsapplication.com/article-view/?id=224
[73] GAURAV KUMAR SINHA, Year 2026 , https://mathsapplication.com/article-view/?id=225
[74] Rasmi Ranjan Choudhury, Design and Implementation of a Tuned Rear Axle Bending Damper for Beat Boom Noise Mitigation in SUV Driveline Systems, 2026 International Conference, https://ieeexplore.ieee.org/document/11746870/authors#authors
[75] Rasmi Ranjan Choudhury, Comparative Performance Analysis of Single-Speed Versus Multi-Speed Transmissions in Battery Electric, 2026 International Conference, https://ieeexplore.ieee.org/document/11746704
[76] Nithesh Gudipuri1, Lalith Chandra Bandaru2, Mohammed Shakeer Bandrevu, SCALABLE DATA ENGINEERING FRAMEWORKS FOR REAL-TIME PROCESSING IN DISTRIBUTED SYSTEMS, Published:2026-10-18,Pages:444-453, https://actainnovations.com/index.php/pub/article/view/1900 ,doi: 10.62441/actainnovations.vi.1900
[77] Mohammed Shakeer Bandrevu; Pratham Patkar; Janardhana Naidu Kola, Trustworthy AI Models for Secure Data Processing Environments, Published in: 2026 International Conference on Natural Language Processing and

Computer Vision (ICNPCV),Date of Conference: 17-18 July 2026, DOI: 10.1109/ICNPCV69434.2026.11708997, https://ieeexplore.ieee.org/document/11708997/authors#authors,

[78] Lalith Chandra Bandaru; Venkateswara Rao Batta; Uday Kumar Reddy Gangula, Hybrid AI and Data Analytics Frameworks for Cyber Resilience, 2026 International Conference on Natural Language Processing and Computer Vision (ICNPCV),Date of Conference: 17-18 July 2026,DOI: 10.1109/ICNPCV69434.2026.11708795, https://ieeexplore.ieee.org/document/11708795

[79] Lalith Chandra Bandaru; Karimulla Syed,Head of Access Management, S&P Global, Raleigh, NC, USA

[80] ; Nishanth Vepachedu; Rajmani Singh Oinam; Raja Vignesh Vaithiyalingam; Quazi Taif Sadat, Advanced Secure Knowledge Extraction Techniques for Large Language Models, 2026 International Conference on Emerging in Computing Technologies (INCECT),Date of Conference: 21-22 August 2026, DOI: 10.1109/INCECT72370.2026.11742788, https://ieeexplore.ieee.org/document/11742788

[81] Mohammed Shakeer Bandrevu, Independent Researcher, Aubrey, TX, USA, ; NagaSatyanarayana Raju Uppalapati; Ramesh Bellamkonda; Srikanth Kamatala, Llm-Augmented Cyber Threat Detection Using Federated Edge Intelligence and Data Mining Techniques, Published in: 2025 International Conference on Computing Technologies & Data Communication (ICCTDC),Date of Conference: 04-05 July 2025, https://ieeexplore.ieee.org/document/11158153 , DOI: 10.1109/ICCTDC64446.2025.11158153

[82] Ankit Pathak, Leveraging AI for Better Data Quality and Insights, Volume 7 Issue 3 (2025), https://doi.org/10.32996/jcsts.2025.7.2.19, https://al-kindipublisher.com/index.php/jcsts/article/view/9345

[83] Uma Rajeswari Batchu, Bharati Reddi, Joshna Rani Surapaneni, Aruna Jangam, Sunil Misra, Anthony Addlagatta & Prakasham Reddy Shetty, Inhibition of Xanthine Oxidase by 1-O-Methyl Chrysophanol, A Hydroxyanthraquinone Isolated from Amycolatopsis Thermoflava ICTA 103, Published: 13 December 2025,Volume 66, pages 1049–1068 (2026) https://link.springer.com/article/10.1007/s12088-025-01534-4

[84] GAURAV KUMAR SINHA, Year 2025, https://visualcompublications.es/SAUC/article/view/6525

[85] GAURAV KUMAR SINHA, Year 2025, https://visualcompublications.es/SAUC/article/view/6526

[86] GAURAV KUMAR SINHA, Year 2025, https://visualcompublications.es/SAUC/article/view/6527

[87] D Benito Johnson 1, Gurinderdeep Singh 2, Deeksha Sharma 3, Venkatesan Natarajan 4, Knv Chenchu Lakshmi 5, Ram C Dhakar 6, Sadhana R Shahi 7, Suresh Velayutham 8, Ruchi Tiwari 9, Exploring Computational Advancements in ADME: Essential Insights for Drug Disposition, PMID: 39496312 DOI: 10.62958/j.cjap.2024.033, https://pubmed.ncbi.nlm.nih.gov/39496312/

[88] Mohammed Shakeer Bandrevu, Service Now Ebonding Methods, Implementation and Key features Published,30 Nov 2024, https://ijcttjournal.org/archives/ijctt-v72i11p125