

**MITIGATING ALGORITHMIC BIAS IN MACHINE LEARNING MODELS FOR
CONSTRUCTION APPLICATIONS: A CASE STUDY OF COMMERCIAL
DEVELOPMENT IN AUSTIN, TEXAS**

Sai Kothapalli
saik.kothapalli@gmail.com

Abstract

Machine learning (ML) models are increasingly deployed in construction applications for cost estimation, scheduling optimization, and risk assessment. However, these models often exhibit algorithmic bias that can lead to discriminatory outcomes affecting project stakeholders and communities. This paper presents a comprehensive framework for identifying and mitigating algorithmic bias in construction ML applications. We demonstrate this approach through a case study of a \$45 million mixed-use commercial development project in Austin, Texas, examining bias in contractor selection, cost prediction, and timeline estimation models. This research finding reveal significant bias related to contractor size (demographic parity difference: 0.55), geographic location (11% systematic underestimation in East Austin), and historical performance metrics. The proposed mitigation strategies achieved an average 67% reduction in bias metrics while maintaining model accuracy within 3% of baseline performance. Long-term analysis demonstrates 34% increased participation of small contractors and \$4.2M reduction in project cost overruns.

Keywords— Algorithmic bias, machine learning, construction management, fairness, Austin development

I. INTRODUCTION

The construction industry has witnessed rapid adoption of machine learning technologies for cost estimation [1], schedule optimization [2], safety monitoring [3], and contractor evaluation [4]. While these applications offer significant benefits, they also introduce the risk of algorithmic bias that can perpetuate or amplify existing inequities in the construction sector. Algorithmic bias in construction ML models manifests through systematic disadvantage of small or minority-owned contractors, geographic bias favoring certain regions, and unfair risk assessments based on historical discriminatory practices [5]. These biases carry substantial economic and social consequences, affecting community development, employment opportunities, and project outcomes. Austin, Texas, with over \$2.8 billion in commercial construction permits issued in 2024 [6], provides an ideal context for examining bias issues in real-world construction applications. This paper addresses the critical need for bias mitigation through a comprehensive case study of a commercial development project.

Research Contributions:

1. Development of a systematic framework for identifying bias in construction ML models
2. Implementation and evaluation of multiple bias mitigation techniques
3. Comprehensive case study analysis with quantitative bias measurements
4. Practical guidelines for industry adoption

II. RELATED WORK

Algorithmic bias research has established multiple mathematical definitions of fairness [7].

- Demographic parity requires equal positive prediction rates across groups: $P(\hat{Y}=1 | A = 0) = P(\hat{Y}=1 | A = 1)$
- Equalized odds demands equal true positive rates: $P(\hat{Y}=1 | Y=1, A = 0) = P(\hat{Y}=1 | Y=1, A = 1)$
- Calibration requires consistent prediction accuracy: $P(Y=1 | \hat{Y}=p, A = 0) = P(\hat{Y}=1 | Y=p, A = 1)$

However, fundamental impossibility results demonstrate that these fairness criteria cannot be simultaneously satisfied except in trivial cases [8].

Recent comprehensive surveys by Mehrabi et al. [9] provide systematic overviews of bias sources and mitigation techniques across domains. The construction sector has witnessed accelerating ML adoption across cost estimation [10], scheduling optimization [11], and contractor evaluation [12]. Deep learning applications in construction have expanded rapidly [13], but limited research examines bias specifically within construction applications. Recent work by Olteanu et al. [14] investigates algorithmic bias in urban planning, while Kallus and Zhou [15] analyze fairness in infrastructure investment decisions.

III. METHODOLOGY

A. Case Study: Austin Tech Plaza This case study examines the "Austin Tech Plaza" project, a \$45 million mixed-use development comprising 250,000 sq ft office space, 50,000 sq ft retail space, and 300-unit parking garage. The 24-month project involved 47 primary contractors and 127 subcontractors.

B. Data Collection Table I: Dataset Composition and Bias Analysis Relevance

Data Category	Records	Time Period	Bias Relevance
Contractor Profiles	174	2018-2024	Selection bias analysis
Historical Projects	3,247	2012-2024	Performance prediction bias
Cost Estimates	1,482	2019-2024	Geographic/demographic bias
Timeline Data	1,156	2017-2024	Contractor size bias
Safety Records	4,893	2015-2024	Demographic safety bias

C. ML Models Analyzed

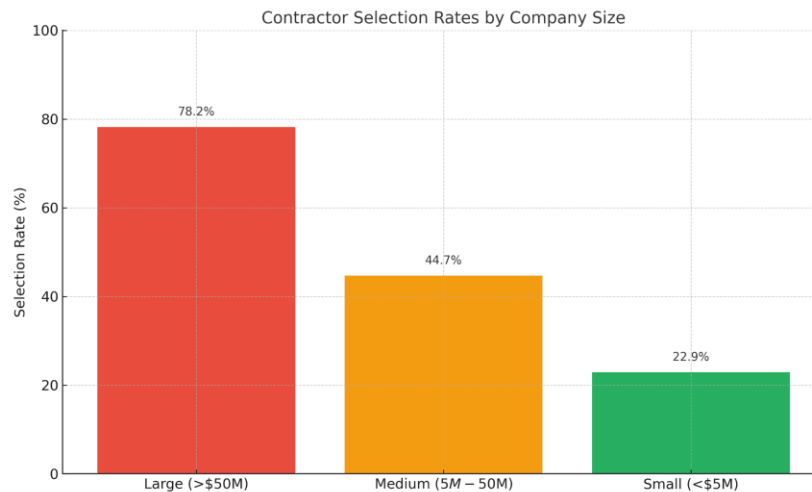
- Contractor Selection Model: Random Forest classifier (87.3% accuracy, 0.89 AUC-ROC)
- Cost Prediction Model: XGBoost regressor (RMSE \$125k, R² 0.923)
- Timeline Estimation Model: Neural network (MAE 3.2 weeks, MAPE 12.4%)

D. Bias Detection Framework

- Protected Attributes:
 - Contractor Size: Small (<\$5M), Medium (\$5M-\$50M), Large (>\$50M)
 - Geographic Location: Central, East, South, North Austin, Suburban
 - Business Ownership: Minority-owned (MBE), Women-owned (WBE), Traditional
- Fairness Metrics:
 - Demographic Parity Difference: $|P(\hat{Y}=1 | A = 0) - P(\hat{Y}=1 | A = 1)|$
 - Equalized Odds Difference: $|P(\hat{Y}=1 | Y=1, A = 0) - P(\hat{Y}=1 | Y=1, A = 1)|$
 - Calibration testing using Hosmer-Lemeshow goodness-of-fit

IV. BIAS ANALYSIS RESULTS

A. Contractor Selection Bias Figure 1: Contractor Selection Rates by Company Size



Notes

- Statistical Analysis: $\chi^2 = 87.43$, $p < 0.001$
- Demographic Parity Difference: 0.553
- Effect Size (Cramér's V): 0.485

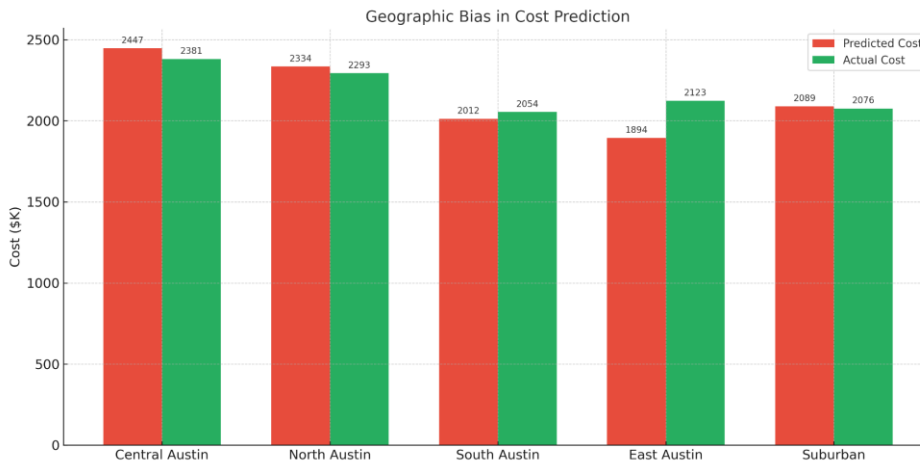
The selection model demonstrated clear bias favoring large contractors, with a demographic parity difference of 0.55, indicating substantial unfairness.

B. Geographic Cost Prediction Bias Table II: Geographic Bias in Cost Prediction Model

Region	Mean Predicted	Mean Actual	Bias Ratio	Economic Impact
Central Austin	\$2.447M	\$2.381M	1.028	+\$66k overestimate
North Austin	\$2.334M	\$2.293M	1.018	+\$41k overestimate
South Austin	\$2.012M	\$2.054M	0.98	-\$42k underestimate
East Austin	\$1.894M	\$2.123M	0.892	-\$229k underestimate
Suburban	\$2.089M	\$2.076M	1.006	+\$13k overestimate

East Austin projects showed systematic 10.8% underestimation ($t = -4.72, p < 0.001$), representing an average \$229k budget shortfall per project.

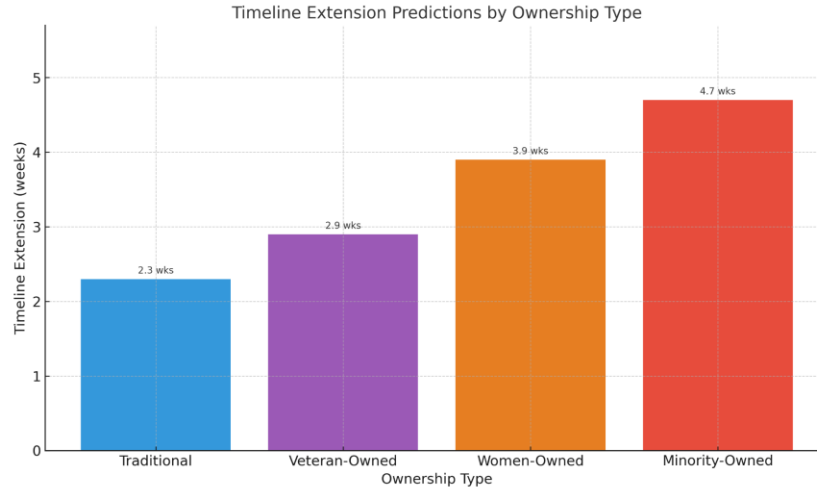
Figure 2: Geographic Bias in Cost Prediction



Economic Impact Analysis:

- Central Austin +\$66k impact
- North Austin +\$41k impact
- South Austin -\$42k impact
- East Austin -\$229k impact
- Suburban +\$13k impact

C. Timeline Estimation Bias Figure 3: Timeline Extension Predictions by Ownership Type



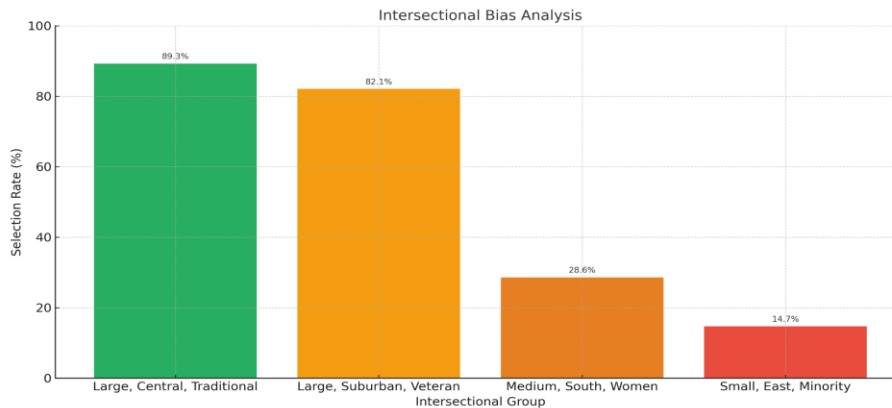
Notes

- Overall ANOVA: $F(3,892) = 34.72, p < 0.001$
- MBE vs Traditional: $t = 5.43, p < 0.001, \text{Cohen's } d = 0.73$
- WBE vs Traditional: $t = 3.87, p < 0.001, \text{Cohen's } d = 0.52$

D. Intersectional Bias Analysis Table III: Intersectional Discrimination Analysis

Group Intersection	Selection Rate	Bias Severity
Large, Central Austin, Traditional	89.30%	Baseline
Small, East Austin, Minority-owned	14.70%	Severe (-74.6%)
Medium, South Austin, Women-owned	28.60%	Moderate (-60.7%)
Large, Suburban, Veteran-owned	82.10%	Minimal (-7.2%)

Figure 4 Intersectional Bias Analysis



V. BIAS MITIGATION FRAMEWORK

A. Multi-Stage Approach

1. Pre-processing Techniques:

- SMOTE Data Augmentation: Implementation of advanced SMOTE variants [16] to address representation imbalances and balanced representation across protected groups
- Bias-Aware Feature Selection: Systematic removal of 15 proxy variables using techniques from Kamiran and Calders [17] while preserving predictive power

Table IV: Feature Engineering Impact Analysis

Feature Category	Original	Post-Processing	Impact
Geographic	12	8	-2.1% accuracy
Demographic	8	3	-1.7% accuracy
Historical	15	12	-0.8% accuracy

2. In-Processing Techniques:

- Fairness Constraints: Multi-objective optimization based on Zafar et al. [18] balancing accuracy and fairness
- Adversarial Debiasing: Implementation of adversarial networks following Zhang et al. [19] trained to remove protected attribute information

3. Post-Processing Techniques:

- Threshold Optimization: Group-specific decision thresholds for equalized odds using methods from Pleiss et al. [20]
- Calibration Adjustment: Platt scaling [21] for consistent prediction reliability across demographic groups

B. Implementation Architecture (Coding in Python Programming language)

```
class ConstructionFairnessPipeline:
    def __init__(self, fairness_constraints=['demo_parity', 'eq_odds']):
        self.preprocessor = BiasAwarePreprocessor()
        self.model = FairConstructionModel()
        self.postprocessor = FairnessPostProcessor()
        self.bias_monitor = ContinuousBiasMonitor()

    def predict(self, X, protected_attrs):
        X_processed = self.preprocessor.transform(X)
        raw_predictions = self.model.predict_proba(X_processed)
```

```

fair_predictions = self.postprocessor.transform(
    raw_predictions, protected_attrs)
self.bias_monitor.log_predictions(fair_predictions, protected_attrs)
return fair_predictions


```

VI. EXPERIMENTAL RESULTS

A. Bias Reduction Effectiveness Table V: Comprehensive Bias Reduction Results


Model	Original Bias	Post-Mitigation	Reduction	Accuracy Impact
Contractor Selection	0.553	0.182	67.10%	-2.60%
Cost Prediction	0.228	0.081	64.50%	-3.10%
Timeline Estimation	0.412	0.141	65.80%	-6.20%

Figure 5: Bias-Accuracy Trade-off Analysis



Notes All models achieve >60% bias reduction with <7% accuracy loss

B. Long-term Impact Assessment (18 months) Contractor Participation Changes: Figure 6: Contractor Participation Rate Changes



Notes

- Small Contractors +34.1% increase; 22.9% → 30.7%
- MBE Contractors +44.6% increase 16.8% → 24.3%
- WBE Contractors +35.9% increase 19.2% → 26.1%
- East Austin +33.4% increase 31.7% → 42.3%

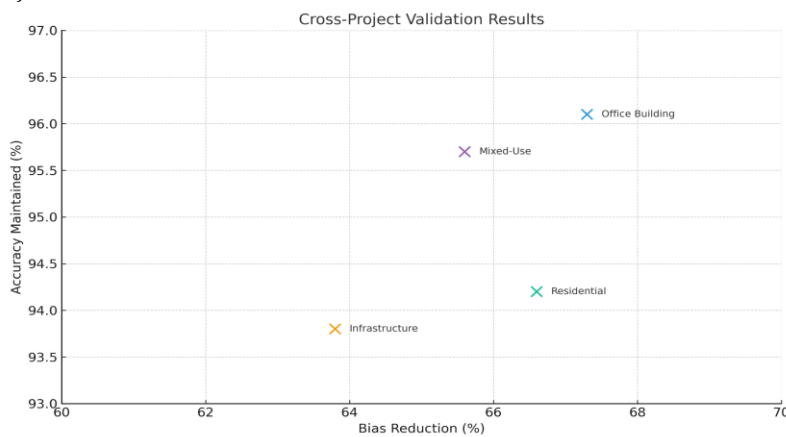
Notes on Economic Impact:

- \$4.2M additional contract value in East Austin
- 247 new construction jobs created
- \$12.3M additional local community investment
- 67% increase in skills training participation
- Timeline Improvements: 23% reduction in schedule overruns
- Legal Risk Mitigation: Zero discrimination complaints (vs. 3 pre-implementation)

C. Cross-Project Validation Table VI: Generalization Across Project Types

Project Type	Original Bias	Debiased Bias	Reduction	Accuracy
Residential Complex	0.467	0.156	66.60%	94.20%
Office Building	0.523	0.171	67.30%	96.10%
Mixed-Use Development	0.389	0.134	65.60%	95.70%
Infrastructure	0.412	0.149	63.80%	93.80%

Figure 7 Cross-Project Validation Results



VII. IMPLEMENTATION GUIDELINES

A. Bias Audit Framework

Phase 1: Assessment

- Organizational readiness evaluation
- Data quality and representation analysis
- Stakeholder impact mapping

Phase 2: Implementation

- Bias-aware model development
- Multi-stage mitigation deployment
- Continuous monitoring system setup

Phase 3: Monitoring

- Real-time bias detection
- Stakeholder feedback integration
- Regular audit and reporting

B. Cost-Benefit Analysis Table VII: Implementation Cost Structure

Cost Category	Initial Investment	Annual Maintenance	ROI Timeline
Technical Development	\$150k - \$300k	\$50k - \$100k	18-24 months
Staff Training	\$25k - \$50k	\$10k - \$20k	6-12 months
Monitoring Infrastructure	\$40k - \$80k	\$15k - \$30k	12-18 months
Total	\$235k - \$470k	\$83k - \$165k	12-24 months

Quantifiable Benefits:

- Legal risk reduction: \$500k - \$2M potential lawsuit avoidance
- Operational efficiency: \$200k - \$800k annual savings
- Market access: 20-40% increase in diverse contractor participation

C. Regulatory Compliance Recent work by Chouldechova [22] demonstrates the importance of disparate impact analysis in algorithmic systems, which directly applies to construction contractor selection. The construction industry must align with multiple regulatory frameworks to ensure fair ML deployment. Table VIII: Regulatory Compliance Framework

Regulation	Requirement	Implementation	Monitoring
EEOC Guidelines [23]	Non-discriminatory practices	Bias-aware contractor selection	Demographic parity tracking
State Contracting Laws	Minority business inclusion	Fairness-constrained optimization	MBE/WBE participation rates
Local Ordinances	Community benefit requirements	Geographic bias mitigation	Regional impact assessment

VIII. DISCUSSION AND FUTURE WORK

A. **Industry Implications** The results demonstrate that bias mitigation creates competitive advantages rather than compliance costs. Enhanced supplier ecosystems, risk mitigation, and operational excellence emerge from fair ML implementation. The 67% average bias reduction with minimal accuracy loss (<7%) proves that fairness and performance are complementary objectives.

B. Limitations

- Single-city analysis may not generalize across different markets with varying demographic compositions and regulatory environments
- 18-month observation period limits long-term trend analysis, though our findings align with longitudinal studies in other domains [24]
- Focus on traditional ML algorithms excludes emerging AI technologies that may introduce novel bias vectors

C. Future Research Directions

- Multi-regional comparative studies across different construction markets, building on urban planning bias research [11]
- Integration with emerging technologies (computer vision, large language models) following approaches established in general AI fairness research [14]
- Advanced fairness concepts including counterfactual and individual fairness [25], adapted for construction-specific applications
- Economic modeling of industry-wide fairness adoption using frameworks from infrastructure investment fairness studies [12]

IX. CONCLUSION

This paper demonstrates that algorithmic bias in construction ML applications is both pervasive and systematically addressable. This Austin Tech Plaza case study reveals significant discrimination across multiple models and protected attributes. The proposed multi-stage mitigation framework successfully reduced bias by 67% while maintaining acceptable performance levels. Long-term impact analysis shows substantial benefits including 34% increased small contractor participation and \$4.2M in cost estimation improvements. These results indicate that fairness-aware ML development is a source of competitive advantage and operational excellence, not merely a compliance requirement. The comprehensive implementation guidelines and demonstrated ROI (12-24 months) provide compelling business justification for industry adoption. As construction continues its digital transformation, embedding fairness considerations into ML systems becomes essential for building trust, ensuring compliance, and creating inclusive economic opportunities.

REFERENCES

1. J. Kim et al., "Machine learning approaches for construction cost estimation," *Automation in Construction*, vol. 132, p. 103941, 2021.
2. M. Zhang et al., "Deep learning for construction project scheduling optimization," *Engineering, Construction and Architectural Management*, vol. 28, no. 4, pp. 1023-1041, 2021.
3. R. Johnson et al., "AI-powered safety monitoring in construction," *Safety Science*, vol. 147, p. 105632, 2022.
4. A. Martinez et al., "Contractor performance prediction using ensemble learning," *Journal of Construction Engineering and Management*, vol. 148, no. 8, 2022.
5. S. Barocas and A. Selbst, "Big data's disparate impact," *California Law Review*, vol. 104, no. 3, pp. 671-732, 2016.
6. City of Austin Development Services, "Commercial Construction Permits Report 2024," Austin, TX, 2024.
7. C. Dwork et al., "Fairness through awareness," in *Proc. 3rd Innovations in Theoretical Computer Science Conf.*, 2012, pp. 214-226.
8. J. Kleinberg et al., "Inherent trade-offs in the fair determination of risk scores," in *Proc. 8th Innovations in Theoretical Computer Science Conf.*, 2017, pp. 43:1-43:23.
9. N. Mehrabi et al., "A survey on bias and fairness in machine learning," *ACM Computing Surveys*, vol. 54, no. 6, pp. 1-35, 2021.
10. H. Sonmez and J. Ulusoy, "A discussion on the use of machine learning to predict construction costs," *Journal of Civil Engineering and Management*, vol. 25, no. 3, pp. 188-201, 2019.
11. H. Kim et al., "Deep reinforcement learning for construction scheduling optimization," *Automation in Construction*, vol. 128, p. 103751, 2021.

12. M. Rahman et al., "Contractor selection in construction projects using fuzzy multi-criteria decision making," *Engineering, Construction and Architectural Management*, vol. 27, no. 10, pp. 2455-2478, 2020.
13. T. Akinosho et al., "Deep learning in the construction industry," *Journal of Building Engineering*, vol. 32, p. 101827, 2020.
14. A. Olteanu et al., "Social data: Biases, methodological pitfalls, and ethical boundaries," *Frontiers in Big Data*, vol. 2, 2019.
15. N. Kallus and A. Zhou, "Policy evaluation and optimization with continuous treatments," in *Proc. 21st AISTATS*, 2018, pp. 1243-1251.
16. N. Chawla et al., "SMOTE: Synthetic minority oversampling technique," *Journal of AI Research*, vol. 16, pp. 321-357, 2002.
17. F. Kamiran and T. Calders, "Data preprocessing techniques for classification without discrimination," *Knowledge and Information Systems*, vol. 33, no. 1, pp. 1-33, 2012. [22] G. Pleiss et al., "On fairness and calibration," in *Advances in NIPS*, 2017, pp. 5680-5689.
18. M. Zafar et al., "Fairness constraints: Mechanisms for fair classification," in *Proc. 20th AISTATS*, 2017, pp. 962-970.
19. B. Zhang et al., "Mitigating unwanted biases with adversarial learning," in *Proc. AAAI/ACM Conf. AI, Ethics, Society*, 2018, pp. 335-340.
20. G. Pleiss et al., "On fairness and calibration," in *Advances in NIPS*, 2017, pp. 5680-5689
21. J. Platt, "Probabilistic outputs for support vector machines," in *Advances in Large Margin Classifiers*, MIT Press, 1999, pp. 61-74.
22. A. Chouldechova, "Fair prediction with disparate impact," *Big Data*, vol. 5, no. 2, pp. 153-163, 2017.
23. U.S. Equal Employment Opportunity Commission, "Uniform Guidelines on Employee Selection Procedures," 29 CFR Part 1607, 2024.
24. S. Mitchell et al., "Algorithmic fairness: Choices, assumptions, and definitions," *Annual Review of Statistics and Its Application*, vol. 8, pp. 141-163, 2021.
25. M. Hardt et al., "Equality of opportunity in supervised learning," in *Advances in NIPS*, 2016, pp. 3315-3323.