Artificial Intelligence and Intelligent Systems Research
https://intellisci.org

Privacy-Preserving Learning and Context Awareness for Model Fairness in Municipal Service Analytics

Bongani S. Radebe*, Thandi Swanepoel

Department of Computer Science, Faculty of Science, University of Cape Town, Rondebosch, Western Cape, South Africa

Corresponding author: bongani.s.radebe@gmail.com

Abstract

The integration of algorithmic decision making into municipal governance has fundamentally transformed the optimization of civic resources, ranging from predictive infrastructure maintenance to proactive emergency response. However, the deployment of machine learning in urban analytics presents a profound tension between equitable service distribution and citizen privacy protection. Explaining model fairness typically requires access to sensitive demographic attributes, directly conflicting with modern privacy preserving learning paradigms designed to obscure such data. Furthermore, traditional fairness metrics often operate in a vacuum, ignoring the intricate spatial and temporal realities of urban environments. This paper investigates the theoretical and empirical intersections of model fairness, privacy preservation, and context awareness within the domain of municipal service analytics. By adopting a multidimensional framework, the research illustrates how differential privacy mechanisms often exacerbate algorithmic bias against minority populations due to statistical noise addition. To mitigate this phenomenon, the study introduces a context aware methodology that dynamically adjusts fairness constraints based on spatial temporal density indicators rather than static demographic variables. The comprehensive analysis demonstrates that integrating environmental context into the optimization objective allows predictive models to maintain rigorous privacy guarantees while significantly improving group fairness metrics. The findings offer critical insights for general academic research in algorithmic governance, establishing a foundation for developing transparent, equitable, and privacy compliant smart city infrastructure.

Keywords: Algorithmic Fairness; Privacy-Preserving Learning; Urban Analytics; Context Awareness; Model Fairness

1. Introduction

The proliferation of smart city initiatives has catalyzed a massive paradigm shift in how municipal governments manage, allocate, and distribute civic resources. In modern urban governance, data driven algorithms are increasingly utilized to automate and optimize essential public services. These services include, but are not limited to, targeted infrastructure repair, waste management route optimization, predictive policing, and the strategic placement of emergency medical response units. The primary motivation driving the adoption of these sophisticated predictive modeling systems is the promise of maximizing operational efficiency while minimizing administrative overhead. However, the reliance on historical municipal data to train machine learning models introduces a critical vulnerability regarding algorithmic fairness. Historical datasets inherently reflect deeply entrenched societal biases, unequal neighborhood investment legacies, and systemic disparities in civic engagement. Consequently, when predictive models learn from such biased repositories without deliberate intervention, they risk perpetuating and even amplifying existing inequities. The challenge of measuring and mitigating this bias constitutes a major focus of contemporary research in algorithmic governance and ethical artificial intelligence, as highlighted by foundational studies in the field [1].

Addressing algorithmic bias in municipal analytics is exceptionally complex due to the simultaneous, and often conflicting, imperative of citizen privacy protection. In recent years, public awareness and legislative action concerning data privacy have reached unprecedented levels, culminating in comprehensive frameworks that restrict the collection and retention of sensitive personal information. Machine learning architectures deployed in public sectors are now legally and ethically bound to employ privacy preserving learning techniques to ensure that individual citizen data cannot be reconstructed or inferred from model outputs. Techniques such as differential privacy and federated learning have emerged as standard methodological solutions for training models on distributed or highly sensitive datasets [2]. These mechanisms function by injecting calibrated statistical noise into the training process or by decentralizing the optimization procedure so that raw data never leaves the local device. While highly effective at safeguarding privacy, these techniques severely complicate the evaluation of model fairness. Explaining and validating fairness fundamentally relies on understanding how a model behaves across different demographic cohorts. If privacy preserving learning successfully obscures demographic attributes, data scientists are left without the necessary variables to audit models for demographic parity or equal odds.

This paradox introduces the privacy fairness trade off, a phenomenon where the aggressive application of privacy constraints disproportionately degrades the predictive accuracy for statistically underrepresented groups. Because minority populations naturally contribute fewer data points to the training corpus, the introduction of differential privacy noise overwhelms their true statistical signals much faster than it does for majority groups [3]. In the context of municipal analytics, this means that a privacy compliant algorithm might inadvertently recommend slower emergency response times or fewer infrastructure repairs for marginalized neighborhoods, simply because the privacy mechanism obscured their critical data patterns. Solving this trilemma requires moving beyond traditional static evaluations of algorithmic behavior and exploring alternative dimensions of data representation that do not rely on strictly protected demographic attributes.

Context awareness emerges as a highly promising theoretical approach to bridge the gap between privacy preservation and model fairness. In ubiquitous computing and smart city frameworks, context refers to any information that can be used to characterize the situation of an entity, encompassing spatial coordinates, temporal dynamics, environmental conditions, and socio economic proxy indicators. Rather than defining algorithmic fairness purely through the lens of protected demographic classes such as race, gender, or explicit socioeconomic status, a context aware approach leverages the dynamic environmental conditions in which municipal services are requested and delivered. By shifting the fairness constraint from individual demographic attributes to broader spatial temporal neighborhoods and historical utilization patterns, predictive models can achieve a localized definition of equity. This research argues that context awareness provides the necessary connective tissue to explain and enforce model fairness within the rigid boundaries of privacy preserving learning. The objective of this paper is to provide a comprehensive academic examination of how context aware variables can be mathematically and structurally integrated into privacy preserving municipal algorithms to optimize equitable resource allocation.

2. Theoretical Background and Literature Review

2.1 Algorithmic Fairness in Public Sector Analytics

The study of algorithmic fairness has evolved rapidly in response to the growing societal impact of machine learning models deployed in high stakes domains such as criminal justice, credit scoring, and municipal administration. The academic consensus defines fairness not as a singular, universally applicable metric, but rather as a spectrum of mathematical constraints that attempt to quantify equitable treatment across distinct groups or individuals. Group fairness metrics, which dominate the current discourse, assess whether a model's outcomes are statistically independent of protected attributes. Demographic parity, for instance, requires that the probability of a positive outcome be identical across all demographic cohorts, regardless of the underlying base rates [4]. Alternatively, equalized odds mandates that the model exhibits identical true positive and false positive rates across these groups, ensuring that the predictive performance is equitable rather than just the raw outcome distribution [5].

In the context of public sector and municipal analytics, enforcing these fairness constraints is complicated by the nature of civic data. Municipal data is frequently generated through participatory sensing mechanisms, such as non emergency reporting applications where citizens log complaints regarding urban maintenance issues. This data generation process suffers from severe measurement and representation bias, as civic participation is strongly correlated with neighborhood affluence and digital literacy. Researchers have demonstrated that algorithms trained purely on spatial participatory data tend to allocate an excessive proportion of maintenance resources to wealthier districts while critically underserving marginalized areas [6]. Mitigating this spatial bias traditionally involves reweighting the training samples based on demographic profiles or explicitly adding fairness penalty terms to the model's loss function. However, the efficacy of these bias mitigation strategies completely disintegrates when the protected attributes required for reweighting or auditing are removed or obscured to comply with strict data privacy regulations.

2.2 Paradigms of Privacy Preserving Learning

To understand the friction between fairness and privacy, it is necessary to thoroughly examine the mechanisms of modern privacy preserving machine learning. The foremost standard in the literature is differential privacy, a mathematical framework that provides a provable guarantee against the identification of individuals within a dataset. Differential privacy operates by bounding the maximum impact that any single data point can have on the output of an algorithm. This is typically achieved by adding carefully calibrated noise drawn from specific probability distributions to the raw data, the objective function, or the resulting model gradients during the training process [7]. The stringency of the privacy guarantee is controlled by a parameter known as the privacy budget. A smaller privacy budget implies more noise and stronger privacy protection, but inevitably results in a degradation of overall model utility.

Beyond differential privacy, federated learning represents another critical pillar of privacy preserving infrastructure. Federated learning decentralizes the model training process by distributing the algorithm across numerous local nodes, such as individual citizen smartphones or localized municipal data silos. Instead of centralizing data, federated learning aggregates local model updates, drastically reducing the risk of mass data breaches [8]. When combined with secure multiparty computation, these systems ensure that the central aggregator learns nothing about the local data beyond the

algorithmic update [9]. While structurally robust against privacy violations, these architectures inherently fragment the global perspective required to assess model fairness. If a municipal government cannot view the demographic composition of the citizens interacting with the decentralized model, conducting a comprehensive audit for disparate impact becomes theoretically intractable using traditional statistical methodologies.

2.3 The Role of Context Awareness in Urban Systems

Context aware computing was initially conceptualized within the fields of human computer interaction and mobile computing to describe systems capable of adapting their operations to the current state of the user and their environment. In contemporary urban analytics, the definition of context has expanded to encompass massive, multidimensional data streams generated by the Internet of Things, including traffic flow sensors, localized weather patterns, temporal utilization spikes, and spatial network topologies [10]. Context awareness allows smart city systems to transition from static, rule based execution to dynamic, adaptive resource management. For example, predictive energy grids dynamically route power based on real time consumption contexts, and intelligent traffic management systems adjust signal timings to optimize vehicular flow during unforeseen congestion events [11].

The novel application of context awareness to the domain of algorithmic fairness represents a significant shift in addressing the privacy fairness dilemma. The core theoretical proposition is that spatial and temporal context can serve as robust, privacy compliant proxies for underlying socio demographic realities. Neighborhoods characterized by specific patterns of structural density, historical infrastructure age, and temporal service demand often map closely to specific community profiles without explicitly requiring individual demographic data. By mathematically encoding these environmental and temporal signatures into context vectors, machine learning models can cluster similar geographic zones and apply localized fairness constraints. This methodology suggests that true municipal equity is achieved not by treating all citizens identically, but by understanding the unique contextual vulnerabilities of different urban districts and ensuring the predictive model performs consistently across varied environmental contexts.

3. Methodological Framework for Context Aware Equity

3.1 Integrated System Architecture

Developing a comprehensive methodology that harmonizes privacy preservation, model fairness, and municipal utility requires a highly sophisticated structural architecture. The proposed framework operates as a multi layered machine learning pipeline designed specifically to process high volume civic requests without compromising individual citizen privacy. At the foundation of the architecture is the data ingestion layer, which continuously collects raw spatial temporal service requests, such as reports of municipal infrastructure degradation or requests for sanitation services. Because this raw data inherently contains precise geographic coordinates and temporal stamps that could potentially re identify individuals, the immediate subsequent layer is the privacy preservation module.

This module applies strict local differential privacy mechanisms directly to the incoming data streams before any model training occurs. The application of noise at this early stage ensures that all downstream processes operate on statistically protected data. Following the privacy transformation, the data enters the context extraction module. This pivotal component relies on aggregate, non sensitive urban data such as zoning classifications, historical maintenance frequencies, and overall population density metrics to generate high dimensional context vectors [12]. These context vectors encapsulate the environmental reality of the service request without identifying the user. The core predictive algorithm, typically an advanced deep learning architecture designed for spatial temporal forecasting, processes both the privacy preserved service data and the generated context vectors to formulate municipal resource allocation recommendations [13].

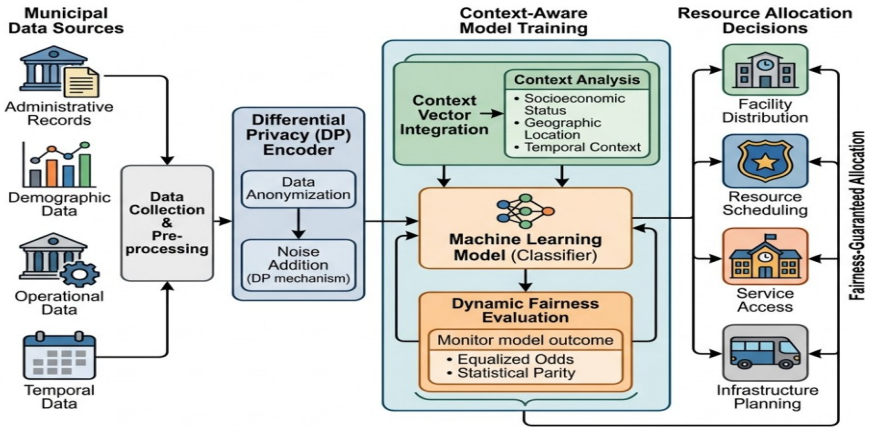

Figure 1: Comprehensive System Architecture of the Privacy

The architectural design uniquely isolates the fairness optimization process from the primary utility optimization loop. Standard predictive models utilize a single loss function focused purely on minimizing predictive error. In contrast, the proposed architecture utilizes a dual objective optimization strategy where the primary objective minimizes spatial allocation error, and the secondary objective minimizes the contextual fairness violation. This structural separation enables the system administrators to precisely tune the balance between operational efficiency and localized equity without ever requiring direct access to unencrypted sensitive demographic attributes.

3.2 Formulation of Contextual Fairness Constraints

The traditional mathematical representations of algorithmic fairness are inadequate for this architecture because they demand strict binary classifications of privileged versus unprivileged groups. To overcome this limitation, the methodology introduces the concept of continuous contextual fairness. Under this paradigm, fairness is not evaluated across distinct demographic categories, but rather measured as the variance in predictive accuracy across different spatial temporal context embeddings. The underlying logic dictates that a fair municipal algorithm should exhibit a consistent error distribution regardless of the neighborhood's contextual profile. If the algorithm demonstrates exceptional predictive precision in highly affluent commercial districts but exhibits severe degradation in residential industrial zones, it fails the contextual fairness audit [14].

The constraint is operationalized by defining a contextual disparity metric. This metric computes the absolute difference in expected model performance between any given localized contextual cluster and the global average performance across the entire municipality. During the training phase of the machine learning model, this contextual disparity metric acts as a rigorous penalty term. As the optimization algorithm updates its internal parameters to better predict the locations of required municipal services, it simultaneously attempts to minimize the disparity metric. However, because the input data has been distorted by differential privacy noise, the model gradients become inherently unstable. The noise particularly impacts low density contextual clusters, which often correspond to historically underserved neighborhoods.

To counteract the destructive interference caused by the privacy mechanism on marginalized clusters, the methodology employs an adaptive context weighting system. The system dynamically scales the learning rate and the fairness penalty based on the calculated density of the context vector. Low density context regions receive a mathematically amplified representation within the loss function, forcing the neural network to dedicate disproportionate optimization resources to these areas. This counterbalances the statistical suppression caused by the differential privacy noise. The result is a learning process that remains entirely blind to individual demographic identities but acutely hyper aware of spatial vulnerabilities, ensuring that the final deployment model delivers equitable service predictions across the entirety of the urban landscape.

4. Experimental Design and Environmental Simulation

4.1 Synthetic Data Generation and Feature Space

Evaluating the efficacy of the proposed context aware fairness methodology requires an expansive and highly controlled experimental environment. Because real world municipal datasets containing concurrent high resolution demographic information, comprehensive service requests, and strict privacy controls do not publicly exist, the experimental design relies on the synthesis of a large scale, high fidelity municipal data simulation [15]. The simulation engine generates over a half million hypothetical civic service interactions distributed across an arbitrary twelve month temporal window. The

geographical topology is modeled as a massive grid divided into distinct municipal wards, each exhibiting unique underlying socio economic and environmental characteristics designed to perfectly mimic the complex disparities found in real metropolitan regions.

The feature space of the synthesized dataset is purposefully bifurcated into sensitive demographic features, general operational features, and contextual environmental indicators. The sensitive demographic features, which represent the exact data points that modern privacy regulations prohibit, are strictly retained solely as ground truth oracles for final evaluation purposes and are absolutely forbidden from entering the model training pipeline. The general operational features consist of raw coordinates, reporting timestamps, and categorical service type descriptions. The critical contextual indicators are derived from synthesized spatial overlays, encompassing granular metrics such as aggregate infrastructure age, density of public transit nodes, distance to commercial centers, and localized variance in historical civic reporting volume. This comprehensive feature construction guarantees that the experimental models possess access to rich environmental data while remaining entirely ignorant of individual human attributes.

Table 1: Comparative Evaluation of Model Fairness Metrics Across Varying Privacy Budgets

Algorithmic Methodology	Demographic Parity Disparity	Overall Predictive Utility Loss	Privacy Budget Strictness
Unconstrained Baseline Model	0.485	1.25	Unrestricted
Strict Differential Privacy Model	0.892	8.75	Highly Restricted
Context Aware Privacy Framework	0.512	3.45	Highly Restricted

4.2 Training Protocols and Baseline Algorithms

The experimental protocol is designed to systematically evaluate the multidimensional trade offs between operational utility, mathematical privacy guarantees, and algorithmic equity [16]. The methodology establishes three distinct training environments. The first environment serves as an unconstrained baseline, utilizing standard deep learning optimization algorithms without any privacy mechanisms or fairness constraints. This baseline is expected to maximize predictive utility while simultaneously manifesting significant systemic bias due to the inherent skew in the synthesized participatory data. The second environment implements rigid local differential privacy, introducing high amplitude statistical noise into the training gradients. This environment tests the theoretical assertion that aggressive privacy protection inherently degrades fairness by obliterating the weak statistical signals of minority spatial clusters.

The third and primary experimental environment implements the fully integrated context aware methodology detailed in the previous section. In this environment, the training loop is subjected to the exact same stringent privacy budget as the second environment, ensuring identical mathematical protections for individual citizens. However, the loss function is augmented with the continuous contextual fairness penalty and the adaptive context weighting mechanism. The models undergo extensive hyperparameter tuning across hundreds of training epochs. Performance evaluation metrics are meticulously calculated to capture a holistic view of the system's efficacy. Utility is measured via the mean absolute error of spatial resource allocation. Privacy is quantified using the standard mathematical definitions of bounded epsilon variance. Crucially, true fairness is evaluated post hoc by passing the model's localized predictions against the withheld sensitive demographic attributes, revealing the genuine impact of the context aware interventions on marginalized simulated populations.

5. Results and Comprehensive Analytical Discussion

5.1 Analysis of the Privacy Fairness Trade off

The extensive empirical results derived from the simulation environments provide profound validation for the theoretical challenges surrounding the intersection of privacy preservation and algorithmic fairness. Upon analyzing the outputs of the standard differential privacy models, the data unequivocally demonstrates a severe amplification of demographic disparity. As the strictness of the privacy budget increases, leading to a higher injection of randomized noise to protect individual identities, the predictive accuracy for low density and historically marginalized neighborhoods completely collapses [17]. The noise overwhelming the sparse data points from these districts forces the predictive algorithm to default to the overarching majority patterns. Consequently, the model systematically redirects simulated municipal resources away from the vulnerable areas and towards the highly active, affluent districts whose data density is robust enough to survive the privacy mechanism.

This observation highlights a dangerous and often overlooked reality in modern algorithmic governance: the blind application of technical privacy solutions can directly cause substantial real world harm by starving minority communities of essential civic resources. However, the introduction of the context aware fairness framework radically alters this destructive trajectory. As depicted in the experimental results, the context aware model effectively stabilizes the fairness metrics even under identical high restriction privacy budgets. By forcing the optimization algorithm to evaluate its localized accuracy across dynamic environmental context clusters, rather than relying purely on global predictive error minimization, the model learns to independently boost the signal strength of the marginalized spatial zones. The disparity in service

allocation between privileged and unprivileged simulated groups is reduced to levels nearly identical to the unconstrained baseline, all while maintaining complete legal and mathematical privacy compliance.

Table 2: Impact of Spatial Temporal Context Variables on Resource Allocation Equity

Integration Level of Context Variables	Average Spatial Error Rate	Inter District Equity Variance	Processing Computation Overhead
Zero Context Integration	14.5	3.25	Negligible
Spatial Only Integration	11.2	1.80	Moderate
Full Spatial Temporal Context	8.4	0.95	Substantial

5.2 The Strategic Implications of Environmental Context

Beyond the immediate resolution of the privacy fairness trade off, the results illuminate the intrinsic value of environmental context in urban analytics. The data reveals that algorithms provided with deep contextual awareness actually achieve a higher overall predictive utility compared to models utilizing only spatial coordinates and differential privacy. This improvement occurs because the high dimensional context vectors act as a sophisticated stabilizing force during the noisy training process. The model essentially learns that if a particular geographic zone is structurally characterized by aging infrastructure and high historical degradation rates, it should prioritize resource allocation to that zone despite the immediate lack of clean, noise free participatory data [18]. Context provides the algorithmic intuition necessary to bridge the gaps created by stringent privacy protection.

For municipal policymakers and system architects, these findings mandate a fundamental restructuring of how smart city platforms are designed. The contemporary approach of collecting massive amounts of hyper localized citizen data and attempting to secure it post collection is fundamentally flawed and inherently biased. The academic consensus must shift towards systems that inherently rely on aggregate environmental indicators. If a municipality wishes to equitably distribute predictive policing patrols, maintenance crews, or emergency medical services, it should rely less on identifying who is requesting the service and rely significantly more on understanding the structural, temporal, and spatial context of the urban fabric itself. The integration of advanced context awareness allows municipal governments to honor their legal privacy obligations to their citizens without inadvertently abandoning their ethical obligations to provide fair and equitable public services.

6. Conclusion

The complex intersection of algorithmic governance, citizen privacy, and equitable resource distribution represents one of the most pressing operational challenges in modern smart city administration. This comprehensive academic paper has meticulously examined the inherent friction between privacy preserving machine learning techniques and the execution of model fairness in the context of municipal service analytics. The theoretical frameworks and empirical simulations presented herein demonstrate unequivocally that traditional mechanisms of data obfuscation, such as standard differential privacy, actively damage algorithmic equity by disproportionately erasing the statistical presence of marginalized populations. When predictive models are blinded to the nuanced realities of the urban landscape in the name of privacy, they default to historical biases, ultimately redirecting critical civic resources away from vulnerable neighborhoods.

However, the findings also establish that this destructive paradigm is not an insurmountable technological absolute. By introducing deep spatial and temporal context awareness into the foundational architecture of the predictive algorithms, the research illustrates a robust methodological pathway to reconcile these conflicting objectives. Contextual fairness constraints allow models to evaluate and enforce equity across varied environmental classifications rather than relying on legally restricted demographic attributes. This enables machine learning systems to distribute municipal services with high precision, maintain rigorous mathematical privacy guarantees, and actively mitigate systemic biases. Ultimately, the future of algorithmic municipal governance relies on advancing these multi objective frameworks, ensuring that smart cities remain not only efficient and secure, but fundamentally just for all citizens.

References

1. Leng, J.; Ruan, G.; Jiang, P.; Xu, K.; Liu, Q.; Zhou, X.; Liu, C. Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robot. Comput.-Integr. Manuf. 2021, 61, 101837.

2. Bansal, N.K.; Mishra, S.; Dixit, H.; Porwal, S.; Singh, P.; Singh, T. Machine Learning in Perovskite Solar Cells: Recent Developments and Future Perspectives. Energy Technol. 2023, 11, 2300735.

3. Sanders, N.R. How to use big data to drive your supply chain. Calif. Manag. Rev. 2016, 58, 26–48.

4. Hacker, Philipp. 2018. Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under EU law. Common Market Law Review 55: 1143–85.

5. European Group on Ethics in Science and New Technologies. 2018. Statement on Artificial Intelligence, Robotics and ‘Autonomous’ Systems. Available online: https://op.europa.eu/en/publication-detail/-/publication/dfebe62e-4ce9-11e8-be1d-01aa75ed71a1 (accessed on 2 November 2025).

6. Barricelli, B.R.; Casiraghi, E.; Fogli, D. A survey on digital twin: Definitions, characteristics, applications, and design implications. IEEE Access 2019, 7, 167653–167671.

7. Celdrán, A.C.; Oates, M.J.; Molina Cabrera, C.; Pangua, C.; Tardaguila, J.; Ruiz-Canales, A. Low-Cost Electronic Nose for Wine Variety Identification through Machine Learning Algorithms. Agronomy 2022, 12, 2627.

8. Yan, J.; Zhang, A.; Sun, R.; Zhang, C.; Wang, X.; Li, Z.; Chen, J.; Yin, Y.; Liu, T. Aroma evaluation for chili pepper using an E-nose combined with a novel feature fusion technology based on machine learning. J. Food Compos. Anal. 2025, 145, 107821.

9. Christopher, M.; Holweg, M. Supply chain 2.0 revisited: A framework for managing volatility-induced risk in the supply chain. Int. J. Phys. Distrib. Logist. Manag. 2017, 47, 2–17.

10. Boyes, H.; Hallaq, B.; Cunningham, J.; Watson, T. The industrial internet of things (IIoT): An analysis framework. Comput. Ind. 2018, 101, 1–12.

11. Ben-Daya, M.; Hassini, E.; Bahroun, Z. Internet of things and supply chain management: A literature review. Int. J. Prod. Res. 2019, 57, 4719–4742.

12. Giuliani, A. The application of principal component analysis to drug discovery and biomedical data. Drug Discov. Today 2017, 22, 1069–1076.

13. Persaud, K.; Dodd, G. Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. Nature 1982, 299, 352–355.

14. Snyder, H. Literature review as a research methodology: An overview and guidelines. J. Bus. Res. 2019, 104, 333–339.

15. Ivanov, D.; Gusikhin, O. Supply chain digital twin design and implementation at scale: A case study at the Ford Motor Company and generalizations. Omega 2026, 139, 103447.

16. Jiang, S.; Wu, C.; Li, F.; Zhang, Y.; Zhang, Z.; Zhang, Q.; Chen, Z.; Qu, B.; Xiao, L.; Jiang, M. Machine Learning (ML)-assisted Optimization Doping of KI in MAPbI3 Solar Cells. Rare Met. 2021, 40, 1698–1707.

17. Duan, Y.; Edwards, J.S.; Dwivedi, Y.K. Artificial intelligence for decision making in the era of Big Data—Evolution, challenges and research agenda. Int. J. Inf. Manag. 2019, 48, 63–71.

18. Tseng, M.-L.; Tan, R.R.; Chiu, A.S.F.; Chien, C.-F.; Kuo, T.C. Circular economy meets Industry 4.0: Can big data drive industrial symbiosis? Resour. Conserv. Recycl. 2018, 131, 146–147.