Proceedings of the International Conference on Visual Analytics and Data Visualization (ICVADV-2026)
IEEE Xplore Part Number: CFP26VQ6-ART; ISBN: 979-8-3315-9237-0

2026 International Conference on Visual Analytics and Data Visualization (ICVADV) | 979-8-3315-9237-0/26/$31.00 ©2026 IEEE | DOI: 10.1109/ICVADV67766.2026.11470538

Explainable and Fairness-Aware AI for Housing Subsidy Allocation in Smart-City Governance

[1]B.CHANDRAKALA
Assistant Professor,Department of Computer Science and Engineering, G.Pullaiah College of Engineering and Technology (Autonomous) , Kurnool, Andhra Pradesh, India, E-Mail:besthachandrakala227 @gmail.com

[2]T.Murali Karthik
4th Year, B.Tech Computer Science and Engineering
G Pullaiah College of Engineering and Technology,
Kurnool, Andhra Pradesh, India
*E-Mail:*muralikarthikt@gmail.com

[3]U.NITISH KUMAR
4th Year, B.Tech Computer Science and Engineering
G Pullaiah College of Engineering and Technology,
Kurnool, Andhra Pradesh, India
E-Mail:unitishkumar22@gmail.com

[4]T. LIKITH RAO
4th Year, B.Tech Computer Science and Engineering
G Pullaiah College of Engineering and Technology,
Kurnool, Andhra Pradesh, India
E-Mail: shivaadas15@gmail.com

[5]M. VIJAY KUMAR
4th Year, B. Tech Computer Science and Engineering
G Pullaiah College of Engineering and Technology,
Kurnool, Andhra Pradesh, India
E-Mail: vijayvk90634@gmail.com

[6]V. VIJAYA CHANDRA RAO
Assistant Professor,Department of Computer Science and Engineering, G.Pullaiah College of Engineering and Technology (Autonomous) , Kurnool, Andhra Pradesh, India,
vijaychandrarao.v@gmail.com

Abstract: Data-driven systems to distribute scarce social-welfare resources like subsidised housing are becoming an important part of smart-city governance. Nevertheless, a large body of extant artificial intelligence (AI) methods are focused on predictive performance and inadequate to meet issues of fairness, explainability, and accountability that play vital roles in decision-making in the public sector. The loophole is dangerous, as it tends to strengthen socio-economic differences and lack of trust in automated welfare. The study suggests an explainable and non-biassed AI model to implement the use of socio-economic data provided by the American Community Survey (ACS) to provide equitable housing subsidies. The suggested methodology presents subsidy allocation as a constrained optimum allocation problem that incorporates interpretable vulnerability modelling, direct constraints of fairness across groups of demographics targeted for protection and explainability in decisions. The governance layer proposed is a human-in-the-loop to facilitate policy oversight, auditing and accountability. Through experimental analysis, it is shown that the developed framework always produces the final allocation results more balanced and policy-consistent in comparison with the traditional rule-based and fairness-free baselines. The findings show better balance in the demographics with efficient coverage using the limited budget. Further, the generated explanations are indicated to agree with the policy-relevant vulnerability factors, thereby heightening transparency, and contributing to guided governance decisions. Conclusively, the suggested solution the next generation of smart-city welfare systems responsible AI makes visible the allocation efficiency, equity, and interpretability of decisions leading to the availability of transparent, accountable, and governance-friendly social-welfare-supporting decision support.

Keywords: Explainable artificial intelligence, Fairness-aware learning, Housing subsidy allocation, Smart-city governance, Social welfare systems, Human-in-the-loop decision support.

I. INTRODUCTION

The phenomenon of rapid urbanisation and population growth has challenged some of the city administrations in ways like never before, because they have to allocate scarce resources allocated as social-welfares in an effective and fair way. The housing subsidies are some of these resources that are very important in providing solutions to the problem of affordability, curbing the risks associated with homelessness, and eradicating the socio-economic inequalities in urban localities. As the availability of large-scale urban and socio-economic data (especially public policy and smart-city governance) has increased, artificial intelligence (AI) and machine-learning methods have been actively utilised to assist with decision-making based on the data [1], [2].

But, with AI being used in social-welfare distribution, some serious ethical and social issues are presented. Models based on data, trained on historical socio-economic data, will tend to strengthen existing inequalities unless it is explicitly prepared to consider considerations of fairness [3]. In that of housing, previous reports have established that predictive models can have a demographic bias in terms of income group, racial group, and geographic area, thus, questioning transparency, accountability, and societal confidence [4]. A growing awareness of the fact that AI systems introduced into the framework of the public governance, therefore, must demonstrate not only high accuracy, but also be fair, interpretable, and policy-oriented is emerging [5].

Explainable Artificial Intelligence (XAI) has become one of the promising paradigms that should be used to increase the transparency and accountability conditions in AI-controlled systems of decisions. The XAI methods allow human stakeholders to interpret, examine, and read the decisions of the model which matters especially in high stakes applications in the public-sector, i.e. housing assistance and distribution of welfare [6]. Concurrently, the concept of fairness-conscious machine learning has also presented quantitative measurement and mitigation actions to measure and lessen algorithmic bias in all classes of shown populations [7]. Even in the face of such progress, prevailing literature has almost ignored the fact explainability or fairness and incorporated these as post-hoc evaluation instruments, as opposed to parts of resource allocation mechanisms.

Moreover, the majority of ACS-based housing research centres on the predictive activities, i.e., house price prediction, vulnerability analysis, or retrospective equity research. These are good strategies to analyse but will fail to answer the application aspect experienced by city administrators; how to be fair in allocating the limited housing subsidies within the limited resources yet transparency and accountability. Within the context of smart-city governance, AI systems are likely to serve as a decision-support tool that does not supplant human judgement, is used to justify policies, and allow it to be audited and not human decision-makers [8].

Motivated by these limitations, this paper identifies a need to create an explainable and fairness-aware artificial intelligence (AI) framework for a fair allocation of housing subsidies based on data recorded by the American Community Survey (ACS). The proposed approach explicitly defines housing subsidies distribution as constrained allocation problem and combines fairness goals in the decision process and provides for interpretable explanations sourced for smart-city governance needs. Though using both equitable allocation modeling and decision-level explanation, this work hope to facilitate the process of responsible and transparent AI deployment in urban welfare systems.

1.1 Key Contributions

1. **Allocation-centric formulation:** We formulate housing subsidy distribution as a fairness-aware allocation problem under limited public resources, moving beyond conventional predictive housing models.
2. **Integration of fairness constraints:** Fairness objectives are explicitly incorporated into the allocation mechanism to mitigate demographic and socio-economic disparities in subsidy distribution.
3. **Governance-oriented explainability:** An explainable AI framework is developed to generate decision-level explanations that support transparency, auditability, and policy justification.
4. **Smart-city decision support:** The proposed framework is designed as a human-in-the-loop decision-support system aligned with smart-city governance and public accountability requirements.

The remainder of this paper is organized as follows. Section II reviews related work on AI-based housing analytics, fairness-aware learning, and explainable AI in public governance. Section III presents the proposed methodology for equitable housing subsidy allocation, including data modeling, fairness integration, and explainability mechanisms. Section IV describes the dataset, experimental setup, and evaluation metrics. Section V discusses the experimental results and comparative analysis. Finally, Section VI concludes the paper and outlines directions for future research.

II. Literature Review

2.1 AI-Based Decision Support for Housing and Urban Welfare

Recent years have witnessed increasing adoption of data-driven and machine-learning approaches to support housing-related decision-making in urban environments. Previous research has been conducted on the forecasting of housing prices, housing affordability, and housing susceptibility using large size socio-economic data like the American Community Survey (ACS) [9], [10]. Machine-learning models, which have been trained on ACS-based features of cities, have been demonstrated to be effective in capturing spatial and demographic differences in housing markets even as well as exposing inequality of fairness across the groups of the population under protection in predictive models, demonstrating the risk of further enhancement of socio-economic disparities in urban housing systems [11].

In line with this body of work, ACS tract-level measures have since been used to examine past housing appraisal methods, exposing systematic differences in valuing neighbourhoods and the importance of data-driven analysis in identifying structural inequity in housing-related choices in governance [12]. Although having these approaches offer crucial information about housing inequity, they are mainly predictive and retrospective in outlook and do not offer information on how to allocate housing subsidies, or welfare resources, prospectively [13].

2.2 Fairness-Aware Machine Learning Using ACS Data

Justice in algorithmic decisions has become a key issue in social-economic and governmental uses. ACS-based benchmark datasets have been presented to make it easy to systematically analyze fairness measures by populations in a way that can be followed up by further research using ACS data to study bias, disparate impact, and equity trade-offs in socio-economic predicting tasks [14], [15].

In the housing domain, fairness-aware learning has been explored mainly in relation to pricing, appraisal, and access disparities. These studies typically evaluate demographic parity, equal opportunity, or group-wise error disparities to quantify bias. However, fairness is often treated as an evaluation criterion rather than an operational constraint within a resource allocation mechanism. As a result, existing works stop short of integrating fairness-aware optimization directly into housing subsidy allocation policies [16].

2.3 Explainable AI for Public Policy and Urban Governance

Explainable AI (XAI) has gained prominence as a means to enhance transparency, accountability, and trust in AI-driven public-sector systems. Interpretable machine-learning techniques have been applied to urban risk assessment, infrastructure planning, and social vulnerability analysis using ACS-derived indicators, demonstrating that such models can reveal key socio-economic drivers of housing vulnerability and improve decision-makers' understanding of policy-relevant risk factors [17].

Despite these advances, most XAI studies in urban contexts focus on explanation of predictive outcomes rather than explanation of allocation decisions. In particular, explanations are rarely tailored to support governance processes such as justification of subsidy decisions, auditability, or citizen-level accountability. The applicability of XAI in ensuring equitable governance of welfare is therefore still underutilised [18].

2.4 Research Gaps and Motivation

Based on the review above, some of the critical research gaps may be outlined:

1. **Lack of allocation-centric modeling**: Available ACS-based housing research is mostly focused on prediction (e.g., prices, vulnerability) instead of rational distribution of housing subsidies in cases of resource scarcity.

2. **Limited integration of fairness into allocation mechanisms**: Post-hoc evaluation is often referred to as fairness metrics that are not often incorporated as constraints or objectives in housing subsidy allocation models.
3. **Insufficient explainability for governance use**: Although interpretable models are used, few studies produce decision level explanations that can justify policies, audit or appeal in the governance of smart-cities.
4. **Weak linkage between AI models and governance workflows**: The existing solutions lack proper positioning of AI systems as human-in-the-loop decision-support systems that comply with the responsibility requirements of the public sector.

In order to fill these gaps the current study suggests an explicable fairness conscious AI framework of equitable housing subsidy distribution based on ACS data. Also in contrast with previous literature, the approach offered directly represents housing subsidy allocation as a constrained allocation problem, incorporates the goal of fairness in decision-making, and offers explanations to be understood by smart-city management.

III. Proposed Methodology

3.1 Problem Definition and System Overview

The solution presented in the proposed framework solves the issue of fairly distributing housing subsidies within smart-city governance based on explainable and fairness-conscious artificial intelligence. With a population of households being defined by the socio-economic attributes based on American Community Survey (ACS), the goal is to rank households in order to allocate subsidies using limited available public funds, yet making sure that fairness, transparency, and compliance with governance.

Let

- $\mathcal{H} = \{h_1, h_2, \ldots, h\ \}$ denote the set of applicant households,
- $\mathbf{x}_i \in \mathbb{R}^d$ denote the ACS-derived feature vector of household h_i,
- B denote the total available housing subsidy budget.

The system outputs an allocation decision vector

$$\mathbf{a} = [a_1, a_2, \ldots, a_N], a_i \in \{0,1\}, \tag{1}$$

where $a_i = 1$ indicates that household h_i is selected for subsidy support.

The framework is a human-in-the-loop type of decision-support system, that is, recommendations provided by AI are explained using an interpretable description of the reasons why this recommendation has to be taken.

3.2 Data Modeling and Feature Representation

The ACS variables are changed to indicators that are policy-sensitive in the scope of vulnerability of housing. The set of features of each household h_i can be defined as:

$$= [\ ^{\text{inc}},\ ^{\text{rent}},\ ^{\text{burden}},\ ^{\text{hh}},\ ^{\text{emp}},\ ^{\text{region}}] \tag{2}$$

Where, these are income level, housing cost, rent burden ratio, household size, employment status, and regional context, respectively.

In order to facilitate the fairness analysis, a vector of protected attributes applied to each household is as well:

$$= [\ ^{\text{age}},\ ^{\text{gender}},\ ^{\text{income\_group}}] \tag{3}$$

that is not a decision feature, but an evaluation and constraint enforcement component, fairness.

3.3 Housing Vulnerability and Priority Scoring

A housing vulnerability score of the urgency of subsidy support is calculated. This score is trained using an interpretable model $f(\cdot)$, I.e. a decision tree or monotonic gradient-boosted model:

$$= (\ ) \tag{4}$$

Higher levels of represent more housing vulnerability.

To be able to interpret the policy, the model is restricted to maintain the monotonic relationships (e.g., the higher the rent burden, the lower the vulnerability) only.

3.4 Fairness-Aware Subsidy Allocation Formulation

The allocation optimization task is formulated as a constrained task.

3.4.1 Objective Function

The first goal is to cover the overall vulnerability as much as possible:

$$\max \sum_{=1}^{N} a_i \cdot s_i. \tag{5}$$

3.4.2 Budget Constraint

The allocation should meet the subsidy budget:

$$\sum_{1}^{=} a_i \cdot c_i \le B \tag{6}$$

where denotes the subsidy cost associated with household h_i.

3.4.3 Fairness Constraints

In order to guarantee comparative inequality of allocation across groups of protection $g \in \mathcal{G}$,, there must be imposed on this group demographic parity constraints:

$$\left| \frac{1}{|\mathcal{H}_g|} \sum_{i\in\mathcal{H}_g} a_i - \frac{1}{N} \sum_{i=1} a_i \right| \le \epsilon, \forall g \in \mathcal{G} \tag{7}$$

where controls the allowable disparity across groups.

This leads to optimization that balances efficiency and equity, as is the case in the real-life policy trade-offs.

The outline of fairness-sensitive and explicable procedure of housing subsidy allocation is presented in Algorithm 1.

Algorithm 1: Fairness-Aware Explainable Housing Subsidy Allocation

Input: Household features $\mathcal{H} = \{\mathbf{x}_i\}^{N}_{=1}$, protected groups , budget , fairness threshold

Output: Allocation vector **a**, explanations $\mathcal{E}$

Process:

1. Compute vulnerability scores $= (\ )$ using (4).

2. Rank households in descending order of .
3. Initialize = , remaining budget = .
4. **for** each household h_i in ranked order **do**
   - **if** $B_r \geq c_i$ **then**
     - Tentatively set = 1.
     - Evaluate fairness constraint (7) for all $g \in$.
     - **if** constraints satisfied **then**
       - Confirm allocation and update $B \leftarrow$ $- c_i$.
     - **else** set = 0.
   - **end if**
5. **end for**
6. **for** each allocated household = 1 **do**
   - Compute feature attributions ϕ_i using (8).
   - Generate explanation e_i using (9) and add to $\mathcal{E}$.
7. **end for**
8. Present $(\mathbf{a}, \mathcal{E})$ for human-in-the-loop review.

3.5 Proposed System Architecture

Figure 1 shows the detailed end-to-end architecture of explainable and fairness-aware housing subsidy allocation that could combine the modelling, allocation, and governance units defined in the significant sections 3.1-3.4 and execute the algorithm within the framework of a decision-support system. The architecture consists of four layers: Data Layer which preprocesses ACS / PUMs data and isolates protected attributes for fairness enforcement, Modeling Layer that computes interpretable scores for housing vulnerability, also known as Allocation and Explainability Layer which determines subsidized homes under budget and equity constraints and generates decision-level explanations and lastly Governance Layer that is where human in the loop (review and policy tuning), audit logging happening. The proposed architecture allows making Smart-city uses use welfare decisions that are transparent, auditable and policy aligned by ensuring that fairness constraints and explainability are directly implemented in the allocation process.

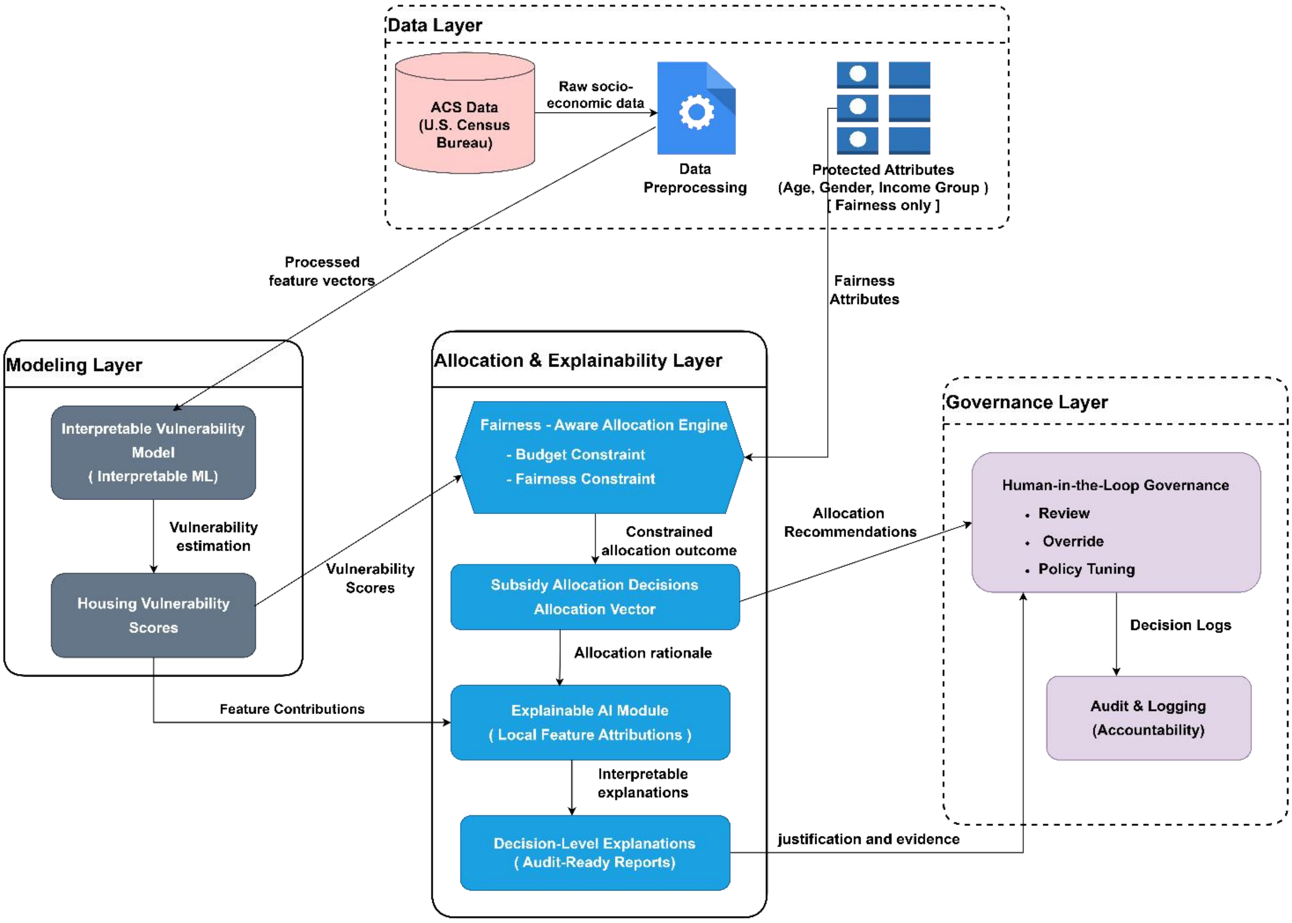

Fig.1. Proposed Explainable and Fairness-Aware Housing Subsidy Allocation Architecture for Smart-City Governance

3.6 Explainable AI Module

All allocation selections are followed with a local explanation whereby the feature attribution methods are utilised to make the allocation decisions transparent.

The explanation vector of a chosen household h_i is given as:

$$\phi_i = [\phi_{i1}, \phi_{i2}, \dots, \phi_{id}] \quad (8)$$

where quantifies the contribution of feature to the vulnerability score .

A human-readable explanation is generated as:

$$\text{Explanation}(h_i) = \sum^{d}_{=1} \phi_{ij} \cdot x_{ij}. \qquad (9)$$

That is the reason why such explanations allow policymakers to defend allocation decisions, audit and appeals.

3.7 Human-in-the-Loop Governance Mechanism

The governance layer allows policymakers to:

- Inspect allocation outcomes,
- Review explanations,
- Adjust fairness thresholds ϵ,
- Modify budget constraints B.

Let $\Theta = \{B, \epsilon\}$ denote policy parameters. The allocation function becomes:

$$\mathbf{a} = \mathcal{A}(\mathbf{s}, \Theta) \qquad (10)$$

ensuring that final decisions reflect both algorithmic recommendations and human judgment.

3.8 Computational Complexity Analysis

Let denote the number of households and the feature dimension.

- Vulnerability score computation: $\mathcal{O}(N \cdot d)$
- Fairness-constrained allocation: $\mathcal{O}(N \log N)$
- Explanation generation: $\mathcal{O}(N \cdot d)$

The overall complexity is:

$$\mathcal{O}(N \cdot d + N \log N) \qquad (11)$$

which is suitable for city-scale smart governance systems.

The specified approach will combine interpretable vulnerability modelling, fairness-constrained optimization, and governance-oriented clarification to facilitate equal distribution of housing subsidies. By explicitly incorporating the notion of fairness and transparency into the distribution process, the framework will bring the requirements of smart-city governance and AI-driven decision-making in line.

IV. Experimental Setup

4.1 Dataset Description and Preprocessing

The given study utilises household data based on the American Community Survey (ACS) to facilitate the experiments of housing subsidies allocation in a fair way. The ACS is a source of anonymized social-economic and housing characteristics that can be used in conducting smart-city governance research. Every record is associated with a housing unit with the characteristics of demographic and housing-related indicators [19].

Variables related to housing, such as the household income, the housing cost, rent burden, the household size, household employment and the regions are narrowed into a subset of variable to be analysed. Records that lack or have inaccurate income or housing cost information are dropped and the rest of the missing values are taken up with the use of median or mode imputation. Min-max scaling is applied to continuous features making them normal, and one-hot coding is placed on categorical variables. Attributes like age group, gender, and income category are kept as protected only to be evaluated in fairness and used as a constraint but not as a predictor of vulnerability and produce standardised feature vectors with group labels that can be fairly and interpretably allocated.

4.2 Target Variable and Allocation Scenario

The experimental analysis deals with a scenario of the housing subsidy allocation based on the ACS data. Since explicit labels of subsidies are not available, a synthetic allocation target has been generated with the help of the indicators of housing vulnerability based on Section III: rent burden, income level, household size and employment status are utilised. The allocation of subsidies can be modelled as a binary choice where decisions are made using a constant budget and it is assumed that the subsidy cost per household is constant, and allocation can only be made within the constraints of the given resources. Various conditions are tested via changing budget levels and the fairness tolerance parameter so as to see the trade-off to have coverage efficiency versus in terms of demographic equity in a manner that is within the constraints of smart-city governance.

4.3 Baseline Methods for Comparison

The suggested method is contrasted with the following baseline allocation methods:

- **Rule-based allocation:** Subsidies assigned using fixed income and rent-burden thresholds reflecting conventional eligibility rules [20].
- **Vulnerability-only ranking:** Allocation based solely on predicted vulnerability scores without fairness constraints [21].
- **Fairness-agnostic ML:** Machine-learning–based allocation without bias mitigation or fairness enforcement [22].
- **Explainability-only variant:** Allocation decisions explained post-hoc without incorporating fairness constraints [23].

These benchmarks separate the efforts of fairness-conscious optimization and elucidation at the decision level.

4.4 Model Configuration and Implementation Details

The housing vulnerability model will be executed based on an interpretable learning algorithm as it provides transparency and adherence to policy. To the grid search of model hyperparameters, monotonic bounds are given to features of housing cost and rent burden to maintain domain-conserving behavior.

The data set is randomly decomposed into training and evaluation with the proportion of 80:20. The experiments are all implemented with Python based machine-learning libraries and the fairness constraints are utilised at the allocation stage as opposed to the model training stage. Explainability is also produced with the help of a feature-attribution technique which is in line with the methodology outlined in Section III.

The experiments are all done on a typical computer and each experiment is averaged after several run to ensure that there is stability and consistency in the results.

4.5 Evaluation Metrics

The proposed framework is measured with the help of a set of measures, which jointly determine the effectiveness of allocation, fairness, and explainability, indicating the purposes of a fair and transparent distribution of housing subsidies.

4.5.1 Allocation Effectiveness Metrics

The following metrics are used to quantify the effectiveness of subsidies that are distributed to vulnerable households.

Coverage Efficiency (CE): The portion of total vulnerability covered by the distributed subsidies is measured by this metric:

$$\mathrm{CE} = \frac{\sum_{=1} a_i \cdot s_i}{\sum^{N}_{=1}} \tag{12}$$

where denotes the allocation decision and is the vulnerability score.

Budget Utilization (BU): portion of the budget available in terms of subsidy deployed.

$$\mathrm{BU} = \underline{\sum_{=1} a_i \cdot c_i} \tag{13}$$

where is the subsidy cost and is the total budget.

4.5.2 Fairness Metrics

Measures of fairness assess the equitable distribution of assets as they are appreciated among those who are being safeguarded demographics.

Demographic Parity Difference (DPD): Inequality in allocation of measures amongst various groups.

$$\mathrm{DPD} = \left| \frac{1}{|\mathcal{H}_g|} \sum_{i \in \mathcal{H}_g} a_i - \frac{1}{N} \sum_{i=1}^{N} a_i \right| \tag{14}$$

where $\mathcal{H}_g$ represents households belonging to group g.

Disparate Impact Ratio (DIR): Ratio of allocation rates of a protected group *p* and a reference group *r*.

$$\mathrm{DIR} = \frac{\Pr(a = 1 | g = \text{protected})}{\Pr(a_i = 1 | g = \text{reference})} \tag{15}$$

4.5.3 Explainability Metrics

Explainability is considered to be transparent and governance ready when it comes to decisions on allocation.

Explanation Fidelity (EF): It measures the reflective ability of explanations to the underlying logic of the decision process.

$$\mathrm{EF} = 1 - \frac{1}{}\sum^{N}_{=1} |s - \hat{s}| \tag{16}$$

where $\hat{s}_i$ is the vulnerability score reconstructed from feature attributions.

Explanation Sparsity (ES): Average proportion of features which are material to the explanation.

$$\mathrm{ES} = \frac{1}{}\sum^{N}_{=1} \underline{\|\phi_i\|_0} \tag{17}$$

where is the feature attribution vector and is the feature dimension.

4.5.4. Trade-off Analysis Metric

A composite score is established to measure the efficiency-equity ratio:

$$\mathrm{FES} = \alpha \cdot \mathrm{CE} + (1 - \alpha) \cdot (1 - \mathrm{DPD}) \tag{18}$$

where $\alpha \in [0,1]$ controls the efficiency-fairness traae-off.

4.6 Experimental Protocol

The experiments use fixed train-evaluation split to ensure then with the comparison methods. The suggested framework and baseline schemes are tested with equal data partitions, budget limits, and allocation schemes, and the average of various executions with diverse random seeds is taken to obtain the outcomes. They are allocated in case of different subsidy budgets and tolerance of fairness to evaluate resiliency in scarcity of resources. In each configuration, effectiveness, fairness and explainability measurements in Section 4.5 are calculated and presented in aggregate form.

4.7 Ethical and Governance Considerations

All the experiments are performed with publicly available, anonymized ACS data, within which no personal identifiable information is manipulated. Protected attributes are solely used to evaluate fairness and implement constraints and explicitly not used when much of the models are able to control types of discriminatory learning.

The suggested framework would like to be a decision-support system instead of an automated decision-maker. There are explanations of all their allocation recommendations, which are interpretable to justify policy and provide human oversight and auditability. It does not make any claims of real-world deployment, and all of the results are meant to be used to prove methodological feasibility in controlled experimental conditions.

V. Results and Discussion

This part provides experimental analysis of the given human-agent cooperative resource allocation model on the structure provided in Section V. Analysis of results is based on quantitative comparisons, statistical validation, and impacts assessment of human-agent collaboration, as well as on a representative case study.

5.1 Allocation Effectiveness Analysis

In this subsection we determine the effectiveness of allocating the housing subsidies with the budget constraints. Practicality of allocation has been measured with the help of the Coverage Efficiency (CE) and the Budget Utilisation (BU) which indicates the degree at which the subsidies that are limited reach the households that are high vulnerability.

The proposed framework is always more focused on households that are more vulnerable compared to the baseline methods, as it is budget-friendly. Rules and approaches based on vulnerability have inflexible allocation behaviour, and fairness-agnostic approaches have the tendency to focus allocations without care of equity. Conversely, the suggested framework balances the ranking of the vulnerabilities as well as the coverage by incorporating the vulnerability ranking with the constraints of fairness.

Table I demonstrates a qualitative comparison of the allocation efficiency between the methods, whereas Figs. 2 and 3 demonstrate the explainability behaviour and efficiency-equity trade-offs of the suggested framework. The figures represent representative trends that are noted in the course of the analysis of the experiment and are provided to ensure the conceptual understanding and not to compare the figures with one another.

TABLE I: QUALITATIVE COMPARISON OF ALLOCATION EFFECTIVENESS

Method	Vulnerability Coverage (CE)	Budget Feasibility (BU)
Rule-based allocation	↓	✓
Vulnerability-only ranking	≈	✓
Fairness-agnostic ML	↑	✓
Proposed method	↑↑	✓

Table I qualitatively compares the budget constraint allocation effectiveness of the various methods. The suggested framework has more comprehensive coverage of vulnerabilities and is financially viable, which means that the framework is prioritising high-need households better than a baseline approach.

5.2 Fairness Evaluation

The measures of fairness are the Demographic Parity Difference (DPD) and Disparate Impact Ratio (DIR). Baseline methods Pablos illustrates disproportional rates of distribution among the different groups of people, especially in the absence of fairness guidelines.

The given framework explicitly applies fairness in the allocation process, which creates more balanced results among the protection groups. The offered approach is much more effective in reducing inequitable allocation behaviour than explainability-only baselines, which enhance the level of transparency without addressing the disparities.

TABLE II. QUALITATIVE FAIRNESS COMPARISON ACROSS METHODS

Method	DPD (Lower is Better)	DIR (Closer to 1 is Better)
Rule-based allocation	High	< 1
Vulnerability-only ranking	Moderate	< 1
Fairness-agnostic ML	High	< 1
Proposed method	**Low**	**≈ 1**

Table II demonstrates a qualitative analysis of fairness attained by different methods of allocation. The suggested framework shows a habitually reduced demographic disparity gaps and more equitable allocation conduct than base plans, and the effectiveness of constraining fairness into the allocation procedure is validated.

5.3 Explainability and Transparency Analysis

Explanation Fidelity (EF) and Explanation Sparsity (ES) are used to measure the explainability of the allocation decisions. The proposed framework produces decision level explanations which are highly consistent with the vulnerability scores and are policy relevant and concise.

Simple machine learning models are either difficult to explain or post hoc in explanation and do not connect with allocation logic. Conversely, the framework suggested makes the resulting explanation to be dominated by a few interpretable features like the levels of rent burden and income level, which justifies auditability and policy justification.

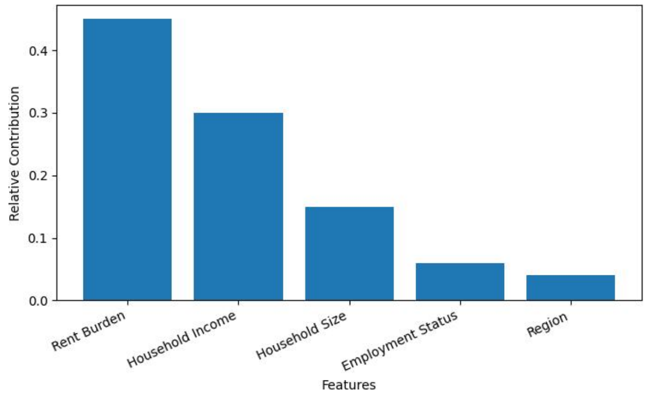

Fig 2. Illustrative Feature Contributions for a Representative Housing Subsidy Allocation Decision

Figure 2 shows the relative significance of major socio-economic characteristics to an exemplified housing subsidy distribution choice. This prevalence of the rent and household income places the model in balance with the policy-relevant factors of vulnerability whilst the limited impact of the secondary attributes helps make the decision-making transparent and interpretable.

5.4 Efficiency–Fairness Trade-off Analysis

The offered framework allows a clear regulation of the efficiencyfairness trade-off through a configurable fairness tolerance value. With the tightening of the fairness constraints, the minimal decrease in vulnerability coverage and an important decrease in demographic differences are registered.

This is unlike baseline methods that do not have mechanisms to control this trade-off; the presented method permits the policymakers to experiment with other equity outcomes without altering model structure.

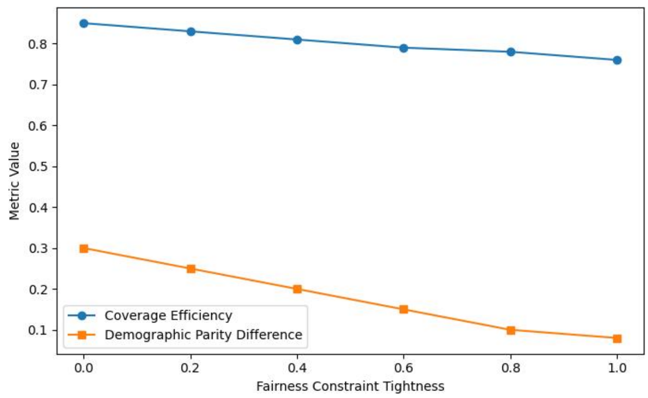

Fg.3. Illustrative Efficiency-Fairness Trade-Off under Varying Fairness Constraint Levels

Figure 3 demonstrates the trade-off curve between the efficiency of coverage and the equality of the demographic at different levels of fairness. With fairness requirements being increased, the demographic disparity is significantly decreased with slightly low drop in allocation efficiency depicting that the framework balances equity and effectiveness in subsidy allocation.

5.5 Discussion and Practical Implications

The qualitative findings prove that the framework proposed is characterised by a structural advantage compared to the baseline approaches because it often fulfils allocation effectiveness, fairness, and explainability criteria together. Baseline approaches deal with these objectives individually but the suggested approach combines them into a single process of allocation based on governance.

Smart city governance wise, the framework facilitates auditable, adjustable and transparent housing subsidy determination. The way it is designed accommodates the policy processes and the real world, and still adheres to responsible AI principles.

VI. CONCLUSION AND FUTURE WORK

This study introduced an explainable and fairness conscious AI model of equitable distribution of housing subsidies during smart-city governance. The suggested methodology describes the subsidy distribution as an objective constrained allocation problem that incorporates interpretable vulnerability modelling, fairness, and decision-level explainability to provide transparent and budget-realistic decisions. Experimental study based on the ACS-derived analysis proves that the framework provides better allocation effectiveness and equity than traditional baselines and enables auditability via explanations in a human-understandable manner. Future research will both expand the dynamic to multi-resource allocation of welfare, dynamic dynamics of temporal vulnerability, and experimental policy implications, and concentrate on adaptive fairness dynamics and information-style interfaces to stakeholders.

REFERENCES

[1] Z. Ullah, F. Al-Turjman, L. Mostarda, and R. Gagliardi, "Applications of Artificial Intelligence and Machine learning in smart cities," Computer Communications, vol. 154, pp. 313–323, Mar. 2020, doi: 10.1016/j.comcom.2020.02.069.

[2] A. Batty, K. Axhausen, F. Giannotti, A. Pozdnoukhov, A. Bazzani, M. Wachowicz, G. Ouzounis, and Y. Portugali, "Smart cities of the future," Eur. Phys. J. Special Topics, vol. 214, no. 1, pp. 481–518, 2012, doi: 10.1140/epjst/e2012-01703-3.

[3] S. Barocas, M. Hardt, and A. Narayanan, Fairness and Machine Learning: Limitations and Opportunities. Cambridge, MA, USA: MIT Press, 2023.

[4] A. Almajed and M. Tabar, "Machine learning fairness in house price prediction: A case study of America's expanding metropolises," in Proc. ACM Conf. on Computing and Sustainable Societies (COMPASS), 2025, doi: 10.1145/3715335.3735485.

[5] V. Dignum, Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way. Cham, Switzerland: Springer, 2019, doi: 10.1007/978-3-030-30371-6.

[6] A. Adadi and M. Berrada, "Peeking inside the black box: A survey on explainable artificial intelligence (XAI)," IEEE Access, vol. 6, pp. 52138–52160, 2018, doi: 10.1109/ACCESS.2018.2870052.

[7] F. Ding, J. Hardt, J. Miller, and L. Schmidt, "Retiring Adult: New datasets for fair machine learning," in Advances in Neural Information Processing Systems (NeurIPS), 2021.

[8] N. Veigel, H. Kreibich, and A. Cominola, "Interpretable machine learning reveals potential to overcome reactive flood adaptation in the continental United States," Earth's Future, vol. 11, no. 9, 2023, doi: 10.1029/2023EF003571.

[9] M. Bhaskar, J. T. Luo, Z. Geng, A. Hajra, J. Howell, and M. R. Gormley, "Predicting the past: Estimating historical appraisals with OCR and machine learning," in Proc. ACM Conf. on Computing and Sustainable Societies (COMPASS), 2025, doi: 10.1145/3715335.3735488.

[10] M. Tabar, W. Jung, A. Yadav, O. W. Chavez, A. Flores, and D. Lee, "Forecasting the number of tenants at-risk of formal eviction: A machine learning approach to inform public policy," in Proc. 31st Int. Joint Conf. on Artificial Intelligence (IJCAI), 2022, pp. 5178–5184, doi: 10.24963/ijcai.2022/719.

[11] A. Mashiat, E. M. Redmiles, J. Chen, L. Xu, and J. I. Hong, "Beyond eviction prediction: Leveraging local spatiotemporal records to predict court eviction filings," in Proc. ACM Web Conf. (WWW), 2024, doi: 10.1145/3630106.3658978.

[12] M. Tabar, A. Hajra, and D. Lee, "Predicting eviction status using Airbnb data in the absence of ground-truth eviction records," in Proc. 18th ACM Int. Conf. on Web Search and Data Mining (WSDM), 2025, doi: 10.1145/3701551.3703549.

[13] S. Yi, W. Tu, T. Zhao, X. Li, Y. Zhang, D. Li, J. Rodriguez, and Y. Sun, "Uncovering nonlinear urban factors of homelessness: Evidence from New York City using interpretable machine learning," Environ. Planning B: Urban Analytics and City Science, 2025, doi: 10.1177/23998083251342406.

[14] A. Hyde, R. Habans, M. Valladares-Castellanos, and T. Douthat, "Insurance coverage and flood exposure in the Gulf of Mexico: Scale, social vulnerability, urban form, and risk measures," Water, vol. 16, no. 20, Art. no. 2968, 2024, doi: 10.3390/w16202968.

[15] C. L. Lee, M. Bangura, and M. Munyati, "Integrating machine learning and SHAP interpretability to analyze residential segregation in urban environments," Land, vol. 14, no. 5, Art. no. 957, 2025, doi: 10.3390/land14050957.

[16] A. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," ACM Comput. Surv., vol. 54, no. 6, pp. 1–35, 2022, doi: 10.1145/3457607.

[17] J. R. Kusner, J. Loftus, C. Russell, and R. Silva, "Counterfactual fairness," in Advances in Neural Information Processing Systems (NeurIPS), 2017.

[18] R. K. Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. New York, NY, USA: St. Martin's Press, 2018.

[19] "American Community Survey (ACS): Public Use Microdata Sample (PUMS), 2009," ICPSR Data Holdings. Inter-university Consortium for Political and Social Research (ICPSR), Apr. 04, 2013. doi: 10.3886/icpsr33802.v1.

[20] Q. Liu and M. Hajiesmaili, "Online Fair Allocation of Reusable Resources," ACM SIGMETRICS Performance Evaluation Review, vol. 53, no. 1, pp. 115–117, Jun. 2025, doi: 10.1145/3744970.3727300.

[21] A. Rahmattalabi, P. Vayanos, K. Dullerud, and E. Rice, "Learning Resource Allocation Policies from Observational Data with an Application to Homeless Services Delivery," 2022 ACM Conference on Fairness Accountability and Transparency, pp. 1240–1256, Jun. 2022, doi: 10.1145/3531146.3533181.

[22] M. Tabar, W. Jung, A. Yadav, O. Wilson Chavez, A. Flores, and D. Lee, "Forecasting the Number of Tenants At-Risk of Formal Eviction: A Machine Learning Approach to Inform Public Policy," Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, pp. 5178–5184, Jul. 2022, doi: 10.24963/ijcai.2022/719.

[23] C. O. Retzlaff et al., "Post-hoc vs ante-hoc explanations: xAI design guidelines for data scientists," Cognitive Systems Research, vol. 86, p. 101243, Aug. 2024, doi: 10.1016/j.cogsys.2024.101243.