

Article

Prioritising the Adaptive Reuse of Closed Schools in Depopulating Regions: Reconciling Urgency and Potential Through a Multi-Criteria Framework

Jinju Jung ¹, Inkwan Paik ², Junhyuk Lim ² and Seunguk Na ^{3,*} 

¹ Department of Architecture, Cheongju University, 298, Daeseong-ro, Cheongwon-gu, Cheongju-si 28503, Chungcheongbuk-do, Republic of Korea; jinjuj@cju.ac.kr

² Department of Architectural Engineering, Cheongju University, 298, Daeseong-ro, Cheongwon-gu, Cheongju-si 28503, Chungcheongbuk-do, Republic of Korea; paik@cju.ac.kr (I.P.); ljh0330@cju.ac.kr (J.L.)

³ Department of Architectural Engineering, Dankook University, 152, Jukjeon-ro, Yongin-si 16890, Gyeonggi-do, Republic of Korea

* Correspondence: drseunguk@dankook.ac.kr

Abstract

The accelerating closure of schools in depopulating regions is leaving a growing surplus of public assets whose reuse must be prioritised, yet systematic and transferable tools for supporting such decisions in advance remain scarce. This study proposes an artificial-intelligence-augmented multi-criteria framework that prioritises the adaptive reuse of closed schools using only openly available demographic and spatial data. Four criteria—regional ageing, building floor area, and proximity to administrative and transport infrastructure—were evaluated for 121 closed schools in Chungbuk Province, South Korea, under two weighting schemes, the subjective analytic hierarchy process and the objective entropy method, with a large-language-model agent added as an explanatory layer to interpret context and recommend reuse types. The data-driven weights proved liable to a structural distortion, elevating the single largest building to first place in urgency on the strength of its size alone, a misjudgement the agent corrected through contextual reasoning over the same data. Examining the schools from the opposed standpoints of intervention urgency and reuse potential further revealed a near-perfect inversion between them (Spearman $\rho = -0.998$), indicating that the most urgent schools are systematically those least able to sustain a market-led conversion. The framework addresses this dilemma not through a single optimal ranking but through spatially differentiated, agent-generated recommendations and is formulated for transfer to other middle-income economies approaching the same demographic transition.

Keywords: adaptive reuse; closed schools; depopulation; multi-criteria decision-making; entropy weighting; large-language-model agent; urban shrinkage; spatial planning



Academic Editor: Alessandra Mobili

Received: 16 June 2026

Revised: 7 July 2026

Accepted: 10 July 2026

Published: 14 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

An empty schoolyard, where the sound of children has given way to the wind, is among the most visible signs of a demographic transition that statistics alone fail to convey. For much of the twentieth century, schools were built as populations concentrated and communities expanded; today, in a striking reversal, they are being closed as populations age, disperse, and decline. A shuttered schoolhouse is therefore not merely a vacant building but the physical condensation of a structural shift that has, until recently, remained an abstraction expressed only in birth-rate curves and long-range population projections.

In the Republic of Korea, this reversal is unfolding with particular speed. In 2025 alone, 49 elementary, middle, and high schools were closed across the country, of which roughly 88% were located in rural areas [1]; over the preceding five years, a total of 137 schools were shuttered for the same reasons of enrolment decline and consolidation [2–4]. The phenomenon is no longer confined to the countryside: even Seoul, long regarded as demographically resilient, has announced a five-year strategy committing approximately KRW 273.2 billion to convert closed and relocated school sites into future educational and community spaces [5]. Each closure leaves behind a sizeable public asset whose subsequent use—or disuse—carries lasting fiscal and social consequences.

Although the Korean case is acute, the underlying challenge is not unique to Korea. Population decline and ageing are spreading well beyond a handful of advanced economies, and a growing number of countries approaching the same demographic turning point will soon confront comparable surpluses of underused public facilities. Advanced economies have begun to respond in a systematic manner; the OECD, for instance, has framed the task as one of “shrinking smartly,” calling for spatial planning that anticipates demographic change rather than merely reacting to it [6]. Many middle-income countries, however, are likely to experience the same transition faster and with fewer fiscal and institutional resources yet remain markedly less prepared for it. In this respect, Korea may be regarded less as an exceptional case than as a forerunner whose experience offers an early and instructive view of a problem that others will soon face.

The stakes of inaction are considerable. Where vacant assets are left unmanaged, they tend to accelerate a self-reinforcing cycle in which shrinking populations lead to underused infrastructure, rising per capita service costs, and a deteriorating built environment, which in turn encourage further out-migration [6,7]. Within this cycle, the question of which closed school should be converted, to what use, and in what order of priority ceases to be a routine administrative matter and becomes a consequential decision bearing directly on public finances and regional equity. The need for sound and defensible decisions is thus both urgent and recurrent.

Despite this need, the reuse of closed schools has to date been approached in a largely ad hoc, case-based, and retrospective manner, and systematic tools for supporting such decisions in advance remain scarce; in one frequently cited survey, only a single school district among twelve was found to possess a formal procedure for repurposing closed schools [8]. Where structured methods have been applied, the literature on adaptive reuse has relied predominantly on multi-criteria decision-making (MCDM), and on the analytic hierarchy process (AHP) in particular, typically in combination with GIS or building information modelling [9–11]. Two limitations recur. First, this body of work has concentrated on heritage and single-building cases rather than on the portfolio-level prioritisation that a wave of closures demands. Second, the weights upon which such methods depend are commonly static and subjective, reflecting expert judgement that is difficult to validate [12]. In parallel, large-language-model (LLM) and agent-based approaches to urban decision support have advanced rapidly [13], yet, to the best of our knowledge, they have not been brought to bear on the reuse of closed schools or comparable surplus public assets under conditions of depopulation. In short, while case examples and individual techniques are by no means absent, systematic and transferable decision frameworks for this specific problem remain rare.

Accordingly, this study proposes an artificial-intelligence-augmented multi-criteria framework for prioritising the adaptive reuse of closed schools in depopulating regions, designed to rely solely on openly available demographic and spatial data. The framework couples data-driven, entropy-based weighting with the explainable reasoning of an LLM-based agent so as to mitigate the static and subjective character of conventional MCDM

while preserving interpretability. A further and deliberate feature of the framework is that it evaluates each school from two opposing perspectives—the urgency of intervention and the potential for successful reuse—so as to expose, rather than to average away, the tension between them. The framework is developed and validated using closed schools in a single Korean region as a test case but is deliberately formulated as a transferable approach whose logic and data requirements can be carried over to other middle-income countries confronting similar demographic pressures.

Thus, in this research, we aim to determine whether the reuse priority of closed schools can be assessed systematically from public data alone, whether the proposed combination of data-driven weighting and agent-based reasoning can reduce the weighting limitations of conventional MCDM, and whether the resulting recommendations are consistent with the utilisation status actually observed. In doing so, the study offers three contributions: it addresses, for the first time at the portfolio level, the prioritisation of closed-school reuse under depopulation; it advances a hybrid MCDM–data-driven–agent method that replaces subjective weighting with data-driven and explainable reasoning and in doing so reveals a structural trade-off between intervention urgency and reuse potential that a single ranking would conceal; and it reframes the surplus of closed schools as an early indicator of the demographic pressures facing middle-income economies. The remainder of this paper is organised as follows. Section 2 reviews related work on the reuse of vacant public facilities, urban shrinkage, and AI-based decision support. Section 3 describes the proposed framework, the data, and the methods. Section 4 presents the results and their validation, and the final sections discuss the implications, limitations, and directions for future research.

2. Related Work

2.1. Urban Shrinkage and the Management of Surplus Public Assets

Population decline and ageing, once regarded as concerns specific to a few advanced economies, are now recognised as a broad structural transition affecting an expanding set of countries [6,14]. A recurring consequence of this transition is the accumulation of underused public infrastructure. The OECD has characterised the resulting challenge as one of “shrinking smartly,” warning that, where it is left unaddressed, depopulation can set in motion a self-reinforcing cycle in which declining populations produce underused infrastructure, rising per capita service costs, and a deteriorating built environment, which together encourage further out-migration [6,15]. Empirical studies have documented this dynamic in concrete terms, for example through GIS-based analyses projecting steep increases in vacant dwellings across rural regions over relatively short horizons [16]. This body of work, however, has concentrated overwhelmingly on high-income and OECD contexts, and the distinctive position of middle-income countries—which are likely to undergo the same transition more rapidly and with fewer fiscal and institutional resources—has received comparatively little attention. Among the surplus assets generated by this process, closed schools are among the most numerous and the most socially significant, since they have historically served not only as places of education but as anchor institutions within their communities [17].

2.2. Adaptive Reuse Decision-Making and Multi-Criteria Methods

The reuse of vacant buildings, including former schools, has to date been approached in a largely ad hoc, case-based, and retrospective manner. Surveys of practice have found that formal procedures for deciding how to repurpose closed schools are rare; in one frequently cited study, only a single district among twelve possessed such a procedure [8]. Much of the existing literature consists of qualitative case studies and best-practice

catalogues—conversions into housing, community hubs, or cultural facilities—rather than systematic methods for deciding among alternatives [17]. Where structured methods have been employed, research on adaptive reuse has relied predominantly on multi-criteria decision-making (MCDM), and on the analytic hierarchy process (AHP) in particular, frequently in combination with GIS or building information modelling [10,11,18]. Two limitations recur across this work. First, it has focused largely on heritage and single-building cases, addressing the evaluation of one asset rather than the portfolio-level prioritisation that a wave of simultaneous closures demands. Second, the criterion weights on which these methods depend are typically static and subjective, derived from expert judgement that is difficult to validate or transfer across contexts [12]. Recent studies have begun to address the latter concern by introducing data-driven techniques to inform criterion selection and weighting [12,19], yet the integration of MCDM with data-driven methods remains at an early stage, and such efforts have generally stopped short of providing any interpretation of the rankings they produce.

2.3. Artificial Intelligence and Agent-Based Decision Support

In parallel, the application of artificial intelligence to decision support in the built environment and urban planning has advanced rapidly. Beyond established uses of computer vision for condition assessment and monitoring, recent work has explored large-language-model (LLM) and agent-based systems for urban management [20]. Such systems have been shown to process heterogeneous urban data, to improve the accuracy of contextual recommendations through retrieval-augmented generation, and to support planning tasks through multi-agent collaboration that can represent diverse stakeholder perspectives [13,21,22]. Reviews of this emerging field describe a progression from AI as a passive knowledge tool towards AI as an active partner in the planning process [13]. To the best of our knowledge, however, these capabilities have not yet been brought to bear on the reuse of closed schools or comparable surplus public assets under conditions of depopulation, where the decision problem is distinctively one of prioritising many heterogeneous assets against multiple, often conflicting, criteria.

2.4. Summary and Research Gap

Taken together, the reviewed literature indicates that the relevant fields are individually mature but remain disconnected at their intersection. The shrinking-city literature establishes the macro-level significance of surplus public assets but rarely descends to the level of specific reuse decisions and seldom addresses middle-income contexts. The adaptive-reuse and MCDM literature offers structured decision methods, but these are oriented towards single heritage buildings and rely on static, subjective weights; where they prioritise, they do so along a single axis, leaving unexamined the possibility that urgency and reuse potential may pull in opposite directions. The AI literature provides powerful data-driven and agent-based tools but has not been applied to the closed-school reuse problem. Consequently, a systematic and transferable decision framework that prioritises the adaptive reuse of closed schools under depopulation—one that combines data-driven weighting with explainable, agent-based reasoning and that confronts rather than conceals the trade-off between urgency and potential—remains absent. The present study is positioned to address this gap.

3. Materials and Methods

3.1. Study Area and Data

Chungcheongbuk-do (hereafter Chungbuk), a landlocked province in central South Korea, was selected as the study area on both substantive and methodological grounds.

Substantively, the province encapsulates the rural–urban demographic divergence that increasingly characterises ageing societies: whereas its administrative core around Cheongju continues to attract population, peripheral counties such as Boeun, Goesan, and Danyang have seen ageing rates climb beyond 42%, a level at which the depopulation of rural communities becomes effectively self-reinforcing [23]. Methodologically, this internal heterogeneity renders Chungbuk a compact analogue of the depopulation processes now unfolding across many middle-income economies so that a framework calibrated on Chungbuk may plausibly inform decision-making in comparable contexts elsewhere [6,23].

The empirical basis of the study combines two complementary sources. School-level attributes—including location, year of closure, school level, building floor area, land area, and current utilisation status—were extracted from the Standard Dataset of Nationwide Closed School Properties disseminated through the national open-data portal [24], while municipal-level demographic indicators were obtained from the resident registration statistics of the Korean Statistical Information Service and aggregated to the si-gun-gu level [25]. Because the school dataset contained no geographic coordinates, each address was geocoded through the Kakao Local API, which resolved all records without exception; straight-line distances from each school to its nearest municipal office, expressway interchange, and medical facility were subsequently derived using the Haversine formula, following the common practice of approximating accessibility by Euclidean proximity where road-network data are unavailable [26].

The analytical sample was established through a sequence of refinements. Of the 1192 closed schools recorded nationwide, 122 fell within Chungbuk; one of these proved to be a duplicate registration of a single building—identical in coordinates, land area, and closure year—and was therefore removed, leaving 121 schools for analysis. Two data-quality issues were addressed before modelling. A land-area value of 5,006,976 m², exceeding the sample median by more than four hundred-fold, was judged a manifest entry error and recoded as missing, as was a small number of building-floor-area values recorded as zero, since a genuine zero could not be distinguished from a non-entry. All such decisions were documented in a cleaning log to preserve the reproducibility of the dataset.

3.2. Evaluation Criteria Selection

The selection of evaluation criteria was guided by the dual requirement of conceptual breadth and statistical independence, the latter being essential because strong intercorrelations among criteria tend to inflate the apparent influence of redundant information in multi-criteria decision analysis [12]. An initial Pearson correlation analysis among the candidate variables exposed two pairs of near-collinear measures—total population with youth population ($r = 0.999$) and the ageing rate with the youth ratio ($r = -0.984$)—indicating that the four population-related variables in fact conveyed no more than two independent dimensions. Here the ageing rate denotes the share of residents aged 65 and over within a municipality and the youth ratio the share of those aged 14 and under, each expressed as a percentage of the total resident population: one capturing the demographic scale of a municipality and the other its age structure. Retaining more than one variable per dimension would therefore have introduced redundancy without adding explanatory value, and a single representative was accordingly chosen from each. Distance to medical facilities was likewise set aside; the dense provision of rural community clinics compressed this variable to a median of only 1.37 km, leaving it with too little variation to discriminate meaningfully among schools.

The four retained criteria were chosen so as to span distinct and policy-relevant dimensions of adaptive reuse rather than to maximise their number. The ageing rate of the surrounding municipality serves as an indicator of regional demographic decline; building

floor area represents the physical scale of the reusable asset; distance to the municipal office captures administrative and service accessibility; and distance to the nearest expressway interchange reflects transport and external connectivity. Together these four measures encompass the demographic, physical, administrative, and transport dimensions of the reuse problem while remaining mutually independent, furnishing a parsimonious yet defensible basis for the analysis that follows [27].

3.3. Two Evaluation Perspectives: Intervention Urgency and Utilisation Potential

A premise underlying the present framework is that the notion of “priority” in adaptive reuse is not singular but contingent on the policy question being posed and that a single ranking may therefore obscure as much as it reveals. Two complementary perspectives were accordingly defined over an identical set of criteria, distinguished only by the direction in which each criterion is taken to influence priority. The first, termed intervention urgency, asks which schools are most exposed to harm if left unmanaged, and under this framing the markers of decline and isolation—a high ageing rate, a large neglected structure, and remoteness from administrative and transport infrastructure—raise a school’s priority. The second, termed utilisation potential, asks instead which schools are most likely to sustain a successful conversion, so that the same large floor area now signals capacity for reuse while proximity and a younger surrounding population, rather than their opposites, elevate priority. The full directional specification of the four criteria under the two perspectives is given in Table 1.

Table 1. Descriptive statistics and directionality of the four evaluation criteria under the two perspectives ($n = 121$).

Criterion	Min.	Median	Max.	Mean (SD)	Urgency	Potential
Ageing rate (%)	18.2	29.8	43.7	32.2 (9.3)	Benefit (+)	Cost (–)
Building floor area (m ²)	40	931	25,278	1634 (2711)	Benefit (+)	Benefit (+)
Distance to municipal office (km)	0.4	11.0	27.7	11.1 (5.0)	Benefit (+)	Cost (–)
Distance to nearest interchange (km)	0.9	9.8	27.7	10.2 (5.4)	Benefit (+)	Cost (–)

Note: Benefit (+) denotes a direction in which a larger value raises priority, and cost (–) one in which a smaller value raises priority, as applied in the normalisation of the technique for order preference by similarity to ideal solution (TOPSIS). The two perspectives share an identical criterion set and differ only in directional assignment; building floor area is the sole criterion treated as benefit-type under both. The pronounced dispersion of building floor area (SD = 2711 m², against a median of 931 m²) is notable, as it becomes central to the divergence between the two weighting conditions reported in Section 4.

The decision to retain both perspectives, rather than to settle on one, is substantive rather than merely descriptive. Considered jointly, the two rankings make it possible to identify schools for which urgency and potential diverge—those most in need of intervention yet least amenable to conversion—a tension that carries direct implications for how scarce public resources ought to be allocated and that no single composite score could surface on its own. This divergence is examined empirically in Section 4.5.

3.4. Weighting and Ranking Methods

To isolate the contribution of the weighting approach itself, two weighting conditions were constructed and then evaluated through a common ranking procedure. The first, designated Condition A, derives criterion weights through the analytic hierarchy process (AHP), in which pairwise comparisons among the criteria are synthesised into a weight vector [10,11,18]. Rather than eliciting these comparisons through a large-scale expert survey, the four members of the research team—all specialists in architectural and construction engineering—jointly constructed the pairwise-comparison matrix to reflect the intervention-urgency rationale developed above, attributing the greatest relative im-

portance to the ageing rate as the most direct expression of regional decline, with the two accessibility measures and the building floor area following in turn. The resulting weights were subsequently reviewed and validated by a panel of ten professors of architectural engineering and architecture, whose concurrence confirmed the plausibility of the derived priorities. Given a comparison matrix $A = [a_{ij}]$ in which a_{ij} expresses the relative importance of criterion i over criterion j , the weight of each criterion was obtained by the row-geometric-mean method,

$$w_i = \frac{\left(\prod_{j=1}^n a_{ij}\right)^{1/n}}{\sum_{k=1}^n \left(\prod_{j=1}^n a_{kj}\right)^{1/n}} \quad (1)$$

and the internal coherence of the matrix was assessed through the consistency ratio $CR = CI/RI$, where $CI = (\lambda_{\max} - n)/(n - 1)$ is the consistency index and RI the corresponding random index. The constructed matrix yielded $CR = 0.0038$, an order of magnitude below the conventional 0.1 threshold, confirming that the preferences were internally coherent [10,11].

The second, designated Condition C, derives weights through the entropy method, which infers the relative informativeness of each criterion directly from the dispersion of its observed values, assigning greater weight to criteria that discriminate more sharply among alternatives [28,29]. For a decision matrix of m alternatives and n criteria, each column is first normalised to $p_{ij} = x_{ij}/\sum_{i=1}^m x_{ij}$; the entropy of criterion j is then

$$e_j = -\frac{1}{\ln m} \sum_{i=1}^m p_{ij} \ln p_{ij} \quad (2)$$

and its weight follows from the degree of diversification $d_j = 1 - e_j$ as

$$w_j = \frac{d_j}{\sum_{k=1}^n d_k} \quad (3)$$

so that criteria exhibiting greater dispersion receive proportionally greater weight. Because it requires no externally supplied preferences or ground-truth labels, the entropy method is particularly suited to the present problem, in which no authoritative prioritisation of closed schools exists against which weights might otherwise be calibrated. The juxtaposition of a subjective, expert-informed scheme with an objective, data-driven one is deliberate, and the contrast between the two is itself a principal object of analysis.

So that any difference in the resulting rankings could be attributed to the weights alone, both conditions were carried through the same aggregation procedure, TOPSIS [30]. Each criterion value was first vector-normalised to $r_{ij} = x_{ij}/\sqrt{\sum_{i=1}^m x_{ij}^2}$ and weighted as $v_{ij} = w_j r_{ij}$, with its orientation set according to the directional specification of the perspective under consideration. The positive and negative ideal solutions, A^+ and A^- , were identified from the best and worst weighted values of each criterion, and for every school the Euclidean separations from them,

$$D_i^+ = \sqrt{\sum_{j=1}^n (v_{ij} - v_j^+)^2}, D_i^- = \sqrt{\sum_{j=1}^n (v_{ij} - v_j^-)^2} \quad (4)$$

were combined into a closeness coefficient

$$C_i = \frac{D_i^-}{D_i^+ + D_i^-} \quad (5)$$

whose descending order yields the final ranking. Holding this procedure constant across the two conditions ensures that the comparison reported below reflects the influence of the weighting method rather than of the aggregation rule. The complete analytical workflow, from data preparation through to the explanatory layer described next, is summarised in Figure 1.




Figure 1. Analytical workflow of the proposed framework.

3.5. LLM Agent Explanatory Layer

The weighting and ranking procedures described above yield ordinal priorities, but they cannot account for the meaning of a particular score nor adjudicate between rankings that disagree. To supply this interpretive capacity, an explanatory layer built upon a large-language-model (LLM) agent was incorporated into the framework. It bears emphasis that the agent plays no part in determining criterion weights; its function is confined to interpreting the context of each school, reconciling discrepancies between the two weighting conditions, and proposing a reuse type commensurate with the school’s circumstances [31].

The agent was governed by a two-part prompt designed to constrain its reasoning to the supplied evidence. A system instruction cast the model in the role of an urban and rural planning analyst and required, in explicit and repeated terms, that its analysis draw solely on the numerical values provided: the model was directed to “analyse strictly on the basis of the data values provided” and instructed not to “infer or add facts that are not present in the data,” with the response further fixed to a five-field JSON schema admitting no free-form text. For each school, a structured user prompt then assembled the school’s basic attributes, the four criterion values—each annotated with its sample range or median so that the model could situate the case within the distribution—and the AHP and entropy ranks together with the signed difference between them.

A notable feature of this construction is a conditional safeguard against the very distortion later identified in Section 4.2. Where a school’s building floor area exceeded a threshold of 2800 m², approximately three times the sample median, a warning was automatically inserted into the prompt requiring the agent to “explicitly examine the possibility that the ranking has been overestimated by the entropy weight on building floor area.” The agent was thereby led to interrogate potential size-driven overestimation as a

matter of design rather than chance—a property that distinguishes the present approach from an unstructured query to a language model and to which the corrective behaviour reported in Section 4.3 may be directly traced. The overall structure of the prompt is shown in Figure 2.




Figure 2. Structure of the LLM agent prompt and its constrained output.

The agent returned, for each school, an adjusted urgency level with its supporting rationale, a judgement as to which of the two rankings was the more defensible where they diverged, and a recommended reuse type with justification, each field required to cite the relevant figures. Malformed responses were re-requested up to twice before being flagged, and the layer was implemented using Claude Sonnet [32]; all 121 schools were processed without parsing or execution failure.

4. Results

4.1. Criteria Weights Under the Two Conditions

The two weighting conditions produced markedly different pictures of which criteria should govern the prioritisation, and the nature of their disagreement proves central to the analysis that follows. Under Condition A, the pairwise comparisons yielded a weight of 0.423 for the ageing rate, 0.227 each for the distances to the municipal office and the nearest interchange, and 0.123 for building floor area, a distribution that mirrors the intervention-urgency rationale in which demographic decline is taken to be the most telling signal of need. Under Condition C, by contrast, the entropy method assigned an overwhelming weight of 0.529 to building floor area, with the ageing rate, interchange distance, and municipal-office distance receiving 0.231, 0.145, and 0.095 respectively (Table 2). The divergence is most striking precisely where the two schemes treat the same variable in opposite terms: the criterion awarded the least importance by expert judgement is the one the data-driven method elevates above all others (see Figure 3). This inversion arises because the entropy method responds to statistical dispersion rather than to substantive meaning, and building floor area, ranging from a few dozen to more than twenty-five thousand square metres across the sample, exhibits by far the greatest variability of the four criteria.

Table 2. Criteria weights under the two weighting conditions.

Criterion	Condition A (AHP)	Condition C (Entropy)
Ageing rate (%)	0.423	0.231
Building floor area (m ²)	0.123	0.529
Distance to municipal office (km)	0.227	0.095
Distance to nearest interchange (km)	0.227	0.145
Sum	1.000	1.000

Note: Condition A weights were derived from expert pairwise comparison and yielded a consistency ratio of CR = 0.0038, well below the conventionally accepted threshold of 0.1. Condition C weights were obtained by the entropy method from the observed dispersion of the data. The two schemes invert the relative standing of building floor area, which expert judgement ranks least important yet the data-driven method elevates above all other criteria.




Figure 3. Comparison of criteria weights under the two conditions. Note: building floor area, ranked least important under the AHP condition, is elevated above all other criteria by the entropy method; AHP CR = 0.0038.

4.2. Rank Divergence and the Limits of Data-Driven Weighting

When the two weight vectors were carried through the common TOPSIS procedure, the resulting urgency rankings agreed substantially in their overall structure, with a Spearman rank correlation of 0.83, yet diverged sharply at points that prove more consequential than the aggregate statistic suggests. The most instructive of these concerns the former Chungbuk Physical Education High School, which the entropy condition ranked first in urgency among all 121 schools, while the AHP condition placed it 108th. The source of this discrepancy is unambiguous: with a building floor area of 25,278 m²—some twenty-seven times the sample median—the school dominates the very criterion to which the entropy method assigns more than half of the total weight and is propelled to the top of the ranking on the strength of its size alone, notwithstanding an ageing rate of 18.2% and a location barely 1.5 km from the municipal centre that together mark it as anything but an urgent case.

This was not an isolated artefact. A recurring pattern emerged in which schools possessing unusually large buildings were systematically over-ranked by the entropy condition relative to the AHP condition: the former Namhangang Elementary School and the former Juseong Middle School, among others, exhibited rank displacements of more than eighty positions traceable to the same mechanism (see Table 3; Figure 4). The pattern indicates that the limitation is structural rather than incidental—a purely data-driven weighting, sensitive to dispersion but indifferent to meaning, will reliably mistake physical magnitude for intervention urgency wherever a single variable dominates the variance.

It is this systematic blind spot, rather than any failure of the entropy method on its own terms, that motivates the explanatory layer introduced below.

Table 3. Principal rank-discrepancy cases between the AHP and entropy conditions.

School (Former)	Floor Area (m ²)	Ageing Rate (%)	Entropy Rank	AHP Rank	Displ.
Chungbuk Physical Education HS	25,278	18.2	1	108	107
Namhangang ES	8501	26.6	11	96	85
Juseong MS	9024	18.2	37	120	83

Note: Ranks are intervention-urgency ranks among the 121 schools under the entropy condition (Condition C) and the AHP condition (Condition A); displacement (Displ.) is their absolute difference. The sample median building floor area is 931 m², against which the Chungbuk Physical Education High School building is roughly twenty-seven times larger. In all three cases an unusually large building inflates the entropy-condition rank through the dominant floor-area weight, despite attributes that mark the school as non-urgent.




Figure 4. Rank comparison between the AHP and entropy conditions. Note: each point is one of the 121 schools by urgency rank under the two conditions. The dashed diagonal denotes the line of identity ($y = x$), along which a school would hold the same rank under both conditions; the shaded band marks a tolerance of ± 15 ranks around it. Points far below the diagonal are over-ranked by entropy relative to AHP, with principal discrepancy cases labelled.

4.3. Contextual Reconciliation Through the Agent Layer

Confronted with the same schools, the agent layer arrived at judgements that the numerical rankings could not. In the case of the former Chungbuk Physical Education High School, the agent assigned an urgency level of “low,” reasoning explicitly from the inputs that an ageing rate of 18.2% lay at the bottom of the observed range and that proximity to both the municipal centre and the expressway rendered the site well served rather than isolated; it identified the first-place entropy ranking as an overestimate driven by floor area

and judged the AHP ranking the more defensible of the two. The agent thus reproduced, through contextual reasoning over the same data, the corrective that the comparison in Section 4.2 required.

Across all 121 schools, the agent's assessments cohered into an interpretable structure. Urgency was distributed across 41 schools rated high, 30 medium, and 50 low, while recommended reuse types concentrated heavily on welfare functions (75.2%), followed by cultural (11.6%) and industrial (10.7%) uses (see Figure 5)—a distribution consonant with a sample dominated by ageing rural communities, for which welfare-oriented conversion represents the most fitting response. Crucially, the agent did not reconcile discrepancies by mechanically favouring one ranking over the other. For the former Boeun Information High School, for instance, it concurred that the entropy ranking overstated the school's standing on account of building size yet nonetheless retained a high urgency rating, reasoning that an ageing rate exceeding 43% constituted an independent and overriding ground for intervention. This capacity to separate the question of which ranking is more reliable from the question of how urgent a school truly is—to weigh, in effect, several considerations at once—is precisely what the numerical methods, taken alone, could not provide.




Figure 5. Distribution of recommended reuse types across the 121 schools. Note: the agent assigned a single category per school—welfare 75.2%, cultural 11.6%, industrial 10.7%, tourism 1.7%, and other 0.8%.

4.4. Consistency with Observed Utilisation Status

If the priorities generated by the framework are to carry practical weight, they ought at a minimum to be consistent with the manner in which schools have in fact been used or left unused. Because the framework was at no point trained on utilisation status, any concordance between the two constitutes independent corroboration rather than a built-in result. Examined in these terms, the agent's urgency assessments align with observed status in the expected direction: among schools recorded as unutilised, 39.3% were rated high in urgency, the largest such proportion of any group, whereas among those already in self-directed use only 23.1% were so rated and fully 69.2% were rated low. This association is statistically significant: a chi-square test of independence between utilisation status and the agent's three-level urgency rating returned $\chi^2(4, N = 121) = 11.72$, $p = 0.020$ (Cramér's $V = 0.22$). Because several expected cell frequencies approached the conventional threshold of five, the result was corroborated by a Fisher exact test on the reduced classification of self-utilized versus other schools against low versus higher urgency, which confirmed the association (odds ratio = 4.43, $p = 0.002$); schools already in self-directed use were thus more than four times as likely to be judged of low urgency. The fact that schools already serving a function should register as least urgent, and those standing idle as most, is precisely the

ordering one would hope a defensible prioritisation to recover. The resulting priorities are not spatially uniform but concentrate in the severely ageing peripheral counties, as the geographic distribution of priority scores makes plain (see Figure 6).




Figure 6. Spatial distribution of the closed schools by intervention-urgency priority score. Note: darker shading denotes higher urgency under the AHP condition; labelled points are the principal rank-discrepancy cases.

The agent's reuse recommendations displayed a comparable sensitivity to context. Among self-utilised schools situated in ageing rural settings—where ageing rates exceeded 40% notwithstanding the buildings' current occupancy—the agent uniformly recommended welfare-oriented conversion and assigned high urgency, reasoning that demographic decline warranted an upgrading of function even where a use was already in place. Among self-utilised schools in the lower-ageing urban districts of Cheongju, by contrast, it consistently recommended cultural or industrial uses at low urgency, reflecting the greater latent demand and superior accessibility of those locations. The contrast suggests that the recommendations are not generic but responsive to the particular demographic and locational circumstances of each school. It should be noted that the difference in high-urgency proportions between unutilised and leased schools was itself modest; the significance reported above rests chiefly on the sharp separation of self-utilized schools from the rest. The concordance is therefore best understood as a statistically supported consistency check rather than as a full external validation, which would require richer outcome data of the kind discussed in Section 5.

4.5. The Urgency–Potential Trade-Off

The starkest finding of the study emerged from setting the two perspectives against one another. When the urgency ranking (Perspective B) was compared with the utilisation-potential ranking (Perspective A), the two were found to be almost perfectly inverted,

with a Spearman correlation of -0.998 under the AHP weights and -0.968 under the entropy weights; not a single school ranked highly on both. The schools most urgently in need of intervention were, with near-mechanical regularity, those least equipped to sustain a successful conversion—remote rural branch schools in severely ageing counties, whose small buildings and distance from administrative and transport infrastructure place them at once at the head of the urgency ordering and at the foot of the potential ordering. Conversely, the schools offering the greatest conversion potential were the large, well-connected facilities of low-ageing urban districts, for which the case for urgent public intervention is correspondingly weak.

This inversion is not a deficiency of the analysis but its most consequential result, for it exposes a structural dilemma that no single ranking, however carefully weighted, could resolve on its own terms. A prioritisation oriented to urgency directs attention to precisely those assets for which market-led reuse is least likely to materialise, while one oriented to potential favours assets that would in many cases be converted without public prompting; to optimise for either criterion alone is thus to misallocate effort. It is at this impasse that the contribution of the explanatory layer becomes clearest. Rather than forcing a choice between the two orderings, the agent responds to the dilemma differentially, recommending survival-oriented welfare functions for the urgent-but-low-potential schools of the depopulating periphery and value-generating cultural or industrial uses for the high-potential schools of the urban core (see Figure 7). The trade-off that defeats a purely numerical prioritisation is, in this way, met not by a single optimal ranking but by context-sensitive recommendations of a kind the agent layer is distinctively suited to provide.




Figure 7. Urgency versus utilisation potential for all 121 schools. Note: each point is positioned by its urgency rank and potential rank, where rank 1 is highest on each axis; the two rankings are near-perfectly inverted (Spearman $\rho = -0.998$ under AHP weighting, -0.968 under entropy weighting), and the empty upper-right region shows that no school ranks high on both.

5. Discussion

5.1. From Ranking to Reasoning: Methodological Implications

The principal methodological lesson of this study is that objectivity in weighting, though desirable, is not by itself sufficient. The entropy method supplied weights that were entirely free of expert preference, yet it was precisely this freedom from substantive judgement that led it to mistake the largest building in the sample for the most urgent case. Whereas prior applications of entropy-based weighting to problems of urban shrinkage and regeneration have generally treated its objectivity as an unqualified strength [33], the present findings show that this same objectivity can mislead where a single criterion dominates the variance. Refining the weighting scheme further would not have remedied the problem, for the difficulty lay not in the accuracy of the weights but in the absence of any capacity to interpret what a given configuration of values means in context. The contribution advanced here is therefore best understood not as a better weighting method but as a shift in the locus of the analysis—from ranking to reasoning. The agent layer does not displace the AHP or the entropy procedures so much as complete them, allotting to the numerical methods the tasks of transparency and reproducibility at which they excel while reserving to itself the contextual interpretation that they cannot perform. This positioning is consonant with the trajectory described in recent reviews of language models in planning, which chart a shift from artificial intelligence as a passive knowledge tool towards an active partner in the planning process [34]; the present framework locates the agent specifically at the interpretive layer, leaving weight determination to transparent numerical methods. In this division of labour lies an answer to one of the questions with which the study began: the coupling of data-driven weighting with agent-based reasoning does indeed attenuate the weighting limitations of conventional MCDM, not by perfecting the weights but by relieving them of interpretive burdens they were never equipped to bear.

5.2. The Urgency–Potential Dilemma as a Policy Problem

The near-perfect inversion of the urgency and potential rankings carries an import that extends well beyond the methodological. That the schools most in need of intervention are, almost without exception, those least capable of sustaining a market-led conversion describes not a statistical curiosity but the central predicament of reuse policy under depopulation. Where urgency and potential coincide, ordinary incentives suffice; it is where they diverge—as here they almost entirely do—that public action becomes indispensable, since the remote and severely ageing communities that head the urgency ordering are exactly those that private initiative will pass over. A framework that renders this divergence visible does more than rank assets: it reframes the policy question itself, displacing the search for a single optimal ordering with the recognition that urgent and promising schools call for categorically different responses. The agent’s differentiated recommendations—survival-oriented welfare provision for the depopulating periphery and value-generating cultural and industrial uses for the urban core—give operational form to this recognition, translating a structural tension that paralyses single-objective optimisation into a tractable basis for spatially differentiated intervention.

5.3. Transferability to Other Middle-Income Contexts

Although the framework was developed and tested on a single Korean province, its design was from the outset oriented towards transferability. The analysis draws exclusively on openly available demographic and spatial data, depends on no feature peculiar to Korean administrative practice, and rests on four criteria—demographic decline, asset scale, and two dimensions of accessibility—that are general rather than context-specific. These properties matter because the demographic transition that Korea is now experienc-

ing in acute form is approaching in many middle-income economies, which are likely to encounter it more rapidly, and with thinner fiscal and institutional reserves, than the advanced economies that have so far shaped the response. For such countries, the value of an anticipatory and systematic instrument—one that prioritises from public data before a backlog of derelict assets accumulates—is correspondingly greater. This is not to suggest that the framework may be transplanted unmodified; the calibration of criteria and the interpretation of recommendations would require adjustment to local data and circumstance. It is to suggest, rather, that the Korean experience offers not merely a cautionary illustration but a portable method and that treating the surplus of closed schools as an early signal of a wider transition may be of use well beyond the case examined here.

5.4. Limitations

Several limitations qualify these findings and delimit the claims that may reasonably be drawn from them. The analysis rests on a single province and 121 schools and is, by design, a demonstration of method rather than a comprehensive evaluation; its generalisation would benefit from application across multiple regions and a larger and more varied sample. Accessibility was approximated by straight-line distance, which, while a standard expedient where road-network data are lacking, cannot capture the travel realities that terrain and infrastructure impose, particularly in the mountainous districts that figure prominently among the urgent cases, where the divergence between Euclidean distance and actual driving time can be substantial. Future regional applications would therefore benefit from replacing Haversine proximity with network-based measures—whether GIS network routing or API-derived travel times—so as to reflect on-the-ground accessibility more faithfully. The concordance between the framework's priorities and observed utilisation status, though encouraging, was modest in magnitude and is offered as an illustrative pattern rather than as formal validation; establishing a firmer correspondence would require richer outcome data, including the specific functions to which self-utilised schools have been put, which the source dataset did not record. Finally, the explanatory layer reflects the output of a single model at a single point in time, and although its reliance on structured inputs and outputs was designed to support reproducibility, its recommendations may vary with model and prompt; systematic examination of this sensitivity remains a task for further work.

5.5. Practical Implications

Beyond its methodological contribution, the framework offers several concrete implications for the authorities charged with managing closed-school portfolios. First, because it runs entirely on openly available demographic and spatial data, it can serve as a low-cost triage instrument that screens an entire regional portfolio before scarce survey and appraisal resources are committed, directing detailed feasibility study towards the schools the framework flags as most urgent. Second, the sharp inversion between urgency and potential implies that a single disposal policy cannot serve both ends: the urgent-but-low-potential schools of the depopulating periphery call for publicly funded, welfare-oriented conversion that the market will not supply, whereas the high-potential schools of the urban core are better suited to revenue-generating cultural or industrial reuse, often through private participation. Recognising this division allows administrators to match policy instrument to asset type rather than applying uniform criteria across a heterogeneous stock. Third, the agent layer's structured, reason-bearing outputs provide a transparent and communicable basis for prioritisation decisions, which may assist public authorities in justifying resource allocation to residents and oversight bodies. Taken together, these features position the framework not as a replacement for professional judgement but as

a decision-support aid that renders large-portfolio prioritisation systematic, defensible, and transferable.

6. Conclusions

This study set out to determine whether the reuse of closed schools in depopulating regions could be prioritised systematically from openly available data and whether the limitations of conventional multi-criteria methods might be eased by coupling data-driven weighting with agent-based reasoning. Working from a sample of 121 closed schools in a single Korean province, it found that purely data-driven weights, for all their objectivity, can be misled by statistical dispersion into mistaking physical scale for intervention urgency—the entropy scheme assigning 52.9% of the total weight to building floor area and thereby elevating the single largest building from 108th to first in urgency—and that this shortcoming is structural rather than incidental. The two evaluation perspectives proved almost perfectly inverted (Spearman $\rho = -0.998$), with not a single school ranking highly on both, while the agent’s urgency assessments aligned significantly with observed utilisation status ($\chi^2 p = 0.020$; Fisher $p = 0.002$), lending empirical support to the framework’s outputs. An explanatory layer founded on a large-language-model agent was shown to correct such misjudgements through contextual reasoning over the same data, to produce recommendations consistent with the manner in which schools have actually been used, and to respond coherently to a trade-off that defeats any single ranking.

Three contributions follow. The study addresses, at the portfolio level rather than the level of the individual building, the prioritisation of closed-school reuse under depopulation, a problem hitherto approached in a largely ad hoc manner. It advances a hybrid method in which subjective weighting gives way to data-driven and explainable reasoning, shifting the locus of multi-criteria analysis from ranking towards interpretation, and in so doing, it brings to light a structural trade-off between urgency and potential that a single ordering would have left hidden. And it reframes the growing surplus of closed schools as an early indicator of demographic pressures that many middle-income economies will shortly confront. Future work might extend the framework across multiple regions, incorporate network-based measures of accessibility, and examine the stability of agent-based recommendations across models so as to consolidate what is offered here into a validated instrument as a transferable approach.

Author Contributions: Conceptualization, S.N. and J.J.; methodology, S.N. and J.J.; software, J.J. and J.L.; validation, I.P., J.L. and S.N.; formal analysis, J.J. and I.P.; investigation, J.J. and J.L.; resources, I.P. and S.N.; data curation, J.J. and J.L.; writing—original draft preparation, J.J. and S.N.; writing—review and editing, S.N. and I.P.; visualisation, J.J. and J.L.; supervision, S.N.; project administration, S.N.; funding acquisition, S.N. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2022R111A01053268).

Data Availability Statement: The data supporting the findings of this study were derived from publicly available sources. The closed-school property data are openly available from the Korean Public Data Portal (data.go.kr) as part of the Standard Dataset of Nationwide Closed School Properties, and the municipal demographic indicators are openly available from the Korean Statistical Information Service (KOSIS, kosis.kr). The processed analytical dataset of 121 closed schools, the criteria-weight matrices, and the agent-generated outputs created during this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Lee, H. 49 Schools in Korea to Close Amid Population Decline. *The Korea Times*, 24 February 2025.
- Jung, M. Over 4000 Schools Shut Down Nationwide as Student Numbers Plunge. *The Korea Times*, 28 December 2025.
- Goulas, S. *Enrollment Declines and Permanent Closures in Public Schools*; IZA Discussion Paper Series; IZA Institute of Labor Economics: Bonn, Germany, 2025.
- OECD. *PISA 2022 Results (Volume II): Learning During—And From—Disruption*; OECD Publishing: Paris, France, 2023.
- Korea JoongAng Daily. Seoul Unveils \$184 M 5-Year Plan to Transform Its Studentless Schools. *Korea JoongAng Daily*, 2026.
- OECD. *Shrinking Smartly and Sustainably: Strategies for Action*; OECD Publishing: Paris, France, 2025.
- OECD. *Adapting Regional Policy in Korea: Preparing Regions for Demographic Change*; OECD Publishing: Paris, France, 2022.
- Dowdall, E.; Warner, S. *Shuttered Public Schools: The Struggle to Bring Old Buildings New Life*; The Pew Charitable Trusts: Philadelphia, PA, USA, 2013.
- Pan, Z.; Guo, F. Digital twin (DT) development strategies in the architecture, engineering, and construction (AEC) industry based on SWOT-AHP analysis. *J. Build. Des. Environ.* **2025**, *3*, 202441. [\[CrossRef\]](#)
- Saaty, T.L. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*; McGraw-Hill: New York, NY, USA, 1980.
- Saaty, T.L. How to make a decision: The Analytic Hierarchy Process. *Eur. J. Oper. Res.* **1990**, *48*, 9–26. [\[CrossRef\]](#)
- Diakoulaki, D.; Mavrotas, G.; Papayannakis, L. Determining objective weights in multiple criteria problems: The critic method. *Comput. Oper. Res.* **1995**, *22*, 763–770. [\[CrossRef\]](#)
- Jiang, Y.; Chao, Q.; Chen, Y.; Li, X.; Liu, S.; Cong, G. UrbanLLM: Autonomous Urban Activity Planning and Management with Large Language Models. *arXiv* **2024**, arXiv:2406.12360.
- Pallagst, K.; Schwarz, T.; Popper, F.; Hollander, J. Planning shrinking cities. *Prog. Plan.* **2009**, *72*, 223–232. [\[CrossRef\]](#)
- Strykiewicz, T.; Jaroszewska, E. The process of shrinkage as a challenge to urban governance. *Quaest. Geogr.* **2016**, *35*, 27–37. [\[CrossRef\]](#)
- Lee, S. Spatial Analysis of Abandoned Houses and Their Influencing Factors in South Korea. *Appl. Sci.* **2021**, *11*, 8576. [\[CrossRef\]](#)
- Jiusto, S. Puerto Rico's Rescued Schools: A Grassroots Adaptive Reuse Movement for Abandoned School Buildings. *Soc. Sci.* **2023**, *12*, 662. [\[CrossRef\]](#)
- Saaty, T.L. Decision making with the analytic hierarchy process. *Int. J. Serv. Sci.* **2008**, *1*, 83–98. [\[CrossRef\]](#)
- Keshavarz-Ghorabae, M.; Amiri, M.; Zavadskas, E.K.; Turskis, Z.; Antucheviciene, J. Determination of Objective Weights Using a New Method Based on the Removal Effects of Criteria (MEREC). *Symmetry* **2021**, *13*, 525. [\[CrossRef\]](#)
- Spencer, B.F., Jr.; Hoskere, V.; Narazaki, Y. Advances in Computer Vision-Based Civil Infrastructure Inspection and Monitoring. *Engineering* **2019**, *5*, 199–222. [\[CrossRef\]](#)
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-T.; Rocktäschel, T.; et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 9459–9474.
- Park, J.S.; Joseph, C.; Cai, C.J.; Morris, M.R.; Liang, P.; Bernstein, M.S. Generative Agents: Interactive Simulacra of Human Behavior. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST), San Francisco, CA, USA, 29 October–1 November 2023.
- Yu, J. Comparative Study on National Policies and Educational Approaches toward Regional Revitalization in Japan and South Korea: Aiming to Achieve the Sustainable Development Goals. *Societies* **2023**, *13*, 210. [\[CrossRef\]](#)
- Ministry of Education. *Standard Dataset of Basic Information on Closed School Properties*; Public Data Portal: Sejong-si, Republic of Korea, 2024. Available online: <https://www.data.go.kr> (accessed on 25 April 2026).
- Ministry of the Interior and Safety. *Resident Registration Population by Administrative District (Si, Gun, and Gu) and Sex*; Korean Statistical Information Service (KOSIS): Daejeon, Republic of Korea, 2026. Available online: <https://kosis.kr> (accessed on 25 April 2026).
- Boscoe, F.P.; Henry, K.A.; Zdeb, M.S. A Nationwide Comparison of Driving Distance Versus Straight-Line Distance to Hospitals. *Prof. Geogr.* **2012**, *64*, 188–196. [\[CrossRef\]](#)
- Triantaphyllou, E. *Multi-Criteria Decision Making Methods: A Comparative Study*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 2000.
- Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [\[CrossRef\]](#)
- Zeleny, M. *Multiple Criteria Decision Making*; McGraw-Hill: New York, NY, USA, 1982.
- Behzadian, M.; Otaghsara, S.K.; Yazdani, M.; Ignatius, J. A state-of-the-art survey of TOPSIS applications. *Expert Syst. Appl.* **2012**, *39*, 13051–13069. [\[CrossRef\]](#)
- Mirzadeh, I.; Alizadeh-Vahid, K.; Shahrokhi, H.; Tuzel, O.; Bengio, S.; Farajtabar, M. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv* **2024**, arXiv:2410.05229.
- Anthropic. Introducing the Model Context Protocol. 2024. Available online: <https://www.anthropic.com/news/model-context-protocol> (accessed on 6 May 2026).

33. Lui, T.; Yan, T. Study on Urban Shrinkage Based on Entropy Weight-TOPSIS and GCA Method. In *Proceedings of the 2024 4th International Conference on Computational Modeling, Simulation and Data Analysis (CMSDA 2024)*; Association for Computing Machinery: New York, NY, USA, 2024.
34. Zheng, Y.; Xu, F.; Lin, Y.; Santi, P.; Ratti, C.; Wang, Q.; Li, Y. Urban planning in the era of large language models. *Nat. Comput. Sci.* **2025**, *5*, 727–736. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.