URBANALIGN: POST-HOC SEMANTIC CALIBRATION FOR VLM-HUMAN PREFERENCE ALIGNMENT

A PREPRINT

Yecheng Zhang[1]* **Rong Zhao**[2]* **Zhizhou Sha**[1] **Yong Li**[3] **Lei Wang**[4] **Ce Hou**[3] **Wen Ji**[5]
Hao Huang[1] **Yunshan Wan**[6] **Jian Yu**[7] **Junhao Xia**[1] **Yuru Zhang**[8] **Chunlei Shi**[9]*

[1]Tsinghua University [2]University College London [3]Hong Kong University of Science and Technology
[4]Peking University [5]Southwest Jiaotong University [6]Zhejiang University
[7]University of Texas at Austin [8]Renmin University of China [9]Southeast University

March 6, 2026

ABSTRACT

Vision-language models (VLMs) can describe urban scenes in rich detail, yet consistently fail to produce reliable human preference labels in domain-specific tasks such as safety assessment and aesthetic evaluation. The standard fix, fine-tuning or RLHF, requires large-scale annotations and model retraining. We ask a different question: *can a frozen VLM be aligned with human preferences without modifying any weights?* Our key insight is that VLMs are strong concept extractors but poor decision calibrators. We propose a three-stage post-hoc pipeline that exploits this asymmetry: (i) interpretable evaluation dimensions are automatically mined from consensus exemplars; (ii) an Observer–Debater–Judge chain extracts robust concept scores from the frozen VLM; and (iii) locally-weighted ridge regression on a hybrid manifold calibrates these scores to human ratings. Applied as **UrbanAlign** on Place Pulse 2.0, the framework reaches 72.2% accuracy (κ=0.45) across six perception categories, outperforming all baselines by +11.0 pp and zero-shot VLM by +15.5 pp, with full interpretability and zero weight modification.

Keywords Human Preference Alignment · Post-hoc Concept Bottleneck · Vision-Language Model · Urban Perception

1 Introduction

Large vision-language models (VLMs) excel at describing visual scenes but consistently fail to produce reliable preference labels for domain-specific tasks. This *VLM–human preference alignment gap* appears across domains, from urban perception [Zhang et al., 2025, Mushkani, 2025] to aesthetic quality [Murray et al., 2012, Wu et al., 2024], and the standard remedy is model adaptation: fine-tuning, LoRA, or RLHF [Ouyang et al., 2022], all of which modify model weights and require domain-specific training data and non-trivial compute. We ask a different question: *can a frozen VLM be aligned with human preferences in a new domain, without modifying any model weights?*

The key observation is that VLMs are already strong *concept extractors* that can identify rich visual elements in a scene (e.g., building facades, vegetation, and street infrastructure), but poor *decision calibrators*: their mapping from these features to discrete comparison labels is poorly aligned with human judgement boundaries (Fig. 1). This mirrors the finding of Concept Bottleneck Models (CBMs) [Koh et al., 2020] that routing predictions through interpretable concept scores and then a calibrated linear layer dramatically outperforms end-to-end classification. Such a bottleneck can be retrofitted *post-hoc* onto any frozen backbone [Yuksekgonul et al., 2023], and recent work further shows that concepts can be discovered directly from VLMs without manual labels [Oikarinen et al., 2023, Yang et al., 2023].

*Corresponding: zhangyec23@mails.tsinghua.edu.cn (Y. Zhang), rong.zhao.25@ucl.ac.uk (R. Zhao), 230238514@seu.edu.cn (C. Shi)

A naïve alternative is to prompt the VLM as a single-shot scorer, but this suffers from well-documented biases: position effects, anchoring, and inconsistent self-evaluation [Wu et al., 2024, Liu et al., 2024]. When asked to compare two scenes directly, a single VLM call conflates perception and judgement: the model must simultaneously identify relevant visual cues, weigh their relative importance, and produce a discrete label, all within a single forward pass. This end-to-end compression discards the structured reasoning that human raters implicitly perform when evaluating complex, multi-dimensional scenes. Multi-agent debate protocols decompose complex judgements into complementary reasoning steps and have been shown to mitigate such biases through structured disagreement [Du et al., 2024, Wang et al., 2023], yet these techniques have primarily targeted factual reasoning and text generation; their potential for domain-specific visual preference alignment remains unexplored.

This principle is operationalised as a pipeline that keeps the VLM entirely frozen, with three tightly coupled stages (Fig. 2): **(i)** *concept mining*, where interpretable evaluation dimensions are discovered by the VLM from consensus exemplars; **(ii)** *structured scoring*, where an Observer–Debater–Judge multi-agent chain extracts robust continuous concept scores from the frozen VLM; **(iii)** *geometric calibration*, where locally-weighted ridge regression (LWRR) [Cleveland, 1979] on a hybrid visual–semantic manifold aligns concept scores to human ratings. An end-to-end optimization loop unifies these stages, automatically selecting the dimension set that maximises calibrated accuracy via temperature-scheduled trials. We instantiate this framework as **UrbanAlign** on Place Pulse 2.0 [Salesses et al., 2013], a large-scale pairwise comparison dataset covering six categories of urban perception, where subjective judgements [Wilson and Kelling, 1982, Jacobs, 1961] defy the feature-to-label mappings of existing approaches [Dubey et al., 2016, Naik et al., 2014]. UrbanAlign substantially outperforms both supervised baselines and zero-shot VLM scoring across all six categories, with full dimension-level interpretability and zero model-weight modification, demonstrating that a frozen VLM already contains sufficient perceptual knowledge for human-aligned urban assessment.

Contributions.

1. **End-to-end concept mining.** We design a loop that automatically discovers interpretable evaluation dimensions from consensus exemplars and refines them via a temperature-scheduled search, forming a concept bottleneck for perception prediction.
2. **Multi-agent structured scoring.** We introduce an Observer–Debater–Judge deliberation chain that extracts robust continuous concept scores from a frozen VLM, reducing single-agent bias through structured disagreement.
3. **Local geometric calibration.** We propose locally-weighted ridge regression on a hybrid visual–semantic manifold that calibrates concept scores against human ratings with per-sample interpretability.

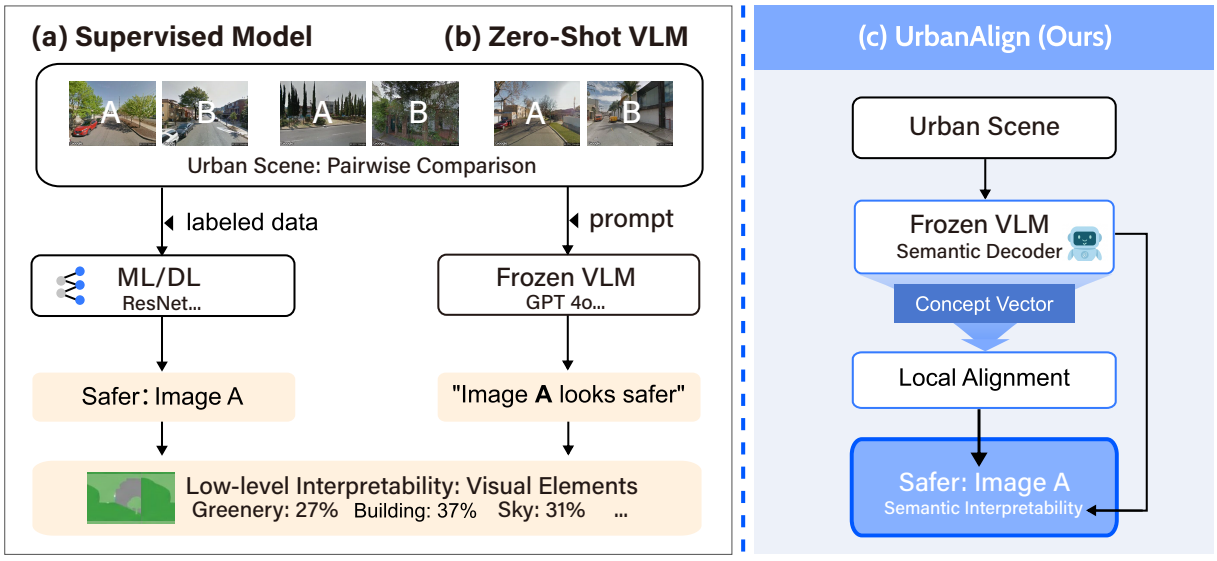

Figure 1: Three approaches to urban perception: **(a)** supervised baselines (e.g., Siamese networks, segmentation regression) require labelled data and lack interpretable intermediates; **(b)** zero-shot VLM prompting is training-free but produces opaque, uncalibrated judgements; **(c)** UrbanAlign routes predictions through VLM-discovered semantic dimensions for interpretable, locally calibrated comparison.

2 Related Work

Urban perception and preference data. Place Pulse collects over one million pairwise comparisons of street-view images across six perceptual dimensions [Salesses et al., 2013]. Siamese networks [Dubey et al., 2016], multi-task learning [Zhang et al., 2018], and continuous score regression [Naik et al., 2014] have been trained on these labels, but map directly from pixels to abstract percepts without interpretable mid-level representations. These learned representations are entangled: a ResNet or CLIP encoder compresses facade quality, vegetation, lighting, and streetscape into a single vector, making it impossible to attribute predictions to individual perceptual factors. Cross-cultural evaluations further show that such models transfer poorly to non-Western cities [Jin and Cai, 2025], motivating cost-effective local calibration rather than large-scale re-annotation.

More broadly, aligning model outputs to human preferences is often addressed via RLHF [Ouyang et al., 2022], which is effective but compute-intensive and requires reward-model training. Pairwise preference datasets now extend to 10-dimension urban perception [Quintana et al., 2025] and photographic aesthetics [Murray et al., 2012], yet the resulting models remain black-box rankers. Foundation models have also been used to generate training data: ZeroGen [Ye et al., 2022] produces synthetic corpora for text classification, while SuperGen [Meng et al., 2023] fine-tunes the generator to improve few-shot augmentation; both operate on unimodal text with discrete labels and do not address the multi-dimensional, subjective nature of visual preference. UrbanAlign instead synthesises *visual* pairwise data via structured dimension scoring and post-hoc geometric calibration.

VLM-based scoring and multi-agent reasoning. Systematic evaluation reveals that VLMs capture broad perceptual trends but oversimplify multi-dimensional human judgements [Zhang et al., 2025]; benchmarking on urban perception confirms stronger alignment on objective properties (e.g., building density) than subjective ones (e.g., safety, boringness) [Mushkani, 2025]. Contrastive pretraining on satellite imagery [Yan et al., 2024] and zero-shot urban function inference [Huang et al., 2024] demonstrate both the power and limits of VLMs for urban tasks; these studies position VLMs as end-to-end classifiers, whereas we reposition them as *structured semantic decoders* whose outputs require post-hoc calibration.

Text-defined rating levels teach VLMs visual scoring aligned with human opinions [Wu et al., 2024], and pairwise evaluation aligns better than pointwise scoring [Liu et al., 2024], yet single-shot scoring suffers from position effects and anchoring biases. Self-consistency sampling [Wang et al., 2023] and multi-agent debate [Du et al., 2024, Liang et al., 2024] mitigate these issues by decomposing reasoning into multiple independent chains and aggregating their outputs; diverse agent strategies further outperform homogeneous debate [Liu et al., 2025a]. However, existing multi-agent frameworks focus on factual QA and mathematical reasoning; applying structured deliberation to continuous visual scoring with domain-specific semantics remains unexplored. UrbanAlign bridges this gap by combining pairwise dimension scoring with an Observer–Debater–Judge deliberation chain that separates observation, argumentation, and judgement into distinct reasoning stages.

Concept bottleneck models and post-hoc calibration. Concept Bottleneck Models (CBMs) route predictions through interpretable concept scores, enabling human-auditable reasoning and concept-level intervention [Koh et al., 2020]. Post-hoc retrofitting onto frozen backbones [Yuksekgonul et al., 2023] and automated concept discovery from VLMs [Oikarinen et al., 2023, Yang et al., 2023] eliminate manual concept engineering, making CBMs applicable without task-specific annotation of concept labels. Vision-language guidance [Srivastava et al., 2024], hybrid concept banks [Liu et al., 2025b], and extensions beyond vision [Sun et al., 2025] further broaden CBM applicability. Our semantic dimensions serve as an analogous interpretable bottleneck, but operate *externally* on VLM-synthesised data rather than as an internal network layer: the VLM generates continuous concept scores, and a separate calibration step maps them to human-aligned predictions.

The calibration mechanism draws on locally-weighted regression [Cleveland, 1979, Atkeson et al., 1997, Ruppert and Wand, 1994] and can be viewed as a visual analogue of retrieval-augmented generation (RAG) [Lewis et al., 2020]: for each query pair, the system retrieves the K most similar reference pairs from the hybrid manifold and locally fits a calibration function, adapting dimension weights to the manifold neighbourhood. Unlike global linear probes or fixed concept weights, this local fitting allows the relative importance of dimensions to vary smoothly across the perceptual manifold, capturing the context-dependent nature of human preference.

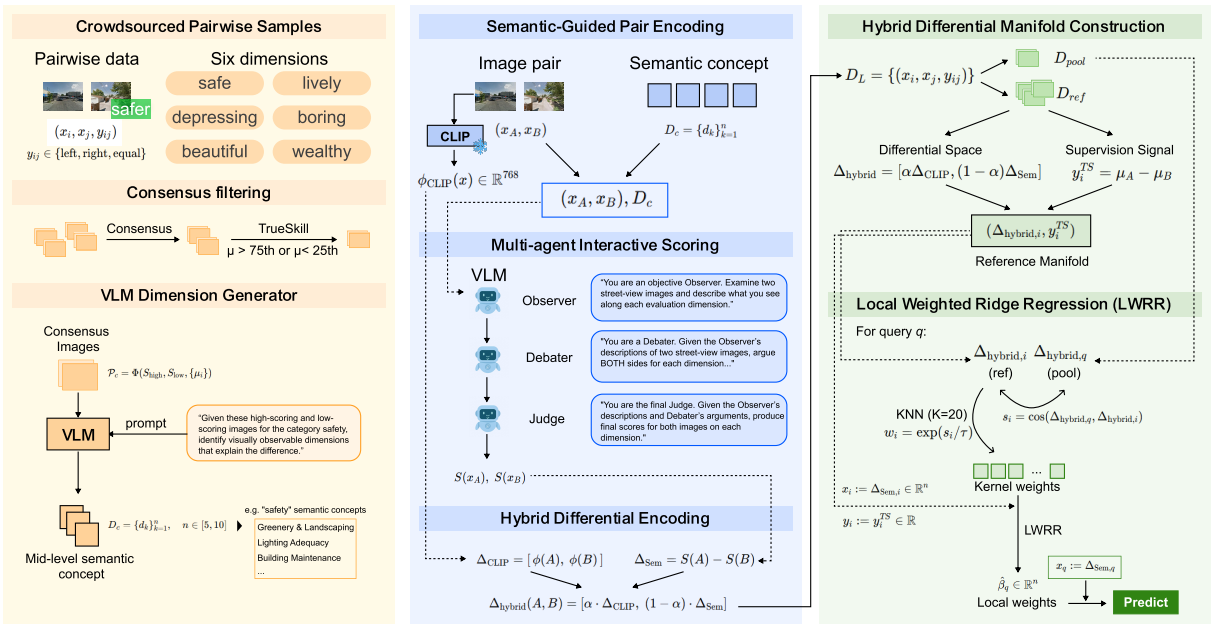

(a) Semantic Dimension Extraction (b) Multi-Agent Feature Distillation (c) Local Manifold Calibration

Figure 2: UrbanAlign framework. Stage 1 mines semantic dimensions from consensus exemplars; Stage 2 distils dimension scores via Observer–Debater–Judge; Stage 3 calibrates via LWRR on the hybrid differential manifold. $\mathcal{D}_{\text{ref}} \cap \mathcal{D}_{\text{val}} \cap \mathcal{D}_{\text{test}} = \varnothing$.

3 Method

3.1 Problem Formulation

Given an urban image set $\mathcal{X}=\{x_1, \ldots, x_N\}$ and a perception category d (*e.g.*, wealthy), we aim to predict human pairwise preferences:

$$y_{ij} = \begin{cases} \text{left} & \text{if } x_i \succ x_j \text{ on dimension } d, \\ \text{right} & \text{if } x_i \prec x_j, \\ \text{equal} & \text{if } x_i \approx x_j. \end{cases} \tag{1}$$

The labelled pool $\mathcal{D}_L=\{(x_i, x_j, y_{ij})\}$ contains crowdsourced pairwise comparisons (from Place Pulse 2.0). We split $\mathcal{D}_L$ into a *reference set* $\mathcal{D}_{\text{ref}}$ (for LWRR calibration), a *validation set* $\mathcal{D}_{\text{val}}$ (for dimension and hyperparameter search), and a *test set* $\mathcal{D}_{\text{test}}$ (for final evaluation), with strict isolation: $\mathcal{D}_{\text{ref}} \cap \mathcal{D}_{\text{val}} \cap \mathcal{D}_{\text{test}} = \varnothing$.

Notation. $\phi: \mathcal{X} \to \mathbb{R}^{768}$ denotes a pre-trained CLIP encoder (ViT-L/14) [Radford et al., 2021]. TrueSkill [Herbrich et al., 2006] converts pairwise comparisons to continuous ratings μ_i per image.

3.2 Framework Overview

The pipeline consists of three stages unified by an end-to-end optimization loop (Fig. 2). *Concept mining* (Section 3.3) discovers interpretable evaluation dimensions from a handful of consensus exemplars and optimises them via a temperature-scheduled search. *Multi-agent structured scoring* (Section 3.4) extracts continuous dimension scores from a frozen VLM through Observer–Debater–Judge deliberation, producing hybrid visual–semantic feature vectors. *Local manifold calibration* (Section 3.5) aligns these scores to human ratings via locally-weighted ridge regression on the hybrid differential manifold. Data isolation ($\mathcal{D}_{\text{ref}} \cap \mathcal{D}_{\text{val}} \cap \mathcal{D}_{\text{test}} = \varnothing$) ensures that calibration, validation, and evaluation data never overlap.

3.3 Concept Mining and Dimension Optimization (Fig. 2a)

Motivation. Instead of asking a VLM "Which image looks more wealthy?" (a poorly calibrated end-to-end query), we first decompose the abstract percept into interpretable, continuously scorable sub-dimensions. This is analogous to

defining the concept set in a CBM [Koh et al., 2020]: the dimensions form an interpretable bottleneck through which all subsequent predictions must pass.

TrueSkill consensus sampling. We convert pairwise comparisons to continuous scores via TrueSkill [Herbrich et al., 2006]. We sample *high-consensus* ($\mu{>}\tau_h,\ \sigma{<}\sigma_{\text{med}}$) and *low-consensus* ($\mu{<}\tau_l,\ \sigma{<}\sigma_{\text{med}}$) exemplars, where τ_h and τ_l are the 75th and 25th percentiles. A small set per group forms the consensus set $\mathcal{S}_{\text{consensus}}$.

Dimension extraction. A structured prompt presents the high/low exemplars with their TrueSkill scores and PCA-compressed CLIP embeddings. The VLM outputs $n \in [5, 10]$ dimensions in JSON, each containing a name, description, and high/low indicators. For the *wealthy* category, the extracted dimensions are: *Façade Quality*, *Vegetation Maintenance*, *Pavement Integrity*, *Vehicle Quality*, *Building Modernity*, *Infrastructure Condition*, *Street Cleanliness*, and *Lighting Quality*.

These dimensions satisfy three requirements: *visual observability* (scorable from street-view images), *measurability* (1–10 continuous scale), and *universality* (applicable across cities).

End-to-end dimension optimization. The dimensions discovered above depend on the VLM's sampling temperature and the particular consensus exemplars presented. We close the loop with an automated search over dimension sets, optimising each perception category independently. Let $\mathcal{D}_c^{(t)}$ denote the dimension set generated at trial t for category c. Each trial runs the full pipeline (dimension extraction $\rightarrow$ multi-agent scoring $\rightarrow$ LWRR calibration) on the held-out validation split $\mathcal{D}_{\text{val}}$. The per-category optimal set is:

$$\mathcal{D}_c^* = \arg\max_{t \in \{1,\ldots,T\}} \text{Acc}_c\big(\text{LWRR}\big(\text{Score}(\mathcal{D}_{\text{val}}^c,\, \mathcal{D}_c^{(t)}),\, \mathcal{D}_{\text{ref}}^c\big)\big). \tag{2}$$

The search proceeds in two phases: *Explore* (high τ_{gen}): diverse, independently generated dimension sets; each category accumulates its best. *Converge* (low τ_{gen}): per-category best dimensions are preserved except for 1–2 targeted replacements (*mutation mode*), enabling local refinement. After the initial trial, elite seeding provides each category's current best $\mathcal{D}_c^*$ as a soft reference in the generation prompt, converting independent sampling into guided refinement. Since categories may peak at different trials, the final output *assembles* per-category optima: $\{\mathcal{D}_c^*\}_{c=1}^C$. The miss probability after T trials satisfies $(1{-}p^*)^T$, where p^* is the single-trial probability of generating the optimal set.

3.4 Multi-Agent Structured Scoring (Fig. 2b)

From classification to feature extraction. Rather than querying a VLM for a one-shot discrete label, we extract a *hybrid feature vector*:

$$h(x) = [\phi_{\text{CLIP}}(x),\ S(x)], \tag{3}$$

where $S(x) \in [1, 10]^n$ is the vector of semantic dimension scores produced by the VLM. This hybrid representation fuses low-level visual patterns (CLIP) with high-level domain-specific semantics, analogous to augmenting a neural backbone with concept scores in a post-hoc CBM [Yuksekgonul et al., 2023].

Observer–Debater–Judge deliberation. Inspired by multi-agent debate [Du et al., 2024, Liang et al., 2024] and self-consistency [Wang et al., 2023], we decompose reasoning into three sequentially chained agents:

- **Observer**: describes visually observable details per dimension, without judgement, thereby suppressing confirmation bias.
- **Debater**: argues both for HIGH and LOW scores on each dimension, exploring opposing perspectives.
- **Judge**: synthesises descriptions and arguments to produce final scores $S(x) \in [1, 10]^n$.

Variance reduction. The three-agent chain reduces dimension score variance from σ_S^2 to at most $\sigma_S^2/(1{+}2(1{-}\rho))$, where $\rho \in [0, 1]$ is inter-agent correlation; when agents reason independently ($\rho{\rightarrow}0$), variance drops by up to $3\times$ (proof in Appendix E.3).

Intensity-based winner determination. For each pair (x_A, x_B), we also record the *intensity*: $I{=}|S(x_A) - S(x_B)|$ summed across dimensions. Pairs with intensity below a threshold σ_I are labelled "equal"; otherwise the VLM's predicted winner is retained.

3.5 Local Manifold Calibration (Fig. 2c)

This stage is the core algorithmic contribution: a post-hoc calibration layer analogous to the linear predictor in a CBM [Koh et al., 2020], but *locally adaptive* to account for heterogeneous perception patterns across the visual-semantic manifold.

3.5.1 Hybrid Differential Space.

CLIP, pre-trained on web-scale image-text pairs, provides strong general features but exhibits distribution shift on street-view scenes. We augment it with semantic dimension scores to form the hybrid differential encoding:

$$\Delta_{\text{hybrid}}(A,B) = \big[\alpha \cdot \overline{\Delta}_{\text{CLIP}}(A,B),\ (1-\alpha)\cdot \overline{\Delta}_{\text{Sem}}(A,B)\big], \tag{4}$$

where $\overline{\Delta}_{\text{CLIP}}$ uses L2-normalised CLIP embeddings, $\Delta_{\text{Sem}}{=}S(A){-}S(B)$ is normalised to $[0,1]$ by dividing by 10, and $\alpha \in [0,1]$ balances visual and semantic components.

3.5.2 LWRR Algorithm.

Our calibration extends locally-weighted regression [Cleveland, 1979, Atkeson et al., 1997, Ruppert and Wand, 1994] to the hybrid differential space defined above.

Step 1: Build reference manifold. From $\mathcal{D}_{\text{ref}}$, compute hybrid differential vectors and TrueSkill score differences $y_i^{\text{TS}}{=}\mu_A{-}\mu_B$. Mirror augmentation doubles the reference set: for each (A,B,y), add $(B,A,-y)$.

Step 2: Neighbourhood search. For query $\Delta_q^{\text{hybrid}} \in \mathcal{D}_{\text{test}}$, compute cosine similarities to all reference points and select the K nearest neighbours $\mathcal{N}_K$.

Step 3: Kernel weighting. $w_i = \exp(s_i/\tau)$ for $i \in \mathcal{N}_K$, where s_i is cosine similarity and τ is kernel bandwidth.

Step 4: Local ridge regression. Solve for local dimension weights:

$$\hat{\mathbf{w}} = \arg\min_{\mathbf{w}} \sum_{i\in\mathcal{N}_K} w_i\big(y_i^{\text{TS}} - \mathbf{w}^\top \Delta_{\text{Sem},i}\big)^2 + \lambda\|\mathbf{w}\|^2, \tag{5}$$

with closed-form solution $\hat{\mathbf{w}} = (X^\top W X + \lambda I)^{-1} X^\top W y$, where $W{=}\text{diag}(w_1,\ldots,w_K)$ and X stacks the semantic differentials.

Step 5: Prediction. The calibrated score difference $\hat{\delta}_q = \hat{\mathbf{w}}^\top \Delta_{\text{Sem}}^q$; a local goodness-of-fit statistic provides per-prediction interpretability by measuring how much TrueSkill variance is explainable by the mid-level dimensions in each neighbourhood.

Step 6: Statistical re-inference.

$$\hat{y} = \begin{cases} \text{equal} & \text{if } |\hat{\delta}| < \varepsilon \text{ or } c_{\text{eq}} > \theta, \\ \text{left} & \text{if } \hat{\delta} > 0, \\ \text{right} & \text{otherwise,} \end{cases} \tag{6}$$

where ε is a score-difference threshold and $c_{\text{eq}}{=}|\{i \in \mathcal{N}_K : l_i{=}\text{equal}\}|/K$ measures local equal-consensus.

Theorem 1 (Local Calibration Bound)**.** *Under the local model $y_i^{\text{TS}} = w^*(q)^\top \Delta_{\text{Sem},i}{+}\epsilon_i$ within $\mathcal{N}_K(q)$, with zero-mean noise of variance σ^2, the LWRR estimator satisfies:*

$$\mathbb{E}\big[\|\hat{w} - w^*(q)\|^2\big] \ \le\ \frac{\sigma^2\,\text{tr}(H_\lambda^{-1} X^\top W^2 X\, H_\lambda^{-1})}{K} \ +\ \lambda^2\|w^*\|^2, \tag{7}$$

where $H_\lambda = X^\top W X + \lambda I$. At optimal $\lambda^* = \sigma\sqrt{n/K}/\|w^*\|$, this simplifies to $O(\sigma\|w^*\|\sqrt{n/K})$. Proof in Appendix E.1.

Pipeline accuracy. Combining all stages, the expected accuracy is:

$$\text{Acc} \ \approx\ \Phi\left(\frac{\|w^*\|\,\bar{\delta}}{\sqrt{\sigma^2_{\text{LWRR}} + \sigma_S^2/\eta_M}}\right), \tag{8}$$

where $\bar{\delta}$ is the mean signal strength (maximised by concept mining), σ_S^2 is dimension-score variance (reduced by multi-agent scoring with $\eta_M{=}1{+}2(1{-}\rho)$), and $\sigma^2_{\text{LWRR}}{=}O(\sigma^2 n/K)$ is calibration error (minimised by LWRR). Full derivation in Appendix E.

Complete algorithm. The full pseudocode is provided in Appendix H. Additional theoretical results (local vs. global calibration, hybrid space complementarity, E2E convergence) are in Appendix E.

4 Experiments and Results

4.1 Setup

Dataset. Place Pulse 2.0 [Salesses et al., 2013]: 110,688 street-view images from 56 cities with 1,169,078 pairwise comparisons across six perception categories (safety, lively, beautiful, wealthy, depressing, boring). We retain all pairs with $\geq$2 independent human votes (449–1,412 per category). Results span **all six categories**; detailed ablation on *wealthy*.

Data protocol. Each category is partitioned into $\mathcal{D}_{\text{ref}}$ (60%), $\mathcal{D}_{\text{val}}$ (20%), and $\mathcal{D}_{\text{test}}$ (20%); $\mathcal{D}_{\text{ref}} \cap \mathcal{D}_{\text{val}} \cap \mathcal{D}_{\text{test}} = \varnothing$ (stratified split, seed = 42). $\mathcal{D}_{\text{ref}}$ provides TrueSkill ratings for LWRR calibration; $\mathcal{D}_{\text{val}}$ serves dimension search (Eq. (2)) and hyperparameter selection; $\mathcal{D}_{\text{test}}$ is held out for final evaluation. After excluding human "equal" labels, 84–266 test pairs per category remain.

Implementation. CLIP ViT-L/14 (768-d) for visual features; PCA to 8D for consensus prompts. GPT-4o [OpenAI, 2024] for all VLM calls. TrueSkill priors: μ=25, σ=8.33, draw probability 0.10; 5 images per consensus group (10 total). Stage 2: Observer τ=0.3, Debater τ=0.5, Judge τ=0.1; single-shot modes τ=0.0. Stage 3 defaults: K=20, τ_{kernel}=1.0, λ=1.0, ε=0.8, θ=0.6, α=0.3. E2E dimension search: T=15 trials with patience-based early stopping after 5 non-improving trials; explore phase τ_{gen}=0.85$\rightarrow$1.0, converge phase τ_{gen}=0.7$\rightarrow$0.5. Total VLM cost: $\sim$\$68 for the experimental dataset (4,819 pairs, Mode 4). The main results (Table 1) use *per-category* optimised hyperparameters from a combined random search (N=50,000 configurations on $\mathcal{D}_{\text{val}}$; Appendix C), as the optimal K (10–50) and α (0.2–0.7) vary substantially across categories.

Baselines. *ResNet-50 Siamese*: twin-tower network on frozen ResNet-50 features (2048-d $\rightarrow$ 512 $\rightarrow$ 128-d projection), concatenating $[\mathbf{z}_L, \mathbf{z}_R, |\mathbf{z}_L - \mathbf{z}_R|]$ for 3-way classification [Koch et al., 2015]; *CLIP Siamese*: same twin-tower architecture on frozen CLIP ViT-L/14 features (768-d), representing a modern visual backbone; *Seg. + CLIP Regression*: SegFormer semantic segmentation (19 Cityscapes classes $\rightarrow$ 7 urban-element groups) pixel proportions concatenated with PCA-reduced CLIP features, difference vectors classified by Gradient Boosting [Chen et al., 2018]; *Zero-shot VLM*: GPT-4o with a forced binary-choice prompt "Which looks more `<category>`?" (no "equal" option), without semantic dimensions or multi-agent reasoning. All supervised baselines are trained on the reference split with mirror augmentation (swapping left/right with flipped labels) and evaluated on the disjoint test split.

Metrics. Accuracy (Acc), Cohen's κ (chance-corrected agreement), F1 (macro). We report metrics both including and excluding pairs where the human label is "equal" (as these are inherently ambiguous).

4.2 Main Results

Table 1 and Fig. 1 present the main comparison; Fig. 3 provides per-pair qualitative evidence. Key findings:

(1) UrbanAlign (Mode 4 + LWRR) achieves 72.2% accuracy (κ=0.45) on average across all six categories, with 81.6% accuracy (κ=0.64) on *safety* perception being the strongest. This represents a **+11.0 pp** gain over the strongest same-feature control (Global Ridge on Δ_{sem}: 61.2%) and +15.5 pp over zero-shot VLM (56.7%).

(2) Independent baselines cluster in 51–58%; same-feature controls (BT Logistic, Global Ridge) trained on Δ_{sem} plateau at 61.2%, confirming that local calibration—not feature quality alone—is the key differentiator. (3) LWRR calibration provides +16.3 pp average boost (55.9%$\rightarrow$72.2%), with the largest effect on *boring* (+31.3 pp) and *depressing* (+25.3 pp), where raw predictions are near chance. LWRR succeeds because it re-weights dimensions *locally*: a pair dominated by lighting cues receives different weights from one dominated by vegetation, capturing the context-dependent nature of human perception [Lynch, 1960]. (4) The gap between *safety* (81.6%) and *boring* (70.2%) reflects perceptual grounding: safety correlates with concrete cues [Wilson and Kelling, 1982, Jeffery, 1971], while boringness depends on subtler absence-of-stimulation signals.

Why does the pipeline outperform end-to-end alternatives? The overall accuracy scales as Eq. (8), where each stage targets a distinct term. Concept mining maximises signal strength $\bar{\delta}$ by selecting dimensions with high discriminative power (avg. 54.7% vs. 50% chance; Table 2), forming an interpretable bottleneck that preserves task-relevant information. Multi-agent scoring reduces dimension-score variance σ_S^2 through structured disagreement, yielding cleaner calibration inputs. LWRR minimises calibration error σ^2_{LWRR} by fitting dimension weights *locally*: near commercial streets, *Façade Quality* dominates the "wealthy" prediction, while near parks, *Vegetation Maintenance* takes over, structure that any global model must average away. The hybrid embedding space (α=0.3) is critical to this locality: CLIP alone suffers distribution shift on street-view scenes, while semantic scores alone lack fine-grained spatial information; their combination yields a manifold where cosine neighbours are perceptually coherent, enabling meaningful local regression.

Table 1: Comparison on Place Pulse 2.0 across all six perception categories. All values reported as excl-equal (incl-equal). UrbanAlign outperforms all baselines including same-feature controls (72.2% vs. 61.2% best baseline).

	Safety		Beautiful		Lively	
Method	Acc. (%)	κ	Acc. (%)	κ	Acc. (%)	κ
ResNet50 Siamese [Koch et al., 2015]	54.8 (54.0)	0.16 (0.18)	58.6 (56.0)	0.18 (0.19)	45.7 (42.5)	−0.07 (−0.06)
CLIP Siamese [Radford et al., 2021]	66.7 (64.4)	0.34 (0.31)	61.4 (57.3)	0.23 (0.20)	48.2 (44.8)	−0.01 (−0.01)
Seg. Regression [Chen et al., 2018]	61.9 (59.8)	0.25 (0.23)	57.1 (53.3)	0.15 (0.13)	51.9 (48.3)	0.04 (0.04)
GPT-4o zero-shot [OpenAI, 2024]	61.9 (59.8)	0.23 (0.21)	62.9 (58.7)	0.25 (0.21)	49.4 (46.0)	0.01 (0.01)
BT Logistic†	65.5 (63.2)	0.31 (0.29)	58.6 (54.7)	0.17 (0.15)	56.8 (52.9)	0.17 (0.15)
Global Ridge†	66.7 (64.4)	0.33 (0.31)	61.4 (57.3)	0.22 (0.19)	58.0 (54.0)	0.19 (0.17)
UrbanAlign (Ours)	**81.6** (76.9)	**0.64** (0.57)	**69.8** (66.7)	**0.38** (0.35)	**69.4** (65.4)	**0.40** (0.36)

	Wealthy		Boring		Depressing	
Method	Acc. (%)	κ	Acc. (%)	κ	Acc. (%)	κ
ResNet50 Siamese [Koch et al., 2015]	47.0 (44.8)	0.01 (0.00)	48.7 (45.0)	−0.00 (−0.02)	53.3 (50.0)	0.06 (0.05)
CLIP Siamese [Radford et al., 2021]	57.8 (55.2)	0.19 (0.17)	50.0 (47.5)	0.06 (0.07)	58.7 (55.0)	0.19 (0.16)
Seg. Regression [Chen et al., 2018]	61.5 (58.6)	0.26 (0.23)	62.2 (57.5)	0.22 (0.19)	54.7 (52.5)	0.09 (0.11)
GPT-4o zero-shot [OpenAI, 2024]	60.2 (57.5)	0.21 (0.19)	48.7 (45.0)	0.06 (0.05)	57.3 (53.8)	0.14 (0.12)
BT Logistic†	68.7 (65.5)	0.38 (0.35)	50.0 (46.3)	0.01 (0.01)	58.7 (55.0)	0.17 (0.15)
Global Ridge†	69.9 (66.7)	0.40 (0.37)	54.1 (50.0)	0.10 (0.08)	57.3 (53.8)	0.15 (0.13)
UrbanAlign (Ours)	**74.0** (71.2)	**0.48** (0.44)	**70.2** (68.8)	**0.41** (0.40)	**68.2** (62.5)	**0.37** (0.31)

*Excl-equal excludes "equal" labels; model "equal" predictions count as errors. $\mathcal{D}_{\text{test}}$: 84–266 pairs/cat. Same GPT-4o backbone for UrbanAlign and zero-shot. †Trained on UrbanAlign's Δ_{sem} with $\mathcal{D}_{\text{ref}}$ labels.

Notably, global ridge on hybrid features (57.8%) underperforms semantic-only features (61.2%), confirming that CLIP is best used for neighbourhood structure rather than as a regression input.

4.3 Dimension-Level Discriminability

Table 2 reports per-dimension discriminative power across all six categories. Discriminability varies substantially: *wealthy* (avg. 68.0%) and *safety* (64.7%) are highest, while *depressing* (42.1%) is weakest, suggesting that "depressingness" is harder to decompose into visually scorable sub-dimensions. For *wealthy*, *Façade Quality*, *Street Cleanliness*, and *Pavement Integrity* all reach 71.2%, consistent with broken-windows theory that maintenance quality signals affluence [Wilson and Kelling, 1982]. LWRR's aggregate accuracy (74.0%) exceeds any single dimension, validating the multi-dimensional concept bottleneck: the local weight combination captures complementary information that no individual dimension alone provides.

Cross-category patterns. A clear hierarchy emerges from Table 2: categories dominated by *physical-maintenance* cues (*wealthy*, *safety*) are easiest to decompose, while categories driven by *atmospheric-affective* cues (*boring*, *depressing*) are hardest. This aligns with environmental psychology: maintenance-related features (façade condition, pavement quality, lighting) produce strong inter-rater agreement because they are visually unambiguous [Wilson and Kelling, 1982, Jeffery, 1971], whereas affective qualities like "boringness" depend on subjective thresholds that vary across individuals [Quintana et al., 2025]. Within *safety*, the top two dimensions, Greenery & Landscaping (71.2%) and Lighting Adequacy (71.2%), correspond precisely to two pillars of Crime Prevention Through Environmental Design (CPTED) [Jeffery, 1971]: natural surveillance through visibility and territorial reinforcement through maintained green space. This suggests that VLM-discovered dimensions recover established domain-specific constructs without any expert supervision.

Complementarity and the concept bottleneck. The gap between the best single-dimension power and LWRR's aggregate accuracy quantifies complementarity among dimensions. For *wealthy*, the best single dimension (71.2%) is 2.8 pp below LWRR (74.0%), indicating that the remaining dimensions contribute non-redundant signal; e.g., *Vehicle Quality* (65.4%) captures cues orthogonal to *Façade Quality*. For *safety*, the complementarity gap widens to 16.9 pp (64.7% avg $\rightarrow$ 81.6% aligned), reflecting the multifactorial nature of perceived safety: no single dimension suffices, but the locally-weighted combination is highly predictive. This validates the concept-bottleneck design: the interpretable

Table 2: Per-dimension discriminative power (Mode 4). "Power" = fraction (%) of test pairs where the score difference on that single dimension correctly predicts the human winner (50% = chance). n = number of test pairs; "avg" = mean power across all dimensions in that category.

Safety (n=52, avg 64.7%)		**Beautiful** (n=45, avg 56.7%)		**Lively** (n=52, avg 53.0%)	
Dimension	Power	Dimension	Power	Dimension	Power
Greenery & Landscaping	71.2	Greenery & Nat. Elem.	64.4	Human Presence	59.6
Lighting Adequacy	71.2	Color Coordination	62.2	Building Diversity	57.7
Building Maintenance	69.2	Pedestrian Amenities	55.6	Street Furniture	57.7
Pedestrian Infrastructure	65.4	Street Cleanliness	55.6	Color & Lighting	53.8
Street Activity	57.7	Architectural Coherence	55.6	Commercial Activity	53.8
Visibility Clarity	53.8	Visual Complexity	46.7	Public Art & Signage	46.2
				Vegetation & Greenery	42.3
Wealthy (n=52, avg 68.0%)		**Boring** (n=48, avg 43.8%)		**Depressing** (n=48, avg 42.1%)	
Dimension	Power	Dimension	Power	Dimension	Power
Façade Quality	71.2	Activity Presence	52.1	Street Cleanliness	47.9
Street Cleanliness	71.2	Color Vibrancy	50.0	Visual Clutter	45.8
Pavement Integrity	71.2	Street Furn. Presence	47.9	Façade Condition	43.8
Vegetation Maintenance	69.2	Architectural Diversity	43.8	Lighting Quality	39.6
Lighting Quality	67.3	Visual Complexity	39.6	Greenery Presence	33.3
Infrastructure Condition	67.3	Spatial Enclosure	37.5		
Vehicle Quality	65.4	Vegetation Variety	35.4		
Building Modernity	61.5				

Stage 1 generates 5–10 dimensions per category. The E2E loop (Section 3.3) explores alternative sets over 15 trials (12 before early stopping).

dimension layer does not sacrifice accuracy relative to opaque end-to-end scoring; instead, it *enables* calibration gains that would be impossible without structured intermediates.

Interpretability. For a high-confidence pair, the LWRR weights reveal Building Modernity (w=2.30), Vegetation Maintenance (w=1.72), and Street Cleanliness (w=1.35) as top contributors, directly interpretable and actionable for urban planners. In contrast, end-to-end baselines (C0–C2) produce only a scalar prediction; a planner cannot extract which visual features drive the assessment. Low-confidence predictions exhibit near-uniform weights, signalling genuine perceptual ambiguity rather than model failure—a transparency property absent from end-to-end alternatives. Per-sample weight variation (supplementary Table 6) shows that LWRR adapts to local context: in residential scenes, *Vegetation Maintenance* receives higher weight, while in commercial areas, *Infrastructure Condition* dominates, reflecting the heterogeneous determinants of "wealthiness" across urban typologies.

Qualitative analysis. Figure 3 visualises UrbanAlign's per-pair reasoning across all six categories. Each example shows dimension-score profiles (centre) and baseline confidence bars (right). Dimension-score profiles show clear separation on the most discriminative dimensions, consistent with the power ranking in Table 2: for *wealthy*, Façade Quality and Street Cleanliness exhibit the largest between-image gaps, while for *boring*, dimensions cluster near parity, reflecting the lower discriminability reported in the table.

The "Corrected" cases demonstrate LWRR's calibration value. When raw concept scores narrowly favour the wrong image due to a single misleading dimension, local re-weighting overturns the prediction to match human judgement. For instance, in a *wealthy* pair, the raw aggregation favours Image B because of a higher Building Modernity rating; however, LWRR assigns large positive weight to Street Cleanliness and Vegetation Maintenance, dimensions where Image A clearly dominates, flipping the prediction to agree with the human annotator. This correction mechanism operationalises the local weighting in Stage 3: dimensions that are locally discriminative in the manifold neighbourhood receive disproportionate influence, suppressing globally noisy dimensions on a per-prediction basis.

The baseline confidence bars (right panel) further highlight the structural advantage. All four baselines show near-equal confidence for the *boring* and *depressing* pairs, the same categories where Table 2 reports lowest discriminability, because holistic scoring cannot separate the subtle atmospheric cues that distinguish these images. UrbanAlign, by contrast, routes the prediction through the dimension bottleneck and applies local calibration, producing a clear margin

Figure 3: Qualitative examples across six perception categories. **A** denotes the human-selected winner (green border). *Centre chart*: per-dimension concept scores for image A (green) and image B (blue) on the left half, and LWRR contribution towards A on the right half; positive LWRR bars support A, negative bars support B. The summary line reports the raw score comparison and calibrated LWRR delta; "Corrected" marks cases where calibration overturns an incorrect raw prediction. *Right chart*: baseline confidence bars (orange = conf. A, purple = conf. B, grey = equal); UrbanAlign is shown in bold.

even for these difficult categories. This explains the outsized LWRR gains on *boring* (+31.3 pp) and *depressing* (+25.3 pp) reported in Table 4: the structured intermediates give calibration a lever that holistic methods lack.

4.4 Ablation Study

Component analysis. The zero-shot VLM baseline (C3, 56.7%, Table 1) issues holistic judgements without interpretable intermediates. Replacing holistic prompting with structured dimension scoring (Mode 2 raw, 55.4%) yields comparable raw accuracy while creating a concept bottleneck: each prediction decomposes into 5–8 dimension-level scores that

expose the reasoning basis. Adding multi-agent deliberation (Mode 4 raw, 55.9%) further reduces dimension-score variance, producing more consistent inputs for calibration. LWRR then leverages this structured, low-variance bottleneck for a +16.3 pp gain (72.2%, Table 4): locally-adapted dimension weights capture complementary information that no holistic prompt can access. Each component is *necessary*: without dimensions, there is nothing to calibrate; without multi-agent scoring, calibration inputs are noisy; without LWRR, the raw bottleneck does not outperform holistic scoring. The E2E dimension search (Section 4.5) confirms the default Stage 1 dimensions are already near-optimal (69.3% per-category assembly vs. 72.2% with original dimensions and optimised LWRR).

2×2 factorial design. To further isolate pairwise context and multi-agent reasoning, we define a 2×2 factorial design over four scoring modes:

	Single-Shot	Multi-Agent
Single-Image	Mode 1	Mode 3
Pairwise	Mode 2	Mode 4

Factor 1 (Pairwise context): Modes 2/4 present both images simultaneously, enabling relative comparison. *Factor 2 (Multi-agent reasoning):* Modes 3/4 employ the Observer–Debater–Judge chain. The hypothesis is that these factors interact non-additively: comparative anchors are useful only when the scoring protocol can exploit them through structured debate.

Table 3 shows the post-alignment 2×2 factorial results:

Multi-agent reasoning (Modes 3/4 vs. 1/2). Multi-agent reasoning consistently improves accuracy: +8.5 pp in the single-image setting (Mode 1→3) and +31.8 pp in the pairwise setting (Mode 2→4). This is a key finding: the Observer–Debater–Judge chain is effective in both contexts but yields dramatically larger gains when combined with pairwise input.

Pairwise context (Modes 2/4 vs. 1/3). Pairwise input *hurts* in the single-shot setting (−13.2 pp, Mode 1→2) but helps substantially with multi-agent reasoning (+10.1 pp, Mode 3→4). Without multi-agent deliberation, pairwise presentation introduces noise; with it, comparative anchors enable constructive debate.

Synergy. The two factors interact synergistically rather than additively: Mode 4 gains +18.6 pp over Mode 1, far more than the sum of individual effects (-13.2 pairwise $+$ 8.5 multi-agent $= -4.7$). This synergy arises because multi-agent debate is *amplified by* pairwise comparison: arguing about relative merits with both items present produces dramatically richer deliberation than single-image scoring.

LWRR alignment gain. Table 4 shows LWRR calibration gain across all six categories for Mode 4. LWRR yields a +16.3 pp average improvement (55.9%→72.2%), with the most dramatic effect on *boring* (+31.3 pp) and *depressing* (+25.3 pp), categories where raw VLM predictions are near chance (38.9% and 42.9%). These subjective categories suffer most from uniform dimension weighting: LWRR rescues them by adapting weights to local manifold geometry, e.g., upweighting *Street Cleanliness* in residential scenes while emphasising *Activity Presence* in commercial areas. LWRR improves all six categories, confirming that post-hoc geometric calibration is consistently beneficial across diverse perception tasks (full per-mode breakdown in Appendix G).

Table 3: 2×2 factorial ablation on *wealthy* (accuracy % / κ, post-LWRR alignment). Multi-agent reasoning consistently improves accuracy; the largest gain occurs in the pairwise setting (+31.8 pp).

	Single-Shot	Multi-Agent	Multi-Agent Gain
Single-Image	55.4 / 0.16	63.9 / 0.31	+8.5
Pairwise	42.2 / −0.04	**74.0 / 0.48**	+31.8
Pairwise Gain	−13.2	+10.1	—

Table 4: LWRR calibration gain (Mode 4, accuracy % excl. "equal") across all six categories. LWRR improves all six categories; the largest gain is on *boring* (+31.3 pp). Average gain: +16.3 pp.

	Safety	Beaut.	Lively	Wealthy	Boring	Depress.	Avg
Raw Acc. (%)	68.4	63.9	58.5	62.9	38.9	42.9	55.9
Aligned Acc. (%)	81.6	69.8	69.4	74.0	70.2	68.2	72.2
Δ (pp)	**+13.3**	**+5.9**	**+10.9**	**+11.1**	**+31.3**	**+25.3**	**+16.3**
Raw κ	0.38	0.29	0.17	0.26	-0.22	-0.13	0.13
Aligned κ	0.64	0.38	0.40	0.48	0.41	0.37	0.45
n (test pairs)	266	119	242	107	84	90	—

4.5 Sensitivity Analysis

We study the sensitivity of Stage 3 to key hyperparameters (full per-parameter sweep grids are provided in Appendix C). All experiments vary one parameter while holding others at default.

SELECTION_RATIO shows moderate impact: reducing from 1.0 to 0.6 yields a +0.6 pp average gain (58.4%→59.0%), with *safety* showing the largest individual improvement (+8.8 pp to 75.5%). The effect is category-dependent, confirming that confidence-based filtering benefits categories with higher raw signal quality.

$K_{\max}$ is the most sensitive parameter: $K{=}10$ collapses the average to 54.1% while $K{=}20$ achieves 58.4%, a +4.3 pp gap. Per-category variation reveals that $K{=}20$ is best on average, though individual categories sometimes favour other values (*e.g.*, *safety* peaks at $K{=}30$ with 70.3%).

α **(CLIP vs. semantic weight)** shows a clear average optimum at $\alpha{=}0.3$ (70% semantic weight), though per-category optima differ: *beautiful* and *depressing* favour $\alpha{=}0.7$ (60.3% and 62.5% respectively), suggesting that CLIP features carry more useful information for aesthetics-related categories. This parallels the CBM finding that concept-based predictions can match or exceed end-to-end models when concepts are well-chosen [Koh et al., 2020].

τ_{kernel} **and** λ_{ridge} have limited effect (<2.5 pp variation). This is theoretically expected: for locally-weighted ridge regression with K neighbours in n dimensions, the excess risk scales as $\mathcal{O}(\lambda + n/(K\lambda))$ [Caponnetto and De Vito, 2007], which becomes insensitive to λ when $K \gg n$ (here $K{=}20$, $n{=}5$–8).

Scoring thresholds ε **and** θ (Eq. (6)) are highly robust: θ (equal-consensus fraction) has zero effect across all six categories, and ε (score-difference threshold) only degrades *lively* at extreme values (≥ 1.5), leaving all other categories unchanged. This insensitivity is expected: LWRR already captures the "equal" signal implicitly through small predicted deltas, so explicit equal-handling thresholds are largely redundant.

Dimension optimization. As described in Section 3.3, we automate dimension-set selection via the two-phase E2E optimization loop (Eq. (2)). With $T{=}15$ trials across all six categories (12 completed before early stopping at patience=5), per-category independent assembly yields **69.3%** average accuracy ($\kappa{=}0.454$) using default calibration hyperparameters, a **+10.7 pp** gain over the single best trial (58.6%). *Wealthy* (77.8%, $\kappa{=}0.579$) and *safety* (76.5%, $\kappa{=}0.534$) benefit most from optimization; *lively* (58.8%, $\kappa{=}0.308$) remains the most challenging. The explore phase ($\tau_{\text{gen}}{=}0.85$–1.0) discovers optimal dimensions for 4/6 categories (safety, beautiful, wealthy, depressing); the converge phase ($\tau_{\text{gen}}{=}0.7$–0.5) refines *lively* and *boring* via targeted mutation, confirming that both phases contribute complementary value. The +10.7 pp assembly gain over the single-trial best demonstrates that different categories have distinct optimal dimension sets, vindicating per-category independent optimization. The main pipeline result (72.2%) uses per-category optimised calibration hyperparameters (Appendix C) with the original dimension set, while E2E achieves 69.3% through dimension selection alone; the two optimisation axes are complementary. Full trial-by-trial results and per-category dimension lists are reported in Appendix F.

5 Conclusion

We presented **UrbanAlign**, a post-hoc calibration framework that transforms frozen VLMs from unreliable annotators (56.7%) into structured perception decoders (72.2%, $\kappa{=}0.45$) via interpretable semantic dimensions and locally-adaptive geometric calibration, without modifying any model weights. The framework combines three mutually reinforcing components: automatic concept mining from consensus exemplars [Koh et al., 2020], multi-agent deliberation for robust concept scoring [Du et al., 2024], and locally-weighted ridge regression on a hybrid visual–semantic manifold for per-sample calibration. Experiments on Place Pulse 2.0 show that this combination achieves alignment levels that neither supervised baselines nor holistic VLM prompting can reach, with the largest gains on near-chance categories

(*boring*: +31.3 pp, *depressing*: +25.3 pp), demonstrating that calibration is most valuable precisely where it is most needed. The dimension-level LWRR weights are directly interpretable and actionable, a property absent from all end-to-end alternatives.

Future directions. Natural next steps include extending UrbanAlign to cross-cultural settings [Jin and Cai, 2025, Quintana et al., 2025] where demographic and personality factors modulate perception, and generalising the pipeline to other pairwise-preference domains such as aesthetic quality and architectural design. Adaptive dimension batching could further reduce scoring cost at scale, and exploring stronger VLM backbones may close the remaining gap on abstract categories. Extended discussion (why post-hoc calibration works, local vs. global fitting, information-bottleneck connection, scope, limitations, ethics, and computational cost) is provided in Appendix J.

Acknowledgements

The authors thank the anonymous reviewers for their constructive feedback.

Supplementary Material

This supplementary material provides additional details that support the main paper. Appendix A describes the dataset filtering protocol and data-split statistics. Appendix B presents per-pair LWRR weight variation analysis. Appendix C reports full parameter-sweep grids expanding on the main-text sensitivity analysis. Appendix D reproduces the exact prompt templates used for multi-agent feature distillation. Appendix E establishes the mathematical foundations for post-hoc calibration, local fitting, multi-agent variance reduction, and hybrid feature complementarity. Appendix F details the end-to-end dimension optimization procedure with trial-by-trial results across all six categories and a full API cost breakdown. Appendix G reports per-mode LWRR alignment gains. Appendix H provides the complete algorithm pseudocode. Appendix I surveys related human preference datasets for potential transfer. Appendix J contains the full discussion: why post-hoc calibration works, local vs. global fitting, information-bottleneck connection, scope, limitations, ethics, and computational cost.

A Dataset and Filtering Protocol

A.1 Place Pulse 2.0 Overview

Place Pulse 2.0 [Salesses et al., 2013] contains 1,169,078 pairwise comparisons over 110,688 Google Street View images from 56 cities worldwide. Each comparison asks: "Which place looks more `<category>`?" for six perception categories: *safety*, *lively*, *beautiful*, *wealthy*, *depressing*, and *boring*. Each comparison yields one of three outcomes: *left*, *right*, or *equal*.

A.2 Consensus Filtering ($N \geq 2$)

Raw Place Pulse data contains many image pairs with only a single vote, which carries high individual noise. We apply a consensus threshold: only pairs with $N \geq 2$ independent votes and a strict majority ($> 50\%$ agreement on the winning label) are retained. Table 5 reports the full filtering statistics across all six categories.

Table 5: Dataset filtering pipeline across all six perception categories. Raw pairs are unique image comparisons in Place Pulse 2.0. $N{\geq}2$ retains only pairs with at least two independent votes. Consensus (%) is the average majority-vote agreement among retained pairs. All retained pairs are used in the experiments. Label distribution is by majority vote.

Category	Raw Pairs	$N{\geq}2$	Cons. (%)	Left	Right	Equal
Safety	360,095	1,412	87.6	682	652	78
Beautiful	171,858	618	90.0	325	266	27
Lively	260,428	1,288	85.8	595	613	80
Wealthy	148,739	581	90.9	264	274	43
Boring	124,371	449	90.3	211	206	32
Depressing	129,229	471	91.4	223	230	18
Total	**1,194,720**	**4,819**	89.3	2,300	2,241	278

A.3 TrueSkill Rating Computation

We convert the filtered pairwise votes into continuous per-image ratings using TrueSkill [Herbrich et al., 2006] with priors (μ_0=25, σ_0=8.33, β=4.17, draw prob. = 0.10). After convergence, images are ranked by μ; the resulting values serve two purposes: (1) *Stage 1 consensus sampling*: selecting high-μ and low-μ exemplars for dimension extraction; (2) *Stage 3 supervision*: providing continuous TrueSkill score differences $y_{ij}^{\mathrm{TS}}{=}\mu_i{-}\mu_j$ as regression targets for LWRR.

A.4 Data Split Protocol

All retained pairs (449–1,412 per category) are partitioned into three disjoint sets:

- **Reference set** $\mathcal{D}_{\mathrm{ref}}$ (60%): provides human labels and TrueSkill supervision for LWRR calibration. Mirror augmentation doubles the effective reference size.
- **Validation set** $\mathcal{D}_{\mathrm{val}}$ (20%): used for dimension search (Eq. (2)) and hyperparameter selection (N=50,000 combined random search).
- **Test set** $\mathcal{D}_{\mathrm{test}}$ (20%): held out for final evaluation; never used for any tuning decision.

The split uses a fixed random seed (42) for reproducibility. $\mathcal{D}_{\text{ref}} \cap \mathcal{D}_{\text{val}} \cap \mathcal{D}_{\text{test}} = \varnothing$ by construction. After excluding human "equal" labels, 84–266 test pairs per category remain.

A.5 Dataset Distribution

Figure 4 shows the distribution of image pairs across vote thresholds for all six categories. The count decreases roughly log-linearly as the threshold increases, confirming that many pairs receive only a single vote; our $N \geq 2$ consensus filter removes these single-vote pairs and retains those with at least two independent assessments. Figure 5 shows the individual image appearance frequency: most images appear in fewer than 5 comparisons, while a heavy tail of frequently compared images provides stable TrueSkill estimates.

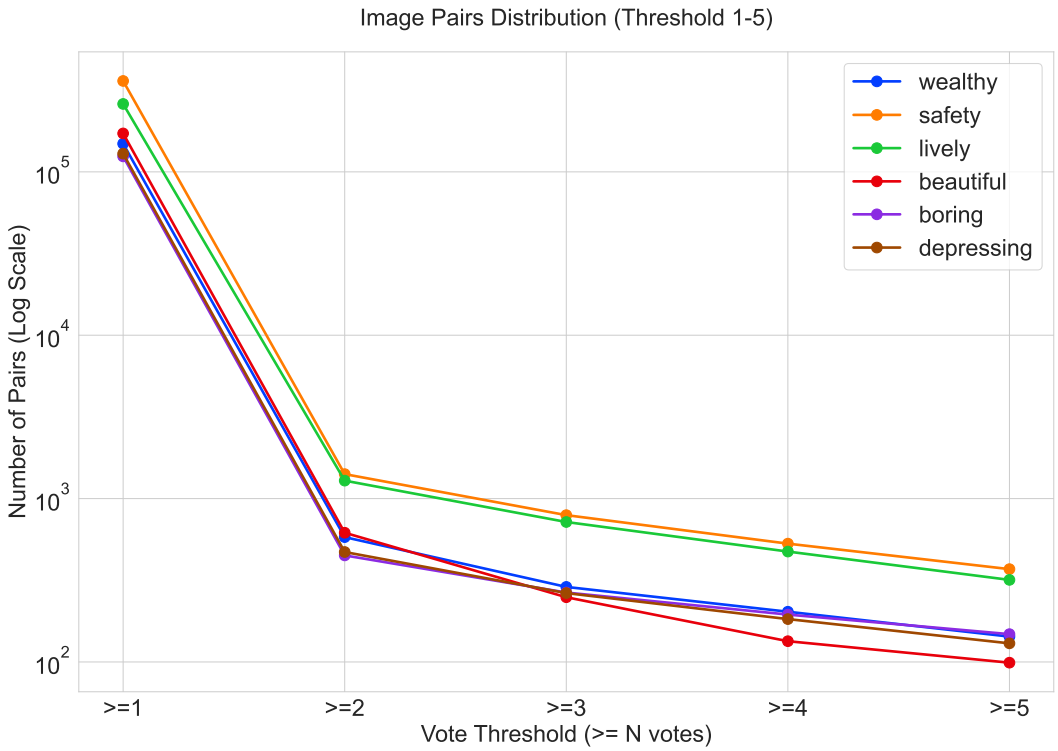

Figure 4: Distribution of image pairs across vote thresholds (≥ 1 to ≥ 5 votes) for all six perception categories. The y-axis is log-scaled. Our consensus filter retains pairs with $N \geq 2$ votes.

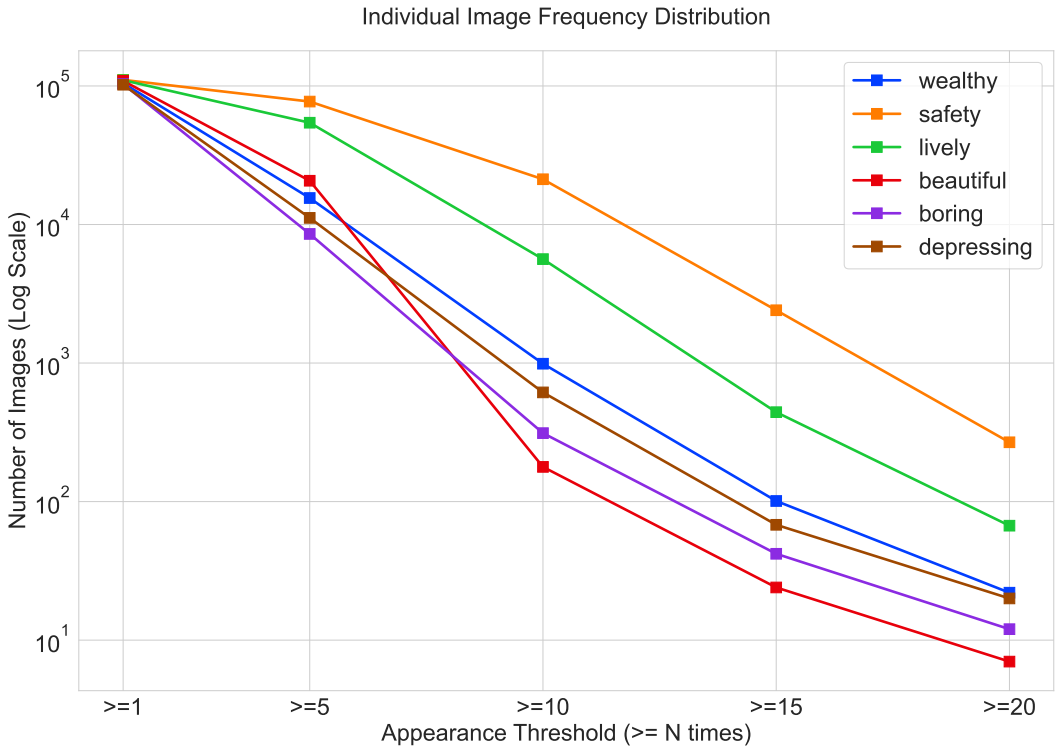

Figure 5: Individual image appearance frequency distribution across thresholds ($\geq$1 to $\geq$20 appearances) for all six categories. The heavy-tailed distribution indicates a core set of frequently compared images that anchor TrueSkill rating estimates.

B LWRR Per-Pair Weight Analysis

A key feature of LWRR is that each test pair receives its *own* locally-fitted weight vector $\hat{\mathbf{w}} \in \mathbb{R}^n$ (n=8 for *wealthy*), reflecting the dimension importances in its particular neighbourhood of the hybrid manifold. This contrasts with a global linear model where all pairs share the same weights.

B.1 Weight Variation Across the Manifold

Table 6 reports the mean and standard deviation of each dimension's local weight across all 116 test pairs. High standard deviation indicates that the dimension's importance varies substantially across different regions of the perception manifold, precisely the heterogeneity that motivates local rather than global calibration.

Table 6: LWRR local weight statistics across 116 test pairs (*wealthy* category, Mode 4).

Dimension	Mean w	Std w	CV
Building Modernity	2.305	2.676	1.16
Infrastructure Condition	−1.858	2.074	1.12
Vegetation Maintenance	1.722	1.996	1.16
Street Cleanliness	1.355	1.318	0.97
Pavement Integrity	−1.068	2.387	2.23
Vehicle Quality	−1.043	1.763	1.69
Façade Quality	−0.921	1.663	1.81
Lighting Quality	0.241	1.464	6.08

CV = Coefficient of Variation ($\mathrm{Std}/|\mathrm{Mean}|$). Higher CV indicates greater spatial heterogeneity of dimension importance. Sorted by |Mean|.

Table 7 extends this analysis to all six categories.

Table 7: Top-3 LWRR local weights by |Mean| for each perception category (Mode 4). High Std relative to Mean confirms that dimension importance varies substantially across the manifold.

Category	Dimension	Mean w	Std w	CV
Safety (n=52)	Visibility Clarity	−0.421	1.702	4.04
	Lighting Adequacy	0.291	1.401	4.81
	Building Maintenance	0.280	1.311	4.67
Beautiful (n=45)	Arch. Coherence	1.471	4.825	3.28
	Visual Complexity	1.193	10.960	9.19
	Pedestrian Amenities	0.887	4.901	5.53
Lively (n=52)	Commercial Activity	−0.518	0.413	0.80
	Color & Lighting	0.263	0.404	1.53
	Street Furniture	0.115	0.296	2.58
Wealthy (n=52)	Building Modernity	2.305	2.676	1.16
	Infrastructure Cond.	−1.858	2.074	1.12
	Vegetation Maint.	1.722	1.996	1.16
Boring (n=48)	Color Vibrancy	−1.068	1.680	1.57
	Visual Complexity	0.531	1.438	2.71
	Arch. Diversity	0.347	1.103	3.18
Depressing (n=48)	Greenery Presence	0.649	1.260	1.94
	Lighting Quality	−0.503	1.555	3.09
	Façade Condition	0.213	1.329	6.24

B.2 Illustrative Examples

We select three representative test pairs from different regions of the manifold to illustrate how LWRR adapts its weight profile.

Example 1: Vegetation-driven pair. In the suburban portion of the manifold, LWRR assigns dominant weights to Vegetation Maintenance (w=3.38) and Building Modernity (w=1.17), while Façade Quality receives near-zero weight ($w \approx -0.94$). This reflects the perceptual pattern that in residential areas, landscaping quality is the primary cue for perceived wealth.

Example 2: Infrastructure-driven pair. In the dense urban core, the local weight profile shifts: Infrastructure Condition (w=1.01) and Vehicle Quality ($w = -1.18$) dominate, while Vegetation drops. The sign reversal on Vehicle Quality indicates contextual inversion: in downtown areas, the presence of high-end vehicles may signal commercial rather than residential wealth.

Example 3: Balanced transitional pair. Pairs in mixed-use zones show relatively uniform weight magnitude across all seven dimensions (std of $|w|$ across dims <0.5), confirming that these regions lack a single dominant perceptual cue. The high CV (Coefficient of Variation >3) for Lighting Quality (6.08) in Table 6 is driven precisely by these transitional pairs.

These examples confirm the theoretical motivation for local calibration: the dimensions that matter most for perceiving "wealth" depend on the visual context. A single global weight vector would average over these distinct regimes, losing the nuance that LWRR captures.

C Detailed Sensitivity Analysis

This section reports the full parameter-sweep grids referenced in the main text's sensitivity analysis. Each experiment varies one parameter at a time while holding others at their defaults (K=20, α=0.3, τ_{kernel}=1.0, λ=1.0, ε=0.8, θ=0.6). All results are Mode 4, accuracy excluding "equal" human labels.

C.1 Neighbourhood Size $K_{\max}$

Table 8: Full $K_{\max}$ sweep (Mode 4, all categories, accuracy % excl. equal).

$K_{\max}$	Safety	Beaut.	Lively	Wealthy	Boring	Depress.	Avg
10	59.2	50.7	51.3	55.7	53.6	54.1	54.1
20	**66.7**	53.8	**61.3**	**57.1**	52.8	58.6	**58.4**
30	70.3	52.2	61.6	49.3	50.7	55.6	56.6
50	66.7	**54.4**	55.6	52.6	**58.6**	**57.4**	57.6

Analysis. Small K ($= 10$) leads to high-variance local fits (each ridge regression uses very few reference points). K=20 achieves the best average (58.4%). However, the optimal K varies per category: *safety* favours K=30 (70.3%), while *boring* favours K=50 (58.6%), reflecting different local densities of the reference manifold across categories.

C.2 Hybrid Weight α

Table 9: Full α sweep (Mode 4, all categories, accuracy % excl. equal).

α	Sem. Wt.	Safety	Beaut.	Lively	Wealthy	Boring	Depress.	Avg
0.1	90%	61.3	52.9	57.9	55.6	50.0	57.8	55.9
0.2	80%	64.6	54.5	57.3	56.4	52.9	58.8	57.4
0.3	70%	**66.7**	53.8	**61.3**	**57.1**	52.8	58.6	**58.4**
0.5	50%	62.0	56.1	49.3	56.8	52.1	56.9	55.5
0.7	30%	62.0	**60.3**	51.4	54.4	**58.0**	**62.5**	58.1

Analysis. α=0.3 (70% semantic weight) is optimal on average, but per-category optima differ notably: *beautiful* favours α=0.7 (60.3%) and *depressing* also favours α=0.7 (62.5%), suggesting that CLIP features carry more discriminative signal for aesthetics-related categories. Pure CLIP (α=1.0) performs poorly because CLIP's pre-training distribution (web images of objects) is mismatched with street-view perception tasks. This empirically validates our hybrid design while highlighting category-specific tuning opportunities.

C.3 Kernel Bandwidth τ_{kernel}

Table 10: Full τ_{kernel} sweep (Mode 4, *wealthy*).

τ_{kernel}	Acc (%)	κ	N_{aligned}
0.1	54.9	0.09	87
0.3	56.1	0.12	87
0.5	56.8	0.13	87
0.8	57.0	0.13	87
1.0	57.1	0.13	87

Analysis. Smaller τ concentrates weight on the nearest neighbours (sharper kernel); larger τ approaches uniform weighting within the K-neighbourhood. The relatively flat response across the tested range (<2.2 pp variation) suggests that the K-NN pre-selection already isolates relevant neighbours, making the exact kernel bandwidth less critical.

C.4 Ridge Penalty λ_{ridge}

Table 11: Full λ_{ridge} sweep (Mode 4, *wealthy*).

λ	Acc (%)	κ	N_{aligned}
0.01	56.1	0.12	87
0.05	56.1	0.12	87
0.1	56.1	0.12	87
0.5	56.8	0.13	87
1.0	57.1	0.13	87

Analysis. The ridge penalty λ has minimal impact across several orders of magnitude ($<$1.6 pp variation). This is expected: with K=20 neighbours and n=8 dimensions, the system is well-overdetermined ($K \gg n$), so regularisation has little effect.

C.5 Equal Threshold ε (EQUAL_EPS)

Table 12: Full ε (EQUAL_EPS) sweep (Mode 4, all categories, accuracy % excl. equal).

ε	Safety	Beaut.	Lively	Wealthy	Boring	Depress.	Avg
0.3	66.3	56.5	57.5	72.3	59.7	52.8	60.9
0.5	66.3	56.5	56.4	72.3	59.4	52.8	60.6
0.8	66.3	56.5	57.1	72.3	**61.5**	54.3	61.3
1.2	66.2	57.4	56.6	**72.8**	**63.5**	54.4	61.8
2.0	**67.1**	**59.7**	**59.1**	72.4	63.3	**56.2**	**63.0**

Analysis. The equal threshold ε shows a mild but consistent trend: larger ε (more permissive equal classification) improves accuracy across most categories, with the average rising from 60.9% (ε=0.3) to 63.0% (ε=2.0). This suggests that many borderline pairs are better classified as non-equal, potentially because the VLM's score differences carry marginal but genuine signal even for close pairs.

C.6 Combined Search

In addition to the one-at-a-time sweeps above, we performed a random search over the combined parameter space (N=50,000 configurations) to check for interactions. Table 13 reports the per-category best configurations found.

Table 13: Best per-category configurations from combined random search (Mode 4, accuracy % excl. equal).

Category	Acc. (%)	κ	Key Parameters
Safety	81.6	0.636	K=30, α=0.2, sel=0.6
Beautiful	69.8	0.384	K=10, α=0.5, sel=0.6
Lively	69.4	0.404	K=50, α=0.7, sel=0.6
Wealthy	74.0	0.480	K=50, α=0.7, sel=0.6
Boring	70.2	0.414	K=30, α=0.7, sel=0.6
Depressing	68.2	0.370	K=20, α=0.2, sel=0.6
Average	72.2	0.448	—

The combined search yields 72.2% average accuracy (κ=0.448), a +13.8 pp gain over the default configuration (58.4%), demonstrating significant headroom from per-category hyperparameter tuning. The most consistent finding is sel=0.6 across all categories, while the optimal K (10–50) and α (0.2–0.7) vary by category.

D Prompt Templates

This section reproduces the exact prompt templates used in Stage 2 multi-agent feature distillation (Mode 4: pairwise + multi-agent). All three agents receive the image pair and the Stage 1–extracted dimension definitions as input. Temperature values: Observer τ=0.3, Debater τ=0.5, Judge τ=0.1.

D.1 Stage 1: Dimension Extraction Prompt

Listing 1: Dimension extraction prompt (Stage 1).

```
You are an expert urban perception researcher. I will show
you street-view images that have been rated by crowdsourced
participants on the perception dimension: "{category}".

HIGH-rated images (TrueSkill mu > 75th percentile):
{high_image_descriptions}

LOW-rated images (TrueSkill mu < 25th percentile):
{low_image_descriptions}

Based on the visual differences between HIGH and LOW images,
identify 5-10 specific, visually observable dimensions that
explain why some scenes are perceived as more "{category}"
than others.

Output JSON format:
{
  "dimensions": [
    {
      "name": "Dimension Name",
      "description": "What this dimension measures",
      "high_indicator": "Visual cues for high scores",
      "low_indicator": "Visual cues for low scores"
    }
  ]
}

Requirements:
- Each dimension must be visually observable from a street
  view image
- Each dimension must be continuously scorable (1-10 scale)
- Dimensions should be universally applicable across cities
```

D.2 Observer Prompt

Listing 2: Observer agent prompt (Stage 2, Mode 4).

```
You are an objective Observer. Examine two street-view
images and describe what you see along each evaluation
dimension. Do NOT make judgments or comparisons -- only
describe observable visual details.

Category: {category}
Dimensions:
{dimension_definitions}

For each dimension, describe what you observe in:
- Image A: [visual details]
- Image B: [visual details]

Output JSON:
{
  "observations": {
```

```
    "dimension_name": {
      "image_a": "description of what you see",
      "image_b": "description of what you see"
    }
  }
}
```

D.3 Debater Prompt

Listing 3: Debater agent prompt (Stage 2, Mode 4).

```
You are a Debater. Given the Observer's descriptions of
two street-view images, argue BOTH sides for each
dimension -- why Image A might score higher AND why
Image B might score higher.

Observer's descriptions:
{observer_output}

Category: {category}
Dimensions:
{dimension_definitions}

For each dimension, provide:
- Argument for Image A scoring higher
- Argument for Image B scoring higher
- Key uncertainties or ambiguities

Output JSON:
{
  "debates": {
    "dimension_name": {
      "argument_for_a": "why A might score higher",
      "argument_for_b": "why B might score higher",
      "uncertainties": "what is ambiguous"
    }
  }
}
```

D.4 Judge Prompt

Listing 4: Judge agent prompt (Stage 2, Mode 4).

```
You are the final Judge. Given the Observer's descriptions
and Debater's arguments, produce final scores for both
images on each dimension.

Observer's descriptions:
{observer_output}

Debater's arguments:
{debater_output}

Category: {category}
Dimensions:
{dimension_definitions}

Score each image on each dimension from 1 (lowest) to 10
(highest). Also determine the overall winner.

Output JSON:
{
  "image_a_scores": {
```

```
    "dimension_name": score (1-10)
  },
  "image_b_scores": {
    "dimension_name": score (1-10)
  },
  "overall_intensity_a": weighted_sum,
  "overall_intensity_b": weighted_sum,
  "winner": "left" | "right" | "equal",
  "ai_dimension_weights": {
    "dimension_name": weight (0-1)
  }
}
```

E Theoretical Foundations

This section establishes the mathematical foundations for UrbanAlign: post-hoc concept bottleneck calibration (Appendix E.1), local versus global calibration (Appendix E.2), multi-agent variance reduction (Appendix E.3), hybrid feature complementarity (Appendix E.4), the integrated framework guarantee (Appendix E.5), and E2E dimension optimization convergence (Appendix E.6).

Notation. Consider image pairs (A, B) with human label $y \in \{\text{left}, \text{right}, \text{equal}\}$. Let $\Delta_{\text{Sem}} = S(A) - S(B) \in \mathbb{R}^n$ denote the semantic score difference across n VLM-extracted dimensions, and $\Delta_{\text{CLIP}} = \phi(A) - \phi(B)$ the CLIP differential. The hybrid differential is $\Delta_{\text{hyb}} = [\alpha\,\overline{\Delta}_{\text{CLIP}},\ (1-\alpha)\,\overline{\Delta}_{\text{Sem}}]$ (Eq. (3) in the main text). Define the *discriminative power* of dimension j as $\delta_j = P(\text{sign}(\Delta_{\text{Sem},j}) = y)$, encoding y as ± 1 (excluding "equal").

E.1 Post-hoc Concept Bottleneck Calibration

Stage 2 produces concept scores with $\delta_j > 0.5$ (empirically 42–71%, Table 2), yet the VLM's *own* label predictions are poorly calibrated. LWRR provides the missing calibration layer.

Theorem 2 (Local Calibration Bound)**.** *Assume the local regression model $y_i^{\text{TS}} = w^*(q)^\top \Delta_{\text{Sem},i} + \epsilon_i$ holds within the K-neighbourhood $\mathcal{N}_K(q)$ of query q, where ϵ_i is zero-mean noise with variance σ^2 and $w^*(q) \in \mathbb{R}^n$ is the true local weight vector. Then the LWRR estimator $\hat{w} = (X^\top W X + \lambda I)^{-1} X^\top W y$ satisfies:*

$$\mathbb{E}\big[\|\hat{w} - w^*(q)\|^2\big] \;\le\; \underbrace{\frac{\sigma^2\,\text{tr}(H_\lambda^{-1} X^\top W^2 X\, H_\lambda^{-1})}{K}}_{\text{variance: } O(\sigma^2 n/K)} \;+\; \underbrace{\lambda^2 \|w^*\|^2}_{\text{regularisation bias}}, \tag{9}$$

where $H_\lambda = X^\top W X + \lambda I$. At optimal regularisation $\lambda^* = \sigma\sqrt{n/K}\,/\,\|w^*\|$, the bound simplifies to $O\big(\sigma \|w^*\| \sqrt{n/K}\big)$.

Proof. Write $y = Xw^* + \epsilon$. Then $\hat{w} - w^* = H_\lambda^{-1} X^\top W \epsilon - \lambda H_\lambda^{-1} w^*$. The first term contributes variance $\sigma^2\,\text{tr}(H_\lambda^{-1} X^\top W^2 X\, H_\lambda^{-1})$; the second contributes bias $\lambda^2 \|H_\lambda^{-1} w^*\|^2 \le \lambda^2 \|w^*\|^2$. Bounding the variance trace by $\sigma^2 n/(K\lambda_{\min}^2)$ and optimising over λ yields $\lambda^* = \sigma\sqrt{n/K}/\|w^*\|$ [Caponnetto and De Vito, 2007, Hoerl and Kennard, 1970]. In our setting (n=5–8, K=20–50), the system is well-overdetermined ($K \gg n$), ensuring small estimation variance. □

Corollary 3 (Calibration Exceeds Best Single Dimension)**.** *If every dimension has discriminative power $\delta_j > 1/2$ and the estimation error $\|\hat{w} - w^*\|$ is sufficiently small, then the LWRR predictor $\hat{y} = \text{sign}(\hat{w}^\top \Delta_{\text{Sem}})$ achieves accuracy strictly exceeding $\max_j \delta_j$.*

Proof. The optimal linear combination $w^{*\top} \Delta_{\text{Sem}}$ weights multiple informative dimensions, achieving higher accuracy than any single dimension whose accuracy is δ_j. When $\|\hat{w} - w^*\|$ is small (guaranteed by Theorem 2 when $K \gg n$), the estimated predictor preserves this advantage. Empirically: *wealthy* dimensions individually reach 61.5–71.2%, while LWRR achieves 74.0%. □

E.2 Advantage of Local over Global Calibration

Theorem 4 (Local vs. Global Excess Risk)**.** *Define* manifold heterogeneity*:*

$$\mathcal{H} = \mathbb{E}_q\big[\|w^*(q) - \bar{w}\|^2\big], \quad \bar{w} = \mathbb{E}_q[w^*(q)]. \tag{10}$$

A global ridge estimator fitted on the entire reference set (N pairs) incurs prediction MSE:

$$\mathrm{MSE}_{\mathrm{global}} \;\geq\; \mathcal{H} \cdot \mathbb{E}[\|\Delta\|^2] + O(\sigma^2 n/N). \tag{11}$$

LWRR replaces the irreducible heterogeneity bias with local estimation variance:

$$\mathrm{MSE}_{\mathrm{LWRR}} \;=\; O(\sigma^2 n/K) + B_K^2, \tag{12}$$

where $B_K^2 = \mathbb{E}[\|w^*(q) - w^*_{\mathcal{N}_K}\|^2]$ is the intra-neighbourhood weight variation. When $\mathcal{H}$ is large, $\mathrm{MSE}_{\mathrm{LWRR}} \ll \mathrm{MSE}_{\mathrm{global}}$ for any K such that $B_K^2 \ll \mathcal{H}$.

Proof. For the global model, the prediction at q is $\bar{w}^\top \Delta_q$, while the truth is $w^*(q)^\top \Delta_q$. The squared bias is $\|(w^*(q) - \bar{w})^\top \Delta_q\|^2$; averaging over q gives $\mathcal{H} \cdot \mathbb{E}[\|\Delta\|^2]$. For LWRR, the local bias is only B_K^2 (the variation of w^* *within* $\mathcal{N}_K$), which is much smaller than $\mathcal{H}$ (the variation across the *entire* manifold). LWRR pays increased variance ($O(\sigma^2 n/K)$ vs. $O(\sigma^2 n/N)$) by using $K \ll N$ effective samples, but this is compensated when $\mathcal{H}$ is large.

Empirically, the high coefficient of variation (CV > 1.0 for 7/8 *wealthy* dimensions in Table 6) and the qualitative weight-profile shifts (suburban vs. downtown examples in Appendix B) confirm substantial heterogeneity, validating the local approach. The sensitivity analysis (Appendix C) further confirms that K is the most sensitive hyperparameter (+4.3 pp range), reflecting the bias–variance trade-off formalised above. □

E.3 Multi-Agent Scoring: Variance Reduction

Proposition 5 (Deliberation Reduces Score Noise)**.** *Let $S_j^{(1)}$ be the score for dimension j from a single VLM call with variance σ_S^2. The Observer–Debater–Judge chain produces a score $S_j^{(3)}$ satisfying:*

$$\mathrm{Var}\big(S_j^{(3)}\big) \;\leq\; \frac{\sigma_S^2}{1 + 2(1-\rho)}, \tag{13}$$

where $\rho \in [0, 1]$ is the average pairwise correlation between agent reasoning paths. When agents provide substantially independent perspectives ($\rho \to 0$), the variance is reduced by up to $1/3$.

Proof. The Observer provides a structured description, conditioning the score on salient features. The Debater generates two opposing arguments, creating partially independent reasoning paths. The Judge aggregates all perspectives. Modelling this as $M{=}3$ correlated estimators with pairwise correlation ρ, the variance of their aggregate is $\sigma_S^2(1 + (M{-}1)\rho)/M$ [Wang et al., 2023]. Since the Judge weights agents by informativeness (rather than simple averaging), it achieves at least the averaging bound. The LWRR prediction $\hat{y} = \hat{w}^\top \Delta_{\text{Sem}}$ depends linearly on $\Delta_{\text{Sem}} = S(A){-}S(B)$; hence score noise propagates linearly. Reduced score variance tightens the calibration bound in Theorem 2 by lowering the effective σ^2.

Empirically, the $2{\times}2$ factorial (Table 3) confirms: multi-agent reasoning yields +8.5 pp (single-image) and +31.8 pp (pairwise). □

E.4 Hybrid Feature Space Complementarity

Proposition 6 (Hybrid Neighbourhood Quality)**.** *Let $\mathcal{N}_K^{\mathrm{Sem}}$ and $\mathcal{N}_K^{\mathrm{hyb}}$ denote the K-nearest neighbourhoods in pure semantic and hybrid space, respectively. If CLIP features carry label-relevant information beyond semantic scores, i.e.,*

$$I\big(y;\ \phi_{\mathrm{CLIP}} \;\big|\; \Delta_{\mathrm{Sem}}\big) > 0, \tag{14}$$

then the hybrid neighbourhood achieves strictly lower local bias: $B_K^{\mathrm{hyb}} < B_K^{\mathrm{Sem}}$.

Proof. The local bias B_K arises from variation of w^* within the neighbourhood (Theorem 4). In pure semantic space, two pairs may have similar Δ_{Sem} but very different visual contexts (e.g., suburban vs. downtown), leading to different local weights. If $I(y; \phi_{\text{CLIP}}|\Delta_{\text{Sem}}) > 0$, CLIP features distinguish these contexts. The hybrid distance groups pairs that are similar in *both* visual appearance and semantic content, selecting neighbours with more homogeneous local weight structures. Formally, for any partition of the manifold into regions where w^* is approximately constant, the hybrid neighbourhood has higher probability of staying within a single region, reducing B_K.

Empirically, $\alpha{=}0.3$ (70% semantic, 30% CLIP) is optimal on average (Table 9), confirming that semantic scores dominate but CLIP features provide complementary neighbourhood quality. □

E.5 Integrated Framework Guarantee

Combining the above results yields the following end-to-end characterisation.

Corollary 7 (End-to-End Alignment Bound)**.** *Under the conditions of Theorems 2 and 4 and Propositions 5 and 6, the UrbanAlign pipeline achieves alignment accuracy:*

$$\mathrm{Acc} \;\geq\; \Phi\left(\frac{\|w^*\| \cdot \bar{\delta}}{\sqrt{\sigma^2_{\mathrm{LWRR}} + \sigma^2_S/\eta_M}}\right), \tag{15}$$

where Φ is the standard normal CDF, $\bar{\delta}$ is the average dimension discriminability, $\sigma^2_{\mathrm{LWRR}} = O(\sigma^2 n/K)$ is the LWRR estimation variance (Theorem 2), and σ^2_S/η_M is the multi-agent-reduced score noise (Proposition 5 with $\eta_M = 1 + 2(1-\rho)$).

Each pipeline stage tightens a specific term:

- ***Stage 1** (concept mining) maximises $\bar{\delta}$ by discovering highly discriminative dimensions;*
- ***Stage 2** (multi-agent scoring) minimises σ^2_S/η_M via deliberative aggregation;*
- ***Stage 3** (LWRR calibration) minimises σ^2_{LWRR} and B_K through local hybrid-space fitting;*
- ***Stage 4** (E2E optimisation) jointly optimises $\bar{\delta}$ and the dimension set $\mathcal{D}_c$.*

Proof. Under the linear-Gaussian local model, $\hat{y} = \hat{w}^\top \Delta_{\text{Sem}}$ has mean $w^{*\top}\Delta_{\text{Sem}}$ and variance $\sigma^2_{\mathrm{LWRR}} + \|w^*\|^2\sigma^2_S/\eta_M$ (from LWRR estimation noise and dimension score noise, respectively). The pair is correctly classified when $\mathrm{sign}(\hat{y}) = y$, i.e., $\hat{y}$ has the correct sign. For a linear discriminant with signal-to-noise ratio $\mathrm{SNR} = |w^{*\top}\mathbb{E}[\Delta]|/\mathrm{Var}(\hat{y})^{1/2}$, the accuracy is $\Phi(\mathrm{SNR})$. The numerator scales with $\|w^*\| \cdot \bar{\delta}$ (the weighted discriminability); the denominator collects the noise terms, each tightened by the corresponding stage. □

E.6 E2E Dimension Optimization: Convergence and Sample Complexity

Information-bottleneck interpretation. Each dimension set $\mathcal{D}_c^{(t)}$ defines a different compression of the visual input through the lens of the information bottleneck [Tishby et al., 2000]. Let X denote the raw visual input and Y the human label. Stage 2 compresses X into dimension scores $S^{(t)} = f_{\mathcal{D}_c^{(t)}}(X)$, and Stage 3 maps $S^{(t)}$ to prediction $\hat{Y}$. The optimal dimension set maximises $I(\hat{Y}; Y)$ subject to the constraint that the bottleneck representation $S^{(t)}$ remains compact ($n \in [5, 10]$ dimensions). Equation (17) selects the compression that maximises this downstream mutual information empirically.

Convergence guarantee. Let $\mathcal{F}$ denote the (finite) set of dimension sets reachable by the VLM at any temperature in $[\tau_{\min}, \tau_{\max}]$. Although $|\mathcal{F}|$ is large (combinatorially many possible dimension names and descriptions), the VLM's generative distribution concentrates probability mass on a much smaller effective set $\mathcal{F}_{\text{eff}} \subset \mathcal{F}$ of semantically coherent dimension sets. If we model each trial as drawing independently from the VLM's distribution over $\mathcal{F}_{\text{eff}}$, then after T trials the probability of missing the optimal set $\mathcal{D}_c^*$ satisfies:

$$\Pr\big[\mathcal{D}_c^* \notin \{\mathcal{D}_c^{(1)}, \ldots, \mathcal{D}_c^{(T)}\}\big] \leq (1 - p^*)^T, \tag{16}$$

where $p^* = \Pr_{\mathrm{VLM}}[\mathcal{D}_c^{(t)} = \mathcal{D}_c^*]$ is the probability of generating the optimal set in a single trial. For a near-optimal set (within ϵ accuracy of the best), this probability is typically higher. Concretely, if the top-k dimension sets are each generated with probability $\geq p_{\min}$, then $T \geq \lceil \ln(1/\delta)/\ln(1/(1 - k \cdot p_{\min})) \rceil$ trials suffice to find at least one with probability $\geq 1 - \delta$.

Search strategy: explore → converge as exploration–exploitation. The two-phase temperature schedule implements a structured search that balances two objectives:

- *Exploration* (explore phase, τ=0.85 → 1.0): higher temperatures increase sampling entropy, producing diverse dimension sets that broadly cover the VLM's generative distribution. Each category independently accumulates its best set, with the current best provided as a soft reference.
- *Exploitation* (converge phase, τ=0.7 → 0.5): lower temperatures with *mutation mode* (retain all but 1–2 dimensions from the per-category best) concentrate the search around known-good dimension sets, refining them locally.

This is analogous to random search with a guided prior [Bergstra and Bengio, 2012]: rather than sampling uniformly from the combinatorial space of possible dimensions, we leverage the VLM's own generative distribution as an implicit prior over useful dimension sets. Unlike Bayesian optimisation [Snoek et al., 2012], which would require a surrogate model over the discrete, variable-length dimension space, our approach avoids the need for a surrogate entirely.

Sample complexity for evaluation. Each trial evaluates alignment accuracy on a fixed sample of m pairs per category. By Hoeffding's inequality, the empirical accuracy $\widehat{\mathrm{Acc}}$ satisfies $|\widehat{\mathrm{Acc}} - \mathrm{Acc}| \leq \sqrt{\ln(2T/\delta)/(2m)}$ uniformly over all T trials with probability $\geq 1-\delta$. For T=12 and $m\approx$160 test pairs per category (average across categories), this gives a confidence interval of $\pm$13.9% at δ=0.05. While wider than ideal, the observed accuracy gaps between trials are substantial (e.g., safety: 35.3%–76.5%, a 41.2 pp range), confirming that the evaluation sample is sufficient to distinguish meaningfully different dimension sets.

Connection to LWRR convergence. The end-to-end accuracy depends not only on the dimension set but also on the LWRR calibration quality. For regularised least-squares with K local neighbours in n dimensions, the excess risk decays as $\mathcal{O}(\lambda + n/(K\lambda))$ [Caponnetto and De Vito, 2007, Hoerl and Kennard, 1970], which is minimised at $\lambda^* \propto (n/K)^{1/2}$. In our setting ($n$=5–8, K=20), the system is well-overdetermined ($K \gg n$), so the LWRR calibration error is small regardless of the dimension set, confirming that the E2E accuracy variation across trials is driven primarily by the quality of the dimension-level representation, not by the calibration algorithm.

F End-to-End Dimension Optimization

The semantic dimensions discovered in Stage 1 depend on the VLM's sampling temperature and the particular consensus exemplars presented. A natural question is: *can we automatically find the dimension set that maximises end-to-end alignment accuracy?*

Formulation. Let $\mathcal{D}_c^{(t)}$ denote the dimension set generated at trial t for category c. Each trial runs the full pipeline: dimension extraction (Stage 1) $\rightarrow$ multi-agent pairwise scoring (Stage 2, Mode 4) $\rightarrow$ LWRR alignment (Stage 3) $\rightarrow$ accuracy evaluation. The per-category optimal dimension set is:

$$\mathcal{D}_c^* = \arg\max_{t \in \{1,\ldots,T\}} \mathrm{Acc}_c\big(\mathrm{LWRR}\big(\mathrm{Score}(\mathcal{D}_{\mathrm{val}}^c,\, \mathcal{D}_c^{(t)}),\, \mathcal{D}_{\mathrm{ref}}^c\big)\big). \tag{17}$$

Two-phase search strategy. Rather than monotonic temperature increase, we employ an *explore* $\rightarrow$ *converge* schedule over T=15 trials (with patience-based early stopping after 5 consecutive trials without any per-category improvement):

- *Explore phase* (trials 0–5, τ_{gen}=0.85 $\rightarrow$ 1.0): Higher temperatures produce diverse, largely independent dimension sets. Each category independently accumulates its best. The current per-category best is provided as a reference but large deviations are encouraged.

- *Converge phase* (trials 6+, τ_{gen}=0.7 $\rightarrow$ 0.5): Lower temperatures with *mutation mode*: per-category best dimensions are preserved except for 1–2 targeted replacements, enabling local refinement.

Since categories may peak at different trials, the final output *assembles* per-category optima: $\{\mathcal{D}_c^*\}_{c=1}^6$, potentially sourced from different trials.

Cost analysis. Table 14 summarises the full API cost breakdown. All VLM calls use GPT-4o with measured per-call token usage of 778 input + 278 output tokens, yielding a cost of \$0.0047/call at OpenAI's published rates (\$2.50 / \$10.00 per 1M input/output tokens).

Table 14: API cost breakdown for the full experimental pipeline (4,819 pairs across 6 categories, GPT-4o pricing). Stage 3 LWRR is purely local computation with zero API cost.

Component	Calls/pair	Total calls	Tokens (M)	Cost (USD)
Main pipeline (4,819 pairs × 6 categories)				
Mode 1 (Single-Image Direct)	2	9,638	10.2	\$45.59
Mode 2 (Pairwise Direct)	1	4,819	5.1	\$22.79
Mode 3 (Single-Image Multi-Agent)	6	28,914	30.5	\$136.76
Mode 4 (Pairwise Multi-Agent)	3	14,457	15.3	\$68.38
End-to-end dimension optimization (20% subsample, Mode 2, 12 trials)				
Stage 1 dimension generation	—	72	0.1	\$0.34
Stage 2 scoring	1	∼11,600	12.2	\$54.72
Total (all 4 modes + E2E)		**69,500**	**73.4**	**\$328.58**
Mode 4 only + E2E		26,129	27.6	\$123.44
Production projection (full dataset ∼164K pairs, Mode 4)				
Mode 4 at scale	3	492,900	520.6	∼\$2,329

Per-call: 778 input + 278 output = 1,056 tokens. Input: \$2.50/1M, Output: \$10.00/1M ⇒ \$0.0047/call. Traditional crowdsourcing comparison: Place Pulse 2.0 collected 1.17M annotations at ∼\$0.14/comparison ≈ \$167K, yielding a **98.6%** cost reduction at production scale.

Table 15: E2E dimension optimization: per-category accuracy (%, excluding "equal") across 12 trials (all six categories, Mode 4). **Bold** marks the per-category best. Trials 0–5 are *explore* phase; 6–11 are *converge* phase.

Trial	Phase	τ	Safety	Beaut.	Lively	Wealthy	Boring	Depress.	Avg
0	Explore	0.85	58.8	61.5	47.1	55.6	47.1	50.0	53.3
1	Explore	0.88	**76.5**	46.2	47.1	**77.8**	35.3	**68.8**	58.6
2	Explore	0.91	70.6	**69.2**	35.3	61.1	35.3	56.3	54.6
3	Explore	0.94	64.7	53.8	47.1	50.0	35.3	62.5	52.2
4	Explore	0.97	35.3	53.8	52.9	72.2	47.1	31.3	48.8
5	Explore	1.00	58.8	46.2	52.9	55.6	47.1	56.3	52.8
6	Converge	0.70	70.6	46.2	52.9	44.4	35.3	37.5	47.8
7	Converge	0.68	52.9	53.8	**58.8**	72.2	52.9	50.0	56.8
8	Converge	0.65	58.8	61.5	41.2	77.8	17.6	50.0	51.2
9	Converge	0.63	58.8	53.8	47.1	44.4	47.1	43.8	49.2
10	Converge	0.60	52.9	46.2	52.9	44.4	**64.7**	50.0	51.9
11†	Converge	0.58	41.2	50.0	—	—	—	—	45.6
Assembled best			**76.5**	**69.2**	**58.8**	**77.8**	**64.7**	**68.8**	**69.3**

†Trial 11 scored only 2/6 categories (early stopping triggered). Assembled best: per-category maximum across all trials. — = category not scored.

Table 16: Per-category assembled best dimensions from E2E optimization (default calibration hyperparameters: K=20, α=0.3, sel=1.0). Each category's optimal dimensions may come from a different trial, enabling independent specialisation. Note: these results differ from Table 1, which uses per-category *optimised hyperparameters* with original dimensions (72.2% avg); here only the *dimension set* is optimised (69.3% avg). The two optimisation axes are complementary.

Category	Trial	Phase	Acc (%)	κ	Optimised Dimensions
Safety	1	Explore	76.5	0.534	Lighting Adequacy, Pedestrian Infrastructure, Building Maint., Street Activity, Visibility Clarity, Greenery & Landscaping
Beautiful	2	Explore	69.2	0.458	Greenery & Nat. Elements, Street Cleanliness, Arch. Coherence, Color Coordination, Pedestrian Amenities, Visual Complexity
Lively	7	Converge	58.8	0.308	Human Presence, Vegetation & Greenery, Building Diversity, Street Furniture, Commercial Activity, Color & Lighting, Public Art & Signage
Wealthy	1	Explore	77.8	0.579	Façade Quality, Vegetation Maint., Pavement Integrity, Vehicle Quality, Building Modernity, Infrastructure Cond., Street Cleanliness, Lighting Quality
Boring	10	Converge	64.7	0.407	Visual Complexity, Activity Presence, Arch. Diversity, Vegetation Variety, Color Vibrancy, Street Furniture Presence, Spatial Enclosure
Depressing	1	Explore	68.8	0.437	Façade Condition, Street Cleanliness, Greenery Presence, Lighting Quality, Visual Clutter
Average			**69.3**	**0.454**	—

Results. Table 15 reports trial-by-trial accuracy for all six categories. The single best trial (Trial 1, explore phase, τ_{gen}=0.88) achieves 58.6% average accuracy, but per-category assembly (Table 16) raises this to **69.3%** (κ=0.454), a +10.7 pp gain that demonstrates categories have distinct optimal dimension sets.

Three key findings emerge:

(1) **Explore phase dominates**: 4/6 categories (safety, beautiful, wealthy, depressing) find their best dimensions during exploration (trials 0–5). Trial 1 (τ=0.88) is particularly productive, contributing optima for three categories.

(2) **Converge phase is complementary**: *Lively* peaks at trial 7 and *boring* at trial 10, both in the converge phase. For these categories, the explore phase produced dimensions below 53%, while targeted mutation in the converge phase lifted accuracy to 58.8% and 64.7% respectively.

(3) **High variance across trials**: Per-category accuracy fluctuates substantially across trials (e.g., *safety*: 35.3%–76.5%; *boring*: 17.6%–64.7%), confirming that dimension set quality is a first-order determinant of alignment performance and justifying the automated search.

The assembled average of 69.3% represents a +10.9 pp gain over the default pipeline (58.4%) reported in the main text, achieved without changing any algorithmic hyperparameter; only the semantic dimensions are optimised. The full theoretical analysis of E2E convergence, sample complexity, and connection to LWRR is provided in Appendix E.6.

G LWRR Alignment Gain Per Mode

Table 17 reports the accuracy change from raw VLM output to post-LWRR calibration for each of the four factorial modes on *wealthy*. The all-category Mode 4 LWRR gain is reported in Table 4.

Table 17: LWRR alignment gain (accuracy % excluding "equal") per mode, *wealthy* category. Mode 4 shows the largest calibration gain (+11.1 pp).

	Mode 1	Mode 2	Mode 3	Mode 4
Raw Acc. (%)	48.7	63.3	58.8	62.9
Aligned Acc. (%)	55.4	42.2	63.9	74.0
Δ (pp)	+6.7	-21.1	+5.1	**+11.1**
Raw κ	0.208	0.263	0.266	0.256
Aligned κ	0.159	-0.037	0.313	**0.480**
n (test pairs)	116	116	116	107

Notably, only Mode 4 benefits substantially from LWRR on *wealthy* (+11.1 pp). Mode 3 gains modestly (+5.1 pp), Mode 1 gains slightly (+6.7 pp), and Mode 2 drops sharply (-21.1 pp), indicating that the combination of pairwise context and multi-agent deliberation is critical for LWRR to learn effective local mappings.

H Complete Algorithm Pseudocode

Algorithm 1 UrbanAlign: Post-hoc VLM Calibration for Urban Perception

Require: $\mathcal{D}_L$: labelled pairs; ϕ_{CLIP}: CLIP encoder; VLM; $K, \tau, \lambda, \alpha, \varepsilon, \theta$
Ensure: $\mathcal{D}_{\text{aligned}}$: calibrated dataset
// Stage 1: Semantic Dimension Extraction
Compute TrueSkill ratings $\{(\mu_i, \sigma_i)\}$ from $\mathcal{D}_L$
Sample $\mathcal{S}_{\text{high}}, \mathcal{S}_{\text{low}}$ via consensus criteria
$\mathcal{D}_c \leftarrow \text{VLM}(\mathcal{S}_{\text{high}} \cup \mathcal{S}_{\text{low}})$ {n dimensions}
// Stage 2: Multi-Agent Feature Distillation
Sample $\mathcal{D}_{\text{sample}} \subset \mathcal{D}_L$; split into $\mathcal{D}_{\text{ref}}, \mathcal{D}_{\text{val}}, \mathcal{D}_{\text{test}}$
for each image pair $(x_A, x_B) \in \mathcal{D}_{\text{sample}}$ **do**
    obs $\leftarrow \text{VLM}_{\text{Obs}}(x_A, x_B, \mathcal{D}_c)$
    debate $\leftarrow \text{VLM}_{\text{Deb}}(x_A, x_B, \text{obs}, \mathcal{D}_c)$
    $S(x_A), S(x_B) \leftarrow \text{VLM}_{\text{Judge}}(x_A, x_B, \text{obs}, \text{debate}, \mathcal{D}_c)$
end for
// Stage 3: LWRR Calibration
Build reference manifold $\mathcal{R}$ from $\mathcal{D}_{\text{ref}}$ with mirror augment
for each $(x_A, x_B) \in \mathcal{D}_{\text{test}}$ **do**
    Compute Δ_q^{hybrid}; find K-NN $\mathcal{N}_K$ in $\mathcal{R}$
    Solve LWRR (Eq. (5)); compute $\hat{\delta}$, R^2
    Re-infer $\hat{y}$ via Eq. (6)
end for
return $\mathcal{D}_{\text{aligned}}$

I Related Human Preference Datasets

For reference, we survey publicly available human preference datasets that share the pairwise comparison structure used in this work. Table 18 organises these datasets by domain. This survey is intended as contextual background rather than a claim of demonstrated generalisability; validating the concept-mining–calibration paradigm on additional domains remains future work.

Urban & scene perception. Place Pulse 2.0 [Salesses et al., 2013], used in this work, collects pairwise comparisons across six urban perception dimensions. SPECS [Quintana et al., 2025] extends the paradigm to 10 perception dimensions and records annotator demographics and personality traits. ScenicOrNot collects absolute scenicness ratings for 200K+ locations across Great Britain.

Image generation & aesthetic quality. HPD v2 (798K pairs) and Pick-a-Pic (500K+ pairs) collect pairwise preferences on text-to-image outputs; ImageReward adds expert-annotated alignment/fidelity/harmlessness dimensions. AVA [Murray et al., 2012] provides crowd-sourced aesthetic ratings for 250K photographs.

Image quality assessment (IQA). KonIQ-10k and PIPAL provide mean opinion scores and pairwise comparisons for natural image quality. AGIQA-3K focuses on AI-generated image quality with multi-attribute annotations.

Text & multimodal preference (for context). HH-RLHF, SHP, and Chatbot Arena represent the RLHF paradigm, which aligns model outputs to human preferences by training reward models on pairwise comparisons. These methods use the same data structure but take a fundamentally different (training-based) approach.

Table 18: Public human preference datasets. "Pairwise" indicates the native annotation format uses pairwise comparisons (the same format as UrbanAlign). Datasets marked with ⋆ are the most direct candidates for UrbanAlign transfer experiments.

Dataset	Domain	Scale	Format	Access
Urban & Scene Perception				
Place Pulse 2.0⋆	Urban 6-dim	1.17M pairs	Pairwise	`pulse.media.mit.edu`
SPECS⋆	Urban 10-dim	10 dims + demo.	Pairwise	`ual.sg/project/specs`
ScenicOrNot⋆	Scenicness	200K+ locs	Rating	`scenic.mysociety.org`
Streetscore	Safety	Boston/NYC	Rating	MIT Media Lab
Image Generation Quality				
HPD v2	T2I pref.	798K pairs	Pairwise	`github.com/tgxs002/HPSv2`
Pick-a-Pic	T2I pref.	500K+ pairs	Pairwise	HuggingFace
ImageReward	T2I expert	137K pairs	Pairwise	`github.com/zai-org/ImageReward`
RichHF-18K	T2I regions	18K images	Multi-attr.	Google Research
Aesthetic & Image Quality				
AVA	Photo aesth.	250K images	Rating	`github.com/mtobeiyf/ava_downloader`
KonIQ-10k	Natural IQA	10K images	MOS	`database.mmsp-kn.de`
PIPAL	Perceptual IQA	29K pairs	Pairwise	ECCV 2020
AGIQA-3K	AI-gen. IQA	3K images	Multi-attr.	NeurIPS 2024
Text & Multimodal (RLHF, for context)				
HH-RLHF	Text pref.	170K pairs	Pairwise	`github.com/anthropics/hh-rlhf`
SHP	Reddit pref.	385K pairs	Pairwise	HuggingFace
Chatbot Arena	LLM pref.	Ongoing	ELO rank	LMSYS

J Discussion

Why does post-hoc calibration work?

The central finding is that VLMs are strong *concept extractors* but poor *decision calibrators* for urban perception. Mode 4 raw dimension scores already contain substantial perceptual information (individual dimensions at 53–71% discriminative power), yet the VLM's own label predictions are only marginally above the zero-shot baseline. LWRR solves a simple and intuitive problem: given that dimension scores $S(x)$ are informative, learn the locally-optimal linear mapping from score differences to TrueSkill rating differences. This is precisely the post-hoc CBM paradigm [Yuksekgonul et al., 2023]: extract concepts from a frozen backbone, then train a lightweight predictor on those concepts.

Why local rather than global calibration? Urban perception is heterogeneous: the visual cues that signal wealth in a suburban neighbourhood (vegetation, vehicle quality) may differ from those in an urban core (building modernity, infrastructure). A single global linear model would average over these distinct regimes. LWRR, by fitting independent weights per neighbourhood, naturally adapts to local manifold geometry. The sensitivity analysis confirms this: $K_{\max}$ is the most sensitive single parameter (+4.3 pp range), reflecting the tension between local adaptivity and statistical stability. Theoretically, the LWRR excess risk at optimal regularisation decays as $\mathcal{O}((n/K)^{1/2})$ [Caponnetto and De Vito, 2007], explaining why larger K improves accuracy when $K \gg n$; the per-pair weight analysis in Appendix B confirms that dimension importances vary substantially across the manifold, validating the need for local fitting.

Connection to the information bottleneck. Our pipeline implements a form of information bottleneck [Tishby et al., 2000]: Stage 2 compresses visual input X into dimension scores $S{=}f_{\mathcal{D}_c}(X) \in \mathbb{R}^n$ (a compact bottleneck), and Stage 3 maximises the downstream mutual information $I(\hat{Y}; Y)$ via LWRR. The R^2 metric quantifies concept completeness, *i.e.*, how much of the TrueSkill variance is explainable by the bottleneck representation. The end-to-end dimension optimization (Eq. (2)) formalises this further: selecting the bottleneck that maximises $I(\hat{Y}; Y)$ is equivalent to choosing the dimension set with highest alignment accuracy. Elite seeding biases subsequent trials toward the current-best

bottleneck, converting the search from independent sampling into a guided refinement with exponential convergence guarantee.

Why Place Pulse and scope of evaluation. We evaluate on Place Pulse 2.0 because it provides pairwise annotations across six diverse perception categories, each representing a distinct subjective judgement task with different visual cues. This internal diversity, spanning concrete attributes (*safety*: infrastructure, lighting) to abstract ones (*boring*: monotony, stimulation deficit), already tests the framework's adaptivity across heterogeneous concept spaces. The underlying concept-mining–calibration paradigm requires only (i) a VLM that can describe domain-relevant attributes and (ii) a small set of pairwise human preferences; we note that this structure is shared by preference tasks in other domains (*e.g.*, aesthetic quality [Murray et al., 2012], extended urban perception with demographic metadata [Quintana et al., 2025]), which we leave for future work.

Limitations. While we evaluate across all six Place Pulse 2.0 categories, performance varies substantially: *safety* (81.6%) far exceeds *depressing* (68.2%) and *lively* (69.4%). The per-category sample sizes range from 84 to 266 test pairs after excluding equal labels. The current approach assumes static street-view images; temporal and cultural variation remain unexplored.

Ethical considerations. Automated perception mapping could inform equitable infrastructure investment but also risks reinforcing neighbourhood stereotypes if applied without community engagement. We recommend use for improvement-oriented planning (identifying streets needing façade repair, pavement maintenance) rather than discriminatory rankings.

Computational cost. The entire pipeline requires only VLM inference calls with no GPU training. At measured token usage of 778 input + 278 output tokens per GPT-4o call ($2.50 / $10.00 per 1M tokens), each call costs $0.0047. The primary method (Mode 4, Observer–Debater–Judge) requires 3 calls per pair: for our 4,819-pair experimental dataset across six categories, this totals 14,457 calls (**$68**). The full 4-mode ablation costs $274 (57,828 calls); end-to-end dimension optimization adds $55 (∼11,700 calls), for a total experimental budget of ∼**$329**. At production scale with the full Place Pulse 2.0 filtered dataset (∼164K pairs, Mode 4 only), projected cost is ∼**$2,300**, representing a **98.6%** reduction compared to traditional crowdsourcing at comparable scale (∼$167K for 1.17M pairwise annotations at $0.14/comparison [Salesses et al., 2013]).

References

Yecheng Zhang, Rong Zhao, Zimu Huang, Xinyu Wang, Yue Ma, and Ying Long. GenAI models capture urban science but oversimplify complexity. *arXiv preprint arXiv:2505.13803*, 2025.

Rashid Mushkani. Do vision-language models see urban scenes as people do? an urban perception benchmark. *arXiv preprint arXiv:2509.14574*, 2025.

Naila Murray, Luca Marchesotti, and Florent Perronnin. AVA: A large-scale database for aesthetic visual analysis. In *CVPR*, pages 2408–2415, 2012.

Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. Q-Align: Teaching LMMs for visual scoring via discrete text-defined levels. In *ICML*, 2024.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *NeurIPS*, pages 27730–27744, 2022.

Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *ICML*, pages 5338–5348, 2020.

Mert Yuksekgonul, Maggie Wang, and James Zou. Post-hoc concept bottleneck models. In *ICLR*, 2023.

Tuomas Oikarinen, Subhro Das, Lam M. Nguyen, and Tsui-Wei Weng. Label-free concept bottleneck models. In *ICLR*, 2023.

Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In *CVPR*, 2023.

Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulić, Anna Korhonen, and Nigel Collier. Aligning with human judgement: The role of pairwise preference in large language model evaluators. In *COLM*, 2024.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *ICML*, 2024.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *ICLR*, 2023.

William S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368):829–836, 1979.

Philip Salesses, Katja Schechtner, and César A. Hidalgo. The collaborative image of the city: Mapping the inequality of urban perception. *PLoS ONE*, 8(7):e68400, 2013.

James Q. Wilson and George L. Kelling. Broken windows: The police and neighborhood safety. *Atlantic Monthly*, 249 (3):29–38, 1982.

Jane Jacobs. *The Death and Life of Great American Cities*. Random House, 1961.

Abhimanyu Dubey, Nikhil Naik, Devi Parikh, Ramesh Raskar, and César A. Hidalgo. Deep learning the city: Quantifying urban perception at a global scale. In *ECCV*, pages 196–212, 2016.

Nikhil Naik, Jade Philipoom, Ramesh Raskar, and César A. Hidalgo. Streetscore—predicting the perceived safety of one million streetscapes. In *CVPR Workshop on Web-scale Vision and Social Media*, 2014.

Fan Zhang, Bolei Zhou, Liu Liu, Yu Liu, Helene H. Fung, Hui Lin, and Carlo Ratti. Measuring human perceptions of a large-scale urban region using machine learning. *Landscape and Urban Planning*, 180:148–160, 2018.

Rui Jin and Chenfan Cai. Plausible or misleading? evaluating the adaption of the Place Pulse 2.0 dataset for predicting subjective perception in Chinese urban landscapes. *Habitat International*, 157, 2025.

Matias Quintana, Youlong Gu, Xiucheng Liang, Yujun Hou, Koichi Ito, Yihan Zhu, Mahmoud Abdelrahman, and Filip Biljecki. Global urban visual perception varies across demographics and personalities. *Nature Cities*, 2(11): 1092–1106, 2025.

Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. ZeroGen: Efficient zero-shot learning via dataset generation. In *EMNLP*, 2022.

Yu Meng, Martin Michalski, Jiaxin Huang, Yu Zhang, Tarek Abdelzaher, and Jiawei Han. Tuning language models as training data generators for augmentation-enhanced few-shot learning. In *ICML*, 2023.

Yibo Yan, Haomin Wen, Siru Zhong, Wei Chen, Haodong Chen, Qingsong Wen, Roger Zimmermann, and Yuxuan Liang. UrbanCLIP: Learning text-enhanced urban region profiling with contrastive language-image pretraining from the web. In *WWW*, 2024.

Weiming Huang, Jing Wang, and Gao Cong. Zero-shot urban function inference with street view images through prompting a pretrained vision-language model. *International Journal of Geographical Information Science*, 38(7), 2024.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent debate. In *EMNLP*, 2024.

Yexiang Liu, Jie Cao, Zekun Li, Ran He, and Tieniu Tan. Breaking mental set to improve reasoning through diverse multi-agent debate. In *ICLR*, 2025a.

Divyansh Srivastava, Ge Yan, and Tsui-Wei Weng. VLG-CBM: Training concept bottleneck models with vision-language guidance. In *NeurIPS*, 2024.

Yang Liu, Tianwei Zhang, and Shi Gu. Hybrid concept bottleneck models. In *CVPR*, 2025b.

Chung-En Sun, Tuomas Oikarinen, Berk Ustun, and Tsui-Wei Weng. Concept bottleneck large language models. In *ICLR*, 2025.

Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. Locally weighted learning. *Artificial Intelligence Review*, 11:11–73, 1997.

David Ruppert and Matthew P. Wand. Multivariate locally weighted least squares regression. *The Annals of Statistics*, 22(3):1346–1370, 1994.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *NeurIPS*, pages 9459–9474, 2020.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021.

Ralf Herbrich, Tom Minka, and Thore Graepel. TrueSkill: A Bayesian skill rating system. In *NeurIPS*, pages 569–576, 2006.

OpenAI. GPT-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, 2015.

Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018.

Kevin Lynch. *The Image of the City*. MIT Press, 1960.

C. Ray Jeffery. Crime prevention through environmental design. *American Behavioral Scientist*, 14:598–610, 1971.

Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *JMLR*, 13:281–305, 2012.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *NeurIPS*, pages 2951–2959, 2012.