FLATLANDS: INDOOR FLOORMAP COMPLETION FROM A SINGLE EGOCENTRIC VIEW

Anonymous authors
Paper under double-blind review

ABSTRACT

A single egocentric image typically captures only a small portion of the floor, yet a complete metric traversability map of the surroundings would better serve applications such as indoor navigation. We introduce FlatLands, a dataset and benchmark for single-view indoor floormap completion in metric bird's-eye view (BEV) conditioned on a single egocentric view of the room. The dataset contains 270,575 observations from 17,656 scene records across six metric 3D datasets, with aligned observation, visibility, validity, and ground-truth maps and both in- and out-of-distribution protocols. We compare training-free, deterministic, ensemble, and generative methods on shared geometry-derived inputs and, without fine-tuning, on maps estimated from real egocentric RGB images. Our results show that predicting floor almost everywhere can achieve high IoU when most evaluated cells are floor, while generative sample sets outperform the best single-output model under a proper scoring rule and support safer risk-weighted planning. FlatLands establishes a rigorous testbed for evaluating both floor-map accuracy and how effectively alternative completions represent the ambiguity of unseen space.

1 INTRODUCTION

Indoor autonomy requires decisions from incomplete, noisy, and viewpoint-limited evidence (Kaelbling et al., 1998; Curtis et al., 2024). In one egocentric image, occlusion and a limited field of view hide much of the floor. Scene completion infers unseen geometry (Song et al., 2017; Cao & de Charette, 2022), yet several layouts fit one image: a doorway behind furniture may or may not exist, and each case implies a different route.

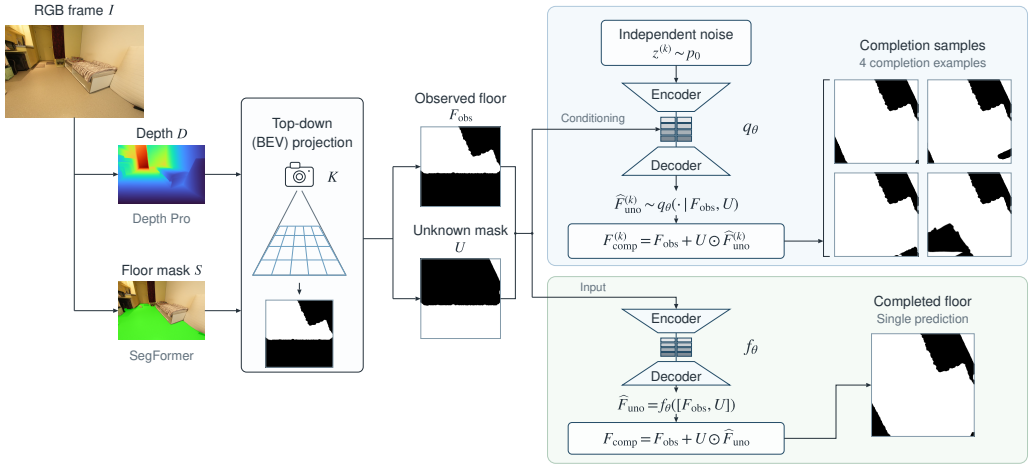


Figure 1: **Single-view floormap completion.** Depth and floor segmentation from RGB are projected using camera intrinsics to obtain (F_{obs}, U). **Upper branch:** these maps condition the generative model, while independent noise draws produce alternative completions. **Lower branch:** the deterministic model receives the concatenated maps directly and predicts a single completion.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

For planar indoor navigation, a metric 2D grid that retains floor layout and connectivity suffices (Ho et al., 2025; Wang et al., 2025): it supports reachability queries, collision checking, and cost-based route selection (LaValle, 2006). Such grids are also economical. Widely used navigation stacks plan on 2D costmaps (Marder-Eppstein et al., 2010; Macenski et al., 2020), whereas volumetric maps need dedicated data structures to remain tractable (Hornung et al., 2013). Our deterministic models complete a 10×10 m floormap in under 0.1 s (Section 5). Anticipating unobserved floor aids exploration (Ramakrishnan et al., 2020). We therefore study *single-view indoor bird's-eye-view (BEV) floor completion*: predicting floor labels in unobserved cells of a metric ground-plane grid while preserving the supplied observed labels.

Generative models (Ho et al., 2020; Lipman et al., 2023) can draw many completions of one observation and increasingly support robotic perception and planning (Reed et al., 2024; Carvalho et al., 2025). Aggregate IoU cannot show whether these samples help navigation: it rewards predicting floor wherever floor predominates and scores one map at a time. FlatLands therefore judges completions individually, as plausible sets, and by the routes they support, with every method given the same evidence.

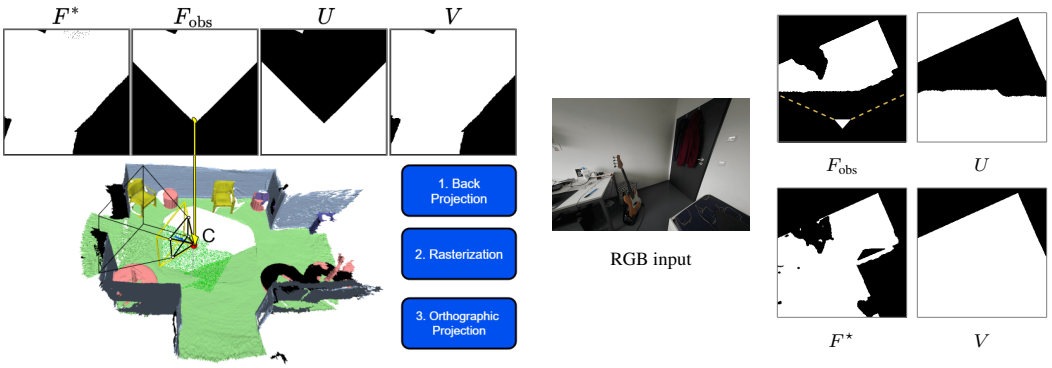
Contributions. **(i)** We formalize single-view metric indoor BEV floor completion and introduce FlatLands, with $270{,}575$ observations from $17{,}656$ metric scene records across six datasets (Chang et al., 2017; Dai et al., 2017; Yeshwanth et al., 2023; Baruch et al., 2021; Wald et al., 2019; Cruz et al., 2021); each observation has four aligned binary maps and provenance metadata. Five sources are in-distribution (ID); ScanNet++ is held out for out-of-distribution (OOD) evaluation. **(ii)** A source-geometry track gives all methods identical geometric evidence; an RGB-to-BEV track evaluates the same trained weights, without fine-tuning, on maps estimated from real images. Both tracks score single outputs, oracle best-of-K candidates, and completion sets, the last with the strictly proper masked energy score (MES; Gneiting & Raftery, 2007; Appendix F); a challenging subset isolates cases with scarce observed floor and a non-floor majority. **(iii)** Across eleven methods, overlap and utility rank methods differently. Predicting floor everywhere scores well on overlap yet collapses when visible floor is scarce. U-Net produces the most accurate single maps, yet every generator scores better as a set, and planning over generated samples yields safer routes (83% Safe Success) than planning on U-Net's maps or on the same FM+XAttn samples collapsed into one map (78% and 80%, respectively; Table 6).

2 Background and Related Work

Mapping and prediction beyond observations. Classical occupancy mapping integrates sensor evidence (Elfes, 1989; Thrun et al., 2005; Hornung et al., 2013). Occupancy anticipation and local 3D synthesis infer unseen geometry (Ramakrishnan et al., 2020; Reed et al., 2024), while unobserved object detection predicts distributions over the categories and locations of occluded or out-of-view objects (Bhattacharjee et al., 2025). Footprints predicts visible and occluded traversable geometry within the camera frustum from one RGB image; it is evaluated indoors and demonstrated in path planning (Watson et al., 2020). FlatLands evaluates metric BEV floor support in occluded and out-of-frustum regions under shared observation and validity masks.

Layouts, occupancy, and BEV representations. Related representations include room layouts (Lee et al., 2017; Zou et al., 2018; Sun et al., 2019), scan-derived floorplans and structured scene descriptions (Liu et al., 2018a; Chen et al., 2019; Avetisyan et al., 2024), and volumetric occupancy (Song et al., 2017; Cao & de Charette, 2022; Zhu et al., 2026). 360BEV maps panoramic RGB-D observations to indoor semantic BEV maps (Teng et al., 2024). Autonomous driving methods construct BEV maps from multiple cameras, sometimes also using previous frames (Philion & Fidler, 2020; Li et al., 2022). FlatLands isolates floor completion by giving all methods the same geometric maps, then tests the same models on maps estimated from RGB images (Section 3).

Generating and scoring multiple completions. Inpainting networks (Liu et al., 2018b; Suvorov et al., 2022; Lugmayr et al., 2022), diffusion models, and flow matching (Ho et al., 2020; Lipman et al., 2023; Tong et al., 2024) can all complete partial maps. Conditional generators produce alternative outcomes (Kohl et al., 2018); ensembles combine independently trained predictors (Lakshminarayanan et al., 2017). A set of completions contains multiple floormaps predicted from the same observation. Energy scores (Gneiting & Raftery, 2007; Rizzo & Székely, 2016) assess their accuracy and diversity jointly, without selecting a map using ground truth.



(a) Constructing a camera-aligned floormap.

(b) An RGB image and its four aligned binary maps.

Figure 2: **From a partial view to a floormap completion task.** (a) Visible floor points are projected onto the floor plane and rasterized in BEV. (b) Each observation provides maps of observed floor F_{obs}, unobserved cells U, complete floor $F^\star$, and valid workspace V; the white marker shows the camera anchor C. The pair (F_{obs}, U) is the input; $F^\star$ and V are used for supervision and scoring.

Navigating with predicted maps. Predicted maps support exploration and navigation through ensemble uncertainty (Georgakis et al., 2022), estimated information gain (Ho et al., 2025; Baek et al., 2025), and multiple plausible layouts (Wang et al., 2025). An occupancy-grid study also examines the value of retaining map uncertainty for route selection (Banfi et al., 2022). SceneSense applies generative priors to 3D occupancy prediction using RGB-D observations and an accumulated map (Reed et al., 2024). FlatLands provides a common single-view floor-support representation across six indoor datasets and includes an evaluation with conditioning derived from a single egocentric RGB image. For controlled comparison, we adapt and train architectures from prior work using shared input and output representations, then assess completion quality and the utility of the predicted maps for point-goal navigation.

3 Benchmark Task and Interface

Task: unobserved floormap completion. A *floormap* is a binary floor-support map on an $H \times W$ camera-aligned ground-plane grid at resolution Δ (m/px), with one denoting floor. The target $F^\star$ supplies labels on cells marked valid by V. The observed and unobserved masks O and $U = \mathbf{1} - O$ have disjoint supports R_O and R_U; unobserved cells are occluded or outside the camera frustum. The map F_{obs} records floor labels on R_O and is zero elsewhere, so a zero denotes observed non-floor only where $O = 1$. The completion target F_{uno} specifies floor labels on valid unobserved cells, selected by $U \odot V$. Multiple layouts can agree with the observed evidence, making completion a conditional inverse problem (Stuart, 2010). The RGB task targets $P(F_{\text{uno}} \mid I)$; the controlled benchmark targets $P(F_{\text{uno}} \mid c)$, where $c = (F_{\text{obs}}, U)$. These distributions can differ because BEV construction discards visual information. A learned conditional distribution produces candidate completions $\hat{F}_{\text{uno}}^{(k)} \sim q_\theta(\cdot \mid F_{\text{obs}}, U)$ for $k = 1, \ldots, K$; for deterministic models q_θ is a point mass and K=1.

Estimating the observed floormap from an egocentric view. The operator Φ estimates metric depth and floor segmentation, $(D, S) := \Phi(I)$, using Depth Pro (Bochkovskiy et al., 2025) for depth and an ADE20K-trained SegFormer (Xie et al., 2021; Zhou et al., 2019) for segmentation. The projection operator Ψ back-projects depth using intrinsics $\mathbf{K}$, aligns the floor plane, and rasterizes the points by top-down orthographic projection, giving $(F_{\text{obs}}, O) := \Psi(D, S; \mathbf{K})$. We use calibrated intrinsics when available and estimate them otherwise. Alignment requires no global pose; the source-geometry track instead derives the same pair offline from scene geometry.

Completion assembly and supervision. After binarization, each prediction $\hat{F}_{\text{uno}}^{(k)}$ is assembled into a full completion as

$$F_{\text{comp}}^{(k)} = F_{\text{obs}} + U \odot \hat{F}_{\text{uno}}^{(k)}. \tag{1}$$

This construction preserves observed labels and fills only unobserved cells, so errors in the estimated observations remain in the completed map. Supervision and completion metrics use $R_{\text{eval}} = U \odot V$ with $|R_{\text{eval}}| > 0$; neither V nor $F^\star$ is given to the completion model (Appendix B.3). RGB completion scores therefore assess valid unobserved cells, excluding errors in the unchanged observed region.

Why not train directly using RGB as input? Direct RGB-to-BEV training combines visual feature extraction and projection into metric BEV (Philion & Fidler, 2020; Li et al., 2022) with prediction of unobserved floor. Synthesizing shared BEV maps from source geometry expands viewpoint coverage beyond the captured images and isolates completion quality. Methods can then be compared under identical geometric evidence, independently of the architecture, input modality, and training data of an RGB perception front end. The RGB-to-BEV track tests the resulting models with image-derived inputs (Section 5).

4 FLATLANDS DATASET

Sources and splits. Table 1 lists six sources and 17,656 scene records, with up to 24 camera poses per scene placed on valid floor; visibility comes from field-of-view and ray-occlusion tests. Curation retains 391,024 of the 423,743 initial observations; the filter $r_{\text{cond}} = |F_{\text{obs}}|/|F^\star| \geq 0.10$ leaves 270,575 (Appendix B.5). Splitting by source scene yields 215,342 training, 26,890 validation, and 28,343 test observations; 16,214 test observations come from held-out ScanNet++ (Table 16, Fig. 9, and Appendices B and B.6). Each observation contains four aligned binary maps (Fig. 2b) with $F_{\text{obs}} \odot U = 0$ and $F_{\text{obs}} \leq V$.

Ground-truth provenance. For five sources, we extract floor labels from reconstructed meshes. For ZInD, we build floormaps from DA^2-*estimated* panoramic depth (Li et al., 2026), scaled using annotated room vertices. Thus 133,096 of the 270,575 filtered observations (49%) have targets scaled from room annotations. The OOD targets come from meshes (Appendix B.1).

Representation of floormaps. All four maps share a floor-aligned egocentric 256×256 BEV grid spanning 10×10 m at 25.6 pixels per meter, with the camera at pixel $(128, 192)$. The median scene floor area is 16.4 m^2 (Appendix B.2 and Figs. 13 and 14). The fraction of floor cells in R_{eval} is 0.802 ± 0.238 over the pooled test split, 0.661 when ID sources are weighted equally, and 0.854 on OOD. All-Floor attains exactly these IoUs under the corresponding averaging conventions (Table 2). Section 6 separately evaluates observations with lower floor prevalence (Fig. 10).

Table 1: **Source datasets in FlatLands.** Obs. (raw) counts observations before quality filtering; Obs. (filtered) counts only observations with $r_{\text{cond}} \geq 0.10$, valid support, observed/target floor agreement, and non-degenerate crops. Curation and crop validation retain 391,024 observations before applying the r_{cond} threshold; split-level counts for this intermediate stage appear in Appendix B.4. ScanNet++ is OOD test-only.

Dataset	Scenes	Obs. (raw)	Obs. (filtered)	Asset Used
ZInD (Cruz et al., 2021)	7,026	168,624	133,096	Panorama-to-PLY point cloud
ARKitScenes (Baruch et al., 2021)	4,803	115,271	40,282	PLY mesh
Matterport3D (Chang et al., 2017)	2,101	50,424	38,004	Room PLY mesh
ScanNet (Dai et al., 2017)	1,508	36,192	24,763	PLY mesh
3RScan (Wald et al., 2019)	1,291	30,984	18,216	OBJ mesh
ScanNet++ (Yeshwanth et al., 2023)	927	22,248	16,214	PLY mesh
FlatLands (Ours)	**17,656**	**423,743**	**270,575**	4× binary 2D maps

5 BENCHMARK PROTOCOL

Baselines. Eleven baselines receive the same inputs (F_{obs}, U) and are scored on R_{eval} (Table 17). Training-free methods fill unobserved cells with zero (All-Obstacle) or one (All-Floor), copy the nearest observed label (NN Propagation), or draw Bernoulli(0.5) labels (Uniform Random).

Learned predictors. U-Net (Ronneberger et al., 2015) and PConv-UNet (Liu et al., 2018b) use masked binary cross-entropy. LaMa (Suvorov et al., 2022) uses masked ℓ_1, adversarial, and feature-matching losses; LaMa-Ensemble combines four independently trained LaMa predictors (Lakshminarayanan et al., 2017). Diffusion predicts noise (Ho et al., 2020), Flow Matching predicts velocity (Lipman et al., 2023; Liu et al., 2023), and FM+XAttn adds conditioning cross-attention (Vaswani et al., 2017). The three generators train with masked mean-squared error. At inference, they use guidance $s = 2.0$, clamp observed labels, and produce $K = 4$ samples. Diffusion uses 50 DDIM steps; Flow Matching and FM+XAttn use 50 and 25 Heun steps, respectively (Song et al., 2021; Süli & Mayers, 2003). Appendix D specifies the objectives and attention resolutions.

Training implementation and inference. All learned completion models train on source-geometry BEV maps from the five ID sources, with a shared schedule on 4×A100 40 GB GPUs. Single-model runs use seed 42; LaMa-Ensemble uses seeds 41–44. The default guidance scale $s = 2.0$ was selected using the training split (Appendix E). Outputs ≥ 0.5 denote floor. On one A100, completion takes 63–81 ms for deterministic models, 136 ms for LaMa-Ensemble, and 1,745–6,661 ms to generate four samples (Appendix E and Table 18).

Reporting protocol. Two tracks differ in how conditioning is constructed. The *source-geometry (canonical) track* uses synthesized BEV maps, as in training (Appendix B.3). Its 28,343 test observations cover held-out scenes of the five ID sources and the held-out source ScanNet++ (OOD; Appendix B.6); results appear in Tables 2 and 3. The *RGB-to-BEV track* estimates metric depth with pretrained Depth Pro and floor segmentation with SegFormer for every frame, then back-projects, aligns the floor, and rasterizes (Appendix C). This test-only track reuses the same completion weights without fine-tuning. Its 1,000 real frames comprise 200 each from held-out test scenes of four ID sources and ScanNet++, excluding ZInD panoramas (Table 5 and Appendix A.2). In either track, ID denotes sources used for training; OOD denotes held-out ScanNet++. Both report floor-class IoU/F1, mismatch rate (UMR), and masked energy score (MES) on the valid unobserved cells R_{eval}, assessing individual predictions and sets of completions separately.

Unobserved-region mismatch rate. The unobserved-region mismatch rate (UMR; lower is better) is the binary error rate on R_{eval}:

$$\text{UMR} = \frac{\text{FP} + \text{FN}}{|R_{\text{eval}}|}, \tag{2}$$

where FP and FN count false-positive and false-negative cells within R_{eval}. For RGB and challenging-subset results, F1 pools TP/FP/FN across observations, whereas IoU and UMR average scores per observation. The reported F1 and IoU therefore cannot be converted into one another.

Masked energy score. For $K > 1$ completions, MES (lower is better) compares them with the target and with one another (Gneiting & Raftery, 2007; Rizzo & Székely, 2016):

$$\widehat{\text{MES}} = \frac{1}{K}\sum_{k=1}^{K} d_{R_{\text{eval}}}\left(F_{\text{comp}}^{(k)}, F^{\star}\right) - \frac{1}{2K(K-1)}\sum_{k\neq\ell} d_{R_{\text{eval}}}\left(F_{\text{comp}}^{(k)}, F_{\text{comp}}^{(\ell)}\right). \tag{3}$$

The first term measures average error against the target; subtracting the second rewards diversity among completions. Distance is $d_{R_{\text{eval}}}(A, B) = 1 - \text{IoU}_{R_{\text{eval}}}(A, B)$: greater floor overlap means smaller distance. This is masked Jaccard distance (Lipkus, 1999); we use $K = 4$ samples (Appendix A.3). For deterministic predictions, the population score reduces to $1 - \text{IoU}_{R_{\text{eval}}}$ (Appendix F), enabling comparison with sets of completions.

Sample selection. Single-output fidelity evaluates one completion. Best-of-K selects the highest-IoU sample using ground truth, giving an oracle upper bound for selection from those candidates. LaMa-Ensemble's single-output score evaluates the mean of its four member probability maps, thresholded at 0.5. IoU_m averages the four IoU scores after thresholding each prediction at 0.5. MES and variance use the same binary completions; variance divides by four (Appendix F).

Pooling, precision, and table conventions. Equal-source ID summaries weight the five datasets equally; observation-pooled summaries weight observations equally. For equal-source summaries, SD is computed across the five source means. These SDs describe variation across sources or observations, not across training runs. Bold marks the best value and ties at the displayed precision within each comparison block; oracle-selected and single-output results are ranked separately. Per-cell variance (Var) summarizes disagreement among samples (Appendices A.1 and F).

Table 2: Fidelity (mean±SD). The upper block evaluates single outputs; the lower evaluates best-of-four predictions selected using ground truth. ID single-output entries are equal-source means (n=5); ID best-of-four and all OOD entries are observation-level, so ID cross-block gains mix selection and weighting. Extended results: Tables 9 and 12.

Method	ID			OOD		
	UMR↓	IoU↑	F1↑	UMR↓	IoU↑	F1↑
Single output: one completion per observation						
All-Obstacle	0.661±0.229	0.000±0.000	0.000±0.000	0.854±0.155	0.000±0.000	0.000±0.000
All-Floor	0.339±0.229	0.661±0.229	0.737±0.210	0.146±0.155	0.854±0.155	0.912±0.116
NN Propagation	0.243±0.078	0.620±0.219	0.708±0.202	0.198±0.156	0.779±0.186	0.861±0.143
Uniform Random	0.501±0.002	0.367±0.104	0.516±0.128	0.502±0.003	0.453±0.057	0.621±0.064
U-Net	**0.122±0.030**	**0.763±0.132**	**0.843±0.098**	**0.114±0.111**	**0.870±0.136**	**0.924±0.094**
PConv-UNet	0.126±0.033	0.758±0.132	0.838±0.098	**0.114±0.113**	**0.870±0.139**	**0.924±0.097**
LaMa	0.209±0.059	0.623±0.161	0.734±0.131	0.283±0.159	0.672±0.190	0.787±0.153
Best-of-four (oracle): four candidates, one selected using the ground truth						
LaMa-Ensemble	0.105±0.099	0.817±0.205	0.882±0.154	0.107±0.094	0.874±0.126	0.927±0.083
Diffusion	**0.100±0.097**	**0.830±0.196**	**0.892±0.143**	**0.094±0.092**	**0.890±0.122**	**0.936±0.082**
Flow Matching	**0.100±0.095**	**0.830±0.192**	**0.892±0.139**	0.102±0.099	0.880±0.129	0.930±0.088
FM+XAttn	**0.100±0.094**	**0.830±0.192**	**0.892±0.139**	0.109±0.092	0.873±0.124	0.927±0.084

6 Results and Discussion

Fidelity. Table 2 separates single-output fidelity from oracle sample selection. A single deterministic prediction is hard to beat on overlap: U-Net leads single-output IoU on ID and ties PConv-UNet on OOD, well above the generators' first samples (Table 8). Sampling closes the gap only with an oracle: on OOD, the best of four samples exceeds U-Net's 0.870 for every generator, with most of the gain coming from the second sample (Table 12). These ground-truth-selected gains measure candidate quality, not a selection rule usable at test time; MES and planning use the full sample set without oracle selection.

Table 3: Evaluating sets of completions, $K = 4$. ID MES and IoU_m average the five source means; Var and all OOD entries are observation-level. MES combines distance to the target with pairwise sample distance; Var averages the cellwise variance of the four binary predictions (divisor four) over R_{eval}. Conventions: Section 5 and Table 10.

Method	ID			OOD		
	MES↓	IoU_m ↑	Var	MES↓	IoU_m ↑	Var
LaMa-Ensemble	0.195±0.083	0.637±0.175	0.076±0.062	0.149±0.093	0.720±0.155	0.084±0.043
Diffusion	0.174±0.088	0.711±0.151	0.044±0.029	0.097±0.102	**0.832±0.124**	0.048±0.026
Flow Matching	0.173±0.083	**0.712±0.150**	0.037±0.029	0.101±0.103	**0.832±0.135**	0.043±0.028
FM+XAttn	**0.161±0.076**	**0.712±0.146**	0.044±0.032	**0.095±0.093**	0.813±0.123	0.061±0.028

Challenging cases. Because All-Floor labels every unobserved cell as floor, its IoU equals floor prevalence, which is why it scores 0.854 on OOD. The challenging subset removes this advantage: on 1,386 cases with little visible floor and a non-floor majority, All-Floor falls to 0.191 while U-Net keeps 0.468, and All-Floor mislabels most evaluated cells (Table 4). Among sampled methods, FM+XAttn leads on MES, mean IoU, and oracle IoU. LaMa-Ensemble's averaged map instead drifts toward All-Floor: members that disagree over broad unobserved regions average above the floor threshold, so its UMR and pooled F1 are worse than those of a single LaMa model.

Evaluating multiple completions. Scored as sets, the generators overtake the deterministic models. A deterministic prediction's MES equals $1 - \mathrm{IoU}$, so U-Net scores 0.237 on ID and 0.130 on OOD; every generator does better on both splits, with FM+XAttn scoring best, whereas LaMa-Ensemble beats U-Net only on ID (Table 3). Removing FM+XAttn's cross-attention changes its scores by only

Table 4: Challenging subset: 1,386 cases with $r_{\text{cond}} \leq 0.20$ and floor prevalence < 0.50 (mean 0.19). UMR and pooled F1 use the first prediction, except for LaMa-Ensemble's mean map (§). IoU evaluates single outputs in the upper block and ground-truth-selected best-of-four predictions in the lower block. Reporting conventions: Section 5. Standard deviations: Table 13; protocol: Appendix A.4.

Method	UMR↓	IoU↑	$F1_{\text{pool}}$ ↑	MES↓	IoU_m ↑	Var
Single output						
All-Floor	0.809	0.191	0.302	–	–	–
NN Propagation	0.337	0.192	0.299	–	–	–
U-Net	**0.111**	**0.468**	**0.609**	–	–	–
PConv-UNet	**0.111**	0.458	0.599	–	–	–
LaMa	0.147	0.317	0.463	–	–	–
Best-of-four (oracle) in the IoU column						
LaMa-Ensemble	0.790§	0.397	0.296§	0.362	0.295	0.174
Diffusion	0.151	0.462	0.558	0.352	0.396	0.055
Flow Matching	0.170	0.464	0.515	0.346	0.394	0.040
FM+XAttn	0.155	**0.479**	0.550	**0.316**	**0.417**	0.046

0.001 (Table 20), and it is the slowest sampler (Table 18). All three generators perform best at the default guidance scale $s = 2.0$ (Table 19).

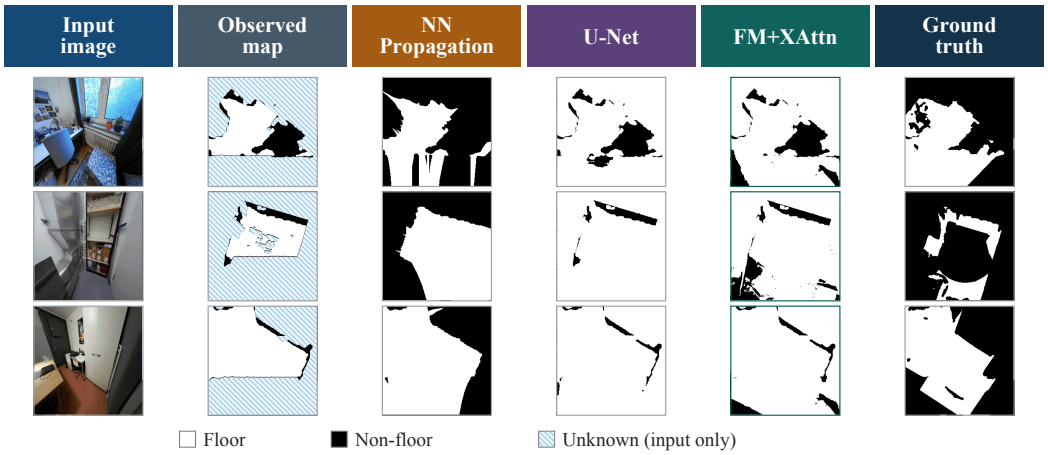


Figure 3: **Three RGB-conditioned benchmark examples.** Each row compares selected baselines on the same input; ground truth at right. FM+XAttn shows its first generated sample. Hatching marks unobserved input cells for display only; the maps remain binary.

Qualitative observations. In Fig. 3, NN Propagation extends visible patches into long bands (top row) and a broad wedge (middle row). The middle observation already marks part of the target's central non-floor region as floor; clamping preserves this error. U-Net and FM+XAttn fill extensive unseen regions, missing the bottom target's angular boundary. FM+XAttn's displayed first draw adds some corner exclusions but still overextends floor. Figure 15 compares all eleven methods on the canonical track; Fig. 18 extends the RGB examples. Appendix A.7 adds a constructed example in which one input has five valid candidate targets.

Boundary-band IoU and interior disagreement. Scoring only cells near floor edges does not escape the prevalence effect: on OOD boundary-band IoU (Cheng et al., 2021), All-Floor (0.820) matches best-of-four Diffusion (0.822) and beats single-output U-Net (0.809). Where samples disagree is more telling: FM+XAttn shows almost no interior disagreement, as its samples rarely differ inside floor regions, unlike LaMa-Ensemble's members (Table 14, Fig. 4, and Appendix A.6).

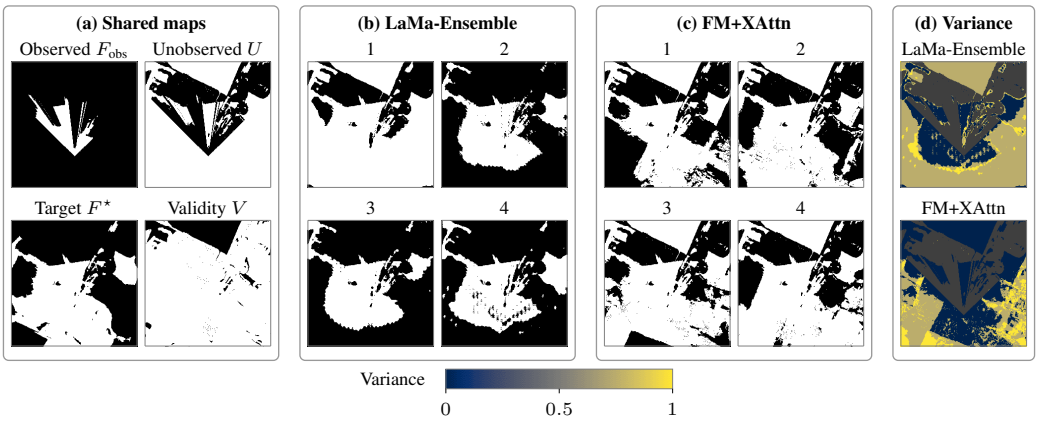


Figure 4: **Four completions and sample disagreement, multi-room ScanNet++ scene.** (a) Conditioning (top) and evaluation maps (bottom). (b,c) Four completions per method. (d) Variance of each method's four binary completions on unobserved cells, with divisor four. Both maps use the same scale: 1 denotes the maximum variance 0.25 (a two-two split).

End-to-end RGB-to-BEV evaluation. The same pattern holds on maps estimated from real images, without fine-tuning (Table 5 and Fig. 3). U-Net again produces the best single map (IoU 0.564), and the best of four generator samples at best roughly matches it, yet every generator scores better as a set, with MES of 0.293–0.308 versus U-Net's 0.436. Diffusion replaces FM+XAttn as the MES leader. Boundary leakage, missing floor, and shape drift persist (Fig. 5).

Table 5: RGB-to-BEV over 1,000 real frames (200 from each of five pinhole sources; ZInD excluded). UMR, IoU_1, and pooled F1 evaluate single outputs: the first sample for generators and the thresholded mean map for LaMa-Ensemble, which can outperform every individual member. IoU_b reports ground-truth-selected best-of-four IoU. IoU_m, MES, and Var use all four outputs, with Var averaged over R_{eval}; All-Obstacle and Uniform Random appear in Table 11. Conventions: Section 5.

Method	UMR↓	IoU_1 ↑	$\mathrm{F1}_{\mathrm{pool}}$ ↑	IoU_b ↑	MES↓	IoU_m ↑	Var
All-Floor	0.559	0.441	0.508	–	–	–	–
NN Propagation	0.467	0.248	0.328	–	–	–	–
U-Net	0.435	**0.564**	**0.650**	–	–	–	–
PConv-UNet	0.437	0.563	0.649	–	–	–	–
LaMa	0.521	0.148	0.223	–	–	–	–
LaMa-Ensemble	0.486	0.439	0.539	0.414	0.302	0.233	0.141
Diffusion	**0.432**	0.497	0.604	0.522	**0.293**	0.504	0.047
Flow Matching	0.449	0.542	0.632	0.564	0.306	0.552	0.038
FM+XAttn	0.440	0.560	0.647	**0.565**	0.308	**0.562**	0.041

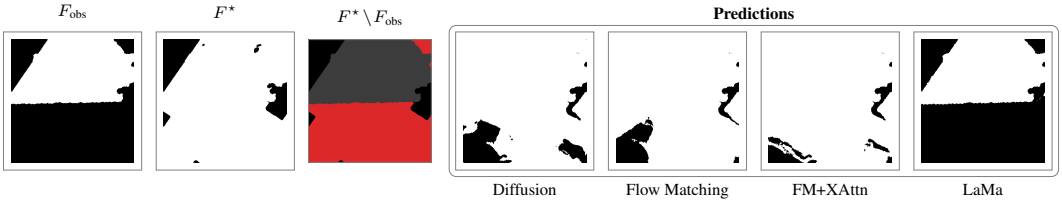


Figure 5: **Failure under wide occlusion.** Red: unseen target floor ($F^\star \setminus F_{\mathrm{obs}}$). Sampled maps cross target boundaries; LaMa predicts almost no floor (displayed IoU 0.003; mean over seeds 41–44: 0.006). Each generator shows its first generated sample; LaMa shows the single model trained with seed 42. More examples: Fig. 17.

Table 6: Offline point-goal planning from RGB-derived BEV conditioning. Found: path returned; Safe Success: route entirely on valid ground-truth floor; SPL also penalizes path length. CIs: 95% Wilson intervals (Wilson, 1927). Bold marks the best learned planner, excluding All-Floor and the oracle. †‡ Each marked pair has identical reported scores.

Planner	Found (%) ↑	CI	Safe Success (%) ↑	CI	SPL ↑
All-Floor	100.00	[99.90, 100.00]	62.68	[61.16, 64.16]	0.610
U-Net / PConv-UNet †	86.85	[85.77, 87.86]	78.20	[76.89, 79.45]	0.780
FM+XAttn mean	90.35	[89.40, 91.23]	79.55	[78.27, 80.77]	0.794
Diffusion / Flow Matching ‡	**96.93**	[96.34, 97.42]	83.05	[81.86, 84.18]	0.800
FM+XAttn risk-weighted	95.73	[95.05, 96.31]	**83.28**	[82.09, 84.40]	**0.829**
Ground-truth oracle	100.00	[99.90, 100.00]	100.00	[99.90, 100.00]	1.000

Offline point-goal planning. We evaluate 4,000 start–goal tasks with feasible ground-truth paths in a separate cohort of 500 RGB-to-BEV test scenes (250 ID, 250 OOD). All planners run Dijkstra search within the valid support. Binary-map planners use unit costs on their maps; risk-weighted planners admit any cell that some sample labels floor and charge more where fewer samples agree (Appendix A.5). Keeping the samples pays off: every risk-weighted generative planner reaches the goal safely more often than the planner using U-Net's map, with non-overlapping confidence intervals (Table 6). The most direct comparison is within FM+XAttn: collapsing its samples into one thresholded map lowers Safe Success from 83.28% to 79.55%, although cell admission and costs both change. All-Floor finds every path, but only 62.68% are safe. The three generators perform similarly, so the benefit is not tied to one model.

7 Conclusion

FlatLands grounds the inverse problem of inferring multiple plausible spatial layouts from partial observations in a concrete robotics task, enabling systematic evaluation. Eleven methods, from training-free rules to generative models, receive identical inputs, masks, and scores across six indoor datasets, a held-out source, and maps estimated from real images. Its answer is that the most accurate single map is not the most useful one. Overlap rewards guessing floor wherever floor is common. Sets of completions, scored with a proper rule and used for risk-aware planning, lead to safer routes than any single prediction. And the way samples disagree is itself informative: a conditional generator such as FM+XAttn is nearly certain inside floor regions, whereas ensemble members are not. These findings hold on the held-out OOD source and on real RGB inputs.

Applications. Because the interface is an abstract binary grid, completion is sensor-agnostic. Completed floormaps can serve as probabilistic spatial priors for belief-space planning (Kurniawati, 2022; Curtis et al., 2024), active perception (Ramakrishnan et al., 2020), SLAM back-ends (Cadena et al., 2016), generative motion planning (Serifi et al., 2024; Bhattacharjee et al., 2026), and planning over multiple map hypotheses (Wang et al., 2025). The natural next step is closed-loop navigation under map uncertainty (Savva et al., 2019; Ho et al., 2025), where a robot acts on its samples and observes the consequences.

Limitations. This benchmark targets *binary* traversability only, without semantic cues, and assumes a simplified pinhole camera model that does not capture the diversity of real robotic cameras. It also omits *graded* traversability (*e.g.* stairs, slopes, *etc.*), which would better reflect the traversability constraints of heterogeneous embodied platforms like mobile, legged, or aerial robots. Failure analysis (Fig. 5) indicates that when most ground-truth floor is unobserved, boundary placement becomes underdetermined; consequently, all methods produce local defects poorly captured by IoU alone.

Future Work. Promising directions include expanding the source datasets to cover more architectural styles and building types, incorporating RGB and semantic cues for conditioning, adopting traversability definitions for different robotic systems, moving beyond single-view input to multi-view settings, and extending baselines to more model architectures (Waseem et al., 2026).

Reproducibility Statement

All learned completion models train on source-geometry BEV maps from the five ID sources. The source-geometry (canonical) track evaluates held-out scenes from these sources and the entirely held-out source ScanNet++ (OOD). The RGB-to-BEV track evaluates real images from held-out test scenes of four ID sources and ScanNet++, using the same completion weights without fine-tuning. Every frame uses depth estimated by Depth Pro and floor segmentation from SegFormer to construct conditioning. Sections 3 and 5 define the shared interface, evaluation tracks, and metrics. Dataset construction and scene splits, model objectives, and training settings appear in Appendices B, D and E, respectively. RGB preprocessing is specified in Appendix C and evaluation-frame selection in Appendix A.2. Appendices A.5 and A.7 specify offline planning and the constructed multi-solution example; Appendix A.1 gives extended comparisons. Code, trained weights, the FlatLands dataset, metadata, and experiment artifacts will be released alongside a public leaderboard. Source assets remain distributed by their providers under the original licenses.

Ethics Statement

FlatLands targets assistive robotics, accessibility mapping, and indoor navigation. RGB imagery and binary floormaps can reveal sensitive spatial information and enable unauthorized mapping, so collection, release, and use require appropriate permissions and privacy safeguards. Source assets are excluded from our release and must be obtained from their providers under the original dataset licenses. Consent, institutional review, and image filtering depend on source protocols; permissions and privacy considerations also apply to derived maps. The emphasis on North American and European interiors may limit generalization across building types, accessibility needs, and floor-prevalence distributions. Physical traversability and closed-loop safety require platform- and user-specific assessment. Compute is approximately 280 A100 GPU-hours for seven learned method families on a shared institutional cluster partly powered by renewable energy.

AI use statement

Large language models assisted with editing prose, developing code, and refining figure presentation. They were not used to formulate the research questions or to develop the methods. Separately, the pipeline itself uses pretrained vision models: DA^2 for ZInD panoramic depth (Section 4), and Depth Pro and SegFormer for depth estimation and floor segmentation in the RGB-to-BEV track (Section 5 and Appendix C). The authors take full responsibility for all claims, results, and artifacts.

References

Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir R. Zamir. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018. doi: 10.48550/arXiv.1807.06757.

Armen Avetisyan, Christopher Xie, Henry Howard-Jenkins, Tsun-Yi Yang, Samir Aroudj, Suvam Patra, Fuyang Zhang, Duncan Frost, Luke Holland, Campbell Orme, Jakob Engel, Edward Miller, Richard Newcombe, and Vasileios Balntas. SceneScript: Reconstructing Scenes with an Autoregressive Structured Language Model. In *Computer Vision – ECCV 2024*, pp. 247–263. Springer Nature Switzerland, 2024. doi: 10.1007/978-3-031-73030-6_14.

Seungjae Baek, Brady Moon, Seungchan Kim, Muqing Cao, Cherie Ho, Sebastian Scherer, and Jeong Hwan Jeon. PIPE Planner: Pathwise Information Gain with Map Predictions for Indoor Robot Exploration. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7684–7691, 2025. doi: 10.1109/iros60139.2025.11246190.

Jacopo Banfi, Lindsey Woo, and Mark Campbell. Is it Worth to Reason about Uncertainty in Occupancy Grid Maps during Path Planning? In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11102–11108, 2022. doi: 10.1109/icra46639.2022.9812431.

H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf. Parametric correspondence and chamfer matching: two new techniques for image matching. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'77, pp. 659–663, San Francisco, CA, USA, 1977. Morgan Kaufmann Publishers Inc. URL https://www.ijcai.org/Proceedings/77-2/Papers/024.pdf.

Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Yuri Feigin, Peter Fu, Thomas Gebauer, Daniel Kurz, Tal Dimry, Brandon Joffe, Arik Schwartz, and Elad Shulman. ARKitScenes: A diverse real-world dataset for 3D indoor scene understanding using mobile RGB-D data. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. URL https://openreview.net/forum?id=tjZjv_qh_CE.

Subhransu S. Bhattacharjee, Dylan Campbell, and Rahul Shome. Believing is Seeing: Unobserved object detection using generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19366–19377, 2025. doi: 10.1109/CVPR52734.2025.01804.

Subhransu S. Bhattacharjee, Hao Lu, Dylan Campbell, and Rahul Shome. MatterDoor: Sampling Zero-shot Spatio-semantic Priors using Generative Models. *arXiv preprint arXiv:2510.11014v2*, 2026. doi: 10.48550/arXiv.2510.11014.

Alexey Bochkovskiy, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan Richter, and Vladlen Koltun. Depth Pro: Sharp Monocular Metric Depth in Less Than a Second. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=aueXfY0Clv.

Mathieu Bouchard, Anne-Laure Jousselme, and Pierre-Emmanuel Doré. A proof for the positive definiteness of the Jaccard index matrix. *International Journal of Approximate Reasoning*, 54(5): 615–626, 2013. doi: 10.1016/j.ijar.2013.01.006.

Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, Jose Neira, Ian Reid, and John J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *Trans. Rob.*, 32(6):1309–1332, December 2016. ISSN 1552-3098. doi: 10.1109/TRO.2016.2624754.

Anh-Quan Cao and Raoul de Charette. MonoScene: Monocular 3D Semantic Scene Completion. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3981–3991, Los Alamitos, CA, USA, June 2022. IEEE Computer Society. doi: 10.1109/CVPR52688.2022.00396.

João Carvalho, An Thai Le, Piotr Kicki, Dorothea Koert, and Jan Peters. Motion Planning Diffusion: Learning and Adapting Robot Motion Planning with Diffusion Models. *IEEE Transactions on Robotics*, 41:4881–4901, 2025. doi: 10.1109/TRO.2025.3593109.

Angel X. Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D Data in Indoor Environments. In *International Conference on 3D Vision (3DV)*, pp. 667–676, 2017. doi: 10.1109/3dv.2017.00081.

Jiacheng Chen, Chen Liu, Jiaye Wu, and Yasutaka Furukawa. Floor-SP: Inverse CAD for Floorplans by Sequential Room-wise Shortest Path. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2661–2670, 2019. doi: 10.1109/ICCV.2019.00275.

Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training Deep Nets with Sublinear Memory Cost. *arXiv preprint arXiv:1604.06174*, 2016. doi: 10.48550/arXiv.1604.06174.

Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog Bits: Generating Discrete Data using Diffusion Models with Self-Conditioning. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/pdf?id=3itjR9QxFw.

Bowen Cheng, Ross Girshick, Piotr Dollár, Alexander C. Berg, and Alexander Kirillov. Boundary IoU: Improving Object-Centric Image Segmentation Evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15329–15337, 2021. doi: 10.1109/cvpr46437.2021.01508.

Lu Chi, Borui Jiang, and Yadong Mu. Fast Fourier convolution. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 4479–4488. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/2fd5d41ec6cfab47e32164d5624269b1-Abstract.html.

Jacob K Christopher, Stephen Baek, and Ferdinando Fioretto. Constrained synthesis with projected diffusion models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=FsdB3I9Y24.

Steve Cruz, Will Hutchcroft, Yuguang Li, Naji Khosravan, Ivaylo Boyadzhiev, and Sing Bing Kang. Zillow Indoor Dataset: Annotated Floor Plans with 360° Panoramas and 3D Room Layouts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2133–2143, 2021. doi: 10.1109/cvpr46437.2021.00217.

Aidan Curtis, George Matheos, Nishad Gothoskar, Vikash Mansinghka, Joshua B. Tenenbaum, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Partially Observable Task and Motion Planning with Uncertainty and Risk Awareness. In *Robotics: Science and Systems (RSS)*, 2024. doi: 10.15607/rss.2024.xx.118.

Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2432–2443, 2017. doi: 10.1109/CVPR.2017.261.

Alberto Elfes. Using Occupancy Grids for Mobile Robot Perception and Navigation. *Computer*, 22 (6):46–57, 1989. doi: 10.1109/2.30720.

Martin A. Fischler and Robert C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24(6):381–395, 1981. doi: 10.1145/358669.358692.

Georgios Georgakis, Bernadette Bucher, Anton Arapin, Karl Schmeckpeper, Nikolai Matni, and Kostas Daniilidis. Uncertainty-driven Planner for Exploration and Navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11295–11302, 2022. doi: 10.1109/icra46639.2022.9812423.

Tilmann Gneiting and Adrian E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Cherie Ho, Seungchan Kim, Brady Moon, Aditya Parandekar, Narek Harutyunyan, Chen Wang, Katia Sycara, Graeme Best, and Sebastian Scherer. MapEx: Indoor Structure Exploration with Probabilistic Information Gain from Global Map Predictions. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 13074–13080, 2025. doi: 10.1109/icra55743.2025.11128862.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Armin Hornung, Kai M. Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees. *Autonomous Robots*, 34(3): 189–206, 2013. doi: 10.1007/s10514-012-9321-0.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99–134, 1998. ISSN 0004-3702. doi: 10.1016/S0004-3702(98)00023-X.

Simon A. A. Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R. Ledsam, Klaus H. Maier-Hein, S. M. Ali Eslami, Danilo Jimenez Rezende, and Olaf Ronneberger. A probabilistic U-Net for segmentation of ambiguous images. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pp. 6965–6975, Red Hook, NY, USA, 2018. Curran Associates Inc. URL https://proceedings.neurips.cc/paper/2018/hash/473447ac58e1cd7e96172575f48dca3b-Abstract.html.

Hanna Kurniawati. Partially Observable Markov Decision Processes and Robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 5(1):253–277, 2022. doi: 10.1146/annurev-control-042920-092451.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 6405–6416, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964. URL https://proceedings.neurips.cc/paper/2017/hash/9ef2ed4b7fd2c810847ffa5fa85bce38-Abstract.html.

Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006. doi: 10.1017/CBO9780511546877.

Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. RoomNet: End-to-End Room Layout Estimation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 4875–4884, Los Alamitos, CA, USA, October 2017. IEEE Computer Society. doi: 10.1109/ICCV.2017.521.

Haodong Li, Wangguandong Zheng, Jing He, Yuhao Liu, Xin Lin, Xin Yang, Ying-Cong Chen, and Chunchao Guo. DA^2: Depth anything in any direction. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=323ximYcsk.

Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. BEVFormer: Learning Bird's-Eye-View Representation from Multi-Camera Images via Spatiotemporal Transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 1–18, 2022. doi: 10.1007/978-3-031-20077-9_1.

Alan H. Lipkus. A proof of the triangle inequality for the Tanimoto distance. *Journal of Mathematical Chemistry*, 26(1):263–265, Oct 1999. ISSN 1572-8897. doi: 10.1023/A:1019154432472.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow Matching Guide and Code. *arXiv preprint arXiv:2412.06264*, 2024. doi: 10.48550/arXiv.2412.06264.

Chen Liu, Jiaye Wu, and Yasutaka Furukawa. FloorNet: A Unified Framework for Floorplan Reconstruction from 3D Scans. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 203–219, 2018a. doi: 10.1007/978-3-030-01231-1_13.

Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image Inpainting for Irregular Holes Using Partial Convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 89–105, 2018b. doi: 10.1007/978-3-030-01252-6_6.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=XVjTT1nw5z.

Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=Skq89Scxx.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting Using Denoising Diffusion Probabilistic Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11451–11461, 2022. doi: 10.1109/cvpr52688.2022.01117.

Steve Macenski, Francisco Martín, Ruffin White, and Jonatan Ginés Clavero. The Marathon 2: A Navigation System. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2718–2725, 2020. doi: 10.1109/IROS45743.2020.9341207.

Eitan Marder-Eppstein, Eric Berger, Tully Foote, Brian Gerkey, and Kurt Konolige. The Office Marathon: Robust Navigation in an Indoor Office Environment. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 300–307, 2010. doi: 10.1109/ROBOT.2010.5509725.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8162–8171. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/nichol21a.html.

Jonah Philion and Sanja Fidler. Lift, Splat, Shoot: Encoding Images from Arbitrary Camera Rigs by Implicitly Unprojecting to 3D. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 194–210, 2020. doi: 10.1007/978-3-030-58568-6_12.

Santhosh K. Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy Anticipation for Efficient Exploration and Navigation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 400–418, 2020. doi: 10.1007/978-3-030-58558-7_24.

Alec Reed, Brendan Crowe, Doncey Albin, Lorin Achey, Bradley Hayes, and Christoffer Heckman. SceneSense: Diffusion Models for 3D Occupancy Synthesis from Partial Observation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7383–7390, 2024. doi: 10.1109/iros58592.2024.10802589.

Maria L. Rizzo and Gábor J. Székely. Energy distance. *WIREs Computational Statistics*, 8(1):27–38, January 2016. ISSN 1939-5108. doi: 10.1002/wics.1375.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10674–10685, 2022. doi: 10.1109/cvpr52688.2022.01042.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 234–241, 2015. doi: 10.1007/978-3-319-24574-4_28.

Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9338–9346, 2019. doi: 10.1109/ICCV.2019.00943.

Agon Serifi, Ruben Grandia, Espen Knoop, Markus Gross, and Moritz Bächer. Robot Motion Diffusion Model: Motion Generation for Robotic Characters. In *SIGGRAPH Asia 2024 Conference Papers (SA '24)*, pp. 1–9, 2024. doi: 10.1145/3680528.3687626.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=St1giarCHLP.

Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas Funkhouser. Semantic Scene Completion from a Single Depth Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 190–198, 2017. doi: 10.1109/CVPR.2017.28.

Andrew M. Stuart. Inverse Problems: A Bayesian Perspective. *Acta Numerica*, 19:451–559, 2010. doi: 10.1017/s0962492910000061.

Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong Chen. HorizonNet: Learning Room Layout with 1D Representation and Pano Stretch Data Augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1047–1056, 2019. doi: 10.1109/CVPR.2019.00114.

Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution–robust large mask inpainting with Fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 3172–3182, 2022. doi: 10.1109/wacv51458.2022.00323.

Endre Süli and David F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003. doi: 10.1017/CBO9780511801181.

Zhifeng Teng, Jiaming Zhang, Kailun Yang, Kunyu Peng, Hao Shi, Simon Reiß, Ke Cao, and Rainer Stiefelhagen. 360BEV: Panoramic semantic mapping for indoor bird's-eye view. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 373–382, 2024. doi: 10.1109/WACV57701.2024.00044.

Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, 2005. ISBN 9780262201629. URL https://mitpress.mit.edu/9780262201629/probabilistic-robotics/.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=CD9Snc73AW.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964. URL https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. RIO: 3D Object Instance Re-Localization in Changing Indoor Environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7657–7666, 2019. doi: 10.1109/ICCV.2019.00775.

Yizhuo Wang, Haodong He, Jingsong Liang, Yuhong Cao, Ritabrata Chakraborty, and Guillaume Adrien Sartoretti. CogniPlan: Uncertainty-guided path planning with conditional generative layout prediction. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 1382–1396. PMLR, 27–30 Sep 2025. URL https://proceedings.mlr.press/v305/wang25d.html.

Ahmad Waseem, Pietro Ruiu, Seth Nixon, Andrea Lagorio, and Massimo Tistarelli. Self-attention as the backbone: A survey on vision transformers. *Computer Vision and Image Understanding*, 269: 104791, 2026. ISSN 1077-3142. doi: 10.1016/j.cviu.2026.104791.

Jamie Watson, Michael Firman, Aron Monszpart, and Gabriel J. Brostow. Footprints and free space from a single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11–20, 2020. doi: 10.1109/CVPR42600.2020.00009.

Edwin B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927. ISSN 01621459, 1537274X. URL http://www.jstor.org/stable/2276774.

Yuxin Wu and Kaiming He. Group Normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018. doi: 10.1007/978-3-030-01261-8_1.

Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 12077–12090. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/64f1f27bf1b4ec22924fd0acb550c235-Abstract.html.

Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. ScanNet++: A High-Fidelity Dataset of 3D Indoor Scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12–22, 2023. doi: 10.1109/iccv51070.2023.00008.

Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ADE20K dataset. *International Journal of Computer Vision*, 127(3):302–321, Mar 2019. ISSN 1573-1405. doi: 10.1007/s11263-018-1140-0.

Hu Zhu, Bohan Li, Xianda Guo, Hongsi Liu, Baorui Peng, Mingqi Yuan, Xin Jin, Wenjun Zeng, and Chang Wen Chen. GEM-Occ: From Visual Geometry Evidence to Embodied Semantic Occupancy Memory. *arXiv:2607.05543*, 2026. doi: 10.48550/arXiv.2607.05543.

Chuhang Zou, Alex Colburn, Qi Shan, and Derek Hoiem. LayoutNet: Reconstructing the 3D Room Layout from a Single RGB Image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2051–2059, 2018. doi: 10.1109/CVPR.2018.00219.

APPENDIX

Reader's guide. The appendix supports each main-paper result with its protocol, extended tables, and qualitative evidence. Table 7 maps each result to its supporting material. Sections are ordered as extended results (Appendix A), dataset construction (Appendix B), the RGB front end (Appendix C), baselines and implementation (Appendices D and E), metrics (Appendix F), and qualitative examples (Appendix G).

Table 7: **Where to find supporting material for each main-paper result.**

Main-paper result	Supporting material
Single-output and best-of-K fidelity (Table 2)	Per-source fidelity, UMR by conditioning ratio, and sample-count sensitivity (Appendices A.1 and A.3)
Sets of completions and MES (Table 3)	Per-source MES (Table 10); MES propriety and estimators (Appendix F)
Challenging subset (Table 4)	Selection criteria and standard deviations (Appendix A.4)
RGB-to-BEV transfer (Table 5)	Cohort design (Appendix A.2); front end (Appendix C); galleries (Fig. 18)
Offline point-goal planning (Table 6)	Task construction, map-dependent costs, and SPL (Appendix A.5)
Interior disagreement (Fig. 4)	Boundary radius, regional variance, and boundary-band IoU (Appendix A.6)
Constructed multi-solution example	Construction and distributional metrics (Appendix A.7)
Dataset, splits, and filtering (Section 4)	Sources, observation synthesis, curation, conditioning tiers, and splits (Appendix B)
Baselines, training, and runtime (Section 5)	Formulations, hyperparameters, runtime, and guidance and cross-attention ablations (Appendices D and E)
Qualitative results and failures (Figs. 3 and 5)	Full canonical comparison, RGB galleries, and failure grid (Appendix G)

A EXTENDED BENCHMARK RESULTS

A.1 EXTENDED COMPARISONS

Table 8 compares single-output and best-of-K IoU; Tables 9 and 10 report source-level fidelity and sample-set scores. Figure 6 stratifies canonical UMR by conditioning ratio.

A.2 END-TO-END EVALUATION DESIGN

The RGB-to-BEV track uses image-derived conditioning from Appendix C and reuses the source-geometry-trained completion weights without fine-tuning. Both tracks share the BEV interface and metrics (Section 5 and Tables 2 and 3). The end-to-end cohort contains $N = 1{,}000$ egocentric frames, with 200 frames from each of ScanNet, ARKitScenes, Matterport3D, 3RScan, and ScanNet++; ZInD panoramas are excluded from this pinhole-camera protocol. Frames come from real captured camera poses rather than synthesized viewpoints. The four ID sources contribute frames only from held-out test scenes, disjoint from completion-model training and validation scenes. ScanNet++ provides held-out-source evaluation within this RGB cohort. The evaluation retains observations satisfying the conditioning-signal threshold $r_{\text{cond}} \geq 0.10$, computed using ground truth. The geometric frame-validity filter in Appendix C applies at test time. Table 11 reports the quantitative results; Fig. 18 shows the associated map structure and boundary errors.

A.3 SAMPLE SIZE AND SENSITIVITY

We fix $K{=}4$ for all stochastic methods (generators and LaMa-Ensemble), balancing estimation quality against compute while matching the ensemble member count. The sensitivity protocol

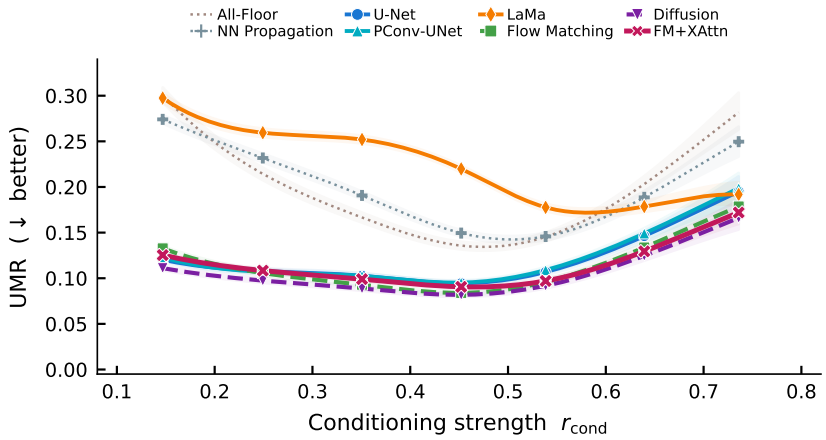


Figure 6: UMR on the canonical track by conditioning ratio. All-Floor UMR equals one minus target-floor prevalence within each conditioning bin. Generator UMR is computed from the first generated completion per observation.

evaluates $K \in \{1, 2, 3, 4\}$ predictions per observation. Table 12 reports best-of-K IoU on both ID and OOD splits. The $K = 1$ column reports single-output fidelity; the larger-K columns report ground-truth-selected candidate quality.

A.4 Structurally Challenging Subset

Floor prevalence measures class balance on R_{eval}; r_{cond} measures the observed fraction of target floor. The *structurally challenging* subset requires both low conditioning signal ($r_{\text{cond}} \leq 0.20$; Learnable tier, Appendix B.5) and low floor prevalence (floor fraction on R_{eval} below 0.50). Thus, at least 80% of target-floor cells are unobserved, and non-floor cells outnumber floor cells in the evaluation region. This yields N=1,386 observations (mean r_{cond}=0.14, floor prevalence 0.19). The subset reuses the canonical-track test predictions.

Table 4 gives the main comparison; Table 13 also reports standard deviations. All-Floor IoU is 0.191, U-Net obtains 0.468, and FM+XAttn has the largest best-of-four IoU (0.479) and lowest MES (0.316) among the sampled methods.

Table 8: **Single-output and best-of-K IoU.** IoU_1 evaluates one output: a generator's first sample or, for LaMa-Ensemble, the mean of its four member probability maps thresholded at 0.5. Best of 2 and IoU_b select the highest-IoU prediction among the first two and four candidates, respectively, using ground truth; ensemble candidates are individual members in seed order $41, 42, 43, 44$. ID IoU_1 (‡) uses equal-source averaging, and all OOD entries use observation-level averaging (Section 5).

	ID‡	OOD		
Method	IoU_1 ↑	IoU_1 ↑	Best of 2 ↑	IoU_b ↑
Deterministic				
U-Net	**0.763**	**0.870**	–	–
PConv-UNet	0.758	**0.870**	–	–
Sampled, K=4				
LaMa-Ensemble	0.661	0.853	**0.867**	0.874
Diffusion	0.656	0.734	0.866	**0.890**
Flow Matching	0.667	0.771	0.862	0.880
FM+XAttn	0.659	0.725	0.842	0.873

Table 9: Per-source single-output fidelity (mean±SD). ID covers the five in-distribution sources; OOD is ScanNet++. Single-model runs use seed 42; LaMa-Ensemble uses seeds 41–44. Each fidelity entry evaluates one output per observation. For LaMa-Ensemble, this is the mean of four member probability maps, thresholded at 0.5 (Section 5). ARKit denotes ARKitScenes; MP3D denotes Matterport3D. RGB rows evaluate a separate cohort of 1,000 ScanNet++ RGB frames; multi-source RGB results appear in Table 11.

Method	Split	Dataset	IoU↑	F1↑
LaMa	ID	3RScan	0.519 ± 0.272	0.640 ± 0.244
LaMa	ID	ARKit	0.420 ± 0.189	0.567 ± 0.187
LaMa	ID	MP3D	0.738 ± 0.197	0.832 ± 0.155
LaMa	ID	ScanNet	0.620 ± 0.198	0.745 ± 0.171
LaMa	ID	ZInD	0.818 ± 0.193	0.884 ± 0.148
LaMa	OOD	ScanNet++	0.672 ± 0.190	0.787 ± 0.153
LaMa-Ensemble	ID	3RScan	0.546 ± 0.372	0.621 ± 0.357
LaMa-Ensemble	ID	ARKit	0.308 ± 0.265	0.414 ± 0.285
LaMa-Ensemble	ID	MP3D	0.829 ± 0.211	0.887 ± 0.177
LaMa-Ensemble	ID	ScanNet	0.765 ± 0.179	0.853 ± 0.134
LaMa-Ensemble	ID	ZInD	0.859 ± 0.197	0.907 ± 0.160
LaMa-Ensemble	OOD	ScanNet++	0.853 ± 0.153	0.911 ± 0.115
LaMa-Ensemble	RGB	ScanNet++ RGB	0.690 ± 0.223	0.792 ± 0.191
Diffusion	ID	3RScan	0.607 ± 0.249	0.722 ± 0.222
Diffusion	ID	ARKit	0.485 ± 0.207	0.626 ± 0.197
Diffusion	ID	MP3D	0.738 ± 0.169	0.837 ± 0.130
Diffusion	ID	ScanNet	0.681 ± 0.157	0.799 ± 0.120
Diffusion	ID	ZInD	0.769 ± 0.165	0.858 ± 0.126
Diffusion	OOD	ScanNet++	0.734 ± 0.148	0.837 ± 0.109
Diffusion	RGB	ScanNet++ RGB	0.721 ± 0.209	0.818 ± 0.169
Flow Matching	ID	3RScan	0.598 ± 0.273	0.707 ± 0.243
Flow Matching	ID	ARKit	0.451 ± 0.209	0.593 ± 0.204
Flow Matching	ID	MP3D	0.770 ± 0.179	0.856 ± 0.136
Flow Matching	ID	ScanNet	0.706 ± 0.161	0.816 ± 0.122
Flow Matching	ID	ZInD	0.810 ± 0.182	0.882 ± 0.137
Flow Matching	OOD	ScanNet++	0.771 ± 0.151	0.862 ± 0.109
Flow Matching	RGB	ScanNet++ RGB	0.704 ± 0.196	0.808 ± 0.162

Table 9: Per-source single-output fidelity (continued). Reporting conventions and cohort definitions follow the first part of this table.

Method	Split	Dataset	IoU↑	F1↑
FM+XAttn	ID	3RScan	0.599 ± 0.252	0.714 ± 0.223
FM+XAttn	ID	ARKit	0.466 ± 0.192	0.612 ± 0.185
FM+XAttn	ID	MP3D	0.746 ± 0.169	0.842 ± 0.128
FM+XAttn	ID	ScanNet	0.675 ± 0.152	0.795 ± 0.118
FM+XAttn	ID	ZInD	0.808 ± 0.179	0.881 ± 0.135
FM+XAttn	OOD	ScanNet++	0.725 ± 0.144	0.832 ± 0.106
FM+XAttn	RGB	ScanNet++ RGB	0.729 ± 0.205	0.824 ± 0.166
U-Net	ID	3RScan	0.694 ± 0.260	0.787 ± 0.215
U-Net	ID	ARKit	0.567 ± 0.209	0.700 ± 0.184
U-Net	ID	MP3D	0.867 ± 0.157	0.919 ± 0.113
U-Net	ID	ScanNet	0.803 ± 0.160	0.880 ± 0.115
U-Net	ID	ZInD	0.885 ± 0.161	0.929 ± 0.119
U-Net	OOD	ScanNet++	0.870 ± 0.136	0.924 ± 0.094
U-Net	RGB	ScanNet++ RGB	0.761 ± 0.219	0.844 ± 0.170
PConv-UNet	ID	3RScan	0.694 ± 0.261	0.787 ± 0.215
PConv-UNet	ID	ARKit	0.560 ± 0.213	0.693 ± 0.187
PConv-UNet	ID	MP3D	0.862 ± 0.164	0.915 ± 0.117
PConv-UNet	ID	ScanNet	0.791 ± 0.171	0.872 ± 0.126
PConv-UNet	ID	ZInD	0.881 ± 0.165	0.926 ± 0.122
PConv-UNet	OOD	ScanNet++	0.870 ± 0.139	0.924 ± 0.097
PConv-UNet	RGB	ScanNet++ RGB	0.766 ± 0.218	0.847 ± 0.170

A.5 Offline Planning Protocol

Task construction and common planner. The evaluation in Section 6 uses 4,000 point-goal tasks from 500 RGB-to-BEV test scenes: 250 ID and 250 held-out-source scenes, with eight tasks per scene. This planning cohort is separate from the 1,000-frame completion cohort in Appendix A.2.

Starts are sampled on observed floor cells, and every planner treats the goal cell as traversable. Valid-support masks constrain all planners. All methods use Dijkstra's algorithm on an 8-connected grid, with map-dependent admissibility and traversal costs defined below; diagonal corner cutting is disallowed. The robot is represented by a point, with no footprint inflation; a route is feasible if it passes only through valid floor cells.

Map-dependent planning costs. The planning comparison measures how completed maps support routes between supplied endpoints. Thresholded and probabilistic occupancy maps provide different ways to use uncertain geometry (Banfi et al., 2022). We combine path length with a clipped negative-log floor-support penalty, defined below.

For Diffusion, Flow Matching, and FM+XAttn, four binary completions define the empirical floor frequency $p_{\text{floor}}(x) = \frac{1}{4}\sum_{k=1}^{4} b_k(x)$. Besides the always-traversable goal cell, their risk-weighted planners admit cells with $p_{\text{floor}}(x) > 0$. Unobserved cells have cost $c(x) = 1 - \log(\max\{p_{\text{floor}}(x), 0.5\})$, and observed floor has $c(x) = 1$. Each allowed move $x \to y$ has positive edge weight $w(x, y) = \|y - x\|_2 c(y)$. Dijkstra therefore minimizes the following objective over admissible paths $\pi = (x_0, \ldots, x_T)$ within valid support:

$$J(\pi) = \sum_{t=1}^{T} \|x_t - x_{t-1}\|_2 \, c(x_t), \tag{A1}$$

with coordinates in grid cells: axial moves have length 1 and diagonals $\sqrt{2}$. The FM+XAttn mean-map comparator admits cells with $p_{\text{floor}} \geq 0.5$ and uses unit cell costs, giving ordinary shortest-path search on the thresholded map.

Table 10: Per-source and pooled sample-set scores (observation-level mean±SD), using all $K = 4$ outputs. ID covers five sources; OOD is ScanNet++. Pooled ID weights observations across all five sources. ID MES and IoU_m in Table 3 average the five source means equally. ARKit denotes ARKitScenes; MP3D denotes Matterport3D.

Method	Split	Dataset	MES↓	IoU_m↑
LaMa-Ensemble	ID	3RScan	0.243 ± 0.148	0.530 ± 0.270
LaMa-Ensemble	ID	ARKit	0.308 ± 0.111	0.398 ± 0.182
LaMa-Ensemble	ID	MP3D	0.132 ± 0.111	0.764 ± 0.179
LaMa-Ensemble	ID	ScanNet	0.188 ± 0.102	0.660 ± 0.159
LaMa-Ensemble	ID	ZInD	0.102 ± 0.116	0.831 ± 0.180
LaMa-Ensemble	ID	Pooled	0.158 ± 0.139	0.713 ± 0.246
LaMa-Ensemble	OOD	ScanNet++	0.149 ± 0.093	0.720 ± 0.155
Diffusion	ID	3RScan	0.215 ± 0.179	0.637 ± 0.276
Diffusion	ID	ARKit	0.306 ± 0.142	0.483 ± 0.211
Diffusion	ID	MP3D	0.107 ± 0.124	0.823 ± 0.160
Diffusion	ID	ScanNet	0.150 ± 0.125	0.763 ± 0.145
Diffusion	ID	ZInD	0.091 ± 0.121	0.849 ± 0.157
Diffusion	ID	Pooled	0.143 ± 0.152	0.763 ± 0.222
Diffusion	OOD	ScanNet++	0.097 ± 0.102	0.832 ± 0.124
Flow Matching	ID	3RScan	0.206 ± 0.161	0.640 ± 0.266
Flow Matching	ID	ARKit	0.298 ± 0.134	0.487 ± 0.206
Flow Matching	ID	MP3D	0.109 ± 0.123	0.824 ± 0.169
Flow Matching	ID	ScanNet	0.156 ± 0.125	0.757 ± 0.157
Flow Matching	ID	ZInD	0.094 ± 0.122	0.853 ± 0.167
Flow Matching	ID	Pooled	0.143 ± 0.148	0.766 ± 0.226
Flow Matching	OOD	ScanNet++	0.101 ± 0.103	0.832 ± 0.135
FM+XAttn	ID	3RScan	0.194 ± 0.156	0.642 ± 0.256
FM+XAttn	ID	ARKit	0.275 ± 0.125	0.495 ± 0.193
FM+XAttn	ID	MP3D	0.101 ± 0.112	0.819 ± 0.154
FM+XAttn	ID	ScanNet	0.146 ± 0.116	0.748 ± 0.144
FM+XAttn	ID	ZInD	0.089 ± 0.115	0.855 ± 0.162
FM+XAttn	ID	Pooled	0.133 ± 0.138	0.766 ± 0.218
FM+XAttn	OOD	ScanNet++	0.095 ± 0.093	0.813 ± 0.123

Table 11: Extended RGB-to-BEV completion results (N=1,000; Table 5). Entries with dispersion show observation-level mean±SD. IoU_1 evaluates a single output: the first sample for generators and the thresholded mean map for LaMa-Ensemble, which can outperform every individual member. IoU_b is ground-truth-selected best-of-four IoU. IoU_m averages four separately computed IoUs; MES and Var use all four outputs, and Var averages the per-cell variance of the four binary outputs over R_{eval}. A dash marks a metric omitted from this table.

Method	IoU_1 ↑	IoU_b ↑	MES↓	IoU_m ↑	Var
All-Obstacle	0.000 ± 0.000	–	–	–	–
All-Floor	0.441 ± 0.228	–	–	–	–
Uniform Random	0.329 ± 0.119	–	–	–	–
NN Propagation	0.248 ± 0.175	–	–	–	–
U-Net	$\mathbf{0.564 \pm 0.195}$	–	–	–	–
PConv-UNet	0.563 ± 0.198	–	–	–	–
LaMa	0.148 ± 0.153	–	–	–	–
LaMa-Ensemble	0.439 ± 0.236	0.414	0.302 ± 0.147	0.233 ± 0.188	0.141 ± 0.063
Diffusion	0.497 ± 0.214	0.522	$\mathbf{0.293 \pm 0.136}$	0.504 ± 0.217	0.047 ± 0.029
Flow Matching	0.542 ± 0.190	0.564	0.306 ± 0.141	0.552 ± 0.192	0.038 ± 0.040
FM+XAttn	0.560 ± 0.177	**0.565**	0.308 ± 0.132	$\mathbf{0.562 \pm 0.185}$	0.041 ± 0.022

Table 12: Sample-count sensitivity ($K \in \{1, 2, 3, 4\}$) on ID and OOD: observation-level mean±SD of best-of-K IoU (↑). The $K = 1$ column gives single-output fidelity, which for LaMa-Ensemble is its mean map. For $K > 1$, ground truth selects the highest-IoU prediction among the first K candidates; LaMa-Ensemble's candidates are individual members in seed order $41, 42, 43, 44$. These scores give an oracle upper bound for selection from those candidates.

Method	Split	K=1	K=2	K=3	K=4
LaMa-Ensemble	ID	0.735 ± 0.304	0.787 ± 0.236	0.806 ± 0.214	0.817 ± 0.205
LaMa-Ensemble	OOD	0.853 ± 0.153	0.867 ± 0.130	0.871 ± 0.127	0.874 ± 0.126
Diffusion	ID	0.699 ± 0.207	0.796 ± 0.220	0.819 ± 0.205	0.830 ± 0.196
Diffusion	OOD	0.734 ± 0.148	0.866 ± 0.122	0.879 ± 0.123	0.890 ± 0.122
Flow Matching	ID	0.722 ± 0.233	0.792 ± 0.223	0.817 ± 0.204	0.830 ± 0.192
Flow Matching	OOD	0.771 ± 0.151	0.862 ± 0.127	0.874 ± 0.127	0.880 ± 0.129
FM+XAttn	ID	0.716 ± 0.222	0.790 ± 0.214	0.818 ± 0.201	0.830 ± 0.192
FM+XAttn	OOD	0.725 ± 0.144	0.842 ± 0.120	0.862 ± 0.121	0.873 ± 0.124

Table 13: Structurally challenging subset (mean±SD, N=1,386): $r_{\text{cond}} \leq 0.20$ and floor prevalence < 0.50 on R_{eval}. IoU evaluates single outputs for $K = 1$ and ground-truth-selected best-of-four predictions for $K = 4$. IoU_m averages four IoUs; MES and Var use all four outputs, and Var averages per-cell binary variance over R_{eval}. UMR and pooled F1 are reported in Table 4, where LaMa-Ensemble uses its thresholded mean map.

Method	K	IoU↑	MES↓	IoU_m ↑	Var
All-Floor	1	0.191 ± 0.168	–	–	–
NN Propagation	1	0.192 ± 0.153	–	–	–
U-Net	1	0.468 ± 0.193	–	–	–
PConv-UNet	1	0.458 ± 0.198	–	–	–
LaMa	1	0.317 ± 0.187	–	–	–
LaMa-Ensemble	4	0.397 ± 0.231	0.362 ± 0.158	0.295 ± 0.216	0.174 ± 0.072
Diffusion	4	0.462 ± 0.198	0.352 ± 0.153	0.396 ± 0.202	0.055 ± 0.034
Flow Matching	4	0.464 ± 0.195	0.346 ± 0.148	0.394 ± 0.199	0.040 ± 0.028
FM+XAttn	4	$\mathbf{0.479 \pm 0.191}$	$\mathbf{0.316 \pm 0.139}$	$\mathbf{0.417 \pm 0.192}$	0.046 ± 0.031

ID = 1,221, OOD = 165.

The binary-map planners (All-Floor, U-Net, and PConv-UNet) admit floor-labeled cells at the same threshold, clamp observed evidence, and use unit cell costs. They share the search algorithm, connectivity, corner-cutting rule, valid-support constraint, and endpoint handling specified above.

Parameter roles. The budget $K = 4$ matches completion evaluation (Appendix A.3), giving floor frequencies $\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$. The mean-map threshold requires at least two votes; positive-probability admission requires one. Clipping at 0.5 caps the added penalty at $\log 2$: one-vote and two-vote cells cost $1 + \log 2$, and unanimous floor costs 1. The unit term charges for path length; the log term penalizes lower predicted floor support. All generators share the unit offset and clipping threshold. The resulting traversal penalty ranks routes without assigning calibrated probabilities of whole-path success. Eight-connected search permits diagonals, while the corner-cutting rule excludes diagonal moves that pass between two blocked cells. Safety is evaluated against ground truth.

Route evaluation. "Found" is the percentage of tasks with a returned path (Table 6). Safe Success requires every visited cell to be valid ground-truth floor. SPL (Anderson et al., 2018) weights success by geometric path efficiency:

$$\mathrm{SPL} = \frac{1}{N} \sum_{i=1}^{N} S_i \frac{\ell_i}{\max\{p_i, \ell_i\}}.$$

Here S_i indicates Safe Success, ℓ_i is the ground-truth shortest-path length, and p_i is the returned-path length; failed tasks contribute zero. Risk-weighted FM+XAttn achieves 3.73 percentage points higher Safe Success and 5.38 percentage points higher Found than its mean-map comparator.

A.6 Boundary Radius and Sample Disagreement

For the boundary-variance analysis, the spatial regions depend on a radius r (in pixels). Given $F^\star$, we extract the 1 px-wide floor edge via morphological erosion (3×3) and subtraction, then dilate with a $(2r+1)\times(2r+1)$ square kernel to form Ω_{bnd} (intersected with U). Unobserved floor cells surviving erosion by the same kernel form Ω_{int}. The boundary band spans the target floor edge and its neighborhood on both sides. Both regions exclude observed cells, so their averages describe variation in the completed part of the map. We use $r = 7$ pixels, corresponding to approximately 0.27 m along each grid axis at the canonical scale. This radius provides 10^3–10^4 boundary cells per scene. The radius sweep gives an ensemble-to-FM+XAttn interior variance ratio of 700–1000 for $r \in \{5, \ldots, 9\}$ and a boundary ratio of 2–4. A wider band ($r = 15$) reduces the observed boundary–interior contrast. The square dilation measures distance in the grid's maximum-coordinate norm.

Interior and boundary disagreement. Binary sample variance uses divisor four (Eq. (A7)). We call the mean per-cell variance over Ω_{int} the *interior disagreement*, and over Ω_{bnd} the *boundary disagreement*. FM+XAttn has interior disagreement 6×10^{-5} and boundary disagreement 0.037; the LaMa-Ensemble values are 0.052 and 0.125. FM+XAttn therefore shows almost no interior disagreement, whereas LaMa-Ensemble's members disagree inside floor interiors and also show greater boundary disagreement (Fig. 4).

Boundary-band and interior IoU. Boundary-sensitive evaluation distinguishes local transitions from floor interiors (Cheng et al., 2021). Table 14 evaluates IoU within the ground-truth boundary band and within the eroded floor interior defined above, each restricted to R_{eval}. For either region Ω, the score is $|F_{\mathrm{comp}} \cap F^\star \cap \Omega| / |(F_{\mathrm{comp}} \cup F^\star) \cap \Omega|$, using the predicted and target floor sets.

The stratified canonical test sample contains $1{,}710$ ID and $1{,}984$ OOD observations (seed 42). Single-output methods use one prediction. For stochastic methods, full-region IoU selects the best of four; the same selected prediction is scored on both regions.

Diffusion has the highest boundary-band IoU on ID (0.751) and OOD (0.822); FM+XAttn scores $0.750/0.804$. LaMa scores $0.605/0.580$. The interior contains only ground-truth floor cells, so All-Floor attains IoU 1.000 there by construction; lower scores indicate omitted floor within that region.

On OOD boundary-band IoU, All-Floor reaches 0.820 against best-of-four Diffusion's 0.822 and single-output U-Net's 0.809. Learned-method interior IoUs range from 0.847 to 0.982 on ID and 0.744 to 0.991 on OOD.

Table 14: Boundary-band and interior IoU on the stratified canonical sample (1,710 ID and 1,984 OOD observations). The boundary band uses radius $r = 7$ pixels. Interior columns report means; boundary columns report mean±SD. The top five methods use single outputs; the bottom four use best-of-four selection by full-region IoU. OOD is ScanNet++. Bold marks the highest boundary-band IoU per split.

	ID		OOD	
Method	Interior↑	Boundary↑	Interior↑	Boundary↑
All-Floor	1.000	0.731 ± 0.171	1.000	0.820 ± 0.131
NN Propagation	0.930	0.604 ± 0.191	0.917	0.639 ± 0.186
U-Net	0.974	0.744 ± 0.154	0.972	0.809 ± 0.124
PConv-UNet	0.971	0.744 ± 0.156	0.975	0.812 ± 0.126
LaMa	0.847	0.605 ± 0.174	0.744	0.580 ± 0.185
LaMa-Ensemble	0.979	0.735 ± 0.167	0.991	0.801 ± 0.140
Diffusion	0.982	$\mathbf{0.751 \pm 0.148}$	0.984	$\mathbf{0.822 \pm 0.121}$
Flow Matching	0.972	0.740 ± 0.158	0.974	0.808 ± 0.129
FM+XAttn	0.975	0.750 ± 0.151	0.964	0.804 ± 0.131

A.7 Constructed Multi-Solution Example

To illustrate structural ambiguity in BEV completion, we construct a single conditioning input with multiple valid ground-truth completions (Fig. 7). We select five observations from different scenes across the dataset whose floormaps overlap in a shared region of the BEV grid. The shared conditioning map $F_{\text{obs}}^{\text{syn}}$ is the intersection of the five observed floormaps, and the maps shown in Fig. 7 are four of the five real ground-truth floormaps, not model outputs.

Construction. Let the five selected observations be indexed by $i \in \{1, \ldots, 5\}$. We combine them into a single test instance:

$$F_{\text{obs}}^{\text{syn}} = \bigcap_i F_{\text{obs}}^{(i)}, \qquad V^{\text{syn}} = \bigcap_i V^{(i)}, \tag{A2}$$

$$U^{\text{syn}} = \Big(\bigcup_i U^{(i)}\Big) \cup \Delta_d, \qquad R_{\text{eval}}^{\text{syn}} = U^{\text{syn}} \odot V^{\text{syn}}, \tag{A3}$$

where Δ_d contains 473 cells that are observed in all five inputs but have conflicting floor labels; promoting them to the unobserved set makes the shared conditioning consistent with every candidate target. The resulting regions contain $|F_{\text{obs}}^{\text{syn}}| = 5{,}998$, $|V^{\text{syn}}| = 11{,}124$, and $|R_{\text{eval}}^{\text{syn}}| = 5{,}126$ cells. We denote the five candidate targets by $\{G^{(j)}\}_{j=1}^{5}$; each is the corresponding floormap restricted to V^{syn}. Pairwise target IoU on $R_{\text{eval}}^{\text{syn}}$ ranges from 0.719 to 0.948, with 268 to 1,438 disagreeing cells per pair, confirming genuine multi-modality rather than annotation noise. By construction, every candidate target matches the shared conditioning on observed valid cells.

Distributional metrics. Given K model samples $\{Y^{(k)}\}_{k=1}^{K}$ and the five candidate targets $\{G^{(j)}\}_{j=1}^{5}$, we define the IoU distance $d(A, B) = 1 - \text{IoU}(A \odot R_{\text{eval}}^{\text{syn}}, B \odot R_{\text{eval}}^{\text{syn}})$ and report symmetric Chamfer distance in IoU space (Barrow et al., 1977):

$$d_{\text{p}\to\text{g}} = \frac{1}{K} \sum_{k=1}^{K} \min_{1 \le j \le 5} d(Y^{(k)}, G^{(j)}),$$

$$d_{\text{g}\to\text{p}} = \frac{1}{5} \sum_{j=1}^{5} \min_{1 \le k \le K} d(G^{(j)}, Y^{(k)}),$$

$$d_{\text{sym}} = \tfrac{1}{2}\left(d_{\text{p}\to\text{g}} + d_{\text{g}\to\text{p}}\right).$$

The first term measures each prediction's agreement with its nearest candidate target; the second measures how well the prediction set represents each target. All distances use the evaluation mask $R_{\text{eval}}^{\text{syn}}$.

Table 15: Distributional evaluation on the constructed multi-solution example. d_{sym}: symmetric Chamfer distance in IoU space (↓); $d_{\text{p}\to\text{g}}$: prediction-to-target distance, a precision measure (↓); $d_{\text{g}\to\text{p}}$: target-to-prediction distance, a recall measure (↓); coverage (↑); diversity (descriptive). Bold marks the lowest symmetric distance and highest coverage among stochastic methods.

Method	K	$d_{\text{p}\to\text{g}}$	$d_{\text{g}\to\text{p}}$	d_{sym} ↓	Cov. ↑	Div.
All-Obstacle	1	1.000	1.000	1.000	0.0	—
All-Floor	1	0.014	0.087	0.051	0.8	—
NN Propagation	1	0.013	0.088	0.051	0.8	—
U-Net	1	0.014	0.087	0.051	0.8	—
PConv-UNet	1	0.014	0.087	0.051	0.8	—
LaMa	1	0.147	0.279	0.213	0.0	—
Uniform Random	1	0.511	0.528	0.519	0.0	—
LaMa-Ensemble	4	0.119	0.099	0.109	0.4	0.188
Diffusion	4	0.021	0.087	**0.054**	**0.8**	0.013
Flow Matching	4	0.020	0.087	**0.054**	**0.8**	0.012
FM+XAttn	4	0.051	0.087	0.069	**0.8**	0.073

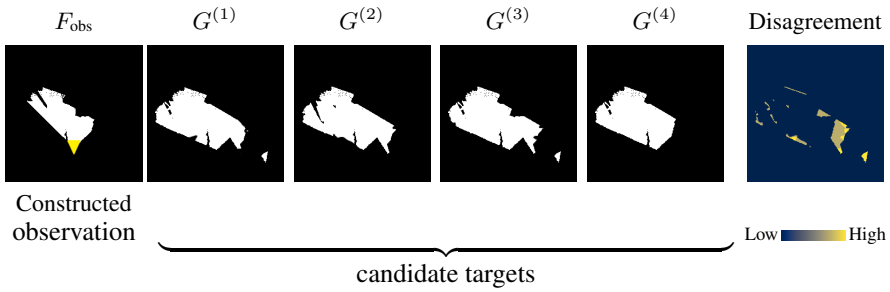


Figure 7: **Constructed multi-solution example.** A shared partial BEV observation F_{obs} is formed as the intersection of five observations from different scenes across the dataset; four of the five real ground-truth floormaps $G^{(j)}$ are shown. White denotes floor and black denotes non-floor. The yellow marker (▼) marks the camera. The final panel shows relative per-cell disagreement among the candidate maps; its color scale is qualitative.

Coverage is the fraction of candidate targets matched by at least one prediction at distance strictly below 0.1:

$$\text{coverage} = \frac{1}{5}\sum_{j=1}^{5} \mathbf{1}\left\{ \min_{1\le k\le K} d(G^{(j)}, Y^{(k)}) < 0.1 \right\}.$$

Diversity averages pairwise distances between distinct predictions:

$$\text{diversity} = \frac{2}{K(K-1)} \sum_{1\le k<\ell\le K} d(Y^{(k)}, Y^{(\ell)}), \qquad K > 1.$$

Diversity applies to $K > 1$; single-prediction entries use a dash. One prediction may match several targets, so coverage need not count distinct recovered modes.

Results. Table 15 summarizes performance on this multi-solution instance. All-Floor, NN Propagation, U-Net, and PConv-UNet obtain $d_{\text{sym}} = 0.051$ with coverage 0.8, matching four of the five candidate targets; U-Net and PConv-UNet predict floor on every evaluated cell here, so their rows equal All-Floor's. Because all five targets are mostly floor, a single all-floor map already lies within 0.1 of four of them, so coverage on this instance cannot separate sample sets from single maps. The missed target is the most distinct one (pairwise IoU as low as 0.719). LaMa deviates substantially ($d_{\text{sym}} = 0.213$, coverage 0.0), matching no target under the 0.1 threshold.

Among the stochastic methods, FM+XAttn has the highest diversity (0.073) of those reaching coverage 0.8. Diffusion and Flow Matching achieve slightly lower d_{sym} (0.054 versus 0.069) because

their samples cluster near the dominant mode ($d_{\mathrm{p}\to\mathrm{g}} = 0.020$–$0.021$), but their diversity is near zero (0.012–0.013): their samples are almost identical.

LaMa-Ensemble has high raw diversity (0.188) but lower coverage (0.4): its independently trained members produce disparate maps, some of which fall outside the match threshold, reflecting seed diversity rather than conditional sampling.

B Dataset Construction and Curation

B.1 Source Datasets

This section extends the construction summary in Section 4. Six indoor datasets are harmonized into a shared metric BEV representation (Table 16 and Fig. 8). For sources with semantic floor annotations (ScanNet, ScanNet++, Matterport3D, 3RScan), floor geometry is extracted directly; for the remainder (ARKitScenes, ZInD), the floor plane is estimated from the 15th percentile of vertex heights (± 0.05 m tolerance).

For ZInD, DA^2 (Li et al., 2026) supplies panoramic depth, which is rescaled to match the human-annotated room layout vertices in ZInD's metric floor-plan frame before point-cloud floor extraction. The annotations determine the metric extent of each target floormap; the depth model supplies within-room structure beyond the layout vertices, including geometry that occludes floor.

Table 16: Source datasets in the FlatLands benchmark. Upstream counts are before filtering. Per-dataset result breakdowns appear in Tables 9 and 10.

Dataset	Scenes	Obs.	Format	Floor source	Axis transform
3RScan	1,291	30,984	OBJ mesh	Semantic labels	Flip $z \leftarrow -z$
ScanNet	1,508	36,192	PLY mesh	Semantic labels	Identity
ARKitScenes	4,803	115,271	PLY mesh	Height percentile	Swap Y/Z
Matterport3D	2,101	50,424	PLY mesh	Semantic labels	Identity
ZInD	7,026	168,624	PLY points	Height percentile	Metric calibration
ScanNet++ (OOD)	927	22,248	PLY mesh	Semantic labels	Identity
Total	**17,656**	**423,743**			

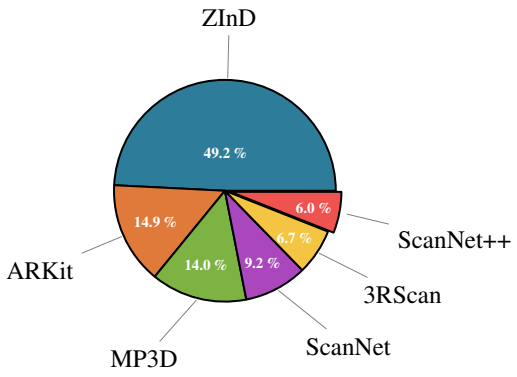


Figure 8: Breakdown of the 270,575 filtered observations across the six indoor source datasets. Source counts and assets are listed in Table 1; construction appears in Fig. 2a.

B.2 BEV Window and Coordinates

Geometry is expressed in an egocentric coordinate frame aligned with the floor plane. The model grid is 256×256 at 25.6 pixels per meter ($\Delta = 0.0390625$ m per pixel), covering a 10×10 m window.

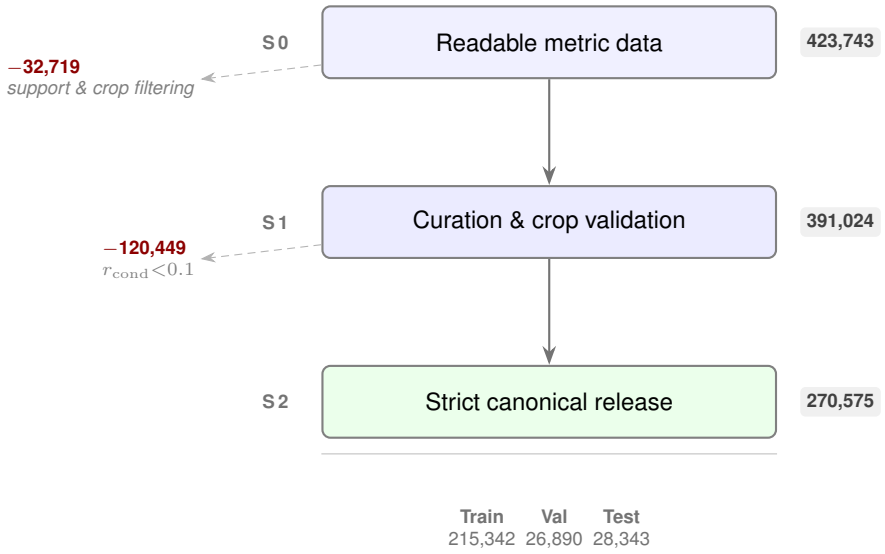


Figure 9: **Dataset curation pipeline.** Each stage applies deterministic quality filters, retaining observations from all six source datasets. Dashed branches indicate removed observations. The strict $r_{\text{cond}} \geq 0.1$ gate accounts for 79% of all removals. Crop validation excludes one observation.

The camera is located at pixel $(128, 192)$, with forward corresponding to decreasing row index. The window therefore extends $5\,\text{m}$ to either side, $7.5\,\text{m}$ forward, and $2.5\,\text{m}$ behind the camera. The observed floor, complete floor target, unobserved-region mask, and validity mask share this spatial alignment. Maps are downsampled from the 512×512 synthesis canvas by average pooling and then binarized.

B.3 Observation Synthesis

Observation synthesis comprises five processing steps:

1. Floor extraction (semantic labels or 15th-percentile height; see Appendix B.1).
2. Removal of geometry above $z_{\text{floor}} + 1.25\,\text{m}$ (robot ceiling).
3. Camera sampling on floor-valid positions (up to 24 observations per scene).
4. Rasterization at 512×512 for four binary channels.
5. Visibility reasoning using field-of-view and ray-occlusion tests.

Camera placement and parameters. Each scene contributes up to 24 observations. Positions are sampled via a spatial-coverage algorithm on a $0.4\,\text{m}$ floor grid with center-biased weighting; headings are drawn uniformly from 36 discrete angles at 10° increments.

The virtual sensor is a 90°-HFOV pinhole camera at $h{=}1.25\,\text{m}$ above z_{floor} with a frame of 640×480. The horizontal field of view matches the Habitat navigation setup (Savva et al., 2019); the height and frame size are specific to our pipeline.

In the synthesis-stage 512×512 BEV canvas, the camera sits at pixel $(256, 384)$, horizontally centered and three quarters down the canvas, with the forward direction along decreasing row index.

Candidates are accepted only when valid cells cover at least 10% of the canvas, the target contains at least 100 floor pixels, and at least 50 floor pixels are observed. Up to 500 candidates are attempted per scene. These synthesis-stage coordinates specify visibility calculations; Appendix B.2 gives the final model grid, metric window, and camera anchor.

Multi-channel rasterization. Each accepted observation produces four aligned binary 512×512 synthesis maps:

- **Floormap** ($F^\star$): complete floor support rasterized in BEV via polygon fill (mesh datasets) or point splatting with morphological closing (ZInD point clouds).
- **Validity mask** (V): cells with defined scene support, including floor and non-floor labels, on which supervision and evaluation are valid. Visibility is represented separately by the observation and unobserved masks.
- **Observed floor** (F_{obs}): visibility-aware floor pixels, computed via z-buffer ray casting (256 samples per ray) that checks for occlusions from walls and furniture.
- **Unobserved** (U): cells where the camera has no direct line of sight, because they are occluded or lie outside the FOV. Scoring intersects this mask with V.

B.4 Filtering, Crop Validation, and Label Balance

Figure 9 summarizes the three-stage curation pipeline, which retains observations from all six source datasets. Of the 423,743 synthesized observations, 391,024 survive curation and crop validation (train/val/test: 312,817/39,099/39,108), and 270,575 pass the strict canonical filter ($r_{\text{cond}} \geq 0.10$; train/val/test: 215,342/26,890/28,343).

Observation validity. Accepted observations have compatible masks, observed floor contained in target floor ($F_{\text{obs}} \subseteq F^\star$), nonempty evaluation support ($|R_{\text{eval}}| > 0$), and non-degenerate crops.

Label balance. Floor prevalence on R_{eval}, pooled across all $N = 28{,}343$ test observations, is 0.802 ± 0.238 (mean $\pm$ standard deviation; Fig. 10). This floor dominance motivates separate reporting on subsets with lower floor prevalence (Appendix A.4). Table 4 gives the corresponding fidelity and sample-set comparisons.

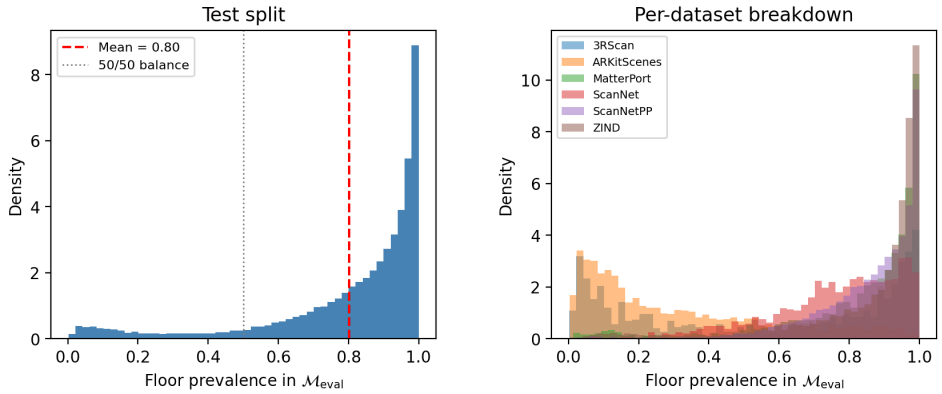


Figure 10: Observation-pooled floor prevalence on R_{eval} across the test observations (Appendix B.4). **Left:** pooled prevalence density with the mean marked by a dashed red line. **Right:** per-source prevalence distributions.

B.5 Conditioning Signal and Learnability

The *conditioning signal ratio* $r_{\text{cond}} = |F_{\text{obs}}|/|F^\star|$ measures the fraction of target floor cells observed, for $|F^\star| > 0$.

We use three descriptive conditioning-signal tiers on the full N=423,743 upstream observations:

- **Easy** ($r_{\text{cond}} > 0.20$): strong conditioning signal. 158,435 observations (37.4%).
- **Learnable** ($0.02 \leq r_{\text{cond}} \leq 0.20$): moderate signal with substantial unobserved structure. 201,455 observations (47.5%).
- **Negligible** ($r_{\text{cond}} < 0.02$): negligible conditioning signal. 63,853 observations (15.1%).

The combined **Easy+Learnable** set comprises 84.9% of observations. Figure 11 shows tier proportions by source.

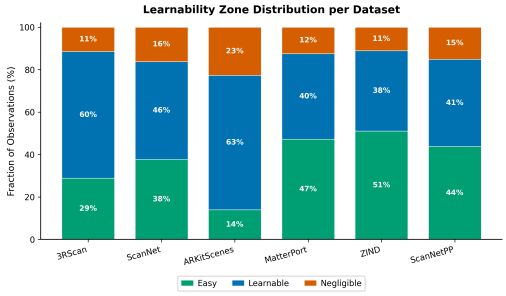


Figure 11: Conditioning-signal tiers across the upstream corpus. Three tiers are defined by r_{cond}: **Easy** (> 0.20), **Learnable** ($0.02 \leq r_{\text{cond}} \leq 0.20$), and **Negligible** ($< 0.02$).

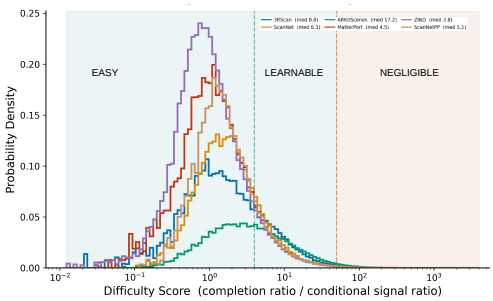


Figure 12: Difficulty-score distribution on a log scale, where $D = (1-r_{\text{cond}})/r_{\text{cond}}$. Tier boundaries are $D < 4$, $4 \leq D \leq 49$, and $D > 49$. The logarithmic plot labels the upper Learnable cutoff as approximately 50. The canonical filter retains $D \leq 9$.

Difficulty score. A scalar difficulty $D = (1 - r_{\text{cond}})/r_{\text{cond}}$, defined for $r_{\text{cond}} > 0$, gives the tier boundaries $D < 4$ (Easy), $4 \leq D \leq 49$ (Learnable), and $D > 49$ (Negligible). Figure 12 shows the distribution. The canonical filter $r_{\text{cond}} \geq 0.1$ retains only $D \leq 9$.

Threshold selection. The minimum conditioning signal $r_{\text{cond}} \geq \tau$ is one of several curation thresholds. We use $\tau = 0.10$, retaining 270,575 observations: 69.2% of the 391,024 post-curation observations and 63.9% of the 423,743 upstream observations.

The filter removes 120,449 post-curation observations (30.8%), all with less than 10% of their target floor observed. Requiring $r_{\text{cond}} > 0.20$ would keep only Easy-tier observations, of which there are 158,435 upstream (37.4%, versus 63.9% retained at $\tau = 0.10$).

Manual inspection found near-degenerate low-signal views, including floor visible through narrow furniture gaps. The threshold balances conditioning strength and retained viewpoint diversity, including observations with limited visible floor. The ratio uses the complete target $F^\star$ for dataset selection and evaluation.

B.6 Split Construction and Stratification

Viewpoints are grouped by source scene to prevent viewpoint leakage. The five ID sources use scene-level train/validation/test partitions; ScanNet++ supplies the held-out-source evaluation. Canonical counts are 215,342 training, 26,890 validation, and 28,343 test observations, with 12,129 ID and 16,214 OOD test observations.

Before the strict canonical filter, each split's Easy/Learnable/Negligible proportions are within one percentage point of $37.4/47.5/15.1\%$, and the median within-scene standard deviation of r_{cond} is 0.08, reflecting viewpoint variation within scenes. The canonical 0.10 filter excludes the Negligible tier. Across 17,656 scene records, the median floor area is $16.4\,\text{m}^2$; per-viewpoint floor coverage, the share of the BEV window that is target floor, ranges from 56% for Matterport3D to 65% for ScanNet (Figs. 13 and 14).

C End-to-End RGB-to-BEV Pipeline

Test-time processing. Figure 1 summarizes the front end and subsequent conditional completion. For each evaluated frame, Depth Pro (Bochkovskiy et al., 2025) estimates metric depth from one egocentric RGB image. We then back-project via the standard pinhole model. Intrinsics come from each source's camera calibration; Depth Pro focal estimates ($f_x{=}f_y$, image-center principal point) are the fallback for uncalibrated frames. The raw point cloud is voxel-downsampled (0.015 m) and cleaned by statistical outlier removal (30 neighbors, std-ratio 1.5). Floor pixels are predicted with SegFormer (Xie et al., 2021); the floor plane is estimated with RANSAC (Fischler & Bolles, 1981) (1000 trials, 5 cm threshold) followed by SVD normal estimation on the inlier set.

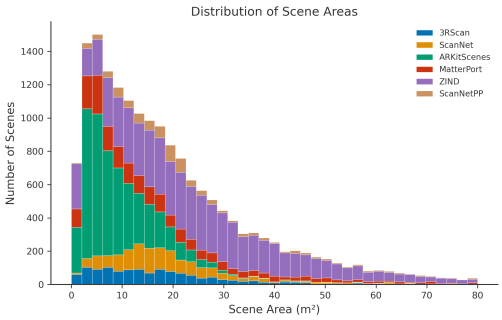


Figure 13: Floor areas across 17,656 scene records (median 16.4 m^2; Appendix B.6).

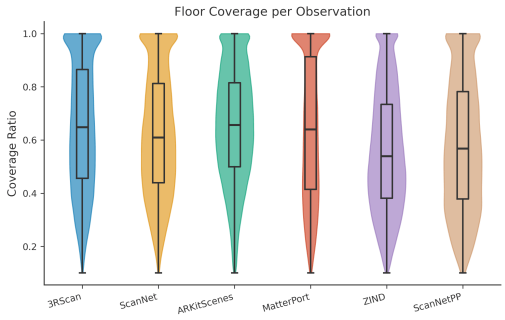


Figure 14: Per-viewpoint floor coverage (share of the BEV window that is target floor) across the six source datasets (Appendix B.6).

BEV conditioning. Rigid alignment places the estimated floor at $y = 0$, with positive y pointing downward. Height filtering retains points with $y \geq -1.25$ m, removing geometry more than 1.25 m above the floor. The retained geometry is rasterized into the common camera-aligned BEV frame in Appendix B.2, yielding the conditioning interface in Section 3. The complete floor target and valid evaluation support come from the source scene and are used only for scoring.

Frame validity filter. Frames are rejected when floor-normal estimation fails (floor pixels < 100, floor coverage $< 5\%$, inlier ratio $< 60\%$, or geometric degeneracy). The end-to-end qualitative examples (Fig. 3) comprise ScanNet++ test-split frames passing this filter. The quantitative cohort contains 1,000 frames across five physical-camera sources; Appendix A.2 specifies source composition and the additional ground-truth conditioning-signal threshold used for evaluation.

D Baseline Details

All learned methods receive the same conditioning pair (F_{obs}, U) concatenated channel-wise. Cellwise reconstruction, noise-prediction, and velocity-prediction losses use $R_{\text{eval}} = U \odot V$; auxiliary losses are specified below. No data augmentation is applied. We adapt the U-Net, PConv-UNet, and LaMa backbones to binary BEV inputs and outputs; their losses are described below. Analog Bits (Chen et al., 2023) provides a precedent for converting continuous generative outputs into discrete bits; the benchmark uses a common binarization threshold of 0.5.

U-Net and PConv-UNet. Both models are trained with masked binary cross-entropy on R_{eval}. At inference, outputs are binarized at threshold 0.5 and hard-clamped to preserve observed evidence.

PConv-UNet additionally propagates a binary partial-convolution mask through its layers using the update rule of Liu et al. (2018b).

LaMa and LaMa-Ensemble. We adapt LaMa (Suvorov et al., 2022) from 3-channel RGB inpainting to 1-channel binary completion: the generator accepts 2 input channels (F_{obs}, U) and produces a single-channel output passed through a sigmoid activation. LaMa-Ensemble trains four copies with different random seeds, producing one output per independently trained member (K=4 total). Ensembles also appear in map-completion exploration (Ho et al., 2025).

Diffusion. We adopt a DDPM (Ho et al., 2020) in pixel space with a cosine variance schedule (Nichol & Dhariwal, 2021) (T=1000). A U-Net backbone ϵ_θ predicts the added noise given the noisy sample x_t, timestep t, and conditioning $c = [F_{\text{obs}}, U]$ (concatenated channel-wise). The backbone is trained with masked MSE on R_{eval}. At inference, classifier-free guidance takes the form

$$\tilde{\epsilon}_\theta = (1{+}s)\,\epsilon_\theta(x_t, t, c) - s\,\epsilon_\theta(x_t, t, \varnothing),$$

where $\varnothing$ denotes the null input, obtained by zeroing the conditioning channels as in conditioning dropout during training. At each denoising step, we project onto the observed-label constraints by restoring the observed cells; projected diffusion provides a general precedent (Christopher et al., 2024). The final assembly also enforces agreement with the supplied observed evidence.

Table 17: *Baseline catalog.* Unless noted, learned methods share inputs (F_{obs}, U), supervision mask R_{eval}, and hard evidence clamping before evaluation.

Method	*Adaptations*	*Inference*
Training-free baselines	All-Obstacle (fill 0), All-Floor (fill 1), NN Propagation (nearest-observed copy), Uniform Random (p=0.5). Parameter-free rules on U; merged with observed evidence via hard clamping (Eq. (1)).	Single pass
U-Net (Ronneberger et al., 2015; Wu & He, 2018)	Masked BCE on R_{eval}; binary BEV input/output.	Thr. 0.5 + clamp
PConv-UNet (Liu et al., 2018b)	Masked BCE on R_{eval}. Partial-conv. mask propagation; adapted to binary BEV completion.	Thr. 0.5 + clamp
LaMa (Chi et al., 2020; Suvorov et al., 2022)	Masked ℓ_1, adversarial, and feature-matching losses; Fourier-conv. generator adapted to binary BEV.	Thr. 0.5 + clamp
LaMa-Ensemble (Lakshminarayanan et al., 2017)	Same as LaMa (4 seeds); one prediction per independently trained member.	K=4
Diffusion (Ho et al., 2020; Song et al., 2021)	Masked noise-prediction MSE on R_{eval}; per-step evidence clamping (Christopher et al., 2024).	DDIM, 50 st., CFG s=2
Flow Matching (Lipman et al., 2023)	Masked velocity-prediction MSE on R_{eval}. Same conditioning/clamping as Diffusion; ODE transport parameterization.	Heun, 50 st., CFG s=2
FM+XAttn (Lipman et al., 2023; Vaswani et al., 2017)	Masked velocity-prediction MSE on R_{eval}. Adds condition encoder on $[F_{\text{obs}}, U]$ with cross-attn. (64×64, 32×32); see (Rombach et al., 2022).	Heun, 25 st., CFG s=2

Flow Matching. Flow Matching (Lipman et al., 2023; Liu et al., 2023; Lipman et al., 2024) learns a velocity field $v_\theta(x_t, t, c)$ along the linear path $x_t = (1-t)\, x_0 + t\, F^\star$ for $t \in [0, 1]$ and $x_0 \sim \mathcal{N}(0, I)$, trained with masked MSE on R_{eval} against the target velocity $u_t = F^\star - x_0$. Conditioning is identical to that of Diffusion (channel-wise concatenation of $[F_{\text{obs}}, U]$).

FM+XAttn (Flow Matching with Cross-Attention). FM+XAttn augments the Flow Matching backbone with an auxiliary condition encoder that processes $[F_{\text{obs}}, U]$ into a sequence of spatial tokens. Cross-attention layers (Vaswani et al., 2017) inject these tokens at the coarse resolutions (64×64 and 32×32) of the U-Net, following cross-attention conditioning in generative models (Rombach et al., 2022). This supplies spatial condition features to the velocity predictor in addition to channel concatenation. The cross-attention ablation in Table 20 compares completion scores with and without these blocks. The training loss is identical to that of Flow Matching (masked MSE on R_{eval}); the architectural additions are the condition encoder and its cross-attention pathway.

E Implementation Details

Hyperparameters. All single-model runs use seed 42; LaMa-Ensemble uses seeds 41–44. For LaMa and LaMa-Ensemble, the VGG and ResNet perceptual losses are disabled (weight=0), because those features were learned from three-channel natural images rather than single-channel binary maps. Their objectives combine masked ℓ_1 reconstruction, hinge adversarial loss, and multi-scale feature matching, with weights λ_{rec}=10, λ_{adv}=10, and λ_{fm}=250 (Table 18). Training runs on 4×A100 40 GB GPUs with distributed data-parallel training, FP32 precision, and activation checkpointing (Chen et al., 2016); peak memory is 30–36 GB per GPU, and training all seven learned method families took approximately 280 A100 GPU-hours. At inference, we integrate the learned ODEs used in Flow Matching models from t=0 to t=1 using the second-order Heun solver (Süli & Mayers, 2003). We use DDIM (Song et al., 2021) with 50 steps for the diffusion model. Classifier-free guidance (CFG) (Ho & Salimans, 2021) at scale s=2.0, conditioning dropout probability 0.1, and per-step

evidence clamping apply to Diffusion, Flow Matching, and FM+XAttn. LaMa-Ensemble produces one prediction per trained member without CFG.

Runtime and sampling cost. Deterministic models use one forward pass per observation; LaMa-Ensemble uses four member predictions. Diffusion and Flow Matching use 50 solver steps per sample; FM+XAttn uses 25. Across $K = 4$ outputs, the total solver-update counts are 200, 200, and 100. Each Heun update evaluates the velocity at two stages, and CFG combines conditional and unconditional network predictions at each stage. Per-observation latencies are 63–81 ms for deterministic models and 1,745–6,661 ms for the four-sample generators (Table 18). FM+XAttn uses fewer solver updates and has the largest latency in this comparison. Model checkpoints are selected by peak validation IoU, computed from counts pooled over all validation observations. Fixed random seeds specify the sampled noise.

Table 18: **Training configuration and inference runtime. Top:** Optimization settings; model-specific entries are marked. **Bottom:** Per-observation latency and throughput on a single A100 (256×256 inputs); generative rows use $K = 4$. Latencies include a fixed per-observation overhead, so LaMa-Ensemble's four member passes take less than four times LaMa's latency. *Updates* counts forward passes for feedforward methods and solver steps summed over outputs for generators; network function evaluations also depend on the solver and guidance.

Training hyperparameters	
Optimizer	AdamW (Loshchilov & Hutter, 2019)
Learning rate	1×10^{-4}
Weight decay	1×10^{-2}
LR schedule	Cosine annealing (Loshchilov & Hutter, 2017) (no restarts)
Effective batch size	64 (16/GPU × 4 GPUs)
Training steps	300,000
Precision	FP32
Gradient accumulation	1 step
BCE positive weight	1.0 (U-Net and PConv-UNet only)
CFG dropout p_{drop}	0.1 (Diffusion and flow models only)
LaMa / LaMa-Ensemble additional	
λ_{rec} / λ_{adv} / λ_{fm}	10.0 / 10.0 / 250.0
Seeds (LaMa-Ensemble)	41, 42, 43, 44

Inference runtime

Method	Sampler	Updates	Latency (ms)	FPS
U-Net	Forward	1	81	12.3
PConv-UNet	Forward	1	68	14.7
LaMa	Forward	1	**63**	**15.9**
LaMa-Ensemble	Forward	4	136	7.4
Diffusion	DDIM	200	1745	0.6
Flow Matching	Heun	200	3377	0.3
FM+XAttn	Heun	100	6661	0.2

Guidance and conditioning ablations.

The default guidance scale $s = 2.0$ was selected on the training split. Table 19 evaluates post-hoc OOD sensitivity at five scales from 0.5 to 5.0, holding all other settings fixed. All three generators reach their highest tested OOD best-of-four IoU at $s = 2.0$.

The architecture ablation compares FM+XAttn with and without cross-attention at fixed $s = 2.0$ and $K = 4$ (Table 20). The ablated variant retrains FM+XAttn without its cross-attention blocks, with all other settings fixed. Removing the blocks changes OOD MES from 0.095 to 0.096 and best-of-four IoU from 0.873 to 0.872.

Table 19: Post-hoc guidance-scale sensitivity on the OOD split, evaluated using best-of-four IoU (↑). Only the classifier-free guidance scale changes. Bold marks the best tested scale for each method.

s	Diffusion ↑	Flow Matching ↑	FM+XAttn ↑
0.5	0.850	0.833	0.819
1.0	0.852	0.816	0.867
2.0	**0.890**	**0.880**	**0.873**
3.0	0.870	0.848	0.835
5.0	0.843	0.804	0.826

Table 20: Conditioning architecture ablation for FM+XAttn at s=2.0, K=4, on the OOD split. MES evaluates the sample set (↓); IoU_b evaluates the best of four candidates (↑).

Variant	MES ↓	IoU_b OOD ↑
FM+XAttn (full; cross-attn at 64×64, 32×32)	**0.095**	**0.873**
FM+XAttn (ablated; no cross-attention)	0.096	0.872

F Evaluation Metric Details

The completion metrics defined here use the interface and evaluation region specified in Sections 3 and 5.

Fidelity metrics. Fidelity metrics are computed on the evaluation mask $R_{\text{eval}} = U \odot V$. Before scoring, we hard-clamp observed evidence into every prediction.

IoU and F1 metrics for the floor class are computed from masked counts of true positives, false positives, and false negatives (TP, FP, FN) restricted to R_{eval}. IoU divides TP by $\mathrm{TP} + \mathrm{FP} + \mathrm{FN}$; F1 divides 2TP by $2\mathrm{TP} + \mathrm{FP} + \mathrm{FN}$. Both operate on floor sets within the scored region.

Mean-of-four IoU. For observation i, let $\hat{F}_i^{(k)}$ denote the kth binary completion, with output values ≥ 0.5 labeled floor and observed evidence clamped. We score each completion separately and average the four IoUs:

$$\mathrm{IoU}_{m,i} = \frac{1}{4} \sum_{k=1}^{4} \mathrm{IoU}_{R_{\text{eval},i}} \left(\hat{F}_i^{(k)}, F_i^{\star} \right). \tag{A4}$$

Population masked energy score. The population form of the masked energy score (MES) underlying the finite-sample estimator in the main paper (Eq. (3)) uses masked Jaccard distance $d_{R_{\text{eval}}}(A, B) = 1 - \mathrm{IoU}_{R_{\text{eval}}}(A, B)$ as the dissimilarity. On masks restricted to R_{eval}, it is a bounded metric when the distance between two empty sets is defined as zero (Lipkus, 1999):

$$\mathrm{MES}_{R_{\text{eval}}}(P, F^{\star}) := \mathbb{E}[d_{R_{\text{eval}}}(Y, F^{\star})] - \tfrac{1}{2}\mathbb{E}[d_{R_{\text{eval}}}(Y, Y')], \tag{A5}$$

with $Y, Y' \overset{iid}{\sim} P$ (Gneiting & Raftery, 2007; Rizzo & Székely, 2016). The first term is the expected distance to the target; the second subtracts half the expected distance between two independent samples. For deterministic predictors, $Y{=}Y'$ a.s. and MES reduces to $1 - \mathrm{IoU}_{R_{\text{eval}}}$.

Propriety on a fixed evaluation mask. Fix a finite, nonempty scored pixel set E and restrict floor sets to E. With $\mathrm{IoU}(\varnothing, \varnothing) = 1$, the Jaccard similarity matrix $\mathbf{J}_E$ over distinct subsets of E is strictly positive definite (Bouchard et al., 2013). For distributions P, Q over these subsets, with probability vectors $\mathbf{p}, \mathbf{q}$,

$$\mathbb{E}_{Z \sim Q}[\mathrm{MES}_E(P, Z) - \mathrm{MES}_E(Q, Z)] = \tfrac{1}{2}(\mathbf{p} - \mathbf{q})^{\top} \mathbf{J}_E (\mathbf{p} - \mathbf{q}). \tag{A6}$$

This quantity is positive when $P \neq Q$, establishing strict propriety for distributions over floor sets on the fixed evaluation mask E.

Finite-sample estimators. The second sum in Eq. (3) contains all ordered pairs of distinct samples. For i.i.d. draws from one conditional distribution, the $K > 1$ expression is an unbiased estimator of the population MES above. Deterministic predictors ($K = 1$) use the population limit given above.

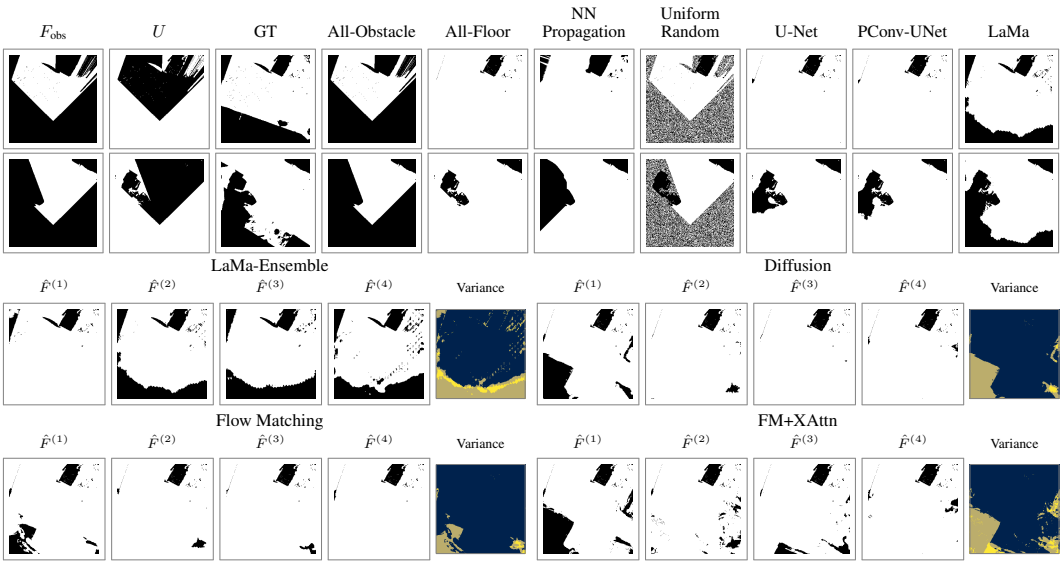


Figure 15: **Completion from source-geometry BEV observations.** Top: single-output baselines for one ID and one OOD scene (rows 1–2), preceded by conditioning and ground truth. Bottom: four outputs from LaMa-Ensemble and each generator for the ID scene, with per-cell variance (brighter means greater disagreement). White/black denote floor/non-floor; the mask marks unobserved cells. Observed labels remain fixed. Extended scenes: Fig. 16.

Binary sample variance. The Var columns and boundary analysis summarize disagreement among four binary completions thresholded at 0.5. For their labels $b_k(x) \in \{0, 1\}$ at cell x, define

$$v(x) = \frac{1}{4}\sum_{k=1}^{4}\big(b_k(x) - \bar{b}(x)\big)^2, \qquad \bar{b}(x) = \frac{1}{4}\sum_{k=1}^{4} b_k(x). \tag{A7}$$

The Var columns average $v(x)$ over R_{eval}; boundary and interior summaries average it within their respective regions (Appendix A.6).

Spatial error. UMR, IoU, and F1 depend on class prevalence and average errors across space; high floor prevalence can obscure local boundary defects. The offline planning study in Section 6 evaluates the resulting maps as route-planning inputs.

G Extended Qualitative Examples

These examples illustrate completions on the canonical track and extend the RGB comparisons and failure cases in Figs. 3 and 5; numerical results are in Appendix A.

Figure 15 gives the canonical comparison referenced from Section 6, and Fig. 16 extends it. Source geometry supplies the BEV evidence used by every completion method. Figure 15 shows single-output baselines and all four outputs from LaMa-Ensemble and each generator. Figure 16 shows one output per method, using best-of-four IoU on R_{eval} for the multi-output methods. Its columns are **O**: observed floor F_{obs}; **M**: evaluation mask R_{eval}; **G**: ground truth; **AO**: All-Obstacle; **AF**: All-Floor; **NN**: NN Propagation; **Un**: Uniform Random; **U**: U-Net; **PC**: PConv-UNet; **LE**: LaMa-Ensemble; **Di**: Diffusion; **Fl**: Flow Matching; and **XA**: FM+XAttn.

Failure cases. In Fig. 17, LaMa obtains IoU ≤ 0.009 throughout. **Row 1**: IoUs are 0.824 (FM+XAttn) and 0.821 (Diffusion). **Row 2**: 0.913 (FM+XAttn), 0.801 (Flow Matching), and 0.788 (Diffusion). **Row 3**: 0.927 (Diffusion) and 0.926 (FM+XAttn).

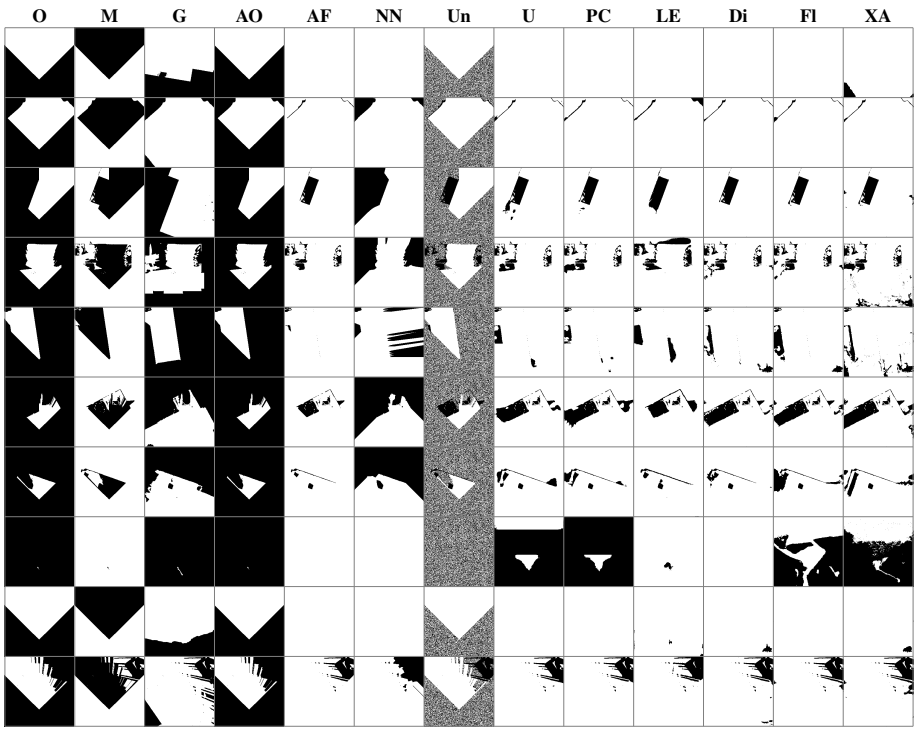


Figure 16: **Extended canonical comparison.** Source-geometry conditioning with single-output baselines and ground-truth-selected best-of-four outputs from the ensemble and generators. Column definitions and selection are specified in Appendix G.

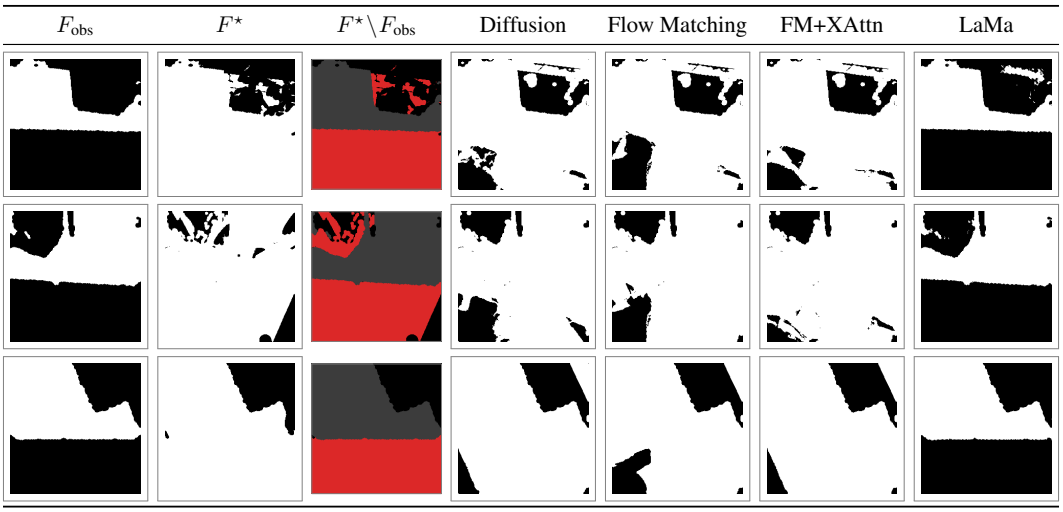


Figure 17: **Supplementary failure cases.** Three additional test scenes show disconnected floor islands (row 1), coarse-shape drift (row 2), and residual boundary artifacts (row 3). The boundary-leakage case is shown in Fig. 5. The red overlay ($F^\star \setminus F_{\text{obs}}$) shows target floor absent from the observation (50–61% across cases). Each generator shows its first generated sample; LaMa shows the single model trained with seed 42. LaMa produces small floor regions in every case.

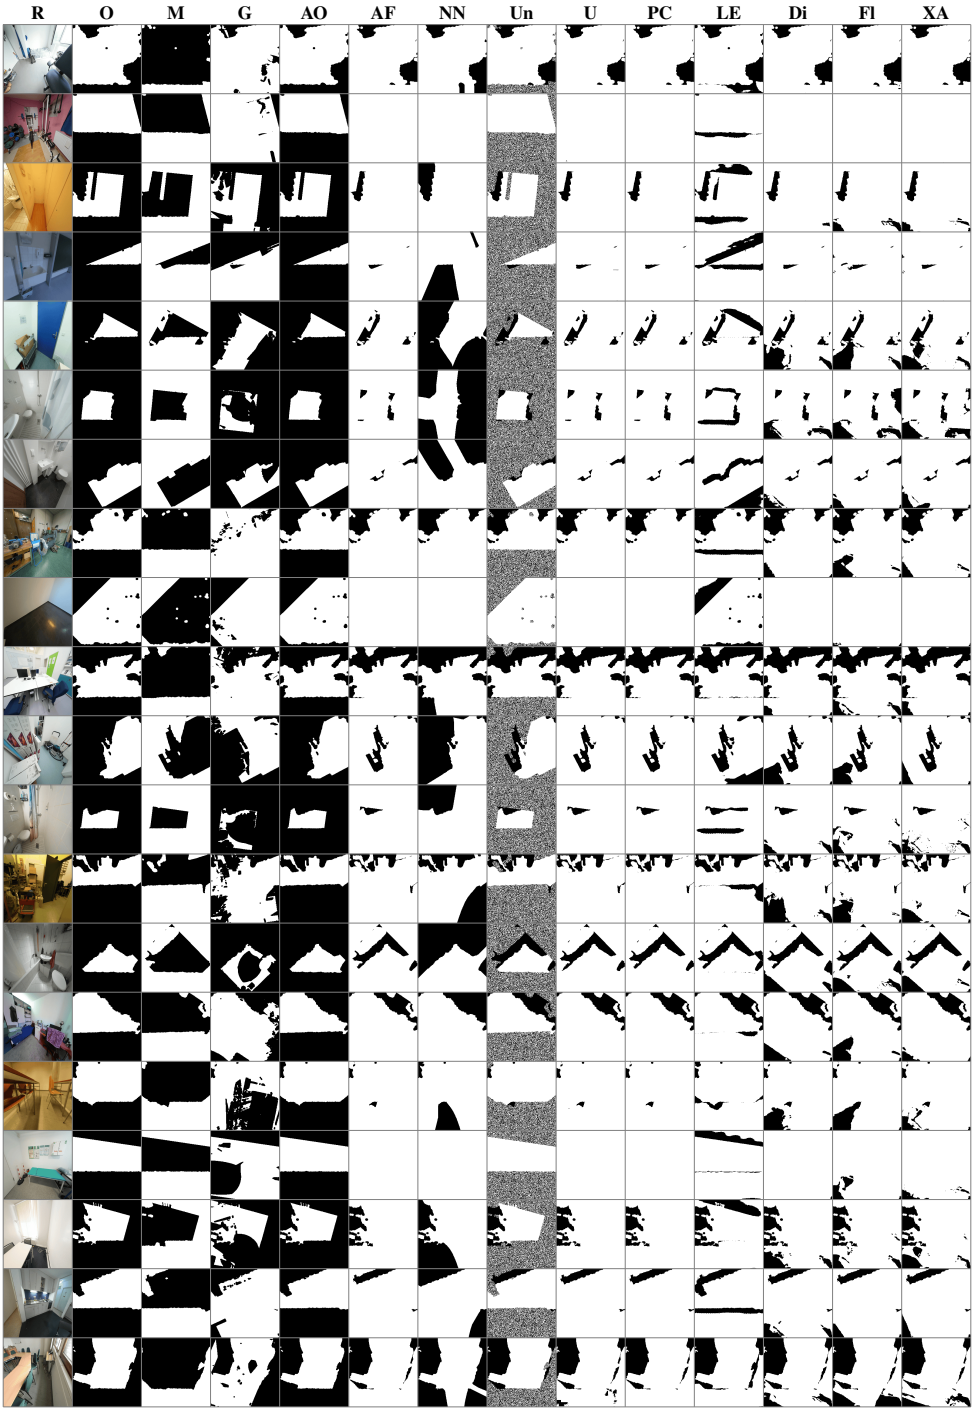


Figure 18: **Extended RGB comparisons (1/2). R**: RGB; **M**: supplied unobserved-region mask; other columns follow Appendix G. Each generator shows its first generated sample; LaMa-Ensemble shows the individual output of its seed-41 member before ensemble averaging. The observed floormap F_{obs} comes from estimated depth, segmentation, and projection. Continued on the next page.

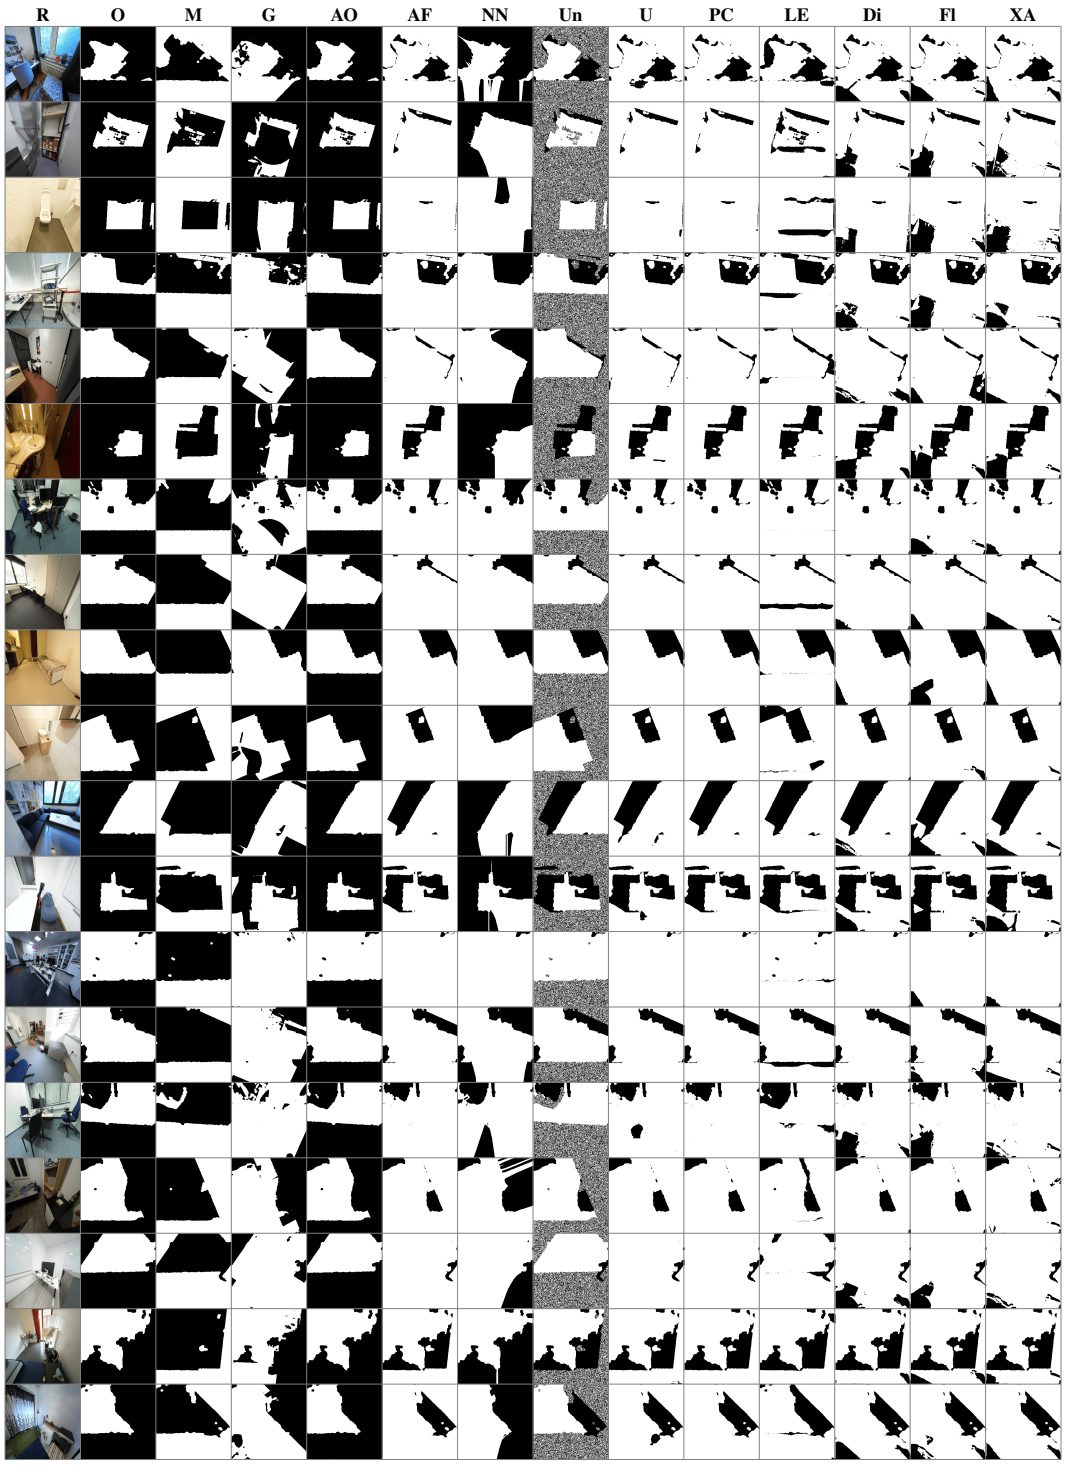


Figure 18: **Extended RGB comparisons (2/2).** Additional frames with the columns and displayed-output rules defined in Fig. 18.