

GroundUp: Rapid Sketch-Based 3D City Massing

Gizem Esra Ünlü¹ Mohamed Sayed² Yulia Gryaditskaya³ Gabriel Brostow¹

¹University College London ²Niantic ³PAI and CVSSP, University of Surrey

Abstract. We propose *GroundUp*, the first sketch-based ideation tool for 3D city *massing* of urban areas. We focus on early-stage urban design, where sketching is a common tool and the design starts from balancing building volumes (masses) and open spaces. With Human-Centered AI in mind, we aim to help architects quickly revise their ideas by easily switching between 2D sketches and 3D models, allowing for smoother iteration and sharing of ideas. Inspired by feedback from architects and existing workflows, our system takes as a first input a user sketch of multiple buildings in a top-down view. The user then draws a perspective sketch of the envisioned site. Our method is designed to exploit the complementarity of information in the two sketches and allows users to quickly preview and adjust the inferred 3D shapes. Our model has two main components. First, we propose a novel sketch-to-depth prediction network for perspective sketches *that exploits top-down sketch shapes*. Second, we use depth cues derived from the perspective sketch as a condition to our diffusion model, which ultimately completes the geometry in a top-down view. Thus, our final 3D geometry is represented as a heightfield, allowing users to construct the city “*from the ground up*”. The code, datasets, and interface are available at visual.cs.ucl.ac.uk/pubs/groundup.

1 Introduction

Urban design has a deep impact on people’s lives, and it epitomizes the opportunities to bring Human-Centered AI for Computer Vision [91] into iterative design [34]. The loop of drawing and discussing buildings, and specifically sketching the buildings’ *masses*, *i.e.* coarse shapes, is the crucial first stage of urban planning [59]. “Architectural design begins with a massing study” [40], where the term “massing” is used for this stage because it locks in the long-term balance between constructed mass versus open space. In pilot interviews, architects said that existing 3D software for urban modeling is too cumbersome for ideation and does not support beginners.

Presently, the ease of sketching is hard to beat. 3D model precision is not the top priority in massing. Rather, urban design aims to satisfy the constraints and desires of whole teams of stakeholders. For example, an architect will play with many massing alternatives, often changing their mind mid-sketch. Currently, they iterate further in 2D with fellow architects on a shortlist of favorites, before re-doing just one or a few designs in 3D software (*e.g.* Rhino or Sketchup), to test out the idea. Our work aims to facilitate the design process by providing the means to quickly preview designs in 3D.

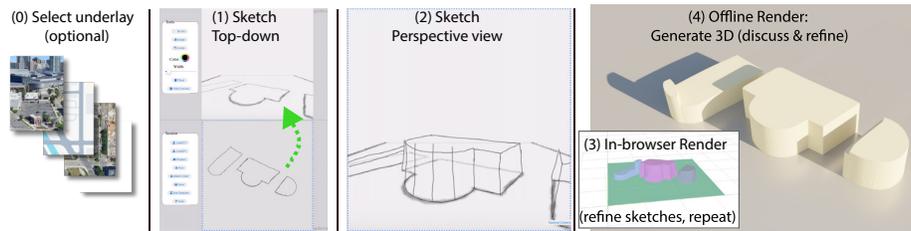


Fig. 1: An illustrative example of our method. (0) Users of our web-based GroundUP system can optionally load registered maps, satellite images, or perspective photographs as underlay layers. These give context for the “massing” process. A blank underlay is used in this example. (1) (bottom) The user sketches the initial footprints of multiple buildings in a top-down view. These strokes are projected into the perspective-view canvas (top). (2) The user sketches a perspective view of the site, and then (3) they trigger our trained model to infer the 3D shape of the sketched buildings. The user can then refine their ideas, iterating between 2D sketching and 3D visuals.

We propose GroundUP, a sketch-based 3D modeling tool for city massing. As shown in Fig. 1 and the video, it works by getting the user to draw and refine their ideas in two views: a top-down “plan” sketch and a perspective sketch. In both views, users can optionally sketch on top of backprojected lines and selected underlay photos. This helps to iterate or when remodeling of an existing site is required. Tightly coupled with this interface, our algorithm quickly infers 3D massing-quality geometry. Such 3D geometry, once approved, can be refined outside of GroundUP and used in the downstream stages of architectural design.

The model intertwined with this interface faces multiple challenges. Compared to photos, sketch lines only provide a sparse signal about the scene. Between a top-down and perspective sketch, it is hard to expect texture regions to match in appearance, making off-the-shelf approaches targeting 3D reconstruction from multi-view images [18, 47, 66] inapplicable to our problem. Additionally, urban areas are inherently complex scenes, so perspective views that convey building heights and roof shapes also suffer from extreme self-occlusions. With many unobserved or partially-observed regions, we turn to diffusion as a generative formulation that could help our method to reconstruct plausible building shapes (Fig. 1(3)).

Critically, the model updates must be responsive for the system to be usable, imposing trade-offs between interactivity and the geometric quality from our adapted latent diffusion model.

Our proposed solution to these challenges offers the following contributions:

- GroundUP is the first system for quick 2D sketch-based iteration on 3D massing design of city blocks.
- Our novel sketch-to-depth prediction network for perspective sketches exploits the top-down view’s cues, and necessitated a bespoke training data process for this important domain.

- We carefully design our top-down diffusion model to handle multiple conditions, integrating cues from both top-down and perspective sketches.

2 Related Work

Sketch-based 3D shape modeling systems greatly facilitate the creation of 3D content, and the first proposed systems came in the 1970s [12,33]. For an in-depth review of existing systems, we refer the reader to the comprehensive reviews [2, 4,6]. Here, we focus on works related to our overall goal of urban reconstruction and papers most related to our method.

2.1 3D Building and City Reconstruction

The exciting works related to the problems of urban reconstruction can be classified into two sets of problems [20]: first of layout generation [1,16,32] and second, of city modeling and rendering [37–39,46,77]. In our work, we aim to provide the user with direct control of the layout via sketching in a top-down view, rather than generating it automatically. Next, many algorithms [43,69,74] for architectural modeling take as input point clouds or Digital Surface Models (DSMs) that contain building height information, obtained with LiDAR (Light Detection and Ranging) or photogrammetry [85]. In contrast, we pursue a different goal of how to obtain buildings’ height information from sparse user-provided sketches.

Multiple works utilize convolutional neural networks (CNNs) for monocular depth estimation from a satellite image [8,9,24,52,54] and building segmentation in a satellite image [7,44,51], or both [44,51]. In the first stage of our method, we also rely on a CNN to obtain a segmentation of a top-down sketch into individual buildings. We then propose to inject this information into a monocular depth estimation network that takes a perspective sketch as an input – the step that we show is paramount in the context of sparse sketch inputs.

2.2 3D from Sketches

3D representations: Sketch to 3D inference has been based on voxel-based representations [14], point clouds [73,88], implicit functions [11,29,89], and 3D diffusion models [3]. Existing methods have a restricted ability to reconstruct details and to scale to larger scenes (*e.g.* multiple objects). We aim for the prompt reconstruction of multiple object shapes within an interactive interface. Our method controls for computational complexity and reduces memory footprint by regressing only 2.5D information, which is subsequently converted to a 3D mesh. We leverage depth and normal maps as intermediate representations. Using intermediate representations such as depth and normal maps is a common approach in sketch-based 3D reconstruction [23,70,76]. Just as we leverage a U-Net architecture [63], several works do so to predict multi-view depth and normal maps [42,49,90]. These methods then fuse the maps to a 3D shape. In contrast, aiming at complex scenes with multiple occlusions, we predict only

one perspective view map and rely on a diffusion model to predict a plausible heightfield, matching perspective and top-down views. Recent work targeting lifting sketches of machine-made shapes to 3D [60], similar to us, first predicts depth. Their full method focuses on the reconstruction of sharp edges. However, it takes about two minutes for single object inference on average. In comparison, our method runs end-to-end in under 2.7 seconds on multi-building scenes.

Ambiguity of 3D reconstruction: For single-view and even sparse multi-view reconstruction, *unobserved* regions create uncertainty, on top of the shape ambiguity of the observed geometric surfaces. Learning of shape category priors is one of the most prominent approaches to dealing with sparse sketch inputs [3, 10, 29, 73, 83, 87, 88]. To further alleviate uncertainty, in single object modeling, symmetry can be leveraged [30, 79]. Other approaches regress parameters of predefined procedural programs [56, 58], or assume availability of additional information [41, 86]. For our task of modeling 3D building shape masses in city neighborhoods, we have more pronounced uncertainty from the multiple layers of occlusions in any perspective view. Additionally, 3D buildings and their groups also have irregular shapes covering diverse geometric configurations. Therefore, we leverage a generative diffusion model that allows us to obtain plausible-looking building *masses* from two sparse sketch constraints (Fig. 1).

Urban modeling: Nishida *et al.* [56] explored data-driven inference of procedural grammars for individual building reconstruction. The reconstruction ability of such methods is limited to what is possible to represent with the considered grammar. Also, their approach cannot infer the shape from a complete drawing and assumes a specific drawing order, matching the grammar used. Liu *et al.* [48] extends procedural modeling to VR sketch inputs. Vitruvio [72] targets individual 3D building reconstruction from input sketches. The paper stresses the importance of *perspective 2D sketch-based modeling in architectural applications for early idea development*. Their method adopts an occupancy network [53] that is either fine-tuned or trained from scratch on synthetic sketches of individual buildings. However, their results show blobby reconstructions with some floating geometry pieces, typical in implicit 3D shape representations such as occupancy grids and signed distance fields [50, 55, 57]. Our heightfield representation allows us to obtain higher quality reconstructions of multiple buildings in one scene.

2.3 Depth Estimation from RGB Images

Sketches are harder for shape inference than RGB images, but we draw lessons nonetheless. For calibrated stereo pairs [13, 35, 82] or unstructured views with known poses [22, 28, 67], cost volumes reveal metric depth by matching photometric appearance between views. Unfortunately, the winning disparities are misleading with our textureless sketches. For depth from a single image, recent methods rely on a learned prior for depth estimation [19, 27, 84]. Follow-ups utilize 3D point networks [81] to combat scale ambiguity, dataset mixing [61] for more generalizable models, classification heads [21] for improved accuracy, or

generative models [17, 36, 65] for sharper depth maps. Recent methods combine the two: cost volumes and strong image priors, to produce sharp metric depths from multiple views [18, 47, 66]. Rather than relying on photometric matching, we utilize a top-down sketch in an occupancy volume to resolve scale ambiguity.

3 Method

Our supervised model is tightly coupled with the user-facing 2D and 3D interface described in Sec. 1 and in Fig. 1. The model has several components, shown and summarised in Fig. 2.

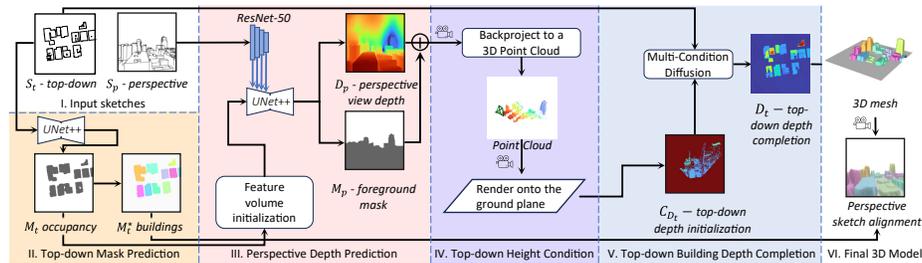


Fig. 2: Reconstruction pipeline overview. (I.) From input sketches, (II.) we estimate the segmentation of the top-down sketch into individual buildings (as detailed in Sec. 3.1). (III.) We then inject the volumetric information about the spaces not occupied by buildings (based on the segmentation result and using a known perspective camera from our interface) into the network that predicts depth and a foreground mask for the perspective sketch view (further detailed in Sec. 3.2). (IV.) From the predicted depth values, we obtain a partial 3D point cloud of the user-envisioned 3D city block. (V) By projecting a sparse 3D prediction into a top-down view, we obtain an initial guess for a top-down view heightfield. Finally, we rely on a diffusion model to obtain a plausible 3D reconstruction that aligns with the perspective and top-down sketches (as shown in V-VI. and detailed in Sec. 3.3).

3.1 Building occupancy mask estimation for top-down sketches

First, given a top-down sketch S_t we aim to obtain building occupancy M_t , and instance segmentation M_t^* maps (Fig. 2 II.): We use a subscript t to denote maps of the top-down views. Our top-down occupancy prediction network follows a UNet++ architecture: an encoder-decoder network with dense and nested skip connections introduced in [92]. As an encoder, we use ResNet-50 [31], initialized with the weights of the model pre-trained on ImageNet [15]. We train our network with a weighted binary-cross entropy (BCE) loss

$$\mathcal{L}_{mask} = -\frac{1}{N} \sum_{i=1}^N \left[\lambda_1 [y_i \log(p_i)] + \lambda_0 [(1 - y_i) \log(1 - p_i)] \right], \quad (1)$$

where y_i and p_i are the ground-truth and predicted mask values for the i^{th} pixel, respectively. λ_0 and λ_1 are the weights for ground and building class predictions, respectively. We empirically found that a bigger weight λ_1 for the building pixels improves mask prediction performance, accounting for class imbalance as buildings occupy a smaller area in the image. We provide further implementation details in the supplemental.

We then segment into individual buildings, M_t^* , by applying Connected-Component Labeling [64] (Fig. 2 II.). We use this building-level segmentation M_t^* for visualization of the 3D reconstruction results in our UI (Fig. 2 VI).

3.2 Depth prediction from perspective sketches

Given a perspective sketch S_p and a top-down building mask M_t , we aim to predict perspective depth maps D_p and foreground masks M_p (Fig. 2 III.) – We use a subscript p to denote maps of the perspective views. In contrast to the masks in Sec. 3.1, M_p labels both the building and ground pixels as foreground, with the background being sky pixels.

Network design Predicting depth from a single sparse sketch is an ill-posed problem. Moreover, in our scenario, each sketch can be quite complex with multiple buildings and occlusions. We design our perspective view depth predictor to handle such complex urban scenes.

Our architecture is inspired by a multi-view depth estimation method [66]. The backbone of this network is a UNet++ architecture, identical to the one we introduced in Sec. 3.1. To reduce ambiguity in a perspective view, we leverage top-down view information. However, applying the multi-view stereo as in [66,78] is not feasible, as our views have little visual overlap so it is infeasible to perform meaningful feature matching between such views. Instead, we exploit the fact that the top-down building occupancy mask, M_t , provides information on whether a location in 3D space is free. We construct a 3D occupancy volume, which is aligned with the perspective view frustum. We construct it by slicing the 3D view frustum of the perspective camera with n depth planes at equidistant intervals between the near d_{near} and far planes d_{far} . We populate the occupancy volume by setting all voxels that fall above non-occupied regions to $-\nu$ and all voxels above occupied regions to ν . We discuss the choice of ν in detail in the supplemental material. Intuitively, we pick ν to be sufficiently large, but within the range of our encoder features. We feed 3D occupancy features as input to the UNet++ encoder. Namely, the 3D occupancy features are of shape $D \times H \times W$, where D is the number of depth planes. When feeding these features into the 2D encoder in the UNet++, we consider depth planes as image feature channels C . Then, similarly to [66], we pass the input sketch through a ResNet-50 encoder to obtain multi-level features. Starting from the first layer of the UNet++, at every second layer, we concatenate output features with corresponding encoded sketch features. This network design allows us to efficiently inject top-down sketch information, resulting in more accurate perspective depth predictions. We provide the ablation study of our design in Sec. 4.

Training During training, we use ground-truth M_t building occupancy masks. We train our depth predictor with a weighted sum of four loss terms, so

$$\mathcal{L}_D = \omega_d \mathcal{L}_{\text{depth}} + \omega_g \mathcal{L}_{\text{grad}} + \omega_n \mathcal{L}_{\text{p,norm}} + \omega_m \mathcal{L}_{\text{mask}}, \quad (2)$$

where ω_* denotes the weight of the corresponding loss component. We introduce each term below.

$\mathcal{L}_{\text{depth}}$ is a multi-scale loss on depth predictions, that was shown to provide sharper depth maps at depth discontinuities than a loss applied only at the final depth map resolution [19, 25, 26, 66]. Following previous works, we predict depths at four resolutions from different levels of our UNet++ decoder, such that at each of the subsequent scales the spatial resolution is doubled. It is defined as

$$\mathcal{L}_{\text{depth}} = \sum_{s=1}^S \|(D_p)_s - (D_p^{gt})_s\|_1 \quad (3)$$

where $\|\cdot\|_1$ is the L_1 -norm and (D_{ps}, D_{ps}^{gt}) are the predicted and ground-truth depth maps at the s^{th} scale.

Similarly, inspired by [45, 66], to encourage smoother gradient changes and sharper depth discontinuities in predicted depth maps, we use a multi-scale loss $\mathcal{L}_{\text{grad}}$ that penalizes differences in depth gradients between the predicted and ground-truth depth map:

$$\mathcal{L}_{\text{grad}} = \sum_{s=1}^S \|\nabla_x R_s\|_1 + \|\nabla_y R_s\|_1, \quad (4)$$

where $R_s = (D_p)_s - (D_p^{gt})_s$.

Following Yin *et al.* [80], who showed that a geometric constraint on normal maps improves monocular depth estimation, we use a loss $\mathcal{L}_{\text{p,norm}}$ between ground-truth N_p^{gt} and predicted N_p normal maps:

$$\mathcal{L}_{\text{p,norm}} = \sum_{i=1}^N (1 - (N_p)_i \cdot (N_p^{gt})_i), \quad (5)$$

where we sum over the dot products of normal vectors $(N_p^*)_i \in \mathbb{R}^3$ in corresponding normal map locations i . We observed that this loss improves the performance in our setting as well. Both N_p and N_p^{gt} are computed on the fly from their corresponding depth maps.

Finally, $\mathcal{L}_{\text{mask}}$ is a weighted BCE loss, defined similarly to the one in Eq. (1). We use it to segment out building and ground pixels.

3.3 Conditional diffusion model for 3D building reconstruction

In the previous section, we described how we obtain a depth estimation for a perspective sketch view. As the next step, we backproject the depth map to

obtain a 3D point cloud. From this point cloud, we initialize a heightfield of the city block aligned with the top-down user sketch. To account for possible inaccuracies in the depth prediction network, we leverage a mask M_t , predicted with our building occupancy mask estimation network as described in Sec. 3.1. We set all heightfield predictions that fall outside the occupied regions to a constant ground-level value. We then use a diffusion model conditioned on the input sketch and the initial heightfield from the perspective sketch view to complete missing depth regions in the top-down view. Since our conditioning relies on both views, the model predicts plausible 3D buildings that align with user sketches. Note that, during training, we initialize heightfields using ground-truth perspective view depth maps.

Network architecture We build on the latent space diffusion model by Duan *et al.* [17], adapting it to handle multiple conditions. We chose a latent diffusion model due to its memory efficiency and inference speed compared to image-space diffusion models.

We map ground-truth depth maps to a latent space using a depth encoder: $z = \mathcal{E}_{\text{depth}}(D_t^{gt})$. Additionally, we encode sketch and depth conditions: $c_{\text{sketch}} = \mathcal{E}_{\text{sketch}}(C_{St})$ and $c_{\text{depth}} = \mathcal{E}_{\text{depth}}(C_{Dt})$, respectively. We initialize $\mathcal{E}_{\text{sketch}}$ and $\mathcal{E}_{\text{depth}}$ with pre-trained weights. Specifically, for $\mathcal{E}_{\text{sketch}}$, we employ a ResNet-50 architecture pre-trained on ImageNet. For the depth encoder, we employ the one from the Stable Diffusion [62]. We pre-train the autoencoder following their strategy and supervise using ground-truth top-down depth maps D_t^{gt} using KL-regularization in the latent space. We fine-tune both latent encoders when training the full model.

To construct the final input to the denoising network, we combine the sketch, c_{sketch} , and depth, c_{depth} , conditions with the noisy depth latent z_k , for a given noise level k . To align features, we pass latent representations c_{depth} and z_k through two separate CNNs, consisting of two convolutional layers. The final denoising network input is created by combining sketch latent features and depth conditions with the noisy depth latent through an element-wise summation.

Training The training objective for the diffusion process is defined as

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{k \sim [1, T], z_k, \epsilon_k} [\|\epsilon_k - \epsilon_{\theta}(z_k, c_{S_t}, c_{D_t}, k)\|]^2, \quad (6)$$

where ϵ_t and ϵ_{θ} are the ground-truth and predicted noise maps, at timestep k .

Additionally, we use auxiliary pixel-based losses to help train the conditioning process. Firstly, L_1 and L_2 losses on predicted D_t and ground-truth D_t^{gt} depth maps are used, defined as

$$\mathcal{L}_{L_1} = \|D_t - D_t^{gt}\|_1 \quad \text{and} \quad \mathcal{L}_{L_2} = \|D_t - D_t^{gt}\|_2. \quad (7)$$

We also use a loss on normal maps $\mathcal{L}_{t, \text{norm}}$, defined similarly to the one in Eq. (5). We find that this loss results in sharper, more uniform depth predictions. We ablate its effect in Sec. 4.

The complete objective loss of our top-down heightfield completion diffusion model is defined as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{diff}} + \mathcal{L}_{L_1} + \mathcal{L}_{L_2} + \mathcal{L}_{t,\text{norm}}. \quad (8)$$

3D mesh: From the ground up Finally, to obtain a 3D mesh, we create a 3D mesh grid $\mathcal{M}^{3D} \in \mathbb{R}^{N \times N \times 3}$ with $N \times N$ vertices, where N is the width/height of the top-down depth map where the horizontal x and vertical y axes map to pixel coordinates. We obtain the height of each vertex v_{ij} in \mathcal{M}^{3D} as

$$v_{ij}^z = d_{\text{ground}} - (D_t)_{ij}, \quad (9)$$

where d_{ground} is the depth value of the ground plane and $(D_t)_{ij}$ is the predicted top-down depth at pixel location (i, j) . We assign d_{ground} to the maximum depth value in D_t .

4 Experiments

In this section, we evaluate our method on synthetic sketches. We first evaluate our perspective depth prediction network and discuss the importance of various design choices. We then assess our complete method, by evaluating our top-down completion network on inputs predicted by the perspective depth network. We compare with a few alternative baselines and ablate our design choices. The details of data generation and splits are provided in the supplemental.

Perspective depth prediction In Tab. 1, we assess our design choices for the perspective depth prediction network and compare against several baselines using standard depth metrics [19]. Briefly, *Abs Diff* is the absolute difference between ground-truth and predicted depth maps, *Abs Rel* is the absolute difference normalized by the ground-truth depth map, *Sq Rel* is the square of *Abs Rel*, *RMSE* is the root mean square error between both depth maps, *Log RMSE* is the root mean square error on logged depths, and *a5* is the ratio of pixels whose depth values have a relative depth error lower than 5%.

Baselines: We train a naive monocular depth predictor baseline from Sayed *et al.* [66] without a cost volume (no source views for multi-view stereo), which we refer to as *Mono*, and compare two image encoder backbones. In Tab. 1, lines [1-2] refer to *Mono_S* for a smaller (EfficientNet [71]) and *Mono_L* for a larger encoder (ResNet-50 [31]). A large image encoder leads to superiority across all depth metrics, with a minimal increase in inference speed – 0.16s on average per sample. Given this, we use this larger backbone for all other experiments.

Ablations: In Tab. 1, *OV* represents our model with the occupancy volume obtained as described in Sec. 3.2. We empirically found $\nu = 50$ to give the best results. This value is close to the mid-point of the range of multi-scale

	Model	Abs Diff ↓	Abs Rel ↓	Sq Rel ↓	RMSE ↓	Log RMSE ↓	a5 ↑
1	<i>Mono_S</i>	6.64	5.33	0.79	10.06	9.41	66.13
2	<i>Mono_L</i>	5.57	4.33	0.53	8.58	7.44	68.82
3	<i>OV_S</i>	6.23	5.24	0.74	9.28	9.55	67.80
4	<i>OV_L</i>	3.49	2.13	0.21	6.54	3.43	89.20

Table 1: Quantitative evaluation of the perspective depth estimation. *Mono* stands for a monocular depth predictor baseline by Sayed *et al.* [66], where subscripts *S* and *L* define a smaller and larger encoder backbones, respectively. *OV* represents our model with the occupancy volume obtained as described in Sec. 3.2. Please see Sec. 4 for the details. All metric values apart from *a5* are scaled up by 10^2 .

image features. We hypothesize that this setting allows the network to leverage the occupancy information most beneficially. We provide a detailed analysis of the choice of ν in the supplemental. Lines [4] vs. [3] show the advantage of the larger encoder backbone. Our complete model then comprises a ResNet-50 encoder backbone, and an occupancy volume with voxels assigned using $\nu = 50$.

Comparison: In Tab. 1, lines [3-4] vs [1-2] show that the *OV* models outperform *Mono* models. We show a qualitative comparison of the *Mono* baseline with our *OV* method in Fig. 3, showing the importance of the proposed occupancy feature volume for correcting for spatial ambiguity from single-view depth estimation.

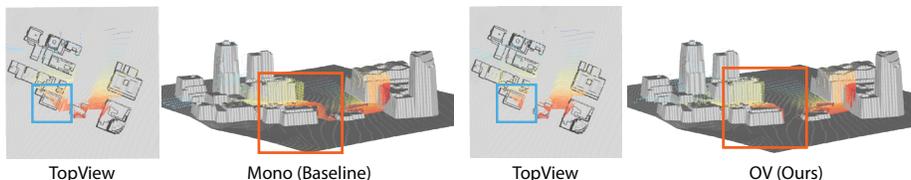


Fig. 3: Qualitative evaluation of the perspective depth estimation. *Mono* stands for a monocular depth predictor baseline by Sayed *et al.* [66]. *OV* represents our model with the occupancy volume, obtained as described in Sec. 3.2. Grey mesh corresponds to the geometry obtained from the ground-truth heightfield. Point clouds represent the estimated depth values from a perspective sketch. Colors encode the distance from a camera. Our prediction visually aligns better with the ground-truth.

Top-down depth completion Our final goal is to infer plausible building geometries from top-down S_t and perspective S_p sketches (Fig. 5 [a,b]). Namely, we rely on the top-down sketch to recover building layouts and on the perspective sketch to estimate buildings’ heights. We obtain height cues with the perspective depth prediction network. Then, the aim of our diffusion model, introduced in Sec. 3.3, is to produce top-down depth maps faithful to the top-down sketch S_t and height cues C_{Dt} .

Model	Abs Diff↓	Abs Rel↓	Sq Rel↓	RMSE↓
S_t+C_{Dt} (fs)	0.1090	0.0221	0.0042	0.1246
S_t+C_{Dt} (pt)	0.1079	0.0218	0.0040	0.1221
S_t+C_{Dt} (pt) + $\mathcal{L}_{t,\text{norm}}$	0.1056	0.0214	0.0039	0.1200
<i>HeightFields</i> [75]	0.1099	0.0225	0.0046	0.1341

Table 2: Quantitative analysis of top-down depth prediction. (fs) denotes training sketch and depth encoders from scratch jointly with the diffusion model. (pt) refers to pre-trained encoders for sketch and depth conditions, as described in Sec. 3.3. S_t+C_{Dt} denotes that we use two conditions: a top-down sketch and a partial top-down depth prediction based on the perspective sketch view. The numbers in the first two lines represent diffusion models trained with the losses defined by equations Eqs. (6) and (7), while the last line represents the model trained with the full loss Eq. (8).

Model	Acc↓	Compl↓	Chamfer↓	Precision↑	Recall↑	F Score↑
S_t+C_{Dt} (fs)	3.89	2.91	3.40	81.6	87.7	84.2
S_t+C_{Dt} (pt)	3.79	2.90	3.34	81.8	88.1	84.4
S_t+C_{Dt} (pt) + $\mathcal{L}_{t,\text{norm}}$	4.35	3.54	3.94	76.1	82.0	78.3
<i>HeightFields</i> [75]	5.35	6.99	6.17	70.4	63.9	66.0

Table 3: Quantitative 3D evaluation of the final reconstructed meshes. The metrics in this table account for the visibility of 3D geometry in a perspective sketch view. Please see Sec. 4 for details. The notation in this table matches the caption of Tab. 2. We details on the metrics: *Completion*, *Accuracy*, *Chamfer Distance*, *Precision*, *Recall*, and *F-Score* can be found in [5].

We first ablate the design of our network and then compare it with a deterministic baseline. For evaluations, we use metrics in 2D (Tab. 2) and 3D (Tab. 3), comparing against the ground-truth. For 2D evaluation, we use metrics similar to the ones in Tab. 1. Since we focus on buildings and not the terrain, we compute all 2D metrics only within buildings’ ground-truth regions, using building masks M_t . We evaluate 3D metrics only for the parts of geometries observed in the perspective sketch viewpoints. This allows us to focus the evaluation on regions for which the perspective sketches provide explicit control of the buildings’ heights. Before computing sampled point cloud distances between predicted and ground-truth meshes, we remove points not in the region around the back-projected ground-truth perspective depth map.

Role of pre-training: Tab. 2 demonstrates the importance of pretraining sketch and depth encoders, $\mathcal{E}_{\text{sketch}}$ and $\mathcal{E}_{\text{depth}}$, respectively. (fs) refers to training the encoders from scratch jointly with the diffusion model and (pt) refers to pre-training latent encoders for sketch and depth conditions.

Role of normal loss: We show the qualitative evaluation of the role of the normal loss in Fig. 4. It shows that the normal loss yields building geometries with sharper corners and flat building tops. Computed on all building regions, 2D losses in Tab. 2 show that the normal loss $\mathcal{L}_{t,\text{norm}}$, defined with Eq. (8), sig-

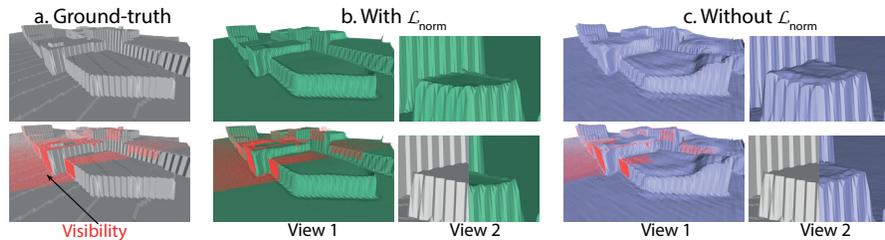


Fig. 4: Role of the normal loss. a) Visibility regions (red points) are computed based on ground-truth geometry and the perspective sketch viewpoint. b) Prediction when the normal loss is used: The red point cloud is riding slightly above the green prediction. As shown in *view 2* the height is slightly underestimated in the visible regions, but the loss results in more even roofs overall. c) Prediction when the normal loss is not used: the model produces bobby building geometry outside the visible regions.

nificantly improves the accuracy of top-down depth-predictions – reflecting on the overall appearance of the buildings. 3D metrics in Tab. 3, computed only on visible regions from the perspective sketch viewpoint, highlight a slight geometry shrinkage, visible in View 2 in Fig. 4. While adding a normal loss hurts quantitative 3D metrics, we advocate its usage as it produces much sharper and smoother surfaces, as shown in Fig. 4 and supported with Tab. 2.

Comparison with a deterministic baseline: Qualitative results of our method are shown in Fig. 5 (d) and (g): We can infer realistic building geometries following input sketches that closely resemble the ground-truth – Fig. 5 (e) and (h). We compare our generative approach against the *HeightFields* [75] baseline – a deterministic model designed for heightfield completion from multi-frame RGB sequences. We train and test it on the same input as our model and visualize the results in Fig. 5 (c) and (f). In particular, the *HeightFields* model’s test time input is the output of our first step: the predicted partial point cloud from a perspective view. This model is not a suitable stand-alone method for the task. To train this model, we also added $\mathcal{L}_{t,\text{norm}}$, as we found it to result in better performance. However, even with this additional loss, the *HeightFields* model fails to produce buildings with correct heights, and produces less plausible building geometries. In particular, it fails to capture sharp details and flat rooftops.

Tabs. 2 and 3 show quantitative comparison of our full model with *HeightFields* [75] baseline. They show the superiority of our diffusion model in all settings, confirming the visual observations.

5 User Study

Modeling Interface: To validate our contributions, we built an interactive user interface in HTML, JavaScript, and Python. The 3D massing system runs real-time on a Titan X and can be used on any touch-screen device through a

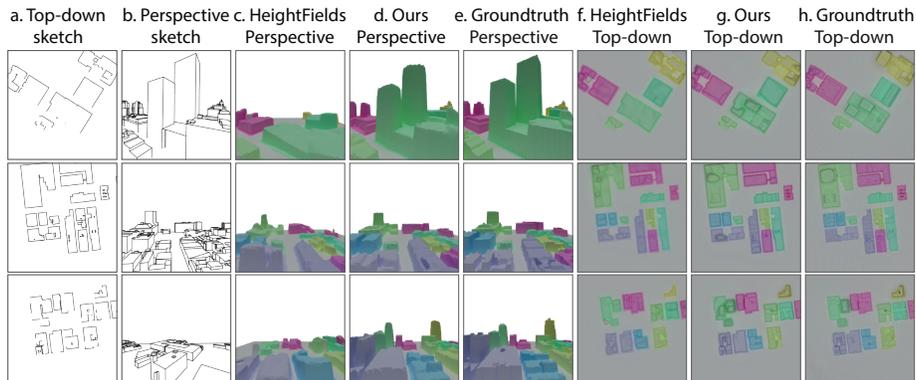


Fig. 5: Qualitative evaluation on synthetic sketches. (a) and (b) show example top-down and perspective sketches. (c) and (f) show reconstruction results obtained with the *HeightFields* [75] method, which is trained and tested on the same data as our method. (d) and (g) show reconstruction results by our method. (e) and (h) show the heightfield of the ground-truth top-down depth map. Note that the colors are assigned according to the ground-truth segmentation of buildings. Please zoom in to better see the alignment of predicted geometries with the ground-truth buildings’ areas.

browser, ideally with a stylus. Broadly, the UI lets users sketch perspective and top-down views on 2D canvases, edit strokes, project a top-down sketch into the perspective canvas to align their sketches, and contains a 3D viewer. An overview of the UI is in Fig. 1 and is described in greater detail in the supplemental.

Evaluation: To validate our system, we run a proof-of-concept user study. For the study, we collaborated with one of the world-leading schools in urban design, the Bartlett School of Architecture, at University College London. We engaged 5 urban design architects: 2 undergraduate students, advanced in their studies, and 3 postgraduates with varying years of professional practice. Additionally, to test how friendly our system is for users with limited modeling and sketching experience, we engaged 5 further volunteers. All users watched a short video tutorial and had 5 minutes to play with the interface before starting the task. To have a concrete qualitative goal in our main study, we chose to provide participants with reference top-down and perspective renderings as underlays (the example screenshot is provided in the supplemental). We selected 9 scenes, randomly distributed between participants. Each participant drew two scenes.

In a post-study questionnaire, all architects indicated that they were able to recreate the building from the reference in under 5 minutes. As expected, it was more challenging for novices, yet, 2/5 were satisfied with the outcome. On a 5-point Likert scale, architects (novices) gave an average score of 0.8 (1.2) on how well the results match the reference, with +2 for matching the reference well and -2 for failing completely. On a 5-point Likert scale, architects gave an average 1.4 score on how likely they are to use such an interface: where -2 for highly unlikely and +2 for highly likely. This analysis shows that overall our

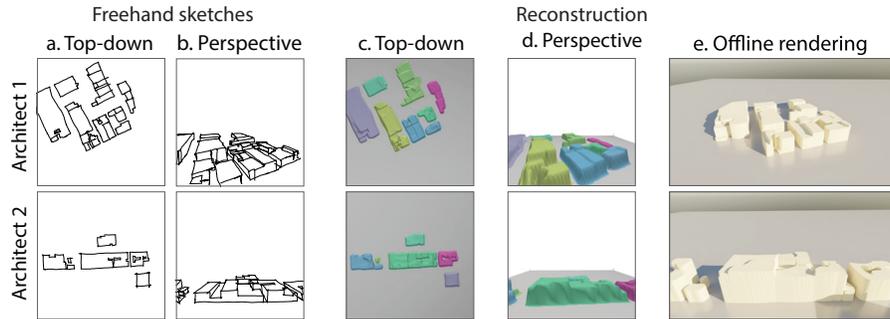


Fig. 6: Freehand sketches and corresponding 3D reconstructions in our user interface, made by urban design architects in our study. (e) shows automatically post-processed results, rendered with an offline rendered, as described in the supplemental.

system achieves a set goal of fast prototyping of building masses, while future work could aim to further improve the reconstruction accuracy. The detailed statistics for the post-study questionnaire are provided in the supplemental.

In our pilot study, architects consistently indicated that it would take them about 10 minutes in Rhino for scenes comparable to the ones we target. The pilot study on SketchUp, documented in the supplemental, similarly showed that it is not suitable for fast prototyping. This, in particular, shows the lack of convenient tools for early design stages and reinforces the motivation for our work.

Two urban design architects also did *freehand modeling* after the main study and completed post-study questionnaires. These sketches are shown in Fig. 6.

6 Conclusion and discussion

We have presented the first sketch-based method for early-stage urban design, aligning it with the Human-Centered AI philosophy [68]. Taking into account design workflows that commonly start from top-down city layouts, we proposed models that, while working in image space, efficiently leverage information from both perspective and top-down sketch views. GroundUp addresses the especially challenging (but not unique) aspects of our problem: complexity and diversity of scene geometries, sparsity of sketch inputs, and incomplete depth cues in user-provided views. While we only show the results for a single perspective sketch, our system is trivially extended to a multi-view setting: by projecting point clouds inferred from extra perspective sketches into the top-down views passed to our diffusion model. We provide numerical experiments in the supplemental. With this work, we have taken a step toward quick building massing. To propel the integration of our tool into design workflows, future work might focus on directly predicting editable mesh representations and supporting finer details. Additionally, it could be interesting to extend this work to trees and terrain, for example, by sketching trunks and contour lines for the terrain.

Acknowledgements

We thank Prof. Tobias Ritschel for his invaluable feedback and help; Natalia Laskovaya for an inspiring and detailed early discussion on design processes in architecture; Sharon Betts for her huge help in making our user study possible. We also thank Kening Guo and all the anonymous participants of the user studies. Gizem Esra Ünlü is funded by a Niantic PhD scholarship.

References

1. Benes, B., Zhou, X., Chang, P., Cani, M.P.R.: Urban brush: Intuitive and controllable urban layout editing. In: *The 34th Annual ACM Symposium on User Interface Software and Technology (2021)*
2. Bhattacharjee, S., Chaudhuri, P.: A survey on sketch based content creation: from the desktop to virtual and augmented reality. *Computer Graphics Forum* **39**, 757–780 (05 2020)
3. Binninger, A., Hertz, A., Sorkine-Hornung, O., Cohen-Or, D., Giryes, R.: Sens: Sketch-based implicit neural shape modeling. *Arxiv preprint -(-)* (06 2023)
4. Bonnici, A., Akman, A., Calleja, G., Camilleri, K., Fehling, P., Ferreira, A., Hermuth, F., Israel, J., Landwehr, T., Liu, J., Padfield, N., Sezgin, T., Rosin, P.: Sketch-based interaction and modeling: where do we stand? *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* **33**, 1–19 (11 2019)
5. Bozic, A., Palafox, P., Thies, J., Dai, A., Nießner, M.: TransformerFusion: Monocular RGB scene reconstruction using transformers. *NeurIPS (2021)*
6. Camba, J.D., Company, P., Naya, F.: Sketch-based modeling in mechanical engineering design: Current status and opportunities. *Computer-Aided Design* **150**, 103283 (2022)
7. Chen, S., Ogawa, Y., Zhao, C., Sekimoto, Y.: Large-scale individual building extraction from open-source satellite imagery via super-resolution-based instance segmentation approach. *ISPRS Journal of Photogrammetry and Remote Sensing* **195** (2023)
8. Chen, S., Shi, Y., Xiong, Z., Zhu, X.X.: Htc-dc net: Monocular height estimation from single remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* **61** (2023)
9. Chen, Z., Zhang, Y., Qi, X., Mao, Y., Zhou, X., Niu, L., Wu, H., Wang, L., Ge, Y.: Heightformer: A multilevel interaction and image-adaptive classification-regression network for monocular height estimation with aerial images. *arXiv preprint arXiv:2310.07995* (2023)
10. Cheng, Z., Chai, M., Ren, J., Lee, H.Y., Olszewski, K., Huang, Z., Maji, S., Tulyakov, S.: Cross-modal 3D Shape Generation and Manipulation, pp. 303–321. Springer (11 2022)
11. Chowdhury, P.N., Wang, T., Ceylan, D., Song, Y.Z., Gryaditskaya, Y.: Garment ideation: Iterative view-aware sketch-based garment modeling. In: *2022 International Conference on 3D Vision (3DV)*. pp. 22–31 (2022)
12. Clowes, M.B.: On seeing things. *Artificial intelligence* **2**(1), 79–116 (1971)
13. Collins, R.T.: A space-sweep approach to true multi-image matching. In: *CVPR (1996)*

14. Delanoy, J., Aubry, M., Isola, P., Efros, A.A., Bousseau, A.: 3d sketching using multi-view deep volumetric prediction. *Proc. ACM Comput. Graph. Interact. Tech.* **1**(1) (Jul 2018)
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. IEEE (2009)
16. Deng, J., Chai, W., Guo, J., Huang, Q., Hu, W., Hwang, J.N., Wang, G.: Citygen: Infinite and controllable 3d city layout generation. *arXiv preprint arXiv:2312.01508* (2023)
17. Duan, Y., Zhu, Z., Guo, X.: Diffusiondepth: Diffusion denoising approach for monocular depth estimation. *CoRR* **abs/2303.05021** (2023)
18. Duzceker, A., Galliani, S., Vogel, C., Speciale, P., Dusmanu, M., Pollefeys, M.: Deepvideomvs: Multi-view stereo on video with recurrent spatio-temporal fusion. In: *CVPR* (2021)
19. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*. pp. 2366–2374 (2014)
20. Feng, T., Fan, F., Bednarz, T.: A review of computer graphics approaches to urban modeling from a machine learning perspective. *Frontiers of Information Technology & Electronic Engineering* **22**(7) (2021)
21. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018)
22. Furukawa, Y., Hernández, C., et al.: Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision* **9**(1-2), 1–148 (2015)
23. Gao, C., Yu, Q., Sheng, L., Song, Y., Xu, D.: Sketchsampler: Sketch-based 3d reconstruction via view-dependent depth sampling. In: *ECCV* (2022)
24. Ghamisi, P., Yokoya, N.: Img2dsm: Height simulation from single imagery using conditional generative adversarial net. *IEEE Geoscience and Remote Sensing Letters* **15**(5) (2018)
25. Godard, C., Aodha, O.M., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. pp. 6602–6611. IEEE Computer Society (2017)
26. Godard, C., Aodha, O.M., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. pp. 3827–3837. IEEE (2019)
27. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: *CVPR* (2017)
28. Goesele, M., Curless, B., Seitz, S.M.: Multi-view stereo revisited. In: *CVPR* (2006)
29. Guillard, B., Remelli, E., Yvernay, P., Fua, P.: Sketch2mesh: Reconstructing and editing 3d shapes from sketches. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021)
30. Hähnlein, F., Gryaditskaya, Y., Sheffer, A., Bousseau, A.: Symmetry-driven 3d reconstruction from concept sketches. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–8 (2022)

31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2016)
32. He, L., Aliaga, D.: Globalmapper: Arbitrary-shaped urban layout generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 454–464 (October 2023)
33. Huffman, D.A.: Impossible objects as nonsense sentences. *Machine intelligence* **6**, 295–323 (1971)
34. Jacoby, S.: *Drawing Architecture and the Urban*. Wiley (2016)
35. Kang, S.B., Szeliski, R., Chai, J.: Handling occlusions in dense multi-view stereo. In: CVPR (2001)
36. Ke, B., Obukhov, A., Huang, S., Metzger, N., Dautt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. arXiv preprint arXiv:2312.02145 (2023)
37. Kelly, T., Femiani, J., Wonka, P., Mitra, N.J.: Bigsur: Large-scale structured urban reconstruction. *ACM Transactions on Graphics* **36**(6) (November 2017)
38. Kelly, T., Guerrero, P., Steed, A., Wonka, P., Mitra, N.J.: Frankengan: Guided detail synthesis for building mass models using style-synchronized gans. *ACM Trans. Graph.* **37**(6), 1:1–1:14 (2018)
39. Kim, S., Kim, D., Choi, S.: Citycraft: 3d virtual city creation from a single image. *The Visual Computer* **36** (2020)
40. Leyton, M.: *A generative theory of shape*. vol. 2145, p. p366. Springer Berlin / Heidelberg, Germany (2001)
41. Li, C., Pan, H., Bousseau, A., Mitra, N.J.: Free2cad: Parsing freehand drawings into cad commands. *ACM TOG* (2022)
42. Li, C., Pan, H., Liu, Y., Tong, X., Sheffer, A., Wang, W.: Robust flow-guided neural prediction for sketch-based freeform surface modeling. *ACM Trans. Graph.* **37**(6) (2018)
43. Li, L., Song, N., Sun, F., Liu, X., Wang, R., Yao, J., Cao, S.: Point2roof: End-to-end 3d building roof modeling from airborne lidar point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing* **193**, 17–28 (2022)
44. Li, X., Wen, C., Wang, L., Fang, Y.: Geometry-aware segmentation of remote sensing images via joint height estimation. *IEEE Geoscience and Remote Sensing Letters* **19** (2021)
45. Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018. pp. 2041–2050. Computer Vision Foundation / IEEE Computer Society (2018)
46. Lin, C.H., Lee, H.Y., Menapace, W., Chai, M., Siarohin, A., Yang, M.H., Tulyakov, S.: Infincity: Infinite-scale city synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2023)
47. Lipson, L., Teed, Z., Deng, J.: Raft-stereo: Multilevel recurrent field transforms for stereo matching. In: 2021 International Conference on 3D Vision (3DV). pp. 218–227. IEEE (2021)
48. Liu, Z., Zhang, F., Cheng, Z.: Buildingsketch: Freehand mid-air sketching for building modeling. In: IEEE International Symposium on Mixed and Augmented Reality (ISMAR). IEEE (2021)
49. Lun, Z., Gadelha, M., Kalogerakis, E., Maji, S., Wang, R.: 3d shape reconstruction from sketches via multi-view convolutional networks. In: International Conference on 3D Vision (3DV) (2017)

50. Luo, L., Chowdhury, P.N., Xiang, T., Song, Y.Z., Gryaditskaya, Y.: 3d vr sketch guided 3d shape prototyping and exploration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
51. Mahdi, E., Ziming, Z., Xinming, H.: Aerial height prediction and refinement neural networks with semantic and geometric guidance. arXiv preprint arXiv:2011.10697 (2020)
52. Mahmud, J., Price, T., Bapat, A., Frahm, J.M.: Boundary-aware 3d building reconstruction from a single overhead image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020)
53. Mescheder, L.M., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. pp. 4460–4470. Computer Vision Foundation / IEEE, Long Beach, CA, USA (2019)
54. Mou, L., Zhu, X.X.: Im2height: Height estimation from single monocular imagery via fully residual convolutional-deconvolutional network. arXiv preprint arXiv:1802.10249 (2018)
55. Nam, G., Khlifi, M., Rodriguez, A., Tono, A., Zhou, L., Guerrero, P.: 3d-ldm: Neural implicit 3d shape generation with latent diffusion models. arXiv preprint arXiv:2212.00842 (2022)
56. Nishida, G., Garcia-Dorado, I., Aliaga, D.G., Benes, B., Bousseau, A.: Interactive sketching of urban procedural models. *ACM Transactions on Graphics (TOG)* **35**(4) (2016)
57. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2019)
58. Pearl, O., Lang, I., Hu, Y., Yeh, R.A., Hanocka, R.: Geocode: Interpretable shape programs. arXiv preprint arXiv:2212.11715 (2022)
59. Pitts, G., Luther, M.: A parametric approach to 3d massing and density modelling. In: Digital Physicality: Proceedings of the 30th eCAADe Conference. pp. 157–165 (2012)
60. Puhachov, I., Martens, C., Kry, P.G., Bessmeltsev, M.: Reconstruction of machine-made shapes from bitmap sketches. *ACM Trans. Graph.* **42**(6) (2023)
61. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* **44**(3), 1623–1637 (2020)
62. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 10674–10685. IEEE, New Orleans, Louisiana, USA (2022)
63. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. Springer (2015)
64. Rosenfeld, A., Pfaltz, J.L.: Sequential operations in digital picture processing. *Journal of the ACM (JACM)* **13**(4), 471–494 (1966)
65. Saxena, S., Kar, A., Norouzi, M., Fleet, D.J.: Monocular depth estimation using diffusion models. *CoRR* **abs/2302.14816** (2023)

66. Sayed, M., Gibson, J., Watson, J., Prisacariu, V., Firman, M., Godard, C.: Simplecon: 3d reconstruction without 3d convolutions. In: *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIII. Lecture Notes in Computer Science*, vol. 13693, pp. 1–19. Springer (2022)
67. Schönberger, J.L., Zheng, E., Frahm, J.M., Pollefeys, M.: Pixelwise view selection for unstructured multi-view stereo. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III 14*. pp. 501–518. Springer (2016)
68. Shneiderman, B.: *Human-centered AI*. Oxford University Press (2022), <https://books.google.co.uk/books?id=YS9VEAAAQBAJ>
69. Stucker, C., Schindler, K.: Resdepth: A deep residual prior for 3d reconstruction from high-resolution satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing* **183** (2022)
70. Su, W., Du, D., Yang, X., Zhou, S., Fu, H.: Interactive sketch-based normal map generation with deep neural networks. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* **1**(1) (2018)
71. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. PMLR (2019)
72. Tono, A., Huang, H., Agrawal, A., Fischer, M.: Vitruvio: 3d building meshes via single perspective sketches. *arXiv preprint arXiv:2210.13634* (2022)
73. Wang, J., Lin, J., Yu, Q., Liu, R., Chen, Y., Yu, S.X.: 3d shape reconstruction from free-hand sketches. In: *Karlinsky, L., Michaeli, T., Nishino, K. (eds.) ECCV Workshops* (2022)
74. Wang, Y., Zorzi, S., Bittner, K.: Machine-learned 3d building vectorization from satellite imagery. In: *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19-25, 2021*. pp. 1072–1081. Computer Vision Foundation / IEEE, Virtual (2021)
75. Watson, J., Vicente, S., Aodha, O.M., Godard, C., Brostow, G.J., Firman, M.: Heightfields for efficient scene reconstruction for AR. In: *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2023, Waikoloa, HI, USA, January 2-7, 2023*. pp. 5839–5849. IEEE (2023)
76. Wu, J., Zhang, C., Zhang, X., Zhang, Z., Freeman, W.T., Tenenbaum, J.B.: Learning shape priors for single-view 3d completion and reconstruction. In: *ECCV* (2018)
77. Xie, H., Chen, Z., Hong, F., Liu, Z.: Citydreamer: Compositional generative model of unbounded 3d cities. *arXiv preprint arXiv:2309.00610* (2023)
78. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: MVSNet: Depth inference for unstructured multi-view stereo. In: *ECCV* (2018)
79. Yao, Y., Schertler, N., Rosales, E., Rhodin, H., Sigal, L., Sheffer, A.: Front2back: Single view 3d shape reconstruction via front to back prediction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020)
80. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. pp. 5683–5692. IEEE (2019)
81. Yin, W., Zhang, J., Wang, O., Niklaus, S., Mai, L., Chen, S., Shen, C.: Learning to recover 3d scene shape from a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 204–213 (2021)
82. Žbontar, J., LeCun, Y.: Stereo matching by training a convolutional neural network to compare image patches. *JMLR* (2016)

83. Zhang, S.H., Guo, Y.C., Gu, Q.W.: Sketch2model: View-aware 3d modeling from single free-hand sketches. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6012–6021 (2021)
84. Zhao, C., Sun, Q., Zhang, C., Tang, Y., Qian, F.: Monocular depth estimation based on deep learning: An overview. *Science China Technological Sciences* **63**(9) (2020)
85. Zhao, L., Wang, H., Zhu, Y., Song, M.: A review of 3d reconstruction from high-resolution urban satellite images. *International Journal of Remote Sensing* **44**(2) (2023)
86. Zheng, J., Zhu, Y., Wang, K., Zou, Q., Zhou, Z.: Deep learning assisted optimization for 3d reconstruction from single 2d line drawings. arXiv e-prints pp. arXiv-2209 (2022)
87. Zheng, X.Y., Pan, H., Wang, P.S., Tong, X., Liu, Y., Shum, H.Y.: Locally attentional sdf diffusion for controllable 3d shape generation. *ACM Trans. Graph.* **42**(4) (2023)
88. Zhong, Y., Gryaditskaya, Y., Zhang, H., Song, Y.: Deep sketch-based modeling: Tips and tricks. In: Struc, V., Fernández, F.G. (eds.) *International Conference on 3D Vision (3DV)*. IEEE (2020)
89. Zhong, Y., Gryaditskaya, Y., Zhang, H., Song, Y.Z.: A study of deep single sketch-based modeling: View/style invariance, sparsity and latent space disentanglement. *Computers & Graphics* **106**, 237–247 (2022)
90. Zhong, Y., Qi, Y., Gryaditskaya, Y., Zhang, H., Song, Y.Z.: Towards practical sketch-based 3d shape generation: The role of professional sketches. *IEEE Transactions on Circuits and Systems for Video Technology* (2020)
91. Zhou, B., Russakovsky, O., Fong, R., Hoffman, J.: *CVPR Tutorial on Human-Centered AI for Computer Vision* (2022), <https://human-centeredai.github.io/>
92. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*. Springer (2018)