[Supplementary Document] Roof-GAN: Learning to Generate Roof Geometry and Relations for Residential Houses

Yiming Qian Hao Zhang Yasutaka Furukawa
Simon Fraser University
{yimingq, haoz, furukawa}@sfu.ca

The supplementary document provides 1) the full architectural specification of Roof-GAN in Table 1; 2) the details of the facet angle and primitive type vectorization algorithms in Sects 1 and 2; 3) rasterization heuristics for GT preparation in Sect 3; and 4) more generated roof models in Sect 4.

1. Facet angle vectorization

Facet angle vectorization is simple thanks to our raster-geometry representation. Suppose a roof is a horizontal hip (Fig. 2c in the paper), we compute the average roof angle inside the green pixels to obtain left/right facet angle, for example. Note that the green region consists of two facets (left and right), but we assume symmetry and the roof "angle" is expected to be the same inside the region. Precisely, we compute the angle for the left/right facets by taking the weighted average of the facet angle image and applying the inverse cosine, where the weight is specified by the left/right channel in the facet orientation image. The same goes for the top/bottom angle.

2. Primitive type vectorization

A roof primitive has four types (Fig. 2 in the paper). First, we binarize the facet orientation image by winner-takes-all at each pixel (*i.e.*, a channel with the highest probability becomes one). Let S_b and S_g denote the sum of the binary image in the top/bottom (blue) and the left/right (green) channels, respectively. We determine the type to be horizontal if $S_b \geq S_g$, otherwise vertical. Next, suppose a primitive is a horizontal type. If the green region is significantly smaller (i.e., $S_g \leq 0.05(S_b + S_g)$), we determine the type to be gable, otherwise hip. The same goes for the vertical case.

3. Rasterization for GT data

For a roof, we first rasterize each primitive using the annotated bounding box and the estimated facet angles from optimization, resulting in the per-primitive facet orientation image. In practice, for each pixel inside the bounding box, we compute its heights using all facets (two for gable, four for hip), and then label the pixel belonging to the facet that has the lowest height.

Next, we combine each primitive's facet orientation image to generate a common orientation image for the entire roof model. For each non-background pixel, we gather all facets that occupy it, and label it with the facet that has the highest heights. Note that coplanar facets are merged and have the same label.

Finally, we extract the polygonal corners of each facet, which are used for our RMMD metric computation. Since our primitives are axis-aligned rectangles, the polygonal facet corners become the endpoints of horizontal and vertical lines in our case. For each binary mask of a facet, we apply the OpenCV function "findContours" while setting the contour approximation mode as "CHAIN_APPROX_SIMPLE". Then, we trace along each pair of consecutive contour points, which are treated as endpoints if their x or y coordinates are equal.

4. More qualitative results

Figures 1, 2, 3, 4, 5, and 6 show more qualitative comparisons of generated models by PQ-Net, House-GAN and Roof-GAN. Figures 7 and 8 show more results of RMMD-based model retrieval evaluations.

Table 1. Network architecture of Roof-GAN. **k** is the kernel size, **s** the stride, **channel** the number of channels of input and output, **input** the input to the **layer**, **output size** the size of the output in the form of $channel \times height \times width$. All convolutional layers are followed by a Leaky ReLU layer (negative_slope $= 0.1$) except for the output layers or the layers before a linear layer. padding $= 1$ is set to all convolutional layers.

network	layer	k	s	channel	input	output size
Generator	linear_reshape_1			128/1024	128d noise	$16 \times 8 \times 8$
	convmpn_1	3	1	[32/32, 32/16, 16/16]	linear_reshape_1	$16 \times 8 \times 8$
	upsample_1	4	2	16/16	convmpn_1	$16 \times 16 \times 16$
	convmpn_2	3	1	[32/32, 32/16, 16/16]	upsample_1	$16 \times 16 \times 16$
	upsample_2	4	2	16/16	convmpn_2	$16 \times 32 \times 32$
	conv_3a	3	1	16/256	upsample_2	$256 \times 32 \times 32$
	conv_3b	3	1	256/128	conv_3a	$128 \times 32 \times 32$
	conv_softmax_3c	3	1	128/3	conv_3b	$3 \times 32 \times 32$
	conv_4a	3	1	16/256	upsample_2	$256 \times 32 \times 32$
	conv_4b	3	1	256/128	conv_4a	$128 \times 32 \times 32$
	conv_sigmoid_4c	3	1	128/1	conv_4b	$1 \times 32 \times 32$
	conv_5a	4	2	32/48	upsample_2 (pair)	$48 \times 16 \times 16$
	conv_5b	4	2	48/64	conv_5a	$64 \times 8 \times 8$
	conv_5c	4	2	64/128	conv_5b	$128 \times 4 \times 4$
	conv_5d	4	2	128/256	conv_5c	$256 \times 2 \times 2$
	conv_reshape_5e	4	2	256/512	conv_5d	512×1
	linear_sigmoid_5f			512/6	conv_reshape_5e	6×1
Geometry discriminator	conv_1a	3	1	6/16	orientation	$16 \times 32 \times 32$
	conv_1b	3	1	16/16	conv_1a	$16 \times 32 \times 32$
	conv_1c	3	1	16/16	conv_1b	$16 \times 32 \times 32$
	conv_2a	3	1	2/16	angle	$16 \times 32 \times 32$
	conv_2b	3	1	16/16	conv_2a	$16 \times 32 \times 32$
	conv_2c	3	1	16/16	conv_2b	$16 \times 32 \times 32$
	convmpn_3	3	1	[64/64, 64/32, 32/32]	conv_1c+conv_2c	$32 \times 32 \times 32$
	conv_3	3	2	32/32	convmpn_3	$32 \times 16 \times 16$
	convmpn_4	3	1	[64/64, 64/32, 32/32]	conv_3	$32 \times 32 \times 32$
	conv_4	3	2	32/32	convmpn_4	$32 \times 8 \times 8$
	conv_5a	3	2	32/256	conv_4	$256 \times 4 \times 4$
	conv_5b	3	2	256/128	conv_5a	$128 \times 2 \times 2$
	conv_reshape_5c	3	2	128/128	conv_5b	128×1
	linear_5d			128/1	conv_reshape_5c (maxpool)	1
Relationship discriminator	linear_reshape_1			6/4096	relation	$4 \times 32 \times 32$
	conv_1	4	2	26/64	orientation+angle+relation	$64 \times 16 \times 16$
	conv_2	4	2	64/128	conv_1	$128 \times 8 \times 8$
	conv_3	4	2	128/256	conv_2	$256 \times 4 \times 4$
	conv_4	4	2	256/256	conv_3	$256 \times 2 \times 2$
	conv_reshape_5	4	2	256/128	conv_4	128×1
	linear_6			128/1	conv_reshape_5	1

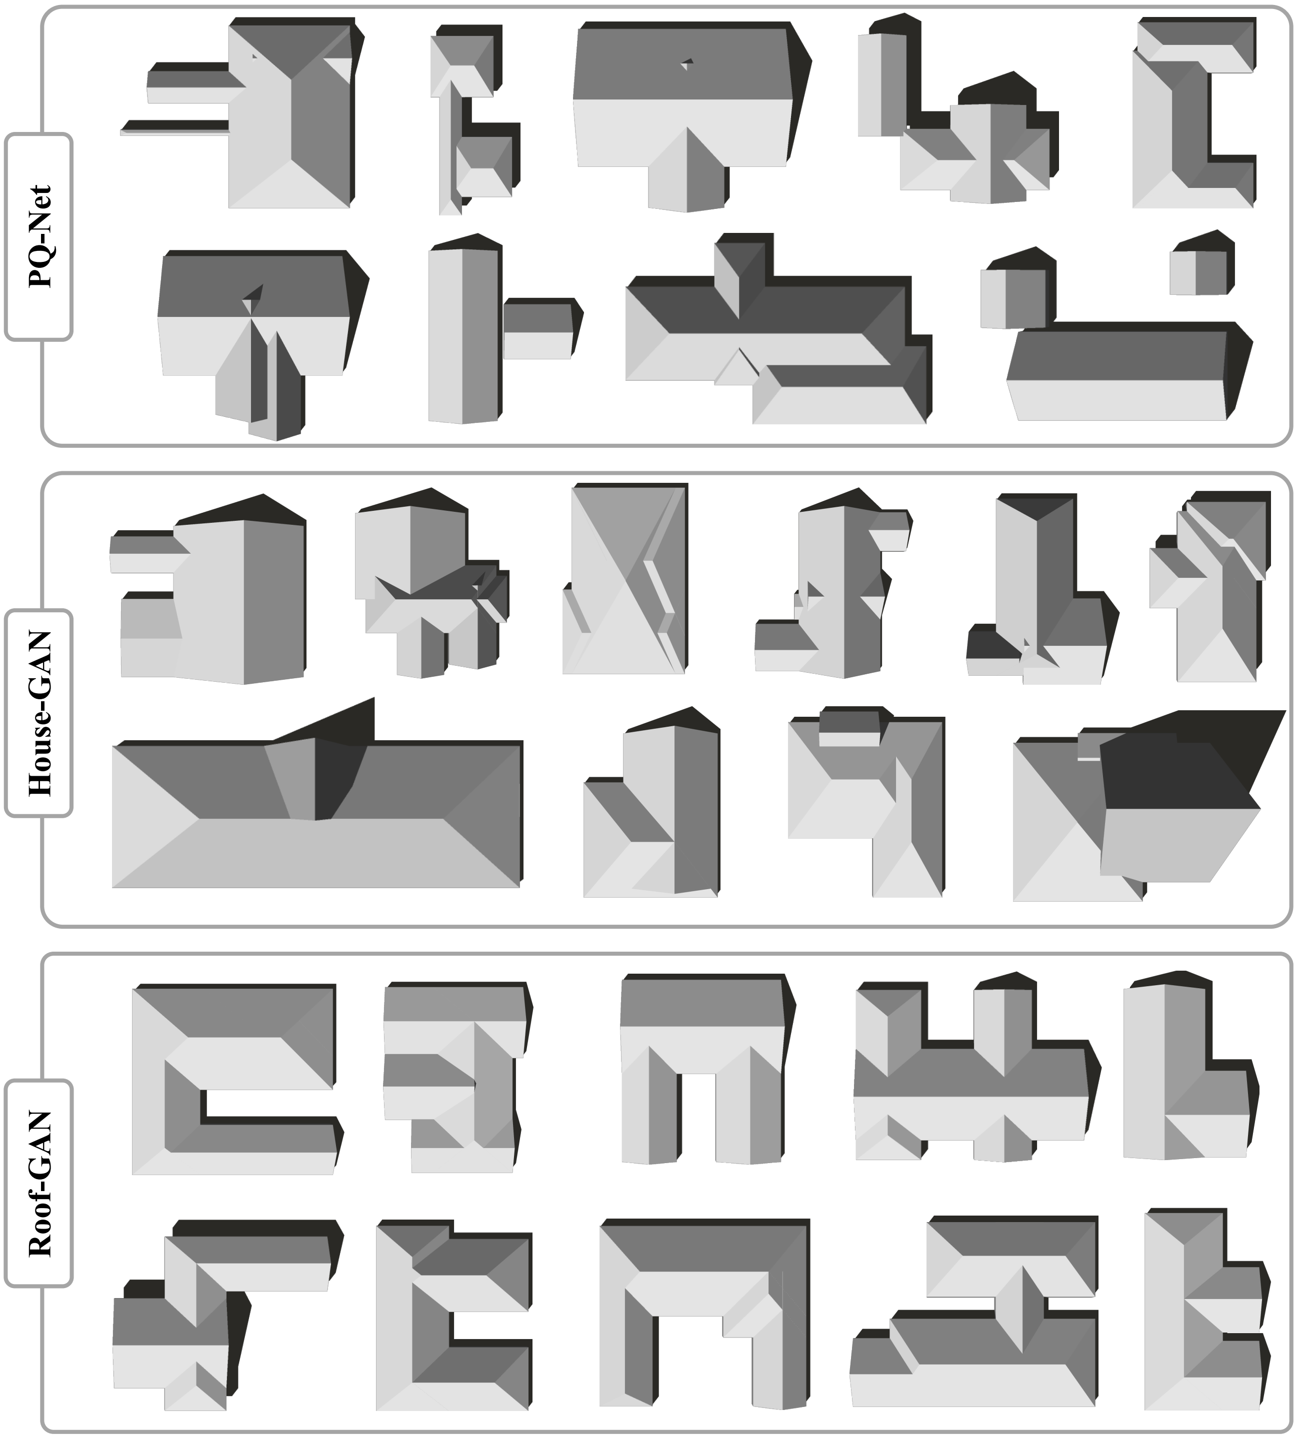

Figure 1. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

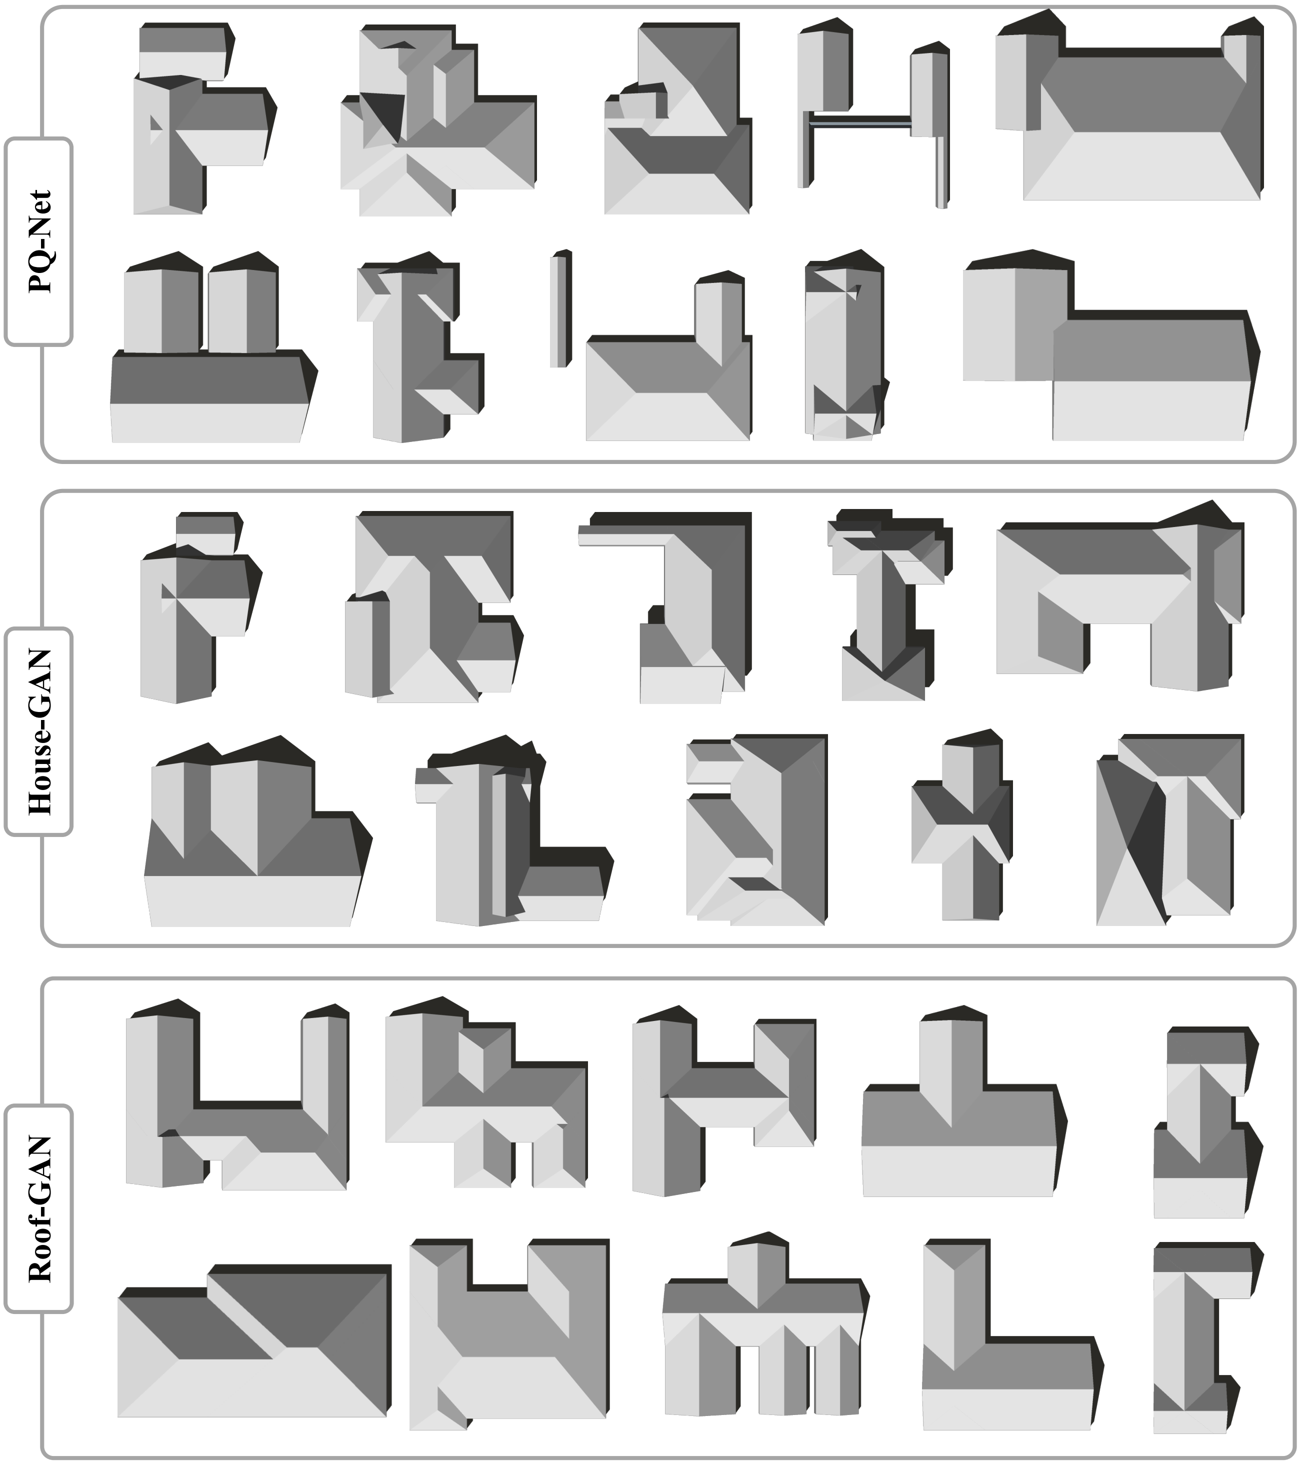

Figure 2. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

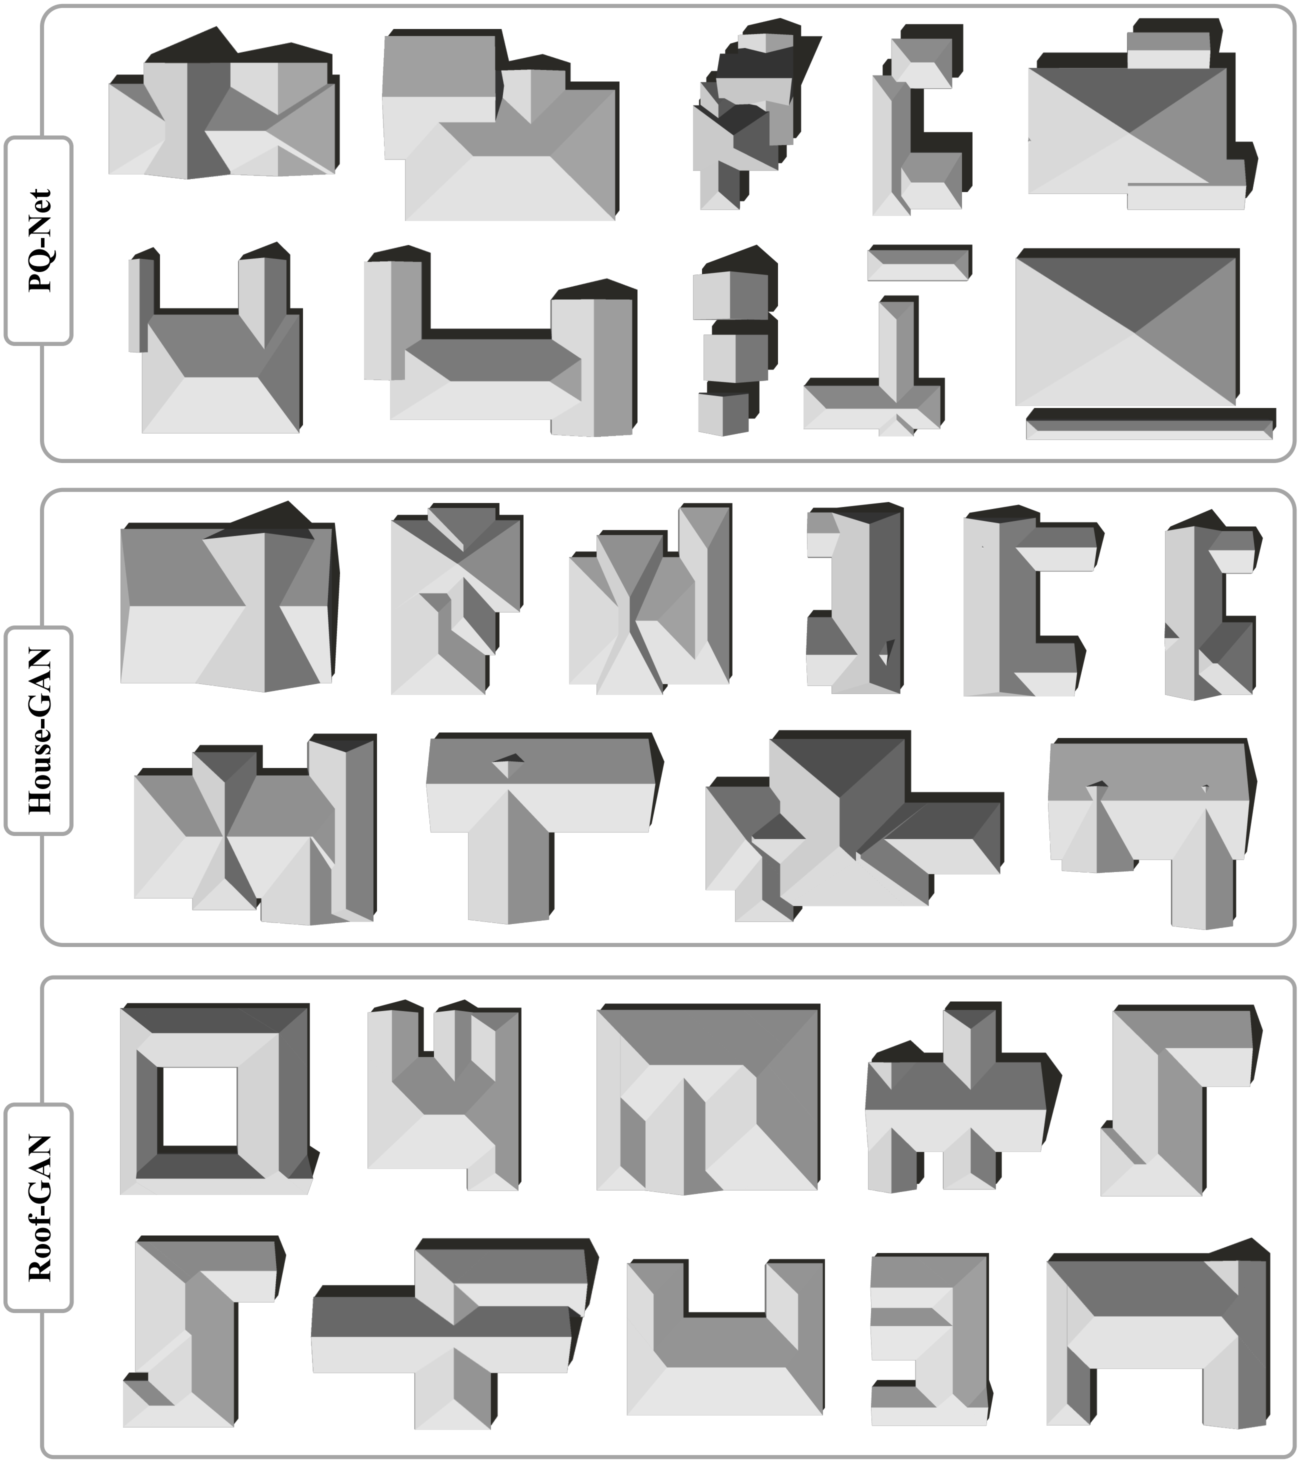

Figure 3. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

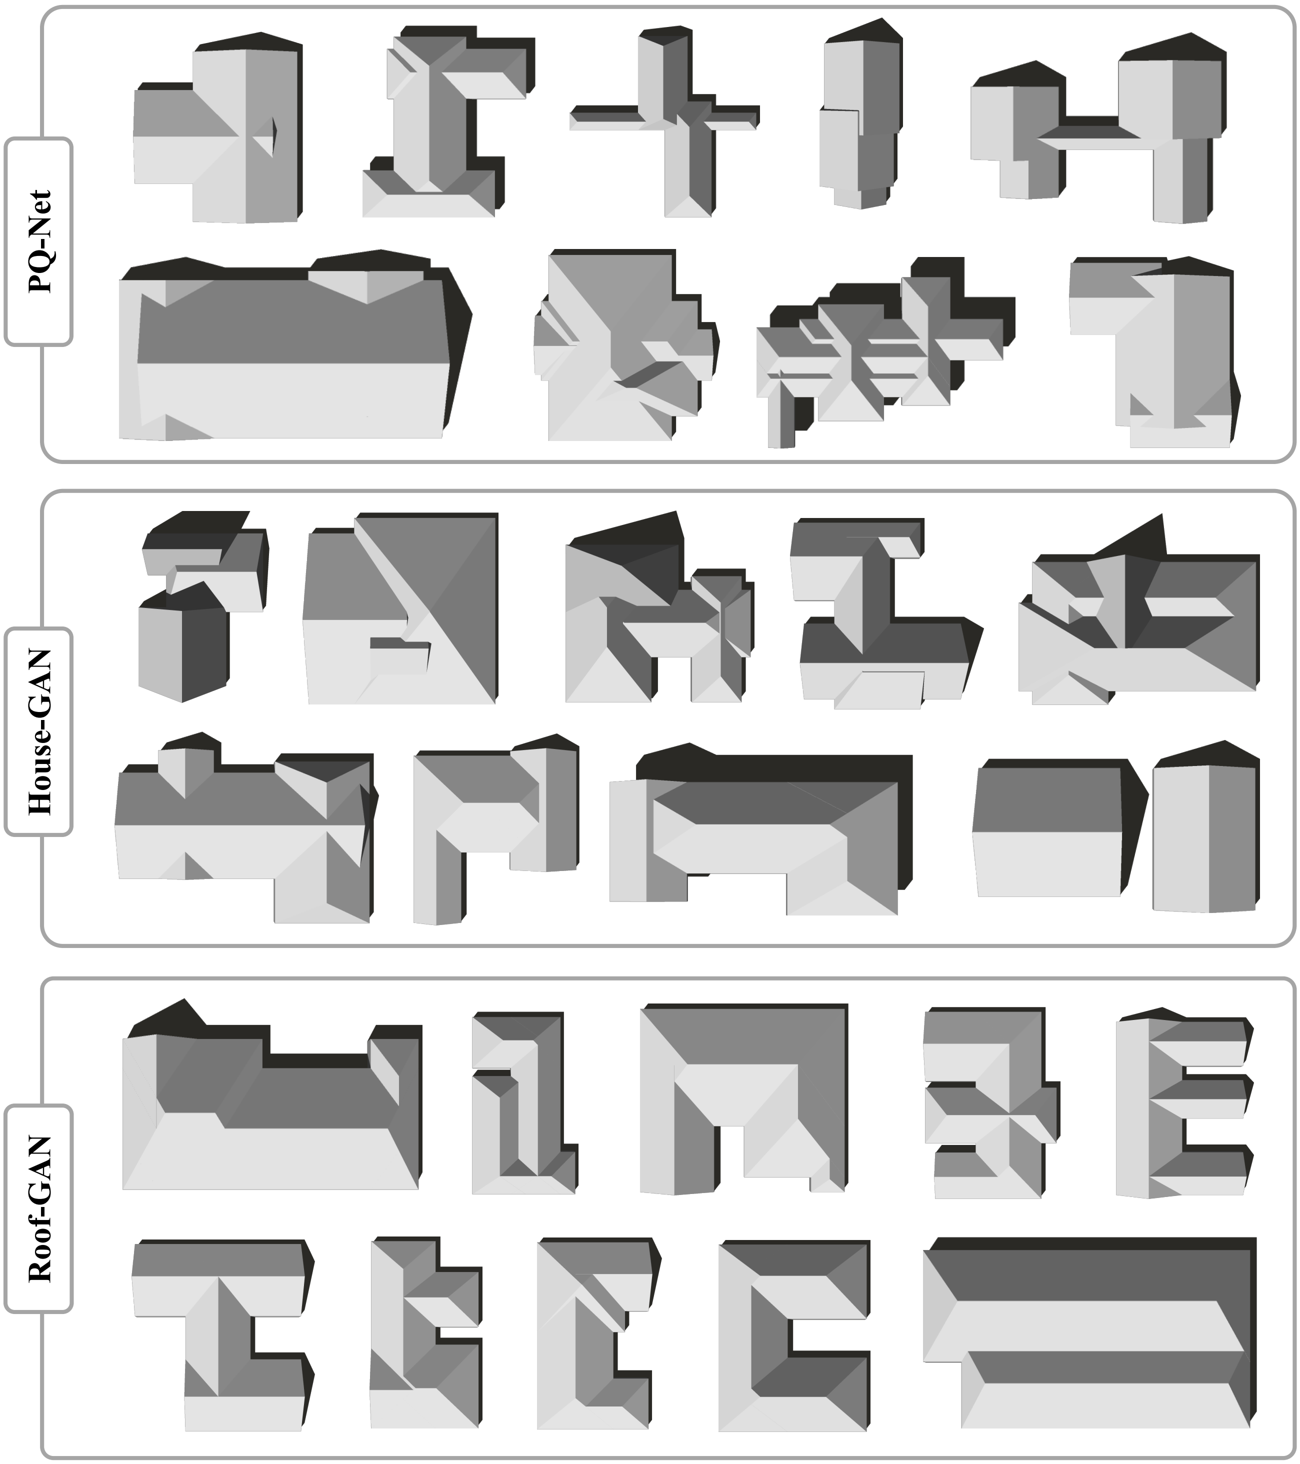

Figure 4. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

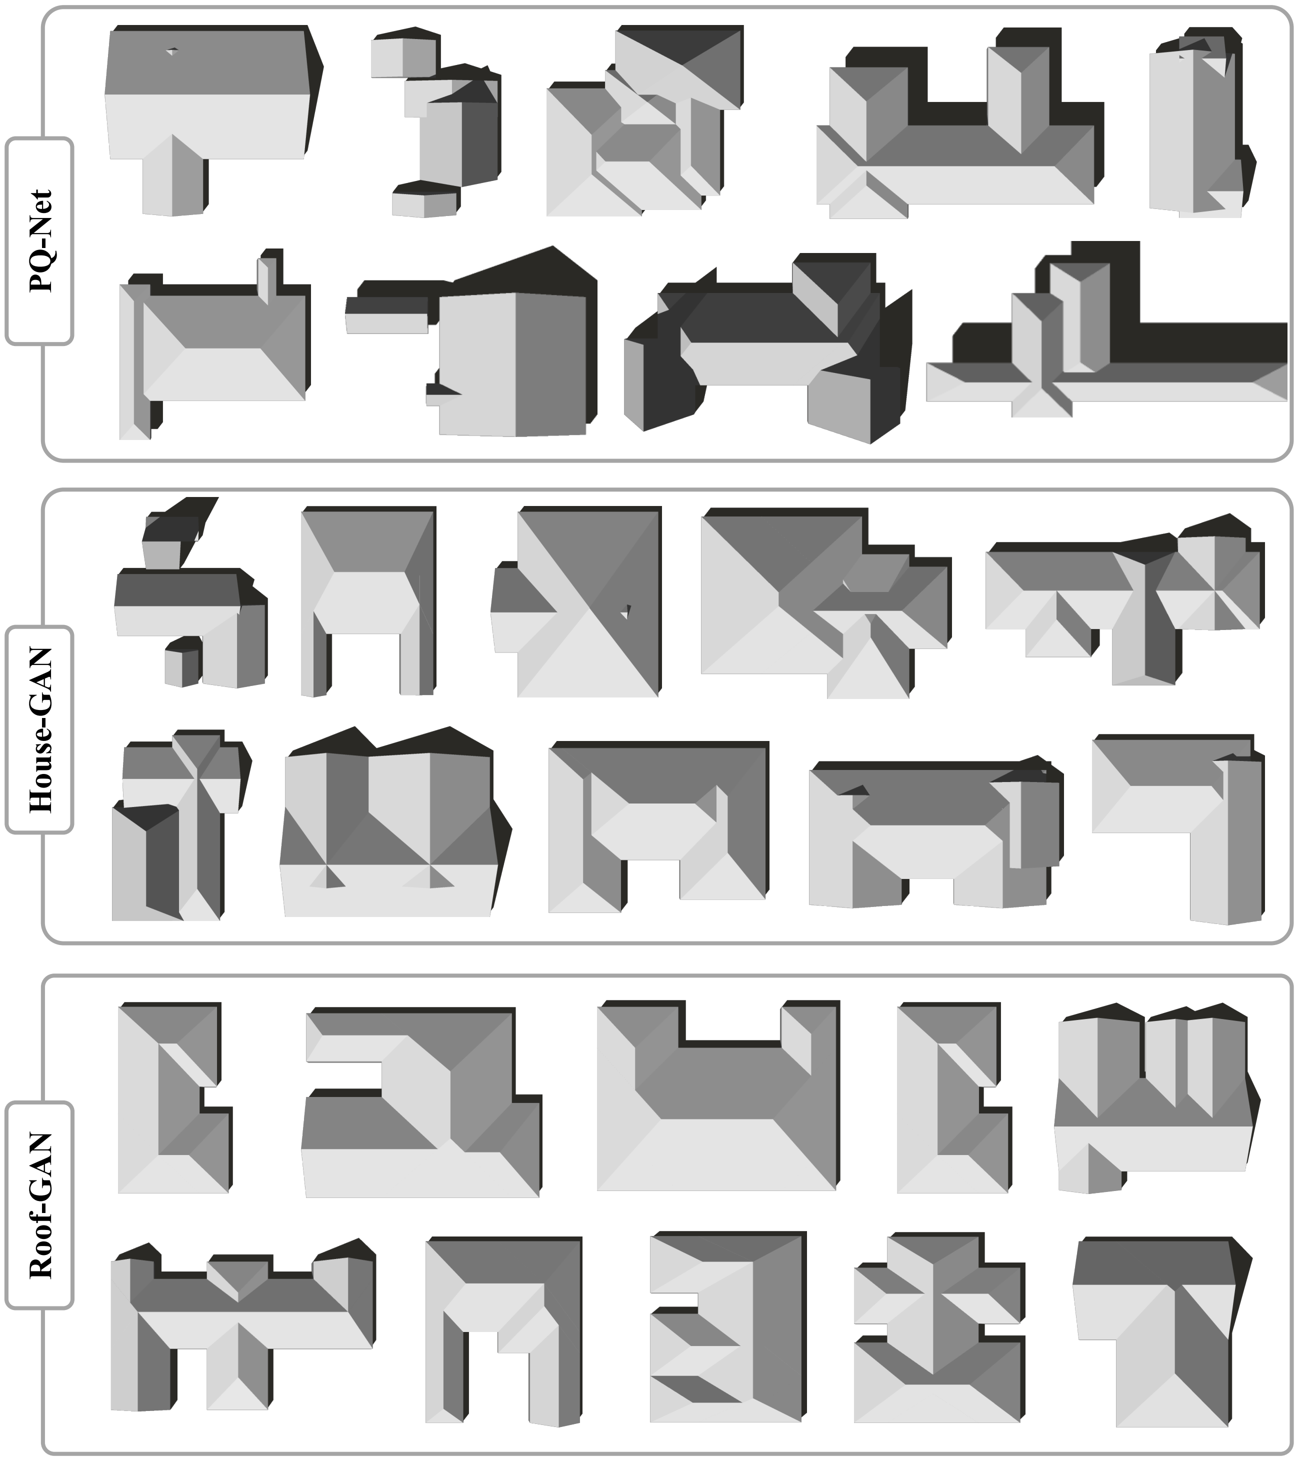

Figure 5. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

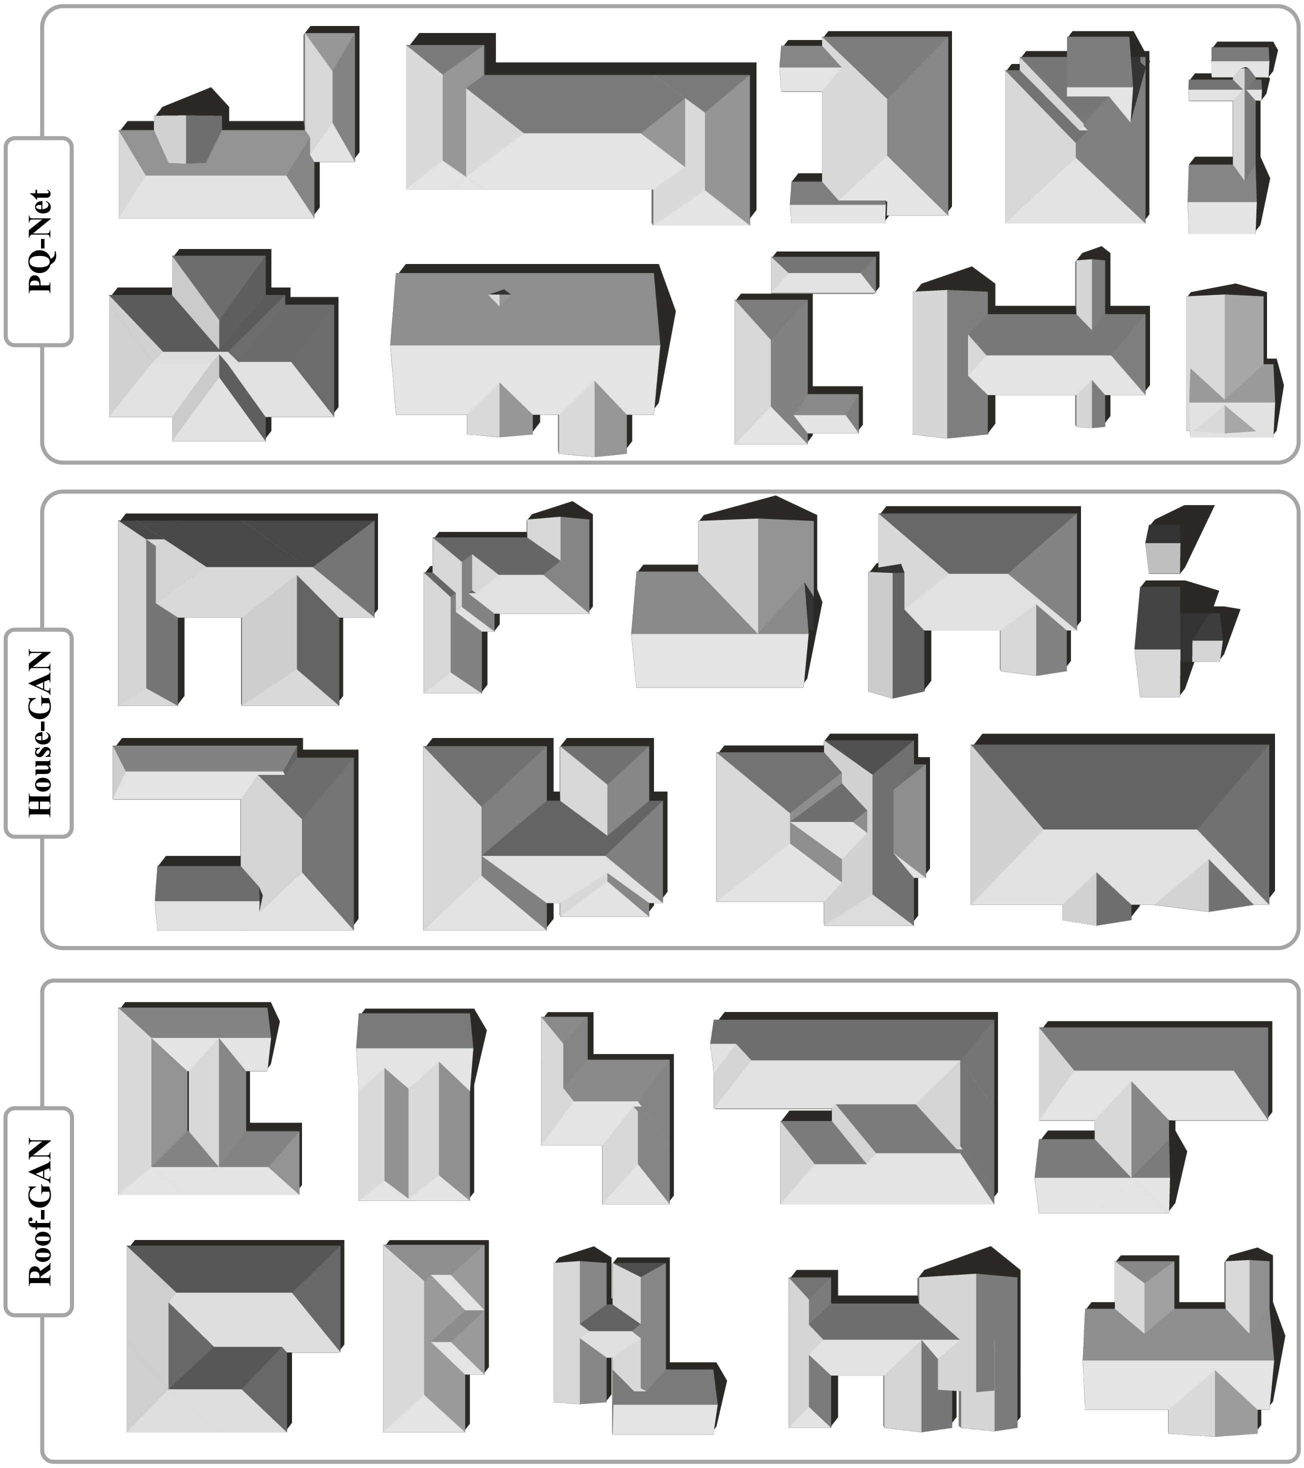

Figure 6. More qualitative evaluations among PQ-Net, House-GAN and Roof-GAN.

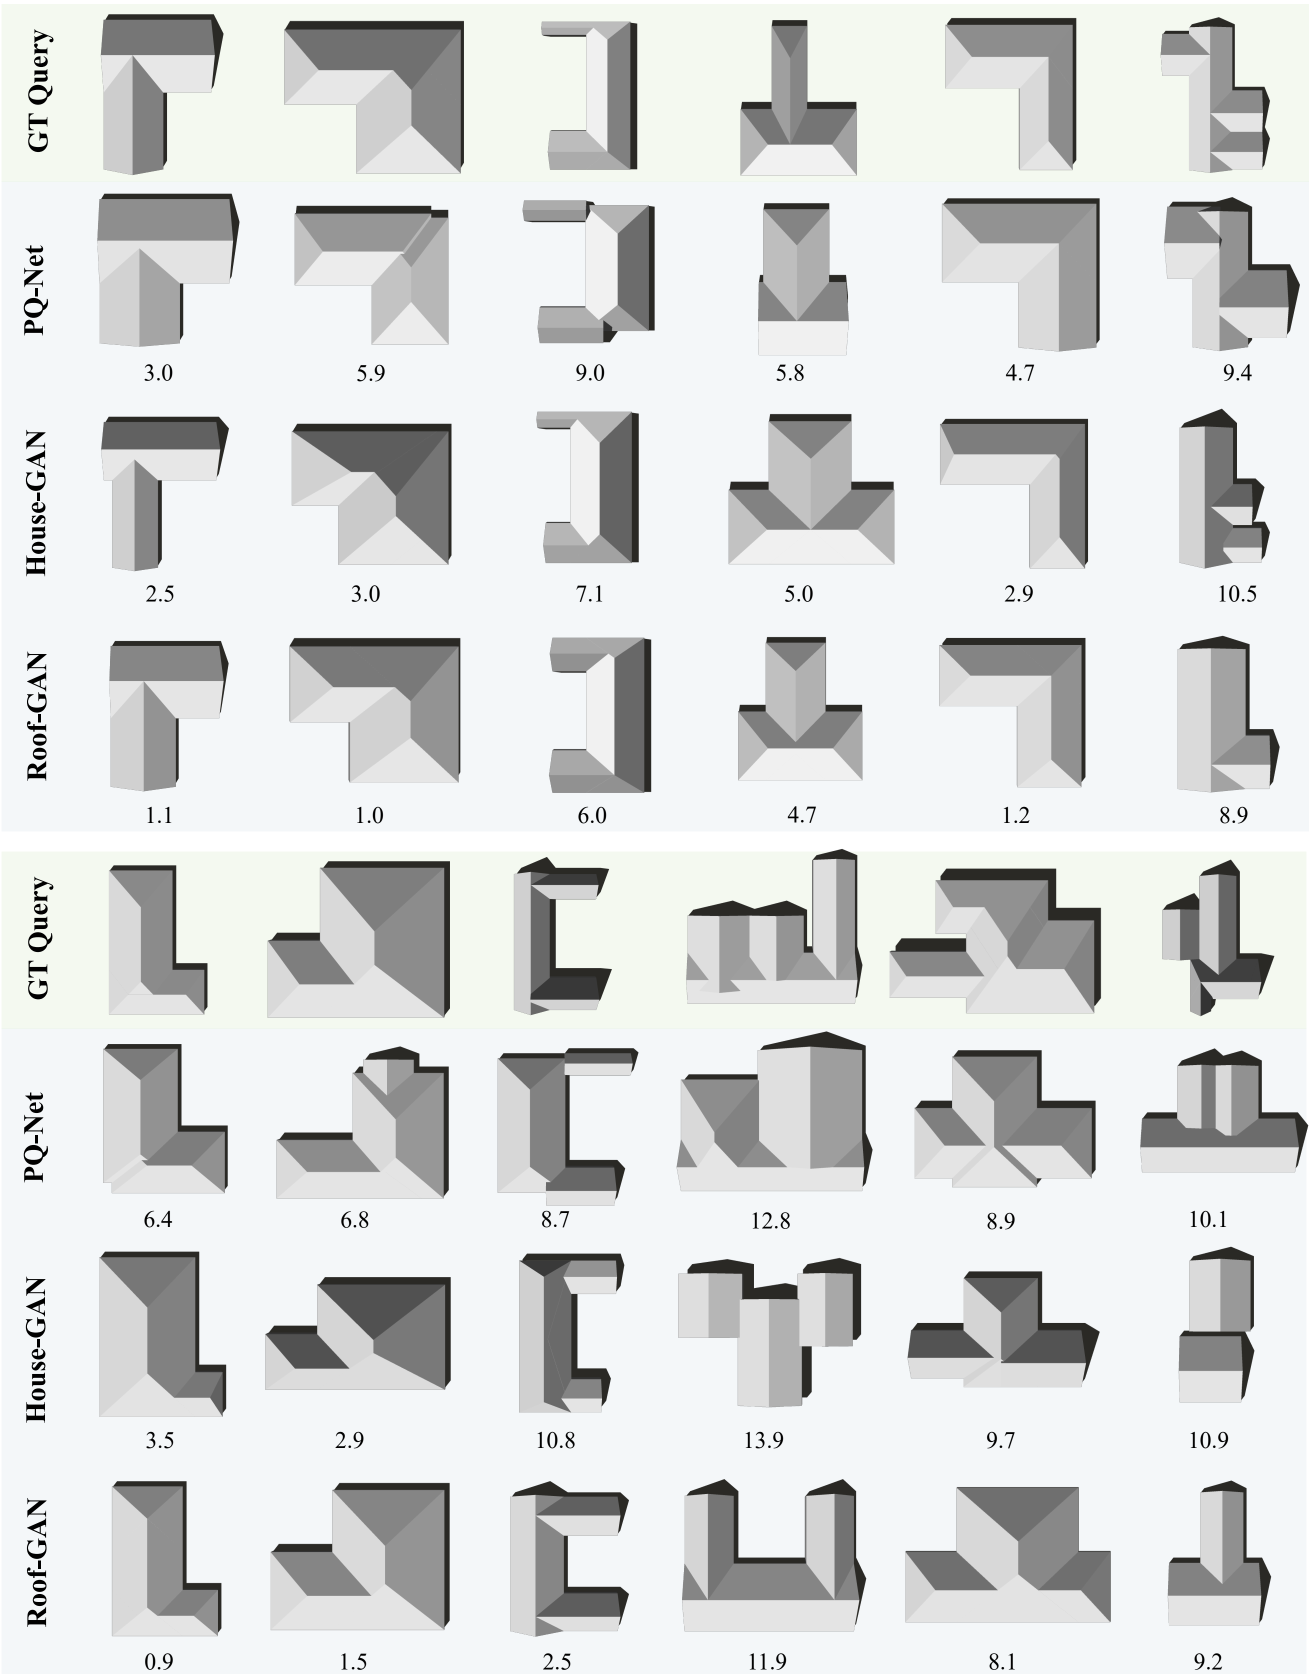

Figure 7. More results of RMMD-based model retrieval evaluations.

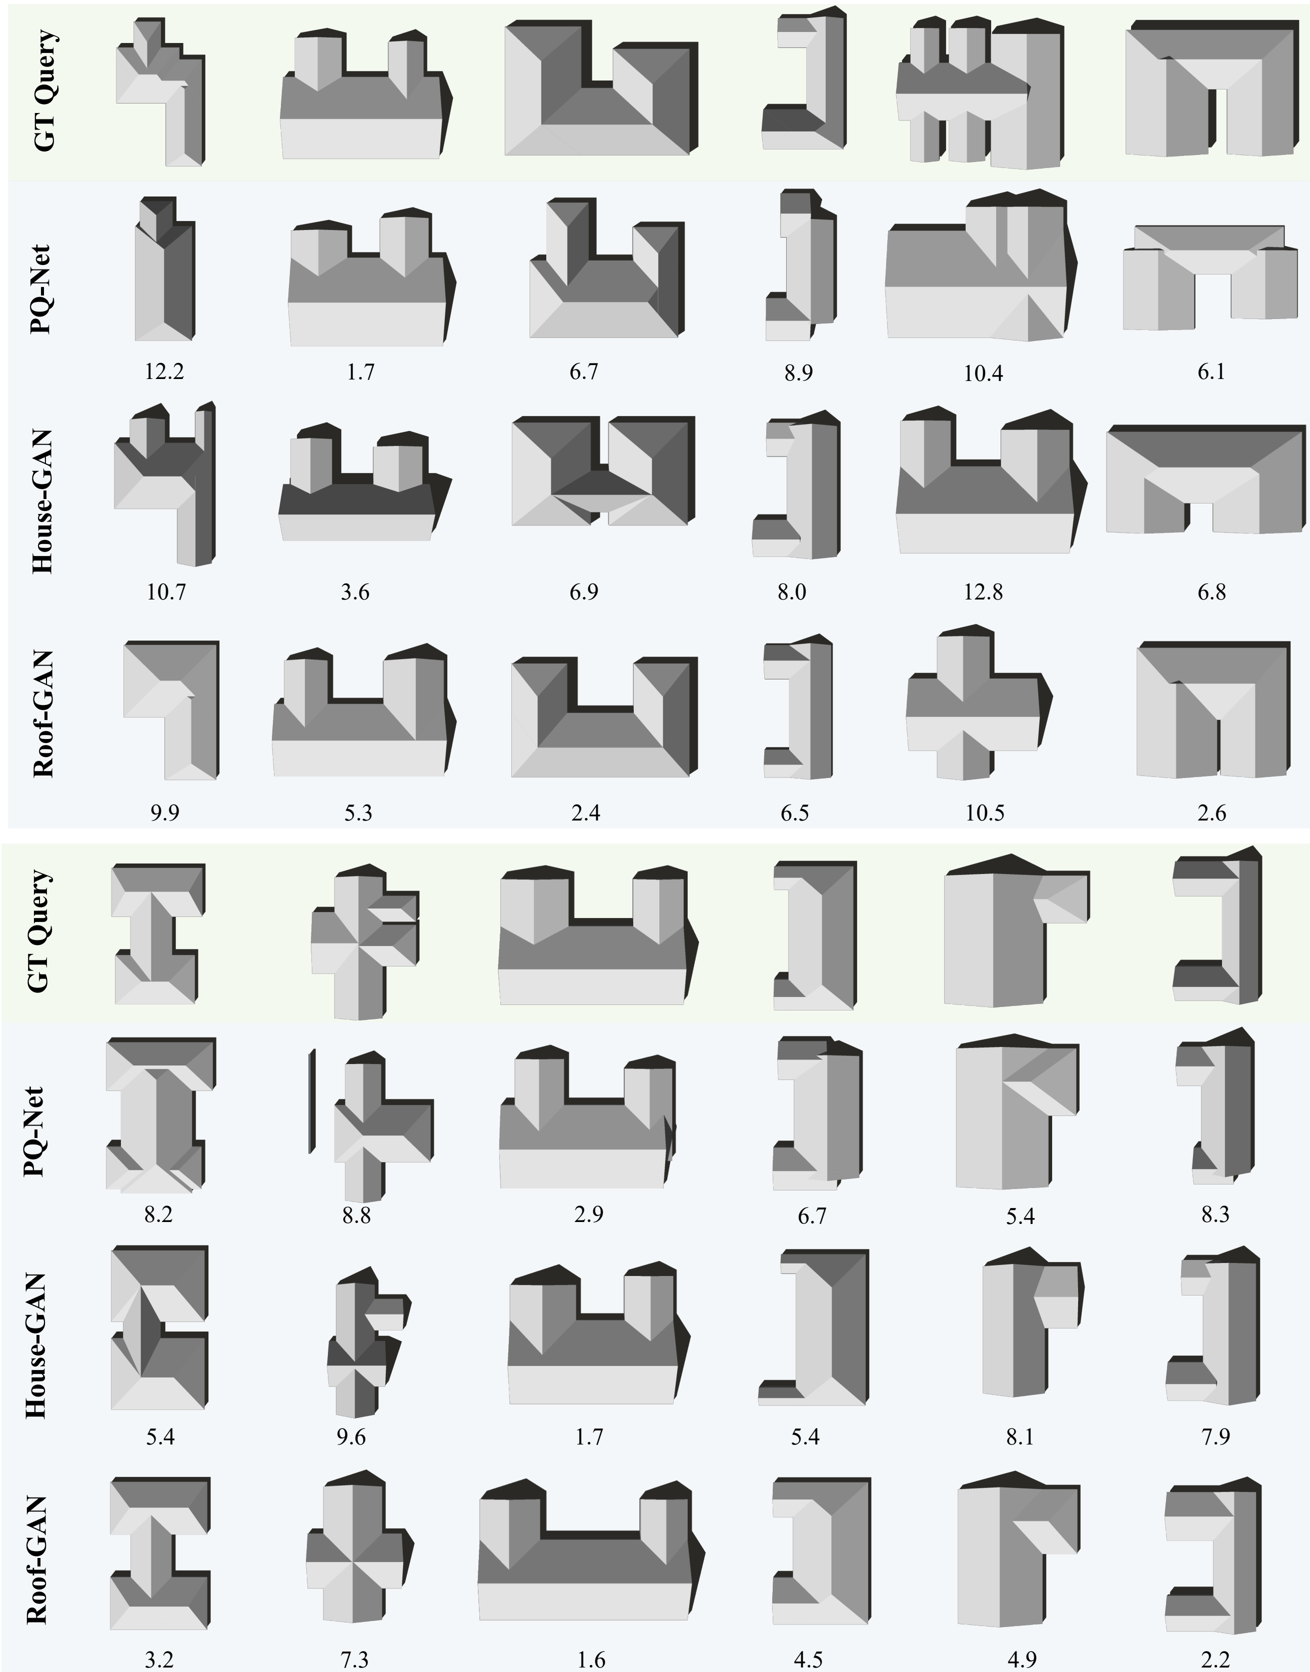

Figure 8. More results of RMMD-based model retrieval evaluations.