Supplementary Material for Paper:

Rent3D: Floor-Plan Priors for Monocular Layout Estimation

Chenxi Liu[1] * Alexander G. Schwing [2,*] Kaustav Kundu[2] Raquel Urtasun[2] Sanja Fidler[2]

[1]State Key Lab. on Intelligent Technology and Systems
Tsinghua Nat. Lab. for Inf. Science and Tech. (TNList)
Department of Automation, Tsinghua University

[2]Department of Computer Science
University of Toronto

chenxi_liu@live.cn, {aschwing, kkundu, urtasun, fidler}@cs.toronto.edu

In the Supplementary Material we first provide the derivations used in our model. We then provide several additional results: both the 3D reconstructions as well as layout and layout-to-floor-plan alignment.

1. Geometry of Layout with Floor-Plans

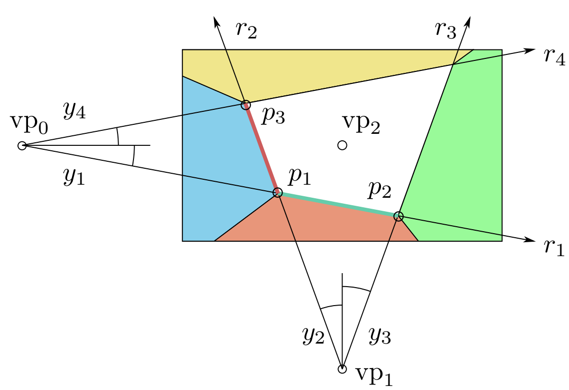

Figure 1. Parameterization of our layout task. Using projective geometry, the parameter y_4 can be computed from the three parameters $y_1, \ldots, y_3$ and a known apartment height h.

It is well known that the camera intrinsic matrix K and the rotation R can be recovered from three orthogonal vanishing points [1]. Given K and R we can compute a homography for each face of the room, i.e., a projection which maps the face defined by two vanishing points vp_i and vp_j to fronto-parallel view:

$$H_{ij} = K \cdot R_{ij} \cdot K^{-1}.$$

Here R_{ij} denotes a column-wise permutation of R, shuffling the columns i and j to be column 1 and 2.

1.1. Computing y_4 from y_1, y_2 and y_3

Let a be a known aspect ratio defined as w/h, where h is the height and w the width of the front wall in the physical world (we know these dimensions from the rental site). Then, given two corner points $\mathbf{p_1}$ and $\mathbf{p_2}$ of a wall, defined by the two intersections of rays y_1, y_2 and y_1, y_3, respectively, we can compute the corner point $\mathbf{p_3}$ defined by the intersection of rays y_2 and y_4. Knowing $\mathbf{p_3}$ we easily obtain the ray y_4. We now show how to compute $\mathbf{p_3}$.

*Denotes equal contribution

We want a point $\mathbf{p_3}$ on a ray $(\mathbf{vp_1}, \mathbf{p_1})$ such that the aspect ratio $\frac{||\hat{\mathbf{p}}_2 - \hat{\mathbf{p}}_1||}{||\hat{\mathbf{p}}_3 - \hat{\mathbf{p}}_1||}$ equals a pre-specified value a. Here $\hat{\mathbf{p}}$ denotes a point in the fronto-parallel view, i.e.,

$$\hat{\mathbf{p}} = \frac{H \cdot \mathbf{p}}{H_{3:} \cdot \mathbf{p}},$$

where $H_{3:}$ denotes the third row of matrix H. Note that $||\hat{\mathbf{p}}_2 - \hat{\mathbf{p}}_1||$ is the width of the wall in the fronto-parallel view, and $||\hat{\mathbf{p}}_3 - \hat{\mathbf{p}}_1||$ refers to its height.

Since $\mathbf{p_1}$ and $\mathbf{p_2}$ lie on a line through $\mathbf{vp_0}$, let's express $\mathbf{p_2}$ as follows:

$$\mathbf{p_2} = \mathbf{p_1} + \lambda(\mathbf{p_1} - \mathbf{vp_0}), \tag{1}$$

where λ is easily computable from the two points.

In rectified coordinates

$$\hat{\mathbf{p}}_2 = \frac{H \cdot \mathbf{p_2}}{H_{3:} \cdot \mathbf{p_2}} = \frac{H \cdot (\mathbf{p_1} + \lambda(\mathbf{p_1} - \mathbf{vp_0}))}{H_{3:}(\mathbf{p_1} + \lambda(\mathbf{p_1} - \mathbf{vp_0}))} \overset{H_{3:}\mathbf{vp_0}=0}{=} \frac{(1+\lambda)H\mathbf{p_1} - \lambda \cdot H\mathbf{vp_0}}{(1+\lambda)H_{3:} \cdot \mathbf{p_1}} \tag{2}$$

$$= \frac{H \cdot \mathbf{p_1}}{H_{3:} \cdot \mathbf{p_1}} - \lambda \frac{H\mathbf{vp_0}}{(1+\lambda)H_{3:}\mathbf{p_1}}, \tag{3}$$

where we have used a shorthand $H = H_{01}$. The width in rectified view is therefore given by

$$w = ||\hat{\mathbf{p}}_2 - \hat{\mathbf{p}}_1|| = \left|\frac{\lambda}{1+\lambda}\right| \frac{||H\mathbf{vp_0}||}{|H_{3:}\mathbf{p_1}|}.$$

Since $\mathbf{p_3}$ lies on a line through $\mathbf{vp_1}$ and $\mathbf{p_1}$, let

$$\mathbf{p_3} = \mathbf{p_1} + \mu(\mathbf{p_1} - \mathbf{vp_1}). \tag{4}$$

As before we can get $\mathbf{p_3}$ in the rectified view:

$$\hat{\mathbf{p}}_3 = \frac{H \cdot \mathbf{p_1}}{H_{3:} \cdot \mathbf{p_1}} - \mu \frac{H\mathbf{vp_1}}{(1+\mu)H_{3:}\mathbf{p_1}},$$

and thus the height is given by

$$h = ||\hat{\mathbf{p}}_3 - \hat{\mathbf{p}}_1|| = \left|\frac{\mu}{1+\mu}\right| \frac{||H\mathbf{vp_1}||}{|H_{3:}\mathbf{p_1}|}.$$

Since we know that $w = a \cdot h$, we obtain the equality

$$\left|\frac{\lambda}{1+\lambda}\right| \cdot ||H\mathbf{vp_0}|| = a \cdot \left|\frac{\mu}{1+\mu}\right| \cdot ||H\mathbf{vp_1}||. \tag{5}$$

Let's define:

$$f_{10} = \frac{||H\mathbf{vp_1}||}{||H\mathbf{vp_0}||},$$

which is a value that can be precomputed prior to any computation with our model. Redefining our aspect ratio via $\tilde{a} = a \cdot f_{10}$, again a value that can be precomputed, we obtain the following rule after simple manipulation of Eq. (5):

$$\mu = \frac{1}{\tilde{a} \cdot \left|1 + \frac{1}{\lambda}\right| - 1}, \quad \mu > 0, \tag{6}$$

$$\mu = \frac{1}{-\tilde{a} \cdot \left|1 + \frac{1}{\lambda}\right| - 1}, \quad \mu < 0. \tag{7}$$

Note that a negative denominator for $\mu > 0$ would result in a rectangle for the front wall which is not valid (it crosses the infinity point). We now have μ, which defines $\mathbf{p_3}$, from which we can easily obtain the fourth ray y_4.

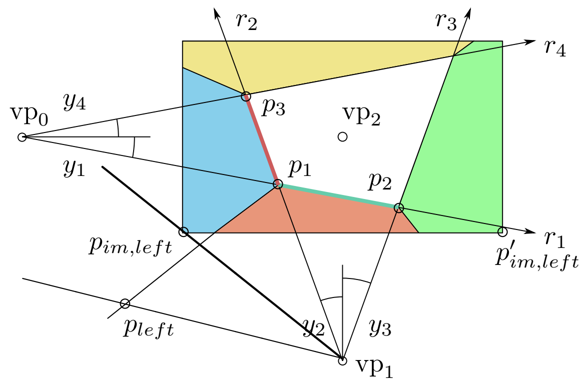

Figure 2. A ray through vp_1 and the point p_{left}, the left back coordinate of the floor, needs to fully lie outside the image in order to define a proper room with known aspect ratios for the walls. This imposes restrictions on where p_2 can be.

1.2. Computing the Left and Right Wall

The previous section gives us the full front wall given the three rays y_1, y_2 and y_3. However, to get a valid room hypothesis we also need to make use of the aspect ratio of the floor given via the floor-plan. A room hypothesis is valid if the two corner points corresponding to the intersection of the (left, ground, back) walls and (right, ground, back) walls, fall outside the image. This imposes a restriction that the camera is placed inside the room.

Let's first compute the left back corner, i.e., the corner $\mathbf{p_{left}}$ corresponding to the intersection of the (left, ground, back) walls (Fig. 2), given the aspect ratio a_{floor}. We define a_{floor} as the aspect ratio of the floor, i.e., the width of the front wall divided by the width of the left wall. A correct configuration must place the ray through $\mathbf{vp_1}$ and $\mathbf{p_{left}}$ to lie outside of the image (to the left of the ray through $\mathbf{vp_1}$ and $[1, \mathrm{imheight}]$).

We again parametrize the point via the ray (line) it lies on, i.e.,

$$\mathbf{p_{left}} = \mathbf{p_1} + \mu_{left}(\mathbf{p_1} - \mathbf{vp_2}).$$

Employing the same derivation as in the previous subsection we obtain

$$\mu_{left} = \frac{1}{\tilde{a}_{floor} \cdot \left|1 + \frac{1}{\lambda}\right| - 1}, \quad \mu_{left} > 0$$
$$\mu_{left} = \frac{1}{-\tilde{a}_{floor} \cdot \left|1 + \frac{1}{\lambda}\right| - 1}, \quad \mu_{left} < 0$$

where:

$$\tilde{a}_{floor} = a_{floor} \cdot f_{12}, \quad \text{and} \quad f_{12} = \frac{||H_{12}\mathbf{vp_1}||}{||H_{12}\mathbf{vp_2}||}.$$

Note that H_{12} is a homography defined by the pair of vanishing points $\mathbf{vp_1}$ and $\mathbf{vp_2}$.

With a very similar derivation, we can also compute the point $\mathbf{p_{right}}$, a corner of the room where the (right, floor, back) walls meet.

To get a valid room all we need to check is whether the ray from $\mathbf{vp_1}$ to $\mathbf{p_{left}}$, and a ray $\mathbf{vp_1}$ to $\mathbf{p_{right}}$ fall outside the image. The easiest to check this is illustrated in the following. Let $\mathbf{p_{im,left}}$ be the corner of the image with $x = 1$, that is the closest to $\mathbf{vp_1}$, and let $\mathbf{p'_{im,left}} = [\mathrm{imwidth}, \mathbf{p_{im,left}}(\mathbf{y})]$. Figure 2 illustrates the notation.

We can compute the line that goes through $\mathbf{vp_1}$ and $\mathbf{p_{im,left}}$ using homogeneous coordinates as follows: $\mathbf{n_{left}} = \mathbf{vp_2} \times \mathbf{p_{im,left}}$. Then our room hypothesis is valid if

$$\mathrm{sign}\Big(\mathbf{n_{left}}^T \cdot (\mathbf{p_{left}} - \mathbf{vp_1})\Big) = -\mathrm{sign}\Big(\mathbf{n_{left}}^T \cdot (\mathbf{p'_{im,left}} - \mathbf{vp_1})\Big). \tag{8}$$

This just checks that both $\mathbf{p_{left}}$ and $\mathbf{p'_{im,left}}$ are on opposite sides of the line that goes through $\mathbf{vp_1}$ and $\mathbf{p_{im,left}}$. A similar constraint follows for the right point.

Note that $\mathbf{n_{left}}$ can be precomputed, and likewise $s := -\mathrm{sign}\Big(\mathbf{n_{left}}^T \cdot (\mathbf{p'_{im,left}} - \mathbf{vp_1})\Big)$. Assume for the moment that $s > 0$. We will also assume $\lambda > 0$, which can easily be imposed in our parametrization of the layout problem. Then, to get a valid room, we want to find a constraint on λ such that

$$\mathbf{n_{left}}^T \cdot (\mathbf{p_{left}} - \mathbf{vp_1}) > 0.$$

Plugging in $\mathbf{p_{left}} = \mathbf{p_1} + \mu_{left}(\mathbf{p_1} - \mathbf{vp_2})$, we obtain:

$$n_{left}(x) \cdot \Big[p_1(x) + \mu_{left}\Big(p_1(x) - vp_2(x)\Big) - vp_1(x)\Big] + n_{left}(y) \cdot \Big[p_1(y) + \mu_{left}\Big(p_1(y) - vp_2(y)\Big) - vp_1(y)\Big] > 0,$$

which is equivalent to

$$\mu_{left} \underbrace{\Big[n_x\Big(p_1(x) - vp_2(x)\Big) + n_y\Big(p_1(y) - vp_2(y)\Big)\Big]}_{\text{Define this as: } c_x} + \underbrace{n_{left}(x)\Big(p_1(x) - vp_1(x)\Big) + n_{left}(y)\Big(p_1(y) - vp_1(y)\Big)}_{\text{Define this as: } c_y} > 0.$$

Note that c_x and c_y depend only on $\mathbf{p_1}$ and are easily computable. All in all, we require

$$\mu_{left} \cdot c_x + c_y > 0.$$

We separate two cases:

- $c_x > 0$

$$\mu_{left} > -\frac{c_y}{c_x}.$$

We plug in: $\mu_{left} = \frac{1}{\tilde{a}_{floor}\left(1+\frac{1}{\lambda}\right)-1}$. Given our parametrization we assume $\mu_{left} > 0$. So if $-\frac{c_y}{c_x} < 0$, then the above constraint is satisfied for all $\lambda > 0$. If $-\frac{c_y}{c_x} > 0$, then we have:

$$\frac{1}{\lambda} < \frac{-\frac{c_x}{c_y} + 1}{\tilde{a}_{floor}} - 1.$$

If $\frac{-\frac{c_x}{c_y}+1}{\tilde{a}_{floor}} - 1 < 0$, then the constraint is never satisfied. If $\frac{-\frac{c_x}{c_y}+1}{\tilde{a}_{floor}} - 1 > 0$, then:

$$\lambda > \frac{\tilde{a}_{floor}}{-\frac{c_x}{c_y} + 1 - \tilde{a}_{floor}}.$$

- $c_x < 0$

$$\mu_{left} < -\frac{c_y}{c_x}.$$

Notice that if $c_y < 0$, then this constraint is never satisfied. If $c_y > 0$,

$$\frac{1}{\lambda} > \frac{-\frac{c_x}{c_y} + 1}{\tilde{a}_{floor}} - 1.$$

If $\frac{-\frac{c_x}{c_y}+1}{\tilde{a}_{floor}} - 1 < 0$, then this constraint is satisfied for all λ. If $\frac{-\frac{c_x}{c_y}+1}{\tilde{a}_{floor}} - 1 > 0$, then:

$$\lambda < \frac{\tilde{a}_{floor}}{-\frac{c_x}{c_y} + 1 - \tilde{a}_{floor}}.$$

Note that in both cases the bound on λ is function of $\frac{c_x}{c_y}$, and both c_x and c_y are linear functions of $\mathbf{p_1}$. It's easy to show that the bound is thus either a convex or a concave function of $\mathbf{p_1}$. This bound can be used to carve the space of possible layouts faster in our branch and bound inference of the layout problem [3, 2]. In particular, at each branching step, we get an interval for y_1 and an interval for y_2. This defines a convex area (a quadrilateral) for the point $\mathbf{p_1}$. Clearly, if at least one corner point of this quadrilateral lies to the left of line that goes through $\mathbf{vp_1}$ and $\mathbf{p_{im,left}}$, then our $\mathbf{p_{left}}$ will always be to the left of this line as well, and thus outside the image. Such cases do not impose any restrictions on λ. If the quadrilateral for $\mathbf{p_1}$ lies full on the right side of this line, then we need to inspect the bound on λ. Since the bound is a convex or concave function of $\mathbf{p_1}$ and the region for $\mathbf{p_1}$ is convex, we only need to compute the bound in the corner points of this region. Since in each step of our branch and bound inference we need to find the biggest possible wall in order to compute the upper bound on our optimization function, and the smaller wall to compute the lower bound, we can either take the least restrictive or the most restrictive bound on λ. This essentially tells us where the point $\mathbf{p_2}$ could possibly be in order to ensure a valid room configuration. We then branch $\mathbf{p_2}$ only within the allowed interval.

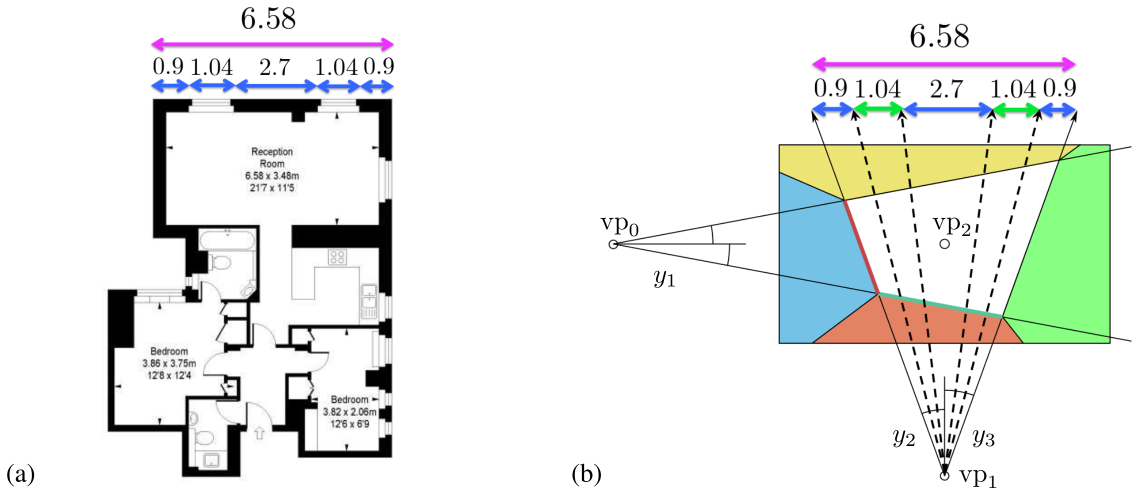

Figure 3. **(a)** Example of a floor-plan with window-to-wall ratios (top), **(b)** The window-to-wall ratios in the floor plan enable us to compute the rays (dashed lines) for the windows given y_2 and y_3 in our layout problem. The windows are denoted with green lines.

1.3. Window rays given p_1 and p_2

To complete our derivations we still need to compute the vertical rays (rays going through vp_1) that define the regions for the windows. This can be done through the floor-plan in which we are given the ratios of the lengths for the window given the wall as illustrated in Fig. 3(a). In order to cast our rays we make use of the cross-ratio perspective invariant as shown in Fig. 3(b).

2. Visualizations

In the following pages we show success cases and failure modes in the following two Tables respectively. Specifically, we illustrate the generated 3D visualizations corresponding to different instantiations of our model. Below each 3D plot we provide the numbers of correctly matched images and walls. Below we also visualize the layout prediction and the estimated layout to floor-plan alignment for each image. Following the main submission, dashed lines in both the image as well as the floor plan depict the groundtruth, while solid lines and semi-transparent image overlays indicate the prediction.

References

[1] C. Rother. A new approach for vanishing point detection in architectural environments. In *Proc. BMVC*, 2000. 1

[2] A. G. Schwing, S. Fidler, M. Pollefeys, and R. Urtasun. Box in the box: Joint 3d layout and object reasoning from single images. In *Proc. ICCV*, 2013. 4

[3] A. G. Schwing and R. Urtasun. Efficient Exact Inference for 3D Indoor Scene Understanding. In *Proc. ECCV*, 2012. 4

2.1. Success cases

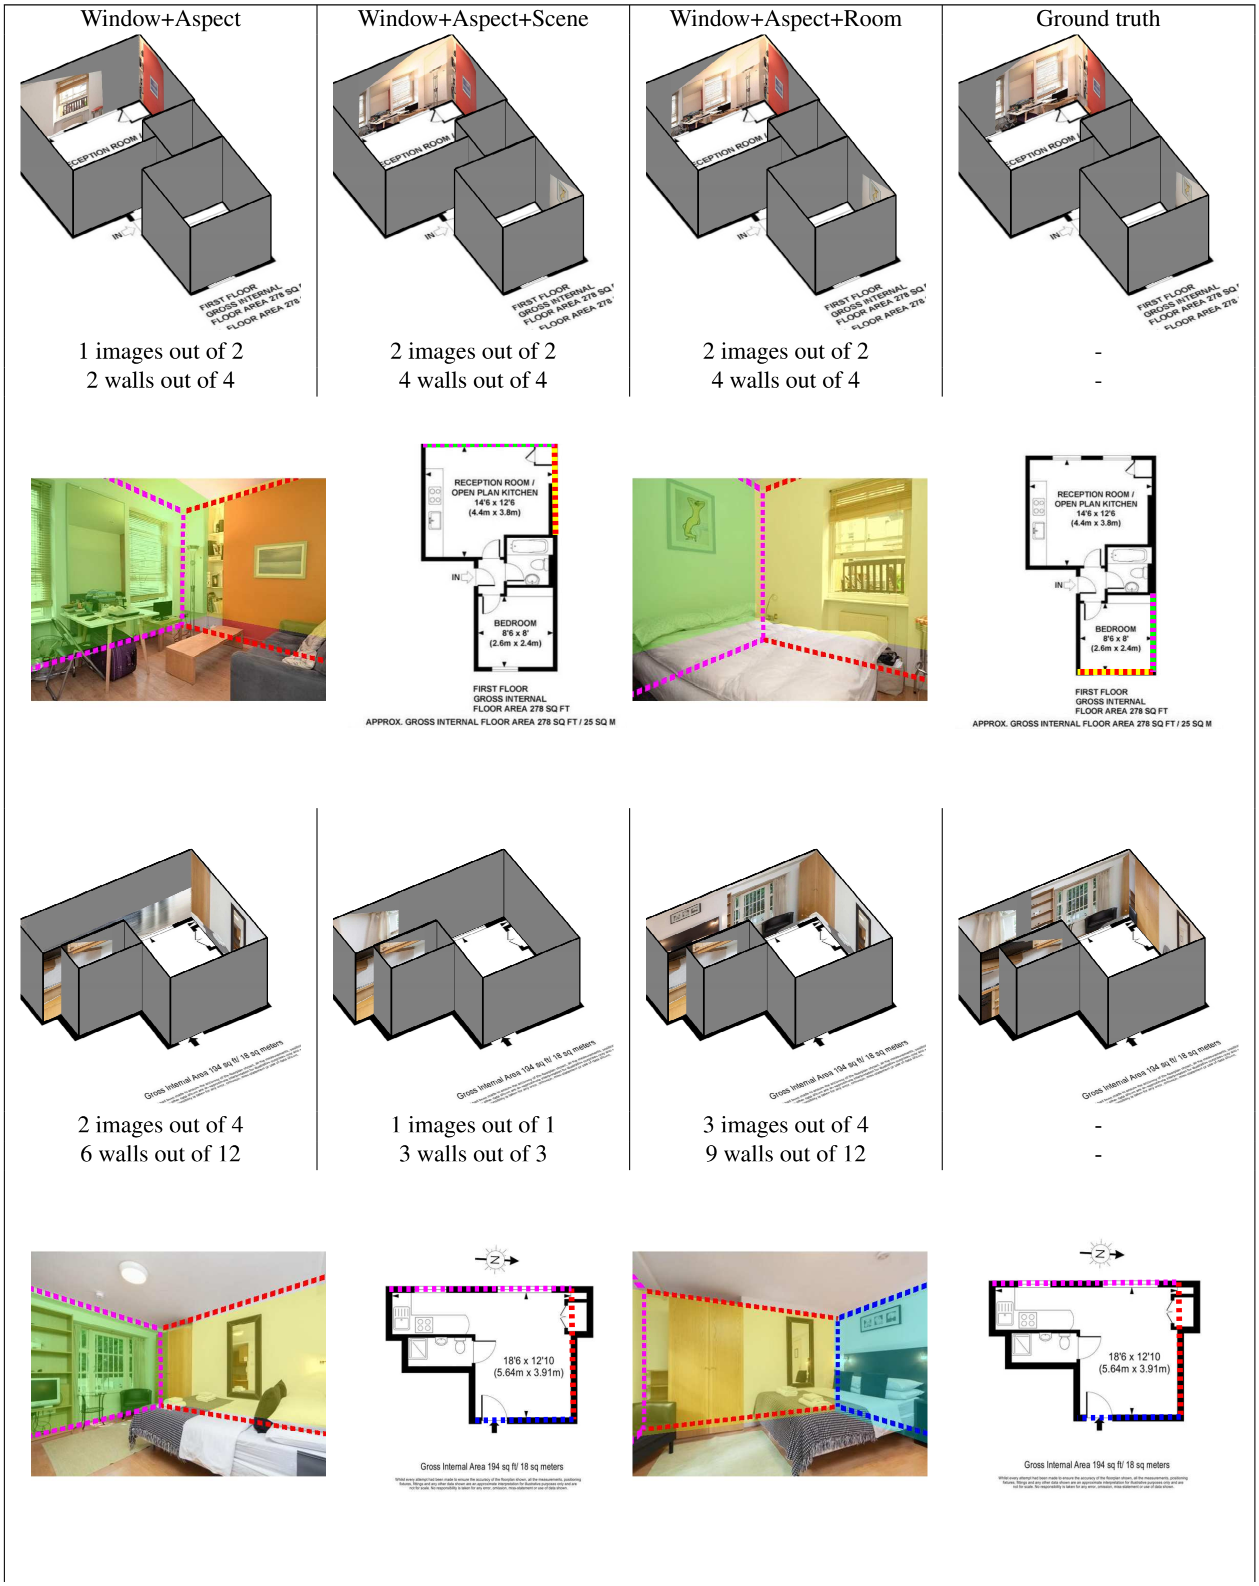

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
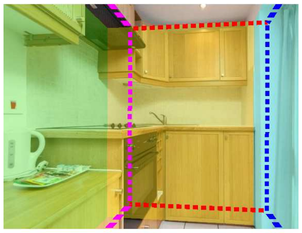	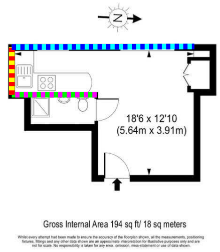	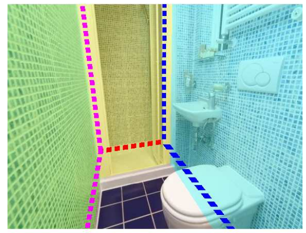	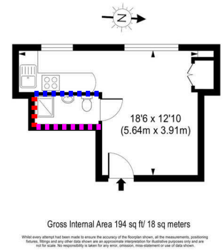
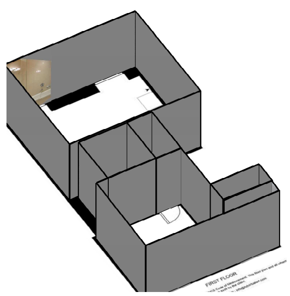 1 images out of 4 2 walls out of 8	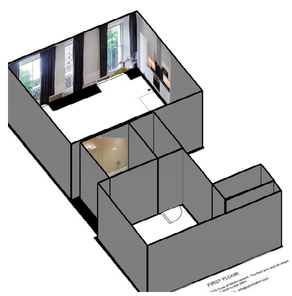 4 images out of 4 8 walls out of 8	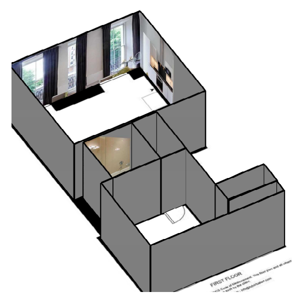 4 images out of 4 8 walls out of 8	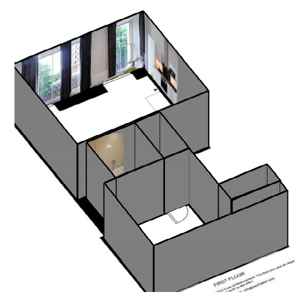 - -
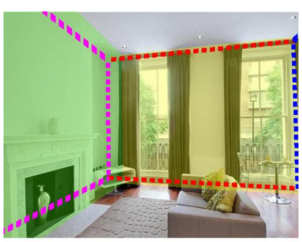	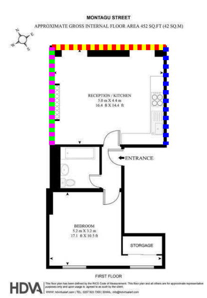		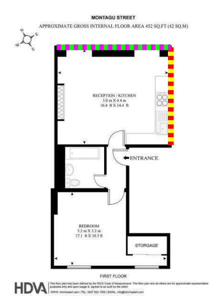
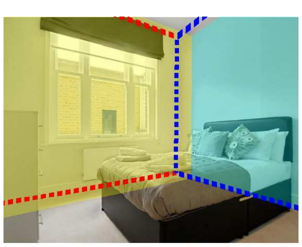			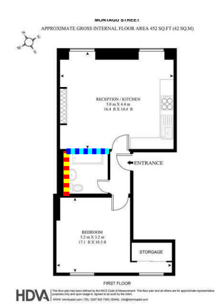

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
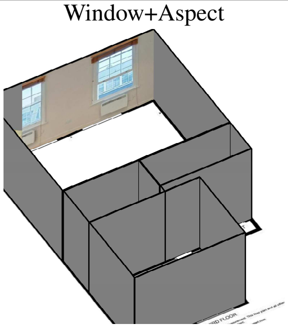	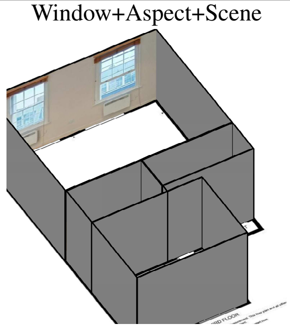	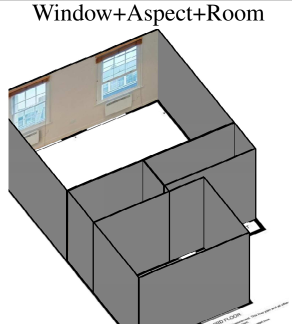	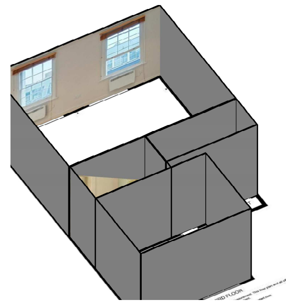
2 images out of 3 4 walls out of 7	3 images out of 3 7 walls out of 7	3 images out of 3 7 walls out of 7	- -
	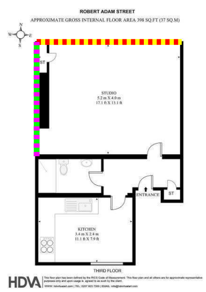		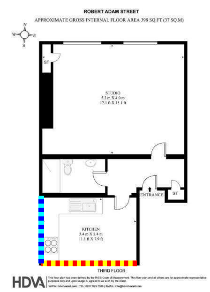
	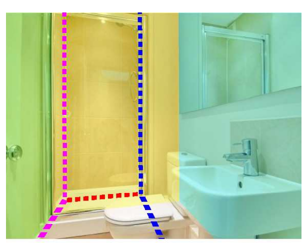	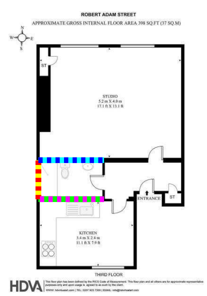	
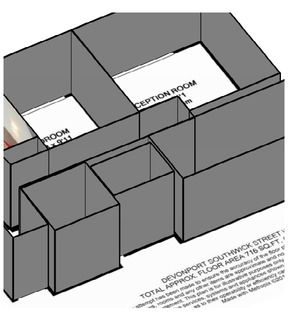	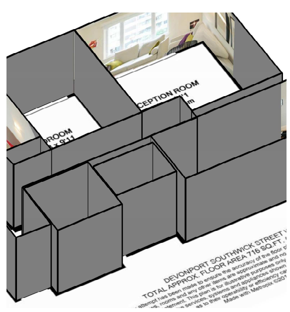	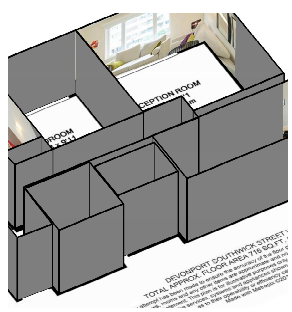	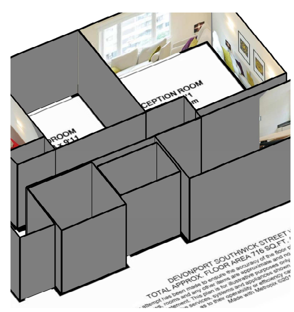
1 images out of 4 3 walls out of 11	3 images out of 4 9 walls out of 11	3 images out of 4 9 walls out of 11	- -


Window+Aspect
Window+Aspect+Scene
Window+Aspect+Room
Ground truth
1 images out of 5
3 walls out of 14
4 images out of 5
11 walls out of 14
4 images out of 5
11 walls out of 14
-
-
THIRD FLOOR
APPROX. GROSS INTERNAL FLOOR AREA 277 SQ FT / 25 SC
BEDROOM
RECEPTION ROOM
KITCHEN
HALLWAY
WARDROBE

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth

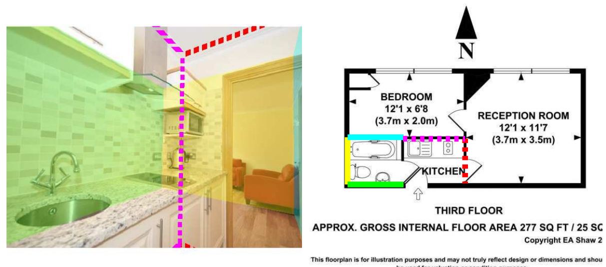

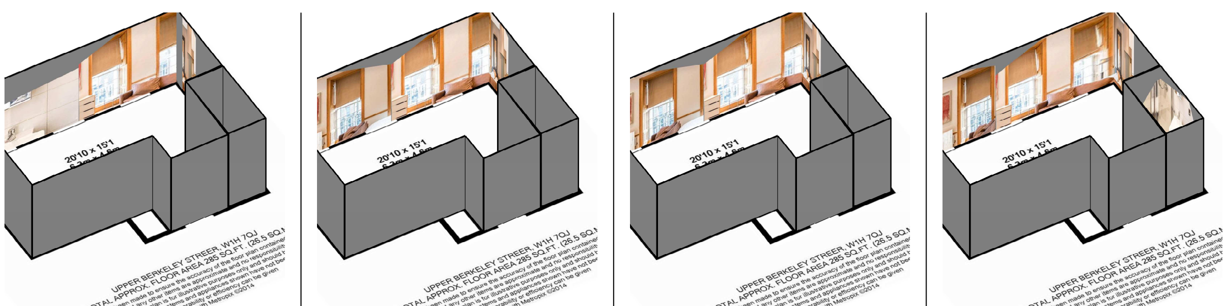

3 images out of 5	3 images out of 4	3 images out of 5	-
6 walls out of 11	6 walls out of 8	6 walls out of 11	-

<table>
<tr><th>Window+Aspect</th><th>Window+Aspect+Scene</th><th>Window+Aspect+Room</th><th>Ground truth</th></tr>
<tr>
<td></td>
<td>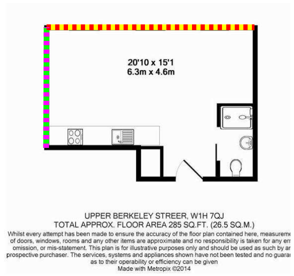</td>
<td></td>
<td>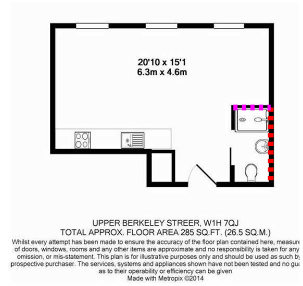</td>
</tr>
<tr>
<td></td>
<td></td>
<td>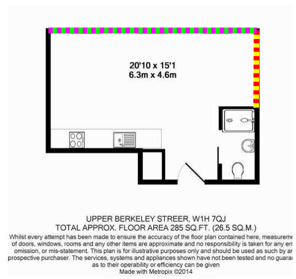</td>
<td></td>
</tr>
<tr>
<td>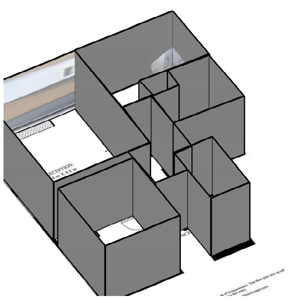
1 images out of 4
3 walls out of 12</td>
<td>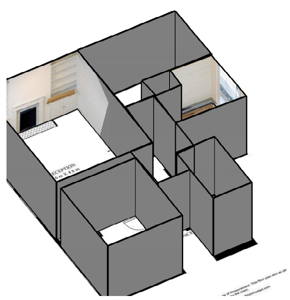
3 images out of 4
9 walls out of 12</td>
<td>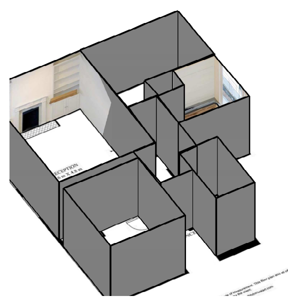
3 images out of 4
9 walls out of 12</td>
<td>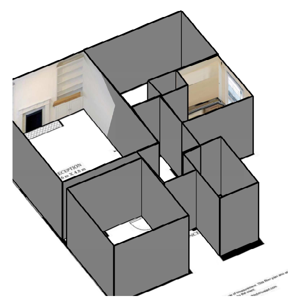
-
-</td>
</tr>
<tr>
<td></td>
<td>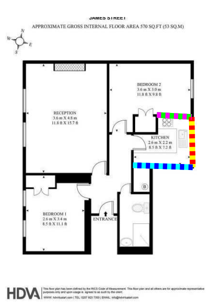</td>
<td></td>
<td>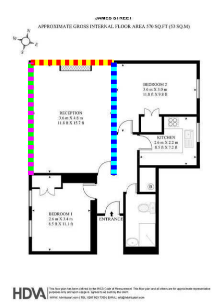</td>
</tr>
</table>

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
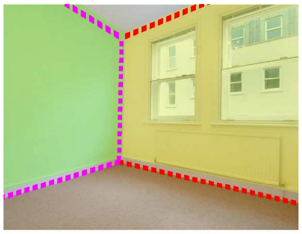		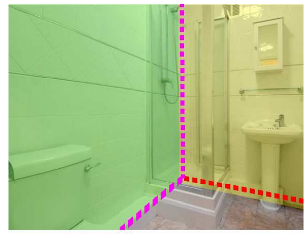	
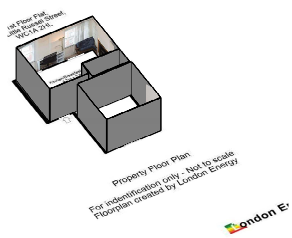	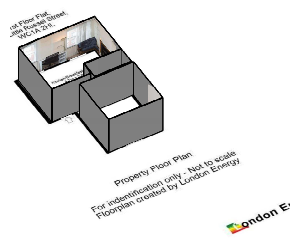	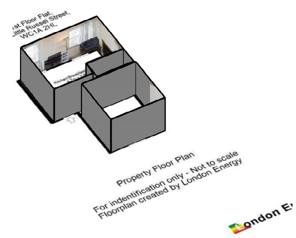	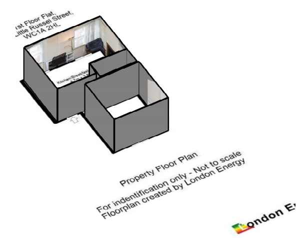
2 images out of 5	4 images out of 5	4 images out of 5	-
4 walls out of 11	8 walls out of 11	8 walls out of 11	-
	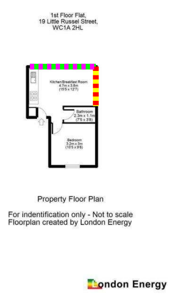		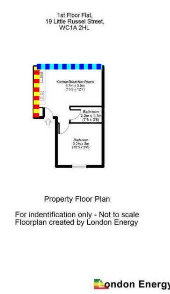
	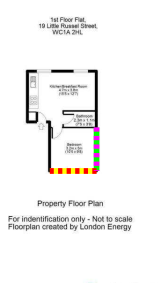		

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
		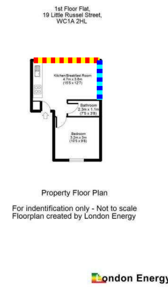	
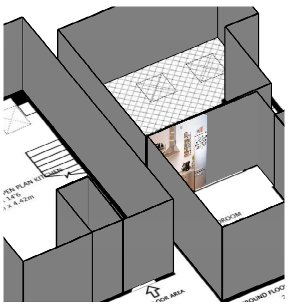	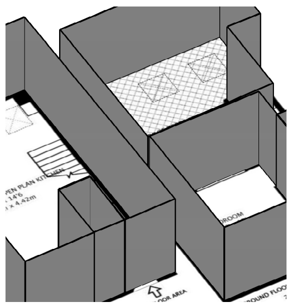	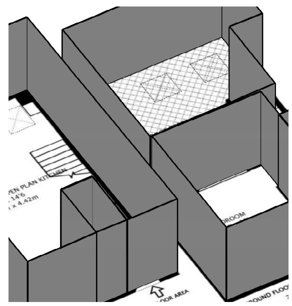	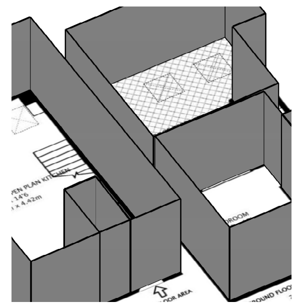
0 images out of 3 0 walls out of 6	2 images out of 3 4 walls out of 6	2 images out of 2 4 walls out of 4	- -
	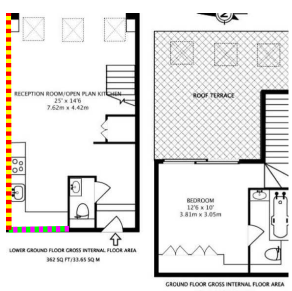		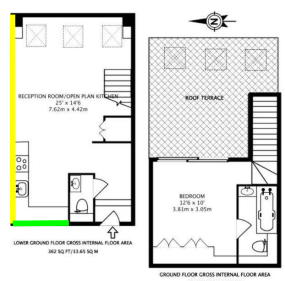
		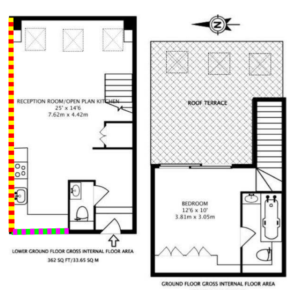	

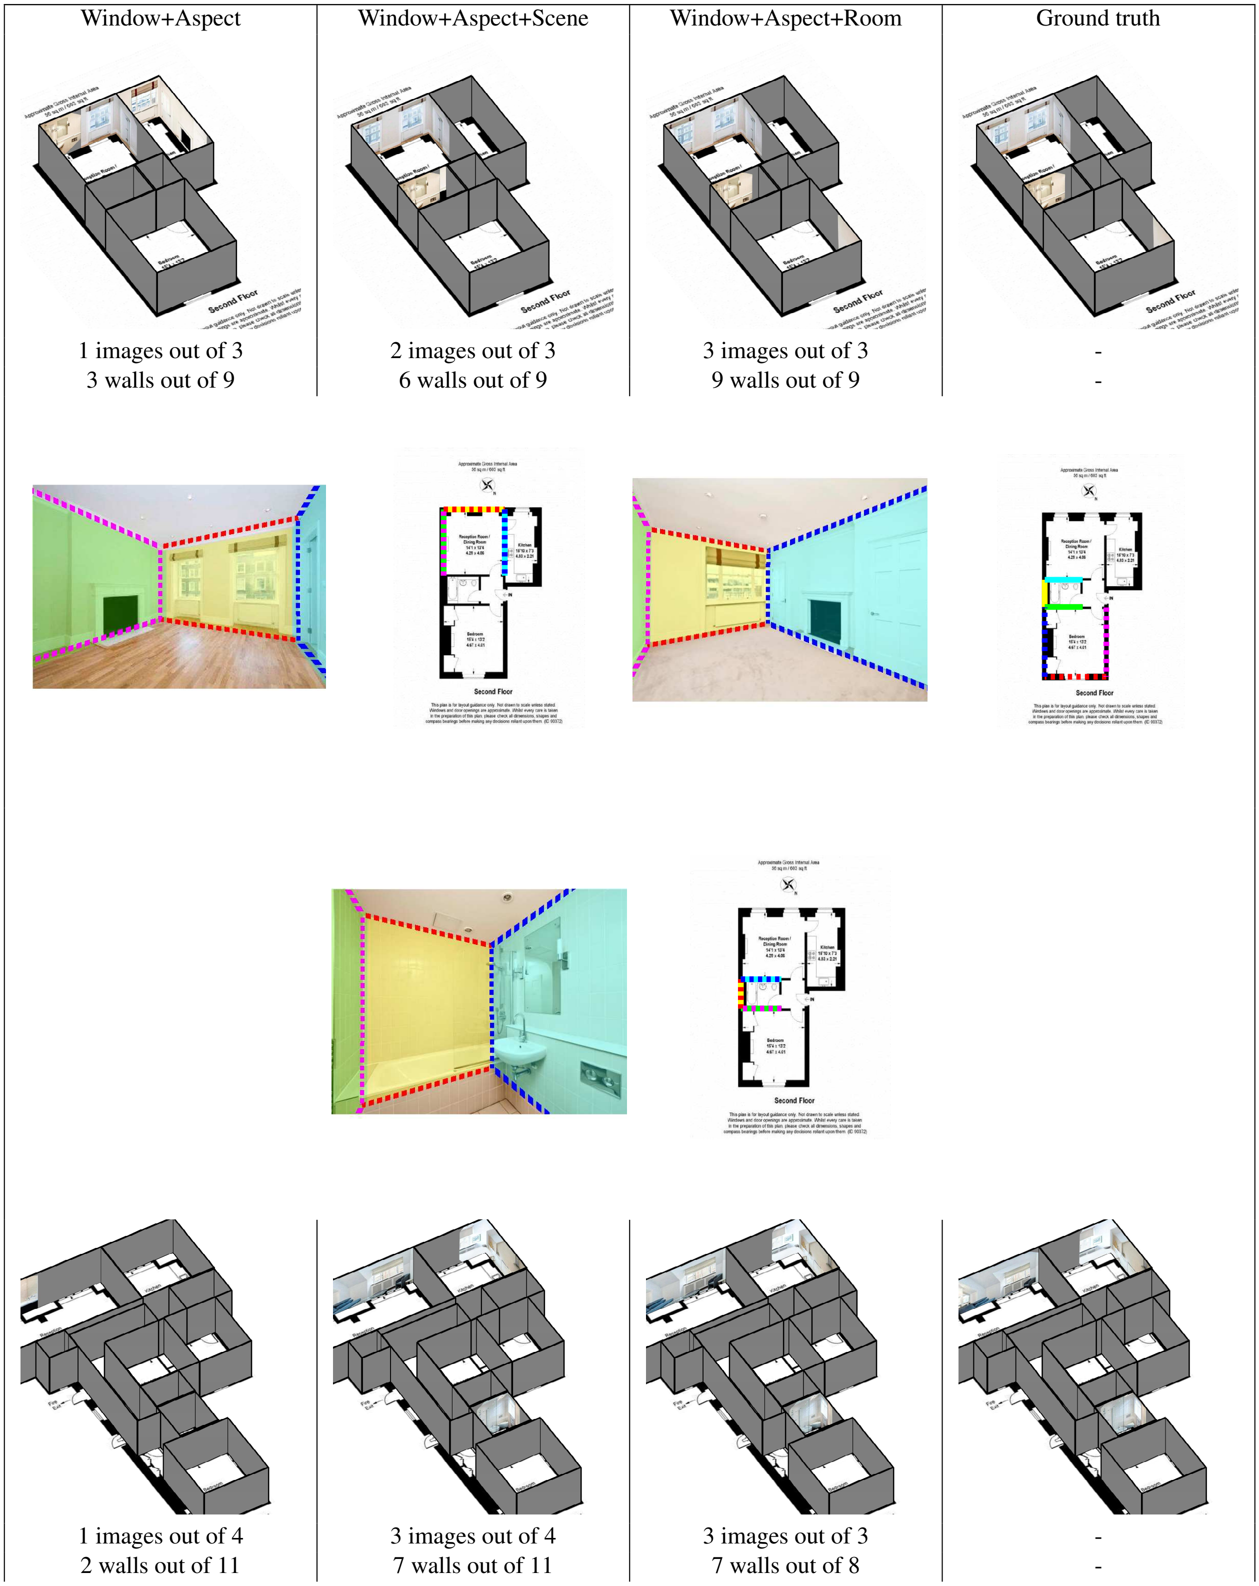
Window+Aspect
Window+Aspect+Scene
Window+Aspect+Room
Ground truth
1 images out of 3
3 walls out of 9
2 images out of 3
6 walls out of 9
3 images out of 3
9 walls out of 9
-
-
1 images out of 4
2 walls out of 11
3 images out of 4
7 walls out of 11
3 images out of 3
7 walls out of 8
-
-

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
	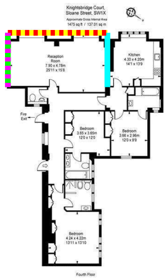		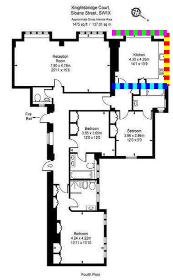
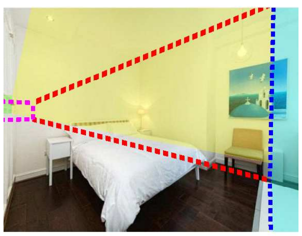	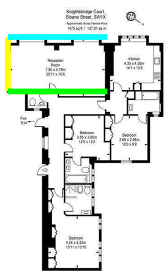		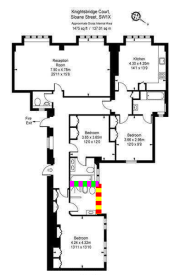
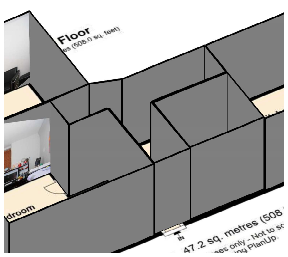	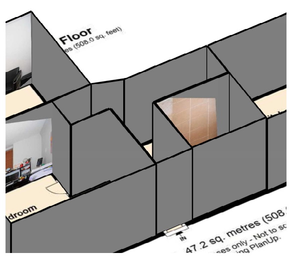	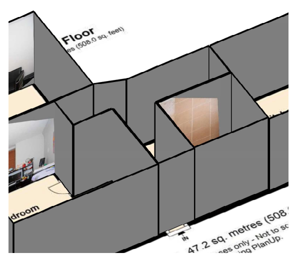	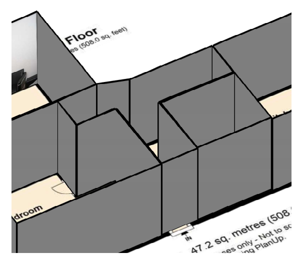
2 images out of 4 6 walls out of 10	2 images out of 4 6 walls out of 10	2 images out of 4 6 walls out of 10	- -
	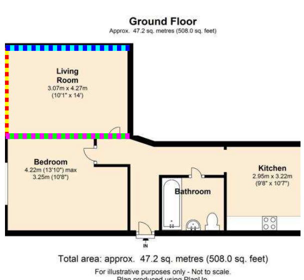		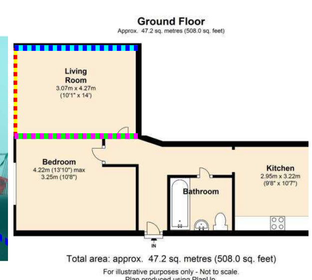

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth

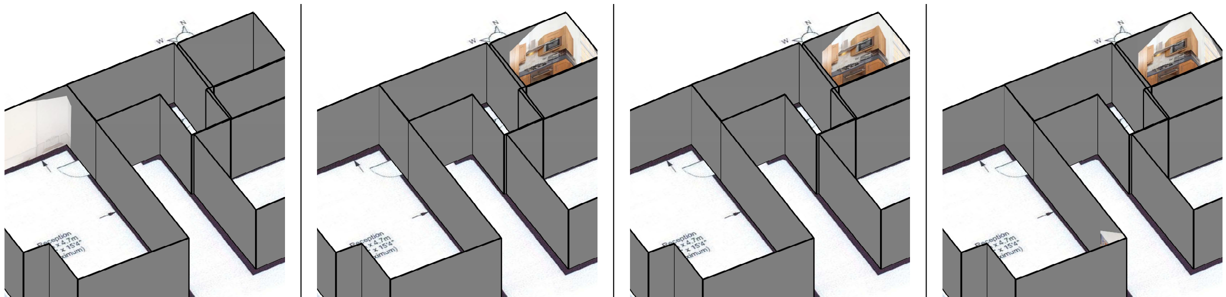

0 images out of 5	3 images out of 5	3 images out of 5	-
2 walls out of 15	9 walls out of 15	9 walls out of 15	-

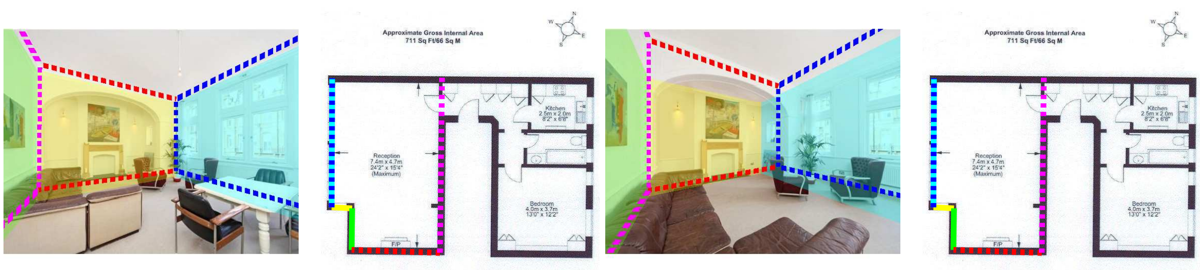

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
		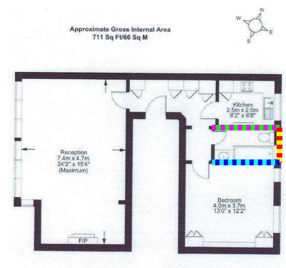	
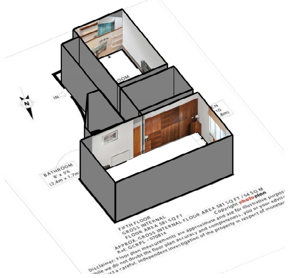 2 images out of 6 5 walls out of 14	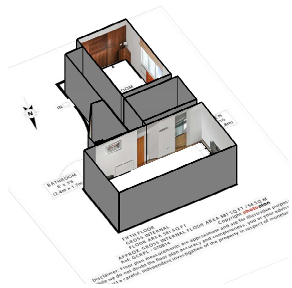 4 images out of 6 8 walls out of 14	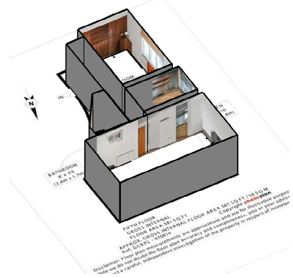 5 images out of 6 10 walls out of 14	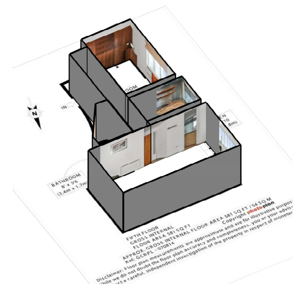 - -
	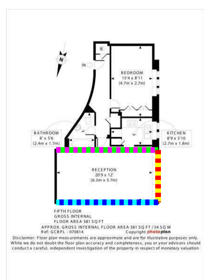		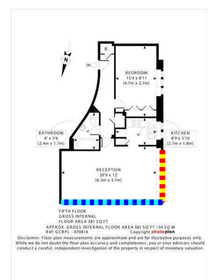
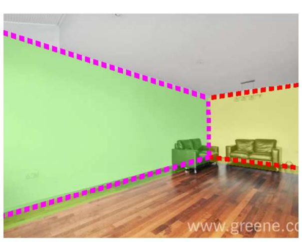	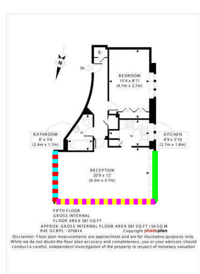		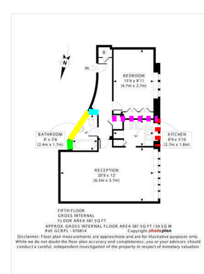

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
	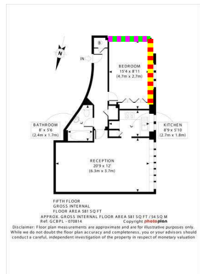		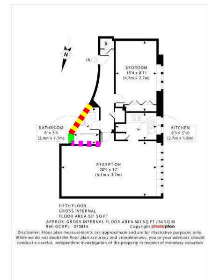

Table 1: For different settings of our model, 'window+aspect,' 'window+aspect+scene', 'window+aspect+room,' as well as for GT we illustrate successful apartment reconstructions on the test set in 3D. Below each we provide the numbers of how many images and walls were matched correctly.

2.2. Failure cases

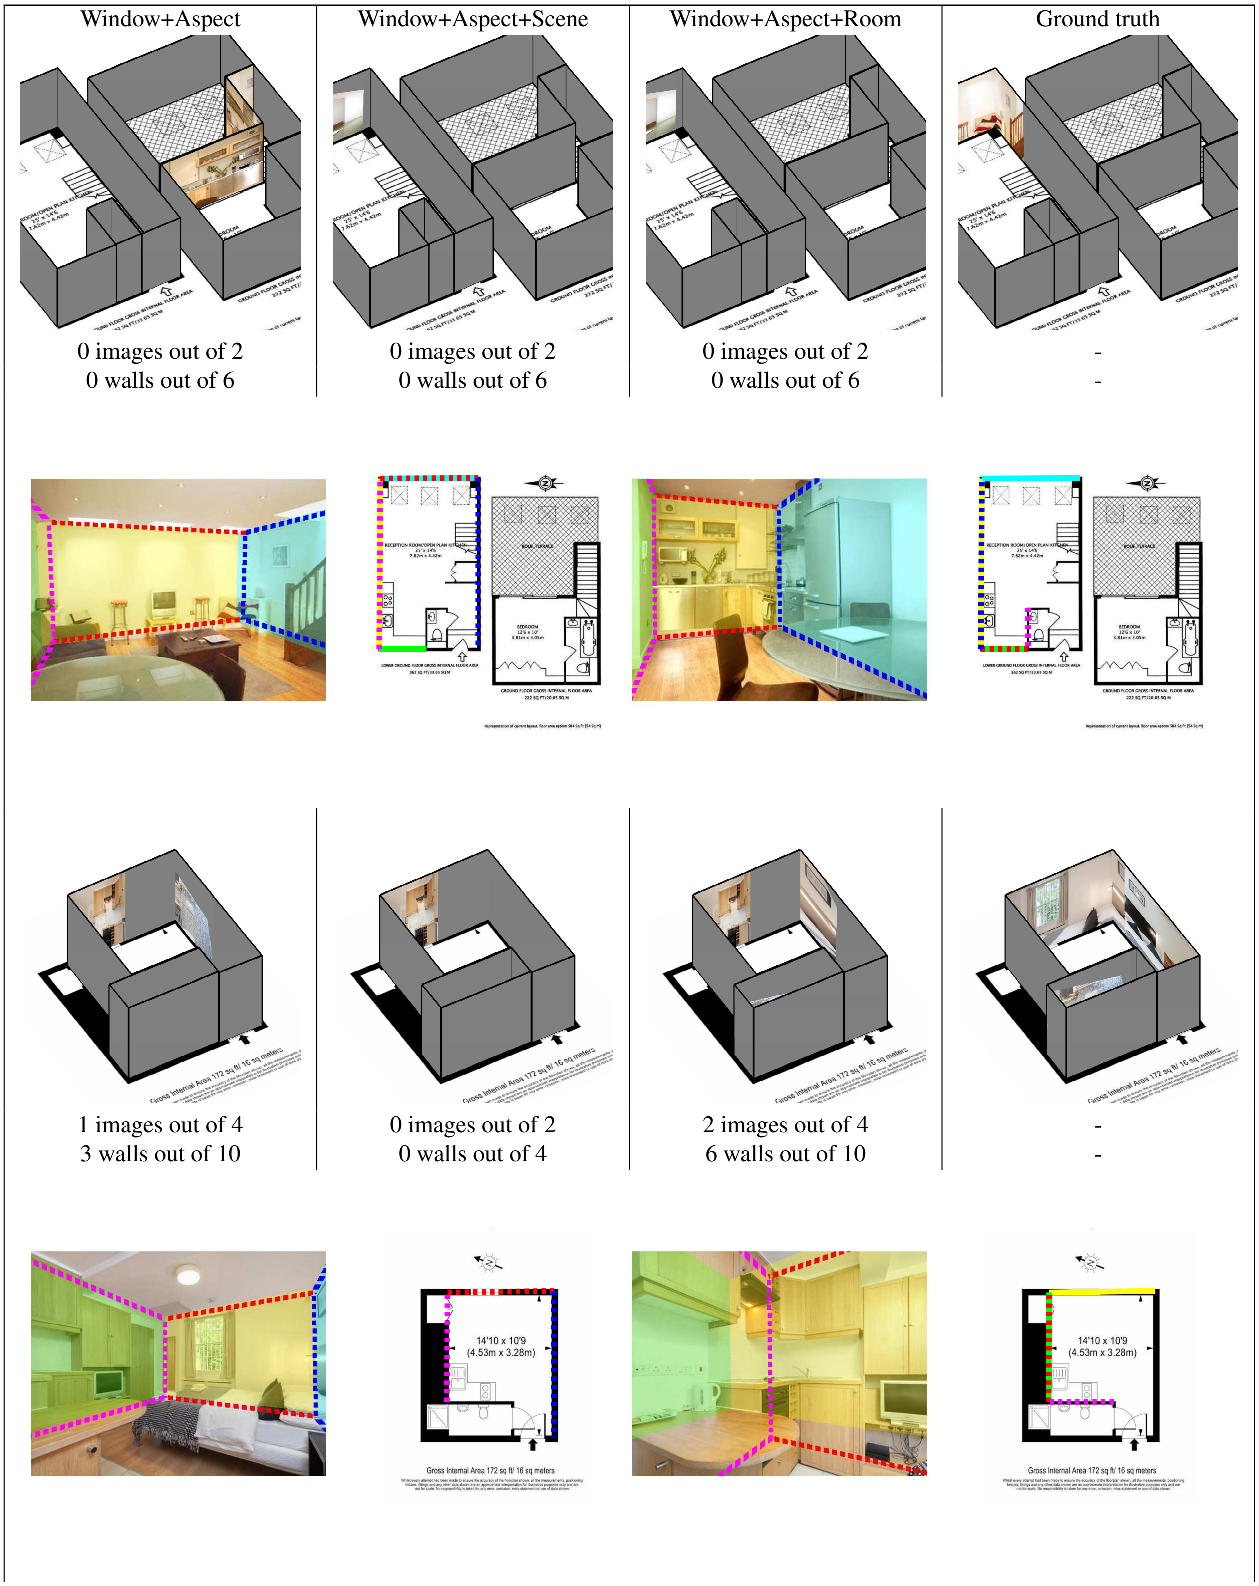

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
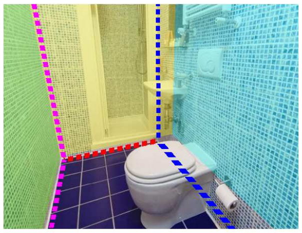	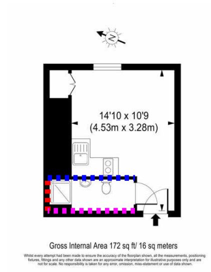	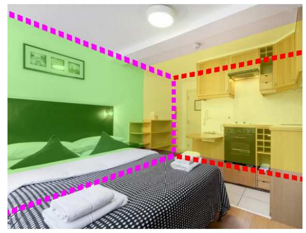	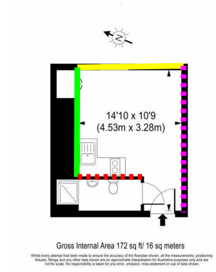
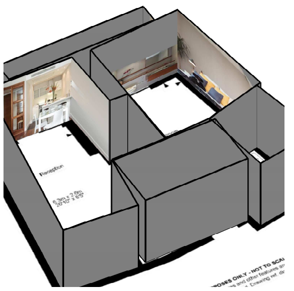	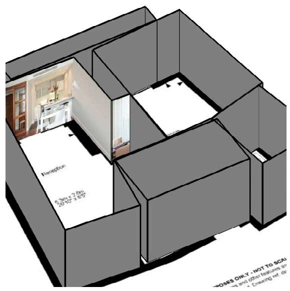	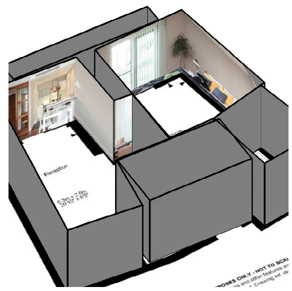	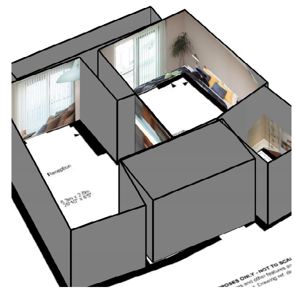
1 images out of 5 3 walls out of 12	0 images out of 5 0 walls out of 12	1 images out of 5 3 walls out of 12	- -
	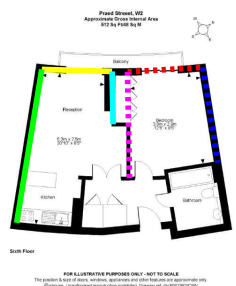		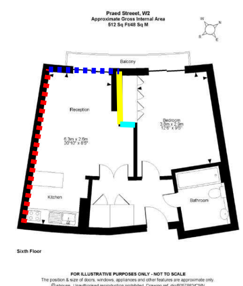
	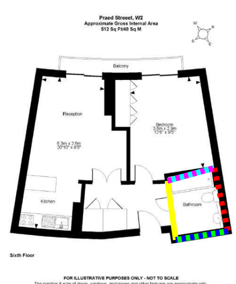	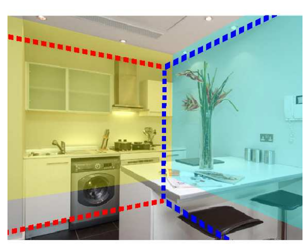	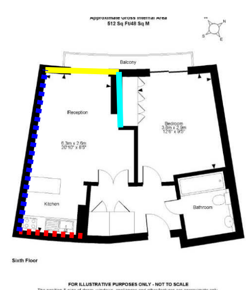

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
			
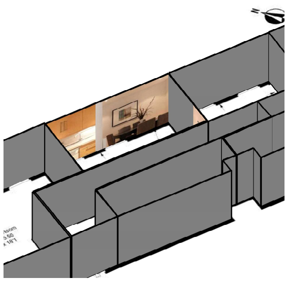	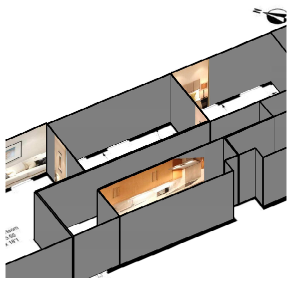	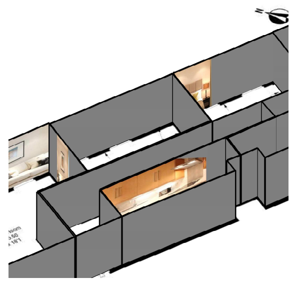	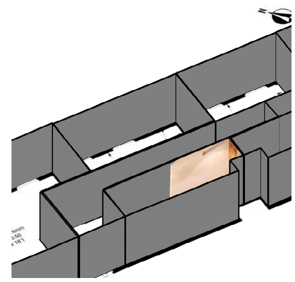
0 images out of 5 0 walls out of 11	0 images out of 5 0 walls out of 11	0 images out of 4 0 walls out of 9	- -
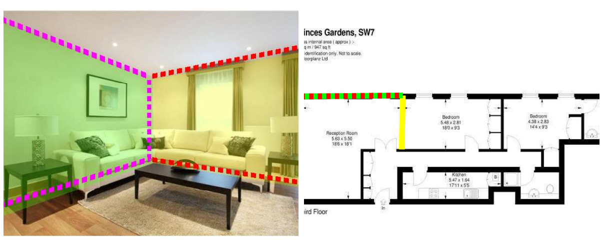		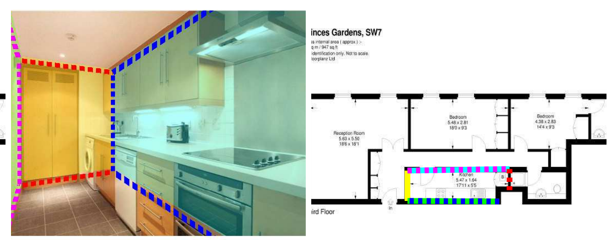	
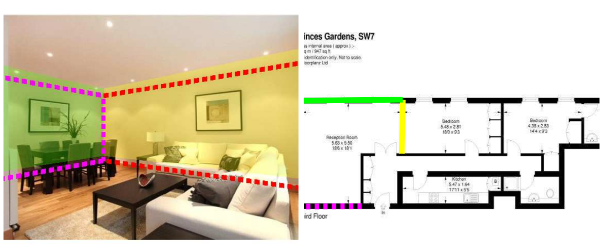			

Window+Aspect	Window+Aspect+Scene	Window+Aspect+Room	Ground truth
		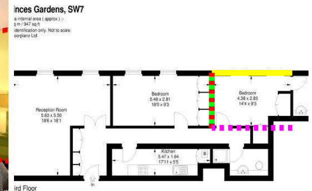	
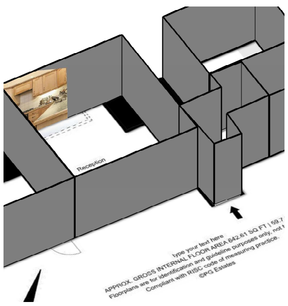	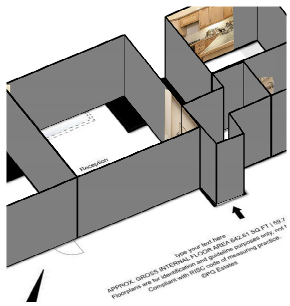	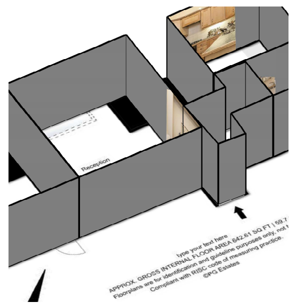	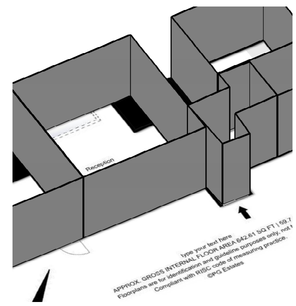
0 images out of 4 0 walls out of 9	0 images out of 4 0 walls out of 9	0 images out of 4 0 walls out of 9	- -
	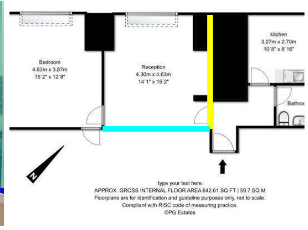		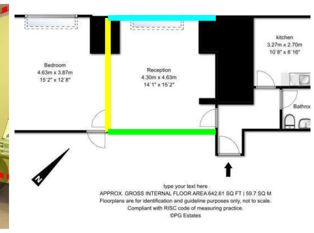
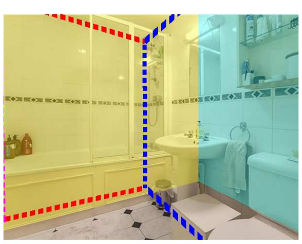	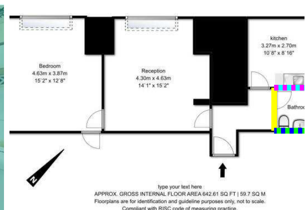		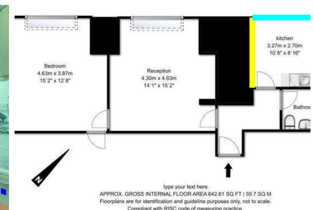

Table 2: For different settings of our model, 'window+aspect,' 'window+aspect+scene', 'window+aspect+room,' as well as for GT we illustrate failing apartment reconstructions on the test set in 3D. Below each we provide the numbers of how many images and walls were matched correctly.