

## Supplemental Material: Field of Junctions: Extracting Boundary Structure at Low SNR

### Table of Contents

<b>1</b>	<b>Additional demonstrations of the model’s parameters</b>	<b>2</b>
<b>2</b>	<b>Complete proof of Theorem 1</b>	<b>4</b>
<b>3</b>	<b>Importance of refinement</b>	<b>6</b>
<b>4</b>	<b>Parametrization of indicator functions during refinement</b>	<b>8</b>
<b>5</b>	<b>Additional results for Figure 1</b>	<b>9</b>
<b>6</b>	<b>Additional comparisons with increasing noise</b>	<b>10</b>
<b>7</b>	<b>Additional experiments with correlated noise</b>	<b>12</b>
<b>8</b>	<b>Runtime and effect of stride</b>	<b>14</b>
<b>9</b>	<b>Complete specification of vertex detector</b>	<b>16</b>
<b>10</b>	<b>Synthetic dataset</b>	<b>17</b>
<b>11</b>	<b>Comparison with the 4-junction model</b>	<b>18</b>
<b>12</b>	<b>Comparison with the linear color model</b>	<b>21</b>
	12.1 Derivation	21
	12.2 Results	22
<b>13</b>	<b>Application to RGB-D images</b>	<b>23</b>
<b>14</b>	<b>Repeatability over change in viewpoint</b>	<b>24</b>
<b>15</b>	<b>Comparison with <math>L_0</math> smoothing</b>	<b>25</b>
<b>16</b>	<b>Additional results on photographs</b>	<b>26</b>

# 1 Additional demonstrations of the model’s parameters

Our model has three important adjustable parameters: the patch size  $R$ , the boundary consistency weight  $\lambda_B$ , and the color consistency weight  $\lambda_C$ . In this section we provide more information on the behavior of each of those parameters, as well as examples demonstrating their effect.

The boundary consistency term, governed by  $\lambda_B$ , promotes agreement on boundaries between overlapping patches. It improves the resilience of the field of junctions to noise, and acts as a curvature regularizer which favors isolated corners and junctions connected by low curvature contours.

The color consistency term, governed by  $\lambda_C$ , promotes agreement on colors between overlapping patches. It provides a means for long-range color sharing between different patches by penalizing color disagreement at the overlap between different patches.

The patch size  $R$  determines the minimal detail scale which the field of junctions can capture. Details that have spatial size smaller than  $R$ , such as two nearby junctions, cannot be described using a single generalized junction, and are suppressed. In addition, the patch size  $R$  is crucial in high-noise usage: it should be large to provide enough information for the initialization algorithm to obtain a good initialization, yet small enough to allow capturing all required structure in the image.

Figure S1 shows an extension of Figure 4 from the main paper. It compares our results with those obtained by the  $\ell_1$ -elastica and  $\ell_2$ -elastica Chan-Vese methods [S5], using various parameters. Both methods use a penalty of the form  $\int (a + b|\kappa(s)|^d)ds$ , where  $d = 1$  for  $\ell_1$ -elastica and  $d = 2$  for  $\ell_2$ -elastica, and  $\kappa(s)$  denotes the curvature at a point parameterized by  $s$ . In addition to the two parameters  $a$  and  $b$ , the implementation in [S5] uses additional parameters corresponding to a set of auxiliary variables in an augmented Lagrangian optimization scheme. The  $\ell_1$ -elastica method introduces four new parameters denoted  $r_1, r_2, r_3$ , and  $r_4$ , and the  $\ell_2$ -elastica model uses three parameters  $r_1, r_2$ , and  $r_3$ . These additional parameters also have an effect on the results along with  $a$  and  $b$ , and we adjust them to yield optimal results separately for each image. Figure S1 shows a number of representative results using the two models compared with our field of junctions.

Figure S2 shows the effect of the patch size  $R$  in the strong regularization regime, *i.e.*, when  $\lambda_B$  is set to a large value. In that case, the only permissible contours are polygons with  $\ell_\infty$  side length of at least  $R$ .

Figure S3 shows the effect of the color consistency term in the high noise regime. The color consistency term acts to share color information between overlapping patches, which helps refine the junction parameters at each location. This significantly improves the resilience of the field of junctions to high noise.

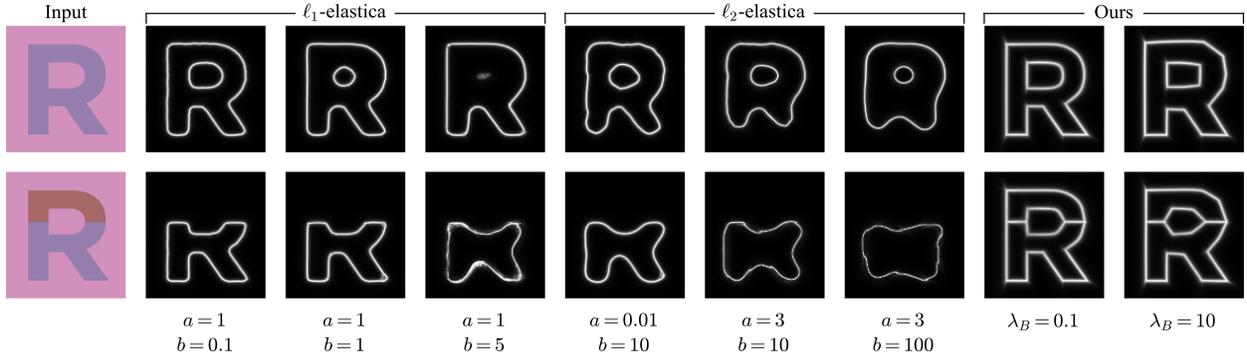


Figure S1: An extended version of Figure 4 of our main paper, showing the curvature regularization capabilities of the field of junctions compared to  $\ell_1$ -elastica and  $\ell_2$ -elastica with different parameters. The additional parameters of the elastica models were tuned to show a wide range of behaviors. We show the results from the field of junctions using two  $\lambda_B$  values, and the other two parameters of the field of junctions were fixed at  $R = 21$  and  $\lambda_C = 0$ .

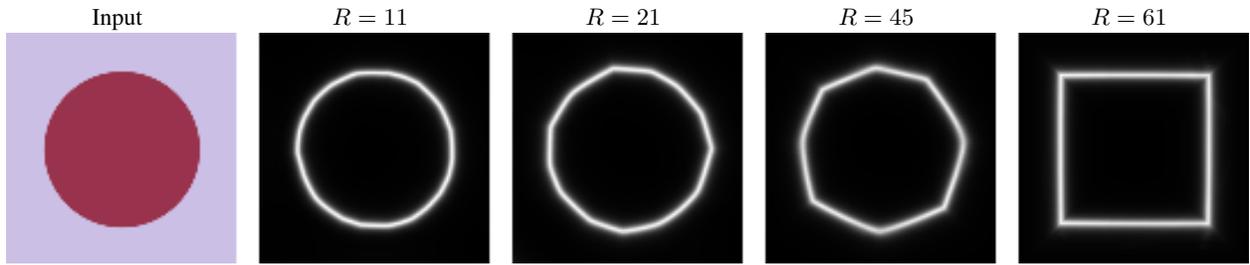


Figure S2: The boundary map obtained by the field of junctions with  $\lambda_B = 100$  and  $\lambda_C = 0$  for increasing values of  $R$ . The boundary consistency term penalizes vertices separated by  $\ell_\infty$  distance smaller than  $R$ , which means that only polygons with side length of at least  $R$  are permissible when  $\lambda_B$  is large.

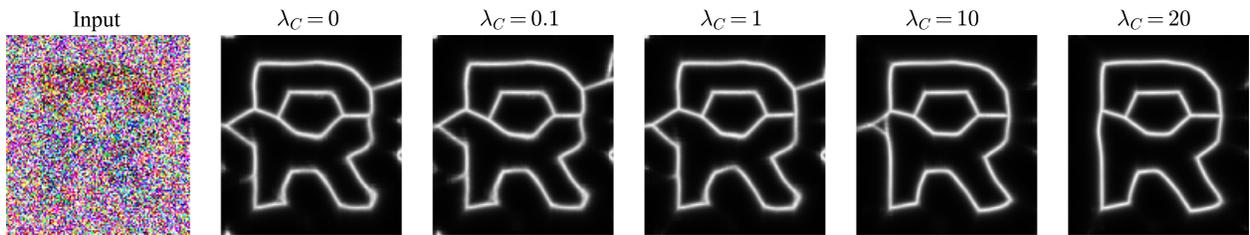


Figure S3: The boundary map obtained by the field of junctions with increasing color consistency strength  $\lambda_C$ . The color consistency term improves the field of junction's resilience to noise by enforcing a global agreement on color values.

## 2 Complete proof of Theorem 1

We present a complete proof of Theorem 1 from the main paper, showing that Algorithm 1 is guaranteed to find the  $M$  true angles of an image (or image patch)  $I(\mathbf{x})$  containing a single junction, in the noiseless case, and when the vertex position is known. In the main paper we show that in practice Algorithm 2 works very well even when there is substantial noise, or when the vertex position is not known.

For completeness we begin by restating the theorem, omitting the patch index  $i$  for convenience:

**Theorem 1.** *For a junction image  $I(\mathbf{x})$  with no noise (i.e.,  $n \equiv 0$  in Eq. 1 of the main paper) and with vertex  $\mathbf{x}^{(0)}$  known, Algorithm 1 is guaranteed to find the globally optimal angles  $\phi$ .*

We begin by first proving a lemma:

**Lemma S1.** *Define  $L(\phi) \triangleq \int_{a(\phi)}^{b(\phi)} w(\phi') \|I(\phi') - c^*(\phi)\|^2 d\phi'$ , where  $a, b, w: [0, 2\pi) \rightarrow \mathbb{R}$  and  $I, c^*: [0, 2\pi) \rightarrow \mathbb{R}^K$  are functions such that  $c^*(\phi) \int_{a(\phi)}^{b(\phi)} w(\phi') d\phi' = \int_{a(\phi)}^{b(\phi)} w(\phi') I(\phi') d\phi'$ . Then we have:*

$$\frac{d}{d\phi} L(\phi) = w(b(\phi)) \|I(b(\phi)) - c^*(\phi)\|^2 \frac{d}{d\phi} b(\phi) - w(a(\phi)) \|I(a(\phi)) - c^*(\phi)\|^2 \frac{d}{d\phi} a(\phi). \quad (\text{S1})$$

*Proof.* Using the Leibniz integral rule we have the desired equality plus an additional term:

$$\begin{aligned} & \int_{a(\phi)}^{b(\phi)} w(\phi') \frac{\partial}{\partial \phi} \|I(\phi') - c^*(\phi)\|^2 d\phi' \\ &= -2 \int_{a(\phi)}^{b(\phi)} w(\phi') \frac{d}{d\phi} c^*(\phi) \cdot [I(\phi') - c^*(\phi)] d\phi' \\ &= -2 \frac{d}{d\phi} c^*(\phi) \cdot \int_{a(\phi)}^{b(\phi)} w(\phi') [I(\phi') - c^*(\phi)] d\phi' = 0, \end{aligned} \quad (\text{S2})$$

where the final equality was obtained using the fact that the integral is zero, due to the construction of  $c^*$  in the lemma.  $\square$

**Lemma S2.** *Under the conditions of Theorem 1, the negative log-likelihood function restricted to the  $j$ th coordinate  $\ell_j(\phi) = \ell(\phi^{(1)}, \dots, \phi^{(j-1)}, \phi, \phi^{(j+1)}, \dots, \phi^{(M)}, x^{(0)}, y^{(0)})$  has no local minima in the open angular interval between any pair of adjacent true angles.*

*Proof.* The junction image  $I(\mathbf{x})$  is radially symmetric around the known vertex  $\mathbf{x}^{(0)}$ , and therefore we can treat it as an angular function  $I(\phi)$  of the angle centered at point  $\mathbf{x}^{(0)}$ .

From Equation 4 in the main paper, and using polar coordinates relative to the vertex  $\mathbf{x}^{(0)}$  we can write the negative log-likelihood:

$$\begin{aligned} \ell(\phi^{(1)}, \dots, \phi^{(M)}, x^{(0)}, y^{(0)}) &= \alpha \sum_{j=1}^M \int_0^{2\pi} \int_0^{R(\phi')} u^{(j)}(\phi') \|I(\phi') - c^{(j)}\|^2 r' dr' d\phi' \\ &= \alpha \sum_{j=1}^M \int_0^{2\pi} \frac{1}{2} R^2(\phi') u^{(j)}(\phi') \|I(\phi') - c^{(j)}\|^2 d\phi' \\ &= \alpha \sum_{j=1}^M \int_{\phi^{(j)}}^{\phi^{(j+1)}} \frac{1}{2} R^2(\phi') \|I(\phi') - c^{(j)}\|^2 d\phi', \end{aligned} \quad (\text{S3})$$

where  $R(\phi')$  is the distance of each point on the boundary of the  $R \times R$  patch at angle  $\phi'$  from the vertex  $\mathbf{x}^{(0)}$ , and  $u^{(j)}(\phi')$  is an angular indicator function returning 1 if  $\phi' \in (\phi^{(j)}, \phi^{(j+1)})$  and 0 otherwise (with the index  $j+1$  computed modulo  $M$ , and assuming without loss of generality that the angles are non-decreasing,  $\phi^{(1)} \leq \dots \leq \phi^{(M)}$ ). The optimal colors  $\{c^{(j)}\}$  are given using Equation 3 in the main paper, or in polar coordinates:

$$c^{(j)} = \frac{\int_0^{2\pi} \int_0^{R(\phi')} u^{(j)}(\phi') I(\phi') r' dr' d\phi'}{\int_0^{2\pi} \int_0^{R(\phi')} u^{(j)}(\phi') r' dr' d\phi'} = \frac{\int_{\phi^{(j)}}^{\phi^{(j+1)}} R^2(\phi') I(\phi') d\phi'}{\int_{\phi^{(j)}}^{\phi^{(j+1)}} R^2(\phi') d\phi'}. \quad (\text{S4})$$

By applying Lemma S1 to Equation S3 and discarding all terms that do not depend on  $\phi^{(j)}$ , the partial derivative of the negative log-likelihood function with respect to the  $j$ th angle in any open interval between two true angles is:

$$\frac{d}{d\phi} \ell_j(\phi) = \frac{1}{2} \alpha R^2(\phi) \left[ \|I(\phi) - c^{(j-1)}\|^2 - \|I(\phi) - c^{(j)}\|^2 \right]. \quad (\text{S5})$$

We finish proving the lemma by contradiction. Let us assume  $\phi$  is a local minimum point of  $\ell_j$  in the open interval between a pair of true angles. Because  $\ell_j$  is smooth within any such interval, we have  $\frac{d}{d\phi} \ell_j(\phi) = 0$ , and from Equation S5 we have:

$$\|I(\phi) - c^{(j-1)}\|^2 - \|I(\phi) - c^{(j)}\|^2 = 0. \quad (\text{S6})$$

Using this expression with the derivative of Equation S5 we obtain that the second derivative at the local minimum is:

$$\begin{aligned} \frac{d^2}{d^2\phi} \ell_j(\phi) &= \alpha R^2(\phi) \left\{ \left[ \frac{d}{d\phi} I(\phi) - \frac{d}{d\phi} c^{(j-1)} \right] \cdot [I(\phi) - c^{(j-1)}] - \left[ \frac{d}{d\phi} I(\phi) - \frac{d}{d\phi} c^{(j)} \right] \cdot [I(\phi) - c^{(j)}] \right\} \\ &= \alpha R^2(\phi) \left\{ -\frac{d}{d\phi} c^{(j-1)} \cdot [I(\phi) - c^{(j-1)}] + \frac{d}{d\phi} c^{(j)} \cdot [I(\phi) - c^{(j)}] \right\}, \end{aligned} \quad (\text{S7})$$

where we used the fact that  $I(\phi)$  is constant for every  $\phi$  in an open interval between two true angles.

Multiplying both sides of Equation S4 by the denominator of the right hand side for indices  $j - 1$  and  $j$  and differentiating with respect to  $\phi$  yields:

$$\begin{aligned} \frac{d}{d\phi} c^{(j-1)} \int_{\phi^{(j-1)}}^{\phi} R^2(\phi') d\phi' + c^{(j-1)} R^2(\phi) &= R^2(\phi) I(\phi), \\ \frac{d}{d\phi} c^{(j)} \int_{\phi}^{\phi^{(j)}} R^2(\phi') d\phi' - c^{(j)} R^2(\phi) &= -R^2(\phi) I(\phi). \end{aligned} \quad (\text{S8})$$

Rearranging we obtain:

$$R^2(\phi) [I(\phi) - c^{(j-1)}] = \frac{d}{d\phi} c^{(j-1)} \int_{\phi^{(j-1)}}^{\phi} R^2(\phi') d\phi', \quad (\text{S9})$$

$$R^2(\phi) [I(\phi) - c^{(j)}] = -\frac{d}{d\phi} c^{(j)} \int_{\phi}^{\phi^{(j)}} R^2(\phi') d\phi'. \quad (\text{S10})$$

Finally, we substitute Equations S9 and S10 into the expression for the second derivative in Equation S7:

$$\frac{d^2}{d^2\phi} \ell_j(\phi) = \alpha \left\{ -\left\| \frac{d}{d\phi} c^{(j-1)} \right\|^2 \int_{\phi^{(j-1)}}^{\phi} R^2(\phi') d\phi' - \left\| \frac{d}{d\phi} c^{(j)} \right\|^2 \int_{\phi}^{\phi^{(j)}} R^2(\phi') d\phi' \right\}, \quad (\text{S11})$$

and the second derivative at  $\phi$  is negative, which is contradictory to our assumption that  $\phi$  is a minimum point, thus concluding our proof of the lemma.  $\square$

The complete proof of Theorem 1 is then obtained by using Lemma S2 and the proof sketch from the main paper:

*Proof.* First, note that  $\ell_j(\phi)$  is continuous and smooth for all  $\phi$  other than possibly a discontinuity in the derivative at any of the true junction angles. If the optimal  $\phi$  is not one of the true junction angles then it must lie in the *open* interval between two such angles, *i.e.*  $\phi \in (\phi^-, \phi^+)$ . From Lemma S2 we have that  $\ell_j(\phi)$  does not have any local minima in  $(\phi^-, \phi^+)$ , and therefore for each angular interval between two true junction angles the cost function must be minimized at one of the endpoints. Therefore repeatedly minimizing  $\ell_j(\phi)$  for  $j = 1, \dots, M$  is guaranteed to provide a globally optimal set of angles  $\{\phi^{(1)}, \dots, \phi^{(M)}\}$ .  $\square$

### 3 Importance of refinement

Figure S4 shows a comparison of the results obtained by a field of junctions with and without the refinement stage on a photograph. In natural images with low noise, the refinement stage mainly cleans up the boundary maps, only maintaining boundaries corresponding to salient edges and junctions in the image. Because the main effect of the refinement stage in this case is to remove unnecessary boundaries around uniform regions of the image, it does not affect the boundary-aware smoothing, as the pixel colors predicted in both cases are similar.

Figure S5 shows the effect of the refinement stage over multiple noise levels. While in the low-noise regime the refinement stage of our algorithm only suppresses superfluous boundaries, in the high-noise regime the spatial consistency enforced by it also significantly improves the boundaries predicted by each junction.

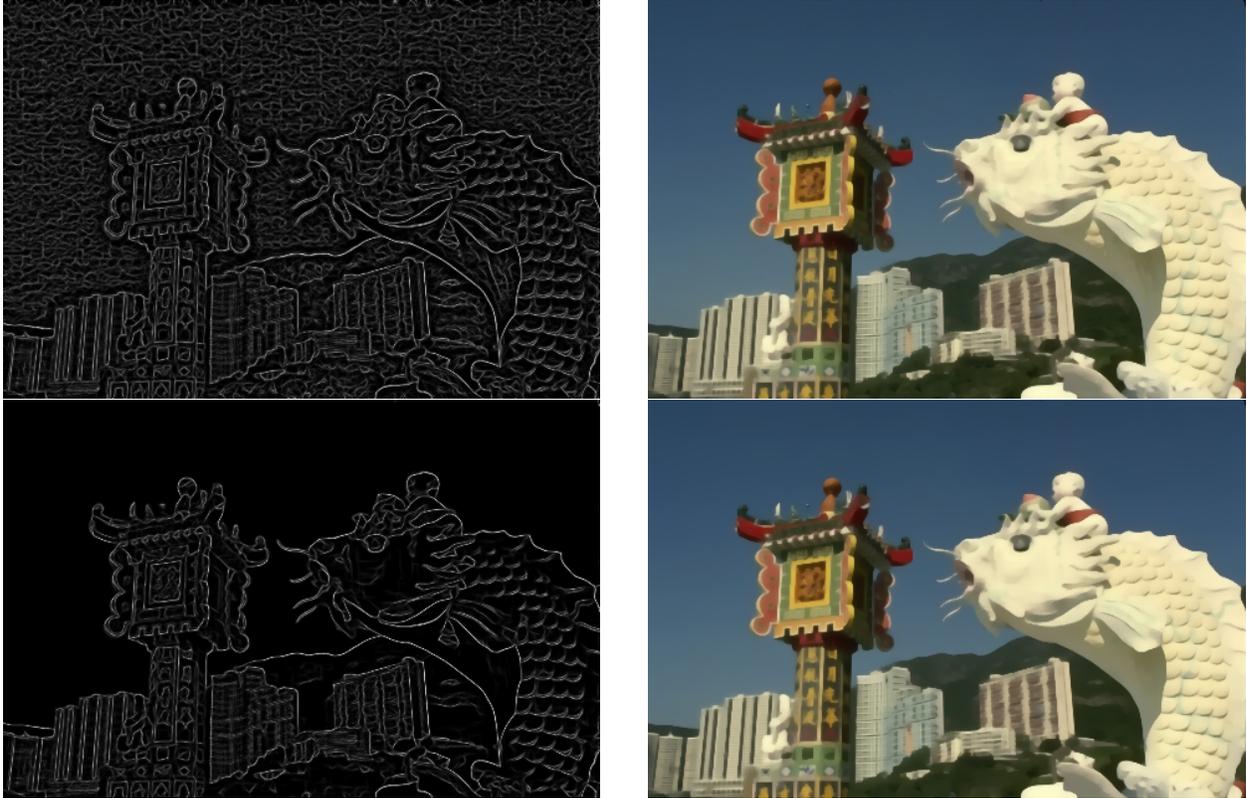


Figure S4: Comparison of our model without the refinement stage (top) and with it (bottom). The refinement stage improves the boundary map by suppressing boundaries in uniform regions and preventing additional spurious boundaries. In the low-noise regime the refinement stage does not significantly impact the boundary-aware smoothing result, meaning that smoothing using a field of junctions can be accelerated significantly by omitting the refinement stage.



Figure S5: Field of junctions obtained with the refinement stage of our algorithm (middle row) and without it (bottom row), on a synthetic image with increasing noise levels. The refinement stage reduces superfluous boundaries across all noise levels, and significantly improves the prediction near true boundaries in the high-noise regime.

## 4 Parametrization of indicator functions during refinement

As described in Section 4.2 of the main paper, we replace the indicator functions  $\{\mathbf{u}_\theta(\mathbf{x})\}$  with smooth functions based on level sets. Figure S6 provides an example of this process for  $M = 3$  with junction angles  $\phi^{(1)} = 0$ ,  $\phi^{(2)} = \pi/2$ ,  $\phi^{(3)} = 5\pi/4$ , and the junction's vertex is chosen at the center of the patch.

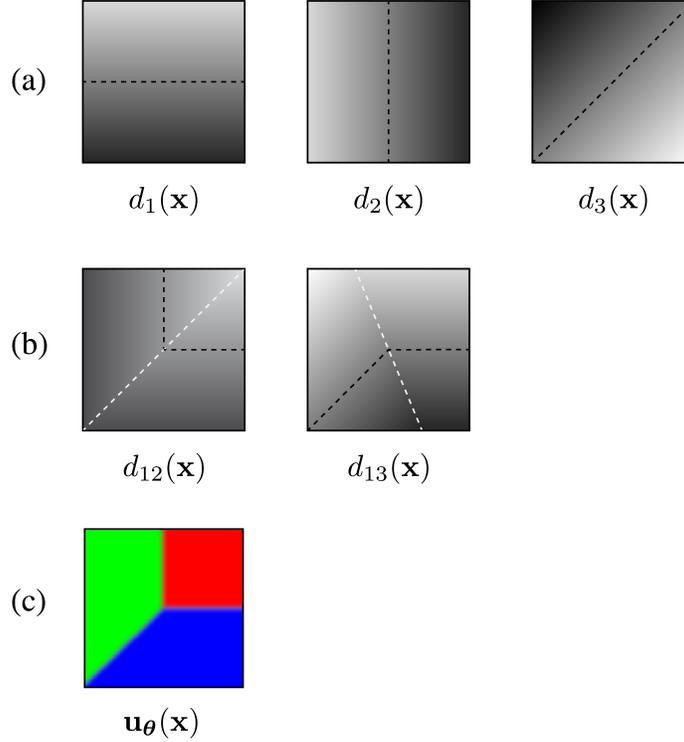


Figure S6: Indicator function relaxation based on level-set functions. (a) For the  $k$ th angle  $\phi^{(k)}$  we define a signed distance function from a line passing through the vertex, oriented at angle  $\phi^{(k)}$ . (b) Each distance function  $d_k(\mathbf{x})$ ,  $k > 1$  is combined with the first distance function  $d_1(\mathbf{x})$ , using Equation 11 from the main paper, to obtain a distance function  $d_{1k}(\mathbf{x})$  that is zero-valued at the boundary of the wedge that spans angle  $\phi^{(1)}$  through  $\phi^{(k)}$ . (c) We combine all distance functions  $\{d_{1k}(\mathbf{x})\}$  using Equation 12 from the main paper into  $M$  smooth indicator functions  $\mathbf{u}_\theta(\mathbf{x})$ , plotted here as a three-channel RGB image, with the blending of the different channels determined by the width of the regularized Heaviside function  $\eta$ . Black dashed lines mark the zero-level set of their respective distance functions, and white dashed lines in (b) show the set of points  $\mathbf{x}$  equidistant from the 1st and  $k$ th lines, *i.e.*, those satisfying  $d_1(\mathbf{x}) = -d_k(\mathbf{x})$ .

## 5 Additional results for Figure 1

We include additional results corresponding to the photograph in Figure 1 of the main paper. Figure S7 depicts the corners/junctions and boundary-aware smoothing computed by the field of junctions, and a comparison to existing corner/junction detectors (Harris [S4], ASJ [S16]) and smoothers ( $L_0$  [S15], Bilateral filter [S14]), all preceded by an optimally-tuned denoiser [S3].

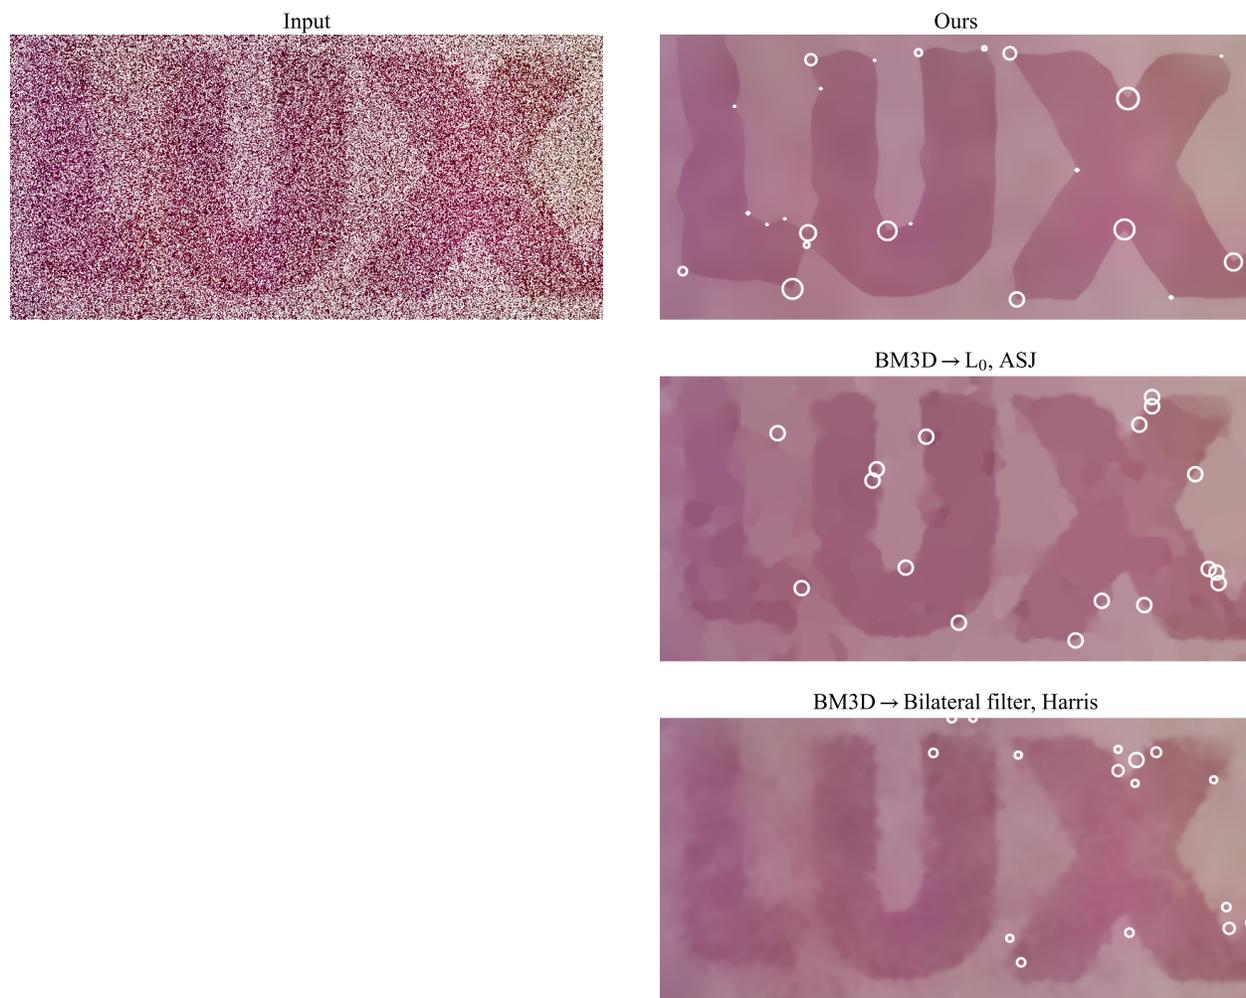


Figure S7: Corner and junction detection, and boundary-aware smoothing, for the input image from Figure 1 of the main paper. The smoothed images are shown in color, and the junctions and corners are marked by white circles. For the field of junctions and the Harris corner detector, the radius of each circle is proportional to the confidence of the vertex it denotes.

## 6 Additional comparisons with increasing noise

Figures S8 and S9 provide an extended comparison of our method to previous methods, similar to Figure 2 of the main paper. The field of junctions is more stable than previous methods at detecting boundaries, corners and junctions, and smooth colors, even when those are preceded by a denoiser [S3] equipped with the true PSNR of the image. While most previous methods are somewhat successful at moderate noise levels when preceded by such a denoiser, when they are run directly on the noisy input image they all fail even at low noise levels (see for example gPb in Figure S8).

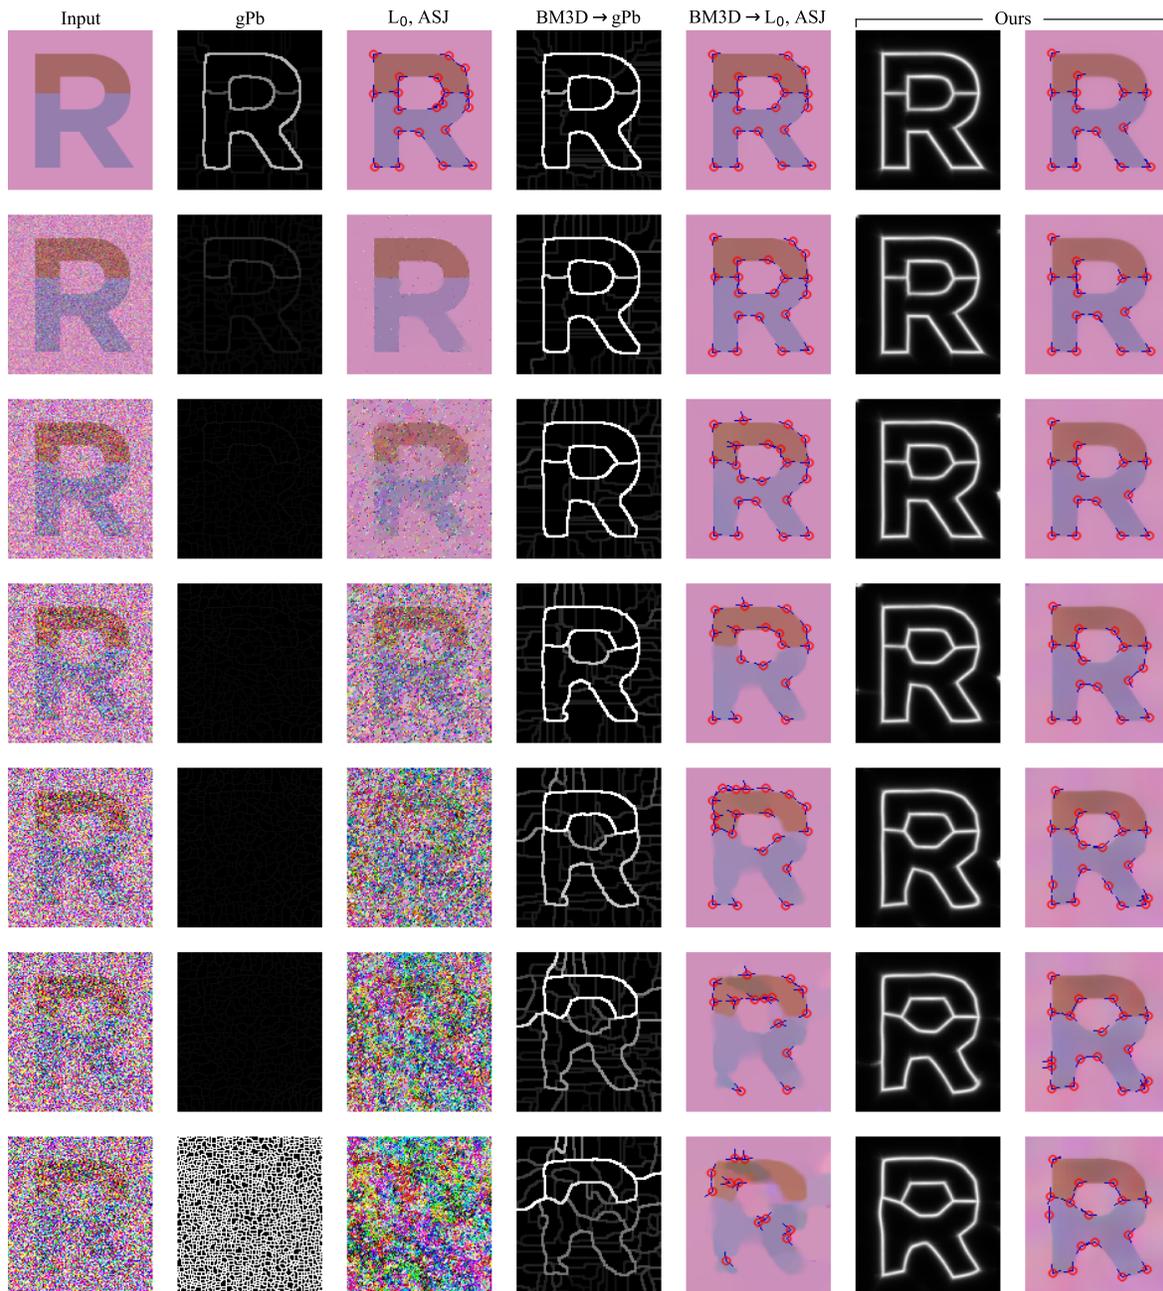


Figure S8: An extended version of Figure 2 from the main paper, over a larger range of noise levels. The field of junctions identifies boundaries (column 6), corners/junctions (circles and angles, column 7), and smooth colors (colors, column 7). These are more stable compared to a contour detector (gPb [S2]), junction detector (ASJ [S16]), and boundary-aware smoother ( $L_0$  [S15]), with and without first applying a denoiser to the input (BM3D [S3]).

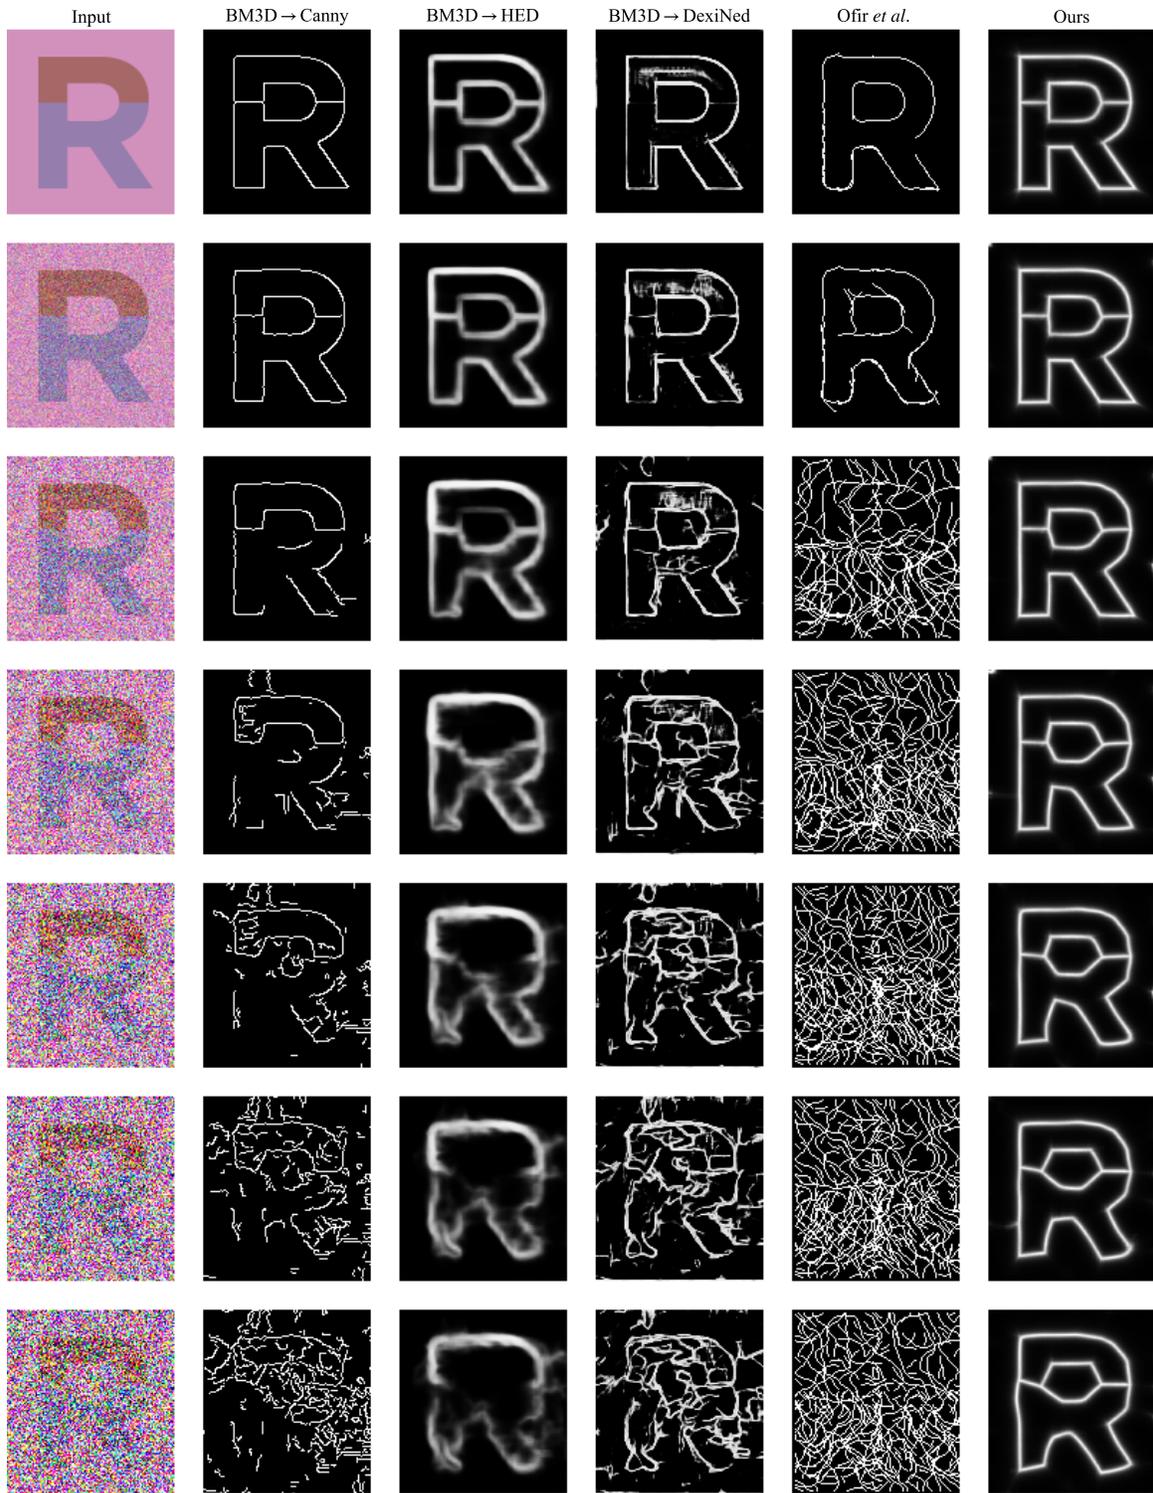


Figure S9: Additional comparisons to previous boundary detectors over a large range of noise levels. The methods above include all methods from Figure 1 of the main paper, except for Chan-Vese which does not work when there are more than two image regions.

## 7 Additional experiments with correlated noise

Our analysis algorithm is developed using an additive white Gaussian noise model, and in the paper we derive a closed-form solution for the optimal wedge colors under that noise model. In practice we find that the algorithm performs well under other spatially-independent noise models as well, including photographic ISO noise as shown in Figure 1 of the main paper and supplement Figure S7.

Here we evaluate performance when the noise is spatially correlated. Figure S9 shows the F-score obtained by our method and previous approaches (similar to the bottom half of Figure 7 in the main paper) on boundary detection in images corrupted by single-scale Perlin noise [S11] at different PSNR values and increasingly large spatial noise scales. The figure insets show examples for each noise level and spatial scale.

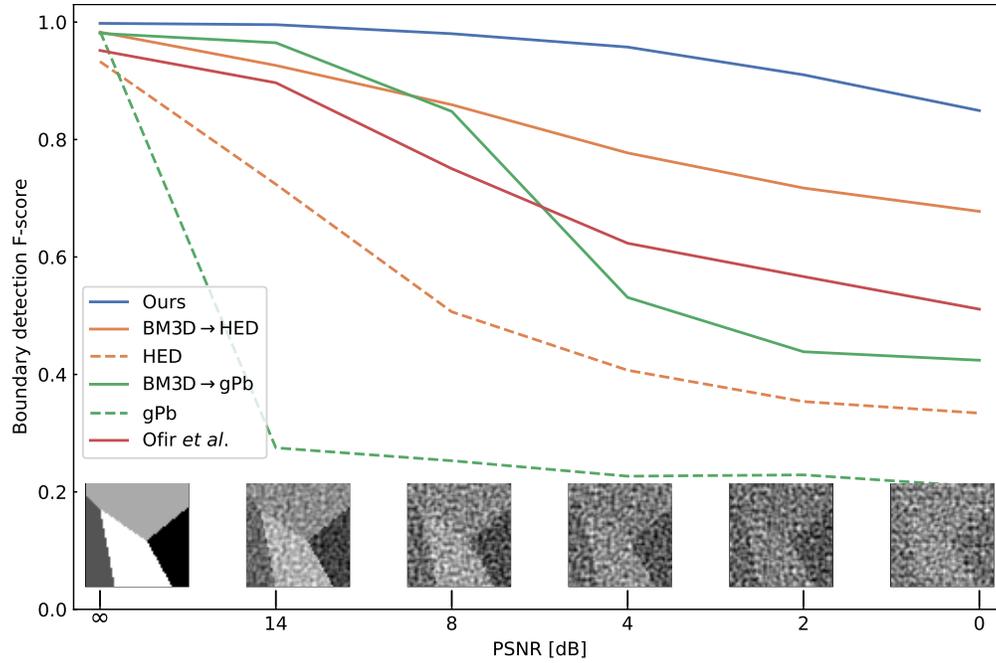


Figure S9

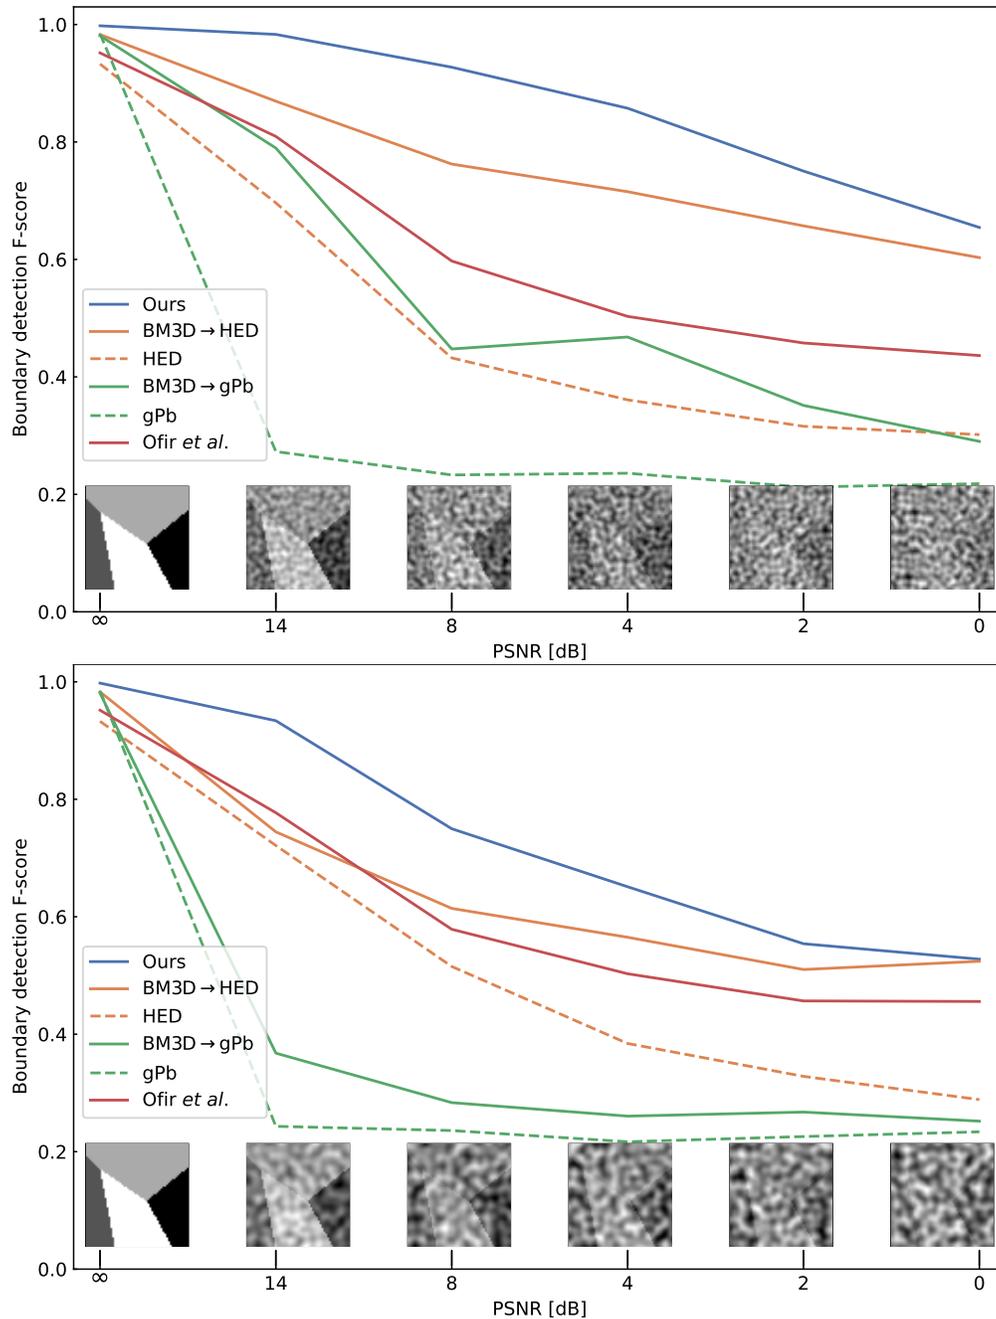


Figure S9: (*cont.*) Boundary detection F-score for our dataset corrupted by increasing levels of Perlin noise, at small spatial scale (top), intermediate spatial scale (center), and large spatial scale (bottom). Our model is comparable to previous edge detectors at low noise but outperforms them significantly at high noise, even when the previous methods are preceded by BM3D denoising. At very high noise levels and large noise scale (bottom), all methods obtain a relatively low F-score, and estimating the underlying boundary structure from the corresponding noisy image patch is extremely difficult even for a human observer. Insets: sample patch at different noise levels and spatial noise scales.

## 8 Runtime and effect of stride

Figure S10 shows the effect of using a strided field of junctions for various stride factors  $s$ . The runtime of our method as a function of  $s$  for a fixed image size and patch size  $R$ , is shown in Figure S11. Setting  $s > 1$  significantly accelerates optimization, and does not significantly affect the results, especially at high or intermediate SNRs. Finally, Figure S12 shows runtime as a function of image size  $N$ . Because our algorithm computes all junction updates in parallel, runtime grows very slowly as the image size  $N$  is increased.



Figure S10: Boundary map obtained from the field of junctions for increasing noise levels, and various strides  $s$ . At high noise levels the boundaries slowly degrade as  $s$  is increased, with a significant runtime improvement (see Figure S11).

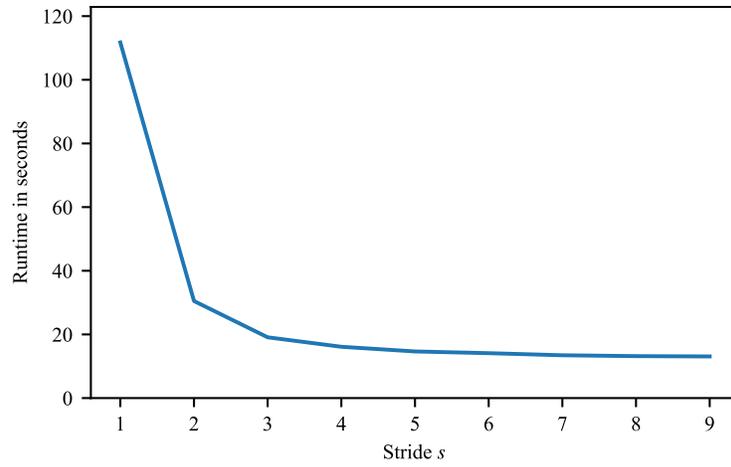


Figure S11: Runtime (in seconds) for different stride values  $s$ . Runtimes were computed using an NVIDIA Tesla V100 GPU on an image of size  $192 \times 192$  with patch size  $R = 21$ .

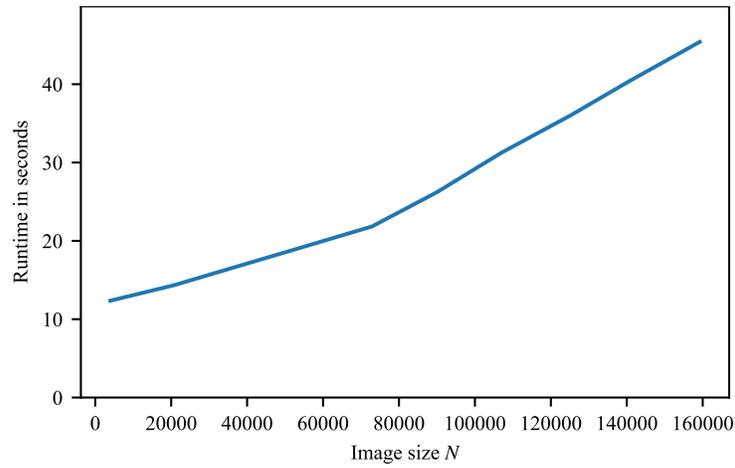


Figure S12: Runtime (in seconds) for different image sizes  $N$ . Runtimes were computed using an NVIDIA Tesla V100 GPU using stride  $s = 3$  and patch size  $R = 21$ .

## 9 Complete specification of vertex detector

The vertex probability for each pixel  $V(\mathbf{x})$  is described in Equation 13 of the main paper. We repeat the expression here with a full explanation of all of its terms. We set:

$$V(\mathbf{x}) = \sum_{i=1}^N w_i \kappa(\mathbf{x} - \mathbf{x}_i^{(0)}), \quad (\text{S12})$$

$$\kappa(\Delta\mathbf{x}) = \exp\left(-\frac{\|\Delta\mathbf{x}\|^2}{2\gamma^2}\right), \quad (\text{S13})$$

$$w_i = \exp\left(-\frac{\|\mathbf{x}_i^{(0)} - \mathbf{p}_i\|^2}{2\nu_d^2}\right) \cdot \max_{k \neq l} \min \left\{ 1, \left[1 + \cos(\phi^{(k)} - \phi^{(l)})\right] \left[1 - \left|\cos(\phi^{(k)} - \phi^{(l)})\right|\right]^{\nu_e} \right\}, \quad (\text{S14})$$

where  $\mathbf{p}_i$  is the position of the center of the  $i$ th patch and  $\gamma, \nu_d, \nu_e$  are parameters.

Each weight  $w_i$  is composed of two terms, such that a junction only contributes to the detector if its vertex position is sufficiently close to the center of the patch, and at least a pair of its wedges has an angle not too close to 0 or  $\pi$ . The first term is simply a Gaussian term in the distance from the vertex to the patch center, and the second is a function that returns 1 if any of the junction's wedges has an angle close to  $\pi/2$  and drops to 0 if all angles are 0 or  $\pi$ . The second term is chosen to be asymmetric around  $\pi/2$  since angles larger than  $\pi/2$  are more likely to be caused by structure in the image with nonzero curvature than those below  $\pi/2$ , and thus should have a lower weight. In all experiments in our paper we use  $\gamma = 2R, \nu_d = R/2, \nu_e = 2$ .

## 10 Synthetic dataset

Figure S13 shows representative images selected from our synthetic dataset. Our dataset is designed to provide control over noise levels and allow access to the ground truth edges and vertices, so that accurate quantitative analysis can be made. Our dataset contains patches that our field of junctions can handle exactly (uniformly-colored regions, edges, corners, and junctions), as well as ones that it can only model approximately (curved edges, curved junctions, *etc.*).

Our dataset contains 300 grayscale images, 100 of each of three types:

1. Images containing two curved objects of gray levels 128 and 255 on a black background (see Figure S13(a)).
2. Images containing two rotated squares of gray levels 128 and 255 on a black background (see Figure S13(b)).
3. Images containing 4 regions of gray levels 0, 85, 170, and 255 separated by two 3-junctions (see Figure S13(c)).

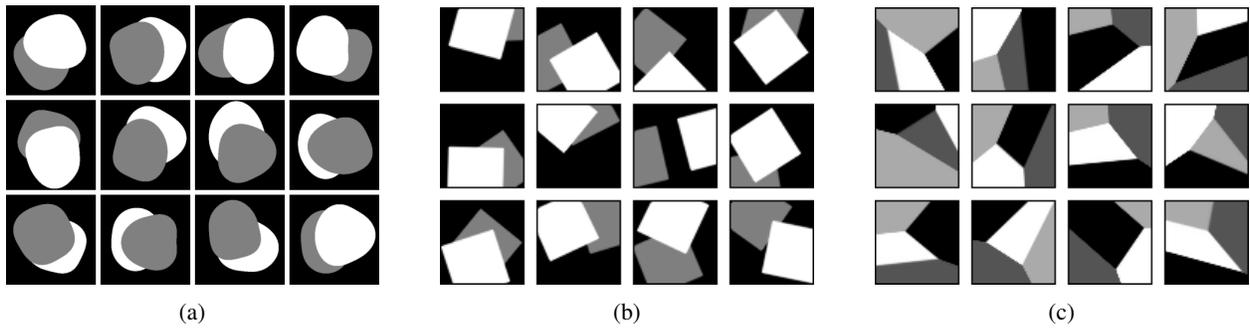


Figure S13: Images from our synthetic dataset.

## 11 Comparison with the 4-junction model

In our main paper all experiments are done using the 3-junction model (*i.e.* we set  $M = 3$ ). However, using a 4-junction model ( $M = 4$ ) can provide an advantage in certain applications which tend to contain X-junctions. Some examples for such cases are images containing specific textures like checkerboards, semi-transparent objects, or reflections from transparent ones [S8]. Figure S14 demonstrates how our model behaves in one such example taken from SIDD [S1].

Figure S15 shows a comparison of a field of junctions of a photograph analyzed using our  $M = 3$  model with that obtained by our  $M = 4$  model. The differences between the two are usually insignificant in photographs, since those typically do not contain many salient X-junctions.

Figures S16 and S17 show a comparison of the two models with increasing noise levels, on synthetic and real images. In images not containing X-junctions both models perform similarly, with the 3-junction model slightly outperforming the 4-junction model at very high noise levels. Because the 4-junction model contains an additional angle it tends to either output more superfluous boundaries than the 3-junction model (as in Figure S17), or in some cases less of them (as in Figure S16).

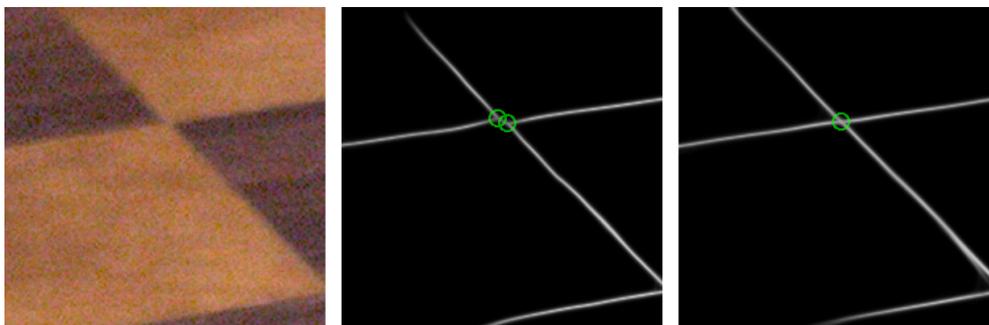


Figure S14: Close-up of the output of our model with  $M = 3$  (center) and with  $M = 4$  (right) on a photograph from SIDD. Here, the latter allows for better localization of the X-junction and the boundaries in its neighborhood



Figure S15: Boundary maps obtained by our field of junctions using the 3-junction and 4-junction models. The boundary maps obtained from the input image (left) using  $M = 3$  and  $M = 4$  are plotted (right) in the red and green channels respectively. Salient boundaries generally appear yellow, corresponding to a perfect agreement between the two models. Some boundaries appear green, corresponding to ones predicted only by the 4-junction model, which tends to suppress small-scale details less than the 3-junction one.

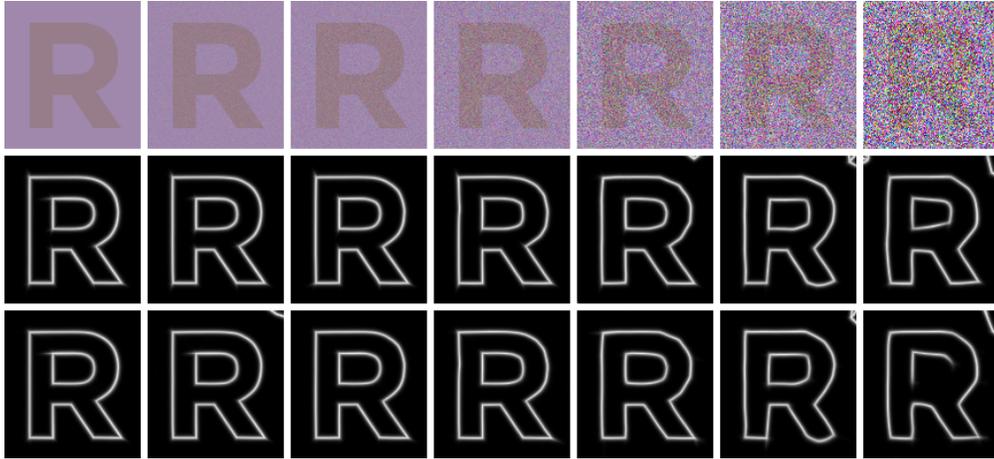


Figure S16: Boundary maps obtained using the 3-junction and 4-junction models on a synthetic image under increasing noise. At low noise both models produce similar results, but at very high noise, the 3-junction model performs slightly better.

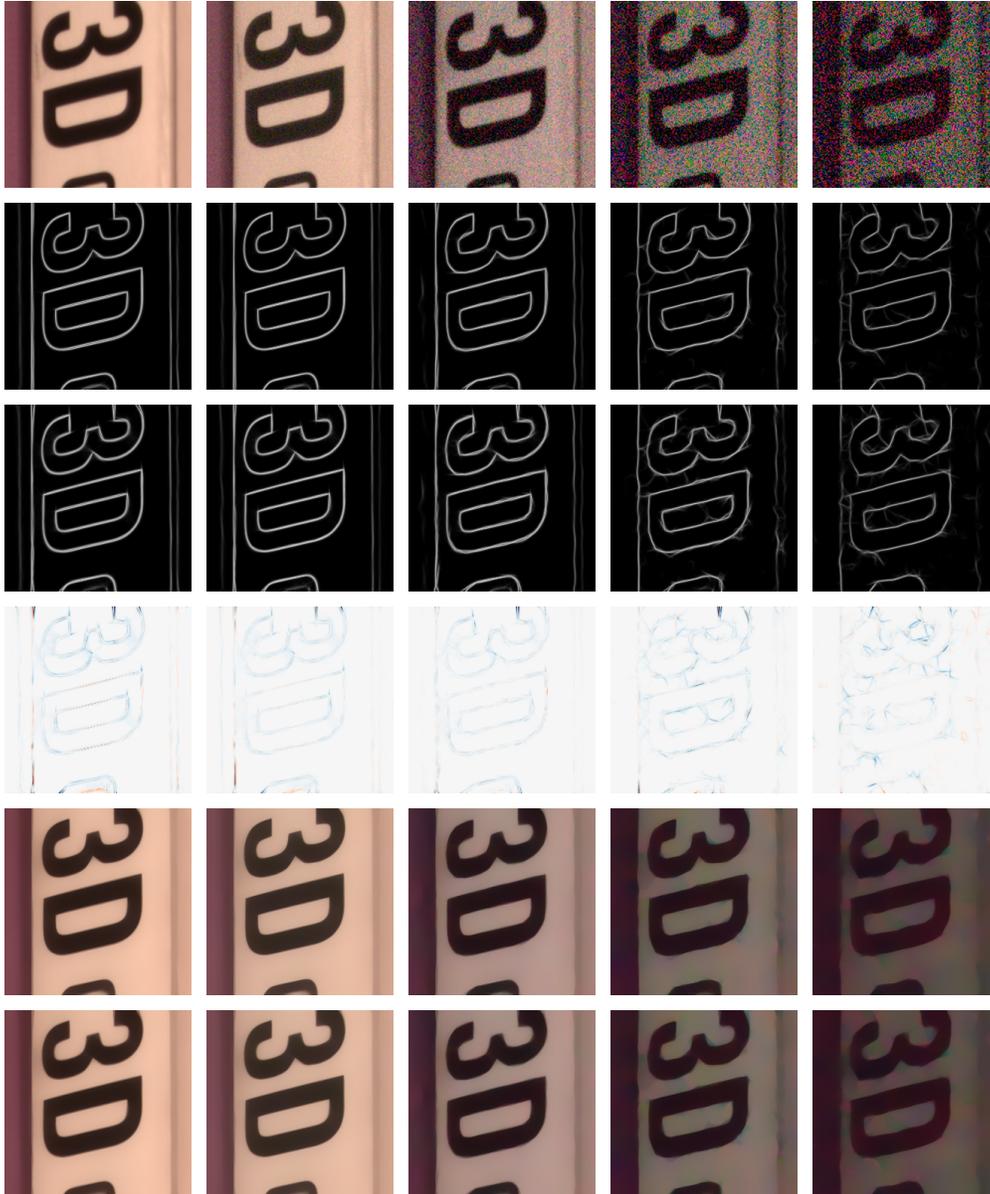


Figure S17: Boundary maps obtained using the 3-junction and 4-junction models on images from SIDD. The boundary maps output by the 4-junction model (2nd row) are similar to those output by the 3-junction model (3rd row), but typically contain more boundaries. The difference between the two boundary maps is shown in the 4th row, where blue pixels are those predicted as boundaries by the 4-junction model and not by the 3-junction model, and red pixels correspond to the opposite case. The boundary-aware smoothing results (using  $M = 3$  in 5th row,  $M = 4$  in 6th row) are not significantly affected by the number of angles.

## 12 Comparison with the linear color model

In our main paper all experiments are done using the constant color model. However, in certain situations higher-order models are beneficial. In fact, any color model where the color functions are linear combinations of functions of the coordinates has a closed-form solution for the color functions. One simple example is a linear color model, where the color in each wedge is a linear function of the coordinates  $\mathbf{x} = (x, y)$ . The linear color model outperforms the constant color model in specific applications. We first present a closed-form solution for the color functions in the linear color model, and we then show an example where the linear color model outperforms the constant one.

### 12.1 Derivation

We begin by deriving a closed-form solution for the optimal color functions under the linear model, corresponding to Equation 10 of the main paper.

Similar to Equation 9 in the main paper, we can find the optimal color functions separately for each patch  $i$  and wedge  $j$ . The optimal color function  $c_i^{(j)}$  at the  $j$ th wedge of the  $i$ th patch is the solution to:

$$\min_{c_i^{(j)}} \int u_{\theta_i}^{(j)}(\mathbf{x}) \left\| c_i^{(j)}(\mathbf{x}) - I_i(\mathbf{x}) \right\|^2 d\mathbf{x} + \lambda_C \int u_{\theta_i}^{(j)}(\mathbf{x}) \left\| c_i^{(j)}(\mathbf{x}) - \hat{I}_i(\mathbf{x}) \right\|^2 d\mathbf{x}, \quad (\text{S15})$$

where the boundary consistency term was dropped, as it does not depend on the color functions  $C$ .

Assuming a linear color model means we can write  $c_i^{(j)}(\mathbf{x}) = a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)}$ , where  $a_i^{(j)}, b_i^{(j)}, d_i^{(j)} \in \mathbb{R}^K$ , for all  $\mathbf{x} = (x, y)$ . Plugging into Equation S15, we get:

$$\min_{a_i^{(j)}, b_i^{(j)}, d_i^{(j)}} \int u_{\theta_i}^{(j)}(\mathbf{x}) \left\| a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)} - I_i(\mathbf{x}) \right\|^2 d\mathbf{x} + \lambda_C \int u_{\theta_i}^{(j)}(\mathbf{x}) \left\| a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)} - \hat{I}_i(\mathbf{x}) \right\|^2 d\mathbf{x}. \quad (\text{S16})$$

The objective in Equation S16 is quadratic in  $a_i^{(j)}, b_i^{(j)}$ , and  $d_i^{(j)}$ , and therefore it suffices to set its gradient with respect to the three coefficients to zero, yielding three equations linear in  $a_i^{(j)}, b_i^{(j)}, d_i^{(j)}$ :

$$\begin{aligned} \iint u_{\theta_i}^{(j)}(\mathbf{x}) \left( a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)} - [I(\mathbf{x}) + \lambda_C \hat{I}_i(\mathbf{x})] \right) x dx dy &= 0, \\ \iint u_{\theta_i}^{(j)}(\mathbf{x}) \left( a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)} - [I(\mathbf{x}) + \lambda_C \hat{I}_i(\mathbf{x})] \right) y dx dy &= 0, \\ \iint u_{\theta_i}^{(j)}(\mathbf{x}) \left( a_i^{(j)}x + b_i^{(j)}y + d_i^{(j)} - [I(\mathbf{x}) + \lambda_C \hat{I}_i(\mathbf{x})] \right) dx dy &= 0. \end{aligned} \quad (\text{S17})$$

The set of three linear equations in S17 can be solved by separately inverting a  $3 \times 3$  matrix and multiplying it by a 3-vector for each of the  $K$  components of  $a_i^{(j)}, b_i^{(j)}$ , and  $d_i^{(j)}$ .

## 12.2 Results

An example for a natural application for the linear color model is when the images contain large regions in which the colors varies linearly with position. Figure S18 shows an example for this type of image. Our linear color model perfectly explains the image, while the constant color model provides a number of additional low-confidence boundaries throughout the image.

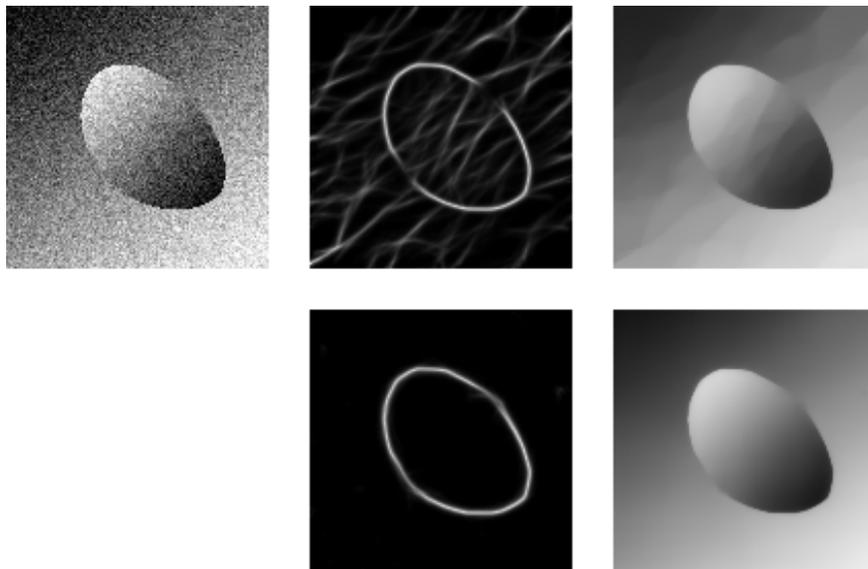


Figure S18: Field of junctions for an image containing linear color regions (left) using the constant color model (top row) compared with the linear color model (bottom row). The linear color model provides a better result in this type of application, but we found the constant color model to be sufficient for the experiments in our paper, and specifically for photographs. Input image taken from [S10].

### 13 Application to RGB-D images

In our main paper all experiments are done using grayscale ( $K = 1$ ) or RGB images ( $K = 3$ ). We find that our field of junctions is also useful for RGB-D images ( $K = 4$ ), allowing recovery of missing depth values from datasets such as the Middlebury Stereo Dataset [S13]. Figure S19 shows the depth values obtained by first analyzing an image into its field of junctions by ignoring pixels with missing depth values, and using the 4-channel color functions  $\{c_i^{(j)}\}$  to compute the global color maps of the input RGB-D image using Equation 8 in the main paper.

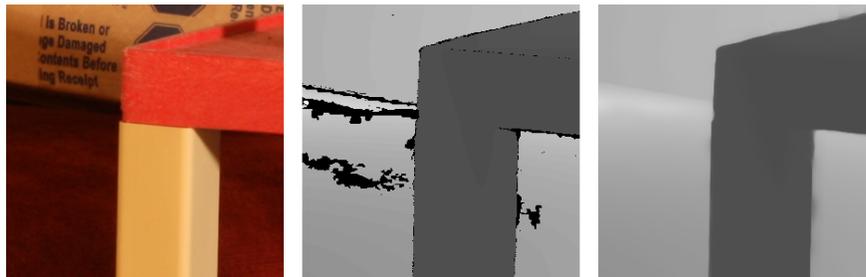


Figure S19: Analyzing an image into its field of junctions can be used for filling missing depth values in stereo datasets such as Middlebury. The input RGB image (left) and depth map with missing values (center) are analyzed into their 4-channel field of junctions. The depth channel of the resulting boundary-aware smoothed RGB-D image (right) smoothly fills the missing values.

## 14 Repeatability over change in viewpoint

We test the repeatability of our vertex detector with respect to change in viewpoint using the Graffiti dataset of [S9], which includes images of a planar surface from angles of 20 degrees to 60 degrees. Figure S20 shows the repeatability of our vertex detector compared to that of other interest point detectors, as well as the output of our vertex detector on three of the images in the dataset.

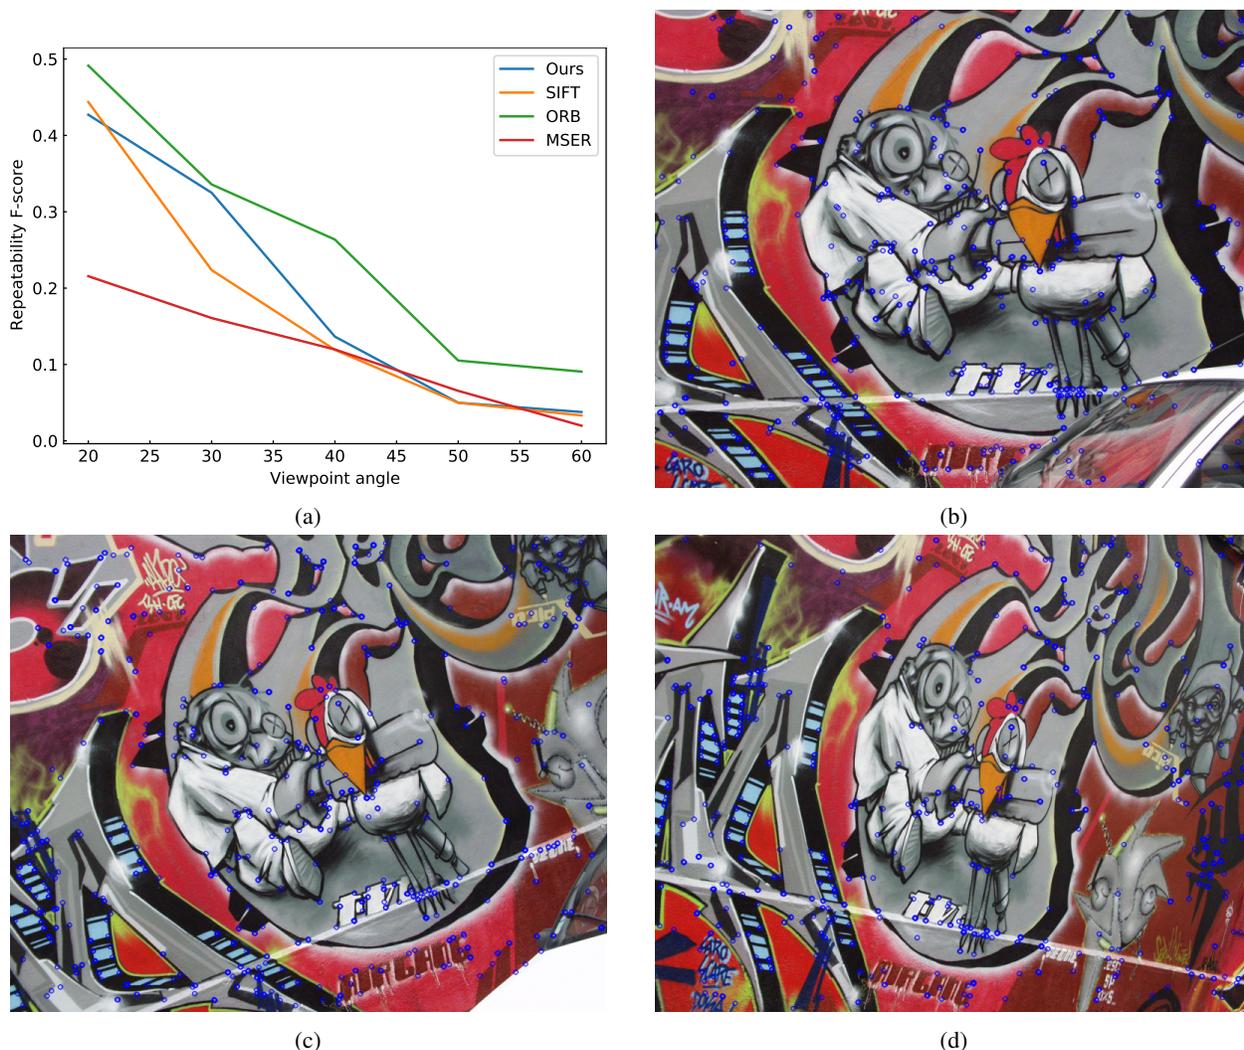


Figure S20: Comparison of repeatability over viewpoint on the “Graffiti” scene from the VGG dataset [S9]. (a) A comparison of our vertex detector with SIFT [S6], ORB [S12], and MSER [S7]; (b-d) Examples of our vertex detector’s output in images captured at the frontal view and at viewing angles 20 and 30 degrees. Our repeatability is comparable to other popular interest point detectors.

## 15 Comparison with $L_0$ smoothing

Figure S21 shows a comparison of our boundary-aware smoothing results with those obtained by  $L_0$  smoothing [S15], increasing each method's structural smoothing parameter (denoted  $\lambda_B$  for the field of junctions and  $\lambda$  in [S15]). Our method also has a parameter for selecting the spatial size of the details in the smoothed image (see for example the smoothed details inside the structure in the bottom example of Figure S21).

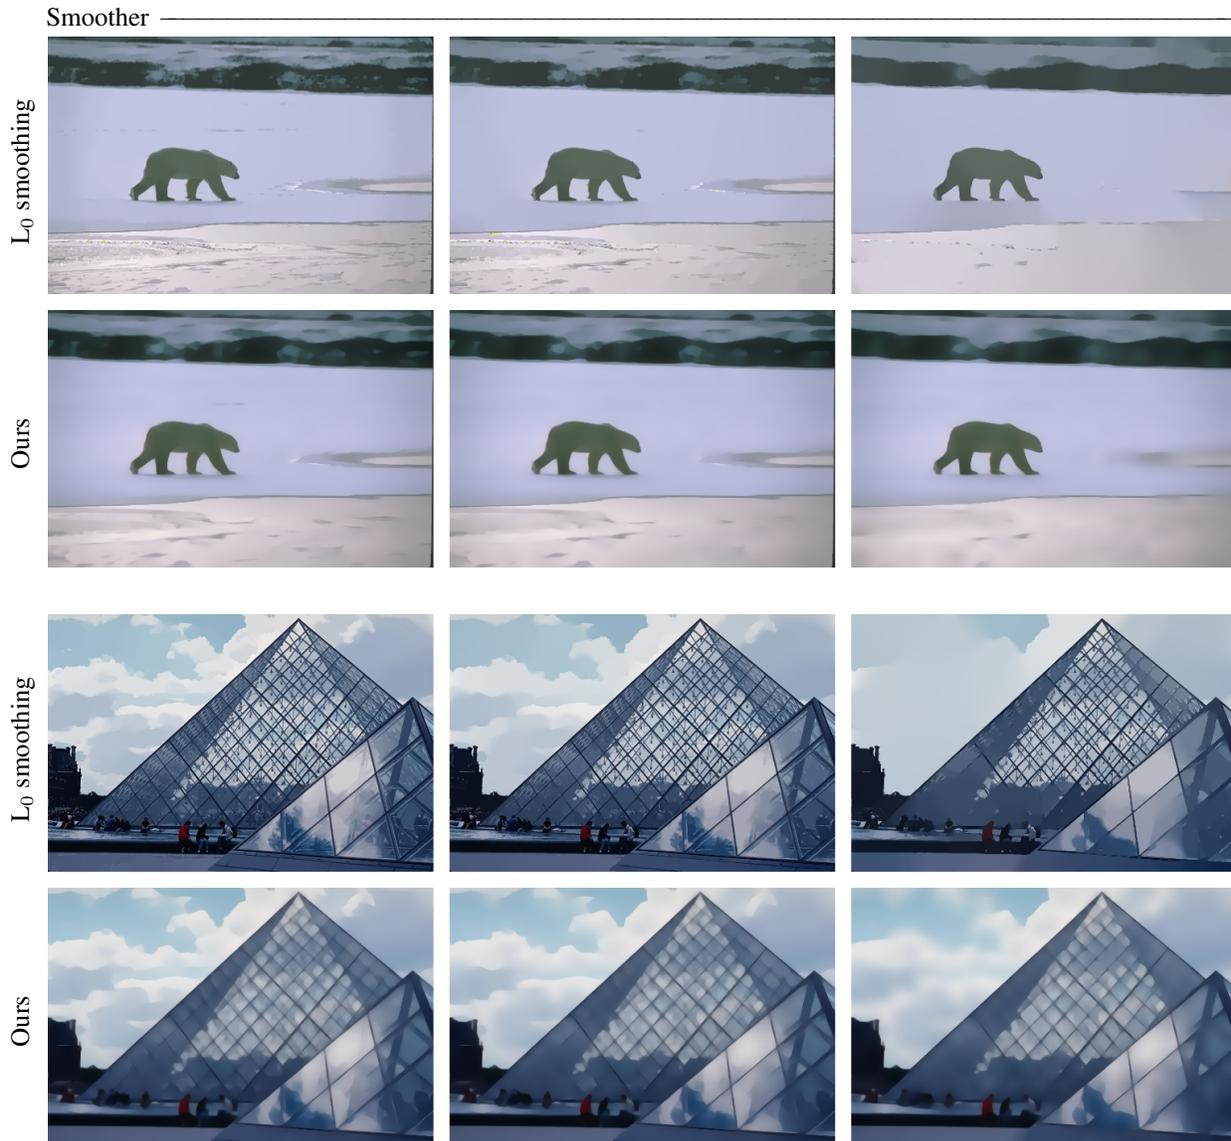


Figure S21: The smoothed results obtained by  $L_0$  smoothing (top row of each example) compared to our boundary-aware smoothing (bottom row). The smoothing parameter is increased from left to right for both methods.

## 16 Additional results on photographs

Figure S22 shows additional results of boundary-aware smoothing and boundary detection on photographs, similar to Figure 6 from the main paper. Figure S23 shows additional results obtained by the field of junctions on a noisy photograph with increasing noise levels. Finally, Figure S24 shows an artistic effect obtained by combining the smoothed image with the boundary map using the field of junctions of a photograph.

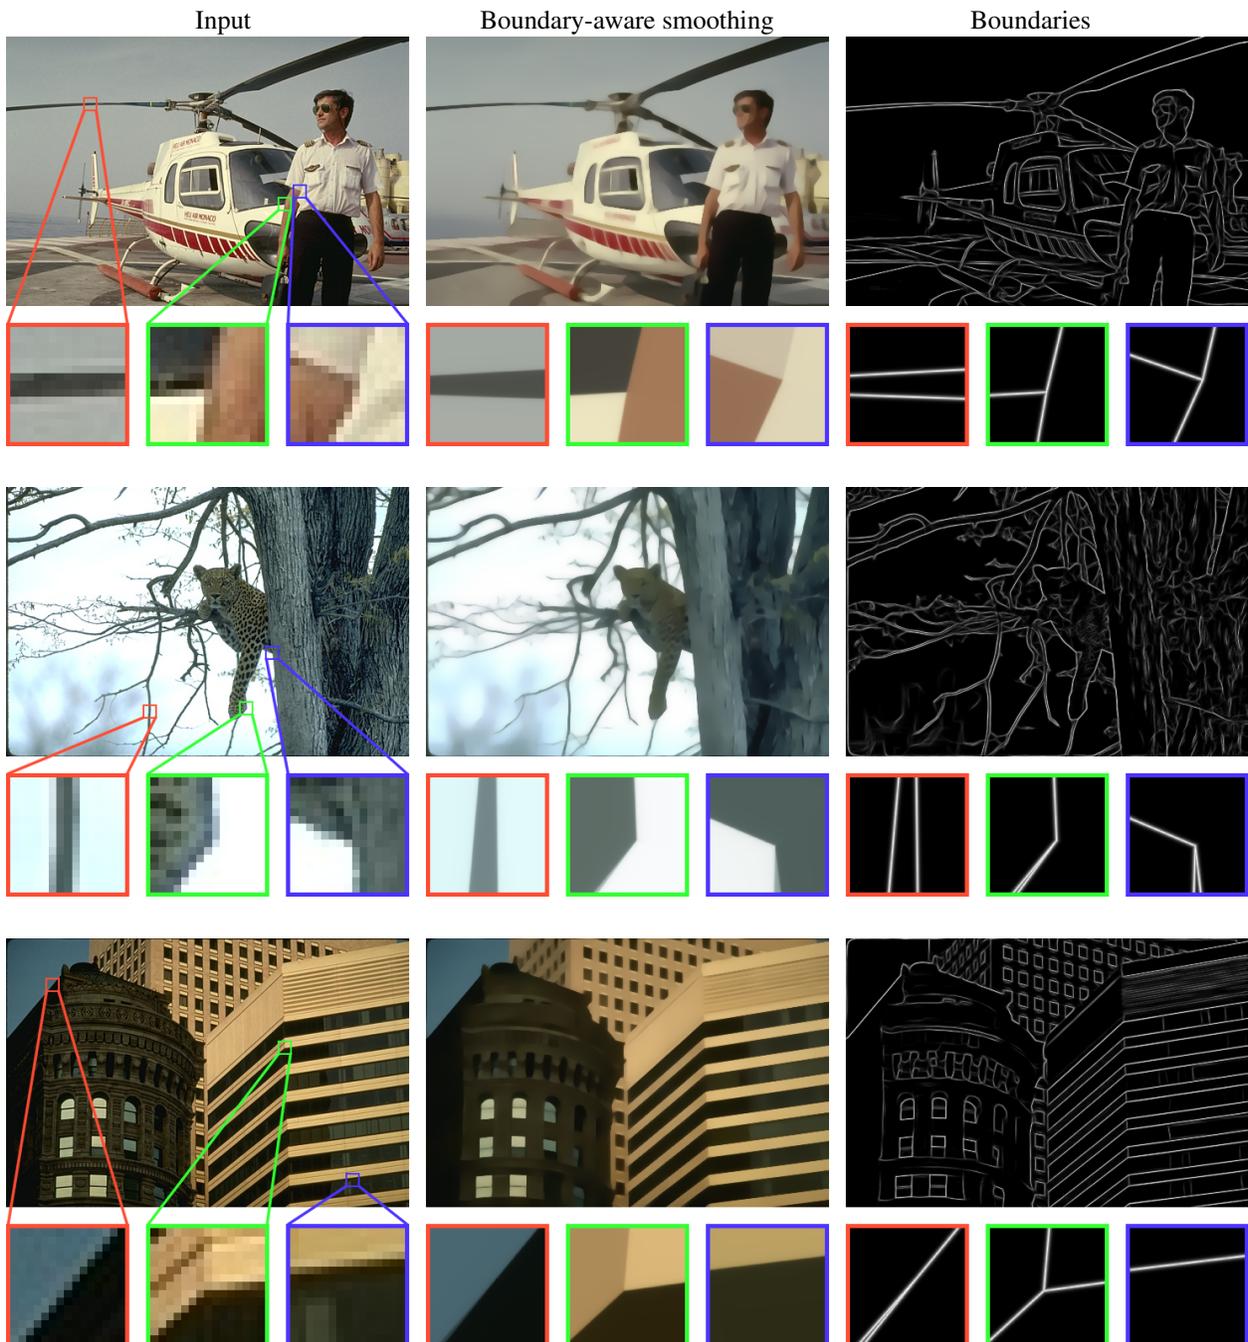


Figure S22: Additional results on photographs.

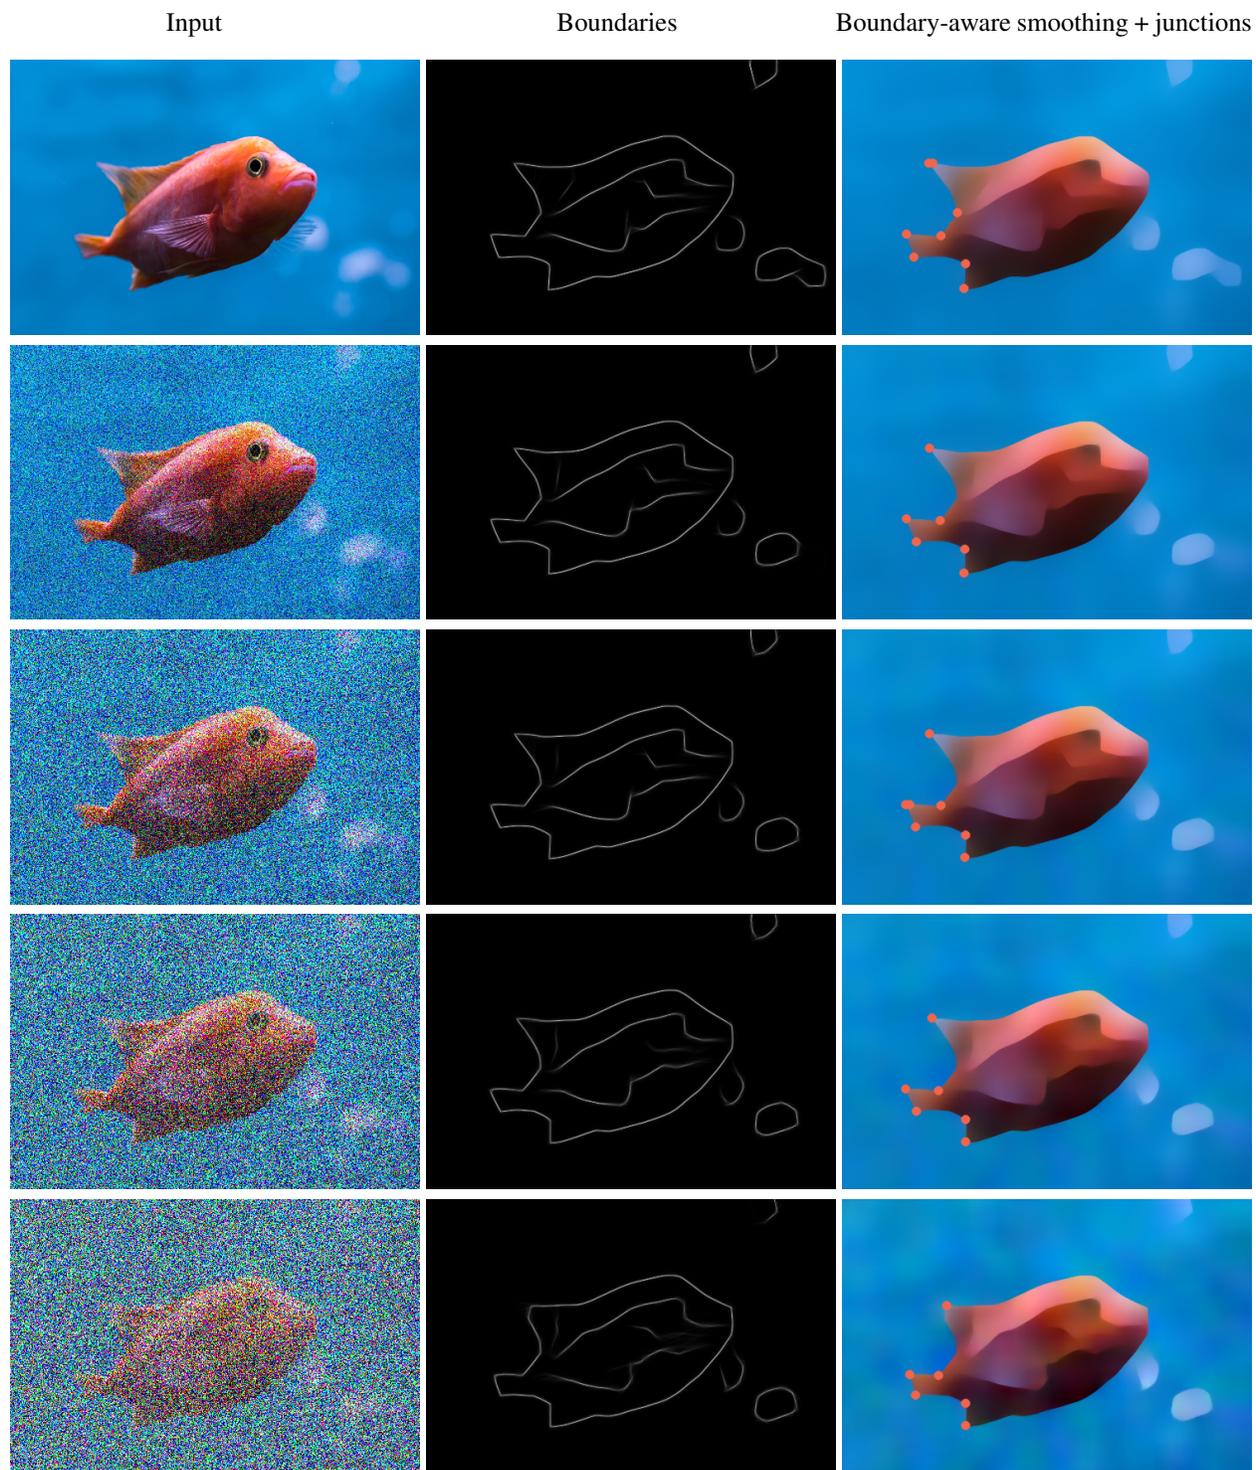


Figure S23: Additional results on an increasingly noisy photograph.

Input



Edited image



Figure S24: Non-photorealistic editing by combining the smoothed image and the boundary map obtained from the field of junctions.

## References

- [S1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1692–1700, 2018. 18
- [S2] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010. 10
- [S3] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 9, 10
- [S4] Christopher G Harris and Mike Stephens. A combined corner and edge detector. In *Alvey vision conference*, volume 15, pages 10–5244. Citeseer, 1988. 9
- [S5] Xuan He, Wei Zhu, and Xue-Cheng Tai. Segmentation by elastica energy with l1 and l2 curvatures: a performance comparison. *Numerical Mathematics: Theory, Methods and Applications*, doi: 10.4208/nmtma.OA-2017, 51, 2019. 2
- [S6] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 24
- [S7] Jiri Matas, Ondrej Chum, Martin Urban, and Tomas Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*, 22(10):761–767, 2004. 24
- [S8] Fabio Metelli. The perception of transparency. *Scientific American*, 230(4):90–99, 1974. 18
- [S9] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, Jiri Matas, Frederik Schaffalitzky, Timor Kadir, and Luc Van Gool. A comparison of affine region detectors. *International journal of computer vision*, 65(1-2):43–72, 2005. 24
- [S10] Suvadip Mukherjee and Scott T Acton. Region based segmentation in presence of intensity inhomogeneity using legendre polynomials. *IEEE Signal Processing Letters*, 22(3):298–302, 2014. 22
- [S11] Ken Perlin. An image synthesizer. *ACM Siggraph Computer Graphics*, 19(3):287–296, 1985. 12
- [S12] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *2011 International conference on computer vision*, pages 2564–2571. IEEE, 2011. 24
- [S13] Daniel Scharstein, Heiko Hirschmuller, York Kitajima, Greg Krathwohl, Nera Nesic, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German conference on pattern recognition*, pages 31–42. Springer, 2014. 23
- [S14] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998. 9
- [S15] Li Xu, Cewu Lu, Yi Xu, and Jiaya Jia. Image smoothing via  $L_0$  gradient minimization. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–12, 2011. 9, 10, 25
- [S16] Nan Xue, Gui-Song Xia, Xiang Bai, Liangpei Zhang, and Weiming Shen. Anisotropic-scale junction detection and matching for indoor images. *IEEE Transactions on Image Processing*, 27(1):78–91, 2017. 9, 10