

C2PD: Continuity-Constrained Pixelwise Deformation for Guided Depth Super-Resolution

Jiahui Kang¹, Qing Cai^{1*}, Runqing Tan¹, Yimei Liu¹, Zhi Liu²

¹Faculty of Computer Science and Technology, Ocean University of China

²School of Information Science and Engineering, Shandong University

{kangjiahui, trq, liuyimei}@stu.ouc.edu.cn, cq@ouc.edu.cn,

liuzhi@sdu.edu.cn

Abstract

Guided depth super-resolution (GDSR) has demonstrated impressive performance across a wide range of domains, with numerous methods being proposed. However, existing methods often treat depth maps as images, where shading values are computed discretely, making them struggle to effectively restore the continuity inherent in the depth map. In this paper, we propose a novel approach that maximizes the utilization of spatial characteristics in depth, coupled with human abstract perception of real-world substance, by transforming the GDSR issue into deformation of a roughcast with ideal plasticity, which can be deformed by force like a continuous object. Specifically, we firstly designed a cross-modal operation, Continuity-constrained Asymmetrical Pixelwise Operation (CAPO), which can mimic the process of deforming an isovolumetrically flexible object through external forces. Utilizing CAPO as the fundamental component, we develop the Pixelwise Cross Gradient Deformation (PCGD), which is capable of emulating operations on ideal plastic objects (without volume constraint). Notably, our approach demonstrates state-of-the-art performance across four widely adopted benchmarks for GDSR, with significant advantages in large-scale tasks and generalizability.

Introduction

Depth maps are widely utilized in fields such as 3D reconstruction (Chen et al. 2020a; Im et al. 2018), autonomous driving (Kerl, Sturm, and Cremers 2013), scene understanding (Hoffman, Gupta, and Darrell 2016), and semantic segmentation (Chen et al. 2020b; Gupta et al. 2014; Weder et al. 2020). Technical and cost limitations has prompted the desire to exploit the RGB images to guide super-resolution of depth maps. Recent research on guided depth super-resolution (GDSR) predominantly centers on deep neural networks, yielding promising outcomes (Zhong et al. 2023b; Guo et al. 2018; Su et al. 2019; Zhou et al. 2023). The strategies of computing each point discretely and expanding feature channels, empower fundamental networks like CNN (LeCun et al. 1998) and Transformer (Vaswani et al. 2017) with outstanding abilities for abstract extraction, which then demonstrate excellent performance in the feature extraction of RGB images (Dong et al. 2015; Liang et al. 2021; Lyu

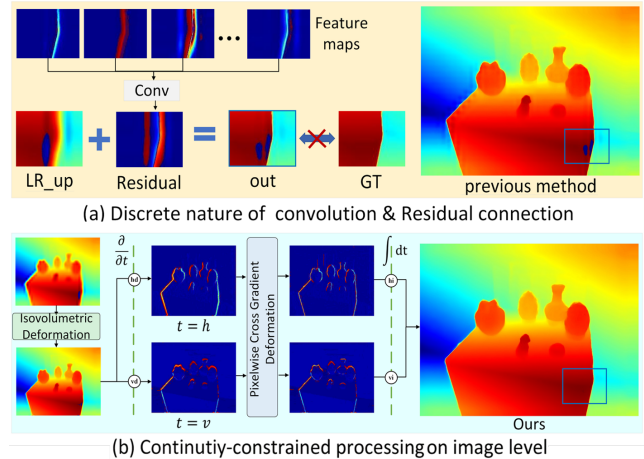


Figure 1: Illustration of some issues of GDSR and our method to address them. (a) After convolution, each point on residual is only obtained by corresponding area of input, which means the relationship of points in output is discrete. Residual connections also increase the difficulty of residual fitting. Those operations treat depth discretely, easy to induce deviation after reconstruction. (b) Our model established a systematic architecture for continuity-constrained deformation with fully integrated operational modules.

et al. 2023; Cai et al. 2022, 2023, 2024). Successes in RGB super-resolution have prompted existing methods to process depth at the feature level similarly, merging features of depth and RGB, ultimately reducing channels to obtain the target residual (Lim et al. 2017) or depth map.

Although existing methods have achieved impressive performance, as shown in Fig. 1, they often approach depth maps as images where each value is computed discretely, thereby inducing deviation. The discrete approaches in previous works led computers to consider depth map as a series of values on a plane, requiring extensive training to learn how to adjust the current value based on other information, lacking the ability to generate holistic spatial awareness akin to humans. In contrast, humans focus on the morphological characteristics rather than specific pixel values. Specifically, when observing low-resolution depth (LR) and ground truth (GT), humans establish associations based on their overall

*Corresponding author

morphological features, unaffected by deviations in object edges or color. This is because our brains can abstract visual information into an object with ideal plasticity, allowing us to imagine the process of LR transforming into GT. Similarly, even when replacing GT with RGB images, we can still envision how LR transforms into GT based on RGB images. Indeed, we can instruct deep learning models to restore LR based on this human awareness. And fortunately, LR, as spatial information already captured, can be directly treated as the initial form of the plasticity model.

Motivated by preceding discussion, we conceptualize the depth as ideal plastic substance, capable of reshaping without volume constraints. To integrate continuity constraints into our model, we consider particle interactions and devise the Continuity-constrained Asymmetrical Pixelwise Operation (CAPO) aimed at reshaping the plastic material while maintaining equal volume under external forces. This operation excels in guiding deformation by strictly adhering to source image. By exploiting the properties of differential, we propose the Pixelwise Cross Gradient Deformation (PCGD), simulate the complete processing of ideal plastic material and maximize module reusability to ensure model generalization. Moreover, the guidance information utilized by our method differs from specific coloration information in previous methods, emphasizing operational guidance for relative changes (analogous to force information), highlighting abstract information like spatial trends. Consequently, adopting CAPO allows U-net to focus solely on extraction and abstraction, overcoming previous constraints (having to balance the abstract extraction with maintaining cues for reconstruction). In summary, our contributions include:

1. A Continuity-constrained Asymmetrical Pixelwise Operation (CAPO) is proposed, which incorporates human perception of continuity into neural networks and can emulate isovolumetric deformation of flexible materials.
2. Based on CAPO, a Pixelwise Cross Gradient Deformation (PCGD) is designed, which emulates deformations on objects with ideal plasticity, breaking volume constraints while preserving continuity.
3. Extensive experiments demonstrate state-of-the-art performance and exhibit increasingly prominent advantages with scale escalation, suggesting new directions for further breakthroughs in large-scale tasks such as x32.

Relative Work

Deep Learning GDSR Methods

The potent extraction capability of deep learning has propelled advancements in GDSR (Riegler et al. 2016; Hui, Loy, and Tang 2016a; Tang et al. 2021; Ye et al. 2020). Early works like DMSG (Hui, Loy, and Tang 2016b), DG (Gu et al. 2017), and DJF (Li et al. 2016, 2019), initially employed CNNs for feature extraction and direct regression towards the target image. (Deng and Dragotti 2019) proposed Deep Coupled ISTA Network based on the Iterative Shrinkage-Thresholding Algorithm (ISTA) (Gregor and LeCun 2010), and CUNet (Deng and Dragotti 2020) for separating common/unique features through sparse coding.

Furthermore, (Kim, Ponce, and Ham 2021) designed a deformable kernel-based filtering network (DKN), adaptively outputting neighborhoods and weights for each pixel. (He et al. 2021) proposed FDSR which exploits octave convolution (Chen et al. 2019) to decompose the RGB feature. (Tang, Chen, and Zeng 2021) introduced implicit neural representation to learn both the interpolation weights and values (JIIF). (Shi, Ye, and Du 2022) introduced a symmetric uncertainty method SUFT, selecting effective RGB information for HR depth recovery while skipping harmful textures. (Zhao et al. 2022) developed a discrete cosine network (DCTNet), extracting multi-modal information through a semi-decoupled feature extraction module. Most recently, (Zhao et al. 2023) proposed spherical space feature decomposition (SSDNet) to separate shared and private features. (Zhong et al. 2023a) proposed a multiscale fusion guided filter framework DAGF to reuse intermediate results in the coarse-to-fine process. (Metzger, Daudt, and Schindler 2023) proposed a hybrid approach DADA combining guided anisotropic diffusion with CNNs for optimization operations on single images. (Wang, Yan, and Yang 2024) proposed a SNet designed to capture high-frequency structure from RGB image by emphasizing gradient and frequency domain.

These methods produce pixel values through independent computations, lacking continuity constraints between adjacent regions. Consequently, they are unable to incorporate the rationality of spatial variations as a human would do.

Comparison with Existing Methods

Existing approaches (Wang, Yan, and Yang 2024; Zhong et al. 2023a; Zhao et al. 2022) typically involve feature extraction from both depth and RGB modalities, posing challenges in combining the features from RGB and depth to reconstruct continuous information. Moreover, alignment of the two modalities is required before fusion, and deviations in fusion process also affect the results. (Wang et al. 2021; Sun et al. 2022) In contrast, our method extracts operational information from RGB, providing guidance for depth variations at a global scale, rather than numerically fusing them as similar types. Thus, deviations in guidance information hardly misled depth during this operation process. Under continuity constraints, depth accepts operational guidance based on its original trend and continuity, meaning that unreasonable textures (may causing by mismatching RGB textures) in the guidance have minimal impact on the depth.

Although DADA (Metzger, Daudt, and Schindler 2023) also attempts to constrain changes in the depth map under guidance of RGB, the process of iterative diffusion is difficult to effectively control, and the resulting images obtained after diffusion cannot be directly used as output. Therefore, DADA ultimately modifies LR_{up} by multiplication, rather than directly performing diffuse operations on the source image, and fails to establish a complete and theoretically controlled system for single-image operations. In contrast, our method can directly obtain controlled outputs through CAPO, and effectively operates on gradients with PCGD, ultimately forming a controlled system of continuous deformation without volume limitation.

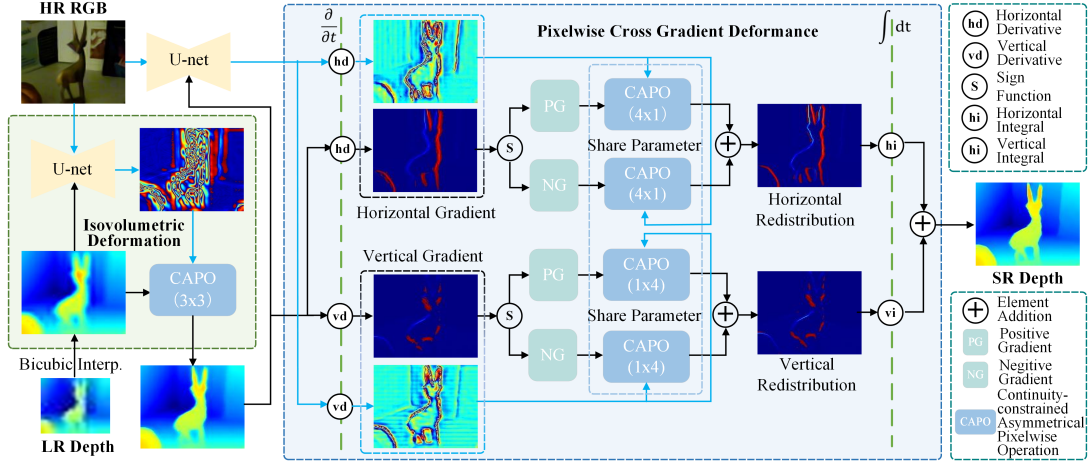


Figure 2: Network architecture of our method. The Isovolumetric Deformation initially uses CAPO to deform the depth to organize pertinent information. Then, depth undergoes further processing through PCGD in the form of gradients.

Method

Fig. 2 illustrates the architecture of our method. Firstly, we construct the Continuity-constrained Asymmetrical Pixelwise Operation (CAPO) using neural network. Subsequently, we derive the formula of the Pixelwise Cross Gradient Deformation (PCGD) and implement its structure.

Continuity-constrained Asymmetrical Pixelwise Operation

CAPO is a cross-modal operation designed to model how depth changes under the influence of guidance information, while adhering to continuity constraints. These constraints are based on human cognition of changes in real-world objects, as illustrated in Fig. 3 (a). In this illustration, the value at a point is predominantly derived from its surrounding locations rather than being directly transferred over long distances. In this process, depth and operational information are treated as input, and CAPO can fuse them in a way of guiding isovolumetric flexible object with force field.

Since CAPO(3x3), CAPO(4x1), CAPO(1x4) all share the same idea (just the shape of the receptive field is different), so next we will mainly analyze the CAPO(1x4) operation. As shown in Fig. 3 (b), the asymmetric structure empowers CAPO to learn the guidance relationship, while the constraints within a region and the effects between regions ensure the continuity of output. CAPO(1x4) can be seen as “ \circ ” of $I_h(x_i) \circ G_h(x_i)$ in formula (9). Here $I_h(x_i)$ and $G_h(x_i)$ are two inputs of CAPO, presenting image to be operated and information for guidance respectively. When considering the horizontal operation, the y coordinate is determined. To simplify the representation, we treat $I_h(x_i) \circ G_h(x_i)$ as a one-dimensional operation on the x-axis (horizontal direction). Among them, $I_h(x_i)$ is used to represent the variation of I along the x direction at position x_i , while $G_h(x_i)$ is used to represent the expected change degree at x_i . If x_i is in the edge transition area on G , then $G_h(x_i)$ will be a relatively large value, which will guide I_h to assign more values

to x_i . Guided by G_h crossing modalities, I_h ought to focus on the relative relationship in local areas of G_h , rather than directly using $G_h(x_i)$ as a value for guiding. Firstly, let us use \mathcal{N}_4^t to represent the four adjacent positions from $x=t-3$ to $x=t$, and use D^t to represent the set of values on I and G at these positions, thereby:

$$D^t = \{z | z = I_h(x) \vee G_h(x), x \in \mathcal{N}_4^t\}, \quad (1)$$

where $\mathcal{N}_4^t = \{x_{t-3}, x_{t-2}, x_{t-1}, x_t\}$, and “ \vee ” means “or” operation. To simulate the interaction of particles in a unit area under external force field (op1 in Fig. 3 (c)), we fit var_i to the target variation at position i within \mathcal{N}_4^t , then:

$$Interaction(D^t) = \{var_3(D^t), var_2(D^t), var_1(D^t), var_0(D^t)\}, \quad (2)$$

where $var_3, var_2, var_1, var_0$ present the target variations of the four positions $x_{t-3}, x_{t-2}, x_{t-1}$, and x_t respectively. Here var_i will decide how many values (particles) to allocate to its own position, based on the relationship between x_i and the entire \mathcal{N}_4^t area, with the guidance of G_h . In order to ensure the principle of volume-conserving constraint, we need to convert var_i to var_j^* :

$$var_i^*(D^t) = var_i(D^t) - \frac{1}{4} \sum_{j=0}^3 var_j(D^t). \quad (3)$$

Following the principle of particle fluidic propagation, we aim for points closer in distance x_i to exert a greater influence on x_i with higher probability. Therefore, we implement op2 in Fig. 3(b) by overlapping regional effects. Then, the value at position x_i can be expressed as:

$$I_h(x_i) \circ G_h(x_i) = I_h(x_i) + \frac{1}{4} \sum_{j=0}^3 var_j^*(D^{i+j}). \quad (4)$$

That is to say, the value at x_i obtained after the $I_h(x_i) \circ G_h(x_i)$ is actually affected by the areas from x_{i-3} to x_{i+3} , with positions closer to x_i having a higher probability of making a substantial impact. (op2 in Fig. 3(c))

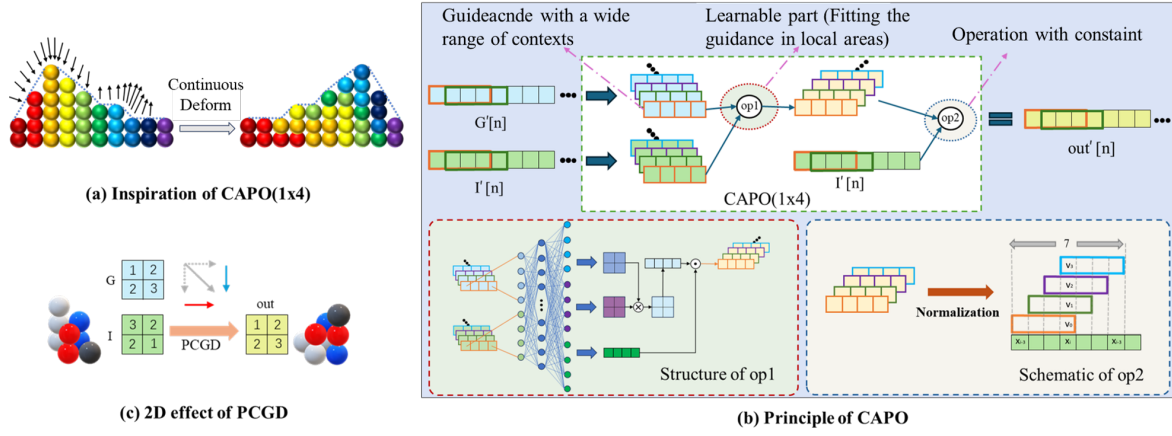


Figure 3: (a) Considering the interaction among particles, which preferentially convey collectively as a cohesive viscous fluid, the variation at each position is influenced by the combined trends in surrounding regions. (b) The principle of CAPO(1x4) with respect to local information. (c) Controlled with PCGD, each point only needs to focus on its own gradient changes, while the 2D macroscopic effect is regulated by guidance information with large contexts.

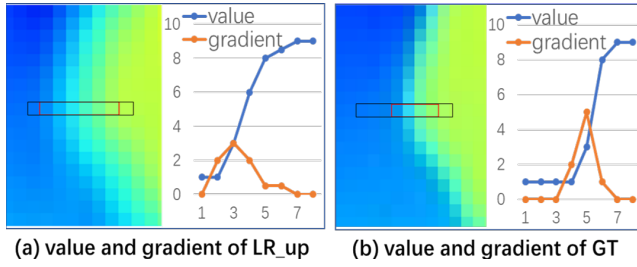


Figure 4: To transfer upsampled LR (LR_up) to GT, we should move the values to right and shrink the transition area. These changes may be shown as the blue lines, where target values in (a) are all less or equal to the current values in (b), which means we cannot achieve GT by isovolumetric deformation. However, it can be completed by transferring gradient showed as orange line so that the values around position 3 in (a) are passed to position 5, with width shrinking.

Pixelwise Cross Gradient Deformation

As shown in Fig. 4, operation on gradient empowers PCGD to genuinely “move” edges, rather than merely “altering” by blending colors. Thus, we will use this theory as inspiration to develop PCGD. Let’s start by considering real-world scenarios where images are perceived as continuous rather than discrete information. A depth image is essentially the distance from a specific plane to the surface of an object, and surface typically exhibit good continuity. Therefore, we can shift our attention towards the variation trend of depth through derivation, and subsequently adjusting these trends.

For a depth map D , since the depth information of each position can be uniquely determined by the abscissa x and ordinate y , we can represent the real depth information corresponding to the image through the binary function $h(x, y)$,

thus the depth map is viewed as a function image of $h(x, y)$:

$$D = h(x, y) = \frac{1}{2} \left[\int \frac{\partial h(x, y)}{\partial x} dx + \int \frac{\partial h(x, y)}{\partial y} dy \right], \quad (5)$$

where $\frac{\partial h(x, y)}{\partial x}$ and $\frac{\partial h(x, y)}{\partial y}$ are gradients of depth in the x and y directions. At this point we can transform the image D into what we want by adjusting $\frac{\partial h(x, y)}{\partial x}$ and $\frac{\partial h(x, y)}{\partial y}$. Now, let us apply the idea in Fig. 4 to explore the method of restoring a blurred LR_up image to a clear high-resolution image. Whether its blur caused by interpolation or other factors, the image that gives us a blurry experience is essentially a combination of valid information and biased or erroneous messages. For the blurred image I , we hope to obtain the guidance image G to guide I with the help of a specific operation. In the same way, we can use functions $f(x, y)$ and $g(x, y)$ to represent I and G respectively:

$$I = f(x, y), G = g(x, y). \quad (6)$$

In other words, the gradients of G , $\frac{\partial g(x, y)}{\partial x}$ and $\frac{\partial g(x, y)}{\partial y}$, can be utilized through an operation “ \circ ” to guide the trend of I , $\frac{\partial f(x, y)}{\partial x}$ and $\frac{\partial f(x, y)}{\partial y}$, thereby directing I convert to clear image I_{out} (In fact, as analysis in Fig. 4, we choose CAPO as “ \circ ”). Using formula (5)(6), we can get the I_{out} :

$$I_{out} = \frac{1}{2} \left[\int \left(\frac{\partial f(x, y)}{\partial x} \circ \frac{\partial g(x, y)}{\partial x} \right) dx + \int \left(\frac{\partial f(x, y)}{\partial y} \circ \frac{\partial g(x, y)}{\partial y} \right) dy \right]. \quad (7)$$

Thus, value of I_{out} at (x_p, y_p) can be expressed as:

$$I_{out}|_{(x_p, y_p)} = \frac{1}{2} \int_0^{x_p} [f'_x(x, y_p) \circ g'_x(x, y_p)] dx + \frac{1}{2} \int_0^{y_p} [f'_y(x_p, y) \circ g'_y(x_p, y)] dy. \quad (8)$$

Unlike continuous images in the real world, image processed in computer is sampled to discrete points, so we need to discretize the above formula. To simplify the expression, we let

Methods	Middlebury			Lu			NYU_v2			RGBDD		
	x4	x8	x16	x4	x8	x16	x4	x8	x16	x4	x8	x16
DKN (IJCV'2021)	1.23	2.12	4.24	0.96	2.16	5.11	1.62	3.26	6.51	1.30	1.96	3.42
FDKN (IJCV'21)	1.08	2.17	4.50	0.82	2.10	5.05	1.86	3.58	6.96	1.18	1.91	3.41
FDSR (CVPR'21)	1.13	2.08	4.39	1.29	2.19	5.00	1.61	3.18	5.84	1.18	1.74	3.05
JiIF (MM'21)	1.09	1.82	3.31	0.85	1.73	4.16	1.37	2.76	5.27	1.17	1.79	2.87
DCTNet (CVPR'22)	1.10	2.05	4.19	0.88	1.85	4.39	1.59	3.16	5.84	<u>1.08</u>	1.74	3.05
SUFT (MM'22)	1.07	1.75	3.18	1.10	1.74	3.92	<u>1.12</u>	2.51	4.86	1.10	1.69	2.71
SSDNet (ICCV'23)	<u>1.02</u>	1.91	4.02	<u>0.80</u>	1.82	4.77	1.60	3.14	5.86	1.04	1.72	2.92
DAGF (TNNLS'23)	1.15	1.80	3.70	0.83	1.93	4.80	1.36	2.87	6.06	1.18	1.82	2.91
DADA (CVPR'23)	1.20	2.03	4.18	0.96	1.87	4.01	1.54	2.74	4.80	1.20	1.83	2.80
SGNet (AAAI'24)	1.15	<u>1.64</u>	<u>2.95</u>	1.03	<u>1.61</u>	<u>3.55</u>	1.10	<u>2.44</u>	<u>4.77</u>	1.10	<u>1.64</u>	<u>2.55</u>
Ours	0.99	1.57	2.80	0.78	1.53	3.11	1.25	2.36	4.48	1.09	1.63	2.41

Table 1: Quantitative comparison (in average RMSE) with existing state-of-the-art methods on four benchmark datasets. The best and the second-best values are highlighted by bold and underline, respectively.

$I_h(x) = f'_x(x, y_p)$, $G_h(x) = g'_x(x, y_p)$, $I_v(x) = f'_y(x_p, y)$,
 $G_v(y) = g'_y(x_p, y)$. Thus we getting:

$$\begin{aligned}
& I_{out}|_{(x_p, y_p)} \\
&= \frac{1}{2} \sum_{i=0}^{x_p} [I_h(x_i) \circ G_h(x_i)] + \frac{1}{2} \sum_{i=0}^{y_p} [I_v(y_i) \circ G_v(y_i)] \\
&= \frac{1}{2} \sum_{i=0}^{x_p} [f(x_{i+1}, y_p) - f(x_i, y_p)] \circ [g(x_{i+1}, y_p) - g(x_i, y_p)] \quad (9) \\
&+ \frac{1}{2} \sum_{i=0}^{y_p} [f(x_p, y_{i+1}) - f(x_p, y_i)] \circ [g(x_p, y_{i+1}) - g(x_p, y_i)].
\end{aligned}$$

The equation (9) represents our mathematical model for deforming the ideal plastic object (depth map). Considering gap between modalities, we use the absolute gradient of guidance information. As shown in Fig. 2, to leverage the reversible discretization of differential operations, we separate the positive and negative gradients, apply the same operation in parallel, finally restore the directional information and add them together to obtain the processed gradient. (Since differentiation is a reversible operation and the information after differentiation is discretized along the differential direction, separately processing the gradients along each direction will not disrupt the continuity after integration.)

As Fig. 3 (c) shows, although each CAPO in PCGD just focus on gradient in one direction, PCGD leverages the context to provide reasonable changes in both horizontal and vertical gradients for each point. The task of CAPO here can be compared to the point-wise multiplication when using an attention map on an image, which can leverage the local information within the attention map. To ensure that improvements stem from our original architecture and operations, we employ U-net (Ronneberger, Fischer, and Brox 2015) in the same manner as DADA (with the exception that our output channels are set to 32, while DADA uses 64) (Metzger, Daudt, and Schindler 2023; He et al. 2016; Deng et al. 2009).

Loss Function

We choose L_1 loss as the training objective, with D_{gt} and D_{out} presenting GT and the predicted depth respectively:

$$\mathcal{L} = \|D_{gt} - D_{out}\|_1. \quad (10)$$

Experiment

Experimental Settings

Datasets and Evaluation Metrics. We conduct experiments on NYU_v2 (Wright et al. 2010), Middlebury (Hirschmuller and Scharstein 2007; Scharstein et al. 2014), Lu (Lu, Ren, and Liu 2014), and RGBDD (He et al. 2021). Consistent with prior studies (Kim, Ponce, and Ham 2021; He et al. 2021; Zhao et al. 2022; Zhong et al. 2023a; Wang, Yan, and Yang 2024), we utilize the first 1000 RGB-D pairs from the NYU-v2 dataset for training, with the remaining 449 pairs reserved for validation. Furthermore, the same pretrained model trained on NYUv2 is evaluated on Middlebury (30 pairs), Lu (6 pairs), and RGBDD (405 pairs) datasets. For the synthetic scenes, low-resolution depth inputs are generated by bicubic(Keys 1981) downsampling of the high-resolution GT. In line with prior approaches (Kim, Ponce, and Ham 2021; Zhong et al. 2023a; Wang, Yan, and Yang 2024), we employ the root mean square error (RMSE) in centimeters as the evaluation metric.

Implementation Details. During the training phase, we randomly crop 256×256 image patches from depths and RGB images as inputs. Following (Zhong et al. 2023a), we augment the training data with random flipping and rotation. Adam optimizer is utilized (Kingma and Ba 2014) with $\beta_1 = 0.9$ and $\beta_2 = 0.999$, employing an initial learning rate of 1×10^{-4} . The model is implemented using PyTorch (Paszke et al. 2017) and trained on one RTX 3090ti GPU. Training typically requires two days for the NYU_v2 dataset.

Comparison with SOTAs

To validate the performance of our method, we compare our method with SOTA methods on $\times 4$, $\times 8$ and $\times 16$ GDSR, including DKN (Kim, Ponce, and Ham 2021), FDKN(Kim, Ponce, and Ham 2021), FDSR (He et al. 2021), JiIF(Tang, Chen, and Zeng 2021), DCTNet (Zhao et al. 2022), SUFT (Shi, Ye, and Du 2022), SSDNet (Zhao et al. 2023), DADA (Metzger, Daudt, and Schindler 2023), DAGF (Zhong et al. 2023a), and SGNet (Wang, Yan, and Yang 2024).

Quantitative Comparison. Tab. 1 demonstrates that C2PD achieves state-of-the-art performance on the Middle-

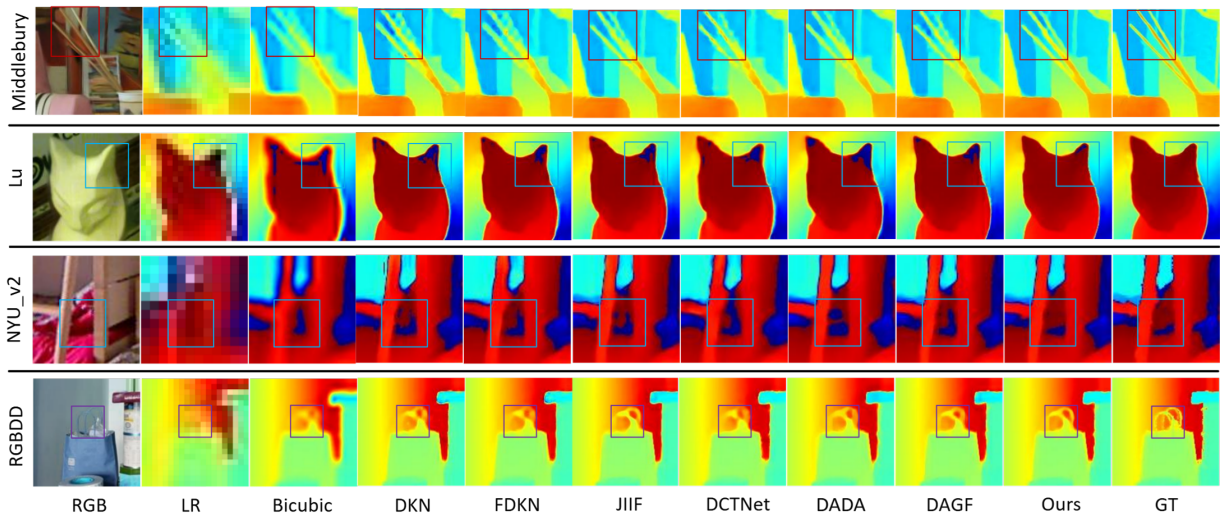


Figure 5: Qualitative comparison of x8 scale factor on four benchmark datasets.

bury, Lu, NYU_v2, and RGBDD datasets. Despite certain results for x4 task are suboptimal (because our method does not use a substantial number of parameters to reconstruct dense effective features), in the x16 task, our method outperforms the second best by 0.15cm (Middlebury), 0.44cm (Lu), 0.29cm (NYU_v2), and 0.14cm (RGBDD) on RMSE (More metrics can be found in supplementary material).

Qualitative Comparison. Fig. 5 displays the visual comparison between our method and others on x8 tasks (Comparisons for more scale factors are provided in supplementary material). It is evident that our method excels not only in edge details but also in preserving the accuracy of internal information. This success is attributed to the capability of PCGD to fully leverage global spatial information.




Figure 6: Visual comparison between DADA and our method on NYU_v2 dataset, with scaling factors of $\times 32$.

Comparison on Large Scale. Tab. 2 and Fig. 6, respectively, present the quantitative and qualitative comparisons on x32 task with DADA, which is demonstrated the best method at the x32 scale. It is evident that our method not only outperforms DADA at the x32 scale but also maintains more detailed and realistic textures without excessive

Methods	Middlebury	Lu	NYU_v2	RGBDD
DADA	6.94	8.47	10.81	4.96
Ours	5.68	6.48	9.03	4.11

Table 2: Quantitative comparison with DADA on x32 scale factor. The best performance is displayed in bold

smoothing observed in DADA.

Complexity Comparison. Tab. 3 shows the complexity comparison between our method with SOTA method. While earlier methods achieved lightweight architectures, recent methods reaching SOTA performance generally incur significantly higher complexity. Our method attains optimal performance with a comparable computational cost.

Methods	Params (M)	FLOPs (G)	Memory (G)	Time (ms)	RMSE(cm)
DAGF	2.28	4623.9	10.77	73	2.87
DADA	31.03	212.58	4.96	975	2.74
SGNet	39.25	4623.9	10.77	73	2.44
Ours	62.05	414.75	8.97	61	2.36

Table 3: Complexity comparison on NYU ($\times 8$)

Ablation Study

Ablation Study for CAPO. To validate the necessity of CAPO, we fairly replace it with residual convolution (with same parameters amount). As shown in Fig. 7 (a), convolution is far from achieving our effect, highlighting the critical role of CAPO (The difference between CAPO and convolution is discussed in supplementary material).

Fig. 7 (b) presents the ablation study on parameter of CAPO, which operates on gradients. It is evident that our method achieves optimal performance when the length is set to 4, with stable performance on remaining datasets. (For lengths of 3, 5, and 6, matrix multiplication is replaced by dot products due to computational constraints, potentially contributing to the superior performance at length = 4.)




Figure 7: Ablation study for CAPO. (a) Replace CAPO with residual block. (b) Parameter of CAPO in PCGD.

Ablation Study for PCGD. As shown in Fig. 8 (a), we test four possible structure for PCGD: operating on vertical gradient, operating on horizontal gradient, operating on both direction without sharing parameter, both direction operated in parallel with parameter sharing. When PCGD sharing parameter in parallel, the performance comes to the best.




Figure 8: (a) Different structures of PCGD on x8 scale. (b) Comparison at different upsampling factors on Middlebury.

Further Analysis

Analysis for Upsampling Factor. As depicted in Fig. 8 (b), our model exhibits a linear degradation trend, indicating that its advantages become more pronounced as the scale increases, confirming that our method can sustain superior performance even with low density of effective information.




Figure 9: Generalization performance of our x32 model.

Generalization Analysis. Since all methods were trained on the NYU_v2, the performance on Middlebury, Lu, and RGBDD reflects the generalization ability. Tab. 1 shows our

model significantly outperforms the second best on these datasets, indicating strong generalization.

Methods	DKN	FDSR	DCTNet	SUFT	SGNet	Ours
RMSE	7.38	7.50	7.37	7.22	7.22	6.68

Table 4: Quantitative comparison on real-world RGBDD.

The robust generalization of our model enhances its adaptability to the intricate scenarios in the real world. Fig. 9 reveals that our trained model (NYU_v2, x32) exhibits excellent generalization across various datasets, and Tab. 4 presents a quantitative comparison of it with state-of-the-art methods on real-world RGB-D data, demonstrating its significant practical value for real-world applications.




Figure 10: Robustness to RGB textures.

Robustness to RGB. Fig. 10 shows the robustness to complex RGB images, which suggest that our approach is not misled by mismatching textures in the guidance. This is because: (1) Freed from the constraint of “clear and accurate”, we employ a deeper network to sufficiently highlight and abstract guidance information, thereby weakening and even eliminating complex textures. (2) As illustrated in Fig. 3 (b), the principle of CAPO is not about merging values but about fitting the deformation pattern of $I'[n]$ using the sequences from $G'[n]$ and $I[n]$. Even if there are slight traces left by irrelevant structural details, “op1” in CAPO can leverage the relationship between the guidance and the depth to weaken erroneous guidance. Additionally, constrained by “op2”, the deformation of $I'[n]$ is still affected by the values of $I[n]$ itself, meaning that trivial textures are unlikely to introduce erroneous textures on the smooth depth map.

Conclusion

In this paper, we propose a novel GDSR method PCGD that integrates human understanding of material properties. By leveraging depth as spatial information, the CAPO operation is proposed, which facilitates continuous constraint deformation on depth, treating it as a flexible object with constant volume. Building upon CAPO, the PCGD is designed, which enable transformations that surpass volume constraints, resulting in a system capable of deforming depth as a substance with ideal plasticity. Our method achieves state-of-the-art performance on four widely used datasets, especially advantageous in large-scale tasks and generalization. This presents a new strategy for further breakthrough in large-scale tasks and real-world application.

Acknowledgements

This work was supported in part by the National Science Foundation of China under Grant62471448, 62102338; in part by Shandong Provincial Natural Science Foundation under Grant ZR2024YQ004; in part by TaiShan Scholars Youth Expert Program of Shandong Province under Grant No.tsqn202312109.

References

- Cai, Q.; Li, J.; Li, H.; Yang, Y.-H.; Wu, F.; and Zhang, D. 2022. TDPN: Texture and detail-preserving network for single image super-resolution. *IEEE Transactions on Image Processing*, 31: 2375–2389.
- Cai, Q.; Li, M.; Ren, D.; Lyu, J.; Zheng, H.; Dong, J.; and Yang, Y.-H. 2024. Spherical pseudo-cylindrical representation for omnidirectional image super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 873–881.
- Cai, Q.; Qian, Y.; Li, J.; Lyu, J.; Yang, Y.-H.; Wu, F.; and Zhang, D. 2023. HIPA: hierarchical patch transformer for single image super resolution. *IEEE Transactions on Image Processing*, 32: 3226–3237.
- Chen, X.; Lin, K.-Y.; Qian, C.; Zeng, G.; and Li, H. 2020a. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4193–4202.
- Chen, X.; Lin, K.-Y.; Wang, J.; Wu, W.; Qian, C.; Li, H.; and Zeng, G. 2020b. Bi-directional cross-modality feature propagation with separation-and-aggregation gate for RGB-D semantic segmentation. In *European Conference on Computer Vision*, 561–577. Springer.
- Chen, Y.; Fan, H.; Xu, B.; Yan, Z.; Kalantidis, Y.; Rohrbach, M.; Yan, S.; and Feng, J. 2019. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3435–3444.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. Ieee.
- Deng, X.; and Dragotti, P. L. 2019. Deep coupled ISTA network for multi-modal image super-resolution. *IEEE Transactions on Image Processing*, 29: 1683–1698.
- Deng, X.; and Dragotti, P. L. 2020. Deep convolutional neural network for multi-modal image restoration and fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10): 3333–3348.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2015. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2): 295–307.
- Gregor, K.; and LeCun, Y. 2010. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 399–406.
- Gu, S.; Zuo, W.; Guo, S.; Chen, Y.; Chen, C.; and Zhang, L. 2017. Learning dynamic guidance for depth image enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3769–3778.
- Guo, C.; Li, C.; Guo, J.; Cong, R.; Fu, H.; and Han, P. 2018. Hierarchical features driven residual learning for depth map super-resolution. *IEEE Transactions on Image Processing*, 28(5): 2545–2557.
- Gupta, S.; Girshick, R.; Arbeláez, P.; and Malik, J. 2014. Learning rich features from RGB-D images for object detection and segmentation. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VII 13*, 345–360. Springer.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- He, L.; Zhu, H.; Li, F.; Bai, H.; Cong, R.; Zhang, C.; Lin, C.; Liu, M.; and Zhao, Y. 2021. Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9229–9238.
- Hirschmuller, H.; and Scharstein, D. 2007. Evaluation of cost functions for stereo matching. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8. IEEE.
- Hoffman, J.; Gupta, S.; and Darrell, T. 2016. Learning with side information through modality hallucination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 826–834.
- Hui, T.-W.; Loy, C. C.; and Tang, X. 2016a. Depth map super-resolution by deep multi-scale guidance. In *Computer Vision—ECCV 2016: 14th European Conference, 11–14, 2016, Proceedings, Part III 14*, 353–369. Springer.
- Hui, T.-W.; Loy, C. C.; and Tang, X. 2016b. Depth map super-resolution by deep multi-scale guidance. In *Computer Vision—ECCV 2016: 14th European Conference, 11–14, 2016, Proceedings, Part III 14*, 353–369. Springer.
- Im, S.; Ha, H.; Choe, G.; Jeon, H.-G.; Joo, K.; and Kweon, I. S. 2018. Accurate 3d reconstruction from small motion clip for rolling shutter cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4): 775–787.
- Kerl, C.; Sturm, J.; and Cremers, D. 2013. Dense visual SLAM for RGB-D cameras. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2100–2106. IEEE.
- Keys, R. 1981. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6): 1153–1160.
- Kim, B.; Ponce, J.; and Ham, B. 2021. Deformable kernel networks for joint image filtering. *International Journal of Computer Vision*, 129(2): 579–600.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *ArXiv preprint arXiv:1412.6980*.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.

- Li, Y.; Huang, J.-B.; Ahuja, N.; and Yang, M.-H. 2016. Deep joint image filtering. In *Computer Vision—ECCV 2016: 14th European Conference, 11–14, 2016, Proceedings, Part IV 14*, 154–169. Springer.
- Li, Y.; Huang, J.-B.; Ahuja, N.; and Yang, M.-H. 2019. Joint image filtering with deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 1909–1923.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1833–1844.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 136–144.
- Lu, S.; Ren, X.; and Liu, F. 2014. Depth enhancement via low-rank matrix completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3390–3397.
- Lyu, J.; Li, G.; Wang, C.; Cai, Q.; Dou, Q.; Zhang, D.; and Qin, J. 2023. Multicontrast mri super-resolution via transformer-empowered multiscale contextual matching and aggregation. *IEEE Transactions on Neural Networks and Learning Systems*.
- Metzger, N.; Daudt, R. C.; and Schindler, K. 2023. Guided depth super-resolution by deep anisotropic diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18237–18246.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch.
- Riegler, G.; Ferstl, D.; R  ther, M.; and Bischof, H. 2016. A deep primal-dual network for guided depth super-resolution. *ArXiv preprint arXiv:1607.08569*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-assisted Intervention—MICCAI 2015: 18th International Conference, 5–9, 2015, Proceedings, part III 18*, 234–241. Springer.
- Scharstein, D.; Hirschm  ller, H.; Kitajima, Y.; Krathwohl, G.; Ne  i  , N.; Wang, X.; and Westling, P. 2014. High-resolution stereo datasets with subpixel-accurate ground truth. In *Pattern Recognition: 36th German Conference, GCPR 2014, M  nster, Germany, September 2–5, 2014, Proceedings 36*, 31–42. Springer.
- Shi, W.; Ye, M.; and Du, B. 2022. Symmetric uncertainty-aware feature transmission for depth super-resolution. In *Proceedings of the 30th ACM International Conference on Multimedia*, 3867–3876.
- Su, H.; Jampani, V.; Sun, D.; Gallo, O.; Learned-Miller, E.; and Kautz, J. 2019. Pixel-adaptive convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11166–11175.
- Sun, W.; Ren, K.; Meng, X.; Yang, G.; Xiao, C.; Peng, J.; and Huang, J. 2022. MLR-DBPFN: A multi-scale low rank deep back projection fusion network for anti-noise hyperspectral and multispectral image fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–14.
- Tang, J.; Chen, X.; and Zeng, G. 2021. Joint implicit image function for guided depth super-resolution. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4390–4399.
- Tang, Q.; Cong, R.; Sheng, R.; He, L.; Zhang, D.; Zhao, Y.; and Kwong, S. 2021. Bridgenet: A joint learning network of depth map super-resolution and monocular depth estimation. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2148–2157.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser,  .; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Wang, T.; Xie, J.; Sun, W.; Yan, Q.; and Chen, Q. 2021. Dual-camera super-resolution with aligned attention modules. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2001–2010.
- Wang, Z.; Yan, Z.; and Yang, J. 2024. Sgnet: Structure guided network via gradient-frequency awareness for depth map super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5823–5831.
- Weder, S.; Schonberger, J.; Pollefeys, M.; and Oswald, M. R. 2020. Routedfusion: Learning real-time depth map fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4887–4897.
- Wright, J.; Ma, Y.; Mairal, J.; Sapiro, G.; Huang, T. S.; and Yan, S. 2010. Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE*, 98(6): 1031–1044.
- Ye, X.; Sun, B.; Wang, Z.; Yang, J.; Xu, R.; Li, H.; and Li, B. 2020. PMBANet: Progressive multi-branch aggregation network for scene depth super-resolution. *IEEE Transactions on Image Processing*, 29: 7427–7442.
- Zhao, Z.; Zhang, J.; Gu, X.; Tan, C.; Xu, S.; Zhang, Y.; Timofte, R.; and Van Gool, L. 2023. Spherical space feature decomposition for guided depth map super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12547–12558.
- Zhao, Z.; Zhang, J.; Xu, S.; Lin, Z.; and Pfister, H. 2022. Discrete cosine transform network for guided depth map super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5697–5707.
- Zhong, Z.; Liu, X.; Jiang, J.; Zhao, D.; and Ji, X. 2023a. Deep attentional guided image filtering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhong, Z.; Liu, X.; Jiang, J.; Zhao, D.; and Ji, X. 2023b. Guided depth map super-resolution: A survey. *ACM Computing Surveys*, 55(14s): 1–36.
- Zhou, M.; Huang, J.; Zheng, N.; and Li, C. 2023. Learned image reasoning prior penetrates deep unfolding network for panchromatic and multi-spectral image fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12398–12407.