

Robust Self-Supervised Cross-Modal Super-Resolution against Real-World Misaligned Observations

Xiaoyu Dong¹, Jiahuan Li^{1,2}, Ziteng Cui¹, and Naoto Yokoya^{1,2}✉

¹ The University of Tokyo, Japan

² RIKEN AIP, Japan

Abstract. Cross-modal super-resolution (SR) on real-world misaligned data is challenging, as only unlabeled low-resolution (LR) source and high-resolution (HR) guide images with complex spatial misalignment are available. Previous methods either rely on simulated training data or adopt suboptimal alignment strategies that overlook cross-modal dependencies, limiting their practical performance. To address these issues, we propose RobSelf, a self-supervised model that jointly optimizes a misalignment-aware feature translator and a content-aware reference filter online. The translator resolves unsupervised cross-modal and cross-resolution alignment via weakly-supervised, misalignment-aware translation, yielding an aligned guide feature. Guided by this feature, the filter performs reference-based discriminative self-enhancement on the source, enabling SR prediction with high resolution and high fidelity. Experiments on synthesized data and collected real-world data demonstrate that RobSelf achieves state-of-the-art performance, outperforming existing self-supervised and supervised methods. Moreover, it achieves superior efficiency, being up to $15.3\times$ faster than prior self-supervised methods. <https://github.com/palmdong/RobSelf>

Keywords: Self-supervised super-resolution · Multi-modal vision




Fig. 1: Real-world misaligned RGB-guided depth SR ($\times 4$). Our model achieves state-of-the-art performance, *requiring no training data, ground-truth supervision, or pre-alignment*. (a) LR source; (b) HR guide; (c) pre-aligned guide by MINIMA [36]; (d) SS-GNet [40] + pre-alignment; (e) SGNet [49] + pre-alignment; (f) RobSelf-Re (Ours).

1 Introduction

Multi-modal images, *e.g.*, RGB, depth, and near-infrared (NIR), capture complementary properties of objects and environments [11, 57, 66, 67]. However, non-visible modalities generally suffer from lower spatial resolution than RGB due to sensing and hardware limits, which hinders downstream tasks and necessitates cross-modal super-resolution (SR) [56, 71, 72].

Cross-modal SR enhances a low-resolution (LR) source image using structural cues from a high-resolution (HR) guide image of another modality. Modern methods fall into supervised and self-supervised categories. Supervised methods rely on large-scale domain-specific training data and ground truth [19, 47, 49, 50, 71, 72]. This reliance hinders their practical generalization, as building such datasets is costly and labor-intensive. Self-supervised methods require neither training data nor ground truth, but instead optimize online on each test pair, making them more data-efficient and generalizable [12, 30, 38, 40].

While promising progress has been achieved, most existing methods assume that the source and guide images are well-aligned [12, 19, 30, 40, 47, 49, 50, 71, 72]. In real-world scenarios, however, spatial misalignment is inevitable in multi-modal images due to inherent cross-sensor discrepancies [28, 51] (*e.g.*, differences in lens distortion, field of view, and physical position) and environmental factors, such as platform-induced viewpoint variation [5, 21] and object motion over time [27, 46]. Although several methods consider misalignment, they either rely on simulated training data [48] or adopt suboptimal alignment strategies [15, 38], limiting their practical performance.

We identify two difficulties in cross-modal SR on *real-world* misaligned data: (I) The scarcity of training data and the absence of SR ground truth for the source [12, 30]; (II) Complex misalignments across modalities and resolutions, along with the lack of alignment ground truth for the guide [7, 61]. These hinder the development of reliable supervised methods. One solution is to perform pre-alignment before SR. However, such a two-stage pipeline may not generalize effectively to real-world data with complex misalignment and resolution gap (Fig. 1). Overall, robust self-supervised cross-modal SR for real-world misaligned data is crucial yet remains an open challenge.

To address these issues, we propose RobSelf (Fig. 2), a self-supervised model that resolves unsupervised cross-modal and cross-resolution alignment via a formulation of joint weakly-supervised, misalignment-aware translation. RobSelf features a misalignment-aware feature translator (Fig. 3) and a content-aware reference filter (Fig. 4). The translator warps and translates the guide feature to mimic the source modality, while yielding an aligned guide feature with inherent redundancy. Guided by this feature, the filter learns content-aware kernels for reference-based discriminative self-enhancement on the source, enabling effective enhancement while mitigating redundancy effects. To evaluate RobSelf, we collected real-world RGB-depth and RGB-NIR data with inherent cross-sensor misalignment, random viewpoint variation, and random object motion. Experiments demonstrate that RobSelf achieves HR and high-fidelity predictions on complex real-world misaligned data, even without training data, ground-truth




Fig. 2: RobSelf is solely supervised by the LR source. Within this framework, the translator is optimized under weak supervision to map F_{guide} into I_{pred}^{Trans} that mimics the source, while yielding $F_{guide}^{Aligned}$. Guided by $F_{guide}^{Aligned}$, the filter performs reference-based discriminative self-enhancement on F_{source} to obtain $F_{source}^{Enhanced}$, from which I_{pred}^{SR} is generated. RobSelf has two variants depending on the the translator design.

supervision, or pre-alignment (Fig. 1). Moreover, it achieves superior efficiency, being up to $15.3\times$ faster than prior self-supervised methods [12, 30, 40].

Contributions: (I) We address the open challenge of robust self-supervised cross-modal SR on real-world misaligned data by proposing RobSelf. (II) We propose a joint weakly-supervised, misalignment-aware translation formulation that effectively handles diverse misalignment conditions and even boundary scenarios with missing guide structures. (III) Our reference-based discriminative self-enhancement strategy enables faithful source enhancement without additional guide processing. (IV) Experiments on synthesized depth SR, real-world depth SR, and real-world NIR SR demonstrate state-of-the-art performance, robustness, and generalizability.

2 Related Work

2.1 Cross-Modal SR

Modern learning-based cross-modal SR methods fall into supervised and self-supervised categories. Supervised methods rely on large-scale domain-specific training data and ground truth [1, 16, 19, 47, 50, 54, 55, 68, 69, 71, 72]. Self-supervised methods [12, 30, 38, 40, 62] are more data-efficient and generalizable, requiring no training data or ground-truth supervision. However, most methods overlook misalignment issues, assuming spatially well-aligned inputs. For example, SFG [58] addresses structural distortion and edge noise in depth SR under aligned RGB guidance. To address misalignment, MOMNet [48], a concurrent supervised method, develops multi-order feature matching modules using simulated training data, which do not fully capture real-world misalignments. CMSR [38]

addresses misalignment via a self-supervised framework built on pseudo-paired data. However, it relies on off-the-shelf RGB deformation layers and overlooks cross-modal dependencies, resulting in suboptimal performance.

In contrast: (I) We address *complex real-world misalignments* arising from sensor discrepancies and environmental factors. (II) Instead of relying on modules designed solely for image or feature alignment [15, 38, 48], we propose a misalignment-aware *feature translator* that, driven by its weakly-supervised translation objective, effectively handles *diverse misalignments and missing guide structures* within a self-supervised SR framework.

Misalignment also arises in guided hyperspectral image (HSI) SR [10, 25, 63], pan-sharpening [8], fusion [26, 35, 52, 65], MRI reconstruction [53], and depth refinement [34]. Unlike these methods that rely on training data and focus on image- or feature-level registration, our online self-supervised model with a joint translation objective achieves better robustness (Sec. 5.4).

2.2 Multi-Modal Image Registration

Multi-modal image registration aims to align images of different modalities into a common coordinate space. Modern data-driven methods estimate transformation parameters or deformation fields in supervised [4, 7, 24, 33, 36, 59, 64] and unsupervised [2, 14, 39, 43, 61] paradigms. For example, Ren *et al.* [36] proposed a supervised model, MINIMA, for unified cross-modal matching. Deng *et al.* [7] developed a supervised model for cross-modal, cross-resolution homography estimation. Recently, Song *et al.* [43] developed an unsupervised registration framework via alternating optimization. However, data-driven methods rely on training data and may not generalize well to complex in-the-wild scenarios.

For our task, training data are scarce and neither SR nor alignment ground truth is available. To address this challenge, we cast unsupervised cross-modal and cross-resolution alignment as a weakly-supervised, misalignment-aware translation subtask within RobSelf for *joint online optimization*. This enables robustness and generalizability across diverse data and misalignments.

2.3 Learnable Image Filtering

Learnable image filtering learns filter rules or weights from data, enabling adaptive content modeling, suppression, and extraction [9, 70]. In high-level vision, learnable filters such as tree [41, 42] and non-local [44, 45] filters have proven effective for modeling dependencies and preserving structures. In low-level vision, non-local filters are effective in exploiting self-similarity priors to preserve structures [13], remove noise [29], and recover details [31]. In cross-modal SR, degradation-aware filters [47], cross-domain adaptive filters [12], and deformable filters [20] have been introduced to facilitate multi-modal feature processing and fusion. These filters are, however, designed for well-aligned data and may fail to transfer guidance information reliably when misalignments occur.

We further observe that the guide, even after alignment (translation), contains both structures essential for source enhancement and redundant content

due to modality discrepancies. Therefore, *rather than filtering or fusing the guide*, we design a filter that learns content-aware kernels to perform reference-based discriminative self-enhancement on the source. This enables HR and faithful SR predictions, free from the effects of redundant guide content.

3 RobSelf

3.1 Overview

As illustrated in Fig. 2, given an input image pair, the LR source $\mathbf{I}_{source}^{LR} \in \mathbb{R}^{h \times w \times \phi}$ is first bilinearly upsampled to match the spatial size of the HR guide $\mathbf{I}_{guide}^{HR} \in \mathbb{R}^{H \times W \times \psi}$. Two extraction layers then produce the source and guide features, $\mathbf{F}_{source}, \mathbf{F}_{guide} \in \mathbb{R}^{C \times H \times W}$. ϕ, ψ , and C denote channel dimensions. The translator maps $\mathbf{F}_{guide}$ to an HR prediction $\mathbf{I}_{pred}^{Trans}$, while yielding the aligned guide feature $\mathbf{F}_{guide}^{Aligned}$. Guided by $\mathbf{F}_{guide}^{Aligned}$, the filter enhances $\mathbf{F}_{source}$ to obtain $\mathbf{F}_{source}^{Enhanced}$, from which the SR prediction $\mathbf{I}_{pred}^{SR}$ is generated by a prediction head. Both $\mathbf{I}_{pred}^{SR}$ and $\mathbf{I}_{pred}^{Trans}$ are supervised by regression to the LR source via a consistency loss [12]:

$$\mathcal{L} = \mathcal{L}_{sr} + \lambda \mathcal{L}_{trans} = \|f_{down}(\mathbf{I}_{pred}^{SR}) - \mathbf{I}_{source}^{LR}\|_1 + \lambda \|f_{down}(\mathbf{I}_{pred}^{Trans}) - \mathbf{I}_{source}^{LR}\|_1, \quad (1)$$

where f_{down} denotes average pooling and λ is a weighting factor. Network details are provided in the supplementary material.

3.2 Misalignment-Aware Feature Translator

As shown in (Fig. 3), driven by the translation objective mimicking the source modality, the translator yields $\mathbf{F}_{guide}^{Aligned}$, which is aligned with the source and contains both essential structures and redundant content.

Misalignment Modeling. The translator first derives a cross-modal deformation field $\mathbf{D}^{G2S}$ that models the misalignment from $\mathbf{F}_{guide}$ to $\mathbf{F}_{source}$ using a multi-level (i -level) estimator. With progressive downsampling and upsampling, this estimator captures misalignments across multiple spatial scales. When $i = 4$,




Fig. 3: Misalignment-aware feature translator. Each encoder layer downsamples by $\times 2$; each decoder layer upsamples by $\times 2$.

for example, the input features are downsampled by 16 after the encoder, enabling misalignment modeling at the regional level. $\mathbf{D}^{\text{G2S}}$ is an $H \times W$ matrix of N -dimensional offset vectors.

Feature Alignment. The alignment layer warps $\mathbf{F}_{\text{guide}}$ according to $\mathbf{D}^{\text{G2S}}$, yielding $\mathbf{F}_{\text{guide}}^{\text{Aligned}}$. It is implemented by a deformable convolution (conv) [6] or a simple spatial resampling operation, resulting in two model variants: RobSelf-De and RobSelf-Re. Formally, at location p :

$$\mathbf{F}_{\text{guide}}^{\text{Aligned}}(p) = \begin{cases} \sum_{i=1}^{k^2} w_i \mathbf{F}_{\text{guide}}(p + p_i + \Delta p_i), & \text{if RobSelf-De,} \\ \mathbf{F}_{\text{guide}}(p + \Delta p), & \text{if RobSelf-Re,} \end{cases} \quad (2)$$

where p_i denotes the i -th relative location in the $k \times k$ conv kernel centered at p , w_i is the learned kernel weight, and the sampling offsets $\Delta p_i = (\Delta x_i, \Delta y_i)$ and $\Delta p = (\Delta x, \Delta y)$ are learned from $\mathbf{D}^{\text{G2S}}$. Accordingly, the dimension N of $\mathbf{D}^{\text{G2S}}$ is $2k^2$ for RobSelf-De and 2 for RobSelf-Re.

Translation Objective. Through a prediction head, the translator finally generates an HR prediction $\mathbf{I}_{\text{pred}}^{\text{Trans}}$, weakly supervised by $\mathbf{I}_{\text{source}}^{\text{LR}}$ to drive effective feature alignment. Such a joint weakly-supervised, misalignment-aware translation formulation enables robustness and generalizability across diverse data and misalignments, as well as the ability to handle missing guide structures even within a self-supervised SR framework (Sec. 5.6).

3.3 Content-Aware Reference Filter

To better exploit essential guidance and mitigate redundancy effects, we design a filter (Fig. 4) that learns content-aware kernels to enhance $\mathbf{F}_{\text{source}}$ using its own pixels, with $\mathbf{F}_{\text{guide}}^{\text{Aligned}}$ only as a reference for weight determination—essentially reference-based discriminative self-enhancement.

Content-Aware Design. We first derive an importance map $\mathbf{M}_{\text{imp}} \in \mathbb{R}^{H \times W}$:

$$\mathbf{M}_{\text{imp}} = \frac{1}{C} \sum_{c=1}^C \|\nabla(\mathbf{F}_{\text{source}}^c)\|_2, \quad (3)$$

where ∇ denotes spatial gradient operator, c is channel index, and C is number of channels. The motivation is: high-gradient source content (*e.g.*, edges and




Fig. 4: Content-aware reference filter.

textures) is important, corresponding to essential guide structures; low-gradient source content (*e.g.*, smooth regions) is less important, corresponding to redundant guide content. Based on content importance, we adopt large and small kernels to discriminatively enhance the source pixels.

Reference-based Filtering. A threshold is computed as:

$$\tau = \eta \mathbb{E}[\mathbf{M}_{imp}], \quad (4)$$

where η is a scaling factor. For important pixels $\mathbf{F}_{source}(i)$ with $\mathbf{M}_{imp}(i) > \tau$, a large kernel $\mathbf{w}_i^{imp} \in \mathbb{R}^{m \times m}$ aggregates more neighbors for stronger self-enhancement under essential guidance. For less important pixels $\mathbf{F}_{source}(j)$ with $\mathbf{M}_{imp}(j) < \tau$, a small kernel $\mathbf{w}_j^{l-imp} \in \mathbb{R}^{n \times n}$ performs lightweight self-update under redundant guidance. We learn correlation-based kernel weights [12], proven effective in inducing pixels to emulate the properties of cross-domain references. The weights of $\mathbf{w}_i^{imp}$, for example, are determined by the correlation between the source pixels within an $m \times m$ neighborhood $\mathcal{N}_i$ centered at $\mathbf{F}_{source}(i)$ and the reference pixel $\mathbf{F}_{guide}^{Aligned}(i)$ with inherent HR properties:

$$w_{i,l}^{imp} = \frac{\exp(\mathbf{F}_{source}(l)^\top \mathbf{F}_{guide}^{Aligned}(i))}{\sum_{l' \in \mathcal{N}_i} \exp(\mathbf{F}_{source}(l')^\top \mathbf{F}_{guide}^{Aligned}(i))}, \quad (5)$$

where $w_{i,l}^{imp}$ is the weight for a specific neighbor $\mathbf{F}_{source}(l)$ of $\mathbf{F}_{source}(i)$, and l' indexes pixels in $\mathcal{N}_i$ for normalization. The self-enhanced pixel is obtained as:


$$\mathbf{F}_{source}^{Enhanced}(i) = \sum_{l \in \mathcal{N}_i} w_{i,l}^{imp} \mathbf{F}_{source}(l). \quad (6)$$

Such reference-based discriminative self-enhancement requires no additional guide processing and enables $\mathbf{F}_{source}^{Enhanced}$ and $\mathbf{I}_{pred}^{SR}$ with high resolution and high fidelity, free from the effects of redundant content in $\mathbf{F}_{guide}^{Aligned}$.

4 Real-World Misaligned Data

To evaluate RobSelf in real-world scenarios, we collected data with two modality combinations, three types of misalignments, real LR data, and diverse scenes and objects (Fig. 5). Following [15, 25], the data were captured using a single multi-modal device (Azure Kinect DK in our case). To avoid bias toward fixed cross-sensor misalignment, we introduced random viewpoint variations and object motions during data acquisition. Collection procedures and data comparisons are provided in the supplementary material.

The RGB-Depth subset provides 52 groups of simple cases with inherent cross-sensor misalignment (caused by factors such as sensors' lens distortion, fields of view, and physical positions), and 60 groups of complex cases with inherent cross-sensor misalignment and random viewpoint variation. Each group



(a) RGB-Depth. Left: Inherent cross-sensor misalignment. Right: Inherent cross-sensor misalignment and random viewpoint variation. **(b) RGB-NIR.** Left: Inherent cross-sensor misalignment. Right: Inherent cross-sensor misalignment and random object motion.

Fig. 5: Examples from our collected real-world misaligned data. LR sources are overlaid on HR guides for better visualization of misalignments.

contains a raw LR depth, a filled LR depth, a $\times 2$ HR depth, a $\times 2$ HR RGB, and a $\times 4$ HR RGB. The data cover small-scale toy scenes and large-scale indoor scenes, with objects such as toys, desks, chairs, and other common indoor items. The RGB-NIR subset provides 50 groups of simple cases with inherent cross-sensor misalignment, and 30 groups of complex cases with inherent cross-sensor misalignment and random object motion. Each group contains an LR NIR, a $\times 2$ HR NIR, a $\times 2$ HR RGB, and a $\times 4$ HR RGB. The data span small-scale toy scenes and medium-scale human-centric scenes, with objects such as real plants, artificial plants, toys, and humans. For both subsets, there is no ground truth for $\times 4$ SR. The resolutions are: LR— 288×320 , $\times 2$ HR— 576×640 , and $\times 4$ HR— 1152×1280 .

5 Experiments

5.1 Settings

Tasks. We conduct three tasks: (I) Synthesized misaligned RGB-guided depth SR ($\times 4$ and $\times 8$); (II) Real-world misaligned RGB-guided depth SR ($\times 2$ and $\times 4$); (III) Real-world misaligned RGB-guided NIR SR ($\times 2$ and $\times 4$). For task I, we use the Middlebury dataset [17, 37], with misalignments synthesized via random translation, rotation, and perspective shifts, and LR data generated following [32]. For tasks II and III, we use the complex cases from our collected data: RGB-depth with inherent cross-sensor misalignment and random viewpoint variation, and RGB-NIR with inherent cross-sensor misalignment and random object motion. Pre-alignment was performed for methods lacking alignment strategies using the pre-trained MINIMA model [36].

Implementation. Our model is implemented in PyTorch on an NVIDIA A100 GPU. To ensure fairness, optimization is conducted per image pair for 1000 iterations following [12], without data augmentation or pre-training. We adopt the Adam optimizer [22] with an initial learning rate of 0.001, decayed by 0.9998 every 5 iterations. The loss weight λ is set to 1. The estimator level i , filter threshold scaling factor η , and kernel sizes $\{m, n\}$ slightly vary with

Table 1: Synthesized misaligned RGB-guided depth SR. For methods lacking alignment strategies, results without and with pre-alignment are reported (wo/w). *Params vary with SR factors ($\times 4/\times 8$) due to level i . The **best** and second best are highlighted.

	P2P [30]	CMSR [38]	DCTNet [69]	MMSR [12]	SSGNet [40]	SGNet [49]	C2PD [19]	DORNet [47]	RobSelf-Re	RobSelf-De
Self-Supervised	✓	✓	✗	✓	✓	✗	✗	✗	✓	✓
Pre-Alignment	✓	✗	✓	✓	✓	✓	✓	✓	✗	✗
Params (M)	0.2	1.18	0.48	0.25	0.31	36.39	65.06	3.05	0.76/0.87*	0.81/0.92*
$\times 4$ (RMSE $\downarrow$)	2.81/2.33	1.91	2.19/2.16	2.22/1.88	2.20/1.92	2.44/2.33	2.12/2.06	2.57/2.46	<u>1.52</u>	1.43
$\times 8$ (RMSE $\downarrow$)	3.33/2.91	3.03	3.34/3.33	2.99/2.79	2.99/2.83	3.48/3.34	3.27/3.15	3.58/3.54	<u>2.57</u>	2.49




Fig. 6: Synthesized misaligned RGB-guided depth SR ($\times 8$). Image patches are shown due to space limit. *Pre-alignment applied.

tasks. Analyses are provided in the supplementary material. RMSE, DSS (depth-specific, no-reference) [18], or NIQE (no-reference) [60] is used for evaluation, depending on the availability of ground truth.

5.2 Synthesized Misaligned RGB-Guided Depth SR

The estimator level i and filter kernel sizes $\{m, n\}$ are set to 4 and $\{7, 5\}$ for $\times 4$ SR and 5 and $\{13, 7\}$ for $\times 8$ SR, with the filter threshold scaling factor $\eta = 0.7$. We compare with four supervised methods (DORNet [47], C2PD [19], SGNet [49], and DCTNet [69]) and four self-supervised methods (SSGNet [40], MMSR [12], CMSR [38], and P2P [30]). The supervised methods are trained on the NYU v2 dataset [23]. Code of [48] is unavailable.

Table 1 reports quantitative results. Our models achieve the best results for both SR factors, despite requiring no training data, ground-truth supervision, or pre-alignment. Figure 6 shows qualitative results. CMSR [38], limited by its suboptimal alignment strategy that overlooks cross-modal dependencies, fails to effectively enhance the source. Two-stage methods (DORNet [47], C2PD [19], SSGNet [40]) perform better due to pre-alignment, but produce some non-faithful structures because their modules are not robust to residual misalignment. In comparison, our models deliver HR and high-fidelity results, approaching the ground truth. These improvements are attributed to our translator and filter, which jointly ensure effective guide alignment and faithful source enhancement. Section 5.5 further analyzes the effect of our techniques.

Table 2: Real-World misaligned RGB-guided depth SR. For methods lacking alignment strategies, results without and with pre-alignment are reported (wo/w). *Params vary with SR factors ($\times 2/\times 4$) due to level i . The **best** and second best are highlighted.

	P2P [30]	CMSR [38]	DCTNet [69]	MMSR [12]	SSGNet [40]	SGNet [49]	DCNAS [72]	DORNet [47]	RobSelf-Re	RobSelf-De
Self-Supervised	✓	✓	✗	✓	✓	✗	✗	✗	✓	✓
Pre-Alignment	✓	✗	✓	✓	✓	✓	✓	✓	✗	✗
Params (M)	0.2	1.18	0.48	0.25	0.31	9.22	0.60	3.05	0.65/0.76*	0.70/0.81*
$\times 2$ (RMSE↓)	3.75/3.51	2.53	3.27/3.14	2.93/2.76	2.93/2.74	3.18/2.98	3.49/3.27	3.24/3.04	<u>2.23</u>	2.18
$\times 4$ (DSS↑)	1.7/2.2	6.7	4.1/4.5	5.0/5.6	5.0/5.5	4.7/5.1	3.6/4.1	4.3/4.9	<u>8.4</u>	8.6




Fig. 7: Real-World misaligned RGB-guided depth SR (upper: $\times 2$; lower: $\times 4$). Image patches are shown due to space limit. *Pre-alignment applied.

5.3 Real-World Misaligned RGB-Guided Depth SR

The estimator level i is set to 3 for $\times 2$ SR and 4 for $\times 4$ SR. The filter threshold scaling factor η and kernel sizes $\{m, n\}$ are fixed to 0.7 and $\{7, 5\}$, respectively. We compare with four supervised methods (DORNet [47], DCNAS [72], SGN [49], and DCTNet [69]) and four self-supervised methods (SSGNet [40], MMSR [12], CMSR [38], and P2P [30]). The supervised methods are trained on the RGB-D-D dataset [16].

Table 2 shows quantitative results. On real-world misaligned RGB-depth data, our models outperform other methods by a large margin across SR factors. Figure 7 shows qualitative results. Two-stage methods suffer from poor boundaries (DORNet [47]), ghosting artifacts (SGNet [49]), or spurious textures (P2P [30]), yielding less enhancement here than on synthesized data. This is

Table 3: Real-World misaligned RGB-guided NIR SR. For methods lacking alignment strategies, results without and with pre-alignment are reported (wo/w). *Params vary with SR factors ($\times 2/\times 4$) due to level i . The **best** and second best are highlighted.

	P2P [30]	CMSR [38]	MMSR [12]	SSGNet [40]	Fusion-ArF [35]	Fusion-CAP [26]	PAN [8]	HSI SR [63]	RobSelf-Re	RobSelf-De
Self-Supervised	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓
Pre-Alignment	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗
Params (M)	0.2	1.18	0.25	0.31	47.39	0.38	7.17	10.97	0.91/1.02*	0.96/1.07*
$\times 2$ (RMSE↓)	18.72/18.91	3.42	8.94/8.91	8.96/9.04	6.58	4.40	5.19	4.52	3.09	3.12
$\times 4$ (NIQE↓)	6.32/ 5.32	8.05	10.95/10.50	10.30/10.23	9.54	8.95	9.24	9.01	7.63	<u>7.60</u>





Fig. 8: Real-world misaligned RGB-guided NIR SR ($\times 2$). Image patches are shown due to space limit. *Pre-alignment applied.

because the pre-alignment struggles to generalize to such real-world data, where the misalignment is more complex due to various cross-sensor discrepancies and viewpoint variation. By contrast, our models achieve robustness and generalizability on real-world data with complex, large misalignments, owing to the joint weakly-supervised, misalignment-aware translation formulation (Sec. 5.6). See the supplementary material for analysis of pre-alignment.

5.4 Real-World Misaligned RGB-Guided NIR SR

The estimator level i is set to 4 for $\times 2$ SR and 5 for $\times 4$ SR. The filter kernel sizes $\{m, n\}$ are both set to 3 to better capture the NIR details, with no threshold applied. We compare with prior self-supervised methods (SSGNet [40], MMSR [12], CMSR [38], P2P [30]) and misalignment-aware, supervised methods for guided HSI SR [63], pan-sharpening [8], and fusion [26, 35]. The supervised methods are trained on EPFL RGB-NIR [3] with synthesized misalignment.

Table 3 and Figure 8 show quantitative and qualitative results, respectively. This task is more challenging due to object motion and the rich NIR details. Other methods degrade numerically and suffer from blur (Fusion-CAP [26], SSGNet [40], P2P [30]) or edge distortion (CMSR [38]). In contrast, our models achieve more effective enhancement and recover finer details, thanks to our translator and filter which together enable robust handling of diverse real-world misalignments and effective reconstruction of fine NIR details. See the supplementary material for more qualitative results.

**Fig. 9:** Effectiveness of our translator and filter.**Table 4:** Effect of our translator and filter. For first two settings, $\{\mathbf{F}_{source}, \mathbf{F}_{guide}\}$ or $\{\mathbf{F}_{source}, \mathbf{F}_{guide}^{Aligned}\}$ are fused using a 1×1 conv. No feature fusion in last two settings.

Translator	✗	✓	✗	✓
Filter	✗	✗	✓	✓
RMSE↓	2.55	2.37	2.31	2.18

5.5 Ablation Study

We analyze the overall effect of our translator and filter using RobSelf-De on task II ($\times 2$ SR). As shown in Tab. 4 and Fig. 9 (upper), when both are removed, the result remains unenhanced since neither the misaligned guide nor self-enhancement is effectively exploited. When only the translator is adopted, the result obtains weak enhancement due to the simple fusion of aligned guidance. When only the filter is adopted, the result obtains naive self-enhancement under misaligned guidance and introduces false structures. When the translator is further used alongside the filter to provide aligned guidance for reference-based self-enhancement, the result exhibits both high resolution and high fidelity. The bare baseline can outperform some previous methods because their modules overly fuse the misaligned guidance, which degrades source structures. This also motivates our design of reference-based self-enhancement.

Figure 9 (lower) visualizes the features from the full model. Our translator, by warping $\mathbf{F}_{guide}$ and driving it to mimic source modality via weakly-supervised translation ($\mathbf{I}_{pred}^{Trans}$), effectively derives $\mathbf{F}_{guide}^{Aligned}$ aligned with $\mathbf{F}_{source}$. Since $\mathbf{F}_{guide}^{Aligned}$ remains in the HR space and is supervised only by LR source, it preserves both essential source-corresponding structures and redundant modality-specific content. The latter can neither be aligned due to modality discrepancies nor removed during translation under weak supervision. Therefore, our filter




Fig. 10: Guide alignment performance: our translator vs. modules in [52]. Visualizations for real-world misaligned RGB-guided depth (left) and NIR (right) SR.



Fig. 11: Our translator can “synthesize” missing guide structures. See the right part of the square pot. Visualizations by RobSelf-De for real-world NIR SR ($\times 2$).

uses $\mathbf{F}_{guide}^{Aligned}$ only as reference to learn content-aware kernels for discriminative self-enhancement, which refines $\mathbf{F}_{source}$ into $\mathbf{F}_{source}^{Enhanced}$, recovering fine details and avoiding redundancy effects. Together, our techniques provide an effective inductive bias for robust self-supervised cross-modal SR under misalignment.

See the supplementary material for more analyses on the translator and filter, as well as limitations.

5.6 Interesting Observations

Compared to off-the-shelf registration modules [52], our translator more effectively handles **diverse misalignments**, *e.g.*, inherent cross-sensor misalignment with large viewpoint variation (Fig. 10, left) and non-rigid object motion (Fig. 10, right). Such robustness stems from its weakly-supervised translation objective, which encourages source-consistent guide features.

Moreover, it can **“synthesize” structures that are missing** from the guide images, thereby still providing effective guidance for source enhancement. This arises when some guide structures are absent due to field-of-view differences and occlusions. As shown in Fig. 11, our translator recovers in $\mathbf{F}_{guide}^{Aligned}$ the right part of the square pot, which is entirely missing from the original guide images and features. With the recovered structures, our filter enhances the source feature and produces promising SR prediction. A similar phenomenon is observed in the second scene of Fig. 7—the left part of the cactus. We attribute this ability to the translator’s tendency to “borrow” and align contextually related content in the guide, driven by the translation objective toward mimicking the source.

In **well-aligned scenarios** (Fig. 12), our model achieves competitive performance, once again demonstrating that our translator and filter facilitate effective




Fig. 12: Performance on well-aligned data: synthesized RGB-guided depth SR ($\times 4$).

Table 5: Runtime comparison of self-supervised methods. *Pre-alignment applied.

	P2P* [30]	MMSR* [12]	SSGNet* [40]	RobSelf-De	RobSelf-Re
Task I ($\times 4$). Guide: $512 \times 576 \times 3$; Upsampled source: $512 \times 576 \times 1$.					
RMSE $\downarrow$	2.33	1.88	1.92	1.43	<u>1.52</u>
Time	325s	137s	151s	100s	71s
Task II ($\times 2$). Guide: $576 \times 640 \times 3$; Upsampled source: $576 \times 640 \times 1$.					
RMSE $\downarrow$	3.51	2.76	2.74	2.18	<u>2.23</u>
Time	982s	164s	181s	122s	89s
Task III ($\times 2$). Guide: $576 \times 640 \times 3$; Upsampled source: $576 \times 640 \times 1$.					
RMSE $\downarrow$	18.91	8.91	9.04	<u>3.12</u>	3.09
Time	982s	164s	181s	<u>97s</u>	64s

guide utilization and source enhancement. See the supplementary material for detailed experiments under different misalignment degrees.

5.7 Runtime Efficiency

Table 5 compares the runtime (online optimization + reference) of self-supervised methods on an NVIDIA A100 GPU. SSGNet [40] and MMSR [12] were optimized for 1000 iterations as in our settings. P2P [30] followed its original settings to ensure performance. CMSR [38] had to be run on a different GPU due to its TensorFlow 1 code base and is therefore excluded. Our models achieve superior efficiency compared to existing self-supervised methods. For example, in task III, RobSelf-Re is up to $15.3\times$ faster than P2P and at least $2.56\times$ faster than MMSR and SSGNet while achieving lower RMSE. Such efficiency stems from our lightweight architecture and the filter, which requires no additional processing (*e.g.*, filtering and fusion) for the guide feature.

6 Conclusions

We achieve robust self-supervised cross-modal SR on real-world misaligned data by proposing RobSelf. Within RobSelf, a misalignment-aware feature translator reformulates and resolves unsupervised cross-modal and cross-resolution alignment, and a content-aware reference filter enables effective and faithful source enhancement. RobSelf demonstrates state-of-the-art performance, robustness, generalizability, and superior efficiency, showing promising applicability to practical scenarios where only unlabeled misaligned data is available.

Acknowledgements

This work was supported by JST FOREST (Grant Number JPMJFR206S), JST NEXUS (Grant Number JPMJNX25CA), and JST CRONOS (Grant Number JPMJCS25K5). JL was also supported by the RIKEN Junior Research Associate (JRA) program.

References

1. Almasri, F., Debeir, O.: Multimodal sensor fusion in single thermal image super-resolution. arXiv preprint arXiv:1812.09276 (2018)
2. Arar, M., Ginger, Y., Danon, D., Bermano, A.H., Cohen-Or, D.: Unsupervised multi-modal image registration via geometry preserving image-to-image translation. In: CVPR (2020)
3. Brown, M., Süssstrunk, S.: Multi-spectral SIFT for scene category recognition. In: CVPR (2011)
4. Cao, S.Y., Zhang, R., Luo, L., Yu, B., Sheng, Z., Li, J., Shen, H.L.: Recurrent homography estimation using homography-guided image warping and focus transformer. In: CVPR (2023)
5. Cui, Z., Liu, S., Dong, X., Chu, X., Gu, L., Yang, M.H., Harada, T.: Unifying color and lightness correction with view-adaptive curve adjustment for robust 3D novel view synthesis. arXiv preprint arXiv:2602.18322 (2026)
6. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: ICCV (2017)
7. Deng, X., Liu, E., Gao, C., Li, S., Gu, S., Xu, M.: CrossHomo: Cross-modality and cross-resolution homography estimation. IEEE TPAMI (2024)
8. Do, J., Kim, S., Youk, G., Lee, J., Kim, M.: PAN-Crafter: Learning modality-consistent alignment for PAN-sharpening. In: ICCV (2025)
9. Dong, J., Pan, J., Yang, Z., Tang, J.: Multi-scale residual low-pass filter network for image deblurring. In: ICCV (2023)
10. Dong, W., Xu, Y., Qu, J., Hou, S.: Learning multi-modal cross-scale deformable transformer network for unregistered hyperspectral image super-resolution. In: AAAI (2024)
11. Dong, X., Yokoya, N.: Understanding dark scenes by contrasting multi-modal observations. In: WACV (2024)
12. Dong, X., Yokoya, N., Wang, L., Uezato, T.: Learning mutual modulation for self-supervised cross-modal super-resolution. In: ECCV (2022)
13. Duan, Y., Wu, X., Deng, H., Deng, L.J.: Content-adaptive non-local convolution for remote sensing pansharpening. In: CVPR (2024)
14. Guo, M.: Unsupervised multi-modal medical image registration via invertible translation. In: ECCV (2024)
15. Gupta, H., Mitra, K.: Toward unaligned guided thermal super-resolution. IEEE TIP (2021)
16. He, L., Zhu, H., Li, F., Bai, H., Cong, R., Zhang, C., Lin, C., Liu, M., Zhao, Y.: Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline. In: CVPR (2021)
17. Hirschmuller, H., Scharstein, D.: Evaluation of cost functions for stereo matching. In: CVPR (2007)

18. Huang, Y., Li, L., Zhu, H., Hu, B.: Blind quality index of depth images based on structural statistics for view synthesis. *IEEE SPL* (2020)
19. Kang, J., Cai, Q., Tan, R., Liu, Y., Liu, Z.: C2PD: Continuity-constrained pixelwise deformation for guided depth super-resolution. In: *AAAI* (2025)
20. Kim, B., Ponce, J., Ham, B.: Deformable kernel networks for joint image filtering. *IJCV* (2021)
21. Kim, J., Baek, S.H.: Pixel-aligned RGB-NIR stereo imaging and dataset for robot vision. In: *CVPR* (2025)
22. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR* (2015)
23. Kohli, P., Silberman, N., Hoiem, D., Fergus, R.: Indoor segmentation and support inference from RGBD images. In: *ECCV* (2012)
24. Kong, L., Qi, X.S., Shen, Q., Wang, J., Zhang, J., Hu, Y., Zhou, Q.: Indescribable multi-modal spatial evaluator. In: *CVPR* (2023)
25. Lai, Z., Fu, Y., Zhang, J.: Hyperspectral image super resolution with real unaligned RGB guidance. *IEEE TNNLS* (2024)
26. Li, H., Yang, Z., Zhang, Y., Jia, W., Yu, Z., Liu, Y.: MulFS-CAP: Multimodal fusion-supervised cross-modality alignment perception for unregistered infrared-visible image fusion. *IEEE TPAMI* (2025)
27. Li, R., Okunev, M., Guo, Z., Duong, A.H., Richardt, C., O’Toole, M., Tompkin, J.: Time of the flight of the gaussians: Optimizing depth indirectly in dynamic radiance fields. In: *CVPR* (2025)
28. Lincetto, F., Agresti, G., Rossi, M., Zanuttigh, P.: Multimodalstudio: A heterogeneous sensor dataset and framework for neural rendering across multiple imaging modalities. In: *CVPR* (2025)
29. Liu, D., Wen, B., Fan, Y., Loy, C.C., Huang, T.S.: Non-local recurrent network for image restoration. In: *NeurIPS* (2018)
30. Lutio, R.d., D’Aronco, S., Wegner, J.D., Schindler, K.: Guided super-resolution as pixel-to-pixel transformation. In: *ICCV* (2019)
31. Mei, Y., Fan, Y., Zhou, Y.: Image super-resolution with non-local sparse attention. In: *CVPR* (2021)
32. Metzger, N., Daudt, R.C., Schindler, K.: Guided depth super-resolution by deep anisotropic diffusion. In: *CVPR* (2023)
33. Mok, T.C.W., Li, Z., Bai, Y., Zhang, J., Liu, W., Zhou, Y.J., Yan, K., Jin, D., Shi, Y., Yin, X., Lu, L., Zhang, L.: Modality-agnostic structural image representation learning for deformable multi-modality medical image registration. In: *CVPR* (2024)
34. Qiu, D., Pang, J., Sun, W., Yang, C.: Deep end-to-end alignment and refinement for time-of-flight RGB-D module. In: *ICCV* (2019)
35. Qu, J., Wu, X., Dong, W., Cui, J., Li, Y.: IR&ArF: Toward deep interpretable arbitrary resolution fusion of unregistered hyperspectral and multispectral images. *IEEE TIP* (2025)
36. Ren, J., Jiang, X., Li, Z., Liang, D., Zhou, X., Bai, X.: MINIMA: Modality invariant image matching. In: *CVPR* (2025)
37. Scharstein, D., Pal, C.: Learning conditional random fields for stereo. In: *CVPR* (2007)
38. Shacht, G., Danon, D., Fogel, S., Cohen-Or, D.: Single pair cross-modality super resolution. In: *CVPR* (2021)
39. Shi, L., Zhao, R., Pan, B., Zou, Z., Shi, Z.: Unsupervised multimodal remote sensing image registration via domain adaptation. *IEEE TGRS* (2023)
40. Shin, J., Shin, S., Jeon, H.G.: Task-specific scene structure representations. In: *AAAI* (2023)

41. Song, L., Li, Y., Jiang, Z., Li, Z., Zhang, X., Sun, H., Sun, J., Zheng, N.: Rethinking learnable tree filter for generic feature transform. In: NeurIPS (2020)
42. Song, L., Li, Y., Li, Z., Yu, G., Sun, H., Sun, J., Zheng, N.: Learnable tree filter for structure-preserving feature transform. In: NeurIPS (2019)
43. Song, S., Lew, J., Jang, H., Yoon, S.: Unsupervised homography estimation on multimodal image pair via alternating optimization. In: NeurIPS (2024)
44. Veksler, O., Boykov, Y.: Sparse non-local CRF. In: CVPR (2022)
45. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: CVPR (2018)
46. Wang, Z., Wu, Y., Li, X., Yan, Z., Yang, J.: Spatiotemporal difference network for video depth super-resolution. In: AAAI (2026)
47. Wang, Z., Yan, Z., Pan, J., Gao, G., Zhang, K., Yang, J.: DORNet: A degradation oriented and regularized network for blind depth super-resolution. In: CVPR (2025)
48. Wang, Z., Yan, Z., Wu, Y., Gao, G., Li, X., Yang, J.: Multi-order matching network for alignment-free depth super-resolution. arXiv preprint arXiv:2511.16361 (2026)
49. Wang, Z., Yan, Z., Yang, J.: SGNet: Structure guided network via gradient-frequency awareness for depth map super-resolution. In: AAAI (2024)
50. Wang, Z., Yan, Z., Yang, M.H., Pan, J., Yang, J., Tai, Y., Gao, G.: Scene prior filtering for depth map super-resolution. IJCV (2026)
51. Xiang, J., Zhu, X., Wang, X., Wang, Y., Zhang, H., Guo, F., Yang, X.: DEPTHOR: Depth enhancement from a practical light-weight dToF sensor and RGB image. In: ICCV (2025)
52. Xu, H., Yuan, J., Ma, J.: MURF: Mutually reinforcing multi-modal image registration and fusion. IEEE TPAMI (2023)
53. Xuan, K., Xiang, L., Huang, X., Zhang, L., Liao, S., Shen, D., Wang, Q.: Multi-modal MRI reconstruction assisted with spatial alignment network. IEEE TMI (2022)
54. Yan, Z., Wang, K., Li, X., Zhang, Z., Li, G., Li, J., Yang, J.: Learning complementary correlations for depth super-resolution with incomplete data in real world. IEEE TNNLS (2022)
55. Yan, Z., Wang, Z., Dong, H., Li, J., Yang, J., Lee, G.H.: Ducos: Duality constrained depth super-resolution via foundation model. In: ICCV (2025)
56. Yan, Z., Wang, Z., Yang, J., Lee, G.H.: DegBins: Degradation-driven binning for depth super-resolution. arXiv preprint arXiv:2605.09628 (2026)
57. Yin, B.W., Cao, J.L., Zhang, X., Chen, Y., Cheng, M.M., Hou, Q.: OmniSegmentor: A flexible multi-modal learning framework for semantic segmentation. In: NeurIPS (2025)
58. Yuan, J., Jiang, H., Li, X., Qian, J., Li, J., Yang, J.: Structure flow-guided network for real depth super-resolution. In: AAAI (2023)
59. Zhang, K., Ma, J.: Sparse-to-dense multimodal image registration via multi-task learning. In: ICML (2024)
60. Zhang, L., Zhang, L., Bovik, A.C.: A feature-enriched completely blind image quality evaluator. IEEE TIP (2015)
61. Zhang, R., Ma, J., Cao, S.Y., Luo, L., Yu, B., Chen, S.J., Li, J., Shen, H.L.: SCPNet: Unsupervised cross-modal homography estimation via intra-modal self-supervised learning. In: ECCV (2024)
62. Zhang, S., Dong, J., Ma, Y., Cai, H., Wang, M., Li, Y., Kabika, T.B., Li, X., Hou, W.: CDF-DSR: Learning continuous depth field for self-supervised RGB-guided depth map super resolution. Information Fusion (2025)
63. Zhang, Y., Lai, Z., Zhang, T., Fu, Y., Zhou, C.: Unaligned RGB guided hyperspectral image super-resolution with spatial-spectral concordance. IJCV (2025)

64. Zhao, Y., Huang, X., Zhang, Z.: Deep Lucas-Kanade homography for multimodal image alignment. In: CVPR (2021)
65. Zhao, Z., Bai, H., Ke, B., Cui, Y., Deng, L., Zhang, Y., Zhang, K., Schindler, K.: A unified solution to video fusion: From multi-frame learning to benchmarking. In: NeurIPS (2025)
66. Zhao, Z., Bai, H., Zhang, J., Zhang, Y., Zhang, K., Xu, S., Chen, D., Timofte, R., Van Gool, L.: Equivariant multi-modality image fusion. In: CVPR (2024)
67. Zhao, Z., Bai, H., Zhu, Y., Zhang, J., Xu, S., Zhang, Y., Zhang, K., Meng, D., Timofte, R., Van Gool, L.: DDFM: Denoising diffusion model for multi-modality image fusion. In: ICCV (2023)
68. Zhao, Z., Zhang, J., Gu, X., Tan, C., Xu, S., Zhang, Y., Timofte, R., Van Gool, L.: Spherical space feature decomposition for guided depth map super-resolution. In: ICCV (2023)
69. Zhao, Z., Zhang, J., Xu, S., Lin, Z., Pfister, H.: Discrete cosine transform network for guided depth map super-resolution. In: CVPR (2022)
70. Zheng, B., Yuan, S., Slabaugh, G., Leonardis, A.: Image demoiréing with learnable bandpass filters. In: CVPR (2020)
71. Zhong, Z., Chen, P., Shen, Q., Li, B., Wang, S.: Dual graph regularized deep unfolding network for guided depth map super-resolution. In: CVPR (2026)
72. Zhong, Z., Liu, X., Jiang, J., Zhao, D., Wang, S.: Dual-level cross-modality neural architecture search for guided image super-resolution. IEEE TPAMI (2025)