

# Unleashing Network Potentials for Semantic Scene Completion

Fengyun Wang<sup>1</sup> Qianru Sun<sup>2</sup> Dong Zhang<sup>3</sup> Jinhui Tang<sup>1\*</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanjing University of Science & Technology

<sup>2</sup>Singapore Management University <sup>3</sup>The Hong Kong University of Science & Technology

E-mail: {fereenwang, jinhuitang}@njust.edu.cn; qianrusun@smu.edu.sg; dongz@ust.hk

## Abstract

Semantic scene completion (SSC) aims to predict complete 3D voxel occupancy and semantics from a single-view RGB-D image, and recent SSC methods commonly adopt multi-modal inputs. However, our investigation reveals two limitations: ineffective feature learning from single modalities and overfitting to limited datasets. To address these issues, this paper proposes a novel SSC framework - Adversarial Modality Modulation Network (AMMNet) - with a fresh perspective of optimizing gradient updates. The proposed AMMNet introduces two core modules: a cross-modal modulation enabling the interdependence of gradient flows between modalities, and a customized adversarial training scheme leveraging dynamic gradient competition. Specifically, the cross-modal modulation adaptively re-calibrates the features to better excite representation potentials from each single modality. The adversarial training employs a minimax game of evolving gradients, with customized guidance to strengthen the generator’s perception of visual fidelity from both geometric completeness and semantic correctness. Extensive experimental results demonstrate that AMMNet outperforms state-of-the-art SSC methods by a large margin, providing a promising direction for improving the effectiveness and generalization of SSC methods. [Our code is available at this link.](#)

## 1. Introduction

Semantic scene completion (SSC) is a crucial task in 3D scene understanding domain that seeks to forecast complete 3D voxel occupancy and semantics from a single-view RGB-D image [17, 20, 21]. Current SSC methods rely on multi-modal inputs like RGB images and depth represented as Truncated Signed Distance Function (TSDF) representations, as shown in Figure 1 (b), which provide complementary cues for scene geometry reconstruction and semantic prediction. These methods typically utilize an encoder-


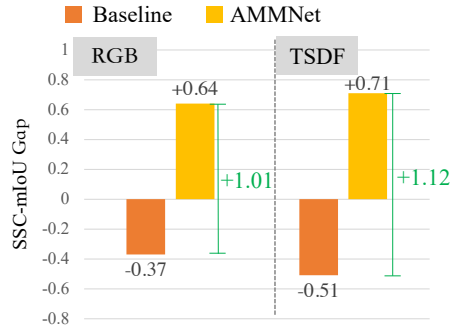


Figure 1. **Comparisons of encoder representation power.** For the multi-modal training in (b), the informative representations in RGB and TSDF are not fully unleashed compared to learning them individually in (a). The proposed Adversarial Modality Modulation Network (AMMNet) in (c) enables a more thorough unleashing of potentials via “cross-modal modulation  $\mathcal{M}$ ” and “adversarial training  $\mathcal{L}_{(D,G)}$ ”.

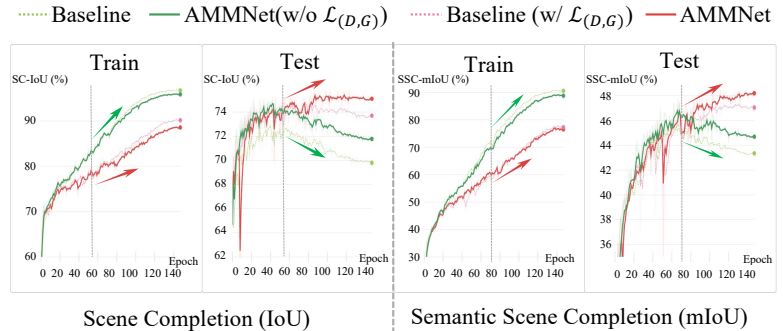
decoder paradigm to exploit multi-modal data, where the RGB image and TSDF are encoded separately and subsequently combined for final predictions. Despite the promising results demonstrated by these multi-modal models on indoor scene completion benchmarks [29], two key observations can still be made from these models [4, 34].

**Observation 1:** *When learned jointly in multi-modal models, the rich information in individual modalities is not sufficiently unleashed compared to single-modal models.* To validate this observation, we conducted a two-stage experiment. In the first stage, we trained uni-modal and multi-modal SSC models using the same dataset and settings. In the second stage, we evaluated separate single-modal networks initialized with encoder weights from stage one models to assess the learned representations. Through controlled comparison, we found degraded representation capabilities of the multi-modal encoders compared to their uni-modal counterparts. Specifically, we examined the representation power of RGB and TSDF encoders by initializing two RGB-only networks and two TSDF-only networks using corresponding encoder weights from uni-modal and multi-modal stage one models. With the encoders frozen, the decoders were then trained. As the only difference was the frozen encoder, performance gaps demonstrated the insufficient unleashing of modalities in multi-modal training.

\*Corresponding author.



(a) Comparison of learning ability of the encoder



(b) SSC-mIoU and SC-IoU over epochs during the training

Figure 2. **Two key observations on multi-modal SSC models.** (a) Performance drops of “multi-modal encoders” compared to single-modal counterparts, validating insufficient unleashing of modalities in joint training. “Our method” demonstrates significantly enhanced encoder capabilities. (b) “Diverging training/validation curves” of baseline and the variant AMMNet, indicating overfitting issues. Under the adversarial training scheme  $\mathcal{L}_{(D,G)}$ , our model alleviates overfitting and achieves “steadily increasing performance”.

As shown in Figure 2 (a), employing the multi-modal RGB encoder led to a performance drop of 0.37% in terms of SSC-mIoU compared to utilizing the single-modal RGB encoder, on the NYU dataset [29]. Similarly, adopting the multi-modal TSDF encoder incurred a 0.51% decrease in SSC-mIoU compared to the single-modal TSDF encoder.

**Observation 2:** *Deep SSC models trained with limited scene data are prone to overfitting.* To validate this observation, we examined the training processes of a baseline model [34] and a variant of our AMMNet without the proposed adversarial training scheme  $\mathcal{L}_{(D,G)}$ . We observed severe overfitting behaviors, where models first reached optimal validation performance but further training led to increasing divergence between training and validation. As illustrated in Figure 2, the baseline model [34] (dotted green line) achieved the best validation score in the middle of training, while later epochs led to a 22.8% SSC-mIoU increase on the training set but a 2.8% SSC-mIoU drop on the validation set. A similar divergence was observed for the AMMNet w/o  $\mathcal{L}_{(D,G)}$  (solid green line), validating that SSC models tend to overfit the training data.

To address these issues, we propose an Adversarial Modality Modulation Network (AMMNet), a novel SSC framework to better unleash the potential by optimizing gradient updating. A conceptual illustration is presented in Figure 1 (c), it consists of two key components designed to address the identified issues. 1) A cross-modal modulation is introduced to better unleash the potentials of individual modalities. With inter-dependent gradient updating across modalities, it can stimulate the encoders to fully unleash the representations of RGB and TSDF in joint training. Specifically, it adaptively recalibrates the RGB features by incorporating information from the TSDF. As illustrated in Figure 2 (a), compared to the encoder in the multi-modal baseline [34] on the NYU [29] dataset, the RGB

and TSDF encoders in AMMNet demonstrate improved capabilities, with 1.01% and 1.12% higher SSC-mIoU, respectively. 2) A customized adversarial training scheme is developed to alleviate overfitting. The minimax competition in the scheme dynamically stimulates the continuous evolution of the models. To provide effective supervision, particularly for SSC, we construct two types of perturbed ground truths: one with disrupted geometric completeness and the other with randomly shuffled semantic categories. These perturbed ground truths are fed to the discriminator as fake samples, explicitly enhancing the discriminator’s ability to recognize flaws in both geometry and semantics. Figure 2 (b) shows that by incorporating the proposed adversarial training scheme, both the baseline model [34] and our AMMNet achieve steadily increasing training and validation accuracy over epochs.

Our main contributions of AMMNet are thus two-fold. 1) we proposed a cross-modal modulation module to better exploit single-modal representations in multi-modal learning. 2) we developed a customized adversarial training scheme to prevent overfitting. Experimental results validated the effectiveness of our AMMNet, showing that it outperforms state-of-the-art SSC methods by significant margins, *e.g.*, improving SSC-mIoU by 3.5% on NYU [29] and 3.3% on NYUCAD [11] compared to previous methods.

## 2. Related Work

**Semantic Scene Completion (SSC).** SSC can be roughly categorized into single-modal methods and multi-modal methods based on the usage of input modalities. Single-modal methods take only TSDF [7, 30] or points [38] converted from depth as input. Volume-based methods like SSCNet [30] adopt 3D CNNs on TSDF. Point-based methods like SPCNet [38] avoid voxel discretization but are

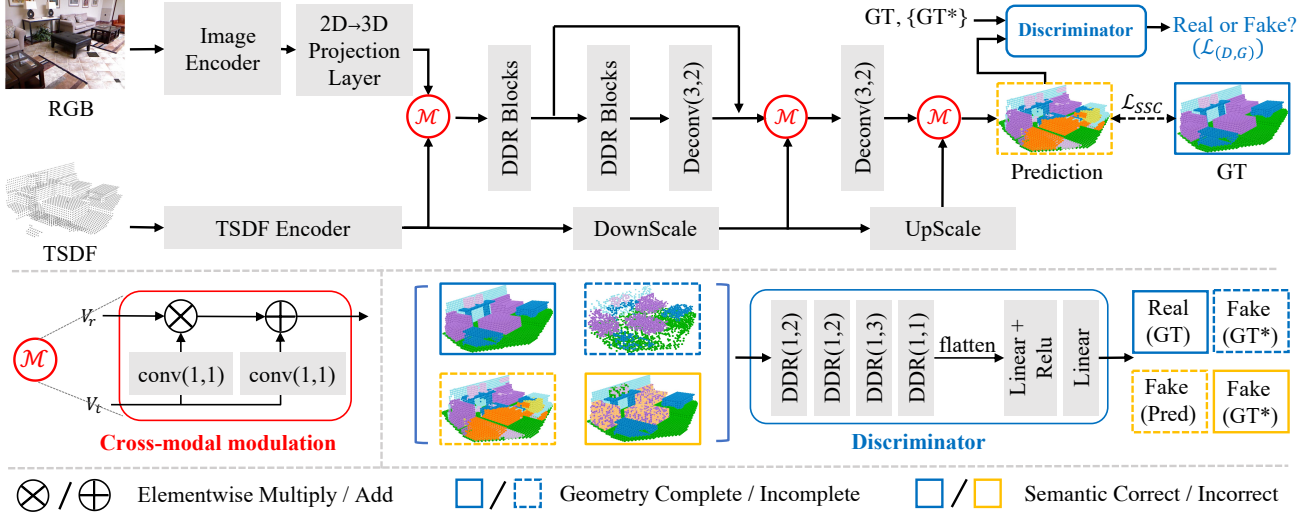


Figure 3. **The overall framework of our AMMNet.** It consists of three components: an image encoder for RGB input, a TSDf encoder for TSDf input, and a decoder for final prediction. It has two novel modules: cross-modal modulations after the encoders and decoder to recalibrate features, and a discriminator that distinguishes real/fake voxels to mitigate overfitting issues. The  $\text{conv}(k, s)$ / $\text{Deconv}(k, s)$  denotes 3D conv/deconv layer with kernel size  $k$  and stride  $s$ , and  $\text{DDR}(d, s)$  denotes DDR layer [17] with dilation  $d$  and stride  $s$ .

prone to noise. Single-modal inputs are limited in complementary cues. Multi-modal methods can be further grouped. Some methods process depth as 2D images using 2D CNNs [15, 17, 18, 20], which are less effective for 3D geometry. Other works encode depth as points [27, 33] or TSDf [4, 34], and adopt dual-branch networks fusing RGB, TSDf, and/or points, achieving state-of-the-art performance. However, they overlook the problems of insufficient encoder feature learning and overfitting as observed in our study. Our work is the first tailored solution addressing these limitations for advanced multi-modal SSC.

**Multi-Modal Learning.** Multi-modal learning has attracted increasing attention owing to the growing availability of multi-modal data. However, it does not always achieve synergistic performance surpassing the sum of individual modalities [25]. The crux lies in the inadequate harnessing of all modal information by most joint training methods [26]. Prior works have explored various directions to address this challenge. Hu et al. proposed a temporal multi-modal deep learning architecture that transforms connected multi-modal restricted Boltzmann machines into probabilistic sequence models, achieving more efficient joint feature learning [16]. Du et al. referred to inferior joint training outcomes as "modality failure" and proposed combining fusion objectives with uni-modal distillation [9]. Instead of designing complex joint training architectures, we aim to enable fuller unleashing of modal representations by proposing simple yet effective solutions - cross-modal modulation and adversarial training.

**Overfitting Problem in Deep Learning.** Various strate-

gies have been explored to alleviate the overfitting problem in deep neural networks. They can be grouped into three categories [1]: 1) *Passive Methods*, like neural architecture search [39] and ensembling [12], aim to determine suitable model configurations before training starts. However, such static configurations lack the flexibility to adapt to dynamic training processes. 2) *Active Methods*, include techniques like Dropout [31], augmentation [28], and Normalization [14]. In this category, the model architecture remains unchanged, but certain model components are unleashed during each training step. 3) *Semi-active Methods*, like pruning [24] and network construction [10], updates the network by adding or removing neurons or connections. Utilizing an adversarial training scheme also falls under this category. Unlike other methods requiring complex network search or design, our adversarial training scheme uses a minimax game to stimulate the model evolution, alleviating overfitting without architectural optimization. To our best knowledge, this is the first attempt at utilizing an adversarial training scheme to mitigate overfitting in SSC.

### 3. Preliminaries

The baseline model [34] for SSC typically contains an image encoder for RGB input, a TSDf encoder for TSDf input, and a decoder for final prediction. Specifically, the RGB encoder extracts a  $D$ -channel 2D feature  $\mathbf{F}_r$  from the RGB input. To obtain 3D representations, a 2D-3D projection layer is applied to map  $\mathbf{F}_r$  to visible surfaces based on per-pixel depth values, while filling other areas with zeros. This results in a 3D RGB feature  $\mathbf{V}_r \in \mathbb{R}^{D \times G_x \times G_y \times G_z}$ ,

where  $G_x$ ,  $G_y$ , and  $G_z$  denote the width, height, and depth of the voxel grid. The TSDF encoder employs 3D convolutions to encode the input voxels to 3D feature maps  $\mathbf{V}_t \in \mathbb{R}^{D \times G_x \times G_y \times G_z}$ . The projected 3D RGB feature  $\mathbf{V}_r$  is fused with  $\mathbf{V}_t$  via element-wise addition as:

$$\bar{\mathbf{V}}_{cm} = \mathbf{V}_r + \mathbf{V}_t. \quad (1)$$

The fused feature  $\bar{\mathbf{V}}_{cm}$  is fed into the decoder, which applies 3D convolutions with skip connection to generate the completed voxel grid  $\hat{\mathbf{Y}} \in \mathbb{R}^{(C+1) \times G_x \times G_y \times G_z}$ , where  $C$  is the number of classes and  $+1$  is an additional channel indicating voxel occupancy, *i.e.*, empty or not.

This structure has been widely adopted as a basic framework for multi-modal SSC models, such as [4, 34]. Such a simple addition fails to fully unleash the potential of each modality as illustrated in observation 1.

## 4. Adversarial Modality Modulation Network

In this section, we will elaborate on the proposed AMM-Net. We introduce two key components in AMMNet to address the aforementioned limitations of existing methods: a cross-modal modulation module and a customized adversarial training scheme.

### 4.1. Cross-Modal Modulation

The cross-modal modulation is expressed as a red circled  $\mathcal{M}$  in Figure 3. It adaptively recalibrates the RGB features by incorporating information from the TSDF, which enables an interdependent gradient updating across modalities. It first transforms  $\mathbf{V}_t$  into a scale  $\mathbf{M}_s$  and a bias  $\mathbf{M}_b$  using two  $1 \times 1 \times 1$  convolutional layers.  $\mathbf{M}_s$  and  $\mathbf{M}_b$  share the same shape as  $\mathbf{V}_r$ . Then, it modulates  $\mathbf{V}_r$  via an element-wise operation, which is expressed as:

$$\hat{\mathbf{V}}_{cm} = \mathbf{V}_r \otimes (1 + \sigma(\mathbf{M}_s)) + \mathbf{M}_b. \quad (2)$$

where  $\otimes$  denotes element-wise multiplication, and  $\sigma$  is the activation function *sigmoid*. The recalibrated feature  $\hat{\mathbf{V}}_{cm}$  is then fed to the decoder. Compared to the element-wise addition operation in the baseline model [34], our modulation enables interdependent gradient updating across modalities. Specifically, the gradient, w.r.t.  $\mathbf{V}_r$ , in the modulation is formulated as:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{V}_r} = \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{V}}_{cm}} \otimes (1 + \sigma(\mathbf{M}_s)) \quad (3)$$

where  $\mathcal{L}$  is the loss function. The gradient w.r.t.  $\mathbf{V}_r$  depends on  $\mathbf{V}_t$  via the scale term  $\mathbf{M}_s$ . Similarly, the gradient  $\mathbf{V}_t$  depends on RGB contexts. In contrast, in the baseline model [34], the gradient  $\mathbf{V}_r$  is expressed as:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{V}_r} = \frac{\partial \mathcal{L}}{\partial \bar{\mathbf{V}}_{cm}}, \quad (4)$$

which is independent of  $\mathbf{V}_t$  and is not desirable.

Meanwhile, the cross-modal modulation enables an improved forward fusion of the modalities. We will validate through ablation studies in Sec. 5.5 that the interdependent gradient updating plays a more critical role than the improved modality fusion. As shown in Figure 3, we apply the cross-modal modulation on the last two layers of the decoder features similarly. To match the resolution with the decoder features, the TSDF feature  $\mathbf{V}_t$  first goes through a DownScale module, resulting in downscaled TSDF features  $\mathbf{V}'_t \in \mathbb{R}^{D \times \frac{G_x}{2} \times \frac{G_y}{2} \times \frac{G_z}{2}}$ . The resulting feature  $\mathbf{V}'_t$  modulates the second last decoder layer output. After that,  $\mathbf{V}'_t$  further goes through an UpScale module to restore the size and is used to modulate the last decoder layer output in the same way.

### 4.2. Adversarial Training

To prevent overfitting, a customized discriminator is designed to introduce dynamic adversarial gradients. By competing against the generator (*i.e.*, the voxel predictor in this work) in a minimax game, the discriminator provides continuous and adaptive supervision besides the static losses. In SSC tasks, complete spatial structures and correct semantic perceptions are vital. Accordingly, we customize the discriminator in three aspects:

**Geometry Completeness.** To make the discriminator more aware of geometry completeness, we construct additional fake samples by perturbing the geometry of ground truth voxels. Specifically, we randomly erase some non-empty voxels in the ground truth maps to “empty”, by:

$$\bar{\mathbf{y}}_i^G = \begin{cases} \text{empty}, & r_i^G |_{r_i^G \sim U(0,1)} \leq p^G \\ \mathbf{y}_i, & \text{otherwise,} \end{cases} \quad (5)$$

where  $r_i^G$  is a random number that follows a uniform distribution  $U(0, 1)$ ,  $p^G$  is the erasing probability,  $\mathbf{y}_i \in \mathbf{Y}$  is a voxel in the ground truth,  $\bar{\mathbf{y}}_i^G \in \bar{\mathbf{Y}}^G$  is the geometrically perturbed counterpart. This operation damages the integrity of the 3D structures randomly. Using these examples (denoted as “fake”) in training, we force the discriminator to learn to pay more attention to the completeness and continuity of geometric shapes, such as to feedback to and guide the SSC voxel predictor to achieve high completeness.

**Semantic Correctness.** To force the discriminator to be more aware of semantic correctness, we construct another set of fake samples by shuffling semantic categories in the ground truth. For a scene with  $m$  non-empty category labels, we randomly select  $n \leq m$  categories to introduce perturbation, where  $n$  is also randomly determined. The perturbation for a particular category  $c_j$  can be formed as:

$$\bar{\mathbf{y}}_{i|\{i \in \{y_i=c_j\}\}}^S = \begin{cases} \text{Map}_{c_j \rightarrow c_k, j \neq k}(\mathbf{y}_i), & r_i^S |_{r_i^S \sim U(0,1)} \leq p_j^S \\ \mathbf{y}_i, & \text{otherwise,} \end{cases} \quad (6)$$



Methods	Inputs	Scene Completion(%)			SSC-mIoU(%)												
		Prec.	Recall	IoU	ceiling	floor	wall	win.	chair	bed	sofa	table	TVs	furn.	objs.	avg.	
SSCNet [30]	D	57.0	<b>94.5</b>	55.1	15.1	94.7	24.4	0.0	12.6	32.1	35.0	13.0	7.8	27.1	10.1	24.7	
CCPNet [37]	D	74.2	90.8	63.5	23.5	<b>96.3</b>	35.7	20.2	25.8	61.4	56.1	18.1	28.1	37.8	20.1	38.5	
DDRNet [17]	RGB+D	71.5	80.8	61.0	21.1	92.2	33.5	6.8	14.8	48.3	42.3	13.2	13.9	35.3	13.2	30.4	
AIC-Net [18]	RGB+D	62.4	91.8	59.2	23.2	90.8	32.3	14.8	18.2	51.1	44.8	15.2	22.4	38.3	15.7	33.3	
3D-Sketch [4]	RGB+D	85.0	81.6	71.3	43.1	93.6	40.5	24.3	30.0	57.1	49.3	29.2	14.3	42.5	28.6	41.1	
FFNet [35]	RGB+D	89.3	78.5	71.8	44.0	93.7	41.5	29.3	36.2	59.0	51.1	28.9	26.5	45.0	32.6	44.4	
CleanerS [34]	RGB+D	88.0	83.5	75.0	46.3	93.9	43.2	33.7	38.5	62.2	54.8	33.7	39.2	45.7	33.8	47.7	
PCANet [19]	RGB+D	89.5	87.5	<b>78.9</b>	44.3	94.5	50.1	30.7	41.8	<b>68.5</b>	56.4	32.6	29.9	53.6	35.4	48.9	
SISNet <sub>DLabv3</sub> [2]	RGB+D	92.1	83.8	78.2	54.7	93.8	<b>53.2</b>	41.9	43.6	66.2	61.4	38.1	29.8	53.9	40.3	52.4	
CVSformer <sub>DLabv3</sub> [6]	RGB+D	87.7	82.1	73.7	46.3	94.3	43.5	<b>42.8</b>	46.2	67.7	66.0	39.2	43.2	53.9	35.0	52.6	
AMMNet	RGB+D	<b>90.5</b>	82.1	75.6	46.7	94.2	43.9	30.6	39.1	60.3	54.8	35.7	44.4	48.2	35.3	48.5	
AMMNet <sub>DLabv3</sub>	RGB+D	88.7	84.5	76.3	<b>49.2</b>	94.2	47.0	41.7	<b>52.4</b>	68.1	<b>66.4</b>	<b>46.4</b>	<b>52.4</b>	<b>58.3</b>	<b>41.1</b>	<b>56.1</b>	

Table 1. Result comparisons on the test set of NYU [29]. Results with “<sub>DLabv3</sub>” denote that these results are based on DeepLabv3 [3] as the backbone network for RGB image feature extraction. Bold numbers represent the best performance.

where  $r_i^S$  is a random number that follows a uniform distribution  $U(0, 1)$ ,  $p_j^S$  is a random perturbation probability for category  $c_j$ ,  $Map(\cdot)$  is an operation that transforms the categorical labels by mapping the class  $c_j$  to a different class  $c_k$ , with  $j, k = 1, 2, \dots, C$ . Note that the category indices  $j, k$  start from 1 instead of 0, as  $c_0$  denotes empty voxels that remain unchanged. By discerning fakes with shuffled semantics, the discriminator is compelled to develop a more robust understanding of inter-class contexts. Its adversarial gradients can thus feedback to and guide the SSC voxel predictor towards the increase of semantic correctness.

**Structurally Lightweight.** We adopt the lightweight DDR layer as a basic building unit to avoid excessive overheads. As illustrated in Figure 3, the discriminator network  $D$  consists of a series of DDR layers with different strides, followed by a few linear layers. It outputs a confidence score indicating whether the input voxel is real or fake, providing an overall understanding of the global structural layout and semantic coherence. The overall adversarial training objective is thus as follows,

$$\begin{aligned} \mathcal{L}_{(D,G)} = \min_G \max_D [ & \mathbb{E}_{\mathbf{Y}}[\log D(\mathbf{Y})] + \mathbb{E}_{\hat{\mathbf{Y}}}[\log(1 - D(\hat{\mathbf{Y}}))] \\ & + \mathbb{E}_{\tilde{\mathbf{Y}}^G}[\log(1 - D(\tilde{\mathbf{Y}}^G))] + \mathbb{E}_{\tilde{\mathbf{Y}}^S}[\log(1 - D(\tilde{\mathbf{Y}}^S))] ] \end{aligned} \quad (7)$$

where  $\hat{\mathbf{Y}}$  is the output of  $G$ , *i.e.*, the voxel predictor. It tries to: 1) maximize the probability of correctly distinguishing real/fake voxels by  $D$ ; 2) minimize the probability of generated voxels by  $G$  being classified as real by  $D$ .

### 4.3. Overall Loss Function

The full training objective of the proposed AMMNet contains two parts: the SSC loss  $\mathcal{L}_{SSC}$  and the proposed adversarial loss  $\mathcal{L}_{(D,G)}$ :

$$\mathcal{L}_{all} = \mathcal{L}_{SSC} + \beta \mathcal{L}_{(D,G)}, \quad (8)$$

where  $\beta$  is a coefficient balancing the two parts.  $\mathcal{L}_{SSC}$  provides supervision for predicting voxels. Following [34], it contains two terms:

$$\mathcal{L}_{SSC} = SCE(\hat{\mathbf{Y}}, \mathbf{Y}) + \lambda SCE(\hat{\mathbf{Y}}_{2D}, \mathbf{Y}_{2D}), \quad (9)$$

where  $SCE$  denotes the smooth cross entropy loss [32]. The first term is applied to the 3D voxel predictions. The second term provides 2D supervision  $\mathbf{Y}_{2D}$  for the image semantic prediction  $\hat{\mathbf{Y}}_{2D}$ . Specifically,  $\hat{\mathbf{Y}}_{2D}$  is obtained by applying an additional convolutional layer on 2D RGB features, and  $\mathbf{Y}_{2D}$  is obtained by back-projecting  $\mathbf{Y}$  to 2D space as described in [34].  $\lambda$  is a coefficient balancing the two loss terms.

## 5. Experiments

### 5.1. Datasets and Evaluation Metrics

**Datasets.** Following [4, 13, 30], we conduct experiments on NYU [29] and NYUCAD [11] datasets which are common indoor benchmarks of SSC. NYU contains 1,449 noisy RGB-D images captured by Kinect sensors. NYUCAD is synthesized from CAD models and has noise-free depth. We use the standard split with 795 images for training and 654 for testing. To select values of hyperparameters, we randomly split one more small set of 100 samples from the training set and take it as a validation set.

**Evaluation Metrics.** Two metrics are used for evaluation: 1) Scene Completion (SC) IoU measuring voxel occupancy prediction accuracy for occluded voxels. It evaluates scene completion capability, and for it, we report precision, recall, and IoU; 2) Semantic Scene Completion (SSC) mIoU measures the accuracy of predicting semantic labels for occupied voxels, where the per-class IoU, as well as the mean IoU (mIoU) across all classes, are reported.

Methods	Inputs	Scene Completion(%)			Semantic Scene Completion(%)												
		Prec.	Recall	IoU	ceiling	floor	wall	win.	chair	bed	sofa	table	TVs	furn.	objs.	avg.	
SSCNet [30]	D	75.4	<b>96.3</b>	73.2	32.5	92.6	40.2	8.9	33.9	57.0	59.5	28.3	8.1	44.8	25.1	40.0	
CCPNet [37]	D	91.3	92.6	82.4	56.2	94.6	58.7	35.1	44.8	68.6	65.3	37.6	35.5	53.1	35.2	53.2	
DDRNet [17]	RGB+D	88.7	88.5	79.4	54.1	91.5	56.4	14.9	37.0	55.7	51.0	28.8	9.2	44.1	27.8	42.8	
AIC-Net [18]	RGB+D	88.2	90.3	80.5	53.0	91.2	57.2	20.2	44.6	58.4	56.2	36.2	9.7	47.1	30.4	45.8	
3D-Sketch [4]	RGB+D	90.6	92.2	84.2	59.7	94.3	64.3	32.6	51.7	72.0	68.7	45.9	19.0	60.5	38.5	55.2	
PCANet [19]	RGB+D	92.1	91.8	84.3	54.8	93.1	62.8	44.3	52.3	75.6	70.2	46.9	44.8	65.3	45.8	59.6	
SISNet <sub>D</sub> Labv3 [2]	RGB+D	94.1	91.2	<b>86.3</b>	63.4	94.4	<b>67.2</b>	52.4	59.2	77.9	71.1	51.8	46.2	65.8	<b>48.8</b>	63.5	
CVSformer <sub>D</sub> Labv3 [6]	RGB+D	94.0	91.0	86.0	<b>65.6</b>	94.2	60.6	<b>54.7</b>	60.4	<b>81.8</b>	71.3	49.8	55.5	65.5	43.5	63.9	
AMMNet	RGB+D	92.4	88.4	82.4	61.3	<b>94.7</b>	65.0	38.9	58.1	76.3	73.2	47.3	46.6	62.0	42.6	60.5	
AMMNet <sub>D</sub> Labv3	RGB+D	<b>92.8</b>	89.0	83.3	65.5	94.6	65.8	52.0	<b>64.9</b>	81.1	<b>78.2</b>	<b>57.6</b>	<b>66.9</b>	<b>67.4</b>	45.1	<b>67.2</b>	

Table 2. Result comparisons on the test set of NYUCAD [11]. Results with “<sub>D</sub>Labv3” denote that these results are based on DeepLabv3 [3] as the backbone network for RGB image feature extraction. Bold numbers represent the best performance.

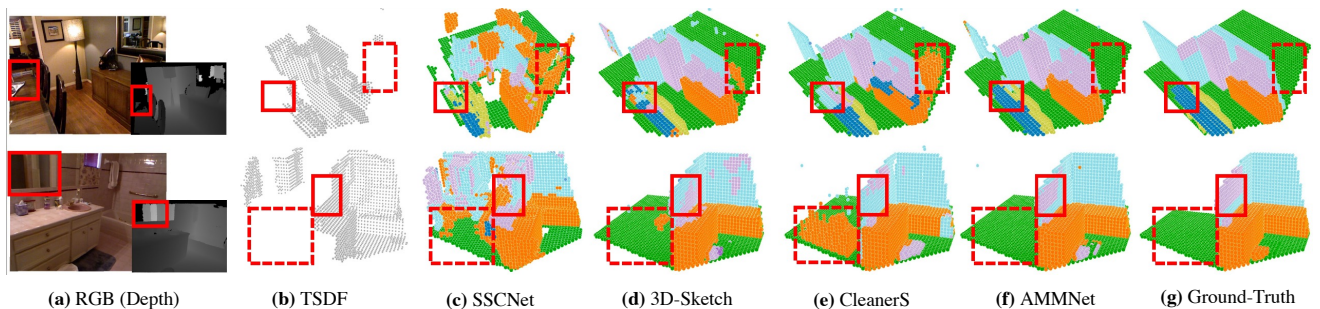


Figure 4. Qualitative comparison on challenging indoor scenes from the test set of NYU [29] with state-of-the-art methods, including SSCNet [30], 3D-Sketch [4], and CleanerS [34].

## 5.2. Implementation Details

**Network Architectures.** For the image encoder, the default setting follows [34] by adopting Segformer-B2 [36] backbone with four MLP layers for feature extraction. The image encoder initializes from ImageNet [5] pre-trained weights and freezes the backbone while keeping the MLP head trainable. To facilitate a fair comparison with CVSformer [6], we alternatively incorporate the pre-trained DeepLabv3 model as the image encoder, which was obtained by training for 1,000 epochs on the RGB image segmentation task and freeze its parameters. The rest of the network trains from scratch. The TSDF encoder consists of three layers of 3D convolution, two DDR [17] blocks, and two layers of 3D deconvolution. The discriminator consists of four DDR [17] layers with stride 2, 2, 3, 1 respectively for downsampling, followed by flattening and two linear layers for prediction. The feature channel size  $D = 256$  and the 3D resolution  $G_x, G_y, G_z$  of prediction is set to (60, 36, 60). In the loss function,  $\lambda$  is set as 0.25, following [34],  $\beta$  is set as 0.005. In the GT perturbations, the probability  $P^G, P_j^S$  are set as random variables uniformly distributed between 0.1 and 0.9, a detailed analysis for the hyperparameters is given in Sec. 5.5.

**Training Settings.** We implement all experiments using

PyTorch on 2 NVIDIA 3090 GPUs. The model is optimized with AdamW [23] using a weight decay of 0.05 and an initial learning rate of 0.001. The learning rate is scheduled in a cosine decay policy [22] with a minimum value of  $1e-7$ . Training is conducted for 150 epochs with a batch size of 4. Common data augmentation is applied, including resize, random crop, flip for 2D images, and random x/z axis flip as well as x-z permutation for 3D volumes, following [8].

## 5.3. Comparisons with State-of-the-Arts.

To evaluate AMMNet, we compare it against several recent state-of-the-art methods, including single-modal methods SSCNet [30] and CCPNet [37], and multi-modal methods—DDRNet [17], AIC-Net [18], 3D-Sketch [4], FFNet [35], CleanerS [34], PCANet [19], SISNet [2], and CVSformer [6].

**Quantitative Comparisons.** 1) **Results on NYU [29].** Table 1 shows AMMNet achieves superior performance on real NYU [29] data, outperforming CVSformer [6], the previous best method, by 3.5% for SSC-mIoU. Moreover, CleanerS [34] has the most relevant setting as our AMMNet, using the same RGB and TSDF inputs, and the same baseline architecture, while relying on clean CAD data and a complex distillation pipeline. In contrast, our approach

Methods	$\mathcal{M}$	$\mathcal{L}_{(D,G)}$	$\mathcal{L}_{SSC}$	SC-IoU	SSC-mIoU
Baseline			✓	71.6%	46.1%
AMMNet	✓		✓	74.0%	47.4%
AMMNet		✓	✓	74.4%	47.7%
AMMNet	✓	✓	✓	<b>75.6%</b>	<b>48.5%</b>

Table 3. Ablation studies of different components in AMMNet on the test set of NYU [29].

Methods	$\mathcal{L}_{(D,G)}$	$\mathcal{M}^{1st}$	$\mathcal{M}^{2nd}$	$\mathcal{M}^{3rd}$	SC-IoU	SSC-mIoU
AMMNet		✓			73.5%	47.1%
AMMNet	✗	✓		✓	74.0%	47.2%
AMMNet		✓	✓	✓	74.0%	47.4%
AMMNet		✓			74.7%	47.8%
AMMNet	✓	✓		✓	74.8%	48.2%
AMMNet		✓	✓	✓	<b>75.6%</b>	<b>48.5%</b>

Table 4. Performance comparison of progressively integrating cross-modal modulation modules in AMMNet on the test set of NYU [29], under settings with and without adversarial training.

simply unleashes the modality potentials, achieving even better results in a simpler manner. 2) **Results on NYU-CAD [11]**. Table 2 presents results on synthetic NYU-CAD [11] dataset. AMMNet demonstrates noticeable advantages, achieving a top SSC-mIoU of 67.2%. Clean CAD environments enable more effective feature learning, thus the benefits of cross-modal modulation become more noticeable, as it can fully exploit the representation power of RGB and TSDF without interference from noise.

**Qualitative Comparisons.** Figure 4 presents two challenging scene completion examples from the NYU [29] dataset to qualitatively compare AMMNet against other methods. AMMNet excels in reconstructing semantically accurate parts, correctly predicting the elongated table and mirror highlighted in red boxes, which other methods fail to capture. Equally importantly, AMMNet completes geometrically plausible structures behind walls, generating empty spaces rather than clutter highlighted in dotted red boxes. Such completion aligns better with natural layouts. Together, these examples showcase AMMNet’s strengths - its ability to produce semantically and geometrically faithful scene reconstructions, exceeding prior arts.

#### 5.4. Ablation Study

To analyze the contribution of each proposed component, we conduct ablation studies on the NYU [29] dataset by adding them incrementally to the baseline model. Following [34], the baseline only contains the basic encoder-decoder architecture without cross-modal modulation ( $\mathcal{M}$ ) and adversarial training ( $\mathcal{L}_{(D,G)}$ ). As presented in Table 3, it achieves 46.1% SSC-mIoU and 71.6% SC-IoU.

**Effectiveness of  $\mathcal{M}$ .** When integrating the proposed  $\mathcal{M}$

Methods	FeatFusion	GradUpdating	SC-IoU	SSC-mIoU
AMMNet*	✓		73.2%	46.5%
AMMNet*		✓	74.1%	46.9%
AMMNet*	✓	✓	<b>73.5%</b>	<b>47.1%</b>

Table 5. Ablation study of two benefits in modulation module  $\mathcal{M}^{1st}$  on the test set of NYU [29], under a simplified AMMNet\* setting where the latter two modulation modules are removed.

Methods	Geometry	Semantic	Discriminator	SC-IoU	SSC-mIoU
AMMNet <sup>†</sup>			✓	74.0%	46.7%
AMMNet <sup>†</sup>	✓		✓	<b>74.4%</b>	46.6%
AMMNet <sup>†</sup>		✓	✓	74.0%	47.1%
AMMNet <sup>†</sup>	✓	✓	✓	<b>74.4%</b>	<b>47.7%</b>

Table 6. Ablation study of the GT perturbations in adversarial training on the test set of NYU [29], under a simplified AMMNet<sup>†</sup> setting where all the modulation modules are removed.

into the baseline model, the performance is improved by 2.4% in SSC-mIoU and 1.3% in SC-IoU, as presented in the 2nd row of Table 3. Meanwhile, regarding the adversarial trained model in the 3rd row as a strong baseline, integrating  $\mathcal{M}$  can still bring an improvement of 1.2% SC-IoU and 0.8% SSC-mIoU. This demonstrates the importance of  $\mathcal{M}$  for enabling superior semantic scene understanding even over a strong baseline.

**Effectiveness of  $\mathcal{L}_{(D,G)}$ .** Similarly, adding adversarial training to the baseline improves the SSC-mIoU by 1.6% and SC-IoU by 2.8% as shown in the 3rd row of Table 3. Without adversarial training in our AMMNet, we observe a performance drop of 1.1% in SSC-mIoU and 1.6% in SC-IoU as in the 2nd row. This proves the regularization effect of adversarial training in preventing overfitting. Furthermore, as depicted in Figure 2(b), the performance curves over training epochs validate that incorporating adversarial training leads to steadily increasing SSC-mIoU and SC-IoU on both train and test sets. This demonstrates that the performance gains of adversarial training stem from enabling stable optimization with dynamic competition in gradient updating, instead of potential overfitting to the training data.

In summary, the ablation studies clearly validate the efficacy of the key components of AMMNet. The combination of both leads to the best performance with an overall improvement of 2.4% in SSC-mIoU and 4.0% in SC-IoU.

#### 5.5. Discussions

**Analysis of Progressive Modulation.** To further analyze the impact of cross-modal modulation, we conduct experiments by progressively adding modulation components. For simplicity, we denote the three modulation modules in Figure 3 from left to right as  $\mathcal{M}^{1st}$ ,  $\mathcal{M}^{2nd}$  and  $\mathcal{M}^{3rd}$  respectively. As shown in Table 4, compared with the base-

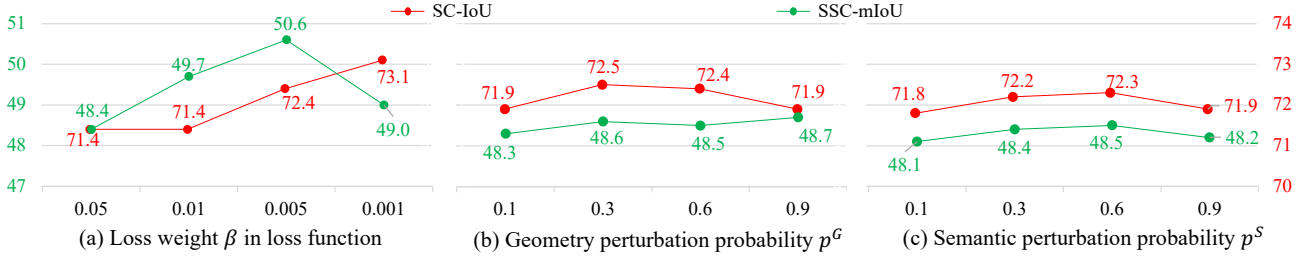


Figure 5. Sensitivity analysis of key hyperparameters based on the random split validation set from NYU [29].

line model, incorporating  $\mathcal{M}^{1st}$  brings substantial gains of 1.0% SSC-mIoU and 1.9% SC-IoU, validating the benefit of enriching uni-modal features with cross-modal contexts. However, further adding  $\mathcal{M}^{2nd}$  and  $\mathcal{M}^{3rd}$  only leads to minor improvements over using  $\mathcal{M}^{1st}$  alone. We hypothesize this is because the model capacity is still limited by overfitting. With the proposed adversarial training scheme, progressively integrating more modulation modules consistently improves the performance, as presented in the lower part of Table 4.

In summary, with proper regularization, the combination of  $\mathcal{M}^{1st}$ ,  $\mathcal{M}^{2nd}$  and  $\mathcal{M}^{3rd}$  leads to the best results by sufficiently modulating both shallow and deep representations. **Analysis of Modulation Benefits.** As discussed in Sec. 4, compared to addition, modulation provides two benefits: 1) inter-dependent gradient updating across modalities, and 2) enhanced forward feature fusion. To analyze their contribution, we simplify AMMNet to only have the first modulation  $\mathcal{M}^{1st}$ , denoted as AMMNet\*. Based on AMMNet\*, we compare two variants: 1) replacing  $\mathcal{M}^{1st}$  with addition but keeping inter-dependent gradient updating; 2) detaching  $\mathcal{M}^{1st}$  gradient while retaining forward feature calibration. As shown in Table 5, detaching inter-dependent gradient causes an SSC-mIoU drop of 0.6%, while replacing it with addition leads to a smaller 0.2% decrease. This suggests inter-dependent gradient updating provides more contribution, while forward calibration also complements the improvements.

**Analysis of Adversarial Training.** To thoroughly analyze the gains from adversarial training, we built an ablated AMMNet, denoted as AMMNet<sup>†</sup>, by removing all cross-modal modulation modules. As shown in Table 6, compared to the baseline model, utilizing the developed discriminator alone (1st row) significantly improves geometric completeness, increasing SC-IoU by 2.4%. Adding fake samples that disrupt geometric completeness (2nd row) further enhances geometric perception, achieving an additional 0.4% SC-IoU gain. However, without guidance on semantic validity, it achieves slightly inferior semantic performance compared to the discriminator-alone setting, with a minor 0.1% SSC-mIoU drop. Introducing fake samples damaging semantic correctness (4th row) provides better semantic supervision,

leading to a noticeable improvement in semantic accuracy, increasing SSC-mIoU by 1.1% over the geometry-only setting. This verifies the complementary effects of the two perturbation strategies in improving scene completion.

**Analysis of Hyperparameter Sensitivity.** We analyze the sensitivity of the key hyperparameters based on the randomly split validation set: the loss weight  $\beta$  in the loss function, and the probabilities  $p^G$  and  $p^S$  of generating geometrically/semantically perturbed fake GT, where the subscript  $j$  in  $p_j^S$  is omitted since the same probability is used for perturbing each semantic category here. As Figure 5 (a) shows, an overly large loss weight for the adversarial training loss like 0.05 overwhelms the task-specific loss  $\mathcal{L}_{SSC}$ , harming the performance. AMMNet selects  $\beta = 0.005$  which achieves optimal results.

In the sensitivity analysis for the perturbation probabilities  $p^G$  and  $p^S$ , we build on the ablated AMMNet<sup>†</sup> with cross-modal modules removed. The experiments show performance remains stable across a wide range of fixed values for both  $p^G$  and  $p^S$ , with only minor variations observed. This indicates the robustness of the proposed adversarial training scheme across different perturbation levels. To avoid hyperparameter tuning burden, we randomly sample probabilities within [0.1, 0.9] for both  $p^G$  and  $p^S$ , avoiding too little noise below 0.1 or too much above 0.9.

## 6. Conclusions

In this work, we identify two limitations of existing RGB-D based semantic scene completion methods: ineffective feature learning and overfitting problems. To address these issues, we propose a new deep learning framework AMMNet. The core techniques include cross-modal modulation to better exploit single-modal representations, and adversarial training to prevent overfitting to the training data. Extensive experiments on NYU and NYUCAD datasets demonstrate AMMNet’s state-of-the-art performance.

## Acknowledgements

This research was supported by the National Natural Science Foundation of China under Grant 61925204.



## References

- [1] Mohammad Mahdi Bejani and Mehdi Ghatee. A systematic review on overfitting control in shallow and deep neural networks. *Artificial Intelligence Review*, pages 1–48, 2021. [3](#)
- [2] Yingjie Cai, Xuesong Chen, Chao Zhang, Kwan-Yee Lin, Xiaogang Wang, and Hongsheng Li. Semantic scene completion via integrating instances and scene in-the-loop. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 324–333, 2021. [5](#), [6](#)
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. [5](#), [6](#)
- [4] Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, and Hongsheng Li. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4193–4202, 2020. [1](#), [3](#), [4](#), [5](#), [6](#)
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. [6](#)
- [6] Haotian Dong, Enhui Ma, Lubo Wang, Miaohui Wang, Wuyuan Xie, Qing Guo, Ping Li, Lingyu Liang, Kairui Yang, and Di Lin. Cvsformer: Cross-view synthesis transformer for semantic scene completion. In *IEEE International Conference on Computer Vision (ICCV)*, pages 8874–8883, 2023. [5](#), [6](#)
- [7] Aloisio Dourado, Teofilo E De Campos, Hansung Kim, and Adrian Hilton. Edgenet: Semantic scene completion from a single rgb-d image. In *International conference on pattern recognition (ICPR)*, pages 503–510, 2021. [2](#)
- [8] Aloisio Dourado, Frederico Guth, and Teofilo de Campos. Data augmented 3d semantic scene completion with 2d segmentation priors. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3781–3790, 2022. [6](#)
- [9] Chenzhuang Du, Tingle Li, Yichen Liu, Zixin Wen, Tianyu Hua, Yue Wang, and Hang Zhao. Improving multimodal learning with uni-modal teachers. *arXiv preprint arXiv:2106.11059*, 2021. [3](#)
- [10] Dajun Du, Kang Li, and Minrui Fei. A fast multi-output rbf neural network construction method. *Neurocomputing*, 73(10-12):2196–2202, 2010. [3](#)
- [11] Michael Firman, Oisín Mac Aodha, Simon Julier, and Gabriel J Brostow. Structured prediction of unobserved voxels from a single depth image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5431–5440, 2016. [2](#), [5](#), [6](#), [7](#)
- [12] Mudasar A Ganaie, Minghui Hu, AK Malik, M Tanveer, and PN Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115:105151, 2022. [3](#)
- [13] Martin Garbade, Yueh-Tung Chen, Johann Sawatzky, and Juergen Gall. Two stream 3d semantic scene completion. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 0–0, 2019. [5](#)
- [14] Christian Garbin, Xingquan Zhu, and Oge Marques. Dropout vs. batch normalization: an empirical study of their impact to deep learning. *Multimedia Tools and Applications*, 79:12777–12815, 2020. [3](#)
- [15] Yuxiao Guo and Xin Tong. View-volume network for semantic scene completion from a single depth image. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 726–732, 2018. [3](#)
- [16] Di Hu, Xuelong Li, et al. Temporal multimodal learning in audiovisual speech recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3574–3582, 2016. [3](#)
- [17] Jie Li, Yu Liu, Dong Gong, Qinfeng Shi, Xia Yuan, Chunxia Zhao, and Ian Reid. Rgb-d based dimensional decomposition residual network for 3d semantic scene completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7693–7702, 2019. [1](#), [3](#), [5](#), [6](#)
- [18] Jie Li, Kai Han, Peng Wang, Yu Liu, and Xia Yuan. Anisotropic convolutional networks for 3d semantic scene completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3351–3359, 2020. [3](#), [5](#), [6](#)
- [19] Jie Li, Qi Song, Xiaohu Yan, Yongquan Chen, and Rui Huang. From front to rear: 3d semantic scene completion through planar convolution and attention-based network. *IEEE Transactions on Multimedia*, 25:8294–8307, 2023. [5](#), [6](#)
- [20] Shice Liu, Yu Hu, Yiming Zeng, Qiankun Tang, Beibei Jin, Yinhe Han, and Xiaowei Li. See and think: Disentangling semantic scene completion. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 263–274, 2018. [1](#), [3](#)
- [21] Yu Liu, Jie Li, Qingsen Yan, Xia Yuan, Chunxia Zhao, Ian Reid, and Cesar Cadena. 3d gated recurrent fusion for semantic scene completion. *arXiv preprint arXiv:2002.07269*, 2020. [1](#)
- [22] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*, 2016. [6](#)
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2018. [6](#)
- [24] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5058–5066, 2017. [3](#)
- [25] Xiaokang Peng, Yake Wei, Andong Deng, Dong Wang, and Di Hu. Balanced multimodal learning via on-the-fly gradient modulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8238–8247, 2022. [3](#)
- [26] Dhanesh Ramachandram and Graham W Taylor. Deep multimodal learning: A survey on recent advances and trends. *IEEE signal processing magazine*, 34(6):96–108, 2017. [3](#)
- [27] Christoph B Rist, David Emmerichs, Markus Enzweiler, and Dariu M Gavrilă. Semantic scene completion using local deep implicit functions on lidar data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7205–7218, 2021. [3](#)

- [28] Connor Shorten, Taghi M Khoshgoftaar, and Borko Furht. Text data augmentation for deep learning. *Journal of big Data*, 8:1–34, 2021. 3
- [29] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision (ECCV)*, pages 746–760, 2012. 1, 2, 5, 6, 7, 8
- [30] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1746–1754, 2017. 2, 5, 6
- [31] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. 3
- [32] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016. 5
- [33] Jiayang Tang, Xiaokang Chen, Jingbo Wang, and Gang Zeng. Not all voxels are equal: Semantic scene completion from the point-voxel perspective. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 2352–2360, 2022. 3
- [34] Fengyun Wang, Dong Zhang, Hanwang Zhang, Jinhui Tang, and Qianru Sun. Semantic scene completion with cleaner self. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 867–877, 2023. 1, 2, 3, 4, 5, 6, 7
- [35] Xuzhi Wang, Di Lin, and Liang Wan. Ffnet: Frequency fusion network for semantic scene completion. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 2550–2557, 2022. 5, 6
- [36] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 12077–12090, 2021. 6
- [37] Pingping Zhang, Wei Liu, Yinjie Lei, Huchuan Lu, and Xiaoyun Yang. Cascaded context pyramid for full-resolution 3d semantic scene completion. In *IEEE International Conference on Computer Vision (ICCV)*, pages 7801–7810, 2019. 5, 6
- [38] Min Zhong and Gang Zeng. Semantic point completion network for 3d semantic scene completion. In *European Conference on Artificial Intelligence (ECAI)*, pages 2824–2831, 2020. 2
- [39] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016. 3