

Y²Seq2Seq: Cross-Modal Representation Learning for 3D Shape and Text by Joint Reconstruction and Prediction of View and Word Sequences

Zhizhong Han,^{1,2} Mingyang Shang,¹ Xiyang Wang,¹ Yu-Shen Liu,^{1*} Matthias Zwicker²

¹School of Software, Tsinghua University, Beijing, China

²Department of Computer Science, University of Maryland, College Park, USA

h312h@umd.edu, smy16@mails.tsinghua.edu.cn, wangxiya16@mails.tsinghua.edu.cn, liuyushen@tsinghua.edu.cn, zwicker@cs.umd.edu

Abstract

Jointly learning representations of 3D shapes and text is crucial to support tasks such as cross-modal retrieval or shape captioning. A recent method employs 3D voxels to represent 3D shapes, but this limits the approach to low resolutions due to the computational cost caused by the cubic complexity of 3D voxels. Hence the method suffers from a lack of detailed geometry. To resolve this issue, we propose Y²Seq2Seq, a view-based model, to learn cross-modal representations by joint reconstruction and prediction of view and word sequences. Specifically, the network architecture of Y²Seq2Seq bridges the semantic meaning embedded in the two modalities by two coupled “Y” like sequence-to-sequence (Seq2Seq) structures. In addition, our novel hierarchical constraints further increase the discriminability of the cross-modal representations by employing more detailed discriminative information. Experimental results on cross-modal retrieval and 3D shape captioning show that Y²Seq2Seq outperforms the state-of-the-art methods.

Introduction

With the development of 3D modeling and scanning techniques, more and more 3D shapes become available on the Internet with detailed physical properties, such as texture, color, and material. With large 3D datasets, however, shape class labels are becoming too coarse of a tool to help people efficiently find what they want, and visually browsing through shape classes is cumbersome. To alleviate this issue, an intuitive approach is to allow users to describe the desired 3D object using a text description. Jointly understanding 3D shape and text by learning a cross-modal representation, however, is still a challenge because it requires an efficient 3D shape representation that can capture highly detailed 3D shape structures.

To overcome this challenge, a 3D-Text cross-modal dataset was recently released in (Chen et al. 2018), where a combined multimodal association model was also proposed to capture the many-to-many relations between 3D voxels and text descriptions. This model employed a shape encoder to compute the embeddings of 3D shapes directly from 3D voxels. However, this strategy is limited to learning from low

resolution voxel representations due to the computational cost caused by the cubic complexity of 3D voxels. This leads to low discriminability of learned cross-modal representations due to a lack of detailed geometry information.

We resolve this issue by proposing to learn cross-modal representations of 3D shape and text from view sequences and word sequences, where each 3D shape is represented by a view sequence. Our deep learning model captures the correlation between 3D shape and text by the joint reconstruction and prediction of view and word sequences from each modality. This strategy aims to get part-level information as much as possible to understand mainly part-related descriptions, which remedies a lack of part prior knowledge in 3D-Text understanding. We call our model Y²Seq2Seq, since it is implemented by two coupled “Y” like sequence-to-sequence (Seq2Seq) structures. In addition, we employ novel hierarchical constraints that can be easily added to extract more detailed discriminative information from different aspects. Specifically, Y²Seq2Seq consists of a 3D shape branch and a text branch. Starting from the encoding of either 3D shape modality or text modality, each branch jointly reconstructs the modality itself and predicts the counterpart modality using one RNN encoder and two RNN decoders. The two branches are coupled by sharing parameters in decoders for reconstruction and prediction, which is effective to bridge the semantic meaning embedded in 3D shape modality and text modality. In addition, we propose to employ the discriminative information at the class level, instance pair level, and instance level as constraints in the training procedure. Our significant contributions are list below.

- We propose a deep learning model called Y²Seq2Seq, which enables to learn cross-modal representations of 3D shape and text from view sequences and word sequences.
- Our novel coupled “Y” like Seq2Seq structures have a powerful capability to bridge the semantic meaning of two sequence-represented modalities by joint reconstruction and prediction.
- Our results demonstrate that our novel hierarchical constraints can further increase the discriminability of learned cross-modal representations by employing more detailed discriminative information.

*Corresponding author: Yu-Shen Liu

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

dicator P^W .

E^V is implemented by a VGG19 (Simonyan and Zisserman 2014) and a RNN. The VGG19 first extracts low-level feature \mathbf{f}_i of each view v_i in the view sequence \mathbf{v} from the last 4096 dimensional fully connected layer. Then, E^V learns the feature \mathbf{F}_s of shape s by aggregating all the N view features, where \mathbf{F}_s is the hidden state at the N -th step of the RNN.

Starting from \mathbf{F}_s , R^V and P^W , which are also implemented by a RNN, jointly reconstruct the view features in sequence \mathbf{v} and predict the word sequence \mathbf{w} . The reconstructed view features \mathbf{f}'_i and predicted word sequence \mathbf{w}' are expected to be close to the ground truth view features \mathbf{f}_i and word sequence \mathbf{w} . Thus, the loss of branch S is formed by the loss L_{V2V} from view to view and the loss L_{V2W} from view to word, as defined below,

$$\begin{aligned} L_S &= \alpha L_{V2V} + \beta L_{V2W}, \\ L_{V2V} &= \frac{1}{N} \sum_{i \in [1, N]} \|\mathbf{f}'_i - \mathbf{f}_i\|_2^2, \\ L_{V2W} &= - \sum_{j \in [1, M]} \log p(w_j | w_{<j}, \mathbf{v}), \end{aligned} \quad (1)$$

where w_j is the j -th word in the word sequence \mathbf{w} , $w_{<j}$ represents the words in front of w_j , $p(w_j | w_{<j}, \mathbf{v})$ is the probability of correctly predicting j -th word according to the previous words $w_{<j}$ and the view sequence \mathbf{v} , and α and β are two balance weights.

Text branch T. T is formed by a description encoder E^W , a description reconstructor R^W and a multi-view predictor P^V . Note that T is coupled with S by sharing weights in R^W with description predictor P^W in S and sharing weights in P^V with multi-view reconstructor R^V in S.

E^W , R^W and P^V are all implemented by a RNN. E^W learns the feature \mathbf{F}_t of the description t of shape s by aggregating the embedding e_j of words w_j in sequence \mathbf{w} . The word embedding e_j is learned with the other parameters in Y²Seq2Seq together. The feature \mathbf{F}_t is represented by the hidden state at the M -th step of E^W .

Starting from \mathbf{F}_t , R^W and P^V jointly reconstruct the word sequence \mathbf{w} and predict the view features in sequence \mathbf{v} . The reconstructed word sequence \mathbf{w}'' and predicted view features \mathbf{f}''_i are expected to be close to the ground truth word sequence \mathbf{w} and view features \mathbf{f}_i . Thus, the loss of branch T is formed by the loss L_{W2W} from word to word and the loss L_{W2V} from word to view, as defined below,

$$\begin{aligned} L_T &= \gamma L_{W2W} + \delta L_{W2V}, \\ L_{W2W} &= - \sum_{j \in [1, M]} \log p(w_j | w_{<j}, \mathbf{w}), \\ L_{W2V} &= \frac{1}{N} \sum_{i \in [1, N]} \|\mathbf{f}''_i - \mathbf{f}_i\|_2^2, \end{aligned} \quad (2)$$

where w_j is the j -th word in the word sequence \mathbf{w} , $w_{<j}$ represents the words in front of w_j , $p(w_j | w_{<j}, \mathbf{w})$ is the probability of correctly reconstructing j -th word according to the

previous words $w_{<j}$ and the ground truth word sequence \mathbf{w} , and γ and δ are balance weights.

Hierarchical constraints. Y²Seq2Seq employs three constraints on \mathbf{F}_s and \mathbf{F}_t to further increase the discriminability of the learned cross-modal representations by employing more detailed discriminative information. These constraints provide hierarchically discriminative information at the class level, the instance pair level, and the instance level.

For the instance pair (s, t) from shape class c , the first constraint makes each instance in the pair be correctly classified using a softmax classifier. Here, we use only one softmax classifier to simultaneously classify \mathbf{F}_s and \mathbf{F}_t in both modalities. This helps push the two modalities together in the cross-modal representation space. The estimated shape class c' should be the ground truth shape class c . Thus, the classification loss is the negative log likelihood as below,

$$L_{C1} = -\log p(c' = c | \mathbf{F}_s) - \log p(c' = c | \mathbf{F}_t), \quad (3)$$

where $p(c' = c | \mathbf{F}_s)$ and $p(c' = c | \mathbf{F}_t)$ are probabilities of correctly classifying 3D shape s and text t , respectively.

We use a triplet constraint (Schroff, Kalenichenko, and Philbin 2015) to leverage the relationships between instance pairs, such as (s^+, t^+) and (s^-, t^-) . The instance pair level constraint ensures the feature of s^+ is closer to the feature of t^+ than the feature of any other s^- or t^- . Thus, our triplet loss is defined as below,

$$\begin{aligned} L_{C2} &= [\|\mathbf{F}_{s^+} - \mathbf{F}_{t^+}\|_2^2 + \|\mathbf{F}_{s^+} - \mathbf{F}_{t^-}\|_2^2 + \mu]_+ \\ &+ [\|\mathbf{F}_{t^+} - \mathbf{F}_{s^+}\|_2^2 + \|\mathbf{F}_{t^+} - \mathbf{F}_{s^-}\|_2^2 + \mu]_+, \end{aligned} \quad (4)$$

where $[\bullet]_+$ is a $\max(0, \bullet)$ function, (s^+, t^+, t^-) and (t^+, s^+, s^-) are two triplets in the training set, μ is a margin that is enforced between instances within pairs and instances across pairs. The two triplets are formed by finding a (s^-, t^-) in another shape class for each (s^+, t^+) .

Finally, we include a joint embedding constraint for discriminative information at the instance level. Although instance level information has been involved in the triplet constraint, two instances s and t in the same pair (s, t) are merely required to be as close as possible there. Our joint embedding constraint further toughly pushes the two features \mathbf{F}_s and \mathbf{F}_t of the instances s and t to the same point in the cross-modal representation space. Thus, the joint embedding loss is defined as below,

$$L_{C3} = \|\mathbf{F}_s - \mathbf{F}_t\|_2^2. \quad (5)$$

Finally, the loss L_C defined for our hierarchical constraints is shown below, where ϕ , φ and ψ are balance weights for the hierarchical constraints at different levels.

$$L_C = \phi L_{C1} + \varphi L_{C2} + \psi L_{C3}, \quad (6)$$

Objective function. Y²Seq2Seq is trained to minimize all the above losses. Thus, the objective function is defined as,

$$\min L_S + L_T + L_C. \quad (7)$$

We learn our learning targets, the cross-modal representations of 3D shape and text, \mathbf{F}_s and \mathbf{F}_t , by iteratively updating

Table 1: The balance weights employed in all the three datasets.

Set	α	β	γ	δ	ϕ	φ	ψ
Prim	1	2	0.001	1	0.0001	$0.1(\mu = 1)$	0.1
Shap	1	1	0.01	1	0.001	$0.01(\mu = 1.5)$	0.1

Table 2: Effect of coupled ‘‘Y’’ like Seq2Seq under primitive subset.

	Metrics	Rec	Pre	R+P	C-S	C-T	C-Y
S2T	RR@1	1.47	1.33	1.33	69.87	73.73	80.13
	RR@5	1.73	2.67	4.27	70.40	81.60	82.53
	NDCG@5	1.44	1.37	1.05	69.87	67.92	80.16
T2S	RR@1	2.27	2.06	1.85	46.92	72.57	92.45
	RR@5	4.70	7.58	3.34	63.37	87.73	95.99
	NDCG@5	1.76	3.01	0.97	41.79	70.21	88.52

them using back propagated gradients with a learning rate ε , i.e., $\mathbf{F}_s \leftarrow \mathbf{F}_s - \varepsilon \times \partial L / \partial \mathbf{F}_s$ and $\mathbf{F}_t \leftarrow \mathbf{F}_t - \varepsilon \times \partial L / \partial \mathbf{F}_t$.

Experiments and analysis

Dataset. We evaluate Y²Seq2Seq under the 3D-Text cross-modal dataset by (Chen et al. 2018), which consists of a primitive subset and a ShapeNet subset. Specifically, the primitive subset contains 7,560 shapes and 191,850 descriptions. The ShapeNet subset contains 15,038 shapes and 75,344 descriptions. We employ the same training/test splitting under both subsets as (Chen et al. 2018).

Experimental setup. We first explore how each element of Y²Seq2Seq affects its performance under the primitive subset and the ShapeNet subset in 3D-Text cross-modal retrieval. Then, we evaluate the performance of Y²Seq2Seq by comparing it with the state-of-the-art methods in cross-modal retrieval and 3D shape captioning under the two subsets, respectively. In all the experiments, the dimension of the cross-modal representations \mathbf{F}_s and \mathbf{F}_t is 128, the dimension of the word embedding e_j is set to 512, and both branches in Y²Seq2Seq are implemented using GRU cells (Cho et al. 2014). All the experiments are conducted with a learning rate ε of 0.00001.

According to the descriptions involved in each subset, we employ dataset-dependent vocabulary. Under the primitive subset, we extract 77 unique words to form the vocabulary. Under the ShapeNet subset, the employed vocabulary is formed by 3587 unique words. All the RNNs for descriptions are dynamic, with varying length according to the length of the descriptions.

Table 3: Effect of coupled ‘‘Y’’ like Seq2Seq under ShapeNet subset.

	Metrics	Rec	Pre	R+P	C-S	C-T	C-Y
S2T	RR@1	0.07	0.07	0.13	1.61	1.74	1.88
	RR@5	0.34	0.34	0.34	6.03	6.17	7.51
	NDCG@5	0.07	0.07	0.08	1.44	1.42	1.65
T2S	RR@1	0.13	0.11	0.07	0.42	0.77	1.04
	RR@5	0.34	0.32	0.35	1.20	3.26	4.25
	NDCG@5	0.24	0.21	0.20	0.82	1.98	2.62

Table 4: Effect of hierarchical constraints under primitive subset.

	Metrics	No	+ L_{C1}	+ $L_{C1} + L_{C2}$	+ L_C
S2T	RR@1	80.13	83.07	88.53	94.13
	RR@5	82.53	85.73	88.80	94.13
	NDCG@5	80.16	82.43	88.33	94.10
T2S	RR@1	92.45	93.20	95.99	96.66
	RR@5	95.99	97.50	97.53	97.57
	NDCG@5	88.52	89.36	95.52	95.87

Under each subset, according to the order of magnitude of each loss in L , we set both α in L_S and δ in L_T to 1, since the losses of reconstructing the view sequence in S and predicting the view sequence in T are similar. Subsequently, all the other balance weights are set based on them. The balance weights used in both subsets are shown in Table 1, where the margin μ involved in the triplet loss is also presented.

In each cross-modal retrieval experiment below, we show the results in two directions to comprehensively evaluate the performance of Y²Seq2Seq. One direction uses 3D shapes as queries and aims to retrieve instances from the set of descriptions. We call this direction shape-to-text, which is abbreviated as ‘‘S2T’’. In contrast, the other direction uses descriptions as queries and aims to retrieve instances from the set of 3D shapes. Accordingly, we call this direction text-to-shape and abbreviate it as ‘‘T2S’’. In addition, the recall rate (RR@ k) (Chen et al. 2018) and NDCG (Jrvelin and Kekäläinen 2002) are used as the metrics. In addition, all the results shown in both the ‘‘T2S’’ and ‘‘S2T’’ directions in the same column in the following tables are obtained with the same trained Y²Seq2Seq.

In each 3D shape captioning experiment below, METEOR (Denkowski and Lavie 2014), ROUGE (Lin 2004), CIDEr (Vedantam, Zitnick, and Parikh 2015), and BLUE (Papineni et al. 2002) are used as the metrics to evaluate the quality of generated descriptions according to the ground truth descriptions, where these metrics are abbreviated as ‘‘M’’, ‘‘R’’, ‘‘C’’, and ‘‘B-1’’, ‘‘B-2’’, ‘‘B-3’’, ‘‘B-4’’, respectively in the following tables.

Effect of the coupled ‘‘Y’’ like Seq2Seq. First, we conduct experiments to show how the coupled ‘‘Y’’ like Seq2Seq structures contribute to the cross-modal representation learning by joint reconstruction and prediction of view and word sequences.

As shown in Table 2 and Table 3, we conduct six experiments under the two subsets, respectively, where RR@1, RR@5 and NDCG@5 are used as metrics in all the experiments. Under each subset, all the six experiments are conducted without hierarchical constraints. The first two experiments are conducted with only reconstruction within the same modality and only with prediction across the modalities, which are indicated as ‘‘Rec’’ and ‘‘Pre’’ in the tables. In these two experiments, the two ‘‘Y’’ like branches are de-generated into line like branches. The third experiment is conducted with both reconstruction and prediction, but the two branches are not coupled, where the decoders of the two branches are not sharing parameters, as indicated as ‘‘R+P’’ in the tables. The fourth and fifth experiments are conducted

Table 5: Effect of hierarchical constraints under ShapeNet subset.

	Metrics	No	$+L_{C1}$	$+L_{C1} + L_{C2}$	$+L_C$
S2T	RR@1	1.88	2.82	3.42	6.77
	RR@5	7.51	10.19	10.59	19.30
	NDCG@5	1.65	2.40	2.59	5.30
T2S	RR@1	1.04	1.83	1.92	2.93
	RR@5	4.25	6.36	6.89	9.23
	NDCG@5	2.62	4.07	4.40	6.05

Table 6: Effect of voxel resolution under ShapeNet subset.

	Metrics	32^3	64^3	128^3
S2T	RR@1	6.77	7.31	7.64
	RR@5	19.30	19.97	20.64
	NDCG@5	5.30	5.43	5.48
T2S	RR@1	2.93	2.37	2.70
	RR@5	9.23	8.81	9.82
	NDCG@5	6.05	5.61	6.27

with coupled branches but without joint reconstruction and prediction, i.e., the decoders of both branches regard either the view sequence or the word sequence as the target, as denoted as “C-S” or “C-T”, respectively. The last experiment is conducted with coupled “R+P”, as indicated as “C-Y” in the tables. Note that we use the same balance weights in Table 1 in all the four experiments.

Both the results of “Rec” and “Pre” in Table 2 and Table 3 show that merely reconstruction within the same modality or prediction across the modalities in the two branches can not learn satisfactory cross-modal representations. This is because the two branches can not bridge the semantic meaning embedded in different modalities. The “R+P” results in the two tables still suffer from the same issue. Even with increasing the interactions across modalities, joint reconstruction and prediction in each branch can only slightly increase the discriminability of the learned cross-modal representations in some metrics. With the two branches coupled, only using the features from one modality as the mapping target of the decoders achieves much better results than the results without coupling, as shown by “C-S” or “C-T”. This means the coupled branches are able to help bridge the semantic meaning embedded in different modalities by pushing the modalities together better. To explore how to learn more satisfactory cross-modal representations, we propose to jointly perform reconstruction and prediction by the coupled two branches through sharing the parameters in the decoders of both branches. As shown by the “C-Y” results, the outperforming of other results demonstrates that the reconstruction, the prediction and the coupling can all contribute to the performance, and both the reconstruction and prediction can only contribute based on the coupling.

Effect of the hierarchical constraints. Next, we explore how each hierarchical constraint contributes to the cross-modal representation learning under the two subsets.

As shown in Table 4 and Table 5, we conduct three experiments under the primitive subset and ShapeNet subset,

Table 7: The comparison in cross-modal retrieval under primitive subset.

	Methods	RR@1	RR@5	NDCG@5
S2T	ML	24.67	29.87	24.38
	DS	80.50	85.87	80.36
	MiViSE	17.87	24.13	16.44
	SLR	1.20	2.80	1.15
	LBAT	5.20	6.13	5.25
	LBAM	89.20	90.53	89.48
	FTST	92.00	92.40	91.98
	FMM	93.47	93.47	93.47
	Our	94.13	94.13	94.10
T2S	ML	25.93	57.24	25.00
	DS	81.77	90.70	81.29
	MiViSE	8.21	15.42	6.84
	SLR	4.08	9.49	2.31
	LBAT	5.06	15.29	5.92
	LBAM	91.13	98.27	91.90
	FTST	94.24	97.55	95.20
	FMM	95.07	99.08	95.51
	Our	96.66	97.57	95.87

respectively, where RR@1, RR@5 and NDCG@5 are used as metrics in all the experiments. Under each subset, each hierarchical constraint is incrementally added starting from “Coupled “Y”” in Table 2 or Table 3, which are indicated as “No” in the three experiments. These experiments are conducted with only the classification constraint, indicated as “ $+L_{C1}$ ”, with the classification and triplet constraints, indicated as “ $+L_{C1} + L_{C2}$ ”, and with all hierarchical constraints, indicated as “ $+L_C$ ”, respectively. Note that we use the balance weights in Table 1 in all the three experiments.

Under both subsets, all the metrics are increasing along with adding an additional hierarchical constraint. The increasing results in terms of all metrics show that each hierarchical constraint is able to contribute to the performance of $Y^2Seq2Seq$. However, the degree of contribution provided by each constraint is slightly different, depending on the dataset. For example, according to the degree of increase, the three constraints contribute almost equally under the primitive subset, while the joint embedding constraint contributes more than the others under the ShapeNet subset.

Effect of voxel resolution. In the experiments above, the involved 3D shapes are represented by view sequences rendered from voxel representations of the 3D shapes with a resolution of 32^3 . However, $Y^2Seq2Seq$ can learn from higher-resolution voxel representations without increasing the memory requirements of the neural networks because it performs view-based deep learning. This resolves the memory issue of methods that perform learning directly from the voxel representations. We explore the effect of voxel resolution on the performance of $Y^2Seq2Seq$ in cross-modal retrieval under ShapeNet subset. We use views captured from 3D shapes represented by voxel grids with a resolution of 32^3 , 64^3 and 128^3 .

The grids with higher resolution can provide more geometry detail on the shape surface, which results in more dis-

Table 8: The comparison in cross-modal retrieval under ShapeNet subset.

	Methods	RR@1	RR@5	NDCG@5
S2T	ML	0.13	0.47	0.11
	DS	0.13	0.60	0.13
	MiViSE	0.20	0.40	0.10
	SLR	0.27	0.40	0.11
	LBAT	0.20	0.80	0.12
	LBAM	0.07	0.34	0.07
	FTST	0.94	3.69	0.85
	FMM	0.83	3.37	0.73
	Our	6.77	19.30	5.30
T2S	ML	0.13	0.61	0.36
	DS	0.12	0.65	0.38
	MiViSE	0.11	0.31	0.20
	SLR	0.11	0.38	0.24
	LBAT	0.04	0.20	0.12
	LBAM	0.08	0.34	0.21
	FTST	0.22	1.63	0.87
	FMM	0.40	2.37	1.35
	Our	2.93	9.23	6.05

Table 9: The comparison in 3D shape captioning under primitive subset.

Model	M	R	C	B-1	B-2	B-3	B-4
SLR-N	0.18	0.44	0.13	0.42	0.31	0.21	0.15
MiV-N	0.35	0.67	0.53	0.66	0.53	0.45	0.39
S2VT	0.47	0.87	0.96	0.88	0.82	0.75	0.70
Our-N	0.70	0.98	1.37	0.98	0.97	0.96	0.96
Our	0.54	0.92	1.21	0.92	0.88	0.84	0.80

crimative 3D shape features. This is verified by the continuously increasing results of “S2T” in Table 6. However, because of the limited information embedded in the description, voxels with higher resolution could not keep contributing to more discriminative text features, as shown by the results of “T2S”, although we have got the best “T2S” results with the highest resolution of 128^3 .

Cross-modal retrieval. To evaluate the performance of $Y^2Seq2Seq$, we compare our method with the state-of-the-art methods under the primitive subset and the ShapeNet subset in cross-modal retrieval. To conduct a fair comparison, our results are obtained from voxel grids of 32^3 resolution under the primitive subset and the ShapeNet subset, which is the same as in the other methods.

Under the two subsets, we compare our method with metric learning (ML) (Song et al. 2016), deep symmetric structured joint embedding (DS) (Reed et al. 2016), MiViSE (Song and Soleymani 2018), SLR (Shen et al. 2018), association learning with only TST round trips (LBAT) (Häusser, Mordvintsev, and Cremers 2017), and several methods from (Chen et al. 2018), such as LBA-MM (LBAM), Full-TST (FTST), and Full-MM (FMM). Among all these methods, MiViSE, SLR and our method are view-based, and the rest are voxel-based.

As shown in Table 7 and Table 8, $Y^2Seq2Seq$ outperforms all the other methods in terms of all metrics. Particu-

Table 10: The comparison in 3D shape captioning under ShapeNet subset.

Model	M	R	C	B-1	B-2	B-3	B-4
SLR-N	0.11	0.24	0.05	0.40	0.17	0.08	0.04
MiV-N	0.16	0.36	0.14	0.61	0.35	0.21	0.12
S2VT	0.21	0.45	0.27	0.67	0.43	0.26	0.15
Our1-N	0.22	0.41	0.29	0.57	0.34	0.22	0.17
Our1	0.29	0.56	0.71	0.80	0.65	0.53	0.46
Our2-N	0.22	0.41	0.30	0.57	0.34	0.23	0.18
Our2	0.30	0.56	0.72	0.80	0.65	0.54	0.46
Our3-N	0.22	0.41	0.31	0.58	0.35	0.24	0.19
Our3	0.29	0.55	0.70	0.80	0.64	0.52	0.44

larly, our results significantly outperform other methods under the ShapeNet subset. For example, our result is about 6 times better in “S2T” retrieval and about 4 times better in “T2S” retrieval than the state-of-the-art FMM. We believe these experimental improvements come from our novel way of bridging the semantic meaning embedded in the cross-modal sequences in the process of joint reconstruction and prediction by the two coupled branches.

Furthermore, we visualize the “S2T” and the “T2S” retrieval results under the primitive subset and the ShapeNet subset in Fig. 2 (a) and Fig. 2 (b) respectively, where the Top-5 retrieved items with their distances to the query are shown. The retrieved 3D shapes or descriptions highly match their description queries or 3D shape queries. The retrieved results show that our learned discriminative representations can jointly represent 3D shapes and descriptions in the same cross-modal space. In addition, they are able to distinguish the subtle difference between shapes and descriptions. For example, the retrieved shapes in “T2S” retrieval under ShapeNet subset are all matching the description query, but there are still some style and appearance differences among them. Similarly, the retrieved descriptions in “S2T” retrieval under the primitive subset are all matching the 3D shape query, but the emphasized characteristics of the query 3D shape are still different.

3D shape captioning. To further evaluate the performance of $Y^2Seq2Seq$, we compare our method with the state-of-the-art methods under the primitive subset and the ShapeNet subset in 3D shape captioning. For the methods which can directly generate descriptions for 3D shapes, we will directly evaluate the generated descriptions according to the ground truth descriptions. Otherwise, we will use the feature of a 3D shape to retrieve the nearest description of the 3D shape as the generated description, where we use the suffix “-N” to denote these results. Note that the captioning results of different methods with either retrieval or generation are produced by the same trained parameter in cross-modal retrieval experiments.

Under the primitive subset and the ShapeNet subset, we compare our method with SLR (Shen et al. 2018), MiV (Song and Soleymani 2018), and S2VT (Venugopalan et al. 2015). We show our results obtained by voxel grids with different resolutions, such as 32^3 , 64^3 and 128^3 under the ShapeNet subset, where these three sets of results are denoted as “Our1”, “Our2” and “Our3” respectively in Ta-





	Query	Retrieved	Query	Retrieved
S2T		(3.331241) @1-The large pyramid is tall narrow yellow. (3.334018) @2-The large pyramid is high narrow yellow. (3.334379) @3-The large pyramid is tall thin yellow. (3.336639) @4-The pyramidal shape is large tall narrow yellow. (3.337065) @5-The large pyramid is high thin yellow.		(3.29) @1-It is a large blue ping pong table. It has a net that can be placed in the center. (3.37) @2-A blue colored table tennis table which has a metal frame. (3.48) @3-A blue frame table with peach top and a glass partition in the center of the table along its length with blue frame work around it. (3.53) @4-A blue flat surfaced unit standing on four short legs, of blue translucent plastic material, molded in one solid piece, having no separate parts. (3.55) @5-This is a blue ping ball table with a blue glass top and brown legs.
T2S	A large short wide green sphere		A wooden and metal chair that is brown and grey in color.	

Figure 2: The demonstration of “S2T” and “T2S” retrieval results under (a) primitive subset and (b) ShapeNet subset. The top 5 retrieved results with their distances to the query are shown.













	The cylindrical shape is large teal.		A rectangular table with a green top and brown leg table.
	The cylindrical shape is large teal.		A table made of brown wood that has a colorful top made of green patterns.
	The green rectangular shape is squat.		A brown, wooden chair with a back rest and a black cushion.
	The green rectangular shape is squat wide.		Brown colored, wooden, rest chair. Four legs and L shaped seat.
	The conical shape is nude colored.		A gray chair with a white cushion and a gray seat, two seater space with back rest and arm rest.
	The large conical shape is nude-colored.		Rectangular shaped wooden chair with sponge grey, black and brown in color, two seater space with back rest and arms rest.

Figure 3: The demonstration of generated descriptions under (a) primitive subset, (b) ShapeNet subset, where the ground truth descriptions are shown in green (second one of each pair of descriptions).

ble 10. We use the description predictor P^W in the 3D shape branch S to generate a description of a 3D shape. In addition, we also show our results obtained by nearest neighbor retrieval.

From the results shown, we can see that all our results obtained by either retrieval or generation under the primitive subset are the best. In addition, our results with different voxel resolutions obtained by either retrieval or generation under the ShapeNet subset are also the best among all competitors.

In addition, although all our results obtained by generation are good, they are always no better than the ones by retrieval under the primitive subset, while this phenomenon is reversed under the ShapeNet subset, as shown by the results with different voxel resolutions. This is because the primitive subset only contains simple primitive shapes that are described by simple sentences with similar patterns. $Y^2Seq2Seq$ obtains good cross-modal retrieval results, as shown in Table 7, which makes the retrieved descriptions almost the same as the ground truth. In contrast, the ShapeNet subset contains complex shapes with complex descriptions. It is not easy to retrieve the right description. However, $Y^2Seq2Seq$ can still generate some words for the key characteristics of the shape, such as color and material, which leads to better results with generation.

Moreover, we find the voxel resolution only slightly increases the results with both retrieval and generation. This

phenomenon is similar to the one in the “T2S” retrieval in Table 6, where higher resolution mainly contributes to “S2T” retrieval but only provides a small contribution to “T2S” retrieval.

Finally, we also visualize the description generated by $Y^2Seq2Seq$ under the primitive subset and the ShapeNet subset in Fig. 3. Compared to the ground truth, the 3D shape branch S of $Y^2Seq2Seq$ can generate high quality description for the characteristics embedded in the view sequence according to the knowledge learned from joint reconstruction and prediction in the two branches. $Y^2Seq2Seq$ generates almost the same descriptions as the ground truth under the primitive subset, while it generates descriptions with similar semantic meanings to the ground truth under the ShapeNet subset. The generated descriptions successfully represent the class, shape, color, and materials of the 3D shapes, which indicates that $Y^2Seq2Seq$ can jointly understand 3D shapes and text quite well.

Conclusions

To resolve the issue of cubic complexity of 3D voxel grids for 3D-Text cross-modal representation learning, we propose $Y^2Seq2Seq$ as a view-based model to learn cross-modal representations by joint reconstruction and prediction of view and word sequences. The novel coupled “Y”-like $Seq2Seq$ structures effectively bridge the semantic meaning embedded in the 3D shape and text modalities by simultaneously reconstructing within the same modality and predicting across modalities. In addition, more detailed discriminative information can be successfully employed to further increase the discriminability of learned cross-modal representations by our novel hierarchical constraints. Our experimental results show that $Y^2Seq2Seq$ is able to learn more discriminative cross-modal representations from high-resolution voxels for 3D-Text cross-modal retrieval and 3D shape captioning compared to the competing state-of-the-art techniques.

Acknowledgments

Yu-Shen Liu is the corresponding author. This work was supported by National Key R&D Program of China

(2018YFB0505400), the National Natural Science Foundation of China (61472202), and Swiss National Science Foundation grant (169151). We thank all anonymous reviewers for their constructive comments.

References

- Chen, K.; Choy, C. B.; Savva, M.; Chang, A. X.; Funkhouser, T.; and Savarese, S. 2018. Text2shape: Generating shapes from natural language by learning joint embeddings. *arXiv preprint arXiv:1803.08495*.
- Cho, K.; van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *SSST@EMNLP*, 103–111.
- Denkowski, M., and Lavie, A. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, 376–380.
- Feng, F.; Wang, X.; and Li, R. 2014. Cross-modal retrieval with correspondence autoencoder. In *Proceedings of ACM International Conference on Multimedia*, 7–16.
- Häusser, P.; Mordvintsev, A.; and Cremers, D. 2017. Learning by association - A versatile semi-supervised training method for neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 626–635.
- Järvelin, K., and Kekäläinen, J. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems* 20(4):422–446.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceeding of ACL workshop on Text Summarization Branches Out*.
- Ngiam, J.; Khosla, A.; Kim, M.; Nam, J.; Lee, H.; and Ng, A. Y. 2011. Multimodal deep learning. In *Proceedings of International Conference on International Conference on Machine Learning*, 689–696.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: A method for automatic evaluation of machine translation. In *Annual Meeting on Association for Computational Linguistics*, 311–318.
- Reed, S. E.; Akata, Z.; Lee, H.; and Schiele, B. 2016. Learning deep representations of fine-grained visual descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 49–58.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.
- Shen, X.; Tian, X.; Xing, J.; Rui, Y.; and Tao, D. 2018. Sequence-to-sequence learning via shared latent representation. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.
- Song, Y., and Soleymani, M. 2018. Cross-modal retrieval with implicit concept association. *CoRR* abs/1804.04318.
- Song, H. O.; Xiang, Y.; Jegelka, S.; and Savarese, S. 2016. Deep metric learning via lifted structured feature embedding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 4004–4012.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 4566–4575.
- Venugopalan, S.; Rohrbach, M.; Donahue, J.; Mooney, R.; Darrell, T.; and Saenko, K. 2015. Sequence to sequence – video to text. In *Proceedings of the IEEE International Conference on Computer Vision*, 4534–4542.
- Wang, B.; Ma, L.; Zhang, W.; and Liu, W. 2018. Reconstruction network for video captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*.