

This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Sequential View Synthesis with Transformer

Phong Nguyen-Ha¹[0000-0002-9678-0886]</sup>, Lam Huynh¹[0000-0002-8311-1288]</sup>, Esa Rahtu²[0000-0001-8767-0864]</sup>, and Janne Heikkilä¹[0000-0003-0073-0866]</sup>

¹ Center for Machine Vision and Signal Analysis, University of Oulu, Finland {phong.nguyen,lam.huynh,janne.heikkila}@oulu.fi ² Computer Vision Group, Tampere University, Finland esa.rahtu@tuni.fi

Abstract. This paper addresses the problem of novel view synthesis by means of neural rendering, where we are interested in predicting the novel view at an arbitrary camera pose based on a given set of input images from other viewpoints. Using the known query pose and input poses, we create an ordered set of observations that leads to the target view. Thus, the problem of single novel view synthesis is reformulated as a sequential view prediction task. In this paper, the proposed Transformer-based Generative Query Network (T-GQN) extends the neural-rendering methods by adding two new concepts. First, we use multi-view attention learning between context images to obtain multiple implicit scene representations. Second, we introduce a sequential rendering decoder to predict an image sequence, including the target view, based on the learned representations. Finally, we evaluate our model on various challenging datasets and demonstrate that our model not only gives consistent predictions but also doesn't require any retraining for finetuning.

Keywords: Sequential view synthesis · Transformer · Multi-view attention

1 Introduction

View synthesis aims to create novel views of an object or a scene from a perspective of a virtual camera based on a set of reference images. It has been an active field of research already for several decades in computer vision and computer graphics due to its various application areas including free-viewpoint television, virtual and augmented reality, and telepresence 12345.

Conventionally, the view synthesis problem has been addressed by using image-based or geometry-based approaches 1. In pure image-based rendering, the novel view is warped from a densely sampled set of reference images or the light field without exploiting any geometric information, which obviously requires a large amount of image data and limits the new viewpoints to a relatively small range. In geometry-based rendering, the novel view is generated using a 3D model that has been first created from the reference views using multi-view stereo or some other image-based modeling techniques. This allows for larger baselines between the views, but also sets high requirements to the quality of

the 3D model. Between these two extremes, free viewpoint depth-image-based rendering (DIBR) uses depth maps associated to the reference views enabling 3D image warping to synthesize the novel view [3]. In practice, all these approaches tend to produce notable artifacts due to missing or inaccurate data, which reduced the quality of the rendered image.



Fig. 1. Overview of our proposed Transformer-based Generative Query Network (T-GQN). Given a set of input images and camera poses $\{(I_n, v_n)\}$, where $n = 1, \ldots, N$, we obtain a set of N scene representations by passing them one-by-one through the encoder (E) network. We use a stack of transformer encoder (T) to perform the multi-view attention learning between these representation. Finally, we use each these attention representation r_n^* to sequentially render the output sequence using the decoder (D) network. The dashed arrow indicates that we initialize states of the decoder network using the computed states of the previous step.

Recently, researchers have adopted deep learning techniques to overcome the inherent limitation of the conventional approaches. This paper focuses on a family of neural rendering methods **678910** that infer the underlying 3D scene structure and faithfully produces the target view even at a distant query pose. These methods use an aggregate function to represent the entire 3D scene as a single implicit representation. Although they manage to successfully render the target view, a large amount of data is required to train, which in turn, takes a long time to converge. We argue that because these methods focus on synthesizing only a single target image, they are inefficient when rendering target images for distant query poses that are subject to strong geometric transformations and occlusions with respect to the reference images.

In this paper, we introduce a Transformer-based Generative Query Network (T-GQN) to address the problem of novel view synthesis in a sequential manner. We train an end-to-end model that sequentially renders a positionally ordered set of nearby-views and then predicts the target view at the final rendering step (as shown in Fig. 1). We claim that if the model is able to render the nearby-views accurately, it will be also capable of predicting the target view correctly.

Since we do not have the nearby-views of the target pose, we train our model to render the input views. Moreover, instead of rendering the viewpoints using a single implicit scene representation, we use multi-view attention learning based on Transformer Encoder Π to learn multiple scene representations. At each rendering step, we modify the decoder network of Generative Query Network (GQN) [6] to share its states through rendering steps. To summarize, our key contributions are as follows:

- 1. We reformulate the problem of single view synthesis into sequential view synthesis.
- 2. Our proposed T-GQN introduces two novel concepts: multi-view attention learning via the Transformer Encoder and a sequential rendering decoder. Our model extends the previously proposed GQN by sequentially rendering a pose-ordered set of novel views.
- 3. We demonstrate that our proposed framework achieves state-of-the-art performance on challenging view synthesis datasets. Moreover, our method reaches convergence faster and requires less computational resource to train.

The source code and the model are available at https://github.com/phongnhhn92/ TransformerGQN .

2 Related Work

The literature related to view synthesis is extensive, and to limit the scope, we focus on a few deep learning based methods in this section that are most relevant to our method. We suggest a state-of-the-art report for an extensive review 12.

Many early solutions use regression to derive the pixel colors of the target view directly from the input images [13]1415[16]17]. In [15], Tatarchenko et al. maps an image of a scene to an RGB-D image from an unknown viewpoint with an auto-encoder architecture and train their model using supervised learning. Instead of synthesizing pixels from scratch, other works explores using CNNs to predict appearance flow [16]. Later work by Sun et al. [17] presents a self-confidence aggregation mechanism to integrate both predicted appearance flow and pixel hallucination to achieve contractually consistent results. Although, those method manage to render plausible novel views, their results are limited to a scene with a single object or a slight change in viewpoints in the autonomous driving situation.

A great amount of effort has been dedicated to incorporate geometric information to the model. For example, 181920212223 apply deep learning techniques to leverage geometry cues and learn to predict the novel view. The deep learning based-light field camera view interpolation 2024 use a deep network to predict depth separately for every novel view. Another line of work [1819]21] cleverly extract a Multiplane image representation of the scene. This representation offers regularization that allows for an impressive stereo baseline extrapolation. Recently, Choi et al. [22] use deep neural network to estimate a depth probability volume, rather than just a single depth value for each pixel of the novel view. Even though, they show promising results but they are limited to synthesizing a middle view among source images or a magnified view from a single input. In contrast, our proposed framework focuses on arbitrary target views and is able to learn from source images that vary in length.

Recent progresses in geometric deep learning proposes to represent a scene as a voxel-grid to enforce the 3D structure. These methods [25]26]27] use 3D convolution layers to learn 3D spatial transformations from the input views to the novel view and then apply GAN training [28] to enhance quality of the output image. Recently, Sitzman et al. [29] presents a continuous, 3D-structureaware scene representation that encodes both geometry and appearance. Their work learns a mapping from world coordinates to feature representation of local scene properties. Another contribution of this paper is that authors empirically demonstrated that their method is able to show generalization across scenes for classes of single objects using an additional MLP network such as HyperNetwork [30]. In our paper, we evaluate our method on the novel view synthesis dataset which contains 2 millions scenes each (multiple combinations of objects types, colors, lightning positions). We argue that training their method to generalize to a large number of instances would be computationally expensive. In fact, this is an open research problem that has not yet been solved.

Recent **neural rendering** methods have introduced a generative model that understands the underlying 3D scene structure and faithfully produces the target view at the distant query pose [31]32]33]34]. Generative Query Network (GQN) [6] and its variant [7]8]9 are incorporating all input observation (images and poses) into a single implicit 3D scene representation to generate the target view. This aggregated representation contains all necessary information (e.g. object identities, positions, colors, scene layout) to make accurate image predictions. Moreover, these methods corresponds to a special case of Neural Processes [35]36]. In this paper, we argue that generating the target view using such compact representation leads to poor predictions and slow training convergence.

3 Approach

In this section, we first provide the reader with a brief background of the Generative Query Network (GQN). Then we convert the problem of **single view synthesis** to the problem of **sequential view synthesis** by introducing our Transformer-based Generative Query Network (T-GQN) that extends the current GQN architecture with two novel building blocks: multi-view attention learning via Transformer Encoder and sequential rendering decoder.

3.1 Generative Query Network

Given the observations that include N images $I_n \in I$ and their corresponding camera poses $v_n \in V$, GQN [6] solves the **single view synthesis** problem by

using a encoder-decoder neural network to predict the target image I_q at an arbitrary query pose v_q .

First, the encoder is a feed-forward neural network that takes N observations as input and produce a single implicit scene representation $R = \sum_{n=1}^{N} r_n$ by performing a element-wise sum of N encoded scene representation r_n . The decoder then takes R and v_q as an input and predicts the new view I'_q from that viewpoint. The decoder network is a conditional latent variable model DRAW [37]38] which includes M pairs of Generation and Inference convolutional LSTM networks. At each generation step, the hidden state of the Generation and Inference LSTM core is utilized to approximate the prior π and posterior distribution q. Since the target view I_q is fed into the Inference sub-network, minimizing the Kullback-Leibler (KL) distance between π and q would help the Generation subnetwork to produce an accurate result. Both the encoder and decoder networks are trained jointly to minimize the ELBO loss \mathcal{L}_{GQN} function:

$$\mathcal{L}_{GQN} = \left[-\ln \mathcal{N}(I_q | I'_q) + \sum_{m=1}^M KL \Big[\mathcal{N}(q_m) || \mathcal{N}(\pi_m) \Big] \right]$$
(1)

In the next section, we will describe how we reformulate the problem of **single view synthesis** to the problem of **sequential view synthesis** to address this issue.

3.2 Sequential view synthesis

As can be seen from Fig. 2 (a), GQN 6 predicts the target view I_q in a single rendering step. If the query pose v_q is distant from all input poses then the target view might look completely different than all input views. In this case, minimizing the above \mathcal{L}_{GQN} loss does not guarantee to generate a plausible target view and it might take a long training time to reach the convergence.

We argue that if the model is able to predict an input view I'_n for n > 1 based on previous input data $\{(I_1, v_1), \ldots, (I_{n-1}, v_{n-1})\}$ then it also renders the target view I_q at the query pose v_q provided that the camera poses $\{v_1, \ldots, v_N, v_q\}$ have been organized as a sequence where the adjacent poses are the closest ones. To achieve such ability, we train our proposed T-GQN model using multiple rendering steps. Each rendering step of our model is identical with GQN except that we use different sets of input observations and query poses. In Fig. 2 (b) we illustrate these sets of input observations at each rendering step with boxes of different colors.

For example, in the first rendering step, we only allow the model to use the information inside the green box which includes the input view I_1 and its camera pose v_1 . Then, our T-GQN model is trained to predict the next input view I'_2 at the viewpoint v_2 . Preventing the model to look at the input view I_2 would encourage the model to reason about the given contexts and produce a global implicit scene representation that extrapolates beyond the given information. In the final rendering step, both GQN and our proposed T-GQN model are allowed to see all N input observations to render the target view but our model can



Fig. 2. An illustration of the single view synthesis (a) compared to our proposed sequential view synthesis (b). Previously proposed GQN [6] directly predicts the target view in a single rendering step, while our proposed T-GQN sequentially predicts novel views from an ordered set of N observations. The output set includes N - 1 nearest views of the target view and the target view itself. Colored boxes indicate which input observations are encoded at each rendering step.

leverage the past experiences from previous rendering steps to have a better approximation I'_q of the target image I_q . Furthermore, training our model in multiple rendering steps enforces our model to make consistent predictions of novel views at different viewpoints in a forward pass. This helps to stabilize the training process by making the network to produce deterministic results and also be able to reach the convergence faster than the former architecture. Note that our only use sequential view synthesis during training. In the testing time, we use randomly ordered sets of context views to have a fair comparison with other methods.

Therefore, the problem of single view synthesis can be redefined as the problem of sequential view synthesis by predicting a sequence of N novel views $S_{out} = \{I'_2, ..., I'_N, I'_q\}$ from a sequence of observations $S_{in} = \{(I_1, v_1), ..., (I_N, v_N)\}$. Since we are having N different target views then it would be beneficial to have N different scene representations. In order to have such multiple implicit scene representations, we use the Transformer Encoder [11] to learn the dependencies between input observations at each rendering step.

3.3 Multi-view attention learning via Transformer Encoder

Recent works in the Language Modeling task **11**,394041 use the self attentionbased Transformer Encoder to effectively learn dependencies between word em-

7



Fig. 3. Illustration of multi-view attention learning using a stack of Transformer Encoder (T) blocks. For the visualization purposes, we show how a single Transformer Encoder block is able to produce multiple implicit scene representations r^* using the input view representations r.

beddings in a sentence. Transformer Encoder takes a set of word embeddings as an input and produces another set of enhanced word embedding representations. Each of these representations reflects the long-range dependencies between input embeddings and they have been proven to be useful for training various language modeling tasks. Therefore, we can also take a set of scene representations r as an input and use the Transformer Encoder to produce another set of enhanced scene representations r^* that are trained to exploit the multi-view dependencies.

Fig. \exists shows an example of how we use a stack of Transformer Encoder blocks to perform the multi-view attention learning and implicitly represent a scene using a set of multiple representations. Within a Transformer Encoder block, the most important component is the Multi-Head Attention layer using the selfattention mechanism \blacksquare . The self-attention function produces an $N \times N$ matrix A of multi-view attention scores so that each row of the matrix A reflects the learned multi-view dependencies at each rendering step. In addition, we apply an attention mask m to the multi-view attention scores A. This attention mask allows us to control which part of the input representation sequence r we would like the model to ignore when computing the attention scores. We leave the implementation details of Transformer Encoder to the supplementary material. In practice, we found that applying this attention mask leads to better performance on datasets that have high similarities between viewpoints. When the input viewpoints are not overlapping, we found that better results are achieved without applying the attention mask.

If the movement of the camera is restricted then input viewpoints overlap with each other. Therefore, masking the attention scores would mean that the model is trained to predict the novel view using only a subset of the input sequence. Fig. 4(a) shows the multi-view attention scores of a camera-restricted example. In the 1st rendering step, we try to render the 2nd input view (green box) as the novel view. Since our model is only allowed to use the 1st view (yellow box) as the input, the predicted image only contains the cyan object. However,

8 P. Nguyen et al.



Fig. 4. Visualization of multi-view attention scores at each rendering step produced by our methods when (a) the camera movement is restricted and (b) the camera is free to move. Note that the order of the input sequence does not affect the learned multi-view attention scores.

our model is able to render the missing orange object by putting higher scores on the 2^{nd} input view in the next rendering step. Since the 1^{st} input view does not have enough information to render the target view, our model learns to give it less attention and produce higher attention scores on another two input views.

When the camera movement is not restricted, the input sequence might contain views which do not necessary include information to render the target view. In this case, we allow the model to use all input observations by not applying the attention mask. In Fig. 4 (b), our model is able to give high attention scores to input observations which are the ground-truth views in the first two rendering steps. Therefore, the first two predicted novel views are identical to the groundtruth because the model has already seen these novel views as inputs. However, our goal is to be able to predict the target view (red box) which has a distant query pose. Among three input observations, the 2^{nd} input is the only view



Fig. 5. Illustration of our proposed sequential rendering decoder. At each rendering step, we use the computed states from the previous rendering step as the initialization for the decoder network. The canvas state U_n^G is utilized to render the novel image I'_n .

which contains information about objects on the scene. As can be seen in the last rendering steps, our model manages to pay most attention on the 2^{nd} input view (green box) and gives low scores to the other two views and successfully predicts the target view despite the distant query pose.

3.4 Sequential rendering decoder

In this paper, we improve the decoder network of GQN [6] by applying a recurrent mechanism between rendering steps. In order to transfer the knowledge between rendering steps, computed LSTM states from previous steps are utilized as initialization for the current step as shown in the Fig. [5] Here, we simplify the notation by defining that C_n^G is the cell state of the m^{th} Generation core at the n^{th} rendering step. We apply this kind of notation to all cell states C, hidden states H and canvas states U of the decoder network. At the n^{th} rendering step, the novel image I'_n is computed as follows:

$$I'_{n} = \operatorname{decoder}(r_{n}^{*}, v_{n}, C_{n-1}^{G}, H_{n-1}^{G}, U_{n-1}^{G}, C_{n-1}^{I}, H_{n-1}^{I})$$

$$(2)$$

Since our model has multiple rendering steps, the prior and the posterior terms, π_n^m and q_n^m , are obtained from the m^{th} LSTM core at the n^{th} rendering step. We train our model using the modified ELBO loss function as follows:

$$\mathcal{L}_{TGQN} = \left[-\sum_{n=1}^{N} \ln \mathcal{N}(I_n | I'_n) + \beta \sum_{n=1}^{N} \sum_{m=1}^{M} \mathrm{KL}\left[\mathcal{N}(q_n^m) | | \mathcal{N}(\pi_n^m)\right]\right]$$
(3)

By adding the β coefficient to the KL term, we emphasize discovering the disentangled latent factors [42]. This technique has proven to be effective to maximize the probability of generating the desired output and keeping the distance between the prior and posterior distribution small [43].

4 Experiments

4.1 Experimental setup

As illustrated in the Fig. 2 to create a sequence of adjacent poses, we could, for example, use the overlap between the viewing frustums of the cameras to

Table 1. Quantitative comparison of results between our T-GQN model, originalGQN6 and E-GQN10.

Model	# parameters (millions)	L1 (pixels)			L2 (pixels)		
		RRC	RFC	RRO	RRC	RFC	RRO
GQN (8 LSTM layers)	381	8.50 ± 7.01	14.23 ± 13.56	14.12 ± 8.19	18.62 ± 12.86	30.28 ± 25.69	21.79 ± 10.23
GQN (12 LSTM layers)	428	7.40 ± 6.22	12.44 ± 12.89	10.12 ± 5.15	14.62 ± 12.77	26.80 ± 21.35	19.63 ± 9.14
E-GQN	N/A	3.59 ± 2.10	12.05 ± 11.79	6.59 ± 3.23	6.80 ± 5.23	27.65 ± 20.72	12.08 ± 6.52
T-GQN without mask	382	3.30 ± 2.05	$\textbf{9.25} \pm \textbf{9.15}$	6.31 ± 2.59	6.65 ± 4.29	12.72 ± 10.56	11.72 ± 5.23
T-GQN with mask	382	$\textbf{2.31} \pm \textbf{1.89}$	11.65 ± 10.25	$\textbf{5.28} \pm \textbf{2.31}$	$\textbf{5.92} \pm \textbf{2.69}$	15.44 ± 13.37	$11.65~\pm~5.13$

measure their adjacency. However, in our experiments we simply reordered all input observations based on the Euclidean distance between the translation vector of the query pose and the input poses that turned out to be sufficient to demonstrate the efficiency of our method. We choose to train our T-GQN model using three input observations to render a target view at an arbitrary query pose. Further details about our architecture are presented in the supplementary material.

4.2 Comparing with state-of-the-art methods

To evaluate our model, we compare against the GQN [6] and an improved variant of E-GQN [10]. We use two datasets: Rooms-Ring-Camera (RRC) and Rooms-Free-Camera (RFC), from [6] and one dataset: Rooms-Random-Objects (RRO) from [10]. The RRC dataset contains 2 millions rendered 3D square rooms that are composed of random objects of various shapes, colors and locations. Moreover, the scene textures, walls and lights are also randomly generated. In this dataset, the camera only moves on a fixed ring and always faces the center of the room. In case of the RFC dataset, the environment is the same with RRC except for the freely moving camera and objects rotated around their vertical axes. We also evaluate our method on the RRO dataset which includes complex 3D scenes with realistic 3D objects from the ShapeNet dataset [44].

Recent works [31]32 on view synthesis often require retraining for any testing scene. Although, they produce visually impressive novel views but they are not able to generalize to unseen data. Therefore, we decide to compare our method against Scene Representation Network (SRN) [29] which is the current top-performing technique for view synthesis which generalizes reasonably well on testing data. Due to computational constraints, we can not compare this method with the full RRC or RFC dataset. We instead use a small subset of 1000 scenes from the Shapard-Metzler-7-Parts (SM7) and 3500 scenes from the RRC dataset to train both our proposed model and SRN model. In this experiment, the ratio between the training set and testing set is 9:1.





Fig. 6. Example of rendered novel views using our proposed T-GQN (with and without masking) and other methods. Overall, our proposed T-GQN model is able to produce better target view than previous methods. More examples can be found in the Appendix section. We urge the reader to zoom in for better visualization.

4.3 Results

Comparing against $GQN[\underline{6}]$ and E- $GQN[\underline{10}]$: In our experiment, we have tested our model to sample several novel views given the same input views from the same scene. In the Table[] we have reported the L1 and L2 distances between the rendered novel views and the ground-truth images across testing scenes. Since our testing models are probabilistic, we run our experiment several times and report the average and standard deviation based on the test runs. In each testing scene, we sample 10 independent viewpoints and use 3 of them as the input to our model to generate the rest of viewpoints as novel images.

Qualitative results shown in Fig. 6 demonstrate that the T-GQN model with or without the attention mask is able to render all novel images much sharper than the former architectures. Meanwhile, predicted target views from GQN 6 are often blurry and not able to get the correct object types, colors and positions. As explained in the Section 3.2 our model is able to predict nearby views of the novel view. Therefore, we include step-by-step renderings in the supplementary material.

As can be seen from Table 1 training our T-GQN model with the attention mask leads to a performance gain on the RRC and RRO datasets. In the case of the RFC dataset, our method without the attention mask performs significantly better than other methods including ours with the attention mask. In the freecamera case, the query poses are often far from the input poses so there may be not enough information from input views to generate the novel view. This explains the small gap of performance between our T-GQN model with masking on the RRC dataset and our T-GQN model without masking on the RFC dataset. In our paper, we apply a zero attention mask in testing the RFC dataset so that the model has access to all possible information. Our model then learns to attend to the relevant input views via our proposed multi-view attention learning described in Section 3.3 This highlights the strength of the Transformer



Fig. 7. Qualitative results between our proposed T-GQN and SRN on (a) SM7 dataset and (b) RRC dataset.

Table 2. Quantitative comparison of results between SRN, our proposed T-GQN modeltested on the small subset of SM7 and RRC dataset.

Method	L1 (pixels)		L2 (pixels)		
	SM7	RRC	SM7	RRC	
SRN T-GQN 2	2.15 2.14 ± 0.21	$3.56 \\ 1.92 \pm 0.15$	$5.65 \\ 5.25 \pm 0.34$	$6.21 \\ 2.96 \pm 0.93$	

architecture that learns to attend to the most relevant piece of information among input views.

Comparing against SRN[29]: We also benchmark novel view synthesis accuracy on few-shot reconstruction to compare our probabilistic T-GQN model with a deterministic view synthesis SRN [29]. As can be seen in Table. [2] and Fig. [7] our proposed method manages to achieve similar results with SRN model on the SM7 dataset. In case of RRC dataset, our model performs significantly better than SRN. We argue that SM7 is an easy dataset which contains only a single rotating colored object. Therefore, learning to generalize between training and testing data of SM7 dataset is easier than with the RRC dataset.

As the number of training scenes increase, training a deterministic approach such as SRN to have good generalization across a large number of testing scenes is difficult and still an open challenge. Qualitative results in Fig. 7 (b) shows that our T-GQN model is not only able to render correct room layouts and object properties but also produce significantly sharper results than the previously proposed SRN method with the demanding RRC dataset. We discuss more about the capabilities of rendering novel views between our method and SRN in the supplementary material.



Fig. 8. Example of generated novel views compared between our proposed T-GQN and variants on the RRC dataset.

Table 3. Quantitative results of T-GQN and its variants on the Rooms-Ring-Camera dataset.

Metric	GQN	SeqGQN	$\begin{array}{l} \text{T-GQN} \\ (\beta = 1) \end{array}$	$\begin{array}{l} \text{T-GQN} \\ (\beta = 250) \end{array}$
L1 (pixels)	7.40 ± 6.22	4.51 ± 2.35	3.7 ± 1.92	$\textbf{2.31} \pm \textbf{1.89}$
L2 (pixels)	14.62 ± 12.77	8.2 ± 4.27	7.1 ± 2.47	$\textbf{5.92} \pm \textbf{2.69}$
SSIM	0.85	0.882	0.905	0.92

5 Ablation Study

To investigate the effectiveness of the proposed modules, we use our sequential rendering decoder to generate novel views using the aggregated scene representation R from [6] and denote this model as SeqGQN. Instead of using different multi-view attention representations at each rendering step, we use a single representation R as an input to all ConvLSTM layers in our proposed sequential rendering decoder network. The qualitative and quantitative results including the rendered novel views using SeqGQN are shown in Fig. 8 and Table 3 We observe that our full model significantly outperforms both SeqGQN and GQN. Although both GQN and SeqGQN are using the same aggregated scene representation R, SeqGQN is able to produce better and more accurate target views than the baseline. This result demonstrates that approaching the neural rendering in the sequential manner leads to more accurate view synthesis. We also find that each learned representation is fully capable of rending the novel view. Therefore, we can use a single scene representation r_n^* to render the novel view in the testing time. Further explanations are included in the supplementary material.

Since we are adding a large β coefficient into the KL loss term (Eq. 3), we are encouraging the model to approximate the posterior distribution q_n^m close to the prior distribution π_n^m as much as possible. If the query pose is far away from given input observations then it is hard to get a good estimate for q_n^m . In [6], the model is prone to mistakes because the target view is rendered in a single step.



Fig. 9. Qualitative(a) and quantitative(b) results on 4 testing scenes of Tank and Temples dataset.

Using the sequential rendering decoder, we solve the problem by generating the nearby views of the target view before rendering the target view. As can be seen from Table 3 training our T-GQN model with $\beta = 1$ leads to a drop of the overall performance. However, this model is still able to outperform both GQN and SeqGQN.

To evaluate our method on the real data, we compare our method with the recently proposed NeRF [31] on the challenging Tanks and Temples dataset. We train our proposed method on 10 randomly selected training scenes and evaluate the performance on 4 unseen testing scenes. Note that our method haven't seen any testing images during training or require any retraining like NeRF. However, as can be seen in Fig. 9 our method outperforms NeRF both qualitatively and quantitatively. This may be explained by the fact that NeRF makes an assumption of inward-facing scenes while our method does not have such limitations. On the flip side, our proposed method is not bounded to this condition and can generate plausible novel views in a wide variety of 3D scenes. This highlights the effectiveness of our method that generalize reasonably well on unseen data.

6 Conclusions

In this paper, we presented a method to synthesize novel views in a sequential manner. Instead of directly rendering the target view, we train our model to predict a sequence of novel views in multiple rendering steps. Using the Transformer Encoder, our proposed multi-view attention learning is able to learn different implicit scene representations for each rendering step. The experimental results demonstrate that our model is able to render more accurate novel views using less training time.

In the future works, we will explore different implicit neural scene representations to further improve the quality of synthesized images. Another interesting approach would be learning to synthesize novel views that are useful for further tasks such as 3D scene reconstruction.

References

- CHANG, Y., WANG, G.P.: A review on image-based rendering. Virtual Reality & Intelligent Hardware 1 (2019) 39 – 54
- Tanimoto, M.: Ftv: Free-viewpoint television. Signal Processing: Image Communication 27 (2012) 555 – 570
- Smirnov, S., Battisti, F., Gotchev, A.P.: Layered approach for improving the quality of free-viewpoint depth-image-based rendering images. Journal of Electronic Imaging 28 (2019) 1 – 17
- 4. Joachimczak, M., Liu, J., Ando, H.: Real-time mixed-reality telepresence via 3d reconstruction with hololens and commodity depth sensors. In: Proceedings of the 19th ACM International Conference on Multimodal Interaction. ICMI '17, New York, NY, USA, Association for Computing Machinery (2017) 514–515
- 5. Kolkmeier, J., Harmsen, E., Giesselink, S., Reidsma, D., Theune, M., Heylen, D.: With a little help from a holographic friend: The openimpress mixed reality telepresence toolkit for remote collaboration systems. In: Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology. VRST '18, New York, NY, USA, Association for Computing Machinery (2018)
- Eslami, S.A., Rezende, D.J., Besse, F., Viola, F., Morcos, A.S., Garnelo, M., Ruderman, A., Rusu, A.A., Danihelka, I., Gregor, K., et al.: Neural scene representation and rendering. Science **360** (2018) 1204–1210
- Rosenbaum, D., Besse, F., Viola, F., Rezende, D.J., Eslami, S.: Learning models for visual 3d localization with implicit mapping. arXiv preprint arXiv:1807.03149 (2018)
- Kumar, A., Eslami, S., Rezende, D.J., Garnelo, M., Viola, F., Lockhart, E., Shanahan, M.: Consistent generative query networks. arXiv preprint arXiv:1807.02033 (2018)
- Nguyen-Ha, P., Huynh, L., Rahtu, E., Heikkilä, J.: Predicting novel views using generative adversarial query network. In: Scandinavian Conference on Image Analysis, Springer (2019) 16–27
- Tobin, J., Zaremba, W., Abbeel, P.: Geometry-aware neural rendering. In: Advances in Neural Information Processing Systems 32. Curran Associates, Inc. (2019)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: Advances in Neural Information Processing Systems 30. Curran Associates, Inc. (2017) 5998–6008
- Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J.M., Nießner, M., Pandey, R., Fanello, S.R., Wetzstein, G., Zhu, J., Theobalt, C., Agrawala, M., Shechtman, E., Goldman, D.B., Zollhöfer, M.: State of the art on neural rendering. CoRR abs/2004.03805 (2020)
- Olszewski, K., Tulyakov, S., Woodford, O., Li, H., Luo, L.: Transformable bottleneck networks. The IEEE International Conference on Computer Vision (ICCV) (2019)

- 16 P. Nguyen et al.
- Chen, X., Song, J., Hilliges, O.: Monocular neural image based rendering with continuous view control. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 4090–4100
- Tatarchenko, M., Dosovitskiy, A., Brox, T.: Multi-view 3d models from single images with a convolutional network. In: European Conference on Computer Vision (ECCV). (2016)
- Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.A.: View synthesis by appearance flow. CoRR abs/1605.03557 (2016)
- Sun, S.H., Huh, M., Liao, Y.H., Zhang, N., Lim, J.J.: Multi-view to novel view: Synthesizing novel views with self-learned confidence. In: European Conference on Computer Vision. (2018)
- Flynn, J., Neulander, I., Philbin, J., Snavely, N.: Deep stereo: Learning to predict new views from the world's imagery. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 5515–5524
- 19. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. In: SIGGRAPH. (2018)
- Kalantari, N.K., Wang, T.C., Ramamoorthi, R.: Learning-based view synthesis for light field cameras. ACM Trans. Graph. 35 (2016)
- Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019, Computer Vision Foundation / IEEE (2019) 175–184
- Choi, I., Gallo, O., Troccoli, A., Kim, M.H., Kautz, J.: Extreme view synthesis. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 7781–7790
- 23. Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020)
- Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. ACM Trans. Graph. 38 (2019)
- Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: Hologan: Unsupervised learning of 3d representations from natural images. In: The IEEE International Conference on Computer Vision (ICCV). (2019)
- Sitzmann, V., Thies, J., Heide, F., Nießner, M., Wetzstein, G., Zollhöfer, M.: Deepvoxels: Learning persistent 3d feature embeddings. In: Proc. Computer Vision and Pattern Recognition (CVPR), IEEE. (2019)
- Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. ACM Trans. Graph. 38 (2019)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q., eds.: Advances in Neural Information Processing Systems 27. Curran Associates, Inc. (2014) 2672–2680
- Sitzmann, V., Zollhoefer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. In: Advances in Neural Information Processing Systems 32. Curran Associates, Inc. (2019) 1121–1132
- Ha, D., Dai, A.M., Le, Q.V.: Hypernetworks. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, OpenReview.net (2017)

17

- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV. (2020)
- Li, Z., Xian, W., Davis, A., Snavely, N.: Crowdsampling the plenoptic function. In: Proc. European Conference on Computer Vision (ECCV). (2020)
- Dupont, E., Bautista, M.A., Colburn, A., Sankar, A., Guestrin, C., Susskind, J., Shan, Q.: Equivariant neural rendering (2020)
- 34. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision (2019)
- 35. Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y.W., Rezende, D., Eslami, S.M.A.: Conditional neural processes. In Dy, J., Krause, A., eds.: Proceedings of the 35th International Conference on Machine Learning. Volume 80 of Proceedings of Machine Learning Research., Stockholmsmässan, Stockholm Sweden, PMLR (2018) 1704–1713
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, S.M.A., Rosenbaum, D., Vinyals, O., Teh, Y.W.: Attentive neural processes. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019, OpenReview.net (2019)
- 37. Gregor, K., Danihelka, I., Graves, A., Rezende, D., Wierstra, D.: Draw: A recurrent neural network for image generation. In Bach, F., Blei, D., eds.: Proceedings of the 32nd International Conference on Machine Learning. Volume 37 of Proceedings of Machine Learning Research., Lille, France, PMLR (2015) 1462–1471
- Gregor, K., Besse, F., Jimenez Rezende, D., Danihelka, I., Wierstra, D.: Towards conceptual compression. In Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R., eds.: Advances in Neural Information Processing Systems 29. Curran Associates, Inc. (2016) 3549–3557
- 39. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, Minnesota, Association for Computational Linguistics (2019) 4171– 4186
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. OpenAI Blog 1 (2019) 9
- 41. Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., Salakhutdinov, R.: Transformer-XL: Attentive language models beyond a fixed-length context. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, Association for Computational Linguistics (2019) 2978–2988
- Kingma, D.P., Welling, M.: An introduction to variational autoencoders. Foundations and Trends in Machine Learning 12 (2019) 307–392
- 43. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M.M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. In: ICLR. (2017)
- 44. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository (2015)