Tensor-to-Image: Image-to-Image Translation with Vision Transformers

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

03-10-2021 / 09-10-2021

CITATION

Gunduc, Yigit (2021): Tensor-to-Image: Image-to-Image Translation with Vision Transformers. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.16727140.v1>

DOI

[10.36227/techrxiv.16727140.v1](https://doi.org/10.36227/techrxiv.16727140.v1)

Tensor-to-Image: Image-to-Image Translation with Vision Transformers

Yiğit Gündüç¹

¹ygunduc@gmail.com

Abstract

Transformers gain huge attention since they are first introduced and have a wide range of applications. Transformers start to take over all areas of deep learning and the Vision transformers paper also proved that they can be used for computer vision tasks. In this paper, we utilized a vision transformer-based custom-designed model, tensor-to-image, for the image to image translation. With the help of self-attention, our model was able to generalize and apply to different problems without a single modification.

1 Introduction

Deep convolutional neural networks have been the state of the art for most of the vision tasks. Image classification, image segmentation, object detection are some of the major applications of convolutional networks. Recently, image-to-image translation became the focus of attention since it has a wide range of applications like photo enhancement, object transfiguration, and semantic segmentation. Image-to-image translation can be applied where an image can be mapped into another one. The success of the Convolutional networks on the majority of vision tasks made them the go-to choice for image-to-image translation too. Most of the models used for image-to-image translation tasks are Autoencoders [1, 2] or U-Net [3] which are based on convolutional networks.

Despite their success, convolutional layers have difficulty modeling long-range relationships on a given input due to the nature of the convolution operation. New concepts, Attention, Transformers, patch, and patch embeddings widened the scope of deep learning and introduced alternative ways to solve problems. Convolution work like local attention with a sliding kernel over the image hence destroys the long-range relation. Transformers can be used to remedy this problem. Transformers use self-attention that is applied to the whole input rather than a portion of it. The model is aware of all of the input and can relate from regions far apart. This enables self-attention models to gain upper hand in this kind of task where seeing the big picture can help the model generate accurate results.

In this work, we have developed a self-attention-based image-to-image translation framework. For code, and examples please see <https://github.com/YigitGunduc/tensor-to-image>

2 Related work

A variety of methods with CNNs have been used for image-to-image translation. Most of the previous solutions [4, 5, 6] have employed on Autoencoders [1, 2] for image-to-image translation. In an Autoencoder, the input is passed through the encoder layers. Encoder downsamples the input until a bottleneck layer. At the bottleneck, the latent vector is generated. After the latent vector has been created, the process has been reversed by a decoder where it upsamples the latent vector until the desired image is created. The problem with these encoder-decoder architectures is that when inputs pass through the bottleneck most of the low-level information which might be helpful while constructing the output image is lost. Some successful methods have utilized U-Net [3] like approaches to overcome this problem. U-Net adds skip connections between layers. To formulate these skip connections between layers is i^{th} layer and the $(n-i)^{\text{th}}$ layer is directly connected, where n is the layer number. U-Net is especially successful because of the skip connections where it concatenates the encoder to the decoder. This allows the low-level information to be preserved and directly transferred to the output.

Transformers [7] were proposed for machine translation and became the state of the art in most of the NLP tasks. Generally, Transformer models are pre-trained on large text corpora and then fine-tuned for specific tasks [8, 9, 10]. On the contrary of the recurrent networks, transformers do not suffer from memory shortages. This is due to their unique self-attention mechanisms and the way that they utilize Multilayer perceptrons(MLPs). Recurrent layers cannot recall tokens after a certain amount of time. Some experiments show that the memory of the recurrent layers can be extended by a certain amount with the help of the attention mechanism but still nowhere near the transformers. There have been attempts to utilize transformers like self-attention-based architectures with CNNs for computer vision tasks [11, 12, 13, 18, 19]. Image GPT(IGPT) [13] is another notable transformer-based architecture that takes pixel sequences as inputs and tries to generate corresponding pixels. Unlike vision transformers, where they take all the images and split them into patches, IGPT treats pixels just like tokens of words. IGPT takes pixels as input and outputs target pixels similar to a text generation model.

Vision transformers (Vit) [14] are a kind of transformer architecture that only uses transformer encoders and custom patch embeddings to work with images. Vit works by splitting an image into patches, flatten them, and produce lower-dimensional linear embeddings from the flattened patches. After adding the positional embeddings, patches are feed those sequences of input to a transformer encoder. Vit models are usually pre-trained on a large dataset and then finetuned on smaller datasets for image classification. Vit was able to achieve an accuracy of 88.55% on ImageNet.

Our work differs from others by its unique way of understanding and creating images. Unlike prior methods, the proposed architecture utilizes vision transformers for image understanding and uses convolutional layers for image generation.

Figure 1: Tensor-to-Image Architecture

3 Method

3.1 Patches and Patch Encoding

Our model, Tensor-to-Image, uses Vit-like patches and patch embedding to understand the image. When the inputs are passed into the model it splits them into patches. The number of patches depends

on the image and patch sizes. The number of patches can be calculated with the following formula $(\text{Imageheight}/\text{patchsize})^2$. After Obtaining all patches, patches are flattened into $1 - D$ sequences. Also, trainable embeddings have been used to transform patches into latent vectors. After latent vectors have been generated, $1 - D$ learnable position embeddings are added to the patches to preserve the positional information of the patches that have been extracted from the image.

3.2 Transformer Layers

The Transformer layers are composed of a stack of N number of identical layers. Each layer has a multi-head self-attention and a feed-forward network. Layer normalization [25] is applied after each layer to keep all the activations in scale with each other. The output of each sub-layer in the network is $\text{LayerNorm}(x + F(x))$, where $F(x)$ is the operation done in the sub-layer itself.

3.3 Multi-head attention

Multi-head attention (MHA) contains H identical heads where each head computes independent scaled dot-product attention. Therefore each head can attend different places of the input. After the calculation of scaled dot-product attention, transformation is applied to concatenate the attention results of different heads before proceeding further in the model.

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O$$

The attention applied in the present work is the same as the transformer [7] attention and the formula,

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where, Q , K , V are the query, key, and the value respectively.

After the dot-product attention is applied, we scale the outputs since the dot-product attention tends to output large values. When we are done with the attention operations we normalize the outputs and finalize the outputs with a feed-forward linear layer.

3.4 Upsampling

To upsample the transformer encoders output residual layers and transpose convolutions have been used side by side. The output of Transformer encoders is passed to the transpose convolutions to generate the image. between all the transpose convolution a residual block is placed. After each layer either transposed convolution or residual block, batch normalization and ReLU have been applied.

3.5 Architecture

As shown in the figure 1 of the proposed architecture, images are split into patches and passed through the patch encoder that converts $2 - D$ images to embedded features. The resulting embedding vectors are passed to the transformer encoders, N number of them stacked on top of each other. Transformer encoder to have an understanding of the context, this case the image. The transformer encoder output is passed to a stack of a residual block followed by a transpose convolution, trailed by LeakyReLU and batch normalization layers combination. The number of repeated layers depends on the complexity of the input image. Finally, the outputs are passed to convolutional layers to generate the resulting image. For code, and examples please see <https://github.com/YigitGunduc/tensor-to-image>

4 Experiments

The model performance is tested on a variety of different tasks namely, image segmentation, depth prediction from a single image, and object segmentation. The same architecture is used for image segmentation, depth prediction from a single image, and object segmentation tasks. Once the architecture is optimized, without any change, three datasets are used for training and testing. The architecture optimization, functions of layers are set using above mentioned datasets.

The results section also includes a performance comparison section. The proposed architecture and two other commonly used architectures, U-Net and Autoencoder, are tested on the same data set. Both visual (Figure 5) and numeric comparisons (Table 1) presented. For numerical comparison three evaluation metrics (details are given in section 4.2).

4.1 Generator Architectures

The search for the best-performing architecture started by experimenting with various architectures. The carefully chosen datasets were tested on different architectures to identify the most adaptable one to obtain the best performance on a wide range of vision tasks. Hence, experiments are performed with various types of generators. The objective of the experiments was not only to find the best architecture but also to identify the faulty ones, where there exists space for improvements. In the end, we conclude with a generator that is capable and stable. In this paper, three of the experimental generators will be introduced. The selected generator, which will be called generator C, does not suffer from modal collapse or generate visual artifacts. The results obtained by using generator C are presented in the results section.

The first architecture is the most bare-bone generator (generator A). For this architecture, we did not use any advanced techniques used in GANs to improve the performance. The following basic steps constitute the backbone of more advanced architectures that we named generator B and generator C. The input image is split into patches, patches are embedded and positional information has been added. Following this step embedded patches are passed to transformer encoders where the model has an understanding of the input image. The transformer layer generates the input for convolutional layers. Transformer layers

output is passed to the through transpose convolutions to upsample until the image has reached the desired size. In the final step, the image is finalized with a convolutional layer with the same number of filters as the desired number of channels.

The generator B is the result of the second iteration of improvements. Generator B is a U-Net-based architecture. All the convolutional layers in the model have skip connections with encoded patches. The biggest hiccup was that embedded patches have a shape of (batch size, patch size, patch size, embedding dimension), and the further we go through the model the outputs of the convolutional layers get bigger since we apply upsampling.

It is not possible to perform a concatenation operation since the tensors have different shapes. We used a custom layer to overcome this problem. This custom layer takes both the embedded patches and the activation of the layer before. takes the embedded patches and upsample them to the output shape of the previous convolutional layer. The upsampling is done with bilinear interpolation. Also, convolution can be applied at will. The final step of this layer is to concatenate both tensors; upsampled patches and the activation of the previous layer. The details of the designed custom layer are presented in the appendix.

The third and last architecture (generator C) uses residual blocks in between the Transpose convolutions. The same architecture that we have used, described in this paper (figure 1). And all the results shown in the present work are generated by Generator C.

In this experimental study, it is observed that the best performance is achieved by using generator C. Generator A generates foggy results which are not sharp and realistic enough. Generator B tends to generate images that are corrupt or have visual artifacts. Another defect of generator B is that it requires way more data than the others. We could not succeed to get rid of those artifacts. Generator C generates sharp and realistic results and can operate with a limited amount of data hence it has been the best choice among the tested three architectures.

4.2 Evaluation Metrics

All the models are tested against three different evaluation metrics; Inception score(IS) [15], Frechet Inception Distance(FID) [17] and Structural Similarity(SSIM) [24]. Inception score is a metric about image quality and Image diversity. The IS works with a pre-trained deep convolutional neural network. Mostly the pre-trained network happens to be the Inception v3 model. Inception Score tries to answer two questions; Do images look like a specific object? Is a wide range of objects generated? Higher the IS indicates better-generated results since image quality is higher and generated images are from a more diverse pool. Frechet Inception Distance Is a metric for evaluation generative models. FID score is proposed as an improved version of the Inception score. FID score calculates the distance between the generated images and the real ones. Since the FID is the distance between the real and generated images a lower FID indicates better-quality images and a higher score indicates a lower-quality image. The structural similarity index is the measure of image similarity. SSIM is used for calculating the similarity between two images. SSIM compares an input image with a reference. The highest possible score of 1 which means two identical images.

4.3 Dataset

The datasets, Oxford-IIIT Pet Dataset [20], RGB-D [21], and Cityscapes [22, 23] are used for the performance evaluation of the model Tensor-to-Image.

4.3.1 Oxford-IIIT Pet Dataset

Oxford-IIIT Pet Dataset [20] is a 37 category pet dataset with 200 images for each class. The dataset contains various poses, scales of pet images under different lighting conditions. All ground truth images have an associated pixel-level trimap segmentation. The model is trained on the Oxford-IIIT Pet Dataset without a discriminator. The evaluation metric for the model to optimize is The Sparse Categorical Crossentropy loss. We inputted the model the ground truth images and expect it to produce pixel-level trimap segmentation maps as output. Figure 3 show the input target and output obtained from the model architecture.

4.3.2 RGB-D Dataset

RGB-D [21] dataset is a dataset of images with associated depth maps. Both indoor and outdoor scenes of various locations exist in the dataset. The dataset contains images from Yonsei and the Ewha Universities. Images of offices, rooms, exhibition centers, streets, the roads constitute the dataset. The model is trained on the RGB-D dataset to test the depth perception performance from a single image. The model successfully generates the depth maps from a single raw input image. Figure 4 shows input image, target and obtained output of the model.

4.3.3 Cityscapes Dataset

Cityscapes Dataset [22, 23] contains a diverse set of street scenes from 50 different cities. This dataset is used to test tasks of semantic urban scene understanding of the model. Figure 2 show image, target and model output for comparison and visual performance evaluation.

For all three datasets the performance of the model, without any tuning or additional layers, is very high, and the output images are close to target image quality. This fact can be observed by visual inspection on the figures 2,3,4.

4.4 Comparison with Other Architectures

We used some of the most advanced image evaluation metrics; Frechet Inception Distance(FID) [17], Structural Similarity(SSIM) [24], Inceptions Score(IS) [15] to evaluate our model and compare it with others. Inceptions score is one of the most popular evaluation metrics for GANs. We have compared our

architectures with U-Net and an Autoencoder which are popular choices for most of the image-to-image translation and image segmentation. We trained an Autoencoder, a U-Net, and our approach on the cityscapes dataset to convert photos to segmentation maps. The following table, (Table 1) measures the generated image quality among the different architectures according to the metrics mentioned above.

Model	FID	IS	SSIM
Tensor-to-Image(Ours)	834	1.267	0.70
U-Net	3946	1.163	0.52
Autoencoder	21182	1.203	0.26

Table 1: Comparerative performances of different architectures

Table 1 shows performance measures of two widely used and one introduced architectures for comparison. The performances of the compared models can also be compared by visual inspection using figure 5. The results presented in table 1, show that evaluation metric FID well explains the performance of Autoencoder architecture. The highest FID score well matches with the observed output images shown in figure 5. Autoencoders have a low-level understanding of the image but not much. Figure 5 shows that Autoencoder architecture is only able to distinguish the land from the sky thus generating unrealistic results that mostly contain two colors one for the sky and the other for the land. Despite U-Net is an Autoencoder with skip connections, both, evaluation metric results (table 1) and figure 5 indicate that its performance is far better than the Autoencoder. U-Net was able to pick up different objects in the scene and fill them with different colors. Nevertheless, its success is also limited. The generated results contain contours of the objects with different colors but there are no shadows or reflections which make an image closer to the photo-realistic level.

The performance of the proposed, Tensor-to-Image model is superior compared with two widely used architectures, namely Autoencoder and U-Net. First of all, evaluation metrics results (Table 1) show that for all three measures, FID, IS, and SSIM, Tensor-to-Image architecture performs better than the other two tested architectures.

5 Conclusion

Our work differs from others by its unique generator. Unlike other prior methods, it utilizes vision transformers for image understanding and uses convolutional layers for image generation. We have applied it to three different datasets to show how generalizable it is and also made a comparison with other popular architectures like autoencoder and u-net where it outperformed both. For code, and examples please see <https://github.com/YigitGunduc/tensor-to-image>

For all three datasets the performance of the model, without any tuning or additional layers, is very high, and the output images are close to target image quality. This fact can be observed by visual inspection on the figures 2,3,4.

This work shows that transformers have an even wider range of applications and they are perfectly

suitable for image to image translation. The model is trained on 3 different datasets and generates promising results. Also, a comparison is made between the U-Net and the Autoencoder. Our method outperformed both.

Figure 2: Cityscapes Dataset – Experiment on Semantic Understanding of the Model.

Figure 3: Oxford-IIIT Pet Dataset: Experiment on contour identification.

Figure 4: RGB-D Dataset: Experiment on Depth Perception of the Model

Figure 5: Visual Comparison of Autoencoder, U-Net, and Tensor-to-Image (Ours) Models

References

- [1] Rumelhart D.E., Hinton G.E., Williams R.J., Parallel Distributed processing: Explorations in the Microstructure of cognition, Vol. 1 Chap. Learning Internal Representations by Error Propagation, pp. 318-362 MIT Press, Cambridge, MA, USA, URL <http://dl.acm.org/citation.cfm?id=104279.104293> (1986).
- [2] Makhzani, A., Shlens, J., Jaitly, N., Goodfellow I.J., Adversarial Autoencoders. CoRR abs/1511.05644 (2015).
- [3] Ronneberger O., Fischer P., Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab N., Hornegger J., Wells W., Frangi A. (eds) Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015. MICCAI 2015. Lecture Notes in Computer Science, vol 9351. Springer, Cham. https://doi.org/10.1007/978-3-319-24574-4_28 (2015).
- [4] Pathak D., Krahenbuhl P., Donahue J., Darrell T., and A. Efros A. Context encoders: Feature learning by inpainting. In CVPR, 2, 3, 13, 17 (2016).
- [5] Wang X. and Gupta A. Generative image modeling using style and structure adversarial networks. In ECCV, 2, 3, 5, (2016).
- [6] Zhou Y. and Berg T. L., Learning temporal transformations from time-lapse videos. In ECCV, 2, 3, 8, (2016).
- [7] Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A. N., Kaiser L., Polosukhin I., Attention Is All You Need, arXiv1706.03762[cs.CL] (2017).
- [8] Devlin J., Chang Ming-Wei, Lee K., and Toutanova K., BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL, (2019).
- [9] Brown T. B., Mann B., Ryder N., Subbiah M., Kaplan J., Dhariwal P., Neelakantan A., Shyam P., Sastry G., Askell A., et al. Language models are few-shot learners. arXiv, 2020.
- [10] Radford A., Narasimhan K., Salimans T., and Sutskever I., Improving language understanding with unsupervised learning. Technical Report, (2018).
- [11] Hu H., Gu J., Zhang Z., Dai J., and Wei Y., Relation networks for object detection. In CVPR, (2018).
- [12] Locatello F., Weissenborn D., Unterthiner T., Mahendran A., Heigold G., Uszkoreit J., Dosovitskiy A., and Kipf T., Object-centric learning with slot attention. arXiv, (2020).
- [13] Chen Y-C. , Li L., Yu L., Kholy A. E., Ahmed F., Gan Z., Cheng Y., and Liu J., UNITER: UNiversal Image-TExt Representation Learning. In ECCV, (2020).
- [14] Buades A., Coll B., and Morel J.-M., A non-local algorithm for image denoising. In CVPR, 1 (2005).
- [15] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. In Advances in Neural Information Processing Systems, pp. 2234-2242, 2016.

- [16] Heusel, M., Ramsauer H., Unterthiner T., Nessler B., Hochreiter S., "GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium". *Advances in Neural Information Processing Systems*. 30. arXiv:1706.08500 (2017).
- [17] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Advances in neural information processing systems*, pp. 6626–6637, (2017).
- [18] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A Simple and Performant Baseline for Vision and Language. In *Arxiv*, 2019.
- [19] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, 2019.
- [20] Omkar M. Parkhi and Andrea Vedaldi and Andrew Zisserman and C. V. Jawahar, "Cats and Dogs, IEEE Conference on Computer Vision and Pattern Recognition, (2012).
- [21] J. Sturm and N. Engelhard and F. Endres and W. Burgard and D. Cremers, A Benchmark for the Evaluation of RGB-D SLAM Systems, *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, 2012
- [22] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] M. Cordts, M. Omran, S. Ramos, T. Scharwächter, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset," in *CVPR Workshop on The Future of Datasets in Vision*, 2015.
- [24] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004
- [25] Jimmy Lei Ba and Jamie Ryan Kiros and Geoffrey E. Hinton, Layer Normalization, arXiv:1607.06450, 2016.

6 Appendix

Network architecture

Following network, architecture is the one that has been used throughout the paper and the one that has generated the result shown in this paper. The model consists of Convolutional transpose layers, Convolutional layers, Residual layers, transformer layers, patch encoder, batch norm, Leaky ReLU. For code, and examples please see <https://github.com/YigitGunduc/tensor-to-image>

P16	-	PE	-	TL	-	TL	-	TL	-	TL	-					
CT512	-	RL512	-	CT256	-	RL256	-	CT64	-	RL64	-	CT32	-	RL32	-	C

- P** $\langle x \rangle$: Splits the image into patches with the patch size of x
- PE** : Encodes split patches and projects them to 64 dimensions.
- TL** : Denotes transformer layer with embedding dimension of 64, number heads of 2, and with 32 neurons.
- CT** $\langle x \rangle$: Convolutional transpose with x number of channels stride of 2 same paddings trailed by batch Normalization of Leaky Relu.
- RL** $\langle x \rangle$: Residual layer with x number of filters.
- C** $\langle x \rangle$: Convolutional layer with desire number of output channels and stride of oned and activation function of tanh.

6.1 Optimization

To optimize our networks we used Adam optimizer, with a learning rate of 0.0002, and momentum parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$.