

Chapter 11

Data Augmentation Using GANs for 3D Applications


Ioannis Maniadis

Information Technologies Institute, Centre for Research and Technology, Hellas, Greece

Vassilis Solachidis

Information Technologies Institute, Centre for Research and Technology, Hellas, Greece

Nicholas Vretos

 <https://orcid.org/0000-0003-3604-9685>

Information Technologies Institute, Centre for Research and Technology, Hellas, Greece

Petros Daras

 <https://orcid.org/0000-0003-3814-6710>

Information Technologies Institute, Centre for Research and Technology, Hellas, Greece

ABSTRACT

Modern deep learning techniques have proven that they have the capacity to be successful in a wide area of domains and tasks, including applications related to 3D and 2D images. However, their quality depends on the quality and quantity of the data with which models are trained. As the capacity of deep learning models increases, data availability becomes the most significant. To counter this issue, various techniques are utilized, including data augmentation, which refers to the practice of expanding the original dataset with artificially created samples. One approach that has been found is the generative adversarial networks (GANs), which, unlike other domain-agnostic transformation-based methods, can produce diverse samples that belong to a given data distribution. Taking advantage of this property, a multitude of GAN architectures has been leveraged for data augmentation applications. The subject of this chapter is to review and organize implementations of this approach on 3D and 2D imagery, examine the methods that were used, and survey the areas in which they were applied.

DOI: 10.4018/978-1-5225-5294-9.ch011

INTRODUCTION

The advances that have been made in the field of deep learning have provided us with ever more potent tools, which can be applied in an increasing number of tasks, computer vision being foremost among them. Concurrently, the potential value of data has become apparent, and so data gathering and mining are now employed in several domains in order to make it possible for deep learning techniques to be applied in those domains. Deep learning models require data for their training which constitute a representative sampling of a given task. When the available data only relate to a subset of cases, a model will only learn to address those cases only and fail in the task overall. For this reason, deep learning models generally require significant amounts of data for their training. Nonetheless, in many cases the available data is not sufficient for training models that generalise adequately. That may occur either because data gathering might be difficult in a given setting (due to scarcity of subject cases or due to difficulties in collecting them) or because the available data are not annotated. In all the above cases, other methods must be employed to make deep learning feasible.

Several techniques have been developed in order to tackle the above-mentioned problems, particularly when dealing with 2D and 3D image data. One point of focus is to develop architectural modifications that make models generalize better in a given task, such as dropout and weight regularization (Sutskever, Hinton, Krizhevsky, & Salakhutdinov, 2014). Another approach suggests expanding the initial dataset by manipulating the existing data and creating new synthetic samples. This approach is usually referred to as data augmentation. The most frequent implementations of data augmentation are the addition of random noise to the data and the application of geometric and/or other transformations (Taylor & Nitschke, 2019). The latter is particularly effective in image data, whose features have spatial properties. Those data augmentation methods however, while suitable for image data, are domain-agnostic, since they apply transformations without taking into account the nature, characteristics and features of the original data, and produce synthetic samples that could deviate from the original distribution.

In order to achieve these two objectives, that is to augment a dataset with meaningfully and significantly diverse samples, a method would be required that augments a dataset in ways specific to its properties, so that the generated samples would cover the largest area of the sample space possible, without deviating from it. GANs (Goodfellow et al., 2014), (Zoumpourlis, Doumanoglou, Vretos, & Daras, 2017), (Shijie, Ping, Peiyi, & Siping, 2017) as is demonstrated in (Shijie, Ping, Peiyi, & Siping, 2017), can be used to augment data in this exact way, and so they are an attractive alternative to the above domain agnostic methods. Ideally, GANs produce samples that belong to the original data distribution, while at the same time they differ from any given sample of that distribution. In this way they fulfill the essential objective of data augmentation, which is to provide to the model a diverse sample pool with which to train, which is representative of a given task. Additionally, the fundamental formulation of GANs has proven to be remarkably flexible, in that GANs can be modified to generate samples in many different ways, and can be combined with a variety of architectures to tackle different data augmentation tasks. However, GANs are also remarkably hard to train (Goodfellow I., 2016), and so have been the subject of intense study in an effort to develop more efficient architectures (Arjovsky & Bottou, 2017), (Salimans et al., 2016). It is important to note at this point that unlike other GAN review papers, such as (Z. Wang, She, & Ward, 2019) and (Pan et al., 2019), this work does not provide a broad review of GAN models. Rather, it limits its purview to cases where GANs have been used in the context of 2D and 3D image data augmentation for the purpose of improved performance in classification, segmentation, object detection/identification

39 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/data-augmentation-using-gans-for-3d-applications/254595

Related Content

Learning Robot Vision for Assisted Living

Wenjie Yan, Elena Torta, David van der Pol, Nils Meins, Cornelius Weber, Raymond H. Cuijpers and Stefan Wermter (2013). *Image Processing: Concepts, Methodologies, Tools, and Applications* (pp. 1232-1255). www.irma-international.org/chapter/learning-robot-vision-assisted-living/77596

A Recurrent Probabilistic Neural Network for EMG Pattern Recognition

Toshio Tsuji, Nan Bu and Osamu Fukuda (2008). *Pattern Recognition Technologies and Applications: Recent Advances* (pp. 370-387). www.irma-international.org/chapter/recurrent-probabilistic-neural-network-emg/28039

Reversible Data Hiding Based on Statistical Correlation of Blocked Sub-Sampled Image

Mona Nafari, Mansour Nejati Jahromi and Gholam Hosein Sheisi (2012). *International Journal of Computer Vision and Image Processing* (pp. 42-59). www.irma-international.org/article/reversible-data-hiding-based-statistical/68003

Recognition of Color Objects Using Hybrid Descriptors

Driss Naji, M. Fakir, B. Bouikhalene and M. Boutaoune (2013). *International Journal of Computer Vision and Image Processing* (pp. 60-68). www.irma-international.org/article/recognition-of-color-objects-using-hybrid-descriptors/103959

A Framework for Interactive 3D Rendering on Mobile Devices

Ekta Walia and Vishal Verma (2014). *International Journal of Computer Vision and Image Processing* (pp. 18-31). www.irma-international.org/article/a-framework-for-interactive-3d-rendering-on-mobile-devices/115837