Doctoral thesis position – 2020

Title: Learning visual texture features for inverse procedural modeling and texture synthesis by example

Host team: IGG (Computer Graphics and Geometry Group), ICube laboratory

Advisor: Jean-Michel Dischler, Professor in Computer Science (dischler@unistra.fr)

Co-advisor: Rémi Allègre, Associate Professor in Computer Science (remi.allegre@unistra.fr)

Starting date: October 2020

Ending date: 3 years from the starting date

Funding: Doctoral contract (about 1500 EUR per month, net salary)

Location: Strasbourg area, France

Keywords: Data science, computer graphics, image processing, textures, machine learning, deep learning

Desired skills:
- Data science and/or computer graphics or image processing
- Basic skills in machine and deep learning
- Knowledge in texture synthesis is a plus

Skills that will be acquired:
- Texture analysis
- By-example texture synthesis
- Inverse procedural modeling
- Machine learning
- Deep learning
- Data augmentation

Abstract

By-example texture synthesis aims at generating and editing visual textures similar to input texture samples, the former being larger than the latter, or of the same size depending on the applications. The two prominent approaches are *procedural modeling* and *data-driven texture synthesis*. Procedural approaches seek to represent textures by a mathematical model, whereas data-driven techniques rely on the matching of pixel neighborhoods or global image statistics on pixel intensities or filter responses. Inverse procedural modeling from arbitrary examples is a scientific bottleneck. It stems from a lack of image analysis tools tailored for the decomposition of the various visual texture features, especially *structure* and *noise* at multiple scales.
The goal of this thesis is to devise novel unsupervised or semi-supervised methods to learn visual texture features in the context of inverse procedural modeling and by-example texture synthesis. Particular attention will be paid to challenging heterogeneous textures exhibiting complex structures. The developed approach will be inspired by representation learning, with a focus on image factorization and attribute learning. Two fields of applications will be considered: texturing of virtual 3D environments and data augmentation for the training of models in histopathological slide image analysis.

1. Context and problem statement

Textures, in the visual meaning, refer to the visual impression produced to human observers by the surface of visible objects. Textures are characterized by local spatial variations of stimuli like color, orientation and intensity in an image. In the fields of image analysis and synthesis, textures are also characterized at multiple scales by spatial properties like stationarity (vs. non-stationarity) and regularity (vs. randomness), as well as by other statistical properties that usually lead to different representation models and computational approaches. Depending on the area of application and considered objects, textures can also be assumed to match the realization of known statistical processes (e.g. a brick wall or a textile), or context-dependent rules (e.g. the anatomy of biological structures in a histological slide).

By-example texture synthesis aims at generating and editing visual textures similar to input texture samples, the former being larger than the latter, or of the same size depending on the applications [WLK+09, RDD+17]. There are multiple areas of application, ranging from texturing of virtual 3D environments to data augmentation for the training of models for medical image analysis. The two prominent approaches are *procedural modeling* and *data-driven texture synthesis*. Procedural approaches seek to represent textures by a mathematical model, which is typically a graph whose nodes are either procedures of functions. Data-driven texture synthesis relies either on the matching of pixel neighborhoods, or on the matching of global image statistics on pixel intensities or filter responses. The procedural representation has several benefits: it is easily controllable, compact and spatially unbounded. However, devising procedural models is a complex task and inverse modeling from exemplars is restricted to homogeneous noise textures [GLL+12, GLM17, GSD+17], or it proceeds by browsing collections of predefined models [LGD18, HDR19]. Inverse procedural modeling from arbitrary texture exemplars is a scientific bottleneck, that requires to develop image analysis tools tailored for the decomposition of the various visual texture features in a relevant way, especially *structure* and *noise* at multiple scales. The structure coincides with the visual features that are invariant for different observations (e.g. several patches of a same texture), whereas the noise corresponds to the random part that gives a specific appearance to each observation. This decomposition problem is difficult to address with state-of-the-art image processing techniques owing to the great variety of textures and the lack of prior knowledge on the scales and features to consider. There is a need of *high-level* image representations, suitable for the purpose of texture modeling and by-example synthesis. *Representation learning* [BCV13] is a promising line of research to tackle this problem, that has recently introduced new techniques to decompose images into multiple independent *factors* in an unsupervised or semi-supervised way. Generative models such as InfoGAN [CDH16], StyleGAN [KLA19] or STGAN [LDX+19] use adversarial neural networks to learn attributes that have systematic and predictable effects on the output, like face orientation or presence/absence of visual features at different scales. To our knowledge, this approach with high potential has never been explored for texture modeling and synthesis.

2. Thesis goal

The goal of this thesis is to tackle the scientific bottlenecks that have just been presented, by developing novel unsupervised or semi-supervised techniques to learn visual texture features for the purpose of inverse procedural modeling and by-example texture synthesis. Heterogeneous textures exhibiting complex structures will be especially considered.

This thesis will start by a review of state-of-the-art texture models and by-example texture synthesis methods. The review will cover recent work on generative models, comprising adversarial neural networks (GANs) and variational autoencoders (VAEs). It will also encompass image factorization and attribute learning methods. The conditions in which these methods are applicable, including the properties of the images and their distributions, will be thoroughly evaluated. Then the next step will be to extend these methods to textures by making necessary adjustments, with applications in the two fields where the IGG group carries out its research: texturing of virtual 3D environments, and data augmentation for the training of models in histopathological slide image analysis.

Textures contribute in a prominent way to the realism of virtual 3D environments, like e.g. natural landscapes of urban environments. The goal of by-example texture synthesis is to facilitate the work of artists who have to cope with the increasing demand of highly detailed digital content in the Computer Graphics industry. In this context, visual texture features will be grasped in connection with generic procedural structure and noise models developed by the IGG group [LSA+16, GSD+17], as well as with a database of structures extracted from a collection of textures by hand.

In the medical image analysis community, deep learning has attracted a lot of attention due to its performance for image classification and as a complement for image interpretation. Some recent results demonstrate that model-based prediction can perform equal or better than image analysis

by human experts [LFK+19, MSG+20, HPA+20]. The methods rely on the availability of large labeled datasets, which are necessary for robust model training. However, collecting and labeling images are a challenge for certain pathologies. The number of available exemplars is often small due to the high level of expertise and cost in time required for both data acquisition (biopsies, resections or ablations of organs), manual segmentation and annotation on images of extremely high resolution (up to 100k×100k). Data augmentation aims at alleviating this issue by generating additional and diversified data samples [SK19]. Histopathological slide images can be considered as heterogeneous textures with rich structure and multiple shape variations. In this context, visual texture features will be connected to known biological descriptions. The validation of the results will be achieved using the datasets and methods provided by the SDC (Data Science and Knowledge) group of the ICube laboratory.
Implementations will be performed in the texture analysis and synthesis library [ASTex] developed by the IGG group.

3. References

[ASTex] Analysis and Synthesis of Textures library, IGG group. https://astex-icube.github.io

[BCV13] Y. Bengio, A. Courville, P. Vincent. Representation Learning: A Review and New Perspectives. IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 35, Issue 8, 2013.

[CDH16] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. Advances in Neural Information Processing Systems, Number 29, Pages 2172-2180, 2016.

[GLL+12] B. Galerne, A. Lagae, S. Lefebvre, and G. Drettakis. 2012. Gabor Noise by Example. ACM Transactions on Graphics, 31, 4, Article 73, July 2012.

[GLM17] B. Galerne, A. Leclaire, and L. Moisan. Texton Noise. Computer Graphics Forum, Volume 36, Number 8, Pages 205-218, 2017.

[GSD+17] G. Guingo, B. Sauvage, J-M. Dischler, M-P. Cani. Bi-Layer textures: a Model for Synthesis and Deformation of Composite Textures. Computer Graphics Forum, pages 111-122, Volume 36, Number 4, 2017.

[HDR19] Y. Hu, J. Dorsey, and H. Rushmeier. A Novel Framework For Inverse Procedural Texture Modeling. ACM Transactions on Graphics, 38, 6, Nov. 2019.

[HPA+20] T.C. Hollon, B. Pandian, B., A.R. Adapa, *et al.* Near real-time intraoperative brain tumor diagnosis using stimulated Raman histology and deep neural networks. Nature Medicine, 2020.

[KLA19] T. Karras, S. Laine, and T. Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.

[LFK+19] Liu, X., L. Faes, A. U. Kale, S. K.Wagner, D. J. Fu, A. Bruynseels, T. Mahendiran, G. Moraes, M. Shamdas, C. Kern, J. R. Ledsam, M. K. Schmid, K. Balaskas, E. J. Topol, L. M. Bachmann, P. A. Keane, and A. K. Denniston (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. The Lancet Digital Health 1(6), e271 – e297.

[LDX+19] M. Liu, Y. Ding, M. Xia, X. Liu, E. Ding, W. Zuo, and S. Wen. STGAN: A Unified Selective Transfer Network for Arbitrary Image Attribute Editing. Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.

[LGD18] J. Liu, Y. Gan, J. Dong, L. Qi, X. Sun, M. Jian, L. Wang, H. Yu. Perception-driven procedural texture generation from examples. Neurocomputing, Volume 291, Pages 21-34, 2018.

[LSA+16] Y. Lockerman, B. Sauvage, R. Allegre, J-M. Dischler, J. Dorsey, H. Rushmeier. Multi-Scale Label-Maps Extraction for Texture Synthesis. ACM Transactions on Graphics, Volume 35, Number 4, pages 140:1-140:12, 2016

[MSG+20] S. M. McKinney, M. Sieniek, V. Godbole, *et al.* International evaluation of an AI system for breast cancer screening. Nature 577**,** 89–94, 2020.

[RDD+17] L. Raad, A. Davy, A. Desolneux, J.-M. Morel. A survey of exemplar-based texture synthesis. Annals of Mathematical Sciences and Applications, Volume 3, 2017.

[SK19] C. Shorten, T. M. Khoshgoftaar. A survey on Image Data Augmentation for Deep Learning. Journal of Big Data, 6, 1, Jul 2019.

[WLK+09] Li-Yi Wei, S. Lefebvre, V. Kwatra, and G. Turk. State of the Art in Example-based Texture Synthesis. Eurographics '09 State of the Art Reports (STARs), 2009.

[YWB19] X. Yi, E. Walia, P. Babyn. Generative adversarial network in medical imaging: A review. Medical Image Analysis, Volume 58, 2019.