# Understanding Mystery Behind Example-Based Image Synthesis

J. Lu,[1] M. Lukáč,[1] and D. Sýkora[2]

[1]Adobe Research, USA
[2]Czech Technical University in Prague, Faculty of Electrical Engineering, Czech Republic

**Figure 1:** *Example-based image synthesis applications explained in this tutorial (from left to right): appearance transfer to fluid simulations [JFA*15], directional texture painting [LFA*15], illumination-driven artistic style transfer to 3D models [FJL*16], and example-based stylization of head portraits [FJS*17].*

**Abstract**

*This tutorial presents a concise overview of development in the field of example-based image synthesis that over the last two decades rapidly evolved into a powerful tool enabling the production of synthetic imagery often indistinguishable from the source exemplar. We discuss not only the basic algorithmic concepts but also their further improvements which lead to significant reduction of computational overhead as well as better visual quality. We also demonstrate numerous applications including texture synthesis, hole-filling, video completion, retargeting, reshuffling, morphing, melding, painting by feature, appearance transfer to fluid animations or artistic style transfer to 3D models and facial animations.*

## 1. Intended Audience

Our tutorial is tailored to students, researchers, and developers who are interested in example-based content creation with a specific emphasis on patch-based synthesis. Although it may seem this topic nowadays is a bit outdated due to the advent of parametric neural-based techniques we demonstrate that non-parametric approaches with proper guidance still deliver state-of-the-art quality and can be a valuable inspiration for further development combining current advances in deep learning with the power of patch-based techniques.

## 2. Background

This tutorial was inspired by a recently published comprehensive survey about patch-based synthesis techniques written by Barnes & Zhang [BZ17] which stems from and extends previous state-of-the-art report by Wei et al. [WLKT09]. Those two seminal publications can serve as a primary reference and guidepost for all participants where they can find the essential information at one place. In this tutorial, we further extend those two surveys by adding new methods which were published after publication of Barnes & Zhang's paper.

## 3. Prerequisites

There are no specific prerequisites to this tutorial. We just assume a basic knowledge of image processing. All necessary principles will be described in sufficient detail. Links to further readings are provided in this tutorial description document for those who are interested in more in-depth understanding of presented methods and their applications.

## 4. Presenters

**Jingwan Lu** is a Research Scientist at the Creative Intelligence lab of Adobe Research. She has a passion for data-driven content creation. She has worked on brush models, stylization, guided texture synthesis, voice synthesis, etc using various data-driven approaches. She has ventured into deep generative models in recent



DOI: 10.2312/egt.20181027

years. Her vision is to harness the power of machine learning in the age of data explosion to invent the next generation image and video editing tools.

**Michal Lukáč** is a Research Scientist at the Creative Intelligence lab of Adobe Research. In his job as a toolmaker, he is working to empower artists to proliferate their creativity.

**Daniel Sýkora** is an Associate Professor at the Department of Computer Graphics and Interaction, Faculty of Electrical Engineering, Czech Technical University in Prague. Daniel's lifelong passion is to eliminate repetitive and time-consuming tasks while retaining the uniqueness of hand-drawn style. He found example-based image synthesis to be a powerful tool that can help to fulfill this goal. For his contribution to the field, Daniel received prestigious Neuron Award for Promising Young Scientists.

## 5. Outline

In this section, we describe the content of the tutorial in more detail. The tutorial is subdivided into three main thematic blocks each of which is governed by one of the speakers.

### 5.1. Basic concepts (Daniel Sýkora)

Research in the field of example-based image synthesis can be divided into two main branches: *parametric* [PS00] and *non-parametric* [EL99]. At the beginning of our tutorial, we describe those two concepts in more detail and compare their performance concerning quality and ability to reproduce different style exemplar. After this initial overview, we focus more on non-parametric techniques which in general better preserve the fidelity of the source exemplar. Finally, we introduce the concept of guided synthesis [HJO*01] which enables to control the synthesis process using a set of user-specified guiding channels (see Fig. 2).

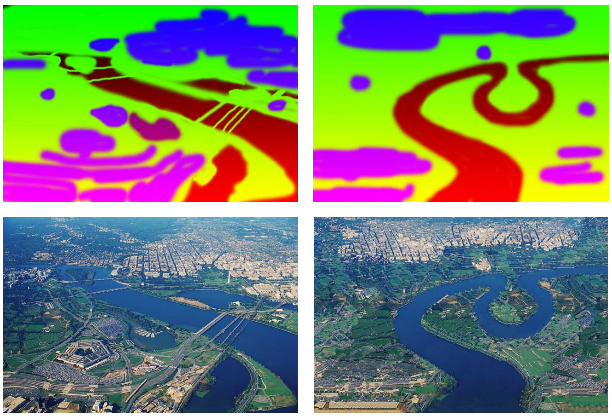

**Figure 2:** *An illustration of the guided synthesis principle [HJO*01]: original photo (bottom left) is annotated with segmentation guide (top left), the user can then provide segmentation guide for the target image (top right), and the algorithm synthesizes image using the original photo as a source while respecting the correspondence between segmentation guides (bottom right).*

After the introduction of basic principles, we will focus more on the energy-based formulation of non-parametric synthesis [KEBK05] and describe a popular algorithm [WSI07] which allows minimization of the proposed energy function efficiently. One of the critical limitations of this approach, however, is the enormous computational overhead connected with expensive nearest-neighbor retrieval which represents a central bottleneck of the whole method. To alleviate it we describe a randomized algorithm (PatchMatch [BSFG09]) which can deliver comparable nearest-neighbor field quality orders of magnitude faster as compared to the deterministic solution.

### 5.2. Applications of EM-based texture synthesis (Jingwan Lu)

Texture synthesis using EM-based optimization finds many graphics and vision applications. We will introduce some most influential work in recent years to give the audience examples of how the basic concepts presented in the first session can be extended in different ways for diverse applications.

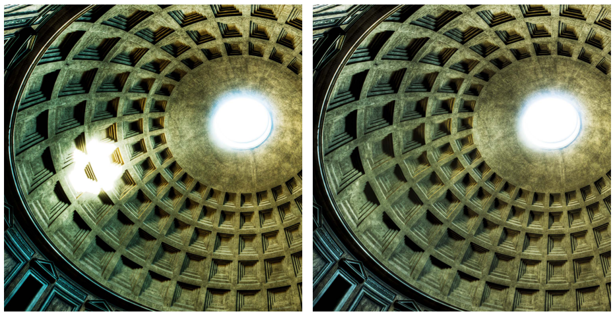

**Figure 3:** *An example of advanced image completion operation (right) which uses a variant of EM-based texture synthesis [DSB*12] while taking into account different orientations and scales of the patches taken from the inpainted source (left).*

We will first introduce how the basic formulation of unguided texture synthesis can be extended to go beyond 2D for synthesizing geometry on surfaces or 3D volumetric textures [LHGM05, KFCO*07, DLTD08, LPKK12]. Then, we will do a deep dive into the recent applications of guided synthesis in the image and video domain. We will present unconstrained or implicitly-guided texture synthesis approaches for image hole filing and completion [BSFG09, CGMP11, DSB*12, GM14, HKAK14] (see Fig. 3), content reshuffling [CBAF08, BSFG09] and retargeting [SCSI08], and extend the concept to incorporate temporal constraints for video inpainting [WSI07, GKT*12, NAF*14]. Finally, we will describe the recent advances in explicitly-guided texture synthesis for two particular application domains, stylization [HJO*01, BCK*13, FLJ*14, FJL*16, FJS*17] and painting [RLC*06, LFB*13, LFA*15] (see Fig. 1).

### 5.3. Addressing problems of EM-based texture synthesis and corresponding applications (Michal Lukáč)

Next, we will have a look at the fundamental limitations of patch-based approaches, proposed ways of dealing with them, and open problems for future work.

In recent years, a common problem of patch-based optimization called "wash-out" has been identified. This problem affects simple as well as guided synthesis and results in a dramatic loss of high-frequency content. We will take a look at its causes, the issues it creates, and how the optimization can be re-formulated to address it. We will explain bidirectional similarity [WHZ*08, SCSI08], color histogram matching [KFCO*07], uniform patch utilization [RCOL09, JFA*15, KNL*15] (see Fig. 4) and adaptive uniform utilization [FJL*16] (see Fig. 1).

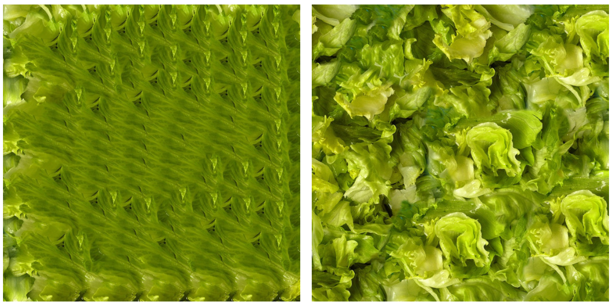

**Figure 4:** *An example of "wash-out" effect: a texture synthesized using standard energy formulation of Kwatra et al. [KEBK05] (left) and a result of improved energy by Kaspar et al. [KNL*15] which takes into account uniform utilization of source patches (right).*

Applications such as stylization take advantage of guided synthesis, but the choice of guidance channels has so far been heuristic. We will consider the heuristics used so far [JFA*15, FJL*16, FJS*17] and compare them with alternatives, such as neural-based guidance [LYY*17].

A fundamental limitation of non-parametric example-based approaches is that they cannot hallucinate content. We will show how other methods have been used to handle this problem [HE07, PKD*16, ISSI17] and how they can tie into patch-based optimization.

Finally, we will examine the question whether in light of the above problems the EM patch-based optimization is really a useful model, and what alternatives there could be.

## References


[BCK*13] BÉNARD P., COLE F., KASS M., MORDATCH I., HEGARTY J., SENN M. S., FLEISCHER K., PESARE D., BREEDEN K.: Stylizing animation by example. *ACM Transactions on Graphics 32*, 4 (2013), 119. 2

[BSFG09] BARNES C., SHECHTMAN E., FINKELSTEIN A., GOLDMAN D. B.: PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics 28*, 3 (2009), 24. 2

[BZ17] BARNES C., ZHANG F.-L.: A survey of the state-of-the-art in patch-based synthesis. *Computational Visual Media 3*, 1 (2017), 3–20. 1

[CBAF08] CHO T. S., BUTMAN M., AVIDAN S., FREEMAN W. T.: The patch transform and its applications to image editing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2008). 2

[CGMP11] CAO F., GOUSSEAU Y., MASNOU S., PÉREZ P.: Geometrically guided exemplar-based inpainting. *SIAM Journal on Imaging Sciences 4*, 4 (2011), 1143–1179. 2

[DLTD08] DONG Y., LEFEBVRE S., TONG X., DRETTAKIS G.: Lazy solid texture synthesis. *Computer Graphics Forum 27*, 4 (2008), 1165–1174. 2

[DSB*12] DARABI S., SHECHTMAN E., BARNES C., GOLDMAN D. B., SEN P.: Image Melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on Graphics 31*, 4 (2012), 82. 2

[EL99] EFROS A. A., LEUNG T. K.: Texture synthesis by non-parametric sampling. In *Proceedings of IEEE International Conference on Computer Vision* (1999), pp. 1033–1038. 2

[FJL*16] FIŠER J., JAMRIŠKA O., LUKÁČ M., SHECHTMAN E., ASENTE P., LU J., SÝKORA D.: StyLit: Illumination-guided example-based stylization of 3D renderings. *ACM Transactions on Graphics 35*, 4 (2016), 92. 1, 2, 3

[FJS*17] FIŠER J., JAMRIŠKA O., SIMONS D., SHECHTMAN E., LU J., ASENTE P., LUKÁČ M., SÝKORA D.: Example-based synthesis of stylized facial animations. *ACM Transactions on Graphics 36*, 4 (2017), 155. 1, 2, 3

[FLJ*14] FIŠER J., LUKÁČ M., JAMRIŠKA O., ČADÍK M., GINGOLD Y., ASENTE P., SÝKORA D.: Color Me Noisy: Example-based rendering of hand-colored animations with temporal noise control. *Computer Graphics Forum 33*, 4 (2014), 1–10. 2

[GKT*12] GRANADOS M., KIM K. I., TOMPKIN J., KAUTZ J., THEOBALT C.: Background inpainting for videos with dynamic objects and a free-moving camera. In *Proceedings of European Conference on Computer Vision* (2012), pp. 682–695. 2

[GM14] GUILLEMOT C., MEUR O. L.: Image inpainting: Overview and recent advances. *IEEE Signal Processing Magazine 31*, 1 (2014), 127–144. 2

[HE07] HAYS J., EFROS A. A.: Scene completion using millions of photographs. *ACM Transactions on Graphics 26*, 3 (2007). 3

[HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURLESS B., SALESIN D. H.: Image analogies. In *SIGGRAPH Conference Proceedings* (2001), pp. 327–340. 2

[HKAK14] HUANG J.-B., KANG S. B., AHUJA N., KOPF J.: Image completion using planar structure guidance. *ACM Transactions on Graphics 33*, 4 (2014), 129. 2

[ISSI17] IIZUKA S., SIMO-SERRA E., ISHIKAWA H.: Globally and locally consistent image completion. *ACM Transactions on Graphics 36*, 4 (jul 2017), 107:1–107:14. 3

[JFA*15] JAMRIŠKA O., FIŠER J., ASENTE P., LU J., SHECHTMAN E., SÝKORA D.: LazyFluids: Appearance transfer for fluid animations. *ACM Transactions on Graphics 34*, 4 (2015), 92. 1, 3

[KEBK05] KWATRA V., ESSA I. A., BOBICK A. F., KWATRA N.: Texture optimization for example-based synthesis. *ACM Transactions on Graphics 24*, 3 (2005), 795–802. 2, 3

[KFCO*07] KOPF J., FU C.-W., COHEN-OR D., DEUSSEN O., LISCHINSKI D., WONG T.-T.: Solid texture synthesis from 2D exemplars. *ACM Transactions on Graphics 26*, 3 (2007), 2. 2, 3

[KNL*15] KASPAR A., NEUBERT B., LISCHINSKI D., PAULY M., KOPF J.: Self tuning texture optimization. *Computer Graphics Forum 34*, 2 (2015), 349–360. 3

[LFA*15] LUKÁČ M., FIŠER J., ASENTE P., LU J., SHECHTMAN E., SÝKORA D.: Brushables: Example-based edge-aware directional texture painting. *Computer Graphics Forum 34*, 7 (2015), 257–268. 1, 2

[LFB*13] LUKÁČ M., FIŠER J., BAZIN J.-C., JAMRIŠKA O., SORKINE-HORNUNG A., SÝKORA D.: Painting by feature: Texture boundaries for example-based image creation. *ACM Transactions on Graphics 32*, 4 (2013), 116. 2

[LHGM05] LAI Y.-K., HU S.-M., GU D. X., MARTIN R. R.: Geometric texture synthesis and transfer via geometry images. In *Proceedings of ACM Symposium on Solid and Physical Modeling* (2005), pp. 15–26. 2

[LPKK12] LEE S.-H., PARK T., KIM J.-H., KIM C.-H.: Adaptive synthesis of distance fields. *IEEE Transactions on Visualization and Computer Graphics 18*, 7 (2012), 1135–1145. 2

[LYY*17] LIAO J., YAO Y., YUAN L., HUA G., KANG S. B.: Visual attribute transfer through deep image analogy. *ACM Transactions on Graphics 36*, 4 (2017), 120. 3

[NAF*14] NEWSON A., ALMANSA A., FRADET M., GOUSSEAU Y., PÉREZ P.: Video inpainting of complex scenes. *SIAM Journal of Imaging Science 7*, 4 (2014), 1993–2019. 2

[PKD*16] PATHAK D., KRÄHENBÜHL P., DONAHUE J., DARRELL T., EFROS A.: Context encoders: Feature learning by inpainting. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2016). 3

[PS00] PORTILLA J., SIMONCELLI E. P.: A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision 40*, 1 (2000), 49–70. 2

[RCOL09] ROSENBERGER A., COHEN-OR D., LISCHINSKI D.: Layered Shape Synthesis: Automatic generation of control maps for non-stationary textures. *ACM Transactions on Graphics 28*, 5 (2009), 107. 3

[RLC*06] RITTER L., LI W., CURLESS B., AGRAWALA M., SALESIN D.: Painting with texture. In *Proceedings of Eurographics Symposium on Rendering* (2006), pp. 371–376. 2

[SCSI08] SIMAKOV D., CASPI Y., SHECHTMAN E., IRANI M.: Summarizing visual data using bidirectional similarity. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (2008). 2, 3

[WHZ*08] WEI L.-Y., HAN J., ZHOU K., BAO H., GUO B., SHUM H.-Y.: Inverse texture synthesis. *ACM Transactions on Graphics 27*, 3 (2008). 3

[WLKT09] WEI L.-Y., LEFEBVRE S., KWATRA V., TURK G.: State of the art in example-based texture synthesis. In *Eurographics (STARs)* (2009), pp. 93–117. 1

[WSI07] WEXLER Y., SHECHTMAN E., IRANI M.: Space-time completion of video. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 3 (2007), 463–476. 2