

Example-Based Image Synthesis via Randomized Patch-Matching

Yi Ren, Yaniv Romano, and Michael Elad, *Fellow, IEEE*

Abstract—Image and texture synthesis is a challenging task that has long been drawing attention in the fields of image processing, graphics, and machine learning. This problem consists of modeling the desired type of images, either through training examples or via a parametric modeling, and then generating images that belong to the same statistical origin. This paper addresses the image synthesis task, focusing on two specific families of images—handwritten digits and face images. This paper offers two main contributions. First, we suggest a simple and intuitive algorithm capable of generating such images in a unified way. The proposed approach taken is pyramidal, consisting of upscaling and refining the estimated image several times. For each upscaling stage, the algorithm randomly draws small patches from a patch database and merges these to form a coherent and novel image with high visual quality. The second contribution is a general framework for the evaluation of the generation performance, which combines three aspects: the likelihood, the originality, and the spread of the synthesized images. We assess the proposed synthesis scheme and show that the results are similar in nature, and yet different from the ones found in the training set, suggesting that true synthesis effect has been obtained.

Index Terms—Image synthesis, patch-matching, hallucination, t-SNE, MNIST, face images.

I. INTRODUCTION

THE task of image synthesis is central in the fields of image processing, graphics, and machine learning. The motivation to study this topic has several origins. First, the availability of a technique for generating images that obey a given patch-based image model establishes an ultimate way for testing local models and their suitability to treat images. Second, a successful synthesis algorithm may improve the performance of image restoration tasks, e.g. denoising [1]–[3], demosaicking [3], inpainting [4], super-resolution [5], and

other ill-posed inverse problems [6], as well as related image processing tasks, such as image analogies [7], [8] or style transfer [9]. Moreover, in case of severe corruption, the conventional image restoration algorithms do not achieve clear and sharp images, while such synthesis methods could be used to provide a plausible solution, out of infinitely many possible ones. Put differently, rather than considering a solution as a direct recovery task, armed with a good synthesis method one may migrate the treatment to have a randomized hallucination flavor. Third, the generated images themselves are interesting, since they are created out of “thin air”, and may be considered as an appealing art. Indeed, this is the effect of the recently introduced synthesized images [9]–[16]. Finally, the synthesis methods can be converted to handle other data sources such as music, thereby enriching the scope of this field of research far beyond its original objectives.

When addressing the general image synthesis problem, one may narrow down the scope to a more specific task to ease the otherwise quite complicated general objective. Indeed, over the years many works focus on texture synthesis [17]–[22], handwritten digits [13]–[15], [23]–[25], and human faces [13], [25]–[28]. These cases are all appealing as test cases for synthesis algorithms, and as a stepping stone towards the ultimate problem of general content image generation. This is also the path we shall take in this work.

Traditionally, in the case of texture synthesis, the generation is regularly done using example-based models [17], [18], [20]–[22]. Newly emerging methods in the field of image synthesis commonly model the probability distribution of the images by neural networks, then randomly sample from it to generate new content [10], [11], [13]–[15], [19], [23]–[26], [29], [30]. These methods have been shown to lead to interesting results for generating digits, faces, textures, and even natural scenes. However, a fundamental drawback of these methods is their tendency to be difficult to train and interpret.

Interestingly, in the realm of image synthesis, little existing work has relied on patch-based image models, despite the demonstrated effectiveness of such techniques in image restoration [1]–[6], image completion and interactive image editing [8], [31]–[34], texture-synthesis [18], and many other tasks [35]. More specifically, modeling of patches, either directly using examples [4], [8], [17], [18], [20]–[22], [31]–[36], or via a parametric form [1]–[3], [5], [6], [37], [38], has been shown to be highly expressive and rich. Perhaps the reason of avoiding such localized methods in image synthesis is the inevitable need to quilt or otherwise aggregate these patches to form the final created image in a way that is globally faithful.

Manuscript received October 11, 2016; revised April 9, 2017 and August 24, 2017; accepted August 30, 2017. Date of publication September 8, 2017; date of current version October 20, 2017. This work was supported in part by the European Research Council through the EU 7th Framework Programme, ERC, under Grant 320649, and in part by the Israel Science Foundation under Grant 1770/14. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Catarina Brites.

Y. Ren and M. Elad are with the Computer Science Department, Technion–Israel Institute of Technology, Technion City, Haifa 32000, Israel (e-mail: yi.ren@alumni.technion.ac.il; elad@cs.technion.ac.il).

Y. Romano is with the Department of Electrical Engineering, Technion–Israel Institute of Technology, Technion City, Haifa 32000, Israel (e-mail: yromano@tx.technion.ac.il).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author. The material provides the parameters used in the authors’ synthesis experiments, along with additional synthesis and assessment results. The total size of the files is 2.06 MB. Contact yi.ren@alumni.technion.ac.il for further questions about this work.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2017.2750419

A successful patch-based image restoration algorithm that tackles the above concern to some extent is the EPLL [6], which builds upon a parametric local Gaussian Mixture Model (GMM). Concretely, EPLL promotes the patches of the final restored image to comply with the local model, by alternating between (1) restoring each patch taken from the previous estimate according to the local prior, and (2) reconstructing the image by applying a patch-averaging step. This iterative process results in a significant reduction of artifacts, leading to state-of-the-art restoration. Incidentally, a similar concept to EPLL has been presented in [18], in the context of texture synthesis. More specifically, while [18] relies directly on example patches from a given texture image (rather than a parametric model), its main objective and way to achieve it are the same as in the EPLL: The goal is getting a synthesized image in which every patch has a close match to the example set, and the way to get this is by a similar iterative process.

In this paper we propose to leverage the example-based model and the EPLL framework, putting forward a multi-scale image generation algorithm, which is intuitive and competitive with the state-of-the-art, while being fully interpretable. Given a very small seed image (e.g. 4×4 pixels) as an input, we first upscale it by factor of 2 in each axis using a simple interpolation method (e.g. bilinear). Then, new content is *hallucinated* via an example-based spatially varying local priors, which are plugged to the EPLL scheme. Clearly, this process can be repeated several times, until the image reaches the desired size. Following previous work [13]–[15], [23]–[27], we test the ability of the proposed algorithm to synthesize images of both handwritten digits (based on MNIST [39]) and human faces, showing that we achieve results comparable to the state-of-the-art. Our proposed algorithm bears some similarities to the texture-synthesis work reported in [18]. We map clearly the differences between [18] and our algorithm in Section III-E. Furthermore, the similarities between our work and PatchMatch [31] and Visio-ization [28] are also explained in detail in Sections III-F and III-G, respectively.

Our second contribution is a framework for assessing the performance of an arbitrary synthesis machine. Previous work has evaluated the performance of the synthesis algorithm by computing the Log-Likelihood (LL) of the test images in the probability distribution associated with their parametric synthesizer [14], [15], [23]–[25]. This measure indeed indicates the generalization power of the trained model. However, since the test set contains only real images, the measure does not necessarily assign low probabilities to undesired images (such as the blank image and ones containing artifacts). Thus, the LL value obtained on the test images does not indicate whether such failed images may be generated. Therefore to evaluate the visual quality of the synthesis outcome, it is indispensable to assess the LL directly on the generated images, independently from the synthesis model or in a non-parametric manner.

However, moving from LL evaluation of the test images to the LL of synthesized ones is not sufficient, as this measure does not reveal the whole picture. Suppose we have a trivial synthesis machine that memorized training images and provide them as its output. In this case, the generated images are of

high LL measure, and they are *spread* evenly over the training images, but *without introducing any new content*. Furthermore, consider another machine producing always the same image which is both novel and of high quality. Here the synthesized images have better *originality*, but are concentrated on a single point, ignoring the distribution of the training images. In these two cases the LL will indicate an excellent performance while the synthesis is both deterministic and degenerated. Motivated by this observation, we propose a complete assessment framework combining three aspects of the performance: the likelihood, the originality and the spread of the synthesized images. Both numerical measures and visualization tools are presented to evaluate these three aspects. Our experiments show that the proposed synthesis algorithm results in high quality images with good likelihood and originality, along with a reasonable spread.

This paper is organized as follows: In Section II we review the EPLL algorithm as we will rely on it in the proposed synthesis scheme, which is described in Section III with all its ingredients. Then, in Section IV we provide various synthesis results of digits and faces that demonstrate the effectiveness of the proposed scheme. Section V presents our way for assessing the goodness of the results, and provides the assessment results of our synthesis outcome, compared to the state-of-the-art. In Section VI we summarize the paper and outline future research directions.

II. BACKGROUND: EPLL VIA ADMM

In this section we present the EPLL with a slightly changed form — rather than relying on the quadratic half-splitting strategy used in [6], we base our derivations on the more accurate Alternating Directions Methods of Multipliers (ADMM) method [40] which has been shown to be very effective in numerous applications.

The core idea behind the EPLL is to regularize the given inverse problem by specifying a prior on patches only. The regularization term averages over the individual patch priors in a form of an expectation, explaining the name given to this method. The restoration problem suggested by EPLL can be formulated as follows:

$$\min_{\underline{X}} \left\{ \frac{\lambda}{2} \|\mathbf{H}\underline{X} - \underline{Y}\|_2^2 - \sum_{i \in I} \log P_i(\mathbf{R}_i \underline{X}) \right\} \quad (1)$$

where

- \mathbf{H} is the matrix representing the degradation operator.
- \underline{Y} is the input image, which is assumed to be corrupted by \mathbf{H} .
- \underline{X} is the output restored image of size $w \times h$.
- I is the set of locations of the fully overlapped patches of size $n \times n$ in image \underline{X} .
- $\mathbf{R}_i \in \mathbb{R}^{n^2 \times (w \cdot h)}$ is a patch extraction operator for each patch location i , e.g. $\mathbf{R}_i \underline{X}$ is the n^2 -dimensional patch of \underline{X} at location i (being reordered as a vector).
- $P_i(\underline{z}) : \mathbb{R}^{n^2} \rightarrow [0, 1]$ represents the prior probability function of a patch \underline{z} located at position i . As we use local patch models, this prior may be space- (and scale-) dependent.

The optimization problem in Equation (1) is equivalent to

$$\min_{\underline{X}, \underline{z}} \left\{ \frac{\lambda}{2} \|\mathbf{H}\underline{X} - \underline{Y}\|_2^2 - \sum_{i \in I} \log P_i(\underline{z}_i) \right\}$$

$$\text{s.t. } \mathbf{R}_i \underline{X} - \underline{z}_i = 0 \quad (\forall i \in I), \quad (2)$$

where $\underline{z} = \{\underline{z}_i\}_{i \in I}$ are the auxiliary variables representing the restored patches. Using the scaled form of ADMM [40], the Lagrangian of the above problem can be written as

$$L(\underline{X}, \underline{z}, \{\underline{u}_i\}_{i \in I}) = \frac{\lambda}{2} \|\mathbf{H}\underline{X} - \underline{Y}\|_2^2 - \sum_{i \in I} \log P_i(\underline{z}_i) + \frac{\rho}{2} \sum_{i \in I} \|\mathbf{R}_i \underline{X} + \underline{u}_i - \underline{z}_i\|_2^2 \quad (3)$$

in which \underline{u}_i plays the role of a Lagrange multiplier for the i^{th} constraint (i.e. $\mathbf{R}_i \underline{X} - \underline{z}_i = 0$), and ρ is the weight of the corresponding penalties. This problem can be solved by the following iterative and alternating steps:

- z -step:

$$\underline{z}_i^{k+1} = \arg \min_{\underline{z}} \left\{ \frac{\rho}{2} \|\mathbf{R}_i \underline{X}^k + \underline{u}_i^k - \underline{z}\|_2^2 - \log P_i(\underline{z}_i) \right\} \quad (4)$$

for all $i \in I$. This local step is formulated as a MAP estimation of \underline{z}_i given the measurement vector $\mathbf{R}_i \underline{X}^k + \underline{u}_i^k$.

- X -step:

$$\underline{X}^{k+1} = \arg \min_{\underline{X}} \left\{ \frac{\lambda}{2} \|\mathbf{H}\underline{X} - \underline{Y}\|_2^2 + \frac{\rho}{2} \sum_{i \in I} \|\mathbf{R}_i \underline{X} - \underline{z}_i^{k+1} + \underline{u}_i^k\|_2^2 \right\}$$

$$= \left(\lambda \mathbf{H}^T \mathbf{H} + \rho \sum_{i \in I} \mathbf{R}_i^T \mathbf{R}_i \right)^{-1} \cdot \left(\lambda \mathbf{H}^T \underline{Y} + \rho \sum_{i \in I} \mathbf{R}_i^T (\underline{z}_i^{k+1} - \underline{u}_i^k) \right). \quad (5)$$

This step merges the estimated patches together to form the global image, while taking into consideration the corrupted image \underline{Y} and the degradation operator \mathbf{H} . This step is essentially built upon a patch averaging procedure, coupled with a Wiener restoration [6].

- u -step:

$$\forall i \in I, \underline{u}_i^{k+1} = \underline{u}_i^k + (\mathbf{R}_i \underline{X}^{k+1} - \underline{z}_i^{k+1}), \quad (6)$$

which updates the Lagrange multipliers vectors according to the Augmented Lagrange method [40].

We should note that [6] chose to apply GMM for the patch prior P_i at all locations $i \in I$, while in this work we will use an example-based prior — see Section III-B for more details.

III. PATCH-BASED IMAGE SYNTHESIS ALGORITHM

The core challenge of applying directly an existing image restoration scheme for synthesis is the lack of randomness in the output, which we overcome by a multi-scale synthesis scheme described in Section III-A, in which we iteratively upscale a given small image via randomized EPLL.

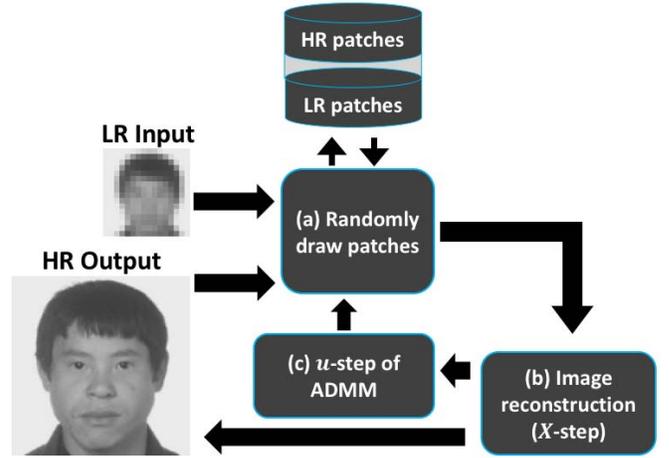


Fig. 1. Block diagram of the *Layer-Synthesis* process, consisting of 3 iterative steps: (a) Randomly sample high-resolution patches from an example-based patch prior (as detailed in Section III-B) based on the low-resolution (LR) input and the current estimate of the high-resolution (HR) output. The sampling is described in Section III-A. (b) Given the sampled patches, reconstruct the image according to the X -step in Equation (5), leading to a new estimate of the HR result. (c) Update the Lagrangian multipliers $\{\underline{u}_i\}_{i \in I}$ according to Equation (6). Once a *Layer-Synthesis* is accomplished, the HR output serves as the LR input of the *Layer-Synthesis* of the next scale, until we obtain a full-sized image. Notice that the initial HR estimate can be obtained by a simple interpolation.

Moreover, it is crucial to choose a traceable patch prior in the EPLL in order to sample sharp and likely patches. We propose for this need a non-parametric example-based prior in Section III-B. In Section III-D we describe several key improvements on our scheme, such as handling the patch overlaps and utilizing patch context.

A. From Restoration to Synthesis

In this subsection we turn to develop the proposed synthesis algorithm. The approach taken is pyramidal, where we suggest repeating the following process: Given a low-resolution (LR) image \underline{Y} , we first upscale it by factor 2 in each axis using a simple interpolation method, leading to the image \underline{X} . Then, the randomized EPLL is utilized to refine the estimation \underline{X} by promoting its patches to comply with a local *spatially varying* patch-prior, while being close to the LR image \underline{Y} . Meanwhile, randomness force is injected into the EPLL (as described below) in order to avoid deterministic results, in contrast to the conventional upscaling/super-resolution methods [36], [41]. The above upscaling process, termed *Layer-Synthesis*, can be repeated several times, starting from an extremely small image containing almost no information (such an image is referred to hereafter as the *seed*), and leading to an image of any desired size, as long as we have adequate local priors to drive the process. This *Layer-Synthesis* procedure is summarized in Figure 1.

We now consider the question of how to convert a super-resolver into a synthesis algorithm with randomness. Given a LR image \underline{Y} and a downsampling operator \mathbf{H} ,¹ the goal of

¹For simplicity, we assume that \mathbf{H} halves the size of the input image in each axis throughout the paper.

Algorithm 1 $\underline{X} = \text{Layer-Synthesis}(\underline{Y}, \{P_i\}_{i \in I})$ **Input:**

- Y : The LR image
- \mathbf{H} : The 2D 2:1 downsampling operator
- n : The height and width of a patch
- I : Set of fully overlapping patch locations
- $\{P_i\}_i$: The patch priors for all patch locations i
- K : The number of iterations of the ADMM
- $\{\lambda_k, \rho_k\}_{k=0}^{K-1}$: EPLL Parameters

Output:

- Twice higher-resolution synthesized image X

- 1: $\underline{u}_i^0 \leftarrow 0$ for $\forall i \in I$;
- 2: $\underline{X}^0 \leftarrow$ interpolation (e.g. bilinear) of Y ;
- 3: **for** $k = 0 : 1 : K - 1$ **do**
- 4: For all patch locations $i \in I$, draw a patch \underline{z}_i^{k+1} random from the conditional distribution given in Equation (7);
- 5: Reconstruct the image \underline{X}^{k+1} from $\{\underline{z}_i^{k+1}\}_{i \in I}$, according to Equation (5);
- 6: For all $i \in I$, update \underline{u}_i^{k+1} using Equation (6);
- 7: **end for**
- 8: $\underline{X} \leftarrow \underline{X}^{K-1}$;

both patch-based super-resolution and synthesis is to obtain a high-resolution (HR) image X such that $\mathbf{H}\underline{X}$ is close to the input \underline{Y} and each patch from \underline{X} is likely under a local patch model. Oftentimes, this process is formulated as an optimization problem, as described in Equation (1), which can be solved efficiently using the EPLL. However, in contrast to the deterministic outcome of a super-resolver, a synthesis method is also expected to generate many different HR images, all being plausible HR versions of Y . Thus, we should allow randomness in our synthesis algorithm in order to diversify the output possibilities. To this end, we change the patch estimation stage in the EPLL (refer to Equation (4) in Section II – the z -step), so that instead of minimizing the function $\frac{\rho}{2} \|\mathbf{R}_i \underline{X}^k + \underline{u}_i^k - \underline{z}\|_2^2 - \log P_i(\underline{z})$ we randomly draw a patch from the posterior distribution of \underline{z} , whose density is given by

$$T(\underline{z}|i, \underline{X}^k, \underline{u}_i^k) = \frac{1}{G_i} \cdot \exp \left\{ \log P_i(\underline{z}) - \frac{\rho}{2} \|\mathbf{R}_i \underline{X}^k + \underline{u}_i^k - \underline{z}\|_2^2 \right\}, \quad (7)$$

where G_i is the partition function, P_i is our example-based patch prior which is described in detail in Section III-B, and \underline{u}_i^k is a scaled dual variable being a by-product of the ADMM. The first term in the exponent encourages the generated patch \underline{z} to align with the local prior by preferring higher $P_i(\underline{z})$ values. The second term enforces \underline{z} to fit to the current estimate of the patch being refined $\mathbf{R}_i \underline{X}^k + \underline{u}_i^k$.

Plugging this randomness force to the EPLL scheme leads to the proposed *Layer-Synthesis* algorithm, which is summarized in Algorithm 1. Armed with this single scale generation step, a seed image can be gradually upscaled by invoking the *Layer-Synthesis* several times, formulating our multi-scale approach as described in Algorithm 2.

Algorithm 2 Multi-scale Synthesis Algorithm**Input:**

- Y : The seed image
- L : The number of up-scalings to perform
- $\{\mathbf{H}_l\}_{l=0}^{L-1}$: The downsampling operators
- n : The height and width of a patch
- $\{I_l\}_{l=0}^{L-1}$: Fully overlapping patch locations at each layer
- $\{P_{i,l}\}_{i,0 \leq l \leq L-1}$: The patch priors per scale and location
- $\{K_l\}_{l=0}^{L-1}$: The number of iterations
- $\{\lambda_{k,l}, \rho_{k,l}\}_{0 \leq k \leq K_l-1, 0 \leq l \leq L-1}$: EPLL Parameters

Output:

- HR synthesized image \underline{X}_0 , up-scaled L times from \underline{Y}

- 1: $\underline{X}_L \leftarrow \underline{Y}$;
- 2: **for** $l = L - 1 : -1 : 0$ **do**
- 3: $\underline{X}_l \leftarrow \text{Layer-Synthesis}(\underline{X}_{l+1}, \{P_{i,l}\}_i)$;
- 4: **end for**

Notice that the proposed method does not create totally arbitrary images as all emerge from a given seed image Y . Using a fixed seed for different synthesis runs, we aim to show the randomness power of our method. Furthermore, as the seed images are very small (4×4 for digits, 8×8 for faces), the images created are expected to be quite diverse. In the experiments, we choose to use seeds taken from known test images, in order to be able to compare the outcomes to the original test image they correspond to. This implies a huge but finite number of possible output images. To break this limit, we can instead use random seeds drawn from a Gaussian or GMM fit to the existing seeds from the dataset. Interestingly, the resulting images are a bit more diverse compared to the ones generated from existing seeds, with the cost of slightly more artifacts.

B. Local Nearest-Neighbor Patch Prior

We now turn to construct the proposed spatially varying example-based synthesis priors, which are plugged to Equation (7). These are formulated as a collection of LR-HR example patch pairs, assembled in dictionaries $D_{LR}^{i,l}$ and $D_{HR}^{i,l}$ for each patch location i and each layer l . The construction of these is elucidated in Section III-B.1. Next, in Section III-B.2 we describe the sampling process from the prior using $D_{LR}^{i,l}$ and $D_{HR}^{i,l}$.

1) *Building Example Dictionaries*: In order to construct the LR-HR patch pair dictionaries $D_{LR}^{i,l}$ and $D_{HR}^{i,l}$, we decompose the training images $\{V_j\}_{j=1}^N$ ($V_j \in \mathbb{R}^{r \times c}$) into Gaussian pyramids [42] (see Figure 2), thereby creating a sequence of layers $V_j^0, V_j^1, \dots, V_j^L$ where

$$\underline{V}_j^0 = \underline{V}_j \quad (\text{original image})$$

and

$$\underline{V}_j^l = \mathbf{H}_{l-1} \underline{V}_j^{l-1}, \quad l = 1, \dots, L.$$

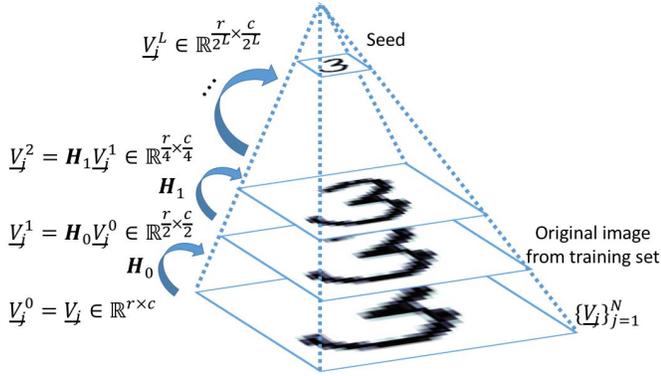


Fig. 2. The building process of the Gaussian pyramid for the image V_j from the set of training images.

The matrices $\mathbf{H}_0, \dots, \mathbf{H}_{L-1}$ are downsampling operators (which apply blur and 2:1 decimation in each axis), as such, each layer $V_j^l \in \mathbb{R}^{(r/2^l) \times (c/2^l)}$ is twice smaller than the previous one $V_j^{l-1} \in \mathbb{R}^{(r/2^{l-1}) \times (c/2^{l-1})}$. Following this rationale, the image $V_j^L \in \mathbb{R}^{(r/2^L) \times (c/2^L)}$ is the smallest one, corresponding to the seed size.

Next, $D_{LR}^{i,l}$ for layer l and location i is the set of LR patches originating from the lower resolution images $\{V_j^{l+1}\}_{j=1}^N$, given by

$$D_{LR}^{i,l} = \{\mathbf{Q}_{\frac{i}{2}, l+1} V_j^{l+1}\}_{j=1}^N,$$

where, following Figure 3, $\mathbf{Q}_{\frac{i}{2}, l+1}$ extracts the patch of size $(n/2) \times (n/2)$ from location $\frac{i}{2}$ of the images $\{V_j^{l+1}\}_{j=1}^N$. Similarly, the set $D_{HR}^{i,l}$ contains the HR patches corresponding to the elements in $D_{LR}^{i,l}$, which are extracted from the larger (higher resolution) images $\{V_j^l\}_{j=1}^N$. This can be formulated by

$$D_{HR}^{i,l} = \{\mathbf{R}_{i,l} V_j^l\}_{j=1}^N,$$

where $\mathbf{R}_{i,l}$ extracts the i -th patch of size $n \times n$ from $\{V_j^l\}_{j=1}^N$, as depicted in Figure 3. Notice that the location $\frac{i}{2}$ is meaningful only for even i 's. Therefore, in synthesis we only treat every other HR patch, instead of all of them.

2) *Sampling Process in Synthesis*: Next, we describe the definition of the patch priors based on $D_{LR}^{i,l}$ and $D_{HR}^{i,l}$, and show how these are used in the sampling process leading to the randomized synthesis effect. Our multi-scale algorithm generates an image by inferring a Gaussian pyramid in the reversed order X_L, \dots, X_0 , from the seed X_L to the full image X_0 . In particular, to generate X_l (larger image) from X_{l+1} (smaller one), the process $\underline{X}_l = \text{Layer-Synthesis}(\underline{X}_{l+1}, \{P_{i,l}\}_i)$ is invoked, where we design the example-based prior $P_{i,l}$ (using $D_{LR}^{i,l}$ and $D_{HR}^{i,l}$ as defined above) to prefer HR patches that fit well to the LR content in \underline{X}_{l+1} . Formally, $P_{i,l}$ is defined only on existing patches from $D_{HR}^{i,l}$, and the probability assigned to the HR patch $D_{HR}^{i,l}(j)$ is dictated by the proximity between its LR version $D_{LR}^{i,l}(j)$ and the LR patch found at location $\frac{i}{2}$

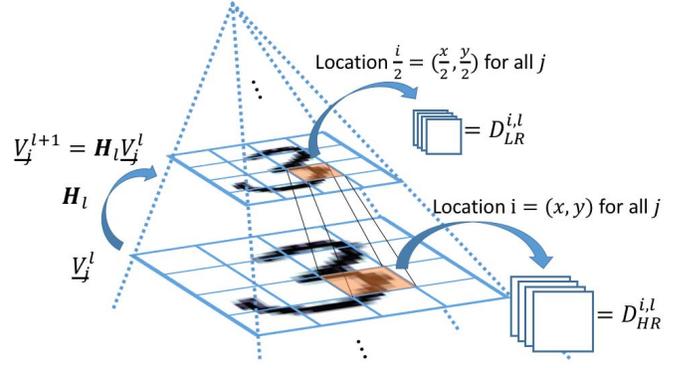


Fig. 3. The building process of the prior $P_{i,l}$ for scale l and for each patch location $i \in I_l$. In the figure $i = (x, y)$ means (x, y) is the central pixel of the HR patch at location i , and $\frac{i}{2} = (\frac{x}{2}, \frac{y}{2})$ means $(\frac{x}{2}, \frac{y}{2})$ is the central pixel of the corresponding LR patch at location $\frac{i}{2}$.

in X_{l+1} , for all $j = 1, \dots, N$:

$$P_{i,l}(D_{HR}^{i,l}(j) | \underline{X}_{l+1}) = \frac{1}{G_{i,l}} \cdot \exp \left\{ -\frac{1}{h} \|\mathbf{Q}_{\frac{i}{2}, l+1} \underline{X}_{l+1} - D_{LR}^{i,l}(j)\|_2^2 \right\}, \quad (8)$$

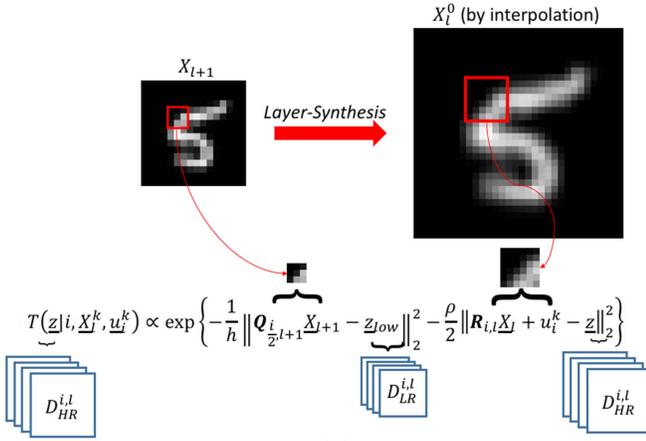
where $G_{i,l}$ is the partition function. Note that we do not sample patches directly from the above prior. Instead, we shall plug it into Equation (7) and draw from the resulting posterior distribution:

$$T(D_{HR}^{i,l}(j) | i, \underline{X}_l^k, \underline{u}_i^k) = \frac{1}{G_i} \cdot \exp \left\{ -\frac{1}{h} \|\mathbf{Q}_{\frac{i}{2}, l+1} \underline{X}_{l+1} - D_{LR}^{i,l}(j)\|_2^2 - \frac{\rho}{2} \|\mathbf{R}_i \underline{X}_l^k + \underline{u}_i^k - D_{HR}^{i,l}(j)\|_2^2 \right\}, \quad (9)$$

where \underline{X}_l^k is the estimation of the final synthesized image at layer l (denoted \underline{X}_l) after iteration k . Now we simply draw one of the patches from $D_{HR}^{i,l}$ to be used in location i , where the probability of each patch is determined by the above posterior. In fact, both the LR and HR measurements are taken into account simultaneously in Equation (9), as illustrated in Figure 4. This corresponds to a patch-matching process where the LR measurement is matched with $D_{LR}^{i,l}$, and the estimated HR content is matched with $D_{HR}^{i,l}$.

Notice that the positive parameter h in Equation (9) controls the randomness of the prior: if $h \rightarrow 0$, then the HR version of the nearest neighbor of $\mathbf{Q}_{\frac{i}{2}, l+1} \underline{X}_{l+1}$ from $D_{LR}^{i,l}$ will be chosen with probability 1, while $h \rightarrow +\infty$ makes the prior uniform over all the patches in $D_{HR}^{i,l}$. Through the iterations of the *Layer-Synthesis*, we can change the parameter h from a large initial value towards a smaller one (see Equation (9)), so as to first allow new information to emerge and then, in later iterations, to make the HR image X_l^{k+1} comply with both the previous estimate X_l^k and the LR version X_{l+1} .

In order to obtain more relevant patches and speed up the algorithm, we sample from the patches corresponding to the nearest neighbors of $\mathbf{Q}_{\frac{i}{2}, l+1} \underline{X}_{l+1}$ (instead of from the whole candidate set $D_{HR}^{i,l}$), according to the probability assigned



(a)

z_{low} $\in D_{LR}^{i,l}$						
z $\in D_{HR}^{i,l}$						
$\log(T(z))$	-6.21	-8.03	-9.31	-10.6	-12.1	

(b)

Fig. 4. (a) The computation of the posterior distribution in Equation (9), from which a patch will be drawn for the refinement of X_l . (b) Example candidates of sampling from the distribution shown in Figure (a). Rightmost are the LR and HR measurements in patch-matching. Notice that the HR candidates are sharper than the current HR measurement, therefore improving the visual quality of the HR image at location i . In addition, as can be seen, the value of $\log(T(z))$ is correlated with the relevance of the match.

by Equation (9). The nearest-neighbor searching functionality required here is provided by the efficient FLANN library [43].

C. Theoretical Issues

In this subsection, we discuss briefly two theoretical issues raised by our algorithm: the distribution of the output and convergence.

Regarding the distribution of the created images, clearly, our goal is to reproduce the true high-dimensional distribution of the images in question, as manifested by the training set. Since this by itself is a daunting task, our algorithm is designed with the hope to get a sampling close to the posterior of the global image statistics we operate on in each scale, while operating locally. Towards this, we replace the full posterior by the expected patch-log-likelihood, which is the core idea behind the EPLL. Therefore, we cannot make theoretical claims regarding the actual distribution created, and thus we offer various empirical evidence in later sections. We note that other work tackling this synthesis goal (e.g. [14], [28]) take the same route and do not give supporting theoretical claims for the sampling obtained.

Regarding the convergence to a fixed point, a simple strategy to guarantee this is to increase the parameter ρ as a function of the ADMM iterations (e.g. as done in [44], and in the EPLL

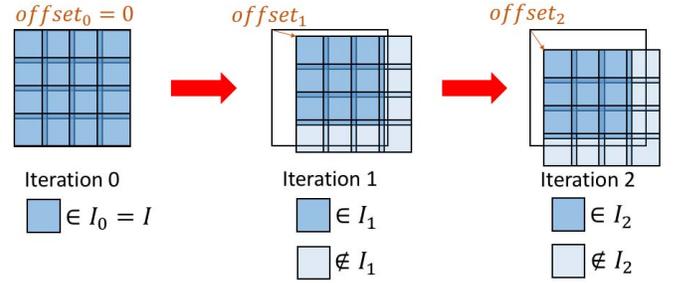


Fig. 5. Illustration of cycle-spinning through iterations for *Layer-Synthesis*. I_k is the set of patches being refined in the k -th iteration. The light-blue patches are not taken into account as they go outside the image support, for simplicity.

work in the context of the variable splitting). However, such a strategy does not guarantee convergence in our case: Once $\rho \rightarrow \infty$ we in fact seek for a patch $z = \mathbf{R}_i X^k + u_i^k$ (See Equation (4)). Following Equation (9), this will not be satisfied unless we will find a HR example in $D_{HR}^{i,l}$ exactly equal to $\mathbf{R}_i X^k + u_i^k$, which is a direct drawback of using a discrete patch prior. As a future work we aim to overcome the convergence problem by replacing the discrete prior with the continuous local GMM while preserving the sharpness of the patches somehow.

D. Extensions

In what follows we present three extensions to the above synthesis scheme so as to improve its outcome.

1) *Cycle-Spinning*: We observe a trade-off between the visual quality and the richness of the generation result when choosing the overlap between patches. Concretely, fully overlapping patches (as in Algorithm 1) avoid artifacts due to the averaging of many samples, while leading to blurry synthesized results which are similar to each other. On the other hand, with small overlaps we obtain sharp and diverse images, with the cost of artifacts that appear especially along the patch-borders. To bypass this trade-off, we suggest using the “cycle-spinning” technique [45] as follows. According to Figure 5, the leftmost image (corresponding to iteration 0) is reconstructed from the patches at locations $i \in I_0 = I$ having small overlaps with each other (meaning that I no longer contains all the patch locations). Then in the next iteration, as appears in the central part of Figure 5, we refine the previous estimate, this time by (randomly) restoring the patches at locations $i \in I$ shifted by offset_1 (the set of the shifted locations is denoted by I_1). Clearly, one can repeat this process with different offsets, thereby improving the visual quality due to the iterated refinement steps, while preserving the sharpness and richness of the result thanks to the small overlaps. Examples of images generated with and without cycle-spinning are shown in Figure 6 to depict the necessity of this step. A down side to this process is the need to hold local priors to all the candidate locations and invoke them in the restoration appropriately.

2) *Neighbor Search Window*: As the training images may be slightly unaligned (e.g. digits and faces), in order to achieve more diverse results, the matching process should seek for

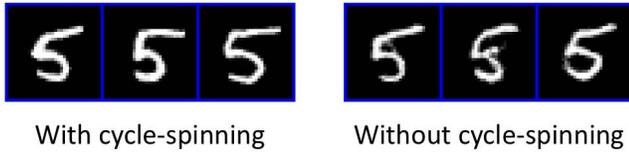


Fig. 6. Examples of synthesized digit images with and without cycle-spinning. The images created with cycle-spinning are in general of better quality and contain less artifacts.

similar patches in the database that are not only located in the same coordinates, but also in a small neighborhood around it. Formally, we propose to enlarge the dictionaries $D_{HR}^{i,l}$ and $D_{LR}^{i,l}$ built in Section III-B.1 to contain patches at location i (as before), as well as from the neighboring locations:

$$D_{HR}^{i,l} = \bigcup_{i' \text{ neighbor of } i} \{\mathbf{R}_{i',l} \mathbf{V}_j^l\}_{j=1}^N,$$

$$D_{LR}^{i,l} = \bigcup_{i' \text{ neighbor of } \frac{i}{2}} \{\mathbf{Q}_{i',l+1} \mathbf{V}_j^{l+1}\}_{j=1}^N,$$

where i' is a neighboring position of i if the offset between i and i' is less than a given number of pixels in both x - and y -axes. This number is denoted as the neighborhood-size $_l$ for layer l . Naturally, since digit and face images tend to be more “aligned” in lower resolution, the window is increased together with the image size (see the tables of parameters in the supplementary material).

In the extreme case, one could use all the patches in the images X^l (larger) and X^{l+1} (smaller) to form the dictionaries $D_{HR}^{i,l}$ and $D_{LR}^{i,l}$, respectively. This leads to a spatially invariant local model (in contrast to the spatially varying one that is defined above), which gives no influence to the locations of the patches. In coherent set of images as treated in this work, this approach is necessarily inferior.

3) *Patch With Context*: When choosing the patch size to work with, the trade-off between image quality and originality arises: larger patch size leads to better visual quality, while limiting the richness of the generated images, and vice versa. The reason of this tendency is the limited size of the training set from which we can draw examples, along with the curse of dimensionality which implies that as patches get larger, relevant neighbors are becoming scarce. To cope with this, the definition of distance between patches can be extended to take the surrounding area into account, as depicted in Figure 7. Intuitively, as the patch associated with its context contains wider-range information, it will ensure that the high-level structure of the generated images will be more realistic, without sacrificing its originality. This approach is inspired by the recent work of Con-Patch [46], which is shown to lead to better patch models. We stress here that this is not equivalent to working with bigger patches.

Later in the experiments, we use two kinds of contexts: the square context and the horizontal one, as illustrated in Figure 7. The square context is used for both digit and face synthesis, while the horizontal context is used for faces to improve their symmetry.

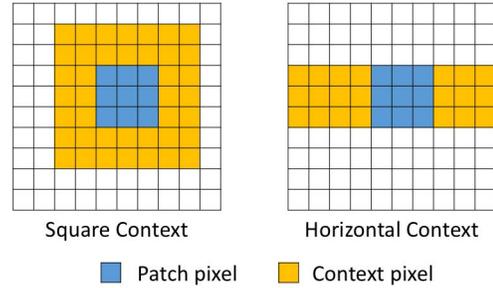


Fig. 7. The square context and the horizontal context of an image patch.

E. Comparison to [18]

While the work reported in [18] addressed a different task of texture-synthesis, it bears some similarities to the above described algorithm, due to the reliance on patch-matching, the operation in multi-scale and more. Here are the key differences between the two works:

- This work deals with image synthesis in a more general way. As such, our treatment relies on local priors that are richer than the representation practiced in [18].
- While in [18] randomness and originality of the result are not the prime goals, they are central in our scheme. In [18] only the initialization stage is random, while in our work the randomness is used in every patch-matching stage.
- The energy functional and minimization method in [18] are similar to the one posed by the EPLL (half-quadratic splitting), and therefore similar to ours as well. However, we introduce the ADMM, making the formulation more “well-formed”.
- Both methods use multi-scale pyramids, but [18] practices also a sweep over the patch sizes. We chose to avoid such feature in our algorithm, due to the anticipated problem of losing relevant neighbors in the consequent search, and the fear of getting large portions of trained images copied to the synthesized image.
- Both these works use small overlap between the patches and for the same reasons. our algorithm adds to this a cycle-spinning shift of the patch positions to avoid border artifacts. Perhaps [18] overcomes these artifacts due to the use of different patch-sizes.
- Last but not least, we accompany our algorithm with a framework for assessing the performance of whole image synthesis, as we will outline in Section V.

As a side note, the above similarities shared by our work and [18] indicate the potential capability of our method in texture-synthesis, which should be further validated by more experiments.

F. Comparison to PatchMatch [31] and PatchTable [8]

Our method also shares some apparent similarities with PatchMatch, in that both deal with patches and randomized selection. However, they are fundamentally different in the following ways:

- The goal of PatchMatch is to approximate the nearest neighbor search between a database image and a query

image, demonstrated on tasks like retargeting, inpainting and reshuffling; On the other hand, our work aims to leverage existing patch (i.e. local) models in order to synthesize different images from “thin air”.

- In PatchMatch, the randomness is used in initialization and the random search step merely to reduce computational cost, whereas in our work the randomness is used to draw a patch from a posterior distribution, which is at the very core of our interest of generating *different* images.

PatchTable is a follow-up of PatchMatch, which pre-computes some data structures for the database image, so that the nearest neighbor search for each individual query image is even faster than PatchMatch. As such, it is able to perform more computationally intensive tasks such as artistic video stylization, light-field super-resolution and multi-image editing. An interesting future work is to use PatchTable as our nearest-neighbor search engine (by stitching all the training images together as a huge database image) in the local patch prior described in Section III-B. This should result in a significant speed up of our image generation process.

G. Comparison to Visio-Lization [28]

The work reported in Visio-ization tackles the problem of human face synthesis, and is highly related to our approach. The main common points and differences of these methods are as follows:

- Both methods use datasets of face images aligned by key points (eyes, mouth etc.).
- Both methods create images in a coarse-to-fine manner to ensure the global consistency of the faces. Our algorithm achieves this by generating a Gaussian pyramid from top to bottom, whereas [28] first draws a blurry face image from a GMM model, then add fine details to it.
- Both approaches use a spatially varying (non-stationary) local patch model for the generation. However, our method draws the patches at each location simultaneously and merges them by EPLL and ADMM (so that all patches are treated equally), while [28] draws the patches sequentially (similarly to [36]) and merges them via Poisson image editing [47].
- Our method may produce asymmetric faces (see Figure 12 for example), while [28] enforces the patches at symmetric locations of the face to come from the same source image, in order to completely avoid this.

IV. SYNTHESIS EXPERIMENTS

We now turn to present the results obtained by the described algorithm, synthesizing digits and face images.

A. MNIST Digits

Following the previous work in handwritten digit synthesis [13]–[15], [23]–[25], we use the MNIST [39] dataset in our experiments, which includes 60000 training images and 10000 test ones of size 28×28 pixels, padded to 32×32 before synthesis. The test images are downsampled to 4×4 ,

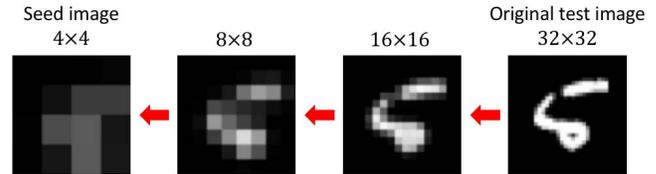


Fig. 8. Downsampling of one original MNIST test image to the size of a seed.

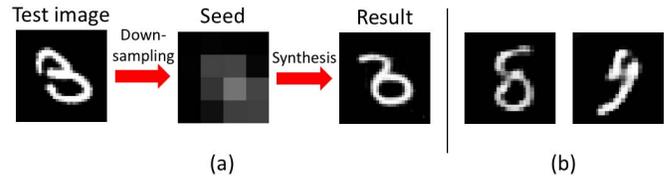


Fig. 9. (a) Example of MNIST synthesis using a patch model built with all digits together. The result is not a digit, and looks like the upper half of a digit “3” merged with the lower half of a digit “8”. Notice that the LR seed image is indistinguishable from “8”. (b) Examples of “non-digit” synthesized by DRAW [15], which also models all digits together.

which serve as the seed images (see Figure 8). It is worth noting that we choose to synthesize each type of digit using only the training images of the same digit (with the very same set of parameters), while existing work commonly builds the synthesis model with all kinds of digits together. Since the classification of MNIST digits is a very well studied problem for which highly accurate classifiers are available, we consider the classification as a simple preprocessing step. When we ignore the class-specific construction of the example set and use all the different digits together, images of “non-digit” might be produced as shown in Figure 9. A similar behavior, presented in the same figure, is observed when generating images by DRAW [15].²

The synthesis consists of a pyramid of 4 layers: 4×4 (seed) $\rightarrow 8 \times 8 \rightarrow 16 \times 16 \rightarrow 32 \times 32$. At each layer, our local example based patch prior (see Section III-B) is used, along with the proposed extensions (as detailed in Section III-D).

The parameters (provided in the supplementary material) are chosen to achieve a good balance between the visual quality and the richness of the generated images. Although there seems to be a lot of parameters to select, each parameter has a clear meaning, and thus its influence is explicit and easy to follow. Moreover, some basic principles are followed, making the overall scheme simple to tune:

- Each of the offsets are used once in a somewhat arbitrary order;
- Most of the parameters (h , λ and ρ) are monotonic through the iterations of the same layer;
- The value of ρ increases through the iterations, so that the high resolution patches conform with the current synthesized image more and more, and finally stabilize.

All these imply that tuning these parameters is a relatively easy and straight-forward task. On the other hand, we do note that a possible future direction would be to tune these parameters automatically (e.g. by evaluating the quality of the generation

²We used an unofficial implementation of DRAW, given in [48].



Fig. 10. Example MNIST digits synthesis results produced by the suggested method. Green frame: original test image used to create the seed. Blue frame: different generation results with the seed in the same row. Red frame: the nearest neighbor of the synthesized image on its left from the training set. The distance between each generated image (blue frame) and its nearest neighbor (red frame on its right) is indicated below the blue frame, and the distance between this neighbor and its own nearest neighbor from the training set is indicated under the red frame for comparison. More generated digits are shown in the supplementary material.

result at each layer using our assessment framework), in order to be able to apply this method more easily to other datasets, while gaining a deeper understanding of the roles played by these parameters.

Figure 10 illustrates various generation results of all kinds of digits. As can be seen, the generated images have high visual quality. In addition, the results are different both from their nearest neighbor in the set of training images and the HR version of the seed image, i.e. our method successfully generated new images that are non-existing in the training set (good originality). Furthermore, as can be observed, we synthesize very diverse digits that originate from the same seed thanks to the randomness force, showing the effectiveness of our synthesis algorithm, and the good spread obtained over the training images. Given a seed image, a diagram of the synthesis process is depicted in Figure 11, showing how three different generations evolve throughout the layers.

As observed in Figure 10, there are no visible artifacts in the generated images. Nevertheless, we encountered several synthesis failures, which occur rarely, as shown in Figure 15 (second row). It is important to note that such failures can

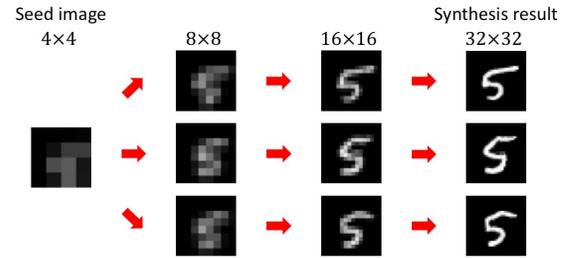


Fig. 11. Different runs of synthesis of the digit “5” using the same seed image.

be discovered and explained by the LL measure, as described later in Section V-D.

B. Faces

Compared to digits, images of human faces are much more challenging to model and synthesize, as they contain richer details and very long range structures (e.g. the two ears should be consistent). Furthermore, humans are extremely sensitive to small unnaturalness of faces, making the generation task even more demanding. We emphasize that the goal of this experiment is to demonstrate the generality of our method by applying it to a completely different category of images with few adaptations, rather than creating perfect and realistic synthetic face images.

In this experiment, we generate faces using a dataset of grayscale human faces from passport photos [49]. This dataset consists of 4500 example images and 500 test ones, all aligned by feature points (e.g. eyes and tip of nose). The alignment is done as in [49], where the feature points of each face are first located automatically, and then moved to the aligned locations by warping the triangles formed by these points using affine transformations. All the images (of size 221×179) are cropped to 179×179 and then resized to 128×128 . The synthesis consists of a pyramid of 5 layers: 8×8 (seed) $\rightarrow 16 \times 16 \rightarrow 32 \times 32 \rightarrow 64 \times 64 \rightarrow 128 \times 128$. We use the very same algorithm as in the digit synthesis, but with different parameters (see the tables of parameters in the supplementary material) due to the disparity between these two classes of images.

Example of generated faces are depicted in Figure 12. Visually, these synthesized images are realistic, containing fine details, and do not have severe artifacts. Furthermore, the ability of the proposed method to generate different faces from the same seed image is demonstrated in Figure 13. This serves as an indication to effective randomness, leading to both rich and high quality results. As can be seen, despite the local nature of the proposed method, we obtain pleasant and realistic results. The symmetric property that is unique to faces is achieved by leveraging the horizontal context. As demonstrated, the eyes and ears in the generated images are aligned and have similar shape per instance. Minor asymmetry in some generated faces is still observed, marked by the symbol Δ in Figure 12.

Differently from digit synthesis, in the case of faces, patches from the same training image tend to appear together in the generated result, forming coherent parts of a face

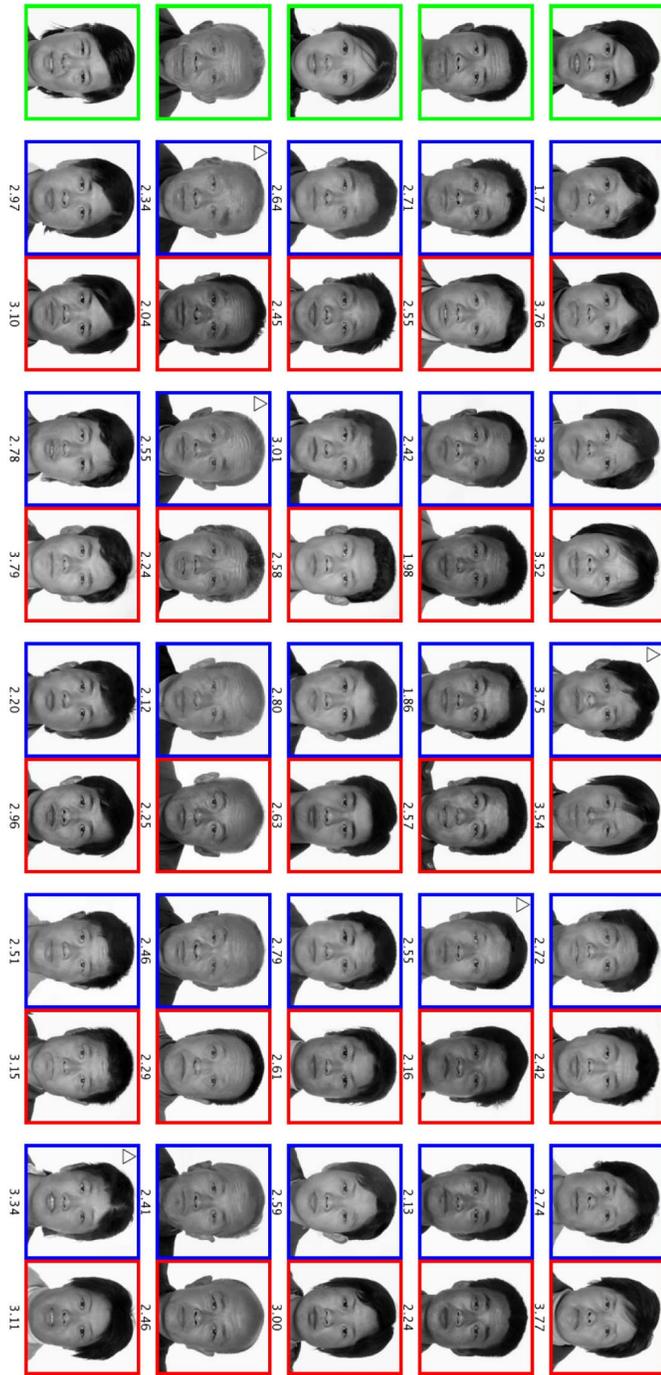


Fig. 12. Example results of aligned faces synthesis. Refer to the caption of Figure 10 for the meaning of the color frames and the numbers. The distance used for searching the nearest neighbor takes only the central pixels into account, as they are the most informative part in the face, while the boundary (e.g. the hair and the clothes) has large variation. The symbol Δ marks the faces that are slightly asymmetric.

(e.g. eye, nose; see Figure 14 for example). This behavior is similar to irregular patch stitching [17] in texture synthesis. This phenomenon leads to appealing results, but may indicate that more training examples are necessary, as we do not have enough freedom in combining the patches. On the down side, the patch stitching might deteriorate the originality of the generated images. However, we will show that the proposed algorithm does not suffer from this limitation in the assessment

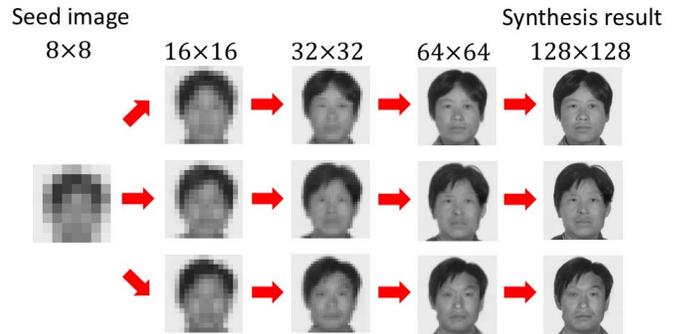


Fig. 13. Different runs of aligned face synthesis using the same seed image.

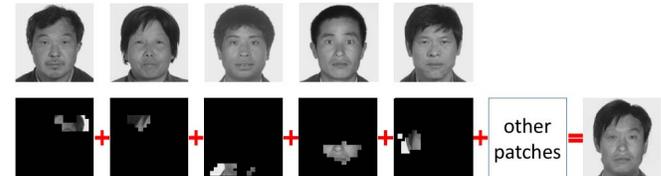


Fig. 14. The result of one aligned face generation, shown as the stitching of patches from different images. First row: The first five training images contributing the most patches in the generated image. Second row: The patches these five images provide, and how they sum up to the synthesized image (with other patches). Notice that the new face is visually very different from the five faces on the left.

of the synthesized faces (see the supplementary material). Notice that it is also possible to add a penalty term in the posterior distribution (Equation (9)) in order to discourage patches from the same source image to be drawn (therefore an order must be specified for the patches). However, we decide not to do so for two reasons: (1) simplicity; (2) we hope to treat all the patches as equally as possible.

V. ASSESSMENT OF SYNTHESIS PERFORMANCE

As we have seen in Section IV, our synthesis algorithm generates visually appealing and realistic digits and human faces. However, to assess the performance of our synthesis machine in a complete fashion, it is necessary to test if the generated images and the training ones represent the same probability distribution. As mentioned in Section I, evaluating the LL for either the test images or the generated images is insufficient for this task, and we suggest a complete assessment framework combining the following three indispensable components: the LL (Section V-A), the originality (Section V-B), and the spread (Section V-C). Then, we apply this framework to our digits synthesis results in Section V-D, showing the strength and effectiveness of the proposed synthesis algorithm. The very same assessment is also applied to the synthesized faces, of which the results are shown in the supplementary material.

When comparing the scores of different generation methods, we believe that the three measures have the following importance order:

$$\text{LL} \rightarrow \text{originality} \rightarrow \text{spread.}$$

In words, if the LL of the generated images is exceedingly low, then they are of low visual quality or even meaningless,

therefore the synthesis does not produce “valid” images, and there is no point to evaluate the originality and spread. Similarly, if the generated images are too similar to the training ones (low originality), then the synthesis power is not effective, failing to fulfil one of the principal objectives of the process. In this case, a good spread score is misleading without high originality, as the generation results may be mostly replicated from the training set.

A. Log-Likelihood

The goal of the LL measure is multiple: (1) to show that the synthesis algorithm minimizes the EPLL objective function (up to some extent due to the randomness force); (2) to ensure the generated images have good visual quality, and detect the failure cases; (3) to have a performance value which is comparable to previous work, for completeness.

We define the likelihood of an image X to be similar to the EPLL objective function as introduced in Equation (1), which we aim to minimize during the synthesis process. Formally, we start by defining the following:

$$LL_{pixel}(X) = \frac{1}{\text{Tr}(\sum_{i \in I} \mathbf{R}_i^T \mathbf{R}_i)} \sum_{i \in I} \log P_i(\mathbf{R}_i \underline{X}), \quad (10)$$

where P_i is the patch prior used in synthesis, and $\sum_{i \in I} \mathbf{R}_i^T \mathbf{R}_i$ is a diagonal matrix that counts for each pixel the number of different estimates emerging from the overlapping patches in I . As such, the term $1/\text{Tr}(\sum_{i \in I} \mathbf{R}_i^T \mathbf{R}_i)$ translates the sum over the LL of the patches $\mathbf{R}_i \underline{X}$ to the expected LL per pixel, denoted by $LL_{pixel}(X)$. Thus, the estimate of the LL of the whole image is given by

$$LL(X) = |\underline{X}| \cdot LL_{pixel}(X), \quad (11)$$

where $|\underline{X}|$ denotes the number of pixels in X . Notice that this formulation is based on the EPLL assumption that all the patches $\mathbf{R}_i \underline{X}$ in the image X are independent (even though they may be overlapping). As such, we sum their LL without conditioning.

While we aim at evaluating the LL of patches $\mathbf{R}_i \underline{X}$ that might not exist in the example patch database (as in Equation (10)), our discrete priors are defined only on the existing patches in a dictionary (see Section III-B and Equation (8) therein). Therefore, we suggest a continuous variant of the proposed discrete priors, formulated as a Parzen window [50] with Gaussian kernel:

$$P_i(\underline{x}) = \frac{1}{|D_{HR}^{i,0}|} \sum_{\underline{y}_j \in D_{HR}^{i,0}} \frac{1}{(2\pi\sigma^2)^{n^2/2}} \exp\left\{-\frac{\|\underline{x} - \underline{y}_j\|_2^2}{2\sigma^2}\right\} \quad (12)$$

where $D_{HR}^{i,0}$ is the set of HR example patches, σ^2 stands for the window width, and n^2 denotes the number of pixels in a patch. Notice that we assume the covariance matrix of each Gaussian has full rank (i.e. equal to $\sigma^2 I$). Nevertheless, this full rank assumption is not valid in general, especially when considering background patches in both digits and faces. Specifically, all the patches in the background of the digits are

TABLE I

AVERAGE LL VALUES OF *Test* IMAGES OF THE MNIST DATASET (THE LL VALUES OF THE *Generated* IMAGES ARE PROVIDED LATER IN SECTION V-D). LARGER (CLOSER TO ZERO) IS BETTER. NOTE THAT THE OTHER METHODS USE THE SAME LL FUNCTION FOR ALL DIGITS, WHEREAS OUR MEASURE USES DIFFERENT PRIORS FOR EACH TYPE OF DIGIT. THE SUGGESTED MEASURE IS COMPARED WITH VARIOUS OTHER METHODS IN THE SUPPLEMENTARY MATERIAL

Model	LL Test
DBM 2hl [24]:	≈ -84.62
EoNADE-5 2hl (128 orderings) [23]:	-84.68
DRAW [15]:	≥ -80.97
Diagonal BiLSTM (7 layers, $h = 16$) [14]:	-79.20
suggested measure	-82.52

totally flat, i.e. the term $\|\underline{x} - \underline{y}_j\|_2^2$ is zero for all j and for the totally flat x , leading to $P_i(\underline{x}) = 1/(2\pi\sigma^2)^{n^2/2}$. Clearly, in this special case we expect to obtain $P_i(\underline{x}) = 1$. This is a direct consequence of the full-rank assumption, which is invalid for all σ values different than $1/\sqrt{2\pi}$. To cope with this singularity, one can suggest to estimate the rank of each covariance matrix (in the example above the rank should be 0), however this approach raises various other difficulties. Another possible solution is to simply choose $\sigma = 1/\sqrt{2\pi}$ so that the actual rank of the local covariance has no importance since $2\pi\sigma^2 = 1$. As a simple justification, we can view this σ as representing a Gaussian noise level. In fact, $\sigma = 1/\sqrt{2\pi}$ corresponds to a noise level of 17/255 per pixel for 6×6 patches (digit) and 13/255 for 8×8 patches (face), which is reasonable. In the supplementary material we provide more justifications for this value of σ .

Table I provides the average LL value computed using Equations (10), (11) and (12) on the *test* images of MNIST dataset. This table also lists the LL measure obtained by previous works. As can be inferred, our LL value is close to the state-of-the-art, indicating that the MNIST images are modeled well by our example-based local priors. Surprisingly, despite the simplicity and traceability of our approach, it is comparable to previous work which tends to be more complex.

As a closing remark, we emphasize that the above LL measure is utilized to evaluate the quality of the *generated* images (see Section V-D), instead of the *test* ones as done in Table I, which merely serves to align our modeling of the LL with the previous work.

B. Originality

In this subsection we propose a score that measures the originality/novelty of the generated images, which, in particular, ensures that the generated images are not simple replications of the training set. Intuitively, we can measure the distance of a generated image X to the training images used for synthesis, denoted by the set $X_T = \{X_i | i \in T\}$. One can conclude that if this distance is very small, the originality of the image X is limited. Formally, this distance is defined by

$$D_G(X, X_T) \equiv \min_{X_i \in X_T} \|X - X_i\|_2.$$

Consequently, the nearest neighbor is denoted by

$$X_{NN} = \text{NN}(X; X_T) \equiv \arg \min_{X_i \in X_T} \|X - X_i\|_2$$

Clearly, if the distance is extremely small the generated image is not novel as it is exceedingly similar to an existing example. However, a natural question that arises is how large this distance should be so one can truly consider this image as a new one. The answer we suggest is to compare $D_G(X, X_T)$ with the following distance:

$$D_T(X, X_T) \equiv \min_{X_i \in X_T \setminus \{X_{NN}\}} \|X_{NN} - X_i\|_2,$$

where $A \setminus B$ denotes the set subtraction operation, and $X_{NN} = \text{NN}(X; X_T)$. The above measures how close the training image X_{NN} is to its own nearest neighbor in the database X_T (which is different from itself). If $D_G(X, X_T) \geq D_T(X, X_T)$ we conclude that the generated image X is indeed novel. On the other hand, $D_G(X, X_T) \ll D_T(X, X_T)$ implies that the synthesis algorithm does not hallucinate new data. This brings us to the definition of originality measure, given by

$$\text{Originality}(X; X_T) = \frac{D_G(X, X_T)}{D_T(X, X_T)}.$$

Thus, to quantify the overall originality of the whole set of generated images $X_G = \{X_i | i \in G\}$, we simply average the originality of all the individual generated images X_i :

$$\text{Originality}(X_G; X_T) = \frac{1}{|X_G|} \sum_{i \in G} \text{Originality}(X_i; X_T).$$

We also propose a visual illustration of this measure, in which the 2-dimensional points ($D_G(X_i, X_T)$, $D_T(X_i, X_T)$) are plotted for each image $X_i \in X_G$ (refer to Figure 17 as an example to such visualization that compares our generation to the results of DRAW).

C. Spread

Consider a scenario in which a synthesis algorithm produces one single image that has high LL and originality. In this case, the overall performance is poor due to the lack of richness and diversity of the results. Motivated by this, we suggest a measure of spread to ensure that the generated images are as “spread out” as the training ones.

We suggest to assess the spread of the generated images mainly based on the t-SNE [51] unsupervised non-linear embedding. The motivation is that t-SNE is widely used and produces state-of-the-art result in MNIST visualization, in which the different types of digits are visually separated into different clusters (while other popular embedding techniques such as Isomap [52] and LLE [53] do not). Therefore, we believe that t-SNE is also able to reveal considerable inconsistencies between the training images and the generated ones, if any.

Next, we review briefly the main principle of the t-SNE applied on an image set $X_D = \{X_i | i \in D\}$. In this technique, the probability of the image $X_j \in X_D$ being the neighbor of $X_i \in X_D$, denoted by $p_{j|i}$, is defined as:

$$p_{j|i} = \frac{\exp(-\|X_i - X_j\|_2^2 / 2\sigma_i^2)}{\sum_{k \in D, k \neq i} \exp(-\|X_i - X_k\|_2^2 / 2\sigma_i^2)} \quad (13)$$

where σ_i is chosen such that

$$H(P_i) = - \sum_k p_{k|i} \log p_{k|i}$$

has a given fixed value for all $i \in D$. In fact, $p_{j|i}$ defines a weighted neighborhood of the point X_i , i.e. any image $X_j \in X_D$ is neighbor of X_i with weight $p_{j|i}$. Next, for the actual embedding task, denote by y_i the 2-dimensional embedded point of X_i , and $q_{j|i}$ the neighborhood probability for the embedded points, similarly to $p_{j|i}$:

$$q_{j|i} = \frac{\exp(-\|y_i - y_j\|_2^2)}{\sum_{k \in D, k \neq i} \exp(-\|y_i - y_k\|_2^2)}$$

As such, the objective function to minimize with regard to the embedded points y_i is given by

$$C = \sum_i KL(P_i || Q_i) = \sum_i \sum_j p_{j|i} \log \left(\frac{p_{j|i}}{q_{j|i}} \right).$$

After minimizing C , the points y_i represent the two most significant components of the geometry of the set X_D , which can be plotted on a 2-dimensional plan for visualization (see the embedding result in Figure 18, for example. Notice that the points are visually separated into 10 clusters, as there are 10 types of digits).

The success of t-SNE shows that $p_{j|i}$ preserves important information of the manifold of images, which should be beneficial for the measure of spread in the context of image synthesis. Denoting by $X_T = \{X_i | i \in T\}$ the set of training images and by $X_G = \{X_i | i \in G\}$ the set of generated ones,³ the proposed spread measure around one training image $X_i \in X_T$ is defined to be the ratio between the density of the neighbors of X_i from X_G and its neighbors from X_T :

$$\text{Spread}(i; X_T, X_G) = \log \left(\frac{\sum_{j \in G} p_{j|i} \|X_j - X_i\|_2^2}{\sum_{j \in T} p_{j|i} \|X_j - X_i\|_2^2} \right),$$

where $p_{j|i}$ is defined as in Equation (13) by merging the training set with the set of generated images, i.e. $D = T \cup G$.⁴ The sum $\sum_{j \in G} p_{j|i} \|X_j - X_i\|_2^2 + \sum_{j \in T} p_{j|i} \|X_j - X_i\|_2^2$ can be seen as the trace of the covariance matrix of all the neighbors of X_i from X_G and X_T .⁵ The term $\sum_{j \in G} p_{j|i} \|X_j - X_i\|_2^2$ is the contribution of X_G to this trace, and $\sum_{j \in T} p_{j|i} \|X_j - X_i\|_2^2$ is the contribution of X_T to it. Intuitively, as $|X_T| \approx |X_G|$, $\text{Spread}(i; X_T, X_G) \approx 0$ indicates that X_i has about the same number of neighbors from X_T and X_G within the radius σ_i , so there is little bias in synthesis from the “point of view” of X_i . If $\text{Spread}(i; X_T, X_G) \gg 0$, then X_G has a much larger density of points around X_i compared to X_T , and vice versa for the case $\text{Spread}(i; X_T, X_G) \ll 0$.

Having the spread measure for each training image defined, we can score the overall performance simply by the average

³For a fair comparison between the two sets, we assume that they have about the same number of elements: $|X_T| \approx |X_G|$.

⁴We stress on the fact that this spread measure does not use the directly the embedding result of t-SNE $\{y_i\}_{i \in D}$. Instead, it is merely based on the pairwise distances of images and the probabilities $\{p_{j|i}\}_{i,j}$.

⁵The neighbor X_j is weighted by $p_{j|i}$. The mean of the neighbors is assumed to be X_i .

TABLE II

SCORES OF THE ASSESSMENT FRAMEWORK ON *Generated* (FIRST TWO ROWS) AND *Test* (THIRD ROW) MNIST DIGITS. THE SCORES OF IMAGES OF DIFFERENT DIGITS ARE AVERAGED TOGETHER.

BOLD INDICATES THE BEST PERFORMANCE IN SYNTHESIS. NOTICE THAT THE FIRST TWO LL VALUES ARE COMPUTED ON *Generated* IMAGES, IN CONTRAST TO THE VALUES SHOWN IN TABLE I, WHICH ARE EVALUATED ON *Test* IMAGES

Model	LL	Originality	Spread
DRAW [15]	-79.33	0.622	1.11
suggested method	-68.20	1.04	1.31
test images	-82.52	1.09	0.911

of the the spread of all training images:

$$\text{Spread}(X_G; X_T) = \frac{1}{|X_T|} \sum_{i \in T} |\text{Spread}(i; X_T, X_G)|,$$

which is expected to be as small as possible.

Finally, to visualize how the synthesized images spread out over the training ones, we simply apply t-SNE on the training images X_T and generated ones X_G altogether, and plot the 2-dimensional embedding result, using different colors for images from X_T and X_G (see Figure 18 for example). This way, one can directly observe if the elements of X_T and X_G overlap well in the embedded space.

D. Assessment of MNIST Synthesis

In this section we apply the assessment framework described above to compare our generated MNIST digits to the ones of state-of-the-art DRAW [15]. Table II shows these scores for the synthesized- and the test-images as a reference performance.

For the LL (the higher the better), our average value is much higher than DRAW and the one obtained on the test images, implying that the generated images are of high visual quality with regard to the training set. This is not surprising, as our method targets the minimization of the sum of the LL of the patches. Notice that the test images have smaller LL than both synthesis methods, for which a possible explanation is that patches of the test images might not exist in the training set, resulting in a lower patch LL.

Notice that low LL score indicates that the image is unique, however, it cannot distinguish whether this unique result is a failure or a truly novel and realistic generated image. This explains why the LL of the test images is lower than the generated ones. On the other hand, we empirically observe that low LL measure is correlated with badly *generated* images, but not with bad *training/test* ones, as depicted in Figure 15.

As such, we use the LL measure to detect and understand the failed generations. In fact, we observe that the failure is mostly related to the seeds: Digits of high quality are generated from most seeds, while some seeds lead more frequently to failures. To explain this difference, we can compute the LL of these two kinds of seeds,⁶ and compare their synthesis results. Figure 16 shows two example seeds of the digit “8”

⁶We model the probability of the 4×4 seeds by the Parzen window defined by the seed images created from the whole training set, using Gaussian kernel.

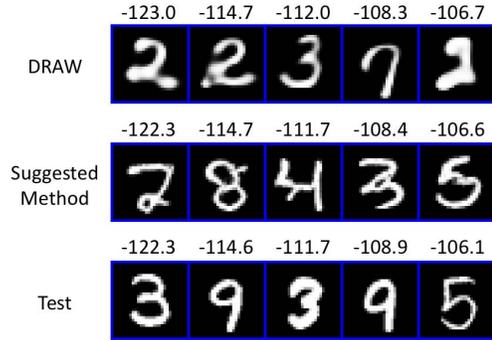


Fig. 15. Examples of MNIST images with low LL values. The LL of each image is indicated above it. From up to down, The images are generated by DRAW, the proposed method and taken from the test set, respectively.

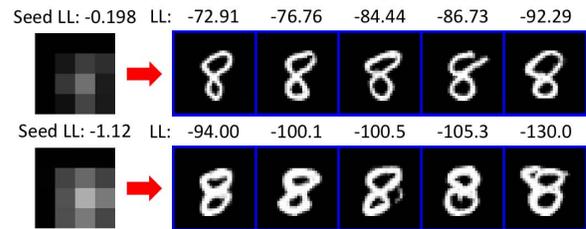


Fig. 16. Two seeds of digit the “8” with different LL and their generation results. The images are sorted by decreasing LL. The LL of a seed image is evaluated under the Parzen window defined with the seeds created by downsampling all the training images. The average of LL of all the seeds of “8” is -0.279.

and different generation runs from them, with the LL of each image. The first seed has a high LL, and its results have good visual quality and contain no visible failures; On the other hand, the second seed has a much lower LL, leading to moderate failures. Intuitively, the LL of the seed indicates its uniqueness, and the difficulty to create likely images from it. Also notice that in the results of the second seed, the visual quality decreases with the LL (from left to right), showing that our LL definition is indeed meaningful.

As for the second assessment score – the originality (higher is better) – our generated images have an originality value close to 1, implying that the proposed method synthesizes rich and novel images which are far from being replicated version of the training images. As a comparison, the images of DRAW have much smaller originality (equal to 0.622), as shown in Table II. Visually, Figure 17 plots the points (D_G, D_T) (as defined in Section V-B) for our results, the results of DRAW and the test images. The points corresponding to images with good originality are those close to the line $D_G = D_T$ or below it. As can be seen, our generated images have high originality as the test images do, which is greater than the originality of DRAW.

Now we turn to the last measure, the spread (lower is better). Following Table II, DRAW has a better spread score than our method, which is confirmed by the t-SNE visualizations in Figure 18. As can be observed, the points corresponding to the generated images of DRAW overlap well the points of the training images (right figure), while the points corresponding to our generated images are “biased” towards the regions more

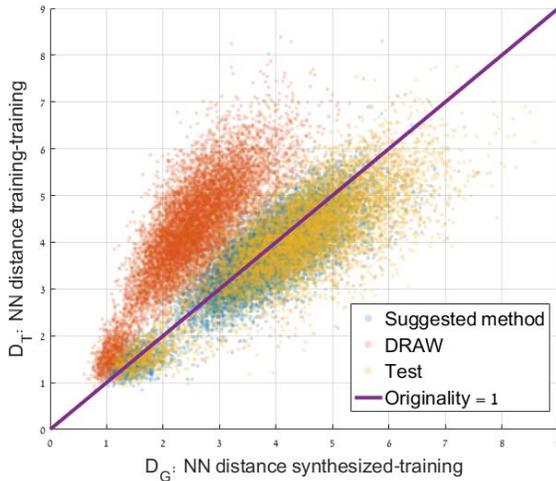


Fig. 17. Visualization of the originality of the MNIST digits synthesized by the suggested method and DRAW, and the MNIST test images. The supplementary material contains a different visualization of the originality.

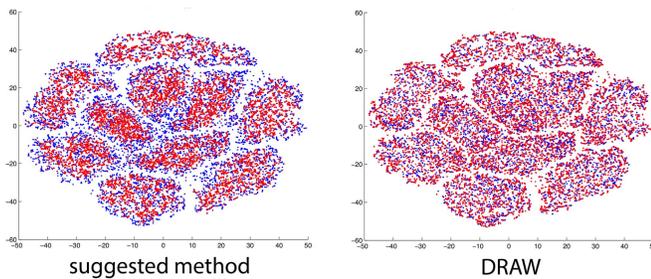


Fig. 18. t-SNE visualization result of the MNIST digits generated by the suggested method (red points on the left) and DRAW (red points on the right), together with part of the training images (blue points). The x and y -axes represent the two most significant geometry components found by t-SNE.

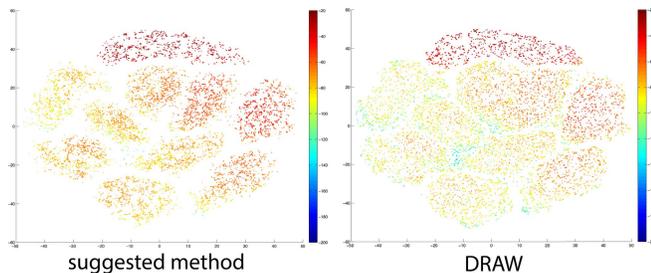


Fig. 19. Visualization of the LL of the MNIST digits generated by the suggested method (left) and DRAW (right). The positions of the points are the same as the red points in Figure 18. The colors of the points correspond to the LL value of the images. More results of LL visualization are provided in the supplementary material.

populated by the points of the training images (left figure). The worse spread of our result can be explained by the LL measure. As shown in Figure 19, the LL value of the images in less populated regions have slightly lower LL values. Since our method maximizes the LL to some extent, the images are “dragged” towards more populated regions during the synthesis process.

In conclusion, our synthesis algorithm produces high quality MNIST images, outperforming DRAW in both LL and originality. Although DRAW has a better spread, it may be due

to its tendency to replicate training images (as its originality score indicates).

VI. DISCUSSION

The synthesis scheme we propose generates high quality MNIST images competitive to the state-of-the-art, and visually appealing face images. Furthermore, it is general since various patch priors can be easily plugged into the generation framework to create completely different kinds of images (e.g. digits and faces). Finally, our proposed prior is non-parametric and intuitive opposed to the majority of recent works, which rely on neural networks as blackboxes, thus lose traceability.

In addition, our evaluation framework makes the first step towards a complete way to assess the synthesis performance, by evaluating the LL, the originality and the spread of the generated images. The framework borrows tools from the well known t-SNE [51], which visualizes the MNIST digits in an impressive way.

One can wonder how other patch priors work in synthesis compared to the example-based one in Section III-B. In fact, a natural alternative is GMM, which provides state-of-the-art restoration results when used with EPLL [6], and has a natural definition of patch LL. However, we believe that GMM tends to prefer smooth patches as their probability is emphasized in the Gaussian distribution, and this is a shortcoming if we are to generate sharp images. On the other hand, our current example-based model does not generate new patches, which limits the novelty of the generated images in small range. Therefore, it will be interesting to find a patch prior able to generate new yet sharp and likely patches.

A promising future direction is to extend our algorithm to work with more complex images, e.g. natural scenes. However, we believe that simple adaptations are not enough, as natural images have rich high-level structures, for example objects and their relations, and these images have weaker locality than digits and faces. Therefore, more delicate work may be needed towards using the proposed scheme for getting appealing natural image synthesis. One possibility is to model the patches with better use of their context, or in a feature space (e.g. as done in [7]).

Another future work direction consists of further improving the proposed assessment framework. Recall that two of the three measures (namely the originality and the spread) are based on the L_2 -distance. This distance is a reasonable choice for intrinsically low dimensional images with enough samples, e.g. digits. However, when it comes to faces or even natural images, L_2 -distance includes less information and the desired properties of natural images (e.g. invariance to shift, rotation and luminance) are largely missing in this metric. Therefore, an appropriate metric is crucial for us to assess generated complex images in a meaningful way. A recent work on synthesis with neural networks [26] suggests trained metrics, showing one possibility of achieving this goal.

REFERENCES

- [1] M. Elad and M. Aharon, “Image denoising via sparse and redundant representations over learned dictionaries.” *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.

- [2] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [3] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. Image Process.*, vol. 17, no. 1, pp. 53–69, Jan. 2008.
- [4] A. Criminisi, P. Pérez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1200–1212, 2004.
- [5] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Trans. Image Process.*, vol. 19, no. 11, pp. 2861–2873, Nov. 2010.
- [6] D. Zoran and Y. Weiss, "From learning models of natural image patches to whole image restoration," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2011, pp. 479–486.
- [7] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin, "Image analogies," in *Proc. 28th Annu. Conf. Comput. Graph. Interaction Techn.*, 2001, pp. 327–340.
- [8] C. Barnes, F.-L. Zhang, L. Lou, X. Wu, and S.-M. Hu, "PatchTable: efficient patch queries for large datasets and applications," *ACM Trans. Graph. (TOG)*, vol. 34, no. 4, p. 97, 2015.
- [9] J. E. Kyrianiadis, J. Collomosse, T. Wang, and T. Isenberg, "State of the 'Art': A taxonomy of artistic stylization techniques for images and video," *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 5, pp. 866–885, May 2013.
- [10] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 2414–2423.
- [11] J. Johnson, A. Alahi, and L. Fei-Fei. (2016). "Perceptual losses for real-time style transfer and super-resolution." [Online]. Available: <https://arxiv.org/abs/1603.08155>
- [12] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, Feb. 2015.
- [13] A. Radford, L. Metz, and S. Chintala. (2015). "Unsupervised representation learning with deep convolutional generative adversarial networks." [Online]. Available: <https://arxiv.org/abs/1511.06434>
- [14] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. (2016). "Pixel recurrent neural networks." [Online]. Available: <https://arxiv.org/abs/1601.06759>
- [15] K. Gregor, I. Danihelka, A. Graves, D. Rezende, and D. Wierstra, "DRAW: A recurrent neural network for image generation," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1462–1471.
- [16] G. Research. (2015). *Inceptionism: Going Deeper Into Neural Networks*. Accessed: Aug. 6, 2015. [Online]. Available: <http://googleresearch.blogspot.ch/2015/06/inceptionism-going-deeper-into-neural.html>
- [17] A. A. Efros and W. T. Freeman, "Image quilting for texture synthesis and transfer," in *Proc. 28th Annu. Conf. Comput. Graph. Interaction Techn.*, 2001, pp. 341–346.
- [18] V. Kwatra, I. Essa, A. Bobick, and N. Kwatra, "Texture optimization for example-based synthesis," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 795–802, 2005.
- [19] L. Gatys, A. S. Ecker, and M. Bethge, "Texture synthesis using convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 262–270.
- [20] M. Ashikhmin, "Synthesizing natural textures," in *Proc. Symp. Interaction 3D Graph.*, 2001, pp. 217–226.
- [21] S. Lefebvre and H. Hoppe, "Appearance-space texture synthesis," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 541–548, 2006.
- [22] A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, Sep. 1999, pp. 1033–1038.
- [23] T. Raiko, Y. Li, K. Cho, and Y. Bengio, "Iterative neural autoregressive distribution estimator (NADE-k)," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 325–333.
- [24] R. Salakhutdinov and G. E. Hinton, "Deep boltzmann machines," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2009, pp. 448–455.
- [25] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [26] A. B. L. Larsen, S. K. Sønderby, and O. Winther. (2015). "Autoencoding beyond pixels using a learned similarity metric." [Online]. Available: <https://arxiv.org/abs/1512.09300>
- [27] C. Liu, H.-Y. Shum, and W. T. Freeman, "Face hallucination: Theory and practice," *Int. J. Comput. Vis.*, vol. 75, no. 1, pp. 115–134, 2007.
- [28] U. Mohammed, S. J. Prince, and J. Kautz, "Visio-ization: Generating novel facial images," *ACM Trans. Graph.*, vol. 28, no. 3, p. 57, 2009.
- [29] C. Li and M. Wand, "Combining Markov random fields and convolutional neural networks for image synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 2479–2486.
- [30] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. (2016). "Image-to-image translation with conditional adversarial networks." [Online]. Available: <https://arxiv.org/abs/1611.07004>
- [31] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman, "PatchMatch: A randomized correspondence algorithm for structural image editing," *ACM Trans. Graph.*, vol. 28, no. 3, p. 24, 2009.
- [32] S.-M. Hu, F.-L. Zhang, M. Wang, R. R. Martin, and J. Wang, "PatchNet: A patch-based image representation for interactive library-driven image editing," *ACM Trans. Graph.*, vol. 32, no. 6, p. 196, 2013.
- [33] F.-L. Zhang, J. Wang, E. Shechtman, Z.-Y. Zhou, J.-X. Shi, and S.-M. Hu, "PlenoPatch: Patch-based plenoptic image manipulation," *IEEE Trans. Vis. Comput. Graphics*, vol. 23, no. 5, pp. 1561–1573, May 2017.
- [34] J. Ho Lee, I. Choi, and M. H. Kim, "Laplacian patch-based image synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 2727–2735.
- [35] C. Barnes and F.-L. Zhang, "A survey of the state-of-the-art in patch-based synthesis," *Comput. Vis. Media*, vol. 3, no. 1, pp. 3–20, Mar. 2017.
- [36] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Comput. Graph. Appl.*, vol. 22, no. 2, pp. 56–65, Mar./Apr. 2002.
- [37] G. Yu, G. Sapiro, and S. Mallat, "Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity," *IEEE Trans. Image Process.*, vol. 21, no. 5, pp. 2481–2499, May 2012.
- [38] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Oct. 2009, pp. 2272–2279.
- [39] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [40] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [41] V. Papan and M. Elad, "Multi-scale patch-based image restoration," *IEEE Trans. Image Process.*, vol. 25, no. 1, pp. 249–261, Jan. 2016.
- [42] D. Forsyth and J. Ponce, *Computer Vision: A Modern Approach Always Learning*. London, U.K.: Pearson, 2012. [Online]. Available: <https://books.google.com/books?id=gM63QQAACAAJ>
- [43] M. Muja and D. G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," *VISAPP*, vol. 2, nos. 331–340, p. 2, 2009.
- [44] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play admm for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 84–98, Jan. 2017.
- [45] I. Ram, M. Elad, and I. Cohen, "Image processing using smooth ordering of its patches," *IEEE Trans. Image Process.*, vol. 22, no. 7, pp. 2764–2774, Jul. 2013.
- [46] Y. Romano and M. Elad, "Con-patch: When a patch meets its context," *IEEE Trans. Image Process.*, vol. 25, no. 9, pp. 3967–3978, Sep. 2016.
- [47] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 313–318, 2003.
- [48] E. Jang. (2016). *TensorFlow Implementation of 'DRAW: A Recurrent Neural Network for Image Generation*. Accessed: Aug. 2, 2016. [Online]. Available: <https://github.com/ericjang/draw>
- [49] M. Elad, R. Goldenberg, and R. Kimmel, "Low bit-rate compression of facial images," *IEEE Trans. Image Process.*, vol. 16, no. 9, pp. 2379–2383, Sep. 2007.
- [50] E. Parzen, "On estimation of a probability density function and mode," *Ann. Math. Statist.*, vol. 33, no. 3, pp. 1065–1076, Sep. 1962.
- [51] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, nos. 2579–2605, p. 85, 2008.
- [52] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, Dec. 2000.
- [53] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, Dec. 2000.



Yi Ren received the B.Sc. degree from the Ecole Polytechnique, France, in 2014, and the M.Sc. degree from the Department of Computer Science, Technion–Israel Institute of Technology, in 2017. He is currently a deep learning Research Engineer. His research interests include signal and image processing, computer vision, machine learning, and deep learning.



Yaniv Romano received the B.Sc. degree from the Department of Electrical Engineering, Technion–Israel Institute of Technology, in 2012, where he is currently pursuing the Ph.D. degree. He received the 2015 Zeff Fellowship, the 2017 Andrew, and Erna Finzi Viterbi Fellowship, and the 2017 Irwin and Joan Jacobs Fellowship.

Since 2011, he has been involved in the industry as an Image Processing Algorithm Developer. His research interests include sparse and redundant representations, deep learning, signal and image modeling, inverse problems, and machine learning.



Michael Elad (F'12) received the B.Sc., M.Sc., and D.Sc. degrees in electrical engineering from the Technion–Israel Institute of Technology in 1986, 1988, and 1997, respectively. After several years in industrial research, he served as a Research Associate with Stanford University from 2001 to 2003, involved closely with Prof. G. Golub (CS), Prof. P. Milanfar (EE-UCSC), and Prof. D. L. Donoho (Stat). Since 2003, he has been holding a faculty position with the Computer-Science Department, Technion–Israel Institute of Technology.

He is currently involved in the field of signal and image processing, specializing in particular on inverse problems and sparse representations. He has authored hundreds of technical publications in leading venues, many of which have led to exceptional impact. He has authored the book *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, in 2010, which is a leading publication in this field.

Dr. Elad has served as an Associate Editor of the IEEE-TIP from 2007 to 2011, the IEEE-TIT from 2011 to 2014, ACHA from 2012 to 2015, and the *SIAM Imaging Sciences-SIIMS* from 2010 to 2015. He held a Senior Editorial role of the IEEE-SPL from 2012 to 2014. Since 2016, he has been serving as the Editor-in-Chief of SIIMS.

Dr. Elad received numerous teaching and research awards and grants. He received the ERC Advanced Grant in 2013. He was a recipient of the 2008 and 2015 Henri Taub Prizes for Academic Excellence, and the 2010 Hershel-Rich Prize for innovation.