

**FASTER UNSUPERVISED SEMANTIC INPAINTING: A GAN BASED APPROACH**Avisek Lahiri<sup>1†‡</sup>Arnav Kumar Jain<sup>2\*‡</sup>Divyasri Nadendla<sup>3\*</sup>Prabir Kumar Biswas<sup>1</sup><sup>1</sup> Dept. of E&ECE, Indian Institute of Technology, Kharagpur, India<sup>2</sup> Microsoft, India<sup>3</sup> Qualcomm, India

avisek@ece.iitkgp.ac.in, arnavkj95@gmail.com, divyasri.vss@gmail.com, pkb@ece.iitkgp.ac.in

**ABSTRACT**

In this paper, we propose to improve the inference speed and visual quality of contemporary baseline of Generative Adversarial Networks (GAN) based unsupervised semantic inpainting. This is made possible with better initialization of the core iterative optimization involved in the framework. To our best knowledge, this is also the first attempt of GAN based video inpainting with consideration to temporal cues. On single image inpainting, we achieve about 4.5-5 $\times$  speedup and 80 $\times$  on videos compared to baseline. Simultaneously, our method has better spatial and temporal reconstruction qualities as found on three image and one video dataset.

**Index Terms**— Generative Adversarial Networks, Semantic Inpainting, Temporal Consistency, Video Inpainting

**1. INTRODUCTION**

Semantic inpainting refers to filling up of missing pixels in a given image by leveraging neighborhood information. Traditional methods [1, 2] were mainly successful when deployed on background scenes and images with repeated textures. However, they fail to learn complex semantic representations and thereby manifest displeasing reconstructions on complex non-repetitive textured objects. With the advent of Generative Adversarial Networks (GAN), there has been a recent surge of interest [3, 4, 5, 6] to solve inpainting with deep generative models. There are mainly two schools of approach:

**Fully unsupervised:** This approach, first proposed by Yeh *et al.* [3] aligns with the concept of the pioneering paper of GAN [7]. In [3], the objective is learn a GAN model to generate realistic images conditioned on noise priors only and inpainting is done by iteratively matching a masked/damaged image to its ‘best matching’ noise prior. This method does not require any paired training set (masked, unmasked) and hence we term it as ‘unsupervised’.

**Hybrid:** These methods [5, 6, 8, 4] in general rely on initial training with usual reconstruction loss on a paired (masked,

unmasked) dataset. Since  $L_2$  loss based reconstructions manifest lack of high frequency components, the next step is to push the solutions nearer to original data manifold with an additional adversarial loss. It is to be noted, that without the initial supervised training phase, these methods fail to work and thus we term this framework as ‘hybrid’ approach.

**Motivation:** For [3], being fully unsupervised comes at a cost of significant inference time due to an iterative search for a matching noise prior. Hybrid methods perform test time inference in one single forward pass and thus research has been dedicated mainly towards this genre. However, going against the trend, we advocate the former method because true potential of GAN is appreciated when there is no source of supervision. The motivation in this paper is to primarily reduce the inference run time of [3], yet achieve better/similar reconstruction performance compared to [3]. Towards this, the paper presents the following contributions:

1. A better initializing method for the iterative optimization of [3] to speedup inference run time on single image inpainting by 4.5-5 $\times$ .
2. First demonstration of totally unsupervised GAN based inpainting on videos (in context of error concealment) with speedup upto 80 $\times$  by leveraging temporal redundancy
3. A group consistency loss for a more temporally consistent sequence reconstruction and thereby leading to more pleasing spatio-temporal experience as ascertained by the MOVIE metric [9]
4. Exhaustive experiments on SVHN, Stanford Cars, CelebA image dataset and VidTIMID video dataset manifest the benefits of our approach

**2. GAN PRELIMINARIES**

A GAN model consists of two deep neural nets, viz., generator,  $G$ , and discriminator,  $D$ . The task of the generator is to create an image,  $x \in \mathcal{R}^{H \times W \times 3}$  with a noise prior vector,  $z \in \mathcal{R}^d$ , as input.  $z$  is sampled from a known distribution,  $p_z(z)$ ; usually  $z \sim \mathcal{U}[-1, 1]^d$ . The discriminator has to distinguish between real samples (sampled from real distribution,  $p_{data}$ ) and generated samples. The game is played on

<sup>†</sup>Funded by Google PhD Fellowship & Qualcomm Fellowship

<sup>‡</sup>Denotes equal contribution

\*Work done while at IIT Kharagpur

$V(D, G)$ :

$$\begin{aligned} \min_G \max_D V(D, G) &= \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] \\ &+ \mathbb{E}_{z \sim p_z(z)} [1 - D(G(z))]. \end{aligned} \quad (1)$$

### 3. METHOD

We build upon the unsupervised inpainting framework of Yeh et al.[3]. Given a masked image,  $I_d = M \odot I$ , corresponding to an original image,  $I$ , and a pre-trained GAN model, the idea is to iteratively find the ‘closest’  $z$  vector (starting randomly from  $\mathcal{U}[-1, 1]^d$ ) which results in a reconstructed image whose semantics are similar to corrupted image.  $z$  is optimized as,

$$\hat{z} = \underset{z}{\operatorname{argmin}} J(M \odot G(z), I_d). \quad (2)$$

where  $M$  is the binary mask with zeros on masked region else unity,  $\odot$  is the pointwise multiplier and  $J(\cdot)$  is the objective function to be minimized. Interesting to note is that the objective function never assumes knowledge of pixel intensities inside the masked region (and thus the term ‘unsupervised’). Upon convergence, the inpainted image,  $\hat{I}$ , is given as,  $\hat{I} = M \odot I + (1 - M) \odot G(\hat{z})$ . The objective function,  $J(\cdot)$  is composed of two components:

**Fidelity Loss:** This loss ensures that the predicted noise prior preserves fidelity between generated image and the original unmasked regions.

$$L_f = |M \odot (I - G_{\theta_G}(\hat{z}))| \quad (3)$$

**Perceptual Loss:** This loss ensures that the inpainted output lies near the original/real data manifold and is measured by the log likelihood of real class assigned by the pre-trained discriminator;

$$L_p = \log(1 - D_{\theta_d}(G_{\theta_G}(\hat{z}))) \quad (4)$$

The overall objective,  $J(\cdot) = L_f + \lambda L_p$ , where  $\lambda$  controls the relative importance of  $L_p$ .

#### 3.1. Better initiation for noise prior search

One of the fundamental drawbacks of Yeh *et al.* is the iterative optimization requirement of Eq.2. A random  $z \sim \mathcal{U}[-1, 1]^d$  usually tends to generate images quite disparate from the concerned masked image and thus the optimization requires multiple updates of  $z$ . In fact, the authors in [3] suggest around 1000 rounds of iterations per image. Our motivation is to initiate  $z$  by respecting some global statistics of the concerned masked image.

**Nearest neighbor search:** After training a GAN, we store (one time offline task) a pool,  $P$ , of  $N$  images by passing  $N$  random noise vectors through the pre-trained generator. For

a given damaged image,  $I_d$ , we perform a nearest neighbor search over the pool,  $P$ , to identify ‘the closest’ matching pair. Specifically, we perform the matching between  $I_d$  and a candidate pooled image (generated from  $z_i$ ),  $p_i = G(z_i)$ , based on a distance metric,  $D(I_d, M \odot p_i)$ . Please note, even during matching we are not exploiting the masked region of  $I_d$ . While formulating  $D(\cdot)$  we want to make sure that we not only match the overall color statistics of the damaged image but also respect the overall structure. Thus  $D(\cdot)$  has got two components:

**Data loss:** This loss penalizes if the pixel intensities of a pooled image,  $p_i$ , deviate from the damaged image,  $I_d$ :

$$L_D = |I_d - (M \odot p_i)|. \quad (5)$$

**Structure loss:** This loss penalizes if the structure (captured in essence with gradients) of pooled image deviates from damaged image. Structure loss,  $L_S$  is defined as:

$$L_S = |\nabla_x I_d - \nabla_x M \odot p_i| + |\nabla_y I_d - \nabla_y M \odot p_i|, \quad (6)$$

where  $\nabla_x$  and  $\nabla_y$  are horizontal and vertical gradient operators. The final matching criterion is,  $L_{nn} = L_D + \gamma L_S$ , where  $\gamma$  controls relative importance of  $L_S$ . Effect of  $\gamma$  is discussed in Fig.1. The initial noise vector,  $z_{init}$  is given by,

$$z_{init} = \underset{z}{\operatorname{argmin}} L_{nn}(I_d, G(z)); z \in \{z_0, z_1, \dots, z_N\} \quad (7)$$

#### 3.2. Video Inpainting: Exploiting temporal redundancy

To our best knowledge, this is the first demonstration of unsupervised GAN based inpainting on videos. By video inpainting we refer to concept of error concealment in video, i.e., to recover damaged/masked portion of a frame. A naive application of [3] would be to apply single image model independently on each frame. This poses two problems, viz, a) such approach does not leverage temporal redundancy among neighboring frames b) independent frame level reconstructions result in temporal inconsistency in a sequence. We propose to address these challenges with two innovations.

**Reuse of predicted  $z$  vector:** It is safe to assume that neighboring frames are coherent in appearance and thus the noise priors. Thus, it makes sense to initiate  $z_{t+1} = z_t$ . This drastically speeds up optimization by almost  $80\times$  compared to vanilla version of Yeh *et al.*[3]. We refer to this proposed method as *Proposed (Re)* in all experiments.

**Group consistency loss:** Even though we initialize ( $z_{t+1}$ ) with  $z_t$ , the final solution for time step ( $t+k$ ) can diverge away appreciably from  $z_t$ . This would mean that there will abrupt changes in scene appearance when the frames are viewed as a sequence. To enforce smooth temporal dynamics we impose a group consistency loss, ( $L_G$ ), by constraining a group of reconstructed frames to be similar. Disparity between two generated images can be expressed with corresponding disparity between the corresponding  $z$  vectors [10].

Specifically, we impose consistency loss over a window of  $W$  neighboring frames,

$$L_G = |z_i - z_k|; \forall i \in [1, W], \forall k \in [1, W] \quad (8)$$

Please note that  $L_G$  is imposed only over a neighboring window of  $W$  frames and not over entire video. Combination of this loss + reuse of  $z$  vector is denoted as *Proposed (Re + G)* in experiments.

#### 4. EXPERIMENT SETTINGS


**Datasets:** For image inpainting we tested on SVHN[11], Stanford Cars[12] @  $64 \times 64$  resolution and CelebA[13] @  $64 \times 64$  and  $128 \times 128$  resolution. For video inpainting, we experimented at  $128 \times 128$  resolution on VidTIMIT[14] dataset.

**Network Architectures:** For fair comparison with our baseline of [3], we borrowed their architectures for both generator and discriminator and followed their paradigm of training GAN. Parameter,  $\lambda$ , was set to 0.01 following [3].

**Balancing data loss and structure loss:** Hyperparameter,  $\gamma$  sets the relative importance of structure loss,  $L_S$ , over data loss  $L_D$ . Setting  $\gamma = 0$  means initial solution will not explicitly preserve edge information and just retrieve nearest image based on raw intensity. On contrary, a high value of  $\gamma$  will enforce only edge preservation without respecting the intensity. See Fig.1 for illustration of these two extreme cases. In either of these cases, initial solution can be appreciably different than desired solution and will thus require longer iterative refinements. We reserved a validation dataset on which we test  $\gamma \in \{0.001, 0.005, 0.01, 0.05, 0.1, 0.5\}$ .  $\gamma = 0.01 - 0.03$  gave peak speedup across datasets and thus  $\gamma$  is set to 0.01 for all experiments. See Fig.2 for some examples of initial solutions with our proposed method v/s random initialization of [3].

**Pool size ( $N$ ) for nearest neighbor search:** Nearest neighbor search with an ideal generator ( $p_G = p_{data}$ ) will be able to retrieve the exact  $z$  vector corresponding to a masked image,  $I_d$ , if we allow,  $N \rightarrow \infty$ . However, for practical viability it is not possible to search over every possible  $z$  vector. On validation set, we experiment with different  $N$  and compare the average MS-SSIM between initial masked solution and masked image.  $N$  is set to 300(on all datasets unless otherwise stated), above which, the MS-SSIM does not increase appreciably.

**Selecting group consistency window,  $W$ :** Setting  $W$  (in Eq.8) to a large value results in over smoothing of a sequence in temporal dimension and thus leads to degraded MOVIE metric[9] due to poor perceptual quality.  $W = 0$ , on the other hand will have no effect in incorporating temporal coherence which again hurts MOVIE metric. On held out validation set of VidTIMIT we experimented with  $W \in [1, 10]$ . See Fig. ?? for variation of MOVIE metric with  $W$  on validation set of VidTIMIT.  $W = 5$  was selected which, on average, yielded best MOVIE metric. So, for a sequence, every alternate  $5^{th}$



**Fig. 1:** Role of data loss,  $L_D$  (Eq.5) and structural loss,  $L_S$  (Eq.6) on retrieving nearest matching initial solution. For each tuple, left column: masked image, middle column: initial solution retrieved with only  $L_S$ , right column: initial solution retrieved with only  $L_D$ . Using  $L_D$  mainly tries to maintain the global color statistics while only  $L_S$  focuses on matching the structure irrespective of absolute intensity concern. It can be appreciated that only  $L_S$  retrieves initial matches by maintaining facial expression(smile), orientation of cars, keystrokes of digits. Thus we apply weighted combination of  $L_D$  and  $L_S$  as objective (Eq.7) for nearest neighbor retrieval.



**Fig. 2:** Benefit of proposed noise prior initialization v/s random initialization of Yeh *et al.* [3]. Top row: Masked image, Middle row: initial solution of [3], Bottom row: initial solution by our proposed nearest neighbor based. It is evident that our initial solutions are much more closer to masked images and thus requires lesser iterative updates of Eq.2. compared to [3].

frame (pivot frame) is inpainted as a single image. Intermediate frames are initiated with inpainted pivot frame and solutions of a group are additionally constrained by  $L_G$ .

**Comparing Methods:** Our primary comparing method is the unsupervised baseline of [3]. However, for the completeness of the paper we also provide comparisons with hybrid frameworks of [8, 4, 5]. Please note, the later approaches are only successful if trained initially with supervised reconstruction loss. If trained with only unsupervised adversarial loss, these method fail drastically.

**Table 1:** Comparing inpainting PSNR (in dB) on different datasets with the unsupervised baseline of [3]. We also compare with hybrid methods of [8, 4, 5].

	[8]	[4]	[5]	[3]	Ours
Cars	14.3	15.3	14.5	13.5	14.1
SVHN	21.5	23.6	23.7	20.4	22.0
CelebA(64)	23.0	24.1	24.2	22.6	23.3
CelebA(128)	20.0	20.9	20.6	17.6	18.8

#### 5. RESULTS

**Speedup in optimization:** With respect to our unsupervised baseline[3], on average, we achieved about  $5 \times$  speedup for single image inpainting. On videos the speedup is almost  $80 \times$ . See Table 2 for speed comparisons.

**Image Inpainting:** In Fig. 3 we show some exemplary inpainting comparison with [3]. Even with appreciable speedup

**Table 2:** Comparison of absolute run times (in seconds) on a NVIDIA K-40 GPU on image (64×64) and video (128×128) inpainting. Time is measured till corresponding loss of a model converges to 95% of saturation value. Notice how application of single image model of [3] naively on higher resolution video booms up the inference time; however, our proposed model (Re + G) appreciably speeds up by almost 80%. In fact, on image and video we achieve speedups 5× and 100× in terms of iterations count. This table also considers the time of nearest neighbor search.

	Image	Video
Yeh <i>et al.</i> [3]	9.0	33.5
Proposed (Re)	1.9	0.36
Proposed (Re + G)	-	0.41

**Table 3:** Average temporal consistency ( $\eta$ ) in dB on test sets of different dataset. Higher value means a model is more temporally coherent.

	[8]	[4]	[5]	[3]	Ours (Re)	Ours (Re + G)
Cars	13.8	14.2	14.5	15.6	18.2	20.0
SVHN	21.3	21.8	22.1	22.8	24.1	24.8
CelebA(64)	23.1	23.2	23.6	24.1	25.6	26.3
CelebA(128)	21.8	20.9	21.6	21.9	22.4	23.5


**Table 4:** Comparison on MOVIE metric[9]  $\in [0, 1]$  on ViDTIMIT video test set. We compare with unsupervised baseline of Yeh *et al.* [3] and also with hybrid methods of [8, 4, 5]. A lower MOVIE metric is better in terms of spatio-temporal effectiveness of inpainting.

[3]	[8]	[4]	[5]	Proposed (Re)	Proposed (Re + G)
0.66	0.63	0.55	0.46	0.52	0.47


we usually achieve better (or similar visual performance) to [3]. We also show some visual comparison with recent hybrid benchmarks in Fig.4. Recently [15, 3, 5] researchers have shown that PSNR metric is not fully justifiable to assess tasks such as inpainting and super resolution. However, for reference we also report PSNR in Table1.

**Pseudo sequences and temporal consistency** Before analyzing effects of our proposed losses on real videos, we study the effects on a simpler case of pseudo sequences. A pseudo sequence of length,  $S$  is basically a single image replicated  $S$  times but masked with different masks. An ideal model would inpaint all the frames identically. We can define temporal consistency,  $\eta$  as,  $\eta = \frac{1}{\binom{S}{2}} \sum_{k=1}^S \sum_{j=1}^S PSNR(\hat{I}_k^i, \hat{I}_j^i) \forall$  combinations of  $(j, k) \in [1, 2, \dots, S]$ . In Table 3 we report temporal consistency over different datasets. It can be seen that proposed initialization technique of  $z_{t+1} = z_t$  improves consistency with further improvement brought by group consistency loss. Even the current hybrid benchmarks [5, 8, 4] manifest greater inconsistency because these methods do not leverage any temporal information. Exemplary visualizations are shown in Fig. 5. Note,  $\eta$  gives an indication of temporal consistency only (which can be studied only on such pseudo sequences). It does not give essence of spatial correctness. For example, a model might reconstruct all blank images which will give high  $\eta$ , but will yield worse MOVIE metric.

**Analyzing real video inpainting performance** PSNR and MS-SSIM are not well suited for judging a reconstructed



**Fig. 3:** Visual comparison of inpainting by the unsupervised baseline of [3](3<sup>rd</sup> column) and our proposed method(4<sup>th</sup> column). We perform equivalently(sometimes better) with 5× less iterations. 1<sup>st</sup> col: Original; 2<sup>nd</sup>: masked image.



**Fig. 4:** Comparing with contemporary hybrid methods of CE[8] and GIP[5].

video since these metrics are agnostic to temporal dimension. We advocate using the MOVIE metric[9] which considers spatial, temporal and spatio-temporal aspects by comparing original and reconstructed sequence. In Table 4 we report average MOVIE metric on VidTIMIT test set. Our proposed modifications to [3] significantly improves reconstructed video quality.



**Fig. 5:** Visualizing benefit of **proposed model** (bottom row) over Yeh *et al.*[3] (middle row) on pseudo sequences. A pseudo sequence is created from a single image of a person but masked with different masks (thereby mimicking a temporal aspect). An ideal sequence inpainting model should result in identical outputs for a given subject. Note that our sequence reconstructions lead to more temporally consistent (notice the lips, eyes) solutions.

## 6. CONCLUSION

In this paper, we first discussed the problem of impractical long inference time of the recent completely unsupervised

inpainting framework of [3]. We then proposed to speedup the iterative optimization of [3] by better initialization technique on images and also leveraging temporal redundancy in videos. In the process, we achieved a significant speedup. Several comparisons were also done with current hybrid benchmarks and we achieved comparable performance. Future work might consider replacing the iterative optimization by learning to project a damaged image into noise prior space.

## 7. REFERENCES

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman, “Patchmatch: A randomized correspondence algorithm for structural image editing,” *ACM Transactions on Graphics (TOG)*, vol. 28, no. 3, pp. 24, 2009.
- [2] James Hays and Alexei A Efros, “Scene completion using millions of photographs,” in *ACM Transactions on Graphics (TOG)*. ACM, 2007, vol. 26, p. 4.
- [3] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do, “Semantic image inpainting with deep generative models,” in *CVPR*, 2017, pp. 5485–5493.
- [4] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa, “Globally and locally consistent image completion,” *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 107, 2017.
- [5] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang, “Generative image inpainting with contextual attention,” in *CVPR*, 2018.
- [6] Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang, “Generative face completion,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, vol. 1, p. 3.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *NIPS*, 2014, pp. 2672–2680.
- [8] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros, “Context encoders: Feature learning by inpainting,” in *CVPR*, 2016, pp. 2536–2544.
- [9] Kalpana Seshadrinathan and Alan Conrad Bovik, “Motion tuned spatio-temporal quality assessment of natural videos,” *IEEE transactions on image processing*, vol. 19, no. 2, pp. 335–350, 2010.
- [10] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros, “Generative visual manipulation on the natural image manifold,” in *ECCV*. Springer, 2016, pp. 597–613.
- [11] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng, “Reading digits in natural images with unsupervised feature learning,” in *NIPS workshop on deep learning and unsupervised feature learning*, 2011, vol. 2011, p. 5.
- [12] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei, “3d object representations for fine-grained categorization,” in *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- [13] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang, “Deep learning face attributes in the wild,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3730–3738.
- [14] Conrad Sanderson and Brian C Lovell, “Multi-region probabilistic histograms for robust and scalable identity inference,” in *International Conference on Biometrics*. Springer, 2009, pp. 199–208.
- [15] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Ballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al., “Photo-realistic single image super-resolution using a generative adversarial network,” in *CVPR*, 2017, vol. 2, p. 4.

*Accepted as a full paper at IEEE ICIP, 2019, Taiwan.*