Computational Visual Media
DOI 10.1007/s41095-xxx-xxxx-x Vol. x, No. x, month year, xx–xx

Article Type (Research/Review)

iSeeBetter: Spatio-temporal video super-resolution using recurrent generative back-projection networks

Aman Chadha[1] (✉), John Britto[2], and M. Mani Roja[3]

© The Author(s) 2020. This article is published with open access at Springerlink.com

arXiv:2006.11161v1 [cs.CV] 13 Jun 2020

Abstract Recently, learning-based models have enhanced the performance of single-image super-resolution (SISR). However, applying SISR successively to each video frame leads to a lack of temporal coherency. Convolutional neural networks (CNNs) outperform traditional approaches in terms of image quality metrics such as peak signal to noise ratio (PSNR) and structural similarity (SSIM). However, generative adversarial networks (GANs) offer a competitive advantage by being able to mitigate the issue of a lack of finer texture details, usually seen with CNNs when super-resolving at large upscaling factors. We present iSeeBetter, a novel GAN-based spatio-temporal approach to video super-resolution (VSR) that renders temporally consistent super-resolution videos. iSeeBetter extracts spatial and temporal information from the current and neighboring frames using the concept of recurrent back-projection networks as its generator. Furthermore, to improve the "naturality" of the super-resolved image while eliminating artifacts seen with traditional algorithms, we utilize the discriminator from super-resolution generative adversarial network (SRGAN). Although mean squared error (MSE) as a primary loss-minimization objective improves PSNR/SSIM, these metrics may not capture fine details in the image resulting in misrepresentation of perceptual quality. To address this, we use a four-fold (MSE, perceptual, adversarial, and total-variation (TV)) loss function. Our results demonstrate that iSeeBetter offers superior VSR fidelity and surpasses state-of-the-art performance.

Keywords super resolution; video upscaling; frame recurrence; optical flow; generative adversarial networks; convolutional neural networks.

1 Introduction

The goal of super-resolution (SR) is to enhance a low resolution (LR) image to a higher resolution (HR) image by filling in missing fine-grained details in the LR image. The domain of SR research can be divided into three main areas: single image SR (SISR) [6, 15, 17, 27], multi image SR (MISR) [9, 10] and video SR (VSR) [3, 16, 23, 38, 43].

Consider an LR video source which consists of a sequence of LR video frames LR_{t-n}, ..., LR_t, ..., LR_{t+n}, where we super-resolve a target frame LR_t. The idea behind SISR is to super-resolve LR_t by utilizing spatial information inherent in the frame, independently of other frames in the video sequence. However, this technique fails to exploit the temporal details inherent in a video sequence resulting in temporal incoherence. MISR seeks to address just that – it utilizes the missing details available from the neighboring frames LR_{t-n}, ..., LR_t, ..., LR_{t+n} and fuses them for super-resolving LR_t. After spatially aligning frames, missing details are extracted by separating differences between the aligned frames

1 Department of Computer Science, Stanford University, 450 Serra Mall, Stanford, CA 94305, USA. E-mail: aman@amanchadha.com; amanc@stanford.edu (✉).

2 Department of Computer Science, University of Massachusetts Amherst, Amherst, MA 01003, USA. E-mail: jnadar@umass.edu.

3 Department of Electronics and Telecommunication Engineering, University of Mumbai, Mumbai, Maharashtra 400032, India. E-mail: maniroja@tsec.edu

Manuscript received: 2020-02-20; accepted: 2020-04-24.

from missing details observed only in one or some of the frames. However, in MISR, the alignment of the frames is done without any concern for temporal smoothness, which is in stark contrast to VSR where the frames are typically aligned in temporal smooth order.

Traditional VSR methods upscale based on a single degradation model (usually bicubic interpolation) followed by reconstruction. This is sub-optimal and adds computational complexity [39]. Recently, learning-based models that utilize convolutional neural networks (CNNs) have outperformed traditional approaches in terms of widely-accepted image reconstruction metrics such as peak signal to noise ratio (PSNR) and structural similarity (SSIM).

In some recent VSR methods that utilize CNNs, frames are concatenated [23] or fed into recurrent neural networks (RNNs) [20] in temporal order, without explicit alignment. In other methods, the frames are aligned explicitly, using motion cues between temporal frames with the alignment modules [3, 33, 38, 43]. The latter set of methods generally render temporally smoother results compared to the methods with no explicit spatial alignment [20, 31]. However, these VSR methods suffer from a number of problems. In the frame-concatenation approach [3, 23, 33], many frames are processed simultaneously in the network, resulting in significantly higher network training times. With methods that use RNNs [20, 38, 43], modeling both subtle and significant changes simultaneously (e.g., slow and quick motions of foreground objects) is a challenging task even if long short-term memory units (LSTMs) are deployed, which are designed for maintaining long-term temporal dependencies [12]. A crucial aspect of an effective VSR system is the ability to handle motion sequences, which are often integral components of videos [3, 34].

The proposed method, iSeeBetter, is inspired by recurrent back-projection networks (RBPNs) [16] which utilize "back-projection" as their underpinning approach, originally introduced in [21, 22] for MISR. The basic concept behind back-projection is to iteratively calculate residual images as reconstruction error between a target image and a set of neighboring images. The residuals are then back-projected to the target image for improving super-resolution accuracy. The multiple residuals enable representation of subtle and significant differences between the target frame and its adjacent frames, thus exploiting temporal relationships between adjacent frames as shown in Fig. 1. Deep back-projection networks (DBPNs) [15] use back-projection to perform SISR using learning-based methods by estimating the output frame SR_t using the corresponding LR_t frame. To this end, DBPN produces a high-resolution feature map that is iteratively refined through multiple up- and down-sampling layers. RBPN offers superior results by combining the benefits of the original MISR back-projection approach with DBPN. Specifically, RBPN uses the idea of iteratively refining HR feature maps from DBPN, but extracts missing details using neighboring video frames like the original back-projection technique [21, 22]. This results in superior SR accuracy.

To mitigate the issue of a lack of finer texture details when super-resolving at large upscaling factors that is usually seen with CNNs [30], iSeeBetter utilizes GANs with a loss function that weighs adversarial loss, perceptual loss [30], mean square error (MSE)-based loss and total-variation (TV) loss [37]. Our approach combines the merits of RBPN and SRGAN [30] – it is based on RBPN as its generator and is complemented by SRGAN's discriminator architecture, which is trained to differentiate between super-resolved images and original photo-realistic images. Blending these techniques yields iSeeBetter, a state-of-the-art system that is able to recover precise photo-realistic textures and motion-based scenes from heavily down-sampled videos.

Our contributions include the following key innovations.

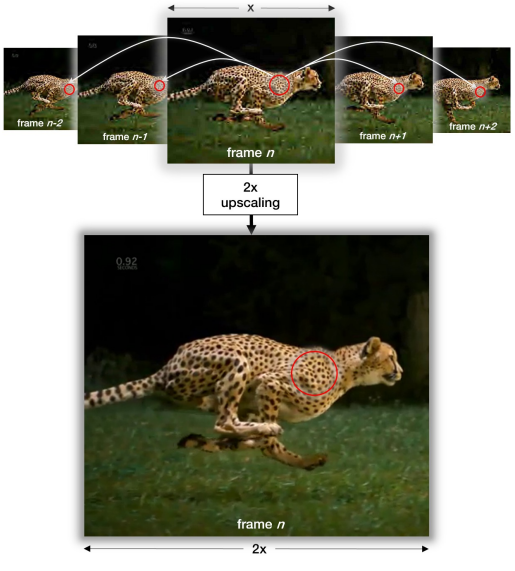

Fig. 1 Temporal relationships between adjacent frames.

Combining the state-of-the-art in SR: We propose a model that leverages two superior SR techniques – (i) RBPN, which is based on the idea of integrating SISR and MISR in a unified VSR framework using back-projection and, (ii) SRGAN, which is a framework capable of inferring photo-realistic natural images. RBPN enables iSeeBetter to extract details from neighboring frames, complemented by the generator-discriminator architecture in GANs which pushes iSeeBetter to generate more realistic and appealing frames while eliminating artifacts seen with traditional algorithms [47]. iSeeBetter thus yields more than the sum of the benefits of RBPN and SRGAN.

"Optimizing" the loss function: Pixel-wise loss functions such as L1 loss, used in RBPN [16], struggle to handle the uncertainty inherent in recovering lost high-frequency details such as complex textures that commonly exist in many videos. Minimizing MSE encourages finding pixel-wise averages of plausible solutions that are typically overly-smooth and thus have poor perceptual quality [2, 7, 24, 36]. To address this, we adopt a four-fold (MSE, perceptual, adversarial, and TV) loss function for superior results. Similar to SRGAN [30], we utilize a loss function that optimizes perceptual quality by minimizing adversarial loss and MSE loss. Adversarial loss helps improve the "naturality" associated with the output image using the discriminator. On the other hand, MSE loss focuses on optimizing perceptual similarity instead of similarity in pixel space. Furthermore, we use a de-noising loss function called TV loss [1]. We carried out experiments comparing L1 loss with our four-fold loss and found significant improvements with the latter (cf. Section 4).

Extended evaluation protocol: To evaluate iSeeBetter, we used standard datasets: Vimeo90K [48], Vid4 [32] and SPMCS [43]. Since Vid4 and SPMCS lack significant motion sequences, we included Vimeo90K, a dataset containing various types of motion. This enabled us to conduct a more holistic evaluation of the strengths and weaknesses of iSeeBetter. To make iSeeBetter more robust and enable it to handle real-world videos, we expanded the spectrum of data diversity and wrote scripts to collect additional data from YouTube. As a result, we augmented our dataset to about 170,000 clips.

User-friendly infrastructure: We built several useful tools to download and structure datasets, visualize temporal profiles of intermediate blocks and the output, and run predefined benchmark sequences on a trained model to be able to iterate on different models quickly. In addition, we built a video-to-frames tool to directly input videos to iSeeBetter, rather than frames. We also ensured our script infrastructure is flexible (such that it supports a myriad of options) and can be easily leveraged. The code and pre-trained models are available at `https://iseebetter.amanchadha.com`.

2 Related work

Since the seminal work by Tsai on image registration [44] two decades ago, many SR techniques based on various underlying principles have been proposed. Initial methods included spatial or frequency domain signal processing, statistical models and interpolation approaches [49]. In this section, we focus our discussion on learning-based methods which have emerged as superior VSR techniques compared to traditional statistical methods.

2.1 Deep SISR

First introduced by SRCNN [6], deep SISR required a predefined up-sampling operator. Further improvements in this field include better up-sampling layers [39], residual learning [42], back-projection [15], recursive layers [28], and progressive up-sampling [29]. A significant milestone in SR research was the introduction of a GAN-powered SR approach [30], which achieved state-of-the-art performance.

2.2 Deep VSR

Deep VSR can be divided into five types based on the approach to preserving temporal information.

(a) Temporal Concatenation. The most popular approach to retain temporal information in VSR is concatenating multiple frames [3, 23, 26, 31]. This approach can be seen as an extension of SISR to accept multiple input images. However, this approach fails to represent multiple motion regimes within a single input sequence since the input frames are simply concatenated together.

(b) Temporal Aggregation. To address the dynamic motion problem in VSR, [33] proposed multiple SR inferences which work on different motion regimes. The final layer aggregates the outputs of all branches to construct SR frame. However, this approach still concatenates many input frames, resulting in lengthy convergence during global optimization.

Tab. 1 Datasets used for training and evaluation

Dataset	Resolution	# of Clips	# of Frames/Clip	# of Frames
Vimeo90K	448 × 256	13,100	7	91,701
SPMCS	240 × 135	30	31	930
Vid4	(720 × 576 × 3) × 2, (704 × 576 × 3) × 2	4	41, 34, 49, 47	684
Augmented	960 × 720	7,000	110	77,000
Total	-	**46,034**	-	**170,315**

(c) Recurrent Networks. RNNs deal with temporal inputs and/or outputs and have been deployed in a myriad of applications ranging from video captioning [25, 35, 50], video summarization [5, 45], and VSR [20, 38, 43]. Two types of RNNs have been used for VSR. A many-to-one architecture is used in [20, 43] where a sequence of LR frames is mapped to a single target HR frame. A many-to-many RNN has recently been used by [38] where an optical flow network to accepts LR_{t-1} and LR_t, which is fed to an SR network along with LR_t. This approach was first proposed by [20] using bidirectional RNNs. However, the network has a small network capacity and has no frame alignment step. Further improvement is proposed by [43] using a motion compensation module and a ConvLSTM layer [40].

(d) Optical Flow-Based Methods. The above methods estimate a single HR frame by combining a batch of LR frames and are thus computationally expensive. They often result in unwanted flickering artifacts in the output frames [37]. To address this, [38] proposed a method that utilizes a network trained on estimating the optical flow along with the SR network. Optical flow methods allow estimation of the trajectories of moving objects, thereby assisting in VSR. [26] warp video frames LR_{t-1} and LR_{t+1} onto LR_t using the optical flow method of [8], concatenate the three frames, and pass them through a CNN that produces the output frame SR_{t+1}. [3] follow the same approach but replace the optical flow model with a trainable motion compensation network.

(e) Pre-Training then Fine-Tuning v/s End-to-End Training. While most of the above-mentioned methods are end-to-end trainable, certain approaches first pre-train each component before fine-tuning the system as a whole in a final step [3, 33, 43].

Our approach is a combination of (i) an RNN-based optical flow method that preserves spatio-temporal information in the current and adjacent frames as the generator and, (ii) a discriminator that is adept at ensuring the generated SR frame offers superior fidelity.

3 Methods

3.1 Datasets

To train iSeeBetter, we amalgamated diverse datasets with differing video lengths, resolutions, motion sequences, and number of clips. Tab. 1 presents a summary of the datasets used. When training our model, we generated the corresponding LR frame for each HR input frame by performing 4× down-sampling using bicubic interpolation. We thus perform self-supervised learning by automatically generating the input-output pairs for training without any human intervention. To further extend our dataset, we wrote scripts to collect additional data from YouTube. The dataset was shuffled for training and testing. Our training/validation/test split was 80%/10%/10%.

3.2 Network architecture

Fig. 2 shows the iSeeBetter architecture that consists of RBPN [16] and SRGAN [30] as its generator and discriminator respectively. Tab. 2 shows our notational convention. RBPN has two approaches that extract missing details from different sources: SISR and MISR. Fig. 3 shows the horizontal flow (represented by blue arrows in Fig. 2) that enlarges LR_t using SISR. Fig. 4 shows the vertical flow (represented by red arrows in Fig. 2) which is based on MISR that computes residual features from (i) a pair of LR_t and its neighboring frames (LR_{t-1}, ..., LR_{t-n}) coupled with, (ii) the pre-computed dense motion flow maps (F_{t-1}, ..., F_{t-n}).

Tab. 2 Adopted notation

HR_t	input high resolution image
LR_t	low resolution image (derived from HR_t)
F_t	optical flow output
H_{t-1}	residual features extracted from (LR_{t-1}, F_{t-1}, LR_t)
SR_t	estimated HR output

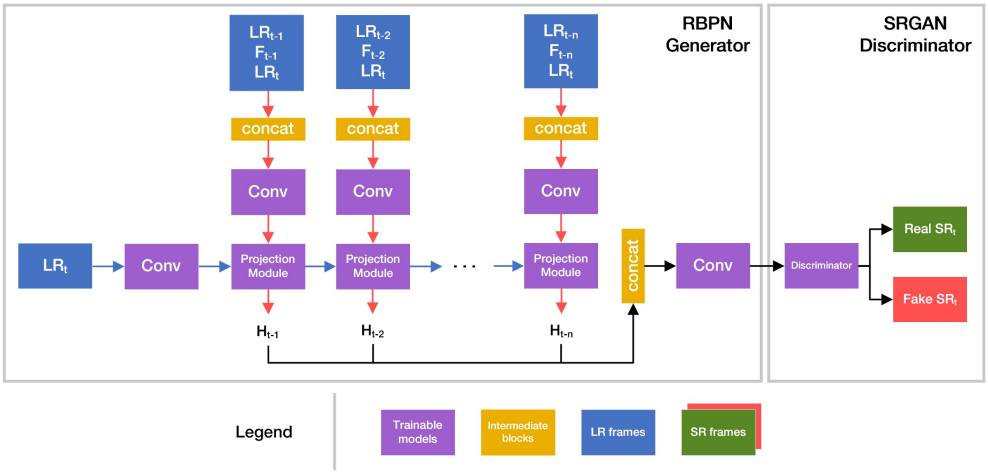

Fig. 2 Overview of iSeeBetter

At each projection step, RBPN observes the missing details from LR_t and extracts residual features from neighboring frames to recover details. The convolutional layers that feed the projection modules in Fig. 2 thus serve as initial feature extractors. Within the projection modules, RBPN utilizes a recurrent encoder-decoder mechanism for fusing details extracted from adjacent frames in SISR and MISR and incorporates them into the estimated frame SR_t through back-projection. The convolutional layer that operates on the concatenated output from all the projection modules is responsible for generating SR_t. Once SR_t is synthesized, it is sent over to the discriminator (shown in Fig. 5) to validate its "authenticity".

3.3 Loss functions

The perceptual image quality of the resulting SR image is dependent on the choice of the loss function. To evaluate the quality of an image, MSE is the most commonly used loss function in a wide variety of state-of-the-art SR approaches, which aims to improve the PSNR of an image [19]. While optimizing MSE during training improves PSNR and SSIM, these metrics may not capture fine details in the image leading to misrepresentation of perceptual quality [30]. The ability of MSE to capture intricate texture details based on pixel-wise frame differences is very limited, and can cause the resulting video frames to be overly-smooth [4]. In a series of experiments, it was found that even manually distorted images had an MSE score comparable to the original image [46]. To address this, iSeeBetter uses a four-fold (MSE, perceptual, adversarial, and TV) loss instead of solely relying on pixel-wise MSE loss. We weigh these losses together as a final evaluation standard for training iSeeBetter, thus taking into account both pixel-wise similarities and high-level features. Fig. 6 shows the individual components of the iSeeBetter loss function.

3.3.1 MSE loss

We use pixel-wise MSE loss (also called content loss [30]) for the estimated frame SR_t against the ground truth HR_t.

$$MSE_t = \frac{1}{WH}\sum_{x=0}^{W}\sum_{y=0}^{H}\left((HR_t)_{x,y} - G_{\theta_G}(LR_t)_{x,y}\right)^2 \quad (1)$$

where, $G_{\theta_G}(LR_t)$ is the estimated frame SR_t. W and H represent the width and height of the frames respectively.

3.3.2 Perceptual loss

[2, 11] introduced a new loss function called perceptual loss, also used in [24, 30], which focuses on perceptual similarity instead of similarity in pixel space. Perceptual loss relies on features extracted from the activation layers of the pre-trained VGG-19 network in [41], instead of low-level pixel-wise error measures. We define perceptual loss as the euclidean distance between the feature representations of the estimated SR image $G_{\theta_G}(LR_t)$ and the ground truth HR_t.

$$PerceptualLoss_t = \frac{1}{W_{i,j}H_{i,j}}\sum_{x=1}^{W_{i,j}}\sum_{y=1}^{H_{i,j}}\left(\begin{array}{l} VGG_{i,j}(HR_t)_{x,y} - \\ VGG_{i,j}(G_{\theta_G}(LR_t))_{x,y} \end{array}\right)^2 \quad (2)$$

where, $VGG_{i,j}$ denotes the feature map obtained by the j^{th} convolution (after activation) before the i^{th} maxpooling layer in the VGG-19 network. $W_{i,j}$ and $H_{i,j}$ are the dimensions of the respective feature maps in the VGG-19 network.

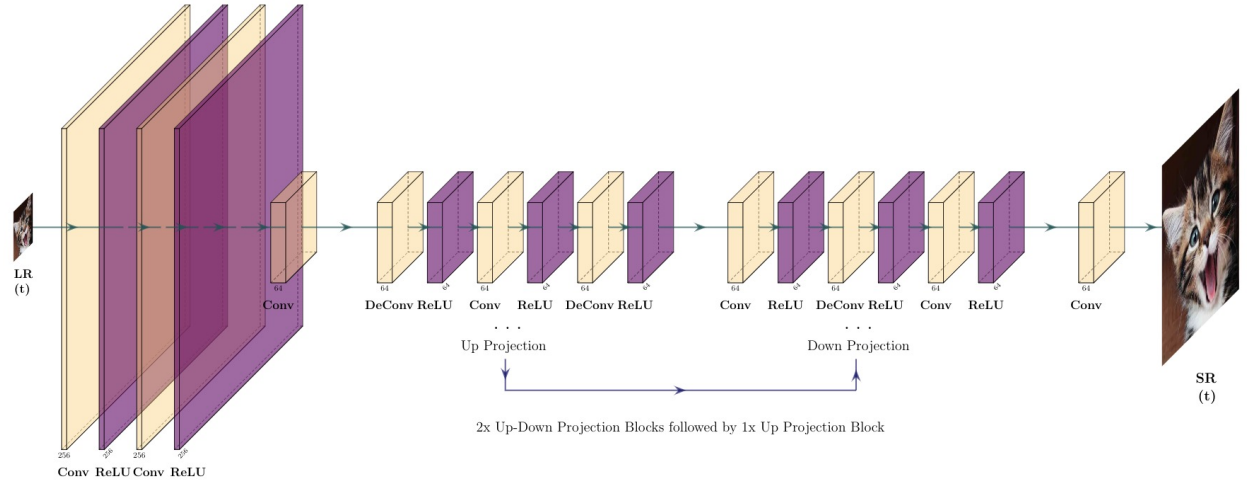

Fig. 3 DBPN [15] architecture for SISR, where we perform up-down-up sampling using 8 × 8 kernels with a stride of 4 and padding of 2. Similar to the ResNet architecture above, the DBPN network also uses Parametric ReLUs [18] as its activation functions.

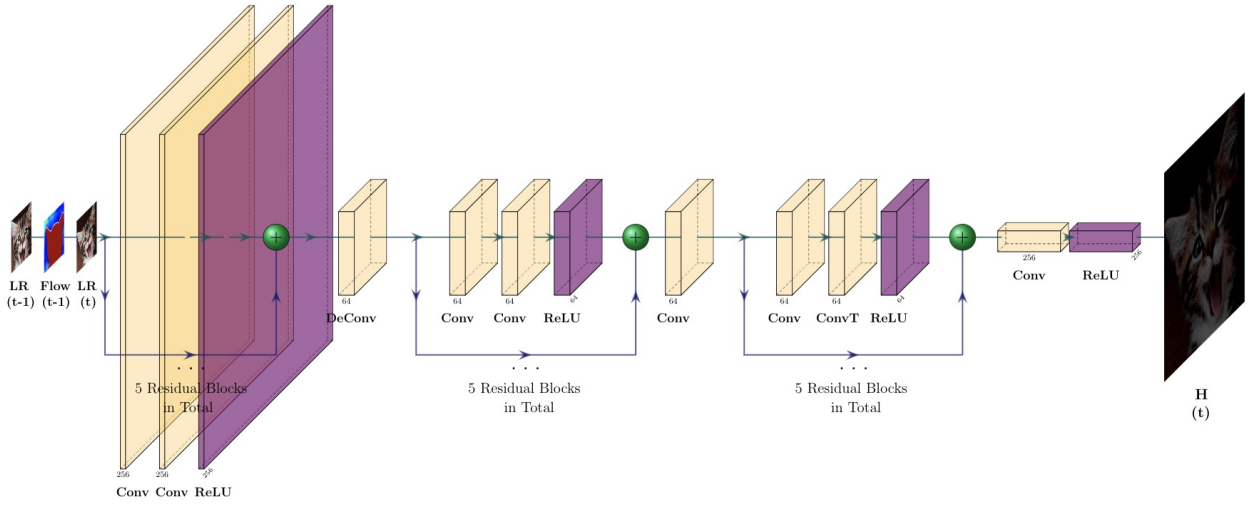

Fig. 4 ResNet architecture for MISR that is composed of three tiles of five blocks where each block consists of two convolutional layers with 3 × 3 kernels, a stride of 1 and padding of 1. The network uses Parametric ReLUs [18] for its activations.

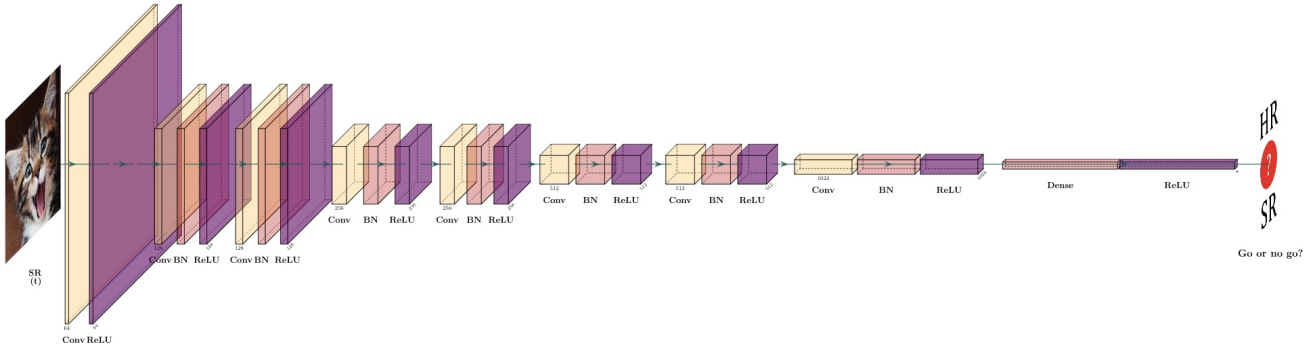

Fig. 5 Discriminator Architecture from SRGAN [30]. The discriminator uses Leaky ReLUs for computing its activations.

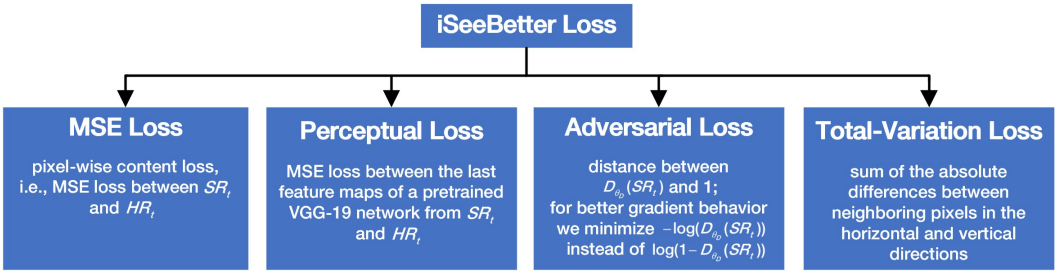

Fig. 6 The MSE, perceptual, adversarial, and TV loss components of the iSeeBetter loss function.

3.3.3 Adversarial loss

We use the generative component of iSeeBetter as the adversarial loss to limit model "fantasy", thus improving the "naturality" associated with the super-resolved image. Adversarial loss is defined as:

$$AdversarialLoss_t = -log(D_{\theta_D}(G_{\theta_G}(LR_t)) \quad (3)$$

where, $D_{\theta_D}(G_{\theta_G}(LR_t))$ is the discriminator's output probability that the reconstructed image $G_{\theta_G}(LR_t)$ is a real HR image. We minimize $-log(D_{\theta_D}(G_{\theta_G}(LR_t))$ instead of $log(1 - D_{\theta_D}(G_{\theta_G}(LR_t))$ for better gradient behavior [13].

3.3.4 Total-Variation loss

TV loss was introduced as a loss function in the domain of SR by [1]. It is defined as the sum of the absolute differences between neighboring pixels in the horizontal and vertical directions [47]. Since TV loss measures noise in the input, minimizing it as part of our overall loss objective helps de-noise the output SR image and thus encourages spatial smoothness. TV loss is defined as follows:

$$TVLoss_t = \frac{1}{WH}\sum_{i=0}^{W}\sum_{j=0}^{H}\sqrt{\begin{array}{l}(G_{\theta_G}(LR_t)_{i,j+1,k} - G_{\theta_G}(LR_t)_{i,j,k})^2 + \\ (G_{\theta_G}(LR_t)_{i+1,j,k} - G_{\theta_G}(LR_t)_{i,j,k})^2\end{array}} \quad (4)$$

3.3.5 Loss formulation

We define our overall loss objective for each frame as the weighted sum of the MSE, adversarial, perceptual, and TV loss components:

$$Loss_{G_{\theta_G}}(SR_t) = \begin{array}{l}\alpha \times MSE(SR_t, HR_t) \\ +\beta \times PerceptualLoss(SR_t, HR_t) \\ +\gamma \times AdversarialLoss(SR_t) \\ +\delta \times TVLoss(SR_t, HR_t)\end{array} \quad (5)$$

where, α, β, γ, δ are weights set as 1, 6×10^{-3}, 10^{-3} and 2×10^{-8} respectively [14].

The discriminator loss for each frame is as follows:

$$Loss_{D_{\theta_D}}(SR_t) = 1 - D_{\theta_D}(HR_t) + D_{\theta_D}(SR_t) \quad (6)$$

The total loss of an input sample is the average loss of all frames.

$$\begin{array}{l}Loss_{G_{\theta_G}} = \frac{1}{N}\sum_{t=1}^{N}(Loss_{G_{\theta_G}}(SR_t)) \\ Loss_{D_{\theta_D}} = \frac{1}{N}\sum_{t=1}^{N}(Loss_{D_{\theta_D}}(SR_t))\end{array} \quad (7)$$

4 Experimental evaluation

To train the model, we used an Amazon EC2 P3.2xLarge instance with an NVIDIA Tesla V100 GPU with 16GB VRAM, 8 vCPUs and 64GB of host memory. We used the hyperparameters from RBPN and SRGAN. Tab. 3 compares iSeeBetter with six state-of-the-art VSR algorithms: DBPN [15], B_{123} + T [33], DRDVSR [43], FRVSR [38], VSR-DUF [23] and RBPN/6-PF [16]. Tab. 4 offers a visual analysis of VSR-DUF and iSeeBetter. Tab. 5 shows ablation studies to assess the impact of using a generator-discriminator architecture and the four-fold loss as design decisions.

Tab. 3 PSNR/SSIM evaluation of state-of-the-art VSR algorithms using Vid4 and Vimeo90K for 4× upscaling. Bold numbers indicate best performance.

Dataset	Clip Name	Flow	Bicubic	DBPN [15]	B_{123} + T [33]	DRDVSR [43]	FRVSR [38]	RBPN/6-PF [16]	VSR-DUF [23]	**iSeeBetter**
Vid4	Calendar	1.14	19.82/0.554	22.19/0.714	21.66/0.704	22.18/0.746	-	23.99/0.807	24.09/0.813	**24.13/0.817**
	City	1.63	24.93/0.586	26.01/0.684	26.45/0.720	26.98/0.755	-	27.73/0.803	28.26/0.833	**28.34/0.841**
	Foliage	1.48	23.42/0.575	24.67/0.662	24.98/0.698	25.42/0.720	-	26.22/0.757	26.38/0.771	**26.57/0.773**
	Walk	1.44	26.03/0.802	28.61/0.870	28.26/0.859	28.92/0.875	-	**30.70/0.909**	30.50/0.912	30.68/0.908
Average		1.42	23.53/0.629	25.37/0.737	25.34/0.745	25.88/0.774	26.69/0.822	27.12/0.818	27.31/0.832	**27.43/0.835**
Vimeo90K	Fast Motion	8.30	34.05/0.902	37.46/0.944	-	-	-	40.03/0.960	37.49/0.949	**40.17/0.971**

5 Conclusions and future work

We proposed iSeeBetter, a novel spatio-temporal approach to VSR that uses recurrent-generative back-projection networks. iSeeBetter couples the virtues of RBPN and SRGAN. RBPN enables iSeeBetter to generate superior SR images by combining spatial and temporal information from the input and neighboring frames. In addition, SRGAN's discriminator architecture fosters generation of photo-realistic frames. We used a four-fold loss function that emphasizes perceptual quality. Furthermore, we proposed a new evaluation protocol for video SR by collating diverse datasets. With extensive experiments, we assessed the role played by various design choices in the ultimate performance of iSeeBetter, and demonstrated that on a vast majority of test video sequences, iSeeBetter advances the state-of-the-art.

To improve iSeeBetter, a couple of ideas could be explored. In visual imagery the foreground recieves much more attention than the background since it typically includes subjects such as humans. To improve perceptual quality, we can segment the foreground and background, and make iSeeBetter perform "adaptive VSR" by utilizing different policies for the foreground and background. For instance, we could adopt a wider span of the number of frames to extract details from for the foreground compared to the background. Another idea is to decompose a video sequence into scenes on the basis of frame-similarity and make iSeeBetter assign weights to adjacent frames based on which scene they belong to. Adjacent frames from a different scene can be weighed lower compared to frames from the same scene, thereby making iSeeBetter focus on extracting details from frames within the same scene – à la the concept of attention applied to VSR.

Tab. 4 Visually inspecting examples from Vid4, SPMCS and Vimeo-90k comparing VSR-DUF and iSeeBetter. We chose VSR-DUF for comparison because it was the state-of-the-art at the time of publication. Top row: fine-grained textual features that help with readability; middle row: intricate high-frequency image details; bottom row: camera panning motion.

Dataset	Clip Name	VSR-DUF [23]	**iSeeBetter**	Ground Truth
Vid4	Calendar			
SPMCS	Pagoda			
Vimeo90K	Motion			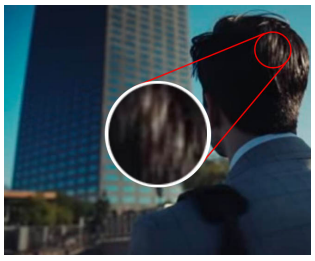

Acknowledgements

The authors would like to thank Andrew Ng's lab at Stanford University for their guidance on this project. In particular, the authors express their gratitude to Mohamed El-Geish for the idea-inducing brainstorming sessions throughout the project.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

[1] H. A. Aly and E. Dubois. Image up-sampling using total-variation regularization with a new observation model. *IEEE Transactions on Image Processing*, 14(10):1647–1659, 2005.

[2] J. Bruna Estrach, P. Sprechmann, and Y. LeCun. Super-resolution with deep convolutional sufficient statistics. 1 2016. 4th International Conference on Learning Representations (ICLR) 2016.

[3] J. Caballero, C. Ledig, A. Aitken, A. Acosta, J. Totz, Z. Wang, and W. Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4778–4787, 2017.

[4] M.-H. Cheng, N.-W. Lin, K.-S. Hwang, and J.-H. Jeng. Fast video super-resolution using artificial neural networks. In *2012 8th International Symposium on Communication Systems, Networks & Digital Signal Processing (CSNDSP)*, pages 1–4. IEEE, 2012.

[5] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.

[6] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015.

[7] A. Dosovitskiy and T. Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in neural information processing systems*, pages 658–666, 2016.

[8] M. Drulea and S. Nedevschi. Total variation regularization of local-global optical flow. In *2011 14th International IEEE Conference on Intelligent Transportation Systems*, pages 318–323. IEEE, 2011.

[9] E. Faramarzi, D. Rajan, and M. P. Christensen. Unified blind method for multi-image super-resolution and single/multi-image blur deconvolution. *IEEE Transactions on Image Processing*, 22(6):2101–2114, 2013.

[10] D. C. Garcia, C. Dorea, and R. L. de Queiroz. Super resolution for multiview images using depth information. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(9):1249–1256, 2012.

[11] L. Gatys, A. S. Ecker, and M. Bethge. Texture synthesis using convolutional neural networks. In *Advances in neural information processing systems*, pages 262–270, 2015.

[12] F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: Continual prediction with lstm. 1999.

[13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[14] J. Hany and G. Walters. *Hands-On Generative Adversarial Networks with PyTorch 1. x: Implement next-generation neural networks to build powerful GAN models using Python*. Packt Publishing Ltd, 2019.

[15] M. Haris, G. Shakhnarovich, and N. Ukita. Deep back-projection networks for super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1664–1673, 2018.

[16] M. Haris, G. Shakhnarovich, and N. Ukita. Recurrent back-projection network for video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3897–3906,

Tab. 5 Ablation analysis for iSeeBetter using the "City" clip from Vid4.

iSeeBetter Config	PSNR
RBPN baseline with L1 loss	27.73
RBPN baseline with MSE loss	27.77
RBPN generator + SRGAN discriminator with adversarial loss	28.08
RBPN generator + SRGAN discriminator with adversarial + MSE loss	28.12
RBPN generator + SRGAN discriminator with adversarial + MSE + perceptual loss	28.27
RBPN generator + SRGAN discriminator with adversarial + MSE + perceptual + TV loss	28.34

2019.

[17] M. Haris, M. R. Widyanto, and H. Nobuhara. Inception learning super-resolution. *Applied optics*, 56(22):6043–6048, 2017.

[18] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[19] A. Hore and D. Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th International Conference on Pattern Recognition*, pages 2366–2369. IEEE, 2010.

[20] Y. Huang, W. Wang, and L. Wang. Bidirectional recurrent convolutional networks for multi-frame super-resolution. In *Advances in Neural Information Processing Systems*, pages 235–243, 2015.

[21] M. Irani and S. Peleg. Improving resolution by image registration. *CVGIP: Graphical models and image processing*, 53(3):231–239, 1991.

[22] M. Irani and S. Peleg. Motion analysis for image enhancement: Resolution, occlusion, and transparency. *Journal of Visual Communication and Image Representation*, 4(4):324–335, 1993.

[23] Y. Jo, S. Wug Oh, J. Kang, and S. Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3224–3232, 2018.

[24] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.

[25] J. Johnson, A. Karpathy, and L. Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4565–4574, 2016.

[26] A. Kappeler, S. Yoo, Q. Dai, and A. K. Katsaggelos. Video super-resolution with convolutional neural networks. *IEEE Transactions on Computational Imaging*, 2(2):109–122, 2016.

[27] J. Kim, J. Kwon Lee, and K. Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.

[28] J. Kim, J. Kwon Lee, and K. Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1637–1645, 2016.

[29] W.-S. Lai, J.-B. Huang, N. Ahuja, and M.-H. Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 624–632, 2017.

[30] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.

[31] R. Liao, X. Tao, R. Li, Z. Ma, and J. Jia. Video super-resolution via deep draft-ensemble learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 531–539, 2015.

[32] C. Liu and D. Sun. A bayesian approach to adaptive video super resolution. In *CVPR 2011*, pages 209–216. IEEE, 2011.

[33] D. Liu, Z. Wang, Y. Fan, X. Liu, Z. Wang, S. Chang, and T. Huang. Robust video super-resolution with learned temporal dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2507–2515, 2017.

[34] O. Makansi, E. Ilg, and T. Brox. End-to-end learning of video super-resolution with motion compensation. In *German conference on pattern recognition*, pages 203–214. Springer, 2017.

[35] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). *arXiv preprint arXiv:1412.6632*, 2014.

[36] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.

[37] H. Ren and X. Fang. Recurrent back-projection network for video super-resolution. *Final Project for MIT 6.819 Advances in Computer Vision*, 2018.

[38] M. S. Sajjadi, R. Vemulapalli, and M. Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6626–6634, 2018.

[39] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.

[40] X. Shi, Z. Chen, H. Wang, D. Yeung, W. Wong, and W. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 1–9, 2015.

[41] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[42] Y. Tai, J. Yang, and X. Liu. Image super-resolution via deep recursive residual network. In *Proceedings of the IEEE conference on computer vision and pattern*

recognition, pages 3147–3155, 2017.

[43] X. Tao, H. Gao, R. Liao, J. Wang, and J. Jia. Detail-revealing deep video super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4472–4480, 2017.

[44] R. Tsai. Multiframe image restoration and registration. *Advance Computer Visual and Image Processing*, 1:317–339, 1984.

[45] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. pages 1494–1504, 2015.

[46] Z. Wang and A. C. Bovik. A universal image quality index. *IEEE signal processing letters*, 9(3):81–84, 2002.

[47] Z. Wang, J. Chen, and S. C. Hoi. Deep learning for image super-resolution: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[48] T. Xue, B. Chen, J. Wu, D. Wei, and W. T. Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019.

[49] J. Yang and T. Huang. Image super-resolution: Historical overview and future challenges. *Super-resolution imaging*, pages 20–34, 2010.

[50] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu. Video paragraph captioning using hierarchical recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4584–4593, 2016.

Aman Chadha has held positions at some of the world's leading semiconductor/product companies. He is currently based out of Cupertino (Silicon Valley), California and is currently pursuing his graduate studies in Artificial Intelligence from Stanford University. He has published in prestigious international journals and conferences, and has authored two books. His publications have garnered about 200 citations. He currently serves on the editorial boards of several international journals including IJATCA, IJLTET, IJCET, IJEACS and IJRTER. He has served as a reviewer for IJEST, IJCST, IJCSEIT and JESTEC. Aman graduated with an M.S. from the University of Wisconsin-Madison with an outstanding graduate student award in 2014 and his B.E. with distinction from the University of Mumbai in 2012. His research interests include Computer Vision (particularly, Pattern Recognition), Artificial Intelligence, Machine Learning and Computer Architecture. Aman has 18 publications to his credit.

John Britto is pursuing his M.S. in Computer Science from the University of Massachusetts, Amherst. He completed his B.E. in Computer Engineering from the University of Mumbai in 2018. His research interests lie in Machine Learning, Natural Language Processing, and Artificial Intelligence.

M. Mani Roja is a full professor in the Electronics and Telecommunication Department at the University of Mumbai since the past 30 years. She received her Ph.D. in Electronics and Telecommunication Engineering from Sant Gadge Baba Amravati University and her Masters in Electronics and Telecommunication Engineering from the University of Mumbai. She has collaborated across the years in the fields of Image Processing, Speech Processing, and Biometric recognition.