



# End-to-End Image Super-Resolution via Deep and Shallow Convolutional Networks

<sup>1</sup>Dr. J Rajaram Professor,

[drrajram81@gmail.com](mailto:drrajram81@gmail.com),

<sup>2</sup>E Krishna Assistant Professor,

[krishna.cseit@gmail.com](mailto:krishna.cseit@gmail.com),

<sup>3</sup>Anishetty Shiva Rama Krishna Associate Professor,

[shiva.rama16@gmail.com](mailto:shiva.rama16@gmail.com),

<sup>4</sup>Banothu Usha Assistant Professor,

[banothuusha@gmail.com](mailto:banothuusha@gmail.com),

Department of CSE Engineering,

Nagole , Institute of Engineering and Technology collage in Hyderabad.

## ABSTRACT

*A novel picture super-resolution (SR) technique based on a Convolution Neural Network (CNN) is being developed as part of this project's research. When learning the feature extraction, upsampling, and high-resolution (HR) reconstruction modules at the same time, a deep convolutional neural network (CNN) is created that can be used to rebuild pictures from any source and is completely trainable. If, on the other hand, you want to train a deep network in a straight line from start to end, this is time-consuming and may provide sub-optimal results since it takes a longer time to converge than other strategies. According to our results, an ensemble of deep and shallow networks should be trained at the same time in order to overcome this difficulty. Its stronger representation power, rather than a lower learning capacity, allows the deep network to capture the high-frequency information contained within visual images, rather than the other way around. When utilised in combination with joint training, the shallow network reduces the complexity of deep network optimization by a factor of two, in part because the shallow network is considerably simpler to optimise than the deep network. High frequency characteristics are rebuilt in a multi-scale manner to further improve the accuracy of HR reconstruction. This allows for the simultaneous integration of both short- and long-range contextual information to be included in the reconstruction, which further improves the accuracy of HR reconstruction. The suggested technique has been carefully examined on a variety of commonly used data sets, and when compared to current best practises, it beats them by a significant margin. Large-scale ablation experiments are carried out to establish the contributions of various network topologies to image SR, which results in the finding of new insights that may be used to future study.*

## 1. Introduction

A low resolution (LR) observation is used to attempt to recover a high resolution (HR) picture with a large number of high-frequency characteristics from a low resolution (LR) observation. Single image super-resolution (SR) attempts to recover a high resolution (HR) picture with a large number of high-frequency characteristics from a low resolution (LR). However, SR is fundamentally ill-posed since there is a lack of appropriate information about the situation, which is particularly true when considering that numerous HR images may be down-sampled into a single lower-resolution image. According to the most recent study, learning-based strategies have been gaining more and more attention, and they have shown to be more effective in image SR than their predecessors. It is the fundamental premise of learning the mapping function from the LR picture to its HR counterpart via the examination of auxiliary data obtained throughout the method that is being discussed. In order to estimate the residual between the HR picture and the bicubic-interpolated LR image, machine



learning algorithms based on the commonly used notion of image SR utilising CNNs are applied. According to the assumptions, the basic structure of the target HR image will be structurally identical to the fundamental structure of the bicubic up sampled LR version. In contrast to the custom-crafted bi cubic interpolation, which was expressly created for this purpose, the custom-crafted bi cubic interpolation may have a negative impact on the final performance. In contrast to the previously disclosed CNN-based tactics that make use of bicubic interpolation, our approach makes use of CNNs to learn a direct mapping from LR to HR pictures, which is both faster and more accurate than the previously stated techniques, as shown in Figure 1. On the basis of our early study, we have learned that it is difficult to train a complicated deep network in an end-to-end manner, and that the final results are often poor in a wide range of conditions. According to our results, an ensemble of deep and shallow networks should be trained at the same time in order to overcome this difficulty. To build deep networks, it is necessary to follow a systematic procedure. There are three basic ways, with the shallow network being the most lightweight (it only has three convolution layers, for example) and simplest to adjust of the two.

It is important to do feature extraction on the original LR picture before mapping it into a deep feature space in order to map it into a deep feature space in LR. Learning filters are used to achieve up sampling of deep features to the appropriate spatial size, and the HR picture is rebuilt by taking into account the multi-scale contextual information included within the up sampled deep features. A shallow network trained in combination with other networks has the potential to converge fast and correctly capture the essential structure of an HR picture, which is mostly made of low-frequency information, in a very short amount of time. As a result, the deep network is only responsible for retrieving high-frequency features that are dependent on the basic picture structure, resulting in a significant decrease in the complexity of the deep network training approach, which is advantageous. While the suggested network ensemble is not nearly as complex as the earlier CNN, which was developed using bicubic

interpolation-based techniques, it is equivalent in that the deep network is designed to learn the high frequency residual information, which is similar to the prior CNN. Because our technique substitutes a shallow network for the bicubic interpolation, it is completely trainable from the beginning to the conclusion, which separates it from the other available solutions. The process of duplicating a single pixel, according to some experts, may be impacted by either short- or long-range contextual information during the generation process. When applied to SR with high up scaling factors, some CNN-based algorithms that employ tiny picture patches to anticipate the centre pixel value perform less well, although they are still useful when applied to SR.

## 2. LITERATURE SURVEY

This work describes a learning-based approach for predicting scenes from photos that may be used to a variety of low-level vision challenges in general, as well as specific vision difficulties. We can construct a completely synthetic universe of events with their associated projected pictures, which we can then play back in real time, by characterising these interactions using a Markov network. It is possible to determine a local maximum of the posterior probability for a scene based on an image by use Bayesian belief propagation. In this particular case, this strategy is quite effective. It is known as VISTA (Vision by Image/Scene Training), which stands for Vision by Image/Scene Training Approach.

With respect to the "super-resolution" job (estimating high frequency information from a low-resolution picture), it is shown that VISTA works well, providing favourable results. As a final demonstration of the method's potential breadth and adaptability, we apply it to two more problem domains, both of which are reduced duplicates of the original problem. Learning to discern between shadow and reflectance differences in a single picture captured under certain lighting circumstances is an important skill to have in one's toolbox of



photographic abilities. A probabilistic approach is used to demonstrate figure/ground discrimination, solution of the aperture issue, and filling-in, all of which are similar to the probabilistic approach used to demonstrate the motion estimate problem in a "blobs world."

Generally speaking, super-resolution algorithms may be classified into two categories: Superresolution techniques include (i) classic multi-image superresolution (which combines photographs acquired at various sub-pixel misalignments), and (ii) Example-Based superresolution (which combines images collected at various pixel misalignments) (learning correspondence between low and high resolution image patches from a database). Our inquiry will benefit from merging these two families of methodologies in conjunction with one another since we will be able to present a cohesive framework. It is shown in further detail below how this combination strategy may be used to attain super resolution from as little as a single shot as feasible using the techniques described above (with no database or prior examples). For patch recognition in natural photographs, we developed an approach based on the fact that patches in a genuine photograph tend to redundantly repeat several times throughout the image, both within the same scale and across other sizes, as well as across different scales. Traditionally, super-resolution is obtained by the repetition of patches within the same picture scale (at subpixel misalignments), but example-based super-resolution is produced through the repetition of patches across several image scales (at subpixel misalignments) (at subpixel misalignments). We want to recover at each pixel the best possible resolution increase based on the patch redundancy within and across scales, while also attempting to recover at each pixel the best possible resolution increase based on the patch redundancy within and across scales.

### 3. SYSTEM ANALYSIS

Image SR methods may be divided into three types, according to their approach: interpolation-based reconstruction-based approaches, learning-based

approaches, and hybrid approaches. A learning-based approach to image SR, with its main concept being that image SR is a nonlinear mapping from low-resolution (LR) to high-resolution (HR) pictures, and that the mapping is learned using auxiliary data in a controlled environment, has recently emerged as one of the most active research areas in the field, becoming one of the most active study areas in recent years. According to Freeman et al., this strategy employs Markov Random Field (MRF) and patchbased external examples to achieve effective magnification by employing Markov Random Field (MRF) and patchbased external examples to generate effective magnification. Several techniques that were inspired by it were developed and put into practise as a result of its publication. A sparse representation algorithm is used to ensure that HR patches have a sparse linear representation over an overcomplete dictionary of patches randomly selected from comparable pictures. One representative approach is based on the sparse representation algorithm, whereas another method is also based on the sparse representation algorithm when using the algorithm. In this work, Yan et al. trains both the LR and HR dictionaries at the same time, with the limitation that both the LR patches and their corresponding HR counterparts have the same sparse representation as the LR patches, as described above. When it comes to training the coarse vocabulary, Orthogonal Matching Pursuit (K-SVD) is employed, and when it comes to training the fine dictionary, Orthogonal Matching Pursuit (K-SVD) is utilised. This work makes use of the Orthogonal Matching Pursuit (OMP), which was invented by and is being used to solve the decomposition issue. The neighbour embedding approach is used to produce super-resolved LR pictures, which are then processed further. According to this strategy, low-dimensional nonlinear manifolds with locally identical shape are used to locate the LR and HR patches. A large number of ideas are offered in order to increase the overall efficiency of computing even more. Yang and Yang develop a simple mapping function for each subspace after partitioning the LR feature space into several



subspaces. This mapping function will be useful in the future.

Several linear regressors are utilised to anchor the neighbours on a local level in order to achieve this goal. The use of precalculated anchors and regressors, which are calculated in advance, allows A+ [11] to improve SR performance in terms of accuracy and speed by using precalculated anchors and regressors. In order to construct another line of image SR approaches, the regression trees or forests approach is used. This method is referred to as the regression trees or forests approach. Through the use of leaf nodes as building blocks, this technique extends the capabilities of linear multivariate regression models that have previously been used. It then linearizes the mapping from LR to HR patches in the vicinity of centroids, using leaf nodes as building blocks to do this. Recently, the use of deep learning-based algorithms to image SR has shown impressive results. Images with strong signal recovery (image SR) are produced by using a CNN with three convolution layers, which is composed of three convolution layers and three convolution layers, respectively. Deep networks may one day be used to reformulate the classic sparse coding-based technique, which has showed some promise in the past but has yet to be fully realised. Using the Gibbs distribution as the conditional model, and the proper statistics predicted by a CNN, Reference is a firm that specialises in the restoration of human-related pictures. Kim et al. introduce a deep network with 20 convolutional layers, which they describe as an extension of the residual prediction algorithms that have been used in previous research. As a result of training the deep network to understand the difference between HR and LR photos, its performance has significantly increased. The authors provide a strongly recursive neural network to aid in the reconstruction of the HR pictures, which they believe will be of additional assistance. Through the usage of this idea, feature maps are retrieved from the LR space and learning is utilised to raise resolution just at the very beginning and end of the network, proving that the upscaling filters that have been learnt may be used to improve the accuracy of prediction even more. After that, there are a variety of different

CNN-based algorithms that are employed in image SR, including densely connected networks, recursive networks, and cascade upsampling networks, among other approaches, among others. In contrast to previous studies, the system we present is completely trainable from the beginning to the conclusion, and it is composed mostly of a combination of deep and shallow neural network components. An additional module is offered, which is a multi-scale high-resolution image restoration module, which is intended to gather information on both short and long-range contextual linkages via the use of photographs. Earlier examinations of these approaches have not been carried out in the same way as they have been carried out in previous investigations.

## DISADVANTAGES

Less accuracy score

Low performance

Unable to predict the resolution

## 3.1 PROPOSED SYSTEMS

Methods for image SR, such as end-to-end deep and shallow networks, often referred to as EEDS, will be discussed in further depth later in this section. Because it offers an overview of the architecture of the network ensemble, which is comprised of a deep convolutional neural network and a shallow convolutional neural network, this part is very important to the success of the project. Even more complicated, the deep CNN may be divided into three modules, each of which operates in parallel, and each of which performs a different task, such as feature extraction, up sampling, and multi-scale reconstruction.

## A. THERE ARE SPECIFICATIONS EXTRACTED FROM THEREIN. B.

The recovery of local aspects of high-frequency information in conventional shallow approaches is accomplished by computing the first and second



order gradients of an image patch, which is comparable to filtering the input picture using high-pass filters that are manually generated. Higher-level techniques extract local features by computing the first and second order gradients of the picture patch, which is analogous to filtering the input image with high-pass filters that have been meticulously developed and built by hand, as described in the paper. Higher-level techniques extract local features by computing the first and second order gradients of the picture patch. The deep learning-based solutions, rather than manually developing these filters, automatically learn these filters from training data, resulting in considerable time and effort reductions on both ends of the spectrum. Some studies, on the other hand, extract features from coarse HR pictures, which are obtained by up sampling the LR images to the HR size and then using bicubic interpolation to achieve the HR size in order to acquire the HR size in order to obtain the HR size. It is our belief that the bicubic interpolation was not particularly intended for this purpose, and that it may even be harmful in some cases.

LR information that can be essential in recouping the expenditures of human resources. As a consequence, in contrast to the previously described technique, the suggested methodology adopts an alternative strategy and executes feature extraction directly on the original LR photos utilising convolution layers, rather than through the convolution layers themselves. A nonlinear mapping function is accomplished by Rectified Linear Units (ReLUs), which are interleaved across three convolution layers in our feature extraction module. This is the structure of our feature extraction module. When connecting the input feature map of the second layer with the output feature map of the third layer, which is expressed as a "residual unit," the use of a shortcut connection based on identity mapping is required. This is accomplished by employing a shortcut connection based on identity mapping to connect the input feature map of the second layer with its output feature map. As previously indicated, the use of a residual unit may effectively assist gradients flow through several layers, thereby speeding up deep network training. Our reconstruction module makes use of structures

that are pretty similar to those depicted in the prior two cases of related structures. Each of the three convolution layers, each with a kernel size of 33 percent, creates feature maps with 64 channels, which are then merged to form a single feature map with all three convolution layers integrated. It is vital to retain the spatial size of the output feature maps; consequently, zero padding is utilised to accomplish this purpose.

## B. THE APPLICATION OF UPSPAMPLING IS ESSENTIAL.

An upsampling method is done on the features that have been retrieved from the original LR images in order to increase their spatial span to the desired HR size after they have been recovered from the original LR photos. The learning-based upsampling method we utilise instead of hand-designed interpolation techniques results in a trainable system from the beginning to the conclusion of the process, which saves time and money. As a consequence, we will analyse two alternative strategies widely applied in CNN for up sampling, namely un pooling and deconvolutions, which are both extensively used in CNN. Un pooling and deconvolutions are both extensively utilised in CNN. Consider the un pooling procedure with an up scaling factor  $e$ . When compared to a conventional pooling operation, the un pooling process with an up scaling factor  $e$  replaces each item in an input feature map with an  $e$  block, where the top left element is set to the value of the input entry and the rest components are set to zero, as illustrated in the image. The unpooling technique yields output feature maps that are both bigger and more sparse than the input feature maps, suggesting that it is more efficient. The values of output values that have been sparsely activated may be transferred to surrounding regions as a result of the convolution layers utilised in the technique. In deconvolution layers with forward and backward propagation of  $s$ , the forward and backward propagation of convolution layers with forward and backward propagation of  $s$  is inverted. This leads in an exponential rise in the size of the input feature maps when employing an output stride of  $s$ , as illustrated in Figure 1. Pooling and deconvolution have diverse



implementations, but they are basically comparable when it comes to up scaling feature maps, and both are well suited to the work at hand, as shown in the following example. We are able to acquire some extremely promising findings as a consequence of the deconvolution layer that has been added.

Option C is Reconstruction on a Multi-Scale Environment.

Due to the fact that similar image patterns may recur across multiple scales in different images from both the training and test sets, accurate inference of the input image should be highly invariant to image scale variations and may rely on the aggregation of multi-scale contextual information with respect to image scale variations. In recent years, several vision-related difficulties, such as image item identification [39], scene recognition, and other analogous tasks [40, 41], have been carefully researched and proved to be successful. For image SR, past research has shown that incorporating multi-scale context may greatly boost HR picture reconstruction in a variety of conditions, including those involving high-resolution pictures. Because it is probable that HR picture restoration will be dependent on both short- and long-range contextual information, we recommend that HR reconstruction be achieved via multi-scale convolutions to explicitly retain multi-context information throughout the reconstruction process. After passing through the R residual units, the input to our HR reconstruction module is eventually delivered to our HR reconstruction module itself. On top of that, a second layer of dimension reduction is done to produce the desired outcome. This layer is made up of a 1 1 convolution that maps the input feature map of 64 channels to the output feature map of 16 channels, resulting in an output feature map with a total of 64 channels as a consequence of the convolution. Following that, there is a multi-scale convolution layer, which consists of four convolution operations with changing kernel sizes: one convolution operation with three kernels, five convolution operations with five kernels, and seven times seven convolution operations with seven kernels. After that, there is a decomposition layer, which consists of four decomposition procedures

with variable kernel sizes. All four convolutions are performed on the input feature map at the same time, resulting in four feature maps with a total of 16 channels in each of the four feature maps.

## ADVANTAGES

Good accuracy score

Good performance

Predict the higher resolution

## 4. IMPLEMENTATION

### ARCHITECTURE ANALYSIS:

For a better understanding of our contributions, we will undertake more testing on many different permutations of the EEDS approach that we have presented in this article, as described in this study. In general, while training all of the methods, we rigorously adhere to the implementation parameters described in Section IV-A, unless differently mentioned in the approach description.

Because it is implemented as an ensemble, our approach is capable of training both a deep and a shallow network at the same time. Using the two networks as a starting point, the proposed EEDS model is divided into two versions, namely EED (end-to-end deep network) and EES (end-to-end shallow network). These two versions of the proposed EEDS model are then used to analyse the impact of the two networks on the overall performance of the system under consideration. All three models' convergence graphs on the Set5 data set are displayed in Fig. 3, with the time scale represented by the x-axis. A shallow network expedites the process of convergent EES, allowing it to be finished in less time than it would otherwise take. In spite of the fact that the EES system has a substantial amount of available capacity, the system's



overall performance is sub-par. Education in EED, on the other hand, may prove to be difficult to master. Throughout the training process, there are frequent swings in training loss, indicating that the mechanism is very instable. However, despite the fact that EED has a higher PSNR than EES after convergence, the result is still unsatisfactory. What is causing this to occur may be connected to the fact that directly mapping LR shots to HR images is a very difficult operation, and EED may eventually settle on a local minimum, but the exact reason for this is unknown. The suggested EEDS technique, which integrates deep and shallow networks into a single ensemble network structure, alleviates this challenge by combining them into one structure. Despite the fact that the shallow network converges far more rapidly than the deep network, it is the shallow network that dominates the performance from the very beginning of the training session. When the shallow network has captured the majority of the HR pictures, direct SR becomes much less difficult to do, resulting in a reduction in the complexity of the direct SR procedure. Therefore, the deep network concentrates on high-frequency input and learns to rectify the flaws created by the shallow network, resulting in the greatest overall performance among the three systems tested. As soon as the shallow network of EDS reaches convergence, the prediction made by the shallow network restores the majority of the content that was previously blurred or artifactually altered. In contrast, when the shallow network of EEDS reaches convergence, the deep network of EEDS learns to predict the residual between the HR image and the output of the shallow network, which is predominantly composed of high-frequency content. Deep and shallow networks are combined using simple addition, and the behaviour of deep and shallow networks is supported and confirmed by the key findings of deep residual networks, which indicate that deep residual learning can be achieved through the addition of subnetworks and that deep networks are easier to optimise. Deep residual networks are used to learn about the behaviour of deep and shallow networks. The inclusion of deep and shallow networks is not only compatible with previous SR approaches, but it is also compatible with previous SR approaches that include learning

the residual between a high-resolution HR picture and a bicubic interpolated LR input, rather than learning the residual between two images. The residual prediction-based technique, for example, is an example of our methodology since it uses a shallow network instead of the fixed bicubic interpolation and trains both deep and shallow networks concurrently. Using a baseline deep CNN (designated as DCNN) with an architecture comparable to SRCNN against a combination of a baseline deep CNN and an SRCNN-like 3-layer shallow CNN (designated as SRCNN) (designated as DSCNN), it is discovered that the benefits of mixing deep and shallow networks can be applied to a wide range of network topologies. As a consequence, the DSCNN regularly outperforms the basic deep CNN over a wide range of different data sets.

## 5. PROBLEM STATEMENT

It is our experience that when we transfer photographs, the resolution of the images is reduced, and as a result, the clarity of the image is reduced as a result. When converting a low quality picture to a high resolution image, we use CNN to enhance the clarity of the image.

## 6. Results

It will be shown in the next part how a sequence of output screens is created, as well as how the actual process of applying CNN takes place.

Figure 1 on the output screen contains information on the images that were used in the process of making it, which is shown in the second figure.



```
Code Text
prepare_images('source/', 2)
(276, 276, 3)
Saving face.bmp
(512, 512, 3)
Saving man.bmp
(512, 512, 3)
Saving pepper.bmp
(512, 512, 3)
Saving babo_0T.bmp
(361, 256, 3)
Saving comic.bmp
(256, 256, 3)
Saving butterfly_0T.bmp
(512, 512, 3)
Saving bridge.bmp
(391, 506, 3)
Saving zebra.bmp
(344, 228, 3)
Saving woman_0T.bmp
(480, 500, 3)
Saving baboon.bmp
(362, 506, 3)
Saving flowers.bmp
(388, 352, 3)
Saving foramen.bmp
(288, 288, 3)
Saving road_0T.bmp
(512, 512, 3)
Saving lemma.bmp
(288, 352, 3)
Saving coastguard.bmp
(576, 720, 3)
Saving barbara.bmp
(656, 529, 3)
Saving ppt3.bmp
(512, 768, 3)
Saving monarch.bmp
(388, 288, 3)
Saving bird_0T.bmp
```



All the images are converted in this format and put in a folder called output

## 7. CONCLUSION

A fully trainable single picture SR system that is totally end-to-end scalable will be constructed in this research with the help of an ensemble of deep and shallow networks as the building blocks, which will be used as the building blocks. Figure 1 illustrates how a shallow network learns to display the primary structure of an HR image due to its lightweight design and ease of optimization, whereas a deep network, which has a higher learning power, is solely responsible for capturing the high frequency features of an HR image due to its higher learning power. Due to this, grouping together to train the network ensemble might potentially greatly reduce the amount of effort necessary for network training while also

providing significantly enhanced performance. A multi-scale method to HR reconstruction is utilised for more accurate restoration of HR pictures since it allows for the incorporation of both short- and long-range contextual information into the same reconstruction. This method provides for more accurate restoration of HR pictures than the previous method. Using experimental data, it has been discovered that the suggested strategy outperforms current state of the art techniques in terms of overall performance and efficiency. The results of this study include comprehensive ablation tests to corroborate the contributions of different network architectures to image SR, as well as further insights into future research.

## REFERENCES

- [1] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *Int. J. Comput. Vis.*, vol. 40, no. 1, pp. 25–47, 2000.
- [2] 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.
- [3] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *Proc. IEEE Int. Conf. Comput. Vis.*, Sep. 2009, pp. 349–356.
- [4] 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.
- [5] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," in *Proc. Int. Conf. Curves Surf. Berlin, Germany: Springer*, Jun. 2010, pp. 711–730.
- [6] G. Freedman and R. Fattal, "Image and video upscaling from local selfexamples," *ACM Trans. Graph.*, vol. 30, no. 2, p. 12, 2011.
- [7] J. Yang, Z. Wang, Z. Lin, S. Cohen, and T. Huang, "Coupled dictionary training for image super-resolution," *IEEE Trans. Image Process.*, vol. 21, no. 8, pp. 3467–3478, Aug. 2012.
- [8] H. Chang, D.-Y. Yeung, and Y. Xiong, "Super-resolution through neighbor embedding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 1, Jun./Jul. 2004, pp. 275–282.
- [9] C.-Y. Yang and M.-H. Yang, "Fast direct super-resolution by simple functions," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 561–568.
- [10] R. Timofte, V. Smet, and L. Van Gool, "Anchored neighborhood regression for fast example-based super-resolution," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 1920–1927