WHAT IS CLAIMED IS:

1. A computer-implemented method for training a neural network to downsample images in a video encoding pipeline, the method comprising:

executing a first convolutional neural network on a first source image having a first resolution to generate a first downsampled image, wherein the first convolutional neural network includes at least two residual blocks and is associated with a first downsampling factor;

executing an upsampling algorithm on the first downsampled image to generate a first reconstructed image having the first resolution;

computing a first reconstruction error based on the first reconstructed image and the first source image; and

updating at least one parameter of the first convolutional neural network based on the first reconstruction error to generate a trained convolutional neural network;

wherein a residual block comprises a portion of the first convolutional neural network that maps the input of the residual block to a residual and then adds the residual to a function of the input of the residual block to generate the output of the residual block; and

wherein each downsampled image has a resolution that is lower than a resolution of a corresponding source image by the first downsampling factor

wherein if the first downsampling factor comprises a non-integer, the method further comprises:

determining a resampling factor numerator and a resampling factor denominator based on the first downsampling factor;

generating an upsampling residual block based on the resampling factor numerator;

generating a downsampling residual block based on the resampling factor denominator; and

appending the downsampling residual block to the upsampling residual block to generate the first convolutional neural network.

2. The computer-implemented method of claim 1, wherein the upsampling residual block includes a deconvolution layer having an output stride equal to the resampling factor numerator and a second upsampling algorithm that implements an

upsampling factor equal to the resampling factor numerator.

3. The computer-implemented method of claim 1, wherein if the first downsampling factor comprises an integer, the method further comprises:

generating a downsampling residual block based on the first downsampling factor; and

appending the downsampling residual block to an identity residual block to generate the first convolutional neural network.

4. The computer-implemented method of claim 1, wherein the upsampling algorithm is differentiable.

5. The computer-implemented method of claim 1, further comprising configuring the upsampling algorithm to implement an upsampling factor equal to the first downsampling factor.

6. The computer-implemented method of claim 1, wherein computing the first reconstruction error comprises computing a mean squared error of the first reconstructed image relative to the first source image.

7. The computer-implemented method of claim 1, wherein updating the at least one parameter of the first convolution neural network comprises:

computing an iteration error based on the first reconstruction error and at least a second reconstruction error, wherein the second reconstruction error is associated with a second source image having a second resolution that is not equal to the first resolution; and

performing at least one of a backpropagation operation or a gradient descent operation on the first convolutional neural network based on the iteration error to update the at least one parameter.

8. The computer-implemented method of claim 1, wherein the first source image comprises at least a portion of a frame included in a training video.

9. One or more computer readable media including instructions that, when

executed by one or more processors, cause the one or more processors to train a neural network to downsample images by performing the method steps of any preceding claim.

10. A computer-implemented method for downsampling images, the method comprising:

executing a first trained convolutional neural network on a first source image having a first resolution to generate a first downsampled image having a second resolution that is lower than the first resolution;

wherein the first trained convolutional neural network includes at least two residual blocks and is associated with a first downsampling factor;

wherein a residual block comprises a portion of the first convolutional neural network that maps the input of the residual block to a residual and then adds the residual to a function of the input of the residual block to generate the output of the residual block; and

wherein each downsampled image has a resolution that is lower than a resolution of a corresponding source image by the first downsampling factor;

wherein the at least two residual blocks include an upsampling residual block that is associated with a numerator of a resampling fraction and a downsampling residual block that is associated with a denominator of the resampling fraction.

11. The computer-implemented method of claim 10, wherein the upsampling residual block includes a deconvolution layer having an output stride equal to the numerator of the resampling fraction.

12. The computer-implemented method of claim 10, wherein the upsampling residual block includes an upsampling algorithm that implements an upsampling factor equal to the numerator of the resampling fraction.

13. The computer-implemented method of claim 10, wherein the first downsampling factor comprises an integer, and wherein the at least two residual blocks include an identity residual block and a downsampling residual block.

14. The computer-implemented method of claim 10, wherein the first source image comprises a frame of a source video, and the first downsampled image comprises a frame of a downsampled video.

15. The computer-implemented method of claim 10, further comprising performing one or more encoding operations on the first downsampled image to generate an encoded image.