What is claimed is :

1. A method comprising:

receiving, at a user device, a media stream including frames of a first resolution generated by a graphics-rendering application;

on the user device, utilizing one or more weight matrices pre-trained in association with the graphics-rendering application to locally upscale each received frame of the media stream to a second resolution greater than the first resolution; and

rendering the upscaled frames of the media stream to a display of the user device.

2. The method of claim 1, wherein each of the frames is upscaled to the second resolution using a select weight matrix of the one or more weight matrices trained in association with the graphics-rendering application, the upscaling of each of the frames providing for supplementation of the frame with additional detail derived based on the associated select weight matrix.

3. The method of claim 1, further comprising:

transmitting, from the user device, a request for a select weight matrix, the request specifying a matrix identifier stored in a remote database in association with a media title for the graphics-rendering application;

in response to transmission of the request, receiving and locally storing the select weight matrix on the user device;

dynamically loading a super-resolution neural net (SRNN)inference engine with the select weight matrix when the rendering operation reaches a predefined checkpoint within the media stream.

4. The method of claim 3, wherein multiple weight matrices are pre-trained in association with the graphics-rendering application and wherein the method further comprises:

loading the SRNN inference engine with a different one of the weight matrices upon reaching each of multiple predefined checkpoints during the presentation of the media stream.

5. The method of claim 1, wherein multiple weight matrices are pre-trained in association with the graphics-rendering application and each of the weight matrices are trained with frames sampled from an associated one of multiple different discrete frame segments of the media stream.

6. The method of claim 5, wherein each of the weight matrices is used to dynamically upscale frames of the associated one of the multiple different discrete frame

segments during media streaming operations.

7. The method of claim 1, wherein the user device is a gaming console and the one or more weight matrices are stored on the console for dynamic loading into a super-resolution neural net (SRNN) inference engine at select times in association with different checkpoints within the media stream.

8. The method of claim 1, wherein the user device is a mobile device and the method further comprises:

receiving and storing metadata for the graphics-rendering application, the metadata specifying matrix identification information uniquely identifying multiple different weight matrices pretrained for the graphics-rendering application and further specifying checkpoint identifiers each identifying a checkpoint location within the media stream at which to load an associated one of the multiple different weight matrices into a super-resolution neural net (SRNN) inference engine.

9. A system comprising:

a super-resolution neural net (SRNN) inference engine stored in memory of a user device and configured to:

receive a media stream including frames of a first resolution generated by a graphics-rendering application and utilize one or more weight matrices pre-trained in association with the graphics-rendering application to locally upscale each of multiple frames received within the media stream from the first resolution a second resolution greater than the first resolution; and

a media viewing application locally executing on the device that presents the upscaled frames of the media stream to a display of the user device.

10. The system of claim 9, wherein upscaling of each frame of the multiple frames received within the media stream includes supplementing each of the frame with additional detail derived based on an associated one of the one or more weight matrices.

11. The system of claim 9, wherein the device further comprises:

an application programming interface (API) configured to:

transmit a request for a select weight matrix during the presentation of the media stream, the request specifying a matrix identifier stored in a remote database in association with a media title for the graphics-rendering application; and

in response to transmission of the request, receive and locally store the requested weight matrix on the user device, the user device being adapted to dynamically load the SRNN inference engine with the select weight matrix responsive to reaching a predefined checkpoint during the presentation of the media stream.

12. The system of claim 9, wherein multiple weight matrices are pre-trained in association with the graphics-rendering application and wherein SRNN inference engine is dynamically loaded with a different one of the weight matrices upon reaching each of multiple predefined checkpoints during the presentation of the media stream.

13. The system of claim 9, wherein multiple weight matrices are pre-trained in association with the graphics-rendering application and each of the weight matrices are trained with frames sampled from an associated one of multiple different discrete frame segments of the media stream.

14. The system of claim 9, wherein each of the weight matrices is used to dynamically upscale frames of the associated one of the multiple different discrete frame segments during media streaming operations.

15. The system of claim 14, wherein the user device is a gaming console and the one or more weight matrices are jointly stored on the gaming console for dynamic loading into the SRNN inference engine in association with different checkpoints within the media stream.

16. One or more computer-readable storage media encoding computer-executable instructions for executing a computer process comprising :

receiving, at a user device, a media stream including frames of a first resolution generated by a graphics-rendering application;

on the user device, utilizing one or more weight matrices pre-trained in association with the graphics-rendering application to locally upscale each received frame of the media stream to a second resolution greater than the first resolution; and

rendering the upscaled frames of the media stream to a display of the user device.

17. The one or more computer-readable storage media of claim 16, wherein each of the frames is upscaled to the second resolution using a select weight matrix of the one or more weight matrices trained in association with the graphics-rendering application, the upscaling of each of the frames supplementing each of the frames with additional detail derived based on the associated select weight matrix.

18. The one or more computer-readable storage media of claim 17, wherein the computer process further comprises :

transmitting, from the user device, a request for a select weight matrix, the request specifying a matrix identifier stored in a remote database in association with a media title for the graphics-rendering application;

in response to transmission of the request, receiving and locally storing the select weight matrix on the user device; and

dynamically loading a super-resolution neural net (SRNN) inference engine with the select weight matrix when the rendering operation reaches a predefined checkpoint within the media stream.

19. The one or more computer-readable storage media of claim 17, wherein multiple weight matrices are pre-trained in association with the graphics-rendering application and wherein the computer process further comprises :

loading the SRNN inference engine with a different one of the weight matrices upon reaching each of multiple predefined checkpoints during presentation of the media stream.

20. The method of claim 1, wherein pretraining of the pretrained one or more weight matrices includes :

inputting an image at the first resolution into an SRNN,

the SRNN configured to use the pretrained one or more weight matrices to upscale images;

receiving, responsive to the inputting, an upscaled version of the image at the second resolution from the SRNN;

comparing the upscaled version of the first resolution image with a high-resolution version of the image;

modifying the pretrained one or more weight matrices based on the comparison.