IN THE CLAIMS:

1. A computer-implemented method for estimating perceived video quality for reconstructed video, the method comprising:

computing a first plurality of feature values corresponding to a plurality of visual quality metrics based on a first reconstructed frame, a first source frame, a first display resolution, and a first normalized viewing distance at which the reconstructed video is viewed by a viewer on a viewing device;

computing a second plurality of feature values corresponding to the plurality of visual quality metrics based on the first reconstructed frame, the first source frame, the first display resolution, and a second normalized viewing distance at which the reconstructed video is viewed by a viewer on the viewing device;

executing a trained perceptual quality model on the first plurality of feature values to generate a first perceptual quality score that indicates a first perceived visual quality level for the first reconstructed frame when viewed at the first normalized viewing distance;

executing the trained perceptual quality model on the second plurality of feature values to generate a second perceptual quality score that indicates a second perceived visual quality level for the first reconstructed frame when viewed at the second normalized viewing distance; and

performing one or more operations associated with an encoding process based on at least one of the first perceptual quality score or the second perceptual quality score.

2. The computer-implemented method of claim 1, wherein performing the one or more operations associated with the encoding process comprises:

computing an overall perceptual quality score for a first reconstructed video sequence that includes the first reconstructed frame based on the first perceptual quality score, wherein the overall perceptual quality score indicates an overall perceived visual quality level for the first reconstructed video sequence; and

computing a tradeoff between quality and bitrate based on the overall perceptual quality score when encoding a first source video sequence.

3. The computer-implemented method of claim 1, wherein performing the one or more operations associated with the encoding process comprises evaluating at least one of an encoder, a decoder, or an adaptive bitrate streaming algorithm based on the first perceptual quality score.

4. The computer-implemented method of claim 1, wherein the first perceived visual quality level approximates the quality of an actual viewing experience where the viewing device has the first display resolution and the viewer views the viewing device at the first normalized viewing distance.

5. The computer-implemented method of claim 1, further comprising: computing a first pixels-per-degree value based on the first display resolution and the first normalized viewing distance; and

computing an absolute perceptual quality score for a first reconstructed video sequence that includes the first reconstructed frame based on the first perceptual quality score and a first offset associated with the first pixels-per-degree value, wherein the absolute perceptual quality score indicates an overall perceived visual quality level for the first reconstructed video sequence.

6. The computer-implemented method of claim 1, wherein the plurality of visual quality metrics comprises at least one of a Detail Loss Metric that is modified to account for different normalized viewing distances and different display resolutions or a Visual Information Fidelity index that is modified to account for different normalized viewing distances.

7. The computer-implemented method of claim 1, wherein computing the first plurality of feature values comprises computing a first feature value by:

computing a pixels-per-degree value based on the first display resolution and the first normalized viewing distance;

computing a plurality of contrast sensitivities based on a plurality of scale indices, the pixels-per-degree value, and at least one contrast sensitivity function; and

computing the first feature value based on the first reconstructed frame, the first source frame, and the plurality of contrast sensitivities.

8. The computer-implemented method of claim 1, wherein computing the first plurality of feature values comprises computing a first feature value by:

computing a noise coefficient based on the first normalized viewing distance; and

computing the first feature value based on the first reconstructed frame, the first source frame, the first display resolution, the first normalized viewing distance, and the noise coefficient.

9. The computer-implemented method of claim 1, wherein computing the first plurality of feature values comprises computing a first feature value by:

generating a Gaussian two-dimensional kernel based on the first normalized viewing distance; and

computing the first feature value based on the Gaussian two-dimensional kernel, the first reconstructed frame, and the first source frame.

10. One or more non-transitory computer readable media including instructions that, when executed by one or more processors, cause the one or more processors to estimate perceived video quality for reconstructed video by performing the steps of:

computing a first plurality of feature values corresponding to a plurality of visual quality metrics based on a first reconstructed frame, a first source frame, a first display resolution, and a first normalized viewing distance at which the reconstructed video is viewed by a viewer on a viewing device;

computing a second plurality of feature values corresponding to the plurality of visual quality metrics based on the first reconstructed frame, the first source frame, the first display resolution, and a second normalized viewing distance at which the reconstructed video is viewed by a viewer on the viewing device;

executing a trained perceptual quality model on the first plurality of feature values to generate a first perceptual quality score that indicates a first perceived visual quality level for the first reconstructed frame when viewed at the first normalized viewing distance;

executing the trained perceptual quality model on the second plurality of feature values to generate a second perceptual quality score that indicates a second perceived visual quality level for the first reconstructed frame when viewed at the second normalized viewing distance; and

performing one or more operations associated with an encoding process based on at least one of the first perceptual quality score or the second perceptual quality score.

11. The one or more non-transitory computer readable media of claim 10, wherein performing the one or more operations associated with the encoding process comprises computing a tradeoff between quality and bitrate based on the first perceptual quality score when encoding a first source video sequence that includes the first source frame.

12. The one or more non-transitory computer readable media of claim 10, wherein performing the one or more operations associated with the encoding process comprises:

computing an overall perceptual quality score for a first reconstructed video sequence that includes the first reconstructed frame based on the first perceptual quality score, wherein the overall perceptual quality score indicates an overall perceived visual quality level for the first reconstructed video sequence; and

evaluating at least one of an encoder, a decoder, or an adaptive bitrate streaming algorithm based on the overall perceptual quality score.

13. The one or more non-transitory computer readable media of claim 10, wherein the trained perceptual quality model comprises a trained support vector regression model, a trained artificial neural network, or a trained regression tree.

14. The one or more non-transitory computer readable media of claim 10, further comprising:

computing a first pixels-per-degree value based on the first display resolution and the first normalized viewing distance; and

computing an absolute perceptual quality score for a first reconstructed video sequence that includes the first reconstructed frame based on the first perceptual quality score and a first offset associated with the first pixels-per-degree value, wherein the absolute perceptual quality score indicates an overall perceived visual quality level for the first reconstructed video sequence.

15. The one or more non-transitory computer readable media of claim 10, wherein the plurality of visual quality metrics comprises at least one of a Detail Loss Metric that is modified to account for different normalized viewing distances and

different display resolutions or a Visual Information Fidelity index that is modified to account for different normalized viewing distances.

16. The one or more non-transitory computer readable media of claim 10, wherein computing the first plurality of feature values comprises computing a first feature value by:

computing a pixels-per-degree value based on the first display resolution and the first normalized viewing distance;

computing a plurality of contrast sensitivities based on a plurality of scale indices, the pixels-per-degree value, and at least one contrast sensitivity function; and

computing the first feature value based on the first reconstructed frame, the first source frame, and the plurality of contrast sensitivities.

17. The one or more non-transitory computer readable media of claim 10, wherein computing the first plurality of feature values comprises computing a first feature value by:

generating a Gaussian two-dimensional kernel based on the first normalized viewing distance; and

computing the first feature value based on the Gaussian two-dimensional kernel, the first reconstructed frame, and the first source frame.

18. The one or more non-transitory computer readable media of claim 10, wherein the first reconstructed frame is in a High Dynamic Range format.

19. A system comprising:

one or more memories storing instructions; and

one or more processors coupled to the one or more memories that, when executing the instructions, perform the steps of:

computing a first plurality of feature values corresponding to a plurality of visual quality metrics based on a first reconstructed frame, a first

source frame, a first display resolution, and a first normalized viewing distance at which the reconstructed video is viewed by a viewer on a viewing device;

computing a second plurality of feature values corresponding to the plurality of visual quality metrics based on the first reconstructed frame, the first source frame, the first display resolution, and a second normalized viewing distance at which the reconstructed video is viewed by a viewer on the viewing device;

executing a trained perceptual quality model on the first plurality of feature values to generate a first perceptual quality score that indicates a first perceived visual quality level for the first reconstructed frame when viewed at the first normalized viewing distance;

executing the trained perceptual quality model on the second plurality of feature values to generate a second perceptual quality score that indicates a second perceived visual quality level for the first reconstructed frame when viewed at the second normalized viewing distance; and

performing one or more operations associated with an encoding process based on at least one of the first perceptual quality score or the second perceptual quality score.