

TECHNIQUES FOR PREDICTING PERCEPTUAL VIDEO QUALITY BASED ON COMPLEMENTARY PERCEPTUAL QUALITY MODELS

CROSS-REFERENCE TO RELATED APPLICATIONS

- 5 [0001] This application claims benefit of United States patent application serial number 15/890,710, filed February 7, 2018, which is hereby incorporated herein by reference.

BACKGROUND OF THE INVENTION

Field of the Invention

- 10 [0002] Embodiments of the present invention relate generally to video technology and, more specifically, to techniques for predicting perceptual video quality based on complementary perceptual quality models.

Description of the Related Art

- 15 [0003] Efficiently and accurately encoding source video content is critical for real-time delivery of high-quality video content. Because of variations in encoded video content quality, it is desirable to implement quality controls to ensure that the visual quality of re-constructed video content decoded from the encoded source video content is acceptable. Manually verifying the visual quality of re-constructed video content is prohibitively time consuming. Accordingly, automated video content quality
20 assessment is often an integral part of the encoding and streaming infrastructure. For example, automated video quality assessment can be employed when evaluating encoders and fine-tuning streaming bitrates to optimize streamed video content.

- [0004] Some approaches to automatically assessing the quality of re-constructed video content involve computing a quality score based on a perceptual quality model
25 that is trained based on human assessments of visual quality. In general, perceptual quality models associate values for features of re-constructed video content with different quality scores. One limitation of predicting the quality of re-constructed video content via a perceptual quality model is that perceptual quality models typically lack robustness. For example, if none of the features associated with a perceptual quality
30 model reflect certain types of distortions, then the perceptual quality model is unable to accurately predict human perception of the quality of re-constructed video content that include such distortions.

[0005] In particular, perceptual quality models are typically unable to reliably predict the perceived visual quality of re-constructed video content included in the public-domain LIVE Video Quality Database. Analysis of the deficiencies of perceptual quality models with respect to the LIVE Video Quality Database indicate that none of the features associated with the perceptual quality model accurately model the perceived quality degradation attributable to many types of complex temporal video distortions. Examples of complex temporal video distortions include, without limitation, motion estimation mismatches, flicker, and ghosting.

[0006] As the foregoing illustrates, what is needed in the art are more effective techniques for predicting the perceived quality of re-constructed video content.

SUMMARY OF THE INVENTION

[0007] One embodiment of the present invention sets forth a computer-implemented method for predicting the perceived video quality of re-constructed video content. The method includes computing a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, where the first model is generated via a first machine learning algorithm; computing a second quality score for the re-constructed video content based on a second set of values for a second set of features and a second model that associates the second set of values with the second quality score; and determining an overall quality score for the re-constructed video content based on the first quality score and the second quality score, where the overall quality score indicates a level of visual quality associated with streamed video content.

[0008] At least one advantage of the disclosed techniques for predicting perceived video quality is that the overall quality score accurately predicts perceived video quality across a diverse range of re-constructed video content. In particular, at least one technological improvement relative to the prior art is that the disclosed techniques aggregate quality scores obtained via different perceptual quality models associated with different sets of features to determine the overall quality score. Because features are distributed between the perceptual quality models, the disclosed techniques decrease the likelihood that increasing the total number of features results in inaccuracies attributable to overfitting.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] So that the manner in which the above recited features of the present invention can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to embodiments, some of which are
5 illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0010] Figure 1 is a conceptual illustration of a system configured to implement one or
10 more aspects of the present invention;

[0011] Figure 2 is a more detailed illustration of the temporal extractor of Figure 1, according to various embodiments of the present invention;

[0012] Figure 3 is a flow diagram of method steps for computing a quality score for re-constructed video content, according to various embodiments of the present
15 invention;

[0013] Figure 4 is a conceptual illustration of an ensemble system configured to implement one or more aspects of the present invention; and

[0014] Figure 5 is a flow diagram of method steps for computing an overall quality score for re-constructed video content based on two different sets of features,
20 according to various embodiments of the present invention.

DETAILED DESCRIPTION

[0015] In the following description, numerous specific details are set forth to provide a more thorough understanding of the present invention. However, it will be apparent to one of skilled in the art that the present invention may be practiced without one or
25 more of these specific details.

[0016] The disclosed techniques efficiently and reliably predict quality scores for re-constructed video content via a perceptual quality model that associates values for features with a quality score. Notably, the perceptual quality model is associated with multiple temporal features, where each temporal feature is associated with a different
30 scale. In operation, a temporal extractor computes per-frame values for a specific

temporal feature based on re-constructed video content, the associated video content, and the scale associated with the temporal feature. For a given frame, the temporal extractor computes a frame difference matrix specifying differences in luminance between the frame and an adjacent frame. The temporal extractor then
5 down-samples the frame difference matrix based on the scale to generate a scaled frame difference matrix. Subsequently, the temporal extractor performs temporal filtering operations and local mean-subtraction operations on the scaled frame difference matrix to compute a local mean-subtracted matrix. The temporal extractor then performs conditioning operations on the scaled frame difference matrix to
10 generate a first weighted entropy matrix that is associated with the frame of the re-constructed video content.

[0017] In an analogous fashion, the temporal extractor derives a second weighted entropy matrix based on the video content and the scale. The temporal extractor performs subtraction operations between the first weighted entropy matrix and the
15 second weighted entropy matrix to compute an entropy difference matrix. Finally, the temporal extractor performs spatial pooling operations on the entropy difference matrix to generate a single value for the temporal feature of the frame of the re-constructed video content. Because computing entropy difference matrices at multiple scales emulates aspects of how the human visual system processes visual
20 information, the perceptual quality model accurately models perceived quality degradation attributable to many types of complex temporal video distortions.

[0018] In some embodiments, an ensemble prediction engine computes a ensemble quality score for re-constructed video content based on two perceptual quality models. Each of the perceptual quality models associates values for a different set of
25 features with a different quality score. An ensemble prediction engine composites quality scores obtained via the two perceptual quality models to obtain an ensemble quality score. Advantageously, distributing features between two prediction quality models decreases the likelihood that increasing the total number of features results in inaccuracies attributable to overfitting.

30 **System Overview**

[0019] Figure 1 is a conceptual illustration of a system 100 configured to implement one or more aspects of the present invention. As shown, the system 100 includes,

without limitation, any number of compute instances 110, a training database 120, an objective database 150, a subjective database 182, and a target feature database 170. For explanatory purposes, multiple instances of like objects are denoted with reference numbers identifying the object and parenthetical numbers identifying the instance where needed. In various embodiments, any number of the components of the system 100 may be distributed across multiple geographic locations or included in one or more cloud computing environments (i.e., encapsulated shared resources, software, data, etc.) in any combination.

[0020] As shown, each of the compute instances 110 includes, without limitation, a processor 112 and a memory 116. The processor 112 may be any instruction execution system, apparatus, or device capable of executing instructions. For example, the processor 112 could comprise a central processing unit (CPU), a graphics processing unit (GPU), a controller, a microcontroller, a state machine, or any combination thereof. The memory 116 stores content, such as software applications and data, for use by the processor 112 of the compute instance 110.

[0021] The memory 116 may be one or more of a readily available memory, such as random access memory (RAM), read only memory (ROM), floppy disk, hard disk, or any other form of digital storage, local or remote. In some embodiments, a storage (not shown) may supplement or replace the memory 116. The storage may include any number and type of external memories that are accessible to the processor 112. For example, and without limitation, the storage may include a Secure Digital Card, an external Flash memory, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing.

[0022] In general, the compute instances 110 are configured to implement one or more applications. For explanatory purposes only, each application is depicted as residing in the memory 116 of a single compute instance 110 and executing on a processor 112 of the single compute instance 110. However, as persons skilled in the art will recognize, the functionality of each application may be distributed across any number of other applications that reside in the memories 116 of any number of compute instances 110 and execute on the processors 112 of any number of

compute instances 110 in any combination. Further, the functionality of any number of applications may be consolidated into a single application or subsystem.

5 [0023] In particular, the compute instances 110 are configured to provide automated video content quality assessment for re-constructed video content. The video content quality assessment may be used to implement quality controls associated with an encoding and streaming infrastructure. For example, to ensure an acceptable viewing experience, the quality of re-constructed video content and, indirectly, the caliber of the elements included in the encoding and streaming infrastructure are typically evaluated at various points in the design and delivery process based on visual quality
10 scores. The visual quality scores are then used to guide the development of applications (e.g., encoders) and the real-time optimization of content delivery, such as stream-switching algorithms that are quality-aware.

[0024] As shown, a training feature engine 130, a training engine 180, a prediction engine 192, an instance of a feature extractor 140, and a temporal pooling engine 194
15 reside in the memory 116 and execute on the processor 112 of different compute instances 110. In alternate embodiments, any number of instances of the training feature engine 130, the training engine 180, the feature extractor 140, the prediction engine 192, the temporal pooling engine 194, the training feature engine 130, and the training engine 180 may reside in any number and types of memory and execute on
20 any number and types of processing devices in any combination.

[0025] During a learning phase, depicted with dotted lines, the training feature engine 130 and the training engine 180 collaborate to generate a perceptual quality model 190. For re-constructed video content, the perceptual quality model 190 associates a set of values for a set of features with a quality score 184 that correlates to subjective
25 human visual experience when viewing the re-constructed video content.

[0026] In operation, the training feature engine 130 acquires a training database 120 and generates the objective database 150. The training database 120 includes, without limitation, any number of sources 122, and any number of re-constructed sources 124. Each of the sources 122 includes, without limitation, any amount and
30 type of video content. Examples of video content include, without limitation, any portion (including all) of feature length films, episodes of television programs, and music videos, to name a few. Each of the re-constructed sources 124 include,

without limitation, re-constructed video content that is derived from the video content included in one of the sources 122.

5 [0027] As a general matter, each of the re-constructed sources 124 replicates the video content included in the associated source 122 as conveyed to a viewer via a streaming infrastructure and a display device. The re-constructed sources 124 may be derived from the sources 122 based on any number of encoding parameters and in any technically feasible fashion. For instance, in some embodiments, three different re-constructed sources 124 are derived from each of the sources 122. For each of three different sets of encoding parameters, an encoding engine (not shown in Figure 1) independently encodes the source 122(x). For each of the three resulting encoded sources, a decoding engine (not shown in Figure 1) independently decodes the encoded source to generate one of the re-constructed sources 124 derived from the sources 122(x).

15 [0028] As shown, the training feature engine 130 includes, without limitation, the feature extractor 140. The feature extractor 140 computes values for features. Each of the features is either a “spatial feature” that measures spatial video distortions or a “temporal feature” that measures temporal visual distortions. Examples of spatial video distortions include, without limitation, compression artifacts, scaling artifacts, and film grain noise. Examples of temporal video distortions include, without limitation, motion estimation mismatches, flicker, and ghosting.

20 [0029] The feature extractor 140 includes, without limitation, a spatial extractor 148 and a temporal extractor 146. For each re-constructed source 124(x) and each spatial feature, the spatial extractor 148 computes a spatial array 172(x) based on the re-constructed source 124(x) and, optionally, the source 122 associated with the re-constructed source 124(x). The spatial array 172(x) includes, without limitation, a different frame value for each frame included in the re-constructed source 124(x).

30 [0030] Similarly, for each re-constructed source 124(x) and each temporal feature, the temporal extractor 146 computes a temporal array 174(x) based on the re-constructed source 124(x) and the source 122 associated with the re-constructed source 124(x). The temporal array 174(x) includes, without limitation, a different frame pair value for each frame included in the re-constructed source 124(x). As referred to herein, a “frame pair” comprises two sequential frames, and the frame pair

value is associated with the first of the two sequential frames. In alternate embodiments, the feature extractor 140 may compute any number of values for a particular feature of a particular re-constructed source 124 in any technically feasible fashion.

- 5 **[0031]** For each spatial feature and each re-constructed source 124(x), the training feature engine 130 performs temporal pooling operations on the spatial array 172(x) associated with the spatial feature to determine a spatial value 162 for the feature of the re-constructed source 124(x). The training feature engine 130 may perform any amount and type of temporal pooling operations in any technically feasible fashion.
- 10 For instance, in some embodiments, the training feature engine 130 could set the spatial value 162 for a given spatial feature of the re-constructed source 124(x) equal to the average of the frame values included in the spatial array 172(x) associated with the feature. Similarly, for each temporal feature and each re-constructed source 124(x), the training feature engine 130 could perform set the temporal value 164 for a
- 15 given temporal feature of the re-constructed source 124(x) equal to the average of the frame pair values included in the temporal array 174(x) associated with the feature.

- [0032]** For each of the re-constructed sources 124, the training feature engine 130 generates a training feature database 152 that includes, without limitation, the spatial values 162 and the temporal values 164 for the features of the re-constructed source
- 20 124. The training feature engine 130 then generates the objective database 150 that includes, without limitation, the training feature databases 152. Consequently, if the training database 120 includes N re-constructed sources 124, then the objective database 150 includes N training feature databases 152. Further, if the feature extractor 140 computes S spatial arrays 172 for S spatial features and T temporal
- 25 arrays 174 for T temporal features, then each of the training feature databases 152 includes S spatial values 162 and T temporal values 164.

- [0033]** The training engine 180 generates the perceptual quality model 190 based on the objective database 150 and the subjective database 182. The subjective database 182 includes, without limitation, a quality score 184 of each of the re-
- 30 constructed sources 142. The quality score 184(x) for the re-constructed source 142(x) is generated based on human viewings of the re-constructed source 142(x). More specifically, as viewers watch the re-constructed source 142(x) on display

devices, the viewers personally rate the visual quality – assigning personalized scores. The personalized scores are then aggregated to generate the quality score 184(x).

5 **[0034]** The training engine 180 implements any number and type of machine learning algorithms in any technically feasible fashion to generate the perceptual quality model 190 based on the objective database 150 and the subjective database 182. Examples of machine learning algorithms include, without limitation, support vector machine algorithms, artificial neural network algorithms, and random forest algorithms, to name a few. In general, for each of the re-constructed sources 124(x),
10 the perceptual quality model 190 is trained based on the association between the spatial values 162 and the temporal values 164 included in the objective database 150(x) and the quality score 184(x) included in the subjective database 182. In alternate embodiments, the training engine 180 may train the perceptual quality model 190 at any technically feasible level of granularity (e.g., based on associating
15 frame values for the features with frame quality scores).

[0035] Notably, irrespective of the granularity at which the training engine 180 trains the perceptual quality model 190, the perceptual quality model 190 associates a set of values for the set of features at any given granularity to a score at the same granularity. More precisely, the perceptual quality model 190 associates a set of S
20 values for the set of S spatial features and a set of T values for the set of T temporal features with a score that predicts an associated perceived visual quality. For example, the perceptual quality model 190 associates a set of S frame values for the spatial features of the first frame and a set of T frame pair values for the temporal features of the first frame to a frame score that represents the perceived visual quality
25 of the first frame.

[0036] During a prediction phase, depicted with solid lines, the feature extractor 140, the prediction engine 192, and the temporal pooling engine 194 collaborate to predict the quality score 184 of a re-constructed target 144 based on the perceptual quality model 190, the re-constructed target 144, and a target 142. The feature extractor
30 140, the prediction engine 192, and the temporal pooling engine 194 are collectively referred to herein as a “prediction application.” The re-constructed target 144 includes, without limitation, re-constructed video content that is derived from the video

content included in the target 142. The target 142 includes, without limitation, any amount and type of video content.

[0037] The re-constructed target 144 replicates the video content included in the associated source 122 as conveyed to a viewer via a streaming infrastructure and a display device. The re-constructed target 144 may be generated in any technically feasible fashion. For instance, in some embodiments, an encoder encodes the target 142 based on a constant rate encoding strategy to generate an encoded target. Subsequently, a decoder decodes the encoded target to generate the re-constructed target 144.

[0038] For each spatial feature associated with the perceptual quality model 190, the spatial extractor 148 computes the spatial array 172 based on the re-constructed target 144 and, optionally, the target 142. For each temporal feature associated with the perceptual quality model 190, the temporal extractor 146 computes the temporal array 174 based on the re-constructed target 144 and the target 142. In this fashion, for each frame included in the re-constructed target 144, the feature extractor 140 computes a different frame value for each spatial feature and a different frame pair value for each temporal feature.

[0039] For each frame included in the re-constructed target 144, the prediction engine 192 generates a frame score based the frame values associated with the frame, the frame pair values associated with the frame, and the perceptual quality model 190. For example, for the first frame included in the re-constructed target 144, the prediction engine 192 determines a first set of values for the features based on the first entries of the spatial arrays 172 and the first entries of the temporal arrays 174. The prediction engine 192 then determines the frame score for the first frame based on the first set of values for the features and the perceptual quality model 190. In alternate embodiments, the prediction engine 192 may compute scores for perceived visual quality at any level of granularity based on values for the features at the same level of granularity and the perceptual quality model 190.

[0040] Subsequently, the temporal pooling engine 194 determines the quality score 184 of the re-constructed target 144 based on the arithmetic mean of the frame scores. More specifically, the temporal pooling engine 184 sums all of the frame scores and then divides the summation by the total number of frames to compute the

quality score 184. Accordingly, the quality score 184 represents an average visual quality across the frames included in the re-constructed target 144. Finally, the temporal pooling engine 194 transmits or displays the quality score 184 of the re-constructed target 144 for evaluation of streamed video content.

- 5 **[0041]** In alternate embodiments, the temporal pooling engine 194 may determine the quality score 184 of the re-constructed target 144 based on the frame scores in any technically feasible fashion. For instance, in some embodiments, the temporal pooling engine 194 may perform hysteresis pooling operations that mimic the relatively smooth variance of human opinion scores in response to changes in video
10 quality. More precisely, the temporal pooling engine 194 could perform both linear low pass operations and non linear (rank-order) weighting operations on the frame scores to generate the quality score 184 of the re-constructed target 144.

- [0042]** In some alternate embodiments, the temporal pooling engine 194 determines a different quality score 184 of each chunk of the re-constructed target 144 instead of a
15 single quality score 184 for the re-constructed target 144. As referred to here, a “chunk” typically corresponds to a scene, and a scene is a sequence of frames where the spatiotemporal characteristics are approximately constant. To compute the quality score 184 of a particular chunk of the re-constructed target 144, the temporal pooling engine 194 computes the arithmetic mean of the frame scores for the subset
20 of the frames that are included in the chunk.

- [0043]** One limitation of predicting the quality of re-constructed video content via a conventional perceptual quality model is that conventional perceptual quality models typically lack robustness. For example, if none of the features associated with a conventional perceptual quality model reflect certain types of distortions, then the
25 conventional perceptual quality model is unable to accurately predict human perception of the quality of re-constructed video content that include such distortions.

- [0044]** A conventional perceptual quality model is typically associated with multiple spatial features and a single temporal feature. The spatial features may include without limitation, any number of additive impairment measures (ADM) features, blind
30 or referenceless image spatial quality evaluator features, and visual information fidelity (VIF) features, to name a few. The temporal feature is a frame difference feature. For a given frame pair, a conventional temporal extractor computes per-pixel

differences in luminance between the two frames. The temporal extractor then sets the frame pair value for the frame difference feature equal to the average of the per-pixel differences in luminance.

5 [0045] However, as demonstrated by benchmarking conducted using the public-domain LIVE Video Quality Database, the frame difference feature does not accurately measure the perceived quality degradation attributable to many types of complex temporal video distortions. Examples of complex temporal video distortions include, without limitation, motion estimation mismatches, flicker, and ghosting. Because the frame difference feature is an unreliable measure of complex temporal
10 video distortions, conventional perceptual quality models are unable to robustly predict the perceived visual quality of re-constructed video content.

Implementing Accurate Temporal Features

[0046] To address the above problems, the system 100 establishes multiple temporal features that emulate aspects of how the human visual system processes visual
15 information. For instance, as described in conjunction with Figure 2, because the human visual system processes visual information at multiple scales, the temporal extractor 146 computes frame pair values for multiple temporal features, where each temporal feature is associated with a different scale. In alternate embodiments, the temporal extractor 146 may compute frame pair values for any number of temporal
20 features in any technically feasible fashion that ensures that multiple scales are represented. In general, the temporal extractor 146 computes frame pair values for temporal features that are perceptually relevant and, together, accurately measure perceived quality degradation attributable to complex temporal video distortions.

[0047] It will be appreciated that the system 100 shown herein is illustrative and that
25 variations and modifications are possible. For example the functionality provided by the training feature engine 130, the training engine 180, the feature extractor 140, the spatial extractor 148, the temporal extractor 146, the prediction engine 192, and the temporal pooling engine 194 as described herein may be integrated into or distributed across any number of software applications (including one) and any number of
30 components of the system 100. Further, the connection topology between the various units in Figure 1 may be modified as desired.

[0048] Figure 2 is a more detailed illustration of the temporal extractor 146 of Figure 1, according to various embodiments of the present invention. As shown, the temporal extractor 146 includes, without limitation, a frame comparison engine 220, a down-sampling engine 240, a spatial filtering engine 250, an entropy engine 270, an entropy comparison engine 280, and a spatial pooling engine 290.

[0049] In general, the temporal extractor 146 generates values for the temporal features of re-constructed video content based on the re-constructed video content and the video content from which the re-constructed video content is derived. Notably, each of the temporal features is associated with a different scale 250. The re-constructed video content may be any of the re-constructed sources 124 or the re-constructed target 144. For explanatory purposes only, the temporal extractor 146 is described in Figure 2 in the context of generating the temporal array 174(1) for a first temporal feature of the re-constructed target 144..

[0050] As shown, the re-constructed target 144 includes, without limitation, frames 210(1)-210(F), where F is any integer greater than 1. In a complementary fashion, the target 142 includes, without limitation, the frames 210'(1)-210'(F). For explanatory purposes only, an object that is associated with the target 142 is distinguished from a comparable object that is associated the re-constructed target 144 with a prime symbol (i.e., ' ') that decorates the reference number. As described in conjunction with Figure 1, each of the re-constructed target 144 and the target 142 are associated with F frames 210 and F frame pairs.

[0051] In operation, the frame comparison engine 220(1) computes frame difference matrices 230(1)-230(F) based on the re-constructed target 144. Each frame difference matrix 230 is associated with a different frame pair. In general, to compute the frame difference matrix 230(x), the frame comparison engine 220(1) performs per-pixel comparison operations between the frames 210(x) and 210(x+1) to determine per-pixel luminance differences. Accordingly, each pixel location included in the frame difference map 230(x) specifies a different per-pixel luminance difference. In an analogous manner, the frame comparison engine 220(2) independently computes the frame difference matrices 230'(1)-230'(F) based on the target 142.

[0052] The down-sampling engine 240(1) independently performs down-sampling operations on each of the frame difference matrices 230(1)-230(F) based on the scale

250(1) associated with the first temporal feature to generate, respectively, scaled frame difference (FD) matrices 242(1)-242(F). As shown, the scale 250(1) is equal to 2. In alternate embodiments, the first temporal feature may be associated with any scale 250 in any technically feasible fashion. The down-sampling engine 240(1) may
5 perform any number and type of down-sampling operations based on the scale 250(1). In general, if each of the frame difference matrices 230 includes N by M pixels and the scale 240(1) is Q, then each of the scaled frame difference matrices 242 includes N/Q by M/Q pixels. In an analogous manner, the down-sampling engine 240(2) independently performs down-sampling operations on each of the
10 frame difference matrices 230'(1)-230'(F) based on the scale 240(1) to generate, respectively, the scaled frame difference matrices 242'(1)-242'(F).

[0053] The spatial filtering engine 250(1) independently processes each of the scaled frame difference matrices 242(1)-242(F) to generate, respectively, local mean-subtracted matrices 262(1)-262(F). To process the scaled frame difference matrix
15 242(x), the spatial filtering engine 250(1) performs a two-step localized spatial filtering process that approximates bandpass filtering of the scaled frame difference matrix 242(x). First, the spatial filtering engine 250(1) filters the scaled frame difference matrix 242(x) with a spatial isotropic Gaussian filter to generate a locally averaged scaled frame difference matrix (not shown). The spatial filtering engine 250(1) then
20 subtracts the locally averaged scaled frame difference matrix from the scaled frame difference matrix 242(x) to generate the local mean-subtracted matrix 262(x). In an analogous fashion, the spatial filtering engine 250(2) independently processes each of the scaled frame difference matrices 242'(1)-242'(F) to generate, respectively, the local mean-subtracted matrices 262'(1)-262'(F).

25 **[0054]** The entropy engine 270(1) independently performs conditioning operations on each of the local mean-subtracted matrices 262(1)-262(F) to generate, respectively, weighted entropy matrices 272(1)-272(F). Notably, the conditioning operations are related to divisive normalization process that is known to occur in the primary visual cortex (area V1). To process the local mean-subtracted matrix 262(x), the spatial
30 filtering engine 250(1) assigns the samples included in the local mean-subtracted matrix 262(x) to overlapping blocks, where each block includes a bxb set of neighboring samples. The spatial filtering engine 250(1) then estimates a bxb covariance matrix based on the overlapping blocks. The spatial filtering engine

250(1) also partitions the samples included in the local mean-subtracted matrix 282(x) into non-overlapping bxb blocks. For each non-overlapping block, the spatial filtering engine 250(1) computes the variance of the non-overlapping block.

[0055] Subsequently, for each non-overlapping block, the spatial filtering engine 250(1) computes an entropy value for the non-overlapping block based on the covariance matrix. Finally, for each non-overlapping block, the spatial filtering engine 250(1) weighs the entropy value of the non-overlapping block by the variance of the non-overlapping block via a logarithmic function to compute a weighted entropy value. Accordingly, the weighted entropy matrix 272(x) includes, without limitation, a weighted entropy value for each non-overlapping block. In an analogous fashion, the spatial filtering engine 250(2) independently processes each of the local mean-subtracted matrices 262'(1)-262'(F) to generate, respectively, the weighted entropy matrices 272'(1)-272'(F).

[0056] The entropy difference engine 280 computes entropy difference matrices 282(1)-282(F) based on the weighted entropy matrices 272(1)-272(F) associated with the re-constructed target 144 and the weighted entropy matrices 272'(1)-272'(F) associated with the target 142. More specifically, to compute the entropy difference matrix 282(x), the entropy difference engine 280 computes the difference between each of the weighted entropy values included in the weighted entropy matrix 272(x) and the corresponding weighted entropy value included in the weighted entropy matrix 272'(x). Accordingly, the entropy difference matrix 282(x) includes, without limitation, an "entropy difference value" for each non-overlapping block.

[0057] The spatial pooling engine 290 performs spatial pooling operations to generate the temporal array 174(1) based on the entropy difference matrices 282. As shown, the temporal array 174(1) includes, without limitation, frame pair values 264(1)-264(x). Each of the frame pair values 264 included in the temporal array 174(1) specifies a value for the first temporal feature of a different one of the frame pairs included in the re-constructed target 144. To compute the frame pair value 264(x), the spatial pooling engine 290 averages the entropy difference values over all the non-overlapping blocks included in the entropy difference matrix 282(x).

[0058] In alternate embodiments, the spatial pooling engine 290 may aggregate the entropy difference values included in the entropy difference matrix 282(x) in any

technically feasible fashion to generate the frame pair value 264(x). Advantageously, entropy differences between re-constructed video content and the video content from which the re-constructed video content is derived correlate with subjective human assessments of the quality of the re-constructed video content.

5 **[0059]** Notably, a second temporal feature is associated with the scale 250(2) of 3 and a third temporal feature is associated with the scale 240(3) of 4. When computing the temporal array 174(2), the down-sampling engines 240(1) and 240(2) perform down-sampling operations based on the scale 240(2) of 3. When computing the temporal array 174(3), the down-sampling engines 240(1) and 240(2) perform down-sampling
10 operations based on the scale 240(3) of 4. The operations of the frame comparison engine 220, the spatial filtering engine 250, the entropy engine 270, the entropy comparison engine 280, and the spatial pooling engine 290 are not altered based on the scale 250. Advantageously, computing temporal features at multiple scales reflects multiscale processing inherent in the human vision system.

15 **[0060]** Note that the techniques described herein are illustrative rather than restrictive, and may be altered without departing from the broader spirit and scope of the invention. Many modifications and variations on the functionality provided by the temporal extractor 146, the frame comparison engine 220, the down-sampling engine 240, the spatial filtering engine 250, the entropy engine 270, the entropy comparison
20 engine 280, and the spatial pooling engine 290 will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments.

[0061] For instance, in various embodiments, the temporal extractor 146 may be configured to compute values for any number of temporal feature based on any of the
25 matrices at any number of scales 250. For example, the temporal extractor 146 could compute frame pair values for a first frame difference feature based on the frame difference matrices 230 and 230'. The temporal extractor 146 could compute frame pair values for a second frame difference feature based on the frame difference matrices 242 and 242' that are associated with the scale 250 of 2. The temporal
30 extractor could compute frame pair values for a first entropy feature based on the entropy difference matrices 282 that are associated with the scale 250 of 2. Finally, the temporal extractor could compute frame pair values for a second entropy feature

based on the entropy difference matrices 282 that are associated with the scale 250 of 3.

5 **[0062]** Figure 3 is a flow diagram of method steps for computing a quality score for re-constructed video content, according to various embodiments of the present invention. Although the method steps are described with reference to the systems of Figures 1-2, persons skilled in the art will understand that any system configured to implement the method steps, in any order, falls within the scope of the present invention.

10 **[0063]** As shown, a method 300 begins at step 302, where the frame comparison engine 220 generates the frame difference matrices 230 based on the re-constructed target 144 and the frame difference matrices 230' based on the target 142. At step 304, the temporal extractor 146 selects the first temporal feature and the scale 250(1) associated with the first temporal feature. At step 306, the down-sampling engine 240 down-samples the frame difference matrices 230 and the frame difference matrices 15 230' based on the selected scale 250 to generate, respectively, the scaled frame difference matrices 242 and the scaled frame difference matrices 242'.

20 **[0064]** At step 308, for each of the scaled frame difference matrices 242, the spatial filtering engine 250 independently performs localized spatial filtering operations on the frame difference matrix 242 to generate the corresponding local mean-subtracted matrix 262. In an analogous fashion, for each of the scaled frame difference matrices 242', the spatial filtering engine 250 independently performs localized spatial filtering operations on the scaled frame difference matrix 242' to generate the corresponding local mean-subtracted matrix 262'.

25 **[0065]** At step 310, for each of the local mean-subtracted matrices 262, the entropy engine 270 performs conditioning operation on the local mean-subtracted matrix 262 to compute the weighted entropy matrix 272. In an analogous fashion, for each of the local mean-subtracted matrices 262', the entropy engine 270 performs conditioning operation on the local mean-subtracted matrix 262' to compute the weighted entropy matrix 272'.

30 **[0066]** At step 312, the entropy difference engine 280 compares the weighted entropy matrices 272 associated with the re-constructed target 144 to the weighted entropy

matrices 272' associated with the target 142 to compute the entropy difference matrices 282. More specifically, to compute the entropy difference matrix 282(x), the entropy difference engine 280 computes the difference between each of the weighted entropy values included in the weighted entropy matrix 272(x) and the corresponding weighted entropy value included in the weighted entropy matrix 272'(x). At step 314, for each of entropy difference matrices 282, the spatial pooling engine 290 performs spatial pooling operations to compute the frame pair value 264(x) included in the temporal array 174 associated with the selected temporal feature.

[0067] At step 316, the temporal extractor 146 determines whether the selected temporal feature is the last temporal feature. If, at step 316, the temporal extractor 146 determines that the selected temporal feature is not the last temporal feature, then the method 300 proceeds to step 318. At step 318, the temporal extractor 146 selects the next feature and the associated scale 250, and the method 300 returns to step 306, where the down-sampling engine 240 down-samples the frame difference matrices 230 and the frame difference matrices 230' based on the selected scale 250.

[0068] If, however, at step 316, the temporal extractor 146 determines that the selected temporal feature is the last temporal feature, then the method 300 proceeds directly to step 320. At step 320, for each spatial feature, the spatial extractor 146 computes frame values based on the re-constructed target 144 and, optionally, the target 142. At step 322, for each of the frames 210, the prediction engine 192 generate a frame score based on the frame values associated with the frame, the frame pair values associated with the frame, and the perceptual quality model 190.

[0069] At step 324, the temporal pooling engine 194 performs temporal pooling operations on the frame scores to determine the quality score 184 of the re-constructed target 144. The temporal pooling engine 194 may perform any number and type of temporal pooling operations to aggregate the frame scores. For instance, in some embodiments, the temporal pooling engine 194 determines the quality score 184 of the re-constructed target 144 based on the arithmetic mean of the frame scores. The temporal pooling engine 194 then transmits or displays the quality score 184 of the re-constructed target 144 for evaluation of streamed video content. The method 300 then terminates.

[0070] Advantageously, any number of the quality scores 184 may be aggregated to generate an aggregate quality score as part of evaluating the visual quality of streamed media content. For example, a dashboard application that monitors a streaming service on a daily basis can indicate the health of the streaming service based on the aggregated quality score. In general, unless a change occurs to the streaming service, the aggregated quality score remains stable. Consequently, a deviation in the aggregated quality score between a particular day and the immediately preceding day indicates a change to the streaming service (e.g., an issue or an improvement) that impacts the visual quality of the streamed media content. By contrast, if a change to the streaming service is not intended to affect the visual quality of the streamed media content, then the aggregated quality score can provide confirmation that the visual quality of the streamed media content is unaffected. For instance, if a new encode for a smart phone saves bandwidth while preserving visual quality, then the aggregated quality score remains stable while a time-weighted bitrate decreases.

[0071] In another example, an A/B test application compares a control group of streaming sessions with one or more treatment groups of streaming sessions. A treatment group differs from the control group in exactly one key parameter. The aggregated quality score indicates whether the visual quality of the streamed media content is impacted by the difference in the key parameter. For instance, if the treatment group uses a new encode and the control group uses a current encode, then the A/B test application could compare the aggregated quality score for the control group and the aggregated quality score for the treatment group. The results of the comparison indicate whether the new encode impacts the visual quality of streamed media content for the treatment groups.

Predicting Visual Quality Based on Complementary Perceptual Quality Models

[0072] As persons skilled in the art will recognize, as the number of features associated with the perceptual quality model 190 increases, the likelihood of overfitting the perceptual quality model 190 also increases. When the perceptual quality model 190 is overfit, the perceptual quality model 190 models idiosyncrasies in the training data (e.g., the objective database 150 and the subjective database 184). Consequently, an overfit perceptual quality model 190 is often unable to accurately predict the quality score 184 of re-constructed video content that is not included in the

training database 120. To reduce the likelihood of overfitting a single perceptual quality model 190, some embodiments of the present invention implement multiple, complementary perceptual quality models 190.

5 [0073] Figure 4 is a conceptual illustration of an ensemble system 400 configured to implement one or more aspects of the present invention. As shown, the ensemble system 400 includes, without limitation, any number of the compute instances 110. As shown, an ensemble training engine 410 resides in the memory 112 and executes on the processor of one of the computer instances 110. An ensemble prediction engine 490 resides in the memory 112 and execute on the processor 116 of another
10 computer instance 110. In alternate embodiments, the ensemble training engine 410 and the ensemble prediction engine 490 may reside in any number and type of memories and execute on any number and type of processing devices in any combination. The ensemble prediction engine 490 is also referred to herein as an “ensemble prediction application.”

15 [0074] The ensemble training engine 410 includes without limitation, a feature assigner 420 and two instances of the training engine 180. The feature assigner 420 receives an overall feature set 412 that includes any number of features and generates two features sets 422(1) and 422(2). More precisely, the feature assigner 420 assigns each of the features included in the overall feature set 412 to at least one
20 of the features sets 422.

[0075] The feature assigner 420 may be configured to assign each of the features included in the overall feature set 412 to at least one of the features sets 422 in any technically feasible fashion. For example, in some embodiments, the feature assigner 420 generates the features sets 422 based on user input received via a graphical
25 user interface (not shown). In other embodiments, the feature assigner 420 generates the feature sets 422 based on any number of heuristics that optimize any number and type of criteria.

[0076] For explanatory purposes only, the overall feature set 412 depicted in Figure 4 includes five spatial features and four temporal features. The five spatial features
30 include one VIF feature and four ADM features (associated with four different scales 250). The four temporal features include one frame difference feature and three entropy features (associated with three different scales 250). The feature assigner

420 assigns the five spatial features to both the feature set 422(1) and the feature set 422(2), the frame difference feature to the feature set 422(1), and the three entropy features to the feature set 422(2).

5 [0077] Subsequently, the ensemble training engine 410 configures the training engine 180(1) to generate the perceptual quality model 190(1) based on the subjective database 182 and the portion of the objective database 150 that specifies values for the features included in the feature set 422(1). By contrast, the ensemble training engine 410 configures the training engine 180(2) to generate the perceptual quality model 190(2) based on the subjective database 182 and the portion of the objective
10 database 150 that specifies values for the features included in the feature set 422(2).

[0078] In alternate embodiments, the ensemble training engine 410 includes two different training engines that implement different machine learning algorithms. The ensemble training engine 410 configures the first training engine to generate the perceptual quality model 190(1) based on the subjective database 182 and the
15 portion of the objective database 150 that specifies values for the features included in the feature set 422(1). The ensemble training engine 410 configures the second training engine to generate the perceptual quality model 190(2) based on the subjective database 182 and the portion of the objective database 150 that specifies values for the features included in the feature set 422(2).

20 [0079] In general, the perceptual quality model 190(1) associates a set of values for the feature set 422(1) with a first quality score, and the perceptual quality model 190(2) associates a set of values for the feature set 422(2) with a second quality score. Advantageously, because neither of the perceptual quality models 190 is associated with all of the features included in the overall feature set 422, the
25 likelihood that the ensemble training engine 410 generates overfit perceptual quality models 190 is mitigated.

[0080] The ensemble prediction engine 490 includes, without limitation, two instances of the prediction engine 192(1) and 192(2), an aggregation engine 440, and the temporal pooling engine 194. The ensemble prediction engine 490 receives the
30 perceptual quality model 190(1), the perceptual quality model 190(2), and the target feature database 170 associated with the re-constructed target 144. In alternate

embodiments, the ensemble prediction engine 490 configures the feature extractor 140 to compute the target feature database 170.

5 [0081] The ensemble prediction engine 490 configures the prediction engine 192(1) to compute a frame prediction database 430(1) based on the perceptual quality model 190(1) and the portion of the target feature database 170 associated with the features included in the feature set 422(1). For each frame 210(x) included in the re-constructed target 144, the frame prediction database 430(1) includes, without limitation, a frame score 434(x) that represents a predicted perceived visual quality of the frame 210. In operation, for the frame 210(x), the ensemble prediction engine 490
10 determines a set of values for the feature set 422(1) based on the target feature database 170. The ensemble prediction engine 490 then configures the prediction engine 192(1) to compute the frame score 434(x) included in the frame prediction database 430(1) based on the set of values for the feature set 422(1) and the perceptual quality model 190(1).

15 [0082] In a complementary fashion, the ensemble prediction engine 490 configures the prediction engine 192(2) to compute the frame prediction database 430(2) based on the perceptual quality model 190(2) and the portion of the target feature database 170 associated with the features included in the feature set 422(2). In operation, for the frame 210(x), the ensemble prediction engine 490 determines a set of values for
20 the feature set 422(2) based on the target feature database 170. The ensemble prediction engine 490 then configures the prediction engine 192(2) to compute the frame score 434(x) included in the frame prediction database 430(2) based on the set of values for the feature set 422(2) and the perceptual quality model 190(2).

[0083] Upon receiving the frame prediction database 430(1) and the frame prediction
25 database 430(2), the aggregation engine 440 generates an aggregated prediction database 450 that includes a different aggregated frame score 454 for each of the frames 210. More precisely, the aggregation engine 440 performs any number of aggregation operations on the frame score 434(x) included in the frame prediction database 430(1) and the frame score 434(x) included in the frame prediction
30 database 430(2) to generate the aggregated frame score 454(x) for the frame 210(x).

[0084] The aggregation engine 440 may compute the aggregated frame scores 454 in any technically feasible fashion. For instance, in some embodiments, the aggregation

engine 440 computes the average of the frame score 434(x) included in the frame prediction database 430(1) and the frame score 434(x) included in the frame prediction database 430(2) to generate the aggregated frame score 454(x) for the frame 210(x). In other embodiments, the aggregation engine 440 computes a
5 weighted average of the frame score 434(x) included in the frame prediction database 430(1) and the frame score 434(x) included in the frame prediction database 430(2) to generate the aggregated frame score 454(x) for the frame 210(x).

[0085] In yet other embodiments, the aggregation engine 440 may perform any number of machine learning algorithms to generate an aggregation model that
10 associates a frame score 434 associated with the feature set 422(1) and a corresponding frame score 434 associated with the feature set 422(2) with the aggregated frame quality score 454. Examples of machine learning algorithms that the aggregation engine 440 may perform include, without limitation, bayes-based methods, bagging, boosting, bayesian parameter averaging, bayesian model
15 combination, bucket of models, stacking, and the like. The aggregation model may or may not implement the same machine learning algorithms that the training engine 180 implements to generate the perceptual quality models 190. Subsequently, the aggregation engine 440 may compute the aggregated prediction database 450 that includes the aggregated frame scores 454 based on the aggregation model, the
20 frame prediction database 430(1), and the frame prediction database 430(2).

[0086] After the aggregation engine 440 generates the aggregated prediction database 450, the temporal pooling engine 194 performs any number of temporal pooling operations on the aggregated prediction database 450 to compute the quality score 184 of the re-constructed target 144. As described in detail in conjunction with
25 Figure 1, some examples of temporal pooling operations include arithmetic mean operations and hysteresis pooling operations. Further, as also described in detail in conjunction with Figure 1, in alternate embodiments, the temporal pooling engine 194 may generate different quality scores 184 for each chunk of the re-constructed target 144.

30 **[0087]** Notably, the reliability of the quality score 184 correlates to the reliability of the perceptual quality model 190(1), the reliability of the perceptual quality model 190(2), and the degree to which the perceptual quality models 190 complement each other.

Accordingly, the feature assigner 420 may be configured to generate the features sets 422 based on any number of heuristics that optimize the accuracy and/or robustness of the quality score 184.

5 **[0088]** For instance, in some embodiments, the feature assigner 420 implements heuristics that optimize the features included in each of the feature sets 422 to capture spatiotemporal aspects of perceptual visual quality, result in an accurate and robust perceptual quality model 190, and are diverse within the feature set 422 as well as across the feature sets 422.

10 **[0089]** Note that the techniques described herein are illustrative rather than restrictive, and may be altered without departing from the broader spirit and scope of the invention. Many modifications and variations on the functionality provided by the ensemble training engine 410, the feature assigner 420, the training engine 180, the ensemble prediction engine 490, the prediction engine 192, the aggregation engine 440, and the temporal pooling engine 194 will be apparent to those of ordinary skill in
15 the art without departing from the scope and spirit of the described embodiments.

20 **[0090]** Figure 5 is a flow diagram of method steps for computing an overall quality score for re-constructed video content based on two different sets of features, according to various embodiments of the present invention. Although the method steps are described with reference to the systems of Figures 1, 2, and 4, persons skilled in the art will understand that any system configured to implement the method steps, in any order, falls within the scope of the present invention.

25 **[0091]** As shown, a method 500 begins at step 502, where for each feature included in the overall feature set 412, the feature assigner 420 assigns the feature to at least one of the feature set 422(1) and the feature set 422(2). Notably, at least one feature included in each of the feature sets 422 is not included in the other feature set 422. At step 504, the ensemble training engine 410 configures the training engine 180(1) to generate the perceptual quality model 190(1) based on the subjective database 182 and the portion of the objective database 150 associated with the feature set 422(1). At step 506, the ensemble training engine 410 configures the training engine
30 180(2) to generate the perceptual quality model 190(2) based on the subjective database 182 and the portion of the objective database 150 associated with the feature set 422(2).

[0092] At step 508, the ensemble prediction engine 490 receives the target 142 and the re-constructed target 144. At step 510, for each frame 210 included in the re-constructed target 144, the ensemble prediction engine 490 computes a first set of values for the feature set 422(1) and a second set of values for the feature set 422(2).
5 More precisely, the ensemble prediction engine 490 configures the feature extractor 140 to compute the target feature database 170 based on the re-constructed target 144 and the target 142. For each of the frames 210, the ensemble prediction engine 490 then identifies the first set of values for the feature set 422(1) and the second set of values for the feature set 422(2) based on the target feature database 170. In
10 alternate embodiments, for each of the frames 210, the ensemble prediction engine 490 may compute the first set of values for the feature set 422(1) and the second set of values for the feature set 422(2) in any technically feasible fashion.

[0093] At step 512, for each frame 210(x), the ensemble prediction engine 490 configures the prediction engine 192(1) to compute a first frame score 434(x) based
15 on the associated set of values for the feature set 422(1) and the perceptual quality model 190(1). At step 514, for each frame 210(x), the ensemble prediction engine 490 configures the prediction engine 192(2) to compute a second frame score 434(x) based on the associated set of values for the feature set 422(2) and the perceptual quality model 190(2).

20 [0094] At step 514, for each frame 210(x), the aggregation engine 440 computes the aggregated frame score 454 based on the first frame score 434(x) and the second frame score 434(x). At step 516, the temporal pooling engine 194 computes the quality score 184 of the re-constructed target 144 based on the aggregated frame scores 454 for the frames 210. The temporal pooling engine 194 may perform any
25 number and type of temporal pooling operations to aggregate the frame scores 434. For instance, in some embodiments, the temporal pooling engine 194 determines the quality score 184 of the re-constructed target 144 based on the arithmetic mean of the frame scores 434 for the frames 210 included in the re-constructed target 144.

[0095] The temporal pooling engine 194 then transmits or displays the quality score
30 184 of the re-constructed target 144 for evaluation of streamed video content. The method 500 then terminates. Advantageously, as described in greater detail in conjunction with Figure 3, any number of the quality scores 184 may be aggregated to

generate an aggregate quality score as part of evaluating the visual quality of streamed media content.

[0096] In sum, the disclosed techniques may be used to efficiently and reliably predict quality scores for re-constructed video content. During a training phase, a training engine generates a perceptual quality model that maps values for features to a quality score. Subsequently, during a prediction phase, a prediction engine computes per-frame quality scores for re-constructed video content based on the perceptual quality model and values for the features of the re-constructed video content. A temporal pooling engine then composites the per-frame quality scores into a single quality score. Notably, during both the training phase and the prediction phase, a temporal extractor computes per-frame values for multiple temporal features, where each temporal feature is associated with a different scale.

[0097] The temporal extractor computes a per-frame value for a specific temporal feature based on re-constructed video content, the video content, and the scale associated with the temporal feature. To compute the per-frame value for a particular frame, the temporal extractor computes a frame difference matrix specifying differences in luminance between the frame and an adjacent frame. The temporal extractor then down-samples the frame difference matrix based on the scale to generate a scaled frame difference matrix. Subsequently, the temporal extractor performs temporal filtering operations and local mean-subtraction operations on the scaled frame difference matrix to compute a first weighted entropy matrix that is associated with the frame of the re-constructed video content. In an analogous fashion, the temporal extractor derives a second weighted entropy matrix that is associated with the corresponding frame of the video content based on the video content and the scale. The temporal extractor performs subtraction operations between the first weighted entropy matrix and the second weighted entropy matrix to compute an entropy difference matrix. Finally, the temporal extractor performs spatial pooling operations on the entropy difference matrix to generate a single frame-specific value for the feature.

[0098] In some embodiments, an ensemble prediction engine computes a quality score for re-constructed video content based on two perceptual quality models. Each of the perceptual quality models maps values for a different set of features to a

different quality score. For each frame included in re-constructed video content, an ensemble prediction engine computes a first frame score based on the first perceptual quality model and a first set of values for the first set of features. For each frame included in re-constructed video content, the ensemble prediction engine computes a
5 second frame score based on the second perceptual quality model and a second set of values for the second set of features. For each frame, the ensemble prediction engine then aggregates the first frame score and the second frame score to generate an aggregated frame score. Finally, the temporal pooling engine generates a single quality score based on the aggregated frame scores for the frames included in the re-
10 constructed video content.

[0099] At least one technological improvement relative to prior art is that computing the entropy difference matrices at multiple scales emulates aspects of how the human visual system processes visual information. Consequently, the temporal features accurately model perceived quality degradation attributable to many types of complex
15 temporal video distortions. Further, computing quality scores based on different perceptual quality models associated with different sets of features reduces the likelihood that increasing the total number of features results in inaccuracies attributable to overfitting.

[0100] 1. In some embodiments, a computer-implemented method comprises
20 computing a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, wherein the first model is generated via a first machine learning algorithm; computing a second quality score for the re-constructed video content based on a second set of values for a second set of features and a second
25 model that associates the second set of values with the second quality score; and determining an overall quality score for the re-constructed video content based on the first quality score and the second quality score, wherein the overall quality score indicates a level of visual quality associated with streamed video content.

[0101] 2. The computer-implemented method of clause 1, wherein the first quality
30 score and the second quality score are associated with a first frame included in the re-constructed video content, and determining the overall quality score comprises determining a first aggregated quality score associated with the first frame based on

first quality score and the second quality score; and performing one or more temporal pooling operations on the first aggregated quality score and a second aggregated quality score associated with a second frame included in the re-constructed video content to compute the overall quality score.

- 5 **[0102]** 3. The computer-implemented method of clauses 1 or 2, wherein determining the first aggregated quality score comprises computing an average or a weighted average of the first quality score and the second quality score.

- 10 **[0103]** 4. The computer-implemented method of any of clauses 1-3, wherein determining the first aggregated quality score comprises associating the first quality score and the second quality score with the first aggregated quality score based on a third model that is generated via the first machine learning algorithm or a second machine learning algorithm.

- 15 **[0104]** 5. The computer-implemented method of any of clauses 1-4, wherein the first set of features includes a first feature that is not included in the second set of features, and the second set of features includes a second feature that is not included in the first set of features.

- 15 **[0105]** 6. The computer-implemented method of any of clauses 1-5, wherein the first set of features includes a temporal feature and a spatial feature.

- 20 **[0106]** 7. The computer-implemented method of any of clauses 1-6, wherein the first set of features comprises at least one of an additive impairment measure feature, a blind or referenceless image spatial quality evaluator feature, and a visual information fidelity feature.

- 25 **[0107]** 8. The computer-implemented method of any of clauses 1-7, further comprising generating the first model based on the first machine learning algorithm and a plurality of subjective scores, wherein each subjective score included in the plurality of subjective scores is associated with different re-constructed test video content.

- 25 **[0108]** 9. The computer-implemented method of any of clauses 1-8, further comprising generating the second model based on the plurality of subjective scores.

[0109] 10. In some embodiments, a computer-readable storage medium includes instructions that, when executed by a processor, cause the processor to perform the steps of computing a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, wherein the first model is generated via a first machine learning algorithm; computing a second quality score for the re-constructed video content based on a second set of values for a second set of features and a second model that associates the second set of values with the second quality score, wherein the second model is generated via the first machine learning algorithm or a second machine learning algorithm; and determining an overall quality score for the re-constructed video content based on the first quality score and the second quality score, wherein the overall quality score indicates a level of visual quality associated with streamed video content.

[0110] 11. The computer-readable storage medium of clause 10, wherein the first quality score and the second quality score are associated with a first frame included in the re-constructed video content, and determining the overall quality score comprises determining a first aggregated quality score associated with the first frame based on first quality score and the second quality score; and performing one or more temporal pooling operations on the first aggregated quality score and a second aggregated quality score associated with a second frame included in the re-constructed video content to compute the overall quality score.

[0111] 12. The computer-readable storage medium of clauses 10 or 11, wherein determining the first aggregated quality score comprises computing an average or a weighted average of the first quality score and the second quality score.

[0112] 13. The computer-readable storage medium of any of clauses 10-12, wherein performing the one or more temporal pooling operations comprises performing at least one of a linear low pass operation, a non-linear rank-order weighting operation, and an arithmetic mean operation on the first aggregated quality score and a second aggregated quality score.

[0113] 14. The computer-readable storage medium of any of clauses 10-13, wherein the first set of features includes a first feature that is not included in the second set of

features, and the second set of features includes a second feature that is not included in the first set of features.

[0114] 15. The computer-readable storage medium of any of clauses 10-14, wherein the first set of features includes a temporal feature and a spatial feature.

5 **[0115]** 16. The computer-readable storage medium of any of clauses 10-15, wherein the first set of features includes a plurality of temporal features, and each temporal feature included in the plurality of temporal features is associated with a different scale.

10 **[0116]** 17. The computer-readable storage medium of any of clauses 10-16, further comprising, prior to computing the first quality score, assigning each feature included in an overall set of features to at least one of the first set of features and the second set of features.

15 **[0117]** 18. The computer-readable storage medium of any of clauses 10-17, further comprising generating the first model based on the first machine learning algorithm and a plurality of subjective scores, wherein each subjective score included in the plurality of subjective scores is associated with different re-constructed test video content.

20 **[0118]** 19. In some embodiments a system comprises a memory storing instructions; and a processor that is coupled to the memory and, when executing the instructions, is configured to compute a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, wherein the first model is generated based on a plurality of subjective scores; compute a second quality score for the re-constructed video content based on a second set of values for a second set of
25 features and a second model that associates the second set of values with the second quality score, wherein the second model is generated based on the plurality of subjective scores; and determine an overall quality score for the re-constructed video content based on the first quality score and the second quality score, wherein the overall quality score indicates a level of visual quality associated with streamed video
30 content.

5 [0119] 20. The system of clause 19, wherein the first quality score and the second quality score are associated with a first frame included in the re-constructed video content, and the processor is configured to determine the overall quality score by determining a first aggregated quality score associated with the first frame based on first quality score and the second quality score; and performing one or more temporal pooling operations on the first aggregated quality score and a second aggregated quality score associated with a second frame included in the re-constructed video content to compute the overall quality score.

10 [0120] Any and all combinations of any of the claim elements recited in any of the claims and/or any elements described in this application, in any fashion, fall within the contemplated scope of the present invention and protection.

15 [0121] The descriptions of the various embodiments have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments.

20 [0122] Aspects of the present embodiments may be embodied as a system, method or computer program product. Accordingly, aspects of the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “module” or “system.” Furthermore, aspects of the present disclosure may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

25 [0123] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access

30

memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer
5 readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0124] Aspects of the present disclosure are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and
10 computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special
15 purpose computer, or other programmable data processing apparatus to produce a machine. The instructions, when executed via the processor of the computer or other programmable data processing apparatus, enable the implementation of the functions/acts specified in the flowchart and/or block diagram block or blocks. Such processors may be, without limitation, general purpose processors, special-purpose
20 processors, application-specific processors, or field-programmable gate arrays.

[0125] The flowchart and block diagrams in the figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may
25 represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may
30 sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that

perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0126] While the preceding is directed to embodiments of the present disclosure, other and further embodiments of the disclosure may be devised without departing
5 from the basic scope thereof, and the scope thereof is determined by the claims that follow.

WHAT IS CLAIMED IS:

1. A computer-implemented method, comprising:
computing a first quality score for re-constructed video content based on a first
set of values for a first set of features and a first model that associates
5 the first set of values with the first quality score, wherein the first model
is generated via a first machine learning algorithm;
computing a second quality score for the re-constructed video content based
on a second set of values for a second set of features and a second
model that associates the second set of values with the second quality
10 score; and
determining an overall quality score for the re-constructed video content based
on the first quality score and the second quality score, wherein the
overall quality score indicates a level of visual quality associated with
streamed video content.
15
2. The computer-implemented method of claim 1, wherein the first quality score
and the second quality score are associated with a first frame included in the re-
constructed video content, and determining the overall quality score comprises:
determining a first aggregated quality score associated with the first frame
20 based on first quality score and the second quality score; and
performing one or more temporal pooling operations on the first aggregated
quality score and a second aggregated quality score associated with a
second frame included in the re-constructed video content to compute
the overall quality score.
25
3. The computer-implemented method of claim 2, wherein determining the first
aggregated quality score comprises computing an average or a weighted average of
the first quality score and the second quality score.
- 30 4. The computer-implemented method of claim 2, wherein determining the first
aggregated quality score comprises associating the first quality score and the second
quality score with the first aggregated quality score based on a third model that is
generated via the first machine learning algorithm or a second machine learning

algorithm.

5. The computer-implemented method of claim 1, wherein the first set of features includes a first feature that is not included in the second set of features, and the
5 second set of features includes a second feature that is not included in the first set of features.

6. The computer-implemented method of claim 1, wherein the first set of features includes a temporal feature and a spatial feature.
10

7. The computer-implemented method of claim 1, wherein the first set of features comprises at least one of an additive impairment measure feature, a blind or referenceless image spatial quality evaluator feature, and a visual information fidelity feature.
15

8. The computer-implemented method of claim 1, further comprising generating the first model based on the first machine learning algorithm and a plurality of subjective scores, wherein each subjective score included in the plurality of subjective scores is associated with different re-constructed test video content.
20

9. The computer-implemented method of claim 8, further comprising generating the second model based on the plurality of subjective scores.

10. A computer-readable storage medium including instructions that, when executed by a processor, cause the processor to perform the steps of:
25 computing a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, wherein the first model is generated via a first machine learning algorithm;
computing a second quality score for the re-constructed video content based
30 on a second set of values for a second set of features and a second model that associates the second set of values with the second quality score, wherein the second model is generated via the first machine learning algorithm or a second machine learning algorithm; and
determining an overall quality score for the re-constructed video content based

on the first quality score and the second quality score, wherein the overall quality score indicates a level of visual quality associated with streamed video content.

- 5 11. The computer-readable storage medium of claim 10, wherein the first quality score and the second quality score are associated with a first frame included in the re-constructed video content, and determining the overall quality score comprises:
- 10 determining a first aggregated quality score associated with the first frame based on first quality score and the second quality score; and performing one or more temporal pooling operations on the first aggregated quality score and a second aggregated quality score associated with a second frame included in the re-constructed video content to compute the overall quality score.
- 15 12. The computer-readable storage medium of claim 11, wherein determining the first aggregated quality score comprises computing an average or a weighted average of the first quality score and the second quality score.
- 20 13. The computer-readable storage medium of claim 11, wherein performing the one or more temporal pooling operations comprises performing at least one of a linear low pass operation, a non-linear rank-order weighting operation, and an arithmetic mean operation on the first aggregated quality score and a second aggregated quality score.
- 25 14. The computer-readable storage medium of claim 10, wherein the first set of features includes a first feature that is not included in the second set of features, and the second set of features includes a second feature that is not included in the first set of features.
- 30 15. The computer-readable storage medium of claim 10, wherein the first set of features includes a temporal feature and a spatial feature.
16. The computer-readable storage medium of claim 10, wherein the first set of features includes a plurality of temporal features, and each temporal feature included

in the plurality of temporal features is associated with a different scale.

17. The computer-readable storage medium of claim 10, further comprising, prior to computing the first quality score, assigning each feature included in an overall set
5 of features to at least one of the first set of features and the second set of features.

18. The computer-readable storage medium of claim 10, further comprising generating the first model based on the first machine learning algorithm and a plurality of subjective scores, wherein each subjective score included in the plurality of
10 subjective scores is associated with different re-constructed test video content.

19. A system, comprising:
a memory storing instructions; and
a processor that is coupled to the memory and, when executing the instructions, is configured to:
15 compute a first quality score for re-constructed video content based on a first set of values for a first set of features and a first model that associates the first set of values with the first quality score, wherein the first model is generated based on a plurality of subjective scores;
20 compute a second quality score for the re-constructed video content based on a second set of values for a second set of features and a second model that associates the second set of values with the second quality score, wherein the second model is generated based on the plurality of subjective scores; and
25 determine an overall quality score for the re-constructed video content based on the first quality score and the second quality score, wherein the overall quality score indicates a level of visual quality associated with streamed video content.

30 20. The system of claim 19, wherein the first quality score and the second quality score are associated with a first frame included in the re-constructed video content, and the processor is configured to determine the overall quality score by:
determining a first aggregated quality score associated with the first frame
based on first quality score and the second quality score; and

performing one or more temporal pooling operations on the first aggregated quality score and a second aggregated quality score associated with a second frame included in the re-constructed video content to compute the overall quality score.

5