The invention claimed is:

1. A computer-implemented method, comprising:

generating multiple sets of encoded chunks for a source video sequence, wherein each set of encoded chunks corresponds to a different shot sequence included in the source video sequence;

generating, for a first set of encoded chunks, a first set of data points based on the first set of encoded chunks;

performing, for the first set of data points, one or more convex hull operations across the first set of data points to compute a first subset of data points that are optimized across at least two metrics;

computing a first slope value between a first data point included in the first subset of data points and a second data point included in the first subset of data points;

comparing the first slope value against one or more other slope values to select a final slope value, wherein each of the one or more other slope values is computed in association with a different set of encoded chunks; and

determining that a first encoded chunk included in the set of encoded chunks associated with the final slope value should be included in a final encoded version of the source video sequence.

2. The computer-implemented method of claim 1, wherein generating the multiple sets of encoded chunks comprises:

identifying within the source video sequence a first sequence of frames that is associated with a first point of capture;

resampling the first sequence of frames at a plurality of different resolutions to generate a resolution ladder of resampled versions of the first sequence of frames; and

encoding each resampled version of the first sequence of frames with a different encoding parameter to generate the first set of encoded chunks.

3. The computer-implemented method of claim 1, wherein generating the first set of data points comprises:

decoding each encoded chunk in the first set of encoded chunks to generate a first set of decoded chunks;

upsampling each decoded chunk in the first set of decoded chunks to a source resolution associated with the source video sequence to generate a first set of upsampled chunks; and

generating a different data point for each upsampled chunk in the first set of upsampled data chunks.

4. The computer-implemented method of claim 3, wherein a specific data point in the first set of data points is generated by:

generating a specific objective quality metric for a specific upsampled chunk in the first set of upsampled chunks;

converting the specific objective quality metric to a specific distortion metric;

computing a bitrate for the specific upsampled chunk; and

combining the specific distortion metric and the bitrate to generate the specific data point.

5. The computer-implemented method of claim 1, wherein performing the one or more convex hull operations across the first set of data points to compute the first subset of data points comprises:

determining a first region that includes the first set of data points;

identifying a first boundary of the first region, wherein no data points in the first set of points reside on a first side of the first boundary; and

discarding any data points that do not reside along the first boundary, wherein each data point that resides along the first boundary optimizes a first metric included in the at least two metrics with respect to a second metric included in the at least two metrics.

6. The computer-implemented method of claim 5, wherein the first metric comprises distortion and the second metric comprises bitrate.

7. The computer-implemented method of claim 1, further comprising:

generating a second set of data points based on a second set of encoded chunks included in the multiple sets of encoded chunks;

performing one or more convex hull operations across the second set of data points to compute a second subset of data points that are optimized across the at least two metrics; and

computing a second slope value between a third data point included in the second subset of data points and a fourth data point included in the second subset of data points.

8. The computer-implemented method of claim 7, wherein determining that the first encoded chunk associated with the final slope value should be included in the final encoded version of the source video sequence comprises determining that the first slope has a greater magnitude than the second slope.

9. The computer-implemented method of claim 7, further comprising determining that a second encoded chunk associated with the fourth data point should be included in another encoded version of the source video sequence based on determining that the second slope value has a greater magnitude than other slope values computed in associated with other sets of encoded chunks.

10. The computer-implemented method of claim 1, wherein the first set of encoded chunks is associated with a first sequence of video frames captured continuously from a first point of capture, and a second set of encoded chunks is associated with a second sequence of video frames captured continuously from a second point of capture.

11. One or more non-transitory computer-readable media storing program instructions that, when executed by one or more processors, configure the one or more processors to perform the steps of:

generating multiple sets of encoded chunks for a source video sequence, wherein each set of encoded chunks corresponds to a different shot sequence included in the source video sequence;

generating, for a first set of encoded chunks, a first set of data points based on the first set of encoded chunks;

performing, for the first set of data points, one or more convex hull operations across the first set of data points to compute a first subset of data points that are optimized across at least two metrics;

computing a first slope value between a first data point included in the first subset of data points and a second data point included in the first subset of data points;

comparing the first slope value against one or more other slope values to select a final slope value, wherein each of the one or more other slope values is computed in association with a different set of encoded chunks; and

determining that a first encoded chunk included in the set of encoded chunks associated with the final slope value should be included in a final encoded version of the source video sequence.

12. The one or more non-transitory computer-readable media of claim 11, wherein the step of generating the multiple sets of encoded chunks comprises:

identifying within the source video sequence a first sequence of frames that is associated with a first point of capture;

resampling the first sequence of frames at a plurality of different resolutions to generate a resolution ladder of resampled versions of the first sequence of frames; and

encoding each resampled version of the first sequence of frames with a different encoding parameter to generate the first set of encoded chunks.

13. The one or more non-transitory computer-readable media of claim 11, wherein the step of generating the multiple sets of encoded chunks comprises:

generating a plurality of values for an encoding parameter based on a plurality of possible values and a maximum number of encoded chunks, wherein the total number of values included in the plurality of values is less than the total number of possible values included in the plurality of possible values; and encoding a plurality of resampled versions of a first sequence of frames based on the plurality of values for the encoding parameter to generate the first set of encoded chunks.

14. The one or more non-transitory computer-readable media of claim 11, wherein the step of generating the first set of data points comprises:

decoding each encoded chunk in the first set of encoded chunks to generate a first set of decoded chunks;

upsampling each decoded chunk in the first set of decoded chunks to a source resolution associated with the source video sequence to generate a first set of upsampled chunks; and

generating a different data point for each upsampled chunk in the first set of upsampled data chunks.

15. The one or more non-transitory computer-readable media of claim 11, wherein the step of performing the one or more convex hull operations across the first set of data points to compute the first subset of data points comprises:

determining a first region that includes the first set of data points; identifying a first boundary of the first region, wherein no data points in the first set of points reside on a first side of the first boundary; and

including any data points that reside along the first boundary in the first subset of data points.

16. The one or more non-transitory computer-readable media of claim 15, wherein a first metric included in the at least two metrics comprises distortion and a second metric included in the at least two metrics comprises bitrate.

17. The one or more non-transitory computer-readable media of claim 11, further comprising the steps of: generating a second set of data points based on a second set of encoded chunks included in the multiple sets of encoded chunks; performing one or more convex hull operations across the second set of data points to compute a second subset of data points that are optimized across the at least two metrics; and

computing a second slope value between a third data point included in the second subset of data points and a fourth data point included in the second subset of data points.

18. The one or more non-transitory computer-readable media of claim 17,

wherein determining that the first encoded chunk associated with the final slope value should be included in the final encoded version of the source video sequence comprises determining that the first slope has a greater magnitude than the second slope.

19. The one or more non-transitory computer-readable media of claim 17, further comprising determining that a second encoded chunk associated with the fourth data point should not be included in another encoded version of the source video sequence based on determining that the second slope value has a smaller magnitude is less than one or more other slope values associated with one or more other subsets of data points.

20. The one or more non-transitory computer-readable media of claim 17, wherein the first set of encoded chunks is associated with a first shot sequence and a second set of encoded chunks is associated with a second shot sequence.

21. A system, comprising:

one or more memories storing a software application; and one or more processors that are coupled to the one or more memories and, when executing the software application, are configured to:

generate multiple sets of encoded chunks for a source video sequence, wherein each set of encoded chunks corresponds to a different shot sequence included in the source video sequence, generate, for a first set of encoded chunks, a first set of data points based on the first set of encoded chunks, perform, for the first set of data points, one or more convex hull operations across the first set of data points to compute a first subset of data points that are optimized across at least two metrics, compute a first slope value between a first data point included in the first subset of data points and a second data point included in the first subset of data points, compare the first slope value against one or more other slope values to select a final slope value, wherein each of the one or more other slope values is computed in association with a different set of encoded chunks, and determine that a first encoded chunk included in the set of encoded chunks associated with the final slope value should be included in a final encoded version of the source video sequence.