WHAT IS CLAIMED IS:

1. A computer-implemented method, comprising:

assigning a first shot sequence included in a source video sequence to a first chunk;

determining that a second shot sequence included in the source video sequence does not satisfy at least one collation criterion with respect to the first chunk;

assigning the second shot sequence or a third shot sequence derived from the second shot sequence to a second chunk;

causing a first encoding instance to generate a first encoded chunk by independently encoding each shot sequence assigned to the first chunk based on encoding parameters determined for the shot sequence; and

causing a second encoding instance to generate a second encoded chunk by independently encoding each shot sequence assigned to the second chunk based on encoding parameters determined for the shot sequence, wherein the first encoded chunk and the second encoded chunk are combined to generate at least part of an encoded video sequence.

2. The computer-implemented method of claim 1, wherein determining that the second shot sequence does not satisfy the at least one collation criterion comprises:

performing an operation that adds the size of the first chunk and the size of the second shot sequence to generate a tentative size; and

determining that the tentative size exceeds a maximum chunk size.

3. The computer-implemented method of claim 2, wherein assigning the second shot sequence or the third shot sequence derived from the second shot sequence to the second chunk comprises:

partitioning the second shot sequence into the third shot sequence and a fourth shot sequence based on the size of the first chunk and the maximum chunk size;

assigning the third shot sequence to the second chunk; and assigning the fourth shot sequence to the first chunk.

4. The computer-implemented method of claim 1, wherein causing the first encoding instance to generate the first encoded chunk by independently encoding each shot sequence assigned to the first chunk comprises generating an encoding task based on the first chunk, and storing the encoding task in a task queue.

5. The computer-implemented method of claim 1, wherein assigning the second shot sequence or the third shot sequence derived from the second shot sequence to the second chunk comprises:

determining that the size of the second shot sequence exceeds a maximum chunk size; partitioning the second shot sequence into the third shot sequence and a fourth shot sequence; and assigning the third shot sequence to the second chunk.

6. The computer-implemented method of claim 1, further comprising, prior to assigning the first shot sequence, performing one or more scene detection operations to identify a plurality of shot sequences included in the source video sequence, wherein the plurality of shot sequences includes the first shot sequence and the second shot sequence.

7. The computer-implemented method of claim 1, wherein the at least one collation criterion specifies a one-to-one mapping of shot sequences to chunks.

8. The computer-implemented method of claim 1, further comprising, prior to assigning the first shot sequence, determining the at least one collation criterion based on at least one of the source video sequence or a resource constraint.

9. The computer-implemented method of claim 1, further comprising, prior to causing the second encoding instance to generate the second encoded chunk by independently encoding each sequence assigned to the second chunk:

determining that a fourth shot sequence included in the source video sequence satisfies the at least one collation criterion with respect to the second chunk; and

assigning the fourth shot sequence to the second chunk.

10. A non-transitory computer-readable storage medium including instructions that, when executed by a processor, cause the processor to perform the steps of:

determining that a first shot sequence included in a source video sequence satisfies one or more collation criteria with respect to a first chunk;

assigning the first shot sequence to the first chunk;

determining that a second shot sequence included in the source video sequence does not satisfy a first collation criterion included in the one or more collation criteria with respect to the first chunk;

assigning the second shot sequence to a second chunk;

causing a first encoding instance to generate a first encoded chunk by independently encoding each shot sequence assigned to the first chunk based on encoding parameters determined for the shot sequence; and

causing a second encoding instance to generate a second encoded chunk by independently encoding each shot sequence assigned to the second chunk based on encoding parameters determined for the shot sequence, wherein the first encoded chunk and the second encoded chunk are combined to generate at least part of an encoded video sequence.

11. The computer-readable storage medium of claim 10, wherein determining that the second shot sequence does not satisfy the first collation criterion comprises:

performing an operation that adds the size of the first chunk and the size of the second shot sequence to generate a tentative size; and

determining that the tentative size exceeds a maximum chunk size.

12. The computer-readable storage medium of claim 11, further comprising, prior to causing the second encoding instance to independently encode each shot sequence assigned to the second chunk:

determining that a third shot sequence does not satisfy the first collation criterion with respect to the second chunk;

partitioning the third shot sequence into a fourth shot sequence and a fifth shot sequence based on the size of the second chunk and the maximum chunk size;

assigning the fourth shot sequence to the second chunk; and

assigning the fifth shot sequence to a third chunk.

13. The computer-readable storage medium of claim 10, further comprising, prior to determining that the first shot sequence satisfies the one or more collation criteria, performing one or more scene detection operations to identify a plurality of shot sequences included in the source video sequence, wherein the plurality of shot sequences includes the first shot sequence and the second shot sequence.

14. The computer-readable storage medium of claim 10, wherein the first collation criterion specifies a one-to-one mapping of shot sequences to chunks.

15. The computer-readable storage medium of claim 10, wherein the first collation criterion specifies whether each shot sequence that is smaller than a chunk size maps to a single chunk.

16. The computer-readable storage medium of claim 10, the steps further comprising, prior to determining that the first shot sequence satisfies the one or more collation criteria, determining the first collation criterion based on at least one of the source video sequence or a resource constraint.

17. The computer-readable storage medium of claim 16, wherein determining the first collation criterion comprises evaluating a plurality of shot sequences included in the source video sequence in conjunction with the resource constraint to determine a chunk size.

18. The computer-readable storage medium of claim 17, wherein the resource constraint comprises a memory bandwidth, the size of a task queue, or a maximum number of encoding tasks.

19. A system, comprising:

a memory storing a shot collation application; and

a processor that is coupled to the memory, wherein, when executed by the processor, the shot collation application causes the processor to:

assign a first shot sequence included in a source video sequence to a first chunk, determine that a second shot sequence included in the source video sequence does not satisfy at least one collation criterion with respect to the first chunk, assign the second shot sequence to a second chunk, cause a first encoding instance to generate a first encoded chunk by independently encoding each shot sequence assigned to the first chunk based on encoding parameters determined for the shot sequence, and cause a second encoding instance that performs encoding operations substantially in parallel to the first encoding instance to generate a second encoded chunk by independently encoding each shot sequence assigned to the second chunk based on encoding parameters determined for the shot sequence, wherein the first encoded chunk and the second encoded chunk are combined to generate at least part of an encoded video sequence.

20. The system of claim 19, wherein the shot collation application causes the processor to determine that the second shot sequence does not satisfy the at least one collation criterion by: performing an operation that adds the size of the first chunk and the size of the second shot sequence to generate a tentative size; and determining that the tentative size exceeds a maximum chunk size.