

CLAIMS

I Claim:

1. A method comprising:
 - accessing an audiovisual composition comprising a target video segment and a source video segment;
 - identifying a first descriptive keyword for the source video segment based on content associated with the source video segment;
 - calculating a primary relevance score for the first descriptive keyword relative to the source video segment;
 - in response to presence of the target video segment and the source video segment in the audiovisual composition:
 - identifying the first descriptive keyword for the source video segment as a first associative keyword for the target video segment; and
 - calculating a secondary relevance score for the first associative keyword relative to the target video segment based on:
 - a temporal position of the source video segment in the audiovisual composition;
 - a temporal position of the target video segment in the audiovisual composition; and
 - the primary relevance score for the first descriptive keyword relative to the source video segment;
 - accessing a textual query comprising the first associative keyword;
 - generating a first query result based on the textual query, the first query result comprising:
 - the source video segment based on the primary relevance score; and
 - the target video segment based on the secondary relevance score; and
 - at a native composition application, rendering a representation of the first query result.
2. The method of Claim 1, wherein identifying the first descriptive keyword for the source video segment based on the content associated with the source video segment comprises:

- executing a taxonomic video classification model on the source video segment to generate a classification vector associated with the source video segment; and
 - extracting the first descriptive keyword based on a classification in the classification vector.
3. The method of Claim 2, wherein calculating the primary relevance score for the first descriptive keyword relative to the source video segment comprises calculating the primary relevance score for the first descriptive keyword based on a confidence score for the classification in the classification vector.
4. The method of Claim 1:
- wherein accessing the audiovisual composition comprises accessing the audiovisual composition comprising:
 - the target video segment;
 - the source video segment; and
 - audio content; and
 - wherein identifying the first descriptive keyword for the source video segment based on the content associated with the source video segment comprises:
 - identifying a concurrent audio segment of the audio content to the source video segment;
 - executing a speech recognition model on the concurrent audio segment to generate a set of audible words; and
 - extracting the first descriptive keyword from the set of audible words.
5. The method of Claim 1, wherein identifying the first descriptive keyword for the source video segment based on the content associated with the source video segment comprises, at the native composition application, receiving a user input designating the first descriptive keyword.
6. The method of Claim 1, wherein identifying the first descriptive keyword for the source video segment based on the content associated with the source video segment comprises:
- at the native composition application, receiving a timestamped comment concurrent with the source video segment; and
 - extracting the first descriptive keyword from the timestamped comment.

7. The method of Claim 1, wherein calculating the secondary relevance score for the first associative keyword relative to the target video segment comprises:
 - calculating an adjacency of the source video segment to the target video segment based on the temporal position of the source video segment in the audiovisual composition and the temporal position of the target video segment in the audiovisual composition; and
 - calculating the secondary relevance score for the first associated keyword based on the adjacency of the source video segment to the target video segment and the primary relevance score for the first descriptive keyword relative to the source video segment.
8. The method of Claim 1, wherein accessing the textual query comprising the first associative keyword comprises, in response to receiving a textual input comprising a first input keyword at the native composition application, appending the first input keyword to the textual query, the first input keyword matching the first associative keyword for the target video segment.
9. The method of Claim 1, wherein accessing the textual query comprising the first associative keyword comprises, in response to receiving a selection of a selected video segment for inclusion in a new audiovisual composition at the native composition application:
 - identifying a second descriptive keyword for the selected video segment; and
 - appending the second descriptive keyword to the textual query, the second descriptive keyword for the selected video segment matching the first associative keyword for the target video segment.
10. The method of Claim 1, wherein accessing the textual query comprising the first associative keyword comprises:
 - in response to receiving a textual input comprising a first input keyword at the native composition application, appending the first input keyword to the textual query;
 - in response to receiving a selection of a selected video segment for inclusion in a new audiovisual composition at the native composition application:
 - identifying a second descriptive keyword for the selected video segment; and
 - appending the second descriptive keyword to the textual query; and

- wherein one of a group consisting of the first input keyword and the second descriptive keyword for the selected video segment matches the first associative keyword of the target video segment.
11. The method of Claim 1:
- further comprising:
 - representing the source video segment as a source video node in a graph database;
 - representing the target video segment as a target video node in the graph database;
 - representing the first descriptive keyword for the source video segment and the first associative keyword for the target video segment as a first keyword node in the graph database;
 - in response to identifying the first descriptive keyword for the source video segment and in response to calculating the primary relevance score for the first descriptive keyword relative to the source video segment, defining a first edge in a graph database connecting the first keyword node and the source video node, the first edge representing the primary relevance score; and
 - in response to identifying the first descriptive keyword for the source video segment as the first associative keyword for the target video segment, defining a second edge in a graph database connecting the first keyword node and the target video node, the second edge representing the secondary relevance score; and
 - wherein generating the first query result based on the textual query comprises:
 - in response to the textual query comprising the first associative keyword, identifying the first keyword node representing the first associative keyword;
 - traversing the first edge from the first keyword node to access the primary relevance score and the source video node;
 - appending the source video segment represented by the source video node to the first query result based on the primary relevance score;
 - traversing the second edge from the first keyword node to access the secondary relevance score and the target video node; and
 - appending the target video segment represented by target video node to the first query result based on the secondary relevance score.

12. The method of Claim 1:

- wherein accessing the audiovisual composition comprising the target video segment and the source video segment comprises accessing the audiovisual composition comprising the target video segment and the source video segment, the target video segment associated with a first user; and
- further comprising:
 - generating a report for the target video segment comprising a set of keywords for the target video segment, the set of keywords comprising the first associative keyword; and
 - serving the report to the first user.

13. The method of Claim 12, wherein generating the report for the target video segment comprises:

- calculating a sentiment score based on the set of keywords for the target video segment; and
- in response to the sentiment score exceeding a threshold sentiment score:
 - generating the report for the target video segment; and
 - serving the report to the first user.

14. A method comprising:

- accessing an audiovisual composition comprising a target video segment represented by a target video node and a source video segment represented by a source video node;
- identifying a first keyword for the source video segment based on content associated with the source video segment, the first keyword represented by a first keyword node;
- calculating a primary relevance score for the first keyword relative to the source video segment;
- updating a first edge connecting the first keyword node and the source video node based on the primary relevance score;
- in response to presence of the target video segment and the source video segment in the audiovisual composition:
 - calculating a secondary relevance score for the first keyword relative to the target video segment based on:
 - a temporal position of the source video segment in the audiovisual composition;

- a temporal position of the target video segment in the audiovisual composition; and
 - the primary relevance score for the first keyword relative to the source video segment; and
 - updating a second edge connecting the first keyword node and the target video node based on the secondary relevance score;
- accessing a textual query comprising the first keyword;
- in response to the textual query comprising the first keyword:
 - identifying the first keyword node representing the first associative keyword;
 - traversing the first edge from the first keyword node to access the source video node;
 - appending the source video segment represented by the source video node to a first query result based on the first edge;
 - traversing the second edge from the first keyword node to access the target video node; and
 - appending the target video segment represented by target video node to the first query result based on the second edge; and
- at a native composition application, rendering a representation of the first query result.

15. The method of Claim 14:

- wherein updating the first edge connecting the first keyword node and the source video node based on the primary relevance score comprises:
 - accessing a first cumulative relevance score associated with the first edge; and
 - updating the first cumulative relevance score based on the primary relevance score;
- wherein updating a second edge connecting the first keyword node and the target video node based on the secondary relevance score comprises:
 - accessing a second cumulative relevance score associated with the second edge; and
 - and
 - updating the second cumulative relevance score based on the secondary relevance score;
- wherein traversing the first edge from the first keyword node to access the source video node comprises traversing the first edge from the first keyword node to access the source video node and the first cumulative relevance score;

- wherein appending the source video segment represented by the source video node to the first query result based on the first edge comprises, in response to the first cumulative relevance score exceeding a threshold cumulative relevance score, appending the source video segment represented by the source video node to the first query result;
- wherein traversing the second edge from the first keyword node comprises traversing the second edge from the first keyword node to access the target video node and the second cumulative relevance score; and
- wherein appending the target video segment represented by target video node to the first query result based on the second edge comprises, in response to the second cumulative relevance score exceeding the threshold cumulative relevance score, appending the target video segment to the first query result.

16. The method of Claim 15:

- wherein updating the first cumulative relevance score based on the primary relevance score comprises summing the first cumulative relevance score and the primary relevance score; and
- wherein updating the second cumulative relevance score based on the secondary relevance score comprises summing the second cumulative relevance score and the secondary relevance score.

17. The method of Claim 15:

- wherein accessing the textual query comprising the first keyword comprises accessing the textual query comprising the first keyword based on input from a user at the native composition application; and
- further comprising:
 - identifying a user node representing the user;
 - traversing a set of user edges connecting the user node to a set of keyword nodes;
 - identifying a first set of video edges connecting the set of keyword nodes to the source video node;
 - identifying a second set of video edges connecting the set of keyword nodes to the target video node;
 - weighting the first cumulative relevance score based on the set of user edges and the first set of video edges; and

- weighting the second cumulative relevance score based on the set of user edges and the second set of video edges.

18. The method of Claim 14:

- wherein updating the second edge connecting the first keyword node and the target video node based on the secondary relevance score comprises updating the second edge connecting the first keyword node and the target video node based on the secondary relevance score, the second edge representing the temporal position of the source video segment relative to the temporal position of the target video segment; and
- wherein rendering the representation of the first query result at the native composition application comprises, at the native composition application:
 - rendering a representation of the source video segment; and
 - rendering a representation of the target video segment positioned relative to the representation of the source video segment based on the second edge representing the temporal position of the source video segment relative to the temporal position of the target video segment.

19. A method comprising:

- accessing an audiovisual composition comprising a target video segment and a source video segment;
- in response to presence of the target video segment and the source video segment in the audiovisual composition:
 - accessing a first keyword associated with the source video segment; and
 - calculating a first relevance score for the first keyword relative to the target video segment based on:
 - a temporal position of the source video segment in the audiovisual composition; and
 - a temporal position of the target video segment in the audiovisual composition;
- accessing a textual query comprising the first keyword;
- generating a first query result based on the textual query, the first query result comprising the target video segment based on the first relevance score; and
- at a native composition application, rendering a representation of the first query result.

20. The method of Claim 19:

- further comprising:
 - identifying the first keyword associated with the source video segment based on content associated with the source video segment; and
 - calculating a second relevance score for the first keyword relative to the source video segment;
- wherein calculating the first relevance score for the first keyword relative to the target video segment comprises calculating a first relevance score for the first keyword relative to the target video segment based on:
 - the temporal position of the source video segment in the audiovisual composition;
 - the temporal position of the target video segment in the audiovisual composition; and
 - the second relevance score for the first keyword relative to the source video segment; and
- wherein generating the first query result based on the textual query comprises generating the first query result based on the textual query, the first query result comprising:
 - the target video segment based on the first relevance score; and
 - the source video segment based on the second relevance score.