CLAIMS

1. A parental control method comprising employing at least one hardware processor of a computer system to:

determine an aggressiveness score, a friendliness score, and a sexual content score of a conversation according to a text content of the conversation, wherein the conversation comprises a sequence of electronic messages exchanged between a first user and a second user, wherein the sequence is ordered according to a timestamp of each message, the aggressiveness score indicating a level of aggressiveness of the conversation, the friendliness score indicating a level of friendliness of the conversation, the sexual content score indicating whether the conversation comprises sexually-explicit language, and wherein at least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation;

determine a first sentiment score according to messages of the conversation sent by the first user, the first sentiment score indicative of an inferred sentiment of the first user;

determine a second sentiment score according to messages of the conversation sent by the second user, the second sentiment score indicative of an inferred sentiment of the second user;

determine whether the first user is bullied by the second user according to the aggressiveness score, friendliness score, sexual content score, and further according to the first and second sentiment scores; and

in response, when the first user is bullied, transmit a parental notification (24) to a parental reporting device (14) identified from a plurality of devices according to the first user, the parental notification indicating that the first user is bullied.

2. The method of claim 1, comprising determining that the first user is not bullied when the aggressiveness score and the sexual content score indicate that the conversation is both aggressive and sexual.

3. The method of claim 1, comprising determining that the first user is not bullied when the aggressiveness score indicates that both the first and second users use aggressive language within the conversation.

4. The method of claim 1, further comprising constructing the conversation in preparation for determining the aggressiveness and friendliness scores, wherein constructing the conversation comprises:

selecting a message from a plurality of electronic messages surreptitiously harvested from a communication device used by the first user; and

including the selected message into the conversation according to whether the message was exchanged between the first user and the second user, and further according to a timestamp of the selected message.

5. The method of claim 1, comprising determining the aggressiveness score according to a count of messages that include aggressive language.

6. The method of claim 1, comprising determining the aggressiveness score according to a count of consecutive messages separating two messages within the conversation that both include aggressive language.

7. The method of claim 1, comprising determining that the first user is not bullied when the aggressiveness and friendliness scores indicate that the conversation is both aggressive and friendly.

8. The method of claim 1, comprising determining that the first user is bullied when the aggressiveness score indicates that a language of the second user is more aggressive than a language of the first user.

9. A computer system comprising at least one hardware processor configured to execute a conversation analyzer (51) and a parental notification dispatcher (59), wherein:

the conversation analyzer (51) is configured to:

analyze a conversation to determine an aggressiveness score, a friendliness score, and a sexual content score of a conversation according to a text content of the conversation, wherein the conversation comprises a sequence of electronic messages exchanged between a first user and a second user, wherein the sequence is ordered according to a timestamp of each message, the aggressiveness score indicating a level of aggressiveness of the conversation, the friendliness score indicating a level of friendliness of the conversation, the sexual content score indicating whether the conversation comprises sexually-explicit language, and wherein at least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation,

determine a first sentiment score according to messages of the conversation sent by the first user, the first sentiment score indicative of an inferred sentiment of the first user,

determine a second sentiment score according to messages of the conversation sent by the second user, the second sentiment score indicative of an inferred sentiment of the second user, and

determine whether the first user is bullied by the second user according to the aggressiveness score, friendliness scores, and sexual content score, and further according to the first and second sentiment scores; and

the parental notification dispatcher (59) is configured, in response to the conversation analyzer determining that first user is bullied, to transmit a parental notification (24) to a parental reporting device (14) identified from a plurality of devices according to the first user, the notification message indicating that the first user is bullied.

10. The computer system of claim 9, wherein the conversation analyzer is configured to determine that the first user is not bullied when the aggressiveness score and the sexual content score indicate that the conversation is both aggressive and sexual.

11. The computer system of claim 9, wherein the conversation analyzer is configured to determine that the first user is not bullied when the aggressiveness score indicates that both the first and second users use aggressive language within the conversation.

12. The computer system of claim 9, further comprising a message aggregator configured to construct the conversation in preparation for determining the aggressiveness and friendliness scores, wherein constructing the conversation comprises:

selecting a message from a plurality of electronic messages surreptitiously harvested from a communication device used by the first user; and

including the selected message into the conversation according to whether the message was exchanged between the first user and the second user, and further according to a timestamp associated with the selected message.

13. The computer system of claim 9, wherein the conversation analyzer is configured to determine the aggressiveness score according to a count of messages that include aggressive language.

14. The computer system of claim 9, wherein the conversation analyzer is configured to determine the aggressiveness score according to a count of consecutive messages separating two messages within the conversation that both include aggressive language.

15. The computer system of claim 9, wherein the conversation analyzer is configured to determine that the first user is not bullied when the aggressiveness and friendliness scores Indicate that the conversation is both aggressive and friendly.

16. The computer system of claim 9, wherein the conversation analyzer is configured to determine that the first user is bullied when the aggressiveness score indicates that a language of the second user is more aggressive than a language of the first user.

17. A non-transitory computer-readable medium storing instructions which, when executed by at least one hardware processor of a computer system, cause the computer system to form a conversation analyzer (51) and a parental notification dispatcher (59), wherein:

the conversation analyzer (51) is configured to:

analyze a conversation to determine an aggressiveness score, a friendliness score, and a sexual content score of a conversation according to a text content of the conversation, wherein the conversation comprises a sequence of electronic messages exchanged between a first user and a second user, wherein the sequence is ordered according to a timestamp of each message, the aggressiveness score indicating a level of aggressiveness of the conversation, the friendliness score indicating a level of friendliness of the conversation, the sexual content score indicating whether the conversation comprises sexually-explicit language, and wherein at least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation,

determine a first sentiment score according to messages of the conversation sent by the first user, the first sentiment score indicative of an inferred sentiment of the first user,

determine a second sentiment score according to messages of the conversation sent by the second user, the second sentiment score indicative of an inferred sentiment of the second user, and

determine whether the first user is bullied by the second user according to the aggressiveness score, friendliness score, and sexual content score, and further according to the first and second sentiment scores; and

the parental notification dispatcher (59) is configured, in response to the conversation analyzer determining that first user is bullied, to transmit a parental notification (24) to a parental reporting device (14) identified from a plurality of devices according to the first user, the notification message indicating that the first user is bullied.