

Anti-Cyberbullying Systems and Methods

RELATED APPLICATIONS

[0001] This application claims the benefit of the filing date of U.S. provisional patent application No. 62/794,856, filed on Jan. 21, 2019, entitled "Parental Control Systems and Methods," the entire contents of which are incorporated by reference herein.

BACKGROUND

[0002] The present invention relates to computer security, and in particular to systems and methods for protecting vulnerable Internet users (e.g., children) against online threats such as cyberbullying, online abuse, grooming, sexual exploitation, and theft of confidential information, among others.

[0003] Bullying is commonly defined as the activity of repeated, aggressive behavior intended to hurt another individual physically, mentally, or emotionally. Bullying behavior may manifest itself in various ways, such as verbally, physically, etc. When bullying occurs via modern means of communication such as electronic messaging and posting on social media, it is commonly referred to as cyberbullying. Successful bullying typically requires an imbalance of power and/or peer pressure, the weak side being at the receiving end of the abuse. Bullying is known to cause serious distress, even leading to suicide in some cases. Some social categories (children, young adults, members of a racial or sexual minority) may be more exposed to such threats than others.

[0004] With the explosive growth of the Internet, children and teens are spending a significant amount of time browsing and communicating online, at a point in their physical and emotional development where they are particularly vulnerable to threats such as bullying, sexual exploitation, and identity theft. The problem is amplified by the fact that the online culture of social media and instant messaging does not easily lend itself to supervision by traditional authority figures (parents, teachers, etc.), either because young users are often more

technologically savvy than their guardians, or because the communication platforms themselves do not allow snooping.

[0005] In recent years, security software has been used successfully to protect computer users from computer threats such as malicious software (malware) and intrusion (hacking). There is currently substantial interest in developing software capable of protecting users against other emerging threats such as cyberbullying, grooming, sexual exploitation, and online harassment, ideally while preserving the privacy of their electronic messaging.

SUMMARY

[0006] According to one aspect, a parental control method comprises employing at least one hardware processor of a computer system to analyze a conversation to determine an aggressiveness score and a friendliness score. The conversation comprises a sequence of electronic messages exchanged between a first user and a second user. The aggressiveness score indicates a level of aggressiveness of the conversation, while the friendliness score indicates a level of friendliness of the conversation. At least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation. The method further comprises employing the at least one hardware processor to determine whether the first user is bullied by the second user according to the aggressiveness and friendliness scores. The method further comprises employing the at least one hardware processor in response, when the first user is bullied, to transmit a parental notification to a parental reporting device identified from a plurality of devices according to the first user, the parental notification indicating that the first user is bullied.

[0007] According to another aspect a computer system comprises at least one hardware processor configured to execute a conversation analyzer and a parental notification dispatcher. The conversation analyzer is configured to analyze a conversation to determine an aggressiveness score and a friendliness score. The conversation comprises a sequence of electronic messages exchanged between a first user and a second user. The aggressiveness score indicates a level of aggressiveness of the conversation, while the friendliness score indicates a

level of friendliness of the conversation. At least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation. The conversation analyzer is further configured to determine whether the first user is bullied by the second user according to the aggressiveness and friendliness scores. The parental notification dispatcher is configured, in response to the conversation analyzer determining that first user is bullied, to transmit a parental notification to a parental reporting device identified from a plurality of devices according to the first user, the notification message indicating that the first user is bullied.

[0008] According to another aspect, a non-transitory computer-readable medium stores instructions which, when executed by at least one hardware processor of a computer system, cause the computer system to form a conversation analyzer and a parental notification dispatcher. The conversation analyzer is configured to analyze a conversation to determine an aggressiveness score and a friendliness score. The conversation comprises a sequence of electronic messages exchanged between a first user and a second user. The aggressiveness score indicates a level of aggressiveness of the conversation, while the friendliness score indicates a level of friendliness of the conversation. At least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation. The conversation analyzer is further configured to determine whether the first user is bullied by the second user according to the aggressiveness and friendliness scores. The parental notification dispatcher is configured, in response to the conversation analyzer determining that first user is bullied, to transmit a parental notification to a parental reporting device identified from a plurality of devices according to the first user, the notification message indicating that the first user is bullied.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The foregoing aspects and advantages of the present invention will become better understood upon reading the following detailed description and upon reference to the drawings where:

[0010] Fig. 1 illustrates an exemplary parental control system wherein a monitored device engaging in electronic messaging is protected against online threats according to some embodiments of the present invention.

5 [0011] Fig. 2-A shows an exemplary data exchange between a monitored device, a security server, and a reporting device according to some embodiments of the present invention.

[0012] Fig. 2-B shows an alternative data exchange between the monitored device, security server, and reporting device according to some embodiments of the present invention.

[0013] Fig. 3 shows exemplary software components executing on the monitored device according to some embodiments of the present invention.

10 [0014] Fig. 4 illustrates the operation of an exemplary parental control application executing on the monitored device according to some embodiments of the present invention.

[0015] Fig. 5 illustrates an exemplary conversation indicator according to some embodiments of the present invention.

15 [0016] Fig. 6 shows an exemplary sequence of steps carried out by the parental control application according to some embodiments of the present invention.

[0017] Fig. 7 shows an exemplary sequence of steps performed by message aggregator to construct a set of conversations according to some embodiments of the present invention.

[0018] Fig. 8 shows exemplary software components executing on the security server according to some embodiments of the present invention.

20 [0019] Fig. 9 shows an exemplary operation of the software components illustrated in Fig. 8.

[0020] Fig. 10 shows an exemplary sequence of steps performed by the security server according to some embodiments of the present invention.

[0021] Fig. 11 illustrates a set of exemplary text processors according to some embodiments of the present invention.

[0022] Fig. 12 shows a set of exemplary image processors according to some embodiments of the present invention.

5 [0023] Fig. 13 illustrates exemplary body parts that an image processor is trained to detect in an image, according to some embodiments of the present invention.

[0024] Fig. 14 shows an exemplary hardware configuration of a computation device configured to carry out parental control operations according to some embodiments of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

10 [0025] In the following description, it is understood that all recited connections between structures can be direct operative connections or indirect operative connections through intermediary structures. A set of elements includes one or more elements. Any recitation of an element is understood to refer to at least one element. A plurality of elements includes at least two elements. Unless otherwise specified, any use of “OR” refers to a non-exclusive or. Unless
15 otherwise required, any described method steps need not be necessarily performed in a particular illustrated order. A first element (e.g. data) derived from a second element encompasses a first element equal to the second element, as well as a first element generated by processing the second element and optionally other data. Making a determination or decision according to a parameter encompasses making the determination or decision according to the parameter and
20 optionally according to other data. Unless otherwise specified, an indicator of some quantity/data may be the quantity/data itself, or an indicator different from the quantity/data itself. A minor is a person under the age of full legal responsibility. A computer program is a sequence of processor instructions carrying out a task. Computer programs described in some embodiments of the present invention may be stand-alone software entities or sub-entities (e.g.,
25 subroutines, libraries) of other computer programs. Computer readable media encompass non-transitory media such as magnetic, optic, and semiconductor storage media (e.g. hard drives, optical disks, flash memory, DRAM), as well as communication links such as conductive cables

and fiber optic links. According to some embodiments, the present invention provides, *inter alia*, computer systems comprising hardware (e.g. one or more processors) programmed to perform the methods described herein, as well as computer-readable media encoding instructions to perform the methods described herein.

5 **[0026]** The following description illustrates embodiments of the invention by way of example and not necessarily by way of limitation.

10 **[0027]** Fig. 1 shows an exemplary parental control system protecting a user of a monitored device against online threats such as cyberbullying, sexual exploitation, and theft of confidential information, among others. In a typical scenario according to some embodiments of the present invention, the protected user (e.g., a minor) employs messaging software executing on a monitored device **10** (e.g., a smartphone) to exchange electronic messages with users of other messaging partner devices **12a-b**. In some embodiments, security software executing on monitored device **10** and/or a remote security server **18** may be used to snoop on such conversations, typically without knowledge of the respective user. Conversations are then
15 analyzed for content. When the security software determines according to the conversation content that the user is subject to an online threat, some embodiments transmit a notification to another party (e.g., parent, teacher, manager, etc.) via a reporting device **14** such as a smartphone or personal computer.

20 **[0028]** Monitored device **10** may comprise any electronic device having a processor and a memory, and capable of connecting to a communication network for exchanging electronic messages with messaging partner devices **12a-b**. Exemplary monitored devices **10** include personal computers, laptop computers, tablet computers, smartphones, gaming consoles, virtual assistant devices, household appliances (e.g., smart TVs, media players, refrigerators), and wearable computer devices (e.g., smartwatches).

25 **[0029]** An electronic message comprises a communication transmitted between two electronic devices, the communication including at least an encoding of a text message between two human users of the respective devices. Electronic messaging is typically carried out using an instant

messaging platform such as FACEBOOK® Messenger®, Instagram® Direct®, Snapchat®, WhatsApp®, etc., via electronic mail (email), and/or via a telephony messaging service such as short message service (SMS). Messaging platforms comprise software configured to enable a user to send and receive electronic messages to/from other users. Messages may vary in format according to the respective platform/service, but in general, an electronic message comprises an encoding of a text part and/or an encoding of a media file (e.g., image, movie, sound, etc.) The text part may comprise text written in a natural language (e.g., English, Chinese, etc.), and other alphanumeric and/or special characters such as emoticons, among others. In a typical configuration, messages are coordinated, centralized, and dispatched by a messaging server **16**, in the sense that electronic messages between monitored device **10** and partner devices **12a-b** are routed via server **16** (client-server protocol). In alternative embodiments, electronic messaging uses a de-centralized peer-to-peer network of connections between monitored devices and their respective messaging partner devices. Monitored device **10**, messaging partner device(s) **12a-b** and messaging server **16** are interconnected by a communication network **15** such as the Internet. Parts of network **15** may include a local area network (LAN), and a telecommunication network (e.g., mobile telephony).

[0030] Threat detection operations may be divided between monitored device **10** and security server **18** in various ways, as shown in detail below. Server **18** generically represents a set of interconnected computers which may or may not be in physical proximity to each other. Figs **2-A-B** show exemplary data exchanges between monitored device **10** and security server **18** according to some embodiments of the present invention. In various embodiments, monitored device **10** may transmit conversation data (represented by conversation indicator **20** in Fig. **2-A**) and/or threat-indicative information (represented by risk assessment indicator **22** in Fig. **2-B**) to security server **18**. At least a part of the conversation analysis/threat detection may then be carried out by components executing on security server **18**. When the analysis indicates a potential threat to a user of monitored device **10**, some embodiments of security server **18** send a parental notification **24** to a reporting device **14** (e.g., mobile telephone, personal computer, etc.) associated with the respective monitored device, thus informing a user of reporting device **14** about the respective threat. The term 'parental' is herein used only for simplicity and is not

meant to be limiting in the sense that the receiver of the respective notification is necessarily a parent, or that the protected user is necessarily a child. Although typical applications of some embodiments are in parental control, a skilled artisan will understand that they can be adapted to monitoring and/or protecting other categories of users/devices. In child monitoring applications, notification **24** may be sent to a teacher, guardian, or any other person charged with supervising the respective child. In other exemplary applications directed at protecting employees against bullying and/or sexual harassment, notification **24** may be delivered to a manager, supervisor, or human resources staff, for instance. Exemplary formats and contents of notification **24** are shown further below.

[0031] Fig. **3** shows exemplary software components executing on monitored device **10** according to some embodiments of the present invention. Operating system **46a** may comprise any widely available operating system such as Microsoft Windows®, MacOS®, Linux®, iOS®, or Android®, among others. OS **46a** provides an interface between other computer programs (represented by applications **48** and **50**) and hardware devices of monitored device **10**.

[0032] Messaging application **48** generically represents any software configured to enable a user of device **10** to exchange electronic messages with other users. Exemplary messaging applications **48** include Yahoo® Messenger®, FACEBOOK®, Instagram®, and Snapchat® client applications, among others. Another exemplary messaging application **48** comprises an email client. Yet another exemplary messaging application **48** comprises software implementing a short message service (SMS) on a mobile telephone. Application **48** may display a content of each electronic message on an output device (e.g., screen) of monitored device **10** and may further organize messages according to sender, recipient, time, subject, or other criteria. Application **48** may further receive text input from a user of device **10** (e.g., from a keyboard, touchscreen, or dictation interface), formulate electronic messages according to the received text input, and transmit electronic messages to messaging server **16** and/or directly to messaging partner device(s) **12a-b**. Message format and encoding may vary according to the messaging platform. Transmitting a message may comprise, for instance, adding an encoding of the respective message to an outbound queue of a communication interface of monitored device **10**.

[0033] In some embodiments, parental control application **50** comprises software configured to access, harvest, and/or analyze a content of message exchanges between monitored device **10** and partner device(s) **12a-b**. Parental control application **50** may be part of a larger computer security software suite comprising anti-malware and intrusion detection tools, among others.

5 Fig. **4** shows exemplary components of parental control application **50** according to some embodiments of the present invention.

[0034] A data grabber **52** is configured to extract message content generated and/or received by messaging application **48**. Extracting message content may comprise identifying individual electronic messages and determining message-specific features such as a sender and/or receiver, a time of transmission (e.g., timestamp), a text of the respective message, and possibly other content data such as an image attached to the respective message. Content extraction may proceed according to any method known in the art. In some embodiments, data grabber **52** surreptitiously modifies a component of messaging application **48** (for instance by hooking) to install a software agent that notifies data grabber when application **48** executes some specific operation such as receiving a communication or receiving user input, and enables data grabber **52** to extract message information. Some embodiments extract message content using built-in features of OS **46a** such as an accessibility application programming interface (API). Accessibility APIs comprise software typically configured to grab information currently displayed on an output device (e.g., screen) of monitored device **10** for the purpose of making such information accessible to people with disabilities. One exemplary application of such accessibility APIs comprises translating on-screen text into audio (spoken text) to enable visually impaired people to use the computer. Some embodiments of data grabber **52** are configured to call specific accessibility API functions to parse data structures such as user interface trees while device **10** is displaying content generated by messaging application **48**, and thus extract information such as message interlocutor names/aliases and a content of individual messages. Yet another embodiment of data grabber **52** may extract message content directly from intercepted network traffic going into messaging application **48** and/or passing via a network adapter(s) of monitored device **10**. Such communication interceptors may implement communication protocols such as HTTP, WebSocket, and MQTT, among others, to parse

communications and extract structured message data. When instant messages are encrypted, some embodiments employ techniques such as man-in-the-middle (MITM) to decrypt traffic for message content extraction.

5 [0035] Some embodiments of the present invention rely on the observation that threats such as bullying and sexual grooming typically involve complex social dynamics, and therefore are more accurately inferred from an extended conversation, as opposed to individual messages. In some embodiments therefore, a message aggregator **54** may aggregate individual messages into conversations consisting of multiple messages exchanged between the same pair of interlocutors (in the case of a one-to-one exchange), or within the same group (in the case of a group chat, for
10 instance). Message aggregator **54** may collaborate with data grabber **52** to identify a sender and/or receiver of each intercepted message, organize a message stream into individual conversations, and output a conversation indicator **20**. The operation of message aggregator **54** is further detailed below.

[0036] An exemplary conversation indicator **20** illustrated in Fig. 5 comprises a user_ID
15 identifying monitored device **10** and/or an individual user of the respective device, and a pair_ID uniquely identifying a pair of interlocutors. In some embodiments, conversation indicator **20** further includes a plurality of message indicators Message_1...Message_n, each corresponding to an individual message exchanged between the respective interlocutors. Individual message indicators may in turn include an identifier of a sender and/or of a receiver, a text content of each
20 message (represented as MessageText_i in Fig. 5), and a timestamp indicating a moment in time when the respective message was sent and/or received. In an alternative embodiment, conversation indicator **20** comprises a concatenation of the text content of all messages in the respective conversation, individual messages arranged in the order of transmission according to their respective timestamp.

25 [0037] Conversation indicator **20** may further include a set of media indicators (represented as MediaFile_j in Fig. 5), for instance copies of image/video/audio files attached to messages belonging to the respective conversation, or a network address/URL where the respective media

file is located. Other exemplary media indicators may include an indicator of a media format (encoding protocol), etc. A skilled artisan will understand that the actual data format for encoding conversation indicator **20** may differ among embodiments; exemplary formats include a version of an extensible markup language (XML), and Javascript Object Notation (JSON), among others.

[0038] Fig. **6** shows an exemplary sequence of steps performed by parental control application **50** according to some embodiments of the present invention. Fig. **7** further illustrates an exemplary algorithm for constructing conversations out of individual messages (step **204** in Fig. **6**).

[0039] Parental control application **50** may represent each conversation as a separate data structure (e.g., an object with multiple data fields). Conversations may be defined according to various criteria, such as length (e.g., total count of messages, total word count) and/or time (e.g., messages exchanged in a pre-determined time interval). In some embodiments, a conversation is considered to be alive as long as its count of messages does not exceed a predetermined value; alternatively, a conversation may be considered alive as long as the time elapsed since its first message does not exceed a predetermined time threshold, and/or as long as a time elapsed since its latest message does not exceed another predetermined time threshold. Conversations which are no longer alive are herein deemed expired. In one example illustrated in Fig. **7**, parental control application **50** monitors multiple live conversations, each conversation identified by a unique conversation ID. A step **212** determines an amount of time elapsed since the latest message of each live conversation. When said amount of time exceeds a pre-determined threshold (e.g., one hour), message aggregator **54** may consider the respective conversation closed/expired and remove it from the set of live conversations. A further step **230** may formulate conversation indicator **20** of the respective conversation and transmit the respective data away for further analysis. A similar flowchart may describe the operation of an alternative message aggregator that considers a conversation to be closed when the count of messages exceeds a pre-determined threshold (e.g., 500).

[0040] Meanwhile, data grabber **52** may listen for new messages (step **216**). When a message is detected, a step **220** may identify the interlocutors of the respective message, for instance by parsing message data, or by analyzing the user interface of messaging application **48** (see above, in relation to using Accessibility APIs). When there is currently at least a live conversation with the respective interlocutors, in a step **226**, aggregator **54** may add data characterizing the current message to a conversation object identified by the current interlocutors (e.g., pair ID). When there is currently no live conversation between the interlocutors of the current message, a step **224** may initialize a new conversation object identified by the current interlocutors/pair ID and may add message data to the newly initialized object. Application **50** may then return to listening for new messages and/or determining whether any live conversation has expired.

[0041] Figs. **8-9** illustrate exemplary software components executing on security server **18**, and an exemplary operation of such components, respectively, according to some embodiments of the present invention. Fig. **10** further details the operation of said components as an exemplary sequence of steps.

[0042] In some embodiments, conversation data is received from message aggregator **54** in the form of conversation indicator(s) **20**. Each indicator **20** may represent a single conversation, which in turn may comprise multiple messages exchanged between the same interlocutors over a specified time period. In some embodiments, conversation indicators **20** accumulate in a queue, awaiting further processing. Such processing may comprise selecting a conversation and removing it from the queue (steps **302-204-306** in Fig. **10**). The selected indicator **20** is then fed to a conversation analyzer **51**, which analyzes a content of the respective conversation to determine a plurality of assessment indicators (e.g., numerical or Boolean scores, category labels, etc.) and output them to a decision unit **53**. In a further step **312**, decision unit **53** may aggregate analysis results received from conversation analyzer **51** and apply a set of decision criteria to determine whether a user of monitored device **10** is subject to an online threat such as bullying, sexual harassment, grooming, etc. In some embodiments, decision unit **53** communicates a unified risk assessment indicator **22** to a notification dispatcher **59**. In a step **314**, notification dispatcher may determine whether a notification condition is satisfied

according to the received assessment indicator. When yes, dispatcher **59** may formulate and transmit parental notification **24** to reporting device **14**. Some of the above steps will be further detailed below.

[0043] In some embodiments, conversation analyzer **51** comprises a set of text processors **56** configured to analyze a text content of a conversation, and/or a set of image processors **58** configured to analyze an image and/or video content of a conversation. Each processor **56-58** may analyze each conversation according to a distinct aspect of the respective conversation and/or according to a distinct algorithm. For instance, each processor **56-58** may determine whether a user is subject to a different type of threat (bullying, sexual harassment, grooming, etc.) In another example, there may be multiple processors detecting the same type of threat, but each processor may use a different criterion or algorithm, or may consider a different aspect of the analyzed conversation. For instance, some text processors may search the analyzed conversation for certain keywords, while others may employ a neural network to produce a score or a label characterizing the respective message or conversation, etc. Other exemplary conversation aspects include aggressiveness, friendliness, and sexual content, among others.

[0044] In some embodiments, a text content of a conversation is normalized in preparation for feeding to at least some of text processors **56** (step **308** in Fig. **10**). Such normalization may include spellchecking, expanding acronyms, detecting and interpreting emojis, URLs, person and/or location names. Normalization may comprise looking up a dictionary of the respective natural language (e.g., English), augmented with slang items and various expressions/acronyms frequently used in instant messaging.

[0045] Some exemplary text and image processors **56-58** are illustrated in Figs. **11-12**, respectively. Each text processor **56** may output a text assessment indicator **26**. Similarly, image processors **58** may output a set of image assessment indicators **28**. One exemplary text assessment indicator **26** includes a numerical score indicative of a likelihood that at least one interlocutor is the subject of an online threat (e.g., bullying) according to a text content of the respective conversation. An exemplary image assessment indicator **28** may indicate whether the

current conversation comprises at least one image belonging to a particular category (e.g., a nude or partially nude picture, a photograph of a credit card, etc.)

[0046] Some embodiments of the present invention rely on the observation that due to the complexity of social interactions, which are especially emotionally charged in childhood and adolescence, a single algorithm/analysis protocol is unlikely to successfully detect threats such as bullying. For instance, children and teens often address each other using demeaning nicknames, insults, and derogatory language, even when they mean no harm. Such “posturing” is simply seen as cool or a fun thing to do. Therefore, a text analysis algorithm merely aimed at detecting insults and/or conflict-indicative language may wrongly classify a cocky exchange between close friends as a word fight or instance of bullying. To avoid such false positives, some embodiments employ multiple natural language processing algorithms to analyze various aspects of each conversation and extract a variety of assessment indicators. Some embodiments then increase the reliability of threat detection by aggregating information provided by multiple individual assessment indicators. Image assessment indicators may be combined with text assessment indicators. For instance, a nude picture may provide an additional clue to a suspicion of sexting, etc.

[0047] Exemplary text processors 56 illustrated in Fig. 11 include, among others, an aggressiveness assessor, a friendliness assessor, a sexual content assessor, a sentiment assessor, and a text confidentiality assessor. Each text processor 56 may output a set of scores, labels, etc. Such scores/labels may be determined for each individual message of the conversation, or may be determined for the respective conversation as a whole.

[0048] An exemplary aggressiveness assessor computes a score for each message of a conversation, the score indicative of a level of aggression indicated by the language of the respective message. The aggressiveness score may be expressed as a binary number (1/0, YES/NO), or as non-binary number which may take any value between pre-determined bounds. Aggressiveness assessors may employ methods such as detecting the presence of certain aggression-indicative keywords, or any other method known in the art. A preferred embodiment

trains a recurrent neural network (RNN) using a vector representation of each word in a dictionary. Exemplary vector representations can be obtained using a version of a word-2-vec and/or Glove families of algorithms. Each message of a conversation may then be represented as a sequence of vectors. The architecture of the aggressiveness assessor may include, among
5 others, a long short-term memory (LSTM) stacked on top of a gated recurrent unit (GRU) layer. Training may enforce particular rules, e.g., only insults formulated in the second person may be labeled as positive/aggressive. In one such example, the phrase “you are so stupid” may score higher for aggressiveness than “he is so stupid”. The output of such a neural network may comprise a score/label determined for each individual message, or a score/label determined for
10 the whole conversation.

[0049] The architecture of an exemplary sexual content assessor may be similar to the one described for the aggressiveness assessor. However, the sexual content assessor may be specifically trained to output a score indicating whether each conversation and/or message contains sexual language. Sometimes sexual and aggressive language co-exist in a conversation,
15 so this is an example wherein having independent assessors for each aspect of a conversation may produce a more nuanced and possibly more accurate classification of the respective conversation. Some embodiments may be further trained to identify other text patterns which may not be sexually explicit, but may nevertheless indicate grooming or sexual predation. For instance, some embodiments may detect whether a message is asking for a meeting, for a
20 personal address, etc. Some embodiments of the sexual content assessor may be trained to distinguish between multiple scenarios and/or categories of sexual content (e.g., grooming, sexting, etc.) In one such example, the sexual content assessor may output a vector of scores, each score corresponding to a distinct category/scenario and indicating a likelihood that the analyzed conversation falls within the respective category/scenario.

[0050] An exemplary friendliness assessor aims to detect phrases that display affection and a friendly attitude towards one or the other of the interlocutors. Since friends often tease each other using offensive language, a friendliness indicator/score may help distinguish true abuse from behaviors that could appear aggressive, but are in fact playful and benign. An exemplary
25

friendliness assessor may employ a rule-based system to identify direct positive phrasing towards a conversation partner (e.g., “I like you”) and/or indirect phrasing, wherein a current message is used to respond positively to a previous one (e.g., “do you like me?”, “God, you ARE stupid. Sure I do. You’re the best”). This is another example wherein text content analysis is correlated
5 across multiple messages of the same conversation, as opposed to analyzing each message separately.

[0051] An exemplary sentiment assessor may employ any method known in the art to determine a numerical or categorical indicator of mood/sentiment of the respective conversation. An exemplary indicator may have positive values when the conversation is deemed happy/relaxed,
10 and negative values when the conversation indicates stress, depression, anger, etc. The value of the respective assessment indicator may indicate an intensity of the respective sentiment. An exemplary sentiment assessor uses a Valence Aware Dictionary and Sentiment Reasoner (VADER) methodology, wherein each token of a message (e.g., each word or phrase) is labelled according to its semantic orientation as either positive or negative, and an aggregate score/label
15 is computed by combining individual token labels. The aggregate score may be computed at the granularity of individual messages, or for the conversation as a whole. In some embodiments, an aggressive conversation wherein only one side is feeling bad/upset is a strong indication that bullying is under way. Such a situation may hence receive a relatively high aggregated bullying-indicative score for the respective conversation.

20 [0052] An exemplary text confidentiality assessor may determine whether a conversation communicates sensitive information which the respective user (e.g., child, employee) should not be sharing with others. Some examples of such information are credit card numbers, social security numbers, and home addresses, among others. One exemplary text confidentiality assessor may use character pattern matching (e.g., regular expressions) to identify data such as
25 credit card numbers and addresses. Other embodiments may train a neural network to detect text patterns that look like credit card information, social security numbers, etc. A text confidentiality assessor may output a vector of scores, each score indicating whether the text of

the current conversation contains a distinct category of confidential data. Such text confidentiality scores may be determined for the conversation as a whole.

[0053] In some embodiments, image processor **58** (see Figs. **5** and **10**) use a variety of methods/algorithms to detect various features of image and/or video data exchanged as part of a conversation. One exemplary image processor **58** comprises a nudity assessor configured to return a score indicative of a likelihood that an image contains nudity. In an alternative embodiment, the nudity assessor may return a plurality of scores, each score indicating a likelihood that the image shows a particular body part (e.g., face, breast, nipple, leg), and/or whether the respective image is likely to belong to a particular type of imagery (sexual activity, sunbathing, etc.). In some embodiments, the nudity assessor is further configured to return an indicator of whether each visible body part is naked or covered.

[0054] Fig. **13** shows a few illustrative body parts **60a-c** that an exemplary nudity assessor is trained to discover in an image file transmitted as part of a conversation according to some embodiments of the present invention. Each body part **60a-c** comprises a part of a human body, such as a head, face, hair, chest, cleavage, breast, nipple, under breast, abdomen, navel, lower waist, crotch, genitals, anus, buttock, sacrum, lower back, middle back, shoulder blade, neck, nape, upper arm, lower arm, hand, thigh, upper leg, lower leg, knee, and foot, among others. Some such body parts may overlap. Some embodiments are further trained to determine whether a body part detected in an image belongs to a man or a woman.

[0055] In a preferred embodiment, the nudity assessor may comprise a set of interconnected artificial neural networks, for instance a stack of convolutional neural networks further feeding into a fully connected layer. The respective nudity assessor may receive the analyzed image as input and may be configured to output a set of scores and/or labels. The neural networks may be trained on a corpus of annotated images. Training a neural network may comprise iteratively adjusting a set of functional parameters (e.g., connection weights) of the respective neural network in an effort to reduce a mismatch between the actual output of the network and a desired output such as the one provided by annotation.

[0056] Another exemplary image processor 58 comprises an image confidentiality assessor configured to return a score indicative of a likelihood that the respective image contains confidential information. Examples of confidential image data include an image of a bank card, an image of an official identification document such as a driver's license, social security card or passport, an image of a car license plate, an image of a user's home/school, etc. Bank cards include credit and debit cards, among others.

[0057] In some embodiments, the image confidentiality assessor comprises a set of interconnected artificial neural networks (e.g., convolutional neural networks) trained to input an image and output a set of scores and/or labels indicative of a likelihood that the image falls within a specific category (for instance that the image shows a specific type of physical object, such as a bank card.) The respective neural networks may be trained on an annotated corpus containing images of various kinds of documents in various contexts, for instance bank cards issued by various banks and having various designs, passports and/or driver's licenses issued by various countries, etc.

[0058] Some embodiments detect the presence of an physical object in an image according to characteristic features of the respective physical object. For instance, to detect the presence of a bank card, an exemplary image confidentiality assessor may be trained to detect an image of a magnetic strip, an image of a handwritten signature located in the vicinity of a magnetic strip, an image of an embedded microchip, an image of 16 digits aligned and divided in groups of four (i.e., the card number), an image of the VISA® or MASTERCARD® logo, etc. In the case of a social security card, the image confidentiality assessor may be trained to determine whether the analyzed image comprises a logo of the Social Security Administration and/or a set of 11 digits aligned and divided into three groups (i.e., the social security number). Driver's licenses and passports may also be identified according to characteristic features, such as a photograph of a human head, and a specific placement of various data on the respective document.

[0059] In some embodiments, the image confidentiality assessor (e.g., neural network) may be trained to output a plurality of scores, each score indicative of a likelihood that the analyzed

image shows a distinct feature of a bank card, social security card, etc. For instance, one score may indicate a likelihood that the image shows an embedded card microchip, another score may indicate a likelihood that the image show a VISA® logo, etc. Such individual scores may then be aggregated by decision unit **53**, for instance using a weighted average or a decision algorithm.

5 **[0060]** Some embodiments of the image confidentiality assessor may be further trained to extract structured data from the analyzed images. For instance, in addition to determining that an image shows a credit card, some embodiments may determine a type of card (e.g., VISA®), an issuing bank, etc. Similarly, in addition to detecting an image of a driver's license, some embodiments may automatically determine a name of the driver, etc.

10 **[0061]** In some embodiments, decision unit **53** (Fig. **9**) inputs individual assessment indicators **26-28** received from text and/or image processors **56-58**, respectively, and outputs an aggregated risk assessment indicator **22** determined according to the individual risk assessment indicators. An exemplary aggregated risk assessment indicator **22** is determined for the conversation as a whole and comprises a set of scores wherein each score indicates a likelihood
15 of a distinct type of threat or scenario (e.g., fighting, bullying, depression, sexual exposure, grooming, loss of confidential data, etc.). Aggregate indicators **22/scores** may be computed using any method known in the art. One example comprises computing a weighted average of individual assessment indicators/scores. In another example, an aggregate score is determined according to a decision algorithm: if score *x* is YES and score *y* is below 0.4, then the aggregated
20 score is 0.8.

[0062] An aggregate score for bullying may be determined according to the following observations. Bullying language typically occurs in scattered bursts, rather than being distributed uniformly throughout the conversation. There are typically multiple such bursts within an abusive conversation. In some embodiments, to qualify as bullying, aggressive
25 language should be persistent within an individual burst (i.e., a single offensive message does not indicate bullying). More often than not, children are using offensive language without the intent to cause harm. Usually in this kind of interactions (i.e., non-bullying), there is evidence of

friendly language, and the general tone of the conversation is rather positive or neutral. Finally, in many cases aggressive language and sexual language overlap.

[0063] Using the above observations, some embodiments scan each conversation with a fixed-length window and a fixed step (i.e., a pre-determined number of consecutive messages at a time). For each such conversation segment and each interlocutor, decision unit 53 may create a vector wherein each element represents a combined score determined for a distinct individual message of the respective conversation. Individual text assessment indicators may be combined as follows:

$$S_i = A_i - F_i - X_i, \quad [1]$$

[0064] wherein S_i denotes a message-specific combined score, A_i and F_i denote an aggressiveness and a friendliness score of the respective message, and X_i denotes a sexual content (e.g., sexting) score of the respective message. Some of the following situations may occur, for instance: if a message is only aggressive, the respective combined score S_i is 1; if the message is detected to be both aggressive and sexual, the combined score S_i is 0 (sexual language cancels aggressive language); if the message is detected to be both aggressive and friendly, the combined score S_i is also 0 (friendly language cancels aggressive language).

[0065] A further step may compute an aggressiveness concentration score for the current conversation segment, for instance using the formula:

$$C = \frac{1}{N} \sum_i \left(S_i + \frac{S_i}{d_i} \right), \quad [2]$$

[0066] wherein N denotes the total number of messages within the respective conversation segment, S_i is the combined score of each message of the respective segment, and d_i denotes a distance (e.g., count of messages) between the current message and the closest aggressive message (e.g., combined score $S_i = 1$). Formula [2] yields a relatively higher value for conversation segments that have closely-spaced aggressive messages compared to other conversation segments. In a subsequent step, the value of C may be compared to a pre-

determined threshold (e.g., 0.5). A value in excess of the threshold indicates an aggressive user for the respective segment of conversation. The calculation may be performed separately for each interlocutor and each segment of the conversation. Some embodiments may then determine a type of relationship between the interlocutors, for instance: normal – none of the interlocutors is excessively aggressive; bullying - one of the interlocutors is substantially more aggressive than the other; fight - both interlocutors are substantially and equally aggressive. For a verdict of bullying, some embodiments may further determine whether the bully is the user of monitored device **10** or not, to enable notification dispatcher **59** to include such information in parental notification **24**.

[0067] In some embodiments, conversation-specific risk assessment indicators/scores **22** are compared with a set of pre-determined thresholds specific to each type of threat. A score exceeding the respective threshold may indicate a presence of the respective threat/scenario. When a score exceeds the respective threshold, some embodiments of notification dispatcher **59** may formulate and send parental notification **24** to reporting device **14**.

[0068] Several exemplary conversation snippets and their associated scoring are shown below.

[0069] Example 1: Bullying

User	Message	Aggressiveness	Friendliness	Sexting
A	hey faggot	1	0	0
B	stop calling me that	0	0	0
A	or what, are u going to run to your fat mom?	1	0	0
A	you lame ass fag	1	0	0
B	stop it	0	0	0

[0070] This example is characterized by a substantial imbalance of aggressiveness: user A is abusing user B, and B is not responding in kind. An aggressiveness concentration indicator determined according to formula [2] yields 2.0 for user A and 0.0 for user B. Some embodiments compare the difference in aggressiveness between the two users to a threshold (e.g., 0.5), and since the difference exceeds the threshold, determine that user A is substantially more aggressive than B. Therefore, A is bullying B.

[0071] Example 2: Non-bullying

User	Message	Aggressiveness	Friendliness	Sexting
A	you being a bitch right now	1	0	0
A	you know I like you, but you're overreacting	0	1	0
B	I'm going to kill you for that	1	0	0
A	hey babe, it was nothing	0	1	0
A	I love you	0	1	0
B	yea I guess you're right	0	0	0
B	I love you too	0	1	0

[0072] In this example aggressive language coexists with friendly language. However, the friendliness score exceeds the aggressiveness score, and formula [2] yields sub-zero values for both interlocutors. Therefore, the conversation is not classified as bullying.

[0073] Example 3: Sexting

User	Message	Aggressiveness	Friendliness	Sexting
------	---------	----------------	--------------	---------

A	maybe you can blow me	1	0	1
A	you sexy mofo	1	0	1
B	I'm getting horny already	0	0	1
B	can you come over?	0	1	0

[0074] In this example, aggressive language coexists with sexting, and therefore the aggressiveness score is cancelled out. The conversation does not qualify as bullying, even though only one side of the conversation is aggressive.

[0075] Thresholds and/or other scoring parameters (e.g., weights given to specific scores) may be tailored and/or adjusted per monitored device, user, and/or category of users, for instance according to a subscription type or service-level agreement, thus providing a degree of vigilance that is customizable. Some embodiments rely on the observation that what is considered 'acceptable behavior' may vary widely between countries, cultures, and even individuals. For instance, in some countries and cultures, women are required to cover their hair in public, so an image of a woman with an uncovered head may be seen as unacceptably revealing, whereas in other cultures it is completely normal. The same is true for other body parts, such as an ankle or an upper arm. Even in Western societies, conservative families are stricter on the behavior of children and teens than more liberal ones. For instance, a short skirt may be considered normal for some, and too revealing for others. Therefore, in some embodiments of the present invention, thresholds and/or score aggregation strategies may be adjusted according to personal choice, cultural criteria and/or according to a geographical location of devices **10** and/or **14**. In one such example, when installing and/or configuring software on monitored device **10** and/or reporting device **14**, a user may be shown a configuration interface and invited to customize a set of criteria for receiving parental notifications. For instance, the user may be invited to select an overall level of vigilance (e.g., on a sliding scale from 0 to 10), and/or to select a subset of scenarios that should trigger notifications from a broader set of exemplary scenarios. In some embodiments, the software may automatically choose a set of parameter values (e.g., notification

scenarios, vigilance level, etc.) according to a country and/or culture of the respective user. Such automatic selection may include, for instance, automatically detecting a geolocation of reporting device **14** and/or monitored device **10** and looking up a database of location-specific settings. Such settings may then be automatically translated into specific threshold values and/or other score aggregation parameters.

[0076] In determining aggregated assessment indicator **22**, decision unit **53** may combine text assessment indicators with image assessment indicators determined for the same conversation. In one such example, output of the sexual content assessor (text) may be combined with output of the nudity assessor (image) to generate an aggregate sexual content score. When the text of a conversation includes sexual content, the respective content may merely represent vulgar language used in a word fight. The score given to the respective conversation by the sexual content assessor may therefore not be high enough to trigger classifying the conversation into a sexual threat category. However, when the respective conversation also includes a revealing image, the score given by the nudity assessor may be combined with the score returned by the text processor, to produce an aggregate score that exceeds the respective threshold. Scores may be combined for instance as a weighted average, wherein each individual weight may reflect a relevance of the respective score to a particular threat/situation. In the example of sexual content, the score produced by the image processor may receive a higher weight than the score produced by the text processor.

[0077] In preparation for sending parental notification **24**, dispatcher **59** may identify the appropriate receiver of the respective notification, i.e., reporting device **14**. In some embodiments, parental control services are provided in accordance with a subscription and/or a service level agreement (SLA). To provide such services, some embodiments maintain a database of subscription or account entries, wherein each entry may be attached to a reporting device **14**, so that notifications generated in relation to the respective subscription/account are delivered to the respective reporting device. Reporting device **14** may be identified for instance according to a network address, or according to a unique identifier generated by a software agent executing on device **14** and configured to collaborate with server **18** in delivering notifications.

The respective subscription/account entry may further indicate a set of monitored devices **10** as targets for collecting and analyzing conversation data. A single subscription/account may service multiple monitored devices **10**. The subscription/account entry may further indicate a monitored user of device **10**, for instance as a username, alias and/or avatar used by the monitored user (e.g., minor) inside messaging application **48**. The association between monitored devices **10**, users, and reporting device **14** enables dispatcher **59** to selectively identify reporting device **14** according to an identity of a monitored device **10** and/or according to an identity of a user of device **10**.

[0078] Delivering notification **24** may proceed according to any method known in the art, for instance by pushing notification **24** to a software agent/application executing on reporting device **14**, including notification **24** an email or SMS message, etc.

[0079] Parental notification **24** may comprise a notification message formulated in a natural language, e.g. English. Notification messages may include an indicator of a detected incident/scenario/threat, e.g., child is bullied, child has sent confidential information, etc. To preserve the monitored user's privacy, some embodiments do not reveal actual message contents to parents/guardians/administrators. Some embodiments further include parenting advice/suggestions of how to address the respective detected scenario or threat, and/or a set of psychology resources (hyperlinks, literature references, etc.) relevant to the respective detected incident/threat. The notifications may be formulated as much as possible to not alarm the parent/guardian, and to not reveal the identity of the parties involved in the respective conversations. When bullying is detected, the notification message may indicate whether the user of monitored device **10** is the perpetrator or the receiver of the abuse. Some examples of notification messages are given below.

Cyberbullying & anti-predator

[0080] Examples of notifications in response to scenarios/threats detected using text analysis:

[0081] Child is in a one-to-one verbal fight: “Jenny is involved in a verbal dispute in a chat named [conversation name]. When talking about online behavior, remind your child that aggression is never 'just a joke.' Even if it seems like everybody is doing it, it is not ok.

[0082] Child is bullied one to one: “Jenny is cyberbullied in a chat named [conversation name].

5 Listen to your child. Find out exactly what happened, how she felt and why. Here are a few pointers on how to start a conversation: [link]”

[0083] Child is in a group chat and is the only one being aggressive: “Jenny is involved in a verbal dispute in a group chat named [conversation name]. Her behavior seems aggressive toward interlocutors. When talking about online behavior, remind your child aggression is never

10 'just a joke.' Even if it seems like everybody is doing it, it's not ok.”

[0084] Child is in a group verbal fight but is not aggressive: “Jenny is involved in a verbal dispute in a group chat named [conversation name]. Her behavior does not seem aggressive toward interlocutors. When talking about online behavior, remind your child aggression is never 'just a joke.' Even if it seems like everybody is doing it, it's not ok.”

15 **Child security, sex predation and grooming**

[0085] Examples of notifications in response to scenarios/threats detected using text analysis:

[0086] Child receives a personal address: “Jenny received a personal address in a conversation named [conversation name]. Remain calm and talk face to face with your child about the importance of privacy. Here are a few pointers on how to start a conversation: [link]”

20 [0087] Child is asked for a face-to-face meeting: “Jenny received a request for a meeting in a chat named [conversation name]. Talk calmly with your child about the meeting request, preferably in person.”

[0088] Child accepts a face-to-face meeting: “Jenny accepted a meeting in a conversation named [conversation name] at 7 pm. It doesn't necessarily mean something is wrong. Maybe it's just

25 good that you know.”

[0089] Child is asked for pictures: “Jenny received a request for photos in a chat named [conversation name]. Remind your kids that sharing everything makes them vulnerable. Both online and offline, some things should remain private.”

5 [0090] Child has a new contact in Social Media: “Jenny started talking online with a new contact in a chat named [conversation name]. Keep a good balance. A new contact doesn't necessarily mean trouble, nor that should you interrogate your child. Nevertheless, you might want to check out the new contact's profile and pay attention to any changes in your child's behavior.”

[0091] Examples of notifications in response to scenarios/threats detected using image or combined text and image analysis:

10 [0092] Child receives a picture that is too revealing. “Jenny received a photo containing nudity in a chat named [conversation name]. If your child has been sent a sexual image or video, advise him/her to delete it immediately, and not to share it with anyone. Have a conversation about this later on.”

15 [0093] Child sends a picture that is too revealing: “Jenny sent a picture containing nudity in a chat named [conversation name]. Keep calm and start a conversation with your child by asking 'If you got into a fight with this person, would you like them to have this photo of you?'.

20 [0094] Child has a revealing picture stored on her device: “Inappropriate media content detected on Jenny's handheld device. If your child has been sent a sexual image or video, advise him/her to delete it immediately, and not to share it with anyone. Have a conversation about this later on.”

Confidentiality, identity theft, and family security

[0095] Examples of notifications in response to events/threats detected using text analysis:

25 [0096] Child is asked for her personal address: “Someone asked Jenny for a personal address in a chat named [conversation name]. Remind your child that your address should only be shared, with your consent, with certain people.”

[0097] Child sends her personal address: “Jenny sent a personal address in a conversation named [conversation name]. Remain calm and talk face to face with your child about the importance of privacy.”

5 [0098] Child is asked for credit card numbers: “Jenny was asked for a credit card number in a chat named [conversation name]. Calmly but firmly advise your child not to disclose such information. Discuss the consequences in further detail, face to face.”

[0099] Child is asked for Social Security Number: “Jenny was asked for a social security number in a chat named [conversation name]. Calmly but firmly advise your child not to disclose such information. Discuss the consequences in further detail, face to face.”

10 [0100] Child sends a credit card number in a conversation: “Jenny sent a credit card number in a conversation named [conversation name]. Remain calm and talk face to face with your child about the risks of giving away financial information.”

[0101] Child sends a social security number in a conversation: “Jenny sent a social security number in a conversation named [conversation name]. Remain calm and talk face to face with
15 your child about the risks of giving away private information”.

[0102] Examples of notifications in response to events/threats detected using image or combined image and text analysis:

[0103] Child sends a photo of a credit card: “Jenny sent a picture of a credit card in a conversation named [conversation_name]. It appears to be a Mastercard issued by Chase. The
20 details of this credit card are now online and in the possession of another person. You should secure the respective card following the bank's procedure. Afterwards, have a calm, yet firm conversation about this with your child.”

[0104] Child sends a photo of a social security card: “Jenny sent a picture of a social security card in a conversation named [conversation_name]. The social security number is now in the
25 possession of another person. Various confidential data is attached to your social security

number. Here are a few steps you can take to check whether the gained information is being used for malicious purposes such as identity theft [link]. Have a calm, yet firm conversation about this with your child.”

5 [0105] Child sends a photo of an identity document: “Jenny sent a picture of an identity document in a conversation named [conversation_name]. The document appears to be an American passport. Have a calm, yet firm conversation with your child about identity theft and/or the risks of giving away personal information online.”

10 [0106] Child has a photo stored on her device, the photo showing a credit card. “A picture of a credit card has been detected on Jenny's handheld device. Remain calm and talk face to face with your child about the risks of giving away financial information.”

[0107] Child has a photo stored on her device, the photo showing an identity card: “A picture of an identity card has been detected on Jenny's handheld device. Remain calm and talk face to face with your child about the risks of giving away personal information.”

15 [0108] Although the above description relates to a configuration as described in Figs. 3 and 8, a skilled artisan will understand that alternative embodiments may use another distribution of software components. For instance, in some embodiments, conversation analyzer 51 and decision unit 53 may execute on monitored device 10 instead of on server 18 as illustrated in Fig. 8. In such configurations, a typical data exchange between device 10 and server 18 is illustrated in Fig. 2-B. Such configurations may have the advantage that all message content
20 stays on monitored device 10, thus ensuring the privacy of the respective user. A disadvantage is that operations of conversation analyzer 51 and/or decision unit 53 are typically computationally expensive, and may put an unacceptable burden on a relatively modest device such as a mobile telephone or tablet computer. Another potential disadvantage of carrying out text and/or image processing at monitored device 10 is the necessity of distributing software updates to all such
25 devices. In contrast, when conversation analysis is carried out at security server 18, a single machine may process conversation data received from multiple (possibly thousands) of monitored devices.

[0109] In yet another alternative embodiment, message aggregator **54** may execute on security server **18** instead of on monitored device **10** as illustrated in Fig. 4. In such configurations, aggregator **54** may receive message indicators from a plurality of devices. Such configurations may allow aggregating conversations across multiple devices, for instance when a user starts a
5 FACEBOOK® messaging exchange on a smartphone, but continues it later from a personal computer.

[0110] Fig. 14 shows an exemplary hardware configuration of a computing device **70** programmed to execute some of the methods described herein. Device **70** may represent any of monitored device **10**, security server **18**, and reporting device **14** in Fig. 1. The illustrated
10 configuration is that of a personal computer; other computing devices such as mobile telephones, tablet computers, and wearables may have slightly different hardware. Processor(s) **72** comprise a physical device (e.g. microprocessor, multi-core integrated circuit formed on a semiconductor substrate) configured to execute computational and/or logical operations with a set of signals and/or data. Such signals or data may be encoded and delivered to processor(s) **72** in the form of
15 processor instructions, e.g., machine code. Processor(s) **72** may include a central processing unit (CPU) and/or an array of graphics processing units (GPU.)

[0111] Memory unit **74** may comprise volatile computer-readable media (e.g. dynamic random-access memory – DRAM) storing data/signals/instruction encodings accessed or generated by processor(s) **72** in the course of carrying out operations. Input devices **76** may include computer
20 keyboards, mice, and microphones, among others, including the respective hardware interfaces and/or adapters allowing a user to introduce data and/or instructions into computing device **70**. Output devices **78** may include display devices such as monitors and speakers among others, as well as hardware interfaces/adapters such as graphic cards, enabling computing device **70** to communicate data to a user. In some embodiments, input and output devices **76-78** share a
25 common piece of hardware (e.g., a touch screen.) Storage devices **82** include computer-readable media enabling the non-volatile storage, reading, and writing of software instructions and/or data. Exemplary storage devices include magnetic and optical disks and flash memory devices, as well as removable media such as CD and/or DVD disks and drives. Network adapter(s) **84**

enable computing device **70** to connect to an electronic communication network (e.g., network **15** in Fig. 1) and/or to other devices/computer systems.

[0112] Controller hub **80** generically represents the plurality of system, peripheral, and/or chipset buses, and/or all other circuitry enabling the communication between processor(s) **22** and the rest of the hardware components of device **70**. For instance, controller hub **80** may comprise a memory controller, an input/output (I/O) controller, and an interrupt controller. Depending on hardware manufacturer, some such controllers may be incorporated into a single integrated circuit, and/or may be integrated with processor(s) **72**. In another example, controller hub **80** may comprise a northbridge connecting processor **72** to memory **74**, and/or a southbridge connecting processor **72** to devices **76**, **78**, **82**, and **84**.

[0113] It will also be apparent to one of ordinary skill in the art that aspects of the invention, as described above, may be implemented in various forms of software, firmware, and hardware, or a combination thereof. For example, certain portions of the invention may be described as specialized hardware logic that performs one or more functions. This specialized logic may include an application specific integrated circuit (ASIC) or a field programmable gate array (FPGA). The actual software code or specialized control hardware used to implement aspects consistent with the principles of the invention is not limiting of the present invention. Thus, the operation and behavior of the aspects of the invention were described without reference to the specific software code – it being understood that one of ordinary skill in the art would be able to design software and control hardware to implement the aspects based on the description herein.

[0114] The exemplary systems and methods described herein allow protecting vulnerable Internet users (e.g., minors) against online threats such as cyberbullying, online abuse, grooming, sexual harassment or exploitation, and theft of confidential information, among others. Such systems and methods typically fall in the category of parental control. However, some systems and methods described herein may extend beyond classical parental control applications, for instance to detecting online abuse such as racist, sexist, or homophobic attacks perpetrated against adults using online messaging services.

[0115] In some embodiments, electronic messages exchanged by a monitored user are automatically and surreptitiously harvested from an electronic device (e.g., smartphone, tablet computer) the respective user is using to communicate. Messages are then selectively aggregated into conversations comprised of messages exchanged between the same interlocutors. Conversation data is then analyzed according to various criteria. When analysis concludes that the monitored user is subject to an online threat, some embodiments transmit a notification message to a supervisor of the respective user (e.g., parent, teacher, manager, etc.).

[0116] Some embodiments rely on the observation that the social dynamics involved in dangerous scenarios such as bullying are relatively complex. Therefore, in determining whether a user is subject to such an online threat, some embodiments aggregate information from multiple messages and/or multiple aspects of a conversation. For instance, some embodiments combine a result of analyzing a text of a conversation with a result of analyzing an image transmitted as part of the respective conversation. In turn, the analysis of the text part may also be multifaceted: some embodiments combine evaluations of an aggressiveness, friendliness, and sexual content of a conversation.

[0117] Other exemplary embodiments combine image analysis with text analysis to determine whether a monitored user is engaging in risky behavior of disclosing confidential information such as credit card data and social security numbers, among others. In one such example, images harvested from electronic messages are analyzed to determine whether they comprise a photograph of a bank card, social security card, driver's license, etc. Discovery of such an image may trigger a parental notification.

[0118] It will be clear to one skilled in the art that the above embodiments may be altered in many ways without departing from the scope of the invention. Accordingly, the scope of the invention should be determined by the following claims and their legal equivalents.

CLAIMS

What is claimed is:

1. A parental control method comprising employing at least one hardware processor of a computer system to:
 - analyze a conversation to determine an aggressiveness score and a friendliness score, wherein the conversation comprises a sequence of electronic messages exchanged between a first user and a second user, the aggressiveness score indicating a level of aggressiveness of the conversation, the friendliness score indicating a level of friendliness of the conversation, and wherein at least one of the aggressiveness score and the friendliness score is determined according to multiple messages of the conversation;
 - determine whether the first user is bullied by the second user according to the aggressiveness and friendliness scores; and
 - in response, when the first user is bullied, transmit a parental notification to a parental reporting device 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, further comprising:
 - determining a first sentiment score and a second sentiment score according to the conversation, the first sentiment score indicative of an inferred sentiment of the first user and the second sentiment score indicative of an inferred sentiment of the second user; and
 - determining whether the first user is bullied further according to the first and second sentiment scores.
3. The method of claim 1, further comprising:

2 determining a sexual content score according to the conversation, the sexual content
3 score indicating whether the conversation comprises sexually-explicit
4 language; and
5 determining whether the first user is bullied further according to the sexual content
6 score.

7

1 4. The method of claim 1, wherein the aggressiveness score indicates whether the first
2 user uses aggressive language within the conversation.

3

1 5. The method of claim 1, comprising determining the aggressiveness score according to
2 whether multiple messages of the conversation include aggressive language.

3

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

3

1 7. The method of claim 5, comprising determining the aggressiveness score
2 according to a count of consecutive messages separating two messages
3 within the conversation that both include aggressive language.

4

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

4

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

4

1 10. A computer system comprising at least one hardware processor configured to execute a
2 conversation analyzer and a parental notification dispatcher, wherein:

3 the conversation analyzer is configured to:

4 analyze a conversation to determine an aggressiveness score and a friendliness
5 score, wherein the conversation comprises a sequence of electronic
6 messages exchanged between a first user and a second user, the
7 aggressiveness score indicating a level of aggressiveness of the
8 conversation, the friendliness score indicating a level of friendliness of the
9 conversation, and wherein at least one of the aggressiveness score and the
10 friendliness score is determined according to multiple messages of the
11 conversation, and

12 determine whether the first user is bullied by the second user according to the
13 aggressiveness and friendliness scores; and

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

18
1 11. The computer system of claim 10, wherein the conversation analyzer is further
2 configured to:

3 determine a first sentiment score and a second sentiment score according to the
4 conversation, the first sentiment score indicative of an inferred sentiment
5 of the first user and the second sentiment score indicative of an inferred
6 sentiment of the second user; and

7 determine whether the first user is bullied further according to the first and second
8 sentiment scores.

9
1 12. The computer system of claim 10, wherein the conversation analyzer is further
2 configured to:

3 determine a sexual content score according to the conversation, the sexual content
4 score indicating whether the conversation comprises sexually-explicit
5 language; and
6 determine whether the first user is bullied further according to the sexual content
7 score.

8
1 13. The computer system of claim 10, wherein the aggressiveness score indicates whether
2 the first user uses aggressive language within the conversation.

3
1 14. The computer system of claim 10, wherein the conversation analyzer is configured to
2 determine the aggressiveness score according to whether multiple messages of the
3 conversation include aggressive language.

4
1 15. The computer system of claim 14, wherein the conversation analyzer is
2 configured to determine the aggressiveness score according to a count of
3 messages that include aggressive language.

4
1 16. The computer system of claim 14, wherein the conversation analyzer is
2 configured to determine the aggressiveness score according to a count of
3 consecutive messages separating two messages within the conversation
4 that both include aggressive language.

5
1 17. The computer system of claim 10, wherein the conversation analyzer is configured to
2 determine that the first user is not bullied when the aggressiveness and
3 friendliness scores indicate that the conversation is both aggressive and friendly.

4
1 18. The computer system of claim 10, wherein the conversation analyzer is configured to
2 determine that the first user is bullied when the aggressiveness score indicates that

3 a language of the second user is substantially more aggressive than a language of
4 the first user.

- 5
1 19. A non-transitory computer-readable medium storing instructions which, when executed
2 by at least one hardware processor of a computer system, cause the computer system to
3 form a conversation analyzer and a parental notification dispatcher, wherein:
4 the conversation analyzer is configured to:

5 analyze a conversation to determine an aggressiveness score and a friendliness
6 score, wherein the conversation comprises a sequence of electronic
7 messages exchanged between a first user and a second user, the
8 aggressiveness score indicating a level of aggressiveness of the
9 conversation, the friendliness score indicating a level of friendliness of the
10 conversation, and wherein at least one of the aggressiveness score and the
11 friendliness score is determined according to multiple messages of the
12 conversation, and

13 determine whether the first user is bullied by the second user according to the
14 aggressiveness and friendliness scores; and

15 the parental notification dispatcher is configured, in response to the conversation analyzer
16 determining that first user is bullied, to transmit a parental notification to a
17 parental reporting device identified from a plurality of devices according to the
18 first user, the notification message indicating that the first user is bullied.
19