



US 20180309710A1

(19) **United States**

(12) **Patent Application Publication**  
**Sutton et al.**

(10) **Pub. No.: US 2018/0309710 A1**

(43) **Pub. Date: Oct. 25, 2018**

(54) **CLASSIFYING SOCIAL ENTITIES AND APPLYING UNIQUE POLICIES ON SOCIAL ENTITIES BASED ON CROWD-SOURCED DATA**

(60) Provisional application No. 61/867,004, filed on Aug. 16, 2013.

**Publication Classification**

(71) Applicant: **Nexgate, Inc.**, Burlingame, CA (US)  
(72) Inventors: **Richard Banks Sutton**, San Mateo, CA (US); **Daniel Oshiro Nadir**, Carlsbad, CA (US); **Devin Hugh Redmond**, Encinitas, CA (US); **Harold Nguyen**, Burlingame, CA (US); **Anthony Lawrence Dorie**, San Francisco, CA (US); **Jorge Iriso**, Alameda, CA (US); **Wenchang Yan**, Burlingame, CA (US)

(51) **Int. Cl.**  
**H04L 12/58** (2006.01)  
**H04L 29/08** (2006.01)  
**G06F 17/30** (2006.01)  
**G06Q 50/00** (2012.01)  
(52) **U.S. Cl.**  
CPC ..... **H04L 51/32** (2013.01); **H04L 67/22** (2013.01); **H04L 51/12** (2013.01); **G06Q 50/01** (2013.01); **G06F 17/30705** (2013.01)

(21) Appl. No.: **16/019,428**

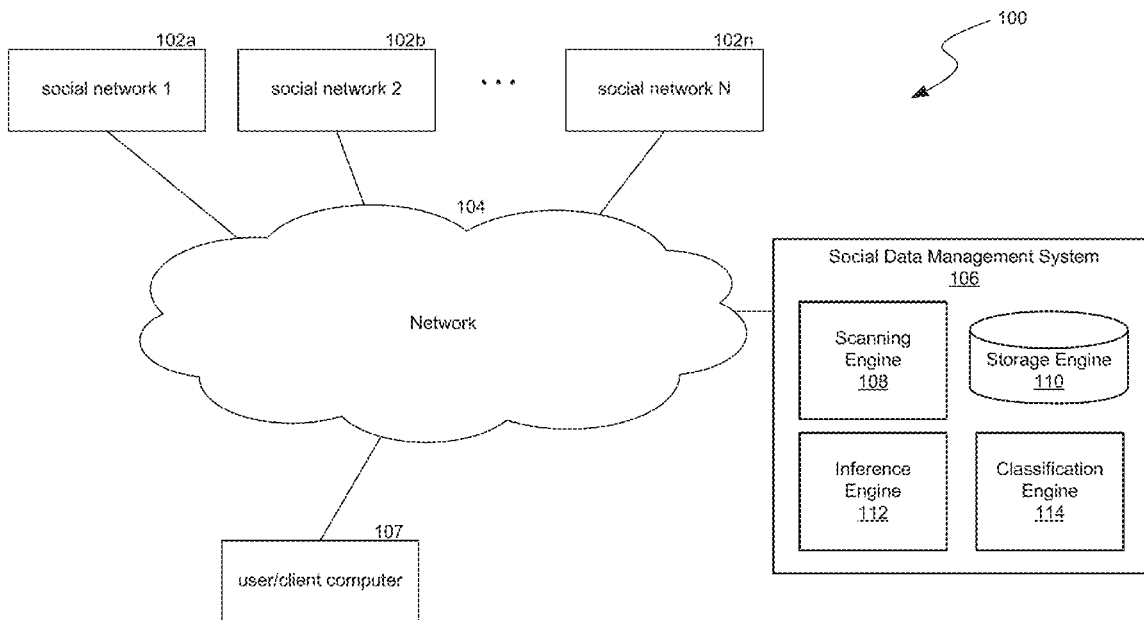
(22) Filed: **Jun. 26, 2018**

**Related U.S. Application Data**

(63) Continuation of application No. 14/306,098, filed on Jun. 16, 2014, now Pat. No. 10,033,684.

(57) **ABSTRACT**

Technology is disclosed for detecting, classifying, and/or enforcing rules on social networking activity. The technology can scan and collect social content data from one or more social networks, store the social content data, classify content data posted to a social network, create and apply a set of social data content rules to future posted social content data.






FIG. 1

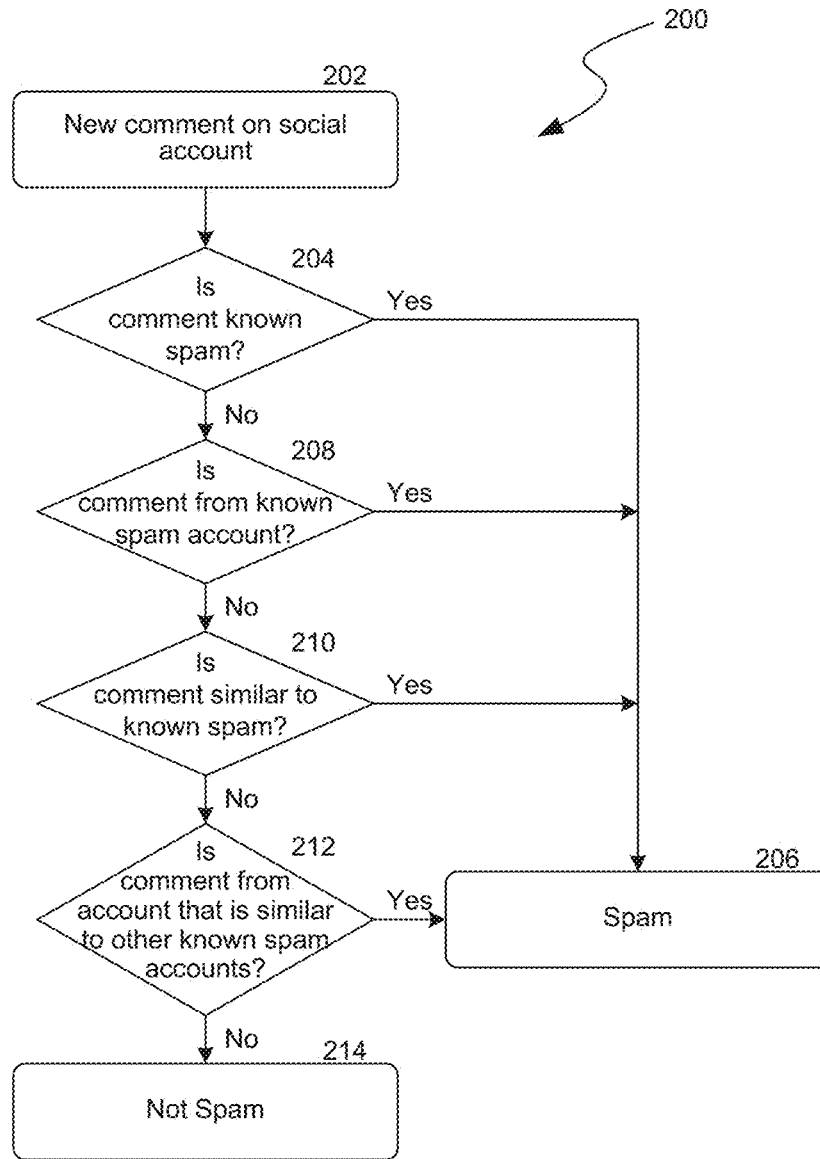


FIG. 2

Accounts Incidents Access Policies App Policies Content Policies Custom Policies Commenter Watchlist Archive


Block users on the Social Media SPAM List across all my accounts

**ADD A USER TO COMMENTERS LIST**

--none selected--

	Blocked Commenters	Watched Commenters
Facebook	2	5
Google+	1	1
YouTube	0	1

Commenters in your Commenter Block List will have their comments automatically removed from your accounts.

**Block List** 






 <b>John Doe</b> Facebook	 <b>Bob Smith</b> Facebook
 <b>Jane Doe</b> Facebook	 <b>Phinneas Bloggs</b> Facebook

FIG. 3

400

**Content Incidents**

Policy Categories: PII, Spam



**Phinneas Bloggs**  
 Comment on 05/02/13 2:41:45 PM

WARNING: Your account is reported to have violated the policies that are considered annoying or insulting Facebook users. System will disable your account within 24 hours if you do not do the reconfirmation. Please confirm your facebook account below. Thanks, The Facebook Security Team Inc., Department 415 P.O. Box 10005 Palo Alto CA 94303

[See less](#)

[View on Facebook](#) | Posted via Web

Detected Action(s):

Show Parent | Actions ▼

- ADD TO COMMENTER WATCH LIST
- ADD TO COMMENTER BLOCK LIST
- POST RESPONSE
- REMOVE FROM FACEBOOK
- IGNORE INCIDENT
- IGNORE ALL LIKE THIS

**FIG. 4**

**CLASSIFYING SOCIAL ENTITIES AND  
APPLYING UNIQUE POLICIES ON SOCIAL  
ENTITIES BASED ON CROWD-SOURCED  
DATA**

CROSS-REFERENCE TO RELATED  
APPLICATION(S)

**[0001]** The present application is a continuation of U.S. application Ser. No. 14/306,098 filed Jun. 16, 2014, which claims the benefit of U.S. Provisional Patent Application No. 61/867,004, filed Aug. 16, 2013, entitled “CLASSIFYING SOCIAL ENTITIES AND APPLYING UNIQUE POLICIES ON SOCIAL ENTITIES BASED ON CROWD-SOURCED DATA,” each of which is incorporated herein in its entirety by reference.

BACKGROUND

**[0002]** Social networking platforms and networks, including, e.g., FACEBOOK, GOOGLE+, LINKEDIN, TWITTER, YOUTUBE, XING, and many others, are commonly employed by millions of users. They have become so pervasive that they are now commonly used even by organizations and other commercial enterprises to advertise as well as communicate with their target audiences on behalf of the organizations. (The social networking platforms named in this application may be registered or common law trademarks of entities unrelated to the applicant. Applicant makes no claim of ownership over any such trademarks.)

**[0003]** A social account (or simply “account”) is an account on one or more social networking platforms (“social network”). An account can be associated with an individual (e.g., a natural person) or it can be associated with an organization (e.g., created by an employee of the organization) for broadcasting information about the organization or its products or services. Creating accounts on these social networks is relatively simple and users can create a “presence” on the social networks for themselves, their products, their brands and/or their initiatives. To present a consistent social identity across social networking platforms, individuals and organizations can create a social account on each social networking platform using similar or identical names, pictures and descriptions. The individual or organization can then broadcast similar or identical content across multiple or even all of their social accounts to reach the audiences or constituents on each social networking platform to which they are connected.

**[0004]** Similarly, an individual or organization can create multiple social accounts on the same social networking platform to represent different products or “personas.” Even though these social accounts may appear to be different, the same person or representative of a company or other entity can control, configure and post content using these accounts.

**[0005]** Individuals or companies attempting to engage in “bad” activity (“fraudulent users”) on a social networking platform can attempt to evade detection by creating multiple social accounts (“fake social accounts”) on this social networking platform that appear to represent non-malicious or real-world identities. Examples of bad activity include, e.g., violating social networking platforms’ rules regarding terms of use, abusive online behavior, violating laws, etc. These malicious individuals or companies can then interact with legitimate social accounts on social networking platforms via one or more of these fake social accounts with malicious

intent, e.g., abusing, bullying, exploiting, harming, infecting (e.g., with malware or viruses), or stealing from the legitimate social accounts.

**[0006]** To operate at a high scale, these fraudulent users may write software programs, commonly called “bots”, that automate the creation of multiple social accounts and the publishing of social content using those accounts. This technique is most often used to spread to legitimate social users spam, malware, or other abusive or malicious content, who then may read or potentially act upon the content.

**[0007]** Social networks generally have programmatic application program interfaces (APIs) that enable third party applications (“social applications” or simply “applications”) to integrate with the social networks and provide additional services for the social networks’ users. Through these APIs, social networks provide additional services for the social network’s users. These APIs provide access to public data within the social network. Access to private data may also be allowed if it is granted by the social account that owns that data.

BRIEF DESCRIPTION OF THE DRAWINGS

**[0008]** FIG. 1 is a block diagram illustrating various components of the disclosed technology.

**[0009]** FIG. 2 is a flow diagram illustrating a routine invoked by the technology in various embodiments to classify postings made to a social network as spam in accordance with the present technology.

**[0010]** FIG. 3 is a user interface diagram illustrating a user interface for enabling users to specify social account participants and social media content rules that can be applied to social data provided by the participants in various embodiments.

**[0011]** FIG. 4 is a user interface diagram illustrating a user interface for enabling users to apply specified content rules to a particular social data entry and/or a commenter in accordance with embodiments of the technology.

DETAILED DESCRIPTION

**[0012]** Technology is disclosed that creates and enforces rules for classifying social data, such as crowd-sourced data, on a social networking platform based on activity observed on that social networking platform or other social networking platforms. In various embodiments, social data comprises social entities and any associated metadata. Social entities are activities by social accounts, attributes associated with social accounts, or social accounts themselves. Examples of social entities can include posts, tweets, account names, account attributes (e.g., email address, phone number, etc.), comments, language used, videos, photos, etc. In various embodiments, the technology observes social data relating to a particular social account associated with a particular social networking platform to compute a classification on an entity observed on that social account, or on a different social account, associated with either the same or different social networking platform. The technology can then employ the computed classification to provide sets of rules and policies that can be applied to the particular social account or multiple social accounts. The rules can identify actions to be taken on social entities, e.g., to report, log, alert, block, remove, and/or remediate the social entities.

**[0013]** Several embodiments of the described technology are described in more detail in reference to the Figures. The computing devices on which the described technology may be implemented may include one or more central processing units, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), storage devices (e.g., disk drives), and network devices (e.g., network interfaces). The memory and storage devices are computer-readable media that may store instructions that implement at least portions of the described technology. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communications link. Various communications links may be used, such as the Internet, a local area network, a wide area network, or a point-to-point dial-up connection. Thus, computer-readable media can comprise computer-readable storage media (e.g., “non-transitory” media) and computer-readable transmission media.

**[0014]** FIG. 1 is a schematic block diagram illustrating an environment 100 in which the technology may operate in various embodiments. The environment can include one or more social networks 102 (individually identified as social network 1 102A, social network 2 102B, and social network 3 102N). Examples of social networks can include, but is not limited to, FACEBOOK, PINTEREST, TWITTER, GOOGLE+, YOUTUBE, LINKEDIN, etc. The social networks 102 may be communicatively coupled via a network 104 (e.g., an intranet, the Internet, etc.) with a social data management system 106 (“system 106”) associated with the technology and one or more client computing devices 107.

**[0015]** In some embodiments, the system 106 includes a: (1) a scanning engine 108, which is composed of one or more processes, that collects social data from a plurality of social networks via the social networks’ APIs; (2) a storage engine 110 that stores social data, e.g., outside of the social network it was observed on; (3) an inference engine 112, which is composed of one or more processes, that analyzes social data in order to build rules to classify new social data; and (4) a classification engine 114, which is composed of one or more processes, that applies the newly inferred rules to new social data while or soon after this new social data is observed. In some embodiments, the scanning engine 108, storage engine 110, inference engine 112 and classification engine 114 may be composed of either a single process or multiple processes that implement a combination of subsets of functionality for each component. The system 106 can further include a user interface for applying rules and policies to a social account or across social accounts leveraging classifications of the system. In operation, the functionality of the engines 108, 110, 112, 114 can be exposed via the user interface for use on a social account or across multiple social accounts.

**[0016]** In some embodiments of the technology, the inference engine 112 builds a list of social accounts that are known to publish social data that are classified as “spam,” e.g., undesirable content. The inference engine 112 evaluates the percentage of overall posts and comments seen from a first social account that were classified as spam, and the number of additional social accounts (e.g., a second social account, a third social account, etc.) on which that spam content was posted. If the percentage of posts classified as spam and/or the number of social accounts the spam content was posted surpass specified (e.g., pre-determined) thresholds, it is determined that there exists an increased likelihood

that the next posting of social data published from the first social account will also be spam.

**[0017]** The inference engine 112 is configured to instruct the classification engine 114 to take one or more actions, and to apply new rules to all social data classified by the system 106 regardless of which social account published the data and/or regardless of which social account the data was published to. In several embodiments, the inference engine 112 can inform the classification engine 114 to (1) consider all social data subsequently posted by the first social account as inappropriate; (2) change sensitivity to spam classification on subsequent social data posted by the first social account; (3) consider subsequent social data that is the same or similar to the social data posted by the first social account to be spam regardless of which other social account posts it (as described in further detail below); and/or (4) consider social data posted by other social accounts (e.g., a second social account, a third social account, etc.) that is the same or similar to the observed first social account as spam (as described in further detail below).

**[0018]** In one embodiment, the classification engine 114 is configured to determine the similarity of multiple social data, either on the same social networking platform or different social networking platforms, by comparing the metadata from those social data and looking for similarities. Such similarities can include, for example, one or more of: (1) lexicographical similarities as in a small Levenshtein distance; (2) chronologically close posting times; (3) use of the same third party publishing application; (4) geo location data that is physically close; (5) use of a similar URL characteristic (e.g., same URL, URLs that redirect to the same final location, URLs that link to the same URL domain or URLs that link to the same content hosted on different URL domains); (6) use of the same or similar images; and/or (7) multiple of these factors used in combination. In instances wherein the classification engine 114 concludes that a set of social data is similar, the classification engine 114 can consider that content equivalent for the purpose of choosing which rules from the inference engine 112 to apply to that content.

**[0019]** In some embodiments, the classification engine 114 can determine the similarity of multiple social accounts, either on the same social networking platform or different social networking platforms, e.g., by comparing the metadata from those accounts and looking for similarities. Such account similarities can include one or more of: (1) the profile images are identical or very similar; (2) the name, handle, or username used on the account is the same or very similar; (3) the date the accounts were created is chronologically close; (4) the geo location of the account or the posts on the account is the same or physical close; (5) identical content is published to each account at the same time; (6) the same third party applications are used to publish content to each account; and/or (7) multiple of these factors used in combination. In instances wherein the classification engine 114 concludes that a set of social accounts are similar, the classification engine 114 can consider those accounts to be equivalent for the purpose of choosing which rules from the inference engine 112 to apply to the content on and from those accounts.

**[0020]** FIG. 2 is a flow diagram illustrating a routine 200 invoked by the technology in various embodiments to classify postings made to a social network as spam in accordance with the present technology. The routine begins at

block **202** when a new comment on a first social account is detected by the classification engine **114**. At decision block **204**, the routine **200** determines whether the new comment is known spam data. As an example, if the comment is recognized as known spam, the classification engine **114** can mark the comment as spam at block **206**. If the comment is not recognized as spam, the routine **200** determines whether the comment is from a known spam social account at decision block **208**. If the routine **200** determines that the first social account posting the comment is a known spam account, the comment can be flagged as spam at block **206**. If the first social account posting the comment is not a known spam account, the routine **200** continues at decision block **210** when the routine determines if the comment includes content data similar to spam data. If the routine **200** determines the posted comment is similar to spam, the comment is marked as spam at block **206**. If the content data of the posted comment is not similar to known spam data, the routine continues at decision block **212**. At decision block **212**, the routine **200** determines whether the first social account posting the comment is similar to another known spam account. If the first social account is similar to a known spam account, the comment is flagged as spam at block **206**. If the first social account posting the comment is not similar to a known spam account, the routine determines that the comment is not spam at block **214**. The routine **200** may be invoked at regular time intervals, random time intervals, upon detection that content (e.g., a comment, posting, etc.) has been uploaded to a social network or is about to be uploaded to a social network, etc. In various embodiments, the routine **200** invokes APIs relating to the social networks or APIs relating to content publication tools. In various embodiments, the technology and the routine **200** may be invoked by social networks or content publication tools, e.g., as part of an extension provided by the social networks or publication tools.

**[0021]** Those skilled in the art will appreciate that the logic illustrated in FIG. 2 and described above may be altered in a variety of ways. For example, the order of the logic may be rearranged, substeps may be performed in parallel, illustrated logic may be omitted, other logic may be included, etc.

**[0022]** Referring back to FIG. 1, the inference engine **112** can also be configured to evaluate the frequency of profanity and/or hateful speech (e.g., abusive speech) present in the social data posted by a first social account to determine if the first social account should be blocked from further postings and/or posting any social data on any other social account (e.g., a second social account, a third social account, etc.). Once the inference engine **112** makes this determination, the inference engine **112** can be configured to inform the classification engine **114** to take one or more actions, and to apply rules to all social data classified by the system **106** as containing profanity and/or hateful speech, as described above with respect to the identification and classification of spam data. For example, the classification engine **114** can classify social data as abusive speech if the social data meets one or more of the following criteria: (1) consider all social data subsequently posted by the first social account as inappropriate; (2) change sensitivity to abusive speech classification on subsequent social data posted by the first social account; (3) consider subsequent social data that is the same or similar to the social data posted by the first social account to contain abusive speech regardless of which other social

account posts it; and/or (4) consider social data posted by other social accounts (e.g., a second social account, a third social account, etc.) that is the same or similar to the observed first social account as abusive speech.

**[0023]** In addition to classifying and applying rules to spam data and abusive speech, the present technology can also be configured to address and classify other forms of social data including, but not limited to, malware, bullying, fraudulent activity, potentially violent speech, sexually predatory activity, etc. For example, in other embodiments of the present technology, the inference engine **112** evaluates linked content (e.g., by URL links) in the social data. If the linked content is determined or classified to be malware, then the inference engine **112** determines that any subsequent social data that also links to the linked malware content is malicious. If the percentage of malicious social data posted by a social account reaches a specified threshold, then the inference engine **112** infers that any subsequent social data or URL links posted from that social account will be malicious and updates instructions to the classification engine **114** accordingly.

**[0024]** In further embodiments of the present technology, the inference engine **112** evaluates activity of other applications (e.g., social publishing apps such as Hootsuite™, Sprinklr®, Spredfast®, or other publisher that posts content on behalf of social network users) that use a social networking platform's APIs. If the social data posted by the application is classified as spam, malicious, abusive or in any other way unwanted (e.g., unwanted data) by users of the social networking platform, or if the social data posted by the application is automated so as to occur at a higher frequency than would be possible if generated by actual user activity, the inference engine **112** can be configured to determine that all subsequent social data posted to any social account from that application is classified as unwanted data and updates instructions to the classification engine **114** accordingly.

**[0025]** FIGS. 3 and 4 are user interface diagrams illustrating user interfaces **300** and **400**, respectively, for enabling users to interact with various embodiments of the environment **100** such as to specify and/or control a set of rules and/or policies to apply to one or more social accounts. In various embodiment, the user interfaces **300**, **400** can be presented in a Web browser, standalone application, applet, etc. The social accounts can be from any social networking platform and are defined as a set of social accounts separately from the rules that combine to form a social data classification policy as described above. For example, FIG. 3 illustrates the user interface **300** in which a commenters list can be controlled by a user of the system for a particular social account or across a set of social accounts operated by the user. In one embodiment, a user of the system can select a social media participant or a participants' social media account (e.g., social media presence) and define one or more rules to associate with that participant (e.g., blocked, watched, allowed, etc.). If the system classifies a social data entry as a content incident (e.g., flags a comment as spam, abusive speech, etc.), a user or social media account owner can interface with the system at user interface **400** (FIG. 4) to specify which action, if any, the user would like to assign to flagged social data entry. For example, the user can take actions with regard to the commentator or social media participant (e.g., add commenter to a watch list, add commenter to a block list), or can take action with respect to a



particular comment or social data entry (e.g., post response, remove response, ignore incident, ignore all similar future incidents, etc.).

**[0026]** In some embodiments, the inference engine 112 can adopt to a user's input to modify an operation. For example, if a user decides to put a commenter's social account on a watch list or block list, the inference engine 112 can be configured to alter rules that are applied to similar subsequent social data classifications made by the system. In another example, if the user elects to reassign a system classification of a social data entry and "ignore" the social data entry, the inference engine 112 can be configured to alter rules that are applied to similar subsequent social data classifications made by the system.

**[0027]** Accordingly, various embodiments of the technology can be adapted to involve or incorporate user interaction or input at any step in the process of scanning, classifying, storing, inferring and/or updating classification rules. In other embodiments, the technology may be automated.

**[0028]** In several embodiments, of the system can be configured to scan and classify social media data and content on a continuous and/or intermittent basis (e.g., over minutes, hours or days). In some embodiments, the social media content may be evaluated, classified, and acted upon in real-time as the scanning engine 108 retrieves the social data from the social networks' APIs. In other embodiments, the social media content can be collected and stored (e.g., by the scanning engine 108 and storage engine 110, respectively), and the content can later be evaluated, classified and acted upon, for example, in an off-line setting.

**[0029]** Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims. Accordingly, the invention is not limited except as by the appended claims.

I/We claim:

1. A system, comprising:

a central processing unit configured to:

collect social data representing content shared on a first social network;

determine a new rule based on a subset of the social data that corresponds to a first social account from the social data, wherein determining the new rule includes:

classifying posts in the subset as undesirable content, calculating a percentage of the undesirable content within the subset, and

determining a new rule for the first social account based on the percentage of the undesirable content, wherein the new rule is determined from a set of social data content rules;

calculate a similarity between a second social account and the first social account, between social data corresponding to the second account and the subset or the undesirable content corresponding to the first social account, or a combination thereof;

apply an action corresponding to the new rule, wherein the action is applied with respect to the second social account based on the similarity; and

a memory, operably coupled to the central processing unit, the memory configured to store the social data.

2. The system of claim 1 wherein the central processing unit is further configured to:

calculate a likelihood of a next published post on the first social account being the undesirable content, wherein the likelihood is calculated based on the percentage of the undesirable content among the posts of the first social account; and

apply the action with respect to the second social account according to the likelihood corresponding to the first social account and the similarity between a new post on the second social account and the next published post on the first social account.

3. The system of claim 1 wherein the social data includes at least one of spam, abusive, speech, and a link to malicious content.

4. The system of claim 1 wherein the action includes blocking a specific social data associated with the social data subset, wherein the specific social data is blocked from posting to or posting by the first social account.

5. The system of claim 4 wherein the central processing unit is further configured to block a specific social data associated with the social data subset based on metadata, wherein the specific social data is blocked from posting to or posting by a second social account.

6. The system of claim 1 wherein the central processing unit is configured to:

determine the new rule for application on a second social network different than the first social network; and

the action includes blocking a social entity attempting to post a specific social data associated with the social data subset, wherein the social entity is blocked from posting the specific social data, additional social data, or a combination thereof on the second social network.

7. A non-transitory computer readable storage medium storing computer-executable instructions, comprising:

instructions for monitoring social account activities operating on one or more social networking platforms;

instructions for classifying one or more social account activities on the first social account, wherein the classification is based on detection of unwanted social data associated with content of the one or more social account activities;

instructions for calculating a percentage of undesirable content among the one or more social account activities;

instructions for calculating a similarity between a second social account and the first social account, between social data corresponding to the second social account and the undesirable content corresponding to the first social account, or a combination thereof; and

instructions for applying an action with respect to the second social account based on the percentage of the undesirable content and the similarity.

8. The non-transitory computer-readable storage medium of claim 7 wherein the unwanted social data includes at least one of spam and abusive speech.

9. The non-transitory computer-readable storage medium of claim 7, further comprising:

instructions for determining the rule based on the classification of a first new social account activity by the first social account; and

instructions for enforcing the rule on a second new social account activity.

**10.** The non-transitory computer-readable storage medium of claim **9** wherein the second new social account activity resides on a different social networking platform than that of the first new social account activity.

**11.** A method, comprising:

scanning social data residing within at least one social network;

determining from at least a subset of the stored social data at least one rule, wherein:

the subset of the stored social data corresponds to a first social account,

the rule is determined based on the content of the subset, wherein determining the rule includes:

classifying posts in the social data subset as undesirable content,

calculating a percentage of the undesirable content among the social data subset corresponding to the first social account,

determining the new rule according to the first social account based on the percentage of the undesirable content;

calculating a similarity between a second social account and the first social account, between social data corresponding to the second social account and the undesirable content corresponding to the first social account, or a combination thereof; and

applying the rule to a first new social account activity associated with the second social account based on the similarity.

**12.** The method of claim **11** wherein calculating the similarity includes calculating a similarity between the second social account and the first social account based on analyzing profile images, names, handles, usernames, creation dates, associated geo locations, publication of identical

contents, usage of overlapping third-party publishing applications, or a combination thereof.

**13.** The method of claim **11** wherein calculating the similarity includes calculating the similarity between the social data of the second social account and the undesirable content corresponding to the first social account based on one or more of a lexicographical similarity, a chronologically close posting time, a same publishing application, a physically close geo location data, a similar URL characteristic, and an image similarity.

**14.** The method of claim **11** wherein the rule comprises one or more actions on the first new social account activity selected from the group consisting of reporting, logging, alerting, blocking, removing and remediating.

**15.** The method of claim **11** wherein the applying occurs in real-time.

**16.** The method of claim **11** wherein the rule is formed when the stored social data includes a type of content selected from the group consisting of malware, bullying, fraud, violence and sexually predatory.

**17.** The method of claim **11** wherein the rule is altered in accordance with user input.

**18.** The method of claim **11** wherein the rule is applied to the first new social on a different social network different than the at least one social network.

**19.** The method of claim **11** further comprising changing sensitivity for classifying subsequent posts for the first social account.

**20.** The method of claim **11** wherein determining the new rule includes determining an increased likelihood that subsequent postings from the first social account are classified as the undesirable content.

\* \* \* \* \*