




(86) Date de dépôt PCT/PCT Filing Date: 2018/06/20
(87) Date publication PCT/PCT Publication Date: 2018/12/27
(45) Date de délivrance/Issue Date: 2023/10/03
(85) Entrée phase nationale/National Entry: 2019/12/20
(86) N° demande PCT/PCT Application No.: US 2018/038639
(87) N° publication PCT/PCT Publication No.: 2018/237098
(30) Priorités/Priorities: 2017/06/20 (US62/522,644);
2017/07/18 (US62/534,172)

(51) Cl.Int./Int.Cl. *G06Q 30/0201* (2023.01),
G06Q 50/00 (2012.01)
(72) Inventeurs/Inventors:
BARASH, VLADIMIR D., US;
KELLY, JOHN W., US
(73) Propriétaire/Owner:
GRAPHIKA, INC., US
(74) Agent: MACRAE & CO.

(54) Titre : PROCEDES ET SYSTEMES D'IDENTIFICATION DE MARQUEURS D'ACTIVITE COORDONNEE DANS DES MOUVEMENTS DE MEDIAS SOCIAUX
(54) Title: METHODS AND SYSTEMS FOR IDENTIFYING MARKERS OF COORDINATED ACTIVITY IN SOCIAL MEDIA MOVEMENTS



(57) **Abstrégé/Abstract:**

Methods and systems generally include determining coordinated activity in social media movements on a social media channel. The method includes identifying a plurality of markers of coordinated activity through analysis of campaign signals from the social media movements. The method includes configuring a data structure of the plurality of markers for a social media campaign on a social media channel. The plurality of markers includes a network dimension for representing how accounts are connected, a temporal dimension for representing patterns of messages over time, and a semantic dimension for representing a diversity of topics and meanings of the social media movements. The method includes analyzing the campaign signals indicative of the coordinate activity of the social media movements in the social media campaign including determining users within the social media campaign, determining clusters of users that make up the social media campaign, and determining relationships between the users participating in the social media movements, and determining propagation patterns across clusters of users, of the social media campaign.



(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau

(43) International Publication Date
27 December 2018 (27.12.2018)



(10) International Publication Number
WO 2018/237098 A1

(51) International Patent Classification:

G06F 17/30 (2006.01) *G06Q 50/00* (2006.01)

(21) International Application Number:

PCT/US2018/038639

(22) International Filing Date:

20 June 2018 (20.06.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/522,644 20 June 2017 (20.06.2017) US
62/534,172 18 July 2017 (18.07.2017) US

(71) Applicant: **GRAPHIKA, INC.** [US/US]; 116 West 23rd
Street, New York City, New York 10011 (US).

(72) Inventors: **BARASH, Vladimir D.**; 7 Lowden Ave, Apt.
1, Somerville, Massachusetts 02144 (US). **KELLY, John
W.**; 365 W 123rd Street, New York City, New York 10027
(US).

(74) Agent: **DOBROWITSKY, Joshua B.**; FMA, LLP, PO
BOX 2534, Birmingham, Michigan 48012 (US).

(81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) Title: METHODS AND SYSTEMS FOR IDENTIFYING MARKERS OF COORDINATED ACTIVITY IN SOCIAL MEDIA
MOVEMENTS

200




Fig. 2

(57) Abstract: Methods and systems generally include determining coordinated activity in social media movements on a social media channel. The method includes identifying a plurality of markers of coordinated activity through analysis of campaign signals from the social media movements. The method includes configuring a data structure of the plurality of markers for a social media campaign on a social media channel. The plurality of markers includes a network dimension for representing how accounts are connected, a temporal dimension for representing patterns of messages over time, and a semantic dimension for representing a diversity of topics and meanings of the social media movements. The method includes analyzing the campaign signals indicative of the coordinate activity of the social media movements in the social media campaign including determining users within the social media campaign, determining clusters of users that make up the social media campaign, and determining relationships between the users participating in the social media

[Continued on next page]



WO 2018/237098 A1

WO 2018/237098 A1

SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

DEMANDE OU BREVET VOLUMINEUX

LA PRÉSENTE PARTIE DE CETTE DEMANDE OU CE BREVET COMPREND PLUS D'UN TOME.

CECI EST LE TOME 1 DE 2
CONTENANT LES PAGES 1 À 95

NOTE : Pour les tomes additionels, veuillez contacter le Bureau canadien des brevets

JUMBO APPLICATIONS/PATENTS

THIS SECTION OF THE APPLICATION/PATENT CONTAINS MORE THAN ONE VOLUME

THIS IS VOLUME 1 OF 2
CONTAINING PAGES 1 TO 95

NOTE: For additional volumes, please contact the Canadian Patent Office

NOM DU FICHIER / FILE NAME :

NOTE POUR LE TOME / VOLUME NOTE:

METHODS AND SYSTEMS FOR IDENTIFYING MARKERS OF COORDINATED ACTIVITY IN SOCIAL MEDIA MOVEMENTS

[0001] Continue to paragraph [0002].

[0002] Continue to paragraph [0003].

BACKGROUND

1. Field

[0003] The present disclosure relates to methods for classifying at least one contagious phenomenon propagating on a network.

2. Description of the Related Art

[0004] Internet-based technologies, and the manifold genres of interaction they afford, are re-architecting public and private communications alike and thus altering the relationships between all manner of social actors, from individuals, to organizations, to mass media institutions. Internet technologies, have enabled shifts in methods and practices of interpersonal, communication. Many-to-many and social scale-spanning internet communications technologies are eliminating the channel-segregation that previously reinforced the independence of classes of actors at these levels of scale, enabling (or more accurately in many cases, forcing) them to represent themselves to one another via a common medium, and increasingly in ways that are universally visible, searchable and persistent.

[0005] Online readers typically navigate hyperlinked chains of related stories, bouncing between numerous websites in a hypertext network, returning periodically to favored starting points to pick up new trails. Hyperlinks result from a combination of choices, from those made by individual,

autonomous authors to those made programmatically by designed systems, such as permalinks, site navigation, embedded advertising, tracking services, and the like. Human authors practice the same kind of information selectivity online that they do offline, i.e., what authors (including those representing organizations) write about and link to reflects somewhat stable interests, attitudes, and social/organizational relationships. The structure of the network formed by these hyperlinks is a product of these choices, and thus large-scale regularities in choices will be evident in macro-level structure. This structure will thus bear the mark of individual preferences and characteristics of designed systems and allows a kind of “flow map” of how the Internet channels attention to online resources. Discriminating among types of links, and the ability to select categories of those which represent author choices, allows structural analytics to discover similarities among authors. Errors, randomness, or noise in linking at the individual level has local, independent causes, and does not bias large-scale macro patterns.

[0006] Thus, in order to understand and leverage the online information ecosystem, there remains a need for systems and methods for structural analytics aimed at identifying clusters of online readers and influential authors, discovering how they drive traffic to particular online resources, and leveraging that knowledge across various applications ranging from targeted advertising and communication to expert identification, and the like. This need includes a need for understanding the role of structures and similarities among authors and readers in situations involving phenomena that follow a pattern of contagion, i.e., where an item of interest, such as a news story, a political topic, a product, an item of entertainment content, or the like, initiates with a single point or a small group, then spreads and grows through the network. Predicting the pattern of spread or contagion, the parties who will take interest in, be involved with, or be influenced by a particular item, and the like may have great value in a range of applications; accordingly, a need exists for methods and systems that assist in or enable such prediction of the behavior of contagious phenomena.

SUMMARY

[0007] In embodiments, methods and systems generally include determining coordinated activity in social media movements on a social media channel. The method includes identifying a plurality of markers of coordinated activity through analysis of campaign signals from the social media movements. The method includes configuring a data structure of the plurality of markers for a social media campaign on a social media channel. The plurality of markers includes a network dimension for representing how accounts are connected, a temporal dimension for representing patterns of messages over time, and a semantic dimension for representing a diversity of topics and meanings of the social media movements. The method also includes analyzing the campaign signals indicative of the coordinate activity of the social

media movements in the social media campaign including determining users within the social media campaign, determining clusters of users that make up the social media campaign and determining relationships between the users participating in the social media movements, and determining propagation patterns across clusters of users of the social media campaign.

[0008] In embodiments, identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is concentrated in the clusters of users. In embodiments, the coordinated activity of the social media campaign is determined from user actions within the social media movements in the social media campaign. In embodiments, identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is distributed among the clusters of users. In embodiments, the plurality of markers includes a day peakedness marker that indicates a percentage of the coordinated activity of the social media campaign that take place on a day identified as most active of the social media campaign. In embodiments, the plurality of markers includes a commitment signal that is computed by averaging a number of subsequent participation actions for each of a plurality of participants in the coordinated activity of the social media campaign. In embodiments, the plurality of markers includes a post regularity commitment signal that represents a deviation of commitment to participation by a user from natural human attention patterns. In embodiments, identifying the plurality of markers includes determining a semantic diversity score for the coordinated activity of the social media campaign by assigning messages in the campaign to topics and calculating a diversity of the topics on a topic distance scale that facilitates determining the semantic diversity score. In embodiments, identifying the plurality of markers includes computing temporal alignment of campaign-related actions for users in the campaign by comparing temporal sequences of campaign-related actions. In embodiments, identifying the plurality of markers includes computing semantic diversity over time to identify co-occurring topics in the social media campaign, wherein a relatively small value of the semantic diversity score is configured to be indicative of fabricated campaigns, wherein a relatively large value of the semantic diversity score is configured to be indicative of spambots, and wherein a semantic diversity score having a value in-between is indicative of normal human activity.

[0009] In embodiments, methods and systems generally include a computer system for determining coordinated activity in social media movements on a social media channel. The system includes a user interface that configures a social media campaign on one or more social media channels and that communicates via a network. The system includes a computing device that identifies a plurality of markers of coordinated activity through analysis of campaign signals from the social media movements and that configures one or more data structures containing the

plurality of markers for the social media campaign on one or more social media channels. The plurality of markers includes a network dimension for representing how accounts are connected, a temporal dimension for representing patterns of messages over time, and a semantic dimension for representing a diversity of topics and meanings of the social media movements. The analysis of the campaign signals indicative of the coordinated activity of the social media movements in the social media campaign includes determining users within the social media campaign, determining clusters of users that make up the social media campaign and determining relationships between the users participating in the social media movements, and determining propagation patterns across clusters of users of the social media campaign. The system includes a storage system that stores one or more of the data structures containing the plurality of markers for the social media campaign on one or more of the social media channels. The system includes

[0010] a processing system that executes computer-readable instructions that cause the processing system to: receive a request from an external system about the coordinated activity of the campaign signals from the social media movements; retrieve at least a portion of one or more data structures containing the plurality of markers for the social media campaign on one or more of the social media channels; and transmit contents of at least a portion of the analysis to the user interface that displays at least a portion of the plurality of markers indicative one of coordinated activity and normal human activity

[0011] In embodiments, identifying the plurality of markers through analysis of campaign signals includes evaluating a degree to which the coordinated activity of the social media campaign is concentrated in the clusters of users. In embodiments, the coordinated activity of the social media campaign is determined from user actions within the social media movements in the social media campaign. The coordinated activity includes a relatively large number of accounts on one or more of the social media channels controlled by a relatively small number of coordinated entities resulting in a relative lack of diversity of similar accounts on one or more social media channels controlled by uncoordinated users. In embodiments, identifying the plurality of markers through analysis of campaign signals includes evaluating a degree to which the coordinated activity of the social media campaign is distributed among the clusters of users.

[0012] In embodiments, the plurality of markers includes a day peakedness marker that indicates a percentage of the coordinated activity of the social media campaign that take place on a day identified as most active of the social media campaign. In embodiments, the plurality of indicators includes a commitment signal that is computed by averaging a number of subsequent participation actions for each of a plurality of participants in the coordinated activity of the social media campaign. In embodiments, the plurality of indicators includes a post regularity

commitment signal that represents a deviation of commitment to participation by a user from natural human attention patterns. In embodiments, identifying the plurality of markers through analysis of campaign signals includes determining a semantic diversity score for the coordinated activity of the social media campaign. Determining a semantic diversity score includes assigning messages in the campaign to topics and calculating a diversity of the topics on a topic distance scale that facilitates determining the semantic diversity score. In embodiments, identifying the plurality of markers through analysis of campaign signals includes computing temporal alignment of campaign-related actions for users in the campaign by comparing temporal sequences of campaign-related actions. In embodiments, identifying the plurality of markers through analysis of campaign signals includes computing semantic diversity over time to identify co-occurring topics in the social media campaign. A relatively small value of the semantic diversity score is configured to be indicative of fabricated campaigns, a relatively large value of the semantic diversity score is configured to be indicative of spambots, and a semantic diversity score having a value in-between is indicative of normal human activity.

[0013] In an aspect of the disclosure, methods and systems are provided that allow characterization of structures and features of networks, such as online networks of creators and consumers of items of content, in turn enabling prediction course of action of actors in such networks and the flow of items, such as items of content, through such networks, including the growth and spreading of contagious phenomena.

[0014] In an aspect of the disclosure, a computer-readable storage medium with an executable program stored thereon, wherein the program instructs a processor to perform the steps of attentive clustering and analysis; may include constructing an online author network, wherein constructing the online author network includes selecting a set of source nodes (S), a set of outlink targets (T) from at least one selected type of hyperlink, and a set of edges (E) between S and T defined by the at least one selected type of hyperlink from S to T during a specified time period; deriving a set of nodes, T', by any one of or combination of a.) normalizing nodes in T, optionally to a selected level of abstraction, b.) using lists of target nodes for exclusion ("blacklists"), and c.) using lists of target nodes for inclusion ("whitelists"); transforming the online author network into a matrix of source nodes in S linked to targets in T'; partitioning the online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and/or at least one set of outlink targets with a similar citation profile to form an outlink bundle; and optionally, generating a graphical representation of attentive clusters and/or outlink bundles in the network to enable interpretation of network features and behavior and calculation of comparative statistical measures across the attentive clusters and outlink bundles; wherein at least one element of the graphical representation depicts a measure of an extent of a type of activity within the

network; and measuring frequencies of links between attentive clusters and outlink bundles enabling identification and measurement of large-scale regularities in the distribution of attention by online authors across sources of information. The element of the graphical representation may use at least one of size, thickness, color and pattern to depict a type of activity. Attentive clusters and their constituent nodes may be differentiated in the graphical representation by at least one of a color (including hue, intensity and saturation), a shape (including 2D or 3D representations), a geometric arrangement, a shading, a transparency and a size. The size of the object representing the clustered nodes in the graphical representation may correlate with a metric. The nodes, targets, and edges may be collected from public and private sources of information. Constructing the matrix may include applying at least one threshold parameter from the group consisting of: maxnodes, targetmax, nodemin, targetmin, maxlinks, and linkmin. Constructing the matrix may include applying a minimum threshold for the number of included nodes that must link to a target to qualify it for inclusion in the matrix. Constructing the matrix may include applying a minimum threshold for the number of included targets that must link to a node to qualify it for inclusion in the matrix. The matrix may be a graph matrix. The method may further include applying any lists specifying inclusion or exclusion of particular nodes.

[0015] It should be understood that, except where context prevents, the term “author,” as used herein, should be understood to encompass human and non-human creators and editors of content (including, without limitation, text, images, video, tweets, animations, multimedia and any combinations or other types of content and including, without limitation, original content, derivative works, commentary, analysis, and other genres of content) that can be consumed (e.g., read or viewed) by others, such as readers or viewers in a network.

[0016] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and collecting clickstream data for the source nodes of the attentive cluster.

[0017] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and collecting clickstream data for the target nodes of the outlink bundle.

[0018] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink

targets with a similar citation profile to form an outlink bundle, and collecting survey data for the source nodes of the attentive cluster.

[0019] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and collecting survey data for the target nodes of the outlink bundle.

[0020] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and collecting geo-location data for the source nodes of the attentive cluster.

[0021] In an aspect of the disclosure, a method of using attentive clustering to steer a further data collection process may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and collecting geo-location data for the target nodes of the outlink bundle.

[0022] In an aspect of the disclosure, a method of metadata tag analysis to facilitate interpretation of an attentive cluster may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, collecting a metadata tag associated with the source nodes in the attentive cluster, and performing a differential frequency analysis on the metadata tags that are associated with the attentive cluster. The method may further include sorting cluster focus scores on a plurality of the metadata tags.

[0023] In an aspect of the disclosure, a method of metadata tag analysis to facilitate interpretation of an attentive cluster may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, collecting a metadata tag associated with the source nodes in the attentive cluster, and performing a differential frequency analysis on the metadata tags that are associated with the outlink bundle. The method may further include sorting cluster focus scores on a plurality of the metadata tags.

[0024] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, forming a density matrix of the attentive cluster and the outlink bundle, determining where there is a higher

density in the density matrix than chance would predict, and identifying patterns of influence of a block of web sites on a block of authors by analyzing the higher density area of the density matrix.

[0025] In an aspect of the disclosure, a method of macro measurement of link density may include constructing an online author network, wherein constructing the online author network comprises selecting a set of source nodes (S), a set of outlink targets (T), and a set of edges (E) between S and T defined by the at least one selected type of hyperlink from S to T during a specified time period, deriving a set of nodes, T', by normalizing nodes in T, transforming the online author network into a matrix of source nodes in S linked to targets in T', and collapsing the matrix to aggregate link measures among clusters of sources and clusters of targets. The aggregated link measure may be at least one of a count of the number of nodes in source cluster S linking to any member of target set T, a density calculated by dividing counts by the product of the number of members in S and the number of members in T; and a standard score that is a standardized measure of the deviation from random chance for counts across each source node-outlink target crossing in the density matrix.

[0026] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and associating the attentive cluster with a real world group of people.

[0027] In an aspect of the disclosure, a method of multi-layer attentive clustering may include partitioning a multi-layered social segmentation into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and monitoring at least one of the attentive cluster and the outlink bundle on at least one layer of the social segmentation. The social segmentation may be an online social media author network. Monitoring may be tracking the growth of an attentive cluster over time. The method may further include examining a source node associated with a specific player in the attentive cluster in order to determine a characteristic. The monitoring may be used to identify a group of people who are susceptible to a message and track downstream activities in response to the message.

[0028] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and analyzing the attentive cluster over time to depict changes in a linking pattern of the attentive cluster over a time period. The outlink bundle may be a list of semantic markers. The semantic marker may be at least one of a text element, a post, a tweet, an online content, and a metadata

tag. Analyzing may involve tracking a semantic marker or set of semantic markers across one or more attentive clusters within the online author network.

[0029] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and calculating a set of cluster focus index (CFI) scores for the attentive cluster, wherein the CFI represents the degree to which a particular outlink target is disproportionately cited by members of a particular attentive cluster as compared to the average citation frequency for all nodes in S. At least one source node may be a high attention source node. The method may further include automatically placing an advertisement at the particular outlink target.

[0030] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and generating a graphical representation of attentive clusters and/or outlink bundles in the network to enable interpretation of network features and behavior and calculation of comparative statistical measures across the attentive clusters and outlink bundles, wherein at least one element of the graphical representation depicts a measure of an extent of a type of activity within the network. The method may further include further segmenting the network using at least one of a text, an item of online content, a link, and an object. The source node in the graphical representation may be represented by an individual dot. The size of the dot may be determined based on the number of other source nodes that link to it.

[0031] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, calculating a set of cluster focus index (CFI) scores (CFI) for the attentive cluster, wherein the CFI represents the degree to which a particular outlink target is disproportionately cited by at least one source node of a particular attentive cluster, and generating a graphical representation of attentive clusters and/or outlink bundles in the network, wherein at least one element of the graphical representation depicts a measure of an extent of a type of activity within the network, wherein the higher the CFI score, the higher the outlink target appears along at least one axis of the graphical representation.

[0032] In an aspect of the disclosure, a method of attentive clustering may include defining a semantic bundle, searching a plurality of candidate nodes for items in the bundle in order to generate a relevance metric for use in selecting high-relevance online authors, partitioning the online author network into at least one set of source nodes with a similar linking history to form

an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and calculating metrics with across clusters for items in the semantic bundle.

[0033] In an aspect of the disclosure, a method may include partitioning an online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and generating a graphical representation of link targets, semantic events, and node-associated metadata scattered in an x-y coordinate space, wherein the dimensions of the graph are custom-defined using sets of attentive clusters grouped to represent substantive dimensions of interest for a particular analysis.

[0034] In an aspect, a computerized search method may include presenting, to a user, a computer interface for specifying one or more search terms for a search query, presenting at least one selectable item corresponding to at least one of an M score and a CFI score filter for the search query, generating an amended search query based on a selected item, and performing a search using the amended search query. The search may be of the Internet. The search may be of a document-corpus. The search may be of a CFI-filtered set of clusters within an online network. The search may be of a set of nodes having an M score greater than a threshold.

[0035] CFI may represent the degree to which an event, characteristic or behavior disproportionately occurs in a particular cluster, or a particular cluster, relative to a network, preferentially manifests an event, characteristic or behavior. M score may be calculated using the formula $M \text{ score} = \text{count}(\alpha) + \text{CFI}(1 - \alpha)$ [normalized 1 to 10], where count is the overall number of members on a cluster focus map that have engaged with a target.

[0036] In an aspect, a computerized search method may include presenting, to a user, a computer interface for specifying one or more search terms for a search query, presenting, to the user, a computer interface for selecting content to search with the search terms, wherein the content is taken from an online creator network partitioned into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle, and performing a search of the selected content using the search query.

[0037] In an aspect, a method to iteratively reduce the scale of a network to its most influential core communities and obtain a sub-graph of maximally connected sub-actors may include assigning a variable, K_{core} , to each individual member of the network, where K_{core} relates to a minimum connectedness based on the number of other nodes in the network to which the individual is connected, removing inactive individuals and individuals with few followers from the network, temporarily removing certain individuals with a large number of followers for later re-joining, restricting the remaining individuals iteratively by removing individuals with the

lowest K_{corr} values first, then removing individuals with the next highest K_{corr} values until a threshold is reached, wherein the threshold is at least one of a number of individuals removed, a number of individuals remaining, and a K_{corr} value, and re-joining the temporarily removed individuals.

[0038] In an aspect, a self-service tool to construct a social media map may include an automated process (e.g., bot) that harvests data (e.g., nodes) and maps the data to one or more clusters/segments, a processor that provides cluster/segment labels and CFI scores for the clusters/segments, and an interface that enables user browsing of clusters/segments and the map, tagging nodes, and re-grouping/re-labeling of clusters/segments. The automated process may also be capable of: automatically refreshing the social media map based on using a relevance score for nodes in the map, positively or negatively weighting at least one cluster based on a CFI score calculation to include positively weighted nodes and exclude negatively weighted nodes from the map, filtering out unwanted nodes, obligatorily including nodes that were not clustered in a first version of the social media map, crowd-sourced information regarding nodes and/or links that drives nodes to bundles, processing social media map usage data for trends/indicators, wherein the usage data relates to one or more of what is ignored, what is further explored, what is used, how are clusters grouped, what name/label is assigned to a cluster, what color is used for a cluster, what order/position is the cluster placed in a report and wherein nodes preferentially interacted with are weighted more heavily, and user-contributed data as metadata for the social media map.

[0039] In an aspect, a method of strategic messaging may include generating a list of targets in a network/cluster/segment, filtering the list by a criteria to limit whom to message in the network/cluster/segment in order to maximize the impact of the message on the cluster/segment, wherein the filter is at least one of CFI score, M score, number of followers, following status, follower status, number of mentions/re-tweets, number of distinct mentions, status of exposure to content, status of exposure to content that has already peaked, footprint, and number of tweets/publication frequency, and ranking the list by the filtered criteria.

[0040] In an aspect, a method of strategic network building may include generating a list of targets in a network/cluster/segment, wherein the list is generated using at least one of CFI, M score, # of followers, mentions/re-tweets, distinct mentions, and number of tweets, and following the targets.

[0041] In an aspect, a method of calculating M score may include calculating a cluster focus index score based on a degree to which a target disproportionately occurs in a particular cluster, or a particular cluster, relative to a network, preferentially engages with a target, determining an overall number of members of the cluster or network that have engaged with that target, and calculating an M score based on the formula: count plus CFI, wherein count is the overall number of members of the cluster that have engaged with that target.

[0042] In an aspect, an M score filter for a list of targets may include taking a cluster focus index (CFI) score based on a degree to which a target disproportionately occurs in a particular cluster, or a particular cluster, relative to a network, preferentially engages with a target, and providing a slider to indicate an M score, wherein the M score is based on the formula: $\text{count}(\alpha) + \text{CFI}(1 - \alpha)$, wherein count is the overall number of members of the cluster or network that have engaged with that target, and wherein the slider is used to indicate the value of alpha between 0 and 1.

[0043] In an aspect, a method of strategic ad placement may include generating a list of targets in a network/cluster/segment representing linkages in a social media environment, filtering the list by a criteria to limit the targets in order to maximize the impact of the ad on the network/cluster/segment, wherein the filter is at least one of CFI score and M score, ranking the list by the filtered criteria, and providing an interface to launch an ad campaign to place ads directly from the environment representing the linkages to the target/website. Ad placement may be done via integration with various products, such as Twitter™ sponsored tweets, Facebook™ ad exchange, Google™ AdSense/Adwords, and third party online ad networks. The method may further include tracking interaction with the ad across social networks.

[0044] In an aspect, a method for using cosine similarity to determine the relationship between one or more clusters may include for each cluster, building a vector based on the CFI scores calculated for a number of items, plotting the vectors in a 3D vector space, determine the cosine of the angle between the vectors as an indication of the relationship between, the clusters, and when a relationship is identified between clusters based on the cosine, automatically labeling the clusters with the same label. If the cosine is small, the confidence that there is a high degree of similarity is high.

[0045] In an aspect, a method may include publishing a map of content as a widget, and tracking interaction with the content in the widget to obtain behavioral data about a user of the map.

[0046] In an aspect, a method may include publishing a map of content as a widget, tracking interactions with the content in the widget to obtain behavioral data about a user of the published map; and analysing the behavioral data in order to at least one of suggest content, track network evolution, modify the network in strategically valuable ways, and measure the success of an ad campaign.

[0047] These and other systems, methods, objects, features, and advantages of the present disclosure will be apparent to those skilled in the art from the following detailed description of the preferred embodiment and the drawings.

[0048] References to items in the singular should be understood, to include items in the plural and vice

versa, unless explicitly stated otherwise or clear from the text. Grammatical conjunctions are intended to express any and all disjunctive and conjunctive combinations of conjoined clauses, sentences, words, and the like, unless otherwise stated or clear from the context.

BRIEF DESCRIPTION OF THE FIGURES

[0049] The structures, methods, systems, inventions and the following detailed description of certain embodiments thereof may be understood by reference to the following figures:

[0050] FIG. 1 depicts a process flow for attentive clustering.

[0051] FIG. 2 depicts a social network map in the form of a proximity cluster map.

[0052] FIG. 3 depicts a social network map in the form of a proximity cluster map highlighting attentive clusters of liberal and conservative U.S. bloggers, and British bloggers.

[0053] FIG. 4 depicts a social network map in the form of a proximity cluster map focused on environmentalists, feminists, political bloggers, and parents.

[0054] FIG. 5 depicts a social network map in the form of a proximity cluster map with a cluster relationship identified.

[0055] FIG. 6 depicts a social network map in the form of a proximity cluster map with a bridge blog identified.

[0056] FIG. 7 depicts a flow diagram for attentive clustering.

[0057] FIG. 8 depicts a Political Video Barometer valence graph.

[0058] FIG. 9 depicts a graph of CFI scores.

[0059] FIG. 10 depicts a graph of CFI scores.

[0060] FIG. 11 depicts a bi-polar valence graph of link targets in the Russian blogosphere.

[0061] FIG. 12 depicts an interactive burstmap interface.

[0062] FIG. 13 depicts a valence graph of outlink targets organized by proportion of links from liberal vs. conservative bloggers.

[0063] FIG. 14 depicts a flow diagram relating to social media maps.

[0064] FIG. 15 depicts a flow diagram relating to refreshing social media maps.

[0065] FIG. 16 depicts a flow diagram relating to social media maps.

[0066] FIG. 17 depicts formation of a ranked target list.

[0067] FIG. 18 depicts Peakedness vs. Commitment by Time Range for two sets of hashtags.

[0068] FIG. 19a depicts Peakedness vs. Commitment by Subsequent Uses.

[0069] FIG. 19b depicts Peakedness vs. Commitment by Commitment by Time Range.

[0070] FIG. 20 depicts a distribution of mention-weighted normalized concentration by topic.

[0071] FIG. 21 depicts a distribution of Cohesion by topic.

[0072] FIG. 22a depicts a chronotope of the #metro29 hashtag.

[0073] FIG. 22b depicts a chronotope of the #samara hashtag.

- [0074] FIG. 22c depicts a chronotope of the #iRu hashtag.
- [0075] FIG. 23 depicts a social media map platform user flow.
- [0076] FIG. 24 depicts a recent activity page for a social media map platform.
- [0077] FIG. 25 depicts a recent activity page for a social media map platform.
- [0078] FIG. 26 depicts an overview page for a social media map platform.
- [0079] FIG. 27 depicts an interactive map for a social media map platform.
- [0080] FIG. 28 depicts an overview page for a social media map platform.
- [0081] FIG. 29 depicts an influencers page for a social media map platform.
- [0082] FIG. 30 depicts an influencer detail for a social media map platform.
- [0083] FIG. 31 depicts a conversation leaders page for a social media map platform.
- [0084] FIG. 32 depicts a tweets page for a social media map platform.
- [0085] FIG. 33 depicts a websites page for a social media map platform.
- [0086] FIG. 34 depicts a key content page for a social media map platform.
- [0087] FIG. 35 depicts a media page for a social media map platform.
- [0088] FIG. 36 depicts a terms page for a social media map platform.
- [0089] FIG. 37 depicts a lists page for a social media map platform.

DETAILED DESCRIPTION

[0090] The present disclosure relates to a computer-implemented method for attentive clustering and analysis. Attentive clusters are groups of authors who share similar linking profiles or collections of nodes whose use of sources indicates common attentive behavior. Attentive clustering and related analytics may include measuring and visualizing the prominence and specificity of textual elements, semantic activity, sources of information, and hyperlinked objects across emergent categories of online authors within targeted subgraphs of the global Internet. The disclosure may include a set of specialized parsers that identify and extract online conversations. The disclosure may include algorithms that cluster data and map them into intuitive visualizations (publishing nodes, blogs, tweets, etc.) to determine emergent clusterings that are highly navigable. The disclosure may include a front end/dashboard for interaction with the clustering data. The disclosure may include a database for tracking clustering data. The disclosure may include tools and data to visualize, interpret and act upon measurable relationships in online media. The approach may be to segment an online landscape based on behavior of authors over time, thus creating an emergent segmentation of authors based on real behavior that drives metrics, rather than driving metrics based on pre-conceived lists. Because the analysis is a structural one, rather than language-based, the analysis is language agnostic. In an embodiment, the segmentation may be global, such as of the English language blogosphere. In an embodiment, the segmentation may involve a relevance metric for every node based on semantic markers and a custom mapping of

high-relevance nodes. The disclosure enables identifying influencers, such as who is authoritative about what to whom.

[0091] One method of obtaining attentive clusters may involve construction of a bipartite matrix, however, any number and variety of flat or hierarchical clustering algorithms may be used to obtain an attentive cluster in the disclosure. In an embodiment, a set of content-publishing source nodes ("authors") may be selected based on a chosen combination of linguistic, behavioral, semantic, network-based or other criteria. A mixed-mode network may be constructed, comprising the set S of all source nodes, the set T of all outlink targets from selected types of hyperlinks, and the set E of edges between them defined by the selected type or types of links from S to T found during a specified time period. A matrix, such as a bipartite graph matrix, may be constructed of source nodes in S linked to targets in T, derived by any combination of a.) normalizing nodes in T, optionally to a selected level of abstraction, b.) using lists of target nodes for exclusion ("blacklists"), and c.) using lists of target nodes for inclusion ("whitelists"). The matrix may represent a two-mode network (or actor-event network) that associates two completely different categories of nodes, actors and events, to build a network of actors through their participation in events or affiliations. In embodiments, the matrix is, in effect, an affiliation matrix of all authors with the things that they link to, wherein the patterns of their linking may be used to do statistical clustering of their nodes.

[0092] The matrix may be processed according to user-selected parameters, and clustered in order to perform one or more of the following: 1.) partition the network into sets of source nodes with similar linking histories ("attentive clusters"); 2.) identify sets of targets (linked-to websites or objects) with similar citation profiles ("outlink bundles"); 3.) calculate comparative statistical measures across these partitions/attentive clusters; 4.) construct visualizations to aid in interpretation of network features and behavior; 5.) measure frequencies of links between attentive clusters and outlink bundles, allowing identification and measurement of large-scale regularities in the distribution of attention by authors across sources of information, and the like. An arbitrary number and variety of flat or hierarchical clustering algorithms may be used to partition the matrix, and the results may be stored in order to select any solution for output generation. The resulting outputs (measures and visualizations) may provide novel, unique, and useful insights for determining influential authors and websites, planning communications strategies, targeting online advertising, and the like.

[0093] In an embodiment, systems and methods for attentive clustering and analysis may be embodied in a computer system comprising hardware and software elements, including local or network access to a corpus of chronologically-published internet data, such as blog posts, RSS feeds, online articles, Twitter™ "tweets," Facebook™ postings, and the like.

[0094] Referring to FIG. 1, attentive clustering and analysis may include: 1.) network selection 102, 2.) partitioning 104, which may include two-mode network clustering in this embodiment, and 3.) visualization and metrics output 108. Network selection 102 may include at least two operations: a.) node selection 110, and b.) link selection 112. Optionally, a third may be applied in which network analytic operations are used to further specify the set of source nodes under consideration for clustering. For example, the operation may be filtering. Filtering may be technology-based, blacklist-based, whitelist-based, and the like.

[0095] In an embodiment, nodes may be URLs, at which chronologically published streams or elements of content may be available. An initial set containing any number of nodes may be selected based on any combination of node-level characteristics and/or calculated relevance scores. Regarding node-level characteristics, there may be a number of different kinds of nodes publishing content online, such as weblogs (blogs), online media sites (like newspaper websites), microblogs (like Twitter™), forums/bulletin boards (like <http://www.biology-online.org/biology-forum/>), feeds (like RSS/ATOM), and the like. In addition to different technical genres of node, nodes may differ according to an arbitrary number of other intrinsic or extrinsic node-level characteristics, such as the hosting platform (e.g., Blogspot, LiveJournal), the type of content published (text, images, audio), languages of textual content (e.g., French, Spanish), type of authoring entity (individual, group, corporation, NGO, government, online content aggregator, etc.), frequency or regularity of publication (daily, regular, monthly, bursty), network characteristics (e.g., central, authoritative, A-list, isolated, un-linked, long-tail), readership/traffic levels, geographical or political location of authoring entity or focus of its concern (e.g., Russian language, Russian Federation, Bay Area Calif.), membership in a particular online ad distribution network (e.g., BLOGADS, GOOGLE™ ADSENSE), third-party categorizations, and the like.

[0096] To support node selection 110 based on relevance to particular issues or actors, or relevance-based node selection 110, lists of relevance markers may be used to calculate composite scores across nodes. These lists may include such items as key words and phrases, semantic entities, full or partial URLs, meta tags embedded in site code and/or published documents, associated tags in third-party collections (e.g., DELICIOUS tags), and the like. For example, tags may be collected automatically, such as by “spidering” sites for meta keywords. The corpus of internet data may be scanned and matches on list elements tabulated for each node. A number of methods may be used to calculate a relevance score based on these match counts. In an embodiment, relevance scores may be calculated by calculating individual index scores for text matches (T), link matches (L), and metadata matches (M), and then summing them. These individual index scores (I) may be calculated for each node by scanning all content published by a node during a specified period of time using a list of j relevance markers:

$I = \text{sum}((x_1 * w_1)/t_1 + (x_2 * w_2)/t_2 \dots (x_j * w_j)/t_j)$, where x is the number of matches for the item, w is a user-assigned weight (a scale of 1 to 5 is typical), and t is the total number of item matches in the scanned corpus. In an example, an initial set of source nodes may include the 100,000 Russian language weblogs most highly cited during a particular time frame. In another example, the initial set may include the 10,000 English language weblogs with the highest relevance scores based on relevance marker lists associated with the political issue of healthcare. In another example, the initial set may include all nodes by Indian and Pakistani authors in whatever language that have published at least three times within the past six months.

[0097] With respect to the link selection 112 component of network selection 102, objects may be particular units of chronologically published content found at a node, such as blog posts, "tweets," and the like. Links, also referred to as outlinks herein, may be hyperlink URLs found within a node's source HTML code or its published objects. Many kinds of links exist, and the ability to choose which kinds are used for clustering may be a key feature of the method. There are links for navigation, links to archives, links to servers for embedded advertising, links in comments, links to link-tracking services, and the like. Link selection 112 may be applied to links that represent deliberate choices made by authors, of which there may also be many kinds. These links may be to nodes (e.g., a weblog address found in a "blogroll"), objects (e.g., a particular YOUTUBE™ video embedded in a blog post), and other classes of entity, such as "friends" and "followers." Some node hosting platforms define a typology of links to reflect explicitly defined relationships, such as "friend," "friend-of," "community member," and "community follower" in LIVEJOURNAL, or "follower" and "following" in Twitter™, Facebook™ and the like. In other cases, informal conventions, such as "blogrolls," define a type of link. Some of these link types are relatively static, meaning they are typically available as part of the interface used by a visitor to a node website, while others are dynamic, embedded within published content objects. Link types may be parsed or estimated and stored with the link data. These links represent different types of relationships between authors and linked entities, and therefore, according to the user's objectives, certain classes of links may be selected for inclusion. Different sorts of links also have time values associated with them, such as the date/time of initial publication of an object in which a dynamic link is embedded, or the first-detected and most recently seen date/time of a static link. Links may be further selected for clustering based on these time values.

[0098] From the parameters defined for node selection 110 and link selection 112, a mixed-mode network X 130 may be constructed, consisting of the set S of all source nodes, the set T of all outlink targets from selected types of hyperlinks, and the set E of edges between them defined by the selected type or types of links from S to T found during a specified time period. The network 130 may be considered "mixed mode" because while it may be formally bipartite, a number of

nodes in S may also exist in T, which may be considered a violation of the normal concept of two-mode networks. Rather than excluding nodes that may be considered either S or T nodes, the systems and methods of the present disclosure consider them logically separate. A particular node may be considered a source of attention (S) in one mode, and an object of attention (T) in the other. Before clustering, the set of nodes may be further constrained by parameters applied to X, or to a one-mode subnetwork X' consisting of the network 130 defined by nodes in S along with all nodes in T that are also in S (or at a level of abstraction under an element in S, collapsed to the parent node). Standard network analytic techniques may be applied to X' in order to reduce the source nodes under consideration for clustering. For instance, requirements for k-connectedness may be applied in order to limit consideration to well-connected nodes.

[0099] In an embodiment, partitioning 104 may include: 1.) specification of node level for building the two-mode network, 2.) assembly of bipartite network matrix 132 using iterative processing of matrix to conform with chosen threshold parameters, and 3.) statistical clustering (multiple methods possible) of nodes on each mode, that is, source node clustering 114 and outlink clustering 118. Outlink clustering 118 to form an outlink bundle may involve identifying sets of web sites that are accessed by the same kinds of people.

[0100] With respect to specification of node level, distinction may be made between “nodes” and “objects,” considering the node as a stable URL at which a number of objects are published. This may result in generation of a straightforward two-level hierarchy (object-node); however, nodes sometimes have a hierarchical relationship among each other (object-node-metanode). Consider the following three URLs:

[0101] 1.) <http://www.bloghost.com/>;

[0102] 2.) <http://www.bloghost.com/users/johndoe/blog/>; and

[0103] 3.) <http://www.bloghost.com/users/johndoe/blog/09/6/21/myblogpost.html->.

[0104] Here, a three-level hierarchy with a metanode [1], node [2], and object exists. In some embodiments, the node URL may correspond very simply to a “hostname” (the part of a URL after “http://” and before the next “/”) or a hostname plus a uniform path element (like “/blog” after the hostname). In other embodiments though, multiple nodes may exist at pathnames under the same hostname. Depending on the objective of the user, a “node level” may be selected for building the two-mode network, such that second mode nodes include (from most general to most specific level) a.) metanodes (collapsing sub-nodes into one) and independent nodes, b.) child, or sub-nodes (treated individually) and independent nodes, or c.) objects (of which a great many may exist for any given parent node). In embodiments, it may be possible to mix node levels according to a rule set based on defining levels for particular sets of nodes and metanodes, or on link thresholds for qualifying objects independently. Furthermore, a node with a webpage URL may

often have one or more associated “feed” URLs, at which published content may be available. These feeds are generally considered as the same logical node as the parent site, but may be considered as independent nodes. If a target URL is not a publishing node, but another kind of website, the level may likewise be chosen, though more levels of hierarchy may be possible, and typically the practical choice may be between hostname level or full pathname level.

[0105] With respect to the assembly of the bipartite network matrix 132 using iterative processing of the matrix 132 to conform with chosen threshold parameters, links may be reviewed and collapsed (if necessary) to the proper node level as described hereinabove, and the two-mode network may be built between all link sources (the initial node set) and all target (second-mode) nodes at the specified node level or levels. Optionally, blacklists and whitelists may be used to, respectively, exclude or force inclusion of specific source or target nodes. From this full network data, an $N \times K$ bipartite matrix M , in which N is the set of final source nodes and K is the set of final target nodes, may be constructed according to user-specified, optional parameters, such as maxnodes, nodemin, maxlinks, linkmin, and the like. An iterative sorting algorithm may prioritize highly connected sources and widely cited targets, and then use these values to determine which nodes and targets from the full network data may be included in the matrix. Maxsources and maxtargets may set the maximum values for the number of elements in N and K . Nodemin may specify the minimum number of included targets (degree) that a source is required to link to in order to qualify for inclusion in the matrix. Linkmin similarly may specify the minimum number of included sources (degree) that must link to a target to qualify it for inclusion in the matrix. Two other optional parameters, nodemax and linkmax may be used to specify upper thresholds for source and target degree as well. Each value (V_{ij}) in M , is the number of individual links from source i to target j .

[0106] With respect to statistical clustering in each mode, that is node clustering 114 and outlink clustering 118, there may be a number of clustering algorithms which may be used to partition the network, including hierarchical agglomerative, divisive, k-means, spectral, and the like. They may each have merits for certain objectives. In an embodiment, one approach for producing interpretable results based on internet data may be as follows: 1.) make M binary, reducing all values >0 to 1; 2.) calculate distance matrices for M and its transpose, yielding an $N \times N$ matrix of distances between sources, and a $K \times K$ matrix of distances between targets. Various distance measures may be possible, but good results may be obtained by converting Pearson correlations to distances by subtracting from 1; 3.) using Ward’s method for hierarchical agglomerative clustering, a cluster hierarchy (tree) may be computed and stored for each distance matrix. Results of an arbitrary number of clustering operations may be saved in their entirety, so that any particular flat cluster solutions may be chosen as the basis for generating outputs.

[0107] In an embodiment, the clustering algorithm may be language agnostic, that is, forming attentive clusters around similar targets of attention without a constraint on the language of the targets. In an embodiment, clustering may make use of metadata that may enable the system to know about the content of various websites without having to understand a language. In another embodiment, the algorithm may have a translator or work in conjunction with a translation application in order to find terms across publications of any language.

[0108] Now that the first two stages of attentive clustering, network selection and two-mode network clustering, have been described we turn to a description of visualization and metrics output. Any particular set of cluster solutions for source nodes (an assignment of each node to a cluster) may be selected by the user in order to generate one or more of the following classes of output: 1.) per-cluster network metrics for source nodes 120; 2.) across clusters comparative frequency measures of link, text, semantic and other node and link-level events, content and features; 3.) visualizations 124 of the partitioned network combined with these measures and other data on node and link-level events, content and features; and 4.) aggregate cluster metrics reflecting ties among clusters taken as groups. Further, any particular set of cluster solutions for target nodes may be selected and used in combination with the set of cluster solutions for source nodes in order to generate: 1.) measures of link frequencies and densities 128 between source clusters and target clusters; 2.) visualization 124 of the previous as a network of nodes representing clusters of sources and targets with ties corresponding to link densities 128; and 3.) visualizations 124 of one-mode calculated (network of target nodes) networks with partition data.

[0109] In one class of output, and with respect to per-cluster network metrics for source nodes 120, in addition to standard network metrics for source nodes that are generated over the entire network, and which reflect various properties important for determining influence and role in information flow, user-selected cluster solutions may be used to generate a set of measures for each node, per-cluster. These measures may represent the node's direct and indirect influence on, or visibility to, each cluster, as well as its attentiveness to each cluster. For every node i , these measures may include the following: same-in: the number of nodes in the same cluster that link to i ; same-out: the number of nodes in the same cluster i links to; diff-in: the number of nodes in other clusters that link to i ; diff-out: the number of nodes in other clusters that i links to; same-in-ratio: the proportion of in-linking nodes from the same cluster; same-out-ratio: the proportion of in-linking nodes from other clusters; w-same-in: same-in scores where value of in-linking blogs is weighted by its centrality measure; w-diff-in: diff-in scores where value of in-linking blogs is weighted by its centrality measure; and per-cluster influence scores: similar scores (raw and weighted) for in-links from, and out-links to, each cluster on the map.

[0110] In another class of output, and with respect to across clusters comparative frequency measures of link, text, semantic and other node and link-level events, content and features, the partitioning of the network into sets of source nodes may allow independent and comparative measures to be generated for any number of items associated with source nodes. These may include such items as: a) the set of target nodes K in M ; b.) any subset of all target nodes, including those on user-generated lists; c.) any set of target objects, such as all URLs for videos on YOUTUBE™, or all object URLs on user-created lists; d.) any other URLs; e.) any text string found in published material from source nodes; f.) any semantic entities found in published material from source nodes; g.) any class of meta-data associated with source nodes, such as tags, location data, author demographics, and the like. For any item i in a set of items associated with source nodes, the following examples of measures may be generated per each cluster: 1.) total count: number of occurrences of item within the cluster (multiple occurrences per source node counted); 2.) node count: number of nodes with item occurrence within cluster (multiple occurrences per source node count as 1); 3.) item/cluster frequency: total count/# of nodes in the cluster; 4.) node/cluster frequency: node count/# of nodes in the cluster; 5.) standardized item/cluster frequency: multiple approaches are possible, including z-scores, and one approach is to use standardized Pearson residuals, which control for both cluster size and item frequency across clusters and items in the set; and 6.) standardized node/cluster frequency: multiple approaches are possible, including z-scores, and one approach is to use standardized Pearson residuals, or Cluster Focus Index scores 122. The higher the CFI score for the item, the greater the degree of its disproportionate use by the cluster. A score of zero indicates that the cluster cites the source at the same frequency as the network does on average. Other detailed data may be possible to obtain, such as the top nodes in each cluster, lists of all nodes in the cluster, lists of relevant Internet sites that each of the clusters link to (which enables identifying target outlinks where a message can be placed in order to reach specific clusters), the relative use of key terms across the clusters (which enables developing specific messages to communicate to each cluster), a hitcount (the raw number of times each outlink and term was found within all the identified nodes), source node and/or cluster geography and demographics, sentiment, and the like.

[0111] For example, differential frequency analysis can be done on meta-data, such as tags, that are associated with different attentive clusters to facilitate cluster interpretation. In the example, by sorting cluster focus scores 122 on the meta-data tags, interpretations of what the clusters are about may be derived without any manual review. The meta-data associated with the clusters may be used to facilitate interpretation of the meaning of the clusters. In an example, the meta-data may be language independent, such as GIS map data.

[0112] In another class of output, and with respect to visualizations of the partitioned network 124, a social network diagram may be generated and used to display link, text, semantic and other node and link-level events, content and features (“event data”), such as that shown in FIG. 2. The network map may be static or it may be the basis of an interactive interface for user interaction via software, software-as-a-service (SaaS), or the like. There may be two components to this process of visualization: 1.) creating a map of source nodes in a dimensional space for viewing; and 2.) use of colors, opacity and sizes of graphical elements to represent clusters, nodes and event data. With the dimensional mapping component, multiple approaches may be possible. One method may be to use a “physics model” or “spring embedder” algorithm suitable for plotting large network diagrams. The Fruchterman-Reingold algorithm may be used to plot nodes in two or three dimensions. In these maps, every node is represented by a dot, and its position is determined by link to, from, and among its neighbors. The size of the dot can vary according to network metrics, typically representing the chosen measures of node centrality. The technique is analogous to a locally-optimized multidimensional scaling algorithm. With the component related to use of colors, opacity and sizes of graphical elements to represent clusters and event data, nodes may be colored according to selected cluster partitions, to allow easy identification of various partitions. This projection of the cluster solution onto the dimensional map may facilitate intuitive understanding of the “social geography” of the online network. This type of visualization may be referred to as a “proximity cluster” map, because proximity of nodes to one another indicate relationships of influence and interaction. Further, projection of event data onto the map may enable powerful and immediate insight into the network context of various online events, such as the use of particular words or phrases, linking to particular sources of information, or the embedding of particular videos. This may be produced as static images, and may also be the basis of software-based interactive tools for exploring content and link behavior among network nodes.

[0113] In another class of output, and with respect to aggregate cluster metrics 128, metrics may be calculated for partitions at the aggregate level. Event metrics may include raw counts, node counts, frequencies (counts/# nodes in cluster), normalized and standardized scores, and the like. Examples typically include values such as: the proportion of blogs in a cluster using a certain phrase; the number of blogs in a cluster linking to a target website; the standardized Pearson residual (representing deviation from expected values based on chance) of the links to a target list of online videos; the per cluster “temperature” of an issue calculated from an array of weighted-value relevance markers; and the like.

[0114] As described above, any particular set of cluster solutions for target nodes may be selected and used in combination with the set of cluster solutions for source nodes in order to generate additional outputs. Visualizations produced may include: 1.) two-mode network diagram of

relationships between clusters of sources and targets, treated as aggregate nodes and with tie strength corresponding to link density measures; and 2.) second-mode (“co-citation”) network diagram, in which targets are nodes, connected by ties representing the number of sources citing both of them, and colors corresponding to cluster solution partitions. Another output may be macro measurement of link density. To reveal and measure large-scale patterns in the distribution of links from source nodes to targets, the matrix M may be collapsed to aggregate link measures among clusters of sources and clusters of targets. A series of $S \times T$ matrices may be used, with S as the set of source clusters (“attentive clusters”) and T as the set of clustered targets (“outlink bundles”). These matrices may contain aggregated link measures, including: counts (c): the number of nodes in source cluster s linking to any member of target set t ; densities (d): c divided by the product of the number of members in s and the number of members in t ; and standard scores (s): standardized measures of the deviation from random chance for counts across each cell. Various standardized measures are possible, with standardized Pearson residuals obtaining good results. Any of these measures may be used as the basis of tie strength for two-mode visualizations described above.

[0115] In an embodiment, a density matrix may be constructed between attentive clusters and outlink bundles. The attentive clusters may be represented as row headers and the outlink bundles may be represented as column headers. The density matrix may allow users to see patterns in attention between certain sets of websites and certain bundles. The density matrix may provide a way to identify similar media sources. Further, the density matrix may provide information about attentive clusters that may be based on particular verticals.

[0116] Having described the process for attentive clustering, we now turn to examples of applications of the technique and various related analytical applications thereof for measuring frequencies of links between attentive clusters and outlink bundles, thus enabling identification and measurement of large-scale regularities in the distribution of attention by online authors across sources of information.

[0117] In an embodiment, and referring to FIG. 2, a social network map of the English-language blogosphere is depicted. The social network map graphically depicts the most linked-to blogs in the English language blogosphere. The size of the icons representing each individual blog may be representative of a network metric, such as the number of inbound links to the blog. This visualization depicts the output from a method for attentive clustering and analysis which identified attentive clusters of linked-to blogs, wherein the attentive clusters included authors with similar interests.

[0118] Referring to FIG. 3, the method for attentive clustering and analysis analyzes bloggers' patterns of linking to understand their interests. The visualization in FIG. 3 highlights liberal and

conservative U.S. bloggers, and British bloggers as attentive clusters. By zooming in on the visualization, subgroups such as conservatives focused on economics or liberals focused on defense may be identified from among the attentive clusters depicted.

[0119] Referring to FIG. 4, the method for attentive clustering and analysis enables building a custom network map. In FIG. 4, the network map features attentive clusters of bloggers attuned to these topics: environmentalists, feminists, political bloggers, and parents. Subgroups within each topic may be delineated by a different color, a different icon shape, and the like. For example, within the parent bloggers, icons representing the liberal parent bloggers may be colored differently than the traditional parent bloggers. Surprising relationships may be discovered among groups of bloggers. For example, in FIG. 5, two parent bloggers with very different social values are closer in the network than either is to political bloggers who share their broader political views.

[0120] Referring to FIG. 6, each attentive cluster may have its own core concerns, viewpoints, and opinion leaders. The method for attentive clustering and analysis enables identification of blogs that are considered bridge blogs, such as the one shown circled, which indicates that the blog is popular among multiple attentive clusters. The method for attentive clustering and analysis enables identification of whose opinions matter, about what, and among what groups.

[0121] Referring to FIG. 7, the steps of attentive clustering and analysis may include constructing an online author network, wherein constructing the online author network includes selecting a set of source nodes (S), a set of outlink targets (T) from at least one selected type of hyperlink, and a set of edges (E) between S and T defined by the at least one selected type or types of hyperlink from S to T during a specified time period 702; deriving a set of nodes, T', by any combination of a.) normalizing nodes in T, optionally to a selected level of abstraction, b.) using lists of target nodes for exclusion ("blacklists"), and c.) using lists of target nodes for inclusion ("whitelists") 704; transforming the online author network into a matrix of source nodes in S linked to targets in T' 708; and partitioning the online author network into at least one set of source nodes with a similar linking history to form an attentive cluster and at least one set of outlink targets with a similar citation profile to form an outlink bundle 710. The steps may optionally include generating a graphical representation of attentive clusters and/or outlink bundles in the network to enable interpretation of network features and behavior and calculation of comparative statistical measures across the attentive clusters and outlink bundles 712, wherein at least one element of the graphical representation depicts a measure of an extent of a type of activity within the network; and optionally measuring frequencies of links between attentive clusters and outlink bundles enabling identification and measurement of large-scale regularities in the distribution of attention by online authors across sources of information 714. The element of the graphical representation may use at least one of size, thickness, color and pattern to depict a type of activity. Attentive clusters may

be visually differentiated in the graphical representation by at least one of a color, a shape, a shading, and a size. The size of the object representing the attentive clusters in the graphical representation may correlate with a metric. The nodes, targets, and edges may be collected from public and private sources of information. Constructing the matrix may include applying at least one threshold parameter from the group consisting of: maxnodes, targetmax, nodemin, targetmin, maxlinks, and linkmin. Constructing the matrix may include applying a minimum threshold for the number of included nodes that must link to a target to qualify it for inclusion in the matrix. Constructing the matrix may include applying a minimum threshold for the number of included targets that must link to a node to qualify it for inclusion in the matrix. Constructing the matrix may include using blacklists to exclude particular nodes, and whitelists to force inclusion of particular nodes. The matrix may be a graph matrix.

[0122] By identifying and measuring the frequencies of links between attentive clusters and outlink bundles, all manner of information about the distribution of attention by online authors across sources of information may be obtained. Various examples of the sorts of information, visualizations, applications, reports, APIs, widgets, tools, and the like that are possible using the methods described herein will be described. For example, two playlists for YOUTUBE™ videos may be identified, one that has fraction with sub-cluster A the other with sub-cluster B. In another example, two RSS feeds may be organized that supply a user with items that have more attention from sub-cluster A versus sub-cluster B. In another example, a valence graph may be constructed that depicts words, phrases, links, objects, and the like that are preferred by one sub-cluster over another sub-cluster; such valence graphs may use aggregated sets of clusters defined by users to display dimensions of substantive interest, such as in FIG. 11. In yet another example, works from authors who are most relevant in a particular cluster may be displayed and then published as a widget, which may be custom-based on a valence graph, as a way of monitoring an ongoing stream of information from that cluster. Clusters may be customizable within the widget, such as via a dialog box, menu item, or the like. Further examples will be described hereinbelow.

[0123] A user may be able to, optionally in real time through a user interface, select a stream of information based on looking at the environment, zoom in based on clustering, figure out a valid emergent segmentation, and then set up monitors to watch the flow of events, such as media objects, text, key words/language, and the like, in real time.

[0124] In an embodiment, differences in word frequency use by attentive clusters may be used to differentiate and segment clusters. For example, the attentive clusters “militant feminism” and “feminist mom” may both frequently use terms associated with feminism in their publications, but additional use of terms related to militantism in one case and maternity in another case may have been used to subdivide a cluster of feminists into the two attentive clusters “militant feminism”

and “feminist mom.” In extending this concept, not just word usage but the frequency of word usage, may also be useful in segmenting clusters. For example, in clusters of parents, the ones actually doing home schooling did not use the term “home school” frequently, but rather used the term “home education” with greater frequency. By identifying the specific language/words used by a cluster, the system may enable crafting messages, brands, language, and the like for particular clusters. In an embodiment, an application may automatically craft an advertisement to be placed at one or more outlinks in an outlink bundle using high frequency terms used by an attentive cluster. Further in the embodiment, the advertisement may be automatically sent to the appropriate ad space vendor for placement at the one or more outlinks.

[0125] In an embodiment, a method of using attentive clustering based on analysis of link structures to steer a further data collection process is provided. The data collection may include collection of web-based data, such as, for example, clickstream data, data about websites, photos, emails, tweets, blogs, phone calls, online shopping behavior, and the like. For example, tags may be collected automatically or manually for every website that is a node. The tags may be non-hierarchical keywords or terms. These tags may help describe an item and may also allow the item to be found again by browsing or searching. In an example, tags may be associated in third-party collections such as DELICIOUS tags, and the like. In another example, there web crawlers may extract meta keywords and tags included within node html. Further, specific keywords and phrases may be exported to a database. In yet another example, the tags may be generated by human coders. Once a cluster partitioning exists, the system may do differential frequency analysis on the tags that are associated with different attention clusters. By sorting cluster focus index (CFI) scores along with the tags, the system can come up with an interpretation of the meaning of a cluster without requiring further analysis of the cluster itself. In an embodiment, the system may apply a further data collection process in order to associate respondents to a survey and their news sources with various corners of the internet landscape. For example, the influence of a particular news outlet across a segmented environment of the online network may be obtained by examining clustering in conjunction with a downstream data collection process, such as obtaining survey research, clickstream data, extraction of textual features for content analysis including automated sentiment analysis, content coding of a sample of nodes or messages, or other data.

[0126] In an embodiment, clustering data may be overlaid on GIS maps, “human terrain” maps, asset data on a terrain, cyberterrain, and the like.

[0127] In an embodiment of the present disclosure, a method of determining a probability that a user will be exposed to a media source given a known media source exposure is provided. The media source may include newspapers, magazines, radio stations, television stations, and the like.

For example, a user who may be exposed to a particular media source may be clustered in a specific attentive cluster. Accordingly, the system may determine that users in that particular attentive cluster are more likely to be exposed to another media source because the second media source may also be present in an outlink bundle preferred by the cluster.

[0128] In an embodiment of the present disclosure, a method of attentive clustering on a meso level is provided. The method may enable identifying emergent audiences (Attentive Clusters) and monitor how messages (as specific as a single article in print; as broad as core campaign themes) traverse cyberspace. The method may involve mapping the attentive clusters where messages have, or are likely to find, receptive audiences. Mapping may enable identifying opinion leaders, and information sources, online and offline, which help shape their views.

[0129] The method may enable identification of the mindset/social trends of a group of users. For example, the system may be able to associate an attentive cluster with a known network, such as a political party, a political movement, a group of activists, people organizing demonstrations, people planning protests, and the like. Via the ability to associate attentive clusters with particular groups of people, the system may be able to track the evolution of a movement or identity over time. Further, if a cluster supports a political movement, the system may track the impact of the political movement of the cluster on society. The system may track if the political movement has been accepted by majority of the people of the society, rejected by the society, if there is debate about the political movement, and the like. Accordingly, the method may enable growth of a brand, sale of a product, conveying a message, prediction of what people care about or do, and the like.

[0130] In an embodiment of the present disclosure, a system and method for multi-layer attentive clustering may be provided. In the system and method, attentive clusters may be tracked across various layers of a social segmentation, such as specific social media networks (Twitter™, Facebook™, Orkut™, and the like), a blogosphere, and the like. The system may be able to track development of an attentive cluster in a single layer or across multiple layers at every stage of the development of the cluster. When different layers of online media (such as weblogs, microblogs, and a social network service) are clustered individually, measures of association may be created between clusters across layers, based on density of hyperlinks between them, common identities of underlying authors, mutual citation of the same sources, mutual preference for certain topics or language, and the like. The system may also track the major players of clusters at every stage of development of the cluster.

[0131] For example, the growth of an attentive cluster supporting a political movement may be tracked back in time and over a period of a time. In the example, once an attentive cluster may be identified, the system may examine the nodes associated with specific players in the attentive

cluster in order to determine characteristics, such as who is talking to whom, identify key nodes or hubs that link many other layers and/or media sources, identify apparent patterns of affinity or antagonism among clusters or other known networks, who may have started the political movement, when the political movement may have started, what messages were used at the forefront of the political movement's establishment, the size of the movement, the number of people who initially joined the political movement, growth of the political movement, influential people from various stages of the political movement, and the like. In this example, all of the analysis may be confined to activity in a single layer of a social segmentation or it may be undertaken across multiple layers. Continuing with the example, the impact of the political movement on society may be examined by tracking the penetration of an attentive cluster or its message across layers or the expansion of the attentive cluster in a single layer. Likewise, attentive cluster analysis may enable predictions. For example, an attentive cluster may be tracked in a single layer, such as by monitoring the number of Twitter™ followers (or other applicable social platforms), the frequency of new followers added, the content associated with that attentive cluster, inter-cluster associations, and the like, to determine if a political movement may be being spawned, expanded, diminished, or the like. In an embodiment, the socio-ideological configuration of the people who spawned the political movement may be evident from analyzing one or more of a blog layer, a social networking layer, a traditional media layer, and the like.

[0132] For example, a Twitter™ (or other applicable platform) map may be formed where each colored dot is an individual Twitter™ account and the position is a function of the “follows” relationship. People are close to people they are following or who are following them. The pattern of the map may be related to the structure of influence across the network.

[0133] In an embodiment, the system may be deployed on a social networking site to identify and track attentive clusters and linkage patterns associated with the attentive clusters. For example, the system for attentive clustering may be applied on Facebook™ to identify attentive clusters in the Facebook™ audience and track the cluster's activity within Facebook™. In an example, the system may be used to identify a group of people who may be susceptible to a message. By identifying and tracking an attentive cluster in the Facebook™ layer that may be susceptible to a message, downstream activities, such as organizing in response to the message, may be examined. For example, an attentive cluster of university students may be presented with a message regarding a proposed law lowering the drinking age. The system may track activity within the cluster related to the message, identify new groups formed around the topic of the message, invitations to other groups regarding the message, opposition from other groups in response to the message, and the like. Indeed, the system may be able to track the formation of new attentive clusters in the Facebook™ layer in response to the message. In this case, the system may identify individuals or

groups that link to one another who share a common interest or target of attention, such as concerned parents opposing the proposed law, anti-government groups supporting the proposed law, child advocate groups opposing the law, and the like. Discoveries related to the original layer may be applied to strongly associated clusters in other layers. For instance, determination about the interests of a cluster in the Facebook™ layer may be used to drive a communications or advertising strategy in associated clusters of other layers such as weblogs or Twitter™.

[0134] Measures for characterizing contagious phenomena propagating on networks may include peakedness, commitment (such as by subsequent uses and time range), and dispersion (including normalized concentration and cohesion) and will be further described herein.

[0135] In other embodiments, two-mode networks may be generated by projecting modes one onto another. For example, certain social networks may not allow handling of individual data, but may allow public page data to be accessed. In this way, data from individuals who comment on public pages may be obtained. Public pages may be treated as a two-mode network that is collapsed to one mode. For example, a two-mode network may be formed from two classes of actors, people and cocktail parties that the people attend. One class of actors could be labeled I-5 and the other A-E to generate a scatter diagram depicting a two-mode network, either a network of cocktail parties attended by the same people or a network of people who attended the same cocktail parties. Likewise, networks may be formed based on who participates in the stream of objects that come from different public pages, the relationship between public pages, such as if there is a direct “like” relationship between public pages, weighted by how many people commented on objects from two or more pages, and the like.

[0136] These data may be clustered as described herein. In embodiments, the weight between public pages indicated by the number of users commenting on object from both pages may be used to visually indicate a stronger connection between pages with higher weights.

[0137] Clustering of this public page data may result in the formation of poles. For example, two poles may form where one set of pages is interacted with by one population and another set of pages interacted with by a very different population. There may be individuals who are interacting with both of these sets of pages at either pole. In any event, in the process of attentive clustering, users who are most tenuously connected to anything are forced to the outer edges of the cluster map.

[0138] In an embodiment of the present disclosure, a method of analyzing attentive clusters over time is provided. The analysis of these attentive clusters may enable the system to depict changes in the linking patterns of attentive clusters over a time period. Further, the analysis may allow depiction of any changes in the structure of the network itself.

[0139] In an embodiment, a time-based reporting method may be used by the system to demonstrate the effects of events/actions throughout a network of attentive clusters for a period of time. In the method, bundles that may be lists of semantic markers, including text elements embedded in a post or tweet, links to pieces of online content, metadata tags, and the like, may be tracked in clusters across a network, such as a blogosphere.

[0140] For example, a bundle of semantic markers related to obesity may be tracked over time to determine how the topic of obesity is being discussed. In the example, a particular bundle (with text, link and meta data elements) can be tracked across clusters to see where they are getting attention or not. The measure of attention may be defined as a “temperature.” The “temperature” is based conceptually on Fahrenheit temperatures (without negatives) as compared to other issues where 100 is very hot and 0 is ice cold. The method may have a tracking report as an output for tracking issues in a map across time. In this example, the tracking report may be focused on a collection of blogs most focused on childhood obesity organized into attentive clusters over a moving 12-month period of time. The blogs may be clustered broadly into policy/politics, issue focus, culture, family/parenting, and food attentive clusters. There may be sub-clusters defined for each of those clusters, such as conservative, social conservative, and liberal sub-clusters under the policy/politics cluster. The report may indicate the issue intensity for each cluster/sub-cluster by assigning it an average temperature per blog of conversation on the broad topic of childhood obesity within each group. The report may indicate the issue distribution for each cluster/sub-cluster by calculating a percentage of childhood obesity conversations taking place on blogs not in the map and within each cluster within the map. Continuing with this example, specific terms may be tracked across the clusters/sub-clusters over time and the method may indicate an average temperature based on the uses of specific terms in blogs within each cluster. In the example, the term “school lunch” has a high “temperature” in certain issue focus clusters, liberal policy clusters, and foodie clusters and steadily increased over the last eight moving 12-month periods. Similarly, the intensity of sites, or the average temperature based on links to specific web sites on blogs within each cluster, may be provided by the report. The intensity of source objects, or the average temperature based on the links to specific web content (articles, videos, etc.), may be provided by the report. The intensity of sub-issues, or the average temperature of conversation on identified issues defined by a set of terms and links, may be provided by the report. In the report, specific terms may be tracked on a monthly and per-cluster basis, specific sites may be tracked on a monthly and per-cluster basis, and specific objects may be tracked on a monthly and per-cluster basis.

[0141] In an exemplary embodiment, the system may identify and track structural changes in a network. For example, during the recent US elections, blogs appeared instantaneously that were

anti-Obama, Pro-Palin, or Pro-McCain but were outside the conservative blogosphere. This rapid change in the network structure may be indicative of a coordinated, synchronized campaign to message and blog.

[0142] In an embodiment of the present disclosure, a method of attentive clustering by partitioning an author network into a set of source nodes with similar adoption and use of technology features is provided. For example, instead of a website being a target of attention for an attentive cluster or around which an attentive cluster forms, a feature or a piece of technology, such as an embedded Facebook™ “Like” button, may be a target of attention or clustering item.

[0143] In an embodiment, a method of creating clusters of people and describing probabilistic relationships with other clusters, such as words, brands, people, and the like, is provided. The system may describe any probability of any relation between them.

[0144] To identify what an attentive cluster links to more than the network average or what words and phrases they use more than the network average, a cluster focus index score (CFI) may be calculated. CFI represents the degree to which an event, characteristic, or behavior disproportionately occurs in a particular cluster, or a particular cluster, relative to the network, preferentially manifests an event, characteristic, or behavior. For example, CFI score could be generated for a particular cluster across a set of target nodes, representing the degree to which a particular target is disproportionately and preferentially cited by members of the particular cluster, or the degree to which the particular cluster, relative to the network, preferentially cites the target. The CFI gives a sense of what is important to an attentive cluster, where they go for their information, what words, phrases and issues they discuss, and the like. FIG. 9 depicts a graph of cluster focus index scores for targets of a conservative-grassroots attentive cluster. The targets circled on FIG. 9 (F through J) are those that everyone in the network links to, according to their CFI. The targets circled in FIG. 10 (A through E) are those that are disproportionately linked to by the conservative-grassroots attentive cluster, according to their CFI.

[0145] In an embodiment, a method of identifying websites with high attention from an identified attentive cluster or author is provided. The method may include determining the websites frequently or preferentially cited by identified authors by examining the websites' cluster focus index (CFI) score. Further, the method may include automatically sending or placing advertisements, alerts, notifications, and the like to the websites. For example, a social network analysis may generate a network map with thousands of nodes clustered into attentive clusters. In an example with bloggers, influence data that results from the network analysis may be influence metrics for sites from across the Internet which bloggers link to, including mainstream media, niche media, Web 2.0, other bloggers, and the like. These are the influential sources (also called outlinks, or targets) used by specific groups of nodes across the map. For example, influencing a

targeted cluster of bloggers can often be accomplished by targeting these sources, “upstream” in the information cycle, rather than going after the bloggers directly. In other embodiments, influence data may be metrics that reveal network influence among bloggers directly. Bloggers are usually thought of as simply being more influential or less, but this data lets the analyst discover which blogs are influential among which online clusters (segments), a far more granular and targeted approach. Each of these data sets can be sorted to examine either influence over the entire map or disproportionate influence over particular clusters (i.e., how to reach particular audiences). Cluster targeting can be further refined to identify which nodes in a specific cluster have influence on any of the other clusters on the map. Because the conversation within social media covers a wide variety of topics, source and network influence alone do not necessarily reflect influence on a specific topic. A relevance index metric for discussion regarding particular topics, events, and the like may be added to a social network analysis to identify which nodes are most focused on this topic.

[0146] For both data sets there are two main sorts metrics representing influence. First are metrics representing the influence of nodes in the one-mode network (set of source nodes S) as a whole, or directly among particular clusters or among specific other nodes. For example, for any given node in S, count (also called in-degree) is the number of other nodes in S that link to it. Count can be calculated across the whole map, or per cluster. Second, score can be calculated that shows the influence of target nodes (nodes in T or T') on clusters of nodes in S. Count can also be used, and CFI scores can be calculated that represent the influence of particular targets on specific attentive clusters. In other words, how specifically interesting or authoritative the target is for that cluster. Relevance index scores may for nodes may also be calculated using lists of semantic markers, to provide further metrics of value for targeting communications, advertising, and the like. Depending on the communications strategy, specific sorts of the data will create lists of likely high-value targets for further action. While count, CFI, and relevance index scores are all important, they can be combined in order to maximize certain objectives. The following use case examples include combining count and relevance into a targeting index, by multiplying their values. Other, more complicated maximization formulas are possible as well. The examples demonstrate specific influence sorts that can be generated from the Russian network data to address each use case. The network data is based on the linking patterns of the nodes in the RuNet map over a nine-month period ending in February 2010.

[0147] Use Case 1 and Use Case 2 involve finding influential sources. Use Case 1 involves identifying sources with the most influence over the entire map by doing a sort using the highest values of count. While extremely influential, and in many cases suitable for advertising

campaigns, these universally salient sites also tend to be much harder to reach out to than sites that are smaller but specifically important to targeted segments.

[0148] Use Case 2 involves identifying sources that reach a targeted cluster by sorting on sources by Cluster Focus Index. CFIs may be sorted for any of the attentive clusters. Count metrics from the map as a whole and from the targeted cluster can be used to further prioritize for action. This sort is the equivalent of identifying traditional media trade press, the go-to sites for the selected segment. Frequently, these include specifically influential bloggers in addition to niche media and other sources.

[0149] Use Cases 3-6 involve finding influential nodes. Use Case 3 involves identifying the greatest network influence by sorting the nodes by indeg (total number of links from other nodes within the entire network). This sort specifically identifies the network's "A-list" nodes, the most influential network members (bloggers). Like prominent sources, these are often more difficult to reach than more targeted niche influentials, but they contribute greatly to spreading viral niche messages across the wider network.

[0150] Use Case 4 involves finding the most targeted influencers for a particular cluster by sorting the Cluster Focus Index scores for a targeted cluster to find nodes with cluster-specific influence. This identifies the nodes with particular influence, interest or prestige among the target cluster. These nodes tend to be much more "on topic" than others, and much easier to reach than map-wide A-list nodes. Cluster-specific influentials are not always from the target cluster itself, which can be very useful for trying to move discussion between particular clusters. Link metrics provide further assistance in deciding targeting priorities.

[0151] Use Case 5 involves following a particular topic at the map level by sorting using topic focus target scores, which combine links (network influence) and topic focus index (issue relevance). Formulas for calculating focus target score can be varied, but the default may be to multiply links by topic focus index. This may allow identification of those nodes in the entire map that discuss the target issue most frequently. These may be monitored to gauge dominant threads of discussion and opinion about the issue, and targeted for outreach.

[0152] Use Case 6 involves targeting a particular cluster's conversation on a topic by sorting within a cluster by the topic focus target score. This may allow members of the target cluster who write about the target issue to be identified for monitoring or persuasion. Variations of the formula for combining influence and relevance metrics into a single targeting metric can be used to bias the sort toward relevance, or toward influence, depending on strategic objective.

[0153] In an embodiment, a proximity cluster map method may be used to visualize 124 attentive cluster-based data and generate a network map. In the method, attentive clusters and their constituent nodes may be displayed in a proximity cluster map. Nodes in the network map may

be represented by individual dots, optionally represented by different colors, whose size is determined based on the number of other nodes on the map link to them. A general force may act to move dots toward the circular border of the map, while a specific force pulls together every pair of nodes connected by a link. In static images or an interactive visualization via software connected to a database, nodes may receive a visual treatment to display additional data of interest. For example, dots representing nodes may be lit or highlighted to represent all nodes linking to a particular target, or using a particular word, with other nodes darkened. In another example dot size may be varied to indicate a selected node metric.

[0154] In an embodiment, a valence graph method may be used to visualize 124 attentive cluster-based data and generate a valence graph. In the method, targets of attention or semantic elements occurring in the output of nodes may be displayed in a valence graph. The valence graph method may be understood via description of how a particular valence graph is built, such as a Political Video Barometer valence graph (FIG. 8) useful for discovering what videos liberal and conservative bloggers are writing about. This particular valence graph may be used to watch and track videos linked-to by bloggers who share a user's political opinions, view clips popular with the user's political "enemies," and the like.

[0155] The videos shown in the Barometer are chosen by queries against a large database built by network analysis engines performing network selection 102. Periodically, a crawler (or "spider") visits millions of blogs and collects their contents and links. Next, the system mines the links in these blogs to perform partitioning 104 and forms attentive clusters based on how the blogs link to one-another (primarily via their blog rolls), and, over time, what else the bloggers link to in common. Attentive clusters may be large or small, and the bigger ones can contain many sub-clusters and even sub-sub-clusters. In embodiments, determining what the blogs have in common may be done by examining meta-data, tags, language analysis, link target patterns, contextual understanding technology, or by human examination of the blogs or a subset thereof. In the example, American liberal bloggers and American conservative bloggers form the two largest sets of clusters in the English language blogosphere, and the Barometer draws upon roughly the 8,000 "most linked-to" blogs in each of these groups to position the videos on the graph by calculating proportions of links to each target by the two political cluster groupings.

[0156] The Barometer may be continually updated by scanning the blogs periodically, looking for new links to videos (or videos embedded right in the blogs). By counting these links, it can be determined what videos political bloggers are promoting. In embodiments, the link count may be displayed on the valence graph using an identifier such as icon or marker. In this example, some videos are linked to almost exclusively by liberal bloggers, some are linked to mostly by conservative bloggers, and a few are linked to more or less evenly by both groups. Once the

system determines that a video has traction in the political clusters, it scans through data from other parts of the blogosphere to count how many “non-political” bloggers link to it as well.

[0157] The Political Video Barometer example illustrates one kind of valence graph and the insight that can be gained and the applications that can be built based on the method and the data obtained by the method. It should be understood that the method may be used to examine any sort of potentially cluster-able data, such as technology, celebrity gossip, the use of linguistic elements, the identification of new sub-clusters of particular interest, and the like. All aspects of the valence graph method, and the underlying attentive clustering analysis, may be customized along multiple variables to enable planning and monitoring campaigns of all kinds.

[0158] In an embodiment, a multi-cluster focus comparison method may enable comparing cluster focus index (CFI) scores of multiple attentive clusters. The CFI score may be a measure of the degree to which a particular outlink is of disproportionate interest to the attentive cluster being analyzed; in other words, the CFI indicates what link targets are of specific interest to a particular cluster beyond their general interest to the network as a whole. In an example, X may be the CFI score for cluster A and Y may be the CFI score for cluster B. The multi-cluster focus comparison method may compare the two clusters, A and B, based on their CFI scores, X and Y. This would allow a user to discern elements of common interest vs. divergent interest between the two clusters. Insights derived from this method would be of great value in creating and targeting advertising and communications campaigns.

[0159] In another embodiment, link targets, semantic events, and node-associated metadata may be scattered in an x-y coordinate space, and the dimensions of the graph may be custom-defined using sets of clusters grouped to represent substantive dimensions of interest for a particular analysis. Elements are plotted on X and Y according to the proportions of links from defined cluster groupings. For example, and referring to FIG. 11, using data from the Russian blogosphere, the top 2000 link targets for Russian bloggers may be plotted such that the proportion of links from “news-attentive” blog clusters vs. links from “non-news attentive” clusters determined the position on Y, while the proportion of links from the “Democratic Opposition” cluster vs. the “Nationalist” cluster determines the position on X, as shown in FIG. 11. In another example, popular outlink targets for the US blogosphere may be displayed with the X dimension representing the proportion Liberal vs. Conservative bloggers linking to them, and the proportion of political bloggers of any type vs. non-political bloggers represented by the Y dimension, as shown in FIG. 13. Various data may be visualized in the graph associated with the clusters of news-attentive and political bloggers, such as meta-data tags, words, links, tweets, words that occur within 10 words of a target word, and the like. These visualizations may be used in

interactive software allowing user-driven exploration of the data graphed in valence space, optionally allowing user-defined sets of clusters to be used in calculating valence metrics.

[0160] In an embodiment, a method of node selection 110 based on node relevance to a defined issue, also known as semantic slicing, is provided. Semantic slicing may involve clustering according to a relevance bundle. A relevance bundle may include one or more of key markers, what the nodes may have linked to, what the nodes have posted, text elements, links, tags, and the like. In essence, semantic slicing involves pre-screened nodes for relevance based on semantic analysis.

[0161] The relevance bundles may be used to sort through all of the network data to select the top high relevance nodes. In an embodiment, a custom-mapping of a sub-set of the link economy may be done.

[0162] In an embodiment, semantic slicing may enable generating a contextualized report of interest to a user on an industry level. Semantic slicing may enable focusing attentive clustering on selected vertical markets. The vertical markets may be a group of similar businesses and customers who may engage in trade based on specific and specialized needs. Lists of semantic markers, such as key words and phrases, links to relevant websites and online content, and relevant metadata tags, are built which represent the relevant vertical market. Relevance metrics are calculated for candidate nodes, and a selection of high-relevance nodes are mapped and clustered. Continuing the example, the semantic slice may be done to analyze an energy policy vertical market by focusing the attentive clustering around one or more selected, highly relevant nodes. Thus, the attentive clusters may be more specific to identified domain interest of interest or vertical market. In this example, instead of just forming an attentive cluster of Conservative bloggers, by focusing attentive clustering on one or more key markers related to energy policy, the attentive clusters discovered include topic-relevant segmentations of particular kinds of Conservative bloggers discussing the issue, such as Conservative-Grassroots and Conservative-Beltway. Additional high-relevance attentive clusters may be identified, such as Climate Skeptics, Middle East policy, and the like. Cluster focus index scores may be used to determine what sites everyone in each cluster links to and which sites are preferred by the cluster. In an embodiment, semantic slicing may be done using a single node, such as a particular website, a particular piece of content, and the like. In an embodiment, semantic slicing may be done over a period of time to enable monitoring the impact of a campaign.

[0163] In an embodiment, a tool, such as software-as-a-service, for enabling users to define one or more semantic bundles for attentive clustering and as the basis of report outputs is provided. The tool may be an on-demand tool that may be used for semantic slicing. In such models, a user may declare a semantic bundle of nodes and/or links prior to attentive clustering.

[0164] In an embodiment, the system may provide an application programming interface (API) for delivering a segmentation to track one or more particular clusters of attention, or track how an audience is interacting with a piece of content, and the like. The data about the various clusters may be collected directly from the API. For example, a user may wish to track a cluster. The user may enter keywords related to the cluster in a search option provided by the API. Thereafter, the tool may track various websites and report back the weblinks and data that may be relevant to the cluster. The API may be used to interact with a valence graph at various resolutions. The API may provide segmentation data and metadata derived from the segmentation to other analytics and web data tracking firms, for use in their own client-facing tools and products. The segmentation and resultant data from attentive clustering provide an additional dimension of high value against which third-party tools and other analytic capabilities such as automated sentiment monitoring may be leveraged.

[0165] In an embodiment, the system may enable real-time selection of elements to visualize based on attentive clustering of social media. The system may facilitate selection of a stream of information based on looking at the environment, zooming in on a data element based on clustering, determining a valid emergent segmentation, and monitoring the flow of events in real time. The events may include media objects, text, key words/language, and the like. For example, the real-time selection of elements may facilitate an analysis of trends/events especially for financial purposes.

[0166] In an embodiment, a search engine may be provided that prioritizes search results being displayed to a user based on a determination of real-time attention including attention from a particular cluster or set of clusters. A user may be able to customize the prioritization of search results, such as by getting real-time attention from a particular cluster, from a particular sub-cluster, and the like.

[0167] In an embodiment, a search engine is provided that searches within only those sites/accounts with high cluster focus for a chosen segment. For example, a GOOGLE™ search may be restricted to the 30 websites with the highest CFI scores for the Dirt Bike racing cluster of OAKLEY's TWITTER™ followers map. Thus, the search may only return results from a list of key influential sites related to the chosen segment. In other embodiments, the search may be restricted to websites (or domains within them) with a particular CFI score, websites (or domains) that meet a threshold CFI score, websites that fall into a range of CFI scores for a chosen segment, websites with a particular M score, and the like. In an embodiment, the search query may restrict the search to particular websites that are identified based on the CFI scores. In an embodiment, the search query may be restricted by CFI score of a website and the CFI score restriction may be indicated in the settings of the search engine. In other embodiments, the CFI score for sites to

search may be indicated in the search string itself. For example, a user may indicate a particular search they want to perform and they may be provided with a slider bar where the user indicates that the search should be restricted to those websites with a CFI score falling into the range selected on the slider bar. The slider may be provided with a normalized scale, such as ascribing 1 to low CFI scores and 10 to high CFI scores, such as using a linear, logarithmic, or other scaling process. The system may then search a database of websites for the range of CFI scores to identify one or more websites to which to limit the search. These websites are then included in a search string that is provided to a search engine.

[0168] Similarly, the search can be restricted to only specific content, or specific content may be promoted to high ranking within a search, leaving other content to the lower ranked results. One way to do this restriction is to utilize the valence mapping functionality of the system. As described herein, a valence graph may be constructed for a chosen segment that depicts words, phrases, links, objects, and the like that are preferred by one cluster over another cluster. Content indicated in the valence graph may be indexed by the system and only that content in the valence graph may be searched by a search engine. Further restriction of the content may be employed, such as by website, CFI score, and the like.

[0169] In an embodiment, attentive clustering and related analyses may result in identifying issues, attitudes and messaging language that may be specific to discourse for a target market, and may be suitable for presentation in a report. For example, in a clustering of bloggers sympathetic to Arts in Schools, by examining intra-cluster linking patterns, it may be determined that most of the bloggers within each cluster tend to keep the discussion within their cluster except for the bloggers in the “Interesting/teachers/educators” cluster who have a tendency to spread conversation to each of the other clusters. This behavior points to an opportunity to work with these bloggers to spread messages across the space. In continuing with the example, by examining clustering related to specific keywords, websites, outlinks, objects, and the like, it may be determined that there is a broader discussion about education and education reform than about arts and arts education. Therefore, a conclusion may be that introducing an arts education message to education discussions has more potential than introducing arts education messages to arts discussions. In the report, various valence graphs may be presented, such as cluster specific term valence maps, maps of sources, outlink maps, term specific maps, issue maps, and the like. Alternatively, the report may be presented as a spreadsheet of data.

[0170] In an embodiment of the present disclosure, the report may feed into a method of generating a campaign blueprint for both social and upstream media sources and a method of identifying influence inter-cluster and intra-cluster in order to plan a campaign. The blueprint may include target audience, demographic details, objectives of the campaign, flow of the

campaign, messaging to use in the campaign, outlinks to target, and the like. Systems and methods for measuring the success of a campaign in various online segments and generating targeted data sets identifying sub-clusters specific to a user's identity or objective are provided.

[0171] In an exemplary embodiment, the campaign tracker may track data from a variety of sources to provide closed-loop return on investment (ROI) analysis. The tool may parse the information of each website accessed by the users, keywords entered, any information about the campaign, and the like. Further, the tool may track how people react to the campaigns and which ones are most successful. The campaign tracker may track and analyze results in real-time to determine the effectiveness of the campaigns.

[0172] In addition, the tool may enable the system to generate reports for clients. The reports may include details about the campaigns such as campaign type, number of people who have viewed the campaign, any feedback from the people, and the like.

[0173] In an embodiment, analyst coding tools (ACT) and a survey integrator may support distributed metadata collection for qualitative analysis to best interpret quantitative findings. The tools may include an interactive visual interface for navigating complex data sets and harvesting content. This interface may contain an interactive proximity cluster map which can display specific node data, metadata, search results, and the like. This proximity cluster map interface may enable the user to click on nodes to see node-specific metadata and to open the node URL in a browser window or external browser. Using the tools, a user can add metadata and view metadata about any given blogger on a map. The tools enable grabbing whole sets of blogs or items to add to semantic lists, and may enable a user to define surveys so a team of human coders can open the website and fill out surveys.

[0174] In an embodiment of the present disclosure, a dashboard may be provided. The dashboard may combine advanced network and text analysis, real-time updates, team-based data collection and management, and the like. In the embodiment, the dashboard may also include flexible tools and interfaces for both "big picture" views and minute-by-minute updates on messages as they move through networks. Using the dashboard, a user may define bundles and track them in the aggregate through networks over time. Using the dashboard, a user may be able to see how specific media objects are doing with a particular cluster over time.

[0175] In an embodiment, the dashboard may provide a burstmap feature in which the history of selected events or sets of events over a timeframe may be displayed using a proximity cluster map. During playback, nodes in the map will light up at a time corresponding to their participation in the selected event or events. For example, at a time in playback representing a certain date, every node which linked to a particular YOUTUBE™ video will light up, allowing the user to see the pattern of linking as it unfolded over time. Optionally, this burstmap feature may include a

timeline view displaying event-related metrics over time, such as the number of nodes linking to a particular video. Optionally, the burstmap feature may include lists of events available for display. An example of a burstmap interface is found in FIG. 12.

[0176] In an embodiment, techniques disclosed herein may be used to generate social media maps that visualize social media relationship data and enable utilization of a suite of metrics on the data. Social media maps may be constructed via clustering of various social media communities including TWITTER™, FACEBOOK™, blogs, online social media, and others. In one embodiment, the clustering technique used may be manual, relationship-based, attentive clustering such as previously disclosed herein, network segmentation, or another analogous technique. The social media maps may be organized in portfolios that are targeted to market segments or relate to an issue/topic campaign. Social media maps may be offered via an API or as raw data to plug into a third party dashboard. Services related to the social media maps that may be offered include robust tools for searching, comparing and generating integrated reports across multiple maps, searchable indexing and map browsing. Pricing for social media maps may be via subscription, for one or more maps, a portfolio of maps, the whole portfolio of maps, the whole portfolio maps save some exclusive/custom items, or the like. Systems and methods for how to generate, utilize, update and offer social media maps will be further described herein.

[0177] A comprehensive catalog of social media maps and network segmentations may be offered and updated on a regular basis. The catalog may include targeted portfolios for key markets, such as consumer goods, media and entertainment, politics and public policy, energy, science and technology, government, and more. The catalog may contain maps for each layer of the social media system, such as blogs, Twitter™, social network services, forums, and the like. It may contain maps for all major languages, countries and regions of the world. Social media map data may be used within partner dashboard systems, so that a range of commercial tools can be leveraged by subscribers and so that the social media map data are “portable” across various tools. In addition, a suite of reporting tools may be used in conjunction with the social media maps.

[0178] In an embodiment, one or more social media maps and network segmentations may be constructed via clustering of data from at least one social media community. The social media map or network segmentation may be offered via an API or as raw data. The social media community may be based on at least one of a social media layer, a language, a country, a region, or the like. In some embodiments, the clustering technique may be attentive clustering, as described previously herein, relationship-based, manual, network segmentation, or the like. Referring now to FIG. 14, relationship-based clustering of data from at least one social media community 1402 is used to construct one or more social media maps and network segmentations using the clustering 1404. One or more social media maps and network segmentations may be

offered via an API 1408 or as raw data 1410. A report may demonstrate the interaction of nodes/links between the maps 1412.

[0179] In embodiments, the maps may be generated by an autonomous process. The autonomous process may create maps based on one or more criteria, a scope definition, an instruction, or the like. For example, a social graph may be generated based on followers of an individual or entity in a social network. In another example, the map criteria may be semantically based, such as based on key words or hashtags. In yet another example, the maps may be geo-based, such as based on which users/nodes are in a territory. In still another example, the maps may be based on previous mappings. In this example, segments in other maps on health and fitness may be used to triangulate or iterate to a mapping of a new category. In another example, the map may be based on an arbitrary set of accounts generated by a third party. One scenario might be a mapping of the social network accounts for all the users of a mobile application. In still another example, the maps may be based on a nomination of individuals based on some criteria, such as demographics. Once generated, the maps may be stored and indexed.

[0180] In embodiments, maps may be based on CFI scores for dynamic data (e.g., YOUTUBE™ videos). However, the amount of data may be increased to obtain a better indication of what the segment is communicating about whether data can be obtained on the influencers of a segment, which may be coming from off the map. In addition to looking at data coming from the segment, the system may be able to access data from social media accounts that have high CFI for that segment (not just the ones that are “in” the segment). Thus, calculating cluster focus for the dynamic data may be improved. CFI scores may be calculated for a first segment. Then, CFI scores may be calculated for those influencers on the first segment. For example, the first segment may be followers of a particular art gallery but the system can also examine the CFI for the first segment's influencers, which may be several well-known Art Gallery aficionados who may or may not be followers of the particular art gallery. In embodiments, certain maps may be based only on the CFI scores calculated for the influencers.

[0181] A searchable index for a catalog of social media maps may be constructed 1414. Further, social media maps in the catalog may be searchable. For example, the maps may be searchable by a keyword, a URL, a semantic marker, and the like. In embodiments, the social media maps may be indexed by one or more of a keyword, URL or semantic marker so as to form a searchable index of social media maps. In embodiments, the searchable index may include metrics to indicate a statistic regarding the social media maps. For example, the statistic may represent a dimension of popularity, relevance, semantic density, or similar feature. For example, a search engine may be enabled to return maps in terms of relevance by using certain statistics in the searchable index.

[0182] For example, a semantic marker may include a keyword, a phrase, a URL (node or object level), a tag (such as those from bookmarking and annotation services, meta keywords extracted from HTML, tags assigned by coders, etc.), and the like. Semantic markers may also include those used in particular social network environments, such as TWITTER™, and may include follows relationships, mentions, retweets, replies, hashtags, URL targets, and the like. Any of these semantic markers may be used to index a social media map.

[0183] Based on at least one of the search terms or the search results, a new social media map subscription may be suggested. For example, if a user searches a social media map index for the terms “Nissan LEAF™,” “electric vehicle,” and leafstations.com, subscriptions to social media maps such as automobiles, eco-friendly products, and California trends may be suggested.

[0184] In an embodiment, a dashboard may be used for browsing, visualizing, manipulating, and calculating metrics for one or more social media maps constructed via clustering of data from at least one social media community. Clustering techniques may include relationship-based, manual, attentive clustering, or the like. In some embodiments, the dashboard may be a third party dashboard that supports visualization of data from clustering, wherein the data may be delivered by a raw data feed, an API plug-in, or any other data delivery method. In embodiments, the data from clustering may be joined with or otherwise integrated with data from other data sources to form a new data set. The new data set may be similarly browsed, visualized, manipulated, and processed by dashboards.

[0185] In an embodiment, APIs, dashboards, and partner tools may be used with social media maps for planning/assessment. For example, social media maps may be used for enterprise resource planning, business insight, marketing, search engine optimization, intelligence, politics, industry verticals, financial industry, and the like. For example, an entertainment promotion company may own a plurality of social media accounts. If they could navigate sector-level mappings related to genres of music, they could use the maps to target music genre-specific messages using the most appropriate of those accounts for maximum effectiveness.

[0186] In embodiments, custom maps may be derived from mashing up sets of social media maps.

[0187] In an embodiment, the social media maps may be constructed via clustering (e.g., relationship-based, manual, attentive, etc.) of data from at least one social media community targeted to a specific market segment. For example, the market segments may include government intelligence, public diplomacy, social media landscapes in other countries, pharmaceuticals, medical, health care, sports, parenting, consumer products, energy, and the like. In these embodiments, the market segment may be used to index the social media maps.

[0188] In an embodiment, a reporting product may leverage social media maps to demonstrate the interaction of nodes and/or links between social media maps. For example, a multi-map report

may be generated comparing the nodes and links in different social media communities in a particular market/environment. The reporting product may be integrated with a dashboard or analytics platform. Multi-map reports generated by the reporting product may be used to demonstrate various phenomena, such as how particular items can be found in particular social media layers. For example, a multi-map report may demonstrate how weblog hosts are having customers driven to them from TWITTER™. In another example, a multi-map report may demonstrate how FACEBOOK™ pages are getting attention from a segment of TWITTER™.

[0189] In an embodiment, information derived from the social media maps, including portions of or the entire map itself, may be published or displayed as a map widget, which may enable monitoring an ongoing stream of information from one or more clusters or one or more maps. Information being displayed that is derived from the social media map may be customizable within the widget, such as via a dialog box, menu item, or the like. A user may be able to, optionally in real time through a user interface, select a stream of information based on looking at the environment, zoom in based on clustering, figure out a valid emergent segmentation, and then set up monitors to watch the flow of events, such as media objects, text, key words/language, and the like, in real time. The published, widgetized map acts as a sensor network to obtain a host of behavioral data and leads that can be leveraged by the map's user or hosts. In embodiments, users may interact with other users' map widgets to discover content and individuals/entities. Using other users' map widgets, users may grow their own networks by engaging with the content and people/entities in the widget, such as to start following a person or to retweet an item.

[0190] There are at least three processes that yield attributes of nodes, including calculating a relevance score, performing a CFI bias weighting, and identifying nodes as "allowed" or "not allowed" (e.g., blacklist/whitelist). Automated social media map refresh may leverage one or more of these processes.

[0191] In an embodiment and referring to FIG. 15, a social media map may be automatically refreshed via calculating a relevance score for nodes or bundles in the map 1502 and re-constructing the map based on a relevance ranking revealed by the relevance score 1504. Semantic/relevance marker bundles may include lists of semantic markers like key words, phrases, relevant link targets, accounts that are followed on TWITTER™, and the like. Semantic markers may be manually curated. In an embodiment, the refresh process may involve performing the relevance search/semantic slice that generated the original map for new relevance/semantic markers. A relevance calculation may be performed on the nodes to calculate a relevance score.

[0192] In another embodiment, a social media map may be automatically refreshed via positively or negatively weighting at least one cluster based on a CFI score calculation 1508 and re-constructing the map to modify the nodes in the clusters 1510. Modifying the nodes may be done

to include positively weighted nodes and exclude negatively weighted nodes. CFI scores for clusters may be leveraged to evolve a map in a certain direction. Clusters in the map that include preferred/wanted nodes/links are positively weighted. Clusters are negatively weighted in they are deemed to not be relevant. Applying weightings to the map may enable pulling in additional nodes that are more relevant. Weighting map clusters for the CFI bias operation may be done by humans.

[0193] In an embodiment, a social media map may be automatically refreshed via filtering out unwanted nodes 1512. In an embodiment, a social media map may be automatically refreshed via obligatorily including nodes that were not clustered in the original map 1514. Semantic markers that are known to not fit based on their relevance ranking or for some other reason are not allowed are filtered out. In embodiments, nodes may be forced into the map whether or not they were identified in the relevance search/semantic slice. Curating black lists of nodes may be done by humans.

[0194] In an embodiment, a social media map may be automatically refreshed via crowd-sourced information regarding nodes and/or links that drive nodes to bundles 1518. In an embodiment, a social media map may be automatically refreshed via processing social media map usage data for trends/indicators 1520. Usage data may relate to one or more of what is ignored, what is further explored, what is used, how clusters are grouped, what name/label is assigned to a cluster, what color is used for a cluster, what order/position the cluster is placed in a report, and the like. Nodes preferentially interacted with may be weighted more heavily.

[0195] In embodiments, community feedback may influence each of the three streams of automated map refresh described herein. Community feedback provides an indication of news, events, information, etc. that may drive addition of nodes to the bundles, such as for example, if a new website is a target link. This sort of feedback may provide feedback or guidance as to the CFI bias operation. For example, if feedback suggests that a cluster is relevant, then that cluster may be positively weighted.

[0196] Feedback and updating may be based on how people are using the maps, such as, understanding what they ignore, what they drill down on, what they use, how they want to group things, what name/label they assign a cluster, what color they use for a cluster, what clusters are most important to a client based on an order/position the client places it in a report, and the like. Refreshing the maps may leverage this captured information.

[0197] In an embodiment, feedback may be received passively from clickable/interactive maps via a built-in feedback system. This feedback system may be used as a naive weighting system. In an embodiment, the map may include a flag available to provide commentary or feedback.

[0198] In an example, a map may include raw clusters and human-made groupings and the attachment of other sort of metadata such as the coloring of a cluster. The example may be that of the Russian blogosphere, which may contain 40 clusters and 7-8 groups, including 5 right wing Russian nationalist groups and a liberal opposition group. Clusters may be processed by human-assigned re-aggregation, and metrics may be run against them to progressively refine the clusters. Different clients, even on a base map, may want to group things differently, name a cluster in an interface differently, color a cluster in an interface differently, and the like. Users need to be able to define groups, re-label clusters, select clusters and the like. Community feedback may provide observations as to how users are grouping the same map and that yields data about which clusters are related to each other that is "crowd-sourced" to the user. Users may define the order in which the data are presented in the reporting. For example, a user may want to place data on preferred clusters higher in a chart. Cluster ordering and positioning information is customizable, which can be harvested as an importance weighting by the community.

[0199] In another example, map users may contribute to map metadata to generate a community data set established and/or expanded by users. For example, users could input the gender of a Tweeter/blogger. The user community itself may be a segmentable population. The user community can contribute to scoping a map for a particular topic. For example, something about a disease might appear in various places: Consumer segments, Politics, Medical/science, Sports, and the like. User feedback may also help scope the size of the map. For example, a user may ask: Should the map be constructed on the first 5,000 targets or should 20,000 targets be used? In an embodiment, user-contributed data may be used to provide metadata for a social media map constructed via clustering (e.g., relationship-based, manual, attentive, or the like) of data from at least one social media community.

[0200] In an embodiment and referring to FIG. 16, data, including user-contributed data, may form a searchable, editable metadata and basic information repository for URLs 1602, such as to form a URLipedia. The repository may be linked to one or more social media maps 1604.

[0201] In an embodiment and referring to FIG. 17, clustering (e.g., relationship-based, manual, attentive, or the like) of data from at least one social media community may be used to generate an actionable targeting list. Targeting lists combine network centrality 1704, issue relevance 1708 and CFI for a cluster 1710 into a ranked target list 1702 that may be used by marketers or other interested parties in order to reach certain nodes in some meaningful order for targeting for strategic communication or other business purpose. The formula of combination may be adjusted to maximize ranking to suit client/user objectives. In an embodiment, network centrality may be a universal score related to how central a node is in the network. For example, daytime talk show hosts may have a network centrality of 100 in the general population, while economists may be a

zero. In an embodiment, a Cluster Focus Index score may be calculated for each cluster. For example, daytime talk show hosts may be a zero CFI for economics, but economists are 100. In an embodiment, an issue relevance score may be calculated for each cluster. For example, the issue relevance related to the budget deficit may be calculated based on a publication frequency score (e.g., # of tweets). Other score techniques may be used to calculate an issue relevance.

[0202] In an embodiment, users may be able to purchase ads or message placements on a target from the targeting list 1712. From the targeting list, users may be enabled to buy an ad placement or message placement on the target site at the click of a button. In an embodiment, the effect, or impact, of the ad/message placements may be tracked for the node and across a social media map. Thus, the system may enable users to identify targets according to a ranked list based on network centrality, CFI, and issue relevance, and then place and track ads/messages on the targets from the lists. In another embodiment, targeting lists may be used in connection with any ad network for ad/message placement. Tracking ads/messages may involve receiving feedback on actions taken with respect to the ads/messages, calculating impact metrics, and the like.

[0203] In an embodiment, a historical data browser may provide a mechanism for visualizing archived, historical social media map data, such as for research or historical purposes. For example, there may be value to academia of accumulating old social media maps and showing the delta between them, such as to explore how the market has evolved over some period of time. Historical social media map data may also be useful for financial industry forensics and intelligence analysis.

[0204] In an embodiment, CFI metrics may be displayed on a social media map. A CFI metric for items in clusters indicates how much attention there is to that item for that cluster. An attention score indicates the relative attention to an item as compared to other items for a cluster for a range of time or for a "point" in time. A higher attention score means the item is more specific to the cluster. Attention scores are non-linear in the sense that anything below two is not significant and greater than two, it is exponentially significant.

[0205] CFI scores may be a metric for measuring search engine optimization and/or advertising effectiveness because it represents cluster specificity. CFI metrics would have to be combined with a more global metric to enable companies to shift from thinking at the execution/implementation layer (e.g., where do I advertise?) to the strategic layer (e.g., where are we going with this community?).

[0206] In an embodiment, a CFI Graph may include CFI scores for sources and nodes on the map. In the upper right of the map are clusters with high focus on the particular cluster, high overall level of attention, and many in-links. On the CFI graph, users can see various items at a glance.

For example, users may find the key players related to a topic or the landscape of players to determine who has influence.

[0207] In an embodiment, a CFI graph may include a Cluster Map Properties Editor/User Interface. The interface enables users to label clusters, assign clusters to a group, and perform group metrics.

[0208] Maps may be generated based on semantic elements, bundles, white lists, black lists, and the like in an automated fashion in some embodiments but labeling the clusters in an automated way, such as when a map update is made, may be difficult. Draft labels may be assigned when the cluster is created or updated based on a previous storehouse of knowledge. A confidence score as to that labeling may be generated. To automate the labeling, members of a cluster may be compared with membership of clusters of past maps and if a high percentage are the same then it is assumed the clusters relate to the same thing and are labeled similarly. In another embodiment, automated labeling is based on a structural equivalence. Labeling a node or an object that has well defined properties may be easier than labeling a cluster, which is a collection of objects. Structural equivalence involves examining the node's outlinks. For example, if people are friends with the same people, then they may have similar interests. In another example, blogs that link to the same sets of things are likely to be similar. In yet another example, if there are two people who have superior relationships to twenty soldiers, chances are that the two people are sergeants or some other form of commander. While this may work at the node level, it is harder to do at the cluster level. CFI scores, which are already generated for clusters, may be used in the generation of labels. For example, for two clusters with numerous links from nodes in these clusters to other nodes, it is difficult to compare the clusters at face value. One might just be larger, more popular, or have more links. However, CFI scores enables a comparison between two items or sets of items that a cluster may be disproportionately paying attention to. For example, Cluster 1 is very interested in horses and baseball, while Cluster 2 is very interested in horses and basketball. Given the CFI scores, vector cosine similarity can be used to determine the relationship between the two clusters. For each cluster, vectors can be built based on the CFI scores calculated for each of the clusters for the same items (e.g., Cluster 1=CFI1(1), CFI1(2) . . . etc.; Cluster 2=CFI2(1), CFI2(2) . . . , etc.). The vectors may be plotted in a 3D vector space. The cosine of the angle between the two vectors may be one indication of the relationship between the two clusters. If the cosine is small, the confidence is high. As maps are updated with new content, clusters in the new map can be compared to clusters of old maps. When there is a match, that is, a small angle between two cluster vectors, the label from the cluster in the old map is assigned to the cluster in the new map. In embodiments, the cosine of the angle may also act as a similarity score. There are a number of

measures for vector distance, including correlation distance, cosine similarity, Euclidian distance, and the like.

[0209] In embodiments, to limit the number of CFI's to include in vector generation the CFI's may be filtered to include only a CFI of two or more on a particular cluster. This effectively reduces the dimensionality of the space.

[0210] In other embodiments, items that are similar may be aggregated in labeling. For example, using outlink bundles rather than an individual CFI score may enable grouping items into target clusters and examining the density of links to the target cluster.

[0211] In an embodiment, an advertising campaign planning tool can enable running a campaign on blogs, and tracking success in other layers (e.g., TWITTER™, FACEBOOK™; segment-specific online forums).

[0212] In an embodiment, URL shorteners included in social media content may be tracked. The system may provide reporting outputs that track the success of a social media campaign including a URL shortener in different layers of the social media system. The system may not only be used to plan the campaign, but may also be used to report on the TWITTER™ bounce from blog activity or the FACEBOOK™ bounce from blog activity, for example.

[0213] In an embodiment, the system may enable campaign planning (e.g., domestic, international, multi-platform, multi-network, etc.) where language is not a required first limitation. For example, the system may enable campaign planning in marketing, such as, for consumer goods, media and entertainment, movie marketing, video games, social games, music, international product launches, talent agencies, public diplomacy, public health, political campaigns, and the like. Campaigns may be tracked, such as with a chronotope analysis, as will be further described herein, to determine a pattern that exists in time and space determined by combining temporal and network features in the analysis of the segments/clusters.

[0214] In an embodiment, the system may marry internal reporting with other reporting tools such as splash, resonance, clicks, transactions, and the like.

[0215] In an embodiment, the system enables analysis and prediction, such as in the financial industry (e.g., market predictions and trading positions), social media firms whose value is built around prediction, and the like.

[0216] In embodiments, third party data and clusters may be used with the mapping techniques described herein.

[0217] In embodiments, models may be built on one or more clusters using tools that can be accessed across clusters.

[0218] In some embodiments, a social media map and network segmentation may be constructed via clustering of data from a single user's social media community. Referring now to FIG. 23, a

user flow for becoming a user and interacting with a map is depicted. Starting from logical block 2300, processing flow proceeds to a login screen at logical block 2302 where users may log in, such as via a social media authorization. If the user is a new user, the user is sent to a sign up page at logical block 2304, where they may sign up or be given additional content to entice a sign-up. If the user is already on a list as having requested access, processing flow proceeds to logical block 2308 to check a wait list status. If the user is a beta user, processing flow proceeds to logical block 2310 where it is determined if the login is a first login. If so, processing flow proceeds to logical block 2312 where a tour may be taken. After the tour, processing flow may proceed to logical block 2318 where a map overview is presented, including a competitive overview, a text description, a cluster power, and the like. If the user is not a beta user, processing flow may proceed to logical block 2314, where the delta since last visit is presented, including new followers, recent activity with map indicators, and the like. Processing flow may then proceed to logical block 2318. From logical block 2318, processing flow may proceed back to logical block 2314 if recent activity is requested again.

[0219] Alternatively, if the user chooses a cluster or group at logical block 2318, processing flow may proceed to logical block 2320 to obtain a cluster overview, including local competitive performance, influencers, conversation, images, videos, recent tweets, and the like. If the user chooses to delve into the entire interactive map, processing flow may proceed to logical block 2322 for clustormap navigation. Processing flow may alternatively proceed to logical block 2324 from logical block 2320 where the user may take action. In an alternative embodiment, processing flow may first proceed to logical block 2328 where the user may first view full lists, and then processing flow may proceed to logical block 2324 where only actions that are relevant to the list being reviewed are displayed at logical block 2324. From logical block 2324, the user may choose to build a network, save one or more clusters as a list, move a message, engage with content, or the like. If choosing to build a network, processing flow may proceed to logical block 2330, where the user is prompted to make a list of influencers. From there, user details may be entered at logical block 2332, and then actions such as engaging one of the users make current logical block 2334 or a follow action may be taken at logical block 2338. From logical block 2330, a follow list may be generated at logical block 2340, or the current view may be saved as a Twitter™ list or some other social media list at logical block 2342. Likewise, if the “Save Cluster as List” action is selected, processing flow may proceed to save the current view as a Twitter™ list or some other social media list at logical block 2342. If the move message action is selected, a list of followers may be made at logical block 2344 and from there the current view may be saved as a Twitter™ list or some other social media list at logical block 2342, or a message may be composed at logical block 2348 which may include content and context and the message. If engage with content is

chosen at logical block 2324, processing flow may proceed to logical block 2350, where a list of content, such as URLs, key content and media, may be made. Users may choose to screen content details at logical block 2352 after which processing flow may proceed to logical block 2360 where a word tweet is generated, logical block 2358, where a re-tweet is generated, or logical block 2354 where tweets by influencers who tweeted the content are found and then potentially re-tweeted at logical block 2358.

[0220] In order to scale the amount of information in the social media maps, clustering techniques may need to be modified. In general, some set of nodes pay attention to some set of targets and the nodes get clustered based on the targets they pay attention to. There are at least two extensions of this general approach. In one embodiment, a very large number of nodes pay attention to a very large number of targets. Thus, for clustering, the number of operations scales at least polynomially (e.g., the cube of the number of nodes). For example, for 10,000 nodes the number of operations is in the billions. To accommodate this scale, computing power may need to be augmented.

[0221] In another embodiment, attentive gravity may be used to scale up the size of the social media maps. Nodes pay attention to targets (input data), however an object may be created where nodes are not discretely assigned to a cluster but are drawn to different poles, such as ideological, thematic, or topical poles. Depending on which nodes a target pays attention to, it can be drawn to one pole, another pole, or the middle. Instead of discrete maps with a plurality of clusters (e.g., 40) in a plurality of colors (e.g., 40), an attentive gravity map may have poles where the nodes are distributed based on how close they are to each pole. A node may have a set of scores which represent a gravitational coefficient for each of the poles of gravity. The gravitational coefficient may be used with other visualizations in order to modify the size, color, or opacity of the cluster representation based on the attentive gravity toward a pole. In another embodiment, the gravitational coefficient may simply be used as a metric on the cluster map previously described herein. The gravitational coefficient provides the degree to which a node matches a segmentation (e.g., a sports weight and a parenting weight for the same node, rather than just sorting the nodes into different clusters/segmentations and throwing out the relationship to other clusters or segmentations).

[0222] Clusters themselves may not really be definitive. For example, a node might not be in just one cluster. Such characteristics may be reflected in mapping technologies.

[0223] One technique may be a Discrimination Function. In an example, 1,000,000 nodes may be clustered. An initial condition may be a seed attentive clustering for a small number of nodes, such as 10,000. To expand the clustering, the centroids of the clusters are used to assign values to the other clusters (the X, Y average of the dots). For example, it can be determined if a new

node is closer to the centroid of one cluster or of another. As many nodes as desired to be incorporated into a map may be clustered via this technique. In this example, this technique applies to nodes 10,001 through 1,000,000.

[0224] Another technique may be to iteratively cluster the 1,000,000 nodes in batches of 10,000. Then, the CFI scores of those clusters may be used to cluster like clusters with each other. The clusters may be combined at a meta-cluster level. To make that work well, how similar some clusters are may need to be tracked across large groups of sub-clusters to see which ones are idiosyncratic and should stand alone versus ones that are somewhat consistent and should be joined.

[0225] In an embodiment, it may be desired to reduce the scale of the map to just those actors connected at a mesoscale while eliminating actors who are not really active members of the network and are just “star” followers. An Influence Network Discovery method may be used to reduce very large networks to their most influential core communities and obtain a sub-graph of maximally connected sub-actors. A variable K_{core} may be assigned to each member of the network, where K_{core} relates to a minimum connectedness, or the number of other nodes in the network to which the individual is connected (e.g., a known measure of connectedness in networks). One way to reduce the network quickly is to restrict the network by K_{core} value. For example, a network may be restricted to only those with a K_{core} of five and up, that is, only those people connected to at least five other people. Another way to reduce the network may be done iteratively. For example, a network of people surrounding the Democratic Party may be reduced iteratively. In a first step, inactive members and members with few followers may be eliminated. Then, certain network members, such as public figures or those who have a lot of followers may be removed temporarily from the network and reserved in a “keep” set. Then, the remaining network may be examined and refined by K_{core} . In the example, members of the network with a K_{core} of one are removed from the network. Removal of these people from the network may change the K_{core} values for the remaining members of the network. The process iterates, removing those network members with the lowest K_{core} values. The process can iterate until a specified number of network members is obtained. At this point, any members in the keep set may be added back to the network. As a second pass, a K_{core} of the keep set members may be done and limited to the node threshold. Based on the follow patterns of the members retained in the map, they may be assigned to a cluster.

[0226] In an embodiment, a delta report may be provided to examine the evolution of a cluster map and capture the most salient points of change in the last interval. The delta report may identify which clusters have grown, which sites are being targeted more by clusters now than before, which topics are being discussed more now than before, which clusters are more active than before, and

the like. The delta report may be provided on a periodic basis, such as weekly, monthly, and the like. Generating the delta report may involve reporting which CFI scores changed the most and which clusters are more active than before. Delta reports may be enabled by organization into a self-updating database with time snapshots. A delta report may be useful in customizing a stream of content. For example, a stream of new objects of interest for clusters in the map can be provided as a delta report and feed to a user.

[0227] In an embodiment, a self-service tool may be designed to let users access the system and initiate generation of a social media map. In an embodiment, a user may log in to the system or, in embodiments, to a social network or other third party website, in order to initiate the map creation process. A bot may be spawned that harvests data and maps the data to clusters. The bot may further provide cluster labels and CFI scores. The output may be a social media map data object with CFI scores. The self-service tool may enable user browsing of clusters and the map, tagging nodes, grouping and labeling clusters, and the like. In an embodiment, a machine learning labeler may suggest cluster labels. The user-generated labels may be fed into the machine learning facility used to label clusters for the social media maps. The focus of the self-service tool may be on actions that strategically build a user's network and strategically message to components of the network. CFIs can be used to determine a similarity among maps so that an existing social media map that is similar to the self-service map may be recommended for review.

[0228] Social media maps may be used to enable users to strategically message components of their network. In an example, a social media map may be created for the Twitter™ followers of a live entertainment company. Certain clusters relate to dense communities around particular stars or particular genres of music. For the live entertainment company, there are relatively few messages that they transmit that everyone in the map cares about; however, using social media maps, clustering enables more discrete message targeting. If the company wants to use Twitter™ to get the word out about a country artist, for example, they can target the country music cluster only with their messaging. If the company wants to target only those nodes within the country music cluster that have the highest influence, CFI scores may be used to limit the messaging in order to maximize the impact on the cluster. Such discrete targeting may be particularly useful in the case where direct messaging to followers may be limited.

[0229] Social media maps may be used to enable users to strategically build their network. For example, in the live entertainment company, the country music cluster may be growing in size. The social media map may be used to identify niche influential nodes for the country music cluster, such as by using segment CFI data to maximize connections with targeted segments/key influencers. Then, the user can start following those influential nodes in hopes that they will follow back. Such a process may help build the network in a desired strategic direction. Users

may be able to see how they are doing against competitors for any given segment by examining the proportion of influencers (high CFI target), who may or may not be in the map, following them versus others.

[0230] In one embodiment, social media maps may be organized and navigated as a map of maps, where each map appears as a node on a larger map. The strength of the connection between maps is the maximum of ratios of how many nodes are in one map versus another map. In navigating and searching the maps for a particular target, an indication may be given when a cluster in one map is very similar to another cluster in another map that may or may not be accessible by the user, for example, if one map relates to diabetes and another relates to obesity, a common cluster may be groups actively modifying lifestyles to avoid both pathologies. In embodiments, the system may provide an interface from the search screen with which the user may purchase the map they do not currently have access to.

[0231] In an embodiment, user segmentation may be used to find segments for targeting as customers. Maps may be automatically generated for the target customer and conversion rates to paying customers may be tracked.

[0232] Described herein is a system for examining social media phenomena, such as hashtags, and how they spread in a network. Patterns of spreading may include salience, commitment, or a combination thereof termed resonant salience, where there is a burst of activity followed by a sustained commitment, or resonance, pattern. By combining temporal and network features in the analysis of the segments/clusters, chronotopes (i.e., patterns that exist in time and space) emerge.

[0233] In an embodiment, a timeline view may be used to examine messages across clusters. The timeline may include the chronotope as the drill down. For example, a primary timeline may be organized in rows by grouping of clusters (e.g., similar clusters are assigned together into a group). There may be several bands for groups (e.g., things for which there is a CFI score). The timeline may be examined for objects of interest that have very high CFI scores at some point. One example may be hash tags in a Twitter network. A dot may be placed at the point in time when the activity (attention) peaked (had the most citations, re-tweets, etc.) for that object of interest. A dot may be placed in the macro timeline for the group (showing the peak points of all objects of interest) where the peaks were for each group (a group corresponds to a band below the main timeline). When the dot that corresponds to the peak of attention to an object of interest for a group/cluster is clicked, the chronotope is revealed. The chronotope for that object of interest may appear in a window below the timeline. The timeline view may include time on the X axis and groups/clusters on the Y axis. Peak interest points for objects may appear as dots at points in time corresponding to the groups that have interest. Clicking on that object reveals the chronotope for that object for all of those groups.

[0234] Interacting with data in the chronotope view may reveal what the object of interest is. In some embodiments, a group of items may be selected at a time period for a certain cluster/group and a word cloud or semantic analysis of proper nouns that appear in those items may be assembled.

[0235] Social media sites enable users to engage in the spread of contagious phenomena: everything from information and rumors to social movements and virally marketed products. For example, Twitter™ has been observed to function as a platform for political discourse, allowing political movements to spread their message and engage supporters, and also as a platform for information diffusion, allowing everyone from mass media to citizens to reach a wide audience with a critical piece of news. Different contagious phenomena may display distinct propagation dynamics, and in particular, news may spread differently through a population than other phenomena. Described herein is a system for classifying contagious phenomena based on the properties of their propagation dynamics, by combining temporal and network features. Methods and systems described herein are designed to explore the propagation of contagious hashtags in two dimensions: their dynamics, that is, the properties of the time series of the contagious phenomena, and their dispersion, that is, the distribution of the contagious phenomena across communities within a population of interest. Further described is a method for simultaneously visualizing both the dynamics and dispersion of particular contagious phenomena. Using this method, particular contagious phenomenon chronotopes, or persistent patterns across time and network structure, may help emerge a taxonomy for contagious phenomena in general.

[0236] Given some contagious phenomenon p , p may be considered to have spread to user u the first time that u engages with p . For simplicity, engagement is measured as mentioning the phenomenon. For news, mentioning is likely a sufficient form of engagement, while for a political movement, stronger evidence of engagement may be preferable (contributing money, attending a rally, etc.). However, in social media sites, higher levels of mentioning often correlate with higher levels of engagement (e.g., users tweet about a political rally), while false indicators of engagement are rare: if a user wishes to mention a political movement to disagree with it, she will often not use a tag or specific name referring to that movement, but use a variant of it (e.g., a Twitter™ user who wants Vladimir Putin out of power may use the tag #Putinout instead of #Putin when tweeting about the prime minister and future Russian president). Therefore, the number of first mentions of p by users in some social media site is used as a proxy for the number of users that p has spread to.

[0237] In an embodiment, measures for characterizing contagious phenomena propagating on networks may include peakedness, commitment (such as by subsequent uses and time range), and dispersion (including normalized concentration and cohesion).

[0238] The peakedness of a contagious phenomenon is a scale-invariant measure of how concentrated that phenomenon is in time. A peak may be defined as a day-long period where total first mentions by day lies two standard deviations above the median first mentions. The specific duration of the peak window and the required deviation can be varied to maximize usefulness for particular kinds of phenomena and for particular social media networks. Median may be used instead of mean because, due to the skewed distribution of first mentions by day for most contagious phenomena, the mean is over-inflated. Contagious phenomena with short lifespans tend to have a sharp peak, when a large number of people mention the phenomenon, but the number of mentions is very small on either side of the peak. In contrast, long-lifespan contagious phenomena tend to grow slowly, with a less pronounced peak of mentions. The peakedness of a contagious phenomenon is the fraction of all engagements with that phenomenon that occur on the day with the most engagements with that phenomenon. A high peakedness means that most of the network's engagement with the phenomenon (e.g., for a social network, people in the network mentioning it) occurs within a short span of time, typically, hours to days. In contrast, a low peakedness means that the network's engagement with the phenomenon is spread over a long period of time, typically, weeks to months. Phenomena with high peakedness, such as news stories, may propagate rapidly through the network and then dissipate just as rapidly in the course of the daily news cycle. Phenomena with low peakedness may include popular websites and videos, which may maintain a slow but steady rate of engagement--individuals in the network are constantly discovering these phenomena, even as others get tired of them and stop engaging.

[0239] Commitment is the measure of the average scope of engagement with a particular contagious phenomenon by nodes in the network, or the staying power of a phenomena. Using the example of people engaging with online content in a social network, the commitment with a particular piece of online content can be the average scope of mentions of that content by pieces of the network. This measure would, for example, differentiate between a political movement that is just a fad, and one that accumulates a number of diehard supporters who keep the movement alive. Scope may be measured in at least two ways, which leads to the following two sub-measures: Commitment by Subsequent Uses and Commitment by Time Range. In social media sites, the cost in terms of time and effort to mention something for the second or third or tenth time is relatively small; therefore, for a second dimension, two quantities may be defined: first, the average number of subsequent mentions (all mentions excluding the first mention of the phenomenon by a user) of a contagious phenomenon among the adopting users; and second, the average time difference (in days) between first and last mention of the phenomenon among the adopting users. While the first measure, "Commitment by Subsequent Uses," is relatively easy to inflate by mentioning the phenomenon multiple times in a short period, the second measure,

“Commitment by Time Range”, indicates long-term commitment to mentioning the phenomenon by a set of users.

[0240] Commitment by Subsequent Uses is the average number of subsequent engagements with a phenomenon after a node's first engagement. For instance, if each person in a social network played an online game at most once, Commitment by Subsequent Uses for that story would be zero. In contrast, if just one percent of the people in a social network played an online game thirty times each, Commitment by Subsequent Uses for that game would be twenty-nine. Phenomena with high Commitment by Subsequent Uses may include online games, which encourage repeat engagements. Other phenomena with high Commitment by Subsequent Uses may include astro-turfed content, where a third party may encourage repeated interest in the content by paying or otherwise endorsing people who engage with it.

[0241] Commitment by Time Range is the average time period between the first and last engagement with a phenomenon by nodes in the network, measured over some large time window (e.g., a year). For example, if each person in a social network read articles on a blog ten times over the course of one day and never visited it again, Commitment by Time Range for that blog would be one day. However, if just one percent of the people in a social network read articles on a blog once every week for ten weeks and then abandoned it, Commitment by Time Range for that blog would be ten weeks. Phenomena with high Commitment by Time Range include blogs with loyal followers who keep coming back for more content. Phenomena with low commitment by Time Range include news stories that, on average, a person reads only once and never sees again.

[0242] In addition to measuring the dynamics of contagious phenomena (the properties of the time series of engagements with a phenomenon), the dispersion of contagious phenomena (the properties of distribution of a contagious phenomenon throughout a population) may be measured. Dispersion is a measure of the distribution of engagements with a contagious phenomenon over the network through which it propagates. Phenomena that are highly dispersed are broadly popular but may have less focused engagement from a particular group; phenomena that are not dispersed are not broadly popular, but may have focused engagement with a particular group. There are many ways of measuring the distribution of engagements with a phenomenon over a network, including the following two sub-measures: Normalized Concentration and Cohesion.

[0243] The Normalized Concentration of a contagious phenomenon presupposes a partition of the underlying network into discrete clusters, which usually represent communities. Given such a partition, the Normalized Concentration of a contagious phenomenon is the fraction of all engagements that come from the cluster that engages most with the phenomenon, or the Majority Cluster. For instance, if a social network were divided into two clusters, one of which engaged

with a particular news story nine times, and the other, only once, the Normalized Concentration for that phenomenon would be 0.9. However, if both clusters had engaged with the story five times, the Normalized Concentration for that phenomenon would be 0.5. Phenomena with high Normalized Concentration tend to be the cause célèbre of a particular community, e.g., political and social movements that have not gained wide traction. Phenomena with low Normalized Concentration may include headline news stories that touch many communities at once. Depending on the size of individual communities, Concentration may or may not correlate inversely with popularity.

[0244] In addition to Normalized Concentration, some aspect of the connections between the engaged users may be measured. For example, it is possible that a contagious phenomenon is widely spread across a number of communities, but diffuses only through strong ties so that the engaged users form a clique. Conversely, it is possible that a contagious phenomenon is confined to a single community, but spreads through weak ties and the engaged users are sparsely interconnected. Therefore, a measure of Cohesion may be defined as the network density over the subgraph on all users engaged in a particular contagious phenomenon. Contagious phenomena that spread over strongly connected sets of users will have a Cohesion close to one, whereas phenomena that spread over weakly connected sets of users will have a Cohesion close to zero. The Cohesion of a contagious phenomenon is the network density of the sub-graph of all nodes engaging with the phenomenon. The network density of a graph is the total number of actual connections between nodes in the graph divided by the total possible number of connections (usually $n*(n-1)/2$ for undirected graphs, where n is the number of nodes in the graph). For example, if only three people read a particular blog, but all those people knew each other, the Cohesion of that blog would be 1.0. In contrast, if ten people read a particular blog, but every one of those ten people knew exactly two of the others (the people were connected in a circle graph), the Cohesion of that blog would be $10/(10*9/2)=10/45=0.22$. Phenomena with high Cohesion may include stories and memes that propagate in an “echo chamber” of people who already know each other and engage with similar kinds of online content. Phenomena with low Cohesion include news and rumors that move between acquaintances, such that, for example, after multiple propagations, the person who hears the rumor and the person who started it may be total strangers.

[0245] In embodiments, phenomena with high Peakedness tend to have low Commitment, making those two measures a natural pair for comparing different online phenomena. For example, FIG. 18 depicts Commitment by Time Range on the Y axis and Peakedness on the X axis for two different sets of data depicted by different icons. In this example, the two datasets are: 1.) 112 Bundled hashtags relating to specific topics shown in red or as icon #1; and 2.) a baseline dataset of the top 500 hashtags for all users shown in black or as icon #2. The bundled hashtags display

a generally lower level of Commitment by Time Range than the top 500 hashtags at the same level of Peakedness. Some of the top 500 hashtags have extreme levels of Commitment, up to 150 days. Hashtags with the highest levels of Commitment are of several sorts, which notably include regional/location tags, tags for particular sports, religion tags (e.g., “Catholic,” “Jewish”), tags for particular news outlets, and general tags related to investing and financial markets. Intuitively, all of these are topics that might engage a stable set of users over a long time.

[0246] Referring to FIG. 19, and in an example, dealing primarily with topics related to Russia, peakedness is plotted for the bundled hashtags against both levels of Commitment: subsequent uses (FIG. 19a) and time range (FIG. 19b). In FIG. 19a, there are several distinct regions of the distribution. On the bottom right, hashtags with high Peakedness and low Commitment by Subsequent Uses are all directly related to salient news events, which in this case are the airport and metro bombings in Russia (#Domodedovo, #explosion, #metro29, #Moscow29). On the bottom left, hashtags with low Peakedness and low Commitment by Subsequent Uses are generally not very popular. Some of them are very generic (#moscow, #metro), and some just never had a peak nor became adopted by a committed user base. Some of these are tags that are similar to popular tags, but reflect less-used variations. On the top left, hashtags with low Peakedness and high Commitment by Subsequent Uses are all regional hashtags (with the exception of the Nashi hashtag that refers to a pro-government political youth movement in Russia). These regional hashtags were tangentially related to the forest fire events, but their main use is likely in talking about local affairs, hence the high commitment of a few users. Finally, on the top right, there are a number of hashtags with both high Peakedness and high Commitment by Subsequent Uses. These tend to be pro-government political hashtags (#iRu and #GoRu are both related to Medvedev's policy of modernization while #ruspioner and #seliger are both related to the Seliger youth camp). This observation suggests that pro-government political hashtags have some event (such as the Seliger camp) that is linked to a sudden burst of popularity, but subsequent to that event, users continue to include the hashtag in their tweets. This suggests that pro-government political hashtags may have “staying power” in the Russian Twitter™ community. Alternatively, or in combination with this, a committed set of users may use the pro-government hashtag both before and after the event, perhaps in an organizational or mobilizing capacity.

[0247] In contrast, and referring to FIG. 19b, some of the same clustering seen in FIG. 19a is depicted, where news is on the bottom right, regional hashtags are on the top left, but the top right group dominated by pro-government hashtags has moved down, indicating that these hashtags do not have staying power over long periods of time; they may be mentioned multiple times, but in a relatively short time range around the peak (days or weeks, not months). In contrast, the hashtags on the top right in FIG. 19b are the regional hashtag #Moscow and the political hashtag #Putinout.

(referring to the anti-Putin movement). It is important to note that #Putinout in particular has relatively long temporal staying power (an average of 50 days between first and last mention by a user in the dataset) but relatively short staying power by mentions (an average of less than six subsequent mentions).

[0248] Referring to FIG. 20 and FIG. 21, measures of dispersion of hashtags are analyzed across a core set of Twitter™ users. In FIG. 20, the distribution across nine topics of Normalized Concentration are plotted by hashtag within each topic. Comparing across all nine topics enables distinctive patterns to emerge; the minimum Concentration among pro-government hashtags in the Seliger and modernization topics is between 0.3 and 0.4. In contrast, the maximum Concentration among opposition hashtags in the Kashin and Russian Drivers' Movement topics, is between 0.4 and 0.5. Pro-government hashtags are on the whole more concentrated within one cluster than opposition hashtags. Hashtags related to news events, such as the Moscow Metro Bombing and the Domodedovo attack, tend to be diffuse, which is in line with the intuition that major news events tend to engage the population as a whole rather than specific communities.

[0249] In FIG. 21, the distribution across nine topics of Cohesion are plotted by hashtag within each topic. For ease of visualizing, the distribution plots are cut off at 0.2 and all hashtags with Cohesion > 0.2 are assigned a value of 0.2. Again, there is a contrast between opposition hashtags, which have extremely small Cohesion of 0.03 and below, and some pro-government hashtags (especially those in the Seliger and modernization topics), that have the much higher Cohesion of 0.10-0.30. Curiously, a few news-related hashtags have very high Cohesion, which suggests that some news-related hashtags may spread through strong ties.

[0250] FIGS. 18 through 21 provide a high-level analysis of hashtag diffusion among the Russian-speaking Twitter™ community, both from the temporal and the spatial (network) perspective. However, this analysis necessarily leaves out the idiosyncrasies of individual hashtags. Referring now to FIG. 22a, FIG. 22b, and FIG. 22c, chronotopes of the #metro29 (a), #samara (b), and #iRu (c) hashtags are depicted. In typical chronotope images, color indicates cluster group, and color brightness indicates volume of engagements. Detailed analysis of individual contagious phenomena enables crossing the dimensions of dynamics (loosely, temporal properties) and dispersion (loosely, spatial properties) of the latter. Therefore, spatiotemporal analyses of contagious phenomena, such as hashtags, may be constructed, and patterns in their diffusion across time and space may be discovered. Such patterns may be called the chronotopes of the hashtags. A chronotope is simply a pattern that persists across a spatiotemporal context, originally used in literary theory to describe genres or tropes.

[0251] In order to discover hashtag chronotopes, the diffusion of individual hashtags is visualized both across different communities and across time. First, a particular hashtag is selected and the

set of engagements of Twitter™ users with this hashtag is binned by day. Next, for each day, the volume of engagements for that day is broken down by cluster group. Finally, a grid where columns correspond to cluster groups and rows correspond to days is created. Each row-column cell of the grid is filled with a color corresponding to the cluster group. A cue as to the volume of engagements corresponding to a particular cell is given via the brightness of the color: the brighter the cell, the more engagements with a hashtag on that day from that cluster group. Black cells correspond to days when a particular cluster group has no engagements with the hashtag.

[0252] FIG. 22 shows three such visualizations: the #metro29 hashtag related to the Moscow Metro bombings on Mar. 29, 2010; the #samara hashtag related to the Russian city of Samara; and the #iRu hashtag, related to President Dmitri Medvedev's policy of modernizing Russia. These three visualizations display three distinctive patterns across space and time; #metro29, in FIG. 22a has a "salience" chronotope, with engagements across the spectrum of cluster groups during the week around March 29. In contrast, #samara in FIG. 22b has a "resonance" chronotope, with consistent engagements from the local cluster group, presumably residents of Samara talking about their city. Finally, #iRu in FIG. 22c has a "resonant salience" chronotope, with an initial cross-group burst of activity in late November 2010 (around the time of Medvedev's announcement of his new policies), followed by consistent engagements from the Pro-Government cluster group over the next month. Note that FIG. 22 does not contrast with FIG. 19, which suggests that pro-government hashtags have low staying power, but instead presents a more subtle picture; the cluster group of pro-government users remains active in the #iRu hashtag over the course of a month, but, as FIG. 19b indicates, individuals within that cluster rarely carry on with adoptions for more than 5 days. There may be a high turnover of users of the #iRu hashtag, with new enthusiasts coming in even as the original adopters lose interest in the topic.

[0253] In embodiments, phenomena with the Salience Chronotope tend to have high Peakedness and low Commitment, while phenomena with the Resonance Chronotope tend to have low Peakedness and high Commitment by Time Range. Phenomena with the Resonant Salience Chronotope tend to have both high Peakedness and high Commitment by Time Range.

[0254] In an embodiment, a flexible algorithm may be used for optimizing a targeted network influence campaign. For example, a user may have a high CFI score, but they may not message their social networks frequently, thus targeting these individuals may not optimize the campaign. The algorithm may output an M Score, which may be calculated from a CFI score plus some other network or behavioral metric. In embodiments, wherever it is described to use the CFI score, the M score may instead be used to maximize campaign effectiveness. In embodiments, the M score may be an interpolation of the number of followers of the target item (influence) and the CFI score of the target item (specificity). This mathematical calculation may result in a normalized score

on a scale, such as a scale from 1 to 10 where 1 is low impact and 10 is high impact. Thus, the M score is a general measure of influence and specificity.

[0255] One way to calculate the M score is to combine CFI and count, where count is the overall number of members on the map that have engaged with that target, in a formulaic way. The formula is $M \text{ score} = \text{count}(\alpha) + \text{CFI}(1 - \alpha)$ [normalized 1 to 10].

[0256] In embodiments, the M-score may be user-tunable, so that there is a choice to prioritize “segment specificity” vs. “global footprint,” and/or “network position” vs. “behavioral profile” (e.g., someone who retweets frequently) when selecting behavioral and/or network metrics to calculate the M score. In an embodiment, for example, a slider 2902 may be provided to users so that can select a target that is more niche or more global. The M score enables optimizing a campaign on network position or on behavior. If the slider is dragged towards “niche,” alpha approaches zero and the M score is near equivalent to just the CFI score of the target item (high specificity). If the slider is dragged towards “broad,” alpha approaches 1 so that the M score is near equivalent to just the number of followers of the target item (high influence). Setting the slider somewhere in between “niche” and “broad” allows users to tune the set of individuals/entities that they want to target.

[0257] In an embodiment, direct ad placement may be enabled by CFI scores/M scores. Using CFI scores and/or M scores, a list of targets/websites may be created and ads may be placed directly on the target/website via integration with various products, such as Twitter™ sponsored tweets, Facebook™ ad exchange, Google™ AdSense/Adwords, third party online ad networks, and the like.

[0258] Referring now to FIG. 24, a recent activity page of a social media map platform provides recent activity, such as new followers, new influencers following the user, an indication of any re-tweets including the number of people who have retweeted an item, changes to the user's cluster groups with links to respective group overview screens, a list of new influencers including their cluster group and their number of followers, the current conversation leaders including their cluster group and their number of followers, a view of all media being shared in the network including the latest influential media and the segments in which the media is influential, links to an overview page, links to a lists page, links to a help and support page, and the like. The user may continue to their map from this screen. Graphics, such as a bar graph, may be included in the changes to the user cluster groups box to indicate the number of users in each cluster group. Graphics, such as a bubble chart, may also be included in the media box to indicate the size of the segments in which the displayed latest media is influential.

[0259] Referring now to FIG. 25, another example of a recent activity page of a social media map platform is shown. In this example, new followers are shown; included in the number of followers

are new influencers and group changes, including a percent change for each cluster group, information on new influencers, such as their name, handle, number of tweets, number of followers, number of people they are following, and a button to message them or follow them. Also on this page are trending terms/URLs, including the number of mentions of a hashtag that is related to the user, trending media and imagery, and latest influencer tweets. Icons may be provided to reply, retweet, favorite a tweet, share or embed a tweet, and the like.

[0260] Referring now to FIG. 26, an overview page is shown. The overview page includes a table of cluster groups, the number of members in the group, the power of the cluster, and the tweet activity. A power score is an indication of which segment is worth engaging with and may be an indication of which segments are most dense and represent the greatest signal of interest. In one embodiment, power may be calculated based on network density: the number of connections divided by the number of possible connections. In another embodiment, power is calculated based on coordinates, such as the average distance from the center of a cluster map. In another embodiment, power may be calculated as the average distance from the centroid of the cluster that emerges in the clustering computation. In embodiments, power is like the segment/cluster version of the M score.

[0261] Continuing with the page on FIG. 26, an individual cluster may be selected and a representation of that cluster in a map may be highlighted. For example, the UK design cluster has been highlighted and a dialog box appears showing more information about the individual group, including number of members and graphics depicting the power and tweet activity associated with the group. When the user clicks the "Read more" link, a box may appear with more information. The map and group information items may remain visible when the page scrolls such that they are in a fixed position. Selecting clearer on the page overview causes the selected row to be cleared and makes all map nodes visible. An alarm icon on the overview page allows the user to review all recent activity including number of tweets from various members of the network. Selecting "view full-screen map" will send the user to a screen such as that shown in FIG. 27. Referring now to FIG. 27, a full-screen map is displayed. In this map, the international cluster has been selected and the South America sub-cluster was selected. The colored nodes in the map may indicate one or both of the selected clusters and sub-clusters. The influencers in a particular sub-cluster may be viewed and when an influencer is selected, the URLs associated with that influencer there may be shown. A node overview may appear including the influencer name, their handle, their location, their URL, when they joined the social network, their number of tweets, their number of followers, the number of people they are following, the groups they are linking in, the number of in-links in each group, as well as any other relevant information.

[0262] Referring now to FIG. 28, an embodiment of an overview page is shown. In this view, a segment or cluster has been selected and data regarding that segment is displayed, such as key influencers, current conversation leaders (mentions), an interactive map, key photos and videos or other media, key tweets/retweets, key websites, key content, latest conversation terms, and the like. Effectively, this page shows an enhanced version of cluster-focused data and makes it more accessible. The power score for the segment is displayed as well as an icon from which the user may take certain actions such as build their network, find content, find media, find tweets, message followers, launch a Twitter™ campaign, launch a Facebook™ campaign, launch a mobile campaign, launch a social media campaign, launch an AdWords campaign, launch an advertisement campaign, and the like. The overview page may be a user interface. Notifications of certain data and data presentation may be made in the user interface, for example, which may be implemented by software and embodied in a tangible medium, such as a mobile device, smartphone, tablet computer, or the like. The user interface may be a touchscreen embodiment, such that to utilize the user interface, a user is required to touch the screen of the device displaying the user interface. The user interface may be accessible on different computing devices and capable of dynamically accessing user specific data stored on a network server and/or local device.

[0263] Referring now to FIG. 29, the “influencers” tab has been invoked. Various ways to filter the influencers are provided such as by follower status (all followers, follows the user, does not follow the user) or by following status (show all, the user follows, the user does not follow). Another way to filter influencers may be by M score, follower count, mentions, name, screen name, and the like. One way to filter by M score is by use of a slider 2902 to obtain more niche or broader individuals/entities as described elsewhere herein. Another way to filter individuals/entities may be by their exposure to particular content. By utilizing this filter, the user may target individuals/entities who have not already been exposed to the content. Users may take action from this page such as to follow selected individuals/entities, save individuals/entities to a Twitter™ list, create a new list, add a selection to a list, send a direct message, send a sponsored tweet, and the like. When saving individuals/entities to a Twitter™ list, a dialog box may appear with list choices for the user, such as a list for my influencers following me, a list for my influencers and not following me, a branding group, and the like. In this example, one action being taken is to follow seven new users. By following individuals/entities and engaging in behaviors that might cause them to be aware of the user, the users network may potentially expand to include the newly followed individuals/entities. Another action that is taken is to compose a message. The compose message screen may include suggested content such as most used hashtags or other media based on a CFI, popular terms, key content such as high M score media, and the like. Influencer information may be leveraged in determining whom to message. The suggested

content may be filtered by the exposure of target individuals/entities to the content. Data related to the content such as its peakedness, first appearance, and the like may be exposed to the user so that the user can decide whether it makes sense to share the content with other individuals/entities. Referring to FIG. 30, users may be able to drill down to the individual influencer level to see in what other segments/clusters the individual is influential, their latest tweets, M score, number of tweets, number of followers, number following, footprint, following/follower status with respect to the user, demographic information, URL, and the like. Icons may be available to follow, act (i.e., add the person to a list, retweet their latest tweet, send a direct message, etc.), view a social media profile, and the like.

[0264] Referring now to FIG. 31, a tab for conversation leaders is displayed. Various ways to filter the conversation leaders are provided such as by follower status (all followers, follows the user, does not follow the user) or by following status (show all, the user follows, the user does not follow). Another way to filter conversation leaders is by peak date such as all, today, past week, past month, custom date range, and the like. Another way to filter conversation leaders may be by M score, follower count, mentions, peak, peakedness, name, screen name, and the like. Another way to filter conversation leaders may be by their exposure to particular content. By utilizing this filter, the user may target individuals/entities who have not already been exposed to the content. Users may take action from this page such as to follow selected individuals/entities, save individuals/entities to a Twitter™ list, create a new list, add a selection to a list, send a direct message, send a sponsored tweet, and the like.

[0265] Referring now to FIG. 32, a tweets tab is displayed. The tweets may be filtered by peak date such as all, today, past week, past month, custom date range, and the like. The tweets may be filtered by M score, re-tweets, original postdate, peak, peakedness, name of poster, screen name of poster, and the like. One way to filter by M score is by use of a slider to obtain an audience that is more niche or broader, as described elsewhere herein. Data regarding each displayed may include an M score the number of influential re-tweets, the number of retweet, the posted date, the peak date, a graphic of the peak pattern, icons with which to take action such as reply/retweet/favorite, name, screen name, and the like. Selecting one of the tweets may cause a drill down box to appear with additional information about the individual/entity who made the tweet, such as M score, number of tweets, number of followers, number following, footprint, number of friends, follower/following status, demographic data, URL, which segments the individual/entity is retweeting in, who have they been retweeted by, icons to social media profiles, icons with which to take actions such as reply/re-tweet/favorite/add to list, and the like.

[0266] Referring now to FIG. 33, a websites tab is displayed. The websites can be sorted by mentions, M score, subpages mentioned, hostname, and the like. One way to filter the websites

by M score is by use of a slider to obtain an audience that is more niche or broader, as described elsewhere herein. Users may take action from this page such as to buy an ad, create a new list, add a selection to a list, and the like. Selecting a website reveals a drill down box for the website. Information about the website in the drill down box may include M score, distinct mentions, mentions, subpages mentioned, excerpt, peak date, a graphic of the peak pattern, segments/clusters the website is mentioned in, who mentioned the website, latest tweets mentioning this URL, a button to take action, and the like.

[0267] Referring now to FIG. 34, a tab for key content may be displayed. Information about the key content included in this view includes the name of the website, name of an article, URL, peak date, a peak pattern, M score, citations, distinct citations, and the like. The key content may be sorted by M score, citations, peak, peakedness, host name, content title and the like. One way to filter by M score is by use of a slider to obtain an audience that is more niche or broader, as described elsewhere herein. The key content may be filtered by peak date such as all, today, past week, past month, custom date range, and the like. Users may take action from this page such as to compose a message, compose a tweet, view a drill down box for the key content, and the like. In the compose message or compose Tweet view, users may be able to select one or more individuals/entities or and influencers/conversation leaders to message with suggested content (most used hashtags, popular terms, key content, etc.). In one embodiment, the individuals/entities may be part of a list such that either certain members of the list or the entire list may be easily included as recipients of the message. Selecting a key content reveals a drill down box for the content. Information about the content in the drill down box may include name of website, title of article, M score, distinct mentions, mentions, subpages mentioned, excerpt, peak date, a graphic of the peak pattern, segments/clusters the content is mentioned in, who mentioned the content, latest tweets mentioning this URL, most used hashtags, a button to take action (tweet this, use in direct message, add list, etc.), and the like.

[0268] Referring now to FIG. 35, a media tab is displayed. Media may be filtered by images, videos, audio, GIFs, and the like. The media may be filtered by peak date such as all, today, past week, past month, custom date range, and the like. The media may be sorted by M score, citations, peak, peakedness, host name, content title and the like. Information about the media in this view may include title, duration, media type, M score, mentions, distinct mentions, peak date, peak pattern, and the like. By selecting one of the media items, a drill down box may appear. Information in the drill down box may include title of media, URL, M score, mentions, distinct mentions, peak date, peak pattern, media type, duration, what segments/clusters the media is mentioned in, most used hashtags, who has mentioned the media, latest tweets mentioning this media, an icon to take action with, and the like.

[0269] Referring to FIG. 36, a tab for terms is displayed. The terms may be filtered by hash tags, one word, 2 words, 3 words, and the like. The terms may be filtered by peak date such as all, today, past week, past month, custom date range, and the like. The terms may be sorted by M score, citations, peak, peakedness, host name, content title and the like. Information about terms in the list may include the term, peak date, peak pattern, M score, mentions, distinct mentions, and the like. Selecting a term may reveal a drill down box where additional information about the term may be displayed including which segments/clusters the term has been mentioned in frequently, what other terms have been mentioned with the selected term, who has mentioned the term, latest tweets mentioning this term, an icon to take action with, and the like.

[0270] Referring now to FIG. 37, a list page of a social media map platform is displayed. In this view, information may be provided in the form of lists, such as lists of influencers, conversation leaders, key content, terms, and the like. Information about each list member may include name, screen name, M score, followers, mentions, follower/following status, and the like. Lists may be sorted/filtered by any of the techniques mentioned herein including by influence, M score (such as with a slider or other user input), and the like. Users may take action from the list view.

[0271] In further embodiments, an analytical framework for a coordinated campaign identification includes proposing a framework for analyzing fabricated social movements. In many embodiments, not only is there the ability to distinguish these movements from truly organic ones, there is also the ability to create a formal method for studying patterns of fabricated, pseudo-grassroots (also, "astroturf") collective action.

[0272] It will be appreciated in light of the disclosure that any such collective action may be required to give the impression of a large group of people coalescing around a movement that is easy to describe and share with others. If the group is not well-connected enough, then it may be logistically difficult for any actor to organize the group's online behavior. If the group is not acting in temporal lockstep, then its message may not achieve a high frequency. In embodiments, low-frequency messages do not appear as global trends; for example, Twitter's "trending" algorithm appears to identify topics that are popular now, rather than topics that have been popular for a while or on a daily basis, to help you discover the hottest emerging topics of discussion on Twitter™. The many examples remain applicable to the myriad social platforms. Finally, if the group behind a fabricated social movement does not promote it with a coherent message, the movement's impact on the general public may be blunted by conflicting information.

[0273] It will be appreciated in light of the disclosure that these constraints suggest a natural set of three dimensions for analyzing fabricated social movements: 1.) the semantic dimension (how messages are formulated), 2.) the network dimension (how accounts within the campaigns are connected to one another), and 3.) the temporal dimension (when messages spread throughout the

campaign). In many embodiments, these dimensions, and their intersections, yield discrete signals that can be used to scrutinize social media operations and assess if they display a suspicious degree of hidden coordination.

[0274] In embodiments, the framework operates on three levels: 1.) Event, the level of an entire social campaign; 2.) Segment, the level of a community of users participating in a social media campaign (e.g., Russian social media troll accounts), and 3.) Actor, the level of an individual user participating in a social media campaign.

[0275] Table 1 below shows examples of the three-dimensional analysis framework in more detail specifically, the signals relevant for particular combinations of level and dimension. It will be appreciated in light of the disclosure that not every combination of level and dimension has corresponding relevant signals.

	Network	Temporal	Semantic
Network	Event: how concentrated is online participation in the movement? Does it cover a broad range of politically / socially / culturally distinct communities, or is it contained in a homogeneous “echo chamber”? Segment: do communities of actors who participate in the movement pay disproportionate attention to each other? Actor: do actors who participate in the movement do so in conjunction with their communities, or independently of them?	Segment: How does participation in the movement vary between different communities and over time? Are particular communities always lagging behind the rest in participation (taking time to formulate a response)? Actor: how long does the average actor participate in the movement?	Segment: How topically diverse is all discourse among communities participating in the movement?
Temporal		Event: Does participation in the movement follow an unusually temporally regular pattern, when compared to spontaneous human posting behavior? Segment: do communities of actors coordinate their activities, even across time zones? Actor: Do some actors behave similarly to pre-identified troll or spambot accounts with regard to their temporal posting patterns?	Event/Segment/Actor: How does the diversity of the discourse among all participants / specific communities / individual actors participating in the movement vary over time?
Semantic			Event/Actor: How topically diverse is the discourse around the movement among all actors / individual actors?

Table 1. Three-Dimensional Analysis Framework

[0276] This framework is a helpful methodological tool, but it would not be useful without operational definitions, which are captured via mathematical metrics of campaign activity. In embodiments, each signal in Table 1 above is mapped to a discrete metric in Table 2. Further detail regarding key definitions for understanding these metrics, and any non-obvious activity metrics are provided herein.

Table 1 Row	Table 1 Column	Table 1 Level	Metric
Network	Network	Event	Entropy E
Network	Network	Segment	Inter-community homophily H
Network	Network	Actor	% of actor's community participating in campaign, by number of individuals or total posts
Network	Temporal	Segment	Time delta between peak date of campaign participation by segment
Network	Temporal	Actor	Commitment by actor M
Network	Semantic	Segment	Semantic Diversity by Segment Ω_s
Temporal	Temporal	Event	Campaign Peakedness P
Temporal	Temporal	Segment	Dynamic Time Warp alignment between Segments DS
Temporal	Temporal	Actor	Dynamic Time Warp alignment between Users DU
Temporal	Semantic	Event/Segment/Actor	Semantic Diversity over time by Event / Segment / Actor $\Omega_{TE}, \Omega_{TS}, \Omega_{TA}$
Semantic	Semantic	Event/Segment/Actor	Semantic Diversity by Event / Actor Ω_E, Ω_A

Table 2. Mapping of Signals to Metrics

Key Definitions

Network

[0277] In many embodiments, the network dimension assumes that actors participating in a campaign are connected to each other in a directed network G (i.e., a connection from user a to user b does not imply the reverse). Twitter™ following networks are an example of directed

networks: many people follow Twitter™ celebrities, but those celebrities do not follow their fans back as a general rule. Other social media platforms and connected platforms are applicable.

Segment

[0278] When calculating metrics at the network level, it is assumed that each actor participating in a campaign belongs to exactly one community c , where c represents a group of actors with similar interests, whether social, political, or otherwise.

Identifying Networks and Communities

[0279] In order to identify relevant networks and communities within those networks, network segmentation technologies are leveraged such as hierarchical agglomerative clustering. In many examples, it may be shown that network segmentation framework, based on hierarchical agglomerative clustering has been tested on more than eight hundred different sociocultural contexts with many academic applications. By way of many examples, the unit of analysis is a “map,” which may be a collection of key social media accounts around a particular social context. A map may be composed of “nodes,” which are the social media accounts in question. Each node may be connected to one or more nodes in the map through “edges” and edges may represent social relationships embedded in the respective social media platform (e.g., “following” for Twitter™, Facebook™, or the like).

[0280] In embodiments, each node in the map may belong to exactly one “segment” and one “group.” By way of these examples, a segment may be a collection of nodes with a shared pattern of interests (e.g., a collection of Twitter™ accounts who all follow US Tea Party politicians). Each segment may have a label (e.g., “Tea Party”). A group may be a collection of segments with similar interest profiles (e.g., a collection of “Tea Party,” “Constitutional Conservatives,” etc. segments into a “Conservative” group). The process for generating segments, groups, labels, and colors for a map be fully or partially automated, as follows: a proprietary clustering algorithm may automatically generate segments and groups for a map; subsequently, the map-making process may use supervised machine learning to generate labels for segments and groups from human-labeled examples. At the end of the automated process, a Subject Matter Expert, an individual well-versed in the topic and/or geographical area covered by the map, may perform a quality assurance check on the segment and group labels.

Key Metrics Explained

[0281] To illustrate metrics in this section, a toy campaign example may be employed. The example consists of 100 users connected in a network G . The network G further breaks down into exactly two communities A and B , each with exactly one half of the total population. The overall number of connections from members of A to any other actor in the network is 500, while the number of connections from members of A to members of B is 200. The campaign proceeds over

the course of ten days, and the first of those days features the highest level of campaign activity, with exactly one quarter of all actors participating.

Entropy E

[0282] This metric is the degree to which a particular campaign is concentrated in one community versus diffused among many different communities. Given a mapping of users to Communities, which is described in more detail below, the entropy of a campaign may be, as known in the art, the information theoretic entropy of the distribution of users active in the campaign among different communities. In the toy example, the Entropy of the campaign may be:

$$E = - \sum_{i=1}^{|c|} p(c(i)) \log_b(c(i)) = -0.5 \log_2(0.5) - 0.5 \log_2(0.5) = 1$$

In general, it may be shown that low values of E represent campaigns concentrated in one community, while high values of E represent campaigns distributed among a wide array of communities.

Inter-community Homophily H

[0283] It is known in the art that the inter-community Homophily H is the degree to which communities active around the campaign are more interconnected than one would expect by random chance. Mathematically, H is calculated for an ordered pair of communities A, B . The quantity $H(A, B)$ is the ratio of the actual number of connections from members of A to members of B , $E(A, B)$, to a normalizing factor ρ that assumes that members of A make their connections to all other nodes at random. In the random baseline, the number of connections from members of A to members of B is the number of all connections from members of A to any other node in the network G , multiplied by the fraction of G that B represents. In the toy example, the Homophily from community A to community B is:

$$H(A, B) = \frac{E(A, B)}{\rho} = \frac{200}{500 * 0.5} = 0.8$$

[0284] Values of H below 1.0 may be shown to represent heterophily, or lower-than-expected interconnectivity between communities. Values of H equal to 1.0 may be shown to represent the baseline random expectation. Values of H above 1.0 may be shown to represent homophily, or higher-than-expected interconnectivity.

[0285] H is superlinear, so a value of $H = 4.0$ is much more than twice as interconnected as $H = 2.0$.

[0286] While the random baseline for Homophily is established in the citation above, it will be appreciated in light of the disclosure that it may be an excessively low baseline for such empirical analyses. Therefore, when possible, H values are used for community pairs where there may be expected low / high values (e.g., ideologically separate / ideologically aligned communities) in the same networked terrain as the case study as a baseline.

Commitment M

[0287] Commitment to a particular campaign is measured in two ways: 1.) M_e , the number of subsequent engagements with the campaign by an actor; or 2.) M_r , the length of time between first and last recorded engagement with the campaign by an actor.

Semantic Diversity Ω

[0288] Semantic diversity of a particular actor's / segments / campaign's messaging is based on the assignment of messages to topics. As known in the art, LDA is a common method for identifying topics in text data. Once messages have been assigned to topics, a semantic diversity score may be calculated for the message set. The authors of the referenced work may represent their measure of semantic diversity as the probability that two documents chosen from the corpus at random with replacement will be on the same topic. By way of these examples, the corpus may be the message set, and the documents may be user Tweet histories, post histories, etc. aggregated by user. In many examples, the LDA algorithm may run for 15 iterations, with a number of topics no less than 20% of the number of documents and no more than 30 iterations and may average semantic diversity over 20 distinct runs of the LDA algorithm on the same corpus to smooth out variations due to the initial conditions for a particular run. For topics that do not co-occur in documents, a topic may be assigned a distance score of 1000.

[0289] In embodiments, versions of Ω are run for individual users (Ω_A), communities (Ω_C), or entire campaigns (Ω_E). These metrics can also be run for all messages within a particular time period (Ω_{T*}) to calculate the change in semantic diversity over time.

[0290] Semantic diversity scores of less than one may represent users who exclusively post about the same topic, characteristic of fabricated campaigns. Semantic diversity scores between 1 and 100 may represent users who post on a variety of topics, characteristic of normal human activity. Finally, semantic diversity scores above 100 may represent users who post on an extremely diverse set of topics, characteristic of spambots or users who bridge different cultural and/or linguistic communities (e.g., users who post in different languages, etc.)

Campaign Peakedness P

[0291] Campaign Peakedness may be defined as the fraction of all activity that occurs in the day with the most campaign-related activity during some time frame. In the toy example, $P = \frac{1}{4} = 0.25$.

Dynamic Time Warp Alignment D*

[0292] The Dynamic Time Warp is an algorithm known in the art for comparing two temporal sequences of activity. In the many embodiments, the Dynamic Time Warp may be used to compare the activities of individual users (*DU*) or entire segments (*DS*). In general, the Dynamic Time Warp between two sequences *S1* and *S2* is the number of warping transformations that are required to change *S1* into *S2*. In many examples, Dynamic Time Warp may be used to identify bots and trolls in a different social media setting.

[0293] In many examples, this framework and these metrics on eighteen case studies of political campaigns have been tested in seven different sociocultural settings, spanning three continents and six years in all. These studies included ten groups of Twitter™ hashtags linked by subject matter experts (SME) to known coordinated campaigns, and eight groups of Twitter™ hashtags linked by SMEs to known spontaneous campaigns. Based on the eighteen case studies, it may be shown that clear differences between coordinated and spontaneous campaigns across sociocultural setting and time for four of the metrics listed above: Entropy *H*, Commitment by subsequent engagements *ME*, Time delta, and Peakedness *P*. The same analysis also showed that at least one especially coordinated campaign showed extremely low values of Semantic Diversity by Event Ω_E and high Dynamic Time Warp alignment *DS* between the activity of different segments.

[0294] In further embodiments, methods and systems are disclosed for identifying markers of coordinated activity in social media movements that may identify a large number of accounts that may be controlled by a small number of coordinated entities that may result in a measurable lack of diversity of a similar number of accounts controlled by uncoordinated individual actors. To facilitate the methods and systems of identifying markers of coordinated activity in social media movements, a framework of signals (or metrics) along at least three dimensions may be constructed and may include, without limitation:

[0295] A Network dimension that may, for example, represent how accounts are connected;

[0296] A Temporal dimension that may represent, for example, patterns of messaging across time; and

[0297] A Semantic dimension that may represent, for example, diversity of topics and meaning.

[0298] From this framework, a plurality of hypotheses may be derived for “signals” exploring potentially hidden coordination on social media movements on a social media channel such as Twitter™, Facebook™ or the like. The exploring potentially hidden coordination on social media movements on a social media channel may occur at the level of the entire campaign (e.g., nine signals), a cluster level of the campaign (e.g., a set of well interwoven accounts), at the individual account level, and the like. In embodiments, the plurality of hypotheses may include twenty-five or more such hypotheses. Empirical evidence associated with these signals can be shown across

a number of case studies of known coordinated (i.e., inorganic, centrally-controlled) and spontaneous (i.e., organic, individually) campaigns. In embodiments, three of the campaign signals may systematically reveal coordination in social media movements on Twitter™, Facebook™ and other platforms. Some signals, either at the cluster or at the individual account level, may facilitate campaign analysis, and some of them may be transformed into campaign-level signals.

[0299] Campaign / Cluster / User – Each campaign may include a set of “seeds” from a specified timeframe that may be, for example, a hashtag, a sentence shared in posts, a URL shared in posts, or the like. In embodiments, clusters may be communities of users active within the campaign. In embodiments, users may be defined by their individual accounts, defined by their Twitter™ handle, Facebook™ identification defined by their user name on other social media platforms, or the like.

[0300] Network Terrain – Campaigns may occur in a specific context referred to as the “network terrain.” In one example, it will be appreciated in the light of the disclosure that the #BlackLivesMatter movement may be better analyzed within its “network terrain,” which displays the US political conversation on Twitter™, Facebook™ or other relevant social media platforms. In a representative model, social media platforms like Twitter™, Facebook™ may constitute a cyber-social “network terrain” formed by the relationships (such as following in Twitter™, Facebook™, or the like among actors. The structure of the network or social media platform may determine who and what may be visible to whom, and thus it may be the social landscape on which the struggle for influence may occur. The methods and systems may include analyzing case study campaigns across specific network terrain maps in order to understand the relationships between participants and the patterns of campaign propagation across specific online communities (e.g., clusters or clusters discovered using machine learning analysis of network relationships and the like).

[0301] Campaign versus Investigatory Signals – Signals measured at the cluster and individual actor (user) levels may facilitate investigating the inner workings of specific campaigns, building a more qualitative understanding of how these campaigns unfolded, and helping form campaign level metrics among other things.

[0302] Case Studies – To date, the methods and systems may include testing signals set on a set of case studies and exemplary campaigns.

SIGNAL SUMMARY

[0303] Exemplary Investigatory Signals – The investigatory signals may operate at the cluster or at the individual level. The investigatory signals may facilitate building a qualitative understanding of the dynamics of a campaign and may provide tools to build campaign-level

signals. [C] indicates a signal operating at the cluster level, and [U] indicates a signal is operating at the user level.

[0304] The following are exemplary priority signals:

[0305] Concentration in Lead Cluster [C];

[0306] Concentration via Entropy [C];

[0307] Day – peakedness [C];

[0308] Temporal coordination per cluster [C];

[0309] Temporal coordination per user [U];

[0310] Client diversity per cluster [C]; and

[0311] Time delta between clusters [C].

[0312] Other signals include:

[0313] Commitment by user [U];

[0314] Commitment by cluster [C];

[0315] Account creation date diversity for cluster [C];

[0316] Homophily [C];

[0317] Language mismatch [C];

[0318] Russian language profile % [C];

[0319] % in cluster also active [C];

[0320] % of hits in own cluster [C];

[0321] Account creation date diversity [C];

[0322] Semantic diversity by user for user Tweets™ (or other postings) [U];

[0323] Semantic diversity by time slice by cluster [C]; and

[0324] Semantic diversity by time slice by user [U].

[0325] In embodiments, a priority signal name is Concentration in Lead Cluster.

[0326] The concentration in lead cluster signal description – Large-scale spontaneous campaigns may be more likely to engage participants from a range of different clusters, whereas coordinated campaigns are typically highly concentrated in a specific cluster of the network or social media platform. The concentration in lead cluster signal (metric) evaluates the degree to which an entire campaign's activity is concentrated in a particular cluster of participants. The concentration in lead cluster signal (metric) may measure by the fraction of all campaign participants who are members of the most campaign-active cluster in the network terrain map.

[0327] The range of score value range of the concentration in lead cluster signal (metric) is zero to 100%. In embodiments, the concentration in lead cluster signal (metric) value is computed by determining the value of the concentration of the fraction of a campaign's participants that are members of the most active community in the campaign. In an example including a 3-community

map, if 50 participants are from community A, 25 from community B, and 25 from community C, then the value of the concentration in lead cluster signal (metric) for the campaign on this map equals 50%. In embodiments, possible values of the concentration in lead cluster signal (or metric) may be between 0 (i.e., not concentrated) and 100% (i.e., fully concentrated in 1 cluster).

[0328] The concentration in lead cluster signal (or metric) may be consistent across a set of campaigns, which may cover a variety of geographies and dates. It will be appreciated in light of the disclosure that coordinated campaigns, on average, may be shown to have larger values of the concentration in lead cluster signal (or metric) than those of spontaneous campaigns. It will also be appreciated in light of the disclosure that there may be some overlap between the coordinated and spontaneous ranges due at least in part to a large number of sociocultural settings and time periods in the data sets.

[0329] An exemplary average value of the concentration in lead cluster signal for coordinated campaigns is 48%.

[0330] An exemplary range of values of the concentration in lead cluster signal score for coordinated campaigns is 20% to 89%. The range here is the full range between the lowest value and the highest value for this category in the campaign.

[0331] An exemplary value of the standard deviation of the concentration in lead cluster signal for coordinated campaigns is 0.21.

[0332] An exemplary average value of the concentration in lead cluster signal for spontaneous (organic) campaigns is 22%.

[0333] An exemplary range of values of the concentration in lead cluster signal score for spontaneous campaigns is 9% to 50%.

[0334] An exemplary value of the standard deviation of the concentration in lead cluster signal for spontaneous campaigns is 0.12.

[0335] In embodiments, the performance of the concentration in lead cluster signal (metric) may be sensitive to the specific terrain map being used because the signal (metric) may be less successful if the terrain map used only captures the active participants in a campaign. The concentration in lead cluster signal (metric) may be more successful when capturing the broader terrain in which the campaign under scrutiny unfolds.

[0336] The methods and systems described herein also include computing the value of the concentration in lead cluster signal (or metric) using actions rather than users and may measure what proportion of the total actions (Tweets™ or the like) in the campaign that came from the most active community. This approach can be shown to be reliable because heavy posters (those who Tweet™ or the like) may create skews in the measurements.

[0337] In embodiments, a priority signal name is Concentration via Entropy.

[0338] The concentration via entropy signal description – The concentration via entropy signal is another approach to measuring concentration that looks at how the participants are distributed among the active communities in the campaign rather than simply looking at how many of them belong to the most prevalent community. The concentration via entropy signal (metric) may be shown to be a useful signal for knowing if more than one community is driving a coordinated campaign, which could be missed relying on the concentration in lead cluster signal (metric) alone. The concentration via entropy signal (metric) may calculate the concentration of distribution among all clusters. In embodiments, coordinated campaigns generally tend to have values of the concentration via entropy signal (metric) that are less than 2.0.

[0339] The concentration via entropy signal value range – Relatively higher values of the concentration via entropy signal (metric) reflect more even distributions of participants between the communities active in the campaign. The lowest score is zero (all participants belong to the same community). The highest score depends on the number of communities active in the map. Because the highest number of communities in an exemplary case study map may be 50, the highest entropy value in this example would be four (assuming a perfectly even distribution of participants amongst the 50 communities).

[0340] How the concentration via entropy signal is computed – The concentration via entropy signal (metric) may be an entropy of the distribution of participants among communities. In an example with a two-community map, the value of the Concentration via Entropy signal would be 1.0 when 50 participants are from community A, 50 participants are from community B, and thus the distribution would be 0.5, 0.5.

[0341] Exemplary formula for the concentration via entropy signal (metric):

$$E = - \sum_{i=1}^{|c|} p(c(i)) \log_b(c(i))$$

[0342] In the formula, $c(i)$ is the count of participants in the i th cluster and $p(c(i))$ is the fraction of all participants coming from the i th cluster.

[0343] In embodiments, the concentration via entropy signal (metric) is based on a logarithmic scale, so a small difference in entropy belies a large difference in the unevenness of the underlying distribution. It will be appreciated in light of the disclosure that a very rough rule of thumb is that a difference of one point in the value of the concentration via entropy signal may be equivalent a change in concentration by a factor of three, so a campaign with the concentration via entropy signal equal to two is three times more concentrated in a few clusters than a campaign with the concentration via entropy signal that is equal to three.

[0344] Analysis in case studies – The concentration via entropy signal (metric) can be shown to

be consistent across campaigns despite the variety of geographies and dates. It will be appreciated in light of the disclosure that coordinated campaigns, on average, have a lower concentration via entropy signal.

[0345] An exemplary average value of the concentration via entropy signal for coordinated campaigns is 1.43.

[0346] An exemplary average range of values of the concentration via entropy signal for coordinated campaigns is 0.46 to 2.19.

[0347] An exemplary standard deviation of the value of the concentration via entropy signal for coordinated campaigns is 0.57.

[0348] An exemplary average value of the concentration via entropy signal for spontaneous campaigns is 2.52.

[0349] An exemplary average range of values of the concentration via entropy signal for spontaneous campaigns is 0.69 - 3.38.

[0350] An exemplary standard deviation of the value of the concentration via entropy signal for spontaneous campaigns is 0.71.

[0351] In embodiments, the concentration via entropy signal (metric) may be useful to analyze “battleground campaigns” where a few clusters fight for control over the social media narrative, e.g., on a dedicated hashtag, where these campaigns may be concentrated in these few communities and simply using a measure focused on the lead community may miss this activity.

[0352] In embodiments, a priority signal name is DayPeakedness.

[0353] The daypeakedness signal description – A coordinated campaign, typically, may exhibit sustained activity by the accounts promoting it. Spontaneous activity, in contrast, is characterized by “bursty” cascades of activity. In embodiments, the daypeakedness signal may detail the percentage of all activity that the busiest day of the campaign may represent.

[0354] The daypeakedness signal (metric) of a campaign is measured as the percentage of campaign actions (Tweets™ or the like) that take place on the most active day of the campaign. It will be appreciated in light of the disclosure that generally spontaneous campaigns appear to be more “bursty” because, for example, spontaneous campaigns exhibit more of a peak (or more of a number of peaks) than coordinated campaigns.

[0355] In embodiments, the range of the values of the daypeakedness signal (metric) is 0% to 100%.

[0356] In embodiments, the value of the daypeakedness signal (metric) is computed by determining the fraction of all activity that occurs on the day with the most campaign-related activity. Examples include a campaign that proceeds over the course of ten days, and the first of those days features the highest level of campaign activity, with one-quarter of all actors

participating. In this example, the value of the daypeakedness signal (metric) is 25%.

[0357] It will be appreciated in light of the disclosure that one-eighth of all activity in coordinated campaigns, on average, happens during peak day, whereas over one-third of all activity for spontaneous campaigns happens during peak day. In embodiments, the daypeakedness signal (metric) can be shown to be consistent across campaigns despite the variety of geographies and dates. By way of this example, coordinated campaigns, on average, may have a lower value of the daypeakedness signal (metric) than spontaneous campaigns. It will be appreciated in light of the disclosure that there may be some overlap between the coordinated and spontaneous ranges due to the large number of sociocultural settings and time periods in the campaign.

[0358] An exemplary average value of the daypeakedness signal for coordinated campaigns is 0.14.

[0359] An exemplary range of values of the daypeakedness signal for coordinated campaigns is 0.08 to 0.22.

[0360] An exemplary standard deviation of the value of the daypeakedness signal for coordinated campaigns is 0.05.

[0361] An exemplary average value of the daypeakedness signal for spontaneous campaigns is 0.41.

[0362] An exemplary average range of values of the daypeakedness signal for spontaneous campaigns is 0 to 0.71.

[0363] An exemplary standard deviation of the value of the daypeakedness signal for spontaneous campaigns is 0.21.

[0364] The daypeakedness signal (metric) may be sensitive to date-boundary/time-zones most notably when the campaign is being analyzed only over the last few days. In embodiments, the sensitivity of the daypeakedness signal (metric) may be improved by allowing it to be less sensitive to time zones.

[0365] It will be appreciated in light of the disclosure that there are other possibly more complex ways to calculate the value of the daypeakedness signal. In embodiments, the peak time may be identified as the median of time stamps of a dynamic phenomenon to be able to observe a logarithmic distribution of volume around the peak. The methods and systems described herein may identify peaks as days when volume exceeds two standard deviations above the median, and may calculate the value of the daypeakedness signal as a fraction of all content that occurred during a 24-hour period. It will be appreciated in light of the disclosure that the median volume may be used instead of mean volume due in part to the observation that volume follows a skewed distribution, so the mean may not be an appropriate statistic to use to characterize it. The measure of peakedness in the methods and systems described herein may be relatively less sophisticated

and, therefore, may be easier to interpret while giving a good initial impression of the utility of the signal from a social media platform for identifying coordinated campaigns.

[0366] In embodiments, the value of the daypeakedness signal (metric) may be affected by the overall time range of a campaign. By way of this example, if a campaign lasts three days, then the value of the daypeakedness signal may not go below 33% but if the campaign lasts 10 days, then the value of the daypeakedness signal cannot go below 10%. In embodiments, campaigns may last as little as one week and may last as long as several months. The value of the daypeakedness signal may be shown to follow the pattern described in the campaign value examples across these time ranges.

[0367] In embodiments, a signal name is Commitment: Average Posts Count in the Campaign.

[0368] The commitment: average posts count in campaign signal description - Campaigns typically feature numerous die-hard supporters who post repeatedly and fewer casual participants who merely chime in. This commitment: average posts count in campaign signal (metric) may capture the degree to which a campaign's body of actors sticks with further posting after their first engagement with the social media platform. In embodiments, the value of the commitment: average posts count in campaign signal (metric) can include the average number of campaign-related posts that participants publish after their first campaign post.

[0369] The range of values of the commitment: average posts count in campaign signal (metric) is bounded by the lowest value being zero which corresponds to a user only posting once about the campaign. In embodiments, the commitment: average posts count in campaign signal (metric) may have a range of values between 0 and 10 posts. It will be appreciated in light of the disclosure that the maximum value of the commitment: average posts count in campaign signal (metric) could be much higher. In one example, participants in a campaign may be very dedicated and may post 100 times about a certain subject during the scope of analysis, and the like.

[0370] To compute the value of the commitment: average posts count in campaign signal (metric), the methods and systems disclosed herein determine the average number of subsequent participation actions, e.g., Tweets™ (or other posting) with campaign hashtag, across all participants in a campaign. In embodiments, participants (i.e., posters) in a campaign can be a smaller subset of participants in a map. In embodiments, the map may capture some of their followers and/or other members of the network terrain when those are highly connected to active participants in the campaign. In order to compute the commitment: average posts count in campaign signal (metric), only participants who actually posted about the campaign are taken into account. For example, when a participant posted through Twitter™, Facebook™, or the like with a campaign-related hashtag twice, their commitment is 1.0. In embodiments, campaign participation can include Tweets™ or the like with campaign-related hashtags (for campaigns

organized around a hashtag), Tweets™ or the like with links to a video or article (for campaigns organized around a video or article), retweets of the above tweets and the like. Examples of out of scope for participation include favorites of tweets with campaign-related hashtags or links or @-replies or @-mentions of Tweets™ (or the like) with campaign-related hashtags or links.

[0371] It will be appreciated in light of the disclosure that participants in spontaneous campaigns post more about their campaigns than participants in coordinated campaigns. It will also be appreciated in light of the disclosure that this pattern may be counterintuitive, as one may expect participants in coordinated campaigns to be extrinsically motivated to hit certain participation targets (e.g., by being paid by number of posts) and thus to post more than participants in spontaneous campaigns, who lack such motivation.

[0372] An exemplary average value of the commitment: average posts count in campaign signal (metric) for coordinated campaigns is 2.52.

[0373] An exemplary average range of values of the commitment: average posts count in campaign signal (metric) for coordinated campaigns is 1.28 to 3.40.

[0374] An exemplary standard deviation of the value of the commitment: average posts count in campaign signal (metric) for coordinated campaigns is 0.84.

[0375] An exemplary average value of the commitment: average posts count in campaign signal (metric) for spontaneous campaigns is 3.53.

[0376] An exemplary average range of values of the commitment: average posts count in campaign signal (metric) for spontaneous campaigns is 1.39 to 6.07.

[0377] An exemplary standard deviation of the value of the commitment: average posts count in campaign signal (metric) for spontaneous campaigns is 1.48.

[0378] In embodiments, the commitment: average posts count in campaign signal (metric) can be analyzed at the community level, at a cluster level, and a participant level. The commitment: average posts count in campaign signal (metric) can be analyzed at the community level to single out communities with participants being particularly committed to a campaign. The commitment: average posts count in campaign signal (metric) can be analyzed at the participant level to represent individuals who have extremely high commitment values, e.g., posting about a campaign one hundred times.

[0379] In embodiments, the commitment: average posts count in campaign signal (metric) is focused on participations after the first post and complemented by a measurement of the proportion of participants in the campaign who have only participated once.

[0380] In embodiments, the commitment: average posts count in campaign signal (metric) may be combined with a commitment: average time range of participation signal (metric) into a commitment: post regularity signal (metric) that may capture the deviation of campaign

participants from natural human attention patterns.

[0381] In embodiments, other statistical properties of the distribution of posts per user may be part of refining the commitment metrics. In embodiments, there may be a natural shape of this distribution for spontaneous campaigns and that natural shape may be skewed. It will be appreciated in light of the disclosure that the commitment: average posts count in campaign signal (metric) may make average post count an inappropriate metric in many long duration situations. Instead, it may be possible to be able to identify coordinated campaigns by a lack of skewness and/or the presence of a second moment at some value above one, which may both be indicative of an unusually large percentage of participants posting multiple times about a campaign, e.g., due to a coordinating body paying these participants per post.

[0382] In embodiments, the commitment: average posts count in campaign signal (metric) may be normalized to take into account average posts per users in order to control for users with a very heavy activity across all campaigns.

[0383] In embodiments, a priority signal name is Commitment: Average Time Range of Participation.

[0384] The commitment: average time range of participation signal description – In the desire to determine whether participants in this campaign are die-hard supporters or just people who chime in, the commitment: average time range of participation signal (metric) may be used to facilitate looking at how long (in days) participants remained engaged in pushing the campaign. In embodiments, the loyalty of participants to the campaign may be measured by time range (in days) for their campaign-related Tweets™ (or other postings) that may be averaged across all participants.

[0385] The range of the values of the commitment: average time range of participation signal (metric) is an unbounded value and therefore can be zero days to the total length of the campaign.

[0386] In embodiments, the commitment: average time range of participation signal (metric) may look at the time frame between first and last participation action that can be averaged across all participants in a campaign. By way of this example, the commitment: average time range of participation signal (metric) may measure whether actors participate in a “one-off” way (one Tweet™ and done) or demonstrate a commitment to the campaign (multiple Tweets™ or other postings over time).

[0387] It will be appreciated in light of the disclosure that participants in coordinated campaigns engage with the campaign over a longer period than participants in spontaneous campaigns. It will also be appreciated in light of the disclosure that participants in coordinated campaigns may be more likely than participants in spontaneous campaigns to receive extrinsic motivation, such as payment, for engaging with the campaign and, as such, the extrinsic motivation may lead to a

longer engagement period than intrinsic motivation.

[0388] An exemplary average value of the commitment: average time range of participation signal (metric) for coordinated campaigns is 7.24.

[0389] An exemplary average range of values of the commitment: average time range of participation signal (metric) signal for coordinated campaigns is 0.08 to 22.33 days.

[0390] An exemplary standard deviation of the value of the commitment: average time range of participation signal (metric) for coordinated campaigns is 9.04 days.

[0391] An exemplary average value of the commitment: average time range of participation signal (metric) for spontaneous campaigns is 1.53 days.

[0392] An exemplary average range of values of the commitment: average time range of participation signal (metric) for spontaneous campaigns is 0 to 3.36 days.

[0393] An exemplary standard deviation of the value of the commitment: average time range of participation signal (metric) for spontaneous campaigns is 1.23 days.

[0394] It will be appreciated in light of the disclosure that the commitment: average time range of participation signal (metric) may be affected by the overall time range of a campaign, e.g., if a campaign lasts three days, then this metric cannot go above a value of three. In embodiments, the commitment: average time range of participation signal (metric) may be combined into a commitment: post regularity signal that may capture the deviation of campaign participants from natural human attention patterns.

[0395] In embodiments, a signal name is Semantic Diversity for all Messages.

[0396] The semantic diversity for all messages signal (metric) description – The semantic diversity for all messages signal (metric) looks to detail how generally on-message is the campaign. The semantic diversity for all messages signal (metric) also looks to determine whether the interaction or activity appears like a diverse conversation covering a range of topics and expressions or may be a fairly uniform campaign with low semantic diversity. It will be appreciated in light of the disclosure that people tend to Tweet™ (or otherwise post) on a variety of topics related to their daily lives, work, and interests. A group trying to promote a coordinated campaign, however, may be interested only in the narrow range of topics relevant to that campaign. In embodiments, bots or propaganda accounts may also be interested in any Tweet™ (or applicable posting) relevant to any campaign they are trying to push, and therefore could be Tweeting™ (or otherwise posting) on an extremely wide range of topics. In embodiments, the semantic diversity for all messages signal (metric) may be measuring the extent to which participants in the campaign are Tweeting™ (or otherwise posting) on an intermediate range of topics, which suggests that their activities are spontaneous and human rather than automated or coordinated to propagate a specific message.

[0397] In embodiments, the range of values of the semantic diversity for all messages signal (metric) is zero to 100%.

[0398] In embodiments, raw values of the semantic diversity for all messages signal (metric) fall into three categories: (i) When the value of the semantic diversity for all messages signal (metric) is <1 (less than one), then it may represent users who exclusively post about the same topic, which may be a characteristic of fabricated campaigns. (ii) When the value of the semantic diversity for all messages signal (metric) is between one and 100, then it may represent users who post on a variety of topics and being characteristic of normal human activity. (iii) When the value of the semantic diversity for all messages signal (metric) is above 100, then it may represent users who post on an extremely diverse set of topics, characteristic of spambots or users who bridge different cultural and/or linguistic communities (e.g., users who post in different languages, etc.). In embodiments, the semantic diversity for all messages signal (metric) may be set to be bounded at 1000 because it may be necessary to fix a maximum value for the “distance” between any pair of topics, for which no document includes terms from both topics. It will be appreciated in light of the disclosure that mathematically the distance should be infinity but, typically, it can be to set the value to 1000. The percentage of users with the semantic diversity for all messages signal (metric) may be greater than or equal to 1.0 and less than 100 and thus varies between zero and 100%.

[0399] How the semantic diversity for all messages signal (metric) is computed – The value of the semantic diversity for all messages signal (metric) of a particular actor’s (or cluster’s, or campaign’s) messaging may be based on the assignment of messages to topics. In embodiments, the computation of the semantic diversity for all messages signal (metric) may use a Latent Dirichlet Allocation algorithm. By way of this example, once messages have been assigned to topics, the semantic diversity for all messages signal (metric) is determined for the message set. In embodiments, the measure of the value of the semantic diversity for all messages signal (metric) is determined as the probability that two documents chosen from the corpus at random with replacement will be on the same topic.

[0400] In the current exemplary case, the corpus is the message set, and the documents may be user Tweet™ (or other posting) histories, aggregated by user. The Latent Dirichlet Allocation (LDA) algorithm may be run for fifteen iterations with a number of topics no less than 20% of the number of documents and no more than 30%. An average value of the semantic diversity for all messages signal (metric) over twenty distinct runs of the LDA algorithm is used on the same corpus to smooth out variations due to the initial conditions for a particular run. In embodiments, a topic distance score of 1000 may be assigned to the semantic diversity for all messages signal (metric) for topics that do not co-occur in documents.

[0401] Because the focus of the many embodiments is differentiating coordinated and/or

automated campaigns from spontaneous and human-driven campaigns, the semantic diversity for all messages signal (metric) as the percentage of all users in a campaign is computed with raw diversity score falling into the range of normal human activity, i.e., the metric being greater or equal to 1.0 but less than 100. In embodiments, the semantic diversity for all messages signal (metric) may refer to all campaign-related messages.

[0402] The values below show the percentage of users with the semantic diversity for all messages signal (metric) greater than or equal to 1.0 and less than 100.0.

[0403] An exemplary average value of the semantic diversity for all messages signal (metric) for coordinated campaigns is 55%.

[0404] An exemplary average range of values of the semantic diversity for all messages signal (metric) for coordinated campaigns is 17% to 90%.

[0405] An exemplary standard deviation of the value of the semantic diversity for all messages signal (metric) for coordinated campaigns is 36.59%.

[0406] An exemplary average value of the semantic diversity for all messages signal (metric) for spontaneous campaigns is 71.3%.

[0407] An exemplary average range of values of the semantic diversity for all messages signal (metric) for spontaneous campaigns is 50% to 98%.

[0408] An exemplary standard deviation of the value of the semantic diversity for all messages signal (metric) for spontaneous campaigns is 21.2%.

[0409] In embodiments, the semantic diversity for all messages signal (metric) may be very sensitive to confounds. By way of this example, news organizations may tend to have low semantic diversity because news organizations may post the same story headlines over and over even though such news organizations are not coordinated actors. Moreover, Tweets™ (or other postings) in one language tend to be more coordinated than Tweets™ (or other postings) in multiple languages, because the Latent Dirichlet Allocation (LDA) algorithm may not translate terms across languages.

[0410] At the same time, the semantic diversity for all messages signal (metric) may point to the differentiation between natural language use and the use of language to push a particular message. It will be appreciated in light of the disclosure that coordination around a message may require that that message may be as clear and simple as possible, whereas natural language can be complex, metaphorical, and even slightly confusing. To that end, coordinated campaigns may, therefore, not wish to increase the semantic diversity of their messages even if the technical or organizational opportunity was available.

[0411] In embodiments, the semantic diversity for all messages signal (metric) includes separating language diversity from semantic diversity either by grouping Tweets™ (or other postings) by

post language prior to analysis or using automated machine translation to pre-convert all Tweets™ (or other postings) to the same language. The semantic diversity for all messages signal (metric) also includes leveraging existing natural language processing approaches to identify certain kinds of low-semantic diversity language that may not be of interest, e.g., news headlines and press releases.

[0412] In embodiments, the semantic diversity for all messages signal (metric) may measure the temporal alignment of campaign-related Tweets™ (or other postings) for all participants. It will be appreciated in light of the disclosure that users generally do not time their Tweets™ (or other postings) to coincide with the Tweets™ (or postings) of others. When the Tweet™ (or other posting) histories of campaign participants follow the same pattern of ebb and flow, especially across time zone boundaries, this may be evidence that an actor is coordinating the activities of participants to create a concentrated temporal burst of engagement. The semantic diversity for all messages signal (metric) may include temporal coordination of Tweets™ (or other postings) between campaign participants measured by alignment of Tweet™ (or other posting) histories across all participants in the campaign.

[0413] In embodiments, the range of the values of the semantic diversity for all messages signal (metric) is between 0% and 100% and represents the percent alignment of two users' temporal normalized sequences of participation in the campaign. Toward that end, 0% alignment may mean that the users' sequences do not match at all, while 100% alignment may indicate a perfect match.

[0414] In embodiments, the semantic diversity for all messages signal (metric) may be computed with a dynamic time warp algorithm for comparing two temporal sequences of activity. In general, the dynamic time warp algorithm between two sequences S1 and S2 is the number of warping transformations that are required to change S1 into S2. The methods and systems described herein may, for example, use the dynamic time warp algorithm to identify bots and trolls in a different social media setting. The number of warping transformations may be normalized by the length of both sequences S1 and S2 and multiplied by 100 to get a percent value. Finally, the normalized number may be subtracted from 100 in order to calculate the percent alignment of S1 and S2.

[0415] In embodiments, a priority signal name is temporal coordination per cluster.

[0416] The temporal coordination per cluster signal (metric) description — The temporal coordination per cluster signal (metric) may look at the communities who participate in this campaign to identify different communities exhibiting very similar patterns of engagement that may be considered as being odd. In embodiments, the pattern of the temporal coordination per cluster signal (metric) may be even odder when postings exist in different time zones. The temporal coordination per cluster signal (metric) is measuring the temporal alignment of campaign-related Tweets™ (or other postings) aggregated at the cluster level. With that in mind,

communities generally do not time their Tweets™ (or other postings) to coincide with the Tweets™ (or other postings) of other communities. When the Tweet™ (or other posting) histories of participating clusters follow the same pattern of ebb and flow, especially across time zone boundaries, this may be evidence that an actor is coordinating the activities of participants to create a concentrated temporal burst of engagement.

[0417] The range of values for the temporal coordination per cluster signal (metric) is zero percent to 100%. The value of the temporal coordination per cluster signal (metric) represents the percent alignment of two users' temporal normalized sequences of participation in the campaign. Toward that end, 0% alignment may mean that the users' sequences do not match at all, while 100% alignment indicates a perfect match.

[0418] The temporal coordination per cluster signal (metric) description – The temporal coordination per cluster signal (metric) is a per-user take on examining temporal coordination, which might be helpful when other metrics are noisy. Temporal coordination per user is technically the temporal coordination between pairs of users. In embodiments, the temporal coordination per cluster signal (metric) may measure the temporal alignment of campaign-related Tweets™ (or other postings) between individual campaign participants. As noted before, users generally do not time their Tweets™ (or other postings) to coincide with the tweets of others. When the Tweet™ (or other posting) histories of campaign participants follow the same pattern of ebb and flow, especially across time zone boundaries, this may be evidence that an actor is coordinating the activities of participants to create a concentrated temporal burst of engagement.

[0419] The temporal coordination per cluster signal (metric), especially its heatmap visualization, may provide a good high-level description of the rate of unusual coordination across the users participating in a campaign. The temporal coordination per cluster signal (metric), however, may suffer from the same overestimation of actual temporal coordination so the algorithm may be adjustable for including in the calculation the average temporal coordination across users.

[0420] In embodiments, a signal name is client diversity per cluster.

[0421] The client diversity per cluster signal (metric) description – The client diversity per cluster signal (metric) may determine how accounts in a given cluster use Twitter™, Facebook™, or other social media platforms. The client diversity per cluster signal (metric) may also determine how Twitter™ users (or other posters or various relevant platforms) go through a mobile device, a computer, or directly access APIs of Twitter™ to Tweet™ (or other social media postings). In one example, some clients may be used to coordinated Tweets™ (or other social media postings) and the client diversity per cluster signal (metric) may be used to determine how coordinate are the Tweets™ (or other social media postings), and are such coordinating Tweets™ (or other social media postings) those that are used heavily in some of the communities who participate in this

campaign. It will be appreciated in light of the disclosure that client diversity per cluster signal (metric) is the same as the client diversity at campaign scale signal (metric) but analyzed at the cluster level.

[0422] There is no specific range of values applicable to the client diversity per cluster signal (metric) because it is a qualitative signal (metric)

[0423] The value of the client diversity per cluster signal (metric) is computed by using the "source" field of the Tweet™ (or other posting) to identify the client used to make the Tweet™ (or other posting), as in the Client diversity at campaign scale signal (metric). Then the Tweets™ (or other postings) are aggregated into clusters of the author of the Tweet™ (or other posting) in the campaign map.

[0424] In embodiments, a signal name is Time Delta between Communities.

[0425] The time delta between communities signal (metric) description – the time delta between communities signal (metric) may identify a community that is engaging with the campaign significantly ahead of others. In one example, this is due to kick-starting that campaign or being significantly behind maybe because there is a need to coordinate talking points before engaging. It will be appreciated in light of the disclosure that the time delta between communities signal (metric) was inspired by qualitative analysis initially done in the Syrian Civil War context such that communities pretending to portray civilians while being led by military intelligence engaged with popular topics with a lag of several hours to days. Toward that end, the time delta between communities signal (metric) may examine when clusters are most active in the campaign. By way of this example, the time delta between communities signal (metric) may measure the distance between a given cluster's peak and the more general peak of the overall campaign.

[0426] In embodiments, the range of values of the time delta between communities signal (metric) represents a number of days. Negative values may indicate that a community's peak of temporal activity happens before the average peak date for all other communities. Positive values may indicate the peak happens after the average peak date for all other communities. A score of zero may indicate a community peaking in sync with the rest of the communities.

[0427] How the time delta between communities signal (metric) is computed – This metric measures the number of days between the peak date of campaign participation in a given cluster and peak date of campaign participation averaged across all other clusters. In one example with three clusters, where activity in cluster A peaks on 25 January 2017, activity in cluster B peaks on 26 January 2017, and activity in cluster C peaks on 27 January 2017, the value of the time delta between communities signal (metric) for A equals -1.5, the value of the time delta between communities signal (metric) for B equals zero, and the value of the time delta between communities signal (metric) for C equals 1.5.

[0428] In embodiments, the time delta between communities signal (metric) may be helpful to analyze disputed hashtags, with both spontaneous and coordinated clusters engaging in the same campaign. In embodiments, the time delta between communities signal (metric) may point to the natural logistical cost of coordinating a message of a campaign in response to a sudden event, such as a late-breaking news story. It will be appreciated in light of the disclosure that even the most sophisticated coordinated campaigns cannot anticipate such events and at the same time, they cannot respond to these events spontaneously as it may distract from their message and may hurt the overall aim of the campaign. It will also be appreciated in light of the disclosure that all coordinated campaigns will need at least a little time to respond to late-breaking events, and their responses will measurably lag behind spontaneous human reactions to the same. In embodiments, the time delta between communities signal (metric) may include automatic identification of sudden events as they happen, e.g., by matching campaign-related terms against Google™ News, other news sources, and the like. A subsequent step may be to automatically track responses to the same events from campaign compared to non-campaign-related clusters.

[0429] In embodiments, a signal name is Commitment by User.

[0430] The commitment by user signal (metric) description – Loyalty of participants to the campaign may be measured by the number of times the participants Tweet™ (or otherwise post) about the campaign and time range (in days) for their campaign-related Tweets™ (or other postings). The commitment by user signal (metric) may be measured by the user. In embodiments, the commitment by user signal (metric) looks at whether individual users are particularly committed to a campaign. In embodiments, the commitment by user signal (metric) may facilitate looking at users and their own commitments by determining whether there are, for example, people who Tweet™ (or otherwise post) exactly 100 times, or some predictable predetermined amount. The value of the commitment by user signal (metric) may facilitate identifying and singling out accounts that might be incentivized to participate x number of times or for x days straight.

[0431] The range of values of the commitment by user signal (metric) are unbounded values starting at zero, i.e., no subsequent actions, zero days pass between first and last action. In embodiments, values for the commitment by user signal (metric) by subsequent actions are between zero and ten actions, those for commitment by time frame are between zero and thirty days.

[0432] In embodiments, there may be users whose commitment by user signal (metric) is extremely high and such behavior may also contribute to higher values associated with the Commitment: average time range of participation signal (metric) noted above.

[0433] In embodiments, a signal name is Commitment by Cluster.

[0434] The commitment by cluster signal (metric) description -- The commitment by cluster signal (metric) may be used to determine whether a specific cluster is particularly committed to a campaign. In embodiments, the commitment by cluster signal (metric) may facilitate looking at clusters and their own commitments. By way of this example, the commitment by cluster signal (metric) may facilitate the determination of whether there are clusters that Tweet™ (or otherwise post) exactly 100 times. In embodiments, the commitment by cluster signal (metric) may be used to single out clusters that might be incentivized to participate a certain number of times or for a certain length of time. In one example, the commitment by cluster signal (metric) may be used to determine whether a group of accounts showed up, Tweeted™ (or otherwise posted) 100 times over five days, and then left.

[0435] In embodiments, the commitment by cluster signal (metric) may look at the loyalty of participants to the campaign that may be measured by the number of times the participants Tweet™ (or otherwise post) about the campaign and time range (in days) for their campaign-related Tweets™ (or other postings). In embodiments, the commitment by cluster signal (metric) may measure the degree to which a body of actors in the campaign stick with it after their first engagement with the campaign. It will be appreciated in light of the disclosure that the value of the commitment by cluster signal (metric) for most human activity is a skewed distribution in measurable contrast to coordinated activity that may include those who participate once with a few die-hard supporters that participate a lot. Deviations from the skewed distribution detailing human activity may, therefore, may reveal coordination. By way of this example, if an actor participates in a campaign exactly 100 times, this may suggest that they were incentivized by a coordinating body to meet that threshold.

[0436] The range of the values of the commitment by cluster signal (metric) are unbounded values starting at zero, i.e., no subsequent actions, zero days pass between first and last action. In embodiments, the value of the commitment by cluster signal (metric) by subsequent actions is between zero and ten actions. In further embodiments, the value of the commitment by cluster signal (metric) by time frame is between zero and thirty days.

[0437] How the value of the commitment by cluster signal (metric) is computed -- There are two commitment metrics: (i) counting the number of subsequent participation "actions" (i.e., Tweets™ or other postings with a campaign hashtag), and (ii) the time frame (in days, can be fractional) between first and last participation action. Both metrics may be averaged across all participants in a campaign. Both metrics may measure whether actors participate in a "one-off" way (i.e., one Tweet™ or other posting and done) or may demonstrate a commitment to the campaign (e.g., multiple Tweets™ or other postings over time).

[0438] In embodiments, a signal name is Account Creation Date Diversity for Cluster.

[0439] The account creation date diversity for cluster signal (metric) description – this signal (metric) may facilitate observing how close in time all accounts participating in a campaign were created. If 90% of participating accounts within a given cluster were created within a span of five days, for example, then such activity may indicate a heavy coordination within that cluster. The account creation date diversity for cluster signal (metric) may be particularly helpful to spot bots, troll farms, and the like on networks using fake accounts generated in bulk.

[0440] The range of values of the account creation date diversity for cluster signal (metric) is zero to 4,015 days. It will be appreciated in light of the disclosure that the maximum range may range from zero to the total day since the founding of Twitter™ or the other applicable social media platforms. The values of the account creation date diversity for cluster signal (metric) in datasets evaluated have included a range of zero to 1,200 days.

[0441] How the account creation date diversity for cluster signal (metric) is computed – Account creation date diversity for a particular cluster and campaign combination is the standard deviation (in days) of Twitter™ (or other applicable social media platform) account creation dates for all accounts in that cluster who engaged with the campaign in question. As a baseline, embodiments may compare account creation date diversity for a particular cluster to account creation date diversity for the entire campaign.

[0442] In embodiments, a signal name is Homophily.

[0443] The homophily signal (metric) description – This signal (metric) may facilitate looking for communities that pay a “disproportionate” amount of attention to one another, for instance across ideologies, language, culture, or the like. In embodiments, the homophily signal (metric) can identify disproportionate attention relationships between clusters measured by a number of following relationships between clusters. When looking at communities (clusters), it will be appreciated in light of the disclosure that it is just as important to understand who the community pays attention to as who is in the community. With this in mind, the homophily signal (metric) may measure deviations from expected patterns of attention in social media. By way of this example, it will be appreciated in light of the disclosure that most people may pay most of their attention to like-minded friends and the vast majority of people may pay most of their attention to friends in the same cultural and linguistic environment or in their affinity. In further examples, the homophily signal (metric) may facilitate the identification of patterns of intense inter-attention across ideologies, culture, and language that may imply evidence for coordination.

[0444] The range of values of the homophily signal (metric) can be shown to be zero to ten.

[0445] How the homophily signal (metric) is computed – The homophily signal (metric) as a telltale of cluster attention is a ratio of the actual number of edges connecting members of the clusters compared to what would be expected under conditions where each cluster paid attention

to every other cluster strictly in proportion to the cluster's size. Typically, the baseline for such a signal (metric) is random connection patterns. In embodiments, the homophily signal (metric) includes relatively more aggressive baselines because no actual human relationships follow a random pattern.

[0446] In embodiments, a signal name is Language Mismatch.

[0447] The language mismatch signal (metric) description – The default language for a new Twitter™ (or other social media) account appears to be English. Users may, however, choose to change their profile language if they want. It will be appreciated in light of the disclosure that users posting frequently in a language that differs from their default Twitter™ (or other social media) profile language may be part of a foreign-language propaganda operation on behalf of some coordinated entity.

[0448] The language mismatch signal (metric) may measure the percentage of a campaign's Tweets™ (or other postings) - at both the cluster and campaign level - that is in a language that differs from the users' default Twitter™ (or other social media) profile language.

[0449] The range of values of the language mismatch signal (metric) is zero to one hundred percent, where one hundred percent would have indicated that all campaign participation actions in this cluster/campaign are Tweeted™ (or otherwise posted) in a language different from their accounts' default profile language.

[0450] How the language mismatch signal (metric) is computed – For each Tweet™ (or other posting) with the campaign-related hashtag, the language mismatch signal (metric) may identify the language of the Tweet™ (or other posting) and the language profile setting in the Twitter™ API or the API of another social media platform. In embodiments, the language mismatch signal (metric) may also aggregate the Tweets™ (or other postings) by the cluster of the author of the Tweet™ (or other posting) in a campaign map. By way of this example, the % of Tweets™ (or other postings) for each cluster whose tweet language did not match the poster language of the Tweet™ (or other posting) may be reported.

[0451] Detailed embodiments of the present disclosure are disclosed herein; however, it is to be understood that the various disclosed embodiments are merely exemplary of the disclosure, which may be embodied in various forms. Therefore, specific structural and functional details disclosed herein are not to be interpreted as limiting, but merely as a basis for the claims and as a representative basis for teaching one skilled in the art to variously employ the present disclosure in virtually any appropriately detailed structure.

[0452] The terms “a” or “an,” as used herein, are defined as one or more than one. The term “another,” as used herein, is defined as at least a second or more. The terms “including” and/or “having,” as used herein, are defined as comprising (i.e., open transition).

[0453] While only a few embodiments of the present disclosure have been shown and described, it will be obvious to those skilled in the art that many changes and modifications may be made thereunto without departing from the spirit and scope of the present disclosure as described in the following claims.

[0454] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software, program codes, and/or instructions on a processor. The present disclosure may be implemented as a method on the machine, as a system or apparatus as part of or in relation to the machine, or as a computer program product embodied in a computer readable medium executing on one or more of the machines. In embodiments, the processor may be part of a server, cloud server, client, network infrastructure, mobile computing platform, stationary computing platform, or other computing platform. A processor may be any kind of computational or processing device capable of executing program instructions, codes, binary instructions, and the like. The processor may be or may include a signal processor, digital processor, embedded processor, microprocessor, or any variant such as a co-processor (math co-processor, graphic co-processor, communication co-processor and the like) and the like that may directly or indirectly facilitate execution of program code or program instructions stored thereon. In addition, the processor may enable execution of multiple programs, threads, and codes. The threads may be executed simultaneously to enhance the performance of the processor and to facilitate simultaneous operations of the application. By way of implementation, methods, program codes, program instructions and the like described herein may be implemented in one or more thread. The thread may spawn other threads that may have assigned priorities associated with them; the processor may execute these threads based on priority or any other order based on instructions provided in the program code. The processor, or any machine utilizing one, may include non-transitory memory that stores methods, codes, instructions, and programs as described herein and elsewhere. The processor may access a non-transitory storage medium, through an interface that may store methods, codes, and instructions as described herein and elsewhere. The storage medium associated with the processor for storing methods, programs, codes, program instructions or other type of instructions capable of being executed by the computing or processing device may include but may not be limited to one or more of a CD-ROM, DVD, memory, hard disk, flash drive, RAM, ROM, cache, and the like.

[0455] A processor may include one or more cores that may enhance speed and performance of a multiprocessor. In embodiments, the process may be a dual core processor, quad core processors, other chip-level multiprocessor and the like that combine two or more independent cores (called

a die).

[0456] The methods and systems described herein may be deployed in part or in whole through a machine that executes computer software on a server, client, firewall, gateway, hub, router, or other such computer and/or networking hardware. The software program may be associated with a server that may include a file server, print server, domain server, Internet server, intranet server, cloud server, and other variants such as secondary server, host server, distributed server, and the like. The server may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other servers, clients, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the server. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the server.

[0457] The server may provide an interface to other devices including, without limitation, clients, other servers, printers, database servers, print servers, file servers, communication servers, distributed servers, social networks, and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the server through an interface may include at least one storage medium capable of storing methods, programs, code and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[0458] The software program may be associated with a client that may include a file client, print client, domain client, Internet client, intranet client and other variants such as secondary client, host client, distributed client, and the like. The client may include one or more of memories, processors, computer readable media, storage media, ports (physical and virtual), communication devices, and interfaces capable of accessing other clients, servers, machines, and devices through a wired or a wireless medium, and the like. The methods, programs, or codes as described herein and elsewhere may be executed by the client. In addition, other devices required for execution of methods as described in this application may be considered as a part of the infrastructure associated with the client.

[0459] The client may provide an interface to other devices including, without limitation, servers, other clients, printers, database servers, print servers, file servers, communication servers, distributed servers, and the like. Additionally, this coupling and/or connection may facilitate remote execution of program across the network. The networking of some or all of these devices

may facilitate parallel processing of a program or method at one or more location without deviating from the scope of the disclosure. In addition, any of the devices attached to the client through an interface may include at least one storage medium capable of storing methods, programs, applications, code and/or instructions. A central repository may provide program instructions to be executed on different devices. In this implementation, the remote repository may act as a storage medium for program code, instructions, and programs.

[0460] The methods and systems described herein may be deployed in part or in whole through network infrastructures. The network infrastructure may include elements such as computing devices, servers, routers, hubs, firewalls, clients, personal computers, communication devices, routing devices and other active and passive devices, modules and/or components as known in the art. The computing and/or non-computing device(s) associated with the network infrastructure may include, apart from other components, a storage medium such as flash memory, buffer, stack, RAM, ROM, and the like. The processes, methods, program codes, instructions described herein and elsewhere may be executed by one or more of the network infrastructural elements. The methods and systems described herein may be adapted for use with any kind of private, community, or hybrid cloud computing network or cloud computing environment, including those which involve features of software as a service (SaaS), platform as a service (PaaS), and/or infrastructure as a service (IaaS).

[0461] The methods, program codes, and instructions described herein and elsewhere may be implemented on a cellular network having multiple cells. The cellular network may either be frequency division multiple access (FDMA) network or code division multiple access (CDMA) network. The cellular network may include mobile devices, cell sites, base stations, repeaters, antennas, towers, and the like. The cell network may be a GSM, GPRS, 3G, EVDO, mesh, or other networks types.

[0462] The methods, program codes, and instructions described herein and elsewhere may be implemented on or through mobile devices. The mobile devices may include navigation devices, cell phones, mobile phones, mobile personal digital assistants, laptops, palmtops, netbooks, pagers, electronic books readers, music players and the like. These devices may include, apart from other components, a storage medium such as a flash memory, buffer, RAM, ROM and one or more computing devices. The computing devices associated with mobile devices may be enabled to execute program codes, methods, and instructions stored thereon. Alternatively, the mobile devices may be configured to execute instructions in collaboration with other devices. The mobile devices may communicate with base stations interfaced with servers and configured to execute program codes. The mobile devices may communicate on a peer-to-peer network, mesh network, or other communications network. The program code may be stored on the storage

DEMANDE OU BREVET VOLUMINEUX

LA PRÉSENTE PARTIE DE CETTE DEMANDE OU CE BREVET COMPREND PLUS D'UN TOME.

CECI EST LE TOME 1 DE 2
CONTENANT LES PAGES 1 À 95

NOTE : Pour les tomes additionels, veuillez contacter le Bureau canadien des brevets

JUMBO APPLICATIONS/PATENTS

THIS SECTION OF THE APPLICATION/PATENT CONTAINS MORE THAN ONE VOLUME

THIS IS VOLUME 1 OF 2
CONTAINING PAGES 1 TO 95

NOTE: For additional volumes, please contact the Canadian Patent Office

NOM DU FICHIER / FILE NAME :

NOTE POUR LE TOME / VOLUME NOTE:

CLAIMS

1. A method for determining a coordinated activity in social media movements on a social media channel, the method comprising:
 - identifying a plurality of markers of the coordinated activity through analysis of campaign data from the social media movements;
 - storing, in a storage associated with the social media channel, a data structure of the plurality of markers for a social media campaign on the social media channel, wherein the plurality of markers includes a network dimension representing how user accounts of the social media channel are connected, a temporal dimension representing patterns of messages associated with the user accounts over time, and a semantic dimension representing a diversity of topics and meanings of the social media movements;
 - analyzing the data structure to identify the coordinated activity of the social media movements in the social media campaign including:
 - computing semantic diversity over time to identify co-occurring topics in the social media campaign,
 - determining users participating in the social media movements,
 - generating clusters of users in the social media campaign based on relationships between the users participating in the social media movements, and
 - determining propagation patterns of the coordinated activity across the clusters of users of the social media campaign:
 - storing, in the storage, the analyzed data structure;
 - receiving a request from an external system about the coordinated activity of the social media movements;
 - retrieving at least a portion of the analyzed data structure of the plurality of markers for the social media campaign; and
 - transmitting the at least portion of the analyzed data structure to a user interface of the external system that displays at least a portion of the plurality of markers indicative of one of a fabricated campaign, a spambots activity, or normal human activity, wherein
 - a predetermined small value of a semantic diversity score is configured to be indicative of the fabricated campaign,

a predetermined large value of the semantic diversity score is configured to be indicative of the spambots activity, and

a value in-between the predetermined small and large values is indicative of the normal human activity.

2. The method of claim 1, wherein the identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is concentrated in the clusters of users.
3. The method of claim 1, wherein the coordinated activity of the social media campaign is determined from user actions within the social media movements in the social media campaign.
4. The method of claim 1, wherein the identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is distributed among the clusters of users.
5. The method of claim 1, wherein the plurality of markers includes a day peakedness marker that indicates a percentage of the coordinated activity of the social media campaign on a day identified as most active of the social media campaign.
6. The method of claim 1, wherein the plurality of markers includes a commitment indicator that is computed by averaging a number of subsequent participation actions for each of a plurality of participants in the coordinated activity of the social media campaign.
7. The method of claim 6, wherein the plurality of markers includes a post regularity commitment indicator that represents a deviation of commitment to participation by a user from natural human attention patterns.
8. The method of claim 1, wherein the identifying the plurality of markers includes determining the semantic diversity score for the coordinated activity of the social media

campaign by assigning messages in the campaign to topics and calculating a diversity of the topics on a topic distance scale that facilitates determining the semantic diversity score.

9. The method of claim 1, wherein the identifying the plurality of markers includes computing temporal alignment of campaign-related actions for the users in the social media campaign by comparing temporal sequences of the campaign-related actions.

10. A computer system for determining a coordinated activity in social media movements on social media channel, the system comprising:

- a user interface that manages a social media campaign on one or more social media channels and that communicates via a network;

- a computing device that:

- identifies a plurality of markers of the coordinated activity through analysis of campaign data from the social media movements,

- stores one or more data structures containing the plurality of markers for the social media campaign on the one or more social media channels, wherein the plurality of markers includes a network dimension representing how user accounts of the one or more social media channels are connected, a temporal dimension representing patterns of messages associated with the user accounts over time, and a semantic dimension representing a diversity of topics and meanings of the social media movements,

- analyzes the one or more data structures to identify the coordinated activity of the social media movements in the social media campaign including:

- computing semantic diversity over time to identify co-occurring topics in the social media campaign,

- determining users participating in the social media movements;

- generating clusters of users in the social media campaign based on relationships between the users participating in the social media movements, and

- determining propagation patterns of the coordinated activity across the clusters of users of the social media campaign;

- a storage system that stores the analyzed one or more of data structures containing the plurality of markers for the social media campaign on the one or more of the social media channels;

a processing system that executes computer-readable instructions that cause the processing system to:

receive a request from an external system about the coordinated activity of from the social media movements;

retrieve at least a portion of the analyzed one or more data structures containing the plurality of markers for the social media campaign on the one or more of the social media channels; and

transmit the at least portion of the analyzed one or more data structures to a user interface of the external system that displays at least a portion of the plurality of markers indicative of one of a fabricated campaign, a spambots activity, and normal human activity, wherein:

a predetermined small value of a semantic diversity score is configured to be indicative of the fabricated campaign,

a predetermined large value of the semantic diversity score is configured to be indicative of the spambots activity, and

a value in-between the predetermined small and large values is indicative of the normal human activity.

11. The system of claim 10, wherein the identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is concentrated in the clusters of users.
12. The system of claim 10, wherein the coordinated activity of the social media campaign is determined from user actions within the social media movements in the social media campaign, wherein the coordinated activity includes a relatively large number of accounts on one or more of the social media channels controlled by a relatively small number of coordinated entities resulting in a relative lack of diversity of similar accounts on the one or more social medial channels controlled by uncoordinated users.
13. The system of claim 10, wherein the identifying the plurality of markers includes evaluating a degree to which the coordinated activity of the social media campaign is distributed among the clusters of users.

14. The system of claim 10, wherein the plurality of markers includes a day peakedness marker that indicates a percentage of the coordinated activity of the social media campaign on a day identified as most active of the social media campaign.
15. The system of claim 10, wherein the plurality of indicators includes a commitment indicator that is computed by averaging a number of subsequent participation actions for each of a plurality of participants in the coordinated activity of the social media campaign.
16. The system of claim 15, wherein the plurality of markers includes a post regularity commitment indicator that represents a deviation of commitment to participation by a user from natural human attention patterns.
17. The system of claim 10, wherein the identifying the plurality of markers through analysis of campaign signals includes determining the semantic diversity score for the coordinated activity of the social media campaign, wherein determining a semantic diversity score includes assigning messages in the campaign to topics and calculating a diversity of the topics on a topic distance scale that facilitates determining the semantic diversity score.
18. The system of claim 10, wherein the identifying the plurality of markers includes computing temporal alignment of campaign-related actions for the users in the social media campaign by comparing temporal sequences of the campaign-related actions.





Fig. 1

Fig. 2



200

300 ↗




Fig. 3




Fig. 4

400



500 ↗




Fig. 5

600 ↗




Fig. 6




Fig. 7

800






Fig. 8





1100
NEWS ATTENTIVE

NON-NEWS ATTENTIVE




Fig. 11

1200




Fig. 12


1300
↗

Fig. 13

1400 ↗




Fig. 14




Fig. 15

1600 




Fig. 16

1700 →





Fig. 17






Fig. 19b




Fig. 19b




Fig. 20




Fig. 21




Fig. 22a



Fig. 22b



Fig. 22c




Fig. 24






Fig. 26




Fig. 27





Fig. 28






Fig. 30




Fig. 31





Fig. 32







Fig. 34







Fig. 37





Overview

Take Action



Click the selected row and make all map nodes visible


3 Results

Clear

Group / segment name	Members	Power	Tweet activity
UK Design	143	5	5
London Design	127	6	6
UK Marketing Pro	124	5	5
UK Agencies	103	4	4
UK Printing	98	4	4
UK Illustration	95	3	3
UK Art	88	3	3
UK Design	73	3	3

UK Day For Design	65	2	2
UK Entrepreneur	60	1	1
UK Illustration	45	1	1
Manchester Design	33	1	1
UK Brand + Design	34	1	1
UK Architecture	18	1	1

International	543	6	5
Soundmark Design	157	4	4
BR Design	155	4	4
High Design	88	3	3



View Fullscreen Map



Influencer Name Followers
2m Tweets: 12,000

Influencer Name Followers
2m Tweets: 12,000

Influencer Name Followers
2m Tweets: 12,000

View all Recent Activity



Design

UK Design

Grid molecule eugae congue et pretium. Vestibulum ante ipsum primis in faucibus orci luctus et ultrices posuere.

Read more

25.2K

Members

Power

Tweet activity

Items remain visible when page scrolls (fixed position)