LEVERAGING SOCIAL NETWORK ANALYSIS AND SUPERVISED MACHINE LEARNING TO STUDY COORDINATION IN ONLINE INFORMATION CAMPAIGNS

Tuja Khaund

Information Science
University of Arkansas at Little Rock, USA

Agenda

- Introduction
- Literature Review
- Motivation
- Methodology
- Analysis &Results
- Conclusion

Introduction

Introduction

- Information diffusion
  - Tactics and Organizational structures
  - Actors
    - Bots
    - Cyborgs
  - Strategies
    - Coordinated campaigns
    - Manipulation of search ranking
    - Cross-media dissemination



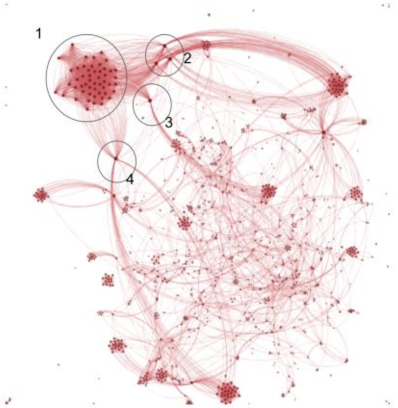

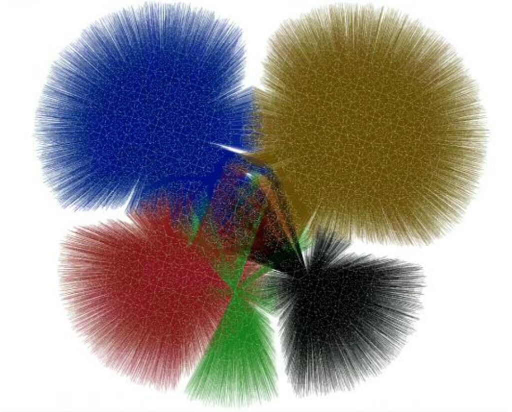

Motivation

- Social bots exploit, mislead, and manipulate social media discourse.
- Information actors coordinate with other actors to spread information faster within the network.

Previous Work

1. Analyzing Social Bots and Their Coordination During Natural Disasters.
2. Analyzing Social and Communication Network Structures of Social Bots and Humans.

Highlights

- Research aims to study the behavior of bots in a social space and analyze their communication and network coordination strategies.
- We analyzed the role of social bots during two different events, natural disaster events and international sporting event.
- Successfully identified distinct network characteristics between bots and humans across different events.

Campaign coordination on Twitter

What is Coordination?

- *“the additional information processing performed when multiple, connected actors pursue goals that a single actor pursuing the same goals would not perform”* [***Thomas W Malone. 1988. What is Coordination Theory?***]

or

- the process of “*managing dependencies between activities*”. [***Thomas W. Malone, and Kevin Crowston. 1994. The Interdisciplinary Study of Coordination. ACM Computing Surveys.***]

Coordination Strategies

1. Set of interdependent actors engaged in the environment;
2. who perform tasks of mapping goals to activities; and
3. to achieve the goals of better coordination performance

Thomas W. Malone. 1987. Modeling Coordination in Organizations and Markets. *Manag. Sci.* 33, 10 (October 1987), 1317–1332.

Research Questions

Research Questions

- How to detect coordination in online campaigns?
- Which social network measures help in assessing coordination?
- How to model coordination based on network measures?

Data Origin and Description

- In November 2017, the House Intelligence Committee released a list of accounts, given to them by Twitter, that were found to be associated with Russia's Internet Research Agency (IRA) and their influence campaign targeting the 2016 U.S. election.

- On October, 2018, Twitter released an archive of tweets shared by accounts from the Internet Research Agency, an organization in St. Petersburg, Russia, with alleged ties to the Russian government's intelligence apparatus.

Tweets	9,041,308
Total users	3667
Earliest Tweet	2009
Latest Tweet	June 21, 2018

Methodology

Coordination Framework

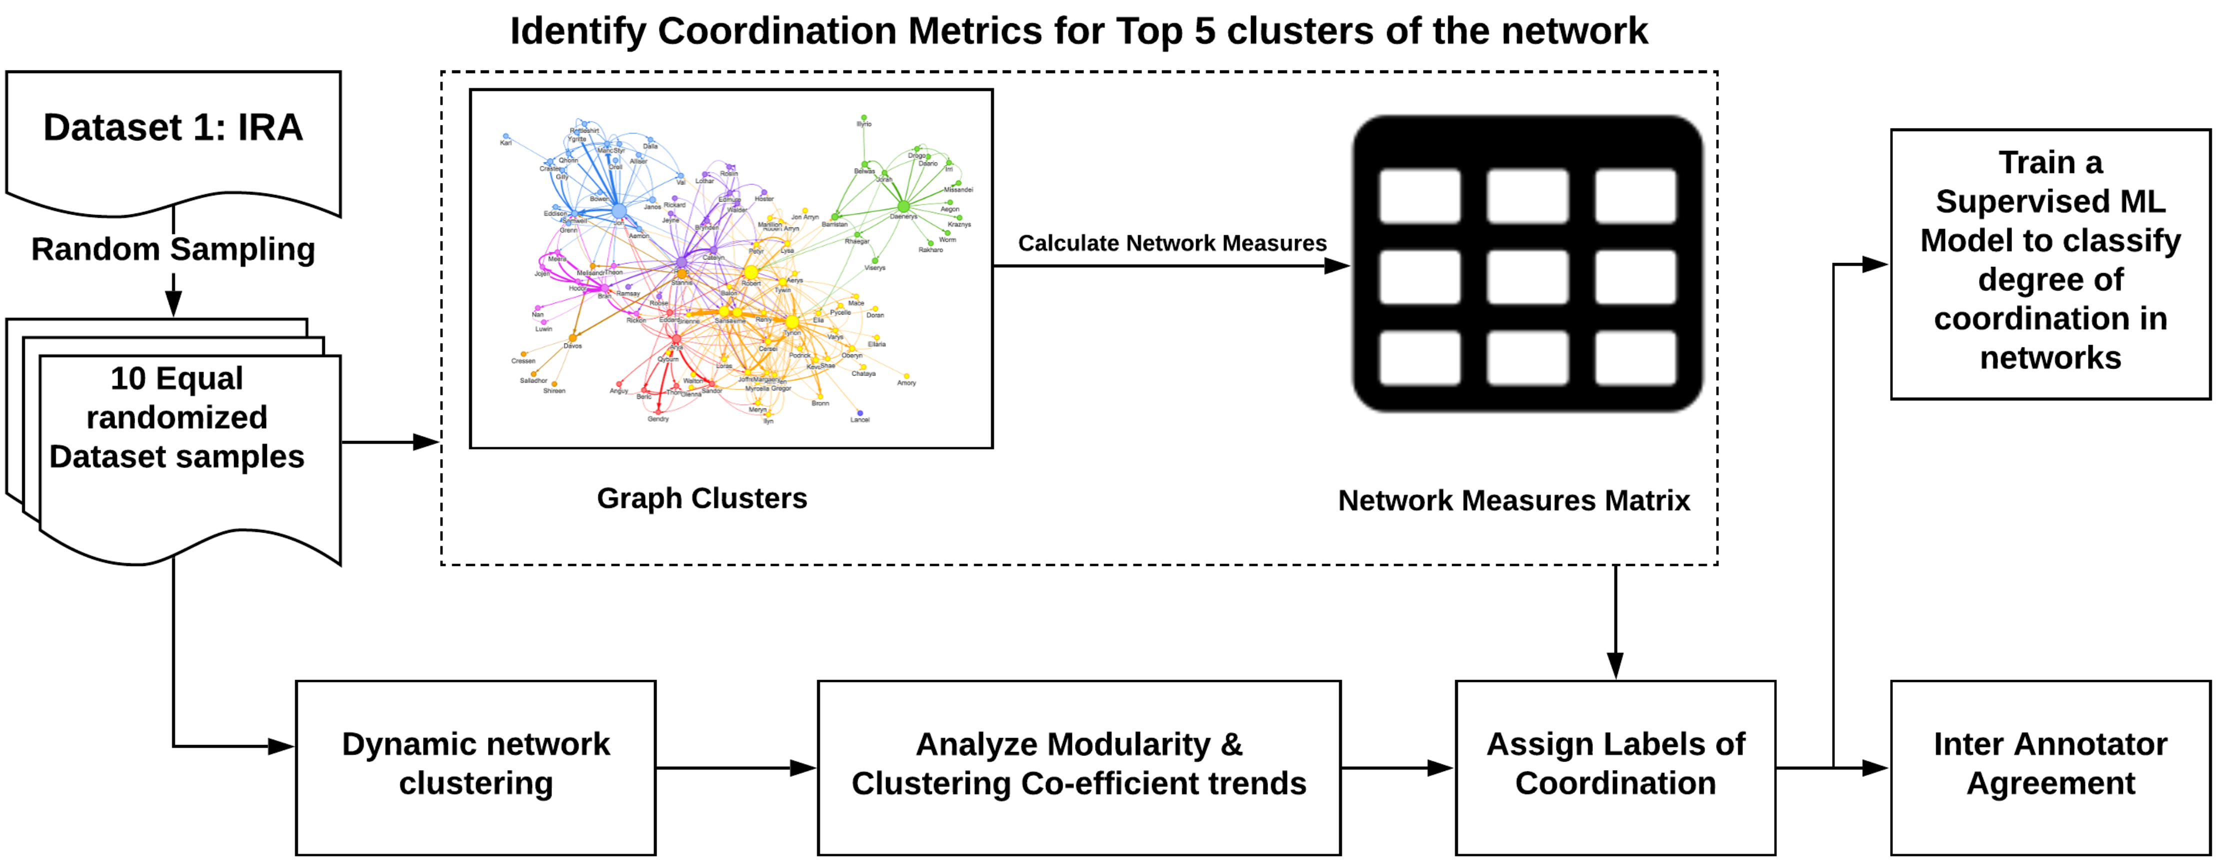

Network Types

1. *Communication network based on retweets and mentions*
2. *Interaction network based on links shared*
3. *Hashtag co-occurrence network*
4. *Hashtag co-usage network*
5. *Identical Tweet network*

Data Preprocessing

Random Sampling

1. Filtered the dataset first by only selecting the rows with hashtags
2. Divide the dataset into 10 equal sized samples with randomized tweets.
3. Each sample consist of ~250K tweets.

Each sample later created user-user shared hashtag networks

Analysis & Findings

RQ1: How to detect coordination in online campaigns?

Bots vs Humans - Social Network

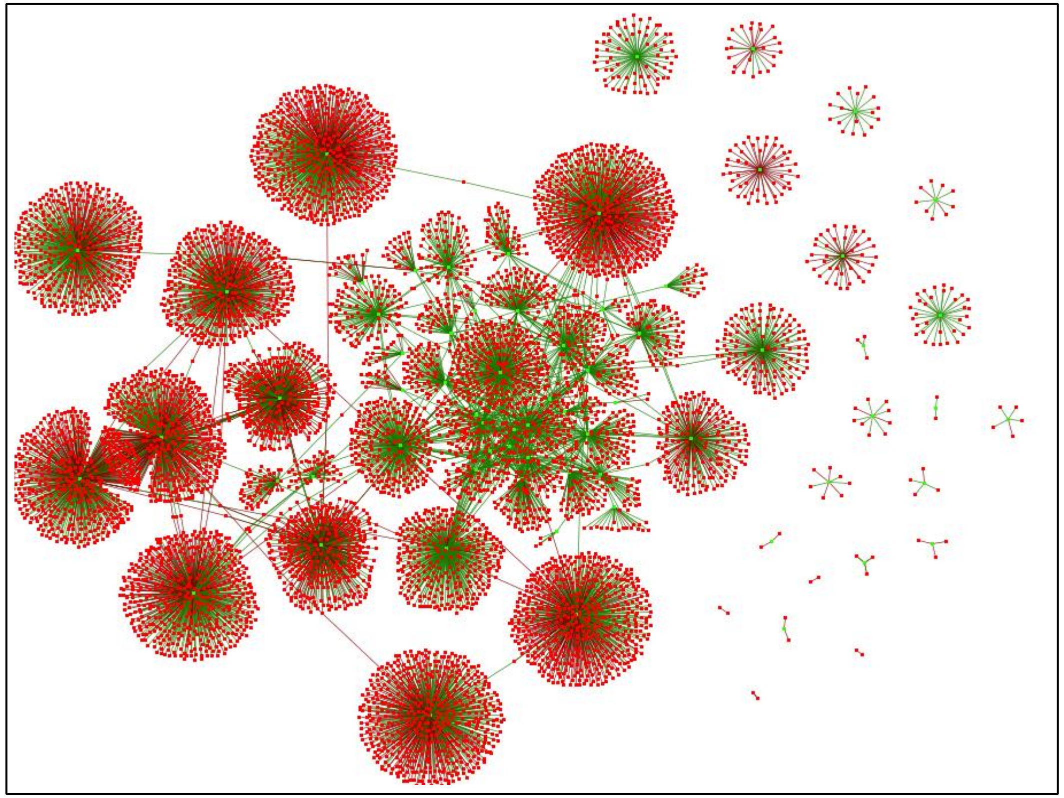

Bots Social Network

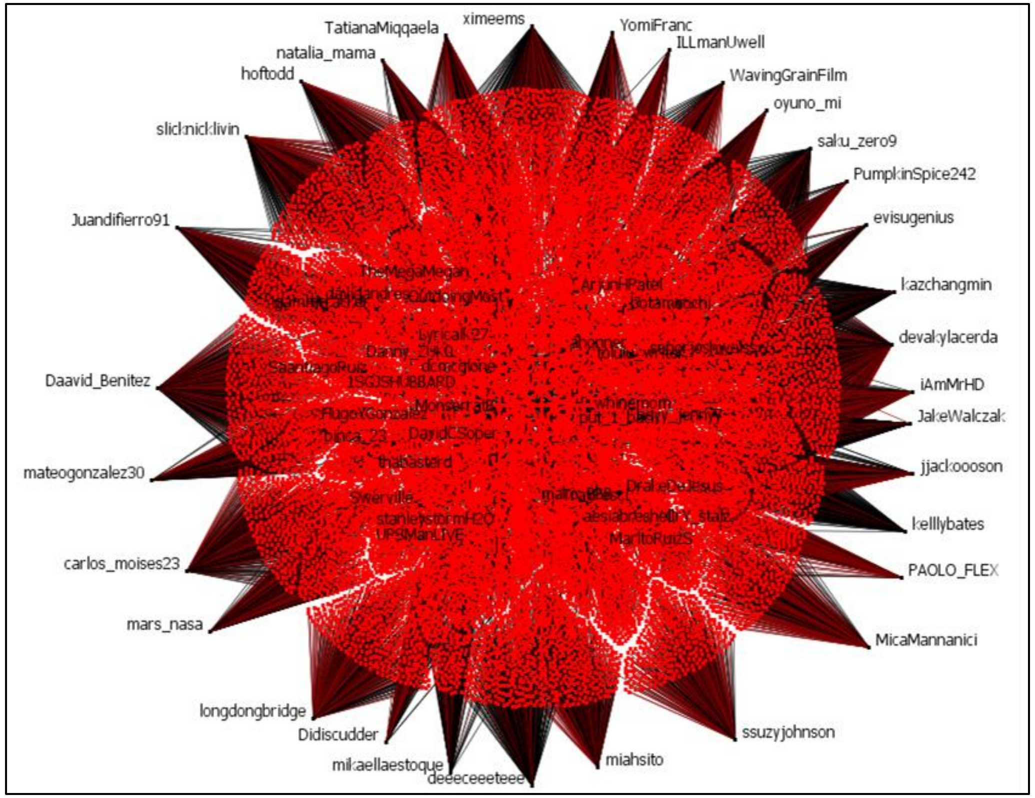

Humans Social Network

Core and peripheral network structure clearly visible in the Bot network.

There is only one large connected component in the Human network.

Content Analysis – 2017 Weather Events

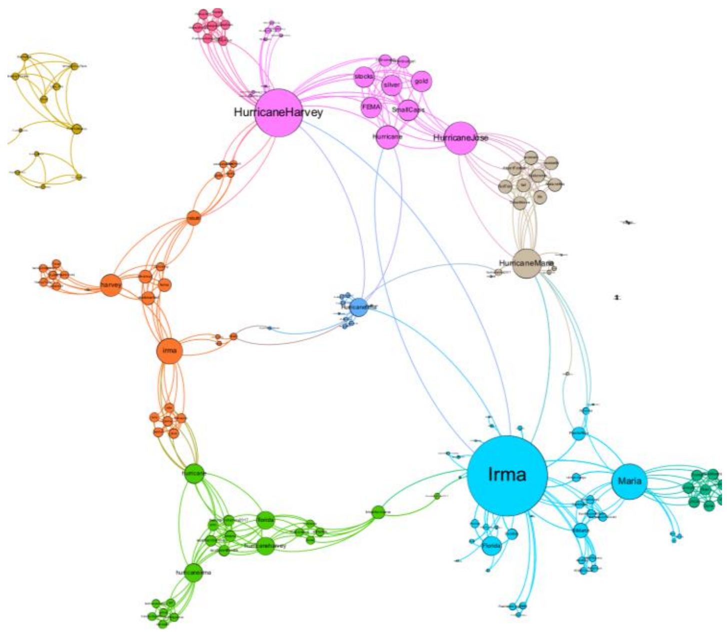

Bots Hashtag Co-occurrence Network

Bots have fewer, larger, and sparser communities.

Humans have more, smaller, and denser communities.

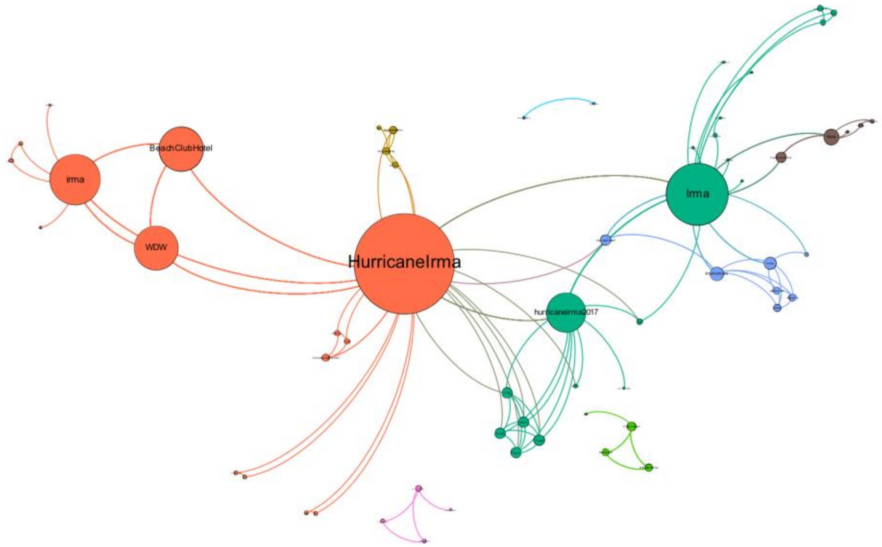

Human Hashtag Co-occurrence Network

Content Analysis – 2018 Sport Events

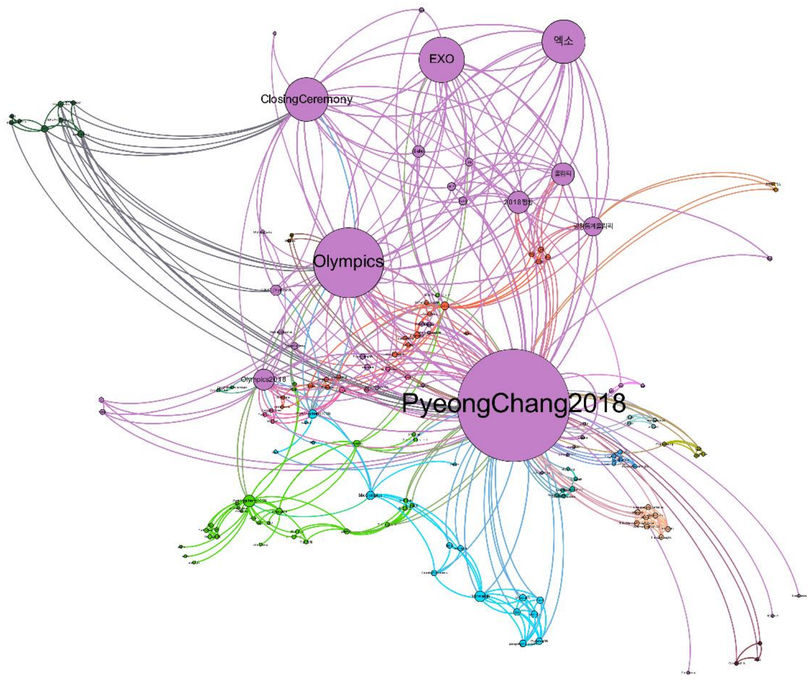

Bots Hashtag Co-occurrence Network

Bots have fewer, larger, and sparser communities.

Humans have more, smaller, and denser communities.

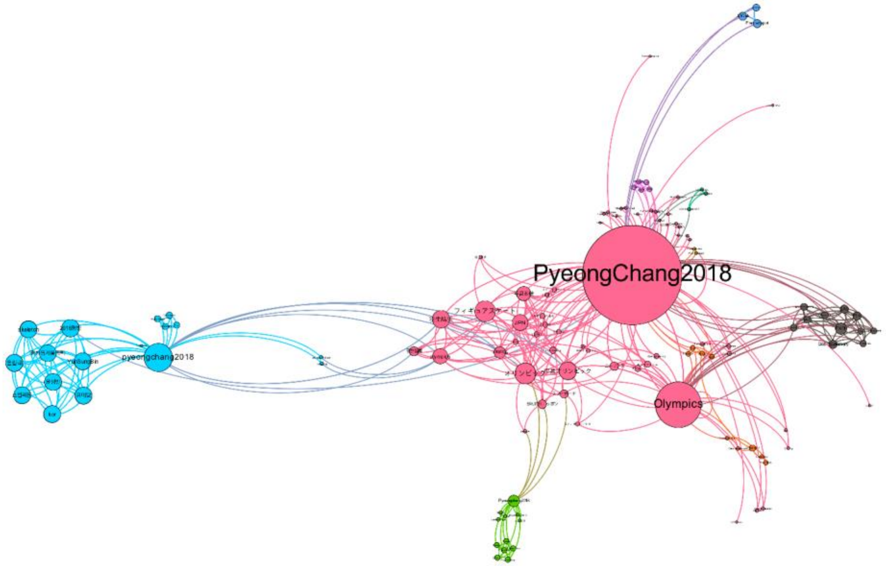

Human Hashtag Co-occurrence Network

Alternative Narratives during Crisis Events

Analyzing hashtags provides a way to track alternative narratives on Twitter. During 2017 weather crisis events, understandably, English and Spanish were the dominant languages for hashtags. However, Arabic, French, and Japanese among several other languages were also observed. On examining non-event related hashtags, **several alternate narratives were found**.

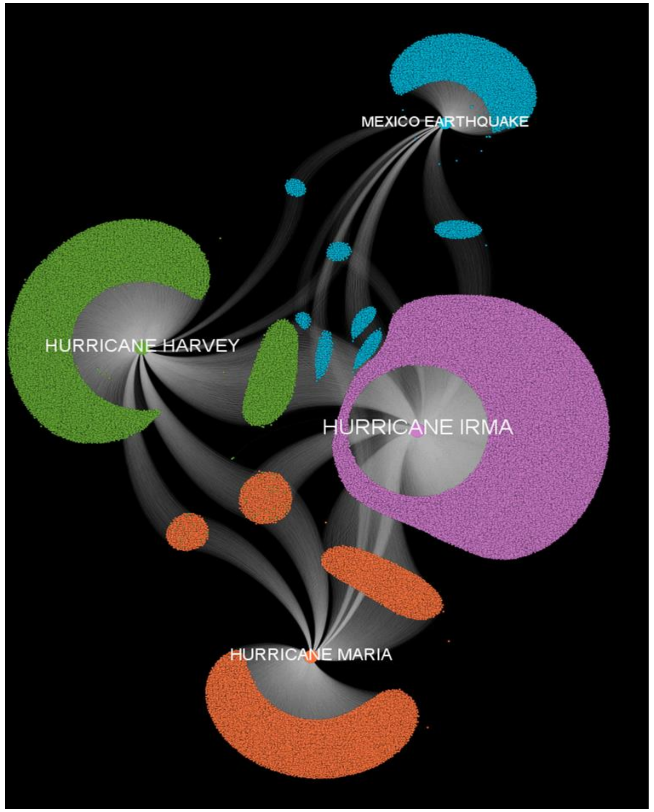

Hashtag Co-occurrence Network

Non-relevant Hashtag Clouds for various Languages. From top-left (clockwise) – English, Arabic, French, Mandarin, and Spanish.

Language	Hashtag (Translation)
English	#DACA, #BlackLivesMatter
Spanish	#VenezuelaDemocraciaYDiálogo (Venezuela Democracy and Dialogue), #Cáncer (Cancer)
Arabic	(The demise Of Israel) #زوال_إسرائيل, (The Jews) #اليهود
French	#Nucléaire (Nuclear), #GendarmerieEnOpération (Gendarmerie Special Operations)
Japanese	金正恩 (Kim Jong-un), 核试验 (Nuclear Test)

Alternative Narratives

Trend of Hoaxes during Crisis Events

During **Hurricane Sandy in 2012,** the following rumors were floating on social media.

- Sharks swimming through waterlogged suburban neighborhoods,
- Statue of Liberty engulfed in ominous clouds, and
- Floor of the New York Stock Exchange flooded.

Source: http://www.nytimes.com/2012/11/01/technology/on-twitter-sifting-through-falsehoods-in-critical-times.html

During **Hurricane Harvey in 2017**, again the “shark on highway” hoax went viral.

Source: https://www.washingtonpost.com/news/the-intersect/wp/2017/08/28/no-the-shark-picture-isnt-real-a-running-list-of-harveys-viral-hoaxes/

Misinformation during Crisis Events

During crisis events, misinformation is rampant. One of the most commonly spread hoaxes is “shark on highways”. Timeline below illustrates this hoax as it is propagated during various hurricanes in the U.S.

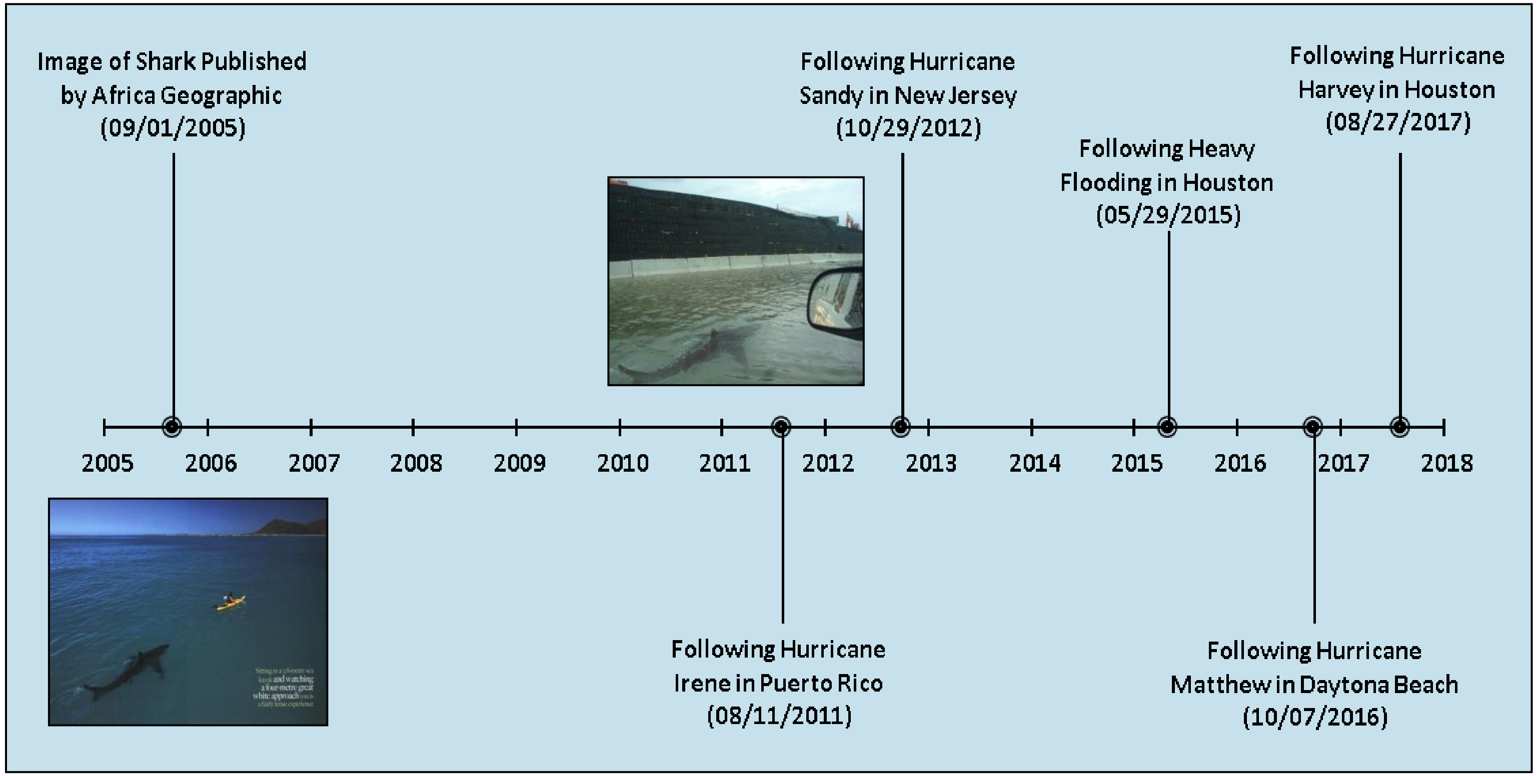

Timeline of activity of the shark hoax

Original image was published in Africa Geographic magazine in September 2005.

Used as hoax during Hurricane Irene in 2011, Hurricane Sandy in 2012, Houston Flood in 2015, Hurricane Matthew in 2016, and finally in our collected datasets for Hurricane Harvey, Hurricane Irma, and Hurricane Maria in 2017.

Hashtag Co-usage Network - IRA

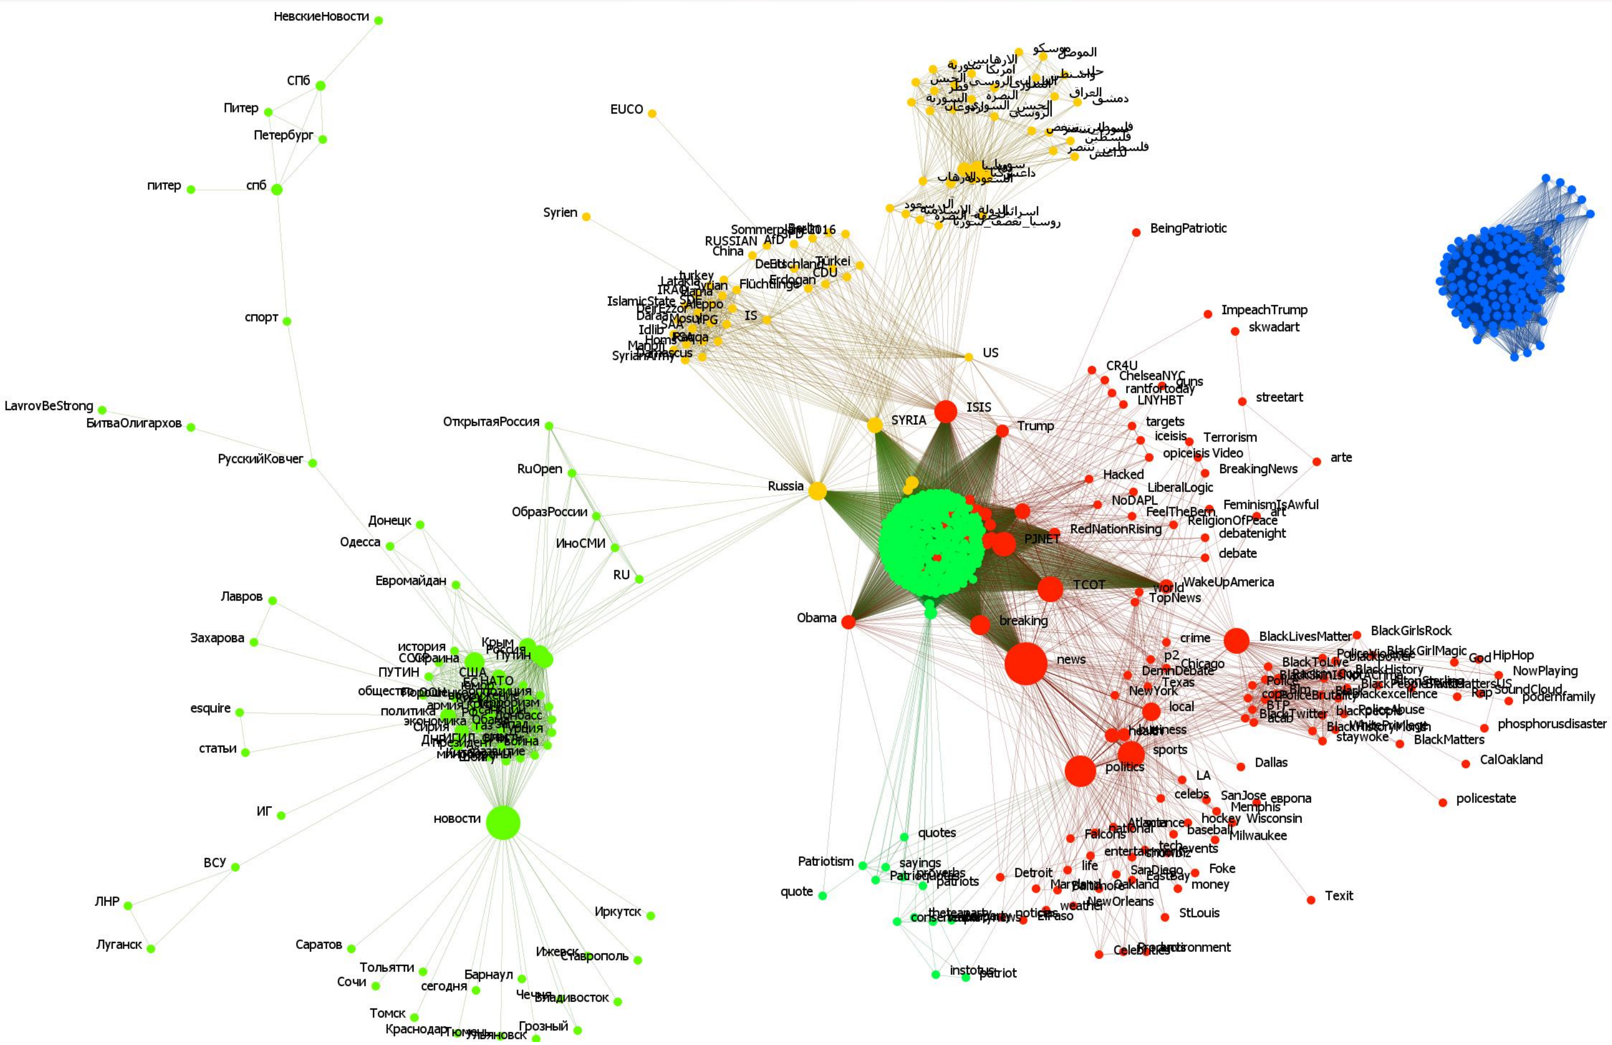

Analysis & Findings

RQ2: Which social network measures help in assessing bot coordination?

Network Metrics – Network Level

- Social Network Analysis (SNA) provides both a visual and a mathematical analysis of human-influenced relationships.
- *Network-level metrics* deal with how users are connected with one another and describe the interaction network among network users.

Measure	Definition
Size	Number of actors in the network
Inclusiveness	Total actors in a network minus the number of isolated actors
Network Diameter	The length of the longest shortest path between two nodes
Average Degree	Average number of links per node
Modularity	Measures the strength of division of a network into modules (also called groups, clusters or communities)
Clustering Coefficient	Measures the extent to which my friends are friends with one another.
Connected Component	A connected subset of network nodes and links
Connectivity (Reachability)	Extent to which actors in the network are linked to one another by direct or indirect ties
Connectedness	Ratio of pairs of nodes that are mutually reachable to total number of pairs of nodes
Density	Ratio of the number of actual links to the number of possible links in the network
Centralization	Difference between the centrality scores of the most central actor and those of all actors in a network is calculated, and used to form the ratio of the actual sum of the differences to the maximum sum of the differences.
Symmetry	Ratio of number of symmetric to asymmetric links
Transitivity	Number of transitive triples divided by the number of potential transitive triples
Clique	The maximum number of individuals in the network who are all directly connected to one another, but are not all directly connected to any additional individuals in the network

Network Level Measures

Network Metrics – Node Level

- *Node-level metrics* deal with how users interact with other users and describe the importance of a single node as opposed to the entire network.

Measure	Definition
Degree	Number of direct links with other actors
In-degree	Number of directional links to the actor from other actors (in-coming links)
Out-degree	Number of directional links from the actor to other actors (out-going links)
Range (diversity)	Number of links to different others (others are defined as different tot the extent that they are not themselves linked to each other, or represent different groups or statuses)
Closeness	Extent to which an actor is close to, or can easily reach all the other actors in the network
Betweenness	Extent to which an actor mediates, or falls between any other two actors on the shortest path between those actors
Centrality	Extent to which an actor is central to a network.

Node Level Measures

Network Measures

Cluster Type	Density	Netwo	Radius	Avg Path Le	Average Degree	Average Weighte	Modularity	Connected c	Clustering (	Degree Ce	Betweenness	Eigenvector	Size (node co	Edge count	Edge cou	Transitivity
											user - user shared #s					
sample 0	0.266	5	0	1.796	798.989	6549.43	0.592	41	0.763	0.458	0.007	0.091	3010	1205488	5065906	0.68
sample 0 top 5 clusters	0.272	5	0	1.796	809.205	6633.17	0.384	3	0.763	0.23	0.005	0.091	2972	1202478	9856892	0.68
sample 0 cluster 1	0.708	3	2	1.295	859.723	9377.274	0.077	1	0.893	0.117	0.007	0.074	1215	522282	5696694	0.875
sample 0 cluster 2	0.273	5	3	1.814	265.636	2768.978	0.327	1	0.775	0.266	0.005	0.269	973	129232	1347108	0.654
sample 0 cluster 3	0.611	4	2	1.403	477.33	3169.23	0.134	1	0.843	0.146	0.008	0.086	782	186636	1239170	0.818
sample 0 cluster 4	NaN	0	0	NaN	0	0	NaN						1	0		
sample 0 cluster 5	NaN	0	0	NaN	0	0	NaN						1	0		
sample 1	0.265	5	0	1.794	795.435	6532.422	0.59	31	0.763	0.445	0.007	0.092	3005	1198146	5045390	0.68
sample 1 top 5 clusters	0.27	5	0	1.794	802.916	6593.862	0.383	3	0.763	0.223	0.003	0.092	2977	1195141	9814964	0.676
sample 1 cluster 1	0.714	4	2	1.289	857.509	9411.045	0.078	1	0.893	0.111	0.007	0.077	1202	515363	5656038	0.877
sample 1 cluster 2	0.249	5	3	1.84	230.744	2290.944	0.246	1	0.783	0.279	0.006	0.304	928	107065	1062998	0.635
sample 1 cluster 3	0.603	4	2	1.408	507.936	3775.559	0.216	1	0.838	0.157	0.006	0.223	844	214349	1593286	0.813
sample 1 cluster 4	1	1	1	1	1	2	0	1	NaN				2	1		
sample 1 cluster 5	NaN	0	0	NaN	0	0	NaN						1	0		
sample 2	0.263	4	0	1.79	792.793	6509.667	0.58	46	0.763	0.467	0.006	0.091	3012	1196959	5039723	0.678
sample 2 top 5 clusters	0.272	4	1	1.79	805.087	6610.622	0.382	3	0.763	0.235	0.003	0.092	2966	1193944	9803552	0.678
sample 2 cluster 1	0.714	3	2	1.288	858.361	9388.276	0.076	1	0.894	0.115	0.007	0.076	1203	516304	5647048	0.878
sample 2 cluster 2	0.255	4	3	1.83	236.885	2323.948	0.245	1	0.771	0.262	0.009	0.304	931	110270	1081798	0.639
sample 2 cluster 3	0.617	4	2	1.387	510.362	3778.444	0.216	1	0.84	0.146	0.006	0.231	828	211290	1564276	0.812
sample 2 cluster 4	1	1	1	1	1	2	0	1	NaN				2	1		
sample 2 cluster 5	1	1	1	1	1	2	0	1	NaN				2	1		

Dynamic Network Analysis

sample_cluster	start_ts	end_ts	modularity	clustering_co-eff	nodes	edges	% nodes growth(cumulative)	% edges growth(cumulative)
s0_c1_1	2011-01-18 21:36:00	2011-07-18 21:36:00	0	0	1	0	0.08230452675	0
s0_c1_2	2011-01-18 21:36:00	2012-01-18 21:36:00	0	0	3	0	0.2469135802	0
s0_c1_3	2011-01-18 21:36:00	2012-07-18 21:36:00	0.11	0.889	21	109	1.728395062	0.02086995148
s0_c1_4	2011-01-18 21:36:00	2013-01-18 21:36:00	0.11	0.908	47	758	3.868312757	0.1451323232
s0_c1_5	2011-01-18 21:36:00	2013-07-18 21:36:00	0.114	0.903	57	1101	4.691358025	0.2108056567
s0_c1_6	2011-01-18 21:36:00	2014-01-18 21:36:00	0.115	0.9	66	1206	5.432098765	0.2309097384
s0_c1_7	2011-01-18 21:36:00	2014-07-18 21:36:00	0.126	0.91	405	55091	33.33333333	10.548133
s0_c1_8	2011-01-18 21:36:00	2015-01-18 21:36:00	0.109	0.804	1036	284649	85.26748971	54.50101669
s0_c1_9	2011-01-18 21:36:00	2015-07-18 21:36:00	0.072	0.886	1107	418513	91.11111111	80.13161472
s0_c1_10	2011-01-18 21:36:00	2016-01-18 21:36:00	0.072	0.885	1169	458660	96.21399177	87.81845823
s0_c1_11	2011-01-18 21:36:00	2016-07-18 21:36:00	0.072	0.887	1186	477482	97.61316872	91.42225847
s0_c1_12	2011-01-18 21:36:00	2017-01-18 21:36:00	0.075	0.891	1204	510150	99.09465021	97.67711696
s0_c1_13	2011-01-18 21:36:00	2017-07-18 21:36:00	0.076	0.893	1210	521177	99.58847737	99.78842847
s0_c1_14	2011-01-18 21:36:00	2018-01-18 21:36:00	0.077	0.893	1214	522274	99.91769547	99.99846826
s0_c1_15	2011-01-18 21:36:00	2018-07-18 21:36:00	0.076	0.893	1215	522282	100	100

Trend of Network Measures

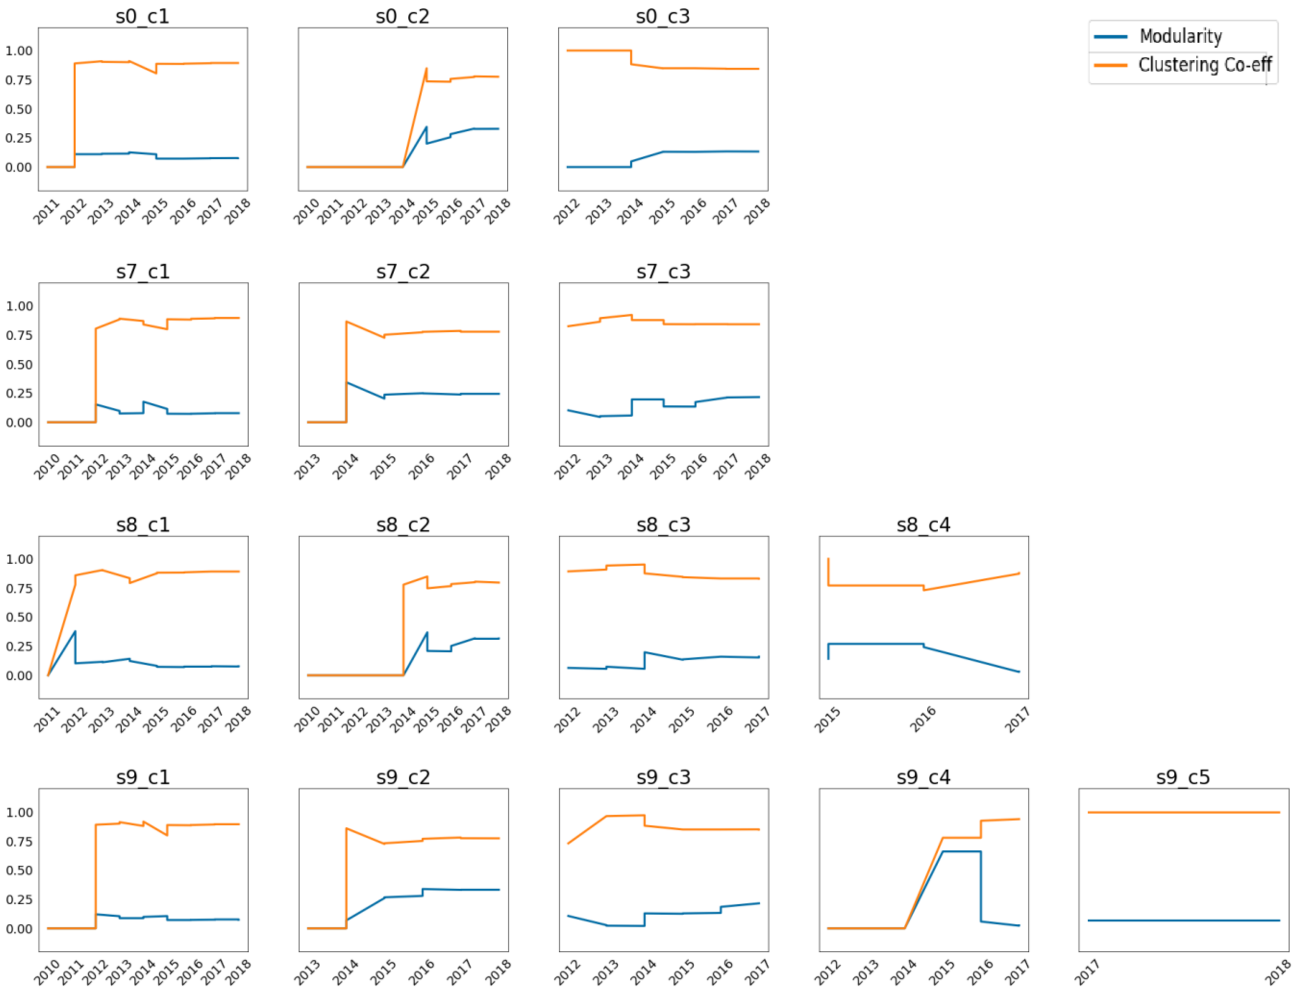

Analysis & Findings

RQ3: How to model coordination based on network measures?

Human Annotations

Sample Clusters	User 1	User 2	User 3	User 4	User 5	User 6
s0_c1	H	H	H	H	H	H
s0_c2	M	M	M	M	M	H
s0_c3	L	L	L	L	L	M
s1_c1	H	M	H	H	H	H
s1_c2	M	M	M	M	M	H
s1_c3	L	L	L	L	L	M
s2_c1	H	M	M	H	H	H
s2_c2	M	M	M	M	L	H
s2_c3	M	L	L	L	L	L
s3_c1	H	M	H	H	H	H
s3_c2	M	M	M	M	L	M
s3_c3	L	L	L	L	L	L
s3_c4	L	L	L	L	L	L
s4_c1	H	H	H	H	H	H
s4_c2	M	M	M	M	M	M
s4_c3	L	M	L	L	L	L
s5_c1	H	H	H	H	H	H
s5_c2	M	M	H	M	L	M
s5_c3	L	L	L	L	L	M
s5_c4	L	L	L	L	L	L
s6_c1	H	H	H	H	H	H
s6_c2	M	M	L	M	L	M
s6_c3	M	L	M	H	M	M
s6_c4	H	M	L	H	H	H
s7_c1	H	H	H	H	H	H

Inter-Annotator Agreement

- Fleiss' kappa is a statistical measure for assessing the reliability of agreement between a fixed numbers of raters when assigning categorical ratings to several items or classifying items.

$$Kappa\ (k) = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}$$

- Where, the factor $1 - \bar{P}_e$ gives the degree of agreement that is attainable above chance

- $\bar{P} - \bar{P}_e$ gives the degree of agreement actually achieved above chance.

- Kappa values ranges from 0 to 1.
  - If raters are in complete agreement then κ=1. If no agreement, then $\kappa \leq 0$.

Inter-Annotator Agreement

- Let,

  N be the total number of subjects (in this case, the network sample clusters),

  n be the number of ratings per subject (in this case, the annotators), and

  k be the number of categories into which assignments are made. (in this case, High, Moderate and Low)

- The subjects are indexed by i = 1,...N and the categories are indexed by j = 1, ... k. Let n_{ij} represent the number of raters who assigned the i^{th} subject to the j^{th} category.
- First , the proportion of all assignments which were to the j^{th} category is calculated:

$$p_j = \frac{1}{Nn}\sum_{i=1}^{N} n_{ij}, \qquad 1 = \sum_{j=1}^{k} p_j$$

Inter-Annotator Agreement

- Then, P_i, the extent to which raters agree for the i^{th} subject is calculated:

$$P_i = \frac{1}{n(n-1)} \sum_{j=1}^{k} n_{ij}(n_{ij}-1)$$
$$= \frac{1}{n(n-1)} \sum_{j=1}^{k} (n_{ij}^2 - n_{ij})$$
$$= \frac{1}{n(n-1)} \left[\left(\sum_{j=1}^{k} n_{ij}^2 \right) - (n) \right]$$

- Now, $\bar{P}$, the mean of P_i's, and $\bar{P}_e$ is computed which is later used to calculate κ:

$$\bar{P} = \frac{1}{N} \sum_{i=1}^{N} P_i$$

$$\bar{P}_e = \sum_{j=1}^{k} p_j^2$$

Inter-Annotator Agreement

- In order to calculate Kappa (κ), the following values are needed:
  - N = number of subjects = 36, n = number of annotators = 6, k = number of categories = 3, N * n = sum of all entries = 216

sample clusters	High	Moderate	Low	P_i
s0_c1	6	0	0	1
s0_c2	1	5	0	0.6666666667
s0_c3	0	1	5	0.6666666667
s1_c1	5	1	0	0.6666666667
⋮	⋮	⋮	⋮	⋮
s8_c2	1	3	2	0.2666666667
s8_c3	0	2	4	0.4666666667
s8_c4	2	3	1	0.2666666667
s9_c1	6	0	0	1
s9_c2	0	5	1	0.6666666667
s9_c3	0	3	3	0.4
s9_c4	3	1	2	0.2666666667
Total	71	72	73	24.6
p_j	0.3287037037	0.3333333333	0.337962963	

Inter-Annotator Agreement

- After calculating and for each row and column respectively, calculate the following:

$\sum_{i=1}^{N} P_i$ = Sum of P_i = 24.6; $\bar{P}$ = $\frac{24.6}{36}$ = 0.6833; $\bar{P}_e$ = Sum of the squares of p_j= 0.3334

$$\text{Kappa } (\kappa) = \frac{0.6833 - 0.3334}{1-0.3334} = 0.525$$

Inter-Annotator Agreement

Interpretation of Kappa

	Poor	Slight	Fair	Moderate	Substantial	Almost perfect
Kappa	0.0	.20	.40	.60	.80	1.0

Kappa	Agreement
< 0	Less than chance agreement
0.01–0.20	Slight agreement
0.21– 0.40	Fair agreement
0.41–0.60	Moderate agreement
0.61–0.80	Substantial agreement
0.81–0.99	Almost perfect agreement

Decision Trees

- Decision Tree is one of the popular tree-based algorithms used for supervised learning scenarios.
- They are easy to understand and visualize with great adaptability.
- Decision Trees consists of a *root node* that represents the entire population or sample and the respective *decision node* further gets divided into two or more homogeneous sets. The *leaf or terminal nodes* are the ones that no longer split.

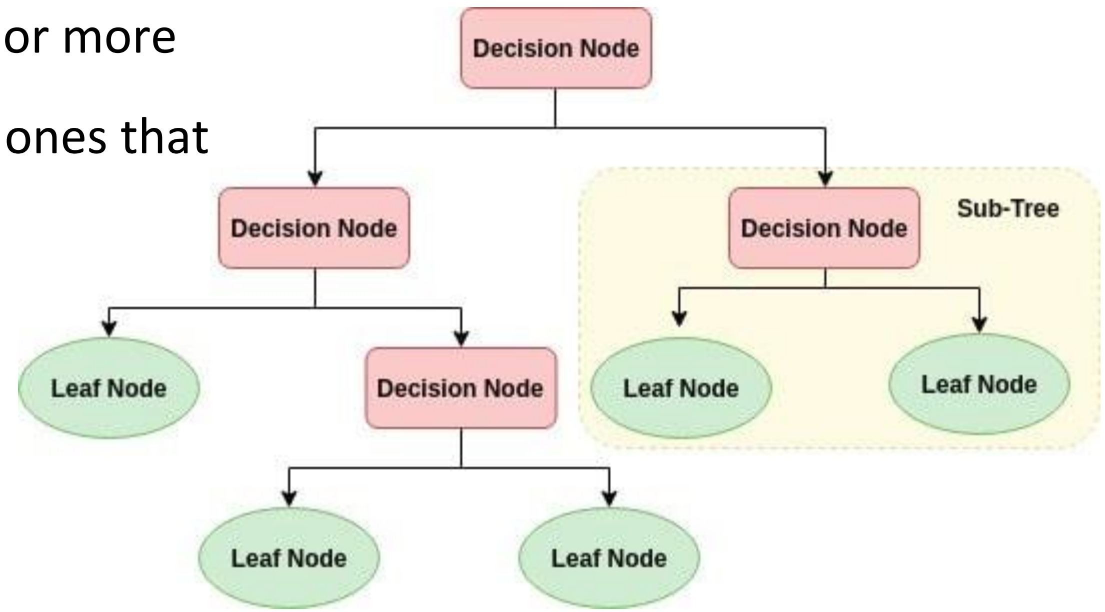

Decision Trees

- The accuracy of Decision Tree depends on effective splitting of the data based on a specific criterion.

1. Gini Index
2. Information Gain

Cluster Type	Density	Diameter	Degree	Modularity	Clustering Coefficient	Transitivity	Label
sample 0 cluster 1	0.708	3	859.723	0.077	0.893	0.875	High
sample 0 cluster 2	0.273	5	265.636	0.327	0.775	0.654	Moderate
sample 0 cluster 3	0.611	4	477.330	0.134	0.843	0.818	Low
sample 1 cluster 1	0.714	4	857.509	0.078	0.893	0.877	High
sample 1 cluster 2	0.249	5	230.744	0.246	0.783	0.635	Moderate

Decision Trees

- The feature variables used are Modularity and Clustering Co-efficient and the target variable is Label.
- The dataset was split into 70% for training and the remaining 30% was used for testing the model's performance.
- The accuracy was calculated by comparing the actual test set values and the predicted values.
  - It uses Gini Index as the splitting criterion and obtained a classification rate of 63.63% which is considered as good accuracy.

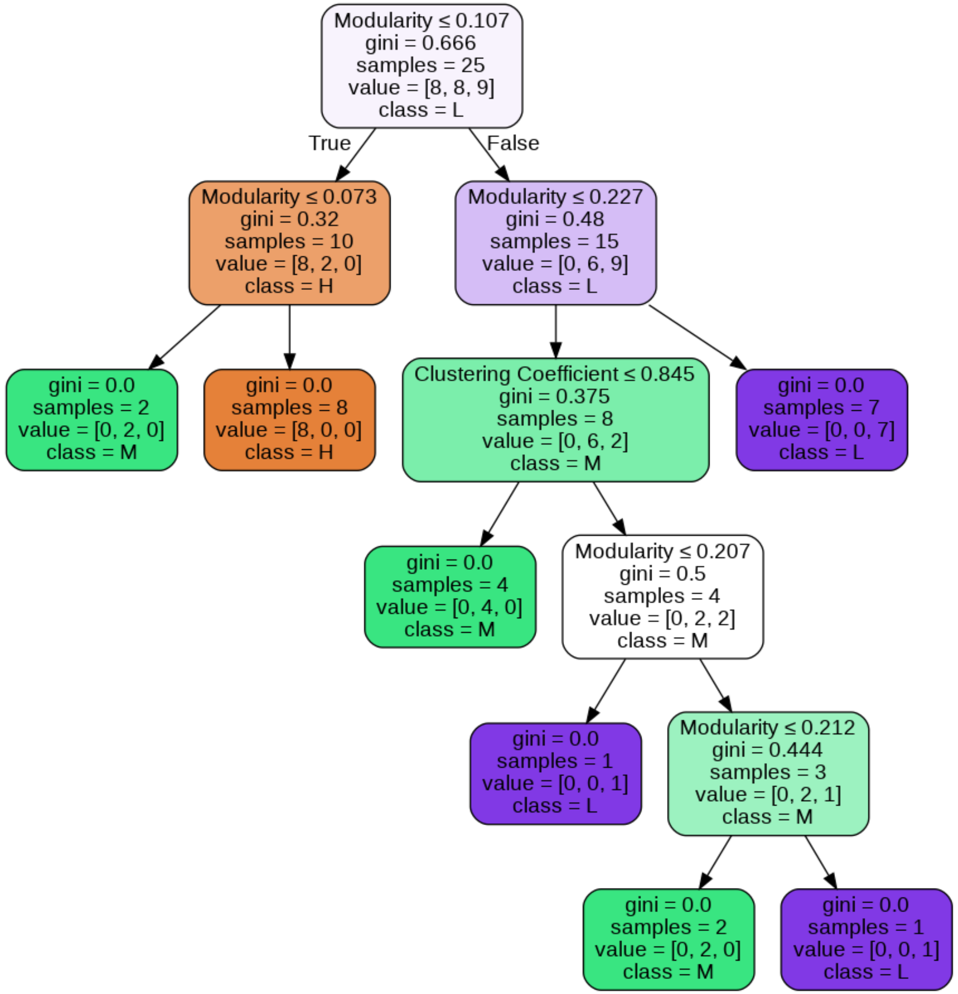

Decision Trees

- The classification rate increased to 81.82% which has better accuracy than the previous model.

- This pruned model is balanced, and easy to understand than the previous decision tree model plot.

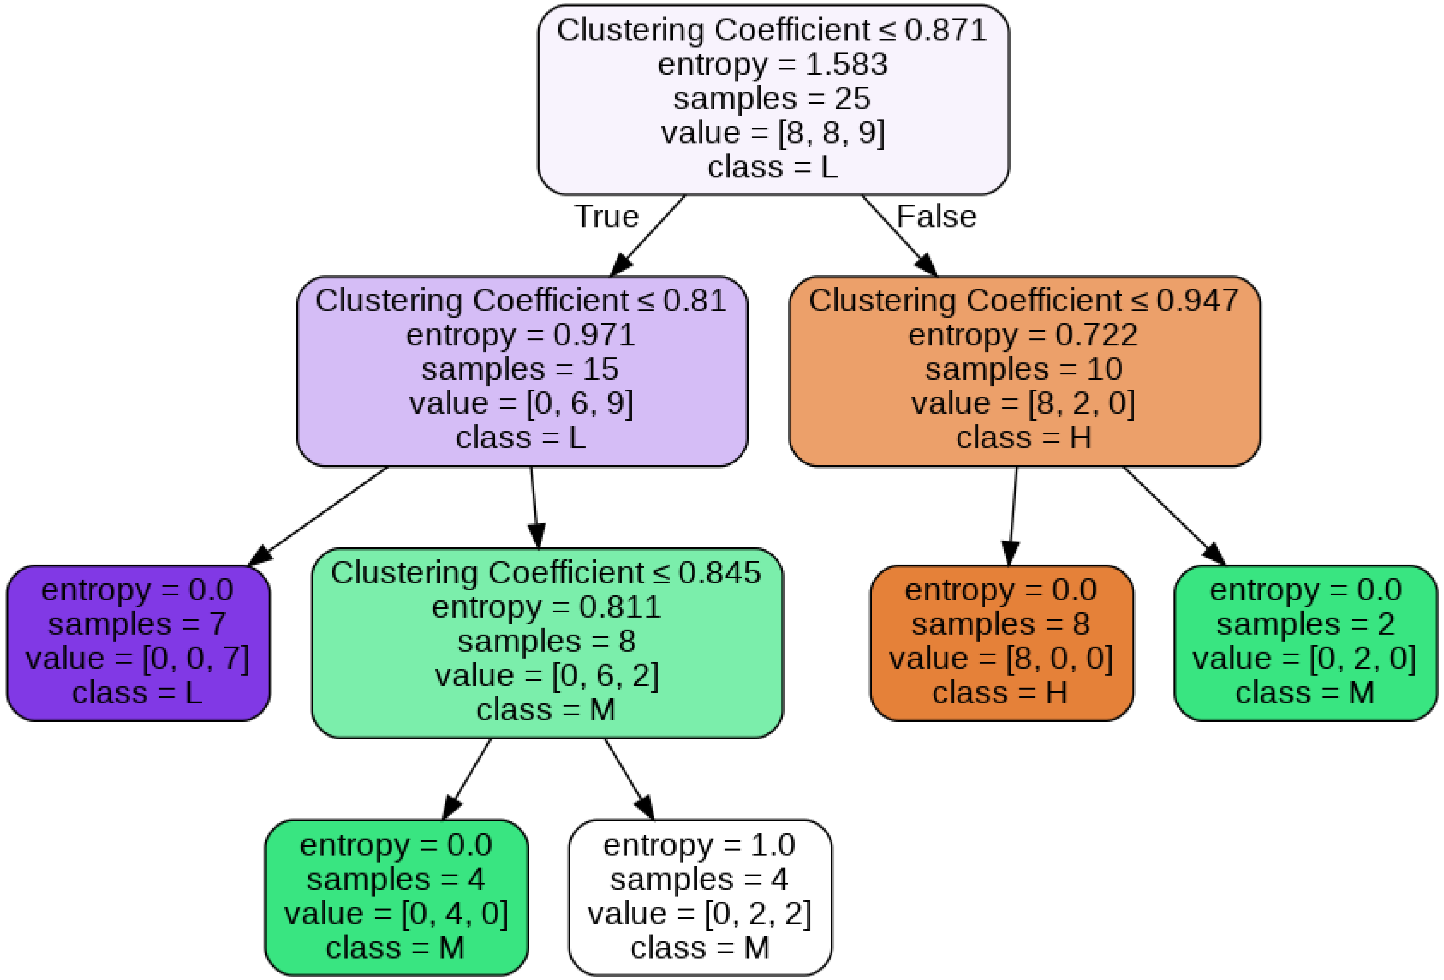

Decision Trees

Criterion	Labels	Precision	Recall	F1 score	Accuracy
Gini Index	High	1	0.50	0.67	0.64
	Moderate	1	0.60	0.75	
	Low	0.33	1.00	0.50	
Information Gain	High	0.80	1	0.89	0.82
	Moderate	1	0.60	0.75	
	Low	0.67	1	0.80	

Conclusion & Future Work

Conclusion

- Coordinated activity is often qualitatively analyzed and reported from a single user interaction networks‘ perspective.
  - One way to detect coordination is by visual analysis of Twitter communication networks to show how users are coordinating either as bots or non bots.

- This research measures as well as characterizes coordination.
  - The coordination behavior is characterized based on three different class labels such as highly, moderately or low coordinating.
  - Coordination is measured by proposing these values through a multi-dimensional classification problem where multiple features are taken into account.

- This research relies on data mining, especially a supervised machine learning model, Decision Tree to make detecting coordination as automated and explainable as possible.

Conclusion

- This research started with empirical observation where instances of coordinated inorganic activity revealed how social network analysis helps assess community structure over a period of time or during online campaigns, real-world events, etc.
- Then, humans annotated these network clusters, based on their various coordinating measures which were later used as input data to train a machine.
- All these efforts make this a social computing research which can be extended to multiple platforms such as YouTube or Facebook.

Limitations

- Random sampling without replacement changes the sample size every time which is why this research adopted a way sampling the dataset into ten equal chunks by randomly shuffling the index.

- This brings repetition in the following processes of the methodology which may discourage reproducibility.

- With two distinct network measures, this research was able to assess coordination but it is not restricted.

Future Work

- This research studies online information campaigns in general where coordination is measured at a group level.

- Future work can include detection of social bots, if present during the campaign, and study the evolution of their network structures over time and assess coordination.

- The labeled dataset can also be improved by using a rank-based labeled assignment and train it across other supervised machine learning models to compare the accuracies.

LEVERAGING SOCIAL NETWORK ANALYSIS AND SUPERVISED MACHINE LEARNING TO STUDY COORDINATION IN ONLINE INFORMATION CAMPAIGNS

Tuja Khaund

Information Science
University of Arkansas at Little Rock, USA