Identifying Coordinated Activities on Online Social Networks Using Contrast Pattern Mining

No Author Given

No Institute Given

Abstract. The proliferation of misinformation and disinformation on social media networks has become increasingly concerning. With a significant portion of the population using social media on a regular basis, there are growing efforts by malicious organizations to manipulate public opinion through coordinated campaigns. Current methods for identifying coordinated user accounts typically rely on either similarities in user behaviour, latent coordination in activity traces, or classification techniques. In our study, we propose a framework based on the hypothesis that coordinated users will demonstrate abnormal growth in their behavioural patterns over time relative to the wider population. Specifically, we utilize the EPClose algorithm to extract contrasting patterns of user behaviour during a time window of malicious activity, which we then compare to a historical time window. We evaluated the effectiveness of our approach using real-world data, and our results show a minimum increase of 10% in the F1 score compared to existing approaches.

Keywords: Coordination detection · Campaign detection · Contrast patterns · Pattern mining · Behavioural patterns.

1 Introduction

The usage of social media has dramatically escalated in the last decade. This is due to the factors such as peer pressure, the presence of communities that have developed on social media, and public interest in following popular and influential people on social media. However, the ease with which fake accounts can be created on social media has heightened the risk of misinformation and disinformation. Since disinformation campaigns are frequently politically driven, large numbers of user accounts are needed to disseminate their ideologies and achieve their desired objectives. Hence, coordination between accounts is required to carry out mass malicious campaigns.

The USA presidential election in 2016 was influenced by information operations carried out by Russia's Internet Research Agency (IRA) on Twitter and Facebook [10]. The Permanent Select Committee on Intelligence [13] identified 3,841 coordinated Twitter accounts and 470 Facebook pages that were affiliated with the IRA in 2017. In 2018, Twitter publicly released tweets and users related to this case. Additionally, in 2019, the UK general elections were influenced by coordinated users that polarized political opinions on Twitter [11].

Even though social media platforms claim to take measures to mitigate malicious campaigns by inspecting the behaviours of individuals, it is harder to identify a campaign as a whole due to its apparent natural growth and organic behaviour. Hence, identifying coordination from a broader perspective is an important step to identify malicious campaigns. We believe that summarizing behavioural patterns within a community is a promising approach to achieve this broad perspective.

Existing techniques to identify coordinating behaviours can be categorized into either (1) content-driven, (2) network structure-driven, or (3) activity trace-driven. Content driven methods [6, 7] are limited to the assumption that coordination is reflected in the theme of posts while other behavioural aspects are overlooked. Network-based approaches [3, 9, 11, 12, 22] tend to define coordination in terms of community detection on user similarity graphs. A major limitation of network-based approaches is that they perform well only when the networks are sufficiently sparse [9]. In contrast, Zhang et al. [23] and Sharma et al. [15] define coordination in terms of the synchronicity of users over time. They try to identify coordinated users using masked self-attention [20] to encode the event history.

An important alternative approach for summarizing changes in a dataset is pattern mining. Although pattern mining has been used in areas such as medicine [14, 16–18], education [5, 19], and computer security [1, 2], pattern mining on social networks has only been used for bot detection [8]. In this work, we utilize pattern mining in order to identify anomalous coordinating activities in online social networks.

The notion of patterns provides a means of encoding the behavioural patterns of users. Our hypothesis is that coordinated users will exhibit unusual correlated activity patterns that are not reflected among the wider population of users over time. If we can compare the frequencies of the behavioural patterns in the present-time behaviour with activity in an earlier reference time period that is assumed to be normal, we can interpret the growth of frequencies compared to the background as anomalous behaviour. Such patterns can be identified as anomalous patterns. This method of comparing patterns from two sets of datasets is known as contrast pattern mining.

Our experiments show that the social media users that are associated with contrasting behavioural patterns compared to their historical behaviours are likely to be coordinated in nature. We achieve F1 scores up to 86% in identifying coordinating users for the IRA dataset, thus supporting our hypothesis. Moreover, the accuracy of our approach in terms of F1 score exceeds the corresponding accuracy of a range of benchmark approaches by more than 10%.

We note the following as our contributions: (1) Formulating the usage of contrast pattern mining for identifying coordinated users, (2) Proposing a framework for making use of contrasting behavioural patterns for real data, (3) Conducting experiments comparing different parameters, attributes and approaches on real-life social network data to establish our claims.

In Section 2, we provide the background and definitions needed for our framework. In Section 3, we formulate our research problem. The methodology is given in Section 4. The experiments, results and analysis for those results are presented in Section 5. Finally, we give a conclusion of our study and identify directions for future research in Section 6.

2 Preliminaries

We refer to an interaction made by a user with the social network as an *event*. An event can be stored using a list of (attribute, value) pairs. Values can be either numerical or categorical. The domain of values for an attribute a is denoted Δ_a. An *item* is an (attribute, value) pair. A *transaction* is a set of items. A transaction is associated with a *transaction id*. For example, a single tweet is a transaction. In that case, the `tweetid` is the transacation id. The item (`username`, @abc) reflects the author of a tweet in a transaction. A list of such transactions is a *transactional dataset*. *Attribute space* $\mathcal{A}$ is the set of all attributes in a given dataset. A *pattern* or an *itemset* is a set of items. We say an itemset X is contained in transaction T iff. $X \subseteq T$. $f_D(X)$, the set of transactions that contain the pattern X is defined as $\{T \in D \mid X \subseteq T\}$. The number of transactions in a dataset D that contain pattern X is the *support count* of that pattern i.e., $SC(X, D) = |f_D(X)|$. The *support* of a pattern X is defined as $supp(X, D) = \frac{SC(X,D)}{|D|}$.

Patterns in the form of itemsets provide an opportunity to encode the behavioural patterns of users. An example of a behavioural pattern encoded as an itemset is, $\{$(*user*, u_1), (*is retweet?*, yes), (*original tweet's author*, BBC News), (*day of week*, Monday), (*time of day*, 8 AM - 10 AM)$\}$. If we compare the support of such itemsets in a historical time window with the support of those itemsets in a subsequent time span containing anomalous activities, we can interpret the growth in support as anomalous behavioural patterns. Such itemsets with high growth can be identified as contrast patterns. The following are some key definitions related to contrast pattern mining. The main dataset that is to be analysed and compared to other datasets is called the *target dataset* D_t. A baseline dataset against which changes in the target dataset are found is called the *background dataset* D_b. The *growth rate* of a pattern is the ratio of its supports between the target and background datasets $gr(X, D_t, D_b) = \frac{supp(X,D_t)}{supp(X,D_b)}$. If $supp(X, D_b) = supp(X, D_t) = 0$, then $gr(X, D_t, D_b) = 0$ and if $supp(X, D_b) = 0$ and $supp(X, D_t) > 0$, then $gr(X, D_t, D_b) = \infty$. *Support delta* is another way of measuring the growth of support of a given pattern, and is defined as $supp_\delta(X, D_t, D_b) = supp(X, D_t) - supp(X, D_b)$. A *contrast pattern* X is a pattern whose support in the target is significantly different from the background. Given a *growth rate threshold* $\rho > 1$ and a *minimum support delta* $\sigma_\delta > 0$, we say pattern X is a contrast pattern iff $gr(X, D_t, D_b) \geq \rho$ or $supp_\delta(X, D_t, D_b) \geq \sigma_\delta$. A contrast pattern p takes the following form: $p = \{(a, v) \mid a \in \mathcal{A}, v \in \Delta_a\}$.

A pattern X is called a *closed pattern* iff there exists no superset Y of X satisfying $SC(Y, D) = SC(X, D)$. A pattern X is a *closed contrast pattern* (CCP) iff $supp(X, D_t) \geq \sigma > 0, gr(X, D_t, D_b) \geq \rho > 1$ and X is a closed pattern in $D_t \cup D_b$. Here, σ is called the *minimum support.*

3 Problem Statement

The challenge we address is how to identify anomalous coordinting behaviour among social media users in a robust manner. This requires a novel approach to identify a combination of features from user posts that succinctly characterise the change in behaviour in contrast to normal behaviour. To address this challenge, we build upon the theory of contrast pattern mining, which provides a robust and scalable method for exploring the huge search space of possible feature combinations.

Let D be a set of posts in an online social network in a time interval $[t_s, t_e]$. Each element of the set D takes the form of a transaction. Say we determine two time intervals $[t_0, t_1]$ and $[t_2, t_3]$ such that $t_s \leq t_0 < t_1 \ll t_2 < t_3 \leq t_e$, $[t_2, t_3]$ presumably contains anomalous activities based on observations, and $[t_0, t_1]$ presumably does not contain anomalous activities. Let D_b and D_t be the subsets of D such that each post in D_b is created in the time interval $[t_0, t_1]$ and each post in D_t is created in the time interval $[t_2, t_3]$.

Problem Definition. *Given two transactional datasets of posts (D_b and D_t), find the behavioural patterns that have a significant growth from a historical time span $[t_0, t_1]$ that is likely to have few anomalous coordinating behaviours to a subsequent time span $[t_2, t_3]$ that is highly likely to have anomalous coordinating behaviours. Find the users that are associated with such behavioural patterns and test the hypothesis that: users who show anomalous behaviour are likely to be associated with contrasting behavioural patterns compared to their historical behaviour and other normal users.*

4 Methodology

This section outlines the proposed framework to solve the problem we identified above. Initially, we determine the background and target time intervals based on observations of the dataset of posts. Subsequently, we pre-process the data to extract relevant attributes and users. We then apply a contrast pattern mining algorithm to the converted transaction tables. Finally, we extract suspicious users by using the contrast patterns that we have obtained. An overview of our framework is presented in Figure 1.

4.1 Overview

Given background and target transactional datasets, we derive the set of contrast patterns $\mathcal{P}$ using a contrast pattern mining algorithm. Given a set of attributes $A \subseteq \mathcal{A}$, the subset of $\mathcal{P}$ where each contrast pattern is associated with every attribute in A is the set of filtered contrast patterns $\mathcal{P}_A$. Define $attribs(p) = \{a \mid$

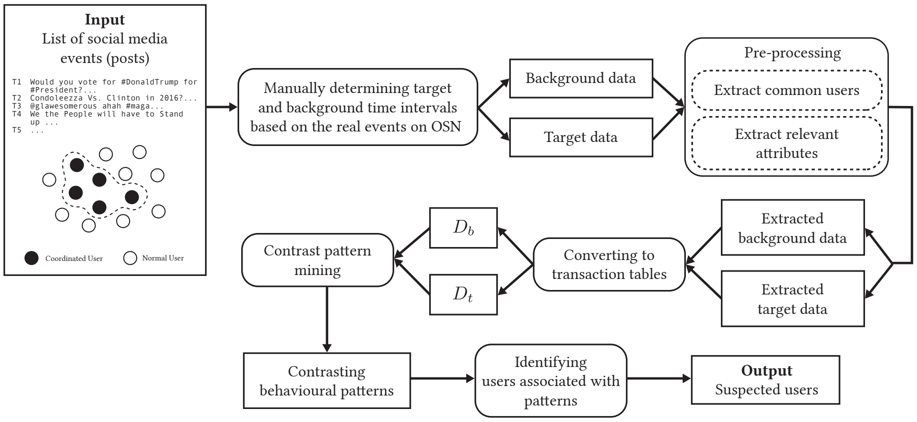

```
T1  Would you vote for #DonaldTrump for
    #President?...
T2  Condoleezza Vs. Clinton in 2016?...
T3  @glawesomerous ahah #maga...
T4  We the People will have to Stand
    up ...
T5  ...
```

Input
List of social media events (posts)
Coordinated User
Normal User
Manually determining target and background time intervals based on the real events on OSN
Background data
Target data
Pre-processing
Extract common users
Extract relevant attributes
Extracted background data
Extracted target data
Converting to transaction tables
D_b
D_t
Contrast pattern mining
Contrasting behavioural patterns
Identifying users associated with patterns
Output
Suspected users

Fig. 1. Overview of the proposed framework to identify suspected coordinated user accounts based on contrast pattern mining. Boxes with rounded corners represent a process while rectangular boxes represent data. The dark circles denote coordinating users and empty circles denote normal users.

$a \in \mathcal{A}, (a, v) \in p\}$, i.e., the set of attributes of a given contrast pattern p. Then, $\mathcal{P}_A = \{p \mid A \subseteq attribs(p)\}$. $\mathcal{P}_A$ can be obtained by filtering patterns in $\mathcal{P}$, which only contains all attributes in A. Since $\mathcal{P}_{\{user\}}$ contains the contrast patterns with emerging behavioural patterns for users, we claim that the users that appear in $\mathcal{P}_{\{user\}}$ are the set of suspicious of users. The set of users $U_{suspicious}$ that appear in $\mathcal{P}_{\{user\}}$ is marked as identified anomalous coordinating users. Formally, $U_{suspicious} = \{u \mid p \in \mathcal{P}_{\{user\}}, (user, u) \in p\}$.

4.2 Pre-processing

- *Extracting common users.* For a user to contribute to a contrast pattern, that user must be active in both D_b and D_t. Hence, common users are extracted for the sake of efficiency.
- *Extracting relevant attributes.* Categorical values are required for the purpose of grouping similar itemsets due to the low likelihood of matching real-valued observations for equality in practice. Binning or other pre-processing steps are needed for numerical values. For a Twitter dataset, the following fields were selected such that the attribute space consists of categorical values only – user id, user reported location, tweet language, tweet time – divided into two fields; day of week and time of day (12 equal sized time slots per day), tweet client, is the tweet a retweet?, author of the original tweet if retweeted, list of segmented hashtags - each hashtag segmented using a Twitter corpora, and list of user mentions. The fields in the form of a list, such as hashtags and user mentions, were flattened out in transactions.

- *Converting to transactional datasets.* Posts in the extracted datasets are converted to lists of (attribute, value) pairs, i.e., transactions. For multivalued attributes, an itemset is generated with the same attribute.

4.3 Mining coordinated users

- *Contrast pattern mining.* EPClose [1] is a fast scalable algorithm that extracts closed contrast patterns during closed pattern generation. It should be noted that applying a threshold to support count (SC) is equivalent to applying a corresponding threshold to support (*supp*) since $supp \propto SC$. By applying the threshold to D_t, the algorithm outputs patterns whose support count in D_b is 0, which is unexpected since we are interested in patterns whose support is growing from a non-zero value in D_b to D_t. Thus, we modify the threshold to support in D_b instead of D_t. Here onwards, this article refers to that threshold as the *minimum support* with the symbol σ.
- *Identifying users associated with patterns.* The users that appear in contrast patterns are extracted and marked as coordinating users. $\mathcal{P}_{\{user\}}$ is constructed using $\mathcal{P}$ from the last step.

5 Experiments

We tested our proposed framework on real-world data. Experiments were devised with the aim of investigating the following questions.

1. How well can the model discriminate coordination and normal behaviour, and how sensitive are the results to the choice of growth rate thresholds and minimum supports?
2. Which attributes reveal coordination the best?

5.1 Data

We experiment on the dataset of the activity of Russia's Internet Research Agency (IRA) influencing the 2016 USA presidential elections [10, 13], which consists of confirmed coordinated activities. This is a widely used dataset for detecting coordination [15, 21–23] due to the availability of ground truth. The dataset consists of 8.76 million tweets posted by 3613 users. Figure 2 shows the distribution of activity across the time.

In order to test the effectiveness of a coordination detection model, we introduce a set of noisy background events to the IRA dataset, since the IRA dataset only contains the set of coordinating users. The criteria that were used to extract noise data were: posted time between 2008 and 2018, marked location anywhere in the USA, contains either one of the following hashtags - *Election2016*, *MAGA*, *MakeAmericaGreatAgain*, *AmericaFirst, DonaldTrump*, *WakeUpUSA*, *Trump*, *TrumpTrain*, *HilaryClinton, Trump2016*, *DrainTheSwamp*, *TrumpPence16*, *tcot*, *POTUS*, *GOP*, *Resist*, *UniteBlue*, *NeverHillary*, *ElizabethWarren*, *WeThePeople*, *IllegalAliens*, *TrumpRussia, ImWithHer*, *GayHillary*, *WakeUpAmerica.* The above set of hashtags were the top-occurring hashtags in the original IRA dataset. The background data of normal users consists of 2.80 million tweets from 333 thousand of users.

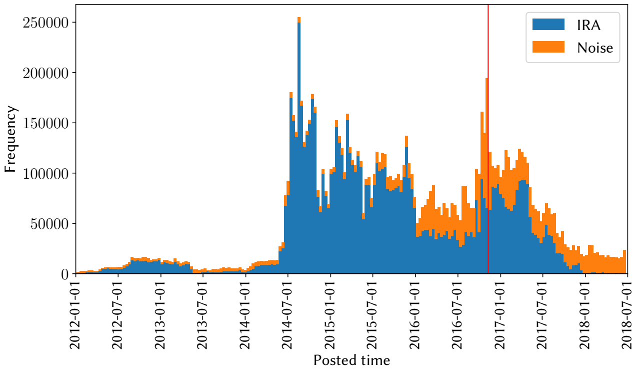

Fig. 2. Stacked distribution of IRA activities and extracted noise tweets across time. The bin size for the x-axis is 1 million seconds (~11.6 days). The red vertical line shows the election date.

5.2 Experimental Setup

Time intervals. Dataset D_t was chosen such that it includes the election time period (November 2016). The range of the posting times of tweets in D_t was the period of four months ending in November 2016. The time interval for the dataset D_b was the period of four months ending in May 2015. This is entirely based on observations.

Extracting top users. In addition to the step described in Section 4.2, for the purpose of evaluation, we further extract the top n_C number of coordinating users and top n_N number of normal users from the common users from the background and target datasets using posting frequency for the sake of performance. For the presented results, n_C and n_N were chosen as 400.

Baselines. We choose the following baselines in order to compare our results.

1. *Tweet frequency growth.* Instead of counting patterns and comparing supports, we formulate a comparison between the frequencies of tweets of each user in order to verify that our results are not due to the general growth of tweeting frequency that we see in the tweet density plots. Say the frequency of a user u posting in dataset D is $freq(u, D)$. Then for a given σ and ρ, we can check the following requirements; (a) $freq(u, D_b) \geq \sigma > 0$.

(b) Say $g(u, D_b, D_t) = \frac{freq(u,D_t)/|D_t|}{freq(u,D_b)/|D_b|}$ and $g(u, D_b, D_t) \geq \rho > 1$. If (a) and (b) satisfies, then we mark them as suspecting users.

2. *Tweet language.* Since most (82%) of the data in the coordinated set of users are in Russian and most (93%) of the data in the noise data are in English, we compare our results with the results of a model that only use the language to determine the coordinated status. This model simply classifies a user to be coordinated if the language is Russian.
3. *LCN+HCC [22].* This approach aims to identify coordinated communities using community detection on user similarity graphs. The temporal aspect is considered by a windowing mechanism. For a fair comparison, we use the dataset $D_b + D_t$ as the input. We use the window size as 10 days as the window length parameter.
4. *QT-LAMP-EP-BH [4].* Replacing EPClose by the above contrast pattern algorithm. In QT-LAMP-EP-BH, false discovery rates (FDR) are controlled using Benjamini-Hochberg (BH) method.
5. *QT-LAMP-EP-BY [4].* Instead of BH, use Benjamini-Yekutieli (BY) method to control FDRs.
6. *AMDN-HAGE.* [15] The SOTA for identifying coordinated users. We use the same set of hyperparameters except the threshold to determine the output influence values. Instead, we maximize the F1 score to determine it.

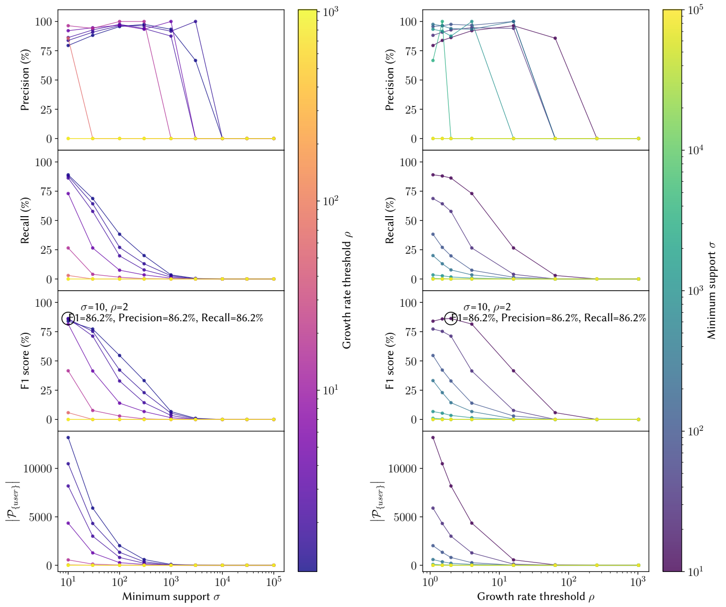

Fig. 3. The variation of precision, recall, F1 score and the number of contrast patterns that are associated with users ($|\mathcal{P}_{\{user\}}|$) with the variation of σ (left) and ρ (right). Target time period: 2016/07 - 2016/11. Background time period: 2015/01 - 2015/05.

5.3 Results

Quantitative Results. The variation of precision, recall, F1 score and the number of contrast patterns associated with users with respect to variation in the growth rate threshold (ρ) and minimum support (σ) is shown in Figure 3. The parameters and the performance metrics corresponding to the experiment with the highest F1 score is circled. There is a general increase in precision when σ and ρ is increased until a failure point. By increasing σ and ρ, less relevant contrast patterns get filtered out. Promising precision values prove that the users that are associated with contrast patterns are likely to be coordinating users, thus supporting our statement in the problem definition. For large enough σ and ρ, there are no contrast patterns, to say nothing of contrast patterns that are associated with users. Hence, the precision and recall is zero after some failure point. When we filter out more and more contrast patterns, the chance of us leaving out more and more relevant users is increased. Because of that, the recall keeps decreasing when σ and ρ is increased. F1 score demonstrates the balance between the precision and recall, and it should be noted that F1 score is maximized for low σ and ρ values. We recommend that $\sigma = 10$ and $\rho \in [1.1, 2]$ are parameters that yield generally good performance.

Baseline Comparisons. A comparison between our approach and the baselines in Section 5.2 is shown in Table 1. Multiple experiments were carried out for each dataset by changing their parameters. The model with the maximum F1 score is displayed in the table. The performance of our model is satisfactory for both small and large datasets compared to the other approaches. The *tweet frequency* baseline reveals that our results are not due to a general growth of frequency in tweets by each user. The missing cells are due to the heavy resource usage of the QT-LAMP-EP-* methods. The dataset $n_C = n_N = 200$ takes $\sim$80 GB of memory for those baseline methods. EPClose needs $\sim$300 MB of memory for contrast pattern mining for the case of $n_C = n_N = 400$. The memory usage and the results for that case demonstrate the scalability of our model.

Table 1. Results for detecting coordinated users using different methods. n_C - number of coordinating users in the dataset, n_N - number of normal users in the dataset.

Method	$n_C = n_N = 100$			$n_C = n_N = 200$			$n_C = n_N = 400$		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score	Precision	Recall	F1 Score
Tweet frequency	92.1%	35.0%	50.7%	98.0%	48.0%	64.4%	91.5%	43.3%	58.7%
Tweet language	64.0%	80.0%	71.1%	66.0%	81.0%	72.7%	66.0%	75.0%	70.2%
QT-LAMP-EP(BH)	59.8%	58.0%	58.9%	84.8%	64.0%	72.9%	-	-	-
QT-LAMP-EP(BY)	59.1%	55.0%	57.0%	85.1%	63.0%	72.4%	-	-	-
LCN+HCC	76.1%	63.0%	68.9%	77.3%	65.0%	70.6%	81.5%	70.4%	75.5%
AMDN-HAGE	50.0%	98.0%	66.2%	50.4%	100%	67.0%	50.6%	100%	67.2%
Our approach	77.1%	81.0%	**79.0%**	88.6%	82.0%	**85.2%**	86.2%	86.2%	**86.2%**

Ablation Study. The impact of the choice of attributes was tested using a greedy approach. We searched the impact of attributes using two methods.

1. *Subtractive.* We start from attribute set $A_0 = \mathcal{A}$. For each attribute a that is in A_i except `userid`, we remove that attribute, and we test multiple σ and ρ values to select the model with the maximum F1 score. Then we determine the attribute a that results in the model with the minimum of that maximum F1 score to be the attribute that has the greatest impact at the stage of A_i. Hence, A_{i+1} is created by removing that attribute from A_i. Define $F1(\sigma, \rho, A)$ to be the F1 score of the model with the attribute set A, minimum support σ, and growth rate threshold ρ. Then, formally,

$$A_{i+1} = A_i - \left\{ \operatorname*{arg\,min}_{a \in A_i - \{\texttt{userid}\}} \left\{ \max_{\sigma,\rho} F1\left(\sigma, \rho, A_i - \{a\}\right) \right\} \right\}$$

   We end the procedure when $A_n = \{\texttt{userid}\}$ where $n = |\mathcal{A}| - 1$.
2. *Additive.* We start from attribute set $A_0 = \{\texttt{userid}\}$. For each attribute a that is not in A_i, we test multiple σ and ρ values to select the model with the maximum F1 score. Then we determine the attribute a that results in that model with the maximum of that maximum F1 score to be the attribute that has the greatest impact at the stage of A_i. Hence, A_{i+1} is created by appending that attribute to A_i. Formally,

$$A_{i+1} = A_i \cup \left\{ \operatorname*{arg\,max}_{a \in \mathcal{A} - A_i} \left\{ \max_{\sigma,\rho} F1\left(\sigma, \rho, A_i \cup \{a\}\right) \right\} \right\}$$

   We end the procedure when $A_n = \mathcal{A}$ where $n = |\mathcal{A}| - 1$.

The purpose of the above methods is to search the locally optimal attribute sets and thus to assess the order of importance for each attribute. Figure 4 shows the variation of F1 scores and the number of contrast patterns when the highest impacting attribute is removed or introduced to the existing attribute set. The subtractive method reveals that `tweet_time`, `day_of_week`, `is_retweet` are the three highest impact attributes in that order. The additive method reveals that `day_of_week`, `tweet_time`, `tweet_client_name` are the three highest impact attributes in that order. Both methods identify that the first two attributes are most important compared to the rest of the attributes. Thus, temporal aspects of the events have played a major role in the behavioural patterns of this set of coordinating users. It is interesting to observe this indication of automation. Further, it is apparent that even for low numbers of attributes, similar to all attribute cases, low σ and ρ values yield the best possible F1 scores.

6 Conclusion

In this work, we proposed a novel approach to identify coordination by exploiting the growth of behavioural patterns of users in an online social network. We

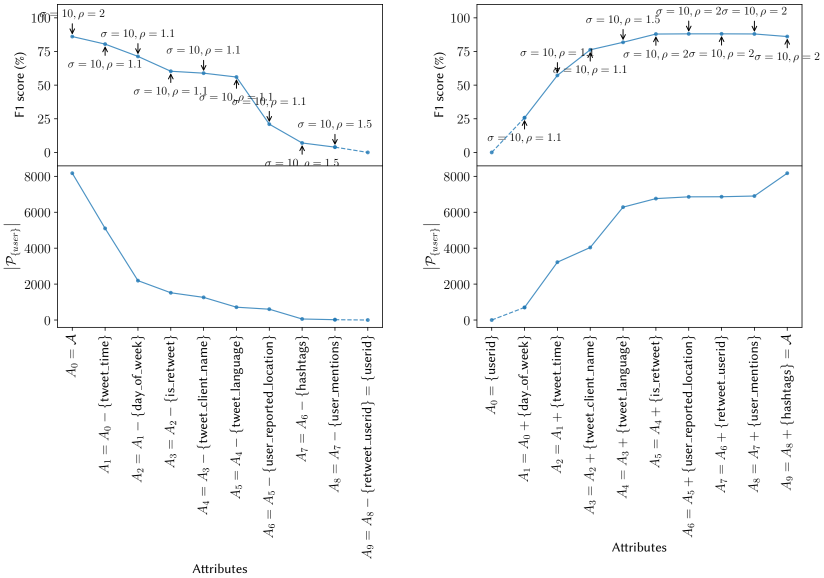

Fig. 4. The variation of F1 scores and the number of contrast patterns that are associated with users ($|\mathcal{P}_{\{user\}}|$) when the highest impacting attribute is removed (left) or added (right) to the attribute set. $n_C = n_N = 400$.

demonstrate that our proposed approach can achieve an increase of at least 10% in the F1 score compared to existing approaches. For future work, we aim to extend our model to utilize derived attributes such as: sentiments, emotions, and topics using the content of the posts. Investigating methods to automatically determine time intervals will also be important for real-time data. Further, we intend to work on a case study to use this approach on other social media datasets.

References

[1] AlipourChavary, E., Erfani, S.M., Leckie, C.: Improving Scalability of Contrast Pattern Mining for Network Traffic Using Closed Patterns (2020)

[2] Hellal, A., Ben Romdhane, L.: Minimal contrast frequent pattern mining for malware detection. Computers & Security (2016)

[3] Hristakieva, K., Cresci, S., Da San Martino, G., Conti, M., Nakov, P.: The Spread of Propaganda by Coordinated Communities on Social Media. In: ACM WEBSCI (2022)

[4] Komiyama, J., Ishihata, M., Arimura, H., Nishibayashi, T., Minato, S.i.: Statistical Emerging Pattern Mining with Multiple Testing Correction. In: Proceedings of the 23rd ACM SIGKDD (2017)

[5] Kong, J., Han, J., Ding, J., Xia, H., Han, X.: Analysis of students' learning and psychological features by contrast frequent patterns mining on academic performance. Neural Computing and Applications (2020)

[6] Lee, K., Caverlee, J., Cheng, Z., Sui, D.Z.: Content-driven detections of campaigns in social media. In: CIKM (2011)

[7] Lee, K., Caverlee, J., Kamath, K.Y., Cheng, Z.: Detecting collective attention spam. In: Proceedings of WICOW/AIRWeb - WebQuality '12 (2012)

[8] Loyola-González, O., Monroy, R., Rodríguez, J., López-Cuevas, A., Mata-Sánchez, J.I.: Contrast pattern-based classification for bot detection on twitter. IEEE access : practical innovations, open solutions **7** (2019)

[9] Magelinski, T., Ng, L.H.X., Carley, K.M.: A Synchronized Action Framework for Responsible Detection of Coordination on Social Media (2021)

[10] Mueller, R.S., Internet Research Agency, L., States, U.: UNITED STATES OF AMERICA V. INTERNET RESEARCH AGENCY LLC (2018)

[11] Nizzoli, L., Tardelli, S., Avvenuti, M., Cresci, S., Tesconi, M.: Coordinated Behavior on Social Media in 2019 UK General Election (2021)

[12] Pacheco, D., Hui, P.M., Torres-Lugo, C., Truong, B.T., Flammini, A., Menczer, F.: Uncovering Coordinated Networks on Social Media: Methods and Case Studies. Proceedings of AAAI ICWSM **15** (2021)

[13] Permanent Select Committee on Intelligence: Exposing Russia's Effort to Sow Discord Online: The Internet Research Agency and Advertisements. https://intelligence.house.gov/social-media-content/default.aspx (2018)

[14] Ryu, K.H., Lee, D.G., Piao, M.: Emerging pattern based prediction of heart diseases and powerline safety. In: Contrast Data Mining: Concepts, Algorithms, and Applications (2013)

[15] Sharma, K., Zhang, Y., Ferrara, E., Liu, Y.: Identifying Coordinated Accounts on Social Media through Hidden Influence and Group Behaviours. In: KDD (2021)

[16] Sherhod, R., Gillet, V.J., Hanser, T., Judson, P.N., Vessey, J.D.: Toxicological knowledge discovery by mining emerging patterns from toxicity data. Journal of Cheminformatics **5**(S1) (2013)

[17] Sherhod, R., Gillet, V.J., Judson, P.N., Vessey, J.D.: Automating Knowledge Discovery for Toxicity Prediction Using Jumping Emerging Pattern Mining. JCIM **52**(11) (2012)

[18] Sherhod, R., Judson, P.N., Hanser, T., Vessey, J.D., Webb, S.J., Gillet, V.J.: Emerging Pattern Mining To Aid Toxicological Knowledge Discovery. JCIM **54**(7) (2014)

[19] Thanasuan, K., Chaisangmongkon, W., Wongviriyawong, C.: Emerging patterns in student's learning attributes through text mining. In: EDM (2017)

[20] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS **30** (2017)

[21] Weber, D., Falzon, L.: Temporal Nuances of Coordination Network Semantics (Aug 2022)

[22] Weber, D., Neumann, F.: Amplifying influence through coordinated behaviour in social networks. SNAM **11**(1) (2021)

[23] Zhang, Y., Sharma, K., Liu, Y.: VigDet: Knowledge Informed Neural Temporal Point Process for Coordination Detection on Social Media (2021)