

# Stop Clickbait: Detecting and Preventing Clickbaits in Online News Media

## Study of Click-Bait Articles

A report submitted to the  
Indian Institute of Technology, Kharagpur  
in complete fulfillment of the requirements for the degree of

**Bachelor of Technology**  
in  
**Computer Science and Engineering**

by  
**Bhargavi Paranjape**

**12CS10056**

advised by  
**Prof. Niloy Ganguly**



**Department of Computer Science and Engineering**  
**Indian Institute of Technology, Kharagpur**

**November 2015**

# Certificate

This is to certify that the report titled **Stop Clickbait: Detecting and Preventing Clickbaits in Online News Media** submitted by **Bhargavi Paranjape (12CS10056)** to the Department of Computer Science and Engineering is a bonafide record of work carried out by her under my supervision and guidance. The report has fulfilled all the requirements as per the regulations of the Institute and, in my opinion, has reached the standard needed for submission.

**Prof. Niloy Ganguly**

Department of Computer Science and Engineering

Indian Institute of Technology, Kharagpur

May 2016

## Abstract

In the online world, most of the media outlets have to generate revenue from the clicks made by their readers, and due to the presence of numerous such outlets, they have to compete with each other for reader attention. To attract the readers to click on an article and visit the media site, they often come up with catchy headlines accompanying the article links, which lure the readers to click on the link. Such headlines are known as *Clickbaits*. While these baits may trick the readers into clicking, in the long-run, clickbaits usually don't live up to the expectation of the readers and leave them disappointed. In this work, we attempt to automatically detect clickbaits and then build a browser extension which warns the users about the possibility of being baited by such headlines in different websites. The extension also offers the users an option to block clickbaits she doesn't want to see and subsequently it automatically blocks similar clickbaits during future visits. We run extensive offline and online experiments across multiple web domains and find that the proposed clickbait detection and blocking approaches perform very well achieving 93% accuracy in detecting and 89% accuracy in blocking clickbaits.

## I. INTRODUCTION

As a large chunk of news consumption is moving online, the media landscape is undergoing a sea change. We can attribute this change to two primary dimensions. First, compared to the traditional offline media, where the readers' allegiances to a particular newspaper were almost static, online media offers the readers a gamut of options ranging from local, national or international media outlets to several niche blogs specializing on particular topics of interest. Second, most of the online media sites do not have any subscription charges and their revenue mostly come from the advertisements on their web pages.

Essentially, in the online world, every media outlet has to compete with many such outlets for reader attention and make their money from the *clicks* made by the readers. To attract the readers to visit the media site and click on an article, they employ various techniques. This often translates into the media outlets coming up with catchy headlines accompanying the article links, which lure the readers to click on the links. Such headlines are known as *Clickbaits*. According to the Oxford English Dictionary, clickbait is defined<sup>1</sup> as “(On the Internet) content whose main purpose is to attract attention and encourage visitors to click on a link to a particular web page.” Examples of such clickbaits include “*This Rugby Fan’s Super-Excited Reaction To Meeting Shane Williams Will Make You Grin Like A Fool*”, “*15 Things That Happen When Your Best Friend Is Obsessed With FIFA*”, “*Which Real Housewife Are You Based On Your*

<sup>1</sup>[oxforddictionaries.com/us/definition/american\\_english/clickbait](http://oxforddictionaries.com/us/definition/american_english/clickbait)

*Birth Month*” or the epic “*They Said She Had Cancer. What Happens Next Will Blow Your Mind*”.

Clickbaits exploit the cognitive phenomenon known as *Curiosity Gap* [1], where the headlines provide forward referencing cues to generate enough curiosity among the readers such that they become compelled to click on the link to fill the knowledge gap. While these baits may trick the readers into clicking, in the long-run, clickbaits usually don’t live up to the expectation of the readers, and leave them disappointed. Mark et. al. [2] argued that clickbait is an enabler of attention distraction. As the readers keep switching to new articles after being baited by the headlines, the attention residue from these constant switches result in cognitive overload, deterring the users from reading more informative and in-depth stories.

Even with all these hue and cry around clickbaits and several news articles [3] highlighting the ill effects of the prevalence of clickbaits, there has been little attempt to devise a systemic approach for a comprehensive solution to the problem. In 2014, Facebook declared that they are going to remove clickbait stories from users’ news feeds<sup>2</sup>, depending on the click-to-share ratio and the amount of time spent on these stories. Yet, Facebook users still complain that they continue to see clickbaits and there is a renewed effort now to clamp down on clickbaits<sup>3</sup>. In a recent research work, Potthast et al. [4] attempted to detect clickbaity tweets in Twitter. The problem with such standalone approaches is that clickbaits are prevalent not only on particular social media sites, but also on many other reputed websites across the web. For example, the section on sponsored stories at the end of the article pages of traditional media organizations like ‘The Guardian’ or ‘Washington Post’ contains many clickbaits. Hence, we need to have a comprehensive solution which can work across the web.

Separately, there has been a few ad-hoc approaches like ‘Downworthy’ [5] which detects clickbait headlines using a fixed set of common clickbait phrases and then converts the headline into something more garbage-ish, hoping that the users are not going to click on them; and ‘Clickbait Remover for Facebook’ [6] which prevents the links to a fixed set of domains from

<sup>2</sup>[newsroom.fb.com/news/2014/08/news-feed-fyi-click-baiting](https://newsroom.fb.com/news/2014/08/news-feed-fyi-click-baiting)

<sup>3</sup>[thenextweb.com/facebook/2016/04/21/facebook-might-finally-kill-clickbait-new-algorithm-tweaks/](https://thenextweb.com/facebook/2016/04/21/facebook-might-finally-kill-clickbait-new-algorithm-tweaks/)

appearing in the users' newsfeed. The problem with having a fixed rule set is they are not scalable and the classification performance is not good in an online setting. Similarly, preventing links to a fixed set of domains will also block links which are not clickbaits.

In this work, we take the first step towards building a comprehensive solution which can work across the web. We first build a classifier which automatically detects whether a headline is clickbait or not. Then we explore ways to block certain clickbaits from appearing in different websites. A survey conducted on 12 regular users of news media sites suggested that the headlines the users would like to block vary greatly across users, and they are influenced by the particular user's interests. Hence, instead of a generalized solution, we develop personalized classifier for individual users which can predict whether the user would like to block a particular clickbait given her earlier block and click history.

We finally build a browser extension, 'Stop Clickbait, which warns the users about the possibility of being baited by clickbait headlines in different websites. The extension also offers the users an option to block certain types of clickbaits she would not like to see during future encounters. We run extensive offline and online experiments across multiple web domains and find that the proposed clickbait detection and blocking approaches perform very well achieving 93% accuracy in detecting and 89% accuracy in blocking clickbaits. We believe that widespread use of such extensions would deter the users from getting lured by clickbait headlines, which in turn would disincentivize the media outlets from relying on clickbaits as a tool for attracting visitors to their sites.

## II. DATASET

We collected extensive data for both clickbait and non-clickbait categories.

**Non-clickbait:** We extracted the titles of news stories from a corpus of 18,513 Wikinews articles collected by NewsReader application [7]. In Wikinews, articles are produced by a community of contributors and each article needs to be verified by the community before publication. There are fixed style guides which specify the way some events need to be reported and presented to the readers. In order to be eligible for publication, each article needs to conform to these style guides. For example, to write the headline of a story, there are a set of guidelines<sup>4</sup> the author needs to follow. Due to these rigorous checks employed by Wikinews, we have considered the headlines of these articles as gold standard for non-clickbaits.

**Clickbait:** For clickbaits, we manually identified the following domains which publish clickbait articles: ‘BuzzFeed’, ‘Upworthy’, ‘ViralNova’, ‘Thatscoop’, ‘Scoopwhoop’ and ‘ViralStories’. We crawled 8,069 web articles from these domains during the month of September, 2015. To avoid false negatives (i.e. the articles in these domains which are not clickbaits), we recruited six volunteers and asked to label the headlines of these articles as either clickbait or non-clickbait. We divided the articles among the volunteers such that no volunteer needs to label more than 4,000 articles and each article is labeled by three volunteers. We obtained a ‘fair’ inter-annotator agreement with a Fleiss’ kappa score of 0.35. Taking the majority vote as ground truth, a total of 7,623 articles were marked as clickbaits. The notable examples of stories the volunteers marked non-clickbaits include the stories in the ‘news’ section on Buzzfeed, most of which are reported like traditional news bodies.

Finally, to have an equal representation of clickbait and non-clickbait articles while comparing them and building the classifier, we randomly selected 7,500 articles from both the categories.

### A. Analysis of clickbait on twitter

Random sample of tweets made available through Twitter API was collected for the month of June 2014. This data was found to contain an estimated 5 million tweets. A random tweet

<sup>4</sup>[en.wikinews.org/wiki/Wikinews:Style\\_guide#Headlines](http://en.wikinews.org/wiki/Wikinews:Style_guide#Headlines)

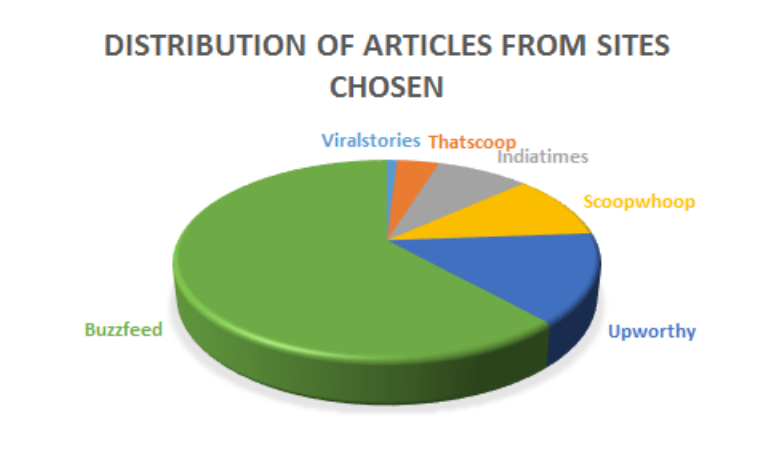


Fig. 1: Distribution of articles from clickbait sites since their time of creation

has the following entities: noitemsep

- Time of creation
- Tweet id
- Tweet text
- Source (platform, domain) of tweet
- User name
- Tweet Id in reply to
- Location
- Language of tweet
- Followers of the user
- Retweet count
- Urls
- Hashtags
- User-mentions

For the random sample of tweets from June 2014 (which consisted of around 5 million tweets), the urls tweeted by users were collected. In particular, we looked at the urls originating from popular clickbait websites and established online news media and the results obtained are shown in Figure 2. BuzzFeed generates as much content as the reputed news body, Wall Street Journal and half as much content as the most popular established news body on twitter, The New York

Times




Fig. 2: No. of Urls mentioned in tweets from news and clickbait domains



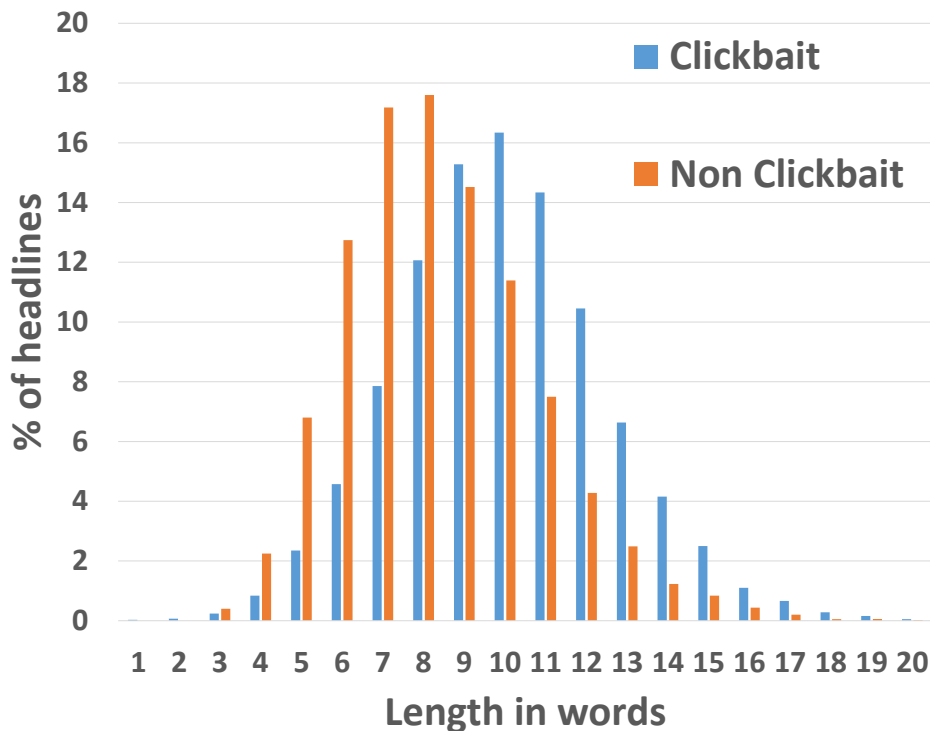


Fig. 3: Distribution of the length of both clickbait and non-clickbait headlines

### III. COMPARING CLICKBAITS AND NON-CLICKBAITS

We carried out a detailed linguistic analysis on the 15,000 headlines both in the clickbait and non-clickbait categories, using the Stanford CoreNLP tool [8]. A closer inspection of the clickbait headlines gives some insight about the semantic and syntactic nuances that occur more frequently in clickbait headlines compared to the traditional non-clickbait headlines.

News headlines occur in Title case on websites, and this was causing a poor performance of the POS tagger, sentiment analysis and parsing. So the words in the titles were selectively capitalised if they were Proper Nouns or Abbreviations. This was done using AlchemyAPI Named Entity Recognizer (NER) that recognises persons, locations, organizations etc. Entities were recognised and accordingly cased to give standard English sentences.

#### A. Sentence Structure

**Length of the headlines:** Figure 3 shows the distribution of the number of words in both clickbait and non-clickbait headlines. It indicates that conventional non-clickbait headlines are

shorter than clickbait headlines. For example, the average length of the clickbait headlines is 10, whereas the average length is 7 for non-clickbait headlines.

Traditional news headlines typically contain mostly *content words* referring to specific persons and locations, while the *function words* are left out for readers to interpret from context. As an example, consider the news headline: “*Visa deal or no migrant deal, Turkey warns EU*”. Here most of the words are content words summarizing the main takeaway from the story, and it has very few connecting function words in between the content words.

On the other hand, clickbait headlines are longer, well-formed English sentences that include both content and function words. One example of such headlines is “*A 22-Year-Old Whose Husband And Baby Were Killed By A Drunk Driver Has Posted A Gut-Wrenching Facebook Plea*”. It contains a generous mix of content and function words.

**Length of the words:** Even though the number of words are more in clickbait headlines, the average word length is shorter. Specifically in our dataset, the average word length of clickbait headlines is found to be 4.5 characters, while the non-clickbait headlines have an average word length of 6.

The reason for shorter word lengths in clickbaits is primarily due to the frequent use of shorter function words and word shortenings. Shortened forms of words like *they’re*, *you’re*, *you’ll*, *we’d* which are informally used in the English language are prevalent in clickbait headlines. On the other hand, they are not commonly found in non-clickbait headlines. As we can see in Figure 4(a), only 0.6% of the traditional news headlines contain word shortenings, whereas nearly 22% of clickbait headlines have such shortened words.

**Length of the syntactic dependencies:** We used the Stanford collapsed-coprocessor dependency parser [9] to identify the syntactic dependencies between all pairs of words in the headlines, and then computed the distance between the governing and the dependent words in terms of the number of words separating them.

Figure 8 shows the distribution of the maximum distance between governing and dependent words in both clickbait and non-clickbait headlines. On average, clickbaits have longer dependencies than non-clickbaits; the main reason being the existence of more complex phrasal




Fig. 4: Percentage of clickbait and non-clickbait headlines, which include (a) Word Contractions, (c) Hyperbolic Words, and (d) Determiners. (b) Percentage of words identified as Stop words in both clickbait and non-click headlines.

sentences as compared to non-clickbait headlines. Consider the example “*A 22-Year-Old Whose Husband And Baby Were Killed By A Drunk Driver Has Posted A Gut-Wrenching Facebook Plea*”, where the subject ‘22-Year-Old’ and the verb ‘Posted’ are separated by an adjective clause, making the length of the syntactic dependency as high as 11.

#### B. Stop Words, Hyperbolic Terms and Common Phrases

**Stop words:** Stop words are defined as the most common words that occur in any corpus of a particular language. Figure 4(b) shows the percentage of words present in both categories of headlines, which are stop words in English. It can be seen that, in clickbait headlines, stop words are used more frequently (e.g. 45% compared to 18% in non-clickbaits) to complete the structure of the headlines. On the other hand, in conventional news reporting, more *content* words are used and inference of stop words is left to the reader. Due to this anomalous proportion of stop words in clickbait headlines and their contribution to sentence semantics, in the subsequent




Fig. 5: Distribution of longest syntactic dependencies between all pair of words in clickbait and non-clickbait headlines.




Fig. 6: One of the longest dependencies found in clickbait sentences where the words posted and 22-year-old are separated by 11 words

n-gram analysis, stop words were retained. This is a diversion from typical n-gram analysis, where stop words are removed before any analysis is performed.

**Hyperbolic terms:** To compare the sentiment values of the constituent words in both clickbait and non-clickbait headlines, we performed sentiment analysis using the Stanford Sentiment Analysis tool [10]. We found that a substantial fraction of clickbait headlines consist of words having ‘Very Positive’ sentiments (e.g. *Awe-inspiring*, *breathtakingly*, *gut-wrenching*, *soul-*

Clickbait	I, dog, everyone, girl, girls, guy, guys, he, here, it, kid, kids, man, men, mom, one, parent, people, photos, reasons, she, signs, something, that, they, thing, things, this, thoughts, times, video, ways, we, what, who, woman, you, images, boys, dads
Non-clickbait	Apple, Australia, bomb, bomber, Bush, Canada, China, court, crash, earthquake, Explosion, fire, France, Google, government, group, House, India, Iran, Israel, Korea, leader, man, minister, NASA, Obama, party, people, PM, police, President, rocket, Russia, Scientists, Senate, Thousands, toll, U.S., UN, US, Woman

TABLE I: 40 most commonly occurring subject words in both clickbait and non-clickbait categories

*stirring, etc.*), which are almost non-existent in non-clickbait headlines. We call these extremely positive words as *hyperbolic terms*. Figure 4(c) shows the percentage of headlines in both categories which include hyperbolic terms. Use of such eye-catching terms in clickbaits strongly urge the reader to pursue the article with a promise of sensational information.

**Internet slangs:** Another class of words commonly found in clickbait headlines are Internet slang words like *WOW, LOL, LMAO, AMA, AF, etc.* Along with hyperbolic terms, use of the slang words also immediately catches the attention of the reader and lure them to read the article.

**Punctuation patterns:** Clickbait headlines also make use of informal punctuation patterns such as *!?, ..., \*\*\*, !!!* – which are not used in conventional non-clickbait headlines.

**Common bait phrases:** Further, several commonly used catch phrases in clickbait headlines exploit the “curiosity gap” of users, such as “*Will Blow Your Mind*”, “*You Won’t Believe*”. We manually compiled a list of most commonly used bait phrases in the clickbait corpus. We further extended this list with the phrases used by Downworthy [5] to detect clickbaits.

### C. Subjects, Determiners and Possessives

**Sentence subjects:** The subject of an English sentence is the person or thing about whom the statement is made. To identify the subject words in the headlines, we used the Stanford syntactic dependency parser [9], and then looked for the dependency relation ‘*nsubj*’ among all the dependency relations found by the parser. For example, the 40 most commonly occurring subject words found in both clickbait and non-clickbait headlines are listed in Table I.

One interesting pattern we observed in clickbaits is the repetition of the popular subjects across many headlines. Nearly 62% of the clickbait headlines contained one of the 40 most common clickbait subject words listed in Table I. On the other hand, only 16% of non-clickbait headlines contained the top 40 non-clickbait subject words.

**Determiners:** Clickbait headlines often employ determiners like *their, my, these, which, etc* to reference particular people or things in the article. Figure 4(d) shows the percentage of headlines in both clickbait and non-clickbait headlines, where determiners are present. It can be seen that the use of determiners is way more in clickbaits compared to non-clickbaits. The use of such determiners is primarily to make the user inquisitive about the object being referenced and persuade them to pursue the article further.

**Possessive case:** Clickbait headlines often address the reader in the first and second person with the use of subject words *I, We, You*. Even third person references are common nouns like *he, she, it, they, man, dog* rather than specific proper nouns. This is in stark contrast with non-clickbait categories, where the reporting is always done in third person.

#### D. Word N-grams, POS tags and Syntactic N-grams

**Word N-grams:** Word N-gram is defined as a contiguous sequence of n words from a given text sample, where N can vary from 1, 2, ..., to the length of the sample. From the clickbait and non-clickbait headlines, we extracted all possible 1, 2, 3, and 4-grams. For example, some commonly occurring 3 and 4-grams in clickbait headlines are *'how well do you', 'can we guess your', 'what happens when', 'how many of these'*; while in non-clickbait headlines, some commonly occurring 3 and 4-grams are *'dies at age', 'kills at least', 'us supreme court', 'found guilty of', 'won the match'*. We have found that the proportion of headlines that contained the top 0.05% of unique n-grams in clickbaits was 65%, whereas in non-clickbaits, the proportion was 19%. Clickbait headlines follow a pattern of phrases repeatedly, while news headlines are inherently factual and unique in their reporting.

**Part of Speech tags:** We tagged the headlines in both categories with the Part of Speech (POS) tags of their constituent words using the 45 Penn Treebank POS tags [11]. We specifically used




Fig. 7: Distribution of Parts-of-Speech (POS) tags for words in both clickbait and non-clickbait headlines

the MaxEnt POS Tagger [12]. As most of the article headlines use title casing, we found some words to be erroneously tagged by the direct application of the POS tagger.

To overcome this limitation, we added a preprocessing step, where we identified named entities using Stanford Named Entity Recognizer tool [13] and retained those words in title case. We converted every other word to lowercase to avoid ambiguity for the POS tagger. After the preprocessing step, the POS tagger identified the POS tags for all words in the headlines. Figure 7 shows the distribution in POS tagging.

From Figure 7, we make the following interesting observations:

(i) Conventional non-clickbait headlines contain much larger proportion of proper nouns (POS tag: NN), indicating more content words and entities, than in clickbaits.

(ii) Clickbait headlines contain more adverbs and determiners (POS tags: RB, DT, WDT) than non-clickbait headlines.

(iii) Clickbaits also have higher proportion of personal and possessive pronouns (POS tags: PRP and PRP\$) like *her*, *his*, *its*, *you* compared to non-clickbaits.




Fig. 8: One of the longest dependencies found in clickbait sentences where the words posted and 22-year-old are separated by 11 words

(iv) Clickbaits and non-clickbaits use verbs in different ways. Overall number of verbs is more in clickbaits as they focus on forming well-formed sentences. Regarding the person and the tense of the verbs, non-clickbait headlines tend to use more verbs in past participle and 3rd person singular form (POS tags: VBN and VBZ), whereas clickbaits use mostly past tense and non-3rd person singular forms (POS tags: VBD and VBP).

**Syntactic N-grams:** Syntactic N-grams (SN-grams) are formed by traversing paths of length  $N$  along the syntactic tree obtained from the collapsed-coprocessor dependency parsing [9] of the headlines. We do a depth-first traversal of the syntactic tree, using every node as a source. For example, in the sentence “A 22-Year-Old Whose Husband And Baby Were Killed By A Drunk Driver Has Posted A Gut-Wrenching Facebook Plea”, we extract the SN-gram “nsubj, acl:relcl, nmod:agent” corresponding to “22-Year-Old Posted Driver Killed”. This way, SN-grams combine syntactic dependencies among non-neighbor words in the headline. The advantages of using SN-grams are two-fold: (i) SN-grams are fewer in number than word n-grams, and (ii) SN-grams help capture linguistic phenomena by linking words that may not be neighbors in the surface structure but are syntactically related.

For instance, in the headlines “Which Disney Song Are You Based On Your Zodiac Sign” and “Which Badass Witch Are You Based On Your Birth Month”, the syntactic bigram (*dobj*, *det*) captures the pattern “**Which** \*\*\*\* **based**”. An estimate of the number of such SN-grams is the height of the syntactic parse tree, which on an average was found to be 10.03 for clickbait and 6.45 for non-clickbaits in our dataset.



#### IV. CLASSIFYING HEADLINES AS CLICKBAITS

The comparative analysis, described in the earlier section, indicates prominent linguistic and structural differences between clickbait and non-clickbait headlines. We attempt to use these differences as features to classify article headlines into clickbait and non-clickbait categories.

##### A. Feature selection

**Sentence structure:** To capture the structural properties of the headline as classifier features, we used the length of the headline, the average length of words, the ratio of the number of stop words to the number of content words and the longest separation between syntactically dependent words of a headline.

**Word patterns:** Word level structural features that we included were the presence of cardinal numbers in the beginning of the headline, presence of unusual punctuation patterns and the number of contracted word forms employed in the headline.

**Clickbait language:** Features that capture the nuances of the language employed, especially in the clickbait headlines, include the presence of hyperbolic terms, common clickbait phrases, internet slangs and determiners. To model the popularity of the subject word in clickbait headlines as a feature, we used the score of a multinomial Naive Bayes classifier over sentence subjects. The score represents the ratio of the probability of assigning clickbait class label to the probability of assigning non-clickbait class label, given the subject word of the sentence. Model parameters for the Naive Bayes classifier were estimated using both datasets.

**N-gram features:** Word n-grams, POS n-grams and syntactic n-grams were used as features. N-gram feature space grows linearly with the size of the dataset. In order to limit the number of n-gram features used based on their frequency of occurrence, we pruned the feature space efficiently by using the sub-sequence Property and an APRIORI-like algorithm [14]. Similar to the case with subject words, we built three multinomial Naive Bayes classifiers for the three sets of pruned features: n-grams, POS n-grams and syntactic n-grams. The scores of these 3 auxiliary classifiers were input as features to the main classifier.

##### B. Classification performance

We used a set of 14 features spanning the structural, word-level, n-gram and linguistic categories, as described above. We experimented with prediction models such as Support Vector

Model	SVM					Decision Tree					Random Forest				
Measure	Acc.	Prec.	Rec.	F1	roc-auc	Acc.	Prec.	Rec.	F1	roc-auc	Acc.	Prec.	Rec.	F1	roc-auc
Sentence Structure	0.77	0.79	0.75	0.77	0.84	0.74	0.77	0.69	0.73	0.8	0.75	0.76	0.73	0.74	0.82
Word/Token Patterns	0.84	0.84	0.83	0.84	0.91	0.84	0.84	0.84	0.84	0.91	0.84	0.84	0.84	0.84	0.91
Clickbait language	0.86	0.88	0.82	0.85	0.88	0.86	0.88	0.82	0.85	0.90	0.86	0.88	0.82	0.85	0.90
N-grams	0.89	0.92	0.85	0.89	0.9	0.89	0.92	0.85	0.89	0.91	0.89	0.92	0.85	0.89	0.91
Overall	0.93	0.95	0.90	0.93	0.97	0.90	0.91	0.89	0.90	0.90	0.92	0.94	0.91	0.92	0.97

TABLE II: Performance of the clickbait classifier using different prediction models.

Machines (SVM) with RBF kernel, Decision Trees and Random Forests. Table 2 shows the 10-fold cross validation performance (specifically accuracy, precision, recall, F1-score and ROC AUC scores) for all three prediction models. It shows the evaluation scores if each category of features were used separately and then for the combined feature set. We can see from Table 2 that SVM performed best with an accuracy of 93%, a precision of 0.95 and recall 0.9.

Finally, as a baseline for comparison, we took the fixed set of rules employed by Downworthy [5] to detect clickbaits, and ran it on our dataset. The 10-fold cross validation performance achieved by Downworthy is 76% accuracy with 0.72 recall and 0.78 precision. Therefore, the proposed classification technique outperforms the baseline with a large margin.

Further note that, even though the current classifier works with titles in English only, the characteristics of clickbait used as features are common linguistic phenomena occurring across languages and the classifier can easily be extended to other languages.

## V. BLOCKING CLICKBAIT HEADLINES

In the earlier section, we saw that the classifier achieves 93% accuracy in detecting clickbaits. With this impressive performance of the classifier, the next logical step is to devise an approach which can block certain clickbaits according to the user's discretion.

Towards that end, we attempted to understand whether there is some uniformity in choice about the type of headlines users would like to block. We conducted a user survey by asking a group of 12 users to review 200 randomly chosen headlines from the clickbait category. The users were asked to mark the headlines they would have clicked while surfing the web, and also mark the headlines they would have liked to block. We then computed the average Jaccard coefficient between all pairs of users, for the headlines to be clicked as well as the headlines to be blocked. Though all 200 headlines show the characteristic features of clickbaits, the average Jaccard coefficients for clicked as well as blocked headlines are 0.13 and 0.15 respectively. These low Jaccard scores indicate that different users interact with different clickbaits very differently, and there is not much commonality in user behaviors towards clickbait headlines. Effectively, users' interpretation of clickbait is subject to their own interests, likes as well as dislikes.

**Blocking as personalized classification:** As users' interaction with clickbait headlines vary drastically, a one-size-fits-all approach can not work in blocking clickbaits. We instead need personalized classifier for individual users to classify the clickbaits into the ones to block and the ones not to block. Essentially the problem is to model the user's interests from the articles she has already read as well as already blocked. Accordingly, for a new clickbait, we need to predict whether the user would like to block the article or not.

The notion of user interests in clickbait articles or lack thereof can be modeled with the help of two interpretations. For instance, if some user decides to block the clickbait headline *Can You Guess The Hogwarts House Of These Harry Potter Characters?*, we can make two conclusions. (i) The user may not be interested in the topic 'Harry Potter' itself, and does not want to read any article related to 'Harry Potter'; or (ii) she may be annoyed by the commonly occurring pattern "*Can you guess ...*" but may click on another 'Harry Potter' related article in

the future. There can also be cases, where both reasons play a role. Hence, we modeled and designed methods to capture both notions of user interests as well as a combination of both factors.

**Blocking based on topical similarity:** Our first approach to block clickbait articles is to first extract a set of topics from a clickbait article, and find the similarity between this set and the topics previously extracted from blocked and clicked articles.

One widely used topic modeling approach is Latent Dirichlet allocation (LDA) [15]. However, LDA performs poorly if the available data is sparse and noisy. Most articles with clickbait headlines have very little textual content, and in many cases they are not relevant to the topic being addressed in the headline. For example, in our dataset, the average length of textual content in Clickbait articles is found to be 300 words compared to 810 words in articles with non-clickbait headlines. Furthermore, Only 70% of the clickbait articles had the content words of the headline present in the article body, contrasting the 96% in non-clickbait articles. Due to these reasons, topic modeling approaches like LDA could not be used with clickbait articles.

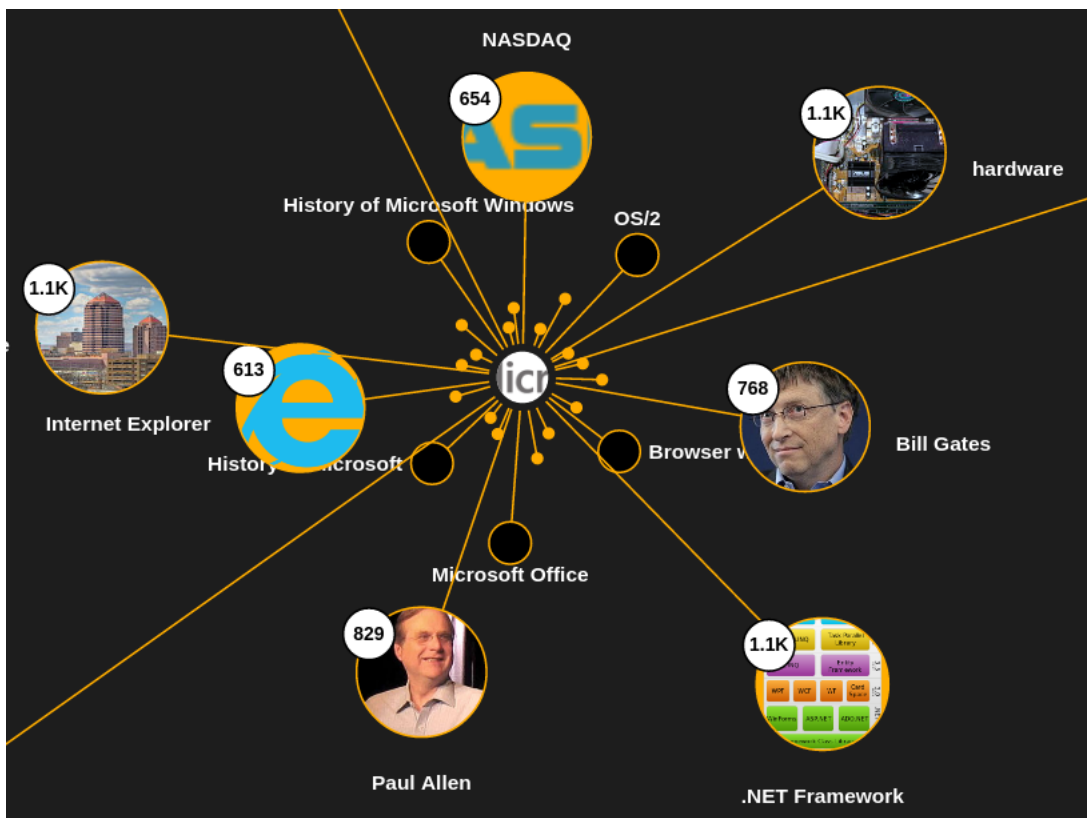
To find topics of interest to the user, we chose to use the content words in the article headline, as well as the article metatags and keywords that occur in the <head> tag in the article’s html source. For instance, in the html source for the article having headline *We Tried The New Starbucks Butterbeer Drink And Dumbledore Would Definitely Approve*, the tags found were *butterbeer, harry potter, hermione, jk rowling, wizarding world*. These tags are given by the developer of the corresponding webpages and they contain topical information regarding the article. They identify the hidden topics of an article that otherwise topic-modeling systems like LDA [15]<sup>5</sup> will take large training data to identify. Tags and keywords extracted from a particular clickbait link are stored as attributes *ClickTags* or *BlockTags*, depending on whether the link has been clicked or blocked.

For a given link, we used BabelNet [16]<sup>6</sup> to expand its *BlockTags* or *ClickTags* sets. We discover the nodes in BabelNet that correspond to these tags. These nodes initially form a self-

<sup>5</sup>LDA here performs poorly as the available data are sparse and noisy

<sup>6</sup>BabelNet is a multilingual semantic network which connects 14 million concepts and named entities extracted from WordNet and Wikipedia. Each node in the network is called a BabelSynset.

Fig. 9: Nodes in BabelNet that are neighbours to Microsoft



contained cluster, called a *Nugget*. We expand the nugget further by adding common hypernyms of member nodes. Two Nuggets are merged when a BabelSynset (i.e. a node) is common to both.

Considering all the links a user has blocked and clicked respectively, we form a user's *BlockNuggets* and *ClickNuggets*. Then for a new clickbait link, the block/do not block decision is predicted based on whether the nugget for the new link is more *similar* to *BlockNuggets* or *ClickNuggets*. Here similarity is computed based on the number of nodes common in two nuggets.

For every user, we use the top 100 blocked and clicked links ordered by timestamps for tag extraction and nugget formation. This is done to limit the data considered for training to the latest user interests, which can change with time. We use 1000 blocked and clicked links ordered by timestamps for the keyword extraction and training. This is done to limit the data considered for training to the latest user interests, which can change with time.

**Blocking based on linguistic patterns:** In the second approach, we identified the linguistic

```

1: procedure TRAIN(BlockNuggets)
2:
3:   BlockTags  $\leftarrow$  Keywords in a blocked article
4:   NewNugget, Merge  $\leftarrow$  Nodes in BlockTags
   and thier Common Hypernymnodes
5:   for  $N \in$  BlockNuggets do
6:     if  $N \cap$  NewNugget  $\neq \emptyset$  then
7:       Merge  $\leftarrow$   $N, Merge$ 
8:     end if
9:   end for
10:  BlockNuggets  $\leftarrow$  BlockNuggets + Merge
11:                                      $\triangleright$  Similarly create ClickNuggets from the set of clicked articles
12: end procedure

```

```

1: procedure DECIDE(BlockNuggets, ClickNuggets, Q)
2:   QueryTags  $\leftarrow$  KeyWords in Test Article Q
3:   BlockVote  $\leftarrow$  0
4:   for  $N \in$  BlockNuggets do
5:     if  $QueryTags \cap N \neq \emptyset$  then
6:       sim  $\leftarrow$   $|QueryTags \cap N|$ 
7:       BlockVote  $\leftarrow$  BlockVote +  $e^{sim}$ 
8:     end if
9:   end for
10:                                      $\triangleright$  Similarly get ClickVote from ClickNuggets
11:
12:   if BlockVote > ClickVote then
13:     return block
14:   else
15:     return donotblock
16:   end if
17: end procedure

```

Fig. 10: Blocking based on topical similarity

```

1: procedure DECIDE(BlockPatterns, ClickPatterns, Q)
2:   BlockVote  $\leftarrow$  0
3:   ClickVote  $\leftarrow$  0
4:   for  $P \in$  BlockPatterns do
5:      $sim(P, Q) = 1 - \frac{edit\ distance(P, Q)}{max(|P|, |Q|)}$ 
6:     BlockVote  $\leftarrow$  BlockVote +  $e^{sim(P, Q)}$ 
7:   end for
8:
   Similarly get ClickVote
9:   if BlockVote > ClickVote then
10:     return block
11:   else
12:     return donotblock
13:   end if
14: end procedure

```

Fig. 11: Blocking based on linguistic patterns

patterns in the articles that the user clicks on or chooses to block. The pattern is formed by normalizing the words in the headlines in the following ways. (i) Numbers and Quotes are replaced by tags <D> and <QUOTE>. (ii) The top 200 most commonly occurring words in the clickbait corpus, including English stop words, were retained in their original form. (iii) Content words such as Nouns, Adjectives, Adverbs and Verb inflections were replaced by their POS tags. For instance, “Which Dead ‘Grey’s Anatomy’ Character Are You”, reduces to “which JJ <QUOTE> character are you” and “Which ‘Inside Amy Schumer’ Character Are You” reduces to “which <QUOTE> character are you”.

We convert each headline into such patterns, and thus we get a set of patterns for both blocked articles and clicked articles. To compute the similarity between two patterns, we use the word-level edit distance. Using the mechanism similar to the topical similarity case, we make the block/do not block decision.

The list of blocked and clicked patterns used to compute score are the top 100 recorded patterns in the order of timestamp values to take dynamic user interests into consideration.

Approach	Accuracy	Precision	Recall
Pattern Based	0.81	0.834	0.76
Topic Based	0.75	0.769	0.72
Hybrid	0.72	0.766	0.682

TABLE III: Performance of different blocking approaches

**Hybrid approach:** We also experimented with an hybrid approach which takes into account both topical similarity and linguistic patterns. For a test article, its tags are extracted, nugget is formed and compared with the *BlockNuggets* and *ClickNuggets* – this gives the topical similarity scores. Similarly, we get the linguistic similarity scores. The hybrid scores are obtained using a weighted combination of both topical and linguistic similarity scores, and finally we make the block/do not block decision based on the hybrid scores.

**Evaluation:** We tested all three approaches using the click and block decisions marked by the 12 users. Table III shows the average accuracy, precision and recall scores for all approaches. Since the pattern based approach yields better results and executes faster than the more involved topic based approach, it is better suited for a real-time environment. Hence, it was integrated into the browser extension ‘Stop Clickbait’ introduced in the next section.



## VI. BROWSER EXTENSION: STOP CLICKBAIT

In the earlier sections, we showed that both the classifier and the blocking approach achieve high accuracy in detecting and blocking clickbaits. Hence, to increase the applicability of the proposed approaches, and to help the users to deal with clickbaits across different websites, we attempt to build an extension ‘Stop Clickbait’ for Chrome browsers. ‘Stop Clickbait’ warns the users about the existence of clickbaits in different webpages. It provides the users the facility to block certain clickbaits whereby it automatically block similar clickbaits in future visits. We next describe the working of the chrome extension.

When a webpage is loaded in Chrome, ‘Stop Clickbait’ scans the Document Object Model (DOM) for anchor elements (`<a href=... >`), and it also keeps listening for dynamic insertion of anchor elements in the DOM. Once an anchor tag is discovered, it proceeds to check if the anchor tag has any anchor text. If the anchor text is available, it is sent to a server, where the text is input to the clickbait classifier. If the anchor text is not available in the DOM, then the url is sent to the server, where it makes a GET request for the webpage title and runs the classifier on the obtained title.

Then, based on the anchor text or the webpage title, the classifier predicts whether the link is clickbait or not. The result of the classification is fed back to the extension, and links that are flagged as clickbaits is marked with a green button (by adding an element into the DOM) and links that are not clickbaits is left unmarked. Figure 14(a) shows the green clickbait indicators in one of the webpages. Finally, when the user clicks on the green indicator button (as shown in Figure 14(b)), two options appear for the user: (i) Block similar content in future, and (ii) the clickbait indicator was placed erroneously, i.e. the classification was wrong. Further, in case, the extension fails to mark a genuine clickbaits in a webpage, the user can click on the link (as shown in Figure 14(c)), and mark it as clickbait. We store these feedbacks in the server and the clickbait classifier is retrained every day with the added feedback data.

**Evaluating the performance of the extension:** We uploaded the extension to the official Google chrome store<sup>7</sup> and also circulated the extension binary in our peer groups. We provided

<sup>7</sup>The extension is available at [chrome.google.com/webstore/detail/stop-clickbait/iffolpdcmebhghbamkgobjjdeejinma](https://chrome.google.com/webstore/detail/stop-clickbait/iffolpdcmebhghbamkgobjjdeejinma)

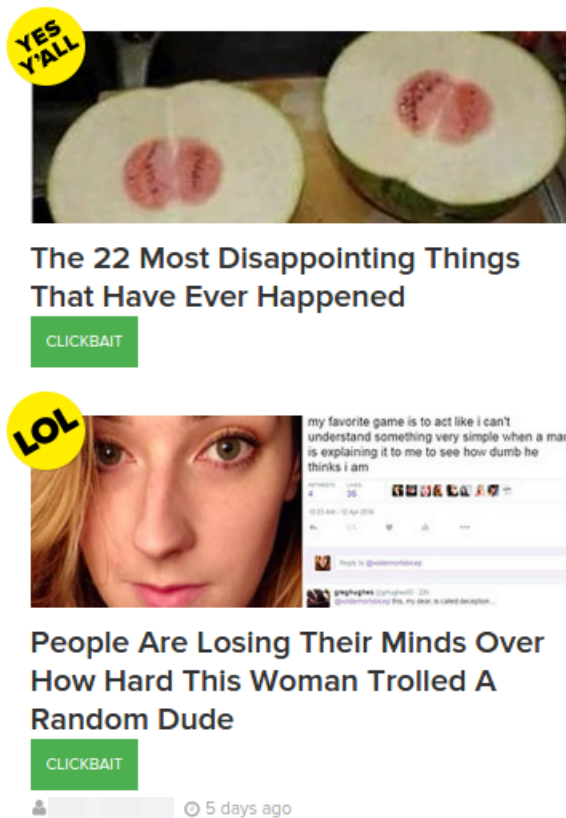


Fig. 12: Green clickbait indicators in webpages

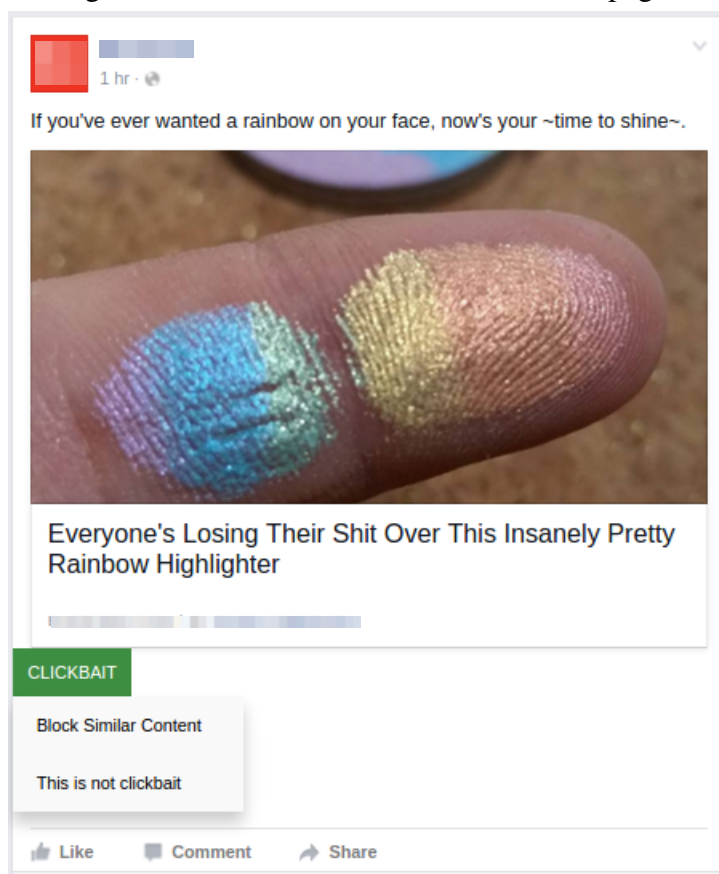


Fig. 13: Option to block a link or report misclassification



Fig. 14: Option to report link which should be marked as clickbait

a randomly generated unique 32 byte identification number for every instance of the extension to facilitate training for each personalized classifier.

Overall 37 people used the extension during the month of April, 2016. These users carried out their regular web activities and visited different websites like Facebook, Twitter, New York Times etc. and they reported both false positives and false negatives. We present in Table IV, the overall performance of the extension across different websites and also individually at 5 different domains. We can see that the classification performance is very good yielding 94% accuracy and F1-score of 0.93 across all websites the users visited.

Out of these 37 users, 16 users explicitly blocked more than 10 clickbaits while browsing different websites. We invoked the personalized classifier as described in the earlier section and blocked different clickbaits during their further visits. We also provided an option to the users

Domain	Precision	Recall	Accuracy
Huffington Post	0.88	0.93	0.92
CNN	0.93	0.99	0.97
Buzzfeed	1.00	0.97	0.98
NY Times	0.83	0.95	0.95
Facebook	0.85	1.00	0.93
Overall	0.92	0.95	0.94

TABLE IV: Performance of the extension at different domains

to check the links which have been blocked by the extension and give feedback on whether the blocking was a right decision or not. According to the user feedbacks, on average, the extension had correctly blocked 89% of the links.

The extension ‘Stop Clickbait’ works across all websites. However, most webpages have large number of extreme hyperlinks which are not links to media article. Also, the aesthetics of every webpage was not under the author’s control and insertion of the element, though non-invasive to the website content, may interfere with the page style elements. However, hyperlinks on Facebook and Twitter are filtered and limited to links appearing in the respective newsfeeds. The extension runs smoothly with both these domains. This bias is because users are exposed to clickbait primarily through the sharing of such articles on social networking sites and will need the extension to work well on these sites. The extension smoothly with Huffingtonpost, Buzzfeed, New York Times and CNN as well.

This was the motivation behind extending the existing clickbait classifier to a personalised classifier that learns based on user behavior patterns as they interact with the extension. The extension listens for clicks on links that the classifier has flagged as clickbait. It also allows the user to specifically mark a clickbait link to be blocked and records blocked links for every user. We provide a randomly generated unique 32 byte identification number for every download of the extension from a single Google account to record this information on the server database.

## VII. RELATED WORK

The origin of clickbaits can be traced back to the advent of tabloid journalism, which started focusing on ‘soft news’ compared to ‘hard news’, and sensationalization rather than reporting in depth and truthful account of the events. There has been many research works in media studies highlighting the problems with tabloidization. For example, Rowe [17] examined how the common tabloid properties like simplification and spectacularization of news, are making its way into the more conventional newspapers and how it is changing the course of professional journalism. Similarly, the current concerns on the prevalence of clickbaits [3] highlight the changing face of journalistic gatekeeping during the abundance of clickbait articles having very low news value.

There has been recent works to understand the psychological appeal of clickbaits. Blom et. al. [18] examined how clickbaits employ two forms of forward referencing – discourse deixis and cataphora – to lure the readers to click on the article links. Chen et. al. [19] argued for labeling clickbaits as misleading content or false news.

However, there has been little attempt to detect and prevent clickbaits. As mentioned earlier, Facebook attempted to remove clickbaits depending on the click-to-share ratio and the amount of time spent on different stories. A recent work by Potthast et al. [4] attempted to detect clickbaity Tweets in Twitter by using common words occurring in clickbaits, and by extracting some other tweet specific features. The browser plugin ‘Downworthy’ [5] detects clickbait headlines using a fixed set of common clickbaity phrases, and then converts them to meaningless garbage. The problems with the above approaches are that they either work on a single domain, or the fixed ruleset does not capture the nuances employed across different websites. In this work, we propose and demonstrate a comprehensive solution which works very well across the web.

## VIII. CONCLUSION

In this paper, we compared clickbait and non-clickbait headlines, and highlighted many interesting differences between these two categories. We then utilized these differences as features to detect clickbaits. We also proposed personalized approaches which can block certain clickbaits according to user interests. Finally, using these two components, we have developed a Chrome extension which warns the users of different media websites about the presence of

clickbaits in these websites. The extension also gives the option to the users to block clickbaits and it automatically blocks similar ones in subsequent visits by the users.

To the best of our knowledge, our work is the first attempt to provide a comprehensive solution to deter the prevalence of clickbaits. However, the job is far from over. Our future work lies in improving the classification and blocking performances further and tune the extension according to further user feedbacks. Finally, it is our belief that combating the prevalence of clickbaits should be a community initiative and towards that end, we plan to make all the data and source code public, so that the researcher and developer communities can come forward, and collectively make the effort a grand success.

## REFERENCES

- [1] G. Loewenstein, "The psychology of curiosity: A review and reinterpretation." *Psychological bulletin*, vol. 116, 1994.
- [2] G. Mark, "Click bait is a distracting affront to our focus," [nytimes.com/roomfordebate/2014/11/24/you-wont-believe-what-these-people-say-about-click-bait/click-bait-is-a-distracting-affront-to-our-focus](http://nytimes.com/roomfordebate/2014/11/24/you-wont-believe-what-these-people-say-about-click-bait/click-bait-is-a-distracting-affront-to-our-focus).
- [3] J. Dvorkin, "Column: Why click-bait will be the death of journalism," [pbs.org/newshour/making-sense/what-you-dont-know-about-click-bait-journalism-could-kill-you/](http://pbs.org/newshour/making-sense/what-you-dont-know-about-click-bait-journalism-could-kill-you/).
- [4] M. Potthast, S. Köpsel, B. Stein, and M. Hagen, "Clickbait detection," in *Advances in Information Retrieval*. Springer, 2016.
- [5] A. Gianotto, "Downworthy: A browser plugin to turn hyperbolic viral headlines into what they really mean," [downworthy.snipe.net/](http://downworthy.snipe.net/).
- [6] W. Markus, "Clickbait remover for facebook," [chrome.google.com/webstore/detail/clickbait-remover-for-fac/hkbhmlgcpmneffdammemapgiiniagj](http://chrome.google.com/webstore/detail/clickbait-remover-for-fac/hkbhmlgcpmneffdammemapgiiniagj).
- [7] V. Piek *et al.*, "Newsreader: How semantic web helps natural language processing helps semantic web," *Special Issue Knowledge Based Systems, Elsevier*.
- [8] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, and D. McClosky, "The stanford corenlp natural language processing toolkit." in *ACL (System Demonstrations)*, 2014, pp. 55–60.
- [9] M.-C. De Marneffe, B. MacCartney, C. D. Manning *et al.*, "Generating typed dependency parses from phrase structure parses," in *Proceedings of LREC*, vol. 6, 2006.
- [10] R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts, "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proceedings of EMNLP*, vol. 1631, 2013.
- [11] B. Santorini, "Part-of-speech tagging guidelines for the penn treebank project (3rd revision)," 1990.
- [12] A. Ratnaparkhi *et al.*, "A maximum entropy model for part-of-speech tagging," in *Proceedings of EMNLP*, vol. 1, 1996.
- [13] J. R. Finkel, T. Grenager, and C. Manning, "Incorporating non-local information into information extraction systems by gibbs sampling," in *Proceedings of ACL*, 2005, pp. 363–370.
- [14] J. Fürnkranz, "A study using n-gram features for text categorization," *Austrian Research Institute for Artificial Intelligence*, vol. 3, 1998.
- [15] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.
- [16] R. Navigli and S. P. Ponzetto, "Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network," *Artificial Intelligence*, vol. 193, pp. 217–250, 2012.
- [17] D. Rowe, "Obituary for the newspaper? tracking the tabloid," *Journalism*, 2011.
- [18] J. N. Blom and K. R. Hansen, "Click bait: Forward-reference as lure in online news headlines," *Journal of Pragmatics*, vol. 76, 2015.
- [19] Y. Chen, N. J. Conroy, and V. L. Rubin, "Misleading online content: Recognizing clickbait as? false news?" in *Proceedings of ACM Multimodal Deception Detection*, 2015.