

“Dummy Grandpa, Do You Know Anything?”: Identifying and Characterizing Ad Hominem Fallacy Usage in the Wild

Utkarsh Patel, Animesh Mukherjee, Mainack Mondal

Indian Institute of Technology, Kharagpur
utkarshpatel@iitkgp.ac.in, animeshm@cse.iitkgp.ac.in, mainack@cse.iitkgp.ac.in

Abstract

Today, participating in discussions on online forums is extremely commonplace and these discussions have started rendering a strong influence on the overall opinion of online users. Naturally, twisting the flow of the argument can have a strong impact on the minds of naïve users, which in the long run might have socio-political ramifications, for example, winning an election or spreading targeted misinformation. Thus, these platforms are potentially highly vulnerable to malicious players who might act individually or as a cohort to breed fallacious arguments with a motive to sway public opinion. *Ad hominem* arguments are one of the most effective forms of such fallacies. Although a simple fallacy, it is effective enough to sway public debates in offline world and can be used as a precursor to shutting down the voice of opposition by slander.

In this work, we take a first step in shedding light on the usage of ad hominem fallacies in the wild. First, we build a powerful ad hominem detector based on transformer architecture with high accuracy (F1 more than 83%, showing a significant improvement over prior work), even for datasets for which annotated instances constitute a very small fraction. We then used our detector on 265k arguments collected from the online debate forum – *CreateDebate*. Our crowd-sourced surveys validate our in-the-wild predictions on CreateDebate data (94% match with manual annotation). Our analysis revealed that a surprising 31.23% of CreateDebate content contains ad hominem fallacy, and a cohort of highly active users post significantly more ad hominem to suppress opposing views. Then, our temporal analysis revealed that ad hominem argument usage increased significantly since the 2016 US Presidential election, not only for topics like Politics, but also for Science and Law. We conclude by discussing important implications of our work to detect and defend against ad hominem fallacies.

Introduction

Today online forums and social media sites facilitate easy collaborative opinion formation for billions of users surpassing geographical boundaries. However, perhaps quite naturally, this process of opinion formation also involves participating in online arguments where multiple parties often present their conflicting views. The caveat here is

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

that the arguments presented in online debates are not always sound. They often contain *deceptive arguments in disguise* (Kennedy 1993). Intuitively, in online forums, the users present informal fallacies, necessitating deep content analysis to identify them (as opposed to the formal ones, which can be examined using logical representations) (Sahai, Balalau, and Horincar 2021).

Amongst different fallacies, *ad hominem* is perhaps the most famous one in the offline world (Macagno 2013; Schiappa and Nordin 2013; Zalta 2004; Woods 2007). *Ad hominem* or *against the person* is a fallacious argument, based on feelings of bias (mostly irrelevant to the argumentation), rather than reality, reason, and rationale. However, despite a long history of dissecting and condemning ad hominem fallacies in the offline world, even online users are no stranger to the usage of ad hominem fallacies (Goodman 2020; Redinger 2020). *Ad hominem* arguments are often personal attacks on someone’s character or motive rather than an attempt to address the reasoning that they presented. People tend to use ad hominem arguments because they want to appeal to others’ emotions rather than reasoning.

Recently, there has been substantial research concerning investigating and countering hate speech, misinformation as well as cyberbullying within the user-generated content posted on social media (Mondal, Silva, and Benevenuto 2017; Mondal et al. 2018; Das et al. 2021; Mathew et al. 2020a,b). In the same vein, although relatively scarce, some very recent works are exploring the detection of ad hominem fallacies in the wild using computational methods (Habernal et al. 2018; Sahai, Balalau, and Horincar 2021). However, these works focused more on the detection of ad hominem (and other) fallacies using automated methods in online forums. There is not much work shedding light on the lay of the land for ad hominem usage over time. We aim to bridge this gap. Specifically, we ask the following research questions:

1. Can we design a practical ad-hominem detector which can uncover ad-hominems in the wild with high accuracy?
2. How does the dynamics of ad hominem argumentation evolve with time in the wild? Who are users that play the key role in posting these ad hominem arguments?

To address these questions, in this work, we present a

data-driven exploration of ad hominem arguments in the wild using CreateDebate¹, an online discussion forum, as an experimental testbed. We used an in-house high-accuracy and high external validity ad hominem detector on a dataset containing more than 18k posts with 265k comments generated by 15k users of CreateDebate. Next, we analyzed the detected large-scale ad hominem arguments to shed light on in-the-wild ad hominem usage. Specifically, we have made three key contributions to this work.

First, to answer the first research question, we developed ad hominem detectors considering two scenarios—when the annotated data is abundant and when it is not, using an annotated dataset from previous work on Reddit (Habernal et al. 2018). Our models significantly improved over the ad hominem detectors reported in prior work and achieved a macro-F1 score more than 0.83. Furthermore, we evaluated the predictions of our detector on the CreateDebate dataset using a user study. Our user study demonstrated that results from our detector (trained on Reddit data) are also externally valid—achieving a 94% accuracy on the CreateDebate data. Overall, unlike previous work (Habernal et al. 2018), leveraging modern transformer models enabled us to build a high-accuracy detector and effectively handle situations with low amount of annotated data.

Subsequently, to answer the second research question, we leveraged this detector on our large-scale CreateDebate data and found that a significant 31.23% (i.e., almost one-third) of all the comments on CreateDebate are ad hominem. On further investigation, we uncovered that a community of highly active users posts a disproportionately high volume of ad hominem fallacies (more than 50% of their total comments).

Third, to dig deeper into our second research question, we checked whether ad hominem arguments were always used in such a high volume, or was it just a recent trend? It appeared that the fraction of ad hominem arguments showed a sharp rise after 2016. This trend was prominent not only in the Politics subforum, but transcended in subforums like Science and Law. We found a striking correlation—the userbase of the Politics subforum of CreateDebate has a high overlap with userbase for subforums like Science and Law. The rise of ad hominem arguments in Politics subforums seems to be triggered by the 2016 US Presidential election, which resulted in the users active in Political subforum posting insulting comments in other forums as well, hence, significantly increasing the ad hominem usage. We also release our model to the research community.²

Related Works

We divide prior works into two important sub-parts—exploration on ad hominem argumentation, especially in online forums like ChangeMyView³ and CreateDebate, and usage of Generative Adversarial Networks (GANs) in NLP.

¹<https://www.createdebate.com/>

²<https://bit.ly/3TJbMPR>

³<https://www.reddit.com/r/changemyview/>

Investigation on Ad Hominem Argumentation

Aristotle first identified that some arguments are indeed *deceptions in disguise* (Kennedy 1993). He called them fallacies. The ad hominem arguments are addressed in most of the follow-up treatises of fallacies (Hamblin 1970; Eemeren and Grootendorst 1987; Boudry, Paglieri, and Pigliucci 2015). Ad hominem arguments are used in a debate for simply attacking the opponents’ traits instead of countering their arguments (Tindale 2007). Naturally, ad hominem arguments are based on feelings of bias rather than reason, often involving personal attacks on someone’s character or motive. Though arguing *against the person* is considered faulty, these arguments are used in online debate forums and social media sites (Habernal et al. 2018; Sahai, Balalau, and Horincar 2021). Ad hominem arguments are multifaceted and use complex strategies, involving not a simple argument, but a cumulation of several combined tactics (Macagno 2013). Interestingly, the issue in some type of ad hominem arguments is not logical inconsistency, but rather what one would call a pragmatic inconsistency. It refers to a kind of inconsistency between asserted statements and personal actions (Walton 1998). However, this type of inconsistency is out of scope for this work. Specifically, the majority of the above research was often aimed at dissecting what constitutes ad hominem in philosophy, rather than evaluating its usage in the real world (Schiappa and Nordin 2013; Macagno 2013; Zalta 2004; Woods 2007).

Recently, the scenario started to change when researchers working on NLP aimed to identify different fallacies in online forums. To that end, Wulczyn, Thain, and Dixon (2017) annotated 38k instances of Wikipedia talk page comments for detecting personal attacks on the forum and Jain et al. (2014) studied principal roles in discussions from the Wikipedia Article for Deletion pages, and extracted several typical roles like ‘idiots’, ‘voices’, ‘rebels’, etc. which might be considered signals for ad hominem fallacies.

Studies on ChangeMyView: More recent work exploited online discussion forums like Reddit, primarily the ChangeMyView subreddit, to detect naturally occurring hate-speech and ad hominem fallacies. Wei, Liu, and Li (2016) studied the impact of different sets of features on the identification of persuasive comments. Tan et al. (2016) developed a framework for analyzing persuasive arguments and malleable opinions. Habernal et al. (2018) investigated ad hominem argumentation at three levels of discourse complexity (arguments in isolation, in direct replies to original post without dialogical context and in a larger inter-personal discourse context). Sahai, Balalau, and Horincar (2021) extended this work and found the types of fallacies. Our work builds on this type of detection methods, yet extends them considerably.

Studies on CreateDebate: Abbott et al. (2016) developed Internet Argument Corpus (v2.0), a collection of corpora for research in political debate on Internet forums, which contains a sample from CreateDebate (3k posts) and includes topic annotations. Wei et al. (2016) analyzed the disputation action in the online debate by labeling a set of disputing argument pairs extracted from CreateDebate and performing annotation studies. Trabelsi and Zaïane (2014) sug-

gested a fine-grained probabilistic framework for improving the quality of opinion mining from online contention texts. Hasan and Ng (2014) exploited stance information for reason classification, proposing systems of varying complexity for modeling stances and reasons. Qiu (2015) modeled user posting behaviors and user opinions for viewpoint discovery and proposed an integrated model that jointly considers arguments, stances, and attributes. Qiu et al. (2015) predicted user stances on a variety of topics and assembled user arguments, interactions, and attributes into a collaborative filtering framework that exploits recently introduced fast inference methods.

Text Classification Using Deep Learning for Our Detector

We explored a number of state-of-the-art transformer architectures like Google’s BERT (Devlin et al. 2019), OpenAI’s GPT-1 (Radford et al. 2018) for our experiments. These architectures are trained over large-scale annotated datasets and only require fine-tuning for a targeted task to achieve high accuracy for various NLP applications. However, one major disadvantage of these architectures is that even fine-tuning them often requires at least thousands of annotated examples for the targeted tasks. However, in our use case, obtaining thousands of annotated ad hominem arguments (and an equal number of non-ad hominem arguments) is costly and might be difficult to obtain. To that end, we leveraged a recent technique—integrating these huge pre-trained models with Generative Adversarial Networks (GANs) (Goodfellow et al. 2014). In GANs, a ‘generator’ is trained to produce samples resembling some data distribution. This training process ‘adversarially’ depends on a ‘discriminator’, which instead is trained to distinguish examples of the generator from the real instances. SS-GANs are an extension to GANs where the discriminator also assigns class labels to each example while discriminating whether or not it was naturally produced (Salimans et al. 2016). Croce, Castellucci, and Basili (2020) proposed ‘GAN-BERT’ that extends the fine-tuning of BERT-like architectures with unlabeled data in a generative adversarial setting using SS-GAN schema. Building on this related work helped us create an accurate yet explainable detector with limited data. Next, we will start with describing our approach to develop the classifier.

Present work: Our work builds a highly accurate detector, beating the models used by Habernal et al. (2018) in macro averaged F1-score for classification and establishes the external validity of the detector on CreateDebate data. We created a novel, significantly improved ad hominem detector and demonstrated its utility in detecting ad hominem in the wild (with very limited annotated sample, which considerably increased the utility of the model). Then using this detector, we measured the prevalence of ad hominem in the wild—we found that the amount of ad hominem in recent times increased manyfold, more than the figures hinted in any of the previous works. Our highly accurate detector relied on the recent advances in text classification using transformer models as discussed above. Our analysis revealed *for the first time*, the extremely worrying prevalence of ad

hominem in the wild—one-third of the posts in the CreateDebate forum were ad hominems and a small cohort of highly active users hurl the largest number of ad hominems. Quite interestingly, hurling ad hominems accelerated at time periods closer to the 2016 US Presidential election. Overall, political debates are found to be at the core of increased ad hominem usage with its effects transcending to other topics like religion, science and law.

Ad Hominem Detection

For our experiments, we used ad hominem argumentation in the ChangeMyView (CMV) dataset (Habernal et al. 2018) as benchmark. ChangeMyView is a popular subreddit in which a user (called OP, original poster) posts an opinion and other users write comments to change the perspective of OP about the posted opinion. OP can acknowledge convincing arguments by giving *delta* points.

Unlike regular debate forums, strict rules are enforced on this subreddit content. Violating these rules results in deletion of the content by moderators. The CMV dataset contains 7242 comments from this subreddit (3622 instances with the label ‘ad hominem’ and 3620 instances with the label ‘none’). The dataset was created maintaining a balance of syntactic and semantic similarity between the instances of the two-class labels. We refer the reader to (Habernal et al. 2018) for additional details about the dataset. We will use this dataset to build powerful classifiers for ad hominem detection.

Habernal et al. (2018) used CNN and 2 Stacked Bi-LSTM models for detecting ad hominem comments. They reported 10-fold cross validation results. We used the BERT model (case-insensitive, base) and carried out the same 10-fold cross validation experiments. The results are presented in Table 1.

Model	Accuracy	Macro-F1
CNN	0.808	0.807
2 Stacked Bi-LSTM	0.781	0.781
BERT	0.834	0.834

Table 1: 10-fold cross validation results on the CMV dataset for classification of ad hominem comments using CNN (baseline), 2 Stacked BiLSTM (baseline) and BERT.

As noted in Table 1, we achieved a 2.6% improvement in accuracy over the baselines. We further investigate the segments that could be potentially responsible for flagging an argument as ad hominem. We cast an explanation of triggers and dynamics of ad hominem argumentation as a supervised learning problem and draw theoretical insights by a retrospective interpretation of the learned model. As [CLS] token is the aggregate representation of the input sequence for classification tasks we use attention scores⁴ for [CLS] token

⁴In a transformer based architecture, attention scores refer to the overall strength of the relationship of a word with the other words in the sequence. Attention scores over input regions or intermediate features are often interpreted as a contribution of the attended unit to the inference made by the model (Rigotti et al. 2022).




Figure 1: Example visualization of attention scores for the [CLS] token. Words like *poor*, *fucking*, etc. get more attention, hence, have more influence on the classification.

to identify key tokens influencing the classification. A sample visualization of the attention scores is shown in Figure 1. We greedily select top three tokens (excluding [CLS] and [SEP] tokens) on the basis of attention scores so that the trigrams centered at those tokens do not overlap and highlight these trigrams in the visualization. A sample visualization of these trigrams is shown in Figure 3. We analyzed these highlighted trigrams for the comments which were flagged as ad hominem and observed that the BERT model is not only able to beat the baselines in terms of accuracy, but also the highlighted trigrams in this case can be easily interpreted to extract linguistic insights into the potential triggers for ad hominem argumentation.

One bottleneck for such studies is the cost of generating an annotated dataset. Annotating ad hominem arguments is a very costly task as they are difficult to comprehensively and objectively define (Sheng et al. 2020). Hence, we simulate a situation where we assume that labeled instances constitute a very small fraction of the dataset. While doing training on different folds, we intentionally drop the class labels of a given fraction of instances and then fine-tune the BERT model on only the labeled instances and evaluate its performance. This is repeated for different fractions of labeled instances while training. We observe that as we increase the fraction of unlabeled instances, the macro-F1 score for classification degrades. This is a major disadvantage of transformer architectures like BERT—they require thousands of annotated examples to achieve state-of-the-art results on the targeted tasks. However, with few annotated examples, they fall apart.

We experimented with GAN-BERT to leverage the unlabeled instances in the training set in a generative adversarial setting, which we simply cannot use while working with BERT. As the results in Figure 2 show, the requirement of labeled instances can be drastically reduced (up to only 50 – 100 labeled instances) for obtaining similar performance as that of BERT. When low fraction of instances in training set are annotated, GAN-BERT beats BERT. One

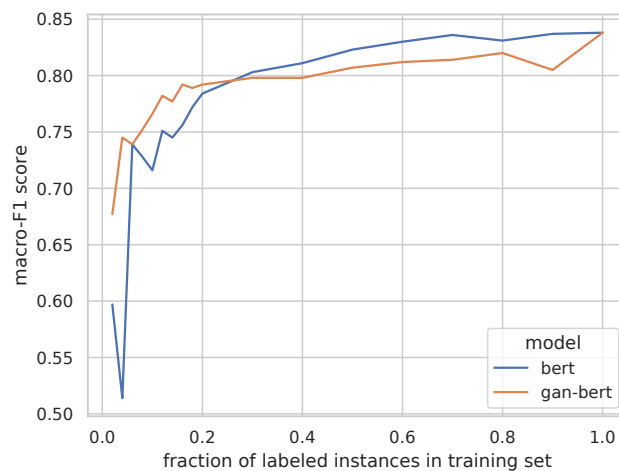


Figure 2: BERT vs. GAN-BERT – Macro-F1 scores for ad hominem classification for different fraction of labeled instances in the training set. When fraction of labeled instances is low, GAN-BERT beats BERT. However, for high fraction, BERT outperforms GAN-BERT.

of the possible reasons could be that as the number of annotated instances decreases, GAN-BERT leverages its GAN architecture to use the unlabeled instances to train its discriminator. As this utility is not present in BERT, its performance is not very good. Naturally, when a larger fraction of annotated instances are available in the training set, BERT beats GAN-BERT since the utility of the GAN architecture is diminished. Figure 3 compares the visualizations obtained by using BERT (when 100% training instances are labeled) and GAN-BERT (when 15% training instances are labeled). We observe that even with much less annotated data, GAN-BERT is able to detect triggers similar to that of BERT. In the rest of the paper we will report results using the BERT-based model, the GAN-BERT models produced similar results.

Detecting Ad Hominem Fallacies in the Wild

Now that we created an accurate and explainable ad hominem detector using the Reddit data, we aimed to test the usage of ad hominem fallacies in discussion forums from the wild. To that end, we chose CreateDebate as our experimental testbed.

Collecting CreateDebate Dataset

CreateDebate is a social networking debate website, launched in 2008. It was built around ideas, discussion, and democracy to help groups of people to sort through issues, viewpoints, and opinions. The discussions in CreateDebate often aim towards reaching consensus and understanding to make better decisions. CreateDebate is very similar to Reddit with a notable exception—its moderation policy is different, only the debate creator can be the debate moderator⁵.

⁵<https://www.createdebate.com/about/faq>

1.3.18 Comment 18 - Ad hominem

super stupid has just said that the mule team has explained to it that women have a 63 percent dis approval of trump super stupid did the mule team tell you that or was it abc clear up your confused world for us

1.3.18 Comment 18 - Ad hominem

super stupid has just said that the mule team has explained to it that women have a 63 percent dis approval of trump super stupid did the mule team tell you that or was it abc clear up your confused world for us

Figure 3: Example visualization of trigrams, which have the largest influence on classification, based on attention scores of [CLS] token for an ad hominem comment using BERT (cyan) and GAN-BERT (orange). Even with very less annotated data, GAN-BERT is able to detect trigrams similar to that of BERT.

Topic	Examples
Politics	STUPOR STUPID you think a lot in the CONFUSED MIND you have ! And by your IGNORANT RESPONSE you further expose you are a consumer of a media that lies to you and you LAP it up like the GOOD LITTLE LEFTIST you are !!!!!
Religion	Tell me all about your love for the Muslims you Party Parrot ? Stupid i thought you were a Jew ? Do you have confusion each day until ya turn on DNC Media ?
World News	You threatening to kill anyone who disagrees with Corbyn is fascism, yes. Christ you're such a liar it's embarrassing. Do you actually believe the words you type you boring fucking fascist troll?
Science	You haven't made a single argument to disprove the existence of a designer .Not even one That is because I am not disputing that there may have been a designer, you retarded imbecile. Do you even read the stuff you reply to?
Law	Ah , the American patriot who never heard of implied consent laws, why not do a bit of reasearch you dummy ? I love it when an idiot like you starts back tracking as you claimed cops could not pull you over now they can. A cop can pull you over for various reasons , you've just done a huge about turn on your lazy assertions as usual
Technology	Shut up you bitter old fraud. You're a twisted, delusional, selfish piece of excrement with precisely zero personal integrity. You are a walking, talking form of ass cancer.

Table 2: Examples of ad hominem arguments across different topical sub-forums of CreateDebate.

Each CreateDebate post is created by a user who also acts as its sole moderator. The forum allows users to write their perspectives as comments on the posts. Other users can support a comment, dispute it or clarify it as replies. The site, like Reddit, doesn't limit the depth of comment nesting. CreateDebate forum is divided into 14 topical forums—Politics, Entertainment, World News, Religion, Law, Science, Technology, Sports, Comedy, Business, Travel, Shopping, Health, and NSFW. The majority of the content in CreateDebate is public for all the forums. We found CreateDebate to be suitable for our investigation since it is a popular discussion-based forum in the wild with weak moderation. In effect, CreateDebate provides us an opportunity to measure the prominence of ad hominem fallacy usage over time.

In this work, we programmatically collected the complete publicly available CreateDebate dataset for all topical CreateDebate forums from the inception of the CreateDebate service. However, for brevity, we will present primary results from the top six CreateDebate forums—Politics, Religion, World News, Science, Law, and Technology. Results from all other topical forums remained the same. We present the general statistics of the dataset in Table 3—in totality, these six forums contain 18, 227 posts with 264, 814

comments made by 14, 961 users uploaded over 14 years. We verified that all posts we collected from CreateDebate were in English. We leveraged this large-scale discussion dataset (posted over the years) collected from CreateDebate to detect ad hominem from online discussions. However, we faced a crucial question—is our detector, trained over the Reddit CMV dataset, extendable to the CreateDebate dataset? We explored this question next.

Validating Our Ad Hominem Detector on CreateDebate Discussions

After collecting the CreateDebate data, we faced a dilemma—our ad hominem detector was fine-tuned on the Reddit CMV dataset (as noted in the previous section), however, CreateDebate is a very different forum with potentially different userbase and linguistic styles (including syntactic and semantic differences with Reddit). Thus, in this section, we will report a real-world data-driven survey that establishes the validity of our detector on the CreateDebate dataset.

Study setup: Our goal was to test the accuracy of our model output on the CreateDebate dataset. To that end, we ran our BERT-based detector on this dataset and randomly

Topic	# Posts	# Comments	# Users
Politics	10,434	119,850	7,686
Religion	2,841	77,418	4,563
World News	2,008	27,418	3,622
Science	1,276	20,691	2,837
Law	759	11,016	1,436
Technology	909	8,421	2,674
Total	18,227	264,814	14,961

Table 3: Basic statistics of our collected CreateDebate dataset. The first post in our dataset was posted on February 20, 2008 and the last post was updated on November 24, 2021.

Topic	% AH comment	% AH users
Politics	37.03 ± 6.36	34.95
Religion	27.47 ± 6.37	37.52
World News	25.21 ± 5.30	32.88
Science	27.90 ± 6.10	35.11
Law	25.41 ± 5.07	33.57
Technology	18.94 ± 5.49	25.28
Overall	31.23 ± 6.12	33.19

Table 4: Basic statistics of ad hominem (AH) content in our collected CreateDebate dataset. Error bands for our estimates for ad hominem content are also included. Users who have 50% or more comments as ad hominem arguments are referred to as ad hominem users here. The third column indicates the percentage of such ad hominem users in the different topical forums.

sampled 50 comments which were classified as ad hominem by our detector and another 50 comments which were non-ad hominem. Next, we created a simple online survey. We used Prolific⁶, a crowdsourcing platform, to recruit participants for our survey. We recruited 18+ years old US nationals who were fluent in English, had a 100% approval rating on the platform and had participated in at least 200 previous studies. In this survey, we presented a set of comments (from our sample of 100 random CreateDebate comments) along with a link to access the original CreateDebate discussion and its replies to each participant. Then we asked the participants to mark each of those comments as ad hominem or non ad hominem. For each comment, we additionally presented (in case the participant deem it to be an ad hominem) top three phrases identified by our model (with highest weights) and enquired if these key phrases indeed make this content ad hominem (the participants could also add their own phrases in a free form text field). This part of our survey was aimed to validate the explainability of our model. In total, 15 participants gave three responses for each of the 100 comments (each participant gave responses for a batch of 20 comments); comparing the annotators across 5 batches yielded substantial inter-annotator agreement (0.73 Fliess’ κ). The average time per participant was 8 minutes and we compensated them with \$1. The survey instrument is provided in Appendix A of our arXiv submission⁷.

⁶<https://www.prolific.co/>

⁷<https://arxiv.org/abs/2209.02062>

Results: We found that for 94% of CreateDebate comments, the labels given by participants were the same as the predicted label by our model, signifying the high validity of our model output even on the CreateDebate dataset. Furthermore, for 94.3% of ad hominem comments, the key phrases identified by our model (with the highest attention scores) exactly matched with the participant-identified phrases. This shows the power of the generalizability of our model.

Characterizing Ad Hominem Fallacy Usage in CreateDebate Discussions

We used our (almost) accurate and explainable detector on the CreateDebate dataset and characterized the ad hominem fallacies. We will start by exploring the volume of ad hominem fallacies in the wild.

Usage of Ad Hominem Fallacies in CreateDebate

We simply run our BERT-based detector on CreateDebate data to find the answer to the question—do CreateDebate users leverage ad hominem fallacy? We present the answer in Table 4. Surprisingly, the percentage of ad hominem comments in the CreateDebate forum is alarmingly high, especially for CreateDebate topical forums related to Politics (37.03%). In fact, a large number of users are using these ad hominem fallacies—34.95% for Politics, demonstrating, ad hominem fallacies are used rampantly in the wild. These numbers contrast with the Reddit CMV forum where Habernal et al. (2018) found only 0.02% posts to be ad hominem. Even for a regular online discussion, only 19.5% of comments under online news articles were found to be incivil (Coe, Kenski, and Rains 2014), much lower than the reported fraction of ad hominems. We show some examples of topical ad hominems posted on CreateDebate in Table 2.

Username	# TLC	Username	# DR
UserA	1,445	UserC	5,371
UserB	1,388	UserB	4,039
UserC	1,078	UserA	2,694
UserD	1,077	UserK	2,513
UserE	1,030	UserD	2,365
UserF	845	UserL	2,345
UserG	770	UserG	1,577
UserH	603	UserM	1,297
UserI	521	UserI	1,200
UserJ	506	UserN	1,069

Table 5: Users who posted most of the top-level comments (denoted by column TLC) and who received most of the direct replies (denoted by column DR) on the Politics subforum of CreateDebate. We omitted usernames for privacy concerns and common users from both the sides are represented in bold font.

Now, we ask an obvious question—why is this percentage alarmingly high in contrast to CMV’s 0.02%, even though there is a mechanism to *report* a comment on CreateDebate? To investigate, we focus on ad hominem posts from ‘Politics’ subforum (owing to its 37.03% ad hominem content). As the statistics in Table 5 show, the Politics sub-forum of




Figure 4: Variation in (a) number of authors in $S(\lambda, \rho)$; reciprocity amongst the users computed using (b) support network and (c) dispute network for different λ and ρ for CreateDebate Politics sub-forum.




Figure 5: Variation in percentage of ad hominem posts posted per month on CreateDebate across different topics (*left*) and percentage of users for whom 50% or more posts are ad hominem per month (*right*) (averaged over 1 year).

CreateDebate follows a heavy tail distribution, where the 10 most active users posted about 26% of the top-level comments, it also shows the users who received the most direct replies to their posts on the subforum. Interestingly, both sides have 6 usernames in common, signifying the possible influence of only a handful of key players. Thus, we ask—do these active users have any role in elevating the fraction of ad hominem arguments in the CreateDebate forum?

Correlation between Activity and Ad Hominem

We wanted to investigate if a handful of users are colluding to upload disproportionately more ad hominem content. To that end, we define $X(\lambda)$ as the set of authors who posted at least λ top-level comments, and $Y(\rho)$ as the set of authors who received at least ρ direct replies to their comments. We define the set $S(\lambda, \rho)$ as $X(\lambda) \cap Y(\rho)$ and attempt to understand the activity of users with different levels of λ and ρ .

Highly active users act as a community while posting:

We start by creating directed graphs showing support and dispute between the authors in $S(\lambda, \rho)$. The weights of the edge from node A to node B in support and dispute networks represents how many times author A agreed/disagreed with author B via direct reply. As the results in Figure 4 show, the number of authors in the set decreases if either λ or ρ

is increased, yet the reciprocity⁸ in support and dispute networks increases with an increase in λ and ρ . Our finding implies that the influential actors, who also happened to write most posts and receive most replies (high λ and ρ), participate in the debates not as an individual but as small-sized communities/cohorts of highly active users.

Community of highly active users post a disproportionately high volume of ad hominem posts: We started with the hypothesis that the alarmingly high ad hominem content present in the CreateDebate forum is mostly generated by small communities of highly active users. These users always post together and reply to each other. To check whether our hypothesis is true, we grouped the users participating in political debates on the CreateDebate forum based on their total comment count. Next, we test the fraction of ad hominem comments in the content generated by each group. The result is shown in Table 6. The set of authors who wrote more than 2000 comments constituted only 0.1% of the users, yet they post around 26.5% of total content with as high as 55.2% of their content flagged as ad hominem. This

⁸“Reciprocity” is defined to be a measure of how likely a pair of vertices are mutually connected (Garlaschelli and Loffredo 2004) (i.e., both $node_1 \rightarrow node_2$ and $node_2 \rightarrow node_1$ edges exist) in a directed network. In our setup, this indicates, that if an author a_1 supports (disputes) author a_2 then what is the likelihood that a_2 would support (dispute) a_1 .

group’s activity is in stark contrast with the users who post less than 10 comments—they posted only 11.8% comments and a meager 16.8% of those comments were ad hominem. This finding confirms our hypothesis and identifies an intriguing pattern of ad hominem posting—these illogical personal attacks are often used by a highly active community or cohort of users, presumably to shut down voices of less active opponents.

Summary: CreateDebate dataset contains a surprisingly high volume of ad hominem—31.23%, which is much higher than the figures reported in earlier works on ad hominem content in other forums. Using the Politics sub-forum, we further showed that these ad hominem comments are largely facilitated by highly active CreateDebate users who collude among themselves to post together on the same discussion-reply threads and reply to each other⁹.

# CC	% Users	% Comments	% Ad hominem
≤ 10	87.8	11.8	16.8
11 – 50	8.1	10.5	24.6
51 – 100	3.6	31.0	26.2
101 – 2000	0.4	20.2	26.1
≥ 2000	0.1	26.5	55.2
Total	100.0	100.0	37.03

Table 6: Grouping users on the basis of their top-level comment count (denoted by column #CC) and analyzing percentage of users and top-level comments for each group. Column %Ad hominem denotes the percentage of ad hominem comments in the given group only.

Characterising Users Posting Ad Hominem Comments

A natural question at this point is whether it is possible to characterise, early in time, the users who have a tendency to post ad hominem comments. In this subsection, we show that characteristic differences exist between those who post ad hominem comments and those who do not. To this end, we utilise the profile page of the users maintained by CreateDebate. The profile page contains various information about an user. These include the *reward points* earned by a user, their *efficiency* while debating, the *number of debates* they participated in, the *number of comments* they posted since they joined the forum. Other additional features consist of *allies*, *enemies* or *hostiles* of any given user¹⁰.

We collected these characteristics for all the users of CreateDebate and partitioned the users into two classes—those who have posted ad hominem comment at least once on the forum (C_1) and those who haven’t (C_2), and then computed the average and standard deviation of these characteristics for these two classes. We also considered the average and

⁹Our observations also hold for other topics; these results are not shown for brevity.

¹⁰Allies are users with whom one has common interests and opinions; enemies are the ones with whom a user has different interests and opinions; a hostile is one who declared the user as their enemy. Please check CreateDebate FAQ page for more information.

standard deviation of reciprocity observed from the support and dispute networks.

We observed that users belonging to class C_1 , on average, tend to post more often, have more reward points and less efficiency while debating, have more number of allies, enemies and hostiles and also show higher reciprocity, both in support and dispute network, when compared to the users belonging to class C_2 . The results are shown in Table 7. It can be observed that the distributions of the two classes are statistically very different from each other (very low p -values using Mann-Whitney U test). In a future work, one can featurize these user profile based characteristics to build an ML model to predict the propensity of a user to post ad hominem comments early in time. However, development of such models would, in turn, need a detailed understanding of the temporal evolution of ad hominem usage in the platform. So, in the next section, we investigate the following—does the huge prevalence of ad hominem has any correlation with time, or was ad hominem always equally prevalent in the platform?

Understanding Temporal Variations in Ad Hominem Usage

CreateDebate was launched in 2008 as a tool ‘that democratizes the decision-making process through online debate’. However, as we have observed, the percentage of ad hominem content on the website is alarmingly high when compared to any other regular debate forum. So, what went wrong? To answer this question, we perform the first temporal analysis of ad hominem usage for the CreateDebate forum. Our first task involved generating temporal snapshots to capture the month-wise activity on the site. As our dataset spans from February 2008 to November 2021, it would be very clumsy to do the day-wise activity analysis, and a year-wise scheme will have only 14 data points. Hence, we chose to study month-wise activities. The variation in the percentage of comments which were flagged as ad hominem and the percentage of users who had been posting such comments is shown in Figure 5 for each month in the period. It can be observed that the plots for all the topical sub-forums follow a similar trend—initially they are stationary, then they show a steep rise and then they fall. In order to gain insights of what exactly triggered this sharp rise and fall, we performed change point detection experiments (Truong, Oudre, and Vayatis 2020) to quantitatively partition the corpus into three sub-corpora—the stationary \mathcal{H}_1 , the rise \mathcal{H}_2 and the fall \mathcal{H}_3 . We used variation in the number of comments posted, percentage of comments which were flagged as ad hominem and percentage of users who were posting such comments for each month across all topical sub-forums as input to the change point detection algorithm which uses dynamic programming to find the optimal partition using RBF kernels as cost function. The cutoffs for the partitions as predicted by the algorithm are *March 2017* and *September 2019*. Hence, the timeline for the sub-corpora are— \mathcal{H}_1 (February 2008 – February 2017), \mathcal{H}_2 (March 2017 – September 2019) and \mathcal{H}_3 (October 2019 – November 2021).

We then compared these sub-corpora by generating word-

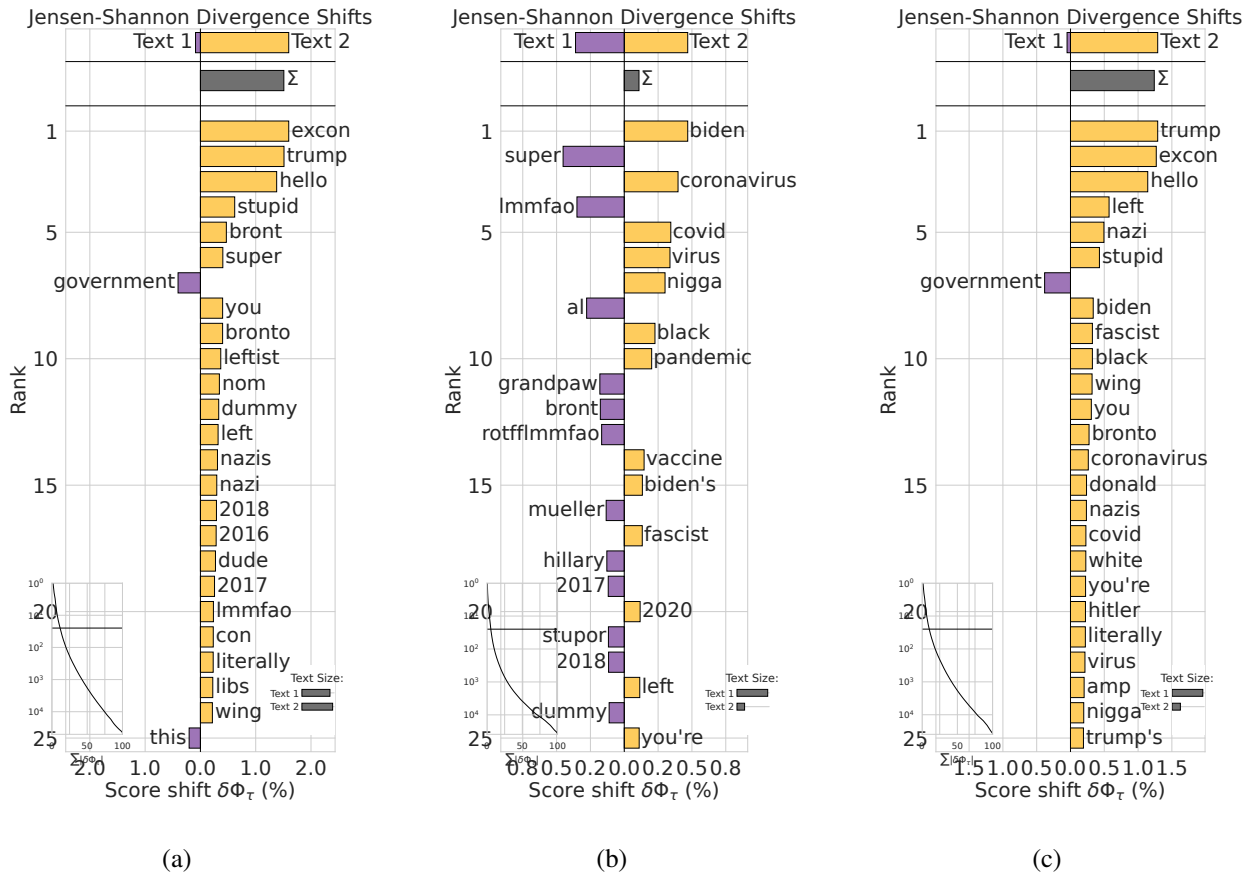


Figure 6: Word shift graphs constructed using Jensen-Shannon divergence for comparing sub-corpora (a) \mathcal{H}_1 and \mathcal{H}_2 , (b) \mathcal{H}_2 and \mathcal{H}_3 , and (c) \mathcal{H}_1 and \mathcal{H}_3 . Ad hominem triggering words are prominent in \mathcal{H}_2 when compared to \mathcal{H}_1 . For \mathcal{H}_3 , these triggers are not prominent.

shift graphs using Jensen-Shannon divergence, which are shown in Figure 6. It can be observed that the use of terms that act as triggers to ad hominem argumentation is very prominent in \mathcal{H}_2 when compared to \mathcal{H}_1 . These triggers are also present in \mathcal{H}_3 , but they are not as dominant as in \mathcal{H}_2 . We constructed the support and the dispute networks for \mathcal{H}_1 and \mathcal{H}_2 (see Figure 7) and observed that the reciprocity has increased significantly between users who wrote at least 100 top-level comments or received at least 200 direct replies, for support as well as dispute networks.

It is very interesting to note that the timelines for the 2016 US Presidential election and the Covid-19 outbreak are very close to the predicted cutoffs of the partitions. We observed that the CreateDebate forum was heavily used for political debates during the 2016 US Presidential election. Our hypothesis is that usage of ad hominem argumentation was accelerated after the 2016 US Presidential debates – forum becoming highly polar, people choosing sides. As it has been observed throughout history, the use of illogical arguments is very prevalent when people are discussing Politics, but to win debates, the use of ad hominem comments skyrocketed on the forum. However, due to the Covid-19 pandemic,

the activity of users on the forum declined significantly, thus curbing the ad hominem usage to some extent.

This is a plausible explanation about rising ad hominem arguments in political debates. But, what about other topical forums like Religion? Why is the ad hominem content increasing for these topics? To understand this complex phenomenon, we partitioned the religious debates using the above scheme and peeked into what users are talking about. We observed that the religious debates that were published before the 2016 US Presidential debates have a negligible political angle, however, those published after the Presidential debates were highly convoluted with Politics. As the results in Figure 8 show, each topical sub-forum on the CreateDebate site has a huge user overlap with Politics, especially the highly active users. For categories like Science and Law, the overlap approaches 100%. This explains why the increase in ad hominem comments and users posting them across different categories show similar trends as Politics (see Figure 5). So, it is likely that political discussions are root cause of the alarmingly high ad hominem content on CreateDebate, which skyrocketed after start of the 2016 US Presidential debates.

Characteristics	Class \mathcal{C}_1	Class \mathcal{C}_2	MWU p value
# posts	49.26 \pm 294.71	2.02 \pm 2.25	0
reward points	169.41 \pm 1070.35	6.89 \pm 14.31	0
efficiency	88.31 \pm 13.07	91.58 \pm 13.55	1.86×10^{-142}
# allies	1.84 \pm 9.11	0.14 \pm 0.83	2.76×10^{-249}
# enemies	0.66 \pm 6.21	0.03 \pm 0.24	3.67×10^{-203}
# hostiles	0.59 \pm 2.99	0.03 \pm 0.20	0
reciprocity (SN)	0.64 \pm 0.06	0.58 \pm 0.02	3.18×10^{-264}
reciprocity (DN)	0.69 \pm 0.05	0.64 \pm 0.03	1.22×10^{-263}

Table 7: Average and standard deviation of different characteristics along with p value computed using Mann–Whitney U test for the distribution of the two classes \mathcal{C}_1 (users who have posted at least one ad hominem comment) and \mathcal{C}_2 (users who have not posted any ad hominem comment). Low p value for MWU test denotes that the distributions of the two classes are statistically very different for each other across all user characteristics. SN denotes support network and DN denotes dispute network.




Figure 7: Reciprocity amongst the users of Politics sub-forum of CreateDebate computed using (a) support network and (b) dispute network constructed from sub-corpora \mathcal{H}_1 , and (c) support network and (d) dispute network constructed from sub-corpora \mathcal{H}_2 . λ and ρ are also varied. Reciprocity has increased significantly between the users for sub-corpora \mathcal{H}_2 when compared to \mathcal{H}_1 .

Concluding Discussion

In this work, for the first time, we shed a data-driven light on ad hominem usage in the wild by leveraging the CreateDebate data posted by tens of thousands of users over a period of more than a decade. We reported creating detectors with high accuracy whose judgment matched that of users for 94% of the cases. Using this detector, we uncover that very surprisingly around *one-third* of all content on CreateDebate is simply ad hominem. We deep-dived to find that a cohort of highly active users is responsible for this high fraction of ad hominems. Moreover, users particularly influenced by the political discourse resorted to this fallacy. While almost all data-driven studies suffer from intrinsic bias, we strongly believe this work is still valuable for understanding the ecosystem of cyber aggression as well as designing novel and more respectful debating platforms. In this final section, we will discuss the limitations as well as key implications of our work.

Limitations

Our work has a couple of limitations. First, our detector is bound by the annotation quality and volume of the CMV dataset. We believe we could have achieved higher accuracy with more data. However, even our detector achieved a significantly high accuracy compared to the prior works and we established that this detector is also valid on CreateDebate—a different forum and dataset, underpinning the efficacy of

the detector. Second, our results might or might not be generalizable beyond Reddit, CreateDebate and in general online forums to share opinions and debate. Even then, it has a lot of societal importance since, as prior works show, such debate forums in themselves are extremely important to shape the Internet and public opinion (Proferes et al. 2021). Third, our findings are often correlations and not causations. However, we believe these correlations show underlying large-scale behavioral patterns which promote ad hominem comments and, in effect, a toxic culture. Thus, in spite of this limitation, our work is still useful as a first attempt to shed light on ad hominem usage patterns.

Implications

We identify three key implications of our work for platform designers as well as platform regulators.

Defending against logical fallacies is becoming important: One very surprising and concerning finding is that our dataset from CreateDebate, a popular opinion-forming and debating forming forum, is filled with ad hominem fallacies. This finding highlights the importance of understanding and defending against logical fallacies which perpetuate toxic culture and attempts to stop any opposing arguments using irrelevant personal attacks. Thus, to counter cyber-aggression and bring back respect to online spaces, the platforms should acknowledge this issue and actively design defenses against such fallacies. The key focus today is on defending against hate speech and misinformation. The alarm-

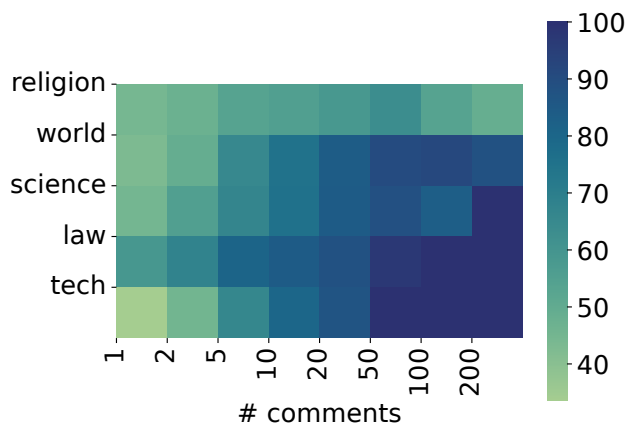


Figure 8: Overlap (%) of user base for different topical sub-forums with Politics sub-forum on CreateDebate. The overlap is computed across users with different comment count (denoted here as #comments). The heat map shows that the overlap is very large across different sub-forums, and it approaches 100% if we only consider the highly-active users.

ing rate of growth of ad hominem necessitates the design and development of dedicated countermeasures for this form of online malice.

It is possible to detect and defend against these fallacies via automated means: Our work demonstrated that we can leverage current extremely powerful techniques like BERT and GAN-BERT to create highly accurate ad hominem detectors, even when only a handful of annotated posts are available. Thus, our work can also be interpreted as a very strong proof of concept about using automated means (e.g., classifiers) by the platforms to protect against such fallacies along with hate speech and misinformation.

Users need to be nudged to reduce the usage of these fallacies: Finally, the results of this work strongly hint at the need of nudging users to reduce logical fallacies. As our results show, highly active user cohorts in CreateDebate are using ad hominem in more than 50% of their comments. Hence, regular online awareness campaigns could be organized to urge users to report about any such suspicious behavior that come to their notice. Furthermore, the substantially high ad hominem usage was rooted in the political climate of 2016 and now, it is spreading through other topics and forums, polluting the online space. The high volume of the affected population possibly even hints that many of these users might not even realize that they are utilizing fallacious arguments. Thus, the current platforms should focus on nudging the users against the usage of potential ad hominem even before they upload a fallacious post. Using the models discussed in this paper, online debate forums and social media sites can nudge users when they are about to post ad hominem comments by making them aware of the ad hominem triggers present in their post. Moderators can also decrease exposure of such comments by pushing them at the bottom of the thread, and even flagging them. Overall, we strongly believe that our findings will help policymakers

and platform developers help detect and defend against ad hominem fallacies in online opinion influencing forums.

Ethical Statement

In this work, we collected and analyzed data from CreateDebate and also conducted an annotation survey for validating our classifier. However, since we were analyzing user-generated data in this work, we tried our best to conduct our study ethically and protect the privacy of the users in our dataset. Specifically, we leveraged the previous work by Eysenbach and Till (2001) to check the ethics of our work. We noted that CreateDebate is a moderate-sized forum with around 15k members, and no registration was necessary to view and collect the CreateDebate data, signifying it was an ‘open’ forum. Finally, the debate topics often revolved around general phenomena (e.g., election), signifying the potentially public nature of our collected dataset. Nonetheless, following the footsteps of previous work by Cook, Ayers, and Horsch (2018), we hashed usernames after data collection to protect the privacy of the users during our analysis. Along the same lines of ethical consideration, for our annotation study, we did not collect any personally identifiable data from our participants to protect their privacy.

References

- Abbott, R.; Ecker, B.; Anand, P.; and Walker, M. 2016. Internet Argument Corpus 2.0: An SQL schema for Dialogic Social Media and the Corpora to go with it. In *LREC'16*.
- Boudry, M.; Paglieri, F.; and Pigliucci, M. 2015. The Fake, the Flimsy, and the Fallacious: Demarcating Arguments in Real Life. *Argumentation*, 29(4): 10–1007.
- Coe, K.; Kenski, K.; and Rains, S. A. 2014. Online and Uncivil? Patterns and Determinants of Incivility in Newspaper Website Comments. *Journal of Communication*, 64: 658–679.
- Cook, N.; Ayers, S.; and Horsch, A. 2018. Maternal post traumatic stress disorder during the perinatal period and child outcomes: A systematic review. *J Affect Disord* ., 2018(225).
- Croce, D.; Castellucci, G.; and Basili, R. 2020. GAN-BERT: Generative Adversarial Learning for Robust Text Classification with a Bunch of Labeled Examples. In *ACL'2020*.
- Das, M.; Saha, P.; Dutt, R.; Goyal, P.; Mukherjee, A.; and Mathew, B. 2021. You too Brutus! Trapping Hateful Users in Social Media: Challenges, Solutions & Insights. arXiv:2108.00524.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL'2019*.
- Eemeren, F. H. V.; and Grootendorst, R. 1987. Fallacies in Pragma-Dialectical Perspective. *Argumentation*, 1(3): 283–301.
- Eysenbach, G.; and Till, J. E. 2001. Ethical issues in qualitative research on internet communities. *BMJ*, 323(7321).
- Garlaschelli, D.; and Loffredo, M. I. 2004. Patterns of Link Reciprocity in Directed Networks. *Phys. Rev. Lett.*, 93.

- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *NIPS'14*, 2672–2680.
- Goodman, J. C. 2020. Four Years of Argument Ad Hominem. <https://www.independent.org/news/article.asp?id=13359>. Accessed: 2021-01-14.
- Habernal, I.; Wachsmuth, H.; Gurevych, I.; and Stein, B. 2018. Before Name-Calling: Dynamics and Triggers of Ad Hominem Fallacies in Web Argumentation. In *NAACL'18*, 386–396.
- Hamblin, C. L. 1970. Fallacies. *Tijdschrift Voor Filosofie*, 33(1): 183–188.
- Hasan, K. S.; and Ng, V. 2014. Why are You Taking this Stance? Identifying and Classifying Reasons in Ideological Debates. In *EMNLP'14*, 751–762.
- Jain, S.; Bhatia, A.; Rein, A.; and Hovy, E. 2014. A Corpus of Participant Roles in Contentious Discussions. In *LREC'14*.
- Kennedy, G. A. 1993. Aristotle "On Rhetoric": A Theory of Civic Discourse. *Philosophy and Rhetoric*, 26(4): 322–327.
- Macagno, F. 2013. Strategies of Character Attack. *Argumentation*, 27: 369–401.
- Mathew, B.; Illendula, A.; Saha, P.; Sarkar, S.; Goyal, P.; and Mukherjee, A. 2020a. Hate begets Hate: A Temporal Study of Hate Speech. arXiv:1909.10966.
- Mathew, B.; Saha, P.; Yimam, S. M.; Biemann, C.; Goyal, P.; and Mukherjee, A. 2020b. HateXplain: A Benchmark Dataset for Explainable Hate Speech Detection. arXiv:2012.10289.
- Mondal, M.; Silva, L.; Correa, D.; and Benevenuto, F. 2018. Characterizing Usage of Explicit Hate Expressions in Social Media. *New Review of Hypermedia and Multimedia*, 24(2): 110–130.
- Mondal, M.; Silva, L. A.; and Benevenuto, F. 2017. A Measurement Study of Hate Speech in Social Media. In *HT'17*.
- Proferes, N.; Jones, N.; Gilbert, S.; Fiesler, C.; and Zimmer, M. 2021. Studying Reddit: A Systematic Overview of Disciplines, Approaches, Methods, and Ethics. *Social Media + Society*, 7(2).
- Qiu, M. 2015. *Mining user viewpoints in online discussions*. Ph.D. thesis, School of Information Systems, Singapore Management University.
- Qiu, M.; Sim, Y.; Smith, N. A.; and Jiang, J. 2015. Modeling User Arguments, Interactions, and Attributes for Stance Prediction in Online Debate Forums. In *SDM*.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training. <https://openai.com/research/language-unsupervised>. Accessed: 2021-01-14.
- Redinger, C. 2020. OPINION: It's not debatable, ad hominem attacks destroy constructive conversation. <https://unothegateway.com/opinion-its-not-debatable-ad-hominem-attacks-destroy-constructive-conversation/>. Accessed: 2021-01-14.
- Rigotti, M.; Mikšović, C.; Giurgiu, I.; Gschwind, T.; and Scotton, P. 2022. Attention-based Interpretability with Concept Transformers. In *ICLR*.
- Sahai, S.; Balalau, O.; and Horincar, R. 2021. Breaking Down the Invisible Wall of Informal Fallacies in Online Discussions. In *ACL'21*, 644–657.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X.; and Chen, X. 2016. Improved Techniques for Training GANs. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 29.
- Schiappa, E.; and Nordin, J. P. 2013. *Argumentation: Keeping Faith with Reason*. Pearson UK.
- Sheng, E.; Chang, K.; Natarajan, P.; and Peng, N. 2020. "Nice Try, Kiddo": Ad Hominem in Dialogue Systems. *CoRR*, abs/2010.12820.
- Tan, C.; Niculae, V.; Danescu-Niculescu-Mizil, C.; and Lee, L. 2016. Winning Arguments. *WWW'16*.
- Tindale, C. W. 2007. *Fallacies and Argument Appraisal*. Cambridge University Press.
- Trabelsi, A.; and Zaiane, O. R. 2014. Finding Arguing Expressions of Divergent Viewpoints in Online Debates. In *LASM'14*, 35–43.
- Truong, C.; Oudre, L.; and Vayatis, N. 2020. Selective review of offline change point detection methods. *Signal Processing*, 167: 107299.
- Walton, D. 1998. *Ad Hominem Arguments*. Studies in Rhetoric and Communication. University of Alabama Press. ISBN 9780817309220.
- Wei, Z.; Liu, Y.; and Li, Y. 2016. Is This Post Persuasive? Ranking Argumentative Comments in Online Forum. In *ACL'16*, 195–200.
- Wei, Z.; Xia, Y.; Li, C.; Liu, Y.; Stallbohm, Z.; Li, Y.; and Jin, Y. 2016. A Preliminary Study of Disputation Behavior in Online Debating Forum. In *ArgMining'16*, 166–171.
- Woods, J. 2007. Lightening up on the Ad Hominem. *Informal Logic*, 27: 109–134.
- Wulczyn, E.; Thain, N.; and Dixon, L. 2017. Ex Machina: Personal Attacks Seen at Scale. In *WWW'17*, 1391–1399.
- Zalta, E. N. 2004. *The Stanford Encyclopedia of Philosophy*. Stanford, CA: The Metaphysics Research Lab.