

# Can we detect the need-for-closure bias in social media behavior?

Matthijs Westera, Mark Rademaker

The questions we entertain can affect how we process incoming information (Grossnickle, 2016). We are more likely to accept information that answers our previous questions, a bias known as need-for-closure (Webster & Kruglanski, 1994). Despite this, the effects of questions in social media on users' adoption of (dis)information have hardly been studied. Our work aims to detect the need-for-closure bias in a newly collected corpus of hundreds of thousands of Reddit posts. Do social media users more readily accept information that answers a question they were previously exposed to?

Through the Reddit API, we obtained the 1000 most recent posts from a predetermined set of subreddits. Among the authors of these posts, we selected user accounts older than three years, expecting enough longitudinal data to detect long-term effects of the need-for-closure bias. For each of these users, we downloaded their top 1000 posts, resulting in a dataset of over 400K posts, with collection ongoing. For each post, we also collected its parent post (if it replies to another post) and any children (replies).

In the resulting dataset, we identify what we call pivots: sentences from posts with which a user has interacted (e.g., replied to) and which potentially affected their beliefs. For each pivot, we check (i) whether it answers a question from a post the user previously interacted with, and (ii) whether the user subsequently adopts the pivot. We hypothesize that pivots that answer a prior question with which the user engaged, tend to be entailed more by the user's subsequent posts: a direct operationalization of the need-for-closure bias. We expect it to be only a small effect, due to the noisy and varying nature of social media text, but an effect nonetheless.

Given the scale of our dataset, we compute the two crucial features with off-the-shelf models for Question Answering (QA; an Albert-xxlarge model trained on the SQuAD v2 dataset (Rajpurkar, 2018)) and Natural Language Inference Model (NLI; an Albert-xxlarge model trained on the datasets MNLI (Williams, 2017) and SNLI (Bowman, 2015)): The first gives us a proxy for the degree to which a pivot answers a prior question with which the user engaged; the second gives us an estimate of whether subsequent posts by the user logically entail (i.e., endorse) the pivot. We worked around some limitations of these models; for instance, we exclude overly long or short sentences, and use additional classifiers to identify and exclude overly subjective or abstract sentences, which we found lead to false positives from the QA and NLI models. Moreover, in our statistical analysis we isolate potential effects of the QA and entailment scores from those of plain relatedness (embedding similarity).

We relate our findings to work on disinformation, work in Linguistics on how questions can be triggered implicitly, and to studies in Psychology on the effects of questions, beyond need-for-closure, on information processing, mental health and social cohesion. We hope to spur additional research into questions as a window onto our individual and collective curiosity, doubts, and vulnerabilities to disinformation.