



Google UK
Central St Giles
1 St Giles High St
London
WC2H 8AG

Julian Knight MP
Digital, Culture, Media and Sport Committee
House of Commons
London
SW1A 0AA

19th June 2020

Dear Mr Knight

Thank you for your time on June 4th and the committee's dedication to fighting COVID-19 misinformation and interest in our work at Google and YouTube.

What is YouTube doing to prevent content that is false surfacing on the home page, in-app notifications and suggested videos? Are there any steps you plan to take in the future?

Connecting our users with the content they love is important to us. We want to help viewers find new interests and passions—such as a new favorite artist, a new creator they can follow or simply the best food recipes. Our goal is for recommendations and for search results to point people to the highest quality, most authoritative information available. This is critical for topics like news and political issues, which are prone to misunderstanding and misinformation.

We have been systematically investing in improving the systems that recommend content across YouTube, both on the home page and via "Up Next." Since January 2019, we've launched over 30 different changes on YouTube to reduce recommendations of content that comes close to violating our Community Guidelines but doesn't quite cross the line, which we call borderline content.

Determining what is harmful misinformation or borderline can be difficult, especially for the breadth and variety of videos that are on YouTube. We rely on external evaluators located around the world to provide critical input on the quality of a video. These evaluators use [public guidelines](#) to guide their work. Each evaluated video receives up to nine different opinions and some critical areas require certified experts. For example, medical doctors provide guidance on the validity of videos about specific medical treatments to limit the spread of medical misinformation.

Continuously improving our recommendations system is a priority, and as we discussed before the committee, is a fundamental pillar of our responsibility work. And in tandem with this work to reduce recommendations of borderline content, we also continue to counterbalance this by raising authoritative voices by spotlighting news sources in search results and on the YouTube homepage. We also use several different types of information panels to highlight authoritative sources and provide users with more information and context when they search for or watch content related to certain issues.

We are also committed to taking a closer look at how we can further reduce the spread of content that comes close to—but doesn't quite cross the line of—violating our Community Guidelines and will continue to make necessary changes to improve the effectiveness of our efforts.

How much is this a function of your systems simply reflecting what other people have searched for?

We are constantly testing, learning and adjusting to recommend videos that are relevant to our users. Our recommendations systems take into account many signals, including a user's watch and search history (if enabled) as well as the channels that they have subscribed to. We also consider a user's context, such as their country and time of day. For example, this helps us show locally relevant news.

Another factor that YouTube's recommendation systems consider is whether others who clicked on the same video watched it to completion—a sign that the video is higher quality or enjoyable—or just clicked on it and shortly after starting to view the video, clicked away.

However, it's important to note that, where applicable, these signals are overruled by the signals relating our efforts to reduce recommendations of borderline content and harmful misinformation. This means in practice that—unless a user subscribes to the channel that posted the content—a video that has been identified as borderline content will not be recommended, even if other signals suggest it would be a viable recommendation.

We also ask users directly about their experience with individual videos and our recommendation systems using random surveys that appear on their homepage and elsewhere throughout the app.

[In June 2019](#) we announced changes that gave users more control over content they're recommended:

- Remove suggestions from channels users don't want to watch: We made it simple for users to tell us to stop suggesting videos from a particular channel, by tapping the three-dot menu next to a video on the homepage or Up Next, then "Don't recommend channel".

- Learn more about why a video was suggested: Sometimes we recommend videos from channels users haven't seen before based on what other viewers with similar interests have liked and watched in the past. When we make suggestions like this, we provide more information on why these videos are surfaced underneath the video in a small box.

We use this direct feedback to fine-tune and improve these systems for all users.

In addition, ['Your data in YouTube'](#) is a feature that makes it easier for users to understand and control what data is saved and how it is used on YouTube and across Google. It gives quick access to YouTube privacy controls, as well as information about how YouTube works with data such as a user's search and watch history. Users can remove specific videos from their watch history and queries from their search history, pause watch and search history, or start fresh by clearing their entire watch and search history. Any videos removed from their history will no longer influence their recommendations across the site.

What is Google doing to prevent instances where junk news, dis/misinformation and low quality search results resurfacing? Can you provide information on experiments, research and other measures you are taking to examine this?

Ranking algorithms are an important tool in our fight against disinformation. Ranking elevates the relevant information that our algorithms determine is the most authoritative and trustworthy above information that may be less reliable. Our ranking system is specifically designed to identify sites with high indicia of expertise, authority, and trustworthiness. We continue to develop and improve our ranking system to ensure we are delivering on our commitment to surfacing high quality search results. This work involves ongoing rigorous testing and evaluation processes to help us benchmark the quality of our results and make sure these meet a high bar all around the world. It also involves actively looking for and taking action against attempts to deceive our ranking systems or circumvent our policies.

As we've mentioned in previous correspondence, in 2019 alone, we ran over 464,065 experiments, with trained external Search Raters and live tests resulting in more than 3620 improvements to Search. To perform these evaluations, we work with Search Quality Evaluators who help us measure the quality of Search results on an ongoing basis. Evaluators assess whether a website provides users who click on it with the content they were looking for, and they evaluate the quality of results based on the expertise, authoritativeness, and trustworthiness of the content. The resulting ratings do not affect the ranking of any individual website in Google Search, but they do help us benchmark the quality of our results, which in turn allows us to build algorithms that globally recognize results that meet high-quality criteria. To ensure a consistent approach, our evaluators study and are tested on our [Search Quality Rater Guidelines](#) before they can begin to provide ratings. These are publicly available guidelines that effectively describe in great detail what our algorithms are trying to do and provide guidance and examples for appropriate ratings.

Over the years we've continued to update our ranking signals and our evaluation methods, including our search quality rater guidelines, to ensure these mechanisms are tuned to help continue to improve the quality of search results. For example, as referenced in "[Disinformation optimized: gaming search engine algorithms to amplify junk news](#)", "[junk news] discoverability abruptly declined beginning in August 2017 following major announcements from Google about changes to its search engine algorithms, as well as other initiatives to combat the spread of junk news in search results. This suggests that Google can, and has, measurably impacted the discoverability of junk news on Search..." (Bradshaw, 2019). The decline referenced occurred after a [series of steps](#) we took to strengthen our understanding of low quality content, enabling us to refine our algorithms. We updated our Search Quality Rater Guidelines and improved our ranking signals to incorporate other factors that allow us to understand page quality and authoritativeness better. The updates we made to our [Search Quality Rater Guidelines](#) more explicitly addressed misleading content by clarifying how high quality results should be evaluated.

Under the updated guidelines, we clarified that high quality results should be "accurate and highly credible" and that we do not think we've met our users expectations of quality when we surface pages which are "deliberately created with factually inaccurate content," or those which "directly contradict well-established scientific or medical consensus...[or]...historical fact". We also provided more detailed examples of low-quality webpages for raters to appropriately flag which can include misleading information, unexpected offensive results, hoaxes and unsupported conspiracy theories. These updated guidelines have continued to help our algorithms demote low-quality content and have allowed us to make additional improvements over time.

In addition to our ongoing work to evolve our algorithms to improve ranking, we also actively look for and take action against attempts to deceive our ranking systems or circumvent our policies. Our [webmaster guidelines](#) clearly spell out actions that are prohibited and state that we will take action against websites engaging in such behaviors. While not all spammers engage in disinformation, many of the malicious actors who try to distribute disinformation (at all levels of sophistication or funding) engage in some form of spam. The tactics they use are similar to those of other spammers. Therefore, our work against spam goes hand-in-hand with our work against disinformation. For most of the 25 billion spammy pages detected each day, we're able to automatically recognize their spammy behavior and ensure they don't rank well in our results. The remaining spam is tackled manually by our spam removal team, which reviews pages (often based on user feedback) and flags them if they violate the Webmaster Guidelines. Our efforts have helped ensure that more than 99% of visits from our results lead to spam-free experiences. We continue to be vigilant regarding techniques used by spammers and remain conscientious of what we share about the ways our ranking systems work so as not to create vulnerabilities they can exploit.

While our search results will never be perfect, we are continuously working to improve our ranking algorithms and prevent bad actors from deceiving our systems in order to surface high quality search results for users all over the world. You can learn more about the systems we

have in place to help ensure our Search algorithms meet high standards of relevance and quality on our How Search Works [website](#). We also recently published a [blog post](#) that provides details on our ongoing efforts to keep spam out of Search, including a link to our annual webspam [report](#) for more details on numbers.

Thank you again for the opportunity to share these details about our fight against disinformation. Please do not hesitate to get in touch if we can help with anything else.

Yours sincerely,

Leslie Miller , VP Government Affairs and Public Policy
Derek Slater, Global Director Information Policy