



16 September 2022

Australian Competition and Consumer Commission
Digital Platforms Inquiry
175 Pitt Street
Sydney NSW 2000

By email: digitalmonitoring@accc.gov.au

Dear Commissioner,

Thank you for the opportunity to make this submission and participate in the Australian Competition and Consumer Commission's (ACCC) public consultation following the release of the most recent *Digital Platform Services Inquiry March 2023 Report on social media services Issues Paper* (the Issues Paper).

Twitter's mission is to serve the public conversation, and we believe that regulation should promote and protect consumer choice and fair competition. Maintaining an Open Internet fosters a diverse ecosystem where consumers have variety and quality in their digital and online choices, as well the means to choose how they connect and find and access information.


Free discourse can only thrive if its underlying architecture is protected. We have a shared responsibility to promote a culture of innovation and ensure that smaller players or new market entrants can compete alongside more established services.

Twitter is committed to working with the government, our industry partners, academia, non-government organisations, and wider civil society as we continue to build our shared understanding of the issues and find optimal ways to approach these together.

We trust this written submission will be a useful input to the ACCC's consultation process. Working with the broader community, we will continue to collaborate to create a safe, open, and secure digital ecosystem.

Thank you again for the opportunity to provide input as part of this important process.

Kind regards,



Kara Hinesley
Director of Public Policy
APAC



Kathleen Reen
Senior Director of Public Policy
JAPAC



Table of Contents

Introduction	4
Our commitment to serving the public conversation and the Open Internet	5
Overview of the supply of social media services	6
Content moderation is more than just leave up or take down	7
Recommendation algorithms should be subject to human choice and control	8
Market structure and trends in the supply of social media services	9
Advertising and social media platforms	9
Relationships between digital platforms and consumers	9
Twitter’s Developer Platform and Commitment to Transparency	11
Conclusion	12



Introduction

Twitter's mission is to serve the public conversation. Twitter shows the world what is happening, democratises access to information, and at its best, provides people with insights into a diversity of perspectives on critical issues in real time.

As we consider the current regulatory landscape in Australia, we must also understand the diverse challenges and competing policy priorities at hand. The potential consequences for hasty policy decisions or balkanised approaches will stretch far beyond today's headlines, and are bigger than any single company. How these challenges are defined, understood, and addressed will affect services of all sizes, the ability of millions of people to share and access information both nationally and globally, and impact the future of the digital economy.

Regulatory approaches to new industries are often shaped by the policy responses designed in response to new technologies. This approach struggles to adapt to the unpredictable and rapidly evolving nature of human use of technology and expression. Designing regulation around the established online services of today can risk chilling innovation and further entrenching already dominant players, with negative consequences for digital participation.

More broadly, the policy issues addressed are often rooted in complex societal challenges that exist in offline, as well as online, contexts. As we continue to work together in good faith on these complex issues, we emphasise that these challenges will not be resolved by the removal of content online alone. Effective solutions demand a whole of society response that recognises the full scope of the problem being addressed. Policy design must also consider the effects, including unintended consequences, that any framework might have on essential democratic rights.

This submission will provide further insight into:

- Twitter's commitment to serving the public conversation and the Open Internet
- Overview of the supply of social media services
- Market structure and trends in the supply of social media services
- Relationships between digital platforms and consumers

Throughout the Issues Paper, the ACCC has recognised that regulatory intervention should be carefully designed to minimise unnecessary burdens and ensure that there are effective mechanisms to manage risks and mitigate potential harm for consumers. Twitter shares this concern.

Any regulatory responses must guard against unintended consequences that limit competition or innovation in fast-moving online markets and entrench dominant players. Noting that the forthcoming report may contemplate new recommendations and rules, we encourage the ACCC to resist 'over intervention' that risks stifling innovation and competition to the detriment of advertisers, platforms, and consumers. In general, we consider that any proposed interventions should be targeted only to those market participants whose dominant position undermines the competitive process, creates or entrenches barriers to entry or expansion, or gives rise to consumer harm.

Regulation should not be a one-size-fits-all approach. For example, regulatory reform or data portability requirements that apply to all firms regardless of their market position and power may inhibit new entrants and smaller market participants from competing effectively and devising and promoting new, competitive business models. Rules or requirements introduced to manage self-preferencing, conflicts of interest, and/or interoperability should be directed and only applicable to firms which the ACCC considers to have a market dominance or power and that engage in problematic or anticompetitive conduct.



We also urge the ACCC to ensure that any enforceable rules or requirements are not so prescriptive as to inhibit innovation as the internet and online markets rapidly change, and not so wide or general as to capture pro-competitive or innovative market behaviour. Any proposed regulation should be informed by Open Internet principles, and tailored to address specific conduct by dominant players.

Finally, we encourage the ACCC to consider a harmonised approach to market regulation. Given the global nature of markets, we consider that any regulatory reform pursued in Australia will only be effective if it is consistent and harmonised with international approaches to similar issues in international jurisdictions. We elaborate on these issues below, in the context of particular interventions identified in the Issues Paper.

Our commitment to serving the public conversation and the Open Internet

Twitter's rules aim to reflect the voice of the people who use the platform. We think it's critical to consider global perspectives, as well as make our content moderation decisions easier to understand.

Our aim is to have comprehensive policies that appropriately balance fundamental human rights and consider the global context in which we operate. We have recognised the role of deep consultation to appropriately address the complexity of online safety issues. Through our policies, products, and partnerships, we have undertaken coordinated efforts to consult with a range of partners, human rights experts, civil society organisations, academics, and the general public, whose feedback is reflected in revisions to the policy frameworks that govern the Twitter platform.

For example, we have called for public feedback in the development of a number of policies, including our synthetic and manipulated media policy, our approach to world leaders, and our hateful conduct policy. These consultations sought a variety of perspectives and voices when considering how to balance harm reduction and freedom of expression, as well as how to avoid any unintended consequences of removing legitimate speech from marginalised groups.

Twitter also recognises the need to balance tackling harm with protecting a free and secure Open Internet. In our advancements to promote online safety and address the policy challenges ahead, we have focused on five overarching principles and how they intersect with issues of competition, content moderation, and the role and responsibilities of services like Twitter.

With regards to the creation of digital policy, we believe:

- (1) The Open Internet is global, should be available to all, and should be built on open standards and the protection of human rights.
- (2) Trust is essential and can be built with transparency, procedural fairness, and privacy protections.
- (3) Recommendation and ranking algorithms should be subject to human choice and control.
- (4) Competition, choice, and innovation are foundations of the Open Internet and should be protected and expanded, ensuring incumbents are not entrenched by laws and regulations.
- (5) Content moderation is more than just leave up or take down. Regulation should allow for a range of interventions, while setting clear definitions for categories of content.

Twitter believes the foundation that underpins robust consumer protections and a strong competitive environment is preserving an Open Internet that is global, available to all, and built on open standards and the protection of human rights.



The Open Internet has been an unprecedented engine for economic growth, cultural development, and self-expression. For these benefits to continue, it must be available to all.¹ The infrastructure of the internet is itself now a geopolitical space.

Governments should prioritise policies, partnerships, and investments at home and abroad that support and defend the Open Internet, both through regulatory and standards bodies, as well as ensuring domestic regulation does not undermine global norms or set dangerous or conflicting precedents. Open standards championed by these bodies will provide for greater interoperability, connection, and competition.

Access is a critical issue. The principle that information should be able to safely and securely move across borders freely as part of a global internet should be core to democratic regulation.

Rhetoric and policies pursuing national data sovereignty should be avoided and scrutinised. Some actors seek to exploit this concept to strengthen control of and access to data, weakening the Open Internet through forced data localisation and limits on the free flow of data.

The principle that data belongs to a person does not mean that all people's data belongs to the state. Policymakers should avoid the use of extra-territorial application of national content standards as this further undermines the global internet and encourages a race to the bottom, with the entire world's open communications imperilled by those actors least committed to freedom of expression.

Both governments and industry should ensure their approaches to addressing online harm are consistent with universally recognised human rights norms, including proportionality and the protection of privacy and freedom of expression.

Overview of the supply of social media services

Twitter believes that competition, choice, and innovation are foundations of the Open Internet and should be protected and expanded, ensuring incumbents are not entrenched by laws and regulations.

A less competitive internet trends towards a less Open Internet. There's a risk that some regulatory interventions will undermine competition and entrench large, incumbent services, reducing consumer choice. It's not unique to the technology sector that incumbents often seek to use new regulations to protect their own market dominance, and in some cases, a less Open Internet may suit certain businesses.

Certain existing legal frameworks and market realities have fostered competition in technology sectors, including among social media platforms, both large and small, incumbent and entrant alike. It will be important moving forward for new and existing regulations to account for the way in which competition depends on these elements.

Intermediary liability protection is a foundation of the global, Open Internet and critical to competition online. Without this foundation, the internet as we know it — allowing speech, interaction, and discovery for billions of people — would cease.

Policymakers should avoid mandating technical means of implementation that have the effect of further entrenching services based on those tools and technologies. Additionally, they should guard against technical requirements that inadvertently benefit those that have the financial and technical means to deploy the particular implementation proposed. Opportunities to expand interoperability and the adoption of open standards will empower people with greater choice and flexibility about how they interact with online services and drive competition.

¹contractfortheweb.org/principle-theme/access/



Finally, the technologies that underpin the ability to address and remove the most harmful content and respond to further harms remain in proprietary silos, become exponentially more effective as businesses scale, and can further enhance dominance and undermine competition. Content moderation technology is one of the most significant barriers to entry, particularly as regulators set ever stricter requirements on the time within which harmful content must be removed. Policymakers should encourage and facilitate the availability of proactive technologies, and the data that underpins them, to enable new services and tools to be made more accessible to a greater range of services, including by providing a robust legal framework for information sharing.

Within the Issues Paper in this context, it is unclear if the definition of ‘social media services’ is the same definition used in the *Online Safety Act 2021* (Cth) (OSA) or a broader designation that encompasses additional products and services. That OSA defines ‘social media services’ quite broadly to encompass services where “the sole or primary purpose of the service is to enable online social interaction between 2 or more end-users; (ii) the service allows end-users to link to, or interact with, some or all of the other end-users; (iii) the service allows end-users to post material on the service”. We would note that further clarification of the scope of this specific inquiry in relation to relevant market participants and related policy proposals would be helpful.

Content moderation is more than just leave up or take down

Legislation and regulation should allow for a range of interventions while setting clear definitions for categories of content they seek to address, with substantive definitions and boundaries and be consistent with human rights standards.

Where the content at issue is lawful, but a government believes there’s a need to intervene, the regulatory framework should clearly distinguish between these types of content. Government requests for the removal of specific pieces of content based on illegality should be based upon valid legal process and provide for transparency about how these powers are used. It’s a fundamental question of due process that a government agency, not a private actor, is responsible for determining criminality. Except in specific, exigent circumstances consistent with prevailing law, companies should be free to provide notice to people around the basis for actions taken against them or their accounts.

Secondly, we believe the regulatory debate needs to reflect how content moderation is now more than just leaving content up or taking it down. For example:

- Providing users with context, whether concerning an account, piece of content, or form of engagement, is more informative to the broader public conversation than removing content.
- Providing settings to people and communities to control their own experience is empowering and impactful.
- Equally, deamplification allows a more nuanced approach to types of speech that may be considered problematic, better striking a balance between freedom of speech and freedom of reach. In the long term, how peoples’ attention is directed is a critical question.

Thirdly, regulatory frameworks that address system-wide processes, as opposed to individual pieces of content, will be able to better reflect the challenges of scale for all modern communications services. The approach is also more flexible and fit for purpose, recognising that the nature of challenges change depending on whether the goal is to protect a certain group, such as young people, or address a particular type of behaviour, such as platform manipulation.²

As has been noted by a range of voices, the combination of significant administrative penalties for individual pieces of content and expected removal of content in short time periods — whether one hour

² blog.twitter.com/en_us/topics/company/2020/new-approach-to-synthetic-and-manipulated-media.html



or 24 hours — creates a significant corporate incentive to over-remove content, particularly in edge cases. This would more acutely impact small companies and new services that have more limited resources to litigate or pay fines. These frameworks must be underpinned with strong, independent processes that are free from political interference and allow for civil society participation.

To avoid incentivising over-removal, regulation that assesses the system-wide performance of how services enforce their terms of service will provide essential flexibility and reduce incentives to over-moderate content, while incentivising robust appeal mechanisms, and investment in technological solutions.

Recommendation algorithms should be subject to human choice and control

Recognising that content moderation and content organisation are two different spheres of work, particularly when content is recommended without a positive signal to seek it out, policymakers should prioritise empowering people to have control over algorithms they interact with and ultimately drive an ability to make our own choices among algorithms. Choice can also help foster greater understanding and awareness of how algorithms impact people’s online experiences, leading to greater digital literacy.

While algorithmic transparency is an important part of deepening understanding of how these systems work, both in terms of process and training data, the focus on source code for algorithms - a literal interpretation of the phrase “algorithmic transparency” - offers flawed and unclear benefits. While in a limited context this may provide a small, highly technical audience with insights, it does little to change the experience of people online.

A clear example of such control, choice and transparency for users of Twitter is on the Home timeline. Since 2018, Twitter has provided those using the service the option to turn off our Home timeline ranking algorithm, returning them to the classic reverse-chronological order of Tweets. This means people can choose between viewing their top Tweets first (i.e. the top Tweets algorithmically ranked based on accounts and content the user interacts with the most), or the most recent Tweets first³. This option can also be both easily identified and executed by tapping the ‘sparkle icon’ in the main app interface. We also actively communicate this choice via official public communication channels, like company blogs, official accounts, and the Twitter Help Centre.⁴

In April 2021, we announced the establishment of Twitter’s Responsible Machine Learning Initiative, led by our internal Machine Learning Ethics, Transparency and Accountability (META) team. This team’s responsibilities include driving transparency about our machine learning decisions, how we arrived at them, and vitally, enabling agency and algorithmic choice for people on Twitter.⁵ As above, we believe it is important to understand the agency held by the individual when using the Twitter service and the choices available to them regarding how algorithms might affect what they see.

For Twitter, Responsible ML consists of: (1) Taking responsibility for our algorithmic decisions; (2) Equity and fairness of outcomes; (3) Transparency about our decisions and how we arrived at them; and (4) Enabling agency and algorithmic choice. Responsible technological use includes studying the effects it can have over time. When Twitter uses ML, it can impact hundreds of millions of Tweets per day and sometimes, a system designed to help can behave differently than was intended. These subtle shifts can

³ Blog.twitter.com. Never Miss Important Tweets From People You Follow. [online] Available at: https://blog.twitter.com/official/en_us/a/2016/never-miss-important-tweets-from-people-you-follow.html <Accessed Sept 2022>.

⁴ Twitter Support, 2018, Available at: <https://twitter.com/TwitterSupport/status/1075506036818104320> <Accessed Sept 2022>; Twitter Help Centre, 2021. Available at: <https://help.twitter.com/en/using-twitter/twitter-timeline> <Accessed Sept 2022>.

⁵ Blog.twitter.com. Introducing our Responsible Machine Learning Initiative, https://blog.twitter.com/en_us/topics/company/2021/introducing-responsible-machine-learning-initiative.html, <Accessed Sept 2022>.



then start to impact the people using Twitter, and so we want to make sure we're studying the changes and using them to build a better product.

This effort is part of our ongoing work to look at algorithms across a range of topics. We shared the findings of our analysis of bias in our image cropping algorithm and how they informed changes in our product.⁶Our teams are also publishing learnings from an in-depth analysis of whether our recommendation algorithms amplify political content.⁷

This research study highlights the complex interplay between an algorithmic system and people using the platform. Algorithmic amplification is not problematic by default – all algorithms amplify. Algorithmic amplification is problematic if there is preferential treatment as a function of how the algorithm is constructed versus the interactions people have with it. Currently, our teams are exploring opportunities this work may unlock for future collaboration with external researchers looking to reproduce, validate and extend our internal research, and will continue to share additional updates.

Looking ahead, Twitter also announced Bluesky, a non-profit initiative to develop a decentralised social network protocol that is currently in research phases.⁸ In the long term, as we envision @bluesky, controls will extend to the choice among ranking algorithms, and built on an open standard for social media to which we hope Twitter will ultimately adhere. The idea of “Protocols not platforms” is instructive not only for the technological potential for standardisation of ranking algorithms but also for the underlying impact this would have on protecting free expression and driving competition.⁹

Market structure and trends in the supply of social media services

Advertising and social media platforms

Further to these public facing mechanisms, we have engaged with rapidly developing, critical subject areas such as algorithmic transparency, choice, and ethics. Advertisers on Twitter are responsible for their Twitter Ads. This means following all applicable laws and regulations, creating honest ads, and advertising safely and respectfully – with Twitter Ads policies applying to all of Twitter’s paid advertising products.¹⁰ In addition to Twitter Ads Policies, advertisers must follow Twitter’s Terms of Service, Twitter Rules, and all the policies on our Help Centre governing use of our services.

Twitter believes political reach should be earned and not bought and we are committed to ensuring that ads on our service are free from bad-faith actors. In 2019, Twitter was the first social media company to end all political and issue advertising and prohibit state-controlled media authorities from purchasing advertisements on Twitter. This means Twitter globally prohibits the promotion of political content. We define political content as content that references a candidate, political party, elected or appointed government official, election, referendum, ballot measure, legislation, regulation, directive, or judicial outcome. Ads that contain references to political content, including appeals for votes, solicitations of financial support, and advocacy for or against any of the above-listed types of political content, are prohibited under this policy. We also do not allow ads of any type by candidates, political parties, or elected or appointed government officials.¹¹

Relationships between digital platforms and consumers

⁶ https://blog.twitter.com/en_us/topics/product/2020/transparency-image-cropping

⁷ https://blog.twitter.com/en_us/topics/company/2021/rml-politicalcontent

⁸ <https://twitter.com/bluesky>

⁹ knights.columbia.org/content/protocols-not-platforms-a-technological-approach-to-free-speech

¹⁰ <https://business.twitter.com/en/help/ads-policies.html>

¹¹ <https://business.twitter.com/en/help/ads-policies/ads-content-policies/political-content.html>



When examining the relationships between digital platforms and consumers, trust is essential and can be built with transparency, procedural fairness, and privacy protections.

There's a deficit in trust with respect to many online services and government functions alike. It's essential every sector works to rebuild trust, beginning with greater transparency. People should understand the rules of online services and the way that governmental legal powers are used. Transparency enables accountability for companies and governments. Without transparency, there can be no accountability.

Just as due process is a core feature of robust judicial systems, procedural fairness should be a core function of online services. These concepts should be at the core of regulation, particularly where governments seek to require services to remove content and companies take action under their terms of service.

Legislators should ensure clear harmonised standards for safeguarding and processing personal data, supplemented by regulatory guidance as new issues emerge, recognising that it's neither feasible nor desirable to legislate for every potential scenario of how personal data is used in primary legislation.

Fragmented and inconsistent frameworks weaken consumer protection and the establishment of industry norms. While many services do collect data to enable them to provide advertising, granular privacy controls balance the functionality of online services with consumer control while serving a desire to allow people who use services to make informed decisions about the data they share. Individuals should know and have meaningful control over what data is being collected about them, how it's used, and when it's shared. In the long run, regulation should protect and encourage services based on a range of business models, not just those built on advertising.

Policymakers should protect the ability for people to use the internet pseudonymously, anonymously, and without having to disclose one's identity or provide detailed personal information or government-issued identification. This is foundational to a universally accessible internet for all, and it's essential to recognise that not all services require the same amount of personal information to be disclosed or verified and nor should they be required to.

Our aim at Twitter is to have comprehensive policies that appropriately balance fundamental human rights and consider the global context in which we operate. We have recognised the role of deep consultation to appropriately address the complexity of online safety issues. Through the development of our policies, products, and partnerships, we have undertaken coordinated efforts to consult with a range of partners, including human rights experts, civil society organisations, academics, the general public, and our Trust & Safety Council, whose feedback is reflected in revisions to the policy frameworks that govern the Twitter platform.¹²

We want the Twitter Rules and Terms of Service (TOS) to ensure all people can participate in the public conversation freely and safely. We take our role in promoting healthy conversation seriously, and as part of that mission and based on feedback from partners, we allow people to use Twitter pseudonymously or as a parody, commentary, or fan account.¹³

Our partners have emphasised that anonymity and pseudonymity can be a critical entry point for people to connect online and build trusted relationships. Trust can be derived from the content someone posts and how they connect with others.¹⁴ Trust is based on what you do, not who you are, and anonymity or pseudonymity provides space for more people to express themselves freely and safely, in ways that actually engender that sort of trust-building connection. Simply put, in addition to providing safety, anonymity and pseudonymity provides people with the agency and control to

¹² <https://about.twitter.com/en/our-priorities/healthy-conversations/trust-and-safety-council>;
https://blog.twitter.com/en_us/topics/company/2019/synthetic_manipulated_media_policy_feedback

¹³ <https://help.twitter.com/en/rules-and-policies/platform-manipulation>

¹⁴ <https://blog.twitter.com/common-thread/en/topics/stories/2021/whats-in-a-name-the-case-for-inclusivity-through-anonymity>



choose how they present themselves. This has been a core tenet of the internet since its inception and is essential to a society that promotes individual choice and freedoms.

To be clear, pseudonymity is not a shield to the Twitter Rules or Terms of Service, nor is it a shield from criminal liability. When a person signs up for a Twitter account, we ask for a verified email or phone number, which can assist our teams in preventing ban evasion, platform manipulation, or compliance with duly executed valid legal requests for account information.¹⁵ Government agencies, law enforcement, and authorised representatives can submit requests for information about an account they're investigating any time via our dedicated Legal Requests Submission site.¹⁶

Notably, various academic studies have already shown that anonymity alone does not lead to harassment, and in many cases, people post harassing, toxic replies under their real name with a photo of their real face.¹⁷ Research into why people might harass others on the internet, and how to most effectively tackle the behaviour, continues to evolve, including by deepening our understanding of the social contexts these problems exist within. Currently, there is not conclusive evidence that requiring the display of names and identities will reduce social problems, and many studies have documented the problems it actually creates, like posing real threats to vulnerable communities.¹⁸

Additionally, the potential consequences involved with digital identification policies might have on online participation, access, and the widening of the 'digital divide' should be considered. Personal identification can pose risks to vulnerable groups who are not able to safely use services under their real name, such as those seeking information or support for domestic violence, whistleblowers, or LGBTQIA+ people. The requirement for digital identification verification could also cause inadvertent repercussions in people accessing services on the internet. According to the World Bank, an estimated 1 billion people worldwide do not have an official form of identification.¹⁹ It is often the marginalised, vulnerable, and impoverished who lack government-issued identification, leading to larger inequalities amongst people being able to access online services.

Twitter's Developer Platform and Commitment to Transparency

Transparency is core to Twitter's approach. We are committed to transparency surrounding our efforts to tackle some of the biggest challenges we all face online.

For the last ten years since 2012, we have provided biannual Transparency Reports, which shine a light on our own practices, including enforcement of the Twitter Rules and our ongoing work to disrupt global state-backed information operations.²⁰ We want the general public and policymakers to be better informed about our actions, which is why our original report evolved into a more comprehensive Twitter Transparency Centre covering a broader array of our transparency efforts. We now include sections covering information requests, removal requests, copyright notices, trademark notices, email security, Twitter Rules enforcement, platform manipulation, and state-backed information operations.

Through initiatives, such as Twitter's open application programming interfaces (APIs), our developer platform, our information operations archive, and our disclosures in the Twitter Transparency Centre and Lumen, we continue to support third-party research about what's happening on Twitter.²¹

We'll continue to build on these efforts and inform the public as we improve Twitter in the open. In the past year, we have:

¹⁵ <https://help.twitter.com/en/rules-and-policies/ban-evasion>

¹⁶ https://legalrequests.twitter.com/forms/landing_disclaimer

¹⁷ <https://jillancyork.com/2021/01/14/everything-old-is-new-part-2-why-online-anonymity-matters/>

¹⁸ <https://coralproject.net/blog/the-real-name-fallacy/>

¹⁹ <https://id4d.worldbank.org/global-dataset>

²⁰ <https://transparency.twitter.com/>

²¹ <https://transparency.twitter.com/>



- **Twitter API for Academic Research:** In early 2021, we launched a dedicated Academic Research product track on the new Twitter API giving qualified researchers access to the entire history of public conversation and elevated access to real-time data for free – the only major service to make public conversation data available via an API for research purposes.²² This raises awareness and increases understanding more widely of the scale and nature of the challenges impacting the integrity of public conversation online. By making this data available, Twitter can also be kept accountable for our own response to these challenges.
- **Algorithmic bias bounty challenge:** When we introduced our commitment to responsible machine learning, we also said, “the journey to responsible, responsive, and community-driven machine learning systems is a collaborative one.” That’s why we introduced the industry’s first algorithmic bias bounty competition to draw on the global ethical AI community’s knowledge of the unintended harms of saliency algorithms to expand our own understanding and to reward the people doing work in this field.²³
- **Twitter Moderation Research Consortium (TMRC):** We announced the creation of a new global expert group of academics, members of civil societies and NGOs, and journalists to study platform governance issues. We look forward to deeper analysis from this range of global experts and expect the collaboration to result in expanded disclosures beyond information operations to include sharing data in areas like misinformation, coordinated harmful activity, and safety.²⁴
- **Launch of an API curriculum:** “Getting started with the Twitter API for Academic Research” is now being used at universities, enabling students and teachers to learn how to use Twitter data for academic research. It is currently starred by over 200 academics on Github.²⁵
- **Developer Platform Academic Research advisory board:** This group of 12 scholars began work with our team this year to better understand how we can enhance the use of the Twitter API for academic research, while increasing meaningful dialogue between the Twitter Academic program and the academic community.²⁶
- **Developer research highlights:** We published and continued to spotlight key research areas Twitter teams are working on today in an effort to inspire even more researchers to pursue these topics.²⁷

As we continue to invite trusted partners and the public to share feedback on ways to make Twitter safe, it’s important to be transparent about how we develop and enforce the Twitter Rules. Our newly formed Content Governance Initiative (CGI) aims to do this by developing a governance framework that provides a consistent and principled approach to the development, enforcement, and assessment of our global rules and policies. To build our governance framework, we’re engaging external stakeholders and have created an additional advisory group on our Trust and Safety Council.²⁸

We’ll continue collaborating with this group and cross-functional teams across Twitter to establish standardised guidelines on policy development, enforcement, and appeals that help drive a common understanding of Twitter’s approach to content moderation. The framework’s principles and guidelines will aim to fulfil the following objectives:

- Build legitimacy and trust through transparency and accountability.
- Deepen our commitment to good governance and human rights.
- Provide additional clarity on Twitter’s content moderation processes.
- Affirm our commitment to serving a diverse and inclusive global community.

²² https://blog.twitter.com/developer/en_us/topics/tools/2021/enabling-the-future-of-academic-research-with-the-twitter-api

²³ https://blog.twitter.com/engineering/en_us/topics/insights/2021/algorithmic-bias-bounty-challenge

²⁴ https://blog.twitter.com/en_us/topics/company/2021-expanding-access-beyond-information-operations

²⁵ <https://github.com/twitterdev/getting-started-with-the-twitter-api-v2-for-academic-research>

²⁶

https://blog.twitter.com/developer/en_us/topics/community/2021/introducing-the-developer-platform-academic-research-advisory-board

²⁷ <https://developer.twitter.com/en/use-cases/do-research/academic-research/research-areas>

²⁸ <https://about.twitter.com/en/our-priorities/healthy-conversations/trust-and-safety-council>



We recognise that achieving these objectives will not be easy. Content moderation at scale is a highly complex and challenging process. This initiative reflects our ongoing commitment to working systematically – in partnership with external stakeholders around the world – to improve the transparency and consistency of our content moderation processes.

As the internet has become a foundational medium by which the right to freedom of speech and the right to reach information is exercised – as well as a central tool in commerce, education, and culture – in order to have equality of access, companies, and institutions alike must provide meaningful transparency and accountability.

Conclusion

Twitter is engaged in open dialogue with governments around the world as we seek to foster collaborative partnerships and continue to drive forward online safety solutions while protecting vital public expression.

The issues raised in the Issues Paper are broad and complex, and as the control of digital infrastructure is increasingly a focus of geopolitical action, these issues cannot be viewed in isolation. It is essential that there is a coordinated, multi-stakeholder strategy to respond to these threats and defend the free, secure, competitive, and global Open Internet.

Our work will never be complete as the threats we face constantly evolve. Going forward, we look forward to continuing to work collaboratively and in good faith with the ACCC, government, civil society, not-for-profits, academia, and industry to address online safety and work to create lasting global solutions to build a safer and Open Internet.