

Search Engines, White Ignorance, and the Social Epistemology of Technology

JOSHUA HABGOOD-COOTE

Abstract

How should we think about the ways search engines can go wrong? Following the publication of Safiya Noble's *Algorithms of Oppression* (Noble, 2018), a view has emerged that racist, sexist, and other problematic results should be thought of as indicative of algorithmic bias. In this paper, I offer an alternative angle on these results, building on Noble's suggestion that search engines are complicit in a racial contract (Mills, 1997). I argue that racist and sexist results should be thought of as part of the workings of the social system of white ignorance. Along the way, I will argue that we should think about search engines not as sources of testimony, but as information-classification systems, and make a preliminary case for the importance of the social epistemology of technology.

1. Introduction

In September of 2021, Google UK released an advert entitled *The more we learn, the closer we get*.¹ Against a moody backing track, the advert shows a series of characteristic images of Modern Multicultural Britain: a white teenager greeting a group of Black teenagers with a cheery 'wagwan'; an East Asian man looking out of a bus window at a group of men performing Salah; a Black man at a Ceilidh, a Black woman looking a group of South Asian people celebrating Diwali in a back garden; a white mechanic noticing his distressed colleague; a Black boy looking at a mural of Marcus Rashford (which had recently been defaced following a missed penalty). Following these images, the narration – read by Rashford – says 'it's not our questions that define us. But what we do with the answers.' Returning to the initial characters, the advert shows a series of search bars being filled in: 'who can say wagwan', 'whats a ceilidh', 'how to check on someone', and 'how can

¹ https://www.youtube.com/watch?v=s1kUMBCVvMY&ab_channel=GoogleUK.

doi:10.1017/S003181912400007X © The Author(s), 2024. Published by Cambridge University Press on behalf of The Royal Institute of Philosophy. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

we understand one another'. The video cuts to a headshot of Rashford, with the voiceover 'because the more we learn, the closer we get.'²

The central claim of this paper is that we should think about Google search as a part of the social institution of white ignorance: an institution which fosters miscognitions to both maintain and obfuscate the existence of White Supremacy (Mills, 1997, 2017). This means that this advertisement is an instance of undermining propaganda (Stanley, 2015), in the sense that it deploys the ideal of mutual understanding in a multi-racial society to advertise a product which actively undermines the realisation of this goal. I take the connection between Search Engines and white ignorance from the work of Jessie Daniels and Safiya Noble, who both draw on Mills's early work in *The Racial Contract* to understand the problems with Google search (Daniels, 2009, pp. 8, 20; Noble, 2018, p. 60; see also Frost-Arnold, 2023, Ch. 4). You can think of this paper as a remix of ideas from Noble's and Daniels' work which makes epistemological issues the central theme.³

I have two background goals. The first is to demonstrate the importance of drawing on critical technology scholars when thinking about the social epistemology of technology (see Frost-Arnold, 2023). The second is to establish the importance of thinking about technological systems as co-constitutive with social systems (see Benjamin, 2018). By thinking about technologies as part of the architecture that scaffolds our social lives, we can think more clearly about the problems of technology, avoiding both the *technodeterminist* view which thinks of technology as having an inexorable power over our lives, and the *technovoluntarist* view that technology is a neutral tool whose uses are determined by our social practices.⁴

² This followed the earlier advert *It all starts with summer*, which was ubiquitous on British television through the summer of 2021 (https://www.youtube.com/watch?v=tEfhyYJgcZ4&ab_channel=GoogleUK).

³ White ignorance is not the only system of ignorance production which serves an unjust political system. Mills's work on the racial contract was inspired by Carol Pateman's *The Sexual Contract* (Pateman, 1988) and Marxian ideology critique. It is plausible that there is a system of *Patriarchal Ignorance* which serves to maintain and obfuscate the patriarchal social system and its associated sexual division of labour and subjugation of women, and a system of *Capitalist Ignorance* which serves to maintain and obfuscate the capitalist economic system and the realities of the labour process. Our discussion below will touch on the way Google Search is operating at the intersection of all three systems of ignorance production, representing women as sexual objects (see section 4.2), and obfuscating the considerable human platform labour (see footnote 9) which is required to produce search results.

⁴ See Kukla (2021) on spatial determinism and spatial voluntarism.

The plan of action is as follows. In the first section, I introduce Mills's work on the Racial Contract and develop the notion of white ignorance as an ignorance-producing social institution. In the second section, I develop a basic picture of the mechanics of search and consider whether we should think about a search engine as a source of testimony, or as a relevance-filtering device (Munton, [forthcoming](#)). In section three we turn to Daniels' and Noble's work, surveying problematic autocomplete, image, and search results. In section four, we consider how we ought to think about these results, arguing that the framework of white ignorance gives us a helpful way to bring together the role of design, user behaviour, and structural features of search engines.

2. White Ignorance and the Racial Contract

There is a tendency within technology criticism to see the internet as a new social space governed by a race-blind contract. John Perry Barlow articulates this view clearly in *A Declaration of the Independence of Cyberspace*:

You claim there are problems among us that you need to solve. You use this claim as an excuse to invade our precincts. Many of these problems don't exist. We are forming our own Social Contract. This governance will arise according to the conditions of our world, not yours. Our world is different. [...] We are creating a world that all may enter without privilege or prejudice accorded by race, economic power, military force, or station of birth. (Barlow, 1996)

This idea is remarkably resilient, despite the abundant evidence about how race is played out through online spaces (Daniels, 2013, 2015).

In *The Racial Contract*, Charles Mills mounts a sustained critique of the raceless social contract, arguing that the social contract tradition within political philosophy has ignored the existence of a Racial Contract which governs the status and entitlements of whites and Blacks (Mills, 1997).⁵ This actual contract consists in a set of agreements between subjects racialised as white that prescribes a social ontology that partitions people into white persons and non-white sub-persons, with the white authors of the contract being marked for privileged access to the bodies, land, and resources of

⁵ In putting the racial contract at the centre of Mills's work on White Ignorance, I am relying heavily on Bain (2018).

non-whites, whilst the non-white subjects of the contract are correspondingly marked as targets for exploitation. Mills argues that despite their deployment of egalitarian social contracts, modern European states have enacted a global system built on colonial exploitation which is governed by an implicit commitment to the Racial Contract. What was once a *de jure* set of legal agreements that established a subhuman status for blacks has now become a *de facto* set of social practices, which maintain social hierarchies in large part by obscuring the history of how these hierarchies came about (Mills 1997, pp. 77–78).

Although white subjects have had – and continue to have – a clear economic interest in maintaining the system of global white supremacy, it is uncomfortable to hold in one’s mind both a commitment to racial hierarchy and a commitment to the ideals of egalitarian humanism. In Mills’s view, this tension is managed via what he calls the *epistemological contract*:

Thus in effect, on matters related to race, the Racial Contract prescribes for its signatories an inverted epistemology, an epistemology of ignorance, a particular pattern of localised and global cognitive dysfunctions (which are psychologically and socially functional), producing the ironic outcome that whites will in general be unable to understand the world that they themselves have created. (Mills, 1997, p. 18)

This contract ensures that both whites and non-whites misperceive the world in a way that obscures the existence of the Racial Contract, its harms to non-whites (Mills, 1997, pp. 98–101), the history of colonial exploitation (1997, p. 77; 2017) and justifies thinking of non-whites as sub-persons (1997, pp. 59–61). Mills initially presents this system of ignorance production as a set of cognitive norms (1997, pp. 17–18), but in later work he highlights the way this contract works through perception, conception, memory, and testimony (Mills, 2017, 60–71), and offers an account of the role of institutions in ignorance production (2017, pp. 66–68). Crucially, the institution of white ignorance obscures its own existence, in part through ‘strategic colour-blindness’ which refuses to engage with questions about racial inequality, and the history of racial exploitation (2017, p. 64).

Like any complex social phenomenon, the system of white ignorance emerges through the interplay of individual habits, social practices, and institutional systems. Given this, there is a question about whether to use the label ‘white ignorance’ to refer to the whole social system, its cognitive substrate, or the misrepresentations which are

the outputs of this system (El Kassar, 2018; Martín, 2021). Mills appears to vacillate between all three uses,⁶ but as our interest is in thinking about the role of technological systems in ignorance-production, it will be helpful to use ‘white ignorance’ to refer to the social system of ignorance production which serves Racial Contract.

Mills discusses several mechanisms which are important to this social system. For our purposes, the most important are the provision of false information, the production of controlling images, and the direction of inquisitive attitudes:

Provision of false information

In *White Ignorance*, Mills focuses on the provision of false information (Mills, 2017). In this essay, he employs Goldman’s veritist approach to social epistemology, arguing that epistemic sources which have traditionally been thought of as sources of knowledge – perception, testimony, memory – can also be sources of racialised ignorance.

Production of controlling images

In *White Ignorance* (2017, pp. 64–65), as well as his discussion of the presentation of wild men and wild spaces in the *Racial Contract* (1997, pp. 41–52), Mills makes clear that he takes conceptual resources to play an important role in the systems of white ignorance. Borrowing a concept from Patricia Hill Collins, we might suggest that *controlling images* of people racialised as Black – which function to other, objectify, and situate as deviant – will play an important role within the system of white ignorance.

Direction of inquisitive attitudes

In a passage riffing on the silence of white European philosophers on the ‘question of race’, Mills asks:

Where is Grotius’s magisterial *On Natural Law and the Wrongness of the Conquest of the Indies*, Locke’s stirring Letter concerning the Treatment of the Indians, Kant’s moving *On the Personhood of Negroes*, Mill’s famous

⁶ In *The Racial Contract*, the epistemological contract appears to mandate a social practice which is a part of the institution of the Racial Contract, in *Global White Ignorance* he describes white ignorance as a cognitive outlook ‘an absence of belief, a false belief, a set of false beliefs, a pervasively deforming outlook’ (2015, p. 217), and in *White Ignorance* he glosses White Ignorance as ‘an ignorance, a non-knowing, that is not contingent, but in which race [...] plays a critical causal role’ (Mills, 2017, p. 56).

condemnatory Implications of Utilitarianism for English Colonialism, Karl Marx and Frederick Engels's outraged Political Economy of Slavery? Intellectuals write about what interests them, what they find important, and – especially if the writer is prolific – silence constitutes good prima facie evidence that the subject was not of particular interest. (1997, p. 94)

This passage is important because it highlights that white ignorance works not only through the propagation of false information in the service of cultivating false belief, but also through the construction of topics as issues for public debate and contestation (see Case, 2018 on the contestation of issues in the so-called 'Age of Questions'). Which knowledge people produce depends on which subjects they investigate, which depends on which subjects are widely taken to be pressing and worthy of investigation (Pepp, Michaelson, and Sterken, 2022).

To underline the aptness of applying the idea of white ignorance to technological systems, it is worth noting that when Mills introduces the notion of an epistemology of ignorance in the passage quoted above, he immediately reaches for a technological metaphor:

To a significant extent, then, white signatories will live in an invented delusional world, a racial fantasyland, a "consensual hallucination", to quote William Gibson's famous characterisation of cyberspace, thought this particular hallucination is located in real space. (Mills, 1997, p. 18)

If we can use the idea of cyberspace to get a grip on the effects of white ignorance, then we might think that we can use the idea of white ignorance to get a grip on the epistemology cyberspace.

3. The Function of Search Engines

To understand the problems of search engines, we need to have a working model of what the functions of a search engine are. At the mechanical level, a search engine is a technosocial system that embodies a function from structured strings of text ('courgette & how to grow') to ranked sets of links to sites. Most commercial search engines also return integrated adverts and snippets of text which purport to provide helpful information (normally stripped from Wikipedia). The ranking of results to a given input is determined by a combination of the PageRank algorithm (Brin and Page,

1998),⁷ the AdWords system (Zuboff, 2018, Ch. 3), some level of personalisation (Levy, 2010),⁸ and a large amount of human labour put into ranking results (Newitz, 2017; MacDonald, 2020).⁹ As a matter of philosophical analysis it is not obvious what *should* determine the ranking of results, a question which opens up more general questions about the function of search engines (Broder, 2002). Google employees typically lean on the idea that search engine results should be *relevant*, but this term is never made fully clear.

The model that is ready to hand is to think about search engines as artificial testifiers (Gunn and Lynch, 2019). This model is a little strained: the inputs to search engines are not typically interrogative sentences, and testimony doesn't ordinarily return ranked lists of answers to a question. It is also notable that Google employees seem not to think about search engines in this way (Metzler, Tay, Bahri, and Najork, 2021; see also Shah and Bender, 2022). We might try to finesse the model, by suggesting that a search engine is a distinctive species of testimony: Simpson suggests that we think about search engines as expert testimony that expresses understanding of a subject matter (Simpson, 2012), and Munton suggests that search engines provide information about the question of where one might find information about a given topic (Munton, [forthcoming](#)).¹⁰ I want to take a different tack. Picking up on Noble's suggestion that the problems with Google Search are akin to the problems with library classification systems (Noble, 2018, Ch. 5), I suggest we think about search engines as information-classification systems.

What is an information classification system? Information scientists present library systems as organising resources by which subject they are *about* (Joudrey, Taylor, and Wisser, 2018, pp. 25–26). Philosophers of language and linguists present *aboutness* as a matter of a representational device (sentence, conversation, book) being associated both with propositional content, and a subject matter (Roberts, 1996; Yablo, 2014; Szabó, 2017). The propositional content can be thought of as a set of worlds which the representational device locates us in, and the subject matter as a set of sets of worlds

⁷ Which is often treated as approximating a Condorcet Jury theorem situation (Masterton, Tolsson, and Angere, 2016).

⁸ On the limited effectiveness of personalisation, see Feuz, Fuller, Stalder (2011) and Hwang (2020).

⁹ On platform work, see Roberts (2019), Gray and Suri (2019), and Jones (2021).

¹⁰ In the terminology of Habgood-Coote (2022a), search engines answer methodological questions rather than object questions.

which the representational device is poised to choose between (we can think about this as the question which the device aims to resolve (Roberts, 1996)). An information classification system sorts representational devices not just by their propositional content, but also by their subject matter. At first pass, we might say that the classification mark ‘Britain’ in a library groups books together which are about the subject matter: *Britain*.

To make good on this suggestion, we need to acknowledge that search engines are rather distinctive information classification systems:

- First, unlike the library classification systems which are relatively static (the shelving of new books notwithstanding), a search engine is a dynamic system. A search engine will take an open-ended set of inputs, will reckon with the torrent of new webpages, and will often output different ranked sets of results for different users.
- Secondly, the lack of a controlled vocabulary for input terms creates a significant problem of underdetermination for search engines. When we input a string in keywordese – ‘Courgettes grow UK’ – the sentence will dramatically underdetermine what topic the user is interested in (*how to grow Courgettes in the UK?, where do Courgettes grow best in the UK?, history of Courgettes in the UK?*). Search engines will not only need to semantically interpret the input string, but bring in contextual information to determine what subject matter to return resources about.
- Thirdly, the fact that search engines rank their outputs raises a tricky question about how the ranking ought to work. Work on subject matters in philosophy suggests a number of metrics: relevance (how similar the subject matter of the webpage is to the subject matter of the search input), informativeness (how much information is provided about the subject matter of the search input, see Groenendijk and Stokhof, 1984, pp. 379–80), accuracy (how many of the sub-questions in the subject matter of the search input have been answered correctly, see Habgood-Coote, 2022b), and quality (how good the information is on certain contextually determined standards; for example, for a search query for recipes how *tasty* a recipe is (on qualitative and quantitative gradability of answers, see Pavese, 2017)). Getting from this bundle of metrics to a single ranking is a tricky problem.¹¹

¹¹ We can think about this problem of combining these different kinds of goodness of resources relative to a subject matter into one ranking as a social choice problem, see D’Ambrosio and Hedden ([forthcoming](#)).

Whereas testimony can go wrong by providing false, unhelpful, or misleading information, an information classification system can go wrong by distorting a topic (Simpson, 2012, pp. 433–37; Munton, forthcoming). If a library has a shelf for books about Britain, but there are only books about England on that shelf, even if each of these books is relevant, informative, accurate, and well written relative to the topic *Britain*, as a whole the classification system has gone wrong. This grouping of resources under *Britain* distorts both *Britain's present* – at the time of writing, Wales, Scotland, Northern Ireland and the oversea territories are part of Britain – and *Britain's past* – leaving out the many places that have been colonised by Britain since 1542.¹² This failure of categorisation will also convey false information about the geography and history of Britain, but I suggest that we think about this primarily as a failure in our handling of subject matters.

4. The Problems of Search Engines

Search engines – especially Google Search – returning problematic patterns of results has been well-known for over a decade. Since 2005, cloaked websites run by white supremacists have been showing up in search results for innocuous terms (Daniels, 2009, Ch. 7). During 2015 and 2016 a series of examples of problematic search results circulated widely on social media, including the query ‘gorillas’ on Google Image Search returning images of two black teenagers (Noble, 2018, pp. 7–9, 81–83, 113).

Noble approaches Google search with the tools of *Critical Discourse Analysis*, an approach to social research which investigates the way in which language and social practice constitute one another, using the close reading of a small set of texts, guided by works of social theory (Recuber, 2016). Noble's corpus is a set of Search and Image results, and she is guided by work in Black Feminist theory. One of the issues with this approach is that it prioritises depth over breadth, and it would be good to understand how widely the results Noble discusses were returned. Answering this question would be difficult to carry out without a more systematic quantitative investigation, so in lieu of that, I have compiled a set of search results

¹² The date of the Tudor invasion of Ireland which led to the appropriation of land and the establishment of the system of corporate-backed plantations across Ireland.

for the queries of interest which provide some corroboration of the patterns of problematic results which Noble discusses.

Our plan of action is as follows: I will discuss Noble's examples of problematic search, image and autocomplete results in turn, situating them within our model of the epistemology of search and Mills's discussion of white ignorance, before highlighting some contemporary examples of similar problems.

4.1 Search Results

Noble focuses on two problematic patterns in search results:

- In 2015, the query 'black on white crimes' returned several cloaked white supremacist websites providing false narratives about the perpetrators of violent crime (Noble, 2018, pp. 113–14; see Daniels, 2009, Ch. 7);
- In 2011 and 2012 the query 'Black girls' returned a series of pornographic websites and adverts providing racialised pornography, as did searches for 'Asian girls', and 'Latina girls' (Noble, 2013, 2018, pp. 64–78).

The combination of algorithmic recommendation and websites that fake their credentials means that the early part of the story of fake news is the story of white supremacist sites (Frost-Arnold, 2023, Ch. 4). There are two problems with cloaked websites: they make lots of false claims, and including them in search results boosts their credibility. Although websites like The Council of Conservative Citizens, and New Nation News which show up on Noble's searches make lots of claims about the topic of crime between racialised groups – meaning that these sites are relevant – these claims are systematically false and support white supremacist propaganda. The inclusion of these results on the first page of Google search – typically the home of authoritative websites like Wikipedia and the Encyclopaedia Britannica – boosts the credibility of these White Supremacist sites and obfuscates their goals and ownership.

By contrast, the problem with the results for 'Black girls' is not that the sites linked are conveying false claims (although we might think that pornography presupposes falsehoods about women's willing subordination (Langton and West, 1999)). There are two ways in which we can think about these results within the model of information classification systems. The first is to think about them as a

mischaracterisation of the topic of Black women. By including a large amount of highly ranked pornographic content, the search results create the impression that Black women are primarily of interest as objects of sexual desire and use. This impression would be problematic for any social group but given the history of sexualised oppression of Black women, and the availability of dehumanising propaganda about Black women, this kind of miscategorisation is particularly harmful (Noble, 2018, pp. 92–104). The second is to think of this result as the problematic *interpolation* of a topic. When Google search returns a set of pornographic results for the query ‘Black girls’, we can read the system as narrowing the topic corresponding to this keyword to something like ‘how can I find pornography featuring Black girls?’. It is as if the search engine is autocompleting ‘porn’ for every search for involving ‘Black girls’. Although it is tempting to think that this distortion of the topic at issue is a simple consequence of the preponderance of searches for pornography on the internet, Noble argues persuasively that these results occur in large part because of Google Search’s advertising model, which allows companies to purchase the results for keywords, through a combination of search engine optimisation and advertising. She characterises this process as a kind of commodification of identity markers (Noble, 2018, 86–92).

At the time of writing, these patterns of results have changed considerably. A search for ‘Black on white crimes’ does not feature cloaked websites and has a number of websites debunking white supremacist tropes. With that said, these results do include a link to an article on the website for the Heritage Foundation – a Right-Wing US think tank – which uses the FBI figures on crime to try to undermine the idea that crime against Black people is due to White Supremacy.¹³ A search for ‘Black girls’ no longer links to preponderance of pornographic content, instead providing information more relevant to Black teenagers. However, it looks likely that Google search has patched rather than fixed the problem. Links to pornographic websites do still show up for ‘Asian girls’ and ‘Latina girls’ on the first page of results, and advertises for ‘Asian dating sites’ feature prominently on the results for ‘Asian girls’.

More worryingly, searches in other languages return patterns of results like those highlighted by Noble. Searches for ‘filles Noires’ (figure 1) and ‘filles Asiatiques’ (figure 2) on [google.fr](https://www.google.fr) on the 27th July 2022 returned a roughly even split between links to

¹³ <https://www.heritage.org/crime-and-justice/commentary/who-suffers-the-most-crime-wave>.

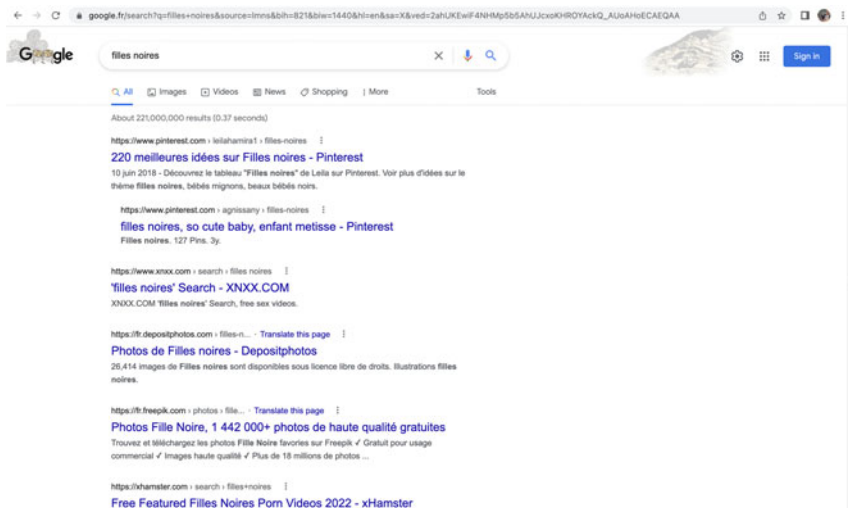


Figure 1. Search for ‘filles noires’ on [google.fr](https://www.google.fr) on 27 July 2022.

pornographic and non-pornographic websites on the first page of results, with the latter including several adverts for dating sites. Several other widely spoken languages (including Russian and Italian) return similar results, although this pattern is by no means universal in non-English languages (perhaps partly due to the ambiguity of ‘girl’).

Further evidence comes from Google’s keywords planner: a tool for advertisers which shows associations between searches. When prompted with the keywords ‘Black girls’, ‘Asian girls’, and ‘Latina girls’ in 2020, the keywords suggested by the tool were overwhelmingly pornographic, while keywords suggestions for ‘white girls’ were simply blocked (Yin and Sankin, 2020).¹⁴

4.2 Image Results

Many of Noble’s most striking results involve Google Image search results. I want to highlight two:

- In 2014, a search for ‘Black girls’ returned a highly sexualised set of images of Black women (Noble, 2018, p. 20);

¹⁴ <https://themarkup.org/google-the-giant/2020/07/23/google-advertising-keywords-black-girls>.

Search Engines, White Ignorance

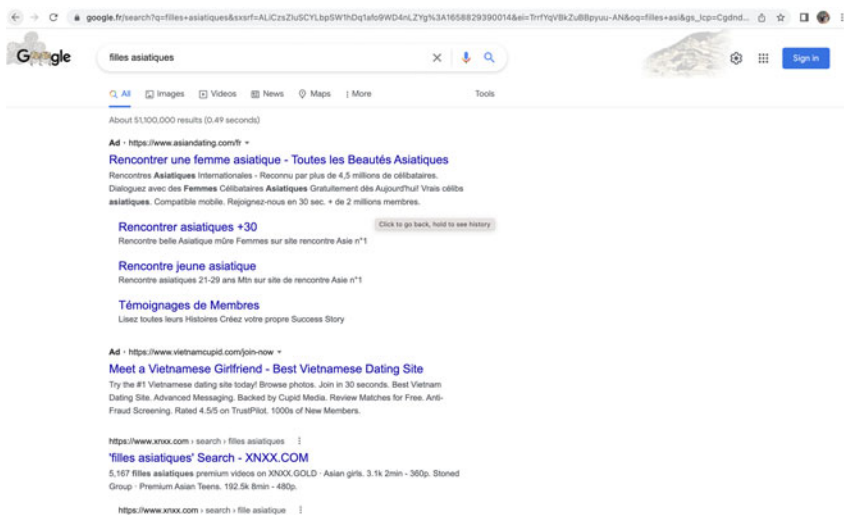


Figure 2. Search for ‘filles asiatiques’ on [google.fr](https://www.google.fr) on 27 July 2022.

- In 2016, a search for ‘unprofessional hair’ returned a set of pictures of Black women, whilst a search for ‘professional hair’ returned a set of pictures of white women (Noble, 2018, p. 83, citing a tweet from @BonKamora).

To think about these results, we need a model of Google Image search. If text search is a social-technical system which embodies a function from queries referring to topics to sets of links to websites which purport to provide information about those topics, we might think of image search as a social-technical system which embodies a function from queries referring to objects, activities, and events to sets of images which purport to represent those objects, activities and events. For example, the query ‘dog’ should – if all is going well – output a set of images of dogs. Some of the complexity of Google image search arises from the fact that its images – sourced from across the internet – are often taken to represent a social visual imaginary, what Noble calls ‘algorithmic conceptualizations’ (Noble, 2018, p. 24). One way to flesh out this idea is to propose that Google Image outputs as a set are typically taken to be *characteristic* or *representative* images of the object, activity, or event in question. It would be bad if a Google Image result for ‘dog’ returned only images of dogs wearing bandanas (adorable, but not

characteristic of the species), or if all of the dogs returned were Leonbergers (majestic, but not representative of the diversity of dog breeds).

With this model, we can start to think through the problems with Noble's examples.

Much as the Google search results for 'Black girls' mischaracterised the topic *Black girls*, the set of sexualised images outputted by Google image for 'Black girls' feeds into a particular false view of Black women. Noble argues that this view is not simply pornographic; it deploys set of visual imagery and stereotypes for representing Black women drawn from this history of the *Jezebel*, *Mammy*, and *Sapphire* images (Noble, 2018, pp. 94–98). She highlights the links to sites advertising *Hot Black Pussy*, alluding to bell hooks' essay on the commodification of Black Women's identity in the media *Selling Hot Pussy* (hooks, 1992).

In *Black Feminist Theory* Patricia Hill Collins introduces the idea of controlling images to make sense of the historically laden stereotypes which shape the way Black women are perceived (Collins, 2000, pp. 76–77). Collins argues that controlling images of Black women function to situate Black women as an Other, to objectify them as a mere object of knowledge, and to present them as deviant category against which the norms of white femininity can be defined. By presenting images that reproduce the Jezebel image in response to 'Black girls', Google images both presents hypersexuality as characteristic of Black women, and situates their sexuality as a deviant form, against which the normal sexuality of white women can be defined.

Collins connects controlling images to beauty norms, arguing that racialised standards of beauty create a social hierarchy within a system that 'elevates whiteness over Blackness' (Collins, 2000, p. 98). This system appears throughout Google image results, both in the search for 'beautiful' which returned exclusively pictures of white women (Noble, 2018, p. 22), and in the results for '(un)professional hair' mentioned above. Associating hairstyles worn by white women with professional roles, and hairstyles worn by Black women with the derogatory label 'unprofessional' supports a racial hierarchy which excludes Black women from professional work, only including them insofar as they approximate the ideals of white beauty. The concern here is not simply that Google image is reproducing historical racist representations, but that by recycling controlling images of Black, Asian, and Latina women it is contributing to the ongoing construction of racial categories, within which women from these

Search Engines, White Ignorance

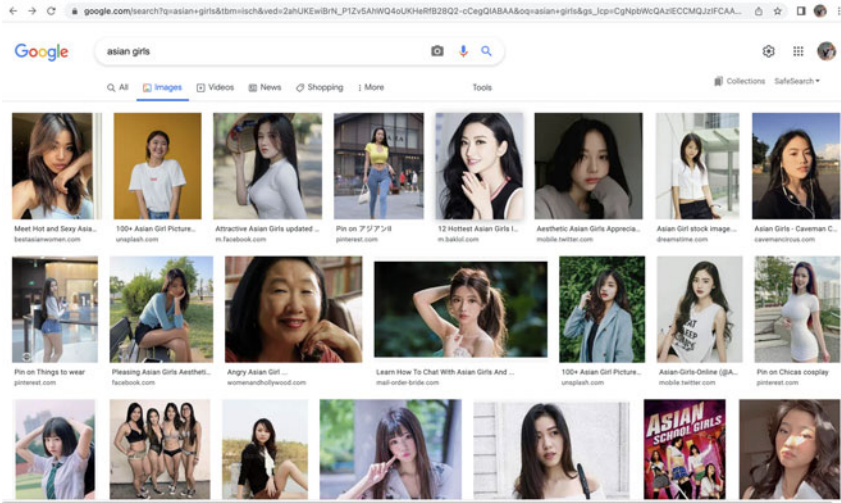


Figure 3. Google images result for ‘Asian girls’ on 28 July 2022.

groups are socially, politically, and economically subordinated (Noble, 2018, p. 84).

Returning to these results in 2022 is salutary. The Google image results for ‘Black girls’ are a neutral collection of headshots, but the results for ‘Asian girls’ (figure 3) includes a large number of sexualised images of young Asian women, many of which – at first pass – appear to deploy the controlling image of the *Lotus Blossom*.¹⁵

As above, the results in other languages continue to replicate the problematic patterns highlighted by Noble. In Italian, the search query ‘ragazze nere’ (figure 4) returns a collection of images of Black women. Although some of these images are fairly neutral, others are highly sexualised and reproduce the Jezebel image. It is worth noting at the time this search was made, the first image link is taken from an explicitly white supremacist site.

Turning to the results for ‘unprofessional hair’ (figure 5), it is worth noting that Google images appears to have been changed to return more diverse images of people in general. Although ‘professional hair’ returns a racially diverse set of pictures of women, ‘unprofessional hair’ is a mix of pictures from articles about the bias in Google’s results and pictures of Black women with natural hair.

¹⁵ On the fetishisation of Asian women and its harms, see Zheng (2016).

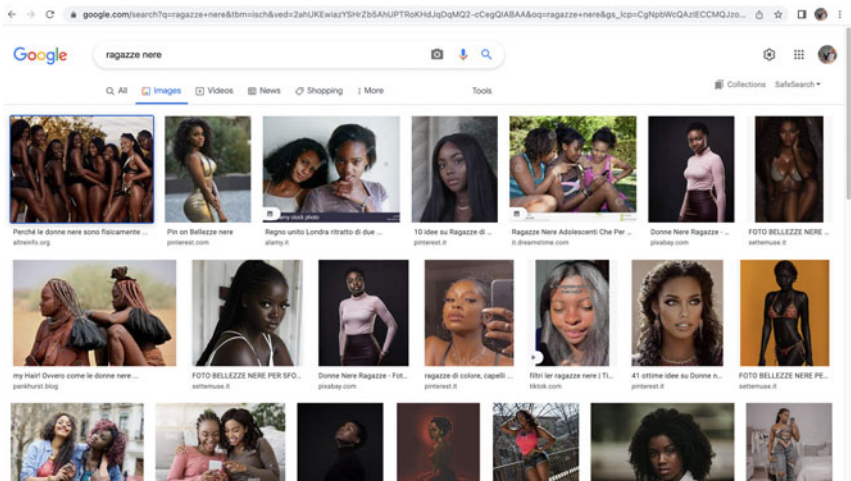


Figure 4. Google images .it results for ‘Ragazze nere’ on 28 July 2022.

A search for ‘cheveux pas professionnel’ on [google.fr](https://www.google.fr/) (figure 6) returns strikingly similar results to those originally highlighted by @BonKamora.

4.3 Autocomplete

Perhaps the most striking example in Noble’s book are the autocomplete results for queries about Black people and Black women. In January of 2013, Noble found that the query ‘why are black people so’ was filled in with the suggestions ‘loud’, ‘athletic’, ‘lazy’, ‘fast’, where ‘why are Black women so’ was filled in with ‘angry’, ‘loud’, ‘mean’, ‘attractive’ (Noble, 2018, pp. 20–21).

I suggest that we think about Google’s autocomplete function in search as providing something in between an automatic text completion function (similar to those found on texting applications), and a recommendation function for queries. When we misspell a word in a query, the text-completion function may be more salient, but in these cases Google search appears to be recommending queries like ‘why are Black women so angry’. The query recommendation function works not by providing information, but by directing users’ inquiries, and shaping their curiosity (Miller and Record, 2017, pp. 1949–50). In many cases these recommendations may be ignored, but plausibly a decent number of people follow them (otherwise Google search would have removed the function). Even without following a recommendation, glancing at autocomplete results can

Search Engines, White Ignorance

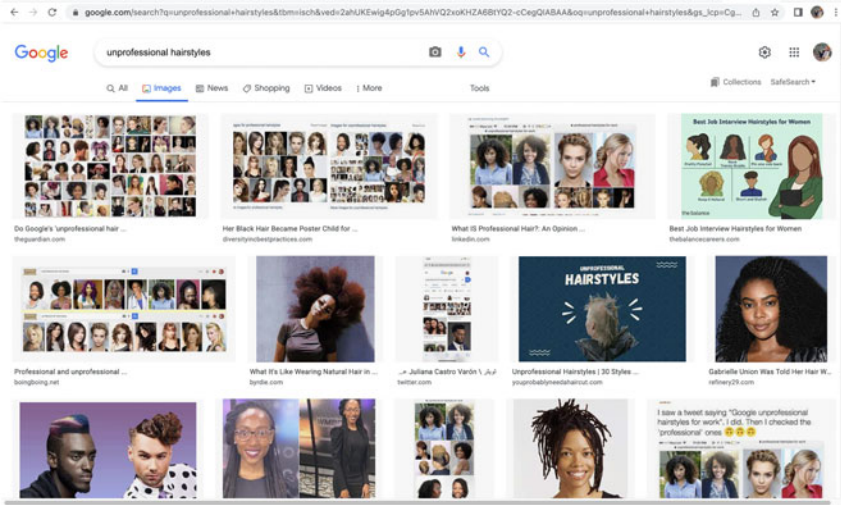


Figure 5. Google images result for ‘unprofessional hair’ on 28 July 2022.

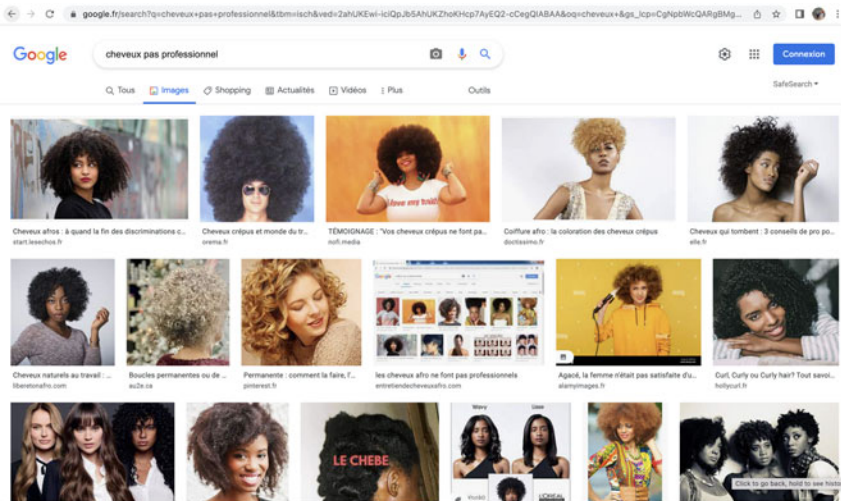


Figure 6. Google.fr image search for ‘cheveux pas professionnel’ on 28 July 2022.

convey a sense of the topics related to one’s query. Exactly how suggestions are generated is a complicated question: plausibly they draw both on data about popular searches in the querist’s area (which also appear in the ‘trending searches’ function), and the use of language models to predict the next word in a string.

Noble's discussion points us toward two problematic features of these autocomplete results.

The first is that these results demonstrate Google's algorithmic conceptualisation of Black people involves negative racial stereotypes. Just as human representations of social groups can become enmeshed with negative characteristics, so too can algorithmic representations – especially when algorithmic representations are derived from data sets produced by humans with implicit biases (Johnson, 2020). The problem is not merely that these stereotyped representations have been produced by a technological system, but that because of the association between algorithmic systems and the epistemic virtue of objectivity (Benjamin, 2019), these racial stereotypes are presented as authoritative and objective. The publication of *Algorithms of Oppression* coincided with other research and investigative journalism that demonstrated algorithmic bias in various important systems – see Boulamwini and Gebru (2018), Angwin *et al.* (2018), Dastin (2018)¹⁶ – and I take it that this idea has been important to the public uptake of the book.

The second problematic feature of these results concerns whose interests are represented by these questions. An important theme in both feminist philosophy of science and Black feminism is that questions are not neutral: a question may be more pressing for one or other group, and the way in which a question is framed may prevent some information from being shared (Noble, 2018, p. 31, quoting Harding, 1987; see also Cooper, 1898; Longino, 1990; Crenshaw, 1991; Anderson, 2004; Haslanger, 2016). The questions being suggested by the autocomplete results for 'why are Black people so' transparently do not promote the interests of Black people. Noble argues that Google search results systematically promote the interests of capital, particularly advertising companies:

Search results reflect the values and norms of the search company's commercial partners and advertisers and often reflect our lowest and most demeaning beliefs, because these ideas circulate so freely and so often that they are normalised and extremely profitable. [...] Google's monopoly status, coupled with its algorithmic practices of biasing information toward the interests of the neoliberal capital and social elites in the United States, has resulted in a provision of information that purports to be credible but is actually a reflection of advertising interests. (Noble, 2018, pp. 35–36)

¹⁶ See Friedman and Nissenbaum (1996) for an important precursor to this work.

If the questions which Google search is asking – both explicitly via its recommendation function, and implicitly via the interpolation of subject matters – do not promote the interests of racialised minority groups, then there may be a good case for developing minority-interest search engines which build the interests of minority groups into the technology from the start (Noble, 2018, pp. 150–51).¹⁷

At the time of writing, Google search appears to have tried to have tried to fix autocomplete results through a combination of blocking the autocomplete function for some key words, and aggressively filtering out problematic suggestions.¹⁸ Typing ‘why are Black women so’ into Google Search returns no suggestions, and ‘why are Black people so’ prompts ‘so good at running’ and ‘tall’ (figure 7).

The shorter query ‘why are Black people’ prompts the suggestion ‘why are Black people attacking Asians’, which appears to reflect a widespread narrative that Black men were responsible for a rise in attacks on Asian-Americans during the COVID-19 pandemic (figure 8).¹⁹

Turning to Google’s competitors, the picture gets worse.²⁰ In Bing, the query ‘why are Black women so’ returns ‘masculine’, ‘sassy’, and ‘obese’ (figure 9). Whereas Google image’s results seemed to be in the grip of the Jezebel image, Bing’s algorithmic conceptualisation manifests the Mammy image (Collins, 2000, pp. 80–88).

Yahoo search (which uses Bing’s search algorithms) has similar results. The query ‘why are Black people so’ suggests ‘ugly’, ‘arrogant’, ‘rude’, and ‘racist’ (figure 10).

While search engines may no longer be returning the unfiltered results of their autocomplete algorithms to users, these results

¹⁷ In an important early paper on the politics search, Introna and Nissenbaum make a related argument against commercial search. They suggest that Pareto’s law applies to search queries, meaning that 80% of queries are directed toward 20% of sites, whereas the remaining 20% of queries seek the other 80% of sites. They suggest that whereas a commercial search engine will cater to majority interests, developing a product designed to find the 20% of most popular sites, there is a public interest in maintaining links to the remaining 80% of sites, in order to cultivate a healthy public sphere (Introna and Nissenbaum, 2000). See Noble’s discussion of the enclosure of the online public sphere (2018, pp. 50–51).

¹⁸ The initial changes seem to have been made in 2016, but were not fully successful since problematic results were still showing up in 2018 <https://www.wired.com/story/google-autocomplete-vile-suggestions/>.

¹⁹ See <https://www.vox.com/22321234/black-asian-american-tensions-solidarity-history>.

²⁰ See <https://www.theverge.com/2018/10/10/17959328/bing-yahoo-offensive-search-suggestions-racism-antisemitism>.

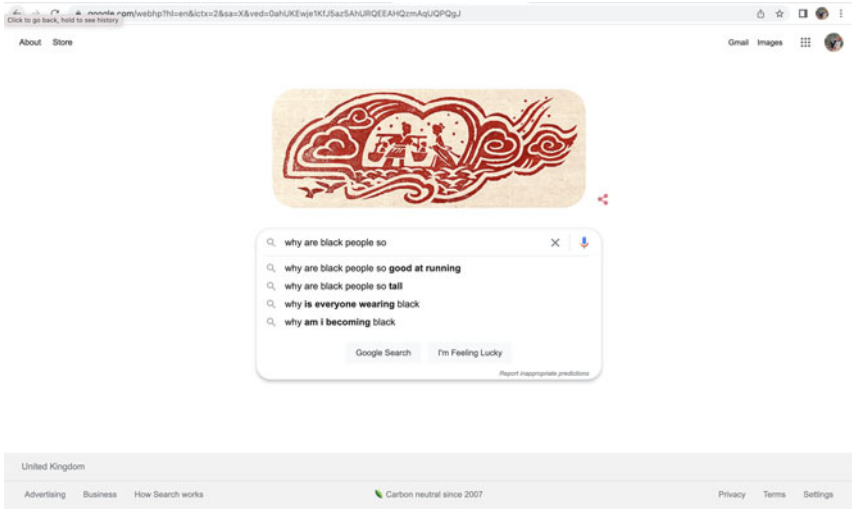


Figure 7. Google autocomplete suggestions for ‘why are black people so’ on 4 August 2022.

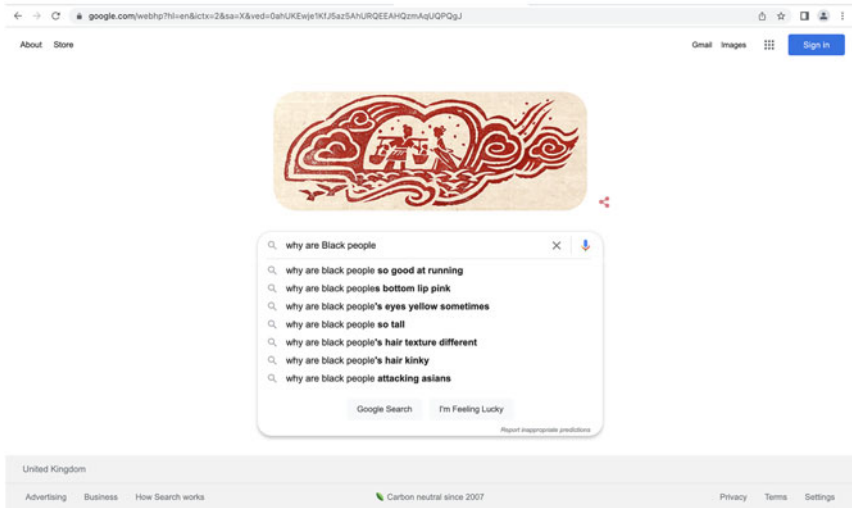


Figure 8. Google autocomplete suggestions for ‘why are black people’ on 4 August 2022.

suggest that their moderation practices are insufficiently aligned to the kinds of racialised harms which can be caused by autocomplete results (see Frost-Arnold, 2023, Ch. 2).

Search Engines, White Ignorance

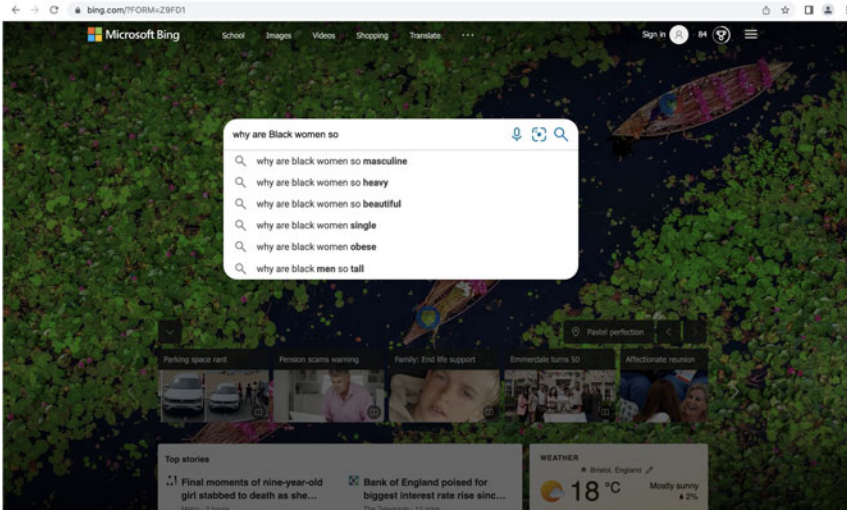


Figure 9. Bing search autocomplete suggestions for ‘why are Black women so’ on 4 August 2022.

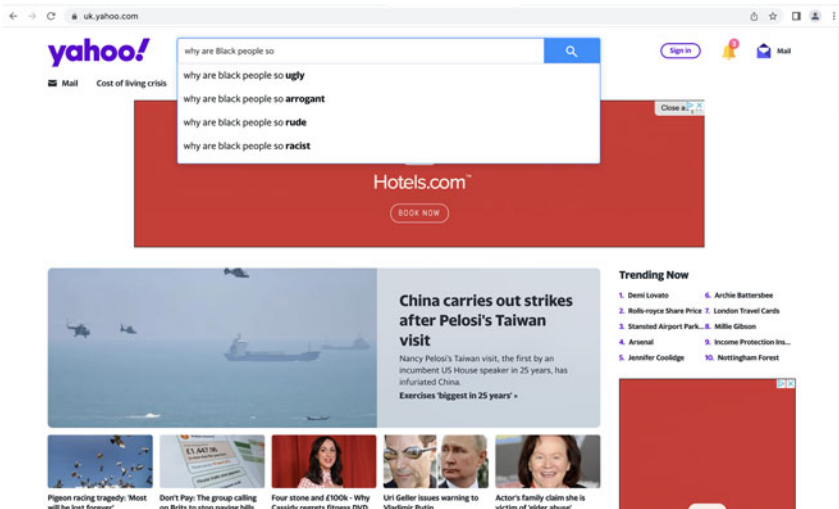


Figure 10. Yahoo search autocomplete suggestions for ‘why are Black people so’ on 4 August 2022.

5. Diagnosing the Problem

The pattern of results that Noble theorises puts considerable pressure on Google’s self-presentation as an objective and neutral service

which provides relevant information that can assuage the ignorance of citizens of a modern multicultural democracy. There are several different ways in which we might think about these results.

The most sympathetic diagnosis is that these patterns of racist and misogynistic results are simply *glitches*, random errors in a system which otherwise provides a reliable navigation tool (see Benjamin, 2019, Ch. 2). I don't think this diagnosis is worth much time; the pattern of problematic results appears to be robust both across different queries associated with racialised groups, and across time as the same kinds of problems re-emerge despite local fixes.

Another is to blame the users of Google search. In a blog post responding to the anti-Semitic site JewWatch appearing as the top-ranked result for the query 'Jew' in 2004, The Google Team claimed:

If you use Google to search for "Judaism," "Jewish" or "Jewish people," the results are informative and relevant. So why is a search for "Jew" different? One reason is that the word "Jew" is often used in an anti-Semitic context. [...] Someone searching for information on Jewish people would be more likely to enter terms like "Judaism," "Jewish people," or "Jews" than the single word "Jew." (Google 2004)

It is difficult to reconstruct the position of this blog, but one thing that The Google Team might be suggesting is that because the majority of users who enter 'Jew' – as opposed to 'Jewish', or 'Judaism' – are interested in finding anti-Semitic sites, the ranking algorithms have pushed anti-Semitic sites up the rankings for 'Jew' in order to meet the presumed needs of users. This diagnosis allows Google to continue to tout the semi-magical power of its search algorithms, whilst avoiding the impression of endorsing any of the sites that they output, and pinning the responsibility for problematic results onto a presumed anti-Semitic minority. It is pretty clear that some unexpected results are due to user behaviour: Google Bombing can lead to irrelevant links having very high rankings, and search engine optimisation is in effect a manipulation of rankings by users. This doesn't mean that sole responsibility lies with users: Google has corporately made a decision to allow the ranking of search results to be determined by a combination of user behaviour, explicit advertising, and search engine optimisation, in order to provide an efficient advertising platform whilst giving (the majority of) users useful results.

A third explanation is that Google is racist, either in the sense that its employees have deliberately produced a technology which causes a racialised pattern of harms, or in the sense that it is a

structurally racist organisation. We shouldn't write off the role of explicit racism in producing these problematic results. Many cloaked websites are produced by self-declared white supremacists, who are innovation opportunists, exploiting the affordances of the latest technology (Daniels, 2018). Although things are improving, African Americans remain under-represented in Google's workforce: in 2015 2% of Google's workforce was Black (Lee, 2016), and by 2023 the percentage of the US workforce which are Black was 5.6% (Google, 2023) (on US census classifications around 12.5% of the US population is African American. Over the same period, women increased from 20% of the workforce to 33% of the worldwide workforce). Support for James Damore's memo *Google's Ideological Echo Chamber* (Damore, 2017) – which makes the case for biologically based differences in men and women's psychology – within the tech sector raises questions about the gender and race politics of Silicon Valley (Noble and Roberts, 2019). There is much more to be said about the politics of Google search the organisation, and its employees. However, given that Google search is a technological system, we need the conceptual tools for thinking about it as a racist technology.

5.1 Google Search and White Ignorance

The suggestion that I want to develop is that Google search – and other search engines – produce patterns of problematic and false results because they are racist socio-technological systems, which are congruent with the wider institution of white ignorance which enacts an inverted epistemology in the service of a white supremacist social order.²¹ To be clear, the claim is not that search engines constitute the whole of the institution of white ignorance but that we should see the pattern of problematic results theorised by Noble as the manifestation of a system of white ignorance.

What is a racist technology? Drawing on work in science and technology studies, Liao and Huebner (2020) carve out a category of oppressive objects. They are interested in the idea that objects are oppressive in the sense that they are *congruent* with oppressive systems. They understand congruence with an oppressive system as having three conditions: i) the object is biased in the same direction as the social system, ii) the object is causally embedded within the

²¹ Search engines are also plausibly implicated in other systems of ignorance production, see footnote 3.

oppressive system, and iii) the object is bi-directionally embedded within an oppressive system, both reflecting the kinds of oppression involved in that system, and guiding and constraining psychological processes and social practices which reproduce it (2020, p. 9). Noble's analysis of Google search results shows that the outputs of Google search have historically tended to produce representations that normalise and justify racial hierarchies, getting us directional bias. The causal embedding of Google search within an oppressive social system is established by the fact that both its users and the websites that it classifies are the products of a society characterised by racist oppression. And the bi-directional links between oppressive social practices and Google search are suggested by the way in which its results are both determined by patterns of racist queries (as with the autocorrect results), and go on to guide racist social practices, including in extreme cases, guiding white supremacist violence (see Noble's discussion of the Dylan Roof case (2018, pp. 110–18)).

Within philosophical discussions, the notion of white ignorance has become associated with factual ignorance about questions about the history of colonialism, which might make one worry whether it is a sufficiently general concept to characterise the diversity of problematic search results. I think that Mills's focus on historical ignorance in *White Ignorance* stems from his diagnosis of the racial politics of the 1990s and 2000s, which combined a commitment to 'colourblind' politics with a refusal to engage with questions about the effects of colonial history, with the effect that existing racial inequalities were perpetuated. When we take into account his wider discussions of inverted epistemologies, including in earlier historical periods, it becomes clear that for Mills, white ignorance works through a diversity of mechanisms, including the propagation of dehumanising conceptual systems, and the norming of spaces and bodies (1997, pp. 41–62).

Mills is clear that the Racial Contract is a dynamic social phenomenon, which adapts to the changing political situations. There are three important features of the way in which the Racial Contract has worked itself into our contemporary technological systems.

The first is the emergence of a form of technologically enabled colour blindness, which Ruha Benjamin calls the *New Jim Code* (Benjamin, 2019). This ideology invests technological systems with the values of neutrality, objectivity, and benevolence, enabling us to see technological systems as the harbingers of progress, which operate outwith the messy human realities of social injustice. The effect of this ideology is to obscure the discriminatory designs which are built into these systems, to hide the ways in which those

systems reproduce existing inequalities, and (perversely) to present technology as the solution to problems of racial injustice. As more and more state systems – welfare, the identification of criminals, child safety – are entrusted to automated systems (Eubanks, 2018), these patterns of harms are both more widespread, and increasingly obfuscated by the authority granted to algorithmic systems.

The second is that technological systems provide an unusual amount of tools for self-conscious white supremacists. Jessie Daniels has argued that white supremacists have long been what she calls innovation opportunists, taking advantages of the affordances of new communication technologies to organise and proselytise (Daniels, 2009, 2018), and the communication possibilities of social media sites, algorithmic recommendation systems, message boards, and video-sharing sites have given contemporary white supremacists unprecedented tools for organising (Marwick and Lewis, 2017).

The third is that technological systems have enabled new forms of racialised economic exploitation. Mills claims that the origins of the Racial Contract lie in the economic exploitation of Blacks, through a combination of dispossession and plantation slavery (Mills, 1997, pp. 32–40), which suggests that he endorses the view that racial classification is functional for the economic system of capitalism (Bright, Gabriel, O'Connor, and Táiwò, 2022). If we look at the new kinds of markets which are enabled by contemporary technology, we see persistent exploitation along racial lines, meaning that Racial Capitalism is a helpful frame for thinking about technological systems (Cottom, 2020). Noble discusses how Google search's advertising market has commodified identity markers for racialised groups, marking Black women as hyper sexualised (2018, pp. 92–104), and the way it obfuscates the labour performed by its users (2018, pp. 56–58). Online labour platforms have created markets segmented along racial lines, including gig workers in the global north – a workforce which tends to be made up of racialised minorities and migrants (Cottom, 2020; Gebriel, forthcoming) – and platform workers and call centre workers in the global south (Gray and Suri, 2019; Roberts, 2019). In each of these cases, racialised exploitation is obfuscated by the logic of opacity (Roberts, 2019) of technologically mediated market, which makes it difficult to see what work is being done, and who is operating in a particular market.

The suggestion that search engines contribute to a system of white ignorance is both more and less radical than it might seem. It is radical in the sense that it accepts the existence of a social system of white ignorance, which produces information, guides inquiry, and structures epistemic resources in a way that promotes and maintains

the system of global white supremacy. It is less radical than it might seem because once we have accepted the existence of this ignorance-producing institution, it should be unsurprising that any particular part of our epistemic architecture contributes to this institution. Once we have recognised that search engines are reproducing a wider political institution, we are in a better position to see how these problematic results might be ameliorated. Merely fixing the results piecemeal, while continuing to trumpet Search's credentials as a source of information about ethnic minorities does not scratch the surface of what is fundamentally a political problem. Recognising that the source of these problematic results is a political institution means that we need to realise that 'an App won't save us' (Noble, 2018, p. 165; see Benjamin, 2019, on design thinking), and we need to engage with political questions.

6. Towards a Social Epistemology of Technology

In this paper I've argued that Noble's diagnosis of patterns of problematic search results supports the idea that Google search – and it would appear, other search engines – are part of a wider system of white ignorance, which produces miscognitions in service of the Racial Contract. I want to draw out two wider lessons from this discussion.

The first concerns the importance of *social* epistemology of technology. Many algorithmic and technological systems function as knowledge-generating systems, aiming to produce knowledge in a similar way to instruments like thermometers, watches and rulers. This means that it is tempting to assess these systems primarily by thinking about the outputs they produce for their users. This framing of the epistemology of technology will highlight many of the ways in which these systems can go wrong, but it will miss out important ways in which technological systems contribute to institutions of ignorance production. Thinking of a search engine as a tool for individual inquirers, a link to a white supremacist site on a query about racial justice is an irrelevance on par with a link to a site with out-of-date statistics. Thinking of a search engine as operating in a context of systems of ignorance production, we can see that a link to a white supremacist site is congruent with a system of ignorance production, and ought to be of much greater concern than a link to out-of-date statistics. Put besides the fact – familiar in science and technology studies – that many putatively automated technological systems are in reality assemblages of technological and social processes, we should be thinking about the epistemology of technology as social in two senses: the

sense that technological systems include social processes, and the sense that they are part of wider social practices.

The second concerns the politics of the social epistemology of technology. If the pattern of problematic results produced by Google search is not the fault of a technological glitch, but a manifestation of a social institution that produces miscognition, then the ultimate remedy is not simply to 'fix' the technological system, but to dismantle the institution of white ignorance. Thinking about google search is a helpful way of diagnosing wider problematic social practices (Flowers, 2019). Insofar as search engines don't merely reflect but are also embedded within the social system of white ignorance, these practices need to be disrupted. Long term, the underlying financial model of Google search is based on producing a majoritarian technology which produces a product which is of use to the majority of people, while being useful to advertisers. The interests of advertisers and users are likely to be in conflict, but ultimately a commercial search engine will be incentivised to prioritise its financial interests over the well-being of its users. There's a harm reduction possibility here: better content moderation. To deal with the kinds of racialised harms produced by an epistemology of ignorance – especially as that system evolves over time, and operates in a specific social context – we cannot rely on automated systems, and will need to rely on human content moderation which is given sufficient space to develop the skills of identifying racialised harms (see Frost-Arnold, 2023, Ch. 2).

Postscript: Large Language Models in Search

When I started writing this paper in 2022, the idea of implementing Large Language Models (LLMs) in publicly available search engines was just a possibility (see Metzler, Tay, Bahri, and Najork, 2021; Shah and Bender, 2022).²² As I finish the paper, several search engines have integrated LLMs into their search functions and (with more or less success) released question-answering chatbots based on LLMs. While it remains to be seen what problems will be displayed by this technology, there is a wealth of evidence for racial and gender bias in the outputs of LLMs (Brown *et al.* [OpenAI], 2020; Abid *et al.*, 2021; Bender *et al.*, 2021), and there is evidence that Chat-GPT3 avoided producing problematic outputs by employing a large underpaid workforce in Kenya to label problematic results (Perrigo, 2023).

²² For a non-technical overview of large language models, see Levinstein (2023).

Acknowledgements

Thanks to Natalie Alana Ashton, Zara Bain, Michael Barany, Karen Frost-Arnold, Nadja El Kassar, Fintan Mallory, Jared Millson, Alessandra Tanesini, Dennis Whitcomb, an anonymous reviewer for this journal, and audiences in Bologna and London. Special thanks to Andrew Peet for comments in London. This research has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 818633).

References

- Abubakar Abid, Maheen Farooqi, and James Zou, 'Persistent Anti-Muslim Bias in Large Language Models', in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (2021), 298–306.
- Elizabeth Anderson, 'Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce', *Hypatia*, 19:1 (2004), 1–24.
- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, 'Machine Bias', *ProPublica*, 23 May 2016, accessed 9 April 2024 at <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Zara Bain, 'Is There Such a Thing as "White Ignorance" in British Education?', *Ethics and Education*, 13:1 (2018), 4–21.
- John Perry Barlow, *A Declaration of the Independence of Cyberspace* (1996), accessed 9 April 2024 at <https://www.eff.org/cyberspace-independence>.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (2021), 610–23.
- Ruha Benjamin, *Race After Technology: Abolitionist Tools for the New Jim Code* (New York: Polity, 2019).
- Liam Kofi Bright, Nathan Gabriel, Cailin O'Connor, and Olufemi Taiwo, 'On the Stability of Racial Capitalism', 24 May 2022.
- Sergei Brin and Larry Page, 'The Anatomy of a Large-Scale Hypertextual Web Search Engine', *Computer Networks and ISDN Systems*, 30:1–7 (1998), 107–17.
- A. Broder, 'A Taxonomy of Web Search', *ACM Sigir Forum*, 36:2 (2002), 3–10.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, ... (Open AI), 'Language Models are Few-Shot Learners', *Advances in Neural Information Processing Systems*, 33 (2020), 1877–1901.

Search Engines, White Ignorance

- Joy Buolamwini and Timnit Gebru, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', in *Conference on Fairness, Accountability and Transparency* (2018), 77–91.
- Holly Case, *The Age of Questions: Or, A First Attempt at an Aggregate History of the Eastern, Social, Woman, American, Jewish, Polish, Bullion, Tuberculosis, and Many Other Questions over the Nineteenth Century, and Beyond* (Princeton: Princeton University Press, 2018).
- Anna Julia Cooper, *A Voice from the South* (Oxford: Oxford University Press, [1898] 2018)
- Patricia Hill Collins, *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment* (Abingdon: Routledge, 2000).
- Tressie McMillan Cottom, 'Where Platform Capitalism and Racial Capitalism Meet: The Sociology of Race and Racism in the Digital Society', *Sociology of Race and Ethnicity*, 6:4 (2020), 441–49.
- Kimberle Williams Crenshaw, 'Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color', *Stanford Law Review*, 43:6 (1991), 1241–99.
- Justin D'Ambrosio and Brian Hedden, 'Multidimensional Adjectives', *Australasian Journal of Philosophy*, (forthcoming).
- James Damore, *Google's Ideological Echo Chamber* (2017), accessed at <https://s3.documentcloud.org/documents/3914586/Googles-Ideological-Echo-Chamber.pdf>.
- Jessie Daniels, *Cyber Racism: White Supremacy Online and the New Attack on Civil Rights* (Lanham, Maryland: Rowman & Littlefield, 2009).
- Jessie Daniels, 'Race and Racism in Internet Studies: A Review and Critique', *New Media & Society*, 15:5 (2013), 695–719.
- Jessie Daniels, "'My Brain Database Doesn't See Skin Color": Color-Blind Racism in the Technology Industry and in Theorizing the Web', *American Behavioral Scientist*, 59:11 (2015), 1377–93.
- Jessie Daniels, 'The Algorithmic Rise of the "Alt-Right"', *Contexts*, 17:1 (2018), 60–65.
- J. Dastin, 'Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women', *Reuters*, 11 October 2018, accessed 9 April 2024 at <https://www.reuters.com/article/idUSKCN1MK0AG/>.
- Martin Feuz, Matthew Fuller, and Felix Stalder, 'Personal Web Searching in the Age of Semantic Capitalism: Diagnosing the Mechanisms of Personalisation', *First Monday*, 16:2–7 (2011).
- Janathan C. Flowers, 'Rethinking Algorithmic Bias through Phenomenology and Pragmatism', *Computer Ethics – Philosophical Enquiry (CEPE) proceedings*, 2019:1 (2019), 27 pp.
- Batya Friedman and Helen Nissenbaum, 'Bias in Computer Systems', *ACM Transactions on Information Systems (TOIS)*, 14:3 (1996), 330–47.
- Karen Frost-Arnold, *Who Should We Be Online: A Social Epistemology for the Internet* (New York: Oxford University Press, 2023).

Joshua Habgood-Coote

- Dalia Gebrial, 'Racial Platform Capitalism: Empire, Migration and the Making of Uber in London', *Environment and Planning A: Economy and Space*, (forthcoming).
- Google, 'Google Diversity Annual Report 2023', accessed 9 April 2024 at https://static.googleusercontent.com/media/about.google/en//belonging/diversity-annual-report/2023/static/pdfs/google_2023_diversity_annual_report.pdf?cachebust=2943cac.
- Mary L. Gray and Siddarth Suri, *Ghost Work: How to stop Silicon Valley from Building a New Global Underclass* (New York: Mariner, 2019).
- Jeroen Groenendijk and Martin Stokhof, *Studies on the Semantics of Questions and the Pragmatics of Answers* (PhD thesis, Amsterdam, 1984).
- Hannah Gunn and Michael P. Lynch, 'Googling', in David Coady and James Chase (eds.), *Routledge Handbook of Applied Epistemology* (New York: Routledge, 2019), 41–53.
- Joshua Habgood-Coote, 'Group Inquiry', *Erkenntnis*, 87:3 (2022a), 1099–1123.
- Joshua Habgood-Coote, 'Knowing More (about Questions)', *Synthese*, 200:1 (2022b), 1–23.
- Sally Haslanger, 'What Is a (Social) Structural Explanation?', *Philosophical Studies*, 173:1 (2016), 113–30.
- bell hooks, 'Selling Hot Pussy', in *Black Looks: Race and Representation* (Boston MA: South End Press, 1992), 61–78.
- Tim Hwang, *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (New York: FSG originals, 2020).
- Lucas D. Introna and Helen Nissenbaum, 'Shaping the Web: Why the Politics of Search Engines Matters', *The Information Society*, 16:3 (2000), 169–85.
- Gabrielle M. Johnson, 'Algorithmic Bias: On the Implicit Biases of Social Technology', *Synthese*, 198:10 (2020), 9941–61.
- Phil Jones, *Work Without the Worker: Labour in the Age of Platform Capitalism* (London: Verso, 2021).
- Daniel N. Joudrey, Arlene G. Taylor, and Katherine M. Wisser, *The Organization of Information*, 4th ed. (Santa Barbara, California: Libraries Unlimited, 2018).
- Nadja El Kassar, 'What Ignorance Really Is. Examining the Foundations of Epistemology of Ignorance', *Social Epistemology*, 32:5 (2018), 300–10.
- Quill R. Kukla, *City Living: How Urban Spaces and Urban Dwellers Make One Another* (Oxford: Oxford University Press, 2021).
- Rae Langton and Caroline West, 'Scorekeeping in a Pornographic Language Game', *Australasian Journal of Philosophy*, 77:3 (1999), 303–19.
- Nancy Lee, 'Focusing on Diversity', 30 June 2016, accessed 9 April 2024 at <https://blog.google/outreach-initiatives/diversity/focusing-on-diversity30/>.
- Ben Levinstein, *A Conceptual Guide to Transformers* (2023), accessed 9 April 2024 at <https://benlevinstein.substack.com/p/a-conceptual-guide-to-transformers?sd=pf>.

Search Engines, White Ignorance

- Steven Levy, 'Exclusive: How Google's Algorithm Rules the Web', *Wired*, 2010, accessed 9 April 2024 at https://web.archive.org/web/20110416062117/http://www.wired.com/magazine/2010/02/ff_google_algorithm/.
- Shen-yi Liao and Bryce Huebner, 'Oppressive Things', *Philosophy and Phenomenological Research*, 103:1 (2020), 92–113.
- Helen E. Longino, *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry* (Princeton: Princeton University Press, 1990).
- Marie MacDonald, 'There's a Massive Scam Hiding Behind Google's Search Results', *Wired*, 2 July 2020, accessed 9 April 2024 at <https://www.wired.co.uk/article/gaming-google-crowdwork>.
- Annette Martín, 'What is White Ignorance?', *Philosophical Quarterly*, 71:4 (2021), 864–85.
- Alice E. Marwick and Rebecca Lewis, 'Media Manipulation and Disinformation Online', *Data and Society*, (2017), accessed 9 April 2024 at <https://datasociety.net/library/media-manipulation-and-disinfo-online/>.
- George Masterton, Erik J. Olsson, and Staffan Angere, 'Linking as Voting: How the Condorcet Jury Theorem in Political Science is Relevant to Webometrics', *Scientometrics*, 106:3 (2016), 945–66.
- D. Metzler, Y. Tay, D. Bahri, and M. Najork, 'Rethinking Search: Making Experts out of Dilettantes', *ArXiv*, (2021).
- Boaz Miller and Isaac Record, 'Responsible Epistemic Technologies: A Social-Epistemological Analysis of Autocompleted Web Search', *New Media and Society*, 19:12 (2017), 1945–63.
- Charles W. Mills, *The Racial Contract* (Ithaca: Cornell University Press, 1997).
- Charles W. Mills, 'Global White Ignorance', in Matthias Gross and Linda McGooney (eds.), *Routledge International Handbook of Ignorance Studies* (Abingdon: Routledge, 2015), 217–27.
- Charles W. Mills, 'White Ignorance', in *Black Rights/White Wrongs: The Critique of Racial Liberalism* (New York: Oxford University Press, 2017), 49–71.
- Jess Munton, 'Answering Machines: How to (Epistemically) Evaluate a Search Engine', *Inquiry*, (forthcoming).
- Annalee Newitz, 'The Secret Lives of Google Raters', *Ars Technica*, 27 April 2017, accessed 9 April 2024 at <https://arstechnica.com/features/2017/04/the-secret-lives-of-google-raters/>.
- Safiya U. Noble, 'Google Search: Hyper-Visibility as a Means of Rendering Black Women and Girls Invisible', *InVisible Culture*, 19 (2013).
- Safiya U. Noble, *Algorithms of Oppression* (New York: New York University Press, 2018).
- Safiya U. Noble and Sarah Roberts, 'Technological Elites, the Meritocracy, and Postracial Myths in Silicon Valley', in Roopali Mukherjee, Sarah Banet-Weiser, and Herman Gray (eds.), *Racism Postrace* (Durham, NC: Duke University Press, 2019), 113–30.
- Carole Pateman, *The Sexual Contract* (Cambridge: Polity Press, 1988).

Joshua Habgood-Coote

- Carlotta Pavese, 'Know-How and Gradability', *Philosophical Review*, 126:3 (2017), 345–83.
- Billy Perrigo, 'Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic', *Time*, 18 January 2023, accessed 9 April 2024 at <https://time.com/6247678/openai-chatgpt-kenya-workers/>.
- Timothy Recuber, 'Digital Discourse Analysis: Finding Meaning in Small Online Spaces', in J. Daniels, and K. Gregory (eds.), *Digital Sociologies* (Bristol: Policy Press, 2016).
- Jessica Pepp, Eliot Michaelson, Rachel Sterken, 'Relevance-Based Knowledge Resistance', in Jesper Strömbäck, Åsa Wikforss, Kathrin Glüer, Torun Lindholm, and Henrik Oscarsson (eds.), *Knowledge Resistance in High-Choice Information Environments* (London: Routledge, 2022).
- Craige Roberts, 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics', *Semantics and Pragmatics*, 5 (1996), 1–69.
- Sarah Roberts, *Behind the Screen* (Yale: Yale University Press, 2019).
- Chirag Shah and Emily M. Bender, 'Situating Search', in *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (2022), 221–32.
- Thomas W. Simpson, 'Evaluating Google as an Epistemic Tool', *Metaphilosophy*, 43:4 (2012), 426–45.
- Jason Stanley, *How Propaganda Works* (Princeton: Princeton University Press, 2015).
- Zoltán Gendler Szabó, 'Finding the Question', *Philosophical Studies*, 174:3 (2017), 779–86.
- The Google Team, 'An Explanation of Our Search Results', [Google.com](https://web.archive.org/web/20040610171859/https://www.google.com/explanation.html) (2004), accessed 9 April 2024 at <https://web.archive.org/web/20040610171859/https://www.google.com/explanation.html>.
- Stephen Yablo, *Aboutness* (Oxford: Princeton University Press, 2014).
- Leon Yin and Aaron Sankin, 'Google Ad Portal equated "Black Girls" with Porn', *The Markup*, 23 July 2020, accessed April 2024 at <https://themarkup.org/google-the-giant/2020/07/23/google-advertising-keywords-black-girls>.
- Robin Zheng, 'Why Yellow Fever Isn't Flattering: A Case Against Racial Fetishes', *Journal of the American Philosophical Association*, 2:3 (2016), 400–19.
- Shoshanna Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (London: Profile books, 2018).

JOSHUA HABGOOD-COOTE (j.habgood-coote@leeds.ac.uk) is a lecturer in Philosophy of Language at the school of Philosophy, Religion and History of Science at the University of Leeds. He works on Epistemology, Philosophy of Language, and the Philosophy of Technology.