

Search and Representation

(situating Algorithms of Oppression by Safiya Umoja Noble)

Munongedzi Mabhoko Clarkson University mabhokm@clarkson.edu

Treat search not as a helpful index but as a form of soft governance. A results page functions like a miniature policy decision that defines who is visible, assigns authority, what counts as knowledge, and which narratives are granted credibility or which stories about people become thinkable. The popular story that search is a neutral conduit survives because the interface is clean and the page is free to use. They ask us to look at the economics and curation underneath the surface, where “relevance” is produced inside ad markets and partnerships, and where the first screen becomes an auctioned gateway to meaning. In that world, neutrality is an appearance sustained by commercial choices, and users are trained to accept the page as if it had been assembled by a kindly librarian rather than by a revenue machine. The most useful way to read this book is not as a catalogue of bias incidents but as a theory of how search governs representation. The opening move in the text is to puncture the notion of neutrality by situating search in an economy of ranking, advertising, and optimization. That economy does not simply reflect social life, it edits it. Search is driven by market forces and paid advertising, so the question to ask when we look at a page of results is not only whether the links are relevant but for whom that relevance has been engineered.

Once neutrality is set aside, the analysis turns to what is called digital redlining. The name echoes a familiar urban practice in which literal lines on maps excluded communities from credit and investment. The instruments differ, but the effect is similar. Authority and opportunity flow to some neighborhoods of knowledge, while other neighborhoods are made difficult to find or are represented primarily through degrading tropes. The shift from paper maps to algorithmic lists hides the line-drawing from public view and places it behind proprietary ranking. The lines are still there. They are inscribed through data licensing, ad auctions, and optimization signals that convert historical prejudice into “user interest.” To reinterpret the opening moves of this argument, begin with classification. Classification is never merely technical. It is cultural power that authorizes an order of the world. When a query about professional hair surfaces images that associate “unprofessional” with Black women’s natural styles, the system is not simply echoing the web. It is rehearsing a long media history that marked Black appearance as deviant, and it is doing so at the very moment many people treat search output as a common reference. Pornographic returns on simple identity terms work the same way. They recast a broad set of people as objects of sexual curiosity and make that portrayal feel like a public consensus because it appears at the top of the page. The page is profitable. It is also pedagogical. It teaches users who belongs where in the social imagination.

These examples matter less as gotchas than as evidence that ranking has a representational logic. It rewards what sells attention and normalizes what advertisers find comfortable. The lesson is that the harms at stake are not only factual errors. They are patterned portrayals that either present people as stereotypes or erase them as experts. The deeper claim is that classification power becomes narrative power. Search does not just retrieve pages. It assembles mini narratives through ordering, snippets, and repetition of similar sources, so that certain explanations appear corroborated by quantity and position. The commercial motive seeks

engagement and conversion, the government motive seeks prediction and control. Both require large, labeled populations and both are served by infrastructures that sort people by alleged risk, desirability, or credibility. In this alignment, search is not peripheral. It is a social control surface through which populations are represented to themselves and to one another. To consider this carefully is to recognize that the harms described in the book are not only incidents of bad data. They are the predictable outcomes of a political economy that monetizes attention and privileges frictionless curation over civic obligations.

One way to test this interpretation is to step outside classical “search” and examine a neighboring domain that also ranks people. Studies of dating platforms for Black women reveal systems that actively shape who is visible and who is ignored. Sorting logics and engagement-based feeds decide which faces appear and in what order, with opacity that prevents users from understanding or contesting the rules. Even when users pay to boost their visibility, whiteness remains the normative baseline, which shows that monetization by itself does not dissolve the underlying hierarchy. These findings are not about romance alone. They expose how ranking regimes convert social bias into an interface that looks like personal choice. The architecture of desirability on these apps resembles the architecture of authority on results pages. The point of importing that study is not to claim that dating and search are the same. It is to show a shared grammar. Both are gatekeepers that convert collective behavior and business goals into the experience of who is available and credible. Both harvest clicks to decide what counts as desirable or authoritative, both teach users to read social order as if it were the outcome of their own selections.

A structural approach is not complete if it stops at race and gender. One of the studies you attached argues that language and religion shape power in the digital sphere yet are often underexamined in mainstream intersectional work. Ever since early internet protocols privileged English and Roman characters, participation and design authority have been unevenly distributed. Even today, customer service workers are coached to present whiteness in voice and name. Large language models amplify these imbalances by baking English defaults into the production of meaning. Religion, too, is a system that organizes belonging and exclusion and has historically underwritten categories of the human that travel into data work. Attending these layers provides a more complete map of how search can misrecognize people long before they arrive at a query.

The text’s most provocative case work concerns two representational patterns that thrive in commercial ranking. The first is sexualization. Identity terms and innocuous searches for girls and women are routed to pornography. The second is securitization. Crime-adjacent queries lead users toward frames that present racialized threat as an obvious fact. Both patterns emerge from the same infrastructure. They monetize desire and fear while outsourcing responsibility to what users click. That perspective pulls focus from individual intentions to the institutional logics that repeat themselves through interfaces and metrics. It foregrounds interlocking hierarchies of race, gender, class, ability, and sexuality and asks how technology sediments those hierarchies into everyday life.

If that structural reading holds, then key questions change. Instead of asking whether a particular query is “biased,” probe the representational distribution a system generates across time for identity-linked terms. Which portrayals recur, which are missing, and which are monetized. Interrogating the value of “relevance” when relevance is inferred from past clicks that were themselves shaped by earlier ranking choices is counterproductive. Consider the dignity floor a

results page should guarantee when the query is about a group of people. If ad engines can implement brand safety filters in milliseconds, the infrastructure clearly supports the rapid demotion of demeaning material. The barrier is not technical capacity. It is the absence of a public mandate that places human dignity above auction outcomes. Search is driven by market forces and paid advertising and the operator already prioritizes for profit, which means the operator could prioritize for public interest when it decides to do so.

Search shapes ordinary sense-making before anyone arrives at the hard edges of policing or security. A model that claims to know what you want to see next becomes a device for governing who appears as credible, desirable, or dangerous. In that light, a search result is not a mirror. It is an instrument for predictive representation. Therefore, policy should recognize that ranking systems have public effects. Transparency around auction dynamics for identity-linked queries would allow independent audits of how money intersects with representation. Restrictions on monetization for dignity-sensitive terms would prevent the direct sale of demeaning portrayals as first-page authority. Most importantly, policy should treat search and recommendation as knowledge infrastructures that carry obligations similar to those of broadcasting and libraries, not as mere conduits. This is consistent with the way recent rights frameworks are beginning to treat algorithmic systems as potential threats to human rights when opacity and market concentration are combined.

This reframing also belongs in classrooms and public literacy. Media literacy has often trained students to spot fake news and to triangulate sources. That remains useful, yet the literacy that fits the present moment is infrastructure literacy. Students should learn to read a results page as an editorial artifact produced by algorithms, ad markets, and negotiated categories. They should learn to ask who benefits from the page they see and which portrayals have been made scarce. Many students still approach search as benign, precisely because it is offered as free and universal. Teaching them to interrogate the page will not produce cynicism. It will produce competent readers of contemporary power.

The book under review makes one more move worth adopting. It widens the notion of the digital divide. Devices and broadband are not the whole story if those most in need of knowledge encounter themselves as caricatures when they look for it. Access without dignity is not equity. The fix is therefore structural. It will require public oversight, alternative infrastructures, and explicit design goals that put representational justice on the same footing as speed and scale. The benefit of this view is not only critique. It provides a method for testing whether design choices shift power or simply launder it through code and dashboards. It is a call to govern the infrastructures that already govern us.

Markets reward attention extraction. Security logics reward population sorting. Search is where these two rationalities meet and naturalize themselves in the ordinary. A politics of representation that stops at individual responsibility cannot answer a configuration like that. A structural reading provides a better test. It asks whether the infrastructure shifts the balance toward dignity and plural authority or whether it continues the long work of othering under a new technical gloss.

To synthesize, think of search as a civic institution that has been privatized. The question is not whether the interface should be perfect. The question is which obligations should bind a system that selects what a society sees first. Read the evidence this way and a thorough line emerges.

Marketized ranking will reproduce the visibility patterns of the market unless interrupted by countervailing values. The responsibility to do so belongs not to a single engineer but to the institutions that license, fund, and regulate knowledge infrastructures. Until those obligations are spelled out and enforced, predictive representation will continue to map old lines onto new screens and call them relevance.

Concentrated control over the gateway to knowledge should be met with competition, reporting requirements, and noncommercial anchors that can be embedded into results for dignity-sensitive topics. The final pages of the book argue for structural change. Better machine learning alone cannot deliver it because the outcomes in question are profitable and reputationally useful to the firms that produce them. The way forward is to recast search as public infrastructure and then to govern it accordingly. That means transparency around how ranking and ad systems treat identity-linked queries, enforceable obligations to avoid degrading portrayals, and the cultivation of noncommercial knowledge spaces that can counterbalance commercial logic. It also means teaching people to see a results page as a designed object whose defaults can be changed. To think about search in this way is to see that the stakes are not limited to what we can find. They include who we get to be to one another when we look.

At several points the text insists that variability in results is part of the evidence. The fact that a company can change what appears after public criticism shows that the levers exist and that curation already happens. That plasticity moves the debate away from the metaphysics of algorithmic objectivity and toward institutional accountability. The hardest questions are therefore not technical. They are constitutional for the information age. Who authorizes the criteria by which public knowledge is ranked? What are the obligations of a gatekeeper that has become a default curriculum for citizens? How should value be counted when the first page of results can do cultural damage even if each link is a legal expression of speech?