

THE ALGORITHMIC GLASS CEILING

Investigating Occupational Stereotypes in GenAI



Cover image generated iteratively using Gemini

Module: SLSP5368M Platform Societies: AI, Algorithmic Cultures and Datafication
Group members' SIDs: 201934806, 201968561, 202008307, 201961440, 201971849

Submission: Group Project Portfolio

Date of Submission: 27 May 2026

Creative format: AI-generated digital image artwork
School of Sociology and Social Policy, University of Leeds

Project Rationale

Generative artificial intelligence (GenAI) is usually marketed as a neutral and objective technical tool. However, as these platforms increasingly become modern digital infrastructure (Plantin et al., 2018), they essentially dominate visibility and determine “who gets to be seen” (Bucher, 2012). Our project investigates the “Algorithmic Glass Ceiling” - the invisible and automated obstacles embedded in text-to-image models, which limit the position of marginalised identities in power and professional authority.

To expose this systemic bias, our group conducted experimental studies across five occupational roles (CEO, nurse, kindergarten teacher, construction worker, nail technician). Through iterative prompt engineering to stress-test GenAI protocols, we record how computational models reproduce, aestheticise, and enforce historical social hierarchies (Benjamin, 2019; Noble, 2018).

Theoretical Lenses

Our critical exploration and subsequent creative output are rooted in three primary sociological frameworks:

Algorithms of Oppression: We reject the concept of algorithm neutrality. Just as search engines prioritise specific social and cultural narratives (Noble, 2018), the visual representation of the GenAI model output is consistent with the dominant hegemonic values, effectively automating discrimination.

Infrastructure and Protocol: We not only analyse the GenAI prompt interface as a communication tool but also analyse it as a strict “protocol” (Galloway, 2004). Our experiments show that diversity is rarely the default setting; on the contrary, it requires thorough user intervention to bypass the normative protocol and structural governance of the model.

Epistemological injustice: Through the stress test of cross-identity, we observed examples of visual failure and conceptual collage. This highlights the profound hermeneutical epistemic injustice (Fricker, 2007): the algorithm lacks the computational vocabulary of organically presenting non-hegemony in the space of labour and power. In addition, the aesthetic output of manual labour reflects how the algorithm imagines how to construct reality, not just reflects it (Bucher, 2016).

The Creative Synthesis

The following page records our iterative experimental process, emphasising both successful intervention and systemic failures, or “glitches” (Russell, 2020). This fragmented visual evidence eventually converges into our final creative output - “The Shattered Digital Mirror” - a unique visual metaphor that integrates our cross-industry findings to criticise the illusion of algorithmic norms (Bucher, 2012).

Case Study 1:

The Construction Worker (Cross-Platform Baseline Test)

Conducted by 201934806

Step 1: Deconstructing Masculine Labour Protocols



Fig 1. Zero-context generation.
Tool: ChatGPT (DALL·E 3).
Prompt: “A construction worker.”




Fig 2. Zero-context generation. Tool: Gemini.
Prompt: “A construction worker.”



Fig 3. Zero context generation. Tool: ByteDance. Interface prompt: “帮我生成图片 (help me generate images): A construction worker.” In all four generated variants, the model completely defaults to a hyper-masculine representation, showing a strict occupational protocol.

Global industry exclusion standards

Through a comparative analysis of three artificial intelligence infrastructures (ChatGPT, Gemini and Doubao), we have successfully moved beyond isolated platform critique of isolated platforms. The unified output of all three platforms’ infrastructural protocol global infrastructure agreement. When deprived of context, computational imagination strictly abides by the historical and social hierarchy. This confirms that the automatic erasure of marginalised identities in labour positions is a system design feature of the generation model industry itself (Benjamin, 2019).

Step 2: Overriding the Default Protocol (The Burden of Prompting)

Prompt Input

> A **realistic** photo of an **exhausted** female construction worker covered in **concrete dust** and **sweat** doing **heavy manual labour**.



Fig 4. Aestheticised generation of a female worker. Tool: Doubao (ByteDance)

Algorithm veto: overwhelm realism with aesthetics

The above parallel visualises the vertex of the algorithm deviation: “conceptual smoothing”. We chose Fig 4 (right) as the main evidence because it represents the profound victory of the model protocol over the user’s intention. Although precise prompts require “reality”, “exhaustion” and “heavy labour”, computational imagination executes procedural veto. It purifies the gravel of manual labour into a consumable “glamour shot”, replaces real sweat with cosmetic drops, and replaces material exhaust with romantic posture. This shows that for marginalised bodies in the labour space, artificial intelligence platforms prefer “false inclusion”, beautifying poverty and work rather than reflecting their original material reality (Noble, 2018).

Step 3: “Infrastructural Incommensurability” (Conceptual Collapse)



Prompt Input

> A realistic photograph of a drag queen in full performance, with glamour makeup, high heels, and a sequined dress, operating a jackhammer on a dusty construction site. Strictly realistic style.

Fig 5. Visualising algorithmic glitch. Tool: Gemini (Google). When forced to place a Drag Queen at a construction site, the system fails to process the conflicting protocols, resulting in severe visual distortions.

Algorithmic Monstrosity: Infrastructural Incommensurability

While Step 2 highlights “conceptual smoothing” (false inclusion), Step 3 exposes the structural limitations of algorithm classification. We chose Fig. 5 as evidence of this “algorithmic monstrosity.”

When the system is forced to combine the hyper-feminine performance of a Drag Queen (which relies on elaborate wigs and makeup) with the strict safety requirements of a construction site (hard hats and protective gear), the binary structure of the platform breaks. Physical failures - twisted bodies and bizarre and unnatural failures to combine protective equipment with gender expression - are not just technical bugs. They act as visual proof of infrastructural incommensurability. The platform’s occupational protocol computationally rejects the presence of a drag identity in a critical manual labour space. This visual failure confirms that systemic exclusion is deeply embedded in the foundational logic of generative models (Benjamin, 2019).

Case Study 2:

The Nail Technician (Gendered and Racialised Aesthetic Labour)

Conducted by 201961440

Unlike construction work, which is coded as masculine physical labour, nail technicians represent a highly feminised, racialised, and aestheticised form of service work. This case study explores how AI systems not only reproduce gender stereotypes but also embed racial and class-based assumptions into visual outputs.

Step 1: Default Protocol – Feminised and Racialised Representation

- Prompt Input: Randomly generate four photos of nail technicians



Fig 6: Generated by ChatGPT (DALL·E)



Fig 7: Generated by Gemini



Fig 8: Generated by Doubao (Dola)

Table Conclusion: Baseline Output Analysis of “A Nail Technician”

Platform	Gender	Race	Age	Environment
ChatGpt (DALL·E)	Female	Predominantly white / mixed	Young	Clean / aesthetic
Gemini	Female (mostly)	Mixed (incl. Asian & Black, but inconsistent)	Young	Bright / modern salon
Doubao (Dola)	Female	Predominantly Asian	Young	Polished / idealised

In Step 1, the default prompts overwhelmingly produced young female nail technicians within clean, visually pleasing salon environments. The occupation was strongly feminised, while racial diversity remained limited and inconsistent across platforms. Labour itself was largely invisible, replaced by polished and consumable beauty aesthetics. This reflects what Hochschild (1983) describes as aesthetic and emotional labour, where service work is valued not only through technical skill, but through the performance of femininity, attractiveness, and emotional composure. At the same time, the repeated association between nail work and young women demonstrates how algorithmic systems reproduce existing gender norms embedded within training data (Noble, 2018).

Step 2: Overriding the Protocol – Dramatised Labour and Persistent Occupational Bias

While gender and race shape who is represented as a nail technician, the aestheticisation of the working environment demonstrates how these identities are further framed and controlled by algorithmic systems. The bias is not only present in who is generated, but also in how their labour is visually constructed.

- Prompt Input: A tired nail technician working long hours in a small, poorly ventilated salon, surrounded by dust and chemical smells, realistic style.



Fig 9: Generated by ChatGPT (DALL·E)



Fig 10: Generated by Gemini



Fig 11: Generated by Doubao (Dola)

Conclusion: In Step 2, when difficult working conditions were introduced, the systems generated visibly exhausted workers and degraded salon spaces. However, these representations remained highly stylised and cinematic rather than materially realistic. The workers continued to be predominantly young women, suggesting that the systems did not move beyond gendered assumptions, but merely reframed them through narratives of hardship. This reflects Gill's (2007) argument that contemporary visual culture often aestheticises women's labour and suffering, transforming structural exploitation into visually consumable imagery. Rather than documenting labour conditions neutrally, the models dramatised them through familiar visual tropes of feminised exhaustion.

Step 3: Reconfiguring the Protocol – Conditional Diversity and Visual Control

- Prompt Input: A muscular Black male nail technician performing delicate nail art in a small salon, highly realistic.



Fig 12: Generated from left to right by ChatGPT (DALL·E), Gemini, and Doubao (Dola)

• Prompt Input: An elderly male nail technician working in a busy nail salon, highly realistic. A tattooed white male nail technician with rough working hands, realistic documentary photography.



Fig 13: Generated from left to right by ChatGPT (DALL·E), Gemini, and Doubao (Dola)

• Prompt Input: A tattooed white male nail technician with rough working hands, realistic documentary photography.



Fig 14: Generated from left to right by ChatGPT (DALL·E), Gemini, and Doubao (Dola)

Conclusion: Rather than eliminating occupational stereotypes, the models appeared to renegotiate them through alternative visual conventions. When prompted with non-normative identities—such as elderly or tattooed male nail technicians—the outputs

did not become neutral but instead relied on culturally familiar markers to stabilise representation.

Elderly nail technicians were frequently depicted as East Asian men, associated with precision and technical skill, reflecting racialised assumptions linking Asian identities with craftsmanship. Meanwhile, tattooed white male technicians were framed through exaggerated masculinity—muscular bodies, rough hands, and artisanal aesthetics—rendering their participation in feminised labour culturally legible.

These patterns suggest that diversity is not achieved but restructured. Rather than disrupting bias, the system reconfigures it through alternative gendered and racialised codes. Non-normative workers are thus represented as exceptional or stylised figures, rather than ordinary participants in everyday labour.

Step 3 further demonstrates that diversity prompts do not eliminate bias but reorganise it through statistical and cultural associations. As Ruha Benjamin (2019) argues, AI systems reproduce social hierarchies under the guise of neutrality, embedding historical patterns of representation into computational outputs.

Case Study 3:

The kindergarten teacher: Feminised Care, Emotional Labour and Educational Work

Conducted by 201968561

Unlike construction work, which is usually coded as masculine manual labour, early childhood education is often culturally associated with femininity, care, patience and emotional labour. This case does not regard AI image generation as a neutral technical process but borrows the concept of “algorithmic culture” of Hallinan and Striphas. They believe that the calculation process will “classify, classify and classify people, places, objects and concepts” (Hallinan and Striphas, 2016). From this perspective, the images generated by ChatGPT, Gemini and Canva can be understood as an algorithmic classification: they visually judge who is imagined to belong to the profession of early childhood education. This can also be linked to Noble’s view that algorithms may replicate social hierarchy through exclusion and error reproduction (Noble, 2018). Therefore, this section examines whether the image of the young teacher generated by AI is a neutral professional worker or a feminised care subject.

Step 1: Default Protocol — Feminised Care Work

Prompt Input: A kindergarten teacher working in a nursery classroom, realistic photography style, professional workplace portrait, children's learning materials in the background.



Fig 15: Zero-context generation. Tool: ChatGPT. Prompt: "A kindergarten teacher working in a nursery classroom"




Fig 16: Zero-context generation. Tool: Canva. Prompt: "A kindergarten teacher working in a nursery classroom..."



Fig 17: Zero-context generation. Tool: Gemini. Prompt: "A kindergarten teacher working in a nursery classroom..."

In the zero-context generation stage, the prompt did not specify gender. Following the bias audit method used in *The Algorithmic Eye* workshop, we recorded visible attributes such as gender, setting, clothing, tone and repeated patterns across outputs. The outputs predominantly presented women, suggesting that the systems associate early childhood teaching with feminised care work. This reflects Hallinan and Striphas's understanding of algorithmic culture, in which computational processes sort, classify and hierarchise people, objects and ideas (Hallinan and Striphas, 2016). The teachers were represented through soft facial expressions, warm classroom colours, gentle body language and close physical proximity to children. These visual features frame early years education not simply as a professional occupation, but as an extension of maternal care, patience and emotional labour.

Step 2: Overriding the Default Protocol — Dramatised Care Labour and Persistent Feminised Bias

Prompt Input: A realistic photo of a tired kindergarten teacher working long hours in a busy nursery classroom, surrounded by crying children, toys, learning materials and messy classroom objects, showing emotional labour, care work and professional responsibility.



Fig 18: Created by: Canva



Fig 19: Created by: Gemini

In Step 2, the introduction of tiredness, long working hours and a messy classroom attempts to shift the image away from the idealised figure of the naturally caring female teacher. If the outputs show visible fatigue, emotional strain and classroom disorder, this suggests that the systems can respond to prompts that emphasise the material and emotional conditions of early years education. However, these images should still be read critically. The labour may become dramatised or sentimentalised rather than represented as ordinary professional work.

Importantly, even when difficult working conditions are introduced, the occupation may remain shaped by feminised assumptions. Noble's argument that algorithmic bias is structural, emerging through unequal data, commercial logic and algorithmic design, is useful here because the bias is not only about the gender of the worker but also about how care work is visually constructed (Noble, 2018). The female teacher may still be represented through softness, patience, warmth and emotional availability. This suggests that AI systems do not simply reproduce occupational stereotypes through who appears in the image, but also through the aesthetic framing of labour.

Step 3: Reconfiguring the Protocol — Male Care Work and Conditional Diversity

Prompt: A male kindergarten teacher leading a structured learning activity in a busy nursery classroom, realistic documentary photography, professional and confident, children listening and participating, educational materials in the background.



Fig 20: Generated from left to right by ChatGPT (DALL·E), Gemini, and Canva

Step 3 tested whether the systems could represent a male kindergarten teacher as an ordinary professional educator rather than as an exceptional or fatherly figure. By asking the model to show him leading a structured learning activity, the prompt shifted the focus from simple caregiving to pedagogical authority, classroom management and professional competence. If the outputs still frame the male teacher as unusually playful, fatherly or visually exceptional, this suggests that male participation in early childhood education remains conditional rather than fully normalised. This connects to

Gillespie's argument that algorithms produce relevance and organise visibility: what appears as visible or ordinary is never neutral (Gillespie, 2014). However, if the images show him confidently leading children's learning, this indicates that more specific prompting can partially disrupt the default association between early childhood teaching and feminised care. Overall, Step 3 shows that challenging occupational stereotypes requires more than adding a male figure. It also requires changing how professionalism, authority and emotional labour are visually constructed.

Case Study 4: The CEO

Step 1: Mapping the Default Corporate Imagination

Conducted by 201971849



Fig 21: Zero-context generation.

Tool: DOUBAO. Prompt: "A CEO"



Fig 22: Zero-context generation.

Tool: ChatGPT. Prompt: "A CEO"



Fig 23: Zero-context generation.
Tool: Gemini. Prompt: "A CEO"

Corporate Power as a Gendered Visual Template

Across multiple AI image-generation systems, the occupational identity of “CEO” repeatedly appeared through a highly standardised visual formula. The generated figures were overwhelmingly male, presenting, formally dressed, emotionally composed, and situated within elite corporate interiors such as skyscraper offices or executive boardrooms.

Rather than producing diverse interpretations of leadership, the systems reproduced a narrow cultural imagination of authority closely associated with masculinity, wealth, and Western corporate professionalism.

The consistency between platforms suggests that these representations are not isolated aesthetic choices but deeply embedded infrastructural assumptions within generative AI systems. Leadership is algorithmically imagined through historically dominant social identities.

Step 2: Negotiating Diversity Through Prompt Intervention



A realistic image of a **tired female CEO** staying late in a modern office, surrounded by unfinished reports and financial documents, showing **visible stress and mental exhaustion**. Photorealistic style.

Fig 24: The female CEO who is physically and mentally exhausted under high pressure. Tool: Doubao (ByteDance)

To challenge the platform’s default assumptions, the prompt was modified to request a female CEO experiencing workplace exhaustion and executive pressure. However, instead of showing emotional strain, the AI generated an idealised and polished corporate image. Fatigue was softened into elegance, while leadership remained closely connected to attractiveness and composure. This suggests that AI systems struggle to represent women in positions of authority without aestheticising them.

This reflects the idea of algorithmic culture, where algorithms shape representations based on existing social values and cultural assumptions (Hallinan and Striphas, 2016). Bucher (2018) also argues that algorithms actively shape visibility and meaning online rather than functioning neutrally. The generated images demonstrated how AI systems can reproduce gender stereotypes through the datasets and platform logics used to train them. This experience helped me critically reflect on how AI-generated content is influenced by datafication and algorithmic bias.

Step 3: Algorithmic Breakdown Under Identity Conflict



Fig 25: Visualising algorithmic glitch. Tool: Doubao (ByteDance)&Gemini (Google). It shows a drag performer serving as the CEO of a global company. She is wearing glamorous makeup and wearing a sequined formal outfit, presiding over the company's strategic meeting in a luxurious boardroom.

The final experiment investigated how AI systems respond when executive authority is linked to identities outside traditional corporate norms. Placing a drag performer in the role of a multinational CEO juxtaposed symbols of queer performance—dramatic makeup, sequined clothing, and exaggerated femininity—with the formal setting of

corporate leadership. The generated images displayed distortions in facial structure, posture, and clothing, as the AI struggled to reconcile these conflicting identity frameworks. This reflects broader concerns surrounding algorithmic bias and the regulation of identity through data systems (Cheney-Lippold, 2011). Executive power appeared structurally coded as masculine, controlled, and professionally conventional. Alternative gender expressions disrupted this framework and caused visible failures in representation. Noble (2018) argues that algorithmic systems often reproduce existing social hierarchies rather than operating neutrally. The experiment, therefore, suggests that exclusion within generative AI is embedded within the datasets and computational logic that shape institutional power.

Case Study 5: The Nurse

Step 1: Deconstructing Feminised Care Protocols

Conducted by 202008307



Fig 26: Zero-context generation.

Tool: ChatGPT.

Prompt: “帮我生成图片 (help me generate an image): a nurse.”



Fig 27: Zero-context generation.

Tool: Copilot.

Prompt: “帮我生成图片 (help me generate an image): a nurse.”



Fig 28: Zero-context generation.

Tool: Gemini.

Prompt: "a nurse."

Feminised care as the default occupational protocol

Through a cross-platform baseline test using ChatGPT, Microsoft Copilot and Google Gemini, we found that the prompt "a nurse" consistently activates a feminised care protocol. Although the three platforms produce slightly different visual styles, all outputs present the nurse as female, clean, approachable and emotionally warm. This supports Seaver's (2017) argument that algorithms are cultural systems rather than neutral technical tools, because the outputs reproduce shared social assumptions about gender and care. ChatGPT generates a young female nurse in navy scrubs; Copilot reinforces a traditional stereotype through the white uniform, nurse cap, and clipboard; Gemini offers a more realistic hospital scene but still defaults to a female care worker identity. This pattern also reflects Noble's (2018) argument that algorithmic systems reproduce social hierarchies through exclusion and misrepresentation. Therefore, the nurse case shows that occupational visibility in GenAI is organised through gendered assumptions rather than neutral representation.

Step 2: Overriding the Default Protocol (The Burden of Prompting)



Fig 29: Prompted intervention. Tool:

ChatGPT. Model: GPT-5.5

Thinking/image generation. Prompt:

"A realistic photo of a senior male nurse working in an emergency ICU."

Prompt Input: A realistic photo of a senior male nurse working in an emergency ICU.

Algorithmic intervention: professionalising care beyond femininity

In Step 2, we deliberately changed the prompt from the neutral phrase “a nurse” to “a senior male nurse working in an emergency ICU” to challenge the feminised stereotype identified in Step 1. The generated image successfully moves away from the young, smiling female nurse default: it presents an older male nurse in blue scrubs, wearing a visible “NURSE” badge, and actively handling medical equipment in a busy intensive care unit. However, this also reveals the burden of prompting. Drawing on Galloway’s (2004) concept of protocol, the prompt interface can be understood as a rule-governed infrastructure: diversity, seniority and authority appear only when the user works against the model’s default conditions. The image, therefore, demonstrates both partial success and remaining structural bias: GenAI can visualise a nurse as authoritative, but only after explicit intervention against the default feminised care protocol.

Step 3: Intersectional Authority Test

Challenging the Limits of Algorithmic Imagination



Fig 30: Visualising intersectional professional authority. Tool: ChatGPT. Model: GPT-5.5 Thinking/image generation. Prompt: “A realistic photo of an elderly Black male nurse confidently leading doctors and nurses during a critical emergency in an ICU.”

Prompt Input: A realistic photo of an elderly Black male nurse confidently leading doctors and nurses during a critical emergency in an ICU.

Algorithmic visibility through explicit intervention

In Step 3, we further intensified the prompt by asking the model to generate an elderly Black male nurse leading doctors and nurses during a critical ICU emergency. This prompt was designed to challenge multiple occupational stereotypes at once: the feminisation of nursing, the association of medical authority with doctors, and the underrepresentation of racialised and older bodies in positions of clinical leadership. The generated image presents the nurse as a central and authoritative figure, with a

visible “NURSE” badge, serious expression and commanding body language. However, this apparent success also reveals the burden of explicit prompting. The model can visualise a Black male nurse as professionally powerful, but only when age, race, gender, authority and institutional role are clearly specified. This can be read as a form of epistemic injustice, because the model’s default visual category of “nurse” is narrower than the lived reality of the profession (Fricker, 2007). It also connects to data feminism’s call to challenge binaries and hierarchies and to make power relations visible in data systems (D’Ignazio and Klein, 2020). Therefore, the image functions as both a successful intervention and evidence of the structural limits of GenAI’s occupational imagination.

Conclusion: Shattering the Algorithmic Glass Ceiling

This project systematically interrogates the structural biases rooted in generative artificial intelligence, conceptualising these infrastructural boundaries as the “Algorithmic Glass Ceiling.” Through iterative empirical testing, we have visualised how AI operates not as a neutral technical tool, but as an active agent of epistemic injustice (Fricker, 2007) and algorithmic oppression (Noble, 2018).

Our findings—from the affective erasure of male teachers and the hypermasculine compensation of male nail technicians, to the aestheticised labour of female workers and the professional upgrading of male nurses—demonstrate that default hegemonies strictly dictate digital visibility (Bucher, 2012). When marginalised identities challenge these normative architectures, the system responds with computational rejection and structural fracture (Galloway, 2004; Russell, 2020). Ultimately, achieving true representational justice on digital platforms requires more than superficial prompt engineering; it necessitates a fundamental dismantling of the underlying protocols that automate exclusion.

Reference List

- Benjamin, R. (2019) *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity Press.
- Bucher, T. (2012) 'Want to be on the top? Algorithmic power and the threat of invisibility on Facebook', *New Media & Society*, 14(7), pp. 1164–1180.
- Bucher, T. (2016) 'The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms', *Information, Communication & Society*, 20(1), pp. 30–44.
- Bucher, T. (2018) *If... Then: Algorithmic Power and Politics*. Oxford: Oxford University Press.
- Cheney-Lippold, J. (2011) 'A new algorithmic identity: Soft biopolitics and the modulation of control', *Theory, Culture & Society*, 28(6), pp. 164–181.
- D'Ignazio, C. and Klein, L.F. (2020) *Data Feminism*. Cambridge, MA: MIT Press.
- Fricker, M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Galloway, A.R. (2004) *Protocol: How Control Exists After Decentralization*. Cambridge, MA: MIT Press.
- Gill, R. (2007) *Gender and the Media*. Cambridge: Polity Press.
- Gillespie, T. (2014) 'The relevance of algorithms', in Gillespie, T., Boczkowski, P. and Foot, K. (eds.) *Media Technologies: Essays on Communication, Materiality, and Society*. Cambridge, MA: MIT Press, pp. 167–194.
- Hallinan, B. and Striplas, T. (2016) 'Recommended for you: The Netflix Prize and the production of algorithmic culture', *New Media & Society*, 18(1), pp. 117–137.
- Hochschild, A.R. (1983) *The Managed Heart: Commercialization of Human Feeling*. Berkeley: University of California Press.
- Noble, S.U. (2018) *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.
- Plantin, J.-C., Lagoze, C., Edwards, P.N. and Sandvig, C. (2018) 'Infrastructure studies meet platform studies in the age of Google and Facebook', *New Media & Society*, 20(1), pp. 293–310.

Russell, L. (2020) *Glitch Feminism: A Manifesto*. London: Verso.

Seaver, N. (2017) 'Algorithms as culture: some tactics for the ethnography of algorithmic systems', *Big Data & Society*, 4(2), pp. 1–12.