Digital inequalities, placemaking and geographical diversity in tweets about Cape Town railways, 2010 - 2022

Marion Walton
Centre for Film and Media Studies
University of Cape Town
marionwalton@uct.ac.za

Abstract

Digital placemaking is implicated with power relations and can expose inequities, exclusions and erasures. A longitudinal archive of 5606 tweets geotagged in Cape Town was collected to explore placemaking practices and inequality in commuter posts about rail travel in the city over the period 2010-2022. Named Entity Recognition (NER) identified 1755 tweets from this archive as a record of digital placemaking about the major train lines in Cape Town. A sustained bias over time in favour of the suburban Southern Line is detected. Strong core-periphery patterns emphasize major interchanges in the collection, while measures of diversity indicate small increases in geographical diversity in the railway system over time. Structural biases in access to the means of digital placemaking thus shape this archive. While not causally related, a similar geographical skew towards socially hegemonic places in the system was apparent while using an LLM to annotate a sample of the tweets (n=350). Algorithmic biases were even more pronounced in a synthetic sample of LLM-generated tweets about the railway system, which overwhelmingly focused on local attractions on the Southern Line. Thus the GenerativeAI tools used to read the archive themselves introduce algorithmic biases and potentially perpetuate socially hegemonic digital defaults.

1 Introduction

Although the right to freedom of movement is constitutionally protected in South Africa, everyday mobilities of people in Cape Town continue to be shaped by spatial consequences of dispossession, segregation, market-led development, gentrification and decades of lip-service to progressive policies. While the extent of segregation has diminished, in 2011 Cape Town remained South Africa's most unequal city (Turok et al., 2021).

The Twitter Train Stories is a longitudinal archive (2010-2022) of 5606 tweets, all geotagged in Cape Town. The archive was collected to explore how placemaking on the Twitter platform can provide insights into the different and unequal worlds of commuter experiences of rail travel in the city.[1]

Commercial GenAI foundation models such as OpenAI's ChatGPT, Google's Gemini and Anthropic's Claude are used globally for everyday tasks, including information retrieval. Researchers are asking to what extent such models are geographically biased in favour of the Global North and what the consequence of such biases might be for user interactions around the world (Liu et al., 2026). This paper explores some local geographic biases of LLMs and NER tools in the Cape Town context, comparing them with the inequities and exclusions identified in the Twitter Train Stories archive to highlight how structural bias, data bias and algorithmic bias are separately implicated in placemaking and digital inequality and how these can also influence tools used in computational social science and digital humanities.

Mapping these records reveals inequalities of participation and uneven distribution of placemaking practices. While Generative AI models are useful in such mapping exercises they can both reflect and perpetuate the hegemonic geographical biases of the broader society.

1.1 Commuter rail, place and space in Cape Town

In South Africa, racialized dispossesion and apartheid urban planning peripheralized the majority to outlying matchbox housing in townships and informal settlements, creating fundamental disparities in access to housing and work (Czeglédy, 2004) and inequalities in mobility which persist today.

[1] LLM prompts, station lists and output are available in the project repository (Walton, 2026). At the time of writing the archive of tweets is undergoing review for publication in pseudonymised form.

This work is licensed under CC BY SA 4.0. To view a copy of this license, visit https://creativecommons.org/licenses/by-sa/4.0/deed.en
The copyright remains with the authors.

The radial design of railways, built in the colonial and apartheid eras, segregated passengers by race in separate carriages and established separate lines serving racially segregated areas. The Southern Line was initially developed in the colonial era to connect Cape Town's CBD with outlying suburbs and the False Bay coast, while the Cape Flats and Central Lines shuttled commuters from racially segregated townships (Pirie, 2013). Trains remain the slowest and most unsafe form of transport form of public transit (Statistics South Africa, 2020; Vanderschuren and Lane-Visser, 2025).

1.2 Making places

The idea of placemaking recognises the significance of the urban conditions of daily life and the production of space through social practice, sensory phenomena, and "products of the imagination" (Lefebvre, 1991). Space is inhabited and constructed as it is read, while meaning is produced and reproduced in space. These material, social and semiotic dimensions create and recreate places. Scholarship in urban design has focused on deliberate strategies in urban design to improve safety and security, aesthetics, liveliness and activity (Chidambara, 2019). In tourism, places are marketed generating multiple versions, attachments and meanings to be consumed, with "destination branding" shaping how both visitors and locals imagine and consume places (Bickford-Smith, 2009). Placemaking is also implicated with people's sense of boundaries, and the power struggles associated with enforcing boundaries, as well as tensions around place and the potential transgressions implicated in tourism, migration, and self-presentation (Halegoua and Polson, 2021).

1.3 Digital placemaking

Digital placemaking practices use digital media to "create a sense of place for oneself and/or others", by creating emotional attachments and negotiating differential access to mobility (Halegoua and Polson, 2021). The process of representing places digitally is implicated with power relations and can expose inequities, exclusions and erasures. For example, by integrating queer perspectives, digital placemaking practices can recreate churches as sites of queer religious belonging both online and offline (Figueroa et al., 2026).

1.4 Social media, location determination and placemaking

Placemaking technologies such as the GeoIDs used for this study are integral to social platforms such as Twitter, which are often used as mobile applications and incorporate location intelligence systems, including location determination technologies, databases, data collection practices, and machine learning algorithms (Özkul, 2021). Özkul (2021) terms location intelligence the "algorithmic fix", in that the goal of such systems is to identify a user's whereabouts in order to sort and classify the person as well as the places they visit. One of the goals of such systems of "personalisation", is to more precisely tune advertising (Carah et al., 2024) in order to target specific users in suitably receptive contexts.

Placemaking technologies are only one part of the picture. People also appropriate spatial signifiers in their own digital placemaking practices, such as using the hashtags such as #egoli to label an image or video on Instagram, or using geolocation to track and share their exercise routines on apps such as Strava (Bosch, 2021).

1.5 Digital placemaking and geographic bias

Everyday digital placemaking practices are used to develop apparently universal global databases of content such as those owned by all of the social media platforms, including Twitter.

Scholars have raised the alert about the serious biases in these databases, such as for example, Wikipedia's early highly disproportionate coverage of articles about the United States and European countries (Graham et al., 2014).While global coverage is improving, there are still disparities (Yu et al., 2025) which relate to the location and placemaking interests of Wikipedia editors. Other sources of geographical bias on Wikipedia come about because of the provenance of the information, such as for example, when articles about places in the Global South and about Africa in particular are created disproportionately by editors who live in countries in the Global North. Western knowledge preservation and sharing conventions have also been found to present many obstacles to local Wikipedia editors, particularly in low-resource language communities (Nigatu et al., 2024). The structural biases in access to the means of digital placemaking produce geographical data biases on the web and in the training data for LLMs, which are perpetuated by

algorithmic biases in algorithmic systems, including language models (Baeza-Yates, 2018).

1.6 Placemaking with Named Entity Recognition and GenAI

Named Entity Recognition (NER) is a Natural Language Processing (NLP) technique used to automate the identification of specific entities in text, and is often conducted today with LLMs and other language embeddings. The representational biases on Wikipedia influence the tools used for popular contemporary techniques for Named Entity Recognition because so-called "foundation" models which underpin popular commercial Generative AI systems are pretrained on vast collections of digital geographic content, including Wikipedia, and on user generated content on social platforms, which themselves reflect global placemaking practices and patterns of activity.

The biases of LLMs are well-documented to the extent that scholars have argued that they can be considered a fundamental property of the technology (Resnik, 2025). Hence it is not surprising that LLMs embed substantial geographic biases. For example, the benchmark datasets used to test the accuracy with which models can answer questions are often derived from resources such as Wikipedia and thus fail to represent a broad and balanced range of geographic regions (Liu et al., 2026). Geographic representational bias is known to be inherent to many benchmarked training datasets, whether textual or visual (Liu et al., 2026).

Interest in documenting fairness and representation across geographic regions often stops at the country level, and model development runs the risk of using benchmarks which perpetuate inequalities as a result of aggregation bias (Liu et al., 2026).

Discussions of geographic bias reference the need to understand geographic diversity of such sources as their richness, or the number of places they reference, as well as the evenness of coverage, or the balance in the amount of information provided about specific places (Liu et al., 2026, 2025).

When tested on factual accuracy using World Bank development data, LLMs were found to have serious shortcomings, with biases based on the geographic region (favouring Western countries) and income level of countries (Moayeri et al., 2024). This investigation elicited particularly poor performance when LLMs were prompted for information about Sub-Saharan Africa and higher error rates were associated with lower income countries in general.

The Large Language Models (LLMs) used in Generative AI perform unevenly because of both representational and aggregation bias.

Representational bias refers to the use of training data from one geographical region for applications in another one. For example, benchmark datasets for computer vision Open Images and ImageNet were shown to be heavily biased in favour of the US, with 60% of the training data coming from six countries in Europe and North America. (Liu et al., 2026) Aggregation bias is inherent to spatial analysis and occurs when generalised assumptions are used across spatially and geographically diverse areas.

Geoparsing disparities introduce additional biases. Geoparsing consists of named entity recognition to identify toponyms (place names) and to resolve them to specific sets of coordinates (geocoding). Toponym recognition involves the detection of place names in text, in other words it is a named entity recognition task (Liu et al., 2026:206), while toponym resolution or geocoding identifies specific co-ordinates associated with the place name.

An LLM-specific form of geographical bias seen in synthetic text is the tendency to produce a limited set of default "prototypical places" (Liu et al., 2026) which leads to geographical narrowness when describing places, identifying instances, or synthesizing visual representations of specific geographic entities.

Language models thus perform unevenly because of a combination of both representational and aggregation bias, as well as a type of algorithmic bias whereby geocoding algorithms are often designed to prefer more populous cities.

Methods

1.7 Data access and sampling

Tweets about trains geo-tagged with places situated within the broader Cape Town area were retrieved via researcher access to the Twitter Researcher API 2.0 in March 2023.

Until 2023, the Twitter Academic Research API2.0 I allowed free access to archive searches for historical tweets and over 10 million tweets per month. A Python (v3.6) script was used to access the Twitter records.

Archive searches were conducted for a range of keywords [train] and [station] where the geoID

[8b9ec16fdc0d7e55] indicated a location in Cape Town, South Africa (the area between coordinates longitude 18.3180332 to 18.6600898 and latitude -34.35839 to -33.8849254).

Additional queries found a small number of additional tweets using the Afrikaans query [trein]. There were no tweets using the isiXhosa keyword [ulolile]. Unfortunately data access was interrupted at this stage with withdrawal of free access to the Twitter Academic Researcher API. Thus while the dataset includes many multilingual tweets, they should not be considered exhaustive. A dataset of 5606 tweets matching the queries was compiled.

1.8 Geo-tagged tweets

Geo-tagged tweets constitute a small subset of tweets where an individual tweet has been tagged by the user with the name of a specific place. The query returned results within a bounding box set to the geoID [8b9ec16fdc0d7e55] which encompasses the broader Cape Town area.

While finely grained geolocation pinpoints the location of a device to a specific position, coarsely grained geolocation such as a city-sized bounding box was sufficient for the purposes of the study, where the goal was to explore place-making, not to pinpoint the exact spatial location of a device. The geoID metadata gives access to a user-created tag and does not reveal the exact spatial location of the user. So for example, the user may well have uploaded a photograph long after visiting the place tagged in their tweet.

Consequently, the Twitter Train Stories archive should not be considered representative of Twitter use at the time or of the entire Twitter discourse about trains in Cape Town. The presence of the geoID reflects a decision by the user to contextualise their tweet by tagging a specific place within the broader Cape Town area. Filtering the results in this way also excludes untagged online discussion of media stories about trains and bots or spam on trending hashtags.

1.9 Named entity recognition

1.10 Station/line sample and LLM annotation

An initial manual review of the entire archive and a number of Named Entity Recognition (NER) techniques were used to compile a purposive sample of tweets discussing or referencing specific stations and lines within the Cape Town railway system. Wikipedia-based DBPedia had serious difficulties recognising place names, particularly if they were multilingual or colloquial versions. 4CAT's NER feature (Peeters and Hagen, 2022) with Python library SpaCY and an open weights LLM (MistralAI, 2025) were also tested. This process of annotation resulted in a dataset of 1755 tweets along with the list of all nodes in the railway network which were mentioned in the text. If tweets mentioned more than one station or line, all were included in the annotation. These nodes were geocoded (see Prompt 3) and mapped in Figure 1.

1.11 Subsample for manual coding

A random stratified subsample was drawn from tweets which mentioned stations on the Northern or Southern Line or one of the major Interchanges. In the cases of the smaller categories (Cape Flats and Central Lines) all available tweets were included in the subsample. Together these produced a subsample (n=350) for more detailed content analysis and annotation (see Table 1). These tweets were again manually categorised as referring to one of the lines in Cape Town. If no line could be identified, tweets were classified as "No line mentioned". As far as possible, tweets were viewed on the web in order to infer what they were referring to according to the communicative context and the geography of Cape Town.

More satisfactory annotation performance was achieved with the commercial LLM Claude Opus 4-6, prompted with a list of all stations and railway lines in Cape Town. In Prompt 1, the LLM was given a list of lines in Cape Town and instructed to return the train line referred to in the tweet (see Appendix). The output for Prompt 1 showed that Opus 4-6 could not reliably identify the railway line, and so the prompt was adapted to be more detailed (see Prompt 2 in Appendix). In ambiguous cases, such as when the station mentioned was an interchange and could have belonged to two or more lines, tweets were allocated to the category "Interchange". Annotations thus represent the Cape Town stations and lines mentioned in or implied by the tweet. Prompt 2 was more detailed and made use of a JSON schema with a list of specific stations and/or train lines (see Appendix and (**?**)).

A classification evaluation with 4CAT was conducted to compare the Claude Opus 4-6 classification with the manual line categorisations (Peeters and Hagen, 2022) to calculate evaluation metrics including accuracy, precision, recall, F1, and Cohen's Kappa. These metrics allowed a comparison

label	accuracy	precision	recall	f1	cohens_kappa	support
overall	0.71143	0.76642	0.7008	0.70736	0.66206	350
Cape Flats Line	0.89714	0.76923	0.70423	0.73529	0.67162	71
Central Line	0.92	0.92424	0.72619	0.81333	0.76335	84
Line not mentioned	0.86857	0.78571	0.35484	0.48889	0.42557	62
Northern Line	0.90857	0.72131	0.74576	0.73333	0.67818	59
Southern Line	0.87429	0.63158	0.97297	0.76596	0.68525	74

Table 1: Classification analysis comparing LLM coding with Claude Opus 4-6 and manually coded tweets.

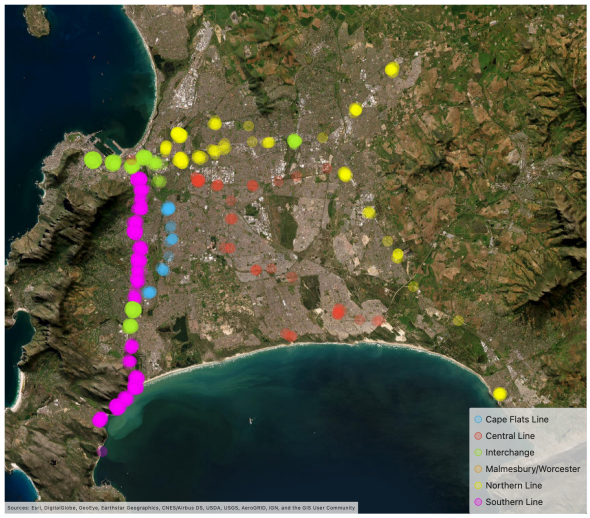


Figure 1: Distribution of stations mentioned in relation to major Cape Town commuter railway lines (2010-2022) n = 1755. Dithering of coordinates and transparency are used to indicate higher concentrations of tweets about a particular station. Map created with Orange Data Mining.

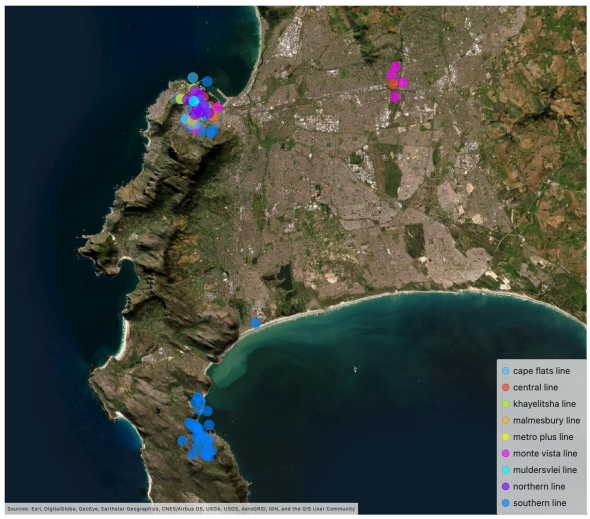


Figure 2: Destinations and lines mentioned in synthetic commuter tweets generated by Claude Opus 4-6. Dithering of coordinates and transparency are used owing to the high concentrations of tweets mentioning the stations. Map created with Orange Data Mining.

of the accuracy of the LLM-generated annotations (Claude Opus 4-6) against the manual coding of tweets by railway line.

1.12 Synthetic tweet generation

Using Prompt 4 (see Appendix), Claude Opus 4-6 was provided with a complete list of 91 stations on all the Cape Town lines, and prompted to generate a sample of synthetic tweets, as written by someone living in Cape Town and tweeting from one of the listed stations on Twitter/X at 8am on Monday morning. The synthetic tweets (N=91) were coded manually according to the selected destination and line and mapped (see Figure 2).

1.13 Findings

Figure 1 maps the geographic distribution of stations mentioned in the Twitter Train Stories archive (n=1755). The points indicate the location of individual stations as mentioned in the tweets in the Station/Line sample, extracted with Prompt 2 (n = 1755). Tweets are densely concentrated around Southern Line stations along the route from Cape Town through Retreat, and then again on the coastal section of the line between Muizenberg and Simonstown (pink icons in Figure 1). Dense clusters appear around the CBD and the other major interchanges (green icons in Figure 1). By contrast, the areas in Cape Town with the highest level of transit dependent population (Khayelitsha - 3045 persons/km2, Mitchells Plain/Gugulethu - 2630 persons/km2 and Blue Downs - 1231 persons/km2) (Vanderschuren et al., 2021) are sparsely represented in the Twitter Train Stories. Only a sprinkling of tweets represent references to stations along the Central Line (orange icons in Figure 2) and the section of the Northern Line (yellow icons in Figure 2) which serves these areas.

In Figure 2, Simpson's diversity index (1-λ), also known as the Gini-Simpson index, is used to quantify the geographic diversity of stations mentioned

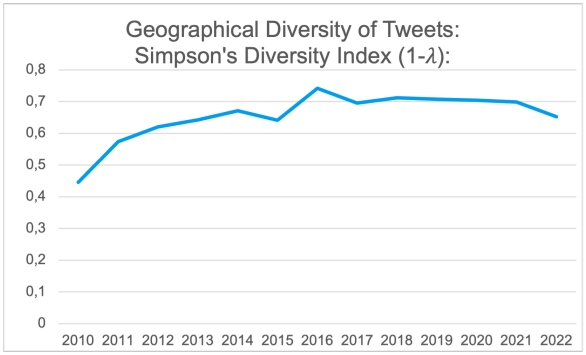


Figure 3: Simpson's diversity index and the distribution of tweets across Cape Town railway lines, 2010-2022 (n = 1755)

over time, categorised by five railway lines and major interchanges in Cape Town (n=1755). Simpson's diversity index is a measure from ecology, which has been used to identify geographic biases and changes in geographic diversity and equity over time (Liu et al., 2025). This measure does not represent the overall volume of tweets but rather their distribution across categories (railway lines).

Table 1 reports the classification analysis which checks the accuracy of automatically coded LLM annotations (output from Prompt 1) against manually coded categories which identify which Cape Town railway line is mentioned in each tweet (n=350).

Figure 4 maps the geolocations of the only four destination stations and lines mentioned in the synthetic tweets created with Prompt 4. Although the tweets were prompted using the name of 91 stations in the Cape Town area, Cape Town, Simonstown, Bellville and Muizenberg were the only stations mentioned in the synthesized tweets.

2 Discussion

2.1 Geographic bias in the Twitter Train Stories collection

The Twitter Train Stories collection displays a strong geographic bias towards major interchanges and places on and around the Southern Line. The core-periphery bias in favour of major interchanges in Figure 1 is to be expected, particularly given the radial design of the Cape Town system and the way in which commuters refer to the terminal station or interchange to clarify which train line they are discussing (e.g. "The Cape Town line" or "The train via Retreat").

Thus the major noteworthy bias in the sample is the geographical cluster of tweets referring to minor stations on the suburban Southern Line. This cluster makes up almost a third of the collection. Available figures for ticket sales and per line passenger figures suggest that this ratio is out of proportion to the level of usage on the line, since the Northern Line and the Central line served the majority of rail transit users in Cape Town at various stages in the period, apart from a short time in 2020.

Figures 1 shows that references to stations on the Central Line and Cape Flats Line are relatively sparse throughout the period.

Using the Simpson Diversity Index to measure mentions of stations on the various lines and the evenness of the spread of activity across the lines, shows the geographic diversity of the Twitter archive increasing from 2010-2013 as the app becomes more widely adopted (Figure 3). For the period 2013 -2020, the archive includes tweets about all four major lines and major interchanges as well as smaller rural lines. In 2022 the diversity index drops again to the lowest point since 2015.

The changes in Simpson's diversity index indicate an increasing level of geographic richness in the places represented as well as sustained inequality in distribution of placemaking practices in the Twitter Train Stories collection. While early adopters of Twitter from the suburban Southern and Northern Lines are over-represented in the early years, after 2013 the collection includes mentions of all of the railway lines, with the representation of tweets across all Cape Town lines gradually improving. By this measure, in 2016 the collection reaches maximum diversity. Although the volume of tweets was dropping at this stage, mentions are spread across all the lines.

The suspension of train services from 2020 makes little impact on the distribution of tweets until 2022, when the diversity index dips, reflecting an increasing proportion of tweets about the Southern Line.

Many public transit users would have been excluded from placemaking with Twitter (and thus from this study) by the need to use smartphones and mobile data. Interviews in 2016 with mobile phone users living below the poverty line in South Africa highlighted the extent to which the "mobile currencies" (airtime and data) are in short supply, often dependent on gifting between friends and relatives, and subject to careful, frugal management, buying data or airtime in very small quantities, and leaving data connections off except to check messages,

rather than the “always on” mode often assumed by app designers (de Lanerolle et al., 2017).

At the time of this study, Whatsapp and Facebook were used by the majority of SA Internet users, being considerably more accessible owing to zero rated access and the ability to buy special "bundle" or discounted data deals for them from the mobile networks. By contrast, less popular platforms such as Twitter, Instagram and Youtube were used by relatively small minorities.

Demographics of commuters aside, other possible reasons for the sustained focus on the Southern Line over the period requires a more careful analysis of the tweets in context than is possible in this paper.

2.2 Named Entity Recognition tools as digital placemakers

Commercial language models have been extensively critiqued for being resource-intensive and unsustainable technologies (Bender et al., 2021) which perpetuate epistemic and social marginalisation (Birhane, 2021). For this reason, open-source and low-resource tools were utilised in the first instance.

2.3 Open-source and low-resource tools

Named Entity Recognition for placenames was attempted with DBPedia, SpaCy, an open-source library for Natural Language Processing in Python, and with Mistral, an open-source EU-developed LLM. DBPedia failed the task of recognising a finite list of Cape Town stations, missing places such as Belhar, Diep River, Fisantekraal, Netreg, Nolungile, Nyanga, Nonkqubela and Unibell. SpaCy was more successful with station names, identifying multilingual or codeswitched variants (Koeberg Weg, Diep River), hashtags (#WaitingatRetreatagain) contractions (Lavis, Mplain), non-standard orthographies (e.g. blackenfell), colloquial placenames and some placenames with locative prefixes e.g eKapa, However, several common variants were not successfully identified (cpt, Khayelitja, eKhayelitsha, Kwa Langa) using these less resource intensive methods. These limitations have been documented in evaluations of existing tools for studies of multilingual discourse on social platforms in the South African context (Shibeshi, Zelalem et al., 2025).

The commercial LLM model Claude Opus 4-6 identified all of the above variants, failing only to identify more complex instances of placemaking, such as circumlocutions (“SA’s murder capital” for Nyanga), generic terms (“town” for Cape Town) or knowledge of the geographical context (The station at CPUT is known as Pentech).

During geocoding, careful disambiguation of several toponyms was required. This challenge can be understood in relation to the prevalence of colonial placemaking practices, where colonial places were named after places in the imperial centre, often in more than one colony. This need for disambiguation affected almost one in four station names in Cape Town.

LLMs are not designed for information retrieval, and RAG (Retrieval Augmented Generation) improves accuracy in informationally-motivated prompting (Moayeri et al., 2024). The presence of openly licensed geographical datasets and models are necessary for developing more accurate and representative RAG systems, and so this analysis suggests the continued importance of investing in the development of such repositories of local geographic information and ensuring that tools accommodate local translanguaging practices (Shibeshi, Zelalem et al., 2025).

2.4 Claude takes the Southern Line?

Identifying which railway line was under discussion proved more challenging for LLM annotation than identifying named places from a provided list. An Accuracy of 0.71 and Cohen’s Kappa of 0.66 are conventionally read as "substantial" agreement between automated and manual coding. Nonetheless, almost half of the roughly 29% of tweets which were coded differently to the manual benchmark were misclassified as the “Southern Line”.

The category with the weakest agreement, "Line not mentioned" (F1 = 0.49, $\kappa = 0.43$) highlights low recall, again misclassifying many ambiguous cases as the "Southern Line". In other words, where a specific line is not mentioned, the LLM over-infers the Southern Line.

Analysis of classification of the Southern Line shows very high recall (0.97) but low precision (0.63). The LLM over-applies this label, identifying nearly all of the tweets which actually did refer to the Southern Line while generating a large number of false positives from tweets about most of the other lines.

The Central Line (F1 = 0.81, $\kappa = 0.76$) is coded with high precision (0.92). There is a recall of only 0.73 for this category because 27% of tweets about the Central Line are incorrectly categorised,

two thirds of them as "Cape Flats Line". In these cases, the train line is confused with the “cape flats area”, such as in this example: The Central Line (F1 = 0.81, $\kappa = 0.76$) is coded with relatively high precision (0.92). There is nonetheless a recall of only 0.73 for this category because 27% of tweets about the Central Line are incorrectly categorised, two thirds of them annotated as the "Cape Flats Line". In these cases, a specific train line is used interchangeably with the “cape flats area”, as in this example:

> **Claude Opus 4-6:** cape flats line. reason: bonteheuwel station is located on the cape flats line, which runs through the cape flats area of cape town. this station is served by the cape flats line connecting cape town to chris hani (khayelitsha) and kapteinsklip (mitchells plain) via bonteheuwel.

The presence of a similar bias in the LLM coding as in the collection of social media posts does not imply that the preference of Twitter users for the Southern Line is directly encoded into the Anthropic model. Instead, both of these cases of geographic bias relate to the slightly different and often unpredictable ways that digital inequalities and socially hegemonic discourses about places in Cape Town play out in locative tweets on social media and in LLM training datasets respectively. Because of the way languages are shaped by power relations and power struggles in society, researchers will continue encountering such imbalances despite efforts at "debiasing".

2.5 Prototypical places in synthetic text

The geographical narrowness of LLM output synthetic text is mapped in Figure 2, showing a more distinct bias even than that found in the Twitter archive and seen in Figure 1. This LLM-specific form of geographical bias is the tendency to produce a limited set of default “prototypical places” (Liu et al., 2026) which leads to a highly restricted geographical repertoire in descriptions of places, identifying instances, or synthesized images representing specific geographic entities.

The severe geographical narrowness of output created by the model when synthesizing tweets about destinations and lines can be seen in the most prevalent digital defaults or “prototypical places” in the dataset of synthetic tweets, namely Cape Town (64%), the Southern Line (43%), and Simon’s Town (30%). Notably, these defaults are an exaggerated version of the biases also found in the social media data, as well as in the LLM annotation task. It is particularly surprising that Simon’s Town (a small outlying suburb) is mentioned in one in three tweets since the prompt mentioned “Monday 8am”, implying the context of an everyday commute.

Automated generation such as Prompt 4 is increasingly used to populate social media content and generate text for websites. Models which generate text featuring these geographically narrow sets of “prototypical” places in both text and images are thus likely to be shaping increasing proportions of online discourse.

Because commercial LLM models such as Claude Opus 4-6 are under the control of a handful of unaccountable “big tech” corporations, our responses need to go beyond attempts to “debias” or augment commercial tools. We also need to critically document the structural inequalities reflected when we use such algorithmic systems and interrogate the asymmetrical power dynamics involved in their use, along with the assumptions they embed regarding knowledge, justice, and technology itself (Birhane, 2021).

2.6 Limitations

The study should not be taken to represent all discourse about trains in Cape Town on Twitter/X, nor does it address people’s actual everyday placemaking practices, as they also involve complex interpersonal and cross-platform interactions. Both qualitative content analysis and qualitative, ethnographic and cross-platform multilingual studies are needed to explore how these and other biases might play out as people in Cape Town use various placemaking technologies in their everyday commutes and excursions across a range of transport modalities.

2.7 Conclusion

This study is an initial quantitative foray documenting placemaking practices and geographical and representational biases on Twitter over a particular period of time.

The project of describing historical unevenness is an encounter haunted by inequalities both old and new. Social platforms record differential access to participation in placemaking activities. GenerativeAI tools used to research inequities are them-

selves shaped by unequal participation in web platforms and by the everyday promotional activities associated with tourism and marketing. As a result, LLM embeddings encode hegemonic meanings and associations of place names, which can perpetuate biases and encode hegemonic meanings for places in GenerativeAI's defaults, both as a discriminatory technology (for content classification) and when used for synthesis of social media discourse.

Commodified public spaces like Twitter reflect and may exacerbate inequities in access to mobility in Cape Town. In the Twitter Train Stories collection, one in three tweets chosen at random from the sample as a whole is likely to feature the Southern Line. Despite gradual increases in equity, peaking in 2016 as Twitter adoption became more widespread, material struggles around spatial inequality and mobility in South Africa reshaped people's everyday mobility practices during the period of this study.

Digital inequality is a dynamic, changing and yet structural feature of the discourse on social platforms and of LLM discourse in South Africa. Along with other inequalities inscribed deeply in this society, digital inequality helps to reproduce the geographical biases documented in this study.

Furthermore, owing to South Africa's history of continued spatial inequalities, the geographical biases documented in this paper are tightly implicated with racial and gendered biases. The growing use of LLMs to generate text and images may already be leading to many unexpected forms of cultural, discursive and linguistic narrowing.

2.8 Acknowledgements

GIS support from Nicholas Lindenberg, University of Cape Town's Digital Library Services. Data analysis relied on 4CAT.uct.ac.za with support from Raymond Maclean and Sal Hagen (University of Amsterdam). Computations were performed using facilities provided by the University of Cape Town's ICTS High Performance Computing team: hpc.uct.ac.za – https://doi.org/10.5281/zenodo.10021612

2.9 List of references

References

D. Baeza-Yates. 2018. Bias on the web. *Communications of the ACM*, 61(6):54–61. Doi:10.1145/3209581.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, pages 610–623, New York, NY, USA. Association for Computing Machinery.

V. Bickford-Smith. 2009. Creating a city of the tourist imagination: The case of cape town, "the fairest cape of them all.". *Urban Studies*, 46(9):1763–1785.

Abeba Birhane. 2021. Algorithmic injustice: a relational ethics approach. *Patterns*, 2(2):100205.

T. E. Bosch. 2021. *Social media and everyday life in South Africa*. Routledge.

Nicholas Carah, Lauren Hayden, Maria-Gemma Brown, Daniel Angus, Aimee Brownbill, Kiah Hawker, Xue Ying Tan, Amy Dobson, and Brady Robards. 2024. Observing "tuned" advertising on digital platforms. *Internet Policy Review*, 13(2).

Chidambara. 2019. Walking the first/last mile to/from transit: Placemaking a key determinant. *Urban Planning*, 4(2):183–195.

A. P. Czeglédy. 2004. Getting around town: Transportation and the built environment in post-apartheid south africa. *City Society*, 16(2):63–92. Doi:10.1525/city.2004.16.2.63.

I. de Lanerolle, M. Walton, and A. Schoon. 2017. *Izolo: Mobile diaries of the less connected [Report]*. The Institute of Development Studies and Partner Organisations.

P. D. Figueroa, M. J. M. Agonos, and J. S. Labor. 2026. *Exploring digital placemaking of queer religiosity: Digital spaces as enabling environments for queer religious practices*. New Media Society. Doi:10.1177/14614448261433601.

M. Graham, B. Hogan, R. K. Straumann, and A. Medhat. 2014. Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers*, 104(4):746–764.

G. Halegoua and E. Polson. 2021. Exploring 'digital placemaking.' convergence (london. *England)*, 27(3):573–578. Doi:10.1177/13548565211014828.

H. Lefebvre. 1991. *The production of space (D. Nicholson-Smith, Trans*, volume 142. Blackwell, Oxford.

Z. Liu, K. Janowicz, G. Mai, S. Gao, and R. Zhu. 2026. Geographic bias and diversity in ai evaluation. In K. Janowicz, R. Zhu, G. Mai, S. Gao, Y. Hu, Z. Wang, L. Cai, and L. Bennett, editors, *Geography According to Foundation Models*. SAGE Publications. Doi:10.3233/FAIA260482.

Z. Liu, K. Janowicz, I. Majic, M. Shi, A. Fortacz, M. Karimi, G. Mai, and K. Currier. 2025. Operationalizing geographic diversity for the evaluation of ai-generated content. *Transactions in GIS*, 29(3):e70057. Doi:10.1111/tgis.70057.

MistralAI. 2025. Ollama - mistral.

M. Moayeri, E. Tabassi, and S. Feizi. 2024. Worldbench: Quantifying geographic disparities in llm factual recall. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24*, pages 1211–1228. Doi:10.1145/3630106.3658967.

Hellina Hailu Nigatu, John Canny, and Sarah Chasins. 2024. A Need Finding Study with Low-Resource Language Content Creators. In *Proceedings of the 4th African Human Computer Interaction Conference*, AfriCHI '23, pages 1–4, New York, NY, USA. Association for Computing Machinery.

Stijn Peeters and Sal Hagen. 2022. The 4CAT Capture and Analysis Toolkit: A Modular Tool for Transparent and Traceable Social Media Research. *Computational Communication Research*, 4(2):571–589. Type: Journal Article.

G. Pirie. 2013. Transport geography in south africa. *Journal of Transport Geography*, 31:312–314. Doi:10.1016/j.jtrangeo.2013.06.002.

Philip Resnik. 2025. Large language models are biased because they are large language models. *Computational Linguistics*, 51(3):885–906.

Shibeshi, Zelalem, Walton, Marion, Sibeko, Johannes, Mthembu, Vusi Maxwell, Evans, Henri-Count, Ngcungca, Nkazimlo, and Schoon, Alette. 2025. Multilingual vibes: Visualising linguistic resources and emoji in Southern African online discourse. *Journal of the Digital Humanities Association of Southern Africa (DHASA)*, 6(1).

Statistics South Africa. 2020. National Household Travel Survey 2020.

Ivan Turok, Justin Visagie, and Andreas Scheba. 2021. Social inequality and spatial segregation in cape town. In *Urban socio-economic segregation and income inequality: A global perspective*, pages 71–90. Springer International Publishing Cham.

M. Vanderschuren, R. Cameron, A. Newlands, and H. Schalekamp. 2021. Geographical modelling of transit deserts in cape town. *Sustainability*, 13(2):997. Doi:10.3390/su13020997.

M. Vanderschuren and T. Lane-Visser. 2025. Transport equity in south africa—how much progress was made over the last two decades? *African Transport Studies*, 3:100057. Doi:10.1016/j.aftran.2025.100057.

Marion Walton. 2026. Twitter train stories repository marionwalton/twitter_train_stories.

Yulin Yu, Xianglong Li, Tianyi Li, Paramveer S. Dhillon, and Daniel M. Romero. 2025. Demographic disparity in Wikipedia coverage: a global perspective. *EPJ Data Science*, 14(1):15.

D. Özkul. 2021. The algorithmic fix: Location intelligence, placemaking, and predictable futures. *Convergence*, 27(3):594–608. Doi:10.1177/13548565211005644.

A Appendix

Table 2: Input tweet and Claude Opus 4.6 annotations with $Prompt2$

Tweet text	**Nodes**
"we passed Langa and now heading to Cape Town"	Langa Station, Central Line; Cape Town Station, Interchange
"View of the mountain from Bonteheuwel"	Bonteheuwel Station, Central Line
"Just came across a 1997 piece still surviving the Cape Flatz train line."	NA, Cape Flats Line

Table 3: Tweet and manual annotations

Tweet text	**Line**
we passed Langa and now heading to Cape Town	Central Line
Crowd moving towards CPT train station. #CosatuMarch HeartFMNews"	No Line mentioned
CT students say their train was delayed on purpose to keep them from attending the meeting #CPUT #Fees2017	Central Line

A.1 Prompt 1

System prompt: You are a named entity recognition expert.User prompt: Given the following tweet about trains in Cape Town [text] Identify which of the following train lines from the list below are being discussed and return the line with a reason for the choice If the train line is at an interchange, and so could be one of several of the above options, return "line not mentioned" If the train line being discussed is not one of the above lines, return "other" If the tweet is not about stations, trains or train lines, return "na"

A.2 Prompt 2

System prompt: You are a named entity recognition expert.User prompt: Given the following tweet about trains in Cape Town [text] You are looking for station names in the tweet in order to categorise it by the train line. Search the text for any of the stations and train lines from the list that follows. Return each station name and its associated train line. If no station name is mentioned, return in csv format "NA" followed by the name of any train line mentioned If no station name or train line is mentioned in the text, return "NA", "NA".

A.3 Prompt 3

You are a geocoding expert. Provide the specific geographic co-ordinates and other details associated with the Cape Town train station named: [station name] train station

A.4 Prompt 4

System prompt: You are living in Cape Town and posting an update on Twitter/X at 8am on Monday morning. User prompt: You are currently at [station name] train station in Cape Town. Select a suitable destination where you are going today by train. Phrase a short tweet which will update everyone on Twitter/X about the experience of using the railway, the destination for your journey today and the railway line you will take to get there. Post only the text of the tweet.