The "Gowisha" Phenomenon: Vernacular Camouflage and Digital Coping Mechanisms Among South African Youth

Samuel Olaniran[1]*, Henri-Count Evans[2]†, Iginio Gagliardone[1]‡, Noma Gumede[1], Sobantu Dlamini[2]
Maya Loon[1]§, Kamohelo Nkobolo[1], Elison Shumba[3], Fritz Ngale Ilongo[2], Lindokuhle Dlamini[1]

[1]University of the Witwatersrand, South Africa, [2]University of Eswatini, Eswatini
[3]University of Johannesburg, South Africa

psalmuel35@gmail.com, henricount@gmail.com, iginio.gagliardone@wits.ac.za, nomandla.gumede1@gmail.com
sobeedlamini@gmail.com, maya.loon@wits.ac.za, kamnkobolo@gmail.com, elsonshumba@gmail.com
nfilongo@uniswa.sz, lindonkosie@gmail.com

Abstract

Digital expression of distress in the Global South often circumvents Western clinical vocabulary through vernacular substitution. This study examines *gowishing*, a South African idiom for severe distress, across TikTok, Instagram, and X. Drawing on a dataset of 78,538 posts and comments, including a census of 10,366 idiom items and 10,940 items coded by large language model coders, it analyses how users deploy code-switching, slang, and emoji as linguistic camouflage. In the census, 88.9% of uses convey distress. An English-centric rule treating comic delivery as evidence against distress erased about two-thirds of measured mental-health content in Sotho-Tswana languages, two-fifths in Nguni languages, and a fifth in English, but none in Afrikaans. Camouflage proves bidirectional and shows signs of commercial capture. The study extends Goffman's stigma theory through boyd's context collapse, updated for algorithmic platforms, and puts both in dialogue with Ubuntu-informed communal distress, treating their unresolved tension as its central contribution.

1 Introduction

When navigating acute psychosocial stressors, such as examination failure, retrenchment, or relationship breakdowns, South African users on Instagram, TikTok and X rarely utilise standard clinical terminology. Instead, expressions of distress frequently centre on *gowisha*, a vernacular idiom without formal recognition in clinical diagnostics. The word combines English *go* with Nguni verbal morphology, and users gloss it as "going through it": enduring severe stress, suffering, or anxiety (users' own etymologies vary). These posts frequently incorporate laughing-crying emojis, code-switching (e.g., isiZulu-English or Afrikaans-English), and comic framing, and often appear on the surface as mundane humour. This linguistic ambiguity functions as a communicative strategy tailored to networked environments. Platform affordances often collapse distinct social groups such as peers, family members, and prospective employers into a single, un-segmentable audience. In this context, vernacular substitution, code-switching, and humour allow users to articulate distress in a manner that remains legible to peers capable of offering support, while appearing ambiguous or innocuous to unintended readers. Consequently, this linguistic camouflage serves both to mitigate stigma and to signal for communal help.

Analysing this phenomenon requires engaging with overlapping theoretical frameworks. Goffman's (1959, 1963) dramaturgical account of stigma management provides a foundational vocabulary for understanding disclosure and concealment; however, it is predicated on bounded interactions rather than un-segmentable digital audiences. boyd's (2010) concept of context collapse addresses this structural shift, explaining how users navigate forced co-presence. Yet, both frameworks generally presuppose an individuated self-managing personal risk. In the South African context, these Western paradigms encounter Ubuntu-informed understandings of relational personhood, where distress is often conceptualised as a disruption in communal harmony requiring a collective response.

To examine how these communicative practices function in digital spaces, this study poses four

*Samuel Olaniran is a Researcher at the University of the Witwatersrand.

†Henri-Count Evans is a Senior Lecturer in the Department of Journalism and Mass Communication, University of Eswatini, and a Research Associate at the University of the Witwatersrand.

‡Iginio Gagliardone is a Professor in Media Studies and the inaugural SA-UK SARChI Chair in the Digital Humanities at the University of the Witwatersrand.

§Maya Loon is a postdoctoral fellow within the SA-UK SARChI in the Digital Humanities at the University of the Witwatersrand.

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.

research questions:

1. How do South Africans deploy code-switching between English, indigenous, and Afrikaans languages as a linguistic camouflage strategy when expressing psychological distress on social media?

2. What role does digital humour, and emoji-mediated affect specifically, play in mitigating stigma while simultaneously signalling a need for support?

3. In what ways does context collapse extend or complicate Goffman's dramaturgical account of stigma management in a platform environment his framework did not anticipate?

4. How does an Ubuntu-informed understanding of communal disclosure and relational personhood trouble the unmodified application of Western stigma theory to South African digital distress expression?

5. This study contributes in three ways. Theoretically, it reads context collapse, updated for algorithmic platforms, alongside Ubuntu-informed models of communal distress, to articulate the intersection of individual stigma management and relational personhood. Empirically, it offers an account of gowishing as a named linguistic phenomenon, resting on a census of 10,366 occurrences within a dataset of 78,538 items, 10,940 of which were coded for language and mental-health relatedness. Methodologically, it documents, and releases the prompts and codebook for, LLM-assisted coding of code-switched, multilingual, emoji-rich text, with its limits and a human validation step stated openly.

2 Literature Review and Theoretical Framework

2.1 Goffman, Stigma, and the Limits of Bounded Interaction

Goffman's (1959, 1963) dramaturgical model distinguishes a "front stage," where a curated self is presented to an audience, from a "back stage," where the performance relaxes, and yields a typology of stigma management: disclosure, concealment, information control, and solidarity-seeking. The typology has been applied to health-related self-disclosure online, including South African work on young women living with HIV (Ahn, 2026). It assumes, however, that the discloser can identify and partly select the audience. On Instagram, TikTok, and X a single post reaches family, employers, partners, and algorithmically curated strangers at once, so the framework needs structural modification to account for digital distress disclosure.

2.2 From Dramaturgy to Context Collapse

boyd's concept of context collapse, developed across her work on networked publics (boyd, 2010; boyd & Ellison, 2007) and elaborated by Marwick and boyd (2011, 2014), supplies the missing structural vocabulary. boyd (2010) identifies four affordances of networked publics, persistence, replicability, scalability, and searchability, that strip away the spatial, social, and temporal boundaries that ordinarily separate audiences. Where those boundaries collapse, the poster addresses an "imagined audience" (Marwick & boyd, 2011). Read through Goffman alone, slang or dark humour in a gowishing post resembles concealment of a distressed self. Read through context collapse, the same choices function as targeted audience segmentation within a single feed: a post layered with vernacular code and laughing-crying emoji is legible as disclosure to a peer and as a joke to an employer or older relative.

These affordances describe a follower-graph era of social media, and they need updating for the platforms studied here. On TikTok, who sees a post is decided less by who follows the poster than by the platform's content graph (Schoon & Bosch, 2024). Visibility is algorithmic as well as social, and users' awareness of it, and of biased moderation, shapes what they say (Marwick, 2023; Owens, 2026). We therefore treat collapse as two regimes that our data bring apart. Social collapse operates on X and Instagram, where the idiom's census life is concentrated. Algorithmic collapse operates in TikTok comment threads, where disclosure appears among strangers whom an algorithm, not the poster, assembled. Camouflage in the second setting addresses human readers and may also pass unnoticed by automated moderation, a hypothesis our data cannot test.

2.3 Ubuntu and Afrocentric Models of Distress

Context collapse addresses the structural limits of Goffman's framework, but an epistemic limita-

tion remains: classical stigma theory presupposes an individuated self protecting personal standing. Ubuntu philosophy instead conceptualises personhood as constitutively relational (*umuntu ngumuntu ngabantu*, a person is a person through other persons). Here distress is less a private attribute than a disruption of relational harmony that calls for a communal response (Nwoye, 2015; Ratele et al., 2018; Van Dyk & Nefale, 2005). Ubuntu-informed social work integrates biopsychosocial models with communal orientations (Mugumbate & Chereni, 2020; Mugumbate & Nyanguru, 2013), and empirical studies link communal self-construal and extended-family care to culturally rooted coping and reduced distress (Constantine et al., 2002; Eaton & Louw, 2000). Through this lens, vernacular camouflage may be not only impression management but an invocation of communal support that avoids the individualising frame of Western clinical disclosure.

Goffman's impression management also covers the reverse case, in which fortune rather than misfortune is the information to be managed. In Southern African anthropology, visible prosperity attracts suspicion and obligation, from the occult economies of the postcolony (Comaroff & Comaroff, 1999; Ashforth, 2005; Niehaus, 2001) to the claims that kin make on earners, discussed as "black tax" (Mhlongo, 2019). Under such conditions, performing hardship can protect a poster from envy and from demands on their resources, while the Ubuntu ethic of mutual obligation makes those demands legitimate. Gowisha is available for this work because it reads as distress and as a joke at once. We therefore expect camouflage to run in both directions, concealing distress and concealing wellbeing, and we examine the second in the results section.

2.4 Code-Switching, Vernacular Registers, and South African Youth Language

Myers-Scotton's (1993) Markedness Model treats code-switching as a socially meaningful choice of marked or unmarked codes that negotiates the rights and obligations of an exchange. South African urban youth languages such as Tsotsitaal have long served as performative registers of status and generational identity (Brookes, 2014, 2020), built from slang lexicons inserted into the grammatical frame of a local language (Mesthrie & Hurst, 2013). These functions persist online, shaping relatability and audience trust (Botha & Coetzee-Van Rooy, 2021; Mushaathoni & Dondolo, 2025). *Gowishing* belongs to this tradition, adapted to the audience dynamics of networked platforms.

Southern multilingualism scholarship adds a caution about the unit of analysis. Work on translanguaging argues that named, bounded languages are a colonial inheritance and that speakers draw on a single repertoire (Pennycook & Makoni, 2020; McKinney, 2024), a view borne out by South African youth and hip-hop registers that mix languages freely (Brookes, 2014; Haupt, 2017). We keep language labels because coders and NLP tools need them, but we treat vocabulary shared across Nguni and Sotho-Tswana languages as evidence of a repertoire, not as an error. Recent work visualises linguistic resources and emoji in Southern African online discourse (Shibeshi et al., 2025), and open South African language corpora and models are emerging (Lombard et al., 2026), but we know of none that targets mental-health text (Eiselen & Puttkammer, 2014; Nekoto et al., 2020).

2.5 Digital Humour, Emoji, and the Aesthetics of Camouflaged Distress

Self-deprecating humour lets youth publics express and commiserate over everyday distress (Ask & Abidin, 2018), and TikTok users narrate grief through dark humour (Eriksson Krutrök, 2024). Emoji are calibrated signals rather than decoration: users choose masking emoji for less intimate audiences to manage emotion (Liu, 2023). Because emoji are interpretively ambiguous (Miller et al., 2016) and reciprocal use fosters warmth (Cavalheiro et al., 2024), they suit a register that needs dual legibility, cueing the in-group while masking severity from outsiders.

2.6 The Vernacular Camouflage Framework

The Vernacular Camouflage Framework proposed in this paper operates at three levels. First, following boyd, camouflage operates as differentiated address under conditions of forced audience copresence, segmenting the feed through linguistic code. Second, following Goffman, camouflage retains a concealment function, protecting the user from the relational risks of explicit clinical disclosure in stigma-sensitive contexts. Third, following Ubuntu-informed scholarship, camouflage simultaneously functions as a communally legible distress signal, inviting relational response from an in-group while declining Western clinical vocabulary. Methodologically, the paper holds this theoretical

tension open and tests how the three operations manifest empirically.

3 Methodology

3.1 Research Design

The study uses a sequential mixed-methods design. A quantitative layer establishes the scale, distribution, and composition of vernacular distress discourse. A qualitative layer, conducted as sociocognitive critical discourse analysis (van Dijk, 2008, 2015), examines how camouflage is discursively accomplished. Throughout, *LLM coder* means a large language model assigning item-level labels, *coded* means labelled by an LLM coder, and *read* is reserved for reading by a researcher. Table 4 (Appendix D) states who or what performed each task, and whether the step is deterministic or probabilistic.

3.2 Data Collection and Dataset

Data were collected between May and July 2026 with Socius, a self-hosted multi-platform system built on browser automation, network response interception, and PostgreSQL storage. Collection targeted South African mental-health discourse through an iteratively developed query strategy combining clinical terminology, vernacular distress expressions (*gowisha*, *umgowo*, *kuningi*, and variants), hashtags, and platform search parameters. South African attribution was made at collection time by an LLM relevance judge on each item's text and metadata. This attribution was not audited against a researcher-coded sample, so some items from outside South Africa may remain (see Limitations). YouTube items were collected but excluded from analysis.

The analysis dataset comprises 78,538 items: 23,005 posts and 55,533 comments from TikTok (42,248; 53.8%), Instagram (18,653; 23.7%), and X (17,637; 22.5%). Each item carries one of four collection cohorts: institutional and NGO accounts (6,284), local vernacular stressors (10,006), acute distress and crisis (8,148), and personal narrative (54,100). TikTok, Instagram, and X were chosen because they are public, text-bearing, comment-enabled platforms widely used in South Africa (Kemp, 2025), and because their content could be collected at scale. Closed channels such as WhatsApp were excluded because their content is not public. Every comment is linked to its parent post by a platform identifier, which allows thread-context analysis.

Most items date from 2025 or 2026 (70.2%), and the idiom census is 80.2% X; temporal coverage and platform composition are detailed in Appendices B and E.

3.3 Sampling Design and the Coded Sample

Items for coding were drawn by stratified sampling. Posts were stratified by platform, indigenous-language candidacy, and length class (24 strata), and comments were stratified the same way (25 strata). Candidacy came from a deliberately over-inclusive orthographic screen used only for stratification, never for labelling. Allocation was proportional, with a floor of 250 items per stratum. Comments are short (about 60% under four words) and heavily TikTok, so the floors lift small but analytically important strata.

The realised coded sample comprises 10,940 items: 8,968 posts (82.0%) and 1,972 comments (18.0%), posted between March 2012 and July 2026. Because strata are sampled with floors, the unweighted shares reported in the results describe the composition of this sample and are not population prevalence. Two further groups complete the data: the 10,366 items of the idiom census, of which 2,996 also fall in the coded sample, and the remaining 60,228 items, which were not coded for language and serve as sampling frame.

3.4 Why Automated Language Identification was Set Aside

Three automated approaches were evaluated and rejected, and the way each failed is itself informative. fastText lid.176 (Joulin et al., 2017), the most widely used language-identification model in computational social science, returned no South African indigenous language at all, because it has no isiZulu, isiXhosa, Sesotho, Setswana, or Sepedi labels. An English-anchored mental-health keyword lexicon excluded 94.6% of 30,956 items later confirmed to contain indigenous-language material. An LLM-generated language field assigned 2,516 of 2,753 Nguni-language posts to an uninformative "mixed" category. These are representational absences, not accuracy shortfalls: a tool with no label for isiZulu does not make errors on isiZulu, it cannot see it. This is the English-centric infrastructure gap of vernacular camouflage framework appearing inside the study's own toolchain, and it is why the coding was done by LLM coders instructed in language-specific diagnostics.

3.5 Coding Procedure

Items were coded in 42 chunk files (24 of posts, 18 of comments) of 93 to 257 items each (median 222), plus a further set of 3,343 posts coded in a separate Claude Sonnet pass. The coders were Claude Sonnet 5 (Anthropic), run between May and June 2026 in an agentic environment at the provider's default sampling settings, with temperature set by the researchers, operating under instructions that prohibited delegation, scripted pattern-matching, and extrapolation, and required incremental, verified output. LLM annotation can rival human annotation on some tasks while varying with prompt and model (Gilardi et al., 2023; Ziems et al., 2024; Reiss, 2023), which is why we report verification and run-to-run stability below and make no claim of human-level reliability.

Coders assigned primary language (the eleven official languages, Other, or Unclassifiable), all languages present (the code-switching field), code-switching status, mental-health relatedness, and comic or weeping register. Because the Nguni languages are agglutinative and closely related, coders received morphological diagnostics rather than word lists (for example isiXhosa *ndi-* against isiZulu *ngi-*; *intliziyo* against *inhliziyo*; *wam/yam* against *wami/yami*), with comparable diagnostics for the Sotho-Tswana group. Each coder also returned a structured report of counts, judgments on negated forms, verbatim metalinguistic quotations, borderline cases, and items flagged for researcher decision (collective distress, bereavement, gender-based violence, culture-bound idioms). The 138 reports, about 48,500 words, were read by the researchers and are the source of candidate extracts for the close readings. They are preserved unedited as a qualitative audit corpus.

3.6 Coding-Scheme Revisions

Four rules were revised during the study. Each earlier rule was an analyst assumption, each was falsified by the corpus, and each revision is reported because it bears on the paper's central claim about imported analytic frameworks.

Humour. The first scheme treated laughing emoji as evidence against distress. The principal investigator, a South African researcher, established that this inverted the phenomenon: *umgowo*, *ukugowa*, and *gowisha* denote going through a great deal, and are characteristically delivered in a comic register. Humour here expresses and endures distress, and does not negate it. Correcting the rule reclassified 1,616 coded items as mental-health related, raising the share among coded items from 49.4% to 64.1%.

Ambiguous Nguni. Early chunks defaulted Nguni material with no diagnostic marker to isiZulu, the largest Nguni language. From post chunk 19 and comment chunk 13 onward, such material was coded Unclassifiable. In the dataset, 8,968 items (82.0%) were coded under the default regime and 1,972 (18.0%) under the strict regime. Because regime is confounded with chunk order and record type, contrasts between the two are indicative only.

Negation. Coders were first told that negated forms such as *Angigowi* ("I am not gowishing") indicate no distress. Close reading showed negation to be frequently ironic, defensive, or escalatory, and the rule was replaced by a context-sensitive typology.

Thread context. Comments were first coded in isolation. Reading threads end to end showed that replies such as short congratulations under a recovery disclosure had been coded as not mental-health, because alone they read as generic praise. Sampled comments with a parent post were therefore re-coded with the parent visible, producing context-aware mental-health labels and a ten-category comment role (for example comfort, shared disclosure, advice, congratulation, joke, spam). Where a TikTok caption was only hashtags, the parent was classified from its replies only when several independent replies converged, and every such inference is recorded.

3.7 Verification and Stability

Verification (deterministic). Coder self-reports were not trusted. Every output file was verified against its source for row count, identifier-set equality, stray identifiers, and label validity, and a text-based check scored each label against orthographic evidence in its own source text to detect block misalignment, where labels drift from their identifiers while row counts still reconcile. Across 95 coding chunks in the full workflow, these checks caught reporting errors, one chunk substituting scripted pattern-matching for reading, and nine dropped or misaligned row blocks. Contaminated output was isolated and re-produced, and 52 of 53 re-coded chunks then matched their sources exactly (Appendix C).

Stability. Across fourteen census chunks coded by independent coders, with no target rate disclosed, distress-rate estimates fell within a five-point band (85.6% to 91.6%). This shows stability across runs, not validity. Because LLM coders are stochastic and not exchangeable with human coders, we do not treat agreement among them as intercoder reliability in the sense of Krippendorff (2004) or O'Connor and Joffe (2020).

Human judgement. No human-coded validation sample was drawn for this paper, so we report no agreement statistics against human coders and do not present the LLM labels as validated measures. Human judgement entered the pipeline at three other points. Researchers read all 138 coder reports and made every rule revision in Section 3.6; the humour revision, initiated by the principal investigator, is itself evidence that the LLM labels diverged from first-language judgement in a systematic and correctable way. Researchers read every extract used in Section 4.5 within its full thread. Glosses of quoted extracts were checked with first-language speakers (Section 3.9). Headline shares are therefore LLM-coded estimates whose structural integrity and run-to-run stability are established, but whose agreement with first-language coders at scale is untested. A stratified validation sample coded by first-language speakers of Nguni and Sotho-Tswana languages is the first task of the follow-up study.

3.8 The Idiom Census

The census comprises 10,366 items (5,942 posts, 4,424 comments) containing the *gowisha* family (*umgowo*, *ukugowa*, *gowa*, *gowisha*, *ngiyagowisha*, *ndiyagowa*, and inflectional variants). Each was coded by LLM coders on three separate fields.

Context is the topical setting of the use: distress disclosure, entertainment or television, relationships, work and money, ambient slang, sport, alcohol and drinking, or other.

Real distress is a yes/no judgment, made in every context, on whether the use conveys the distress of a real person. A use in an entertainment context can convey a viewer's real distress, and a small share of uses classed as distress disclosure do not (1.4%).

Experiencer is whose distress is conveyed: self, another real person, a collective, a fictional character, or no one. It is not the addressee of the post.

Purely referential uses (brand and event names, song titles, scripted characters) were coded as non-distress. Real participants in reality television were treated as real persons. Only the 2,996 census items that also fall in the coded sample carry language labels.

3.9 Critical Discourse Analysis

Sociocognitive critical discourse analysis treats discourse as mediated by mental models and context models, the participants' subjective representations of the communicative situation, including who the audience is (van Dijk, 2008, 2015). We use it to ask how a poster's model of a collapsed audience is indexed through code choice, register, and emoji. Candidate extracts came from the coder reports, the thread-context re-coding, and census items flagged for researcher decision. Researchers read each extract with its full thread and analysed three levels: linguistic form (morphology, code, emoji placement), audience construction (addressee, referent, and overhearer), and the action performed (disclosure, appeal, or witnessing). Interpretations were checked with first-language speakers. Because LLM reports guided candidate selection, extracts reflect what the coders surfaced, so we use them to show how patterns work, not how common they are. Prevalence claims rest on counts alone.

3.10 Analysis and Estimation

Shares reported for the coded sample are unweighted, describe its composition, and carry Wilson 95% confidence intervals. They are not population prevalence. Contrasts that do not depend on weights, such as default against strict regime and comic against weeping register within the census, are reported within sample. Prevalence is never computed against hashtag-matched volume, because mental-health hashtags are attached to unrelated dance, comedy, and commercial content to gain reach. Headline figures are reported with and without the institutional cohort.

4 Results

4.1 The Linguistic Landscape

The 10,940 coded items come from three differently built sets: 5,625 posts and 1,972 comments from the chunk files, and 3,343 posts from the separate Sonnet set. Pooled, English is the primary language of 68.2% of items (95% CI 67.4 to 69.1), followed by Afrikaans (13.0%), isiZulu (7.2%), isiXhosa (5.3%), Setswana (1.2%), and Sesotho

(0.7%). Sepedi, siSwati, Tshivenda, Xitsonga, and isiNdebele together account for 0.2%, Other for 0.4%, and 3.8% were unclassifiable. The sets differ (Table 5): comments are 17.5% unclassifiable and the Sonnet set is 25.8% Afrikaans. The pooled shares therefore describe this sample and are not population prevalence.

The shared-vocabulary register is evidenced mainly by the coders' reports and by recurring items such as *wena*, *kodwa*, *kubi*, *sisi*, and *hayi* (Nguni) and *batho*, *ke*, and *mara* (Sotho-Tswana), which carry no orthographic marker of one language. Quantitatively we can give only a floor. Under the strict regime, 35 of 1,972 items (1.8%) were left unclassified because of Nguni or Sotho-Tswana ambiguity, and 45 more (2.3%) because text was unresolved. The floor is low because coders recorded a reason only when they abstained. The regimes also differ in ways consistent with general conservatism, not isiZulu-specific over-attribution: from default to strict, isiZulu falls from 7.7% to 4.6%, but Afrikaans, whose diagnostics the rule never touched, falls from 14.2% to 7.9%, and Unclassifiable rises from 2.6% to 9.2%. We read the shared vocabulary as a signature of a supra-ethnic register that is legible to an in-group and hard for classifiers and outsiders to attribute.

4.2 The Idiom at Scale and the Semantic Commons

The census holds 10,366 items (5,942 posts, 4,424 comments). It is 80.2% X, 11.3% Instagram, and 8.4% TikTok, and 56.8% of its items date from 2025 or 2026. Of the census items, 88.9% (9,214) convey the distress of a real person. Distress disclosure is the context of 61.4% of uses, followed by entertainment and television commentary (12.7%), relationships (6.5%), and work or financial stressors (5.6%). The experiencer is the poster in 43.5% of uses, another real person in 25.3%, and a collective such as a family, fandom, or community in 17.4%. Nearly a fifth of the idiom's work therefore concerns a group, which suggests relational witnessing as well as individual disclosure.

The idiom is a semantic commons that users maintain. Participants define it ("Mgowo means to GO THROUGH!!!!"), etymologise it ("it is to go but add a Xhosa syllable"), and police its boundaries against outsiders ("stop jumping on the 'gowisha' train… translating izinto"). This policing shows metapragmatic awareness that the register's value depends on restricted legibility.

4.3 Humour as a Constitutive Register

Comic delivery is the unmarked case for this vernacular, not a mitigation of it. It is unevenly distributed across language groups (Table 6, Appendix F). Among census items coded for language, laughing markers accompany 40.3% of Sesotho and 31.4% of Setswana uses but 11.8% of isiXhosa uses, which lean towards weeping markers (41.6%). The Sotho-Tswana groups are small (67 and 118 items), and their intervals are wide, but the contrast with isiXhosa is clear.

Users articulate five functions for comic framing. It serves expression ("Just realized gore ka gowisha and the only thing i can do about that is to make fun of my current situation"), endurance ("Keep em jokes coming when ndi gowisha please. They keep me insane [sane]"), and concealment ("personally I jst post funny reels on my statuses … Soze ndibonise abantu ukuba ndiya gowa NEVER", 'I will NEVER show people that I am going through it'). It also serves substitution ("Ave ngiwuziba umgowo bakithi … I don't even cry anymore", 'I just ignore the umgowo, my people') and regulation ("i now need to watch a funny video fast fast so l can feel better otherwise ndiyagowisha"). Some users treat the loss of humour as a warning sign: "I have subtle sense of humour, if I lose it, just know that ka gowisha". This is an emic claim and we do not test it as a clinical indicator.

The first coding rule treated laughing emoji as evidence against distress, and it did not fail evenly (Table 1). It missed 68% of the corrected mental-health content in Sotho-Tswana languages (95% CI 62 to 74), 41% in Nguni languages, and 18% in English, and none in Afrikaans. Correcting it reclassified 1,616 items. An English-centric rule thus erased most of the signal where comic delivery is most common.

4.4 The Temporal Grammar of Negation

In 74% of cases, negated forms of the idiom convey distress through recognisable subtypes. These include escalation by substitution ("we're not gowishing we're fokofing") and self-contradictory reversal ("andigowi mna 😭 ndya gowisha", 'I'm not going through it 😭 I am going through it'). The register also marks distress onset and duration ("kunini ngithi yhoo", 'how long have I been saying yhoo'), exit ("umgowo uphelile"), and renewal ("After umgowo there's that feeling of renewal").

Language group	n	MH-related, first rule %	Corrected rule %	Missed by first rule % (95% CI)
English	7,466	57	70	18 (17 to 19)
Afrikaans	1,425	25	26	0 (0 to 2)
Nguni (isiZulu, isiXhosa, siSwati, isiNdebele)	1,367	48	81	41 (38 to 44)
Sotho-Tswana (Sesotho, Setswana, Sepedi)	222	29	91	68 (62 to 74)

Table 1: Mental-health content missed by the first coding rule, by primary language group (n = 10,940).

4.4 Close Readings: Operations of the Camouflage

Bidirectionality. The camouflage does not only mask distress; it can also mask wellbeing to manage envy. One user explains: "Ankere bare we're bragging when we do, so we just pretend like siya 'gowisha' 😂😂😂" ('because people say we are bragging... we perform gowishing instead'). The poster expects good fortune to be taken as boasting, and chooses a register in which hardship is the safe thing to display. One explicit statement shows that the practice exists, but our data do not show how common it is. It connects to the obligations and suspicions discussed under the models of distress.

Simultaneous operations. A single post can conceal, segment its audience, and appeal for communal help at once. One post reads: "Cela ningiThandazela ndiyaGowisha kaloku. Abantu bakwatele ukuthi angifuni ukubapha iMali yam. Kunini ngithi yhoo" ('Please pray for me, I am gowishing. People are angry that I won't give them my money. How long have I been saying yhoo'). The opening is a plural imperative addressed to a praying public, which makes the appeal communal (Ubuntu). The poster discloses a family stressor (Goffman), but only through the anonymous category *abantu*, 'people', with no name or tag. The angry relatives are the topic of the post, not its addressee. Code does not keep them out, because they probably share the language and can read it. What the post does is deny them the position of addressee and give the poster deniability about who is meant. Segmentation here works by address, not by language alone (boyd).

Ambient disclosure. Re-coding comments with the parent post visible more than doubled measured mental-health content, from 5.4% to 11.7%, which indicates that disclosure often occurs in threads unrelated to mental health. A user may write "we are hungry shem 😭😭" under a student check-in, or a minimal vernacular unit such as "Kuningi" ('it is much') under a comedy skit. In such threads the audience was assembled by the platform, and the poster treats it as one that can respond.

Communal reception. In mental-health threads, comfort (26.6%) and reciprocal disclosure (25.6%) form most replies. In some threads, strangers also point posters towards formal care, which is the Ubuntu-informed pathway to the clinic that the framework anticipates, although we cannot say how often it succeeds.

5 Discussion

5.1 Code as Address and the Relational Target

The findings demonstrate that code-switching in *gowishing* posts operates as a highly specific form of audience address. The idiom itself is fossilized code-switching, an English verb inflected with Nguni verbal morphology. By building a supra-ethnic register from shared Nguni and Sotho-Tswana vocabulary, users create a signal that is legible to a localised, generational in-group while remaining opaque to outsiders. If the idiom functioned solely as an individualistic stigma-management device, its experiencers would be overwhelmingly the poster. Instead, in about two-fifths of uses the experiencer is someone other than the poster (25.3% another real person, 17.4% a collective). This is about whose distress is conveyed, not who is addressed, and it is consistent with a relational reading of the idiom.

5.2 The Functions and Limits of Comic Framing

Comic delivery is the register's unmarked baseline, not a mitigation of distress. The corpus differentiates this dark humour beyond generic coping categories into expression, endurance, concealment, substitution, and affect regulation, and participants' own theory inverts the usual diagnostic: the *absence* of comic framing is what they read as deterioration. Its distribution differs by language group (Table 6), so assumptions about humour and vulnerability cannot be applied uniformly even within one national context. The first coding rule shows the cost: it missed two-thirds of the corrected signal

in Sotho-Tswana languages and none in Afrikaans (Table 1). Users' own view that the loss of humour signals deterioration is an emic account, and we do not test it as a clinical marker.

5.3 Mapping the Theoretical Boundaries: Stigma, Collapse, and Community

The Vernacular Camouflage Framework does not force a choice between dramaturgical stigma management, context collapse, and Ubuntu-informed relational personhood; the data map where each operates. Goffman's (1963) framework governs specific relational risks: the masking quotes ("behind the camera the heart is hurting") defend a "back stage" from the collapsed audience, and impression management describes the register's bidirectionality, in which fortune is camouflaged as distress to manage envy, consistent with work on occult economies (Ashforth, 2005; Comaroff & Comaroff, 1999; Niehaus, 2001).

The corpus suggests that context collapse can be an affordance as well as a hazard. Minimal vernacular units dropped into unrelated threads imply that some users treat an undifferentiated public as one that may respond, and Ubuntu supplies the expectation of obligation. We hold this as an interpretation, since the thread-level evidence is limited to the comments re-coded with their parents. Our two collapse regimes also matter here: the socially collapsed audiences of X and Instagram and the algorithmically assembled audiences of TikTok comment threads need not work the same way.

5.4 The Wasting Asset of Vernacular Camouflage

The efficacy of vernacular camouflage relies entirely on restricted legibility. The corpus documents that this restriction is actively eroding. Brands now routinely enter the collapsed audience, deploying the idiom to sell products, while users voice genuine occupational distress to corporations that reward vernacular performance with engagement. Because its mainstreaming and commercial capture are already underway, the specific *gowisha* camouflage is a wasting asset. However, the underlying communicative strategy, encoding suffering so that only the appropriate in-group can interpret it, will likely persist through the continuous morphological innovation typical of South African youth languages.

Platform interface features can undo restricted legibility. Automated translation of comments can expose vernacular posts to outsiders, and can mistranslate them in ways that have caused real harm elsewhere (Hern, 2017), so camouflage cannot assume a fixed audience. Making the register legible to computation is also double-edged: the same tools that could direct support to people in distress could serve surveillance or moderation without consent.

6 Conclusion

Young South Africans make distress speakable on public platforms by encoding it in a vernacular that some readers can decode and others will take as a joke. In a dataset of 78,538 posts and comments, most uses of *gowisha* convey a real person's distress. The register works through fossilised code-switching, comic delivery, and negation that usually still conveys suffering. Theoretically, the evidence supports holding three accounts together instead of choosing one. Goffman explains why distress is concealed from particular readers. Context collapse explains why code must do the work that audience selection cannot. Ubuntu explains why the same post can summon a communal response. Collapse also takes different forms here: social on X and Instagram, where the idiom is most visible, and algorithmic in TikTok comment threads.

The practical lesson concerns measurement. A coding rule that treated laughter as evidence against distress missed about two-thirds of the mental-health content in Sotho-Tswana languages, and none in Afrikaans. English-centric tools therefore fail unevenly, and worst for some languages. Tools for detecting distress need validation with first-language speakers. Whether such detection should exist at all is a question for the communities concerned, because the register's protection lies in its limited legibility. These findings rest on LLM-coded text from three platforms. Human validation, ethnographic work, and analysis of audio and video are the next steps. Because the idiom is already being taken up commercially, the strategy of encoding distress for an in-group may outlast the word.

Limitations

TikTok supplies 53.8% of the dataset, but 96% of that is comments; it is 14.7% of the coded sample and 8.4% of the idiom census. Our evidence for the idiom comes mostly from X and Instagram text, while TikTok is audiovisual, and a text analysis

misses disclosures in video and audio. The coded sample is stratified, so its unweighted shares describe composition, not prevalence, and three coded sets were built differently (Table 5, Appendix E). Platform search favours recent items, which is why most items date from 2025 and 2026. Labels were assigned by LLM coders, not humans. Their output was verified structurally and is stable across runs, but it was not validated against first-language human coders, and South African attribution at collection was not audited. Glosses remain open to correction by speakers of isiZulu, isiXhosa, Sesotho, and Setswana. Boundary decisions on collective distress, bereavement, and culture-bound idioms (for example *ukuthwasa*) remain under researcher review. Discourse analysis of text also has limits. Ethnographic and observational methods would add what posts cannot show, such as how posters understand their audiences and how interface features such as automated translation change who can read them (Schoon et al., 2020; Bosch, 2022). Our claims about audience segmentation are inferences from posts, not observations of readers.

Ethical Considerations

The study received clearance from the Wits University Human Research Ethics Committee. All data were publicly accessible at collection, in line with the AoIR guidelines, which treat anonymisation and quotation as case-by-case decisions (franzke et al., 2020). We did not contact posters or attempt to identify them. Author handles were used only to identify institutional accounts and are removed from the released dataset.

The register studied here is itself a collective privacy practice: users encode distress so that only an in-group can read it. Our analysis makes that encoding legible to a wider audience, which is a harm the paper must own. We quote verbatim only where linguistic form is the datum, because paraphrase would erase the code, the emoji, and the self-contradiction under analysis. Verbatim text can still be found by searching platforms, so quotes disclosing bereavement, abuse, or health conditions are lightly altered, and handles, family names, and other identifying details are removed, including those embedded in quoted text. Residual re-identification risk remains, and we do not claim otherwise.

Item text was processed by a third-party LLM service under retention, training use and was sent without author handles. LLM reading of vernacular distress at scale is the institutional reading the register is built to evade, and the same capability could serve surveillance or moderation without consent. We therefore used LLM labels only for aggregate analysis, never to flag, contact, or make decisions about any account, and we do not release a trained detector. LLM labels are fallible, as the humour-rule error shows. The tension between detection and protection is discussed in the findings.

The codebook, verification scripts, and aggregate tables will be deposited in the SADiLaR repository (https://sadilar.org). The released dataset has no author handles, replaces platform identifiers with hashed identifiers (Instagram post identifiers otherwise embed the poster's numeric user identifier), and reports dates to the month. Item text is available only under a data-use agreement, with @mentions and links masked, because verbatim text is searchable. The coder reports contain verbatim quotations and are treated the same way. Researchers may request access via email to the authors or by contacting SADiLaR. Platform terms limit redistribution of content.

References

Ahn, S. (2026). Digital commentary activities for managing spoiled identities among young South African women and girls (16–24) living with HIV: A sociologically informed content and thematic analysis. *Digital Health*. https://doi.org/10.1177/20552076261433085

Ashforth, A. (2005). *Witchcraft, violence, and democracy in South Africa*. University of Chicago Press.

Ask, K., & Abidin, C. (2018). My life is a mess: Self-deprecating relatability and collective identities in the memification of student issues. *Information, Communication & Society*, 21(6), 834–850. https://doi.org/10.1080/1369118X.2018.1437204

Bosch, T. (2022). Decolonizing Digital Methods. *Communication Theory*, 32(2), 298 to 302. https://doi.org/10.1093/ct/qtac005

Botha, W., & Coetzee-Van Rooy, S. (2021). Language mixing in the social network of a South African student. *World Englishes*. https://doi.org/10.1111/weng.12473

boyd, d. (2010). Social network sites as networked publics: Affordances, dynamics, and implications. In Z. Papacharissi (Ed.), *A networked self: Identity, community, and culture on social network sites* (pp. 39–58). Routledge.

boyd, d., & Ellison, N. B. (2007). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13(1), 210–230.

Brookes, H. (2014). Urban youth languages in South Africa: A case study of Tsotsitaal in a South African township. *Anthropological Linguistics*, 56(3–4), 356–388.

Brookes, H. (2020). Youth language in South Africa: The role of English in South African tsotsitaals. In R. Hickey (Ed.), *English in multilingual South Africa* (pp. 176–195). Cambridge University Press.

Cavalheiro, B. P., Prada, M., & Rodrigues, D. L. (2024). Examining the effects of reciprocal emoji use on interpersonal and communication outcomes. *Journal of Social and Personal Relationships*, 41(8), 2147–2168. https://doi.org/10.1177/02654075231219032

Comaroff, J., & Comaroff, J. L. (1999). Occult economies and the violence of abstraction: Notes from the South African postcolony. *American Ethnologist*, 26(2), 279–303. https://doi.org/10.1525/ae.1999.26.2.279

Constantine, M. G., Donnelly, P. C., & Myers, L. J. (2002). Collective self-esteem and Africultural coping styles in African American adolescents. *Journal of Black Studies*, 32(6), 698–710. https://doi.org/10.1177/00234702032006004

Eaton, L., & Louw, J. (2000). Culture and self in South Africa: Individualism–collectivism predictions. *The Journal of Social Psychology*, 140(2), 210–217. https://doi.org/10.1080/00224540009600461

Eiselen, R., & Puttkammer, M. (2014). Developing text resources for ten South African languages. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)* (pp. 3698–3703). European Language Resources Association.

Eriksson Krutrök, M. (2024). TikTok memes of grief: Playfulness and dark humour in youth digital grief narratives. In C. Traher & L. J. Breen (Eds.), *The Routledge international handbook of child and adolescent grief in contemporary contexts* (pp. 286–294). Routledge. https://doi.org/10.4324/9781003221692-30

Franzke, A. S., Bechmann, A., Zimmer, M., Ess, C., & the Association of Internet Researchers. (2020). *Internet research: Ethical guidelines 3.0*. Association of Internet Researchers. https://aoir.org/reports/ethics3.pdf

Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *PNAS*, 120(30)

.

Goffman, E. (1959). *The presentation of self in everyday life*. Doubleday.

Goffman, E. (1963). *Stigma: Notes on the management of spoiled identity*. Prentice-Hall.

Haupt, A. (2017). Mix en meng it op: Emile YX?'s alternative race and language politics in South African hip-hop. *M/C Journal*, 20(1).

Hern, A. (2017, 24 October). Facebook translates 'good morning' into 'attack them', leading to arrest. *The Guardian*.

Insel, T. R. (2017). Digital phenotyping: Technology for a new science of behavior. *JAMA*, 318(13), 1215–1216.

Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2017). Bag of tricks for efficient text classification. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, 427–431.

Kemp, S. (2025, 8 November). *Digital 2026: South Africa*. DataReportal. https://datareportal.com/reports/digital-2026-south-africa

Krippendorff, K. (2004). Reliability in content analysis: Some common misconceptions and recommendations. *Human Communication Research*, 30(3), 411–433.

Liu, M. (2023). Are you really smiling? Display rules for emojis and the relationship between emotion management and psychological well-being. *Frontiers in Psychology*, 14, Article 1035742. https://doi.org/10.3389/fpsyg.2023.1035742

Lombard, A., Aina, T., Wolff, E., Norvick, E., Gumede, S., Mawere, S., ... & Buys, J. (2026). MzansiText and MzansiLM. *Proceedings of LREC 2026*, 1394 to 1408.

Marwick, A. E. (2023). *The private is political: Networked privacy and social media*. Yale University Press.

Marwick, A. E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1), 114–133.

Marwick, A. E., & boyd, d. (2014). Networked privacy: How teenagers negotiate context in social media. *New Media & Society*, 16(7), 1051–1067.

McKinney, C. (2024). *The Routledge handbook of multilingualism* (2nd ed.). Routledge.

Mesthrie, R., & Hurst, E. (2013). Slang registers, code-switching and restructured urban varieties in South Africa: An analytic overview of tsotsitaals with special reference to the Cape Town variety. *Journal of Pidgin and Creole Languages*, 28(1), 103–130.

Mhlongo, N. (Ed.). (2019). *Black tax: Burden or ubuntu?* Jonathan Ball Publishers.

Miller, H., Thebault-Spieker, J., Chang, S., Johnson, I., Terveen, L., & Hecht, B. (2016). "Blissfully happy" or "ready to fight": Varying interpretations of emoji. *Proceedings of the International AAAI Conference on Web and Social Media*, 10(1), 259–268. https://doi.org/10.1609/icwsm.v10i1.14757

Mugumbate, J. R., & Chereni, A. (2020). Editorial: Now, the theory of Ubuntu has its space in social work. *African Journal of Social Work*, 10(1), v–xv.

Mugumbate, J., & Nyanguru, A. (2013). Exploring African philosophy: The value of ubuntu in social work. *African Journal of Social Work*, 3(1), 82–100.

Mushaathoni, M., & Dondolo, B. (2025). Reactions of social media users to trendsetters' code-switching: An exploration of linguistic choices in digital spaces in South Africa. *Forum for Linguistic Studies*, 7(12), 1236–1252. https://doi.org/10.30564/fls.v7i12.11255

Myers-Scotton, C. (1993). *Social motivations for codeswitching: Evidence from Africa*. Oxford University Press.

Nekoto, W., Marivate, V., Matsila, T., Fasubaa, T., Fagbohungbe, T., Akinola, S. O., Muhammad, S., Kabongo Kabenamualu, S., Osei, S., Sackey, F., Niyongabo, R. A., Macharm, R., Ogayo, P., Ahia, O., Berhe, M. M., Adeyemi, M., Mokgesi-Selinga, M., Okegbemi, L., Martinus, L., … Bashir, A. (2020). Participatory research for low-resourced machine translation: A case study in African languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020* (pp. 2144–2160). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.findings-emnlp.195

Niehaus, I., with Mohlala, E., & Shokane, K. (2001). *Witchcraft, power and politics: Exploring the occult in the South African lowveld*. Pluto Press.

Nwoye, A. (2015). African psychology and the Africentric paradigm to clinical diagnosis and treatment. *South African Journal of Psychology*, 45(3), 305–317. https://doi.org/10.1177/0081246315570960

O'Connor, C., & Joffe, H. (2020). Intercoder reliability in qualitative research: Debates and practical guidelines. *International Journal of Qualitative Methods*, 19.

Owens, E. Conceptualising TikTok in Youth and Media Research. *JAYS* (2026). https://doi.org/10.1007/s43151-026-00223-3.

Pennycook, A., & Makoni, S. (2020). *Innovations and challenges in applied linguistics from the Global South*. Routledge.

Ratele, K., Cornell, J., Dlamini, S., Helman, R., Malherbe, N., & Titi, N. (2018). Some basic questions about (a) decolonizing Africa(n)-centred psychology considered. *South African Journal of Psychology*, 48(3), 331–342. https://doi.org/10.1177/0081246318790444

Reiss, M. V. (2023). Testing the reliability of ChatGPT for text annotation and classification. arXiv:2304.11085.

Schoon, A., & Bosch, T. (2024). TikTok and memetic activism against racism in South Africa. *Information, Communication & Society*, *27*(16), 2760-2775.

Schoon, A., Mabweazara, H. M., Bosch, T., & Dugmore, H. (2020). Decolonising digital media research methods. *African Journalism Studies*, 41(4), 1 to 15.

Shibeshi, Z., Walton, M., Sibeko, J., Mthembu, V. M., Evans, H.-C., Ngcungca, N., & Schoon, A. (2025). Multilingual vibes: Visualising linguistic resources and emoji in Southern African online discourse. *Journal of the DHASA*, 6(1). https://doi.org/10.55492/dhasa.v6i01.6732

van Dijk, T. A. (2008). *Discourse and context: A sociocognitive approach*. Cambridge University Press.

van Dijk, T. A. (2015). Critical discourse studies: A sociocognitive approach. In R. Wodak & M. Meyer (Eds.), *Methods of critical discourse studies* (3rd ed., pp. 466–485). Sage.

Van Dyk, G. A. J., & Nefale, M. C. (2005). The split-ego experience of Africans: Ubuntu therapy as a healing alternative. *Journal of Psychotherapy Integration*, 15(1), 48–66.

Ziems, C., et al. (2024). Can large language models transform computational social science? *Computational Linguistics*, 50(1).

A Coding Instructions and Codebook

A.1 Coding instructions

The verbatim prompts are not reproduced here. Their operative content is the codebook in A.2 together with the standing instructions summarised below, which were common to the four prompt variants (default regime, strict regime, thread context, and census). Model and run details are given in Section 3.5. Coders read every item and were prohibited from delegating, from scripted pattern-matching, and from extrapolating labels across items, and they wrote output incrementally and checked it against the source chunk before reporting. Primary language was assigned from morphological diagnostics, not word lists (Section 3.5). The default and strict regimes differed only in how Nguni material without a diagnostic marker was

treated: isiZulu under the default regime, Unclassifiable under the strict regime (Section 3.6). The humour and negation rules were revised during the study as described in Section 3.6. The thread-context variant showed each comment with its parent post and added the ten-category comment role. The census variant coded context, real distress, and experiencer as defined in Section 3.8. Every variant required the structured coder report described in Section 3.5.

A.2 Codebook

B Temporal Coverage of the Idiom Census

Collection ran in 2026, but items were posted between July 2011 and July 2026, with TikTok items from March 2021. Of all items, 70.2% date from 2025 or 2026. Of the 10,366 census items, 56.8% do. We attribute this concentration to recency in platform search and to growth in use, and we do not read it as a prevalence trend (Table 3). We checked for event-driven spikes. Monthly census counts in 2020 ranged from 28 to 64 (median 48; peak March 2020), and in 2023 from 62 to 107 (median 74; peak April 2023), so idiom use shows no pronounced 2020 or 2023 spike. The dataset as a whole spikes in April 2023 (893 items), but 251 of those are a single-day cluster of TikTok items on 25 April 2023, and the spike does not appear in the census.

Table 3 gives census items by year and platform.

Year	X	Instagram	TikTok	Total
2016	60	7	0	67
2017	31	13	0	44
2018	30	29	0	59
2019	281	73	0	354
2020	492	81	0	573
2021	486	135	20	641
2022	830	138	31	999
2023	698	130	93	921
2024	610	147	66	823
2025	2,859	269	273	3,401
2026	1,939	153	391	2,483

Table 3: Census items by year and platform (n = 10,366; 2026 runs to 22 June; one further item dates from 2015).

C Verification Log

Across 95 coding chunks, verification and coder self-checks caught: two chunks reporting statistics for more rows than they had labelled; one chunk substituting scripted pattern-matching for reading; nine dropped or misaligned row blocks, including a deterministic nine-row loss at a pass boundary that recurred in four independent chunks; and several identifier transcription errors. All contaminated output was isolated and re-produced. Of 53 re-coding chunks, 52 matched their sources exactly, and one was short a single row of 153.

D Division of Labour

Task	Performed by	Nature
Collection and storage (Socius)	Software	Deterministic
South African attribution at collection	LLM relevance judge; not independently audited	Probabilistic
Stratified draw of items for coding	Software (drawn identifiers released)	Reproducible from released identifiers
Item-level coding (language, code-switching, mental-health relatedness, register)	LLM coders (Claude Sonnet 5)	Probabilistic
Idiom census coding (context, real distress, experiencer)	LLM coders	Probabilistic
Integrity verification of coder output	Scripts	Deterministic
Review of coder reports; revision of coding rules; boundary decisions	Researchers	Human judgement
Checking of glosses in quoted extracts	First-language-speakers	Human judgement
Critical discourse analysis	Researchers	Human judgement

Table 4: Division of labour.

E Sample Composition by Platform and Coded Set

Most unclassifiable items are contentless, not ambiguous: 318 of 415 had no usable text, 59 had unresolved text, and only 38 (0.3% of coded items) were set aside because Nguni or Sotho-Tswana vocabulary could not be attributed. A South African language is primary in 27.6% of coded items (3,017), and an indigenous language in 14.6% (1,592). Code-switching is flagged in 25.5% (2,787). The Sonnet set carries no language lists for 3,084 items, so mixing is described for the 7,856 items that have them: 1,186 (15.1%) mix languages. Among these, English with isiZulu accounts for 45.1%, English with isiXhosa 21.6%, Afrikaans with English 10.5%, and English with Sesotho 2.9%, and 41 items combine three lan-

Field	Values	Notes
`llm_coded`	True / False	Renamed from human_read. True for the 10,940 coded items
`coding_protocol`	default, strict	Ambiguous-Nguni regime
`source_chunk`	Chunk file or Sonnet set	Provenance of each coded item
`primary_language`	Eleven official languages, Other, Unclassifiable	Diagnostics-based
`languages`, `n_languages`	Every language present; count	Code-switching field
`code_switched`	True / False	
`mh_related`	True / False	Before humour correction. mh_related_read duplicates it and is dropped
`mh_related_v2`, `recoded`	True / False	After humour correction; recoded marks the 1,616 reclassified coded items (plus 7,370 census-only items)
`comic_register`, `weeping_register`	True / False	Laughing or weeping emoji and markers
`nguni_ambiguous`, `sotho_ambiguous`, `unclassifiable_reason`		No diagnostic in either direction
`sa_indigenous`, `sa_any`	True / False	Derived from primary language: indigenous only; any South African language including Afrikaans
`gowa_family`	True / False	Member of the idiom census
`cohort`	Four cohorts	Data collection section

Table 2: Codebook fields as named in the released dataset.

guages. Counting primary language alone understates indigenous-language presence: isiZulu appears in 11.3% of these items but is primary in 6.2%, and isiXhosa appears in 6.2% but is primary in 4.3%.

The three sub-samples differ sharply. TikTok supplies 53.8% of the dataset, but 96% of its TikTok items are comments. It supplies 14.7% of the coded sample and 8.4% of the idiom census. The census is 80.2% X and 11.3% Instagram, and the coded sample is 48.5% Instagram and 36.8% X. The idiom's textual life is therefore concentrated on X and Instagram, while TikTok contributes the comment threads in which ambient disclosure appears.

F Register Markers by Language

Primary language	n	Laughing % (95% CI)	Weeping % (95% CI)
English	1,712	20.7 (18.9 to 22.7)	30.8 (28.6 to 33.0)
isiZulu	470	23.0 (19.4 to 27.0)	37.0 (32.8 to 41.5)
isiXhosa	507	11.8 (9.3 to 14.9)	41.6 (37.4 to 46.0)
Setswana	118	31.4 (23.7 to 40.2)	23.7 (17.0 to 32.2)
Sesotho	67	40.3 (29.4 to 52.3)	26.9 (17.7 to 38.5)

Table 6: Register markers in census items coded for language (n = 2,996).

Set	n	English	Afrikaans	isiZulu	isiXhosa	Setswana	Sesotho	Unclassifiable
All coded items	10,940	68.2	13.0	7.2	5.3	1.2	0.7	3.8
Chunk files, comments	1,972	54.6	6.6	12.7	5.6	1.1	0.9	17.5
Chunk files, posts	5,625	82.0	7.7	4.0	3.9	0.7	0.5	0.7
Sonnet set	3,343	53.2	25.8	9.2	7.3	2.2	0.9	0.9

Table 5: Primary language by set, % of coded items.