



CASE STUDY REPORT



University of
Nottingham
UK | CHINA | MALAYSIA



Table of Contents

1	Organisation	2
1.1	Context and Ecosystem embedding	2
1.2	Case Study Methodology	3
1.3	Persons	3
1.3.1	Attributes	4
1.4	Use of AI and Expected Benefits	4
2	Challenges	6
3	Addressing the Challenges of Generative AI	14
3.1	Mitigation strategies	14
3.2	Organisational AI policies	15
4	Actionable Guidance.....	16
4.1	Requirements	16
5	Conclusion.....	17

Case Study Report for JournalismAI

1 Organisation

JournalismAI¹ is a Zimbabwean multi-media organisation and civic technology hub based in Bulawayo. JournalismAI was founded in 2016 but actual integration of generative AI started in 2023. The organisation's origins trace back to 2015 with a focus on documenting the Gukurahundi genocide². It has since evolved into a digital-first media house that leverages technology to disseminate information on political, social, and economic issues, with a particular focus on marginalised communities. JournalismAI is committed to investigative journalism, fact-checking, and holding public officials accountable. They are known for their "promise tracker" which monitors campaign promises made by politicians.

The organisation is a key player in Zimbabwe's media landscape, utilising a variety of digital platforms including YouTube, Twitter, and WhatsApp to engage with their audience. They produce a range of content from daily news shows and investigative reports to voter education materials. JournalismAI has also been a pioneer in the use of Artificial Intelligence in Zimbabwean journalism, employing an AI presenter to deliver news bulletins. In the lead-up to the 2023 general elections, JournalismAI deployed over 50 journalists to cover the elections across all 10 of the country's provinces, demonstrating their extensive reach.

The organisation operates a not-for-profit model supported by project-based grants and funding from international organisations. The organisation has received support from various international bodies, including the South Africa Media Innovation Program (SAMIP) and Meedan³.

1.1 Context and Ecosystem embedding

JournalismAI operates in a complex sociotechnical and economic landscape shaped by Zimbabwe's evolving digital media ecosystem, constrained civic space, and infrastructural limitations. As a media and civic technology organisation based in Bulawayo, JournalismAI plays a pivotal role in democratising access to information, fostering civic engagement, and promoting public accountability through innovative storytelling formats. The organisation is embedded in a national context where

¹ The name of the company has been changed to safeguard privacy and confidentiality.

² <https://www.ebsco.com/research-starters/history/gukurahundi-genocide>

³ <https://meedan.com/programs/check-global>

traditional journalism often struggles under political pressure, limited resources, and restricted media freedoms. Zimbabwe is ranked 116th out of 180 countries in the 2024 World Press Freedom Index⁴, with the media landscape heavily dominated by state-controlled outlets that align with the ruling ZANU-PF party⁵. Independent journalism is constrained by restrictive legislation, such as the Cyber Security and Data Protection Act and the Criminal Law (Codification and Reform) Act, often referred to as the "Patriot Act," which can be used to penalise journalists and suppress dissent⁶.

The organisation's embedding within the civic tech ecosystem allows it to circumvent some of the controls placed on traditional print and broadcast media, reaching audiences directly through social media and online platforms. Against this backdrop, JournalismAI's adoption of generative AI technologies, particularly AI avatars and automated content creation tools represents a strategic response to workforce constraints, audience engagement challenges, and the need for cost-effective news production.

JournalismAI ecosystem includes local journalists, students, civic actors, digital rights advocates, and a diverse Zimbabwean audience with varying degrees of digital literacy. Their current solutions are designed to enhance newsroom efficiency, provide accessible news summaries, and introduce creative, visual content.

1.2 Case Study Methodology

The report is centred on the RAISE project case study methodology⁷ which outlines the aims and objectives of the case studies, research hypothesis, case study design, data collection and data analysis. The report also outlines the systematic approach used by the project to study the AI ecosystem.

1.3 Persons

As part of the case study, interviews were carried out with 3 of JournalismAI's staff. This enabled the project to gather insights from a range of perspectives, in this case, a user, technical, and top-level executive staff. The interviews were carried out in October/November 2024 (virtually via Microsoft Teams). This also provided key insights around the organisational structure of the case and how many years of experience people had interacting with generative AI. Additionally, interviews were also carried out during site visits by the RAISE project team which occurred in November 2024. This

⁴ <https://rsf.org/en/country/zimbabwe>

⁵ <https://rsf.org/en/country/zimbabwe>

⁶ <https://www.amnesty.org/en/latest/news/2023/06/zimbabwe-parliaments-passing-of-patriotic-bill-is-a-grave-assault-on-the-human-rights/>

⁷ <https://raise-project.uk/wp-content/uploads/2025/04/RAISE-Project-Case-Study-Methodology.pdf>

provided additional insights which focused on the gaps, which were identified from an initial analysis of the interviews.

1.3.1 Attributes

Attributes of the personnel identified/interviewed in the case study are as follows

- **Director**
- **Innovation and Research Product Manager.**
- **Social Media Editor**

1.4 Use of AI and Expected Benefits

JournalismAI aims to reimagine local journalism through the integration of generative artificial intelligence (AI). Faced with limited newsroom resources and a rapidly evolving media landscape, JournalismAI has embraced AI not merely as a technical tool, but as a catalyst for content innovation, audience personalisation, and financial sustainability, “We are using AI to help us have automated news,” one team member explained, noting that the organisation is currently applying AI at multiple points in the content pipeline from news gathering and production to distribution and audience engagement. JournalismAI core AI solution involves the integration of AI avatars. A key avatar name “Julia⁸” has been at the forefront of their AI solution. The avatars which are virtual newsreaders developed using platforms such as Synthesia, FlexClip, and Play.ht which are used to produce daily news programme summarising stories generated by human journalists “We have a programme which is called the Brief Bulletin, whereby we use an AI avatar to actually present a summary of the stories that would have been written by our journalists,” one participant stated. These avatars also present weather reports and are being explored for new formats such as political candidate profiles, especially during high-volume events like national elections.

This has significantly reduced production costs and time, replacing a traditionally resource-heavy workflow involving videographers, editors, makeup teams, and presenters with a more agile, single-operator model “We just programme the presenter. They already have their makeup done. The only thing that you need to do is to just command it and edit the programme...and then we're done,” one staff member explained, highlighting both the creative opportunity and cost-saving imperative behind the initiative. “With AI, it's just one person collecting the stories and using the AI to do a news bulletin. So, it's faster, it's more efficient, it takes less people.”

In parallel JournalismAI is also developing an in-house WhatsApp chatbot to enhance audience engagement and automate personalised content delivery. This solution aims to automate the delivery of personalised content to JournalismAI’s primary audience who are Zimbabweans who access news via WhatsApp. The chatbot leverages platforms like Meta’s Llama, and third-party services such as Luminate to draw user insights, segment audiences, and tailor content, including targeted advertising. This is especially

⁸ Name pseudonymised for confidentiality and privacy

critical in Zimbabwe where conventional platforms offer limited analytics, and where funding models are shifting due to a decline in donor support. “We’ve noticed that we’ve got websites, we’ve got YouTube, got Facebook...but most of our users are found on WhatsApp,” a participant noted. “We are trying to figure out... how AI can help us to send specific content to our readers.” The chatbot will segment users, learn their content preferences, and deliver targeted updates. It also forms part of a monetisation strategy aimed at shifting from donor dependency to commercial sustainability. “We are now transitioning to a for-profit media organisation... so we want to use this AI to serve targeted ads,” the team said. “The chatbot... will learn the habits of the user and then it will be sending targeted ads whereby our business development manager can... tell advertisers how many people viewed your advert.”

Analytics also play a vital role in JournalismAI’s AI approach. Initially reliant on platforms like Google Analytics, the team now employs tools that aggregate data from WhatsApp, Telegram, and other channels to inform editorial and commercial decision-making. “The most important thing for me as a tech person is the analytics,” one participant reflected. “Now you can try and get it from WhatsApp... [and] that is a big game changer. We now know exactly where our traffic is going and where it's coming from.”

Beyond efficiency and reach, JournalismAI sees AI as a creative force that strengthens journalism. “When journalists use some of these tools like ChatGPT or Perplexity, it gives them other aspects on how to create a news story,” a participant noted. “It has made them better journalists... to ask more questions.”

The expected benefits of JournalismAI’s AI initiatives are threefold:

1. Operational efficiency – Reducing time and human resources required to produce high-quality, engaging content.
2. Audience-centric engagement – Using behavioural data to deliver tailored news and services through dominant local channels like WhatsApp.
3. Financial sustainability – Enabling revenue generation through personalised advertising

JournalismAI’s ambition is not only to enhance internal workflows but to *democratise access* to information across Zimbabwean communities. As one team member remarked, “We want to use AI to serve those communities where people cannot have access... How do we then build around chatbots so someone can access information via WhatsApp?”. Above all, JournalismAI views AI as a strategic equaliser. “The tools are there... and you can also try and compete with the bigger organisations,” they noted. “Yes, they are like maybe 100 steps ahead of you, but you are like maybe 40 steps behind them you’re always following up.”

2 Challenges

The adoption of generative AI by JournalismAI has been met with several economic, ethical, legal, technical, organisational and social challenges. This section discusses some of the key challenges encountered as identified through the analysis of interviews and site visits.

Economic Challenges

While the potential of generative AI to transform journalism in Africa is widely acknowledged, the economic realities of adopting and maintaining these technologies present considerable challenges particularly for small and transitioning news organisations such as JournalismAI. As a participant from JournalismAI remarked, *“although we are a non-profit media organisation, we are now transitioning to a for-profit media organisation because the donor world is shrinking, and we need to be able to make our own money.”* This shift from donor dependency to self-sufficiency is driven not only by changing funding landscapes but also by the escalating costs associated with deploying AI tools effectively.

At the centre of these economic concerns is the high cost of AI development and operations. Participants highlighted that even relatively modest innovations such as creating an AI avatar for news delivery can strain newsroom budgets. *“AI is quite expensive to run,”* one interviewee admitted, *“we had to create our own AI avatar, and it's quite expensive to do that... training and teaching the words and making the eye movement—that was a bit of a challenge.”* These costs are further compounded by the need to hire videographers, secure green screens, and manage other technical processes typically outside the scope of small newsrooms.

One of the most revealing insights came from JournalismAI’s efforts to build a WhatsApp-based AI chatbot, a promising tool for audience engagement and news personalisation. What began as a seemingly affordable venture quickly spiralled. *“So, you know how all these services start—there is the small business and premium. They say it's \$100 a month and you say OK, I can afford that. Then the more people that jump on the platform... you're forced to go into the business tier, which goes up to \$300 a month... then you need like \$5,000 for five hours,”* the participant explained, calling it *“a trap.”*

Beyond monthly subscription costs, integration with APIs (application programming interfaces) and other platforms also introduced hidden costs. *“When you're trying to link the API to the website... there's also a specific charge,”* a participant noted. *“And then all of a sudden you need to redevelop your website so that they can talk nicely to the API. So you need a specific developer.”* Such costs often unanticipated during the initial

planning stages can lead to financial setbacks, delayed implementation, or even project suspension.

In fact, the ongoing costs of operating AI systems have forced JournalismAI to consider pausing its chatbot development entirely. *“The one that has forced us to pause is the WhatsApp chatbot. It’s expensive to run on a monthly subscription... we had thought it was going to be a simple budget, but we’ve blown our budget,”* one member admitted. The team is now attempting to monetise its online ad space to sustain its AI projects, *“trying to figure out ways on how we can make money from the ads on our websites to pay for this AI.”* Even one time expenses were described as prohibitive. The cloning of a voice for a customised avatar, for instance, was quoted at \$4,000 prompting the team to seek out a less expensive platform, albeit at the cost of performance. *“The voice cloning is not really perfect as it should be... it does not talk as much as you’d want it to,”* one participant explained.

These financial concerns are not isolated to custom-built tools. Off-the-shelf generative AI platforms also require sustained investment. *“Most effective AI tools are not free... cloning our own members of staff to have our own AI presenter cost \$1,000 a year for the subscription,”* said one team member. For African newsrooms already grappling with sustainability, even seemingly small monthly costs accumulate quickly. *“To have a newsroom use a GPT—you’ll be paying maybe \$25 to \$40 a month, and imagine how many other software subscriptions we need.”*

Compounding the economic burden is the uncertainty around audience readiness and revenue return. As one participant warned, *“People just jump in without knowing their audience or their market... You might be using AI for doing your products, but you might not have the audience to justify it.”*

Ethical Challenges

As newsrooms across Africa explore the integration of generative AI into journalistic practice, the participants from JournalismAI caution that this journey is fraught with not just economic hurdles, but also deep ethical concerns. While AI offers the promise of enhanced productivity, creativity, and cost-efficiency, these gains cannot come at the expense of foundational journalistic values or cultural accuracy.

A key concern voiced during the discussion relates to the perceived sidelining of ethics in the race to adopt AI-driven solutions. One participant observed a growing tension between innovation and ethical oversight, noting that *“you have scholars coming and say ethics, policies... so you ethics guys stay at the back.”* This sentiment underscores a

worrying trend in some tech-led initiatives where speed and novelty take precedence over careful consideration of impact and integrity.

Participants also expressed apprehension over the erosion of journalistic ethics in the face of increasingly powerful AI tools. With the proliferation of deepfakes and synthetic media, the threat of misinformation and impersonation looms large. *“We are living at a time when there's a lot of fake news... with AI, someone can clone your voice and pretend it's you,”* one journalist warned. In such a climate, participants argued that journalistic standards must not be abandoned: *“I don't think we should abandon our basic journalism ethics when it comes to AI.”*

Beyond the question of misinformation, the discussion highlighted cultural and racial biases embedded in generative AI systems, many of which reflect dominant Western paradigms. A striking example was raised concerning visual representations generated by AI image models: *“If I go on and just say draw a picture of a journalist carrying a camera, it's always a white person. If it's a black person, then they have to have something with African attire. So what the hell told these guys that everyone in Africa walks around with African attire?”* This points to the broader issue of algorithmic bias, where underrepresented populations are misrepresented or stereotyped due to training data that fails to capture global diversity.

The language barrier was also raised as a significant concern. AI tools, largely trained on English and other dominant global languages, often struggle with African vernaculars and indigenous linguistic patterns. *“How do we teach these AI models our own local languages?”* one participant asked, stressing that the offline inequalities and marginalisations found in traditional media are being replicated in AI systems. *“Remember the biases that we find offline—you also find them on AI.”*

Though the team acknowledges the risks including hallucinations and the imperfect nature of AI-generated content they emphasise a pragmatic, iterative approach. *“We make the AI write the story and then have our editor edit the story because it might not be fully perfect,”* one interviewee said.

The insights from JournalismAI serve as an important reminder that the integration of generative AI in journalism must go hand in hand with ethics, inclusivity, and critical reflection. While the technology offers undeniable potential, its development and deployment should be informed by diverse cultural contexts, linguistic pluralism, and a steadfast commitment to truth and fairness in reporting.

Legal challenges

Insights from the JournalismAI participants reveal a spectrum of attitudes from cautious adaptation to bold defiance towards legal compliance and regulatory frameworks when deploying AI technologies in journalism.

One of the central challenges highlighted is the incompatibility of global frameworks with local realities. Participants reported that while they are aware of legal standards such as the General Data Protection Regulation (GDPR), full implementation can alienate audiences unfamiliar with such protocols. As one interviewee explained, *“Some people don’t understand they need to sign up for the GDPR issue. If we try and put that into our systems, people get turned off our audiences. So we try and adopt some of them and localise them. We take what we want and then what we do not like, we leave it as it is.”* This practical approach adopting what is useful and discarding the rest highlights the lack of contextual adaptation in many international AI governance tools.

Yet beyond these adaptation challenges, there are deeper ethical-legal dilemmas about the use of generative AI in content creation. One particularly contentious issue relates to image generation. In the absence of visual footage from remote areas, some team members debated the possibility of using AI-generated images to simulate accidents or events. The editorial team ultimately rejected this idea: *“We had a very [heated] debate... Do we use a generative AI to create these images and then say it was an accident? So we have said no, we won’t do that because that is misrepresentation.”* Their decision underscores a commitment to avoiding potential misinformation and legal liability through visual manipulation.

However, legal awareness and adherence were not uniformly prioritised among participants. One contributor took a defiant stance, shaped by a background in unlicensed community radio. *“If you ever want not to do anything, start reading about the law... because you will end up not doing anything. So I ask myself: If I do this thing, am I going to die?”* This provocative sentiment reflects a broader frustration with restrictive legal environments that may stifle innovation, especially in underfunded or alternative media spaces.

In fact, some participants even argued that legal frameworks are often misaligned with creative experimentation and public need. *“The most popular things are not the most legal ones... unless and until you adopt the mind of a drug dealer,”* one participant remarked, suggesting that legal nonconformity is sometimes a prerequisite for media innovation in contexts where traditional infrastructures or permissions are inaccessible.

That said, the issue is far from black and white. The complexity of transnational audiences especially for outlets with a strong diaspora following introduces further legal tensions. One interviewee noted, *“Our audience is a bit varied... most of them are in the*

United Kingdom, in America, Australia and South Africa... automatically, when they go to Europe... we are in a peculiar situation [with GDPR compliance].” This underscores the difficulty for local media to navigate data protection rules that apply across jurisdictions.

The novelty of generative AI itself poses a legal grey area. *“When it comes to AI, it’s about pushing boundaries,”* said one participant. *“That’s what I find fun.”* Yet, this attitude while reflective of the experimental culture around new technologies raises serious questions about accountability, transparency, and legal risk.

Organisational Challenges

The integration of generative AI into journalism is not merely a matter of adopting new tools; it is a multifaceted organisational transformation. For small, resource constrained media outfits such as the one represented in the JournalismAI interviews this transformation is marked by growing pains, internal resistance, skill gaps, and structural limitations that make the shift particularly difficult to manage.

A key theme emerging from the interviews is the absence of clear rules and frameworks to guide early adoption. As one participant put it, *“Since this is a new field, there are no simple rules. Also then there’s a bit [of] learning, there’s not yet [anything] cast on stones.”* This lack of clarity has led to operational mishaps, including plagiarism concerns and ineffective use of AI-generated content. *“One or two [times] we were caught with our pants down because I think one of our team members would just copy and paste stuff,”* recalled one team member, revealing a vulnerability in editorial oversight during the AI integration phase.

Rather than immediately improving efficiency, the introduction of generative AI often increased the workload due to necessary quality checks. One respondent noted, *“Instead of making it easier for us, it was like double the job because you would make AI do everything for you. And then now you need to read it and proofread that whatever it’s done is correct.”* This speaks to the need for clear workflows and editorial policies when introducing AI into content production.

Organisationally, the team also grappled with identity and transparency issues around AI-generated content. A critical moment of ethical reflection arose around whether to present an AI-generated output as human-written: *“We were contemplating the ethics: do we let it go and pretend it’s one of the reporters or do we tell them that it’s our reporter with the help of an AI?”*

Resistance to change also emerged as a strong theme. The human tendency to maintain routines means that newsroom staff may push back sometimes passively against shifts in editorial workflow. One participant observed: *“People tend to resist it... [they] are*

creatures of habit... internally, of course, in the newsroom, you can see passive resistance.” This resistance complicates internal alignment around AI tools and undermines their successful deployment.

The team’s small size and voluntary structure also imposes constraints on experimentation and growth. *“Our team is small. We cannot keep on hiring people. And again, with some of the talent that is applied, we cannot afford them.”* This limited capacity not only affects technical development but also contributes to missed opportunities for structured research and testing. One team member admitted that failure to conduct sandbox testing during chatbot deployment caused the system to crash: *“I had to redo it again because I did not do a lot of... R&D... it broke along the way.”* As such, the lack of research and development (R&D) department is a notable organisational weakness.

At a broader strategic level, participants expressed concern over brand identity erosion due to AI-generated assets. One person noted that their AI avatar was now appearing in other products: *“We tended to notice that it's now being used in other products... we said no, this is gonna be a problem because it will end up doing things that we do not identify with.”*

Despite these challenges, there was a strong desire for self-reliance and organisational autonomy. As one respondent noted, *“We are trying to work on a model where we are self-reliant.”* This reflects a grassroots innovation ethos, though it is often constrained by scale and access to funding.

Social Challenges

For news organisations like JournalismAI who aim to be a pioneer in the use of generative AI for news dissemination the transition is fraught with misunderstandings, public resistance, ethical anxieties, and deep-rooted cultural dynamics as discovered in the insights from the case study.

A recurrent issue expressed by the participants is public distrust and confusion about what AI actually is and how it functions. As one team member described, *“AI in general, there was... a majority of our audience [who] did not understand what AI is.”* This knowledge gap necessitated continuous efforts to explain their AI tools to audiences, even embedding clarifying statements within content: *“The fifth sentence of the AI presenter would be ‘Hi, my name is Dash Dash Dash and I’m an artificial intelligence presenter.’”* Despite such efforts, misunderstanding persisted, particularly around the human-like qualities and roles of AI avatars.

Concerns over job displacement emerged as a dominant social tension. The use of AI avatars as news presenters triggered emotional and political backlash from audiences and journalists alike. One participant recalled, *“Most people were not amused by us using a generative AI news reader... They were saying we’ve got a very huge population who is unemployed... why are we using AI instead of hiring our local kids who graduated from the university?”* This sentiment extended into the newsroom itself, where both younger and older journalists expressed scepticism and resistance: *“The younger ones are complaining that you are taking our jobs and giving them to AI,”* while others argued, *“I don’t believe this thing of AI.”*

There was also discomfort with cultural and linguistic mismatches, which fuelled alienation among local audiences. One of the most sensitive flashpoints was pronunciation and representation. *“We faced a lot of backlash... in terms of the pronunciation of some of the indigenous words,”* one participant shared. Another elaborated: *“Our audiences are complaining about how the AI presenter... is mispronouncing some of the local names.”* This challenge reflects the limitations of globally trained AI models when applied in local African contexts with rich linguistic diversity and unique cultural norms.

The visual appearance and gendered perception of AI avatars introduced additional layers of complexity. In one striking example, the news team was forced to change their avatar's design due to audience criticism: *“We thought she looked more like an African... but most people were not quite amused and said she’s not pleasing to the eye.”* The second avatar, deemed more attractive, triggered inappropriate public reactions: *“Men started asking for her number... Some people believed she was real... they wanted to visit her.”* This scenario highlights the need for critical reflection on gender dynamics and digital embodiment in the design of AI tools.

Audience expectations also clashed with AI limitations. *“Some people ask you: why is it that your AI presenter doesn’t change clothes?”* said one respondent, noting the difficulty in managing assumptions when avatars appear “too human.” The anthropomorphism of AI systems blurred the line between reality and simulation, raising ethical concerns about deception and informed consent in media consumption.

Another layer of social complexity is the moral and religious framing of AI technology. In a humorous yet telling observation, one participant suggested that public acceptance might require a form of evangelism: *“Why can’t we start our own religious movement in AI... AI is the truth, the way and the life, whoever doesn’t use AI shall perish.”* This

hyperbolic metaphor captures the uphill struggle faced by innovators attempting to promote AI tools in communities with little technological literacy or trust.

There were also practical, day-to-day social issues, such as WhatsApp group users receiving irrelevant content due to the chatbot's limited personalisation capacity: *"People have got diverse views and needs... it's a learning process and there are a lot of hiccups."* This reinforces that social acceptance also hinges on whether AI can genuinely meet audience needs.

Technical Challenges

The integration of generative AI in African newsrooms is marked by a host of technical challenges that reflect infrastructural constraints, platform limitations, and a lack of localised expertise. Participants from JournalismAI spoke candidly about the difficulties they face when attempting to implement AI solutions in their journalism workflows.

A significant issue raised is the limitation of platform analytics, particularly with regard to WhatsApp, which remains the dominant medium for news distribution in many African contexts. As one participant explained, *"In this part of the world in Africa, people consume their content through WhatsApp and we do not have the analytics – WhatsApp does not give you analytics."* The absence of user insights hampers the ability to measure audience engagement, understand consumption behaviour, or tailor content effectively. Another participant elaborated, *"They have been complaining that yes, you've got the numbers, but the issue is how many do we know that open [a link]? Because using WhatsApp you can't."*

Beyond platform limitations, training AI tools and infrastructure development emerged as a key concern, especially in relation to cost and capacity. The process of training AI models or bots is not only time-consuming but also financially demanding due to lack of localised and custom infrastructure. One participant described the compounding cost *"Ideally, it would be nice if I were to build my own GPU, build my own rig... but getting the components down here is slightly expensive."*

In the longer term, this group is considering building local hardware to manage AI workloads more affordably. As the interviewee put it, *"If we're to do it, it might cost like \$5,000... but moving forward it's our rig – no more hours to pay for. It's all done locally."* This solution, while promising, illustrates the tension between ambition and resource availability in small news organisations.

Another central issue is the scarcity of local AI developers and technical talent. A participant noted, *"Finding local developers is an issue – when I say local, I mean*

Zimbabwean developers.” Because of this gap, organisations are forced to outsource to developers in countries such as India, which adds further costs and reduces the localisation of AI systems. *“We do not have as many people who have the technical know-how when it comes to AI locally here in Africa,”* one contributor stated.

Language localisation remains one of the most persistent technical challenges. Generative AI models often struggle with the pronunciation and processing of African languages, undermining their credibility with local audiences. *“AI will struggle with those words,”* said one interviewee. *“We need to develop our own language modules here that have our AI reading our news in Swahili, in Zulu.”* The lack of linguistic datasets and local language training modules hampers the ability of AI systems to reflect cultural and linguistic diversity. As another participant added, *“Definitely we need more data... to have our AI pronouncing and being able – yeah, just to pronounce some of the words.”*

Ultimately, the technical challenges faced by JournalismAI are deeply interwoven with broader systemic issues – limited infrastructure, weak local technical ecosystems, and poorly contextualised AI tools. These problems not only affect performance and sustainability but also hinder the trust and usability of AI solutions in news production and distribution.

Attempts to integrate AI with existing tools, such as WhatsApp newsletters, also proved technically cumbersome and procedurally off-putting. For example, data collection for newsletters was met with user fatigue due to compliance requirements: *“You needed to comply with those issues. So that was a bit of a turn off for some of our users.”*

3 Addressing the Challenges of Generative AI

Although the previous section highlights several challenges faced by JournalismAI, the analytical insights also identified some measures that have been implemented by JournalismAI to address some of the recognised challenges. Furthermore, the RAISE project has been able to make an impact on the development of its guidelines through its engagement with JournalismAI. This section outlines some of the actionable steps taken through the adoption of mitigation strategies and collaboration with the RAISE project.

3.1 Mitigation strategies

3.2 Organisational AI policies

The integration of generative AI into journalism is unfolding in real time, often outpacing the development of formal organisational policies and ethical frameworks. As expressed by participants from JournalismAI, many small and experimental newsrooms find themselves adopting AI tools first and then retroactively formulating governance structures. One participant candidly remarked, *“You know what is quite interesting is that we jumped in this thing and we thought of policies later, so we are now working back.”*

This admission captures a broader trend among digital-first and resource limited organisations: the use of generative AI is often driven by necessity, creativity, or ambition, while regulatory or ethical guidelines are added later. *“We are currently in the process of developing an AI policy for the organisation,”* one participant shared. *“When we started, the idea was, ‘OK, let’s experiment and see what works.’”* This pragmatic, exploratory approach has enabled these organisations to remain agile and responsive, but also raises questions about transparency, consistency, and accountability in AI-assisted journalism.

There is some scepticism within the group about the role of policy. One contributor reflected, *“Sometimes policies limit... There’s more creativity when there are no boundaries. The moment we start thinking of boundaries, then they are limitations.”* This attitude reveals a tension between innovation and institutionalisation: while formal policies offer clarity and accountability, they are sometimes perceived as stifling to the fluid, iterative nature of digital experimentation.

That said, transparency remains a shared value, especially regarding content production. Several participants stressed the importance of disclosing AI use in editorial processes. As one explained, *“When it comes to writing stories using AI, we actually state that this story was written with the assistance of artificial intelligence... to avoid those kinds of challenges.”* Another emphasised the principle of user consent and integrity, noting, *“Tell people this story was generated by AI and give them a chance to say whether they want to read it or they don’t... Let’s not manipulate the footage using AI; let’s use AI to just enhance the user experience.”*

While a comprehensive policy may not yet be in place for some, many are already drawing from existing external frameworks. *“We looked around everywhere... mainly for newsrooms, if they’ve got any AI policies,”* said one respondent. *“There are some initiatives like AI in journalism where they give you frameworks and guidelines. So we copy and paste them into our own needs.”* This form of adaptive policymaking tailoring global principles to suit local realities is especially vital in regions where media organisations must contend with distinct technological, cultural, and regulatory landscapes.

Ethical continuity also surfaced as a key concern. Despite the novelty of the tools, one participant emphasised the need to stick to foundational journalistic values: *“We still apply the same ethics, you know. Do no harm. Be open. Be transparent in the way you do stories.”* This underscores that while the tools may be new, the moral compass remains anchored in long-standing editorial norms.

4 Actionable Guidance

To provide actionable guidance the RAISE project captured what the case study/SME requires for the responsible adoption and use of generative AI.

4.1 Requirements

Across the case interviews, participants identified a combination of technical, organisational, financial and ethical capabilities that they considered necessary for implementing generative AI responsibly. A central requirement was access to technical expertise and internal R&D capacity. C3 currently relies substantially on externally available AI tools but identified a need to develop greater internal capability to customise and integrate AI into its workflows. As C3P3 explained, the organisation needs support to “build our own assistance”, integrate AI with its content management systems and develop solutions suited to its own operational requirements. The proposed establishment of an AI laboratory reflects this need for sustained experimentation and learning, particularly around newsroom tools and locally relevant applications. Participants also emphasised the importance of localisation, particularly the development of language capabilities that can accurately represent African languages, names and contexts. C3P3 argued that C3 needs to “develop our own language modules” because existing systems frequently struggle with local pronunciation and reproduce Western cultural assumptions. Alongside technical capacity, participants highlighted the need for appropriate governance and ethical guidance. Although C3 initially experimented with AI before developing formal rules, it was subsequently developing an AI policy addressing transparency, disclosure, acceptable uses of synthetic media and the avoidance of misleading content. Importantly, this governance was expected to provide direction without unnecessarily restricting experimentation: the organisation sought policies that would “guide our people, but not to limit the creativity”. Finally, responsible implementation requires sustainable financial resources and access to skilled people. Participants noted that effective AI tools, subscriptions, development and integration can be costly, while smaller African organisations often operate under severe resource constraints. Consequently, C3’s requirements extend beyond adopting individual AI applications towards developing an organisational capability that combines appropriate funding, technical expertise, R&D, localised infrastructure and language

resources, and adaptive governance. This is consistent with C3P3's broader principle that organisations should first understand “what problem” they are trying to solve and retain established ethical principles, including openness, transparency and avoiding harm, when applying AI.

5 Conclusion

This case study provides compelling insights which underpin the complexities, benefits and challenges involved in the integration of generative AI into business processes in the Journalism sector. The use of AI in the Journalism sector provides many opportunities and holds great promise in revolutionising Journalism which is one of the key areas of application by the case study.

The RAISE project has been able provide critical support to the case study through series of activities including workshops, interviews, consultations and discussions. These activities have provided a positive impact in the adoption generative AI by JournalismAI. The RAISE project will continue to work with JournalismAI and other case studies in the RAISE ecosystem to develop frameworks and tools to provide practical guidance to SMEs and startups looking to adopt generative AI.