"Just Send Us Pictures of Your Feet": Consent and Biometric Extraction in the AI Supply Chain

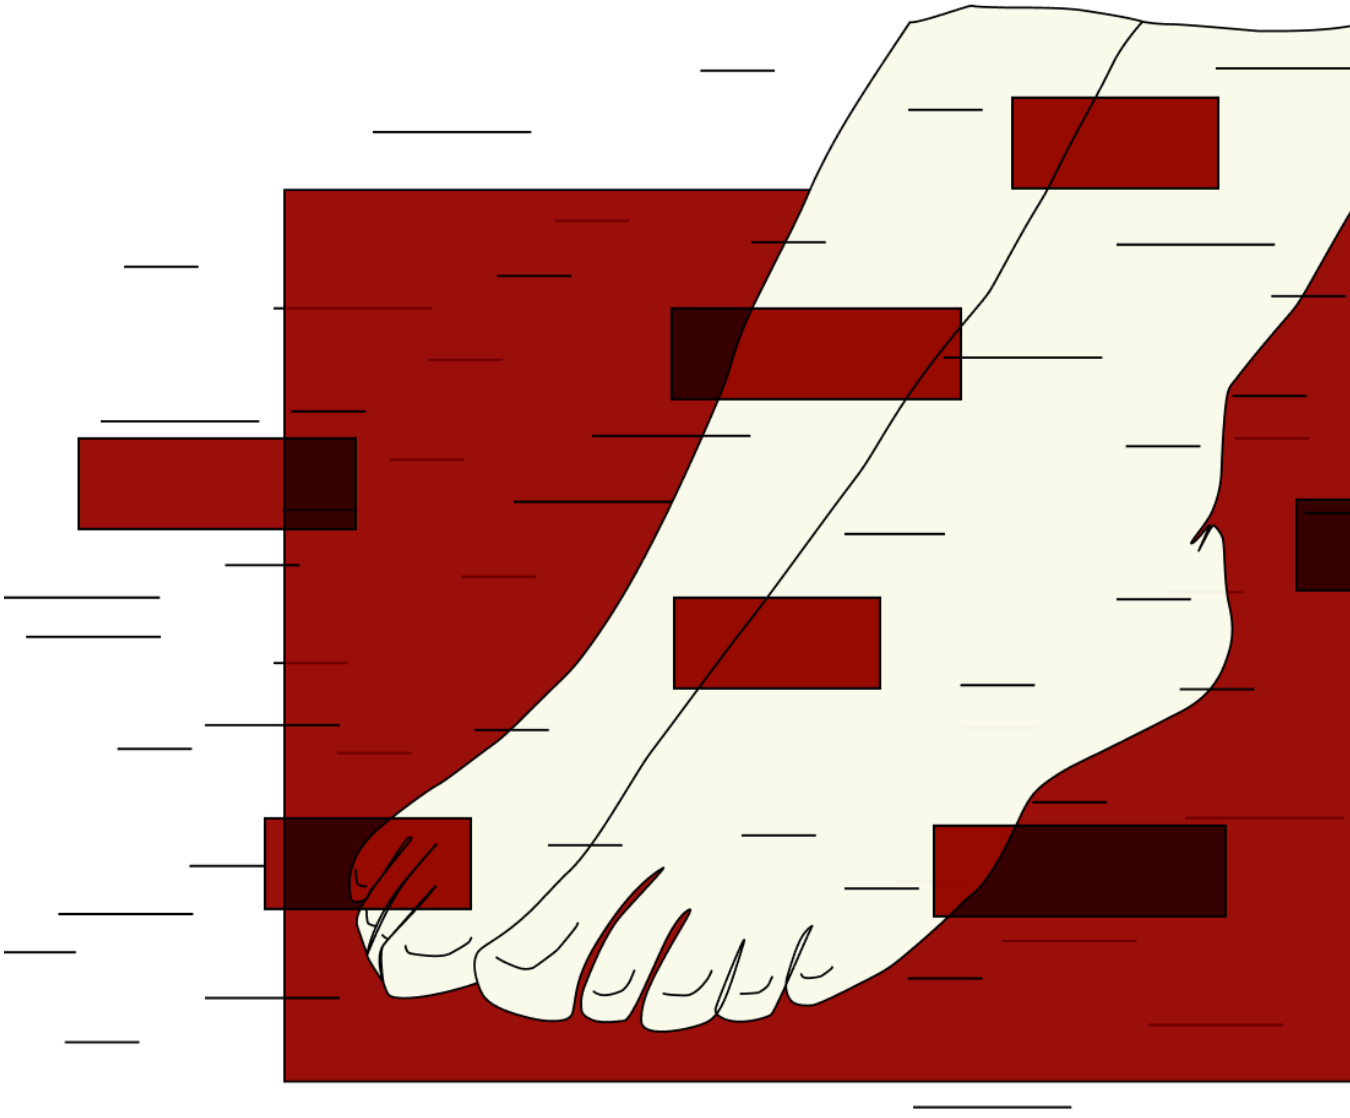

James Mojez Oyange

Suggested citation: Oyange, J. M. 2026, '"Just Send us Pictures of Your Feet": Consent and Biometric Extraction in the AI supply chain'. In: M. Miceli, A. Dinika, and C. Salim Wagner (eds.) *Data Workers' Inquiry*. Creative Commons BY 4.0. Retrieved from https://data-workers.org/Mojez.

A Report by Data Workers' Inquiry

With support from

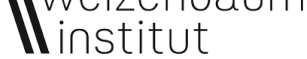

Acknowledgments

First, I want to thank DAIR and the extended team at DWI for helping me share this work. I have always carried this idea with me, writing about it and looking for an outlet, and DWI made that possible.

Secondly, I would like to thank my comrades for finding time and opening up their homes for me to document the in-depth data extraction that happens in their work. Their selflessness in sharing their stories goes a long way in helping the wider audience better understand the challenges that come with this work.

Lastly, I dedicate this report to the growing community at the Data Safety Alliance. As we create our core story in these first months of our organization, this research is a key part of our mission. To the advocates, allies, and partners who will join us, thank you for your shared commitment to making sure the human work behind AI is no longer hidden.

Background

My name is James Oyange—everybody calls me Mojez—and for a long time I have worked in the invisible tech space as a TikTok content moderator and a freelance data annotator, experiencing firsthand the challenges that come in these tech spaces.

My interest in data privacy violations by the companies we work for started with my cousin. He came to visit me one day in early 2024 when he was still in college, and he was constantly on his computer, seated next to me. I was doing my own thing, browsing Twitter, but I kept glancing in his direction because I could see him making facial expressions, tilting his head in different directions.

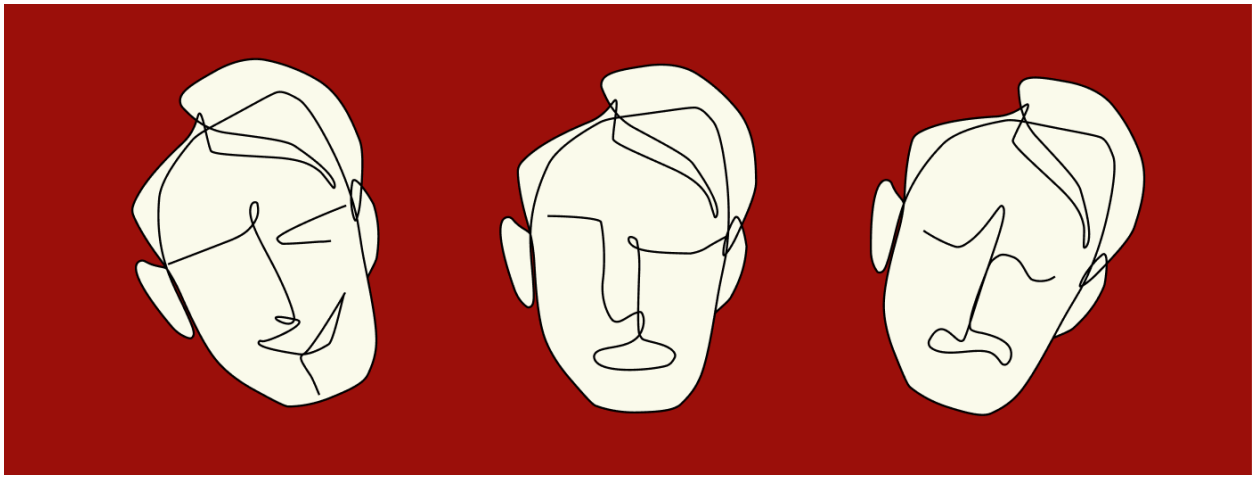

It piqued my interest. When I asked, he showed me the work he was doing: uploading pictures of himself, taken through his webcam, to a shared Google Photos drive and then sending them somewhere else. On the screen I could see the different pictures he had already taken; some of him smiling, frowning, looking up, down, sideways. I found it weird and appalling. I was well aware of how facial recognition technology is used by companies for their benefit. I asked him if he knew that his facial features might be used for profit or fraud. He brushed it off by showing me the amount of money he got in a week after submitting multiple tasks. It was USD 30, a good amount for a university student in Nairobi, but I asked myself: at what cost? It bothered me, and so I dug a little deeper into where these pictures were going. I later found out that the "somewhere" he had mentioned was Telus Digital, a business process outsourcing company which offers data-related services for AI training through offices

located in several countries, including Kenya. It supplies work related to content moderation and data annotation to big tech companies like Meta and OpenAI.

I then asked my comrades who worked in similar outsourcing companies if this was the norm, and to my surprise, it was. I was surprised because during my university days (2017-2021), when I was in and out of the annotation industry, I never came across projects where one sent selfies. It seemed to me that it was a new norm and I believe it came about with the advancements in AI. This terrified me. If my cousin's and comrades' data were being harvested, what about mine?

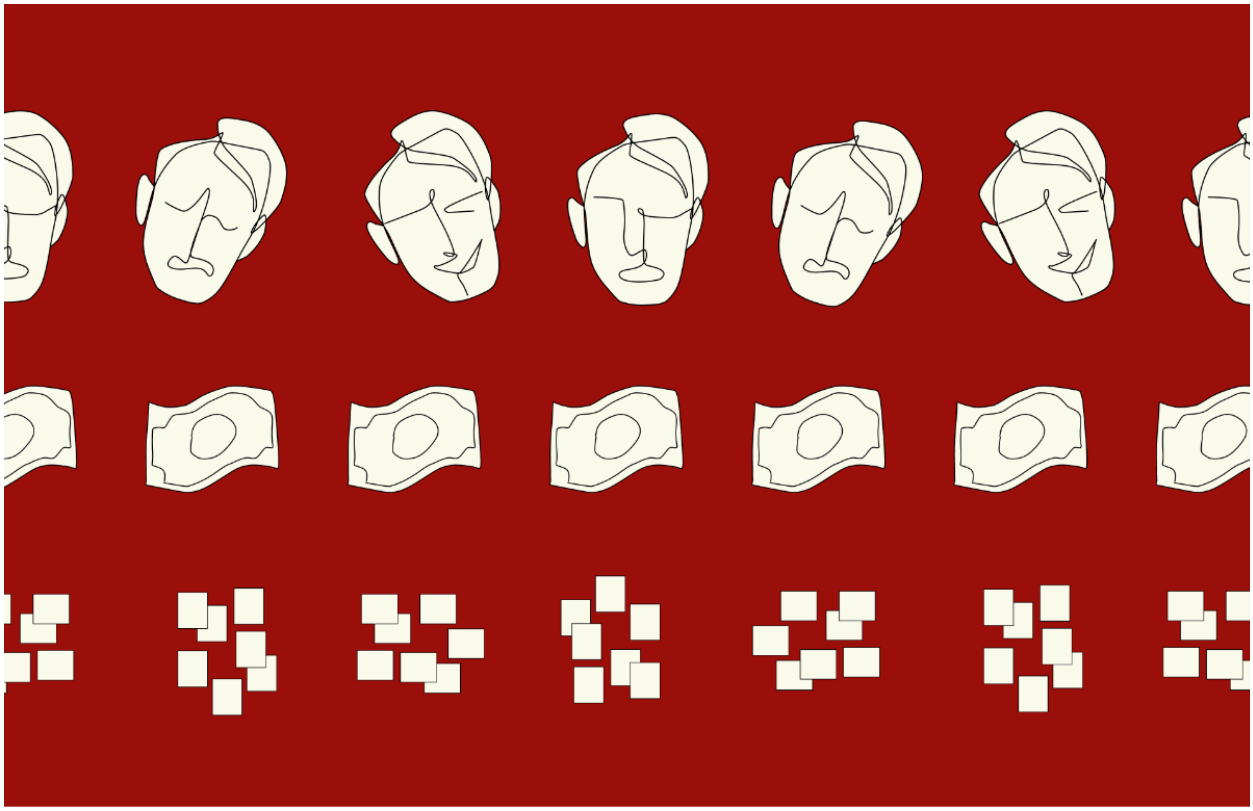

I tried to trace where my own data had been stored during my time as a content moderator at Majorel, now Teleperformance. All the biometrics, addresses, ID pictures—they had to go somewhere, right? I wrote a data subject access request to Teleperformance, wanting to know all the information they held on me, where it was being stored, and how it was being used after my employment ended. I did not receive any response. I eventually linked up with a lawyer friend who sent a follow-up email on my behalf, and that is what made the company respond. Their response was vague. They stated only that they had collected my name and address and that they were

done processing it. I had expected a detailed explanation. I was disappointed but not surprised. I now understood that companies will only respond to a worker if they have some sort of legal representation, or someone who knows the data protection industry.

That realization is what pushed me to enroll in a GDPR Certified Data Protection Officer (CDPO) training course with PECB, a globally accredited institution. I passed and received my certification in 2024. But the course gave me more than a certificate. For the first time, I had the proper language to name what had been happening to us. As a worker and a user, you have a right to access your own data, a right to know how it is being used. There is a principle called data minimisation; companies are only supposed to collect what they actually need. Reading through the course material, I kept recognising things Teleperformance had done to me and my colleagues. Brushing off my access request. Collecting biometrics with no clear explanation of why. These were not just bad practices, they were violations. And once you know something is a violation, you cannot unknow it. I came out of that course with two things: the tools to help me and my colleagues protect our data, and a much bigger question I needed to answer: how widespread was this?

Introduction

When my cousin Stu showed me how he sent pictures of his own face to train AI systems for $0.25 per task, I realized that the term “ghost workers” isn’t just a metaphor. The narrative of artificial intelligence often brings up images of autonomous, self-sufficient machines. Yet, this concept of autonomous machines is nothing but a cover. The real foundation is a worldwide human chain. Millions of "ghost workers" located across the globe are involved in the critical work of tagging images, converting speech into text, and screening content to guarantee that the algorithms are correct, secure and convenient. These individuals, often based in the Global South, perform repetitive and sometimes psychologically harmful tasks to train machine learning systems.

I needed to understand something specific: why were my fellow workers being asked to submit their own faces, voices, and even pictures of their children as part of their jobs?

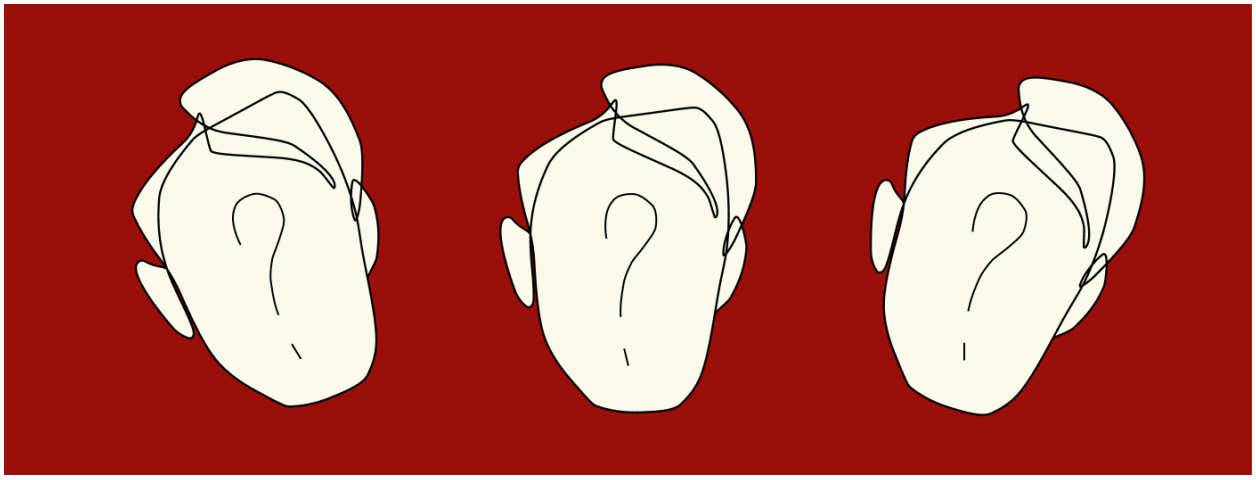

To answer these questions, I conducted interviews with past and present data workers. What became clear from these interviews is that this data is collected as raw material for AI training. These are not end consumers handing over data in exchange for a free service. These are employees and freelance data workers. Yet their most personal information is being collected in ways that would be considered objectionable, and in many cases illegal, in any regulated employment environment in the United States or Europe. The worker becomes the commodity being processed. The line between labour and resources disappears.

Methodology

This report asks: **In what ways is the personal data of AI data workers exploited within AI development pipelines, and how does the lack of informed consent and weak legal enforcement under data protection laws create new forms of digital exploitation?**

I conducted interviews with nine participants who have worked in data annotation and content moderation for Meta, TikTok, OpenAI and various AI annotation platforms like Appen, Remotask and Telus, some working from home and others employed directly from the outsourcing companies. I wanted to hear their stories in their own words and to understand the complete lifecycle of the workers' engagements when providing this data, from how they got the job, to what role they played in the aftermath. Every single person I spoke with worked for a foreign company—American or European firms that had outsourced their AI training to Kenya. That was the system. And I wanted to understand how that system worked from the inside: what data was being collected and how, whether workers were told what was happening to their information, whether they knew they had rights, what it costs them mentally and physically and what they thought needed to change.

The people in this report are known to me. They are my comrades. Some of them I worked alongside. Others I met through organising. I went to them, sat with them in their homes and workplaces, and listened. In order to protect their privacy and prevent potential retaliation, all interview participants in this inquiry are strictly anonymous.

The Pattern Emerges

Scope of Data Collected

The first thing I noticed, across almost every conversation, was what was being asked of these workers. Not just their time. Not just their attention. But their bodies, their homes, their children.

The mid-afternoon bustle of Nairobi CBD filtered into the restaurant we were in, a mix of clanking ceramic and the distant chaotic sounds from traffic. Across the table sat Joe, a data worker for three years at Sama. He seemed like any other professional taking a tea break, but his years in the data industry had left him with a weary perspective on digital privacy.

"During onboarding, it usually includes standard paperwork," he explained, "Identification, signatures, the usual." But as he settled into the role, the boundaries began to blur. He realized he wasn't just working on a computer, he was being monitored by it. Keystrokes, screenshots and performance metrics became silent shadows in his workspace. Then came the surreal request, the ones that did not fit the 'data worker' job description. The first time Joe told me they'd asked for pictures of his feet, I thought I'd misheard. I wanted to hear it in person, and he said it again: "My feet. No explanation. Just send us pictures of your feet," he said. The absurdity of it felt sharp: a worker selling his labour, only to have his most personal, mundane movements harvested for a purpose not known to him.

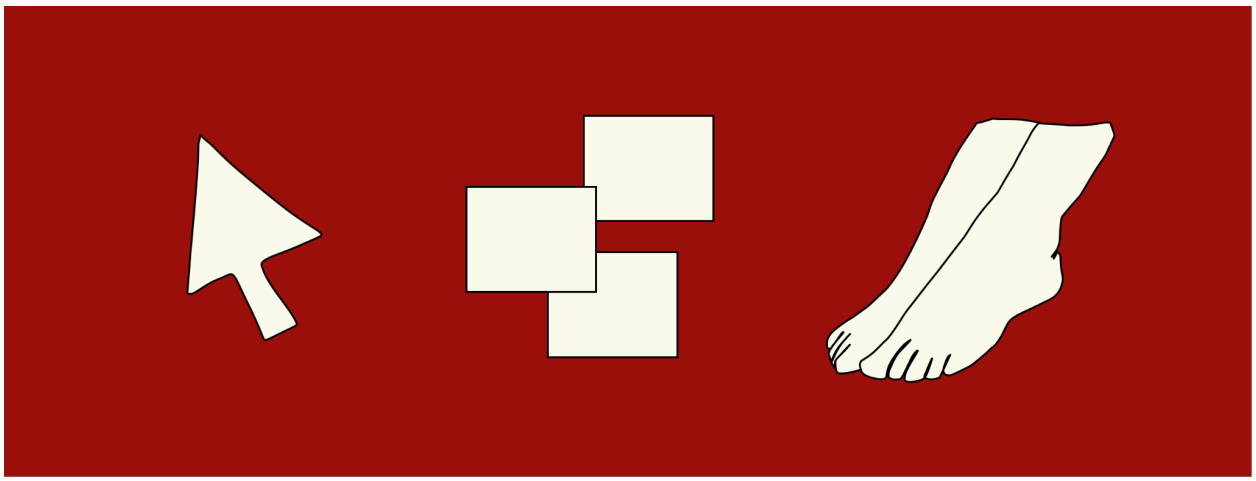

The atmosphere in Jenny's home was defined by the particular relaxation that settles in when a baby is finally asleep. Before she slept, I was shown her new toys, and told how they are meticulously dressed. I played ‘Uncle Mojez’ the protector, a knight protecting a group of small toys from a bigger one. After four years in data work, working for platforms like Remotask and Appen, Jenny navigates her home with the practiced efficiency of someone who knows they are always being recorded. Pointing to her covered webcam, she described her laptop camera not as a tool but as a silent supervisor. She detailed how eerie her routine had become, filming herself doing house chores, turning herself into a dataset. The most striking thing she told me was about taking pictures of her own child while asleep inside a baby cot. "I didn't even know what a baby cot was," she said. "I searched online and had one made so that I could do the tasks, because there were a lot that required pictures of babies sleeping." There was a haunting quality to the way she described it, capturing those private, vulnerable moments of motherhood and sending them as data points to a client who was invisible to her. Jenny's home was both a sanctuary and a production floor.

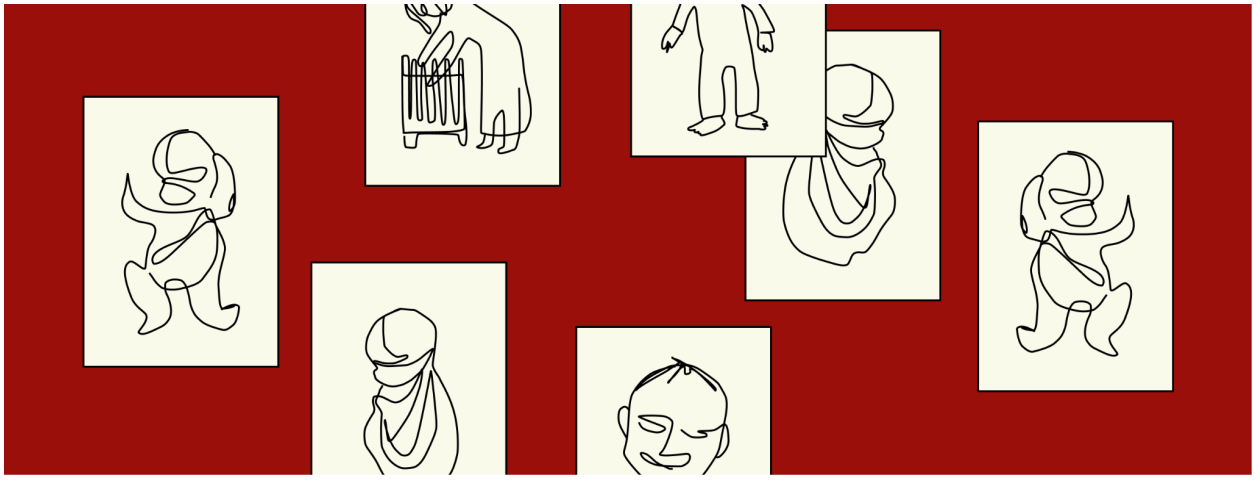

I met Stu after his shift from interning at a mechanic shop. The smell of engine oil from his work clothes still lingered in the air. He walked me through the whole process, pulling up Google Photos on his phone. "See," he said, "my own data, my face, is what I send for each task." He wasn't showing me a manual or a diagnostic report. He was showing me a gallery of his own face. "Look at my Google Photos. See all these? It's just page after page of me. I spent a whole afternoon taking photos of myself, looking up, down, tilting sideways, giving a big smile, then a serious one. Think of it like setting up your Face ID on your phone."

"I keep them all backed up and ready to go. When a task drops on the platform, I don't even have to take new ones. I just hit a link in the app, it hooks directly into my Google Photos, I pick ten different angles of my face, hit 'OK,' and I'm 25 cents richer. I don't even have to be somewhere controlled, even from under a car I can repeat this process and continue doing other things. The money sounds small, but when you already have the photos ready, it's fast money, and it compounds in the long run."

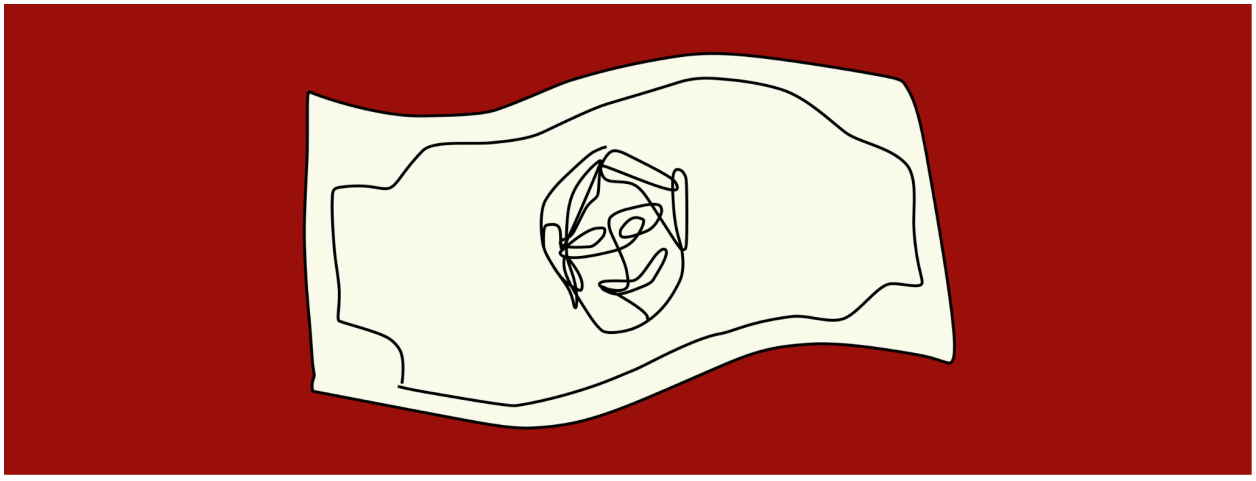

He paused. "The weirdest part, though, right before you finish, a prompt pops up. It asks if I want my photos sold or kept for 'personal use' by Telus. I always wonder about that. Like, if I have already sent them my face, who really owns it once I hit send?"

The cafe near Alexie's shop was loud, and her eyes kept darting to the clock. As a shop attendant, her time is measured in thin slices, and that day I only had 15 minutes before she had to disappear back behind the counter. But even in this hurried break, the story she told of Hugo Outsourcing, a global business process outsourcing (BPO) firm operating in Africa, felt like something out of a science fiction script. The company provides dedicated, university-educated professionals for customer support, e-commerce management, and AI data training. Alexie described a one-day gig that felt less like work and more like an audition for a ghost. For six hrs she had to act differently in various staged rooms, in a corporate office, a fake gym, a simulated kitchen, a mock sitting room and even a synthetic garden. "Once you have done signing up, you are escorted to different rooms where they have different backdrops simulating different environments, one room has a background of a gym, another room has a background of a garden, another one has a kitchen. In each room you are given

the clothes for the different background, i.e, gym clothes to simulate you are in a gym," she explained, leaning over to bite her chicken pie. "They take hundreds. Different angles, in every room they have built."

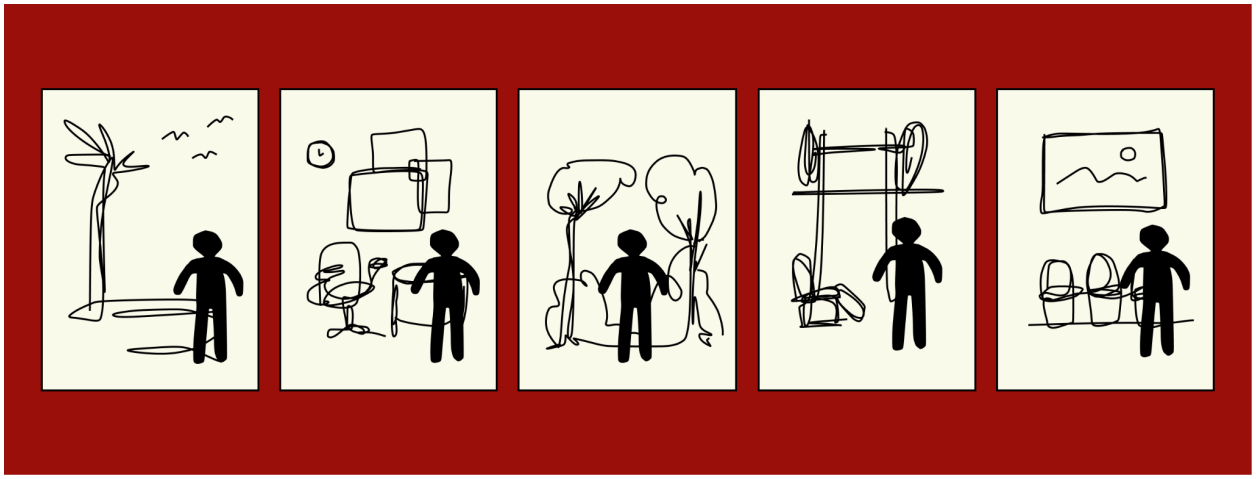

The process was clinical: she surrendered her name, age and ID and in exchange she spent half a day performing "existence" in these manufactured sets. And what did she get in return for sharing her likeness? USD 11 and a breakfast served in a disposable plastic pack. If you thought the staged rooms were unsettling, wait until you hear the recruitment. Alexie was urged to bring in more people—children, teenagers and the elderly. Every age bracket had a different price tag, and every referral put a little more money in her pocket. Sitting there as I took all this in, the reality was stark. By offering referral pay for children and the elderly, the company turned Alexie into a recruiter for her own community's biometrics!

The Consent That Never Happened

The air in Machakos feels thinner and calmer, quieter than the chaos of Nairobi. For Omwami, this move wasn't just about saving money, it was about escaping the pressure of a city that mirrored the high-pressure environment of the AI firms he left behind. After six years at Sama, cycling through roles as an English-Swahili translator, AI data trainer and content moderator, he spoke with the steady clarity of someone who has finally stepped out of the machine. Sitting at home with his family nearby, he reflected on the work he did

He stated that he had been informed that productivity metrics (e.g., task completion rates and quality scores) would be collected. However, the extent of

personal data collection was not always made explicit. When I asked him if he remembered signing consent forms, he paused. He'd been at Sama for six years, longer than almost anyone. “No,” he said firmly. “I do not recall being formally asked. Most data collections appeared to be automatic and tied to performance systems.”

I met Ethan, a data worker who for seven years served across multiple global outsourcing platforms namely, Sama, TELUS, CloudFactory and Appen. When he spoke about his entry into the data world, he didn't use words like ‘permission’ or ‘agreement.’ Instead, he kept coming back to a single word: *verification*. “It’s not that they do not ask,” Ethan explained, leaning back as he dissected his memory. “But they do not call it ‘giving away your data.’ They call it ‘verifying your identity.” He describes the high stakes of this ‘verification’, a process that required him to hand over the most sensitive anchors of his personhood. It wasn't just a signature on a page; it was government issued IDs and high resolution facial footage. In Ethan’s view, the consent was buried in the necessity of the task. You couldn’t ‘verify’ yourself without giving up your image, and you couldn't work without ‘verifying.’ By rebranding surveillance as security, the platform turned a massive data grab into a standard HR hurdle.

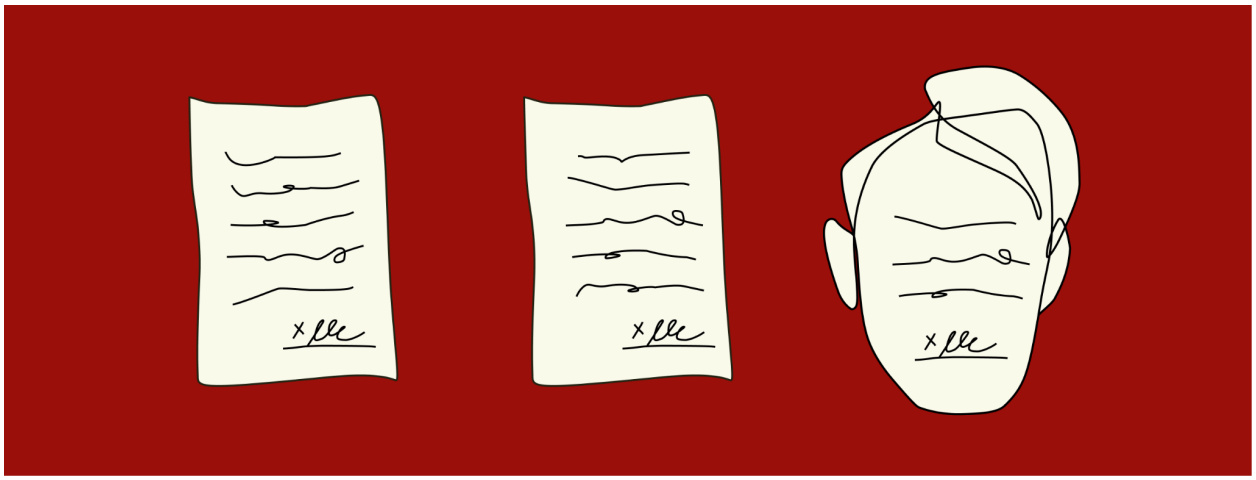

Stu, on the other hand, believed that consent to the use of his data was buried somewhere in the platform's terms and conditions. "There was no onboarding," he said. "Once you sign up to Data Compute, you start doing the tasks. There was a terms and conditions page that stated that our images would be taken." He said this so matter-of-factly, like it was just how things worked. But it's not supposed to work this way. Joe added that there was "no explicit, informed consent process regarding the collection of performance logs, photos, or other data beyond signing a general

employment contract. Meanwhile Jenny stated that no consent was asked for while doing the work.

Talking to these workers, I began to understand something: consent wasn't just absent, it had been weaponized against us. But for consent to be legally valid, it must be freely given. And once you look at the economic reality of data work, that condition is fundamentally unmet. Ethan called it verification, like he was being confirmed as a real person. But verification required his face, his government ID, everything. Could he refuse? Not if he wanted the job. That's not consent. That's coercion. The tool that is supposed to empower workers had been turned into an instrument of legal fiction.

Minimal Awareness of Legal Rights

Kibra sits right inside Nairobi, a massive, restless maze that holds the title of the largest urban slum in Africa. Locals call it "Chocolate City." If you stand on one of the hills and look down, you see thousands of corrugated iron sheets stretched out as far as the eye can see, deeply rusted and scarred.. Kibra is a place of relentless movement, but inside Kharim 's home, the air was thick with the quiet energy of new parenthood. His wife was resting nearby with their two-week-old baby. "That's my son right there," he said, his eyes lighting up as he glanced toward the bed. He had been in data work for four years, working at Sama for Meta related projects. When I asked Kharim about his rights, like the GDPR or the Kenyan Data Protection Act of 2019, he gave a dry, knowing smile. These were new terms to him. "I have been doing this for four years," he said, keeping his voice low so as not to wake the newborn. "But those laws? I have never seen them mentioned. Not once. The company doesn't put those protections in bold letters. They don't highlight your rights in the onboarding slides."

When I was with Stu, I pulled up the Data Protection Act on my phone. "This law says your data can't be collected without your informed consent." Stu stared at the screen. "I'm hearing this for the first time," he said. "Right now. From you." Jenny said she had heard of the GDPR, but was never given any information about whether her employer followed it. "As a worker, I was not informed of my rights or protections in relation to my data," she added.

Omwami was different. When I asked him the same question, he said, "Yes. I am aware of the Kenyan Data Protection Act 2019 and the EU GDPR. Rights such as

consent, access and deletion of personal data. However, these protections were not clearly applied or communicated to workers like me."

That contrast stayed with me. Kharim, Jenny and Stu had spent years doing this work with no idea that the law had anything to say about how their data was being handled. Omwami knew, but knowing hadn't protected him either. Legal awareness wasn't something the industry built into the job. It was something workers stumbled into on their own, usually after years, sometimes after harm had already been done. The companies were not going to tell you. You had to find out for yourself.

Unequal Access to Data Policies

Even among the workers I spoke to, there were differences. Omwami, Ethan and Joe had all worked in office environments, under contract, with health benefits. They had received at least some orientation on data policies, even if it was thin and incomplete. Stu and Jenny, working remotely on platforms, had received nothing at all. No information about how their data was stored, who it was shared with or what happened to it after they submitted it.

What struck me was that these workers were doing the same kinds of work, for the same kinds of companies. The difference was not the job. It was how visible they were. Remote workers were easier to ignore. The further you were from the office, the less the company felt obligated to tell you anything. And the less you knew, the easier it was to take from you.

Legal Gaps and Weak Enforcement

The office at Nzili & Sumbi Advocates hummed with the intensity of a firm challenging some of the largest corporations on earth. Shelves lined with thick legal volumes, the air charged with the meticulous work of building cases. Across the desk sat Kimani Njenga, a paralegal with a sharp, analytical focus. He didn't talk about "data work." He talked about "risk shifting."

"Are these practices legal?" I asked him. He leaned forward and grabbed a copy of the Data Protection Act 2019, tapping it. To Kimani, the problem isn't a lack of law. It's a calculated evasion of it. "The DPA is a strong baseline," he said, his tone clinical and direct. "But in the gig economy, law often gets lost in the cross-border shuffle. Big tech firms in Silicon Valley use local subcontractors in Nairobi—Sama, Teleperformance, smaller annotation firms—to create a compliance gap. They get the cheap labour, but they leave the local firms to shoulder the legal risks." He pointed to the hollow nature of the "consent" I had seen with Stu and Ethan. "Section 29 makes privacy notices mandatory, yet workers never see them. Subcontractors are under so much pressure to deliver fast and cheap that they simply cut corners on things like data minimisation."

I told him about Ethan's "verification." About Stu not knowing he had rights until I showed him his phone. Kimani nodded. He had seen this pattern before. There were already cases in Kenyan courts that established exactly what these companies were doing wrong. In 2020, when the High Court ruled that the government could not roll out the Huduma Namba digital ID without a robust data protection framework, the court was explicit that biometric collection must be necessary, proportionate and supported by clear legal safeguards. In May 2025, the same logic was applied to Worldcoin, whose biometric data collection was ruled unlawful. The court ordered the company to cease operations and permanently delete all collected data. The principle, Kimani said, applied equally to what annotation companies were demanding of their workers. If you cannot demonstrate necessity, you have no legal basis for collecting someone's face.

Then Kimani said something that has stayed with me. "The legal frameworks for data protection are individual-centric. They don't give workers a collective right to bargain over the terms of data collection or the use of AI systems. Unions and worker representatives are often excluded from these discussions entirely. And when workers

try to organise, we have seen what happens." He mentioned Daniel Motaung.[1] I already know Motaung personally. In 2022, the former content moderator filed a landmark lawsuit in Kenya against Meta and its subcontractor Sama, alleging exploitative working conditions, inadequate mental health support and union-busting. He had tried to organise his colleagues. The response from the company had been to shut it down. Hearing his name in Kimani's office, in the context of everything I had now learned, I felt the weight of it differently. Motaung had done what me and my comrades in the African Content Moderators Union were trying to do. The companies had shown clearly what they thought of that. It wasn't a precedent in a legal textbook. It was a warning that had been sent to all of us.

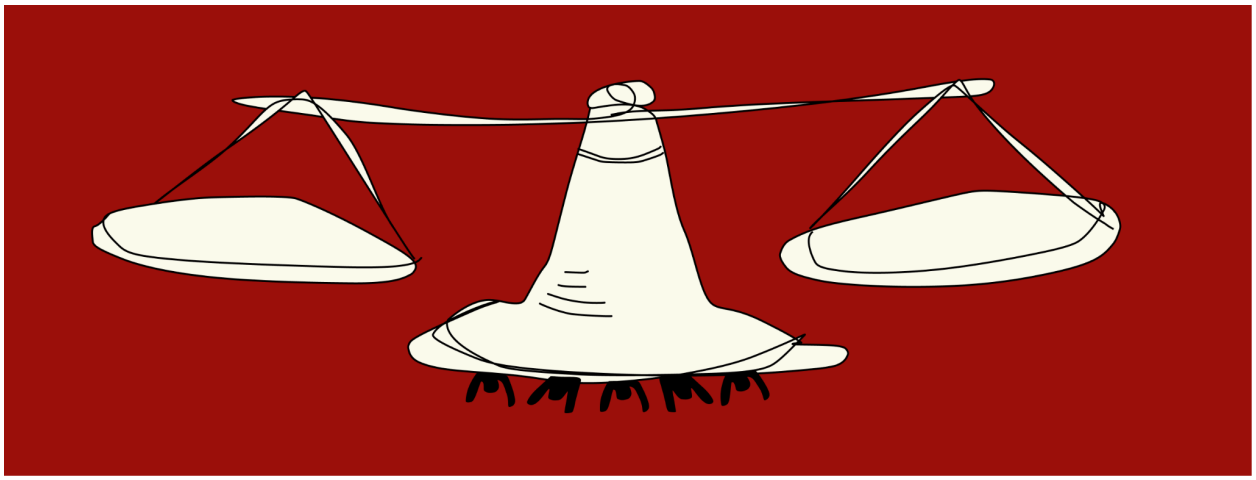

This, Kimani explained, was the deeper problem. The GDPR carries extraterritorial obligations, companies processing the data of people in Europe must comply regardless of where they operate. Kenya's Data Protection Act carries similar provisions for foreign entities processing Kenyan citizens' data. But enforcement requires resources, political will and workers who know their rights well enough to report violations. Right now, almost none of those conditions are met.

Later, I pulled up Article 9 of the GDPR on my phone, the same way I had pulled up the Data Protection Act for Stu. Processing of biometric data is strictly prohibited unless specific safeguards and explicit consent are provided. I read it three times. Everything Stu had described— the face photos, the prompt asking if Telus could keep

[1] Daniel Motaung launches world-first case to force Facebook to finally put content moderators' health and wellbeing ahead of cash.

his images, the $0.25 tasks—was illegal. Everything Ethan had called "verification" was illegal. Everything Joe had done without being told why was illegal. The law existed. It just had never been enforced where we were.

Psychological Harm

As I talked to more workers, something else emerged that I hadn't anticipated. The data was not the only thing being extracted. So was something harder to name and harder to return.

Joe was the one who made me understand the full weight of it. He told me he had suffered secondary trauma and stress-related health effects from prolonged exposure to harmful content, and that there had been no meaningful mental health support from his employer. Then he showed me a medical report. I read it carefully. He had been formally diagnosed with severe PTSD. He didn't want the details shared, and I will honour that. But I will say this: the document was real, it was clinical and it described conditions that had been caused by work he had done for a company that had since moved on without a backward glance.

Ethan put it simply: "I am a victim of the data I have worked on and I still have the scars to date, both mental and physical."

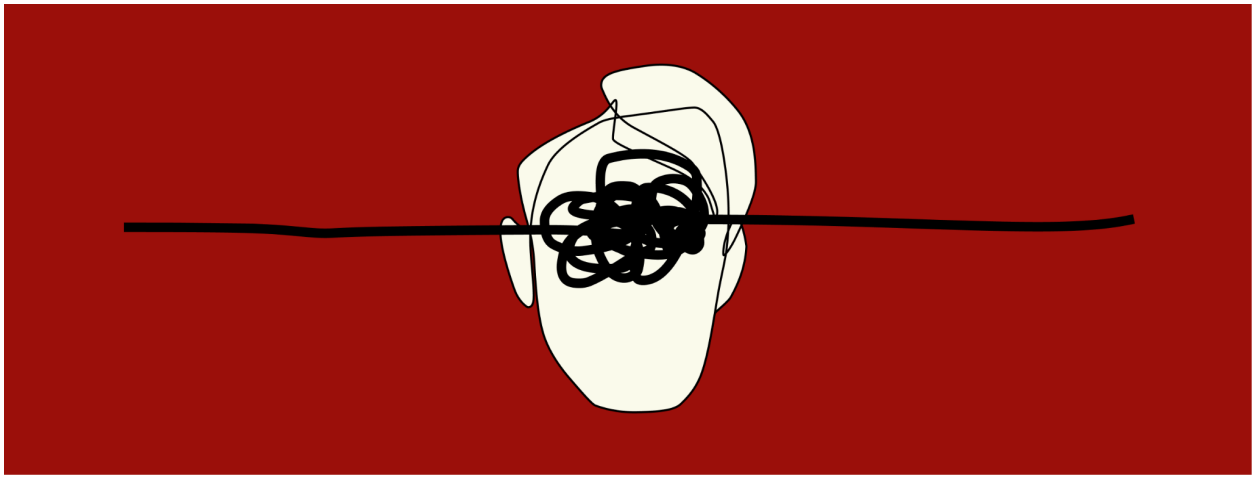

Stu's harm was different, quieter, but persistent. "I am usually worried about what my images are being used for out there," he told me. "They could be used to borrow for some aid money, or to train AI, I couldn't know. The harm comes from this worry." Jenny described becoming uneasy, mentally, thinking about how her data was

being used. "It caused stress and a sense of vulnerability," she said. Omwami said the constant monitoring had contributed to strain, and that exposure to harmful content during moderation had taken its own toll. "While not reputational, these practices affected my wellbeing," he said, and I noticed the careful, almost formal way he said it, like he was still trying to contain something.

What I kept hearing, beneath all of it, was a particular kind of loss; the loss of control over something deeply personal. These workers had sent their faces, their voices, their children's sleeping bodies, into systems they could not see, operated by people they would never meet. And once it was gone, it was gone. That uncertainty did not end when the contract ended. It followed them home.

The System Revealed

After nine interviews, the pattern was unmistakable. This was a system, and it had been designed to steal from workers.

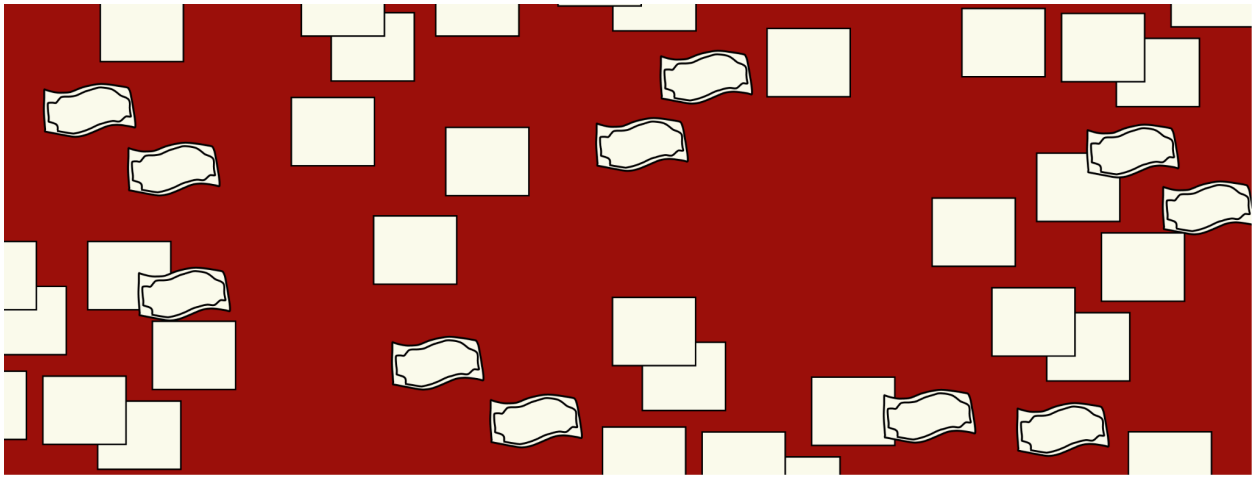

What I kept coming back to was: *these workers were being exploited twice.* The first time for their labour, the time and effort and attention they brought to the work, for which they were paid as little as $0.25 a task. The second time for their data, their faces, their voices, their children's bodies, their biological and behavioural characteristics, which the company appropriated as a permanent asset without their knowledge or meaningful consent. Unlike traditional labour, where the product of work is separate from the person doing it, here the worker's own identity was the product. That is something different. That is something that does not end when the shift does.

Because once that data is gone, it is gone. It can be sold, licensed, breached, repurposed. It could end up training commercial advertising systems. It could end up training state surveillance technologies that might one day be used against the very communities these workers live in. The worker cannot reclaim it. There is no mechanism for that. What was taken was not time or effort;those can be compensated. What was taken was a permanent piece of who they are.

And the system had been constructed to make sure they didn't know it was happening. Remote workers received no policy orientation. Consent was buried in terms and conditions nobody had time to read, and nobody had the power to negotiate. Workers were classified as independent contractors, which meant the company could argue they fell outside the strictest protections of both labour law and data protection law. In Kenya, this gap is particularly sharp. The country has a modern Data Protection Act, but labour laws that were never built to account for the gig economy. The result is a legal vacuum, and the companies know exactly where it is.

I went back to Article 13 of the GDPR and Section 29 of the Data Protection Act. Both require data controllers, the outsourcing firms, to inform workers of their rights, the purposes of processing, and with whom their data is being shared. In practice, not one of the nine workers I spoke with had received that information. Article 7 and Section 32 require that consent be informed, specific and freely given. Bundling it into a lengthy terms and conditions document, that a worker must accept before they can start earning, is not consent. It is a legal fiction dressed up as a formality.

And I do not believe this was carelessness. I believe it was intentional. Because why else would you collect biometric data from your own workforce without proper onboarding, without clear purpose, without telling people what you are taking and why? The answer is that it is cheaper and easier not to. The companies made a calculation. They decided that the cost of compliance was higher than the risk of getting caught. And in a context where enforcement is weak, workers are isolated and most people don't know they have rights, that calculation has been paying off.

Moving Forward

DSA and AI Workforce Demands

I did not come out of these nine interviews feeling helpless. I came out of them feeling certain. Certain that what had been done to us was wrong, that it was illegal, and that it was not going to stop on its own. We need an informed workforce. So let me be clear about where I stand and what we are asking for.

The stories of Joe, Jenny and Kharim are evidence of a structural void where the human body is harvested as data without recourse. To bridge this gap, I founded Data Safety Alliance (DSA), an organization committed to transforming platform and AI workers from passive datasets into informed advocates. We focus on education, training and rights based advocacy to move beyond the 'automatic' surveillance described by Omwami and the legal silence felt in the alleys of Kibra. Through strategic collaborations with labour unions and associations, digital rights activists, and data protection experts, DSA ensures that the human labour powering platforms and AI is no longer invisible. We are building a future where those behind the veil of AI possess the knowledge to protect their own dignity, ensuring that the progress of technology is never purchased at the cost of human privacy.

This report is part of that work, not a document handed to institutions in the hope that they will act, but a record of what has been done to us—in our own words—that we will use to organise and to fight.

These are our demands.

1. **To the companies - Telus, Sama, Teleperformance, Data Compute, and every other firm in this supply chain**

Stop collecting personal data from workers as a condition of employment. Stop burying consent in terms and conditions that workers have no power to negotiate. Stop classifying workers as independent contractors to avoid your legal obligations. Stop exposing workers to traumatic content without genuine, independent mental health support. These are not requests for improvement. They are the minimum requirements of operating lawfully and ethically. You have been falling short of them for years while your clients, the AI companies at the top of this supply chain, have looked the other way.

1. **To the AI companies at the top of the supply chain - the ones whose names appear in the marketing but not in the contracts, Meta, Open AI, ByteDance, Anthropic**

You do not get to claim ignorance of conditions you have designed. The subcontracting structure is not an accident, it is how you access cheap labour while insulating yourself from accountability. That ends when you take responsibility for every worker in your data pipeline, regardless of how many layers of outsourcing separate you from them. Pay a living wage. Provide stable contracts. Uphold the right to organise. Publish your supply chains so that workers and regulators can see them.

2. **To the Kenyan government and the Office of the Data Protection Commissioner**

The Data Protection Act 2019 is a strong law. Enforce it. Conduct proactive audits of BPOs and annotation firms operating in Kenya. Issue clear guidance that consent obtained from workers in conditions of economic precarity does not meet the legal standard of freely given consent. Update the Employment Act, it is nearly two decades old and it was not built for the gig economy. Workers in this industry need a legal category that gives them the protections they are currently denied.

3. **To fellow workers**

You have rights. The companies will not tell you this, because it is not in their interest for you to know. But the law exists. A data subject access request is your right.

Knowing what data has been collected on you, and demanding to know how it is being used, is your right. Organising with your colleagues is your right. If you are in Kenya or are a platform worker and you need support, the Data Safety Alliance exists for you.

I started this research because my cousin showed me his face on a Google Photos grid and didn't understand what he was giving away. I finish it knowing that what happened to him, to Joe, Jenny, Stu, Ethan, Kharim and Omwami. None of it was inevitable. It was a choice that companies made because they believed we didn't know enough, and weren't organised enough, to stop them. We are working on both of those things.

Support a session

To those who want to support us: From Inquiry to Action

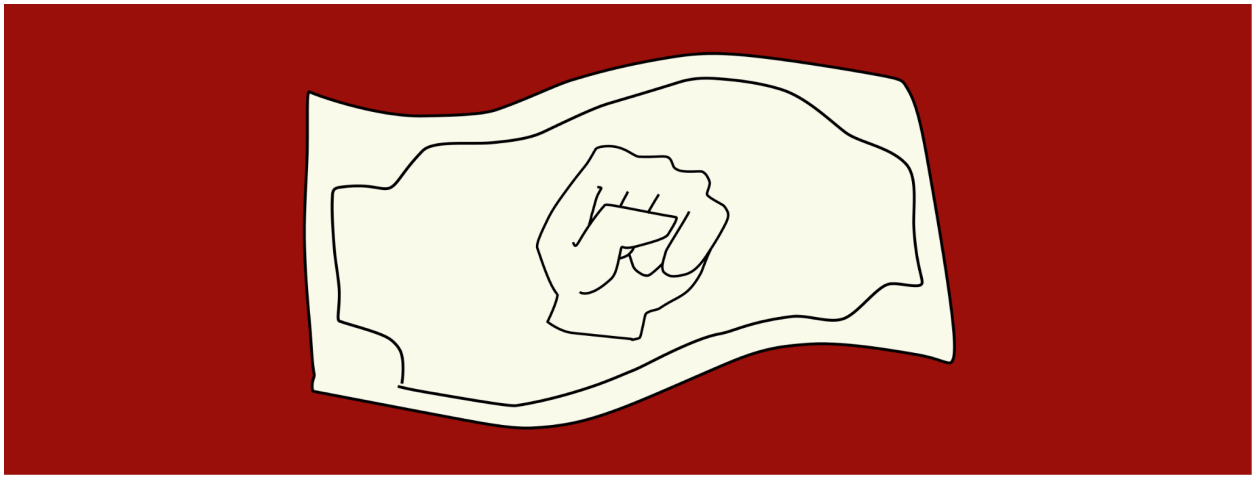

Just knowing about the problem is not enough, and this report makes that clear. If we want to stop our data from being taken, we need to take real action where it matters most. That's why Data Safety Alliance is growing our grassroots data rights clinics in Nairobi.

I began this work quietly and informally as I didn't have the proper qualification and of retaliation from the companies doing this work.. ButI recently became a Certified Data Protection Officer and now have the confidence to run the clinics openly. The problem is huge and we are growing our team to face it directly. Our team now includes two local community organizers and myself. We meet the workers where they are, run

hands on sensitization clinics on matters of data protection and help them file Subject Requests under the Data Protection Act. We support them in taking back their data from these platforms.

Right now we are raising $11,000 to launch our expanded Nairobi Pilot. With this support, our team will run 20 focused clinics and help 400 workers move from being unaware to becoming an informed, legally protected group.

ITEM	**COST**
Clinic Logistics (Venue hire, High-Speed Data for SARs, Handbooks)	$3,000
Organising Team Support (Lead DPO, Community Organizers)	$5,000
Participant Facilitation	$3,000
Total (20 Clinics)	$11,000