

No. 261/26

9th October 2026

TO ALL BRANCHES

Dear Colleague,

CWU Tech / UTAW – AI Workers Inquiry

The CWU / UTAW Tech Workers' Inquiry was officially launched at Labour Party Conference last week with speakers including Dave Ward, GS, Neil Duncan-Jordan MP and UTAW members and branch officials.

The inquiry outlines how generative AI is changing work across the tech sector, written with the perspectives and experiences of CWU members who build, deploy, manage, evaluate and use this technology. It warns about a world of work made worse in every way by AI and demands worker-first solutions.

The report contains proposals which expand upon [Composite Motion 13](#) carried at TUC Congress 2026 - *C13 Artificial Intelligence at work – a just transition through worker voice for the benefit of workers and society*.

The CWU are leading the way when it comes to the trade union movement debate on AI, with our composite motion unanimously passing at TUC Congress, as well as the only Labour Party Conference motion on AI being moved by the CWU.

The website for the inquiry can be found [here](#) and the report pdf is attached below.

Yours sincerely,

John Chadfield
National Officer – Tech



Final PDF Doc -
AI_Workers_Inquiry_

TECH WORKERS' INQUIRY:

AI

VOL 01
SEP 26



EXECUTIVE SUMMARY	4
POLICY PROPOSALS	7
INTRODUCTION	13
ISSUES THE INQUIRY SURFACED	18
GLOSSARY OF TERMS	28
METHODOLOGY	30
ENDNOTES	32
ACKNOWLEDGMENTS	35

AN ALTERNATE BRAND

END

THE END OF THE LINE

U.T.A.W.

UNITED
TECH & ALLIED
WORKERS

EXECUTIVE SUMMARY

TECH WORKERS' INQUIRY: AI					5	
AI is widely used to describe a disparate, and in many cases contradictory, set of technologies. While the term defies consensus, a clear picture of its impacts is emerging. Our participants described it as a largely negative force reorganising their work: increasing output expectations, changing skills requirements, expanding monitoring, degrading expertise, shifting accountability, and creating divisions between colleagues. AI applications rarely remove work; they redistribute and intensify it. Workers stressed that positive uses depend on their ability to decide for themselves how to engage with AI.¹ The central problem is not simply whether AI works well, but how it is introduced, controlled, measured and used by employers. Put simply, it is useful only when one gets to choose if and how to use them. The problems identified are not derived from any technical limits of today's models. Poor quality outputs, hallucinations and weak code do have an impact, but they are only part of the picture. This is because employers use AI to intensify work, rationalise redundancies, eliminate junior hiring, increase surveillance, deskill and downskill workers, shift risk onto employees and weaken their agency. Even as models improve, the underlying workplace issues will remain if workers do not gain enforceable rights over how AI is used, at work and in society at large. Tech workers are deeply concerned over AI's harmful impacts on the environment, in military and state surveillance contexts, and on mental health. This is why the Inquiry's			recommendations focus less on functionality and more on who controls and benefits from it. Workers must be actively empowered to defend and extend their rights. This includes enforceable power over how AI is introduced, used, monitored and evaluated at work. Workers must be given the final decision over how, or if, they use AI tools at work. There must not be a penalty for choosing not to use it. Consultation and transparency matter only if they lead to real accountability and effective action. Disclosure without consequence only whitewashes. The emerging policy directions are: • Transparency alone would not be enough; workers need meaningful consultation and negotiation over AI deployments. • AI use must not be mandatory; workers need legal protections from detriment for refusing AI use, and employers limited in AI-driven management decisions. • Safeguards against AI externalities; strong regulations to ensure the AI we build does not harm people and planet. • Stop the brain drain; establish and defend AI-free learning pathways for junior and graduate workers. • Workers get the final say; bake in collective bargaining over AI at work.			
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES



POLICY PROPOSALS

These proposals expand upon the composite motion carried at TUC Congress 2026 titled Artificial Intelligence at work, moved by the CWU, seconded by Aegis and supported by: UNISON; National Education Union; USDAW.²

NEW REGULATORY
AUTHORITY ON AI AT WORK

We need a regulatory framework with teeth to hold employers and technology companies accountable for their use of AI, along the following lines:

1. Establish a new statutory authority for AI at work, with powers to investigate, enforce and sanction employers and technology providers.
2. The authority should cover all AI systems, including genAI tools, automated decision systems, monitoring tools, recruitment systems, productivity systems and performance management tools.
3. Workers and their unions should be able to report AI-related harms directly to the authority, including unsafe deployment, discriminatory outcomes, excessive monitoring, misleading productivity claims, work intensification, deskilling and downskilling, job displacement, and ethical and environmental concerns.
4. The authority should have the power to require employers to disclose AI systems, pause deployment, carry out impact assessments, consult workers and their unions, remedy harm and withdraw deployed systems.
5. The authority's governance board should include expert worker and trade union representation equal to the number of representatives of employers, civil servants and academics.

SKILLS, TRAINING AND
JUNIOR PATHWAYS

Skills and development must be prioritised where AI is deployed to eliminate AI-induced deskilling and downskilling:

1. Protect junior roles and learning pathways by ensuring AI does not replace training, mentoring, supervision or foundational skill development.
2. Employers should not use AI as a substitute for subject-matter expertise, proper staffing or professional development.
3. Workers should have protected time to learn, practise, review and develop skills without being forced to rely on AI tools.
4. Employers must provide funding and time for reskilling and upskilling where AI changes the nature of work.
5. Workers should not be required to use AI in ways that undermine their professional judgement, craft, confidence or ability to learn.

**RIGHT TO REFUSE,
QUESTION AND CHALLENGE**

Workers must retain the right to refuse and be protected from retaliation in doing so:

1. Codify a right to choose whether and how to use genAI tools where there are professional, ethical, environmental, equality, accessibility, safety, religious, mental health or quality concerns.
2. Workers should be protected from detriment for refusing, questioning, criticising or reporting concerns about AI use or development.
3. No worker should be treated as resistant, underperforming or unsuitable for promotion because they raise concerns about AI.
4. Workers should have the right to human review where AI is used in decisions affecting their work, pay, performance, promotion, discipline, hiring or redundancy.

**ENVIRONMENTAL AND
SOCIAL ACCOUNTABILITY**

The negative externalities of AI must be thoroughly assessed and disclosed:

1. Require employers to assess and disclose the environmental impact of workplace AI, including energy use, water use, data-centre dependency, infrastructure demand and supply-chain impacts.
2. Environmental reporting must not be reduced to narrow carbon accounting. It should include the wider social and material costs of AI expansion.
3. Public policy on AI must include worker voice, environmental accountability and democratic control, not only innovation and productivity claims.
4. Classify genAI models as a dual-use technology and strengthen military exporting license criteria enforcement.

TECH WORKERS' INQUIRY: AI						10
HEALTH AND SAFETY Health and Safety legislation must be updated to take into account the new types of psychological harm that AI brings to the workplace: 1. Require impact assessments before deployment and whenever AI systems are substantially changed. Impacts must consider workload, wellbeing, equality, accessibility, job security, skills, training, environmental and ethical impacts. 2. Legislate a right to disconnect 3. Where AI genuinely creates productivity gains, require employers to ensure that those gains benefit workers proportionately (e.g. through reduced working time at the same pay, paid time off for training, pay rises).			MONITORING, PERFORMANCE AND SURVEILLANCE Workers must be protected from disciplinary and performance processes which are based on how much they're using (or not using) genAI tools to complete their work: ³ 1. AI usage data must not be used for discipline, promotion, redundancy, performance scoring or productivity benchmarking unless this has been clearly disclosed, collectively agreed and independently assessed. This includes AI leaderboards, token rankings or usage dashboards that pressure workers to use AI. 2. Require employers to prove that AI-related metrics measure useful, safe and high quality work, not only activity, cost, volume or speed. 3. Workers must not be penalised for not using AI enough. 4. Prohibit opaque AI-driven or AI-assisted performance evaluation where workers cannot challenge the data, method or outcome.			
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES

PROTECTION AGAINST
REDUNDANCY AND
OFFSHORING

Redundancy protections must be strengthened, employers cannot use AI adoption as a convenient cover for offshoring and outsourcing:

- 1. Employers should not be allowed to use AI adoption, productivity gains or automation forecasts as a shortcut to redundancy without independent evidence, consultation and negotiated safeguards.
- 2. Strengthen redundancy protections by increasing minimum consultation periods and expand the definition of “business unit” to the whole employer for collective processes.
- 3. Introduce stronger penalties where employers rationalise redundancies with AI while later rehiring, outsourcing or redistributing the same work.

CONSULTATION,
TRANSPARENCY AND
DISCLOSURE

Employers must be compelled to involve workers meaningfully in implementation decisions:

- 1. Introduce a statutory duty to meaningfully consult workers before AI tools are introduced, expanded, made mandatory or substantially changed in any workstream.
- 2. Where collective bargaining exists, employers must have a duty to negotiate over AI deployment and its effects on monitoring, productivity expectations, staffing levels and changes to job design.
- 3. Where no collective bargaining is in place, workers must have the right to request a pause, review or withdrawal of AI systems where there are reasonable concerns about detrimental impact, with an escalation pathway to the regulator.
- 4. Employers must disclose AI systems used in the workplace in documentation available to staff, including their purpose, reason for adoption, what data is collected, who can access the data, how long it is kept.



INTRODUCTION

This report investigates AI from the point of view of tech workers who build, deploy, evaluate, manage and use these systems. It is a worker-led and worker-generated intervention into public debate on AI, based on the direct experience of tech workers across the tech workforce.

Unhelpfully, the term AI does not refer to a single technology. It is an umbrella term encompassing different technologies marketed together. As our report will show, some of these technologies have beneficial use-cases and some do not. Some do not even exist and may never, despite forming a major part of public debate on AI.

Our inquiry starts from a basic premise: tech workers are not only affected by AI, they also know how these systems are built, introduced, monitored and used in practice. Their experience is essential to any serious discussion about AI, work and regulation.

WHY A TECH WORKER INQUIRY ON AI?

Who has shaped what we think we know about AI? Almost invariably, it is non-technical media pundits and the tech entrepreneurs who own AI companies. Sometimes, the latter started off as software engineers and scientists; always, they are super-rich businesspeople standing to profit from AI proliferation.

To our knowledge, this report is the first systematic synthesis of the experience of AI from a broad swathe of the tech workforce. Our inquiry covers the UK tech sector supply chain, including customer support agents, research scientists, software developers, cybersecurity analysts, tech educators and more besides.

Our worker inquiry leverages the technical expertise of employees throughout the supply chain. The approach is no-nonsense, focusing on concrete daily experiences, cross-pollinates insights from a diversity of workers, and interrogates underlying dynamics.

As a class, the workforce which develops and deploys AI is arguably the ultimate authority on the topic. As a group of unionised and unionising workers, the voices in this report speak in terms of progress, fairness, and defence of the vulnerable. Tech C-suites, representing capital's avaricious vision for genAI technology, have always been clear: move fast and break things. There is a long line of "things" who have been broken by the barons of Silicon Valley, from teens like Molly Russell to war-ravaged people in the eastern Congo. Tech workers, the intellectual and manual labourers behind this new set of technologies, are concerned that AI is not compatible with people and planet.

WHAT DO WE MEAN BY "AI"?

AI has been described as a "suitcase word", so roomy you can put anything you like in it. Such imprecise language is unhelpful as, in practice, AI encompasses a range of very different technologies, with different risks, environmental impacts and benefits. To be specific, when one speaks about the dangers posed by AI to mental health, cybersecurity or job automation, one is speaking of large language models (LLMs) and agents (LLM-based systems). These are the most energy intensive.

When one speaks of AI breakthroughs and beneficial applications, for example the Nobel Prize-winning Alpha Fold or weather forecasting models, one is talking of a completely different technology. This tech is task-specific (i.e. narrow, well-defined tasks) and uses smaller models whose energy requirements are negligible.

Using the same word to refer to both technologies muddies discussion of their respective risks and benefits. The term generative AI (genAI) is a more accurate catch-all for LLMs, vision-language models (VLMs), and any other machine learning technology which synthesises output based on training data (including audio and video).⁴ In this report, we aim to be as precise as possible in our reference to AI.

TECH WORKERS' INQUIRY: AI					15	
PHILOSOPHICAL TRENDS IN AI Most of the debate about AI centres on artificial general intelligence (AGI), an ill-defined theoretical technology that doesn't exist.⁵ This debate encompasses many questions (e.g. how to define AGI, AI doomers vs AI accelerationists, whether to pursue AGI at all) which we will not deal with here as they are largely a distraction from the real-world applications of genAI. The lack of a definition for AGI pushes discussion of its risks to an indefinite future moment, ignoring the immediate and near-term risks posed by genAI and its increasing integration into the systems that underpin society. Perhaps of some relevance is the debate between AI optimists and pessimists. The former believe genAI (and by extension the elusive AGI) can solve the world's greatest problems, the latter do not. While tech workers do not share a unified opinion here, their response to the bleeding edge of genAI implementations is generally pessimistic. GenAI in practice is a technology of automation, and as with all automating technologies of the past, the chief beneficiaries of efficiencies are typically the super-rich company owners. On the one hand, employees whose work is being transformed by genAI report a variety of harms to their well-being.⁶ On the other, the tech boom of the past 30 years and the AI bubble of the past decade have enriched an exceedingly small number of billionaires and turbo-charged their (decidedly far-right) political influence. Meanwhile the majority of working people have faced year-on-			year income loss for several years running and inequality has drastically been worsening for decades. It is fair to say, therefore, that genAI is a technology which benefits billionaires now far more than it promises to benefit (or wipe out, depending on which type of investment they are angling for) everyone else at some undefined point in the future, and that this reality is reflected in who is hyping AI and who is suffering from its deployment. Workers and small business owners are both facing the same enemy: the Big Tech owners of frontier models, at once forcing genAI onto the labour force and renting their models out at increasingly exorbitant prices. Additionally, both supporters and critics speak of genAI as a technology that is here to stay. This inevitability framing benefits C-suites and their shareholders by transferring political agency from people to "the market", aka the super-rich owners of the technology. Why should workers accept this narrative uncritically? GenAI is a force with social, economic and mechanical impacts that can be understood, reckoned with, regulated, and prohibited where necessary. Any particular genAI deployment, and its very existence as a widespread technology, should not be taken as a given.			
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES

A NOTE ON AI-LED REDUNDANCIES

There have been productivity gains widely touted by AI optimists. We prefer to call it what it is: work intensification. Some commentators have described this as the Jevons paradox.⁷ The more efficiency is gained, the more job creation is stimulated. This half truth (efficiency gains) is a great lie, to paraphrase Benjamin Franklin.

Companies are reducing headcounts and avoiding hiring, with the remaining work distributed to fewer workers under tighter deadlines. At first glance, companies are carrying out redundancies (both mass layoffs and small regular redundancy exercises) and claiming this is because “AI is replacing the need for workers.” Workers on the inside report that layoffs are more often to maintain revenue in the face of rising inflation, to increase company valuation by manipulating financial indices, or even to cover the cost of genAI platform subscriptions.

The element of work which the Jevons paradox actually applies to is text synthesis. If a worker’s only job is to write as much copy as possible with no regard to quality, their job could truly be automated. These so-called AI efficiencies are an excuse to obfuscate a reality where skilled staff are replaced with new tools that do not work well while workers who survive redundancy pick up the ever-intensifying slack.

Taken together, genAI adoption means harder work on less enjoyable tasks, ironically the opposite of AI boosters’ promise that “AI will do the boring tasks and leave the fun ones for humans.”

WORKER CONTROL IN THE FINAL ANALYSIS

Ultimately, the best-placed people to regulate genAI are the workers who build and use it, because they are the ones who most immediately enjoy any benefits and suffer any consequences. As ordinary working people, tech workers are also not insulated from the technology’s impacts on other parts of society – unlike the billionaire owners of the technology. We reject the debunked premise of “pro-worker AI”⁸ and in this report have put forward proposals that are worker-centric.

The logical conclusion of this report is not a new one, and was eloquently put forward by Wendell Phillips over 150 years ago: labour is entitled to all it creates. If we continue to allow tech barons to arrogate to themselves the power and profit, we will continue to see the worst that genAI can do to exploit and dominate the majority. On the flip side, if we place power over AI technologies in the hands of the workers who create and deploy them, society will reap any rewards while avoiding the harms. Workers should be empowered to control AI at work.



ISSUES THE INQUIRY SURFACED

“
I started
enjoying my
job again as
soon as I
stopped using
AI to code
”

Software Engineer

DEGRADING WORK SATISFACTION

Workers described being moved away from intellectually stimulating tasks and into prompting, reviewing and correcting low quality output.

Employees who code are experiencing a qualitative transformation of their work from problem-solving to supervising LLMs. Many workers said they no longer feel “in the flow” of coding, instead supervising models and evaluating outputs. More often than not, so much reviewing and correction is required to avoid approving poor output that time is lost rather than saved.

Software Engineer

“

The more you use a coding assistant, the worse you get at doing it yourself.

”

DESKILLING AND DOWNSKILLING

Downskilling was a major concern across the board. Workers reported having to look up how to do things they previously would not have needed to after prolonged genAI use. This was especially visible in coding work, but not limited to it.

Junior workers and those without strong foundational skills are becoming overly reliant on tools such as agents and assistants. There is no incentive or time for juniors to learn the necessary technical and critical thinking skills; work intensification means they are increasingly being denied a sounding board to chat through problems. Lacking expertise, they are badly placed to judge LLM output quality.

While some senior workers also report downskilling when pushed to use LLMs, some reported feeling more confident in using LLMs as a tool to support their coding, rather than a substitute for their own thinking. Due to their skills being built up before AI

proliferation, they are more able to resist cognitive surrender⁹. As with juniors however, they also recognised that their ability to review and write the code themselves was degrading with LLM use.

“Because AI aggregates and distills information, often missing important details and nuance, there is no opportunity to learn through exposure and mistakes.”
– Data Analyst

Employers are using LLMs as a substitute for training their employees, obscuring gaps in subject-matter expertise and skills. Where LLM research assistants are being mandated, these replace the learning process, leading to unrealistic expectations that someone can become an expert in short order.

The long-term trajectory is a loss of expertise and skills across the sector. Additionally, proliferation of LLM output in open source libraries means that documentation volume is increasing while quality is degrading, a further obstacle to effective learning for all career levels.

Junior Software Engineer

“
AI is baked into our KPI’s now.
Remuneration and bonuses are
contingent on AI use.
”

**MONITORING, DISCIPLINE
AND PROMOTION**

Monitoring of genAI tool usage by management is widespread, at times to the point of being described as surveillance, and takes diverse forms. Workers report being penalised for not using them enough and for using them too much, sometimes in the same company. Tokens have replaced output quality as a performance metric, linking career progression to tool use at the expense of actual competency. In some cases, work done using genAI is also monitored and evaluated by LLM-based evaluators whose inability to accurately understand human answers leads to unfair performance outcomes.

“We were told to use as many tokens as possible, then my colleague was slapped on the wrist for using too many.”
– Senior Data Engineer

In most cases, participants reported opaque data collection processes, with a lack of clarity over what usage data is being collected and how it is being analysed. This deepens the asymmetry between staff and managers in performance review and subsequent disciplinary processes, with workers shooting blindly at a bullseye they cannot see.

EXTERNALITIES AND
WIDER HARMS

Participants repeatedly challenged the idea that what is good for business is good for society, and asserted that the environmental, psychological and political impacts of adoption must be included in company decisionmaking. Even where a particular tool’s adoption is rational and evidence-based, workers say that externalities are ignored.

Many participants also reported feeling guilt over being mandated to use genAI at work when it contradicts their personal beliefs due to the harms caused by its development and maintenance.¹⁰ This cognitive dissonance causes stress that extends beyond the working day for many.

Participants from frontier labs, who develop cutting edge models, report a deep anxiety over how their work is used in military settings. Despite multiple internal escalations, external whistleblowing and public protests, research scientists and engineers describe a senior leadership that persistently denies or wilfully ignores the existence and scale of this problem¹¹.

“I want the AI I build to benefit humanity, not to facilitate a genocide.”
– Research Scientist

Workers raised concerns that the environmental impact of data centre compute (as opposed to human intelligence) is never considered. Workers questioned the underlying logic of burning more fossil fuels so that they can do more work faster and to a poorer standard.

“
Even when it is good,
it is not clear that it
isn’t still bad overall.
”

Software Engineer

IT'S NOT
JUST A
MILITARY
TRAP

IT'S A
TRAP

IT'S A
TRAP

IT'S A
TRAP

IT'S A
TRAP

IT'S A
TRAP

IT'S A
TRAP

UNILATERAL
DECISION MAKING

Consultation was absent in most adoptions (with some notable exceptions), both where usage is trivial and where it is deeply integrated across work processes. Participants said that senior management teams do not understand the tools they are mandating, using staff as a testing ground for tools that often slow down or worsen the work they claim to speed up or improve.

Participants described genAI tools introduced via executive enthusiasm, investor or client pressure and market panic (i.e. “Everyone else is doing it”). Even when consultation did take place, concerns raised by workers were either not taken

seriously or wholly dismissed. In one notable instance, cybersecurity experts in a company raised concerns over a weakening of the company’s cybersecurity due to LLM-generated code but were ignored.

“Tool adoption is often “vibes-based” with no evidence of improving quality or efficiency. Senior leadership won’t consult their own experts before introducing new tools.”
– Cybersecurity Expert

Workers are not asking to simply be told which AI systems are being used. They are experts asking to be involved before decisions are made, and to have the power to change, pause, limit or reject unproven deployments.

“
Nobody was consulted.
The new CEO at an all-hands
announced 40% layoffs
at the same time as
announcing going full AI.
”

Program Manager

INCREASING SOCIAL DIVISIONS AT WORK

Widespread adoption is creating a number of divisions in company workforces. Firstly, by increasing the technical expertise gap between juniors and seniors. Secondly, workers at the lower tiers of workplace hierarchies are subjected to more monitoring and evaluation than those at higher levels. Thirdly, between those producing output faster using LLMs and those who must check, correct, repair or absorb the consequences. Finally, reliance on LLM tools and agents is reducing collaborative work, with participants reporting increasing isolation at work.

“We struggle to meet performance expectations when we’re inundated with poor quality pull requests from junior colleagues and monitored by non-technical product owners who ignore bug proliferation.”
– Data Engineer

Additionally, participants report a deepening division between project/program managers and those in technical roles that carry out the work. As LLMs are increasingly used to estimate task duration or viability, they are often at odds with the experience and expertise of those tasked with the actual work. This leads to an “us versus them” feeling between those who use LLMs to ideate and those who deliver the project.

“

My level of trust in work from colleagues has dropped a lot. It means I put a lot more effort and time into reviewing and understanding other peoples’ code.

”

Senior Software Engineer

WORKLOAD

Across the Inquiry, workers did not describe genAI as reducing workload. They reported work intensification: increased output volumes, tighter deadlines and higher productivity expectations.

“Working with agents is like using a slot machine, output could take a minute or an hour to arrive, and you don’t even know if it’ll work.”
– Third Line Support Engineer

As already outlined, the quality of work is transforming from productive problem-solving tasks into LLM supervision tasks; code is generated much faster than a human can check it. Workers described additional labour reviewing code, debugging issues, checking outputs and repairing problems produced by LLMs. In most cases, this additional labour is hidden, unseen by management and not reflected in KPIs.

Participants in education and academic work blame genAI misuse by students or colleagues for creating

unpaid extra work, such as writing detailed reports to prove misuse, handling appeals/resits, or correcting low-quality output material. In software and cybersecurity, genAI produces more code, tickets, reports or incident write-ups than skilled workers have time to inspect and correct. In customer support areas, text production speed similarly creates additional editing and quality-control work. The temptation to further delegate these supervisory tasks to genAI is powerful, but the burden of responsibility for poor outcomes lies with the worker, never with the tool.

The overall effect described is of work intensification through a positive feedback loop of high productivity expectations generating volumes of unreliable outputs which must be reviewed by human or assistant tools. This loop itself creates a catch 22 situation: the conscientious employee can break the loop only by slowing down to review outputs properly, risking her own performance metrics. But if she follows the the path of least resistance, she risks taking the blame for poorly reviewed LLM outputs.

“
No sooner is the problem stated
than you’re expected to have a
solution within an hour.
”

Software Engineer

Project Manager

“
Stress is up, confidence
is down, attention is down.
I stopped going to all-hands
because there’s always some
exec stating we’re not
working hard enough.
”

HEALTH AND SAFETY

There is increasing evidence of the psychological damage caused by heavy LLM use which, when raised by staff, is ignored or derided by senior management despite its documentation in academic literature, and this is validated by participants’ experiences.¹²

“I love writing code and I loved my job but I know that the joy I had will be gone and will never come back if we continue down this road, not only because the quality is poorer but the language is dying.”
– Software Engineer

In addition to heightened stress brought on by work intensification (which is not a unique phenomenon), participants reported a decrease in self-confidence due to downskilling, fragmentation of their attention due to the sporadic nature of agentic workflows, and negative impacts on work-life balance as the waiting time for outputs is unpredictable. All in all, workers forced to use genAI tools are taking a hit to their mental health.



GLOSSARY OF TERMS

Agent: An LLM-based system designed to carry out tasks autonomously.

Assistant: An LLM-based system designed to support users with tasks by accepting direct inputs.

AI: Artificial Intelligence: In this report, AI is used as a broad umbrella term for technologies that automate, predict, classify, generate, recommend or support decisions. This report avoids treating all AI systems the same.

AI Psychosis: Shorthand for AI-related mental health risks. Typically used to describe the effects of prolonged exposure to AI sycophancy but may include the stress, anxiety, confidence loss, attention fragmentation, ethical conflicts or work intensification that users experience when using genAI.

Compute: Compute is the term used to describe the processing power and physical (hardware) resources that a program needs to run. AI models require a massive amount of compute to train and maintain which leads to the increasing creation of new data centres that consume large amounts of our (currently majority carbon) energy and water supplies. Construction of these data centres also requires a significant amount of minerals, such as coltan and cobalt, which are used in the creation of semiconductors. These minerals are mined mostly in the Democratic Republic of Congo at great expense to both the environment and those who mine the minerals (many of whom are children).¹³

Deskilling: Skill degradation. When a worker loses their aptitude for a job or task due to reliance on technology that prevents skills acquisition.

Downskilling: Skill degradation. When a worker loses their aptitude for a job or task due to reliance on technology that prevents skills acquisition.

Externality: A cost that affects an “uninvolved” third party as a result of involved parties’ behaviour. Relevant examples: data centre carbon burn and water depletion, worker mental health, civilian mass casualties in Gaza. These costs are “externalised” because management ignores them when making decisions on AI development or deployment.

GenAI: Generative AI: A category of LLM- and VLM-based tools that synthesise text and graphic information.

LLM: Large Language Model: A type of machine learning technology trained on large amounts of textual input data that can synthesise or manipulate natural language text. Examples include ChatGPT, Claude, and Gemini.

ML: Machine learning: A broad category of statistical algorithms trained on input data to generate probabilistic predictions as output data. Artificial neural networks are one kind of machine learning algorithm and underpin genAI technology.

Sycophancy: The proclivity of LLMs to output what the model predicts the user expects to hear, rather than what is correct or appropriate.

Tokens: A token is a unit of text processed by an LLM. Tokens are used by companies to measure cost, usage levels or activity, but do not necessarily measure quality or productivity.

VLM: Vision-language Model: A system that can process both images and texts. For example, it may describe or synthesise images, analyse screenshots, or respond to virtual prompts.


METHODOLOGY

TECH WORKERS' INQUIRY: AI						31
The UTAW AI Workers' Inquiry was carried out by members across the UTAW branch of the CWU in August and September of 2026. Workshop sessions took place both in person and online over video calls, where qualitative data was collected by semi-structured interviews between pairs of workers Since even before our official foundation as a union, we were a working group of London Tech Workers Coalition, and have been conducting workers inquiries based on the Marxist tradition. Miceli et al. (2025) note that “The global AI industry is fueled by hidden and precarized labor, yet worker voices remain largely absent from the research that studies them”.¹⁴ This Workers' Inquiry is a study that has emerged as part of this wider trend of workers inquiries in this fast-moving economy.						QUESTIONNAIRE During the workshops, a list of interview questions¹⁵ was given to each pair, which were then used to inform their breakout interviews. Participants could choose to expand on questions that most interested them. After breakouts were held, participants came back to the wider group to briefly discuss their conversations and expand on themes that were emerging, while notes were taken by facilitators. Following this data collection, the questionnaires were anonymised and collated to form this write up. WORKER PROFILES UTAW is made up of workers from all walks of the tech sector, from those in technical roles at tech and non-tech companies to non-technical roles within tech companies as well as workers engaged in educating others around technology. All members of our union were invited to take part in the inquiry regardless of their role, so our research is informed by voices from across the sector. This is also seen in the wide range of company types that participants are employed in, including self-employed workers, small and medium enterprises, and large and multinational Big Tech corporations. Participant roles include those in software development, cybersecurity, IT support, data science and data engineering, customer service, coaching and education, research, and more.
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES

ENDNOTES

1. Doctorow, C. (2026) The Reverse Centaur’s Guide to Life After AI. versobooks.com [p5] 2. C13 Artificial Intelligence at work – a just transition through worker voice for the benefit of workers and society. TUC Congress, 2026. Received from CWU, Aegis congress.tuc.org [p7]				3. This is distinct from simply using a tool incorporating AI-enabled functionality, such as HR tools or “smart” kanbans. [p10] 4. Bender, E.M and Hanna, A. (2025) The AI Con: How to Fight Big Tech’s Hype and Create the Future We Want: Bodley-Head. thecon.ai [p14]		
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES

TECH WORKERS' INQUIRY: AI					33	
5. Hanna, A. and Bender, E.M. (2025) 'The Myth of AGI', Tech Policy Press, 3 June. [Accessed 11 September 2026]. techpolicy.press [p15] 6. OECD (2024) The impact of Artificial Intelligence on productivity, distribution and growth: key mechanisms, innitial evidence and policy challenges. OECD Artificial Intelligence Papers, No. 15. Paris: OECD Publishing. See also, International Labour Organization (2026) The impact of GenAI on jobs, productivity and work organization. a review of the empirical evidence. Geneva: ILO. ilo.org [p15] 7. Jevons, W.S. (1865) The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines. London: Macmillan. libertyfund.org [p16] 8. Gans, J. (2026) WTF is Pro-Worker AI joshuagans.substack.com/p/wtf-is-pro-worker-ai [p16] 9. Shaw, S. D., & Nave, G. (2026, January 12). Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. osf.io [p19] 10.UK Parliament POST (2026) What are data centres and how sustainable are they? POSTnote 762. parliament.uk; The Guardian (2026) 'Google DeepMind in talks with UK unions amid staff concern over US and Israel's AI use', 20 May. [Accessed 16 September 2026] theguardian.com [p21] 11. Concerned AI Workers (2026) Statement in Support of Anthropic's Actions to Ensure Their AI Systems Are Not Used to Harm. [Accessed 10 September 2026] concernedaistaff.org [p21] 12.Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X-H., Beresnitzky, A.V., Braunstein, I. and Maes, P. (2025) 'Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task' arxiv.org [p26] 13.Reitner, L. and Lutz, S. (2025) What direct risks does AI pose to the climate and environment?; lse.ac.uk; The Oakland Institute (2025) The Devastating Toll of Blood Minerals on People oaklandinstitute.org [p29] 14.Miceli, M., Dinika, A.-A., Kauffman, K., Salim Wagner, C., Sachenbacher, L., Hanna, A., & Gebru, T. (2025). Methodological Considerations for Centering Workers' Epistemic Authority in AI Research Proceedings of the AAAI ACM Conference on AI, Ethics, and Society, 8(2), 1698–1710. doi.org [p31] 15.Full questionnaire used during workers' inquiry. go.utaw.tech/aiwi26 [p31]						
EXECUTIVE SUMMARY	POLICY PROPOSALS	INTRODUCTION	ISSUES THE INQUIRY SURFACED	GLOSSARY OF TERMS	METHODOLOGY	ENDNOTES



** ACKNOWLEDGMENTS **

Big up to all the union members who participated, bringing their (and their colleagues') lived experience to increase our collective knowledge as workers.

Commendations to the UTAW national committee for their steer on the need for a thorough AI investigation and their feedback on this report.

Thanks to eeddwiiinn for the sick site build and Adam Evans for the graphic design and print layout.

This report was authored and edited by Lynda Ouazar, Eleanor Payne and Lamian Pheres, who also designed and ran the workers inquiry. If you find mistakes, don't @ us.

