



THE ALGORITHMIC PANOPTICON: NAVIGATING THE LEGAL AND ETHICAL VACUUM OF AI SURVEILLANCE

Gurnoor Kaur*

Abstract

The rapid integration of artificial intelligence (AI) into digital surveillance systems represents a new era of monitoring and governance. While AI-powered tools such as facial recognition, predictive policing, and biometric identification promise enhanced efficiency for state and private actors, they simultaneously pose profound ethical and legal challenges to foundational civil liberties. This paper provides a comprehensive analysis of these challenges, focusing on critical concerns related to privacy erosion, algorithmic bias, the absence of meaningful consent, and significant accountability gaps. Through a comparative examination of legal frameworks and landmark case studies across India, the United States, the United Kingdom, and China, this paper demonstrates a global trend toward a "technological panopticon" where the pursuit of security and efficiency conflicts with the protection of human rights. The analysis reveals a complex landscape characterised by reactive, fragmented legal responses in democratic nations, in stark contrast to the proactive, state-centric models of authoritarian regimes. This paper concludes by advocating for the implementation of a human-centred governance framework grounded in principles of mandatory human oversight, algorithmic transparency, and strengthened legal safeguards to ensure that technological innovation serves to augment, rather than erode, justice and equity.

Keywords: Artificial Intelligence, Technology, Surveillance, Ethical Duties, Algorithmic Transparency

* Student- B.A.LL.B. (Hons.) 4th Year: University Institute of Laws, Panjab University Regional Centre, Ludhiana. Email Id- gurnoor.kchawla@gmail.com, Contact No.- 94658 18322

I Introduction

THERE will come a time when it isn't they're spying on me through my phone anymore. Eventually, it will be my phone is spying on me.

—*"Philip K. Dick"*¹

1.1 The Duality of AI in Governance: Efficiency vs. Ethical Duties

The advent of AI has ushered in a transformative era in state and corporate governance, providing unprecedented capabilities for data analysis, crime prevention, and service delivery. AI-powered systems, including facial recognition, predictive policing, and biometric authentication, are being implemented worldwide to create safer, more efficient societies. Law enforcement agencies, for instance, are utilising these tools to analyse extensive datasets to predict crime hotspots, optimise resource allocation, and identify individuals in real time. This technological shift is fundamentally reshaping the relationship between citizens and the state, transitioning from traditional targeted surveillance to a more widespread and automated form of monitoring.

However, this rapid integration exposes a fundamental dilemma, as technological efficiency often comes into direct conflict with ethical obligations and constitutional safeguards. The same qualities that make AI highly effective for surveillance, including its speed, large-scale reach, and ability to operate discreetly, also create serious risks to personal privacy, individual autonomy, and human dignity. As AI-driven

¹ Quotes about Surveillance, *available at*: <http://www.goodreads.com/quotes/tag/surveillance> (last accessed on September 28, 2025)

monitoring becomes more prevalent, from public spaces to personal devices, it risks making traditional ideas of privacy outdated and creates a landscape where individuals are reduced to "disembodied data."

1.2 Statement of the Research Problem: Navigating the Legal and Ethical Vacuum

The central research problem addressed in this paper is to analyse how different jurisdictions are grappling with the legal and ethical challenges posed by the integration of AI into digital surveillance. The inquiry extends beyond a simple catalogue of these challenges to a deeper investigation of the legal frameworks, policy principles, and on-the-ground consequences in varied geopolitical contexts. The analysis reveals that the rapid deployment of AI has created a legal and ethical vacuum, where existing laws prove inadequate to manage novel technological threats.

Key questions guiding this analysis focus on how legal systems, especially within democratic frameworks, respond to the gradual erosion of privacy and the resulting chilling effect created by pervasive surveillance, with growing concern particularly in India. The discussion also examines the crucial issue of accountability, highlighting how responsibility becomes diffused within public-private AI surveillance models, especially when algorithmic systems are developed or outsourced to private vendors. At a broader level, this analysis reflects an emerging global pattern in which the cross-border adoption of AI by states acts as a powerful accelerator, contributing to authoritarian tendencies and significantly transforming the relationship between

citizens and the state into one increasingly shaped by continuous monitoring.

II Conceptual Framework: Defining the Technological Panopticon

2.1 From Monitoring to Prediction: The Evolution of Surveillance

The current generation of AI surveillance transcends simple monitoring; it shifts state capability from observing past actions to predicting future behaviour. AI systems, encompassing algorithmic surveillance, predictive policing, automated decision-making, and censorship, are becoming powerful instruments for monitoring, influencing, and suppressing dissent, a phenomenon aptly termed 'AI repression' or 'algorithmic repression'. These AI applications are evolving rapidly, increasing the risk of misuse not only by authoritarian regimes but also by democratic states under pressure from security concerns.

This transformation accelerates a global trend toward a system of "digital unfreedom". This paper utilises the concept of the Technological Panopticon to describe this system. The Technological Panopticon is a state where the potential for pervasive, continuous monitoring by state actors, facilitated by linked data and complex algorithms, forces constant behavioural self-correction among citizens. This state of perpetual scrutiny, whether real or merely perceived, erodes fundamental civil liberties and undermines the necessary foundations for democratic discourse.

2.2 Core Surveillance Technologies and Associated Risks

The ethical and legal concerns addressed in this paper primarily centre on two related technological categories:

Facial Recognition Technology (FRT) and Live Facial Recognition (LFR): FRT and LFR offer high-efficiency solutions for identifying individuals in real-time, often without their knowledge or consent. While proponents highlight their security benefits, ethical governance demands strict adherence to principles of fairness, non-discrimination, and respect for human rights in their deployment.² This requires rigorous data protection measures, strict adherence to data minimisation principles, and careful consideration of data retention policies to safeguard individual privacy against unauthorised access or misuse.

Predictive Policing Algorithms: These sophisticated, location-based algorithms utilise existing crime data to forecast areas and times with a statistically higher risk of future criminal activity. Proponents argue that this assists police departments in optimising resource allocation.

III The Erosion of Foundational Rights: Critical Ethical and Legal Challenges

3.1 Privacy Erosion, Function Creep, and the Chilling Effect

The integration of AI surveillance fundamentally challenges established privacy norms. The traditional legal model of "notice and consent" is profoundly inadequate when confronted with mass biometric surveillance, such as LFR deployed in public spaces, where meaningful consent is practically impossible to obtain. While organisations are generally permitted to use personal information for purposes reasonably

² Biometrics Institute Calls for Consistent Facial Recognition Guidance, *available at*: <https://www.biometricsinstitute.org/biometrics-institute-calls-for-consistent-facial-recognition-guidance/> (last visited on October 29, 2025).

expected by the consumer, AI radically alters what is "reasonably expected".

Early judicial steps toward recognising privacy as a right were seen in India in the 1960s and 1970s. The case of *Kharak Singh v. State of Uttar Pradesh*³ did not explicitly recognise privacy as a fundamental right but focused on the broader interpretation of personal liberty, implicitly pointing to the importance of protecting individual dignity and freedom from unwarranted state surveillance. This foundation was expanded in *Maneka Gandhi v. Union of India*.⁴ (1978), which ruled that personal liberty under Article 21 of the Constitution must encompass a broader range of personal decisions that reflect an individual's autonomy and dignity. These cases paved the way for the 2017 *Puttaswamy* ruling, which declared the Right to Privacy a fundamental right.⁵

Function creep is the gradual expansion of data use beyond its original, stated purpose. AI exacerbates this phenomenon because its technological capability makes the cross-referencing of vast, disparate datasets—for instance, combining routine traffic camera data with public health records or social media history—seamless and instantaneous. This technological simplicity pushes the boundaries of use limitation principles, which endeavour to ensure personal information is only used for the purpose for which it was collected.⁶

³ AIR 1963 SC 1295.

⁴ AIR 1978 SC 597.

⁵ Shubham V., "Maneka Gandhi and the Evolution of Constitutional Privacy in India," 26(4) *J. Constitutional Law* 77 (2019).

⁶ TANDFONLINE, Function creep is a phenomenon familiar to most scholars, *available at*: https://www.tandfonline.com/doi/full/10.1080/_17579961.2021.1898299 (last visited on September 29, 2025).

The collective consequence of pervasive, opaque monitoring is the chilling effect on civil liberties. In environments where citizens believe they are under constant, algorithmic scrutiny, they limit their expression, association, and movement to avoid being flagged by potentially biased systems or targeted by state actors. This self-censorship, a hallmark of the Technological Panopticon, fundamentally undermines the necessary preconditions for democratic discourse and freedom of action.

3.2 Systemic Bias and Disparate Impact in Algorithmic Policing

AI algorithms, particularly those used in policing, are known to exacerbate existing biases within the criminal justice system. The bias originates not from malicious design but from the historical inequality embedded within the training data. Algorithms are fed "bad data," derived from biased sources such as historical 911 calls, subjective officer reports, and arrest records, which produce inherently skewed results.

The use of this flawed input data consequently entrenches discriminatory policing practices. Studies on common location-based predictive policing algorithms, such as PredPol, have demonstrated a significant disparate impact on marginalised communities and people of colour. For example, one analysis found that if implemented in Indianapolis, Black and Latino communities would experience between 150% and 400% greater patrol presence compared to white communities.⁷

⁷ MFIA Yale Law School, "Algorithmic Accountability Report" (2018), available at: <https://law.yale.edu/sites/default/files/area/center/mfia/document/infopack.pdf> (last visited on October 15, 2025).

3.3 The Absence of Meaningful Consent in Biometric Data Collection

The widespread implementation of automated recognition technologies, which measure biometrics like facial identity, gait, fingerprints, and voice in public spaces, challenges the core privacy tenet of informed consent. In a mass public environment, Live Facial Recognition (LFR) operates in real-time, often capturing, processing, and identifying individuals without their knowledge of the technology's existence, the purpose of the data collection, or the specific algorithms being utilised.

The Indian Digital Personal Data Protection Act (DPDP Act) 2023 requires consent to be "free, specific, informed, unconditional and unambiguous with a clear affirmative action"⁸. The Act further grants individuals the right to withdraw this consent at any time. However, this requirement is nearly impossible to meet in the context of mass biometric surveillance like LFR in public spaces.

This means the traditional legal construct of "consent" for biometric data capture in public settings is effectively moot. Instead of relying on a broken consent model, regulators must transition toward an enforced risk-management and transparency paradigm, demanding pre-deployment assessments and stringent use limitations to safeguard individual rights.

3.4 Accountability Gaps and the Public-Private Nexus

A major structural governance challenge is the absence of clear liability in the event of AI failure or harm. When an AI system generates

⁸ The Digital Personal Data Protection Bill, 2023 (Bill 113 of 2023), cl. 6.

misleading, discriminatory, or harmful outputs, it is often unclear whether responsibility lies with the government agency as the end-user, the vendor that provided the proprietary software, or the developer that trained the initial model. Organisations must clearly define liability to mitigate these risks.

Government-tech company exchanges, particularly those related to national security and law enforcement, are often shrouded in secrecy, creating an environment "ripe for abuse" and making public oversight difficult. There is often little transparency regarding the ethics underpinning the algorithmic decision-making provided by tech platforms.

This structural accountability gap permits government agencies to deploy AI tools while legally distancing themselves from the inherent technical flaws of proprietary, black-box algorithms. By outsourcing these functions, agencies gain operational capability without incurring full legal liability or transparency duties, further eroding public oversight of critical policing and security practices. For executive leadership, governing AI risk is no longer merely an optional technology issue; it is a critical governance and boardroom imperative demanding robust third-party risk governance, continuous monitoring, and mandatory audits to hold vendors accountable for compliance and transparency.

IV Comparative Jurisprudence: Global Divergence

Established democratic nations have struggled to establish comprehensive, proactive regulatory frameworks for AI surveillance, resulting in reliance on reactive judicial processes or localised, fragmented political interventions.

4.1 The United States: Fragmentation and Local Resistance

The United States currently lacks a unified federal regulatory framework governing AI surveillance, leading to a patchwork of inconsistent state and local responses. A growing number of US cities, including Santa Cruz and Boston, have taken proactive local measures by outright banning the use of FRT and predictive policing technology by police and public authorities. Santa Cruz holds the distinction of being the first US city to implement an outright ban on predictive policing.⁹

However, this reliance on fragmented, reactive legislation proves insufficient against determined state actors. The New Orleans Police Department (NOPD) continued utilising AI-powered facial recognition technology even after the New Orleans City Council unanimously passed an ordinance banning its use in December 2020. Although a subsequent 2022 ordinance allowed FRT for investigating violent crimes under strict guardrails (including human review and mandatory reporting to the City Council), the NOPD consistently violated these rules. The NOPD received real-time automated facial recognition alerts from a network of over 5,000 external cameras operated by Project NOLA, a non-profit partner, which scanned public streets and sent alerts directly to officers' phones.¹⁰

Law enforcement agencies utilised external partners and state/federal fusion centres to bypass local ordinances and oversight requirements.

⁹ DIGITAL FREEDOM FUND, Santa Cruz and Boston Ban Facial Recognition (2020), available at: <https://digitalfreedomfund.org/taking-police-tech-to-court/> (last visited on September 24, 2025).

¹⁰ James Mill, "New Orleans Police Continue Using Facial Recognition Despite City Restrictions", *Criminal Legal News* 1st July 2025, available at: <https://www.criminallegalnews.org/news/2025/jul/1/new-orleans-police-continue-using-facial-recognition-despite-city-restrictions/> (last visited on September 9, 2025).

This process of systemic evasion indicates that local bans are easily nullified by inter-agency collaboration and outsourcing, underscoring the urgent need for comprehensive, national-level procurement policies that mandate transparency and compliance across all partners, regardless of where the data processing or analysis occurs.

4.2 The United Kingdom: Judicial Intervention and Legal Deficiencies

The United Kingdom provides another key example of reactive governance, where the judiciary was forced to intervene due to a legislative lag. Initial judicial deference was observed in the 2019 High Court decision concerning *R (Bridges) v Chief Constable of South Wales Police*¹¹, which categorised Live Facial Recognition (LFR) as "non-intrusive." The court initially ruled that only intrusion upon an individual's residence or bodily integrity required specific statutory authorisation.

This initial ruling was, however, overturned in 2020 by the Court of Appeal in a landmark decision. The appeal court found that the police's use of Automated Facial Recognition (AFR) technology breached privacy rights, data protection laws, and equality laws. Although the court found no evidence that the specific AFR system used was biased, it stressed that the police force failed to take steps to satisfy itself as to whether bias would manifest in practice, thus violating equality duties.

The Bridges case reveals that even in robust democracies with established data protection frameworks, AI technology advanced faster

¹¹ *R. (Bridges) v. Chief Constable of South Wales Police*, EWCA Civ 1058 (2020).

than the legal system could adapt. The court's role became fundamentally reactive, identifying failures within a legal structure that was not designed to handle real-time, mass biometric capture. While the judgment restored rights, it did not create the necessary comprehensive statutory authorisation or regulatory architecture that the technology demands, leaving the UK in a continued state of legislative lag.

4.3 China: Algorithmic Repression and Total Social Control

The People's Republic of China represents the global apex of state-centric AI control, utilising integrated technology to suppress opposition and enforce total social conformity—a system of algorithmic authoritarianism.

The Chinese state employs AI surveillance and predictive policing extensively in Xinjiang, targeting the Uyghur Muslim minority under the pretext of counter-terrorism. Initiatives such as the Sharp Eyes programme for urban monitoring and the Integrated Joint Operations Platform (IJOP) are used to collect vast amounts of data to monitor behaviours deemed indicative of potential social instability. This extensive network allows authorities to utilise predictive models to detect potential dissidents before any concrete act is committed, leading to severe limitations on liberty, widespread self-censorship, and detentions in re-education camps based on opaque AI analyses.¹²

A parallel structure enforcing conformity nationwide is the Social Credit System (SCS). The SCS is not a singular score but a range of

¹² European Parliament, "Artificial intelligence (AI) and Human Rights: Using AI as a weapon of repression and its impact on human rights" (Policy Department for External Relations, June, 2024).

interconnected municipal policies and technologies that use sophisticated algorithms and machine learning to evaluate and rank citizens and corporations. The integration of these interlinking data points across all spheres of daily life demonstrates the ultimate risk of unchecked AI integration: the creation of a total data web used as an instrument for algorithmic behaviour modification. This model is explicitly designed to preserve the political status quo and suppress dissent, and there are concerns regarding the adoption of these surveillance methods by other illiberal and authoritarian regimes globally.¹³

V The Indian Context: From Puttaswamy to the DPDPA 2023

5.1 India: Statutory Erosion of Constitutional Safeguards

India's regulatory approach demonstrates a significant movement toward proactive state-centric permissiveness under the guise of national security. The country's much-anticipated data protection regime, the Digital Personal Data Protection Act (DPDPA) of 2023¹⁴, has been widely criticised for favouring state interests over individual rights.

The Act was developed following the Supreme Court's Puttaswamy judgment (2017), which affirmed privacy as a fundamental right and highlighted the need for a legal framework to protect personal data. The DPDPA aims to align India's laws with international standards by establishing rights such as access, correction, erasure (the "right to be

¹³ Steven Feldstein, "The Road to Digital Unfreedom: How Artificial Intelligence is Reshaping Repression" 30(1) *Journal of Democracy* 40-52 (2019).

¹⁴ The Digital Personal Data Protection Act, 2023 (Act 22 of 2023).

forgotten"), and data portability, while also requiring strict data security for Data Fiduciaries.

However, the most detrimental feature of the DPDPA is the "National Interest" Loophole found in Section 17. This clause grants the government sweeping power to exempt any government agency from the operation of the entire Act for purposes as broad as "National Interest"¹⁵. This statutory permissiveness directly undermines the Supreme Court's landmark direction established in the Puttaswamy (Aadhaar) case (2019). The Puttaswamy judgment mandates that any state surveillance must strictly meet the tripartite tests of legality, necessity, and proportionality.

The DPDPA, by granting excessive discretion to state actors—especially law enforcement agencies—under broad national security language, essentially creates a legal mechanism to nullify constitutional checks on state power. Furthermore, the DPDPA is "devoid of substantive safeguards on surveillance" and fails to establish an independent and appropriately autonomous data protection authority. The permissive framework enabling the deployment of AI-powered surveillance systems across India without rights-based protections demonstrates a "sad drifting away from the constitutional framework," legalisation of unchecked power, and an institutionalised path toward digital authoritarianism.

¹⁵ Tripti Bhushan, "AI, Surveillance and Privacy in India: Human Rights in the age of Technology", *Oxford Human Rights Hub* 3rd September 2025, available at: <https://ohrh.law.ox.ac.uk/ai-surveillance-and-privacy-in-india-human-rights-in-the-age-of-technology/> (last visited on October 7, 2025).

5.2 The Pegasus Litigation: Judicial Review in the Age of Cyber Surveillance

The Pegasus controversy marked India's first major constitutional confrontation with military-grade cyber-surveillance technology. Following international investigative reports alleging the use of Pegasus spyware against journalists, activists, and political actors, multiple petitions were filed before the Supreme Court alleging unconstitutional surveillance.

In *Manohar Lal Sharma v Union of India*¹⁶ (Pegasus case, 2021-2022 proceedings), The Court adopted three constitutionally significant positions:

(i) National Security Cannot Completely Exclude Judicial Review

The Court rejected the government's argument that national security automatically prevents disclosure or judicial scrutiny. This reaffirmed constitutional supremacy over executive secrecy claims.

(ii) Privacy Applies to All Citizens, Not Just Targeted Individuals

The Court emphasised that in a constitutional democracy, citizens possess a reasonable expectation of privacy, and indiscriminate surveillance cannot be justified without statutory safeguards.

(iii) Independent Technical Oversight as Constitutional Remedy

¹⁶ (2022) 1 SCC 381.

The Court appointed an independent technical committee rather than allowing a government-led inquiry, reinforcing the principle that surveillance allegations require independent investigation.

5.3 Committee Findings and Constitutional Implications

The Supreme Court-appointed technical committee examined 29 devices but found no conclusive evidence of Pegasus infection; however, malware was detected in some devices. Nonetheless, the Court observed a constitutionally significant concern: the Central Government did not fully cooperate with the committee's investigation. From a constitutional law perspective, this resulted in three doctrinal developments:

(a) Recognition of the Structural Risk of Covert Digital Surveillance

Even without conclusive spyware proof, the Court acknowledged the seriousness of allegations of unauthorised digital surveillance.

(b) Reinforcement of Procedural Accountability in Surveillance

The Court emphasised that surveillance must comply with “procedure established by law,” linking digital surveillance to Article 21 due process jurisprudence.

(c) Expansion of Privacy Jurisprudence into Cyber Infrastructure

The case extended privacy jurisprudence from data protection to cyber-intrusion and spyware governance.

VI Prescriptive Solutions: A Human-Centred Framework

The fragmentation, reactivity, and state-centric failures identified across the countries necessitate a global governance framework grounded in clear, mandatory, and auditable principles that prioritise human dignity, autonomy, and justice.

6.1 Mandatory Human Oversight and Control

The foundational principle of any responsible AI deployment must be the safeguarding of human autonomy, protecting the independence and dignity of all affected individuals—from law enforcement personnel to citizens, suspects, and victims.

Human control and oversight must be mandatory, defined as the ability and opportunity for humans to properly supervise, engage with, and intervene in an AI system during both its development and use. This is essential for high-risk AI systems that influence critical societal sectors such as healthcare, infrastructure, or legal outcomes. Law enforcement agencies must verify that any AI systems they acquire are specifically designed with functionalities that guarantee humans remain in control and that the necessary organisational structures and training protocols are established to support this oversight.

The operational standard must mandate the ability to override, correct, or halt automated processes based on ethical or legal risk assessments, ensuring human judgment remains the final deciding authority rather than serving as a rubber stamp for algorithmic output.

6.2 Algorithmic Transparency and Explainability

Transparency is paramount to managing the risks of algorithmic bias and ensuring accountability. This requires three interlocking mechanisms:

First, there must be a mandate for Algorithmic Impact Assessments (AIAs). Federal and governmental departments must be required to conduct mandatory AIAs for all AI deployments, particularly those designated as high-risk, to proactively identify potential risks and biases associated with automated decision systems. Furthermore, organisations could be required to conduct Data Protection Impact Assessments (DPIAs) before undertaking new data processing activities that might impact individuals' privacy.

Second, Documentation and Code Review must be mandatory. Transparency demands comprehensive documentation of the design and deployment processes of AI systems, establishing a basis for accountability and clarity in decision-making. Technical safeguards should be implemented, including third-party code review, thorough testing of outputs against known inputs, and validation of key variables to ensure the production of meaningful, unbiased information.

Third, citizens affected by automated decision systems must be afforded the Right to Explanation, allowing them to file complaints and receive clear, meaningful explanations regarding how an automated decision was reached.

6.3 Strengthening Legal Safeguards and Accountability

Generalised data protection frameworks often prove insufficient when the state grants itself sweeping exemptions, as seen with the DPDP in India. Therefore, targeted legislation is urgently needed, specifically for surveillance practices and policing algorithms to promote transparency and oversight. The implementation of the Digital Personal Data Protection Act (DPDP Act) 2023 in India establishes legal rights that must be fully enforced to align with constitutional principles. These rights include the right to access, rectification, and the Right to Erasure. The Act also regulates Cross-Border Data Transfers, requiring that adequate protection is in place for data transferred outside India, a measure designed to safeguard personal information in global transactions. Furthermore, the legal framework must clearly define liability by establishing unambiguous lines of responsibility between developers, vendors, and end-users regarding harmful or misleading outputs. This structural clarity is vital to address the accountability gaps inherent in public-private partnerships. The framework also mandates the establishment of data protection or AI governance authorities endowed with genuine autonomy, sufficient resources, and robust enforcement powers—a recognised deficiency in many jurisdictions. Finally, strict adherence to the Principle of Data Minimisation and Use Limitation must be mandated across all agencies, ensuring data collection is strictly limited to necessity and retention is minimised, thereby actively preventing function creep and upholding privacy safeguards.

VII Conclusion: Augmenting Justice in the Age of AI

The integration of artificial intelligence into digital surveillance infrastructures has precipitated a global governance crisis, marked by an enduring tension between expanding technological capabilities and binding human rights obligations. A clear regulatory divergence is observable across jurisdictions. Democratic systems such as the United States and the United Kingdom have largely relied on reactive, legally fragmented regulatory models, often resulting in post-facto judicial intervention and regulatory gaps that can be bypassed through jurisdictional or technical workarounds. In contrast, states such as China and India have increasingly adopted proactive, codified surveillance architectures, integrating AI tools into formal governance frameworks that significantly expand state monitoring capacity and centralised authority. The core conflicts identified are structural in nature. Algorithmic systems risk perpetuating systemic bias embedded within training datasets; meaningful individual consent becomes functionally unattainable in environments of mass biometric data capture; and accountability mechanisms are weakened by opaque public-private surveillance partnerships that diffuse responsibility across institutional boundaries. The proliferation of AI repression technology, facilitated in part by the global export of surveillance tools, creates an urgent and shared global responsibility to implement governance measures that prioritise human rights. Only a human-centred governance framework, grounded in mandatory structural transparency, enforceable human control mechanisms that permit interference, and comprehensive, targeted legal safeguards, can ensure that the integration of AI serves to augment justice and equity. Failure to implement such measures risks accelerating the erosion of civil liberties toward a fully realised, globally connected Technological Panopticon.