

Between Security and Fundamental Rights: A Critical Analysis of the Fragmentation, Exceptions, and Protective Gaps in European Union Biometric Surveillance Regulation

Victor-Andrei Cărcăle¹, Liana-Teodora Pascariu²

¹“Ștefan cel Mare” University, Suceava, Romania, victorcarcale@litere.usv.ro

²“Ștefan cel Mare” University, Suceava, Romania, liana.pascariu@fdsa.usv.ro

Abstract: In the EU, which currently does not have a general regulation on biometric surveillance technologies, one security-related regulation allows mass-identification by law enforcement and border authorities, the other concerns itself with basic rights and risks of always being identified without our knowledge or consent. This article reviews and critiques the literature on EU regulation of biometric surveillance to argue that the current normative architecture - including, but not limited to, the GDPR, Directive 2016/680, related CJEU and ECtHR case law, as well as the recently adopted AI Act - remains patchy, exception-riddled, and unevenly protective. However, legal and academic attention has disproportionately focused on high-profile applications of these technologies, such as public-space facial recognition and police remote biometric identification, while under-theorizing equally significant practices, including database retention for the purpose of advancing biopolitical governance (e.g., ETIAS, the Common Identity Repository), surveillance of protest and assembly, bio-surveillance in audience studies, and so forth. The paper charts important areas of academic contention, situates the proportionality test as both a foundational pillar and a battleground, and identifies research gaps in emotion recognition, biometric categorization, and dynamic databases. In conclusion, the AI Act is indeed a substantial piece of legislation, but the work still left to be done in its architecture has not addressed the fundamental conflict between security and rights protection; instead, it has simply shifted that conflict into an opaque maze of exceptions and derogations.

Keywords: Biometric Surveillance, Facial Recognition, AI Act, Proportionality, Interoperability

Introduction

Biometric surveillance represents a unique and more contentious subcategory of EU technology law. Other forms of data processing are mediated (and abstracted) through specific representations of what that data is, but biometric identification operates on the physical, very raw as compared to other types: face, gait, voice, and iris all converted into search-, store-, and matchable digital objects. While this feature enables biometric processing, it is the specific capacity for mass identification, often without either the knowledge or true consent of the data subject, that has catapulted biometric surveillance technologies - namely facial recognition systems - into almost a defining regulatory rivalry within modern EU law.

Over the past decade, valid legal instruments to govern these technologies have emerged at breakneck speed from a data protection regime centred largely on the EU General Data Protection Regulation (GDPR) and Law Enforcement Directive, through an extensive body of case law rooted in constitutional and human rights jurisprudence, culminating recently with adoption of the AI Act that introduced for the first time ever an EU-wide legislated risk-based framework specifically aimed at biometric identification systems. Despite this proliferation of legal instruments, the academic literature presents a stark and counterintuitive finding: the overarching challenge analysed is how to reconcile security objectives and basic rights in light of biometric surveillance technologies - yet they come to similar conclusions across studies that EU law currently lacks coherence, rife with exceptions and delivering uneven safeguards across use cases (Barkane, 2022:159-160; Kagkli & Angelaki, 2026, pp. 52-53; De Hert and Bouchagiar, 2021, pp. 210-212).

This article takes that conclusion as a premise and expands it into a formal critique. It is based on the premise that—despite the conceptually broad acceptance of legal standards governing biometric surveillance, namely, the necessity and proportionality test, which have long been operationalized in both European Court of Human Rights case law and its EU equivalent (the great majority but not all of national courts have adopted same solution yet; also called 'legitimate aim')—many papers assert that Law Enforcement access, including facial recognition that may use remote biometric identification or function as interoperable databases, could move ahead of clarity on safeguards, remedies, and oversight (Mobilio, 2023, pp. 23-24). This creates a regulatory pathway where the governing principles are already set and widely accepted, but where operational norms remain relatively low, those norms that should develop those principles into constraints against actual technologies.

Throughout, the article offers a critical synthesis rather than a basic description. It does more than summarize the literature: it examines why the evidence is distributed as it is, what this distribution leaves out, and what these omissions mean for the future of EU biometric surveillance law.

Mapping the Evidence Base: Coverage Across Regulatory Dimensions

The first question someone conducting a scholarly or judicial critique of any area of legal scholarship must ask is not simply what that scholarship holds, but where it has focused and where it has remained blind. This mapping of scholarly coverage in itself is a finding: it identifies the technologies and regulatory mechanisms that have animated legal academia, policymakers, and civil society, as well as those that remain comparatively obscure yet potentially no less influential. This distribution across five regulatory dimensions—prohibitions, oversight mechanisms, data protection safeguards, law enforcement use, and cross-border systems - and across the five technological and operational domains of biometric surveillance is summarized in Table 1.

Table 1. Evidence coverage across biometric surveillance regulation topics, scored by relative density of scholarly and regulatory attention

	Prohibitions	Oversight	Data Protection	Law Enforcement Use	Cross-Border Systems
Remote biometric identification	10	8	7	10	3
Facial recognition in public spaces	8	6	7	9	2
Biometric databases and retention	3	4	9	6	5
Protest and assembly surveillance	2	2	3	4	GAP
Border and migration systems	4	4	5	5	7

The most visually obvious aspect of this distribution is the asymmetry between the succession of rows for visible, public-facing technologies and those at the base representing structurally embedded but largely invisible infrastructures. Remote biometric identification and facial recognition in public settings score highly on nearly all of those dimensions, except for cross-

border systems, where their coverage plummets. Biometric databases and retention practices also exhibit an inverse pattern: thin coverage of prohibitions and oversight compared to other domains, yet the highest score of all for data protection, since database retention has primarily been analyzed through a data protection lens - not through the lens of constitutional prohibition or institutional oversight. The weakest coverage is evident across every dimension of protest and assembly surveillance, including a clear gap in cross-border systems, even though this case covers one of the use cases with the most immediate consequences for how political freedoms are exercised.

This is not a coincidence, and it does not merely reflect the order in which technologies were ready or the ones that are more broadly deployed. Instead, it mirrors the course of EU legislative and judicial focus since the early 2010s. The most covered is law enforcement facial recognition, particularly debates over the necessity and exceptions to police use in the AI Act. The explanations for this concentration are not hard to trace: facial recognition by police in public space was the technology that elicited the most front-and-center civil society mobilization, national news coverage, and ultimately direct legislative interest during the debate around the AI Act, when remote biometric identification became a single-issue point of high political contention in an otherwise multifaceted process.

In contrast, coverage on protest surveillance and the cross-border implications of interoperable systems is thinnest of all, yet the papers flagged serious risks to privacy, assembly, and procedural rights (Gabrielli, 2025, p. 541). Those numbers shouldn't be a sideshow but taken seriously. The comparative scarcity of legal scholarship focusing on protest surveillance is especially striking since the chilling effect associated with biometric identification of protesters is, almost by definition, difficult to observe and measure: an individual avoiding a demonstration for fear of being identified and profiled does not give rise to a court case or data protection authority complaint nor appears as a headline in any paper. The damage is structural and behavioral rather than discrete and reportable, which likely helps explain why it has received less sustained scholarly scrutiny than the much more visible and well-being-affecting (but litigably associated) harms connected with misidentification by facial recognition.

The reason for this asymmetry, as highlighted in the literature, is more institutional than substantive. That gap is probably still the case because EU scholarship has focused empirically on headline controversies over facets of the AI Act and police facial recognition, while what are now emerging as major governance architectures through such examples as ETIAS (European Travel Information and Authorization System), CIR (Common Identity Repository), and interoperable databases have only just begun to be subjected to a more substantial legal analysis (Thommandru et al., 2025, pp. 1, 8-9). The ETIAS and the CIR are both rather new features of the EU institutional landscape, so that only recently have their full operational impacts - especially those surrounding interlinking previously siloed national and EU-level databases into a single searchable infrastructure with interoperable standards - begun to be addressed in academic commentary (indeed, the volume here has not yet caught up being techniques such as facial recognition which were being widely debated at least from around 2018).

In the case of the field, this leads to normative claims being strongest where public-space surveillance is most visible and direct, but the evidence base for systems that are less visible yet fundamentally important for data retention, sharing, and cross-system matching (Mollo, 2023, pp. 1104-05) is much thinner. This observation bears a methodological implication that echoes across the rest of this article: legal scholarship on biometric surveillance needs to be read, not just for what it says about the biopolitical technologies it examines in detail, but also for what its silences can tell us about which forms of state power have managed, so far, to escape sustained critical attention.

The Best-Covered Domain: Facial Recognition, Remote Biometric Identification, and the AI Act

Necessity and Proportionality at the Core

The necessity and proportionality test is the doctrinal core of EU biometric surveillance regulation, as it concerns the use of facial recognition or remote biometric identification. This test - largely developed in the European Court of Human Rights' case-law under Article 8 ECHR, and later absorbed also into EU law through the Charter of Fundamental Rights, and reinforced by case-law from the Court of Justice - states that relating to any interference with the right to privacy/data protection must meet a necessity/proportionality standard but which requires consideration whether there is an appropriate legal basis within national or Union law for affecting this right. The core legal test throughout the corpus is necessity and proportionality. Many papers argue that any actual safeguards, remedies, and oversight stay considerably behind facial recognition, distant biometric identification, interoperable databases, and police access (Akhtar, 2020, pp. 325, 340-42).

The enduring significance of this test as the leading lens through which to analyze the relevant constitutional issue rests on its doctrinal status rooted in two themes: part intrinsic and part extrinsic, so to speak. The point is that since proportionality analysis is a balancing exercise, not a categorical one, it can, in principle, cover an enormous range of factual situations, ranging from a single identification check at the border to sweeping, real-time scanning of all faces within reach of microphones at a political rally. While this flexibility has become the most powerful asset of the test, it also serves as its greatest liability for the protection of rights: By allowing courts and legislators to adjust the degree of intrusion that is appropriate to the factual context, but at a price of creating considerable uncertainty – not only for individuals and civil society organisations but even law enforcement authorities themselves about where exactly what constitutes lawful biometric processing ends in each concrete case.

The AI Act with the Remote Biometric Identification Component

The Artificial Intelligence Act is the most comprehensive piece of legislation in this area to date, establishing an EU-wide framework for remote biometric identification systems that categorizes usages within a risk-based hierarchy and (in very specific circumstances) outright bans certain deployments. The negotiations over these provisions were among the most heated in the AI Act legislative process, prompting considerable scholarship on whether the final compromise adequately protects fundamental rights or legitimizes ongoing police use of biometric identification under carefully crafted exceptions.

Various authors in the literature argue for this compromise and maintain that a legal framework that is effectively prohibitory in principle but provides narrow, judicially facilitated exceptions is preferable to either an outright ban or a free-for-all. Other authors advocate using proportional regulation rather than an outright ban, assuming that the conditions of strict necessity, court or independent review, and limited purposes are met. An absolute ban is thus both politically unworkable, given the genuine counter-terrorism purposes that biometric identification can serve in the sort of rare circumstances where national security necessitates it (e.g., finding victims of serious crime or preventing an imminent terrorist attack), and self-defeating because it might simply push biometric use into more opaque and lesser regulated channels.

Other scholars adopt a considerably more skeptical stance toward the final design of the AI Act, arguing that while its prohibition provisions may seem rigorous in the abstract, their effectiveness is undermined by the breadth and vagueness of numerous exceptions. Some critics contend the AI Act is still too soft in allowing exceptions or loopholes, especially for systems involving remote biometric identification, emotion recognition, and categorization.

This is a central line of argument animating the critique: namely, that the legal test contained in the exceptions – usually revolving around notions like 'serious crime', 'specific threat' or operational 'necessity' – gives law enforcement authorities (and in some Member State implementations, private contractors acting on their behalf) significant interpretive leeway directly at odds with what has become a hallmark of measures interfering with fundamental rights according to the Court of Justice's exacting standard for foreseeability.

Police Facial Recognition: The Challenge of Operational Oversight

Aside from whether the AI Act's substantive prohibitions are drawn fairly clearly, another important part of the literature questions how effective the oversight mechanisms accompanying authorized uses of facial recognition are in practice at meaningfully limiting police discretion. Indeed, the controversy over private facial recognition databases developed by scraping virtually any image posted in public places to make them available to law enforcement across the country has become something of a case study in this literature. This controversy is situated within the broader debate over whether existing oversight mechanisms are sufficient when private commercial actors, rather than state agencies, aggregate and operate the biometric infrastructure that police subsequently rely on in the Clearview case, from a European perspective.

This trajectory of assessment feeds directly into the larger overarching finding across the literature concerning the disconnect between formal legal protections and on-the-ground practice. The GDPR and Law Enforcement Directive have created an entire apparatus of data protection predicated on a fairly unambiguous allocation of responsibilities between controllers and processors, and a fairly transparent chain of custody for personal data. Private mass scraping of facial recognition databases - subsequently licensed to public authorities - further complicates this allocation by injecting difficult questions about whom, if anyone, is legally responsible for protecting the flammable yet essential right to receive materially accurate data in the first place, and thus whether downstream police use of an unlawfully compiled database can ever be rendered legal through ex post procedural safeguards.

Discrimination, Bias, and the Social Costs of Facial Recognition

A second axis of scholarship on facial recognition and remote biometric identification concerns not the formal legality of many deployments of such technologies, but rather how these technologies differentially impact various sub-populations. Many papers also highlight the technical and social risks, particularly discrimination, stigmatization, false matches, and bias against minorities (Tical, 2025, pp. 251, 254-255)

This body of work plays an important corrective role in the wider literature. Without it, proportionality and necessity analysis risks becoming a purely procedural exercise, detached from the substantive adequacy and differential impact of the technology in use. From this perspective, a proportionality assessment that ignores well-documented differences in algorithmic accuracy across demographic groups cannot convincingly claim to balance the relevant interests, since the privacy intrusion caused by a false match falls disproportionately on already disadvantaged groups.

The Under-Examined Domains: Protest Surveillance and Cross-Border Interoperability

Surveillance of Peaceful Protest Using Biometrics

While facial recognition by police in ordinary criminal investigations has received the most consistent scholarly attention, facial recognition and similar technologies for monitoring peaceful protest and political assembly remain comparatively understudied despite the weighty rights involved. This gap is situated in Gabrielli's (2025, p. 240) situational analysis of facial recognition technologies, in the context of peaceful protest, indicating the wider issue that biometric identification of demonstrators identifiably comes with a unique risk of facilitating mass

surveillance practices, which can have immediate and serious ramifications for fundamental human rights such as freedom of assembly.

The challenge in analysing protest surveillance law is that this practice sits at the intersection of several legal regimes that interact only loosely with one another. These include data protection law, which governs biometric data as a special category under the GDPR. They also include freedom of assembly and association, as protected under Article 11 ECHR and its parallel provision in the Charter. And they include public order policing regimes, which in a significant number of EU Member States grant police substantial discretion in managing public gatherings. The literature is consistent in holding that biometric surveillance poses an acute risk to fundamental rights, owing to its capacity to identify individuals at scale and its potential to reinforce state overreach. This risk's chilling effect is most pronounced and least fully accounted for in relation to freedom of assembly (Mobilio, 2023). States have been shown to exploit this regulatory gap, creating a deterrent effect on future participation in protests, since many prospective participants come to understand identification as the very point of the exercise.

The comparatively thin academic coverage of this application is not without risk. If anything, the structural features that the literature flags as most troubling in other contexts—the difficulty of challenging an identification decision after the fact, and retention of biometric templates beyond their immediate operational need—rather than engaging in suspected criminal activity.

ETIAS, the Common Identity Repository and Interoperability architecture

The second major gap identified by this synthesis concerns the long-term consequences of governing an increasingly large and elaborate web of interoperable databases established by EU institutions, especially regarding implications for rights and justice (e.g., ETIAS & CIR). Such systems present a qualitatively different type of biometric governance from the facial recognition deployments that have predominated earlier academic discussions: instead of real-time identification coupled with the monitoring (the live comparison of an image with a previously-collected one at the point of surveillance) and/or tracking an individual detained under rules governing public space access, we are faced with systematic linkages between distinct national and EU-level databases governing immigration, asylum and security data into what would effectively be one queryable infrastructure able to cross-reference otherwise compartmentalized biometric and biographic information typically held in authentication silos that were never designed to interoperate.

Thommandru and colleagues' analysis of algorithmic profiling and facial recognition in EU border control scrutinizes the ETIAS decision-making architecture and highlights some of the most important unresolved tensions between the operational logic of this system against core data protection principles (Thommandru et al., 2025).

Interoperability architectures put pressure on purpose limitation and storage limitation, two of the foundational principles of the GDPR, requiring that personal data be collected for specified, explicit, and legitimate purposes only and retained for no longer than necessary for those purposes. Another key underexplored area concerns the longer-term governance of interoperable EU systems, despite warnings. This structural problem is simple to articulate and find it very difficult to engineer around using conventional data protection mechanisms: in an environment where biometric data collected for one purpose (say, applying for a visa) becomes queryable through an interoperable system (say, criminal investigations), the principle of limitation to specific purposes comes under intense pressure because the practical effect of interoperability is that the contextual boundaries that purpose limitation is designed to protect tend to disappear.

The remedies dimension of this issue/or is, if anything, even more acute. That their biometrical profile has been processed through a cross-system interoperable query—a query

that elicits zero in-person exchange with any individual other than the identity holder, even though there could easily be hundreds of thousands of people queried about them at once—as opposed to the visible interaction (at least in theory) accompanying facial recognition cameras placed in public squares, ex ante notice requirements for which have dominated litigation and policy discussions. Without a trigger event that might lead to the need for interoperable database queries, data protection law's conventional remedy architecture—right of access, right to complain, right to an effective remedy before a tribunal—is ironically much more difficult to implement in practice; an individual cannot exercise these rights in relation to processing she has no reason to even slightly suspect took place.

Closing the Gap: Visible Salience and the Politics of Legal Attention

There is an important reason worth spending more time on: why those two areas of protest surveillance and cross-border interoperability have received less sustained academic attention than other forms of security studies, and how the field can best direct its future research efforts. This structure is a common element across both domains of harm: the harm they cause is diffuse, anticipatory, and behavioral rather than acute, individualized, and immediately litigable. A misidentification by facial recognition resulting in the wrongful arrest would also present a clear, actionable, and legally actionable harm that we can sue over, write about, and analyze as a matter of case law. An ordinary citizen who determines not to go to a protest or whose biometric data resides in an interoperable database that will perhaps be queried at some later date for something wholly separate from what its original purpose was suffers animus, both real and immutably hermetically sealed from traditional documentation of harm within the law and academia.

This asymmetry in the visibility and litigability of harm has an iterative quality: legal scholarship, like litigation itself, is responsive to that which is publicly visible and procedurally tractable, thus directing legislative attention to some technologies over others to determine which whole realms of technology incorporate the most detailed statutory treatment, generating further case law and further scholarship on those same technologies. Counteracting this cycle will require a measured methodological commitment by legal scholarship to study structurally meaningful yet procedurally blind forms of biometric governance, without the dramatic scandal that propelled academic and legislative urgency around police facial recognition.

Regulatory Tensions and Points of Scholarly Disagreement

The Existing Legal Structure and the Perceived Insufficiency of Laws

To address the challenges created by biometric surveillance, an EU response has emerged that draws on a range of separate legal instruments: first, the GDPR itself provides the general data protection contours and classifies biometric data processed for unique identification as special categories of genetic or biometric information requiring additional safeguards; next, while similarly situated, the Law Enforcement Directive creates a parallel but separate regime governing processing by competent authorities for purposes preventing detecting investigating or prosecuting criminal offences; then there is extensive constitutional case law from the CJEU and ECtHR that also weighs on this issue; and finally most recently comes a risk based regulatory model captured in the AI Act. The combination of the GDPR, the Law Enforcement Directive, court doctrine in practice, and, most recently, an AI Act with a similar risk-based model has been viewed by many authors as too abstract or incomplete to enable concrete deployments (Hert and Bouchagiar 2022, pp. 210-212).

The accusation that this framework is too abstract warrants partial unpacking, as it is not entirely abstract at any level. At the level of legislative drafting, both the GDPR and the Law Enforcement Directive made a virtue of drafting technology-neutral provisions intended to encapsulate an enormous range of data processing activities, from standard commercial

profiling through to the most sensitive forms of biometric identification. This technology neutrality—an important strength for ensuring legislative robustness—means that neither instrument gives detailed, technology-specific operational guidance sufficiently clear that a police force deploying a facial recognition system, or even a border authority managing activation of interoperable database query capabilities, could confidently assess whether a specific proposed deployment meets the relevant legal test.

A second and related criticism regards the degree of real discretion which this abstraction necessarily devolves upon Member States and upon law enforcement bodies themselves. A major point of criticism is that the law grants broad discretion to Member States and law-enforcement authorities, resulting in fragmented national rules and creating uncertainty about the lawful acquisition, retention, and onward use of biometric data. This fragmentation leads to practical implications far beyond simple doctrinal cleanliness: a biometric surveillance practice lawful in one Member State under extensive independent oversight might be unlawful or face an entirely different set of procedural safeguards at the level of another neighbouring Member State.

Prohibition Versus Proportional Regulation

When it comes to a simple consensus about what constitutes the basic regulatory strategy EU law should pursue with regard to biometric surveillance, one and only one point of disagreement runs across much of the literature: whether or not EU law should establish a categorical ban upon the most intrusive forms of biometric surveillance or whether instead we need something more like a proportionality-based framework that allows for (some) such applications, provided they are subject to sufficiently-precise safeguards. This is more than just a theoretical dispute about regulatory technique; it represents a fundamental disagreement over whether procedural protections will suffice to address the structural risks created by mass biometric identification, or whether some uses of technology are so inherently rights-violating that no amount of necessary measures can make them acceptable.

As we will see later in this review, one line of the literature defends the proportionality-based model as a normatively justified and manageable position. If strict necessity, judicial or independent oversight, and narrow purposes are fulfilled, some authors argue for proportional regulation rather than an outright ban. The best-known variant of this argument contends that a blanket ban would be under-inclusive because it would not cover related but technically distinct identification technologies that lie outside its very specific statutory definition, and over-inclusive because it would rule out legitimate and actually proportionate biometric identification applications in the exceptional cases of emergencies such as finding missing children or identifying victims of trafficking, where the security gain arguably outweighs the privacy harm.

The competing strand of the literature, not per se denying that proportionality analysis is internally coherent in theory, argues that the practical conditions for legal exceptions based on proportionality have generated a patchwork legislative setting under the AI Act, full of qualifications, carve-outs, and vague phrases that work de facto as a permissive regime rather than a restrictive one. It has been claimed that the Act retains numerous exceptions and loopholes, in particular for remote biometric identification (*RBI*), emotion recognition, and categorization systems.

The Structural Problem of Database Persistence and Interoperability

A third major strand of the literature moves away from a focus on the act of biometric identification as such and instead examines what happens to this data once it has been generated: issues of its storage, its potential reuse for purposes other than those originally intended, and its use across institutional and national borders. Many of those papers address database retention and

interoperability as the fundamental issue, because storage, repurposing, and sharing can make surveillance eternal beyond a specific identification event.

This reframing is analytically significant because it shifts the focus from the relatively well-litigated question of whether this singular instance of biometric identification was lawful to the considerably less litigated and much more opaque question of what happens to biometric data once the identification event has come to a close. Retention of a biometric template that has been properly collected in the context of an appropriate, lawful and judicially sanctioned identification form will not render it safe from subsequent intrusions into privacy if by being retained indefinitely, linked into an interoperable database and then used for a completely different (and potentially more weakly justified) purpose; what is required at that latter stage is a fresh assessment of justification outside of which little protection can be derived from the proportionality judgement carried out when biometrics were originally collected.

This could have significant consequences for the essence of the proportionality test, as it directly affects its interpretation and application. On this understanding, a genuinely rigorous proportionality analysis cannot be a one-off test carried out merely at the original point of data collection; it must instead be seen as a continuing obligation that requires fresh assessment whenever and wherever data is repurposed, shared, or retained beyond its original prediction. The fact that this sort of recurring proportionality assessment is poorly supported under the current legal framework—organized less around the whole data lifecycle and more around isolated instances of collection and processing—helps explain why a number of authors argue that database retention and interoperability are the root problems, which are ultimately much deeper than the more visible controversies over specific facial recognition deployments.

The Competitiveness Tradeoff

The final, and somewhat different, tension throughout the literature is between fundamental rights protections, on the one hand, and broader EU policy aims regarding technology development and global competitiveness in artificial intelligence, on the other. A number of authors also refer to what they consider a policy tradeoff with innovation and competitiveness, but still locate necessity, proportionality, and rights protection at the heart of EU enforcement.

It expresses a broader and well-documented aspect of the overall legislative philosophy underlying the AI Act: to make the European Union a leader in trustworthy artificial intelligence without burdening European AI developers with regulatory requirements that could jeopardize their competitiveness relative to jurisdictions with less strict regulations. This tension appears particularly acute in the context of technologies for large-scale biometric surveillance where arguments against strong regulations as regards remote methods of biometric identification warn that such prohibitions may harm European law enforcement tech providers competing with firms developing indistinguishable systems designed to be exported to states outside the EU legal framework, and may also simply shift significant resources away from research and development on biometric surveillance capabilities by actors within (and covering) the regulatory perimeter into projects deployed either commercially or under color of state authority in political jurisdictions permitting far brinkier implementations -- all without enhancing the protection against future exposure EU citizens face when those same technologies are deployed in non-EU commercial or national contexts.

However, interestingly enough, even the most sensitized authors with regard to this competitiveness dimension do not tend to argue that it should override necessity and proportionality as the leading normative pole for EU biometric surveillance law; rather than being a point of substitution, these analyses treat the competitiveness consideration as one element among others that should be taken into account within –not instead of– a rights-protective structure. This reflects, at least in part, a limited consensus across the literature

regarding the proper further demarcation of fundamental rights protection as the normative core of EU biometric surveillance regulation, despite other notable disagreements over whether and how effectively that constitutional principle can be meaningfully operationalized within existing legal infrastructure.

Research Gaps and Priorities for Future Scholarship

Emotion Recognition and Biometric Categorization

One of the biggest gaps that emerged in the literature concerns emotion recognition and biometric categorization systems (as opposed to identification systems, which have dominated the debate so far), which deserve greater scrutiny from researchers and legislators alike. Current evidence is the strongest regarding facial recognition in policing and public spaces, but weaker for emotion recognition (or recognized emotions), biometric categorization systems, and other emerging modalities that some authors argue have reached the level of performance that should indicate explicit prohibition (Utegen & Rakhmetov, 2023, pp. 834-35).

Importantly, the distinction between identification and categorization has been elaborated upon even less than it deserves in the literature. Identification systems respond to the question of who this person is, through comparison of a biometric sample with the database of known individuals. Instead, they claim to do things like answer the question: ‘What emotion is this person experiencing? Or ‘what demographic or behavioral segment does this person belong to? without actually even naming the individual. This difference is important from a legal perspective, because most existing proportionality jurisprudence and much of the prohibition architecture of the AI Act are built with identification - not categorisation - as the paradigmatic case, and therefore prompts consideration regarding whether or not the legal tests developed for identification transfer well to categorisation, or if it can be argued that categorisation actually entails an identifiable constellation of risks - especially stigmatisation risk, scientifically unfounded inference risk, and discriminatory profiling risk - that require a different legal framework altogether rather than simply a version adapted from the one for identification.

Debate over the scientific validity of emotion recognition technology itself is not especially new, but this controversy between members of the psychological and computer science literature has substantial legal implications: A regulatory framework that presumes proportionality essentially assumes the performance of regulated technologies with respect to their specific claims relatively simply because otherwise a technology proven incapable or unreliable in reliably carrying out its claimed function at all likely fails any necessity test even under circumstances where provisional safeguards are shown to be aloof. Rather than treating the question of whether computers can accurately identify human emotion as a settled question that is something other than legal analysis, future legal scholarship would be less cavalier and more expansive if it tried to engage a bit more directly with this underlying scientific controversy.

Long-Term Sustainable Governance of Interconnected European Systems

A further key long-term research question, widely discussed in Section 4 above, relates to the governance of EU interoperability architectures. Similarly, long-term governance of interoperable EU systems seems underresearched, despite warnings that ETIAS, CIR, and affiliated infrastructures jeopardize purpose limitation, storage limitation, and effective remedies. Future work in this area needs to progress beyond purely descriptive mapping of how these systems are technically organized which is the main contribution of existing scholarship to a more normative project that analyzes what forms of accountability and governance - be they via judicial oversight, parliamentary supervision, or independent administrative review - would be commensurate with

mitigating the cumulative privacy risk created by interoperability as more databases come online and cross-system queries increase over time.

This gap seems particularly salient because it pinpoints an important institutional bottleneck in the EU regulatory process: constitutional courts are comparatively well positioned to identify broad normative principles (necessity, proportionality, requirement of foreseeably and accessible rules), but comparatively poorly positioned to detail sector specifications (e.g., what maximum retention periods related to particular types of biometric data would be appropriate) or precise certification obligations regarding privately assembled biometric databases subsequently used by law enforcement—conditions which would need to exist for these lofty normative principles to have concrete operational effect.

Comparative and Empirical Research

Finally, the literature synthesized in this article is overwhelmingly doctrinal and normative, focusing on what the law does or should require rather than empirically tracing how biometric surveillance technologies are deployed in practice, how often the exceptions and safeguards currently available are invoked in a given population of deployments, or what specific empirical effects - misidentification rates, successful versus unsuccessful legal challenges - these deployments engender. Conversely, fusing the doctrinal sophistication of existing literature with rigorous empirical investigation into actual enforcement practice across Member States would greatly strengthen the evidentiary basis for any future legislative reform of the AI Act and its implementing measures.

Conclusions

In doing so, this article has distilled and critically aggregated existing legal literature on EU regulation of biometric surveillance technologies around (1) the uneven distribution of scholarly and regulatory attention across varying infrastructural domains and regulatory dimensions; (2) key points of contention among scholars; and (3) pathways for future scholarship to address inadequacies in both law and theory.

The analysis confirms and expands on the key finding that European regulation of biometric surveillance technologies oscillates between security and fundamental rights, due to clear governing principles but ambiguous operating rules and exceptions that leave considerable space for intrusive practices with respect to rights. Although doctrinally established and widely accepted in the literature as the right yardstick for measuring biometric surveillance regulation, the necessity and proportionality test has not been formulated with sufficient precision to delimit the enforcement discretion of law enforcement and border authorities, or of individual machine learning systems, in specific instances. While the AI Act is a real and monumental legislative accomplishment as the first comprehensive EU instrument that directly addresses biometric identification systems, it does not resolve this underlying tension—if anything, the density of exceptions, derogations, and operationally undefined provisions created by its prohibition framework has pushed this tension between security and rights protection from legislative principle into administrative or judicial interpretation, where it remains substantively unresolved.

The lopsided distribution of scholarly attentiveness uncovered in Section 2—weighing heavily in the direction of facial recognition and remote biometric identification from law enforcement, and far thinner on protest surveillance and cross-border interoperability—should be viewed not simply as an academic oddity but as a key limitation on having a high-quality policy debate that should be informing current legislative reform efforts and those forthcoming. Those legislators or regulators who depend overwhelmingly on the extant academic literature for guidance on its risks will, as a direct result of this distribution, have much more precise and rich guidance regarding the risks of police facial recognition than they do about arguably equally critical—indeed, in some respects, more structurally embedded—

risks associated with interoperable database architectures and pro-active surveillance of political assembly.

Thus, future legal scholarship, and the legislative reform such scholarship ought to inform, are well advised to concentrate greater attention on three of the issues substantiated in this analysis: development of a separate literature on emotion recognition and biometric categorization systems from the identity-focused framework that has dominated thus far; development of governance mechanisms commensurate with the cumulative privacy risks entailed by interoperable EU databases; and translation well-established constitutional principles of necessity and proportionality into legislation operating at sufficiently granular level (for example: retention period length, database certification system, onward use by public-private actors) before moving further. Policing these more neglected dimensions must surely be the way to transition EU law beyond a ‘regulatory architecture’ in which governing principles are straightforward, but abstract, towards one where those principles take genuinely concrete and enforceable effect across all practices of biometric surveillance.

References

- Akhtar, M. (2020). Police use of facial recognition technology and the right to privacy and data protection in Europe. *NAVEIN REET: Nordic Journal of Law and Social Research*, 1(9), 324–344. <https://doi.org/10.7146/njlsr.v1i9.122165>
- Barkane, I. (2022). Questioning the EU proposal for an Artificial Intelligence Act: The need for prohibitions and a stricter approach to biometric surveillance. *Information Polity*, 27, (2), 147–162. <https://doi.org/10.3233/ip-211524>
- De Hert, P. & Bouchagiar, G. (2021). Facial recognition, visual and biometric data in the US. Recent, promising developments to regulate intrusive technologies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3944536>
- Gabrielli, G. (2025). The Use of Facial Recognition Technologies in the Context of Peaceful Protest: The Risk of Mass Surveillance Practices and the Implications for the Protection of Human Rights. *European Journal of Risk Regulation*, 16, (2), 514–541. <https://doi.org/10.1017/err.2025.26>
- Hert, P. and Bouchagiar, G. (2022). Visual and biometric surveillance in the EU. Saying ‘no’ to mass surveillance practices? *Information Polity*, 27(2). 193–217. <https://doi.org/10.3233/ip-211525>
- Kagkli, A. & Angelaki, A. (2026). EU AI Act: Regulating the Gaze: The AI Act’s Approach to Remote Biometric Identification in Public Spaces. *International Journal of Social Science and Economic Research*, 11(2), 43-56. <https://doi.org/10.46609/ijsser.2026.v11i02.034>
- Kolosov, A. V. (2021). Legal regulation of the turnover of biometric data of citizens in the European Union. *Siberian Law Herald* 3(94), 89–96. <https://doi.org/10.26516/2071-8136.2021.3.89>
- Mobilio, G. (2023). Your face is not new to me – Regulating the surveillance power of facial recognition technologies. *Internet Policy Review*, 12(1). <https://doi.org/10.14763/2023.1.1699>
- Mollo, F. (2023). Il trattamento dei dati biometrici nel quadro della strategia europea dei dati, fra rischi di mass surveillance e tutela dei diritti fondamentali. *Nuovi Autoritarismi e Democrazie: Diritto, Istituzioni, Società (NAD-DIS)*. <https://doi.org/10.54103/2612-6672/22173>
- Thommandru, A., Mone, V., Shokhijakhon, F. & Mirzayev, G. (2025). Algorithmic Profiling and Facial Recognition in EU Border Control: Examining ETIAS Decision-Making, Privacy and Law. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 15(2). <https://doi.org/10.1002/widm.70013>
- Ţical, G.-M. (2025). Facial Recognition and Biometric Systems: Benefits and Challenges for Law Enforcement. *Land Forces Academy Review*, 30, 249–259. <https://doi.org/10.2478/raft-2025-0024>
- Utegen, D. & Rakhmetov, B. (2023). Facial Recognition Technology and Ensuring Security of Biometric Data: Comparative Analysis of Legal Regulation Models. *Journal of Digital Technologies and Law*, 1(3), 825–844. <https://doi.org/10.21202/jdtl.2023.36>