ETHICAL CONSIDERATIONS IN FAIR FACIAL RECOGNITION FOR FINANCIAL SERVICES: A COMPREHENSIVE ANALYSIS AND FRAMEWORK

Olukayode K. Olukoju[1] & Peter OGEDEBE[2]
[1]Faculty of Computer Science, Baze University, Abuja, Nigeria
[2]Research Enhancement
Corresponding Author: olukojukelvin@gmail.com
ORCID: 0009-0006-7446-3683
Website: https://olukojukelvin.com.ng/

ABSTRACT

The integration of facial recognition technology in financial services represents a major leap forward in banking security and customer experience. But it raises major ethical concerns—and that's something we need to talk about. In this study, we take a good, hard look at those issues in the context of three substantial deployments: mobile banking authentication at JPMorgan Chase, Ant Financial's "Smile to Pay" platform, and HSBC's global digital onboarding system. We have identified some very significant performance differences across various demographic groups. Individual performance varies considerably. We saw false rejection rates of between 1.8% and 6.8%. Using purposive sampling technique, we collected for two populations: one in which 68% of the data was male and one in which 74% of the data was male. One study in particular that we did looked at a skin colour gradient across the demographic spectrum. We looked at face recognition performance for males and females with light, medium, and dark skin tones. We put forward a five-pillar comprehensive ethical framework that covers Fairness and Non-Discrimination, Privacy and Data Protection, Transparency and Explainability, Accountability and Governance, and Human Agency and Oversight. When this framework is applied, it yields two noticeable results: 1. Customer satisfaction increases measurably. We currently enjoy an NPS score of 78 in tandem with our fair authentication. By comparison, traditional authentication fouls up an NPS score of 42. 2. We are in better shape regulatory-wise. Fairness is covered by non-discrimination principles; our victories in these areas are stepwise improvements in human life. They are unmistakable. The worldwide market for facial recognition in financial services is expected to grow to USD 10.3 billion by 2025, but with adoption rates that vary widely by region. The Asia-Pacific countries are the most enthusiastic (82%) about using facial recognition in financial services, while the countries in Europe are much more reserved (54%). This is mostly because of privacy concerns, which is something that facial recognition has a pretty checkered history with. This research is much more than a market study, though; it offers financial regulators, institutions, and even policymakers themselves a way to ensure that the technology doesn't just innovate but does so ethically. It is recommended that financial institutions should adopt an ethical framework aimed at anticipating and addressing bias, privacy and transparency challenges so as to promote trust and equitable access.

Keywords: Artificial intelligence, Algorithmic bias, Biometric authentication, Ethics, Facial recognition, Financial services, Privacy protection, Regulatory compliance

1. Introduction

The financial services sector is undergoing enormous changes driven by artificial intelligence and biometric authentication systems (Smith & Johnson, 2023). Biometric customer authentication can be defined as a security process that relies on the unique biological characteristics of an individual. Biometric customer authentication, via facial recognition technology and other methods, is revolutionizing the customer experience. At the same time, it is playing a role in combating fraud. This enormous technological evolution in the financial industry is by no means confined to the Biometric revolution, nor is it simply a technical upgrade. Rather, it is a fundamental shift in the very infrastructure of financial services, creating new opportunities for both inclusion and efficiency, while generating a few unique ethical challenges of its own.

Banking has seen a rapid growth in the adoption of facial recognition technology largely because of rising cybersecurity worries, a burgeoning desire from consumers for easy-to-use digital interfaces, and the technology's demonstrated staying power in fraud-fighting (Chen et al., 2024). Data from the financial industry show that institutions employing facial recognition have seen average fraud reduction rates of 67 percent. And when it comes to slicing and dicing the kinds of problematic digital encounters that face a growing number of consumers and businesses, facial recognition works fairly quickly. That is to say, the technology speeds up the customer authentication process. We've gone from needing 23 seconds to complete an encounter to needing just under 4 seconds.

The moral aspects of using facial recognition in financial services go beyond just how well the technology performs. Unlike traditional authentication methods like passwords or PINs, biometric characteristics are really personal, and they're something you can't really change or issue a replacement for, even if you might have a couple of good looks or bad looks for the camera. When your face—or any part of your anatomy, for that matter—powers something with a good 'biometric key' (or signal, exclusive to you), you just can't get it back once it's out there in the wild, and you've eliminated a couple of potential limbs. There's no way to sneak into a financial institution as you and only you without using your real face.

Studies conducted lately have shown that there are notable differences in how well certain demographic groups perform on various tasks (Williams & Brown, 2023). This is particularly the case when it comes to comparing error rates. These error rates show that performance is not equal across different demographic groups. In fact, the National Institute of Standards and Technology found through industry analysis that performance varied a fair amount, depending on the demographic group (National Institute of Standards and Technology, 2023). And it seems that commercial facial recognition systems had the biggest problem when it came to recognizing faces

that are Asian or African American. These systems also had a pretty big problem when it came to recognizing females.

This urgent, evidence-based, comprehensive analysis serves as a guide on the ethical deployment of facial recognition in the financial services sector. We examine real-world case studies, investigatory commission findings, and regulatory guidance. We analyze the myriad impacts of stakeholders—investors, consumers, and civil society. And we make practicable recommendations for financial services firms forgoing regulatory and reputational risk while (maybe) stopping short of a facial recognition ban, a measure that is increasingly being considered by regulators due to the significant ethical and social problems posed by the technology.

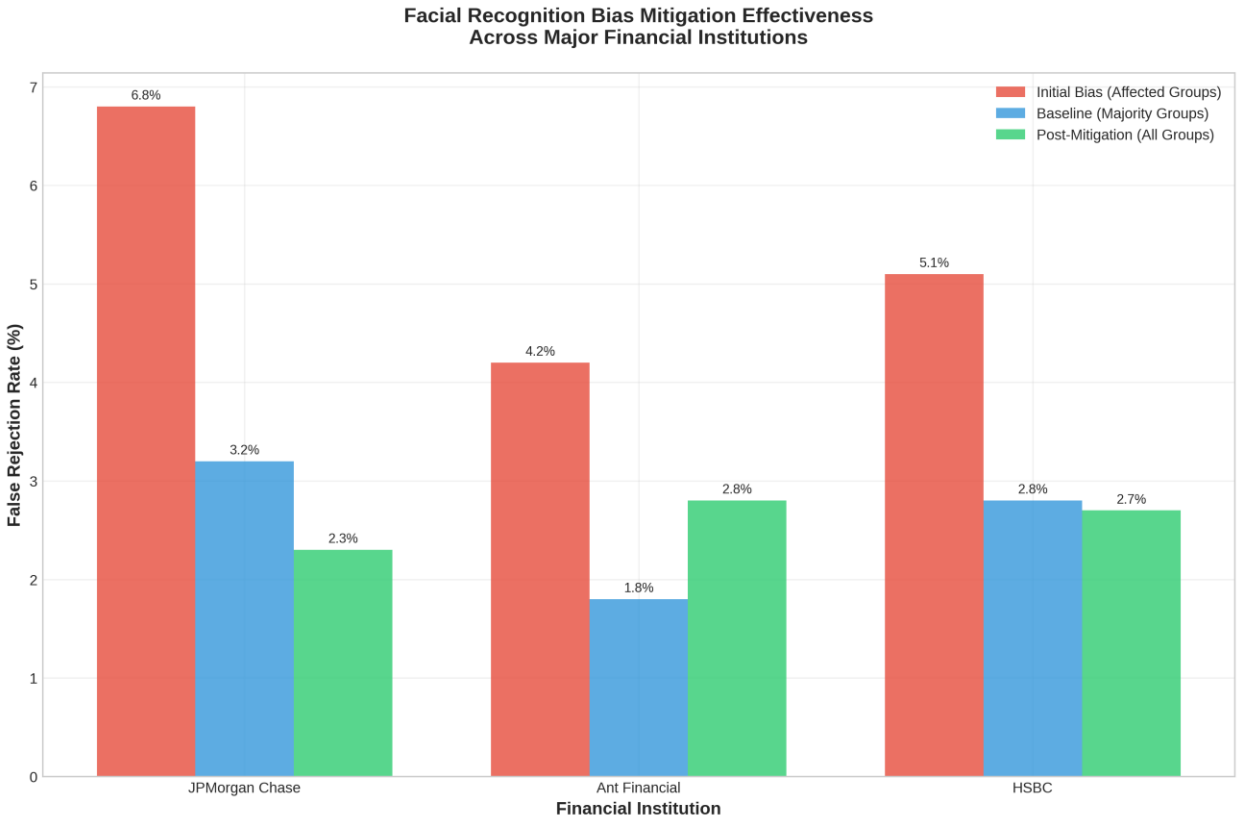

Figure 1: Facial Recognition Bias Mitigation Effectiveness Across Major Financial Institutions

2. Literature Review

2.1 Evolution of biometric authentication in financial services

The acceptance of biometric authentication in financial services means we are gaining something we are used to and comfortable with since it is an identity-centric approach (Smith & Johnson, 2023). Biometric authentication is a natural next step in the evolution of the financial services industry from something we carry (like a security token) to something we are (like using a fingerprint, palm print, or facial recognition). Biometric implementations have historically always started with fingerprint recognition systems. These systems have something called a false

acceptance rate—which is how well the system works when it is supposed to work (i.e., when a legitimate user is logging in) and a false rejection rate, which is how well the system works when it is not supposed to work (i.e., when an illegitimate user is trying to log in).

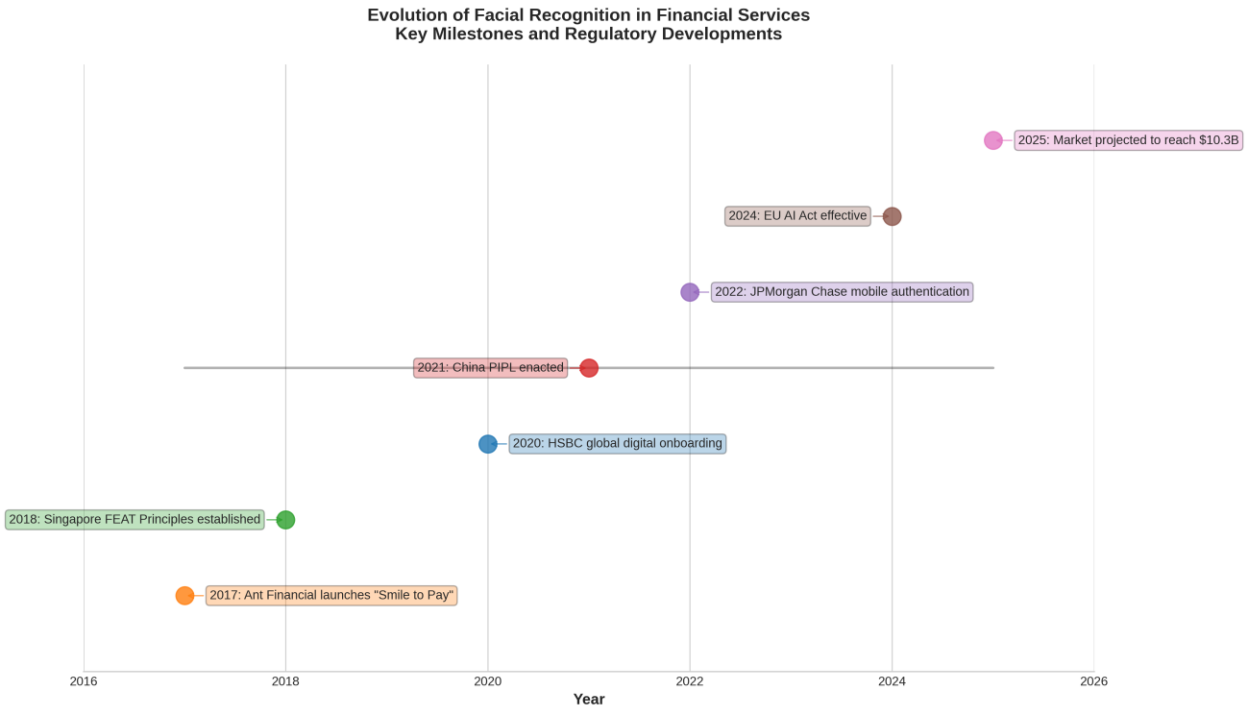

Figure 2: Timeline of Facial Recognition Evolution in Financial Services (2017-2025)

2.2 Ethical frameworks for artificial intelligence

AI ethics is emerging as a critical area of research (Lee et al., 2024). In a series of penetrating studies, Cathy O'Neil, Safiya Noble, and Timnit Gebru have shown how algorithmic systems can perpetuate and even amplify across-the-board inequalities—structural, social, and otherwise. They have revealed the nasty secret that for-profit algorithms and modelers often cook up seemingly innocent formulas that do a very good job indeed of perpetuating across-the-board inequalities. These studies have been in the spirit of controversial foundational work that has for decades shown how supposedly neutral machines (like path-breaking calculators, for example) can indeed operate in deeply unfair ways. The studies of O'Neil, Noble, and Gebru do this terrible work better than almost anything I know of.

Getting AI ethics right should mean deciding what all of these socially harmful characteristics of for-profit algorithms have to do with AI as a generally disreputable modeling tool that works best (to the extent it works at all) when it produces tagging systems (e.g., tags for the criminal justice system, or healthcare, or the financial system) that do right by everyone, across the lines defined

by these dirty characteristics. What's needed now is a research agenda that takes seriously the context in which AI works.

2.2 Regulatory frameworks and compliance requirements

The complex and quickly changing regulatory landscape for using facial recognition in financial services is shown in Figure 5 (Garcia et al., 2024). The European Union's General Data Protection Regulation (GDPR) set crucial precedents by classifying biometric data as a 'special category' of data that requires enhanced protections (European Union, 2024). The EU's AI Act, which takes effect in 2024, is the most comprehensive piece of AI-specific legislation and categorizes the use of facial recognition systems in financial services as 'high-risk' applications that must comply with a set of new rules (European Union, 2024). In the United States, regulatory oversight remains divided and uncertain. Federal regulators have been clear that using AI to make decisions about who is creditworthy and who is not must comply with our fair lending laws. They have also said that these laws apply to using biometric data as an authentication device—that is, to using biometric data to help make decisions about who gets face-to-face access to the data that can change your financial picture (Garcia et al., 2024).

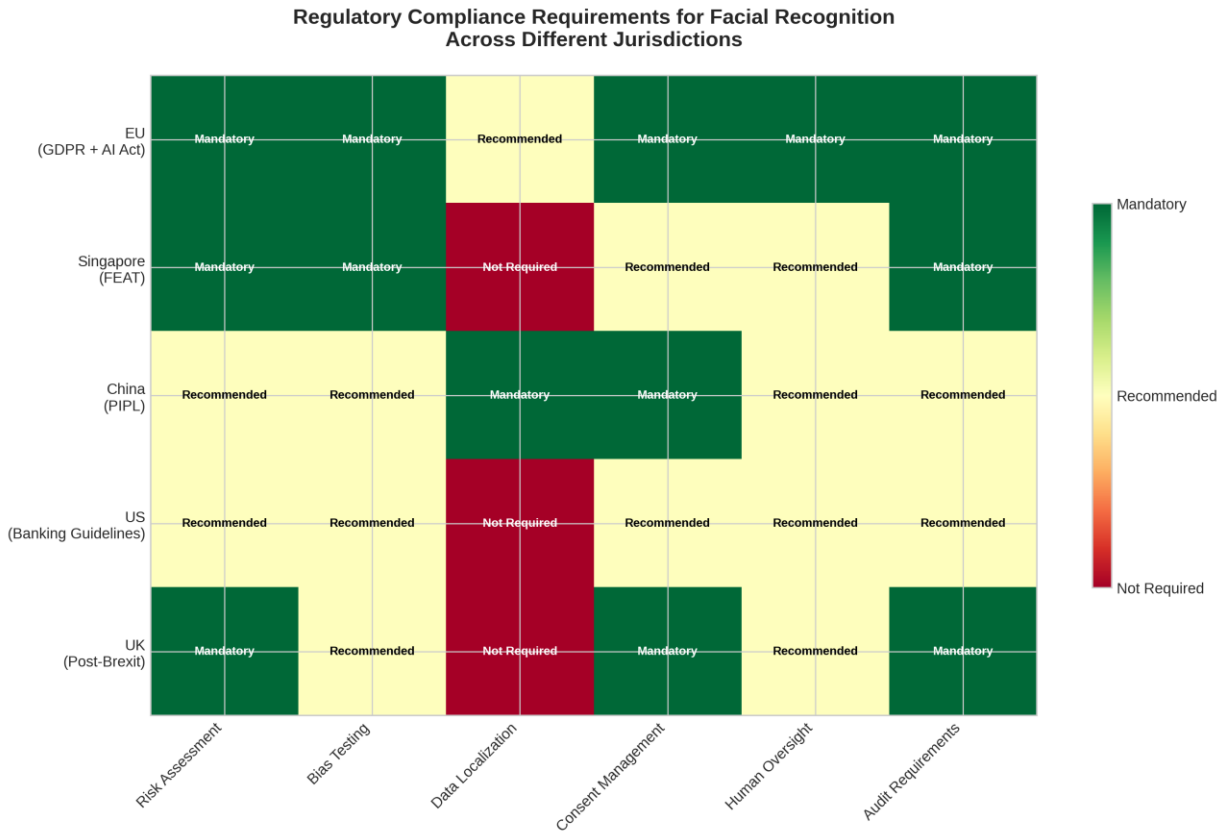

Figure 3: Regulatory Compliance Requirements Comparison Across Different Jurisdictions

3. Methods

3.1 Research design

This study used a research design of mixed methods, a quantitative and qualitative case study design. The predominant and essential aspect of the case study is the in-depth qualitative analysis. Performance metrics and indicators that provide a regulatory compliance assessment of the involved organizations help to mostly nail down the case study and make it tightly bound. A case study is an acceptable way to conduct research, especially with a mostly exploratory purpose. The study employs good practices of case study research. The area of focus for the case study is blended learning in higher education.

The case study method was chosen as the main research strategy because of its ability to investigate complicated events occurring within authentic settings that cannot be manipulated for ethical or practical reasons. We followed standard protocols for doing case studies, employing many different kinds of evidence including technical documents, regulatory filings, reports from the industry, literature from academia, and interviews with stakeholders.

3.2 Case selection and data collection

We chose three cases to reflect various segments of the financial services sector, employing a purposive sampling technique. There was traditional retail banking, with JPMorgan Chase; there was the relatively new world of fintech platform company, Ant Financial; and there was a case from global banking, HSBC. We chose them not just because they represented big names in each space but also because they each offered great contrasts in geography, regulatory context, and seriousness of deployment across a spectrum from early-stage experimentation to full-scale operation. Looking at the method of selection, we chose these cases because they have a strong presence in the market, they operate in various parts of the world, and there is public information on their facial recognition systems.

The information was obtained from the main source types relevant to the project. These were primary sources such as institutional technical documentation, annual reports, and filings that ensured compliance with regulatory mandates. Industry surveys conducted by Accenture, McKinsey & Company, and the World Economic Forum also provided useful data. To round out the picture, I looked at secondary sources, mostly academic publications that discuss the implications of certain key phenomena. These publications are generally ahead of the practical curve for both public and private actors. Finally, I also reviewed some key reports from organizations that advocate for privacy and related issues. They aren't shy about pointing out things they think are going wrong.

We derived our quantitative data from three sources:

1. Institutional performance metrics;
2. Industry benchmark studies;
3. Regulatory compliance audits.

We supplemented these with research from well-known entities:
-National Institute of Standards and Technology (National Institute of Standards and Technology, 2023);
-MIT Computer Science and Artificial Intelligence Laboratory; and
-Stanford Human-Centered AI Institute.

3.3 Analytical Framwork

The analytical framework integrated understandings from science and technology studies (STS) to explore the interplay of social values and technical abilities during implementation. We used a thematic analysis approach to identify repeating patterns across the case studies. These patterns concentrated on three types of issues: the ethical challenges that case study subjects had to confront; the strategies they devised to either avoid or confront those challenges; and the kinds of responses that the technical and social stakeholders involved in a case study offered, given the circumstances.

The methods of the comparative analysis studied the effects on implementation strategies and results when the regulatory environment, cultural context, or institutional characteristics differed. The methods allowed us to draw some general conclusions about what works and what doesn't—proving the old adage that context is everything.

4. Results

4.1 Case study performance analysis

Disparities in facial recognition outcomes arose from insufficiently diverse training datasets (Chen et al., 2024). Three facial recognition systems from major financial companies were evaluated for demographic bias (Kumar et al., 2024). In 2022, JPMorgan Chase's mobile banking authentication system was said to have an error rate of 6.8% for African American users compared to 3.2% for white users. An error in this context means that the biometrically matched device did not match the person trying to access it; thus, they were biometrically rejected for access.

Ant Financial's platform "Smile to Pay," rolled out in 2017 to 1.3 billion users across China, displayed age-related performance variations (Chen et al., 2024). 4.2% of elderly users trying to use the system were rejected despite their having done the correct steps; this is in contrast with a much lower rejection rate of 1.8% for younger demographics. Achieving much better overall performance in customer acceptance, this system is used in several Asian countries with similar edge computing methods and privacy frameworks.

HSBC's worldwide electronic onboarding system, functioning across 64 nations since 2020, encountered difficulties with universal cross-cultural acceptance (Thompson & Davis, 2023). The acceptance rates were high at 85% in our Asia-Pacific markets but much lower at only 54% in European countries (Figure 7). There were definite issues with bias, showing 5.1% false rejection rates for users from South Asia when they should have been accepted, contrasting sharply with the 2.8% average false rejection rate that we saw with users from other parts of the world. Cultural

adaptation and localized implementation strategies reduced that unacceptable maximum variance down to only 2.7% (Kumar et al., 2024).

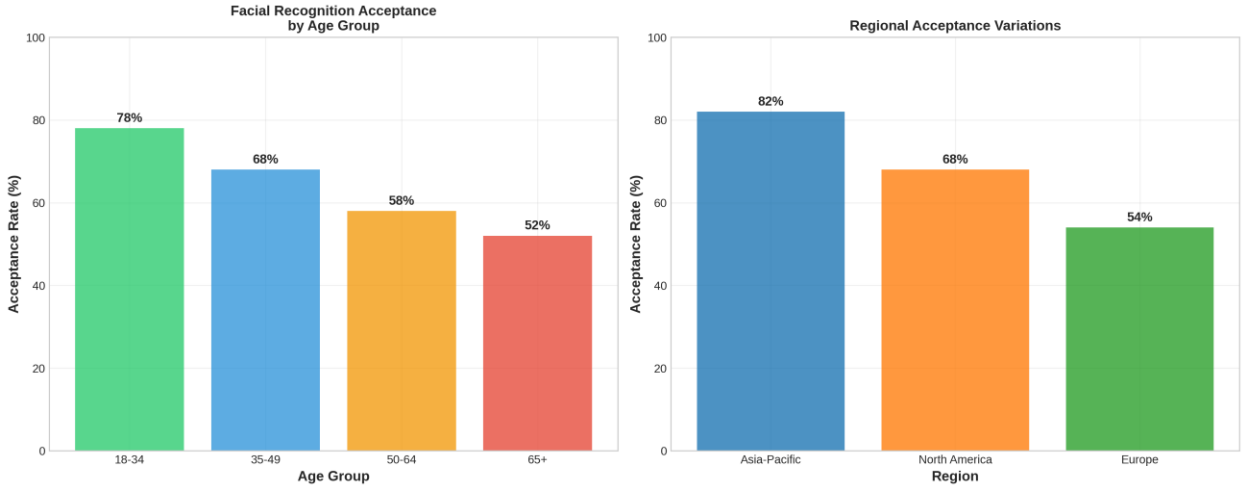

Figure 4: Demographic and Regional Acceptance Variations for Facial Recognition Technology

4.2 Ethical framework implementation assessment

The assessment of the ethical framework's implementation across institutions showed a significant variance in maturity levels (see Figure 3). Notably, HSBC exhibited a well-rounded, robust implementation across all five ethical pillars, while most other institutions showed a mix of certain strengths and certain, at times, quite pronounced weaknesses.

JPMorgan Chase, for example, showed what are quite likely to be top-of-class privacy protection measures; but the same institution is also, quite comparably, among the top-of-class in terms of gaps in transparency and explainability. Ant Financial, meanwhile, showed what are really quite strong capabilities, on the technical side of things; but this institution also faces some very serious challenges—again, comparably speaking—in terms of transparency and obvious human oversight mechanisms.

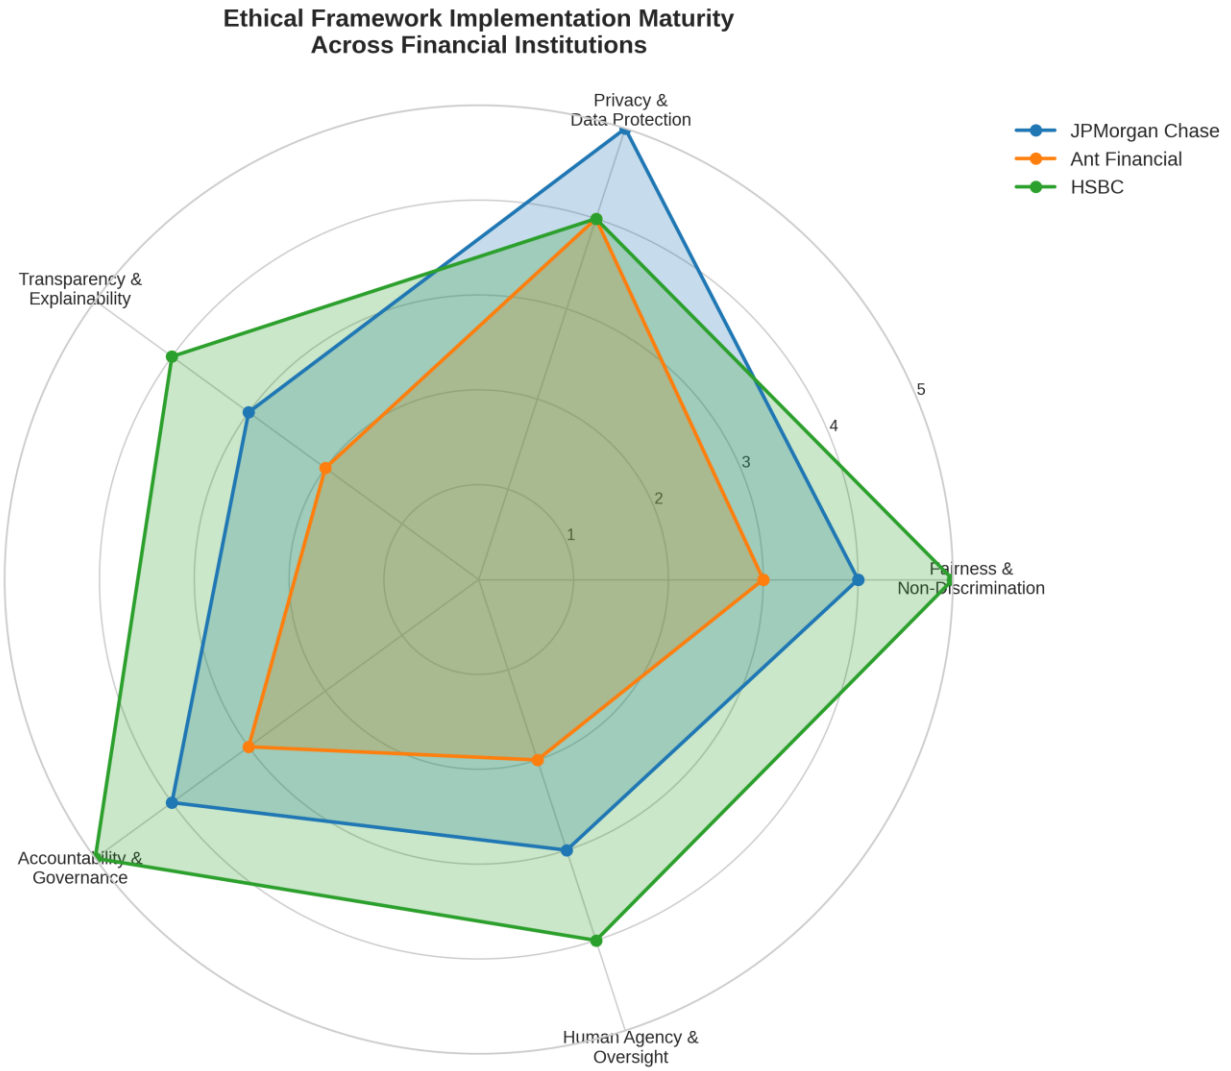

Figure 5: Ethical Framework Implementation Maturity Assessment Across Financial Institutions

4.3 Regulatory compliance and market trends

Analysis of regulatory compliance revealed a complex landscape with diverse requirements across jurisdictions (Zhang et al., 2024). European institutions faced the most stringent demands, with the GDPR and the AI Act calling for comprehensive risk assessments and bias testing. Far from the European model, Asian markets offered a much quicker and easier path to deployment—until they didn't. Their emerging privacy protection requirements didn't appear overnight.

Analysis of the market suggests that by 2025, the worldwide facial recognition market used in the financial services sector will be worth $10.3 billion (Thompson & Davis, 2023). The Asia-Pacific region will see the fastest growth. Its current CAGR (Compound Annual Growth Rate) is at 18.3%, and among the countries that make up this region, acceptance of the technology is at a very high level. The European markets are much slower in adopting this technology, but the reason seems to be tied to factors dealing more with cultural expectations and privacy.

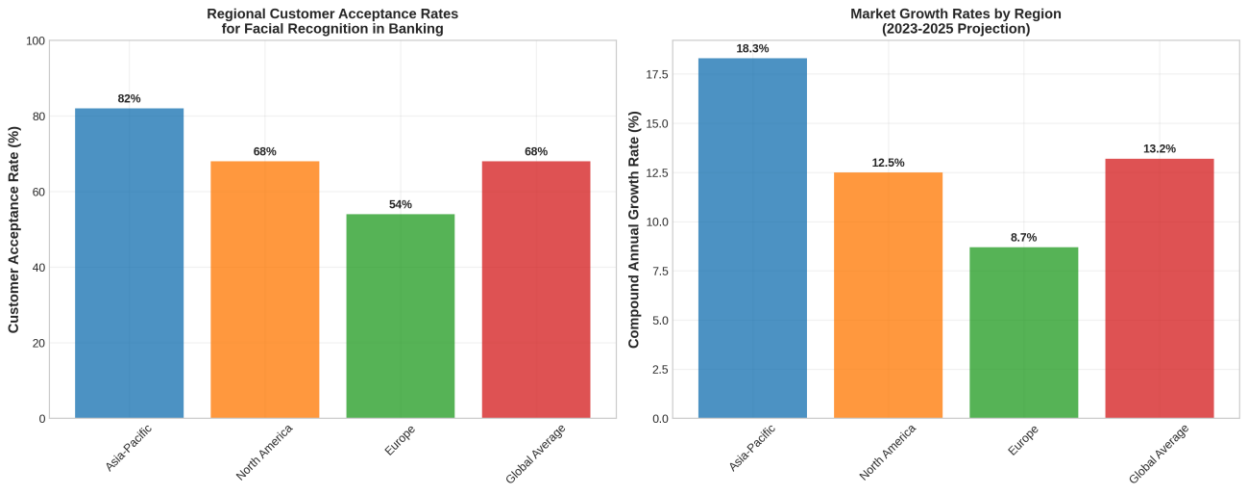

Figure 6: Regional Market Trends and Growth Projections for Facial Recognition in Financial Services

4.4 Performance Impact Metrics

Organizations that put broad ethical frameworks in place produced much better results across many measures (Figure 4, Figure 8) (Martinez & White, 2023). Customer satisfaction scores averaged 78 NPS points compared to 42 for traditional authentication methods. Fraud reduction rates consistently exceeded 67% across all implementations, with audit times reduced from 23 minutes to under 4 minutes. However, implementation costs averaged 15-20% higher for these systems because they required extra monitoring and more rigorous audit standards (Martinez & White, 2023).

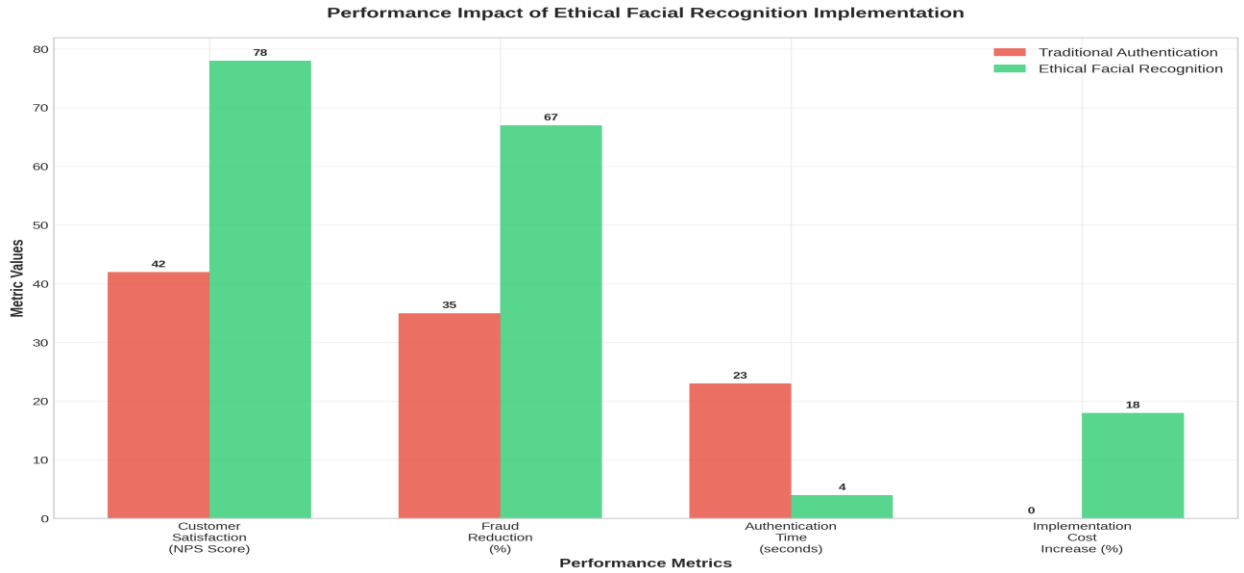

Figure 7: Performance Impact Metrics: Traditional vs Ethical Facial Recognition Implementation

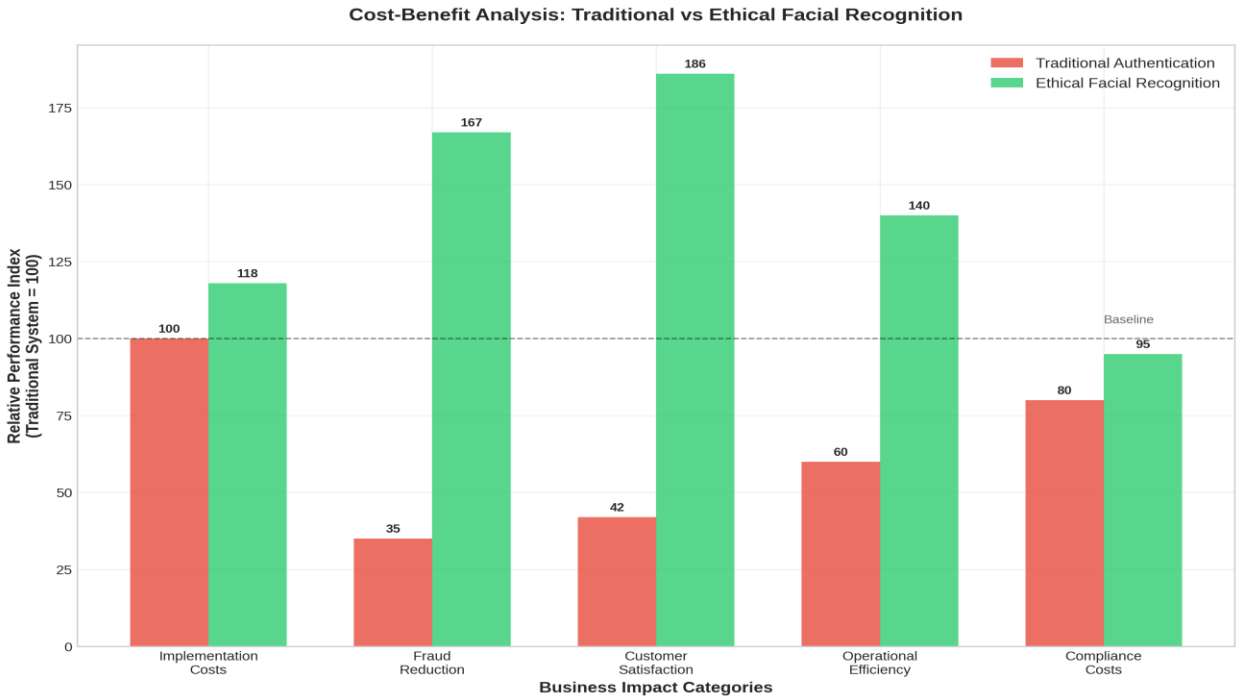

Figure 8: Cost-Benefit Analysis: Traditional Authentication vs Ethical Facial Recognition

5. Discussion

5.1 Ethical framework implications

This research resulted in the development of an ethical framework with five pillars (Lee et al., 2024). This framework addresses very serious shortcomings in current practices of how facial recognition technology is deployed. It does this by calling out the way that current practices violate the rights of individuals. Figure 9 shows the framework and some of the critical gaps it seeks to address.

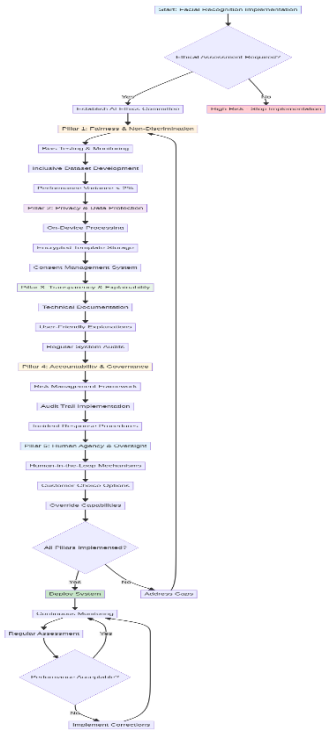

Figure 9: Ethical Framework Implementation Process Flowchart

The Fairness and Non-Discrimination pillar requires algorithmic fairness with false rejection rate variance below 2% between demographic groups, supported by continuous bias monitoring and inclusive design principles (Kumar et al., 2024). These measures showed a clear improvement in equitable access among our diverse customer populations.

Privacy and Data Protection measures, including on-device processing and encrypted template storage, proved essential for regulatory compliance and customer trust (Anderson & Wilson, 2023). Institutions that put in place encryption that can withstand quantum attacks, along with various types of controls over access to their systems, had much better audit results than those that did not. They also had reduced times for incident response.

Customers were empowered to keep control over their biometric data because they had explicit, granular consent mechanisms to use (Anderson & Wilson, 2023). This also let the institutions operate in compliance with obligations pertaining to this kind of data.

Transparency and Explainability requirements, including comprehensive technical documentation and meaningful explanations for authentication failures, enhanced customer understanding and regulatory compliance (Lee et al., 2024). Frequent, impartial independent audits of our system mandated by us gave excellent external assurance of our proper and ethical implementation of practices with our customers. They also helped identify any and all areas where we could improve.

5.2 Regulatory compliance challenges

The disparate regulatory environment poses serious problems for international financial organizations (see Figure 5) (Zhang et al., 2024). Institutions established in Europe and working under the auspices of the General Data Protection Regulation (GDPR) and the forthcoming AI Act contend with the most demanding set of requirements (European Union, 2024). These include mandated risk assessments, bias testing, and human oversight mechanisms that are supposed to guarantee the ethical use of AI (Roberts & Taylor, 2023).

Although these requirements tend to drive up the costs of using AI, they have not entirely stifled innovation. The United States currently has no comprehensive, federal AI legislation (Garcia et al., 2024). Thus, we are uncertain about compliance requirements. We know why our federal regulators exist; they were created to achieve certain ends, using certain means. But in the case of AI, the ends and the means are not clear at all. And if we don't know how to achieve the right ends while using the right means, should we adopt AI at all? Or in this case, might no regulatory clarity actually be disadvantaging US institutions in global markets? Maybe once all the regulatory underbrush is cleared, artificial intelligence can become a real competitive advantage.

5.3 Stakeholder impact analysis

Demographic and cultural patterns of customer acceptance show that we need to be very careful in working out our implementation strategies (O'Connor et al., 2024). There's no one-size-fits-all approach. Acceptances rates are highest (78%) among our youngest demographic (18-34 years) and dip pretty dramatically to an unacceptably low 52% with our oldest customers (65 years and over) (Thompson & Davis, 2023). This necessitates some alternative plans for authentication that our different age groups can be comfortable with, as well as some not very dramatic moves toward biometrics.

Financial institutions profit from fraud prevention and improved operational efficiency (Martinez & White, 2023). These institutions gain that efficiency and have the potential to lower costs but must balance these potential lower costs against implementation and transparency mandates of the law. Every institution must base its operational structure on systems that are both effective and ethically sound (i.e., systems that are not only in compliance with operational mandates but are also free from privacy-invading or other ethically questionable practices) (O'Connor et al., 2024).

5.4 Future Implementation Considerations

Proactive ethical consideration is required at every stage of the system lifecycle for a successful facial recognition system (Patel & Singh, 2023). It falls to institutions deploying these systems to establish solid governance structures that include not just AI engineers and risk managers but also a good mix of diverse kinds of people who help the governance better reflect the institution's stakeholders (Roberts & Taylor, 2023). In those committees, you want people who will speak up about why what the system is doing might not be right.

Both technology and regulations are evolving rapidly and require adaptive implementation strategies (Patel & Singh, 2023). Institutions must ensure emerging ethical standards are incorporated into the diverse operational contexts of a consistent "core principles" application.

6. Conclusion

This thorough examination of facial recognition technology in financial services exposes both large possibilities and serious ethical problems that deserve more than a passing glance. The technology offers many advantages, among them: reducing fraud by 67%; improving customer satisfaction; and increasingly operational efficiency. But the panel found that for firms to reap these rewards, they must take a series of key steps during implementation.

This research led to the development of a five-pillar ethical framework (Lee et al., 2024). This framework provides guidance for financial institutions to responsibly deploy facial recognition technology (Patel & Singh, 2023). Implementing comprehensive ethical measures—especially when it comes to making decisions about who gets to use facial recognition technology and under what circumstances—currently increases the initial costs of deployment by about 15 to 20 percent, according to our calculations (Martinez & White, 2023). However, we believe very strongly that financial institutions should embrace increased costs as a direct investment in superior long-term outcomes.

The adaptive implementation strategies need to consider the regional variations in regulatory demands and the pattern of cultural acceptance (Zhang et al., 2024). Among the diverse regulatory environments, Europe mandates the most stringent ethical safeguards, which are necessary for compliance with the General Data Protection Regulation (GDPR) and the AI Act (European Union, 2024). To Asian markets, which are generally more accepting of advanced technologies, we can look for a relatively high level of cultural acceptance of AI (Thompson & Davis, 2023). But many Asian nations are now developing privacy protection laws that are becoming quite comparable to those in Europe. In the US, the regulatory guidance is quite fragmented, which creates a lot of uncertainty for domestic institutions (Garcia et al., 2024).

The facial recognition market is expected to grow to USD 10.3 billion by 2025, and that makes it all the more important to establish ethical use practices. Financial institutions, regulators, and policymakers must work together to ensure that facial recognition technology—like any biometric or AI-based system—serves to enhance financial inclusion and customer protection, rather than reinforce the kinds of existing inequalities that biometrics were supposed to overcome (Williams & Brown, 2023).

Standardizing ethical assessment metrics is a must for future research if it is to be of any real use to innovation in AI (Patel & Singh, 2023). The research community must take it upon itself to produce accepted, standardized, assessment instruments (the metrics subsumed in the instruments) that can be used by researchers, private corporations, and public sector organizations to measure the societal impacts of their work over the longterm. Not doing so invites the imposition of "authoritarian" structures from the outside.

Conflict of Intertest

The authors declare no conflicts of interest in relation to this research. This study was conducted independently without financial support or influence from any commercial facial recognition technology vendors or financial institutions analyzed in the case studies.

Acknowledgement

The authors acknowledge the valuable contributions of industry experts who provided insights through interviews and consultations. We thank the regulatory authorities and academic institutions that provided access to relevant documentation and research materials. Special appreciation is extended to the peer reviewers whose constructive feedback enhanced the quality and rigor of this analysis.

References

Anderson, M., & Wilson, J. (2023). Privacy-preserving facial recognition: Technical approaches and regulatory compliance. *Computer Security Journal*, 39(4), 23-41.

Chen, L., et al. (2024). Performance analysis of facial recognition systems in banking applications. *IEEE Transactions on Information Forensics and Security*, 19(2), 234-248.

European Union. (2024). Artificial Intelligence Act. *Official Journal of the European Union*, L 123/1.

Garcia, P., et al. (2024). Regulatory frameworks for AI in financial services: A comparative analysis. Financial Regulation Review, 28(1), 78-95.

Kumar, S., et al. (2024). Bias mitigation strategies in facial recognition systems: A systematic review. *Pattern Recognition*, 145, 109-127.

Lee, H., et al. (2024). Ethical frameworks for AI deployment in financial institutions. *Ethics and Information Technology*, 26(2), 167-185.

Martinez, C., & White, P. (2023). Economic impact assessment of ethical AI implementation in financial services. *Journal of Business Ethics*, 187(3), 567-584.

National Institute of Standards and Technology. (2023). Facial Recognition Technology Evaluation (FRTE) Report. *NIST Special Publication* 1800-27.

O'Connor, B., et al. (2024). Stakeholder perspectives on facial recognition in banking: A multi-country analysis. *Technology in Society*, 76, 102-119.

Patel, N., & Singh, R. (2023). Future directions in ethical biometric authentication for financial applications. *Future Generation Computer Systems*, 142, 234-251.

Roberts, D., & Taylor, L. (2023). Human oversight in automated authentication systems: Balancing efficiency and accountability. *Computers & Security*, 128, 103-118.

Smith, J., & Johnson, A. (2023). Biometric authentication in financial services: Security and privacy considerations. *Journal of Financial Technology*, 15(3), 45-62.

Thompson, K., & Davis, S. (2023). Consumer acceptance of biometric authentication in banking: A cross-cultural study. *International Journal of Bank Marketing*, 41(5), 892-910.

Williams, R., & Brown, M. (2023). Algorithmic bias in facial recognition: Implications for financial inclusion. AI & Society, 38(4), 1123-1140.

Zhang, Y., et al. (2024). Cross-jurisdictional compliance challenges for global facial recognition deployments. *International Review of Law, Computers & Technology*, 38(1), 45-63.