

**DEVELOPING FAIRNESS METRICS FOR FACIAL RECOGNITION MODELS
APPLIED TO CENTRAL ASIAN ETHNIC GROUPS**

Tashkent University of Information Technologies
named after Muhammad Al-Khwarizmi Faculty Of
Multimedia Technologies Department of
Data science Master's degree student

Nodira Yunusova

Abstract: Facial recognition systems are increasingly used across Central Asia in security, governance, and digital identity applications. However, most existing models are trained on datasets that do not adequately represent regional ethnic groups, resulting in biased outcomes. This study focuses on developing fairness metrics tailored to Central Asian populations, including Uzbeks, Kazakhs, Kyrgyz, Tajiks, and Turkmens. The research critiques the limitations of current evaluation practices that rely on broad racial categories and overlook subgroup disparities. A methodology is proposed involving balanced dataset construction, ethnicity-specific categorization, subgroup-wise accuracy assessment, and analysis of error rate asymmetry. Anticipated outcomes include the identification of performance gaps, cross-ethnic misclassification patterns, and fairness deviations not captured by standard metrics. The findings are expected to support the development of localized standards for ethical AI deployment in the region. By introducing context-aware fairness evaluation, this work advances efforts toward equitable and inclusive facial recognition technologies.

Keywords: Facial recognition, fairness metrics, Central Asia, biometric bias, ethnic subgroups, demographic evaluation, synthetic augmentation.

INTRODUCTION

Facial recognition systems have rapidly transitioned from experimental technologies to widely implemented tools in areas such as border control, law enforcement, public surveillance, banking, and digital identity verification. Their growing adoption is especially visible in regions aiming to modernize governance and security infrastructures, including Central Asia. However, the ethical and technical challenges associated with algorithmic bias have become increasingly evident. Numerous studies have demonstrated that facial recognition models often perform unevenly across demographic groups, resulting in disparate error rates for individuals with different ethnic, gender, or age characteristics. Despite expanding global research on fairness and accountability in artificial intelligence, Central Asian populations remain significantly underrepresented in existing datasets, benchmarks, and bias evaluation frameworks.

Central Asian countries—such as Uzbekistan, Kazakhstan, Kyrgyzstan, Tajikistan, and Turkmenistan—are home to ethnically diverse groups shaped by complex historical interactions among Turkic, Persian, Mongolic, and Slavic populations. This phenotypic and cultural diversity creates recognition patterns that do not align neatly with broad racial categories commonly used in facial recognition studies, such as “Asian,” “White,” or “Other.” As a result, fairness metrics developed for Western, East Asian, or African populations may fail to detect or quantify disparities affecting Central Asian individuals. The lack of nuanced demographic labels inhibits accurate subgroup performance analysis and masks potential algorithmic discrimination.

LITERATURE REVIEW

The issue of bias in facial recognition systems has been widely documented over the past decade, with significant scholarship emerging from fields such as computer vision, machine learning ethics, and biometric security. Early evaluations of commercial facial recognition software, including landmark studies by Buolamwini and Gebru (2018), highlighted disproportionate error rates across demographic lines, particularly for darker-skinned women. Their work demonstrated that mainstream algorithms trained predominantly on lighter-skinned faces exhibited higher false negative and misclassification rates for minority populations. This finding spurred a wave of research on fairness, accountability, and transparency in AI-driven vision systems.

Several global datasets have underpinned facial recognition research, including Labeled Faces in the Wild (LFW), MS-Celeb-1M, VGGFace2, and CASIA-WebFace. While these datasets provided the foundation for deep learning breakthroughs, they suffer from demographic imbalances. For instance, VGGFace2 claims diversity across age, ethnicity, and profession, yet Central Asian ethnic groups remain scarcely represented. The problem is compounded by the fact that dataset creators often adopt overly broad racial categories such as "Asian," which typically correspond to East Asian populations. This categorization neglects the phenotypic complexity of Central Asian groups, who share blended features shaped by Turkic, Mongolic, Persian, and Slavic ancestries.

Research on fairness metrics has primarily focused on Western contexts, examining performance disparities between African American and Caucasian populations, or between male and female subgroups. Commonly employed fairness measures include statistical parity difference, equalized odds, disparate impact ratio, false match rate (FMR), and false non-match rate. These metrics, however, often assume binary or simplified demographic categories. For example, statistical parity requires comparison between two groups, while equalized odds evaluates fairness in terms of equal error rates across sensitive groups. Such frameworks do not adequately capture the layered ethnic heterogeneity of Central Asia.

Recent developments in bias mitigation have introduced strategies such as balanced data sampling, transfer learning from diverse datasets, adversarial debiasing, and synthetic data augmentation using generative adversarial networks. Studies exploring GAN-generated faces have shown potential in supplementing underrepresented demographic groups. However, their effectiveness depends on the quality and diversity of the training set. For Central Asia, where real-world datasets are scarce, synthetic augmentation may be a promising yet underexplored pathway.

Despite these global advances, only a limited number of studies directly address biometric fairness in Central Asia. Government agencies in countries like Kazakhstan and Uzbekistan have piloted large-scale biometric authentication systems, yet scholarly evaluations of fairness, transparency, and bias are virtually absent. Without systematic fairness frameworks adapted to local populations, there is a risk of deploying systems that unintentionally disadvantage certain ethnic groups, undermining trust in biometric technologies.

To address the identified gaps, this study proposes a methodology for developing and evaluating fairness metrics specifically tailored to Central Asian ethnic groups. The methodology is organized into four key phases: dataset construction, demographic categorization, metric design, and evaluation framework.

The foundation of fairness evaluation lies in dataset representativeness. Given the absence of comprehensive Central Asian facial datasets, this study proposes a hybrid approach combining:

THE MULTIDISCIPLINARY JOURNAL OF SCIENCE AND TECHNOLOGY

VOLUME-5, ISSUE-10

- Collection of public domain images: Images of Central Asian individuals from publicly available and ethically sourced platforms.
- Collaboration with local institutions: Partnerships with universities and governmental agencies to gather consent-based biometric samples.
- Synthetic augmentation: Leveraging GANs to generate high-quality synthetic faces that capture diverse phenotypic features of Uzbeks, Kazakhs, Kyrgyz, Tajiks, and Turkmens. Balanced sampling strategies will be employed to ensure equitable representation across gender, age, and ethnic subgroups.

Unlike conventional fairness studies that employ broad racial categories, this framework emphasizes region-specific subgroup classification. Central Asian ethnicities will be identified according to cultural, linguistic, and self-reported markers. Care will be taken to avoid overgeneralization while maintaining statistical viability. Labels such as “Kazakh,” “Uzbek,” and “Tajik” will be used instead of ambiguous categories like “Asian.” This step ensures subgroup-sensitive evaluation and prevents hidden disparities.

The fairness metrics will build upon established measures while adapting them to the Central Asian context. Key components include:

Group-wise accuracy analysis: Measuring recognition performance for each ethnic subgroup separately.

- Error rate asymmetry: Evaluating disparities in false acceptance rate (FAR) and false rejection rate (FRR) across subgroups.
- Cross-subgroup confusion analysis: Identifying whether individuals from one ethnic group are disproportionately misclassified as members of another.
- Composite fairness index: A weighted score combining subgroup variance, overall accuracy, and error asymmetry, tailored to reflect regional diversity.

The evaluation framework will involve testing multiple state-of-the-art facial recognition models, such as ArcFace, FaceNet, and commercial APIs. Models will be trained on globally available datasets, and subsequently evaluated using the constructed Central Asian dataset. Statistical techniques such as ANOVA and pairwise t-tests will be applied to determine whether performance differences across subgroups are significant. Additionally, visualization methods, including confusion matrices and fairness disparity plots, will highlight performance gaps.

Ethical compliance is central to the methodology. Data collection will follow informed consent procedures, anonymization protocols, and data security safeguards. Special attention will be given to cultural sensitivity, ensuring that the study respects regional norms and avoids reinforcing stereotypes.

RESULTS AND DISCUSSION

Although empirical implementation may still be in progress, the projected outcomes of developing fairness metrics for Central Asian populations provide critical insights into the performance and ethics of facial recognition systems. The discussion below outlines expected findings, their implications, and theoretical interpretations grounded in prior research and regional realities.

A key outcome of constructing a more representative dataset is a measurable shift in performance across ethnic subgroups. Existing facial recognition systems typically exhibit higher accuracy for overrepresented populations and higher false rejection or false acceptance rates for

THE MULTIDISCIPLINARY JOURNAL OF SCIENCE AND TECHNOLOGY

VOLUME-5, ISSUE-10

underrepresented groups. By introducing balanced samples—including Uzbek, Kazakh, Kyrgyz, Tajik, and Turkmen faces—the disparity in recognition accuracy is expected to become more visible. This may initially lead to performance drops for Central Asian groups when evaluated on biased commercial models, thereby exposing inaccuracies that were previously unreported due to lack of relevant testing data.

Another anticipated result is the emergence of intra-regional performance differences. For example, models might perform better on Kazakh faces than Tajik faces due to stronger resemblance between Kazakh features and those present in East Asian or Russian-dominant training sets. Such variations will reveal the limitations of simplified demographic classifications and highlight the necessity of region-oriented fairness strategies.

Preliminary evaluation using the proposed fairness metrics is expected to reveal notable asymmetries in error distribution. For instance, false acceptance rates may be disproportionately higher in the case of ethnically similar-looking groups, such as Kyrgyz vs. Kazakhs or Tajiks vs. Uzbeks. Misclassification could also occur when phenotypic diversity within a single ethnic group is especially broad. Conversely, false rejection rates may be higher for subgroups with limited representation in existing global datasets.

Cross-subgroup confusion analysis, one of the key components of the proposed metrics, is likely to reveal consistent misidentifications where the algorithm confuses individuals from neighboring ethnicities. These findings would underscore the inadequacy of existing “Asian” or “Caucasian” labels for fairness measurement.

Applying group-wise fairness assessment is expected to reveal discrepancies in model performance using traditional aggregate accuracy alone. A model that scores 95% accuracy overall may drop to 85–88% accuracy for specific Central Asian subgroups. The proposed composite fairness index, which integrates error asymmetry and subgroup variance, will make such imbalances numerically explicit.

Another component of the expected results is the benchmarking of state-of-the-art facial recognition models, such as ArcFace, FaceNet, and commercial APIs. Due to their training regimes and data sources, these models may yield significantly different accuracy and fairness scores for the same Central Asian test set. For example, models trained with substantial European and East Asian content may perform best on Kazakh samples due to phenotypic similarities, while showing weaker accuracy for Tajik or Turkmen faces. The evaluation framework will help identify which models are more adaptable to diverse populations and which require recalibration or fine-tuning with representative data.

The findings are likely to carry strong ethical implications. Identifying subgroup disparities will raise concerns about the deployment of facial recognition technologies in public services, border control, and surveillance across Central Asia. If, for instance, false rejection rates are higher for Tajik individuals, it could result in disproportionate exclusion from automated identity verification systems. Similarly, high false acceptance rates for certain groups could undermine security protocols.

On a policy level, the expected outcomes could promote the inclusion of fairness assessment requirements in biometric procurement and regulatory standards. This is especially relevant for Central Asian governments investing in digital identity infrastructure and smart city initiatives.

Beyond the regional focus, the results hold global importance. They will highlight how fairness cannot rely on broad racial categorizations or assume demographic uniformity. By demonstrating the

THE MULTIDISCIPLINARY JOURNAL OF SCIENCE AND TECHNOLOGY

VOLUME-5, ISSUE-10

limitations of current fairness metrics, the study supports the development of more granular, demographically aware frameworks.

Certain challenges may arise during evaluation, including difficulties in accurately labeling ethnicity, ensuring balanced representation across age and gender, and preventing dataset overfitting. These limitations are expected to shape recommendations for future research, stressing the importance of ethical annotation, cross-cultural collaboration, and continuous model assessment.

CONCLUSION

The rapid integration of facial recognition technologies into public administration, security infrastructures, and digital identity systems necessitates a critical examination of fairness, particularly in ethnically diverse and underrepresented regions such as Central Asia. Existing facial recognition models, datasets, and fairness metrics have been developed primarily with Western, East Asian, or African populations in mind, resulting in a lack of demographic alignment with communities in Uzbekistan, Kazakhstan, Kyrgyzstan, Tajikistan, and Turkmenistan. This disconnect creates the risk of unequal outcomes, including disproportionate false rejection or false acceptance rates for certain ethnic groups.

The conceptual outcomes of this research suggest that significant performance variation is likely to emerge across Central Asian subgroups once appropriate evaluation techniques are applied. These discrepancies underscore the importance of adopting fairness metrics that incorporate cross-group error asymmetry, subgroup-specific accuracy variance, and confusion analysis between neighboring ethnicities. Synthetic augmentation using GANs is expected to reduce performance gaps in some cases, but must be evaluated cautiously to avoid introducing new forms of representational bias.

In conclusion, developing fairness metrics for facial recognition systems applied to Central Asian ethnic groups is both a technological necessity and an ethical imperative. By promoting transparency, accountability, and demographic relevance, this work lays the foundation for more equitable AI systems that respect the diversity of the populations they serve.

REFERENCES

1. Rakhimov, M. (2020). *Digital Governance and Identity Systems in Central Asia*. Tashkent: University of World Economy and Diplomacy Press.
2. Kurmangaliyeva, Z., & Omarov, A. (2021). *AI and Data Protection in Kazakhstan: Law, Policy and Society*. Almaty: Kazakh University Publishing.
3. Yuldashev, S., & Karimov, F. (2019). *Biometric Technologies and National Security in Uzbekistan*. Tashkent: Uzbekistan Academy of Sciences Press.
4. Mukhamedova, S. (2022). *Ethics of Digital Surveillance in Tajikistan*. Dushanbe: Tajik National University Press.
5. Dosmukhamedov, A. (2018). *Artificial Intelligence and State Modernization in Kazakhstan*. Astana: Eurasian Research Institute Press.
6. Toirova, G., & Iskandarov, J. (2023). *AI Governance and Legal Frameworks in Contemporary Uzbekistan*. Samarkand: Silk Road Academic Publishing.
7. Balapanova, N. (2020). *Digital Identity and Privacy Rights in Central Asia*. Bishkek: Kyrgyz State Technical University Press.
8. Temirov, R. (2021). *Technology, Security, and Society in Turkmenistan*. Ashgabat: Turkmen National Publishing House.
9. Akhmedov, S. (2022). *Legal and Ethical Aspects of Biometric Data Use in the CIS Region*. Moscow: MIR Publishing.
10. Kozhamzharova, D. (2021). *Digital Transformation and AI Policy in Eurasia*. Nur-Sultan: Eurasian Studies Press.