Facial Demography Analysis of the LAION Dataset

IRIS DOMINGUEZ-CATENA[1], Public University of Navarre, Spain
DANIEL PATERNAIN, Public University of Navarre, Spain
MIKEL GALAR, Public University of Navarre, Spain

arin
Artificial Intelligence & Machine Learning Research

upna
Universidad Pública de Navarra
Nafarroako Unibertsitate Publikoa

Fourth European Workshop on Algorithmic Fairness (EWAF'25)
Eindhoven, June 30–July 2, 2025

[1]iris.dominguez@unavarra.es

Introduction

- Large-scale image-text datasets (e.g., LAION[2]) are critical for modern AI systems.
- **Problem**: Dataset reutilization leads to pervasive biases across models (T2I and others).
- **Goal**: Analyze LAION for:
  - General representational biases (age, gender, race).
  - Intersectional biases (e.g., age-gender, race-age).
- **Why it matters**:
  - Most works focus on model bias, analyzing each model individually[3].
  - Identifying these biases early can open new routes for bias mitigation.

[2]Releasing Re-LAION 5B: Transparent Iteration on LAION-5B with Additional Safety Fixes. https://laion.ai/blog/relaion-5b

[3]Yixin Wan, Arjun Subramonian, Anaelia Ovalle, Zongyu Lin, Ashima Suvarna, Christina Chance, Hritik Bansal, Rebecca Pattichis, and Kai-Wei Chang. 2024. Survey of Bias In Text-to-Image Generation: Definition, Evaluation, and Mitigation. https://doi.org/10.48550/arXiv.2404.01030

Methodology

1. **Dataset**:
   - Random sample: 500,000 URLs from *ReLAION-2B-en*.
2. **Tools**:
   - **RetinaFace**[4]: 37,000 faces detected.
   - **FairFace**[5]: Automatic demographic estimation (age, gender, race).
3. **Bias Analysis**:
   - Representational bias through demographic group proportions.
   - Intersectional bias through the *Ducher's Z* metric.

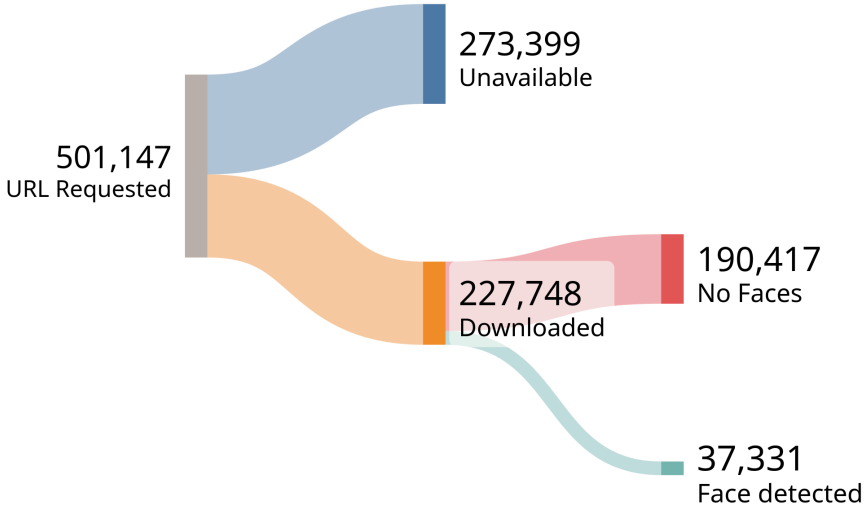

[4]Sefik Serengil and Alper Ozpinar. 2024. A Benchmark of Facial Recognition Pipelines and Co-Usability Performances of Modules. J. Inf. Technol. 17, 2 (2024), 95–107. https://doi.org/10.17671/gazibtd.1399077

[5]Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In 2021 IEEE Winter Conf. Appl. Comput. Vis. WACV. IEEE, Waikoloa, HI, USA, 1547–1557. https://doi.org/10.1109/WACV48630.2021.00159

Methodology

1. **Dataset**:
   - Random sample: 500,000 URLs from *ReLAION-2B-en*.
2. **Tools**:
   - **RetinaFace**[6]: 27,000 faces detected
   - **FairFace**[7]: A
     (age, gender
3. **Bias Analysis**:
   - Representa
     group propo
   - Intersectio

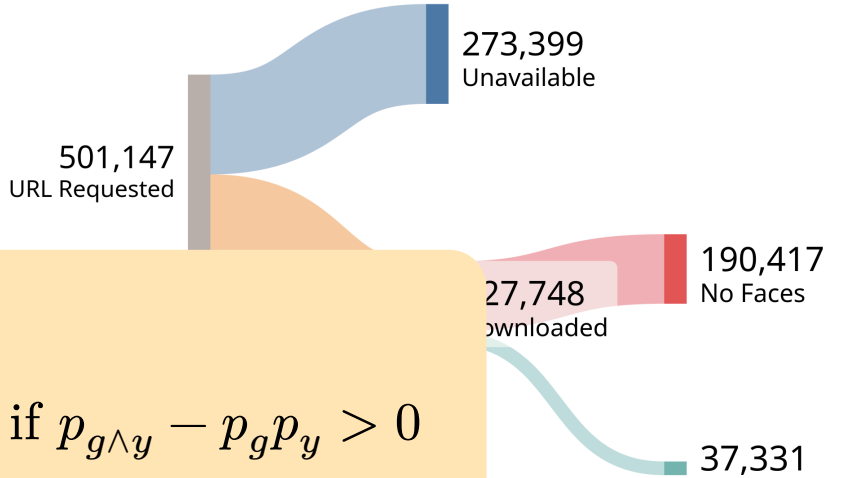

Ducher's Z:

$$Z(X,g,y) = \begin{cases} \frac{p_{g\wedge y} - p_g p_y}{\min[p_g, p_y] - p_g p_y} & \text{if } p_{g\wedge y} - p_g p_y > 0 \\ \frac{p_{g\wedge y} - p_g p_y}{p_g p_y - \max[0, p_g + p_y - 1]} & \text{if } p_{g\wedge y} - p_g p_y < 0 \\ 0 & \text{otherwise} \end{cases}$$

[6]Sefik Serengil and Alper Ozpinar. 2024. A Benchmark of Facial Recognition Pipelines and Co-Usability Performances of Modules. J. Inf. Technol. 17, 2 (2024), 95–107. https://doi.org/10.17671/gazibtd.1399077

[7]Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In 2021 IEEE Winter Conf. Appl. Comput. Vis. WACV. IEEE, Waikoloa, HI, USA, 1547–1557. https://doi.org/10.1109/WACV48630.2021.00159

Demographic Distribution

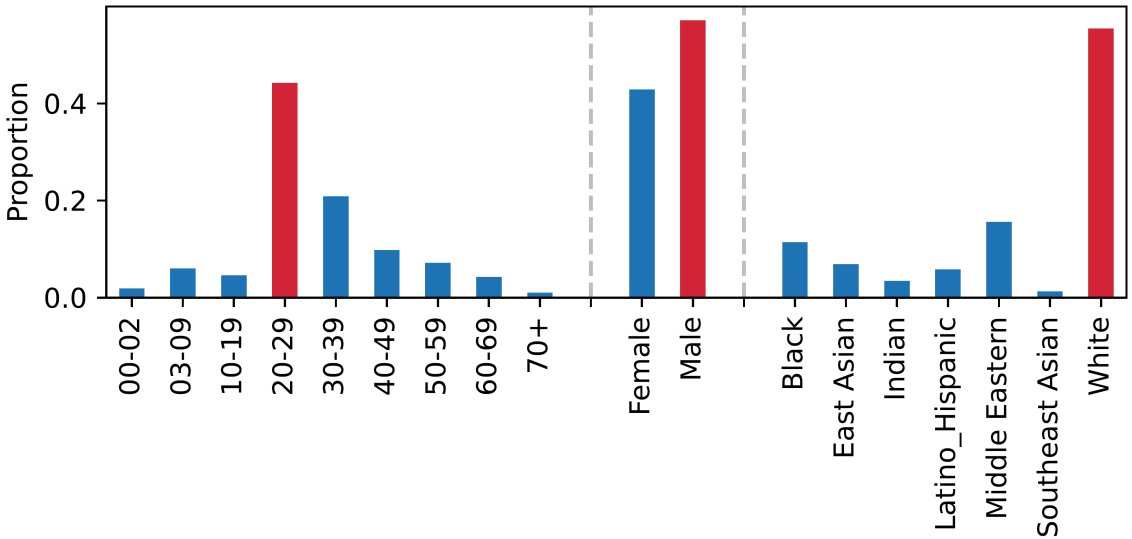

- **Age**: Overrepresentation of individuals aged 20–29 (44%).
- **Gender**: Male bias (57% of faces).
- **Race**: White individuals overrepresented (55%); Southeast Asian underrepresented (1%).

Intersectional Bias

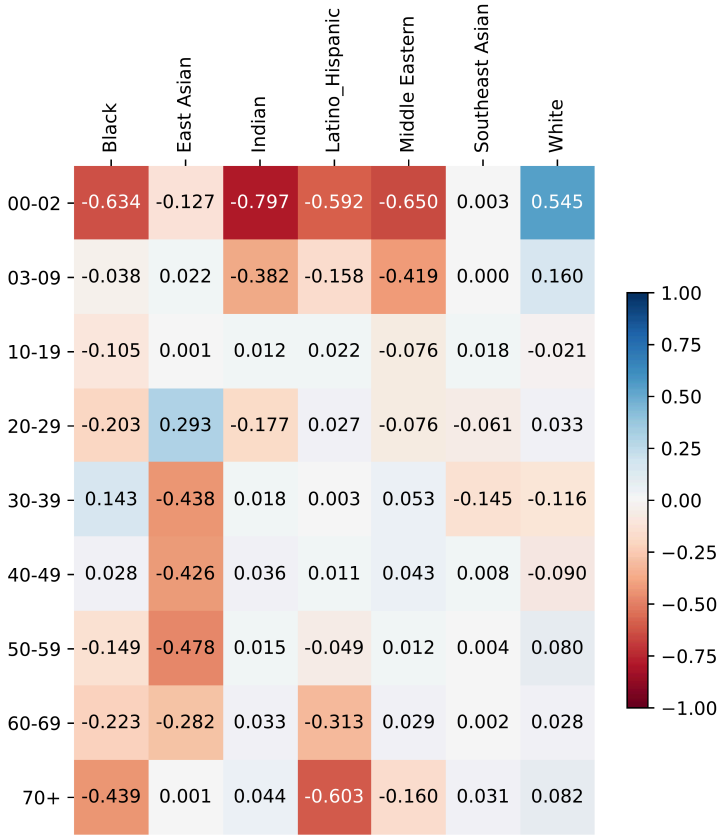

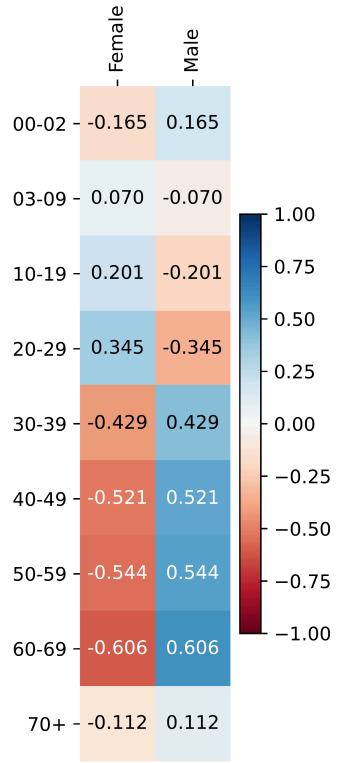

- **Race-Gender**: no significative biases found.
- **Race-Age**: Strong underrepresentation of non-White infants, middle-aged East Asian and elderly Black and Latino Hispanic individuals.
- **Age-Gender**: Middle-aged women underrepresented, young women overrepresented.

Conclusion

- **Findings**:
  - Severe demographic imbalances in the independent demographic categories.
  - Bias issues at multiple intersections of demographic categories (i.e. underrepresentation middle-aged women).
- **Implications**:
  - Biases could propagate to T2I AI models trained on LAION, missrepresenting individuals and populations.
- **Limitations**:
  - Reliance on auxiliary models (RetinaFace, FairFace).
  - Predefined demographic categories miss nuances (e.g., multiracial and gender identities).

Thanks for your attention!

Thanks for your attention!

⚠️ And beware of excesive dataset reutilization ⚠️