GINN: Towards Gender InclusioNeural Network

Original
GINN: Towards Gender InclusioNeural Network / Berta, M.; Vacchetti, B.; Cerquitelli, T.. - (2023), pp. 4119-4126. (Intervento presentato al convegno 2023 IEEE International Conference on Big Data (BigData) tenutosi a Sorrento (IT) nel 15-18 december 2023) [10.1109/BigData59044.2023.10386328].

Availability:
This version is available at: 11583/2987325 since: 2024-03-26T16:10:14Z

Publisher:
IEEE

Published
DOI:10.1109/BigData59044.2023.10386328

(Article begins on next page)

GINN: Towards Gender InclusioNeural Network

Matteo Berta
Department of Control
and *Computer Engineering Dept.*
Politecnico di Torino
Turin, Italy
s295040@studenti.polito.it

Bartolomeo Vacchetti
Department of Control
and *Computer Engineering Dept.*
Politecnico di Torino
Turin, Italy
bartolomeo.vacchetti@polito.it

Tania Cerquitelli
Department of Control
and *Computer Engineering Dept.*
Politecnico di Torino
Turin, Italy
tania.cerquitelli@polito.it

***Abstract*—Today's data-driven systems and official statistics often oversimplify the concept of gender, reducing it to binary data, with far-reaching implications for policy development and equitable access to services. This simplification can lead to misclassification and discrimination against individuals who identify as non-binary.**

We are working to advance our research in this area to develop new, more equitable approaches that can avoid discrimination based on gender identity. Within this research framework, our primary focus is on mitigating the problem of underrepresentation and, in some cases, the complete absence of non-binary individuals in data collection.

With this goal in mind, we present the GINN **Gender InclusioNeural Network. This is our first attempt to develop an equitable neural network that accurately identifies gender in a multiclass context and includes individuals whose gender identity does not fall on the binary spectrum. To achieve this goal, we conducted a comprehensive comparative analysis of several fine-tuned neural network models. Our goal was to gain a deep understanding of the crucial distinguishing features in gender identify classification and to highlight the limitation of current methods using explainable AI techniques.**

The initial results are promising and demonstrate the effectiveness of a fine-tuned EfficientNetB0 model in accurately categorizing images of individuals into their self-reported gender, but we are skeptical about the application in a real-world scenario because of the amount of data available about non-binary people at the moment.

***Index Terms*—Gender Equality, Neural Network, Automatic Gender Recognition, Non-Binary**

I. Introduction and Motivation

Gender identity refers to a person's internal understanding of their own gender, which they privately experience in their sense of self. The definition of gender identity can be a complex and highly personal matter, as it may be consistent with or different from the sex assigned at birth based on physical sex characteristics. The term "non-binary" serves as an inclusive umbrella category that encompasses identities outside of traditional male and female gender identities or those that fall on a spectrum in between. Each person's gender identity is unique and lies on a broad spectrum. Whereas "genderqueer" is a term that signifies a rejection of rigid gender binaries and, in some cases, a direct challenge to societal institutions that reinforce such binaries, as discussed in [1]. Despite some progress in social acceptance, non-binary individuals still struggle to gain visibility and representation, as discussed in [2] [3]. The problem of invisibility also extends to policy development and practice, where non-binary and genderqueer individuals may be marginalized due to the prevailing binary gender categorization in health surveillance. Artificial Intelligence and Big Data applications have traditionally treated gender as a binary attribute, leading to the exclusion of individuals whose gender identity does not conform to this binary norm, as highlighted in [4].

Dealing with a binary attribute in gender identity has profound implications, Joni Seager pointed out that "without the right categories for recording data gender, the right kind of gender data cannot be collected, and without the right kind of data there can be no social change". This approach can be understood in the context that "gender is a historical situation rather than a natural fact," meaning that a person's gender is a product of actions and experiences over a lifetime rather than an innate characteristic (see [5] for more information).

Younger generations have embraced a more inclusive vocabulary for expressing gender identity and have demonstrated a stronger commitment to gender equality, diversity, and the rights of gender and sexual minorities, as seen in [6].

Automatic gender recognition systems have often attempted to predict gender and struggle to account for genders outside the binary spectrum. Similarly, facial recognition systems have difficulty identifying the gender and sexual orientation of individuals who do not conform to traditional binary norms. A more comprehensive gender classification system is needed to appropriately and respectfully account for individuals with gender identities that do not fit within the binary framework. Gender plays an important role in various applications, such as access control, human-computer interaction, and video surveillance. Accurate prediction of gender is also critical for demographic studies, as misclassification of individuals based on gender can bias research and decision-making statistics, exacerbating problems of underrepresentation. The ability to accurately determine an individual's gender is also important for expanding advertising categories. For example, Facebook has increased its gender options for English speakers in the UK and US from 2 to 58, but still groups all genders into a binary system within advertising categories, highlighting the challenge of visibility in this context (see [4]for more details)

The goal of our work is to try to include people whose gender is along the gender spectrum in a common computer vision task called Automatic Gender Recognition, so as to try to make the existing technology used on a daily basis with a wide variety of use cases fairer and less discriminatory. In order to do so we conducted a comparative analysis of different models on a self-built dataset with images of males, females and non-binary individuals. We analyzed the results through an explainability model, so as to understand which features are the most important in the classification for each class. The results obtained and the analysis of the various features have made us realize that there is room for the creation of an inclusive and fair Automatic Gender Recognition system, but that it is necessary to work on data collection. The paper is divided as follows. In Section II there is the literature review while in section III the GINN methodology is presented. Section IV contains a comparative analysis of different models in a non-binary gender classification task. Finally section V contains the conclusions and future works.

After an initial literature review, in which we tried to have a better understanding of the topic from a sociological perspective, the focus of the paper shifts to the methodology of GINN, focusing on the creation of the used dataset and on the pipeline. After that we moved on to the analysis of the results obtained, divided into quantitative results, to compare the performance of the models, and a qualitative one through the use of explainable AI, which is very important for the purposes of our research. Thanks to the insights gained from the results we were able to construct a deep discussion merging the technical and sociological aspects described in the other sections of the paper.

II. Literature Review

In this section we want to analyze different aspects of the task starting from the sociological aspect and the related literature. After analyzing the task of gender classification, its meaning and the use cases, we focused on the documented bias related to the specific task and then we performed an analysis of Automatic Gender Prediction systems from an ethical point of view.

A. Gender Classification Applications

Gender could be considered a soft biometric data, a subsection of biometric data that are relatively less distinct or unique to traditional biometric data such as fingerprints. Soft biometrics rely on behavioral, descriptive or social traits that can be used for identification or verification purposes. Some examples of soft biometric data are height, hair color and style, clothing and accessories, age or behavioral patterns.

Soft biometric data could be particularly useful in surveillance systems, demographic research, online advertising and human-computer interaction. In video surveillance systems, the extraction of soft biometric data can be leveraged to classify attributes like gender or ethnicity, providing benefits in collecting forensic evidence and enhancing security measures. For instance, gender identification can serve as a preliminary step to expedite the search process within extensive databases. In other contexts gender classification is useful to enhance the performance of the advertising systems [7]. Another important field of application is human-computer interaction (HCI) in order to improve personalization, customization and fairness. For example, in an accessibility software gender classification can be used to adapt the interface and communication to be inclusive with regard to people outside of the binary gender representation. [8]

B. Bias in Gender Classification

Facial recognition systems and correlated gender classification systems slightly increased their diffusion in modern days and for this reason is crucial to consider the ethical impact of these systems and, consequently, it is fundamental to enhance the performance of these models when they're working on minorities. As highlighted in [9] the concept of *intersectionality* is pivotal to the understanding of social identities and oppression, particularly with respect to minority groups. Individuals belong to various social groups simultaneously, for example a person might identify as a woman and a person of color. It is mandatory to underline and understand that the systems of oppression and discrimination are interconnected, so people may face unique forms of discrimination when their multiple identities intersect. In order to limit the propagation of this discrimination and oppression, work aimed at minimizing algorithm bias is needed and our objective is to contribute to that purpose. To confirm the difficulties explained above the results shown in [10] are perfect. Darker females have the highest error rates for all gender classifiers ranging from 20.8% to 34.7%. Another confirmation is that accuracy of facial recognition systems used by law enforcement is lower for people labeled female and black [11].

C. "The Misgendering Machines"

Several studies seek to shed light on the importance of not perpetuating misgendering ([1], [2], [3]). The majority of Automatic Gender Recognition systems (AGR), as highlighted in [12], treat the gender as binary without questioning about the existence of genderqueer or trans people. This suggests that there is a lack of awareness on the topic among researchers that do not question the physiological and social basis of their technologies, perpetrating a bias that is already present in our social environment.

The technology can be made to work in a technically inclusive way, but it is difficult to analyze and understand whether this kind of technology can actually be considered ethical because gender, as used in AGR, is something assigned, and researchers should question the fact that it might be a non-consensual way of defining a individual's gender [12]. During our work we kept in consideration the analysis proposed in [12]. Having this kind of awareness led us to perform a deeper analysis of the final results, introducing an explainability model to analyze which features were the most important for the classification.

D. Related works

Despite the fact that [13] already performed a similar analysis, our focus is to extend this analysis presenting a comparative work of different computer vision models and expanding it with an explainability model. The results obtained by the explainability model are suitable for receiving valuable insights into how computer vision models work in the context of gender prediction, in order to avoid creating further bias in the future.

Figure 1. Sample of images from our dataset

III. The GINN Methodology

In the following section, we outline the pipeline adopted for our research described in Figure 3.

Our initial step involves the creation of a comprehensive dataset, bringing together random images from the CelebA dataset [14], balanced by gender, and a collection of 2000 images featuring non-binary individuals.

Subsequently, we undertake data preparation, which includes the division of our dataset into distinct subsets for training, testing, and validation. Furthermore, we implement data augmentation techniques exclusively on the training split, enhancing the robustness of our model.

To tailor a neural network to our specific task, we engage in fine-tuning a pre-trained model. This process leverages the knowledge and features acquired from a broader dataset, enhancing the model's performance on our gender-inclusive dataset.

Finally, we employ LIME (Local Interpretable Model-Agnostic Explanations)[15] for both prediction and evaluation. LIME allows us to scrutinize and interpret the model's predictions, shedding light on its decision-making processes and highlighting areas for improvement.

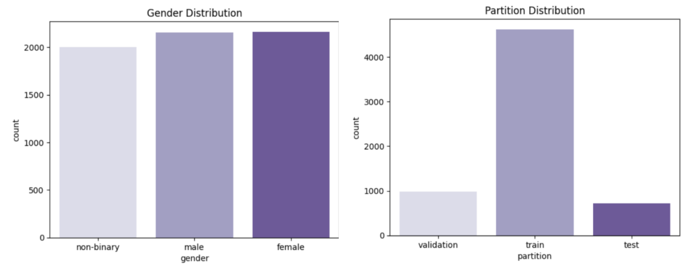

Figure 2. Composition of the used database

Throughout this section, we will delve into the details of each stage, elucidating the methodologies and techniques employed, and presenting the findings and insights obtained at each step of our pipeline.

A. Creation of Non-Binary Database

To extend the binary classifier, it is necessary to create a new database containing images of non-binary and genderqueer people. Wikipedia offers a list of people with non-binary gender identities [16], that are all public figures. The collection of images was performed by [13] and is publicly available for other researchers on GitHub. Our database comprises 2,000 images, each associated with one of 67 distinct identities. To ensure a balanced representation of gender, we augmented this dataset by adding 4,000 images, evenly distributed between males and females, all of different individuals, randomly selected from the Celeb-A database [14]. This augmentation resulted in a collection of 6,000 diverse images, as we can see in Figure 2

As we can see in Figure 1 the images are portraits and close-ups of celebrities in various contexts.

Given that the *non-binary* category is represented by only 67 unique individuals, we partitioned our database into two subsets. The first subset, used for training and validation, includes images from 55 distinct individuals, while the second subset, suited for testing, contains 12 individuals. We opted for this split between individuals because in order to have 80% of different individuals in the train and validation split and the remaining 20% in the test split. We also applied a similar division to the images sourced from Celeb-A to guarantee a balanced composition.

We utilized images of openly declared gender-diverse celebrities, but this approach has limitations due to the scarcity of available images and their limited diversity.

B. Computer Vision architectures

In 2012, the task of image recognition took a significant turn with the performance of a Convolutional Neural Network introduced by Krizhevsky, Sutskever and Hinton [17] in the ImageNet Recognition Challenge, bringing new interest in neural networks for gender prediction from images. Over the past decade, the task of image recognition has achieved considerable results, as numerous research teams have dedicated themselves to enhancing the performances of neural networks.

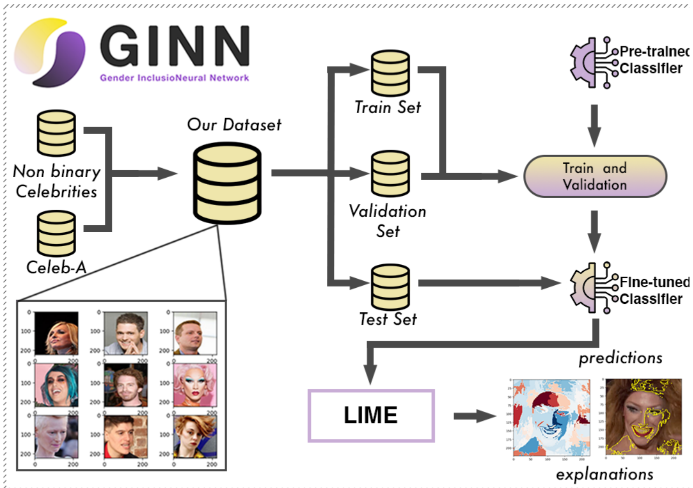

Figure 3. The GINN pipeline

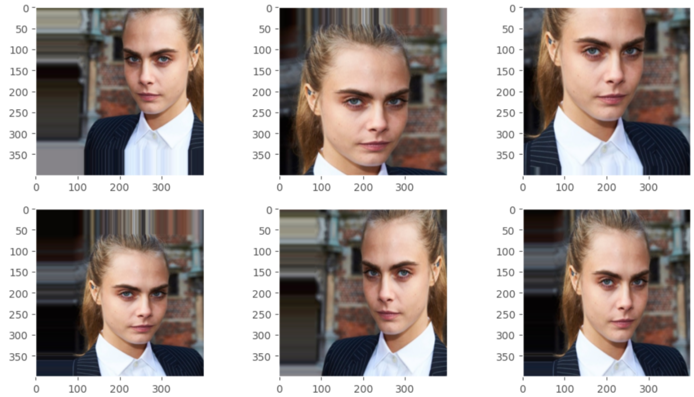

Figure 4. Examples of performed data augmentation

A crucial breakthrough came in 2015 with the introduction of ResNet [18]. This innovative neural network architecture made training deep networks a manageable task thanks to the idea of skipped layers, designed to skip connections and facilitate the direct flow of information across various network depths, leading to deeper neural networks and, consequently, to excellent results on the ImageNet dataset in various tasks. In the same year, another noteworthy development emerged: the Inception models [19]. Their proficiency in capturing multiclass scale features through the utilization of inception modules, this approach achieved remarkable accuracy across a range of image datasets.

In 2020, the EfficientNet [20] was introduced as a series of neural network architectures with the aim of enhancing both model accuracy and computational efficiency. This achievement is realized through a compound scaling approach that methodically enhances the model's depth, width, and resolution. By carefully balancing these dimensions, EfficientNet produces models that are highly efficient and accurate.

MobileNetV2[20] is another interesting benchmark, because it is a lightweight and efficient deep neural network architecture designed for mobile and embedded devices, specifically tailored for tasks like image classification and object detection. It utilizes depthwise separable convolutions and inverted residual blocks to reduce computational demands while maintaining high accuracy, making it well-suited for real-time applications with limited computational resources.

C. Transfer Learning Pipeline

Given the size of our database, we opted to use the transfer learning technique. It consists in applying a model pre-trained on a large-scale database (in our case, ImageNet[21]) or pre-trained for a different task to a different but related task. The objectives of this operation are better results and faster training. We selected multiple models for image classification and we used the same transfer learning pipeline to obtain comparable results. The selected models are: ResNet50[18], InceptionV3[19], MobileNetV2[20] and EfficientNet[22].

To enhance the performances of the training phase we pre-processed the input data using different operations, including rescaling the pixel values to a range of [0, 1] for numerical stability, applying a shear transformation with a range of 0.2, a zoom transformation with a range of 0.2, and horizontal flipping as we can see in 4. These transformations should contribute to a more robust and varied training dataset.

We employed transfer learning by initializing our deep learning model with the various architecture described before, pre-trained on the ImageNet dataset. To adapt the pre-trained model for our specific task, we freezed the last 5 layers, making them untrainable in order to let the model retain the general features learned from ImageNet. Subsequently, we extended the model with custom layers, including a flattening layer, followed by two fully connected layers.

The first fully connected layer comprises 1024 neurons with ReLU activation, serving as a feature extractor, and the final layer consists of 3 neurons with softmax activation, that fit with our specific task, that consists in classifying images into three different categories.

Table I
HYPERPARAMETERS FOR TRANSFER LEARNING

Parameter	**Value**
Epochs	5
Batch Size	64
Optimizer	adam
Loss	categorical crossentropy
Image Size	(224, 224)

We performed the experimental assessment using stratified K-Fold cross-validation technique, in order to ensure that the model's performance evaluation is robust and less dependent on a specific data split. This is crucial because of the composition of our database.

The training of the various models was performed using the hyperparameters shown in table I. We opted for *adam* as optimizer, because it is known for its effectiveness in optimizing deep neural networks and helps to achieve faster convergence, while *categorical crossentropy* is a commonly loss function employed for multiclass classification tasks.

The idea behind the choice of the small number of epochs and the image size, however, lies in the trade-off between computational speed and accuracy.

IV. Comparative analysis

In this section, we present the findings of our analysis, highlighting the significant conclusion and the limitation of this kind of architecture.

Our study aimed to perform the task of Automatic Gender Recognition, expanding the task from a binary problem, as is too often addressed, to a 3-class classification problem, namely *non-binary*, *male* and *female*. The objective is to observe how different models perform in terms of results evaluating it with quantitative indices (see section IV-B) and then, through the use of explainable AI techniques, understand how different models learn to classify (see section IV-C)

A. Results

In this section, we present and discuss the outcomes of our study. We address two primary objectives. First, we provide an in-depth analysis of the predictive performance of the models employed for gender recognition, utilizing quantitative indices, such as accuracy and F1-score, as metrics to evaluate their capabilities.

Subsequently, we focus on the explainable AI techniques applied to gain insights into the specific regions or features within the images that exert the most significant influence on the classification process. Through this comprehensive examination of both model performance and interpretability, we aim to offer a holistic view of the effectiveness of gender recognition models while elucidating the visual cues that play a pivotal role in the classification decision.

B. Quantitative Analysis

It is interesting to see how different models performed in such a different way. The results obtained by the model are shown in table II. We used two different metrics to evaluate the performance of the model. The first one is accuracy, that is a metric used to evaluate the overall performance of a classification model. It calculates the ratio of correct predictions to the total number of predictions made by the model. We also used F1-Score, defined as:

$$F_1 = \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

to evaluate the classification of the various classes. It is a metric that combines precision and recall to provide a more balanced assessment of a model's performance, particularly in cases of class imbalance. Precision represents the ratio of true positives to all positive predictions, while recall (sensitivity) represents the ratio of true positives to all actual positives in the dataset, because in the evaluation of the classes we had a lot of variance, so we wanted to maintain a balance between false positives and false negatives. We can observe how *InceptionV3* and *EfficientNet* outperformed the others.

EfficientNet comes out with the best results, a promising 94.4% of accuracy on our dataset.

Table II
Accuracy of the different models

	ACCURACY	**F1 SCORE**		
MODELS	**Overall**	**Male**	**Female**	**Non Binary**
ResNet50	68.8%	0.47	0.70	0.84
MobileNetV2	64.6%	0.48	0.71	0.68
InceptionV3	92.2%	0.96	0.91	0.90
EfficientNetB0	94.4%	0.92	0.93	0.99

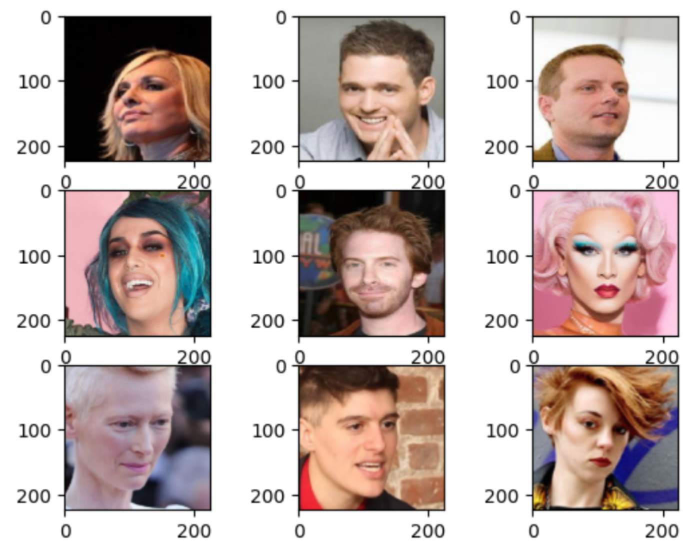

Figure 5. Samples of missclassified images

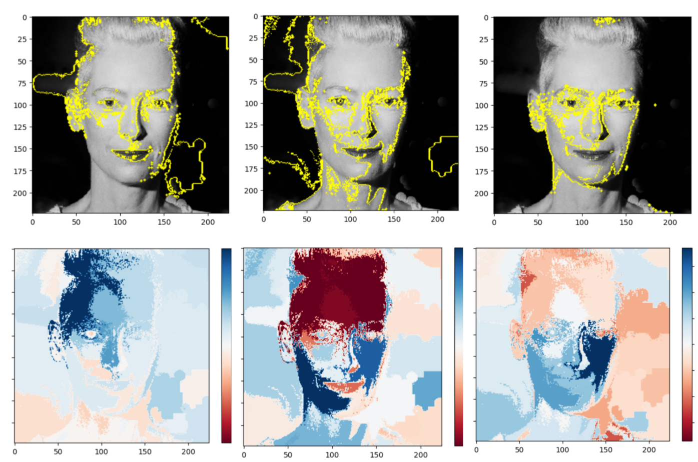

Figure 6. From left to right we can see how ResNet, InceptionV3 and EfficientNet perform classification. The heatmap show to us how InceptionV3 and EfficientNet work in a similar way, instead ResNet keep in consideration different features. The bluest areas in the heatmap are the more used to classify images.

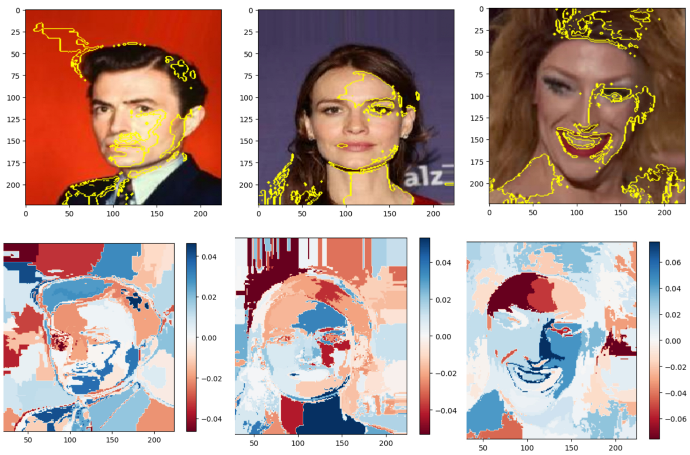

Figure 7. From left to right we can see how EfficientNet perform classification on male, female and non-binary individuals. The bluest areas in the heatmap are the more used to classify images.

C. Qualitative Analysis

To comprehend the fundamental distinctions among the models, we employed an explainability AI algorithm, known as Lime [15]. This was done to gain insights into the rationale behind the classification of images, enabling us to discern why a particular classification is assigned and identify the facial features and components that carry greater significance in the classification process. The heatmap depicted in Figure 6 serves as an entry point for comprehending how different features contribute to the classification of a facial image by various models. This visualization provides valuable insights into the areas of the face that these models prioritize when making their judgments.

In particular, let's examine the ResNet model, which is represented in the first image on the left. It is evident that ResNet places significant importance on the region of the forehead and the initial portion of the hair.

Conversely, when we turn our attention to the InceptionV3 and EfficientNet models, which happen to be our most accurate models, a distinct pattern emerges. These models exhibit little to no emphasis on the forehead. Instead, their focus is predominantly directed towards the cheeks, cheekbones, eyes, nose, and lips. These facial features, which intuitively serve as more descriptive regions of the face, appear to be the primary points of interest for InceptionV3 and EfficientNet.

In essence, this comparison underscores how different deep learning models prioritize and leverage distinct facial regions for accurate classification. While ResNet leans towards the forehead and hair region, InceptionV3 and EfficientNet appear to rely on the more expressive areas of the face, such as the eyes, nose, and lips, to enhance their classification capabilities.

The Figure 7, in a similar way, depicted the differences regarding the areas of the face used in the classification of male, female and non-binary. It is interesting because the heatmap for the male gender category highlights the significance of the jawline, forehead, and cheekbones. These features appear to be key determinants in the gender recognition process for males. In the female category, the heatmap indicates that the under-eyes region, skin of the forehead, cheeks and neck play pivotal roles in gender classification. For the non-binary category, the heatmap reveals a more interesting pattern. It suggests that gender recognition in non-binary individuals relies on a combination of features without a strong emphasis on traditional gender-specific characteristics. The eyes and lips, in particular, stand out as important areas.

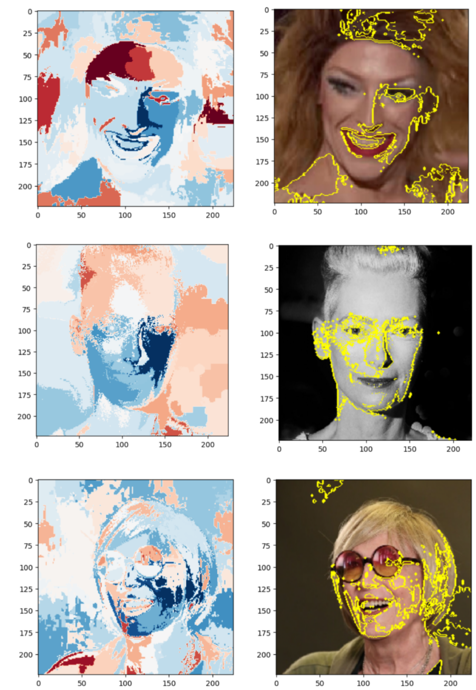

Figure 8. Non-Binary individuals well classified by EfficientNet, but with different recognition patterns

As we can see more precisely in Figure 8 we can suppose that make-up lips and make-up eyes could influence the model recognition of non-binary people. The interpretation of these heatmaps may rely on the concept of gender as *performative act* [5]. Judith Butler redefines traditional notions of gender by proposing that is not an inherent or fixed attribute but a social construct that individuals continually enact and express through their actions, language, and behavior. This theory suggests that people perform their gender roles and that these performances are integral to the construction of one's gender identity. Starting from the definition of gender as an historical situation, theorized by Simone de Beauvoir [23], she emphasizes the role of societal norms in reinforcing or challenging gender expectations and allows a more fluid understanding of gender, transcending binary categories. This concept has had a profound influence on queer theory and discussions of non-binary and transgender identities, offering a more inclusive and dynamic perspective on the nature of gender. The same concept is carried forward by West and Zimmerman in their work *"Doing Gender"* [24]. In a similar way they challenge the notion of gender as a fixed and inherent characteristic, emphasizing that it is a social process enacted and reinforced through everyday interactions. Understanding gender in this way partly explains why certain types of facial features, not necessarily related to biological sex, are very important in trying to understand a person's gender.

D. Discussion

Although EfficientNet has shown promising results on our dataset, it is essential to shed light on the limitations of this approach and the general constraints of this technology when aiming for fairness and inclusivity in real-world applications.

The most important problem relates to data availability and diversity. The lack of data collected on individuals whose gender identity does not conform to a binary framework is a critical problem. In our test set, the average cosine similarity between pairs of images of non-binary individuals is 0.69, while the average cosine similarity between other images in the test set is 0.61. This discrepancy arises from the fact that the test set contains only 12 different non-binary individuals. Relying on the results of such a small and insufficiently varied data set poses a risk to the generalizability and practical relevance of the results.

The Misgendering Machines section of this article underscores the potential pitfalls associated with these technologies. Attempting to determine a person's gender identity solely from an image is a complex task fraught with challenges and risks.

One major ethical concern revolves around the potential for misgendering individuals. Misgendering occurs when a machine incorrectly attributes a gender identity to a person based on facial features or other visual cues. This misclassification can be not only inaccurate but also deeply impactful for the individuals involved. Being misgendered by a machine has real-world consequences, as it can contribute to the perpetuation of gender stereotypes and reinforce biases that already exist in society. [12]

Moreover, the act of misgendering can have emotional and psychological effects on individuals. It may lead to feelings of frustration, discomfort, or alienation, particularly for those whose gender identity may not align with societal expectations or norms. In the broader context, the incorrect use of such technology has the potential to erode trust in automated systems and artificial intelligence, especially if these systems consistently make inaccurate gender attributions.

The ethical implications extend to the responsibility of technology users. It is crucial for individuals and organizations to be aware of the limitations and potential biases of gender inference technologies. Users should exercise caution and deploy such technologies only when absolutely necessary, ensuring that their use aligns with principles of fairness and respect for individuals' rights. Additionally, developers and policymakers need to prioritize the implementation of safeguards and regulations to mitigate the risks associated with misgendering and to promote the responsible use of gender inference technologies.

In conclusion, the ethical considerations surrounding the inference of gender identity from images highlight the need for a thoughtful and cautious approach. As technology continues to advance, it is imperative that developers, users, and policymakers collaborate to address these ethical concerns, foster transparency, and ensure that these technologies are deployed in a manner that respects the dignity and rights of individuals.

V. Future Works

In the future, we would like to extend the comparative analysis evaluating Vision Transformer, a class of deep learning models that apply the Transformer architecture, originally designed for natural language processing, to computer vision tasks, allowing for end-to-end image understanding and classification. In this way we will be able to have a more comprehensive view of the various computer vision techniques, understanding how they work in practice with non-binary individuals.

Another main goal will be to move to a broader perspective of gender recognition and to recognize gender as a continuous spectrum. The concept of the gender continuum, as discussed in "Gender Continuum" [25] recognizes that individuals everywhere can identify themselves in a diverse spectrum of gender expressions, transcending the limitations of the binary model. Such characteristics include gender identity, sexual orientation, scalable personality traits such as assertiveness, inquisitiveness, empathy, and kindness, the observance of culturally defined behaviours such as dress, social interaction, and social roles and the physical traits associated with biological sex—including external genitalia, sex chromosomes, sex-related gene expression, sex hormones, and brain structure and activity—which vary along a continuum ranging from male to intersex to female, as well explained in Encyclopedia Britannica [26]. Categorical systems for recognizing gender are inherently unable to accurately represent the richness of gender diversity, risking misclassification and erasure of non-binary and gender-diverse individuals.

A potentially more ethical approach to the use of automatic gender recognition systems might be to classify people along a continuum that reflects the social spectrum between masculinity and femininity, rather than placing them in a few rigid categories. This approach, while not without its problems, could be of particular interest for purposes such as advertising, as it allows for a more nuanced and accurate representation of people's gender identities.

References

1. Monro, S. Non-binary and genderqueer: An overview of the field. *International Journal of Transgenderism* **20,** 126–131 (2019).
2. Taylor Zalewska, G. An exploration of the lived experiences of non-binary individuals who have presented at a gender identity clinic in the United Kingdom. *The international journal of transgenderism,* 1195–204 (2018).

3. Harrison Grant, H. A Gender Not Listed Here: Genderqueers, Gender Rebels, and OtherWise in the National Transgender Discrimination Survey. *LGBTQ Public Policy Journal at the Harvard Kennedy School* (2012).
4. Pareek, V. Non-binary Gender and Data. https://cis.pubpub.org/pub/non-binary-gender-data (Dec. 2019).
5. Butler, J. Performative Acts and Gender Constitution: An Essay in Phenomenology and Feminist Theory. *Theatre Journal* **40,** 519–531. ISSN: 01922882, 1086332X. http://www.jstor.org/stable/3207893 (2023) (1988).
6. Sara Bragg Emma Renold, J. R. & Jackson, C. 'More than boy, girl, male, female': exploring young people's views on gender diversity within and beyond school contexts. *Sex Education* **18,** 420–434 (2018).
7. Wu, Y., Zhuang, Y., Long, X., Lin, F. & Xu, W. *Human Gender Classification: A Review* 2016. arXiv: 1507.05122 [cs.AI].
8. Khan, S., Ahmad, M., Nazir, M. & Riaz, N. A comparative analysis of gender classification techniques. *Middle - East Journal of Scientific Research* **20,** 1–13 (Jan. 2014).
9. Crenshaw, K. "Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics". *University of Chicago Legal Forum* (1989).
10. Buolamwini, J. & Gebru, T. *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification* in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (eds Friedler, S. A. & Wilson, C.) **81** (PMLR, 23–24 Feb 2018), 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html.
11. Klare, B., Burge, M., Klontz, J., Vorder Bruegge, R. & Jain, A. Face Recognition Performance: Role of Demographic Information. *Information Forensics and Security, IEEE Transactions on* **7,** 1789–1801 (Dec. 2012).
12. Keyes, O. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proceedings of the ACM on Human-Computer Interaction* **2,** 1–22 (Nov. 2018).
13. Wu, W., Protopapas, P., Yang, Z. & Michalatos, P. Gender Classification and Bias Mitigation in Facial Images. *CoRR* **abs/2007.06141.** arXiv: 2007.06141. https://arxiv.org/abs/2007.06141 (2020).
14. Liu, Z., Luo, P., Wang, X. & Tang, X. *Deep Learning Face Attributes in the Wild* in *Proceedings of International Conference on Computer Vision (ICCV)* (Dec. 2015).
15. Ribeiro, M. T., Singh, S. & Guestrin, C. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. *CoRR* **abs/1602.04938.** arXiv: 1602.04938. http://arxiv.org/abs/1602.04938 (2016).
16. Wikipedia contributors. *List of people with non-binary gender identities — Wikipedia, The Free Encyclopedia* [Online; accessed 17-October-2023]. 2023. https://en.wikipedia.org/w/index.php?title=List_of_people_with_non-binary_gender_identities&oldid=1180070054.
17. Krizhevsky, A., Sutskever, I. & Hinton, G. E. *ImageNet Classification with Deep Convolutional Neural Networks* in *Advances in Neural Information Processing Systems* (eds Pereira, F., Burges, C., Bottou, L. & Weinberger, K.) **25** (Curran Associates, Inc., 2012). https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.
18. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *CoRR* **abs/1512.03385.** arXiv: 1512.03385. http://arxiv.org/abs/1512.03385 (2015).
19. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the Inception Architecture for Computer Vision. *CoRR* **abs/1512.00567.** arXiv: 1512.00567. http://arxiv.org/abs/1512.00567 (2015).
20. Howard, A. G. *et al.* MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *CoRR* **abs/1704.04861.** arXiv: 1704.04861. http://arxiv.org/abs/1704.04861 (2017).
21. Deng, J. *et al. ImageNet: A large-scale hierarchical image database* in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009), 248–255.
22. Tan, M. & Le, Q. *EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks* in *Proceedings of the 36th International Conference on Machine Learning* (eds Chaudhuri, K. & Salakhutdinov, R.) **97** (PMLR, Sept. 2019), 6105–6114. https://proceedings.mlr.press/v97/tan19a.html.
23. De Beauvoir, S., Borde, C. & Malovany-Chevallier, S. *The Second Sex* ISBN: 9780307277787. https://books.google.it/books?id=_hywlrNuYvIC (Knopf Doubleday Publishing Group, 2011).
24. West, C. & Zimmerman, D. H. Doing gender. *Gender Society* **1,** 125–151. http://journals.sagepub.com/doi/10.1177/0891243287001002002 (1987).
25. Gülgöz, S., Edwards, D. & Olson, K. Between a boy and a girl: Measuring gender identity on a continuum. English (US). *Social Development* **31,** 916–929. ISSN: 0961-205X (Aug. 2022).
26. Brian, D. gender continuum. English (US). *Encyclopedia Britannica.* https://www.britannica.com/topic/gender-continuum (Oct. 2023).