Chapter 32
Examining Diversity and Fairness Interconnection in Data-Driven Personas

Reham Al Tamime, Joni Salminen, Soon-Gyo Jung, and Bernard Jansen

Abstract As organizations seek to reach and engage diverse social media audiences, creating personas that accurately represent this diversity is becoming progressively important. While creating more personas generally increases the representation of demographic diversity, little is known about the dynamics between persona set size, diversity, and fairness. To study this, we collect user statistics from three organizations' YouTube and Facebook channels and generate persona sets with varying size. Analyzing the personas' gender, age, and country, we find that increasing the persona set size correlates with diversity in gender, age, and country. Increasing the number of personas also correlates with age fairness, while the level of diversity correlates with fairness in gender, age, and country. The results imply that organizations would benefit from using data-driven approaches to generate many personas in order to strengthen the outlooks of diversity and fairness, promoting a broader targeting strategies and inclusive design in systems and artifacts.

Keywords Inclusive design · Persona set size assessment · Diversity and fairness trade-offs · Human-computer interaction

32.1 Introduction

Personas are a human-centered design technique introduced by Cooper (1999) for human-computer interaction (HCI) and other fields that design or develop products for end users. Personas are typically fictional characters that might include a name, photo, background story, basic demographics, personal attributes,

R. Al Tamime (✉) · S.-G. Jung · B. Jansen
Qatar Computing Research Institute, Doha, Qatar
e-mail: realtamime@hbku.edu.qa; sjung@hbku.edu.qa; bjansen@hbku.edu.qa

J. Salminen
University of Vaasa, Vaasa, Finland
e-mail: joni.salminen@uwasa.fi

A. Nagar et al. (eds.), *Intelligent Sustainable Systems*, Lecture Notes in Networks and Systems 1179, https://doi.org/10.1007/978-981-97-9327-3_32

general goals, needs, motivations, and challenges or pain points (Nielsen et al. 2013). Using personas for representing social media users involves several benefits, including differentiating between user groups who use products and platforms, gaining empathy for the user groups whom the personas represent, and improving communication between team members working on design tasks such as social media content creation, targeting, and analyzing engagement results (Meissner and Blake 2011; Jansen et al. 2020, 2021). Personas are used to create artifacts and systems that suit the interests of different people with different use cases (Cutting and Hedenborg 2019).

With the proliferation of Web and social media user analytics, data on social media user populations, such as audiences on YouTube, Twitter, and Facebook, brings about opportunities for creating "data-driven personas." Data-driven personas, generated using data science algorithms and models, represent different audience segments based on factors such as demographics and behaviors (Jansen et al. 2021; An et al. 2018). Creating data-driven personas offers more advantages including enabling the creation of personas from large datasets that are more representative of the user population and easy to adapt to the user needs and preferences (Jansen et al. 2020).

Despite personas involving many benefits for design, shortcomings can arise. In particular, personas can be found disconnected from the user segment due to a lack of consideration for diverse user needs. In some circumstances, personas can introduce stereotypical thinking and reinforce existing biases because of the lack of representation of specific kinds of users (Turner and Turner 2011). Systems and artifacts can inherit harmful biases and stereotypes, thereby lacking inclusivity in different dimensions such as gender. This calls for supporting inclusive design that embraces both the similarities and differences in the user base (Hoque et al. 2022). As digital inclusion is becoming more important and global consumer products and engineered systems adopt increasingly digital interfaces, there is increased attention to addressing these shortcomings (Salminen et al. 2022c; Hoque et al. 2022).

The concept of *diversity* refers to the representation of different demographic groups in the output. For example, a diverse persona set would ideally contain 50% male and 50% female personas (when assuming two genders). This could be understood as different demographic groups' *opportunity* to be included in the generated persona set. There are use cases where maximally diverse user representations might be sought to understand the existing user base of a channel or system, especially in inclusive design (Goodman-Deane et al. 2018) whose goal is to take different groups into consideration.

Fairness refers to the correspondence of the generated data-driven personas with the demographically aggregated source data (Salminen et al. 2022a). For example, if the baseline data about the audience of a given social media channel indicates that 80% of the channel's audience is male, then a fair persona set would also contain 80% male personas (and 20% of personas belonging to some other gender). The fairness described in this example is called "representative fairness" (Abbasi et al. 2021). It ensures that the distribution of output, such as the selection of personas, reflects the distribution of the population or baseline data. In this case, the proportion of male personas in the set is equal to the proportion of males in the channel's audience, maintaining the demographic balance.

It is, therefore, clear to see that optimizing for these two goals – fairness (based on truthfulness to the baseline data) and diversity (based on the equal opportunity of each demographic group to be included) – might impose a trade-off for data-driven persona generation, particularly for social media channels that have a demographically imbalanced audience. Especially given that such channels would benefit from a degree of diversity, as understanding their minority users might result in valuable insights for broadening the overall audience base.

Accordingly, the problem we address is how the dimensions of fairness and diversity in data-driven persona creation relate to each other. The relationship between diversity and fairness has received increased attention in the field of personalization and recommendation (Schelenz 2021; Grgić-Hlača et al. 2022). Previous work has mainly focused on understanding fairness in terms of justice or the absence of discrimination. For example, presenting diverse sources and representations of race and gender will increase fairness toward populations that have previously been rendered invisible (Schelenz 2021)—but the focus has not been on automatically generated personas. So, the connection between diversity and fairness has not been explored in the field of design, HCI, and persona generation.

This work addresses this research gap by studying the interplay between diversity and fairness in data-driven persona creation. More specifically, limited literature explores whether increasing diversity leads to diminishing or increasing fairness in data-driven persona creation. We focus on understanding diversity in terms of *actually* representing the various demographic characteristics of the baseline data in the output by capturing both the number of different items present in the baseline data and commonness or rarity of each item. We focus on understanding fairness in terms of the truthfulness of demographics to the baseline data. Subsequently, we designed an experiment with varying the number of personas and investigated whether creating more personas helps in increasing demographic fairness or diversity. This is an important problem because it directly affects whether persona creators who use data science algorithms would need to deliberately choose between achieving diversity and fairness or whether they can achieve both. In other words, designers are typically obliged to select whether they prefer more diversity or fairness when generating personas related to their project. However, whether they can be granted diverse and fair sets of personas is obscure. More profoundly, it is also valuable to know the minimum and the maximum number of personas (threshold) that should be generated to have diverse and fair sets of personas. With increasing calls for diversity in the design, it is crucial to understand which dimensions of diversity are associated with fairness.

More specifically, we aim to investigate whether creating more personas helps in increasing demographic fairness or diversity. Accordingly, this research aims to understand the relationship between (1) the number of personas and diversity, (2) the number of personas and fairness, and (3) diversity and fairness trade-offs in the context of data-driven personas. We specifically focus on the following research questions (RQs):

– RQ1: What is the relationship between the number of personas and diversity for gender, age, and country?
– RQ2: What is the relationship between the number of personas and fairness for gender, age, and country?
– RQ3: What is the relationship between diversity and fairness for gender, age, and country?

Addressing these RQs, we make the following empirical contributions. First, we provide evidence that increasing the number of personas is valuable to increasing the diversity in gender, age, and country at specific threshold levels. Second, we provide evidence that increasing the number of personas is valuable to increasing the fairness in age at specific threshold levels. Third, we show that achieving a high level of diversity does not compromise fairness, as data-driven persona generation can harmoniously preserve both.

In the remainder of this chapter, we introduce the background and related work on personas and demographic representation in gender, age, and location. We then explain the research design and outline the results. After that, we report the results by each demographic attribute, organization, and social media channel. Finally, we discuss the answers to our research questions and their implications and suggest potential directions for future work.

32.2 Literature Review

The literature review covers background information about persona, data-driven personas, demographical attributes, and diversity in HCI and persona creation. We focus on these three demographic attributes (gender, age, and geographical location) because they are collected on nearly all online platforms and are common in almost all personas. Besides, these three demographic attributes typically serve as "skeletal personas" (Zhu et al. 2019) and are vital in persona development. Therefore, the results are impactful for a wide range of persona research.

32.2.1 The Development of Data-Driven Personas

Incorporating personas into user-centered design (UCD) practices has become increasingly popular. A persona is a fictional character that represents a group of target users who share common needs, characteristics, and goals (Miaskiewicz and Luxmoore 2017). The fictional character is typically represented using a picture, name, and other type of information depending on each use case (Miaskiewicz and Luxmoore 2017; Salminen et al. 2022d). Figure 32.1 shows an example of a data-driven persona. The several advantages of incorporating personas into design practices have been largely cited. First, personas can be used as a tool to

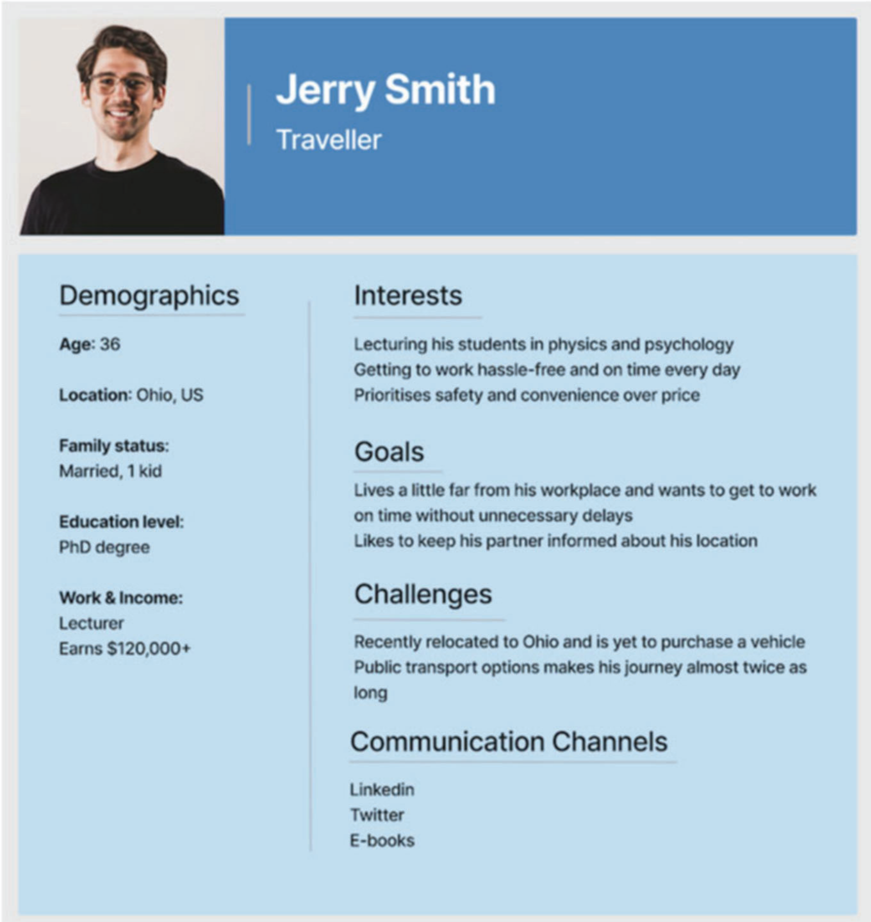

Fig. 32.1 An example of a persona that includes a picture, demographical information, and other information. The persona has been created by manually collecting information about the user

facilitate communication about the target users within the design team and other stakeholders (Miaskiewicz and Luxmoore 2017). Second, personas can help to focus on the needs of the target users when making design and development decisions (Miaskiewicz and Luxmoore 2017). Personas have been incorporated into many fields including content creation, marketing, product design, ergonomics, health, and software development (Jansen et al. 2021). Personas are a widely established interaction design method intended to make design processes more user-centric and are used by researchers and large businesses (Cutting and Hedenborg 2019). Despite the wide adoption of persona, persona generation has been criticized in the literature for various reasons including lack of scalability (i.e., difficulty in creating personas on a large user population datasets), and nonrepresentative data

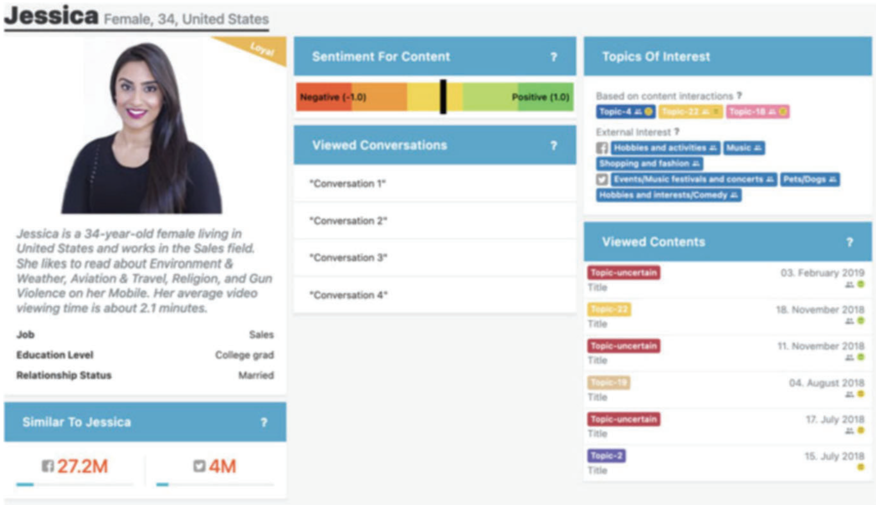

Fig. 32.2 An example of a data-driven persona that includes a picture, demographical information, and other information. The persona has been created by automatically collecting social media data about the user

(i.e., creating personas that do not represent the whole user base or reflect the variations within the datasets) (Jansen et al. 2021).

The proliferation of online data, social media platforms, and data science tools and algorithms enabled the creation of a revived kind of persona called: data-driven persona (Jansen et al. 2021). Data-driven personas are personas that are created from persona templates using quantitative data about a given user population and derived using statistical techniques, including data science, and machine learning algorithms (Salminen et al. 2020b). Data-driven personas rely on automatic data collection and data-driven analysis methods to create and update personas rapidly. Figure 32.2 shows an example of a data-driven persona. Data-driven personas combine the advantages and offset the shortcomings of personas that are manually created using tools and algorithms that are not specifically related to data science (Jansen et al. 2021). These advantages consist of making it possible to create personas from large datasets that are easy to change and adapt to the user needs and preferences (Jansen et al. 2020, 2021). Data-driven personas have been applied to enrich several design tasks. For example, Molenaar (2017) analyzed 400,000 clickstreams collected over a period of three months and then grouped and classified them into common workflows. Similarly, Zhang et al. (2016) analyzed clickstream data and generated five data-driven personas. Creating data-driven personas from clickstreams has been proven to be a scalable way to create personas with little interference from human interviewers or noisy self-reports (Zhang et al. 2016).

Data-driven personas can be created using different methodologies including affinity diagrams, decision trees, exploratory factor analysis (EFA), hierarchical clustering, k-means clustering, latent semantic analysis (LSA), multidimensional scaling analysis (MSA), and weighted graphs (Jansen et al. 2021). A systematic review of quantitative persona creation identified five general approaches for creating data-driven personas (Jansen et al. 2021). These are cluster analysis (methods that consist of grouping a dataset using a predetermined number of clusters k-means and hierarchical clustering), latent Dirichlet allocation (statistical model that models data as a finite mixture over an underlying set of patterns), latent semantic analysis (data analysis algorithm that uses singular value decomposition to detect hidden semantic relationships between words), principal component analysis (linear dimension-reduction algorithm used to extract information by removing nonessential elements with a relatively small variation), and nonnegative matrix factorization (NMF; matrix factorization method in which matrices are constrained as nonnegative and decomposed to extract sparse and meaningful features) (Jansen et al. 2021). Nonnegative matrix factorization has been applied widely as a method for creating data-driven personas, as it is applicable to user analytics datasets and scales from small to large datasets (Jansen et al. 2020; Salminen et al. 2022a; Jung et al. 2018). In the context of social media studies, personas have been evoked for multiple purposes to understand social media users including studies such as Netzorg et al. (2021), Marwick (2013), Jung et al. (2018), Allison and Bussey (2020), Torres-Lugo et al. (2022), and Kong et al. (2022).

On the other hand, challenges arise in the sense that personas can, in some circumstances, increase stereotypical thinking and reinforce existing biases about the channel's followers. This could result in detrimental outcomes, especially for marginalized audience segments that may in some cases be overlooked by the algorithm building the personas (Salminen et al. 2020c). In this regard, one of the main challenges in developing data-driven personas is that for specific attributes like age, gender, or location, persona generation algorithms tend to select only one dominant value (Salminen et al. 2020c). As such, it is important to consider solutions that help increase the variability of personas by an attribute of interest.

32.2.2 Gender in Personas

Gender has been investigated previously in HCI and persona creation (Vorvoreanu et al. 2019; Stumpf et al. 2020). The gender-inclusive design seeks to facilitate individuals' use of technologies and reflects on those possibilities toward more aware and inclusive practices (Breslin and Wadhwa 2018). There are different motivations for investigating gender in inclusive design, including economic, ethical, and political ones (Stumpf et al. 2020). Research around gender-inclusive design reflects on how gender is encoded and operationalized as well as on gendered stereotypes. HCI methods, such as GenderMag (Vorvoreanu et al. 2019; Shekhar and Marsden 2018), have been developed to detect gender biases in software.

Marsden et al. (2017) ensured that the female perspective is represented in the design process by developing female personas of software users. Different studies, such as Hill et al. (2017), found that designing personas with multiple pictures of males and females may contribute to expanding the persona users' understanding of the persona's diverse gender user segments rather than evoking gender stereotypes. Research on persona and gender stereotypes found that increased transparency of the female persona by providing users with clearly understandable explanations on how the information in the persona profiles are generated has a positive effect on completeness (i.e., how well the persona profile captures essential information about the users) and empathy (i.e., how well the participant relates to the persona) (Salminen et al. 2020a). Along these lines, Salminen et al. (2022b) assessed how to obtain a better representation of users' gender by increasing the number of personas.

32.2.3 Age in Personas

The element of age has been investigated previously in HCI and persona creation (Breslin and Wadhwa 2018). Developing technology that can cope well with age diversity and the related challenge of intergenerational codesign (Breslin and Wadhwa 2018). Research has focused on persona creation for children (Antle 2008) and seniors (Bowen et al. 2020; Wöckl et al. 2012). Furthermore, Salminen et al. (2022b) assessed how to better represent users' age by increasing the number of personas.

32.2.4 Geographical Locations or Country in Personas

Facing expanding digitalization over the globe, the interest in investigating the geographical locations, such as country, region, and area dimensions in HCI keeps rising (Himmelsbach et al. 2019). Anvari et al. (2019) developed and assessed personas with diverse countries and cultures for a large cohort of computing undergraduate students to understand the different approaches to learning specific subjects. Other studies found that designers emphasize many similarities among users across countries, but they still find it essential to illustrate cultural differences in the persona descriptions (Jansen et al. 2017). Therefore, international personas must be developed while keeping these similarities and differences in mind (Jansen et al. 2017). Research such as Salminen et al. (2022b) assessed how to better represent users' countries by increasing the number of personas.

32.2.5 Diversity and Inclusivity in Data-Driven Persona

Several authors have discussed the difficulty of creating a representative persona and risking creating a persona that nobody can relate to Fuglerud et al. (2020).

Research has touched on the issue of the need to collect and represent the needs of marginalized user groups to support the universal and inclusive design and consider challenges related to gaining deep insight into the deep experience (Fuglerud et al. 2020). In addition, research has encouraged considering the complexity of people's identities, such as gender, job roles, age, and other categories, to reduce biases and stereotypes (Marsden and Pröbster 2019). Thus, personas could become a tool that shapes design practices for diversity and opens new perspectives for empowerment (Marsden and Pröbster 2019). Research has argued that creating more personas is vital to present different end user types and offer several choices that allow users to select personas from a variety of options (Salminen et al. 2022c). Remarkably, Salminen et al. (2022c) found that age diversity, gender diversity, and country diversity are all higher when applying a larger number of personas. Accordingly, Salminen et al. (2022c) confirmed that applying 15 personas relative to 5 often leads to a higher demographic diversity of user representation. Simultaneously, Salminen et al. (2022b) acknowledged that creating more personas certainly improves the demographic diversity of the user segments and is a key to obtaining inclusivity in user representation (Salminen et al. 2022b). After running a user study, Salminen et al. (2022b) demonstrated that creating 40 personas (and not only 4 as previously advocated for) contributes to gain in relative diversity and larger coverage of demographic features, which is particularly important when dealing with large and heterogeneous users.

32.2.6 Research Gap

Beyond studies that focus on the relationship between the number of personas and diversity only (Salminen et al. 2022b), little is known about the relationship between the number of personas and fairness and the relationship between diversity and fairness. In this line, little is known about the minimum and the maximum number (thresholds) of the number of personas that could be generated to obtain a diverse set of personas, a fair set of personas, or both. Notably, recent studies have recognized the growing interdependence of diversity and fairness in the field of algorithmic design (Schelenz 2021; Grgić-Hlača et al. 2022), but the connection between diversity and fairness has been largely overlooked in the context of persona design, systems design, and HCI. Motivated by the call for inclusive and fair technology design overall, and specifically for personas (Goodman-Deane et al. 2018, 2021), we seek to fill this gap by examining three different organizations and two different social media channels to (1) find the relationship between the number of personas and diversity, the number of personas and fairness, and diversity and fairness and (2) identify the threshold that required to have a significant correlation between the number of personas and diversity, the number of personas and fairness, and diversity and fairness.

32.3 Methodology

This empirical analysis aims to reveal how diversity and fairness are affected when increasing the number of data-driven generated personas, particularly in relation to baseline data about a channel's social media users.

We increase the number of personas using a series with a multiplier of two: 5, 10, 20, 40, 80, and 160. Then, we added 5 to each element from 10 to 80, so we have a final series that consists of 5, 10, 15, 20, 25, 40, 45, 80, 85, and 160 personas. This interval is chosen because this enables us to generate personas from both broad (multiplier of two) and narrow range (addition of 5) of internal values. The minimum number of personas is five because this aligns with the number of personas often used in the literature (Salminen et al. 2022a,b).

32.3.1 Data Collection

We collected data about content engagement and users' demographics from three organizations that publish content on social media:

1. **Organization 1 (ORG01)**: International media company (publishes content focused on news, speech, discussions, and current affairs)
2. **Organization 2 (ORG02)**: Nonprofit organization focused on education and science (publishes content focused on public relations)
3. **Organization 3 (ORG03)**: Commercial airline company (publishes content focused on promotions, recruiting, and airplane information)

Although the organizations vary by the number of subscribers and the type of content they produce, the common denominator is that all of them publish content for a large, heterogeneous online audience. Therefore, they are classified as *internationally oriented online content publishers*. These three organizations target audiences of all genders, age groups, and countries.

For each organization, we collected data from two social media platforms (YouTube and Facebook). The data was collected via application programming interfaces (APIs) with the organizations' permission.

The data does not contain personally identifiable information because online platforms commonly aggregate this data to protect the privacy of individual users. For example, this aggregated data tends to include information about gender (e.g., female),[1] age groups (e.g., 25–34),[2] and country[3] (e.g., USA). For each dataset, we combined the demographic groups (location-age-gender) and aggregated the

[1] Please note that we consider the self-reported gender provided by the user (mostly has been reported as binary: female or male).

[2] Please note that we consider the self-reported age provided by the user.

[3] Please note that we consider the self-reported country provided by the user.

Table 32.1 Description of the baseline datasets: organizations-channels (column 1) is the organization number and social media channel, data collection dates (column 2) is the start and end data collection dates, content engagement (column 3) is the total number of video views, and content items (column 4) is the number of unique videos, demographic groups (column 5) are how many unique combinations of (location-age-gender) exists in the datasets

Organizations-channels	Collection dates	Content engagement	Content items	Demographic groups
ORG01 YouTube	Jul 2016–Oct 2022	911.23M	15,470	2491
ORG01-Facebook	Nov 2016–Oct 2022	32,885,732.37M	8045	96
ORG02-YouTube	Jan 2017–Oct 2022	8.37M	60	203
ORG02-Facebook	Jan 2017–Oct 2022	574,210.75M	980	2184
ORG03-YouTube	Oct 2017–Oct 2022	323.10M	793	1683
ORG03-Facebook	Oct 2017–Sep 2022	13,909,402.67M	1072	1974

content's engagement (number of views per content) for each group. For example, the demographic group [USA, 45–54, Female] can have [1000] instances of content engagement or views for specific content [Video X].

For each of the six datasets, we created an interaction matrix or a pivot table that captures the count of content engagement per demographic group. Table 32.1 describes each dataset in terms of data collection dates, content engagement, the number of unique content items, and the number of demographic groups. Figure 32.3 shows the gender distribution of the baseline data from YouTube and Facebook, while Fig. 32.4 shows the age distribution of the baseline data from YouTube and Facebook. The engagements for all the organizations and channels are dominated mainly by males. We think we got a primary male audience in our sampling because a larger portion of Youtube users worldwide are male.[4] Besides, the engagement mostly comes from 25 to 34 years old, followed by a 35 to 44 audience.

For ORG01, the average engagement on YouTube is 869 engagements. Overall, 77% of demographic groups have engagement above the mean, while 23% are below the mean. Moreover, the average engagement per Facebook content is = 31,362,504.

For ORG02, the average engagement on YouTube is 7,857 engagements. We observe that 57% of demographic groups have engagement above the mean, while there are 43% below the mean. Additionally, the average engagement per Facebook content is 1,900,039. Also, 24% of demographic groups have engagement above the mean, while there are 76% below the mean.

For ORG03, the average engagement on YouTube is 7,378 engagements. Additionally, 56% of demographic groups have engagement above the mean, while there are 44% below the mean. Furthermore, the average engagement per Facebook content is 45,614,919. One notes that 59% of demographic groups have engagement above the mean, while there are 41% below the mean. It is noticeable that the average engagement on Facebook is higher than on YouTube in all organizations.

[4] https://www.statista.com/statistics/1287032/distribution-youtube-users-gender/

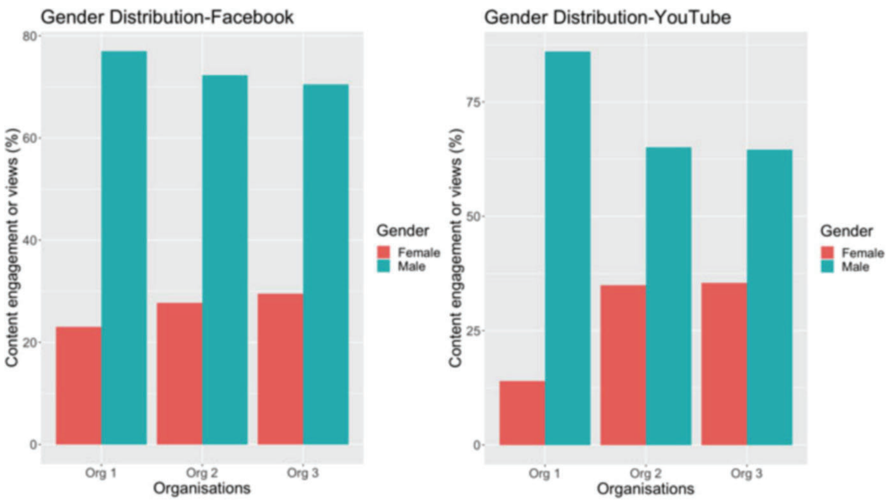

Fig. 32.3 Gender distributions of the baseline data from YouTube and Facebook. The Y-axis shows the engagement (video views) percentage, and the X-axis shows the organizations and gender. The engagements for all the organizations and channels are dominated mainly by males

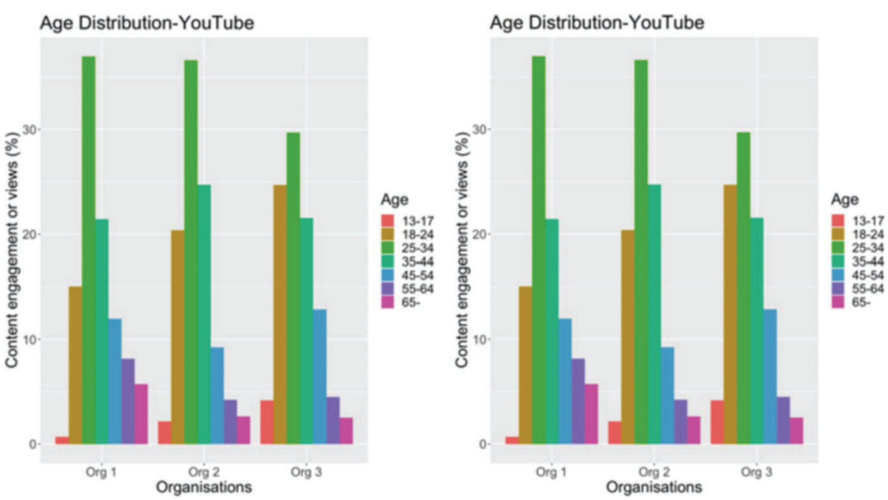

Fig. 32.4 Age Distributions of the Baseline Data from YouTube and Facebook. The Y-axis shows the engagement (views) percentage, and the X-axis shows the organizations and age groups. The engagements for all the organizations and channels are dominated mainly by 25–34 years old audience which, according to our experience, tends to be the case for many social media channels

32.3.2 Persona Generation

We used the method described by An et al. (2018) to generate personas automatically from social media data. Using the interaction matrix as the input, the nonnegative matrix factorization (NMF) algorithm (An et al. 2018) is applied to detect the latent patterns representing the user segment's digital content preferences. We selected NMF because it has been applied widely as a method for creating data-driven personas, and it is applicable to user analytics and both small to large datasets (Jansen et al. 2020; Salminen et al. 2022a; Jung et al. 2018). The NMF algorithm chooses a representative demographic group based on the highest pattern value. To illustrate, we applied the following steps to generate data-driven personas:

1. Collect user statistics automatically via online analytics and social media application programming interfaces such as YouTube Analytics (under the channel owner's permission and following the terms of services by the source platform).
2. Prepare the dataset (also known as interaction matrix (An et al. 2018)) by mapping items (e.g., videos, webpages, ads, and products) with engagement (e.g., views, visits, clicks, and purchases) from different demographic user groups. These demographic groups communicate non-personally identifiable, aggregate information, and are thus the privacy-robust form of user analytics (An et al. 2018). In a typical fashion, online analytics platforms (including platforms such as YouTube Analytics and Facebook Insights) provide behavioral interaction statistics that are grouped by demographic attributes of age bucket, gender, and country (e.g., 25–34, Female, United States).
3. Detect engagement patterns based on matrix factorization conducted on the interaction matrix, applying the well-known nonnegative matrix factorization (NMF) approach by Lee and Seung (1999). This decomposes the interaction matrix into matrices that, respectively, contain information on user groups and items relative to the engagement patterns identified by using the matrix factorization algorithm.
4. Detect top-ranking demographic groups from the decomposed matrix containing information about the demographic groups and items—this is simply done by sorting the demographic groups' factorization scores in descending order for each engagement pattern and user group, choosing the highest-ranking group as the representative demographic group for a given engagement pattern (An et al. 2018).
5. Detect top-ranking content for each representative demographic group to build a skeletal persona that now has (a) demographic attributes and (b) most engaged items.
6. Proceed to enrich the skeletal personas with other information such as social media comments retrieved from the persona's top-ranking content, a picture chosen from a manually curated database, and a name assigned by a probabilistic naming algorithm (Jung et al. 2021).

7. Present the persona profiles in a Web UI to persona users (e.g., social media managers of the channel whose data was used for the persona generation).

As described above, the method considers the demographic aspects of persona generation, which can be considered desirable. In a similar vein, demographics play an important role in making personas lifelike and immersive (Salminen et al. 2021). So, considering both these features in the persona generation process is advantageous. To put it more precisely, our data-driven persona generation method using NMF approach has a broad range of applicability, spanning (1) various web analytics and social media data sources that provide demographically grouped data on user behavior, (2) diverse items such as videos, web pages, ads, products, and more, and (3) different engagement metrics like views, clicks, downloads, purchases, and so forth.

We applied the NMF algorithm consistently to generate each persona set in the series (N=10) and each organization (N=3) and channel (N=2). In total, we have 60 datasets of data-driven generated personas (the number of personas ranges from 5 to 160). We selected the highest patterns value and the corresponding demographic groups for each dataset. For example, if we have five personas in the dataset, we find the highest NMF values for each of the five personas and the five affiliated demographic groups.

32.3.3 Diversity and Fairness

Diversity is defined as how much each demographic variable (gender, age, and country) is represented in the generated persona dataset. The more gender, age, and country are represented, especially fringe audience demographical attributes in the generated persona, the higher the diversity. To measure diversity, we applied the Shannon diversity index for each data-driven personas dataset and the baseline dataset. Equation 32.1 shows the calculation.

$$H = \sum[(pi)ln(pi)] \tag{32.1}$$

In the Shannon diversity index, p is the proportion (n/N) of individuals of one particular demographic variable found in each of the persona datasets(n) divided by the total number of demographic variables (N) in the same dataset; it is the natural log, and $\sum$ is the sum of the calculations. A higher H value means higher diversity.

The Shannon diversity index has been selected because it is a widely used diversity measure that takes both richness and evenness into account (Kingrani et al. 2015). Precisely, the Shannon diversity index measures the diversity of an item by considering not only the number of different items present (richness) but also the commonness or rarity of each item while knowing the relative abundance (evenness) (Kingrani et al. 2015). For more illustrations and examples about the Shannon diversity index, please refer to this Wiki Page.[5] Using the Shannon diversity index measure, we ensure that we *actually* capture the diversity of generated personas by measuring how much they genuinely represent the various characteristics of the user base in a given dataset and consider both the richness and evenness of the dataset.

Fairness is defined as how much each of the demographic variables (i.e., gender, age, and country) is represented in the generated persona dataset *compared to* the baseline dataset by percentage. If a demographic attribute is highly prevalent in the baseline data, higher fairness means this specific attribute is also highly prevalent in the generated personas. To measure fairness, we applied a formula based on statistical parity, which is a metric that calculates the difference in the ratio between each unique demographic group in the persona dataset and the baseline file. For each demographic variable, the statistical parity (SP) calculation is shown in Eq. 32.2.

$$SPi = (Pi/P) - (ni/N) \tag{32.2}$$

For example, if five personas are from the Netherlands in a persona set of 10, then 50% of the personas are from the Netherlands. Given that 30% of the views in the baseline dataset are from the Netherlands, the SPi value for countries = 0.5 − 0.3 = 0.2. The fairness described in this example is called "representative fairness." It takes into consideration that the output, such as the selection of personas, reflects the distribution of the population or baseline data. In other words, it captures the idea that groups should be represented equally in the output as well as in baseline data. In this case, the proportion of personas from the Netherlands in the set is equal to the proportion of the audience from the Netherlands in the main channel, ensuring the demographic balance. Research including Salminen et al. (2022a) and Abbasi et al. (2021) have embraced this particular definition of fairness.

For all the demographic variables, we calculated the sum of each SPi value for age, gender, and country and subtracted it from 1. The statistical parity formula statistical parity (SPi-all) is shown in Eq. 32.3.

$$SPi - all = 1 - \sum(|SPi|) \tag{32.3}$$

Subsequently, we conduct correlation analysis and threshold analysis to understand the association between (1) the number of personas and diversity, (2) the number of personas and fairness, and (3) the diversity and fairness. For statistical testing, we applied $\alpha = 0.05$ as the threshold value for significance.

[5] https://en.wikipedia.org/wiki/Species_evenness

32.4 Correlation Analysis

32.4.1 The Relationship Between the Number of Personas and Diversity

To answer RQ1, we measured the Pearson correlation (N = 10) (N is 10 because the persona series has ten values ranging from 5 to 160 personas) between the number of personas and diversity for the three organizations and two social media channels for gender, age, and country of the audience.

32.4.1.1 Gender

For ORG01, we observed a positive correlation between increasing the number of personas and gender diversity for YouTube and Facebook channels. The correlation is substantially stronger for the YouTube channel ($r(10) = .64$) than for the Facebook channel ($r(10) = .37$). Also, the correlations are statistically significant ($p < .05$).

For ORG02, we found a positive correlation between increasing the number of personas and gender diversity for the YouTube and Facebook channels. The correlation is slightly stronger for the Facebook channel ($r(10) = .69$) than for the YouTube channel ($r(10) = .56$). Additionally, the correlations are statistically significant ($p < .05$).

For ORG03, we observed a positive correlation between increasing the number of personas and gender diversity for the YouTube and Facebook channels. The correlation is substantially stronger for the Facebook channel ($r(10) = .76$) than for the YouTube channel ($r(10) = .19$). Furthermore, the correlations are statistically significant ($p < .05$).

As such, the results indicate that there is a positive relationship between the number of personas and gender diversity, but the strength of this correlation varies among social media channels.

32.4.1.2 Age

For ORG01...3, we found a strong positive correlation ($r(10) > .70$) between the number of personas and age diversity for YouTube and Facebook channels. The correlations are statistically significant across all the organizations and channels ($p < .05$). Therefore, the results indicate that as we increase the number of personas, the age diversity increases, irrespective of the organization and social media channel.

32.4.1.3 Country

For ORG01...3, we found a strong positive correlation ($r(10) > .59$) between the number of personas and the diversity of users' countries for YouTube and Facebook channels. The correlations are statistically significant across all the organizations and channels ($p < .05$), except for the Facebook channel for ORG02, where we have a substantial trend toward significance ($p = .07$). Therefore, the results indicate that as we increase the number of personas, the country diversity increases.

Overall, the results indicate a positive correlation between the number of personas and diversity for gender, age, and country.

32.4.2 The Relationship Between the Number of Persona and Fairness

To address RQ2, we measured the Pearson correlation ($N = 10$) between the number of personas and fairness for the three organizations and two social media channels for gender, age, and country.

32.4.2.1 Gender

For ORG01...3, we observed a positive correlation between increasing the number of personas and gender fairness for YouTube and Facebook channels. The strength of the correlation $r(10)$ varies between .16 and .84. Nevertheless, only some of the correlations are statistically significant. In particular, ORG01 for both channels and ORG02 and ORG03 for YouTube channels exhibit a $p > .05$. Therefore, the results are not conclusive to report that there is a positive relationship between the number of personas and gender fairness.

32.4.2.2 Age

For ORG01...3, we observed a positive correlation ($r(10) > .46$) between increasing the number of personas and age fairness for YouTube and Facebook channels. The YouTube channels for ORG01 and ORG02 exhibit nonsignificant correlations ($p > .05$). Therefore, the results are not conclusive to report that there is a positive relationship between the number of personas and age fairness.

32.4.2.3 Country

For ORG1...3, we observed a positive correlation (r(10) >.29) between the number of personas and country fairness for YouTube and Facebook channels with the exclusion of ORG02, the YouTube channel, where we noted a weak negative correlation. The correlations are not all statistically significant. Therefore, the results are not conclusive to state that there is a relationship between the number of personas and country fairness.

To summarize, the results indicate mixed evidence on the correlation between the number of persona and fairness for gender, age, and country. In most cases, there is a positive (but nonsignificant) relationship.

32.4.3 The Relationship Between Diversity and Fairness

To address RQ3, measure the Pearson correlation (N = 10) between diversity and fairness for the three organizations and two social media channels for gender, age, and countries.

32.4.3.1 Gender

For ORG01...3, we found a strong positive correlation (r(10) >.91) between the gender diversity and fairness of YouTube and Facebook channels. All correlations are statistically significant. Therefore, the results indicate that increasing gender diversity also increases fairness when generating different numbers of personas.

32.4.3.2 Age

For ORG01...3, we found a strong positive correlation (r(10) >.88) between the age diversity and fairness of YouTube and Facebook channels. All correlations are statistically significant. Therefore, the results indicate that increasing age diversity also increases fairness when generating different numbers of personas.

32.4.3.3 Country

For ORG01...3, we found a positive correlation (r(10) >.25) between the country diversity and the fairness of YouTube and Facebook channels. All correlations are statistically significant except for both channels of ORG02. Therefore, the results are not conclusive to report that there is a relationship between country diversity and fairness.

Table 32.2 The Pearson correlation coefficient (r) results between the number of persona and diversity (Sect. 32.4.1). Significant values (p <0.05) are in bold. Overall, the results indicate a positive correlation between the number of personas and diversity for gender, age, and country

	Number of Persona—Diversity					
	ORG01		ORG02		ORG03	
Variables	Youtube	Facebook	Youtube	Facebook	Youtube	Facebook
Gender	**r = .64** **p = .05**	**r = .37** **p = .05**	**r = .56** **p = .05**	**r = .69** **p = .03**	**r = .19** **p = .05**	**r = .76** **p = .01**
Age	**r = .71** **p = .02**	**r = .83** **p = .002**	**r = .81** **p = .003**	**r = .86** **p = .001**	**r = .90** **p = .0003**	**r = .82** **p = .003**
Country	**r = .89** **p = .0005**	**r = .87** **p = .001**	**r = .89** **p = .0005**	r = **.59** **p = .067**	**r = .79** **p = .005**	**r = .76** **p = .009**

Table 32.3 The Pearson correlation coefficient (r) results between the number of persona and fairness (Sect. 32.4.2). Significant values (p <0.05) are in bold. Overall, the results indicate inconclusive evidence on the correlation between the number of persona and fairness for gender, age, and country

	Number of Persona—Fairness					
	ORG01		ORG02		ORG03	
Variables	Youtube	Facebook	Youtube	Facebook	Youtube	Facebook
Gender	r = .37 p = .29	r = .37 p = .28	**r = .58** **p = .07**	**r = .65** **p = .03**	r = .16 p = .66	**r = .84** **p = .002**
Age	r = .46 p = .18	**r = .81** **p = .004**	r = .6 p = .06	**r = .77** **p = .008**	**r = .87** **p = .0009**	**r = .86** **p = .001**
Country	r = .60 p = .06	r = .29 p = .42	r = −.14 p = .69	**r = .63** **p = .05**	**r = .73** **p = .01**	r = .52 p = .12

Table 32.4 The Pearson correlation coefficient (r) results between diversity and fairness (Sect. 32.4.3). Significant values (p <0.05) are in bold. Overall, the results indicate that there is a positive correlation between diversity and fairness for gender and age, but not necessarily for country

	Diversity—Fairness					
	ORG01		ORG02		ORG03	
Variables	Youtube	Facebook	Youtube	Facebook	Youtube	Facebook
Gender	**r = .91** **p = .0002**	**r = .99** **p = .0002**	**r = .95** **p = .0002**	**r = .96** **p = .0002**	**r = .98** **p =.0002**	**r = .978** **p = .0002**
Age	**r = .88** **p =.0002**	**r = .98** **p = .0002**	**r = .94** **p = .0002**	**r = .96** **p = .0002**	**r = .99** **p = .0002**	**r = .97** **p = .0002**
Country	**r = .87** **p = .001**	**r = .68** **p = .03**	r = .25 p = .48	r = .57 p = .08	**r = .99** **p = .0002**	**r = .92** **p = .0002**

Overall, the results indicate that there is a positive correlation between diversity and fairness for gender and age, but not necessarily for country (i.e., the results for country are mixed). Tables 32.2, 32.3, and 32.4 show a summary of the correlation results.

32.5 Threshold Analysis

We investigated whether any correlations change after a given number of personas. Specifically, we want to understand whether there is a threshold for the number of personas for obtaining a statistically significant positive correlation between (1) the number of personas and diversity (RQ1), (2) the number of personas and fairness (RQ2), and (3) diversity and fairness (RQ3). The threshold analysis quantifies precisely how much any of the correlations change when changing the number of personas. For this, we measured the correlations for RQ1, RQ2, and RQ3 between the number of personas from 5 till reaching each [15, 20, 25, 40, 45, 80, 85, and 160] personas or *thresholds*. Then we compared the correlations across gender, age, and countries for the three organizations and two channels.

32.5.1 Finding the Threshold for a Statistically Significant Positive Correlation Between the Number of Personas and Diversity

We modified the number of personas to different sets (from 5 to 15, 5 to 20, 5 to 25, 5 to 40, 5 to 45, 5 to 80, 5 to 85, and 5 to 160). Then, we measured the correlation between the number of personas and diversity (RQ1) eight times (corresponding to each set).

32.5.1.1 Gender

We did not find any threshold affecting the relationship between increasing the number of personas and gender diversity to make it a statistically significant positive correlation. To illustrate, the number of personas should reach at least 85 for ORG02 (Facebook) and ORG03 (Facebook) to exhibit a statistically significant positive correlation. The number of personas for ORG02 (YouTube) should range between 45 and 160 personas, while it should vary between 40 and 45 personas. For ORG01 (YouTube), the number of personas should reach 160 until we can observe a statistically positive correlation between the number of personas and gender diversity. Overall, we could not locate a specific threshold or a constant number of personas across *all* organizations and channels that would obtain a positive, statistically significant correlation between the number of personas and gender diversity.

32.5.1.2 Age

Across the three organizations and two channels, we observed variation in the minimum number of personas to obtain a statistically significant positive correlation between the number of personas and age. For example, ORG01 (YouTube) and ORG03 (Facebook) should have a minimum of 45 personas before a statistically significant positive correlation appears, while ORG01 (Facebook) should have a minimum of 25 personas. Moreover, we observed that ORG02 (Facebook) should have a minimum of 40 personas, while ORG03 (YouTube) should have a minimum of 80 personas. Uniquely, ORG02 (YouTube) does not have a minimum threshold for personas.

To determine a threshold or the constant number of personas across organizations and channels to have a statistically positive correlation between the number of personas and age diversity, we select the highest minimum threshold. We propose that the number of personas should be *at least 80* to have a statistically significant positive correlation between the number of personas and the age diversity across *all* three organizations and two channels.

32.5.1.3 Country

We observed variation in the minimum number of personas to obtain a statistically significant positive correlation between the number of personas and countries. ORG01 (YouTube and Facebook) should have a minimum of 25 personas to exhibit a statistically significant positive correlation, while ORG02 (YouTube) should have a minimum of 40 personas. ORG02 (Facebook) and ORG03 (YouTube and Facebook) do not have a minimum threshold for personas.

To determine a specific threshold or the constant number of personas across organizations and channels to have a statistically positive correlation between the number of personas and country diversity, we select the highest minimum threshold. Consistently, we propose that the number of personas should be *at least 40* to indicate a statistically significant positive correlation between the number of personas and the country diversity across *all* three organizations and two channels. As such, the results indicate that it is possible to determine a minimum number of personas to obtain a significant positive correlation between the number of personas and diversity for age and country.

32.5.2 Finding the Threshold for a Statistically Significant Positive Correlation Between the Number of Personas and Fairness

We modified the number of personas to different sets (from 5 to 15, 5 to 20, 5 to 25, 5 to 40, 5 to 45, 5 to 80, 5 to 85, and 5 to 160). Then, we measured the correlation between the number of personas and fairness (RQ2) eight times (corresponding to each set).

32.5.2.1 Gender

No single number of personas would obtain a positive and statistically significant relationship between the number of personas and gender fairness. We found a statistically significant positive correlation between the number of personas and gender fairness at 85 personas for ORG01 (YouTube) and 45 personas for ORG02 (YouTube). Furthermore, the number of personas should reach 80 for ORG02 (Facebook) and 85 for ORG03 (Facebook). The correlation remains negative before reaching 160 personas for ORG03 (YouTube). These results indicate that there is no specific threshold or a constant number of personas across *all* organizations and channels to obtain a positive, statistically significant correlation between the number of personas and gender fairness. While the previous section reported that the results are not conclusive enough to state a relationship between the number of personas and gender fairness, this section confirms that this does not change as we change the number of personas or thresholds.

32.5.2.2 Age

To observe a statistically significant positive correlation between the number of personas and age fairness, the number of personas should be at least 25 for ORG01 (YouTube) and ORG01 (Facebook), 40 for ORG02 (Facebook) and ORG03 (Facebook), and 80 for ORG03 (YouTube). Also, for ORG02 (YouTube), the number of personas should range between 20 and 85 to observe a significant positive correlation between the number of personas and age fairness.

Based on these results, we select a specific threshold or a constant number of personas across organizations and channels to have a statistically positive correlation between the number of personas and age fairness. Accordingly, we propose that the number of personas should be at *either 80 or 85* to obtain a statistically significant positive correlation between the number of personas and the age fairness across *all* three organizations and two channels. While the previous section found that the results are not conclusive enough to report a statistically significant positive relationship between the number of personas and age fairness, this section illustrates the value of the threshold analysis in determining the values of the possible number of personas.

32.5.2.3 Country

No specific number of personas threshold would yield a positive, statistically significant relationship between the number of personas and country fairness. More precisely, the number of personas should be exactly 45 for ORG02 (YouTube) to obtain a statistically significant positive correlation between the number of personas and country fairness. The minimum number of personas for the other two organizations and channels varies. These results indicate a lack of a specific threshold or a constant number of personas across *all* organizations and channels to obtain a statistically significant positive correlation between the number of personas and country fairness. While the previous section reported that the results are not conclusive to state a relationship between the number of personas and country fairness, this section confirms that this does not change as we change the number of personas or thresholds.

As such, the results indicate that it is possible to determine a certain number of personas to obtain a significant positive correlation between the number of personas and fairness for age only.

32.5.3 Finding the Threshold for a Statistically Significant Positive Correlation Between Diversity and Fairness

We modified the number of personas to different sets (from 5 to 15, 5 to 20, 5 to 25, 5 to 40, 5 to 45, 5 to 80, 5 to 85, and 5 to 160). Then, we measured the correlation between diversity and fairness (RQ3) eight times (corresponding to each set).

32.5.3.1 Gender

We found that for most organizations and channels, there is no specific threshold or minimum number of personas to obtain a statistically significant positive correlation between gender diversity and fairness. Still, for ORG02 (Facebook), the threshold should be at least 20 personas to exhibit a statistically positive correlation between the two variables (diversity and fairness). Considering that, we suggest that the minimum number of personas or consistent threshold should be *at least 20* to obtain a significant positive correlation between gender diversity and fairness across *all* organizations and channels.

32.5.3.2 Age

To observe a statistically significant positive correlation between age diversity and fairness, the minimum number of personas has to be at least 25 for ORG01 (YouTube). The minimum number of personas should be *at least 25* to have a significant positive correlation between age diversity and fairness. Considering this, we suggest that the minimum number of personas or consistent threshold should be *at least 25* to have a significant positive correlation between age diversity and fairness across *all* organizations and channels.

Table 32.5 The threshold for the number of personas for obtaining a statistically significant positive correlation between the number of personas and diversity (column 1), the number of personas and fairness (column 2), and diversity and fairness (column 3) for gender, age, and country. Acknowledging that the results in Sect. 32.4 were not conclusive enough to report a statistically significant positive relationship between the number of personas and age fairness, this section illustrates the value of the threshold analysis to identify a statistically significant positive relationship between the two variables at a specific number of personas

	Number of Persona—Diversity	Number of Persona—Fairness	Diversity—Fairness
Gender	Different thresholds—nonspecific	No threshold has been identified	At least 20 personas
Age	At least 80 Personas	At 80 or 85 Personas	At least 25 Personas
Country	At least 40 Personas	No Threshold has been Identified	At 85 Personas

32.5.3.3 Country

To observe a statistically significant positive correlation between country diversity and fairness, the minimum number of personas has to range between 40 and 85 for ORG02 (YouTube), while for ORG02 (Facebook), the minimum number of personas has to be 85. As such, we recommend that the consistent threshold or the number of personas should be *exactly 85* to have a significant positive correlation between country diversity and fairness across *all* organizations and channels.

As such, the results indicate that it is possible to determine a minimum or exact number of personas to obtain a significant positive correlation between diversity and fairness for gender, age, and country. Table 32.5 shows a summary of the threshold analysis results.

32.6 Discussion

As organizations aim to design systems and artifacts and create content that reaches a diverse and global audience, developing personas that fairly represent their audiences becomes crucial. Related to this overarching goal, we investigated whether creating more personas helps achieve diversity and fairness in personas.

The correlation results indicate that creating more personas is beneficial in increasing diversity in gender, age, and country. The correlation results indicate that creating more personas is beneficial in increasing diversity in gender, age, and country. Precisely, our approach of measuring diversity using the Shannon diversity index (richness and evenness) ensures that we actually capture the extent of how much various characteristics (i.e., gender, age, and country) whether they are rare or common in baseline datasets are represented in the persona datasets. As such, we found that we obtain higher diversity indices when increasing the number of personas as more gender groups, more age groups, and more countries are represented in the persona output. Correspondingly, a larger number of personas can accommodate a greater diversity of gender, age, and country. Still, the threshold analysis suggests that creating at least 80 personas to achieve age diversity and at least 40 personas to achieve country diversity is recommended. Our findings align with previous research that more personas are needed to describe diverse and international audiences (Nielsen et al. 2013; Salminen et al. 2019).

On the other hand, the correlations and the threshold results indicate that it is recommended to set the number of personas to either 80 or 85 to achieve age fairness. Although measuring the correlation did not uncover the relationship between the number of personas and age fairness, conducting threshold analysis and computing the correlation for a specific number of personas has exposed the link between these two variables. For the other demographic attributes (gender and country), the relationship between the number of personas and fairness remains uncertain. Furthermore, our results indicate that increasing the demographic diversity in personas leads to more fairly representing the data. It is suggested to create at least 20 personas to obtain the balance between gender diversity and fairness, to create at least 25 personas to obtain the balance between age diversity and fairness, and to create exactly 85 personas to obtain the balance between country diversity and fairness.

Our findings corroborate prior work that advocated for creating a larger number of personas to represent the user population better and cover a wide range of demographic features (Salminen et al. 2019, 2022c,b). In particular, as discussed by Salminen et al. (2022b), creating more than 40 personas is vital to achieving gender diversity, age diversity, and country diversity. Our results contribute to previous findings by showing the optimal number of personas that are advised to achieve not only diversity but also fairness in human experience design. Furthermore, our findings acknowledge the importance of conducting threshold analysis to identify critical points or boundaries where a significant change in relationship occurs between variables, such as in our case between the number of personas and age fairness. This confirms that testing various thresholds can be used reliably for future work to affirm inconclusive correlations between the number of persona and fairness in gender and country (whether positive or negative correlations).

The study has implications for the creation of data-driven personas (as described in Sect. 32.2.1) and, more specifically, data-driven personas derived from large datasets of users' data (as described in Table 32.1):

First: Persona Proliferation: A Pathway to Inclusive and Diverse Design The fact that creating more personas yields a richer demographic diversity can nudge content creators and technology developers toward diversity and promote more inclusive design. Demographically diverse user representation can help mitigate designer biases and ensure that no individual or community is harmed or left behind in terms of technology use (Goodman-Deane et al. 2021). Diversity sensitivity puts the users and their unique attributes at the center, thus enabling human-centered and context-aware technology design (Himmelsbach et al. 2019). Diversity sensitivity also helps to promote inclusive design and avoid biases and stereotypes that can be passed down (intentionally or unintentionally) to systems and artifacts (Hoque et al. 2022).

Second: Persona Proliferation: A Pathway to Fair Design Pushing toward increasing demographic diversity in personas is key to representing the audience fairly. A persona needs to represent a group of people with sufficient fairness to say with some certainty that if the persona can use the product, then the audience will be able to as well (Goodman-Deane et al. 2018). Fair persona creation is considered a crucial dimension in the ethical use of personas and user datasets (Salminen et al. 2020c).

Third: Persona Proliferation: Hitting Both Diverse and Fair Design with One Stone We provide evidence that achieving a high level of diversity does not compromise fairness, as data-driven persona generation can preserve both in harmony.

Our study has several strengths; we have studied multiple organizations and social media channels to support our findings. Further, we have generated copious numbers of personas ranging from 5 to 160. Also, we have extended our analysis to include a threshold analysis that helps organizations and persona creators determine the minimum number of personas to achieve the maximum benefits of increasing the number of personas.

Despite the above strengths, there are limitations concerning the findings of our study. Initially, there is a data collection limitation in terms of the YouTube AI algorithm inferring the users' demographics (Narayanan 2023). YouTube assigns demographics to YouTube channels and content based on the platform's observation of the audience demographics that most often visit a channel or content, and this is used to assign demographics to users who visit the channel or content but whose demographics are not observable. In essence, anonymous users are assigned demographics based on the demographic characteristics of the users' most observed engaging with the content.

Also, based on these inferred demographics assigned to users, YouTube recommends and ranks the search results of users, and this potentially means that users associated with demographics who are more likely to have already viewed content might be more likely to view the content. As such, our findings are only applied to demographics that were inferred or reported by YouTube. Besides, the data that we collected using the API was limited to content engagement per gender, age, and country. This may suggest that our dataset was limited due to ethical consent received from three different organizations allowing us to retrieve only non-identifiable information related to gender, age, and country from social media channels' API. Additionally, the API collects these three demographic variables intrinsically. Therefore, we chose to examine diversity and fairness in data-driven persona based on these three demographic variables. Future work aims to collect data related to other variables such as race or ethnicity and enrich our dataset.

In addition, we have increased the number of personas based on a broad range (multiplier of two) and a narrow range (addition of five), so we are unaware if changing this range would produce different results or thresholds. Future work should be directed to alter this range and add different parameters. Future work will also aim to run a study with designers to see how they use 5 personas baseline versus 80 personas to maintain diversity and fairness. This leads to improved personas to inform their design iterations and evaluate whether the "inclusiveness" of their design improved as a result of using more personas.

Furthermore, the results did not reveal a consistent threshold for the minimum number of personas created across gender, age, and country. Thus, it is a best practice to develop a guideline for organizations or filtering option based on our threshold analysis that if they need to increase gender diversity, create an X number of personas, while if they need to increase age diversity, create an X number of personas, and so on. In principle, it is possible to iterate across the different demographic attributes and choose the value for the number of personas that satisfies the requirements for all the attributes.

It is also important to acknowledge that fairness is a complex topic that involves many different contexts and multifaceted sociocultural concepts (Mulligan et al. 2019; Caton and Haas 2023). Recent work in HCI has looked at finding a standard definition of fairness, but they only found that people understand fairness differently (Kasinidou et al. 2021). Accordingly, future work will aim to define and measure fairness differently and compare the findings with the ones uncovered in this research. Additionally, we acknowledge that diversity is a complex topic (Himmelsbach et al. 2019) and that the dimensions we examined in this study were limited to gender, age, and location. Considering that we obtain fruitful results, our next plan is to explore other dimensions of diversity, such as race/ethnicity and socioeconomic status and corroborate our datasets with other datasets such as census information or surveys. Future work should define and examine multiple definitions and dimensions of diversity and fairness and make the interplay between the two variables much clearer.

We aim to verify the research findings with designers who would utilize the data, assessing the extent to which the proposed practices can uphold the initial promises of diversity and fairness. Additionally, we intend to evaluate key design elements of our recommendations based on different approaches (Ledo et al. 2018) covering feasibility, usability, usefulness, and desirability of our research recommendations.

Finally, the organizations we studied are considered large with millions of views or content engagement, so our findings are reflective of large organizations that deal with international audiences and broad content bases. Future work should be carried out to validate the results on organizations of different sizes and different levels of engagement and see if the results vary compared to this sample. Organizations could also be varied by their degree of demographic bias—for example, dividing them into low, medium, and high bias, and then analyzing by group. Doing so would help in asserting how sensitive the suggested persona numbers are to organizational differences. The challenge here, as illustrated by our datasets, is that organizations' audiences often are demographically imbalanced.

32.7 Conclusion

The study reveals that increasing the number of personas also increases demographic diversity, providing a better representation of various user groups. We also show that the higher number of personas means that specific demographic personas are more fairly represented. Moreover, the results imply that attempts to accelerate the level of diversity also lead to accelerating the level of fairness. We pinpointed that organizations need to be conscious of different threshold values of the number of personas to retain specific demographic attributes' diversity, fairness, and both. In conclusion, more personas should be generated and considered when trying to understand the users and potentially design more inclusive technology.

32.8 Ethical Statement

The data was retrieved from three different organizations' Facebook and YouTube channels using programmatic data collection, also known as application programming interface (API). We received full consent from these three organizations before collecting the data. We also complied with the terms and conditions of the social media platforms when accessing their APIs. We stored the data in a secured database to which only the research team members have authorized access. Although social media data could be sensitive in some cases, in our case, the anonymity and privacy of the social media users whose behavior the data reflects is maintained because the data is aggregated in demographic groups. In other words, there is no individual-level data in our dataset, and the dataset contains no personally identifiable information. Finally, we surmise that the findings of this study are likely to pose no harm to the users. The study aims to increase the benefit for society by promoting more inclusive design practices.

References

Abbasi M, Bhaskara A, Venkatasubramanian S (2021) Fair clustering via equitable group representations. In: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency. FAccT'21. Association for Computing Machinery, New York, pp 504–514

Allison K, Bussey K (2020) Communal quirks and circlejerks: a taxonomy of processes contributing to insularity in online communities. In: Proceedings of the international AAAI conference on web and social media, vol 14, pp 12–23. ISSN: 2334-0770

An J et al (2018) Imaginary people representing real numbers: Generating personas from online social media data. ACM Trans Web 12(4):27:1–27:26. ISSN: 1559-1131

Antle AN (2008) Child-based personas: need, ability and experience. Cogn Technol Work 10(2):155–166. ISSN: 1435-5566

Anvari F et al (2019) Teaching user centered conceptual design using cross-cultural personas and peer reviews for a large cohort of students. In: 2019 IEEE/ACM 41st international conference on software engineering: software engineering education and training (ICSE-SEET), pp 62–73
Bowen J et al (2020) Personas revisited: extending the use of personas to enhance participatory design. In: Proceedings of the 11th nordic conference on human-computer interaction: shaping experiences, shaping society. NordiCHI'20. ACM, New York, pp 1–12
Breslin S, Wadhwa B (2018) Gender and human-computer interaction. In: The Wiley handbook of human computer interaction. Wiley, United States, pp 71–87
Caton S, Haas C (2023) Fairness in machine learning: a survey. ACM Comput Surv. Just Accepted. ISSN:: 0360-0300
Cooper A (1999) The inmates are running the asylum. In: Arend U, Eberleh E, Pitschke K (eds) Software-Ergonomie'99: Design von Informationswelten. Vieweg+Teubner Verlag, Wiesbaden, p 17
Cutting K, Hedenborg E (2019) Can personas speak? Biopolitics in design processes. In: Companion publication of the 2019 on designing interactive systems conference 2019 companion. DIS'19 Companion. Association for Computing Machinery, New York, pp 153–157. ISBN: 978-1-4503-6270-2
Fuglerud KS et al (2020) Co-creating persona scenarios with diverse users enriching inclusive design. In: Antona M, Stephanidis C (eds) Universal access in human-computer interaction. Design approaches and supporting technologies. Lecture notes in computer science. Springer International Publishing, Cham, pp 48–59. ISBN: 978-3-030-49282-3
Goodman-Deane J et al (2018) Evaluating inclusivity using quantitative personas. In: DRS biennial conference series
Goodman-Deane JA-L et al (2021) Developing personas to help designers to understand digital exclusion. In: Proceedings of the design society, vol 1. Cambridge University Press, United States, pp 1203–1212
Grgić-Hlača N et al (2022) Dimensions of diversity in human perceptions of algorithmic fairness. In: Equity and access in algorithms, mechanisms, and optimization. EAAMO'22. Association for Computing Machinery, New York, pp 1–12. ISBN: 978-1-4503-9477-2
Hill CG et al (2017) Gender-inclusiveness personas vs. stereotyping: can we have it both ways? In: Proceedings of the 2017 CHI conference on human factors in computing systems. CHI'17. ACM, New York, pp 6658–6671
Himmelsbach J et al (2019) Do we care about diversity in Human Computer Interaction: a comprehensive content analysis on diversity dimensions in research. In: Proceedings of the 2019 CHI conference on human factors in computing systems. CHI'19. ACM, New York, pp 1–16
Hoque MN, Ghai B, Elmqvist N (2022) DramatVis personae: visual text analytics for identifying social biases in creative writing. In: Proceedings of the 2022 ACM designing interactive systems conference. DIS'22. Association for Computing Machinery, New York, pp 1260–1276. ISBN: 978-1-4503-9358-4
Jensen I et al (2017) Developing international personas: a new intercultural communication practice in globalized societies. J Intercult Commun 43. ISSN: 1404-1634
Jansen BJ, Salminen JO, Jung S-G (2020) Data-driven personas for enhanced user understanding: combining empathy with rationality for better insights to analytics. Data Inf Manag 4(1):1–17. ISSN: 2543-9251
Jansen BJ et al (2021) Data-driven personas. Synthesis lectures on human-centered informatics. Springer International Publishing, Cham. ISBN: 978-3-031-01103-0 978-3-031-02231-9
Jung S-G et al (2018) Automatically conceptualizing social media analytics data via personas. In: Proceedings of the international AAAI conference on web and social media, vol 12(1). ISSN: 2334-0770
Jung S-G, Salminen J, Jansen BJ (2021) All about the name: assigning demographically appropriate names to data-driven entities. ISBN: 978-0-9981331-4-0
Kasinidou M et al (2021) I agree with the decision, but they didn't deserve this: future developers' perception of fairness in algorithmic decisions. In: Proceedings of the 2021 ACM conference on

fairness, accountability, and transparency. FAccT'21. Association for Computing Machinery, New York, pp 690–700. ISBN: 978-1-4503-8309-7

Kingrani SK, Levene M, Zhang D (2015) Diversity analysis of web search results. In: Proceedings of the ACM web science conference. WebSci'15. Association for Computing Machinery, New York, pp 1–2. ISBN: 978-1-4503-3672-7

Kong Q et al (2022) Slipping to the extreme: a mixed method to explain how extreme opinions infiltrate online discussions. In: Proceedings of the international AAAI conference on web and social media, vol 16, pp 524–535. ISSN: 2334-0770

Ledo D et al (2018) Evaluation strategies for HCI toolkit research. In: Proceedings of the 2018 CHI conference on human factors in computing systems, pp 1–17

Lee DD, Seung HS (1999) Learning the parts of objects by non-negative matrix factorization. Nature 401(6755):788–791. ISSN: 1476-4687. Nature Publishing Group

Marsden N, Pröbster M (2019) Personas and identity: looking at multiple identities to inform the construction of personas. In: Proceedings of the 2019 CHI conference on human factors in computing systems. CHI'19. Association for Computing Machinery, New York, pp 1–14. ISBN: 978-1-4503-5970-2

Marsden N, Hermann J, Pr¨obster M (2017) Developing personas, considering gender: a case study. In: Proceedings of the 29th Australian conference on computer-human interaction. OZCHI'17. ACM, New York, pp 392–396

Marwick A (2013) They're really profound women, they're entrepreneurs': conceptions of authenticity in fashion blogging. In: 7th international AIII conference on weblogs and social media (ICWSM), vol 2011, pp 1–8

Meissner F, Blake E (2011) Understanding culturally distant end-users through intermediary-derived personas. In: Proceedings of the South African institute of computer scientists and information technologists conference on knowledge, innovation and leadership in a diverse, multidisciplinary environment. SAICSIT'11. ACM, New York, pp 314–317

Miaskiewicz T, Luxmoore C (2017) The use of data-driven personas to facilitate organizational adoption – a case study. In: Design J 20(3):357–374. Routledge. ISSN: 1460-6925

Molenaar L (2017) Data-driven personas: generating consumer insights with the use of clustering analysis from big data (Visited on 17 Aug 2023)

Mulligan DK et al (2019) This thing called fairness: disciplinary confusion realizing a value in technology. In: Proceedings of the ACM on human-computer interaction (CSCW 2019), vol 3, pp 119:1–119:36

Narayanan A (2023) Understanding social media recommendation algorithms. http://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms (Visited on 14 May 2023)

Netzorg R et al (2021) PopFactor: live-streamer behavior and popularity. In: Proceedings of the international AAAI conference on web and social media, vol 15, pp 432–442

Nielsen L et al (2013) Going global with personas. In: Kotzé P et al (eds) Human-computer interaction – INTERACT 2013. Lecture notes in computer science. Springer, Berlin/Heidelberg, pp 350–357

Salminen J, Jung S-G, Jansen BJ (2019) Detecting demo-graphic bias in automatically generated personas. In: Extended abstracts of the 2019 CHI conference on human factors in computing systems. CHI EA'19. ACM, New York, pp 1–6

Salminen J et al (2020a) Persona transparency: analyzing the impact of explanations on perceptions of data-driven personas. Int J Hum–Comput Interact 36(8):788–800. ISSN: 1044-7318

Salminen J et al (2020b) Rethinking personas for fairness: algorithmic transparency and accountability in data-driven personas. In: Degen H, Reinerman-Jones L (eds) Artificial intelligence in HCI. Lecture notes in computer science. Springer International Publishing, Cham, pp 82–100. ISBN: 978-3-030-50334-5

Salminen J et al (2020c) The ethics of Data-driven personas. In: Extended abstracts of the 2020 CHI conference on human factors in computing systems. CHI EA'20. ACM, New York, pp 1–9