



How Can Natural Language Processing and Generative AI Address Grand Challenges of Quantitative User Personas?

Joni Salminen¹(✉), Soon-gyo Jung², Hind Almerekhi³, Erik Cambria⁴,
and Bernard Jansen²

¹ School of Marketing and Communication, University of Vaasa, Vaasa, Finland
jonisalm@uwasa.fi

² Qatar Computing Research Institute, Hamad Bin Khalifa University,
Doha, Qatar

³ Qatar National Research Fund, Doha, Qatar

⁴ School of Computer Science and Engineering, Nanyang Technological University,
Singapore, Singapore

Abstract. Human-computer interaction (HCI) and natural language processing (NLP) can engage in mutually beneficial collaboration. This article summarizes previous literature to identify grand challenges for the application of NLP in quantitative user personas (QUPs), which exemplifies such collaboration. Grand challenges provide a collaborative starting point for researchers working at the intersection of NLP and QUPs, towards improved user experiences. NLP research could also benefit from focusing on generating user personas by introducing new solutions to specific NLP tasks, such as classification and generation. We also discuss the novel opportunities introduced by Generative AI to address the grand challenges, offering illustrative examples.

Keywords: User personas · Quantitative user personas · Natural language processing · Generative AI

1 Introduction

This article is intended for researchers working on the intersection of AI (specifically, NLP technologies) and HCI (specifically, quantitative user personas). NLP refers to *natural language processing*, defined as a computer's ability to comprehend spoken and written human language [40]. HCI refers to *human-computer interaction*, which is a multidisciplinary research field that focuses on the interaction between computers and their users, towards the design of more usable and helpful technology, also referred to as user-centered design (UCD).

AI refers to *artificial intelligence*, defined as the “simulation of human intelligence processes by machines” [9]. While this term is often used to refer to applied

machine learning (ML) instead, we use it here because of the current convention in academic fields, including HCI and NLP, essentially referring to applied ML and various pseudo-intelligent algorithms and automata. NLP is commonly considered as a subfield of ML [55]. *Generative AI* is a subfield of AI that produces outputs, including text, images, and videos. *Large language models* (LLMs) are part of the Generative AI technology, and they can be programmed to generate text based on instructions (i.e., prompts) given by the user [2, 26].

Personas (also known as user personas, design personas, marketing personas, and so on) are fictitious characters that aim to represent real user types [15] (also known as segments, clusters, groups... i.e., groups of people that are similar to each other and different from other groups), providing valuable information to designers and others working on UCD tasks (collectively referred to as ‘stakeholders’ in this study) to more empathically connect with users (see example in Fig. 1). *Persona profiles* typically contain demographic information (age, gender, location), as well as various information reflecting the user’s needs, wants, and pain points. *Quantitative user personas* (QUPs), in particular, are personas generated using algorithms; hence, they are sometimes referred to as *algorithmically-generated personas* [28, 62]. Another common term for QUPs in the literature is *data-driven personas* [44] and, with the rise of AI, the QUP terminology has also incorporated *AI personas* [23].

Regardless of this conceptual plurality, all QUP methods share the common trait of using algorithms in the process of persona generation – either fully, ranging from data collection to segmentation and enrichment – or partially, as a

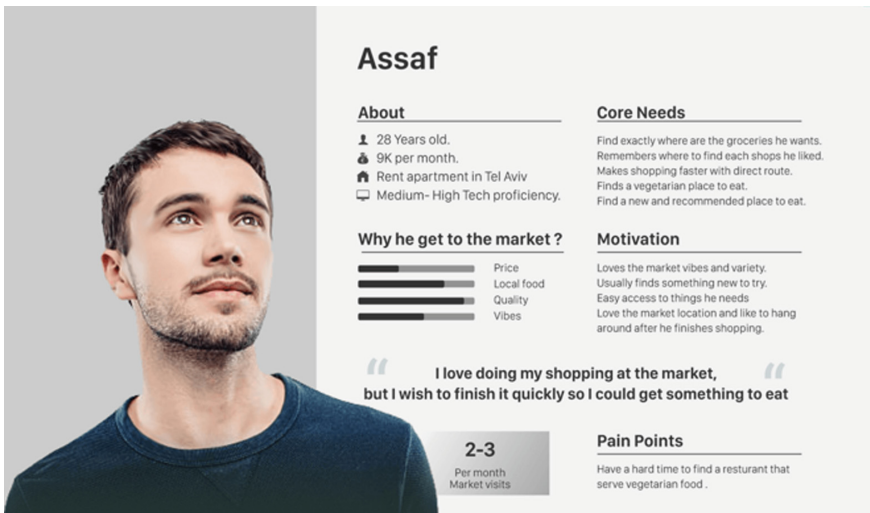


Fig. 1. An example of a persona profile. Personas typically contain demographic information (age, gender, location), and information reflecting the user’s needs, wants, and pain points, often illustrated using direct quotes. Image source: <https://www.justinmind.com/blog/user-persona-templates/>.

part of the overall process. Overall, this process of algorithmic persona generation involves specific challenges where NLP technologies can certainly be useful. We explore these opportunities.

The purpose of this research is to outline, based on prior research, four grand challenges for QUPs, and suggest research designs based on prominent NLP technologies to address these challenges. We focus specifically on QUPs, even though there are other persona types, such as qualitative personas. The reader might be interested in further motivation of this work. *Why study the grand challenges of QUPs? Moreover, why study them using NLP?* These are both worthwhile questions, and we address them in the following.

First, studying the grand challenges of QUPs is important because, despite the popularity of personas in HCI, design, marketing, health informatics, education, social media analytics, and so on [33, 47, 52], there is a consensus that many crucial challenges in persona creation and application still linger [13, 60, 64]. Some of these challenges were observed already in the early 2000s [22, 50], but have not been resolved, lending credence to their difficulty.

This then leads us to the second motivational question - *why apply NLP to QUPs?* Due to two reasons: first, as mentioned, many persona problems remain to be solved; logically, it then makes sense to seek potential solutions from other fields, as these fields may contain technologies and tools that can be applied to personas. NLP, in particular, is a field of rapid growth and progress, as evidenced by the massive increase in the field's popularity among academics, as well as the recent breakthroughs such as transformer language models [17, 38].

Second, the field of HCI is inherently interdisciplinary and applied - it seeks inspiration from other fields. Therefore, translating HCI problems as NLP problems is a logical and likely fruitful activity. In summary, the grand challenges in this study discuss known problems in the development of QUPs that have yet to be solved, partially due to the lack of interaction between HCI and NLP (i.e., HCI does not fully understand NLP advances, and NLP is not that familiar with QUPs or personas in general).

2 Related Work

2.1 User Personas in HCI

Personas can be created using qualitative, quantitative, or mixed methods [27]. Due to the strong progress in data science tools, libraries, and frameworks that help automate data collection and analysis tasks, QUPs are becoming more common in the field [59]. There are systems such as *Automatic Persona Generation* [31], that provide replicable and automatically updating personas for stakeholders in many fields. Such systems allow researchers to generate personas based on behavioral data of sizeable volumes for quantitative analysis [4].

Regardless of the creation method, there are some general design goals for developing good personas. Among others, these design goals include (1) *ability to evoke empathy among the stakeholders*, (2) *accuracy* (i.e., correctness, validity

of the persona information), and (3) *relevance* (i.e., the persona contains information that helps stakeholders make decisions about the group of people the persona represents).

First, empathy refers to stakeholders identifying with the persona’s human qualities [21]. A good persona engages its observer and facilitates the process of perspective taking, i.e., seeing the world from the eyes of the persona [52].

Second, a good persona does not contain bogus information that would mislead stakeholders about the true characteristics of the user segment it represents [13]. In other words, the persona is *data-driven* [25, 44]; i.e., based on actual data about users, whether quantitative, qualitative, or mixed data [27].

Third, a good persona is relevant, containing the necessary information to be useful for decision making but nothing more – the information should not be distracting or redundant. A stakeholder has to be able to take action based on what they learn from the persona [13].

The persona creation process tends to follow two stages: the data about users is first *segmented* using either an algorithm or a human analyst, and the resulting barebone (or skeletal) persona segments are then enriched with personified information to create complete, rounded persona profiles [3, 71]. This *personification* transforms the nameless, faceless segments into identifiable personas that describe a behavioral and demographic pattern in the data [31]. Personification aims to result in realistic, well-rounded personas that stakeholders can identify with and that contain information of value [51, 63].

Because there are demanding technical challenges in QUPs [12], we believe that partnering with NLP researchers can help reach the broader mission of providing stakeholders with personas that help them create usable products. In other words, high-quality personas serve stakeholders by facilitating their jobs and users by enabling usable and user-friendly products that serve real needs.

2.2 Use of NLP in Persona Research

NLP is a field that combines AI, linguistics, and computer science to model human language, especially to analyze and process large amounts of text [54]. From this definition, we can observe at least two guiding principles (GPs):

- a) **GP01:** NLP is an applied field - whenever language is used, NLP can be of help.
- b) **GP02:** NLP is particularly designed to handle large volumes of textual data.

Both notions are compatible with QUPs. Language is a vehicle for storytelling which, in turn, is a route for empathy and learning [41]. Throughout history, people have learned lessons from others through stories. To this end, personas have been likened to narratives or metaphors that tell a story [50]. Therefore, there is an innate compatibility between the *purposes* of NLP and personas.

Moreover, the ability to handle large volumes of text is a major advantage of NLP. While QUPs are typically generated from structured tabular data, such as clickstreams, CRM records, or audience statistics [3, 76, 77], there is tremendous

potential in using *unstructured* textual data, essentially mining the text for information about users and using that to feed information into the persona profile. For example, people express their pain points [66], opinions, needs, and wants in online posts [10]. This textual data could be utilized to generate insightful personas.

Despite the above-described potential, using NLP for QUPs has been tentative [23], with the main efforts summarized in the following *five use cases*.

First, previous research has applied NLP for classifying personas' topics of interest. For this, different methods have been applied, including latent Dirichlet allocation (LDA) [4], which is a form of unsupervised topic modeling [7], as well as supervised ML (text classification) in which a topical taxonomy is first developed together with the organization using the personas, followed by annotating data and using it to train topic classifiers [8]. Researchers have also applied zero-shot classification [53] that uses a pre-trained language model for determining a match between an array of topic words and the persona's content (e.g., comments). This approach has been found useful for small organizations with limited in-house data for training purposes.

Second, researchers have applied sentiment analysis lexicons, such as the EmoLex [48], to assign a sentiment score to a persona based on the social media comments associated with the persona [30]. Sentiment aims to indicate the persona's general attitude towards the channel or organization - however, aspect-based sentiment analysis has not been implemented, which would indicate how the persona feels about a specific topic [35] or product [72].

Third, QUPs generated from social media datasets may include abusive or toxic comments that are then included automatically in the generated persona profiles. Prior research addresses how these comments can be detected and then, if the stakeholder wants, removed from the persona profiles [65]. This is done to improve the user experience (UX) of stakeholders using the personas. This process involves using NLP, more specifically text classification trained on social media comments.

Fourth, personas designed for dialogue systems imitate real conversations by interactively reacting to stakeholder inputs [1, 24]. This is similar to using a chatbot. Persona, in this context, is defined as a linguistic type (personality) that remains consistent over one or multiple user sessions [34]. The purpose is to improve the UX for those engaging with the dialogue interface. Unlike other efforts, this problem describes a specific subclass of generation problems.

Fifth, NLP techniques can also be applied in the segmentation process itself; i.e., the process preceding personification. This involves identifying latent patterns [5] or clusters [46] from textual data, typically using latent semantic spaces and associated techniques (e.g., singular value decomposition, cosine similarity, and so on). When using NLP for segmentation, the process tends to require the use of external information, such as domain experts' participation in the personification process [45].

To summarize, topic modeling, sentiment analysis, toxicity detection, and user segmentation are all text classification problems in NLP terminology, while

dialog systems is a text generation problem. While the above use cases demonstrate how NLP can be of service to QUPs. However, there is room for expanding this interdisciplinary collaboration, given specific unsolved challenges. In the next section, we discuss four grand challenges of QUPs and outline ideas on how NLP can help address them.

3 Grand Challenges for Quantitative User Personas

Compatible with previous research [69, 70], we define a grand challenge as a *difficult but not unsolvable problem that represents a roadblock or bottleneck for a field to progress and evolve into the next stage*. A classic example of grand challenges is the list of 23 problems that David Hilbert proposed to the international society of mathematicians in 1900. This list is said to have “galvanized the efforts of mathematicians for the next century” [42]. Inspired by this, the grand challenges presented here are based on the authors’ experience on researching personas (multiple years, more than a dozen publications, a textbook on the topic) as well as on NLP research both in and out of the persona domain. The research team comprises people from diverse backgrounds, including those focused on NLP (2 people) and those focused on HCI/QUPs (3 people). As such, we aim to offer researchers thought-provoking problems that (hopefully) help foster interdisciplinary collaboration between HCI and NLP.

When formulating the grand challenges, one has to ask, “What is a good grand challenge like?”. According to above definition, it is unsolved and difficult, otherwise, it would not be “grand” or “challenge”. In addition, it is vital that the challenge is *interesting* to the community to which it is presented. This implies that the grand challenges need to be translated to technical problems to attract the interest of NLP researchers. Otherwise, it is unlikely that NLP researchers will invest any serious effort into addressing the challenges.

In our case, we engaged in an internal discussion between the research team members specialized in HCI (problem space experts) and NLP (solution space experts). This discussion represented a “sanity check” to ensure that the challenges were framed in a way that (1) is likely to advance persona creation, (2) corresponds to a technical NLP problem, and (3) interests both HCI and NLP researchers. These assessments are, of course, subjective determinations; other scholars might add or remove elements from this list. Nonetheless, the following grand challenges (GCs) passed this internal validation, and we are thus confident to propose them. They are introduced as follows, in no particular order of importance. For each challenge, we include (a) a description, (b) illustrative examples, (c) pointers on how to address it, (d) pointers on how to evaluate any solutions to it, and (e) reflections on how Generative AI could be of help.

GC1: Frankenstein Personas (i.e., Consistency Problem) *Description.* Susanne Bødker, a Danish computer scientist, first identified this problem, referring to it as ‘Frankenstein personas’ [6] (p. 98). The problem is that the creation

of personas involves combining isolated datapoints into a persona profile that may appear inconsistent, i.e., contain conflicting information.

The general version of the challenge is:

GC1: How to detect and correct inconsistent information in the persona profile?

This problem exists for QUPs because the persona creation process can rely on data collected from multiple sources, processed using different algorithms at different steps of the process. This can result in a patched-up persona profile that can contain inconsistent information.

Example Cases. In a system that generates QUPs automatically using social media data, a male persona says, “us women should stick together” (indicating wrong sex); there is an Indian persona that says, “You Indians cook really nice curry” (indicating nationality mismatch); and a Saudi Arabian persona refer to themselves as Tanzanian (see Fig. 2). The ramification is that the stakeholder using the persona becomes confused about the persona’s true identity [61].

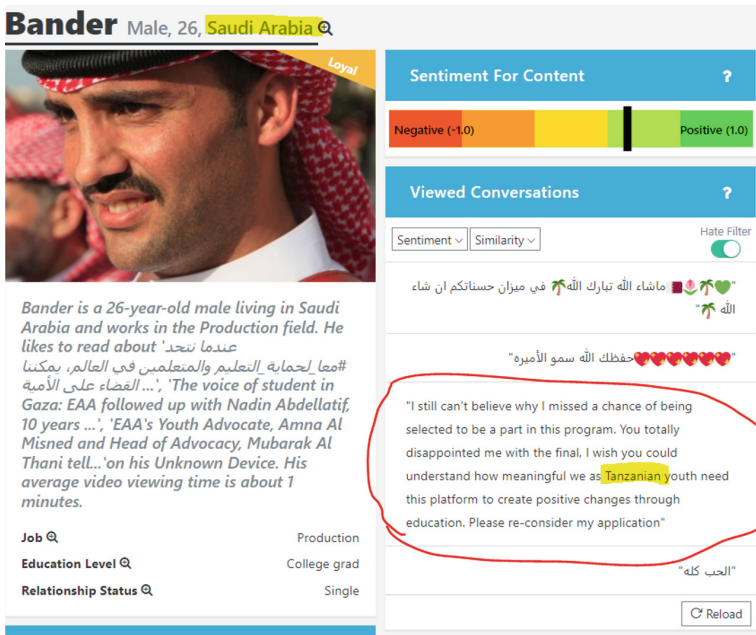


Fig. 2. Bander is from Saudi Arabia, but a random comment from videos of Bander’s behavioral viewing type says, “I wish you could understand how meaningful we as Tanzanian youth need this platform to create positive changes (...)”

How to address? As established, the problem is that the information in the persona profile may be inconsistent due to a failure to verify consistency between different information elements. There is no process for ensuring consistency of each information element relative to other elements. So, developing such processes and mechanisms would address this issue.

Here, NLP can be useful. For example, a body of NLP work focuses on inferring demographic attributes from social media texts [74]. This work typically aims to identify differences in how demographic groups express themselves [74]; e.g., differences between males and females, or the youth and the elderly, and then use the identified differences for prediction. Thus, one approach is predicting the gender, age, and nationality of a user writing a given social media comment. This, without any additional meta-data and only relying on feature extraction and engineering from the text, qualifies as an extremely challenging problem.

The challenge is exacerbated by several factors, including (a) the lack of publicly available training data, (b) the large number of demographic group combinations (gender \times age \times nationality yields thousands of subgroups to consider on a global scale), and the (c) separation of identity and “true” demographics (e.g., some males may identify with female characteristics; an older and younger person might share an interest in the same popular culture, etc.). Given these constraints, it might be appropriate, in some cases, to relax the requirements of including specific demographic information in the persona profiles; for example, omitting gender and age and focusing on other attributes, such as psychographics [67] or simply a narrative description focusing on user needs. Nevertheless, regardless of what information is chosen into the persona profile, the information must form a coherent whole. When the information originates from multiple sources and undergoes multiple diverging algorithmic manipulations, ensuring consistency becomes a potential issue, requiring consistency verification for QUPs. Here, the role of NLP to address this issue remains open.

How to evaluate? Ascertaining the success of solving Bødker’s Frankenstein personas can take the form of technical and stakeholder evaluation. Technical evaluation would, presuming some test data, apply ML performance metrics that pertain to classification tasks, such as the F1 and ROC-AUC scores. Stakeholder evaluation, on the other hand, would relate to whether stakeholders would observe inconsistencies between the personas’ different information elements.

How could Generative AI contribute? We conducted a simple test with Open AI’s GPT-3.5, the state of the art large language model (LLM), also referred to as a type of Generative AI [2]. We provided the following prompt:

Your task is to evaluate and report any inconsistent information in the following persona description. Inconsistent information means that the details about the persona are not aligned; something does not match. Here is the persona description:

Bander is a 26-year-old male from Saudi Arabia. Here is a comment from him: “I wish you could understand how meaningful we as Tanzanian youth need this platform to create positive changes.”

Is there any inconsistent information in the persona description?

As evidenced by Fig. 3, GPT was able to detect the inconsistency, showing promise in addressing this grand challenge. However, a more formal and systematic evaluation would be needed to draw more definite conclusions.

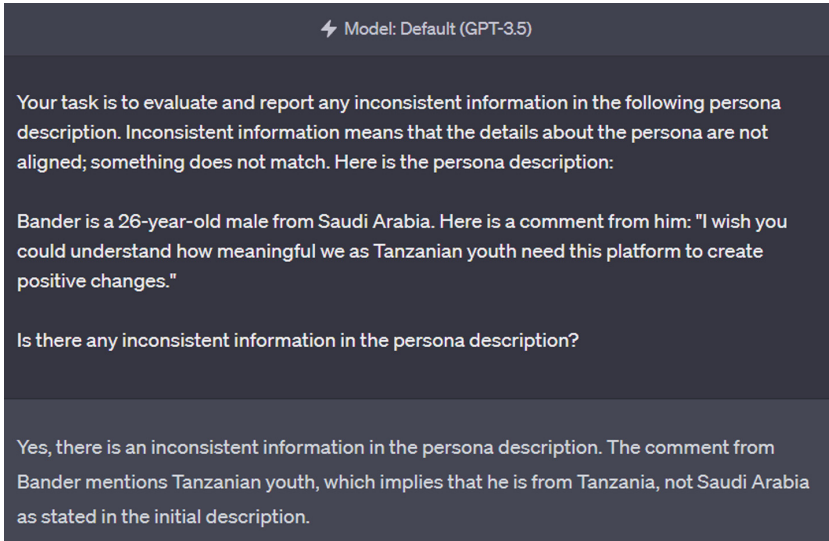


Fig. 3. ChatGPT correctly identified the inconsistency in the persona information: “Yes, there is an inconsistent information in the persona description. The comment from Bander mentions Tanzanian youth, which implies that he is from Tanzania, not Saudi Arabia as stated in the initial description.”

GC2: Extrapolate Beyond the Immediate User Data Description. A general advantage of QUPs is their use of aggregated (i.e., non-personally identifiable) data, which protects individual users’ privacy [3]. On the other hand, the dependency on the user data at hand – often a fairly limited one when compared to the full range of human thoughts – hinders the ability to understand the persona’s needs and wants beyond the specific contexts available in the dataset. For example, the dataset can contain comments about cars, therefore giving us information how the people classified under the persona feel about cars. However, if the dataset is missing information about motorbikes, we do not know how the persona feels about them. Therefore,

GC2: What does the persona think of Topic x ? (Where x is not part of the training data)

How to address? A potential solution could involve generative tasks that apply transfer learning via language models. Modern language models in NLP

are trained on large-scale datasets (up to billions of document samples) and advanced algorithms (e.g., transformers). One major advantage is that, to some extent, these models can be fine-tuned using a smaller annotated dataset, thereby adjusting the model’s broad knowledge to a specific narrow problem. There are language models that perform prediction, generation (i.e., natural language generation (NLG), which is a subfield of NLP) or both [16]. There’s also a growing body of work on “data synthesis” for use cases that require generating textual data to perform causal inference [75].

Essentially, if we are able to represent a type of user with textual corpus that represents “who this person is”, then, in theory, we might be able to fine-tune a language model to query this user type’s opinions in real-time in an accurate manner. Conceptually, this is compatible with the notion of the persona being a ‘mental model’ [52] and a surrogate for the real user [15]—an LLM might be able to capture this mental model and simulate it plausibly and accurately.

In QUPs, the user type can be represented by latent statistical object that generates language according to grammaticality (‘conformity to language usage’) and the patterns observed in a group of users’ language use [5]. The intuition is that, ideally, the language model contributes to grammaticality and the fine-tuning to the stability of the persona. The stability can be divided into two sub-parts: linguistic consistency (i.e., the style of speaking remains the same) and character consistency (i.e., the persona remains the same person throughout the dialogue — e.g., when the persona is asked, “Do you like ice-cream?” and the user group it is based on hates ice cream, it is able to respond, “No.”). Furthermore, to address the scarcity of information, systems such as SenticNet [11] can be helpful as they analyze multiple dimensions (e.g., sentiment, emotions, personality) from the same text and therefore can be used for extracting more information from the same content and, possibly, for building more complete user models by leveraging the interactions between the extracted information variables.

How to evaluate? A myriad of technical metrics measure linguistic consistency [57]. In addition, we recommend a user-focused evaluation, e.g., developing a conversational user interface through which a stakeholder can chat with or talk to the persona; a variation of the Turing test [73]. Herein lies an important observation; while the Turing test requires a human to assess whether the machine is able to execute human behaviors that are indistinguishable from real humans, the test can be passed by an efficient “impersonator” that speaks fluently, answers rationally, but tells things that are untrue, yet convincing. There are a myriad of cases where the human evaluator lacks the expertise to assess the veracity of an LLM; personas being of these.

This is why both aspects - realism (the dialogue persona produces answers that *appear* correct) and accuracy (the produced answers actually reflect the thinking of the group the persona represents) - need consideration. Otherwise, a situation may emerge where the given answer appears to make sense (i.e., the mechanism by which the algorithm generates text works well in terms of generating grammatically correct text), but the meaning of the text does not correspond to the specific and distinct group of people that the persona is.

Due to opaqueness of how the text is generated by LLMs, the stakeholder is able to evaluate realism but not accuracy - i.e., they need to trust the system,

and the system developers need to verify, in a robust manner, that the meanings given by the persona are consistent and truthful.

Again, we want to stress the importance of the problem of mimicry - with the current NLP models, it is *not* difficult to produce grammatically correct text that appears sensible. The challenge is not to do that - the challenge is to get one step closer to intelligence by simulating what the people belonging to the persona's group *do* and *would* think about topics, past or present. Grammatically correct text generation is only the first step in this journey. Generating persona dialogue that passes both scientific and user scrutiny poses a tremendous challenge for NLP because it requires that the persona remains *consistent* in its role [18].

How could Generative AI contribute? As previously, we explore the potential of GPT-3.5 to address this challenge. Below, we show the dialogue with the model: *Bander is a 26-year-old male from Saudi Arabia. He likes to follow social media content on education and youth development. Using your general knowledge about people like Bander, tell me, what else does he like to follow?*

As can be seen from the exhibit, based on its general knowledge, GPT is able to locate related topics. However, this information cannot be relied on without some form of triangulation. How to do such triangulation efficiently and at scale remains an important aspect of resolving this challenge.

Based on general knowledge about individuals interested in education and youth development, there are several other areas that Bander might be inclined to follow. Here are some possibilities:

1. Current Affairs and News: Bander might be interested in staying updated on local and international news to understand the broader context of educational and youth-related issues.
2. Personal Development: Bander may follow content related to personal growth, motivation, self-improvement, and success stories to enhance his own skills and mindset.
3. Technology and Innovation: Bander might have an interest in keeping up with the latest technological advancements and innovations that can impact education and youth development.
4. Nonprofit Organizations and NGOs: Bander may follow the work of nonprofit organizations and non-governmental organizations (NGOs) focused on education, youth empowerment, and social change.
5. Leadership and Entrepreneurship: Bander might be interested in content related to leadership development, entrepreneurship, and business strategies to gain insights into effective ways of driving positive change.
6. Social Issues and Advocacy: Bander might follow content related to social justice, equality, and human rights, as these topics often intersect with education and youth development.

It's important to note that these suggestions are based on common interests and general knowledge, and individual preferences may vary.

GC3: Engage with the Persona in Real-Time, Asking It Anything You Need to Know for Your Decision-Making Task (i.e., Meaningful Engagement) *Description:* The persona does not need to be presented as profile at all. We have already hinted at the dialogue-based persona systems, where one can directly interact with the persona using conversational means [14, 29, 58]. These UIs are likely to replace or complement the traditional persona profiles as LLMs further evolve. The grand challenge here is, in fact, a series of challenges for such systems:

GC3: How to generate personas using LLMs? How to save personas using LLMs? How to update personas using LLMs?

How to address? The obvious way to address this is to focus on the goal of creating personas using LLMs. Fortunately, there are efforts to this end [18]. However, the major challenge there is that the NLP and HCI research communities work on the problem separately, without concentrated efforts and coordination.

How to evaluate? The ultimate method of evaluation for the successful creation of LLM personas is end-user adoption. The technology is quickly reaching a degree of maturity where it can be deployed in real persona systems for commercial and non-profit organizations – the feedback from these entities determines the success of the research endeavors.

How could Generative AI contribute? As previously, we explore the potential of GPT-3.5 to address this challenge. Below, we show the dialogue with the model: “A persona represents a segment of people that share the same characteristics. Generate five personas that describe NEET youth (not in education, employment, or training).”

Sure! Here are five personas that describe NEET youth:

1. Lucas, the Aspiring Artist:

Age: 20

Background: Lucas is a talented artist who has recently graduated from high school. He dreams of pursuing a career in fine arts but struggles to find opportunities to showcase his work. He comes from a lower-middle-class family and lacks the financial means to pursue higher education or professional training. Lucas spends most of his time at home, honing his artistic skills and searching for freelance gigs online.

2. Sarah, the Discouraged Dropout:

Age: 19

Background: Sarah dropped out of high school due to personal issues and a lack of interest in academics. She comes from a disadvantaged background, with limited family support and financial resources. Sarah feels discouraged and disconnected from society, finding it challenging to secure stable employment without a high school diploma. She spends her days aimlessly browsing the internet and occasionally takes up short-term jobs to make ends meet.

3. Raj, the Technophile:

Age: 22

Background: Raj is a tech-savvy individual who possesses excellent computer skills and knowledge about the latest technological trends. However, due to a lack of formal education or certifications, he struggles to find meaningful employment in the tech industry. Raj spends most of his time at home, engaging in online communities, learning coding languages, and exploring new technologies. He hopes to land a remote job or freelance projects that utilize his expertise.

4. Emily, the Caregiver:

Age: 24

Background: Emily had to drop out of college to take care of her ailing parent. She is a compassionate and responsible individual, dedicating her time and energy to provide support at home. Emily's caregiving responsibilities make it difficult for her to commit to a full-time job or educational program. She often feels isolated from her peers and finds solace in online support groups for caregivers. Emily explores online courses and remote work opportunities to gain skills and earn income from home.

5. Alex, the Lost Graduate:

Age: 23

Background: Alex recently completed a bachelor's degree in a field that is currently experiencing a lack of job prospects. Despite the effort put into acquiring a degree, Alex faces significant challenges in finding employment relevant to their qualifications. Lacking guidance and career support, Alex spends their days feeling demotivated, switching between odd jobs and searching for entry-level positions. They are eager to gain relevant experience and willing to undergo further training if the opportunity arises.

Remember, these personas are fictional representations based on common characteristics of NEET youth and should not be taken as specific profiles of real individuals.

To assess whether it is possible to continue dialogue with one of the personas, we decided to engage with Emily:

Now, assume the role of Emily, and answer the following question: Emily, how could we motivate you to continue your studies?

As Emily, to motivate me to continue my studies, consider:

1. Flexible study options that accommodate my caregiving responsibilities.
 2. Financial support such as scholarships or grants for caregivers.
 3. Providing career guidance and mentorship to help me navigate my future path.
 4. Tailored education programs integrating practical experiences related to caregiving.
 5. Establishing a community of caregivers for support and networking.
- These factors would encourage me to pursue further education while managing my caregiving duties.

Again, GPT shows realistic potential in addressing this challenge. However, we must bear in mind that the information of "Emily" is based on general knowledge the model has about NEET youth – this knowledge may be flawed or biased [2, 26]. The information would need to be triangulated with domain experts and/or using a model finetuned with a domain-specific dataset.

GC4: Explain Algorithmic Persona Creation *Description.* Persona creation involves a degree of complexity that can make stakeholders question the resulting personas’ validity, applicability, and usefulness [43]. In the case of qualitative analysis, it is extremely hard to explain precisely how certain information was chosen for the persona, and it is nearly impossible to replicate the analysis with an independent researcher and obtain identical results. In the case of QUPs, the situation appears better at first glance, as one can apply statistical analysis and other algorithms on a fixed data and achieve identical (or at least highly similar) results on each run. However, these processes are opaque to a stakeholder, so the question about how the personas were created still remains topical. Furthermore, an individual piece of persona information is difficult to trace back to its constituents, even though stakeholders often pay attention to details in the persona profiles and ask about how they were derived. So, given these trust and transparency issues, the challenge is:

GC4: How can the quantitative persona creation process be explained?

How to address? To this end, there has been progress in AI Explainability [36, 49]. In its simplest form, explanation is a visualization that shows feature importance of a machine learning task like classification, for example, which words led to the model to predict a certain class. For example, if the topic of interest of a persona is “Cars”, then how did the model come to this conclusion? Explainability frameworks, such as *LIME* [56] and *SHAP* [39], exist for both local and global inferences, thereby potentially addressing some aspects of this challenge. However, in the case of non-tabular data, both LIME and SHAP show major limitations in global explanations. These approaches are generally available for NLP, but the explainability of more advanced deep-learning models is an open research area [37], partly because these models tend to rely on non-human interpretable features (i.e., they create a numerical representation of a text sample that is difficult to trace back to individual words or phrases). Therefore, we invite research on NLP model interpretability towards explaining how persona information is derived.

How to evaluate? The success of transparency relies on two primary factors: (a) whether the provided explanations are reliable, i.e., describe truthfully how a certain decision was made, and (b) whether the stakeholders using the personas find the explanations satisfactory, informative, and comprehensible. Specific techniques for assessing concepts [32] and influence [20] can help explain global patterns and trace their provenance through the model’s processes.

How could Generative AI contribute? For this task, we initially tested GPT-3.5. However, its explanations were not of optimal quality, so we tested GPT-4 which provided a technically more correct answer to the prompt, “Explain how NMF can be used for creating personas.” The explanation given by GPT-4 can be viewed in the online link¹. What we find particularly interesting is GPT-4’s ability to personalize the explanation (see Fig. 4). we asked GPT-4 to make two

¹ <https://chat.openai.com/share/52da0742-1a00-48c7-988f-1c954570e148>.

variations of this explanation, one for children and one for first-year marketing students. For each, we got satisfactory responses (due to space limitations, these responses can be viewed online²).

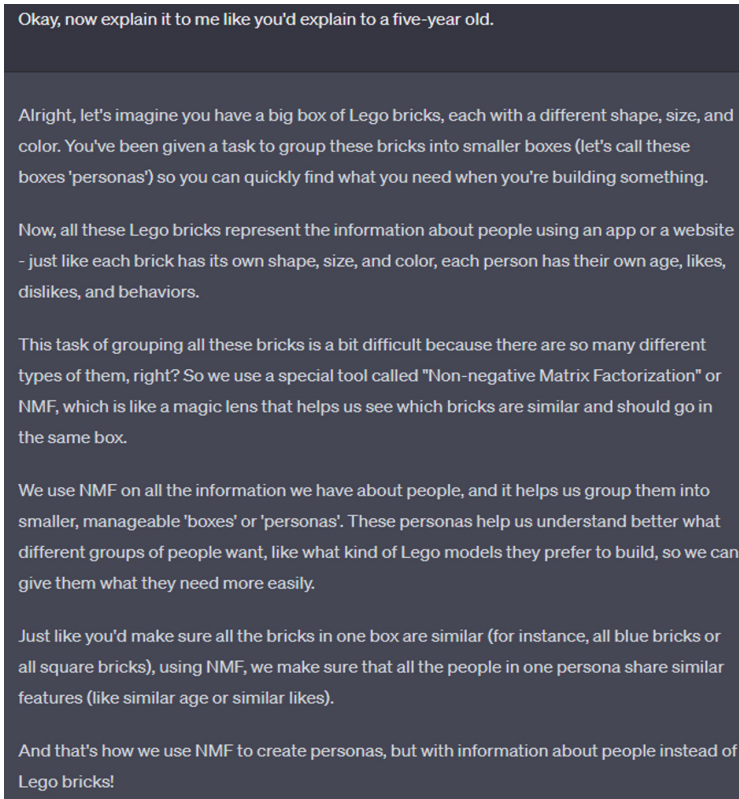


Fig. 4. ChatGPT explaining persona creation using non-negative matrix factorization to a five-year-old.

As with the other challenges, the explanations given by GPT show great promise, especially in terms of personalization. However, more testing is needed.

4 Discussion

Although the use of AI and ML for personas is not novel [25], applying NLP in this domain is nascent and still at its early stages. As such, there is a lack of guidelines and vision for how to apply NLP in a productive manner to address persona creation challenges. For this, the challenges need to be first defined.

² *ibid.*

This work represents an attempt to that end. While we do not claim to present a fully comprehensive list of neither the problems nor the potential solutions, we nonetheless posit that the proposed challenges are relevant for persona theory and practice, and the proposed NLP research directions are feasible.

The potential impact of addressing these challenges is substantial, especially for algorithmic persona generation at scale - the result could be personas that have consistent and complete information, can be understood through discussion, and offer stakeholders the possibility to understand and scrutinize the (currently) opaque information selection process. Better personas, in turn, could lead into designing more user-friendly products, which is beneficial for all.

We suggest that the potential solutions in this space are evaluated using a mixed approach - technical metrics (as customary in NLP studies) and user metrics (e.g., perceptions, subjective feedback; as customary in HCI studies). This would enable arguing for novelty and value of contribution from multiple sides. An example of validating stakeholder feedback on personas is the *Persona Perception Scale* [68], an instrument that measures several impressions, including the persona's perceived credibility, consistency, clarity, and completeness. Human assessment is also vital for rating the quality of NLP outputs [19].

The NLP community possesses unique capabilities for addressing the challenges. Therefore, much depends on its mobilization. For this, it is instrumental to frame the challenges as NLP problems, simply because NLP researchers seek to contribute to the field of NLP. This mobilization can be seen as a grand challenge in itself as cross-disciplinary collaboration can be challenging to create. Nonetheless, the mutual exchange of experiences, methods, directions, and insights is vital for addressing grand challenges [12]. Specific ideas for mobilizing NLP researchers include (1) *workshops and special issues* (e.g., ACL, CHI); (2) *competitions* in Kaggle and other data science platforms; and (3) *reachouts* and evangelization via personal correspondence and research seminars. QUPs should be advertised as a prominent theme.

5 Concluding Remarks

The potential in NLP for personas is plentiful. Persona profiles can be enriched with additional information such as sentiment, topics, of interest, pain points, affinities, and quotes; inferred from text pertaining to the identified persona segment (e.g., tweets, forum posts) or extrapolated using pre-trained language models. While the persona information can be artificial to some extent, it needs to be realistic and truthful, i.e., representative of the persona's unique nature among other personas. Even though current QUPs are mostly static, interactivity could be added using NLP technologies (chat and voice interfaces), thus improving the accessibility and UX of personas. These opportunities require the mobilization of NLP research that integrates with HCI requirements for personas. Other scholars are invited to contribute to the discussion on the grand challenges of QUPs. Particularly, GPT models show promise in addressing each challenge, calling for further evaluation studies.

References

1. Ali Amer Jid Almahri, F., Bell, D., Arzoky, M.: Personas design for conversational systems in education. *Informatics* **6**(4), 46 (2019). <https://doi.org/10.3390/informatics6040046>
2. Amin, M.M., Cambria, E., Schuller, B.W.: Will affective computing emerge from foundation models and general artificial intelligence? a first evaluation of Chat-GPT. *IEEE Intell. Syst.* **38**(2), 15–23 (2023)
3. An, J., Kwak, H., Jung, S., Salminen, J., Jansen, B.J.: Customer segmentation using online platforms: isolating behavioral and demographic segments for persona creation via aggregated user data. *Soc. Netw. Anal. Min.* **8**(1), 1–19 (2018). <https://doi.org/10.1007/s13278-018-0531-0>
4. An, J., Kwak, H., Salminen, J., Jung, S.G., Jansen, B.J.: Imaginary people representing real numbers: generating personas from online social media data. *ACM Trans. Web (TWEB)* **12**(3), 1–26 (2018)
5. Bamman, D., O'Connor, B., Smith, N.A.: Learning latent personas of film characters, p. 10. Bulgaria, Sofia (2013)
6. Bødker, S., Christiansen, E., Nyvang, T., Zander, P.O.: Personas, people and participation: challenges from the trenches of local government. In: *The 12th Participatory Design Conference*, p. 91. ACM Press, Roskilde (2012). <https://doi.org/10.1145/2347635.2347649>. <http://dl.acm.org/citation.cfm?doid=2347635.2347649>. Accessed 31 Mar 2020
7. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *J. Mach. Learn. Res.* **3**(Jan), 993–1022 (2003)
8. Brosas, H., Lim, E., Sevilla, D., Silva, D., Ong, E.: Classifying and extracting data from facebook posts for online persona identification. In: *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation* (2018)
9. Burns, E., Laskowski, N., Tucci, L.: What is artificial intelligence (ai)? definition, benefits and use cases (2022). <https://www.techtarget.com/searchenterpriseai/definition/AI-Artificial-Intelligence>. Accessed 01 Sept 2022
10. Cambria, E.: Affective computing and sentiment analysis. *IEEE Intell. Syst.* **31**(2), 102–107 (2016)
11. Cambria, E., Liu, Q., Decherchi, S., Xing, F., Kwok, K.: Senticnet 7: a commonsense-based neurosymbolic ai framework for explainable sentiment analysis. In: *Proceedings of LREC 2022* (2022)
12. Candello, H., et al.: Cui@chi: mapping grand challenges for the conversational user interface community. In: *CHI 2020: CHI Conference on Human Factors in Computing Systems*, pp. 1–8. ACM, Honolulu (2020). <https://doi.org/10.1145/3334480.3375152>. <https://dl.acm.org/doi/10.1145/3334480.3375152>. Accessed 09 June 2022
13. Chapman, C., Milham, R.P.: The personas' new clothes: methodological and practical arguments against a popular method. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, no. 5, pp. 634–636 (2006). <https://doi.org/10.1177/154193120605000503>
14. Chu, E., Vijayaraghavan, P., Roy, D.: Learning personas from dialogue with attentive memory networks. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2638–2646. Association for Computational Linguistics (2018)
15. Cooper, A.: *The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity*, 1st edn. Sams - Pearson Education, Indianapolis (1999)

16. Cummings, P., Mullins, R., Moquete, M., Schurr, N.: Hello World! I am Charlie, an Artificially Intelligent Conference Panelist (2021). <http://hdl.handle.net/10125/70656>. Accessed 01 Sept 2022
17. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
18. Dinan, E., et al.: The second conversational intelligence challenge (ConvAI2) (2019). <http://arxiv.org/abs/1902.00098>
19. Fadel, A., Al-Ayyoub, M., Cambria, E.: Justers at semeval-2020 task 4: evaluating transformer models against commonsense validation and explanation, p. 535–542 (2020)
20. Garima, L.F., Kale, S., Sundararajan, M.: Estimating training data influence by tracing gradient descent. In: NIPS 2020, pp. 19920–19930. Curran Associates Inc., Red Hook (2020). Accessed 01 Sept 2022
21. Grudin, J.: Why personas work: the psychological evidence. In: Pruitt, J., Adlin, T. (eds.) *The Persona Lifecycle*, pp. 642–663. Elsevier (2006). <https://linkinghub.elsevier.com/retrieve/pii/B9780125662512500137>. <https://doi.org/10.1016/B978-012566251-2/50013-7>
22. Grudin, J., Pruitt, J.: Personas, participatory design and product development: an infrastructure for engagement, p. 8. Sweden (2002)
23. Holzinger, A., Kargl, M., Kipperer, B., Regitnig, P., Plass, M., Müller, H.: Personas for artificial intelligence (AI) an open source toolbox. *IEEE Access* **10**, 23732–23747 (2022). <https://doi.org/10.1109/ACCESS.2022.3154776>
24. Hwang, S., Kim, B., Lee, K.: A data-driven design framework for customer service chatbot. In: Marcus, A., Wang, W. (eds.) *HCI 2019*. LNCS, vol. 11583, pp. 222–236. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23570-3_17
25. Jansen, B., Salminen, J., Jung, S.G., Guan, K.: *Data-Driven Personas, Synthesis Lectures on Human-Centered Informatics*, vol. 14, 1st edn. Morgan & Claypool Publishers, San Rafael (2021)
26. Jansen, B.J., Jung, S.G., Salminen, J.: Employing large language models in survey research. *Natural Lang. Process. J.* 100020 (2023). <https://doi.org/10.1016/j.nlp.2023.100020>. <https://www.sciencedirect.com/science/article/pii/S2949719123000171>
27. Jansen, B.J., Jung, S.G., Nielsen, L., Guan, K.W., Salminen, J.: How to create personas: Three persona creation methodologies with implications for practical employment. *Pac. Asia J. Assoc. Inf. Syst.* **14**(3) (2022). <https://doi.org/10.17705/1pais.14301>. <https://aisel.aisnet.org/pajais/vol14/iss3/1>
28. Jansen, B.J., Jung, S.G., Salminen, J.: Finetuning analytics information systems for a better understanding of users: evidence of personification bias on multiple digital channels. *Inf. Syst. Front.*, 1–24 (2023)
29. Jiang, H., Zhang, X., Cao, X., Kabbara, J., Roy, D.: Personallm: investigating the ability of gpt-3.5 to express personality traits and gender differences. arXiv preprint [arXiv:2305.02547](https://arxiv.org/abs/2305.02547) (2023)
30. Jung, S.G., Salminen, J., Jansen, B.J.: Giving faces to data: creating data-driven personas from personified big data. In: *IUI 2020*, pp. 132–133. Association for Computing Machinery, Cagliari (2020). <https://doi.org/10.1145/3379336.3381465>. Accessed 29 Apr 2020
31. Jung, S.G., Salminen, J., Kwak, H., An, J., Jansen, B.J.: Automatic persona generation (apg): a rationale and demonstration. In: *CHIIR 2018*, pp. 321–324. Association for Computing Machinery, New York (2018). <https://doi.org/10.1145/3176349.3176893>. Accessed 01 Sept 2022

32. Kim, B., et al.: Interpretability beyond feature attribution: quantitative testing with concept activation vectors (tcav), pp. 2668–2677. PMLR (2018). <https://proceedings.mlr.press/v80/kim18d.html>. ISSN: 2640-3498
33. Korsgaard, D., Bjørner, T., Sørensen, P.K., Burelli, P.: Creating user stereotypes for persona development from qualitative data through semi-automatic subspace clustering. *User Model. User-Adap. Inter.* **30**(1), 81–125 (2020). <https://doi.org/10.1007/s11257-019-09252-5>
34. Li, Y., Kazemeini, A., Mehta, Y., Cambria, E.: Multitask learning for emotion and personality traits detection. *Neurocomputing* **493**, 340–350 (2022)
35. Liang, B., Su, H., Gui, L., Cambria, E., Xu, R.: Aspect-based sentiment analysis via affective knowledge enhanced graph convolutional networks. *Knowl.-Based Syst.* **235**, 107643 (2022)
36. Liao, Q.V., Gruen, D., Miller, S.: Questioning the ai: informing design practices for explainable ai user experiences, pp. 1–15 (2020)
37. Liu, H., Yin, Q., Wang, W.Y.: Towards explainable nlp: a generative explanation framework for text classification. [arXiv:1811.00196](https://arxiv.org/abs/1811.00196) (2018)
38. Liu, Y., et al.: Roberta: a robustly optimized bert pretraining approach. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692) [cs] (2019). [arXiv: 1907.11692](https://arxiv.org/abs/1907.11692)
39. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions, vol. 30. Curran Associates, Inc. (2017). <https://papers.nips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>. Accessed 01 Sept 2022
40. Lutkevich, B., Burns, E.: What is natural language processing? an introduction to nlp (2021). <https://www.techtarget.com/searchenterpriseai/definition/natural-language-processing-NLP>, Accessed 01 Sept 2022
41. Madsen, S., Nielsen, L.: Exploring persona-scenarios - using storytelling to create design ideas. In: *Human Work Interaction Design: Usability in Social, Cultural and Organizational Contexts*, pp. 57–66. IFIP Advances in Information and Communication Technology (2010). https://doi.org/10.1007/978-3-642-11762-6_5
42. Maryland, O.G.C.F.S.W.U., Baltimore, S.O.S.W.W.R.S., Email, A.M.P.E.S.U.: Grand challenges for social work (2022). <https://grandchallengesforsocialwork.org/about/history/>. Accessed 01 Sept 2022
43. Matthews, T., Judge, T., Whittaker, S.: How do designers and user experience professionals actually perceive and use personas?. In: *The 2012 ACM Annual Conference* p. 1219. ACM Press, Austin (2012). <https://doi.org/10.1145/2207676.2208573>. <http://dl.acm.org/citation.cfm?doid=2207676.2208573>. Accessed 31 Mar 2020
44. McGinn, J.J., Kotamraju, N.: Data-driven persona development, p. 1521–1524. ACM, Florence (2008). <https://doi.org/10.1145/1357054.1357292>
45. Miaskiewicz, T., Kozar, K.A.: Personas and user-centered design: how can personas benefit product design processes? *Des. Stud.* **32**(5), 417–430 (2011)
46. Miaskiewicz, T., Sumner, T., Kozar, K.A.: A latent semantic analysis methodology for the identification and creation of personas, pp. 1501–1510. ACM (2008). <http://dl.acm.org/citation.cfm?id=1357290>
47. Minichiello, A., Hood, J.R., Harkness, D.S.: Bringing user experience design to bear on stem education: a narrative literature review. *J. STEM Educ. Res.* **1**(1–2), 7–33 (2018)
48. Mohammad, S.M., Turney, P.D.: Crowdsourcing a word-emotion association lexicon. *Comput. Intell.* **29**(3), 436–465 (2013)

49. Mueller, S.T., Hoffman, R.R., Clancey, W.J., Emery, A.K., Klein, G.: Explanation in human-ai systems: a literature meta-review synopsis of key ideas and publications and bibliography for explainable ai. Technical report (2019). <https://apps.dtic.mil/sti/citations/AD1073994>
50. Nielsen, L.: Engaging personas and narrative scenarios, PhD Series, vol. 17. Samfundslitteratur (2004)
51. Nielsen, L., Hansen, K.S., Stage, J., Billestrup, J.: A template for design personas: Analysis of 47 persona descriptions from Danish industries and organizations. *Int. J. Sociotechnol. Knowl. Dev.* **7**(1), 45–61 (2015). <https://doi.org/10.4018/ijskd.2015010104>
52. Nielsen, L., Storgaard, H.K.: Personas is applicable: a study on the use of personas in Denmark, pp. 1665–1674. ACM (2014)
53. Pamungkas, E.W., Basile, V., Patti, V.: A joint learning approach with knowledge injection for zero-shot cross-lingual hate speech detection. *Inf. Process. Manag.* **58**(4), 102544 (2021). <https://doi.org/10.1016/j.ipm.2021.102544>
54. Priyadarshini, S.B.B., Bagjadab, A.B., Mishra, B.K.: A brief overview of natural language processing and artificial intelligence. In: *Natural Language Processing in Artificial Intelligence*, p. 14. Apple Academic Press (2020)
55. Raina, V., Krishnamurthy, S.: Natural language processing. In: *Building an Effective Data Science Practice*, pp. 63–73. Springer, Heidelberg (2022). https://doi.org/10.1007/978-81-322-3972-7_19
56. Ribeiro, M.T., Singh, S., Guestrin, C.: “why should i trust you?”: explaining the predictions of any classifier. In: *KDD 2016*, pp. 1135–1144. Association for Computing Machinery, New York (2016). <https://doi.org/10.1145/2939672.2939778>. Accessed 01 Sept 2022
57. Sai, A.B., Mohankumar, A.K., Khapra, M.M.: A survey of evaluation metrics used for nlg systems. Technical report (2020). Accessed 08 June 2022
58. Salewski, L., Alaniz, S., Rio-Torto, I., Schulz, E., Akata, Z.: In-context impersonation reveals large language models’ strengths and biases. *arXiv preprint arXiv:2305.14930* (2023)
59. Salminen, J., Guan, K., Jung, S.G., Jansen, B.J.: A survey of 15 years of data-driven persona development. *Int. J. Human-Comput. Interact.* **37**(18), 1685–1708 (2021). <https://doi.org/10.1080/10447318.2021.1908670>
60. Salminen, J., Jansen, B.J., An, J., Kwak, H., Jung, S.G.: Are personas done? evaluating their usefulness in the age of digital analytics. *Pers. Stud.* **4**(2), 47–65 (2018). <https://doi.org/10.21153/psj2018vol4no2art737>
61. Salminen, J., Jung, S.G., An, J., Kwak, H., Nielsen, L., Jansen, B.J.: Confusion and information triggered by photos in persona profiles. *Int. J. Human-Comput. Stud.* **129**, 1–14 (2019). <https://doi.org/10.1016/j.ijhcs.2019.03.005>
62. Salminen, J., Jung, S.G., Jansen, B.: Developing persona analytics towards persona science. In: *27th International Conference on Intelligent User Interfaces, IUI 2022*, pp. 323–344. Association for Computing Machinery (2022). <https://doi.org/10.1145/3490099.3511144>
63. Salminen, J., Jung, S.G., Jansen, B.J.: The future of data-driven personas: a marriage of online analytics numbers and human attributes, pp. 596–603. SciTePress, Heraklion (2019). Accessed 22 Aug 2019
64. Salminen, J., Jung, S.G., Jansen, B.J.: Are data-driven personas considered harmful?: diversifying user understandings with more than algorithms. *Pers. Stud.* **7**(1), 48–63 (2021). ISBN: 2205-5258

65. Salminen, J., Jung, S.G., Santos, J., Jansen, B.J.: Toxic text in personas: an experiment on user perceptions. *AIS Trans. Hum.-Comput. Interact.* **13**(4), 453–478 (2021). <https://doi.org/10.17705/1thci.00157>
66. Salminen, J., Mustak, M., Corporan, J., Jung, S.G., Jansen, B.J.: Detecting pain points from user-generated social media posts using machine learning. *J. Interact. Mark.* 10949968221095556 (2022). <https://doi.org/10.1177/10949968221095556>
67. Salminen, J., Rao, R.G., Jung, S., Chowdhury, S.A., Jansen, B.J.: Enriching social media personas with personality traits: a deep learning approach using the big five classes. In: Degen, H., Reinerman-Jones, L. (eds.) *HCI 2020. LNCS*, vol. 12217, pp. 101–120. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50334-5_7
68. Salminen, J., Santos, J.M., Kwak, H., An, J., Jung, S.G., Jansen, B.J.: Persona perception scale: development and exploratory validation of an instrument for evaluating individuals' perceptions of personas. *Int. J. Hum.-Comput. Stud.* **141**, 102437 (2020). <https://doi.org/10.1016/j.ijhcs.2020.102437>
69. Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., Elmqvist, N., Diakopoulos, N.: Grand challenges for HCI researchers. *Interactions* **23**(5), 24–25 (2016)
70. Stephanidis, C., et al.: Seven HCI grand challenges. *Int. J. Hum.-Comput. Interact.* **35**(14), 1229–1269 (2019)
71. Stevenson, P.D., Mattson, C.A.: The personification of big data. In: *Proceedings of the Design Society: International Conference on Engineering Design*, vol. 1. no. 1, pp. 4019–4028 (2019). <https://doi.org/10.1017/dsi.2019.409>
72. Tan, H., Peng, S., Liu, J.X., Zhu, C.P., Zhou, F.: Generating personas for products on social media: a mixed method to analyze online users. *Int. J. Hum.-Comput. Interact.* **38**(13), 1255–1266 (2021). <https://doi.org/10.1080/10447318.2021.1990520>
73. Turing, A.M.: Computing machinery and intelligence. *Mind* **59**(236), 433–460 (1950)
74. Volkova, S., Wilson, T., Yarowsky, D.: Exploring demographic language variations to improve multilingual sentiment analysis in social media. In: *EMNLP 2013*, pp. 1815–1827. Association for Computational Linguistics, Seattle (2013). <https://www.aclweb.org/anthology/D13-1187>. Accessed 27 Dec 2019
75. Wood-Doughty, Z., Shpitser, I., Dredze, M.: Generating synthetic text data to evaluate causal inference methods. Technical report (2021). <http://arxiv.org/abs/2102.05638>. <https://doi.org/10.48550/arXiv.2102.05638>
76. Zhang, X., Brown, H.F., Shankar, A.: Data-driven personas: constructing archetypal users with clickstreams and user telemetry. In: *CHI 2016*, pp. 5350–5359. ACM, San Jose (2016). Accessed 04 Nov 2017
77. Zhu, H., Wang, H., Carroll, J.M.: Creating persona skeletons from imbalanced datasets - a case study using U.S. older adults' health data. In: *DIS 2019*, pp. 61–70. ACM, New York (2019). <https://doi.org/10.1145/3322276.3322285>. Accessed 01 Dec 2021