Transparent Persona Generation With LLMs: An Evidence-based and Traceable Method for User-centred Design

Bojan Blažica[1,*,†], Manca Topole[1,†] and Marko Debeljak[1,†]

[1] *Jožef Stefan Institute, Department of Knowledge Technologies, Jamova cesta 39, Ljubljana, Slovenia*

Abstract

Personas are a cornerstone of user-centred design, but traditional methods for developing them are difficult to validate, prone to bias and labour-intensive. Data-driven approaches have improved scalability, but often lack the narrative richness and empathy that make personas effective. We present a methodology that uses large language models (LLMs) to accelerate the creation of personas while underpinning and constraining the results with contextual and empirical data. Our approach emphasises transparency and traceability: each generated persona attribute can be linked to its source material, including project documentation, workshop transcripts, survey results or other contextual corpora. By combining the narrative strengths of LLMs with the rigour of an evidence-based foundation, the method generates personas that are both descriptive and verifiable. We present a five-step workflow methodology: (1) generation of persona candidates from contextual data using LLMs, (2) iterative refinement to ensure representativeness of personas, (3) selection of the most relevant profiles through expert evaluation, (4) design of detailed persona profiles, and (5) enrichment with empirical evidence to ensure traceability and validation. The methodology is illustrated with a case study from the field of soil health, but can also be applied to other design contexts where alignment between different stakeholders is crucial. We argue that this approach positions LLMs not as a substitute for human expertise, but as an accelerator of persona work that improves accountability, reduces bias and facilitates communication in collaborative design processes.

Keywords

personas, large language model, traceability, user-centered design, decision support systems

1. Introduction

Personas are an established method in user-centred product design to capture archetypal user needs, their motivations and challenges in a lively, narrative form [1]. They help designers maintain empathy with users and facilitate communication within design teams, but have been criticised for being difficult to validate and for reflecting the biases of the experts who create them, while also being time-consuming to produce [2]. These limitations have motivated interest in computational and data-driven persona approaches that use analytics and algorithms to speed up the creation of personas and ensure that profiles are based on real user data [3]. While quantitative data offers efficiency and representativeness, it alone cannot capture the attitudes and motivations that give personas depth. Recent work suggests that large language models (LLMs) can complement such approaches by analysing qualitative inputs at scale and enriching data-driven personas with contextual nuance and empathy-inducing qualities [4].

The rapid emergence of LLMs has opened up new possibilities for persona generation. LLMs can create plausible and detailed personas from prompts that are either completely synthetic or based on empirical data [5]. Compared to previous algorithmic methods, LLMs are attractive because of their ability to generate coherent narratives, integrate different attributes and adapt to different design contexts. Evaluations suggest that LLM-generated personas can match or even outperform human-generated personas in terms of perceived consistency, informativeness and credibility [6, 7].

Human-Computer Interaction Slovenia 2025, October 13, 2025, Koper, Slovenia

*Corresponding author.

†These authors contributed equally.

bojan.blazica@ijs.si (B. Blažica); manca.topole@ijs.si (M. Topole); marko.debeljak@ijs.si (M. Debeljak)

0000-0003-4597-5947 (B. Blažica); 0009-0004-5404-7549 (M. Topole); 0000-0003-4334-4295 (M. Debeljak)

https://doi.org/10.26493/978-961-293-559-7.15

However, risks remain: LLM-generated personas are prone to stereotyping and may encode demographic or cultural biases inherited from training data [8, 9]. They often lack provenance and traceability, as links between persona attributes and underlying user data are rarely made explicit [3, 4]. They are also vulnerable to hallucinations, generating attributes not grounded in evidence [8, 9], and often produce oversimplified, generic or stereotypical profiles that miss contextual nuance and lived complexity [9, 2]. As a result, scholars emphasise the importance of combining LLMs with human expertise, using them as accelerators rather than replacements in persona workflows [4].

New research has begun to assess the use of LLMs in persona development. Salminen and colleagues [10] show that most published persona prompts produce concise, structured results rather than richer, empathetic narratives, highlighting the lack of established best practices for prompt design. De Paoli [11] and Schuller et al. [6] further explore prompting strategies such as role-play, one-shot examples and gradual refinement and show how these decisions influence the quality and credibility of the generated personas. This work suggests that synthetic personas can support rapid prototyping, simulation and pilot studies, but also indicates that human expertise and participatory input remain important to ensure their validity, transparency and relevance in HCI[4, 5].

The literature suggests that LLM-based persona generation is a promising but unsettled field. Compared to traditional and data-driven personas, LLMs offer unprecedented efficiency and flexibility, but their reliability and inclusivity remain contested [6, 2]. This evolving debate emphasises the potential of LLMs to reshape persona-based design [11, 3] while highlighting the need for critical reflection on how such tools are embedded in design practices and decision-making processes [5, 12]. A sentiment echoed in relation to LLM based tools in general; [13] propose a roadmap with four common approaches that the community has taken to achieve transparency: model reporting, publishing evaluation results, providing explanations, and communicating uncertainty.Similarly, recent work on multi-layered human-centered explainability frameworks [14] introduces an architecture that combines (1) foundational models with built-in explainability mechanisms, (2) a human-centered explanation layer tailored to users' cognitive load and expertise, and (3) a dynamic feedback loop that refines explanations through real-time interaction. Together, these perspectives highlight that transparency extends beyond technical implementation as it must also account for how people interpret and use explanations.

Among the unresolved challenges highlighted in previous work, our study focuses on traceability - the ability to link persona attributes to concrete sources. We argue that without traceability, LLM-generated personas run the risk of being dismissed as arbitrary, especially in collaborative contexts where they serve as shared artifacts among different partners. To address this problem, we propose a transparent method that combines the narrative capabilities of LLMs with contextual data, enabling the creation of personas that are both descriptive and verifiable. Our aim is to demonstrate this approach using a case from the field of soil health, although the method is also transferable to other fields.

2. Material and Methods

This section describes the process used to develop LLM-accelerated and traceable personas. It is divided into three parts. Section 2.1 outlines the overall methodology, from the initial generation of a broad pool of candidate personas to the refinement of the final, evidence-based profiles. Section 2.2 presents the prompt design and the models used. Section 2.3 describes the data sources used to ground and calibrate the personas.

Together, these steps form a reproducible process in which LLMs are used to accelerate the creation of personas, while contextual and empirical data ensure that the resulting personas remain relevant and credible.

2.1. Methodology

The methodology consists of five main steps that structure the process of creating, selecting and refining personas with LLM support.

2.1.1. Step 1: Creation of candidate personas

- Objective: Create an initial pool of personas based on contextual data.
- Description: An LLM was used to generate a wide range of candidate personas driven by contextual data to provide meaningful starting examples. In our case, this contextual data was project documentation (described in section 2.3.1). However, the method is not limited to documents. Other data sources, such as user-generated content, relevant literature (scientific or non-scientific) or other domain materials could equally serve to guide the generation of initial candidates. - Extend the methodology section to include discussion of traceability not only for workshop data but also for document-based persona generation. At this stage, traceability is not explicitly applied, as the contextual documents serve only to provide general thematic direction and domain framing to the LLM. The goal is to obtain diverse but plausible first drafts rather than to attribute specific persona statements to individual sources; detailed evidence linking and traceability are introduced in later steps.

2.1.2. Step 2: Iterative refinement

- Objective: Ensure representativeness.
- Description: The candidate pool was iteratively expanded to ensure representativeness for the domain and the task at hand. In our case, the development of a soil health dashboard meant ensuring coverage of all user types (land steward, private actor, public actor, knowledge actor, citizen), land uses (agriculture, forestry, urban) and scales (local, landscape, regional, European). This was achieved through manual iteration in chat sessions with the LLM, for example, by asking for "more public actor persona examples," or through domain-specific prompts such as "what about NGOs and activists as examples of citizen personas".

2.1.3. Step 3: Internal selection

- Objective: Identification of the most relevant personas for the design task.
- Description: Domain experts and project team members were asked to rank the personas created according to their relevance to the project objectives, using a scale from 1 (not at all relevant) to 5 (highly relevant). This ranking helped to focus attention on the most important personas. The top 10% of personas, as determined by ranking were selected as the narrowed pool to be further developed in the next step.

2.1.4. Step 4: Drafting persona profiles

- Objective: Create a set of persona drafts that reflect key contextual dimensions (scale, land use, stakeholder role).
- Description: Persona drafts were developed from the narrowed pool obtained in Step 3 (Section 2.1.3). Additional personas were introduced where stakeholder types were missing based on editorial decision and what is needed by the project. We've structured persona drafts around the Soil Health Dashboard's main functionalities (monitoring, assessing, and managing soil health) and included stakeholder-relevant sections (e.g., context, primary goals and tasks, pain points). We've further elaborated drafts with the help of iterative LLM chat sessions, which resulted in additional relational elements such as hobbies, motivational quotes, and persona images. These elements were included to support relatability in design and use contexts, while remaining secondary to the functional sections.

2.1.5. Step 5: Evidence-based refinement

- Objective: Refine the persona profiles using empirical data from workshops to ensure that they reflect the actual needs, challenges, goals and priorities of soil health stakeholders.

- Description: The persona drafts created in the previous step (Section 2.1.4) were systematically enriched with empirical data from stakeholder workshops. Integration was carried out using LLM, with each refinement cycle combining a detailed and structured prompt with the draft persona, five Excel files with filtered stakeholder- and land-use-specific data (as described in 2.3.2) and a sample persona file to ensure consistent formatting and structure. The resulting persona profiles were then manually reviewed for internal consistency, clarity and logical alignment. This process ensured that each persona could be traced back to empirical evidence while maintaining coherence.

2.2. LLM Use and Prompts

At the time of developing this methodology, it was not possible to feed proprietary data directly into the OpenAI interface; therefore, we used an internal deployment of Onyx (formerly Danswer) for Step 1 (Section 2.1.1) and Step 2 (Section 2.1.2. Onyx is a knowledge-augmented chat tool that combines contextual retrieval with LLM chat integration, allowing documents to be indexed locally and used in conversation with an LLM. The project documentation was entered into Onyx and connected to GPT-4 via an API. In all steps, LLMs were used to generate text and, in the final step, also to generate images for persona profiles. All interactions were carried out either via a chat interface, although automation via API would have been possible.

In Step 1 (Section 2.1.1) and Step 2 (Section 2.1.2), GPT-4 was prompted through Onyx to generate candidate personas in open dialogue sessions. Each request asked for the following information: user, user type, land use, scale, and persona description. In our case, the contextual data came from the project documentation (see Section 2.3.1). A single researcher conducted these sessions, iteratively expanding the pool to improve the coverage of user types, land uses and scales.

Persona drafts were developed through iterative chat sessions with GPT-4o, performed directly in the OpenAI interface, which by then allowed direct grounding with user-supplied data. Prompts were applied in a structured yet conversational manner to elaborate persona drafts while maintaining alignment with the predefined framework (candidate persona pool) established in earlier steps and the project documentation.

In Step 5 (Section 2.1.5), persona refinement was carried out with GPT-5 (auto setting). Here, the process of prompt development was more elaborate: several cycles of testing and comparison were used to arrive at a stable version that provided consistent, evidence-based results. The final prompt included explicit instructions on how to map empirical workshop data to persona sections, how much data to include, and how to cite individual stakeholder statements to ensure traceability. Additional instructions refined during testing specified selection boundaries, narrative style, and formatting rules to keep the profiles consistent and concise. This final version was applied to all completed personas, and the full prompt text is included in Appendix A.

Default parameters for seed and temperature parameters where used as it is not possible to set them in the web-based interface. Comparing various LLMs for persona generation was not in the scope of this work.

2.3. Data

To mitigate the limitations of LLM-based personas mentioned in the Introduction, we grounded our persona development in contextual data. This grounding occurred in two stages: first, a light, broad contextual seeding in the initial persona generation; then later, a deeper grounding for the final profiles using domain-specific empirical data.

2.3.1. Preliminary grounding via contextual documents

In the initial phase, we engaged in conversational prompting with the LLM while supplying project documentation as context. In our case, the project is the EU funded research project BENCHMARKS (https://soilhealthbenchmarks.eu/), which aims to develop soil health indicators across different land uses, stakeholders, and spatial scales. The intended personas are to be used in user-centered design and development of a Soil Health Dashboard: a tool that integrates local, landscape, regional, and European data (e.g., statistical indicators, remote sensing, soil sampling) to help users choose indicators suited to their objectives, land use, and context. By feeding the model the project's DoA (description of the action, also description of work - DoW or Annex 1) document (which covers objectives, scientific rationale, work packages, partner roles, and more), we enabled the LLM to anchor persona suggestions in the project's domain, goals, and structure.

2.3.2. Final profile calibration data

The project workshops involved over 500 stakeholders across 11 European countries and 9 languages. Each of the 24 local workshops combined plenary and group sessions, where participants worked with sticky notes and posters to record inputs. Activities were designed to capture stakeholder perspectives on soil health. This empirical input provided stakeholder-grounded information essential for developing a Soil Health Dashboard that is both relevant and useful for its intended users. While the workshops were not originally conducted for persona development, their results proved highly suitable for this purpose, as integrating them into this process further ensured that the final tool addresses stakeholder needs and priorities.

The workshops comprised different activities in which stakeholders explored soil health challenges, their objectives and management practices, needs, and prioritized soil functions. Outputs from these activities were translated into English by local facilitators and transcribed into standardized Excel sheets. We extracted data subsets corresponding to each persona's stakeholder and land-use type. This resulted in four datasets linked to personas, namely their needs, objectives, practices, and prioritized soil functions. Challenges were also extracted but could not be linked to specific stakeholder types due to the structure of the workshops. Each dataset was stored in a dedicated Excel file, forming the empirical evidence base for Step 5 (Section 2.1.5) of our persona development methodology.

3. Results

The proposed methodology yielded intermediate and final results. Intermediate outputs show the evolution from a large pool of seed personas to a smaller set of prioritized candidates. The final output consists of fully elaborated and traceable persona profiles, calibrated with contextual and empirical data.

3.1. Seed personas

After the initial generation and iterative refinement (Sections 2.1.1 and 2.1.2), we obtained a list of 100 seed personas (Table 1). The list was representative of our domain and balanced across land use, user type, and scale (Table 2).

3.2. Ranked personas

In the internal selection (Section 2.1.3), eight project participants rated each persona on a scale from 1 (not relevant) to 5 (very relevant). Based on the average scores, the top 10% were selected as candidates for further development (Table 3). The average score was 3.50 with a standard deviation of 0.92. After this step, some of the user types (private actors, citizens) and scales (landscape) were not represented, so an editorial decision had to be made as to whether or not they should be included in the final selection for profile development.

3.3. Final personas

The final stage produced a set of 9 calibrated personas representative of the project. Each profile was generated in two steps: first, a draft including general personal data (Section 2.1.4), and second,

Table 1
Excerpt of 5 rows from the full set of 100 seed personas showing the simplified version of personas used in the first step of persona creation.

n	User Type	User	Land Use	Scale	Persona Description
1	Land Steward	Farmer	Agriculture	Local	John Doe; 45, Male; Goal: Improve crop yield; Hobbies: Gardening, Reading
18	Private Actor	Agribusiness Manager	Agriculture	Landscape	Emily Brown; 42; Female; Goal: Integrate soil health into business practices; Hobbies: Traveling, Cooking
37	Public Actor	Regional Policymaker	Agriculture	Regional	Anna Black; 40; Female; Goal: Implement soil health policies; Hobbies: Volunteering, Yoga
59	Knowledge Actor	Soil Scientist	Agriculture	Europe	Dr. Alan Brown; 60; Male; Goal: Research soil health indicators; Hobbies: Writing, Lecturing
82	Citizen	Citizen Scientist	Agriculture	Local	Alex Grey; 30; Non-binary; Goal: Contribute to soil health data collection; Hobbies: Community gardening, Blogging

Table 2
Summarised distribution of personas across land use, user type and scale, showing balanced coverage

User Type		Land Use		Scale	
Land Steward	17	Agriculture	24	Local	42
Private Actor	19	Forestry	20	Landscape	14
Public Actor	22	Urban	23	Regional	17
Knowledge Actor	23	All	33	Europe	27
Citizen	19				

refinement and calibration with empirical data (Section 2.1.5). This process resulted in a persona profile containing two distinct categories of information:

- **General persona data** (personal details that set the broad context and help readers empathize with the persona): name, age & nationality, primary scale & land use, role & production system, farm & climate context, tech proficiency & access, decision horizon & authority, hobbies and motivational quote.
- **Empirical, domain-specific data** (attributes grounded in empirical evidence from stakeholder workshops): primary goals, pain points, needs, dashboard expectations and scenario of use.

To illustrate, *Ana Martínez Sánchez* (Figure 1), an organic arable farmer from Castilla-La Mancha from Spain, was developed into a complete profile. The following examples illustrate how empirical

Table 3
The short list of personas with their average vote, standard deviation and overall ranking. Ranking resulted in missing stakeholder types in the selection calling for an editorial decision to complete the list.

n	User Type	Land Use	Scale	Min	Max	Avg	Std	Rank
3	Land Steward	Agriculture	Local	4	5	4,75	0,43	4
8	Land Steward	Agriculture	Local	4	5	4,88	0,33	2
9	Land Steward	Forestry	Local	4	5	4,75	0,43	4
37	Public Actor	Agriculture	Regional	3	5	4,50	0,71	10
38	Public Actor	Agriculture	Regional	4	5	4,75	0,43	4
39	Public Actor	Agriculture	Europe	4	5	4,88	0,33	2
41	Public Actor	All	Europe	3	5	4,50	0,71	10
52	Public Actor	Forestry	Regional	4	5	4,50	0,50	9
54	Public Actor	Urban	Europe	3	5	4,63	0,70	8
59	Knowledge Actor	Agriculture	Europe	5	5	5,00	0,00	1
60	Knowledge Actor	Agriculture	Europe	4	5	4,75	0,43	4
64	Knowledge Actor	All	Europe	3	5	4,50	0,71	10

evidence shaped her persona:

- **Primary goals** – safeguard soil fertility and resilience under dryland organic farming, grounded in workshop objectives on soil fertility and nutrient cycling.
- **Pain points** – recurrent drought limiting yields, linked to challenges identified by workshop participants.
- **Needs** – validated protocols to measure soil improvements, based on stakeholder needs for objective, standardized assessment methods.
- **Dashboard expectations** – scenario simulations comparing crop rotations and cover crops, connected to workshop objectives and practices on nutrient cycling and water management.

A full version of Ana's persona, including both general and empirical elements with explicit source references, is provided in Appendix B.

4. Discussion

The aim of this study was to address several problems in the development of personas: lack of tractability, bias, oversimplification, limited empirical foundation, and the challenge of balancing narrative richness and methodological rigour. Our approach addresses these issues by combining the generative power of LLMs with structured contextual and empirical data and iterative expert supervision. By anchoring persona characteristics in data, including project documentation (contextual data) and data from stakeholder workshops (empirical data), we enhanced transparency and credibility while reducing the risk of arbitrary or hallucinated results. Each persona characteristic could be linked to a verifiable source, strengthening accountability in collaborative design processes.

A second important innovation of the presented research work lies in the structured, multi-step methodology that integrates generation, refinement, validation and evidence-based calibration into a single workflow. Unlike traditional persona methodologies, which are often static and subjective, our approach accelerates the creation of personas while maintaining contextual relevance, stakeholder alignment and motivational depth. Compared to existing data-driven approaches that focus on scalability at the expense of empathy and tone, our method creates personas that are both representative and meaningful for user-centred design. Recent research has begun to tackle aspects of this challenge. For example, De Paoli [11] proposed iterative workflows combining LLM-based analysis and persona drafting, Shin et al. [4] studied how to distribute tasks between humans and AI in persona generation, and Jung et al. [3] introduced structured data-to-persona pipelines to improve transparency. These efforts, however, focus on isolated stages, whereas our approach integrates generation, refinement, validation, and calibration into a single workflow. This structured methodology is not only novel in itself but also lays the foundation for practical integration into decision support systems (DSSs).

The integration of personas into the development of DSSs is expected to bring several practical benefits. We anticipate that personas can act as a bridge between technical development and user expectations by helping to recognise inconsistencies early on and guiding adjustments to the design of the user interface and data presentation. They are also expected to facilitate communication between the various stakeholders and support the negotiation of trade-offs between complexity and usability. Most importantly, we see the potential for personas to serve as recurring reference points throughout the project, influencing design priorities and evaluation criteria. In future work, we plan to evaluate whether personas indeed evolve from static artifacts into dynamic tools that support decision-making and increase user confidence during the continued development of the DSS.

Figure 1: Persona illustration of Ana Martínez Sánchez, an organic arable farmer from Castilla-La Mancha, Spain. The image was generated with an LLM and is intended to support empathy and visualization rather than depict a real individual.

However, our approach is also associated with some limitations. The quality and representativeness of the contextual data had a strong influence on the results. Biases or gaps in the source material were sometimes reproduced, even if they were more transparent. Human oversight also remained crucial. Careful, timely development, iterative testing and expert validation were required to maintain contextual accuracy, which increased resource requirements. Nevertheless, each validation cycle improved the quality of the persona, but also increased the complexity, and the human judgment required inevitably led to a degree of subjectivity. This reflects the fact that incorporating human contexts into design artifacts requires interpretive involvement in addition to automatic generation.

These considerations highlight both the potential and the current limitations of LLM-based persona generation and point to future methodological refinements as briefly outlined in the conclusions.

5. Conclusions and future work

This study presents a transparent, evidence-based methodology for creating user personas that combines the narrative capabilities of LLMs with contextual and empirical data. The approach addresses key limitations of existing methods by improving tractability, reducing bias, and preserving narrative richness while accelerating the development of DSSs. Applied to the soil health domain, the method shows potential to strengthen user-centred development of DSS by improving stakeholder alignment and system relevance. We expect that, beyond their direct design function, personas can also serve to connect project partners by providing a shared understanding of stakeholder groups and by linking different project objectives to the same shared personas.

Future work will evaluate these expected benefits in practice and will aim to automate provenance tracking, integrate dynamic data, and include additional stakeholders to improve the scalability and impact of persona-driven design. Replicating the workflow across multiple LLM families and diverse application domains will allow us to validate its consistency, broaden its relevance, and strengthen its contribution to transparent, evidence-based persona generation.

In addition, drawing on established persona evaluation methodologies such as the Persona Perception Scale [15], future studies should compare evidence-grounded LLM–human personas with human-only and LLM-only baselines to assess perceived realism, usefulness, and empathy.

In line with earlier work comparing algorithmically generated and human-crafted personas [**?**], future studies will also conduct comparative evaluations contrasting fully synthetic LLM personas, the proposed evidence-grounded LLM + human workflow, and human-only baselines to understand how personas are perceived by persona users and how likely they will adopt them. Finally, since transparency is central to our contribution, further research should also explore more objective and automated methods for verifying footnote-to-source correctness, as the current manual review process, while effective, remains labour-intensive.

Acknowledgments

This study was supported by the Slovenian Research Agency (grant P2-0103) and by the European Union's Horizon Europe research and Innovation program, under Grant Agreement: 101091010, Project BENCHMARKS. We gratefully acknowledge the partners of the BENCHMARKS project and all participants of the stakeholder workshops. Special thanks to Fabio Volkmann, who led the workshops where the data were collected.

Declaration on Generative AI

During the preparation of this work, the author(s) used OpenAI GPT-5 in order to: assist with literature review, text drafting, rephrasing, and grammar correction. After using this tool, the author(s) reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] T. Miaskiewicz, K. A. Kozar, Personas and user-centered design: How can personas benefit product design processes?, Design studies 32 (2011) 417–430.
[2] C. Lazik, C. Katins, C. Kauter, J. Jakob, C. Jay, L. Grunske, T. Kosch, The impostor is among us: Can large language models capture the complexity of human personas?, arXiv preprint arXiv:2501.04543 (2025).
[3] S.-G. Jung, J. Salminen, K. K. Aldous, B. J. Jansen, Personacraft: Leveraging language models for data-driven persona development, International Journal of Human-Computer Studies 197 (2025) 103445.
[4] J. Shin, M. A. Hedderich, B. J. Rey, A. Lucero, A. Oulasvirta, Understanding human-ai workflows for generating personas, in: Proceedings of the 2024 ACM Designing Interactive Systems Conference, 2024, pp. 757–781.

[5] M. Prpa, G. Troiano, B. Yao, T. J.-J. Li, D. Wang, H. Gu, Challenges and opportunities of llm-based synthetic personae and data in hci, in: Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing, 2024, pp. 716–719.
[6] A. Schuller, D. Janssen, J. Blumenröther, T. M. Probst, M. Schmidt, C. Kumar, Generating personas using llms and assessing their viability, in: Extended abstracts of the CHI conference on human factors in computing systems, 2024, pp. 1–7.
[7] J. Salminen, C. Liu, W. Pian, J. Chi, E. Häyhänen, B. J. Jansen, Deus ex machina and personas from large language models: Investigating the composition of ai-generated persona descriptions, in: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, 2024, pp. 1–20.
[8] H. A. Haxvig, Concerns on bias in large language models when creating synthetic personae, arXiv preprint arXiv:2405.05080 (2024).
[9] L. Sun, T. Qin, A. Hu, J. Zhang, S. Lin, J. Chen, M. Ali, M. Prpa, Persona-l has entered the chat: Leveraging llms and ability-based framework for personas of people with complex needs, in: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 1–31.
[10] J. Salminen, D. Amin, B. Jansen, Using ai for user representation: An analysis of 83 persona prompts, arXiv preprint arXiv:2508.13047 (2025).
[11] S. De Paoli, Improved prompting and process for writing user personas with llms, using qualitative interviews: Capturing behaviour and personality traits of users, arXiv preprint arXiv:2310.06391 (2023).
[12] R. Y. Pang, H. Schroeder, K. S. Smith, S. Barocas, Z. Xiao, E. Tseng, D. Bragg, Understanding the llm-ification of chi: Unpacking the impact of llms at chi through a systematic literature review, in: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 1–20.
[13] Q. V. Liao, J. W. Vaughan, Ai transparency in the age of llms: A human-centered research roadmap, arXiv preprint arXiv:2306.01941 (2023).
[14] C. De Silva, T. Halloluwa, D. Vyas, A multi-layered research framework for human-centered ai: Defining the path to explainability and trust, arXiv preprint arXiv:2504.13926 (2025).
[15] J. Salminen, J. M. Santos, H. Kwak, J. An, S.-g. Jung, B. J. Jansen, Persona perception scale: development and exploratory validation of an instrument for evaluating individuals' perceptions of personas, International Journal of Human-Computer Studies 141 (2020) 102437.

Appendices

A. Prompt used for final persona calibration

```
0.      Canonical inputs (will be provided)
        Prompt document: the canonical instructions. If anything is left unclear,
    consult me before making a decision when amending the persona profile.
        Profile template document (Word, finalized persona; e.g., Ana): copy fonts,
     sizes, spacing, section sequence, bullets, headings from this template.
    Persona draft document (Word) for the new persona (includes persona type which
    specifies personas stakeholder type and land use type + includes information
    section (from the Benchmarks milestone reports) which sets the framework for
    the persona profile)
    Five Excel workshop files filtered to the personas stakeholder type  land use:
    i.  Challenges (land-use only)
    ii. Needs (stakeholder  land use)
    iii.        Objectives (stakeholder  land use)
    iv. Practices (stakeholder  land use)
    vi. Priority soil-function votes (stakeholder  land use)

1.      File-to-section mapping (How to integrate workshop data to user persona
    profile)
```

Challenges Pain points
Needs Needs (general requirements); also inform Dashboard expectations
Objectives Primary goal, Dashboard expectations
Practices Dashboard expectations, Scenario of use
Soil-function votes Weave naturally into Primary goal / Dashboard expectations / Scenario of use (not a separate section)

2. Pre-processing & checks
i. Open the Persona draft and identify sections:
Keep (in this order) unless contradicted: title (name and a few main details), persona type, Inspiration, Age & nationality, Primary scale & land use, Context (title of this section differs from persona to persona, depending on persona type), Tech proficiency & access, Decision horizon & authority, Hobbies, Motivational quote
Replace (or blend per step 6): Primary goals, Pain points, Needs, Dashboard expectations, Scenario of use
ii. Open Excels and confirm ID columns:
challenge_id, need_id, objective_id, practice_id.
In Soil-function document use rank_in_slice column and soil_function_name; treat rank as the canonical identifier.
iii. Multi-item rows: if a cell contains multiple statements (newlines/ semicolons), treat each as a separate statement unless clearly one phrase. You will cite only the fragment used later.
iv. Slice correctness: only use rows matching the personas stakeholder type land use (for Challenges: land use only).

3. Selection heuristics (what to keep)
Objectives: choose the 2 to 4 most relevant to the personas ambition and the product (dashboard).
Pain points: 3 to 4 max, concrete and role-true (systemic barriers for public/ knowledge actors; on-farm constraints for land stewards).
Needs: focus on dashboard-relevant needs (analytics, indicators, comparability, reporting, scenarios, data/monitoring, evidence flows), possibly use others as well for broader context
Practices: pick 1 to 3 that the dashboard can model/compare/monitor.
Soil functions: take the top 3 by rank_in_slice for this personas slice.

4. Integration style (how to write)
Bold any sentence/phrase grounded in workshop data.
Blend naturally: do not say workshops say. Let the narrative sound like the persona; the footnotes carry the traceability.
Soil functions in text: refer to them naturally (e.g., habitat provision, carbon & climate regulation, water regulation), without explicit (1st/2nd/3rd priority) unless the request explicitly asks for showing ranks.
Footnote(s) will still point to Soil function (priority X) in the Evidence Table.
Bullets only for Pain points; all other sections are short narrative paragraphs.

5. Superscript footnotes (no brackets in text)
Footnotes are in the form of superscript numbers
For new workshop-grounded statements you add, insert a new superscript number.
Footnotes must be numbered sequentially across the whole profile

6. Section-by-section procedure
i. Primary goal (blend or replace):
- If the draft goal fits the Inspiration, blend workshop objectives/soil-function priorities naturally into 2 to 3 sentences.
- If misaligned, replace with a short narrative merging: 2 to 3 Objectives + (naturally referenced) soil functions + (optionally) 1 dashboard-relevant Need.
- Keep to ~70 to 120 words max.
ii. Pain points:
- 3 to 4 bullets from Challenges, concise and role-true.
- Each bullet grounded in a challenge row (footnote).

iii. Needs:

- 1 short paragraph (2 to 4 sentences), focus on dashboard-relevant needs (3 core ideas).
- Examples: indicators/thresholds, scenario analysis, monitoring/benchmarking, data comparability, evidence export/sharing.

iv. Dashboard expectations:

- 1 short paragraph (3 to 4 ideas), combining Objectives + Practices + (naturally referenced) soil functions.
- Speak in terms of what the dashboard should enable (compare, simulate, monitor, export).

v. Scenario of use:

- 3 to 5 sentences. A simple story of how the persona interacts with the dashboard, weaving 1 to 2 practices and at least one soil-function aspect (naturally).
- Avoid feature laundry lists.

vi. Other draft sections: keep as is unless contradicted by data.

7. Evidence building (footnotes table)

Create an Evidence Table on a new page after the profile text.
Columns: Footnote | Source | Workshop quotation.
For each superscript number, map to its source:
Objective N quote objective_text.
Need N quote need_text.
Challenge N quote challenge_text.
Practice N quote practice_text.
Soil function (priority X) quote soil_function_name for that rank from the personas slice.
If a cell had multiple statements: include only the fragment used and append (excerpt).
If text was used verbatim: show plain (no label).
Maintain the same numbering order as they appear in the profile.

8. Formatting (must mirror template)

Use the Profile template document for:
Fonts, sizes, line spacing, margins, headings, bullet style.
Section sequence exactly as in the template.
Superscripts for footnote markers (not inline digits).
No duplicate sections: remove any draft versions of the five updated sections.

9. Length control (1-page profile)

The profile text (from Primary goal through Scenario of use) should fit within ~1 page in the templates formatting.
If over limit, trim in this order:

1. Dashboard expectations (compress to 3 ideas)
2. Needs (down to 3 tight items)
3. Primary goal (keep to 2 sentences)
4. Pain points (max 3 bullets)

Do not trim the Evidence Table for length; it belongs after the page break.

10. Quality & consistency checks (before exporting)

No contradictions across sections.
All bolded claims have a footnote.
Footnote numbers are sequential and each appears in the Evidence Table.
Soil functions appear naturally in text; priority is represented in footnotes/table, not parenthetical numbers in narrative (unless explicitly requested).
Needs are dashboard-relevant (funding/bureaucracy excluded by default; minor policy context allowed for public actors).
Fits ~1 page profile text in template formatting.

11. When to ask

Ask the user before proceeding if:
The prompt doc and this instruction conflict and its unclear which rule to apply.

```
    The personas slice (stakeholder  land use) is ambiguous or missing in an Excel.
    An objective/need/practice cannot be matched to a credible row.
    Fitting to 1 page would require removing substantive content.
12.     Output
Word document (.docx) that:
- Mirrors the template formatting and section order,
- Contains the amended profile with bold workshop-grounded statements and
    superscript footnotes,
- Ends with an Evidence Table mapping each footnote to the exact workshop quotation
     (or (excerpt)).
```

B. Final persona profile with traceable data

Name	Ana Martínez Sánchez
Age and nationality	44, Spanish (Castilla-La Mancha)
Primary scale and land use	Agriculture, Local (2̃0 ha certified organic arable land)
Farm and climate context	Mediterranean climate — hot, dry summers; mild, wetter winters; occasional intense autumn rains typical of Spain's central plateau
Role and production system	Owner-manager of a 20 ha rainfed (dryland) organic farm. Three-year rotation – year 1: barley for grain; year 2: chickpeas for export; year 3: vetch/cover crops grown as green manure and grazed by a neighbour's ewes. Production relies solely on seasonal rainfall (no irrigation), shallow non-inversion tillage, and composted sheep manure. Her aim is to conserve soil moisture, build organic matter, and keep weeds in check with mechanical tools rather than herbicides.
Tech proficiency and access	Basic desktop PC in her home office; smartphone in the field; intermittent 4G — prefers offline ready exports.
Decision horizon and authority	Seasonal; full decision power over field operations and budget.
Hobbies	Runs a stall at the village cooperative market; experiments with heirloom tomato varieties in her kitchen garden; makes preserves with her two teenage children; and occasionally hikes with friends in the nearby sierras.
Motivational quote	"Take care of what takes care of you — soil, people, and community."
Primary goals	Ana's main ambition is to safeguard soil fertility and resilience under dryland organic farming, while staying economically viable and compliant with organic certification[1]. She wants to compare how practices such as cover crops, shallow tillage, and compost affect yields, soil organic matter, and erosion control[2]. Maintaining soil functions like nutrient cycling[3] and water regulation and purification[4] is central to her vision of long-term stewardship.
Pain points	Recurrent drought limiting yields[5]; Erosion susceptibility on some fields[6]; Administrative burden pulling time from field management[7]
Needs	She values hands-on, validated guidance she can trust. She needs objective, validated measurement protocols to assess soil improvements[8], plus stronger connections between farmers and researchers for practical problem-solving[9]. Peer learning matters: exchange of good practices and materials with nearby farmers[10], and active spaces to talk through needs and experience[11] help turn ideas into confident action.

Footnote	Source ID	Workshop transcript quote
1	Objective 4	Primary production, profitability, crisis preparedness, intergenerational (excerpt)
1	Objective 48	Enrich the soil by improving the soil content of NPK, $CaCO_3$, SO_3 and trace elements, organic matter (excerpt)
2	Objective 2	In the nutrient cycling, the goal is as closed cycles as possible (excerpt)
2	Objective 3	From land drainage to water management (excerpt)
3	Soil function (priority 1)	Nutrient cycle
4	Soil function (priority 2)	Water regulation and purification
5	Challenges 5, 39	Drought
6	Challenge 31	Acidity and erosion susceptibility (soil properties)
7	Challenge 56	The burden of bureaucracy and administration takes away resources and resources from managing the country
8	Need 59	Objective protocols are validated for measuring soil health, Present and disseminate the results obtained on experimental sites for the use of organic waste products, accompany visits and tests on soil quality
9	Need 52	Interconnection offer between farmers and research to share experiences on agricultural techniques and carry out activities with these audiences (excerpt)
10	Need 47	Promotion of a reception area on the Farm for inter-stakeholder meetings (excerpt)
11	Need 10	Active interaction, talk about needs and experiences (excerpt)
12	Objective 2	In the nutrient cycling, the goal is as closed cycles as possible (excerpt)
12	Objective 3	From land drainage to water management (excerpt)
12	Objective 4	Primary production, profitability, crisis preparedness, intergenerational (excerpt)
13	Objective 48	Enrich the soil by improving the soil content of NPK, $CaCO_3$, SO_3 and trace elements, organic matter (excerpt)
14	Practice 4	Include grasses/clovers in the crop rotation (excerpt)
14	Practice 5	Diversifying crop rotations (excerpt)
15	Practice 4	Include grasses/clovers in the crop rotation (excerpt)
16	Practice 1	Diverse crop rotation depending on the situation (excerpt)

Dashboard expectations

She expects the dashboard to integrate field data with clear, easy-to-interpret soil health indicators (e.g., organic matter, erosion risk)[8], and to simulate scenarios comparing different crop rotations, cover crop mixes, or tillage intensities[12]. It should highlight the contribution of practices to water retention and nutrient availability[13], and allow her to download concise, certification-ready reports[14].

Scenario of use

After harvest, Ana logs her tillage operations and uploads soil organic carbon results. The dashboard visualises the organic matter trend and alerts her to compaction and erosion risk[15]. Curious about alternatives, she runs a scenario comparing barley–chickpea with a more diverse rotation including oats and lupins[16], seeing how soil functions like water holding capacity respond. She exports a PDF summary to share with her farm advisor before deciding on next season's rotation.