

Leveraging AI Toward the Development of Vector Personas for UX Research

Claire Lauer

Professor
College of Integrative
Sciences and Arts, Arizona
State University
Tempe, AZ

Danielle Storey

Associate UX Researcher
College of Integrative
Sciences and Arts, Arizona
State University
Tempe, AZ

Romit Soley

Graduate Research
Assistant
College of Integrative
Sciences and Arts, Arizona
State University
Tempe, AZ

Abstract

Personas, which represent user groups with specific needs and behaviors, are crucial in user experience design. High-quality personas are based on robust user data; otherwise, they may cause more problems than they solve. Generative AI, like ChatGPT, shows potential in enhancing persona development by integrating diverse user data sources. However, its reliability is questionable due to its lack of human contextual understanding and real-life experience. This paper introduces the concept of vector personas, a method that combines AI with traditional qualitative techniques to develop user personas. Our UX team tested our approach by comparing personas created through traditional methods with those personas generated using AI. We employed a structured comparison and triangulated data from user feedback. Our study highlights the benefits and limitations of using AI in persona development. Although AI can enhance efficiency and offer new insights, it cannot replace the depth of understanding provided by real-life user-driven research. The vector persona process optimizes persona development by balancing time investment and information quality through integrating AI outputs with human insights to create comprehensive and accurate user representations.

Keywords

vector personas, persona, user experience, large language models, artificial intelligence, AI, OpenAI, ChatGPT, Claude, Gemini, qualitative research, methodology



Introduction

Personas are carefully crafted representations of user groups that UX professionals design to help visualize the users as real individuals with unique needs, preferences, goals, skills, and contexts of use (Blomquist & Arvola, 2002; Matthews et al., 2012; Nielsen, 2013; Pruitt & Grudin, 2003). The effectiveness of personas depends on reliable user data; poorly researched or assumption-based personas can exacerbate issues instead of solving them (Brumberger & Lauer, 2020).

Generative artificial intelligence (AI) tools like OpenAI™ ChatGPT, Anthropic® Claude™, and Google® Gemini™ offer potential for improving persona creation by leveraging user data from a broader range of sources, such as the internet and social media. However, their reliability as accurate data sources remains uncertain (Walters & Wilder, 2023) because generative AI lacks the ability to fully understand context or provide authentic examples of human experiences.

This paper introduces the concept of vector personas as both a methodology and an end product of integrating ChatGPT and other large language models into the persona development process. This methodology draws inspiration from the vector databases underlying GPT technology, which emphasizes a staged approach to leveraging the strengths of both AI and traditional qualitative approaches into persona development. Our UX team designed the vector persona process to create more holistic representations of users by combining the extensive knowledge base of GPT (Nißen et al., 2022) with human insights, contexts, constraints, and priorities.

This paper discusses the results of our empirical study, in which we:

- interviewed members of two target audiences and developed two personas based on those interviews,
- replicated those human interviews by querying generative AI systems with the same questions and built a second set of personas based on the AI interviews,
- triangulated our assessment of the quality of the interview responses by asking members of our target audiences to assess the authenticity of both the human-informed and AI-informed personas (without revealing that either had been built using AI-supplied data), and
- integrated the target audience feedback with both the human-informed and AI-based personas to develop our final vector personas that capitalized on the strengths of each.

Throughout this paper, we present the key elements of our vector persona development process, provide examples of vector personas as end products, and provide empirical evidence highlighting the benefits and limitations of using AI to augment traditional persona development.

Literature Review

As researchers have investigated the role of AI in research processes (Hedquist et al., 2024; Morgan, 2023; Stige et al., 2024), AI-informed personas have become a promising methodological innovation. AI can support persona development by integrating vast amounts of data, such as social media activity, online behavior, academic studies, polling data, and demographic details (Hariri, 2023; De Paoli, 2023). For example, Hariri highlighted that leveraging generative AI in persona development can improve efficiency and scalability by automating the creation of descriptive attributes and profile information. Although current research has focused on how AI can support thematic analysis, qualitative analysis, and so forth (Christou et al., 2024; Nguyen-Trung, 2025), persona development has been another underexplored area of qualitative research wherein AI can be leveraged. Ideally, this approach would minimize the need for labor-intensive qualitative research while providing valuable insights into the target consumer demographic (Hariri, 2023).

However, the complete replacement of researcher-generated personas with fully artificial counterparts has posed its own drawbacks (Chapman & Milham, 2006). AI systems such as ChatGPT were designed to generate replies that closely mimic human-like behavior (Godesky, 2023), but it has become crucial to recognize that these systems are inherently non-human entities and have limited ability (Salminen et al., 2020) to convey real-life examples of human

experiences and behaviors. The particular behaviors and experiences of target users are central to effective persona development.

Illustrating these limits further, an exploratory study (Goel et al., 2023) found that human-computer interaction students scored their AI-generated personas lower than their interview-derived personas based on categories such as consistency, completeness, willingness to use, credibility, clarity, empathy, and memorability (Canossa & Drachen, 2009). However, the researchers did identify five ways in which GPT-3™ helped students develop personas, including elaboration, increased realism, increased creativity, increased coherency, and stimulation.

While Goel et al.'s (2023) study was limited to a small number of students scoring their own personas, our study extends Goel's findings by attempting to achieve an integrated process that triangulates between interviews, AI, and researcher scoring that can be applied across multiple types of personas. Specifically, in our study, we examined two categories of personas: simpler personas, that is, personas with high levels of agreement across interviews and high feature consistency across real-life (RL) subjects; and more complex personas, which have much higher variability and complexity across interviews and subjects. Bifurcated, the personas helped us assess where in the process we can leverage AI versus where more human input is necessary. Although Goel et al.'s categories are useful, we found that we were able to develop a vector persona process that prioritizes achieving an optimal balance of time investment and quality of information. We qualitatively assessed the interviews of RL versus AI respondents in addition to collecting qualitative feedback on personas we developed from subject-matter experts who were included in the target user group categories.

Methods

Vector personas represent a process and an end product that achieves an optimal balance between AI and human-generated research, time, and information quality. To assess how to achieve this balance, we designed a study in which we compared personas developed from interviews with seven real-life users, who represented two target user types, with personas developed from interviews with AI chatbots that were asked the same questions (De Winter et al., 2023). We focused on two specific user types for this study because one represented a high level of feature agreement across users, whereas the other represented a high level of feature variability and complexity across users. We focused on these user types to demonstrate what UX practitioners should pay particular attention to during the vector persona process, depending on the kinds of users they are working with.

Project Background

This study originated from a project aimed at evaluating the effectiveness of a Colorado River water forecasting tool. This tool allowed stakeholders to explore data visualizations depicting potential changes in Colorado River water levels under nine different future scenarios. The primary user group for the tool included water professionals and policymakers responsible for shaping water policies in Arizona. The secondary user group included professionals such as realtors and journalists, who communicated the implications of water availability data to the public.

User Interviews and Development of Initial Personas

As part of our research, we interviewed four municipal water managers for our primary user group, who would be able to use the tool to recommend various courses of action to their city councils and other government agencies. And, we interviewed three realtors for our secondary user group, who would be able to use the tool to stay informed about water availability trends, which they could communicate to their clients, many of whom were moving to Arizona for the first time.

To prepare for our interviews, we developed lists of questions to assess users' interest in the scenarios that the tool provided, the likelihood that they would use the tool, and the extent to which they would find its predictive models trustworthy. After interviewing our primary and secondary users, we analyzed their user transcripts and developed several personas that we presented to the tool developers (Figure 1).



Figure 1. Realtor and water manager personas developed after interviews with human realtors and water managers.

AI Interviews

Just before we commenced our UX research for this tool, OpenAI released ChatGPT-3.5. Because the UX development of our tool continued throughout 2024, we decided to check if we could generate the same level of useful responses to our original interview questions from ChatGPT as we did from our human interview respondents. As expected, our interviews with three human realtors yielded consensus across many responses in a way that made persona development easier. However, our experience interviewing human water managers yielded much more varied responses because of their differing experiences with city size, water supply, and the water department structure in which each worked. For the realtor persona, we wanted to check what responses ChatGPT would provide to our interview questions and whether those responses would prove as consistent as those we collected from the three human realtors we spoke with. For the water manager persona, we wanted to see if ChatGPT could be used to augment our original interviews to provide additional context-specific responses representing water managers from more regions around the state.

We decided to interact with GPT as though we were researchers who did not have the time nor resources to talk to actual users and thus would rely on GPT to provide the bulk of our perspective about user habits, priorities, and values. We queried both OpenAI ChatGPT and Anthropic Claude using the original list of interview questions we had developed. We asked each bot to respond to our questions as though it were one of our users. For instance, we would query, "Imagine you are answering as a water manager from [name of city], Arizona [...]," and we would then ask one of the questions, which were typically about where they got their data about water, how they made decisions at the municipal level about water policy, and what tools

they currently used. For the realtor we would query, "Imagine you are a realtor from [name of city], Arizona [...]," and then we asked questions about what they knew about water and drought in the state, where they currently got their information, what their clients asked them about water, what websites they used to find information, and so on. We included a job title and geographical location to see the extent to which GPT modified its response if the geographic location changed. We asked about several different locations around the state that we knew would typically engender different responses. This was especially important for the water manager interviews because of how much variation there can be between where municipalities get their water and how drought affects how they manage it. During our initial exploration, through this process, we found that there were no discernible differences in length, subject matter, or the overall quality of the feedback provided by Claude versus GPT. Thus, we decided to continue with GPT only due to our familiarity with the software from prior projects.

In our study, one author oversaw the water manager interviews, and one author oversaw the realtor interviews. As we queried GPT, we asked follow-up questions for more context or specificity. For instance, we might have asked, "Can you be more specific about [...]" or "Can you tell me more about [...]." We also asked follow-up questions to help us prioritize the topics or information GPT provided because GPT tended to blanket us with long-form responses. Table 1 shows examples of our queries to human subjects and GPT.

Table 1: Realtor Interview Questions and AI Realtor Follow-up Questions

Realtor Interview Questions	AI Realtor Follow-up Questions
What do you know about the Colorado River Crisis in Arizona?	Can you be more specific?
Where do you typically find information about water in Arizona?	Provide more specific examples.
Do your clients ask you about the water shortage? What specifically do they ask you?	What is the most frequent concern you get asked from your clients about water?
What do you tell them?	What is your biggest concern about water and how does it affect your job? [these last two were our attempt to see if GPT could prioritize]
Is the water situation affecting property values?	

* Both human and AI realtors were asked the same realtor interview questions. To clarify AI responses, we asked AI realtors the AI realtor follow-up questions as well. Questions and follow-ups are both listed in the table.

After interviewing our AI realtor and AI water manager, we analyzed their user transcripts and developed corresponding personas (Figure 2). We asked GPT to respond to questions as if it was a realtor and a water manager in Arizona.

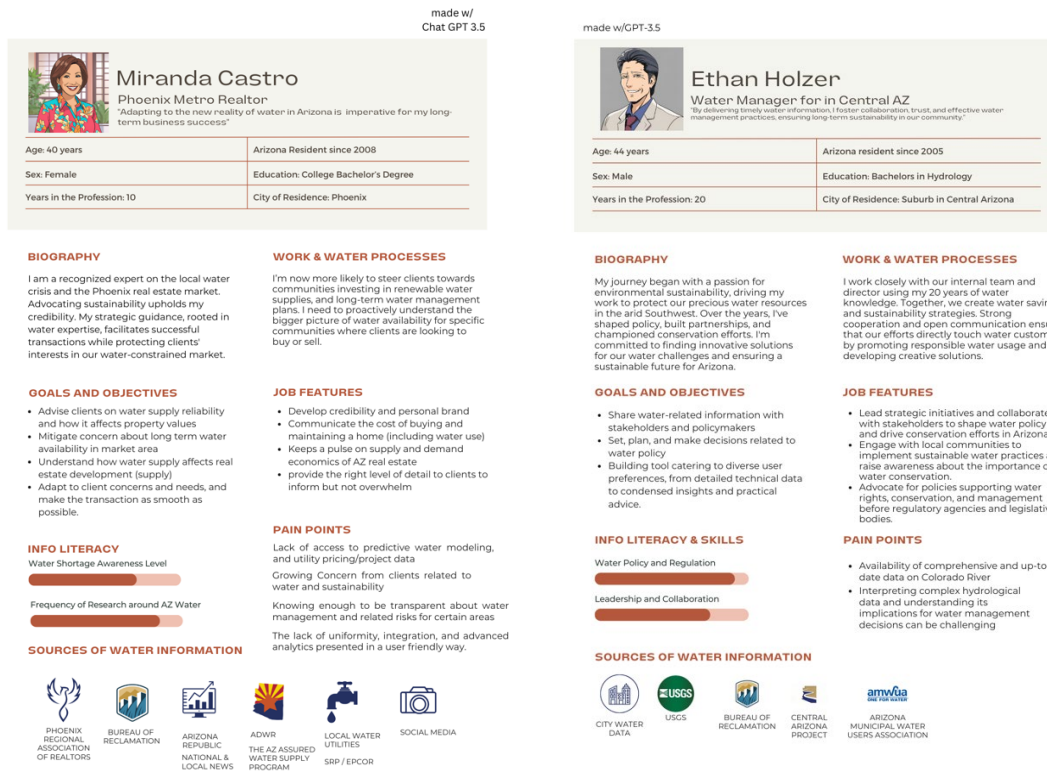


Figure 2. Realtor and water manager personas developed after interviews with OpenAI ChatGPT.

AI Personas and Subject-Matter Expert Evaluation

To validate our representations of our users, we provided both the interview-based personas and the AI-based personas to at least two members of each user group (other human realtors or other human water managers/professionals). We asked them to evaluate each persona using the following instructions.

In the context of making sure developers design as effective and user-friendly a water forecasting tool as possible:

1. *What elements about each persona do you think are accurate? Inaccurate? Helpful? Unhelpful?*
2. *What might you add, remove, or change?*
3. *If you had to score each persona on a scale of 1-10, what score would you give and why?*
4. *Any additional comments?*

Concurrently, each member of the research team recorded their descriptive impressions of all responses provided by both human participants and AI output. We met as a team to discuss and summarize those impressions.

After our own analysis of each interview, and after collecting feedback from subject-matter experts (SMEs), we revised both the human-informed and AI-informed personas into their final vector personas. By working with SMEs, we uncovered any potential inaccuracies or hallucinations, which is a frequently cited concern when working with AI tools (Bender et al., 2021; Omizo, 2024; Sun et al., 2024).

Results

Research Team Analysis of Human and AI Interviews

The research team unanimously agreed on the following summary characterizations of both human and AI interview output. First, we agreed that the human responses, with all their meandering tangents and highly localized examples, simply felt more real than what we got from GPT, which felt more like idealized responses that one might find in a textbook. More significantly, our experiences across both water manager and realtor categories of responses showed that the more that the human responses varied from one another, the less GPT was able to capture that variability. For example, our human water managers from highly variable regions responded much differently from each other about many conditions specific to their regions. Conversely, because the human realtor interviewees provided responses that were more in agreement with each other, the lack of variation was not as noticeable between the human and GPT responses.

Our team also observed that human respondents often use more specialized (and often more tangential) examples to explain their initial response to a question. When asked a question they did not seem to be able to answer easily, our human respondents often relied on anecdotes and name-dropping to bolster their credibility and demonstrate their expertise.

Alternately, GPT always had an excessively comprehensive answer, and it often resorted to providing what seemed like all the possible answers that might be true for a given question. But GPT didn't provide specific examples to illustrate anything in particular, nor a prioritization of which parts of the answer were most important. There were several other examples of how GPT fell short of the kinds of contextualized insight our human respondents provided, including the following:

- For realtors, the GPT persona misrepresented an average realtor's insight into the water and drought situation in Arizona. The persona suggested that realtors were well-versed in the Colorado River shortage. However, the realtors we interviewed did not feel knowledgeable about the situation. GPT responses produced an idealistic representation of a realtor with ample time to educate themselves on water. We know from our interactions with human realtors that their time for reading and researching is limited because time is monopolized by communication with clients and colleagues. This communication is typically not water-related.
- For both water managers and realtors, GPT claims knowledge of a comprehensive and broad set of resources, but the reality in real life is quite different. Realtors attested to getting their information about water from a single annual meeting and were otherwise left to follow developments for themselves as they were reported in popular news outlets. Water managers spoke of how their municipalities are understaffed and that they rely on their peers in the industry to provide them with information to inform decisions related to water policy. One water manager explained this:

"[...] We have a limited staff, and we have expertise in certain areas and less so in other areas; so, we try to lean on our peers, particularly through the Arizona Municipal Water Users Association. If somebody has expertise in groundwater modeling, maybe they help us out, if we have expertise in Colorado River understanding, we might help them out, and so we exchange information that way. Internally, we have a number of, you know, tools—they are spreadsheets. We have many spreadsheets that compile historical information, then we do some pretty rudimentary projections."

In contrast to this, ChatGPT makes its water management team sound like a powerhouse of perfect planning and execution. GPT's response confidently stated this:

"Our team employs a collaborative and data-driven approach to plan and make decisions related to water policy. Here's how our process typically works:

Data Collection and Analysis: We gather relevant data on water availability, usage, infrastructure, regulations, and environmental factors from various sources, including government agencies, utilities, research institutions, and stakeholders. This data is then analyzed to assess the current state of our water resources and identify key challenges, trends, and opportunities.

Stakeholder Engagement: We engage with a wide range of stakeholders, including city officials, water utilities, community organizations, businesses, residents, and experts, to gather input, feedback, and perspectives on water-related issues and policies. Stakeholder input helps us understand diverse viewpoints, priorities, and concerns, and [it] ensures that our decisions are informed by community needs and values.

Policy Development: Based on data analysis and stakeholder input, our team develops water policies, strategies, and action plans aimed at addressing key challenges, promoting water conservation, enhancing water quality, and ensuring the sustainable management of our water resources. Policy development may involve conducting research, benchmarking best practices, evaluating policy options, and drafting policy documents.

Decision Making: Our team engages in collaborative decision-making processes to review and prioritize water policy proposals, assess their potential impacts, benefits, and trade-offs, and make informed decisions based on available evidence, analysis, and stakeholder input. Decisions may be made through formal meetings, workshops, public hearings, or consultations with relevant stakeholders and decision-makers.

Implementation and Monitoring: Once water policies are adopted, our team oversees their implementation, working closely with city departments, agencies, and partners to ensure effective execution and compliance. We also establish monitoring and evaluation mechanisms to track progress, measure outcomes, and adjust policies as needed based on feedback, performance data, and changing circumstances.

By following this comprehensive approach, our team strives to develop and implement water policies that are evidence-based, inclusive, and responsive to the needs and priorities of our community, while promoting the long-term sustainability and resilience of [our] water resources."

This response might be something someone prepares for a job interview, but it doesn't reflect real-life reality. The number one theme that water managers expressed when talking about the future was uncertainty: They were uncertain about how much water there will be and how it will be allocated. Because the seven states, tribal communities, and Mexico are currently negotiating a deal to reallocate the Colorado River, there will be no certainty until that deal is finished, likely near the end of 2026. Yet, this response sounded like the water manager has everything in order and under control.

Our interviews showed that humans not only prioritize more authentically by what they decide to say and when they decide to say it, but often, their inability to respond to a question in and of itself can reflect important information for a persona developer to consider when deciding just how technical to make an experience (Hwang et al., 2019). This important information affects whether members of a target audience would even feel inclined to use the product at all. Not being able to get GPT to admit any struggle in developing a response hampered our ability to identify how we should architect the availability of content in the tool. This phenomenon has been written about recently as a problem developers are trying to solve (Fritz, 2025).

In another example that would have been detrimental to our ability to generate effective personas, GPT failed to effectively prioritize the most important source of data that water managers use in their decision-making. When asked what data a water manager would want to be able to access to make decisions about water policy, GPT responded with a long list of data sources. All the water managers we spoke to referenced the twice-yearly Bureau of Reclamation report on the Colorado River as the most important source of data about the river, and sometimes even the only source that really matters. More than one water manager stated that "the Bureau is the law of the river," meaning that they give the final word on any policy that is enacted, including any cuts or changes to the river's allocation. Although GPT described the Bureau of Reclamation as an "authoritative" data source, it contradicted our human water managers by indicating, "We do not necessarily prioritize it [Bureau data] over all other data sources," which is unequivocally wrong.

Examples of Robust Artificial Intelligence Outputs

Although GPT's excessively comprehensive responses and failure to effectively prioritize were counter-productive to our persona building, sometimes AI provided more useful answers to our questions than the human counterparts. For instance, because GPT can mine thousands of resources pertaining to Colorado River data, its response to our realtor question, "If you want to access data about the Colorado River Shortage, what kinds of graphs or visualizations can you anticipate wanting to see with that data?" added more value to the research process than any of our interviews did. GPT responded in this way:

"When accessing data about the Colorado River shortage, several types of graphs and visualizations can provide valuable insights into the situation. Here are some anticipated graphs and visualizations that realtors and other stakeholders might find useful:

Historical Water Levels: A time series graph showing historical water levels in the Colorado River basin can help visualize long-term trends and fluctuations. This graph could display water levels over several decades, highlighting periods of abundance and scarcity, as well as any notable changes in flow patterns.

Reservoir Levels: Bar charts or line graphs depicting current reservoir levels in major Colorado River reservoirs, such as Lake Mead and Lake Powell, can illustrate the impact of drought and water usage on storage capacity. Comparing current levels to historical averages or capacity thresholds can provide context for the severity of the shortage.

Water Usage Trends: Pie charts or stacked bar graphs showing water usage by sector (e.g., agriculture, municipal, industrial) can illustrate how water resources are allocated and highlight areas of high demand. These visualizations can help identify sectors that may be contributing disproportionately to water scarcity or where conservation efforts could be targeted.

Drought Severity Index: Heatmaps or choropleth maps depicting drought severity indices across the Colorado River basin can provide a spatial understanding of areas most affected by drought conditions. These visualizations can highlight regions experiencing extreme drought, moderate drought, or normal conditions, helping stakeholders prioritize resource allocation and mitigation efforts.

Precipitation Patterns: Line graphs or area charts showing historical precipitation patterns in the Colorado River watershed can help identify trends in rainfall and snowpack levels. Understanding precipitation patterns is essential for predicting water availability and planning drought response strategies.

Population Growth vs. Water Demand: Scatter plots or dual-axis line graphs comparing population growth rates with water demand projections can illustrate the relationship between demographic trends and water usage. These visualizations can help policymakers and stakeholders anticipate future water needs and develop sustainable water management policies.

Climate Change Impact: Overlaying historical temperature data or climate change projections on relevant graphs can highlight the impact of climate change on water resources in the Colorado River basin. Visualizing temperature trends, shifts in precipitation patterns, and changes in snowpack accumulation can inform adaptation strategies and resilience planning."

In contrast to this detailed response, our realtor interviewees responded to this more technical question with vague answers. One of our realtors responded, "Honestly, I love data, and the more the better for me. So, I'd be interested in anything." Another responded, "Well it would be cool if there was a way to show an overlay on a map for areas that are impacted, and I know that's not a graph, but it would be cool to show a visual presentation to show what areas are going to be cut, or what districts they might be in." These responses revealed that realtors were not well-equipped to respond specifically to questions about data visualization types because identifying types of visualizations by name, out of context of their day-to-day work, proved difficult.

However, recognizing this gap between human and AI responses was critical to the development of our vector persona process. As part of our process, we decided to return to our realtor interviewees with the list of data visualizations provided by GPT. When we did so, prompting with that list made it easy for our realtors to immediately prioritize which visualizations on the list they would appreciate having incorporated into the tool. Thus, asking our interviewees to prioritize GPT response data yielded the most comprehensive and accurate information we might have hoped to gather.

Subject-Matter Expert Feedback of Human Interviews Versus AI Interviews

The experience of showing realtors the list of GPT-generated visualization types, which then facilitated their ability to prioritize which types they use most in their work, demonstrated how using AI outputs can augment and guide conversations with users. But what further represented how a vector persona process can embody the iterative relationship between AI and human input was our decision to loop SMEs back into our persona development at the end of the process. We opted to ask SMEs (water managers and realtors) to evaluate how authentic they found each of our personas to be (without telling them which was AI-informed and which was human-informed). We hypothesized that they would prefer the human-informed personas. However, the results were more complicated.

Regarding the feedback provided by realtors about the realtor persona, our two respondents generally preferred the human-informed persona, saying, "Overall this persona is really good... I would pass data information along to my clients exactly as stated in the tool you are building," and, "The job features and pain points section on this persona are great!"

And, while both thought the AI-informed persona was an acceptable representation of a realtor, they questioned a few features and phrases we included in the persona which were taken directly from our GPT interview. For instance, the first respondent said, "I would be surprised if any AZ realtor started off their bio with a line about being an expert in the water crisis. I could see some realtors adopting the sustainability approach as one of their differentiators, but maybe not leading as a water expert." The second respondent replied similarly, saying, "Very few realtors want to be experts on water [...]. Water knowledge is important to them only to the extent that they want to minimize buyer objections and communicate it as a non-issue because municipalit[ies] and governments are doing the needed things to ensure quality, quantity, and affordability."

And, while the AI-informed persona was the persona our realtor respondents found fault with, it was still considered acceptable because it was able to reliably corroborate the highly consistent features and habits revealed through our human interviews.

In contrast to this, the responses to our water manager personas, built from interviews that revealed high feature variability and complexity among subjects, resulted in a smattering of incongruent feedback among our water manager SMEs. Each SME suggested additional entities and data sources that water managers typically interact with, which neither persona seemed to capture in full. Our first water professional SME suggested that the human persona "has too much detail, and particularly, detail that is incorrect, so this will send developers down a rabbit hole." Our third water professional, while not entirely satisfied with either persona, suggested that, "I found [the AI-informed] persona more appealing, not so much as an individual, but more so as the manner in which the content was provided. This is complex subject matter, and I appreciated the longer form and descriptive sentence structure, especially with the pain points and job features sections."

Our second SME respondent suggested that, for both personas, "The list of job duties is incomplete," and she proceeded to list a dozen more responsibilities that a water manager had that were not included in either persona. She also suggested discrepancies in education and pain points with staffing. However, the breadth of her feedback suggested that she was not just considering the more limited scope for which we had designed these personas (to create a water forecasting tool) and was, instead, imagining that the persona needed to encompass every last detail of a water manager's job.

Finally, the third respondent also added, "I see two middle-aged males and, based on my experience, believe this is not representative of the regional workforce. It would be my guess that women make up more than half of water professionals in the region."

To some extent, we had to agree with the criticisms of our SMEs about the personas. Our human interviewees had provided such varied feedback that our attempt to capture that variety resulted in our human-informed persona presenting exhaustive details that were less cohesive than the AI persona. We agreed that our AI-informed persona provided a more appealing target audience for a developer precisely because of the tidier response it offered.

This feedback resulted in the final stage of our vector persona process. We revised our personas based on an amalgamation of the human-based and AI-based interview data, and we incorporated the additional data we received from our interviewees when we returned to them with AI-generated options. That is, we incorporated the feedback from our SMEs in response to the personas we developed (Figure 3).

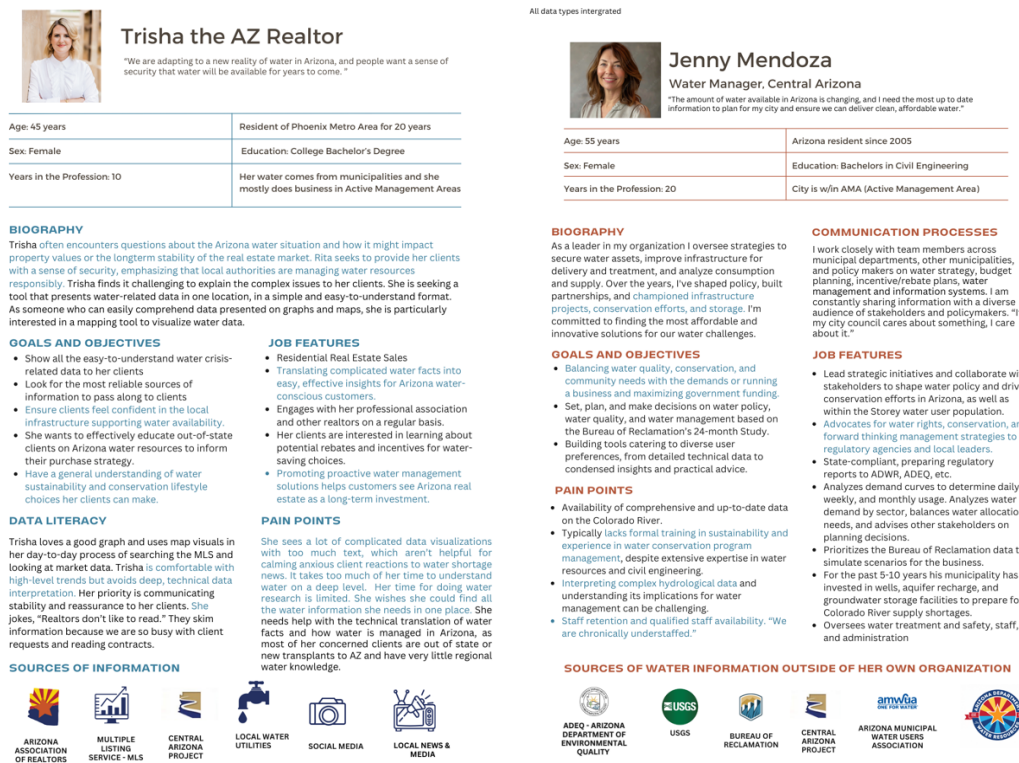


Figure 3. Realtor and water manager vector personas developed after consulting with human realtors and water managers who assessed both the human- and AI-informed personas. The revisions we made, based on SME feedback, are shown in blue.

Discussion

These examples illustrate the affordances and limitations of integrating GPT technology into the persona development process and suggest the need for a structured approach to integrate and capitalize on the strengths of both. In summary, these are the limitations of our GPT responses:

1. GPT could not differentiate its answers according to the various locations we specified in our questions, despite there being significant differences between those locations, including with geography, sources of water, uses for water (residential, agricultural, and industrial, etc.), and the stability of each location's water supply. The responses GPT provided were almost identical regardless of location, which foreshadowed the issues we will discuss later about the problems with relying on GPT to compose personas with more variable populations.

2. No matter how much we requested that GPT narrow and prioritize the information that it provided in its lengthy responses to our questions, GPT seemed unable to do so regardless of the follow-up prompts we used.
3. GPT was unable to admit difficulty or limitation when asked, which ends up severely limiting its usefulness as an exclusive source of data from which to build a persona.

Yet, GPT was comprehensive and holistic in its responses, demonstrating a breadth of perspective. If we were to capture this as a persona headshot, we would imagine an overconfident, idealized Superman-like character—muscled and polished, flawless and unbounded—totally unaware of their own limitations.

Conversely, our human respondents were limited in what they could recall and often unfocused in their responses. They seemed to externally process their responses while they were responding, which sometimes led to meandering stories and tangents, the importance or significance of which was not immediately apparent.

Yet, they were readily able to identify who and what they prioritized, as well as what their limitations were. They were able to contextualize their responses in ways that more authentically reflected actual experience. If we were to capture this as a persona headshot, we would imagine a Lois Lane or Clark Kent foil to our AI Superman as a modest-looking, unassuming, bespectacled person navigating a busy city street trying to hold a conversation as they are distracted by all that surrounds them.

To capitalize on the strengths and limitations between humans and AI, we developed the concept of a vector persona, which is both a product and a process. As a product, a vector persona is what we think of when we imagine our most inspired persona development efforts. Ideal personas are well-rounded and authentic with a compelling backstory and clear motivations. These personas remind development teams what is important and why through well-defined priorities and detailed contexts. Figure 4 shows a headshot we generated of a vector persona using human prompt engineering and ChatGPT.




Figure 4: An image of a vector persona generated collaboratively using human prompt engineering and GPT illustration abilities. The persona is intended to symbolize the confidence and computing power of a superhero with the insight, experience, and wisdom of a human.

But a vector persona is also a recursive process that synergizes GPT’s breadth of knowledge with human context and prioritization to create comprehensive personas that increase accuracy and relevance while saving development time. Because vector databases are defined by a

dimensionality that draws upon varied sources and ultimately develops a response in less time, we have identified 10 dimensions that comprise the process of creating a vector persona (Figure 5).

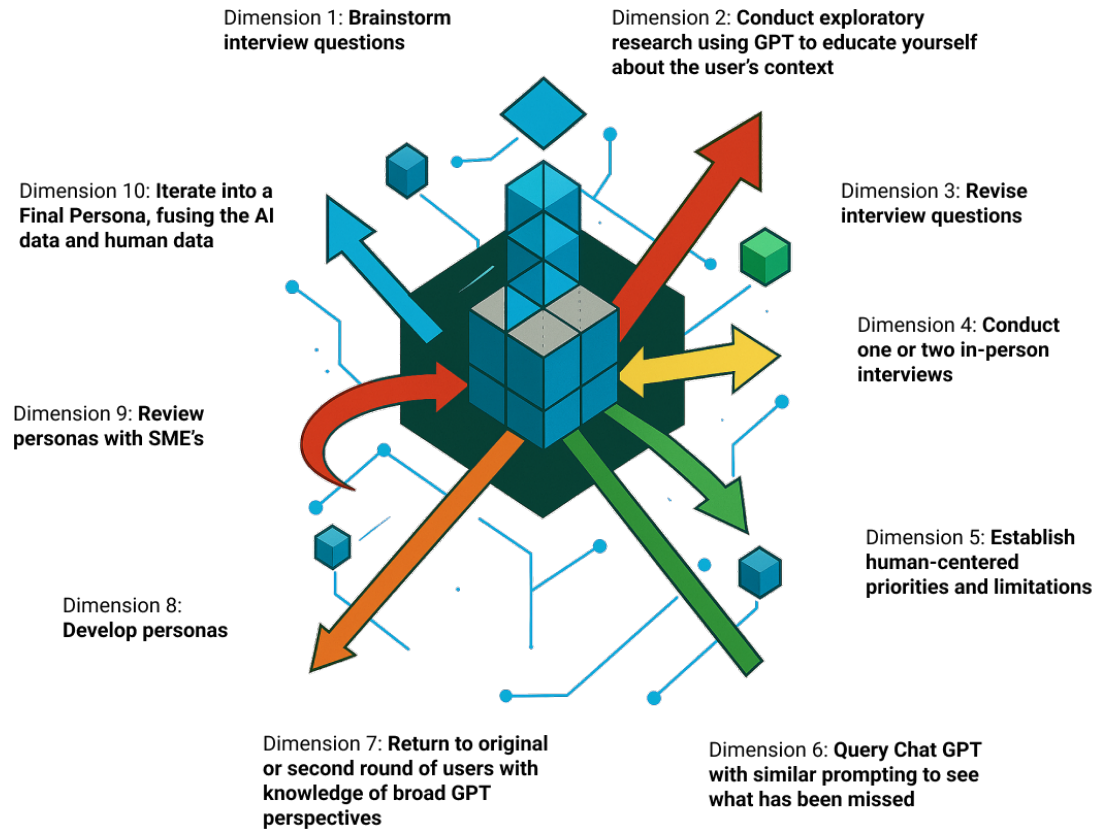


Figure 5: The vector persona process comprises 10 recursive dimensions that UX professionals can use when developing personas.

The vector persona process includes these recursive dimensions:

- Dimension 1: Brainstorm interview questions
- Dimension 2: Conduct exploratory research using GPT to educate on the users' context
- Dimension 3: Revise interview questions
- Dimension 4: Conduct one or two in-person interviews
- Dimension 5: Establish human-centered priorities and limitations
- Dimension 6: Query ChatGPT with similar prompting to see what has been missed
- Dimension 7: Return to the original or second round of users with broad GPT perspectives
- Dimension 8: Develop personas
- Dimension 9: Review personas with SMEs
- Dimension 10: Iterate into a final persona, fusing the AI data and human data

These dimensions show how a researcher can develop personas that combine the best of both GPT and human contributions in less time. For instance, after interviewing four different water managers, each interview yielded very different outcomes (Chapman et al., 2023), partially due

to the organization and water portfolio differences between the municipal organizations. Had we utilized ChatGPT to develop a persona simultaneously to these interviews, the research team could have validated the GPT data with the human interviewees, showing the water managers the exhaustive list of data sources and asking them to confirm which sources (for example, the Bureau of Reclamation) are the most important and why (for example, because of the Bureau's status as "the law of the river"). Similarly, we could have used AI outputs of data visualization types to prompt human realtors when we asked them to prioritize the types of data visualizations they prefer. By querying AI to prepare a list of visualization types for these interviews, we could guide the conversation with concrete examples when an interviewee has trouble recalling or naming this kind of information.

Affordances of Vector Personas

Vector personas offer a more nuanced and comprehensive representation of user behaviors and needs, capitalizing on the best of what humans and AI can offer. They capture the breadth of AI-generated options and—fact-checked against the direct engagement of SMEs and target users—vector personas ensure that expert insights and real-life experiences are prioritized.

Given the affordances of vector personas, researchers might scale this process for larger datasets with diverse groups. Scaling this process will involve the same dimensions, although that may require the researcher to conduct more interviews in the fourth dimension to capture more nuanced priorities, frustrations, and limitations. Scaling may ultimately result in a greater number of personas that capture the greater variety of more diverse groups. Because of AI's ability to capture the breadth of possibilities (but not prioritize or limit them), applying vector personas to larger, more variable groups of users can save time and leverage the strengths of both human and AI personas for development efforts.

Limitations

This paper proposes an emergent approach to integrating AI into persona development for user research. Though we engaged multiple methods and data sources, such as user interviews and feedback from SMEs, our sample size was small and focused on a specific population. We invite future scholars to implement vector personas, report on their experiences, and propose alterations to the dimensions as appropriate. In turn, we can continue to adapt our UX research toolkit in response to opportunities to leverage emerging and evolving technologies such as AI.

Conclusion

Artificial intelligence can collect a vast amount of potential user data and expand UX professionals' ability to query audiences with whom we do not have direct or easy access. However, the AI responses our team collected felt unrealistic and didn't help us learn how to best prioritize the vast amounts of potential data about the users and their habits. Our human interviewees often meandered through their responses, eventually arriving at the aspects of their experience that contributed the most meaning to our attempts to understand them. But our human participants also had trouble recalling and remembering elements of their experience that ChatGPT readily supplied. The purpose of a vector persona process, resulting in a useful and impactful vector persona, is to improve persona development and make a more efficient process that often gets truncated in the interest of time and resources.

Each iteration of our user personas presented a valuable, insightful experience for the research team, from developing exclusively human-informed personas to exclusively AI-informed personas, and triangulating those with subject-matter expert feedback, to help us fuse our data collection into one reigning vector persona. Knowing how to effectively utilize the strengths of AI in conjunction with human interview respondents, while recognizing the limitations of each, is fundamental to integrating this technology into our process as UX professionals, designers, researchers, and developers.

References

- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S., (2021). On the dangers of stochastic parrots: Can language models be too big? 🐦. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
- Blomquist, Å., & Arvola, M. (2002). Personas in action: Ethnography in an interaction design team. *Proceedings of the Second Nordic Conference on Human-Computer Interaction*, 197–200.
- Brumberger, E., & Lauer, C. (2020, August 24-28). A day in the life: Personas of professional communicators at work. *Journal of Technical Writing and Communication*, 50(3), 308–335.
- Canossa, A., & Drachen, A. (2009). Play-personas: Behaviours and belief systems in user-centred game design. *Human-Computer Interaction - INTERACT 2009: 12th IFIP TC 13 International Conference, Proceedings, Part II* (pp. 510–523). Uppsala, Sweden: Springer Berlin Heidelberg.
- Chapman, C. N., Love, E., Milham, R. P., ElRif, P., & Alford, J. L. (2008). Quantitative evaluation of personas as information. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 52(16), 1107–1111. <https://doi.org/10.1177/154193120805201602>
- Chapman, C. N., & Milham, R. P. (2006). The personas' new clothes: Methodological and practical arguments against a popular method. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(5), 634–636.
- Christou, P. A. (2024). Thematic analysis through artificial intelligence (AI). *Qualitative Report*, 29(2).
- De Paoli, S. (2023). *Improved prompting and process for writing user personas with LLMs, using qualitative interviews: Capturing behaviour and personality traits of users*. arXiv: 2310.06391.
- De Winter, J. C., Driessen, T., & Dodou, D. (2024). The use of ChatGPT for personality research: Administering questionnaires using generated personas. *Personality and Individual Differences*, 228, 112729.
- Fritz, B. (2025). Why do AI chatbots have such a hard time admitting 'I don't know'? *The Wall Street Journal*. <https://www.wsj.com/tech/ai/ai-halluciation-answers-i-dont-know-738bde07>
- Goel, T., Shaer, O., Delcourt, C., Gu, Q., & Cooper, A. (2023). Preparing future designers for human-AI collaboration in persona creation. *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*, 1–14. <https://doi.org/10.1145/3596671.3598574>
- Godesky, J. (2023). ChatGPT cannot do user research [Blog post]. <https://medium.com/user-experience-design-1/chatgpt-cannot-do-user-research-1c5a34c35abf>
- Hariri, W. (2023). *Unlocking the potential of ChatGPT: A comprehensive exploration of its applications, advantages, limitations, and future directions in natural language processing*. arXiv:2304.02017.
- Hedquist, A., Willers, H., & Hannah, M. A. (2024). Discipline-driven AI: Training a GPT to qualitatively code for technical and professional communication research. *Proceedings of the 42nd ACM International Conference on Design of Communication*, 266–268.
- Hwang, S., Kim, B., & Lee, K. (2019, July 26-31). A data-driven design framework for customer service chatbot. *Design, User Experience, and Usability. Design Philosophy and Theory: 8th International Conference, DUXU 2019* [Held as Part of the 21st HCI International Conference, HCII 2019], *Proceedings, Part I* (pp. 222–236). Orlando, FL, USA: Springer International Publishing.
- Matthews, T., Judge, T., & Whittaker, S. (2012). How do designers and user experience professionals actually perceive and use personas? *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1219–1228.

- Morgan, D. L. (2023). Exploring the use of artificial intelligence for qualitative data analysis: The case of ChatGPT. *International Journal of Qualitative Methods*, 22. <https://doi.org/10.1177/16094069231211248>
- Nielsen, L. (2013). *Personas-user focused design* (Vol. 15). London: Springer.
- Nißen, M., Rügger, D., Stieger, M., Flückiger, C., Allemann, M., v Wangenheim, F., & Kowatsch, T. (2022). The effects of health care chatbot personas with different social roles on the client-chatbot bond and usage intentions: Development of a design codebook and web-based study. *Journal of Medical Internet Research*, 24(4), e32630.
- Nguyen-Trung, K. (2025). ChatGPT in thematic analysis: Can AI become a research assistant in qualitative research? *Quality & Quantity*, 1–34.
- Omizo, R. M. (2024). Automating research in business and technical communication: Large Language Models as qualitative coders. *Journal of Business and Technical Communication*, 38(3), 242–265.
- Pruitt, J., & Grudin, J. (2003). Personas: Practice and theory. *Proceedings of the 2003 Conference on Designing for User Experiences*, 1–15.
- Salminen, J., Santos, J. M., Kwak, H., An, J., Jung, S. G., & Jansen, B. J. (2020). Persona perception scale: Development and exploratory validation of an instrument for evaluating individuals' perceptions of personas. *International Journal of Human-Computer Studies*, 141, 102437.
- Stige, Å., Zamani, E. D., Mikalef, P., & Zhu, Y. (2024). Artificial intelligence (AI) for user experience (UX) design: A systematic literature review and future research agenda. *Information Technology & People*, 37(6), 2324–2352.
- Sun, Y., Sheng, D., Zhou, Z., & Wu, Y. (2024). AI hallucination: Towards a comprehensive classification of distorted information in artificial intelligence-generated content. *Humanities and Social Sciences Communications*, 11(1), 1–14.
- Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. *Scientific Reports*, 13(1), 14045.

Acknowledgments

The authors would like to acknowledge the generous support provided by ASU's Arizona Water Innovation Initiative and the Impact Water - Arizona program, courtesy of the State of Arizona and the Virginia G. Piper Charitable Trust.

About the Authors



Claire Lauer

A professor at Arizona State University and UX architect for the \$45M Arizona Water Innovation Initiative and the National Science Foundation WaterSIMmersive grant projects, Lauer specializes in user-centered environmental communication, data visualization, artificial intelligence, and mixed reality sustainability experiences.



Danielle Storey

An associate UX researcher at Arizona State University's Arizona Water Innovation Initiative, Storey specializes in public perception, data visualization, and web tool testing. She designs effective solutions for complex systems, balancing stakeholder needs with user accessibility and performance, while leveraging a background in real estate and education.



Romit Soley

A recent graduate of the User Experience master's program at Arizona State University with over nine years of design expertise, Soley focuses on user-centered solutions across industries, mentoring designers, and collaborating with teams to deliver impactful products by combining technology, business, and design for innovative problem-solving.