Digital Diversity:
Crafting Inclusive AI Narratives

Pr. Nr: 2025-1-DE02-KA210-VET-000354956

Focus-Group Interviews and Analysis
Bulgaria

Co-funded by
the European Union

Focus-Group Interviews and Analysis Bulgaria, 17.10.2025

Digital Diversity: Crafting Inclusive AI Narratives (D2CIN)

Project Number: 2025-1-DE02-KA210-VET-000354956

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

Project	Digital Diversity: Crafting Inclusive AI Narratives (D2CIN) Project Number: 2025-1-DE02-KA210-VET-000354956
Work Package / Activity	A1.2 – Focus-Group Interviews and Analysis
Session	Bulgaria – UX & Accessibility Designers
Date & Time (Local)	Live Session: 17 October 2025; 16:00 – 17:30 Online Session: 17 October 2025; 18:00 – 19:00
Location / Platform	Live Session: Sofia University “St. Kliment Ohridski”Faculty of Journalism and Mass Communication, Hall 17Online Session: Google Meet
Moderator / Note-taker / Accessibility Lead	Asst. Prof. Yordan Karapenchev, PhD Project Manager at Budakov Films Ltd.

Participant Overview

Category	Number of Participants	Background / Role	Experience & Affiliation
Graphic Designers (UX Focus)	3	Professional graphic designers with specialization in UX design	Each has between 7–10 years of experience in UX and digital product design
University Students	4	Students studying Public Relations and Graphic Design	Enrolled at the University of Sofia, representing emerging perspectives in creative communication and design
UX Expert (Private Sector)	1	UX specialist from a private technology company	Focused on digital interface optimization and user accessibility

Category	Number of Participants	Background / Role	Experience & Affiliation
Market Research Expert	1	Professional researcher from a company specializing in UX and accessibility surveys	Brings experience in user feedback analysis, product accessibility assessment, and usability studies
PhD Student (AI Field)	1	Doctoral researcher specializing in Artificial Intelligence	Provides an academic and technical perspective on AI applications in design and learning contexts
Graphic Designer (IT Sector)	1	Graphic designer working in an IT company developing mobile applications	Experienced in UX for mobile interfaces and accessibility in app design

Consent obtained from all participants

Yes

Recording method:

The focus group sessions were audio recorded using a digital recording device to ensure accurate capture of discussions. Notes were also taken by the moderator to support transcription and thematic analysis.

Data storage location & retention:

Audio files and transcripts are securely stored on a password-protected institutional drive managed by Budakov Films Ltd. Access is restricted to the project research team only.

Session Summary (≤150 words)

The focus group in Sofia took place in two parts — a live discussion with four participants at Sofia University and an online panel with seven participants on Google Meet. Both sessions were moderated by the same facilitator and followed the same set of guiding questions. Despite being held separately, the conversations revealed strikingly similar views.

Participants in both groups voiced concerns about the potential misuse of AI-based avatars, particularly the risk of promoting unethical or manipulative behavior. At the same time, they spoke passionately about the power of UX design — how the smallest design choices can shape trust, empathy, and the overall experience of interacting with AI. Many agreed that responsible design isn't just about accessibility or aesthetics, but about ethics and impact.

Together, these sessions offered rich, first-hand perspectives on how UX can guide the development of inclusive and trustworthy AI systems.

1. Bias Mitigation & Ethics

Participants in both sessions spoke passionately about the ethical challenges and hidden biases that come with AI-driven design tools. Several noted that artificial intelligence often mirrors the values and blind spots of the people who create it — especially when it's trained on limited or non-diverse data. As one participant put it, these systems "reflect the world we feed them," making it crucial for designers to stay alert to what those choices imply.

There was strong agreement that UX designers have a key role in spotting and addressing bias, not simply leaving that responsibility to developers or AI engineers. Participants felt that ethics should be built into every design stage — from early concept sketches to prototype testing. They also emphasized how much transparency and user education matter for trust. "If people don't understand how an AI makes its choices," one participant noted, "they won't feel safe using it."

Illustrative Quotes

- "If the data behind the AI isn't diverse, the product can't be fair." — *(P3, 00:06:42)*
- "We talk about inclusivity, but most design tools still assume one type of user." — *(P7, 00:06:05)*
- "AI can mirror our blind spots — that's why UX needs to question every assumption." — *(P9, 00:08:33)*

Implications for Design / Training:

- Build bias-awareness modules into UX and AI design training.
- Foster collaboration between designers, data scientists, and ethicists early in the design process.
- Develop project tools (e.g., the *Personalised Feedback Navigator*) that can flag potential bias indicators and provide clear ethical design checklists.
- Use real-world case studies to help designers see how data or interface decisions can unintentionally exclude or mislead users.

2. Inclusive Data Management

The topic of data management and personalization led to one of the most engaged parts of the discussion. Participants began questioning what happens when AI platforms — especially those generating avatars — start remembering users' preferences for gender, religion, or physical appearance. The tone quickly shifted from curiosity to concern.

Several participants agreed that personalization can improve accessibility and comfort, but only if handled with clear consent and strong safeguards. Others warned that storing such data could expose users to discrimination or stereotyping, particularly if this information were shared across platforms or used for marketing. "Once the system remembers your choices," one designer said, "you lose control over how that memory is used."

There was also a thoughtful exchange about implicit profiling — how algorithms might infer sensitive traits from user behavior, even without explicit disclosure. A few participants noted that such systems could unintentionally reinforce social biases or make users feel surveilled, rather than supported.

Ultimately, the group agreed that inclusivity in AI should never come at the expense of privacy or autonomy. Users should know *what* is remembered, *why*, and *for how long*.

Illustrative Quotes (Pseudonyms + Timestamps):

- "Preferences are personal data — they tell more about us than we think." — *(P5, 00:11:14)*
- "It's not wrong to personalize, it's wrong [*for the platform*] to remember without permission." — *(P8, 00:12:02)*
- "What if somebody starts profiling you as a racist if your avatar is always white?" — *(P4, 00:13:10)*

Implications for Design / Training:

- Require explicit, informed consent before storing or reusing user preferences related to identity, religion, or appearance.
- Implement data minimization principles, collecting only what's essential for functionality.
- Provide clear privacy dashboards where users can view, edit, or delete stored personalization data.

3. Inclusive UX & Digital Storytelling

When the discussion shifted toward AI-generated avatars used to deliver training content, the room grew livelier. Some participants leaned forward, curious; others hesitated before speaking. The topic felt personal for many people.

Participants discussed what might happen if an avatar looked or behaved differently from the expected norm. What if the digital trainer appeared to belong to a visible minority group, wore religious symbols like a hijab or a cross, or expressed nontraditional gender traits? Would learners react with empathy, discomfort, or bias? No one had a simple answer, but people agreed the issue mattered deeply.

A few felt that such representations could spark unwanted controversy or distract from the content itself. Others took the opposite view, arguing that avoiding diversity altogether creates its own form of exclusion.

A partial consensus amongst participants formed on the point that diversity should not be about being politically correct, but about designing systems that respect the complexity of real people. Participants suggested offering customization options, so users can choose or adapt avatars according to their comfort and cultural context.

- Quotes:

“If an avatar wears a hijab, people will notice it first and get distracted from what is said.” — (P4, 00:09:17)

“We can’t just make avatars look neutral on all human things.” — (P2, 00:30:01)

“If an avatar acts too feminine or too masculine, it instantly changes how people respond to the message.” — (P7, 00:11:26)

“Diversity isn’t the problem — it’s how we frame it and give users the choice to engage.” — (P6, 00:12:03)

Implications for Design / Training:

Develop clear representation guidelines for avatar design that balance diversity with sensitivity.

Provide user customization tools for avatar appearance, language, and behavior.

Include bias-awareness and cultural sensitivity training for AI and UX teams.

Conduct testing with diverse audiences to assess perceptions of inclusion and comfort.

Reinforce the principle that inclusive design doesn’t mean neutrality.

4. Accessibility Standards (WCAG, assistive tech)

Several participants noted that most current AI avatars are not optimized for assistive technologies such as screen readers, captioning tools, or voice navigation. Others pointed out that accessibility should go beyond technical compatibility — avatars should also be visually inclusive and cognitively clear, using body language, tone, and pacing that make information easier to follow for different learners.

Some professionals shared that accessibility is often treated as a "final layer," added after production rather than integrated from the start. One designer summed it up sharply: "We need to think about accessibility as part of the script — not the subtitles we add later."

The group agreed that improving accessibility for AI-driven avatars will require both design awareness and technological innovation. That includes generating avatars with natural, diverse voices, clear articulation, adaptive gestures, and text alternatives for all key content.

Illustrative Quotes:

- "If an avatar's gestures are too fast or not aligned with captions, it becomes impossible to follow." — *(P2, 00:24:12)*
- "Accessibility isn't just captions; it's how the avatar moves, speaks, and engages." — *(P5, 00:15:03)*
- "Most AI avatars can't be used with screen readers or alternate input devices — that's a real barrier." — *(P6, 00:25:45)*
- "We should design avatars that *adapt* — slower speech, simpler gestures, maybe even adjustable visual contrast." — *(P8, 00:26:22)*

Implications for Design / Training:

- Develop adaptive avatar settings that let users adjust voice speed, tone, contrast, and gesture intensity.
- Ensure AI-generated video content is compatible with assistive technologies, including subtitles, captions, and screen-reader support.
- Integrate accessibility awareness modules in UX and AI training to address multimodal inclusion — visual, auditory, and cognitive.
- Treat accessibility as a core part of storytelling, not a post-production add-on.

5. Speech Evaluation

Prompting participants to reflect on pacing, tone, and pronunciation and discuss if these can make or break a learner's experience. Most agreed that if an avatar speaks too quickly or with inconsistent rhythm, comprehension drops sharply. One designer described it as "listening to someone read without breathing," and several nodded in agreement.

Everyone supported the idea that users should be able to adjust the avatar's speech speed and clarity, making learning more adaptive. The group also talked about multilingual delivery, noting that speech should remain neutral — free from dialects, regional accents, or colloquial expressions — to ensure the content feels inclusive and universally understandable.

At this point, the uncanny valley emerged as a key concern. Some participants said that when the voice sounds almost human but not quite natural, it creates unease and distracts from learning. Others observed that overly expressive tones or exaggerated pauses can feel artificial and break the user's emotional connection. "The closer we get to sounding human," one participant remarked, "the easier it is to notice what's missing."

- Quotes:

"Speech should sound neutral — too much accent or emotion, and it gets creepyvalley." — (P3, 00:21:32)

"It's not about sounding human, it's about being understandable and consistent." — (P7, 00:22:40)

Implications for Design / Training:

- Monitor uncanny valley effects during user testing — avoiding overly "human-like" vocal synthesis that feels artificial or eerie.
- Include speech perception and cognitive load studies in avatar development to balance realism with clarity.
- Emphasize that intelligibility and comfort are central to ethical, accessible AI learning design.

6. Adaptive Content

When the moderator introduced the topic of adaptive content — AI systems that adjust training materials based on learner behavior or progress — the group's reactions were mixed. Participants appreciated the idea in theory but expressed uncertainty about how far adaptation should go and what trade-offs it might introduce. As the discussion progressed most participants veered toward being against the use of adaptive content.

While it could personalize the learning pace, repeat difficult topics, or simplify explanations in real time. Others, however, worried that too much automation might reduce transparency and learner agency. "If the AI changes the content without telling me," one participant asked, "how do I know what I'm missing?"

Quotes:

"If the AI changes the content without telling me," one participant asked, "how do I know what I'm missing?" (P1, 00:40:02)

Implications for Design/Content

- The notion of “adaptive content” is new to the respondents. Proceeding with its use must be done with caution.

7. Explainability & User Control

Many emphasized that explainability is essential — users should understand why the avatar behaves the way it does, especially when it adapts tone, speed, or content. “If the AI changes something,” one participant said, “I should know why it did that.” It should be noted that these answers were

At the same time, several participants cautioned against overloading users with too many personalization options. They explained that while customization can distract people. Also – it is possible that this may infact stimulate the creation of more biased visuals and content.

Quotes:

“Transparency doesn’t mean showing the code; it means explaining choices in human terms.” — (P4, 00:36:45)

Implications for Design / Training:

- Integrate clear feedback mechanisms that explain when and why AI avatars adapt their behavior or content.
- Provide user-friendly control panels with essential settings only — prioritizing simplicity over excessive customization.
- Develop tiered explainability models offering both quick, plain-language insights and deeper technical details for advanced users.

Following other discussions, participants actually were divided on the importance of user contrals. Reinforcing effective control supports learning, while overpersonalization can reduce clarity, focus, and trust. Some emphasized that transparency is essential — users should understand *why* the AI behaves the way it does, especially when it adapts tone, speed, or content. “If the AI changes something,” one participant said, “I should know why it did that.”

At the same time, several participants cautioned against overloading users with too many personalization options. They explained that while customization can create a sense of ownership, excessive choice might distract from learning or make the interface feel complicated. “There’s a fine line,” one designer noted, “between empowerment and fatigue.” Others agreed that personalization should be guided and purposeful, not endless or cosmetic.

The group also returned to the emotional side of control — understanding builds trust, and trust helps learners stay engaged. When people can see how and why an AI makes decisions, they feel less manipulated and more respected.

Quotes:

- “Transparency doesn’t mean showing the code; it means explaining choices in human terms.” — *(P4, 00:40:45)*

- “Don’t just assume I’ll notice.” — *(P1, 00:41:18)*
- “There’s a fine line here — too many sliders, and people give up.” — *(P3, 00:39:50)*

Implications for Design / Training:

- Integrate clear feedback mechanisms that explain when and why AI avatars adapt their behavior or content.
- Provide user-friendly control panels with essential settings only — prioritizing simplicity over excessive customization.
- Develop tiered explainability models offering both quick, plain-language insights and deeper technical details for advanced users.
- Include training on cognitive load and user decision fatigue for designers developing personalization features.
- Reinforce that effective control supports learning, while overpersonalization can reduce clarity, focus, and trust.

Navigator Criteria and Practical Strategies

Participants identified three criteria that could be used to guide the development of the Personalized Feedback Navigator: 1) bias detection, 2) accessibility integration, and 3) adaptive content delivery. They agreed that the tool should not only flag potential ethical or usability issues but also explain *why* something might be problematic. This will offer readers context, not just correction. In their view, the Navigator should function as both a diagnostic and an educational resource, helping designers and educators understand how to make AI-driven learning more inclusive.

Handbook Development Priorities

Participants provided recommendations for shaping the forthcoming Handbook: “Accessible AI for Education – A Comprehensive Guide to Inclusive Design and Training.” They emphasized that the handbook should serve as a practical guide for educators and trainers. This is especially true for people who work with AI-based learning tools, bridging theory with hands-on application.

From a pedagogical perspective, participants recommended focusing on how AI can support differentiated instruction, learner engagement, and ethical classroom integration. They called for simple frameworks that help educators evaluate when and how to use AI responsibly, without over-reliance on automation.

From an accessibility standpoint, the handbook should include step-by-step guidance on integrating universal design principles — covering topics such as multimodal delivery, inclusive representation in AI avatars, and compatibility with assistive technologies.

Finally, from a UX and design perspective, participants urged the inclusion of visual heuristics, design checklists, and case studies that illustrate inclusive practices in real-world contexts. They stressed that the handbook should remain readable, modular, and adaptable across cultural and linguistic contexts. Overall, participants envisioned a resource that empowers professionals to make AI in education more ethical, transparent, and accessible through practical, evidence-based design.