




**Digital Diversity:
Crafting Inclusive AI Narratives**

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

Handbook ‘Accessible AI for Education: A Comprehensive Guide to Inclusive Design and Training’



**Co-funded by
the European Union**

Handbook

'Accessible AI for Education: A Comprehensive Guide to Inclusive Design and Training'

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.



Document Name	Handbook 'Accessible AI for Education: A Comprehensive Guide to Inclusive Design and Training'
Project Activity	Activity 2: Accessible AI for Education: A Comprehensive Guide to Inclusive Design and Training
Revision Type	Final
Revision Date	01.04.2026
Authors	Mimic Productions (Germany), Budakov Films (Bulgaria)

Declaration on copyright:



This document is protected through the Creative Commons

Attribution-Non-Commercial-Share-Alike 4.0 International License. You are free to:

- Share — copy and redistribute the material in any medium or format
- Adapt — remix, transform, and build upon the material under the following terms:
- Attribution — you must give appropriate credit, provide a link to the license, and indicate if changes are made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- Non-commercial — You may not use the material for commercial purposes.
- Share Alike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

Any unauthorized use or reproduction of the contents of this training module will be considered a violation of copyright law and subject to legal action.

Table of Content

Executive Summary.....	5
Introduction: Purpose & Usage Guide.....	7
Chapter 1: Bias Recognition and Mitigation.....	9
Chapter 2: Inclusive UX Design.....	15
Chapter 3: Accessibility Standards in AI.....	23
Chapter 4: AI Ethics and Compliance.....	29
Chapter 5: Digital Storytelling for Inclusive Education.....	35
Conclusion.....	41

Executive Summary

The Handbook, "Accessible AI for Education: A Comprehensive Guide to Inclusive Design and Training," constitutes the primary intellectual output of Activity 2. It serves as a strategic response to the digital skills gap identified within the European creative VET sector, specifically addressing the intersection of Artificial Intelligence, ethics, and inclusive pedagogy.

Context and Evidence Base

The development of this Handbook is not theoretical; it is grounded in rigorous evidence gathered during Activity 1.

- **The Internal Gap:** Our initial Needs Analysis revealed that 89% of target staff (at Mimic Productions and Budakov Films) lack formal training in inclusive AI practices, specifically in bias mitigation and accessibility standards [Source: Proposal].
- **The External Validation:** The content is further refined by the Interdisciplinary Focus Group Analysis conducted in Germany (Inclusive Education Experts) and Bulgaria (UX & Accessibility Designers). These sessions confirmed that while educators prioritize "pedagogical fairness," designers prioritize "cognitive load" and "trust," creating a dual necessity for a handbook that bridges these two disciplines [Source: Insights & Recommendations Report].

The Strategic Goal

The overarching goal of this document is to translate the research findings of Activity 1 into a tangible, "how-to" manual for VET trainers, creative freelancers, and instructional designers. It aims to move beyond compliance with the EU AI Act to the practical application of "Privacy-by-Design" and "Representational Fairness" in daily workflows.

Core Pillars of the Handbook

To ensure a holistic approach to digital inclusion, the Handbook is structured around five critical pillars, validated by the focus group results:

1. **Bias Recognition and Mitigation:** Moving from abstract concepts to practical "Dataset Audits" to prevent historical and representational bias in AI avatars [Source: Insights Report, Section 6].
2. **Inclusive UX Design:** providing heuristics to manage "Cognitive Load" and avoid the "Uncanny Valley" effect for neurodivergent learners [Source: Insights Report, Section 6].

3. Accessibility Standards in AI: Going beyond WCAG 2.2 AA to ensure "Multimodal Accessibility" (compatibility with screen readers, adaptive pacing) specific to VET environments [Source: Insights Report, Section 5.2].
4. AI Ethics and Compliance: Operationalizing "Tiered Explainability," ensuring educators can explain why an AI model made a specific decision to learners [Source: Insights Report, Section 5.3].
5. Digital Storytelling for Inclusive Education: Utilizing "Narrative Agency" to create diverse, authentic scenarios that avoid tokenism and foster emotional connection.

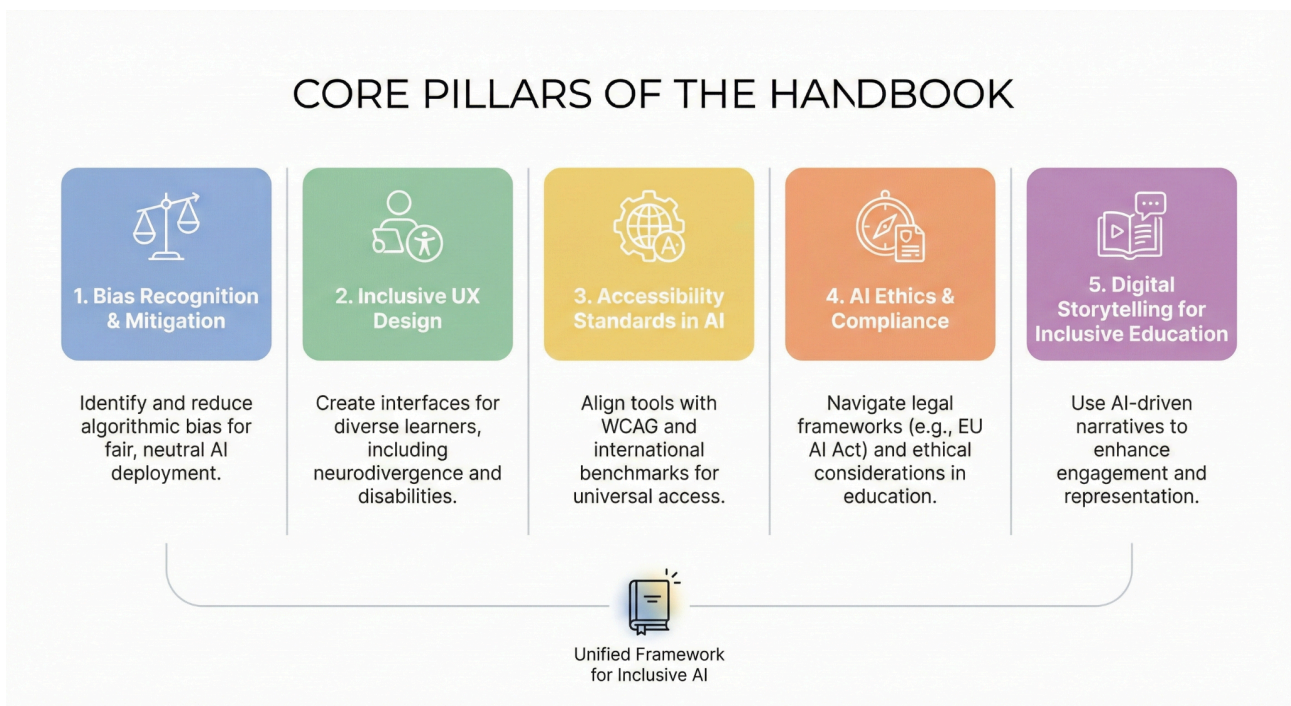


Image: Core Pillars of the Handbook

Intended Impact

By adopting the frameworks within this Handbook, VET professionals will be equipped to:

- Audit their own AI tools for exclusion and bias before deployment.

- Design learning experiences that adapt to diverse cognitive and physical needs.
- Certify their new competencies through the Europass framework, contributing to the "Digital Readiness and Resilience" priority of the Erasmus+ programme.

Introduction: Purpose & Usage Guide

Purpose of this Handbook This Handbook, Accessible AI for Education, is the primary intellectual output of Activity 2. It functions as a strategic bridge between technical AI deployment and pedagogical inclusivity.

It is not merely a theoretical text but a remedial toolkit designed to address the specific skills gaps identified in the creative VET sector, where our research indicates 89% of staff lack formal training in bias mitigation and inclusive design. Its core objective is to empower VET professionals to create, train, and deploy AI avatars and narratives that are:

- **Ethically Sound:** Adhering to the principles of "Privacy-by-Design" and data minimization.
- **Legally Compliant:** Aligned with the transparency requirements of the EU AI Act and GDPR.
- **Culturally Representative:** Utilizing "Representational Fairness" to avoid tokenism and stereotypes in educational content.

How to Use This Handbook This document is engineered as an active implementation tool, not a passive reading resource. It is designed to be used alongside the Personalized Feedback Navigator (Activity 1) to audit and improve real-world projects.

For Developers & Creative Technologists:

- **As an Auditing Protocol:** Use the "In-Practice Scenarios" as mandatory quality assurance checks. Before releasing an AI module, verify it against the Bias & Accessibility Checklists provided in Chapters 1 and 3.
- **For Design Specifications:** Consult the "Inclusive UX" heuristics in Chapter 2 during the wireframing phase to prevent the "Uncanny Valley" effect and manage cognitive load for neurodivergent users.

For Educators & Instructional Designers:

- As a Decision-Support System: Use the "Core Theoretical Concepts" and "Case Studies" to navigate ethical dilemmas in the classroom (e.g., Why did the AI grade this student poorly?).

For Organizational Leaders:

As a Standardization Framework: Use the handbook to establish a unified "Inclusive AI Policy" for your institution, ensuring all digital training materials meet the high standards of the Erasmus+ programme regarding diversity and inclusion.

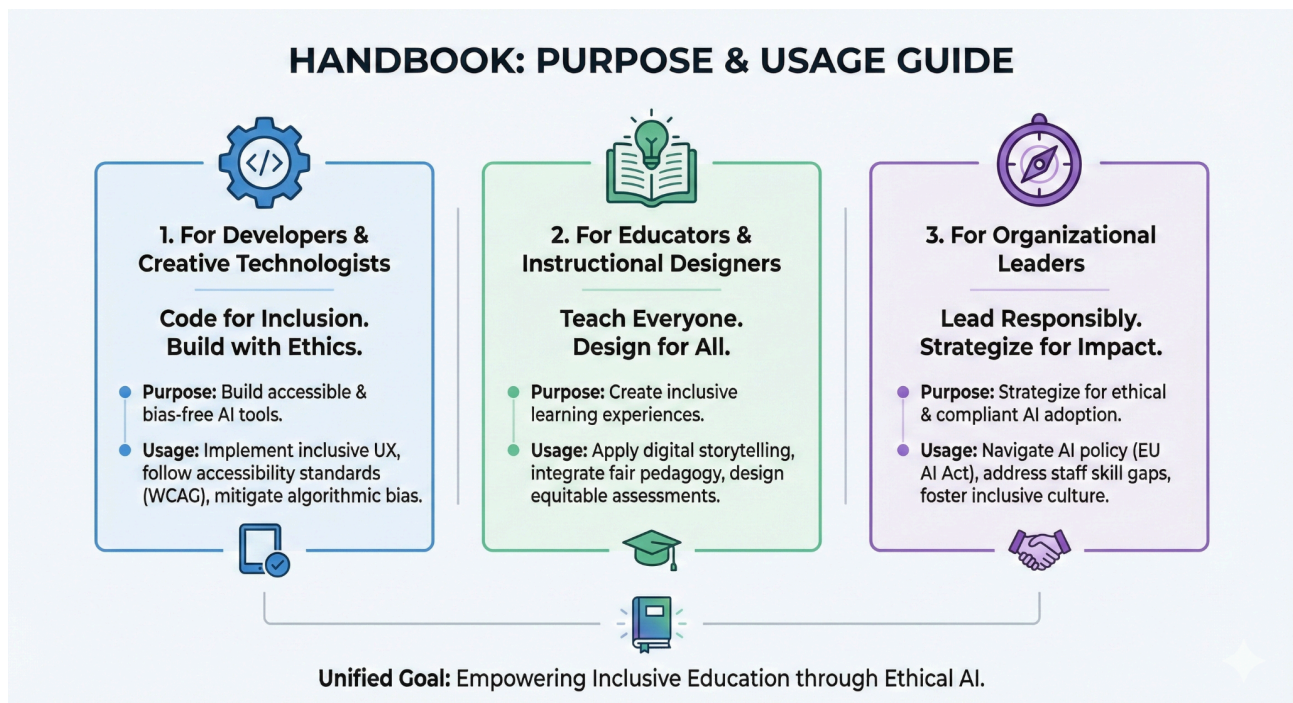


Image: Purpose & Usage Guide

Chapter 1: Bias Recognition and Mitigation

Context & Rationale

The Problem: The D2CIN Needs Analysis revealed that a majority of target staff struggle to produce unbiased conversational AI. This is not simply a technical oversight; it is a reflection of a deeper competency gap. When asked to audit their own tools, staff found it difficult to identify the root causes of skewed AI output. Furthermore, the Interdisciplinary Focus Groups, particularly the Inclusive Education Experts in Germany, emphasized that AI systems inevitably reflect the values, assumptions, and blind spots of their creators. When an AI avatar is trained on historically non-representative data, it does not just produce “neutral” output — it actively reinforces old societal hierarchies, marginalising learners from non-dominant backgrounds and eroding their confidence. (Source: D2CIN Insights Report).

The Goal: This chapter addresses the critical inability to detect and correct bias before it reaches the learner. It equips VET trainers and developers with a clear taxonomy of bias types and a practical, step-by-step dataset audit methodology. By the end of this chapter, the reader will be able to proactively identify discriminatory patterns in AI training content and apply Representational Fairness principles to create a more equitable educational experience for all learners.

2. Learning Objectives (LOs)

- **LO1:** Identify and categorize common forms of algorithmic bias (historical, representation, and measurement) within VET datasets to prevent the reinforcement of stereotypes.
- **LO2:** Apply "Representational Fairness" strategies to ensure AI avatars and voices reflect diverse demographics without resorting to tokenism.

3. Core Theoretical Concepts

To effectively audit AI tools for bias, VET professionals must first understand the vocabulary of the problem. Bias does not emerge from malicious intent; it emerges from uncritical data choices. The following three concepts are the foundational framework for this chapter.

3.1 Data Bias vs. Output Bias: The Root and the Fruit

Data bias refers to systematic skewing in the training datasets that an AI model learns from. If a dataset for a film production AI is composed entirely of Oscar-winning films from Hollywood between 1980 and 2010, the AI has been taught that “quality cinema” looks, sounds, and is directed in a very specific way. This is data bias - it is embedded before the AI ever generates a single response.

Output bias is what the learner experiences: the AI consistently grading experimental, non-Western aesthetic choices as “below standard,” or a career guidance avatar that confidently recommends different professional pathways to learners based on their gender. The output is where the harm occurs, but the root cause is always upstream in the data.

In a practical context, this means developers must audit their source materials with the same critical eye an editor uses on a manuscript. The question to ask before training any AI model is: Who created this content? Who is it about? And crucially, who is missing from it entirely?

3.2 Cultural Neutrality: Balance Without Erasure

Cultural neutrality is often misunderstood as the removal of all cultural specificity from AI content. This is both impossible and undesirable. True cultural neutrality in AI education means achieving a pedagogical balance: providing feedback and content that is globally applicable and actionable, without being culturally prescriptive or implying that one cultural standard is the universal benchmark of excellence.

In a graphic design module, for example, an AI trained primarily on Western Minimalism may rate typography inspired by Japanese Wabi-Sabi aesthetics as “over-designed.” This is not a neutral assessment, it is a culturally-biased one. A culturally neutral AI would instead evaluate the work against its own stated aesthetic intention. For VET professionals operating in a European, multicultural context, this distinction is essential for creating tools that honour the full breadth of their learners’ creative heritages.

3.3 Stereotype Threat: The Psychological Consequence of Biased Feedback

Stereotype Threat, a concept developed by psychologists Claude Steele and Joshua Aronson, describes the psychological risk where a person, aware of a negative stereotype about their group, experiences additional anxiety that impairs their performance. In the context of AI education, if a female learner receives consistent AI feedback that subtly - or overtly - steers her away from technical leadership roles, she does not simply disagree with the suggestion. Research indicates she may begin to internalise the bias, leading to disengagement and reduced professional ambition.

A biased AI is not a neutral tool but an active participant in a learner’s self-conception as a professional. This gives bias mitigation a moral urgency that goes beyond legal compliance. It is a matter of educational equity.

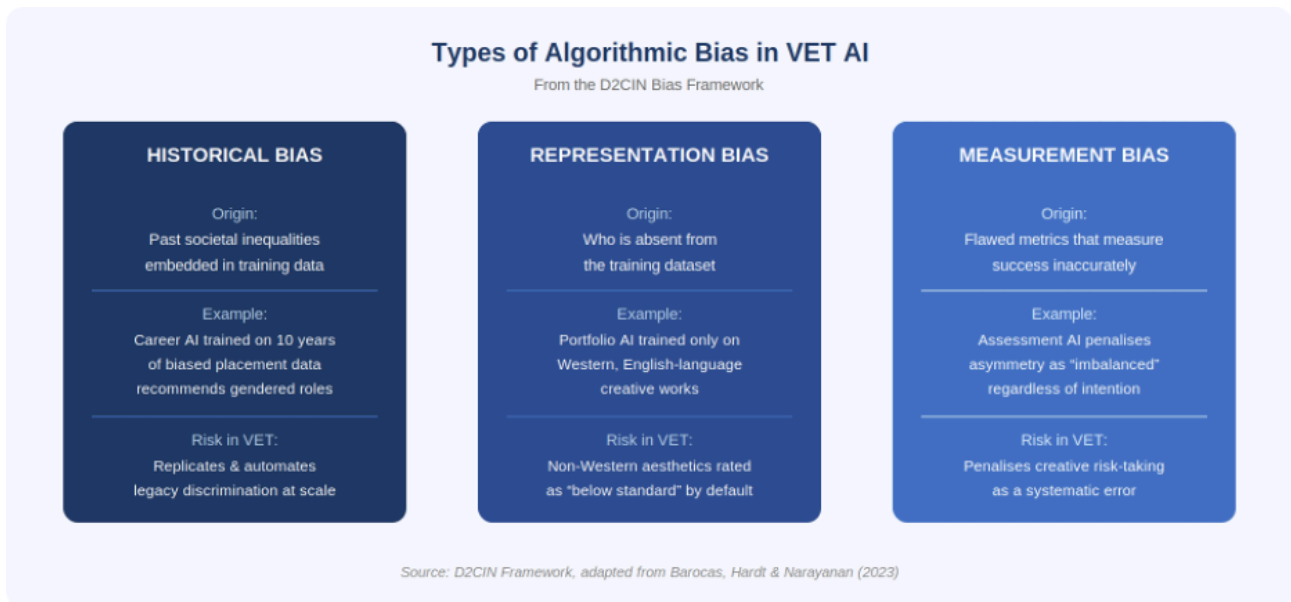


Figure 1.1 The Three Types of Algorithmic Bias in VET AI (D2CIN Framework)

4. Case Study “The Gendered Career Advisor”

The Scenario

A German VET institution specialising in digital media production deployed an AI-powered career guidance assistant designed to help students identify their professional niche within the creative sector. The tool was trained on ten years of successful graduate placement data from the institution. Within the first semester of use, a pattern emerged: the AI consistently recommended “Project Management” and “Client Relations” roles to female-presenting students, while directing male-presenting students towards “Directing,” “Technical Direction,” and “Lead Development.” The system was never explicitly programmed to make gender-based distinctions but it learned them. Because the institution’s historical placement data reflected a decade of unconscious gender bias in its own hiring networks, the AI was, in effect, trained to replicate a systemic problem.

Why it Failed

This case illustrates **Historical Bias** in training data. The AI did not generate discriminatory outputs because it was poorly coded but it did so because the data it learned from was itself the product of a biased professional environment. This is the central paradox of algorithmic bias: the AI was doing

its job perfectly. It was identifying patterns and optimising recommendations - but the pattern it learned to optimise for was gender-based career segregation.

Key Takeaway

Without a proactive dataset audit before deployment, the institution embedded and automated a legacy of inequity into a tool that was presented to learners as objective career guidance.

5. In-Practice Scenario (The "After" Application)

The Representational Fairness Dataset Audit

This audit is a structured quality assurance exercise to be conducted before finalising any AI training dataset or content script. Its purpose is to make invisible patterns of under-representation visible, so that they can be corrected before the AI model is trained.

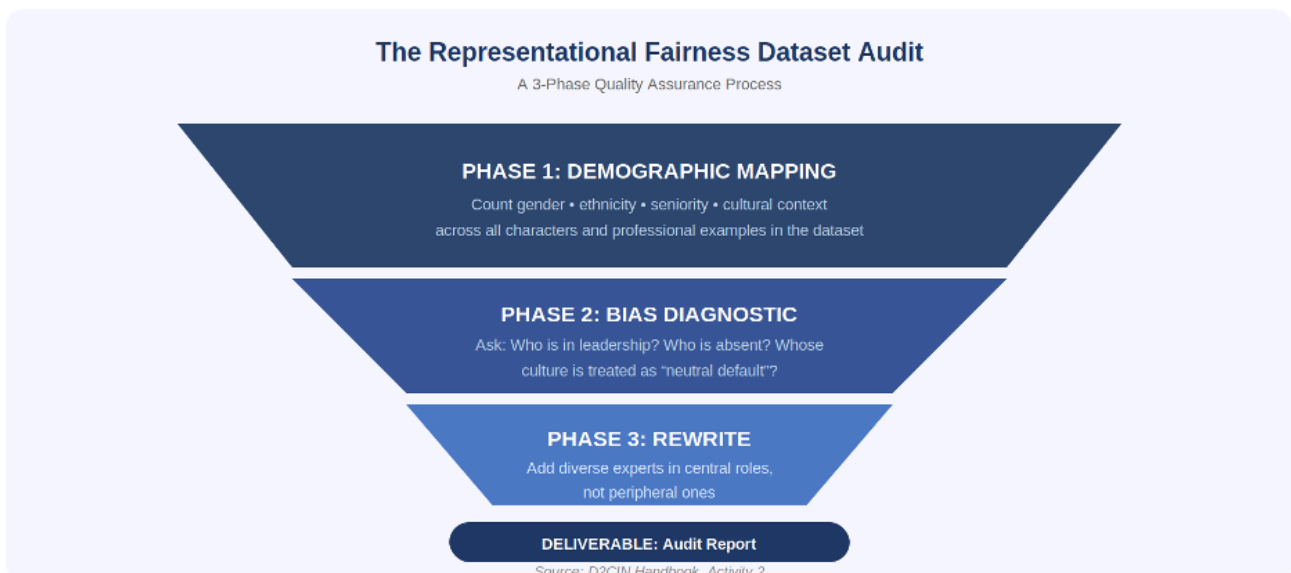


Figure 1.2 - The 3-Phase Representational Fairness Audit Process

• Instructions

Phase 1: The Demographic Mapping (The Count)

Review your training dataset, script, or case study library. For every named character, example professional, or role model referenced, record the following attributes in a simple spreadsheet:

1. Apparent Gender (Male / Female / Non-binary / Not Specified)

2. Apparent Ethnicity or Cultural Background (as inferable from names, locations, and contexts)
3. Role/Seniority Level (Learner / Junior / Senior / Leadership / Expert)
4. Geographic or Cultural Context (European / Non-European / Universal)

Phase 2: The Bias Diagnostic (The Analysis)

Once you have completed the demographic mapping, answer the following diagnostic questions:

5. Are all characters in leadership or expert roles of the same apparent gender or ethnic background?
6. Are characters from non-Western cultural contexts represented only as “problems to solve” or “examples of failure” rather than sources of expertise?
7. Is the dataset missing entire demographic groups entirely? Who is simply not present?
8. Are there any disciplines, art forms, or professional practices that are implicitly treated as “standard” while others are treated as “alternative” or “specialised”?

Phase 3: The Rewrite (The Correction)

For each “Yes” answer in Phase 2, you have identified a representational failure. Your task is to revise the relevant section of your dataset or script. Rewrite a minimum of two case study characters or professional examples to introduce greater demographic diversity. Critically, the diverse characters you add must occupy expert or leadership roles — not supporting or peripheral ones. This is the distinction between Representational Fairness and Tokenism.

• Deliverable

A completed Representational Fairness Audit Report, containing your original demographic mapping spreadsheet, your diagnostic analysis highlighting identified biases, and your revised content with a brief written justification explaining the choices made to improve representational balance.

6. Reflective Question

If we train an AI exclusively on portfolios, graduate data, and professional examples from the last 10 years, whose artistic styles, professional pathways, and cultural perspectives are we systematically and invisibly excluding from the AI’s definition of “excellence”?

7. References

Steele, C.M. & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, 69(5), 797–811.

European Commission. (2021). Proposal for a Regulation laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act). Publications Office of the EU.

D2CIN Project Consortium. (2025). Insights & Recommendations Report (Activity 1.2). Mimic Productions & Budakov Films.

Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.

Chapter 2: Inclusive UX Design

Context & Rationale

The Problem: The integration of Artificial Intelligence in Vocational Education and Training (VET) often focuses heavily on technical capabilities, neglecting the psychological experience of the learner. The D2CIN Needs Analysis revealed a critical competency gap: 89% of target staff lack formal training in inclusive AI practices and design. Furthermore, our Interdisciplinary Focus Groups - specifically the UX & Accessibility Designers across Germany and Bulgaria - highlighted that micro-level UX decisions (such as an avatar's tone of voice, pacing, and facial symmetry) significantly influence user trust and comfort.

The Goal: This chapter addresses the lack of formal training in designing interfaces that cater to neurodiverse learners. It equips VET trainers and developers with actionable design heuristics to manage cognitive bandwidth and foster psychological safety, ensuring AI interfaces facilitate learning rather than creating barriers.

2. Learning Objectives (LOs)

- **LO1:** Design AI interfaces that actively manage "Cognitive Load" to prevent overwhelming learners, ensuring content is digestible for diverse cognitive profiles
- **LO2:** Apply specific design heuristics to minimize the "Uncanny Valley" effect, ensuring AI avatars are perceived as trustworthy and supportive rather than "creepy" or unnerving,

3. Core Theoretical Concepts

To design truly inclusive AI, VET professionals must move beyond basic "usability" and understand the deeper psychological relationship between the learner and the machine. This chapter is grounded in three theoretical pillars identified during the D2CIN research phase, offering a roadmap for creating educational tools that respect the cognitive and emotional needs of diverse learners.

3.1. Cognitive Load Theory in AI: Managing "Mental Bandwidth"



Cognitive Load Theory (CLT) refers to the limited amount of working memory resources available to a learner at any given time. In the context of AI interfaces, it is crucial to distinguish between Intrinsic Load—the effort required to understand the actual lesson—and Extraneous Load, which is the wasted effort spent navigating a confusing interface or processing chaotic visuals. For neurodivergent learners, particularly those with ADHD or dyslexia, high extraneous load can trigger immediate disengagement. An interface that simultaneously presents a moving avatar, scrolling text, and flashing notification icons creates a "sensory bottleneck" that impedes learning.

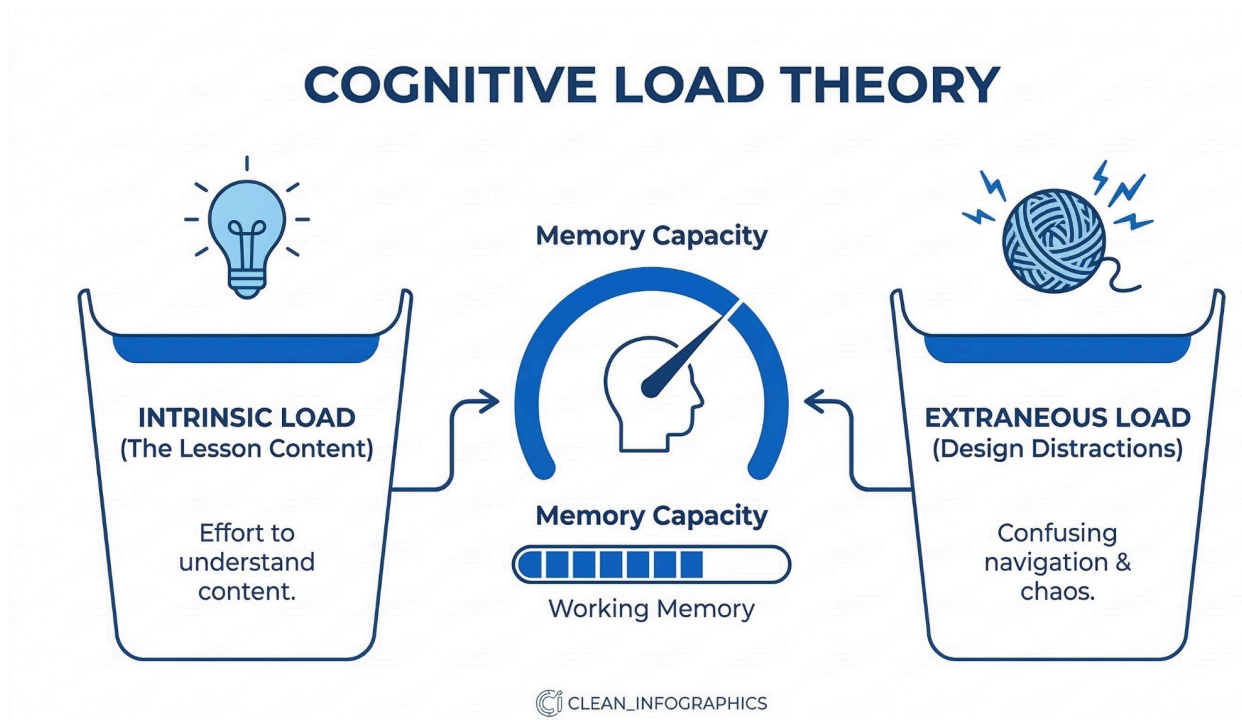


Image: Cognitive Load Theory. Source: Gemini, 2025.

To mitigate this, D2CIN research supports the "Quiet Interface" imperative. As noted by UX experts during the focus groups, good design does not demand attention; it facilitates focus. Developers should avoid cognitive overload caused by excessive customization options. Instead of offering an overwhelming array of settings, the interface should provide a clear choice between a "Rich Media Mode" (Avatar + Voice + Text) and a "Focus Mode" (Text Only + Simplified Layout). This approach empowers the learner to regulate their own mental bandwidth based on their immediate sensory needs.

3.2. The Uncanny Valley: Navigating Visual Discomfort

The "Uncanny Valley" is a psychological phenomenon where a humanoid object, such as an AI avatar, bears a near-identical resemblance to a human but fails in subtle, non-verbal ways. Flaws such as stiff eye movements, delayed lip-syncing, or a lack of micro-expressions can cause a sensation of unease, revulsion, or "creepiness" in the user. In a vocational training scenario, trust is paramount; if an avatar falls into the Uncanny Valley, the learner stops processing the educational content and becomes fixated on the avatar's artificiality. The "teacher" effectively becomes a source of distraction rather than support.

Designers must therefore navigate the "Realism Trap." While there is often a push for hyper-realism, experts caution against this pursuit if the budget does not allow for perfection. Insights from the project suggest that stylized or abstract avatars—such as illustrated characters—often outperform imperfect photorealistic ones. By lowering the user's expectation for strictly human-like behavior, these stylized forms bypass the Uncanny Valley entirely. Consequently, when resources are limited, prioritizing behavioral consistency over visual realism creates a more comfortable and effective learning environment.

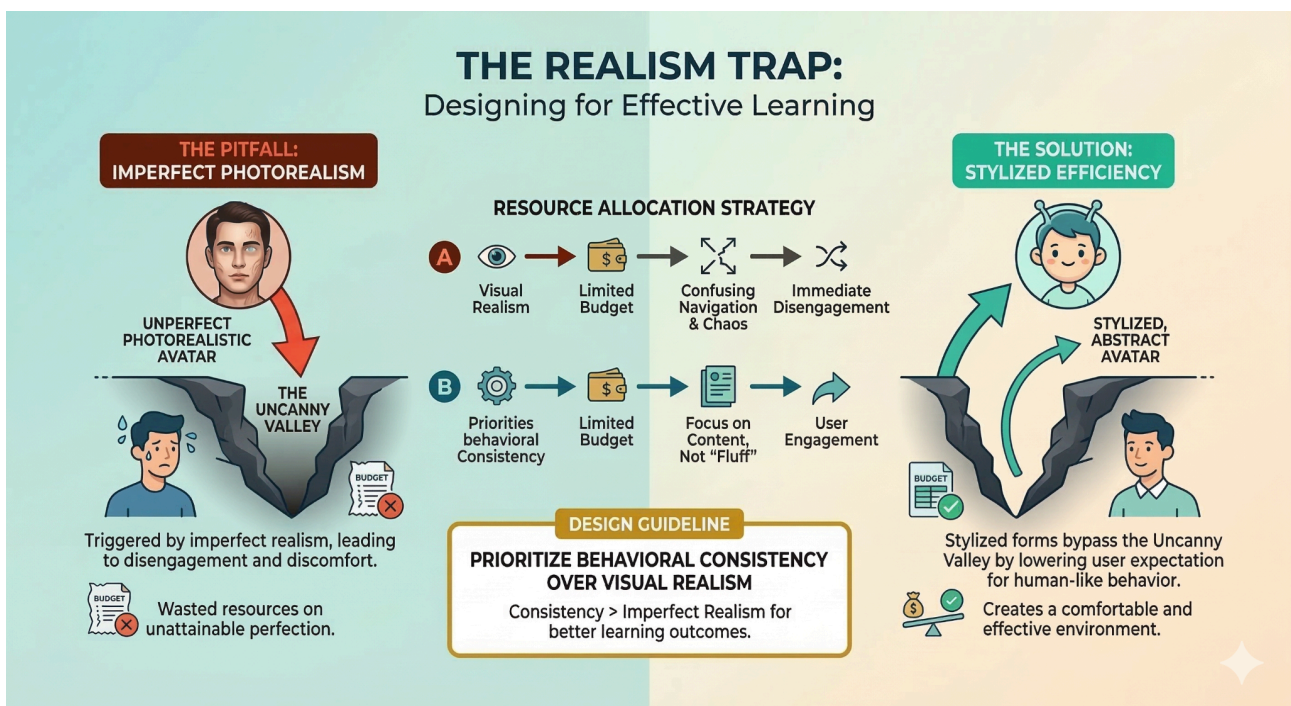


Image: The 'Realism Trap'. Source: Gemini, 2025

3.3. Predictable Interactions & Tiered Explainability

This concept bridges UX Design and AI Ethics, positing that learner trust is built through two key mechanisms: Predictability (the AI behaves consistently) and Explainability (the user understands why the AI acted in a certain way). AI systems in education frequently use adaptive content algorithms to alter difficulty levels or feedback styles in real-time. However, if these changes occur without context, learners—especially those prone to anxiety—may feel they are being judged unfairly or manipulated.

To foster "Psychological Safety," the interface must avoid sudden, unexplained shifts in tone or difficulty. The solution lies in "Tiered Explainability," a design strategy that offers layered transparency. At the primary level, the learner receives a simple, non-technical explanation for the AI's adaptation, such as, "I noticed you paused for a long time, so I switched to simpler vocabulary." At the secondary level, a detailed technical log is available for educators to audit the AI's decision-making process. Ultimately, adaptation must never be hidden; the learner should always retain the agency to override the AI's suggestions, ensuring they remain in control of their educational journey.



Image: Psychological Safety in AI

4. Case Study The "Unblinking Tutor"

The Scenario

To impress stakeholders, a vocational school created a photorealistic AI avatar for a Customer Service module. It looked incredibly human, but technical limits caused two subtle flaws: the avatar maintained intense, unbroken eye contact, and its lip-syncing lagged by just 200 milliseconds.

The Result

During the pilot, students didn't learn. Instead of focusing on the script, they felt "watched" and anxious. Their brains fixated on the avatar's unnatural stare and out-of-sync mouth.

Why It Failed (The Learning Science)

This is a classic example of the Uncanny Valley. When an AI looks almost human but fails in small ways (like not blinking), it stops being a "teacher" and becomes a source of unease.

From a cognitive perspective, this created an Extraneous Cognitive Load. The learners' mental energy was wasted processing the visual errors rather than the lesson content.

Key Takeaway

Better Abstract than Imperfectly Real. If you can't achieve 100% realism, use a stylized or illustrated character. Consistency builds trust; "creepy" realism breaks it.

5. In-Practice Scenario (The "After" Application) *[Actionable Tool (200-300 words). This is the core value of the Handbook. Provide a specific exercise, checklist, or audit tool the reader can use immediately.]*

Objective

To transition from theory to practice, this section provides a standardized quality assurance (QA) protocol. Before finalizing any AI-driven educational tool, developers and instructional designers must conduct this Neuro-Inclusive Interface Audit. This process ensures that the tool supports, rather than overwhelms, the learner's cognitive processing.

Instructions

Review your current AI prototype against the following four criteria. For every "No," apply the suggested remediation strategy.

Phase 1: The Audit Checklist

1. The "Control" Check (User Agency & Pacing)

The Question: Does the learner have absolute control over the avatar's delivery? Can they pause, replay, or slow down the speech rate (e.g., 0.75x speed) at any moment?

The Rationale: Anxiety often stems from a lack of control. Allowing learners to regulate the information flow reduces the "fear of missing out" on key content.

Action: Implement a visible, always-on media control bar.

2. The "Visual Noise" Check (Split-Attention Effect)

The Question: Is the avatar performing "ambient" movements (nodding, shifting weight, blinking rapidly) while critical text is displayed on the screen?

The Rationale: The human brain struggles to process moving images and static text simultaneously. Ambient motion distracts the eye, increasing extraneous cognitive load.

Action: Program the avatar to enter a "neutral/still state" whenever text-heavy slides are presented.

3. The "Tone" Check (Emotional Adaptability)

The Question: Does the interface allow the user to toggle the avatar's persona or communication style? (e.g., switching from "Strict/Professional" to "Supportive/Casual").

The Rationale: Different learners require different emotional scaffolding. A neurodivergent learner may find a direct, authoritative tone overstimulating, whereas a confident learner may prefer efficiency.

Action: Include a simple "Persona Setting" in the onboarding phase.

4. The "Uncanny Valley" Check (Stylization Review)



The Question: Does the avatar’s visual realism match its technical performance? Specifically, is the lip-sync delayed by more than 100ms?

The Rationale: High visual realism creates high expectations for behavior. If the sync is off, the realism becomes a liability (the Uncanny Valley).

Action: If technical latency cannot be fixed, downgrade the visual realism. Switch to a stylized or illustrated character to align user expectations with technical reality.

Phase 2: The Deliverable

Output: The UX Optimization Report

Upon completing the audit, the design team must produce a brief report documenting at least three specific design changes made to accommodate neurodivergent learners. This ensures that inclusion is an active part of the development log, not an afterthought.

Practical Application: The Neuro-Inclusive Interface Audit

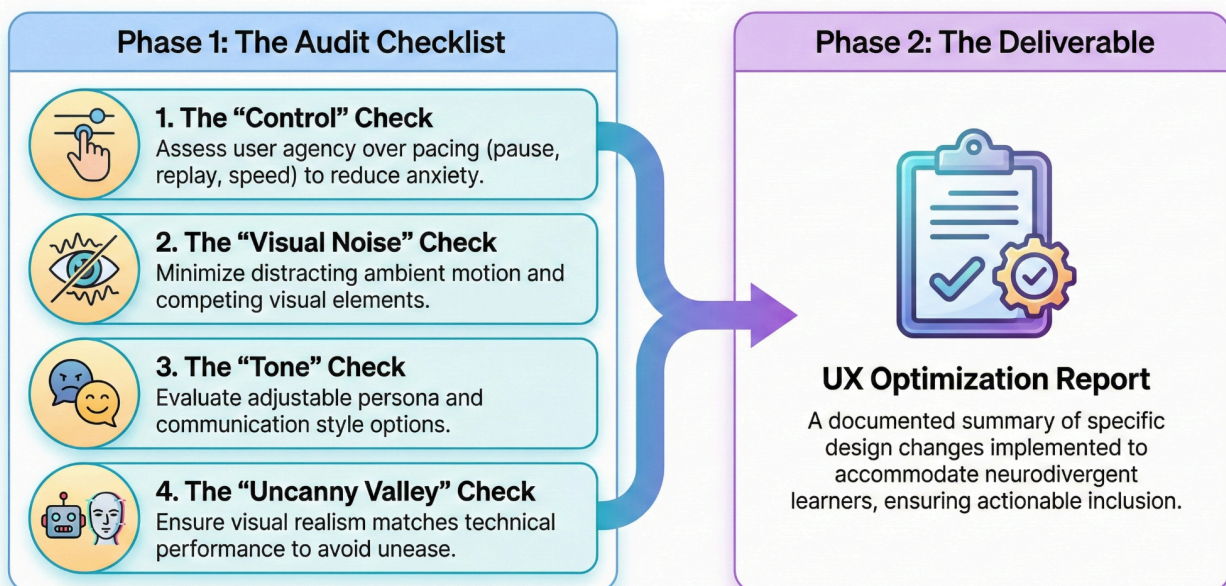


Image: The Neuro-inclusive interface audit.

6. Reflective Question

Does the visual realism of your AI avatar add actual educational value, or is it merely a technical showpiece that distracts the student from the learning material?

7. References

Daniel Kahneman, D. (2011). Thinking, fast and slow. New York, NY: Farrar, Straus and Giroux.

Finale Doshi-Velez, F., & Been Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608. <https://arxiv.org/abs/1702.08608>

Chapter 3: Accessibility Standards in AI

Context & Rationale

The Problem: The D2CIN Needs Analysis identified that many creative VET professionals struggle to meet international accessibility standards when developing conversational AI tools. Furthermore, focus group interviews with UX and accessibility experts highlighted a critical challenge: accessibility in vocational education must go far beyond basic WCAG compliance. Designers often overlook multimodal delivery, the need to support diverse cognitive and linguistic profiles, and the reality of low-bandwidth learning environments.

The Goal: This chapter equips VET trainers and developers with practical methodologies to implement comprehensive multimodal accessibility. It solves the compliance gap by providing actionable steps to ensure AI avatars and adaptive content are ethically designed, transparent, and fully accessible to all learners, regardless of physical, cognitive, or environmental barriers.

2. Learning Objectives (LOs)

LO1: Evaluate and align AI-driven educational tools with WCAG 2.2 AA standards and multimodal accessibility principles (e.g., integrating synchronized captions, keyboard navigation, and screen reader compatibility).

LO2: Apply transparent design strategies to adaptive AI content, ensuring that automated adjustments to a learner's curriculum remain visible, explainable, and under the user's control.

3. Core Theoretical Concepts

Concept A: Multimodal Accessibility

The Simple Definition: Multimodal accessibility means never relying on just one way to communicate information. It is the practice of delivering the same educational content through multiple sensory channels at the same time—such as visual (seeing), auditory (hearing), and cognitive (reading/processing).

Relevance to AI in VET Education: Think of this as the "redundancy rule." In a vocational classroom, you might have learners with hearing impairments, or simply a learner trying to study on a noisy train. If your AI avatar speaks an instruction aloud, the system must simultaneously provide synchronized, accurate closed captions (text). Furthermore, the avatar's visual gestures must support the message without being distracting or overwhelming. True multimodal design also means that if a high-resolution 3D avatar cannot load due to a slow internet connection (low bandwidth), the system automatically provides a lightweight, text-based alternative so the learning can continue uninterrupted.

Concept B: Assistive Technology (AT) Compatibility

The Simple Definition: Assistive Technology refers to the external software or hardware that learners with disabilities already use to navigate the digital world (e.g., screen readers, voice command software, or specialized keyboards). Compatibility means your AI tool "plays nicely" with these existing devices without breaking them.

Relevance to AI in VET Education: It is not enough for an AI tool to be smart; it must meet international standards like WCAG 2.2 AA. For example, if a visually impaired learner uses a screen reader to navigate your VET module, the AI avatar's generated voice must not overlap and "shout over" the screen reader's audio. Furthermore, the entire AI conversation interface must be easily navigable using only a keyboard (without a mouse), ensuring that learners with motor disabilities can fully participate in the training.

Concept C: Transparent Adaptive Content

The Simple Definition: "Adaptive content" means the AI automatically changes the lesson based on how well the student is doing—like lowering the difficulty of a quiz if a student is failing, or changing its pacing. "Transparency" means the AI never does this in secret; it always tells the learner what it is changing and why.

Relevance to AI in VET Education: AI should never be a "black box" that manipulates a learner's educational path without their knowledge. Hidden adaptation can cause frustration and a loss of trust. VET learners must understand why their content changed. For instance, the AI should explicitly say, "I noticed you were struggling with this design concept, so I simplified the text for you." Most importantly, the system must balance personalization with user autonomy—meaning the learner must always have a visible button to override the AI's adjustments and choose their own difficulty level. This aligns with "Tiered Explainability," providing clear, human-readable explanations to students while keeping detailed technical logs available for educators.




Image: Multimodal Accessibility, Source: Gemini 2025

4. Common Case Study (The "Before" Scenario)

• **Title:** The Silent Commuter

• **Scenario:** A creative VET institution developed an advanced, AI-powered conversational avatar to guide graphic design students through complex 3D modeling tutorials. To showcase their technological capability, the developers relied heavily on high-quality voice synthesis and real-time speech evaluation. However, the application was built without closed captions, keyboard navigation, or a text-only alternative. During the pilot phase, a student attempting to complete the module on a noisy commuter train without headphones completely failed to hear the instructions. Another learner, who relies on a screen reader, found that the interactive avatar's graphical overlay blocked critical navigation buttons, rendering the app unusable.

• **Analysis:** This scenario illustrates a severe failure in Multimodal Accessibility and Assistive Tech Compatibility. The developers designed exclusively for a "perfect" sensory environment and an

able-bodied user. By ignoring situational disabilities (a noisy train) and permanent disabilities (visual impairment), they created a tool that actively excluded learners, undermining the core goal of inclusive VET training.

5. In-Practice Scenario (The "After" Application)

Activity Type: Multimodal VET Accessibility Audit Checklist

Instructions: Before finalizing or deploying your AI training module, conduct this 5-point stress test to evaluate its inclusivity. You must physically test your prototype under the following constrained conditions:

1. The "Volume Off" Test: Turn off all audio output. Can the learner successfully complete the entire AI interaction using only synchronized captions or text transcripts?
2. The "Unplugged Mouse" Test: Disconnect your mouse. Can you navigate to the AI avatar, initiate a conversation, and submit answers using exclusively the Tab, Enter, and Space keys?
3. The "Screen Reader" Test: Activate a basic screen reader (e.g., NVDA, VoiceOver). Does the software read out the AI's text prompts clearly, or does it clash with the avatar's generated voice audio?
4. The "Low-Bandwidth" Test: Simulate a slow 3G network connection. If the high-resolution 3D avatar fails to load, is there a lightweight, text-based fallback option automatically provided to the learner?
5. The "Adaptation Transparency" Test: If the AI lowers the difficulty of a quiz based on poor performance, does a clear notification pop up to explain why this happened, alongside a button offering the user the choice to revert?



Image: 5-point physical accessibility audit

Deliverable: A verified Accessibility Compliance Report detailing any failed tests, accompanied by a list of the specific technical adjustments required to meet holistic VET inclusivity standards.

6. Reflective Question

If your AI's most advanced feature (such as 3D rendering or voice recognition) suddenly breaks or cannot load, can the most vulnerable learner in your classroom still achieve the core learning objective?

7. References

World Wide Web Consortium. (2023). Web Content Accessibility Guidelines (WCAG) 2.2.

<https://www.w3.org/TR/WCAG22/>

Shadi Abou-Zahra. (2018). How people with disabilities use the web. World Wide Web Consortium (W3C). <https://www.w3.org/WAI/people-use-web/>

European Commission. (2024). Ethics guidelines for trustworthy AI. Publications Office of the European Union. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

Batya Friedman, B., & David G. Hendry, D. G. (2019). Value sensitive design: Shaping technology with moral imagination. Cambridge, MA: MIT Press.

Chapter 4: AI Ethics and Compliance

Context & Rationale

The Problem: With the EU AI Act entering into force, ethical and legal compliance in AI is no longer optional but a professional obligation. Yet the D2CIN Needs Analysis identified a significant gap: the majority of target staff lack formal training in AI ethics. Rather than an abstract philosophical deficiency, it manifests in real, daily decisions. Practitioners do not know what data they are legally permitted to collect, how to communicate AI limitations to learners, or how to respond when an AI produces an outcome a student perceives as unfair. Furthermore, the D2CIN Focus Groups - particularly ethics and education specialists in Germany - emphasised that the single most corrosive force in AI-assisted education is opacity. When learners cannot understand why an AI made a decision about their work or their pathway, trust collapses entirely (Source: D2CIN Insights Report, Section 5.3).

The Goal: This chapter analyses AI ethics from an abstract compliance framework into a set of concrete, daily practices. It equips VET professionals to apply two foundational principles - Tiered Explainability and Privacy-by-Design - directly within their project workflows. By the end of this chapter, the reader will be able to communicate AI decisions transparently to learners, draft legally sound data consent communications, and design AI interactions that build, rather than erode, institutional trust.

2. Learning Objectives (LOs)

- **LO1:** Apply "Tiered Explainability" to help learners and administrators understand why an AI model gave a specific grade or recommendation.
- **LO2:** Implement "Privacy-by-Design" principles to minimize data collection and manage learner consent regarding voice and behavioral data.

3. Core Theoretical Concepts

Ethical AI in education is not a compliance checklist to be completed once before launch. It is an ongoing design philosophy that must be embedded into every decision about how an AI tool collects data, communicates decisions, and manages the power dynamic between the system and the learner.

3.1 Tiered Explainability: Earning Trust Through Transparency

Tiered Explainability is a design principle that requires AI systems to offer layered, audience-appropriate explanations for their decisions. The premise is straightforward: different stakeholders need different depths of information to feel that an AI decision is legitimate and fair.

In practice, this creates two tiers of explanation. The primary tier is designed for the learner: a simple, non-technical, human-readable statement that explains the AI's action in plain language. The secondary tier is designed for educators and administrators: a detailed, auditable log of the specific data points, weighting criteria, and model parameters that produced the output.

Without tiered explainability, an AI grade is not feedback but a verdict. Learners who cannot understand why they received a particular assessment cannot improve. More critically, without a secondary audit log, educators cannot identify when the AI has made a systematic error and cannot correct it.

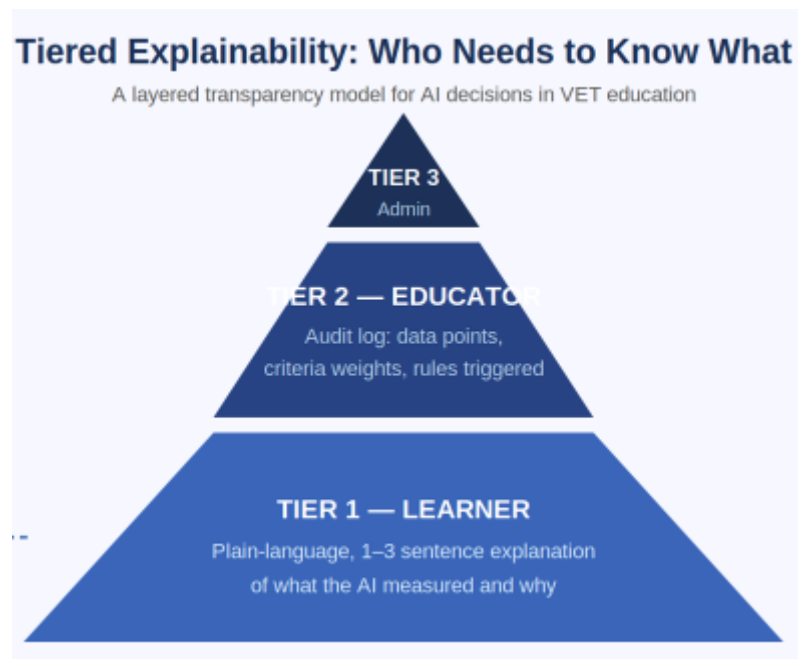


Figure 4.1 - Tiered Explainability: Three Levels of AI Transparency for VET

3.2 Data Minimisation: Collecting Only What Is Necessary

Data Minimisation is a core principle of the EU General Data Protection Regulation (GDPR) and is echoed in the EU AI Act's risk-based framework. It holds that an AI system should collect only the data that is strictly and demonstrably necessary for its stated educational function. Every additional data point collected is an additional ethical and legal liability.

In the context of VET AI tools, this principle has immediate practical implications. A pronunciation coaching AI that listens to a learner’s voice to provide feedback on stress and intonation does not require a recording of that voice to be stored permanently. The acoustic analysis can be performed in real time, the score generated, and the audio discarded.

Many VET developers collect extensive behavioural data under the assumption that more data produces better AI. Data Minimisation challenges this assumption and asks: is the educational improvement worth the privacy cost?

3.3 Transparency: Maintaining the Human-AI Boundary

Transparency, in the context of AI ethics, refers to the obligation to clearly and consistently inform users that they are interacting with an automated system, not a human being. The EU AI Act specifically addresses this in its provisions on AI systems that interact directly with natural persons, requiring that users be informed they are engaging with AI unless this is “obvious from the context.” In an educational setting, context is rarely obvious.

Transparency is not merely a legal obligation but a foundational element of the learner’s informed consent. An AI that presents itself as a human mentor, even implicitly, through design choices such as a human name, a photorealistic face, and emotional language, is manipulating the trust relationship in a way that undermines the integrity of the educational experience.

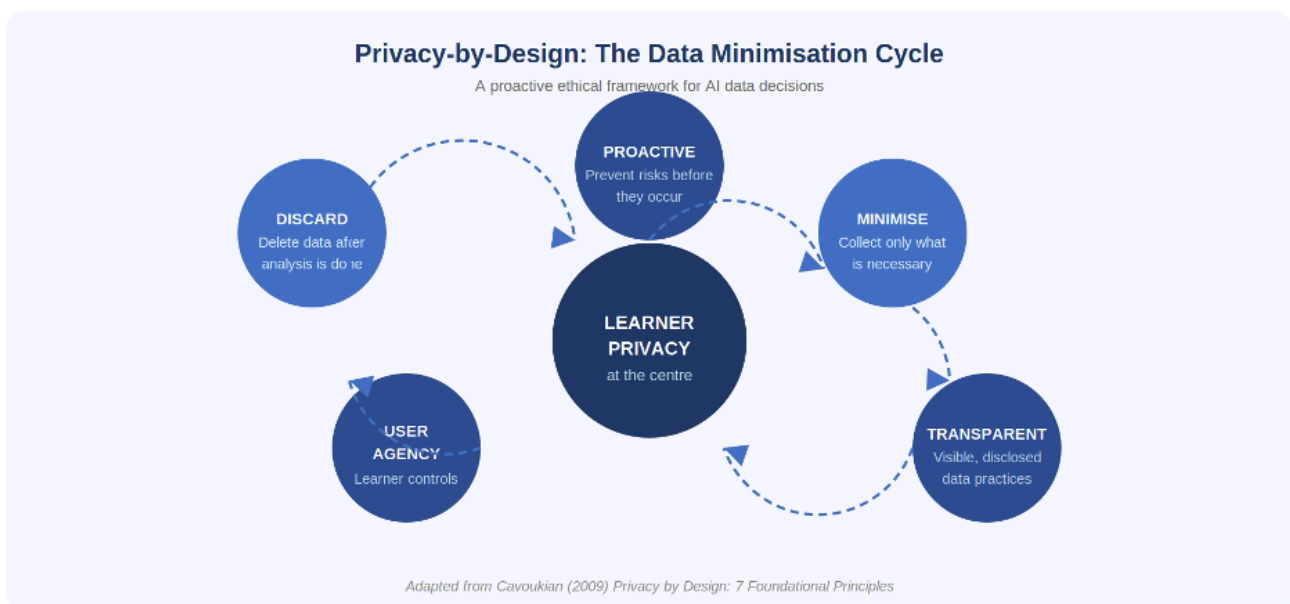


Figure 4.2 The Privacy-by-Design Cycle: Five Principles for Ethical Data Management

4. Case Study (The "Before" Scenario)

- **Title:** "The Black Box Grade:

Scenario

A vocational design school in Bulgaria introduced an AI-powered assessment tool to evaluate student graphic design projects. At the end of each module, students received a numerical score and a single-word rating: "Excellent," "Proficient," or "Developing." No specific feedback was provided. No criteria weighting was disclosed. One student, whose project had been singled out by their human tutor as particularly innovative for its unconventional use of negative space, received a "Developing" rating from the AI. Unable to understand why an approach their tutor praised had generated a failing grade, the student submitted a formal complaint. A review revealed that the AI's composition analysis module penalised asymmetry as "imbalanced" regardless of intentionality - a design bias baked into its training data. However, because there was no explanation layer, this systematic error had gone undetected for three full semesters.

Why this Failed

This scenario highlights the dual failure of an AI system without Tiered Explainability. First, the learner was harmed: they received a grade they could not understand, could not learn from, and could not challenge. Second, the institution was harmed: a systematic AI error, one that penalised creative risk-taking, arguably the most valuable skill in a creative VET context, went undetected for months because there was no audit trail. An unexplained AI decision is not simply unhelpful but ethically indefensible.

5. In-Practice Scenario (The "After" Application)

The Ethical AI Communication Design Exercise

This two-part exercise requires the reader to design the transparency and consent infrastructure for an AI-powered VET learning tool. It operationalises both Tiered Explainability and Privacy-by-Design into concrete, writable outputs.

- **Instructions**
- 

Part 1: The Tiered Explainability Protocol

You are developing an AI assessment tool that evaluates student video productions on three criteria: narrative coherence, technical execution, and accessibility compliance (e.g., the inclusion of subtitles). For each of the three assessment criteria, write two versions of the feedback output:

- a. The Student-Facing Explanation (Tier 1): Write a plain-language statement of no more than three sentences that tells the student exactly what the AI measured, what it found, and what a specific improvement step would look like. The statement must not use technical jargon and must be written in a constructive, non-punitive tone.
- b. The Educator Audit Log Entry (Tier 2): Write a structured log entry for the same assessment result. This entry should include the specific data point the AI evaluated, the weighting assigned to this criterion, and the model rule that was triggered. This log is the audit trail that allows the educator to review and challenge any AI decision.

Part 2: The Data Consent Communication

You are launching the AI assessment tool described above. Before a student submits their first video project, the system must present a data consent notification. Using the Data Minimisation principle as your guide, write a plain-language consent pop-up that:

- a. Clearly identifies all data points the AI will collect from the student's submission.
 - b. Explicitly states what data is analysed but NOT retained (e.g., "We analyse your audio track to check for narration pace, but we do not store the audio recording after analysis is complete.").
 - c. Clearly states what data IS retained and for how long.
 - d. Includes a clearly labelled option for the student to opt out of specific forms of data collection.
- **Deliverable:** A completed Ethics Communication Package, containing: six feedback statements (three Tier 1 student explanations and three Tier 2 audit log entries) and one plain-language consent pop-up text. This package demonstrates that ethical AI communication is a layered design challenge requiring different writing registers for different audiences.

6. Reflective Question

Is the learner in your AI tool explicitly and clearly informed at the moment of first interaction that they are engaging with an automated system - or does the interface's design, name, and conversational tone subtly lead them to believe they are speaking with a human expert?

7. References

Cavoukian, A. (2009). Privacy by Design: The 7 Foundational Principles. Information and Privacy Commissioner of Ontario.

D2CIN Project Consortium. (2025). Insights & Recommendations Report (Activity 1.2). Mimic Productions & Budakov Films.

Doshi-Velez, F. & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv:1702.08608.

European Parliament. (2024). Regulation (EU) 2024/1689 — Artificial Intelligence Act. Official Journal of the European Union.

European Parliament & Council. (2016). Regulation (EU) 2016/679 (GDPR). Official Journal of the European Union.

Chapter 5: Digital Storytelling for Inclusive Education

Context & Rationale

The Problem: Educational technology often treats diversity as a visual specification: include a character with a headscarf, add a wheelchair user to the background, ensure the avatar's skin tones span a gradient. This approach, which the D2CIN focus groups explicitly identified as "tokenism," is not merely ineffective but it is actively harmful. It signals to learners from underrepresented groups that they are afterthoughts in the educational narrative, not its protagonists. Furthermore, the integration of AI into VET storytelling has introduced new risks: AI-generated characters and scenarios, when trained on mainstream media, tend to reproduce dominant cultural stereotypes at scale and with algorithmic authority (Source: D2CIN Insights Report).

The Goal: This chapter provides VET content creators and trainers with a structured creative framework for developing inclusive digital narratives that use AI as a tool of empowerment, not homogenisation. It moves practitioners from passive diversity - the presence of diverse characters - to active inclusion: diverse characters who drive narratives, possess expertise, and grant learners meaningful agency over their educational journey.

2. Learning Objectives (LOs)

- **LO1:** Create diverse narratives that reflect complex identities, avoiding "tokenism" (superficial diversity) in VET scenarios.
- **LO2:** Design Narrative Branching that gives learners agency to control the story's direction or the identity of their AI mentor.

3. Core Theoretical Concepts

Inclusive digital storytelling is not simply a matter of casting. It requires a foundational shift in how we understand the relationship between narrative, identity and learning. The following three concepts provide the creative and ethical framework for this chapter.

3.1 Authentic Representation: Depth Over Decoration

Authentic Representation is the practice of creating characters and scenarios that reflect the full complexity of human experience, rather than relying on shorthand visual or cultural stereotypes.

An authentically represented character has a specific professional history, a coherent set of values, a cultural context that informs their decision-making, and - critically - expertise that is central to the narrative's progression. Authentic representation is the opposite of "set decoration," where diverse characters are visible but voiceless.

In AI-driven VET scenarios, characters function as proxy mentors. If every AI mentor who demonstrates expertise is of the same demographic profile - typically Western, male, and able-bodied - the learner internalises that profile as the norm of professional competence. Authentic representation in AI educational content is, therefore, a form of vocational aspiration-setting with measurable impact on learner ambition and self-efficacy.

3.2 Narrative Agency: Moving Beyond the Linear Script

Narrative Agency is the degree to which a learner can meaningfully influence the direction and outcome of an educational scenario. True narrative agency occurs when the learner's choices produce meaningfully different and equally valid educational outcomes. In a conflict resolution scenario, choosing to de-escalate through empathy and choosing to de-escalate through establishing clear professional boundaries should both be treated as legitimate, context-appropriate strategies. The AI's role is not to judge which is "correct," but to simulate realistic consequences of each approach.

Narrative agency is particularly important for learners from diverse cultural backgrounds whose problem-solving instincts may diverge from the "standard" approach encoded in a linear AI script. An AI that penalises culturally-informed, perfectly valid professional responses because they do not match its pre-defined optimal path is not neutral assessment but a cultural bias operationalised as educational technology.

3.3 Tokenism: Recognising the Difference Between Presence and Inclusion

Tokenism is the practice of making only a superficial or symbolic effort toward diversity - including members of underrepresented groups in a narrative in ways that are visible but not meaningful. In digital storytelling, tokenism manifests when diverse characters are present in a scenario but occupy peripheral roles: they are silent witnesses, supportive sidekicks, or the subject of the problem rather than the agent of the solution.

In AI-driven educational content, tokenism is easy to produce and difficult to detect. The diagnostic question for VET developers is not "Are diverse characters present?" but "Are diverse characters driving the plot, demonstrating expertise, and modelling professional success?"



Figure 5.1 From Tokenism to Authentic Inclusion: A Four-Level Spectrum

4. Common Case Study (The "Before" Scenario)

- **Title:** "The Token Teammate"
- **Scenario:** A VET institution in Germany developed an AI-powered simulation for a Project Management module. The scenario placed learners in a fictional creative agency, navigating a high-pressure product launch. The team was composed of four AI-generated characters: Daniel (Team Leader, white male, 40s), Marcus (Technical Lead, white male, 30s), Sophie (Client Liaison, white female, 30s), and Aiko (Visual Designer, East Asian female, 20s, described only as "the creative one"). Throughout the seven-scenario module, Aiko appeared in four scenes. In each, she nodded silently while Daniel explained the creative strategy, delivered a single line of agreement ("I can make that work"), and was never given a scene in which her visual expertise drove a decision. When learners were asked to evaluate the team's performance, several noted they were "not sure what Aiko's role actually was."

Why this Failed

Aiko is a textbook case of tokenism. She is present, she satisfies a demographic diversity check, but she is not included. Her professional identity is reduced to a single adjective and her narrative function is purely decorative: to make the team visually diverse without shifting any narrative

power toward her. This has two harmful consequences. First, learners from similar backgrounds to Aiko see a model of professional participation that confirms their marginalisation. Second, all learners internalise a model of creative project management in which visual expertise, implicitly coded as non-Western and female, is peripheral to strategic authority.

5. In-Practice Scenario (The "After" Application)

The Inclusive Branching Script Workshop

This creative writing exercise requires the reader to design a short, branching narrative scenario for a VET context that applies Authentic Representation and Narrative Agency simultaneously.

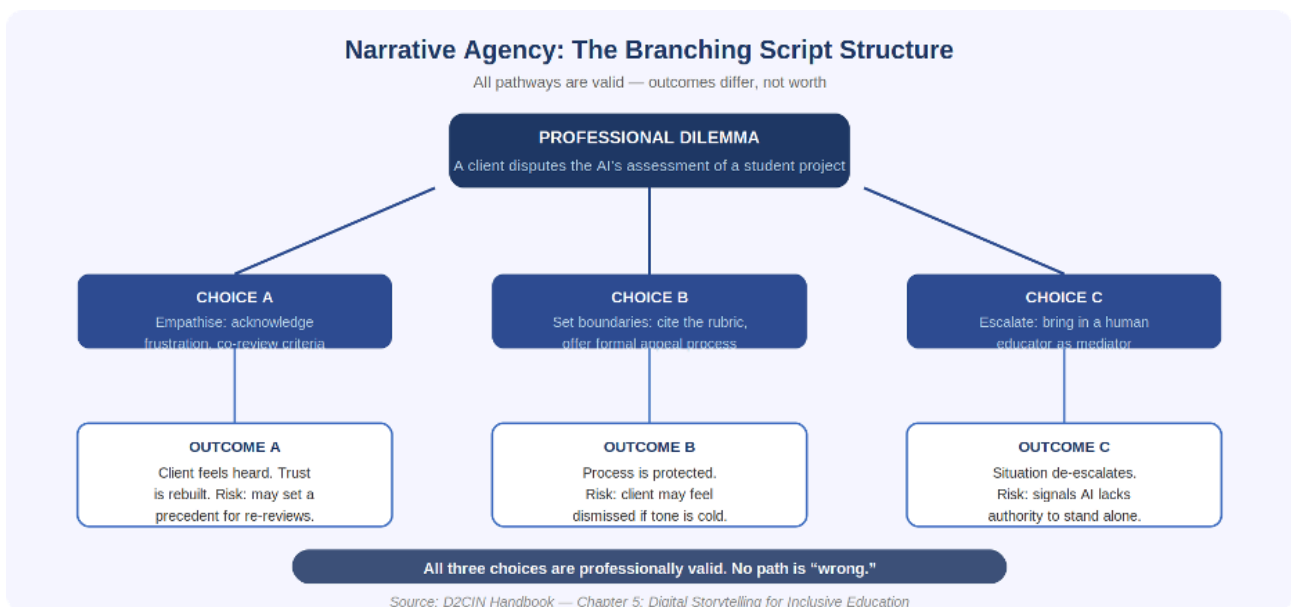


Figure 5.2 The Branching Script Structure: All Pathways Are Valid

• Instructions

Phase 1: Character Design (Applying Authentic Representation)

Design the two primary characters your learner will interact with in the scenario. For each character, you must define the following:

- a. Name, age, and professional background: Give the character a specific, believable professional history - not just a job title.

- b. Cultural context: Identify one aspect of the character's cultural background that meaningfully informs how they approach the professional challenge in the scenario.
- c. Domain of expertise: Define the specific knowledge this character contributes that no other character in the scenario can provide.
- d. Failure mode: Define one way this character's approach, however legitimate, might not be appropriate in every context.

Phase 2: Scenario Design (Applying Narrative Agency)

Write a short scenario (200–300 words) set in a VET-relevant professional context. The scenario must:

- a. Present the learner with a genuine professional dilemma that has no single objectively correct answer.
- b. Offer the learner a minimum of three meaningfully different response options - none of which should be framed as "obviously wrong."
- c. Include a brief description (2–3 sentences each) of the different narrative consequences that follow from each choice.

Phase 3: The Tokenism Audit (Self-Review)

Once you have drafted your scenario, apply the following four-question audit before finalisation:

- a. Do any of my diverse characters exist only to be helped, advised, or rescued by the other characters, rather than to drive the plot forward with their own expertise?
- b. Is the cultural context I assigned to my diverse characters a genuine aspect of their professional identity, or is it a visual or linguistic shorthand?
- c. Does my branching structure genuinely allow for multiple valid professional approaches, or does it subtly reward one cultural norm of professional behaviour over others?
- d. If I removed my diverse characters from the scenario entirely, would the core narrative or learning objective change? (If not, they are set decorations.)

• **Deliverable:** A completed Inclusive Scenario Package, containing: two character design profiles, the full branching scenario script (including all three choice pathways and their consequences), and a completed Phase 3 Tokenism Audit with written responses to all four diagnostic questions.

6. Reflective Question

In the narrative scenario you have designed, do the characters from underrepresented groups drive the plot forward with their specific expertise, or are they present simply to make the scene look diverse? If you removed them from the story entirely, would the learning objective change?

7. References

Hooks, b. (1994). *Teaching to Transgress: Education as the Practice of Freedom*. Routledge.

Steele, C.M. (2010). *Whistling Vivaldi: How Stereotypes Affect Us and What We Can Do*. W. W. Norton & Company.

D2CIN Project Consortium. (2025). *Insights & Recommendations Report (Activity 1.2)*. Mimic Productions & Budakov Films.

Murray, J.H. (1997). *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. The Free Press.

Conclusion

This Handbook was built on a conviction: that inclusive, ethical AI in vocational education is not a future aspiration but it is a present professional obligation. The five chapters that precede this conclusion were not designed as theory to be absorbed and filed away. They were engineered, on the basis of rigorous evidence gathered across two countries and seven disciplinary perspectives, as a practical toolkit for people who build and teach with AI every day.

The journey through this Handbook has moved from the invisible to the actionable. Chapter 1 made bias visible - tracing it from skewed training data through to the career recommendations a learner receives, and equipping practitioners with the Representational Fairness Audit to catch and correct it before deployment. Chapter 2 made the learner's psychological experience the design priority, providing heuristics for managing cognitive load and navigating the Uncanny Valley in AI avatar design. Chapter 3 moved beyond the checkbox of WCAG compliance to a holistic model of Multimodal Accessibility, insisting that an AI tool that fails the "Volume Off" test is not an accessible tool, regardless of its certification. Chapter 4 took the opacity out of AI ethics, replacing abstract principles with writable, deployable outputs: Tiered Explainability protocols, Data Minimisation consent forms, and audit logs that give learners and educators the transparency they are legally and pedagogically owed. And Chapter 5 gave inclusive design its creative dimension - demonstrating that diverse representation is not a casting decision but a narrative craft, with specific, learnable techniques for building stories where all characters drive the plot.

From Handbook to Practice: A Commitment to Action

This Handbook must not be treated as a passive reading resource. It is an active implementation tool, designed to be used alongside real projects, real prototypes, and real classrooms. The five chapters map directly onto the workflow of a VET professional: audit before you train; design before you build; test before you deploy; explain before you assess; and tell stories that empower rather than decorate.

For content creators and developers, this means using the Bias Recognition checklist (Chapter 1) and the Multimodal Accessibility Audit (Chapter 3) as mandatory quality assurance gates before any AI module reaches a learner. For educators and instructional designers, it means using the Tiered Explainability protocols (Chapter 4) and the Inclusive UX heuristics (Chapter 2) to evaluate and improve the tools they are handed, not merely to accept and deploy them. For storytellers and narrative designers, it means applying the Branching Script Workshop and the Tokenism Audit (Chapter 5) to every scenario, every character, every choice structure - asking not just whether diverse characters are present, but whether they are essential.

These are not aspirational goals for an ideal future. They are the minimum standard for AI-assisted VET in the European context, where the EU AI Act, GDPR, and the Erasmus+ programme's priorities around Digital Readiness and Resilience have together established a clear mandate: AI in education must be ethical, explainable, and inclusive by design.

The Road Ahead

This Handbook is the foundation, not the ceiling. It will serve as the primary curricular resource for the D2CIN Interactive AI Workshop Series (Activity 3), where its guidelines will be tested and refined by 40 external participants across Germany and Bulgaria. The pilot workshops will provide a final validation layer, checking not only whether the frameworks work in theory, but whether they hold up in the hands of the VET practitioners who need them most.

Beyond the project period, the methodology embedded in this Handbook is designed to scale. The Representational Fairness Audit, the Neuro-Inclusive Interface Audit, the Tiered Explainability Protocol, and the Inclusive Branching Script Workshop are not proprietary to D2CIN. They are transferable tools, offered under the Creative Commons Attribution-Non-Commercial-Share-Alike 4.0 International License, and intended to be adapted, translated, and refined by any VET institution committed to building AI that serves all learners equitably.

As AI technology continues to evolve - as avatars become more convincing, as adaptive algorithms become more powerful, as data collection becomes more pervasive - the frameworks in this Handbook will require updating. The specific methods will change; the underlying principles will not. Bias is not a technical bug to be patched. Accessibility is not a compliance checkbox to be ticked. Transparency is not a marketing commitment to be managed. And inclusion is not a visual specification to be met.

They are, each of them, ongoing professional practices - habits of mind and habits of craft that must be renewed with every dataset, every design decision, and every story told.

“This can’t be just made by tech people, actual teachers must oversee it.” —
German Focus Group participant, D2CIN Activity A1.2, October 2025

That principle - that inclusive AI cannot be built in isolation, that it requires the sustained, critical involvement of the educators, learners, and communities it is designed to serve - is

the final conclusion of this Handbook. The tools are here. The evidence is documented.
The obligation is shared.

Now the work begins.