Natural Language as the Bridge Between Pedagogy and AI

Paper 2 of 4: Sustainable Frameworks For Audience Tailoring

Dr Joe Watkins

Grassroots Intelligence Limited, Canterbury, UK

Correspondence: joe@grassroots-intelligence.com

Abstract

Paper 1 introduced a lightweight framework for encoding pedagogical expertise as discrete elements with natural language parameters, demonstrating how this structure enables constrained LLM generation of educational content at scale. This paper extends that architecture to audience tailoring, showing how the same pedagogy can produce fundamentally different learning experiences for different learners, without any change to the underlying pedagogical design.

We introduce the concept of a ‘variant profile’, which is a structured set of learner-facing natural language rules that can be applied to a pedagogical content item to tailor its output for a specific individual or group. Related to this is our preferred methodology of creating a ‘primary variant’ and then using variant profiles to tailor this. Although not the main subject, this paper addresses the professional debate about how far tailoring should go, and argues that the framework makes that debate more tractable by making tailoring decisions explicit, testable, and reversible. Paper 3 extends the variant profile concept beyond text to visual and spatial design, covering deterministic and non-deterministic trade-offs in considerably more detail. Paper 4 extends the framework further to temporal sequencing and narration.

Keywords: audience tailoring, variant profiles, differentiation, EAL learners, neurodiversity, constrained LLM generation, prompt engineering, educational technology.

1 Introduction

Paper 1 made the case that pedagogical expertise, expressed as labelled elements with natural language parameters, can serve as structured constraints for LLM content generation. The framework was deliberately audience-agnostic, with the pedagogies in Paper 1 meant to specify *what the content should do*, but not *who it is for*, and this separation was fully intentional.

A pedagogy that encodes a teacher’s approach is, in principle, applicable to any learner. If audience parameters are woven into the pedagogy itself, the same teacher then needs a separate version for every learner group they serve, which makes the framework less manageable and loses precisely the benefit of having a reusable, shareable pedagogical structure in the first place.

This paper, like Paper 1, follows the tradition of design-based research (Barab & Squire, 2004). The observations and design claims presented here are grounded in sustained practitioner experience delivering AI professional development with teachers across primary, secondary, further, and adult education. They reflect what we have consistently observed through extensive, sustained work with real educators producing real content.

Section 2 recaps the key definitions and the central example from Paper 1. Section 3 surveys the research on what tailoring can achieve and why it matters, and addresses the professional debate about how far tailoring should go. Section 4 introduces the simplest form of tailoring, setting audience as a single variable, and describes batch tailoring strategies with ready-to-use prompts, focusing in particular on our definition of the ‘primary variant’ as a standalone pedagogical content item. Section 5 introduces the full ‘variant profile’ architecture. Section 6 describes the novel workflows that emerge from this architecture, including the distinction between non-deterministic and deterministic tailoring and the ‘generate-and-select’ workflow. Section 7 closes with a reflection on empirical testing and the role students can play in shaping their own variant profiles.

2 Recap: Elements, Parameters, and the Pedagogy-First Principle

Paper 1 argued from two foundational pillars. First, LLMs produce substantially better output when given structured, rule-based constraints than when given open-ended instructions, a finding supported by a growing body of research in prompt engineering and constrained generation (Sahoo et al., 2024; Banerjee et al., 2025; Qian, 2025). Second, pedagogical expertise is frequently expressed in rule-like, natural language terms, where experienced teachers describe their approaches in exactly the kind of structured directives that LLMs are designed to interpret and follow. The framework presented in Paper 1 connects these two observations, providing a grammar for capturing a teacher's pedagogical rules and passing them to an LLM as structured constraints on content generation.

For convenience, we reproduce the core definitions here, as they are used consistently throughout this series.

Term	Definition
Element	A discrete, individually-constrained component of a pedagogy. Each element encodes a specific part of the pedagogical approach and is given a descriptive label: **[Definition]**, **[Example]**, **[Surprising Fact]**, **[Retrieval Question]**, and so on.
Parameter	Any individually-adjustable property within a parameter set: its rule, its word count, its relationship to other elements, often expressed in natural-language terms.
Pedagogy	Any combination of elements and parameters using the definitions above.
Variable	Input information applied to the whole pedagogy rather than individual elements. Topic and source material are the canonical examples from Paper 1. This paper extends the concept significantly.
Pedagogical content	Any populated pedagogy, where a topic or other input variables have been entered.
Source material	Subject-matter content that a teacher may optionally provide alongside the pedagogy: lesson notes, schemes of work, textbook passages, or any other trusted reference. We treat this as an optional variable and do not cover this topic extensively in this paper.

Paper 1 introduced four dimensions along which pedagogies can be tested and compared: controlled variation (holding all parameters constant except one), granularity (isolating a single word count, inter-element rule, or structural decision), scale (applying any variation across a full topic library in a batch operation), and shareability (encoding a pedagogy as a transferable artefact that colleagues can adopt and adapt). These dimensions apply equally to audience tailoring, and they underpin the batching strategies described in Section 4.

To ground the discussion throughout this paper, we reproduce the central three-element pedagogy from Paper 1, which serves as the running example across the series. This is meant to be a simple example of a complete pedagogy that is expressed in the architecture described above.

Pedagogy A: “Definition + Example + Surprising Fact”

Element	Parameter 1: Rule	Parameter 2: Max words
[Definition]	A clear, plain-English definition of the key term.	30
[Example]	A concrete, everyday example that illustrates the concept defined in **[Definition]**. Must use a tangible, physical scenario the learner can visualise.	40
[Surprising Fact]	A surprising or counterintuitive fact that challenges or complicates the picture established by **[Definition]** and **[Example]**. Must not simply add information; must make the learner reconsider.	35

Throughout Paper 1, we noted that this pedagogy does not specify a topic, and it also carries no audience information whatsoever. It says nothing about the topic to be covered, nor the learner’s age, language, reading level, cultural context, or prior knowledge. That omission is not an oversight but the principle that makes the framework scalable.

2.1 Implementing Pedagogies with LLMs

Before turning to the question of audience, it is worth dwelling on what the framework enables for teachers working at the pedagogy level itself. A teacher who has encoded their approach to introducing a scientific concept has done something that previously had no formal expression. By taking the time to define – and refine – their understanding in the form of a pedagogy, they have made their craft knowledge visible, portable, and testable. The process of encoding that knowledge is itself a form of what Schön (1983) called reflection-on-action, since the teacher is articulating decisions that were previously implicit. In Paper 1 we argued that it is this step of encoding expertise in the form of simple, natural-language rules which readies this craft knowledge for the sheer computational power and flexibility of LLMs.

The following prompt illustrates Pedagogy A above, with a topic entered as a variable at the end of the prompt.

> 📋 **COPY-PASTE PROMPT: Example three-element pedagogy with topic**
>
> *Please take the following pedagogy and topic and generate as described:*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> *Topic:* **<<TOPIC>>**

Written once, this prompt can then be deployed in many different forms that are highly accessible for teachers. Straightforward examples include local prompt dictionaries, or shared prompts in custom aspects of popular platforms, for example the Prompt Gallery in Microsoft Copilot. Other solutions which are only marginally more technical include CustomGPTs and Copilot Agents, all of which can be created quickly and shared with minimal – often zero – extra cost to existing licenses.

With this Rubicon crossed, it is easy to see how quickly a Department could take an example pedagogy as the one above, and use it as a method to help standardise the approaches to whole aspects of their curriculum. In practical terms, this can happen for zero additional human effort. Paper 1 gave several examples of such batch workflows, a category of content generation that could not have existed prior to the popularisation of LLMs. We revisit these workflows in Section 6, where we explore how LLM-enabled processes can maximise the application of teacher expertise while minimising the time required.

3 Why Tailoring Matters: Research and Professional Debate

The case for tailoring educational content to the learner is well-established across several traditions. This section briefly surveys the research across the five dimensions most relevant to this framework, before addressing the professional debate about how far tailoring should go. The purpose is not to rehearse a literature review in full, but to establish the practical significance of the capabilities described in Sections 4 and 5, as well as reflecting that there *is* substantial debate about this among teachers and education researchers.

A substantial proportion of students in English schools are learning in their second or an additional language (DfE, 2024). Cummins' (1981) distinction between basic interpersonal communication skills (BICS) and cognitive academic language proficiency (CALP) is well-documented. In particular, students may appear fluent in social contexts while still experiencing significant difficulty with content-dense academic text. Linguistically adapted materials that reduce vocabulary complexity, shorten sentence length, and avoid idiom consistently support comprehension without reducing cognitive demand.

Reading level and cognitive load research provides a principled account of why these adjustments matter. Cognitive Load Theory (Sweller, 1988; Sweller et al., 2011) distinguishes between intrinsic load (the inherent complexity of the material) and extraneous load (the demand imposed by how it is presented). Simplified language reduces extraneous load, freeing working memory for the conceptual content rather than the decoding effort. Vocabulary breadth is a reliable predictor of comprehension, and sentence-length constraints are consistently effective at improving accessibility without lowering intellectual expectations. Research on readability measures in educational contexts suggests that even modest reductions in lexical complexity produce measurable gains in comprehension for lower-attaining readers (Flesch, 1948).

Cultural context and example grounding matter because examples are, by design, intended to connect abstract concepts to the learner's existing experience. What constitutes a relatable everyday scenario varies enormously across cultural contexts. For example, a physics example grounded in central heating will not land for a learner in a warm climate, and a mathematics example using sterling amounts is increasingly less universal. Gay's (2010) framework for culturally responsive teaching argues that the relevance of instructional examples to a student's lived experience is not merely motivational but epistemological, i.e. students learn better from content that connects with what they already know. More recent work on situational interest (Hidi & Renninger, 2006) supports this, finding that contextually relevant examples reliably elevate engagement and subsequent retention.

Neurodiversity research makes the case for structural as well as lexical adaptation. Universal Design for Learning (UDL) principles (CAST, 2018) advocate for multiple means of representation as a baseline design practice rather than a specialist accommodation. For learners with dyslexia, short sentences, active voice, and avoidance of dense paragraph structure are well-evidenced supports (British Dyslexia Association, 2018). For learners on the autistic spectrum, explicit structure, avoidance of ambiguous language, and clear signalling of what each part of a content item is doing can significantly reduce processing demand. A Year 8 SEND coordinator in one of our CPD workgroups noted that encoding these preferences as a named profile meant that the same accessibility adjustments could be applied consistently across every subject, rather than depending on each individual teacher's awareness and willingness to adapt.

Learner interest and motivational grounding represent perhaps the most immediately appealing application of audience tailoring, and the one that teachers respond to most enthusiastically in our professional development sessions. Grounding examples in a learner's area of interest engages them in a fundamentally different way than a generic scenario does. Hidi and Renninger's (2006) four-phase model of interest development identifies situational interest triggered by personally relevant material as a precursor to deeper individual interest, suggesting that interest-grounded examples may do more than increase momentary engagement. Critically, interest-based tailoring is most powerful at the individual level: a class of thirty students may share a reading level while being unlikely to share a preferred reference domain, and this is one of the motivating observations behind the more structured approach to audience description developed in Section 5.

3.1 How Far Should Tailoring Go?

The capability to tailor content to individual learners raises a question that practising teachers will recognise immediately: how much should you tailor? This is a professional and pedagogical question, not a technical one, and it is the subject of genuine and lively debate in staffrooms, CPD sessions, and education research.

The case for tailoring is long-standing. Learners come with different prior knowledge, different linguistic backgrounds, different processing styles, and different areas of interest, and content that fails to account for these differences is less effective than content that does. The Universal Design for Learning framework (CAST, 2018) articulates this as a design principle: building multiple means of representation, engagement, and expression into the content itself, rather than retrofitting for individual need after the fact. Differentiated instruction frameworks (Tomlinson, 2001) similarly position responsiveness to learner variance as a hallmark of effective teaching rather than an optional enhancement.

A counter-argument, with its own substantial research base, holds that excessive differentiation can be counterproductive. Ability-based tailoring can lower expectations for particular groups when it becomes conflated with deficit thinking, and it can consume teacher time that might be better spent on higher-impact interventions. The measured effect sizes for many individualisation strategies are also more modest than popular discourse suggests. A further consideration specific to LLM-based tailoring is the risk of bias within the AI systems themselves: LLMs trained on large corpora may encode and reproduce cultural, linguistic, or socioeconomic biases in ways that are not immediately visible in the output but that systematically shape the content generated for different audience groups. A variant profile specifying a particular cultural or linguistic background may interact with these biases in ways that a teacher reviewing the output in their own language would not readily detect, making the primary variant approach and the quality-control practices described in Section 4.3 all the more important.

Hattie's (2009) synthesis is instructive here. The effect size for individualised instruction is positive but not among the highest-ranked interventions, and whole-class instructional quality accounts for more of the variance in attainment than learner-by-learner adaptation. The implication is not that tailoring is valueless but that it is not a substitute for high-quality pedagogy. The framework described here is consistent with this view, since tailoring operates *on top of* a well-designed pedagogy rather than in place of one, and the pedagogy-first, audience-second architecture enforces this structurally.

We do not take a position on how much tailoring is appropriate; that is a professional judgment that belongs with teachers, departments, and schools. What the framework does is change the terms of the debate in a practically useful way. When tailoring decisions are encoded as explicit parameters, they become visible, reversible, testable, and deliberate. A department head can see exactly what audience assumptions are built into a content library, and removing an assumption is as simple as deleting a variable and regenerating. Matched versions at different tailoring levels can be produced and compared at scale, so that the decision to tailor or not becomes an active professional choice captured in the parameters rather than an invisible default shaped by habit or time pressure.

4 Audience as a Variable: The Simplest Form of Tailoring

The most immediate form of audience tailoring is also the simplest: adding a single audience variable to an existing pedagogy, much like in the example prompt we gave earlier, where the topic was added as a single variable. This requires no change to the pedagogical design at all. The teacher specifies who the content is for, and the LLM adjusts accordingly.

4.1 Setting the Audience Variable

At the simplest implementation, tailoring to an audience can be done with a simple variable entered at the end of the prompt. We give the example below of adapting to another language, which is a task that LLMs are consistently adept at.

> 📋 **COPY-PASTE PROMPT: Pedagogy A with a single audience variable**
>
> *Please generate a three-element educational content item using the pedagogy below, adapted to the audience given.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>>** *Density*
>
> **<<AUDIENCE>>** *Polish*

There is an important practical consideration when the audience variable specifies a different language. If a teacher generates content directly in Polish, or Urdu, or Somali, they have typically lost the ability to verify the output against their own subject knowledge. A secondary science teacher who speaks only English cannot judge whether a Polish-language definition of density is accurate, well-phrased, or pitched at the right level. This is a significant quality-control limitation, and one of the reasons why we advocate the primary variant approach developed later in this paper, in which the teacher generates the content first in a language they can review, verifies it against their professional judgment, and only then applies audience tailoring, including translation, as a second step.

We do not – and would never – claim that an LLM will consistently produce perfect translation of materials with total accuracy. However, the results are consistently impressive and lead us ultimately to the central philosophy of this paper. Using the example above, **a Polish-speaking student who receives a near-perfect translation of a carefully checked English original is substantially better served than one who receives no translated materials at all**. The fact that this can be done at scale, efficiently, and inexpensively with LLMs is, in our view, reason alone to consider audience tailoring as a scalable methodology. Our choice of the primary variant workflow ensures that the pedagogical quality of the original has been verified before any translation takes place.

The audience variable and nothing else is sufficient to shift the output substantially. Here are a few of the ways in which different audiences can be served using a single audience variable, as already discussed at length in Section 3.

- **Language.** Declaring a target language is sufficient to produce output in that language. No separate translation step is needed; the LLM generates directly in the target language, constrained by the same pedagogical parameters.
- **Reading level.** Specifying a younger or less experienced audience produces markedly simpler vocabulary and shorter sentences, while specifying an advanced audience produces appropriately demanding prose, all from the same pedagogy throughout.

- **EAL register.** Identifying learners at an early stage of English acquisition prompts the LLM to favour short sentences, high-frequency vocabulary, and avoidance of idiomatic expressions.
- **Neurodiverse learners.** A description such as ‘students with dyslexia; prefer short sentences, active voice, avoid abstract metaphors’ gives the LLM clear directions that meaningfully shape output without requiring specialist knowledge on the teacher’s part.
- **Cultural grounding.** Specifying a learner context shifts example choices away from UK-centric assumptions toward scenarios appropriate for that setting, without any change to the pedagogical parameters.
- **Learner interest.** Naming a learner’s area of interest directs the LLM toward scenarios in that domain, engaging the learner in a fundamentally different way than a generic scenario would.

It is important to acknowledge that we are not suggesting all accessibility or learning needs can be met by a single audience variable. It would be insensitive, for example, to suggest that neurodiversity as a blanket concept could be adequately addressed in this way. The examples above are deliberately simple, and the variant profile architecture described in Section 5 offers a considerably more nuanced and adaptable approach.

A reasonable question to ask at this point is whether the tailored output is good enough, and whether a teacher who cannot verify the suitability of the materials for every audience is acting responsibly in distributing them. We think the answer lies in recognising what the alternative is. The ability to tailor content in this way has never been possible before at any practical scale. A teacher who serves a class with eight different home languages has, until now, had no realistic way to provide linguistically adapted materials for each group. The choice was not between a perfect hand-crafted Polish version and an LLM-generated one; it was between an LLM-generated Polish version and no Polish version at all. A Polish-speaking student who receives a pedagogically structured content item in their home language, even one that is slightly imperfect in register or idiom, is in a fundamentally better position than one who receives nothing.

The framework does not claim to produce flawless tailoring, but rather to make tailoring tractable and the decisions involved visible and revisable. Where a teacher can verify the output, they should, and where they cannot, the structured nature of the generation still produces output that is considerably more purposeful than an unconstrained prompt. The primary variant workflow described in Section 4.3 ensures that the underlying pedagogical quality has been checked before any audience-specific version is committed. Any teacher reviewing the tailored output may naturally want to refine it further, and the fact that refinement can be done so quickly and inexpensively is part of what makes the approach practical. The variant profile architecture described in Section 5 develops this idea considerably further.

4.2 Batch Tailoring: One Pedagogy, Many Topics, Many Audiences

The batching strategies from Paper 1 extend naturally into audience tailoring. Since the audience is treated as a single variable, a teacher can straightforwardly create a prompt that iterates over multiple audiences. In the example below, the teacher includes all required audiences in the initial batch prompt. This produces the complete differentiated library in a single interaction and is the simpler option when the pedagogy is well-established and the audiences are fixed:

> 📋 **COPY-PASTE PROMPT: Workflow 1 – Batch with audiences included**
>
> *Please generate an example of the following pedagogy for the three given audiences.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*

> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>>** *Density*
>
> **<<AUDIENCE A>>** *Polish*
>
> **<<AUDIENCE B>>** *Students with interests in sports*
>
> **<<AUDIENCE C>>** *Students with dyslexia; short sentences, active voice, avoid abstract*

The LLM iterates the pedagogy across the topic and the three audience variables, producing a pedagogical content item for each. There is no reason why these audience variables cannot be more elaborate than the simple versions given here, and it is this observation that leads to the variant profile architecture developed in Section 5.

4.3 The Primary Variant: Generate, Review, Then Tailor

Section 4.1 introduced the concept of considering the audience as a variable in the initial prompt and generating tailored content in a single step. This works well when a teacher is producing content for one known audience, is confident in the pedagogy, and is happy to judge the tailored output directly. The greater the number of audiences that a teacher wishes to tailor to, the less likely they will have the expertise to do so. To provide a framework which counters this, as well as for many other reasons, we advocate a second workflow that introduces an explicit quality-control step, and it is this workflow that we call the ‘primary variant’ approach.

Workflow 1: Generate with audience (single step). As shown in the previous examples, the teacher includes the audience variable in the initial prompt and the LLM generates tailored content directly. This is the simpler option and works well for quick, single-audience generation where the teacher is confident enough in using the process and gauging the outputs.

Workflow 2: The ‘primary variant’ approach. The teacher generates the pedagogical content item using the pedagogy alone and the topic entered as a variable, with no audience variable. This untailored output is the primary variant, becoming the version the teacher reviews, judges, and if necessary adjusts before any audience tailoring takes place. It is worth noting that the primary variant is not entirely untailored in practice, since the LLM will naturally adjust language, tone, and content based on the topic provided, so a degree of implicit tailoring has already taken place, but in a way that should place the material firmly within the teacher’s familiar reach of expertise. Only once the primary variant has been approved does the teacher apply explicit audience tailoring in a follow-up step.

The following two prompts demonstrate the primary variant approach in practice. The first prompt generates the primary variant, which is an untailored pedagogical content item that the teacher reviews for pedagogical quality.

> 📋 **COPY-PASTE PROMPT: Step 1 – Generate base content**
>
> *Please generate a three-element educational content item using the pedagogy below.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>>** *Density*

The output of Step 1 is the primary variant, and the teacher reviews it before proceeding. If the content is satisfactory, the teacher moves to Step 2; if not, they adjust the pedagogy or topic and regenerate using natural-language commands. This checkpoint is what distinguishes the primary variant approach from single-step generation. In this method, the teacher's professional judgment is applied to the pedagogical quality of the content before any tailoring is committed, whether that tailoring is a single follow-up prompt or a batch operation across multiple audiences.

> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Step 2 – Apply audience tailoring**
>
> *Now please produce tailored versions of the content item above, one for each audience below. Present all three versions (the original and both variants) side by side.*
>
> **<<AUDIENCE A>>** *Polish*
>
> **<<AUDIENCE B>>** *Students with interests in sports*
>
> **<<AUDIENCE C>>** *Students with dyslexia; short sentences, active voice, avoid abstract*

The primary variant approach is the linchpin of quality control in any content-production-and-tailoring system that operates at scale. Without it, a pedagogical error in the base content is silently replicated across every audience-specific version, and the cost of catching that error after tailoring is proportional to the number of variants produced. With it, the error is caught once, at the primary variant stage, before any tailoring has taken place.

In our experience, the most effective strategy for the primary variant is to generate the most technically precise version first and then simplify in audience-specific variants, rather than the reverse. A primary variant written at a high level of accuracy gives the teacher a clear standard against which to judge whether a simplified version has lost something essential. A second useful strategy is to generate the primary variant in the teacher's own language, even when the eventual audience will receive the content in translation, for the verification reasons discussed in Section 4.1. Both strategies reflect the same underlying principle: the primary variant should be the version the teacher is best equipped to judge.

The following prompt demonstrates how the primary variant approach can be expressed as a single prompt. Rather than splitting the process across two separate prompts, the instruction asks the LLM to treat the generation as a two-step process within one conversation, presenting the primary variant for review before proceeding with audience tailoring.

> 📋 **COPY-PASTE PROMPT: Primary variant with audience tailoring (two-step, single prompt)**
>
> *Please generate a three-element educational content item using the pedagogy below.* ***Treat this as a two-step process. First, generate the primary variant with no audience tailoring and present it for my review. Then, once I confirm, produce tailored versions for each audience below.***
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>>** **Density**
>
> **<<AUDIENCE A>>** *Polish*
>
> **<<AUDIENCE B>>** *Students with interests in sports*
>
> **<<AUDIENCE C>>** *Students with dyslexia; short sentences, active voice, avoid abstract*

Popular LLMs are typically adept at following multi-step instructions, even if this is included within the statement of one complete prompt, as in the example above. Hence the primary variant approach can have the same user

experience for a teacher, requiring the same effort (a single prompt from a prompt gallery, or similar) to achieve a much wider range of impact.

Another reason for choosing the ‘primary variant’ approach is with the consideration in mind of applying these techniques to large-scale content production systems. We have found that the primary variant approach maps far more sustainably into such systems, than any single-step workflow can. When the primary variant exists as a discrete, approved artefact, it becomes a stable node in a content architecture. It can be versioned, updated when the curriculum changes, reprocessed when a new audience is added, and audited when quality is reviewed. Once a correction to the primary variant of any pedagogical content item has been approved, it can be propagated to every audience-specific version derived from it. The primary variant, viewed through this lens, reinforces our earlier points about the need for expert review in automated pipelines, ensuring that the teacher’s judgment remains embedded in the system even when the workflow is substantially automated.

The primary variant approach naturally allows the consideration and design of more exotic proliferation techniques that would be impossible in a single-step workflow. Because the primary variant is a stable, approved base, any update to a pedagogical parameter can be proliferated automatically through every audience-specific version in the library, and any new variant profile can be applied retrospectively to every existing primary variant without regenerating them.

To give a concrete example: a science department maintaining 200 primary variants and four variant profiles can regenerate all 800 tailored items in one pass by updating a single element parameter, and adding a fifth profile for a newly enrolled EAL cohort produces 200 additional items from the same approved base with no new pedagogical work. This kind of systematic responsiveness to changing needs is a direct consequence of the primary variant architecture, and it is one of the reasons we regard this workflow as essential rather than optional for any content system operating at scale. The manual two-step workflow described above reflects the same logic and gives the practitioner the same quality-assurance benefit at any scale.

5 Variant Profiles: Learner Descriptions as Reusable Natural Language Rules

A single audience variable, as described in Section 4, is powerful but coarse. It treats the learner as a single description applied uniformly across every element, leaving the specific form of the tailoring largely to the training data and architecture of the LLM itself. In practice, a teacher who knows their students well knows far more than a single parameter can capture. They know that one student reads confidently but processes abstract concepts slowly, that another needs culturally grounded examples but is perfectly comfortable with technical vocabulary, that a third engages more when examples come from sport. A single audience may capture some group-level characteristics reasonably well but it cannot capture this kind of individual variation with much precision.

A counter to this is the highly flexible idea of a ‘variant profile’, which extends the single-variable approach into a small, structured set of learner-facing parameters governed by natural language rules, written in the same grammar as the pedagogy itself and stored, shared, and reapplied in exactly the same way. Just as Paper 1 gave simple examples of pedagogies to demonstrate how the framework would clearly allow much more specific pedagogies, we will adopt the same approach here as we begin to demonstrate variant profiles as a technical idea. Namely, we will give simplistic versions of variant profiles as a means of demonstrating the detail to which teachers (or, in our estimate, more likely their students) may wish to take this idea. In this paper, the variant profile is a set of natural language rules about how the learner prefers to receive text content. In Paper 3, we extend the concept to visual and spatial parameters (layout, colour, image style, and typographic choices), encoded in the same grammar and applied in the same way. In Paper 4, the profile extends further to narration and audio-visual synchronisation.

5.1 What a Variant Profile Looks Like

In practice, variant profiles can range from a few simple parameters to highly sophisticated documents with many parameters addressing language, vocabulary, cognitive scaffolding, cultural context, tone, and more. The examples in this paper are deliberately simple for illustration, but the framework places no upper limit on the complexity of a profile.

To be clear, **a variant profile uses the same architecture of elements and parameters** that we defined for pedagogies in Paper 1. It could be applied alongside the pedagogy in a single prompt, as described above, but in our view it works best when applied after the primary variant has been generated and verified. Unlike the elements of a pedagogy, the parameters of a variant profile may be left blank where they are not relevant.

For the sake of consistency and comparability, the profiles throughout this paper use a common set of four parameters: Language/register, Vocabulary, Examples, and Tone. These are deliberately simplified for illustration and, in practice, teachers will certainly develop profiles with different or additional parameters suited to their own context, and the framework places no constraint on what those parameters may be. Blank entries signal that the profile has a principled reason for what it constrains and what it leaves open, rather than representing gaps or oversights.

Here are three example variant profiles. It is important to be clear that we do not present these as definitively effective profiles, but as principled starting points that illustrate the architecture. A teacher or specialist with domain expertise will be able to specify far more precise and nuanced parameters, and we expect that students themselves may prove even more effective at articulating what they need and how they prefer to receive information, a possibility we explore in Section 7.

Variant Profile 1: EAL Early Acquisition

Parameter	Rule
[Language/register]	Generate in English. Use short sentences (maximum 12 words). Avoid idioms, metaphors, and culturally specific references.
[Vocabulary]	Tier 1 and common Tier 2 words only. If a technical term is unavoidable, gloss it in brackets immediately.
[Examples]	Use physical, universal scenarios that do not assume familiarity with UK cultural context (shops, weather, transport, etc.).
[Tone]	—

Variant Profile 2: Dyslexia-Friendly

Parameter	Rule
[Language/register]	Short sentences. Active voice. One idea per sentence. No nested clauses.
[Vocabulary]	Prefer concrete, familiar words. Avoid multiple synonyms for the same concept.
[Examples]	—
[Tone]	—

Profile 2 constrains language and vocabulary only, without naming any specific element, which makes it portable – it can be applied alongside any pedagogy without modification, because nothing in its rules depends on a particular element being present. Profile 3 below takes the same approach, keeping its rules at the level of the content item as a whole rather than targeting individual elements. This portability is deliberate, since a profile that avoids naming specific elements can be applied to any pedagogy, present or future, without needing to be updated when the element structure changes.

Variant Profile 3: High Interest – Football

Parameter	Rule
[Language/register]	—
[Vocabulary]	Assume the learner knows standard football and sport terminology. Do not gloss terms like 'pressing', 'xG', or 'set piece'.
[Examples]	Ground examples in football or sport scenarios wherever this can be done without forcing an inappropriate analogy. Where the content includes a surprising or counterintuitive fact, connect it to sporting performance, athletic science, or sports history where possible.
[Tone]	Enthusiastic but accurate. Write as if the topic genuinely matters to athletic performance.

Each variant profile is a small, focused document of natural language rules whose parameters constrain how the content should be shaped for a particular learner, rather than specifying the content itself. Although a variant profile can in principle refer directly to named elements from the pedagogy it is applied alongside, we found this approach limiting in practice. A variant profile tethered to specific elements of a particular pedagogy is fundamentally less portable, and rules targeting the nature of what an example should contain tend to reflect learner preferences that extend beyond any single pedagogical structure. For this reason, the profiles in this paper keep their rules at the level of the content as a whole, describing how the learner receives information in general terms so that the same profile can be applied alongside any pedagogy, present or future, without modification.

Because a variant profile is just a small collection of natural language rules, a structured data source like a spreadsheet is a natural home for a library of them. A teacher can maintain a simple spreadsheet with one row per profile, with the profile name in the first column and one column per parameter. The profile is then ready to use in two ways. The teacher can read the relevant row and paste the rules directly into a prompt alongside the pedagogy and topic. Alternatively, for LLMs that accept file attachments, the entire spreadsheet can be uploaded alongside the prompt and the instruction simply references the particular variant profile. The spreadsheet does not need to contain any personal identifying information and hence could be held centrally for use by multiple teachers, once a variant profile is found to be effective for a particular student. A department can maintain one shared spreadsheet, version-control it like any other planning document, and distribute it to colleagues with a few clicks.

Within the framework, we treat the variant profile in exactly the same way as any other variable – as a named input declared at the bottom of the prompt, after the pedagogical parameters. The following prompt demonstrates the primary variant approach combined with variant profiles from an attached spreadsheet.

> 📋 **COPY-PASTE PROMPT: Primary variant with variant profiles (two-step, single prompt)**
>
> *Please generate a three-element educational content item using the pedagogy below. Treat this as a two-step process. First, generate the primary variant with no audience tailoring and present it for my review. Then, once I confirm, apply each variant profile from the attached spreadsheet to produce tailored versions.*

> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILES>>** 📎 ***[variant_profiles.xlsx attached]***

For most teachers, we expect that the initial natural infrastructure for deploying these prompts at scale will be a prompt library with agreed variant profiles. However, it is a very short technical leap from this to an agent-based solution, which can radically improve the user experience by centralising and standardising the information and approach used. A Custom GPT or a Copilot Studio agent can encode many approved pedagogies and a set of variant profiles directly in the system prompt, presenting the teacher with a simple work flow such as: enter a topic, choose a pedagogy, approve a primary variant, select a variant profile, and receive output.

5.2 Voice, Character, and Humour as Variant Parameters

A variant profile does not have to describe a learner's cognitive or linguistic characteristics. It can describe how the content sounds, covering voice, register, tone, and humour, and this is in our experience one of the most immediately engaging aspects of the framework for teachers who are new to it. Consider a teacher who has cultivated a particular classroom persona over many years: warm, slightly irreverent, prone to dry understatement. That persona is part of what makes their teaching effective, because students respond to a consistent, recognisable voice. Under the current model, that voice exists only in the classroom and in hand-written materials. With the framework, the teacher can encode their persona as a small set of natural language rules and apply it consistently across an entire content library. Using the common four-parameter structure, a voice profile might look like this:

Variant Profile: Enthusiastic Science Communicator

Parameter	Rule
[Language/register]	Informal but precise. Contract where natural (it's, don't, you'll). Address the reader directly as 'you'.
[Vocabulary]	—
[Examples]	—
[Tone]	Enthusiastic and curious. Write as if you find the topic genuinely surprising. Occasional dry understatement is welcome. Avoid forced jokes. All definitions must read like something you would tell a friend at a dinner table, not a textbook aside.

Profiles of this kind work particularly well in combination with the techniques that we discuss in Papers 3 and 4. When narration is included, the variant profile is heard as well as read, and a well-crafted voice profile applied consistently across a narrated video series produces a recognisable presenter character that learners come to associate with the content. In this context, the variant profile functions as a character brief.

6 Novel Workflows for Content Production and Curation

The preceding sections have described what the framework can produce. This section turns to how teachers and content teams work with it in practice, and in particular to the workflows that emerge naturally when pedagogies and variant profiles are used together at scale, underpinned by the generative power of LLMs. The primary variant approach described in Section 4.3 established a quality-control checkpoint at the heart of the process, and the workflows described here extend that principle into more sophisticated patterns of content creation, curation, and maintenance. These patterns are consistent with the broader direction of LLM practice, where prompt chaining and agentic workflows increasingly structure the interaction between human judgment and machine generation (Wu et al., 2022), and with the design science tradition that positions teaching as iterative, reflective design work (Laurillard, 2012; Goodyear, 2015).

In particular, we have found that batching workflows are an effective mode of working that is only possible with the capabilities of modern LLMs. Teachers are accustomed to balancing competing demands on their time, and content creation competes with every other aspect of the role. The workflows described here aim to maximise the application of a teacher's expertise while minimising the time that application requires.

With the correct infrastructure, whether batching prompts of the kind given in the previous sections or agent-based solutions such as Custom GPTs and Copilot agents, the production of candidate material becomes essentially effortless. The teacher's contribution shifts from generation to focused review, a mode of working that in our experience benefits from sustained, concentrated sessions rather than ad hoc interactions.

6.1 Non-Deterministic and Deterministic Tailoring

All of the tailoring described so far has been non-deterministic, in the sense that the LLM generates both the primary variant and the audience-appropriate content afresh in response to the variant profile. As discussed in Paper 1's Appendix B, this is a fundamental characteristic of LLMs, and the same prompt will produce different output on different runs. For most audience tailoring this is entirely acceptable, and in many cases actively desirable, since the teacher wants a suitable piece of content for a given audience and *any* high-quality output would serve.

For some tailoring operations, however, the non-deterministic nature of LLMs is far less desirable and, as we will argue shortly, a further reason to use the primary variant approach. If a teacher has a **[Definition]** of density that they have carefully written, verified against the exam board specification, and used successfully for years, they would reasonably not want the LLM to regenerate this in future work, especially if that definition is meant to act as part of a primary variant that will be tailored for different audiences. A definition that changes slightly each time it is generated is not a stable foundation for a multilingual content library.

This is the distinction between *non-deterministic tailoring*, in which the LLM generates audience-shaped content, and *deterministic tailoring*, in which fixed, human-verified content is substituted directly (an action which can be handled reliably by LLMs, if they are instructed to do so). Both have their place in a well-designed content production system and we explore this trade-off in considerably more detail in Paper 3, where the distinction becomes central to the visual and spatial layer of the framework, and then again in Paper 4, where we investigate various types of temporal sequencing.

6.2 The Generate-and-Select Workflow

Paper 1 observed that the framework shifts the teacher's role from creator to curator, and nowhere is this more explicit than in the workflow described here. The bridge between non-deterministic and deterministic tailoring is a workflow we call generate-and-select, which we have encountered in various forms across a range of LLM practitioners. In this workflow, a teacher uses the LLM to generate a gallery of candidate outputs for a particular element and then applies professional judgment to select the best candidate, lightly editing it if needed and promoting it to a fixed constant. The creation is handled by the LLM, while the curation belongs entirely to the expertise and judgement of the teacher.

This workflow is fast and effective at soliciting the focused expertise of the teacher, fitting into the category of ‘batching’ workflows we have referred to several times. For example, generating eight candidate definitions for a topic takes seconds using any LLM and, for an experienced teacher, reviewing them takes less time than writing one from scratch. The process draws on the teacher’s expertise at the point of judgment rather than the point of generation, and it produces content that combines the fluency of LLM generation with the accuracy of human, expert verification and the local context that a teacher will be able to apply. The non-determinism that is normally a limitation becomes an asset in this context, with each run producing a different candidate and giving the teacher a genuine range of options from which to choose.

The degree of variation between candidate outputs can also be influenced through prompt wording, so that instructions such as ‘produce six substantially different versions’ or ‘vary the approach but keep the register consistent’ give the teacher a measure of control over how much the candidates diverge. This is, in effect, a way of controlling the creative temperature of the system through natural language rather than technical settings, and it is an intuitive lever that teachers can easily pull.

> 📋 **COPY-PASTE PROMPT: Generate a definition gallery for selection**
>
> *I need a fixed* **[Definition]** *element for a content item on the topic below. Please generate six substantially different versions of this element, numbered clearly. Each must be no more than 20 words, using only short, common English words accessible to all learners. I will select the best to use as my fixed definition across all future pedagogical content items.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept above. Must use a tangible, physical scenario the learner can visualise. Maximum 40 words.*
>
> **[Surprising Fact]***: A surprising or counterintuitive fact that challenges or complicates the picture established by* **[Definition]** *and* **[Example]***. Must not simply add information; must make the learner reconsider. Maximum 35 words.*
>
> **<<TOPIC>> Density**

Once the teacher has selected and edited their preferred version, that text becomes a constant, inserted deterministically into every variant of the content item for that audience. A teacher may of course update it at any time, but a verified technical definition is a clear example of a pedagogical element that should be preserved deterministically, immune from the variation that an LLM would otherwise introduce.

6.3 Batch Workflows and Sustained Content Management

We have found in our own practice that managing content production systems of this kind leads naturally to batch workflows such as generate-and-select, in which content is created, tailored, updated, and reviewed in managed, sustained periods of work rather than as a series of isolated ad hoc interactions. Section 4.2 introduced the simplest form of this pattern, where a single pedagogy is applied across many topics and many audiences in one operation. The primary variant architecture described in Section 4.3 extends the same principle further, since an approved primary variant can serve as a stable base from which any number of audience-specific versions are derived. Once a teacher has a library of approved primary variants and a set of variant profiles, the natural unit of work is no longer the individual content item but the batch.

The proliferation properties alluded to in Section 4.3 are what make this sustainable over time. A pedagogy update, such as tightening the word count on a definition element or adjusting an inter-element rule, does not require the teacher to revisit every content item individually. The update is made once to the primary variant and, once verified, proliferated through the entire content library via established variant profiles. New variant profiles can be

applied retrospectively to existing primary variants without regenerating them, and existing content can be reprocessed with this term's updated profile library in a single operation. The result is a content system that is responsive to changing needs in a way that hand-authored materials could never be.

So far, we have implied that in such a content production system, that the majority of pedagogical content items would be generated non-deterministically, with key points such as the primary variant preserved and hence effectively treated deterministically. However, there is certainly utility in workflows where the majority of the content is deterministically preserved, with only a minority of elements generated with an LLM in a non-deterministic way. For example, a pedagogical content item could have many deterministically-fixed elements to explain a concept, with non-deterministic elements for tailored or contemporary examples. The possibilities are vast and we expect that we have explored relatively few of these.

7 Testing at Scale, and the Student as Co-Designer

7.1 Empirical Status of the Framework

It is important to be clear about what this paper claims and what it does not. The variant profiles presented in Section 5 are not presented as empirically validated interventions, nor as definitive specifications of what effective profiles should contain. These profiles are offered as principled starting points, being plausible illustrations of the architecture, informed by professional experience and the relevant literature. A teacher or specialist with direct knowledge of a particular learner group will certainly produce more precise and effective profiles, and we expect that students themselves, given the opportunity to articulate their own preferences within the framework's natural language grammar, may prove the most effective profile authors of all.

This is the honest position, and it is also the right position for a design-based framework at this stage. The appropriate claims here are claims about design itself. Specifically, we firmly assert that the variant profile provides a tractable structure for teachers in which both pedagogical decisions and audience tailoring decisions can be made explicit, stored, shared, and revised. The framework that we propose separates pedagogical decisions from audience decisions in a way that is practically useful, and it makes the generation of matched, controlled content variants tractable at a scale that hand-authoring never could. Those claims are well-supported by the practitioner experience that underpins this series.

There is, however, one area where the anecdotal evidence is strong enough to mention directly: language-to-language translation. Practitioners working with frontier LLMs report consistently that translation quality is impressively high, capturing nuance, respecting domain terminology, and producing output that reads naturally in the target language. This is consistent with the rapid improvement in multilingual LLM capability over the past two years.

Beyond translation, there are other good reasons to expect that the variant profile approach will generalise well even before formal validation. Reading level adaptation is among the better-understood levers in educational content design, and the relationship between vocabulary complexity, sentence length, and comprehension is well-established (Sweller, 1988). Cultural grounding of examples is consistent with the motivational literature on situational interest (Hidi & Renninger, 2006).

What the framework principally offers, empirically, is control and tractability. A teacher or a researcher who wants to know whether their EAL profile actually produces better comprehension than a standard version can now generate matched content, deploy both, and compare outcomes at a scale that would have been impossible with hand-authored materials. The parametric structure means every variable is controlled except the one under investigation. We are actively planning our first formal empirical work with education partners and welcome collaborators who wish to use the framework as an infrastructure for their own investigations.

7.2 Students as Profile Co-Designers

We want to close with an observation that has struck us repeatedly in our work with teachers, and that we believe will be fundamental in truly harnessing the power of LLMs within the education sector. The idea of students as active participants in shaping their own learning is well-established in the literature on student voice (Flutter & Rudduck, 2004; Cook-Sather, 2006), and the framework described here gives that idea a concrete, operational form.

The variant profile architecture operates entirely in natural language and hence there is nothing in that grammar that requires a teacher to write it. A student who can articulate how they learn can write their own variant profile. A student who says ‘I understand things better when there’s an example I can picture, and I find it easier when definitions are just one sentence’ has already written two parameters of their own profile. A student who adds ‘I like it when there’s something surprising at the end that makes me think’ has described a preference with equal eloquence to that of a teacher, and with a greater degree of self-understanding.

Under the right conditions, students could generate candidate profiles for themselves using the gallery approach from Section 6.2 and share the ones that work with their teachers. A teacher who asks a class of thirty to draft a brief profile in response to a structured prompt has acquired, in the time it takes to run that activity, thirty pieces of information about learner preference that would otherwise take weeks of observation to accumulate. Entered into a structured format (for example, a spreadsheet, as already discussed), a teacher then has the ability to tailor one primary variant of a pedagogical content item to all variant profiles and for no extra effort. The entire primary variant workflow and set of variant profiles can be included within a single, reusable prompt.

To make this practical, here is an activity that a teacher could use in the classroom to elicit student-authored profiles:

> 📋**Classroom Activity: Student self-authored variant profile activity**
>
> *Think about how you prefer to learn new things. Please answer the following questions in one or two sentences each:*
>
> *1. What kind of language do you find easiest to understand? (For example: short sentences, longer explanations, formal, informal.)*
>
> *2. When you learn a new word or idea, what helps you remember it? (For example: a simple definition, seeing it used in context, hearing it explained in everyday language.)*
>
> *3. What kinds of examples help you understand best? (For example: examples from sport, cooking, gaming, music, everyday life.)*
>
> *4. What tone of writing do you prefer? (For example: serious, friendly, funny, calm, enthusiastic.)*

Many learners, regardless of age, may find it difficult to answer these questions directly, since reflecting on the subtleties of one’s own learning is not a natural exercise for most students. To help with this, we have found that LLMs can be used to orchestrate a Socratic questioning exercise in which the student is presented with short examples of different styles and asked to respond to what they see, rather than describe their preferences in the abstract. The following prompt uses this approach.

> 📋 **COPY-PASTE PROMPT:** ***Socratic variant profile activity (LLM-guided)***
>
> *You are helping me as a* **<<Year X>>** *student to discover how I prefer to learn. Please guide me through a short ‘Socratic questioning’ conversation using the approach below on my favourite topics of* **<<Topics A, B, and C>>.** *For each question, please show two or three short examples from one of my preferred subjects to help gain an understanding of how I like to learn.*
>
> *Here is what I would like to understand about my own learning.*
>
> *Step 1 – Language and register.*

> *Step 2 – Vocabulary and definitions.*
>
> *Step 3 – Examples.*
>
> *Step 4 – Tone.*
>
> *Please ask ten questions and, at the end and present your findings in a table with these four rows. Please keep the output fairly short and neat, and formatted in a way that my teacher would thank me for.*

The responses map directly to the four-parameter structure used throughout this paper which, we reiterate, is arbitrary and for the sake of demonstration. However, a teacher who collects these responses has, in effect, a first draft of a variant profile for every student in the class. Those profiles will not be perfect but the raw material is there, expressed in the student's own natural language expressions, and the teacher's professional judgment can shape it into a working profile that reflects genuine learner input.

The framework provides the natural language grammar in which both can be expressed, combined, and tested. Whether students can define what they need as effectively as teachers can define what they believe students need is one of the more interesting questions this framework opens up, and we hope that researchers will explore it. We regard the development of variant profiles in tandem with students as likely to be a particularly fruitful area of discourse, one that sits at the intersection of learner voice, differentiation research, and the practical capabilities that the framework now makes available.

8 Conclusion

This paper has extended the element-and-parameter architecture of Paper 1 into the dimension of audience tailoring. The primary variant approach (generating a base content item first, reviewing it for pedagogical quality, and only then applying audience tailoring) keeps pedagogical and audience decisions cleanly separated, allows quality checks before tailoring is committed, and maps directly to the architecture of a production system capable of applying multiple profiles to a single approved base in one operation. Because the primary variant is a stable, versioned pedagogical content item, updates to the pedagogy or the addition of new variant profiles can be proliferated through an entire content library without regenerating the approved base. This architectural property makes the primary variant approach essential rather than optional for any system operating responsively and at scale.

We then introduced the variant profile architecture, which allows learner descriptions to be captured as small, reusable sets of natural language rules that sit alongside the pedagogy without touching it. Using a common parameter structure across all profiles makes profiles comparable, combinable, and straightforward to maintain. We discussed briefly how non-determinism interacts with such a system of tailoring, and described some of the exciting workflows that interact well with these approaches to creating and managing pedagogical content.

We would like to reiterate that we believe that the structured prompts this series produces are well-suited to institutional deployment. A pedagogy and its associated variant profiles can be encoded in a Custom GPT system prompt or a Copilot Studio agent, presenting teachers with a simple interface that encapsulates the full framework without requiring them to manage the prompt text directly. At a lighter-touch level, prompt dictionaries and shared prompt galleries provide a natural home for the named pedagogies and profiles developed using this approach.

We welcome contact from researchers who wish to use the variant profile architecture as infrastructure for empirical studies of differentiation, and from practitioners who wish to build on the framework in their own contexts. We will soon be releasing an open-source pedagogical content production system which uses all of the techniques developed in this series of papers, as well as many more, and we welcome collaborators.

References

Banerjee, D., Suresh, T., Ugare, S., Misailovic, S., & Singh, G. (2025). CRANE: Reasoning with constrained LLM generation. arXiv preprint arXiv:2502.09061.

Barab, S., & Squire, K. (2004). Design-based research: Putting a stake in the ground. The Journal of the Learning Sciences, 13(1), 1–14. https://doi.org/10.1207/s15327809jls1301_1

British Dyslexia Association. (2018). Dyslexia style guide. https://www.bdadyslexia.org.uk

CAST. (2018). Universal design for learning guidelines version 2.2. http://udlguidelines.cast.org

Cook-Sather, A. (2006). Sound, presence, and power: “Student voice” in educational research and reform. Curriculum Inquiry, 36(4), 359–390. https://doi.org/10.1111/j.1467-873X.2006.00363.x

Cummins, J. (1981). Bilingualism and minority-language children. Ontario Institute for Studies in Education.

Department for Education. (2024). Schools, pupils and their characteristics: Academic year 2023/24. https://explore-education-statistics.service.gov.uk

Flesch, R. (1948). A new readability yardstick. Journal of Applied Psychology, 32(3), 221–233. https://doi.org/10.1037/h0057532

Flutter, J., & Rudduck, J. (2004). Consulting pupils: What’s in it for schools? Routledge. https://doi.org/10.4324/9780203464380

Gay, G. (2010). Culturally responsive teaching: Theory, research, and practice (2nd ed.). Teachers College Press.

Hattie, J. (2009). Visible learning: A synthesis of over 800 meta-analyses relating to achievement. Routledge. https://doi.org/10.4324/9780203887332

Hidi, S., & Renninger, K. A. (2006). The four-phase model of interest development. Educational Psychologist, 41(2), 111–127. https://doi.org/10.1207/s15326985ep4102_4

Qian, Y. (2025). Prompt engineering in education: A systematic review. Journal of Educational Computing Research. https://doi.org/10.1177/07356331251365189

Sahoo, P., Singh, A. K., Saha, S., Jain, V., Mondal, S., & Chadha, A. (2024). A systematic survey of prompt engineering in LLMs. arXiv preprint arXiv:2402.07927.

Schön, D. A. (1983). The reflective practitioner: How professionals think in action. Basic Books.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4

Sweller, J., Ayres, P., & Kalyuga, S. (2011). Cognitive load theory. Springer. https://doi.org/10.1007/978-1-4419-8126-4

Tomlinson, C. A. (2001). How to differentiate instruction in mixed-ability classrooms (2nd ed.). Association for Supervision and Curriculum Development.

Appendix A: Further Variant Profile Examples

The following profiles are offered as worked examples that practitioners can adapt. Each uses the common four-parameter structure (Language/register, Vocabulary, Examples, Tone) introduced in Section 5.1. Blank entries indicate parameters not relevant to the particular profile's purpose, by design rather than oversight. All profiles can be combined with any pedagogy from this series and layered with one another where appropriate.

Variant Profile 4: Advanced Extension

Parameter	Rule
[Language/register]	Formal and precise. Technical vocabulary is expected, not avoided.
[Vocabulary]	Introduce precise technical vocabulary where it exists. Do not simplify at the expense of accuracy. **[Definition]** should include a formal statement where one exists (e.g. a mathematical expression or scientific law).
[Examples]	—
[Tone]	**[Surprising Fact]** should introduce a current research frontier, an unresolved question, or a counterexample that complicates the standard account.

Variant Profile 5: Adult Returner (ESOL)

Parameter	Rule
[Language/register]	English, accessible to ESOL learners at Entry Level 3. Short sentences, active voice, Tier 1 vocabulary.
[Vocabulary]	—
[Examples]	Adult-relevant scenarios: household, workplace, or community contexts. Avoid school-based scenarios.
[Tone]	Treat the learner as an adult with life experience. Do not condescend or use child-directed language.

Variant Profile 6: Dry Wit – Secondary Science

A voice profile, not a learner-need profile. This profile describes how the content sounds rather than who it is for, and can be combined freely with any audience profile.

Parameter	Rule
[Language/register]	Informal but never sloppy. Contractions fine. No exclamation marks. No 'amazing!' or 'fascinating!'.
[Vocabulary]	—
[Examples]	—
[Tone]	Confident, dry, occasionally self-aware. One moment of gentle

	understatement per content item, maximum. Never a punchline. Never at the learner's expense. **[Surprising Fact]** delivered as if the writer only just remembered it, and is slightly annoyed it was not mentioned sooner.

Voice profiles become increasingly deterministic as a content series develops. A creator who runs the generate-and-select workflow across ten topics will identify the specific phrases, transitions, and cadences that best represent their voice and fix those as constants. The non-deterministic LLM layer then operates within the space defined by those fixed anchors, producing new content that is recognisably in the same voice without requiring the creator to write every word from scratch.