Natural Language as the Bridge Between Pedagogy and AI

Paper 3 of 4: Visual and Spatial Design – Extending Variant Profiles to Markup and Rendering

Dr Joe Watkins

Grassroots Intelligence Limited, Canterbury, UK

Correspondence: joe@grassroots-intelligence.com

Abstract

Papers 1 and 2 established a framework for encoding pedagogical expertise as discrete elements with natural language parameters, and extended it to audience tailoring through variant profiles and the primary variant approach. This paper demonstrates that the same architecture, without any structural modification, accommodates visual and spatial design decisions as further parameters of the variant profile. These visual parameters are not interpreted by the LLM but consumed directly by a markup renderer, making them overwhelmingly deterministic. Once a variant profile has been designed and agreed, it can be applied to every pedagogical content item in a library with no extra use of an LLM, with identical results on every run and no need for further review. The paper opens with accessibility, where the research literature on dyslexia, the Web Content Accessibility Guidelines (WCAG), and related standards provides specific, evidence-based visual recommendations that map directly onto variant profile parameters. It then extends the same approach to general visual design and cultural adaptation, includes a worked example of translating a content item into other languages, and addresses image generation as the non-deterministic exception where the generate-and-select workflow from Paper 2 remains essential. Paper 4 applies the same architecture to the temporal dimension, confirming that the framework requires no further structural extension beyond what Papers 1 and 2 established.

Keywords: visual design, markup rendering, accessibility, WCAG, dyslexia, deterministic generation, variant profiles, HTML, image generation, educational content design.

1 Introduction and Recap

By this paper, the architecture we have developed across the series is essentially complete. Papers 1 and 2 established the core constructs (elements, parameters, pedagogies, variables, pedagogical content items, primary variants, and variant profiles), along with a family of workflows including generate-and-select, batch tailoring, and the two-step primary variant approach. This paper introduces no new constructs at all. Instead, it shows that the same architecture comfortably accommodates a quite different rendering target, simply by treating visual and spatial design decisions as further parameters of the variant profile. The same will be true of Paper 4, where the rendering target is temporal rather than spatial. From this point in the series the framework is being applied and interpreted, and not revised in any sense.

The one genuinely new observation is that the parameters of a variant profile now fall into two kinds. Some are non-deterministic: the text-level rules of Paper 2 that an LLM interprets afresh on every run. Others are deterministic: rules about how the content should look when rendered, applied directly by a rendering engine and identical every time. We develop this distinction in the recap below and put it to work from Section 2 onward.

This paper follows the same tradition of design-based research (Barab & Squire, 2004) as its predecessors. The remainder of this section recaps the terminology and the running example used throughout. Section 2 addresses accessibility, the clearest case because its parameters are almost entirely deterministic. Section 3 treats visual parameters more broadly. Section 4 turns from the look of a content item to its language, with worked examples of translation. Section 5 covers rendering into markup and the content generation engines that let the framework work at scale. Section 6 turns to variant profiles and student dialogue. Section 7 addresses image generation as the non-deterministic exception. Section 8 covers scope and limitations.

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

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. Paper 2 extended this concept significantly.
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. Treated as an optional variable.
Pedagogical content	Any populated pedagogy, where a topic or other input variables have been entered. We also refer to a single instance informally as a pedagogical content item.
Primary variant	The initial, audience-neutral pedagogical content item generated from a pedagogy and topic before any audience tailoring is applied. The primary variant serves as the quality-controlled base from which all audience-specific versions are derived. See Paper 2.
Variant profile	Introduced in Paper 2. A structured bundle of parameters applied alongside a pedagogy to tailor its output for a particular learner or group. The parameters need not be natural-language rules, as some are numeric or fixed values. In this paper they are often deterministic, applied directly by a rendering engine, or non-deterministic, interpreted by an LLM.

No new terms are introduced here. The extension this paper describes is a matter of what the existing constructs can already express, rather than anything that needs to be added to them. To ground the discussion we use the same two ingredients throughout, and in the order the primary variant approach of Paper 2 implies: first the pedagogy, which produces the content, and then the variant profile, which tailors it. The pedagogy is the central three-element example carried through the whole series.

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

A variant profile is the second ingredient, applied to the content the pedagogy produces. By way of reminder, here are two example variant profiles that could be used once a primary variant has been generated, with each parameter being fundamentally non-deterministic, as shown. We stress again that these variant profiles are arbitrary, and that any teacher, or language specialist, might prefer a quite different set of parameters and rules. We use the Polish learner deliberately, because the same learner returns in Section 4, where the content is given a full, fluent Polish localisation.

Variant Profile: Polish EAL (Early Acquisition of English)

Parameter	Rule	Type
[Language/register]	Generate in Polish, for a student in the early stages of English acquisition. Use short sentences (maximum 12 words).	**Non-deterministic**
[Vocabulary]	Tier 1 and common Tier 2 words only.	**Non-deterministic**
[Examples]	Use physical, universal scenarios that do not assume familiarity with UK cultural context.	**Non-deterministic**
[Tone]	(left open)	**Non-deterministic**

The above variant profile is intended to take a primary variant and tailor it fundamentally towards another language, in this case with other specific needs of the student expressed as well. The exact same set of parameters could be given quite different rules, however, and so achieve an effect intended instead to suit a student who has a high level of interest in a particular subject. In the example below we show how we could apply the parameters of this variable profile to a student with a significant interest in football. Our experiences is that such variant profiles typically work very well.

Variant Profile: High Interest – Football

Parameter	Rule	Type
[Language/register]	(left open)	**Non-deterministic**
[Vocabulary]	Assume the learner knows standard football and sport terminology. Do not gloss terms like 'pressing', 'xG', or 'set piece'.	**Non-deterministic**
[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.	**Non-deterministic**
[Tone]	Enthusiastic but accurate. Write as if the topic genuinely matters to athletic performance.	**Non-deterministic**

What this paper adds to the existing architecture is the observation that the variant profile naturally accommodates two fundamentally different kinds of parameter, consumed by two different parts of the system. The first kind, which Paper 2 covered extensively, consists of rules that shape the LLM's generation of text. These are non-deterministic, in that the LLM interprets them and the output varies on each run. The second kind, introduced here, consists of rules that describe how the generated content should look when rendered. These are typically deterministic, with a rendering engine applying them and producing an identical output every time.

Both kinds sit side by side in the same variant profile, are applied in the same operation, and can be stored and shared as a single artefact. The text rules are interpreted by the LLM. The visual rules are interpreted by the rendering engine, which the LLM can call directly as one of the many tools that frontier models are provided with as standard. The practical consequence is that the most expensive, unpredictable, and review-intensive component of the system plays no part in the visual layer at all. This is a further decoupling, and so a further means of sustainably managing and updating content libraries. The visual elements of a variant profile, once designed, can be reused indefinitely, across every pedagogical content item in a library, every topic, and every subject.

2 Accessibility as the Lead Case

In our view, the strongest and most evidence-based case for treating visual elements as variant profile parameters is accessibility. The research literature on visual presentation and readability is strong, particularly in the area of visual accessibility and dyslexia, where controlled experiments have produced specific, quantitative recommendations. Many of these recommendations can be expressed as a deterministic parameter in a variant profile, applied automatically by a rendering engine, and reused across an entire content library. The practical need here is a real one, since most schools simply lack the resources to produce consistently accessible materials across a whole curriculum. Universal Design for Learning treats such accessible representation as a baseline design principle rather than a specialist accommodation (CAST, 2018), and the variant profile gives that principle a concrete, reusable form.

We do not claim that the variant profiles presented here constitute clinical accessibility specifications. They are principled starting points, informed by the research literature and by the practical recommendations of bodies such as the British Dyslexia Association and the Royal National Institute of Blind People (RNIB), and we would encourage any school adopting them to refine them in consultation with their own SEND specialists. As with the text-based variant profiles in Paper 2, the claim is about the efficacy and potential of the architecture rather than the specific content of any individual profile.

Every worked example in this paper follows the primary variant workflow of Paper 2, and we present each one in the single-prompt form Paper 2 settled on. One prompt runs the whole two-step process, in that it asks the model first to generate the primary variant with no tailoring and present it for review, and then, once the teacher confirms, to apply a variant profile. The teacher therefore sees and approves the underlying content before any styling or tailoring is committed, but never has to manage more than one prompt. We use the Polish EAL profile from Section 1 as the first illustration.

> 📋 **COPY-PASTE PROMPT: Polish EAL learner (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 styling or tailoring and present it for my review. Then, once I confirm, apply the variant profile below and render the result as a complete, self-contained HTML page.*
>
> **[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]**. *Maximum 35 words. Must not simply add information; must make the learner reconsider.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILE>>** 📎 ***[Polish EAL (Early Acquisition of English) profile attached]***

A teacher who has used Paper 2 will find this entirely familiar. From here on every example takes the same shape, and only the profile changes. We begin with WCAG, the clearest case, because every one of its parameters is deterministic. Later in this paper we mix in some non-deterministic parameters as well, applying a variant profile that carries both. We should note that we consider it entirely believable that teachers would prefer to manage this as a three-stage process, in which the primary variant is generated first, the non-deterministic aspects of a variant profile are applied second, and the deterministic, visual elements are applied last, in a third transformation. This is not our preferred approach, as explained in detail in Section 5.

2.1 WCAG and Related Standards

The Web Content Accessibility Guidelines (WCAG), published by the World Wide Web Consortium (W3C, 2018; W3C, 2023), are the primary international standard for web accessibility. They also provide measurable, testable criteria that map naturally to variant profile parameters. WCAG defines three conformance levels: Level A (minimum accessibility), Level AA (the standard most commonly referenced in legislation), and Level AAA (the highest conformance level).

Colour contrast provides a clear illustration. WCAG 2.1 specifies a minimum contrast ratio of 4.5:1 between text and its background for Level AA and 7:1 for Level AAA. These can all be expressed as a single rule in a variant profile, which is applied deterministically by the rendering engine. Several related standards are also worth noting:

- EN 301 549 (CEN, CENELEC, & ETSI, 2021) is the European accessibility standard for ICT, referencing WCAG 2.1.
- Section 508 of the US Rehabilitation Act (United States Access Board, 2017) serves a similar function in the United States.
- RNIB Clear Print Guidelines (RNIB, n.d.) provide practical recommendations for partially sighted readers.
- ISO 14289 (PDF/UA) (International Organization for Standardization, 2014) governs PDF accessibility.
- The W3C Cognitive Accessibility Guidance (COGA; W3C, 2021) addresses cognitive and learning disabilities more broadly.

Our framework gives teachers and institutions a way to express compliance with any of these standards as a reusable, shareable variant profile, written once and then reused across every pedagogical content item in a library. For such standards there is no need to consider the non-deterministic parameters of a variant profile, since the text of the primary variant is not altered at all. This changes slightly later in the paper, as we consider possible variant profiles for improving accessibility for dyslexia and for other languages, where we use a mixture of deterministic and non-deterministic transformations.

Both accessibility profiles we use immediately below are purely deterministic, since every parameter is applied by the rendering engine and the content of the written text is left exactly as the primary variant wrote it. We define a WCAG AA profile and a stricter WCAG AAA high-contrast profile, and apply the first in the two-step prompt below. A teacher who has to satisfy a particular standard can keep the matching variant profile in their spreadsheet and reach for it by name, without holding the underlying criteria in mind each time.

Variant Profile A: WCAG AA

Parameter	**Rule**	**Type**
[Contrast]	4.5:1 for body text; 3:1 for text 18pt and larger	**Deterministic**
[Font]	Sans-serif, 16pt	**Deterministic**
[Background]	White (#FFFFFF)	**Deterministic**
[Layout]	Single column, maximum width 700px	**Deterministic**
[Element styling]	1px solid #767676 borders; 16px padding per element	**Deterministic**
[Structure]	Semantic HTML headings (h1 to h6); ARIA landmarks	**Deterministic**
[Colour use]	Never signalled by colour alone; always paired with text or a label	**Deterministic**

Variant Profile B: WCAG AAA High Contrast

Parameter	**Rule**	**Type**
[Contrast]	7:1 for body text; 4.5:1 for text 18pt and larger; #000000 on #FFFFFF	**Deterministic**
[Font]	Verdana, 18pt; bold headings	**Deterministic**
[Background]	White (#FFFFFF)	**Deterministic**

[Layout]	Single column, maximum width 550px	**Deterministic**
[Element styling]	3px solid #000000 borders; 20px padding per element	**Deterministic**
[Structure]	Semantic HTML headings (h1 to h6); ARIA landmarks	**Deterministic**
[Colour use]	Never signalled by colour alone; always paired with text or a label	**Deterministic**

Maintaining the paradigm from Paper 2, we can then write an example prompt that produces a primary variant on a given topic and then casts it deterministically into a different visual form. The prompt specifically requests the output as an HTML page, which is easily and quickly achievable for all frontier LLMs which contain these rendering engines. We have chosen to attach the variant profile in a deliberately abstract and open form, since teachers are likely to have different preferences for expressing these. An Excel document in which the named variant profiles are rows and the parameters are columns would be a perfectly suitable way of storing this information, and is easily tractable for an LLM.

> 📋 **COPY-PASTE PROMPT: WCAG AA (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 styling or tailoring and present it for my review. Then, once I confirm, apply the variant profile below and render the result as a complete, self-contained HTML page.*
>
> **[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]***. Maximum 35 words. Must not simply add information; must make the learner reconsider.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILE>>** 📎 ***[WCAG AA profile attached]***

2.2 Rendering Deterministic Parameters into HTML

It is worth pausing on what just happened, because it is the mechanism the rest of the paper relies on. A deterministic parameter, such as a font, a colour, or a line height, calls for no interpretation or judgement. Instead, such parameters have a single value which can be simply transcribed into the markup, which in our case is an HTML page with internal CSS. In the copy-paste workflow above it is the LLM that does this transcription as it writes the HTML, but it is making no creative choice, and a faithful model reproduces the same value on every run by employing an internal HTML rendering engine. In our experience this functions to an extremely high, repeatable standard. In a dedicated pedagogical content engine, of the kind described in Section 5, the same job is done by a rendering layer with no LLM involved at all. It is also worth noting that HTML is a dynamic medium rather than a static one, so the same rendering target can carry interactive elements, for instance a schema for multiple-choice questions or an interactive demonstration the learner can manipulate, each specified once and rendered as reliably as any static element.

One caveat follows directly from this. In the copy-paste workflow above, a teacher who wants a whole library to look consistent should reuse the same LLM conversation thread when rendering new items, rather than starting afresh each time and letting the model improvise small differences. We must be clear that the initial HTML output from one thread is likely to be slightly different from any other thread given the non-deterministic nature of an LLM, but once this output is set the rendering will be entirely consistent. We present a number of other steps to control this aspect later in the paper (such as providing an LLM with an example HTML output as part of the input source material, as well as discussing the matter in a wider sense in Section 5.

2.3 Dyslexia and Typography Research

Around 10% of the population has dyslexia (British Dyslexia Association, 2018), and the visual presentation of text has a well-documented effect on readability for these learners. Rello and Baeza-Yates (2013) conducted the first large-scale eye-tracking study of font type and dyslexia, testing 12 fonts with 48 dyslexic readers, and found that sans-serif, monospaced, and roman font styles significantly improved reading performance over serif, proportional, and italic alternatives. Fonts including Helvetica, Courier, Arial, and Verdana performed best, and the study recommended against italic and serif fonts for dyslexic readers.

Font size matters independently of font type, and Rello, Pielot, Marcos, and Carlini (2013) found that 18-point text was optimal for dyslexic readers on screen in the context of actual web reading. Zorzi et al. (2012), in a study published in the Proceedings of the National Academy of Sciences, demonstrated that extra-large inter-letter spacing significantly improved both reading speed and accuracy in dyslexic children. Beyond typography, Wilkins (2003) demonstrated that tinted backgrounds reduce visual stress for a significant proportion of readers, and the British Dyslexia Association (2018) reflects this in its practical recommendation to avoid pure white backgrounds in favour of cream or pastel alternatives.

Each of these typography findings is a deterministic parameter, applied by a rendering engine exactly as in the WCAG case. A dyslexia profile, though, is usually not purely deterministic, since alongside the typography it tends to want the text restructured in specific ways, and that is a non-deterministic, text-level request the LLM must interpret. We therefore define the profile with its deterministic parameters first and its non-deterministic ones last. We reiterate that we do not claim this variant profile would be ideal for every – or indeed any - student with dyslexia. It is, again, a principled starting point, drawn from our own admittedly naïve reading of the literature, and we would expect it to be refined in dialogue with the students it serves and with the SEND specialists and teachers who know them best. Since each non-deterministic parameter is simply a natural-language rule, its effect can be put to the controlled-variation testing we set out in Papers 1 and 2, so a profile of this kind is not only adjustable but empirically checkable.

Variant Profile C: Dyslexia-Friendly

Parameter	Rule	Type
[Font]	Arial, 16pt; no italics	**Deterministic**
[Spacing]	Line height 1.8; a little extra letter spacing	**Deterministic**
[Background]	Pale cream (#FFF8E7)	**Deterministic**
[Contrast]	Dark grey (#333333) on cream	**Deterministic**
[Element styling]	Soft rounded boxes; element name bold above	**Deterministic**
[Language]	Short sentences; active voice; one idea per sentence	**Non-deterministic**
[Vocabulary]	Concrete, familiar words; avoid needless synonyms	**Non-deterministic**

The ordering of the Dyslexia-Friendly profile is deliberate. Its first five parameters are deterministic: the font, spacing, background, contrast, and element styling are applied verbatim by the rendering engine, and they carry the weight of the typography research described above. The last two, language and vocabulary, are non-deterministic: they are instructions the LLM interprets as it generates the text, and they sit at the end precisely because they are a different kind of transformation. We include them because short sentences and familiar vocabulary are plausibly helpful, and because modern LLMs handle such requests reliably, but we want to be clear that we offer them only as starting points. We are not aware of any study establishing how an LLM should rewrite text specifically for dyslexic readers, and that particular combination of evidence does not, as far as we know, yet exist. The deterministic typography parameters rest on real experiments, while the non-deterministic text parameters rest on professional judgement, and the latter should be treated as a sensible default to be refined rather than a validated intervention. As discussed in Paper 2, these variant profiles could serve as a starting point for discussions with students as to their preferences, and honed to a point where they can be reliably re-used thereafter.

In the second step of the prompt below, this profile therefore does two jobs at once, in that the LLM rewrites the text for the non-deterministic parameters and renders the result for the deterministic ones, all from the same variant profile. We note, however, that the prompt is essentially identical to the one used above for high-contrast accessibility, only with a different variant profile attached, which illustrates the reusability of the primary variant and variant profile approach.

> 📋 **COPY-PASTE PROMPT: Dyslexia-friendly (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 styling or tailoring and present it for my review. Then, once I confirm, apply the variant profile below and render the result as a complete, self-contained HTML page.*
>
> **[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]**. *Maximum 35 words. Must not simply add information; must make the learner reconsider.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILE>>** 📎 ***[Dyslexia-Friendly profile attached]***

Here the LLM does two different jobs at once. In one side it interprets the text rules, which it does afresh on each run, and on the other side it carries the visual rules, which a renderer then applies identically. The dyslexia-friendly rendering appears as profile C on the right of Figure 2, where the simplified wording is plainly visible.

2.4 Comparing Variant Profiles Side by Side

One of the dimensions introduced in Paper 1 was controlled variation, where all parameters are held constant except one (or, in the case below, several). This applies directly to variant profiles, and is most easily seen by rendering the same primary variant through different profiles and placing the results side by side. We label the standard classroom rendering profile A, the WCAG AAA high-contrast rendering profile B, and the dyslexia-friendly rendering profile C, and we take them in that order.

In profile A we deliberately render the element labels in the same purple used for **[Definition]**, **[Example]**, and **[Surprising Fact].** Then, Profiles B and C restyle those labels along with everything else, and the contrast is part of what the eye picks up when the three are set together. Figure 1 places profile A beside profile B. Both are purely deterministic, so the text is word for word identical and only the look changes. Indeed, we see that profile B switches to a large sans-serif font, pure black on white, and heavy borders for maximum contrast.

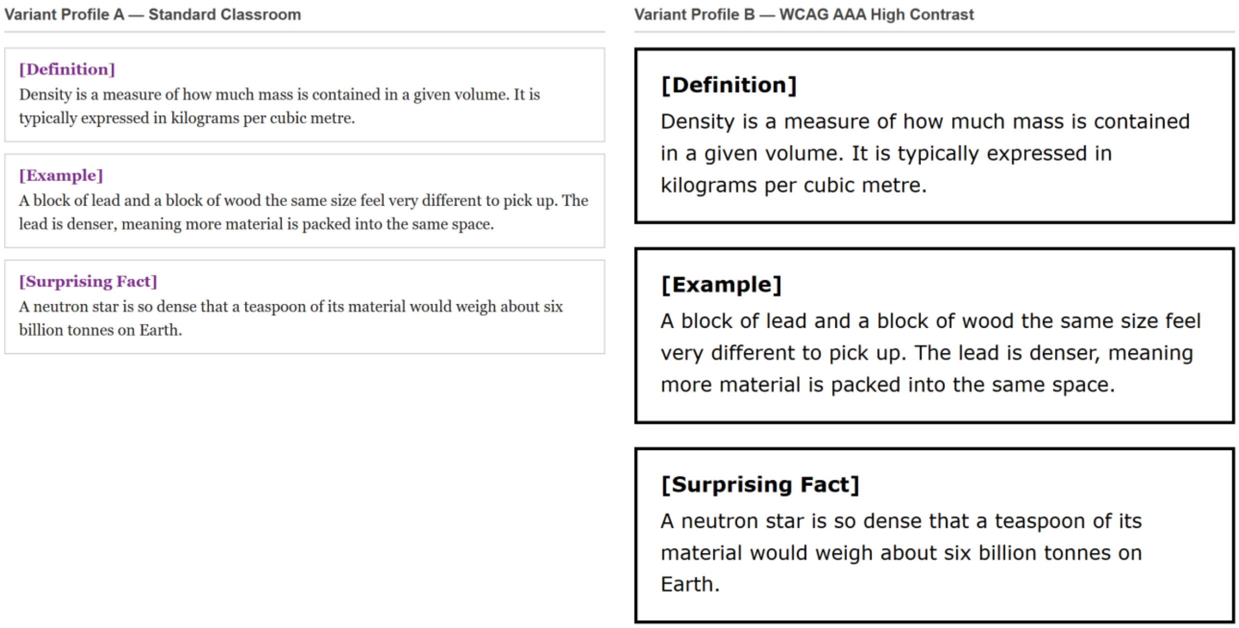

Figure 1. *The same density primary variant under profile A (standard classroom, left) and variant profile B (WCAG AAA high contrast, right).*

Continuing, Figure 2 below places profile A beside profile C, the dyslexia-friendly profile. Here the difference is not only visual. Because profile C also carries non-deterministic text parameters, the LLM has simplified the wording as well, breaking the sentences up and choosing more familiar words, so the right-hand column genuinely reads differently from the left rather than merely looking different.

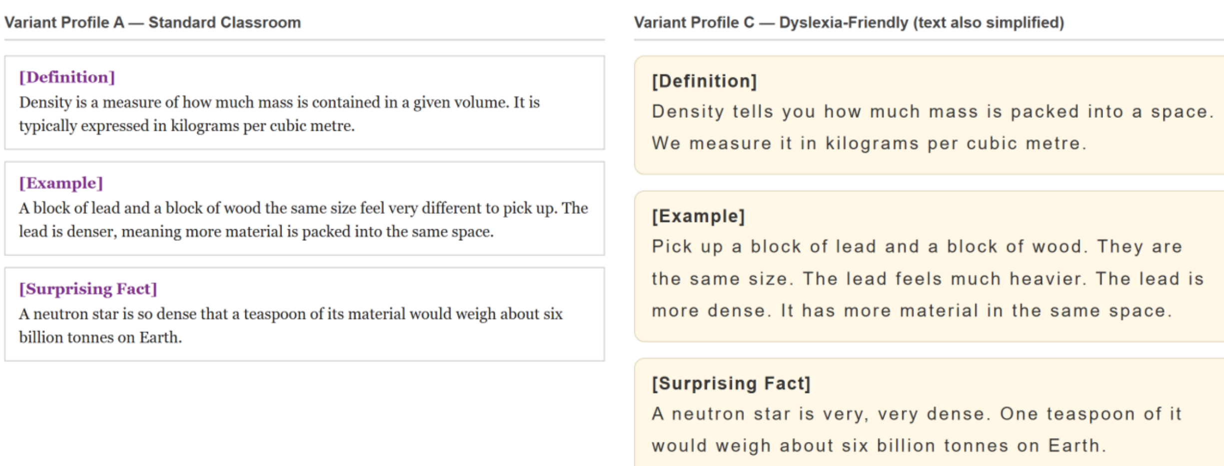

Figure 2. *The same primary variant and then again under variant profile C (dyslexia-friendly, right). Profile C applies both the deterministic typography and the non-deterministic text rules.*

Taken together, the two figures separate the framework's two kinds of parameter cleanly. The move from profile A to profile B is entirely deterministic and free to apply across a whole library. The move from profile A to profile C adds two non-deterministic elements, which transform the written text.

Each example so far has named its profile in the prompt. When a teacher has many profiles to apply at once, they need not be named one by one, and can instead be drawn from a shared spreadsheet, exactly as Paper 2 showed. The two-step structure is unchanged, and the attached-spreadsheet convention is that of Paper 2 exactly, as shown below.

> 📋 **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 the variant profiles below and render the results as complete, separate, self-contained HTML pages.*
>
> **[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]**. *Maximum 35 words. Must not simply add information; must make the learner reconsider.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILES>>** 📎 ***[variant_profiles.xlsx attached]***

A teacher who keeps a spreadsheet of agreed profiles can run this single prompt whenever they need a new content item, review the primary variant once, and receive the whole tailored set in return. Alternatively, they can save themselves the need to copy and paste the prompt, by encoding the instruction above in the equivalent form of a CustomGPT or Copilot Agent, or any similar device offered by a frontier LLM.

3 Visual Parameters Beyond Accessibility

3.1 Typography, Layout, and Institutional Consistency

Accessibility is the most evidence-backed case, but any visual design decision that a teacher, department, or institution wants to apply consistently can be expressed as a variant profile parameter. Font family, size, weight, line spacing, margins, text alignment, background colour, and accent colours are all deterministic and all expressible simply. Many markup and CSS engines also offer the option to animate text and images, but this adds an extra degree of complication that we reserve for Paper 4, where we treat temporal elements in the form of animation, narration and video.

The simplicity of this framework will, we hope, make it attractive for institutional consistency. A department that wants everything it produces to share one visual identity can agree a single house-style profile, deterministic in every parameter, and apply it across its whole library. If the branding were to change, the simple deterministic parameters of the house style change with it and everything is re-rendered deterministically. Markup languages like HTML also offer far more sophisticated styling than the handful of elements we have introduced here, so a very high level of visual finesse can be agreed once and then applied throughout. This can also include temporal elements such as animation or narration, but this is covered more fully in Paper 4.

3.2 Cultural and Contextual Variant Profiles

Paper 2 discussed cultural context as a dimension of audience tailoring. The visual layer extends this to colour, typography, and spatial design. Reinecke and Bernstein (2011) demonstrated that adapting visual design elements to cultural preferences significantly improved user performance, perceived usability, and aesthetic ratings, with improvements of up to 27% in user satisfaction across five cultural regions. Although this research was conducted in web interface contexts rather than education specifically, Gay's (2010) argument that learners engage more effectively with culturally connected material applies to visual presentation as naturally as it does to textual examples. Research on situational interest similarly suggests that contextually resonant presentation can heighten engagement and support later retention (Hidi & Renninger, 2006).

A much stronger caveat is needed here than anywhere else in this paper, and it is worth stating in the plainest possible terms. The accessibility profiles of Section 2 rest on a substantial evidence base, namely that the dyslexia and WCAG literature gives specific, tested values for font, size, spacing, and contrast, and a SEND specialist or teacher who follows them is on firm ground, for the deterministic elements at least. There is no comparable basis for mapping a particular colour scheme or typographic style to a particular language or culture. The choices in our examples below and the colours we use later for other languages … are entirely arbitrary and purely illustrative. We are emphatically not suggesting that they are correct, expected, or in any way meaningful for those audiences, and a reader should not infer that any one culture maps onto any palette and corresponding variant profile. Where accessibility tells us what to do, cultural visual styling at most tells us what we can do. Any real deployment should be shaped by the people it is meant to serve, drawing on their knowledge rather than on an author's or a model's assumptions. This is one of the clearest points in the framework where the responsible move is to treat the parameters as a question for the community rather than a default set from outside. We reiterate once more that we believe that such variant profiles should be agreed in conjunction with the student(s) that they are meant to serve.

To make this concrete, we define two cultural variant profiles below, one tuned to a UK setting and one to a Polish one. Both are purely deterministic in that they differ only in colour scheme, accent colour, and header styling, and the text is left untouched.

Variant Profile: UK

Parameter	Rule	Type
[Font]	Arial, 13pt	**Deterministic**
[Colour scheme]	Navy (#1B2A4A), white, light-blue accent (#D6EAF8)	**Deterministic**
[Background]	White (#FFFFFF)	**Deterministic**
[Element styling]	Clean bordered boxes with navy headers	**Deterministic**

Variant Profile: Poland

Parameter	Rule	Type
[Font]	Arial, 13pt	**Deterministic**
[Colour scheme]	Crimson (#D4213D), white, soft red accent (#FBE0E4)	**Deterministic**
[Background]	White (#FFFFFF)	**Deterministic**
[Element styling]	Clean bordered boxes with crimson headers	**Deterministic**

We then essentially repeat exactly the same prompt as already used within this paper, with the two variant profiles attached as a single file, presumably as an Excel document or other tabular structure of the kind discussed above.

> 📋 **COPY-PASTE PROMPT: Cultural 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 styling or tailoring and present it for my review. Then, once I confirm, apply the variant profiles below and render the results as complete, separate, self-contained HTML pages.*
>
> **[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]**. *Maximum 35 words. Must not simply add information; must make the learner reconsider.*
>
> **<<TOPIC>>** *Density*
>
> **<<VARIANT PROFILES>>** 📎 ***[UK and Poland variant profiles attached]***

The result is shown in Figure 3. The text content is identical in both columns and the pedagogical structure is unchanged, and only the deterministic visual parameters differ, applied by the HTML rendering engine.

Figure 3. *The same density content item rendered with two cultural variant profiles. Text is identical in both columns.*

We are, at this point, still leaving the language entirely untouched. Converting the words themselves is a different, non-deterministic operation and is the subject of Section 4.

4 Localising the Language, Not Just the Look

The variant profiles of Section 3 change how a pedagogical content item looks while leaving its words untouched, although we did touch upon how non-deterministic transformations might be used in tandem to transform the written content in addition to the visual styling. Genuine localization for other cultures or languages will clearly need to go further than this, not only directly translating the language but also localising any terms or examples used in the primary variant. Translation is a text-level, non-deterministic operation, and so it belongs to the tailoring aspect of Paper 2 rather than to the visual aspect that has been the main subject of this paper.

As Paper 2 described, in our framework any linguistic translation could follow the primary variant / variant profile workflow. In this workflow the teacher generates and verifies a primary variant first, checking its pedagogical quality in a language they can read, and only then translates it, so that the pedagogy is sound before a single word moves beyond the teacher's reach.

It helps to see the localisation itself as a variant profile, of exactly the kind introduced in Section 2. Like the Dyslexia-Friendly profile in Section 2.3, it carries both deterministic and non-deterministic parameters, in the same order. There, the non-deterministic parameters had the LLM rewrite the wording into shorter sentences and more familiar vocabulary, the text transformation we actioned alongside the typography. Translation is the same kind of non-deterministic, text-level operation, only taken further, in that instead of simplifying the English the model now converts it into another language entirely. The deterministic parameters still fix the look, with a font that renders the target script cleanly, plus the background and the element styling, all applied as part of the same operation. The non-deterministic ones, the translation and the register, are handled by the LLM. The example below is a candidate Polish localisation profile of this kind, and we apply it in the worked example that follows.

Variant Profile: Polish Localisation

Parameter	Rule	Type
[Font]	Arial, 14pt (renders Polish diacritics cleanly)	**Deterministic**
[Background]	Pale cream (#FFF8E7)	**Deterministic**
[Element styling]	Soft bordered boxes; element name bold above	**Deterministic**
[Language]	Translate the approved primary variant into natural, fluent Polish; preserve every quantity and unit	**Non-deterministic**
[Register]	Pitch at the same age group; avoid idiom that will not carry across	**Non-deterministic**

4.1 A worked example: Polish

The prompt below applies the Polish Localisation variant profile to the approved English primary variant. Please note that, from this point onward, our examples provide the approved primary variant as an attachment alongside the variant profile, in whatever format suits (in our case, pre-rendered and approved HTML). These attachments are simply the 'source material' mechanism of Paper 1 reused once more. In other words, these are trusted reference content supplied alongside the pedagogy for the model to draw on rather than invent, already in the form of the desired rendered output. In our experience this much more closely mirrors how the primary variant and variant profile framework is used at scale, where transformations are driven increasingly by simple prompt wording, and the substance is kept in the attached primary variant and variant profile. This is most certainly the case in dedicated content generation engines, as described in Section 5. The abstraction at this stage is deliberate, and it will remain throughout the rest of the paper.

> 📋 **COPY-PASTE PROMPT: Localise into Polish (two-step, single prompt)**
>
> *Below is an approved English primary variant of a three-element content item. Please apply the variant profile below, rendering the result as a complete, self-contained HTML page.*
>
> **<<PRIMARY VARIANT>>** 📎 ***[Primary variant attached]***
>
> **<<VARIANT PROFILE>>** 📎 ***[Polish Localisation profile attached]***

Since the teacher approved the primary variant first, the pedagogy is already known to be sound. What an English-speaking teacher cannot fully verify is the Polish itself. This is exactly the limitation Paper 2 raised, and the framework's answer is the same. The translation is a single, inspectable step layered on a checked original, easily reviewed by a Polish-speaking colleague, teaching assistant, or family member, and far better for the learner than no Polish version at all. However, in our anecdotal experience, language-to-language translation as undertaken by an LLM is typically excellent, and easily usable to a high-level in most applications. Certainly, it functions better for a recipient student than having no translated version at all.

4.2 A worked example: Hindi

Now working in Hindi sharpens the same point, because it uses a different script and so sits even further beyond the reach of an English-speaking teacher's own checking. The example is hence an exact parallel of the Polish one we gave above. We define a Hindi localisation variant profile that carries the same deterministic and non-deterministic parameters, in the same order, with appropriate adjustment and we then apply it to the same approved primary variant. The abstraction we are now enforcing around the prompt wording, together with our reiterated view that variant profiles are principled starting points, is a crucial one for making the leap into the more sustainable pedagogical content engines described in Section 5.

Variant Profile: Hindi Localisation

Parameter	Rule	Type
[Font]	Nirmala UI, 14pt (renders Devanagari cleanly)	**Deterministic**
[Background]	Pale cream (#FFF8E7)	**Deterministic**
[Element styling]	Soft bordered boxes; element name bold above	**Deterministic**
[Language]	Translate the approved primary variant into natural, fluent Hindi; preserve every quantity and unit	**Non-deterministic**
[Register]	Pitch at the same age group; avoid idiom that will not carry across	**Non-deterministic**

Using the more-abstracted and sustainable version of the prompting method that we have generated, we then see how readily such a variant profile can be applied.

> 📋 **COPY-PASTE PROMPT: Localise into Hindi (two-step, single prompt)**
>
> *Below is an approved English primary variant of a three-element content item. Please apply the variant profile below, rendering the result as a complete, self-contained HTML page.*
>
> **<<PRIMARY VARIANT>>** 📎 ***[Primary variant attached]***
>
> **<<VARIANT PROFILE>>** 📎 ***[Hindi Localisation profile attached]***

Reproduced side by side, the three versions show how much, and how little, changes between languages. The element structure, the named quantities, and the pedagogical arc are constant, and only the language and visual elements are converted. Indeed, such is our faith in the quality of the translated outputs by the LLM that we have felt no need to have them checked for accuracy by an expert.

Figure 4. *The same pedagogical content item translated into Polish and Hindi, rendered as styled cards. Here it is the language that changes, not the pedagogy. The panel colours are arbitrary, chosen only to tell the three versions apart, and no cultural meaning should be read into them (see the caveat in Section 3.2). The Polish and Hindi shown here are LLM-generated drafts and would require a fluent speaker's check before use.*

There are two cautions with an example like this. First, the output is non-deterministic, so a teacher who runs the prompt again will receive slightly different wording each time, which is why a chosen translation, once verified, is often best fixed once it has been established. This then allows even more elaborate possibilities which mix variant profiles. For example, a teacher could pair a translated primary variant with a variant profile, producing, say, a Hindi content item rendered with the Dyslexia-Friendly profile, or a Polish one in the school's house style. We will, however, leave this giddying possibility to future discussion and to be experimented with by any interested readers.

5 Pedagogical Content Generation Engines

The examples so far have all used copy-paste prompts. That is, in our view, a helpful way to learn the framework and to prototype a pedagogy or a variant profile, and we have leaned on it throughout precisely because it needs no technical infrastructure beyond perhaps a prompt dictionary and a sensible storage mechanism for approved primary variants and agreed variant profiles. It is not, however, how a department would build and maintain an archive of hundreds or thousands of content items, especially if these incorporate many variant profiles. Doing that by pasting prompts is, in our experience, a brutal grind and very batch-heavy. At any serious scale the framework needs to be embodied in pre-coded infrastructure. At its weakest that might be a Custom GPT or a Copilot agent holding the pedagogies and profiles in its system prompt and being called at will. More sustainably, it is a purpose-built content generation engine. This section covers both the lightweight rendering an LLM can do directly, useful for authoring and prototyping, and the shape of such an engine, which is what scale actually requires.

5.1 Frontier LLMs as Direct Rendering Engines

As already discussed, many frontier LLMs can now render markup directly. Claude's artifact system, ChatGPT's canvas feature, and equivalent capabilities in other platforms mean that a teacher can design a variant profile with visual parameters and see the rendered output in the same conversation. In this workflow, a teacher defines a pedagogy, enters a topic and some number of variant profiles, and then asks the LLM to generate the pedagogical content items as HTML. If the visual treatment needs adjustment, the teacher modifies a parameter and regenerates. The design-test-refine cycle that Laurillard (2012) and Goodyear (2015) describe as central to teaching as a design science happens within a single conversational interaction and at little cost to the teacher's time, while extracting their valuable expertise.

We want to be cautious about overstating how reliably LLMs currently produce well-formed, visually polished markup, since results vary across models and prompts, and more complex layouts sometimes require a degree of iteration.

5.2 Moving Towards Dynamic Content Generation Engines

In the copy-paste workflow it is the LLM that calls a renderer to author the HTML markup. Hence a teacher who wants a whole library to look the same has to reuse the same thread, or the same saved prompt, when rendering new items, and trust the model not to improvise small differences from one item to the next (or a Custom GPT or Copilot agent, as discussed). These approaches, while generally suitable for ensuring consistency in terms of the visual style, are brittle in that they cannot easily be updated in response to any variant profile. Furthermore, the pedagogies themselves cannot be updated on an element-by-element basis without regenerating significant amounts of content.

Consider a teacher who has built up two hundred approved primary variants over several years. Re-rendering them, even if expedited by the methods above, would still take time to action and oversee, no matter how efficient the prompting, Custom GPTs, Copilot agents, or the supporting storage and file structure for primary variants and agreed variant profiles. However, if we wish to update on a per-element or per-parameter basis at any point, then this is clearly not suitable or sustainable. Furthermore, if we wish to undertake any of the empirical quests laid out in Paper 1 that are afforded by our framework, we clearly are limited in terms of scale by the human effort that is required.

The natural conclusion is to decouple from the idea of using manually-inputted prompts and instead to reimagine our approach as one that fits most naturally to our framework in full. Namely, we will imagine the 'element' and 'parameter' structure as being implemented in the most atomic way, as individual elements of an HTML rendering engine that can be adjusted, if needed, with individual API calls to a suitable LLM. By rendering in this way the deterministic elements such as colour and font choice are held separate, say in CSS definitions, with non-deterministic transformations made as and when needed. This follows much more of a software engineering approach, in particular the CI / CD pipeline aspect that is prevalent in modern paradigms. This is a liberating step to take, allowing many new methods for updating pedagogical content items. For example, an element can have one of its parameters adjusted and then proliferated to all other elements in the pedagogical content items which have been created using the variant profiles.

To reiterate, the deterministic parameters of a variant profile are compiled directly into the page's markup, where the style variables become CSS and the layout and element styling are templated, so applying a profile to a page is a rendering operation with no model involved at all. The non-deterministic parameters are the only things that reach an LLM, and they do so through specific, well-defined API calls: for example, one to generate or regenerate a single element, one to produce the variants of a page, one to translate a whole page in a single pass. Because that markup is HTML, which is inherently dynamic rather than static, the same engine can render interactive content directly, for instance a schema for multiple-choice questions or an interactive demonstration the learner can manipulate, handled as a further templated element alongside the visual ones.

This separation is what makes an archive maintainable. Since the deterministic layer is just markup, changing a profile, tightening a colour or swapping a font, re-renders the entire library instantly and at no model cost. Since the non-deterministic layer is reached only through named calls to the LLM, each generation can be reviewed once and then stored, rather than being repeated every time a page is opened.

There is a spectrum of ways to embody the framework, and a school can sit anywhere along it. A shared prompt library is the lightest. A Custom GPT or a Copilot agent carrying a few approved pedagogies and profiles in its system prompt is a modest step up, and is often enough for an individual teacher. A purpose-built engine of the kind described here is the most capable, and is what we would recommend for any institution producing content at real scale. An individual teacher might never need more than a prompt library and a Custom GPT, whereas a teacher who finds themselves regenerating the same materials every term is the one who will feel the pull towards a fully fledged engine for producing and updating pedagogical content items.

None of this changes the underlying framework we have presented so far in these three papers. The copy-paste prompts in this paper and its predecessors remain the right tool for learning, prototyping, and small batches, and an engine is simply the same architecture made durable. The reliability gains are considerable. The expensive, non-deterministic, review-heavy work is confined to a small number of explicit calls. The bulk of what a learner actually sees, the layout, the typography, the accessible colour scheme, is deterministic, free to apply, and identical on every view.

6 Variant Profiles and Student Dialogue

Paper 2 introduced the idea that students themselves can author variant profiles, and that the natural language grammar of the framework makes this accessible to any learner who can articulate how they prefer to receive information. The visual layer extends this dialogue in a way that we believe is particularly powerful, because visual preferences are often more immediate and more personal than linguistic ones. A student who might struggle to articulate whether they prefer "short sentences with active voice" can often say with confidence that they find it easier to read light text on a dark background, or that they prefer content with colour and clear visual emphasis rather than plain black and white. Visual preferences of this kind are tangible, visible, and easily testable. They open up a conversation about learning that is grounded in what the student can see and respond to, rather than what they can describe in the abstract.

This points to a route we think is particularly valuable, and that mirrors the student-as-co-designer argument of Papers 1 and 2. A SEND coordinator or specialist teacher is often better placed than a subject teacher to understand a particular learner's access needs, and is used to working with students who may not easily advocate for themselves. Working one to one, such a specialist can sit with a student, try a few concrete visual treatments, and co-design a variant profile that genuinely suits them, capturing in a handful of parameters what might otherwise take a class teacher months to notice. This variant profile then becomes a reusable artefact that travels with the learner.

It is also worth noting again that variant profiles can be made extraordinarily specific. The profiles presented in this paper use broad, principled categories such as "dyslexia-friendly" or "WCAG AA compliant" because these are the most evidence-based and the most widely applicable. However, the architecture that we have presented places no constraint on the level of specificity a profile can reach. Take a student who has discovered, through experimentation, that they read best with black Verdana 15pt on a pale green background (#E8F5E9), with 1.6 line spacing and generous paragraph spacing. All of that information can be encoded in a single variant profile and applied to every pedagogical content item they receive. The student has, in effect, designed their own visual learning environment, and the teacher has gained a concrete, reusable specification of that student's visual preferences that travels with the learner across subjects and across years. As a perturbation of this workflow, a SEND specialist or teacher could prepare a range of candidate variant profiles, from which a student could pick the most appropriate.

The Socratic questioning approach from Paper 2 adapts naturally to the deterministic visual parameters, although it can easily be extended to the non-deterministic parameters as discussed in that paper. Rather than asking a student to describe their visual preferences in the abstract, an LLM can present the same pedagogical content item rendered with two or three different visual treatments and ask the student which they prefer and why, building up a variant profile through a series of concrete comparisons. The following prompt demonstrates this approach.

> **COPY-PASTE PROMPT: Socratic variant profile activity (LLM-guided)**
>
> *You are helping a student discover their visual learning preferences. Use the following pedagogical content item on the topic of density as your base text. At each step, render the same text with two different visual treatments as HTML, ask the student which they prefer and why, and use their answer to refine the next comparison. Cover the following in four to five steps:*

> *Step 1. Font and size. Show the same content in a serif font (Georgia, 12pt) and a sans-serif font (Arial, 14pt). Ask which feels easier to read.*
>
> *Step 2. Background colour. Show the preferred font on a white background and on a pale cream (#FFF8E7) background. Ask which is more comfortable.*
>
> *Step 3. Spacing. Show the preferred combination with standard line spacing (1.4) and generous spacing (1.8). Ask which feels less crowded.*
>
> *Step 4. Colour. Show the preferred combination with and without coloured accents on the element headers. Ask which feels more engaging.*
>
> *At the end, present the student's variant profile as a table with their preferred values for Font, Spacing, Background, Contrast, and Element styling, formatted so their teacher can use it directly.*

The result of this activity is a variant profile authored by the student, grounded in their own preferences as expressed through concrete visual comparisons, and ready to be combined with the text-based variant profile from Paper 2's equivalent activity. The combined profile, covering both linguistic and visual preferences, can then be applied to every pedagogical content item the student receives, across every subject, with no further work beyond the initial design. We regard this kind of student-led visual customisation as one of the more promising applications of the framework, and one where the deterministic nature of visual parameters is a particular advantage, since the student can see immediately and consistently what their choices look like in practice.

7 Image and Diagram Generation

The visual parameters discussed so far are overwhelmingly deterministic. Image generation is the exception, sitting firmly at the non-deterministic end of the spectrum described in Paper 2. Current LLMs produce different images from the same prompt on every run, and the quality and accuracy of results vary in ways that matter for educational content. A scientific diagram that misrepresents a concept, or a set of images that are stylistically inconsistent across a library, can undermine the pedagogical value of the material.

Generating several candidate images for every element in a library is possible, but it is computationally heavy and rarely necessary. We recommend a lighter approach, which we call the 'gallery' method. In this workflow the teacher generates one image for every item in the library in a single pass, hence creating a gallery either for the primary variant or with a following variant profile. The teacher then reviews the gallery as a whole and regenerates only those few, keeping the teacher firmly in the role of curator. Their expertise is spent at the point of judgement, deciding which images pass and which need another attempt, rather than at the point of generation.

The first prompt below generates the gallery. One illustration is produced for each topic in the library, all sharing the same style constraints, and the set is returned together for review.

> 📋 **COPY-PASTE PROMPT: Generate an image gallery for a topic library**
>
> *Please generate one illustration for each topic listed below, all adhering to the shared style constraints, and present them together as a single gallery I can review at a glance.*
>
> **<<TOPICS>>** *Density, Forces, Pressure, Energy Transfer, Waves*
>
> ***STYLE CONSTRAINTS:***
>
> > - *Flat, clean illustration style suitable for secondary school science*
> > - *Colour palette: deep blue (#1B3A5C), warm orange (#E8913A), light grey (#F0F0F0), white (#FFFFFF)*
> > - *No text or labels within the image*
> > - *Aspect ratio: 16:9*

Where some images miss the mark, the teacher simply regenerates just those with the normal natural language follow-up prompts and referring to the images which were successfully generated to the standard of the teacher. This is the point at which it is worth asking for a small spread of options to choose from.

> **COPY-PASTE FOLLOW-UP PROMPT: Regenerate the few that missed**
>
> *Those mostly work well. Please regenerate only the items below, keeping the same style constraints, and give me three distinct options for each so I can choose.*
>
> **<<REGENERATE>>** *Forces, Waves*

Once the teacher is satisfied, each approved image becomes a deterministic asset in the content library and is usually not regenerated on subsequent passes. It is attached to its primary variant and reused identically across every variant profile and every rendering pass. Since the image is tied to the primary variant rather than regenerated from a variant profile, it does not change when a profile changes, in just the way Paper 2 promoted a verified definition to a fixed element.

The one real weakness of generating each image independently is stylistic consistency. A gallery assembled from separate generations can drift in palette, line weight, and overall look from one image to the next. Some image generation tools are beginning to close this gap. The image generation service Recraft, for example, offers a style_id mechanism which allows a high level of visual control and consistency. In this, a user defines a visual style once, from some number of reference images, and every later generation reuses it, so that images produced for quite different topics still share the same style. This does not remove the non-determinism, since the content of each image still varies from run to run, but it does help to drastically pin down the style, which is the part that most needs to stay constant across a library. Combined with the gallery method, a style_id lets a teacher generate a stylistically coherent gallery for a whole topic library in one pass, review it in a single sustained session, regenerate the few that miss, and fix the approved images as deterministic constants.

We do not dwell on this aspect overly in this paper, partially due to the sheer computational expense of producing images using LLMs, as well as the inherent lack of reliability. We also do not believe our own methods of controlling this variation are yet robust enough to publicly offer any further advice than that given in this section.

8 Scope and Limitations

Several limitations should be acknowledged, some familiar from Papers 1 and 2 and some specific to this paper.

First, the visual examples use HTML as the rendering target. While HTML is widely supported, it may not align with every school's existing infrastructure. The visual parameters are format-agnostic natural language rules and apply to any rendering target, but the specific implementation will naturally differ.

Second, the variant profiles presented in Section 2 onwards are principled starting points informed by the research literature, not clinical specifications. A teacher or school adopting them should treat them as a starting point rather than a finished answer, and develop their own variant profiles in consultation with SEND or EAL specialists and, wherever possible, with the learners themselves.

Third, the cultural visual adaptation discussed in Section 3.2 draws on the UX and interface design literature rather than on education-specific research. Whether culturally adapted visual design improves educational outcomes is an open question that the framework makes testable.

Fourth, this paper does not treat animated or interactive content in depth. Animations can be specified as variant profile parameters and rendered in HTML. They can also be disabled for learners with photosensitive or attentional needs. A comprehensive treatment of temporal elements is deferred to Paper 4.

9 Conclusion

The architecture established in Papers 1 and 2 accommodates visual and spatial design without any structural modification. Visual design decisions are further parameters of the variant profile, expressed in the same grammar of labelled parameters as every other part of the framework, whether those parameters are natural-language rules or fixed values such as a font size or a colour. The critical difference is that visual parameters are consumed by a rendering engine. That engine may be a tool the LLM calls directly, or a separate component managed alongside it. Either way the parameters are deterministic, consistent, and, once designed, free to apply across every pedagogical content item in a library.

The strongest case is accessibility, where the dyslexia typography research (Rello & Baeza-Yates, 2013; Zorzi et al., 2012), WCAG (W3C, 2018), and practical guidance from the British Dyslexia Association and the RNIB provide evidence-based recommendations that translate directly into variant profile parameters. A school that agrees on a dyslexia-friendly variant profile has, in a single design decision, implemented the recommendations of a substantial body of research, doing so consistently, at scale, and without adding to the cost of any individual item.

Between these two poles sits the bulk of the paper. General visual design and institutional house styles, cultural and contextual adaptation, the localisation of language into Polish and Hindi, the co-design of variant profiles with students themselves, and the move from copy-paste prompts into dedicated content generation engines … all turn out to be the same move in different clothing. Each is a bundle of parameters, some applied deterministically by a renderer and some interpreted by an LLM, attached to an approved primary variant and reused without further cost. The practical payoff is constant throughout. A school agrees a profile once, informed by evidence where evidence exists and by the people it serves where it does not, and then applies it across an entire library, in every subject and for every learner, at no additional generation cost and with identical results on every view. It is worth emphasising, too, that HTML is a dynamic medium rather than a static one, so the same deterministic rendering opens the door to genuinely interactive content, such as schemas for multiple-choice questions and interactive demonstrations a learner can manipulate, which we expect to develop in future work.

Image generation sits at the non-deterministic end of the visual layer, where the generate-and-select workflow from Paper 2 remains essential. A teacher who has worked through Paper 2 already has everything they need to apply these variant profiles, and nothing here asks them to learn a new way of working, only to add a second kind of parameter to a structure they already know. Paper 4 applies the same architecture to the temporal dimension and it will not introduce new constructs. The framework is complete, and what remains is to demonstrate its application across further rendering targets such as animation, and educational video with (or without) narration.

We welcome contact from researchers and practitioners who wish to use the framework in their own contexts.

References

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

CEN, CENELEC, & ETSI. (2021). *EN 301 549 V3.2.1: Accessibility requirements for ICT products and services*. European Telecommunications Standards Institute.

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

Goodyear, P. (2015). Teaching as design. *HERDSA Review of Higher Education, 2*, 27–50.

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

International Organization for Standardization. (2014). *ISO 14289-1:2014. Document management applications: Electronic document file format enhancement for accessibility, Part 1: Use of ISO 32000-1 (PDF/UA-1)*. ISO.

Laurillard, D. (2012). *Teaching as a design science: Building pedagogical patterns for learning and technology*. Routledge. https://doi.org/10.4324/9780203125083

Reinecke, K., & Bernstein, A. (2011). Improving performance, perceived usability, and aesthetics with culturally adaptive user interfaces. *ACM Transactions on Computer-Human Interaction, 18*(2), Article 8. https://doi.org/10.1145/1970378.1970382

Rello, L., & Baeza-Yates, R. (2013). Good fonts for dyslexia. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '13)*. ACM. https://doi.org/10.1145/2513383.2513447

Rello, L., & Baeza-Yates, R. (2016). The effect of font type on screen readability by people with dyslexia. *ACM Transactions on Accessible Computing, 8*(4), Article 15. https://doi.org/10.1145/2897736

Rello, L., Pielot, M., Marcos, M. C., & Carlini, R. (2013). Size matters (spacing not): 18 points for a dyslexic-friendly Wikipedia. In *Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility (W4A '13)*. ACM. https://doi.org/10.1145/2461121.2461127

Royal National Institute of Blind People. (n.d.). *Clear print guidelines*. https://www.rnib.org.uk

United States Access Board. (2017). *Information and communication technology (ICT) standards and guidelines* (Section 508 of the Rehabilitation Act). https://www.access-board.gov

W3C. (2018). *Web Content Accessibility Guidelines (WCAG) 2.1*. https://www.w3.org/TR/WCAG21/

W3C. (2021). *Making content usable for people with cognitive and learning disabilities (COGA)*. https://www.w3.org/TR/coga-usable/

W3C. (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*. https://www.w3.org/TR/WCAG22/

Wilkins, A. J. (2003). *Reading through colour*. Wiley.

Zorzi, M., Barbiero, C., Facoetti, A., Lonciari, I., Carrozzi, M., Montico, M., Bravar, L., Kelly, F., Butterworth, B., & Ziegler, J. C. (2012). Extra-large letter spacing improves reading in dyslexia. *Proceedings of the National Academy of Sciences, 109*(28), 11455–11459. https://doi.org/10.1073/pnas.1205566109

Appendix A: Further Variant Profiles

The following profiles extend the worked examples in the body. Each combines deterministic parameters, applied by the rendering engine, with non-deterministic ones, interpreted by the LLM, and each is marked accordingly in the Type column. As with the profiles in the main text, they are offered as principled starting points for a teacher or specialist to adapt rather than validated specifications, and any of them can be layered onto a pedagogy from anywhere in this series. A teacher could keep the ones that prove useful in the same shared spreadsheet that already holds the profiles from Paper 2.

Variant Profile: Low Vision (Large Print)

Parameter	Rule	Type
[Font]	Sans-serif, minimum 24pt; bold headings; never italics or all-capitals	**Deterministic**
[Spacing]	Line height 2.0; wide letter and word spacing	**Deterministic**

[Contrast]	Black on white, or white on black on request; never colour alone	**Deterministic**
[Layout]	Single column, maximum width 480px, generous margins	**Deterministic**
[Element styling]	A strong rule between elements; element name large and bold	**Deterministic**
[Language]	Expand abbreviations and symbols into words on first use	**Non-deterministic**

The Low Vision profile is almost entirely deterministic: large type, generous spacing, and strong contrast are exactly the kind of parameters a rendering engine can guarantee on every item, identically and at no further cost. Its single non-deterministic rule, expanding abbreviations and symbols into words, is a text-level judgement the LLM makes as it writes.

Variant Profile: Autism-Friendly (Clear Structure)

Parameter	Rule	Type
[Font]	Sans-serif, 16pt	**Deterministic**
[Background]	Muted and low-glare (off-white #FBFBF7); no patterns behind text	**Deterministic**
[Layout]	One consistent layout reused on every item, in a predictable order	**Deterministic**
[Element styling]	Each element clearly bounded and labelled with what it is for	**Deterministic**
[Language]	Literal and explicit; avoid idiom, irony, and ambiguous pronouns	**Non-deterministic**
[Examples]	Concrete and unambiguous, introducing one idea at a time	**Non-deterministic**
[Tone]	Calm and matter-of-fact; signal clearly when something is a question	**Non-deterministic**

The Autism-Friendly profile leans on a deterministic strength that is easy to overlook, which is consistency. Because the layout is fixed in the profile, every content item a learner receives has the same predictable structure, and that predictability is itself part of the support. The non-deterministic rules then steer the language towards the literal and the explicit. A specialist who knows a particular learner well could refine these rules with that learner in mind, and the result would travel with the learner across every subject.

Variant Profile: Dry-Wit Science Voice

Parameter	Rule	Type
[Font]	Clean serif, 12pt	**Deterministic**

[Element styling]	Understated; a single thin accent rule and nothing flashier	**Deterministic**
[Language]	Informal but precise; contractions are fine; no exclamation marks	**Non-deterministic**
[Examples]	Everyday, and chosen to be a little unexpected	**Non-deterministic**
[Tone]	Confident, dry, occasionally self-aware; one moment of understatement per item, never a punchline, and never at the learner’s expense	**Non-deterministic**

The Dry-Wit profile is the mirror image of the first: almost all of its work is non-deterministic, in the register and tone the LLM adopts, with only a light deterministic touch in the typography. It illustrates that a variant profile need not be about access or audience at all. It can simply be a teacher’s voice, captured once and applied across a whole library so that everything sounds recognisably like them. As Paper 2 observed, a teacher who has cultivated a particular classroom persona over years can encode it here and stop having to perform it by hand in every resource.