Natural Language as the Bridge Between Pedagogy and AI

Paper 1 of 4: Encoding Pedagogy –
Elements and Parameters for Constrained LLM Generation

Dr Joe Watkins

Grassroots Intelligence Limited, Canterbury, UK

Correspondence: joe@grassroots-intelligence.com

Abstract

Pedagogical expertise is often expressed in rule-like, natural language terms. This has direct application to LLM content generation: the same natural language rules that describe a pedagogy can serve as structured constraints that guide LLM output. This paper presents a lightweight framework for encoding this expertise as elements and pedagogies. The framework decomposes educational content into discrete *elements*, each guided by *parameters* that specify both per-element constraints and inter-element relationships. These two constructs (elements and parameters) provide a grammar for expressing a wide range of approaches as a set of natural language rules that LLMs can interpret and follow.

The framework has been developed through direct professional experience delivering AI professional development and sustained work groups with practising teachers and content creators across primary, secondary, further, and adult education, and reflects the practical constraints and opportunities that educators encounter when working with these tools. We believe our approach to be agnostic, subject-independent, and accessible to any educator who can articulate what they want their content to do. The framework is illustrated through worked examples that teachers can use in any LLM to begin applying the framework.

This is the first of four papers; subsequent papers extend this foundational architecture to audience-specific tailoring (Paper 2), visual and spatial design (Paper 3), and temporal sequencing for video production (Paper 4). Taken together, the series describes how this framework can support structured content generation at considerable scale, with the ability to tailor to individual student needs, while maintaining the pedagogical consistency that comes from encoding expertise as explicit, reusable rules.

Keywords: content generation, constrained LLM generation, prompt engineering, educational technology, natural language processing, element-based content design.

1 Introduction

The invention of LLMs such as ChatGPT and Copilot has made it possible for anyone to generate educational content at scale. This capability is transformative, but it carries a fundamental risk: without structured constraints, LLM-generated educational content is often fluent but pedagogically incoherent. The model produces text that may read well but which follows no instructional logic, maintains no cognitive arc, and embodies no particular theory of learning.

To see this risk concretely, consider a prompt that a teacher might reasonably try: "Write three short paragraphs about density for secondary school students." A typical LLM output might be:

> *Density is a measure of how much mass is contained in a given volume. It is usually expressed in units such as kilograms per cubic metre (kg/m³) or grams per cubic centimetre (g/cm³). Different materials have different densities, which is why some objects feel heavier than others even when they are the same size.*
>
> *In everyday life, density explains many things we observe. For example, oil floats on water because oil is less dense than water. Similarly, a balloon filled with helium rises because helium is less dense than the surrounding air. Understanding density helps us make sense of why materials behave the way they do.*
>
> *Density is an important concept in science and engineering. It is used in designing ships, aircraft, and buildings. Scientists also use density to identify unknown substances by comparing their density to known values. Learning about density gives you a foundation for understanding many areas of physics and chemistry.*

The three paragraphs cover the topic adequately. However, there is no arc, no deliberate progression, and no moment where the learner is challenged to think again. The paragraphs could be reordered without any loss of coherence, because there is no structural relationship between them. A teacher reading this would recognise it as the sort of content that *looks* like educational material but does none of the work that makes content actually effective.

We believe that the opportunity of our framework – or any equivalent framework – is vast, and can be understood along three axes. The first reason is **scale**: a well-structured pedagogical design, expressed as labelled elements with parameters, can generate content across hundreds of topics in the time it previously took to produce one. The second is **tailoring**: the same content can be adapted for different learners (different languages, reading ages, cultural contexts, neurodiversity needs) while preserving the structure. The third, less immediately obvious but potentially most significant, is **measurability**: when decisions are encoded as explicit rules, they become shareable, comparable, and testable in ways that tacit pedagogical expertise never can be.

This paper draws on several years of experience running teacher CPD workgroups in AI at both local and national level, as well as working with content creators across a range of sectors including education, corporate training, and public engagement. The approach follows the tradition of design-based research (Barab & Squire, 2004): the framework has been iteratively developed and refined through sustained use with practising teachers in authentic professional contexts. The observations and design claims presented here are grounded in that practitioner experience. While they are not yet supported by controlled empirical data, they reflect what we have consistently observed through extensive, sustained work with real educators producing real content. We are now beginning our first set of collaborative research projects and are open to collaborators who wish to make use of our own proprietary system to help conduct their own pedagogical studies using our framework.

2 The framework as a complement to teaching

Many teachers are understandably reluctant to use LLMs in their practice. Concerns about accuracy, quality, data privacy, the environmental cost of large-scale computation, and the commercial interests of technology providers are all routinely observed and entirely legitimate. Especially pertinent are the concerns many teachers have about the over-reliance of students on technology like LLMs. Some teachers have faced institutional pressure to adopt technology without evidence of improved outcomes and are understandably sceptical of yet another tool promising transformation. This paper takes these concerns seriously.

Consider a teacher's working life. They plan lessons, deliver them, respond to students, adapt on the fly, mark work, and provide feedback. Content production such as worksheets, slides, revision materials, extension tasks, and homework is simply a collection of essential but time-consuming tasks. This competes with time for the more rewarding aspects of teaching, including human interaction, responsive teaching, relationship-building, and reflection.

The introduction of this framework aims to empower teachers to begin producing pedagogical content items as a consequence of teaching, not as a separate activity, and in so doing harnessing and reflecting upon their expertise and curated knowledge. A teacher who has effectively taught the causes of World War One for ten years has accumulated lesson plans, explanations, examples, analogies, common misconceptions, and corrective strategies. Under the current model, this expertise will yield a few dozen handmade resources. With the framework that we present, the same expertise readily produces hundreds of content items that embody that expertise and can reach audiences in a way that was previously not possible.

Encoding a pedagogical approach as explicit rules is a useful exercise that encourages teachers to reflect on what they do and why. Questions about sequencing explanations or the order of examples become more conscious. The framework encourages this reflective abstraction, what Schön (1983) called “reflection-on-action”, as a natural consequence of using the system. We hope that teachers who use this framework will begin to encode their practice, empowering them to test their theories at scale and then quickly amplify and share their successes with colleagues and the wider education sector.

3 Core argument

When an experienced teacher plans a lesson, they tend to follow well-tested patterns. For example, they might cover the most important concept first, then give a concrete example, and then check understanding before moving on. These patterns may not always be consciously articulated as rules, but they function as rules, shaped by training, experience, and pedagogical tradition. This paper argues that making such patterns explicit and encoding them as structured constraints for LLMs produces educational content that is dramatically better than unconstrained LLM generation, and that this encoding process is accessible to any teacher without technical knowledge.

The argument rests on two foundational observations, which together create the opportunity that this paper and the wider series describe.

3.1 Pillar 1: LLMs produce better output when given structured constraints

There is now substantial research demonstrating that LLMs perform better when given structured constraints rather than open-ended instructions (Sahoo et al., 2024; Wei et al., 2022). The “sculpting” approach (constrained, rule-based prompting) combines step-by-step reasoning with explicit restrictive rules, and has been shown to outperform standard chain-of-thought prompting on mid-tier models (Khan, 2025). The CRANE framework demonstrates up to 10% improvement on reasoning benchmarks when constraints are carefully designed to preserve reasoning while maintaining correctness (Banerjee et al., 2025). Liu et al. (2024) found from a user-centred study that structured output constraints are widely considered necessary for reliable LLM integration, with professionals identifying scenarios where constraints are "necessary or advantageous" across content generation, classification, and domain-specific tasks.

The key insight from this literature is not simply that constraints help, but *why* they help. Setting constraints for an LLM will reduce the search space for the model, making it easier to produce an output that meets specific quality criteria. When the constraint is pedagogical (e.g. "cover the most critical concept in under 40 words"), the model's output is shaped by available knowledge encoded as natural language rules.

Qian (2025) provides a systematic review of prompt engineering specifically in education, finding that structured prompting (using predefined constraints) was the most frequently utilised strategy across 22 studies. The review found that "core components of prompt engineering involve specifying roles, limiting the scope, and emphasizing the format of requested outputs" and that "LLMs can better interpret user queries when prompts are crafted clearly." Patil et al. (2024) demonstrated domain-specific prompt templates with constraints producing consistent, high-quality content in healthcare education. Taken together, this body of work suggests that the more precisely an LLM is constrained, the more reliably it produces output that meets the intended quality criteria.

3.2 Pillar 2: Pedagogical expertise is often already expressed in rule-like terms

Much of what constitutes established pedagogical thinking can be understood as sets of rules or principles. Rosenshine (2012) expressed his Principles of Instruction as a numbered list of directives (present new material in small steps, ask questions frequently, provide models, guide student practice). Cognitive Load Theory (Sweller, 1988; Sweller et al., 2011) provides principles about what to present simultaneously and what to separate. Bloom's Taxonomy (Anderson & Krathwohl, 2001) offers a sequencing framework for cognitive demand. Hattie's Visible Learning framework encodes effect sizes as decision heuristics for which strategies to prioritise (Hattie, 2009).

Merrill's First Principles of Instruction (Merrill, 2002) decompose effective instruction into activation, demonstration, application, and integration: these are, in effect, labelled components with an implicit sequence. Gagné's Nine Events of Instruction (Gagné, 1985) go further still, specifying a nine-step sequence from gaining attention through guided practice to enhancing retention, each step with a defined role and relationship to the others. What Shulman (1986) termed *pedagogical content knowledge* (PCK), the distinctive expertise that teachers hold about how to teach particular content, is precisely the kind of knowledge that these frameworks attempt to codify and that our framework makes accessible to a wide audience. Dillenbourg (2002) demonstrated that similar rule-based structuring can be applied to collaborative learning through structured activity sequences with defined roles and constraints.

This is not to claim that all teaching can or should be reduced to frameworks and rules. Much of what makes a teacher effective is what Polanyi (1966) called *tacit knowledge*: the kind of expertise that is difficult to articulate but that manifests in practice. Teaching involves judgment, improvisation, relationship, and context-sensitivity that no rule set can fully capture. But the *structural decisions* that teachers make about content organisation, sequencing, and presentation are frequently rule-like in character, and many teachers, when asked to describe their approach, naturally express it in terms that sound very much like rules: "I always start with a concrete example," "I never introduce more than one new concept per slide," "I make sure the checking question requires them to combine what they've just learned."

This is a crucial observation for the framework described here, because these natural language expressions of intent are *exactly the form* that LLMs are designed to interpret and follow. The entire interaction between teacher and system takes place in the same natural language that a teacher would use to explain their approach to a colleague, brief a trainee, or give feedback on a student teacher's lesson plan. A teacher who says "I want the definition first, then a concrete example, then something that makes them think again" is, in effect, writing a prompt that any modern LLM can interpret and act on. To illustrate: a teacher could type exactly this into any LLM …

> *"Write a short piece on the water cycle. Give me: first, a clear definition; second, an example a child could visualise; third, a surprising fact that challenges what the first two established."*

... and receive a structured output. Although simple, the above *is* a statement of teaching, ultimately encoding a learning journey. Even an approach of this simplicity will typically yield results that are vastly more tailored and preferable than any generic prompt.

The framework described in this paper simply shows how this process may be made systematic, reusable, verifiable, and scalable. A teacher who reviews the LLM's output and says "the example is too abstract, please make it more concrete in future" is issuing a natural language edit that refines the output within the same conversational interaction, based on their expertise and judgement. The entire cycle is conducted in the teacher's own language, at the teacher's own level of specificity, with no prohibitive syntax or technical translation required. A central philosophical position of this paper is that we view the expression of this expertise as a pedagogy in its own right. It may be simple, and it may only be applicable in particular situations or for particular students, but it is a reusable piece of rule-based instructional design grounded in years of expertise and refinement.

3.3 The convergence

When these two pillars are placed together, the opportunity becomes clear. LLMs benefit from structured, rule-based constraints (Pillar 1). Teachers naturally express their expertise in rule-like, natural language terms (Pillar 2). The framework described in this paper simply connects the two: it provides a lightweight architecture for capturing a teacher's pedagogical rules and passing them to an LLM as structured constraints on content generation. The result is LLM-generated educational content that is guided by the teacher's own expertise, produced at scale, and adaptable across audiences. The entry point is, quite literally, how teachers already talk about what they do.

It is worth noting that this convergence is increasingly well-supported by developments in LLM practice. The growing use of prompt chaining (where the output of one constrained generation step becomes the input to the next) and agentic workflows (where LLM systems follow multi-step, rule-guided procedures with conditional logic) demonstrates that the field is moving toward structured, sequential constraint architectures. The framework described here is consistent with this direction: a content item may pass through multiple constrained generation steps, each shaped by rules that the educator has defined. The approach is, in effect, a pedagogically-driven prompt chain.

Crucially, modern LLM tools make the sharing of codified expertise remarkably straightforward. A teacher who defines a structure as a prompt can share it as a custom GPT, a Copilot agent, a Google Gem, or simply as text in a prompt gallery or shared document. A department head who encodes their preferred approach can distribute it to every colleague by sharing a link. It has never been easier for teachers to share practice once codified, and the framework described in this paper provides the language in which that codification happens.

As a final note, our choice of syntax and structure is mostly for legibility and simplicity. In our testing it is the case that LLMs tend to benefit from *any* structured approach, and the specifics of that structure are usually less important. We elaborate on this later in the paper, but we would encourage practitioners to experiment with their own methods.

4 Elements and variables as a grammar of pedagogical encoding

The central claim of this paper, and the idea that is pervasive throughout the series, is this: **many forms of pedagogical content can be usefully decomposed into discrete elements, where each element encodes a distinct part of the approach, and where the relationships between those elements are expressed as rules within the elements' own parameters.**

This decomposition is what makes the framework possible: it allows individual parts of the approach to be constrained independently, relationships to be specified explicitly, variants to be generated while preserving structure (Paper 2), visual decisions to be applied per-element (Paper 3), and temporal sequencing to be specified at the element level (Paper 4). The concept of decomposing educational content into reusable, reconfigurable components has precedent in the literature (Wiley, 2000; Merrill, 2002), where questions of granularity, reusability, and composability have been extensively studied. Our framework can be understood as extending this tradition: elements are reusable components, but their parameters ensure they cohere as a sequence rather than merely coexisting. Here are the foundational definitions of this series:

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. The set of constraints that guide a single element is expressed through its *parameters*.

Parameter. Any individually-adjustable property within a parameter set. Parameters may include a natural language *rule* describing the element's role (e.g. "a clear, plain-English definition") as well as quantitative constraints such as *word count* (e.g. max 30 words). Within this paper, we deliberately keep to a small number of parameters, although in practice this can be far more extensive.

Variable. A piece of input information that applies to the whole pedagogy, rather than individual elements. The most common variables are *topic* and *source material*. This concept is extended significantly in Paper 2 to the idea of 'variant profiles'.

Pedagogy. Any combination of elements and parameters, using the definitions above.

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.

These terms are used consistently throughout the series; subsequent papers will introduce additional terminology as needed.

5 Examples of simple pedagogies

To make these definitions concrete, we will initially consider a simple three-element pedagogy that a teacher might use: '*give students a clear definition, then a relatable example, then a surprising fact that makes them think again'*. Here is how that pedagogy would be expressed in the framework.

Simple three-element pedagogy:

Below is a simple three-element pedagogy of this type with two parameters. We have not yet specified any variables for this pedagogy: neither the *topic* nor any *source material*. We are currently treating this as an abstract pedagogy that is topic-agnostic, and we are not yet concerning ourselves with the intended audience. The word counts that we have included are purely for the sake of demonstration, although including a maximum word count is a time-tested method for applying structure to any LLM output.

Element	Parameter 1: Rule	Parameter 2: Max words
[Definition]	A clear, plain-English definition of the key term	30
[Example]	A concrete, everyday example	40
[Surprising Fact]	A surprising or counterintuitive fact	35

These parameters are perfectly valid as they stand. Each describes its own element without referencing any other. An LLM given these three parameters and a topic would produce three pieces of content that individually satisfy their constraints. However, we should notice what is missing: there is no guarantee that **[Example]** actually illustrates **[Definition]**, or that **[Surprising Fact]** has any relationship to what came before. The three elements might be excellent individually but they are not necessarily related as a sequence. If an LLM were used to generate these elements separately from their parameters and any given topic, then there would be no guarantee that they would cohere.

Throughout this paper, we present pedagogies as single prompts in which all elements and their parameters are submitted together, as opposed to a set of prompts entered in sequence. Current frontier models can comfortably process an entire element set in a single interaction, and doing so is preferable because the model has access to the full context of every element and inter-element rule simultaneously, which helps the output cohere as a sequence rather than a collection of independent pieces. There is also a practical advantage to the all-in-one approach: it is simpler. A teacher pastes one prompt and receives one complete content item, rather than managing a sequence of separate calls and stitching the results together.

Simple three-element pedagogy with inter-element parameters:

Here is a more refined version of the pedagogy above, where the first 'rule' parameter is allowed to refer to other elements. This pedagogy is significantly stricter than above, but equally straightforward for a teacher to author.

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 more information; must make the learner reconsider.	35

The inter-element rules are now part of each element's own parameters, expressed as natural language references to other elements by label. This is the framework's mechanism for ensuring what Biggs (1996) called *constructive alignment*: the principle that all components of an instructional sequence should be aligned with each other and with the intended learning outcome. Gagné (1985) made a similar point through his sequencing of instructional events, where each event must support and build upon the previous one. In this pedagogy, the alignment is enforced by the parameters themselves.

5.1 Same structure, new pedagogies

The "definition + example + surprising fact" approach is one teacher's method for a simple three-element pedagogy to introduce a topic. However, the same fundamental structure can encode entirely different strategies simply by changing the parameters. Consider Pedagogy A as a reminder:

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

Now consider two alternatives that use the same three-element pedagogy but encode entirely different approaches simply by varying the parameters.

Pedagogy B: "Hook + Explanation + Check"

Element	Parameter 1: Rule	Parameter 2: Max words
[Hook]	An intriguing question or puzzle that provokes curiosity	25
[Explanation]	The explanation that directly answers and resolves **[Hook]**	50
[Check]	A checking question that requires application of **[Explanation]**, not repetition	20

Pedagogy C: “Misconception + Correction + Evidence”

Element	Parameter 1: Rule	Parameter 2: Max words
[Misconception]	A common misconception, stated as a student might phrase it	25
[Correction]	A clear correction that directly contradicts **[Misconception]**	40
[Evidence]	Concrete evidence supporting **[Correction]**; must not simply restate it	35

Each of these is a complete, self-contained pedagogy expressed as three labelled elements with two parameters. Applied to the same topic, they will produce fundamentally different learning experiences. The structure of pedagogy A creates a cognitive arc from understanding to puzzlement, whereas pedagogy B starts with curiosity and resolves it. Separate from these, pedagogy C confronts a wrong belief and replaces it with evidence. These are three different pedagogies, with three distinct learning journeys, all expressed with straightforward natural language rules in the form of parameters.

The power and flexibility of expressing pedagogies in such natural language is immediately visible here: changing the overall pedagogical journey is as simple as rewriting the two parameters. It is also worth noting that a single 'rule' parameter of this form is not required by any means, but is in our view the simplest way for a teacher to quickly describe the overall essence of an element. The maximum word count is another parameter whose effect is easy to understand and to test.

5.2 A note on audience tailoring

The pedagogies above deliberately avoid any parameters that tailor content to a specific audience. There are no constraints specifying reading age, difficulty level, or target year group. Translating the same pedagogy into another language, for example, is simply a matter of adding one audience constraint, and the pedagogical structure remains identical. For neurodiverse learners, variations in vocabulary, sentence length, and processing demand can produce substantially different learning experiences. Cultural context and prior knowledge similarly shape how information is received and understood. By focusing solely on pedagogical intent, we can create a more universally applicable framework that can be adapted to diverse audiences after generation.

To see why, consider adding audience parameters to Pedagogy A:

Element	Parameter 1: Rule	Parameter 2: Max words	Parameter 3: Audience
[Definition]	A clear, plain-English definition of the key term	30	Accessible to Key Stage 4 Polish native speakers.
[Example]	A concrete, everyday example that illustrates **[Definition]**. Must use a tangible scenario the learner can visualise.	40	Accessible to Key Stage 4 Polish native speakers.

[Surprising Fact]	A surprising fact that challenges **[Definition]** and **[Example]**. Must make the learner reconsider.	35	Accessible to Key Stage 4 Polish native speakers.

Notice that the audience parameters sit awkwardly alongside the pedagogical ones. They are a different kind of constraint; they describe who the content is for and not what the content should do. Furthermore, they create a scaling problem. If the same approach needs to serve Key Stage 3, A-level, and adult learners, we would need three separate pedagogies that differ only in their audience column, which remains constant across the different applications. A key principle of our approach is that audience tailoring should be dealt with, effectively, as variables, rather than parameters. This is the subject of Paper 2, where we develop this idea in full.

5.3 Applying a pedagogy to a topic

To see the framework in action, we now take Pedagogy A from Section 5.1 and apply it to a specific topic. The following prompt can be pasted directly into any LLM as a fully natural language expression of the table version given in Section 5.1:

> 📋 **COPY-PASTE PROMPT: Pedagogy A applied to a topic**
>
> *Please generate a three-element educational content item on the following topic.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept defined in* **[Definition]***. Maximum 40 words. Must use a tangible, physical scenario the learner can visualise.*
>
> **[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**

Notice that the prompt given to the LLM *is* the pedagogy, plus a topic entered as a variable. The parameters from Section 5 have been expressed as natural language instructions and the topic has been entered as a variable.

The practical power of this becomes clear when we consider that applying the same pedagogy to a different topic is as simple as changing the variable above. This gives immediate advantages in scale and repeatability: a teacher who is satisfied with the pedagogy can apply it across an entire curriculum without reworking the design. At its simplest, this means changing one word in a prompt. At scale, the same principle can be applied programmatically across hundreds of topics in a single operation, as we describe in the batching strategies of Section 6.1.

5.4 Source material: quality in, quality out

Section 4 defined ‘source material’ as a variable containing subject-matter content the teacher provides alongside the pedagogy such as lesson notes, schemes of work, textbook passages, exam-board specifications, or any other trusted reference. Where source material is provided, the LLM draws preferentially from it and where it is not, the model falls back on its general training knowledge. As with specifying the topic, source material is treated as a variable.

To illustrate, the following extends the Pedagogy A prompt from Section 5.3 with source material:

> 📋 **COPY-PASTE PROMPT: Pedagogy A with source material**
>
> **<<SOURCE_MATERIAL>>** 📎 **[GCSE_physics_notes_density.pdf attached]**
>
> *Using the attached source material as your primary reference, generate a three-element content item on the following topic.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words. Draw phrasing from the attached notes where possible.*
>
> **[Example]***: A concrete, everyday example that illustrates* **[Definition]***. Maximum 40 words. Must use a tangible, physical scenario the learner can visualise.*
>
> **[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**

The pedagogy is identical to Section 5.3; the only addition is the source material as a variable. The output will be grounded in the teacher's own notes rather than the LLM's general knowledge.

Including source material brings several practical benefits. The LLM grounds its output in the teacher's own content rather than general training data, reducing hallucination and ensuring factual alignment with what students have actually been taught. Source material also anchors the output to a particular curriculum, exam board, or scheme of work, so that a **[Definition]** drawn from a teacher's lesson notes will tend to use the same phrasing and emphasis students have already encountered. Content that visibly draws on the teacher's own materials is easier to verify and more likely to be adopted, which builds trust in the framework. Finally, when the same source material is used across a batch, the resulting content library maintains a consistent register, terminology, and level of detail. Taken together, these benefits mean that a teacher who invests time in curating their source material will consistently produce better results than one relying on the LLM's general knowledge alone.

5.5 Pedagogy inference using LLMs

For teachers who are new to working with LLMs, we have found that one of the most effective starting points is to work from existing materials, which are themselves a product of the teacher's own expertise. A teacher does not need to define their elements and pedagogies from scratch, given that modern LLMs can be effective at *inferring* structure from existing materials. A teacher who pastes an existing worksheet, lesson plan, or set of revision notes into an LLM and asks it to identify the underlying structure will often receive a usable first draft of labelled elements with parameters that they can then refine with natural language commands.

> 📋 **COPY-PASTE PROMPT: Pedagogy inference**
>
> **<<SOURCE_MATERIAL>>** 📎 **[tectonic_plates_worksheet.pdf attached]**
>
> *This is a worksheet I use with my Year 9 class on tectonic plates.*
>
> *Analyse its structure and tell me:*
>
> *1. What distinct elements does it contain? Give each a descriptive label in square brackets.*
>
> *2. For each element, what rules seem to guide it? (e.g. word count, level of complexity, type of content). Refer to these as 'parameters'. There should be one main rule to determine the overall function of each element within the pedagogy.*
>
> *3. For elements that relate to each other, include inter-element references within each element's parameters where relevant.*
>
> *Present your analysis as a table of labelled elements, each with parameters that includes any rules about how it relates to other elements.*

This is, in effect, reverse engineering with LLMs. The teacher provides a finished product and asks the LLM to recover the pedagogy that produced it. The result is a set of labelled elements with inferred parameters, ready to be reviewed, adjusted, and reused. It can be a very quick start to answering the question, "how do I write down what I did here?"

Due to the non-deterministic nature of LLMs (see Appendix B), the output will be different from the above, and a teacher might prefer to be more or less specific in the initial prompt. They might prefer a different set of syntax entirely to the one our prompt suggests. However, from whatever response is offered, the teacher can then use their own expertise to quickly judge the output and adjust it using the common tool of natural language expressions.

A teacher who has been producing effective materials for years already has a library of implicit pedagogies. LLM-based inference makes those pedagogies explicit, turning years of accumulated craft knowledge into reusable, shareable, testable structures in a matter of minutes. Combined with the source material approach described in Section 5.4, a teacher can provide both the *content* they trust and have the *structure* they already use inferred automatically, giving them a complete, working pedagogy without writing a single parameter from scratch.

6 Measuring pedagogies: from craft knowledge to testable hypotheses

The framework makes pedagogical decisions measurable by transforming teaching methods into explicit, testable components. Beyond measurability, the framework also makes it straightforward for a teacher to take an idea that works and scale it rapidly. A pedagogy that produces effective content for one topic can be applied across an entire curriculum in minutes, something that would previously have required hours of manual effort. Since every element and parameter is visible and individually adjustable, the framework creates the conditions for systematic comparison in a way that has not previously been practical. The batching strategies described in Section 6.1 formalise this process.

Some of the key factors that can be explored readily with this framework are as follows:

Controlled variation: holding everything constant except one variable. A teacher who wants to know whether ending with a **[Surprising Fact]** produces better learning than ending with a **[Check]** can generate matched versions of the same content item and compare the results.

Granularity: variables can be isolated at any level of detail. A teacher could vary a single word count, a single inter-element rule (does **[Example]** that *illustrates* **[Definition]** work better than one that *contradicts* it?), or an entire element structure while holding the rest constant.
Scale: any variation can be applied across an entire topic library as a batch operation, producing hundreds of matched pairs as a by-product of the framework's normal operation.
Shareability: a pedagogy is a transferable artefact. A teacher in Manchester who develops a particularly effective five-element structure for teaching can share that pedagogy instantly with a collection of teachers in Lagos, who can apply it to different topics in a different language and consider it on the basis of their own expertise. Over time, libraries of pedagogies, not just content, could accumulate and be refined through collaborative use (Laurillard, 2012; Goodyear, 2015).

Here is a reminder of a pedagogy that we have already seen, applied to the topic of density.

> 📋 **Pedagogy A applied to Density**
>
> *Please generate a three-element educational content item on the following topic.*
>
> **[Definition]***: A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]***: A concrete, everyday example that illustrates the concept defined in* **[Definition]***. Maximum 40 words. Must use a tangible, physical scenario the learner can visualise.*
>
> **[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**

Having generated this content item, the teacher can then explore each of the four variation types described above by adding a follow-up prompt. This could be included as part of the original prompt if desired, although for a first pass it is normally more revealing to judge the initial output before requesting variations.

> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Controlled variation**
>
> *Now repeat the above, but generate two versions. In Version A, cap* **[Definition]** *at 15 words. In Version B, allow up to 60 words. Present both side by side so I can compare how definition length affects overall quality.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Granularity**
>
> *Now repeat the above, but change the rule for* **[Example]** *so that it must contradict* **[Definition]** *rather than illustrate it. Present both versions so I can compare how this single rule change affects the learning journey.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Element substitution**
>
> *Now repeat the above, but replace* **[Surprising Fact]** *with* **[Retrieval Question]***: a question that requires the learner to recall and apply what* **[Definition]** *and* **[Example]** *established. Present the original and modified versions side by side.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Pedagogy comparison**

> *Now generate the same topic using this* **<<ALTERNATIVE PEDAGOGY>>**. *Present the outputs side by side so I can compare which approach works better.*

Each follow-up prompt isolates a single design decision, and each can be run as a continuation of the same interaction. The teacher ends up with matched outputs that differ in only one respect, giving them a concrete basis for comparing approaches.

In these methods the non-deterministic nature of the LLMs themselves can be an advantage. Such a framework also makes it straightforward to generate many different instances of the same pedagogy for comparative testing. Given that every parameter is explicit and individually adjustable, a teacher can produce matched sets of content that differ in only one respect, isolating the effect of a single design decision across a full topic set.

6.1 Batching strategies

We give two examples below. The first demonstrates the most fundamental batching operation: one pedagogy applied across an entire section of curriculum. The second demonstrates parameter tweaking, generating matched variants that differ in a single constraint.

Example 1: One pedagogy, many topics

The following prompt applies Pedagogy A from Section 5.1 over a range of topics. In this instance, we have not even listed the topics, instead choosing a quicker, implied method, by asking for all topics from a set area of a subject.

> 📋 **COPY-PASTE PROMPT: Curriculum batch**
>
> **THE PEDAGOGY:**
>
> **[Definition]**: *A clear, plain-English definition of the key term. Maximum 30 words.*
>
> **[Example]**: *A concrete, everyday example that illustrates* **[Definition]**. *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.*
>
> *I would like one each for each* **<<TOPIC>>** *in the Key Stage 3 curriculum on powers and indices. Please present each in an easy format to download and share.*

This prompt will produce many pieces of structured content in a single interaction and the iteration over many topics is itself a batching strategy that can have considerable value. The pedagogy itself is identical for every topic; only the *topic* variable changes and is here specified implicitly rather than explicitly. The output is a complete, pedagogically consistent content library for that area of the curriculum, produced in the time it would previously have taken to draft one or two items by hand.

In this case, follow-up prompts can be used equally as well to refer to the batched output, as we show below.

> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Parameter tweaking**
>
> *Now repeat the above, but generate two versions of each topic. In Version A, cap* **[Definition]** *at 15 words. In Version B, allow up to 60 words. Present both side by side so I can compare how definition length affects overall quality.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Element substitution**
>
> *Now repeat the above, but replace* **[Surprising Fact]** *with* **[Retrieval Question]***: a question that requires the learner to recall and apply what* **[Definition]** *and* **[Example]** *established. Present the original and modified sets side by side.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Inter-element rule change**
>
> *Now repeat the above, but change the rule for* **[Example]** *so that it must contradict* **[Definition]** *rather than illustrate it. Present both versions so I can compare how this affects the learning journey.*
>
> 📋 **COPY-PASTE FOLLOW-UP PROMPT: Pedagogy comparison**
>
> *Now generate the same set of topics by comparison with the* **<<ALTERNATIVE PEDAGOGY>>**. *Present the outputs side by side so I can compare which approach works better for each topic.*

Each follow-up prompt isolates a single design decision across the full topic set, and each can be run in a single interaction. The result is a set of matched content libraries that differ in only one respect, giving the teacher a concrete, evidence-informed basis for comparing approaches.

Once recognised at the local (or higher) level, pedagogies and prompts that prove effective are easy to preserve and share. A teacher can email that to a colleague, or save a successful pedagogy as a text file in a prompt gallery, embed it in a custom GPT or Copilot agent, or store it in a shared document for colleagues to reuse. The pedagogy becomes a reusable tool rather than a one-off interaction. Equally, the rules within a pedagogy are easy to update in response to feedback: if a teacher finds that **[Example]** consistently produces scenarios that are too abstract, they can tighten the rule ("must involve a named, everyday object the learner can hold") and regenerate the entire library in minutes. The pedagogy evolves with the teacher's judgement, and every iteration is captured in the parameters themselves. Appendix A describes the full implementation spectrum, from saving effective prompts in a gallery through custom agents to vibe-coded production systems.

7 Scope and limitations

This paper presents a framework, not a finished product. Several limitations should be acknowledged.

First, the framework is theoretical and design-based. We have a working implementation and have used it extensively in teacher training and content production, but we have not yet conducted controlled empirical studies measuring learning outcomes. The claims about output quality are based on practitioner experience, not experimental data. At the time of writing, we are actively planning controlled studies with education partners to measure the effect of different pedagogies on learning outcomes. The parametric structure described in Section 6 makes such validation unusually tractable: because every element and parameter is explicit, pedagogical variants can be generated, deployed, and compared systematically at a scale that would have been impractical with hand-authored materials.

Second, LLM output is non-deterministic: the same prompt will not produce identical output on repeated runs. This has implications for reproducibility, version control, and empirical comparison. Appendix B discusses the nature of LLM non-determinism and its practical consequences for the framework.

Third, the worked examples in this paper are deliberately simple. We have used substantially more complex pedagogies in practice (five to seven elements, multi-layered inter-element rules with conditional logic), but these are harder to present concisely. Appendix C demonstrates a six-element Rosenshine-style lesson plan that gives a better sense of the framework's capacity.

Fourth, this paper concerns itself entirely with text: what each element says, how many words it uses, how it relates to other elements. Audience tailoring is deferred to Paper 2. Visual and spatial design is the subject of Paper 3. Temporal sequencing for video production is addressed in Paper 4. Each paper builds on the elements and parameters established here.

Finally, we acknowledge that the framework risks encouraging a view of teaching as the application of fixed rules. Effective teaching involves judgement, improvisation, and responsiveness that no rule set can fully capture. The framework encodes the *structural* decisions teachers make about content, not the full richness of the teaching act. For a detailed discussion of how the framework complements rather than replaces professional practice, see Section 2.

8 Conclusion

The examples in this paper are intentionally simple starting points. They are the sort of structures a teacher might sketch in five minutes. We hope that the framework is accessible enough for a brief sketch to produce something useful, and we expect that practitioners who engage with it will develop pedagogies far more sophisticated than anything demonstrated here. The pedagogy inference approach described in Section 5.5, which allows teachers to work from their existing materials, is one route we have explored to lower the barrier to entry.

The design space of the element-and-parameter architecture is large. Any teaching approach that can be described in natural language can, in principle, be decomposed into labelled elements with rule-based parameters. A Socratic questioning sequence, a problem-based learning cycle, a spaced-repetition schedule, or an inquiry-based investigation can all be expressed using the same grammar. The framework does not prescribe what good pedagogy looks like; it provides a notation in which any pedagogy can be written down, tested, and shared.

We would like to draw particular attention to the role of students. Students may not know the research on cognitive load or the principles of sequencing, but the framework operates through natural language, and students have the same linguistic capability as teachers to express what they want their learning to look like. A student who says "I want the definition first, then an example from a sport I actually play, then a question that makes me think" is articulating a pedagogy that expresses their own preferences about how they learn best. Their priorities may differ from their teacher's, and those differences are not deficiencies but a rich source of data. 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.

We have built and extensively tested multiple working implementations of the system described in this paper. We are keen to support researchers who wish to build on this framework in their own contexts, and we welcome contact from anyone interested in collaborating on empirical studies, implementation projects, or setting up their own content creation systems based on the principles described here. The inline prompts throughout Sections 5 and 6 are working examples a reader can use today, and the framework is designed to be extended and we have a working implementation that covers all principles across all principles described in this four paper series.

References

Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives* (Complete ed.). Longman.

Atil, B., Aykent, S., Chittams, A., Fu, L., Passonneau, R. J., Radcliffe, E., Rajagopal, G. R., Sloan, A., Tudrej, T., Ture, F., Wu, Z., Xu, L., & Baldwin, B. (2025). Non-determinism of "deterministic" LLM settings. *arXiv preprint arXiv:2408.04667v5*. https://arxiv.org/abs/2408.04667

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

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, *32*(3), 347–364. https://doi.org/10.1007/BF00138871

Dillenbourg, P. (2002). Over-scripting CSCL: The risks of blending collaborative learning with instructional design. In P. A. Kirschner (Ed.), *Three worlds of CSCL: Can we support CSCL?* (pp. 61–91). Open Universiteit Nederland.

Gagné, R. M. (1985). *The conditions of learning and theory of instruction* (4th ed.). Holt, Rinehart and Winston.

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

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

Khan, I. (2025). You don't need prompt engineering anymore: The prompting inversion. *arXiv preprint arXiv:2510.22251*. https://arxiv.org/abs/2510.22251

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

Liu, L., Liu, Y., Shin, S., Peng, S., & Zhao, L. (2024). We need structured output: Towards user-centered constraints on LLM output. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (CHI EA '24). ACM. https://doi.org/10.1145/3613905.3650756

Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development*, *50*(3), 43–59. https://doi.org/10.1007/BF02505024

Patil, R., Heston, T. F., & Bhuse, V. (2024). Prompt engineering in healthcare. *Electronics*, *13*(15), 2961. https://doi.org/10.3390/electronics13152961

Polanyi, M. (1966). *The tacit dimension*. University of Chicago Press.

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

Rosenshine, B. (2012). Principles of instruction: Research-based strategies that all teachers should know. *American Educator*, *36*(1), 12–19.

Sahoo, P., Singh, A. K., Saha, S., Jain, V., Mondal, S., & Chadha, A. (2024). A systematic survey of prompt engineering in LLMs: Techniques and applications. *arXiv preprint arXiv:2402.07927*. https://arxiv.org/abs/2402.07927

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

Shulman, L. S. (1986). Those who understand: Knowledge growth in teaching. *Educational Researcher*, *15*(2), 4–14. https://doi.org/10.3102/0013189X015002004

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

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in LLMs. In *Advances in Neural Information Processing Systems*, *35*, 24824–24837. https://arxiv.org/abs/2201.11903

Wiley, D. A. (2000). Connecting learning objects to instructional design theory: A definition, a metaphor, and a taxonomy. In D. A. Wiley (Ed.), *The instructional use of learning objects*. Agency for Instructional Technology.

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, *17*(2), 89–100. https://doi.org/10.1111/j.1469-7610.1976.tb00381.x

Appendix A: Implementation: from prompt to production

This appendix describes the practical implementation spectrum: from a copy-paste prompt in a free LLM to a full production system, all encoding the same architecture. The distance between these endpoints is, in 2026, remarkably short. Here are several layers of accessibility, depending on interest and expertise.

A prompt in an LLM. The simplest use of the framework is a prompt pasted into any LLM: ChatGPT, Claude, Copilot, Gemini, or any equivalent. The teacher defines their pedagogy as labelled elements with parameters, adds a *topic* and optionally *source material*, and generates. The prompts throughout Sections 5 and 6 are working examples of this approach.

A saved prompt. A prompt that works well should not be retyped. Most LLM interfaces now offer a prompt gallery or prompt library feature where a teacher can store a pedagogy as a reusable template. The *topic* becomes a blank to fill in, since the pedagogy is pre-loaded.

A custom agent. When a pedagogy has been tested and found effective, it can be embedded in a reusable agent: a Custom GPT (OpenAI), a Copilot agent (Microsoft) etc. Once the mechanics are understood, which in our experience is normally quick, agents of this type can quickly be tailored to include *archives* of successful pedagogies or reference materials, available to use for a wide range of teachers across many - or perhaps all - subjects. The possibilities for both uniformity and teacher-tailoring are vast and relatively unexplored.

Production systems. Beyond interactive use, the same pedagogies can be integrated into automated workflows that generate content across entire curricula, apply audience tailoring (Paper 2), render visual outputs (Paper 3), and compile narrated video walkthroughs (Paper 4). The degree of automation, customisation, and control that such a system offers is limited only by the engineering effort applied.

Appendix B: Non-determinism in LLM generation

There is an observation to make here about the nature of LLM-generated content that has important implications for how this framework should be understood and evaluated.

LLMs are fundamentally non-deterministic. Even with identical inputs, each generation will produce slightly different output. The same pedagogies applied to the same topic will yield a different definition, a different example, a different surprising fact each time the system runs. This is true even at very low temperature settings; research has demonstrated that LLMs exhibit measurable output variation across identical runs even when configured for maximum determinism, due to floating-point non-associativity, GPU parallelism, and architectural factors such as mixture-of-experts routing (Atil et al., 2025). (Atil et al., 2025)

This non-determinism is not a flaw to be engineered away. It is a fundamental characteristic of the medium, and proponents of this framework must be willing to accept its implications. The framework asks educators to trade *deterministic* control over content for *statistical* reliability, audience reach, and scale. A hand-written definition of density will be the same every time it is read. An LLM-generated definition of density, guided by the same parameters, will be slightly different each time it is generated. For a teacher, this non-determinism can also be a practical advantage. A pedagogy designed to produce question-practice examples will yield a different set of questions each time it is run, giving the teacher a growing bank of varied exercises from a single design. Combined with the audience tailoring described in Paper 2, the same pedagogy can simultaneously produce content at different reading levels, in different languages, or with different cultural reference points, all from one set of parameters.

We focus more on the trade-off between non-deterministic and deterministic transformations in the remaining papers. Paper 2 discusses non-determinism in much more detail as we consider tailoring pedagogies to specific audiences, even positing that certain elements that always require the same output (such as topic titles or formal definitions) are set deterministically, that is, without being generated by an LLM.

As Papers 3 and 4 will show, the visual and temporal layers of the framework are overwhelmingly deterministic. This distinction between deterministic and non-deterministic operations has significant implications for cost, reliability, and the degree of review required at each layer. Where a deterministic method can achieve the same result as a non-deterministic one, the deterministic method is almost always preferable: it is cheaper, faster, more predictable, and requires no review.

Appendix C: A content item as a lesson plan

The pedagogies in Section 5 are deliberately simple, containing three elements, with two parameters each. They demonstrate the framework's grammar. This appendix demonstrates that the same grammar scales to structures of real professional complexity, using a six-element lesson plan as the example.

A Rosenshine-style lesson plan as a pedagogy

Consider a secondary science teacher who wants to encode a lesson structure that reflects their school's expectation of Rosenshine-style instruction, adapted to their own teaching personality. The teacher's informal description might be: *start with quick-fire recall from last lesson, introduce the new concept with an analogy, work through an example showing common mistakes, give scaffolded practice that gradually removes support, set independent questions from different angles, and finish with a quick check that they've got it.*

Expressed in the framework, this becomes a six-element pedagogy with richer parameters than the three-element examples above. Note that the parameter structure here differs from the earlier examples. Where the three-element pedagogies used a simple rule-and-word-count pairing, this lesson plan uses a rule-and-constraints architecture, with the second parameter carrying richer, more specific requirements. This is one of the flexibilities of the system: because parameters are expressed in natural language, they can take whatever form best suits the pedagogy.

Pedagogy D: "Rosenshine lesson plan"

Element	Parameter 1: Rule	Parameter 2: Constraints
[Starter Retrieval]	Three quick-fire questions testing prior knowledge from the previous lesson. Questions must require recall, not recognition. Informal, slightly playful tone.	Must reference content from **[Independent Practice]** of the *previous* content item in the series.
[Key Concept]	Clear explanation of the single most important new idea. Must include an analogy from everyday life. Teacher's characteristic humour encouraged.	Maximum 60 words. Must use language accessible to the lowest reading age in the target class.
[Worked Example]	A fully worked-through example demonstrating the concept from **[Key Concept]**, showing every step. Must include common errors and why they are wrong.	Must not skip steps. Must flag where students typically go wrong.
[Guided Practice]	Two to three scaffolded questions that the learner attempts with support. First question closely mirrors **[Worked Example]**; second removes one scaffold; third requires independent application.	Difficulty must increase across questions. Part (c) must not be solvable by copying **[Worked Example]**.
[Independent Practice]	Three questions for the learner to attempt without support. Must cover the concept from different angles. One question must require the learner to combine this lesson's concept with prior knowledge.	Must not repeat **[Guided Practice]**'s structure. At least one question must be novel in form.
[Exit Ticket]	A single question that tests whether the learner has understood **[Key Concept]**.	Must require application, not recall. Must be different in form from all previous elements.

	Must be answerable in under 30 seconds.	

Several features of this pedagogy are worth noting. The inter-element rules are far more extensive than in the three-element examples with almost-every element referencing at least one other, encoding a deliberate pedagogical arc from prior-knowledge activation through to independent application, consistent with the scaffolding model described by Wood, Bruner, and Ross (1976). The teacher's personality is also part of the pedagogy, with tone parameters such as "slightly playful" or "teacher's characteristic humour encouraged" ensure the generated content sounds like this teacher's lesson, not a generic template. And the pedagogy is entirely subject-agnostic. Nothing in the parameters is specific to science, for example, and the same structure could encode a history, mathematics, or languages lesson.

The practical consequence is that a school or department could maintain its own library of encoded pedagogies as part of its scheme of work: a Rosenshine-style lesson plan for new content, a retrieval-practice structure for revision, an inquiry-based design for practical work. Each would be defined once, applied across the curriculum, and refined over time as teachers share what works and adjust what does not. As we discuss in Paper 2, which introduces an architecture for tailoring existing pedagogies to the needs and preferences of that teacher and hence their students.