STIAEM

Slow Tech Innovation

AI, Ethics and New Teaching Methods

Syntactic Structure of Effective Prompts: Role, Context, Task, and Constraints

Co-funded by the European Union

2025-1-IT02-KA220-SCH-000365213

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them... Project reference code: 2025-1-IT02-KA220-SCH-000365213

Università di Foggia

Studi Umanistici, Lettere, Beni Culturali, Scienze della Formazione

İSTANBUL - ESENLER İLÇE MİLLİ EĞİTİM MÜDÜRLÜĞÜ

1474

STIAEM
Slow Tech Innovation
AI, Ethics and
New Teaching Methods

Co-funded by
the European Union

2025-1-IT02-KA220-SCH-000365213

eduOpen

CC BY-NC-SA 4.0 International

This license requires that reusers give credit to the creator. It allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, for noncommercial purposes only. If others modify or adapt the material, they must license the modified material under identical terms.

Università di Foggia
Studi Umanistici, Lettere, Beni Culturali,
Scienze della Formazione

OER 4.1.1: Syntactic Structure of Effective Prompts: Role, Context, Task, and Constraints
Reference Module: Module 4 - Prompt Engineering for Inclusive Education
Related Video-Pill: Video 4.1 - Anatomy of the Structured Prompt
Authoring partner: Directorate of Secondary Education of Dodecanese - Greece
Topic / Thematic Focus: Syntactic prompt structure, role-context-task framework and operational constraints

Introduction and Learning Objectives

This resource accompanies, together with OER 4.1.2, video-pill 4.1 «Anatomy of the Structured Prompt» and constitutes its operational in-depth analysis. Its subject is the syntax of the effective prompt: the minimal structure - role, context, task and constraints - capable of turning a generic request into a controllable and verifiable teaching artefact.
Why devote an entire text to the form of a request? Because between a badly posed question and the same question well constructed lies the difference between an unusable output and a material that the teacher can adopt in class. In terms of references, the text is placed in the wake of DigComp 3.0 (Cosgrove & Cachia, 2025), the European Framework of digital competences in which AI-related competence runs, explicitly or implicitly, through all twenty-one competences, and of the «AI pedagogy» pillar of the UNESCO Framework for teachers (Miao & Cukurova, 2024), which invites us to build the «pedagogical why» of every technological adoption. The underlying ethical perspective remains the Slow Tech one dear to the project: to empower human thought without uncritically delegating decisions to it.
By the end of the reading, the teacher will be able to:

- recognise the four pillars of the structured prompt - role, context, task and constraints - and the specific function of each;
- turn a generic request into a complete prompt, verifying that none of the four pillars remains implicit;
- reflect critically on the role, which orients the machine's discourse but does not certify the truth of what it produces.

1. Why two identical requests give different results

Let us imagine two teachers who open the same artificial intelligence assistant and type almost identical words: «set me some exercises in logical analysis». The first receives five disconnected sentences, with no indication of the level of difficulty; the second, after a few more minutes of writing, receives an exercise consistent with the syllabus, calibrated on the class and accompanied by solutions. Yet they were using the same tool. What made the difference? Not the power of the model, but the structure of the instruction.
The paradigm of contemporary language models explains why. In the foundational survey by Liu and colleagues, prompting is described as the way in which downstream tasks are reformulated to resemble those encountered during training:
«In this paradigm, instead of adapting pre-trained LMs to downstream tasks via objective engineering, downstream tasks are reformulated to look more like those solved during the original LM training with the help of a textual prompt.» (Liu et al., 2021, 1)
The consequence is twofold. On the one hand, the model is not retrained for every new task: the formulation of the request is enough. On the other hand, precisely for this reason, the quality of the request becomes decisive, to the point that the authors speak explicitly of the

need for prompt engineering, that is, of the search for the most suitable prompt so that the model can solve the task in question (Liu et al., 2021, 1).
The review by Sahoo and colleagues confirms the point with a formula that is particularly clear for those who teach: prompting exploits task-specific instructions without modifying the model's internal parameters. Translated into the classroom: there is no need to «programme» the machine, one needs to tell it precisely what to do. And precision, as we shall see, is a professional competence of the teacher, not a gift of the tool.

2. The role: who plays the machine

The first pillar is the role, or persona. Assigning a role means telling the model who it must play while responding: a Socratic tutor, a strict examiner, a proofreader of drafts, a populariser for fourteen-year-olds. It is not a theatrical trick: the role orients the register, the tone and even the type of content that the model chooses to highlight.
It is one of the patterns catalogued by White and colleagues in their catalogue of prompt patterns, which describes its intent as follows:
«The intent of this pattern is to give the LLM a “persona” that helps it select what types of output to generate and what details to focus on.» (White et al., 2023, §III.E, Persona Pattern)
The same idea appears in the review by Schulhoff and colleagues under the name of role prompting: the model is assigned a specific role already within the text of the prompt, and this directs the output towards the expected style and content (Schulhoff et al., 2024). In classroom practice the role is valuable when a certain type of gaze is needed: asking for the role of «tutor who never gives the solution, but asks guiding questions» is different from asking for the role of «examiner who lists the errors». The first builds a scaffolding of thought, the second a diagnosis.
A warning, however, is necessary from the outset: the role orients the discourse, it does not make it true. A machine that plays the role of the expert teacher will continue to produce plausible, though not guaranteed, statements. Assuming the tone of authority is not the same thing as possessing it.

3. The context: for whom the machine works

The second pillar is the context. If the role says who the machine plays, the context says for whom it works. Without a briefing, the model does not know the class, the syllabus, the prerequisites, the recurring errors: and so, to fill the void, it imagines. It usually imagines an average student, competent and generic, who very probably does not exist.
It is worth fixing here what we mean, exactly, by prompt. The definition in the review by Schulhoff and colleagues is essential:
«A prompt is an input to a Generative AI model, that is used to guide its output.» (Schulhoff et al., 2024, 1.1)
An input that guides the output, then: and the guidance is all the more precise the more context it receives. What must be stated? At least four elements: who the students are (age, class, pathway); what knowledge they already possess; what difficulties recur; in what teaching context the activity is set (a test, a catch-up session, an in-depth study). Adding a line of context costs a few seconds and radically changes the relevance of the result.
On the technical level, context is what makes it possible to «attach» a pre-trained model to a specific task without retraining it:

«Rather than updating the model parameters, prompts allow seamless integration of pre-trained models into downstream tasks by eliciting desired model behaviors solely based on the given prompt.» (Sahoo et al., 2024, Abstract)
It is the meeting point between the machine and the class: the same technology responds to different needs if the context orients it. And since context is always a choice of the teacher, it is also an act of design, not a negligible detail.

4. The task: the operative verb

The third pillar is the task, the action we ask of the machine. Here the most widespread error is the vagueness of the verb. «Tell me about...» or «help me with...» leave the model an enormous margin; operative verbs such as classify, translate, rewrite, refute, synthesise or generate variants instead delimit precisely what we expect.
White and colleagues recall that prompts are not merely requests for text, but instructions intended to control the output:
«Prompts are instructions given to an LLM to enforce rules, automate processes, and ensure specific qualities (and quantities) of generated output.» (White et al., 2023, Abstract)
That «enforce rules» - imposing rules - is the clue to how the task must be formulated: an instruction, not a wish. A good task says what to produce, in what form and with what level of detail. If the verb is «classify», the expected output is a reasoned classification; if it is «synthesise», the output is a summary with declared criteria. The verb, in other words, prefigures the structure of the result.
The definition in the survey by Sahoo and colleagues insists on the same point, speaking of task-specific instructions:
«This approach leverages task-specific instructions, known as prompts, to enhance model efficacy without modifying the core model parameters.» (Sahoo et al., 2024, Abstract)
For the teacher, the practical consequence is a small discipline: writing the task as if it had to be entrusted to a colleague who does not know the class. If a human colleague would not understand what we are asking of them, it is likely that the machine will not understand it either.

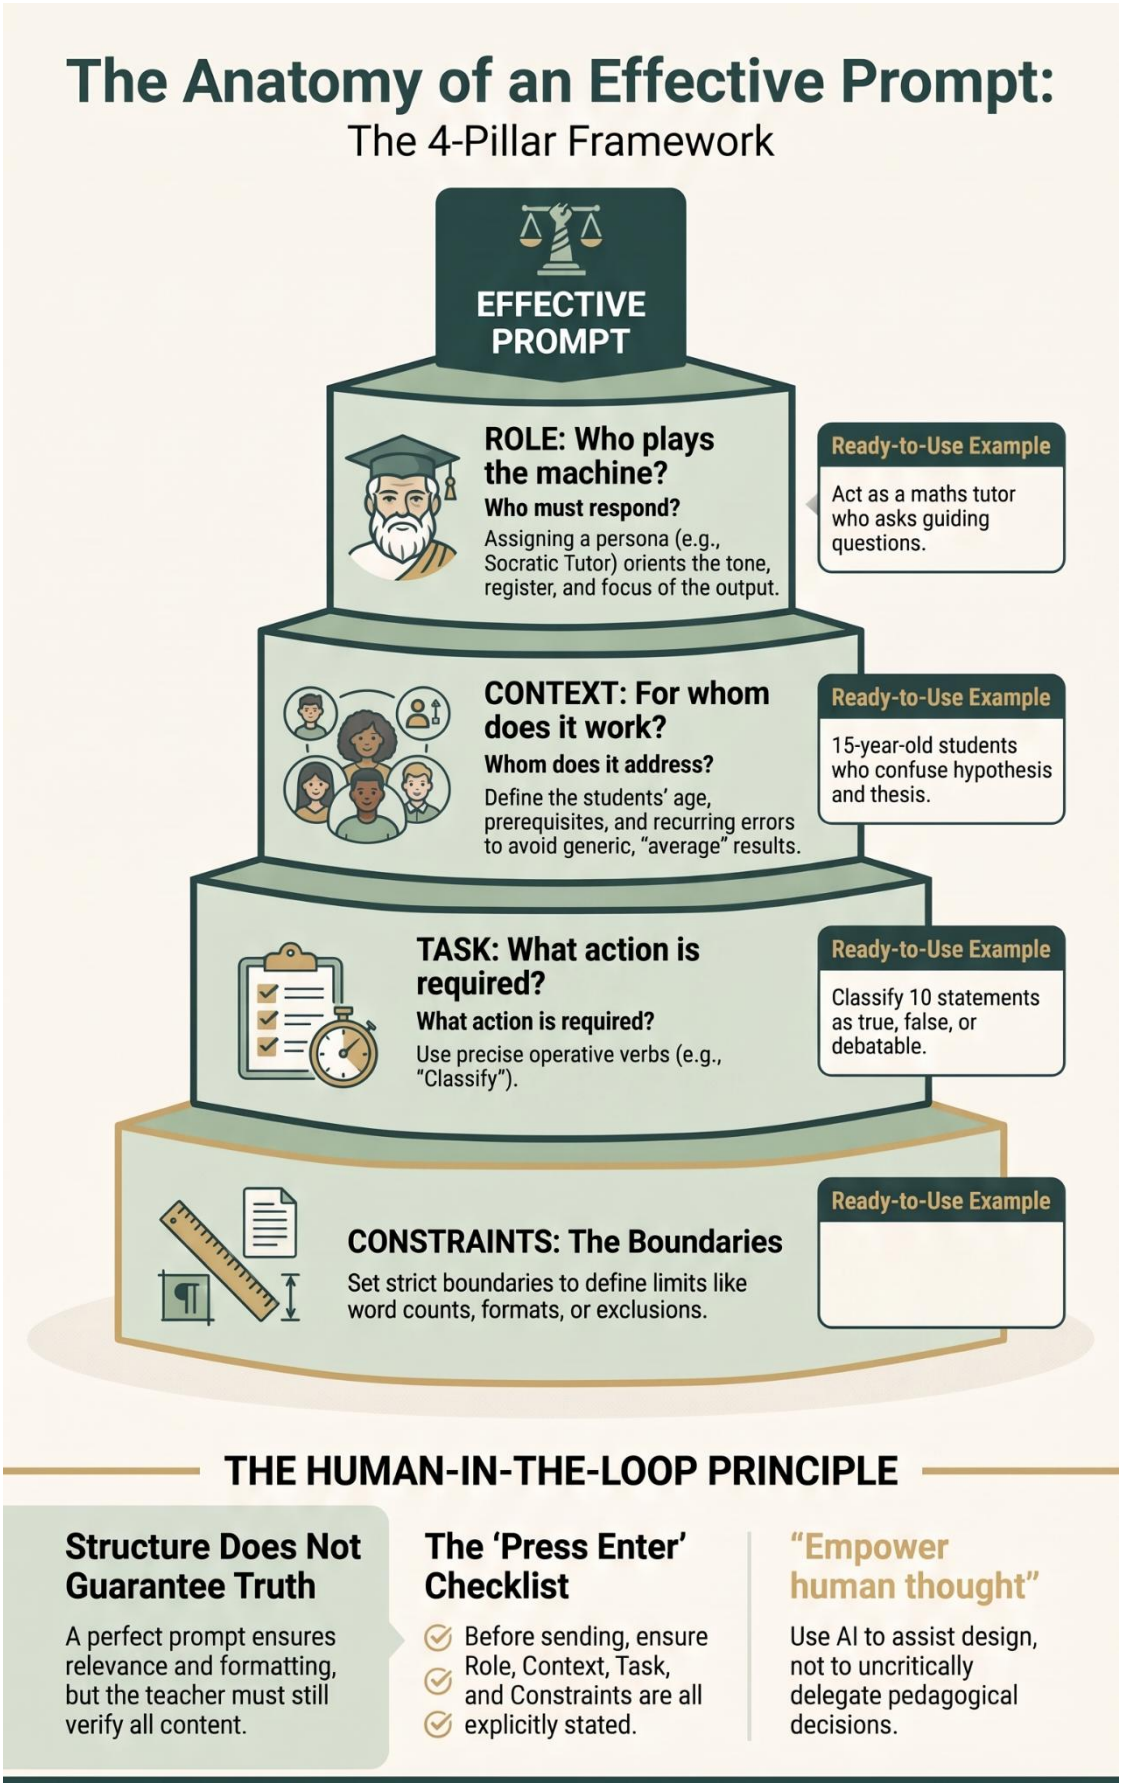
5. Constraints and caution: from control to verification

The fourth pillar gathers constraints, that is, the boundaries within which the output must remain: number of words, format (list, table, continuous text), linguistic register, permitted vocabulary, quantity of examples. Constraints reduce the space in which the model can wander and make the result more predictable and easier to verify.
White and colleagues summarise this function well with an effective image: the prompt is also a form of programming, which customises the output and the interaction with the model. Imposing constraints is, in this sense, a light form of linguistic programming.
«Prompts are also a form of programming that can customize the outputs and interactions with an LLM.» (White et al., 2023, Abstract)
But precisely here the limit must be reiterated, which pill 4.1 repeats several times: structure increases control, not truth. An impeccable prompt does not guarantee that the content produced is correct; it guarantees that it is pertinent, formatted as we wanted and compliant with the constraints. Correctness remains the object of human verification. This is the reason why the teacher does not simply «use» the output: they read it, compare it with the sources, correct it and decide whether to adopt it.

We can summarise the four pillars in a simple rule, useful to bring into the classroom: before pressing enter, check that the prompt declares who plays the machine (role), for whom it works (context), what it must do (task) and within what boundaries (constraints). If one of the four answers is missing, the prompt is incomplete.

Pillar	Function	Guiding question	Ready-to-use example
Role	Defines who plays the model: register, tone, type of *feedback*	«Who must respond?»	«Act as a maths tutor who never gives the solution, but asks guiding questions.»
Context	Declares students, prerequisites and recurring errors	«Whom does it address, and with what knowledge?»	«Third year, 15-year-old students; they confuse hypothesis and thesis; they use colloquial vocabulary.»
Task	Fixes the action with a precise operative verb	«What action is required?»	«Classify ten statements as true, false or debatable, giving a one-line justification.»
Constraints	Limit words, format, register and quantity	«Within what boundaries?»	«Maximum 120 words; two-column table; vocabulary suitable for the second year of upper secondary school.»

Visual Scaffolding



Conclusions and Metacognitive Reflection Prompts

This resource has presented the syntax of the effective prompt as a grammar of four pillars: role, context, task and constraints. The role defines who plays the machine and orients its register; the context declares for whom it works and determines its relevance; the task fixes the action with an operative verb; the constraints delimit its form and quantity. Together, the four pillars turn a generic request into a controllable teaching artefact.

The common thread is a principle of responsibility: the structure of the prompt is a professional choice of the teacher, not an automatism. The machine is as powerful as the question put to it; and the question, in the classroom, is always posed by the teacher. At the same time, no structure guarantees the truth of the output: verification of the content remains human.

To transfer these core ideas into one's own practice, the following reflection prompts are proposed:

Take a prompt you habitually use: which of the four pillars remains implicit, and what changes in the output when you make it explicit?

How would you explain to the class the difference between a role that «orients the discourse» and an authority that «certifies the truth»?

At what point in your design is it best to write the prompt: before defining the objective or after clarifying it precisely?

Technical Metadata Sheet

Metadata Field	Specification / Parameter Value
Identifier	OER 4.1.1
Title	Syntactic Structure of Effective Prompts: Role, Context, Task, and Constraints
Creator	Directorate of Secondary Education of Dodecanese - Greece
Description	The resource presents the syntax of the effective prompt: the minimal structure of role, context, task and constraints that turns a generic request into a controllable and verifiable teaching artefact. Written for non-technical teachers, it analyses the specific function of each pillar and shows how to verify that none remains implicit. The critical focus is on the role, which orients the machine's discourse but does not certify the truth of what it produces.
Keywords	prompt engineering; structured prompt; role; context; task; constraints; prompt syntax; verification
Language	EN (English)
Affiliation	Directorate of Secondary Education of Dodecanese - Greece
Contributor	STIAEM Transnational Scientific Board (peer review, Rhodes validation protocol)
Media Format	Text; Infographics
Genre	Training Manual; Educational Guide
URL	
Citation	Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition). Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149 Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2021). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv. https://arxiv.org/abs/2107.13586 Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. UNESCO. https://doi.org/10.54675/ZJTE2084

	Sahoo, P., Singh, A. K., Saha, S., Jain, V., Mondal, S., & Chadha, A. (2024). A systematic survey of prompt engineering in large language models: Techniques and applications. arXiv. https://arxiv.org/abs/2402.07927 Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., Li, Y., Gupta, A., Han, H., Schulhoff, S., Dulepet, P. S., Vidyadhara, S., Kieley, J., Li, A., Wang, X., ... & Resnik, P. (2024). The prompt report: A systematic survey of prompt engineering techniques. arXiv. https://arxiv.org/abs/2406.06608 White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with ChatGPT. arXiv. https://arxiv.org/abs/2302.11382
Educational Level	Secondary School
Associated Module / MOOC Video Lesson	Module 4 - Prompt Engineering for Inclusive Education / Video 4.1 - Anatomy of the Structured Prompt
Learning Objectives	recognise the four pillars of the structured prompt - role, context, task and constraints - and the specific function of each; turn a generic request into a complete prompt, verifying that none of the four pillars remains implicit; reflect critically on the role, which orients the machine's discourse but does not certify the truth of what it produces.
Rights	Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)
Date Available	14/09/2026
Date Modified	

References and Webography

Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition). Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149

Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2021). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv. https://arxiv.org/abs/2107.13586

Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. UNESCO. https://doi.org/10.54675/ZJTE2084

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

Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., Li, Y., Gupta, A., Han, H., Schulhoff, S., Dulepet, P. S., Vidyadhara, S., Kieley, J., Li, A., Wang, X., ... & Resnik, P. (2024). The prompt report: A systematic survey of prompt engineering techniques. arXiv. https://arxiv.org/abs/2406.06608

White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with ChatGPT. arXiv. https://arxiv.org/abs/2302.11382