DELIVERABLE 5.1:

Project Acronym: HITS
TITLE: High Intelligent Tutoring System for Academic Success
Project Number: **000152340**

Defining similarity and styles in achieving learning objectives and in tutoring

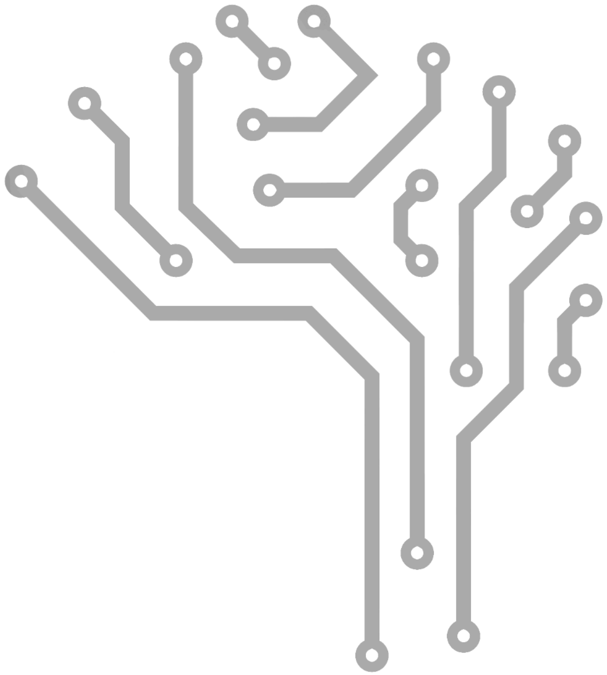

Description: Report on similarity and styles in achieving learning objectives and in tutoring

Lead Party for deliverable: UNIVERSITA DEGLI STUDI DI CAMERINO
Document type: Report
Due date of deliverable: 28/02/2026
Dissemination level: *PU*

Authors:

Andrea Marconi
Pietro Borghi
Sašo Karakatič
Marta Agueda Maroñas
Erika Vaiginienė
Paulina Spanu

HITS PARTNERS:

University of Camerino

Centro Studi Pluriversum

University of Maribor

University of Valladolid

Vilnius University

Universitatea Politehnica din Bucureşti

Summary

T.5.1 Defining similarity and styles in achieving learning objectives and in tutoring

1. Purpose of the Deliverable

This deliverable sets out the analytical rationale that informed the design of the HITS Chatbot within Work Package 5. Its purpose is not simply to describe a technological solution, but to reconstruct the reasoning process through which that solution was conceived. In this respect, T.5.1 occupies a foundational position within WP5: it explains why the chatbot took its present form, which educational assumptions guided its development, and how its architecture derives from prior analysis rather than from generic technological design.

The central claim of the deliverable is that the HITS Chatbot was not conceived as a general-purpose AI assistant. Its adaptive, modular, Retrieval-Augmented Generation (RAG)-based structure emerged from an earlier effort to understand how students differ in the ways they approach learning, manage academic tasks, and respond to institutional support. The system was therefore designed on the basis of analysis, not added afterwards as a neutral interface. In other words, within the logic of WP5, architecture followed analysis. This point was explicitly recognised by the partner validation process as one of the defining strengths of the deliverable.

The work carried out in T.5.1 focused on identifying recurring patterns of similarity across three interrelated dimensions of academic behaviour: cognitive, behavioural, and motivational. The cognitive dimension concerns how students engage with knowledge, materials, and learning tasks. The behavioural dimension concerns temporal and procedural patterns of engagement, including interaction rhythms, responsiveness, and proximity to deadlines. The motivational dimension captures recurrent signals of uncertainty, confidence, disengagement, or persistence. Considered together, these dimensions made it possible to outline a set of learner profiles and, from them, a corresponding set of differentiated tutoring styles. The architecture of the chatbot was then aligned with these patterns so that personalization would not remain superficial or cosmetic, but would become structurally embedded in the system itself.

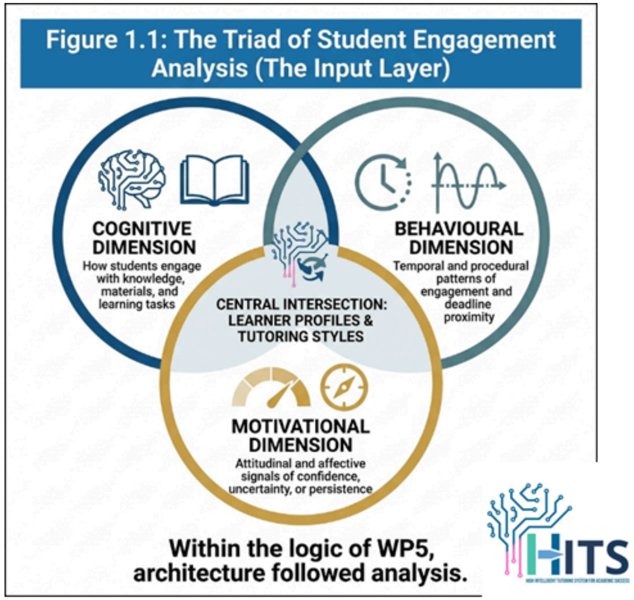
Figure 1.1: The Triad of Student Engagement Analysis (The Input Layer)

This analytical grounding had direct implications for system design. The decision to adopt a RAG architecture, for example, was linked to the need to ensure reliability, auditability, and trust across different learner profiles and institutional contexts. Likewise, the use of adaptive nudging and the inclusion of routing mechanisms toward human tutors were not introduced as generic support features, but as responses to observed differences in student engagement, self-regulation, and support needs. In this sense, the deliverable makes explicit the pedagogical and behavioural logic that underpins the system's technical configuration.

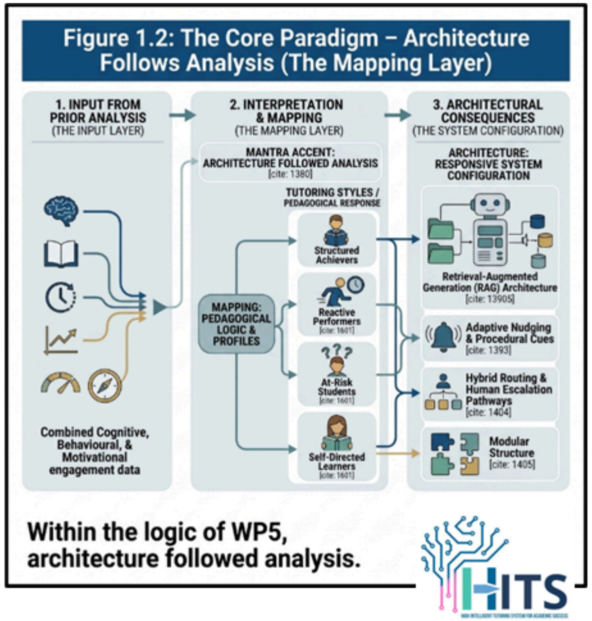

Figure 1.2: The Core Paradigm – Architecture Follows Analysis (The Mapping Layer)

T.5.1 therefore serves as a conceptual and methodological bridge within WP5. It provides the analytical basis for T.5.3, which formalises similarity patterns and clustering logics; it informs T.5.4, where differentiated nudging strategies are developed; and it supports T.5.7, where the predictive AI model is implemented. By clarifying the relation between similarity analysis, learner profiling, tutoring logic, and system architecture, the deliverable reinforces the internal coherence of the work package and strengthens the evidence-based rationale of the AI Learning Library as a whole. This bridging function was consistently acknowledged in the partner validation questionnaires as one of the document's key merits.

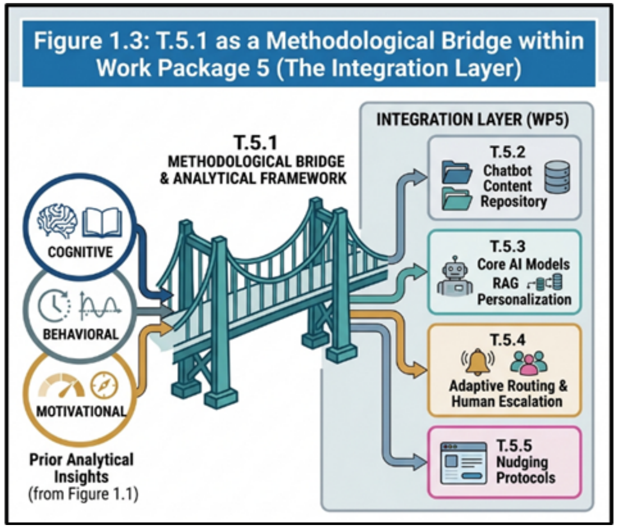

Figure 1.3: T.5.1 as a Methodological Bridge within Work Package 5 (The Integration Layer)

At the same time, the analytical categories introduced in this deliverable should not be interpreted as fixed labels attached to individual students. The learner profiles derived from similarity analysis are intended as provisional and revisable interpretive constructs. Their purpose is to identify recurring patterns in academic engagement and support needs, not to reduce learners to static identities. This distinction is essential both methodologically and ethically, since the wider HITS framework explicitly seeks to avoid deterministic classification and to preserve the possibility of revision as new evidence accumulates over time.

The sections that follow develop this argument progressively. Section 2 presents the theoretical and technical foundations of the chatbot, bringing together behavioural theory and system reliability. Section 3 defines the notion of similarity in academic behaviour and explains the three dimensions on which the analysis is based. Section 4 introduces the learner profiles derived from this analysis, while Section 5 translates those profiles into tutoring styles. Section 6 then explains how these insights informed the structure of the chatbot itself. Finally, Sections 7 and 8 address validation, ethical safeguards, and the broader European relevance of the framework. Through this progression, the deliverable demonstrates that the HITS Chatbot should be understood not as an isolated digital tool, but as the operational expression of a prior pedagogical and analytical model.

2. Theoretical and Technical Framework

The design of the HITS Chatbot rests on a dual foundation. On the one hand, it is informed by a behavioural and pedagogical framework that explains how students engage with academic tasks, where they encounter obstacles, and how support can be structured in a meaningful way. On the other hand, it is implemented through a technical architecture intended to ensure reliability, traceability, and institutional trust. These two dimensions are not parallel or optional components of the system. They are constitutively linked. A theoretically informed intervention that lacks technical robustness risks producing unreliable or misleading guidance. Conversely, a technically sophisticated system without pedagogical grounding risks becoming efficient but educationally irrelevant.

Figure 2.1: The Interlocking Dual Foundation Matrix

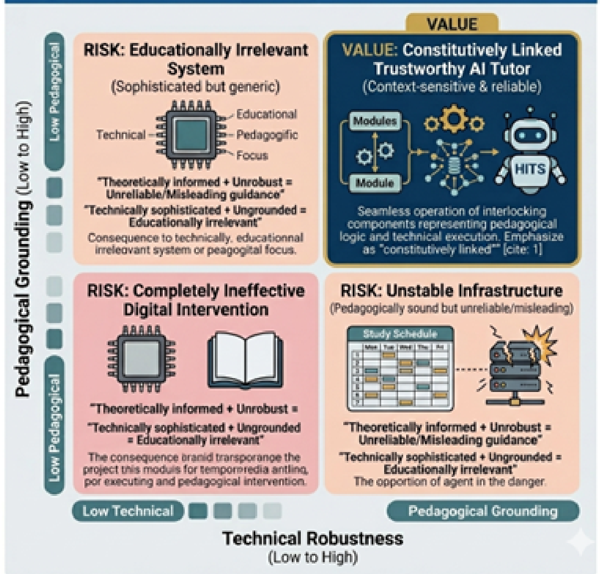

For this reason, the development of the HITS Chatbot did not begin with implementation choices alone. It began with an analysis of how academic behaviour unfolds and where breakdowns typically occur. The project first sought to understand the conditions under which students encounter difficulties in sustaining continuity, interpreting institutional requirements, or mobilising support. Only on that basis was a technical solution designed. The theoretical framework, centred primarily on the COM-B model and complemented by principles of digital nudging and choice architecture, helped to clarify the kinds of barriers that students may face. The technical framework, developed through a Retrieval-Augmented Generation (RAG) architecture and a controlled data ingestion process, was then shaped so as to respond to those barriers in a reliable and institutionally grounded way.

The present section develops this dual foundation in two steps. It first situates the HITS Chatbot within the pedagogical and institutional context of contemporary higher education, showing why behavioural interpretation is necessary before technological intervention can become meaningful. It then explains why the project adopted a RAG-based architecture and how this choice contributes to trust, accuracy, and the reduction of unsupported outputs. In this way, the section makes clear that the HITS Chatbot is neither simply a behavioural support system nor simply a technical information-retrieval tool. It is the product of their deliberate integration.

2.1 Context and Pedagogical Rationale

The HITS Chatbot must be understood against the wider context of contemporary higher education, in which universities are increasingly expected to promote academic continuity, reduce dropout risk, and provide timely support to students with highly differentiated needs. Yet these expectations coexist with a structural limitation: traditional tutoring systems, while pedagogically valuable, are often constrained by time, staffing, and limited capacity for continuous follow-up. Students' support needs, however, rarely arise in stable or predictable ways. They are often intermittent, context-dependent, and closely tied to specific moments of academic pressure, uncertainty, or disorientation.

This structural mismatch between available support and actual patterns of need provided one of the key starting points for the development of HITS. The chatbot was not conceived as a replacement for tutors, but as a complementary support layer capable of offering continuity where human systems cannot always guarantee immediacy or sustained availability. Its pedagogical value lies precisely in this intermediate position: it operates alongside existing institutional support, extending access to guidance while preserving the centrality of human judgement in more complex or sensitive situations.

The main theoretical lens used to interpret these support needs is the COM-B model, which understands behaviour as emerging from the interaction of **capability**, **opportunity**, and **motivation**. Within academic settings, this model makes it possible to move beyond a narrow view of success as a matter of knowledge alone. Students may know what is required and still fail to act effectively, not because of lack of intelligence or commitment, but because one of the necessary conditions for sustained academic behaviour is compromised. A capability-related difficulty may concern ineffective study strategies or limited understanding of how to approach tasks. An opportunity-related difficulty may arise when administrative rules, digital infrastructures, or institutional pathways are difficult to interpret or access. A motivation-related difficulty may emerge when a student understands what should be done but struggles to maintain engagement over time. The COM-B framework therefore offers a parsimonious but powerful way of interpreting educational barriers as multidimensional rather than purely cognitive.

This perspective also makes clear why a one-size-fits-all digital tool would be insufficient. If students encounter different kinds of barriers, then support must also vary in timing, content, and mode of delivery. This is where principles of digital nudging and choice architecture become relevant. These principles do not imply manipulation or behavioural control. Rather, they draw attention to the educational significance of how information is structured, when it is delivered, and how much friction students must overcome in order to act upon it. In academic contexts, reminders, prompts, suggested next steps, or simplified access to relevant information may either support or hinder engagement depending on their clarity, timing, and contextual appropriateness. The HITS Chatbot was therefore designed not merely to provide answers, but to do so in ways that respond to identifiable patterns of academic behaviour and differentiated support needs.

Figure 2.2: The COM-B Academic Barrier Mapping Engine

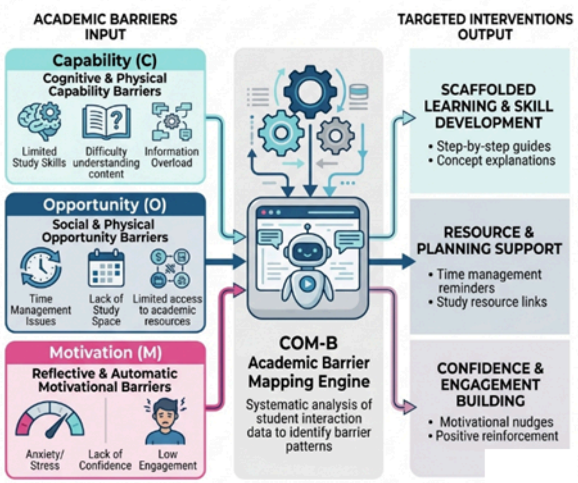

Illustrating how the COM-B framework is operationalized within the HITS system to identify specific academic challenges and deliver personalized interventions.

At the same time, behavioural relevance alone cannot sustain educational trust. Even the most pedagogically well-calibrated support mechanism loses value if the information it provides is inaccurate, unverifiable, or institutionally unreliable. For this reason, the project treats pedagogical interpretation and technical trustworthiness as inseparable. The behavioural model identifies what kind of support is needed; the technical architecture must ensure that the support delivered is correct, traceable, and anchored in legitimate institutional sources. This requirement leads directly to the architectural choice discussed in the following subsection.

2.2 Trust, Accuracy, and the RAG Architecture

In academic environments, reliability is not an added value; it is a basic requirement. Incorrect information regarding deadlines, prerequisites, assessment rules, or procedural obligations can have immediate and serious consequences for students. Any digital system intended to support academic decision-making must therefore be designed with a clear awareness of the risks associated with unsupported or fabricated responses. This requirement is especially important in the case of large language models, whose fluency may create the impression of authority even when the content produced is uncertain or incorrect. The well-known problem of hallucination is, in this context, not a marginal technical issue but an educational and institutional one.

The HITS Chatbot addresses this challenge through a Retrieval-Augmented Generation architecture. In simple terms, this means that language generation is constrained by retrieved documentary evidence rather than relying exclusively on the model’s general pre-trained knowledge. Instead of producing answers from an open-ended statistical model of language alone, the system first searches within a curated institutional corpus and then formulates its response on the basis of the relevant material retrieved from that corpus. This shift—from open generation to grounded synthesis—has important implications. It reduces the risk of unsupported claims, improves auditability, and anchors responses in documentation that can be checked, updated, and governed by the institution.

The corpus used by the system is constructed through a controlled ingestion process. Official institutional documents, including regulations, course syllabi, academic handbooks, and selected web content, are collected, cleaned, standardised, and transformed into a searchable knowledge base. The documents are segmented into manageable textual units and indexed in a way that supports semantic retrieval. Rather than relying on exact word matching alone, the system represents both documents and queries through computational structures that make it possible to retrieve passages on the basis of contextual similarity in meaning. In this sense, retrieval becomes sensitive not only to vocabulary, but to underlying intent. This is particularly important in student support contexts, where the same need may be expressed in highly different ways.

Figure 2.3: The Epistemic Trust Pipeline (The RAG Architecture)

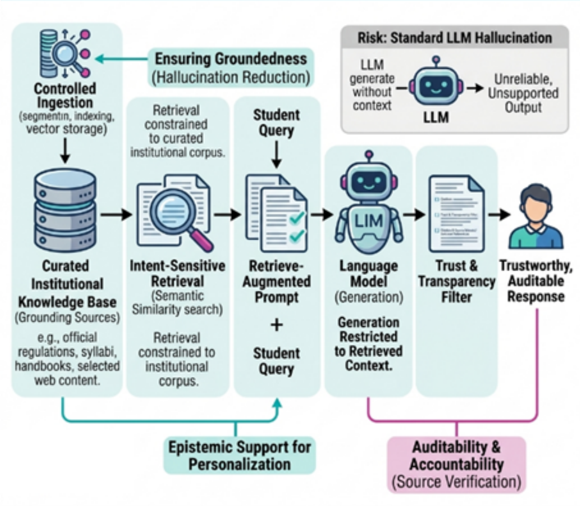

The result is not a live search across the open web, but a controlled institutional knowledge environment. The chatbot operates within defined informational boundaries. When the relevant information is present in the indexed corpus, it can be retrieved and used to generate a grounded response. When it is absent, the system is expected to acknowledge that limitation rather than invent an answer. This distinction is central. Trust is not achieved by giving an answer at all costs; it is achieved by ensuring that answers remain tied to what the institutional corpus can legitimately support. In this way, the RAG architecture becomes not merely a technical solution, but the structural means through which academic reliability is made compatible with adaptive support.

This architectural choice also reflects a broader pedagogical commitment. Student engagement is influenced not only by the availability of support, but by confidence in its legitimacy. A support system that is quick but unreliable will ultimately undermine motivation rather than sustain it. By contrast, a system that can connect its guidance to verified institutional documentation contributes to a more credible support environment. The significance of the RAG architecture therefore exceeds the reduction of hallucinations in a narrow technical sense. It establishes the epistemic conditions under which personalised support may be considered educationally and institutionally defensible. Without such grounding, adaptive tutoring would risk becoming arbitrary. With it, personalization becomes compatible with trust, accountability, and pedagogical care.

Taken together, the behavioural and technical dimensions described in this section define the framework within which the HITS Chatbot was conceived. Behavioural theory identifies where and why support is needed. Technical architecture determines how that support can be delivered responsibly. Their integration is what allows the system to function as a context-sensitive, reliable, and pedagogically grounded academic support instrument rather than as a generic AI interface. It is from this integrated framework that the following analysis of similarity, profiles, tutoring styles, and architectural consequences becomes intelligible.

3. Defining Similarity in Academic Behaviour

Understanding student diversity was a necessary precondition for the design of the HITS Chatbot. If support is to be meaningfully personalised, it cannot rest on generic assumptions about student needs. It must be grounded in observable patterns of academic engagement. For this reason, T.5.1 introduced the notion of **similarity in academic behaviour** as an analytical instrument through which recurring configurations of need, difficulty, and response could be identified. The purpose of this construct is not to classify students once and for all, nor to reduce them to static categories. It is to provide a structured way of interpreting how different learners engage with academic tasks and institutional support over time.

In this context, similarity refers to the extent to which students display comparable tendencies in the way they study, interact with academic structures, and respond to guidance. The emphasis is therefore placed not primarily on final outcomes, but on the processes that precede them. Students may arrive at similar grades through very different forms of engagement, just as comparable difficulties may have different cognitive, behavioural, or motivational roots. A useful support system must be able to recognise such variation. Similarity analysis serves precisely this purpose: it identifies recurring patterns without transforming them into rigid labels. The profiles derived from it remain provisional and revisable, open to recalibration as new evidence accumulates.

Figure 3.1: The Process-Over-Outcome Evaluation Model

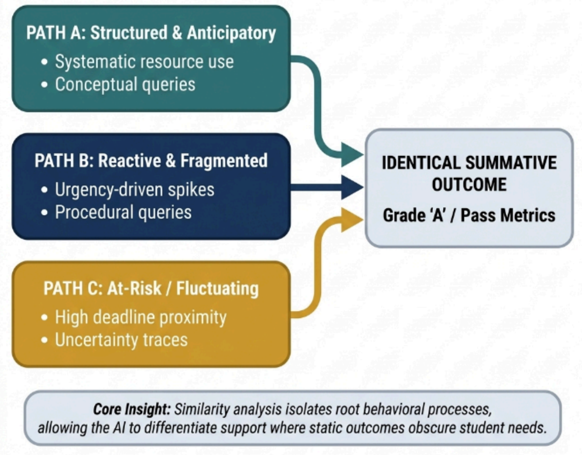

The analytical model adopted in T.5.1 is organised around three interrelated dimensions: **cognitive**, **behavioural**, and **motivational** similarity. These dimensions were selected because, taken together, they offer a sufficiently rich but still manageable picture of academic engagement. The cognitive dimension concerns how students approach content, organise study, and formulate questions. The behavioural dimension concerns rhythms of action, including regularity, responsiveness, and temporal proximity to deadlines. The motivational dimension concerns the attitudinal and affective signals that emerge through language and interaction over time. None of these dimensions functions in isolation. Their educational relevance lies precisely in their interaction. Cognitive difficulty may become more consequential when accompanied by irregular behaviour; behavioural fluctuation may become more meaningful when combined with repeated signals of uncertainty or discouragement. It is through this intersection that recognisable patterns of academic engagement emerge.

At a methodological level, these dimensions are not inferred impressionistically. They are supported by a combination of content analysis, interaction indicators, and longitudinal observation. Student queries and resource-use patterns help identify thematic and strategic regularities in how learners engage with academic material. Temporal data illuminate rhythms of interaction, distribution of effort, and responsiveness to reminders. Linguistic and continuity signals provide indirect but useful indicators of motivational orientation. The technical implementation of these dimensions involves semantic representations, similarity measures, and clustering procedures; however, the present deliverable is concerned less with the mathematical formalism itself than with the educational logic that such procedures make possible. What matters here is that similarity is treated as a structured, evidence-based lens through which recurring support needs may be identified and interpreted.

The sections that follow develop the three dimensions in turn. Each is presented first in conceptual and pedagogical terms, before showing how it contributes to the wider analytical framework. Through this progression, similarity becomes the bridge between individual interaction traces and the more stable interpretive patterns from which learner profiles and tutoring styles can later be derived.

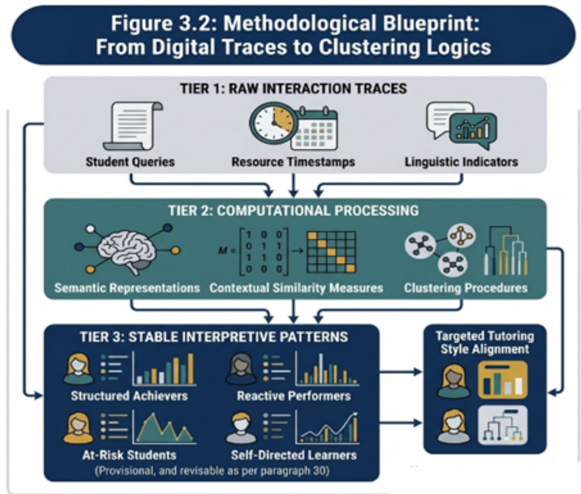

Figure 3.2: Methodological Blueprint: From Digital Traces to Clustering Logics

3.1 Cognitive Similarity

Cognitive similarity concerns recurring patterns in the ways students engage with academic content and organise their learning activity. It does not measure intelligence, ability, or academic worth. Rather, it seeks to capture how students habitually approach tasks, what kinds of resources they rely on, what forms of difficulty recur in their questions, and how they orient themselves in relation to course material. In this sense, the cognitive dimension is concerned with **modes of engagement** rather than with performance alone.

A first important aspect of cognitive similarity lies in students' study strategies. Some learners approach academic work in a relatively structured and anticipatory way. Their engagement with materials follows the progression of the course, and their questions tend to arise as part of an ongoing effort to consolidate understanding. Others study in a more reactive or fragmented way, concentrating their effort at points of urgency or in response to immediate assessment demands.

These differences are pedagogically significant because they affect not only what students know, but also how and when they are likely to benefit from support.

A second aspect concerns the use of didactic materials. Students differ in the extent to which they consult official documents, lecture notes, recommended readings, and supplementary resources. Some show sustained and systematic use of institutional material, while others rely more selectively on summaries, clarifications, or partial references. These patterns provide insight into how students build understanding and where informational friction may arise. The issue is not simply whether students access materials, but how that access is distributed, sequenced, and integrated into their study process.

A third indicator is the nature of the academic queries students formulate. This is especially relevant in the context of a chatbot-based system. Some questions are primarily conceptual: they indicate a wish to understand content, deepen comprehension, or connect ideas. Others are predominantly procedural: they concern deadlines, requirements, prerequisites, or organisational matters. Neither form is inherently more valuable than the other, but they point to different orientations toward academic work and, consequently, to different support needs. Conceptual queries may call for richer didactic scaffolding; procedural queries may signal uncertainty about institutional navigation or course organisation.

Taken together, these elements allow the project to speak of cognitive similarity not as a fixed cognitive type, but as a recurring pattern in the ways students engage with knowledge and academic tasks. At a technical level, such patterns can be supported by semantic analysis of student queries and by indicators derived from material-use behaviour. At an educational level, however, their importance lies in enabling the system to distinguish between qualitatively different forms of academic engagement. This distinction is essential if retrieval logic, tutoring style, and support timing are to remain pedagogically appropriate rather than merely responsive.

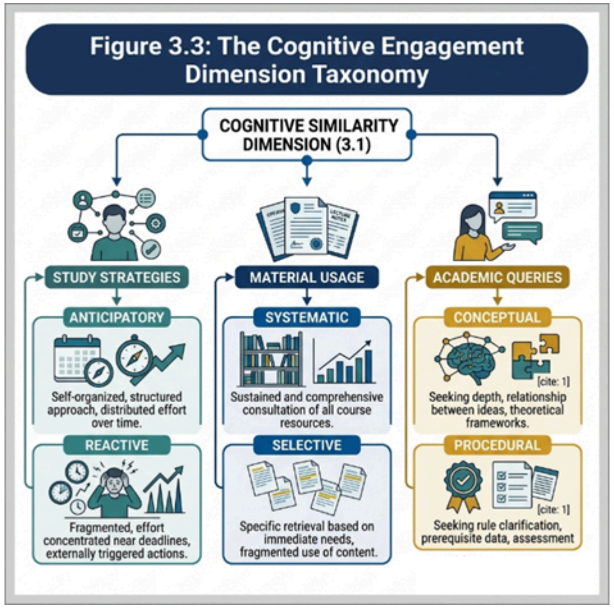

Figure 3.3: The Cognitive Engagement Dimension Taxonomy

3.2 Behavioural Similarity

Where cognitive similarity concerns ways of engaging with content, behavioural similarity concerns the **observable rhythms of engagement**. It asks when students act, how regularly they do so, how closely their activity is tied to institutional deadlines, and how they respond to prompts or reminders. These temporal and procedural patterns are often highly informative. They may reveal forms of self-regulation, organisational strain, or emerging risk long before formal outcomes make such dynamics visible.

A central indicator of behavioural similarity is the frequency and timing of interaction. Some students engage with academic resources in a relatively distributed and continuous way. Their activity is spread across the semester, suggesting stable routines and some degree of anticipatory planning. Others interact more episodically, with clear peaks of activity clustered around deadlines, assessments, or moments of urgency. Such patterns are not incidental. They indicate different relationships to academic time and different modes of coping with institutional demands. A support system that fails to attend to these rhythms risks offering guidance at the wrong moment or in the wrong form.

A second important indicator concerns adherence to deadlines and procedural requirements. Timely submission of assignments, regular consultation of updates, and sustained compliance with course expectations are signs of relative stability in the alignment between intention and action. By contrast, repeated delay, concentration of effort at the last moment, or fluctuating patterns of response may suggest organisational difficulty, competing pressures, or irregular self-regulation. Behavioural similarity makes it possible to detect these tendencies not as isolated events, but as patterns that acquire meaning through recurrence and timing.

A third indicator is responsiveness to reminders and prompts. In an adaptive support environment, this is especially relevant. Some students respond readily to light-touch prompting: a reminder or cue is sufficient to trigger action. Others show limited responsiveness, even when reminders are timely and clear. This difference is pedagogically important because it helps determine whether minimal scaffolding is sufficient or whether a more structured intervention may be needed. Importantly, the interpretation of such behaviour must remain cautious. A single missed reminder or brief period of inactivity does not warrant a strong conclusion. What matters is the persistence of the pattern across time.

Behavioural similarity therefore provides a way of understanding not only whether students engage, but how their engagement is temporally organised. In the HITS framework, this dimension is essential because the timing of support is itself pedagogically consequential. Study guidance, reminders, and escalation pathways cannot be treated as neutral delivery mechanisms; their effectiveness depends on how well they are aligned with actual rhythms of engagement. Behavioural similarity gives the system the basis for making such alignment possible.

3.3 Motivational Similarity

Motivational similarity concerns the attitudinal and affective dimension of academic engagement. While it is more difficult to observe directly than cognitive or behavioural patterns, it nevertheless leaves meaningful traces in the language students use, in the consistency of their participation, and in the kinds of reassurance they seek. This dimension is not intended to provide a psychological diagnosis of learners. Its purpose is more modest and more educationally useful: to identify recurring interactional signals that suggest different forms of confidence, uncertainty, persistence, or disengagement.

One of the most relevant indicators is the presence of linguistic signals of uncertainty or anxiety. Student queries often communicate more than their informational surface. Formulations such as “I don’t know where to start,” “I think I’m falling behind,” or “What happens if I fail?” may indicate not only a request for clarification, but also a broader experience of academic insecurity. When such expressions recur, they point to a support need that is not exhausted by providing more content alone. In such cases, the tone, structure, and sequencing of guidance become as important as the information itself.

Closely related are repeated requests for confirmation or reassurance. Some students frequently seek validation regarding deadlines, procedures, or correctness, even when the relevant information is already available. This does not necessarily indicate incapacity. It may instead reflect low confidence, fear of error, or a fragile sense of control over academic demands. Other students show a much lower need for confirmation and engage with institutional information more autonomously. These differences matter because they call for different forms of response. For one learner, concise clarification may be sufficient; for another, the same answer may need to be framed in a way that stabilises confidence and reduces uncertainty.

A further indicator of motivational similarity lies in continuity of engagement over time. Stable and regular interaction often suggests a relatively secure motivational orientation, even when difficulties are present. By contrast, marked oscillations—periods of intensive activity followed by prolonged silence, or rapid shifts between urgency and disengagement—may indicate motivational volatility. Such patterns become particularly meaningful when they appear alongside uncertainty signals or weak response to reminders. Here again, what matters is not a single episode, but the accumulation of evidence across multiple interactions.

Motivational similarity thus contributes an essential dimension to the overall analytical framework. It helps explain why students with apparently similar cognitive or behavioural patterns may nevertheless require different forms of support. One learner may need concise procedural guidance; another may need reassurance, pacing, and reduced friction in order to remain engaged. By incorporating this dimension into the analysis, the HITS framework avoids treating academic engagement as purely informational or behavioural. It recognises that persistence in higher education is also shaped by how students experience their own relation to study, uncertainty, and support.

Conclusion

The three dimensions outlined above—cognitive, behavioural, and motivational—constitute the analytical foundation upon which the rest of the deliverable is built. They do not produce fixed classifications, nor do they claim to exhaust the complexity of academic life. Their function is to identify recurring and educationally meaningful patterns through which differentiated support may be designed more responsibly. By moving from isolated traces to structured similarity patterns, the framework creates the conditions for the next step of the argument: the identification of learner profiles and the derivation of corresponding tutoring styles. In this sense, similarity analysis does not conclude with description. It prepares the ground for pedagogical and architectural differentiation.

4. Identified Learning Profiles

The analysis of similarity patterns described in the previous section led to the identification of four recurring profiles of academic engagement. These profiles do not represent fixed categories, nor are they intended to label students in a deterministic or exhaustive way. Their function is analytical rather than classificatory. They make visible recurrent combinations of cognitive, behavioural, and motivational tendencies that can help explain why different students may require different forms of support, even when they share a similar academic context.

The profiles should therefore be understood as provisional interpretive constructs. They emerge from recurring patterns observed across interaction traces, study habits, temporal rhythms, and linguistic signals, and remain open to revision as additional evidence accumulates. This point is methodologically and ethically important. The purpose of the framework is not to attach enduring identities to learners, but to recognise recurring configurations of need and engagement in ways that may support more proportionate and pedagogically coherent intervention.

Four profiles were identified with particular consistency across the analytical dimensions: **Structured Achievers**, **Reactive Performers**, **At-Risk Students**, and **Self-Directed Learners**. These profiles are not mutually exclusive in a rigid sense, nor do they imply that all learners can be neatly assigned to one profile alone at all times. Rather, they indicate prevailing patterns of academic behaviour that may be more or less visible at different moments in a learner's trajectory. Their usefulness lies not in simplification for its own sake, but in the possibility of relating recurring forms of engagement to differentiated tutoring logic and, ultimately, to architectural choices within the chatbot system.

4.1 Structured Achievers

Structured Achievers are characterised by relative stability across the three dimensions of analysis. Their engagement tends to be organised, anticipatory, and distributed over time. They interact with academic materials regularly, follow the expected progression of courses, and generally maintain continuity in relation to deadlines and institutional requirements. Their queries, when they arise, are often oriented toward refinement, clarification, or deepening of understanding rather than toward urgent procedural uncertainty.

What distinguishes this profile is not simply academic success in a narrow evaluative sense, but the consistency of engagement. Study routines appear sufficiently internalised, reminders are not heavily relied upon to trigger action, and expressions of confusion or disorientation are relatively infrequent. This does not mean that these students do not encounter difficulties. Rather, when difficulties arise, they tend to occur within an otherwise stable and self-regulated academic trajectory.

From an analytical perspective, this profile is significant because it shows that a support system must be able to recognise when intensive intervention is unnecessary. In such cases, the main educational issue is not the correction of instability, but the maintenance of relevance and precision. Structured Achievers illustrate that academic support may need to function less as behavioural regulation and more as selective academic reinforcement. Their profile therefore indicates a need for support that remains light, accurate, and appropriately non-intrusive.

4.2 Reactive Performers

Reactive Performers display a different configuration of engagement. Their academic activity is often organised around moments of urgency rather than distributed evenly across time. Interaction tends to intensify as deadlines approach, and periods of relative inactivity are often followed by concentrated bursts of study or information-seeking behaviour. This profile is not defined by low ability or lack of commitment, but by a tendency to mobilise effort reactively in response to external academic pressure.

Behaviourally, this profile is visible in clustered patterns of interaction, greater dependence on institutional cues, and a stronger role of reminders or deadlines in triggering action. Procedural questions often become more frequent near assessment points, and engagement appears closely tied to the immediate demands of the academic calendar. Motivationally, the tone of interaction in this group is often practical and task-oriented rather than openly anxious, although the concentration of effort around deadlines may create greater fragility when unexpected obstacles occur.

The educational significance of this profile lies in its temporal organisation of effort. Reactive Performers are not necessarily disengaged, but their relation to academic time is more compressed and externally triggered than in the case of more stable profiles. This makes them particularly sensitive to the timing, visibility, and clarity of support. Their profile suggests that some learners benefit not from more information in general, but from support that helps make academic requirements visible before they become urgent.

4.3 At-Risk Students

The profile of At-Risk Students emerges when irregularities across cognitive, behavioural, and motivational dimensions begin to converge. Engagement appears fragmented or unstable, and patterns of activity often fluctuate sharply across relatively short periods. Delayed or inconsistent responses to reminders, recurring uncertainty about course structure or requirements, and repeated expressions of confusion or discouragement may appear together. No single trace is sufficient to define the profile; what matters is the recurrence of these signals across time.

Cognitively, this profile may involve uncertainty not only about specific content, but also about the broader organisation of academic demands. Questions may concern matters that have already been communicated or basic structural elements of the course, suggesting difficulty in orientation rather than isolated misunderstanding. Behaviourally, engagement may be marked by silence, last-minute activity, or weak continuity. Motivationally, linguistic expressions of disconnection, frustration, or being overwhelmed tend to appear more frequently. These patterns do not imply lack of potential or deficit; rather, they indicate a position of increased academic vulnerability.

The analytical value of this profile lies in its preventive significance. It makes visible the forms of fragility that may precede more serious disengagement if left unaddressed. At-Risk Students therefore represent a profile for which timeliness and proportionate responsiveness are especially important. Their presence within the framework reflects the recognition that academic risk is rarely the product of one variable alone; it usually emerges through the interaction of organisational strain, motivational fatigue, contextual pressure, and uncertainty in navigating academic demands.

4.4 Self-Directed Learners

Self-Directed Learners are characterised by a relatively high degree of autonomy in their relation to academic work. Their engagement is less driven by deadlines or repeated prompting and more closely tied to conceptual interest, self-initiated inquiry, or longer-term academic orientation. Their questions often concern clarification, interpretation, or intellectual extension rather than procedural reassurance, and their interaction with resources tends to be selective, purposeful, and relatively stable over time.

From a cognitive perspective, this profile is marked by the prominence of conceptual inquiry. Behaviourally, it is associated with regular but non-dependent forms of interaction: these students do not usually require extensive prompting in order to remain engaged. Motivationally, their activity tends to reflect a more stable orientation toward learning itself, even when difficulties or uncertainties are present. This does not mean that they are immune to challenge, but rather that their engagement is more often internally sustained than externally triggered.

The importance of this profile lies in the fact that it highlights a form of academic engagement for which support is valuable not because it regulates behaviour, but because it facilitates precision, depth, and efficient access to authoritative materials. Self-Directed Learners illustrate that digital tutoring does not always need to intervene by increasing structure. In some cases, its value lies in accompanying inquiry, reducing informational friction, and supporting the learner's own initiative without displacing it.

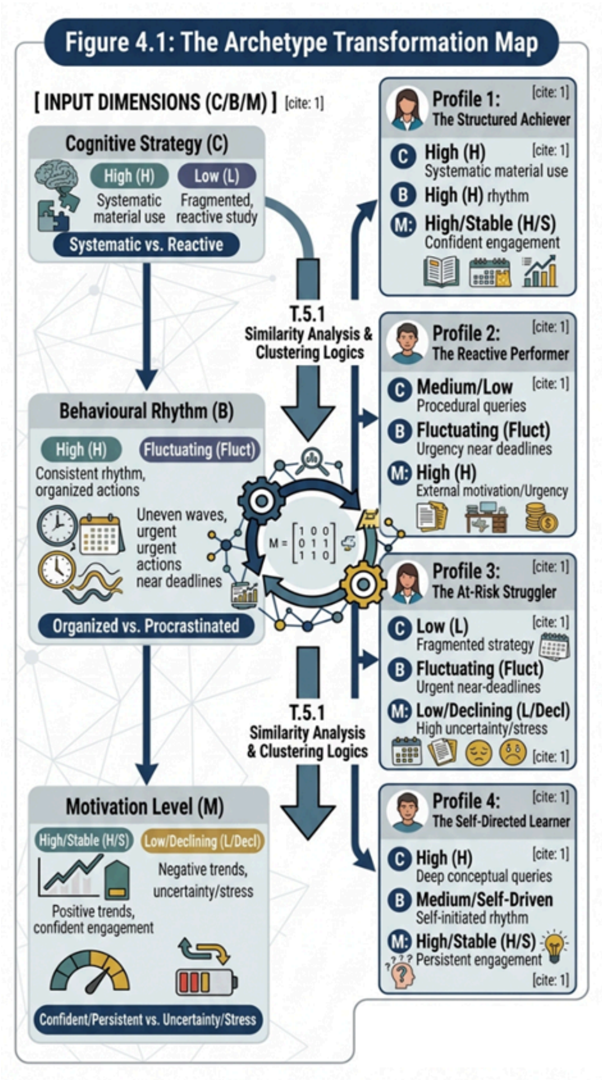

Conclusion

Taken together, these four profiles provide a structured description of recurring forms of academic engagement within the HITS framework. Their value lies in making diversity analytically legible without transforming that diversity into rigid typologies. Each profile highlights a different configuration of stability, urgency, uncertainty, and autonomy, and therefore points toward a different educational logic of support. The next section builds on this descriptive foundation by showing how these profiles inform differentiated tutoring styles. In this way, the movement from similarity analysis to pedagogical response becomes more explicit: profiles remain the interpretive bridge between observed behavioural patterns and the design of proportionate tutoring strategies.

5. Derived Tutoring Styles

The tutoring styles developed within WP5 did not emerge independently of the analytical work presented in the preceding sections. They are the pedagogical consequence of the similarity patterns and learner profiles identified in T.5.1, and they are further shaped by the ethical and methodological commitments articulated in Deliverable 4.3, especially with regard to longitudinal validation, adaptive fidelity, and the avoidance of epistemic freezing. Tutoring, in this framework, is not understood as a generic layer of assistance applied uniformly to all learners. It is conceived as a differentiated and revisable response to identifiable forms of academic engagement.

For this reason, the styles described below should not be read as fixed intervention templates mechanically attached to learner profiles. Rather, they express the forms of educational response that become appropriate when certain patterns of behaviour, uncertainty, stability, or autonomy recur with sufficient consistency. The relation between profile and tutoring style is therefore interpretive rather than automatic. The same student may move from one prevailing pattern to another over time, and the tutoring logic must remain capable of recalibration. This principle is essential to the HITS framework: support should become more precise and proportionate as evidence accumulates, but it should never harden into immutable classification.

The tutoring styles derived from the analysis may be grouped into four broad forms: **structured coaching**, **adaptive reminder systems**, **preventive intervention with attentive monitoring**, and **human escalation under ethical safeguards**. These forms correspond to different educational needs, different levels of intervention intensity, and different degrees of dependence on human judgement. What unifies them is the commitment to proportionality. Support should be neither excessive nor absent, neither intrusive nor indifferent. It should be calibrated to the learner's evolving relation to academic tasks, institutional demands, and available support.

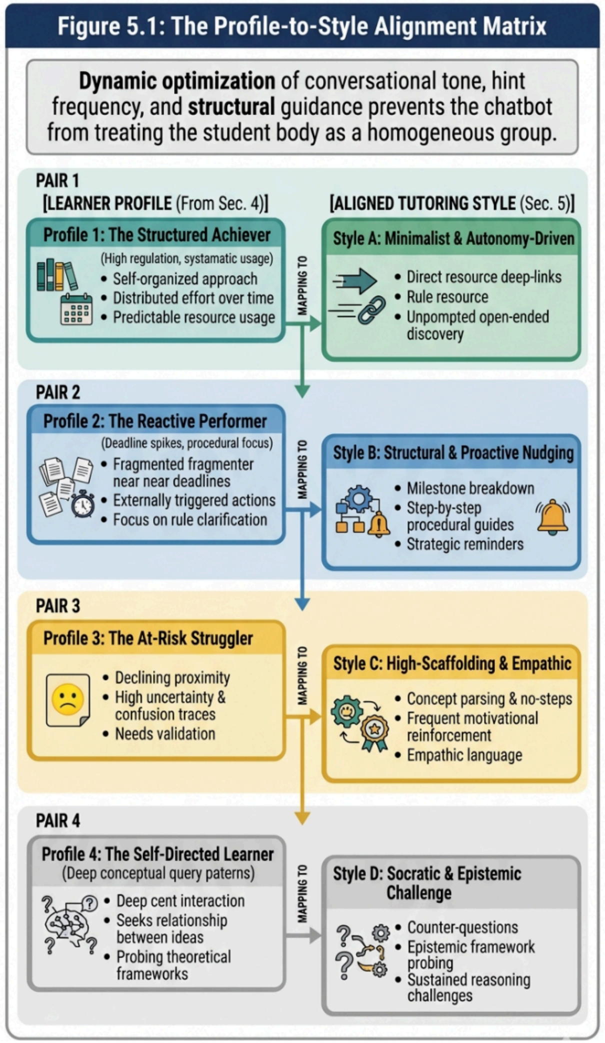

Figure 5.1: The Profile-to-Style Alignment Matrix

5.1 Structured Coaching

Structured coaching is the tutoring logic most clearly associated with students whose academic engagement is stable, anticipatory, and internally regulated. In these cases, the main pedagogical need is not behavioural correction but informed accompaniment. Learners already display continuity in their interaction with course materials, deadlines, and institutional expectations. What they require is support that preserves autonomy while enhancing the relevance, precision, and depth of academic guidance.

Within this tutoring style, the role of the system is therefore selective rather than corrective. The chatbot does not need to intensify reminders or impose a strong prompting rhythm. Instead, it functions as a resource for accurate clarification, contextual retrieval, and occasional enrichment. Precision becomes more important than frequency. Support is offered in response to initiative rather than in place of it, and the value of the system lies in reducing informational friction while respecting the learner's capacity for self-organisation.

This tutoring style is pedagogically important because it recognises that effective support does not always take the form of stronger intervention. In some cases, educational care is best expressed through restraint. A system that over-prompts stable learners risks diminishing perceived usefulness and may even become counterproductive. Structured coaching therefore represents the logic of support appropriate when continuity is already present and when what matters most is the refinement of learning rather than the restoration of engagement.

At the same time, stability is not treated as permanent. Longitudinal monitoring remains active, though unobtrusive. If sustained patterns of declining responsiveness, increased uncertainty, or irregular engagement emerge, the tutoring logic may shift accordingly. Structured coaching thus rests on a principle of confidence without complacency: it acknowledges the learner's autonomy while remaining open to change over time.

5.2 Adaptive Reminder Systems

A second tutoring style is appropriate where engagement is present but temporally compressed and strongly shaped by external academic triggers. In such cases, the main issue is not lack of participation as such, but the instability created by reactive rhythms of effort. Students may mobilise effectively under deadline pressure, yet still remain vulnerable to overload, misunderstanding, or delayed action. The educational task is therefore to stabilise academic timing without replacing student agency.

Adaptive reminder systems respond to this need by functioning as a form of temporal scaffolding. Their role is not to discipline or monitor students continuously, but to make relevant future requirements visible in time for them to become manageable. Reminders, prompts, and procedural cues are calibrated in relation to patterns of prior engagement. Where learners tend to activate only close to deadlines, the system introduces earlier signals designed to redistribute effort more evenly across the academic term. The emphasis is on anticipation rather than reaction.

The pedagogical significance of this tutoring style lies in the way it reduces planning friction. Students who rely heavily on urgency often do not need more information in absolute terms; they need information at a moment when it can still be acted upon effectively. In this sense, the tutoring logic operates not at the level of content alone, but at the level of timing. It makes academic demands more predictable and therefore more governable. The aim is not to impose external control, but to create conditions under which more stable study habits may gradually emerge.

Here too, longitudinal sensitivity is essential. A missed reminder or delayed interaction is not, in itself, sufficient to justify intensified intervention. The tutoring logic remains cautious and accumulative. It responds to recurring rhythms, not to isolated events. This restraint is pedagogically and ethically important because it prevents the system from treating temporary fluctuation as evidence of entrenched difficulty. Adaptive reminder systems therefore combine responsiveness with measured interpretation, supporting continuity without narrowing the space for self-directed recovery.

5.3 Preventive Intervention and Attentive Monitoring

A more explicit and structured tutoring response becomes appropriate when recurring signs of uncertainty, irregularity, and weakened continuity begin to converge. In such cases, support can no longer rely primarily on optimisation or timing alone. The educational priority becomes prevention: not in the sense of controlling students before failure occurs, but in the sense of recognising emerging vulnerability early enough for support to remain effective and proportionate.

Preventive intervention is characterised by greater attentiveness to signs of disorientation, disengagement, or repeated difficulty in translating academic requirements into action. Guidance becomes more explicit, steps are broken down into smaller and more manageable units, and relevant resources are surfaced with minimal friction. Responses may also incorporate a stronger reassuring dimension, not because the system takes on a therapeutic role, but because academic insecurity often requires more than factual clarification. When uncertainty becomes recurrent, the form in which guidance is delivered becomes educationally consequential.

This tutoring style depends heavily on the principle of measured escalation. One of the main insights carried over from Deliverable 4.3 is that early traces should not solidify too quickly into stable representations of the learner. Preventive intervention therefore requires recurrence and pattern confirmation. Brief disorientation, temporary delay, or isolated expressions of doubt are not treated as sufficient grounds for a strong interpretive shift. What matters is the accumulation of evidence across multiple interactions and over time. Only then does the system move from light support toward more clearly structured accompaniment.

In this sense, preventive intervention is as much an ethical stance as it is a tutoring style. It recognises vulnerability without converting it into a label. It intensifies support where needed, but does so with interpretive caution. Its function is to shorten the interval between the emergence of academic fragility and the provision of meaningful guidance, while still preserving the possibility that learners may regain continuity without being defined by a passing state of difficulty.

5.4 Human Escalation and Ethical Safeguards

There are situations in which even a well-designed adaptive system should not remain the final layer of response. Digital tutoring is effective in identifying recurring patterns, clarifying procedures, sequencing next steps, and lowering informational barriers. It is less capable of grasping the full complexity of personal circumstances, emotional nuance, or contextual factors that may shape a learner’s trajectory. For this reason, the HITS framework includes a tutoring logic based on human escalation.

Human escalation does not imply that the system has failed. On the contrary, it represents one of the conditions of its responsible use. When signs of sustained disengagement, persistent uncertainty, or repeated non-response exceed the scope of automated support, the chatbot is expected to activate a pathway toward human review. In practical terms, this means that the system can support triage: it can organise relevant traces of interaction, clarify the type of emerging difficulty, and direct the learner toward appropriate tutors or services. What it does not do is transform these signals into definitive judgements. Final interpretation remains with human actors.

This tutoring style gives concrete form to the human-in-the-loop principle already emphasised in the broader HITS framework. Automated indicators are treated as prompts for attention, not as substitutes for professional discernment. This is a crucial ethical safeguard. Without it, personalization could drift into surveillance, and probabilistic signals could be misread as stable truths about the learner. Human escalation ensures that interpretation remains dialogical, revisable, and open to contextual explanation.

Educationally, this also preserves the relational dimension of tutoring. Some forms of academic difficulty cannot be responsibly addressed through automated guidance alone. The function of the chatbot is therefore not to replace human accompaniment, but to help ensure that human intervention occurs at a stage when it can still be proportionate and constructive. The value of escalation lies precisely in this balance: automation supports continuity and responsiveness, while human review protects fairness, nuance, and learner dignity.

5.5 Longitudinal Adaptivity as a Structural Principle

Across all the tutoring styles described above, one principle remains constant: support is not derived from a one-time classification, but from an ongoing interpretation of evolving evidence. Learner profiles and tutoring responses are provisional, revisable, and longitudinally updated. This principle is not an ancillary methodological choice. It is one of the structural conditions that allows personalization to remain educationally legitimate.

Longitudinal adaptivity means that support intensity, timing, and form are adjusted in relation to recurring patterns rather than isolated events. Early signals may later lose weight if subsequent behaviour points in a different direction. Conversely, apparently minor traces may acquire significance if they recur consistently over time. In this way, the tutoring framework resists epistemic freezing and preserves the possibility of change. Learners are not reduced to fixed categories; they are accompanied through trajectories that remain open to revision.

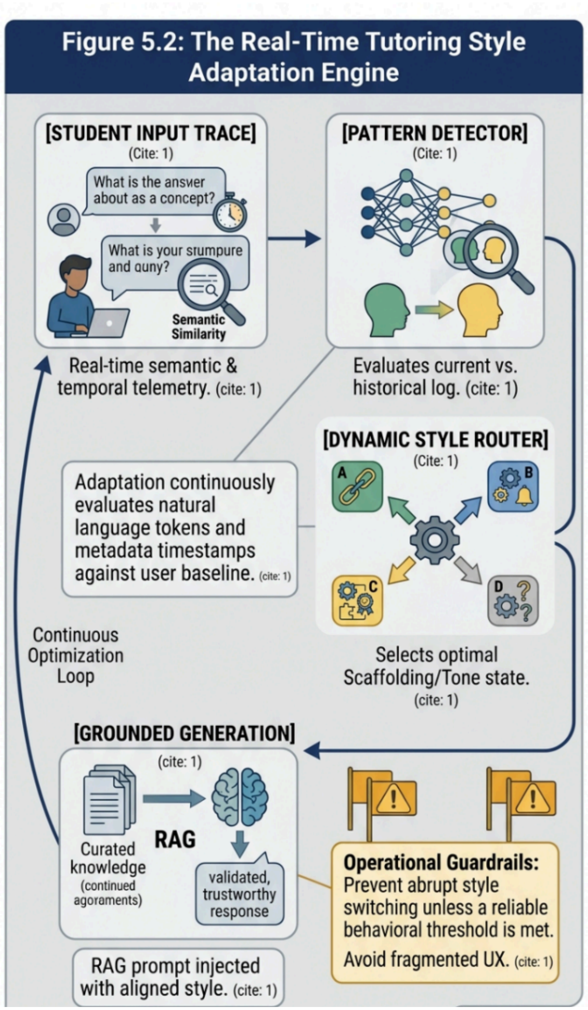

Figure 5.2: The Real-Time Tutoring Style Adaptation Engine

This principle gives coherence to the tutoring framework as a whole. Structured coaching, adaptive reminders, preventive intervention, and human escalation are not separate mechanisms applied in isolation. They are part of a continuum of pedagogical response governed by proportionality, revisability, and ethical restraint. It is this longitudinal logic that allows the HITS framework to combine analytical precision with educational care.

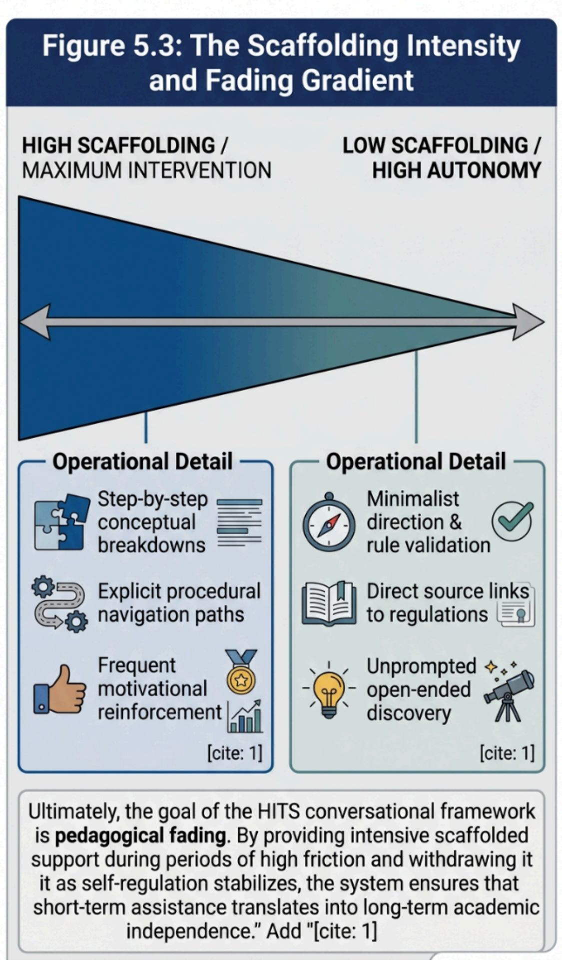

Figure 5.3: The Scaffolding Intensity and Fading Gradient

Conclusion

The tutoring styles outlined in this section translate the learner profiles identified earlier into differentiated forms of pedagogical response. They make clear that within the HITS framework, support is neither uniform nor fixed, but modulated in relation to stability, urgency, vulnerability, and autonomy. At the same time, the section also shows that intervention remains bounded by a broader methodological and ethical principle: interpretation must remain revisable, and support must remain proportionate. This operational logic provides the bridge to the next section, where the consequences of these analytical and pedagogical choices for chatbot architecture are examined more directly.

6. How Similarity Shaped the Chatbot Architecture

The architecture of the HITS Chatbot was conceived as a structural response to learner diversity. The analysis of cognitive strategies, behavioural rhythms, and motivational signals did not remain at the level of pedagogical interpretation alone. It informed the way the system was built. In this sense, the architecture should not be understood as a neutral technical container within which personalization was later inserted. Rather, personalization is embedded in the core logic of the system itself. The knowledge base, retrieval mechanisms, adaptive prompting, and routing pathways were all designed in ways that reflect the differentiated forms of academic engagement identified in the earlier analytical work.

This point is important because it clarifies the role of architecture within the broader argument of T.5.1. The chatbot is not adaptive because it offers different types of answers in a superficial way, nor because it contains isolated personalization features layered onto a general conversational model. It is adaptive because its internal structure was designed to support differentiated forms of academic assistance from the outset. Similarity analysis made it possible to identify what kinds of support learners might need. Architecture provided the means through which those needs could be addressed reliably, proportionately, and within institutional boundaries.

The following subsections explain how this translation from behavioural insight to technical design took place. They show, first, why a RAG architecture became necessary as a mechanism of trust; second, how semantic retrieval supports a more contextual understanding of student intent; third, how adaptive logic and hybrid routing respond to differentiated rhythms of engagement; and fourth, how modularity allows the system to remain transferable across institutional settings without sacrificing local specificity.

6.1 RAG as a Structural Trust Mechanism

One of the clearest implications of the similarity analysis was that trust could not be treated as a secondary quality of the system. Students who engage regularly with academic support expect precision. Students who experience uncertainty or disorientation depend even more strongly on reliable institutional guidance. In both cases, a small inaccuracy may have disproportionate consequences. Incorrect information about deadlines, prerequisites, examination rules, or available services may not merely inconvenience learners; it may undermine their confidence in the support environment itself.

The architectural response to this requirement was the adoption of a Retrieval-Augmented Generation framework. Instead of allowing the language model to generate responses solely from its general statistical training, the chatbot retrieves relevant passages from a curated institutional corpus and formulates its answer on that basis.

This architecture does not eliminate linguistic generation, but it constrains it within a documentary environment that is governed, auditable, and institutionally legitimate. In this way, the system separates fluency from authority: responses are linguistically produced, but their informational legitimacy depends on retrieved evidence.

The significance of this choice is both technical and pedagogical. Technically, it reduces the risk of unsupported claims and makes the response process more traceable. Pedagogically, it ensures that the support offered to students remains grounded in the institutional realities within which their academic decisions must be made. The system can explain procedures, clarify rules, or retrieve course-relevant information not because it “knows” the university in an abstract sense, but because it has access to a defined body of authoritative material. If the relevant information is not present in that body of material, the system is expected to acknowledge the limit rather than fabricate certainty. This capacity to recognise its own boundaries is not a weakness. It is part of what makes the system trustworthy.

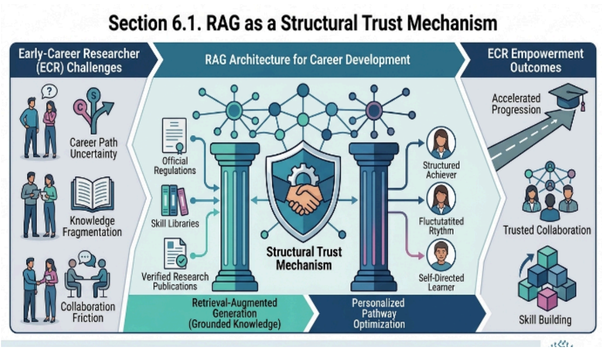

The corpus itself is produced through a controlled ingestion process. Institutional regulations, syllabi, handbooks, and selected web materials are collected, cleaned, standardised, and transformed into a searchable documentary environment. What matters here is not only the existence of such a corpus, but the fact that it remains bounded and governable. The chatbot does not operate by roaming the open web during interaction. It works within an institutional knowledge space whose contents can be reviewed, updated, and administratively controlled. This is what allows trust to become structural rather than merely rhetorical.

6.2 Semantic Retrieval and Intent Sensitivity

If trust depends on the quality of the knowledge base, relevance depends on the way that knowledge is retrieved. Student queries do not arrive in a uniform or standardised form. Two students may express very similar needs through very different kinds of language. One may ask directly about an examination requirement; another may formulate a broader and more uncertain narrative about not knowing how to proceed. A retrieval mechanism based only on literal word matching would struggle to capture such variation. What is needed instead is a way of identifying underlying relevance beyond surface vocabulary.

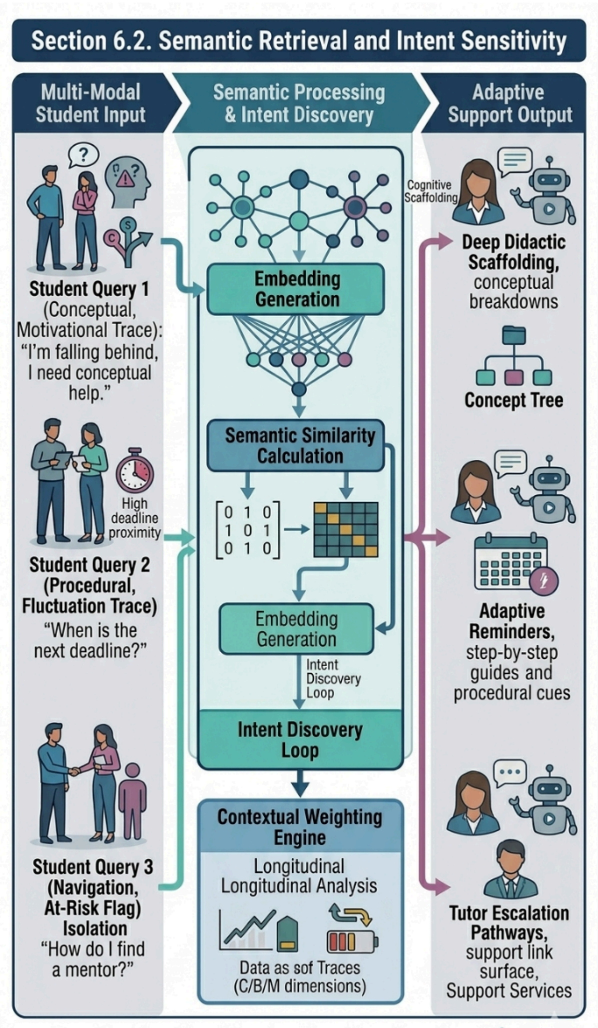

For this reason, the HITS architecture relies on semantic retrieval. Institutional documents and student queries are represented in a shared semantic space in which contextual similarity can be computed. At a technical level, this is enabled through embeddings and a vector-based retrieval layer. At the pedagogical level, what matters is that the system is able to retrieve material on the basis of meaning rather than exact phrasing alone. A student who writes “I am falling behind” may not explicitly mention tutoring, planning, or counselling services, yet the underlying support need can still be interpreted through contextual proximity. Similarly, a detailed conceptual question about a course may appropriately retrieve didactic content rather than administrative guidance, even when both concern the same curricular setting.

This is where the architectural consequences of similarity analysis become especially visible. Cognitive similarity helps the system distinguish between conceptual inquiry and procedural clarification. Motivational similarity influences the way in which uncertainty or reassurance needs are recognised. Behavioural traces can provide further contextual weight, indicating whether a given question should be interpreted as part of a stable learning trajectory, a deadline-driven request, or an emerging pattern of disorientation. In this sense, retrieval is not simply a matter of information access. It is part of the system’s interpretive responsiveness to differentiated academic situations.

Importantly, this responsiveness does not require the system to assign fixed labels to learners in advance. Personalization emerges through contextual interpretation at the level of interaction rather than through static categorisation. The retrieval process remains dynamic: it situates each query within a semantic and institutional environment informed by prior analysis, but without reducing the learner to a rigid profile before the response is generated. This is one of the reasons why semantic retrieval is so central to the HITS architecture. It allows the system to remain sensitive to difference without converting difference into inflexible classification.

6.3 Adaptive Logic and Hybrid Routing

The analysis of behavioural similarity shaped not only what information the system retrieves, but also how support is timed, structured, and, when necessary, escalated. The HITS Chatbot was designed as part of a hybrid service model in which digital support and human tutoring remain complementary rather than mutually exclusive. This model recognises that some academic needs can be addressed efficiently through timely and well-grounded automation, while others require relational interpretation, contextual judgement, or emotional nuance that remain beyond the proper scope of an automated system.

At the first level, the chatbot is able to manage a substantial range of high-frequency, low-complexity interactions. It can clarify deadlines, retrieve procedural information, explain aspects of course organisation, and guide students toward relevant documents or services. This form of automation is not trivial. It reduces friction, extends the availability of support beyond office hours, and relieves academic staff from repetitive informational work. In doing so, it strengthens institutional responsiveness without displacing human tutoring where it is truly needed.

At a second level, the system makes use of adaptive logic informed by behavioural patterns. Students whose interaction tends to cluster around assessments may receive more anticipatory prompts. Those whose engagement is already stable and distributed encounter much lighter scaffolding. Where sustained irregularity appears, the form of support becomes more explicit: tasks may be more clearly sequenced, expectations more visibly restated, and access to relevant resources simplified. What changes here is not the basic function of the system, but the intensity and timing of its pedagogical response.

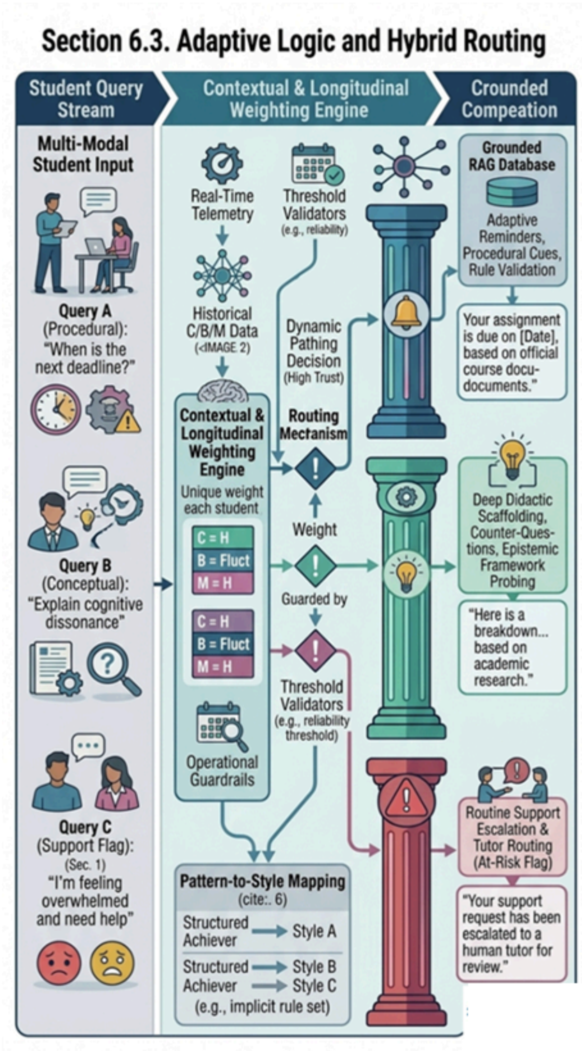

This adaptive logic remains longitudinal and cautious. A single missed interaction, delayed response, or temporary drop in activity is not interpreted as sufficient evidence for intensified support. The system responds to recurring tendencies rather than isolated deviations. This is essential both educationally and ethically. It prevents temporary fluctuations from being overinterpreted and preserves the possibility that students may recover continuity without premature escalation. In this sense, adaptive logic is inseparable from interpretive restraint.

Where interaction patterns move beyond the legitimate scope of automated support, routing mechanisms make escalation possible. The system can help identify when repeated uncertainty, persistent non-response, or more complex forms of vulnerability suggest the need for human follow-up. It may support this process by structuring recent interaction evidence and directing the learner toward the appropriate service or tutor. The final interpretive responsibility, however, remains with human actors. Hybrid routing therefore protects the balance between efficiency and fairness: the chatbot enhances responsiveness, but it does not replace professional pedagogical judgement.

6.4 Modularity and Institutional Customization

The HITS framework was not designed for a single academic setting alone. From the outset, it had to remain compatible with institutional variability across European higher education. Regulations, curricular structures, languages, digital infrastructures, and governance practices differ significantly from one university to another. A chatbot architecture that assumed uniformity would therefore be pedagogically and operationally fragile. For this reason, the HITS system adopts a modular design.

The principle of modularity applies above all to the relation between knowledge ingestion and the core conversational architecture. Institutional documents can be integrated into the system through a standardised pipeline without altering the general logic of retrieval, generation, or routing. This allows each university to define and maintain the boundaries of its own knowledge base while preserving architectural consistency across deployments. In practical terms, the system can therefore be adapted to different institutional ecologies without becoming a different system in conceptual terms.

Modularity also supports flexibility at the infrastructure level. Embedding models, storage environments, and deployment choices may need to vary depending on local privacy requirements, technical resources, or language needs. Some institutions may prioritise local hosting and stronger data-control guarantees; others may emphasise scalability and integration with existing systems. The architecture is designed to accommodate such variation without requiring changes to the pedagogical logic that underpins it. This is important because it ensures that the behavioural framework derived from similarity analysis remains transferable even when institutional conditions differ.

The educational significance of modularity lies in the fact that it prevents transferability from becoming homogenisation. The HITS system does not impose a single content environment or a fixed institutional script. Instead, it offers a structured method through which local institutional knowledge can be transformed into a responsive and trustworthy support layer. In this way, the architecture remains both stable and adaptable: stable in its conceptual logic, adaptable in its documentary and infrastructural instantiation.

6.5 Architecture as the Translation of Behavioural Insight

Taken together, the architectural choices discussed above make visible the central argument of this deliverable: the HITS Chatbot is not simply a conversational interface augmented with a few personalization features. It is a system whose technical organisation reflects prior behavioural interpretation. Similarity analysis provided the framework through which differences in engagement could be understood. The architecture translated that framework into retrieval logic, intervention timing, support routing, and institutional grounding.

This is what gives the architecture its pedagogical meaning. RAG establishes the conditions of epistemic trust. Semantic retrieval allows queries to be interpreted contextually rather than superficially. Adaptive logic makes support responsive to rhythms of engagement, while hybrid routing preserves the role of human judgement where automation reaches its appropriate limit. Modularity, finally, ensures that the same pedagogical logic can be adapted across diverse institutional settings without losing coherence. In this sense, architecture becomes the technical expression of a prior educational claim: that academic support must be reliable, context-sensitive, proportionate, and ethically bounded if it is to respond meaningfully to learner diversity.

Conclusion

Section 6 has shown that the architecture of the HITS Chatbot is not external to the analytical work of T.5.1, but is one of its principal outcomes. The system's trust model, retrieval logic, adaptive prompting, routing pathways, and modular design all derive from the need to translate behavioural insight into responsible technical form. The next section addresses a further question that naturally follows from this architecture: how such a system may be evaluated in ways that are technically rigorous, educationally meaningful, and ethically proportionate.

7. Validation and Key Indicators

The value of the HITS Chatbot cannot be inferred from architectural sophistication alone. A system may be technically robust, conceptually elegant, and institutionally well grounded, yet still fail to contribute meaningfully to students' academic continuity. For this reason, validation is not treated in the HITS framework as a final checkpoint applied once development is complete. It is conceived as a continuous and reflexive dimension of the system itself. Evaluation must accompany design, implementation, and refinement throughout the project lifecycle. This position is fully consistent with the broader methodological orientation already articulated in Deliverable 4.3, where measurement, interpretation, and ethical responsibility are treated as mutually dependent rather than separable concerns.

The validation framework adopted here is therefore intentionally layered. It does not rely on a single metric, nor does it reduce system quality to user volume or interface efficiency. Instead, it considers four interrelated dimensions. First, technical validation concerns the reliability and epistemic integrity of the system: whether responses are grounded, accurate, and appropriately bounded. Second, behavioural validation concerns the degree to which support mechanisms are responsive to evolving engagement patterns without overreacting to isolated signals. Third, academic validation concerns whether the system contributes, even indirectly, to more stable participation, reduced friction, and earlier support in situations of difficulty. Fourth, ethical validation concerns fairness, contextual integrity, and the maintenance of human oversight. Taken together, these dimensions offer a framework through which the educational value of the chatbot can be assessed without collapsing it into purely technical or purely quantitative terms.

This layered approach also answers an important methodological need. The reviewers' feedback made clear that the validation section of the original draft was conceptually convincing, but insufficiently explicit in operational terms. In response to that concern, the present version makes more visible the kinds of indicators that can support evaluation in practice. These indicators should not be understood as rigid benchmarks fixed in advance for all contexts. They constitute an initial evaluative scaffold, to be refined through pilot use, institutional adaptation, and iterative review. Their role is not to close interpretation, but to make the validation logic more transparent and actionable.

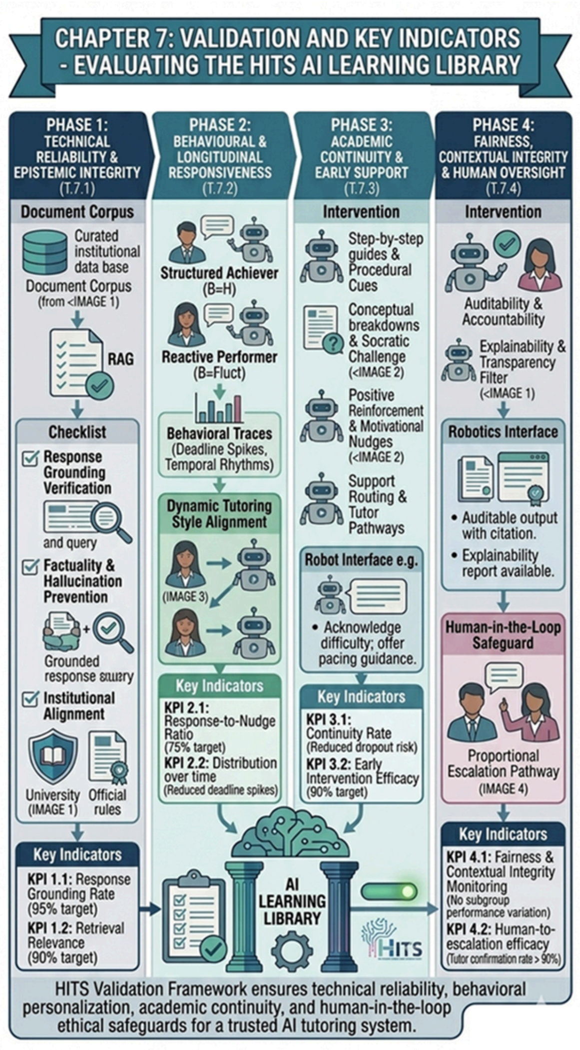

7.1 Technical Reliability and Epistemic Integrity

In academic environments, informational reliability is a foundational condition of trust. If a system provides misleading guidance about deadlines, assessment procedures, prerequisites, or rights and obligations, it risks undermining not only its own credibility but also the institutional confidence on which student support depends. For this reason, technical validation begins with the verification of epistemic integrity. The central question is whether the chatbot produces responses that are adequately grounded in authoritative institutional sources and whether it respects the limits of what those sources can support.

At a practical level, this form of validation includes checking retrieval relevance and response grounding. Outputs may be periodically compared against the corresponding institutional documentation in order to verify whether the retrieved material is contextually appropriate and whether the generated answer remains semantically faithful to that source. In this context, it is useful to observe not only cases of error, but also cases in which the system appropriately signals that the requested information is not available in the corpus. Such responses should not be counted as failures. On the contrary, they indicate that the system is operating responsibly within its epistemic boundaries.

Illustrative technical indicators may include the proportion of responses correctly grounded in retrieved source material, the rate of retrieval judged relevant by expert review, the incidence of unsupported or hallucinatory responses, and the frequency with which the system correctly defers or redirects queries when the relevant information is unavailable. These indicators are especially important because they anchor trust in observable evidence rather than in impression alone. A pedagogically useful system must first be a reliable one.

7.2 Behavioural and Longitudinal Responsiveness

If technical validation addresses whether the system is trustworthy, behavioural validation addresses whether it is educationally responsive. The question here is not simply how often students interact with the chatbot, but how that interaction evolves over time and whether support is calibrated to meaningful patterns of engagement rather than to isolated events. This is one of the areas in which the HITS framework most clearly departs from static and purely transactional models of evaluation.

Behavioural responsiveness is interpreted longitudinally. Engagement counts by themselves are of limited value if they are not read in relation to timing, continuity, and distribution across the academic term. A student who interacts frequently only in the days immediately preceding an assessment may present a very different support profile from one who interacts less often but in a stable and anticipatory way. The key issue is therefore not volume alone, but rhythm. Validation should ask whether engagement becomes more distributed, whether support reaches students before urgency peaks, and whether prompting contributes to continuity rather than simple last-minute activation.

Particular attention should also be given to the **response-to-nudge ratio**, that is, the extent to which prompts or reminders are followed by meaningful academic action. Here again, the interpretation must remain cautious. One ignored reminder or one delayed response should not lead to a strong conclusion about the learner or about the system. What matters is persistence over time. The purpose of validation is not to identify static “responsive” and “non-responsive” students, but to examine whether the adaptive logic of the system is helping to stabilise engagement patterns while avoiding premature over-interpretation.

Illustrative behavioural indicators may therefore include response-to-nudge ratio, distribution of interaction across the semester, variation in deadline clustering, frequency of prolonged inactivity after support prompts, and the proportion of cases in which repeated signals lead to recalibrated rather than fixed system responses. These are not merely behavioural analytics in a narrow sense. They are indicators of whether the system is supporting continuity in a measured and context-sensitive way.

7.3 Academic Continuity and Early Support

Technical reliability and behavioural responsiveness matter because of their contribution to an educational objective that is broader than either of them considered separately: the preservation of academic continuity. The HITS Chatbot is not intended simply to answer questions more quickly than existing systems. Its deeper educational purpose is to reduce friction in academic navigation, support the early recognition of emerging difficulty, and contribute to more sustainable forms of participation. Validation must therefore also ask whether the system is helping students remain engaged with academic requirements before problems become entrenched.

This does not mean that the impact of the chatbot should be interpreted in simplistic causal terms. Academic continuity is shaped by many factors, and no digital intervention should claim sole responsibility for progression or retention outcomes. Nevertheless, there are meaningful indicators that can illuminate whether the system is supporting more stable educational trajectories. These include improvements in the regularity of interaction, more timely adherence to deadlines, reduced recurrence of unresolved procedural uncertainty, and earlier detection of patterns that may warrant human follow-up. Such effects may be modest and indirect, but they are educationally significant.

The validation framework should therefore pay attention to whether the chatbot contributes to shifting the timing of support. A key measure of its usefulness lies in whether signs of fragility are recognised early enough for intervention to remain proportionate and constructive. Here, the notion of “early support” is preferable to a simplistic language of prediction. The system does not function by foretelling failure. It helps make emerging need visible in time for support to be mobilised more effectively.

Illustrative academic indicators may include changes in assignment punctuality, increased consultation of relevant materials before critical deadlines, reduction in repeated confusion about institutional procedures, and the proportion of potential risk situations identified early enough to allow non-punitive intervention. These indicators are best interpreted as signs of support effectiveness rather than as proxies for academic worth. Their purpose is to show whether the system is helping to sustain continuity, not to rank students.

7.4 Fairness, Contextual Integrity, and Ethical Safeguards

No validation framework can be considered adequate if it ignores the ethical conditions under which digital support is provided. Within HITS, fairness is not treated as an external add-on to technical evaluation. It is part of the criteria by which the system itself must be judged. A support system may be accurate in aggregate and still produce unequal effects across linguistic, cultural, institutional, or accessibility-related differences. Validation must therefore remain attentive not only to performance, but also to contextual integrity.

This implies periodic review of how the system behaves across heterogeneous groups and settings. Interaction patterns, linguistic markers, or retrieval outcomes may not carry identical meaning across all users. What appears to be disengagement in one context may reflect a different communication style or institutional structure in another. For this reason, fairness monitoring should not be reduced to abstract statistical neutrality. It must also involve interpretive awareness of context and the possibility that indicators themselves may require recalibration when transferred across settings.

Bias monitoring procedures should therefore be integrated into validation cycles. Where disproportionate effects or recurring anomalies appear, system parameters and underlying assumptions should be reviewed. Similarly, human oversight must remain active throughout the process. Automated signals should be treated as prompts for reflection rather than as definitive judgements. The role of tutors and institutional actors is crucial precisely because they can reintroduce contextual knowledge, dialogue, and interpretive nuance into situations where algorithmic indicators alone are insufficient.

Illustrative fairness and ethics indicators may include subgroup variation in retrieval relevance, differential response patterns across contexts, frequency of tutor overrides or reinterpretations, and the proportion of escalations that are confirmed or revised after human review. These indicators are not intended to create another layer of technical control. They are intended to preserve the central ethical principle of the HITS framework: that digital support must remain accountable, revisable, and humanly interpretable.

7.5 Toward a Reflexive Validation Culture

The validation framework proposed here is best understood as reflexive rather than static. Its purpose is not to certify once and for all that the chatbot “works,” but to establish the conditions under which its educational function can be examined, challenged, and progressively improved. Technical audits, behavioural interpretation, academic continuity indicators, fairness checks, and human review do not constitute separate layers that can be applied independently. They are mutually reinforcing dimensions of a broader evaluative culture in which evidence is accumulated gradually and interpreted with caution.

This reflexive orientation is particularly important in a project that explicitly rejects rigid profiling and deterministic classification. If learner profiles are provisional and revisable, then the same must be true of the evaluative assumptions used to judge the system itself. Validation is therefore not simply the measurement of outcomes. It is part of the ethical and pedagogical design of the system. It ensures that support remains trustworthy, proportionate, and open to recalibration as new institutional, technical, and educational evidence emerges.

In this sense, validation becomes an educational practice in its own right. It protects trust by grounding support in accountable evidence; it protects learners by resisting premature judgement; and it protects institutions by ensuring that innovation remains aligned with pedagogical responsibility. Through this approach, the HITS framework seeks not only to deploy a technically capable chatbot, but to establish a model of digital academic support whose effectiveness is inseparable from its ethical and methodological discipline.

Conclusion

The present section has shown that validation within HITS must be understood as a multi-dimensional and ongoing process. Technical reliability, behavioural responsiveness, academic continuity, and fairness together provide the basis for evaluating whether the chatbot contributes meaningfully to student support. At the same time, the section has made clear that evaluation cannot be separated from the ethical commitments of the project. This naturally leads to the final thematic area of the deliverable: the broader ethical and European framework within which the HITS Chatbot is situated.

8. Ethical and European Dimension

The HITS Chatbot is developed within a European higher education context in which technological innovation cannot be separated from questions of responsibility, transparency, and the protection of individual rights. Its design is therefore not guided by technical ambition alone. It is also shaped by an explicit ethical commitment: to ensure that digital support enhances trust, preserves student dignity, and remains accountable to the institutional and pedagogical environments in which it operates. Within this perspective, innovation is meaningful only insofar as it remains compatible with fairness, interpretability, and respect for the learner as a participant in an educational relationship rather than as an object of data extraction or behavioural management.

This ethical orientation is not external to the architecture described in the previous sections. It is embedded in the very principles through which the system has been conceived. The commitment to grounded retrieval reflects a concern with informational integrity. The emphasis on longitudinal adaptivity reflects a rejection of rigid profiling and deterministic interpretation. The inclusion of human escalation pathways reflects the recognition that automated systems, however useful, cannot legitimately assume full authority in complex educational situations. In this sense, the ethical and European dimension of the HITS framework is not an appendix to the technical design. It is one of the conditions through which that design becomes pedagogically and institutionally acceptable.

The present section considers this dimension under five closely connected aspects: data protection and confidentiality; transparency and accountability; equity and bias mitigation; human oversight and professional development; and, finally, the question of European transferability in relation to institutional autonomy. Taken together, these aspects define the normative environment within which the HITS Chatbot can operate as a trustworthy digital support system.

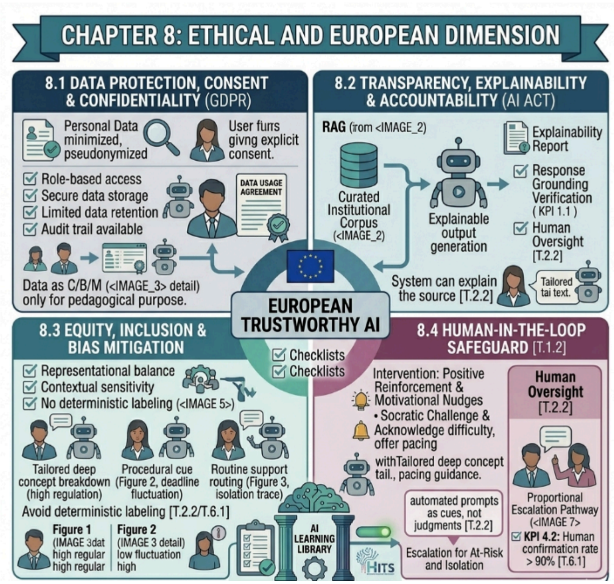

8.1 Data Protection, Consent, and Confidentiality

Respect for data protection is a structural condition of the HITS framework. Because the chatbot operates in relation to academic trajectories, support needs, and interaction traces, its legitimacy depends on a clear governance model regarding what data are processed, for what purposes, under what safeguards, and with what limits. Compliance with the General Data Protection Regulation is therefore not treated as a formal requirement to be addressed after design. It is part of the design logic itself.

Within this framework, data collection is restricted to what is necessary for academic support purposes. Interaction data, query traces, and related indicators are used insofar as they contribute to the functioning of the support system, the interpretation of recurring engagement patterns, and the responsible calibration of tutoring logic. Where possible, personal data are pseudonymized, and access is restricted according to clearly defined roles. Retention periods, governance responsibilities, and conditions of re-use must remain institutionally specified and subject to review. This is particularly important in a context where similarity analysis and adaptive intervention require the use of longitudinal traces. Such traces may be educationally useful, but they must remain governed by necessity, proportionality, and role-based accountability.

Consent, in this context, must also be understood substantively rather than procedurally. Students should not merely be informed that an AI-supported system exists; they should be able to understand in clear terms what kinds of support it provides, what categories of data are involved, and what the limits of the system are. They must also retain the possibility of declining participation or relying on alternative support pathways without academic disadvantage. Informed consent, therefore, is not a peripheral administrative step. It is one of the conditions under which digital tutoring can remain compatible with learner autonomy.

Confidentiality and data security complete this first ethical layer. Access to student-related information must remain limited to authorised institutional actors, and the technical environment must support secure storage, protected transmission, and regular review of security procedures. These measures do not merely ensure legal compliance. They contribute directly to the cultivation of trust, which is indispensable if students are to engage with the system without perceiving it as an opaque or intrusive mechanism.

8.2 Transparency, Explainability, and Accountability

A second ethical requirement concerns the intelligibility of the system. Digital support cannot claim pedagogical legitimacy if its operations remain entirely opaque to those who are affected by them. For this reason, the HITS framework gives particular importance to transparency and explainability, though always in forms appropriate to educational practice rather than in a narrowly technical sense. The aim is not to expose every computational layer of the system in formal detail, but to ensure that students, tutors, and institutional actors can understand how support is generated, what kinds of sources it draws upon, and where the limits of automated assistance lie.

The architectural use of a curated institutional corpus contributes directly to this form of transparency. Because responses are grounded in identifiable documentary sources, it becomes possible to trace the basis of institutional guidance, verify the relevance of retrieved material, and audit the system's informational behaviour. This does not eliminate the need for interpretation, but it ensures that support is not detached from institutional accountability. In practical terms, explainability is strengthened when the system can indicate the kind of source on which a response is based and when institutional staff can review or refine the corpus from which guidance is drawn.

Accountability, however, extends beyond source traceability. It also concerns the allocation of responsibility. Automated signals, prompts, or escalations must not be treated as self-sufficient decisions. Institutional actors remain responsible for the interpretation of support-relevant patterns and for the consequences of any action taken on that basis. This is especially important where patterns of vulnerability, disengagement, or repeated uncertainty are concerned. The system can assist in identifying such patterns, but it cannot legitimately replace professional judgement in understanding their meaning. Accountability therefore requires a clear distinction between the system's role in structuring information and the human role in interpreting educational situations.

Regular review mechanisms are part of this broader accountability framework. Feedback from students, tutors, and institutional services should inform iterative refinement of the system, especially where misunderstandings, unintended effects, or ambiguous outputs emerge. Explainability, in this sense, is not a one-time feature of system communication. It is part of an ongoing institutional practice of making AI-supported tutoring intelligible, reviewable, and corrigible.

8.3 Equity, Inclusion, and Bias Mitigation

The HITS framework also treats equity as a design principle rather than as an afterthought. A digital support system may be technically reliable and pedagogically coherent, yet still reproduce or amplify inequality if it assumes that all students interact with technology, language, and institutional structures in the same way. The ethical legitimacy of the chatbot therefore depends in part on its ability to remain attentive to heterogeneity in learner circumstances, accessibility needs, linguistic practices, and institutional opportunity structures.

This attention to equity operates at several levels. At the level of interface and access, the system should remain usable by students with varying degrees of digital familiarity and, where relevant, should be compatible with accessibility principles and alternative forms of presentation. At the level of content, institutional materials included in the knowledge base should be reviewed with attention to representational balance and inclusiveness. At the level of analytics, retrieval and support mechanisms must be monitored for differential behaviour across contexts, especially when linguistic or cultural differences may affect the interpretation of student queries or engagement signals.

Bias mitigation, accordingly, cannot be reduced to a purely technical problem. It involves reviewing the assumptions embedded in the corpus, in the interpretation of behavioural indicators, and in the thresholds that trigger different forms of support. A pattern that appears to signal disengagement in one context may carry a different meaning in another. What matters, therefore, is not the illusion of a perfectly neutral system, but the presence of procedures through which asymmetries and unintended effects can be identified and addressed. This is why fairness review should be integrated into regular validation cycles rather than reserved for exceptional cases.

In this perspective, learner feedback is also ethically significant. Students should not be treated only as subjects of system interpretation, but as participants in the ongoing refinement of the support environment. Their experience of the system—especially where it concerns relevance, fairness, intelligibility, and perceived respect—provides an indispensable source of insight into whether digital tutoring is functioning as intended. Equity is therefore sustained not only through design choices, but also through responsiveness to the lived experience of those who use the system.

8.4 Human Oversight and Professional Development

A further ethical condition of the HITS framework lies in the preservation of human oversight. The chatbot operates as a support instrument, not as a substitute for pedagogical or institutional responsibility. This principle has practical implications throughout the system. Complex cases, sustained patterns of vulnerability, or ambiguous forms of difficulty must remain subject to human interpretation. The role of the system is to assist in identifying, structuring, and, where appropriate, routing support needs. It is not to decide unilaterally what those needs mean or what consequences should follow from them.

Human oversight is important not because digital systems are intrinsically untrustworthy, but because educational situations are irreducibly contextual. Personal circumstances, emotional states, competing obligations, and institutional nuance cannot be fully captured through interaction traces alone. Tutors and support professionals remain essential precisely because they can place signals within a relational and dialogical frame. In this way, human oversight acts as both a pedagogical and an ethical safeguard: it protects learners from being reduced to data patterns, and it protects institutions from allowing procedural automation to displace professional judgement.

This principle also has implications for staff development. If AI-supported tutoring is to remain responsible, those who work alongside such systems must be supported in developing the relevant forms of professional literacy. This includes familiarity with AI ethics, data governance, limits of automated interpretation, and appropriate uses of system outputs in academic support contexts. Professional development is therefore not ancillary to implementation. It is part of what makes implementation ethically sustainable. A digitally assisted tutoring environment remains educationally sound only if the professionals who inhabit it are able to engage with it critically and reflectively.

8.5 European Transferability and Institutional Autonomy

The ethical framework of HITS is also connected to the question of transferability across European higher education. The chatbot was conceived within a project context that spans different institutional environments, regulatory expectations, and educational cultures. Any claim to transferability must therefore avoid two opposite risks:

excessive abstraction, which would make the system indifferent to local context, and excessive localism, which would make the framework impossible to reuse beyond its original setting. The modular architecture of HITS responds to this challenge by combining a stable pedagogical logic with institutional adaptability.

The system does not impose a single documentary corpus, language regime, or operational infrastructure across institutions. Each university can integrate its own regulations, curricular structures, and institutional resources into the knowledge environment without altering the underlying logic of grounded retrieval, adaptive support, and human oversight. This flexibility is ethically significant because it protects institutional autonomy. Transferability is achieved not through content uniformity, but through a shared method for transforming local institutional knowledge into a trustworthy support environment.

This orientation is fully consistent with wider European priorities in the field of digital education and educational AI, where responsible innovation is increasingly understood in terms of accountability, contextual sensitivity, proportionality, and respect for diversity of implementation. Within this landscape, the HITS framework seeks to contribute not only a functional tool, but also a methodologically and ethically defensible model of how AI may be integrated into academic support. The emphasis on modularity, governance, and institutional customization ensures that transferability does not come at the cost of pedagogical appropriateness or legal compatibility.

Conclusion

The ethical and European dimension of the HITS Chatbot cannot be separated from its pedagogical and technical design. Data protection, transparency, fairness, human oversight, and modular transferability are not external constraints placed upon the system after its development. They are conditions through which the system becomes educationally legitimate and institutionally trustworthy. In this sense, the ethical framework presented in this section does not stand apart from the earlier discussion of similarity, tutoring, architecture, and validation. It completes it. The HITS Chatbot is intended not simply as an efficient digital support tool, but as a form of responsible educational infrastructure: one capable of combining analytical precision, pedagogical care, and ethical accountability within diverse European higher education settings.