Project Acronym: HITS

TITLE: High Intelligent Tutoring System for Academic Success

Project Number: 000152340

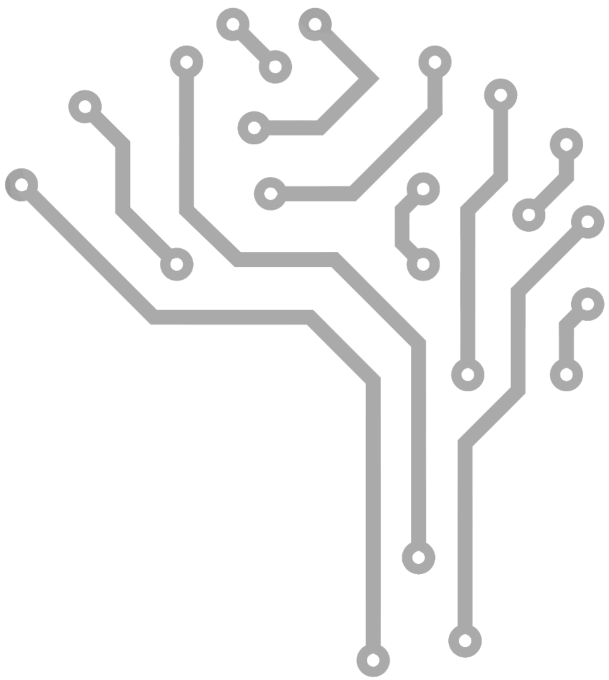

DELIVERABLE 5.2:

Identifying Patterns Related to Digital Academic Commitment and Awareness

Description: Report on set of patterns specifically related to digital academic commitment and awareness interventions

Lead Party for deliverable: CENTRO STUDI PLURIVERSUM SRL

Document type: Report

Due date of deliverable: 31/03/2026

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

1. Purpose of the Deliverable

This deliverable defines the conceptual and methodological framework through which the HITS project interprets **digital academic commitment** and **digital awareness** in AI-supported higher education. Its purpose is not limited to proposing new terminology. Rather, it seeks to clarify how meaningful forms of student engagement in digitally mediated learning environments can be identified, interpreted, and later operationalised in ways that are pedagogically relevant, ethically responsible, and technically credible. In this sense, Deliverable T.5.2 occupies a strategic position within WP5: it provides the conceptual bridge between behavioural interpretation, support calibration, and the later development of predictive and adaptive tutoring mechanisms.

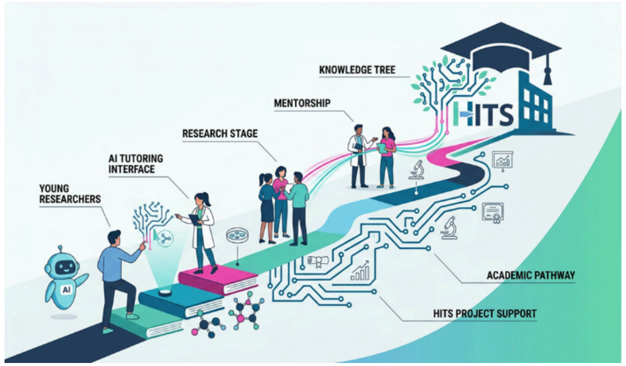

A central aim of the deliverable is to move beyond reductive understandings of digital participation based solely on activity traces or platform presence. The framework developed here starts from the assumption that logging in, clicking, or accessing resources does not in itself constitute meaningful academic engagement. For this reason, the document distinguishes **digital academic commitment** from simple platform usage and defines commitment as a more sustained, proactive, and structured relation to study. This distinction is foundational to the entire argument of the deliverable and was repeatedly recognised by partners as one of its most significant strengths, precisely because it protects the project from misleading interpretations of digital behaviour.

Alongside commitment, the deliverable introduces **digital awareness** as a second essential dimension of engagement. Awareness is understood here not as a purely technical competence, but as a reflective and critical disposition in relation to AI-mediated academic support. It concerns the learner's capacity to recognise how digital systems operate, what kinds of limits and assumptions they involve, how institutional grounding shapes their reliability, and when automated support should be complemented by human interpretation. Through this move, the framework avoids reducing student engagement to behavioural regularity alone. It instead situates participation within a wider field of critical and responsible interaction with digital infrastructures. This reflective dimension was also consistently acknowledged in the partner feedback as one of the distinctive conceptual contributions of the deliverable.

The two constructs are brought together through the notion of **trust**, which the deliverable treats as the catalyst linking technical reliability to meaningful academic behaviour. Within the HITS framework, trust is not regarded as a vague attitude of preference or familiarity. It is a structured condition that emerges when digital support is experienced as accurate, institutionally grounded, contextually relevant, and traceable to legitimate sources. Only under such conditions can AI-supported tutoring become educationally meaningful rather than merely technically functional. The centrality of this trust framework was strongly validated across the questionnaires, where partners repeatedly recognised it as one of the most convincing and coherent aspects of the document.

A further contribution of T.5.2 lies in the fact that it does not remain at the level of conceptual definition. It also proposes a formal interpretive model through which commitment and awareness can be translated into measurable and revisable analytical dimensions. This includes the development of a **four-pattern typology**, the specification of a **measurement model** based on explicit indicators, and the integration of a **longitudinal recalibration logic** designed to avoid static learner representations. The purpose of this model is not to produce fixed student categories or automated rankings. It is to support the calibration of educational support in a way that remains probabilistic, contextual, and open to revision as evidence accumulates over time. This non-deterministic and ethically restrained orientation was among the elements most positively received by the partners.

For this reason, the framework presented in the deliverable should not be interpreted as a classificatory system in the traditional sense. The patterns it identifies are intended as **diagnostic and support-oriented constructs**, not as stable identities attributed to learners. Their role is to make visible recurring configurations of commitment and awareness so that tutoring interventions, nudging strategies, and predictive mechanisms can be designed more responsibly. Final pedagogical judgement, however, remains with human actors. The deliverable therefore explicitly situates its analytical model within a **human-in-the-loop** logic, in which AI-supported pattern identification informs educational action without replacing professional interpretation. This point, too, emerged clearly in the partner feedback, especially in the request to make the human-supervised nature of the framework more visible from the outset.

Within WP5, T.5.2 performs a bridging role of particular importance. It connects the earlier work on behavioural profiling and tutoring logic with the later tasks devoted to adaptive nudging, pattern-based support calibration, and predictive AI development. In doing so, it gives conceptual coherence to the wider HITS architecture by showing how behavioural constructs, ethical safeguards, trust formation, and measurement logic can be brought into a single interpretive framework. The document therefore contributes not only to the theoretical advancement of the work package, but also to its methodological integration. Several partners explicitly recognised this bridging function and valued the way the deliverable clarifies its relation to the broader workflow of the AI Learning Library.

The sections that follow develop this argument progressively. Section 2 defines trust as the behavioural catalyst that transforms technical reliability into educationally meaningful support. Section 3 clarifies the notion of digital academic commitment, while Section 4 examines digital awareness as its reflective counterpart. Section 5 introduces the four-pattern typology derived from the interaction of these dimensions, and Section 6 presents the measurement model through which the framework becomes analytically operational. Section 7 addresses longitudinal recalibration and the ethics of educational analytics, while later sections connect these elements to architecture, safeguards, and implementation logic. Through this structure, the deliverable seeks to show that digital engagement in AI-supported higher education must be understood not merely as interaction with a system, but as a complex educational relation shaped by behaviour, reflection, trust, and institutional mediation.

Finally, it should be stressed that the framework developed in T.5.2 is intended as analytically grounded but empirically revisable. The constructs, indicators, and pattern logic proposed here form a basis for subsequent implementation, pilot validation, and contextual refinement within WP5. Their value lies not in fixing a closed taxonomy of student engagement, but in offering an interpretable and ethically coherent model through which

AI-supported tutoring can become more responsive, more accountable, and more educationally meaningful across different institutional contexts.

2. Conceptual Framework: Trust as the Behavioural Catalyst

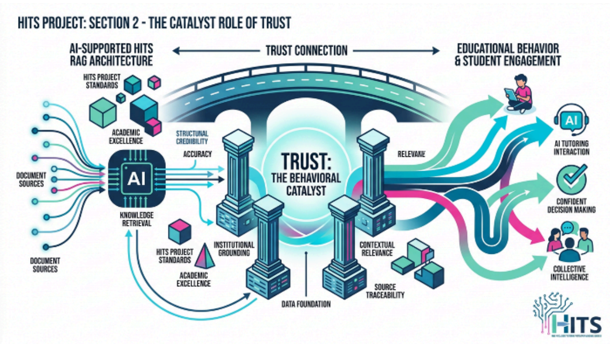

A central argument of this deliverable is that digital academic commitment cannot be understood solely through observable interaction traces, nor can it be adequately supported by technical functionality alone. Between the existence of a digital support system and the student's willingness to rely on it lies a mediating condition that is both behavioural and relational: **trust**. Within the HITS framework, trust is not treated as a vague disposition of preference or familiarity. It is understood as the condition through which technically reliable support becomes pedagogically credible and behaviourally consequential. Without trust, digital systems may be available, accurate, and even sophisticated, yet still fail to enter meaningfully into students' academic decision-making. With trust, by contrast, digital support can become part of the learner's working environment in ways that sustain continuity, reduce friction, and support more deliberate engagement.

This framing is especially important in the context of AI-supported tutoring. Students do not encounter such systems as abstract technical artefacts. They encounter them as sources of guidance in situations that may involve uncertainty, pressure, planning difficulty, or the need for procedural clarification. In such contexts, what matters is not only whether a system can produce a plausible answer, but whether that answer is perceived as dependable, contextually appropriate, and legitimately connected to the academic environment in which the student is acting. The concept of trust therefore allows the framework to connect architecture and behaviour without reducing either one to the other. It explains why technical reliability matters educationally, and why educational uptake cannot be assumed on the basis of technical performance alone.

Within the deliverable, trust is defined through four interrelated dimensions: **accuracy**, **institutional grounding**, **contextual relevance**, and **source traceability**. These dimensions do not operate independently; together, they describe the conditions under which students may come to experience a digital support system as worthy of reliance. Accuracy concerns the correctness of the information provided. Institutional grounding concerns the extent to which that information is anchored in the actual documentary and procedural environment of the university. Contextual relevance concerns whether the guidance offered corresponds to the learner's concrete academic situation rather than remaining generic or decontextualised. Source traceability concerns whether the origin of that guidance can be made visible, checked, and, where necessary, challenged or verified. Several partners explicitly highlighted the clarity and strength of this four-dimensional treatment of trust, confirming that it provides one of the most solid foundations of the deliverable.

Among these dimensions, **institutional grounding** is particularly important within HITS because it links the behavioural concept of trust to the system's technical architecture. In the HITS framework, trust is not secured by rhetorical assurance, but by design. The use of a Retrieval-Augmented Generation architecture means that the system does not rely on open-ended generation detached from institutional context. Instead, responses are constrained by a curated body of official academic sources. In this sense, institutional grounding is not merely a conceptual aspiration. It is operationalised through the architecture itself. The student can rely on the system not because it sounds authoritative, but because its guidance is derived from documentary environments that belong to the institution and can be reviewed within it. This makes trust structurally linked to the RAG logic and explains why source traceability and contextual relevance are so central to the educational use of the system.

A simple example may help clarify this point. A student asking whether a course prerequisite has been met, whether an assessment rule applies, or how a deadline should be interpreted is not simply looking for an answer in the abstract. The student is looking for guidance that can be acted upon within the university's actual academic framework. An answer becomes trustworthy when it is not only linguistically plausible, but recognisably connected to institutional sources and appropriate to the student's context. In this sense, source-groundedness reduces not only the technical risk of hallucination, but also the educational risk of hesitation, confusion, or misplaced reliance. Trust thus affects behaviour because it changes whether and how students are willing to act on digital support.

The role attributed to trust in this deliverable is further clarified through its connection with the **COM-B model**. This connection was strongly validated by partners, who consistently judged it clear and convincing. Within COM-B, behaviour is understood as emerging through the interaction of **capability**, **opportunity**, and **motivation**. Trust contributes to all three dimensions at once. It supports capability insofar as reliable guidance helps students interpret information and act with greater clarity. It supports opportunity insofar as institutionally grounded and contextually relevant digital support reduces barriers to access and navigation. It supports motivation insofar as confidence in the support environment increases the likelihood that students will continue to engage with it over time rather than withdraw from it. Trust, in this sense, is not external to the behavioural model. It is one of the conditions through which digital support becomes behaviourally effective.

This interpretation is particularly important because it prevents a simplistic reading of AI in education. The framework does not assume that the mere presence of a technically advanced tool will automatically improve engagement. Nor does it assume that students' willingness to use the system can be explained by convenience alone. Rather, it suggests that trust transforms availability into reliance. A student may have access to the chatbot, but still avoid using it if the system appears generic, opaque, or weakly grounded. Conversely, when the system is experienced as accurate, relevant, and institutionally anchored, it may become part of the learner's repertoire of legitimate academic supports. Some partners suggested that this behavioural transformation could be made even more tangible through practical examples or through minor dialogue with literature on AI literacy, self-regulated learning, or technology acceptance. Such suggestions are compatible with the present framework and reinforce its educational relevance without requiring any change to its core argument.

At the same time, the framework remains careful not to treat trust as a naïve or totalising concept. Trust does not mean uncritical acceptance of AI outputs. On the contrary, within HITS, trust is linked to awareness, interpretability, and the possibility of verification. A trustworthy system is not one that asks for passive compliance, but one that makes grounded use possible while leaving room for reflection, confirmation, and, where appropriate, human intervention. This is why the section on trust must be read together with the later sections on digital awareness and ethical safeguards. Trust, within the logic of the deliverable, supports reliance without suspending judgement. It enables engagement, but does not replace responsibility.

The conceptual importance of trust within T.5.2 therefore lies in its mediating role. It connects technical reliability to educational use; it connects RAG architecture to academic behaviour; and it prepares the transition from abstract system design to the more specific constructs of commitment and awareness. For this reason, the notion of trust does not remain confined to one chapter of the deliverable. It functions as a transversal condition of the whole framework. The following sections develop this more fully by showing how trust operates within the behavioural construct of digital academic commitment and within the reflective construct of digital awareness.

3. Defining Digital Academic Commitment

A central objective of this deliverable is to define **Digital Academic Commitment** in a way that is educationally meaningful, analytically robust, and clearly distinguishable from simple platform activity. This distinction is essential. In digitally mediated higher education, visible interaction with tools and environments can easily be mistaken for genuine engagement. Yet frequency of access, number of clicks, or duration of presence on a platform do not in themselves reveal whether a student is studying with continuity, acting with purpose, or integrating digital support into an effective academic trajectory. For this reason, Digital Academic Commitment is defined here not as presence in a system, but as a **sustained, proactive, and structured form of academic engagement mediated through digital environments**.

This definition implies that commitment must be understood as a multi-layered construct. It is not exhausted by behaviour alone, nor can it be reduced to intention or attitude in the abstract. Within the HITS framework, Digital Academic Commitment is analysed through three interrelated layers: **behavioural**, **cognitive**, and **relational**. These layers are distinct for analytical purposes, but they acquire their full meaning only in combination. Together, they make it possible to identify whether a student's use of digital academic support reflects episodic activity or a more stable and educationally significant pattern of engagement. Partners consistently recognised this three-layer structure as one of the strengths of the deliverable, precisely because it allows the concept of commitment to remain richer and more discriminating than a simple usage metric.

The behavioural layer concerns the observable organisation of digital engagement over time. It includes such aspects as regularity of interaction, continuity across the academic term, responsiveness to prompts, and relation to deadlines or institutional milestones. Behaviour is relevant because commitment necessarily leaves traces in action. Students who engage consistently with relevant resources, return to support tools at appropriate times, and respond to structured reminders in ways that sustain continuity display a different form of academic relation from those whose activity is sporadic, last-minute, or purely reactive. At the same time, the interpretation of behavioural indicators must remain cautious. Not every reduction in visible activity is a sign of disengagement. In some cases, temporary absence may reflect off-platform study, deliberate reflection, or a strategic pause within a broader pattern of sustained commitment. Behavioural traces are therefore meaningful only when read longitudinally and in relation to other dimensions of engagement. This refinement is particularly important in light of partner feedback warning against overly narrow readings of visible digital activity.

The cognitive layer concerns the extent to which digital engagement is connected to purposeful academic work. This includes not only whether students access resources, but also how they use them, how deeply they engage with them, and whether their interaction suggests orientation toward understanding rather than mere procedural compliance. A student who consults materials selectively but in close relation to academic tasks, formulates queries that reveal effort toward clarification, or returns to content in a structured way may show a stronger form of commitment than one who accesses many resources superficially or only under immediate pressure. The cognitive layer therefore allows the framework to distinguish between engagement that is academically generative and activity that remains thin, fragmented, or merely performative. This is one of the reasons why partners regarded the distinction between commitment and simple usage as so valuable: the model avoids equating intensity of activity with quality of engagement.

The relational layer introduces a further and indispensable dimension. Digital academic commitment is not only a matter of what students do or how they process information. It also concerns whether they are prepared to **rely** on the support environment as part of their academic action. In the HITS framework, this reliance is closely connected to trust. A student may access a system frequently and still not truly incorporate it into academic decision-making if the system is perceived as generic, opaque, or weakly grounded. Conversely, a student may interact less often but in ways that indicate confidence in the system's reliability and relevance. Commitment, in this sense, includes a relational disposition toward digital support: a readiness to treat the system as a legitimate aid within the academic environment. Several partners validated this connection strongly, particularly where the deliverable links commitment to trust as reliance. At the same time, partner feedback suggests that this relational layer becomes more convincing when it explicitly reflects the hybrid nature of tutoring. Commitment may involve trust both in the AI-supported system itself and in the human-mediated educational feedback through which that system is interpreted, contextualised, or, where necessary, overridden.

A simple contrast may help illustrate the value of this three-layer framework. A student may log into a learning environment repeatedly during the days preceding an assessment, consult multiple pages briefly, and pose only minimal procedural questions. From a surface perspective, the student appears active. Yet such activity may reflect urgency rather than structured engagement, and it may have limited cognitive depth or relational stability. Another student may interact less frequently, but do so in a regular and purposeful way: consulting relevant materials in advance, asking clarifying questions, and returning to the system when decisions need to be made. The second case represents a stronger form of commitment, even if raw usage metrics are lower. This is precisely why the deliverable insists on a richer conceptualisation of engagement. Digital Academic Commitment must capture not just visible activity, but the quality, structure, and educational significance of that activity.

For analytical and operational purposes, the framework therefore treats commitment as a composite construct supported by multiple indicators rather than by a single behavioural proxy. The indicators associated with the behavioural layer help illuminate continuity, stability, and responsiveness. Those associated with the cognitive layer help illuminate depth of engagement and academic orientation. Those associated with the relational layer help illuminate whether students are treating the support environment as a trustworthy and meaningful part of their academic practice. The purpose of this measurement logic is not to produce a rigid score that defines the learner once and for all. It is to provide an interpretable basis on which recurring patterns of engagement may be identified and later related to awareness, typology, and support calibration. In this respect, commitment is both a conceptual and a methodological construct: it gives shape to the educational meaning of engagement while also providing the basis for its responsible interpretation.

It is equally important to emphasise what this construct is **not** intended to do. Digital Academic Commitment is not a proxy for merit, effort in a moral sense, or academic worth. Nor is it intended to classify students into immutable categories. The indicators associated with commitment are to be interpreted probabilistically and longitudinally, in line with the broader methodological position of the HITS framework. This means that temporary dips, unusual rhythms, or changes in visible activity should not be over-interpreted. Commitment is inferred through recurring patterns, not isolated traces. This is one of the reasons why the section should be read together with the later discussion of recalibration and ethics: the concept gains its full educational meaning only when coupled with revisability, caution, and human oversight.

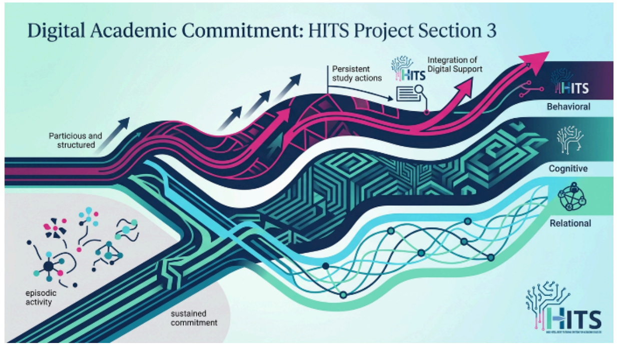

The definition developed here prepares the ground for the next step of the argument. Once commitment has been distinguished from simple activity and articulated across behavioural, cognitive, and relational layers, it becomes possible to examine the second major construct of the deliverable: **Digital Awareness**. If commitment concerns structured participation in digitally mediated academic work, awareness concerns the reflective and critical dimension through which such participation becomes more informed, responsible, and educationally autonomous. The two constructs are analytically distinct, but they are designed to function together. It is from their interaction that the pattern typology later presented in the deliverable becomes possible.

4. Defining Digital Awareness

If digital academic commitment concerns the structured and sustained way in which students engage with digitally mediated academic work, **Digital Awareness** concerns the quality of reflection that accompanies such engagement. Within the HITS framework, awareness is not treated as a marginal or secondary attribute. It is a constitutive dimension of meaningful participation in AI-supported higher education. A student may use digital tools regularly and still do so in an unreflective or overly dependent way. Conversely, a student may engage more selectively but with a clearer understanding of how the support system works, what its limits are, and how its outputs should be interpreted in relation to institutional knowledge and human guidance. For this reason, Digital Awareness is defined here as the **reflective, critical, and interpretive dimension of digital academic engagement**.

This definition is important because it prevents the framework from equating effective use of AI-supported tools with mere technical familiarity. Awareness, in the present context, does not simply mean knowing that a chatbot exists or being able to operate it at a functional level. It means understanding enough about the system's nature, conditions, and limitations to use it responsibly within one's academic work. This includes recognising how AI-supported answers are generated, being able to distinguish institutionally grounded guidance from generic output, and appreciating when human confirmation or pedagogical mediation is appropriate. Several partners identified this reflective conception of awareness as one of the major conceptual strengths of the deliverable, precisely because it connects digital engagement to critical judgment rather than to passive consumption.

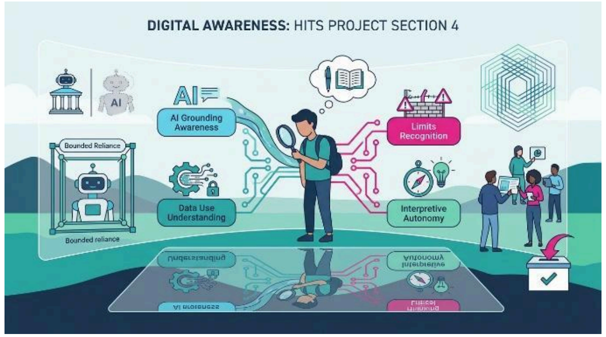

The framework developed in T.5.2 treats Digital Awareness as involving at least four closely related areas of understanding. The first concerns **awareness of AI grounding**. Students need not master the technical details of system architecture, but they should be able to recognise the difference between a tool that produces open-ended, generic responses and one that operates on the basis of curated institutional sources. In the HITS context, this distinction is particularly important because the educational value of the system depends precisely on the fact that it is grounded in an institutional corpus through a RAG-based architecture. Awareness, therefore, includes some understanding of why a source-grounded answer is different in status from a merely plausible one. This distinction between **generic AI** and **institution-grounded AI** was one of the most consistently praised elements in the validation questionnaires and should remain explicit in the final text.

The second area concerns **awareness of data use and interpretation**. Students engaging with an AI-supported tutoring system should be able, at least in broad terms, to understand that their interaction traces may contribute to the adaptation of support. This does not require technical expertise in analytics or modelling, but it does imply a degree of literacy about how digital systems can interpret recurring patterns of engagement. Such awareness is educationally relevant because it supports informed participation rather than opaque dependence. It helps students understand that the system is not merely answering isolated questions, but may be using patterns of interaction to calibrate prompts, detect support needs, or guide routing decisions within a broader tutoring logic. Awareness in this sense is closely linked to transparency and consent, and therefore also to the ethical legitimacy of the system.

The third area concerns **awareness of limits**. A digitally aware student is not simply one who trusts the system, but one who understands that trust must remain bounded by recognition of the system's scope. AI-supported tutoring may be reliable, well grounded, and educationally useful, yet still remain limited in what it can interpret, resolve, or decide. Students should therefore be able to recognise that not every academic or personal difficulty can or should be addressed through automated support alone. Some situations require human interpretation, contextual judgment, or pedagogical dialogue. Awareness of limits is thus not a sign of distrust in the negative sense. It is a condition of appropriate reliance. The more trustworthy a system is, the more important it becomes that students understand where its legitimate authority ends.

The fourth area concerns **awareness of one's own interpretive role**. In this respect, Digital Awareness is closely connected to autonomy and responsible digital citizenship. Students are not expected to behave as passive recipients of AI-generated guidance. They remain participants in an educational process that requires judgement, reflection, and, where necessary, verification. A digitally aware student is therefore one who is able not only to use the system, but also to question it, contextualise it, and decide when to rely on it more directly and when to seek additional confirmation. In a practical sense, this may involve recognising that a source-grounded procedural answer can be acted upon with relative confidence, while a more interpretive or ambiguous issue may require discussion with a tutor or academic advisor. Several partners explicitly validated the section's treatment of autonomy and responsible digital citizenship, confirming that awareness is correctly framed as a condition of reflective engagement rather than of passive compliance.

A simple illustration may help clarify the educational significance of this construct. Two students may both consult the chatbot to clarify an academic issue. The first accepts the answer immediately, without considering its grounding, its scope, or whether the question might require institutional or human confirmation. The second notices that the answer is based on institutional material, understands why this matters, and is also able to recognise when further clarification would be appropriate. Both students are using the same tool, but only the second demonstrates a stronger form of Digital Awareness. What distinguishes the two is not technical skill in the narrow sense, but the quality of reflection that accompanies use. Awareness, in this sense, transforms digital interaction into a more informed and responsible educational practice.

It is also important to note that Digital Awareness is not defined here as an individual moral virtue or a fixed competence that students either possess or lack. As with commitment, awareness should be understood as a dynamic and developable dimension of engagement. It may vary across time, academic contexts, and prior experiences with digital systems. Some students may display high awareness in relation to institutional source-grounding but remain less reflective about data use or system limits. Others may show intuitive caution without being able to articulate the reasons for it clearly. The framework therefore treats awareness not as a rigid literacy threshold, but as a structured interpretive construct that helps explain how students position themselves in relation to AI-supported tutoring environments. This dynamic treatment was appreciated by partners, especially where the deliverable links awareness to trust, AI literacy, and responsible educational use without hardening it into a simplistic competence model.

The analytical value of Digital Awareness becomes particularly clear when considered alongside Digital Academic Commitment. Commitment captures the extent to which students engage in a sustained and structured way; awareness captures the extent to which such engagement is accompanied by critical understanding and responsible interpretive use. The two constructs are therefore distinct but complementary. A student may be highly committed but weakly aware, relying intensively on digital support without sufficiently understanding its conditions or limits. Another may be highly aware but less consistently committed, showing reflection without yet translating it into sustained academic engagement. It is precisely from the interaction between these two dimensions that the pattern typology developed in the following section becomes possible.

For this reason, the section on Digital Awareness should not be read as an optional ethical appendix to the more “substantive” construct of commitment. It is part of the same analytical architecture. If commitment tells us how students engage, awareness tells us how they understand and govern that engagement in relation to AI-supported academic assistance. Together, the two constructs allow the HITS framework to move beyond behavioural observation alone and toward a more educationally adequate interpretation of digitally mediated academic support. The next section develops this logic by showing how different configurations of commitment and awareness give rise to the four interpretive patterns that structure the rest of the deliverable.

5. Pattern Typology

The interaction between **Digital Academic Commitment** and **Digital Awareness**, as defined in the preceding sections, makes it possible to identify a set of recurring interpretive patterns that are relevant for AI-supported tutoring. The purpose of this typology is not to produce a rigid classification of students, nor to reduce the complexity of academic engagement to a small number of fixed categories. Rather, it offers a structured heuristic through which different combinations of sustained engagement and reflective awareness may be interpreted in educationally meaningful ways. Within the HITS framework, the typology serves as an intermediate layer between conceptual analysis and pedagogical calibration. It helps translate abstract constructs into a form that can support differentiated tutoring, measurement, and later predictive modelling, while remaining open to recalibration as evidence accumulates.

The four patterns are derived from the intersection of two analytically distinct dimensions. The first concerns the degree of commitment students show in their relation to digital academic work: whether engagement is structured, sustained, and proactive, or whether it remains partial, unstable, or weakly integrated into study practices. The second concerns the degree of awareness students display in relation to the digital and AI-supported environment itself: whether their engagement is accompanied by reflective understanding of system grounding, limits, interpretive responsibility, and responsible use. When these two dimensions are considered together, four broad configurations emerge:

- **high commitment–high awareness**,
- **high commitment–low awareness**,
- **low commitment–high awareness**,
- **low commitment–low awareness**.

Partners consistently found this structure clear and pedagogically meaningful, which suggests that the typology succeeds as an interpretive instrument.

At the same time, the typology must be read with caution. It is not intended as a closed taxonomy of learner types. Its categories are probabilistic and dynamic, not deterministic and permanent. A student may move between patterns over time, may display features of more than one pattern simultaneously, or may deviate from the typology in ways that remain educationally significant. The value of the framework lies precisely in its provisional nature: it allows recurring configurations of engagement to become visible without turning those configurations into fixed identities. This non-deterministic and revisable logic was explicitly recognised by partners as one of the strengths of the deliverable and should remain central to how the typology is presented.

5.1 Pattern A: High Commitment – High Awareness

The first pattern combines sustained academic commitment with a high degree of digital awareness. Students who fall predominantly within this configuration tend to engage with digital academic support in a regular, purposeful, and reflective way. Their interaction with AI-supported tools is not merely frequent or efficient; it is also informed by an understanding of how such systems operate, what kind of grounding supports their answers, and when human confirmation or institutional interpretation may still be necessary. In this sense, Pattern A represents the most educationally integrated form of digital engagement within the framework.

The pedagogical value of this pattern lies in its balance. These students are not only active; they are capable of using digital support critically and autonomously. Their reliance on the system is accompanied by reflective judgement rather than passive dependence. Such learners are likely to benefit from digital tutoring environments because they can integrate them into their academic work without overestimating their scope. Trust, in this pattern, is neither absent nor uncritical. It takes the form of informed reliance.

Pattern A should not be romanticised as an ideal type to which all students must conform. Its importance lies in showing that sustained engagement and critical awareness can reinforce one another. The pattern illustrates what it means for digital support to be educationally appropriated rather than simply used. At the same time, the framework must remain cautious: even high-commitment, high-awareness students may encounter contextual difficulties, fluctuating demands, or discipline-specific challenges that alter how their engagement is expressed over time. The pattern therefore serves as an interpretive guide, not as a stable achievement status.

5.2 Pattern B: High Commitment – Low Awareness

The second pattern combines strong or sustained engagement with limited digital awareness. Students in this configuration may appear highly active and academically committed in behavioural and cognitive terms, yet show relatively low reflection on how AI-supported systems operate, what their limits are, or how their outputs should be interpreted critically. They may rely extensively on the system, respond regularly to prompts, use resources intensively, and incorporate digital tools into their study practices, while remaining insufficiently attentive to the conditions that make those tools trustworthy or to the points at which their outputs require verification or contextual mediation.

This pattern is particularly important because it reveals that engagement alone does not guarantee responsible educational use. A student may appear well integrated into the digital support environment and still interact with it in ways that remain overly dependent, insufficiently critical, or weakly aware of the distinction between grounded and generic AI support. In practical terms, this may mean treating the system as authoritative without fully understanding why some answers are more reliable than others, or failing to recognise when automated outputs should be supplemented by institutional or human interpretation.

Pedagogically, Pattern B is significant because it identifies a form of digital participation that is strong in continuity but weaker in reflexivity. Within the HITS framework, such a configuration suggests the need not for more activation, but for stronger cultivation of awareness. It points to the importance of helping students develop a more explicit understanding of AI grounding, source traceability, limits of automation, and the hybrid nature of support. Partners found this pattern meaningful precisely because it avoids the simplistic assumption that "more use" necessarily indicates a better or more mature relation to digital support.

5.3 Pattern C: Low Commitment – High Awareness

The third pattern presents a different kind of asymmetry. Here, students display relatively strong awareness of how AI-supported systems should be used, yet their academic commitment remains unstable, limited, or insufficiently translated into sustained action. In other words, they may understand the conditions of trustworthy and responsible use, appreciate the importance of institutional grounding, or recognise the limits of automation, while nevertheless failing to engage with academic work in a consistent and structured way.

This configuration is analytically valuable because it shows that awareness and action do not always develop in parallel. A student may be reflective, cautious, and relatively literate in relation to digital systems, yet still struggle with continuity, timing, or structured academic participation. Such a learner may know that the system is useful, understand why certain forms of support are more reliable than others, and still not manage to convert that understanding into proactive engagement. In this respect, Pattern C highlights a gap between interpretive competence and enacted commitment.

Pedagogically, this pattern is important because it challenges the idea that increased awareness automatically leads to improved academic behaviour. It suggests instead that reflective understanding must sometimes be accompanied by other forms of support—such as structured pacing, motivational reinforcement, or reduced friction in task initiation—if it is to become behaviourally effective. The pattern is therefore not a sign of contradiction in the framework, but evidence of its usefulness: it captures the fact that students may be critically aware without yet being fully able to sustain the rhythms of academic engagement that support requires.

5.4 Pattern D: Low Commitment – Low Awareness

The fourth pattern combines low levels of academic commitment with limited digital awareness. Within the framework, this is the configuration most closely associated with vulnerability in the context of AI-supported tutoring. Students in this pattern may show unstable or weak engagement with digital academic support while also lacking the reflective understanding needed to interpret such support critically, use it effectively, or recognise when additional guidance is required. Activity may be sparse, reactive, fragmented, or weakly integrated into academic practice, while awareness of source grounding, system limits, or the role of human mediation remains limited.

This pattern must be treated with particular care. Its analytical purpose is not to identify a group of "poor users" or to isolate students for negative judgement. On the contrary, it is precisely in this area that the ethical commitments of the HITS framework become most important. Pattern D does not denote failure in an essential sense. It signals a configuration in which both support uptake and reflective orientation appear weakened, and in which educational intervention may need to be especially sensitive, non-stigmatising, and proportionate. The wording used to describe this pattern should therefore remain cautious and student-centred, as several partners explicitly noted in their validation of the typology.

The pedagogical significance of Pattern D lies less in its descriptive precision than in its practical implications for support calibration. It points to the likelihood that both engagement-oriented and awareness-oriented interventions may be needed together: not only more accessible or timely support, but also help in understanding what the system is, how it should be used, and why it should not be interpreted as a self-sufficient authority. This pattern also highlights the importance of human oversight, since learners in this configuration may be least well served by automated support when used in isolation. Here, more than elsewhere, digital analytics must remain the beginning of pedagogical attention rather than its conclusion.

5.5 Interpretive Value and Limits of the Typology

Taken together, the four patterns provide a framework through which the interaction of commitment and awareness can be interpreted in ways that are relevant to AI-supported academic tutoring. Their value lies not in simplifying student diversity into a closed grid, but in making certain recurrent configurations of engagement visible enough to inform later support logic. They offer a structured vocabulary for discussing how behaviour, reflection, trust, and use of digital support may combine differently across learners and across time.

At the same time, the limits of the typology must remain explicit. The patterns are heuristic, not exhaustive; dynamic, not static; probabilistic, not deterministic. They support interpretation and calibration, but they do not define students permanently. This point becomes especially important when the framework is transferred across institutions, disciplines, and linguistic or cultural contexts. One partner rightly noted that the typology, however useful, may still oversimplify the diversity of actual student behaviour in varied educational environments. This observation should not be read as a weakness of the framework, but as a reminder of the conditions under which it should be used responsibly. A typology becomes educationally useful not when it is treated as universally complete, but when it remains open to revision in light of contextual evidence.

Conclusion

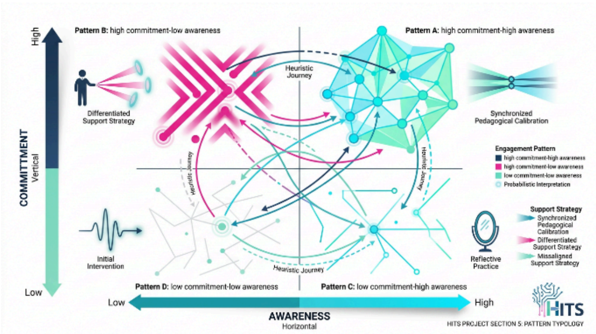

The pattern typology developed here translates the conceptual relation between commitment and awareness into a structured interpretive model. In doing so, it prepares the ground for the next stage of the deliverable: the definition of a measurement logic through which these dimensions and patterns can be operationalised in a transparent, revisable, and ethically bounded way. If the present section identifies the recurrent configurations of engagement that matter educationally, the next section addresses how those configurations may be measured without collapsing them into fixed scores or simplistic metrics.

6. Measurement Model

The analytical framework developed in the preceding sections requires a corresponding measurement logic if it is to support later clustering, tutoring calibration, and predictive modelling within WP5. For this reason, Deliverable T.5.2 does not stop at conceptual definition. It also proposes a **measurement model** through which Digital Academic Commitment and Digital Awareness can be translated into interpretable and operational indicators. The purpose of this model is not to reduce complex educational phenomena to a single score, nor to substitute measurement for pedagogical judgement. Rather, it is to provide a structured and transparent basis for identifying recurring patterns of engagement in ways that remain comparable, revisable, and educationally meaningful.

This measurement model is one of the most distinctive contributions of the deliverable and was explicitly recognised by partners as a major strength. What appears to have been especially valued is the fact that the model combines conceptual coherence with operational feasibility. It does not rely on one-dimensional proxies or opaque scoring logics. Instead, it proposes a multi-indicator approach in which the constructs of commitment and awareness are supported by distinct but related components. In this respect, the model reflects the broader methodological position of the HITS framework: educationally significant patterns must be inferred from multiple traces, interpreted probabilistically, and kept open to recalibration rather than fixed through rigid classificatory thresholds.

At the highest level, the model is organised around two composite dimensions: a **Commitment Index** and an **Awareness Index**. These do not function as isolated variables. They are analytical devices designed to capture, in a structured way, the recurring traces through which sustained academic engagement and reflective digital use become visible. The indices do not claim to measure an inner state directly. They provide an interpretable approximation of how such states may be inferred through patterns of behaviour, inquiry, interaction, and reflective orientation. This distinction is important. The model is not epistemologically naïve: it does not assume that educational commitment or awareness can be read off directly from data. It assumes instead that certain traces, when combined and interpreted cautiously, can support a more reliable understanding of how students relate to AI-supported academic environments.

6.1 The Commitment Index

The **Commitment Index** is intended to capture the structured, sustained, and proactive quality of digital academic engagement. In line with the conceptual framework outlined earlier, it does not treat activity alone as evidence of commitment. Instead, it brings together several components that make it possible to distinguish episodic usage from more meaningful patterns of study-related participation. The index therefore reflects the conviction that commitment is visible not only in the quantity of interaction, but in its organisation, continuity, and educational purpose.

Within the model, three kinds of indicators play a particularly important role. The first concerns **engagement stability**, that is, the extent to which interaction with the support environment is distributed across time in a relatively sustained way rather than concentrated only at moments of pressure. The second concerns **response-to-nudge ratio**, which helps illuminate whether support prompts or reminders are followed by meaningful action. The third concerns **resource depth index**, which points to the degree to which student interaction reflects engagement with academically relevant material rather than superficial or fragmented access. Together, these components provide a richer picture of commitment than any simple count of visits or clicks could offer.

At the same time, the interpretation of these indicators must remain careful. A period of reduced visible activity is not automatically evidence of weak commitment. Some students may study extensively off-platform, pause interaction strategically while processing material, or temporarily reduce their use of the system without withdrawing from academic work. This point, raised in the validation feedback, is particularly important because it protects the framework from reproducing the very simplifications it seeks to avoid. The Commitment Index is therefore not intended to reward constant visibility. It is intended to identify recurring forms of structured engagement through a combination of traces interpreted over time.

6.2 The Awareness Index

The **Awareness Index** is designed to capture the reflective and critical dimension of digital engagement. If the Commitment Index addresses how students engage with academic work in digitally mediated settings, the Awareness Index addresses how they understand and interpret the AI-supported environment within which that engagement takes place. As with commitment, the aim is not to measure awareness as an internal psychological state, but to identify observable indicators that reasonably suggest reflective use, source sensitivity, and understanding of system limits. The model proposes several components for this purpose. These include indicators related to **source verification**, **meta-queries**, and **ethical awareness responses**. Source verification concerns whether and how students attend to the grounding of the information they receive. Meta-queries concern whether students ask questions not only about academic content or procedure, but also about the system itself—its provenance, reliability, or operational limits. Ethical awareness responses concern traces that suggest attention to appropriate use, human confirmation, or the responsible interpretation of AI-supported guidance. Considered together, these indicators make it possible to approach awareness as a structured component of digital engagement rather than as an abstract normative aspiration.

Here too, interpretation must remain contextual and proportionate. Awareness does not necessarily manifest through explicit technical language, and students may show reflective caution without always verbalising it in formal ways. The index therefore should not be treated as a literacy test in the narrow sense. Its function is to illuminate whether students' interaction with digital academic support remains merely instrumental, or whether it includes some degree of reflection on grounding, limitations, and appropriate reliance. This is especially relevant because the deliverable explicitly distinguishes institution-grounded AI support from generic AI use and treats awareness as one of the conditions of responsible educational uptake.

6.3 Weighting Logic and Interpretive Transparency

A measurement model based on multiple indicators necessarily raises the question of weighting. If commitment and awareness are each represented through several components, on what basis are those components combined? The partner feedback makes clear that the weighting logic was generally found transparent and convincing, but also that a short note on its governance would strengthen the section further. This is a useful observation. A strong model should be not only analytically sound, but also institutionally intelligible.

Within the HITS framework, weighting should be understood as **empirically informed and institutionally governable**, not as fixed once and for all. Initial weighting may be established at project level in order to support comparability across contexts and to provide a coherent basis for pilot implementation. However, these initial values should remain open to recalibration in light of pilot evidence, institutional variation, linguistic differences, and observed educational relevance. Weighting is therefore not a hidden technical choice external to the framework. It is part of the model's interpretive governance and should be documented, reviewable, and, where necessary, adjustable.

This point matters because the educational legitimacy of the model depends partly on the transparency of its construction. If the relative importance assigned to certain indicators remains opaque, the framework risks appearing more authoritative than interpretable. By contrast, when weighting is presented as a structured but revisable choice, the model remains aligned with the broader methodological principle of the deliverable: patterns are inferred, not discovered as fixed truths; and inference must remain accountable to evidence and open to revision.

6.4 Probabilistic Pattern Classification

The indices developed in this section do not culminate in deterministic assignment to a fixed category. Instead, they support a **probabilistic pattern classification** in which students may be interpreted as showing a greater or lesser fit with one of the four patterns outlined in the previous section. This probabilistic logic is essential. It allows the framework to identify meaningful configurations of commitment and awareness without converting them into rigid identities. Partners consistently valued this feature and regarded it as one of the ethical strengths of the deliverable.

A probabilistic approach is appropriate for at least two reasons. First, academic engagement is dynamic. Students' patterns may shift across time, across courses, or in response to changing contextual pressures. Second, the indicators on which the model relies are themselves partial and interpretive. No set of traces can fully exhaust the complexity of a learner's academic situation. For these reasons, pattern classification should be treated as an informed estimate of the current configuration of engagement, not as a definitive representation of the learner. The system does not "discover" what a student is in any essential sense; it estimates how current traces align with one or another support-relevant pattern.

This also explains why probabilistic assignment must be read together with longitudinal recalibration and human oversight. A pattern estimate may become more or less reliable as additional evidence accumulates. Human tutors may also have contextual knowledge that qualifies or reinterprets the model's reading. The function of the classification, therefore, is not to close judgement but to support it. It provides an analytically grounded indication of likely configuration, always subject to refinement and pedagogical review.

6.5 Reproducibility, Annexes, and Operational Use

The measurement model is supported by a set of annexes that formalise its logic further. This supplementary material is not peripheral. Several partners explicitly praised the annexes for their clarity and usefulness, particularly in relation to clustering, predictive pipelines, and institutional reproducibility. Their function is to ensure that the model remains not only conceptually clear in the main text, but also sufficiently explicit for later technical implementation and cross-partner comparison.

For this reason, the main text should guide readers more clearly toward the annexes where the measurement model is elaborated in greater detail. The formalisation of indicator structure, standardisation logic, weighting assumptions, and calibration notes is part of what makes the framework usable beyond a purely conceptual level. At the same time, the main body of the deliverable should preserve the interpretive meaning of the model and avoid becoming excessively technical. The relation between the two levels is therefore complementary: the main text explains why the model matters educationally; the annexes explain how it may be implemented and reviewed more formally.

From an operational point of view, the model should be regarded as a structured but revisable scaffold for later WP5 work. It enables commitment and awareness to enter the next phases of clustering, nudging, and predictive development in an interpretable form, while preserving the ethical and methodological caution already built into the framework. This is one of the reasons why the measurement section was so positively received in the validation process: it gives the deliverable methodological weight without sacrificing educational meaning.

Conclusion

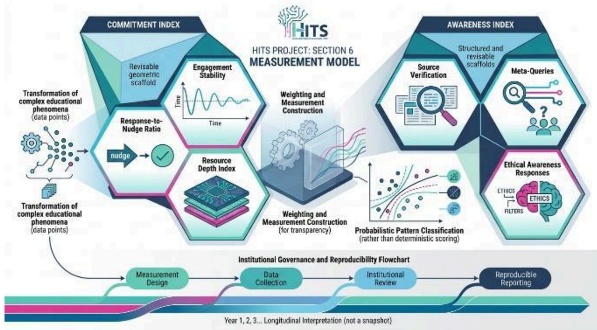

The measurement model presented in this section translates the conceptual architecture of commitment and awareness into an operational framework that remains interpretable, probabilistic, and open to revision. It does not reduce educational engagement to a mechanical score, nor does it seek to eliminate the need for contextual and human interpretation. Rather, it provides the structured basis through which support-relevant patterns can be inferred more responsibly and later used in a calibrated way within the HITS ecosystem. The next section develops one of the conditions that makes such use ethically legitimate: the principle of longitudinal recalibration and the wider ethics of educational analytics.

7. Longitudinal Recalibration and the Ethics of Educational Analytics

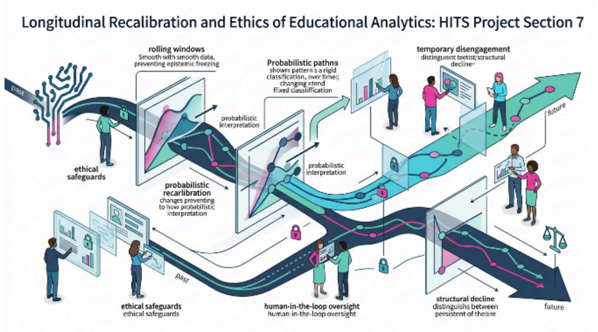

A framework that seeks to identify recurring patterns of digital academic commitment and awareness must also address a fundamental methodological and ethical question: **how should such patterns be interpreted over time**? This question is decisive because educational analytics can become harmful precisely when they transform provisional signals into apparently stable truths about the learner. If early traces are treated too quickly as evidence of enduring characteristics, then digital support risks reproducing fixed assumptions rather than remaining responsive to educational change. For this reason, the HITS framework places **longitudinal recalibration** at the centre of its approach to educational analytics. Pattern identification is not conceived as a one-time classificatory act, but as a revisable interpretive process governed by temporal evidence, contextual caution, and human oversight.

This principle is directly connected to the project's wider ethical concern with avoiding what the framework describes as **epistemic freezing**. Epistemic freezing occurs when data-derived representations of the learner cease to function as open and revisable interpretations and instead begin to operate as fixed identities. In educational settings, such a shift is especially problematic. Students' engagement patterns are rarely stable in a linear or permanent sense. They fluctuate across time, courses, deadlines, personal circumstances, and institutional conditions. A temporary decline in visible commitment may reflect overload, illness, competing responsibilities, or a phase of off-platform study rather than a deeper structural disengagement. Conversely, short periods of high activity may not necessarily indicate stable commitment. The ethical task of educational analytics is therefore not to make rapid judgement appear objective, but to preserve the interpretive openness required by the dynamic nature of learning.

Within the HITS framework, longitudinal recalibration serves precisely this purpose. The model does not treat a learner's current pattern estimate as definitive. It assumes instead that commitment and awareness are inferred through traces that may shift in meaning as additional evidence becomes available. Early estimates must therefore remain provisional, and the significance of observed indicators must be allowed to change over time. This principle is not only methodological; it is also pedagogical. It reflects the recognition that students are not static behavioural profiles but evolving participants in educational processes. Any analytic system that aims to support them must remain capable of reinterpreting its own earlier readings.

One way in which the framework operationalises this principle is through the use of **rolling temporal windows** and related smoothing logics. These mechanisms make it possible to interpret patterns across sequences of interaction rather than through isolated events. Their purpose is not to create statistical sophistication for its own sake, but to ensure that momentary deviations do not immediately trigger strong interpretive shifts. A missed prompt, a short period of reduced activity, or an isolated expression of doubt should not, on its own, lead to the conclusion that commitment has collapsed or that awareness has weakened structurally. What matters is recurrence, persistence, and the relation of one trace to others across time. Several partners explicitly validated the clarity and importance of this longitudinal logic, and one of the deliverable's key strengths lies precisely in the fact that it resists snapshot-based interpretation.

This also explains why the framework insists on distinguishing **temporary disengagement** from **structural decline**. The distinction is both analytically necessary and pedagogically consequential. Temporary disengagement refers to short-term fluctuations in behaviour that may be contextually explainable and reversible without requiring strong intervention. Structural decline, by contrast, refers to more persistent and corroborated weakening of engagement across indicators and time. The educational risk of failing to distinguish the two is considerable. If temporary fluctuation is read as structural decline, students may be over-identified as vulnerable, over-prompted, or prematurely escalated. If persistent decline is dismissed as temporary, meaningful support may arrive too late. The function of recalibration is therefore to maintain a proportionate interpretive threshold between caution and responsiveness.

It is at this point that the role of **human oversight** becomes especially important. The partner validation questionnaires make clear that the framework is appreciated for its human-in-the-loop orientation, but they also suggest that this principle would benefit from slightly stronger practical articulation. This is a justified observation. The ethical legitimacy of longitudinal analytics does not depend only on the existence of probabilistic models or recalibration windows. It also depends on the continued authority of human pedagogical judgement to interpret, contextualise, and, where appropriate, override AI-supported readings. The HITS framework therefore treats automated pattern identification as an input to educational reflection, not as a substitute for it.

This point is especially important in cases where AI-supported interpretation and tutor judgement do not align. Such divergence should not be treated as a system failure to be automatically corrected in favour of one side or the other. Rather, it should trigger contextual review. The tutor may have access to information that is not visible in digital traces: knowledge of the student's wider situation, of discipline-specific demands, of institutional timing, or of recent interaction not captured by the model. Conversely, the model may detect a pattern that is not yet fully visible in direct interaction. The value of the human-in-the-loop principle lies precisely in allowing these two sources of insight to enter into relation without collapsing one into the other. Final pedagogical judgement remains human, but it remains so in dialogue with structured digital evidence rather than in ignorance of it. This is the practical meaning of oversight within the HITS model.

In this sense, longitudinal recalibration should be understood not only as a property of the model, but as a property of the educational process in which the model is embedded. Recalibration involves technical updating, but it also involves institutional interpretation. The system may revise pattern probabilities in light of new traces, but tutors and services must also remain capable of revising their own pedagogical understanding in light of evolving evidence. This double movement is what protects the framework from becoming either mechanically automated or purely impressionistic. Analytics contribute structure; human actors contribute context and judgement. Ethical educational support depends on the interaction between the two.

The section also has an important preventive function. By insisting on revisability, the framework protects students from being reduced to what might be called a persistent "data double": a stabilised digital representation that continues to influence support even after the underlying educational situation has changed. This is the concrete danger of epistemic freezing. Once such a representation becomes naturalised, the system may begin to guide attention, expectation, and intervention in ways that no longer correspond to the learner's present reality. The HITS model explicitly resists this possibility by treating pattern identification as temporary, evidence-sensitive, and always open to reconsideration. This is why the ethical discussion in this section is not external to the measurement model. It is one of the conditions of its responsible use.

A further implication of this approach concerns decision-making impact. One partner noted that the section would be strengthened by becoming slightly more oriented toward practical and decision-related consequences. This is an important suggestion. The value of recalibration lies not only in methodological caution, but in its capacity to guide proportional action. In educational terms, this means that stronger interpretive confidence should correspond not to stronger classification, but to more appropriate forms of support. If evidence of declining commitment becomes more stable over time, this may justify more explicit prompting, closer tutor attention, or earlier human follow-up. If later traces indicate recovery, then support intensity should also recede. Recalibration therefore acts as a safeguard against both overreaction and neglect. It enables a model of support that is responsive because it is restrained.

For this reason, the ethical dimension of educational analytics within HITS does not lie only in what the system avoids doing. It also lies in what it makes possible. By resisting fixed scoring, by distinguishing temporary from structural patterns, and by preserving the interpretive authority of tutors, the framework opens the possibility of support that is both more timely and more respectful of learner complexity. It does not promise certainty. It offers a disciplined form of provisionality. This is, in many respects, the most appropriate epistemic posture for AI-supported tutoring in higher education: not classification without remainder, but interpretation under conditions of responsibility.

The significance of this section extends beyond ethics in a narrow sense. Longitudinal recalibration is also what makes the overall framework educationally credible. Without it, commitment and awareness patterns would quickly harden into labels, and the project's concern with trust, autonomy, and student-centred support would be compromised. With it, the measurement model remains aligned with the broader principles of the HITS framework: interpretability, revisability, proportionality, and humanly accountable use. The next section builds on this logic by showing how the whole commitment-and-awareness framework connects to the architecture of the AI system itself and to the role of grounded retrieval in ensuring that engagement measurement remains trustworthy.

8. Integration with the RAG Architecture

The conceptual framework developed in this deliverable would remain incomplete if it were not connected to the technical conditions that make its educational claims credible. Digital academic commitment and digital awareness cannot be meaningfully interpreted in an AI-supported environment unless the environment itself meets certain epistemic requirements. In the HITS framework, these requirements are addressed through the adoption of a **Retrieval-Augmented Generation (RAG) architecture**, which plays a structural role in the relation between trust, support, and measurement. The relevance of this architectural choice is not limited to accuracy in the narrow technical sense. It is what allows digital support to become institutionally grounded, source-traceable, and therefore educationally reliable enough to sustain meaningful engagement over time.

This point is central to the logic of T.5.2. The deliverable argues that trust functions as the behavioural catalyst through which technical reliability becomes academic commitment. That claim, however, depends on a prior condition: the system must be able to justify the reliability on which trust is built. This is precisely where the RAG architecture becomes decisive. By constraining responses to a curated body of institutional documentation rather than relying on open-ended generation alone, the system transforms **institutional grounding** from a conceptual aspiration into an operational property. In other words, the architecture does not merely host the framework developed in this deliverable; it helps make that framework valid in practice.

The educational significance of this architectural grounding becomes especially clear when one considers the kinds of academic interaction the HITS Chatbot is designed to support. Students may ask about deadlines, procedural obligations, course requirements, examination rules, or available services. They may also seek academic guidance in moments of uncertainty, overload, or hesitation. In all such cases, the usefulness of the system depends not only on whether it can produce an answer, but on whether that answer is **anchored in legitimate and traceable institutional sources**. A system that generates plausible but unsupported guidance may appear helpful in the short term while undermining trust in the longer term. A system that grounds its responses in authoritative documentation, by contrast, creates the conditions under which students may begin to rely on it responsibly. This is why the architecture matters behaviourally as well as technically.

From the perspective of measurement, the importance of RAG is equally substantial. The framework proposed in T.5.2 seeks to identify patterns of commitment and awareness through indicators that include responsiveness, reflective behaviour, source sensitivity, and trust-related traces. Yet such indicators can only be interpreted meaningfully if the support environment itself is sufficiently stable and credible. If the system's answers were weakly grounded or inconsistent, then student hesitation, non-use, or repeated confirmation-seeking could not be interpreted in a straightforward way. It would remain unclear whether such behaviour reflected the learner's disposition or the unreliability of the system. The use of RAG therefore supports not only better tutoring responses, but also more credible analytics. It stabilises the support environment sufficiently for behavioural interpretation to become more meaningful. This is one of the reasons why partners regarded the connection between architecture, measurement credibility, and tutoring logic as convincing and well developed.

This relation may be described as a **trust–accuracy loop**. The more clearly the system can ground its responses in authoritative institutional sources, the more likely students are to perceive it as reliable. The more reliable and intelligible it becomes, the more likely students are to use it in ways that reveal meaningful patterns of engagement. Those patterns, in turn, can support more calibrated tutoring and more responsible measurement. In this sense, architecture and behaviour are not separate layers of the framework. They enter into a reciprocal relation. Trust depends on accuracy, but accuracy also shapes the conditions under which commitment and awareness can be observed, sustained, and interpreted. The validation questionnaires repeatedly recognised the strength of this logic, especially in the way the deliverable connects source traceability to student trust and continued use.

The architecture also matters because it reinforces the distinction between **generic AI** and **institution-grounded AI**, which the partners identified as one of the strongest aspects of the deliverable's conceptual framing. A generic AI tool may be fluent, responsive, and superficially useful, yet still lack the institutional specificity needed for trustworthy academic guidance. A RAG-based system differs precisely because it is designed to work within a bounded and curated documentary environment. The system's legitimacy does not derive from general fluency, but from its capacity to connect student queries to relevant institutional content. This distinction is crucial for digital awareness as well. Students can only develop a reflective understanding of the support environment if the environment itself actually differs from generic AI usage in a meaningful way. The RAG architecture is therefore part of what makes digital awareness a coherent and educationally relevant construct.

A practical illustration may clarify the point. Consider a student who asks whether a particular procedural step must be completed before an assessment deadline. If the system responds with information that is clearly rooted in an institutional source, then the student's subsequent reliance on that answer becomes intelligible as a form of trust in grounded support. If the same student later checks the source, recognises the institutional basis of the answer, and adjusts action accordingly, then both commitment and awareness become visible in relation to a trustworthy support environment. By contrast, if the system merely offers plausible but unsupported advice, then neither the student's reliance nor the student's caution can be interpreted in the same way. The architecture thus helps determine not only what the system can say, but what student interaction with the system can mean.

For these reasons, the integration of the commitment-and-awareness framework with the RAG architecture is not an external technical add-on. It is part of the deliverable's conceptual coherence. The indices, patterns, and ethical safeguards proposed in T.5.2 depend on the existence of a support environment whose outputs are sufficiently grounded to justify both educational reliance and analytical interpretation. Without such grounding, the trust construct would weaken, the behavioural meaning of support traces would become less stable, and the distinction between responsible academic engagement and generic digital use would become harder to sustain. With it, the entire framework gains a stronger methodological foundation.

At the same time, this integration should not be misunderstood as a claim of total architectural sufficiency. RAG supports trust, traceability, and interpretive credibility, but it does not eliminate the need for human judgment, critical awareness, or ethical oversight. A grounded architecture makes meaningful engagement possible; it does not guarantee it automatically. This is why the architecture section must be read together with the earlier discussion of commitment, awareness, recalibration, and human-in-the-loop governance. The system becomes educationally valuable not simply because it retrieves grounded sources, but because those sources are integrated into a wider support logic that remains reflective, revisable, and accountable.

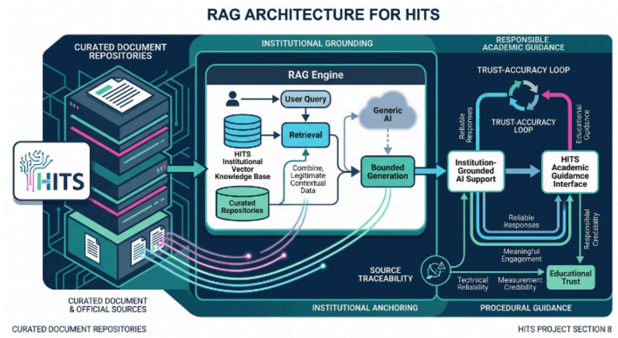

In this sense, Section 8 closes an important circle in the argument of the deliverable. The earlier sections established the conceptual distinction between commitment and awareness and showed how trust mediates between reliability and behaviour. The present section clarifies why that mediation is not merely theoretical. It is materially supported by the architecture through which institutional grounding, contextual relevance, and source traceability become operational. The result is not simply a technically stronger chatbot, but a more credible basis for interpreting digital academic engagement in educational terms.

9. Ethical Safeguards

The analytical and architectural framework developed in this deliverable can only be educationally legitimate if it remains bounded by clear ethical safeguards. This requirement is not secondary to the system's design. On the contrary, it is one of the conditions under which digital academic commitment and digital awareness can be interpreted in ways that remain pedagogically defensible. A model that measures patterns of engagement, identifies recurrent configurations, and supports differentiated tutoring interventions necessarily operates close to questions of representation, judgement, and institutional power. For this reason, the HITS framework treats ethics not as an external compliance layer, but as a constitutive dimension of responsible educational analytics.

A first and foundational safeguard concerns the explicit rejection of **permanent digital scoring** and **fixed learner labels**. The framework does not seek to assign students stable identities based on data traces, nor does it aim to translate engagement into a durable score that could follow the learner across contexts as if it were an intrinsic characteristic. Such practices would be incompatible with the project's methodological and pedagogical premises. The constructs proposed in T.5.2 are interpretive and probabilistic; they are designed to identify evolving configurations of commitment and awareness, not to stabilise students into categories that acquire institutional permanence. This principle was strongly validated by partners and remains one of the clearest ethical strengths of the deliverable.

The rejection of fixed scoring is closely linked to the broader concern with **epistemic freezing**, already discussed in the previous section. Once a learner's pattern is treated as definitive, educational analytics no longer function as a support tool; they begin to function as a mechanism of closure. A student who is persistently interpreted through the lens of low commitment or low awareness may find that future interactions are subtly shaped by that prior representation, even when the underlying educational situation has changed. Ethical safeguards are therefore necessary not only to prevent technical misuse, but also to protect the learner's openness to change. The HITS framework safeguards this openness by treating pattern identification as provisional, revisable, and bounded by human interpretation.

A second safeguard concerns the role of **human-in-the-loop review**. The partner questionnaires repeatedly confirm that this principle is valued as essential, and some partners explicitly ask for its practical meaning to be made more visible. Within HITS, human oversight is not understood as a symbolic final approval appended to an otherwise automated process. It is a substantive condition of the framework's legitimacy. AI-supported pattern identification may help surface recurring configurations of engagement, suggest likely support needs, or indicate the need for closer attention. It may not, however, replace educational judgement. The final interpretation of a student's situation remains with human tutors, support staff, and institutional actors capable of situating data-derived signals within a fuller pedagogical context.

This point matters especially where there is divergence between model output and tutor judgement. Ethical oversight requires that such divergence be treated as an occasion for review rather than as an error to be automatically resolved in favour of the system. A tutor may know that a temporary reduction in activity reflects illness, overload, or a strategic shift in study routine; the model may not. Conversely, the model may detect a recurring pattern not yet visible in direct interaction. Responsible use of the framework depends on keeping both possibilities open. The system contributes structured evidence; the human actor contributes contextual understanding and pedagogical discernment. Neither should be collapsed into the other. In this sense, human oversight is not merely a corrective to automation. It is part of the interpretive ecology through which educational analytics become ethically usable.

A third safeguard concerns **bias monitoring and contextual sensitivity**. Because the framework is intended for use across different countries, institutions, languages, and educational environments, it cannot assume that indicators will function identically in all contexts. The same behavioural trace may carry different meanings depending on institutional organisation, disciplinary culture, linguistic practice, or patterns of access and participation. An ethically responsible model must therefore remain attentive to the possibility that apparently neutral indicators may produce asymmetrical effects when transferred across settings. Several partners explicitly validated the importance of this cross-contextual attention, particularly in relation to institutional and linguistic variation.

Bias mitigation, in this framework, should not be understood in a narrowly technical sense alone. It involves ongoing review of how indicators behave across groups and settings, whether certain forms of interaction are systematically misread, and whether recalibration is needed in light of contextual evidence. This may concern language-specific formulations of doubt or confidence, discipline-specific rhythms of study, or institutional differences in how academic procedures are navigated. Ethical safeguards therefore require not simply statistical vigilance, but contextual interpretive vigilance. The framework remains fair not by pretending to be universally neutral, but by remaining open to the possibility that its own assumptions must be reviewed when they travel across environments.

A fourth safeguard concerns **transparency toward students**. If the HITS framework aims to support responsible and trust-based use of AI-supported tutoring, then students must be able to understand, in practical terms, what the system is doing and what it is not doing. Transparency here does not mean exposing every technical layer of the model in formal detail. It means ensuring that learners are clearly informed that pattern identification is used to calibrate support rather than to rank them, that engagement traces may contribute to adaptation logic, that the system remains grounded in institutional sources, and that human actors remain responsible for final pedagogical interpretation. Partners generally judged this aspect positively, but its practical importance deserves continued visibility in the text.

This transparency has a specifically pedagogical value. Students who understand that the system is designed to support them—rather than to monitor or score them in an irreversible way—are more likely to engage with it as a legitimate academic resource. Conversely, opacity risks undermining the very trust on which the framework depends. Ethical communication is therefore part of the support architecture itself. It sustains responsible participation by making the conditions of AI-mediated support intelligible and contestable. In this respect, transparency is closely related to awareness: the framework does not merely measure students' reflective understanding of AI-supported tutoring; it must also help make such understanding possible.

A fifth and final safeguard concerns the **purpose limitation** of the framework. The patterns and indicators proposed in T.5.2 are intended for **support calibration**, not for ranking, sanction, surveillance, or exclusion. This distinction is crucial. Educational analytics may remain ethically legitimate only if the interpretations they produce are tied to proportionate and supportive action. Once the same patterns are used for performance ranking, selection, or irreversible administrative judgement, their educational status changes fundamentally. The HITS model explicitly rejects such uses. Its purpose is to help identify where support may be needed, how trust and awareness interact with engagement, and how tutoring environments may respond more appropriately. Several partners explicitly validated the clarity with which this distinction is drawn, and it should remain one of the most visible normative statements in the final deliverable.

Taken together, these safeguards give the framework its ethical coherence. They ensure that the measurement of commitment and awareness remains tied to the principles already developed in earlier sections: revisability, interpretability, proportionality, grounded trust, and humanly accountable use. They also make clear that the HITS model is not simply concerned with what educational analytics can do, but with what they should be allowed to do in order to remain compatible with student dignity and pedagogical responsibility. This is one of the main reasons why the deliverable was received so positively by partners: the ethical framework does not appear as a later correction of an otherwise technical model, but as one of the elements that holds the whole architecture together.

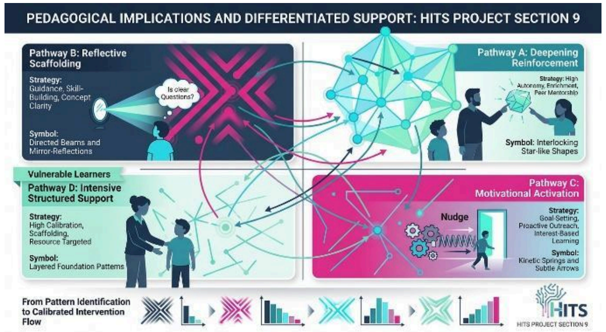

The present section therefore completes an important line of argument running throughout the deliverable. Digital academic commitment and digital awareness are meaningful constructs only if they are measured cautiously, interpreted contextually, and used in ways that remain educationally supportive rather than classificatory. Ethical safeguards are what make that possible. They prevent the framework from collapsing into either naïve automation or unaccountable profiling, and they preserve the project's broader ambition: to develop AI-supported tutoring in higher education that is not only technically effective, but also pedagogically justifiable and institutionally trustworthy.

10. Role in WP5 Progression

The identification of patterns related to **digital academic commitment** and **digital awareness** is not an isolated analytical exercise within the HITS project. It occupies a pivotal position in the internal logic of Work Package 5. Deliverable T.5.2 functions as a methodological hinge between conceptual modelling and later technical and pedagogical implementation. Its significance lies precisely in the fact that it transforms general behavioural constructs into interpretable and operational dimensions that can be used in subsequent tasks. Without such a framework, later work on clustering, adaptive nudging, and predictive modelling would risk proceeding on the basis of diffuse or weakly grounded behavioural assumptions.

This bridging role is one of the clearest contributions of the deliverable and was explicitly recognised in the validation process. Partners valued the way T.5.2 clarifies its connection to the broader WP5 workflow and demonstrates that commitment and awareness are not merely descriptive concepts, but structured variables capable of informing later stages of the AI Learning Library. In this respect, the deliverable contributes not only conceptual depth, but also work-package coherence. It ensures that subsequent tasks remain connected to a shared interpretive foundation rather than developing as disconnected technical modules.

A first and immediate contribution concerns **T.5.3**, where clustering techniques are used to identify strong patterns related to academic success. Clustering procedures require meaningful feature spaces. They do not become educationally useful merely by detecting statistical proximity in raw traces. The framework developed in T.5.2 provides the structured feature set that allows clustering to operate on variables that are pedagogically interpretable rather than analytically arbitrary. Indicators such as engagement stability, response-to-nudge ratio, resource depth, source verification, and meta-query behaviour become the dimensions along which similarity can be computed in ways that remain conceptually grounded. In this sense, T.5.2 protects clustering from becoming a purely computational exercise detached from educational meaning. It ensures that the clusters identified in T.5.3 can be read not as opaque technical groupings, but as recurrent configurations of commitment and awareness relevant to academic support.

A second contribution concerns **T.5.4**, where the project defines differentiated nudging interventions. The patterns articulated in T.5.2 make clear that nudging cannot be designed as a uniform mechanism applied identically to all learners. Its pedagogical appropriateness depends on the configuration of commitment and awareness through which student engagement is being interpreted. A learner displaying high commitment but low awareness, for example, may benefit less from activation prompts and more from interventions that strengthen reflective understanding of AI-supported guidance. A learner displaying low commitment and low awareness may require more structured, supportive, and carefully paced intervention. The role of T.5.2 in this respect is therefore decisive: it gives nudging a behavioural target and prevents support strategies from becoming generic or context-blind. In doing so, it reinforces one of the key ambitions of WP5, namely that AI-mediated support should remain differentiated, proportionate, and pedagogically justified.

A third major contribution concerns **T.5.7**, devoted to the predictive modelling layer of the system. Predictive algorithms require structured input variables if they are to produce results that are not only statistically robust, but also educationally intelligible. Raw interaction logs, taken on their own, offer limited interpretive value. They become educationally meaningful only when transformed into indicators that capture sustained engagement, responsiveness, reflective use, and reliance on grounded digital support. The commitment and awareness indices defined in T.5.2 provide precisely this transformation. They make it possible for predictive models to operate on behaviourally interpretable variables rather than on undifferentiated activity traces alone. In this sense, T.5.2 gives the predictive layer of WP5 a conceptual grammar. It allows prediction to remain connected to educational meaning and support logic rather than drifting toward technically impressive but pedagogically superficial pattern recognition.

The absence of T.5.2 would therefore weaken the whole progression of WP5. Clustering would risk identifying statistical regularities without sufficiently meaningful variables. Nudging would risk becoming generic rather than behaviourally calibrated. Prediction would risk operating on data that are abundant but poorly contextualised. The purpose of the deliverable is precisely to prevent such drift by ensuring that commitment and awareness become formally articulated constructs that can support the next stages of the project while remaining aligned with pedagogical interpretation and ethical safeguards. In this sense, T.5.2 does not merely precede later tasks chronologically. It underpins them methodologically.

This role becomes even clearer when viewed from the perspective of the wider HITS architecture. WP5 is not concerned with AI implementation in abstraction. It seeks to build an educational support ecosystem in which analysis, measurement, intervention, and technical infrastructure remain coherent. T.5.2 contributes to this coherence by ensuring that the measurement of digital academic engagement is conceptually aligned with the behavioural framework introduced earlier in the work package and technically compatible with the later stages of clustering and prediction. It also ensures that the framework remains ethically constrained: the same probabilistic, revisable, and human-supervised logic that governs commitment and awareness measurement carries forward into later tasks. This continuity is one of the deliverable's most important contributions, even where it remains partly implicit.

The deliverable also plays a significant role in relation to the project's broader concern with **interpretability**. One of the recurrent risks in AI-supported higher education is that later technical stages may become increasingly opaque as they move toward clustering and machine learning. T.5.2 acts as an important counterweight to this tendency. By defining commitment and awareness in explicitly educational terms and linking them to measurable indicators, it helps ensure that later analytical outputs remain interpretable by educators, tutors, and institutional actors. This is especially important for trust. If future clustering or prediction outputs are to be used responsibly, they must remain legible in relation to constructs that educators can understand and contest. T.5.2 is thus also a safeguard against methodological opacity within WP5.

A further implication concerns institutional transferability. Because the deliverable formalises commitment and awareness through a combination of conceptual clarity and feature-based logic, it helps create a common behavioural language across partner institutions. This does not eliminate contextual differences, nor should it. But it does provide a shared structure through which different universities can contribute to the broader AI Learning Library while preserving comparability. In this sense, T.5.2 supports not only downstream technical tasks, but also the collaborative integration of evidence across the consortium. This role is especially important in a European project context, where methodological coherence must coexist with institutional diversity.

The progression enabled by T.5.2 is therefore cumulative. Commitment and awareness are first defined conceptually; they are then translated into indicators and patterns; these patterns become features for clustering; the resulting behavioural profiles inform adaptive nudging; and the same variables later support predictive modelling aimed at early detection and proportionate intervention. At each stage, the value of the framework lies in preserving interpretability and educational relevance. Rather than moving from theory to engineering through abrupt discontinuity, WP5 develops through successive transformations of the same behavioural logic. T.5.2 is the point at which that logic becomes sufficiently formalised to travel across tasks without losing coherence.

For this reason, the role of T.5.2 in WP5 progression should be understood as both methodological and strategic. Methodologically, it provides the indices, feature logic, and pattern structures needed for later computational and pedagogical tasks. Strategically, it ensures that the AI Learning Library evolves as an integrated educational system rather than as a set of partially connected technical components. This integrative role is one of the reasons the partner questionnaires so consistently judged the deliverable to be mature, coherent, and strongly aligned with the broader objectives of the project.

The significance of this section, then, lies in making visible what might otherwise remain implicit: that the conceptual work of T.5.2 is indispensable to the later operationalisation of HITS. The deliverable does not simply define patterns; it prepares the system to use them responsibly. In doing so, it helps ensure that future clustering, nudging, and predictive components remain behaviourally grounded, ethically bounded, and educationally interpretable. This is the sense in which T.5.2 becomes one of the key structuring deliverables of WP5. It links analysis to action without sacrificing the methodological caution and pedagogical responsibility on which the project depends.

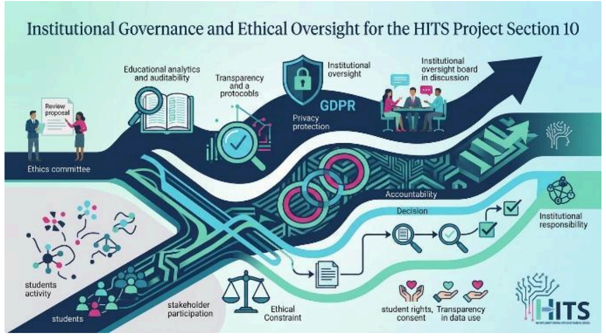

11. Conclusion

Deliverable T.5.2 has sought to establish a framework through which digitally mediated academic engagement can be interpreted in a way that is educationally meaningful, methodologically structured, and ethically responsible. Its central contribution lies in the articulation of two complementary constructs—**Digital Academic Commitment** and **Digital Awareness**—through which the use of AI-supported tutoring systems can be understood not merely as technological interaction, but as a complex academic relation shaped by continuity, reflection, trust, and institutional grounding. In doing so, the deliverable responds to a key challenge of contemporary higher education: how to move beyond superficial measures of digital activity and toward a richer interpretation of what it means for students to engage meaningfully with AI-supported academic environments.

A major achievement of the framework is that it avoids two equally problematic reductions. On the one hand, it does not collapse engagement into platform usage, log frequency, or simple digital presence. On the other hand, it does not retreat into an abstract discourse of reflection detached from measurable educational practice. Instead, it proposes a model in which commitment is treated as sustained, proactive, and structured engagement, while awareness is treated as the reflective and critical dimension through which that engagement becomes more informed and responsible. The interaction of these two constructs allows the framework to remain both conceptually nuanced and operationally useful. This balance was repeatedly recognised in the partner validation process as one of the principal strengths of the deliverable.

The deliverable has also argued that these constructs can only be understood adequately if they are linked through **trust**. Trust, in the HITS framework, is not treated as an incidental psychological factor, but as the behavioural condition through which technical reliability becomes educationally consequential. Only when students can experience digital support as accurate, institutionally grounded, contextually relevant, and source-traceable can such support begin to shape academic behaviour in a stable and meaningful way. This trust-based perspective also explains why the deliverable insists so strongly on the importance of **RAG architecture**, not as a purely technical choice, but as the architectural condition that gives source grounding and traceability practical substance. In this respect, the document makes a broader claim: AI-supported tutoring becomes educationally legitimate not through fluency alone, but through grounded reliability and accountable design.

Another important contribution of T.5.2 lies in its effort to translate conceptual distinctions into a usable methodological framework. The **four-pattern typology**, together with the **Commitment Index** and **Awareness Index**, provides a structured means of identifying recurrent configurations of engagement without reducing them to deterministic categories. This is a decisive point. The deliverable does not propose fixed learner types, permanent scores, or automated labels. It proposes a

probabilistic and revisable interpretive model, intended to support the calibration of tutoring rather than the classification of students as ends in themselves. The insistence on longitudinal recalibration, rolling interpretation, and the avoidance of epistemic freezing ensures that the framework remains aligned with the fluid and contextual nature of educational trajectories. This methodological and ethical caution was among the elements most positively valued by project partners.

At the same time, the deliverable has made clear that no responsible educational analytics framework can operate without **human oversight**. AI-supported pattern recognition may help identify support-relevant configurations, but it cannot replace pedagogical interpretation. The role of tutors, advisors, and institutional actors remains central precisely because educational situations are irreducibly contextual. This is why the deliverable consistently situates its measurement logic within a **human-in-the-loop** framework. Patterns are intended as diagnostic and support-oriented aids, not as autonomous decisions. Final pedagogical judgement remains with human actors, and the interpretive authority of that judgement becomes especially important where AI-supported outputs and contextual educational knowledge do not fully coincide. This human-supervised logic is one of the main safeguards through which the framework retains both ethical legitimacy and pedagogical credibility.

Within WP5, the significance of T.5.2 extends beyond its conceptual contribution. As the deliverable has shown, the framework developed here underpins several downstream tasks. It provides the behavioural and analytical variables needed for clustering in **T.5.3**, offers the interpretive basis for differentiated intervention in **T.5.4**, and supplies structured feature logic for predictive modelling in **T.5.7**. In this sense, T.5.2 performs an essential integrative function. It ensures that later technical and pedagogical developments do not proceed on the basis of unexamined assumptions about engagement, but remain anchored in a shared and interpretable model. Several partners explicitly recognised this bridging role and valued the way the deliverable clarifies its place within the wider workflow of WP5 and the AI Learning Library.

The validation evidence also suggests that the framework has achieved an appropriate balance between maturity and openness. Partners consistently judged the deliverable to be conceptually robust, methodologically credible, and ready for finalisation with only minor edits. At the same time, some partners usefully reminded the project that the present framework should remain open to **empirical validation**, **pilot testing**, and **institutional recalibration**. This is fully consistent with the internal logic of the deliverable itself. The constructs, indicators, and typologies defined here are not meant to close interpretation. They are meant to provide a grounded basis for later testing, refinement, and adaptation across contexts. Their strength lies precisely in being sufficiently formalised to guide implementation, while remaining sufficiently flexible to learn from it.

For this reason, the conclusion of T.5.2 should not be read as the final closure of a conceptual phase, but as the completion of a necessary methodological foundation. The deliverable gives the HITS project a shared language through which digital academic engagement can be analysed, measured, and related to trustworthy AI-supported tutoring. It also provides the safeguards needed to ensure that such analysis remains ethically bounded and educationally oriented. What follows, in the later stages of WP5, is not the replacement of this framework, but its testing in practice: through clustering, adaptive intervention, predictive logic, and iterative human review. In that sense, T.5.2 is both a conclusion and a point of departure. It closes the phase of conceptual definition while opening the phase of empirical and pedagogical operationalisation.

Ultimately, the wider value of the deliverable lies in the kind of educational AI model it makes possible. Rather than treating students as passive data sources or treating AI as a neutral efficiency tool, T.5.2 proposes a different logic: one in which engagement is interpreted through commitment and awareness, support is mediated through trust, and measurement remains accountable to ethics, context, and human judgement. In this respect, the deliverable contributes not only to the technical development of the HITS ecosystem, but also to a more responsible conception of AI in higher education—one in which analytical precision, pedagogical care, and institutional accountability are not separate ambitions, but mutually reinforcing conditions of meaningful innovation.