Lightweight Institutional Analytics: A Construct-Oriented Workflow for Course-Level LMS Indicators

Ariel Ortiz-Beltrán[1,*,†], Davinia Hernández-Leo[1,†]

[1]Universitat Pompeu Fabra, Plaça de la Mercè 10–12, 08002 Barcelona, Spain

Abstract

Learning Management Systems (LMSs) generate large volumes of trace data, but transforming these data into course-level views that are readable, lightweight, and replicable remains a challenge for Institutional Analytics. This paper contributes a construct-oriented analytical workflow that takes a minimal set of LMS-derived indicators and produces course-level views organized around three pedagogical constructs — Distributed Learning, Workload, and Active Learning — together with an assessment of how those views relate to the empirical structure of the underlying indicators and how alternative synthesis strategies produce different summaries. We illustrate the workflow with a case study of 75 university courses, deriving seven course-level indicators from standard LMS fields and examining them through descriptive analysis, Principal Component Analysis (PCA), threshold-based profiling, clustering, and robustness checks. The construct-oriented views preserve variation across courses while reducing the number of dimensions considered simultaneously. The empirical structure of the seven indicators partially aligns with the construct grouping while retaining more than one dimension of variation. The two synthesis strategies produce partially convergent summaries, with clustering more sensitive to filtering decisions than the descriptive and PCA-based layers of the workflow. The workflow is replicable by institutional analysts from standard LMS exports, and produces course-level views designed to support structured reflection by stakeholders in Institutional Analytics.

Keywords

Institutional analytics, LMS trace data, Course-level indicators, Construct-oriented analytics, Dimensionality in learning analytics

1. Introduction

Learning Management Systems (LMSs) generate large volumes of interaction data that can support reflection on teaching and learning in higher education. Learning Analytics (LA) has developed methods for extracting behavioural patterns from digital traces, and Learning Design (LD) has provided concepts for interpreting educational activity in relation to pedagogical intent. A long-standing argument in the field is that these two perspectives should be connected, so that analytics can be interpreted in relation to the educational designs they are expected to inform [1, 2, 3, 4, 5, 6]. In practice, however, much institutional use of LMS data still depends either on raw traces or on isolated indicators that are difficult to read at course level.

This problem is especially relevant for Institutional Analytics, where the goal is not only to describe activity, but to support reflection and decision-making across many courses under realistic constraints on data availability and analytical capacity [7]. Recent work on human-centred and transparent learning analytics has reinforced that indicators are more likely to be useful when their logic can be understood, discussed, and related to stakeholders' needs [8, 9, 10, 11]. At the same time, current debates have highlighted the need for analytical approaches that remain cautious about transferability and generalisability across contexts [12, 13].

Previous work has responded to part of this challenge by showing that standard activity-level LMS fields can be transformed into a concise set of indicators aligned with three pedagogically significant constructs: Distributed Learning, Workload, and Active Learning [14]. Moving from that indicator

LASI Spain 26: Learning Analytics Summer Institute Spain 2026, Zamora, May 07–08, 2026

*Corresponding author.
†These authors contributed equally.
ariel.ortiz@upf.edu (A. Ortiz-Beltrán); davinia.hernandez-leo@upf.edu (D. Hernández-Leo)
0000-0003-1522-2362 (A. Ortiz-Beltrán); 0000-0003-0548-7455 (D. Hernández-Leo)

set to interpretable course-level views, however, raises a distinct set of methodological questions. How do construct-driven groupings of indicators relate to the empirical covariance structure of the indicators themselves? And once construct-oriented views are computed, how sensitive are the resulting course-level summaries to the synthesis strategy used?

This paper contributes a transparent analytical workflow that addresses these questions, together with a case study illustration using 75 university courses. Specifically, the paper contributes: (i) a replicable workflow, runnable from standard LMS exports, that produces construct-oriented course views for stakeholder use; (ii) an empirical characterization of how the construct grouping relates to the PCA-based structure of the indicator set; and (iii) a comparison of two methods for summarising the views at course level — tercile-based profiling and k-means clustering — that makes the effect of analytical choices visible. The paper is guided by:

- **RQ1.** How does the construct-based grouping of the indicators compare with their empirical structure, as revealed by Principal Component Analysis?
- **RQ2.** When the construct-oriented views are summarised at course level using two alternative methods — tercile-based profiling and k-means clustering — how far do the two sets of summaries agree?

Each RQ is addressed with a method chosen for its fit. PCA surfaces the covariance structure of the indicators without reference to the construct grouping, so the construct-based grouping can be compared against a construct-independent alternative (RQ1). A side-by-side comparison of tercile profiling and k-means clustering summarises the construct-oriented views under different assumptions —independent per-view thresholds vs. joint multivariate similarity — and so allows the sensitivity of course-level summaries to be examined (RQ2).

2. Data and analytical approach

We use three terms consistently throughout the paper. A *course-level indicator* is a numerical summary computed per course from LMS traces (e.g., Course Interaction Span). A *construct-oriented view* is an aggregate summary per course obtained by averaging the standardized values of a pre-specified subset of indicators associated with a construct; the three views used here correspond to Distributed Learning, Workload, and Active Learning. The *empirical structure of the indicator set* refers to the covariance and principal-component structure of the seven indicators, obtained without reference to the construct grouping. A *synthesis strategy* is a procedure that maps the three construct-oriented views of a course onto a single summary, such as a profile label or a cluster assignment.

2.1. Dataset and preprocessing

The workflow operates on six core LMS fields commonly available in standard exports: an anonymized learner identifier, a course identifier, an activity or resource identifier, a visit count, a first-access timestamp, and a last-access timestamp. Restricting the workflow to these fields favours transferability across institutional contexts and is consistent with prior work seeking to bridge Learning Design and Learning Analytics through indicators derived from standard LMS traces rather than from institution-specific infrastructures [1, 2, 4, 14]. The case study below uses an instantiation of these fields drawn from the Moodle platform of a Spanish university.

The workflow specifies a small set of preprocessing rules applied before indicator construction: duplicated rows are removed, records with missing values in required fields are excluded, non-positive visit counts are discarded, and cases in which the recorded last access precedes the first access are filtered out. Where multiple records exist for the same learner–course–activity combination, visits are summed, the earliest first-access date and latest last-access date are retained, and categorical fields are collapsed using their most frequent value. Applied to the case study, this aggregation step produced an analytical dataset suited for course-level indicator construction while preserving the temporal and interaction information required by the analysis. Time-spent was retained only as supplementary information and excluded from the indicator set, consistent with prior work [14].

Regarding ethics and data use, the institutional LMS records were authorized for internal research use under the institution's data-governance framework. All learner and course identifiers were anonymized prior to analysis, and no personal attributes, grades, or demographic variables were accessed. The analysis operates on activity-level traces only and does not attempt to infer individual student characteristics. We note the limits of what LMS traces can legitimately evidence: they capture recorded interaction with the platform, not the full pedagogical experience of students or teachers, and the construct-oriented views proposed in this paper should be read within that constraint.

2.2. Indicator framework

The analytical workflow uses a set of minimal LMS-derived indicators based on prior work that transformed standard activity-level LMS fields into course-level measures linked to pedagogically meaningful constructs [14]. In this study, these indicators are used as the starting point for building interpretable course-level views. The purpose is not to infer the full complexity of teaching and learning from LMS traces alone, but to create a concise and transparent representation of course activity that can support institutional reflection.

The indicator set includes seven measures. Course Interaction Span (CIS) captures the number of days between the first and last recorded interaction in a course. Course Interaction Frequency (CIF) summarizes the average number of visits per active course day. Daily Access Frequency (DAF) represents the number of days in which interaction was observed. Diversity Activity/Resource Index (DAI) captures the number of distinct activities or resources accessed in the course. Activity/Resource Interaction Span (AIS) summarizes the average duration of interaction with individual activities or resources. Activity/Resource Interaction Frequency (AIF) captures the average visit rate at the activity/resource level. Finally, First-time Access Interval (FAI) represents the average interval between the first accesses to different activities or resources.

These three constructs are drawn from established strands of the educational literature, and are used here as interpretive lenses rather than as direct measurements. Distributed Learning draws on the distributed-practice and spacing literature, which has documented that spaced rather than concentrated practice is associated with better retention in many learning tasks [15]. Workload is informed by research on student course workload in higher education, where it has been treated both as a subjective experience and as an observable property of course structure [16]. Active Learning draws on the instructional tradition that emphasizes students' engagement with learning activities, which has been empirically associated with learning outcomes in prior research [17]. We do not claim that LMS traces measure these constructs directly; rather, we use the three labels to organize indicators that plausibly carry partial evidence about each dimension, in line with recent human-centred LA work on indicator interpretability [8, 9].

2.3. Construct-oriented course views

To move from isolated indicators to course-level interpretation, the seven indicators were reorganized into three construct-oriented views: Distributed Learning, Workload, and Active Learning. The mapping (Table 1) follows the conceptual content of each construct as introduced in Section 2.2. Distributed Learning groups indicators that capture temporal spread and pacing of engagement (CIS, DAF, AIS, FAI). Workload groups indicators that capture the volume and intensity of demand placed on learners (CIF, DAF, DAI). Active Learning groups indicators that capture sustained and varied engagement across activities (DAI, CIS, FAI, DAF). The three groupings are deliberately non-orthogonal. Pedagogically, the constructs are not independent: a course with sustained interaction across varied activities is plausibly informative about all three simultaneously. This is reflected in shared indicators — for instance, DAF (Daily Access Frequency) appears in all three constructs because the number of days with recorded interaction is informative about temporal spread, accumulated volume, and sustained engagement. The mapping is therefore interpretive rather than measurement-derived, and the overlapping structure is itself a meaningful property of how LMS traces relate to pedagogical constructs. Each view was computed by standardizing the corresponding indicators and averaging the standardized values within the construct, giving equal weight to each indicator.

Table 1
Construct–indicator mapping with brief indicator definitions. Indicators marked with * appear in more than one construct.

Construct	Indicator	Captures
Distributed Learning	CIS*	Days between first and last interaction
	DAF*	Days with observed interaction
	AIS	Avg. duration of activity-level interaction
	FAI*	Avg. interval between first accesses
Workload	CIF	Visits per active course day
	DAF*	Days with observed interaction
	DAI*	Distinct activities/resources accessed
Active Learning	DAI*	Distinct activities/resources accessed
	CIS*	Days between first and last interaction
	FAI*	Avg. interval between first accesses
	DAF*	Days with observed interaction

2.4. Comparative analytical strategies

After constructing the course-level views, we applied five analytical steps [12, 13]. The first examined the descriptive structure of the indicator set as a baseline. The second applied PCA to address RQ1, producing a construct-independent empirical structure for comparison against the construct grouping. The third and fourth — tercile profiling and k-means clustering — addressed RQ2 under different assumptions: independent per-view thresholds and joint multivariate similarity. The fifth was a robustness analysis testing stability under alternative filtering decisions.

To address RQ1, we applied Principal Component Analysis (PCA) to the seven course-level indicators as a construct-independent check on the empirical structure of the indicator set. Indicators were standardised before decomposition, components were inspected using both the eigenvalue-one rule and cumulative variance, and loadings were compared against the construct grouping. PCA was chosen over factor analysis because no latent-factor model is claimed; PCA is used as a descriptive redundancy check, not as a measurement model.

To address RQ2, we compared two strategies for summarising the construct-oriented views. The first used dataset-relative terciles: within each view, courses falling in the bottom, middle, or top third of the standardized distribution were labeled *low*, *medium*, or *high*, respectively, and the three labels were combined into *compact profiles*. Terciles were chosen as a simple design choice consistent with the broader emphasis on transparency and readability in human-centred learning analytics [8, 10]. In practical terms, they are easy to explain to non-technical stakeholders, produce balanced class sizes by construction, and avoid assuming a parametric distribution of the views. The second strategy applied k-means clustering to the three views as a data-driven alternative that groups courses by joint similarity rather than by per-view position; hierarchical clustering was additionally computed as a secondary method used for agreement checking via the Adjusted Rand Index (ARI). Profiles and clusters were then compared to examine how the two strategies represented the same construct-oriented space. Finally, we conducted a robustness analysis by repeating key steps after excluding very small or sparse courses and top-end outliers, to assess whether the broader interpretation remained stable under reasonable filtering decisions.

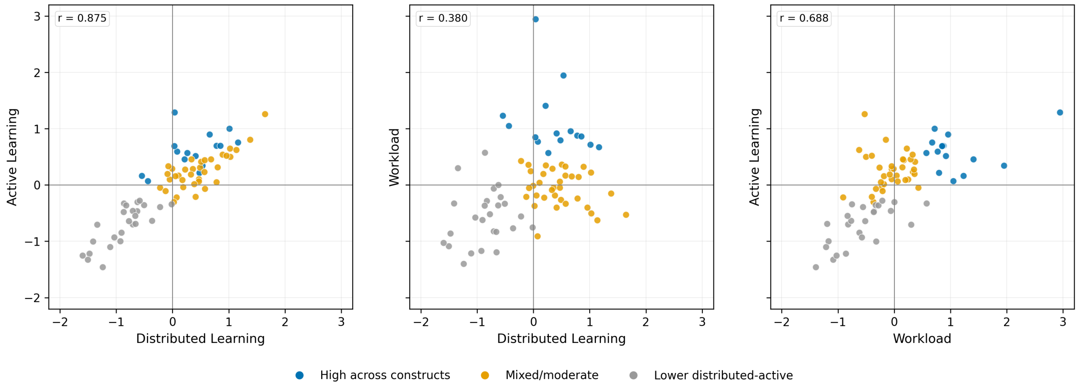

Figure 1: Course-level distribution across the three construct-oriented views, shown pairwise and colored by k-means cluster (k = 3). Each point represents one course. The panel labels report the corresponding Pearson correlations.

3. Results

3.1. Descriptive structure of the indicator set

The workflow described in Section 2.4 was instantiated on the case study to illustrate its outputs and behaviour. The case study included 75 courses after preprocessing. Table 2 reports the main course-level descriptive statistics. The ranges in Table 2 show substantial variation in course size, number of activities/resources, total visits, and temporal span, providing a heterogeneous case study for illustrating the workflow.

Table 2
Selected descriptive statistics of the case study at course level.

Measure	Mean	Median	Range
Number of students	41.8	38	1–98
Number of activities/resources	36.8	30	1–139
Total visits	3916.1	2107	26–30704
CIS	357.6	400	75–529
CIF	10.4	6.1	0.30–68.84
DAF	345.7	391	70–529
DAI	36.8	30	1–139

3.2. Construct-oriented course views

Reorganizing the seven indicators into construct-oriented course views yielded three course-level summaries corresponding to Distributed Learning, Workload, and Active Learning. These views retained variation across the case study while reducing the number of analytical dimensions considered simultaneously.

The three views were related, but not identical. The strongest association was observed between Distributed Learning and Active Learning (r = 0.875). Workload showed a weaker association with Distributed Learning (r = 0.380) and a moderate association with Active Learning (r = 0.688). Some asymmetry between these pairs is consistent with the indicator-construct mapping itself — Distributed Learning and Active Learning share three of their four indicators (Table 1), while Workload shares fewer with each — but the specific magnitudes are properties of the case study rather than predictable from the mapping alone. The standard deviations of the three views were 0.755 for Distributed Learning, 0.748 for Workload, and 0.617 for Active Learning, indicating that all three retained dispersion across

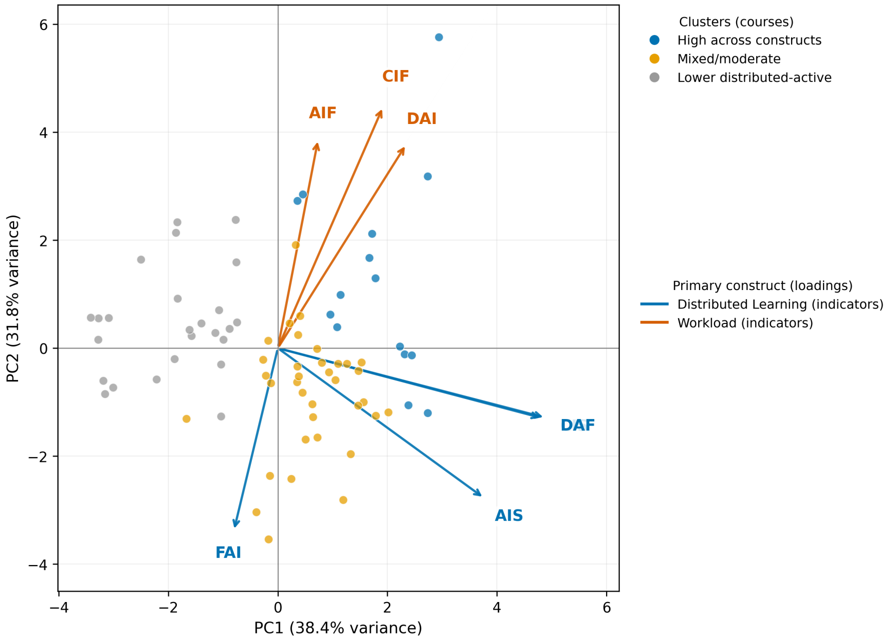

Figure 2: PCA biplot of the seven course-level indicators. Points represent courses and are colored by k-means cluster (k = 3). Arrows represent indicator loadings on PC1 and PC2 and are colored according to their primary construct grouping. Loading arrows were scaled to the score cloud for visibility.

courses, with Workload showing a broader spread than Active Learning. Figure 1 shows the three pairwise relationships between construct-oriented views, colored by k-means cluster.

3.3. Empirical structure of the indicator space

To address RQ1, we compared the construct grouping against the empirical structure of the indicator set by applying a Principal Component Analysis (PCA) to the seven course-level indicators, without reference to the construct assignments. The PCA was used as a descriptive dimensionality-reduction step within the workflow. The first principal component explained 38.4% of the total variance, and the second explained 31.8%, for a cumulative 70.2% captured by the first two components.

The loadings indicated that the first component was mainly associated with Course Interaction Span (CIS, 0.570), Daily Access Frequency (DAF, 0.577), and Activity/Resource Interaction Span (AIS, 0.445). The second component was more strongly associated with Course Interaction Frequency (CIF, 0.528), Activity/Resource Interaction Frequency (AIF, 0.457), Diversity Activity/Resource Index (DAI, 0.446), and First-time Access Interval (FAI, −0.400). In answer to RQ1, the empirical structure of the indicator set partially, but not fully, recovers the construct grouping. The first principal component loads most strongly on CIS, DAF, and AIS — three of the four indicators assigned to Distributed Learning — while the second component loads on CIF, AIF, DAI, and FAI, mixing indicators from Workload and Active Learning. The construct grouping and the PCA-derived structure are therefore related but not equivalent: the empirical structure aligns more cleanly with the Distributed Learning grouping than with the separation between Workload and Active Learning.

Figure 2 visualizes this result as a biplot: courses (points) are colored by k-means cluster, and indicator loadings are shown as arrows from the origin. The *High across constructs* cluster separates from the others along PC2, where CIF, AIF, and DAI load; the *Lower distributed-active* cluster separates along PC1, where CIS, DAF, and AIS load. Because the Active Learning view is defined through indicators also used in Distributed Learning and Workload, it is interpreted through this overlap rather than through a separate arrow colour.

3.4. Comparing synthesis strategies

Two alternative strategies were compared for summarising the construct-oriented views: tercile-based profiling and k-means clustering (see Section 2.4). For clustering, the highest silhouette value was obtained for the two-cluster solution (0.486), while the three-cluster solution yielded a silhouette of 0.395. For the three-cluster solution, the agreement between k-means and hierarchical clustering was moderate (ARI = 0.528). The three k-means clusters are labelled by their mean construct profile: a *high across constructs* cluster, a *lower distributed-active* cluster, and a *mixed/moderate* cluster.

Table 3
Construct-level composition of the three k-means clusters. Each row shows the percentage of courses in the cluster whose tercile level is *High* or *Low* for each construct-oriented view.

	Distributed		**Workload**		**Active**	
Cluster (N)	High %	Low %	High %	Low %	High %	Low %
High across constructs (15)	46.7	13.3	100.0	0.0	80.0	0.0
Lower distributed-active (25)	0.0	92.0	8.0	72.0	0.0	96.0
Mixed/moderate (35)	51.4	0.0	22.9	20.0	37.1	2.9

Table 3 shows how each cluster is composed in terms of the tercile profiles of the three construct-oriented views. Workload shows the sharpest contrast across clusters (100% High in the *High across constructs* cluster, 8% in the *Lower distributed-active* cluster, 22.9% in the *Mixed/moderate* cluster), whereas Distributed Learning and Active Learning separate the lower-activity cluster from the other two but do not distinguish between them. In answer to RQ2, the two synthesis strategies produce partially convergent summaries: tercile profiling captures a course's relative position within each view independently, while clustering captures joint position. Threshold-based profiling and clustering consequently carry overlapping but non-interchangeable information about the construct-oriented space.

3.5. Robustness under filtering decisions

Key steps of the workflow were repeated under three alternative filtering decisions: excluding small or sparse courses, excluding top-end outliers, and combining both restrictions. The construct-oriented correlations remained broadly stable across subsets. In the full sample, the correlation between Distributed Learning and Active Learning was 0.875; after excluding small/sparse courses it remained 0.871, and after excluding top 1% outliers it was 0.885. The weaker relation between Distributed Learning and Workload also remained similar, varying between 0.352 and 0.380 across subsets.

The PCA results were similarly stable: the cumulative variance explained by the first two components was 70.2% in the full sample, 69.2% after excluding small/sparse courses, and 67.9% after excluding top 1% outliers. Clustering was more sensitive in the full sample but converged under filtering: the agreement between k-means and hierarchical clustering at $k = 3$ rose from $ARI = 0.528$ in the full sample to 0.801 after excluding top 1% outliers, and to 1.000 when both small/sparse courses and top 1% outliers were excluded. This is consistent with part of the clustering instability being driven by a small number of atypical courses, though the three-cluster structure remains exploratory.

4. Discussion

This paper builds on previous work on minimal LMS-derived indicators by moving from indicator extraction to course-level interpretation. Its contribution is a transparent analytical workflow that organizes minimal LMS indicators into construct-oriented course views aligned with Distributed Learning, Workload, and Active Learning, and that examines how those views relate to the empirical structure of the indicators and to alternative synthesis strategies. The construct lens and the empirical structure are shown to agree in part and diverge in part, providing a transparent characterisation of how minimal LMS data can be organised for institutional reflection [1, 5, 6].

The empirical structure of the indicator set carries an interpretable asymmetry. Workload and Active Learning share indicators by construction (DAF and DAI appear in both), so a degree of empirical entanglement between them is built into the indicator-construct mapping itself and would be expected in any dataset. What the PCA adds is a measurement of how strongly that built-in entanglement actually surfaces, alongside the relative separability of Distributed Learning. Distributed Learning is directly encoded in the temporal geometry of LMS traces — days of activity, intervals between first accesses, span of interaction — which makes it largely separable in the PCA. Workload and Active Learning, by contrast, tap dimensions that are entangled in the indicator set used here: a course with many varied activities is simultaneously high-workload and high-engagement under these measures. The asymmetric alignment between the construct grouping and the PCA structure is therefore not a defect of the indicators but a structural property of how the indicator set encodes the three constructs.

From the perspective of Institutional Analytics, this is relevant because many institutions need approaches that are not only technically feasible, but also readable and actionable for practitioners and stakeholders. Recent work has emphasized that useful analytics depend not only on computational validity, but also on transparency, trust, adoption conditions, and alignment with users' needs [10, 11]. The construct-oriented views are designed to function for stakeholders as an intermediate representation between isolated indicators and more synthetic summaries; whether practitioners experience them as more readable than the underlying indicators is a question for empirical evaluation. Rather than expecting practitioners to interpret multiple course-level indicators independently, the workflow offers a smaller number of descriptive views that can support reflection while remaining easy to communicate. As one concrete application, a teaching-and-learning support unit could use the construct-oriented views both to identify courses whose activity patterns warrant closer inspection — for example, courses with high Workload but low Distributed Learning — and to surface candidate examples of good practice, such as courses in the *High across constructs* cluster or with profiles combining high Distributed Learning and high Active Learning without extreme Workload. In either case, compact profiles and cluster labels serve as conversation starters with course coordinators rather than as evaluative judgments. The workflow does not prescribe the decision; it surfaces a structured reading of the course landscape that can anchor such decisions. At the same time, descriptive workflows of this kind can have downstream effects on teachers and course coordinators when embedded in institutional decision-making, and should therefore be deployed to anchor conversations among stakeholders rather than to score or rank courses.

A further value of the workflow is that it makes analytical choices visible. Construct-oriented views remained stable across the robustness checks, while synthesis strategies varied with specification and filtering. This suggests that in this case study, the observed clustering instability was diagnosable and localisable, a useful pattern to check for in Institutional Analytics deployments, where practitioners need to know whether a fragile summary reflects a genuine ambiguity in the data or a small number of edge cases. For the same reason, course-level summaries should be read as analytical representations derived under specific assumptions, not as neutral descriptions of courses [10, 12]. The scope of the paper should be read accordingly. The case study illustrates how the workflow can be operationalized on a single institutional dataset; the construct-oriented views are analytical summaries derived from minimal LMS traces, not validated measures of pedagogy, and the resulting course-level summaries are not intended to transfer unchanged to other contexts.

5. Conclusions and future work

This paper proposed a transparent analytical workflow for transforming minimal LMS-derived indicators into construct-oriented course-level views, and characterized how those views relate to the empirical structure of the indicators and to alternative synthesis strategies. The contribution is methodological: the workflow turns minimal LMS data into course-level representations that analysts can compute transparently and that stakeholders can read without specialised technical knowledge, while making analytical choices visible [14]. Future work should examine the workflow across additional institutional contexts and datasets to assess transferability, complement LMS traces with learning design information and assessment structure, and study how practitioners interpret and use the resulting course-level views in real decision-making processes [12, 13, 18].

Acknowledgments

This work has been supported by PID2023-146692OB-C33, 2023-26 and CEX2021-001195-M by MICIU/AEI/10.13039/501100011033. DHL, Serra Húnter, acknowledges the support of ICREA Academia.

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude (Anthropic) in order to: perform grammar and spelling checks, provide content review and feedback, and assist with LaTeX generation. After using these tools, the author(s) reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] L. Lockyer, S. Dawson, Learning designs and learning analytics, in: Proceedings of the 1st International Conference on Learning Analytics and Knowledge, Association for Computing Machinery, New York, NY, USA, 2011, pp. 153–156. URL: https://dl.acm.org/doi/10.1145/2090116.2090140. doi:10.1145/2090116.2090140.

[2] D. Persico, F. Pozzi, Informing learning design with learning analytics to improve teacher inquiry, British Journal of Educational Technology 46 (2015) 230–248. doi:10.1111/bjet.12207.

[3] Y. Mor, R. Ferguson, B. Wasson, Editorial: Learning design, teacher inquiry into student learning and learning analytics: A call for action, British Journal of Educational Technology 46 (2015) 221–229. doi:10.1111/bjet.12273.

[4] A. Bakharia, L. Corrin, P. de Barba, G. Kennedy, D. Gašević, R. Mulder, D. Williams, S. Dawson, L. Lockyer, A conceptual framework linking learning design with learning analytics, in: Proceedings of the Sixth International Conference on Learning Analytics & Knowledge (LAK '16), Association for Computing Machinery, 2016, pp. 329–338. doi:10.1145/2883851.2883944.

[5] L. P. Macfadyen, L. Lockyer, B. Rienties, Learning design and learning analytics: Snapshot 2020, Journal of Learning Analytics 7 (2020) 6–12. URL: https://learning-analytics.info/index.php/JLA/article/view/7389. doi:10.18608/jla.2020.73.2.

[6] D. Hernández-Leo, R. Martinez-Maldonado, A. Pardo, J. A. Muñoz-Cristóbal, M. J. Rodríguez-Triana, Analytics for learning design: A layered framework and tools, British Journal of Educational Technology 50 (2019) 139–152. doi:10.1111/bjet.12645.

[7] L. M. H. De Silva, M. J. Rodríguez-Triana, I.-A. Chounta, A needs analysis for institutional analytics, Studies in Higher Education (2025). URL: https://www.tandfonline.com/doi/full/10.1080/03075079.2025.2596129. doi:10.1080/03075079.2025.2596129.

[8] M. A. Chatti, A. Muslim, M. Guesmi, F. Richtscheid, D. Nasimi, A. Shahin, R. Damera, How to design effective learning analytics indicators? A human-centered design approach, in: Addressing Global Challenges and Quality Education: Proceedings of the 15th European Conference on Technology Enhanced Learning (EC-TEL 2020), Lecture Notes in Computer Science, Springer, 2020. doi:10.1007/978-3-030-57717-9_22.

[9] R. Alfredo, V. Echeverría, Y. Jin, L. Yan, Z. Swiecki, D. Gašević, R. Martinez-Maldonado, Human-centred learning analytics and AI in education: A systematic literature review, Computers and Education: Artificial Intelligence 6 (2024) 100215. doi:10.1016/j.caeai.2024.100215.

[10] S. Joarder, M. A. Chatti, Human-centred development of indicators for self-service learning analytics: A transparency through exploration approach, Journal of Learning Analytics 13 (2026) 163–189. URL: https://learning-analytics.info/index.php/JLA/article/view/8921. doi:10.18608/jla.2026.8921.
[11] D. E. Tzimas, S. N. Demetriadis, Implementing learning analytics in education: Enhancing actionability and adoption, Computers 15 (2026) 56. URL: https://www.mdpi.com/2073-431X/15/1/56. doi:10.3390/computers15010056.
[12] M. Khalil, P. Prinsloo, The lack of generalisability in learning analytics research: why, how does it matter, and where to?, in: Proceedings of the 15th International Learning Analytics and Knowledge Conference, Association for Computing Machinery, 2025, pp. 170–180. URL: https://dl.acm.org/doi/10.1145/3706468.3706489. doi:10.1145/3706468.3706489.
[13] Z. Xu, J. Olson, N. Pochinki, Z. Zheng, R. Yu, Contexts matter but how? course-level correlates of performance and fairness shift in predictive model transfer, in: Proceedings of the 14th International Conference on Learning Analytics and Knowledge, Association for Computing Machinery, 2024, pp. 713–724. URL: https://dl.acm.org/doi/10.1145/3636555.3636936. doi:10.1145/3636555.3636936.
[14] A. Ortiz-Beltrán, D. Hernández-Leo, Extracting institutional analytics features from lms data: Towards bridging learning design analytics and learning analytics, in: Proceedings of the Learning Analytics Summer Institute Spain 2025 (LASI Spain 2025), volume 4148 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2025. URL: https://ceur-ws.org/Vol-4148/Paper2.pdf.
[15] J. Dunlosky, K. A. Rawson, E. J. Marsh, M. J. Nathan, D. T. Willingham, Improving students' learning with effective learning techniques: Promising directions from cognitive and educational psychology, Psychological Science in the Public Interest 14 (2013) 4–58. doi:10.1177/1529100612453266.
[16] D. Kember, Interpreting student workload and the factors which shape students' perceptions of their workload, Studies in Higher Education 29 (2004) 165–184. doi:10.1080/0307507042000190778.
[17] S. Freeman, S. L. Eddy, M. McDonough, M. K. Smith, N. Okoroafor, H. Jordt, M. P. Wenderoth, Active learning increases student performance in science, engineering, and mathematics, Proceedings of the National Academy of Sciences 111 (2014) 8410–8415. doi:10.1073/pnas.1319030609.
[18] Z. A. Mohseni, I. Masiello, Co-designing, developing, and implementing multiple learning analytics dashboards for data-driven decision-making in education: a design-based research approach, Educational Technology Research and Development (2025). URL: https://link.springer.com/article/10.1007/s11423-025-10577-9. doi:10.1007/s11423-025-10577-9.