LEHRPULS · RESEARCH KIT

LehrPuls Research Kit

Data model, column contract, metrics, sample analyses, and notes on ethics and data protection for researchers in higher-education teaching and for theses. As of November 8, 2026, LehrPuls 1.21.1.

1. Purpose of This Kit

LehrPuls is a program for live questions in the lecture hall: the instructor asks questions from their own laptop, participants answer on their smartphones, and every answer is stored locally in a SQLite database together with its timestamp, response time, and grading. Over a semester this produces a data set that is of interest to research on higher-education teaching: the same question across several runs, first and second answers in Peer Instruction, response times per person, confidence ratings, and a topic assignment for every question.

This kit describes how to obtain these data, what every column means, how LehrPuls computes its metrics (so that your analysis agrees with the figures shown in the program), and which rules apply to consent, anonymity, and deletion. It includes a synthetic sample data set and one analysis script each in R and Python that run without a LehrPuls installation.

Everything in this document is taken from the source code of LehrPuls 1.21.1 and from the manual (sections 4.5, 8.1, 8.12, and 10.9).

2. Three Ways to the Data

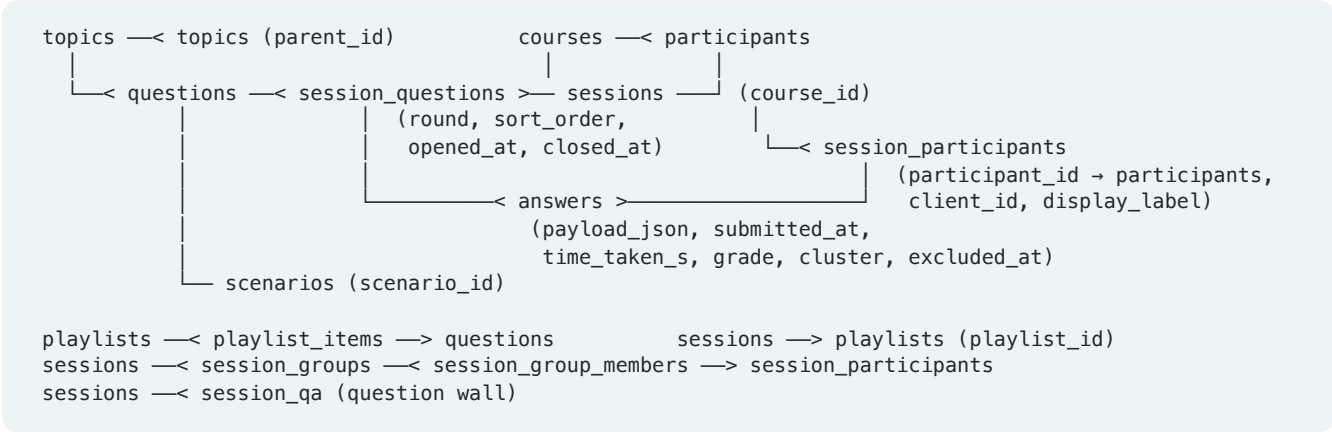
Way	What you get	When to use it
Raw data export (file)	Excel file with three sheets (answers, questions, sessions) or a CSV with the answers; one row per answer (long format)	One-time hand-over of a semester to a research project, archiving, SPSS, jamovi, JASP
Read-only API `/api/v1`	The same three tables with the same column names as JSON or CSV, plus three metric tables; read-only, key with expiry	Ongoing data collection, scripts that fetch after every lesson, dashboards
Sample data set of this kit	Synthetic CSV files in the format of the raw data export	Developing the analysis plan and scripts before real data exist; teaching methods

Both real ways deliver the same columns in the same order. Column names are fixed and language-independent, a contract that later versions only extend, never change (new columns are appended at the end). In LehrPuls the raw data export and the read-only API are available in expert mode only (Settings).

The database file `lehrpuls.db` itself is an ordinary SQLite file and opens in any SQLite tool. Its schema (section 3) is not a contract, however: it changes between program versions, answers are stored there as ungraded JSON, and grading (correct, points) is computed at export time against the current state of the question. Work with the export or the API.

3. Data Model

3.1 Relationship Diagram



Reading: `A ──< B` means “one A has many B”. The chain for every analysis is `answers → session_questions → questions → topics` (content) and `answers → session_participants → participants → courses` (person and course). Topics form a tree via `parent_id`; a topic analysis always aggregates the subtree.

3.2 Entities

Columns follow the LehrPuls 1.21.1 database schema (schema version 56). Listed are the columns that matter for analyses; internal columns (sort order, sample-data flag, translations `i18n_json`, hand-over identifier `uid`) are omitted.

topics (tree)

Column	Meaning
`id`	identifier
`parent_id`	parent topic; empty for a top-level topic
`name`	name; path in the export as “Main > Sub”
`archived`	1 = archived

questions (question pool)

Column	Meaning
`id`	identifier; stays the same across all uses (basis of the longitudinal view)
`topic_id`	topic; empty for ad-hoc questions without a topic
`qtype`	question type, one of 19 values (Appendix A)
`title`, `body_md`	title and question text (Markdown)
`options_json`	options, solution, tolerances per question type
`solution_md`	explanation of the solution
`difficulty`	instructor's estimate (integer); the program does not compare it with the measured difficulty
`time_budget_s`	planned duration in seconds
`tags`, `source`	keywords; origin of imported questions

Column	Meaning
`archived`	1 = archived; questions that have been used can only be archived, not deleted

courses and **participants** (course list)

Column	Meaning
`courses.id` , `name` , `semester`	A course belongs to exactly one semester (text, e.g. “WS 2026/27”); the same class in the next semester is a new course
`courses.purge_after`	deletion deadline (date): from then on LehrPuls reminds the instructor to remove the names (section 9)
`participants.id` , `course_id` , `display_name`	person on the course list; `import_ref` is the identifier from the imported list (e.g. student number)
`participants.active`	0 = left the course

sessions (one teaching unit)

Column	Meaning
`id`	identifier
`course_id`	course; empty for a session without a course
`join_mode`	`anonymous` , `pseudonym` , or `named` ; with `join_verified` = 1 (name plus student number checked) the export shows `verified`
`started_at` , `ended_at`	start and end, UTC
`playlist_id` , `playlist_name`	playlist the questions came from; the name survives deletion of the playlist
`question_lang`	language in which the questions were asked
`probe`	1 = trial run (not shown in analyses)

session_participants (one person's participation in a session)

Column	Meaning
`id`	identifier
`participant_id`	reference into the course list; empty for anonymous participation or guests
`display_label`	displayed name or pseudonym
`client_id`	device identifier; `plenum` is the instructor's plenary answer, `group–…` a group. Neither counts as a person in any statistic

session_questions (one question asked in a session)

Column	Meaning
`id`	identifier (`session_frage_id` in the export)
`question_id`	the question from the pool
`round`	round; 1 = first vote, 2 = second round in Peer Instruction (its own row with its own position)
`sort_order`	position in the session; in the export `frage_nr` = `sort_order` + 1
`opened_at` , `closed_at` , `revealed_at`	question opened, closed, solution revealed (UTC); open time = `closed_at` – `opened_at`

Column	Meaning
`excluded_at`	set when the instructor excluded the question from analysis
`shuffle`	set when the answer options were shown shuffled in this session

answers (one answer)

Column	Meaning
`session_question_id` , `session_participant_id`	question and person; at most one answer per pair (unique)
`payload_json`	raw answer as JSON, e.g. `{"choice":"a"}` , `{"choices":["a","c"]}` , `{"value":38.6}` , `{"text":"…"}`
`submitted_at`	time of submission (UTC)
`time_taken_s`	response time measured on the person's device; may be missing
`grade` , `grade_comment` , `graded_at`	manual grade in percent with justification (open formats)
`cluster`	answer group into which the instructor sorted an open answer
`excluded_at`	set when this answer was excluded from analysis

Further tables: `playlists` and `playlist_items` (question blocks and their order), `playlist_time_log` (stopwatch per playlist node), `session_groups` and `session_group_members` (group phases), `session_qa` (question wall with timestamp, reach, votes), `scenarios` (case descriptions shared by several questions), `einstellungen` (key-value settings, including the minimum group size of the analyses).

3.3 What the Export Makes of the Model

The export joins the chain of section 3.1 into one flat table, grades every answer against the **current** state of the question (`korrekt` , `punkte`), and replaces persons with stable pseudonyms. Two properties are decisive for research designs:

- **Grading against the current state of the question.** Correcting a solution afterwards retroactively changes the statistics of all earlier uses. For before-and-after comparisons of a revised question, create the new version as a new question (new `frage_id`).
- **Pseudonyms are stable only through the course list.** The anonymizer assigns the numbers T1, T2, … in order of first appearance; the key is the person's course-list identifier (`participant_id`). For anonymous participation or guests the key is the session's participation row, so the same person gets a new number in every session. Longitudinal analyses **per person** therefore require named participation (`named` or `verified`); longitudinal analyses **per question** work in every join mode.

---

4. Column Contract of the Raw Data Export

The column names are a fixed contract of the raw data export. `/api/v1/antworten` , `/api/v1/fragen` , and `/api/v1/sessions` deliver the same columns. Column names are German and identical in every language of the user interface.

4.1 Table “Antworten” (answers, one row per answer)

No.	Column	Type	Meaning
1	`nutzer_id`	text	Random identifier of this LehrPuls installation (UUID). Sessions, questions, and participants are unique only together with it when you merge exports of several instructors
2	`nutzer`	text	License name; empty with “user anonymized”
3	`session_id`	number	session
4	`session_start_utc`	text	start of the session, `YYYY-MM-DD HH:MM:SS` in UTC
5–7	`kurs_id` , `kurs` , `semester`		course; empty without a course
8–9	`playlist_id` , `playlist`		playlist; empty for ad-hoc questions
10	`beitrittsmodus`	text	join mode: `anonymous` , `pseudonym` , `named` , or `verified`
11	`teilnehmer_id`	text	stable pseudonym `T1` , `T2` , … (anonymized) or key `P<participant_id>` / `S<session_participant_id>` (export with real names)
12	`teilnehmer`	text	“Participant 1” … or the displayed name; for groups the group name
13	`gruppe`	0/1	1 = the row is a group answer, not a person
14	`session_frage_id`	number	identifier of the question as asked in this session (each round has its own)
15	`frage_nr`	number	position in the session, from 1
16	`runde`	number	round: 1 or 2 (Peer Instruction)
17	`frage_id`	number	question in the pool; identical across all sessions
18	`fragetyp`	text	question type, Appendix A
19	`frage_titel`	text	title
20–21	`thema_id` , `thema`		topic; `thema` as path “Main > Sub”; empty without a topic
22	`antwort`	text	answer, readable (option texts, number with unit, free text, image point in percent)
23	`antwort_zahl`	number	numeric, estimate, and scale questions only: the numeric value
24	`antwort_roh`	text	raw answer as JSON (`payload_json`)
25	`korrekt`	0/1	correct, all-or-nothing; empty for question types that cannot be graded
26	`punkte`	number	score 0 to 1 with partial credit (multiple selection, order, matching, cloze); follows `wertung_hand` where set
27	`dauer_s`	number	response time in seconds; may be empty
28	`gesendet_am_utc`	text	time of submission, UTC
29	`ausgeschlossen`	0/1	1 = question or answer is excluded from analysis. The export keeps the row; filter to 0 when computing yourself
30–31	`wertung_hand` , `kommentar`		manual grade in percent and justification; empty = none
32	`skizze_svg`	text	sketch of an open answer as SVG; otherwise empty

No.	Column	Type	Meaning
33	`optionen_reihenfolge`	text	with shuffled options: option identifiers in the order shown, e.g. `c,a,d,b`
34	`datei`	text	name of an attached file; the file itself is in the media store
35–36	`bilanz_richtig`, `bilanz_gewertet`	number	the person's tally in the whole session (section 6.7); identical in every row of the person; empty for groups
37	`antwortgruppe`	text	answer group into which the instructor sorted an open answer
38	`konfidenz`	text	confidence: `sicher` (sure), `unsicher` (unsure), or `geraten` (guessed); empty if the question did not ask or nothing was chosen
39	`segment`	text	the person's attribute from the session's attribute question (e.g. “ja” for “attended the lecture”); empty without one
40 ff.	`fragensprache`, `frage_titel_en`, `antwort_en`, `thema_en`, `playlist_en`, then per further language `frage_titel_<code>`, `antwort_<code>`, `thema_<code>`, `playlist_<code>`	text	language of the session and versions in other languages; English is always present, further languages only if an exported question has a version in them

4.2 Table “Fragen” (codebook, one row per question that occurs in the answers)

`nutzer_id`, `nutzer`, `frage_id`, `fragetyp`, `titel`, `text`, `optionen`, `loesung` (solution), `thema_id`, `thema`, `vorgesehene_dauer_s` (planned duration), `schwierigkeit` (instructor's estimate), `tags`, `titel_en`, `text_en`, `optionen_en`, `loesung_en`, `mischen` (0/1, “shuffle answers per session”; empty for types without options), `alles_oder_nichts` (0/1, all-or-nothing, multiple selection only), then per language `thema_<code>`.

4.3 Table “Sessions” (one row per completed session)

`nutzer_id`, `nutzer`, `session_id`, `start_utc`, `ende_utc`, `kurs_id`, `kurs`, `semester`, `playlist_id`, `playlist`, `beitrittsmodus`, `teilnehmende` (persons, excluding plenary and groups), `fragen` (questions actually asked), `fragensprache`, then per language `playlist_<code>`.

4.4 File Formats

	Raw data export (file)	Read-only API `?format=csv`
Separator	semicolon `;`	comma `,`
Decimal mark	follows the display language: comma for German, period for English	always period
Encoding	UTF-8 with byte order mark, CRLF line endings	UTF-8
Scope	XLSX: three sheets answers, questions, sessions; CSV: answers only	one table per request
Empty values	empty cell	empty cell (CSV) or `null` (JSON)

In R: `read.csv(file, sep = ";", dec = ",", fileEncoding = "UTF-8-BOM")` for the German export; in Python: `pd.read_csv(file, sep=";", decimal=",", encoding="utf-8-sig")`. The scripts in this kit detect the decimal mark themselves.

5. The Read-Only API `/api/v1`

Source: manual 10.9.

Operation. The API is off until the instructor turns it on in Settings (expert mode) and answers only while LehrPuls is running. It listens on `127.0.0.1` on the first free port from 8716 to 8719; with the switch “reachable on the network” also on the computer's other addresses, unencrypted. IT can disable it entirely with the environment variable `LEHRPULS_API_SPERRE=1`. It is read-only: GET requests only.

Access. An access key (`lp_…`) is valid for 1, 7, 30, or at most 90 days; LehrPuls stores only its SHA-256 hash. The key goes into the header of every request: `Authorization: Bearer lp_…`. It is not accepted in the URL. Without a valid key the API answers 401, for an unknown table 404, for invalid parameters 400, each with a JSON object containing `code` and `message`. CORS is enabled, so browser pages may call it directly.

Defaults when switched on: participants anonymized **on**, user anonymized off, include samples off, apply minimum group size off, timestamps exact, reachable on the network off.

Endpoints

Address	Content
`/api/v1`	program name, version, `api: 1`, key expiry (`gueltig_bis_unix`), whether anonymized, list of tables and parameters
`/api/v1/sessions`	one row per completed session (columns as in 4.3)
`/api/v1/fragen`	one row per question occurring in the answers (columns as in 4.2)
`/api/v1/antworten`	one row per answer (columns as in 4.1)
`/api/v1/fragen-kennzahlen`	one row per asked question with item quality (below)
`/api/v1/themen`	one row per topic with mastery and coverage
`/api/v1/bilanz`	one row per person and session with correct out of graded

Parameters (after the question mark, joined with `&`): `format=json|csv`; `session_id=<id>|letzte|heute` (also `session=`, `last`, `latest`, `today`); `von=`, `bis=` (from, to) as `YYYY-MM-DD`, `heute` (today), or `gestern` (yesterday); `tage=<1…3660>` (the last N days including today); `kurs=<name>` or `kurs_id=`; `thema=<name>` (including subtopics, in any language) or `thema_id=`. Unknown parameters are rejected with 400, as are ambiguous topic names (with the identifiers listed). `fragen-kennzahlen` and `themen` compute over a period and do not accept `session_id`.

Metric tables (columns are a contract):

- `fragen-kennzahlen`: `frage_id`, `frage_titel`, `fragetyp`, `thema_id`, `thema`, `einsaetze` (sessions in which the question ran), `antworten` (answers), `schwierigkeit_p` (difficulty, 0–1, three decimals), `trennschaerfe` (discrimination, −1 to 1), `tote_distraktoren` (dead distractors), `pi_p_runde1`, `pi_p_runde2` (empty without a second round), `zuletzt_gestellt` (last asked).
- `themen`: `thema_id`, `thema`, `ebene` (level, 0 = top), `antworten`, `reifegrad_pct` (mastery), `gespielte_fragen` (questions played), `aktive_fragen` (active questions), `abdeckung_pct` (coverage). The row with an empty topic is the catch-all “without topic”.
- `bilanz`: `session_id`, `session_start_utc`, `kurs_id`, `kurs`, `teilnehmer_id`, `teilnehmer`, `richtig` (correct), `gewertet` (graded), `quote_pct`.

The protection switches “omit questions below the minimum group size” and “round timestamps” apply to `antworten`, `sessions`, and `bilanz` (section 9.3); the other two metric tables only deliver aggregates anyway.

Ready-made examples in R, Python, Power Query, and as a browser page are on the download page of lehrpuls.com under “Tools for your own analyses” (`lehrpuls-kennzahlen.R` , `lehrpuls-kennzahlen.ipynb` , `lehrpuls-kennzahlen.pq` , `lehrpuls-api-demo-en.html`). The scripts of this kit build on them: with `schnittstelle` as the first argument they fetch the answers table directly and compute the same metrics as on the sample data set.

---

6. Metrics: Definitions and Formulas

This is how LehrPuls computes. Apply the same rules and your values agree with the “Questions” tab of the analyses and with `/api/v1/fragen-kennzahlen` . The basis is always: `ausgeschlossen = 0` , `gruppe = 0` , no plenary answers, completed sessions only, no trial runs.

Two thresholds apply throughout: the **minimum base** n = 10 answers, below which the cockpit grays out metrics as “too little data” and gives no hints (constant `MINDESTBASIS`), and the **minimum group size** for distributions (default 5 persons, adjustable 1 to 100), below which the analyses show no answer distribution per question and which acts as a protection switch at export.

6.1 Difficulty P

P = number of correct answers ÷ number of gradable answers, **round 1 only**, **all-or-nothing** (`korrekt`), over all sessions in the period. An answer is gradable when `korrekt` is set. Cockpit hints: P > 0.9 “too easy”, P < 0.3 “too hard”. A high value means an easy question (probability of solving, not difficulty in the literal sense).

6.2 Discrimination (upper-lower index)

Per session: each person's total score = number of correct answers in round 1. Persons are sorted by score (ties by participation identifier); the upper and lower groups each comprise g = round(0.27 · n) persons, at least 1, at most n ÷ 2. For every question in the session: D = proportion correct in the upper group − proportion correct in the lower group. The question's discrimination is the **mean of D over all sessions** in which the question ran. Warning sign: D ≤ 0 with n ≥ 10 (“the strong students fail”: question ambiguous or wrong key).

For research purposes the upper-lower index is coarse: in short sessions (8 questions) the question itself contributes heavily to the total score. The scripts additionally compute the **point-biserial correlation** between `korrekt` and the person's rest score in the session (total score without this question).

6.3 Dead Distractors

A wrong option of a choice question counts as dead when it was chosen in round 1 less often than max(1, 2 % of the gradable round-1 answers), and only if $n_1 \geq 10$. `tote_distraktoren` counts these options per question. Frequently chosen wrong options (typical misconceptions) can be read from `antwort_roh` or `antwort` .

6.4 Peer Instruction Gain

`pi_p_runde1` = proportion correct over all gradable answers to the question in round 1 in the period, `pi_p_runde2` the same for round 2; reported only for questions with answers in both rounds. Note that round 1 here also counts sessions in which no second round took place. For a paired analysis per person, join the rows with `runde = 1` and `runde = 2` of the same `session_id` , `frage_id` , and `teilnehmer_id` and count the switches wrong → right and right → wrong (McNemar test); this is what the scripts do. In LehrPuls a second round is a separately asked question (`session_frage_id` , `frage_nr`) with `runde = 2` .

6.5 Score and Topic Mastery

The **score** of an answer is `punkte` (0 to 1, with partial credit). The **mastery** of a topic (`reifegrad_pct`) is the mean score of all gradable answers to questions in the topic's subtree in the period, in percent. Cockpit levels: below 50 % “building up”, 50–75 % “in progress”, from 75 % “consolidated”; grayed out below n = 10. **Coverage** (`abdeckung_pct`) = active questions played ÷ active questions in the subtree × 100. Both figures deliberately stand side by side and are not combined into one index.

6.6 Response Times

Three notions of time: **response time** = `dauer_s` per answer (measured on the device from opening the question to submitting; analyze robustly with the median, discard missing values), **open time** = `closed_at` – `opened_at` of the asked question (database and read-only API only, not in the raw data export), **timestamp** = `gesendet_am_utc` , which gives the answer's position within the unit relative to `session_start_utc` . Response times are right-skewed; compute on the logarithmic scale.

6.7 Participants' Tally

`bilanz_richtig` out of `bilanz_gewertet` per person and session: graded questions are those that were revealed and are gradable or were graded by hand; in Peer Instruction round 2 counts; excluded items do not count. `quote_pct` = correct ÷ graded × 100. The tally covers the whole session even when the export is filtered by topic.

6.8 Question Types That Cannot Be Graded

A `korrekt` value exists for single and multiple choice, number, order, matching, cloze, formula, the image question (only with target regions), and messages with a yes/no expectation. All other types — open question, poll, estimate, word cloud, ideas, scale, ranking, point budget, 2×2 grid — carry no `korrekt` (except with a manual grade). For them, **participation** (answers ÷ participants of the session) replaces correctness; evaluate estimate questions as deviation from the reference value (`antwort_zahl`).

---

7. Sample Analyses

The scripts `analyse.R` (base R only, no packages) and `analyse.py` (pandas, matplotlib) load the sample data set or your own export, recompute the metrics of section 6, and write `kennzahlen-fragen.csv` , `laengsschnitt.csv` , and a four-panel figure (`analyse-R.png` , `analyse-Python.png`). Both scripts were run on the sample data set on November 8, 2026 (R 4.6.1; Python 3.14 with pandas via `uv run --with pandas --with matplotlib analyse.py`) and deliver identical values for P, Peer Instruction gain, and the longitudinal view; discrimination differs slightly (question 102: −0.27 in R, −0.23 in Python) because R and Python break ties differently when sorting total scores. LehrPuls itself breaks ties by participation identifier. Output labels of the scripts are German.

Analysis	What is computed	Result on the sample data set
Item analysis (quality map)	P and discrimination per question, hints as in the cockpit, point-biserial correlation	40 questions; one question with discrimination around −0.25 (wrong key), one too easy (P = 0.91), three too hard
Peer Instruction gain	`pi_p_runde1` / `pi_p_runde2` as in LehrPuls; paired per person with switch counts and McNemar test	16 questions with a second round; mean paired gain +0.22 (proportion correct)
Longitudinal view of the same question	P per question and semester (round 1 only), range across semesters	two questions become harder every semester (drift 0.70 → 0.43), one easier
Response time distributions	median per question type; single choice by round and correctness; histogram on log scale	median 17 s single choice, 73 s open question; wrong answers take longer (24 s vs. 17 s); round 2 is faster

Further analyses the data set is prepared for: **calibration** (proportion correct per `konfidenz` level `sicher`, `unsicher`, `geraten`, Brier score; six questions per course ask for confidence), **attribute comparison** (`segment` "ja"/"nein" for "attended the lecture" in the third semester of the first course), **participation trajectory** (answers per `frage_nr` relative to `teilnehmende` from the sessions table), **distractor distribution** from `antwort_roh`.

Invocation (in the `skripte` folder): `Rscript analyse.R` or `python3 analyse.py`; with the path to your own CSV as argument; or `LEHRPULS_SCHLUESSEL=lp_… Rscript analyse.R schnittstelle "kurs=Physics 1&tage=120"` to fetch directly from a running API (not tested against a running instance; the fetch logic matches `lehrpuls-kennzahlen.R`).

---

8. The Sample Data Set

`beispieldaten/antworten.csv` (18,870 rows, 44 columns), `fragen.csv` (40 questions), and `sessions.csv` (30 sessions) in the format of the raw data export with German user interface. Generated by `erzeuge-beispieldaten.mjs` with a fixed seed (`node erzeuge-beispieldaten.mjs` produces byte-identical files).

Structure: two classes ("Wirtschaftsinformatik 1", join mode `named`; "Rechnungswesen 2", join mode `verified`) over three semesters (WS 2025/26, SS 2026, WS 2026/27), one course per semester with 62 to 118 participants (533 persons in total, pseudonyms T001 to T533), five sessions per course and semester two weeks apart, eight questions per session rotating through a pool of 20 questions per class (14 single choice, 2 multiple selection, 2 numeric, 1 open, 1 estimate). Every question therefore runs two to three times per semester and six to eight times overall.

Generating model: ability θ per person normally distributed with a small cohort offset per semester; probability of solving logistic in a·(θ − b) with difficulty b and discrimination a per question. Eight single-choice questions per class are Peer Instruction candidates; a second round takes place when 30 to 70 % were correct in round 1 (the manual's rule of thumb). In round 2 a correct answer stays correct with 92 %, a wrong one becomes correct with a probability that rises with the proportion of correct first answers. Response times are lognormal, shorter with high ability and correct answers, shorter in round 2. Confidence correlates with correctness.

Built-in effects to rediscover: question 102 (Normalformen) has negative discrimination; 104 (SELECT mit WHERE) is too easy; 107 (NULL in SQL) and 207 (Zuschlagssatz) too hard; 109 and 209 become harder every semester, 111 and 211 easier (drift); 113 and 213 have a dead distractor; 106 and 206 a particularly attractive wrong distractor; about 0.3 % of rows are excluded (`ausgeschlossen = 1`); 30 % of open answers carry a manual grade.

Limits: all content is fictitious (question texts without subject content, options “Option A” to “Option D”, solution always A). There are no group answers, sketches, files, or languages beyond English. Pseudonyms are written with leading zeros (T001); the real anonymized export writes T1, T2, …

---

9. Ethics and Data Protection

9.1 What LehrPuls Stores and Where

All data reside on the instructor's computer (SQLite file in the data folder); LehrPuls transmits nothing to the manufacturer or third parties. In online teaching the connection runs end-to-end encrypted through a relay that cannot read the content. Participants need no account. Personal data are: names from the course list (`participants.display_name`, `import_ref`), pseudonyms people choose themselves, and whatever people write in open answers, questions on the question wall, or uploaded files. Linking answers to persons is possible only with join mode `named` or `verified`.

9.2 Consent

Taking part in live questions is teaching, not research. As soon as answers are to be analyzed or published for a study, informed consent according to your institution's rules is required. Four recommendations: obtain consent in the dependent relationship instructor–students through a **third person**; **seal the linking list** until after grades are issued; disclose the **conflict of interest** when developers or vendors of LehrPuls take part in the study; **preregister the analysis plan** (e.g. OSF). LehrPuls provides a sample privacy notice in the header of the course page for adaptation (manual 4.5); it is a template, not legal advice.

9.3 Anonymization and Protection Before Hand-Over

Before data leave the computer, LehrPuls offers four levels, all of which also apply to the API:

1. **Participants anonymized** (default on): names are replaced by “Participant 1”, “Participant 2” …, identifier `T1`, `T2` …; the same person keeps their number across all sessions of the export provided they joined through the course list. Numbers are reassigned per export, so two exports cannot be linked via T numbers.
2. **User anonymized**: the license name (instructor) stays empty; the random installation identifier remains.
3. **Omit questions below the minimum group size**: answers to asked questions (`session_frage_id`) with fewer answers than the analyses' minimum group size (default 5) are dropped.
4. **Round timestamps** to the hour or the day: an exact timestamp together with the timetable can reveal a person. Rounded are `session_start_utc`, `gesendet_am_utc`, `start_utc`, and `ende_utc`; the response time stays.

Free text is not anonymized: a name in an open answer, a comment, or a file stays there. Review open answers before hand-over, or drop the columns `antwort`, `antwort_roh`, `kommentar`, `skizze_svg`, and `datei` for open questions.

9.4 Minimum Group Size in Publications

Do not report any metric based on fewer than the minimum base; LehrPuls itself shows no hints below n = 10 and grays out values. For subgroups (segment, confidence level, course × semester) the threshold applies per cell. For planning studies, guide values from a power analysis apply (α = 0.05, power 0.80): paired comparison with d = 0.3 about 90 persons, independent groups with d = 0.5 about 64 per group, correlation r = 0.3 about 85 persons; multilevel models with sessions or courses as the second level from 50 units.

9.5 Deletion Deadlines

A deletion deadline can be set per course (`purge_after` , typically the end of the semester). Once reached, the start page reminds the instructor; “Delete now” removes the names from the course list, replaces them in the sessions with placeholders, strips the sender from questions on the question wall, deletes the course's group compositions, and empties the recycle bin for this course. **Sessions and answers remain, anonymized**, and so do longitudinal views per question. What participants wrote or uploaded themselves stays unchanged. Plan research exports **before** the deletion deadline if you need to link persons across sessions; afterwards linking is possible only within the placeholder numbers; the placeholders are tied to the course seat and therefore remain stable across sessions.

---

10. Open Research Questions

Ten research questions that can be addressed most directly with LehrPuls data. Each entry names the columns that answer it.

1. **Peer Instruction outside the natural sciences** (replication of Smith et al. 2009 in business informatics, business administration, law): `runde` 1 and 2 per `teilnehmer_id` and `frage_id` , `korrekt` , plus isomorphic follow-up questions as their own `frage_id` .
2. **Response time as a learning indicator and the “80 % rule”** for closing the vote: `dauer_s` per answer in both rounds, `korrekt` , `gesendet_am_utc` relative to `session_start_utc` .
3. **Metacognitive calibration at item level**: `konfidenz` (sure, unsure, guessed) against `korrekt` per question, change over the semester and between `runde` 1 and 2.
4. **Psychometrics of formative live items**: `schwierigkeit_p` , `trennschaerfe` , `tote_distraktoren` from `/api/v1/fragen-kennzahlen` , distractor choice from `antwort_roh` , stability across `einsaetze` ; with consent, relation to exam performance via `teilnehmer_id` .
5. **AI-generated versus self-written questions**: the questions' `source` column (origin file of imported questions) against P, discrimination, distractor function, `dauer_s` , and Peer Instruction gain.
6. **Showing the distribution before discussion or not**: switching behavior between `runde` 1 and 2 by condition; the condition itself currently has to be logged outside LehrPuls; per question only the time of revealing is stored (`revealed_at`), from which you can tell whether the distribution was on the wall before round 2.
7. **Attention trajectory across the unit**: participation (answers per `frage_nr` ÷ `teilnehmende`), `dauer_s` , and `korrekt` as a function of position in the session (`gesendet_am_utc` – `session_start_utc`), across many sessions with `playlist` as a control.
8. **Anonymous versus named participation**: `beitrittsmodus` (anonymous, pseudonym, named, verified) as a session condition in a crossover against participation, dispersion of answers, `dauer_s` , length of open answers.
9. **Drift of the same question across cohorts**: `frage_id` × `semester` (round 1 only), separating cohort effects (`kurs` , `semester`) from item drift (question becoming known, outdated context).
0. **Revision cycle of questions**: discrimination and Peer Instruction suitability (30–70 % correct in round 1) before and after a revision; `antwortgruppe` of open answers as a source of new distractors; revised version as a new `frage_id` .

11. Collaboration Offer

We support research projects and theses working with LehrPuls data: with test licenses for the duration of data collection, with explanations of the data model and metrics, with feedback on analysis plans, and, where a project needs a product addition (such as logging an experimental condition per question), with checking whether it fits into a next version. Data remain with the institution; we receive no raw data unless you explicitly give us some.

Contact: office@lehrpuls.com. Please state the research question, discipline, number of courses and participants, and the planned period.

Appendix A: Question Types

`qtype` (19 values, schema 1.21.1): `mc_single` (single choice), `mc_multi` (multiple selection), `numeric` (numeric answer), `estimate` (estimate), `scale` (scale), `open` (open question, also with sketch or file), `cloze` (cloze text), `order` (order), `ranking` (ranking), `matching` (matching), `hotspot` (image question), `poll` (poll), `poll_multi` (poll with multiple selection), `info` (message), `wordcloud` (word cloud), `budget` (distribute points), `ideas` (collect ideas), `grid` (grid), `formula` (formula).

Appendix B: Value Lists

Column	Values
`beitrittsmodus`	`anonymous`, `pseudonym`, `named`, `verified`
`konfidenz`	`sicher`, `unsicher`, `geraten`, empty
`runde`	1, 2
`korrekt`, `gruppe`, `ausgeschlossen`, `mischen`, `alles_oder_nichts`	0, 1 (or empty)
timestamps	`YYYY-MM-DD HH:MM:SS` in UTC; rounded `YYYY-MM-DD HH:00:00` or `YYYY-MM-DD`
rounding in the API	`keine` (none), `stunde` (hour), `tag` (day)

Appendix C: Limits of the Data

- Answers are graded against the current state of the question (section 3.3).
- `dauer_s` is measured on the person's device and may be missing; connection drops lengthen it.
- Ad-hoc questions without a topic are missing from every topic statistic (catch-all “without topic”).
- The instructor's plenary answers and group answers are not persons.
- A person can answer each asked question only once per session; corrections overwrite the answer (`updated_at` in the database).
- The open time of a question (`opened_at`, `closed_at`) is available in the database only, not in the export.
- T numbers are not stable between exports; within one export only for participation through the course list.