

Appendix D

AI-Use Documentation Template

Dr. Jose Mendoza, Academic Director and Clinical Associate Professor

Version 1.0 · August 2026

Except where otherwise noted, this appendix is licensed under CC BY 4.0.

Appendix Information

PURPOSE

This appendix is the standard record for AI-assisted work throughout the guide and the course. It exists so that an analysis produced with assistance can be audited afterward by someone who was not present: what was asked, what came back, what was predicted before it was run, what was kept, what was rejected, what check was performed, and what the analyst decided that the tool could not decide. It supplies two levels of record — a compact one that lives beside the work in a notebook, and a full one submitted with graded work — built on one field set that does not change between them.

VERSION AND DATE

Version 1.0 · August 2026 · Language: English (United States)

SUGGESTED CITATION

Mendoza, J. (2026). AI-use documentation template. In *Applied business analytics for marketing decision-making: Business analytics and data visualization* (Appendix D, Version 1.0) [Open educational resource]. CC BY 4.0.

LICENSE AND RIGHTS

Except where otherwise noted, this appendix is licensed under a Creative Commons Attribution 4.0 International License. Copyright © 2026 by Jose Mendoza.

Google Colab, Google Drive, Google Docs, and Google Sheets are products of Google LLC. Microsoft Word and Microsoft Excel are products of Microsoft Corporation. Tableau Desktop and Tableau Prep are products of Salesforce, Inc. ChatGPT is a product of OpenAI, Claude is a product of Anthropic, Gemini is a product of Google LLC, and GitHub Copilot is a product of GitHub, Inc. Product names are used for identification only and do not imply endorsement. StyleCraft Collective is a fictional company created for instruction.

GENERATIVE AI USE

Generative artificial intelligence and other AI-assisted tools were used in the research, writing, revision, and production of this appendix, including outlining, preliminary drafts, revision of prose, and document formatting. These tools were used under the author's direction and are not credited as authors, researchers, or sources. The author determined the appendix's scope, boundaries, field set, and content, and reviewed and approved all AI-assisted material: every field in the published record was traced to the chapter that requires it, and every numerical value in this appendix and its companion worked example was recomputed from the published arithmetic of the chapter that owns it and confirmed. Responsibility

for the accuracy, originality, and final form of this appendix rests entirely with the author. A fuller statement appears in the front matter of the complete guide.

COMPANION FILES

Editable forms in Word, spreadsheet, and Markdown, all carrying the same field names, together with one assignment documented from beginning to end, are in the [Applied Business Analytics companion repository](#).

How to Use This Appendix

This appendix has nine numbered sections. Sections D.1 and D.2 explain what the record is for and which of its two levels a given piece of work needs. Section D.3 defines the field set once, and every form in this appendix and in the companion repository is built from it. Sections D.4 and D.5 define the two record levels; editable versions of both are in the companion repository. Section D.6 sets a weak record beside a strong one. Sections D.7 and D.8 cover evidence links and the three file formats. Section D.9 states what each chapter asks you to file.

Two conventions carry through. Placeholders appear in square brackets and capital letters — [TOOL], [EXCHANGE ID], [EXPECTED TOTAL] — and a submitted record that still contains one has not been completed. And every field here is answerable in a phrase or a sentence; a field taking a paragraph usually means the exchange was two exchanges.

The record is also written during the work rather than after it. That is the only way the prediction fields can be honest, and it is why the compact form of Section D.4 exists at all. A record reconstructed from memory the night before a deadline looks like documentation and contains none of the evidence documentation is for.

Table D.1

What to complete, by what you are submitting

If you are...	Complete	Section
Working in a notebook, cell by cell	The compact record, beside the cell	D.4
Submitting a graded assignment or lab	The full record	D.5
Submitting a group project	One consolidated full record, with the contribution table	D.5.8
Submitting a dashboard and memo (Chapter 13)	The full record, filed inside the memo's appendix	D.5, D.9
Deciding how much of this a small exchange needs	The proportional rule	D.2.3
Unsure what a field means	The canonical field set	D.3

If you are...	Complete	Section
Unsure what counts as verification	The evidence types	D.5.5
Working in Word, a spreadsheet, or Markdown	The format your submission needs	D.8

D.1 What the Record Is For

Chapter 1 introduced four habits for AI-assisted work — specify, predict-then-verify, explain, document — and gave the fourth a single sentence: documentation should record the tool used, the prompt, the output, the revisions made, the errors or limitations found, and the verification performed. This appendix is that sentence turned into a form, and it is worth being clear about why the form is shaped the way it is.

The obvious purpose — disclosure — is the least important one. A reader who learns only that an assistant was used has learned almost nothing, because the interesting question is never whether AI contributed. It is whether anyone checked. A record that says “used ChatGPT for the code” discloses a fact and establishes nothing; a record that says which cell, what was predicted before it ran, what the row count actually did, and what changed as a result establishes that the work was audited by a person who can be asked about it.

In other words, this appendix documents verification, and mentions AI because AI is what made the verification necessary. That is why the fields carrying the most weight are the prediction, the error found, and the check performed, and why the field naming the tool is one of eleven rather than the point of the exercise. It is also why this record is not a citation: citing a generative tool is a separate obligation with its own conventions, which change as the tools do (McAdoo, 2023). Underneath all of it is a division of labor the guide has been building since Chapter 1. AI changes which parts of analytical work are cheap and which are scarce, and the scarce part is judgment about what an output means and whether it can be trusted (Davenport et al., 2020). Fields D6 through D11 are where that judgment leaves a trace.

The second purpose is narrower. Analytical work is revisited: a number is questioned in a meeting, a model is reused for a decision it was never graded on, a chart is quoted with a stronger verb than it can carry. The question then is always some form of “how do we know this?” The record is the answer, and it has to be legible to someone who was not in the room — which, six weeks later, includes you.

CONCEPT

The Record Is Evidence, Not Disclosure

A completed record does not establish that the work is correct. It establishes that the work was checked, by a named person, against something specific, before it was submitted. Those are different claims, and only the second one a form can support.

The practical test is Chapter 1's: could you answer, without rerunning anything, why the submitted file differs from what the assistant returned? If yes, the record is doing its job. If no, the record is a receipt.

Source: Course concept developed for this guide, informed by the documentation and human–AI configuration practices in National Institute of Standards and Technology (2024).

D.2 The Two Levels

The record comes in two levels because AI-assisted work happens at two scales. A student debugging a `KeyError` needs a line of record beside the cell; a student submitting a churn model with a threshold derived from stated costs needs a document. One form cannot serve both without being too heavy for the first case or too thin for the second.

Both levels use the same field names, and the compact record is a strict subset of the full one. A notebook full of compact records therefore supplies the evidence needed to assemble the exchange log of Section D.5.4 without reconstructing the exchanges from memory, which is the whole reason the field names are held constant across the two.

Table D.2

The two levels compared

	Level 1: compact notebook record	Level 2: full submission record
Where it lives	In a text cell beside the code, chart, or paragraph it produced	A separate document, or an appendix to the deliverable
Unit	One exchange	One assignment or project, covering every meaningful exchange
Fields	Exchange ID plus Chapter 1's four questions; Prediction recorded first from Chapter 4 onward	All eleven fields of Table D.4, plus tools, privacy, verification, dispositions, and the signature
Written	While the work is happening	Assembled at submission from the compact records
Introduced in	Chapter 1, Table 1.6	Chapter 1, Section 1.7

	Level 1: compact notebook record	Level 2: full submission record
Required by	Every lab from Chapter 4 onward	Every AI-assisted homework submission

D.2.1 Level 1: The Compact Notebook Record

Level 1 is Table 1.6 of Chapter 1: four questions, answered in a text cell next to the work. It exists because the fields that matter most decay fastest. Nobody remembers, two weeks later, what they predicted the row count would be, and a prediction reconstructed after the fact is not a prediction.

From Chapter 4 onward the labs add a fifth line, because those chapters require the prediction to be written before a delegated step runs. Section D.4 gives the form.

D.2.2 Level 2: The Full Submission Record

Level 2 is the document submitted with graded work. It consolidates the meaningful exchanges rather than every conversational turn, adds the tools and the data declaration, collects the verification evidence in one table, requires an account of what was rejected, and ends with a signed statement. Section D.5 gives the form, section by section.

“Meaningful” is doing real work in that sentence. Three rounds of the same debugging exchange are one entry, provided the final prompt, the correction, and the check are preserved. A threshold you accepted from an assistant is one entry on its own, however short the exchange was, because a threshold is a cost assumption and Chapter 9 requires it to be re-derivable.

D.2.3 The Proportional Documentation Rule

How much record a given exchange needs follows from the consequence of getting it wrong, not from how long the exchange was. A twenty-minute conversation about how to phrase an axis label may need one line; a three-line exchange that set a decision threshold needs the full treatment.

Table D.3

The proportional documentation rule

The exchange...	Record	Because
Suggested a syntax correction that changed no number	Compact record, one line, with the check that confirmed nothing else moved	The failure mode is a silent side effect, and one rerun catches it
Drafted a step that changed a count, a total, or a grain	Compact record with the predicted and actual figures, carried into the exchange log	Chapter 4 makes the before-and-after counts non-negotiable

The exchange...	Record	Because
Proposed a definition, a threshold, an eligibility rule, or a window	Full entry, with the choice independently re-derived or checked against the governing source, business rule, or arithmetic	These are the choices the tool is least entitled to make and the meeting is most likely to question
Produced a model, a forecast, an experiment analysis, or a scored list	Full entry, plus its own rows in the verification table	The audit is the deliverable in Chapters 8 through 11
Wrote or compressed a sentence that carries a finding	Full entry, with the clause-by-clause comparison Chapter 13 requires	Compression removes intervals and conditions first
Was one of several near-identical debugging turns	One grouped entry, preserving the final prompt, the correction, and the check	Grouping is permitted for low consequence and never for high

CONCEPT

Documentation Depth Follows Consequence

Ask what would go wrong if this exchange were wrong and nobody noticed. If the answer is a broken cell, one line is enough. If the answer is a budget, a customer treated differently, or a sentence a director will repeat, the answer is the full entry.

Grouping low-consequence turns is allowed and sensible. Grouping so broadly that the verification disappears is not, and it is the most common way a long record becomes an empty one.

Source: Course concept developed for this guide.

D.3 The Canonical Field Set

Table D.4 defines every field once. The Word form, the spreadsheet, and the Markdown version in the companion repository are generated from this table, so a field cannot be renamed in one format and not the others, and a record completed in one format can be read by someone who knows only another.

The right-hand column shows where each field comes from. Four of the eleven are Chapter 1's Table 1.6 questions, reproduced in the same words. Three more are the fields the note under Table 1.6 promises this appendix would add. The remaining four — the identifier, the stage, the prediction, and the change to the final work — are what the labs from Chapter 4 onward actually require, and they are the reason a record assembled from compact notes is auditable rather than merely complete.

Table D.4

The canonical field set, and where each field comes from

Field	What goes in it	Chapter 1 source
Exchange ID	A short stable label — E1, E2, E3 — written beside the cell, view, or paragraph the exchange produced, so the record and the work point at each other.	—
Stage	Specify, prepare, analyze, visualize, explain, document, or review. One word; it is what makes a long log searchable.	—
Tool and version	The product name and whatever version string the interface displays. Record the access mode too when it is not the obvious one.	Which AI tool did you use, and for which step?
Prompt as sent	The prompt in the words you actually sent, not a tidied reconstruction. Long prompts may be attached and referenced by location.	What prompt did you submit?
Output received	What came back, summarized in a sentence or two, with the location of the full response if you kept it.	— (promised by the Table 1.6 note)
Prediction recorded first	What you wrote down before the output was run or read: the row count, the total, the shape, the direction, the band. Blank is a finding.	—
Decision	Accepted, modified, rejected, or used only to generate questions. One of the four; not a paragraph.	What did you accept, revise, or reject?
Error or limitation found	Technical, analytical, factual, interpretive, or communication. Name the class and the instance.	— (promised by the Table 1.6 note)
Verification performed	The check you ran and the result: the reconciled total, the hand calculation, the rerun, the baseline, the source consulted, the reader tested.	How did you verify the final result?
Change to the final work	What is different in the submitted artifact because of this exchange. If nothing changed, say so and say why you kept it.	—
Analyst judgment, and what it taught you	The decision you made that the tool could not make, and the one thing the exchange changed about how you will prompt or check next time.	— (promised by the Table 1.6 note)

One field deserves a note. **Prediction recorded first** is blank on exactly two kinds of exchange: those that produced no output to predict, and those where the analyst forgot. The form does not distinguish them, so write “none required” in the first case. A log with many blank predictions is a finding about the analyst’s process, and it is better to see it than to fill it in afterward.

D.4 Level 1: The Compact Notebook Record

Table D.5 is the form. Copy it into a text cell directly beneath the cell, chart, or paragraph the exchange produced, and give it an exchange identifier that the full record can refer to later.

The labels are the canonical ones from Table D.4, not a second vocabulary, which is what lets a notebook full of these supply the exchange log directly rather than from memory. Chapter 1 poses them as four questions, and those questions are the gloss rather than the field names: *Which AI tool did you use, and for which step?* is **Tool and version**; *What prompt did you submit?* is **Prompt as sent**; *What did you accept, revise, or reject?* is **Decision**, whose permitted values are the four of Table D.4 — modified is the word for revised; and *How did you verify the final result?* is **Verification performed**. **Prediction recorded first** is the fifth line, required from Chapter 4 onward wherever a delegated step changes a count, a total, or a grain.

Table D.5

Level 1, the compact notebook record

Field	Your response
Exchange ID	[E1]
Tool and version	
Prompt as sent	
Prediction recorded first	
Decision	
Verification performed	

A completed one looks like this, from a Chapter 4 cleaning step on the practice file. It is six lines and it took under a minute to write.

Example D.1. A completed compact record, in a notebook text cell

```
Exchange ID: E3
Tool and version: general-purpose assistant, version not
  displayed; drafted the channel-standardizing cell below.
Prompt as sent: "Standardize the channel labels in
  transactions_raw to Online, App, Store. Print the value counts
  before and after and state the row count you expect. Change
  nothing else."
Prediction recorded first: 5 labels collapse to 3; row count
  unchanged at 14; revenue total unchanged at $767.40.
Decision: Modified.
Verification performed: the response mapped "web" to Online,
  which is right, and silently dropped two rows whose channel was
  missing, which I did not ask for; I removed that operation.
After the correction: labels 5 -> 3, row count 14 on both
  sides, revenue $767.40 on both sides.
```

Note what makes it useful, and note that all its field labels come from the full record. The prompt is quoted rather than described. The prediction names three quantities and the verification names the same three. The decision is one of the four permitted words and nothing else; the defect it responds to — a silent row drop, which Section 4.10 names as the first of the three cleaning failure modes — sits in the field built to carry it. Someone reading the notebook can reproduce the check in one cell.

D.5 Level 2: The Full Submission Record

The full record has eight parts, completed in the order given. The first three establish who did the work, with what, on what data. The fourth is the log. The fifth is the evidence. The sixth is what did not survive intact. The seventh is the signature, and the eighth applies only to group work.

The tables that follow define each part, field by field, rather than presenting a form you write into on the page. The fillable versions — in Word, spreadsheet, and Markdown, carrying these field names and nothing else — are in the companion repository, and Section D.8 says which to use.

D.5.1 Assignment and Analyst Information

This section is short and is completed first, because a record that cannot be attached to a specific version of a specific artifact is not evidence about anything.

Table D.6

Assignment and analyst information

Field	Response
Student name	

Field	Response
Group name and members, if applicable	
Course and section	INTG1-GC 2300, Section [SECTION]
Assignment or project	
Submission date	
Notebook, workbook, dashboard, or report filename	
Version of the submitted artifact	

D.5.2 AI Tools Used

One block per tool. If you used three assistants for three different purposes, that is three blocks, and the purposes matter more than the names: a tool used for adversarial review carries a different risk profile from the same tool used to draft prose that will carry a finding.

Table D.7

AI tools used, one block per tool

Field	What goes in it
Product name	Whatever the interface calls itself.
Model or version	The version string displayed, if any. Write “not displayed” rather than guessing.
Access mode	Public web tool, institutional account, embedded in Colab, embedded in another application, or local.
Dates used	First and last date, or the single date.
Purpose categories	One or more of the nine categories in Table D.8.

Table D.8 lists the nine purpose categories the guide uses, with the chapters in which each is used or assessed. Use these words rather than inventing your own, so that a term-long log can be read across assignments.

The categories are not the same thing as the **Stage** field of Table D.4, and the two are easy to conflate. Stage identifies where in the workflow the exchange occurred; purpose category identifies what kind of contribution the assistant made. One tool used across a whole assignment carries several purpose categories in this block, while each exchange in the log carries exactly one stage.

Table D.8

Purpose categories, and where each is used or assessed

Category	What it covers	Chapters
Problem framing	Turning a request into a specification; red-teaming a specification you wrote	1–2
Data preparation	Drafting a cleaning or joining step whose effect on counts and totals you predicted first	4
Code drafting	Writing a cell you specified, in a language you can read	1, 4–11
Debugging	Explaining an error message and proposing the smallest correction	4–11
Method explanation	Explaining what a method does, in terms you can check against the chapter	1, 3–11
Analytical interpretation	Proposing readings of an output you have already verified	5–11
Visualization design	Proposing chart specifications, encodings, layouts, or titles	12–13
Writing or editing	Drafting or compressing prose that carries a finding	13
Adversarial review	Arguing against your own specification, model, chart, or recommendation	2, 7–13

D.5.3 Data and Privacy Declaration

Section 1.13 states the rule the course actually enforces: confidential and personally identifiable data are never to be placed into an external AI tool. The declaration is where you record what you supplied and under what conditions. Chapter 3 adds the reason it belongs on the form rather than in a policy nobody rereads — pasting customer-level data into an external assistant is a custody decision rather than a convenience, and the granularity of what is collected and where it travels is itself a privacy design choice (Martin & Murphy, 2017).

Table D.9

Data and privacy declaration

Question	What a complete answer contains
What information was supplied to the tool?	Name the tables, columns, and row counts, or the schema alone. “The data” is not an answer.

Question	What a complete answer contains
What kind of information was it?	Synthetic, public, course-provided, de-identified, or confidential. The StyleCraft files are synthetic and are safe to paste.
Were any personal identifiers, credentials, proprietary records, or restricted data supplied?	Yes or no. If yes, name the approved environment that permitted it, or report the incident to your instructor.
What safeguards were used?	Schema only, safe sample, synthetic substitute, de-identification, aggregation, or an approved enterprise tool.

The section then closes with exactly one of the three statements in Table D.10, signed. They are mutually exclusive on purpose. A single confirmation qualified by *except as described above* would let an incident be described and then signed away in the same sentence, which is the opposite of what a declaration is for.

Table D.10

The three declarations. Sign exactly one

Statement	Sign it when
1. No restricted information was supplied.	Nothing confidential, personally identifiable, proprietary, or otherwise restricted was placed into an AI tool at any point in this work.
2. Restricted information was used only within the approved environment named above.	The environment, the permission under which it was used, and the information supplied are all recorded in the table above.
3. A possible incident occurred and has been reported.	Describe it above, report it to your instructor, and do not sign statement 1 or 2. Reporting an incident is not a penalty; concealing one is a different matter.

The StyleCraft files are synthetic, so for coursework the honest answer is almost always the first. The habit is what transfers — and so is the third, because an organization learns about an exposure from the person who caused it or from someone else, and the second route is always worse.

D.5.4 The Exchange Log

One entry per meaningful exchange, using the eleven fields of Table D.4. The shape follows the format rather than the other way round. In a spreadsheet, one exchange is one row and the eleven fields are columns, which is what makes a long log sortable. In Word, one exchange is a vertical two-column block, field beside entry; an eleven-column table on a portrait page is unreadable and worse to complete. In Markdown either shape works, and short logs read better as blocks. Section D.8 covers the formats, and the field names do not change between them.

Three rules govern the log. Preserve the prompt as sent, because a polished reconstruction hides how underspecified the original request was — which is the whole point of the deliberately loose prompts. Chapters 4 through 13 ask you to submit verbatim. Match the identifier in the log to the identifier in the notebook, workbook, or memo; a log that cannot be tied to the work it describes is a parallel document. And keep the decision field to one of the four permitted words, putting the reasoning in the fields built for it.

Several chapters run a two-prompt pattern: a first prompt that produces mechanics and ends by forbidding narration, then the analyst’s audit, then a second prompt that narrates within the audited results. File both, as adjacent rows sharing a stem — E7a and E7b — with the audit recorded between them in the verification table. Filing them as one row loses the audit, which is the part the pattern exists to protect.

D.5.5 Verification Evidence

This is the table a reader turns to first, and in most submissions it is where the grade is. It carries the checks that mattered, not every `assert` in the notebook — six well-chosen rows beat forty mechanical ones.

Table D.11

The verification evidence table

Field	What goes in it
Claim or output checked	The specific number, table, chart, or sentence under test — not the assignment.
Prediction, recorded first	What you expected, written before the check ran. A prediction written afterward verifies nothing.
Verification method	Reconciliation, hand calculation, rerun, baseline comparison, holdout, source check, reader test.
Expected value or condition	The figure or the condition that would count as a pass, stated in advance.
Actual result	What the check returned, in the same units.
Verdict	Pass, fail, or unresolved. Unresolved is a permitted verdict; a silent one is not.
Action taken	What you did about it. A failed check with no action is an unrepaired defect, not a completed record.
Evidence location	Where the evidence lives: cell, sheet, page, commit, file name.

CONCEPT

The Prediction Must Be Older Than the Result

The second column of Table D.11 is the one that makes the table worth keeping. A check whose expected value was written after the actual value was seen cannot fail, and a check that cannot fail is not a check.

This is Section 1.7's second habit, written down. Predict the number, the range, the shape, or the direction; run the step; compare; and when they disagree, record the disagreement rather than the resolution alone. Chapter 12 makes the same point from the other side: a fault you predicted and did not find is as informative as one you missed.

Source: Course concept developed for this guide, informed by the verification discipline of Chapman et al. (2000) and Provost and Fawcett (2013).

The evidence the guide's labs actually produce takes many forms, and the table is built to hold all of them. Table D.12 names the kinds, with the chapter that introduces each, so that "verification performed" does not collapse into "I looked at it."

Table D.12

Kinds of verification evidence this table is built to hold

Kind	What the row records	Introduced in
Row-count and total reconciliation	Counts and totals before and after a step, and the difference explained	Chapter 4
Hand calculation	A figure computed with a calculator on a miniature, beside the computed one	Chapters 3 and 5
Known-truth comparison	The answer the dataset was built to contain, recovered or not	Chapters 4, 6, and 11
Baseline comparison	The opponent's score on the same data, and the margin	Chapters 8 and 10
Holdout or backtest	A sealed evaluation opened once, after a declared selection rule	Chapters 8, 9, and 10
Confusion-matrix arithmetic	The four cells re-added by hand, summing to the evaluated population	Chapter 9
Threshold re-derivation	The cut recomputed from the stated costs, with its sensitivity	Chapter 9
Experiment assignment check	The flow from randomization to outcome, planned against realized	Chapter 11

Kind	What the row records	Introduced in
Chart audit	Grain, axis floor, encoding, filter state, and the reading each changed	Chapter 12
Reader test	A three-second or ninety-second test, with the reader's words recorded verbatim	Chapter 13
Clause-by-clause comparison	A compressed paragraph diffed against the original, each removal classified	Chapter 13

D.5.6 Rejected, Corrected, or Bounded Output

Most AI-assisted submissions in this course will carry at least one entry here, and several chapter exercises require one in their own words: at least one delegated step you corrected and why, at least one summary whose denominator you had to fix, or at least one case where the assistant's prediction and yours disagreed. Table D.13 is where they go, and its second field is what keeps the section honest about what actually happened.

Table D.13

Rejected, corrected, or bounded output

Field	What goes in it
Output, claim, limitation, or suggestion	Quote it. A paraphrase usually removes the thing that made it wrong, or the ambiguity that made it worth recording.
Disposition	Rejected, corrected, limited, or declined. One of the four.
Reason for disposition	The specific defect, ambiguity, or boundary, named.
Evidence used	What you checked it against.
What appears in the final work	The replacement, the correction, the stated limit, or the fact that nothing from this exchange was used.

The four dispositions are why this section is not called *errors*. **Rejected** means the output never entered the work; **corrected**, that it entered after a specific repair; **limited**, that it entered with a stated boundary on what it may be used for; **declined**, that you considered a suggestion and did not take it. So the purpose is not to manufacture an error, and a fabricated one is worse than none. If nothing material went wrong, record the limitation you worked around or the ambiguity the assistant resolved without saying it was choosing — Chapter 3's exercise puts that question well: which definition did it choose, and did it state that it was choosing?

Where an exercise does not require an entry and nothing was rejected, corrected, limited, or declined, the section is not left blank. Write *none identified* and name the checks that would have detected a material defect. That converts an empty section into a claim someone can evaluate: "I

found nothing” is a finding when the checks behind it are listed and an absence otherwise. Chapters 11, 12, and 13 make the same point from the other direction — an exchange in which the assistant performed badly and you caught everything is a better submission than one in which it performed well — and none of them asks you to go looking for something to reject.

D.5.7 The Analyst-of-Record Statement

The record ends with the statement below, adapted to your own work. It covers five things: what the tool contributed, what you contributed, how the final work was verified, what limitations remain, and that you can explain what you submitted. The bracketed clause is the one that changes every time; the rest is a fixed form so that a reader can find the limitations in the same place in every submission. Then your name and the date.

Statement D.1. The analyst-of-record statement

I remain responsible for the submitted analysis. I reviewed the AI-assisted output identified above, made the decisions recorded in this record, and verified the final work using the evidence listed. I can explain the code, the calculations, the visualizations, and the recommendations in this submission. The remaining limitations are [LIMITATIONS].

Name: [NAME]

Date: [DATE]

CONCEPT

Sign What You Can Explain

Chapter 1 set the standard in one line: do not send work to a manager that you cannot explain. The statement above is where that becomes a signature, and the clause that carries the weight is the third one.

Before signing, run Section 1.9’s four questions on the finished artifact. Does the output answer the question I specified? Does one row still represent what I said it represents? Can I reproduce at least one number by hand? Can I state in one sentence what this result does not show? If any answer is no, the work is not ready and the statement is not yet true.

Source: Course concept developed for this guide.

D.5.8 Group Contribution Record

Group projects submit one consolidated record rather than several disconnected forms, with the contribution table below appended. The point is traceability, not accounting: a reader should be able to take any exchange in the log and find the member who ran it and the member who checked it.

Table D.14

Group contribution record, one row per member

Field	What goes in it
Member	Name.
AI-assisted task performed	Which exchanges, by ID.
Verification performed	Which checks in the verification table are this member's.
Final artifact contribution	The cells, views, sections, or paragraphs that carry this member's work.
Peer confirmation	The initials of a second member who has read the row and agrees with it.

Two group-specific rules follow from the rest of this appendix. Exchange identifiers are unique across the group, so E7 means one thing in the consolidated log. And the member who ran an exchange should not be the member who signs off on its verification wherever the work allows otherwise — Chapter 1's second-reviewer practice, applied to a team that already has a second reader available.

D.6 A Weak Record and a Strong One

The difference between a record that helps and a record that occupies space is almost always specificity. Table D.15 sets one against the other, on the same exchange: a Chapter 5 summary table drafted by an assistant on the fourteen-line miniature Lab 5.1 certifies at eight orders and \$660.00 of revenue.

Table D.15

The same exchange, recorded weakly and recorded well

Field	Weak	Strong
Tool and version	ChatGPT	General-purpose assistant; version not displayed; public web tool. Drafted the channel-summary table in cell 12
Prompt as sent	Asked it to summarize the data	"Summarize this sales data and tell me what it means" — the deliberately vague prompt of Exercise 5.6, sent verbatim
Prediction recorded first	—	Revenue subtotals sum to \$660.00; eight orders; average order value near \$82.50, not the average of the group means
Decision	Used it	Modified

Field	Weak	Strong
Error or limitation found	Some small issues	Analytical: reported an overall average order value of \$105.00 by averaging the two group means, and stated no denominator for the channel shares
Verification performed	Checked it	Recomputed the pooled figure by hand: $\$660.00 \div 8 = \82.50 . The weighted recombination returns \$82.50; the averaged average returns \$105.00
Change to the final work	—	Replaced the headline with the pooled \$82.50, added the denominator to every share, and struck the causal verbs from its narration
Analyst judgment, and what it taught you	—	Pooled, because the question is about the chain rather than the average branch. Next time the prompt states the denominator before the assistant picks one

The weak record is not shorter by accident. Every one of its entries is a place where a specific answer existed and a general one was written instead: “some small issues” had a name, “checked it” had an arithmetic, and “used it” concealed the fact that the headline number was wrong and was fixed. Notice also what the strong record costs — about a hundred and twenty words on a two-minute exchange, every figure of which was already on the screen when it was written. That is the practical argument for writing the compact record during the work. At that moment it is nearly free; an hour later it is archaeology.

D.7 Retention and Evidence Links

The record points at evidence; it does not contain all of it. A prompt that runs to two pages is attached and referenced by location, and a reconciliation is cited by the cell that printed it rather than transcribed. Table D.16 gives the link types the course accepts and what each is good for.

Table D.16

Evidence links, and what each is good for

Link type	Good for, and what to watch
Notebook cell	Printed counts, totals, assertions, and charts. Cell numbers move, so cite the exchange ID you wrote beside the cell rather than its position
Saved conversation export	A long prompt, or a response you needed in full. Export where the tool and the course permit it, and never export restricted data
Workbook or Tableau worksheet	Reconciliation tables and view-level provenance. Name the sheet; “the workbook” is not a location

Link type	Good for, and what to watch
Version-control commit	Showing what changed and when. Cite the short hash and the file, not the branch
Report or memo page	A claim, a caveat, or a reversal condition. Page numbers move between drafts, so cite the version too
Screenshot	An interface state that leaves no other trace, such as a filter. A screenshot is evidence of an appearance, not of a number

One rule overrides the table. Evidence must travel with the submission or remain reachable through the course-approved environment. A link to a conversation on a personal account, a file in a drive nobody else can open, or a page that may not exist next term is not evidence; it is a promise. Where a tool does not permit export, summarize the response in the log and say that no export was available.

Retention has two halves that pull against each other. Keep the record and its evidence available for as long as the work can be questioned — through the course’s grading and appeal period here, and for whatever period the organization requires in professional work. Then delete what you no longer need, particularly exported conversations and any sample of data that was a convenience rather than a requirement. A record is evidence for a period; an unnecessary export of restricted data is a liability indefinitely.

D.8 Completing the Record in Word, a Spreadsheet, or Markdown

The same record can be completed in any of three formats, and the choice is a question of what the submission is rather than of preference. What does not change is the field set: the same eleven exchange fields, the same eight verification fields, the same declaration, the same statement. The companion repository carries all three, generated from one source.

The three files share one stem, `Appendix_D_AI_Use_Record_FORM`, and differ only in extension. Use the `.docx` when the submission is a narrative assignment, a memo appendix, or a project deliverable. Use the `.md` when the record lives inside a Colab notebook or a repository: it survives export to `.ipynb`, renders in the notebook itself, and diffs cleanly under version control. Use the `.xlsx` — or the CSV files corresponding to its worksheets — for an assignment with more than about a dozen exchanges, where one row per exchange sorts, filters, and counts in a way a Word table does not. Group work with simultaneous editors can open the Word and spreadsheet files in Google Drive without changing a field name.

Whichever format you use, the submitted record should be a file rather than a screenshot of a file, for the same reason Section A.15 asks for notebooks rather than pictures of notebooks. A record that cannot be searched cannot be audited.

D.9 What Each Chapter Asks You to File

Every chapter of the guide requires this record, and each one narrows it to the exchange its own work makes most consequential. Table D.17 states what each chapter asks for and which level satisfies it. Where a chapter asks for something the general form does not name — a count of UNSPECIFIED rows, a graded prediction including false alarms, a clause-by-clause diff — it belongs in the verification table, as a row whose method is that chapter’s audit.

Table D.17

What each chapter asks you to file

Ch.	What the chapter asks to be filed	Level
1	Section 1.7 and the Review Before You Rely box; Exercise 1.6 asks for the compact record in the notebook and the full record for submission	1 and 2
2	Section 2.8’s adversarial review of your own specification; Exercise 2.7 asks for the full record	2
3	Section 3.9’s drafted dictionary entries and the audit of them; Exercise 3.7 asks for the full record	2
4	Every delegated cleaning step, with the assistant’s predicted effect on counts beside yours; Exercise 4.5 requires at least one recorded disagreement	1 in the notebook, 2 for submission
5	Every drafted summary, with the denominator, the center, and the reconciliation audited; Exercise 5.5 requires at least one corrected summary	1 in the notebook, 2 for submission
6	The feature build and the cluster audit; Exercise 6.6 requires the full exchange and which audit step changed your assessment most	2
7	Both prompts of the two-prompt pattern, with the verbs graded against Section 7.2’s registry; Exercise 7.6 requires the full exchange	2
8	Both exchanges, plus the five-point audit written out in order; Exercise 8.6 requires every performance claim marked earned or not earned	2
9	Both exchanges, plus the classification supplement; Exercise 9.6 requires the loose-prompt response and the re-prompted one, compared	2
10	Both exchanges, plus the leaderboard band you predicted before reading the response; Exercise 10.6 requires the comparison of the two prompts	2
11	Four moves — registration, analysis, adversarial review, drafting — each filed; Exercise 11.6 requires the UNSPECIFIED count and the marked-up recommendation	2

Ch.	What the chapter asks to be filed	Level
12	Both exchanges of the five-step routine, plus the predicted faults, the graded prediction including false alarms, and a note per repair	2
13	The exchanges — title generation and prose compression are separate tasks — with the candidates' verification lines and the clause-by-clause diff; the record travels inside the memo's appendix	2

Two patterns run through the table. The first is that the unit of filing grows with the work. Chapters 7 through 10 run the two-prompt pattern — mechanics, the analyst's audit, then narration — and ask for both prompts rather than the final one, because the audit between them is what is being graded. Chapter 11 expands that into four moves: registration, analysis, adversarial review, and drafting, each filed. Chapters 12 and 13 run a five-step visual and communication routine in which three of the five steps are analyst work rather than exchanges, and the prediction is graded including its false alarms. The second pattern is that from Chapter 8 onward the record stops being an attachment and becomes part of the deliverable: the audit memo is the lab's final deliverable in Chapters 8 through 11, and in Chapter 13 the record is filed inside the recommendation memo's appendix. That progression is this appendix's real argument. In Chapter 1 the record is a form you complete because the course asks; by Chapter 13 it is what lets a reader see not only what you concluded, but how you came to be entitled to conclude it.

References

Every work below is cited above. Chapter 1 remains the guide’s own statement of the four habits, the compact record, and the privacy rule.

Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). *CRISP-DM 1.0: Step-by-step data mining guide*. SPSS Inc.

Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>

Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. *Journal of the Academy of Marketing Science*, 45(2), 135–155. <https://doi.org/10.1007/s11747-016-0495-4>

McAdoo, T. (2023, April 7). *How to cite ChatGPT*. APA Style. <https://apastyle.apa.org/blog/how-to-cite-chatgpt>

National Institute of Standards and Technology. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.600-1>

Provost, F., & Fawcett, T. (2013). *Data science for business: What you need to know about data mining and data-analytic thinking*. O’Reilly Media.