

Mining the Audience:

AI-Assisted Analysis of Traveller Reviews in the Tourism and Hospitality Industry

Research Project in Marketing | KU Eichstätt-Ingolstadt | Joachim Büschken

1. Course overview

This course is a project-based market research seminar. Student teams investigate a self-chosen management question in the tourism and hospitality industry by analysing large corpora of consumer-generated reviews — from platforms such as TripAdvisor, Booking.com, Google Reviews, Airbnb, and Yelp — using AI-assisted text analysis. No programming is required. The analytical pipeline runs entirely through the upload interfaces of Claude, ChatGPT, or Gemini.

The tourism and hospitality industry is one of the richest sources of unsolicited, aspect-level consumer feedback in the world. Travellers describe what met and failed their expectations in detailed, structured language — rating location, cleanliness, service, value, food, atmosphere, and staff consistently across platforms and property types. This corpus is ideal for management-relevant text analysis: the aspects are well-defined, the management decisions are direct, and the data is abundant.

The course has a single output: a research project that takes a real management question, builds a corpus of traveller reviews relevant to it, applies a systematic AI-assisted coding scheme to identify aspect-level sentiment patterns, and interprets the findings in terms of actionable decisions for a named stakeholder — a hotel general manager, an Airbnb host, a tourism board, an airline, or a restaurant group. The project is developed in five stages across fourteen weeks, with three assessed milestones and a final written report.

2. Learning outcomes

On successful completion of this course, students will be able to:

- **Formulate a management-relevant research question** in the tourism and hospitality context that is answerable through aspect-level text analysis of consumer reviews.
- **Design and justify a corpus** of traveller reviews — specifying platform, property type or destination, time window, and review volume — appropriate to their research question.
- **Develop and apply an operationalised coding scheme** covering the key aspects of traveller experience (location, cleanliness, service, value, food, atmosphere, staff) and document decision rules for ambiguous cases.
- **Conduct AI-assisted aspect-based sentiment analysis** at scale using LLM upload interfaces, with validation and a cross-tool reproducibility check.
- **Interpret coded review patterns** in terms of concrete management recommendations for a named decision-maker in the industry.
- **Report research findings** in a professionally formatted written report that distinguishes corpus description, methodology, results, and management implications clearly.

3. Industry context and management relevance

Tourism and hospitality is a review-driven industry. A one-star difference in TripAdvisor rating has been estimated to shift RevPAR by several percentage points. A pattern of complaints about cleanliness across a hotel portfolio signals a capital investment priority. A systematic gap

between what a destination's marketing promises and what travellers report experiencing generates negative word-of-mouth that no promotional budget can outrun.

The management questions accessible through traveller review analysis are correspondingly direct:

- Which aspects of the guest experience drive overall satisfaction — and which are most associated with complaints?
- How does a property's review profile compare to its competitive set on specific dimensions such as value, service responsiveness, or food quality?
- Do business travellers and leisure travellers evaluate the same hotel aspects differently — and what does this imply for segmented service design?
- How has a destination's perceived safety, accessibility, or environmental quality shifted in traveller discourse over time?
- What aspects of the flight experience generate the strongest negative sentiment, and how does this differ across economy, business, and premium economy reviews?
- Which aspects of Airbnb host communication and property condition most consistently distinguish Superhosts from standard hosts in review language?

These are the kinds of questions this course equips students to answer with systematic, documented, and validated evidence from review text — rather than from intuition or aggregate star ratings.

4. Data sources

The tourism and hospitality sector has exceptionally well-documented review corpora. The following sources are available without programming:

- **Kaggle pre-scraped datasets:** TripAdvisor hotel reviews, Booking.com datasets, Yelp restaurant reviews, Airbnb review corpora, and airline passenger experience datasets are all available on Kaggle as CSV files for direct download. These are the recommended starting point. Review volumes typically range from 10,000 to 500,000+ reviews per dataset.
- **Hugging Face Datasets:** Several tourism-relevant NLP benchmark datasets are available, including hotel and restaurant review corpora with aspect annotations.
- **Google Maps / TripAdvisor manual export:** For focused corpora around specific properties or destinations, review pages can be accessed directly and copy-pasted into a structured spreadsheet for small-scale analysis (up to approximately 300 reviews).
- **Airbnb and Booking.com review text:** Publicly visible on property pages; manually collectable for focused single-property studies.

The no-code constraint means students will not be writing web scrapers. All corpora must be assembled from downloadable datasets or through manual collection for small-scale pilot work. The corpus scoping decision — which platform, which property type, which geographic scope, which time window — is one of the key methodological choices students will make and justify.

5. Typical review aspects in tourism and hospitality

Unlike entertainment reviews — where aspects such as plot, pacing, and acting vary by genre and are not always present — traveller reviews have a relatively stable aspect structure that appears across platforms and property types. The core aspects students will encounter are:

Aspect	What reviewers discuss	Aspect	What reviewers discuss
--------	------------------------	--------	------------------------

Location / accessibility	Proximity to attractions, transport links, parking, neighbourhood safety	Value for money	Whether price paid matches experience received; hidden charges
Cleanliness	Room, bathroom, common areas, bedding, maintenance standards	Food and beverage	Breakfast quality, restaurant, bar, menu variety, dietary options
Service and staff	Friendliness, responsiveness, check-in/out, problem resolution	Atmosphere / ambience	Design, décor, noise levels, general feel of the property
Room quality / comfort	Bed comfort, size, furnishing, temperature, noise, view	Facilities and amenities	Pool, gym, spa, Wi-Fi, parking, business centre

Specific sub-sectors introduce additional aspects — for airline reviews: seat comfort, legroom, in-flight entertainment, meal service, and punctuality. For Airbnb: host communication, accuracy of listing description, and check-in process. Students will adapt the core framework to their chosen sub-sector through the inductive discovery phase in Stage 2.

6. Assessment

Component	Weight	Submission	Description
Stage 1 Research Proposal	15%	Week 4	Research question, corpus scope, management relevance, analytical approach, and aspects. Two pages maximum. AI tutor transcript appended as Exhibit A.
Stage 2 Progress Presentation	15%	Week 8	20-minute group presentation covering finalised coding scheme, pilot results, aspect frequency distribution, validation agreement rate, and at least one preliminary finding.
Stage 3 Final Presentation	25%	Week 13	30-minute group presentation of complete analysis, findings, and management recommendations. Includes Q&A.
Stage 4 Written Report	45%	Week 14	Full research report: introduction, literature review, methodology, results, management implications, limitations. 4,000–6,000 words excluding appendices.

Group composition: Groups of 2–3 students. All group members receive the same grade for presentations and the written report. Individual contributions must be documented in the report's author note.

7. Week-by-week schedule

Week	Topic	Content	Deliverable
1	Industry overview: Tourism and hospitality	The structure of the tourism and hospitality industry: hotels, OTAs, airlines, destinations, vacation rentals. The role of review platforms (TripAdvisor, Booking.com, Google, Airbnb, Yelp) in traveller decision-making and property management. Why traveller reviews are an exceptionally rich corpus for management research. Key industry trends: sharing economy, sustainability, experience economy, digital nomadism.	
2	Management research questions and text as data	What makes a good management research question in tourism and hospitality. Five question types: descriptive, exploratory, explanatory, predictive, comparative. Text as data — what traveller reviews contain (experience description, aspect evaluation, overall judgment, comparison with expectations). Challenges: sampling bias, platform effects, extreme-response bias, non-English corpora. Introduction to the Stage 1 AI tutor prompt. Critical principle: a research question must be developed in dialogue with the data — groups are required to conduct a preliminary scan of 100–200 reviews from their intended corpus before finalising their RQ. The RQ is confirmed only after this scan, not before it.	
3	Data acquisition and corpus scoping	Data sources for traveller review analysis: Kaggle datasets (TripAdvisor, Booking, Yelp, Airbnb), Hugging Face, manual collection. What a review dataset looks like — column structure, required metadata (ReviewID, PropertyName, ReviewText, Rating, Date, Platform). Corpus scoping decisions: platform, property type or destination, geographic scope, time window, review volume. Data quality checks: duplicates, missing text, language filtering, rating distribution, review-bombing.	
4	Stage 1 presentations and feedback	Each group presents their research proposal (5 minutes) and receives structured peer and instructor feedback. Focus on: is the management question answerable through review text? Is the corpus realistically accessible? Is the	Stage 1 Research Proposal (15%)

		management relevance specific to a named decision-maker?	
5	Model-based text analysis — approaches and tools	Five analytical approaches executable through LLM upload interfaces: (A) sentiment analysis — overall and aspect-level; (B) thematic/topic coding; (C) aspect extraction and coding scheme design; (D) comparative analysis across groups; (E) pattern synthesis and management interpretation. Prompt design for tourism and hospitality corpora — handling multi-aspect reviews, differentiating experience description from evaluation. The inductive discovery sequence: pilot scan → coding scheme → full corpus coding. Quality control: validation, stability check, cross-tool reproducibility.	
6	Coding scheme design workshop	Supervised in-session work on coding scheme development. Each group uploads a pilot sample (100–150 reviews) to their chosen LLM tool and conducts an inductive thematic scan. Groups identify which aspects appear in their corpus, draft definitions and decision rules, and present their provisional scheme for feedback. Particular attention to boundary cases: does a comment about Wi-Fi code as 'facilities' or 'value'? Does 'the view was stunning' code as location or atmosphere?	
7	Pilot coding and validation	Groups apply their finalised coding scheme to a pilot batch of 50–100 reviews. Two group members independently validate 20–30 of the coded outputs and calculate inter-rater agreement. Groups run a stability check (same batch submitted twice) and document disagreements. Common failure modes reviewed in plenary: model classifying experience descriptions as evaluations, conflation of 'service' and 'cleanliness', handling of ironic or highly qualified reviews.	
8	Stage 2 progress presentations	Each group presents their coding scheme, pilot results (aspect frequency distribution, sentiment breakdown), validation agreement rate, and at least one preliminary finding. Discussion of what the patterns reveal and what they do not — distinction between	Stage 2 Progress Presentation (15%)

		descriptive frequency and management relevance.	
9	Full corpus coding — execution and troubleshooting	Groups scale up to full corpus coding using their validated scheme. Batching strategy: 50–100 reviews per upload, consistent prompt across all batches, output saved and labelled per batch. Common troubleshooting issues: model drift across batches, inconsistent output format, handling reviews in multiple languages. Cross-tool reproducibility check on a 100-review subsample (Claude vs. Gemini or ChatGPT).	
10	Comparative analysis and pattern interpretation	Aggregating coded outputs: constructing aspect × sentiment frequency tables, calculating the proportion of reviews mentioning each aspect, visualising distributions. Comparative analysis across groups defined by the research question: property type (budget vs. luxury), traveller type (business vs. leisure), platform (Booking.com vs. TripAdvisor), time period (pre-pandemic vs. post-pandemic), geography (domestic vs. international guests). Interpreting patterns in terms of management decisions. Mandatory exploratory sub-task: each group must go beyond the pre-specified RQ and ask what else the data reveals. Required prompts for this task: (1) What is the most surprising pattern in your data that your RQ does not explain? (2) Which aspects appear at unexpectedly high or low rates, and why might that be? (3) Are there subgroups in your corpus — by rating, time period, or platform — that behave differently from the overall pattern? (4) What would a follow-up study ask, and could your current dataset answer it? AI interaction for this sub-task should use open-ended exploratory prompts, not the pre-specified coding scheme.	
11	Management implications and report structure	From patterns to recommendations: what makes a management implication actionable? The decision-maker test — can the finding inform a specific decision by a named stakeholder? Literature on traveller satisfaction, expectation disconfirmation, online reputation management, and service quality in	

		hospitality (SERVQUAL, importance-performance analysis). Report structure: introduction, literature, methodology, results, implications, limitations.	
12	Writing workshop — methodology and results sections	Supervised writing workshop. Groups draft their methodology section (corpus description, coding scheme, validation procedure) and their results section (aspect frequency tables, comparative findings). Peer review of drafts with structured feedback protocol. Common errors: over-reliance on LLM-generated synthesis, underdocumented methodological choices, findings stated without reference to the coded data.	
13	Stage 3 final presentations	30-minute group presentations of complete analysis, findings, and management recommendations. Each presentation is followed by 10 minutes of Q&A. Assessors evaluate: research question clarity, methodological rigour, quality of analysis, specificity and feasibility of management recommendations.	Stage 3 Final Presentation (25%)
14	Report submission and course wrap-up	Submission of written reports. Course retrospective: what worked, what didn't, what would change. Discussion of how AI-assisted text analysis sits within the broader toolkit of market research — when it is the right tool, and when it is not. Reflections on the tourism and hospitality industry as a research context.	Stage 4 Written Report (45%)

8. The Stage 1 AI tutor prompt

A Socratic tutor prompt is provided on ILIAS to guide each group through developing their research question. The prompt runs entirely through the ChatGPT, Claude, or Gemini upload interface — no installation required. It works as follows:

1. **Upload the prompt file** to a new AI session before typing anything else. The file is available on ILIAS under 'Stage 1 Prompt'.
2. The tutor will open with a question about your personal connection to the tourism and hospitality industry — a hotel stay that surprised you, a destination that did not meet expectations, a platform experience that puzzled you. Start there, not with a theory.
3. Work through corpus scoping (which platform, which property type, which geographic scope), management angle (who would use the findings, for which decision), and question type (descriptive, exploratory, explanatory, predictive, comparative) in conversation.
4. **Type COMPLETE** only when the tutor has confirmed that your research question meets all five quality criteria (corpus identified, management relevance, answerable through review text, non-trivial, type-complete). The tutor will then generate a structured closing output.

5. Save the full conversation transcript and the closing output. Both must be submitted as Exhibit A alongside your Stage 1 proposal.

The five quality criteria your research question must satisfy before you type COMPLETE:

- **Corpus identified** — platform, property type or destination, and scope are concrete and feasible (1,000–5,000 reviews).
- **Management relevance** — a real decision is addressed; a specific, plausible user of the findings can be named (e.g. 'the revenue manager of a mid-scale hotel chain in Bavaria').
- **Answerable through review text** — the question can be investigated by identifying aspect-level patterns in traveller reviews, without requiring survey data, booking figures, or operational records.
- **Non-trivial** — at least two plausible outcomes exist; the answer is not already known.
- **Type-complete** — all required elements of the chosen question type are present and clearly named.
- **Data scan completed** — a sample of at least 100–200 reviews from the intended corpus has been read inductively, and the RQ has been confirmed as feasible and well-matched to what those reviews actually discuss. A RQ that outpaces the data will not receive sign-off.

9. AI use policy

AI tools (Claude, ChatGPT, Gemini) are authorised for use in this project under §12 Abs. 7 APO and the KU Handreichung KI in der Hochschullehre (September 2025). The full AI Use Policy Addendum is available as a separate document on ILIAS and forms part of this syllabus. The core principle is stated here:

AI tools are research instruments, not substitutes for intellectual work. The research question, the analytical design, the validation of outputs, the interpretation of findings, and the written conclusions are the students' own work. AI tools assist with the labour-intensive tasks of text classification and pattern generation — but the thinking that surrounds those tasks remains with the researcher.

Uses explicitly encouraged include: the Stage 1 Socratic tutor conversation, inductive thematic scanning of a sample, aspect-level sentiment coding of the full corpus, cross-tool reproducibility checks, stress-testing the coding scheme, and interpreting difficult or ambiguous reviews. Uses not permitted include: generating the research question without the tutor prompt, using a single AI prompt to produce findings from the full corpus, and using AI to write the findings, discussion, or conclusions sections of the report.

10. Written report requirements

The final report should be 4,000–6,000 words excluding references and appendices. It must follow this structure:

6. **Introduction (400–600 words)** — industry context, management problem, research question, brief overview of the study.
7. **Literature review (600–900 words)** — relevant prior research on traveller satisfaction, online reviews in tourism and hospitality, and the management decision addressed. Minimum 8 peer-reviewed sources.
8. **Methodology (700–1,000 words)** — corpus description (platform, property type, geographic scope, time window, review volume, data quality procedures), coding scheme (all aspects with definitions and decision rules), analytical procedure (batching strategy, prompt documentation, validation, stability check, cross-tool reproducibility), and limitations.

9. **Results (900–1,400 words)** — aspect frequency distribution, sentiment breakdown, comparative findings as appropriate to the research question. All claims must reference the coded data explicitly.
10. **Management implications (500–800 words)** — three to five specific recommendations for the named decision-maker, directly grounded in the findings. Each recommendation must be actionable — specifying what the decision-maker should do differently and why.
11. **Limitations and conclusion (300–500 words)** — honest account of what the corpus, coding scheme, and analytical approach cannot tell us, and directions for future research.

Mandatory appendices: Exhibit A — Stage 1 closing output and conversation transcript. Exhibit B — full coding scheme with definitions and decision rules. Exhibit C — all prompts used verbatim. Exhibit D — validation results (sample size, agreement rate, documentation of disagreements). Eigenständigkeitserklärung with AI use declaration.

11. Illustrative research questions

The following examples illustrate the range and type of research questions suitable for this course. They are provided for orientation only — your group must develop its own question through the Stage 1 tutor conversation.

- **Comparative / hotels:** How do guest evaluations of service and cleanliness differ between budget and upscale hotels on Booking.com, and what does this suggest for service investment priorities at different price points?
- **Explanatory / satisfaction:** Which aspects of the guest experience in urban boutique hotels — location, atmosphere, staff, or value — are most strongly associated with overall positive sentiment in TripAdvisor reviews?
- **Comparative / traveller type:** Do business travellers and leisure travellers evaluate the same hotel aspects differently in their reviews, and which aspects show the largest sentiment gap between the two groups?
- **Descriptive / destination:** Which aspects of the visitor experience in Munich dominate TripAdvisor restaurant reviews, and how does the sentiment distribution across food quality, service, value, and atmosphere compare to that of Vienna?
- **Exploratory / vacation rental:** What aspects of the Airbnb guest experience most frequently appear in reviews that explicitly mention 'expectations' or 'as described', and what does this reveal about the sources of expectation misalignment?
- **Temporal / recovery:** How did the sentiment profile of airline passenger reviews on Yelp shift between 2019 and 2023, and which aspects show the most pronounced changes in the post-pandemic recovery period?
- **Comparative / platform:** Do reviews of the same hotel chain on TripAdvisor and Booking.com differ systematically in which aspects are discussed and how positively they are evaluated?

For questions about the course structure, the Stage 1 prompt, or data access, contact the course instructor before the Week 3 session. The AI Use Policy Addendum (separate document on ILIAS) provides detailed guidance on all permitted and non-permitted uses of AI tools in this project.