

Generating Legal Commentaries from Case Databases via Retrieval, Clustering, and Generation^{*}

Max Prior¹, Niklas Wais¹ and Matthias Grabmair¹

¹Technical University of Munich, Boltzmannstraße 3, 85748 Garching near Munich

Abstract

We present a fully automated pipeline that transforms large collections of court decisions into legal commentaries for statutes — without providing any handcrafted doctrinal framework. Using 4.555 decisions of the German Federal Court of Justice that cite §§ 242, 280, 812 and 823 of the German Civil Code (BGB), we extract paragraph-level chunks, summarize their reasoning, and derive keywords, which are embedded and clustered. For each cluster, an LLM generates headings and synthesizes citation-rich sections, which are then merged into coherent commentaries by four state-of-the-art LLMs. We evaluate along five dimensions – topical relevance, heading-match, citation faithfulness, cluster distinction and logical ordering – using both a human expert and an “LLM-judge”. Our results show that commentary-like argument mining from court decisions to generate reports that can be refreshed within minutes at minimal cost is feasible, yet they highlight limitations arising from restricted sources and the normativity of legal reasoning.

Keywords

Argument Mining, Legal Commentaries, Generative NLP

1. Introduction

Common and Civil Law systems are often contrasted by the claim that the former relies on precedent, while the latter centers on statutes. While Civil Law judges are not formally bound by case law and begin with statutes, the notion of judges as mere *bouche de la loi* (“mouth of the law”; for a critical assessment of *Montesquieu*’s actual position, see [1]) has long been proven to be a deliberate misrepresentation invented by the self-proclaimed anti-formalist movements of the early 20th century [2]. Judicial decision-making necessarily involves interpretation, and prior rulings — although generally not binding — inform the consistent application of statutes. Legal commentaries systematize this practice: they collect and analyze the application of statutes by the courts as well as proposals regarding their application from papers, books, and other commentaries. In other words, **they mine arguments** from court decisions and the legal literature.

1.1. Structure of Legal Commentaries

Commentaries typically focus on one statute book, e.g., the German Civil Code (BGB), and are following its structure. In the case of a BGB commentary, there is one “chapter” per legal provision of the BGB (§ 1 to § 2385). After repeating the wording of the respective legal provision, each chapter is structured around questions arising from the application (i.e., interpretation) of the provision. For example, when faced with a case potentially involving liability for injury under § 823 (1) BGB, a judge might question whether witnessing a traumatic event can constitute a violation of “Gesundheit” (health) as intended by the provision. Since the statutory wording itself leaves room for interpretation, the judge will seek arguments supporting or rejecting the view that psychological trauma from witnessing an event should qualify as a health violation within the meaning of § 823 (1) BGB. This involves identifying definitions proposed by legal scholars or established by court rulings concerning the concept

Proceedings of the First Argument Mining and Empirical Legal Research Workshop (AMELR 2025), June 20, 2025, Chicago, United States

✉ max.prior@tum.de (M. Prior); niklas.wais@tum.de (N. Wais); matthias.grabmair@tum.de (M. Grabmair)

🌐 <https://www.cs.cit.tum.de/lt/tum-legal-tech-working-group/> (M. Grabmair)

🆔 0009-0005-7066-996X (M. Prior); 0009-0008-6459-721X (N. Wais); 0000-0001-6586-2486 (M. Grabmair)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

of "Gesundheit" (health) under § 823 (1) BGB. It also involves reviewing prior court decisions addressing similar circumstances. Regarding the question, both pro and contra arguments are collected and systematized in the § 823 BGB chapter of a commentary.

1.2. Computer-Generated Commentaries

Legal Commentaries are massive books written by multiple authors. Their creation is a complicated and costly process. Given the recent advances in Natural Language Processing brought about by LLMs, researchers have started to investigate the automated generation of such commentaries based on collections of court rulings. *Engel and Kruse* present a prototype of an automated pipeline that lets GPT-4 draft a commentary on Article 8 of the German Constitution. Python code fetches 125 rulings of the German Constitutional Court, extracts passages mentioning Article 8, and feeds them – with an elaborate prompt that entails constitutional law doctrine – to GPT-4, which classifies, summarizes and groups every citation under typical headings. Compared to five leading hand-written commentaries, the machine-generated version cites far more decisions, provides pinpoint references and can refresh in hours at a cost of approximately \$33, yet it still misclusters some material, and misses doctrinal shifts. It was created with heavily engineered prompts that provide the doctrinal framework of German constitutional law to the model, which structures the generation into a form comparable to human-written commentaries. The authors present it as a complement, not a substitute [3]. They also use a similar approach with rulings from the European Court of Human Rights (ECtHR) [4]. *Santosh et al.* developed the "LexGenie" model, which is not explicitly positioned as a legal commentary, but an interactive tool to generate structured reports from ECtHR rulings based on keywords provided by the user. Offline, they index at the paragraph level via Mistral-7B key-phrase embeddings and store them in a FAISS database. Online, they retrieve passages via Maximum Marginal Relevance (MMR) to the keywords and organize them using BERTopic and HDBSCAN. Afterwards, GPT-4o-mini refines headings and incrementally drafts citation-rich texts to build multi-case legal guides. Limitations include occasional misclustering and sparse cross-section links [5]. Similar functionality is now being explored by systems like OpenAI's Deep Research, which provide grounded and multi-document syntheses to support advanced academic workflows [6]. Like [5], they do not rely on hard-coded legal knowledge when generating legal reports and do not follow the structure of a legal commentary. A key difference compared to [5] is that they first generate a structured plan and then retrieve documents, rather than generating a structure from retrieved documents.

1.3. Our Contribution

We combine the approaches of [5] and [3] by generating commentary-like reports for legal provisions (not keywords like [5]) without providing doctrinal context (unlike [3]). Instead, we rely solely on the texts of court rulings and statutes. This generalization makes our approach easily extendable to any number of legal provisions. For demonstration purposes, we focus on the *German Civil Code* (BGB), which codifies central aspects of private law and consists of more than two thousand provisions. When compared to the German Constitution or the ECHR, it is of high practical importance and particularly challenging due to the amount of court rulings that refer to it. We also compare the performance of multiple state-of-the-art models for the task at hand, introduce a sophisticated pipeline and provide quantitative as well as qualitative evaluation. Lastly, we contribute to the legal theory discussion on the benefits and limitations of computer-generated commentaries.

2. Data and Methods

Figure 1 illustrates the steps from downloading the court decision up to generating the final commentary. In step 1, we select a set of practically relevant BGB provisions: § 242 *Good faith*, § 280 *Damages for breach of duty*, § 812 *Unjust enrichment*, and § 823 *Tort liability*. They are the most cited provisions in court rulings and provide some diversity: §§ 280, 812, 823 BGB form the basis for obligation-related

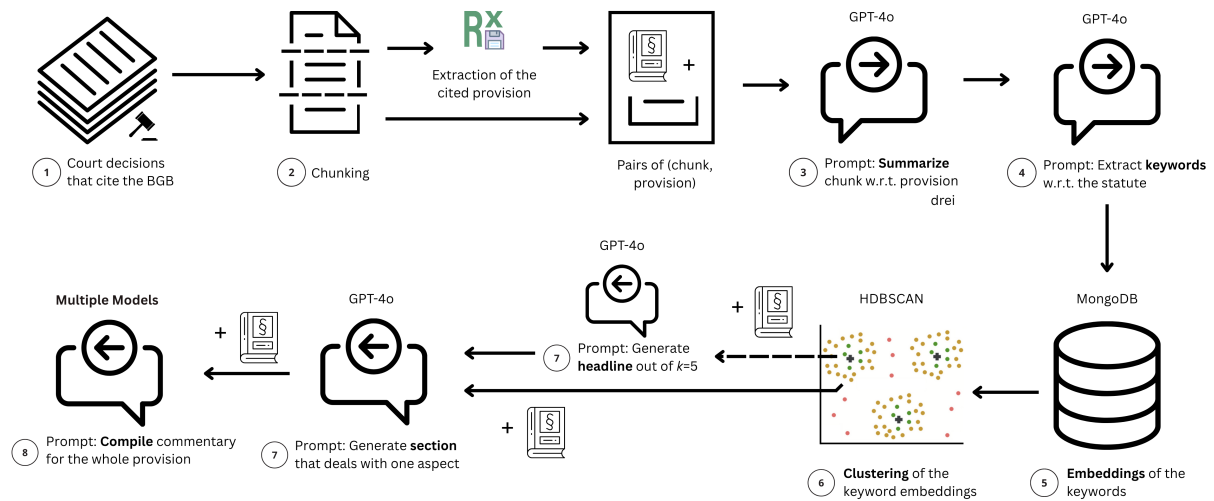


Figure 1: Our pipeline. In the clustering step (6), green circles represent the chunks used to generate headlines, while brown circles represent the broader cluster of chunks that is used to generate the section text for every headline. Red circles are regarded as noise.

damages and tort claims, § 242 BGB acts as a general guiding principle of good faith and fair dealing, which applies to all aspects of contract law and obligations. We rely on court rulings by the *Federal Court of Justice* (BGH), Germany’s highest court for civil and criminal matters. It hears appeals from lower courts to ensure uniform interpretation of the law. We downloaded all 4.555 publicly available BGH decisions that cite at least one BGB provision from the German “Case Law Online” portal.¹ A provision is cited in approximately 11 decisions on average, but the median is only 3, i.e. half of all BGB provisions appear in ≤ 3 court decisions. The four chosen legal provisions were cited far more frequently: in 509 (§ 823), 484 (§ 280), 357 (§ 242), and 260 decisions (§ 812). In step 2, we extract the “Reasons for Decision” section from every court decision, split it into paragraph-level chunks, and filter out paragraphs shorter than 100 characters (they mainly contain signatures or headings). This yields the following corpus sizes: § 280 (3.191 paragraphs; 927.333 tokens), § 242 (2.872; 824.678), § 823 (2.513; 785.704), and § 812 (2.390; 663.587). In step 3, we summarize each paragraph with GPT-4o, instructing the model to focus on how the paragraph applies the cited provision. Our pipeline automatically adds the text of the provision as context; in the case of multiple cited provisions, we provide all texts to enable the model to determine whether the chunk actually deals with the provision in question. From the summaries, we use GPT-4o to extract keywords that reflect the step of the provision’s application, again automatically providing the text of the provision as context (step 4). The keywords are embedded with `text-embedding-3-large` (step 5) and clustered per provision using HDBSCAN [7] (step 6) – an approach inspired by [5]. The records that form the clusters have the structure: `<keyword, summary, embedding>`. Clustering happens in embedding space and is formed based on the semantic similarity of the keywords. Each cluster must consist of at least 20 records, and records that cannot be assigned to a cluster, shown as red dots in Figure 1, are considered as outliers and not further processed. Each cluster has a centroid, marked as a black cross. Orange dots depict regular cluster members, green dots highlight the five records whose keywords are closest to the centroid. These keywords are used to generate the headlines, while the summaries of all the records within a cluster are used to generate the paragraph of this section (step 7). All headline–paragraph pairs are fed to the generative models GPT-4o, GPT-4.1, GPT-4.5, and the reasoning model o3 with instructions to merge them into a well-structured commentary (step 8). For details, see Prompt 1 in the Appendix. We applied this pipeline to generate four commentaries per model (§§ 242, 280, 812, and 823 BGB); however, the framework is generally applicable to all BGB provisions and can easily be extended to other statute books.

¹<https://www.rechtsprechung-im-internet.de/>

§ 242 BGB					
Model	Topical Relevance	Heading-Match	Citation-Faithfulness	Cluster-Distinction	Logical Ordering
GPT-4.1	3 5	4 5	3 4	4 4	3 4
GPT-4.5-preview	5 5	5 5	4 4	4 5	4 4
GPT-4o	3 4	3 5	4 4	3 3	3 3
o3	4 4	5 5	3 2	5 4	5 4
§ 280 BGB					
Model	Topical Relevance	Heading-Match	Citation-Faithfulness	Cluster-Distinction	Logical Ordering
GPT-4.1	3 5	4 4	4 4	3 5	3 5
GPT-4.5-preview	3 5	5 4	3 4	5 5	5 5
GPT-4o	2 5	4 5	4 2	3 5	3 5
o3	3 5	5 4	3 5	4 4	5 5
§ 812 BGB					
Model	Topical Relevance	Heading-Match	Citation-Faithfulness	Cluster-Distinction	Logical Ordering
GPT-4.1	4 4	5 3	4 2	2 3	3 2
GPT-4.5-preview	5 4	5 5	4 2	4 3	4 4
GPT-4o	3 3	4 2	3 1	2 1	3 2
o3	3 4	5 5	3 3	5 5	5 5
§ 823 BGB					
Model	Topical Relevance	Heading-Match	Citation-Faithfulness	Cluster-Distinction	Logical Ordering
GPT-4.1	4 4	4 5	4 3	4 5	3 3
GPT-4.5-preview	5 5	5 4	4 3	5 4	4 4
GPT-4o	3 2	3 5	3 4	2 2	2 2
o3	3 4	5 5	4 3	5 5	4 5
Average score across legal provisions					
Model	Topical Relevance	Heading-Match	Citation-Faithfulness	Cluster-Distinction	Logical Ordering
GPT-4.1	3.50 4.50	4.25 4.25	3.75 3.25	3.25 4.25	3.00 3.50
GPT-4.5-preview	4.50 4.75	5.00 4.50	3.75 3.25	4.50 4.25	4.25 4.25
GPT-4o	2.75 3.50	3.50 4.25	3.50 2.75	2.50 2.75	2.75 3.00
o3	3.25 4.25	5.00 4.75	3.25 3.25	4.75 4.50	4.75 4.75

Table 1

Evaluation scores for § 242 BGB, § 280 BGB, § 812 BGB, § 823 BGB and their overall average (1-5, the higher the better). For every model–criterion pair, human-annotated scores appear on the left, LLM-based scores on the right (human | LLM).

3. Evaluation

We use both LLM-as-a-judge (see Prompt 2 in the Appendix) and human evaluation to assess the overall quality of our generated commentaries (3.1). In addition, we provide an in-depth (qualitative) legal analysis of the generated texts (3.2).

3.1. Overall Quality

We evaluate the overall quality based on five criteria: First, we examine topical relevance, verifying that sections align with the legal provision. Second, we check heading-matching to ensure each heading accurately reflected its section’s content. Third, we test citation faithfulness by confirming that every citation genuinely supports the statement it followed. Fourth, we inspect cluster distinction, assessing whether the sections remained clearly separated without overlap. Fifth, we review logical ordering to determine whether the document maintains a coherent, continuous narrative. LLM-as-judge scores are provided by Gemini 2.5 Flash, which was not used for generation. Human evaluation is carried out by a member of our team with legal education in a blind process. The results of both are shown in Table 1.

GPT-4.5-preview leads with an average human evaluation score of 4.4 across all criteria and provisions, closely followed by o3 (4.2), while GPT-4o trails in every category with an average score of 3. The human judge rates GPT-4.5-preview highest in topical relevance and award both GPT-4.5-preview and o3 scores ranging from 4 to 5 for heading-match, cluster-distinction, and logical order. However, all models plateau around 3.5 for citation faithfulness. Gemini 2.5 Flash consistently scores commentaries higher by approximately half a point in topical relevance, cluster-distinction, and ordering. The closest alignment between model and human evaluations appears in citation faithfulness, where the automated

model scores slightly lower. Divergence between human and automated scores is most pronounced for § 812 BGB, with high human scores contrasting sharply with low LLM ratings. §§ 242 and 280 BGB show closer human-model agreement. The evaluation highlights three findings: First, GPT-4.5-preview and GPT-4o generate high-quality commentaries but have shortcomings in citation faithfulness. Second, the reasoning model o3 trails behind in topical relevance, but shows significant improvements in cluster distinction and logical ordering. Here, GPT-4o occupies the middle ground between GPT-4o and o3, pointing towards a sweet spot in a potential trade-off between the quality of the structure and the content. Third, the substantial gap between human and LLM scores casts doubt on the reliability of model-based evaluations. We have not measured inter-rater agreement, which we leave for future work.

3.2. Legal Analysis

An in-depth qualitative comparison of the generated texts with traditional legal commentaries reveals similarities and peculiarities.² First, the structure of the reports for §§ 280, 823, 812 BGB on the one hand, and § 242 BGB on the other hand, diverge – like in traditional commentaries. This is because the former are bases for claims, while the latter is a general provision with broad use cases. Provisions that function as the basis for claims provide the structure for their examination, which is reflected in traditional commentaries. All models mimic this quite accurately for § 280 BGB, with occasional hickups in the ordering. § 823 BGB and § 812 BGB are more challenging in this regard, because they combine multiple claims, each with diverging structures. GPT-4.1 and GPT-4o struggle with their separation, while GPT-4.5-preview manages to account for them, although not perfectly. This is where o3 shines. It presents two different structures for § 823 BGB and comes up with a surprising solution for § 812 BGB: After presenting key points, it switches to forming case groups. Breaking away from the dogmatic analysis appears reasonable in an area of the law where even the BGH acknowledges that certain cases “cannot be assessed in a general way, but must rather be determined on a case-by-case basis, taking into account the specific circumstances of the respective case” [8]. The problem with o3, however, is that it places too much emphasis on structure; the result sometimes resembles a collection of claim structures rather than a commentary. GPT-4.5-preview does a better job at incorporating more substance in the form of dogmatic explanations and example cases. Legal practitioners will find the results to be closer to a traditional commentary. Nevertheless, even GPT-4.5-preview remains on the “legal surface”. In particular, all generated texts lack discussions of specific problems of interpretation, in which different opinions (and arguments supporting them) are weighed against each other. This is what commentaries are frequently consulted for.

4. Discussion and Limitations

When compared to traditional, hand-written legal commentaries, their new generated counterparts display some promising features. They are able to handle massive amounts of sources. Furthermore, their creation is way cheaper, which allows developers to offer them at a lower price point; we publish our results at no cost. In contrast, hand-written commentaries are expensive. Physical copies have to be re-bought with new editions. While subscription-based online versions exist, their actualization still requires manual labor and is therefore slow. In contrast, our automated commentary can be re-generated with every new decision within minutes. Unlike [3], its creation does not require any specific legal knowledge and can therefore be easily extended to all areas of law. In comparison to [5], it nevertheless manages to recognize the dogmatic structures found in conventional commentaries. Does the future of legal commentaries therefore lie in the machine-made processing of court rulings? In light of our findings, we argue that LLMs can be a great help for collecting as well as evaluating sources and the general structuring of legal commentaries. If one takes into account the ever-growing body of statutes and court rulings, their usage might even become necessary. The proposed multi-step setup with chunking, abstract summarizations, keyword generations, clustering, and text generation

²The generated reports are available at <https://github.com/amelr250501/icail>

has led to impressive results in our experiments. However, this can only form a starting point and is not without limitations. As such, we identify the practical problem of limited sources (4.1) and a fundamental problem of value judgments (4.2).

4.1. Limited Sources

Like [1] and [2], we mine arguments from court decisions (whose judge authors presumably consulted legal commentaries, which in turn evaluated the legal literature). Directly mining the legal literature would provide an unfiltered, more diverse set of arguments. Their absence becomes obvious in the lack of any reported conflicts of opinion in the commentaries generated by us. The underlying problem is a practical one caused by the limited access to commercial databases in Germany. *Engel* and *Kruse* as well as *Santosh et al.* are less affected by it, because both the court rulings from the Federal Constitutional Court of Germany and the European Court of Human Rights are binding – a rare case in Civil Law. Even if a ruling is considered to be law-creating (e.g., by *Kelsen*’s “Pure Theory of Law”), the generated norm is individual; it does not become a general norm unless the legal system explicitly says otherwise. This is the case for both courts (§ 31[1] BVerfGG and Art. 46[1] EMRK). Conversely, court rulings of the BGH or arguments put forward in them carry, in theory, the same weight as arguments put forward by any lower court or legal scholar. While practitioners will nevertheless be guided by the “case law” of the BGH, arguments for contestation put forward in the literature are helpful where a – possible – deviating outcome is sought. Also, the literature deals with cases that have not yet been decided by a court. Well-founded concepts proposed by legal scholars can form the starting point of the decision-making process if the respective cases emerge in practice.

4.2. Value Judgments

This points to a conceptual problem that arises from the fact that, to be helpful in practice, commentaries are actually *expected* to be opinionated. *Christoph Engel* and *Johannes Kruse* argue that using LLMs with predefined inputs renders the creation of commentaries more objective [3]. Since our commentaries present their reports as mere summaries (i.e., “detached observations” [9]), one can go further and call this approach “scientific” in the strict sense of *Hans Kelsen*. Computer-generated commentaries like ours adhere to the ideal he has set for unauthentic interpretation provided by legal commentaries – they come close to what he regarded as “scientific commentaries” [10] that refrain from “committed arguments” [9]. However, traditional legal commentaries are opinionated in their reasoning *because court rulings and pleadings themselves are opinionated* in their reasoning. Commentaries serve as a practical tool; whether, like *Kelsen*, one wishes to deny them scientific character on that basis is a separate question. Practitioners seek from them a well-founded interpretation, fully aware that alternative interpretations may exist. Switching the style of the generated output from “detached observations” to “committed arguments”, however, raises a serious question: When mining arguments from multiple sources, whose value judgments form the basis of the proposed “right” solution? In this context, the use of the supposedly intelligent and objective LLM could even revive the long overcome belief that the one “right” solution can be found by technical means.

5. Conclusions

This paper demonstrates that when used in carefully designed pipelines, LLMs can create doctrine-oriented legal commentaries at the speed with which new court decisions appear. Starting from 4.555 rulings of the German Federal Court of Justice, our system is capable of creating reports for §§ 242, 280, 812 and 823 of the German Civil Code without handcrafted doctrinal scaffolding – and can be extended to any number of provisions. In a blind review, our expert rates the best variant (GPT-4.5-preview) at 4.4/5 across relevance, structure and logical order. Citation grounding remains the chief weakness. Compared with traditional hand-written commentaries, the machine-generated versions can cover

more decisions, can be regenerated within minutes and cost orders of magnitude less. Yet, they do not contain the "committed arguments" legal practitioners will typically expect.

Acknowledgements

This work was carried out within the project "Generatives Sprachmodell der Justiz (GSJ)", a joint initiative of the Ministry of Justice of North Rhine-Westphalia (Ministerium der Justiz des Landes Nordrhein-Westfalen) and the Bavarian State Ministry of Justice (Bayerisches Staatsministerium der Justiz), with the scientific partners Technical University of Munich (Technische Universität München) and University of Cologne (Universität zu Köln). The project is financed through the Digitalisierungsinitiative des Bundes für die Justiz.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] K. Schönfeld, *Rex, lex et judex: Montesquieu and la bouche de la loi revisited*, *European Constitutional Law Review* 4 (2008) 274–301. doi:10.1017/S1574019608002745.
- [2] K. I. Schmidt, *Law, modernity, crisis: German free lawyers, american legal realists, and the transatlantic turn to "life," 1903–1933*, *German Studies Review* 39 (2016) 121–140. URL: <http://www.jstor.org/stable/24809061>.
- [3] C. Engel, J. Kruse, *Kommentar ohne autor: Können sprachmodelle das kommentieren übernehmen?*, *JuristenZeitung (JZ)* 79 (2024) 997–1007. doi:10.1628/jz-2024-0304.
- [4] C. Engel, J. Kruse, *Professor GPT: Having a Large Language Model Write a Commentary on Freedom of Assembly*, Technical Report 2024/14, Max Planck Institute for Research on Collective Goods, 2024. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4994131. doi:10.2139/ssrn.4994131.
- [5] T. Y. S. S. Santosh, M. Aly, O. Ichim, M. Grabmair, *Lexgenie: Automated generation of structured reports for european court of human rights case law*, 2025. URL: <https://arxiv.org/abs/2503.03266>. arXiv:2503.03266.
- [6] OpenAI, *Openai deep research*, <https://openai.com/deep-research>, 2025. URL: <https://openai.com/deep-research>, accessed: 2025-05-02.
- [7] R. J. G. B. Campello, D. Moulavi, J. Sander, *Density-based clustering based on hierarchical density estimates*, in: *Advances in Knowledge Discovery and Data Mining (PAKDD 2013)*, volume 7819 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 160–172. URL: <http://dblp.uni-trier.de/db/conf/pakdd/pakdd2013-2.html#CampelloMS13>.
- [8] Bundesgerichtshof, *Urteil vom 16. märz 2006 – iii zr 62/05, 2006*. URL: <https://www.bundesgerichtshof.de>, "Die Frage, ob ein Anspruch auf Herausgabe wegen ungerechtfertigter Bereicherung im Mehrpersonenverhältnis besteht, entzieht sich einer pauschalen Beurteilung und ist vielmehr im Einzelfall unter Berücksichtigung der konkreten Umstände des jeweiligen Falles zu entscheiden."
- [9] G. P. Fletcher, *Two modes of legal thought*, *Yale Law Journal* 90 (1981) 970–1003. URL: https://scholarship.law.columbia.edu/faculty_scholarship/244.
- [10] H. Kelsen, *The Law of the United Nations: A Critical Analysis of Its Fundamental Problems*, Stevens Sons Limited, London, 1951.

Appendix

You are a German attorney.

Consider the preceding legislative text ("normative text") only to the extent necessary to correctly anchor definitions, structure, and the telos of the norm; do not quote it extensively or summarize it literally.

Revise the following draft linguistically, restructure it logically, number and name the headings consistently, so that each section logically builds upon the previous one. Connect transitions clearly and avoid redundancy. Avoid generic headings (e.g., "*Concept*," "*Practical Case*"). Do not include concluding summaries.

The commentary should reflect the abstract structure of the application of the norm but should also include concrete examples if they appear in the draft. Such examples should only relate to selected aspects of the norm. Do not invent examples; only use those included in the draft.

Write in an objective, formal style. This commentary is intended for trained legal practitioners, not for students. The output may be longer than the input. Avoid bullet points; instead, formulate their content in full sentences.

IMPORTANT: Each occurrence of `ObjectId('...')` corresponds to a reference and must remain as such in the final text if retained.

Return **only the commentary**.

Du bist ein deutscher Rechtsanwalt. Berücksichtige den vorangestellten Gesetzestext („Normtext“) lediglich in angemessenem Umfang, um Definitionen, Systematik und Telos der Norm korrekt zu verankern; zitiere ihn nicht ausführlich und fasse ihn nicht wörtlich zusammen.

Überarbeite den folgenden Entwurf sprachlich, strukturiere ihn logisch um, nummeriere und benenne die Überschriften konsistent, sodass jeder Abschnitt sinnvoll auf den vorangehenden aufbaut. Verbinde Übergänge, vermeide Redundanz. Vermeide generische Überschriften (z.B. "*Begriff*", "*Praxisfall*"). Verzichte auf abschließende Zusammenfassungen.

Der Kommentar soll die abstrakte Struktur der Anwendung der Norm widerspiegeln, aber auch konkrete Beispiele beinhalten, wenn diese im Entwurf vorkommen. Solche Beispiele sollen sich nur auf ausgewählte Aspekte der Norm beziehen. Erfinde keine Beispiele, sondern greife nur in dem Entwurf enthaltene Beispiele auf.

Schreibe in einem sachlichen, formalen Stil. Es handelt sich um einen Kommentar für ausgebildete Rechtsweniger, nicht etwa für Studierende. Der Output darf länger sein als der Input. Verzichte auf Stichpunkte, formuliere deren Inhalt stattdessen aus.

WICHTIG: Jede Stelle `ObjectId('...')` entspricht einer Referenz und muss als solche im finalen Text erhalten bleiben, falls der Text übernommen wird.

Gib **nur den Kommentar** zurück.

Prompt 1: Generation of the final commentary, German original and translation

Evaluation Criteria

Critically evaluate the text of a German legal commentary and assign a score from **1 (barely satisfactory)** to **5 (very well fulfilled)** for each of the following criteria:

1. **Topical Relevance:** Do the headings cover all undefined terms of the underlying norm?
2. **Heading-Match:** Does each paragraph fully meet the content promised by its heading?
3. **Citation-Faithfulness:** Do the cited references genuinely support the statements (no hallucinations)? Are all referenced documents locatable?
4. **Cluster-Distinction:** Is the content clearly distinct with minimal or no overlap with other sections within the text (clear thematic demarcation)?
5. **Logical Ordering:** Does the placement of each section logically fit into the overall structure (coherent thread, comprehensible sequence)?

Return only the result in JSON format—without any additional text.

Bewertungsrichtlinien

Bewerte den Text eines deutschen juristischen Kommentars **kritisch** und vergebe für jedes der folgenden Kriterien einen Wert von **1 (befriedigt kaum)** bis **5 (sehr gut erfüllt)**:

1. **Topical Relevance:** Decken die Überschriften alle unbestimmten Begriffe der zugrunde liegenden Norm ab?
2. **Heading-Match:** Entspricht jeder Absatz inhaltlich vollständig dem Versprechen der Überschrift?
3. **Citation-Faithfulness:** Stützen die angegebenen Fundstellen die Aussagen tatsächlich (keine Halluzinationen)? Werden alle referenzierten Dokumente gefunden?
4. **Cluster-Distinction:** Deckt sich der Inhalt nicht oder nur minimal mit anderen Abschnitten innerhalb des Texts (klare thematische Abgrenzung)?
5. **Logical Ordering:** Passt die Position jedes Abschnitts in die Gesamtstruktur (roter Faden, nachvollziehbare Reihenfolge)?

Gib ausschließlich das Ergebnis im JSON-Format zurück – ohne weiteren Text.

Prompt 2: Evaluation of the commentary. German original and translation