

Adapting Requests for Production for AI GLLM Assessment

*Craig Ball**

Copyright 2024. All rights reserved.
Reprinted with permission.

* More of the author's work can be found at <https://craigball.net>.

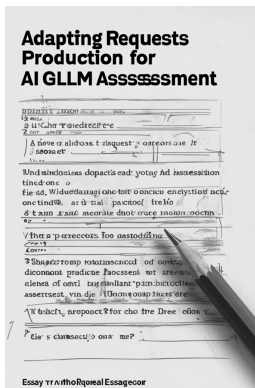


Adapting Requests for Production for AI GLLM Assessment

19 Monday, Aug 2024

POSTED BY CRAIGBALL IN UNCATEGORIZED

≈ 6 COMMENTS



The integration of Generative Large Language Models (GLLMs) into the discovery process is transforming how documents are reviewed for relevance and responsiveness. These AI models, which excel at processing large document collections, offer significant efficiency improvements. However, to harness their full potential, requests for production (RFPs) must evolve to reflect the unique capabilities and limitations of AI systems. Traditional language in RFPs, which relies heavily on human intuition, needs to be adjusted to accommodate AI's reliance on clear instructions, context, and precision. This post explores how to adapt requests for production to optimize GLLM usage, addressing both business disputes and tort claims. I'll provide examples to illustrate how common RFPs can be refined for AI-assisted document review.

Effective AI Prompts in Discovery

AI systems like GLLMs function best with well-structured prompts. In the context of discovery, this means adjusting RFPs to emphasize clarity, specificity, and relevance. The key elements for constructing effective AI prompts in legal discovery are:

1. **Clarity and Specificity:** Ambiguity can cause AI systems to miss important documents or misclassify irrelevant ones. Specific requests guide the AI more effectively.
2. **Contextual Guidance:** AI relies on context to assess relevance. Providing additional background or specifying the purpose of certain requests helps refine the search.
3. **Keyword Precision:** GLLMs rely on keywords to understand and evaluate document content. Choosing precise terms helps reduce the retrieval of irrelevant documents.
4. **Examples:** AI systems can better identify relevant documents if examples are provided within the RFP, as they offer patterns for the system to follow.

Incorporating these principles into RFPs ensures that AI models can make the most accurate assessments during document review.

Adapting Requests for Production: Business Dispute and Tort Claim Examples

Both business disputes and tort claims involve unique types of documents and keywords. Adapting RFPs to suit GLLM's capabilities involves providing detailed instructions for both.

1. Refining RFPs for Business Disputes

Business disputes typically involve contracts, communications, financial records, and project-related documents. A traditional RFP in this context can be too broad, but when refined for AI, it becomes much more effective.

Traditional RFP:

"Produce all documents related to the termination of the contract between the plaintiff and the defendant."

Adapted for GLLM:

"Produce all communications, including emails, letters, and meeting notes, regarding the termination of the contract between the plaintiff and defendant, specifically mentioning breach of contract, missed deadlines, or failure to perform. Additionally, provide any internal communications among defendant's employees discussing the reasons for contract termination."

Explanation: The original request is too broad for AI processing, potentially leading to the inclusion of irrelevant documents. The adapted request clarifies the type of documents (emails, meeting notes, etc.) and key topics of interest (breach of contract, missed deadlines), which helps the AI focus on the relevant aspects of the dispute.

2. Refining RFPs for Tort Claims

Tort claims, on the other hand, often involve medical records, accident reports, witness statements, and other documentation related to personal injury, negligence, or harm. To leverage AI's strength in these cases, RFPs should focus on the specific nature of the injury or event and include key contextual information.

Traditional RFP:

"Produce all documents related to the accident on March 10, 2022, involving the plaintiff."

Adapted for GLLM:

"Produce all accident reports, witness statements, photographs, medical records, and insurance communications related to the accident on March 10, 2022, involving the plaintiff. Specifically include any internal communications from defendant's employees discussing liability, cause of the accident, or settlement offers."

Explanation: The original request does not provide enough guidance for AI to identify the most relevant documents. The refined version specifies categories of documents (accident reports, witness statements) and includes critical keywords like "liability" and "settlement offers," which help the GLLM to filter out irrelevant documents and focus on those pertinent to the tort claim.

Providing Context for AI in Business and Tort Claims

In both business disputes and tort claims, providing contextual guidance can improve AI's performance. AI systems are good at picking up on patterns when given context that defines what is relevant in a particular dispute.

Business Dispute Example

Traditional RFP:

"Produce all documents related to the defendant's performance of its obligations under the agreement."

Adapted for GLLM:

"Produce all project reports, timelines, performance evaluations, internal communications, and external correspondence related to the defendant's performance of its obligations under the agreement, including any discussion of delays, failures to meet contractual benchmarks, or disputes over quality between January 1, 2019, and December 31, 2021."

Explanation: By specifying the categories of documents (project reports, timelines, performance evaluations) and key terms (delays, failures, disputes), the adapted request gives the AI more clarity on what aspects of performance are of interest, thus reducing irrelevant results.

Tort Claim Example

Traditional RFP:

"Produce all documents related to the plaintiff's injury."

Adapted for GLLM:

"Produce all medical records, doctor's notes, diagnostic reports, witness statements, accident scene photographs, and insurance communications related to the plaintiff's injury, specifically discussing the severity of the injury, treatment options, or any references to negligence by the defendant."

Explanation: The broad language of the traditional request may lead AI to include irrelevant medical records or discussions not related to the injury in question. By specifying categories of documents and important keywords like "negligence" and "severity," the adapted RFP allows AI to better focus on documents that are directly relevant to the tort claim.

Keyword Precision in Business and Tort Cases

Keywords are crucial in helping AI distinguish relevant documents from irrelevant ones. In both business disputes and tort claims, identifying precise keywords is essential to avoid overly broad results.

Business Dispute Example

Traditional RFP:

"Produce all documents related to the pricing agreement between the parties."

Adapted for GLLM:

"Produce all emails, pricing models, rate sheets, invoices, and contract amendments related to the pricing agreement between the parties, including any discussions regarding price increases, discounts, or fee disputes from January 1, 2020, to December 31, 2021."

Explanation: The refined request identifies specific document types (emails, pricing models, rate sheets) and includes key terms such as “price increases” and “fee disputes” that will help AI focus on relevant pricing information rather than capturing unrelated financial documents.

Tort Claim Example

Traditional RFP:

“Produce all documents related to the plaintiff’s medical treatment.”

Adapted for GLLM:

“Produce all medical treatment records, diagnostic reports, physician evaluations, hospital discharge summaries, and communications with healthcare providers related to the plaintiff’s medical treatment following the accident on March 10, 2022, particularly those discussing long-term prognosis, disability, or permanent impairment.”

Explanation: The traditional request may capture too many irrelevant medical records. By specifying types of records (treatment records, diagnostic reports) and relevant keywords (long-term prognosis, disability), the GLLM is more likely to return documents pertinent to the plaintiff’s specific injury and treatment following the accident.

Using Examples in Requests

Providing examples within RFPs can guide AI systems in identifying relevant documents. Examples help the AI model detect patterns and apply those to the rest of the document collection.

Business Dispute Example

Traditional RFP:

“Produce all communications related to the parties’ negotiations.”

Adapted for GLLM:

“Produce all emails, letters, and meeting notes related to the parties’ contract negotiations, particularly any communications discussing changes to the payment terms or performance obligations. For example, any document similar to the email from Jane Smith to John Doe on April 15, 2021, discussing contract amendments should be included.”

Explanation: By referencing a specific email, the AI can identify patterns in that communication and search for similar documents, thereby improving accuracy.

Tort Claim Example

Traditional RFP:

“Produce all documents related to settlement negotiations.”

Adapted for GLLM:

"Produce all emails, letters, settlement proposals, and internal communications related to settlement negotiations in the tort claim, particularly any documents discussing potential offers, counteroffers, or liability assessments. For example, any document similar to the settlement offer email sent by John Smith on June 1, 2023, should be included."

Explanation: The use of an example helps AI learn from the context of that specific communication and apply it to other similar documents, increasing the chances of accurately capturing settlement-related discussions.

Conclusion

As GLLMs become more prevalent in the discovery process, the drafting of RFPs must adapt to these AI capabilities. Traditional broad requests fail to provide the clarity and specificity necessary for GLLMs to perform optimally. By incorporating elements like clear categories of documents, contextual guidance, keyword precision, and examples, RFPs can be tailored to the strengths of AI, ensuring a more efficient and effective document review. Whether in business disputes or tort claims, these adaptations enable legal teams to take full advantage of AI's capabilities, improving the relevance and accuracy of the documents produced during discovery.

Postscript

This essay is unlike anything I've ever posted here. This morning, I read [a post at ComplexDiscovery](#), my friend Rob Robinson's excellent blog resource, discussing the shift toward GLLM's. Producing parties are using new tools and technique to save money and time, so I'm always wondering how requesting parties can take steps to ensure that the focus on faster and cheaper doesn't defeat the legitimate ends of discovery, *viz.* getting to the most relevant, probative evidence in reasonable ways. So, I had a minor flash of inspiration and tweeted (X'ed?) [this](#): **"With generative large language models (AI) used to assess information for responsiveness, wise counsel frame requests for production to include elements of effective AI prompts. If you want their AI to find what you need, tell it how to discriminate."**

It's fine to say that and quite another thing to implement the idea.

That got me thinking ("a dangerous pastime, I know"). How might I flesh out that notion for my readers? The post at Complex Discovery carried the byline "*ComplexDiscovery Staff*," an amorphous attribution that, to my suspicious mind, smacks of 'AI-generated.' So, I should tell you that the essay above was the work of "*Ball in your Court Staff*," meaning me prompting ChatGPT and then reading the result to see if I can stand behind it. It felt a bit like cheating, hence this confessional postscript. I used the following prompt: *"Write an essay of under approximately 1,000 words describing specific ways to adapt requests for production in discovery to incorporate the salient elements of effective prompts when AI GLLMs are used to assess document collections for relevance and responsiveness. Supply examples of tweaks to common requests in business disputes and tort claims to be better adapted to this usage."* What emerged isn't bad.; tepid to be sure, not my voice and unlikely to be embraced by counsel ever-apprehensive of framing a request too-narrowly. Still, it faithfully fleshes out the idea and—hopefully—gets **you** thinking about how to take better control of the delegation of your requests to our robot overlords.

thoughts on "Adapting Requests for Production for AI

GLLM Assessment”

1. Pingback: [Adapting Requests for Production for AI GLLM Assessment - TQT Group](#)

2. *said:*[davidkeithtobin](#)

[August 19, 2024 at 11:34 AM](#)

great stuff – and for free! thank you!

REPLY

3. Pingback: [Week 34 – 2024 – This Week In 4n6](#)

4. Pingback: [Adapting Requests for Production for Generative AI](#)

5. Pingback: [AI Generated or AI Assisted? There is a Difference](#)

6. *said:*[John Tredennick](#)

[August 27, 2024 at 8:31 AM](#)

Nice post Craig. I plan to share it with my senior prompt engineer. We are dealing with these issues every day and it never ceases to be interesting. On the one hand, I am constantly amazed at how these LLMs can understand my requests, however poorly formed. At the same time I realize and agree that context and specificity are critical.

One small point regarding this statement: “These AI models, which excel at processing large document collections . . .

I don’t think they excel at processing large document collections. To the contrary they are at their least efficient (and most expensive) when plowing through large document collections. A more effective model is to use other forms of AI to narrow down the population and use GenAI to analyze the most likely relevant files. But that is another blog post.

Thanks

JT

REPLY

[Website Powered by WordPress.com.](#)

6