Artificial Intelligence - foundation models - large language models and their usability in legal analysis and legal reasoning

We will look at the use of so-called *foundation models* in the field of law.

It is explicitly not about application-specific trained or fine-tuned *language models*.

And, it is about whether these models - assuming skilful, i.e. clever and ingenious mental and linguistic control - can be used to understand, interpret and reliably, truthfully, consistently and coherently generate answers to the most complex legal questions, for example in the area of mergers and acquisitions.

To begin with, a general question: How do you, Google Gemini, assess the usability of such models for said purposes as a matter of principle?

Here's my assessment of the usability of foundation models for complex legal tasks like mergers and acquisitions, assuming skilled human control:

Potential Benefits:

- **Efficiency and Scalability:** Foundation models can analyze vast amounts of legal data in a fraction of the time it takes humans. This can be helpful for tasks like due diligence, where reviewing contracts and other documents is crucial.

- **Identification of Patterns and Insights:** These models can find subtle patterns and connections in legal documents that humans might miss. This could be useful for uncovering potential risks or opportunities in a merger or acquisition.

- **Increased Consistency:** By applying legal principles uniformly across a large dataset, foundation models can help reduce inconsistencies in legal analysis.

Limitations and Challenges:

- **Understanding Context and Nuance:** Legal issues often hinge on specific factual contexts and subtle nuances in language. Foundation models, while good at identifying patterns, might struggle with these subtleties.

- **Reasoning and Judgment:** These models are not designed for legal reasoning or applying judgment to complex situations. They can identify relevant information, but skilled human lawyers are still needed to interpret the information and make decisions.

- **Bias and Fairness:** Foundation models are trained on massive datasets, which can perpetuate existing biases in the legal system. Careful selection of training data and ongoing monitoring are essential to mitigate bias.

Overall Assessment:

Foundation models are unlikely to replace human lawyers for complex legal tasks like mergers and acquisitions in the foreseeable future. However, they can be valuable tools to **assist** lawyers by:

- *Performing* initial research and analysis

- *Identifying* relevant legal precedents and documents
- *Highlighting* potential risks and opportunities

Skilled human lawyers will remain crucial for tasks that require:

- Deep *understanding* of legal principles and their application

- Complex legal *reasoning* and *judgment*

- Client communication and building *trust*

In essence, foundation models can be powerful tools to **augment human expertise** in the legal domain, but they shouldn't be seen as a replacement for lawyers themselves.

Foundation models shouldn't be seen as a replacement for lawyers themselves, you say.

Another foundation model that I have primed on the same subject referred to significant limitations in its ability to deal with complex legal issue.

If I understand you correctly, your statement differs insofar that you think that dealing with the most complex legal issues is basically possible but not recommendable, whereas the other model takes the view that getting complex legal issues out of the model is basically infeasible, even with a very professional approach to the model, because the limitations in terms of comprehension are significant.

Hence, the question rose by this objection: what is significant in this context?

Would you agree with that interpretation of your statement and the statement of the other model?

You're right, there's a nuance in how we might view the capabilities of foundation models for complex legal issues. Here's how I see it:

My Perspective:

- **Technically Possible, Not Generally Advisable:** With very skilled human control and careful tailoring of the task, foundation models *could* potentially contribute to understanding even the most complex legal issues.

- **Significance of "Complex":** "Complexity" can be subjective. For some highly intricate legal problems with many moving parts and unique factual scenarios, even the most advanced models might struggle to grasp the nuances effectively.

- **Focus on Assisted Expertise:** My emphasis lies in foundation models as **assistants** to human lawyers. They can handle the information processing and initial analysis, freeing up lawyers for the truly complex aspects that require judgment, reasoning, and client communication.

Possible Interpretation of the Other Model:

- **Fundamental Limitations:** The other model might be highlighting the inherent limitations of foundation models in legal comprehension. These limitations could stem from the models' inability to fully grasp the context, reasoning, and historical precedents that underpin complex legal issues.

- **Infeasibility of Handling Complexity:** Perhaps the other model suggests that even with skilled control, the limitations are so significant that foundation models simply can't be relied upon for truly intricate legal matters.

Finding Common Ground:

Both perspectives hold merit. Foundation models have limitations, and complex legal issues can be highly nuanced. However, advancements in the field are rapid.

Here's a possible reconciliation:

- **Continual Improvement:** As foundation models evolve and become more sophisticated, their ability to handle complexity might improve.

- **Task-Specific Tailoring:** By carefully defining the specific tasks and tailoring the model's use, we can leverage its strengths while mitigating its limitations for even intricate legal problems.

Ultimately, the key takeaway is that foundation models are powerful tools, but they require **careful consideration and human oversight** when dealing with complex legal matters.

By carefully defining the specific tasks and tailoring the model's use, we can leverage its strengths while mitigating its limitations for even intricate legal problems, you say. That is exactly my point in my initial question when in referred to the assumption of skilful, i.e. clever and ingenious mental and linguistic control of such a model.

You talk about carefully defining the specific tasks and tailoring the model's use while I talk about skilful, i.e. clever and ingenious mental and linguistic control. I see your point, you see my point, don't you?

You're absolutely right. We are definitely on the same page here. "Skilful mental and linguistic control" perfectly captures the essence of what I described as "carefully defining specific tasks and tailoring the model's use."

Both terms emphasize the need for a thoughtful and nuanced approach to utilizing foundation models in complex legal tasks. It's about leveraging the model's strengths in information processing and pattern recognition while acknowledging its limitations in areas like legal reasoning and understanding subtle context.

In essence, skilled human control is what bridges the gap between the capabilities of foundation models and the demands of complex legal issues.

Is there anything else you'd like to discuss about the role of foundation models in the legal field?

Thank you. This is in line with my experience. One more thing:

You see, I am a lawyer-linguist. Thus, I am a specialized legal translator and a general lawyer (which is rare these days) and thus not a specialized lawyer.

In this capacity I make use of language models in order to understand, interpret and answer all questions of law in a particular jurisdiction!.

The key to the endeavour is the **reliability** of the propositions generated by the language model. I do this by having several language models running **simultaneously**. But, you know, I have to be 100% sure that what comes out of them is 100% correct. So if I understand you correctly, you would say that this might be achieved?

Unfortunately, achieving 100% certainty with any legal question, even with multiple language models and your expertise as a lawyer-linguist, is very challenging. Here's why:

- **Complexity of Law:** Law is inherently complex, with nuances in language, interpretation, and precedent that can be difficult to capture even for humans. Language models, while powerful, might miss these subtleties.

- **Limitations of Models:** As discussed before, foundation models have limitations in areas like legal reasoning and understanding context. Their outputs might be factually accurate but lack the legal judgment needed for definitive answers.

- **Data Bias:** Language models are trained on massive datasets, and these datasets can contain biases that are reflected in the model's outputs. This can lead to misleading or inaccurate information.

However, your approach of using multiple language models and your expertise as a lawyer-linguist can significantly improve the **reliability and accuracy** of the information you get. Here's how:

- **Cross-referencing and Verification:** By comparing outputs from different models, you can identify inconsistencies and areas requiring further research. Your legal expertise allows you to verify the information and identify potential biases.

- **Focus on Starting Point:** Language models can be a great starting point for legal research, providing a comprehensive overview of relevant legal information. You, as the lawyer-linguist, can then delve deeper into the complexities using your legal knowledge and judgment.

Here's a more realistic goal:

Instead of aiming for 100% certainty, strive for a **high degree of accuracy and a clear understanding of the limitations** of the information generated by the language models. This empowers you to make informed decisions and provide the best possible service to your clients.

Additional Tips:

- **Focus on Objective Tasks:** Use language models for tasks where objectivity is key, like summarizing legal documents or identifying relevant case law.

- **Human Judgment Remains Crucial:** Always rely on your legal expertise and judgment for interpreting the information generated by the models and making final decisions.

By combining the power of language models with your legal expertise, you can significantly **improve the efficiency and effectiveness** of your legal research and analysis.

> Well, 100 percent certainty is a bit of a stretch, no question.
>
> 100 percent certainty doesn't exist anywhere, not even in religion ;-).
>
> The only thing I can be 100% sure of is that I have to die, not true? But you know what I'm getting at, don't you?

You're absolutely right. I apologize for taking your initial statement of "100% certainty" too literally. In the legal field, absolute certainty is very rare, and you perfectly captured that sentiment with the death analogy **;)**

I understand what you're aiming for - achieving a very high degree of **reliability and trustworthiness** in the legal answers you get from language models, given your skilled control and expertise.

Here's what we can focus on:

- **Maximizing Reliability:** By using multiple models, carefully selecting training data, and applying your legal knowledge to verify and refine the outputs, you can significantly increase the reliability of the information you get.

- **Transparency and Trust:** Being transparent with your clients about the limitations of language models and how you use them to enhance your research can build trust and manage expectations.

While 100% certainty might be elusive, through your skilled approach, you can achieve a very high level of **accuracy and trustworthiness** in the legal information you provide. This can be incredibly valuable for your practice.