



**HL INFLUENCERS:
DIGITAL TRANSFORMATION
TRANSCRIPT
Alexander Kardos
Nyheim, Thomson
Reuters**

Leo von Gerlach	Hello everybody and welcome to another edition of <i>The Influencers</i> , our podcast conversation on digital transformation and law. I'm Leo von Gerlach and with me today is Alexander Kardos Nyheim. Alex is the epitome of one of those whiz kids of the emergent AI economy. So when Alex left Cambridge Law, he founded a legal tech company, Safe Sign Technologies, which focused on the optimization of large language models for law. And once that was successful, He sold it incredibly well to Thomson Reuters, where he is also now continuing on the same topic, leading Thomson Reuters research into machine learning. With that, Alex, terrific to have you. Welcome to the show.
Alexander Kardos Nyheim	Pleasure to be here, Leo.
Leo von Gerlach	Alex, I didn't mention that you were also a trainee lawyer with A.O. Sherman. at the time you started out after graduation. And I understand it was at the time that you had the inspiration to found Safe Sign Technologies, probably because there were some stuff that you saw that needed to be fixed. But perhaps you can just give us a little bit of sense of what triggered you to set up Safe Sign Technologies in the first place when you were at Big Law?
Alexander Kardos Nyheim	I first had the idea to found Safe Sign in my first seat of my training contract at Allen & Overy. I was sitting in David Wakeling's department of the Markets Innovation Group, which at the time was focused on and remains focused on AI. They were really the first team, I think, of any BigLaw firm to adopt Harvey. So I was very lucky to be at the cutting edge of BigLaw's adoption of these new technologies as Harvey & Co were emerging. What was missing though, from my perspective, was that the underlying technology behind the likes of Harvey was not specialized on law. And just from an intuitive perspective, I thought it does not seem right in principle that tools that are being used for high stakes, accuracy critical use cases in law are using ultimately an engine, if you like, that is not specialized on law. Even if in the layers above the core LLM, the core large language model, there is optimization for law, the large language model itself will have a performance ceiling because of its lack of training focus on law. That will mean that it is only so good at the law. And for me, I did not have a technological background, but I thought that seems like a potential gap, a potential

	<p>structural weakness of these systems. And it was on that basis that I founded Safe Sign with the idea of developing a legal enhanced large language model. So a model that is taken from an open-source provider and then specialized on law with high quality training data and other algorithmic innovations. The idea being, and this remains a hypothesis that has not been definitively proven, but the idea being that if you can ground a model in law, in legal reasoning, in a certain way of approaching a question and answering a question that we as lawyers are trained to do, It will eventually, and in the long term, outperform even very powerful new iterations of general large language models that come out from the big frontier labs. This is, again, not the end of the battle. I am very much fighting this on a day-by-day basis against the likes of Anthropic, Google DeepMind, OpenAI. But that's my contention, and it was on that basis that my team and I founded Safe Sign.</p>
<p>Leo von Gerlach</p>	<p>We will definitely get back to that contention and dive a little bit deeper into those conflicting, I think, propositions, technology propositions that are out there at the moment. But just for one second, stay with your just intriguing career because what I would also like to know and kind of understand better, when you are deep into your startup, there are all sorts of survival questions. What does it make successful? And do you have kind of a recommendation and understanding of the magic sauce? Is it what everybody calls product market? So to have kind of the right insight, is it the data access or some kind of business advantage? Is it the integration of your workflow and the execution of how you set up your organization? Or what is it - a combination or perhaps you give us some insight of the success story?</p>
<p>Alexander Kardos Nyheim</p>	<p>I by no means have a formula that would always work. And my context is quite specific because we did not develop a final end product per se. We were developing a core foundational large language model. So that's my first caveat. And therefore, we were really an AI research company more than a legal technology product company. I would caveat it with that. What allowed us to succeed was, in our case, the quality of the scientific expertise we had on the team. I was very lucky to be joined by people from Google DeepMind, Harvard, MIT, some of the top AI labs in the world where they were working on the deeper scientific problems that were limiting the safety, robustness, and reliability of LLMs. I had people actually coming over from medical schools where they had been working on AI systems in the drug discovery context. So there was a lot of very interesting connections that were made and very powerful and good scientific expertise that made us able to publish papers and be, in a very narrow context, a leading research company. And I think that is what fits with Thomson Reuters, because they were looking for, I think, really a research capability, the strategic capability to develop their own proprietary, insofar as it can be proprietary, and you would know this in a much more authoritative way, but to develop their own capability to improve large language models on the things they care about. So our product market fit was almost the research capability we brought to a very large legal tech player, and that is the basis on which we were acquired. The other bit I would say is we had very high-quality legal data</p>

	<p>that we created, and we had very high-level legal expertise that we benefited from. Rather than taking maybe a more traditional approach to training at LLM, which is large amounts of sometimes low grade training data, we focused on very small amounts of high impact, very high quality training data, which had a significant impact in improving the model, surprisingly significant impact, more so than large amounts of lower quality data. So we had an emphasis on quality over quantity on the data side. And that is now being brought to bear at scale because, of course, Thomson Reuters has a lot of its own high-quality legal data that we're employing to improve the model further. So I would say it was a combination of scientific talent and very high-quality legal data and legal domain expertise.</p>
<p>Leo von Gerlach</p>	<p>So that's very interesting. You just mentioned right at the beginning that you were research heavy and very introspect in a sense, and that may be a similar situation of deep mind from where you have acquired people has experienced once long time ago they joined a big ship like Google and all those considerations we could read in the newspapers when OpenAI at the time when some Altman was temporarily on garden leave, was considering to join Microsoft because you have then this kind of big jump from a research organization, which is small and very focused, to a large business organization. So that gives you some additional insights of how this shift actually works. And that's also something I'm extremely interested in to see how that came down in your own experience.</p>
<p>Alexander Kardos Nyheim</p>	<p>I agree that in an AI research company, it is the closest thing in the business world that you can have to an academic research lab. Of course, there are still commercial objectives to the research, but the team is given the oxygen, the peace and quiet to conduct quite deep and high-quality work without having to rush in sprints to meet certain product objectives. That was a very powerful factor as to why we were able to develop such good technology as a startup. And it was one of our key conditions on acquisition that we would still be allowed to operate effectively as a startup, effectively independently with the resources we needed to continue. Of course, as I'm sure DeepMind found out at a much larger scale when it joined Google, the reality is not always that. Of course, we have to navigate the big organizational, institutional hurdles. There is still, of course, some bureaucracy and all the rest, but I have really tried to insulate my team from the noise that comes with joining such a large organization. And as really my team is an academic team, coming out from many of these universities, I want to allow them to continue to publish their papers, to attend the leading AI conferences. And one of my collaborators, Dr. Jonathan Schwartz, has a very big focus on allowing his team to publish as much research as possible. And in fact, we've just established a partnership with Imperial College London, where we jointly develop research together, and then we'll be publishing and even possibly releasing models through this Imperial College partnership. So I think key to good research and good outcomes is allowing your team of academics to operate in the same way, though, of course, understanding</p>

	<p>that they have some commercial objectives they need to serve too, and allowing academics to be academics.</p>
<p>Leo von Gerlach</p>	<p>Speaking about research and academics, I think that brings us back to the beginning of what you told me about your business proposition, your idea of where the sweet spots for large language models and their application in law should be. I mean, let's dive deeper into this just incredibly interesting question. And I think it has now, with all the discussion about Claude, co-work, and all it has brought about so much more intense and heated and clearly I understand there are the foundational models, there are certain wrapper solutions. There are then the idea of optimizing the foundation model itself and probably all types of gradients in between and certain combinations with application levels, give us a little bit more granular impression of how all these different pieces relate to one another. And I think we then may go into the question of how does it all compare?</p>
<p>Alexander Kardos Nyheim</p>	<p>So I'll skip over the structural pieces that exist across the whole of AI. So the likes of NVIDIA and the whole computing angle and the cloud providers. And I'll go first to the underlying foundation models themselves. I think as you very clearly already indicated, there are a handful of research companies, of labs in the world that release frontier foundational models. There is an almost circular situation where at any given time, one of these labs has released the best model. Now, what do they mean by the best model? They are thinking about billions of users. There is a recent increased focus on coding and STEM and wider scientific applications, but there's also a renewed focus on what we would call reasoning. So we see a new family of large language models coming out that are called reasoning models, which appear to show they're working if you like, and to apply a level of scientific rigor when answering questions. There is a whole debate as to whether the reasoning that they show is in fact their reasoning or whether it's an illusion. But we don't have to go into that. So we have an ever-widening array of powerful, large language models that are getting better and better at reasoning. And we are also seeing not only the big American companies releasing these, but increasingly some European players like Mistral and Chinese players like Alibaba releasing some quite powerful but light reasoning models in the Quen model family. So there's that layer, which is the underlying foundational model. Now on top of that, you have how that is applied to specific contexts. And that is where I would bring in the likes of Harvey Legora co-counsel in the legal context, where there is a huge effort to find the best models for specific tasks. And then those models are optimized through the best prompting, through being shown in a live context, very high-quality legal data. And then they are guided through these systems to answer legal questions well, to do legal tasks well. And we are familiar with the likes of Harvey Legora with good user interfaces and user experiences designed for certain practice groups. So there's the underlying foundation model layer. And then there is a very intensive and not to be overlooked layer of applying those models to very specific legal tasks. I wouldn't dismiss the likes of Harvey and Legora as wrapper solutions, though there are many wrapper solutions that simply take an off-the-shelf</p>

	<p>model and do very little with it other than put a few buttons on it and then deploy it as a so-called legal tech tool. So there are some quite sophisticated applications of these models within these legal workflows through the likes of Harvey Legora, Co-Counsel, and others. I would say at a very high level, that's how I see the models fitting in with the applications. There's much more granularity as well.</p>
<p>Leo von Gerlach</p>	<p>All right, I understand. And from that starting point, where would you position the own solution that you are now continuing to develop with Thomson Reuters? On the foundational model side, on the optimization side, similar to what we've seen as with Harvey and Legora, somewhere in between, conceptually different. Just give us a sense of what you are pursuing.</p>
<p>Alexander Kardos Nyheim</p>	<p>So if we imagine this pyramid where we have the underlying foundational models, and then above those sit the tools that use them, we are operating at the foundational model level. Now, we are not necessarily developing our own models from the ground up. So we are not necessarily training them on the basics of the English language and putting them through pre-training. But what we are doing is taking very powerful, off-the-shelf open-source models, and then we are putting them through an intensive training pipeline, which establishes a delta, a material improvement on top of those latest reasoning models in our case. That means that we effectively piggyback off very powerful models and add our algorithmic innovations and training data, which can lead to very significant improvements of 10 to 15%, or in lawyer terms, something that is noticeably much better at the law than simply an untrained general foundational model. So our competitors are Anthropic, OpenAI, Google DeepMind, not Harvey Legora, as far as I'm concerned. In fact, I've often been an advocate for Harvey and Legora being potential customers in the distant future rather than competitors as far as my group is concerned, because we're providing that underlying motor, that specialized legal large language model. Now, the question is, why not simply wait for the latest models to come out? And the models are improving. We have new iterations coming from the best and best funded labs every few months. Why not just take those as they are, show them legal data in maybe a RAG setting, optimize them through good prompt engineering, and then allow them to serve lawyers? Why fine-tune and train these models further on the law? I think the existing model developers do not have a sufficient focus on the skills that a model needs to answer legal questions very well. There is an effort to prioritize, I would say, fluency over authority in a way. These models are designed to produce very plausible answers, but not answers that necessarily are anchored in authoritative sources. When answering a legal question, there is much more at play than simply giving an answer that makes logical coherency. You need to effectively follow a rule book that we as lawyers are taught in law school, but which these large language models do not do on their own. And there are certain foundational legal principles and legal ways of approaching questions that if a model is not taught, it is going to struggle to answer a legal question, I think, in a satisfactory way. And this is reflected in how these big model developers train the models, that they are training them for</p>

	different tasks, coding ability, broad knowledge, conversational engagement, helpfulness, rather than being conservative during uncertainty, rather than structuring reasoning to sources and tying it to sources. So these are pretty fundamental differences in how lawyers approach questions as opposed to many other industries. And therefore, we have a training focus that prioritizes the very unique context that lawyers have to operate in.
Leo von Gerlach	Very briefly, is my understanding correct that the actual starting point is to do another layer of pre-training of the model as opposed to a fine-tuning and prompt engineering and all that you have at the later back end, just to get your hands deeper into the model itself to shape it earlier so that it can be better attuned to just legal stuff?
Alexander Kardos Nyheim	Exactly. While we do fine-tuning as well, and fine-tuning can mean a lot of things, but we do, I think, the most intensive form of fine-tuning. We have a system of continuous pre-training that shows the model the universe of Thomson Reuters Westlaw practical law data. And we find that quite fundamentally improves not just the general knowledge, but the fundamental approach a model takes to answering legal questions. Not just, by the way, for English law problems. We have found very interesting transferability of legal skills across linguistic and jurisdictional barriers.
Leo von Gerlach	Very interesting and taken that as face value, what you're saying, would you think that the foundational models developer will also go into this direction that on top of their core models will branch out more sophisticated models like you do now with Thompson LLM, let's say ChatGPT Legal LLM or something like that with Claude or with Gemini or whatever you may have and with medical applications in mind or with other scientific applications in mind, is that a vertical integration that you see down the line or how do you think the market will develop?
Alexander Kardos Nyheim	In short, yes, I think we are going to see the large model developers moving very decisively into law. Law began as an interesting testing ground for the latest models because the model providers knew that law is such a demanding context and problem set that it was a good way of them evaluating whether their models are improving on their core reasoning abilities. So we started with the law being a testing ground for the latest model providers. We have now seen the law becoming a very clear commercial focus of these providers, but not yet in the intense legal-specific model training way you have described. So far, we have seen Microsoft, for example, acquire a part of Robin AI. We have seen OpenAI and Anthropic release effectively plugins, where their general models are optimized at a prompt engineering level to answer certain legal questions well. And we saw the fluctuations in the stock market that it caused. And we have also seen the likes of Google, DeepMind doing model training, custom model training with some large law firms. So yes, the intention is there and it's already in motion. I think the next logical step will be for one of the big players to release a model that they package as a custom trained general purpose

	<p>legal model. I think that is over the horizon. It will be interesting to see who does it because OpenAI has the Harvey dilemma. It is a main backer of Harvey and it does not necessarily want to cannibalize Harvey's work. But it's inevitable, I think. And I think it's good news. It means that law is getting more attention from a training perspective, and lawyers ultimately will have access to more reliable technology. Though from my perspective, it means more intense competition.</p>
Leo von Gerlach	<p>Alex, that has been most fascinating. Thank you so much. That was a terrific conversation. Absolutely insightful.</p>
Alexander Kardos Nyheim	<p>Thank you very much, Leo.</p>
Leo von Gerlach	<p>And thank you everybody for listening. And please join us again for the next edition of The Influencers that will be coming up soon. Until then, take care. Goodbye.</p>