Johner Institute

The great LLM Dream of RA

With Dr. Jochen Jäger, Prof. Dr. Christian Johner

Audio File: REG-ANFORDERUNGEN.MP3

Transcript

00:00:00 Speaker 1

But if there are good documents, then it is quite feasible to extract requirements from good regulatory documents.

00:00:06 Speaker 1

You can build up a database, I can map the products, I can do the Medical Device Insights, a podcast by the Johner Institute for medical device manufacturers, authorities and notified bodies.

00:00:22 Speaker 2

We are currently in a phase where, on the one hand, more and more regulation is being bombarded.

00:00:28 Speaker 2

I had already reported in previous podcast episodes that we are currently monitoring 10000 regulatory documents and have to make around 3000 or more than 3000 regulatory changes per year.

00:00:41 Speaker 2

And of course, this can hardly be achieved by a normal regulatory affairs manager.

00:00:48 Speaker 2

At the same time, we have

00:00:50 Speaker 2

the A.I., especially these L.L.M.s, who are now somehow offering to support us in exactly such tasks.

00:00:57 Speaker 2

And accordingly, several manufacturers are now in the process of trying it out, finding these regulatory requirements, detecting changes, finding out what this means in concrete terms for the individual pro-

duct or for the individual process in Q.M.

00:01:12 Speaker 2

system.

00:01:13 Speaker 2

And

00:01:14 Speaker 2

One company that is also doing this is the company Roche and today I invited Doctor Jochen Jäger and we now want to report in this podcast episode how far we have come with this work.

00:01:26 Speaker 2

So, we don't claim to have the finished solution yet, but perhaps it will help if we report on how far we have come, where we are perhaps still stagnating, i.e. where there are still problems and how we are perhaps confident that we will solve these problems after all.

00:01:41 Speaker 2

Yes, and that brings us to you, Jochen.

00:01:43 Speaker 2

I think it would be.

00:01:44 Speaker 2

ideal for the higher ones, if you would describe your role very briefly and then maybe also, what does it have to do with regulatory requirements and maybe you can now report on it without revealing secrets, of course, how do you usually go about this investigation and what does that mean in terms of effort?

00:02:04 Speaker 1

Yes, I am very happy to do so, thank you very much for the invitation, Christian.

00:02:07 Speaker 1

Jochen Ege is my name, I am at Roche Diagnostik Rotkreuz in Switzerland, now for a good 17 years in various roles,

00:02:13 Speaker 1

was in R&D, was in project management, digitalization, quality and regulatory and am currently in the role of Product Quality Manager.

00:02:24 Speaker 1

In the end, this means that we have to somehow ensure that the product quality is maintained or improved.

00:02:32 Speaker 1

And one of the big influences of product quality is of course external external regulations and that's where it comes from.

00:02:40 Speaker 1

How do we know now when regulations will change?

00:02:43 Speaker 1

or regulatory requirements and how do we get them into our products as quickly as possible.

00:02:47 Speaker 1

Of course, we have a process for this in a large company and the process is as follows: We first have to know what has changed at all or what is currently changing.

00:02:59 Speaker 1

That's what we call monitoring, where we just go in and see which regulations change globally and locally.

00:03:07 Speaker 1

If we now know that a regulation has changed, then the first thing comes to us, is a triage, where we look, does this change affect us?

00:03:14 Speaker 1

So let's do a screening first.

00:03:18 Speaker 1

Is it at all possible for the I.B.D.

00:03:20 Speaker 1

or M.D.

00:03:20 Speaker 1

But it is still possible that there is no product at all that is now affected by the regulation.

00:03:26 Speaker 1

If it affects us and this is already the first clarification, then it goes to the next stage, which is then the more detailed assessment, what are the changes,

00:03:34 Speaker 1

how big are the changes, by when do we have to do what and who needs to be evaluated.

00:03:38 Speaker 1

Now that I know more precisely which product portfolio is, then it's time to clarify in detail what exactly are the requirements that are changing.

00:03:46 Speaker 1

I need a requirements engineer from the projects, who then goes in and looks, I have to adjust my requirements.

00:03:52 Speaker 1

And you can already hear it, it's a longer process, but a more elaborate process that runs through several stages and runs through several people and therefore also relatively much

00:04:03 Speaker 1

time expenditure and also costs.

00:04:05 Speaker 2

The desire to automate this is correspondingly great.

00:04:08 Speaker 2

We are also working together there.

00:04:11 Speaker 2

The ultimate dream would be to have a large global database in which we have all the global requirements in atomic form and where we can query piecemeal or product-specific which of these requirements now apply to this product, to its product type, to this category, to these technologies

00:04:32 Speaker 2

or, if applicable, for the processes associated with it.

00:04:35 Speaker 2

And if you think about how you can approach something like that, then it's essentially 3 steps first, then we might have to zoom in a bit, each.

00:04:43 Speaker 2

So the first step would be to extract all the requirements, so that we have these individual atoms or someone once called it, molecules, requirement molecules, that we have them individually and then also store them,

00:04:58 Speaker 2

then we would need a similarity search, because many regulatory requirements have very overlapping, sometimes even identical requirements.

00:05:07 Speaker 2

And then we would need an integration into a product, such as our Regulatory Radar, where you can then ask very specifically for products, for technologies, what do I have to pay attention to here.

00:05:19 Speaker 2

So those would be the big steps, extraction, transfer to a database.

00:05:23 Speaker 2

So we already know, this is a vector database

00:05:25 Speaker 2

and then the integration into the product.

00:05:28 Speaker 2

And if we now look at this first step, namely the extraction, then you can break it down a notch.

00:05:34 Speaker 2

I now see several steps again.

00:05:37 Speaker 2

The first step is that we have the goal in the first place.

00:05:40 Speaker 2

And this goal will mean that I am able to extract atomic requirements.

00:05:47 Speaker 2

And that brings us to one of the problems later.

00:05:50 Speaker 2

How do we measure that?

00:05:52 Speaker 2

Well, I might outline the problem.

00:05:56 Speaker 2

We could say we determine the percentage of regulatory requirements that are recognized.

00:06:03 Speaker 2

That would be a specificity.

00:06:05 Speaker 2

But if you now optimize for that, then we would also have a lot of false positives.

00:06:10 Speaker 2

So the next thought might be, ah, then let's go to the accuracy, so to speak, to the sum of the identified.

00:06:19 Speaker 2

regulatory and non-regulatory requirements.

00:06:21 Speaker 2

In other words, that you look at it again in total.

00:06:23 Speaker 2

However, this also has certain disadvantages and what we have also found is that we would first have to know and agree on when a text in a regulatory document is a requirement at all.

00:06:35 Speaker 2

Jochen, maybe you can say more about that right away.

00:06:39 Speaker 2

I'll go through the process very briefly.

00:06:40 Speaker 2

I had said at this top level it was extracting, saving the vector database, integration, product.

00:06:46 Speaker 2

And in the first step, I said, first set goals and determine them with metrics, because we also want to be able to evaluate the progress of our experiments.

00:06:55 Speaker 2

Then in the first step, the second sub-step would be to think about how we can do it.

00:07:00 Speaker 2

So the strategy, so we can do it via prompting, do we need advanced context engineering, do we need structured data ontologies, which we need for this.

00:07:09 Speaker 2

That would be the question of strategy, so to speak.

00:07:11 Speaker 2

And then we would collect data in the third sub-step, for example, build up these ontologies, and then actually only in the fourth step, namely start with the actual experiments.

00:07:22 Speaker 2

Yes, before it can go into production later.

00:07:25 Speaker 2

And we are now exactly in this first sub-step.

00:07:28 Speaker 2

So, we have given a lot of thought, first of all, what is the goal, how, how do we notice, how do we measure that we can determine this goal.

00:07:35 Speaker 2

We have also worked with ontologies and are in the process of experimenting.

00:07:40 Speaker 2

and Jochen, if you are perhaps very short or not short, as long as you need, would describe, what were the difficulties we have now encountered, what are we struggling with right now and maybe also, what could be possible solutions?

00:07:54 Speaker 1

Yes, I like to do it.

00:07:55 Speaker 1

So of course you have to explain that regulatory documents are generally written by people, so they are texts and texts are

00:08:06 Speaker 1

As always, interpretable, I'll say now.

00:08:10 Speaker 1

And the difficult thing is that if I now read the same text twice, I get the same result.

00:08:17 Speaker 1

Since we both tried independently with GPTs to automatically extract requirements from regulatory documents, if I do this twice, I come up with different statements with a GPT.

00:08:30 Speaker 1

Then you think, well, the GPT, maybe it's hallucinating or it's not deterministic, that's clear.

00:08:35 Speaker 1

Then we did that and said, we'll do it ourselves and we'll go through only two or three pages of the text and try to read the text manually and determine what the requirements are.

00:08:48 Speaker 1

I can do that.

00:08:50 Speaker 1

If I do this twice as a human being and I read the two pages, then maybe even with a text, well, is that really a requirement now and how deep is it?

00:08:58 Speaker 1

And you're not quite sure yourself.

00:09:00 Speaker 1

So there you have a certain deviation, which is now called intrapersonal deviation.

00:09:04 Speaker 1

not be clear with oneself what the text means exactly.

00:09:09 Speaker 1

So, you understand the text, but is this a requirement that I have to pass on to a product?

00:09:15 Speaker 1

Is it an optional one?

00:09:16 Speaker 1

Is it something that is a superordinate one anyway, or is it actually poorly written and it's just a remark.

00:09:23 Speaker 1

And if we do that.

00:09:25 Speaker 1

now with several people, then of course you can also say, if I've done the text 34 times now, I now have a consensus, these are my requirements and I'll compare it with another person, that's what we both did, Christian and I, then I still see that there are certain deviations.

00:09:40 Speaker 1

Then I can see that some people may have a tendency with a requirement, we can coordinate and have always managed to do that and said, yes, we can convince ourselves that this might be a requirement, but that shows how difficult it is.

00:09:54 Speaker 1

And from that we said, out of the problem, we have to somehow grasp it better, record it mathematically and define a metric in order to recognize at all, what is a requirement and how good is the evaluation, whether all the requirements have been recognized are in a text.

00:10:13 Speaker 1

And only when I have that, we called that the ground truth or the general goodness of a document, if I knew that.

00:10:22 Speaker 1

Then I can now use the A.I.

00:10:24 Speaker 1

or leave several A.I.s on and can say, try to extract the requirements and then I can use the A.I.

00:10:29 Speaker 1

and say that this is right or wrong.

00:10:32 Speaker 1

But it is extremely dependent on which regulatory documents I use.

00:10:38 Speaker 1

Some are very well written, since it is relatively easy to extract requirements and others are written very prosaically and a bit complicated and a bit

00:10:47 Speaker 1

the elaborately written, where it is not quite clear, it is now a requirement, does he want to cover 3 with it, is it more of an explanation or is it something else.

00:10:56 Speaker 1

So the problem is, to sum it up, it is not so clear to identify requirements in text documents and that independently of persona or K.

00:11:10 Speaker 1

I.

00:11:11 Speaker 1

or other things.

00:11:12 Speaker 1

And to that we said, we have to see if we can

00:11:15 Speaker 1

would be able to do something about metrics, about definitions, in order to better define what is a requirement anyway?

00:11:22 Speaker 2

Would you very briefly outline the metric we worked with, with which we hope to measure the quality once from us, of course, but also afterwards from K.

00:11:31 Speaker 2

I.

00:11:32 Speaker 2

.

00:11:33 Speaker 1

Yes, so we said, for example, that we first have to identify what are the smallest units, we called the chunks or

00:11:41 Speaker 1

in this case now sections, text sections, that could be less than a sentence.

00:11:44 Speaker 1

Then we said, if we have identified the chunks, then the metric is, now this is a requirement, if he has it right, then we have awarded certain points, for example 3 points for the correct requirement, minus 3 points if this was detected incorrectly.

00:11:57 Speaker 1

Then we said there might be some where it could be optional.

00:12:01 Speaker 1

So that's not a shell, that's a can and that would be an optional requirement, which isn't wrong if he recognizes it, but also not

00:12:10 Speaker 1

absolutely must be.

00:12:12 Speaker 1

There are fewer points if he recognizes it correctly or even if he recognizes it incorrectly.

00:12:16 Speaker 1

And then there are some that may only be prose, i.e. only explanations of it.

00:12:20 Speaker 1

He should not recognize this as a requirement.

00:12:22 Speaker 1

And then there are examples.

00:12:24 Speaker 1

We just provided them with points and then said, this is our metric and we then calculated the total points via the document after we agreed on what exactly falls into which category.

00:12:37 Speaker 2

Although that was also an iterative process.

00:12:38 Speaker 2

Well, when we noticed that there were both these intrapersonal and interpersonal deviations, yes completely without A.I., we thought about how we could determine that and then we also noticed in the metric that we came to different scores.

00:12:52 Speaker 2

In other words, maximum scores, so to speak, which shows how difficult this topic is.

00:12:56 Speaker 2

And the starting point of this problem, however, is, as you have already said, the sometimes more than questionable quality of some regulatory documents.

00:13:05 Speaker 2

So there are

00:13:06 Speaker 2

2 types of difficulties.

00:13:08 Speaker 2

One is the man-made difficulties and you have now formulated it very diplomatically, but sometimes it is just a babble what is in it and you ask yourself, why does it have to be in it, the person is explaining this topic to himself.

00:13:23 Speaker 2

But that's not a requirement and it's not an example, it's just drivel.

00:13:28 Speaker 2

That means we really have a signal noise issue that we have with us.

00:13:32 Speaker 2

So there's the reason to

00:13:35 Speaker 2

lack of mental acuity on the part of the authors.

00:13:38 Speaker 2

But now there is a second problem that we have encountered, is that if we have individual chunks, it is not just a question of whether we have a regulatory requirement here or not, but now there are other points as well.

00:13:52 Speaker 2

1 is that we have, for example, texts where bulleted lists follow.

00:13:57 Speaker 2

This means that we have to form a kind of cross-product between the

00:14:02 Speaker 2

of the requirements and their expression in the bullet list.

00:14:05 Speaker 2

Yes, so that means that the complete requirement is perhaps the sentence at the front.

00:14:08 Speaker 2

So now an example would be, the manufacturer has to pay for the I.

00:14:12 Speaker 2

T.

00:14:12 Speaker 2

Safety pay attention to and then or must check the I.

00:14:16 Speaker 2

T.

00:14:16 Speaker 2

Security regarding the following aspects and then comes just, what do I know, authorization, authentication and so on.

00:14:22 Speaker 2

That means that a requirement would now be that the manufacturer must, i.e. an atomic requirement, that is, the manufacturer must.

00:14:28 Speaker 2

which was carried out by I.

00:14:29 Speaker 2

T.

00:14:29 Speaker 2

Security, check the authorization mechanisms, for example.

00:14:33 Speaker 2

This means that we can't always take the sentences 1 to 1, but we actually have to form the individual molecules from them first.

00:14:41 Speaker 2

That's an issue.

00:14:42 Speaker 2

And the second thing that cannot be blamed on the authors is that we have a hierarchy of requirements.

00:14:48 Speaker 2

So they are not all on the same level of granularity and concreteness.

00:14:54 Speaker 2

This means that we must also

00:14:56 Speaker 2

ultimately end up with actually according to hierarchy of requirements.

00:15:00 Speaker 2

This means that the result of this database will not be a flat list of requirements, but it will be a multi-level hierarchy of requirements that we can work with later and where you will have to query this database at different hierarchical levels, depending on the use case.

00:15:16 Speaker 2

So these are the ones, it almost goes in the direction of the solution,

00:15:22 Speaker 2

Well, you have already described another solution.

00:15:25 Speaker 2

That was this junking, which is necessary, because otherwise it wouldn't work at all.

00:15:30 Speaker 2

Unfortunately, it is not the case that the authors make demands sentence by sentence, but sometimes they need several sentences until they get to the point.

00:15:39 Speaker 2

And sometimes they mix 2 completely independent requirements into one sentence and think they can combine that by one and easily.

00:15:46 Speaker 2

So that means what we have to do, the A.I.

00:15:48 Speaker 2

to determine the correct chunks here.

00:15:51 Speaker 2

So, we're on it and that seems to help.

00:15:53 Speaker 2

The second is this dissolution of these cross products that I have just described.

00:15:59 Speaker 2

The third approach we have is no longer to assume that we end with a list of requirements, but we are already working with hierarchies.

00:16:09 Speaker 2

and a fourth solution point that we're seeing that is important is working with structured data.

00:16:15 Speaker 2

So, for example, you know, well, we have also modeled this, what requirements aspects do we have in a certain domain?

00:16:24 Speaker 2

I just had the I.

00:16:25 Speaker 2

T.

00:16:26 Speaker 2

Security, you don't have to completely change the K.

00:16:28 Speaker 2

I.

00:16:29 Speaker 2

but you can give it away.

00:16:31 Speaker 2

You can't tell the K.

00:16:32 Speaker 2

I.

00:16:33 Speaker 2

also provide structured data in the sense that

00:16:36 Speaker 2

that it learns more easily, is that a process or a product requirement?

00:16:39 Speaker 2

You can't tell the A.I.

00:16:40 Speaker 2

about structured data, is that more of a target or a must requirement?

00:16:45 Speaker 2

Well, we now have structures that tend to comprise hundreds of rows and which of course also help to minimize the variability in the output or, in other words, to increase reproducibility.

00:17:01 Speaker 2

And with that, we might almost get back to the beginning.

00:17:05 Speaker 2

what you just said, Jochen, the accusation that the A.I.

00:17:08 Speaker 2

likes to do, she hallucinates.

00:17:10 Speaker 2

But we have just established that we are at least as hallucinated human beings and as long as we ourselves do not know what the goal is, it makes no sense at all to talk about the goodness of the A.I.

00:17:22 Speaker 2

to decide.

00:17:23 Speaker 2

And maybe a second thought, because this is now slowly appearing in the audits, when an auditor thinks he is demanding that the A.I.

00:17:32 Speaker 2

must somehow deliver fully reproducible results, then he or she would also have to make this demand on people.

00:17:39 Speaker 2

And we really know from our own best experience that we will fail terribly at it.

00:17:45 Speaker 2

So we humans have this variability or non-reproducibility too.

00:17:50 Speaker 2

Jochen, also your thought, sums it up correctly.

00:17:54 Speaker 1

Yes, exactly.

00:17:55 Speaker 1

So it is, it's written by people, you have to guess the intention of the person in part,

00:18:01 Speaker 1

That's exactly the quality of the documents.

00:18:03 Speaker 1

Can I get to the heart of something so that I have to have as little prior knowledge as possible?

00:18:08 Speaker 1

And some authors just don't bring that with them.

00:18:11 Speaker 1

And then you have to, you have to guess, what did they really mean and what was the intention.

00:18:15 Speaker 1

And one of the ways to get out of it would certainly be to make it a little more mathematical.

00:18:21 Speaker 1

You have already said that in one of your earlier lectures.

00:18:24 Speaker 1

That would be the Regulation is Code topic.

00:18:26 Speaker 1

Can I document regulations in any other way than via a free text?

00:18:31 Speaker 1

and there are approaches that might also help, certainly defuse the problem a bit and I don't know whether there are, whether there are already authorities or certain regulatory units that are already working in this area, you certainly know more.

00:18:50 Speaker 2

So E.

00:18:52 Speaker 2

U.

00:18:52 Speaker 2

Commission, than we ever presented a list of projects to improve the regulatory system,

00:18:58 Speaker 2

that was voted number 1.

00:19:00 Speaker 2

So they said, well, we should please tackle this, because we have already shown this with examples.

00:19:05 Speaker 2

You don't start 0 anymore.

00:19:07 Speaker 2

So there is also a lot of research on this and that's why you should start with it.

00:19:12 Speaker 2

Of course, it is always the question afterwards and who is doing it now and for this it is precisely the people who write these legal texts that are needed.

00:19:19 Speaker 2

It needs the authors of the norms and that would perhaps be our wish now to all who listen,

00:19:25 Speaker 2

especially if they collaborate on such documents.

00:19:28 Speaker 2

So try to check their statements to see if they could be formalized.

00:19:35 Speaker 2

And if they are not, then you know, then you have written sentences that are just a bit soft and where you actually have an A.I.

00:19:44 Speaker 2

in order to peel the hard core out of it afterwards.

00:19:46 Speaker 2

And that's actually unnecessary work, because if we had it formally, we could also check it formally and contradictions, gaps and so on.

00:19:54 Speaker 2

by the computer and afterwards not by the poor manufacturers.

00:19:57 Speaker 2

So that would be our wish, so to speak, to those who listen, if there are now companies that also want to get involved, but now not as spectators, but really actively participate, are of course also welcome to get in touch.

00:20:09 Speaker 2

Yes, and Jochen and the 2 of us just continue to work on it and get back to you.

00:20:14 Speaker 1

We continue and see if we can work with the G.P.T.

00:20:18 Speaker 1

maybe some other one after all

00:20:21 Speaker 1

I can take a step forward and I think you can, the G.P.D.s are getting better.

00:20:26 Speaker 1

I don't think you can guess what the intention of all the people was, because it's really guessing, but if

there are good documents, then it's quite feasible to extract requirements from good regulatory documents.

00:20:39 Speaker 1

You can build up a database, I can map my products, can merge them with my product requirements, can then certainly tighten the sum of the requirements for the products and ultimately I can

00:20:51 Speaker 1

faster, develop better products, and that's exactly what we want.

00:20:54 Speaker 2

This is a wonderful final word, Jochen and with that it remains to say thank you very much.

00:20:59 Speaker 2

I'm looking forward to continuing on this journey later, to these experiments and I have the same high confidence, because it's not an A.I.

00:21:08 Speaker 2

Limitation that we get stuck on, but rather the limitation that we get inadequate input and have to interpret it.

00:21:15 Speaker 2

But I think we can and in that sense, Jochen, thank you very much.

00:21:19 Speaker 1

Yes, thank you, Mr. Hau.