

The role of AI reimagined – Part 1

[00:00:00]

BRADLEY HOWARD (BH): My name is Bradley Howard and today I'm hoping to discuss the very interesting topic of artificial intelligence. Joining me are Radu Orghidan, VP Cognitive Computing Strategy at Endava, hello Radu.

RADU ORGHIDAN (RO): Hello Bradley, thanks for having me here.

[00:00:26]

(BH): It's a pleasure. And Boris Cergol, Head of AI at Comtrade Digital Services, which has recently become part of the Endava family, hello Boris.

BORIS CERGOL (BC): Hi Bradley, really happy to be here.

[00:00:35]

(BH): Excellent, and welcome to the show. So, welcome to both of you, let's dive right in. AI has become even more of a buzz word and a hot topic in the last 12 months, why do you think that is? Let's start with Radu.

[00:00:47]

(RO): Well, I think that artificial intelligence is here because of the progress that we've seen for a long time. Imagine we're now picking the fruits of the work that happened during decades, and now we're gathering everything in a form of ABIs or nicely delivered formats. But also, there are not only benefits to be entirely honest. So, briefly, the high-level AI is associated with systems that can see and talk. So, these concepts are easily to understand from science fiction books and movies, but at the deeper level, we have an interesting mix between the concepts of artificial intelligence and cognitive computing, which deserves a special topic on its own.

[00:01:38]

(BH): And Boris, anything else to add on that?

[00:01:40]

(BC): Yeah, I mean I will also underline that AI has been a buzz word since about 1950s, I would say, but what has really been happening in the last year, or the last few years, was that it is kind of starting to deliver on the promises that it made in decades past. And we saw some really, really great achievements this year, especially in the area of language modelling, and applications in healthcare.

[00:02:08]

(BH): So, Radu, you talked about cognitive computing before, so how can businesses distinguish whether they need to be using AI or cognitive computing to solve complex problems, and what's the difference between the two of them?

[00:02:21]

RADU ORGHIDAN (RO): OK, a great question, thanks. So, AI, artificial intelligence is basically looking at how to implement the best algorithms to solve a problem. A classic programme, or a programmed algorithm chip can also be called AI. But cognitive computing goes a step beyond and tries to mimic human intelligence and winds up by analysing some facts extracted from data. At the end of the day, cognitive computing methods aim at augmenting human capabilities through three things; classification, prediction and prescription abilities, that's how I like to see them. And I

really like the definition of IBM, that IBM gives to cognitive computing, which states that these are advanced these terms that learns at scale, reason with purpose and interacts with humans in a natural form, so this is something else.

[00:03:23]

(BH): And have you got any examples of where there might be a difference between cognitive computing and AI?

[00:03:29]

(RO): I can think of something from our recent experience. We could build a solution for knowledge management for one of our customers using NLP and computer vision for classifying documents. And this can be combined with Q&A both that can have a verbal conversation with users looking for specific information in digitised documents. And on top of everything, since the solution has direct access to users typical typing or talking, whatever, it can also perform sentiment analysis, so all these building bricks form a cognitive computing solution but each of those can be seen as a form of artificial intelligence, a basic building block of the whole solution.

[00:04:19]

(BH): And Boris, have you got anything to add on that?

[00:04:20]

(BC): Yeah, I mean Radu already mentioned, you know, solving the problems like classification, prediction and prescription, what I would add is that, you know, in general, artificial intelligence is also the goal to develop systems that would not only augment our abilities as humans in our work, but also to be systems that are actually entirely autonomous. So, systems to which you could basically set a goal to follow and then they would find ways to achieve this goal.

[00:04:55]

(RO): I wanted to add on top of what Boris is saying, that not always businesses need to look at cognitive computing or AI to solve their problems, maybe the best approach would be first to understand what they need, and not try to match the technology with an imaginary need.

[00:05:14]

(BC): Yeah, yeah, I totally agree with that. So, one of the things, the first things to consider for a business when implementing systems such as this, is actually what is the problem that they are actually trying to solve. So, I don't think it's a good approach to maybe just start with an implementation of an AI system, simply because, you know, they were enamoured by media buzz words in this area. Also, the results are the best when we set firm KPIs that the system should deliver and then iteratively optimise the system to achieve the best result.

[00:05:48]

(BH): And I really like that, that you're talking about starting with the business challenge rather than the solution, whether that does or doesn't involve AI or cognitive computing. But do you, when you speak to customers, do you find that sometimes people have unrealistic expectations of AI?

[00:06:07]

(BC): I would say that it really depends on the maturity that the business has in previous implementations of systems such as this. So, I definitely find that businesses who are just starting out with this technology often have unrealistic expectations. So, in a way, it actually makes our

work much easier when we can deal with someone who has already done some steps in this area, and we can help them build on that.

[00:06:36]

(BH): And Radu?

[00:06:37]

(RO): I would also highlight the one phenomenon that I sometimes find with experience the businesspeople that we face confirmation bias, and over confidence bias.

[00:06:50]

(BH): So, it's interesting that you mention about bias, late last year I did a course online about managing machine learning projects with Google Cloud, and one of the first modules on that course is around ethics and some of the challenges that it raises. What do you think are the most important areas of ethics in AI that we need to solve first, Radu?

[00:07:11]

(RO): So, I think bias is one of the very important things in AI in general, and when talking about AI we're mainly considering machine learning. So, machine learning models are those algorithms that learn from data. So, data becomes their textbooks, therefore, machine learning models reflect our own biases, this output may range from relatively trivial too high significant problems. For example, we can have models that don't match our preferences at the restaurant, but they can, in the end, deny a credit or a bail. So, this is why bias is really, really an issue and there are many types of biases.

As I said before, the confirmation bias but then there is overfeeding, when you have too many data, or underfeeding biases, there are the sample biases when you train a model using only day images and it doesn't work during the night. Or you have prejudice or stereotype bias, where you deal with gender classification using a data set of pictures that are not entirely covering the whole range of data. There are systematic valid distortion and many types of biases that, in the end, are reflected in the artificial intelligence solution.

[00:08:43]

(BH): So, an example I've heard you give before Radu, is if you show a machine lots of pictures of black cats, then the machine will think that all cats are black at that point.

[00:08:52]

(RO): Exactly. And following the same line, if you train an autonomous driving system to drive only on streets during the day, it will face obviously, some difficulties in working in the night. So, this is a sample bias.

[00:09:10]

(BH): And Boris, the whole area of job security, which is something that always comes up whenever we talk about AI, how do you feel about that?

[00:09:18]

(BC): Yes, I'm definitely on the side of the argument that this is a very, very real concern, so I don't see AI just simply augmenting people in their work, but I think that there would be a lot of jobs that actually will be replaced, and I also think that it will start happening much, much sooner than many predictions set it out to actually happen. So, I think that when talking about the ethics in AI we should really put a lot of attention to taking care of the societal changes that we have to implement in order to basically reap the benefits of AI systems, or I should say, that these benefits are

actually enjoyed by a large part of the population. Because one of the problems that is currently happening is that, you know, most of the used cases of AI, or at least a large part of them, are simply increasing the imbalance of power between, you know, different individuals and organisations or, you know, individual different socialising and so on.

So, yeah, this is definitely a big concern. What I would also like to add to what Radu has talked about the bias. So, one area where I see there's a lot of critique in applications of AI systems, and of course, it's related to issues with bias. I think that we should separate this discussion in two areas. So, the first area is whether we should use an AI system to solve a certain problem at all. And the second one is the quality of the system that we are using to solve that problem. Because you might actually simply have a system whose quality was not checked, or who's development was not using the best practises that are already established to handle issues such as bias. But often these two get kind of entangled, and for example, an AI system might be critiqued simply because its quality was not good enough, not just that the problem it is solving is one that shouldn't be solved with it.

[00:11:30]

(BH): So, how do we track the accountability of decisions that are made by AI machines?

[00:11:35]

(RO): Well, there is a parallel between the issue of bias and the opacity of AI systems, and as Boris mentioned before, sometimes the system is not critiqued or doesn't perform as expected because we don't understand what happened there. If somebody understands the reasoning behind, they will accept much better the results of the system. So, regarding who has to take the responsibility for the results of a system, there is, of course, a complex discussion because it's not anymore, the operator, like in the case of the driver of a car. But it's the manufacturer and the software behind it. But I think the opacity of AI should be seen as an object, like the transparency of it mainly, to achieve transparent and explainable AI systems.

[00:12:37]

(BC): We're actually fortunate that, you know, by now that there is a whole field called explainable AI, which is trying to solve that problem specifically and it has made great progress in the recent years. However, I do feel like in the general, you know, the general population and even in, maybe in people in businesses aren't really that aware that this thing exists and that, you know, we have methods that are capable of providing explanations for the predictions that are made and that we can get a lot of insight into actually how these models work.

So, I mean definitely this is one area where I do hope that, you know, as much people as possible who are working with AI systems in any way get more aware of it and that we see more implementations of this algorithms in actual production.

[00:13:31]

(BH): Oh, thanks for that. So, where do you think that AI will be in 10 years' time, let's start with you Boris?

[00:13:38]

(BC): Yeah, so I'm quite optimistic in this regard, but also, sometimes a bit scared or worried, so I think that we have reached now a point where we might actually see something akin to a road map emerging, so I wouldn't look at AI just fundamentally changing how we interact with our machines, or how the business processes are optimised by it. But also, for its transformation or ability in, you know, scientific discovery. We've already seen some of the results of this even in last year, but this will be greatly increased.

So, I think it's actually quite problematic to talk about where we will get with it in 10 years, maybe, you know, better to look at what might already be achieved in this year and maybe one to three years ahead of that.

[00:14:39]

(BH): That's a good answer deflection going on there. Radu, do you want to be braver, where do you think that AI will get to in the longer term future?

[00:14:48]

(RO): I think AI will be an intrinsic part of our lives, they will be embedded in our daily activities through different smart objects that will be able to communicate between themselves and augment our capabilities and help us in doing our work better. But there are also, and I'll be here the devil's advocate, usually I'm very optimistic, but we have to take a look also at what are the potential problems. And one of the things is that it operates in unregulated environments, and we have to have better regulations and understand where, how to allocate the responsibility.

Then it operates in markets that have a winner takes all characteristic, and this is the perfect receipt for monopolies, and I think this is another issue that has to be really regulated. And finally, there is the new economy that's developing on top of the digital service industries, which is based on intangible assets, which is also called capitalism without capital. In the end, we cannot be opposed to the progress and this will all become part of our lives.

[00:16:05]

(BC): I'd like to add two points to that, so the first one Radu has mentioned regulation, so I also definitely think that this is something that we have to work on. But, in a way, I also think that it's really not enough to really help get this thing under control. So, I would wish that we would be developing more systems which actually use AI to try and, you know, counteract the negative effects of other AI systems, like to give an example, there is a lot of talk about how, you know, exposing people to information using social network algorithms is causing a whole variety of problems, it will be really good, for example, if filters were developed that would actually follow the objectives that are good for the individual.

And maybe help filter away the unnecessary information that is causing damage. The second notion is to get back to you Bradley a bit about not being very concrete. So, let me just comment on some of the things that I think will happen with the new language models that are being developed. So, I think that we will really see a surge of used cases relative to all sorts of creativity, so basically, systems that will help people drive various things but also, drive things that are quite compelling by themselves. We might see very soon systems that are able to produce a lot of quite useful computer code, we will see systems that help us connect to other people.

So, maybe when we say something to transfer what we are saying to a language that is more understandable to the person that we are trying to convey the information to, and I'm not speaking here just about, you know, whether we are speaking English or German, but actually the complexity of what we're trying to express. So, also, you know, the whole thing with the intelligent agents, what we are seeing now with like systems like Siri and Alexa, it's really just the beginning.

I think that we will have the ability very soon to have really interesting conversations with those agents, I had the opportunity to have some of them myself when being a part of the better testing group for one of those models, and it's really amazing things that are just about to happen in the area.

[00:18:46]

(BH): And Radu, your comment about AI being like electricity, that we'll just use it, and we won't even know that it exists as a separate form. One could argue that every time we do a Google search, we're already using lots of different forms of AI, just in that Google search, you know, whether it's localised to how we use language, or to local businesses and what have you, it's already using lots of AI inside that search.

[00:19:11]

(RO): Correct, and this is just the beginning. Because there are already, maybe not so visible changes, but I'm amazed as a few years ago, if you wanted to build an application you had to set up servers, install everything, do all the build up a team and so on, and now the technology lets you do everything with just a credit card, and an internet connection, through the cloud providers. So, as these things are becoming more and more a commodity, this is how the algorithms and the models of machine learning will also become part of our daily routine. As you said, Google search is an example.

[00:19:58]

(BH): And that probably brings us back to the very first question about AI has become such a hot topic in the last year or so, because it's so accessible for anyone, or any business, to start using it whenever they wish.

So, let's move onto a quick fire round, I'm going to ask a question to each of you, and then if you can give me a short answer to those that would be great. So, let's start with you Boris on each of those questions, so what movie or book portrays AI the most accurately?

[00:20:28]

(BC): Yeah, so I quite like the movie Her, about a man who is developing a kind of a personal, at least in his case, a bit of a romantic attachment to an intelligent agent in his phone, so I think that we will see quite a bit of that in the future. Also, in a way, a similar movie, Frank and Robot, where there is this elderly gentleman who receives a robot to help them from their family, and it shows what sort of interaction they have, so I think these are the ways in which actually AI systems will enter the lives of people in the near future.

[00:21:09]

(BH): And Radu?

[00:21:10]

(RO): Well, I think AI in movies goes toward singularity which is another interestingly related to AI. And I like WALL-E, the movie, the cartoons. And then I, Robot, which are both dealing with the notion of singularity.

[00:21:28]

(BH): For the benefit of our listeners, would you like to describe what singularity is?

[00:21:32]

(RO): Singularity is opposed to narrow AI, is a form of intelligence that knows, understands the context and knows how to solve difficult problems, not just narrow ones like we are doing nowadays. For example, a narrow AI algorithm will solve the problem of recognising a cat or a dog in a picture, but the singularity would be able to act and mimic the human behaviour, the context, the cultural understanding of a situation that a human has.

[00:22:13]

(BH): Thank you. And back to you Boris, what problem would you most like AI to solve?

[00:22:19]

(BC): Oh, if I had to choose personally, I would definitely choose something related to increasing human longevity, so basically applications related to healthcare that would serve that goal. But this is related to simply wanting to actually increase my chance of seeing AI solve all the other, or at least most of the other problems in other areas.

[00:22:41]

(BH): That's a really interesting answer, isn't it, that if we had an AI machine to make some recommendations, how different might those recommendations be to the advice we currently have, which some of us choose to ignore. But anyway, that's probably another podcast all in itself. And Radu, what problem would you most like AI to solve?

[00:22:58]

(RO): Well, I think I am also biased by the latest success of Google Mind with AlphaFold, and I think I would like to see it pushed forward and see how we can solve the health issues and maybe attain immortality with the creation of organs or seeing how we can improve health problems.

[00:23:24]

(BC): Maybe I just add one thing to that when we are discussing immortality and singularity, and topics like that. I really wish we found a way to solve the problem of the alignment of AI systems. Because, of course, if we develop systems with superhuman capabilities, we really want to find a way that their goals are actually well aligned with ours, otherwise they may lead to something that's rather opposed to the notion of immortality for us.

[00:23:55]

(BH): Moving on swiftly to the next and the final question in our quick-fire round, Boris, who is the most inspiring person or company in AI today?

[00:24:06]

(BC): Yeah, I think that many people would give the same answer as me, so two companies Open AI and DeepMind, which is owned by Google, I think that the research that they do is really, really, really great. And it's interesting that they're following quite different approaches.

And in terms of people, I would say, maybe Richard Sutton is quite inspiring for me, because he's a person who has a kind of explained his observation that the progress, a lot of the progress that we make in AI actually gets down to increases in computational power, and that it is maybe better to focus on algorithms that are such, that can scale very well when we add additional computational power.

[00:24:57]

Rather than focusing on trying to build something very specifically to incorporate in very specific ways the knowledge that we have.

[00:25:06]

(BH): And Radu, who's the most inspiring person or company in AI for you?

[00:25:10]

(RO): I agree with Boris, I like most Open AI, given the advances they have, also DeepMind and yeah, as persons, there is one guy that I'm following quite often, his name is Lex Friedman, who does a series of podcasts which goes through many discussions, interesting discussions, technical and philosophical.

[00:25:37]

(BH): Thank you both Radu and Boris for joining me for this very interesting discussion on such an exciting topic. And listeners, if you have any more questions about AI for Radu or Boris, please send them to podcast@endava.com.

If you've been wondering about how Boris and Radu ended up in a career focused on AI, and what parts of it inspired them the most, please join us for part two.