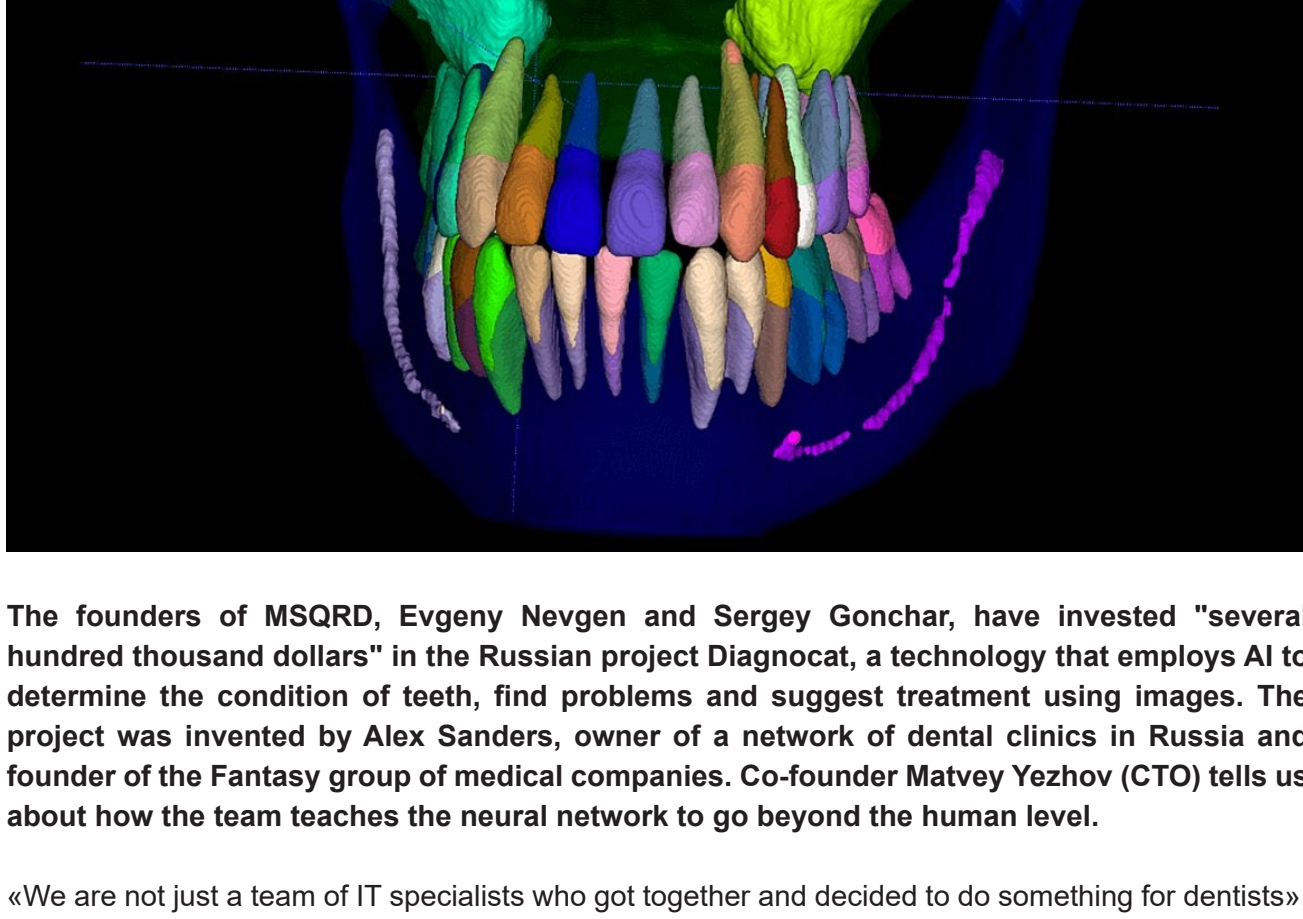


«When you are into everything, you can't build a product that requires a lot of expertise» The CTO of a startup in which the founders of MSQRD have invested talks about «dental» AI with a multi-billion-dollar plan



The founders of MSQRD, Evgeny Nevgen and Sergey Gonchar, have invested "several hundred thousand dollars" in the Russian project **Diagnocat**, a technology that employs AI to determine the condition of teeth, find problems and suggest treatment using images. The project was invented by **Alex Sanders**, owner of a network of dental clinics in Russia and founder of the Fantasy Group of medical companies. Co-founder **Matvey Yezhov** (CTO) tells us about how the team teaches the neural network to go beyond the human level.

«We are not just a team of IT specialists who got together and decided to do something for dentists», says Matvey Yezhov. «We have a dentist, our founder Alex Sanders, who realized that dentistry lacks expertise in the rapid analysis of images, in visualization. And then he gathered a team of IT specialists. He has a network of dental clinics in Moscow, and he himself practiced as a dentist. Our second expert in the team is Evgeny Shumilov, also a dentist. For a long time, he was selling CT scanners. Due to their contacts, we have a huge amount of external expertise: radiologists, consultants, including from other countries.»

The Diagnocat technical team employs six engineers, among them one CTO generalist — that's Matvey. One full stack developer, three researchers and a data engineer. Plus the team of managers and data labelers.

«Not to just save a few minutes of the dentist's routine work, but to do something fundamentally new»

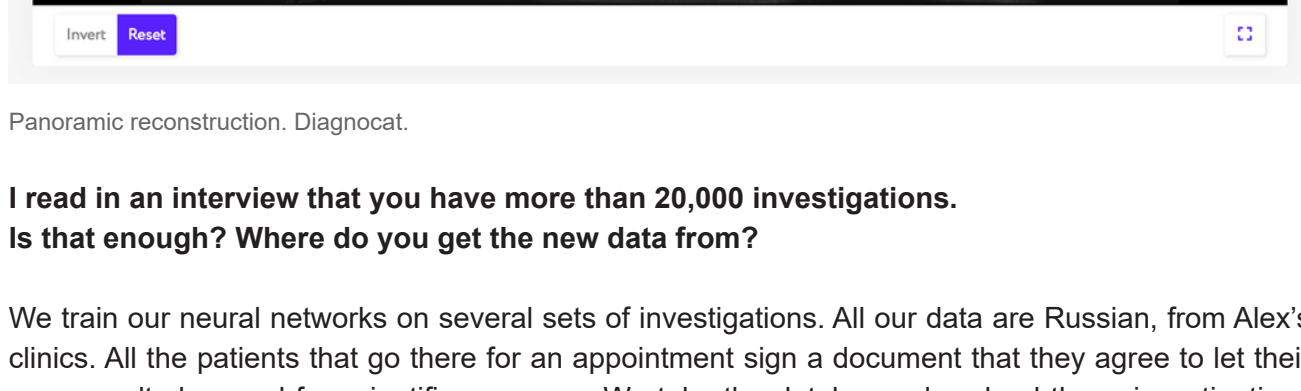
What are you working with today?

The essential images in dentistry are a 3D scan of the entire maxillofacial region, a panoramic 2D scan and targeted images where only a few specific teeth are visible on the scan.

We work with a 3D image. It's a three-dimensional scan; you can't just pick it up and look at it. The information contained inside the voxels (the volumetric equivalent of pixels—ed.) is not visible from the outside, and you need to "crawl" inside the image, slice it. Such a procedure is very time consuming and requires special training and software. There are very few specialists who can do this; there is one radiologist to every thousand dentists. Our technology will take a lot of work off them and allow them to concentrate on really difficult and interesting cases.

The Diagnocat system consists of four separate neural networks that work in two areas: localization:

- finding anatomical structures inside the skull-teeth, bones, canals, sinuses, etc.,
- diagnostics: determining what is wrong with a tooth, canal, bone: caries, periodontitis, other pathologies, and how to treat it.



Panoramic reconstruction. Diagnocat.

I read in an interview that you have more than 20,000 investigations. Is that enough? Where do you get the new data from?

We train our neural networks on several sets of investigations. All our data are Russian, from Alex's clinics. All the patients that go there for an appointment sign a document that they agree to let their scan results be used for scientific purposes. We take the database, download these investigations and depersonalize them. After that, we tag it, that is, we assign certain tags to voxels. This is a tooth, this is caries.

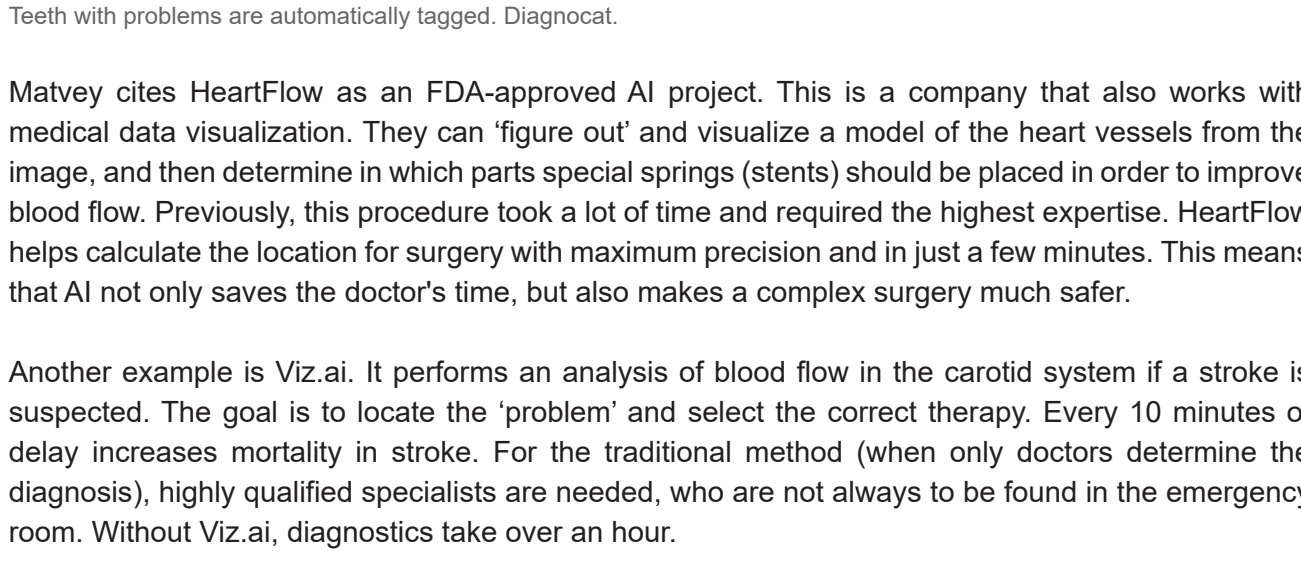
20,000 (and now much more) is the pool of unlabeled investigations. That is, these are the results of 3D images without description, without segmentation, without diagnoses. We are now constantly pulling out a new investigation from this pool and tagging it in different ways.

We have about seven thousand tagged investigations. We do different tagging levels. Some are tagged superficially; for example, we know that such and such a tooth has caries and that's it. Others are tagged more complexly: each voxel inside the scan is assigned a specific tag. There may be various numbers of these tags, and each will indicate some anatomical facts.

«If such a startup brings money in Russia at the seed stage, then investors would value it at \$1.5 million, and in the USA at \$10-12 million»

«Tagging is something we constantly wrack our brains over: how to tag it, for what purpose, and how to train and organize the people who do it. The goal is to extract as much relevant information for the dentist as possible from the investigation and pack it in a format that will be most useful and convenient for him and will help him communicate with patients. It's difficult to say how much data we need; there is no endpoint. As soon as we finish tagging what we are doing now, we will simply start tagging something else», says Matvey.

Diagnocat has already begun patenting the technology in the US. The next big step is FDA (Food and Drug Administration) approval.



Teeth with problems are automatically tagged. Diagnocat.

Matvey cites HeartFlow as an FDA-approved AI project. This is a company that also works with medical data visualization. They can 'figure out' and visualize a model of the heart vessels from the image, and then determine in which parts special springs (stents) should be placed in order to improve blood flow. Previously, this procedure took a lot of time and required the highest expertise. HeartFlow helps calculate the location for surgery with maximum precision and in just a few minutes. This means that AI not only saves the doctor's time, but also makes a complex surgery much safer.

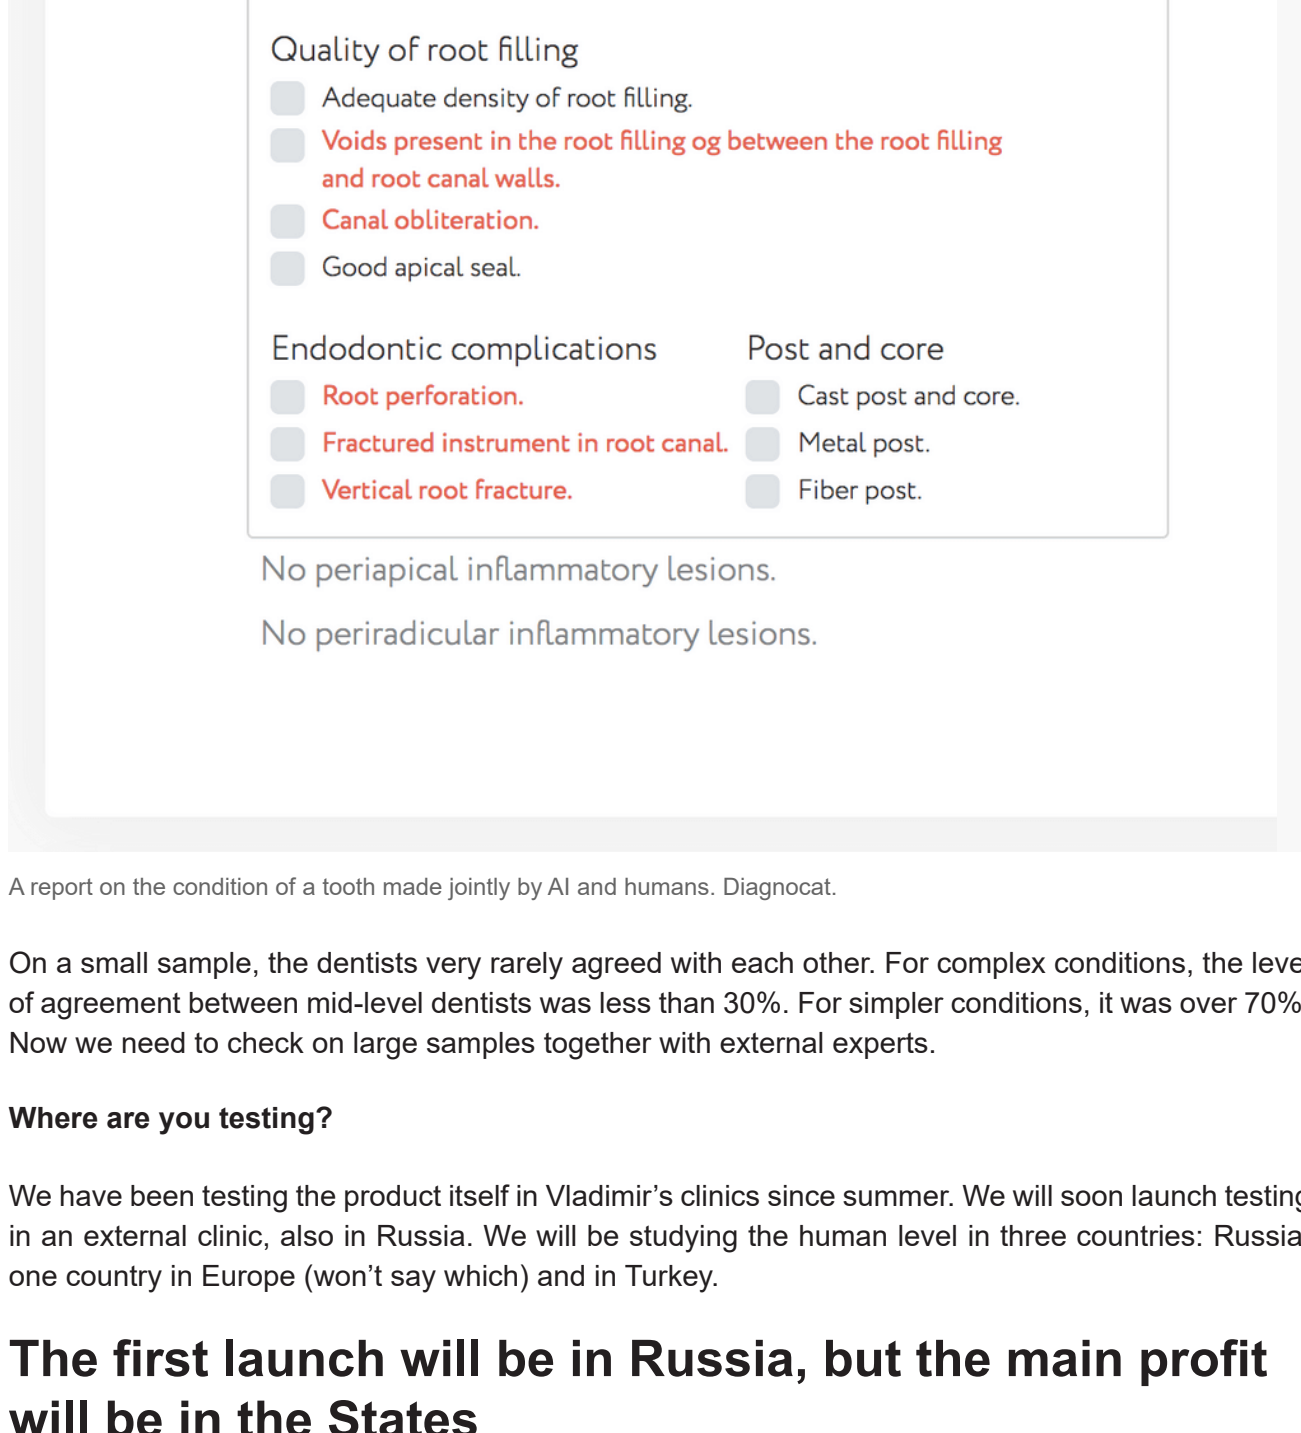
Another example is Viz.ai. It performs an analysis of blood flow in the carotid system if a stroke is suspected. The goal is to locate the 'problem' and select the correct therapy. Every 10 minutes of delay increases mortality in stroke. For the traditional method (when only doctors determine the diagnosis), highly qualified specialists are needed, who are not always to be found in the emergency room. Without Viz.ai, diagnostics take over an hour.

«We at Diagnocat are also striving to apply technology in this way. Not just to save a few minutes of the dentist's routine work, but to do something fundamentally new which was previously impossible. For example, before us, no one could divide and visualize (segment) all anatomical regions in the head—jaws, teeth, canals, sinuses—from an image and build exact models of them. We are going to do that. Diagnocat is not just about teeth. We strive to identify all the elements that are involved in dental operations», says Matvey.

Better than a human

We are working to achieve results better than the human level. With regard to localization, we are already overtaking humans. Diagnosis is more difficult. There are purely technical difficulties. For example, there may be a dark spot in the image, and it is not always possible to clearly determine whether this is caries or simply shading from the instrument. There are shortcomings in the investigations; so far we diagnose rare cases much worse than humans can do it, as it is difficult to assemble a representative data set. Our consultants have compiled a list of several dozen different conditions that are associated with cancer, odd formations and developmental anomalies. Here we are learning.

But that's not all. It is not enough just to see the condition; you need to prove it. This is separate work. We are now starting a large study on measuring the human level in different countries. It goes like this: take dentists, give them a diagnostic task and measure how much they agree with each other's results and with our results. If dentists agree more often with each other than the algorithm does with the dentists, we are worse than a human at determining the diagnosis. If it's the other way around, we are better.



A report on the condition of a tooth made jointly by AI and humans. Diagnocat.

On a small sample, the dentists very rarely agreed with each other. For complex conditions, the level of agreement between mid-level dentists was less than 30%. For simpler conditions, it was over 70%. Now we need to check on large samples together with external experts.

Where are you testing?

We have been testing the product itself in Vladimir's clinics since summer. We will soon launch testing in an external clinic, also in Russia. We will be studying the human level in three countries: Russia, one country in Europe (won't say which) and in Turkey.

The first launch will be in Russia, but the main profit will be in the States

How are Evgeny Nevgen and Sergey Gonchar participating in the project?

Gonchar helps a lot with technical expertise in 3D and in the engineering part in general. Nevgen helps with the product, including design. And with investor relations; he helps bring them in. They both also help find techies for the team. They were the ones who brought me in.

MSQRD was developed in Belarus. There are benefits here, personnel, it's an IT country. Have you thought about moving Diagnocat here?

No, we're not planning to. But we may open an office here. We have several big specialists from Belarus who consult with us; perhaps then they will join the project and we will make a "department" for them in Minsk. But, generally speaking, we are a distributed company—we are not tied to a place.

Are you expected on the US market? What does market research show?

It shows that we are of interest in the States. The heads of clinics and our consultants with whom we communicate—everyone says they want such a product.

Will you be launching in the CIS, or is that not profitable?

Definitely. The first launch will be in Russia. Then we'll go to the States. As for profit... All products launched in more than one country have a personalized pricing policy. In Russia and the CIS, the price may be lower than in other regions. Not only because the standard of living is lower, but also because people here simply don't want to pay. Our product will be sold in Russia at reduced prices, while in the States, even now the cost of one X-ray report (from a dentist) is \$70-150 per investigation.

How are dentists responding to the technology? There's no rejection of the idea that some kind of technology should replace their expertise?

As a rule, very big specialists who see shortcomings in the program say that. But they don't sharply criticize either. They say: "Yes, this program is not for us, I wouldn't use it, but for a less experienced dentist it will be very useful."

With regard to «replacing their expertise», I said before that radiologists are happy because they can focus on important and interesting cases. And dentists are happy, because their main work is still manual, but it would be nice to find help in diagnostics. So, no matter what, it works for everyone.

All in all, the reaction to the product is not as harsh as we thought it would be (laughs).

What competitors do you have?

There are the guys from Ocutri (graduates of the Faculty of Computational Mathematics and Cybernetics of Moscow State University). They are compared to us, but there is little information about them. Judging by the site, they are into everything. And when you are into everything, you have no chance of building a product that requires a lot of medical expertise. We have direct competitors in Israel; they are doing the exact same thing. Orca Dental Ai. We visited them, and they were trying to convince us how difficult it was and that we shouldn't be working in this area (laughs).

But all in all, our market is so large and incomprehensible that we don't mind the competition.

In the spring of 2019, you are going to present Diagnocat at the International Dental Show in Cologne. What do you personally expect from the trip?

I'm expecting users, traffic on servers—an operational pain in the neck, that's what I'm expecting (laughs). In fact, it will be a global launch, there will be many users, feedback will flow. There will be results, metrics, it's very cool.

What is the most negative feedback on Diagnocat so far?

There has been no negative feedback so far. When a dentist opens an image, uploads it to Diagnocat, gives our diagnosis without checking (although we always ask them to check) and screws up, then there will be negative feedback. For now, it's too early.

Will the final project also provide a second opinion or can it completely replace the radiologist?

Depends on how good we are. If we show very good performance and prove it to the regulators, we can make an autonomous version. We have a plan for more than ten years ahead, where we build a multi-billion-dollar company. Of course, I'm not gonna share it with you; it's a cunning plan about how to take over the whole world.



From left to right:

Sergey Gonchar (MSQRD), Alex Sanders (Diagnocat), Evgeny Nevgen (MSQRD) and Matvey Yezhov (Diagnocat).

According to Alex Sanders, over the course of a year and a half, the founders of the project—himself, Evgeny Nevgen and Sergey Gonchar—have invested several hundred thousand dollars of their own funds in Diagnocat:

«The project is to become large-scale, and it is difficult to finance it further with our own funds. We have started looking for investors among venture funds in Europe and the USA. Regarding the amount of required investments, I will say this: if such a startup brings money in at the early (seed) stage in Russia, then investors would value it at \$1.5 million, in Europe at \$5-7 million, and in the USA at \$10-12 million.»