
# INTERVIEW

INTRODUCING

## NATO'S FIRST MACHINE-LEARNING TOOL: AI FELIX

**"If Ray Kurzweil's prediction is correct, before 2030, humans and AI will merge into a super-intelligence and thus reach the next stage of human evolution. Whatever happens: remember that today's AI is the worst you'll ever use!"**

Recently, the Joint Warfare Centre Public Affairs Officer Ms Inci Kucukaksoy interviewed **Mr Simon Purton** and **Dr Arnau Pons** to discuss an exciting and timely addition to NATO's toolbox: AI FELIX. Simon Purton is the Capability Requirements Section Head at Headquarters Supreme Allied Commander Transformation (HQ SACT) in Norfolk, Virginia. He is a mathematician with an avid interest in machine learning. His HQ SACT colleague Dr Arnau Pons is an AI Product Manager and Operations Research Analyst. His background is in aeronautics and astronautics.

**JWC:** Simon, in 2019 you wrote an article for The Three Swords called "Learning about Learning Machines." What has happened since then?

**Purton:** When I wrote the article, I wanted to capture some of the basic concepts about learning machines. Plus, I wanted the reader to understand the difference between how humans make decisions and how decisions may get made otherwise — it's in this aspect that artificial intelligence has some advantages over us. However, I still maintain that while machines have some advantages, that does not make them superior decision-makers compared with people. When I was writing the article, I was also proposing to HQ SACT leadership that they fund an artificial intelligence experiment. I couldn't really get any traction at the time, so I submitted a proposal as part of the ACT Innovation Challenge. The proposal won.

Mr Simon Purton (left) and Dr Arnau Pons. Photo by HQ SACT PAO

**Pons:** Simon describes himself as one of the parents of AI FELIX, which is the name of the bundle of AI applications we have in HQ SACT.

**Purton:** I refer to myself as one of the parents, because there were a few of us involved in bringing AI FELIX into being. A team effort.

**JWC:** You said you couldn't get any traction with your leadership. Why was that?

**Purton:** As we are NATO's Strategic Command for transformation, our leadership's scale of ambition is very large. The proposal was not very large. In fact it was quite modest. I wanted to experiment with AI to automate some of the very important, but very repetitive work that some of our staff have to do here. The proposal wasn't about winning wars or delivering new wonder weapons. It was about fixing some of the knowledge management issues that we all have. And it was focused on something called the Command Read Board.

**JWC:** Can you explain the Command Read Board to our readers who may not be familiar with the term?

**Pons:** Just about every command has some point of entry for incoming correspondence. In HQ SACT, that is called the Command Read Board. It's a function that ingests reports, personal letters, taskings from NATO HQ... anything formal.

**Purton:** We did a podcast on this. I learned that for two people in our command — and across the NATO Command Structure for around 40 people — their daily job was to read incoming correspondence, to upload the file into our document repository, add metadata, and determine if there were any tasks. And if necessary, they created a new entry in our tasking system for review by our executives' committee. All this needed to be done for 100 to 200 documents a day, every day. If we had a holiday and Europe didn't, then they would have twice as much work after returning from

“We have built an AI that leverages the knowledge and experience of humans to help them in their daily work by making it faster — but more importantly, the work is more meaningful.”

their day off. A very, very important job, but also very repetitive work.

**Pons:** By the end of 2019, we had automated this job. Removed humans from being "in the loop" and put them "on the loop," reviewing the recommendations of the computer models.

**JWC:** And you achieved this by harnessing AI?

**Pons:** Yes, absolutely. The tool has an understanding of the directives that we use and we developed a model of that, but we also trained it on around 20,000 pieces of incoming correspondence, because sometimes a person reads a document and makes a judgement that the directive doesn't cover. So the AI learned its own way of judging incoming documents. We put our tool up against people to determine how good it was. It's no worse than people at making judgements on what a document is and who should read it, and it is much better at populating the metadata fields because it can do that very quickly. With AI FELIX, we achieved an 80%-time reduction in processing the daily incoming correspondence.

**JWC:** What are these metadata fields that need populating for every entry?

**Pons:** They are the data you have to enter, such as title and author, classification, keywords and topics. They are also the data our search engines use to find documents. We took our AI tool and ran that on the entire back catalogue and brought all the documents in the archive up to the same level with their metadata. And now our regular search tools are much more effective because our metadata is better.

**JWC:** And this is "AI FELIX"?

**Pons:** Yes. The two people that Simon mentioned who were working on our Command Read Board? One of those was Yeoman First Class Felix. We named the tool after him. This is the human side of the AI FELIX story, and the core of our mission. AI FELIX stands for "Artificial Intelligence Front End Learning Information Execution." We have built an AI that leverages the knowledge and experience of humans to help them in their daily work by making it faster — but more importantly, the work is more meaningful.

**Purton:** The Command Read Board was a starting point. Automating the ingestion of incoming correspondence was one part of our vision for artificial intelligence. It was an example of how AI could support a small process and release staff from menial tasks to allow them to focus on more important things. So less of "what is this document?" and rather more of "what should this command do about this document?" However, we always had bigger ideas. We wanted not just to automate staff work, but to create a digital army of staff officers. That was the next step. We were working on digitizing the knowledge base and then we figured that every staff officer could have a digital assistant to support them in their work.

**JWC:** And is this the AIDA project I have heard about?

**Pons:** No, AIDA comes later. First, we digitized every single file in our entire public library. I should clarify that we have public and private folders. If you put something in your public folder, then you are sharing that with everyone in your organization. If you put something in your public folder, anyone [in NATO] can see that. We used artificial intelligence to put all these public files into a massive database. Simon calls it the "Knowledge Universe." Everything in it is positioned relative to everything else by similarity.

Currently, the Knowledge Universe has around 600,000 things in it. And that's not just documents; it's people. It reads your job description and you can add your interests and skills. And it's also our taskers. If you search for something, it goes to that part of

►►►

the Knowledge Universe and it shows you the documents, the people and the tasks.

**JWC:** So it's like a super search function?

**Pons:** Search is part of it, but more importantly it uses a recommender function, so if someone uploaded a file somewhere that was of interest to you, the tool brings that to your attention. It's like Netflix — if you watch a lot of Arnold Schwarzenegger movies, it starts recommending Sylvester Stallone movies, or Jean-Claude van Damme. And if you read the article and like it, then the Knowledge Universe refines its model of what it thinks you want to see.

The next time a document that you might want to see is uploaded, it will inform you. So the artificial intelligence keeps track of every user's interests. I think we have around 600 users at the moment, so it's juggling little AI versions of all those people in its memory.

**JWC:** So to sum it up: you basically automated the Command Read Board and created a Netflix-style file recommender for all of HQ SACT's staff?

**Purton:** AI FELIX is running in HQ SACT, Supreme Headquarters Allied Powers Europe, Allied Air Command, Joint Force Command Norfolk... And by the end of 2024, it will be integrated with the new roll-out of a staff officer's basic knowledge management toolset. It should be in the hands of over 15,000 staff across the NATO Command Structure.

**JWC:** That is an impressive feat.

**Purton:** And it is not even the best part yet. I had been getting AI FELIX going, but Arnau really started going exponential! We unleashed the AI tool across the enterprise, which led to our cataloging of all of the files — oh, and I know what people are going to think: "Need To Know!" AI FELIX respects the need-to-know principle. It cannot show you something to which you shouldn't have access. This is akin to a law of physics in our universe. Anyway, as I was saying: Arnau wanted to really push the envelope. Tell them about AIDA!

**Pons:** So searching is fine, and recommendations are fine. But why search? Why not engage in a discussion? Why not ask for the answer to your question? AIDA has a large language model (LLM) that leverages NATO domain knowledge to provide answers to NATO questions in different languages. We utilize retrieval-augmented generation (RAG) in order for the LLM to extract the relevant context to your question. This tool resides on the NATO classified network and draws from hundreds of thousands of documents across different security classifications.

NATO produces hundreds of documents every day and its document repositories contain millions of documents, so it is impossible for a person to always find the right information. With AIDA, users can query this vast knowledge universe and obtain relevant answers that include the citations to the references used to produce the response. This ensures traceability of the responses and reduces the risk of LLM "hallucinations."

Our next step is to expand AIDA's functionality by coupling it with AI agents that can conduct tasks such as producing PowerPoint presentations, writing documents by following specific templates, querying internal data-

## HOW AI WILL SHAPE OUR PROFESSIONAL LANDSCAPE

**➲ Automating Repetitive Tasks:**
AI will take over mundane, repetitive tasks, freeing up human workers to focus on more creative, strategic, and value-added activities.

**➲ Enhancing Decision-Making:**
AI-driven analytics will provide data-driven insights for better decision-making. Whether in finance, marketing, or operations, AI will sift through vast datasets to identify patterns and trends. Predictive models will guide strategic choices, risk assessments, and resource allocation.

**➲ Personalized Experiences:**
AI will tailor experiences for employees and customers alike. Chat bots will handle queries, while recommendation engines will suggest relevant content. In human resources, AI can personalize learning paths, career development, and performance feedback.

**➲ Collaboration with AI:**
Humans and AI will collaborate seamlessly. Imagine brainstorming sessions with AI-generated ideas, followed by human refinement. AI-powered virtual assistants will schedule meetings, manage calendars, and even draft emails.

**➲ Upskilling and Reskilling:**
As AI automates certain tasks, workers will need to acquire new skills. Organizations will invest in upskilling and reskilling programmes. Learning platforms will adapt to individual needs, suggesting relevant courses based on career goals.

**➲ Ethical Considerations:**
AI will force us to grapple with ethical dilemmas, such as, how do we ensure fairness, transparency, and accountability in AI systems?

**➲ Remote Work and Flexibility:**
AI-powered tools will facilitate remote work. Virtual collaboration platforms, AI-driven project management, and automated workflows will enable seamless remote collaboration.

**➲ Healthcare and Well-Being:**
AI will enhance employee well-being. Wearables will track health metrics, AI chat bots will provide mental health support, and personalized wellness plans will become common.

**➲ Supply Chain Optimization:**
AI will optimize supply chains, predicting demand fluctuations, managing inventory, and ensuring timely deliveries.

**➲ Creativity and Innovation:**
Contrary to fears of AI stifling creativity, it will actually amplify it. AI will accelerate innovation by automating research, simulations, and design iterations.

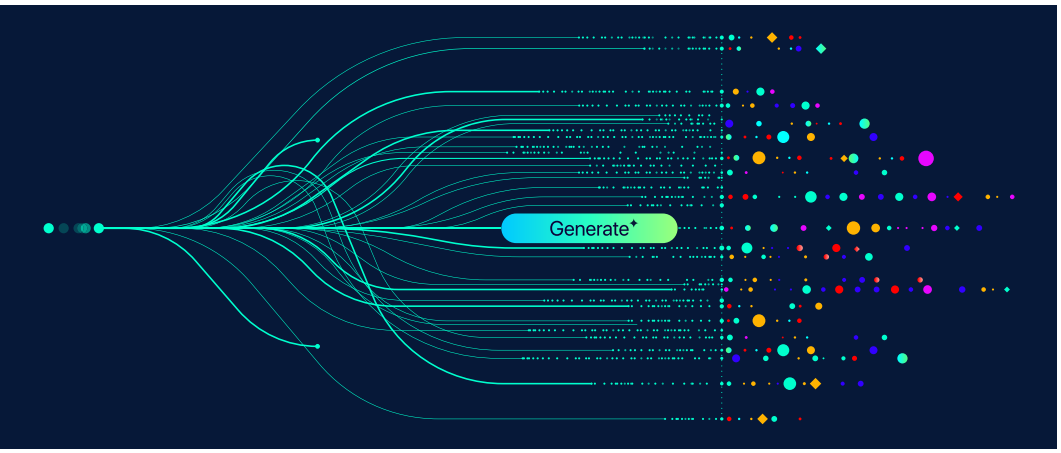


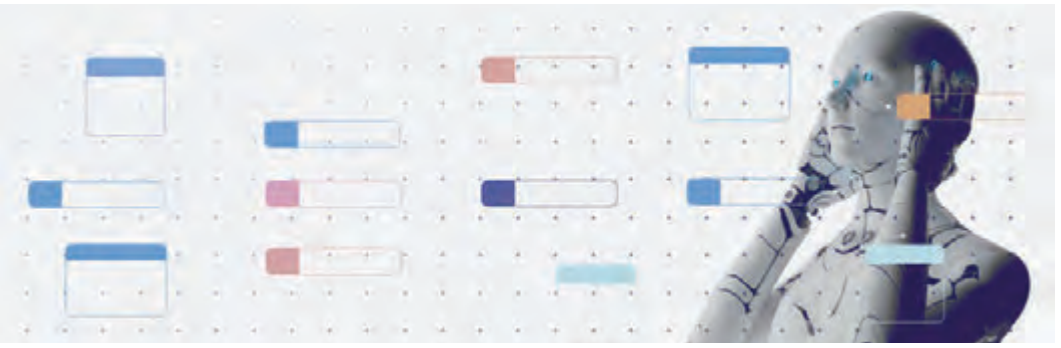

bases, and interacting with other NATO tools.

**JWC:** What does AI mean for our future, and for the future of NATO in particular?

**Pons:** Ideologically, I am a "techno-optimist." I am eager to leverage the benefits of AI. So yes, there will be drones conducting warfare autonomously in the near future, but I do not see [the artificial superintelligence] Skynet from the "Terminator" films taking over humanity. There are also very important concerns about job displacements and impact on society due to a wave of automation powered by AI. As it occurred in past technological waves, our jobs will evolve and adapt to these incredible new AI tools that we will have at our disposal.

Hopefully, we will use AI to make our work more meaningful by focusing more on adding value rather than performing manual tasks. The goal is to exploit our potential by doing what we do best and letting the machines do the work that can be automated. Exercising our judgment, understanding a completely new scenario by connecting the dots with our experience, unleashing our ingenious creativity — those are the skills at which we excel as humans.

It would be too cliché at this point to ask an LLM to summarize the most consequential trends in the field of AI. But actually, LLMs will struggle with the exponential nature of this technology since it is trained with human data. As humans, we are infamously bad at seeing exponential trends and making accurate projections over time. However, human imagination is boundless and incredibly creative, so let's use imagination to envision our future enabled by AI. In the short term, large language models will gradually converge in accuracy and quality as they are mostly trained on the same data (the Open Internet). This will lead to a commoditization of LLMs, which will combine with another key trend: open-source models catching up and democratizing AI space. With ever more powerful open-source models, defence organizations will have more flexibility to leverage LLMs in secure enclaves that maintain the privacy of information. This has indeed been our approach to AIDA at ACT.

**JWC:** So what's next?

**Pons:** Artificial intelligence agents. With LLMs we can now speak to computers using natural language. However, LLMs can provide much more than written responses; they can code, they can interface with other applications, they can reason and devise an action plan taking multiple steps using different tools to complete a task. That's the point when we'll be able to ask an AI agent to use our data, external applications and open-source information to book our upcoming business trip, write a strategic point paper, create a personalized training course, or draft a plan.

Looking further, as LLMs and AI agents get better, more accurate, and more consistent, we will be able to delegate more tasks to them and use agents as personal executive assistants. Every staff officer will head a team of multiple digital agents to conduct specific tasks, and those agents in turn will work collaboratively to succeed.

For instance, one digital agent will conduct research on a topic using multiple sources, another will draft a presentation for decision-makers, another will fact check the presentation, another will red-team the conclusions, share with colleagues to incorporate their comments (including emailing them and applying their inputs), and finally doing your review before submission. This personal executive assistant will have intimate knowledge of how you want tasks to be done, when to act, and how to prepare you in anticipation of incoming work, so that you can confidently delegate tasks on a regular basis.

The ultimate goal is to automate menial tasks, augment your intelligence and unleash your creativity. If you're a fan of the "Ironman" comics and the movies, you see how Tony Stark works with his AI, J.A.R.V.I.S. Jarvis does a lot of heavy lifting, but Tony Stark is providing the direction and the inspiration.

If [U.S. computer scientist and inventor] Ray Kurzweil's prediction is correct, before 2030 we will reach general artificial intelligence, meaning that humans and AI will merge into a super-intelligence and thus reach the next stage of human evolution. Speculations on what could happen after that may seem to belong more in the realm of science fiction, but 19th-century readers thought the ideas of Jules Verne would always remain fiction, too... Whatever happens: remember that today's AI is the worst you'll ever use! ✦