AI FELIX, a Quest for Innovation in Artificial Intelligence in NATO

Arnau Pons, Simon Purton, Michael Martin, and Sean Lewis
Headquarters Supreme Allied Command Transformation
Norfolk, VIRGINIA, 23551-2490
UNITED STATES

arnau.pons@act.nato.int

ABSTRACT

In a few years, virtually all areas of the economy will in some way or another be impacted by Artificial Intelligence (AI). This Fourth Industrial Revolution is characterized by a transformation wave through digital technology. We are living in the era of Big Data, Cloud Computing, mobile devices and many other technologies which have the potential to change our society in ways never experienced before.

These technologies are spurring a growing demand from users across the NATO enterprise to harness the power of artificial intelligence and big data analytics to manage and synthetize increasing amounts of data received by staff. The Artificial Intelligence Front End Learning Information Execution (AI FELIX) is a machine learning tool developed at HQ SACT that operates on the NATO classified system. Each day it processes hundreds of incoming documents received at HQ SACT; makes recommendations on their distribution; automatically uploads them with a full suite of searchable metadata to enhance knowledge management; and assesses whether the incoming correspondence is likely to require action by ACT staff officers and, if necessary, it will generate draft taskings.

This paper presents the evolution of AI FELIX from its development as a prototype, followed by refinement through continuous iteration, certification through NATO's Independent Verification and Validation process, and finally spiralling out to an expanded pool of users and capabilities of the system. This experience has begun to coalesce into a roadmap to employ DevSecOps, agile development, and continuous iteration based on users' feedback as a more flexible approach to incorporate innovative AI and Machine Learning capabilities in NATO.

1.0 INTRODUCTION

In the information age of the 21st century, organisations that can effectively receive, triage, disseminate information, and generate and track work arising from receipt of that information will be in a position to optimize their workforce and fully leverage Emerging Disruptive Technologies. Given the potential capabilities of machine learning tools, Headquarters Supreme Allied Command Transformation (HQ SACT) developed an experimental AI tool to solve its knowledge management problem. The Tasking and Knowledge Management, and Analysis of Alternatives Branches, in partnership with the NATO Communication and Information Agency (NCIA) and Industry, have developed a game-changing technological surrogate that automates previously monotonous, resources intensive, and manual processes.

AI FELIX was an experiment that leverages artificial intelligence, machine learning, and advanced analytics technologies to significantly reduce staff resources assigned to action daily incoming correspondence received at HQ SACT and the related tasking, monitoring, collaboration and performance metrics. Potentially, AI offers significant time-savings as it can process in seconds that which takes people hours. Furthermore, AI can increase the accuracy of information life-cycling within an organisation. AI FELIX was designed to be a modular and agile application that can be easily adapted to other NATO organisations, which was demonstrated when it was successfully exported to Supreme Headquarters Allied Powers Europe (SHAPE) to maintain their mission effectiveness under Covid-19 related working practices.

In essence, the AI FELIX approach aims to automate tedious work so that people can focus on more complex tasks that require human judgement and reasoning. Nowadays, even with the impressive capabilities of artificial intelligence, computer machines have not yet surpassed the intelligence and reasoning capability of humans. Through a Graphical User Interface (GUI), users are able to link AI FELIX with NATO's information and knowledge management tools.[1] Using the GUI, users are able to provide feedback, which allows the AI to improve its predictions, a process known as supervised learning.

Artificial Intelligence is poised to play an ever increasing role in the development of future defence capabilities, so it is critical that NATO embraces the latest developments and techniques and understands their potential to have an impact on the Alliance's defence and security. The Alliance can maintain its technological and military edge in the future warfare environments if it leverages the potential of AI technologies.

The AI FELIX project demonstrated the feasibility of applying AI in a NATO classified environment. In this regard, AI FELIX was unprecedented. This pathfinding effort provided a platform to experiment with new technologies as well as a vehicle to put in practice agile development methodologies. The intersection between Emerging Disruptive Technologies, state-of-the-art development techniques, organisational culture, and procurement processes, generated a wealth of lessons identified on how to exploit the potential of artificial intelligence in NATO.

2.0 TRENDS OF ARTIFICIAL INTELLIGENCE

Over the last ten years, the Artificial Intelligence industry has experienced explosive growth permeating a wide range of sectors. Among the different business segments, artificial intelligence software and autonomous vehicles sectors are forecasted to grow at an annual growth rate of over 40% [1]. This outstanding growth rate and increase in capital investment shows that Artificial Intelligence technologies are not only the future, but also the present. After a few years of strong development, the Artificial Intelligence sector has definitely taken off. In turn, the rapid growth underscores the fast pace of change in the Artificial Intelligence industry and the increasingly shorter development cycles. As it occurred with the IT industry in the past, the versatility and transversal character of Artificial Intelligence result in a fertile ground for growth as this technology can be applied to multiple sectors. From a geographical perspective, the US and China lead the technological advancements in artificial intelligence followed by Japan [2].

"It's not who has the best algorithm that wins. It's who has the most data".

Andrew Ng

In Artificial Intelligence, scale matters. The accuracy of machine learning techniques can depend heavily on the amount of data used in the computation, as shown in Figure 1. Although some algorithms perform better than others, all supervised algorithms exhibit the same general trend, with accuracy improving with larger training datasets. Internet companies and other machine learning practitioners have strong incentives to gather as much data from users as possible in order to enhance the results from their algorithms, maximizing their profits.

Since the Artificial Intelligence research field started to blossom in the late 90's, there has been an impressive progress in the development of new algorithms that allow the machines to perform extremely complex tasks, and approach one step closer to the capabilities of the human brain. For Artificial Intelligence to take off, there needed to be a convergence of three factors: availability of large amounts of data, sufficient computational power, and innovative algorithms. With the advent of widespread use of broadband internet

[1] NATO has a number of knowledge management tools, such as the Enterprise Document Management System (EDMS), Tasker Tracker Plus (TT+), and the NATO Information Portal (NIP).

and cell phones in the early 2000's, the availability of data from websites, internet services, e-commerce, industry, academia, and many other sectors, has grown enormously with the rate of growth accelerating yearly as information technologies expand. However, large datasets increase the computing power necessary to efficiently train and evaluate machine learning models. This need for more computing power has been accomplished by taking advantage of High Performance Computing centres that leverage parallel computing, which allows to perform very demanding computer calculations with efficiency and speed.

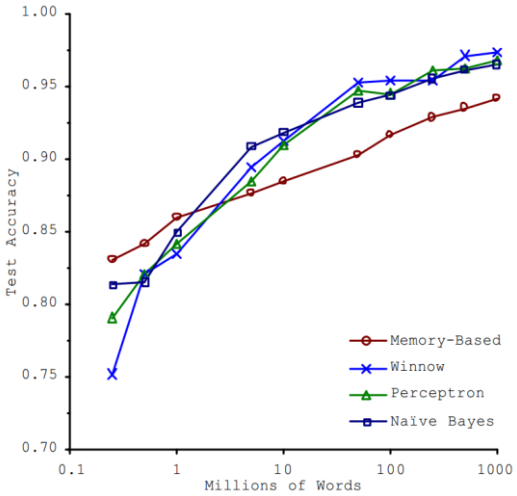

Figure 1: Learning curves for different algorithms for Confusion Set Disambiguation, which is one class of natural language processing problem [3]. The accuracy improves with increasing training dataset size.

Although some advanced machine learning algorithms were actually developed in the 80's, they could not be used to their full capacity due to insufficient amounts of data and lacking computational power. Today, entry barriers are decreasing, and with the increased popularity of open-source code libraries, proliferation of online courses on machine learning, ultra-fast internet connections, and cloud computing, the conditions are ripe for greater participation in the artificial intelligence technological wave.

3.0 AI FELIX

The AI FELIX project officially started after winning the HQ SACT i3 (Initiate, Innovate, Imagine) week Innovation Challenge in November 2018. The goal was to demonstrate how HQ SACT could use artificial intelligence, machine learning, and advanced analytics technologies to enhance the speed, efficiency, and overall quality of command's information and knowledge management efforts. The Tasking and Knowledge Management (TKM), and Analysis of Alternatives (AoA) Branches, in collaboration with the NATO Communication and Information Agency (NCIA) and Industry, collaborated to develop a Minimum Viable Product (MVP) to replace and enhance the existing Command Read Board (CRB) capability.

The Command Read Board is the digital library that stores the incoming official correspondence to Headquarters Supreme Allied Command Transformation. Every weekday, the Registry of HQ SACT receives more than 100 official documents, comprising thousands of pages, from across NATO, which require classification, distribution and consolidation into the CRB. This task was performed manually by staff, by reading, sorting and reviewing hundreds of pages of documents per each day. This particular task represented a suitable application for machine learning; it had bounds, a limited number of users, a clear process, as well as abundance of data. In this way, the team embarked on a journey to develop AI FELIX as an artificial intelligence application to enable ACT to maximize the knowledge, experience and accrued judgement of its military and civilian personnel. The goal was to shift repetitive and monotonous tasks to a machine while empowering humans to synthesize the outputs of AI technology to make decisions

proactively, rapidly and more precisely, essentially putting people "on the loop" as opposed to being in it.

The AI FELIX CRB application is able to read documents and extract metadata such as dates, addressees, security classification, topics, and keywords. Throughout its operation, AI FELIX has demonstrated high levels of accuracy, achieving 90-100% for several fields, and in many cases has shown to be more accurate than people. The tool is capable of determining whether a document contains actions and tasks, and draft taskers (orders) to be assigned to staff officers. As part of the process, incoming documents are more thoroughly tagged with more complete and richer metadata, which enables more effective knowledge management and increased searchability.

Figure 2 depicts the dynamic learning diagram of the AI FELIX Command Read Board application. In the first step, AI FELIX receives the incoming correspondence. As users verify and correct the metadata presented by AI FELIX, the system acquires more data points. Periodically, the tool retrains its algorithms with these growing datasets, further increasing the accuracy of the system over time.

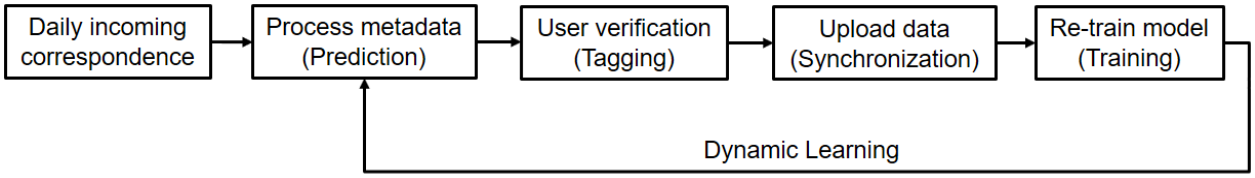

Figure 2: AI FELIX Command Read Board application machine learning cycle.

3.1 Evolution of the Project

AI FELIX project development work was set to start in January 2019. However, before the real code development work could start, there was a period of preparation of datasets and documentation, and the establishment of working protocols for developing machine learning code in a NATO classified environment. This initial stage already previewed some of the challenges of implementing Emerging Disruptive Technologies within a capability development process better tailored for long-term programmes that evolve slowly over time. Furthermore, cybersecurity and software development regulations complicate the effort of developing AI capabilities that require agility and constant evolution to thrive. Indeed, the pace of change of AI technology is so fast that it may be very difficult for any organisations to evolve their doctrines in order to leverage the capabilities of these systems while complying fully with cybersecurity, legal and ethical considerations. From the beginning, the AI FELIX project was a path-finding effort that helped NATO evolve policies to facilitate development of AI capabilities.

By November 2019, version 1.0 of the AI FELIX CRB tool became operational and started to be utilized. The employment of AI FELIX as an operational tool for official business represented a remarkable leap of faith for HQ SACT, albeit a leap informed by lessons from fielding a mature prototype. The management of official incoming correspondence to a military headquarters represents a delicate application because the CRB is the mechanism that distributes documentation and official tasks to all departments. In live operations, AI FELIX consistently delivered accurate results with an average processing time of 27 seconds per document, with the net effect of an 80% reduction in processing time with respect to the legacy method.

Rather than freezing the development of the tool, the team strived for continuously improving the tool based on user feedback. In the software community, this agile management approach is known as Continuous Iteration and Continuous Deployment (CI/CD). Within this model, software is continuously evolved adding new features and enhancing its functionalities by periodically releasing new updates that are informed by user feedback. In comparison with other software tools that follow more rigid development approaches with

little user feedback, with CI/CD users benefit because tools improve quickly based on their requests. Therefore, rather than setting the AI FELIX requirements in stone and trying to perfectly anticipate what users might like, this software development approach allows AI FELIX to be a living tool that continuously adapts to users' needs.

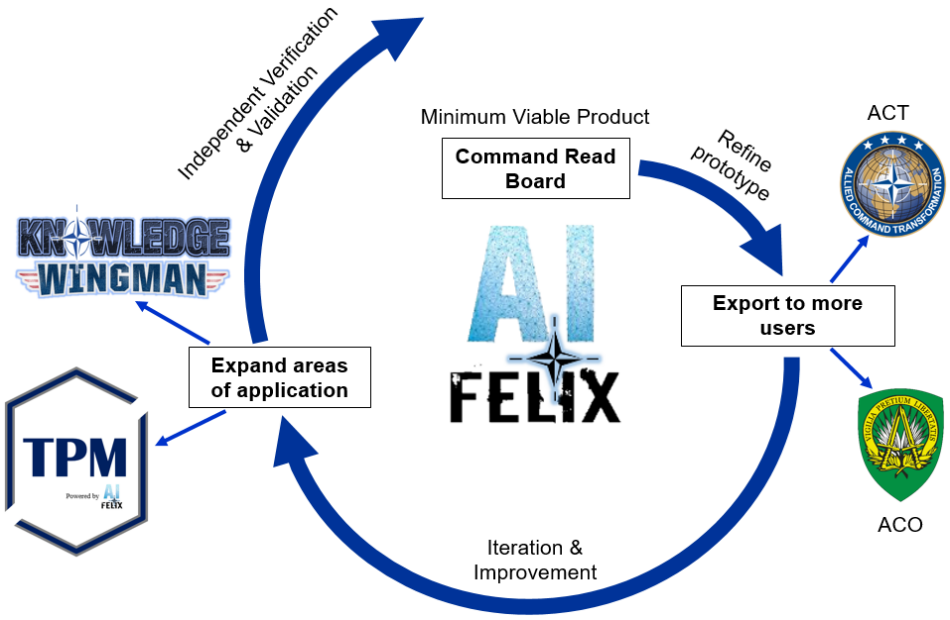

Figure 3: Evolution of the of the AI FELIX project.

As shown in Figure 3, the AI FELIX project has evolved in an outwards spiral. In the first step, the initial AI FELIX tool, the Command Read Board, was continuously improved based on user feedback. Furthermore, this initial prototype was enhanced to include new and more advanced functionalities. After refining the CRB tool and successfully using it operationally for a few months, it became sufficiently mature to consider its distribution to other users across the NATO enterprise. Within the NATO enterprise, the AI FELIX CRB tool has the potential to enhance the daily work of up to 50 personnel. In this way, one of the great advantages of machine learning is scalability; once a tool is developed for a small set of users, it is generally easy to scale the system up to a larger number of users. The natural progression of this artificial intelligence project led to the export of the tool to other NATO locations expanding the number of users. With the COVID-19 crisis limiting the amount of personnel working on site, the AI FELIX team exported the CRB tool to NATO Supreme Headquarters Allied Powers Europe (SHAPE). Within two months, the AI FELIX CRB tool was adapted, tested and made operational in SHAPE. The rapid deployment of AI FELIX to a headquarters with more personnel than HQ SACT demonstrated that the system can be tailored to specific users while leveraging larger amounts of data, thereby paving the way for future exploitation in other locations.

The AI FELIX CRB tool featured several natural language processing functionalities such as extracting topics and keywords from documents, determining actions or tasks, comparing documents, recommending information, all the while synchronizing data with EDMS and TT+. Using these functionalities in a modular fashion, new tools could be devised targeting other information management areas. Indeed, one of the overarching goals of AI FELIX was to assist NATO personnel in managing the ever larger amounts of information. In January 2020, a new phase of AI FELIX started with the aim of expanding the machine learning, big data and advanced analytics tools to other areas of information and knowledge management. Rather than a single tool, AI FELIX then became a suite of AI-enabled applications with three components: Command Read Board, Tasking Portfolio Management, and Knowledge Wingman.

Tasking Portfolio Management

Tasking Portfolio Management (TPM) is the second application from the AI FELIX suite, which has been operational since April 2020. This tool provides a portfolio management analysis of all tasks within Tasker Tracker Plus (TT+) for a particular NATO command. This application allows users to more effectively manage in near-real time the progression of multiple taskers at any one stage. The tool provides quantitative and performance metrics such as predicting the Gantt timeline of any tasker lifecycle, the risk of a tasker being late, and other schedule and performance metrics. In addition to extracting data from TT+, TPM gathers information from other available sources to connect particular taskers to Lines of Effort and the Capability Development Programmes. The application is also able to automatically send alerts and reminders to personnel to mitigate the risk of a tasker being late. The principle of TPM is to switch tasking management from reactive to preventive. Furthermore, TPM can inform leadership with objective and accurate metrics of status of work in a NATO command.

Knowledge Wingman

Knowledge Wingman (KW) is an AI-enabled staff assistant currently in development that fuses multiple isolated data sources to enhance information flow and knowledge management. KW is able to compare a document that the staff officer is working on with other documents from NATO enterprise document libraries, as well as presenting the user with similar relevant documents. In addition, this personalized assistant can identify other authors that are working on similar products, and displays their contact information to foster collaboration. Furthermore, the tool generates a document tree that searches for the referenced documents along with their version history over time. KW includes an auto-tagging capability for individual or multiple document upload that suggests machine learning predicted metadata for the user. Based on the user interests, job position, and document library, Knowledge Wingman provides a reading list of documents that are of interest to the user as well as related training courses. Furthermore, the tool includes a daily digest email with three Command Read Board documents related to the interests of the user. The recommendation systems from KW are based on machine learning algorithms that adapt dynamically to the evolution of the user preferences of the user over time.

Independent Verification and Validation

Although the AI FELIX tools are being used for official business at HQ SACT, they remain, in a sense, in a prototype phase. The prototype status implies that these tools cannot be deployed in the NATO operational network since they have not been yet added to the official catalogue of tools. To do so, applications must pass the NCIA Independent Verification and Validation (IV&V) process, which ensures that any new tool complies with NATO cybersecurity, compatibility and software regulations. Currently, the AI FELIX CRB application is in the IV&V process, which should be finalized at the beginning 2021. After this crucial step, the CRB application will become part of NATO official tools and it will be much easier to export to other NATO entities. As the other AI FELIX tools become more mature after a period of testing and evolution through user feedback, the same process can be followed to deploy these in other locations and reap the benefits of AI-enabled applications for more users within NATO.

3.2 Lessons learned

As a pathfinding effort, AI FELIX has demonstrated the challenges of innovating with artificial intelligence within the Alliance. The conventional capability development process is optimised to support large procurements, but is it overly bureaucratic when one considers applying them in a more agile space such as prototyping, generally the capability development has low risk tolerance, and many other organisational processes that exist for sound reasons, can complicate and slow down the integration of artificial intelligence

technologies in the day-to-day operations. The list below presents relevant lessons that can be useful to organisations attempting to implement artificial intelligence projects:

- **Funding for AI-enabled ventures:** the AI FELIX project started with *"seed"* funding awarded as part of NATO HQ SACT's Innovation Challenge. This initial funding allowed for the development of an independent project exclusively dedicated to delivering a prototype designed to leverage AI, Big Data and advanced analytics for information management applications. To spur innovation, it is necessary to consistently invest in innovative projects knowing that not all of them will succeed. If a project is promising and delivers results, it should be supported with funding and resources to foster its growth. Conversely, if a project does not succeed within a given time frame, the team should then extract lessons learned and move on to other projects with renewed strategies.
- **Start small, aim big:** it is often tempting to seek the all-encompassing solutions that solve all the problems for all users. However, such an approach can be overwhelmingly difficult as it requires balancing the orthogonal or even contradictory needs and requirements from a wide user community, while at the same time combining multiple systems and data sources. As a result of this, tool development drags on for a much longer period of time with the consequent cost overshoot and delay in deployment. In contrast, the agile management model consists of quickly developing and deploying a prototype followed by a continuous iteration process to deliver system updates and improvements. This approach is known in software engineering as Continuous Iteration and Continuous Deployment. Rather than seeking a perfect solution, this model relies on prompt and frequent user feedback and new technologies to continuously enhance the product. Furthermore, the flexibility of the agile management approach allows for an easier spiral out of the tool to larger user communities or into other areas of application.
- **Embracing fluid requirements:** when delivering innovative solutions is the goal, it is very difficult to know in advance the requirements of the project. In the initial stages, the project managers may have a vague idea of the specific goals of the project. In the traditional capability development process, the requirements are elicited at the beginning of the project, and then the programme continues with frozen requirements until delivery. The disadvantage of this approach is that by the time the capability is delivered years later, the assumptions and technologies on which the project was based are already obsolete. For this reason, the traditional deliberate processes are less well suited for modern software development, artificial intelligence, and big data applications because they lack the flexibility to adapt to user needs and evolving technology. Equally, software development efforts often do not have significant physical upfront expenditures, so the cost of changing one's mind is less dramatic.
- **User at the centre:** throughout the development of AI FELIX, the developer and the user often differed in their visions of where the tool development should lead to. There is a more dynamic interplay between what the user thinks they need, what the developer can offer, and the project tolerances. Furthermore, views and requests from the users evolve as they interact with partial solutions. In essence, users need to operate the prototype to consolidate their views on the requirements. It becomes critical that the user be placed at the centre of development, providing feedback at all times, testing the system, and guiding the developers on how to evolve the software application.
- **Consolidate and clean the data:** in machine learning the accuracy of the results is only as good as the quality of the data. Before undertaking a project in artificial intelligence, it is paramount to analyse the quality and structure of the data. For supervised learning algorithms to work effectively, it is necessary to organise the data in a consistent manner while at the same time maintaining a dataset of at least tens of thousands of data points.
- **Failure as an option:** just as it happens in the venture capital community, not all innovation projects succeed. Delivering innovative solutions means that some efforts will result in failure. For

innovation to thrive, it is critical to foster a risk tolerant environment that allows development teams to test new ideas. When organisations reward teams that try new approaches, even though a percentage of these ideas will fail, the aggregate result is a collective that is more innovative and resilient to technological changes.

- **On site talent with Big Data/AI skills:** although outsourcing innovation in AI has its own advantages, whenever possible it is very useful to count on personnel on site with expertise on artificial intelligence. The proximity of developers and users greatly facilitates the process of iterating and improving an AI software prototype.
- **Integration of all stakeholders:** in order to speed up the development, iteration, production, deployment and certification of innovative AI tools it is necessary to engage all the stakeholders from the beginning. The AI FELIX project learned that delaying the engagement with regulators or other stakeholders may lead to re-doing parts of the software prototype to comply with regulations or system constraints. Even though production and certification of software may seem very remote in the initial prototype stage, it is very useful to engage with the departments responsible for deploying, maintaining and certifying the tool. This approach will allow the developer team to embrace best practices from the beginning and design the system for its eventual deployment and certification, thereby saving time and resources later on.
- **Need to share vs. need to know:** within an organisation, different users have different needs, and therefore, they may need different tools to carry out their functions. However, even with different tools, users need to access data across tools. Furthermore, in order to avoid double data entry and outdated datasets in the different tools, systems should be integrated and share data. When developing new software tools, especially in the realm of machine learning, it is critical to strive for data fusion and federation of applications. With this approach, the different user communities can continue using their specific tools that are highly specialized for their function while at the same time benefiting from and contributing to the datasets from the other collectives of users within the organisation.

4.0 INNOVATING WITH AI IN NATO

NATO is not alone in facing the challenge of adapting to and embracing Emerging Disruptive Technologies, and in particular Artificial Intelligence. Across society and within the commercial sector, organisations are now faced with the challenge of navigating this evolving technological landscape with the obsolescence Damocles sword hanging over them. In the business sector, the principle of *"disrupt or be disrupted"* is driving organisations to exploit and heavily invest in Emerging Disruptive Technologies to achieve a commercial edge over competitors **Error! Reference source not found.**.

The consolidation of the lessons identified through AI FELIX coalesced into a roadmap proposal to employ DevSecOps, agile development, CI/CD and user feedback as a more flexible approach to incorporate innovative artificial intelligence capabilities in NATO.

4.1 DevSecOps

DevSecOps, which stands for Development, Security, and Operations, is a software development methodology that integrates security at every phase of the Software Development Life Cycle (SDLC). This methodology encompasses people, process and tools to enable teams to design, develop, test, integrate, deploy, deliver, support, and update software "faster, safer, and sooner". DevSecOps combines agile management with development, quality assurance, operations, and security, in order to foster collaboration and seamless integration between the different stakeholders of the SDLC. By "baking" security into the process, the code is developed within a security framework that leads to easier and more cost effective development. Figure 4 depicts the DevSecOps cycle as a continuously iterative process.

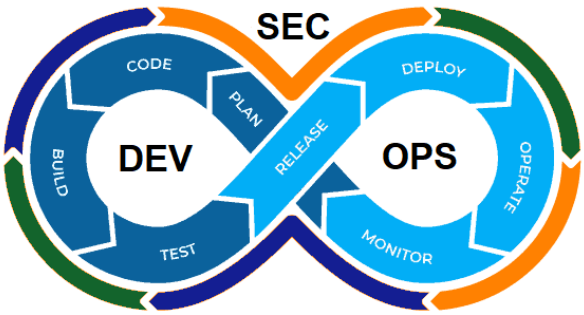

Figure 4: DevSecOps cycle.

4.2 Framework for Innovation in AI in NATO

This section outlines a framework to innovate with artificial intelligence in NATO, with a particular focus on Allied Command Transformation. The model is therefore specific to NATO, but several characteristics may be adapted to other organisations aiming to integrate AI technologies into their day-to-day operations. As the Transformation Command, ACT delivers innovative solutions that enhance NATO's Capability Development by leveraging advanced technologies such as Artificial Intelligence. Within the realm of Information Technologies, the NATO Communications and Information Agency (NCIA) is in charge of managing NATO's communication and information systems.

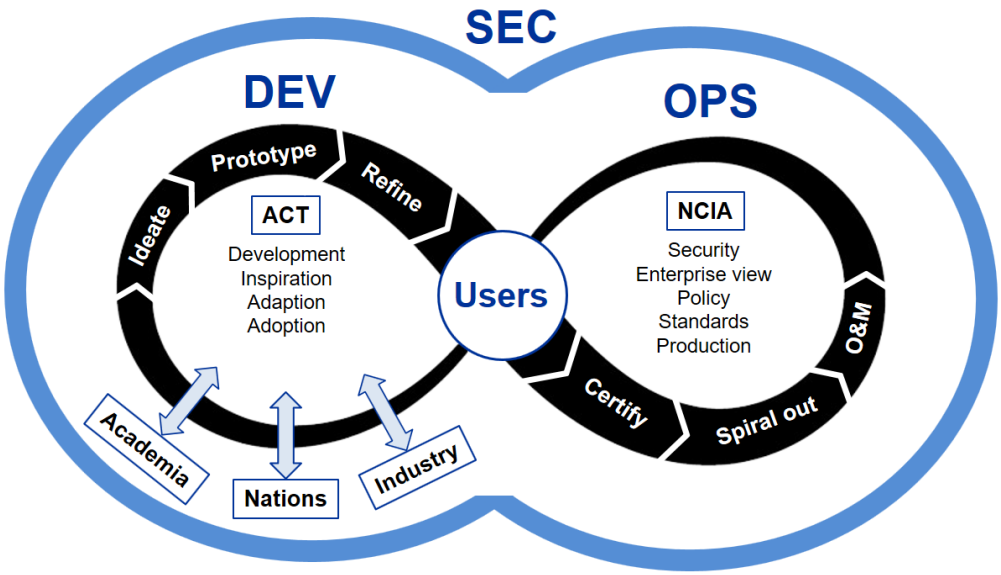

Figure 5: Proposed framework for innovating in Artificial Intelligence in NATO.

1. **Ideate:** the first phase consists of generating ideas for AI-enabled applications to serve a community of users. These ideas can come internally from ACT as well as from Academia, Industry and Nations. Academia may bring different perspectives on the implications and potential of AI in NATO; Industry may fully inculcate the tremendous advancements of AI technologies; and Nations can also leverage the idea and seek partnership initiatives to develop the capability further.

2. **Prototype:** develop a Minimum Viable Product of an application that leverages Artificial Intelligence. Test the prototype with users as soon as the MVP has a reasonable level of maturity and robustness. To speed the development, it is preferable to test sooner rather than later.

3. **Refine:** conduct a continuous iteration process to improve the application based on user feedback as well as debug any code error. Start using the application in daily business within NATO.

4. **Certify:** work with NCIA to conduct the IV&V process to certify the tool for official use in the NATO production network.

5. **Spiral out:** with the knowledge and lessons identified from developing the application, expand the number of users and locations, as well as continue adding more functionalities.

6. **Operations and Maintenance (O&M):** while operating the application, gather continuous user feedback and research new technologies in order to periodically upgrade the capabilities of the tool.

Programme tolerances, on schedule, costs, and requirements, should be set with both specific sprints and the overall objectives in mind, but there should be an emphasis on delivering updates that incorporate the newest technologies.

5.0 CONCLUSIONS

At its core, the challenge of implementing and integrating Emerging Disruptive Technologies is an organisational culture problem more than a technical difficulty. The conventional approach of eliciting requirements at the start, freezing them, developing the product without constant user input, all the while following rigid and cumbersome procurement processes, results in a capability delivery methodology that lags behind new developments and eschews the value that digital information age technologies can provide. The application of traditional capability development processes to artificial intelligence software development leads to deliver tomorrow's capability with yesterday's technology to unsatisfied and disengaged users. To remedy this outcome, development processes need to be adapted, procurement and cybersecurity regulations streamlined, team management flexibility expanded, engagement with users fostered, and continuous change embraced.

In a World where organisations compete under the principle "*disrupt, or be disrupted"*, the only certainty is the continuous change of technology at an ever-increasing pace [4]. Rather than fighting this change, organisations have the opportunity to embrace artificial intelligence in order to become more efficient and effective through enhanced information and knowledge management. The AI FELIX project has successfully demonstrated within NATO the benefits of innovating in artificial intelligence while leveraging agile development methodologies.

6.0 REFERENCES

[1] Morgan Stanley, https://www.morganstanley.com/ideas/fourth-industrial-revolution

[2] WIPO (2019). WIPO Technology Trends 2019: Artificial Intelligence. World Intellectual Property Organization, Geneva, 2019.

[3] Banko, M., Brill, E., *Scaling to Very Very Large Corpora for Natural Language Disambiguation.* Proceedings of the 39th Annual Meeting on Association for Computational Linguistics, pp. 26-33, Toulouse, France, 26-11 July, 2001.

[4] JP Morgan, *Disrupt, or be disrupted. Corporate strategies for an increasingly disruptive world*, September 2017.