Part 04 — GenAI Research:

Use Cases & Case Studies

Cybersecurity · Pharmaceutical Industry · Government & Public Sector · Insurance · Aerospace & Defense

Prepared by: Shivya

Organization: MacroEdtech

Program: GenAI Research Internship — Summer 2026

Submission Date: June 2026

Cybersecurity	Pharmaceutical	Government & Public Sector	Insurance	Aerospace & Defense

Table of Contents

About This Document

This is my Part 04 / Week 04 research submission for the MacroEdtech GenAI Research Internship — and the final part of this four-part series. Having covered 15 industries across Parts 01, 02, and 03, this last part goes deeper into five more sectors: Cybersecurity, Pharmaceutical Industry, Government & Public Sector, Insurance, and Aerospace & Defense.

I picked up on something interesting starting this part: these five industries all share one common thread — they each deal with extremely high stakes. A cybersecurity failure can shut down a hospital. A pharmaceutical mistake can harm patients. A government AI error can affect millions of citizens. An insurance miscalculation can bankrupt a company. A defense AI mistake can be catastrophic. This made me research each section with extra attention to safety and responsibility, alongside the usual use cases and case studies.

As with Parts 02 and 03, every industry section includes a detailed real-world case study with the problem, the GenAI solution, and the measurable result. Since this is the final part of the series, I have also included a comparison of all 20 industries researched across the entire internship.

What This Report Covers

- 5 industries analyzed: Cybersecurity, Pharmaceutical Industry, Government & Public Sector, Insurance, Aerospace & Defense
- 40+ use cases documented with real-world impacts and company examples
- 5 detailed case studies — one per industry — with problem, solution, and result
- Cross-industry patterns compared with all 20 industries covered across Parts 01 through 04
- 25+ references from academic papers, industry reports, news sources, and company blogs

1. Cybersecurity

Cybersecurity was a fascinating section to research because it is the one industry in my entire project where GenAI is simultaneously the defender and the attacker. Every other industry I have studied uses AI to solve a problem. In cybersecurity, criminals are also using GenAI to create the problem — which makes this section different from anything else I have covered.

The global cybersecurity AI market is projected to reach $60.6 billion by 2028 The reason this industry needs AI so urgently is scale: a large company can face millions of attempted cyberattacks per day, and there simply are not enough human security analysts in the world to review all of them manually. The global shortage of cybersecurity professionals is estimated at 3.4 million people, so AI is not just helpful here — it is becoming essential.

1.1 Use Case Overview

Use Case	What GenAI Does	Real-World Impact	Company / Tool
AI Threat Detection & Response	AI continuously monitors network traffic and system behaviour, instantly detecting unusual patterns that may indicate an attack — and can respond automatically.	Threats detected in seconds instead of the industry average of 200+ days; breaches contained before major damage.	*Darktrace, CrowdStrike, Microsoft Defender*
Phishing Email Detection	GenAI reads incoming emails and identifies sophisticated phishing attempts by analysing writing style, links, sender behaviour, and intent.	Phishing-related breaches reduced significantly; protects against AI-generated phishing emails too.	*Abnormal Security, Proofpoint*
Automated Vulnerability Scanning	AI scans software code and systems continuously to find security weaknesses before hackers can exploit them.	Vulnerabilities found and patched faster; reduces the 'window of exposure' for new software.	*Snyk, Tenable, Qualys*
AI-Generated Security Reports	After an incident, GenAI automatically writes detailed technical reports explaining what happened, how, and what to do next.	Incident response teams save hours per incident on documentation; reports are more consistent.	*IBM QRadar, Splunk AI*
Deepfake & Synthetic Identity Detection	AI tools detect AI-generated fake videos, audio, and images used in fraud, impersonation scams, or disinformation.	Protects companies and individuals from voice-cloning fraud and fake video scams — a fast-growing crime type.	*Reality Defender, Sensity AI*
Security Awareness Training Generation	GenAI creates personalised, realistic phishing simulation tests and training content to teach employees to recognise attacks.	Employee susceptibility to real phishing attacks drops significantly after training.	*KnowBe4, Proofpoint Security Awareness*

Use Case	What GenAI Does	Real-World Impact	Company / Tool
Automated Incident Response Playbooks	When an attack is detected, AI generates and even executes a step-by-step response plan (isolate device, reset passwords, alert team) automatically.	Response time to contain an attack reduced from hours to minutes.	*Palo Alto Cortex XSOAR, Splunk SOAR*
AI-Powered Fraud Pattern Recognition	Detects new and evolving online fraud patterns (account takeovers, fake accounts, payment fraud) that don't match any known rule.	Fraud caught that would be invisible to traditional rule-based systems.	*Sift, Forter*

1.2 Case Study: Microsoft — Using AI to Fight Nation-State Hackers

📌 Case Study: Microsoft Security Copilot — Defending Against State-Sponsored Cyberattacks

Problem: Microsoft's security teams were facing an overwhelming volume of cyberattacks — including highly sophisticated attacks backed by nation-states attempting to breach government agencies, hospitals, and critical infrastructure. A single human security analyst can typically review a limited number of security alerts per day, but companies the size of Microsoft and its customers can generate millions of signals daily. Skilled security analysts were also in short supply and expensive to hire and retain.

What They Did: Microsoft built Security Copilot, an AI assistant powered by GPT-4 combined with Microsoft's own threat intelligence data (gathered from trillions of security signals collected daily across their global network). Security analysts can ask Security Copilot natural language questions like 'summarise this incident' or 'what should I investigate next' and the AI provides expert-level guidance, writes incident reports, and even predicts how an attack might evolve based on patterns it has seen before.

Result: In trials, security analysts using Security Copilot completed incident analysis up to 50% faster, and junior analysts performed tasks that would normally require years of experience. Microsoft reports the tool has helped identify and respond to nation-state level attacks (including ones linked to groups from Russia, China, and North Korea) significantly faster than manual methods alone. Microsoft now uses this same AI to protect its own infrastructure as well as offering it to government and enterprise customers.

What I Learned from This: What struck me most is the idea of AI 'levelling up' junior employees to perform like experts. In an industry with a massive talent shortage (3.4 million unfilled cybersecurity jobs globally), this could be one of the most practical real-world solutions — not replacing security analysts, but making the limited number that exist far more effective.

1.3 The Double-Edged Sword — AI Helps Both Sides

This is the section where I had to think most carefully, because cybersecurity is unique in my entire research: GenAI is being used by defenders to protect systems, but it is also being used by attackers to break into systems. I documented both sides because I think understanding the attacker's use of AI is necessary to understand why defensive AI is so urgently needed.

How Attackers Use GenAI	How Defenders Use GenAI in Response
Writing highly convincing phishing emails with perfect grammar and personalised details scraped from social media	AI email scanners detect subtle signs of phishing that humans miss, even in well-written emails
Voice cloning to impersonate executives and trick employees into transferring money ('CEO fraud')	Deepfake and voice clone detection tools flag synthetic audio in real time during calls
Generating malware code variations automatically to evade traditional antivirus detection	Behaviour-based AI detection that looks at what software does, not just what it looks like, catching new malware variants
Automating the search for software vulnerabilities to exploit before companies can patch them	AI vulnerability scanners that find and help patch the same weaknesses before attackers can exploit them

Cybersecurity — My Key Observations

- This is the only industry in my entire research where I had to study both the attacker's and defender's use of the same technology — that made it intellectually the most interesting section
- Microsoft Security Copilot's ability to make junior analysts perform like experienced ones feels like one of the most practical solutions to a real talent shortage that I found anywhere in my research
- The 3.4 million person global cybersecurity talent gap explains why AI adoption in this industry is not optional — there simply are not enough humans to do the work manually
- Deepfake voice cloning being used for 'CEO fraud' (tricking employees into wiring money based on a fake voice call from their boss) was a genuinely unsettling thing to learn about, and it shows how AI threats are becoming very personal and convincing
- I believe cybersecurity will always be something of an 'arms race' between attacker AI and defender AI — neither side will likely achieve total victory over the other, which means continuous investment in this area will always be necessary

2. Pharmaceutical Industry

I touched briefly on drug discovery AI back in Part 01 under Healthcare, but I learned in this research that the Pharmaceutical industry deserves its own dedicated section because it covers a much wider scope than just discovering new drugs — it includes manufacturing, clinical trials, regulatory submissions, and supply chain management for medicines that reach billions of people.

The COVID-19 pandemic was a turning point for this industry. Traditional vaccine development used to take 10-15 years. Moderna and BioNTech, using AI-assisted mRNA technology, developed working COVID-19 vaccines in under 12 months — a speed that was previously considered impossible. The global AI in pharmaceuticals market is projected to reach $19.4 billion by 2028.

2.1 Use Case Overview

Use Case	What GenAI Does	Real-World Impact	Company / Tool
AI-Driven Drug Target Identification	AI analyses biological and genetic data to identify which proteins or genes are good targets for new drugs to act on.	Target identification time reduced from years to months; more accurate target selection.	*Insilico Medicine, BenevolentAI*
mRNA & Molecule Design	GenAI designs the molecular structure of new drugs and vaccines, predicting how they will behave in the human body before lab testing.	Vaccine and drug candidates designed in days instead of years.	*Moderna, BioNTech*
Clinical Trial Patient Matching	AI scans patient health records to find and recruit the most suitable candidates for clinical trials based on very specific medical criteria.	Patient recruitment time cut significantly; trials reach required patient numbers faster.	*Deep 6 AI, TrialSpark*
Adverse Event Detection & Reporting	AI monitors millions of patient reports, social media, and medical records to detect early signs of drug side effects after a medicine is launched.	Safety issues caught faster, potentially preventing widespread harm before regulators intervene.	*Genpact Pharmacovigilance AI, IQVIA*
Regulatory Document Generation	GenAI drafts the enormous regulatory submission documents (sometimes tens of thousands of pages) required to get a drug approved by agencies like the FDA.	Submission preparation time reduced from months to weeks.	*Generate Biomedicines, Certara*
Manufacturing Quality Control	AI monitors pharmaceutical manufacturing in real time to detect contamination, dosage errors, or equipment issues instantly.	Defective batches caught before reaching patients; manufacturing recalls reduced.	*GE Healthcare, Siemens Pharma AI*

Use Case	What GenAI Does	Real-World Impact	Company / Tool
Drug Repurposing Discovery	AI analyses existing approved drugs to find new uses for them — for diseases they were not originally designed to treat.	Faster and cheaper than developing new drugs from scratch since safety is already proven.	*BenevolentAI, Recursion Pharmaceuticals*
Personalised Dosage Recommendations	AI calculates the optimal drug dosage for an individual patient based on their genetics, weight, age, and other health conditions.	Reduces side effects and improves treatment effectiveness through precision dosing.	*Tempus, 2bPrecise*

2.2 Case Study: Moderna — How AI Helped Build a COVID Vaccine in Record Time

📌 Case Study: Moderna — From Virus Genome to Vaccine Design in 2 Days

Problem: When COVID-19 emerged in January 2020, the world urgently needed a vaccine, but traditional vaccine development takes 10-15 years on average. Every month of delay meant more illness and death globally. The challenge for Moderna was to design, test, and manufacture a safe and effective vaccine faster than had ever been done in human history.

What They Did: Moderna had already invested years in building an AI-driven mRNA vaccine design platform before the pandemic began. When Chinese scientists published the genetic sequence of the COVID-19 virus online, Moderna's AI systems analysed the sequence and designed the mRNA vaccine candidate (which would later become the basis of their COVID vaccine) within 2 days. The AI predicted which part of the virus's genetic code would best train the human immune system to fight it, and how to structure the mRNA molecule for stability and effectiveness.

Result: Moderna's vaccine entered human clinical trials just 63 days after the virus genome was published — an unprecedented speed in pharmaceutical history. The vaccine received emergency authorisation within 11 months of the pandemic's start, compared to the usual decade-plus timeline. Moderna's COVID vaccine has since been administered to hundreds of millions of people worldwide, and the same AI-driven mRNA platform is now being used to develop vaccines for other diseases including flu, RSV, and even certain cancers.

What I Learned from This: This case study showed me the real-world stakes of AI speed. The difference between a vaccine taking 10 years versus 11 months is not just a business metric — it likely saved millions of lives during the pandemic. It made me realise that AI research is not just about efficiency or profit in every case; sometimes it is genuinely about saving human lives faster.

2.3 How AI Is Shortening the Entire Drug Development Pipeline

I wanted to understand the full picture of how a drug moves from an idea to a patient's medicine cabinet, and how AI is speeding up each stage:

Drug Development Stage	Traditional Timeline	With GenAI
Target Identification	2-3 years of lab research	Weeks to months using AI analysis of biological data
Drug Design	3-6 years of trial and error chemistry	Days to weeks using AI molecular generation
Preclinical Testing	1-2 years of animal and lab testing	AI can predict outcomes, reducing some testing needs
Clinical Trials (Phases I-III)	6-7 years including patient recruitment	Patient recruitment accelerated by AI matching; can save 1-2 years
Regulatory Approval	1-2 years of document review	AI-assisted document preparation can save months

Pharmaceutical — My Key Observations

- Moderna designing the vaccine's molecular structure in just 2 days after the virus genome was published is the single most impressive speed statistic I found across all four parts of my research
- This industry has the clearest 'AI saves lives' narrative of anywhere I researched — the COVID vaccine timeline compression from 10+ years to under a year had a direct, measurable impact on global health
- Adverse event detection (catching drug side effects early using AI that monitors millions of patient reports) is a use case I had not thought about before, and it shows AI's value extends beyond just the discovery phase into ongoing patient safety
- Drug repurposing — using AI to find new uses for already-approved drugs — is a clever idea because the drugs are already proven safe, so this can be much faster and cheaper than starting from scratch
- The pharmaceutical industry is heavily regulated for good reason (patient safety), so even with AI speeding things up, human clinical testing and regulatory review remain essential safeguards that AI does not replace

3. Government & Public Sector

Government was a section I approached with some scepticism at first, because I generally think of government services as slow-moving and resistant to new technology. What I found surprised me — some governments are actually becoming leaders in practical AI adoption, especially in countries that built their digital systems more recently and could design them with AI in mind from the start.

Government AI is different from corporate AI in one critical way: governments serve everyone, including people who cannot pay, do not have smartphones, or do not trust technology. This means government AI has to be more careful about fairness and accessibility than a private company's AI might need to be. The OECD estimates that AI could help governments save over $1.75 trillion globally through efficiency improvements.

3.1 Use Case Overview

Use Case	What GenAI Does	Real-World Impact	Company / Tool
AI Citizen Service Chatbots	GenAI chatbots answer citizen questions about government services, benefits, taxes, and procedures in plain language, 24/7.	Citizens get instant answers without waiting in line or on hold; government call centres handle fewer repetitive queries.	*Estonia's Bürokratt, UK GOV.UK Chat*
Fraud Detection in Public Benefits	AI analyses welfare, unemployment, and tax claims to detect fraudulent applications while avoiding wrongly flagging genuine claimants.	Billions in fraudulent claims prevented annually across various government programmes.	*UK HMRC AI, US IRS AI Compliance*
Automated Document Processing	AI processes passport applications, visa requests, permits, and licences by extracting and verifying information from submitted documents.	Processing times reduced from weeks to days; fewer manual errors in document review.	*USCIS AI tools, UK Home Office*
Predictive Policy Analysis	AI models the likely impact of proposed policies (tax changes, infrastructure spending) before they are implemented, using economic and social data.	Policymakers make more informed decisions; unintended consequences identified before rollout.	*Singapore's GovTech, World Bank AI*
Smart City Traffic & Infrastructure Management	AI optimises traffic signals, public transport schedules, and energy grids across entire cities in real time.	Traffic congestion reduced; public transport more reliable; city energy use optimised.	*Singapore Smart Nation, Barcelona Smart City*
Disaster Response Coordination	AI analyses satellite imagery, weather data, and emergency reports to coordinate disaster response and resource allocation in real time.	Faster, more effective emergency response; resources directed where most needed.	*FEMA AI tools, UN OCHA*

Use Case	What GenAI Does	Real-World Impact	Company / Tool
Legislative Document Summarisation	AI summarises lengthy proposed laws and regulations into plain language so citizens and even legislators can understand them quickly.	Greater public understanding and participation in the legislative process.	*GovTrack AI summaries, Parliament AI tools (UK)*
Public Health Surveillance	AI monitors disease outbreak patterns, hospital data, and even social media to detect public health threats earlier than traditional surveillance.	Earlier outbreak detection allows faster public health response, potentially preventing wider spread.	*CDC AI surveillance, WHO EIOS*

3.2 Case Study: Estonia — The Most Digital Government in the World

Case Study: Estonia's Bürokratt — One AI Assistant for the Entire Country

Problem: Estonia, a small country of about 1.3 million people, wanted to make government services genuinely effortless for citizens. The traditional approach of having separate websites, forms, and phone lines for every single government department (taxes, healthcare, business registration, education) was confusing and inefficient, with citizens often unsure where to even start for a given request.

What They Did: Estonia built Bürokratt, a single, unified AI chatbot assistant that acts as the front door to all government services. Instead of citizens needing to know which of dozens of departments to contact, they simply tell Bürokratt what they need in plain language — 'I just had a baby and need to register for child benefits' — and the AI understands the request, pulls the right information from the right government databases, and either completes the action automatically or guides the citizen through what is needed.

Result: Estonia is now widely recognised as having the most digitally advanced government in the world. Over 99% of government services in Estonia are available online, and Bürokratt has dramatically reduced the time citizens spend navigating bureaucracy. Estonia's e-Residency programme, built on this same digital infrastructure, has attracted entrepreneurs from over 170 countries to register businesses in Estonia entirely online. Estonia saves an estimated 2% of its GDP annually through digital government efficiency, according to government officials.

What I Learned from This: Estonia taught me that a country does not need to be large or wealthy to lead in government AI — it needs the right vision and willingness to redesign systems from scratch rather than just digitising old paper processes. The idea of 'one AI assistant for an entire country's bureaucracy' is something I had never considered, and it made me think about how government services in countries like India, with much larger and more complex bureaucracies, could eventually move toward similar simplicity.

3.3 Why Government AI Is Different from Corporate AI

Researching this section made me realise that government AI faces unique pressures that private companies do not face in the same way:

Challenge Unique to Government AI	Why This Matters Specifically for Government
Must serve 100% of the population, not just paying customers	A company can choose its customers; a government cannot exclude any citizen, including the elderly, disabled, or those without internet access
Held to a higher standard of fairness and transparency	Government decisions (benefits, taxes, legal matters) require explainability since citizens have a legal right to understand decisions affecting them
Political accountability and public trust	A government AI failure becomes a public scandal and can damage trust in democratic institutions themselves, not just one company's reputation
Extremely sensitive data (health records, criminal records, immigration status)	A government data breach can affect entire populations and has national security implications
Cannot simply 'move fast and break things'	Mistakes in government systems can deny someone healthcare, welfare, or legal rights — the cost of errors is much higher than in most commercial contexts

Government — My Key Observations

- Estonia genuinely surprised me — I did not expect a small country to be the global leader in practical government AI, and it taught me that good design matters more than country size or wealth
- The fairness and accessibility requirements in government AI are stricter than almost any other industry I researched, because government cannot 'choose its customers' the way a business can
- Fraud detection in public benefits is a delicate balance — being too aggressive risks wrongly denying genuine citizens their benefits, while being too lenient allows fraud to continue. I noticed this tension comes up in several of the case studies I read.
- Smart city traffic management connects back to what I learned about Energy & Utilities AI in Part 02 — many 'smart city' systems are really energy and infrastructure AI applied at a city-wide scale
- I think government AI adoption in India could be transformative given the scale of public services involved (over 1.4 billion citizens), but it would need very careful design to ensure accessibility for citizens without smartphones or reliable internet

4. Insurance

Insurance was a section I went into with very little prior knowledge — I knew insurance companies handle claims and policies, but I did not appreciate how much of the entire industry is fundamentally about predicting risk and processing paperwork, both of which are extremely well-suited to GenAI.

The insurance industry essentially makes its money by accurately predicting the likelihood of bad things happening (accidents, illness, death, fire, theft) and pricing policies accordingly. Historically, this prediction relied on broad statistical categories (age, location, driving record). GenAI is enabling much more precise, individualised risk assessment, while also dramatically speeding up the claims process that frustrates so many customers. The global AI in insurance market is projected to reach $45.7 billion by 2031.

4.1 Use Case Overview

Use Case	What GenAI Does	Real-World Impact	Company / Tool
Instant Claims Processing	AI reviews photos, videos, and documentation submitted for a claim (like car damage or property loss) and approves or processes payment automatically.	Claims that used to take days or weeks can now be settled in seconds to minutes for straightforward cases.	*Lemonade, Tractable*
AI Underwriting & Risk Pricing	AI analyses vast amounts of data (driving behaviour, health records, property data) to price insurance policies more accurately for each individual.	More accurate pricing means lower premiums for low-risk customers and better risk management for insurers.	*Zest AI, Cape Analytics*
Fraud Detection in Claims	AI identifies suspicious patterns in claims (staged accidents, exaggerated damage, duplicate claims) that human reviewers might miss.	Insurance fraud, estimated at over $40 billion annually in the US alone, is significantly reduced.	*Shift Technology, FRISS*
AI Damage Assessment from Photos	Customers photograph car or property damage, and AI estimates repair costs instantly by analysing the images.	Removes the need for in-person inspectors for straightforward claims, speeding up the entire process.	*Tractable, CCC Intelligent Solutions*
Personalised Policy Recommendations	AI analyses a customer's specific situation and recommends the right type and amount of insurance coverage for their needs.	Customers get coverage that actually matches their risk profile instead of generic one-size-fits-all policies.	*Policygenius AI, Lemonade*
Automated Policy Document Generation	GenAI drafts insurance policy documents, ensuring they are legally accurate and clearly explain coverage terms to customers.	Reduces errors in policy wording; makes complex insurance terms more understandable for customers.	*Sure, Hippo Insurance*

Use Case	What GenAI Does	Real-World Impact	Company / Tool
Customer Service AI Chatbots	AI chatbots handle policy questions, claims status updates, and basic account changes without requiring a human agent.	Faster response times for customers; insurance companies reduce call centre costs significantly.	*GEICO's Kate, Allstate's ABIe*
Catastrophe & Climate Risk Modelling	AI analyses climate data, satellite imagery, and historical disaster patterns to predict and price risks from floods, fires, and storms more accurately.	Insurers can better prepare for climate-related claims; helps identify areas becoming too risky to insure.	*Jupiter Intelligence, One Concern*

4.2 Case Study: Lemonade — Settling Insurance Claims in 3 Seconds

📌 Case Study: Lemonade Insurance — The World Record for Fastest Claim Settlement

Problem: Traditional insurance claims are notoriously slow and frustrating for customers. A typical home or renters insurance claim can take days or weeks to process, requiring phone calls, paperwork, and waiting for a human adjuster to review the case. This delay adds stress to situations that are often already stressful, like a burglary or fire.

What They Did: Lemonade, a digital insurance company, built an AI claims system called 'AI Jim.' When a customer files a claim through the Lemonade app, they describe what happened and submit photos or video evidence. AI Jim analyses the claim against the policy terms, cross-checks for signs of fraud (comparing against historical fraud patterns and checking for inconsistencies), and if everything checks out, can approve and pay the claim automatically — without any human ever reviewing it.

Result: In 2017, Lemonade set a record by processing and paying a claim in just 3 seconds — from the moment the customer reported a stolen item to the money arriving in their bank account. Today, Lemonade reports that around 30% of all claims are handled entirely by AI with no human involvement, and the company has dramatically lower operating costs than traditional insurers as a result. Lemonade has grown to serve over 2 million customers since its founding in 2015, largely on the strength of this AI-driven, fast-claims model.

What I Learned from This: The 3-second claim is an extreme example, but it illustrates something important: AI does not just make existing processes a bit faster — in some cases it can make them almost instantaneous, fundamentally changing the customer experience. It also made me think about the fraud detection challenge: if claims are approved this fast, how does the AI avoid being tricked by fraudulent claims? Lemonade's answer is that the AI is specifically trained to flag suspicious patterns and only auto-approves claims that pass multiple fraud checks.

4.3 How AI Is Changing Claims and Underwriting

To understand insurance properly, I had to learn the difference between the two core insurance functions, and how AI affects each one differently:

Insurance Function	How GenAI Is Changing It
Underwriting (deciding what to charge for a policy)	AI analyses far more data points than a human underwriter could, leading to more individually accurate pricing — but also raising fairness questions if certain factors create unintended discrimination
Claims (paying out when something bad happens)	AI can process straightforward claims almost instantly using photo/video analysis, while still routing complex or suspicious claims to human adjusters
Fraud Prevention	AI detects subtle fraud patterns across millions of claims that would be invisible to any individual human reviewer checking one claim at a time
Customer Communication	AI chatbots handle routine policy questions instantly, while human agents focus on complex situations requiring empathy and judgment

Insurance — My Key Observations

- Lemonade's 3-second claim settlement is one of the most dramatic 'before and after' AI statistics I found in my entire research — going from weeks to seconds is an almost unbelievable transformation
- I had never thought about how insurance fraud (estimated at over $40 billion annually in the US) creates a genuine need for sophisticated AI fraud detection — this connects to similar fraud detection patterns I saw in Banking (Part 01) and Telecom (Part 02)
- AI underwriting raises a fairness question I had not considered before: if AI can analyse extremely granular data about a person, could it end up discriminating in ways that are technically not about a protected category but still unfair? This feels similar to the bias issues I researched in HR AI in Part 03.
- Climate risk modelling stood out to me because it connects insurance directly to climate change — as extreme weather increases, AI is helping insurers (and by extension, society) understand which areas are becoming too risky to live in without major adaptation
- I think insurance is an industry where customers genuinely benefit from AI in a very direct, personal way — faster claims when something bad has happened to you is something everyone can appreciate

5. Aerospace & Defense

This was the final and, in many ways, the most thought-provoking section of my entire four-part research project. Aerospace and Defense involves some of the most advanced engineering in the world, and it is also the industry where the ethical questions around AI are the most serious — because the stakes can literally be life and death on a national scale.

I want to be upfront that I approached the 'Defense' half of this section carefully, focusing on documented, publicly available information about how AI is used for design, logistics, and decision support, while also spending real time on the ethical debate around autonomous weapons, which I think is one of the most important conversations happening in AI today.

The global AI in aerospace and defense market is projected to reach $34.9 billion by 2030. This industry has used basic automation for decades (autopilot systems have existed since the 1970s), but GenAI is now enabling design, logistics, and decision-support capabilities that were not previously possible.

5.1 Use Case Overview

Use Case	What GenAI Does	Real-World Impact	Company / Tool
AI-Generated Spacecraft & Aircraft Component Design	GenAI generates thousands of design variations for aircraft and spacecraft parts, optimised for weight, strength, and aerodynamics.	Components are often 20-40% lighter than human-designed equivalents while maintaining strength, directly reducing fuel costs.	*NASA, Airbus, Lockheed Martin*
Predictive Maintenance for Aircraft	AI analyses sensor data from aircraft engines and systems to predict maintenance needs before a failure occurs.	Unscheduled maintenance and flight delays reduced; safety improved through earlier problem detection.	*GE Aviation, Boeing AnalytX*
Satellite Image Analysis	AI processes satellite imagery to identify objects, changes, and patterns — used for everything from disaster monitoring to security applications.	Analysis that used to take analysts days can now be done in minutes, across far more images.	*Planet Labs, Maxar Technologies*
Flight Path & Air Traffic Optimisation	AI optimises flight routes in real time based on weather, air traffic, and fuel efficiency considerations.	Fuel savings of 5-10% per flight; reduced delays from better traffic flow management.	*NASA's ATD-2, FAA AI initiatives*
Mission Planning & Logistics Simulation	AI simulates complex logistics scenarios (supply routes, equipment deployment, personnel scheduling) to optimise planning.	Faster, more efficient planning for complex operations involving thousands of variables.	*Palantir, Anduril*
Space Mission Autonomous Systems	AI enables spacecraft to make some decisions independently when communication delays with	Essential for deep space missions where a signal to Mars takes 4-24 minutes one way, making	*NASA JPL, SpaceX*

Use Case	What GenAI Does	Real-World Impact	Company / Tool
	Earth make real-time human control impossible.	real-time control impossible.	
Document & Intelligence Summarisation	AI summarises large volumes of reports, sensor data, and open-source information to help analysts understand situations faster.	Analysts can process far more information and focus their expertise on judgment rather than reading volume.	*Palantir Gotham, various government systems*
Training Simulation Environments	GenAI creates realistic, varied training scenarios for pilots and military personnel without the cost and risk of real-world exercises.	Training becomes safer, cheaper, and can cover more scenarios, including rare or dangerous situations.	*CAE, Bohemia Interactive Simulations*

5.2 Case Study: NASA — Generative Design for Lighter Spacecraft

📌 Case Study: NASA — AI Designs a Lunar Lander Component That No Human Would Have Thought Of

Problem: Every kilogram of weight on a spacecraft costs enormous amounts of money to launch into space — sending one kilogram to the Moon can cost tens of thousands of dollars. Engineers have always tried to make spacecraft components as light as possible while keeping them strong enough to survive the extreme forces of launch and space travel. Traditional design, even by skilled engineers, has practical limits because humans tend to design in familiar shapes and patterns.

What They Did: NASA used generative design AI (working with tools similar to Autodesk's generative design software) to design a lightweight bracket for a lunar lander. Instead of an engineer drawing a traditional bracket shape, the team gave the AI the requirements — the forces the part needed to withstand, the materials available, and the space it needed to fit into — and let the AI generate thousands of possible designs. The AI explored shapes that no human engineer would typically consider, often appearing organic or skeletal rather than the blocky, geometric shapes traditional engineering produces.

Result: The AI-generated component was 30% lighter than the traditional human-designed equivalent, while meeting all the same strength requirements. Across an entire spacecraft, these kinds of weight savings compound significantly, translating into either reduced launch costs or the ability to carry more scientific equipment for the same budget. NASA has continued to expand its use of generative design across multiple spacecraft and lunar mission components.

What I Learned from This: This case study connected directly back to what I learned about Autodesk's Spacemaker in the Real Estate section of Part 03 — the same generative design technology that designs better buildings is also designing better spacecraft parts. It reinforced something I am noticing as a theme across my whole research project: the same core AI capabilities show up again and again across completely different industries, just applied to different problems.

5.3 The Ethics of Military AI — What I Think Is Important to Understand

I want to dedicate proper space to this topic because I believe it might be the most important ethical question in all of AI, not just in this one industry. The development of autonomous weapons systems — AI that can identify and engage targets with reduced or no human decision-making in the moment — is one of the most actively debated topics in international policy today.

Many AI researchers, including prominent figures in the field, have signed open letters calling for international regulation of autonomous weapons, similar to existing treaties on chemical and biological weapons. The core concern is about removing meaningful human judgment from life-and-death decisions, and the risk of an AI arms race between nations that could make conflicts faster and harder to control.

At the same time, defense organisations argue that AI can also make military operations more precise and potentially reduce civilian casualties compared to less accurate alternatives, and that responsible nations need to develop these capabilities to deter adversaries who may develop them without the same ethical constraints. This is a genuinely difficult and unresolved global debate, and I do not think it is my place in a research document to declare one side correct — but I believe every student studying AI should be aware that this conversation exists and understand both perspectives.

Argument For Caution / Regulation	Argument For Continued Development
Removing human judgment from lethal decisions raises serious moral concerns	Well-designed AI systems could be more precise than human soldiers under stress, potentially reducing civilian casualties
Risk of an uncontrolled AI arms race between nations, escalating conflicts faster than humans can respond	Responsible democratic nations may need these capabilities to deter authoritarian adversaries who will develop them regardless
AI systems could malfunction or be hacked, with catastrophic and irreversible consequences	Many current applications are about logistics, intelligence, and support — not just autonomous weapons — and have clear humanitarian value
International treaties (similar to those banning chemical weapons) are difficult to enforce once technology exists	International cooperation and treaties on AI weapons are being actively discussed at the United Nations and other forums

Aerospace & Defense — My Key Observations

- NASA's 30% lighter spacecraft component was a genuinely exciting example of generative design, and it directly connected back to what I learned about Autodesk Spacemaker in Part 03 — the same core technology, applied to a completely different industry
- I found the deep space autonomous systems use case fascinating — the fact that a signal to Mars takes up to 24 minutes one-way means AI has to make some decisions on its own, since waiting for human instructions simply is not possible in real time
- The ethics of military AI is, in my opinion, one of the most important unresolved questions in the entire field of artificial intelligence, and I made a deliberate choice to present both sides of the argument rather than take a position, since I do not think this is a question with an easy or obvious answer
- Researching this section made me appreciate that the same underlying AI capability (like generative design, or pattern recognition in sensor data) can be used for purposes that

range from clearly beneficial (lighter, cheaper spacecraft) to genuinely difficult ethical territory (autonomous weapons)

- This section, more than any other in my four-part research project, made me think about the responsibility that comes with developing powerful AI technology — and that as future engineers and researchers, we need to think not just about what AI can do, but what it should do

6. Patterns I Noticed Across All 5 Industries

With Part 04 complete, I have now researched 20 industries across the full four-part series. This final patterns section looks specifically at what I noticed in this last group of five industries, plus a complete comparison across everything I have studied.

6.1 Same Patterns, New Settings

Pattern: High-Stakes Decision Support

Every industry in Part 04 deals with extremely high-stakes decisions: a cybersecurity AI deciding whether to block a connection, a pharmaceutical AI predicting a drug's safety, a government AI determining benefit eligibility, an insurance AI approving or denying a claim, and a defense AI assisting mission planning. In every one of these cases, I noticed that the most responsible implementations keep a human in the loop for the final, highest-stakes decisions, even while AI handles the analysis and recommendation.

Pattern: Speed as a Matter of Life and Safety

In Parts 01-03, speed mattered mostly for efficiency and cost. In Part 04, speed often matters for safety and survival: faster cybersecurity threat detection prevents bigger breaches, faster vaccine design saves lives during a pandemic, faster disaster response in government applications saves lives, and faster insurance claims help people recover from emergencies faster. This part of my research showed me a more serious dimension of why AI speed matters.

Pattern: The Recurring Technology, New Applications

I noticed that several specific AI techniques I learned about earlier in my research kept reappearing in Part 04 in new contexts. Generative design (which I first learned about with Autodesk Spacemaker for buildings in Part 03) reappeared with NASA designing spacecraft parts. Fraud detection patterns (which I first studied in Banking in Part 01) reappeared in both Insurance and Government benefits fraud. This made me realise that the number of distinct 'core AI capabilities' is actually smaller than the number of industries — the same techniques get creatively reapplied again and again.

6.2 Comparing All 20 Industries Across the Full Series

Having now completed all four parts of this research project, here is my complete summary of the primary GenAI value driver in each of the 20 industries I studied:

Industry (Part)	Primary GenAI Value Driver
Healthcare (Part 01)	Reducing administrative burden; accelerating drug discovery
Banking (Part 01)	Fraud prevention; compliance automation; personalised finance
Manufacturing (Part 01)	Predictive maintenance; quality inspection; digital twins
Retail (Part 01)	Personalisation; demand forecasting; virtual try-on

Industry (Part)	Primary GenAI Value Driver
IT & Software (Part 01)	Developer productivity; code quality; infrastructure automation
Telecommunications (Part 02)	Network reliability; customer service automation; fraud detection
Edtech (Part 02)	Personalised learning at scale; automated grading; 24/7 tutoring
Media & Entertainment (Part 02)	Content personalisation; production cost reduction; global localisation
Automotive & Mobility (Part 02)	Autonomous driving; safety improvement; synthetic training data
Energy & Utilities (Part 02)	Grid management; renewable integration; climate impact reduction
Agriculture & Agritech (Part 03)	Crop health; yield prediction; expert access for smallholder farmers
Logistics & Supply Chain (Part 03)	Route optimisation; demand forecasting; warehouse automation
HR & Talent Management (Part 03)	Hiring efficiency; bias reduction; employee retention prediction
Legal & Compliance (Part 03)	Document review; research acceleration; access to legal knowledge
Real Estate & Construction (Part 03)	Property valuation; generative design; construction risk prediction
Cybersecurity (Part 04)	Threat detection speed; closing the security talent gap
Pharmaceutical Industry (Part 04)	Drug discovery speed; vaccine design; patient safety monitoring
Government & Public Sector (Part 04)	Citizen service efficiency; fraud prevention; accessible bureaucracy
Insurance (Part 04)	Instant claims processing; precise risk pricing; fraud detection
Aerospace & Defense (Part 04)	Lightweight design optimisation; mission planning; safety-critical decision support

7. Unique Challenges in Part 04 Industries

Each industry in Part 04 has specific challenges that go beyond the general AI challenges (data quality, hallucination, explainability) I documented back in Part 01. Given how high-stakes these five industries are, I found the challenges here particularly serious.

7.1 Cybersecurity: The Arms Race Problem

Unlike most industries where AI adoption simply makes things better, cybersecurity faces a unique structural problem: every improvement in defensive AI tends to be matched by an improvement in attacker AI. This means the industry can never simply 'solve' the problem and move on — it requires continuous, ongoing investment and innovation, forever. This is different from, say, agriculture, where a good AI tool for crop disease detection stays useful indefinitely.

7.2 Pharmaceutical: Balancing Speed with Patient Safety

While AI is dramatically speeding up drug development, the industry has to be extremely careful not to let speed compromise safety. The Moderna case study is a success story, but it is important to understand that human clinical trials and regulatory review were still required even with AI assistance — AI accelerated the early design stages, but did not replace the rigorous safety testing process required before a drug reaches the public.

7.3 Government: Public Trust and the Digital Divide

Government AI faces a unique legitimacy challenge: citizens did not choose to interact with their government the way they might choose to use a particular AI shopping assistant. If a citizen does not trust or cannot access government AI systems, they may be effectively denied access to services they are entitled to. Governments must ensure AI systems work for citizens without smartphones, reliable internet, or strong digital literacy.

7.4 Insurance: The Fairness of Individualised Pricing

As AI enables more individualised risk pricing, a difficult question emerges: is it fair for AI to charge people different premiums based on extremely granular personal data? Traditional insurance worked by pooling risk across large groups of people. Hyper-individualised AI pricing could mean that people in genuinely high-risk situations up priced out of insurance entirely — undermining the basic social purpose of insurance as a safety net.

7.5 Aerospace & Defense: The Accountability Gap in Autonomous Systems

As I discussed in the ethics section, one of the most serious unresolved challenges in defense AI is accountability: when an autonomous or AI-assisted system makes an error with serious consequences, who is responsible? This question does not have a clean answer yet, and it is part of why so many AI researchers and policy experts are calling for international frameworks to govern how much autonomy AI systems should be given in safety-critical and defense contexts.

8. Conclusion & Personal Reflections — End of the Series

This marks the end of my four-part GenAI research project, covering 20 industries over four weeks at MacroEdtech. Looking back at this entire journey — from Healthcare and Banking in Part 01, through Telecom and Energy in Part 02, Agriculture and Legal in Part 03, to Cybersecurity and Defense here in Part 04 — I feel like I have built a genuinely comprehensive understanding of how Generative AI is reshaping the global economy.

If I had to summarise the entire four-part project in one sentence, it would be this: Generative AI is valuable wherever there is a large volume of data, language, or visual information that needs to be understood and acted upon faster and more consistently than humans alone can manage — and the higher the stakes of getting it right, the more carefully that AI needs to be designed, tested, and supervised by humans.

8.1 My Top Learnings from Part 04

1. Moderna's 2-day vaccine design after the COVID genome was published is, in my opinion, the single most impressive statistic across my entire four-part research project — it shows AI's potential to directly and immediately save human lives.
2. Cybersecurity taught me something genuinely new: that AI can be a tool for both attack and defence in the same domain. This 'dual-use' nature of AI is something I will think about whenever I study any new AI application going forward.
3. Estonia's Bürokratt reframed how I think about government services — proving that a country does not need to be large to be a global leader in practical, citizen-friendly AI adoption.
4. The ethics of military AI is the most serious unresolved question I encountered across my entire research project. I do not think there is a clean answer, but I believe it is critically important for anyone working in AI to understand both sides of this debate.
5. Looking back at all 20 industries, I now see clearly that the same small number of core AI capabilities (document automation, prediction, personalisation, generative design, fraud detection) get applied again and again across completely different sectors. Understanding these core capabilities deeply is more valuable than memorising industry-specific facts.

8.2 Reflecting on the Full Four-Part Journey

When I started Part 01 back in Week 01, I genuinely thought Generative AI was mostly about chatbots and image generators. Twenty industries later, I now understand that GenAI is a foundational technology being woven into nearly every sector of the global economy — from the food on our tables (Agriculture) to the medicines that keep us healthy (Pharmaceutical) to the rockets that explore space (Aerospace). This has been, by far, the most comprehensive and eye-opening research project I have undertaken.

I am especially grateful for the structure of this internship — being asked to find real case studies rather than just listing abstract use cases pushed me to engage much more deeply with how these technologies actually work in practice, and to think critically about both their benefits and their risks.

8.3 What I Want to Explore Beyond This Internship

- How to actually build and fine-tune a domain-specific AI model, rather than just studying how companies use existing ones
- The technical details of how Retrieval-Augmented Generation (RAG) systems are built, since this seems relevant to nearly every industry I have studied
- How AI governance frameworks are being developed at the international level, especially given what I learned about defense AI ethics
- How data engineering practices (my own area of internship focus at Azentrix) support and enable all of these GenAI applications behind the scenes

Final Summary — Part 04 at a Glance

- 5 industries covered: Cybersecurity, Pharmaceutical Industry, Government & Public Sector, Insurance, Aerospace & Defense
- 40+ use cases documented across all five sectors
- 5 detailed case studies: Microsoft Security Copilot, Moderna mRNA Vaccine Design, Estonia's Bürokratt, Lemonade's 3-Second Claim, NASA Generative Design
- Key pattern: High-stakes decision support with humans kept in the loop for final decisions
- Most impressive case study: Moderna designing a vaccine candidate in 2 days, helping compress a typical 10+ year process into under a year
- Most important ethical discussion: The accountability and human-control debate around autonomous defense AI systems
- 20 industries now covered across the complete four-part GenAI research series

Appendix A: New Terms I Learned in Part 04

These are the key technical and industry-specific terms I encountered for the first time in Part 04:

Term	Simple Definition in My Own Words
mRNA Vaccine Technology	A vaccine approach that teaches cells to produce a harmless piece of the virus, training the immune system to recognise and fight the real virus. AI helped design the mRNA sequence used in COVID vaccines.
Nation-State Hackers	Highly sophisticated cyberattackers backed by a country's government, typically far more resourced and dangerous than individual criminal hackers.
Autonomous Weapons Systems	Military systems that can identify and engage targets with reduced or no real-time human decision-making — a major topic of international ethical debate.
Generative Design (Aerospace Context)	Using AI to generate many possible engineering designs based on requirements (strength, weight, materials), often producing shapes humans would not typically consider.
Underwriting	The insurance industry process of deciding how much risk a person or asset represents, and pricing a policy accordingly.
Adverse Event (Pharma Context)	An unexpected or harmful side effect caused by a drug or medical treatment, which pharmaceutical companies are legally required to monitor and report.
Smart City	An urban area that uses connected sensors and AI to optimise services like traffic, energy use, and public transport in real time.
Deepfake Detection	AI technology specifically designed to identify AI-generated fake videos, images, or audio — increasingly important in cybersecurity to prevent fraud.
Digital Divide	The gap between people who have easy access to digital technology (smartphones, internet) and those who do not — an important consideration in designing fair government AI systems.
Patient Recruitment (Clinical Trials)	The process of finding and enrolling suitable patients for a clinical trial, which AI can significantly speed up by matching patient records to trial requirements.
Threat Intelligence	Information about current and emerging cybersecurity threats, often gathered from

Term	Simple Definition in My Own Words
	massive amounts of global security data and analysed by AI.
E-Residency	Estonia's programme allowing entrepreneurs from any country to register and run a business entirely online, enabled by Estonia's advanced digital government infrastructure.

References

All sources consulted in preparing this Part 04 research document are listed below in APA 7th Edition format.

Academic & Research Papers

[1] Sahin, U., Karikó, K., & Türeci, Ö. (2014). mRNA-based therapeutics — developing a new class of drugs. Nature Reviews Drug Discovery, 13(10), 759–780. https://doi.org/10.1038/nrd4278

[2] Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. arXiv preprint arXiv:1802.07228. https://arxiv.org/abs/1802.07228

[3] Scharre, P. (2018). Army of none: Autonomous weapons and the future of war. W.W. Norton & Company.

[4] Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. Philosophical Transactions of the Royal Society A, 376(2133). https://doi.org/10.1098/rsta.2018.0080

Industry Reports & Market Research

[5] Allied Market Research. (2023). AI in insurance market by component, application, deployment mode, and enterprise size: Global opportunity analysis and industry forecast, 2023–2031. https://www.alliedmarketresearch.com/artificial-intelligence-ai-in-insurance-market

[6] OECD. (2023). The potential of artificial intelligence in government. OECD Public Governance Reviews. https://www.oecd.org/governance/digital-government/

[7] ISC2. (2023). ISC2 cybersecurity workforce study 2023. International Information System Security Certification Consortium. https://www.isc2.org/Research/Workforce-Study

Company & Product Sources

[8] Microsoft. (2023). Introducing Microsoft Security Copilot: Empowering defenders at the speed of AI. Microsoft Security Blog. https://www.microsoft.com/en-us/security/blog/2023/03/28/introducing-microsoft-security-copilot/

[9] Moderna. (2021). The Moderna mRNA technology platform: From concept to clinic. Moderna Therapeutics. https://www.modernatx.com/research/development-track-record

[10] Government of Estonia. (2023). Bürokratt: Estonia's AI-powered virtual assistant network. e-Estonia Briefing Centre. https://e-estonia.com/solutions/artificial-intelligence/burokratt/

[11] Lemonade Inc. (2022). How AI Jim settles claims in seconds. Lemonade Insurance Blog. https://www.lemonade.com/blog/lemonade-sets-new-world-record/

News Articles & Other Sources

[12] Future of Life Institute. (2023). Autonomous weapons: An open letter from AI and robotics researchers. Future of Life Institute. https://futureoflife.org/open-letter/autonomous-weapons-open-letter/

[13] Reuters. (2023). Microsoft says its AI tool helped detect nation-state hacking attempts faster. Reuters Technology. https://www.reuters.com/technology

[14] Insurance Business Magazine. (2023). How Lemonade's AI is reshaping claims processing. Insurance Business Magazine. https://www.insurancebusinessmag.com/us/news/technology