



How AI is Automating Industrial Inspection:

Reducing Costs, Turnaround Time and Risk

Introduction

Artificial Intelligence (AI) is at the center of controversies around the theme of jobs and employment. Will AI take away all our jobs, leaving us unemployed? While some jobs will indeed change or disappear – as they have always and will always throughout history to the present and into the future – AI can be a human force multiplier in business applications where we, as a species, are still under- delivering.

One of these recently introduced set of jobs potentially “threatened” to be greatly improved by AI consists of collecting, storing and analyzing terabytes of drone footage (from conventional RGB cameras to LIDAR and other sensors), portraying assets ranging from telecommunication towers to pipelines, solar panels, bridges, and other structures, where gathering data can be treacherous.

In other words, these are drone-bound inspections, where the drone is essentially a payload carrying machine, flying around with vision sensors, versus a human inspector asked to climb a wind turbine, for example.

While flying drones may have some limited recreational value, the repetitive operation of flying a drone, and analyzing the copious data collected by it, is a job prone to fatigue and errors. Pilots aside, analysts tasked with finding a crack in the 351st insulator inspected during the 25th drone mission of the day may not appeal very much as an activity. If cracks on insulators are not all you live for, chances are you can make a fatal mistake and lose money to your enterprise.

That is where AI comes in: software emulations of human nervous systems that can be trained to emulate, and partially offset, dull and repetitive tasks that human analysts need to perform relentlessly on huge datasets, day after day, where the drone “borrows” the human brain to deliver its final value to the enterprise.

After initial learning, there are at least three ways in which AI can be deployed in the inspection workflow to make them more efficient and precise and extract more economic value out of every analytics flight. In this white paper, we will discuss:

- 1 Edge AI: biased analytics
- 2 Cloud AI: analyzing data for insights
- 3 Predictive AI: knowing when to act

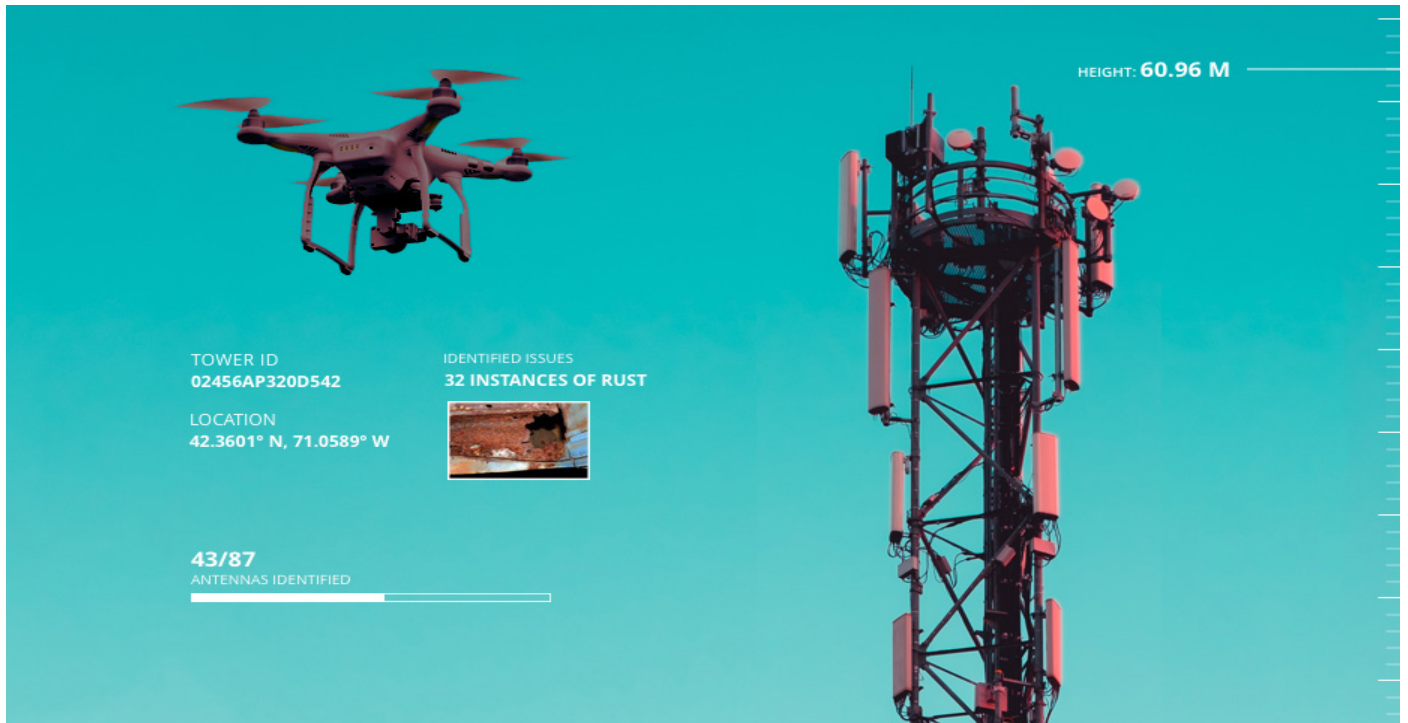
In each application domain, AI radically reduces costs and increases throughput.



Edge AI: biased analytics

Think about how humans inspect a structure: they go from cursory looks to the overall structure to more thoughtful and prolonged analysis of critical areas, which they know from experience or visual inspection are areas of interest. In a sense, humans are biased. These are good biases to have, driving these quick decisions on how to use, or not to use, your precious compute cycles.

On the other hand, think now how drones collect data today: they are flown over a structure in a way that is almost always independent of the actual data collected. In other words, a drone may spend the same three seconds collecting data from a normal area in a structure, from a tree around a structure, or a damaged area. Unless the operator has immediate visual feedback on what the drone “perceives,” the drone will collect its data in an unbiased - instead of biased - way. The result may be the need to collect additional data post-flight.



AI at the compute edge, ideally on the drone itself, can bias data collection by directing the drone to the pertinent structure, then sub-areas in the structure, and then specific, probable anomalies. This can go as far as increasing the time spent collecting data from interesting areas to engaging other sensors and alerting operators for additional actions.

Fortunately, today's drones are built with enough compute power to accommodate edge intelligence: from powerful Snapdragon processors to even beefier NVIDIA GPUs. These processors can all be substrates for real-time operating AI to power this important mission-critical ability. For example, using Neurala's AI, a brain can learn quickly for deployment, and continue to learn while analyzing infrastructures and other assets.




Cloud AI: analyzing data for insights

Once data is collected, ideally good, quality data driven by an intelligent, AI-powered drone and deployed, the hard analytics task begins.

Analytics tasks vary between inventory of structures – namely, finding what is on that structure – and mapping the visual appearance of an object (e.g., an insulator) to more detailed information, such as model and year of production, or to finding defects on that object, including corrosion and damage.

Each mission typically collects thousands of still images or high-definition videos containing a few hundreds of milliseconds to many seconds or minutes per frame. Every time a human analyzes a frame, it has not only an economic cost but also a certain probability of error associated with it.



Given the large number of frames collected and the potentially multiple layers of information, including RGB and LIDAR, the possibility that the human analyst misses a key frame, where the specific item appears or the defect is most visible, increases.

Unlike humans, AI does not experience fatigue or complain about performing this continuous, unending job. AI systems can be trained to perform multiple tasks, from inventory to classifying a 3D image and finding finer details in that object, such as tiny defects. Of course, performing these tasks successfully is predicated on good AI, trained with good data, which is not a given.



Predictive AI: knowing when to act

One of the reasons structures are inspected is for maintenance. Being able to detect a crack or corrosion will tell the analyst that, in the next 6 to 12 months that component may fail, and maintenance should be planned to avoid downtime.

This requires additional, purpose-driven AI, which is able to take into account historical data and relate this data to specific outcomes, with a temporal flavor. Similar to a medical doctor who remarks, "If you keep eating a box of french fries per day, you may gain a few pounds here and there," predictive AI is able to diagnose future states of a system based on current and past data. Applications of AI to temporal predictions abound, from medical data to financial time series and network attacks.

AI will find a stable job in this department as well.

I want AI: how do I get it?

The data problem

We are now convinced that AI is a game changer in drone-bound inspections. All drones and software backends that have to do with the collection and analysis of drone data should have AI working under the hood. Now, how does one get such AI?



The first step is building. The branch of AI systems that is widely used for these purposes, neural networks or deep learning systems, requires a massive amount of data for learning. Unfortunately, this data is not always available. Even worse, when it is available, even adding a simple additional example to the dataset implies retraining the entire AI from scratch, a process that can take anywhere from a few hours to days or weeks.

The next step is deployment. Processors are still being optimized and developed to handle the processing power needed for traditional AI. Lightweight AI is quickly being developed and deployed on low cost and easily available processors to get AI out in the field.

The final step is analytics. Often, there is not enough data to fully train AI, and the prospect of having to constantly update your AI system is daunting. Adding knowledge to your AI day after day to make it smarter in real-time is the most realistic path to deployment.

Fortunately, a new class of AI systems, called Lifelong-DNN™ is being introduced to eliminate this problem. Contrary to conventional AI, L-DNN systems learn like humans, one example at the time in 2-3 seconds, and continue learning after initial training. L-DNN closes the loop of Learn-Deploy-Analyze, to create a scalable solution for you.

How will AI most likely look in a SaaS product for inspection? It will be seamlessly integrated in the visualization tool that inspectors use to annotate images and will constantly learn on each action and annotation of the human analyst until the AI is able to do what the analyst does. Such acquired knowledge will improve continuously and will be available across all devices connected with such a centralized, “always on and always learning” AI.



What to expect (and what not to) from AI in inspections: a practical roadmap

While we would all want AI to be available on drones and cloud tomorrow, a more realistic look at the inspection ecosystem highlights that some application areas may come sooner than others.

While edge AI is appealing in its rationale and value-add, the first application domain of AI will be post-processing, one reason being that this AI won't require particular hardware to be available on the drone. It will be an "AI plugin" running on the enterprise software infrastructure.

Edge AI will be implemented after 1) the realization that post-processing AI is only as good as the data that is initially ingested and 2) beefier, lightweight processors have been developed for edge device, such as drones.

Once the two AI applications above are fielded, predictive AI will come into play, delivering a full, AI-powered software pipeline that maximizes data collection, insights and actionable intelligence for the enterprise.

Drone inspections are finally here, and, with AI for a brain, they are here to stay.



51 Sleeper Street, 3rd Floor, Boston, MA 02210
1-617-418-6161 • info@neurala.com • www.neurala.com