LANDING AI

A Deep Learning Playbook for System Integrators

How to quickly and efficiently evaluate, solve, and support AI/deep-learning machine-vision applications for manufacturing

Contents

Introduction

Manufacturing companies across industries of all types have long turned to automation technologies such as machine vision, robotics, and more recently, artificial intelligence (AI), deep learning (DL), and machine learning (ML) to **improve efficiency, increase throughput, and ultimately, drive revenue**. Increasingly, manufacturers are relying on expert system integrators to develop solutions for difficult and dangerous manufacturing jobs, improving the work conditions of human workers in every industry on the planet.

With a trusted system integrator on call, the small-, medium- and even large-enterprise customer doesn't need to constantly work to recruit, develop, and retain hard-to-find engineering talent. This mutually beneficial relationship has stood the test of time for decades, beginning with discrete manufacturing equipment; then modern material handling, robotics, and machine vision; and now AI and deep learning.

With a trusted system integrator on call, the small-, medium- and even large-enterprise customer doesn't need to constantly work to recruit, develop, and retain hard-to-find engineering talent.

$13 trillion

Estimated GDP growth between now and 2030
attributed to AI/DL technologies

LANDING AI

The Realities of AI and Deep Learning

Turn on the television, and you are likely to see ads for financial, medical, and even consumer services and products that mention AI. Most people outside of technical industries associate AI with artificial general intelligence (AGI), where a human can't tell the difference between AGI and a human worker in a blind test. Today's AI reality is very different from this sci-fi depiction. In current manufacturing environments, **AI — including deep learning and machine learning — leverages new programming methods to develop computer-based solutions to problems that previously could not be efficiently and cost-effectively automated**. Today's reality may seem like a glass-half-full comparison to the hype of AGI. However, informed manufacturers realize the massive potential that today's commercially available AI/DL solutions offer.

For example, between now and 2030, AI/DL technology is expected to create an estimated $13 trillion of GDP growth. While some of that value has already been realized by leading technology companies such as Google, Baidu, Microsoft, and Facebook, much of the additional waves of value creation will go beyond the software sector. And, manufacturing is expected to be one of the largest beneficiaries of AI/DL technology.

This *Deep Learning in Manufacturing Playbook* draws on insights gleaned from Andrew Ng, developer of the Google Brain team, the Baidu AI Group, and now Landing.ai. The playbook delivers insights into what manufacturing companies should consider prior to the deployment of AI and how a data-centric, machine-learning-development (MLOps) approach to AI/DL system design reduces risk, increases efficacy, and enables the type of continuous learning that is so important to manufacturers who must always **add new products and adapt to changing production environments**.

THE PLAYBOOK

1 GET INTELLIGENT ABOUT AI/DL

Successful SIs have educated themselves and leveraged vendor-sourced information to educate their customers on the true capabilities and best-use cases for AI/DL technology.

2 LEVERAGE DESIGN PLATFORM TO EVALUATE MACHINE-VISION APPLICATIONS QUICKLY AND EFFICIENTLY

Choose a platform that can evaluate, facilitate, automatically identify, simplify, and easily integrate

3 ADOPT COLLABORATIVE DESIGN TOOLS THAT BOOST EFFICIENCY

An MLOps data-centric AI/DL tool set for system design will accelerate application evaluation and solution design

4 WORK WITH THE CUSTOMER TO PREPARE THEM TO MAXIMIZE THEIR RETURN

Help curate information about company management and data structures that will help transform the enterprise into a modern manufacturing organization

5 LEVERAGE AN INTELLIGENT AI/DL PLATFORM FOR FOLLOW-UP ON DEVELOPMENT AND NEW PRODUCT INTRODUCTIONS AFTER POST-COMMISSIONING SUPPORT

Adopting an MLOps AI/DL platform with the right functionality will improve the technical and financial performance of automation technology integrators, machine builders, and line builders

1 Get intelligent about AI/DL

System integrators (SIs) have an abundance of design skill and business knowledge; it's what makes them so valuable to manufacturing companies. **What SIs lack, however, is time — specifically, engineering time.** Whether the SI is a sole proprietor or has a team of 40 machine-vision engineers, there is never enough time to answer every customer application call.

Many designers complain that the problem was made worse with the advent of commercial AI/DL solutions in the last decade. Throwing terms such as "AI" and "DL" into the manufacturing toolbox and unleashing the technology on a manufacturing sector that is uninformed about the technologies' true capabilities led to a stampede of customers asking for AI solutions. For SIs, manufacturers' fears of falling behind replaced quality specifications and scope-of-work documents. In response, SIs chased the unicorn. They spent hours on proofs of concept (PoCs) that are not suited to AI/DL solution sets.

The reasons for these early failures were numerous. In a few cases, the AI/DL software wasn't up to the task. But in many cases, the SI and customer did not have a clear understanding of how AI/DL technology differs from traditional machine-vision solutions. **In short, the biggest difference is that traditional machine-vision software programmers define defects directly and numerically while AI/DL defines defects using a statistical sampling of images already reviewed by expert inspectors who learn from their actions rather than mathematical definitions**. If it can be easily defined using arithmetic, solving the application with traditional vision software makes the best sense. Without this understanding, some SIs wasted time by trying to solve vision applications with AI/DL technology that are better solved using traditional machine-vision solutions.

Follow these four steps when conducting an initial evaluation of potential applications for machine vision, deep learning, or some hybrid approach:

1. **Choose a project that will achieve a quick win**. It should be possible for a new or external deep-learning expert, such as a systems integrator, to partner with internal teams that have deep domain knowledge to build AI solutions that start showing traction within weeks to months.

2. **Choose a project with a good ROI.** Make sure the customer has a clearly defined and measurable objective that creates business value. Don't pursue customer applications where AI/DL is deployed because it's a hot topic in machine vision. Deploy it because the customer needs an automated solution that traditional rules-based machine vision cannot offer.

3. **Evaluate the AI/DL model** against human inspectors to prove its value before deployment to a production line, especially if the line is a test for global deployments. This is generally referred to as Human-in-the-Loop (HIL).

4. **The project should be technically feasible.** Many companies seek to deploy deep learning in applications where it simply won't work. Having a trusted SI perform a feasibility study up front will increase the likelihood of success and avoid wasted time, effort, and costs.

Successful SIs have educated themselves and leveraged vendor-sourced information to educate their customers on the true capabilities and best-use cases for AI/DL technology. Fortunately, today there are many digital platforms to help SIs learn about AI/DL solutions. These include MOOCs (massive open online courses such as through Coursera or deeplearning.ai), trade publications, webinars, virtual conferences, e-books, and YouTube videos, all of which can help train large numbers of employees in AI and deep learning. A smart SI will curate rather than create content and then establish processes to ensure that employees are fully trained, and customers understand the framework of a successful AI/DL solution. Of course, having the right AI/DL design platform overcomes all of the early adoption challenges for AI/DL machine-vision inspection systems.

2 Leverage design platform to evaluate machine-vision applications quickly and efficiently

As previously stated, the one resource SIs always lack is time. Many designers learned the hard way that working with inadequate AI/DL software or trying to solve a machine-vision application with the wrong approach can siphon off more of their most precious resource: time. Therefore, SIs should choose AI/DL solutions designed specifically to:

1. **Evaluate** candidate applications in minutes, not days, for applicability with AI/DL solutions.
2. **Facilitate** data set collection, tagging, and client communications to refine AI/DL solutions.
3. **Automatically identify** data outliers that adversely affect AI/DL solution efficacy for faster builds.
4. **Simplify** management of data sets and AI/DL model creation regardless of the number of clients, deployments, or location.
5. **Easily integrate** with traditional machine-vision independent design environments (IDEs) from major machine-vision software vendors because most applications require a mix of traditional machine-vision features — such as "find edge" and "centroid" — plus AI/DL advanced defect classification capability for improved overall efficacy.

Today's SIs face dozens of different AI/DL software design packages, but which is the best fit?

Too often the final selection is made based on familiarity with the vendor's other offerings or simple proximity. A smarter approach looks at the daily challenges facing SIs and chooses the development platform that quickly and efficiently evaluates each application; puts designers back in control of AI/DL data sets and modeling; and simplifies design, commissioning, and after-launch support.

For example, during any given month, an SI may receive 10 different application queries, each specifically asking for an AI system solution. This is in addition to the half-dozen applications in development on the shop floor. Assuming the client submitted a good specification that identifies key regions of interest and includes images of sample defects as well as manufacturing process information, the SI's next step is to evaluate the application for AI/DL compatibility, quickly determining whether or not the application is viable before wasting resources.

The SI loads the customer's tagged data set into the AI/DL design package and initiates the training protocol; a computationally intensive effort that can take minutes, hours, or days depending on the size of the data set. The AI/DL design platform performs statistical analysis on the images and outputs a model with little-to-no programming from the designer beyond the accurate location and labeling of defects in the sample images. While this process leans in to the "artificial intelligence" concept, it also leads to a common criticism of AI/DL solutions, referring to them as "black boxes" because the software does the heavy lifting without significant programmer engagement.

Today, only a few AI/DL design platforms guide designers to improved model performance. This is a critical feature not just for optimizing models for today's production line but also for expanding the model to include tomorrow's product variations.

According to the theory of AI...

10% ▸ 1.88x

If just 10% of data is mislabeled, manufacturers need 1.88 times as much new data to achieve a certain level of accuracy.

30% ▸ 8.4x

If 30% of the data is mislabeled, manufacturers need 8.4 times as much data compared to a situation with clean data.

While a few of the AI/DL models will achieve maximum inspection accuracy after the first training session, usually the model needs improvement. In some cases, designers will try to improve the model by directly manipulating numerical values inside the AI/DL model. In most cases, this has the opposite effect because the computer is better at analyzing the images than humans. In other cases, designers will try to add more images to the data set in hopes that it will improve efficacy. However, according to the theory of AI, if just 10% of data is mislabeled, manufacturers need 1.88 times as much new data to achieve a certain level of accuracy. If 30% of the data is mislabeled, manufacturers need 8.4 times as much data compared to a situation with clean data. Based on this data, it quickly becomes apparent that a data-centric AI/DL design platform that is machine-learning-operations (MLOps) compliant can save considerable time and energy while placing the responsibility for system optimization back onto the customer who needs to provide quality data.

In some cases, poorly labeled data in sufficient quantity can improve an AI/DL solution's accuracy because with large data sets, it is easier to spot trends. But this isn't possible with small data, which requires a clean, non-noisy data set. And in most cases, customers come to SIs with a dozen or even a handful of images for each defect class. To increase data set sizes, most AI/DL design packages include image simulation tools that create additional defect images by copying the image and moving the defect around the image. However, while this approach can improve performance in general, this is not always the case, depending on how the data augmentation/simulation was done.

The best solution uses a design workflow that includes data labeling and organization tools that identify labeled data outliers, including unclear defect definitions, unclear labeling instructions, and labeler inconsistencies. In short, clean data ensures that small data is not the hurdle.

3 Adopt collaborative design tools that boost efficiency

As illustrated above, adopting an MLOps-compliant, data-centric AI/DL tool set for system design can accelerate application evaluation and solution design while reducing the friction caused by asking clients for hundreds or thousands of additional defect images to improve inspection accuracy.

For example, recently a steel manufacturer struggled to inspect rolled steel for defects as part of an automated product grading process. Manual inspection was only managing to capture about 70% of the defects, causing the manufacturer to sell high-grade steel below the potential market price. After spending months attempting to solve the problem, the manufacturer turned to an MLOps-compliant AI/DL design platform and engaged Andrew Ng's team at Landing.ai. **The team was able to reduce system design time by 67%, data labeling time by 50%, and increase accuracy by more than 90%.**

As a result, an AI/DL software program effectively increased design capacity by 67%, improving the SI's productivity and easing engineering labor constraints.

4 Work with the customer to prepare them to maximize their return

As illustrated above, the right AI/DL platform can accelerate machine-vision inspection system design and deployment, but to help guarantee follow-up orders from existing customers, smart SIs help their customers fully leverage the tools, growing enterprise data sets and AI/DL technology to maximum advantage.

For example, SIs can share curated information about company management and data structures that will help transform the enterprise into a modern manufacturing organization. In the AI era, a key moment for many companies is the formation of a centralized AI team. These teams typically report to the CTO, CIO, or CDO (Chief Data Officer or Chief Digital Officer). It could also be led by a dedicated CAIO (Chief AI Officer). Key responsibilities of an AI unit might include:

1. **Build up** internal deep-learning expertise to support the whole company.
2. **Execute** an initial series of high-return projects to gain visibility among key business groups with the assistance of trusted SIs.
3. **After completing the initial projects, set up** repeated processes to simplify, expedite, and support future projects.
4. **Develop** company-wide platforms that are useful to multiple divisions/business units and unlikely to be developed by an individual division. For example, consider working with the CTO/CIO/CDO to develop unified data-warehousing standards.

In many companies, multiple business units report to the CEO. With a new AI unit, customers can leverage AI talent into different divisions to drive cross-functional projects and create new business opportunities for supporting SIs.

5 Leverage an intelligent AI/DL platform for follow-up on development and new product introductions after post-commissioning support

Adopting an MLOps-compliant AI/DL platform with the right functionality can significantly improve the technical and financial performance of automation technology integrators, machine builders, and line builders. The same functionality that identifies data outliers during AI/DL training sessions also allows integrators to add new product variations to the model as the manufacturer's production mix evolves over time.

Best-in-class design platforms also include easy-to-use library management functions for connecting data set iterations to specific model outputs. Then, when dirty data adversely affects model performance, the designer can quickly revert to a better model state and send the new data sets to expert inspectors for additional consideration. This functionality can also help an SI or enterprise customer manage thousands or more AI/DL models across hundreds of production lines anywhere around the globe.

Building AI Momentum in Manufacturing

While AI/DL technology can be simple to develop and use when working on the right development platform, manufacturers will continue to lean on trusted SIs to help them keep up with the accelerating pace of change around the world and across the global manufacturing supply chain.

The following are some final thoughts for building momentum for AI/DL solutions among manufacturing clients:

When Landing.ai's Andrew Ng led the Google Brain team's adoption of AI technology, there was significant skepticism within the company — and around the world — surrounding deep-learning technology. To demonstrate value and gain momentum, Ng chose the Google Speech team as his first internal customer. He worked closely with the customer group to make speech recognition more accurate. By making the Google Speech team more successful using deep learning, Ng gave other teams more faith in AI/DL technology, which allowed the Google Brain team to gain momentum.

The second major internal customer was Google Maps, which used deep learning to improve map data quality. With these two successes, Ng began conversations with the advertising team. Building up momentum gradually led to more successful AI projects. This process is a repeatable model that SIs can use with any customer in any manufacturing industry from automotive and electronics to medical devices.

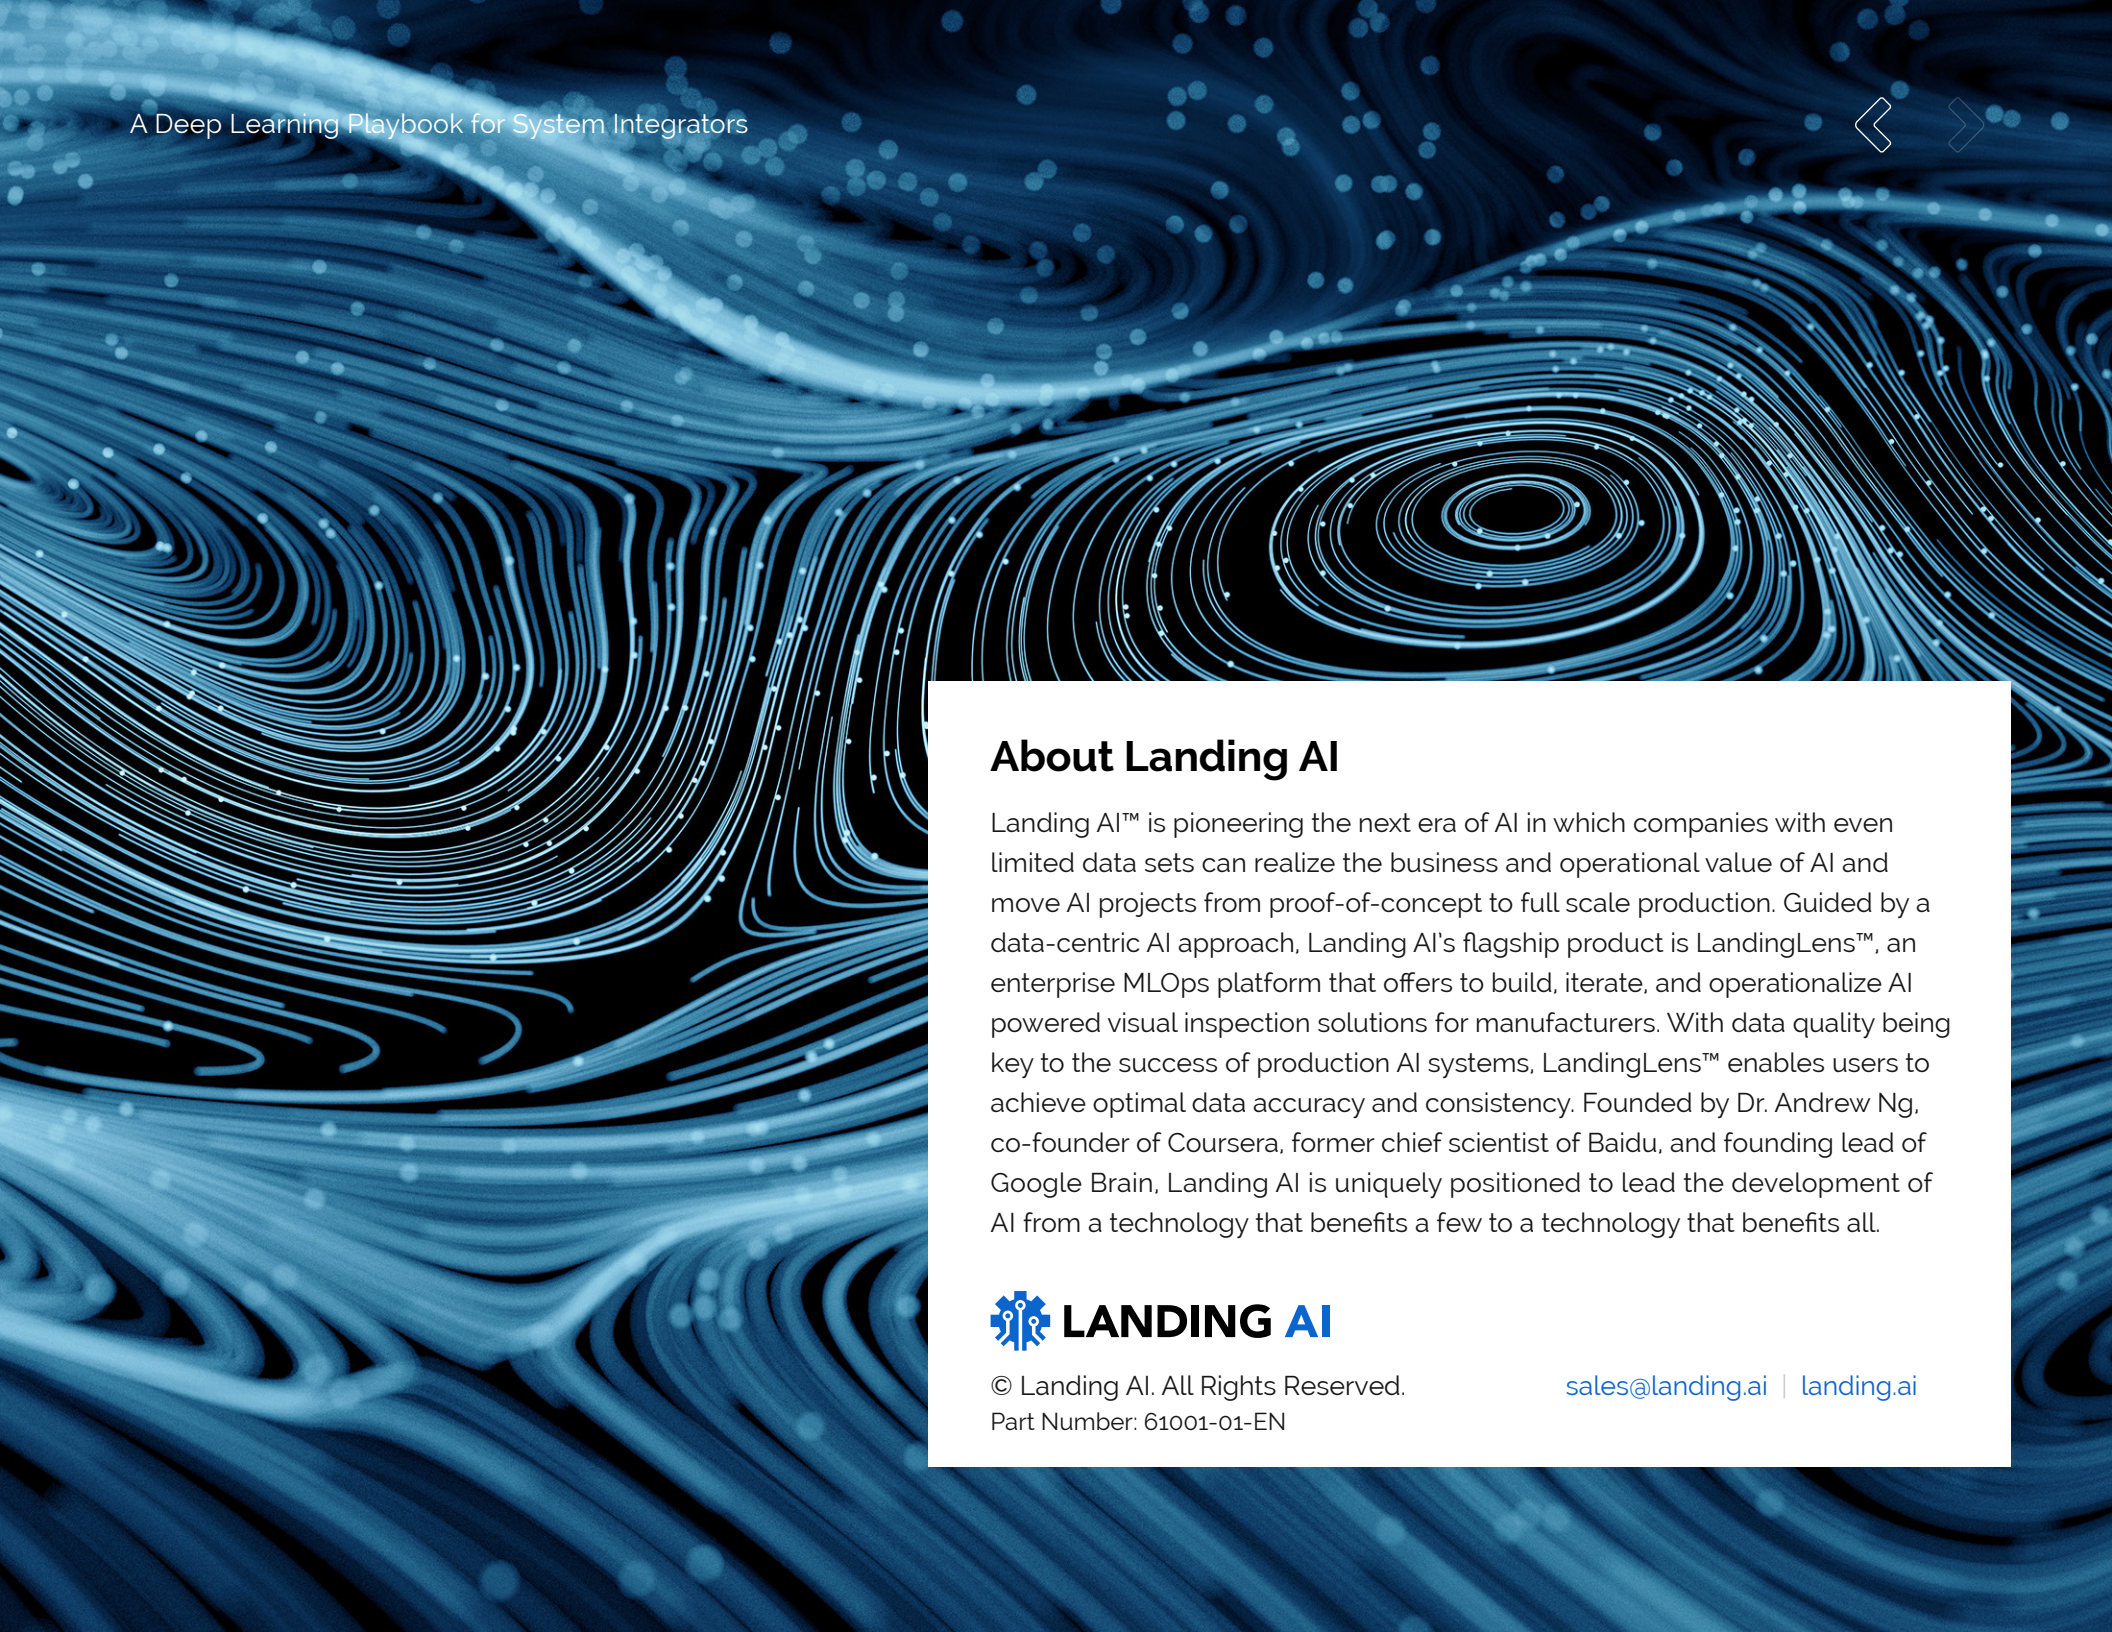

About Landing AI

Landing AI™ is pioneering the next era of AI in which companies with even limited data sets can realize the business and operational value of AI and move AI projects from proof-of-concept to full scale production. Guided by a data-centric AI approach, Landing AI's flagship product is LandingLens™, an enterprise MLOps platform that offers to build, iterate, and operationalize AI powered visual inspection solutions for manufacturers. With data quality being key to the success of production AI systems, LandingLens™ enables users to achieve optimal data accuracy and consistency. Founded by Dr. Andrew Ng, co-founder of Coursera, former chief scientist of Baidu, and founding lead of Google Brain, Landing AI is uniquely positioned to lead the development of AI from a technology that benefits a few to a technology that benefits all.

LANDING AI

sales@landing.ai | landing.ai

Part Number: 61001-01-EN