REPORT

Trends in Visual AI 2025

Data driven insights into how leading enterprises are deploying Al

roboflow

roboflow.com



THE NEW COMPETITIVE EDGE

Visual AI is transforming industries.

For leading enterprises, the next chapter in operational excellence is bringing artificial intelligence to the physical world.

An exciting trend is the growing adoption and availability of visual AI, which is bridging the gap between digital and physical processes in manufacturing, logistics, services, retail, and other industries.

This report examines how enterprises are deploying purpose-built visual AI and solving big challenges more quickly than ever.

Data drawn from thousands of projects

Roboflow supports the largest community of enterprises developing and deploying visual AI in the world. This report is based on anonymized data drawn from thousands of real-world AI projects deployed by organizations over the past six months.



Visual AI has become accessible to a broader range of businesses.

Organizations now have the option to develop powerful yet extremely efficient visual Al applications, capable of solving nuanced challenges in the cloud or at the edge.

• Training custom visual AI models is less data intensive than you expect.

With just a few hundred images, businesses can develop custom visual AI models capable of solving specialized challenges, greatly reducing development time and cost.

• Visual Al infrastructure has matured, reducing time to deploy.

With Al-accelerated labeling and dedicated toolsets, organizations are creating new visual Al models and getting them into production immediately.

• Powerful Al now runs on compact edge devices.

With vision models optimized for edge devices, organizations can deploy sophisticated, costeffective solutions in moving vehicles or remote environments.

NOTE FROM THE EDITOR Two concurrent revolutions in AI

While many associate cutting-edge AI with massive, data center-reliant models, a parallel revolution is happening: the rise of purpose-built AI.

These purpose-built vision models are unlocking solutions to previously intractable business challenges. They're enabling visual Al on edge devices, in moving vehicles, and in remote environments without internet connectivity. This report focuses on the wave of enterprise-deployed, purpose-built visual Al, as it represents a significant shift in practical Al applications.

Of course, the massive AI models with billions of parameters remain extremely vital. Later in the report, we'll explore how they are strategically combined with lighter weight models to solve complex challenges and bring new levels of automation and competitive advantage.

Enterprise-ready models can be trained with fewer images than you think.

Custom AI solves the toughest business challenges

When applying AI to specialized processes, like evaluating color accuracy of products, off-the-shelf AI models struggle. In these cases, custom-built vision models, trained on a collection of real world images from a real-world production environment, often perform better.

Developing custom AI is more accessible than ever

While it may sound daunting, the process of training custom AI has become streamlined. Modern tools and techniques allow for creating tailored AI systems with less effort and expertise.

While complex tasks may require larger datasets, many successful enterprise models are built with just a few hundred images. We've seen enterprises analyzing millions of images each week with AI systems trained on a little over 100 images.

Achieve high accuracy with smaller datasets

In our survey, 43% of enterprise models with top accuracy scores were trained on fewer than 1,000 images. This highlights the potential for rapid AI development and deployment with minimal data requirements.



Train accurate models with fewer images

Almost half of the models with high accuracy scores were trained using less than 1,000 images.

Less than 1,000 images in training set

1,001 to 10,000 images in training set

Over 10,001 images in training set

Based on 635 models developed for enterprise usage with accuracy scores over 85%

Models with millions of frames processed

Purpose-built vision models are trained on a range of dataset sizes, from a couple hundred to tens of thousands, depending on the task

SMALLEST MODEL

LARGEST MODEL

109 Images in training set

40,000 Images in training set

Visual AI models with varying accuracy scores are delivering value

Tailor AI accuracy to your specific needs.

While most AI models demonstrate high accuracy in controlled tests, the ideal level of accuracy depends on how you'll be using it in real-world situations.

Small improvements can make a big difference.

Al is tackling previously unsolvable problems. For example, manufacturers are using visual Al to catch product defects that were previously being missed. By identifying 80% of these defects, companies can reduce their quality issues by 80%, leading to significant improvements in product quality and substantial cost savings.

Combining different AI models enhances accuracy.

By linking multiple AI models, each with its own strengths, companies can achieve remarkably high accuracy, even if the individual test scores aren't perfect. This approach allows for the development of highly accurate and flexible solutions.



Different levels accuracy for different cases

While high accuracy is common, there are lower accuracy models seen in testing and production.



- 90 percentile
- 📕 80 percentile
- 70 percentile



Based on 406 AI models used in enterprise production during the sample period.

Models are retrained often and get better over time.

Increase model accuracy with minimal human effort

A major benefit of visual AI is its ability to continuously learn and improve. Once your initial AI model is up and running, you can retrain it with new data.



Top models with greatest accuracy improvement

The below models, developed by different organizations, saw the largest improvement in accuracy during the sample period. This chart shows the number of times the model was retrained, the number of images added to the dataset, and the resulting accuracy increase.



Based on 406 AI models used in enterprise production during the sample period.

Predicting equipment maintenance with AI trained on 400 sample images.



Facing challenges related to downtime and equipment maintenance, a global supplier of minerals and materials turned to visual AI for a solution. They developed a purpose-built vision model to monitor equipment condition and predict maintenance needs, aiming to improve operational efficiency and reduce costs.

The model was developed using a small, high-quality dataset of a few hundred sample images. To enhance the model's robustness, the team employed data augmentation techniques, duplicating and altering images to increase the training data variety.

The new model monitors and tracks the dimensions of raw materials along the production line, providing valuable insights into equipment condition and wear patterns. By detecting anomalies early on, the company can proactively perform maintenance, reducing unplanned downtime and surprise costs. BEHIND THE MODEL

8M frames processed each week

400-800 images in model training set

83% accuracy test score

6 model iterations

"Accuracy tests are useful when comparing model versions, but the realworld application is more important when evaluating the usefulness of an AI model.

"If you're tracking the wear and tear of industrial machinery, and you need to understand the variation in rock sizes over a period of days, you won't need a model with a huge number of images and 'perfect' synthetic test scores. The key is training on a high-quality dataset that accurately reflects the real-world conditions and challenges.

"This approach allowed us to develop a highly effective model with a small dataset and minimal effort, demonstrating the power of practical, purpose-built Al for solving real-world industrial challenges."

Jim Duffy, Head of Field Engineering, Roboflow



Deploying visual AI now takes hours or days instead of months.

Deploying AI is faster and more accessible than ever.

Developing custom AI no longer requires a team of machine learning experts and months of development time. Our survey found that the majority of new visual AI models were deployed within the same week they were trained.

Dedicated tools facilitate efficient AI development

Instead of building infrastructure and the software stack from scratch, businesses now have access to ready-made tools and platform-agnostic solutions that have reduced time necessary to deploy their custom AI systems and start seeing value immediately.



Train and deploy in the same week

Most organizations are able to quickly deploy a new vision model within the same week it was trained.



LESS TIME LABELING, BETTER QUALITY DATA

Al-accelerated annotation is reducing time spent developing models.

Developing accurate AI models no longer requires endless hours of manual data labeling. AI-powered tools and techniques have streamlined the process, enabling faster development and improved model quality. This allows organizations to focus on refining models instead of spending hours labeling images by hand.

Jumpstart your labeling process with four or five images.

Begin by labeling just a few representative images. The Al then uses these examples to suggest labels for the remaining images, dramatically accelerating the process.

Automate labeling with an initial "training" model.

After training an initial model on a small set of images, you can use this model to assist with labeling the rest of your dataset. The AI suggests labels, and you simply review and correct them, significantly reducing manual effort.

• Continuously improve your model with active learning.

Even after deployment, you can refine your model by reviewing its predictions on new images and adding those to the training data. This iterative process allows your Al to continuously learn and improve, ensuring long-term accuracy and effectiveness.

Al reduces labeling time 55%

When preparing larger datasets, users cut time spent labeling in half by using AI-assisted tools.



"It's now possible to develop sophisticated vision models, test the results in a real environment, and put them to work in a matter of hours or days. This unlocks new possibilities for industrial applications of visual AI."

Brad Dwyer, Co-founder & CTO, Roboflow



Ó

CASE STUDY: AI & FOOD PROCESSING

Inspecting product packaging with powerful AI on small edge devices.



2.7M frames processed each week

BEHIND THE MODEL

2,700 images in model training set

93% accuracy test score

26 model iterations

A national food supplier sought a solution to detect packaging defects early in the production process. Their goal was to improve quality control and reduce waste.

The high-speed production environment presented a challenge, requiring an AI solution with low latency to effectively detect defects in real-time. Additionally, occasional internet outages meant the system needed to operate reliably on-site without relying on cloud connectivity.

To address these challenges, they developed a purpose-built visual Al system that runs on a compact edge device. The system accurately detects packaging defects in real-time and sends immediate alerts to staff, enabling them to intervene promptly, avoid costly mistakes, and prevent issues from reaching consumers.



AI AT THE EDGE

The system was deployed on an NVIDIA Jetson Nano, a powerful yet compact edge computing device. A highspeed network camera captures images of the packaging, and the AI model processes them in real-time.

Combining multiple Al models to achieve complex tasks increases reliability

Greater efficiency and accuracy with multi-stage Al workflows

By combining specialized AI models, businesses can create powerful, multifaceted solutions for complex tasks. This modular approach simplifies development, enables faster adaptation to changing needs, and often leads to significantly improved performance and accuracy.

Consider a system designed to scan freight identification information. A single OCR model might capture all text in the image, leading to extraneous or incorrect data. However, a multi-model approach allows for greater precision. One model can identify the region containing the relevant information, and a second model can then accurately extract the specific text.

By strategically combining the strengths of different AI models, businesses can efficiently address niche challenges without the need to build every component from the ground up. This approach saves time, resources, and accelerates the deployment of effective AI solutions.

BEFORE: SINGLE STAGE

Applying one AI model to an image can result in extraneous predictions, such as this example of scanning all the text on a vehicle.



AFTER: MULTIPLE STAGES

By using a two-stage process, with one model finding the vehicle ID and another scanning text, results are more accurate



Bring the sense of sight to your entire business.

Enterprises are integrating visual AI with existing infrastructure and software, which is unlocking new ways to automate processes, get real-time alerts, and gain valuable insight.



STEP 1 Live video & imagery Collect input from network cameras, drones, or vehicle-mounted feeds.



STEP 2 Visual Al analysis

Connect multiple AI models to video feeds for real-time comprehensive.



STEP 3

Integration

Send output to a range of machinery, databases, or trigger alerts.

Accurate data = Better visibility = Competitive advantage.



Many of the largest freight and logistics companies in the world are automating yard management with visual AI.

With vehicle-mounted cameras – powered by visual AI – these organizations can identify containers and update the yard inventory system in real time. By connecting the data with GPS coordinates of the vehicle, they are able to record the precise location of containers in the yard.

Having complete and accurate view of yard inventory enables these organizations to streamlined operations and use less resources. Additionally, having granular data about yard utilization over time, they can better forecast utilizations and make more informed decisions. These AI-powered insights provide critical business advantages. BEHIND THE MODEL

280K frames processed each week

20K images in model training set

99% accuracy test score

14 model iterations

"Projects like these aren't just about automating inventory, saving time for crew members, and saving costs here and there.

"A key value for organizations comes from gaining a deeper understanding into their our operations. By connecting visual AI with precise location data, you can get incredibly accurate data and visibility into the status of business assets, like yards, warehouses, containers, vehicles, and more.

"This granular data allows you to optimize space utilization, observe trends, and make data-driven decisions about future investments. Ultimately, this translates to improved efficiency and competitive advantage in the long run."

Jeremy Powers, Head of Sales, Roboflow



Visual AI is transforming every industry

Throughout this report, we've explored the transformative potential of visual Al across various industries – how the barriers to developing powerful, enterprise-specific Al are falling.

• Visual AI is more accessible

With streamlined tools and platforms, businesses can quickly develop and deploy custom AI models without extensive expertise or resources.

• Edge Al unlocks new possibilities

Deploying AI on compact edge devices enables real-time analysis and automation in remote or challenging environments.

• Al drives enterprise-wide value

Connecting visual AI with existing systems unlocks new levels of efficiency, automation, and data-driven decision-making.

Visual AI is bridging the gap between digital and physical worlds. Let's see how it can transform your organization.

Contact us

roboflow