EUREKA
ROBOTICS

VOL. 01 · 2026

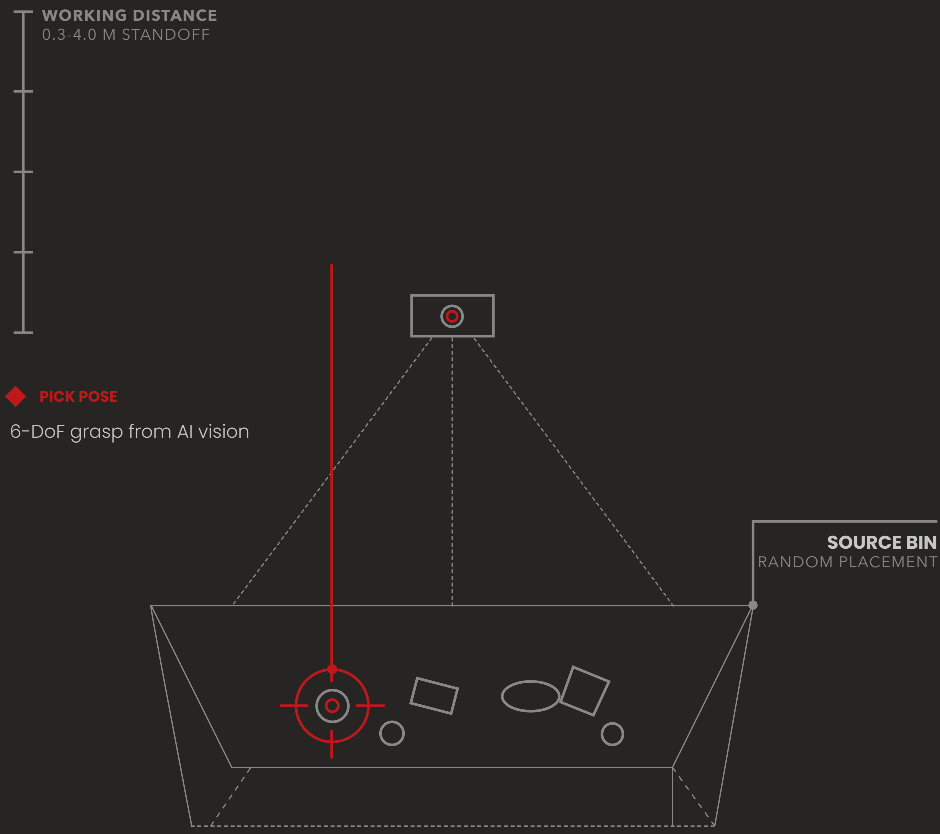

A FIELD GUIDE FOR ENGINEERS

07 CH • 20PP

Bin Picking

From concept to production

EXPERIENCE	COMPLETED PICKS	CUSTOMERS
8YRS	30M+	Fortune 500

What’s inside

This field guide is for engineers keeping up with physical AI, and vision-guided picking in particular. It covers the core concepts, the failure modes that show up in production, and what to look for when buying a system. Where you want more depth, the QR links point to a fuller resource.

BUILT ON PRODUCTION DATA

The content draws on more than **30 million production picks** at **99.5% reliability** across customers including:

Eureka Robotics · eurekarobotics.com

Table of contents

01

Is bin picking the right solution?

Bin picking is the right tool when parts arrive in random positions or orientations, or when known reference positions vary enough to require visual adjustment. If parts are in a fixed pattern, a programmed sequence without vision is faster, cheaper, and more reliable.

READ MORE
Full bin-picking field guide
eurekarobotics.com/blog/bin-picking-guide

Decision matrix

Bin picking isn't always the right answer. Several common scenarios are served better, and more cheaply, by a simpler approach.

SCENARIO	BETTER SOLUTION
Parts random in a bin or container	**Bin picking**
Parts in fixed, repeatable pattern	Programmed sequence, no vision
Loading time small fraction of process time	Manual loading into fixture
Upstream process can deliver parts on a tray	Indexed input tray
Few SKUs, small simple non-fragile parts	Bowl feeder
High-precision, complex parts under 10mm	Flexible feeder + camera

If bin picking is the right approach, the next question is whether a specific application is feasible. Part and bin characteristics determine more than any software choice.

FACTOR	HIGH CHANCE OF SUCCESS	CHALLENGING BUT POSSIBLE	NOT RECOMMENDED
Surface	Matte, flat, smooth	Shiny, translucent, uneven	Fully transparent
Placement	Bulk, no indexing	Accurate w/ indexing	N/A
Geometry	Simple shape	Complex shape	N/A
Cycle-time target	5–8 s	< 5 s	< 4 s
Required success	~98%	99.9%	100%
SKU count	Single	Multiple	> 20 diverse

EARLY WARNING SIGNS A PROJECT WILL STRUGGLE

- Cycle target under 4 s on complex parts needing precise placement
- "100% success rate" written into the requirements
- Hundreds of diverse SKUs with no shared structure
- Entangling components: springs, hooks, wires

02

Anatomy of a bin picking system

A visual reference for what goes into a Eureka AI Vision cell. The exact requirements depend on the application, but the classic bin-picking cell is surprisingly simple: an arm, eyes, a brain, and some lighting.

A complete cell

The camera sits above the source bin and feeds the controller, which runs the AI models and drives the arm. The arm picks from the random pile with a suction or finger gripper and places into the destination. Lighting and safety fencing complete the cell.

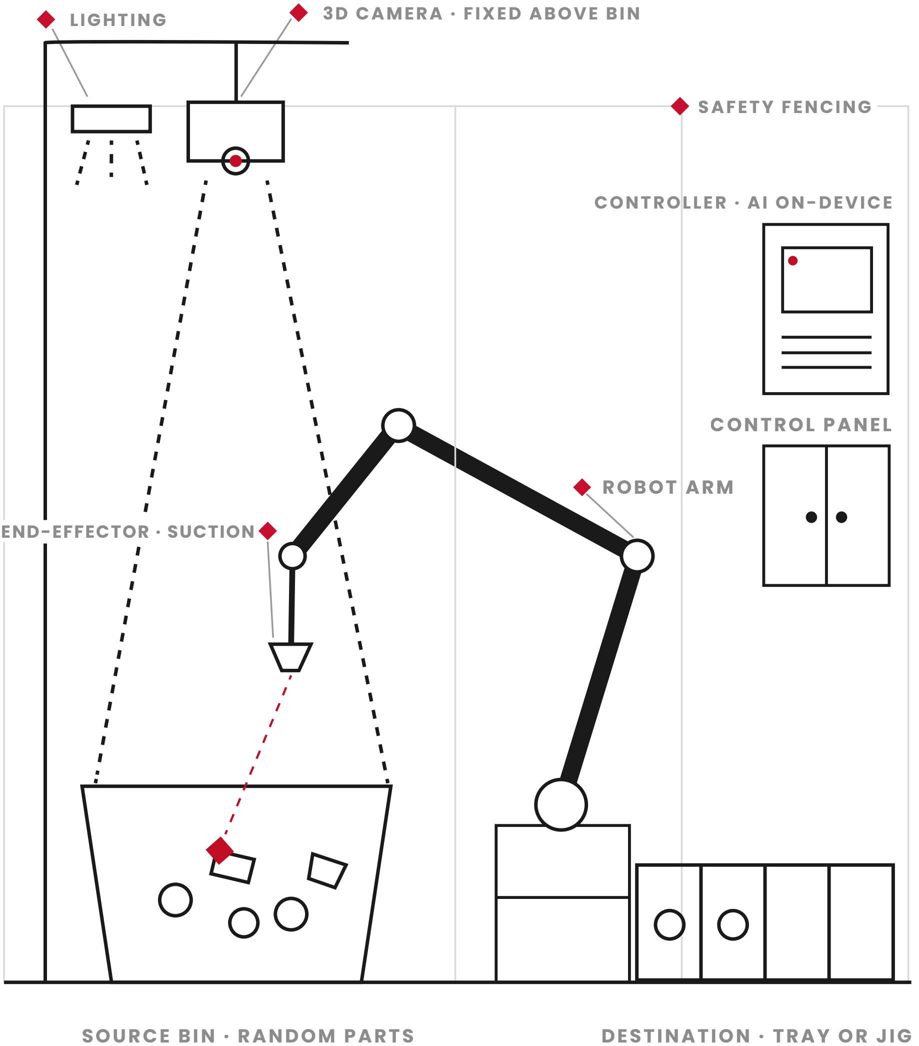

Not to scale. Safety may be physical fencing or light curtains, depending on the risk assessment.

03

How the vision system works

In theory, a bin-picking cell is simple: a vision system finds the parts and tells the robot what to do.

This chapter compares the 3D sensing technologies that do the seeing, and the role AI now plays in turning a camera image into a pick.

TERMS YOU SHOULD KNOW

POINT CLOUD The 3D camera's output: thousands of measured surface points in space.

6-DOF POSE A part's full position and orientation (x, y, z + roll, pitch, yaw).

HAND-EYE CAL. Camera-to-robot calibration. Set in the field; the one that drifts. Intrinsic calibration is factory-set.

EYE-IN-HAND Camera on the robot wrist, vs. fixed above the bin.

SIM-TO-REAL Generating rendered, annotated images from a CAD model to train a model, then deploying it to the vision system to pick real parts.

READ MORE
How ML changes the vision problem
eurekarobotics.com/blog/machine-learning-bin-picking

3D sensing technologies, compared

Time-of-flight and structured light used to be the dominant approaches for 3D picking. With AI in the picture, it's become possible to solve the so-called stereo matching problem: extracting 3D depth from two 2D images. That makes stereo cameras, the cheaper and more robust solution, viable.

	TIME-OF-FLIGHT	STRUCTURED LIGHT	STEREO VISION
How it works	Return time of light pulses	Patterns + triangulation	Two cameras + pixel matching
Resolution	Lower	Good	Good (ML stereomatch)
Hardware complexity	High	High	Low
Robustness	Moderate	Moderate	High
Cost	Higher	Mod-high	Lower

The role of AI in bin picking

AI models in the vision system handle two jobs: recognizing parts and deciding how to grip them. Replacing rule-based vision with learned models is what takes deployment from months to hours, since there is little to program by hand. At Eureka, two models work in sequence. A **detection** model takes the image from the 3D camera and finds the parts in the bin. Its output passes to a **picking** model that works out how to pick the next part and sends the instruction to the robot controller.

PIPELINE

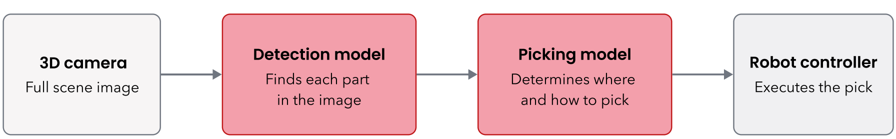

04

How AI models get trained

Most providers have historically trained vision models from CAD. That works until it doesn't: a CAD model may not exist, real parts may not quite match it, and manufacturers are often hesitant to share CAD files early in a project.

That's why Eureka developed three training approaches: **masterless, CAD-based, and annotation-based.** One of them will fit whatever your project has on hand.

READ MORE
How bin-picking AI models are trained
eurekarobotics.com/blog/how-bin-picking-ai-models-are-trained

Three training approaches

	MASTERLESS	CAD-BASED	ANNOTATION
What it is	Pre-trained foundational model	Synthetic data from 3D model	Real images, human-annotated
Effort	None	Low	Moderate
Time to first pick	Hours	1-2 days	1-2 days
CAD required	No	Yes	No
Part complexity	Simple, distinct	Mod-complex	Mod-complex
Placement	Bulk only	Precise	Precise

Masterless uses a model already trained on millions of objects, so there is no training step at all. It's the right fit for POCs and for applications where you can't predict the SKUs, like warehouse fulfillment. Bulk placement only.

CAD-based is the preferred route when accurate CAD exists. A simulation tool takes the model, renders hundreds of scenes with varied lighting, clutter, and noise, and trains on them automatically. Because the tool knows where every virtual part sits, the annotations are exact, which is what makes this both fast and accurate.

Annotation-based works the same way without CAD: you capture 30–50 real images of parts in the bin and annotate them, with smart segmentation tools doing most of the tracing. Use it when CAD is missing, parts don't match it, or surfaces (oil, texture, reflectivity) don't simulate cleanly.

Which to use: masterless for POCs and unpredictable SKUs; CAD-based when CAD exists and placement must be precise; annotation-based when there's no CAD or reality diverges from the model.

CONTINUOUS IMPROVEMENT

When the system fails on a specific case, you add that case to the training set and retrain. It takes **minutes**, and it runs on the controller, on-premises, so CAD and images never leave the facility. The new model handles that scenario and similar ones it hasn't seen.

05

Good to know

Field notes from production deployments: cell design, integration, performance, and reliability.

Cell design

A fixed camera is the default. It scans between picks, so vision costs no cycle time. An in-hand camera adds 1–2 s per cycle because the robot has to stop to scan; use it only when one camera must cover multiple bins or get close to small parts.

Suction first. If the part has a flat, smooth, clean surface, suction is simpler and 1–2 s faster than fingers. Use a two-finger parallel gripper when suction won't hold; anything more complex adds failure points.

Size and weight are relative. A 2 mm part in a 10 cm tray is manageable; in a 1 m bin, very difficult. A heavy carton picks easily by suction; a heavy, small, irregular part may have no gripper option at all.

Cobots cost 3-4 s per pick compared to an industrial robot. Fencing is usually the cheaper trade.

Plan for bin edges. Walls and corners limit approach angles, so collision checking has to include the bin and the gripper body, not just the part. The last parts in a bin always take longest.

Integration

Two bottlenecks dominate schedules. Outdated or proprietary robot and PLC software (closed APIs add days) and weak networks. Use wired CAT6 or better, and never put Wi-Fi between cell components.

Confirm protocols before committing. Modern controllers, including the Eureka Controller, work best with open architectures and standard protocols: EtherNet/IP, Modbus/TCP, Modbus/RTU. Eureka integrates with every major robot brand.

Most operations need a system integrator. A good one brings mechanical, electrical, and software capability, plus local presence for first-line support. Only go without one if you have strong internal teams in all three.

Budget engineering time for the edge cases. Most of it goes to testing, troubleshooting, and on-site fine-tuning, not initial integration or software setup. This is the line item schedules underestimate most.

READ MORE

Eureka setup: From unboxing to production

eurekarobotics.com/blog/eureka-setup-from-unboxing-to-production-in-an-afternoon

Performance

Cycle time runs 4 to 8 s in well-designed production cells: about 5 s for simple parts picked by suction and placed in bulk, 7 to 8 s for complex parts that need precise placement.

Pick success rate sits at 98 to 99.9% in a tuned cell. Treat a 100% requirement in a spec as a red flag, not a target.

Accuracy ≈ 0.5% of container size. Around 0.5 mm on a 10 cm tray, around 4 mm on an 80 cm bin. Set placement requirements from that number, not from a spec sheet.

Demos understate production. A demo rarely pushes cycle time hard, and success rates keep improving as tuning accumulates in the weeks before go-live.

Do the exception math. 99.5% still means 5 misses per 1,000 picks. Decide before go-live what happens on a miss: rescan and retry from a new angle, agitate the bin, or call an operator. Cells that plan for exceptions run unattended; cells that don't generate support calls.

Three reliability metrics matter: pick success rate, downtime to debug and fix, and downtime to add new SKUs. The third one compounds over the system's life. If every new SKU needs the vendor, that cost quietly eats the ROI of a multi-SKU cell; if your team can retrain independently, it doesn't.

Reliability in the field

Failures look like software; causes are usually physical. Wear, vibration, calibration drift, lighting. Check those before you start debugging code.

Lighting is the most common field problem. Overexposure from sunlight or new fixtures, glare off reflective parts, underexposure in deep bins. It shows up in factories far more often than in development.

Watch the wear items. Suction cups, finger pads, pneumatic lines, and cables loosening at the wrist from millions of repetitive motions.

Calibration drift is the hidden one. It produces gradual placement errors or intermittent misses that are hard to reproduce. If miss rates climb over weeks or months, recalibrate before assuming a software bug.

A REAL EXAMPLE

A cleaning crew wiped a camera housing and shifted it by less than 0.1°. Mounted over 1 m above the bin, that produced about 1 mm of localization error, which caused a placement failure once every few hundred cycles on laser lenses going into tight-clearance jigs. Nobody considered that the camera had been touched. **Takeaway:** write a clear SOP for what staff can and cannot do near the camera.

Prevention that pays for itself: a cell maintenance SOP with explicit camera rules, quarterly preventive maintenance, and multi-day stress testing before go-live. None of this is novel. It's just what gets skipped first when schedules tighten.

06

Choosing a vendor

As the market matures and more providers enter, solutions get harder to compare. Most providers fit into one or two of these categories:

- **Pure 3D camera makers:** provide hardware only.
- **Vision software vendors:** usually paired with third-party cameras.
- **Integrated solution providers:** camera, vision software, and robot control as one coordinated stack. This is where Eureka Robotics sits.
- **System integrators:** purchase the hardware and software components and deliver the cell on your floor, typically working closely with the vision vendor.

WHAT ACTUALLY DIFFERENTIATES

Time-to-stable-pick, not feature lists

On paper, most providers look similar. An integrated camera-controller-software system like Eureka's AI Vision System cuts integration time significantly, since the components are designed to work together, and gives you one point of contact for support instead of three.

READ MORE

The Eureka AI Vision System

eurekarobotics.com/technology

07

Before you call a vendor

Gather this before the first conversation. Any serious feasibility assessment needs it, and arriving prepared shortens every later step.

WORKPIECE

- [ ] Images of each workpiece
- [ ] CAD data, if available
- [ ] Total SKU count
- [ ] Frequency of new SKUs

SOURCE BIN

- [ ] Images, empty and loaded
- [ ] Randomness: fully 3D, 2.5D, or partially structured
- [ ] Typical face-up orientation

DESTINATION

- [ ] Images of target bin or jig
- [ ] Placement accuracy required (mm)
- [ ] Fixed or variable position

OPERATIONAL

- [ ] Required cycle time
- [ ] ROI target / payback expectation
- [ ] Gripper constraints (designated or excluded grip positions)
- [ ] Safety & HMI requirements

Questions for the vendor meeting

Ask these to get a full picture of what a vendor actually offers.

TECHNOLOGY

- [ ] Where exactly is ML applied: vision, motion, or both?
- [ ] What 3D sensing technology, and why?

TRAINING

- [ ] Which approaches do you support: masterless, CAD-based, annotation-based?
- [ ] Is CAD required to add a new SKU?
- [ ] Can our team train new SKUs without you?
- [ ] Does training run locally or in the cloud?

INTEGRATION

- [ ] Which robot brands and protocols are supported out of the box?
- [ ] Realistic engineering effort for a typical cell?

COMMERCIAL

- [ ] CAPEX or subscription? Per-pick fees?
- [ ] What does the maintenance contract cover and cost?
- [ ] Cost to add a new SKU after go-live?

SUPPORT

- [ ] Will you demo with our actual parts before we commit?
- [ ] Post-go-live support: who provides first-line response?
- [ ] Reference customers in a similar application?

THE STRONGEST SUCCESS FACTOR

More than any technical factor, success correlates with a customer-side point of contact who understands both the production process and what AI and robotics can realistically do. Struggling projects hand off; successful projects share ownership.

Where to start. The most productive first step is a lab demo with your actual parts: dirty, oily, end-of-batch parts, not cleaned samples. It surfaces the real engineering questions while they're still cheap to address.

NOTES

SEND US YOUR PARTS

We'll show you what's feasible

Eureka offers complimentary feasibility checks. Within 72 hours of your parts arriving at our Atlanta office, we'll send back a written verdict and, where it's feasible, a video of our robot picking them.

WEB

eurekarobotics.com

LAB TRIAL

contact@eurekarobotics.com

EUREKA ROBOTICS