published by
International Federation of Robotics
Frankfurt, Germany
February 2022

IFR
International
Federation of
Robotics

POSITION PAPER

Artificial Intelligence in Robotics

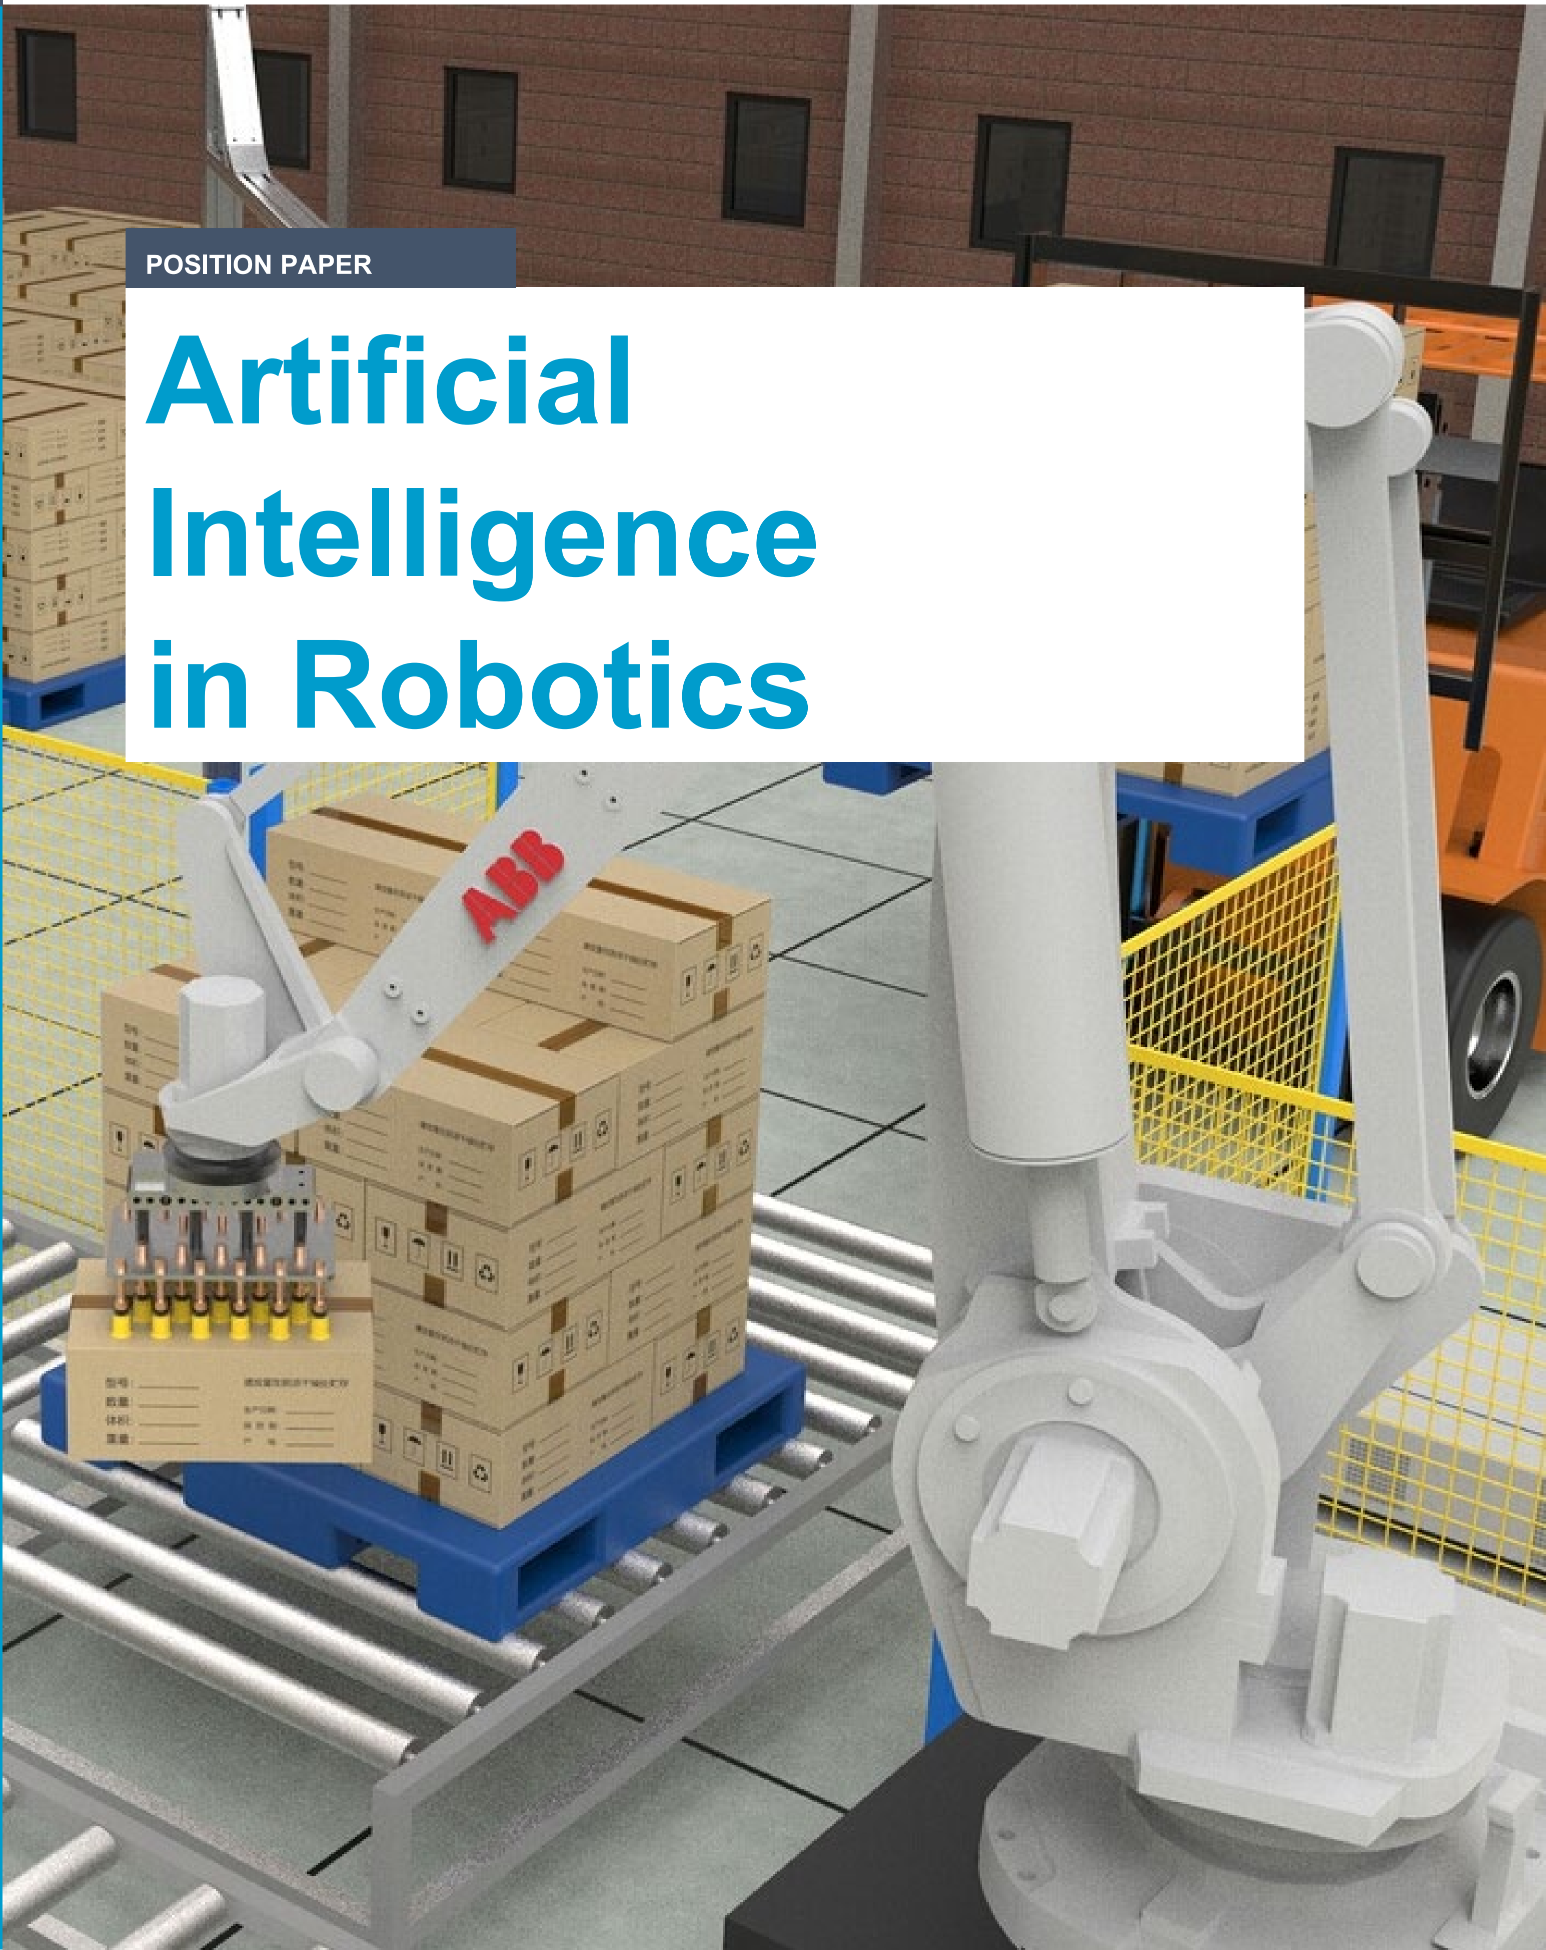

IFR International Federation of Robotics

Contents

Executive Summary 3
Introduction 4
Definitions 4
- Artificial Intelligence 4
- Robots 6
- Artificial intelligence in robotics 6

Applications 8
- Sense and respond 8
- Performance Optimization 14
- The impact of AI in robotics on work and jobs 17

Future trends 17
- Market adoption 17
- Research directions 18

Safety standards and certification for AI in robotics 19
- Safety 19
- Certification 20

Conclusion 20
Case studies and videos 21

Cover picture: MechMind

Executive Summary

Artificial Intelligence (AI) holds great potential for robotics, enabling a range of benefits in sectors as diverse as manufacturing and healthcare. Though AI is already making its mark on robotics, it is at a much slower pace and in a far narrower field of application than is commonly assumed. This paper summarizes the most common applications of AI in robotics currently in commercial use and provides an overview of market potential over the next 5 to 10 years.

The main aim of using AI in robotics is to better **manage variability and unpredictability in the external environment**, either in real-time, or off-line. This offers benefits for manufacturers, logistics providers and retailers dealing with frequently changing products, orders and stock in so-called 'high mix/low-volume' environments. It also helps robots to function in public environments – from supermarkets to hospitals - which are inherently unpredictable.

AI is not necessarily a prerequisite for dealing with variability and unpredictability. For example, simple pick and place applications with variance in product placement, but not in the product itself, can be achieved without AI. Also, robot mobility does not require AI. However, the greater the variability and unpredictability of the environment, the more likely it is that AI algorithms will provide a cost-effective and fast solution – for example for manufacturers or wholesalers dealing with millions of different products that change on a regular basis. AI is also useful in environments in which mobile robots need to distinguish between the objects or people they encounter and respond differently.

We differentiate between two main categories of application:

- **Sense-and-respond** applications in which the robot identifies and responds autonomously to its external environment in a real-time closed loop
- **Performance optimization** in which AI is used to **optimize process design** and robot programming as well as to **improve quality inspection and maintenance.**

Pick-and-place applications are the most widely adopted type of sense-and-respond application in manufacturing and logistics. Pick-and-place covers a broad range of applications including palletizing, packaging and machine tending. Other manufacturing applications in very early stages of AI adoption are **assembly and welding**.

AI has considerable potential to **speed up design and programming** of robotic automation, though this is in early stages of development. AI could help significantly lower the overall cost of the installation and re-programming of a robotic application – by 50% according to some experts. This helps make automation economically viable for small-to-medium sized manufacturers and larger companies with high variability. The use of AI in design and programming of robot applications is still at a very early stage however. Currently, the most established performance optimization application is **robotic quality inspection.** AI enables faults that may be undetectable for humans to be identified at each stage of the production cycle. This improves product quality and also minimizes waste as faulty parts can be taken out of the line before being worked on further.

The manufacturing and logistics sectors are leading the use of AI in robotic applications. However, there are rapid developments in the use of AI in healthcare and professional service assistance robots.

Safety in robotic applications using AI is currently ensured through hard-coded, deterministic algorithms that always take priority, ensuring the robot stops when encountering an obstacle. This may change as AI becomes more ubiquitous in robotics, and researchers are currently looking at ways to enable the use of 'trustworthy' AI in safety-critical components of robotic applications.

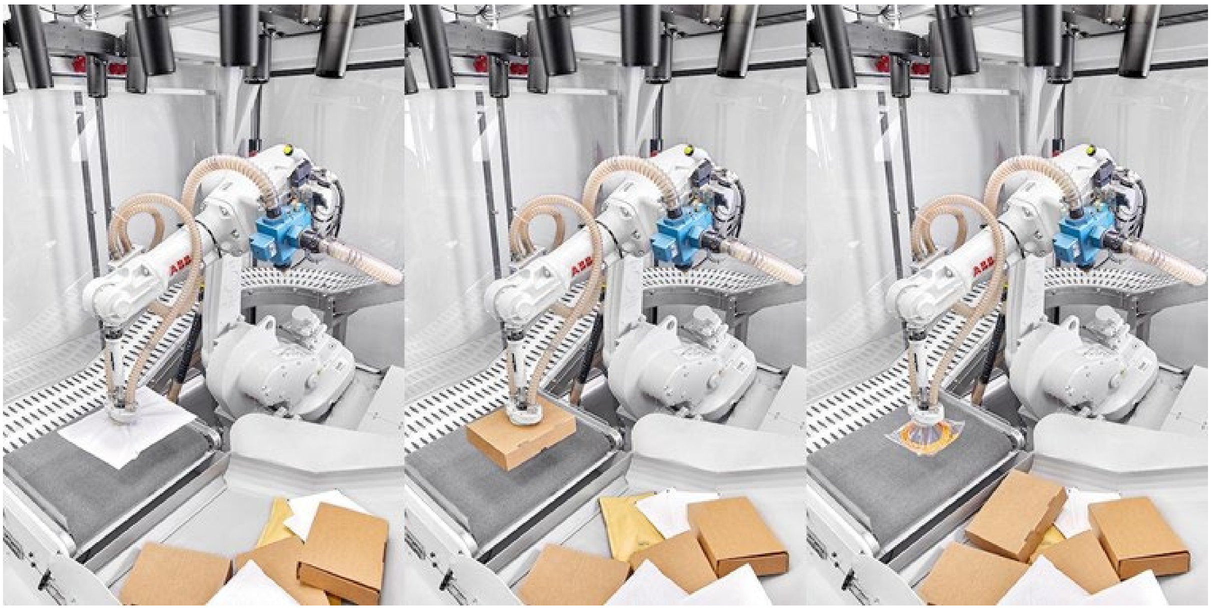

Figure 1: Pick and place using AI, image credit: ABB.

Introduction

There is much confusion around the role of artificial intelligence (AI) in robotics. Images from Hollywood movies from 'Transformer' to 'Ex Machina' portray a future in which robots have achieved human-like intelligence and physical capabilities. Terms such as 'robotic process automation' are used to describe software programs that have no physical robotic component. And there is no standard definition of which algorithms constitute AI.

AI is still in its infancy in real-world applications of robotics. Traditionally, automotive and electronics manufacturers - the main adopters of industrial robots – have looked for precision, speed and predictability from their robot. Achieving these characteristics does not require AI. Increasingly however, manufacturers operate in an environment of high variability, producing small quantities of a large range of products, often ordered at short notice. Re-programming robots each time to fit new production or packaging lines is costly. Meanwhile, robots have moved out of factories and into public domains such as hospitals and supermarkets. And logistics robots, which fetch-and-carry parts and packages around factories and warehouses, are booming. These environments are characterized by high levels of variability and unpredictability, where AI can offer significant benefits, enabling robots to adapt to a changing external environment in real-time without recoding. AI is also starting to be used in optimizing process design and robot programming and is increasingly adopted in robotic quality inspection.

This paper looks at the most common uses of AI in commercial robotic applications and discusses market trends. It also discusses safety standards and regulation for AI in robotics.

Definitions

Artificial intelligence

Existing definitions of AI vary in scope and there is no standard. For example, researchers Stuart Russell and Peter Norvig[1] define AI as 'an intelligent agent, where "Agent" means a software system which perceives its environment through sensors and acts upon that environment through actuators and "Intelligence" means the ability to select an action that is expected to maximize a

[1] *Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. Pearson Education, 3rd edition, 2009*

performance measure. The European Commission, meanwhile, defines AI as 'Systems that display intelligent behavior by analyzing their environment and taking actions – with some degree of autonomy – to achieve specific goals' [2].

There is also no agreement on which algorithms or methods constitute AI. Part of the difficulty with setting a concrete definition is that AI algorithms and models (such as neural networks) are a combination of mathematical functions that, alone, could be considered 'traditional' algorithms. At a certain point, when these functions are combined – such as in a neural network - they achieve a level of autonomy that can be characterized as AI, or they could simply be considered as a combination of traditional mathematical functions.

There is general agreement that machine learning counts as AI. In contrast to a deterministic algorithm, a machine-learning algorithm learns and improves itself through experience. There is less agreement on which algorithms for other areas such as reasoning and decision-making, path planning and performance optimization count as AI. Definitions have changed over time. Algorithms that were considered AI twenty years ago are now not, leading to the joke among AI experts that 'AI is just algorithms we don't understand yet.'

Within the category of machine learning, we refer in this paper to four types described below:

- **Supervised:** The algorithm learns from being presented with hundreds or thousands of labelled examples of what it is tasked with identifying.
- **Semi-supervised**: The algorithm trains itself from a combination of labelled examples and unlabeled examples.
- **Self-supervised**: The algorithm uses labels that already exist in the input data (for example, specific sounds or colors) to create a training model.
- **Reinforcement learning**: The algorithm is given a goal and is rewarded for making steps towards that goal and so learns how to complete a task effectively through trial-and-error learning.

Table 1: Robot definitions according to ISO 8373-2021 [3].

Term	Definition
Industrial Robot	▪ A automatically controlled, reprogrammable multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or fixed to a mobile platform for use in automation applications in an industrial environment ▪ Industrial robots include the manipulating portion(s) of mobile robots, where a mobile robot consists of a mobile platform with an integrated manipulator or robot.
Service Robot	▪ Robot in personal use or professional use that performs useful tasks for humans or equipment. (Tasks in personal use include handling or serving of items, transportation, physical support, providing guidance or information, grooming, cooking and food handling, and cleaning. Tasks in professional use include inspection, surveillance, handling of items, person transportation, providing guidance or information, cooking and food handling, and cleaning.)

[2] *https://digital-strategy.ec.europa.eu/en/library/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines*

[3] *ISO 8373:2021 - Robotics — Vocabulary*

Mobile Robot	▪ Robot able to travel under its own control ▪ A mobile robot can be a mobile platform with or without manipulators. In addition to autonomous operation, a mobile robot can have means to be remotely controlled.

Robots

Robot definitions are provided by the International Standardization Organization (ISO) in ISO 8373-20213 (see Table 1). IFR mainly distinguished between industrial and service robots. Mobile robots are considered as subset of service robots, including those used in manufacturing settings.

Artificial intelligence in robotics

The main aim of using AI in robotics is to better manage variability and unpredictability in the external environment, either in real-time, or off-line.

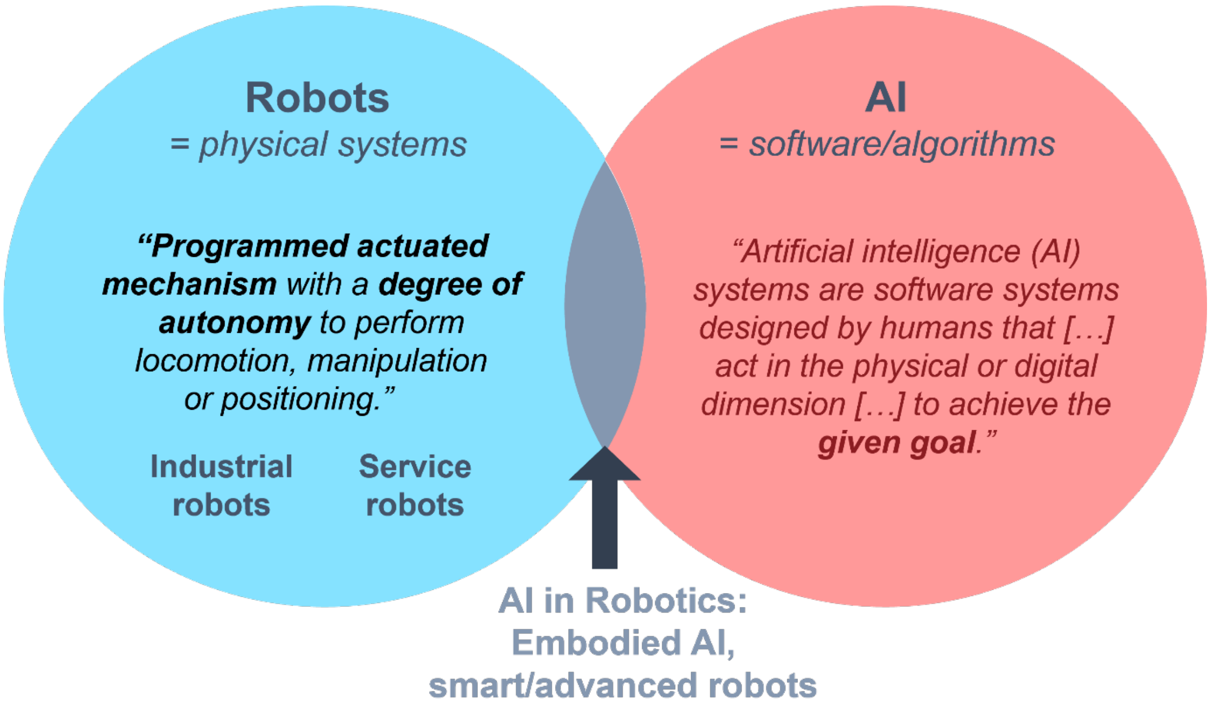

Figure 2: Differentiation and overlap between Robotics and Artificial Intelligence[2, 3].

It's important to note that:

- AI is not per se required to enable the above tasks but can bring significant performance benefits and enable specific functionality. A general rule of thumb is that the more variability and unpredictability there is in the external environment, the more useful AI will be.
- In real-world robotic applications, AI algorithms are combined with non-AI software programs as well as hardware such as sensors and cameras to execute the required task. In contrast to software applications such as recommendation engines, AI is only ever one component of a robotic application. This is particularly relevant to the question of robot safety, which to date is generally governed by hard-coded, deterministic algorithms.

IFR International Federation of Robotics

Background Information

Autonomy in Robotics

Until recently, most robots were hard-coded to execute a task according to a pre-defined trajectory and with a pre-defined level of force. These robots are oblivious to their external environment. This means that the object the robot works on (such as a car part) must always be presented in exactly the same position, and the robot is not able to adjust its force or motion – for example to stop if something or someone gets in its way.

Over the last decade, ‘co-bots’, with either built-in or add-on force-torque sensors, have given robots a limited ability to sense and respond to their external environment. For example, co-bots can recognize movement in a designated zone and adjust speed – or stop – accordingly. This has enabled robots to be integrated into production lines alongside humans, with benefits we discuss in the IFR paper ‘Demystifying Collaborative Robots’.

Force-torque sensors also enable co-bots to adjust to minor variance in the external environment – for example adjusting force and to a flexible part or in a sanding process where the amount of force required reduces during the process as the surface becomes smoother.

The past five years have seen strong growth in autonomous robots which are able to adjust to far greater variability in their external environment. For example, as we discuss in the IFR paper ‘A Mobile Revolution: How Mobility is Reshaping Robotics’, autonomous mobile robots can not only stop if they encounter an object in their path, they can also re-plan their route and adjust their path in real-time. Autonomy does not necessarily require AI. However, the higher the level of autonomy, the greater the chances of AI algorithms being employed to categorize an unfamiliar environment and to determine the best way to interact with that environment to achieve the application’s goal (for example picking up a bottle from an unsorted bin and placing it in a rack).

The IFR has defined the following five levels of autonomy in robotics:

THE FIVE LEVELS OF AUTONOMY IN ROBOTS

	NO AUTONOMY, REMOTE CONTROL	NO AUTONOMY, NO SENSE AND RESPOND	NO AUTONOMY BUT SENSE AND RESPOND	AUTONOMY	ADVANCED AUTONOMY
DESCRIPTION	Robot is operated manually and remotely	Robot is not aware of its internal state or external surroundings.	Robot can sense force and motion and respond according to pre-determined parameters	Robot can sense its external environment and calculate a response	Robot can identify characteristics of its external environment
EXAMPLES	Remoteoperated surgical robot	Weldingrobotin automotivefactory	Collaborative assembly robot in manufacturing	Pick-and-place robot feeding a machine with parts from an unsorted bin	Assistant robot that responds differently to a child than an adult
SAFETY	User ensures safety	Sensors on cages	Settings for force and speed	Hard code overrides probabilistic algorithms	Hard code overrides probabilistic algorithms

NO AI DETERMINISTIC ALGORITHMS → PROBABILISTIC ALGORITHMS *HIGH AI*

Figure 3: Packaging boards for IKEA Slovakia, image credit: Photoneo.

Applications

In this section, we look at some of the most common applications of AI in robotics in commercial use. We differentiate between two types of application:

- **Sense-and-respond** applications in which the robot identifies and responds autonomously to its external environment in a real-time closed loop: Pick-and-place applications are the most common robotic applications in this category
- **Performance optimization** in which AI is used to optimize process design and robot programming, as well as in quality inspection and maintenance.

Sense and respond

This covers applications in which the robot identifies and responds in real-time to its external environment to complete a given task.

The 'sensing' part requires on-robot and/or external sensors and may include cameras to help the robot locate either itself, or an object it is tasked with manipulating. Machine learning algorithms of some form are typically employed here. The 'respond' part generally includes some kind of reasoning and decision-making algorithms to determine the best course of action – such as the correct force and gripper position with which to pick up an object, and the correct path to place it in correctly at the target destination. These algorithms may or may not be AI – or will most likely be a combination.

Generally, the more variability there is in the input ('sensing') data, the more useful AI will be in planning the appropriate response, such as the correct path and force with which to approach and grip an object that is randomly placed. The algorithms for sensing and responding may run on the robot or robot system (for example, a robot and vision system), or in some form of higher control system, or in a mixture of these. If the response does not require coordination with other robots, the algorithm may run on the robot. However, in decisions involving multiple robots, the algorithm will run in the control system. In mobile robots, for example, fleet management systems are used to control groups of robots. AI might be used in these systems to determine which robot to send to perform which task.

Below are the most common sense-and-respond applications.

Pick-and-place

As the name suggests, pick-and-place applications require the robot to pick up an object from one location and place it in another. This can be done with traditional, deterministic, programming using exact coordinates if the object to be picked is always the same, arrives in the same position and is placed in the same position, for example on a feed conveyor. Deterministic programming can also be used to pick objects that are the same but are unsorted.

Video 1: Machine tending at Allied Moulded with ActiNav.

For example, Universal Robots' ActivNav system uses statistical matching to pick similar, but unsorted, parts from bins for machine tending. The system uses a CAD model of the part to be picked as input. Service technicians assess the best pick points based on the object and the type of gripper used – a vacuum gripper requires a pick-point where suction can be applied over the whole surface of the pick-point for example. ActiNav uses a laser-range finder to produce a point cloud that enables identification the object. When an unsorted bin arrives, a 3D picture is taken and a point cloud generated. Statistical matching determines the likelihood of a potential object identified in the point cloud being the object to be picked. Path-planning algorithms determine the best path for the robot arm to approach and pick the object.

The general rule that AI – in this case machine learning – offers benefits as the level of variability increases applies here. For example, transparent, thin or flat objects are difficult to identify as the 3D camera will categorize them as flat surfaces rather than as 3D objects. AI also offers benefits in situations in manufacturing, logistics and retail in which parts (or parcels) to be picked are mixed and unsorted in bins,or loaded differently each time on pallets. AI also enables vision systems to correctly identify the outlines of tightly-packed, similar objects such as packages of the same color and surface or textured surfaces such as wood. Finally, AI is useful in environments with high variance in lighting conditions.

There are a number of applications within the broad category of 'pick-and-place' which require slightly different capabilities:

- **Pick-and-drop**: The robot picks items from one unsorted bin and drops them into another
- **Put-wall**: The robot picks objects from a mixed bin, scans the barcode on the object which indicates which cubbyhole it should be placed in, and then places accordingly
- **Induction sorting**: The robot picks products or packages and places them on an empty spot on a conveyor
- **Order fulfilment**: The robot picks items from multiple bins or cartons and assembles them into the order

For example, Obeta, an electrical supply wholesaler based in Berlin, has implemented AI from Covariant to compile orders from cartons on a conveyor. Obeta's warehouse receives hundreds of new products each week, so the company needed a system that could easily adapt to new items without disruption. Products come in different packaging and quantities and have to be carefully gripped and packed into dispatch containers or cardboard boxes, both quickly and continually. Obeta has struggled to find enough workers, particularly during the COVID-19 pandemic, and the robot, which can pick 600 objects per hour, ensures continuity of operations.

Figure 4: Order fulfilment at Obeta.

IFR International Federation of Robotics

Background Information

Developments in machine learning drive adoption of sense-and-respond applications

Machine-learning algorithms that power sense-and-respond robotic applications are advancing rapidly. Initially, algorithms had to be trained on hundreds or thousands of images of each target object in order to recognize it correctly. This is a time-consuming process that accounts for a substantial portion of the cost of the application. The algorithm must be retrained with tagged examples each time a new product is introduced. This is not economically viable for companies with high product variability – for example wholesalers with millions of products, many of which change on a daily basis.

Significant development effort has been therefore directed by AI software and systems providers to enable algorithms to learn with fewer, or no, tagged examples. This is achieved through semi-supervised or self-supervised learning methods. Semi-supervised learning takes a small set of tagged images as a base, from which the algorithm trains itself. Self-supervised learning is more complex and uses labels that already exist in the input data (for example, specific sounds or colors). The knowledge accrued by the system integrator or software provider in their pick-and-place applications, and by the end-user, naturally expands over time, reducing the training and installation time of an AI-based pick-and-place application for a specific domain, for example packaged food.

The ultimate aim is to enable an algorithm to quickly generalize from what it has already learned, applying existing knowledge to recognizing and manipulating new objects. For example, the more exposure an algorithm has to objects tagged or identified as 'bottle', the more quickly it can identify and pick bottles of different shapes and materials that it has not encountered before from a mixed bin, applying the right force and grasp. Training on different objects, for example clothing, can also strengthen the neural network (the underlying computing model) for bottles, increasing the 'generalizability' of the algorithm.

Though we will therefore see increasingly generalizable algorithms (powered by richer neural networks) this does not imply that AI-driven robotic applications can be easily transferred from one environment – such as bin-picking – to another – such as welding. Robots comprise many different hardware and software components which must be integrated, and robots are themselves typically integrated with enterprise-wide systems such as fleet management, enterprise-resource planning, or manufacturing execution. The interfaces between components of the robot and external systems must all be programmed or adapted for a new application.

Picking technology is advancing rapidly in terms of both speed and reliability. To be economically viable, picking robots need to be faster than human workers, while most companies require an accuracy of around 99.5%. It is still difficult for robots to pick objects that are not rigid – for example, goods in plastic wrapping, fruit and vegetables in nets, or floppy materials – with an accuracy and speed that is commercially viable, but we can expect this to change over the next 5 years.

Additional robotic applications using AI use elements of the pick-and-place functionality described above, tailored to the specific application requirements.

Palletizing, de-palletizing and packaging

Robots are used to take objects off pallets on arrival, and load pallets and packages with objects for shipment. Robots have traditionally been used to palletize and de-palletize objects of standard sizes, while employees unpack and load non-standard goods. Using AI algorithms for object recognition and path planning, robots are starting to be used for palletizing and de-palletizing of non-standard goods. Palletizing objects of variable weight is particularly tricky as it requires the algorithm work out the best order of placement, so that heavier items, even if smaller, are loaded on the bottom layer of a mixed pallet.

For example, LQ Group, a fast-moving consumer goods company, is using AI technology from MechMind to de-palletize goods and place them on a conveyor. LQ receives thousand kinds of cartons of various sizes, colors, patterns, shades, which would be costly or impossible to program individually. In the MechMind solution, once the pallet arrives at a certain position, the host control system sends the details of the cartons on the pallet including size, weight and quantities to the robot. As soon as the robot receives the arrival signal from the host system, it triggers the camera to take a picture and sends the picture to the vision control system, which combines the coordinates and poses of cartons into grasping points. These are then processed by AI algorithms which send instructions back to the robot's gripper including how many suction cups have to be activated in order to pick a specific number of cartons, which specific cartons have to be picked on the same layer and the coordinates for the position of the box when it is put on the conveyor (https://ifr.org/case-studies/smart-3d-vision-solutions).

Video 2: AI-assisted palletizing and de-palletizing with Mech-Mind AI solution.

Similarly, vision systems provider Photoneo uses machine learning for de-palletizing cartons of different sizes, that may be tightly packed or have shiny surfaces. The cartons can be picked individually, versus the traditional process of de-layerization, where the pallet contents are picked up as one and placed on the conveyor. The pallet is scanned with a 3D scanner and the algorithm, which is pre-trained on 5,000 box shapes and constantly learns new ones it encounters. A more sophisticated convolutional neural network (CNN) algorithm can be used to recognize boxes with problematic surfaces, including varying textures, shiny or reflecting material, protruding tapes, patterns, or cartons with black covering, which can also cause problems. A command is then sent to the robot which uses a specially developed universal gripper enabling it to pick with an accuracy of +-3 mm. This way, it is able to unload 1,000 boxes in our hour, with 99.7% pick-rate accuracy.

Video 4: De-palletizing using an ABB robot and Photoneo 3D vision system with machine learning

Machine tending

AI can be used to enable robots to feed parts from unsorted bins to machines and is particularly useful when there is high variability in the application environment, for example lighting, bin sizes and position and mixed bins.

Welding

Industrial robots are widely used for welding when the parts are the same and the weld points therefore in the same position. However, the automation of welding has traditionally been too costly for manufacturers with high part or product variability. Machine learning can be used to identify welding points and weld path.

Assembly

AI is also being trialed in the assembly process in manufacturing. Assembly tasks such as snapping parts into place, gluing parts or inserting a windscreen into a car, are typically carried out with hard-coded parameters but AI can be used to enable these tasks in variable environments. The same technology described for pick-and-place and welding applications can be applied to identifying the coordinates for snapping a part into place and planning the correct path and force for example.

Assembly tasks are generally more complicated than pick-and-place and machine tending, often involving multiple steps. It's likely that AI will first be applied to very specific tasks - for example, enabling a robot to locate the exact position for placing a screw - within an assembly process. IFR members believe larger-scale adoption of AI in assembly will take between five and ten years.

Recycling

Robots are starting to be used in the recycling of waste materials, improving productivity, and saving workers from unpleasant and often dangerous tasks. Some, such as those from AMP Robotics, use AI to recognize objects in a conveyor of mixed objects, for sorting according to type of waste.

Figure 5: Waste recycling, image credit: AMP Robotics.

Mobile Robots

As we discuss in our information paper, '[A Mobile Revolution: How mobility is reshaping robotics]' there is currently a boom in mobile robots, particularly autonomous mobile robots (AMRs) which navigate autonomously and are used in a wide variety of industry sectors including manufacturing, logistics, retail and healthcare. These robots perform a range of tasks, from fetching and carrying goods and parts in manufacturing, transporting linens and medicines in hospitals, tracking stock in supermarkets and serving as 'mobile assistants' in shops, hospitals and public spaces, providing information and enabling remote interaction with specialists.

Manufacturers and logistics providers with high product variability and turnover ideally want to move from fixed racks - where the same product is stored in the same place - to dynamic packing spaces where the size of the packing area, and the products within it, vary continuously and intelligent mobile robots are key to executing this vision.

Autonomous navigation means these robots locate themselves within, and simultaneously create or update, a map of their surroundings. They also use path planning algorithms to determine the best route to their target destination. AI is not needed for autonomous navigation or path planning but AMR providers such as MiR and Fetch Robotics are building AI into their AMRs to improve obstacle recognition and response.

AMRs use sensors to detect if an obstacle is in their path, in which case they will slow or stop. The speed of slow-down or stop is the same for all obstacles. Using 3D cameras and machine learning algorithms enables AMRs to recognize the object in front of them and tailor their response accordingly. For example, the AMR would stop and give priority to an automated guided vehicle that might either hit it or cause an obstacle if it stopped in the AMRs path. On the other hand, it might slow down but not stop if a human approaches and is able to walk round the AMR. If it recognizes another AMR, it can predict that vehicle's motion and adjust its own path, but not slow down.

Academic researchers are also looking at the use of machine learning to optimize path planning, but this is in early stages.

Precision agriculture

There is a nascent market for robots in commercial agriculture. Vision systems using machine learning algorithms are a key component of many agricultural robotic systems. There are a number of start-up companies in precision agriculture focused on planting, weeding, watering and the targeted application of fertilizer and pesticide. Machine learning is used to distinguish between weeds and plants so that only the weeds are sprayed by the robot. Manufacturers of these systems claim precision spraying reduces carbon emissions as it uses 95% less herbicide than conventional spraying, which covers the whole cultivated area [4].

Figure 6: The Avo precision weeding robot from Ecorobotix.

The use of robots to harvest produce is still in very early stages but AI plays a key role. Harvesting applications are a form of pick and place except the produce is no longer in a bin, but unstructured on a plant. Machine learning is used to identify the produce to harvest, as well as to determine whether it is ripe. AI is also used for path planning for the robot's gripper, enabling it to pick the produce with the right grasp and force and place it in a bin or other container.

Assistance robots

Robots are increasingly making an appearance in public environments such as airports, supermarkets and shopping malls, hospitals and care homes, providing a range of assistive functions. Robots in retail spaces can provide production information and lead shoppers to the right aisle for the product in question, for example. In airports, robots can update passengers on gate information, and also guide them where necessary through the airport. In hospitals, robots can enable remote consultations and undertake simple diagnostic procedures such as measuring temperature or blood oxygen levels. These applications all involve mobile robots and are described in more detail in the IFR paper '[A Mobile Revolution: How mobility is reshaping robotics]'.

The extent to which AI is used in these applications depends on the exact functionality required. Even within one sub-task, AI algorithms are likely to be combined with traditional algorithms to achieve the required sense-and-respond functionality.

For example, the following video shows the Lio personal assistance robot which supports care-home residents and staff in a number of tasks such as greeting patients, grasping and carrying objects, offering drinks, reminding patients of, and accompanying them to, upcoming appointments.

Video 3: Lio assistance robot

These tasks use AI algorithms as part of larger programs to accomplish specific tasks. For example, the task of serving a drink comprises a number of sub-functions including: navigation; obstacle avoidance; object recognition to distinguish between a person and a table; facial recognition to serve the right drink to the right resident; placement tasks – recognizing and locating the table on which the drink is to be placed, locating where to place and planning the path of the robot arm and; general behavioral tasks such as recovering from unplanned obstacles or activity.

F&P Robotics, which has developed the Lio, uses AI within some but not all of these sub-functions. For example, AI algorithms are used within the object recognition model. Open-source, pre-tagged object models are used to train the algorithm to recognize common objects such as a person and a table, while others are trained specifically for the application, and the robot is programmed to ignore objects it cannot recognize. AI is

[4] *See https://www.ecorobotix.com/en/avo-autonomous-robot-weeder/*

used for facial recognition, trained on video data of the resident to ensure the input to the algorithm comprises images from different angles. The Lio also has speech synthesis and voice recognition capability, using AI to enable the robot to understand the instruction given. The placement function – which bears some similarity to the pick-and-place functionality described above – also uses AI for placement identification, but not for planning the path of the robot arm to the identified spot on which to place the drink. AI is used extensively in the behavioral model which enables the robot to respond appropriately to the situation it encounters. For example, if the robot identifies a crowd of people, it might say something, but not if it cannot identify a person in the vicinity. If the robot approaches a door, AI algorithms are used to enable it to identify whether the door is open and then take the correct course of action – opening it if closed, for example, starting with locating the handle.

In many of the above tasks, some of the required functionality could be 'hard-coded' – for example, the question of how to respond to encountering a crowd of people could be coded through traditional 'if/ else' algorithms. However, AI can provide greater flexibility for classifying a given situation and determining the appropriate response. The choice of an AI algorithm is therefore specific to the task and influenced by the level of complexity in the environment and the number of potential response options.

Performance optimization

AI is used to optimize process design, robot programming and robot maintenance, as well as in robotic quality inspection. In general, these applications run off-line – the data is collected from the robot and other machines, analyzed, and then the robot program is adjusted. We explore these in more detail below.

Process design

Simulation is often used to design optimal automation processes before they are implemented. Simulation programs are typically based on computer-aided-design (CAD) programs of the machines involved. Machines, including robots, generally perform differently in the real-world to in model environments due to external effects such as vibration from other machines in the production cell. Traditionally, programmers have to observe simulations when transferred to the real world and adjust the simulation in an iterative loop until the real-world applications works properly. AI can be used in designing the initial simulation to predict the optimal robot path and other parameters, and then to identify the best performing simulation from multiple prototypes. Simulations can be run in parallel, and speeded up, providing the AI algorithms with far more data, much faster, than could be achieved by relying only on data from real-world tests. This reduces the number of iterations and fine-tuning required in real-world testing. Companies can also simulate multi-robot systems in which AI is used to optimize robot programs based on the data from all of the robots in the system. Using AI for simulating automation scenarios is still at a very early stage but holds potential, both in industrial robot applications and also in service robot applications, where human-robot interaction can be simulated in a completely risk-free environment.

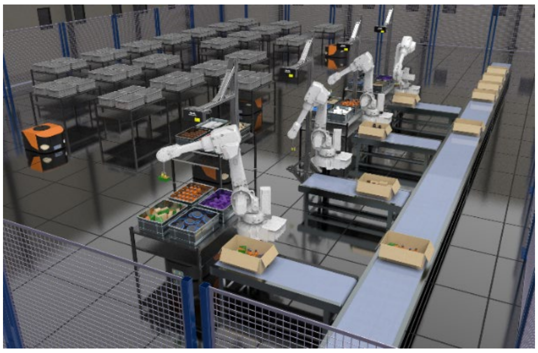

Figure 7: Order picking simulation, image credit: MechMind.

Robot programming

Experts see promise in using AI to help program robots. Programming and integration account for 50-70% of the cost of a robot application. Re-programming costs have traditionally made robotic automation too costly for many manufacturers with short production runs of a wide range of products. Methods to enable faster programming and program re-use are therefore key to robot adoption in variable environments. Experts estimate that AI could halve the resources required for programming.

While industrial robot applications requiring extremely high levels of precision and speed are typically coded by a programmer, many collaborative robot applications that do not have very high cycle time or precision requirements can also be programmed by demonstration.

The robot arm is moved through the steps to be performed to create a program which is then fine-tuned through a tablet-based interface with intuitive touch-based controls. AI can be used in the fine-tuning process to determine the required movements of the robot to perform the task optimally. Ultimately, this means a robot operator with no robot programming skills will be able to teach a robot from any manufacturer.

Some robot manufacturers are already offering this functionality [5]. Neura Robotics is trialing the use of vision systems to program by demonstration. For example, Neura Robotics is working on an application for welding in which an operator simply points to the weld locations and the robot's algorithms detect the exact welds and plan the weld path. This saves around a day's programming time. Additionally, once the algorithm has been trained, it can be applied to parts of different sizes and weld positions [6]. Neura Robotics also uses Natural Language Processing to enable voice instruction. The algorithm is trained using a demonstration of the task in which the demonstrator describes what they are doing – for example picking up a box of biscuits – so that the robot can execute this task when given a verbal instruction.

AI algorithms for tasks such as for path planning, collision detection and picking are also embedded by some robot systems providers, such as MechMind, into their programming interfaces to enable code-free or code-light programming [7]. AI could also be used in future to help match the right program from a suite of applications to a task specified by the user.

Figure 8: Intuitive tabled-based creation and optimization of record paths, image credit: Neura Robotics.

Optimizing existing programs

AI is used to optimize robot programs once running. Data on the robot's movement collected from sensors on and around the robot can be analyzed with AI algorithms to detect the optimal movements to achieve the task. For example, FANUC has introduced 'AI path control function' to improve the precision of cutting and welding with robots, adjusting for external factors such as vibration. FANUC's AI path control function estimates a robot's path from an acceleration sensor and the amount of deviation from the command path. It provides appropriate compensation to achieve high-precision circular and straight-line paths – a task that previously required hours of trial and error by skilled operators [8].

ABB is applying AI algorithms to the analysis of presses and robot behavior in press and stamping lines to minimize equipment waiting times. Using a holistic approach, a control algorithm identifies bottlenecks and manages the start and stop times of robots and presses, thus making lines more stable and predictable[9].

---

[5] *See for example Universal Robots' ActiNav.*

[6] *See http://neura-robotics.com/*

[7] *See MechMind's Mech-Viz Intelligent Robot Programming Environment*

[8] *See https://www.fanuc.co.jp/en/product/new_product/2020/202003_aipathcontrol.html*

[9] *See https://new.abb.com/news/detail/56162/innovation-highlights-2020*

Finally, AI can be used to analyze data from multiple robots performing the same task to optimize the robot program for all of them.

Predictive maintenance

As the name suggests, the goal of predictive maintenance is to predict when maintenance on a part or machine will be required before the part or machine underperforms or fails.

Figure 9: Automotive assembly with an ABB robot and vision system.

Traditionally, maintenance is conducted either on a timed basis, regardless of the actual state of the machine, or when a machine breaks down. Machine downtime can cost large manufacturers over a million dollars per hour. In predictive maintenance, sensors on the machine and externally, if necessary, capture data on a continuous basis about the machine's condition – for example, vibration, noise, and gear speed – which can be analyzed to predict the timing and scope of maintenance. Machine learning is used to establish a model of 'normal' performance so that variations from this norm can be identified, flagged and addressed. Again, AI is not required per se to enable predictive maintenance – parameters for values representing acceptable ranges of performance could be hard-coded, for example. However, AI brings benefit in identifying these parameters where they are not known exactly, and when they may vary in different situations such as external temperature.

A number of robot manufacturers, such as ABB [10] offer predictive maintenance services (sometimes also termed condition-based maintenance) for their robots.

A more recent development is prescriptive maintenance in which AI is used to work out what is most likely to be wrong once a potential maintenance issue or other performance anomaly has occurred. Prescriptive maintenance uses predictive algorithms to determine the solution most likely to be effective. Prescriptive maintenance applications are already used by operators of large plants but are not yet established in robotics.

Quality inspection

Quality inspection is, alongside pick-and-place, the most frequently adopted use of AI in robotics currently. Using 3D vision and machine learning algorithms, robots can inspect parts during the production process. Faulty parts can be taken out of the production line before they are moved on to the next station, saving costs in unnecessary further processing.

It is difficult to code exact parameters that would enable the detection of all instances of flaws such as scratches that are different each time. AI algorithms, in contrast, can generalize from many different examples. The algorithm continues to fine-tune itself as it encounters more examples. In some cases, 3D vision can detect product flaws that are too small for the human eye to identify.

The same functionality can be applied to error-proofing which reduces wastage by detecting and correcting faults at the beginning of a run, before a lot of parts are produced. FANUC offers an error-proofing application [11] - shown in this video - that checks, for example, whether a nut that a robot should tighten is actually in place.

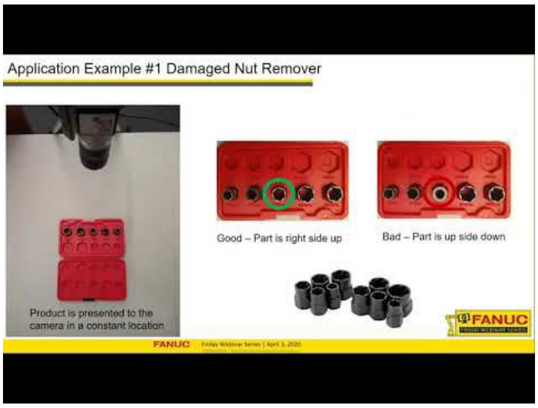

Video 4: Error-proofing with FANUC application.

[10] See https://new.abb.com/news/detail/73275/prsrl-abb-launches-condition-based-maintenance-service-for-fleet-and-individual-robot-assessments

[11] See https://www.fanucamerica.com/news-resources/articles/new-artificial-intelligence-error-proofing-features-machine-learning-technology

Some forms of quality inspection are difficult to automate – particularly the quality of finishes such as polish or paint. Though scratches can easily be detected using AI, it is difficult to train an algorithm to distinguish qualities such as overall smoothness, which looks different in different lighting conditions and for which human quality inspectors would use touch. Digitalizing this kind of sensory or haptic data to provide input to an AI algorithm has so far proved difficult. We can therefore expect final inspection of many finished surfaces to remain a manual operation for now.

Plant inspection

Video 5: Mobile robot inspection at BASF.

A number of chemicals and oil & gas companies such as BASF in Germany and Shell [12] are trialing applications for plant inspection using robots equipped with vision systems. A 3D camera is mounted on a mobile robot (which can be legged, with tracks, aquatic, or a drone), which moves through the plant looking for anomalies such as leaking pipes or gauges showing an abnormal reading. Subsea inspection using robots is already well established, but the use of legged or tracked mobile robots for plant inspection is newer. ANYbotics [13] and Energy Robotics [14], for example, offer AI-enabled plant inspection using legged robots. The Sprint Robotics Collaborative [15], a consortium of chemical and petrochemical companies and inspection robot providers, was launched in January 2015 to accelerate the development of robotic inspection applications for production plants.

The impact of AI in robotics on work and jobs

Despite initial headlines predicting the number of jobs that will be lost to automation, the majority of economic studies support a very different narrative – that automation leads to a net increase in jobs within a country and to an increase in gross domestic product (GDP). It has also been recognized that very few jobs can be fully automated in an economically viable way. Rather, specific tasks within them can be automated, altering the nature of the job but not eliminating it in the majority of cases. The changes in skills requirements driven by automation apply at all levels of the workforce.

AI expands the potential for robots to share tasks or processes with workers, taking on those parts of the task or process that are unergonomic and repetitive, such as lifting, fetching and carrying. These applications do not necessarily require AI, but, as we have discussed above, AI technologies enable the robot to work effectively in unpredictable or rapidly changing environments.

See the [IFR's positioning papers](#) on 'The Impact of Robots on Productivity, Employment and Jobs' and 'Robots and the Workplace of the Future' and 'Next Generation Skills: Enabling Today's and Tomorrow's Workforce to Benefit from Automation' for more information on how robots affect workers and jobs.

Future trends

Market adoption

As discussed, AI is currently most established in pick-and-place applications. While market

[12] See https://www.shell.com/energy-and-innovation/digitalisation/digital-technologies/robotics.html

[13] https://www.anybotics.com/anymal-autonomous-legged-robot/

[14] https://www.energy-robotics.com/

[15] See https://www.sprintrobotics.org/

adoption is difficult to predict accurately, experts consulted for this paper foresee the following broad timeline for the adoption of AI in commercial robotic applications:

Next 5 years

- Pick-and-place applications
- Machine-tending
- Welding
- Quality inspection
- Agriculture (precision application of pesticide and fertilizer)
- Process optimization (e.g. fine-tuning path planning based on external environment)
- Medical screening: Robots with screening devices (for example ultrasound) can be used to scan the required body part, with AI algorithms supporting the robot's path planning, as well as analyzing results of the scan
- Recycling

5 to 10 years

- Assembly tasks in manufacturing
- Clothing manufacturing
- Agriculture (picking of crops)
- Professional service assistant robots
- Construction robots
- Laboratory automation

Longer-term

- Personal /household service assistant robots
- Quality inspection for functions requiring haptic feedback (e.g. smoothness)

Research directions

There is a huge body of research directed to using AI to advance all of the functionality discussed in this paper as well as enabling new functionality. As discussed earlier, much research is being devoted to training AI algorithms faster and enabling them to generalize better so that the applications using them can be quickly and cost-effectively adapted to new situations within the same broad category. Research is also going into enabling robots to generate their own program by watching a demonstration of the task they should complete as well as understanding voice commands.

Another important research area is semantic intelligence aimed at enabling the robot to understand the properties and characteristics of the object or person it is interacting with and make appropriate decisions. For example, a mobile assistance robot might in the future be able to respond differently to an elderly person than it would to a child. Semantic intelligence has most potential for applications in which mobile service robots come into contact with objects and people in public environments where it cannot be known in advance what and who the robot will encounter. However, it involves an extremely complex series of steps. The robot must first identify the properties of the object - for example, person who is a child - then detect the object or person's intention – for example passing by or approaching with the intent to interact – and then work out how to respond appropriately. It will be decades before full-scale scenarios of this kind are commercially viable. However, robot manufacturers are already addressing first steps – for example, a mobile robot that can distinguish between an object and a person therefore plan the speed at which it has to slow down and whether to issue a verbal warning if a person is in its path. Semantic intelligence also has potential in some industrial applications, for example plant inspection. A robot with semantic intelligence would be able to identify whether a pool of liquid was spilt oil or water, taking different actions, such as sending an alert, if it identified oil.

Some research is focused on 'swarm' robotics, where robots coordinate their movements among themselves, without input from a higher control system. Target applications are ones in which humans are not in the area of operation, such as excavation of disaster sites or agriculture, rather than indoor industrial environments. There are few if any commercial applications to date that use swarm models with no central control.

Safety standards and certification for AI in robotics

There is currently much discussion on how to ensure AI applications do not lead to undesirable outcomes, particularly bias in algorithms that power recommendations such as who is a suitable job candidate. In most of these instances it is important to be able to establish what data the algorithm was trained on to ensure the data is representative as well as how the algorithm reaches its conclusions. Much of the proposed regulation around AI – such as the European Commission's recent 'Proposal for a Regulation on a European approach for Artificial Intelligence[16]' – is focused on ensuring fairness and 'explainability' in AI algorithms.

Robots are not currently used in situations that involve bias. The main concerns for robot manufacturers and users when it comes to AI are safety and certification which we discuss below. Additionally, robot manufacturers using vision systems must be able to ensure that any data transferred from the vision system does not contravene data privacy laws, for example, by enabling the identification of a specific individual.

Safety

Because the output of an AI algorithm is not known in advance, users may be understandably cautious about the impact on safety. However, AI can never be used alone in a robot, as there are many other program layers needed to control aspects of the physical robot system – such as the movements of the robot's axes. This means robot programmers have to determine which layer in the 'stack' has priority over others. Priority can be given to the hard-coded, deterministic layer responsible for actions such as ensuring a robot stops if 10cm away from any object, no matter what the AI algorithm would otherwise determine to be the best course of action. This is almost always the case in the robotic applications using AI in operation today. However, as applications advance - especially in applications where the robot is physically interacting with a person whose position and movement are not pre-determined or in unpredictable environments such as public spaces - we can expect to see AI employed in safety-critical components of the application.

Enabling more variance in the robot's reaction to safety-critical situations could also offer productivity benefits in industrial environments. For example, if a mobile robot can identify that a worker is within a safety-critical zone for the robot, but on the other side of the robot arm, it could keep moving, whereas in a hard-coded application, it would always stop. Research is being directed at how to ensure safety in these situations. For example, some research focuses on 'learning controllers' which continuously process data to determine the robot's next move, followed by processing of the data generated as a result of that action, in order to determine the next move.

Currently, ISO safety standards exist for industrial robots[17] – including a technical specification for industrial robots used in collaborative applications[18] – certain service robots[19] and mobile robots[20]. These standards enable robot manufacturers to certify their robot as inherently safe. Since the safety of a robot application is dependent on many other factors – how well-lit is the production cell, is the robot manipulating a part with sharp edges – users must carry out risk assessments of their applications and are

---

[16] *See https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-european-approach-artificial-intelligence*

[17] *ISO 10218-1:2011 and ISO 10218-2:2011 - Safety requirements for industrial robots*

[18] *ISO/TS 15066:2016 - Robots and robotic devices — Collaborative robots*

[19] *ISO 13482:2014 - Robots and robotic devices — Safety requirements for personal care robots*

[20] *ISO 3691-4:2020 - Industrial trucks — Safety requirements and verification — Part 4: Driverless industrial trucks and their systems. Please also refer to ANSI/RIA R15.08 for safety of industrial mobile robots.*

responsible for the safety of their workers, governed by national health and safety laws.

It is currently unclear how safety standards and regulation will evolve for scenarios where AI is used in safety-critical components of the robot application.

Certification

Certification is important for robot manufacturers and systems integrators to demonstrate that specific safety or other standards are met and, as a result, to differentiate themselves from non-certified competitors.

There is currently considerable debate and research on the possibility and benefits of software certification or verification for programs partly or entirely composed of AI algorithms. Certification relies on certainty that the desired outcome will be met, which is difficult in the case of AI where the outcome may not be known in advance, or possible to explain after the fact.

Beyond certification, understanding why an algorithm reached a decision, and what data was important in making the decision, could also support better re-usability of algorithms for other purposes. Again, research is being devoted to this. For example, it would be possible – but computationally expensive – to determine and monitor parameters for the neural network governing the robot's motion, ensuring the robot's actions remain within a pre-determined safety range, even when the precise action that the robot will take is not pre-determined or explainable. The Fraunhofer Institute for Manufacturing and Automation is devoting research to 'dependable AI' in robotics [21].

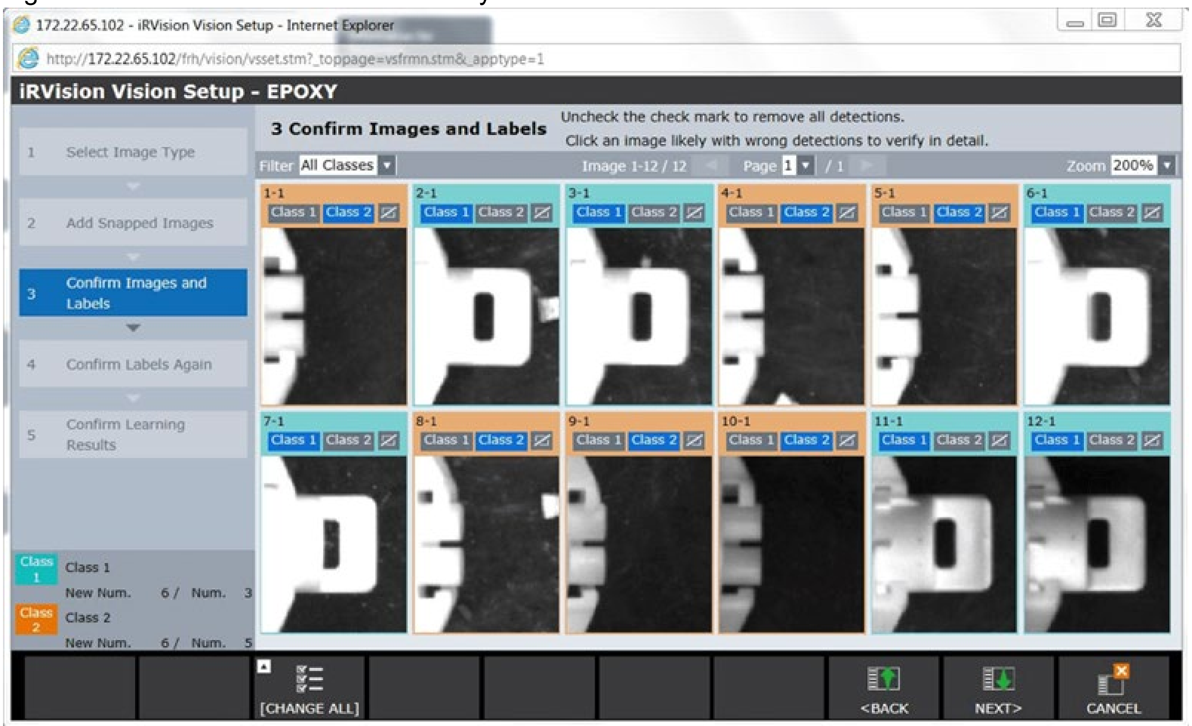

Figure 10: Image classification for supervised learning. Image credit: FANUC.

Conclusion

Artificial intelligence opens up new possibilities for robotic automation, particularly in environments with high variability. In manufacturing, AI is enabling the automation of a number of tasks involving the picking, placing and manipulating of objects, from machine tending to assembly, that have

[21] See https://www.ipa.fraunhofer.de/en/about-us/guiding-themes/ai/Dependable_AI.html

previously only been possible to do manually. Employees are spared strenuous, heavy or unergonomic tasks. AI is also becoming well-established for robotic quality inspection.

AI-driven robotic applications bring greater efficiency to logistics and retail, enabling companies in these sectors to cope better with peaks in orders, high product variability and an often-unreliable labor supply.

Robots are increasingly making their mark in public domains, from hospitals to shopping malls. In future, AI will enable better interaction between robots and the people and objects they encounter. AI will also help drive robotic automation in sectors such as agriculture which have not previously adopted robots.

AI will help reduce the resources and cost required to program and re-task a robot, opening up the possibility of automation to many companies for which automation has not previously been economically viable.

However, it's important to note that these changes will take time. There are significant advancements in enabling increasing generalizability in AI algorithms - such as in the types of objects the robot can recognize when performing a particular task. However, a complete robot application involves many more program components and interfaces that are specific not only to the task, but also to the broader, company-specific automation architecture. For each application to be commercially viable, the cost of automation has to be outweighed by productivity or other gains. In many cases it will therefore be years before developments in research labs gain widespread commercial adoption.

Artificial intelligence is attracting increasing scrutiny from regulators and advocacy groups, particularly regarding the question of ensuring the avoidance of bias in AI software. In robotics, the key issue is ensuring safety. In most commercial robot applications using AI today, the robot either does not come into contact with people, or uses deterministic algorithms that override any AI in order to protect humans, for example to ensuring that the robot stops if the distance from an object or person falls below a certain threshold. Nevertheless, efforts to create frameworks and models for 'explainable AI' in industrial applications are important, particularly in enabling certification of applications using AI algorithms.

The notion of 'general artificial intelligence' in which a robot would be able to apply learning from one task – such as opening a door – to another - such as shutting a cupboard – without further input, as a human would be able to, is unrealistic in commercial environments. As we have noted in this paper, a robot application must be tailored not only to a specific task, but also to the specific environment. The AI algorithms used are therefore 'narrow AI' and the application will always need to be programmed, even if it re-uses portions of other applications or existing blocks of code.

However, as we have discussed, programming is becoming easier, faster, and more intuitive. We can expect strides in reducing the overall time, and skills level, needed to create a tailored robot application, or re-task an existing one. AI will play an increasingly important role in enabling faster application development and re-tasking, and will expand the range of tasks it is economically viable for a robot to perform. This opens up the prospect of automation to new industry sectors, and to many small-to-medium sized companies.

Case studies and videos

Pick and place

- Bin picking solution with FANUC robot and Soft Robotics gripper: https://youtu.be/N5fyTIv1WgQ
- Order fulfilment at Obeta with Covariant AI: https://youtu.be/bQKgbWJrGFY

Palletizing and de-palletizing

- Photoneo and ABB de-palletizing solution: https://youtu.be/vYLGsphsG5E
- Palletizing and de-palletizing with MechMind AI: https://youtu.be/aJhUUG4-Ir4

Machine tending

- Machine tending at Allied Moulding with Universal Robots: https://youtu.be/Fbne-Hc9g2g

Recycling

- Waste recycling with AMP Robotics: https://youtu.be/C1PEsXWl-ZM

Precision agriculture

- The Avo precision weeding robot from Ecorobotix: https://youtu.be/5vvQqqc1zHM
- Harvesting robot: https://youtu.be/WPBnRcQ0NiU

Assistance robots

- The Lio assistance robot: https://youtu.be/i47LFCMn_48

Predictive maintenance

- Mobile inspection at BASF: https://youtu.be/NOcnTFEclSo

Robot programming

- Natural Language Processing from Neura Robotics https://youtu.be/TorR-nV1EZw

Error-proofing

- FANUC error-proofing application: https://youtu.be/RCvTyKwmk-g