

DR. BERTALAN MESKÓ

---

# 100 QUESTIONS AND ANSWERS ABOUT AI IN HEALTHCARE



THE MEDICAL FUTURIST

---

# TABLE OF CONTENTS



- 4. WELCOME MESSAGE
- 5. BASICS OF AI
- 14. THE NEW ERA OF GENERATIVE AI
- 18. THE EXCITING FIELD OF LARGE LANGUAGE MODELS: THE CHATGPT REVOLUTION
- 22. AUTOMATION'S IMPACT ON HEALTHCARE
- 28. DEPLOYING AI INTO THE PRACTICE OF MEDICINE
- 33. THE BUSINESS OF HEALTHCARE AI
- 38. THE ETHICS OF AI
- 43. YOUR GATEWAY TO INTERACTING WITH AI
- 47. REGULATING HEALTHCARE AI WORLDWIDE
- 50. WHERE HEALTHCARE AI IS HEADING



© DR. BERTALAN MESKÓ  
ANDREA KONCZ  
(THE MEDICAL FUTURIST)

DESIGN: ADAM MORONCSIK  
(STUDIO VAN)

2024

# Welcome Message from The Medical Futurist

In an era marked by groundbreaking technological advancements, the healthcare industry is at the forefront of a transformative revolution. The integration of artificial intelligence (AI) and automation into healthcare and medicine has the potential to redefine the way we approach healthcare delivery, patient care, and the very fabric of our medical ecosystem.

Welcome to “100 Questions and Answers About Artificial Intelligence in Healthcare” an informative and accessible guide that aims to shed light on the intersection of AI and healthcare.

Whether you’re a business leader seeking innovative solutions, a healthcare professional navigating the evolving landscape, a policymaker shaping the future of medicine, or a curious patient eager to understand the impact of AI on your well-being, this e-book is tailored to provide you with quick, insightful, and link-rich answers to 100 essential questions.

The Medical Futurist team, driven by a passion for the possibilities of AI in healthcare, has meticulously curated this guide to offer you a comprehensive overview of the transformative potential of AI and automation. Each



question delves into a specific aspect of AI in healthcare, providing you with concise yet thoughtful answers backed by the latest research and real-world examples.

As you navigate the pages of this e-book, you’ll embark on a journey to uncover:

- The key drivers behind the adoption of AI in healthcare.
- How AI is revolutionising patient diagnosis and treatment.
- The ethical considerations surrounding AI in medicine.
- The impact of AI on healthcare policy and regulation.
- The empowerment of patients through AI-driven tools and technologies.
- Promising AI applications in medical research, drug discovery, and beyond.

Our goal is to equip you with a resource that you can return to whenever you seek clarity on AI’s role in the future of healthcare. We understand that your time is precious, so we’ve designed this guide to be both quick to read and rich in valuable insights. Each answer is accompanied by relevant links, enabling you to dive deeper into topics of interest.

The future of healthcare is here, and AI is at the heart of it. We invite you to embark on this journey of exploration, learning, and understanding. Together, we’ll demystify the complexities of AI in healthcare, and by the time you reach the final page, you’ll be well-equipped to navigate the exciting, AI-driven future of medicine.

Let’s take the first step into the future of healthcare with AI. Enjoy your enlightening journey through “100 Questions and Answers About Artificial Intelligence in Healthcare”.

Dr. Bertalan Meskó

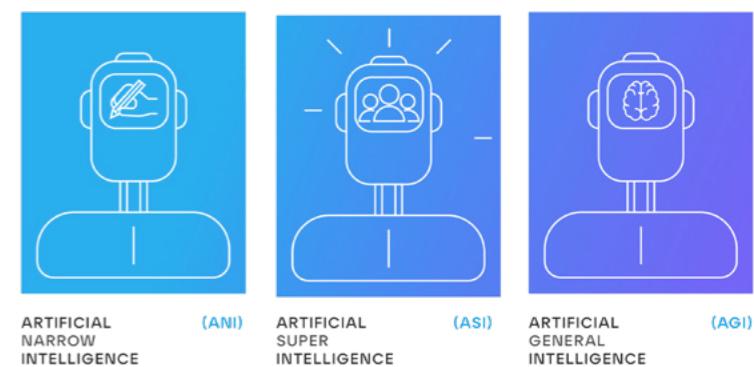
# 1 BASICS OF AI

## 1.1 What is Artificial Intelligence? 1.2 What are the levels of AI from a long-term perspec- tive?

The simplest way to define AI is this: it is intelligence demonstrated by machines. Or, to be a little more precise, the simulation of human intelligence in machines that are programmed to think and learn like humans. The term can also apply to any machine that exhibits traits associated with a human mind such as learning and problem-solving. AI is a broad field that involves various technologies, including machine learning, natural language processing, robotics, and computer vision, among others.

The goal of AI is to create systems that can perform tasks that would typically require human intelligence. These tasks include understanding natural language, recognizing patterns, making decisions, and solving problems. AI systems can range from simple algorithms solving specific problems to complex systems capable of learning and adapting to new situations over time.

We can classify AI in two distinct ways. One way is to define “how smart” the algorithm is, in which we can follow the guidance of University of Oxford philosopher Nick Bostrom. This gives us an answer we can use in the long term.



**Artificial Narrow Intelligence (ANI)**  
ANI algorithms are designed to do one, precisely defined task. This task can be anything from playing chess to spotting cancer cells on chest

X-rays. While they may do their task better than anyone in the world, their intelligence remains zero. ANI already has incredible pattern recognition abilities in huge data sets, which makes it perfect for solving text, voice, or image-based classification and clustering problems.

### Artificial General Intelligence (AGI)

this is the level of human intelligence. AGI one day could have a human being's comprehensive and total cognitive capacity, and reason, argue, memorise and solve issues like you do. Despite the advancement of AI in the past years, we are still quite far from reaching this level.

You  
Let's play a game! We create words with emojis. The first letter of the emojis spells the word. I start: 🙋‍♀️👋‍♀️🦁🦁🦉  
Can you guess?

ChatGPT  
Your sequence of emojis spells out "HELLO" (Hand, Eggplant, Lion, Lion, Owl). That's a fun and creative way to spell out words with emojis!  
Now, it's my turn. Let's see if you can guess this one: 🤖appleappleegg  
What word am I spelling?

You  
EAAE?

ChatGPT  
It looks like my hint wasn't clear enough! I intended for the emojis to represent the word "EASE" (Elephant, Apple, Apple, Egg), but I can see how the representation might not have been straightforward. Let's try another round or play a different game if you'd like!

### Artificial Superintelligence (ASI)

ASI theoretically could have humanity's combined cognitive capacity or even more. Humanity, obviously, would not be able to grasp its knowledge and understand its reasoning. Many organisations work hard to avoid ever reaching this stage.

## 1.3

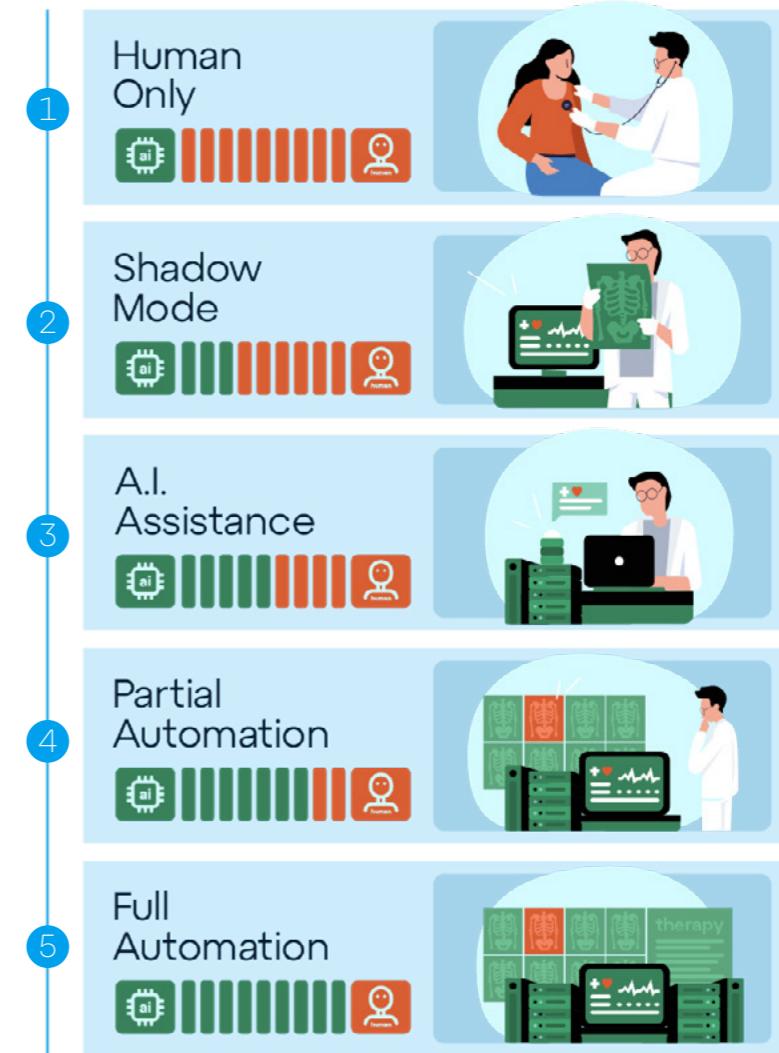
# What are the levels of AI from a short-term perspective?

We have another way of classifying AI, and it offers a more practical, more useful-in-short-term approach: to differentiate between various "autonomy" levels of algorithms. Here we can define five levels of automation. These provide an easy-to-grasp scale regarding various AI models we can use in healthcare (and in other industries).

#### Level 1 - Human only: no AI is involved.

Humans are doing the work at this base level whether it's manual work or inputting data to generate a process. These can also involve simple algorithms but not artificial intelligence. We don't need to go far to imagine this stage as most medical procedures are currently done manually. Whether it's a surgeon performing a laparoscopy or a medical researcher gathering data for a meta-analysis, humans are at the forefront with no assistance from an AI.

## The 5 Levels of Automation In Medical Procedures



#### Level 2 - Shadow mode: the physician teacher and the AI student

Similarly to how medical students follow their designated teacher around the hospital to take notes, ask questions, and do things under supervision, algorithms can also undergo such a "training" pro-

cess with a so-called “shadow mode”. For example, while a physician makes a diagnosis based on an X-ray, a “trainee” AI follows the process without interfering with it. The algorithm thereby takes notes, checks the physician’s accuracy and logs everything that can support a future diagnostic decision made by the AI itself. This can be used to further develop AI technologies that will move them along the automation spectrum.

#### Level 3 - The AI assistant.

At this stage of the automation spectrum, the AI system supports physicians in clinical decision-making via suggestions. For example, after scanning a database of chest CT scans, the AI considers the chest CT results of a patient being investigated and highlights suspicious signs. These signs are then further investigated by the physician. We see a good example of this in this large-scale Swedish trial of AI-assisted mammography, which found that the algorithm can lower the screen reading workload of radiologists by 44% without compromising diagnostics efficiency.

#### Level 4 - Partial automation.

With partial automation, an AI system can come up with its diagnosis; but if it’s not confident enough about it, the AI turns to physicians for help. Several companies are working on such solutions today. The AI-based system from Behold.ai, red dot, classifies chest X-rays and localises its findings. It can even identify abnormal chest X-rays of COVID-19 patients. It can help in ‘instant triage’ to accelerate diagnosis and allocate resources accordingly. Palo Alto-based Nines developed an AI-system that can identify potential cases of intracranial haemorrhage and mass effect from CT scans. It then flags those cases for radiologists to review.

#### Level 5 - Full automation.

As the name suggests, full automation processes are performed by an AI alone and do not involve human input. For example, a Level 5 system could analyse a mammogram on its own and request for subsequent testing without consulting a human physician for this

decision. Similarly, some scientists speculate that some ophthalmological surgeries can be fully automated since some are already partially automated. Nevertheless, other researchers believe that reaching Level 5 automation in any medical setting is “unlikely to be safely achieved in the near term.” So we have to think of it as a long-term eventuality, but such levels of automation fuels fears of A.I. replacing physicians. However, it is more likely that such AI systems will excel at a specific task and healthcare professionals will increasingly interact with them.

## 1.4

# What is machine learning?

Machine learning is a branch of artificial intelligence that empowers computers to learn from data and improve their performance over time without being explicitly programmed for specific tasks. It’s like teaching a computer to become smarter through experience.

By feeding it vast amounts of data, machine learning algorithms can detect patterns, make decisions, and predict outcomes. This approach is crucial for tasks too complex for traditional programming, such as from recognizing faces in photos to making medical diagnoses. At its core, machine learning thrives on data: the more quality data it receives, the better it gets at its assigned tasks.

It spans various techniques, including supervised learning, where the algorithm learns from labelled data; unsupervised learning, which finds hidden patterns in data without explicit instructions; reinforcement learning, where it learns through trial and error; and deep learning, which uses neural networks to process and interpret complex data structures. Machine learning is reshaping how we interact with technology, making it an essential pillar of modern AI development.

## 1.5

# How does machine learning spot cats in photos?

If you stop for a second to figure out how we can teach an algorithm what a cat is, you will very soon recognize that our typical human thinking patterns will be useless here. Unlike humans, who recognize cats through characteristics like soft fur, two triangle-shaped ears, and graceful movements, algorithms operate without such intuitive frames of reference. They perceive images not as cohesive wholes but as arrays of pixels, where what we term as “colour” translates into numerical codes.

Here is where data annotation (or labelling) enters the picture. Machine learning identifies cats in photos through a sophisticated process of pattern recognition, fundamentally rooted in the analysis of pixel data. At the heart of this process are algorithms trained on vast datasets of images, meticulously labelled by data annotators to indicate which photos contain cats.

And of course, not just cats, but also tumours on medical images. [These annotators are the unsung heroes of the AI revolution](#), providing the high-quality, accurately labelled data essential for training reliable machine learning models. Without their painstaking work, algorithms would lack the foundational learning material to distinguish cats from an array of other subjects in photos. The training involves teaching the algorithm to recognize patterns in the pixel data that are characteristic of cats, such as the texture of fur, the shape of ears, or the outline of eyes, without any innate understanding of these features as we humans perceive them.

The algorithm processes images as collections of pixels, with each pixel represented by numerical values that denote colour intensity. Layer-by-layer the machine learning model recognises more and more complex patterns in a hierarchical composition: lines, edges, curves, later gradients of simple shapes, finally the last layers are sensitive to colour-dependent complex patches of real-world images (e.g. cat faces). This hierarchy is created during the neural networks' training process from the data, without the explicit guidance from the developers.

Interestingly, the fact that a hierarchy of visual elements emerges is similar to the way the human vision works: dedicated regions of the brain detect simpler and more complex patterns on the way to recognise our kitten.



## 1.6

# What sub-types of machine learning are the most commonly used ones?

Machine learning has many subtypes and combined methods, we could hardly list them all here, so we'll only mention a few common (and important) methods.

Picture this: we've got a trio of main paths—supervised learning, unsupervised learning, and reinforcement learning—plus a secret tunnel, deep learning. They all lead to different results.

**Supervised learning** is the neat, sunlit path where everything is labelled and organised. It's like having a GPS that tells you exactly where to go. You're teaching a computer to recognise cats by showing it thousands of cat pictures, each tagged "cat." The computer learns to spot a cat anywhere, making it a perfect path for clear-cut tasks. Every-

thing is straightforward, and there's a right answer to aim for.

**Unsupervised learning** is the untamed forest trail. Here, you're dropped in the middle of nowhere and told to find your way out. No labels, no signposts. The computer looks at tons of data—say, social media profiles—and starts to notice patterns on its own. It is very similar to being told to sort through an old, dusty attic full of various objects: books, toys, photographs, and clothes, where there's no inventory list, no labels, just piles and piles of stuff. Your task is to organise these items into categories that make sense.

**Reinforcement learning** is the adventurous mountain climb. You make a move, see if you fall or find a foothold, and learn bit by bit. It's like teaching a dog new tricks with treats. The computer tries different actions, gets rewarded or scolded, and learns the best strategy to win the game or solve a problem. It's perfect for navigating complex situations where the right path isn't clear from the start.

And then, there's the **deep learning** tunnel, burrowing deep into the mountain's core. This is where the magic happens. With layers upon layers of neural networks, the computer can process and understand vast amounts of data—images, sounds, texts—mimicking the way our brains work. It's the realm of self-driving cars and voice assistants, a place where AI starts to feel less like a machine and more like a part of our world.