CAIR4

EU AI Act Playbook

Understand. Deepen. Apply.

Oliver M. Merx

AI REGULATION IS A SERIOUS MATTER,

and this Playbook* is a serious publication. Its aim is to make the essentials of the EU AI Act easier to grasp – using striking visuals to bring the core ideas to life. But why is this necessary at all? Isn't it enough if lawyers understand the AI Act – studying, teaching, and applying it through statutes and legal checklists?

The answer is: no. AI concerns everyone, not just lawyers. And the EU AI Act is the key instrument for building public trust in the responsible use of AI. But what good is the most sophisticated regulation if it is understood only by a few? If, for everyone else, it appears so abstract and complex that it confuses rather than reassures?

This is the dilemma: to ensure legal certainty, the AI Act must follow established legal principles. At the same time, it has to capture the rapid evolution of AI and translate it into rules that are both flexible and durable. That is no easy task.

What deserves to be highlighted is this:

- The AI Act is legally well-structured.
- It strikes a fair balance between the opportunities and risks of AI.
- It has the potential to guide Europe safely into the AI future.

Still, the fact remains: companies, public authorities, and their staff need to understand what obligations and rights they have. And end-users want to know why they can place their trust in AI within the EU.

That is where this Playbook comes in. It uses bold, playful methods to explain the AI Act. Playful in style – serious in intent.

Its purpose is to make the AI Act more accessible, and to make AI itself more trustworthy.

* Parts of the Playbook were translated from German with the help of AI. Legal references beyond the EU AI Act relate to German law.

Oliver M. Merx
German qualified lawyer & computer scientist

WHO IS WHO:

These AI protagonists and their interplay are explained in the Playbook:

1. AI system
2. AI components
3. AI model
4. AI provider / deployer
5. GPAI system
6. GPAI model
7. GPAI actors
8. Prohibited AI practices
9. High-risk AI system
10. Medium-risk AI system
11. Low-risk AI system
12. AI literacy
13. Training, validation and testing data
14. System data
15. Poor-quality data

LEARNING THE AI ACT THROUGH PLAYFUL METHODS

The EU AI Act is the heart of AI law. Understanding it – and applying it with confidence – requires not only legal knowledge but also a basic grasp of artificial intelligence.

Even the central question of what exactly an AI system is can be difficult to answer. The same goes for distinguishing between a quasi-provider and a so-called downstream provider: both roles are important, both sound similar – yet they mean very different things.

So, what is the difference?

Such issues are usually explained by lawyers in a dry, technical manner: through the wording of the law, abstract criteria, and legal checklists. While correct in substance, this approach often leaves key questions unresolved.

The Playbook complements the traditional legal method. Drawing on established learning techniques, it helps overcome the abstract norms of the EU AI Act and makes it easier to place AI topics in their proper legal context – largely without resorting to jargon.

Only at the end does the Playbook point to the legal literature, with references to specific provisions and paragraphs. Before diving into which symbols and terms correspond to which norms of the AI Act, readers are encouraged to first en-gage with the Playbook itself – and let its metaphors sink in.

At the core of the Playbook is a symbolic language: addressing questions of AI and AI law often requires interdisciplinary exchange between lawyers, managers, engineers, and end-users. Each speaks a different professional language. In international projects, foreign languages add yet another layer.

Supporting this demanding communication with clear, visual methods is one of the Playbook's central aims. It is therefore intended for anyone who wants to understand the basics of the EU AI Act – or who is tasked with putting it into practice.

Now let us begin the Playbook journey. The 15 AI protagonists and their interplay are presented across the following eight stations.

EIGHT STATIONS

1. METHOD

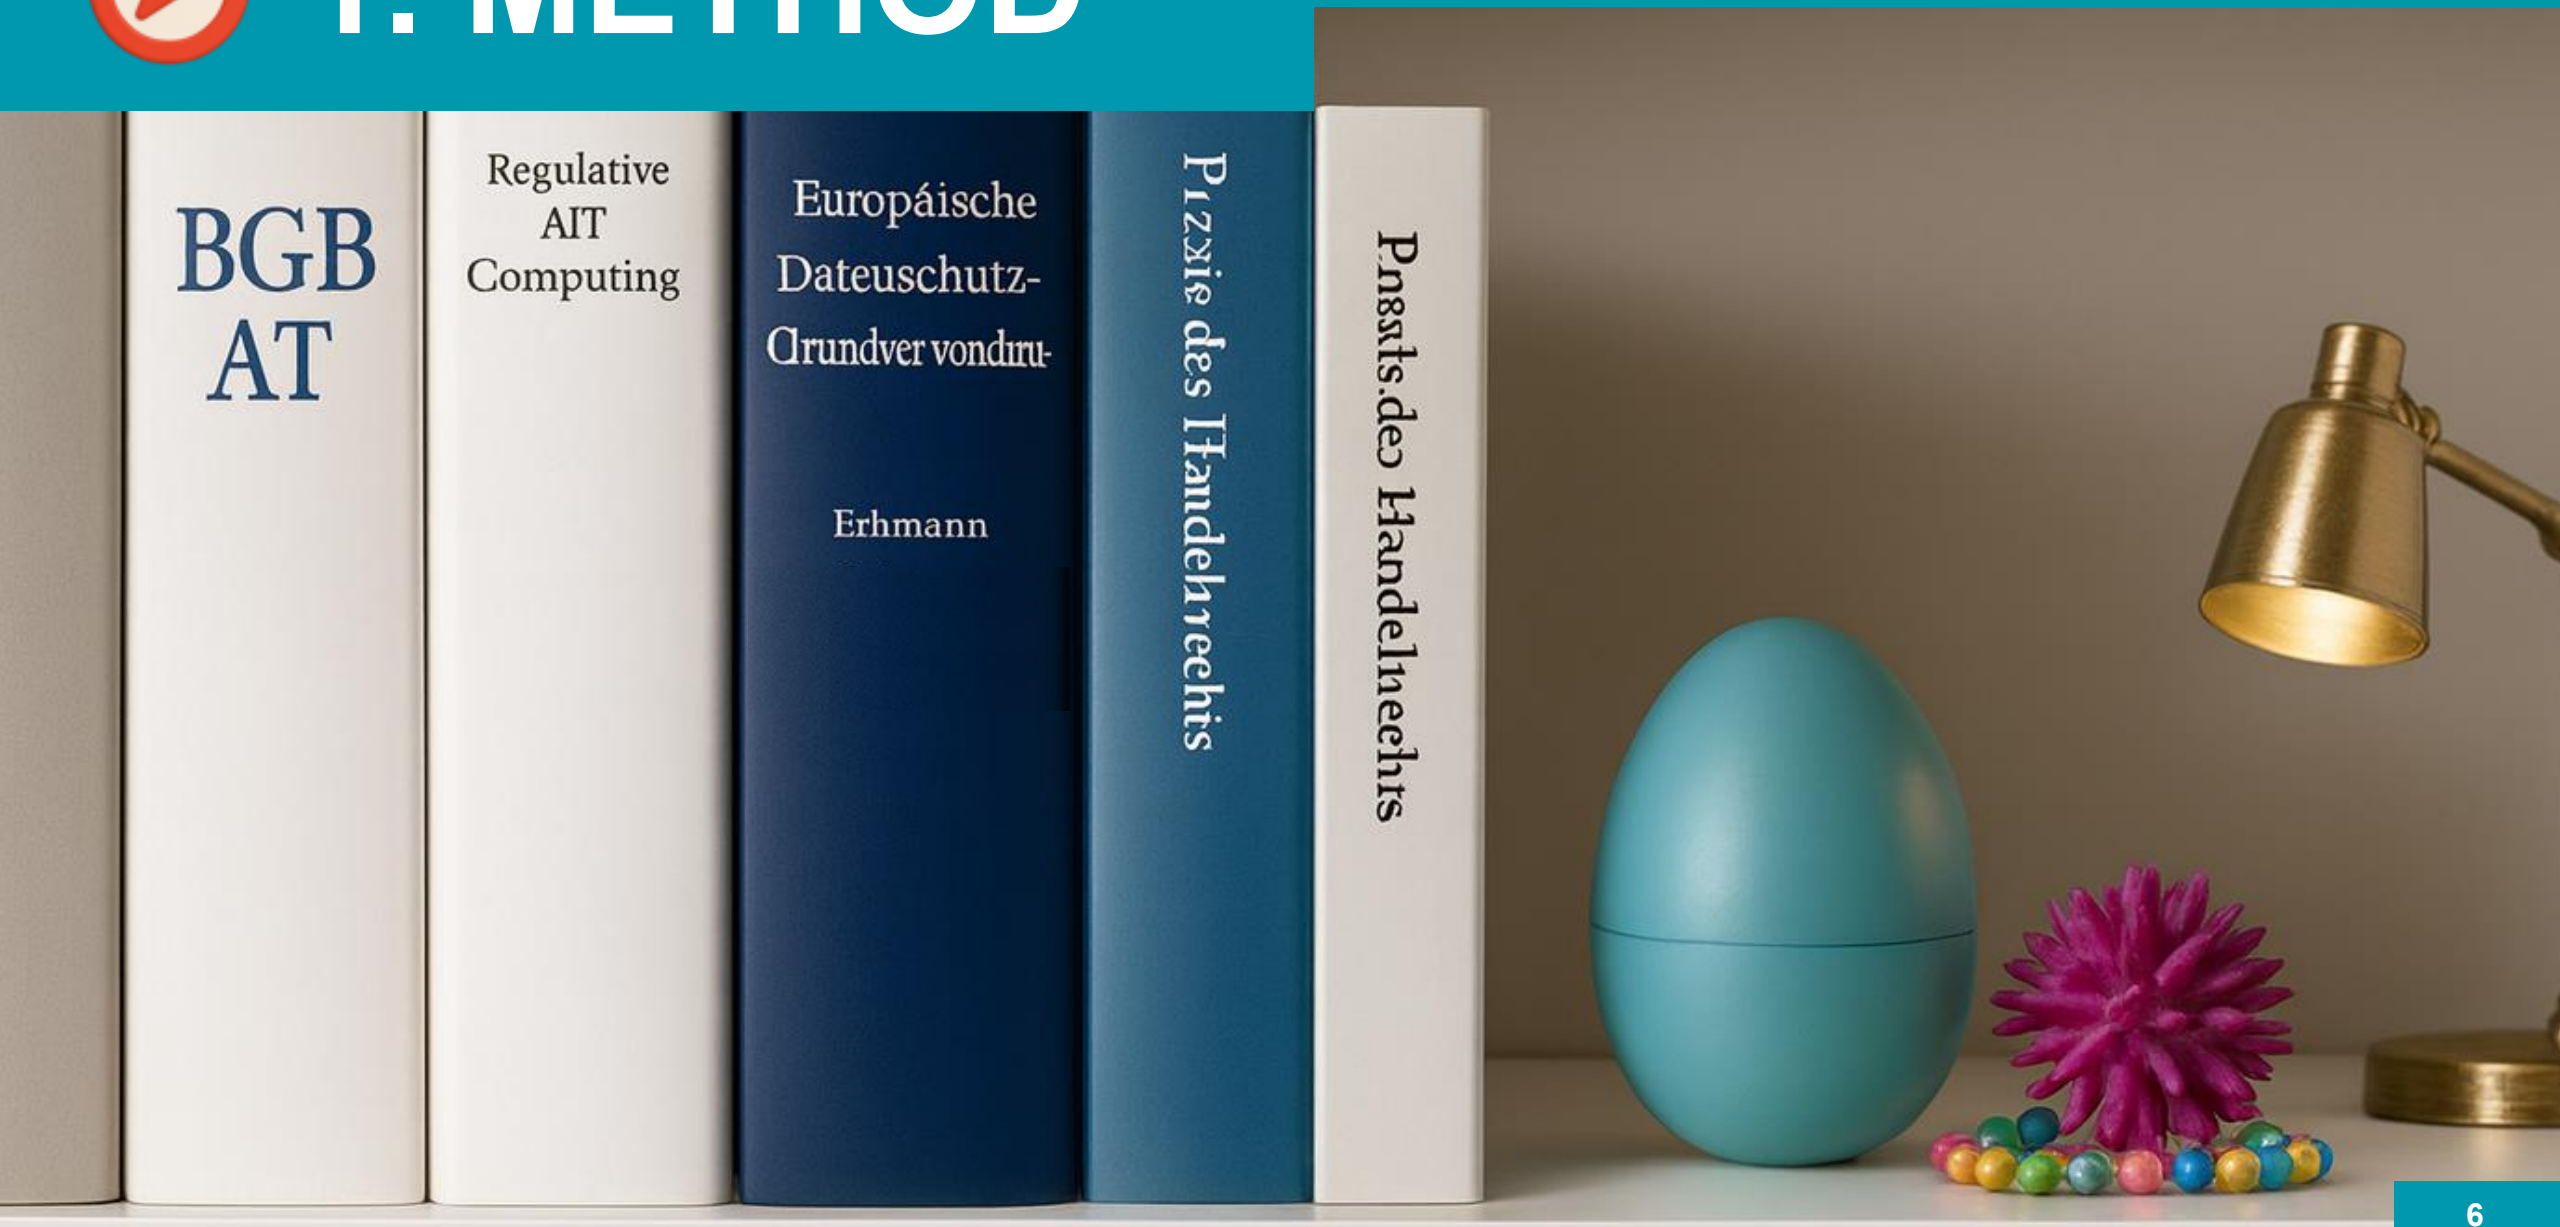
BGB
AT
Regulative
AIT
Computing
Europäische
Dateuschutz-
Erhmann

LEARNING LIKE A MEMORY CHAMPION

Christiane Stenger, multiple world memory champion, uses sophisticated mnemonic techniques to memorize vast amounts of abstract information.

Her method is simple but effective:

- Emotionalize pieces of information
- Connect the content through stories
- Transform, store, and recall

She translates complex information into imaginative images. These images are then woven into unique (and often funny) stories. Through emotional coding and association, content can be understood, retained, and reactivated when needed.

Law tutors have long made use of such techniques to make abstract legal concepts vivid and memorable. For example, the “chest transaction” is a metaphor for self-dealing under § 181 BGB (german civil law). Or the “Nutella theory,” a mnemonic for formal expropriation under Art. 14(3) GG (german basic law): *“Only when it says expropriation on the label, is it really expropriation inside.”*

Learn more about Christiane Stenger: Instagram link

The list could go on endlessly

Against this backdrop, the Playbook reimagines an AI system as a mechanical egg, GPAI models as colorful corals, data as strings of pearls, and prohibited AI practices as a T-Rex.

The Playbook adopts this technique of emotionalization and playful association to make abstract information tangible – through symbols, characters, and short yet serious stories.

Learning this way goes far beyond memorizing words: complex concepts and relationships can be recalled internally within seconds – whether in an exam or in practice.

Techniques like those of Christiane Stenger are highly effective. In this Playbook, they are combined with another method – one especially familiar in the digital economy:

The *serious play* with the products of our childhood…

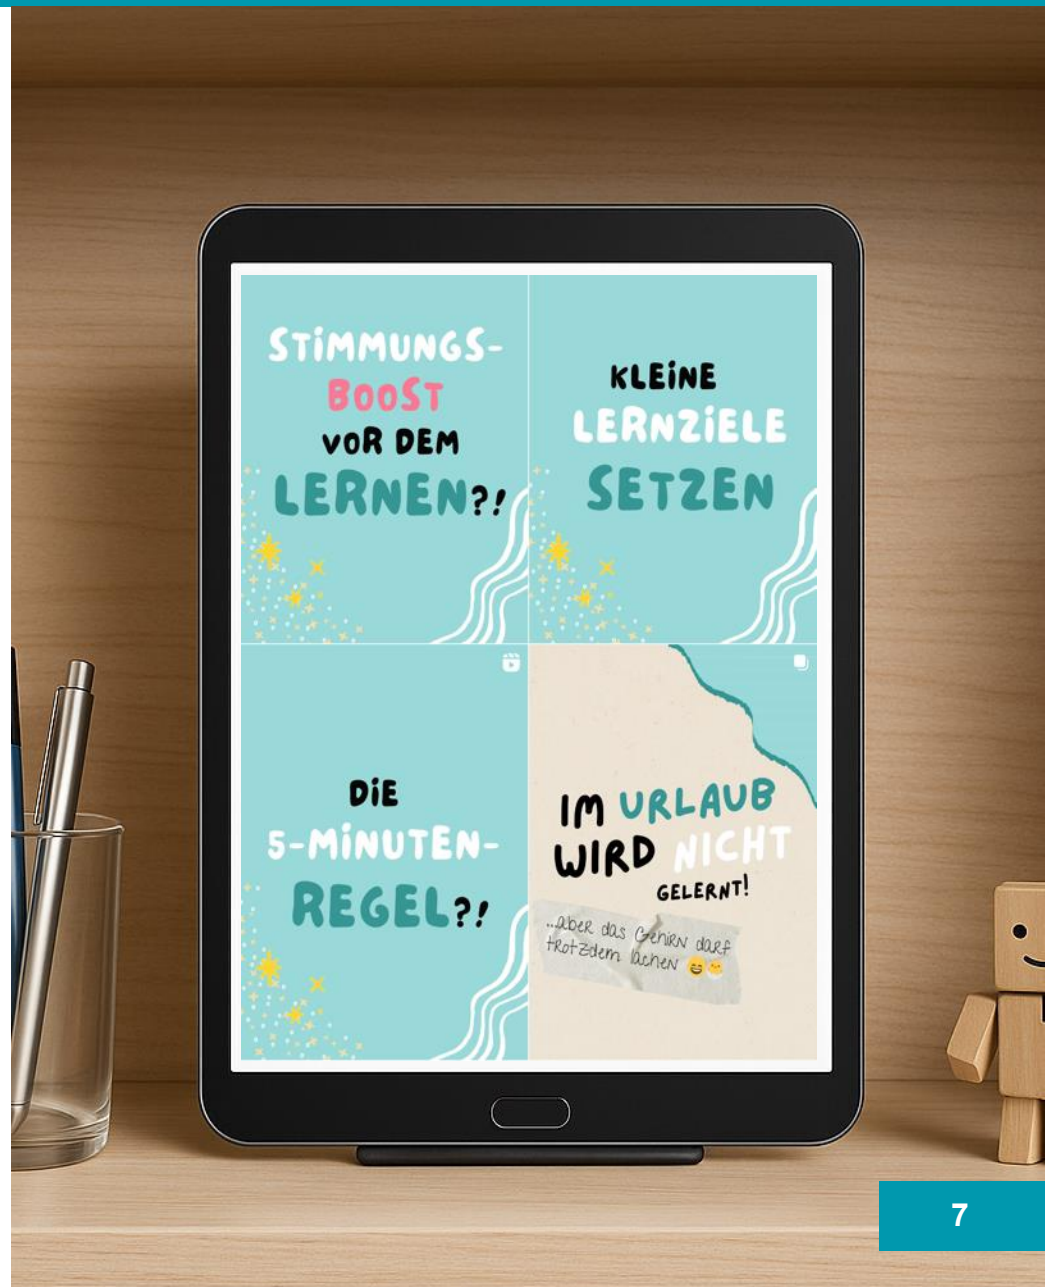

SERIOUS PLAY – MORE THAN COLORFUL FIGURES

LEGO made it famous: so-called *Serious Play*. Highly skilled experts from corporations, organizations, or start-ups gather in a room. With colorful LEGO bricks – or Playmobil Pro figures – they build disruptive solutions, and pay considerable sums for such workshops.

Not without reason: behind the playful surface lies a sophisticated format for solving complex problems. Especially those that cannot be addressed with slides, spreadsheets, or legal provisions alone. In the virtual sphere in particular, serious play shows its strength: physical building blocks break through abstraction, and playfulness lowers the barriers to interaction in teams.

The method enables simple yet vivid concepts. These evolve into stories. And those stories develop into solutions based on interdisciplinary perspectives and expertise.

The same applies to AI law: before regulatory issues can be assessed, it must first be clear what we are talking about. When is a piece of software an AI system or an AI model? When is a service a GPAI model – or a GPAI system?

Only when such similar-sounding terms are clearly distinguished can the appropriate provisions be applied within complex AI value chains.

The Playbook deliberately builds on the method of Serious Play. It uses its own symbolic world – featuring mechanical surprise eggs, corals, pearls, jewelry boxes, little hats, and even pebbles.

Symbols and stories have been proven to make abstract concepts tangible. They create a space in which lawyers, engineers, managers, and users can share a common symbolic language – one that helps them deploy AI successfully and master AI law.

More on LEGO Serious Play: Wikipedia
More on Playmobil Pro: Playmobil Pro

COMBINING PLAY AND LEGAL CONTENT

through play. Whether in print or digital form, language remains the foundation of legal work. Lawyers must know, read, and analyze statutes, commentaries, and case law. The same applies to the EU AI Act.

Play is no substitute for the traditional legal craft. Yet AI itself is abstract and opaque: even IT and technical experts sometimes struggle to define AI clearly or to explain how it works transparently. The AI Act does not make things easier – it sets out, from a legal perspective, what AI is, and what it is not.

This is where the Playbook comes in. It complements the accompanying german legal script *Fundamentals of AI Law*, which explains legal aspects in detail. The Playbook translates these legal contents into vivid images, metaphors, and mnemonics.

AI and AI law thus become *tangible*:

- Understanding through images and language
- Retaining through memory techniques
- Applying through legal text and Playbook

The combination of script and Playbook enables agile legal understanding that goes beyond rote learning: the content remains legally precise, while becoming vivid and memorable.

This is crucial. AI systems are not physical machines; they consist of virtual components: interaction interfaces, AI models, and system data. They receive inputs and generate outputs.

But where does one end and the other begin? The AI Act requires precise differentiation. It contains over 60 definitions. Visualizing the most important ones – and illustrating their interplay in a striking way – is the primary goal of this EU AI Act Playbook.

Once the central definitions of the EU AI Act are firmly grasped, the high quality of the regulation itself becomes apparent. The Act is well-designed, but complex. That is precisely why it should be explored and reinforced through play.

More on the script „Fundamentals of AI Law“ (german only): grundwissen-ki-recht.de

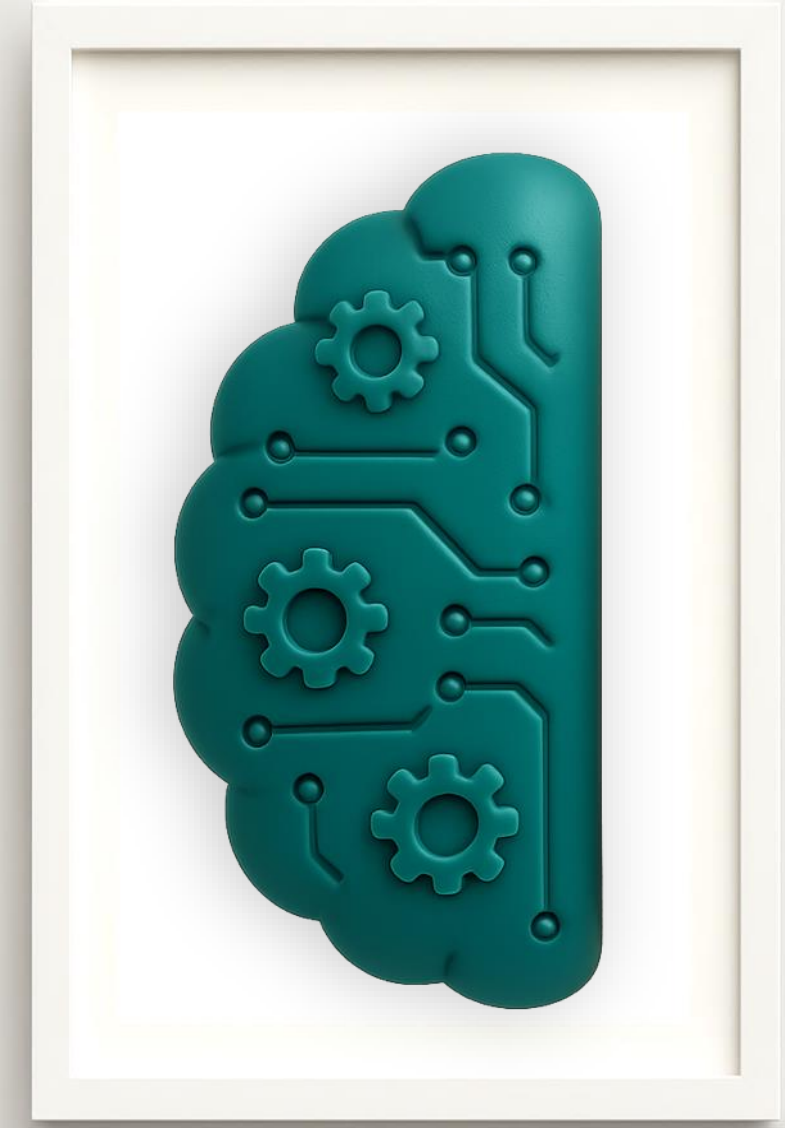

SCRIPT & PLAYBOOK

Put differently**:** the German script *Fundamentals of AI Law* forms the legal foundation. It provides precise definitions, examination steps, overviews, and references to articles of the EU AI Act as well as other laws.

It delivers the legal toolkit that is indispensable for case analysis – both in education and, even more so, in practice. These legal details are processed by the logical *left hemisphere* of the brain.

The Playbook, in turn, translates key aspects of the EU AI Act into concise symbols, characters, and scenes. It adds emotion and concreteness where the script remains abstract. It makes it easier to enter a complex subject matter and creates an overview. In a playful way, it stimulates the creative *right hemisphere* of the brain.

The result is an interplay of play and logic – making it easier to grasp and apply the AI Act, and to design AI that complies with the law.

LET US RECAP STATION ONE – THE METHOD:

1. THIS PLAYBOOK COMPLEMENTS LEGAL LITERATURE – ITS PRIMARY ROLE IS TO SUPPORT AND DEEPEN ITS UNDERSTANDING.
2. CLASSICAL LEGAL METHODS REMAIN INDISPENSABLE – BUT THEY REACH THEIR LIMITS WHEN IT COMES TO AI.
3. MNEMONIC TECHNIQUES HELP MAKE ABSTRACT TERMS SUCH AS *GPAI MODEL* OR *QUASI-PROVIDER* EASIER TO REMEMBER.
4. SERIOUS PLAY HELPS TO GRASP THE COMPLEX INTERPLAY OF AI TYPES, RISKS, AND ACTORS.
5. LEGAL LITERATURE AND THE PLAYBOOK ADDRESS THE BRAIN IN DIFFERENT WAYS.

NOW WE MOVE ON TO STATION TWO: THE AI SYSTEM

2. AI SYSTEM

THE AI SYSTEM: A MECHANICAL EGG

The AI system is the central pivot of the AI Act. A mechanical egg serves as our symbol to identify and understand the many criteria and legal specifics of AI systems – in two ways.

In this and the following chapter we will see:

- What an AI system *is* in positive terms.
- And what it *is not* in negative terms.

Often, negative exclusion is a faster way to clarity than positive definition – especially when uncertainty exists.

Example: imagine you have lost your house key. The first step is a negative check:

Where is the key most likely *not*?

- Not in the apartment – because you left it locked.
- Not in the office – you locked that with the key ring.
- Not in the car – you came on foot … and so on.

By ruling things out, searches – and legal definitions – become more focused and efficient.

Without negative exclusion, defining AI systems could become lengthy and unfocused.

The combination of *negative exclusion* and *positive definition* helps capture the legal essence of AI systems.

And that is why, to illustrate this dual determination under the AI Act, the AI system is represented as a mechanical egg.

In the following section, we will not only highlight the egg's positive features: it will also be visually distinguished from the AI model.

While the names sound similar, the functions are entirely different – and so are the symbols.

Remember: in the sense of the AI Act, the AI system is a mechanical egg.

In short:

AI system = egg

DISTINGUISHING AI SYSTEMS FROM PRODUCTS

AI systems and products are two different things. Mechanical eggs can be embedded in an astonishing variety of products: a pacemaker, an electricity meter, or a toy. Even software is considered a product under the EU's new Product Liability Directive. If the mechanical egg is integrated into a product, we speak of *embedded AI*.

Products are often subject to additional regulations: toys, for example, must be safe for children even without AI. Thus, the manufacturer of a doll that contains an AI system must comply with the rules applicable to toys in general. In addition, the integrated AI system must comply with the EU AI Act. This leads to overlapping rules and responsibilities, which will gradually be harmonized – for critical infrastructure, automobiles, medical products, and toys.

Such overarching regulation makes sense. What matters to children and parents alike is that both are safe: the toy itself, and the AI system integrated into it.

An AI system, however, remains an AI system – whether or not it is integrated into a product.

It is characterized, in all variants, by features that must be observed under the AI Act:

- It is a machine capable of interacting with its environment – whether that environment is a human, a doll, or a piece of software.
- If the magical egg receives information as input, it independently derives its output from data. In that sense, it is often a kind of surprise egg.
- No one knows exactly what output the egg will produce. Often the egg itself does not know. This unpredictability makes AI difficult to manage.
- Some AI systems are capable of learning. They collect data to improve results – but that is not a necessary feature.

Remember: the egg, in the sense of the AI Act, is a machine that interacts with its environment. If it cannot interact – or if it does not act in a machine-like manner – it is not an AI system in the legal sense.

TRUST IN AI: THE PRIMARY GOAL

The previous examples have shown how differently AI systems can be used. Depending on the field of application, opportunities and risks must be assessed in different ways.

Users of AI – and those affected by its use – must always be able to rely on AI systems being trustworthy and used responsibly. For this reason, the AI Act defines different risk classes.

This means, among other things, that AI systems deemed unacceptably dangerous may not be placed on the market or used at all. For example, toys that could manipulate children through AI are prohibited.

But closer examination is essential:

- Some AI systems appear harmless at first glance. The little birds are especially cute – but precisely they might manipulate young children.
- By contrast, the snake instinctively evokes unease in many people. Yet in reality it is often harmless – and even a symbol of medicine.

Appearances can be deceptive. But that is not all: risks arise not only from “malicious” AI, but often from AI that is complex, opaque, or misapplied.

Operation is often where the greatest danger lies: who can blame a helpful bee for stinging when provoked? That is why the training of those who work with AI is one of the most important aspects of the AI Act: AI must – and can – be operated competently.

By fostering AI literacy, users can employ AI correctly and better assess its risks.

Yet this is not unlimited. How an AI system behaves under given circumstances, and what exactly is happening inside it, can scarcely be judged “from the outside.”

This is partly because AI systems contain many components that interact – including other AI systems.

To illustrate this interplay, we now turn to the *matryoshka principle* of AI systems.

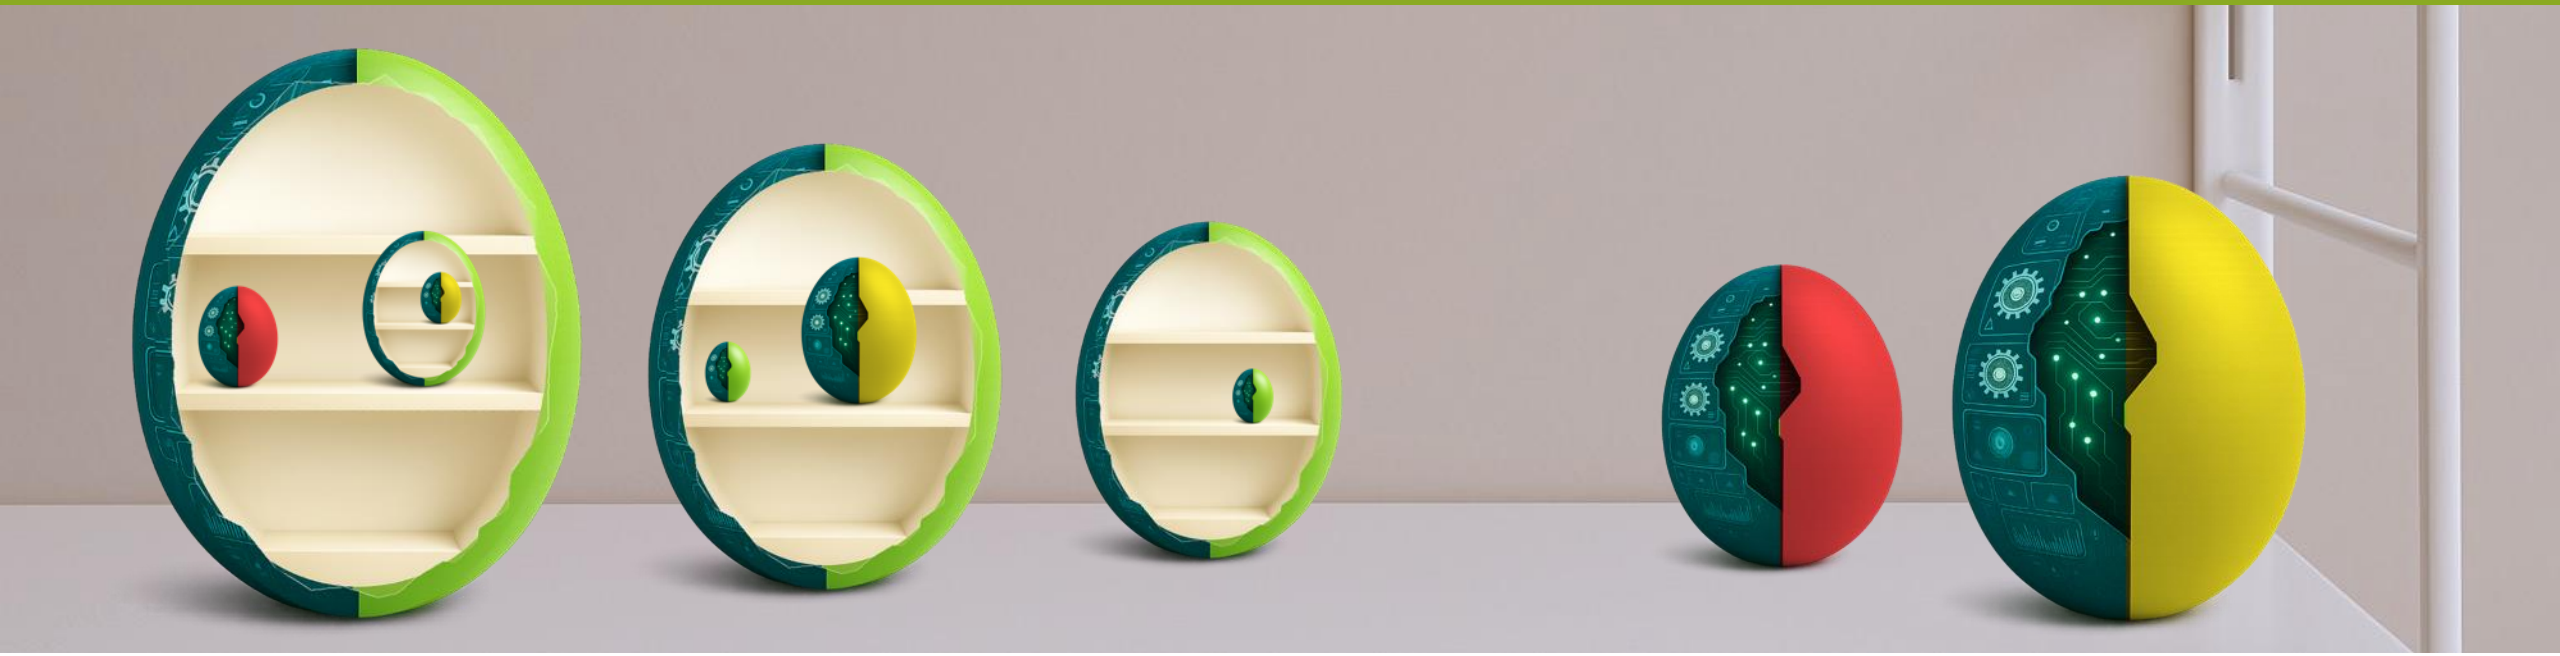

THE MATRYOSHKA PRINCIPLE: EGG IN EGG IN EGG …

To understand the nested architectures of AI systems, it helps to look inside.

There we encounter an interesting phenomenon: AI systems can integrate other AI systems. A mechanical egg may in turn contain many additional eggs – each of them an AI system in its own right. Much like matryoshka dolls, the AI eggs can be nested within one another.

Visually speaking: an egg inside an egg inside an egg.

This makes things demanding, because each AI system within a complex AI value chain normally has to be assessed separately. Taken together, however, they form a new whole – and this whole, with all its integrated components, must comply with the requirements of the EU AI Act. Complex, yes – but real and important.

For example, a medical egg might consist of a language model with a specialized dictionary, an image-recognition AI, a diagnostic AI, and other components. And each of these components may in turn contain further AI systems. The advantage is clear: modular construction allows particularly powerful AI systems to be built.

It is therefore crucial to keep an overview when dealing with a value chain that contains many components.

To make this easier, let us look more closely at the egg itself. It has three compartments designed to simplify the overview of its components.

THREE COMPARTMENTS FOR COMPONENTS

Many AI systems are organized in a division of labor. For this purpose, they contain three compartments:

- The top compartment is for interaction with the environment. The egg must be able to receive information as input and generate output to shape its environment. Examples include a keyboard, a microphone, a joystick, or a sensor.
- The middle compartment contains the intelligence. This resides in the AI model – the “brain” of the system. It stores a large amount of trained data. With this model knowledge, a generative AI system, such as a chatbot, can already answer many questions independently – without needing further data.
- After training ends, the AI model may not always reflect current reality. For instance, it might still “believe” Angela Merkel is chancellor. To remain current, the bottom compartment can perform real-time searches or store specialized system data. This may also include user data or session memory.

Distinguishing between the three compartments is especially important for data protection reasons. It makes a big difference whether user data end up in the second or the third compartment. If stored in the AI model, removal is often difficult and resource-intensive. If stored in a database, in many cases a simple click is enough to delete it.

We will come back to the bottom compartment later. For now, let us look more closely at the middle one. Here we discover two different forms of integrated intelligence:

- On the left: intelligence may come from another integrated AI system (egg within egg).
- On the right: it may stem directly from an AI model. The AI Act differentiates between several types of models that must be distinguished.

The so-called GPAI model (the colorful coral on the middle-right) is particularly important. It is specifically regulated in the EU AI Act. That is the subject of Station Three.

LET US RECAP STATION ONE – THE AI SYSTEM:

1. THE CENTRAL GOAL OF THE AI ACT IS TO BUILD TRUST IN AI SYSTEMS – OF ANY KIND.
2. AN AI SYSTEM MAY BE EMBEDDED IN A PRODUCT OR BE A PRODUCT IN ITSELF. EITHER WAY: IT FALLS UNDER THE EU AI ACT.
3. AI SYSTEMS CAN CONTAIN OTHER AI SYSTEMS – LIKE MATRYOSHKA DOLLS. THIS MAKES ASSESSMENT DEMANDING.
4. AN AI SYSTEM IS A MACHINE. IT REQUIRES INPUT, PROCESSES IT AUTONOMOUSLY, AND GENERATES OUTPUT.
5. THE MECHANICAL EGG HAS THREE COMPARTMENTS: ONE FOR INTERACTION, ONE FOR INTELLIGENCE, AND ONE FOR ADDITIONAL DATA.

NOW WE MOVE ON TO STATION THREE: THE GPAI MODEL

3. GPAI MODEL

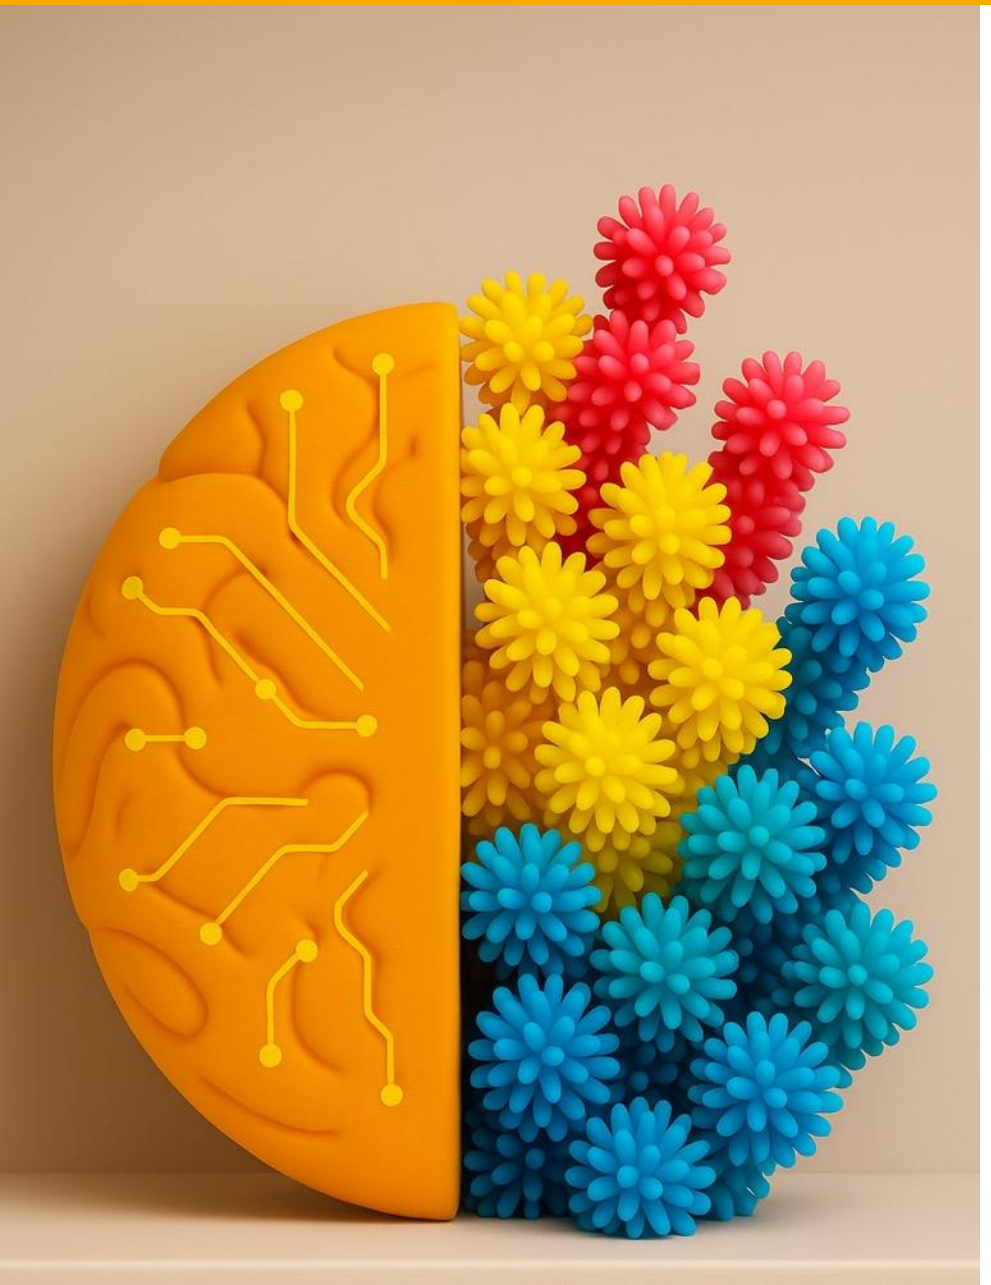

THE GPAI MODEL: CORAL AND CORAL REEF

If the AI system is a mechanical egg, then the GPAI model is a colorful, living coral: a small ecosystem.

The coral is the versatile, intelligent “brain” of generative AI systems. It is what thinks and produces content – texts, images, audio files, or deceptively realistic videos.

GPAI” stands for General Purpose AI. In this context, this refers in particular to a “general-purpose AI model.” This particularly powerful type of AI model is the only one explicitly defined and regulated in the AI Act.

By contrast, the smaller, more specialized AI model is not defined in the Act. In comparison to the coral, it is a polyp: the tiny individual organism from which corals – and ultimately even an entire coral reef – are built.

This makes clear how polyp and coral interact: over time, a polyp can grow larger, more branched, more autonomous, more versatile. At that point, the polyp becomes a coral. And the coral itself can continue to grow – into a coral reef.

A coral reef is beautiful, but not without danger. That is why the EU AI Act devotes specific provisions to it: requirements for *GPAI models with systemic risks*.

Let us sum up:

1.Simple AI model = polyp

2.Versatile GPAI model = coral

3.GPAI model with systemic risks = coral reef

All three are not AI systems. They are like a brain without a body or senses. They lack the top compartment of the mechanical egg. Their place is in the middle compartment – where they provide intelligence to the egg.

One more time: polyps are not defined and scarcely regulated by the AI Act. Corals and coral reefs, by contrast, are subject to specific requirements.

So let us remember the symbol for generative intelligence:

GPAI Model = Coral

CORAL IN A COLORFUL EGG

The coral was added to the EU AI Act rather late – after the “aha moment” when ChatGPT was released. Many provisions of the AI Act are mainly, or even exclusively, linked to the GPAI model. Not without reason!

It is the engine of the GPAI system, which is a particularly powerful kind of AI system:

- Its brain is not a small, specialized polyp, but the large, versatile coral.
- The egg of the GPAI system is therefore as colorful as the coral – like a rainbow.
- The many colors represent the versatility of the GPAI system: it can perform a wide range of tasks with great competence.

We will soon see how useful it is to distinguish between the different egg types and the polyps, corals, or coral reefs inside them.

For now, it is important to know that there are different AI systems (a standard AI system and a GPAI system) and that there are several types of models (a specialized AI model, a GPAI model, and a GPAI model with systemic risks).

Great – you are learning fast!

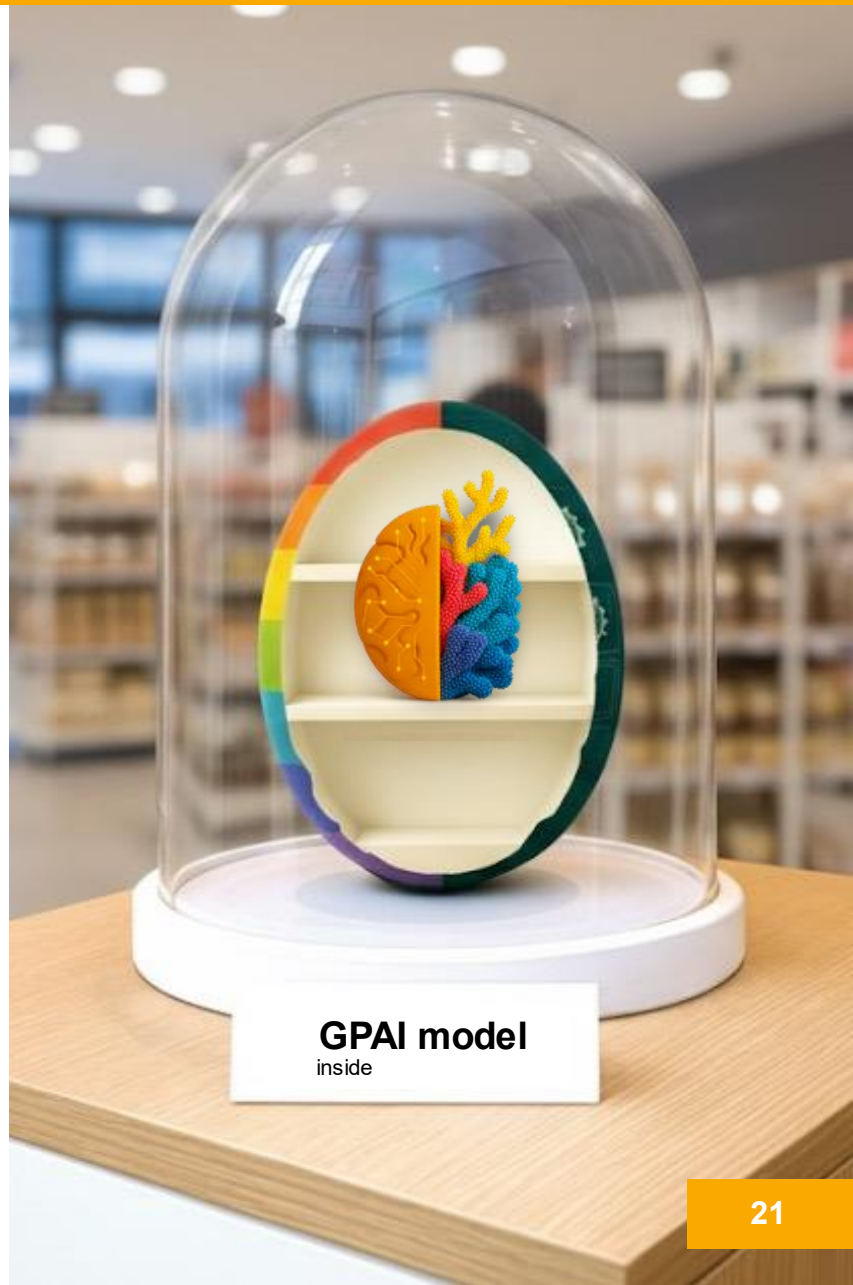

THE AI MODEL FAMILY

The EU AI Act thus recognizes several AI models, and it also defines different AI systems. A toy, for example, usually uses a simple AI system. It runs on batteries and can perform a few, very specific functions – such as making a doll speak or changing its eye color. Inside its magical egg, a mini AI model is at work.

A factory robot is similar, but also different. With a specialized AI system that includes a self-learning AI model, it can accomplish astonishing things. The EU AI Act does not regulate the polyps – that is, the specialized AI models. There are simply too many possible AI techniques at the model level. The Act seeks to remain technologically neutral, since innovation in AI models advances rapidly. What it regulates is the AI system – the mechanical egg – into which a polyp has been integrated. That is sufficient.

Now, if a GPAI model (the coral) is integrated into an AI system, then almost automatically a GPAI system emerges. The most common case – as shown in the illustration on the right – is an AI chatbot. It can be incredibly versatile: generating text, composing music, and interacting with people. The coral alone cannot do this, because it lacks the top compartment for interaction. Nevertheless, the coral itself is regulated, and those who provide it have transparency obligations. More on that shortly.

Now let us discover how the coral turns the GPAI system into a true all-rounder!

VERSATILE AND COMPETENT

It is the coral that makes the surrounding GPAI system a true all-rounder:

- It can draw on its world knowledge to answer all kinds of questions.
- It can recommend cooking recipes and give personalized beauty tips.
- It can describe foreign countries or suggest vacation trips.
- Often it can even create images – and in some cases sounds and videos – with striking perfection.

All this is made possible by the many capabilities of the colorful coral! Its training data allow for connections across modalities – for example, describing an image after it is uploaded, or modifying it on the basis of text prompts.

These multimodal GPAI models, capable of handling images, language, and much more, create enormous opportunities:

- For individuals
- For companies
- For organizations

The true all-rounder, then, is the GPAI model. Yet without the surrounding GPAI system, the coral remains idle: it receives no inputs and generates no outputs. It is like an engine without a chassis. And the reverse is equally true: without the GPAI model, the surrounding GPAI system is nothing more than ordinary software – software not subject to the AI Act, because a magical egg without a coral, a polyp, or a vast coral reef in its middle compartment is not truly intelligent.

Speaking of coral reefs: let us take a closer look at this type of GPAI model. It is an especially large AI model with general-purpose capabilities, but it also carries systemic risks. These risks arise from factors such as very high performance or particularly wide deployment.

Coral reefs are maximally versatile. That is why they are often highly prized. But this very versatility leads to unique risks. To understand them, we must explore the ecosystem of the coral reef.

A fascinating world awaits!

BEAUTIFUL YET RISKY

It seems obvious: the larger and more diverse the corals grow, and the more they develop into a vast, beautiful coral reef, the greater their attraction becomes. More and more fish gather around.

But therein also lies a danger:

- What began as a colorful habitat becomes a highly networked system in which a single error can have far-reaching consequences.
- A vast reef brings not only diversity but also vulnerability – for example, the risk of coral bleaching.

If something tips, it can suddenly affect everything – the entire ecosystem. This is the nature of systemic risks. They are invisible, complex, and spread across the entire value chain.

That is why particularly large and widely deployed GPAI models are subject to special safety requirements in the EU. This makes life more demanding for providers, but it makes society as a whole safer. And that is a good thing.

So let us remember: GPAI models with systemic risks are subject to additional obligations. These safeguards ensure that coral bleaching does not destroy entire ecosystems.

But this raises the question: how does a vast coral reef actually fit inside a mechanical egg?

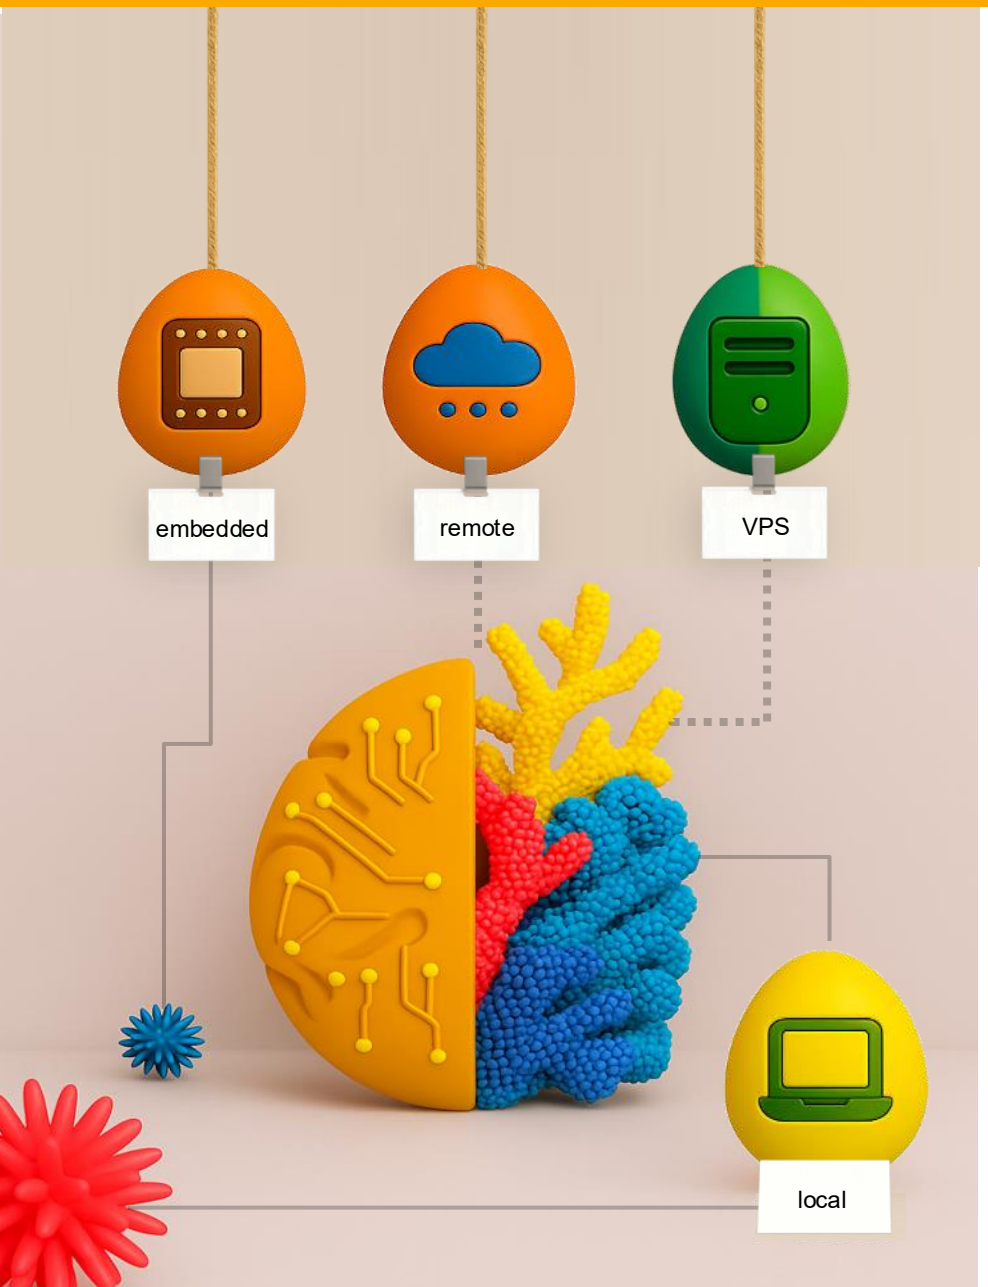

INTEGRATION OF AI MODELS INTO AI SYSTEMS

AI models can be connected to AI systems in different ways. The type of integration determines, among other things, how flexible, powerful, or controllable the overall system is.

At least the following four variants should be known:

a) Remote (via a public cloud)

- The model runs on a remote server (e.g., OpenAI, Google, AWS, Azure).
- The system sends requests via an interface (API with API key).
- Advantage: high computing power, up-to-date models, excellent infrastructure.
- Disadvantage: dependency and data protection concerns.

b) VPS (Virtual Private Server)

- The model runs on a private but remote server.
- Access is more controlled than with public cloud services.
- Advantage: good control and scalability.
- Disadvantage: maintenance effort and the need for strong AI expertise.

c) Local (on a local computer or server)

- The model operates on the device or network of the AI system itself.
- No internet connection is required.
- Advantage: strong data protection, low external dependency.
- Disadvantage: limited computing power, complex updates.

d) Embedded (integrated into a device or chip)

- The model is permanently embedded in the AI system, e.g., in a robot or IoT device.
- Advantage: real-time capability, extremely low response time, energy efficiency.
- Disadvantage: low flexibility, updates often impossible or costly.

Although the AI Act does not directly regulate integration, the way a model is connected can, among other things, determine who qualifies as a provider. More on this in Station Eight.

So: AI models can be integrated into an AI or GPAI system in different ways!

But now, a few words about typical pitfalls – especially those of corals.

THE PITFALLS OF CORALS

Now for an especially important point: the coral has much in common with the human brain. First, its folds look similar. And like a real brain, the coral can:

1. think logically, and

2. create imaginatively.

Visually speaking, corals have two hemispheres. These can complement each other, but they can also come into conflict. When that happens, surprising outputs and answers emerge – commonly called *hallucinations*.

The GPAI model then invents things that are objectively untrue but sound plausible. For example, it may describe a historical event in great detail – yet give the wrong date.

And if the coral, despite clear indications, stubbornly insists on its claim and justifies it with far-fetched arguments, we encounter the *Muenchhausen effect*: the model continues to "make things up."

This matters greatly, because anyone who integrates a coral into their AI system and places it on the market also bears responsibility if the coral makes mistakes. These errors are attributed to the mechanical egg. The same applies to copyright violations: the GPAI model cannot itself deliver works. It lacks the top compartment with the instruments of interaction.

For this reason, and many others, anyone who integrates a coral into an AI system – thereby creating a GPAI system – depends on key information, such as:

- How was the coral trained?
- With which data and methods?
- With which capabilities (image/text)?

To guarantee this information, the EU AI Act sets out rights and obligations that are tied to specific actors.

Against this backdrop, let us first look at the actors, and then at the risk classes and the obligations arising from the AI Act.

LET US RECAP STATION ONE – THE GPAI MODEL:

1. THE BRAIN OF AN AI SYSTEM IS THE AI MODEL. THE EU AI ACT DISTINGUISHES THREE TYPES: POLYPS, CORALS, AND CORAL REEFS.
2. SMALL, SPECIALIZED AI MODELS ARE POLYPS. LARGE, VERSATILE GPAI MODELS ARE COLORFUL CORALS OR CORAL REEFS.
3. THE CORAL REEF CARRIES SYSTEMIC RISKS: IT MUST THEREFORE BE ESPECIALLY WELL PROTECTED AGAINST CORAL BLEACHING.
4. AI MODELS CAN BE INTEGRATED INTO AI SYSTEMS IN DIFFERENT WAYS: REMOTE, VPS, LOCAL, OR EMBEDDED.
5. CORALS CAN HALLUCINATE AND MAKE MISTAKES. THAT IS WHY TRANSPARENCY ABOUT HOW THEY WORK IS ESSENTIAL.

NOW WE MOVE ON TO STATION FOUR: THE ACTORS

4. ACTORS

ACTORS WEARING MULTIPLE HATS

Now to the two most important actors under the AI Act: the *provider* and the *deployer*.

The AI Act also names many other actors – for example, the product manufacturer, importer, distributor, or authorized representative. But these are less frequent and more specific.

Even with just providers and deployers, things are demanding enough! For the roles under the AI Act are not static. They can change: one actor may wear the provider hat and the deployer hat for an AI system – and at the same time wear another hat as a downstream provider of a GPAI model.

What better metaphor, then, than headgear to distinguish the key actors?

We will see that hats can be both a useful and entertaining element to tell actors apart. In this sense, all actors defined in the AI Act wearing multiple hats. Whoever wears a hat carries not only responsibility, but also, in some cases, rights.

Before we can start playing with colors, however, we must first learn the rules. And that means: we must know not only the roles themselves, but also their relation to AI eggs, AI corals, and coral reefs.

For example, the term *provider* under the AI Act appears in four different variants:

- As provider of an AI or GPAI system.
- As quasi-provider or secondary provider of an AI or GPAI system.
- As provider of a GPAI model.
- As downstream provider of a GPAI or other AI model.

Phew – that makes communication really complicated! You cannot simply say "provider." You must also add *of what* someone is a provider – and at *what point in time*.

And this is where the hats help us bring order and clarity.

Let us note: actors under the EU AI Act wear colored hats!

Actors = multiple hats

Example: OpenAI

AN ACTOR WITH A TRIPLE AI TIARA

We now know that actors under the AI Act wear little hats – and that we must distinguish between several roles.

For this reason, it is possible for a single actor to wear several hats at once. The image on the left shows an AI company such as Google, OpenAI, Mistral, DeepSeek, or Anthropic:

These actors offer both:

- GPAI systems, and
- various GPAI models

The GPAI systems are the chatbots with names like Gemini, ChatGPT, or Claude. We recognize them as the colorful egg on the right side of the picture.

And the GPAI models? At OpenAI, for example, they are called GPT-5, GPT-4o, GPT-4-mini, and so on. The figure holds them in her hand as a coral on the left side of the picture.

Now to the triple tiara:

- The yellow hat symbolizes the role as provider of a GPAI model.
- The orange hat symbolizes the role as provider of a GPAI system.
- The red hat signals the additional role as deployer of the GPAI system.

Strictly speaking, companies like OpenAI would not only have to wear one yellow hat, but one for each GPAI model – for every coral or coral reef. This matters because each coral has different capabilities.

For now, let us note that an actor may have just a single role – then only one hat, e.g., a pink or a red hat as deployer.

But an actor may also wear many hats – like here: as provider of a GPAI model (yellow), provider of a GPAI system (orange), and deployer of the GPAI system (red).

Principle clear? Then let's move on!

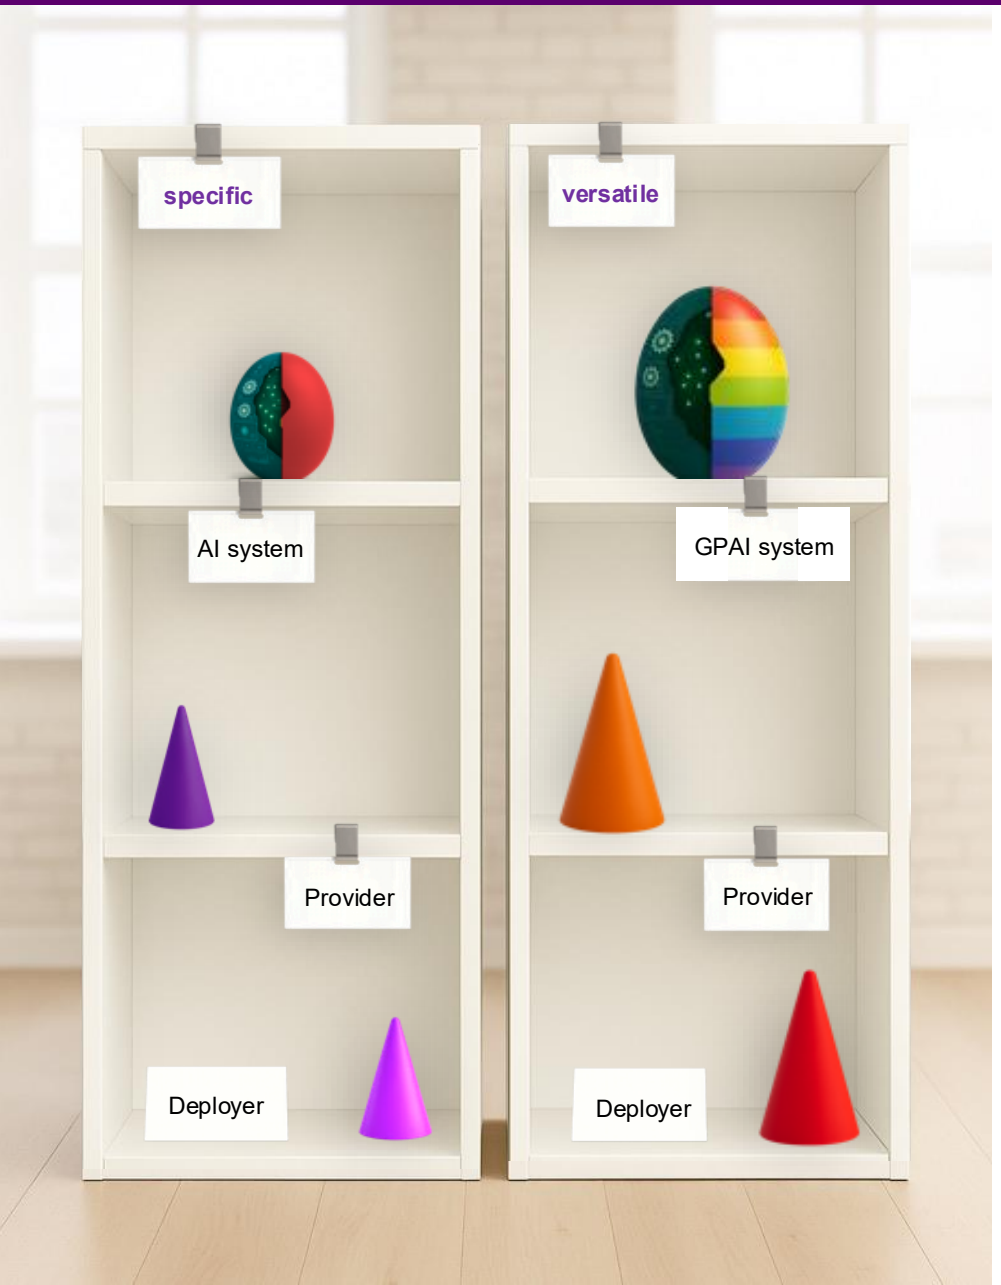

DEPLOYERS EXIST ONLY FOR SYSTEMS

Let us begin with the actors for the two system types:

- On the left, the “normal,” often specialized AI system – symbolized by a smaller, single-colored egg.
- On the right, the versatile, colorful egg – shown here slightly larger – as the symbol for GPAI systems.

For both variants there are roles as *provider* and as *deployer*. This is an important point, because on the next page we will learn that there are also providers for GPAI models.

Deployers, however, exist only for AI and GPAI systems – never for models.

Now to the hats for providers and deployers: we find them in the middle and bottom compartments. The provider of a simple AI system wears a purple hat, and the deployer wears a pink hat. Both hats are somewhat smaller than those for GPAI systems – because they are specialized.

For corals, the provider hat is orange and the deployer hat is red. So we know: whoever wears a red hat is the deployer of a GPAI system, and whoever wears a purple hat is the provider of a specialized AI system.

Anyone already familiar with the AI Act might ask why the two AI system types use different colors for the roles of provider and deployer. The Act itself does not prescribe this. But in practice, certain obligations apply almost exclusively to providers and deployers of GPAI systems – and rarely to those of simple AI systems. That is why we deliberately distinguish them with colors.

Once more, to be clear:

- If a figure wears an orange hat and a red hat, it is both provider and deployer of a GPAI system.
- That means this actor also carries double obligations – as provider and deployer. More on this in Station Six.

Now let us turn to the model types – for them there are two further colored hats: yellow and blue!

(DOWNSTREAM) CORAL PROVIDERS

The AI Act stipulates that for AI models there are only providers of corals or coral reefs. For small, specialized polyps, by contrast, the Act does not define a provider role.

That is why the hat is missing on the left side under the polyp: in principle, this hat does not exist. But only in principle… It is a little complicated, because the missing hat simply means that the provider of an AI system does not use an external AI model.

Still, there can be a downstream provider of the AI model. This happens when a simple AI system uses an AI model (a polyp) that was created by someone else. In that case, the provider of that AI model becomes the *downstream provider* of the polyp integrated into the AI system.

The same applies to GPAI models. Consider a company that deploys its own chatbot, but relies on a GPAI model from Gemini, OpenAI, or Mistral:

- Then OpenAI, Google, and Mistral are the providers of the GPAI model. They wear the yellow hat.
- The company that uses these models is then the downstream provider of the GPAI model – and wears the blue hat.

This distinction matters, because the downstream provider of GPAI models has a right to extensive information about the GPAI model: how it was trained, with which data, etc.

So, the blue hat means: as a provider of an AI or GPAI system, I do not use my own AI model, but an external one.

If you integrate a GPAI model, you hold claims against its provider – who, for GPAI models, wears the yellow hat.

To show this and the previous page together in context, let us now take a look at an AI value chain. That will make everything clearer.

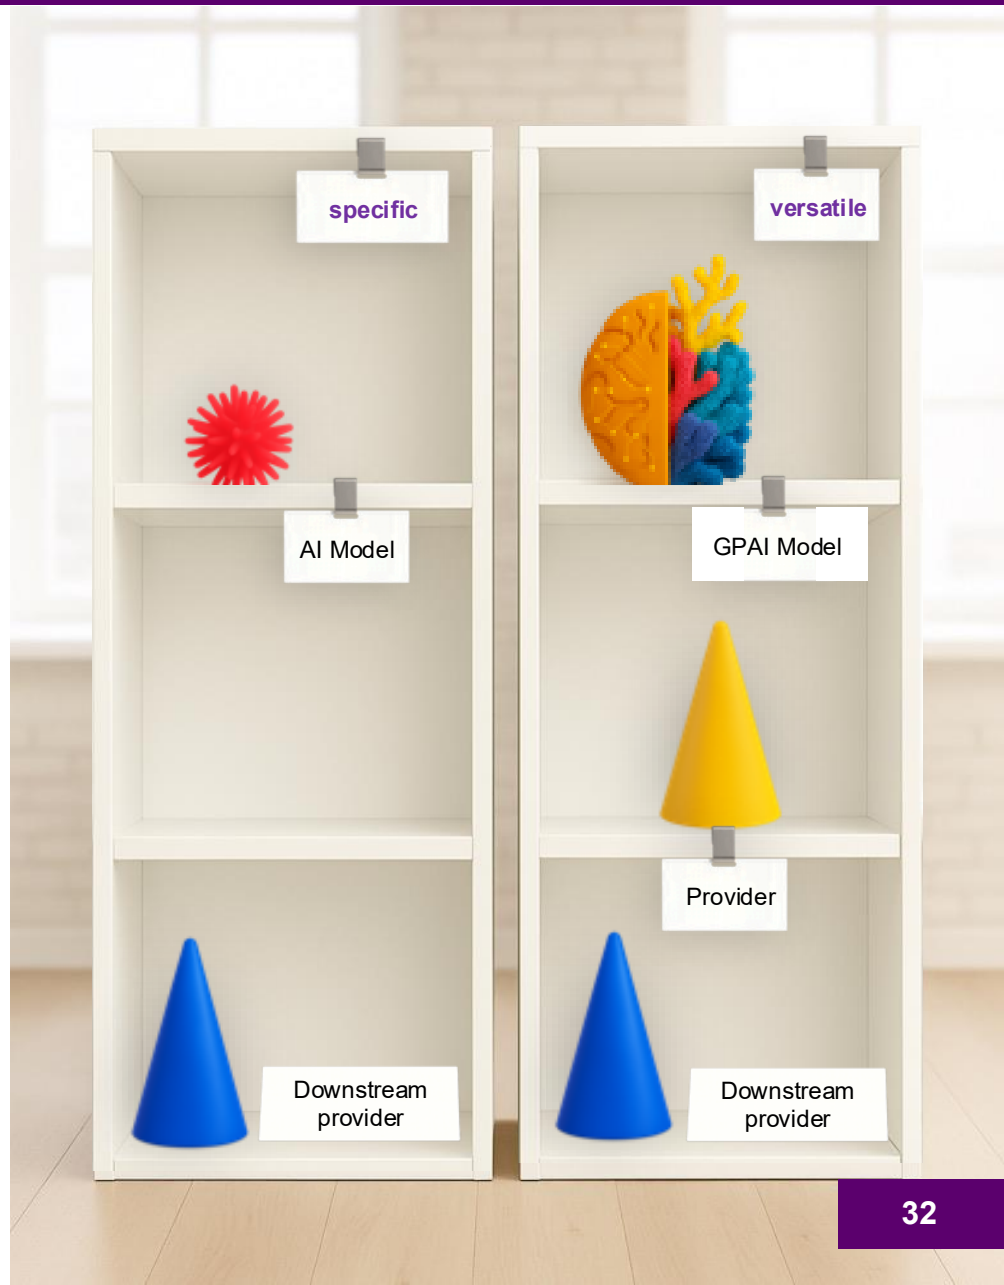

THE GPAI HAT CHAIN – FROM MODEL PROVIDER TO USER

Figure 1 is the provider of a GPAI model – for example, Google, OpenAI, or Mistral. Hence the yellow hat for the coral. Figure 2 is the provider of the colorful GPAI egg and at the same time the downstream provider of the GPAI coral.

This figure therefore wears two hats: blue & orange. Figure 3 is the deployer of the GPAI system. Accordingly, it wears the red hat.The many users of the system (initially) wear no hats at all. They have no legally defined role.

The different hat colors not only help us recognize the roles of various actors. They also illustrate *placing on the market* and *putting into service*. Why?

That is what we will learn on the next page.

PLACEMENT ON THE MARKET, PUTTING INTO SERVICE, AND RESPONSIBLE USE

Let us stay with the previous image for a moment. *Placing on the market* can already occur when a provider of a coral, a GPAI system, or another AI system hands it over for distribution and sale.

Putting into service can already occur when an AI or GPAI system is used internally within an intranet. Through placement on the market, actors become providers; through autonomous use, actors additionally become deployers.

The many colors of the hats not only help us recognize the roles of different actors. They also illustrate the global spread of these roles in the "world atlas" of AI actors. And that world atlas is what we will now explore.

THE WORLD ATLAS OF AI ACTORS

The AI Act applies only within the EU. It concerns actors who offer or deploy AI technology in the EU. But there are also actors from outside the EU who create AI systems and place them on the EU market or put them into service there.

The world map of AI actors illustrates several things at once:

- Companies from outside the EU must comply with the rules and safety standards just like European companies …
- … if they place AI on the European market or if its results are used there.

And this leads us to an important message of the actor world atlas: it shows that most providers and deployers in the EU are linked to specialized AI systems – that is, systems that are not GPAI.

The major providers of GPAI models, as well as the providers and deployers of globally used GPAI systems, are almost all based in the United States and China.

As a result, there are hardly any yellow hats (GPAI model providers) in the EU.

But there are many pink hats (deployers of specialized AI systems). You can find them across nearly all industries. Added to these are the many purple hats – providers of specialized AI systems from the EU.

Orange and red hats are also common, because many companies and public authorities have developed their own chatbots, which they both provide and deploy.

But: most users still rely on GPAI models and systems from the US and China. And precisely for that reason, their widely used services are subject to special criteria: because of their scale, they count as coral reefs – and carry the risk of coral bleaching. This is what the EU seeks to prevent.

So, at the end of this station, things get truly exciting: when actors change their hats – or acquire new ones.

Now for a hat trick!

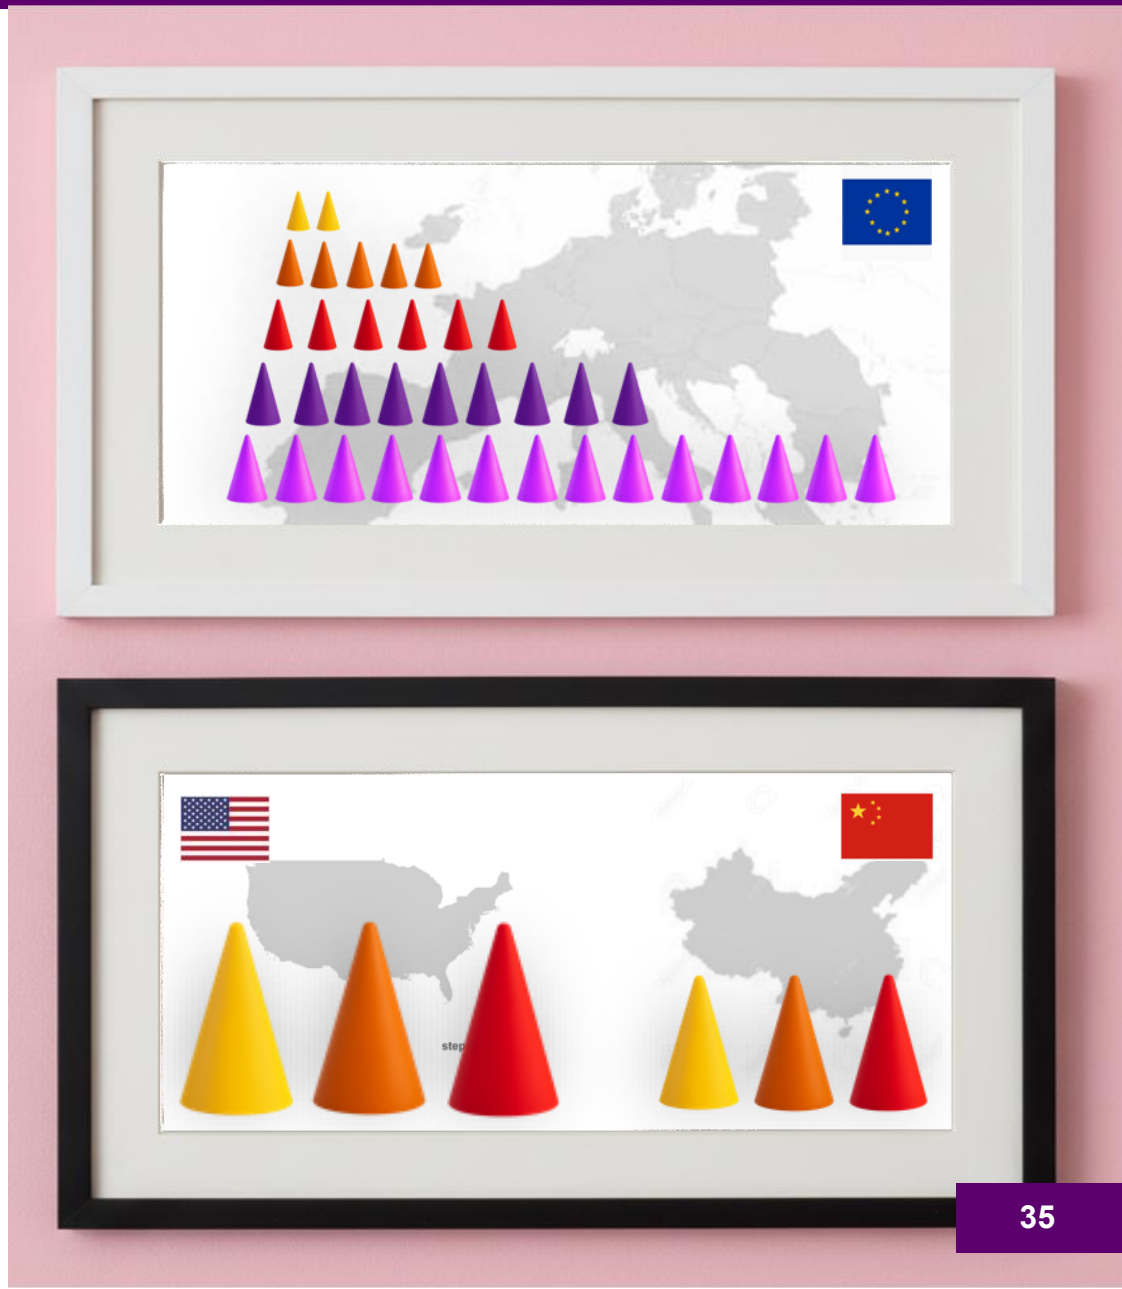

SUBSEQUENT PROVIDERSHIP

Let us take a closer look at the actor on the left:

- At the bottom, it wears a pink hat – so primarily it is the deployer of a specific AI system, shown here as a green egg.
- Above that, however, the figure also wears a purple hat – so in addition, it is also the provider of the AI system.

Hmm. Something seems odd here, because normally the order is reversed:

- First comes the provider's purple hat …
- … and only afterwards the deployer's hat.

After all, before an AI system can be put into service, it must first be placed on the market, right?

Correct. The usual case is: provider hat first, then deployer hat. The same applies to GPAI systems (orange & red).

Here, however, the provider's purple hat was added later – at a point in time when the figure was already a deployer.

How is that possible? Quite simply: this is what is known as a quasi-provider or secondary provider. Evidence for this is given by the screwdriver and the matching logo on both the figure and the egg:

- The screwdriver shows that the deployer has tinkered with the egg, making substantial modifications.
- The logo demonstrates that the deployer gives the outward impression of being the provider – for example, by running the system under its own URL and name on the internet or intranet.

Both lead to quasi-providership:

- Substantial modification of an AI system, and
- Use under one's own name or brand.

This is especially important in cases of high-risk AI. And with that, we move on to Station Five: the risk classes.

LET US RECAP STATION FOUR – THE ACTORS:

1. THE ACTORS UNDER THE EU AI ACT HAVE ROLES. EACH ROLE IS REPRESENTED BY A COLORED HAT.
2. AN ACTOR MAY WEAR SEVERAL HATS – FOR EXAMPLE, AS PROVIDER (PURPLE) AND DEPLOYER (PINK) OF A SPECIFIC AI SYSTEM.
3. THE HATS ILLUSTRATE DIFFERENT ROLES ACROSS THE AI VALUE CHAIN – FROM PROVIDER TO USER.
4. IN THE EU, YELLOW HATS FOR GPAI MODEL PROVIDERS ARE RARE. PINK HATS FOR DEPLOYERS, HOWEVER, ARE VERY COMMON.
5. QUASI-PROVIDERS ARE DEPLOYERS WHO LATER BECOME PROVIDERS – SO THEY WEAR TWO HATS.

NOW WE MOVE ON TO STATION FIFE: THE RISK CLASSES

5. RISK CLASSES

THE RISK CLASSES: SERIOUS BUSINESS – WITH ANIMALS!

AI offers incredible opportunities. But it also carries risks. The main goal of the AI Act is to identify these risks and mitigate them.

So far, so clear. But what do the parrots and that strange green creature on the right have to do with AI risks? And what about the little bird, the turtle, and the T-Rex we saw on the previous page?

What's their connection to the two eggs in the nest?

The answer is simple: they symbolize different risks that may hatch from machine eggs. It's all about the output.

From an egg, you might get:

- a bird (including a parrot),
- a turtle,
- a snake,
- a crocodile,
- … or even a dinosaur.

The T-Rex – one of the most dangerous land animals of all time – is extinct. That's precisely why it represents those risks that must be banned at all costs: risks that must themselves go extinct.

The idea is clear: eggs can produce very different outputs – some harmless, others dangerous.

Some so dangerous that they must be prohibited altogether.

And so we arrive at the bigger picture:

- There are prohibited AI practices.
- There are also risk classes for high-risk AI, as well as for limited-risk and minimal-risk AI applications.
- These risk classes in the EU AI Act apply to AI systems – the machine eggs – not to the models inside them!

All the animals listed above hatch from eggs: they stand for the different levels of risk.

Key takeaway: Risk class = the animal that hatches from the egg.

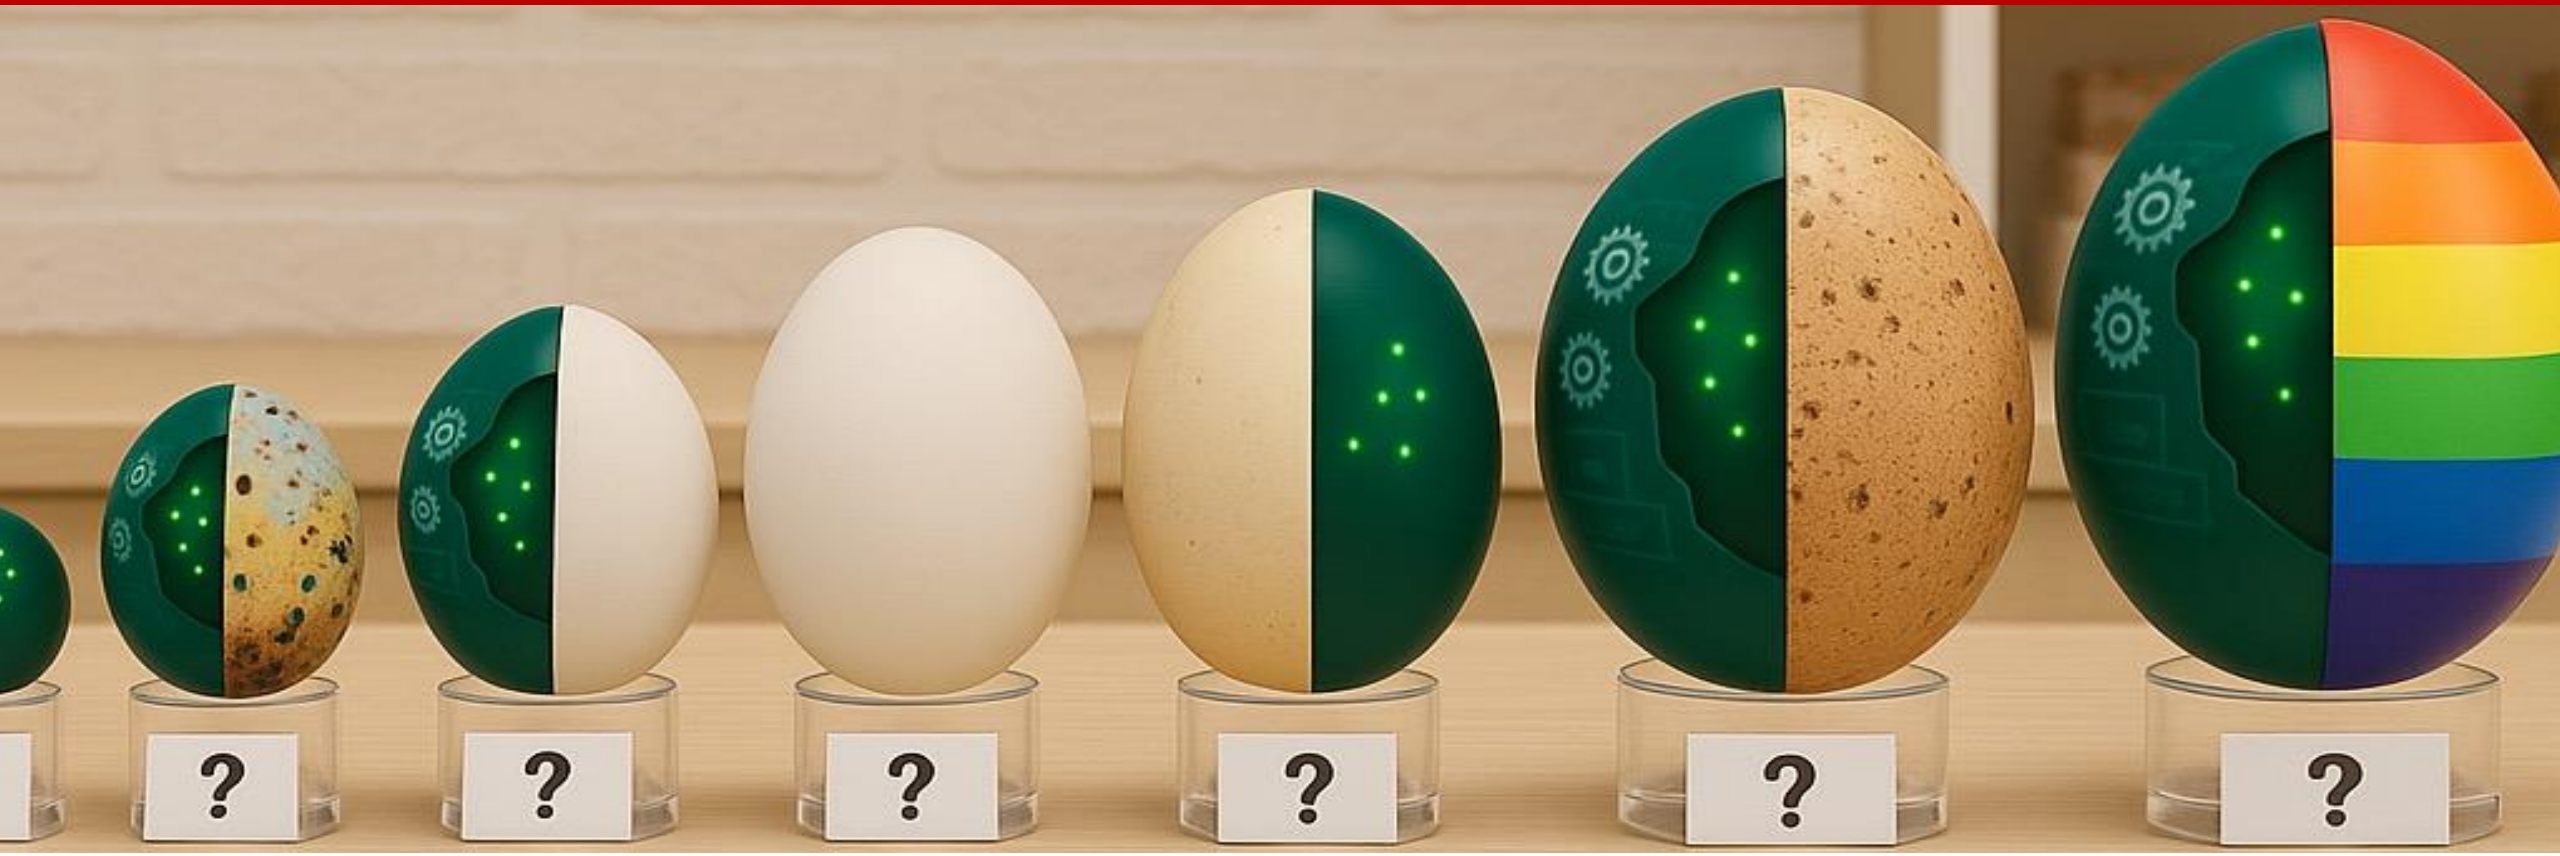

WHICH EGG IS HARMLESS? WHICH ONE IS RISKY?

AI systems are in many ways like surprise eggs: you never know exactly what will hatch from them. From the outside, it's almost impossible to tell. So how can we trust them? How do we know which egg carries which risks?

That's why the AI Act introduces a kind of quality grading system for AI eggs. It distinguishes prohibited AI practices and additional risk classes. The good news: the vast majority of AI eggs are trustworthy! They fall into Class 1 – and are therefore considered fully reliable. The same applies to all other risk classes, provided their specific safeguards are correctly implemented.

The only exception, of course, are the prohibited AI practices. And those are what we will look at next.

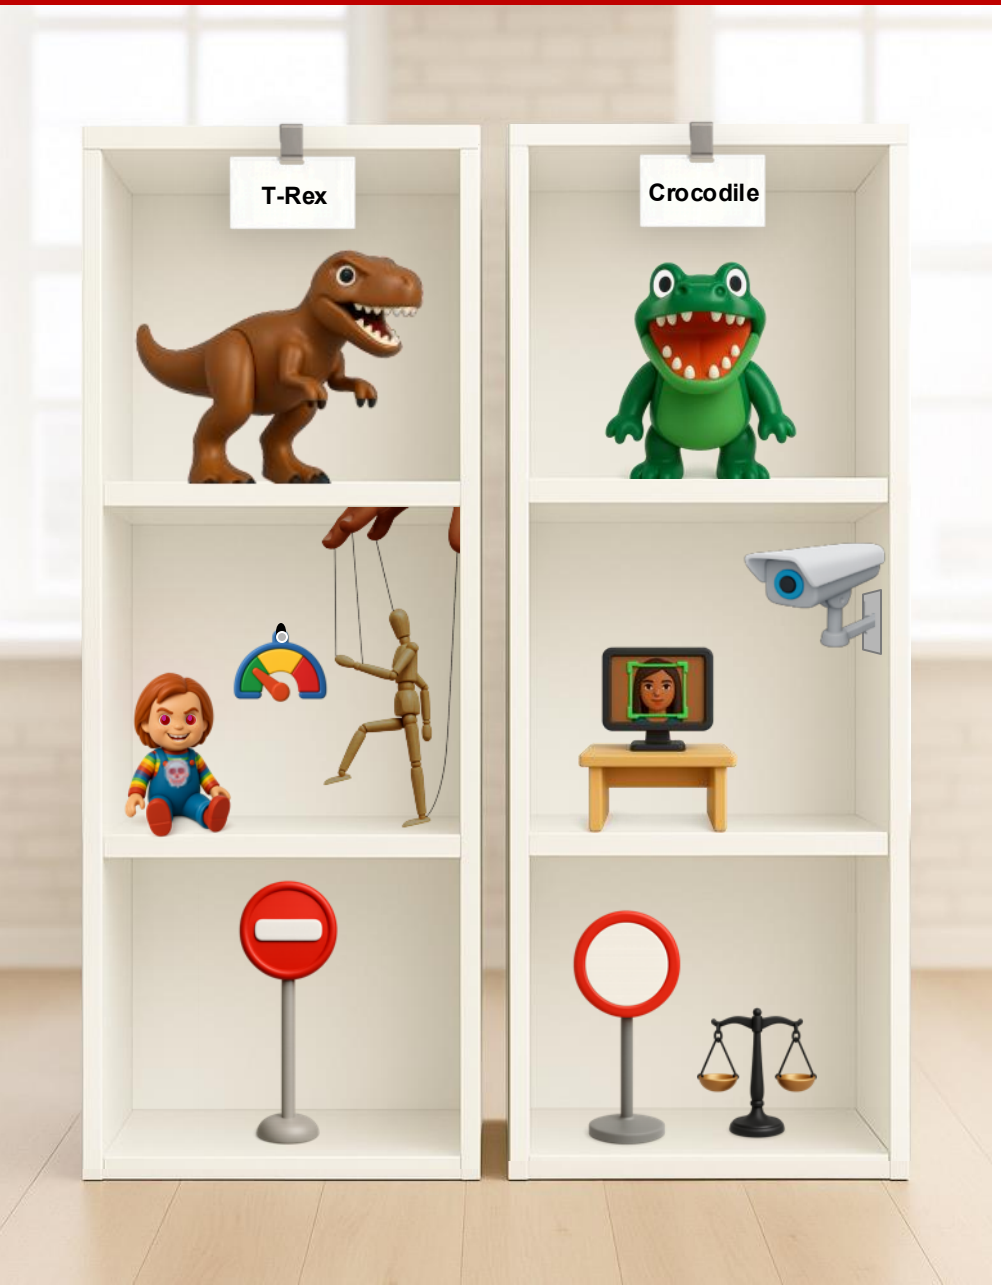

PROHIBITED AI PRACTICES: T-REX VS. CROCODILE

The AI Act distinguishes between two types of prohibited AI practices:

- Those that are absolutely forbidden ("T-Rex practices"), and
- Those that are generally prohibited, but may be allowed under strict conditions or specific exceptions ("Crocodile practices").

The key word here is *practices*: it's about how AI systems are used. The practice determines the prohibition of the AI behind it.

Let's start with the T-Rex practices. Their ban protects especially important values. Vulnerable people and children must not be manipulated by AI. Social scoring is likewise strictly forbidden. And there are further bans of this type.

With T-Rex prohibitions, an AI system may not even be placed on the market in the first place. The goal is that the risk never materializes – this is prevention at its strongest. These prohibitions apply to all users, whether public or private.

Now to the crocodiles. Unlike the T-Rex, they are not extinct. But they're certainly not cuddly animals either.

Here, the ban is fundamental, but exceptions or conditions exist:

- Emotion recognition in the workplace is prohibited – but there are exceptions, for instance in medical contexts.
- Certain forms of biometric remote identification are also prohibited – but may be permitted under strict requirements.

Some crocodile prohibitions apply only to the state: public authorities may use risky AI only under specific circumstances.

In short: the EU AI Act makes sure that neither T-Rex nor crocodiles roam freely. This strengthens basic trust in the responsible use of AI.

To build even more trust, the AI Act also regulates use cases that typically carry high risks. We'll turn to those on the next page.

HIGH-RISK AI IS OFTEN SPECIFIC

When we talk about high-risk use cases, one thing must be emphasized: they are allowed.

Many high-risk AI applications are also highly beneficial:

- for example, medical AI,
- AI in energy supply,
- or intelligent emergency call assessment.

So, high-risk AI is not “evil” by default. In many cases, the point is to ensure that all functions work correctly and that no errors occur which could harm health, safety, or fundamental values such as democracy.

Now let’s take a look at the shelf on the left. Here we see the snake symbol. But it’s a very positive kind: an Aesculapian snake. Since antiquity, it has been the symbol of medicine, health, and the medical profession.

Below it, in the middle compartment, we see symbols for other typical high-risk use cases: finance, critical infrastructure, and aviation.

Of course, there are many more examples: especially in the workplace, great care must be taken when using AI, not least to prevent discrimination.

Now let’s move to the right-hand side of the shelf. This makes clear how the hats matter in the context of risk. High-risk AI is often linked to specific AI systems (the red egg). It is less common to see high-risk AI combined with a GPAI system (the colorful egg) – for instance, in medical advisors or HR processes.

The special obligations for providers and deployers of high-risk AI are highlighted differently in the right-hand shelf: that’s why the red egg is larger, and the hats for its providers and deployers are also larger than those for the colorful egg – simply because this combination is comparatively rare.

The same principle applies in reverse for medium-risk AI: as we will soon see, the colorful egg appears far more frequently there.

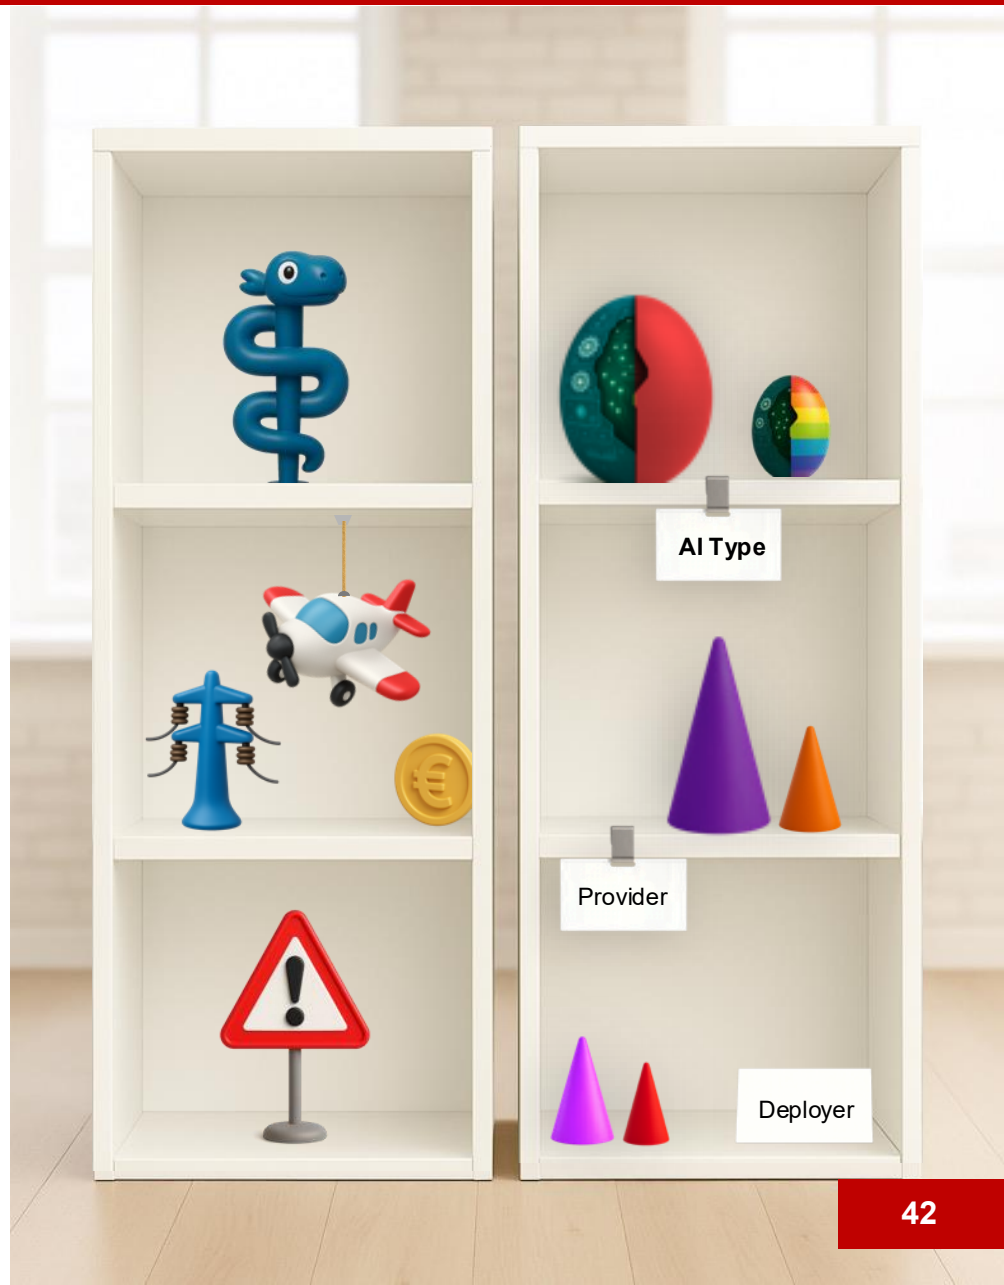

MEDIUM RISKS: MOSTLY A CASE FOR GPAI

As we've seen, certain types of AI systems are more likely to appear in one of the three permitted risk classes.

For the medium-risk class, this means: here we most often find GPAI systems – the colorful eggs.

The reason is simple: to create deepfakes, spread fake news, or run chatbots that are almost indistinguishable from humans, you need enormous amounts of data. And that's exactly what you'll find in AI systems that have integrated a colorful coral – or even an entire coral reef.

In other words: GPAI systems!

That's why it makes sense for providers and deployers of GPAI systems to have their own hat colors. Colorful eggs are especially subject to transparency obligations! They must make it clear when you're dealing with AI and when you're dealing with a human – say, in a customer service interaction. Human voices and AI voices are often hard to tell apart. And that can be confusing.

Because parrots can also mimic human speech, they are the perfect symbol for this risk class. They even match the colorful egg in appearance – don't you think?

The positive news for those offering or deploying a specific AI system (a single-colored egg with a polyp inside): transparency obligations of the medium-risk class rarely apply to them. Though not impossible, it's uncommon.

So here we've learned not just about risks, but also about their different relevance for providers and deployers of colorful vs. single-colored eggs.

And with that, we turn to the last of the four risk classes: minimal risk.

TURTLES: BY FAR THE MOST COMMON

Now for the best news: the vast majority of AI use cases fall into the minimal-risk class. This means that as consumers and users, we can trust them – even without major obligations for their providers and deployers.

This matters because the use cases in this risk class are not only widespread but also extremely diverse. The range of possibilities is endless – for example:

- smart spam filters,
- translation tools,
- self-learning industrial robots,
- clever toothbrushes,
- AI in home entertainment – and the list goes on …

That's why the middle section of the left shelf is filled with colorful building blocks: a symbol of the sheer variety in this risk class. For these AI use cases, the message is: clear road ahead!

This is mirrored on the right shelf. There we see a particularly large green egg. Specific AI systems are disproportionately frequent in this group. That's because GPAI systems – the colorful eggs – are quite rare here:

- They usually fall under the medium-risk class we just discussed.
- As such, they are subject to transparency obligations. They only appear here when they are used for very simple, highly specific tasks. That's rare indeed!
- And in those cases, there's no risk of confusing them with a parrot.

Even better news: providers and deployers in this group have very few obligations – just one essential one. Like all actors in the other risk classes, they must ensure that AI literacy is adequately promoted.

We'll get into this and other obligations in the next station. For now, remember: T-Rex, snake, parrot, turtle …!

… got it!

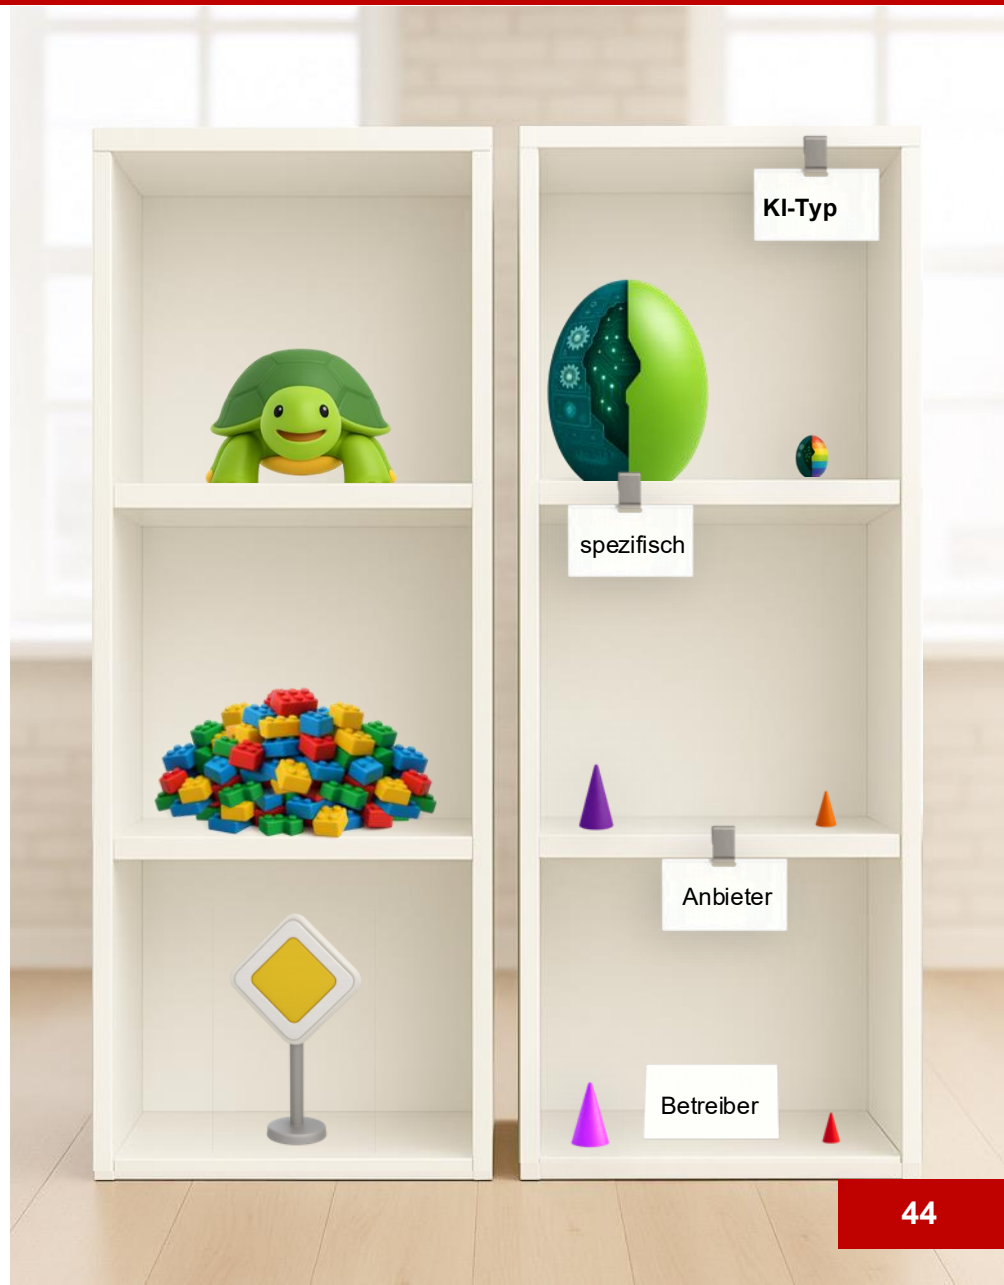

LET US RECAP STATION FIFE – THE RISK CLASSES:

1. THERE ARE FOUR RISK CLASSES: PROHIBITED AI PRACTICES, AND USE CASES WITH HIGH, MEDIUM, OR MINIMAL RISK.
2. PROHIBITIONS ARE RARE. SOME APPLY ABSOLUTELY (T-REX). OTHERS ALLOW EXCEPTIONS OR COME WITH CONDITIONS (CROCODILES).
3. HIGH-RISK USE CASES ARE OFTEN BENEFICIAL (AESCULAPIAN SNAKE). BECAUSE OF THIS, THEY MUST WORK SAFELY AND RELIABLY.
4. MEDIUM RISKS ARE MOSTLY LINKED TO GPAI. LIKE A PARROT, AI CAN IMITATE – AND IT MUST ALWAYS BE CLEAR WHEN THAT'S THE CASE.
5. MINIMAL-RISK USE CASES ARE THE MOST DIVERSE AND THE MOST COMMON. HERE, WE REMEMBER THE FRIENDLY TURTLE.

NOW WE MOVE ON TO STATION SIX: THE OBLIGATIONS

6. OBLIGATIONS

KEEP A COOL HEAD!

First things first: relax! Everywhere you hear about the countless obligations that are supposedly about to fall on all types of actors. At times, pressure is stirred up – often overdramatized by a few black sheep eager to do business – that in reality affects only a small group.

This does not downplay the importance of obligations. On the contrary: if you have obligations, you must take them seriously and implement them diligently.

But the AI Act is not a sprint – it's a marathon. Energy must be paced wisely, and unnecessary early sprints should be avoided, or else motivation and resources will be wasted.

The key is to focus on what really matters: what absolutely must be observed in terms of obligations. And let's not forget – the AI Act doesn't just impose obligations, it also grants important rights. Especially in the complex AI value chains we've already explored.

So here's the serious recommendation: the AI Act must be observed, but everything also takes time – and above all, calm. Rushing only creates frustration and unnecessary mistakes.

The only ones who need to be truly alert right now are the T-Rex and the crocodile. If you spot one – or a crocodile without a license – it should be reported to the authorities immediately!

And yes: the parrot must also be kept in mind early on. But as we learned on the previous pages, most AI use cases are turtles. They only have a few obligations to fulfill – and that's a good thing.

And what about the useful Aesculapian snake? True, more obligations will apply here. But step by step, with appropriate transition periods.

So let's approach obligations this way: focused, calm, and confident that most of them are manageable.

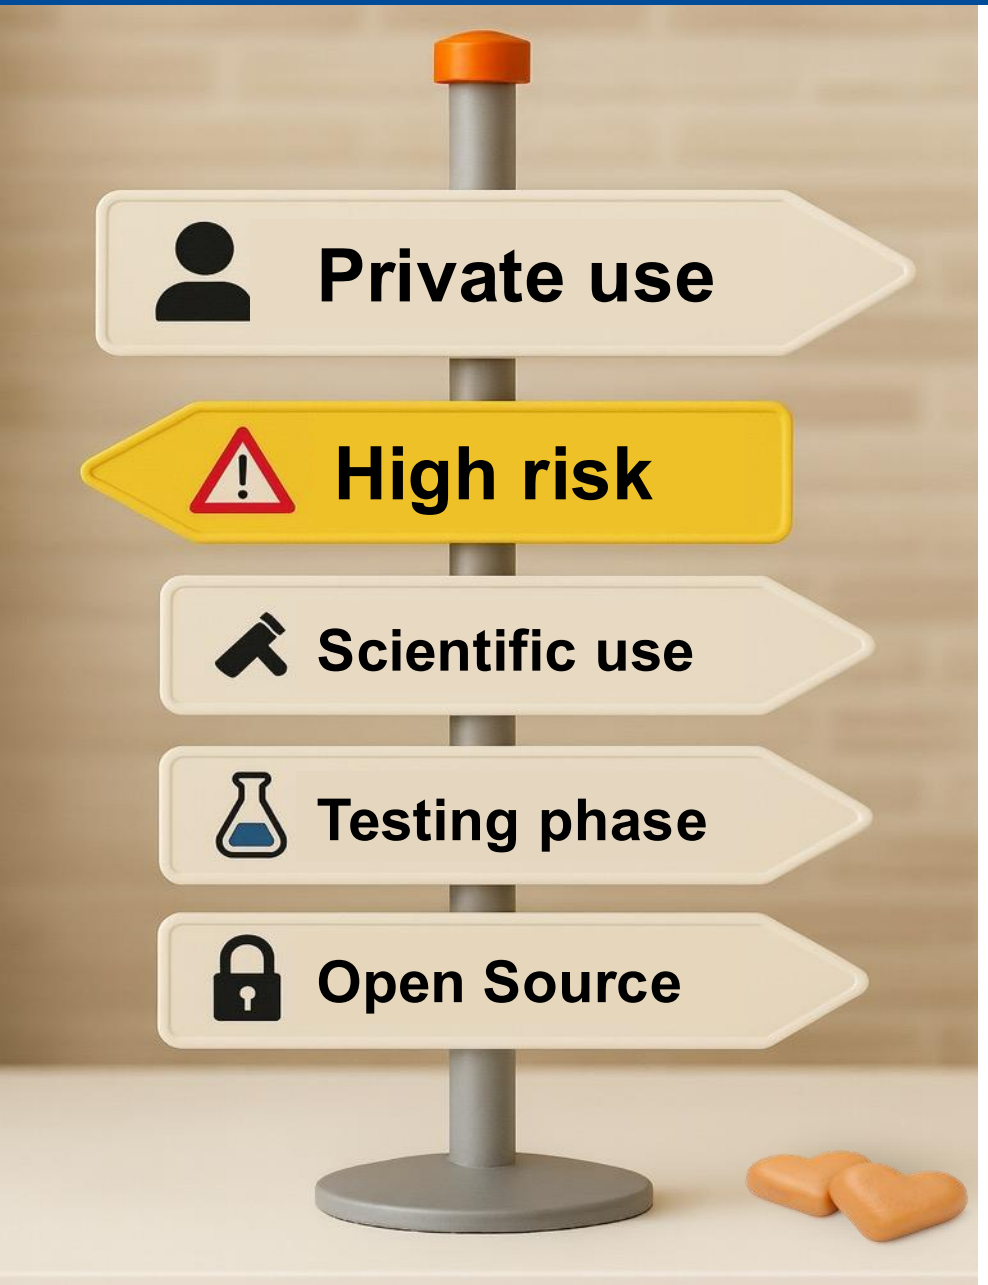

PRIVATE USE AND OTHER EXCEPTIONS

Why was our actor on the previous page so relaxed? Ah yes: he wasn't wearing any actor's hat at all. He is a purely private AI user. And that puts him under one of the several exceptions of the EU AI Act.

Anyone using an AI system purely for private purposes doesn't really need to worry about obligations. And that's a good thing – here, full relaxation applies.

But private use really does mean 100% private use! Whoever uses their private AI account for business purposes automatically becomes an deployer – and must comply with certain obligations, depending on whether, for example, a high-risk AI case is involved or transparency requirements must be met.

Things are similar, yet different, with other exceptions under the AI Act. For example, the purely scientific use of AI systems and GPAI models:

- They also fall under the exceptions of the EU AI Act.
- But here too, it must be genuine free research.

The same applies to the testing phase of AI systems – meaning trials before they are placed on the market.

Finally, there are also exceptions for open-source AI systems and GPAI models.

However, no matter which exception may apply under the AI Act:

- In the case of high-risk AI, it may be different again. For example, open source does not count as an exemption here.
- A chatbot can also act like a colorful parrot and restrict the open-source privilege.
- And even in the testing of high-risk AI, certain rules must be observed.

But as we learned earlier, most use cases are harmless turtles. With them – and with many other applications – one thing remains key:

Building AI literacy! And that's exactly what we'll look at next.

THE HOUSE OF AI LITERACY

Perhaps the most important task of all providers and deployers is to ensure sufficient AI literacy. Literacy that enables as many users as possible to use an AI application safely and with confidence.

This is based on an empirical insight: misuse or incorrect handling of AI is one of the most common causes of AI-related risks – both for AI systems and, to some extent, for AI models.

Let's recall the example of the bee: it is a very useful creature, but if you tease it or treat it wrongly, it may sting – painfully so. We need to learn how to handle it properly!

But how do you acquire AI literacy?

On the one hand, every AI system or model is unique. On the other hand, users vary widely in their background knowledge and in the tasks they want to achieve.

That means: building AI literacy is not a "one-size-fits-all" exercise. Rather, every relevant person should get their own "AI literacy house." And this house is made up of different building blocks.

The picture on the right illustrates which building blocks these are. Depending on risk class, use case, and role, the result will be small colorful houses – and they look different for everyone.

What matters is this: The AI Act requires AI literacy to be promoted for systems of all risk classes. Even for those with low risk. Because if used incorrectly, even seemingly harmless AI can suddenly become problematic.

In short, AI literacy is a very important and at the same time highly individualized task. But once you know the right building blocks, you can build the right house for everyone – and only in the case of high-risk AI does it sometimes have to be a tall competence tower.

Key takeaway:

AI literacy = colorful house

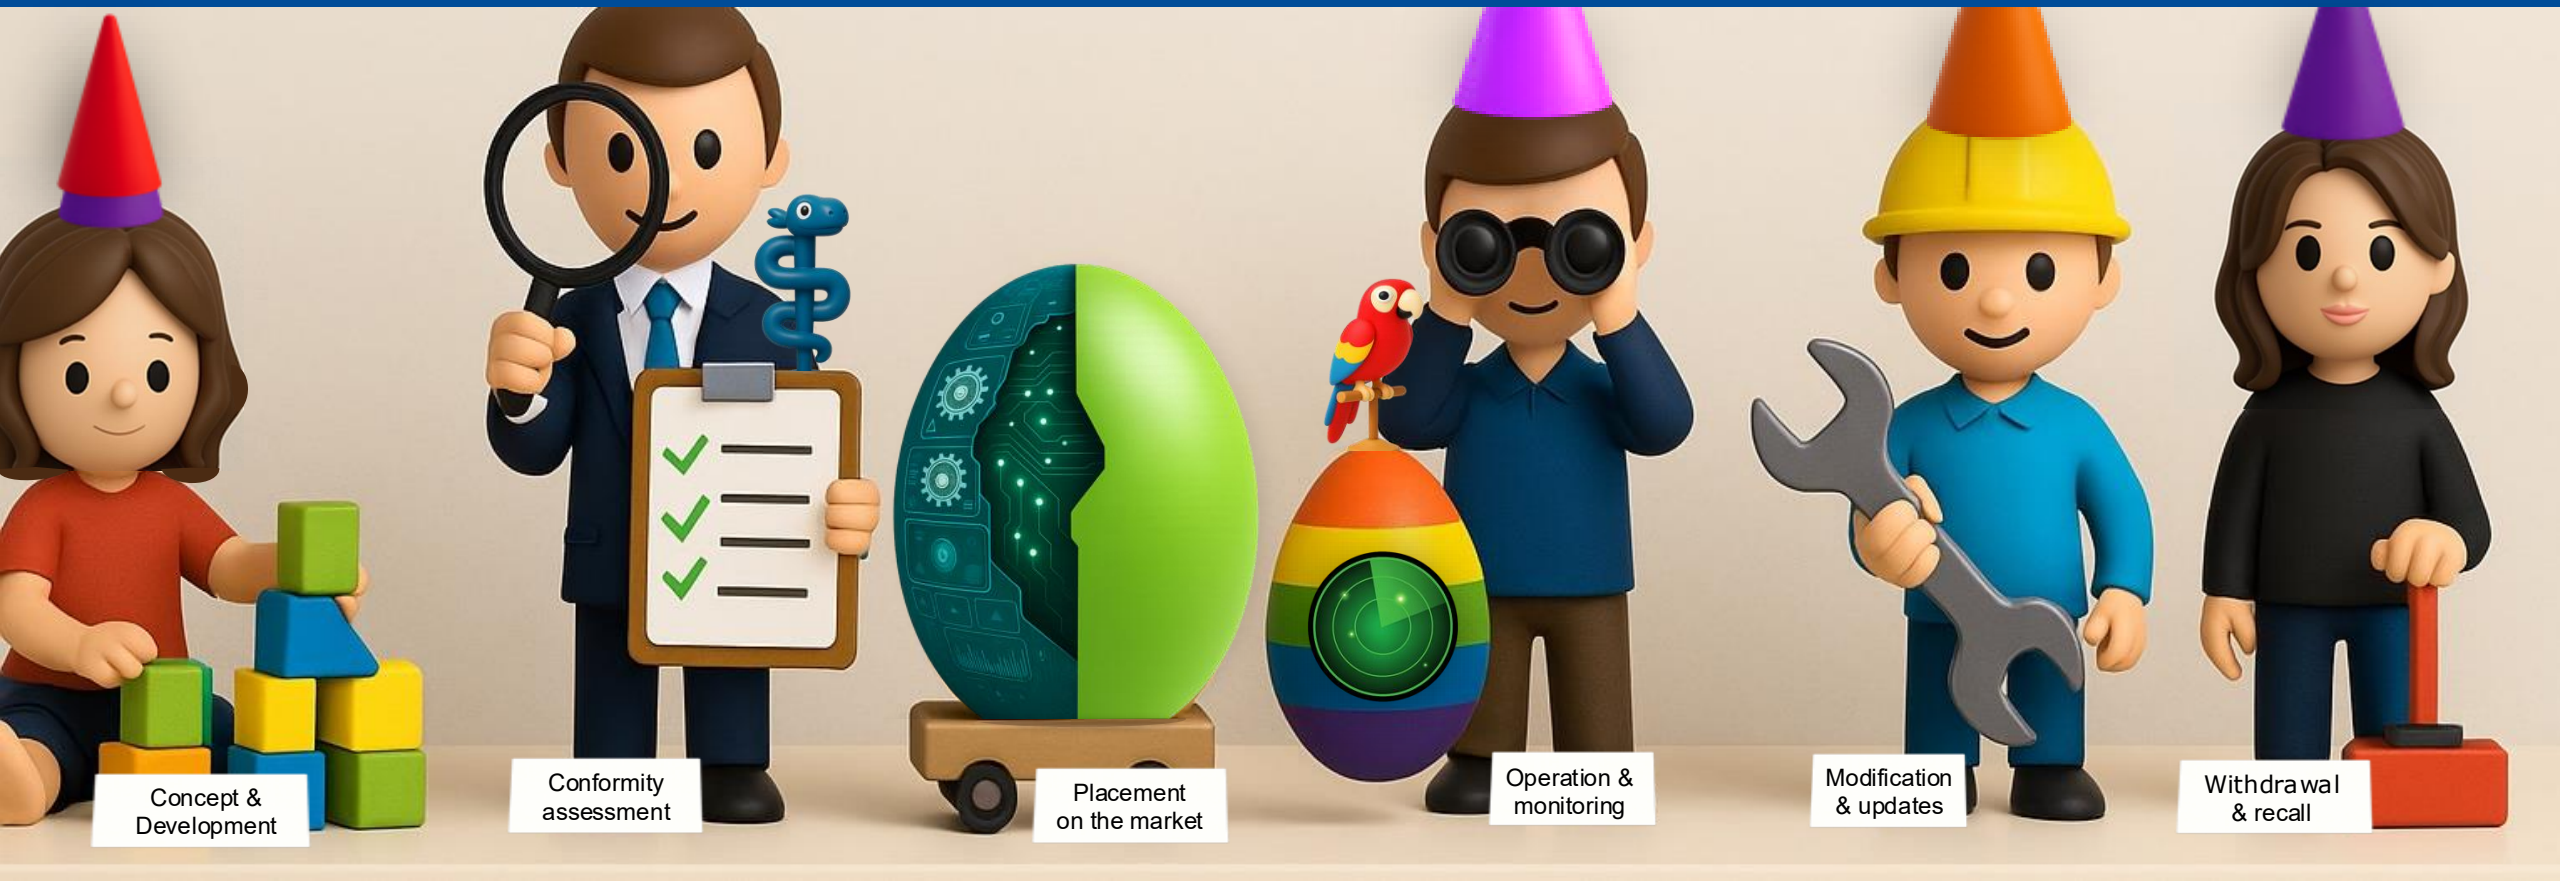

OBLIGATIONS: DEPENDING ON RISK CLASS, ROLE, AND LIFE CYCLE

What do we need to know about the specific obligations under the AI Act? Above all, that they depend on a wide variety of factors.

A key element is the life cycle of an AI system. The term fits well with the egg symbol – after all, an egg develops over time, too. The illustration shows typical phases in the development of an AI system. Certain steps, such as conformity assessment, apply only to high-risk AI systems.

Other obligations, in turn, depend on the type of AI as well as the role of the actor – symbolized by their colored hats.

Most importantly: obligations are not a one-off starting activity. They extend across the entire life cycle of AI.

DEADLINES AND SANCTIONS

The EU AI Act has already entered into force. However, for some obligations (such as those applying to high-risk AI) there are transitional periods. There is also protection for existing AI systems that have not been modified. In other words: not everything that has already been successfully deployed must be rebuilt from scratch.

Nevertheless, it is important to take the obligations seriously! Failure to comply can become very costly. And the chances of being caught if you ignore the rules are not small:

- The AI Act includes what is known as whistleblower protection.
- The figure with the referee's whistle symbolizes this.
- Whistleblowers can bring things to light if someone tries to cheat the system.

And if someone deliberately violates particularly important obligations, an AI system can even be withdrawn from the market …

So: non-compliance with obligations is no trivial matter. Let's remember the cash register, the hourglass, and the whistle!

LET US RECAP STATION SIX – THE OBLIGATIONS:

1. YES, THERE ARE OBLIGATIONS FOR MANY ACTORS. BUT THE KEY IS TO APPROACH THEM CALMLY AND WITH CONTROL!
2. REMEMBER THE EXCEPTIONS: PERSONAL USE, SCIENTIFIC RESEARCH, TESTING, AND OPEN SOURCE.
3. TAILORED AI LITERACY IS ALWAYS ESSENTIAL! LET'S KEEP IN MIND THE COMPETENCE HOUSE.
4. OBLIGATIONS CAN STRETCH ACROSS THE ENTIRE AI LIFECYCLE. THEY ARE NOT A ONE-OFF TASK!
5. LET'S REMEMBER: THE HOURGLASS (DEADLINES), THE CASH REGISTER (SANCTIONS), AND THE WHISTLE (WHISTLEBLOWERS).

NOW WE MOVE ON TO STATION SEVEN: DATA

7. DATA

A LESSON IN DATA PEARLS

Data is the new oil?" Not quite: Data are pearls! They come in many different shapes, colors, and – most importantly – qualities. Especially in AI, the quality of the pearls matters most. Let's start with the coral pearls, since AI models are made up of polyps or corals. This type of data is therefore embedded within the AI model – often rich, but sometimes unstructured. Different from that are the system pearls in the middle: these are structured system data. And on the right? Those are just worthless pebbles – and there are far more of them than we'd like!

That's why it's essential to distinguish between different types of data right from the start: coral pearls and system pearls!

REAL AND SYNTHETIC CORAL PEARLS

Coral pearls are painstakingly trained into a GPAI model – or into individual polyps. But not all pearls are good ones.

Quite often, pebbles slip in as well. They add noise and disorder to the swirl of coral pearls. That's why these sometimes wild coral pearls should be handled with care!

The same applies to synthetic coral pearls: training data that are artificially generated. They look similar to real pearls but are tinted blue, green, or yellow – different from natural coral pearls, which are usually red or beige. Synthetic pearls are useful when too few real pearls are available. They can be helpful and plausible – but they are not real. And they can distort a model.

The origin of coral pearls can also be problematic: some may come from "data nature reserves," creating copyright issues.

Personal data should likewise not be used as training pearls. Once embedded, they are extremely hard to remove.

Whoever integrates a coral or coral reef into their AI system becomes a downstream provider (blue hat). They depend on information from the original provider (yellow hat): Which pearls were used? How were they trained?

The EU AI Act defines specific transparency obligations for corals, coral reefs, and their pearls.

It's similar – but stricter – for high-risk AI: here, the quality of coral pearls must meet defined criteria. Biases, for example, must be examined, since they can lead to faulty outputs and even harm.

Throughout the lifecycle, coral pearls must be checked again and again: new pebbles may have slipped in … and those need to be found and removed.

Remember:

Model data = Coral pearls

CONNECTED SYSTEM PEARLS

Let's recall the three compartments of the egg: in the middle sits intelligence, the coral. We described that with the coral pearls on the previous page.

In the lower compartment, however, we find different pearls – not coral pearls, but pearls from databases, stored in special jewel boxes. And unlike the irregular coral pearls, these are beautifully uniform and structured: true mother-of-pearl pearls.

These data are also subject to strict quality requirements in high-risk AI.

Take, for example, an AI system that automatically screens job applicants and evaluates their suitability. Here, many data points come directly from a database, such as:

- The job profile as a requirements list,
- the prioritization of criteria, and
- other factors considered particularly important – or potentially disqualifying.

These are not coral pearls embedded in the AI model. They come from databases directly connected to the system.

Other examples of system pearls include domain-specific knowledge like legal texts or form templates. They supplement the data stored within the AI model.

Whenever precise reproduction of information is critical, system pearls tend to be far more reliable and accurate than coral pearls.

From a regulatory perspective, outside of high-risk AI there is much greater flexibility regarding system pearls. However, another legal framework becomes especially important here: the GDPR, and thus the protection of personal data. These must be handled with care and remain deletable. When stored in databases, system pearls can usually be deleted quite easily.

As for actor roles, it is important to note that an deployer who connects their own data to the system in a way that substantially modifies it may themselves become a (quasi-) provider.

So let's keep the distinction clear:

System data = real pearls

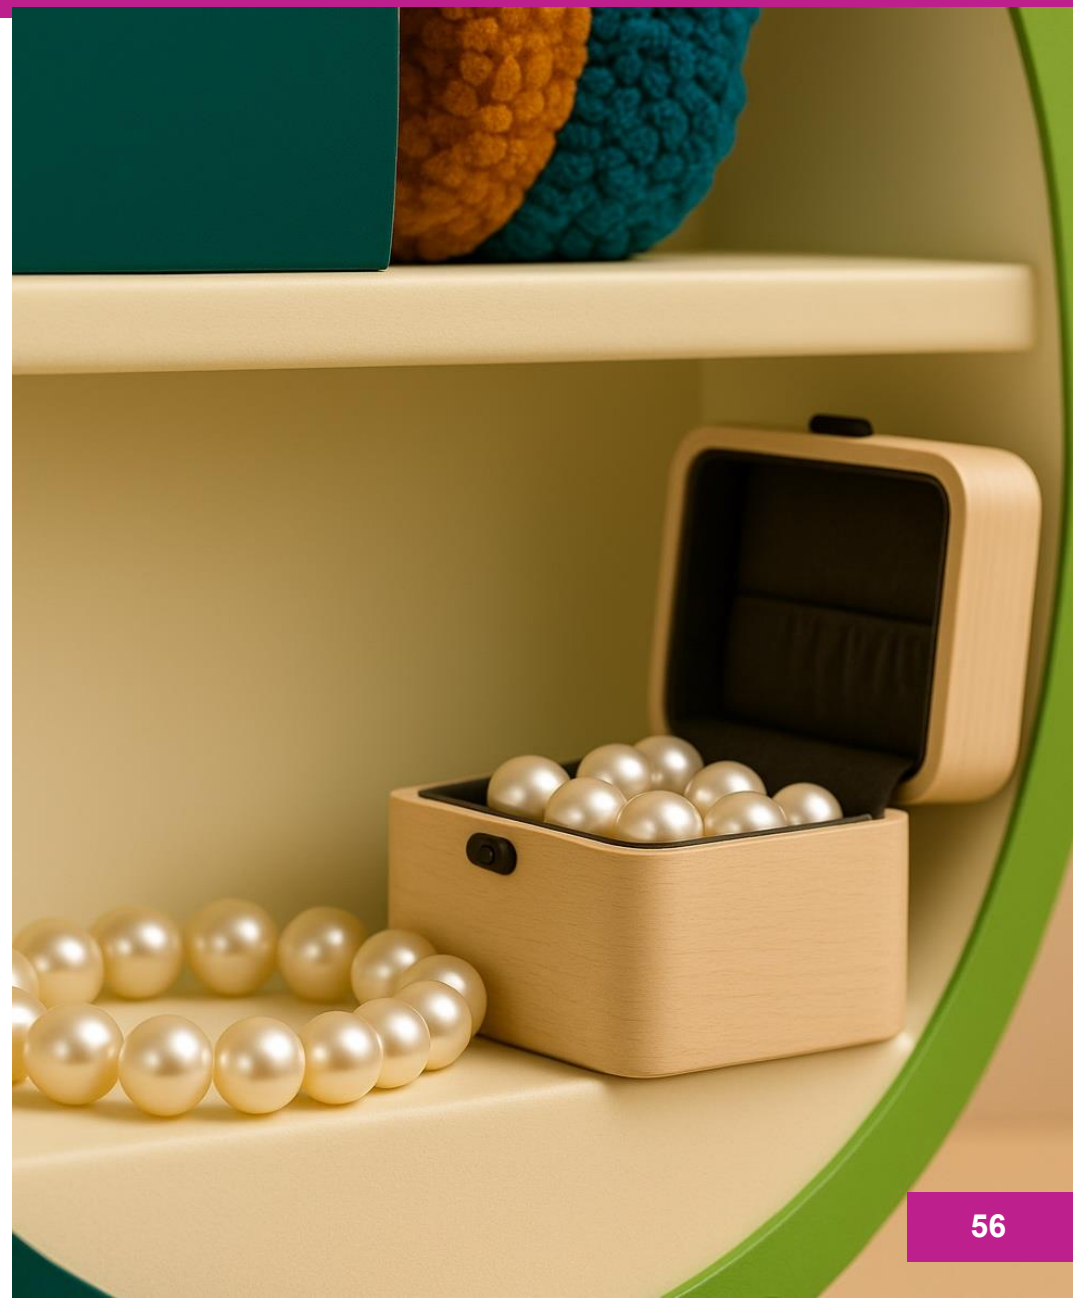

PEBBLES: DATA WASTE AND BAD PROMPTS

Unfortunately, many AI providers, deployers and users believe they possess the finest data pearls. But all too often, they are only pebbles. And if you feed the best GPAI model (trained with the finest coral pearls) with pebbles as input, the output will also be pebbles – not pearls.

That's why the phrase applies so well: "Garbage in, garbage out." If you put in data waste, you'll get data waste back.

Pebbles, in this sense, are data garbage. They can arise in many ways, for example:

- Typos or misspellings in system data,
- missing or incorrect values in datasets that are not properly handled,
- outdated or poorly structured data,
- incorrect linking or integration of data, and so on.

Especially critical are errors in training data: for instance, in HR software that can lead to biased hiring decisions, or in medical datasets that can produce dangerous misdiagnoses.

But even if training and system data are of good quality, poor input data (prompts) can generate continuous new data waste.

From a regulatory perspective, this is highly relevant. It shows that human behavior and a lack of AI competence are often to blame when AI systems produce poor – and potentially risky – results.

That's why building AI competence is so important. One element of the "Competence House" we discussed earlier is learning how to use AI correctly – during training, when connecting system data, and above all when formulating input prompts.

On the left, we see the three data garbage bins: even if we sort and separate waste, garbage remains garbage.

Let's remember:

Pebbles = data waste

CAUTION WHEN RETRAINING AND FINE-TUNING WITH CORAL PEARLS

When specific AI models or GPAI models are retrained by the deployer or when fine-tuning of their content takes place, there is a risk of a role shift.

- On the left, we see the normal role distribution: the provider wears the orange hat, the deployer the red one.
- In the middle, the deployer adds coral pearls to the GPAI model. This substantially alters the system as a whole. As a result, the deployer gains an additional orange hat. Alongside the original provider, they now become a quasi-provider of the GPAI system. We now have two providers.
- The third variant is special: in the case of high-risk AI, a substantial modification triggers a complete role swap. From that point on, there can only be one provider – the deployer, who through fine-tuning becomes the final provider.

So: Caution with retraining or fine-tuning!

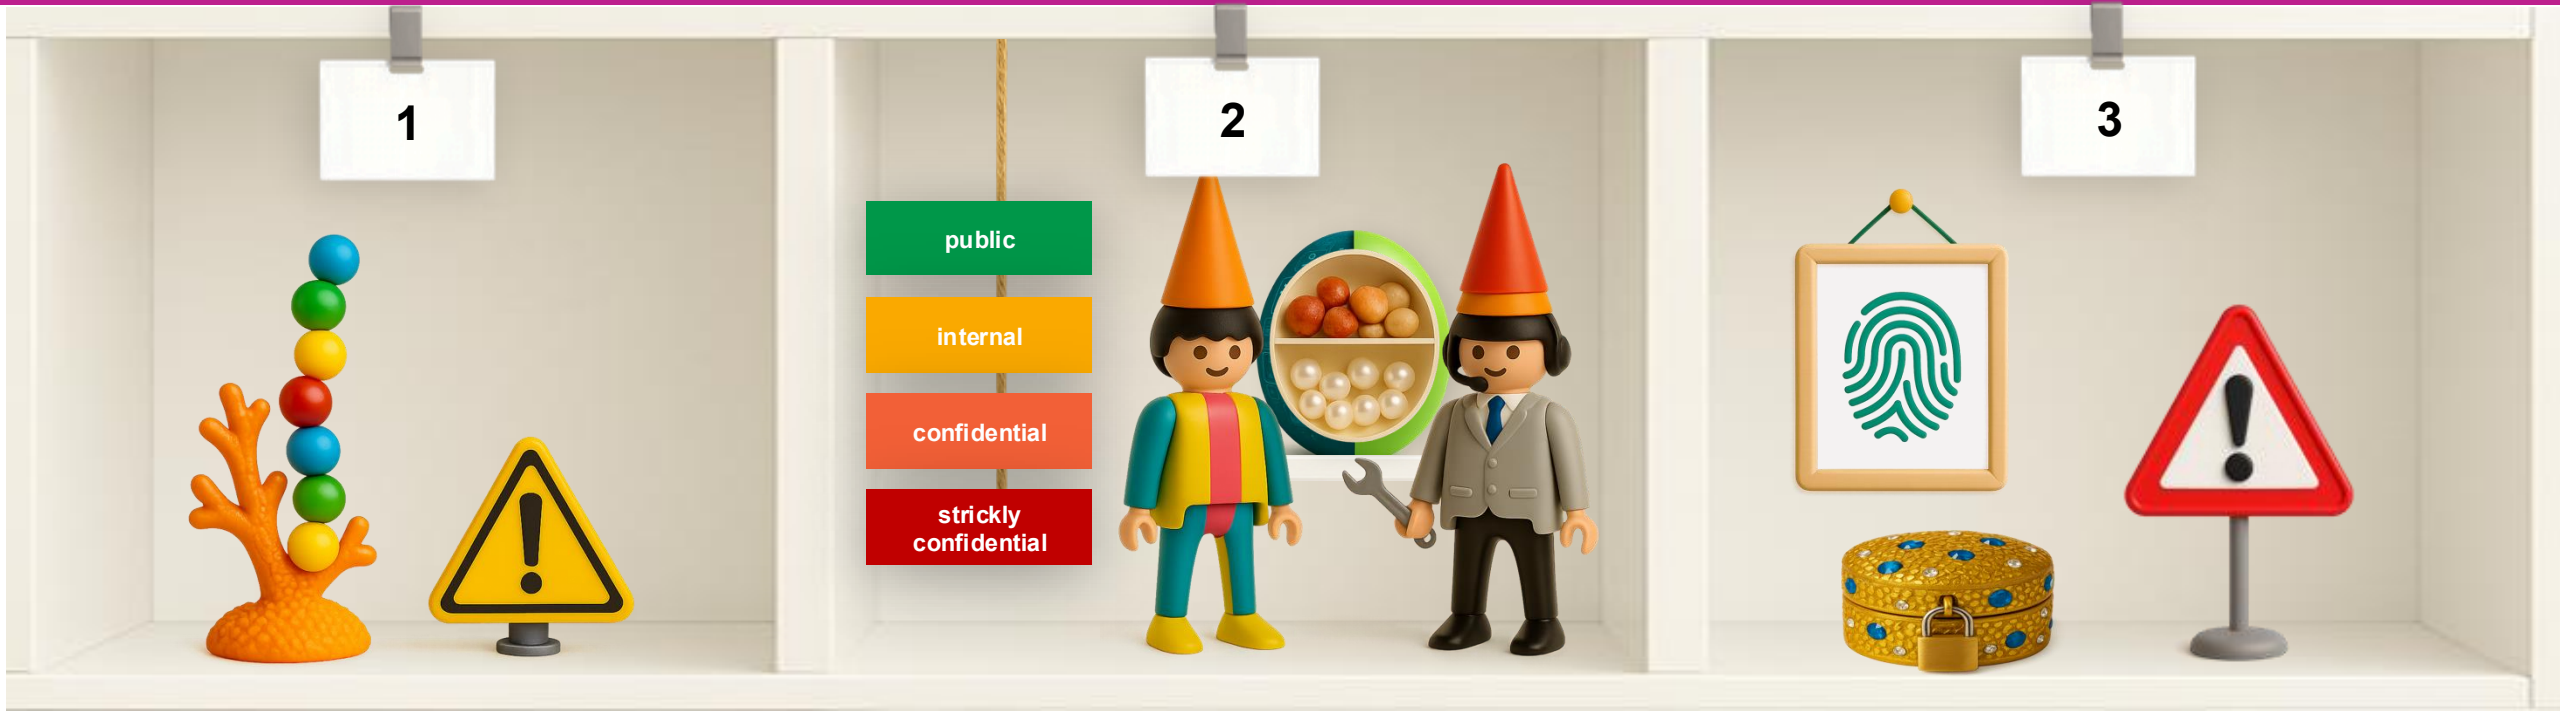

THE RIGHT HANDLING OF DATA IS ALWAYS ESSENTIAL

The AI Act highlights the importance of data in two specific areas:

- For high-risk AI, data governance is explicitly regulated.
- In addition, confidentiality obligations apply to public authorities when they review data from actors.

The first point should also be taken seriously by providers and deployers of AI systems in the medium- and low-risk categories. While violations may not lead to sanctions in these cases, there are still general duties of care that apply independently of the AI Act. On top of that, the GDPR must always be respected.

In practice, this means:

1. Training data should always be carefully checked for quality and continuously documented.
2. System data must be regularly verified for accuracy and timeliness, and classified in terms of confidentiality. These are processed by the AI model and thereby “leave the system.”
3. Personal data requires particular caution: it must be erasable, and its use must be properly logged.

So, regardless of the risk class, responsible data management is always crucial.

LET US RECAP STATION SEVEN – DATA:

1. DISTINGUISH BETWEEN MODEL DATA (CORAL PEARLS) AND SYSTEM DATA (SYSTEM PEARLS)!
2. DOWNSTREAM PROVIDERS (BLUE HAT) ARE ENTITLED TO TRANSPARENCY ABOUT THE DATA USED IN GPAI MODELS.
3. SYSTEM DATA (REAL PEARLS) COME FROM DATABASES – THEIR QUALITY MUST BE CAREFULLY CHECKED, ESPECIALLY IN HIGH-RISK AI.
4. PEBBLES REPRESENT BAD DATA – WHETHER IN MODEL TRAINING, AS SYSTEM DATA, OR AS INPUT DATA.
5. RETRAINING AND FINE-TUNING CAN EVEN LEAD TO A ROLE SWAP. WHENEVER DATA IS INVOLVED: HANDLE WITH CARE!.

LET'S START WITH A QUICK RECAP –
THEN MOVE ON TO DEEPER INSIGHTS & APPLICATIONS.

8. DEEPEN

WHO IS WHO?

Do we still remember which protagonist represents which AI element?

1. Machine Egg
2. Three Compartments
3. Polyp
4. Purple & Pink Hat
5. Colorful Egg
6. Coral / Coral Reef
7. Yellow, Orange, Red Hat
8. T-Rex & Crocodile
9. Aesculapian Snake
10. Parrot
11. Turtle
12. Personal House
13. Coral Pearls
14. Real Pearls
15. Pebbles

FIRST: A CLOSER LOOK AT THE LEGAL NORMS

On the following five pages, the provisions of the EU AI Act are presented alongside the various symbols. The norms are linked so that they can be opened and read directly in a browser.

It should be noted that many definitions are contained in Article 3 of the EU AI Act and are numbered. For example, the provision for a “general-purpose AI model” (= GPAI model) is found in Art. 3 No. 63 EU AI Act. After opening the link to Art. 3 EU AI Act, one must therefore scroll down within the article to the corresponding number.

Reading the original legal text is recommended not only for lawyers: the often technically demanding formulations reveal how difficult it is to describe highly dynamic AI topics in a timeless way. The use of numerous vague legal terms is therefore almost unavoidable.

The structure of the risks, for example, is also revealing. Legally regulated in a direct sense are only the high-risk topics, including their conditions and obligations. Medium and low risks are not explicitly regulated. However, they can be derived from the overall structure of the AI Act.

A closer look at the legal provisions should definitely be undertaken, so that the key articles of the AI Act become familiar.

8. DEEPEN – AI SYSTEM, GPAI MODEL AND ACTORS

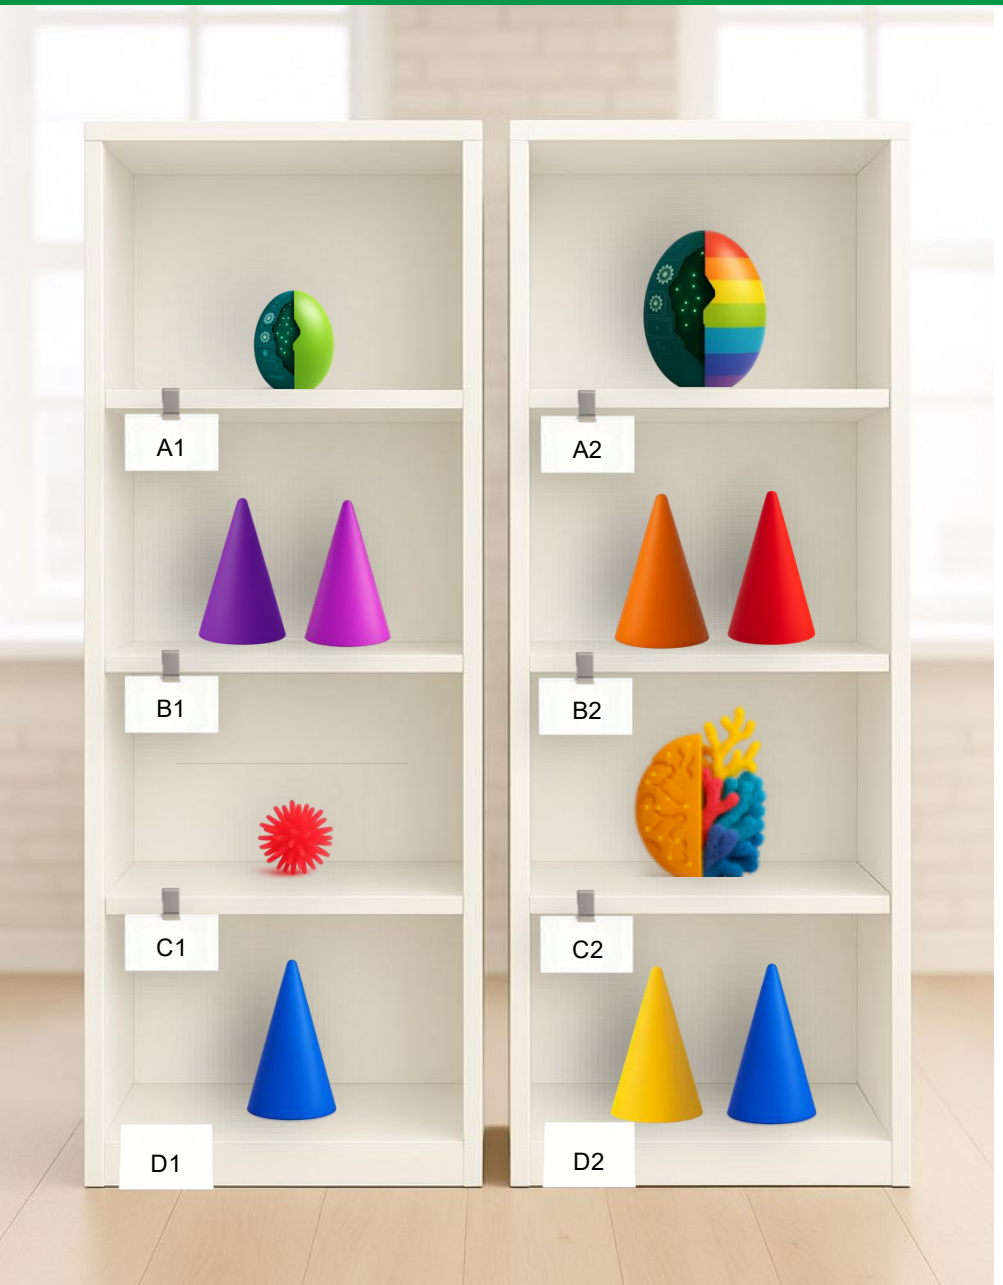

Standards in the EU AI Act	Explaination
A1) AI system (in general): • Art. 3 No. 1 EU AI Act Link to standard **A2) GPAI system:** • Art. 3 No. 66 EU AI Act Link to standard	The AI Act distinguishes between two types of AI systems: the general AI system (usually specific) and the GPAI system (an AI system with general-purpose use). Both have their own legal definitions.
B1) Provider & deployer (in general): • Art. 3 No. 3 & No. 4 EU AI Act Link to standard **B2) Provider & deployer (GPAI system):** • Art. 3 No. 3 & No. 4 EU AI Act Link to standard	The roles of provider and deployer apply equally to both types of AI systems. Nevertheless, a color distinction is useful: obligations under Art. 50 EU AI Act often concern GPAI systems, while high-risk obligations more frequently affect specific AI systems.
C1) AI Modell (in general): • No definition **C2) GPAI Model:** • Art. 3 No. 63 EU AI Act Link to standard	It is important to note: there is no legal definition for the AI models of specific AI systems. This is partly due to the technology-neutral approach of the AI Act. The definition for GPAI models therefore explicitly refers to versatility.
D1) Provider (in general & downstream): • Art. 3 No. 3 & No. 68 EU AI Act Link to standard **D2) Provider (GPAI Model & downstream):** • Art. 3 No. 3 & No. 68 EU AI Act Link to standard	The role of provider is defined only for GPAI models. However, for both types of models (specific and general-purpose), there is the notion of the down-stream provider. This role is particularly relevant within the AI value chain.

8. DEEPEN – DATA

Standards in the EU AI Act	Explaination
E) Prohibited practices: • Art. 5 EU AI Act Link to standard • Add on: Annex II Link to Annex II	Prohibited practices are primarily regulated in Art. 5 EU AI Act. The specific requirements for public authorities are set out in Annex II. This annex therefore applies only to the “crocodiles,” not to the “T-Rex prohibitions.” The latter are always prohibited and thus do not require (criminal law) exceptions.
F) High risk AI: • Art. 6 EU AI Act Link to standard • Add on: Anhang I / III Link to Annex I, Link to Annex III • Obligatoins: Art. 8 & Art. 71 et seq. EU AI Act Link to Art. 8; Link to Art. 71;	High-risk AI has several variants: those linked to products listed in Annex I, and additional use cases specified in Annex III. The obligations are laid down in Art. 8 et seq. EU AI Act as well as Art. 71 et seq. EU AI Act. Importantly, high-risk AI requires a conformity declaration (Art. 47 in conjunction with Annex V).
G) Medium risk AI: • No definition of „medium risks“ • But transparency obligations for providers and deployers of GPAI systems • Art. 50 EU AI Act • Link to standard	It should be noted that the medium-risk category is not explicitly mentioned as such in the EU AI Act. Its existence, however, follows from Art. 50, which establishes specific obligations for “certain” AI systems. These are usually GPAI systems. The most important obligations are the transparency requirements of Art. 50 and the duty to ensure AI literacy under Art. 4.
H) Low risk AI: • No definition of „low risks“ • But transparency obligations for providers and deployers concerning AI literacy • Art. 4 & Art. 3 Nr. 56 EU AI Act • Link to Art. 4, Link to Art. 3	Similarly, the low-risk category is not explicitly named in the Act. It can, however, be derived from Art. 4. According to this provision, sufficient AI literacy must be ensured for all AI use cases – including low-risk systems. If this is not done, it may amount to a breach of due diligence obligations.

8. DEEPEN – ACTIVITIES

Standards in the EU AI Act	Explanation
I) Placing on the market: • Art. 3 No. 9 EU AI Act Link to standard • Further: making available on the market, Art. 3 No. 10 • for AI systems/GPAI systems and GPAI Models	Placing on the market is a process carried out by providers of an AI/GPAI system or of a GPAI model. The key element is the *making available on the market*, which – in addition to the definition in Art. 3 No. 9 – is further regulated in Art. 3 No. 10. According to this, handing over the system to distribution is already sufficient to constitute placing on the market. The provision, whether for remuneration or free of charge, must take place in the context of a commercial activity.
J) Putting into service: • Art. 3 No. 11 EU AI Act Link to standard • for AI systems/GPAI systems only	Putting into service is directly linked to the role of the deployer, which arises once the system is put into service: either by the provider itself, when it deploys the AI/GPAI system, or by transferring the system to a third party, who thereby becomes the deployer. The "first use" marks the threshold at which the system is effectively operated in a real context. From this point onward, obligations resulting from the intended purpose apply.
K) Intended purpose: • Art. 3 No. 12 EU AI Act Link to standard • Art. 3 No. 11 EU AI Act Link to standard • Important for high risk ai, e.g. Art. 7, 8, 10 and 25 EU AI Act • Relevant for change of provider	The intended purpose is significant in several respects: first, it relates to the putting into service of an AI system *in accordance with its intended purpose*. If an AI system is deployed beyond that purpose, this may lead, among other things, to a re-assessment of its risk classification and, in certain cases, to the deployer becoming a (quasi-)provider due to the modified intended purpose or use.

8. DEEPEN – RISK CLASSES

Standards in the EU AI Act	Explanation
L) Training data (GPAI Models): • Art. 3 No. 63 EU AI Act Link to standard • Transparency obligations, Art. 53 EU AI Act Link to standard • Concerning high risk: Art. 10 & Art. 72 EU AI Act Link to Art. 10; Link to Art. 72	The AI Act refers to training data for GPAI models in Art. 3 No. 63 only indirectly. Such data sets are typically very large (“large amounts”) and enable the general-purpose functionality. In particular, for GPAI models that are used locally or on virtual private servers (VPS), retraining may occur. However, retraining (unlike in the case of AI systems) does not result in quasi-provider status for a GPAI model. The transparency obligations under Art. 53 must nevertheless be observed.
M) System data: • Art. 3 No. 29 et seq. EU AI Act Link to standard • Training data, Validation data, input data, personal data • High risk AI: Art. 10 & Art. 72 EU AI Act Link to Art. 10; Link to Art. 72	With regard to system data, the AI Act differentiates between various types. For example, the internal AI model of a system may be trained; in this case, the model is considered system-specific. By contrast, when user preferences are stored in a database, these are generally not training data but input-based personal data. They are typically stored in a database and, as a rule, can be deleted.
N) Poor quality-data: • no definition • Concerning high risk AI: Art. 10 u. Art. 72 EU AI Act Link zu Art. 10; Link zu Art. 72 • Possible general due diligence obligation	Poor-quality data are not explicitly designated as such in the AI Act, neither with regard to AI models nor with regard to system data. However, data quality is addressed in Art. 10 in the context of high-risk systems. It is particularly relevant during annotation and data cleaning. Importantly, data quality must be regarded as a general due diligence obligation for *all* AI systems, not just high-risk AI. In high-risk cases, however, non-compliance can lead to sanctions.

8. DEEPEN – MONITORING, INCIDENTS & CODES OF CONDUCT

Standards in the EU AI Act	Explanation
O) Post-market monitoring: • Obligation, Art. 72 EU AI Act Link to standard • At high risk: Art. 10 u. Art. 72 EU AI Act Link to Art. 10 • Relevant for providers • Possible general due diligence obligation	Providers of high-risk AI systems must have in place a post-market monitoring system (Art. 72). This system must regularly assess and document relevant data (cf. Art. 10). Publicly available data on (new) risks must also be taken into account. This obligation can, in a more limited form, also be understood as a general duty of care applicable to AI systems of all risk classes.
P) Reporting of serious incidents: • Obligation, Art. 73 EU AI Act (high risk) Link to standard • Definition of „serious“: Art. 3 No. 49 Link to Art. 3 • Relevant for providers and deployers	The obligation to report serious incidents applies to both providers and deployers of high-risk AI systems. The definition of a “serious incident” within the meaning of Art. 73 is provided in Art. 3(49). It includes threats to health, critical infrastructure, fundamental rights, as well as severe property or environmental damage. Importantly, reporting obligations are subject to strict deadlines, which also apply to follow-up measures by supervisory authorities.
Q) Voluntary codes of conduct: • Promotion under Art. 95 Link to Art. 95 • Voluntary assumption of obligations designed for high-risk AI systems – also possible for medium- and low-risk classes.	From Art. 95 it can be inferred that requirements formulated for high-risk AI may also be relevant for AI systems of medium and lower risk classes. Their application on a voluntary basis means that non-compliance cannot be sanctioned. At the same time, this makes clear that many of the risks addressed in the high-risk regime are of a more general nature. Accordingly, equivalent safeguards can be interpreted as a general due diligence requirement across all AI systems.

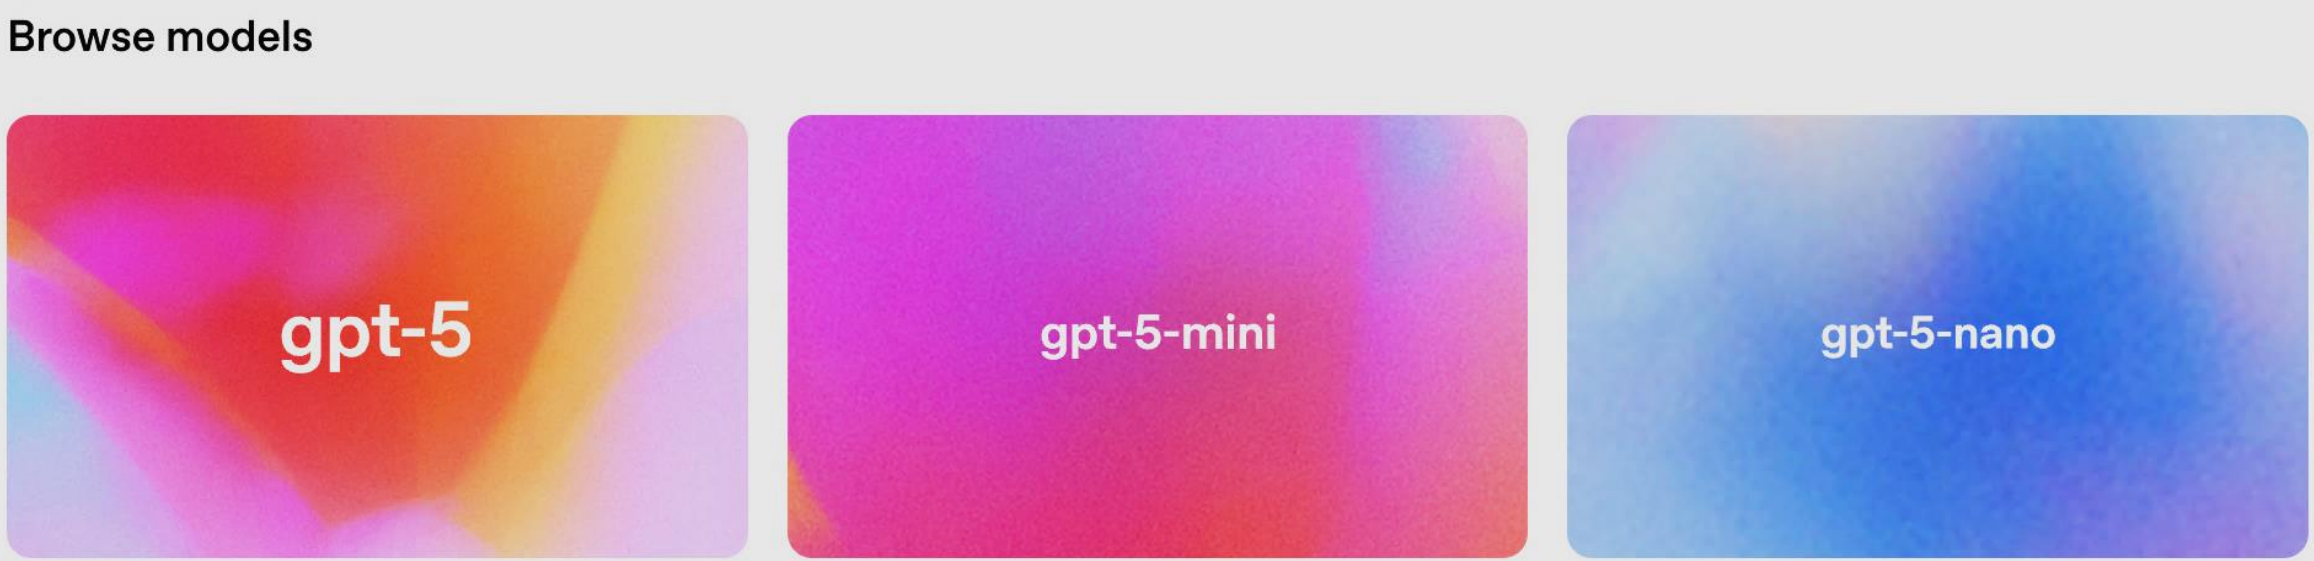

OPENAI: A REAL-WORLD EXAMPLE

The following scenario illustrates a situation that frequently occurs in practice:

- Company U wants to operate its own chatbot within the company intranet.
- U commissions the software firm S to develop an interaction interface for the chatbot that can be integrated into the intranet.
- The chatbot should be able to answer employee questions about U's business content as well as many other topics.
- For the required AI services, U obtains its own license directly from OpenAI.
- S decides to connect two AI services:
  - a large language model, and
  - an AI service that incorporates internet references (gpt-5 and gpt-4o-search-preview).
- S only develops the interaction layer and connects it to the AI services licensed by U.
- Both AI services are accessed with the same API key provided by OpenAI.

This raises the question: How should these AI services be assessed under the AI Act? And what roles do the different actors assume in this scenario?

To answer this, we first look at OpenAI's website to better understand the AI services used by U.

OpenAI serves here as an example provider – the principle is what matters.

AI SYSTEM OR GPAI MODEL?

OpenAI – just like Google, Anthropic, or Mistral – offers a wide range of AI services on its platform. These remote services are often simply labeled as “models.” While this terminology is unproblematic in the U.S., it can be misleading under the AI Act.

What is actually offered remotely, often without clear distinction, may include:

- AI systems as defined in Art. 3 No. 1 AI Act,
- general-purpose AI models (GPAI models) as defined in Art. 3 No. 63 AI Act, and
- GPAI systems as defined in Art. 3 No. 66 AI Act.

Only if the respective service is classified correctly can the legal consequences be determined properly:

- The roles (provider, deployer, user),
- the risk classes (prohibited, high-, medium-, or low-risk),
- the corresponding obligations, sanctions, and deadlines.

To make this case more tangible, we will now use the symbols of the Playbook to illustrate the scenario. Two services are relevant here:

- **A large language model → *gpt-5*,**
- **An internet search service → *gpt-4o-search-preview*.**

Link to openAI-Services: https://platform.openai.com

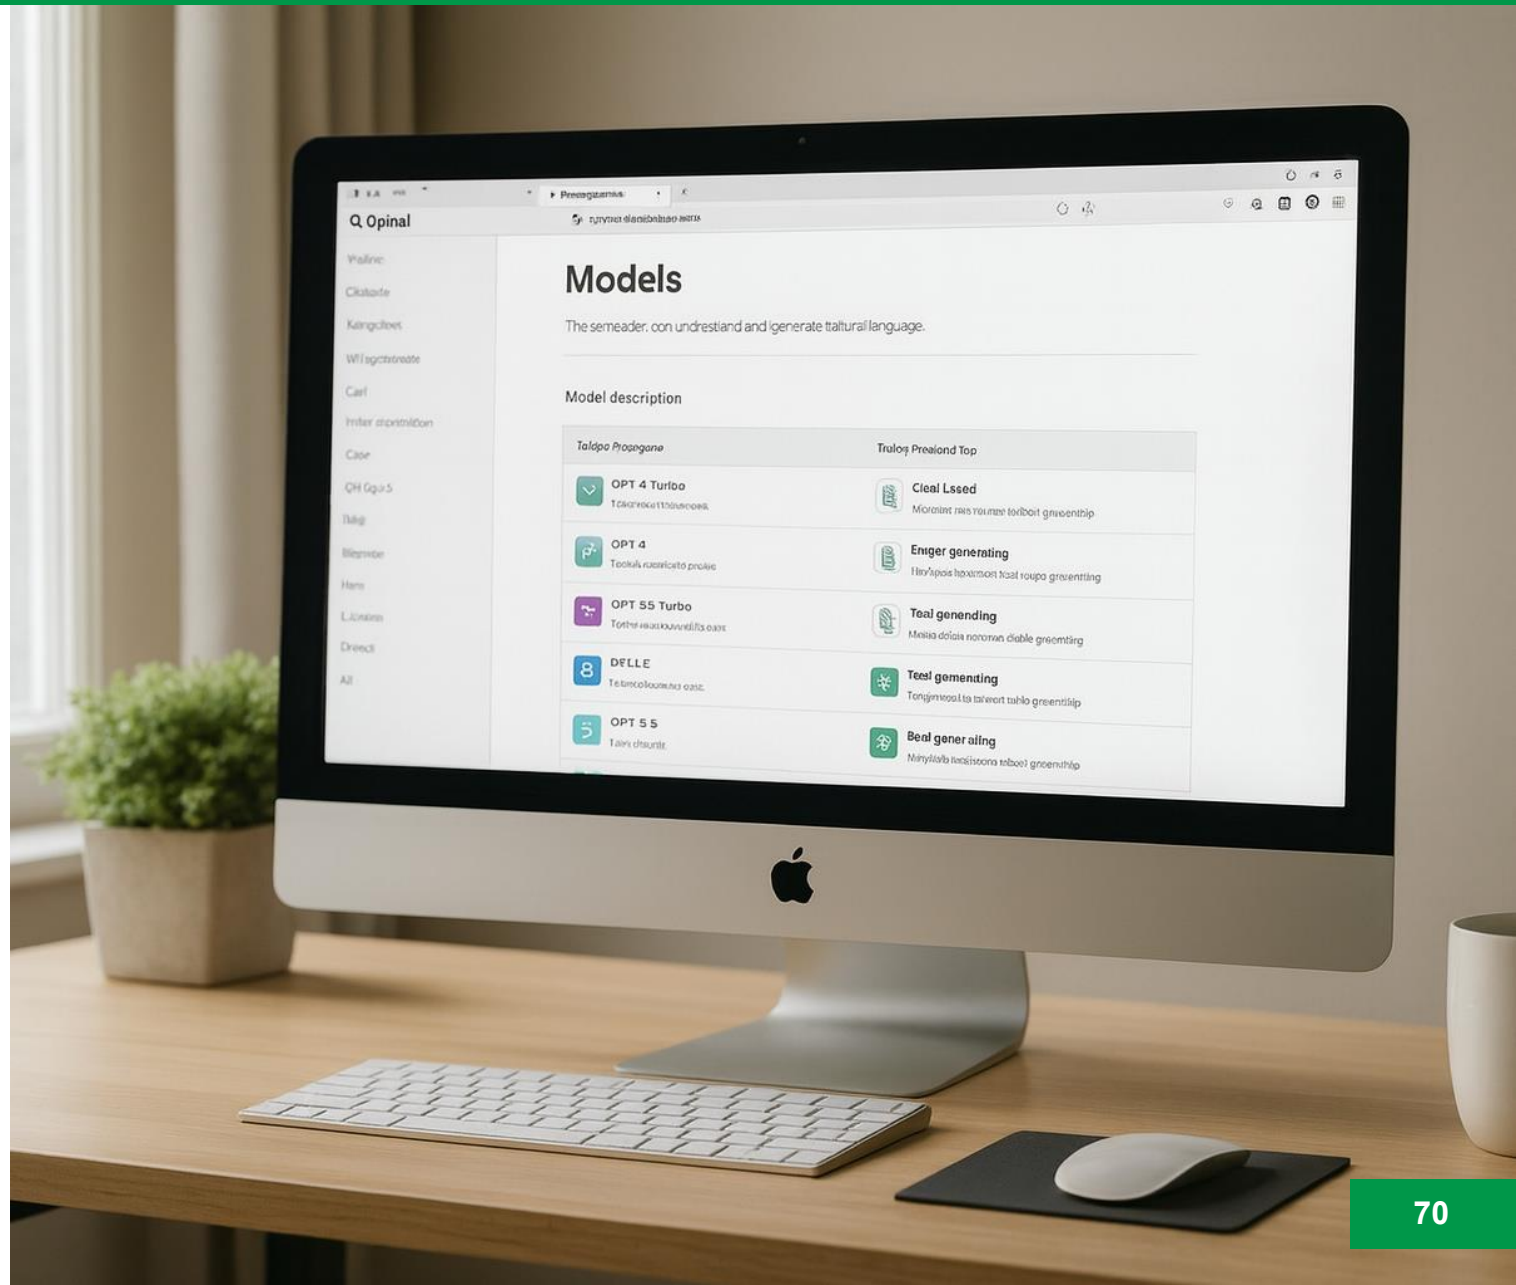

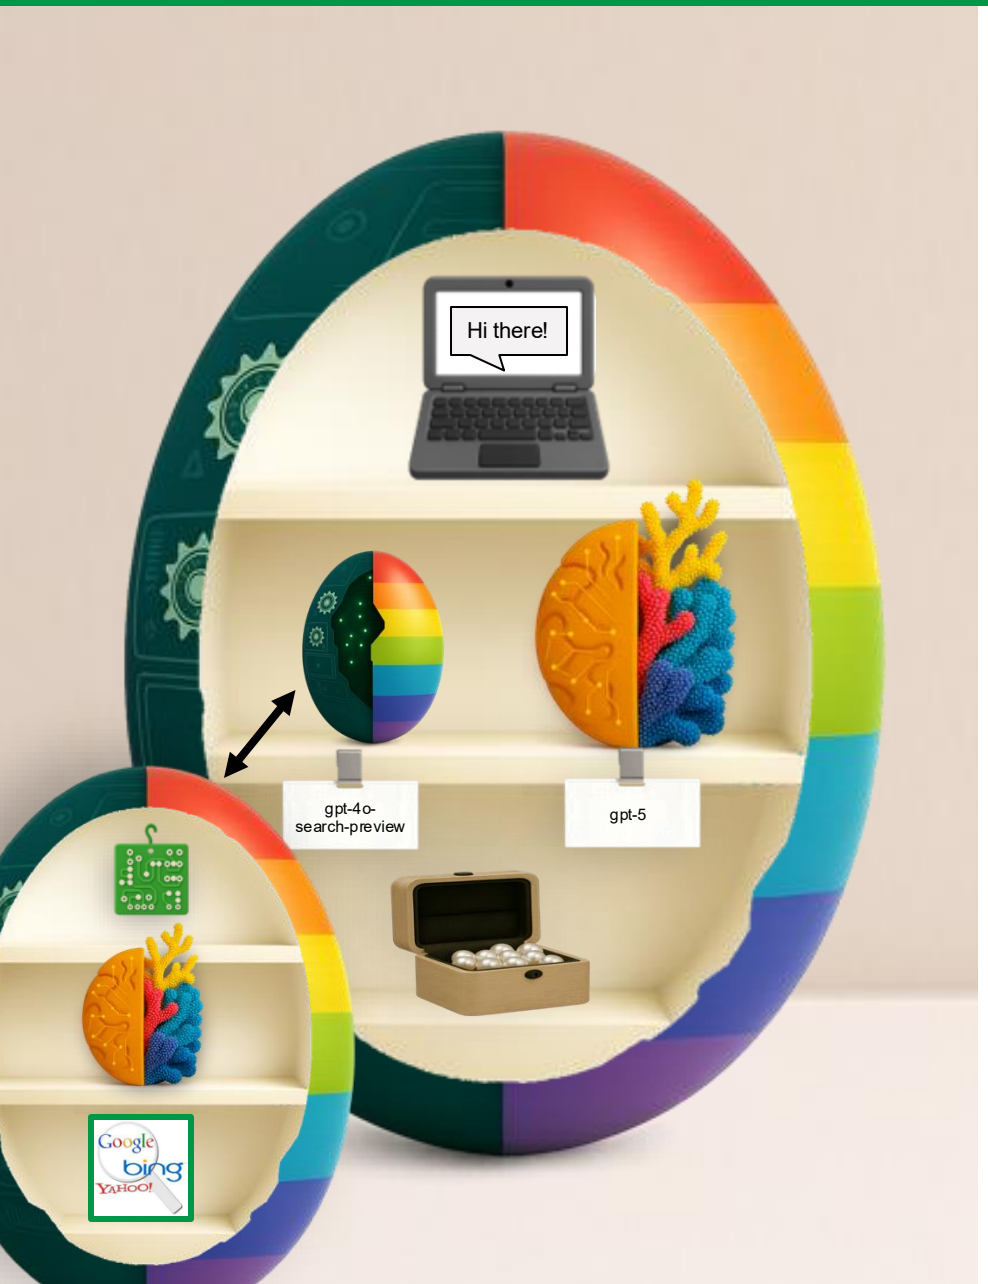

THE CHATBOT FROM THE INSIDE

U's chatbot qualifies as a GPAI system, since it can be used for many different purposes. Its symbol, therefore, is a colorful egg – just as explained on page 21.

- It contains an interaction interface – here shown as a laptop in the upper compartment.
- Within this interface, users can choose the AI service: either a large language model (gpt-5) or a web search service (gpt-4o-search-preview). The interface itself remains the same.
- In the middle compartment – the "intelligence" – the egg contains:
- On the right: a GPAI model (symbol for the LLM gpt-5).
- On the left: a GPAI system, also connected remotely via API key.
- This GPAI system is the web search service gpt-4o-search-preview. It is shown again as a separate colorful egg at the bottom left:
  - In the middle: a colorful coral, the GPAI model "4o."
  - In the upper compartment: an interface chip, enabling the system to process and respond to search queries from the chatbot.
  - In the lowest compartment: the connection to search data from the internet (Google, Bing, Yahoo). This enables real-time outputs, merging external data with the model knowledge.
  - Thus, the chatbot's answers no longer come solely from a GPAI model, but also from a GPAI system inside a GPAI system – the Matryoshka principle (see page 16).

The chatbot typically also contains a session memory and domain-specific data connections, symbolized by the treasure chest in the lower compartment.

The crucial point: On the OpenAI platform, the left egg is labeled as a "model," although in fact it is a system within the meaning of the AI Act.

This directly affects the roles of the actors involved and the legal obligations that arise!

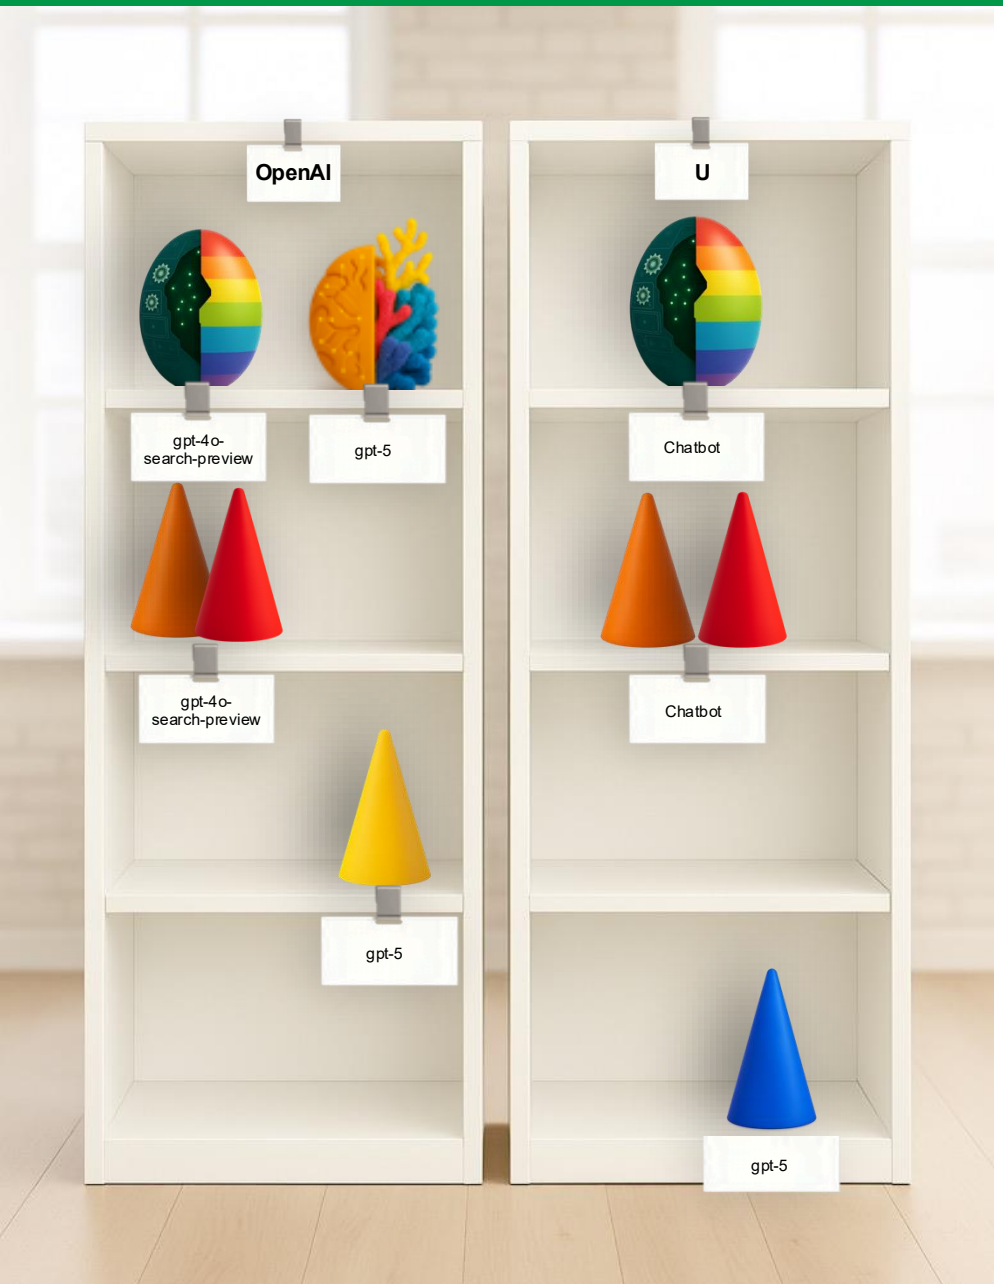

THE DISTRIBUTION OF ROLES

Let's start with the role of S. As a software service provider, S only built the interaction interface – without supplying any intelligence itself. S therefore has no role under the AI Act. S is merely a service provider to U.

U, on the other hand, assumes the roles of both provider and deployer of the chatbot (Art. 3 Nos. 3 & 4, orange/red).

U partly commissioned the chatbot's development and, by using its own API key, also actively created it. Only through the connection to an intelligence service (remote) does the chatbot become a GPAI system (Art. 3 No. 66). U has placed the chatbot on the market and also put it into operation under its own authority (Art. 3 No. 9 –11).

The fact that the chatbot uses two different AI services (one LLM, one search service) does not change its qualification. However, it does mean that U becomes a downstream provider with respect to the LLM (Art. 3 No. 68, blue). This gives U a right to transparency from OpenAI regarding the LLM (here gpt-5). This right does not extend to the search service!

Now let's look at the left side of the shelf: OpenAI provides both the search service gpt-4o-search-preview (egg) and the LLM gpt-5 (coral/coral reef).

- With respect to the search service, OpenAI is both provider and deployer of gpt-4o-search-preview.
- With respect to the LLM (gpt-5), OpenAI is additionally the provider.

Key learnings from this case:

- It does not matter how AI services are labeled (e.g. "model"), but what their true legal character is.
- As a provider/deployer, you must assess independently the character of the AI services you integrate.
- Using the symbols of this Playbook helps make such distinctions clearer – especially when similar terms otherwise create confusion.

This example shows that the use of illustrative "serious play" methods has a very real and practical legal background!

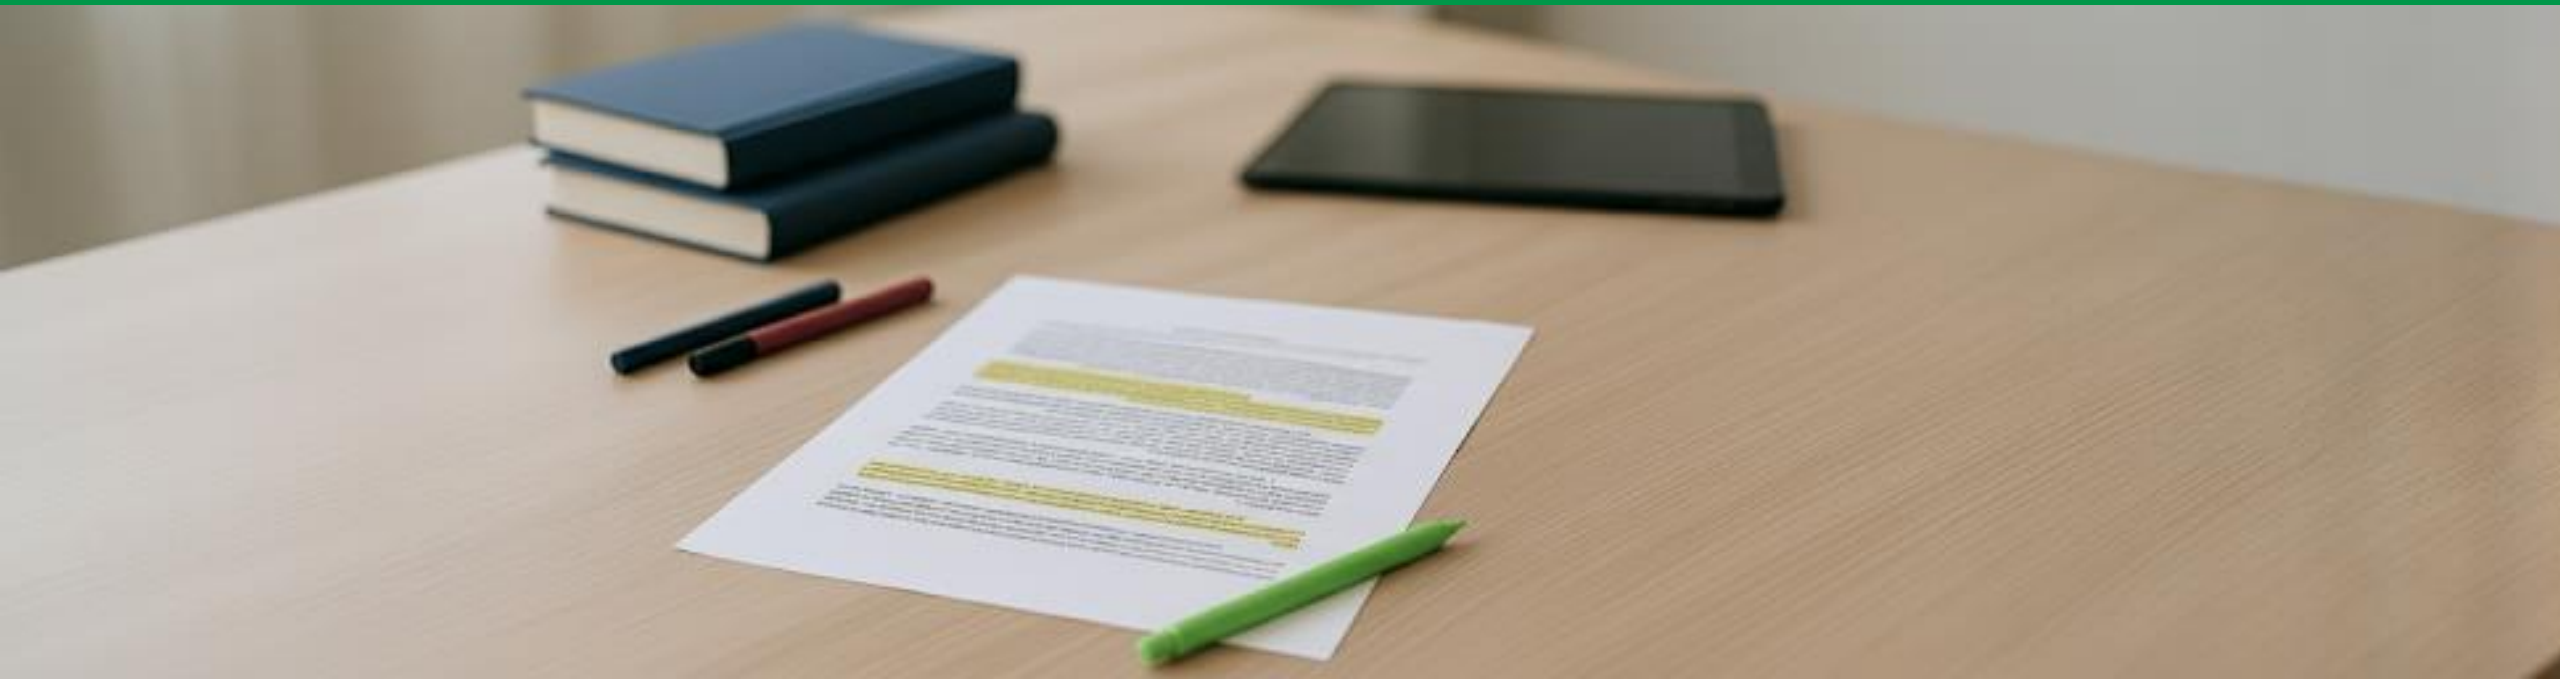

RECOMMENDATION FOR FURTHER READING

In the **German version of the Playbook**, some of the topics presented here are further explored in greater legal depth, based on the companion Script *Grundwissen KI-Recht (Foundations of AI Law)*.

Such legal detail is essential for a conclusive assessment of cases like the example just outlined. This applies even more strongly to individual cases, which ultimately must be solved with the “classic legal toolbox.”

To highlight the connection between Playbook and Script, the German edition therefore includes selected examples from the Script.

In other language versions, this part is not included in the EU AI Act Playbook, since the Script *Grundwissen KI-Recht* has been published only in German.

It is nevertheless recommended to consult legal literature available in the respective national language to further deepen the topics introduced in this Playbook.

Link to the script Grundwissen-Skript (german only):
http://www.grundwissen-ki-recht.de

ANNEX

SAFE USE OF STEP LADDERS






USE OF THE CIRCULAR SAW ON CONSTRUCTION SITES






SUN PROTECTION DURING CONSTRUCTION WORK






IMPRINT

EU AI ACT Playbook
Version 1.0, September 22, 2025

Autor:
Ass. Iur. & AI Officer (AIO)
Oliver M. Merx
83229 Aschau, Bavaria/Germany

Contact:
https://www.linkedin.com/in/oliver-m-merx-83777b/

Internet:
https://cair4.eu
https://grundwissen-ki-recht.de/
https://grundwissen-ki-recht.de/playbook/

Note:
The image on the previous page shows the safety sign of a construction site. The original is located at a new highway bridge project for the A8. After completing the script *“Grundwissen KI-Recht”*, the author came across the symbolic language used on the sign during an e-bike tour. It struck him how this visual language can be understood in a similar way by construction workers of different backgrounds. This moment became the inspiration for developing the symbolic language of this Playbook.