


Deloitte.



Generative AI
is all the rage

Deloitte AI Institute

About the Deloitte AI Institute

The Deloitte AI Institute helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the “Age of With”.

The Deloitte AI Institute aims to promote a dialogue for and development of artificial intelligence, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases.

Combined with Deloitte’s deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, deliver impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you’re in; whether you’re a board member or a C-Suite leader driving strategy for your organization, or a hands on data scientist, bringing an AI strategy to life, the Deloitte AI institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet ups and live events. Let’s explore the future of AI together.

www.deloitte.com/us/AllInstitute

Let's take a moment to cut through the hype.

The AI field took a turn with the release of powerful Generative Artificial Intelligence (AI) models, and as a result, the world is seeing the automation of some skills around creativity and imagination sooner than many expected. For some organizations, Generative AI holds valuable potential for higher order opportunities, like new services and business models. **Deloitte offers a method for selecting Generative AI use cases, as well as some next steps for business leaders in the Age of With™.**



The rise of Generative AI

Generative AI has captured attention in global media and the public square, prompting questions and discussions around this transformative technology.

Businesses, research organizations, and even lay users are experimenting with Generative AI, and given the excitement and interest, it is important to look more closely at the potential capabilities and implications for business.

Generative AI is a subset of artificial intelligence in which machines create new content in the form of text, code, voice, images, videos, processes, and even the 3D structure of proteins. Some forms of Generative AI have been well established in this decade, but it was a large language model (LLM) powering an easily accessible chat interface that enabled Generative AI to have its breakthrough moment and surprise even specialists in the field.

As with other types of AI before it, this new poster child of AI is stimulating the imagination as organizations and individuals consider how to use this tool to benefit both business and society. Generative AI can help in incremental digitization and basic productivity use cases (e.g., more effective text-based channels). But its grander potential is in the higher order opportunities, such as new services or business models that were previously uneconomical.

Generative AI in general and LLM-powered chatbots in particular are not without risks, and this is prompting discussions around things like the potential for job losses and legal questions around intellectual property and ownership. What is more, because the chatbot mimics coherent human phrasing, it may give some the impression that the AI understands the prompts to which it responds, which can lead users to anthropomorphize the chatbot (i.e., the ELIZA effect, as seen in the work of computer scientist Joseph Weizenbaum).

Deloitte is working on a variety of projects exploring the opportunities and business value Generative AI can create for our clients. From experiences and conversations thus far, the clear path ahead, as with all AI, is to attempt to discover and capitalize on capabilities while also responsibly managing the risks that are already emerging.



In this article, we look closely at the potential benefits and limitations of Generative AI, introduce a method to qualify if, where and how these cognitive tools could be used, and offer important factors for business leaders to weigh when adopting Generative AI.

In a prior article, “[Implications of Generative AI for businesses](#),” Deloitte offered a deep dive on the qualities and capabilities of Generative AI, the state of the market, and what that means for organizations wading into this fast-evolving technology field. And in forthcoming articles, we will discuss questions from legal, ethics, risk, and talent and technology perspectives and provide insights into industry use cases.

While this is still the beginning, it's moving fast.

Among organizations across industries, there is interest in differentiating AI use cases, from public service applications to addressing climate change and transforming business functions (see [Deloitte's AI Dossier](#)). AI is viewed as a tool that can automate skills and tasks performed by humans, and AI can be so successful in this regard that humans can forget skills that have been automated. Examples include writing assistants, home automation, and automotive navigation systems. Would most people have the ability to navigate a new city without a mobile phone?

We have seen these kinds of automations emerge across a variety of areas and skillsets. The assumed roadmap for AI was that, in the shorter term, AI is most valuable as a way to automate operational skills, and creative skills will remain the exclusive province of human thinking for the foreseeable future. With Generative AI, this roadmap has taken an unexpected turn.

In late 2022, with the release of an easy-to-use Generative AI chatbot, more people began to discover and imagine how this new technology can be used in the creative space. The chatbot use case opened the door for thinking more broadly about how Generative AI can be used for tasks, ranging from writing copy to generating 3D structures and to outputting organizational processes. As such, we are now seeing the automation of some skills around creativity and imagination sooner than many expected.

There is a lot left to discover. In this Age of With™, the era of humans working with intelligent machines to achieve things greater than what either could do alone, Generative AI will impact the future of work and become a common tool throughout various aspects of our daily lives. In some cases, the applications may be clearly visible, but more often than not, they may be working in the background.



The evolution of Generative AI

The ability of Generative AI to create a convincing (albeit low-quality and greyscale) image of a human face emerged in 2014. Since then, the image quality has increased, and today, almost anything that can be described in words can also be generated as an image, using a textual description called a “prompt.” Throughout 2022, social media users tinkered with Generative AI platforms and shared the results. We have seen avocado armchairs and photorealistic images of astronauts riding horses on the Moon. Cosmopolitan magazine was the first to publish a cover page created by an AI-based image generation tool, and there has even been a case of a user who submitted an AI-generated image to a fine art competition—and won first place. Today, we are seeing similar improvements in other kinds of Generative AI. You may even have found this article via a chat with an AI system that integrates with a search engine.



Images created with Generative AI.

How Generative AI works: Understanding the basics

To understand how Generative AI will impact business and life, we need to understand what it is, what it can do, and what it cannot do, *yet*.

Machine learning has dominated the field of AI for decades. Generally, this approach to AI development is rooted in the concept of learning from examples, rather than following explicitly programmed rules. This is important as there are many tasks that humans perform based on tacit knowledge (and thus can provide examples) but cannot describe the underlying rules to do so. For example, humans know how to recognize a face, but the rules that would instruct an AI system to do the same are much less clear. The approach of learning from examples has led to the development of powerful tools that can identify intricate patterns in complex data.

In a process referred to as training, the algorithm is supplied with a large dataset of input/output examples to extract patterns from the input, which allow conclusions about the expected output. Spam filters, for example, use these patterns to identify similarities in data points and relate those to different classes (i.e., sorting email to a spam folder). While the input data has become more complex over time, from simple arrays of numbers to high-resolution photos, the output side of a model has to this point been mostly limited to categories like “spam” or “not spam,” “cat” or “dog,” or numerical values such as 7°C or \$29. This approach powers nearly all AI that has been deployed so far, the result is “single purpose” applications that can only perform one task.



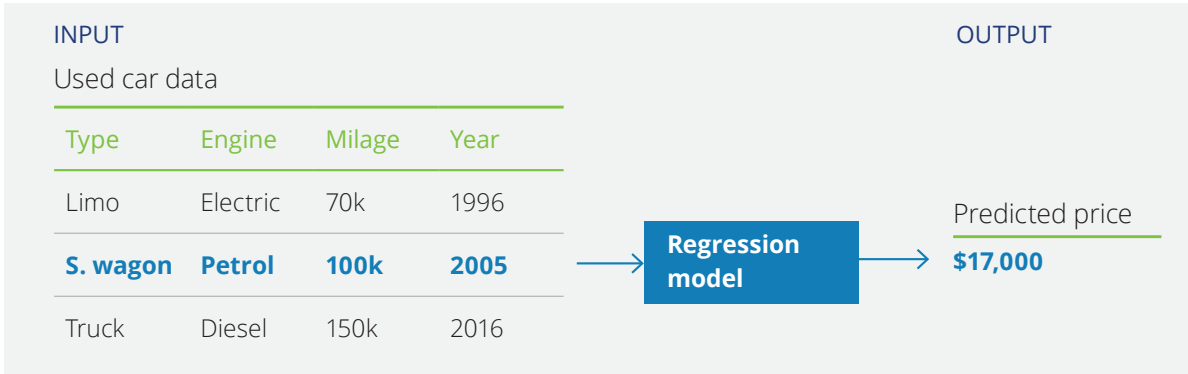


Figure 1: An example of a single purpose machine learning model, using a regression model to predict the resale value for a given car.




Figure 2: An example of how a single purpose model can be used to sort e-mails by “spam” or “not spam.”

Enter Generative AI

The main difference between “traditional AI” and Generative AI is that in the latter, the output is of a higher complexity. Rather than just a number or a label, the output can be an entire high-resolution image, a full page of newly written text (which is generated word by word), or any other digital artifact. This introduces an interesting new element: There is usually more than one possible correct answer. This results in a large degree of freedom and variability, which can be interpreted as creativity.

Generative AI models are typically large and resource hungry. Creating them requires terabytes of high-

quality data processed over weeks on large-scale, GPU-enabled, high performance computing clusters.

Only a few institutions have the necessary resources and talent to build such models. Running a model also requires a lot of compute, which is why access to these kinds of models is often provided via an application programming interface (API). This allows developers to use the models with their existing software products without need for additional infrastructure. These models are versatile and can be fine-tuned for specific tasks, hence they are called Foundation Models. Unlike single-purpose AI, they are suited for multi-purpose tasks.




Figure 3: With Generative AI, user prompts lead to artifacts that can contain a large degree of freedom and variability.

Regarding risks and limitations

Current Generative AI models have limitations. Perhaps the most well-known is termed “hallucination,” which refers to a high-confidence response that is not grounded in the training data. In other words, the response is fictional. For some applications, like art generation, this is a non-issue and perhaps even a desired “creative” feature of Generative AI. For other applications, however, such as copywriting or computer code generation, hallucinations can result in artifacts that are not entirely valid or true, which undercuts the potential value of Generative AI.

Another limiting factor is that today’s Generative AI models generate artifacts based on the model itself and the input prompt. Other additional sources and datasets cannot currently be integrated directly into the model’s internal information processing without costly re-training or fine-tuning, which means Generative AI models can only access information extracted from the data on which they were trained. For similar reasons, they cannot provide references and sources for the generated content, which complicates validation. Furthermore, current models have a context window of a few thousand words, which is the limit for the size of the combined input and output. However, Generative AI models can be combined with other systems (e.g., search, conversational AI) to leverage the benefits of both parts. For example, with a chatbot, a conversational AI system can serve as an orchestration layer between the Generative AI model, a search engine, and the user, which helps to amplify the user experience.

Similar to other AI models, Foundation Models can reproduce latent bias in the training data, and of course, they lack comprehension and the ability to reason as humans do. This has implications for the broader concept of [Trustworthy AI™](#). After all, they are language models, image models, or voice models but not knowledge models.

Despite limitations, Foundation Models can function at such a high quality that many new use cases become possible.

Some known limitations of current Generative AI



Hallucination | Generative AI systems create confident responses that cannot be grounded in any of its training data.



Bias | Similar to other learned models, Foundation Models inherit the bias contained in the training data.



Lack of human reasoning | Generative AI systems are based on statistical features, which is not a solid foundation for logical reasoning.



Limited context window | Current models have a context window of a few thousand words, which is the limit for the combined input and output of the model.

Generating revenue using Generative AI

Using this technology for business benefit can be conceived along two distinct approaches.

First, the models can be used as they are available today, a simple interface that allows near-direct access to the underlying model in the form of a chat for text or an image generation tool. The second approach is to integrate Generative AI with other technologies to automate processes. For example, Generative AI can allow for human-level expressive interactions, while a conversational AI system (i.e., a chat- or voicebot) controls the flow and ensures factual accuracy. An example is an automated, Generative AI-powered call center. We expect the second approach will provide the most value.

A good start to identifying use cases is to find processes or tasks where a digital artifact of some kind is created or processed. This could range from

a job advertisement or a floor plan, all the way to the 3D model of an engine part, a molecule with certain properties, or a workflow, to name a few. Use cases with high usage frequency are preferred, as there will be more example data to fine-tune and improve a model, and subsequently a more substantial impact. Other factors to consider in selecting high-value use cases are existing skill and cost bottlenecks with human generated artifacts. The quality of the artifact may in some cases require human effort, but if it can be created with Generative AI to a commensurate quality, then the human can be liberated to work on higher quality tasks. By turning lower-level creative tasks over to Generative AI, we could see things like databases providing stock content (e.g., images, sounds, or texts) turned upside down as these digital artifacts can be created instantly with a prompt.

If a task requires effort to execute but is easy to validate, it might be a good use case.

Deloitte has designed a Digital Artifact Generation/Validation method to help innovation leaders determine whether an idea can be turned into a beneficial use case leveraging Generative AI. At the core of this method are two of the most critical elements to consider: the human effort required to complete

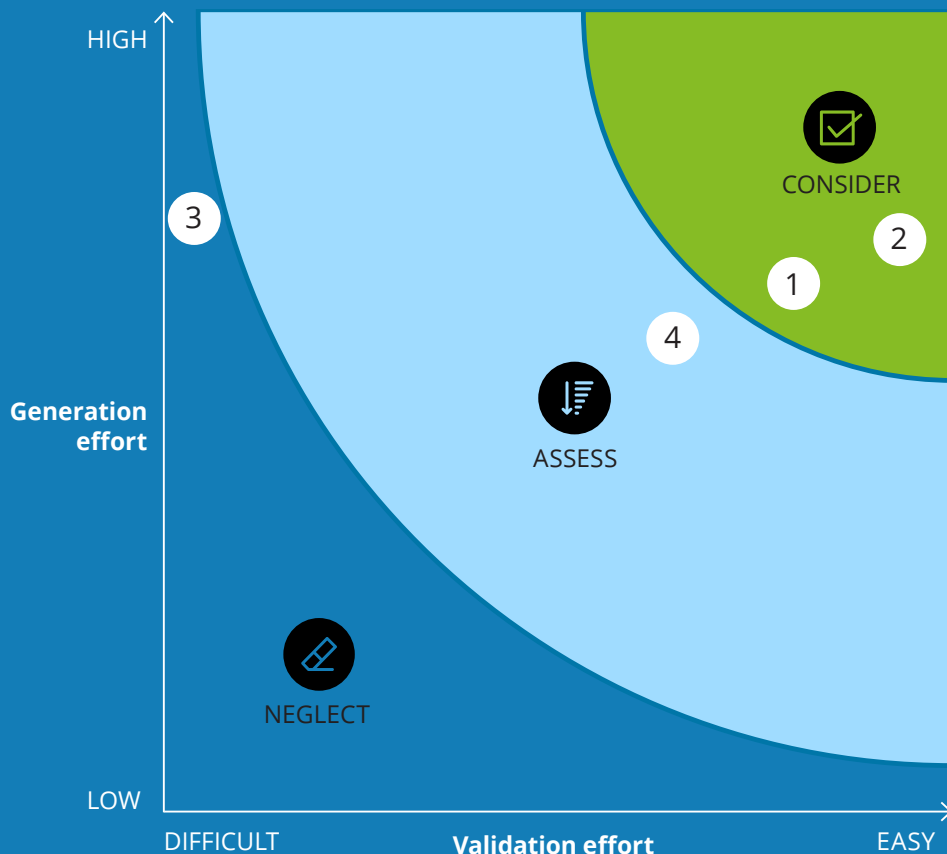


a task without Generative AI; and the necessary effort to validate or fact check the output from the Generative AI. This leads to a two-dimensional classification, categorizing use cases based on the required human effort and the ability of the user to validate the results.

There is a sweet spot for Generative AI use cases

Generative AI is useful where the artifact generation effort is high and validation is easy

Digital Artifact Generation/Validation method



Identifying desirable use cases

Generative AI assists best in use cases where human effort is high, while validation is easy.

Generation effort

How much human effort is required to achieve the desired result

Validation effort

How much human effort is required to check the plausibility or correctness of the result

Examples plotted above

- 1 Create a joke**
 While creating a good joke requires some effort into designing the punch-line and best storytelling-style, it is easy to validate simply by reading it.
- 2 Draw an image of an elephant under a palm tree**
 Drawing any sophisticated image requires reasonable effort for most people regardless of the tools available. On the other hand, validation is easy since you can just look at the picture.
- 3 Draft a contract (without legal expertise)**
 If you do not have legal expertise, drafting a contract is very hard and validating it difficult. Generative AI is useful where the artifact generation effort is high and validation is easy.
- 4 Draft a contract (with legal expertise)**
 If you do have legal expertise, drafting a contract still requires effort, but validating it is significantly easier.

Generative AI is all the rage

For example, re-writing text can be a daunting and time-consuming task for a human. Generative AI tools can take original text and quickly produce a re-written result, a shorter text, a summary, or even a different writing style. A user who is familiar with the original content can validate the accuracy or correctness of the output. Thus, this could be a promising application of Generative AI.

Yet, if the user is reviewing outputs that are outside of their area of expertise, validation becomes more complicated.

The Generative AI output may read as coherent and convincingly accurate, but the potential for a “hallucination” remains. If users lack the knowledge to validate the output and spot hallucinations, the use case is revealed to require high levels of effort for validation and mitigating the risks from hallucinations.

There is an additional complication that should be considered. If the model outputs are consistently correct, users may, over time, become less rigorous in fact checking. Inevitably, however, the model will make an error, and part of the challenge is that the errors may not be obvious, particularly when Generative AI is used to create more complex things, like programming code. Thus, when assessing the ease of validation, weigh the importance of ongoing attention to fact checking.










Insights from Deloitte projects on Generative AI: Reaping benefits from Generative AI requires more than identifying a good use case

Identifying use cases is only part of the challenge. Whenever a transformative technology emerges, some organizations are spurred to experiment for the sake of its novelty, which can lead to “random acts of digital” that do not deliver the anticipated return. Driving business results with Generative AI requires a strategy and collaboration from a cross-disciplinary team. In addition, with a technology that is advancing and maturing as quickly as Generative AI,

avoid the temptation to go forward alone and instead find support and knowledge from partners, colleagues, and third-party organizations operating in this space.

The inherent complexity in current projects leveraging Generative AI requires a cross-disciplinary team to guide and govern the AI lifecycle. Professionals from a variety of domains can help the business answer critical questions, including:

DOMAIN						
						
Ideation & Product Development	Business Operations	Customer & Marketing	Enterprise Technology	Human Capital	Risk Management	Regulations & Laws
STAKEHOLDERS						
Creatives, designers	CEO, COO, Line of Business leaders	CMO	CIO, CTO, IT	CHRO	Risk officers	Legal & Compliance
KEY QUESTIONS						
What can Generative AI permit that reduces human effort and can be rapidly validated?	How does the Generative AI fit into and enhance existing processes and enterprise strategy?	How can the use case be leveraged to build customer engagement, and how much transparency is appropriate?	Can the existing MLOps-tech stack and platform licenses fuel Generative AI, or are third-party services required?	Does the workforce possess the skills to use Generative AI, and what are the implications for talent acquisition and upskilling?	What risks emerge when deploying Generative AI (e.g., jailbreaks, prompt-spooing), and how do these risks impact Generative AI value?	What current and expected laws and regulations concern the use of Generative AI, and are existing governance and MLOps processes sufficient to meet those laws and regulations?

Based on our observations and experience, we recommend business leaders avoid jumping head-first into the hype. Instead, we encourage decision makers to:

1



Develop a strategy for Generative AI and integrate and harmonize it with the enterprise's existing AI strategy. The same principles that guide an AI-fueled organization apply to the use of Generative AI (e.g., access to curated enterprise data; AI governance; process transformation to leverage cognitive workers, etc.). With a technology evolving this quickly, avoid the temptation to go forward alone. Find support and knowledge from partners and third-party organizations operating in this space.

2



Become familiar with the underlying technologies that make Generative AI possible, as well as the current capabilities and limitations. Educate your workforce in the usage, risks, and capabilities of AI to establish a baseline of knowledge through training. Also monitor over time how the technology advances and the impact on business risks and opportunities, as they emerge. This article series may support your efforts.

3



Bring together a cross-disciplinary team of people with the domain knowledge to think creatively about potential use cases. When business leaders, technology leaders, and creatives work together with external experts, they are able to identify valuable applications and also design Generative AI deployments, to mitigate business and cyber risks and meet applicable laws and regulations.

4



Leverage Deloitte's Digital Artifact Generation/Validation method to identify where Generative AI might impact your value chain, with incremental digitization from basic productivity use cases to higher order opportunities, such as new, differentiating services or business models.

5



Ensure the collection and curation of proprietary data, as this is key for tailored use cases that provide a differentiator or competitive advantage.

6



Assess use cases against Trustworthy AI™ principles, including challenges around bias and misinformation, attribution, transparency, and enterprise accountability for the impact from Generative AI.

Deloitte is excited to move into the future with our clients and colleagues, as well as with our connections in academia and the broader AI ecosystem around the world.

The discussions so far show that there is a need for a deeper understanding of Generative AI, from the underlying technology to its impact on the future of work. As such, it is important to look closely at the implications for risk, trust, and governance, which is investigated in a forthcoming article, "Proactive risk management in Generative AI." There are also legal considerations for Generative AI, which we plan to cover in "Legal implications of using Generative AI (What the AI System won't tell you)."

There is a lot to cover and the conversations are far from over. Deloitte is a trusted advisor as we push beyond the initial buzz around this new technology and into how Generative AI can be used for good in the Age of With™.



Reach out for a conversation.




Beena Ammanath
Global Deloitte AI Institute
Leader
Deloitte AI Institute
United States, Lead
Deloitte Consulting, LLP
bammanath@deloitte.com



Wessel Oosthuizen
Deloitte AI Institute Africa,
Lead
Deloitte Africa
woosthuizen@deloitte.com



Dr. Kellie Nuttall
Deloitte AI Institute
Australia, Lead
Deloitte Australia
knuttall@deloitte.com



Jefferson Denti
Deloitte AI Institute Brazil,
Lead
Deloitte Brazil
jdenti@deloitte.com



Audrey Ancion
Deloitte AI Institute
Canada, Lead
Deloitte Canada
aancion@deloitte.ca



Jan Hejtmanek
Deloitte AI Institute
Central Europe, Lead
Deloitte Central Europe
jhejtmanek@deloitte.com



Roman Fan
Deloitte AI Institute China,
Lead
Deloitte China
rfan@deloitte.com



Anne Sultan
Deloitte AI Institute France,
Lead
Deloitte France
asultan@deloitte.com



Dr. Bjoern Bringmann
Deloitte AI Institute
Germany, Lead
Deloitte Germany
bbringmann@deloitte.com



Prashanth Kaddi
Deloitte AI Institute India,
Lead
Deloitte India
pkaddi@deloitte.com



Masaya Mori
Deloitte AI Institute Japan,
Lead
Deloitte Japan
mmori@deloitte.com



Nicholas Griedlich
Deloitte AI Institute
Luxembourg, Lead
Deloitte Luxembourg
ngriedlich@deloitte.com



Naser Bakhshi
Deloitte AI Institute
Netherlands, Lead
Deloitte Netherlands
nbakhshi@deloitte.nl



Tiago Durao
Deloitte AI Institute
Portugal, Lead
Deloitte Portugal
tdurao@deloitte.com



Sulabh Soral
Deloitte AI Institute
United Kingdom, Lead
Deloitte United Kingdom
ssoral@deloitte.com

Special thanks to contributors:

Jakob Nikolaus Moecke, Senior Consultant, Deloitte AI Institute Germany
Philipp Joshua Wendland, Senior Consultant, Deloitte AI Institute Germany
Alexander Mogg, Lead Partner, Monitor Deloitte Germany
Kate Fusillo Schmidt, Senior Manager, Deloitte AI Institute US
Erica Dodd, Senior Manager, Deloitte AI Institute Australia
Jessica Carius, Senior Consultant, Deloitte AI Institute Australia

Elon Allen, Partner, Monitor Deloitte Australia
Bram den Hartog, Partner, Monitor Deloitte Australia
Aisha Greene, Senior Manager, Deloitte AI Institute Canada
Anke Joubert, Senior Manager, Deloitte AI Institute Luxembourg
Dr. Gordon Euchler, Director, Deloitte Germany



This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the “Deloitte” name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.