Breaking creative boundaries: Generative AI and its applications

pwc

Generative AI is no more the 'technology of the future'. It is already here and has extensively reshaped our lives, work environments and the way we interact with technology. It has, in fact, emerged as one of the most important advancements in technology – as fundamental as the creation of the personal computer, the internet or even the mobile phone.

Sumit Srivastav **Partner and Leader – Intelligent Automation, PwC India**

The rise of generative AI is going to reshape life as we know it today. From cognitive chatbots to large language models (LLMs), technology has evolved rapidly within a short span of time and has found relevance across most industries and functions.

As organisations continue to pursue their business priorities, creating a technology ecosystem together with human intelligence and creativity will play a pivotal role in carving out the future.

Although there are apprehensions with regard to AI, generative AI is being deemed as one of the most transformative revolutions of the past few decades – if it is leveraged in the right manner. With continuous advancements in fields such as robotic process automation, intelligent document processing, machine learning (ML), deep learning and natural language processing (NLP), intelligent automation has helped organisations ease complex decision-making processes, optimise workflows, enhance operational efficiency, and amplify customer/end user experience.

At its core, generative AI consists of ML models that ingest enormous amounts of training data, analyse it for correlations and patterns, and use these patterns to make predictions about future states. AI can be leveraged for conversational intelligence, predictions, self-learning and improving productivity.

Additionally, the introduction of intelligent virtual assistants like chatbots in recent years has transformed customer interactions and service delivery considerably. These AI-powered agents can understand natural language, provide recommendations and assist customers in real time. Such virtual assistants have become an integral part of many businesses, enhancing customer experiences and increasing operational efficiencies.

While technology has made tremendous strides in recent years, the ability to be creative was always a big part of what separated human beings from technology. However, the emergence of generative AI backed by LLMs is changing this scenario, giving these technologies, the ability to efficiently generate original content that is sometimes indistinguishable from the work of a human being.

Infusing creativity into technology: Role of LLMs in generative AI

Human creativity and imagination often stem from making connections. Our senses help us form an opinion or idea which is then translated into a poem, song or other work of art.

Generative AI's ML algorithms have extended this human ability to be creative and produce content that can be far more detailed, sophisticated and in-depth, by studying a large amount of training data. Using this trained data, it learns to create something new based on an input (also known as a 'prompt').

It is important to note that exceptionally large data sets are crucial for the success of generative AI. Language models were introduced with the intent of representing words in the form of numbers, and one of the simpler ways to do the same included taking a count of each word in a sentence and creating a vocabulary of such representations over a large data set.

However, in recent years, LLMs have progressed significantly with considerable changes, complexities and large datasets being fed into the models over the last six years. These LLMs now act as the basis for generative AI – which is the broad term used to describe any type of AI that is used to create text, images, audio, video, code or synthetic data.

One of the most advanced AI models of recent times – OpenAI's ChatGPT – has ushered in 'accessible-to-all' AI models, which ultimately impact day-to-day business-as-usual activities. ChatGPT stands out for its ability to generate human-like conversation in response to a user's input. It uses NLP along with LLMs, which allows it to generate natural-sounding text and reimagine the way users interact with AI. These capabilities fall under generative AI, underlining the shift in AI from the statistical approach to neural language models.

Artificial intelligence landscape
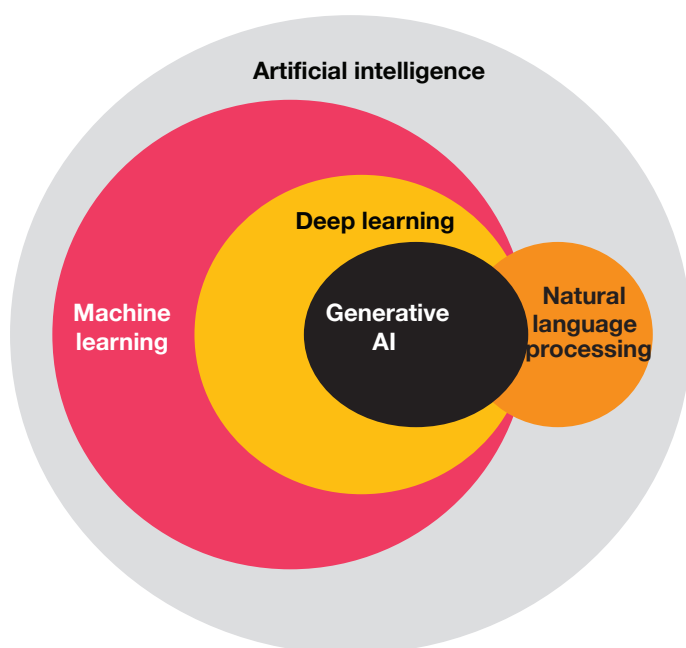
Source: Intelligent Automation, PwC India

However, no matter how impressive a computer-created output might be, it is important to remember that it is always built from the blocks carved out of the data that was used to train it.
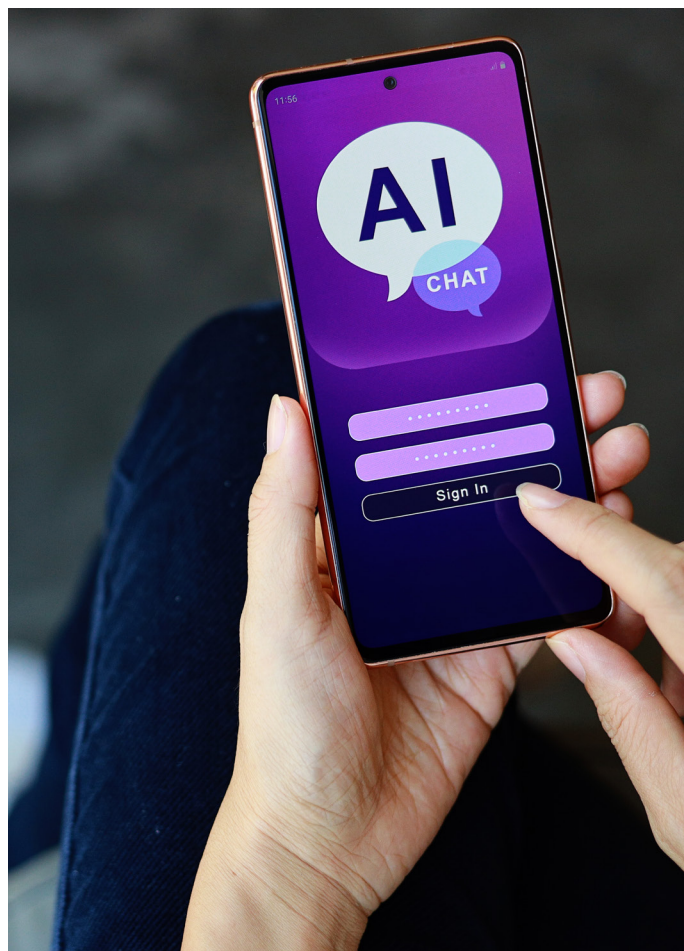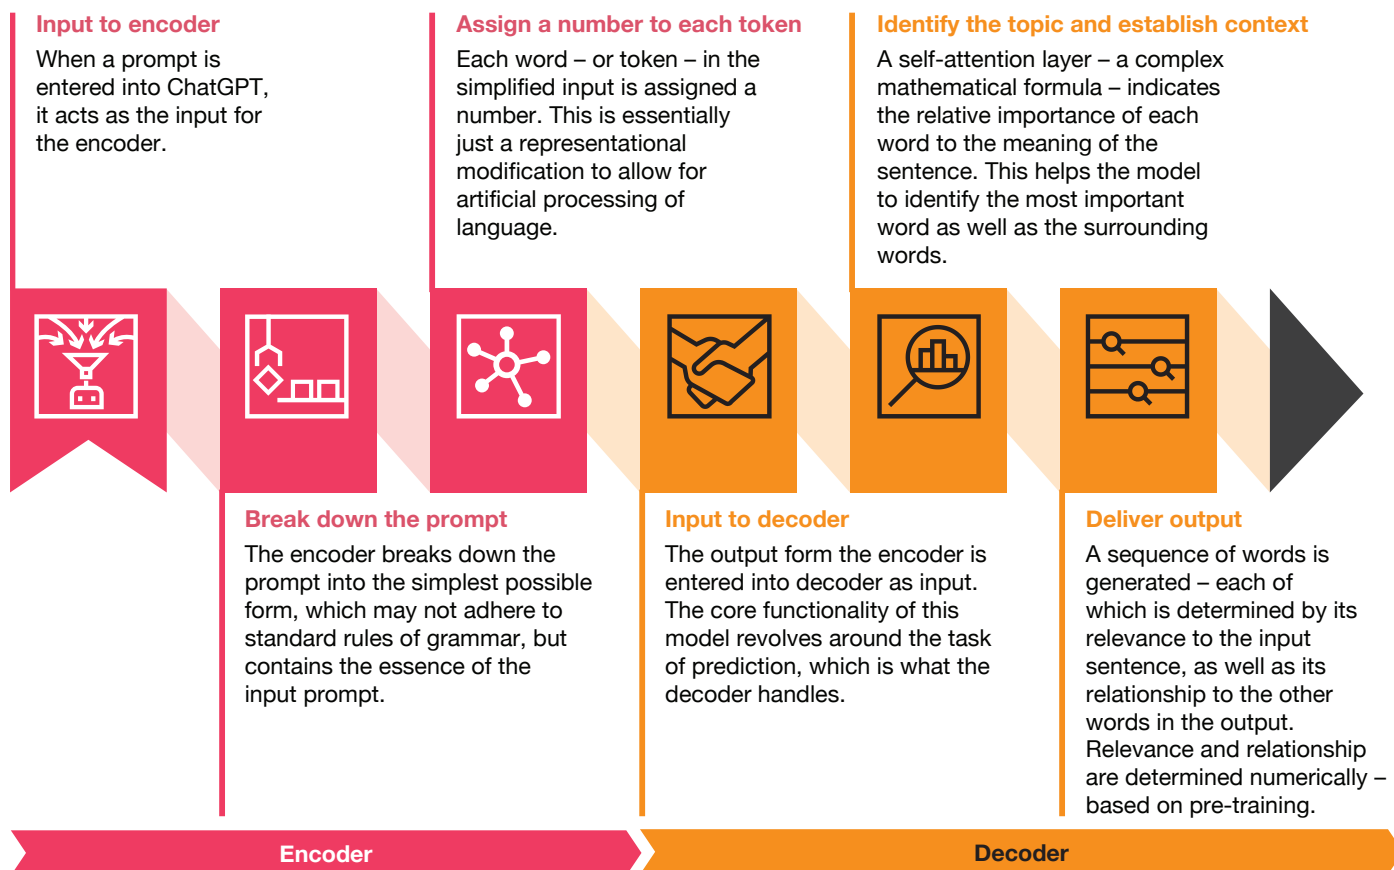
Training LLMs – a closer look at ChatGPT

OpenAI's ChatGPT is one of the most popular tools that harnesses LLMs today. It is estimated to have reached 100 million monthly active users just two months after its launch in November 2022, making it the fastest-growing consumer application in history.[1]

ChatGPT is built on generative pre-trained transformer (GPT)-4, to make responses balanced and more user-aligned. The dataset used for training GPT consists of a massive amount of content scraped from across the internet.

Demystifying GPT: Understanding language processing

Input to encoder
When a prompt is entered into ChatGPT, it acts as the input for the encoder.

Assign a number to each token
Each word – or token – in the simplified input is assigned a number. This is essentially just a representational modification to allow for artificial processing of language.

Identify the topic and establish context
A self-attention layer – a complex mathematical formula – indicates the relative importance of each word to the meaning of the sentence. This helps the model to identify the most important word as well as the surrounding words.

Break down the prompt
The encoder breaks down the prompt into the simplest possible form, which may not adhere to standard rules of grammar, but contains the essence of the input prompt.

Input to decoder
The output form the encoder is entered into decoder as input. The core functionality of this model revolves around the task of prediction, which is what the decoder handles.

Deliver output
A sequence of words is generated – each of which is determined by its relevance to the input sentence, as well as its relationship to the other words in the output. Relevance and relationship are determined numerically – based on pre-training.

Encoder

Decoder

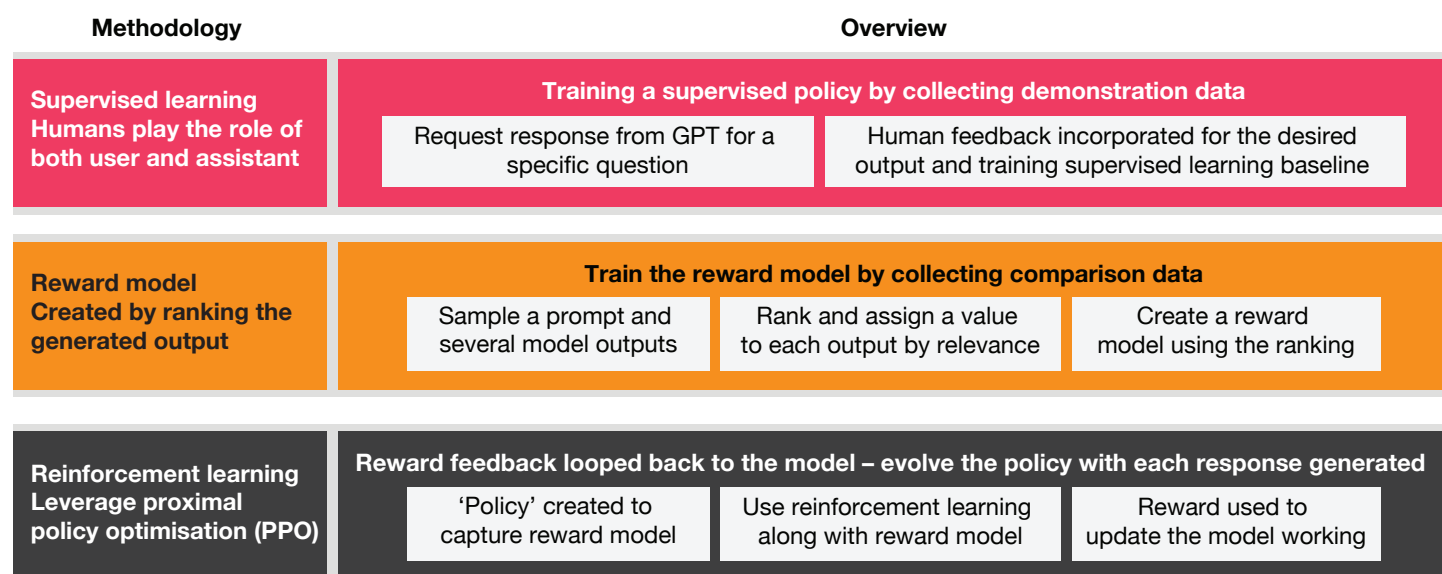**Source:** Intelligent Automation, PwC India

GPT is trained using reinforcement learning from human feedback (RLHF)[2] – a technique that involves extensive training and multi-stage deployment. Human preferences are brought in through a reward model, for which comparison data is collected. It is then trained on question–answer pairs, allowing it to learn which questions and answers are compatible and sensible. Reinforcement learning is the second stage of pre-training, where GPT is given an input and asked to choose the correct output from four options.

1. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/
2. _https://en.wikipedia.org/wiki/Reinforcement_learning_from_human_feedback

Below is a snapshot of the training approach.

Methodology	Overview		
Supervised learning **Humans play the role of both user and assistant**	**Training a supervised policy by collecting demonstration data**		
	Request response from GPT for a specific question	Human feedback incorporated for the desired output and training supervised learning baseline	
Reward model **Created by ranking the generated output**	**Train the reward model by collecting comparison data**		
	Sample a prompt and several model outputs	Rank and assign a value to each output by relevance	Create a reward model using the ranking
Reinforcement learning **Leverage proximal policy optimisation (PPO)**	**Reward feedback looped back to the model – evolve the policy with each response generated**		
	'Policy' created to capture reward model	Use reinforcement learning along with reward model	Reward used to update the model working

Source: Intelligent Automation, PwC India

Broadly, the use cases of LLMs include generating, summarising, rewriting and extracting. In an ideal scenario, the more parameters a model has, the higher the amount of data required to train the model. The vast amount of data sets and multi-modality enables generative AI tools to generate high-quality outputs and unique content.

In the sections below, we have explored various use cases of generative AI and highlighted its potential applications.

Applications of LLM solutions

LLMs have been instrumental to the development of many solutions in the generative AI space. There are different LLM products available to businesses with varied applications, such as generation of text, creation of audio, video. The table below highlights some of the most common applications of these tools.

Application	Description
Text creation	**Generate** • Write product descriptions based on product name and relevant keywords. • Write chatbot/conversational AI responses. • Develop question–answering interface. • Draft emails – given the purpose/command. • Write headlines, articles, brochures based on keyword input by users. **Summarise** • Document customer support interaction. • Generate paper abstracts. • Document dialogues and transcripts. • Collate reporting and compliance documents. **Rewrite** • Paraphrase chat responses. • Redact sensitive information. • Add extra context/background. • Simplify/contextualise communication according to the audience. **Extract** • Parse key information from a document. • Extract information from an image. • Extract document metadata information. • Create tags for blogs.
Image creation	• Generate advanced digital images and artistic impressions based on text inputs. • Perform image-to-image conversions such as drawing-to-image translation, resolution enhancement, erasing parts from a picture, inpainting and outpainting.
Video creation	• Generate long-form videos with automatic subtitles and customisable sounds with a script.
Audio creation	• Generate new and edit existing audio/music based on text prompts.
Image captioning	• Write captions/descriptions for images.
Code creation	• Generate code from scratch based on a natural language prompt. • Offer autocomplete-style suggestions. • Translate code across various programming languages and debug code.
Synthetic data creation	• Create large and diverse sample datasets from an existing dataset as an input.
3D assets creation	• Change text and 2D image prompts to 3D visuals. • Generate human motion.

Gen AI – potential applications across industries

The game changer for businesses

Enterprises across various industries/sectors can use generative AI solutions across various departments to enhance both employee and customer experience. Here are a few examples of its usage:

Customer experience: As chatbots become increasingly human-like in their conversation with the help of LLMs, they are able to handle more nuanced and varied customer interactions, and provide more streamlined answers. In addition, LLMs can be used to improve market strategies and deliver personalised services. With their ability to understand context, they also alter the way search algorithms work, allowing for more relevant search results. Furthermore, tools like Viable[3] and Enterpret[4] organise and summarise massive amounts of customer feedback, providing useful product insights.

Content creation: LLMs can be used to create grammatically correct and contextually relevant textual or visual content, including articles, summaries, and images. From an organisation's perspective, generative AI can be used to create brochures, draft emails, summarise meeting proceedings and build websites or blogs.

Job assistant: The recent decision by Microsoft to embed OpenAI's ChatGPT into popular Microsoft 365 applications like Word, Excel and PowerPoint brings generative AI technology to mainstream productivity tools. Named Copilot, this embedding of ChatGPT-4 allows users to generate a presentation from natural language commands, summarise documents as they're being written, and automatically identify trends and insights from spreadsheets. In addition, Microsoft's search engine, Bing, has also been revamped to include conversational AI functionality, utilising GPT-4 to power Bing Chat.

Software development: Applications like GitHub's Copilot, Amazon's CodeWhisperer[5] and OpenAI's Codex[6] assist coders by generating, completing and reviewing code, including features like bug fixing and refactoring.[7] CodeWhisperer and Codex also have the capabilities to convert natural language to a program. Moreover, GitHub is rolling out Copilot X, which integrates OpenAI's GPT-4 model and brings a ChatGPT-like chatbot experience to code editors. In addition to being able to recognise and explain code, it can suggest optimisations and bug fixes.

Research: With their ability to efficiently summarise large texts, LLMs could significantly reduce the time spent on reading research papers. In addition, the massive datasets allow for greater accuracy in natural language translation, ensuring that the knowledge presented in diverse languages can be accessed and understood by the research community at large.

Education and academia: Generative text tools like ChatGPT are being used as virtual assistants by students while creating educational content. These models can also be used to suggest syntactic or grammatical corrections, assess and streamline writing styles, generate practice problems and quizzes, and help in lesson planning. For example, the new-age online tutoring service, Khan Academy, has released its AI chatbot Khanmigo using GPT-4, through which students can get hints, try out a similar problem or be quizzed on the spot.[8]

---

3. https://www.askviable.com/

4. https://www.enterpret.com/

5. https://aws.amazon.com/codewhisperer/

6. https://openai.com/blog/openai-codex

7. https://research.aimultiple.com/generative-ai-coding/

8. https://www.khanacademy.org/khan-labs

Listed below are some of the other potential applications of generative AI across different industries/functions.

Industry/business functions	Applications	Generates
Finance[9]	**Investment advice portfolio management** – generation of investment scenarios on the basis of historical financial data, to help identify optimal wealth and asset management	Text
	Fraud detection – training algorithms using generated synthetic data of fraudulent transactions for easy identification	Synthetic data
Healthcare	**Medical imaging** – training systems to automatically flag abnormalities in scans and X-rays, alerting doctors to potential issues	Image
	Research – analysing reporting and findings in previous research to suggest new areas of study and uncover trends	Text, image, synthetic data
	Clinical decision making – providing accurate and informed patient diagnosis, as well as customised medication and care, based on real-time patient data	
	Preventive healthcare – generating synthetic data to model new pandemics while developing preventive measures for the future, and enabling enhanced drug discovery and development	
Retail and e-commerce	**Sales and promotions –** generating ideas and concepts for advertisements, jingles and marketing	Text, image
	Product recommendation – providing customers with personalised suggestions, generated automatically on the basis of past purchases and interests	Text
Manufacturing[10]	**Product development and design** – cutting down time to market by generating new designs, identifying shortcomings, and suggesting improvements while keeping in mind efficiency and material considerations	Text, 3D assets
	Inventory management – improving supply chain efficiency and predicting appropriate stock levels through creation and analysis of inventory situations	Text
Customer service	**Recommendations** – using generative AI-based customer service tools to provide recommendations based on order history	Text
	Personalised support – creating more natural conversation flows and ensuring updated website information for users	Text
Energy and utility	**Asset management** – generating and simulating new scenarios and their impact to improve decision making, thus allowing for lower costs and greater efficiency	Text
	Remote monitoring – analysing data from sensors to identify imminent problems and alert companies	Text
Telecommunications and network	**Device troubleshooting** – using generative AI to diagnose device issues and suggest troubleshooting activities	Text
	Network monitoring – keeping track of network traffic and usage, and optimising planning by predicting demand and identifying capacity requirements	Text

9. https://research.aimultiple.com/generative-ai-finance/

10. https://www.xenonstack.com/blog/generative-ai-manufacturing

Application	Use Case	Generates
Education	**Gamification** – creating dynamic tests, which are personalised for students	Text, image
	Test preparation – carrying out mock exams, quizzes, and summarisation of important topics using generative AI tools	
	Teaching assistants – enabling learners to ask questions and receive immediate, relevant answers, creating lesson guides, and transcribing spoken class lessons	Text
Government and public services, legal	**Regulatory compliance** – assisting compliance teams through risk assessment and analysis	Text
	Policy assistance – drafting and reviewing policies, as well as summarising and retrieving information from them	
	Legal data inferencing – researching, simplifying and aligning on defined laws and regulation. Drawing out applicable precedents, previous judgements, defining better arguments and strategically structuring them	Text
Information technology	**Technical support** – facilitating personalised support for troubleshooting common issues	Text, image
	Software development – providing tools for programming assistance and suggestions using generative AI to improve speed and optimisation of code	Text, code

Looking to the future

Until recently, AI has largely remained in the experimental domain. The development of LLM-backed generative AI tools represents a huge breakthrough, as it has made generative AI more accessible. Non-technical users too can enjoy the benefits of generative AI without needing to go into the specifics of the underlying technology. This dynamic form of AI holds immense potential

and is being applied in industry-specific domains as well as across horizontal processes to introduce new ways of working.

The future of generative AI appears promising. There is potential to develop even more sophisticated and innovative tools that can revolutionise numerous industries and significantly improve various business processes. By leveraging generative AI, enterprises can unlock new avenues of creativity, efficiency and problem solving, paving the way for exciting possibilities in the realm of technology-driven advancements.

Continuing this series, in our next thought paper, we will discuss LLM adoption models, challenges faced by organisations during their implementation, required governance and ways of addressing these challenges.

About PwC

At PwC, our purpose is to build trust in society and solve important problems. We're a network of firms in 152 countries with over 328,000 people who are committed to delivering quality in assurance, advisory and tax services. Find out more and tell us what matters to you by visiting us at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

Contact us

Sumit Srivastav
Partner and Leader
Intelligent Automation – PwC India
sumit.srivastav@pwc.com

Mahesh Parab
Partner
Intelligent Automation – PwC India
mahesh.parab@pwc.com

Ritesh Jain
Partner
Intelligent Automation – PwC India
ritesh.k.jain@pwc.com

Hariprasad Gajapathy
Executive Director
Intelligent Automation – PwC India
hariprasad.gajapathy@pwc.com

Contributors

Jahnavi Sivaram

Jasnain Singh

Ashwin Guptha

Mahesh Parab

Shardul Raut

Disha Thadani

Hariprasad Gajapathy

Sumit Srivastav

pwc.in